



Protect Oracle databases

SnapCenter Software 4.7

NetApp

February 12, 2024

This PDF was generated from https://docs.netapp.com/us-en/snapcenter-47/protect-sco/concept_what_you_can_do_with_the_snapcenter_plug_in_for_oracle_database.html on February 12, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Protect Oracle databases. 1
 - Overview of SnapCenter Plug-in for Oracle Database. 1
 - Install SnapCenter Plug-in for Oracle Database 8
 - Install SnapCenter Plug-in for VMware vSphere 35
 - Prepare for protecting Oracle databases 35
 - Back up Oracle databases 37
 - Mount and unmount database backups 67
 - Restore and recover Oracle databases 69
 - Clone Oracle database 86
 - Manage application volumes 107

Protect Oracle databases

Overview of SnapCenter Plug-in for Oracle Database

What can you do with the Plug-in for Oracle Database

The SnapCenter Plug-in for Oracle Database is a host-side component of the NetApp SnapCenter Software that enables application-aware data protection management of Oracle databases.

The Plug-in for Oracle Database automates the backup, cataloging and uncataloging with Oracle Recovery Manager (RMAN), verification, mounting, unmounting, restore, recovery, and cloning of Oracle databases in your SnapCenter environment. The Plug-in for Oracle Database installs SnapCenter Plug-in for UNIX to perform all the data protection operations.

You can use the Plug-in for Oracle Database to manage backups of Oracle databases running SAP applications. However, SAP BR*Tools integration is not supported.

- Back up datafiles, control files, and archive log files.

Backup is supported only at container database (CDB) level.

- Restore and recovery of databases, CDBs, and pluggable databases (PDBs).

Incomplete recovery of PDBs are not supported.

- Create clones of production databases up to a point-in-time.

Cloning is supported only at CDB level.

- Verify backups immediately.
- Mount and unmount data and log backups for recovery operation.
- Schedule backup and verification operations.
- Monitor all operations.
- View reports for backup, restore, and clone operations.

Features of Plug-in for Oracle Database

The Plug-in for Oracle Database integrates with the Oracle database on the Linux or AIX host and with NetApp technologies on the storage system.

- Unified graphical user interface

The SnapCenter interface provides standardization and consistency across plug-ins and environments. The SnapCenter interface enables you to complete consistent backup, restore, recovery, and clone operations across plug-ins, use centralized reporting, use at-a-glance dashboard views, set up role-based access control (RBAC), and monitor jobs across all plug-ins.

- Automated central administration

You can schedule backup and clone operations, configure policy-based backup retention, and perform restore operations. You can also proactively monitor your environment by configuring SnapCenter to send email alerts.

- Nondisruptive NetApp Snapshot copy technology

SnapCenter uses NetApp Snapshot copy technology with the Plug-in for Oracle Database and Plug-in for UNIX to back up databases. Snapshot copies consume minimal storage space.

The Plug-in for Oracle Database also offers the following benefits:

- Support for backup, restore, clone, mount, unmount, and verification workflows
- Automatic discovery of Oracle databases configured on the host
- Support for cataloging and uncataloging using Oracle Recovery Manager (RMAN)
- RBAC-supported security and centralized role delegation

You can also set the credentials so that the authorized SnapCenter users have application-level permissions.

- Support for Archive Log Management (ALM) for restore and clone operations
- Creation of space-efficient and point-in-time copies of production databases for testing or data extraction by using NetApp FlexClone technology

A FlexClone license is required on the storage system where you want to create the clone.

- Support for consistency group (CG) feature of ONTAP as part of creating backups in SAN and ASM environments
- Nondisruptive and automated backup verification
- Capability to run multiple backups simultaneously across multiple database hosts

In a single operation, Snapshot copies are consolidated when databases in a single host share the same volume.

- Support for physical and virtualized infrastructures
- Support for NFS, iSCSI, Fibre Channel (FC), RDM, VMDK over NFS and VMFS, and ASM over NFS, SAN, RDM, and VMDK
- Support for the Selective LUN Map (SLM) feature of ONTAP

Enabled by default, the SLM feature periodically discovers the LUNs that do not have optimized paths and fixes them. You can configure SLM by modifying the parameters in the `scu.properties` file located at `/var/opt/snapcenter/scu/etc`.

- You can disable this by setting the value of `ENABLE_LUNPATH_MONITORING` parameter to false.
- You can specify the frequency in which the LUN paths will be fixed automatically by assigning the value (in hours) to `LUNPATH_MONITORING_INTERVAL` parameter. For information about SLM, see the [ONTAP 9 SAN Administration Guide](#).
- Support for non-volatile memory express (NVMe) on Linux
 - NVMe util should be installed on the host.

You must install NVMe util to clone or mount to alternate host.

- Backup, restore, clone, mount, unmount, catalog, uncatalog, and verification operations are supported on the NVMe hardware except for the virtualized environments like VMDK and RDM.

The above operations are supported on devices without partitions or with single partition.



You can configure multipathing solution for NVMe devices by setting the native multipathing option in the kernel. Device Mapper (DM) multipathing is not supported.

+ For the latest information about supported versions, see the [NetApp Interoperability Matrix Tool](#).

Storage types supported by Plug-in for Oracle Database

SnapCenter supports a wide range of storage types on both physical and virtual machines. You must verify the support for your storage type before installing the SnapCenter Plug-ins Package for Linux or SnapCenter Plug-ins Package for AIX.

SnapCenter does not support storage provisioning for Linux and AIX.


Storage types supported on Linux

The following table lists the storage types supported on Linux.

Machine	Storage type
Physical server	<ul style="list-style-type: none"> • FC-connected LUNs • iSCSI-connected LUNs • NFS-connected volumes
VMware ESXi	<ul style="list-style-type: none"> • RDM LUNs connected by an FC or iSCSI ESXi HBAScanning of host bus adapters (HBAs) might take long time to complete because SnapCenter scans all the host bus adaptors present in the host. <p>You can edit the LinuxConfig.pm file located at <code>/opt/NetApp/snapcenter/spl/plugins/scu/scucore/modules/SCU/Config</code> to set the value of the SCSI_HOSTS_OPTIMIZED_RESCAN parameter to 1 to rescan only those HBAs that are listed in <code>HBA_DRIVER_NAMES</code>.</p> <ul style="list-style-type: none"> • iSCSI LUNs connected directly to the guest system by the iSCSI initiator • VMDKs on VMFS or NFS datastores • NFS volumes connected directly to the guest system

Storage types supported on AIX

The following table lists the storage types supported on AIX.

Machine	Storage type
Physical server	<ul style="list-style-type: none">• FC-connected and iSCSI-connected LUNs. <p>In a SAN environment, ASM, LVM, and SAN file systems are supported.</p> <div> NFS on AIX and filesystem is not supported.</div> <ul style="list-style-type: none">• Enhanced Journaled File System (JFS2) <p>Supports inline logging on SAN filesystems and LVM layout.</p>

The [NetApp Interoperability Matrix Tool](#) contains the latest information about the supported versions.

Prepare storage systems for SnapMirror and SnapVault replication for Plug-in for Oracle

You can use a SnapCenter plug-in with ONTAP SnapMirror technology to create mirror copies of backup sets on another volume, and with ONTAP SnapVault technology to perform disk-to-disk backup replication for standards compliance and other governance-related purposes. Before you perform these tasks, you must configure a data-protection relationship between the source and destination volumes and initialize the relationship.

SnapCenter performs the updates to SnapMirror and SnapVault after it completes the Snapshot copy operation. SnapMirror and SnapVault updates are performed as part of the SnapCenter job; do not create a separate ONTAP schedule.



If you are coming to SnapCenter from a NetApp SnapManager product and are satisfied with the data protection relationships you have configured, you can skip this section.

A data protection relationship replicates data on primary storage (the source volume) to secondary storage (the destination volume). When you initialize the relationship, ONTAP transfers the data blocks referenced on the source volume to the destination volume.



SnapCenter does not support cascade relationships between SnapMirror and SnapVault volumes (**Primary** > **Mirror** > **Vault**). You should use fanout relationships.

SnapCenter supports the management of version-flexible SnapMirror relationships. For details about version-flexible SnapMirror relationships and how to set them up, see the [ONTAP documentation](#).



SnapCenter does not support **sync_mirror** replication.

Minimum ONTAP privileges required for Plug-in for Oracle

The minimum ONTAP privileges that are required vary according to the SnapCenter plug-ins you are using for data protection.

All-access commands: Minimum privileges required for ONTAP 8.3.0 and later
<ul style="list-style-type: none">• event generate-autosupport-log
<ul style="list-style-type: none">• job history show• job stop
<ul style="list-style-type: none">• lun• lun attribute show• lun create• lun delete• lun geometry• lun igroup add• lun igroup create• lun igroup delete• lun igroup rename• lun igroup show• lun mapping add-reporting-nodes• lun mapping create• lun mapping delete• lun mapping remove-reporting-nodes• lun mapping show• lun modify• lun move-in-volume• lun offline• lun online• lun persistent-reservation clear• lun resize• lun serial• lun show

All-access commands: Minimum privileges required for ONTAP 8.3.0 and later

- snapmirror policy add-rule
- snapmirror policy modify-rule
- snapmirror policy remove-rule
- snapmirror policy show
- snapmirror restore
- snapmirror show
- snapmirror show-history
- snapmirror update
- snapmirror update-ls-set
- snapmirror list-destinations

- version

All-access commands: Minimum privileges required for ONTAP 8.3.0 and later

- volume clone create
- volume clone show
- volume clone split start
- volume clone split stop
- volume create
- volume destroy
- volume file clone create
- volume file show-disk-usage
- volume offline
- volume online
- volume modify
- volume qtree create
- volume qtree delete
- volume qtree modify
- volume qtree show
- volume restrict
- volume show
- volume snapshot create
- volume snapshot delete
- volume snapshot modify
- volume snapshot rename
- volume snapshot restore
- volume snapshot restore-file
- volume snapshot show
- volume unmount

- vserver
- vserver cifs
- vserver cifs shadowcopy show
- vserver show


- network interface
- network interface show

- metrocluster show

Install SnapCenter Plug-in for Oracle Database

Installation workflow of SnapCenter Plug-in for Oracle Database

You should install and set up the SnapCenter Plug-in for Oracle Database if you want to protect Oracle databases.



Prerequisites for adding hosts and installing Plug-ins Package for Linux or AIX

Before you add a host and install the plug-ins packages, you must complete all the requirements.

- If you are using iSCSI, the iSCSI service must be running.
- You must have enabled the password-based SSH connection for the root or non-root user.

SnapCenter Plug-in for Oracle Database can be installed by a non-root user. However, you should configure the sudo privileges for the non-root user to install and start the plug-in process. After installing the plug-in, the processes will be running as an effective root user.

- If you are installing the SnapCenter Plug-ins Package for AIX on AIX host, you should have manually resolved the directory level symbolic links.

The SnapCenter Plug-ins Package for AIX automatically resolves the file level symbolic link but not the directory level symbolic links to obtain the JAVA_HOME absolute path.

- Create credentials with authentication mode as Linux or AIX for the install user.
- You must have installed Java 1.8.x, 64-bit, on your Linux or AIX host.

For information to download JAVA, see:

- [Java Downloads for All Operating Systems](#)
- [IBM Java for AIX](#)

- For Oracle databases that are running on a Linux or AIX host, you should install both SnapCenter Plug-in for Oracle Database and SnapCenter Plug-in for UNIX.



You can use the Plug-in for Oracle Database to manage Oracle databases for SAP as well. However, SAP BR*Tools integration is not supported.

- If you are using Oracle database 11.2.0.3 or later, you must install the 13366202 Oracle patch.



UUID mapping in the /etc/fstab file is not supported by SnapCenter.

Linux Host requirements

You should ensure that the host meets the requirements before installing the SnapCenter Plug-ins Package for Linux.

Item	Requirements
Operating systems	<ul style="list-style-type: none"> • Red Hat Enterprise Linux • Oracle Linux <div> <p>If you are using Oracle database on LVM in Oracle Linux or Red Hat Enterprise Linux 6.6 or 7.0 operating systems, you must install the latest version of Logical Volume Manager (LVM).</p> </div> <ul style="list-style-type: none"> • SUSE Linux Enterprise Server (SLES)
Minimum RAM for the SnapCenter plug-in on host	1 GB
Minimum install and log space for the SnapCenter plug-in on host	2 GB <div> <p>You should allocate sufficient disk space and monitor the storage consumption by the logs folder. The log space required varies depending on the number of the entities to be protected and the frequency of data protection operations. If there is no sufficient disk space, the logs will not be created for the recently run operations.</p> </div>
Required software packages	Java 1.8.x (64-bit) Oracle Java and OpenJDK flavors <p>If you have upgraded JAVA to the latest version, you must ensure that the JAVA_HOME option located at /var/opt/snapcenter/spl/etc/spl.properties is set to the correct JAVA version and the correct path.</p>

For the latest information about supported versions, see the [NetApp Interoperability Matrix Tool](#).

Configure sudo privileges for non-root users for Linux host

SnapCenter 2.0 and later releases allow a non-root user to install the SnapCenter Plug-ins Package for Linux

and to start the plug-in process. You should configure sudo privileges for the non-root user to provide access to several paths.

What you will need

- Sudo version between 1.8.7 and 1.8.19P2.
- Ensure that the non-root user is part of the Oracle installation group.
- Edit the `/etc/ssh/sshd_config` file to configure the message authentication code algorithms: MACs hmac-sha2-256 and MACs hmac-sha2-512.

Restart the sshd service after updating the configuration file.

Example:

```
#Port 22
#AddressFamily any
#ListenAddress 0.0.0.0
#ListenAddress ::
#Legacy changes
#KexAlgorithms diffie-hellman-group1-sha1
#Ciphers aes128-cbc
#The default requires explicit activation of protocol
Protocol 2
HostKey/etc/ssh/ssh_host_rsa_key
MACs hmac-sha2-256
```

About this task

You should configure sudo privileges for the non-root user to provide access to the following paths:

- `/home/SUDO_USER/.sc_netapp/snapcenter_linux_host_plugin.bin`
- `/custom_location/NetApp/snapcenter/spl/installation/plugins/uninstall`
- `/custom_location/NetApp/snapcenter/spl/bin/spl`



If you are managing a RAC setup, the non-root user should be an Oracle user and it cannot be just any non-root OS user.

Steps

1. Log in to the Linux host on which you want to install the SnapCenter Plug-ins Package for Linux.
2. Add the following lines to the `/etc/sudoers` file by using the visudo Linux utility.

```

Cmnd_Alias SCCMD = sha224:checksum_value== /home/
SUDO_USER/.sc_netapp/snapcenter_linux_host_plugin.bin,
/opt/NetApp/snapcenter/spl/installation/plugins/uninstall,
/opt/NetApp/snapcenter/spl/bin/spl
Cmnd_Alias PRECHECKCMD = sha224:checksum_value== /home/
SUDO_USER/.sc_netapp/Linux_Prechecks.sh
SUDO_USER ALL=(ALL) NOPASSWD:SETENV: SCCMD, PRECHECKCMD
Defaults: SUDO_USER env_keep=JAVA_HOME
Defaults: SUDO_USER !visiblepw
Defaults: SUDO_USER !requiretty

```

SUDO_USER is the name of the non-root user that you created.

You can obtain the checksum value from the **oracle_checksum.txt** file, which is located at *C:\ProgramData\NetApp\SnapCenter\Package Repository*.

If you have specified a custom location, the location will be *custom_path\NetApp\SnapCenter\Package Repository*.



The example should be used only as a reference for creating your own data.

Best Practice: For security reasons, you should remove the sudo entry after completing every installation or upgrade.


AIX Host requirements

You should ensure that the host meets the requirements before installing the SnapCenter Plug-ins Package for AIX.



SnapCenter Plug-in for UNIX which is part of the SnapCenter Plug-ins Package for AIX, does not support concurrent volume groups.

Item	Requirements
Operating systems	AIX 6.1 or later
Minimum RAM for the SnapCenter plug-in on host	4 GB

Item	Requirements
Minimum install and log space for the SnapCenter plug-in on host	1 GB <div>  <p>You should allocate sufficient disk space and monitor the storage consumption by the logs folder. The log space required varies depending on the number of the entities to be protected and the frequency of data protection operations. If there is no sufficient disk space, the logs will not be created for the recently run operations.</p> </div>
Required software packages	Java 1.8.x (64-bit)IBM Java <p>If you have upgraded JAVA to the latest version, you must ensure that the JAVA_HOME option located at <code>/var/opt/snapcenter/spl/etc/spl.properties</code> is set to the correct JAVA version and the correct path.</p>

For the latest information about supported versions, see the [NetApp Interoperability Matrix Tool](#).

Configure sudo privileges for non-root users for AIX host

SnapCenter 4.4 and later allows a non-root user to install the SnapCenter Plug-ins Package for AIX and to start the plug-in process. You should configure sudo privileges for the non-root user to provide access to several paths.

What you will need

- Sudo version between 1.8.7 and 1.8.19P2.
- Ensure that the non-root user is part of the Oracle installation group.
- Edit the `/etc/ssh/sshd_config` file to configure the message authentication code algorithms: MACs hmac-sha2-256 and MACs hmac-sha2-512.

Restart the sshd service after updating the configuration file.

Example:

```
#Port 22
#AddressFamily any
#ListenAddress 0.0.0.0
#ListenAddress ::
#Legacy changes
#KexAlgorithms diffie-hellman-group1-sha1
#Ciphers aes128-cbc
#The default requires explicit activation of protocol
Protocol 2
HostKey/etc/ssh/ssh_host_rsa_key
MACs hmac-sha2-256
```

About this task

You should configure sudo privileges for the non-root user to provide access to the following paths:

- /home/AIX_USER/.sc_netapp/snapcenter_aix_host_plugin.bsx
- /custom_location/NetApp/snapcenter/spl/installation/plugins/uninstall
- /custom_location/NetApp/snapcenter/spl/bin/spl



If you are managing a RAC setup, the non-root user should be an Oracle user and it cannot be just any non-root OS user.

Steps

1. Log in to the AIX host on which you want to install the SnapCenter Plug-ins Package for AIX.
2. Add the following lines to the /etc/sudoers file by using the visudo Linux utility.

```
Cmnd_Alias SCCMD = sha224:checksum_value== /home/
AIX_USER/.sc_netapp/snapcenter_aix_host_plugin.bsx,
/opt/NetApp/snapcenter/spl/installation/plugins/uninstall,
/opt/NetApp/snapcenter/spl/bin/spl
Cmnd_Alias PRECHECKCMD = sha224:checksum_value== /home/
AIX_USER/.sc_netapp/AIX_Prechecks.sh
AIX_USER ALL=(ALL) NOPASSWD:SETENV: SCCMD, PRECHECKCMD
Defaults: AIX_USER !visiblepw
Defaults: AIX_USER !requiretty
```

AIX_USER is the name of the non-root user that you created.

You can obtain the checksum value from the **oracle_checksum.txt** file, which is located at *C:\ProgramData\NetApp\SnapCenter\Package Repository*.

If you have specified a custom location, the location will be *custom_path\NetApp\SnapCenter\Package Repository*.



The example should be used only as a reference for creating your own data.

Best Practice: For security reasons, you should remove the sudo entry after completing every installation or upgrade.

Set up credentials

SnapCenter uses credentials to authenticate users for SnapCenter operations. You should create credentials for installing the plug-in package on Linux or AIX hosts.

About this task

The credentials are created either for the root user or for a non-root user who has sudo privileges to install and start the plug-in process.

For information, see: [Configure sudo privileges for non-root users for Linux host](#) or [Configure sudo privileges for non-root users for AIX host](#)

Best Practice: Although you are allowed to create credentials after deploying hosts and installing plug-ins, the best practice is to create credentials after you add SVMs, before you deploy hosts and install plug-ins.

Steps

1. In the left navigation pane, click **Settings**.
2. In the Settings page, click **Credential**.
3. Click **New**.
4. In the Credential page, enter the credential information:

For this field...	Do this...
Credential name	Enter a name for the credentials.

For this field...	Do this...
User name/Password	<p>Enter the user name and password that are to be used for authentication.</p> <ul style="list-style-type: none"> Domain administrator <p>Specify the domain administrator on the system on which you are installing the SnapCenter plug-in. Valid formats for the Username field are:</p> <ul style="list-style-type: none"> <i>NetBIOS\UserName</i> <i>Domain FQDN\UserName</i> Local administrator (for workgroups only) <p>For systems that belong to a workgroup, specify the built-in local administrator on the system on which you are installing the SnapCenter plug-in. You can specify a local user account that belongs to the local administrators group if the user account has elevated privileges or the User Access control feature is disabled on the host system. The valid format for the Username field is: <i>UserName</i></p>
Authentication Mode	<p>Select the authentication mode that you want to use.</p> <p>Depending on the operating system of the plug-in host, select either Linux or AIX.</p>
Use sudo privileges	<p>Select the Use sudo privileges check box if you are creating credentials for a non-root user.</p>

5. Click **OK**.

After you finish setting up the credentials, you might want to assign credential maintenance to a user or group of users on the **User and Access** page.

Configure credentials for an Oracle database

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

About this task

You should review the different authentication methods supported for Oracle database. For information, see [Authentication methods for your credentials](#).


If you set up credentials for individual resource groups and the user name does not have full admin privileges, the user name must at least have resource group and backup privileges.


If you have enabled Oracle database authentication, a red padlock icon is shown in the resources view. You must configure database credentials to be able to protect the database or add it to the resource group to perform data protection operations.



If you specify incorrect details while creating a credential, an error message is displayed. You must click **Cancel**, and then retry.

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 **Database** from the **View** list.
3. Click , and then select the host name and the database type to filter the resources.

You can then click  to close the filter pane.

4. Select the database, and then click **Database Settings > Configure Database**.
5. In the Configure database settings section, from the **Use existing Credential** drop-down list, select the credential that should be used to perform data protection jobs on the Oracle database.




The Oracle user should have sysdba privileges.

You can also create a credential by clicking .


6. In the Configure ASM settings section, from the **Use existing Credential** drop-down list, select the credential that should be used to perform data protection jobs on the ASM instance.



The ASM user should have sysasm privilege.

You can also create a credential by clicking .

7. In the Configure RMAN catalog settings section, from the **Use existing credential** drop-down list, select the credential that should be used to perform data protection jobs on the Oracle Recovery Manager (RMAN) catalog database.

You can also create a credential by clicking .

In the **TNSName** field, enter the Transparent Network Substrate (TNS) file name that will be used by the SnapCenter Server to communicate with the database.

8. In the **Preferred RAC Nodes** field, specify the Real Application Cluster (RAC) nodes preferred for backup.

The preferred nodes might be one or all cluster nodes where the RAC database instances are present. The backup operation is triggered only on these preferred nodes in the order of preference.

In RAC One Node, only one node is listed in the preferred nodes, and this preferred node is the node where the database is currently hosted.

After failover or relocation of RAC One Node database, refreshing of resources in the SnapCenter Resources page will remove the host from the **Preferred RAC Nodes** list where the database was earlier hosted. The RAC node where the database is relocated will be listed in **RAC Nodes** and will need to be manually configured as the preferred RAC node.

For more information, see [Preferred nodes in RAC setup](#).

9. Click **OK**.

Add hosts and install Plug-ins Package for Linux or AIX using GUI

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.

About this task

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 of the nodes of the cluster. For Oracle RAC One Node, you should install the plug-in on both active and passive nodes.

You should be assigned to a role that has the plug-in install and uninstall permissions, such as the SnapCenter Admin role.





You cannot add a SnapCenter Server as a plug-in host to another SnapCenter Server.

Steps


1. In the left navigation pane, click **Hosts**.
2. Verify that the **Managed Hosts** tab is selected at the top.
3. Click **Add**.
4. In the Hosts page, perform the following actions:

For this field...	Do this...
Host Type	<p>Select Linux or AIX as the host type.</p> <p>The SnapCenter Server adds the host, and then installs the Plug-in for Oracle Database and the Plug-in for UNIX if the plug-ins are not already installed on the host.</p>

For this field...	Do this...
Host name	<p>Enter the fully qualified domain name (FQDN) or the IP address of the host.</p> <p>SnapCenter depends on the proper configuration of the DNS. Therefore, the best practice is to enter the FQDN.</p> <p>You can enter the IP addresses or FQDN of one of the following:</p> <ul style="list-style-type: none"> • Stand-alone host • Any node in the Oracle Real Application Clusters (RAC) environment <div>  <p>Node VIP or scan IP is not supported</p> </div> <p>If you are adding a host by using SnapCenter and the host is part of a subdomain, you must provide the FQDN.</p>
Credentials	<p>Either select the credential name that you created or create new credentials.</p> <p>The credential must have administrative rights on the remote host. For details, see the information about creating credentials.</p> <p>You can view details about the credentials by positioning the cursor over the credential name that you specified.</p> <div>  <p>The credentials authentication mode is determined by the host type that you specify in the Add Host wizard.</p> </div>


5. In the Select Plug-ins to Install section, select the plug-ins to install.

6. (Optional) Click **More Options**.

For this field...	Do this...
Port	<p>Either retain the default port number or specify the port number.</p> <p>The default port number is 8145. If the SnapCenter Server was installed on a custom port, that port number will be displayed as the default port.</p> <div>  <p>If you manually installed the plug-ins and specified a custom port, you must specify the same port. Otherwise, the operation fails.</p> </div>
Installation Path	<p>The default path is <i>/opt/NetApp/snapcenter</i>.</p> <p>You can optionally customize the path.</p>
Add all hosts in the Oracle RAC	<p>Select this check box to add all the cluster nodes in an Oracle RAC.</p> <p>In a Flex ASM setup, all the nodes irrespective of whether it is a Hub or Leaf node, will be added.</p>
Skip preinstall checks	<p>Select this check box if you already installed the plug-ins manually and you do not want to validate whether the host meets the requirements for installing the plug-in.</p>

7. Click **Submit**.

If you have not selected the Skip prechecks checkbox, the host is validated to verify whether the host meets the requirements for installing the plug-in.



The precheck script does not validate the plug-in port firewall status if it is specified in the firewall reject rules.


Appropriate error or warning messages are displayed if the minimum requirements are not met. If the error is related to disk space or RAM, you can update the web.config file located at *C:\Program Files\NetApp\SnapCenter WebApp* to modify the default values. If the error is related to other parameters, you should fix the issue.



In an HA setup, if you are updating web.config file, you must update the file on both nodes.

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.



SnapCenter does not support ECDSA algorithm.



Fingerprint verification is mandatory even if the same host was added earlier to SnapCenter and the fingerprint was confirmed.

9. Monitor the installation progress.

The installation-specific log files are located at `/custom_location/snapcenter/logs`.

After you finish






All the databases on the host are automatically discovered and displayed in the Resources page. If nothing is displayed, click **Refresh Resources**.

Monitor installation status

You can monitor the progress of SnapCenter plug-in package installation by using the Jobs page. You might want to check the progress of installation to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page and indicate the state of the operation:

-  In progress
-  Completed successfully
-  Failed
-  Completed with warnings or could not start due to warnings
-  Queued

Steps

1. In the left navigation pane, click **Monitor**.
2. In the Monitor page, click **Jobs**.
3. In the Jobs page, to filter the list so that only plug-in installation operations are listed, do the following:
 - a. Click **Filter**.
 - b. Optional: Specify the start and end date.
 - c. From the Type drop-down menu, select **Plug-in installation**.
 - d. From the Status drop-down menu, select the installation status.
 - e. Click **Apply**.
4. Select the installation job and click **Details** to view the job details.
5. In the Job Details page, click **View logs**.

Alternate ways to install Plug-ins Package for Linux or AIX

Install on multiple remote hosts using cmdlets

You should use the `Install-SmHostPackage` PowerShell cmdlet to install the SnapCenter Plug-ins Package for Linux or SnapCenter Plug-ins Package for AIX on multiple hosts.

What you will need

You should be logged in to SnapCenter as a domain user with local administrator rights on each host on which you want to install the plug-in package.

Steps

1. Launch PowerShell.
2. On the SnapCenter Server host, establish a session using the *Open-SmConnection* cmdlet, and then enter your credentials.
3. Install the SnapCenter Plug-ins Package for Linux or SnapCenter Plug-ins Package for AIX using the *Install-SmHostPackage* cmdlet and the required parameters.

You can use the *-skipprecheck* option when you have already installed the plug-ins manually and do not want to validate whether the host meets the requirements to install the plug-in.



The precheck script does not validate the plug-in port firewall status if it is specified in the firewall reject rules.

4. Enter your credentials for remote installation.

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

Install on cluster host

You should install SnapCenter Plug-ins Package for Linux or SnapCenter Plug-ins Package for AIX on both the nodes of the cluster host.

Each of the nodes of the cluster host has two IPs. One of the IPs will be the public IP of the respective nodes and the second IP will be the cluster IP that is shared between both the nodes.

Steps

1. Install SnapCenter Plug-ins Package for Linux or SnapCenter Plug-ins Package for AIX on both the nodes of the cluster host.
2. Validate that the correct values for `SNAPCENTER_SERVER_HOST`, `SPL_PORT`, `SNAPCENTER_SERVER_PORT`, and `SPL_ENABLED_PLUGINS` parameters are specified in the `spl.properties` file located at `/var/opt/snapcenter/spl/etc/`.

If `SPL_ENABLED_PLUGINS` is not specified in `spl.properties`, you can add it and assign the value `SCO,SCU`.

3. On the SnapCenter Server host, establish a session using the *Open-SmConnection* cmdlet, and then enter your credentials.
4. In each of the nodes, set the preferred IPs of the node using the *Set-PreferredHostIPsInStorageExportPolicy* sccli command and the required parameters.
5. In the SnapCenter Server host, add an entry for the cluster IP and corresponding DNS name in `C:\Windows\System32\drivers\etc\hosts`.
6. Add the node to the SnapCenter Server using the *Add-SmHost* cmdlet by specifying the cluster IP for the host name.

Discover the Oracle database on node 1 (assuming the cluster IP is hosted on node 1) and create a backup of the database. If a failover happens, you can use the backup created on node 1 to restore the database on node 2. You can also use the backup created on node 1 to create a clone on node 2.



There will be stale volumes, directories, and lock file if the failover happens while any other SnapCenter operations are running.

Install Plug-ins Package for Linux in silent mode or console mode

You can install the SnapCenter Plug-ins Package for Linux either in console mode or in silent mode by using the command-line interface (CLI).

What you will need

- You should review the prerequisites for installing the plug-ins package.
- You should ensure that the DISPLAY environment variable is not set.

If the DISPLAY environment variable is set, you should run `unset DISPLAY`, and then try to manually install the plug-in.

About this task

You are required to provide the necessary installation information while installing in console mode, whereas in silent mode installation you do not have to provide any installation information.

Steps

1. Download the SnapCenter Plug-ins Package for Linux from the SnapCenter Server installation location.

The default installation path is *C:\ProgramData\NetApp\SnapCenter\PackageRepository*. This path is accessible from the host where the SnapCenter Server is installed.
2. From the command prompt, navigate to the directory where you downloaded the installation file.
3. Depending on your preferred mode of installation, perform one of the following step.

Install mode	Steps
Console mode	<div>a. Run: <pre>./SnapCenter_linux_host_plugin.bin -i console</pre></div> <div>b. Follow the on-screen prompts to complete the installation.</div>

Install mode	Steps
Silent mode	<p>Run:</p> <pre>./SnapCenter_linux_host_plugin.bin-i silent-DPORT=8145- DSERVER_IP=SnapCenter_Server_FQDN- DSERVER_HTTPS_PORT=SnapCenter_Server_P ort- DUSER_INSTALL_DIR==/opt/custom_path</pre>

4. Edit the `spl.properties` file located at `/var/opt/snapcenter/spl/etc/` to add `SPL_ENABLED_PLUGINS=SCO,SCU`, and then restart the SnapCenter Plug-in Loader service.



The installation of the plug-ins package registers the plug-ins on the host and not on the SnapCenter Server. You should register the plug-ins on the SnapCenter Server by adding the host using the SnapCenter GUI or PowerShell cmdlet. While adding the host, select “None” as the credential. After the host is added, the installed plug-ins are automatically discovered.

Install Plug-ins Package for AIX in silent mode

You can install the SnapCenter Plug-ins Package for AIX in silent mode by using the command-line interface (CLI).

What you will need

- You should review the prerequisites for installing the plug-ins package.
- You should ensure that the `DISPLAY` environment variable is not set.

If the `DISPLAY` environment variable is set, you should run `unset DISPLAY`, and then try to manually install the plug-in.

Steps

1. Download the SnapCenter Plug-ins Package for AIX from the SnapCenter Server installation location.

The default installation path is `C:\ProgramData\NetApp\SnapCenter\PackageRepository`. This path is accessible from the host where the SnapCenter Server is installed.

2. From the command prompt, navigate to the directory where you downloaded the installation file.
3. Run

```
./snapcenter_aix_host_plugin.bsx-i silent-DPORT=8145-
DSERVER_IP=SnapCenter_Server_FQDN-DSERVER_HTTPS_PORT=SnapCenter_Server_Port-
DUSER_INSTALL_DIR==/opt/custom_path-
DINSTALL_LOG_NAME=SnapCenter_AIX_Host_Plug-in_Install_MANUAL.log-
DCHOSEN_FEATURE_LIST=CUSTOMDSPL_USER=install_user
```

4. Edit the `spl.properties` file located at `/var/opt/snapcenter/spl/etc/` to add `SPL_ENABLED_PLUGINS=SCO,SCU`, and then restart the SnapCenter Plug-in Loader service.



The installation of the plug-ins package registers the plug-ins on the host and not on the SnapCenter Server. You should register the plug-ins on the SnapCenter Server by adding the host using the SnapCenter GUI or PowerShell cmdlet. While adding the host, select “None” as the credential. After the host is added, the installed plug-ins are automatically discovered.

Configure the SnapCenter Plug-in Loader service

The SnapCenter Plug-in Loader service loads the plug-in package for Linux or AIX to interact with the SnapCenter Server. The SnapCenter Plug-in Loader service is installed when you install the SnapCenter Plug-ins Package for Linux or SnapCenter Plug-ins Package for AIX.



About this task

After installing the SnapCenter Plug-ins Package for Linux or SnapCenter Plug-ins Package for AIX, the SnapCenter Plug-in Loader service starts automatically. If the SnapCenter Plug-in Loader service fails to start automatically, you should:

- Ensure that the directory where the plug-in is operating is not deleted
- Increase the memory space allotted to the Java Virtual Machine

The `spl.properties` file, which is located at `/custom_location/NetApp/snapcenter/spl/etc/`, contains the following parameters. Default values are assigned to these parameters.

Parameter name	Description
LOG_LEVEL	Displays the log levels that are supported. The possible values are INFO, DEBUG, TRACE, ERROR, FATAL, and WARN.
SPL_PROTOCOL	Displays the protocol that is supported by SnapCenter Plug-in Loader. Only the HTTPS protocol is supported. You can add the value if the default value is missing.
SNAPCENTER_SERVER_PROTOCOL	Displays the protocol that is supported by SnapCenter Server. Only the HTTPS protocol is supported. You can add the value if the default value is missing.
SKIP_JAVAHOME_UPDATE	By default, the SPL service detects the java path and update JAVA_HOME parameter. Therefore the default value is set to FALSE. You can set to TRUE if you want to disable the default behavior and manually fix the java path.

Parameter name	Description
SPL_KEYSTORE_PASS	<p>Displays the password of the keystore file.</p> <p>You can change this value only if you change the password or create a new keystore file.</p>
SPL_PORT	<p>Displays the port number on which the SnapCenter Plug-in Loader service is running.</p> <p>You can add the value if the default value is missing.</p> <div>  <p>You should not change the value after installing the plug-ins.</p> </div>
SNAPCENTER_SERVER_HOST	Displays the IP address or host name of the SnapCenter Server.
SPL_KEYSTORE_PATH	Displays the absolute path of the keystore file.
SNAPCENTER_SERVER_PORT	Displays the port number on which the SnapCenter Server is running.
LOGS_MAX_COUNT	<p>Displays the number of SnapCenter Plug-in Loader log files that are retained in the <i>/custom_location/snapcenter/spl/logs</i> folder.</p> <p>The default value is set to 5000. If the count is more than the specified value, then the last 5000 modified files are retained. The check for the number of files is done automatically every 24 hours from when SnapCenter Plug-in Loader service is started.</p> <div>  <p>If you manually delete the <i>spl.properties</i> file, then the number of files to be retained is set to 9999.</p> </div>
JAVA_HOME	<p>Displays the absolute directory path of the JAVA_HOME which is used to start SPL service.</p> <p>This path is determined during installation and as part of starting SPL.</p>
LOG_MAX_SIZE	<p>Displays the maximum size of the job log file.</p> <p>Once the maximum size is reached, the log file is zipped, and the logs are written into the new file of that job.</p>

Parameter name	Description
RETAIN_LOGS_OF_LAST_DAYS	Displays the number of days up to which the logs are retained.
ENABLE_CERTIFICATE_VALIDATION	<p>Displays true when CA certificate validation is enabled for the host.</p> <p>You can enable or disable this parameter either by editing the spl.properties or by using the SnapCenter GUI or cmdlet.</p>

If any of these parameters are not assigned to the default value or if you want to assign or change the value, then you can modify the spl.properties file. You can also verify the spl.properties file and edit the file to troubleshoot any issues related to the values that are assigned to the parameters. After you modify the spl.properties file, you should restart the SnapCenter Plug-in Loader service.

Steps

1. Perform one of the following actions, as required:

- Start the SnapCenter Plug-in Loader service as a root user:

```
`/custom_location/NetApp/snapcenter/spl/bin/spl start`
```

- Stop the SnapCenter Plug-in Loader service:

```
`/custom_location/NetApp/snapcenter/spl/bin/spl stop`
```



You can use the -force option with the stop command to stop the SnapCenter Plug-in Loader service forcefully. However, you should use caution before doing so because it also terminates the existing operations.

- Restart the SnapCenter Plug-in Loader service:

```
`/custom_location/NetApp/snapcenter/spl/bin/spl restart`
```

- Find the status of the SnapCenter Plug-in Loader service:

```
`/custom_location/NetApp/snapcenter/spl/bin/spl status`
```

- Find the change in the SnapCenter Plug-in Loader service:

```
`/custom_location/NetApp/snapcenter/spl/bin/spl change`
```

Configure CA certificate with SnapCenter Plug-in Loader (SPL) service on Linux host

You should manage the password of SPL keystore and its certificate, configure the CA certificate, configure root or intermediate certificates to SPL trust-store, and configure CA signed key pair to SPL trust-store with SnapCenter Plug-in Loader service to activate the installed digital certificate.



SPL uses the file 'keystore.jks', which is located at '/var/opt/snapcenter/spl/etc' both as its trust-store and key-store.

Manage password for SPL keystore and alias of the CA signed key pair in use

Steps

1. You can retrieve SPL keystore default password from SPL property file.

It is the value corresponding to the key 'SPL_KEYSTORE_PASS'.

2. Change the keystore password:

```
keytool -storepasswd -keystore keystore.jks
```

3. Change the password for all aliases of private key entries in the keystore to the same password used for the keystore:

```
keytool -keypasswd -alias "<alias_name>" -keystore keystore.jks
```

Update the same for the key SPL_KEYSTORE_PASS in spl.properties file.

4. Restart the service after changing the password.



Password for SPL keystore and for all the associated alias password of the private key should be same.

Configure root or intermediate certificates to SPL trust-store

You should configure the root or intermediate certificates without the private key to SPL trust-store.

Steps

1. Navigate to the folder containing the SPL keystore: */var/opt/snapcenter/spl/etc*.
2. Locate the file 'keystore.jks'.
3. List the added certificates in the keystore:

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

4. Add a root or intermediate certificate:

```
keytool -import -trustcacerts -alias <AliasNameForCertificateToBeImported>
-file /<CertificatePath> -keystore keystore.jks
```

5. Restart the service after configuring the root or intermediate certificates to SPL trust-store.



You should add the root CA certificate and then the intermediate CA certificates.

Configure CA signed key pair to SPL trust-store

You should configure the CA signed key pair to the SPL trust-store.

Steps

1. Navigate to the folder containing the SPL's keystore `/var/opt/snapcenter/spl/etc`.
2. Locate the file 'keystore.jks'.
3. List the added certificates in the keystore:

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

4. Add the CA certificate having both private and public key.

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

5. List the added certificates in the keystore.

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

6. Verify that the keystore contains the alias corresponding to the new CA certificate, which was added to the keystore.
7. Change the added private key password for CA certificate to the keystore password.

Default SPL keystore password is the value of the key `SPL_KEYSTORE_PASS` in `spl.properties` file.

```
keytool -keypasswd -alias "<aliasNameOfAddedCertInKeystore>" -keystore
keystore.jks
```

8. If the alias name in the CA certificate is long and contains space or special characters (`"**", ",", "`), change the alias name to a simple name:

```
keytool -changealias -alias "<OriginalAliasName>" -destalias  
"<NewAliasName>" -keystore keystore.jks
```

9. Configure the alias name from the keystore located in `spl.properties` file.

Update this value against the key `SPL_CERTIFICATE_ALIAS`.

10. Restart the service after configuring the CA signed key pair to SPL trust-store.

Configure certificate revocation list (CRL) for SPL

You should configure the CRL for SPL

About this task

- SPL will look for the CRL files in a pre-configured directory.
- Default directory for the CRL files for SPL is `/var/opt/snapcenter/spl/etc/crl`.

Steps

1. You can modify and update the default directory in `spl.properties` file against the key `SPL_CRL_PATH`.
2. You can place more than one CRL file in this directory.

The incoming certificates will be verified against each CRL.

Enable CA Certificates for plug-ins

You should configure the CA certificates and deploy the CA certificates in the SnapCenter Server and the corresponding plug-in hosts. You should enable the CA certificate validation for the plug-ins.

What you will need

- You can enable or disable the CA certificates using the run `Set-SmCertificateSettings` cmdlet.
- You can display the certificate status for the plug-ins using the `Get-SmCertificateSettings`.





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

Steps

1. In the left navigation pane, click **Hosts**.
2. In the Hosts page, click **Managed Hosts**.
3. Select single or multiple plug-in hosts.
4. Click **More options**.
5. Select **Enable Certificate Validation**.

After you finish

The Managed Hosts tab host displays a padlock and the color of the padlock indicates the status of the connection between SnapCenter Server and the plug-in host.

-  indicates that the CA certificate is neither enabled nor assigned to the plug-in host.
-  indicates that the CA certificate is successfully validated.
-  indicates that the CA certificate could not be validated.
-  indicates that the connection information could not be retrieved.



When the status is yellow or green, the data protection operations completes successfully.

Import data from SnapManager for Oracle and SnapManager for SAP to SnapCenter

Importing data from SnapManager for Oracle and SnapManager for SAP to SnapCenter enables you to continue to use your data from previous versions.

You can import data from SnapManager for Oracle and SnapManager for SAP to SnapCenter by running the import tool from the command-line interface (Linux host CLI).

The import tool creates policies and resource groups in SnapCenter. The policies and resource groups created in SnapCenter correspond to the profiles and operations performed using those profiles in SnapManager for Oracle and SnapManager for SAP. The SnapCenter import tool interacts with the SnapManager for Oracle and SnapManager for SAP repository databases and the database that you want to import.

- Retrieves all the profiles, schedules, and operations performed using the profiles.
- Creates a SnapCenter backup policy for each unique operation and each schedule attached to a profile.
- Creates a resource group for each target database.

You can run the import tool by executing the `sc-migrate` script located at `/opt/NetApp/snapcenter/spl/bin`. When you install the SnapCenter Plug-ins Package for Linux on the database host that you want to import, the `sc-migrate` script is copied to `/opt/NetApp/snapcenter/spl/bin`.



Importing data is not supported from SnapCenter graphical user interface (GUI).

SnapCenter does not support Data ONTAP operating in 7-Mode. You can use the 7-Mode Transition Tool to migrate data and configurations that are stored on a system running Data ONTAP operating in 7-Mode to an ONTAP system.

Configurations supported for importing data

Before you import data from SnapManager 3.4.x for Oracle and SnapManager 3.4.x for SAP to SnapCenter, you should be aware of the configurations that are supported with the SnapCenter Plug-in for Oracle Database.

The configurations that are supported with the SnapCenter Plug-in for Oracle Database are listed in the [NetApp Interoperability Matrix Tool](#).

What gets imported to SnapCenter

You can import profiles, schedules, and operations performed using the profiles.

From SnapManager for Oracle and SnapManager for SAP	To SnapCenter
Profiles without any operations and schedules	A policy is created with default backup type as Online and backup scope as Full.
Profiles with one or more operations	<p>Multiple policies are created based on a unique combination of a profile and operations performed using that profile.</p> <p>The policies created in SnapCenter contain the archive log pruning and retention details retrieved from the profile and corresponding operations.</p>
Profiles with Oracle Recovery Manager (RMAN) configuration	<p>Policies are created with the Catalog backup with Oracle Recovery Manager option enabled.</p> <p>If external RMAN cataloging was used in SnapManager, you must configure the RMAN catalog settings in SnapCenter. You can either select the existing credential or create a new credential.</p> <p>If RMAN was configured through control file in SnapManager, then you do not have to configure RMAN in SnapCenter.</p>
Schedule attached to a profile	A policy is created just for the schedule.
Database	<p>A resource group is created for each database that is imported.</p> <p>In a Real Application Clusters (RAC) setup, the node on which you run the import tool becomes the preferred node after importing and the resource group is created for that node.</p>



When a profile is imported, a verification policy is created along with the backup policy.

When SnapManager for Oracle and SnapManager for SAP profiles, schedules, and any operations performed using the profiles are imported to SnapCenter, the different parameters values are also imported.

SnapManager for Oracle and SnapManager for SAP parameter and values	SnapCenter parameter and values	Notes
Backup Scope <ul style="list-style-type: none"> • Full • Data • Log 	Backup Scope <ul style="list-style-type: none"> • Full • Data • Log 	

SnapManager for Oracle and SnapManager for SAP parameter and values	SnapCenter parameter and values	Notes
Backup Mode <ul style="list-style-type: none"> • Auto • Online • Offline 	Backup Type <ul style="list-style-type: none"> • Online • Offline Shutdown 	If the backup mode is Auto, then the import tool checks the database state when the operation was performed, and appropriately sets the backup type as either Online or Offline Shutdown.
Retention <ul style="list-style-type: none"> • Days • Counts 	Retention <ul style="list-style-type: none"> • Days • Counts 	SnapManager for Oracle and SnapManager for SAP uses both Days and Counts to set the retention. In SnapCenter, there is either Days <i>OR</i> Counts. So, the retention is set with respect to days as the days get preference over counts in SnapManager for Oracle and SnapManager for SAP.
Pruning for Schedules <ul style="list-style-type: none"> • All • system change number (SCN) • Date • Logs created before specified hours, days, weeks, and months 	Pruning for Schedules <ul style="list-style-type: none"> • All • Logs created before specified hours and days 	SnapCenter does not support pruning based on SCN, Date, weeks, and months.
Notification <ul style="list-style-type: none"> • Emails sent only for successful operations • Emails sent only for failed operations • Emails sent for both success and failed operations 	Notification <ul style="list-style-type: none"> • Always • On failure • Warning • Error 	The email notifications are imported. However, you must manually update the SMTP server using the SnapCenter GUI. The subject of the email is left blank for you to configure.

What does not get imported to SnapCenter

The import tool does not import everything to SnapCenter.

You cannot import the following to SnapCenter:

- Backup metadata
- Partial backups

- Raw device mapping (RDM) and Virtual Storage Console (VSC) related backups
- Roles or any credentials available in the SnapManager for Oracle and SnapManager for SAP repository
- Data related to verification, restore, and clone operations
- Pruning for operations
- Replication details specified in the SnapManager for Oracle and SnapManager for SAP profile

After importing, you must manually edit the corresponding policy created in SnapCenter to include the replication details.

- Cataloged backup information

Prepare to import data

Before you import data to SnapCenter, you must perform certain tasks to run the import operation successfully.

Steps

1. Identify the database that you want to import.
2. Using SnapCenter, add the database host and install SnapCenter Plug-ins Package for Linux.
3. Using SnapCenter, set up the connections for the storage virtual machines (SVMs) used by the databases on the host.
4. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
5. In the Resources page, ensure that the database to be imported is discovered and displayed.

When you want to run the import tool, the database must be accessible or else the resource group creation fails.

If the database has credentials configured, you must create a corresponding credential in SnapCenter, assign the credential to the database, and then re-run discovery of the database. If the database is residing on Automatic Storage Management (ASM), you must create credentials for the ASM instance, and assign the credential to the database.

6. Ensure that the user running the import tool has sufficient privileges to run SnapManager for Oracle or SnapManager for SAP CLI commands (such as the command to suspend schedules) from SnapManager for Oracle or SnapManager for SAP host.
7. Run the following commands on the SnapManager for Oracle or SnapManager for SAP host to suspend the schedules:
 - a. If you want to suspend the schedules on the SnapManager for Oracle host, run:

- `smo credential set -repository -dbname repository_database_name -host host_name -port port_number -login -username user_name_for_repository_database`
- `smo profile sync -repository -dbname repository_database_name -host host_name -port port_number -login -username host_user_name_for_repository_database`
- `smo credential set -profile -name profile_name`



You must run the `smo credential set` command for each profile on the host.

b. If you want to suspend the schedules on the SnapManager for SAP host, run:

- `smsap credential set -repository -dbname repository_database_name -host host_name -port port_number -login -username user_name_for_repository_database`
- `smsap profile sync -repository -dbname repository_database_name -host host_name -port port_number -login -username host_user_name_for_repository_database`
- `smsap credential set -profile -name profile_name`



You must run the `smsap credential set` command for each profile on the host.

8. Ensure that fully qualified domain name (FQDN) of the database host is displayed when you run `hostname -f`.

If FQDN is not displayed, you must modify `/etc/hosts` to specify the FQDN of the host.

Import data

You can import data by running the import tool from the database host.

About this task

The SnapCenter backup policies that are created after importing have different naming formats:

- Policies created for the profiles without any operations and schedules have the `SM_PROFILENAME_ONLINE_FULL_DEFAULT_MIGRATED` format.

When no operation is performed using a profile, the corresponding policy is created with default backup type as online and backup scope as full.

- Policies created for the profiles with one or more operations have the `SM_PROFILENAME_BACKUPMODE_BACKUPSCOPE_MIGRATED` format.
- Policies created for the schedules attached to the profiles have the `SM_PROFILENAME_SMOSCHEDULENAME_BACKUPMODE_BACKUPSCOPE_MIGRATED` format.

Steps

1. Log in to the database host that you want to import.
2. Run the import tool by executing the `sc-migrate` script located at `/opt/NetApp/snapcenter/spl/bin`.
3. Enter the SnapCenter Server user name and password.

After validating the credentials, a connection is established with SnapCenter.

4. Enter the SnapManager for Oracle or SnapManager for SAP repository database details.

The repository database lists the databases that are available on the host.

5. Enter the target database details.

If you want to import all the databases on the host, enter all.

6. If you want to generate a system log or send ASUP messages for failed operations, you must enable them either by running the *Add-SmStorageConnection* or *Set-SmStorageConnection* command.



If you want to cancel an import operation, either while running the import tool or after importing, you must manually delete the SnapCenter policies, credentials, and resource groups that were created as part of import operation.

Results

The SnapCenter backup policies are created for profiles, schedules, and operations performed using the profiles. Resource groups are also created for each target database.

After importing the data successfully, the schedules associated with the imported database are suspended in SnapManager for Oracle and SnapManager for SAP.



After importing, you must manage the imported database or file system using SnapCenter.

The logs for every execution of the import tool are stored in the */var/opt/snapcenter/spl/logs* directory with the name *spl_migration_timestamp.log*. You can refer to this log to review import errors and troubleshoot them.

Install SnapCenter Plug-in for VMware vSphere

If your database is stored on virtual machines (VMs), or if you want to protect VMs and datastores, you must deploy the SnapCenter Plug-in for VMware vSphere virtual appliance.

For information to deploy, see [Deployment Overview](#).

Deploy CA certificate

To configure the CA Certificate with SnapCenter Plug-in for VMware vSphere, see [Create or import SSL certificate](#).

Configure the CRL file

SnapCenter Plug-in for VMware vSphere looks for the CRL files in a pre-configured directory. Default directory of the CRL files for SnapCenter Plug-in for VMware vSphere is */opt/netapp/config/crl*.

You can place more than one CRL file in this directory. The incoming certificates will be verified against each CRL.

Prepare for protecting Oracle databases

Before performing any data protection operation such as backup, clone, or restore operations, you must define your strategy and set up the environment. You can also set up the SnapCenter Server to use SnapMirror and SnapVault technology.

To take advantage of SnapVault and SnapMirror technology, you must configure and initialize a data protection relationship between the source and destination volumes on the storage device. You can use NetAppSystem Manager or you can use the storage console command line to perform these tasks.

Before you use the Plug-in for Oracle Database, the SnapCenter administrator should install and configure the SnapCenter Server and perform the prerequisite tasks.

- Install and configure SnapCenter Server. [Learn more](#)
- Configure the SnapCenter environment by adding storage system connections. [Learn more](#)



SnapCenter does not support multiple SVMs with the same name on different clusters. Each SVM registered with SnapCenter using either SVM registration or cluster registration must be unique.

- Create credentials with authentication mode as Linux or AIX for the install user. [Learn more](#)
- Add hosts, install the plug-ins, and discover the resources.
- If you are using SnapCenter Server to protect Oracle databases that reside on VMware RDM LUNs or VMDKs, you must deploy the SnapCenter Plug-in for VMware vSphere and register the plug-in with SnapCenter.
- Install Java on your Linux or AIX host.

See [Linux host requirements](#) or [AIX host requirements](#) for more information.

- You should set the time out value of the application firewall to 3 hours or more.
- If you have Oracle databases on NFS environments, you must have configured at least one NFS data LIF for primary or secondary storage to perform mount, clone, verification, and restore operations.
- If you have multiple data paths (LIFs) or a dNFS configuration, you can perform the following using the SnapCenter CLI on the database host:
 - By default, all the IP addresses of the database host are added to the NFS storage export policy in storage virtual machine (SVM) for the cloned volumes. If you want to have a specific IP address or restrict to a subset of the IP addresses, run the `Set-PreferredHostIPsInStorageExportPolicy` CLI.
 - If you have multiple data paths (LIFs) in SVM, SnapCenter chooses the appropriate data path (LIF) for mounting the NFS cloned volume. However, if you want to specify a specific data path (LIF), you must run the `Set-SvmPreferredDataPath` CLI. The command reference guide has more information.
- If you have Oracle databases on SAN environments, ensure that the SAN environment is configured as per the recommendation mentioned in the following guides:
 - [Recommended Host Settings for Linux Unified Host Utilities](#)
 - [Using Linux Hosts with ONTAP storage](#)
 - [Host Settings Affected by AIX Host Utilities](#)
- If you have Oracle databases on LVM in Oracle Linux or RHEL operating systems, install the latest version of Logical Volume Management (LVM).
- If you are using SnapManager for Oracle and want to migrate to SnapCenter Plug-in for Oracle Database, you can migrate the profiles to policies and resource groups of SnapCenter by using the `sccli` command `sc-migrate`.
- Configure SnapMirror and SnapVault on ONTAP, if you want backup replication

For SnapCenter 4.1.1 users, the SnapCenter Plug-in for VMware vSphere 4.1.1 documentation has information on protecting virtualized databases and file systems. For SnapCenter 4.2.x users, the NetApp Data Broker 1.0 and 1.0.1, documentation has information on protecting virtualized databases and file systems using the SnapCenter Plug-in for VMware vSphere that is provided by the Linux-based NetApp Data Broker virtual appliance (Open Virtual Appliance format). For SnapCenter 4.3.x users, the SnapCenter Plug-in for

VMware vSphere 4.3 documentation has information on protecting virtualized databases and file systems using the Linux-based SnapCenter Plug-in for VMware vSphere virtual appliance (Open Virtual Appliance format).

Find more information

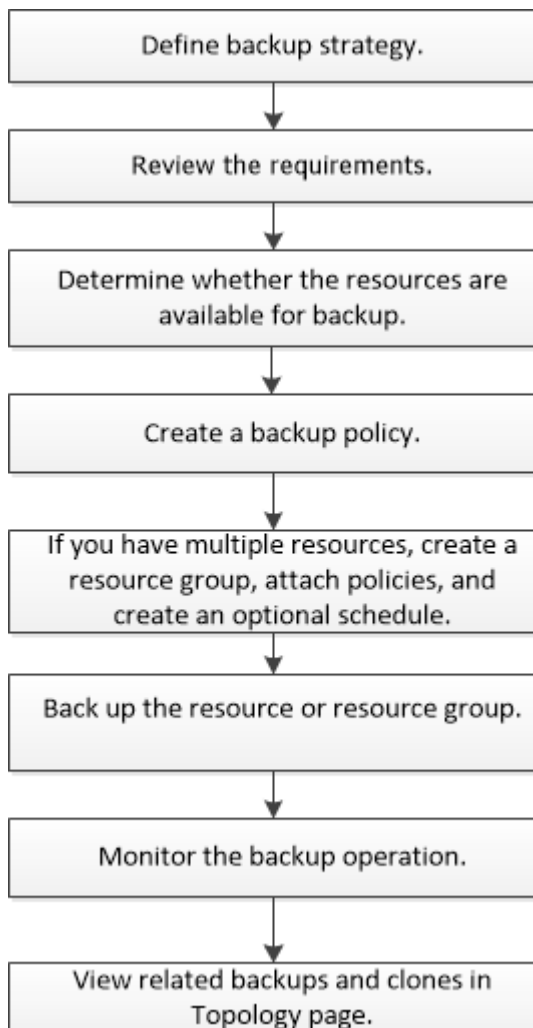
- [Interoperability Matrix Tool](#)
- [SnapCenter Plug-in for VMware vSphere documentation](#)
- [Data protection operation fails in a non-multipath environment in RHEL 7 and later](#)

Back up Oracle databases

Backup workflow

You can either create a backup of a resource (database) or resource group. The backup workflow includes planning, identifying the resources for backup, creating backup policies, creating resource groups and attaching policies, creating backups, and monitoring the operations.

The following workflow shows the sequence in which you must perform the backup operation:



While creating a backup for Oracle databases, an operational lock file (`.sm_lock_dbsid`) is created on the Oracle database host in the `$ORACLE_HOME/dbs` directory to avoid multiple operations being executed on the database. After the database has been backed up, the operational lock file is automatically removed.

However, if the previous backup was completed with a warning, the operational lock file might not get deleted, and the next backup operation gets into the wait queue. It might eventually get canceled if the `.sm_lock_dbsid` file is not deleted. In such scenario, you must manually delete the operational lock file by performing the following steps:

1. From the command prompt, navigate to `$ORACLE_HOME/dbs`.
2. Delete the operational lock: `rm -rf .sm_lock_dbsid`.

Define a backup strategy for Oracle databases

Defining a backup strategy before you create your backup jobs ensures that you have the backups that you require to successfully restore or clone your databases. Your service-level agreement (SLA), recovery time objective (RTO), and recovery point objective (RPO) largely determine your backup strategy.

An SLA defines the level of service that is expected and addresses many service-related issues, including the availability and performance of the service. RTO is the time by which a business process must be restored after a disruption in service. RPO defines the strategy for the age of the files that must be recovered from backup storage for regular operations to resume after a failure. SLA, RTO, and RPO contribute to the data protection strategy.

Supported Oracle database configurations for backups

SnapCenter supports backup of different Oracle database configurations.

- Oracle Standalone
- Oracle Real Application Clusters (RAC)
- Oracle Standalone Legacy
- Oracle Standalone Container Database (CDB)
- Oracle Data Guard standby

You can only create offline-mount backups of Data Guard standby databases. Offline-shutdown backup, archive log only backup, and full backup are not supported.

- Oracle Active Data Guard standby

You can only create online backups of Active Data Guard standby databases. Archive log only backup and full backup are not supported.



Before creating a backup of Data Guard standby or Active Data Guard standby database, the managed recovery process (MRP) is stopped and once the backup is created, MRP is started.

- Automatic Storage Management (ASM)
 - ASM standalone and ASM RAC on Virtual Machine Disk (VMDK)



Among all the restore methods supported for Oracle databases, you can perform only connect-and-copy restore of ASM RAC databases on VMDK.

- ASM standalone and ASM RAC on Raw device mapping (RDM) You can perform backup, restore, and clone operations on Oracle databases on ASM, with or without ASMLib.
- Oracle ASM Filter Driver (ASMFD)



PDB migration and PDB cloning operations are not supported.

- Oracle Flex ASM

For the latest information about supported Oracle versions, see the [NetApp Interoperability Matrix Tool](#).

Types of backup supported for Oracle databases

Backup type specifies the type of backup that you want to create. SnapCenter supports online and offline backup types for Oracle databases.

Online backup

A backup that is created when the database is in the online state is called an online backup. Also called a hot backup, an online backup enables you to create a backup of the database without shutting it down.

As part of online backup, you can create a backup of the following files:

- Datafiles and control files only
- Archive log files only (the database is not brought to backup mode in this scenario)
- Full database that includes datafiles, control files, and archive log files

Offline backup

A backup created when the database is either in a mounted or shutdown state is called an offline backup. An offline backup is also called a cold backup. You can include only datafiles and control files in offline backups. You can create either an offline mount or offline shutdown backup.

- When creating an offline mount backup, you must ensure that the database is in a mounted state.


If the database is in any other state, the backup operation fails.

- When creating an offline shutdown backup, the database can be in any state.

The database state is changed to the required state to create a backup. After creating the backup, the database state is reverted to the original state.

How SnapCenter discovers Oracle databases

"Resources" are Oracle databases on the host that are maintained by SnapCenter. You can add these databases to resource groups to perform data protection operations after you discover the databases that are available. You should be aware of the process that SnapCenter follows to discover different types and versions of Oracle databases.

For Oracle versions 11g to 12cR1	For Oracle versions 12cR2 to 18c
<p>RAC database: The RAC databases are discovered only on the basis of /etc/oratab entries.</p> <p>You should have the database entries in the /etc/oratab file.</p>	<p>RAC database: The RAC databases are discovered using the srvctl config command.</p>
<p>Standalone: The standalone databases are discovered only on the basis of /etc/oratab entries.</p> <p>You should have the database entries in the /etc/oratab file.</p>	<p>Standalone: The standalone databases are discovered based on the entries in the /etc/oratab file and the output of the srvctl config command.</p>
<p>ASM: The ASM instance entry should be available in the /etc/oratab file.</p>	<p>ASM: The ASM instance entry need not be in the /etc/oratab file.</p>
<p>RAC One Node: The RAC One Node databases are discovered only on the basis of /etc/oratab entries.</p> <p>The databases should be either in <i>nomount</i>, <i>mount</i>, or <i>open</i> state. You should have the database entries in the /etc/oratab file.</p> <p>The RAC One Node database status will be marked as renamed or deleted if the database is already discovered and backups are associated with the database.</p> <p>You should perform the following steps if the database is relocated:</p> <ol style="list-style-type: none"> 1. Manually add the relocated database entry in the /etc/oratab file on the failed-over RAC node. 2. Manually refresh the resources. 3. Select the RAC One Node database from the resource page, and then click Database Settings. 4. Configure the database to set the preferred cluster nodes to the RAC node currently hosting the database. 5. Perform the SnapCenter operations. <div>  <p>If you have relocated a database from one node to another node and if the oratab entry in the earlier node is not deleted, you should manually delete the oratab entry to avoid the same database being displayed twice.</p> </div>	<p>RAC One Node: The RAC One Node databases are discovered using the srvctl config command only.</p> <p>The databases should be either in <i>nomount</i>, <i>mount</i>, or <i>open</i> state. The RAC One Node database status will be marked as renamed or deleted if the database is already discovered and backups are associated with the database.</p> <p>You should perform the following steps if the database is relocated:</p> <ol style="list-style-type: none"> 1. Manually refresh the resources. 2. Select the RAC One Node database from the resource page, and then click Database Settings. 3. Configure the database to set the preferred cluster nodes to the RAC node currently hosting the database. 4. Perform the SnapCenter operations.



If there are any Oracle 12cR2 and 18c database entries in the /etc/oratab file and the same database is registered with the srvctl config command, SnapCenter will eliminate the duplicate database entries. If there are stale database entries, the database will be discovered but the database will be unreachable and the status will be offline.

Preferred nodes in RAC setup

In Oracle Real Application Clusters (RAC) setup, you can specify the preferred nodes on which the backup operation will be performed. If you do not specify the preferred node, SnapCenter automatically assigns a node as the preferred node and backup is created on that node.

The preferred nodes might be one or all of the cluster nodes where the RAC database instances are present. The backup operation will be triggered only on these preferred nodes in the order of the preference.

Example: The RAC database cdbbrac has three instances: cdbbrac1 on node1, cdbbrac2 on node2, and cdbbrac3 on node3. The node1 and node2 instances are configured to be the preferred nodes, with node2 as the first preference and node1 as the second preference. When you perform a backup operation, the operation is first attempted on node2 because it is the first preferred node. If node2 is not in the state to back up, which could be due to multiple reasons such as the plug-in agent is not running on the host, the database instance on the host is not in the required state for the specified backup type, or the database instance on node2 in a FlexASM configuration is not being served by the local ASM instance; then the operation will be attempted on node1. The node3 will not be used for backup because it is not on the list of preferred nodes.

In a Flex ASM setup, Leaf nodes will not be listed as preferred nodes if the cardinality is less than the number nodes in the RAC cluster. If there is any change in the Flex ASM cluster node roles, you should manually discover so that the preferred nodes are refreshed.

Required database state

The RAC database instances on the preferred nodes must be in the required state for the backup to finish successfully:

- One of the RAC database instances in the configured preferred nodes must be in the open state to create an online backup.
- One of the RAC database instances in the configured preferred nodes must be in the mount state, and all other instances, including other preferred nodes, must be in the mount state or lower to create an offline mount backup.
- RAC database instances can be in any state, but you must specify the preferred nodes to create an offline shutdown backup.

How to catalog backups with Oracle Recovery Manager

The backups of Oracle databases can be cataloged with Oracle Recovery Manager (RMAN) to store the backup information in the Oracle RMAN repository.

The cataloged backups can be used later for block-level restore or tablespace point-in-time recovery operations. When you do not need these cataloged backups, you can remove the catalog information.

The database must be in mounted or higher state for cataloging. You can perform cataloging on data backups, archive log backups, and full backups. If cataloging is enabled for a backup of a resource group that has multiple databases, cataloging is performed for each database. For Oracle RAC databases, cataloging will be performed on the preferred node where the database is at least in mounted state.



If you want to catalog backups of a RAC database, ensure that no other job is running for that database. If another job is running, the cataloging operation fails instead of getting queued.

By default, the target database control file is used for cataloging. If you want to add external catalog database, you can configure it by specifying the credential and Transparent Network Substrate (TNS) name of the external catalog using the Database Settings wizard from the SnapCenter graphical user interface (GUI). You can also configure the external catalog database from the CLI by running the `Configure-SmOracleDatabase` command with the `-OracleRmanCatalogCredentialName` and `-OracleRmanCatalogTnsName` options.

If you enabled the cataloging option while creating an Oracle backup policy from the SnapCenter GUI, the backups are cataloged using Oracle RMAN as a part of the backup operation. You can also perform deferred cataloging of backups by running the `Catalog-SmBackupWithOracleRMAN` command. After cataloging the backups, you can run the `Get-SmBackupDetails` command to obtain the cataloged backup information such as the tag for cataloged datafiles, the control file catalog path, and the cataloged archive log locations.

If the ASM disk group name is greater than or equal to 16 characters, from SnapCenter 3.0, the naming format used for the backup is `SC_HASHCODEofDISKGROUP_DBSID_BACKUPID`. However, If the disk group name is less than 16 characters, the naming format used for the backup is `DISKGROUPNAME_DBSID_BACKUPID`, which is the same format used in SnapCenter 2.0.



The `HASHCODEofDISKGROUP` is an automatically generated number (2 to 10 digit) unique for each ASM disk group.

You can perform crosschecks to update outdated RMAN repository information about backups whose repository records do not match their physical status. For example, if a user removes archived logs from disk with an operating system command, the control file still indicates that the logs are on disk, when in fact they are not. The crosscheck operation enables you to update the control file with the information. You can enable crosscheck by running the `Set-SmConfigSettings` command and assigning the value `TRUE` to the `ENABLE_CROSSCHECK` parameter. The default value is set to `FALSE`.

```
sccli Set-SmConfigSettings-ConfigSettingsTypePlugin-PluginCodeSCO-ConfigSettings
"KEY=ENABLE_CROSSCHECK, VALUE=TRUE"
```

You can remove the catalog information by running the `Uncatalog-SmBackupWithOracleRMAN` command. You cannot remove the catalog information using the SnapCenter GUI. However, information of a cataloged backup is removed while deleting the backup or while deleting the retention and resource group associated with that cataloged backup.



When you force a deletion of the SnapCenter host, the information of the cataloged backups associated with that host are not removed. You must remove information of all the cataloged backups for that host before forcing the deletion of the host.

If the cataloging and uncataloging fails because the operation time exceeded the time out value specified for the `ORACLE_PLUGIN_RMAN_CATALOG_TIMEOUT` parameter, you should modify the value of the parameter by running the following command:

```
/opt/Netapp/snapcenter/spl/bin/sccli Set-SmConfigSettings-ConfigSettingsType
Plugin -PluginCode SCO-ConfigSettings
"KEY=ORACLE_PLUGIN_RMAN_CATALOG_TIMEOUT,VALUE=user_defined_value"
```

After modifying the value of the parameter, restart the SnapCenter Plug-in Loader (SPL) service by running the following command:

```
/opt/NetApp/snapcenter/spl/bin/spl restart
```

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

Backup schedules

Backup frequency (schedule type) is specified in policies; a backup schedule is specified in the resource group configuration. The most critical factor in determining a backup frequency or schedule is the rate of change for the resource and the importance of the data. You might back up a heavily used resource every hour, while you might back up a rarely used resource once a day. Other factors include the importance of the resource to your organization, your Service Level Agreement (SLA), and your Recover Point Objective (RPO).

An SLA defines the level of service expected and addresses many service-related issues, including the availability and performance of service. An RPO defines the strategy for the age of the files that must be recovered from backup storage for regular operations to resume after a failure. The SLA and RPO contribute to the data protection strategy.

Even for a heavily used resource, there is no requirement to run a full backup more than once or twice a day. For example, regular transaction log backups might be sufficient to ensure that you have the backups you need. The more often you back up your databases, the fewer transaction logs SnapCenter has to use at restore time, which can result in faster restore operations.

Backup schedules have two parts, as follows:

- Backup frequency

Backup frequency (how often backups are to be performed), called *schedule type* for some plug-ins, is part of a policy configuration. You can select hourly, daily, weekly, or monthly as the backup frequency for the policy. If you do not select any of these frequencies, then the policy created is an on-demand-only policy. You can access policies by clicking **Settings > Policies**.

- Backup schedules

Backup schedules (exactly when backups are to be performed) are part of a resource group configuration. For example, if you have a resource group that has a policy configured for weekly backups, you might configure the schedule to back up every Thursday at 10:00 PM. You can access resource group schedules by clicking **Resources > Resource Groups**.

Backup naming conventions

You can either use the default Snapshot copy naming convention or use a customized naming convention. The default backup naming convention adds a timestamp to Snapshot copy names that helps you identify when the copies were created.

The Snapshot copy uses the following default naming convention:

```
resourcegroupname_hostname_timestamp
```

You should name your backup resource groups logically, as in the following example:

```
dtst1_mach1x88_03-12-2015_23.17.26
```

In this example, the syntax elements have the following meanings:

- *dts1* is the resource group name.
- *mach1x88* is the host name.
- *03-12-2015_23.17.26* is the date and timestamp.

Alternatively, you can specify the Snapshot copy name format while protecting resources or resource groups by selecting **Use custom name format for Snapshot copy**. For example, `customtext_resourcegroup_policy_hostname` or `resourcegroup_hostname`. By default, the time stamp suffix is added to the Snapshot copy name.

Backup retention options

You can choose either the number of days for which to retain backup copies or specify the number of backup copies you want to retain, up to a ONTAP maximum of 255 copies. For example, your organization might require that you retain 10 days of backup copies or 130 backup copies.

While creating a policy, you can specify the retention options for the backup type and the schedule type.

If you set up SnapMirror replication, the retention policy is mirrored on the destination volume.

SnapCenter deletes the retained backups that have retention labels that match the schedule type. If the schedule type was changed for the resource or resource group, backups with the old schedule type label might still remain on the system.



For long-term retention of backup copies, you should use SnapVault backup.

Verify backup copy using the primary or secondary storage volume

You can verify backup copies on the primary storage volume or on either the SnapMirror or SnapVault secondary storage volume. Verification using a secondary storage volume reduces load on the primary storage volume.

When you verify a backup that is either on the primary or secondary storage volume, all the primary and the secondary Snapshot copies are marked as verified.

SnapRestore license is required to verify backup copies on SnapMirror and SnapVault secondary storage volume.

Predefined environment variables for backup specific prescript and postscript

SnapCenter allows you to use the predefined environment variables when you execute the prescript and postscript while creating backup policies. This functionality is supported for all Oracle configurations except for VMDK.

SnapCenter predefines the values of the parameters that will be directly accessible in the environment where the shell scripts are executed. You do not have to manually specify the values of these parameters when executing the scripts.

Supported predefined environment variables for creating backup policy

- **SC_JOB_ID** specifies the job ID of the operation.

Example: 256

- **SC_ORACLE_SID** specifies the system identifier of the database.

If the operation involves multiple databases, the parameter will contain database names separated by pipe.

This parameter will be populated for application volumes.

Example: NFSB32|NFSB31

- **SC_HOST** specifies the host name of the database.

For RAC, host name will be the name of the host on which backup is performed.

This parameter will be populated for application volumes.

Example: scsmohost2.gdl.englabe.netapp.com

- **SC_OS_USER** specifies the operating system owner of the database.

The data will be formatted as <db1>@<osuser1>|<db2>@<osuser2>.

Example: NFSB31@oracle|NFSB32@oracle

- **SC_OS_GROUP** specifies the operating system group of the database.

The data will be formatted as <db1>@<osgroup1>|<db2>@<osgroup2>.

Example: NFSB31@install|NFSB32@oinstall

- **SC_BACKUP_TYPE** specifies the backup type (online full, online data, online log, offline shutdown, offline mount)

Examples:

- For full backup: ONLINEFULL
- data only backup: ONLINEDATA
- For log only backup: ONLINELOG

- **SC_BACKUP_NAME** specifies the name of the backup.

This parameter will be populated for application volumes.

Example: DATA@RG2_scspr2417819002_07-20-2021_12.16.48.9267_0|LOG@RG2_scspr2417819002_07-20-2021_12.16.48.9267_1|AV@RG2_scspr2417819002_07-20-2021_12.16.48.9267

- **SC_BACKUP_ID** specifies the backup ID.

This parameter will be populated for application volumes.

Example: DATA@203|LOG@205|AV@207

- **SC_ORACLE_HOME** specifies the path of the Oracle home directory.

Example:

NFSB32@/ora01/app/oracle/product/18.1.0/db_1|NFSB31@/ora01/app/oracle/product/18.1.0/db_1

- **SC_BACKUP_RETENTION** specifies the retention period defined in the policy.

Examples:

- For full backup: Hourly|DATA@DAYS:3|LOG@COUNT:4
- For on-demand data only backup: Ondemand|DATA@COUNT:2
- For on-demand log only backup: Ondemand|LOG@COUNT:2

- **SC_RESOURCE_GROUP_NAME** specifies the name of the resource group.

Example: RG1

- **SC_BACKUP_POLICY_NAME** specifies the name of the backup policy.

Example: backup_policy

- **SC_AV_NAME** specifies the names of the application volumes.

Example: AV1|AV2

- **SC_PRIMARY_DATA_VOLUME_FULL_PATH** specifies the storage mapping of SVM to volume for data files directory. It will be the name of the parent volume for luns and qtrees.

The data will be formatted as <db1>@<SVM1:volume1>|<db2>@<SVM2:volume2>.

Examples:

- For 2 databases in the same resource group:
NFSB32@buck:/vol/scspr2417819002_NFS_CDB_NFSB32_DATA|NFSB31@buck:/vol/scspr2417819002_NFS_CDB_NFSB31_DATA
- For single database with data files spread across multiple volumes:
buck:/vol/scspr2417819002_NFS_CDB_NFSB31_DATA,herculus:/vol/scspr2417819002_NFS

- **SC_PRIMARY_ARCHIVELOGS_VOLUME_FULL_PATH** specifies the storage mapping of SVM to volume for logs file directory. It will be the name of the parent volume for luns and qtrees.

Examples:

- For single database instance: buck:/vol/scspr2417819002_NFS_CDB_NFSB31_REDO
- For multiple database instances:
NFSB31@buck:/vol/scspr2417819002_NFS_CDB_NFSB31_REDO|NFSB32@buck:/vol/scspr2417819002_NFS_CDB_NFSB32_REDO

- **SC_PRIMARY_FULL_SNAPSHOT_NAME_FOR_TAG** specifies the list of Snapshots containing storage system name and volume name.

Examples:

- For single database instance:
buck:/vol/scspr2417819002_NFS_CDB_NFSB32_DATA/RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,buck:/vol/scspr2417819002_NFS_CDB_NFSB32_REDO/RG2_scspr2417819002_07-21-2021_02.28.26.3973_1
- For multiple database instances:

NFSB32@buck:/vol/scspr2417819002_NFS_CDB_NFSB32_DATA/RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,buck:/vol/scspr2417819002_NFS_CDB_NFSB32_REDO/RG2_scspr2417819002_07-21-2021_02.28.26.3973_1|NFSB31@buck:/vol/scspr2417819002_NFS_CDB_NFSB31_DATA/RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,buck:/vol/scspr2417819002_NFS_CDB_NFSB31_REDO/RG2_scspr2417819002_07-21-2021_02.28.26.3973_1

- **SC_PRIMARY_SNAPSHOT_NAMES** specifies the names of the primary Snapshots created during the backup.

Examples:

- For single database instance: RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,RG2_scspr2417819002_07-21-2021_02.28.26.3973_1
- For multiple database instances: NFSB32@RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,RG2_scspr2417819002_07-21-2021_02.28.26.3973_1|NFSB31@RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,RG2_scspr2417819002_07-21-2021_02.28.26.3973_1
- For consistency group Snapshots that involves 2 volumes: cg3_R80404CBEF5V1_04-05-2021_03.08.03.4945_0_bfc279cc-28ad-465c-9d60-5487ac17b25d_2021_4_5_3_8_58_350

- **SC_PRIMARY_MOUNT_POINTS** specifies the mount point details which are part of the backup.

The details include the directory on which volumes are mounted and not the immediate parent of the file under backup. For an ASM configuration, it is the name of the disk group.

The data will be formatted as <db1>@<mountpoint1,mountpoint2>|<db2>@<mountpoint1,mountpoint2>.

Examples:

- For single database instance: /mnt/nfsdb3_data,/mnt/nfsdb3_log,/mnt/nfsdb3_data1
- For multiple database instances: NFSB31@/mnt/nfsdb31_data,/mnt/nfsdb31_log,/mnt/nfsdb31_data1|NFSB32@/mnt/nfsdb32_data,/mnt/nfsdb32_log,/mnt/nfsdb32_data1
- For ASM: +DATA2DG,+LOG2DG

- **SC_PRIMARY_SNAPSHOTS_AND_MOUNT_POINTS** specifies the names of the snapshots created during the backup of each of the mount points.

Examples:

- For single database instance: RG2_scspr2417819002_07-21-2021_02.28.26.3973_0:/mnt/nfsb32_data,RG2_scspr2417819002_07-21-2021_02.28.26.3973_1:/mnt/nfsb31_log
- For multiple database instances: NFSB32@RG2_scspr2417819002_07-21-2021_02.28.26.3973_0:/mnt/nfsb32_data,RG2_scspr2417819002_07-21-2021_02.28.26.3973_1:/mnt/nfsb31_log|NFSB31@RG2_scspr2417819002_07-21-2021_02.28.26.3973_0:/mnt/nfsb31_data,RG2_scspr2417819002_07-21-2021_02.28.26.3973_1:/mnt/nfsb32_log

- **SC_ARCHIVELOGS_LOCATIONS** specifies the location of the archive logs directory.

The directory names will be the immediate parent of the archive log files. If the archive logs are placed in more than one location then all the locations will be captured. This also includes the FRA scenarios. If

softlinks are used for directory then the same will be populated.

Examples:

- For single database on NFS: /mnt/nfsdb2_log
- For multiple databases on NFS and for the NFSB31 database archive logs that are placed in two different locations: NFSB31@/mnt/nfsdb31_log1,/mnt/nfsdb31_log2|NFSB32@/mnt/nfsdb32_log
- For ASM: +LOG2DG/ASMDB2/ARCHIVELOG/2021_07_15

- **SC_REDO_LOGS_LOCATIONS** specifies the location of the redo logs directory.

The directory names will be the immediate parent of the redo log files. If softlinks are used for directory then the same will be populated.

Examples:

- For single database on NFS: /mnt/nfsdb2_data/newdb1
- For multiple databases on NFS:
NFSB31@/mnt/nfsdb31_data/newdb31|NFSB32@/mnt/nfsdb32_data/newdb32
- For ASM: +LOG2DG/ASMDB2/ONLINELOG

- **SC_CONTROL_FILES_LOCATIONS** specifies the location of the control files directory.

The directory names will be the immediate parent of the control files. If softlinks are used for directory then the same will be populated.

Examples:

- For single database on NFS: /mnt/nfsdb2_data/fra/newdb1,/mnt/nfsdb2_data/newdb1
- For multiple databases on NFS:
NFSB31@/mnt/nfsdb31_data/fra/newdb31,/mnt/nfsdb31_data/newdb31|NFSB32@/mnt/nfsdb32_data/fra/newdb32,/mnt/nfsdb32_data/newdb32
- For ASM: +LOG2DG/ASMDB2/CONTROLFILE

- **SC_DATA_FILES_LOCATIONS"** specifies the location of the data files directory.

The directory names will be the immediate parent of the data files. If softlinks are used for directory then the same will be populated.

Examples:

- For single database on NFS: /mnt/nfsdb3_data1,/mnt/nfsdb3_data/NEWDB3/datafile
- For multiple databases on NFS:
NFSB31@/mnt/nfsdb31_data1,/mnt/nfsdb31_data/NEWDB31/datafile|NFSB32@/mnt/nfsdb32_data1,/mnt/nfsdb32_data/NEWDB32/datafile
- For ASM: +DATA2DG/ASMDB2/DATAFILE,+DATA2DG/ASMDB2/TEMPFILE

- **SC_SNAPSHOT_LABEL** specifies the name of the secondary labels.

Examples: Hourly, Daily, Weekly, Monthly, or custom label.

Supported delimiters

- `:` is used to separate SVM name and volume name

Example: `buck:/vol/scspr2417819002_NFS_CDB_NFSB32_DATA/RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,buck:/vol/scspr2417819002_NFS_CDB_NFSB32_REDO/RG2_scspr2417819002_07-21-2021_02.28.26.3973_1`

- `@` is used to separate data from its database name and to separate the value from its key.

Examples:

- `NFSB32@buck:/vol/scspr2417819002_NFS_CDB_NFSB32_DATA/RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,buck:/vol/scspr2417819002_NFS_CDB_NFSB32_REDO/RG2_scspr2417819002_07-21-2021_02.28.26.3973_1|NFSB31@buck:/vol/scspr2417819002_NFS_CDB_NFSB31_DATA/RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,buck:/vol/scspr2417819002_NFS_CDB_NFSB31_REDO/RG2_scspr2417819002_07-21-2021_02.28.26.3973_1`
- `NFSB31@oracle|NFSB32@oracle`
- `|` is used to separate the data between two different databases and to separate the data between two different entities for `SC_BACKUP_ID`, `SC_BACKUP_RETENTION`, and `SC_BACKUP_NAME` parameters.

Examples:

- `DATA@203|LOG@205`
- `Hourly|DATA@DAYS:3|LOG@COUNT:4`
- `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 volume name from its Snapshot for `SC_PRIMARY_SNAPSHOT_NAMES` and `SC_PRIMARY_FULL_SNAPSHOT_NAME_FOR_TAG` parameters.

Example: `NFSB32@buck:/vol/scspr2417819002_NFS_CDB_NFSB32_DATA/RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,buck:/vol/scspr2417819002_NFS_CDB_NFSB32_REDO/RG2_scspr2417819002_07-21-2021_02.28.26.3973_1`

- `,` is used to separate set of variables for the same DB.

Example: `NFSB32@buck:/vol/scspr2417819002_NFS_CDB_NFSB32_DATA/RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,buck:/vol/scspr2417819002_NFS_CDB_NFSB32_REDO/RG2_scspr2417819002_07-21-2021_02.28.26.3973_1|NFSB31@buck:/vol/scspr2417819002_NFS_CDB_NFSB31_DATA/RG2_scspr2417819002_07-21-2021_02.28.26.3973_0,buck:/vol/scspr2417819002_NFS_CDB_NFSB31_REDO/RG2_scspr2417819002_07-21-2021_02.28.26.3973_1`

Determine whether Oracle databases are available for backup

Resources are Oracle databases on the host that are managed by SnapCenter. You can add these databases to resource groups to perform data protection operations after you

discover the databases that are available.

What you will need

- You must have completed tasks such as installing the SnapCenter Server, adding hosts, creating storage system connections, and adding credentials.
- If the databases reside on a Virtual Machine Disk (VMDK) or raw device mapping (RDM), you must deploy the SnapCenter Plug-in for VMware vSphere and register the plug-in with SnapCenter.

For more information, see [Deploy SnapCenter Plug-in for VMware vSphere](#).

- If databases reside on a VMDK file system, you must have logged in to vCenter and navigated to **VM options > Advanced > Edit configuration** to set the value of *disk.enableUUID* to true for the VM.
- You must have reviewed the process that SnapCenter follows to discover different types and versions of Oracle databases.

About this task



After installing the plug-in, all of the databases on that host are automatically discovered and displayed in the Resources page.

The databases should be at least in the mounted state or above for the discovery of the databases to be successful. In an Oracle Real Application Clusters (RAC) environment, the RAC database instance in the host where the discovery is performed, should be at least in the mounted state or above for the discovery of the database instance to be successful. Only the databases that are discovered successfully can be added to the resource groups.

If you have deleted an Oracle database on the host, SnapCenter Server will not be aware and will list the deleted database. You should manually refresh the resources to update the SnapCenter resources list.

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 **Database** from the **View** list.

Click , and then select the host name and the database type to filter the resources. You can then click the  icon to close the filter pane.

3. Click **Refresh Resources**.

In a RAC One Node scenario, the database is discovered as the RAC database on the node where it is currently hosted.

Results

The databases are displayed along with information such as database type, host or cluster name, associated resource groups and policies, and status.



You must refresh the resources if the databases are renamed outside of SnapCenter.

- If the database is on a non-NetApp storage system, the user interface displays a Not available for backup message in the Overall Status column.

You cannot perform data protection operations on the database that is on a non-NetApp storage system.

- If the database is on a NetApp storage system and not protected, the user interface displays a Not protected message in the Overall Status column.
- If the database is on a NetApp storage system and protected, the user interface displays an Available for backup message in the Overall Status column.



If you have enabled an Oracle database authentication, a red padlock icon is shown in the resources view. You must configure database credentials to be able to protect the database or add it to the resource group to perform data protection operations.

How to prevent SnapCenter from discovering non-database entries

You can prevent SnapCenter from discovering non-database entries added in the `oratab` file.

Steps

1. After installing the plug-in for Oracle, the root user should create the **`sc_oratab.config`** file under the directory `/var/opt/snapcenter/sco/etc/`.

Grant the write permission to the Oracle binary owner and group so that the file could be maintained in future.

2. Database administrator should add the non-database entries in the **`sc_oratab.config`** file.

It is recommended to maintain same format defined for the non-database entries in the `/etc/oratab` file or the user can just add the non-database entity string.



The string is case sensitive. Any text with `#` in the beginning is treated as a comment. The comment can be appended after the non-database name.

For example:

```
-----  
# Sample entries  
# Each line can have only one non-database name  
# These are non-database name  
oratar # Added by the admin group -1  
#Added by the script team  
NEWSPT  
DBAGNT:/ora01/app/oracle/product/agent:N  
-----
```

3. Discover the resources.

The non-database entries added in the **`sc_oratab.config`** will not be listed in the Resources page.



It is always recommended to take a backup of the `sc_oratab.config` file before upgrading the SnapCenter plug-in.

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. A backup policy is a set of rules that governs how you manage, schedule, and retain backups. You can also specify the replication, script, and backup type settings. Creating a policy saves time when you want to reuse the policy on another resource or resource group.

What you will need

- You must have defined your backup strategy.
- You must have prepared for data protection by completing tasks such as installing SnapCenter, adding hosts, discovering databases, and creating storage system connections.
- If you are replicating Snapshot copies to a mirror or vault secondary storage, the SnapCenter administrator must have assigned the SVMs to you for both the source and destination volumes.

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 Backup Type page, perform the following steps:
 - If you want to **create an online backup**, select **Online backup**.

You must specify whether you want to back up all the datafiles, control files, and archive log files, only datafiles and control files, or only archive log files.

- If you want to **create an offline backup**, select **Offline backup**, and then select one of the following options:
 - If you want to create an offline backup when the database is in mounted state, select **Mount**.
 - If you want to create an offline shutdown backup by changing the database to shutdown state, select **Shutdown**.

If you are having pluggable databases (PDBs), and want to save the state of the PDBs before creating the backup, you must select **Save state of PDBs**. This enables you to bring the PDBs to their original state after the backup is created.

- Specify the schedule frequency by selecting **On demand**, **Hourly**, **Daily**, **Weekly**, or **Monthly**.



You can specify the schedule (start date and end date) for the backup operation while creating a resource group. This enables you to create resource groups that share the same policy and backup frequency, but enables you to assign different backup schedules to each policy.



If you have scheduled for 2:00 a.m., the schedule will not be triggered during daylight saving time (DST).

- If you want to catalog backup using Oracle Recovery Manager (RMAN), select **Catalog backup with Oracle Recovery Manager (RMAN)**.

You can perform deferred cataloging for one backup at a time either using the GUI or using the SnapCenter CLI command `Catalog-SmBackupWithOracleRMAN`.



If you want to catalog backups of a RAC database, ensure that no other job is running for that database. If another job is running, the cataloging operation fails instead of getting queued.

- If you want to prune archive logs after backup, select **Prune archive logs after backup**.



Pruning of archive logs from the archive log destination that is unconfigured in the database, will be skipped.



If you are using Oracle Standard Edition, you can use `LOG_ARCHIVE_DEST` and `LOG_ARCHIVE_DUPLEX_DEST` parameters while performing archive log backup.

- You can delete archive logs only if you have selected the archive log files as part of your backup.



You must ensure that all the nodes in an RAC environment can access all the archive log locations for the delete operation to be successful.

If you want to...	Then...
Delete all archive logs	Select Delete all archive logs .
Delete archive logs that are older	Select Delete archive logs older than , and then specify the age of the archive logs that are to be deleted in days and hours.
Delete archive logs from all destinations	Select Delete archive logs from all the destinations .
Delete the archive logs from the log destinations that are part of the backup	Select Delete archive logs from the destinations which are part of backup .

☒ Prune archive logs after backup

Prune log retention setting

☐ Delete all archive logs

☒ Delete archive logs older than



Prune log destination setting

☐ Delete archive logs from all the destinations

☒ Delete archive logs from the destinations which are part of backup

7. In the Retention page, specify the retention settings for the backup type and the schedule type selected in

the Backup Type page:


If you want to...	Then...
Keep a certain number of Snapshot copies	<p>Select Total Snapshot copies to keep, and then specify the number of Snapshot copies that you want to keep.</p> <p>If the number of Snapshot copies exceeds the specified number, the Snapshot copies are deleted with the oldest copies deleted first.</p> <div><div></div><div>The maximum retention value is 1018 for resources on ONTAP 9.4 or later, and 254 for resources on ONTAP 9.3 or earlier. Backups will fail if retention is set to a value higher than what the underlying ONTAP version supports.</div></div> <div><div></div><div>You must set the retention count to 2 or higher if you plan to enable SnapVault replication. If you set the retention count to 1, the retention operation might fail because the first Snapshot copy is the reference Snapshot copy for the SnapVault relationship until a newer Snapshot copy is replicated to the target.</div></div>
Keep the Snapshot copies for a certain number of days	Select Keep Snapshot copies for , and then specify the number of days for which you want to keep the Snapshot copies before deleting them.



You can retain archive log backups only if you have selected the archive log files as part of your backup.

8. In the Replication page, specify the replication settings:

For this field...	Do this...
Update SnapMirror after creating a local Snapshot copy	Select this field to create mirror copies of the backup sets on another volume (SnapMirror replication).
Update SnapVault after creating a local Snapshot copy	Select this option to perform disk-to-disk backup replication (SnapVault backups).

For this field...	Do this...
Secondary policy label	<p>Select a Snapshot label.</p> <p>Depending on the Snapshot copy label that you select, ONTAP applies the secondary Snapshot copy retention policy that matches the label.</p> <div>  <p>If you have selected Update SnapMirror after creating a local Snapshot copy, you can optionally specify the secondary policy label. However, if you have selected Update SnapVault after creating a local Snapshot copy, you should specify the secondary policy label.</p> </div>
Error retry count	Enter the maximum number of replication attempts that can be allowed before the operation stops.



You should configure SnapMirror retention policy in ONTAP for the secondary storage to avoid reaching the maximum limit of Snapshot copies on the secondary storage.

9. 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.

You must store the prescripts and postscripts either in `/var/opt/snapcenter/spl/scripts` or in any folder inside this path. By default, the `/var/opt/snapcenter/spl/scripts` path is populated. If you have created any folders inside this path to store the scripts, you must specify those folders in the path.

You can also specify the script timeout value. The default value is 60 seconds.

SnapCenter allows you to use the predefined environment variables when you execute the prescript and postscript. [Learn more](#)

10. In the Verification page, perform the following steps:
 - a. Select the backup schedule for which you want to perform the verification operation.
 - b. In the Verification script commands section, enter the path and the arguments of the prescript or postscript that you want to run before or after the verification operation, respectively.

You must store the prescripts and postscripts either in `/var/opt/snapcenter/spl/scripts` or in any folder inside this path. By default, the `/var/opt/snapcenter/spl/scripts` path is populated. If you have created any folders inside this path to store the scripts, you must specify those folders in the path.

You can also specify the script timeout value. The default value is 60 seconds.

11. Review the summary, and then click **Finish**.

Create resource groups and attach policies for Oracle databases

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.

About this task

You should ensure that the database having files on the ASM disk groups should be either in "MOUNT" or "OPEN" state to verify its backups using the Oracle DBVERIFY utility.

You should attach one or more policies to the resource group to define the type of data protection job that you want to perform.

The following image illustrates the relationship between resources, resource groups, and policies for databases:



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	<div>Enter a name for the resource group.</div> <div><div><div>i</div></div><div>The resource group name should not exceed 250 characters.</div></div>
Tags	<div>Enter one or more labels that will help you later search for the resource group.</div> <div>For example, if you add HR as a tag to multiple resource groups, you can later find all resource groups associated with the HR tag.</div>

For this field...	Do this...
Use custom name format for Snapshot copy	<p>Select this check box, and enter a custom name format that you want to use for the Snapshot copy name.</p> <p>For example, customtext_resource group_policy_hostname or resource group_hostname. By default, a timestamp is appended to the Snapshot copy name.</p>
Exclude archive log destinations from backup	Specify the destinations of the archive log files that you do not want to back up.

4. In the Resources page, select an Oracle database host name from the **Host** drop-down list.



The resources are listed in the Available Resources section only if the resource is discovered successfully. If you have recently added resources, they will appear on the list of available resources only after you refresh your resource list.

5. Select the resources from the Available Resources section and move them to the Selected Resources section.



You can add databases from both Linux and AIX hosts in a single resource group.


6. In the Policies page, perform the following steps:

- a. Select one or more policies from the drop-down list.



You can also create a policy by clicking .

In the Configure schedules for selected policies section, the selected policies are listed.


- b. Click  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* window, configure the schedule, and then click **OK**.

Where, *policy_name* is the name of the policy that you have selected.

The configured schedules are listed in the Applied Schedules column.

Third party backup schedules are not supported when they overlap with SnapCenter backup schedules.

7. On the Verification page, perform the following steps:

- a. Click **Load locators** to load the SnapMirror or SnapVault volumes to perform verification on secondary storage.
- b. Click  in the Configure Schedules column to configure the verification schedule for all the schedule types of the policy.

c. In the Add Verification Schedules policy_name dialog box, perform the following actions:

If you want to...	Do this...
Run verification after backup	Select Run verification after backup .
Schedule a verification	Select Run scheduled verification and then select the schedule type from the drop-down list.

d. Select **Verify on secondary location** to verify your backups on secondary storage system.

e. Click **OK**.

The configured verification schedules are listed in the Applied Schedules column.

8. 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. If you want to attach the report of the operation performed on the resource group, select **Attach Job Report**.



For email notification, you must have specified the SMTP server details using either the GUI or the PowerShell command Set-SmSmtServer.

9. Review the summary, and then click **Finish**.

Requirements for backing up an Oracle database

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


- You must have created a resource group with a policy attached.
- If you want to back up a resource that has a SnapMirror relationship with a secondary storage, the ONTAP role assigned to the storage user should include the “snapmirror all” privilege. However, if you are using the “vsadmin” role, then the “snapmirror all” privilege is not required.
- You must have assigned the aggregate that is being used by the backup operation to the storage virtual machine (SVM) used by the database.
- You should have verified that all data volumes and archive log volumes belonging to the database are protected if secondary protection is enabled for that database.
- You should have verified that the database that has files on the ASM disk groups should be in either “MOUNT” or “OPEN” state to verify its backups using the Oracle DBVERIFY utility.
- You should have verified that the volume mount point length does not exceed 240 characters.
- You should increase value of RESTTimeout to 86400000 ms in *C:\Program Files\NetApp\SMCore\SMCoreServiceHost.exe.config* file in the SnapCenter Server host, if the database being backed up is large (size in TBs).

While modifying the values ensure that there are no running jobs and restart the SnapCenter SMCore service after increasing the value.

Back up Oracle resources

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

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 **Database** from the **View** list.
3. Click , and then select the host name and the database type to filter the resources.

You can then click  to close the filter pane.

4. Select the database that you want to back up.

The Database-Protect page is displayed.

5. In the Resource page, perform the following actions:


For this field...	Do this...
Use custom name format for Snapshot copy	Select this check box, and then enter a custom name format that you want to use for the Snapshot copy name. For example, customtext__policy_hostname or resource_hostname. By default, a timestamp is appended to the Snapshot copy name.
Exclude archive log destinations from backup	Specify the destinations of the archive log files that you do not want to back up.

6. In the Policies page, perform the following steps:
 - a. Select one or more policies from the drop-down list.



You can also create a policy by clicking .

In the Configure schedules for selected policies section, the selected policies are listed.


- b. Click  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* window, configure the schedule, and then click **OK**.


policy_name is the name of the policy that you have selected.

The configured schedules are listed in the Applied Schedules column.

7. In the Verification page, perform the following steps:
 - a. Click **Load locators** to load the SnapMirror or SnapVault volumes to perform verification on secondary

storage.

- b. Click  in the Configure Schedules column to configure the verification schedule for all of the schedule types of the policy.
- c. In the Add Verification Schedules *policy_name* dialog box, perform the following actions:

If you want to...	Do this...
Run verification after backup	Select Run verification after backup .
Schedule a verification	<div><div>Select Run scheduled verification, and then select the schedule type from the drop-down list.</div><div><div>In a Flex ASM setup, you cannot perform verification operation on Leaf nodes if the cardinality is less than the number nodes in the RAC cluster.</div></div></div>

- d. Select **Verify on secondary location** to verify your backups on secondary storage.
- e. Click **OK**.

The configured verification schedules are listed in the Applied Schedules column.

8. 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. If you want to attach the report of the backup operation performed on the resource, and then select **Attach Job Report**.



For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command Set-SmSmtServer.

9. Review the summary, and then click **Finish**.

The database topology page is displayed.

10. Click **Back up Now**.

11. In the Backup page, perform the following steps:

- a. If you have applied multiple policies to 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.

- b. Click **Backup**.

12. Monitor the operation progress by clicking **Monitor > Jobs**.

After you finish

- In AIX setup, you can use the `lkdev` command to lock and the `rendev` command to rename the disks on which the database that was backed up was residing.

Locking or renaming of devices will not affect the restore operation when you restore using that backup.

- If the backup operation fails because database query execution time exceeded the timeout value, you should change the value of the `ORACLE_SQL_QUERY_TIMEOUT` and `ORACLE_PLUGIN_SQL_QUERY_TIMEOUT` parameters by running the `Set-SmConfigSettings` cmdlet:

After modifying the value of the parameters, restart the SnapCenter Plug-in Loader (SPL) service by running the following command `/opt/NetApp/snapcenter/spl/bin/spl restart`

- If the file is not accessible and the mount point is unavailable during the verification process, the operation might fail with error code DBV-00100 specified file. You should modify the values of the `VERIFICATION_DELAY` and `VERIFICATION_RETRY_COUNT` parameters in `sco.properties`.

After modifying the value of the parameters, restart the SnapCenter Plug-in Loader (SPL) service by running the following command `/opt/NetApp/snapcenter/spl/bin/spl restart`

- In MetroCluster configurations, SnapCenter might not be able to detect a protection relationship after a failover.
- If you are backing up application data on VMDKs and the Java heap size for the SnapCenter Plug-in for VMware vSphere is not large enough, the backup might fail.

To increase the Java heap size, locate the script file `/opt/netapp/init_scripts/scvservice`. In that script, the `do_start` method command starts the SnapCenter VMware plug-in service. Update that command to the following: `Java -jar -Xmx8192M -Xms4096M`.

Find more information

- [Unable to detect SnapMirror or SnapVault relationship after MetroCluster failover](#)
- [Oracle RAC One Node database is skipped for performing SnapCenter operations](#)
- [Failed to change the state of an Oracle 12c ASM database](#)
- [Customizable parameters for backup, restore and clone operations on AIX systems](#)



Back up Oracle database resource groups

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

You can back up a resource group on demand from the Resources page. If a resource group has a policy attached and a schedule configured, then backups occur automatically according to the schedule.

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 **Resource Group** from the **View** list.

You can search the resource group either by entering the resource group name in the search box, or by clicking , and then selecting the tag. You can then click  to close the filter pane.

3. In the Resource Groups page, select the resource group that you want to back up, and then click **Back up Now**.



If you have a federated resource group with two databases and one of the database has datafile on non-NetApp storage, the backup operation is aborted even though the other database is on NetApp storage.

4. In the Backup page, perform the following steps:
 - a. If you have associated multiple policies 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. Click **Backup**.
5. Monitor the operation progress by clicking **Monitor > Jobs**.

After you finish

- In AIX setup, you can use the `lkdev` command to lock and the `rendev` command to rename the disks on which the database that was backed up was residing.

Locking or renaming of devices will not affect the restore operation when you restore using that backup.

- If the backup operation fails because database query execution time exceeded the timeout value, you should change the value of the `ORACLE_SQL_QUERY_TIMEOUT` and `ORACLE_PLUGIN_SQL_QUERY_TIMEOUT` parameters by running the `Set-SmConfigSettings` cmdlet:

After modifying the value of the parameters, restart the SnapCenter Plug-in Loader (SPL) service by running the following command `/opt/NetApp/snapcenter/spl/bin/spl restart`

- If the file is not accessible and the mount point is unavailable during the verification process, the operation might fail with error code DBV-00100 specified file. You should modify the values of the `VERIFICATION_DELAY` and `VERIFICATION_RETRY_COUNT` parameters in `sco.properties`.

After modifying the value of the parameters, restart the SnapCenter Plug-in Loader (SPL) service by running the following command `/opt/NetApp/snapcenter/spl/bin/spl restart`

Back up Oracle databases using UNIX commands

The backup workflow includes planning, identifying the resources for backup, creating backup policies, creating resource groups and attaching policies, creating backups, and monitoring the operations.

What you will need

- You should have added the storage system connections and created the credential using the commands *Add-SmStorageConnection* and *Add-SmCredential*.
- You should have established the connection session with the SnapCenter Server using the command *Open-SmConnection*.

You can have only one SnapCenter account login session and the token is stored in the user home

directory.



The connection session is valid only for 24 hours. However, you can create a token with the `TokenNeverExpires` option to create a token that never expires and the session will always be valid.

About this task

You should execute the following commands to establish the connection with the SnapCenter Server, discover the Oracle database instances, add policy and resource group, backup and verify the backup.

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

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user: *Open-SmConnection*
2. Perform host resources discovery operation: *Get-SmResources*
3. Configure Oracle database credentials and preferred nodes for backup operation of a Real Application Cluster (RAC) database: *Configure-SmOracleDatabase*
4. Create a backup policy: *Add-SmPolicy*
5. Retrieve the information about the secondary (SnapVault or SnapMirror) storage location : *Get-SmSecondaryDetails*

This command retrieves the primary to secondary storage mapping details of a specified resource. You can use the mapping details to configure the secondary verification settings while creating a backup resource group.

6. Add a resource group to SnapCenter: *Add-SmResourceGroup*
7. Create a backup: *New-SmBackup*

You can poll the job using the `WaitForCompletion` option. If this option is specified, then the command continues to poll the server until the completion of the backup job.



8. Retrieve the logs from SnapCenter: *Get-SmLogs*





Monitor Oracle database backup operations

You can monitor the progress of different backup operations by using the `SnapCenterJobs` page. You might want to check the progress to determine when it is complete or if there is an issue.


About this task

The following icons appear on the Jobs page and indicate the corresponding state of the operations:


-  In progress
-  Completed successfully

-  Failed
-  Completed with warnings or could not start due to warnings
-  Queued
-  Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the Monitor page, click **Jobs**.
3. In the Jobs page, perform the following steps:
 - a. Click  to filter the list so that only backup operations are listed.
 - b. Specify the start and end dates.
 - c. From the **Type** drop-down list, select **Backup**.
 - d. From the **Status** drop-down, select the backup status.
 - e. Click **Apply** to view the operations completed successfully.
4. Select a backup job, and then click **Details** to view the job details.



Though the backup job status displays  , when you click on job details you might see that some of the child tasks of the backup operation are still in progress or marked with warning signs.

5. In the Job Details page, click **View logs**.


The **View logs** button displays the detailed logs for the selected operation.

Monitor data protection operations in the Activity pane

The Activity pane displays the five most recent operations performed. The Activity pane also displays when the operation was initiated and the status of the operation.

The Activity pane displays information regarding backup, restore, clone, and scheduled backup operations. If you are using Plug-in for SQL Server or Plug-in for Exchange Server, the Activity pane also displays information about the reseed operation.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. Click  on the Activity pane to view the five most recent operations.

When you click one of the operations, the operation details are listed in the Job Details page.

Cancel backup operations of Oracle databases

You can cancel backup operations that are either running, queued, or non-responsive.

You must be logged in as the SnapCenter Admin or job owner to cancel backup operations.

About this task

When you cancel a backup operation, the SnapCenter Server stops the operation and removes all the Snapshot copies from the storage if the backup created is not registered with SnapCenter Server. If the backup is already registered with SnapCenter Server, it will not roll back the already created Snapshot copy even after the cancellation is triggered.


- You can cancel only the log or full backup operation that are queued or running.
- You cannot cancel the operation after the verification has started.

If you cancel the operation before verification, the operation is canceled, and the verification operation will not be performed.

- You cannot cancel the backup operation after the catalog operations has started.
- You can cancel a backup operation from either the Monitor page or the Activity pane.
- In addition to using the SnapCenter GUI, you can use CLI commands to cancel operations.
- The **Cancel Job** button is disabled for operations that cannot be canceled.
- If you selected **All members of this role can see and operate on other members objects** in Users\Groups page while creating a role, you can cancel the queued backup operations of other members while using that role.

Step

Perform one of the following actions:

From the...	Action
Monitor page	a. In the left navigation pane, click Monitor > Jobs . b. Select the operation and click Cancel Job .
Activity pane	a. After initiating the backup job, click  on the Activity pane to view the five most recent operations. b. Select the operation. c. In the Job Details page, click Cancel Job .

Results

The operation is canceled, and the resource is reverted to the original state.

If the operation you canceled is non-responsive in the canceling or running state, you should run the Cancel-SmJob -JobID <int> -Force to forcefully stop the backup operation.




View Oracle database backups and clones in the Topology page

When you are preparing to back up or clone a resource, you might find it helpful to view a graphical representation of all backups and clones on the primary and secondary storage.

About this task

In the Topology page, you can see all of the backups and clones that are available for the selected resource or resource group. You can view the details of those backups and clones, and then select them to perform data protection operations.

You can review the following icons in the Manage Copies view to determine whether the backups and clones are available on the primary or secondary storage (Mirror copies or Vault copies).

-  displays the number of backups and clones that are available on the primary storage.
-  displays the number of backups and clones that are mirrored on the secondary storage using SnapMirror technology.
-  displays the number of backups and clones that are replicated on the secondary storage using SnapVault technology.

The number of backups displayed includes the backups deleted from the secondary storage. For example, if you created 6 backups using a policy to retain only 4 backups, the number of backups displayed is 6.



Clones of a backup of a version-flexible mirror on a mirror-vault type volume are displayed in the topology view, but the mirror backup count in the topology view does not include the version-flexible backup.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, either select the resource or resource group from the **View** drop-down list.
3. Select the resource either from the resource details view or from the resource group details view.

If the resource is protected, the Topology page of the selected resource is displayed.

4. Review the Summary card to see a summary of the number of backups and clones available on the primary and secondary storage.

The Summary Card section displays the total number of backups and clones and total number of log backups.

Clicking the **Refresh** button starts a query of the storage to display an accurate count.

5. In the Manage Copies view, click **Backups** or **Clones** from the primary or secondary storage to see details of a backup or clone.

The details of the backups and clones are displayed in a table format.

6. Select the backup from the table, and then click the data protection icons to perform restore, clone, mount, unmount, rename, catalog, uncatalog, and delete operations.



You cannot rename or delete backups that are on the secondary storage.

- If you have selected a log backup, you can only perform rename, mount, unmount, catalog, uncatalog, and delete operations.
- If you have cataloged the backup using Oracle Recovery Manager (RMAN), you cannot rename those cataloged backups.

7. If you want to delete a clone, select the clone from the table, and then click .

If the value assigned to `SnapmirrorStatusUpdateWaitTime` is less, the Mirror and Vault backup copies are not listed on the topology page even if data and log volumes are successfully protected. You should increase the value assigned to `SnapmirrorStatusUpdateWaitTime` using `Set-SmConfigSettings` PowerShell cmdlet.

The information regarding the parameters that can be used with the command and their descriptions can be obtained by running `Get-Help command_name`.

Alternatively, you can also refer to the [SnapCenter Software Command Reference Guide](#) or [SnapCenter Software Cmdlet Reference Guide](#).

Mount and unmount database backups

You can mount a single or multiple data and log only backups if you want to access the files in the backup. You can either mount the backup to the same host where the backup was created or to a remote host having same type of Oracle and host configurations. If you have manually mounted the backups, you should manually unmount the backups after completing the operation. At any given instance, a backup of a database can be mounted to any one of the host. While performing an operation, you can mount only a single backup.



In a Flex ASM setup, you cannot perform mount operation on Leaf nodes if the cardinality is less than the number nodes in the RAC cluster.

Mount a database backup

You should manually mount a database backup if you want to access the files in the backup.

What you will need

- If you have an Automatic Storage Management (ASM) database instance in an NFS environment and want to mount the ASM backups, you should have added the ASM disk path `/var/opt/snapcenter/sco/backup*//*/*` to the existing path defined in the `asm_diskstring` parameter.
- If you have an ASM database instance in an NFS environment and want to mount the ASM log backups as part of a recovery operation, you should have added the ASM disk path `/var/opt/snapcenter/scu/clones*//*/*` to the existing path defined in the `asm_diskstring` parameter.
- In the `asm_diskstring` parameter, you should configure `AFD:*` if you are using ASMFD or configure `ORCL:*` if you are using ASMLIB.



For information on how to edit the `asm_diskstring` parameter, see [How to add disk paths to asm_diskstring](#).

- You should configure the ASM credentials and the ASM port if it differs from that of the source database host while mounting the backup.


- If you want to mount to an alternate host, you must verify that the alternate host meets the following requirements:
 - Same UID and GID as that of the original host
 - Same Oracle version as that of the original host
 - Same OS distribution and version as that of the original host
 - For NVMe, NVMe util should be installed
- You should ensure that the LUN is not mapped to the AIX host using iGroup consisting of mixed protocols iSCSI and FC. For more information, see [Operation fails with error Unable to discover the device for LUN](#).

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 either from the database details view or from the resource group details view.

The database topology page is displayed.

4. From the Manage Copies view, select **Backups** either from the primary or secondary (mirrored or replicated) storage system.

5. Select the backup from the table, and then click .

6. In the Mount backups page, select the host on which you want to mount the backup from the **Choose the host to mount the backup** drop-down list.

The mount path `/var/opt/snapcenter/sco/backup_mount/backup_name/database_name` is displayed.

If you are mounting the backup of an ASM database, the mount path `+diskgroupname_SID_backupid` is displayed.

7. Click **Mount**.

After you finish

- You can run the following command to retrieve the information related to the mounted backup:

```
./sccli Get-SmBackup -BackupName backup_name -ListMountInfo
```

- If you have mounted an ASM database, you can run the following command to retrieve the information related to the mounted backup:

```
./sccli Get-Smbbackup -BackupNamediskgroupname_SID_backupid-listmountinfo
```

- To retrieve the backup ID, run the following command:

```
./sccli Get-Smbbackup-BackupNamebackup_name
```

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


Unmount a database backup

You can manually unmount a mounted database backup when you no longer want to access files on the backup.

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 either from the database details view or from the resource group details view.

The database topology page is displayed.


4. Select the backup that is mounted, and then click .
5. Click **OK**.

Restore and recover Oracle databases

Restore workflow

The restore workflow includes planning, performing the restore operations, and monitoring the operations.

The following workflow shows the sequence in which you must perform the restore operation:



Define a restore and recovery strategy for Oracle databases

You must define a strategy before you restore and recover your database so that you can perform restore and recover operations successfully.

Types of backups supported for restore and recovery operations

SnapCenter supports restore and recovery of different types of Oracle database backups.

- Online data backup
- Offline shutdown data backup
- Offline mount data backup



If you are restoring an offline shutdown or offline mount data backup, SnapCenter leaves the database in offline state. You should manually recover the database and reset the logs.

- Full backup
- Offline-mount backups of Data Guard standby databases
- Data-only online backups of Active Data Guard standby databases



You cannot perform recovery of Active Data Guard standby databases.

- Online data backups, online full backups, offline mount backups, and offline shutdown backups in a Real Application Clusters (RAC) configuration
- Online data backups, online full backups, offline mount backups, and offline shutdown backups in an Automatic Storage Management (ASM) configuration

Types of restore methods supported for Oracle databases

SnapCenter supports connect-and-copy or in-place restore for Oracle databases. During a restore operation, SnapCenter determines the restore method that is appropriate for the file system to be used for restore without any data loss.



SnapCenter does not support volume-based SnapRestore.

Connect-and-copy restore

If the database layout differs from the backup or if there are any new files after the backup was created, connect-and-copy restore is performed. In the connect-and-copy restore method, the following tasks are performed:

Steps

1. The volume is cloned from the Snapshot copy and the file system stack is built on the host using the cloned LUNs or volumes.
2. The files are copied from the cloned file systems to the original file systems.
3. The cloned file systems are then unmounted from the host and the cloned volumes are deleted from ONTAP.



For a Flex ASM setup (where the cardinality is less than the number nodes in the RAC cluster) or ASM RAC databases on VMDK or RDM, only connect-and-copy restore method is supported.

Even if you have forcefully enabled in-place restore, SnapCenter performs connect-and-copy restore in the following scenarios:

- Restore from secondary storage system and if Data ONTAP is earlier than 8.3
- Restore of ASM disk groups present on nodes of an Oracle RAC setup on which database instance is not configured
- In Oracle RAC setup, on any of the peer nodes if the ASM instance or the cluster instance is not running or if the peer node is down
- Restore of control files only

- Restore a subset of tablespaces residing on a ASM disk group
- Disk group is shared between data files, sp file, and password file
- SnapCenter Plug-in Loader (SPL) service is not installed or not running on the remote node in a RAC environment
- New nodes are added to the Oracle RAC and the SnapCenter Server is not aware of the newly added nodes

In-place restore

If the database layout is similar to the backup and has not undergone any configuration change on the storage and database stack, in-place restore is performed, wherein the restore of file or LUN is performed on ONTAP. SnapCenter supports only Single File SnapRestore (SFSR) as part of the in-place restore method.



Data ONTAP 8.3 or later supports in-place restore from secondary location.

If you want to perform in-place restore on the database, ensure that you have only datafiles on the ASM disk group. You must create a backup after any changes are made to the ASM disk group or in the physical structure of the database. After performing in-place restore, the disk group will contain the same number datafiles as at the time of backup.

The in-place restore will be applied automatically when disk group or mount point matches the following criteria:

- No new datafiles are added after backup (foreign file check)
- No addition, deletion, or recreation of ASM disk or LUN after backup (ASM disk group structural change check)
- No addition, deletion, or recreation of LUNs to LVM disk group (LVM disk group structural change check)



You can also forcefully enable in-place restore either using GUI, SnapCenter CLI, or PowerShell cmdlet to override the foreign file check and LVM disk group structural change check.

Performing In-place restore on ASM RAC

In SnapCenter, the node on which you perform restore is termed as primary node and all other nodes of the RAC on which ASM disk group resides are called peer nodes. SnapCenter changes the state of ASM disk group to dismount on all the nodes where the ASM disk group is in mount state before performing the storage restore operation. After the storage restore is complete, SnapCenter changes the state of ASM disk group as it was before the restore operation.

In SAN environments, SnapCenter removes devices from all the peer nodes and performs LUN unmap operation before storage restore operation. After storage restore operation, SnapCenter performs LUN map operation and constructs devices on all the peer nodes. In a SAN environment if the Oracle RAC ASM layout is residing on LUNs, then while restoring SnapCenter performs LUN unmap, LUN restore, and LUN map operations on all the nodes of the RAC cluster where the ASM disk group resides. Before restoring even if all the initiators of the RAC nodes were not used for the LUNs, after restoring SnapCenter creates a new iGroup with all the initiators of all the RAC nodes.

- If there is any failure during prerestore activity on peer nodes, SnapCenter automatically rolls back the ASM disk group state as it was before performing restore on peer nodes on which prerestore operation was successful. Rollback is not supported for the primary and the peer node on which the operation failed. Before attempting another restore you must manually fix the issue on the peer node and bring the ASM

disk group on the primary node back to mount state.

- If there is any failure during restore activity, then the restore operation fails and no roll back is performed. Before attempting another restore, you must manually fix the storage restore issue and bring the ASM disk group on the primary node back to mount state.
- If there is any failure during postrestore activity on any of the peer nodes, SnapCenter continues with the restore operation on the other peer nodes. You must manually fix the post restore issue on the peer node.

Types of restore operations supported for Oracle databases

SnapCenter enables you to perform different types of restore operations for Oracle databases.

Before restoring the database, backups are validated to identify whether any files are missing when compared to the actual database files.

Full restore

- Restores only the datafiles
- Restores only the control files
- Restores the datafiles and control files
- Restores datafiles, control files, and redo log files in Data Guard standby and Active Data Guard standby databases

Partial restore

- Restores only the selected tablespaces
- Restores only the selected pluggable databases (PDBs)
- Restores only the selected tablespaces of a PDB

Types of recovery operations supported for Oracle databases

SnapCenter enables you to perform different types of recovery operations for Oracle databases.

- The database up to the last transaction (all logs)
- The database up to a specific system change number (SCN)
- The database up to a specific date and time

You must specify the date and time for recovery based on the database host's time zone.

SnapCenter also provides the No recovery option for Oracle databases.



The plug-in for Oracle database does not support recovery if you have restored using a backup that was created with the database role as standby. You must always perform manual recovery for physical standby databases.

Limitations related to restore and recovery of Oracle databases

Before you perform restore and recovery operations, you must be aware of the limitations.

If you are using any version of Oracle from 11.2.0.4 to 12.1.0.1, the restore operation will be in hung state when you run the *renamedg* command. You can apply the Oracle patch 19544733 to fix this issue.

The following restore and recovery operations are not supported:

- Restore and recovery of tablespaces of the root container database (CDB)
- Restore of temporary tablespaces and temporary tablespaces associated with PDBs
- Restore and recovery of tablespaces from multiple PDBs simultaneously
- Restore of log backups
- Restore of backups to a different location
- Restore of redo log files in any configuration other than Data Guard standby or Active Data Guard standby databases
- Restore of SPFILE and Password file
- When you perform a restore operation on a database that was re-created using the preexisting database name on the same host, was managed by SnapCenter, and had valid backups, the restore operation overwrites the newly created database files even though the DBIDs are different.

This can be avoided by performing either of following actions:

- Discover the SnapCenter resources after the database is re-created
- Create a backup of the re-created database

Limitations related to point-in-time recovery of tablespaces

- Point-in-time recovery (PITR) of SYSTEM, SYSAUX, and UNDO tablespaces is not supported
- PITR of tablespaces cannot be performed along with other types of restore
- If a tablespace is renamed and you want to recover it to a point before it was renamed, you should specify the earlier name of the tablespace
- If constraints for the tables in one tablespace are contained in another tablespace, you should recover both the tablespaces
- If a table and its indexes are stored in different tablespaces, then the indexes should be dropped before performing PITR
- PITR cannot be used to recover the current default tablespace
- PITR cannot be used to recover tablespaces containing any of the following objects:
 - Objects with underlying objects (such as materialized views) or contained objects (such as partitioned tables) unless all the underlying or contained objects are in the recovery set

Additionally, if the partitions of a partitioned table are stored in different tablespaces, then you should either drop the table before performing PITR or move all the partitions to the same tablespace before performing PITR.

- Undo or rollback segments
- Oracle 8 compatible advanced queues with multiple recipients
- Objects owned by the SYS user

Examples of these types of objects are PL/SQL, Java classes, call out programs, views, synonyms, users, privileges, dimensions, directories, and sequences.

Sources and destinations for restoring Oracle databases

You can restore an Oracle database from a backup copy on either primary storage or secondary storage. You can only restore databases to the same location on the same database instance. However, in Real Application Cluster (RAC) setup, you can restore databases to other nodes.

Sources for restore operations

You can restore databases from a backup on primary storage or secondary storage. If you want to restore from a backup on the secondary storage in a multiple mirror configuration, you can select the secondary storage mirror as the source.

Destinations for restore operations

You can only restore databases to the same location on the same database instance.

In a RAC setup, you can restore RAC databases from any nodes in the cluster.

Predefined environment variables for restore specific prescript and postscript

SnapCenter allows you to use the predefined environment variables when you execute the prescript and postscript while restoring a database.

Supported predefined environment variables for restoring a database

- **SC_JOB_ID** specifies the job ID of the operation.

Example: 257

- **SC_ORACLE_SID** specifies the system identifier of the database.

If the operation involves multiple databases, this will contain database names separated by pipe.

Example: NFSB31

- **SC_HOST** specifies the host name of the database.

This parameter will be populated for application volumes.

Example: scsmohost2.gdl.englabe.netapp.com

- **SC_OS_USER** specifies the operating system owner of the database.

Example: oracle

- **SC_OS_GROUP** specifies the operating system group of the database.

Example: oinstall

- **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_BACKUP_ID** specifies the ID of the backup.

This parameter will be populated for application volumes.

Examples:

- If the database is not running in ARCHIVELOG mode: DATA@203|LOG@205
- If the database is running in ARCHIVELOG mode: DATA@203|LOG@205,206,207

- **SC_RESOURCE_GROUP_NAME** specifies the name of the resource group.

Example: RG1

- **SC_ORACLE_HOME** specifies the path of the Oracle home directory.

Example: /ora01/app/oracle/product/18.1.0/db_1

- **SC_RECOVERY_TYPE** specifies the files that are recovered and also the recovery scope.

Example:

RESTORESCOPE:usingBackupControlfile=false|RECOVERYSCOPE:allLogs=true,noLogs=false,untiltime=false,untilscn=false.

For information about delimiters, see [Supported delimiters](#).

Requirements for restoring an Oracle database

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

- You should have defined your restore and recovery strategy.
- The SnapCenter administrator should have assigned you the storage virtual machines (SVMs) for both the source volumes and destination volumes if you are replicating Snapshot copies to a mirror or vault.
- If archive logs are pruned as part of backup, you should have manually mounted the required archive log backups.
- If you want to restore Oracle databases that are residing on a Virtual Machine Disk (VMDK), you should ensure that the guest machine has the required number of free slots for allocating the cloned VMDKs.
- You should ensure that all data volumes and archive log volumes belonging to the database are protected if secondary protection is enabled for that database.
- You should ensure that the RAC One Node database is in "nomount" state to perform control file or full database restore.
- If you have an ASM database instance in NFS environment, you should add the ASM disk path `/var/opt/snapcenter/scu/clones/*/*` to the existing path defined in the `asm_diskstring` parameter to successfully mount the ASM log backups as part of recovery operation.

- In the `asm_diskstring` parameter, you should configure `AFD:*` if you are using ASMFD or configure `ORCL:*` if you are using ASMLIB.



For information on how to edit the `asm_diskstring` parameter, see [How to add disk paths to `asm_diskstring`](#)

- You should configure the static listener in the **listener.ora** file available at `$ORACLE_HOME/network/admin` for non ASM databases and `$GRID_HOME/network/admin` for ASM databases if you have disabled OS authentication and enabled Oracle database authentication for an Oracle database, and want to restore the datafiles and control files of that database.
- You should increase value of `SCORestoreTimeout` parameter by running the `Set- SmConfigSettings` command if the database size is in terabytes (TB).
- You should ensure that all the licenses required for vCenter are installed and up to date.

If the licenses are not installed or up to date, a warning message is displayed. If you ignore the warning and proceed, restore from RDM fails.

- You should ensure that the LUN is not mapped to the AIX host using iGroup consisting of mixed protocols iSCSI and FC. For more information, see [Operation fails with error Unable to discover the device for LUN](#).

Restore and recover Oracle database

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.

About this task

Recovery is performed using the archive logs available at the configured archive log location. If the database is running in ARCHIVELOG mode, Oracle database saves the filled groups of redo log files to one or more offline destinations, known collectively as the archived redo log. SnapCenter identifies and mounts optimal number of log backups based on the specified SCN, selected date and time, or all logs option. If the archive logs required for recovery are not available at the configured location, you should mount the Snapshot copy containing the logs and specify the path as external archive logs.

If you migrate ASM database from ASMLIB to ASMFD, then the backups created with ASMLIB cannot be used to restore the database. You should create backups in the ASMFD configuration and use those backups to restore. Similarly, if ASM database is migrated from ASMFD to ASMLIB, you should create backups in the ASMLIB configuration to restore.

When you restore a database, an operational lock file (`.sm_lock_dbsid`) is created on the Oracle database host in the `$ORACLE_HOME/dbs` directory to avoid multiple operations being executed on the database. After the database has been restored, the operational lock file is automatically removed.




Restore of SPFILE and Password file is not supported.

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 from either the database details view or the resource group details view.


The database topology page is displayed.


4. From the Manage Copies view, select **Backups** from either the primary or the secondary (mirrored or replicated) storage systems.
5. Select the backup from the table, and then click  .
6. In the Restore Scope page, perform the following tasks:
 - a. If you have selected a backup of a database in a Real Application Clusters (RAC) environment, select the RAC node.
 - b. When you select a mirrored or vault data:
 - if there are no log backup at mirror or vault, nothing is selected and the locators are empty.
 - if log backups exist in mirror or vault, the latest log backup is selected and corresponding locator is displayed.



If the selected log backup exists in both mirror and vault location, both the locators are displayed.

- c. Perform the following actions:

If you want to restore...	Do this...
All the datafiles of the database	<p>Select All Datafiles.</p> <p>Only the datafiles of the database are restored. The control files, archive logs, or redo log files are not restored.</p>
Tablespaces	<p>Select Tablespaces.</p> <p>You can specify the tablespaces that you want to restore.</p>
Control files	<p>Select Control files.</p> <div><p>While restoring control files, ensure that the directory structure either exists or should be created with the correct user and group ownerships, if any, to allow the files to be copied to the target location by the restore process. If the directory does not exist, the restore job will fail.</p></div>

If you want to restore...	Do this...
Redo log files	<p>Select Redo log files.</p> <p>This option is available only for Data Guard standby or Active Data Guard standby databases.</p> <div>  <p>Redo log files are not backed up for non Data Guard databases. For non Data Guard databases the recovery is performed using archive logs.</p> </div>
Pluggable databases (PDBs)	Select Pluggable databases , and then specify the PDBs that you want to restore.
Pluggable database (PDB) tablespaces	<p>Select Pluggable database (PDB) tablespaces, and then specify the PDB and the tablespaces of that PDB that you want to restore.</p> <p>This option is available only if you have selected a PDB for restore.</p>

- d. 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.


The various states of a database from higher to lower are open, mounted, started, and shutdown. You must select this check box if the database is in a higher state but the state must be changed to a lower state to perform a restore operation. If the database is in a lower state but the state must be changed to a higher state to perform the restore operation, the database state is changed automatically even if you do not select the check box.

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

- e. 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, 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 (System Change Number) .

If you...	Do this...
Want to recover to a specific data and time	<p>Select Date and Time.</p> <p>You must specify the date and time of the database host's time zone.</p>
Do not want to recover	Select No recovery .
Want to specify any external archive log locations	<p>If the database is running in ARCHIVELOG mode, SnapCenter identifies and mounts optimal number of log backups based on the specified SCN, selected date and time, or all logs option.</p> <p>If you still want to specify the location of the external archive log files, select Specify external archive log locations.</p> <p>If archive logs are pruned as part of backup, and you have manually mounted the required archive log backups, you must specify the mounted backup path as the external archive log location for recovery.</p> <div>  <p>You should verify the path and contents of the mount path before listing it as an external log location.</p> </div> <ul style="list-style-type: none"> • NetApp Technical Report 4591: Database Data Protection Backup, Recovery, Replication, and DR • Operation fails with ORA-00308 error

You cannot perform restore with recovery from secondary backups if archive log volumes are not protected but data volumes are protected. You can restore only by selecting **No recovery**.

If you are recovering a RAC database with the open database option selected, only the RAC instance where the recovery operation was initiated is brought back to the open state.



Recovery is not supported for Data Guard standby and Active Data Guard standby databases.

8. In the PreOps page, enter the path and the arguments of the prescript that you want to run before the restore operation.

You must store the prescripts either in the `/var/opt/snapcenter/spl/scripts` path or in any folder inside this path. By default, the `/var/opt/snapcenter/spl/scripts` path is populated. If you have created any folders inside this path to store the scripts, you must specify those folders in the path.

You can also specify the script timeout value. The default value is 60 seconds.

SnapCenter allows you to use the predefined environment variables when you execute the prescript and

postscript. [Learn more](#)

9. In the PostOps page, perform the following steps:

- a. Enter the path and the arguments of the postscript that you want to run after the restore operation.

You must store the postscripts either in `/var/opt/snapcenter/spl/scripts` or in any folder inside this path. By default, the `/var/opt/snapcenter/spl/scripts` path is populated. If you have created any folders inside this path to store the scripts, you must specify those folders in the path.



If the restore operation fails, postscripts will not be executed and cleanup activities will be triggered directly.

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

After restoring a container database (CDB) with or without control files, or after restoring only CDB control files, if you specify to open the database after recovery, then only the CDB is opened and not the pluggable databases (PDB) in that CDB.

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



After restoring a user tablespace with control files, a system tablespace with or without control files, or a PDB with or without control files, only the state of the PDB related to the restore operation is changed to the original state. The state of the other PDBs that were not used for restore are not changed to the original state because the state of those PDBs were not saved. You must manually change the state of the PDBs that were not used for restore.

10. In the Notification page, from the **Email preference** drop-down list, select the scenarios in which you want to send the email notifications.

You must also specify the sender and receiver email addresses, and the subject of the email. If you want to attach the report of the restore operation performed, you must select **Attach Job Report**.



For email notification, you must have specified the SMTP server details by using the either the GUI or the PowerShell command `Set-SmSmtServer`.

11. Review the summary, and then click **Finish**.

12. Monitor the operation progress by clicking **Monitor > Jobs**.

For more information

- [Oracle RAC One Node database is skipped for performing SnapCenter operations](#)
- [Failed to restore from a secondary SnapMirror or SnapVault location](#)
- [Failed to restore from a backup of an orphan incarnation](#)
- [Customizable parameters for backup, restore and clone operations on AIX systems](#)

Restore and recover tablespaces using point-in-time recovery

You can restore a subset of tablespaces that has 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.

What you will need

The backups that are required to perform PITR of tablespaces should be cataloged and mounted.

About this task

During PITR operation, RMAN creates an auxiliary instance at the specified auxiliary destination. The auxiliary destination could be a mount point or ASM disk group. If there is sufficient space in the mounted location, you can reuse one of the mounted locations instead of a dedicated mount point.

You should specify the date and time or SCN and the tablespace is restored on the source database.

You can select and restore multiple tablespaces residing on ASM, NFS, and SAN environments. For example, if tablespaces TS2 and TS3 reside on NFS and TS4 reside on SAN, you can perform on single PITR operation to restore all the tablespaces.



In a RAC setup, you can perform PITR of tablespaces from any node of the RAC.

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) either from the database details view or the resource group details view.

The database topology page is displayed.

4. From the Manage Copies view, select **Backups** from either the primary or the secondary (mirrored or replicated) storage systems.

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

5. Select the cataloged backup, and then click .
6. In the Restore Scope page, perform the following tasks:
 - a. If you have selected a backup of a database in a Real Application Clusters (RAC) environment, select the RAC node.
 - b. Select **Tablespaces**, and then specify the tablespaces you want to restore.



You cannot perform PITR on SYSAUX, SYSTEM, and UNDO 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, perform one of the following actions:
 - If you want to recover to a specific System Change Number (SCN), select **Until SCN** and specify the SCN and auxiliary destination.
 - If you want to recover to a specific date and time, select **Date and Time** and specify the date and time and the auxiliary destination.

SnapCenter identifies and then mounts and catalogs the optimal number of data and log backups required to perform PITR based on specified SCN or the selected date and time.

8. In the PreOps page, enter the path and the arguments of the prescript that you want to run before the restore operation.

You should store the prescripts either in the `/var/opt/snapcenter/spl/scripts` path or in any folder inside this path. By default, the `/var/opt/snapcenter/spl/scripts` path is populated. If you have created any folders inside this path to store the scripts, you must specify those folders in the path.

You can also specify the script timeout value. The default value is 60 seconds.

SnapCenter allows you to use the predefined environment variables when you execute the prescript and postscript. [Learn more](#)

9. In the PostOps page, perform the following steps:
 - a. Enter the path and the arguments of the postscript that you want to run after the restore operation.



If the restore operation fails, postscripts will not be executed and cleanup activities will be triggered directly.

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

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

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.

What you will need

The backups that are required to perform PITR of a PDB should be cataloged and mounted.



In a RAC setup, you should manually close the PDB (changing the state to MOUNTED) on all the nodes of the RAC setup.

About this task

During PITR operation, RMAN creates an auxiliary instance at the specified auxiliary destination. The auxiliary destination could be a mount point or ASM disk group. If there is sufficient space in the mounted location, you can reuse one of the mounted locations instead of a dedicated mount point.

You should specify the date and time or SCN to perform PITR of the PDB. RMAN can recover READ WRITE, READ ONLY, or dropped PDBs including datafiles.

You can restore and recover only:

- one PDB at a time
- one tablespace in a PDB
- multiple tablespaces of the same PDB



In a RAC setup, you can perform PITR of tablespaces from any node of the RAC.


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) either from the database details view or the resource group details view.

The database topology page is displayed.

4. From the Manage Copies view, select **Backups** from either the primary or the secondary (mirrored or replicated) storage systems.

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

5. Select the cataloged backup, and then click .
6. In the Restore Scope page, perform the following tasks:
 - a. If you have selected a backup of a database in a Real Application Clusters (RAC) environment, select the RAC node.
 - b. Depending on whether you want to restore the PDB or tablespaces in a PDB, perform one of the actions:

If you want to...	Steps...
Restore a PDB	<ol style="list-style-type: none"> i. Select Pluggable databases (PDBs). ii. Specify the PDB you want to restore. <div>  <p>You cannot perform PITR on PDB\$SEED database.</p> </div>
Restore tablespaces in a PDB	<ol style="list-style-type: none"> i. Select Pluggable database (PDB) tablespaces. ii. Specify the PDB. iii. Specify either a single tablespace or multiple tablespaces you want to restore. <div>  <p>You cannot perform PITR on SYSAUX, SYSTEM, and UNDO tablespaces.</p> </div>

- 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, perform one of the following actions:

- If you want to recover to a specific System Change Number (SCN), select **Until SCN** and specify the SCN and auxiliary destination.
- If you want to recover to a specific date and time, select **Date and Time** and specify the date and time and the auxiliary destination.

SnapCenter identifies and then mounts and catalogs the optimal number of data and log backups required to perform PITR based on specified SCN or the selected date and time.

8. In the PreOps page, enter the path and the arguments of the prescript that you want to run before the restore operation.

You should store the prescripts either in the `/var/opt/snapcenter/spl/scripts` path or in any folder inside this path. By default, the `/var/opt/snapcenter/spl/scripts` path is populated. If you have created any folders inside this path to store the scripts, you must specify those folders in the path.

You can also specify the script timeout value. The default value is 60 seconds.

SnapCenter allows you to use the predefined environment variables when you execute the prescript and postscript. [Learn more](#)

9. In the PostOps page, perform the following steps:

- a. Enter the path and the arguments of the postscript that you want to run after the restore operation.



If the restore operation fails, postscripts will not be executed and cleanup activities will be triggered directly.

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

In a RAC setup, the PDB will be opened only on the node where the database was recovered. You should manually open the recovered PDB on all the other nodes of the RAC setup.

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

Restore and recover Oracle databases using UNIX commands

The restore and recovery workflow includes planning, performing the restore and recovery operations, and monitoring the operations.

About this task

You should execute the following commands to establish the connection with the SnapCenter Server, list the backups and retrieve its information, and restore the backup.

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

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user: *Open-SmConnection*
2. Retrieve the information about the backups that you want to restore: *Get-SmBackup*
3. Retrieve the detailed information about the specified backup: *Get-SmBackupDetails*

This command retrieves the detailed information about the backup of a specified resource with a given backup ID. The information includes database name, version, home, start and end SCN, tablespaces, pluggable databases, and its tablespaces.

4. Restore data from the backup: *Restore-SmBackup*







Monitor Oracle database restore operations

You can monitor the progress of different SnapCenter restore operations by using the Jobs page. You might want to check the progress of an operation to determine when it is complete or if there is an issue.


About this task

Post-restore states describe the conditions of the resource after a restore operation and any further restore actions that you can take.

The following icons appear on the Jobs page, and indicate the state of the operation:

-  In progress
-  Completed successfully
-  Failed
-  Completed with warnings or could not start due to warnings
-  Queued
-  Canceled


Steps

1. In the left navigation pane, click **Monitor**.
2. In the Monitor page, click **Jobs**.
3. In the Jobs page, perform the following steps:
 - a. Click  to filter the list so that only restore operations are listed.
 - b. Specify the start and end dates.
 - c. From the **Type** drop-down list, select **Restore**.
 - d. From the **Status** drop-down list, select the restore status.
 - e. Click **Apply** to view the operations that have been completed successfully.
4. Select the restore job, and then click **Details** to view the job details.

5. In the Job Details page, click **View logs**.

The **View logs** button displays the detailed logs for the selected operation.



After the volume based restore operation, the backup metadata is deleted from the SnapCenter repository but the backup catalog entries remain in SAP HANA catalog. Though the restore job status displays , you should click on job details to see the warning sign of some of the child tasks. Click on the warning sign and delete the indicated backup catalog entries.

Cancel Oracle database restore operations

You can cancel restore jobs that are queued.


You should be logged in as the SnapCenter Admin or job owner to cancel restore operations.

About this task

- You can cancel a queued restore operation from either the **Monitor** page or the **Activity** pane.
- You cannot cancel a running restore operation.
- You can use the SnapCenter GUI, PowerShell cmdlets, or CLI commands to cancel the queued restore operations.
- The **Cancel Job** button is disabled for restore operations that cannot be canceled.
- If you selected **All members of this role can see and operate on other members objects** in Users\Groups page while creating a role, you can cancel the queued restore operations of other members while using that role.

Step

Perform one of the following actions:

From the...	Action
Monitor page	<ol style="list-style-type: none">a. In the left navigation pane, click Monitor > Jobs.b. Select the job and click Cancel Job.
Activity pane	<ol style="list-style-type: none">a. After initiating the restore operation, click  on the Activity pane to view the five most recent operations.b. Select the operation.c. In the Job Details page, click Cancel Job.

Clone Oracle database

Clone workflow


The clone workflow includes planning, performing the clone operation, and monitoring the

operation.

You might clone databases for the following reasons:

- To test functionality that has to 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.

The following workflow shows the sequence in which you must perform the clone operation:



Define a clone strategy for Oracle databases

Defining a strategy before cloning your database ensures that the cloning operation is successful.

Types of backups supported for cloning

SnapCenter supports cloning of different types of backups of Oracle databases.

- Online data backup
- Online full backup
- Offline mount backup
- Offline shutdown backup
- Backups of Data Guard standby databases and Active Data Guard standby databases
- Online data backups, online full backups, offline mount backups, and offline shutdown backups in a Real Application Clusters (RAC) configuration
- Online data backups, online full backups, offline mount backups, and offline shutdown backups in an Automatic Storage Management (ASM) configuration



SAN configurations are not supported if user_friendly_names option in the multipath configuration file is set to yes.



Cloning of archive log backups is not supported.

Types of cloning supported for Oracle databases

In an Oracle database environment, SnapCenter supports cloning of a database backup. You can clone the backup from primary and secondary storage systems.

The SnapCenter Server uses NetApp FlexClone technology to clone backups.

You can refresh a clone by running the "Refresh-SmClone" command. This command creates a backup of the database, deletes the existing clone, and creates a clone with the same name.



The clone refresh operation can only be performed using the UNIX commands.

Clone naming conventions for Oracle databases

From SnapCenter 3.0, the naming convention used for clones of file systems is different from the clones of ASM disk groups.

- The naming convention for SAN or NFS file systems is `FileSystemNameofsourcedatabase_CLONESID`.
- The naming convention for ASM disk groups is `SC_HASHCODEofDISKGROUP_CLONESID`.

`HASHCODEofDISKGROUP` is an automatically generated number (2 to 10 digits) that is unique for each ASM disk group.

Limitations of cloning Oracle databases

You should be aware of the limitations of clone operations before you clone the databases.

- If you are using any version of Oracle from 11.2.0.4 to 12.1.0.1, the clone operation will be in hung state when you run the *renamedg* command. You can apply the Oracle patch 19544733 to fix this issue.
- Cloning of databases from a LUN that is directly attached to a host (for instance, by using Microsoft iSCSI Initiator on a Windows host) to a VMDK or an RDM LUN on the same Windows host, or another Windows host, or vice versa, is not supported.
- The root directory of the volume mount point cannot be a shared directory.
- If you move a LUN that contains a clone to a new volume, the clone cannot be deleted.

Predefined environment variables for clone specific prescript and postscript

SnapCenter allows you to use the predefined environment variables when you execute the prescript and postscript while cloning a database.

Supported predefined environment variables for 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_AV_NAME** specifies the names of the application volumes.

Example: AV1|AV2

- **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

For information about delimiters, see [Supported delimiters](#).

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 using SnapCenter.

You should have successfully created online data and log backups or offline (mount or shutdown) backups for the cloning operation to succeed.

- If you want to customize the control file or redo log file paths, you should have preprovisioned the required file system or Automatic Storage Management (ASM) disk group.

By default, redo log and control files of the cloned database are created on the ASM disk group or the file system provisioned by SnapCenter for the data files of the clone database.

- If you are using ASM over NFS, you should add `/var/opt/snapcenter/scu/clones/*/*` to the existing path defined in the `asm_diskstring` parameter.
- In the `asm_diskstring` parameter, you should configure `AFD:*` if you are using ASMFD or configure `ORCL:*` if you are using ASMLIB.

For information on how to edit the `asm_diskstring` parameter, see [How to add disk paths to asm_diskstring](#).

- If you are creating the clone on an alternate host, the alternate host should meet the following requirements:
 - SnapCenter Plug-in for Oracle Database should be installed on the alternate host.
 - The clone host should be able to discover LUNs from primary or secondary storage.
 - If you are cloning from primary storage or secondary (Vault or Mirror) storage to an alternate host, then make sure that an iSCSI session is either established between the secondary storage and the alternate host, or zoned properly for FC.
 - If you are cloning from Vault or Mirror storage to the same host, then make sure that an iSCSI session is either established between the Vault or Mirror storage and the host, or zoned properly for FC.
 - If you are cloning in a virtualized environment, ensure that an iSCSI session is either established between the primary or secondary storage and the ESX server hosting the alternate host, or zoned properly for FC.
For information, refer to [host utilities documentation](#).
 - 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 disk group should be provisioned prior to the clone operation if you want to place archive log files of the cloned database in a dedicated ASM disk group.
 - The name of the data disk group can be configured, but ensure that the name is not used by any

other ASM disk group on the host where the clone will be performed.

Data files residing on the ASM disk group are provisioned as part of the SnapCenter clone workflow.

- For NVMe, NVMe util should be installed
- The protection type for the data LUN and the log LUN, such as mirror, vault, or mirror-vault, should be the same to discover secondary locators during cloning to an alternate host using log backups.
- You should set the value of `exclude_seed_cdb_view` to `FALSE` in the source database parameter file to retrieve seed PDB related information for cloning a backup of `12_c_` database.

The seed PDB is a system-supplied template that the CDB can use to create PDBs. The seed PDB is named `PDB$SEED`. For information about `PDB$SEED`, see the Oracle Doc ID 1940806.1.



You should set the value before backing up `12_c_` database.

- SnapCenter supports backup of file systems that are managed by the autofs subsystem. If you are cloning the database, ensure that data mount points are not under the root of the autofs mount point because the root user of the plug-in host does not have permission to create directories under the root of the autofs mount point.

If control and redo log files are under data mount point, you should modify the control file path, and then redo log file path accordingly.



You can manually register the new cloned mount points with the autofs subsystem. The new cloned mount points will not be registered automatically.

- If you have a TDE (auto login) and want to clone the database on the same or alternate host, you should copy wallet (key files) under `/etc/ORACLE/WALLET/$ORACLE_SID` from the source database to the cloned database.
- You should set the value of `use_lvmetad = 0` in `/etc/lvm/lvm.conf` and stop the `lvm2-lvmetad` service to successfully perform cloning in storage area network (SAN) environments on Oracle Linux 7 or later or Red Hat Enterprise Linux (RHEL) 7 or later.
- You should install the 13366202 Oracle patch if you are using Oracle database 11.2.0.3 or later and the database ID for the auxiliary instance is changed using an NID script.
- You should ensure that the aggregates hosting the volumes should be in the assigned aggregates list of the storage virtual machine (SVM).
- For NVMe, if any target port has to be excluded from connecting, you should add the target node name and port name in the `/var/opt/snapcenter/scu/etc/nvme.conf` file.

If the file does not exist, you should create the file as shown in the example below:

```
blacklist {
nn-0x<target_node_name_1>:pn-0x<target_port_name_1>
nn-0x<target_node_name_2>:pn-0x<target_port_name_2>
}
```

- You should ensure that the LUN is not mapped to the AIX host using iGroup consisting of mixed protocols iSCSI and FC. For more information, see [Operation fails with error Unable to discover the device for LUN](#).

Clone an Oracle database backup

You can use SnapCenter to clone an Oracle database using the backup of the database.

About this task

The cloning operation creates a copy of the database data files, and creates new online redo log files and control files. The database can be optionally recovered to a specified time, based on the specified recovery options.



Cloning fails if you try to clone a backup that was created on a Linux host to an AIX host or vice-versa.

SnapCenter creates a stand-alone database when cloned from an Oracle RAC database backup. SnapCenter supports creating clone from the backup of a Data Guard standby and Active Data Guard standby databases.

During cloning, SnapCenter mounts the optimal number of log backups based on SCN or dat and time for recovery operations. After recovery, the log backup is unmounted. All such clones are mounted under `/var/opt/snapcenter/scu/clones/`. If you are using ASM over NFS, you should add `/var/opt/snapcenter/scu/clones/**` to the existing path defined in the `asm_diskstring` parameter.

While cloning a backup of an ASM database in a SAN environment, udev rules for the cloned host devices are created at `/etc/udev/rules.d/999-scu-netapp.rules`. These udev rules associated with the cloned host devices are deleted when you delete the clone.





In a Flex ASM setup, you cannot perform clone operation on Leaf nodes if the cardinality is less than the number nodes in the RAC cluster.

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 either from the database details view or from the resource group details view.

The database topology page is displayed.

4. From the Manage Copies view, select the backups either from Local copies (primary), Mirror copies (secondary), or Vault copies (secondary).
5. Select the Data backup from the table, and then click .
6. In the Name page, perform one of the following actions:

If you want to...	Steps...
Clone a database (CDB or non CDB)	<p>a. Specify the SID of the clone.</p> <p>The clone SID is not available by default, and the maximum length of the SID is 8 characters.</p> <div>  <p>You should ensure that no database with the same SID exists on the host where the clone will be created.</p> </div>
Clone a pluggable database (PDB)	<p>a. Select PDB Clone.</p> <p>b. Specify the PDB that you want to clone.</p> <p>c. Specify the name of cloned PDB. For the detailed steps to clone a PDB, see Clone a pluggable database.</p>

When you select a mirrored or vault data:



- if there are no log backup at mirror or vault, nothing is selected and the locators are empty.
- if log backups exist in mirror or vault, the latest log backup is selected and corresponding locator is displayed.






If the selected log backup exists in both mirror and vault location, both the locators are displayed.

7. In the Locations page, perform the following actions:

For this field...	Do this...
Clone host	<p>By default, the source database host is populated.</p> <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>

For this field...	Do this...
Datafile locations	<p>By default, the datafile location is populated.</p> <p>The SnapCenter default naming convention for SAN or NFS file systems is FileSystemNameofsourcedatabase_CLONESID.</p> <p>The SnapCenter default naming convention for ASM disk groups is SC_HASHCODEofDISKGROUP_CLONESID. The HASHCODEofDISKGROUP is an automatically generated number (2 to 10 digits) that is unique for each ASM disk group.</p> <div data-bbox="873 646 928 703">  </div> <div data-bbox="987 604 1435 743"> <p>If you are customizing the ASM disk group name, ensure that the name length adheres to the maximum length supported by Oracle.</p> </div> <p>If you want to specify a different path, you must enter the datafile mount points or ASM disk group names for clone database. When you customize the datafile path, you must also change the control file and redo log file ASM disk group names or file system either to the same name used for data files or to an existing ASM disk groups or file system.</p>
Control files	<p>By default, the control file path is populated.</p> <p>The control files are placed in the same ASM disk group or file system as that of the data files. If you want to override the control file path, you can provide a different control file path.</p> <div data-bbox="873 1325 928 1381">  </div> <div data-bbox="987 1318 1383 1388"> <p>The file system or the ASM disk group should exist on the host.</p> </div> <p>By default, the number of control files will be same as that of the source database. You can modify the number of control files but a minimum of one control file is required to clone the database.</p> <p>You can customize the control file path to a different file system (existing) than that of the source database.</p>

For this field...	Do this...
Redo logs	<p>By default, the redo log file group, path, and their sizes are populated.</p> <p>The redo logs are placed in the same ASM disk group or file system as that of the data files of the cloned 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..</p> <div>  <p>The new file system or the ASM disk group should exist on the host.</p> </div> <p>By default, the number of redo log groups, redo log files, and their sizes will be same as that of the source database. You can modify the following parameters:</p> <ul style="list-style-type: none"> • Number of redo log groups <div>  <p>A minimum of two redo log groups are required to clone the database.</p> </div> <ul style="list-style-type: none"> • Redo log files in each group and their path <p>You can customize the redo log file path to a different file system (existing) than that of the source database.</p> <div>  <p>A minimum of one redo log file is required in the redo log group to clone the database.</p> </div> <ul style="list-style-type: none"> • Sizes of the redo log file

8. On the Credentials page, perform the following actions:

For this field...	Do this...
Credential name for sys user	<p>Select the Credential to be used for defining the sys user password of the clone database.</p> <p>If SQLNET.AUTHENTICATION_SERVICES is set to NONE in sqlnet.ora file on the target host, you should not select None as the Credential in the SnapCenter GUI.</p>

For this field...	Do this...
ASM Instance Credential name	<p>Select None if OS authentication is enabled for connecting to the ASM instance on the clone host.</p> <p>Otherwise, select the Oracle ASM credential configured with either “sys” user or an user having “sysasm” privilege applicable to the clone host.</p>

The Oracle home, user name, and group details are automatically populated from the source database. You can change the values based on the Oracle environment of the host where the clone will be created.


9. In the PreOps page, perform the following steps:

- a. Enter the path and the arguments of the prescript that you want to run before the clone operation.

You must store the prescript either in `/var/opt/snapcenter/spl/scripts` or in any folder inside this path. By default, the `/var/opt/snapcenter/spl/scripts` path is populated. If you have placed the script in any folder inside this path, you need to provide the complete path up to the folder where the script is placed.

SnapCenter allows you to use the predefined environment variables when you execute the prescript and postscript. [Learn more](#)

- b. In the Database Parameter settings section, modify the values of prepopulated database parameters that are used to initialize the database.

You can add additional parameters by clicking .

If you are using Oracle Standard Edition and the database is running in Archive log mode or you want restore a database from archive redo log, add the parameters and specify the path.

- LOG_ARCHIVE_DEST
- LOG_ARCHIVE_DUPLEX_DEST



Fast recovery area (FRA) is not defined in the prepopulated database parameters. You can configure FRA by adding the related parameters.



The default value of `log_archive_dest_1` is `$ORACLE_HOME/clone_sid` and the archive logs of the cloned database will be created in this location. If you have deleted the `log_archive_dest_1` parameter, the archive log location is determined by Oracle. You can define a new location for archive log by editing `log_archive_dest_1` but ensure that the file system or disk group should be existing and made available on the host.

- c. Click **Reset** to get the default database parameter settings.

10. In the PostOps page, **Recover database** and **Until Cancel** are selected by default to perform recovery of the cloned database.

SnapCenter performs recovery by mounting the latest log backup that have the unbroken sequence of archive logs after the data backup that was selected for cloning. The log and data backup should be on primary storage to perform the clone on primary storage and log and data backup should be on secondary storage to perform the clone on secondary storage.


The **Recover database** and **Until Cancel** options are not selected if SnapCenter fails to find the appropriate log backups. You can provide the external archive log location if log backup is not available in **Specify external archive log locations**. You can specify multiple log locations.



If you want to clone a source database that is configured to support flash recovery area (FRA) and Oracle Managed Files (OMF), the log destination for recovery must also adhere to OMF directory structure.

The PostOps page is not displayed if the source database is a Data Guard standby or an Active Data Guard standby database. For Data Guard standby or an Active Data Guard standby database, SnapCenter does not provide an option to select the type of recovery in the SnapCenter GUI but the database is recovered using Until Cancel recovery type without applying any logs.

Field name	Description
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. The cloned database is recovered till the missing or corrupt log file.
Date and time	<p>SnapCenter recovers the database up to a specified date and time. The accepted format is mm/dd/yyyy hh:mm:ss.</p> <div> <p>The time can be specified in 24 hour format.</p> </div>
Until SCN (System Change Number)	SnapCenter recovers the database up to a specified system change number (SCN).
Specify external archive log locations	<p>If the database is running in ARCHIVELOG mode, SnapCenter identifies and mounts optimal number of log backups based on the specified SCN or the selected date and time.</p> <p>You can also specify the external archive log location.</p> <div> <p>SnapCenter will not automatically identify and mount the log backups if you have selected Until Cancel.</p> </div>

Field name	Description
Create new DBID	<p>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.</p> <p>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.</p>
Create tempfile for temporary tablespace	<p>Select the check box if you want to create a tempfile for the default temporary tablespace of the cloned database.</p> <p>If the check box is not selected, the database clone will be created without the tempfile.</p>
Enter sql entries to apply when clone is created	Add the sql entries that you want to apply when the clone is created.
Enter scripts to run after clone operation	<p>Specify the path and the arguments of the postscript that you want to run after the clone operation.</p> <p>You should store the postscript either in <code>/var/opt/snapcenter/spl/scripts</code> or in any folder inside this path. By default, the <code>/var/opt/snapcenter/spl/scripts</code> path is populated.</p> <p>If you have placed the script in any folder inside this path, you need to provide the complete path up to the folder where the script is placed.</p> <div>  <p>If the clone operation fails, postscripts will not be executed and cleanup activities will be triggered directly.</p> </div>

- 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. If you want to attach the report of the clone operation performed, select **Attach Job Report**.



For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command `Set-SmSmtServer`.

- Review the summary, and then click **Finish**.



While performing recovery as part of clone create operation, even if recovery fails, the clone is created with a warning. You can perform manual recovery on this clone to bring the clone database to consistent state.

13. Monitor the operation progress by clicking **Monitor > Jobs**.

Result

After cloning the database you can refresh the resources page to list the cloned database as one of the resource available for backup. The cloned database can be protected like any other database using the standard backup workflow or can be included in a resource group (either newly created or existing). The cloned database can be further cloned (clone of clones).

After cloning, you should never rename the cloned database.



If you have not performed recovery while cloning, the backing up of the cloned database might fail due to improper recovery and you might have to perform manual recovery. The log backup can also fail if default location which was populated for archive logs is on a non-NetApp storage or if the storage system is not configured with SnapCenter.

In AIX setup, you can use the `lkdev` command to lock and the `rendev` command to rename the disks on which the cloned database resided.

Locking or renaming of devices will not affect the clone deletion operation. For AIX LVM layouts built on SAN devices, renaming of devices will not be supported for the cloned SAN devices.

Find more information

- [Restore or cloning fails with ORA-00308 error message](#)
- [Failed to recover a cloned database](#)
- [Customizable parameters for backup, restore and clone operations on AIX systems](#)


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) either from the database details view or from the resource group details view.

The database topology page is displayed.

4. From the Manage Copies view, select the backups either from Local copies (primary), Mirror copies (secondary), or Vault copies (secondary).
5. Select the backup from the table, and then click .

6. In the Name page, perform the following actions:


- a. Select **PDB Clone**.
- b. Specify the PDB that you want to clone.



You can clone only one PDB at a time.

- c. Specify the name of the clone PDB.

7. In the Locations page, perform the following actions:

For this field...	Do this...
Clone host	<p>By default, the source database host is populated.</p> <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>
Target CDB	<p>Select the CDB where you want to include the cloned PDB.</p> <p>You should ensure that the target CDB is running.</p>
Database State	<p>Select the Open the cloned PDB in READ-WRITE mode checkbox if you want to open the PDB in READ-WRITE mode.</p>
Datafile locations	<p>By default, the datafile location is populated.</p> <p>The SnapCenter default naming convention for SAN or NFS file systems is FileSystemNameofsourcedatabase_SCJOBID.</p> <p>The SnapCenter default naming convention for ASM disk groups is SC_HASHCODEofDISKGROUP_SCJOBID. The HASHCODEofDISKGROUP is an automatically generated number (2 to 10 digits) that is unique for each ASM disk group.</p> <div><p>If you are customizing the ASM disk group name, ensure that the name length adheres to the maximum length supported by Oracle.</p></div> <p>If you want to specify a different path, you must enter the datafile mount points or ASM disk group names for clone database.</p>

The Oracle home, user name, and group details are automatically populated from the source database. You can change the values based on the Oracle environment of the host where the clone will be created.

8. In the PreOps page, perform the following steps:

- a. Enter the path and the arguments of the prescript that you want to run before the clone operation.

You should store the prescript either in `/var/opt/snapcenter/spl/scripts` or in any folder inside this path. By default, the `/var/opt/snapcenter/spl/scripts` path is populated. If you have placed the script in any folder inside this path, you need to provide the complete path up to the folder where the script is placed.

SnapCenter allows you to use the predefined environment variables when you execute the prescript and postscript. [Learn more](#)

- b. In the Auxiliary CDB clone database parameter settings section, modify the values of prepopulated database parameters that are used to initialize the database.


9. Click **Reset** to get the default database parameter settings.


10. In the PostOps page, **Until Cancel** is selected by default to perform recovery of the cloned database.

The **Until Cancel** option is not selected if SnapCenter fails to find the appropriate log backups. You can provide the external archive log location if log backup is not available in **Specify external archive log locations**. You can specify multiple log locations.



If you want to clone a source database that is configured to support flash recovery area (FRA) and Oracle Managed Files (OMF), the log destination for recovery must also adhere to OMF directory structure.

Field name	Description
Until Cancel	<p>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.</p> <p>The log and data backup should be on primary storage to perform the clone on primary storage and log and data backup should be on secondary storage to perform the clone on secondary storage. The cloned database is recovered till the missing or corrupt log file.</p>
Date and time	<p>SnapCenter recovers the database up to a specified date and time.</p> <div> The time can be specified in 24 hour format.</div>
Until SCN (System Change Number)	SnapCenter recovers the database up to a specified system change number (SCN).
Specify external archive log locations	Specify the external archive log location.

Field name	Description
Create new DBID	<p>By default Create new DBID check box is not selected for the auxiliary clone database.</p> <p>Select the check box if you want to generate a unique number (DBID) for the auxiliary cloned database differentiating it from the source database.</p>
Create tempfile for temporary tablespace	<p>Select the check box if you want to create a tempfile for the default temporary tablespace of the cloned database.</p> <p>If the check box is not selected, the database clone will be created without the tempfile.</p>
Enter sql entries to apply when clone is created	Add the sql entries that you want to apply when the clone is created.
Enter scripts to run after clone operation	<p>Specify the path and the arguments of the postscript that you want to run after the clone operation.</p> <p>You should store the postscript either in <code>/var/opt/snapcenter/spl/scripts</code> or in any folder inside this path.</p> <p>By default, the <code>/var/opt/snapcenter/spl/scripts</code> path is populated. If you have placed the script in any folder inside this path, you need to provide the complete path up to the folder where the script is placed.</p> <div>  <p>If the clone operation fails, postscripts will not be executed and cleanup activities will be triggered directly.</p> </div>

- 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. If you want to attach the report of the clone operation performed, select **Attach Job Report**.



For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command `Set-SmSmtServer`.

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

After you finish

If you want to create a backup of the cloned PDB, you should backup the target CDB where the PDB is cloned because backing up only the cloned PDB is not possible. You should create a secondary relationship for the target CDB if you want to create the backup with secondary relationship.

In a RAC setup the storage for cloned PDB is attached only to the node where the PDB clone was performed. The PDBs on the other nodes of the RAC are in MOUNT state. If you want the cloned PDB to be accessible from the other nodes, you should manually attach the storage to the other nodes.

Find more information

- [Restore or cloning fails with ORA-00308 error message](#)
- [Customizable parameters for backup, restore and clone operations on AIX systems](#)

Clone Oracle database backups using UNIX commands

The clone workflow includes planning, performing the clone operation, and monitoring the operation.

About this task

You should execute the following commands to create the Oracle database clone specification file and initiate the clone operation.

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

Steps

1. Create an Oracle database clone specification from a specified backup: *New-SmOracleCloneSpecification*



If secondary data protection policy is unified mirror-vault, then specify only `-IncludeSecondaryDetails`. You do not have to specify `-SecondaryStorageType`.

This command automatically creates an Oracle database clone specification file for the specified source database and its backup. You must also provide a clone database SID so that the specification file created has the automatically generated values for the clone database which you will be creating.



The clone specification file is created at `/var/opt/snapcenter/sco/clone_specs`.

2. Initiate a clone operation from a clone resource group or an existing backup: *New-SmClone*

This command initiates a clone operation. You must also provide an Oracle clone specification file path for the clone operation. You can also specify the recovery options, host where the clone operation to be performed, prescripts, postscripts, and other details.

By default, the archive log destination file for the clone database is automatically populated at `$ORACLE_HOME/CLONE_SIDs`.

Split an Oracle Database Clone

You can use SnapCenter to split a cloned resource from the parent resource. The clone

that is split becomes independent of the parent resource.

About this task


- You cannot perform the clone split operation on an intermediate clone.

For example, after you create clone1 from a database backup, you can create a backup of clone1, and then clone this backup (clone2). After you create clone2, clone1 is an intermediate clone, and you cannot perform the clone split operation on clone1. However, you can perform the clone split operation on clone2.

After splitting clone2, you can perform the clone split operation on clone1 because clone1 is no longer the intermediate clone.

- When you split a clone, the backup copies of the clone are deleted.
- For information about clone split operation limitations, see the [ONTAP 9 Logical Storage Management Guide](#).
- Ensure that the volume or aggregate on the storage system is online.

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 **Database** from the **View** list.
3. Select the cloned resource, (for example, the database or LUN) and then click .
4. Review the estimated size of the clone that is to be split and the required space available on the aggregate, and then click **Start**.
5. Monitor the operation progress by clicking **Monitor > Jobs**.

The clone split operation stops responding if the SMCore service restarts and the databases on which the clone split operation was performed are listed as clones in the Resources page. You should run the *Stop-SmJob* cmdlet to stop the clone split operation, and then retry the clone split operation.

If you want a longer poll time or shorter poll time to check whether the clone is split or not, you can change the value of CloneSplitStatusCheckPollTime parameter in SMCoreServiceHost.exe.config file to set the time interval for SMCore to poll for the status of the clone split operation. The value is in milliseconds and the default value is 5 minutes.

For example,

```
<add key="CloneSplitStatusCheckPollTime" value="300000" />
```



The clone split start operation fails if backup, restore, or another clone split is in progress. You should restart the clone split operation only after the running operations are complete.

Split clone of a pluggable database

You can use SnapCenter to split a cloned pluggable database (PDB).


About this task

If you created a backup of the target CDB where the PDB is cloned, when you split the PDB clone, the cloned PDB is also removed from all the backups of the target CDB containing the cloned PDB.



The PDB clones are not displayed in the inventory or resources view.

Steps







1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. Select the source container database (CDB) from the resource or resource group view.
3. From the Manage Copies view, select **Clones** either from the primary or secondary (mirrored or replicated) storage systems.
4. Select the PDB clone (targetCDB:PDBClone) and then click .
5. Review the estimated size of the clone that is to be split and the required space available on the aggregate, and then click **Start**.
6. Monitor the operation progress by clicking **Monitor > Jobs**.

Monitor Oracle database clone operations


You can monitor the progress of SnapCenter clone operations by using the Jobs page. You might want to check the progress of an operation to determine when it is complete or if there is an issue.

About this task

The following icons appear on the Jobs page, and indicate the state of the operation:

-  In progress
-  Completed successfully
-  Failed
-  Completed with warnings or could not start due to warnings
-  Queued
-  Canceled

Steps

1. In the left navigation pane, click **Monitor**.
2. In the Monitor page, click **Jobs**.
3. In the Jobs page, perform the following steps:
 - a. Click  to filter the list so that only clone operations are listed.
 - b. Specify the start and end dates.
 - c. From the **Type** drop-down list, select **Clone**.
 - d. From the **Status** drop-down list, select the clone status.
 - e. Click **Apply** to view the operations that are completed successfully.

4. Select the clone job, and then click **Details** to view the job details.
5. In the Job Details page, click **View logs**.

Refresh a clone

You can refresh the clone by running the *Refresh-SmClone* command. This command creates a backup of the database, deletes the existing clone, and creates a clone with the same name.



You cannot refresh a PDB clone.

What you will need

- Create an online full backup or an offline data backup policy with no scheduled backups enabled.
- Configure the email notification in the policy for backup failures only.
- Define the retention count for the on-demand backups appropriately to ensure that there are no unwanted backups.
- Ensure that only an online full backup or an offline data backup policy is associated with resource group which is identified for refresh clone operation.
- Create a resource group with only one database.
- If a cron job is created for the clone refresh command, ensure that the SnapCenter schedules and the cron schedules are not overlapping for the database resource group.

For a cron job created for the clone refresh command, ensure that you run *Open-SmConnection* after every 24hrs.

- Ensure that the clone SID is unique for a host.

If multiple refresh clone operations use the same clone specification file or use the clone specification file with same clone SID, existing clone with the SID on the host will be deleted and then the clone will be created.

- Ensure that the backup policy is enabled with secondary protection and the clone specification file is created with “-IncludeSecondaryDetails” to create the clones using secondary backups.
 - If the primary clone specification file is specified but the policy has secondary update option selected, the backup will be created, and update will get transferred to secondary. However, the clone will be created from the primary backup.
 - If the primary clone specification file is specified and the policy does not have secondary update option selected, the backup will be created on primary and clone will be created from primary.

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user: *Open-SmConnection*
2. Create an Oracle database clone specification from a specified backup: *New-SmOracleCloneSpecification*



If secondary data protection policy is unified mirror-vault, then specify only -IncludeSecondaryDetails. You do not have to specify -SecondaryStorageType.

This command automatically creates an Oracle database clone specification file for the specified source

database and its backup. You must also provide a clone database SID so that the specification file created has the automatically generated values for the clone database which you will be creating.



The clone specification file is created at `/var/opt/snapcenter/sco/clone_specs`.

3. Run *Refresh-SmClone*.

If the operation fails with the "PL-SCO-20032: canExecute operation failed with error: PL-SCO-30031: Redo log file +SC_2959770772_clmdb/clmdb/redolog/redo01_01.log exists" error messages, specify a higher value for *-WaitToTriggerClone*.

For detailed information on UNIX commands, see the [SnapCenter Software Command Reference Guide](#).

Delete clone of a pluggable database


You can delete the clone of a pluggable database (PDB) if you no longer require.

If you created a backup of the target CDB where the PDB is cloned, when you delete the PDB clone, the cloned PDB is also removed from the backup of the target CDB.



The PDB clones are not displayed in the inventory or resources view.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. Select the source container database (CDB) from the resource or resource group view.
3. From the Manage Copies view, select **Clones** either from the primary or secondary (mirrored or replicated) storage systems.
4. Select the PDB clone (targetCDB:PDBClone) and then click .
5. Click **OK**.

Manage application volumes

Add application volumes

SnapCenter supports backing up and cloning of application volumes of Oracle database. You should manually add the application volumes. Auto discovery of application volumes is not supported.

Steps

1. In the left navigation pane, click **Resources**, and then select the Oracle Database plug-in from the list.
2. Click **Add Application Volume**.
3. In the Name page, perform the following actions:
 - In the Name field, enter the name of the application volume.
 - In the Host Name field, enter the name of the host.
4. In the Storage Footprint page, enter the storage system name, select one or volumes, and specify the

associated LUNs or Qtrees.



You can add multiple storage systems.

5. Review the summary, and then click **Finish**.
6. In the Resources page, select **Application Volume** from the **View** list to view all the application volumes that you have added.

Modify application volume

You can modify all the values that you specified while adding the application volume, if no backups are created. If the backup is created, you can only modify the storage system details.

Steps


1. In the left navigation pane, click **Resources**, and then select the Oracle Database plug-in from the list.
2. In the Resources page, select **Application Volume** from the **View** list.
3.  Click  to modify the values.

Delete application volume

When you delete an application volume, if there any backups associated with the application volume, the application volume will be put into maintenance mode and no new backups will be created and no earlier backups will be retained. If there are no backups associated, all the metadata will be deleted.

If required, SnapCenter allows you to undo the delete operation.

Steps

1. In the left navigation pane, click **Resources**, and then select the Oracle Database plug-in from the list.
2. In the Resources page, select **Application Volume** from the **View** list.
3. Click  to modify the values.

Backup application volumes


Back up application volume

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

About this task

By default, consistency group (CG) backups are created. If you want to create volume based backups, you should set the value of **EnableOracleNdvVolumeBasedBackup** to true in the *web.config* file.

Steps

1. In the left navigation pane, click **Resources**, and then select the Oracle Database plug-in from the list.
2. In the Resources page, select **Application Volume** from the **View** list.
3. Click , and then select the host name and the database type to filter the resources.

You can then click  to close the filter pane.

4. Select the application volume that you want to back up.

The Application volume-Protect page is displayed.

5. In the Resource page, perform the following actions:

For this field...	Do this...
Use custom name format for Snapshot copy	Select this check box, and then enter a custom name format that you want to use for the Snapshot copy name. For example, customtext__policy_hostname or resource_hostname. By default, a timestamp is appended to the Snapshot copy name.
Exclude archive log destinations from backup	Specify the destinations of the archive log files that you do not want to back up.


6. In the Policies page, perform the following steps:

- a. Select one or more policies from the drop-down list.



You can also create a policy by clicking .

In the Configure schedules for selected policies section, the selected policies are listed.

- a. Click  in the Configure Schedules column for the policy for which you want to configure a schedule.
- b. In the Add schedules for policy *policy_name* window, configure the schedule, and then click **OK**.

policy_name is the name of the policy that you have selected.

The configured schedules are listed in the Applied Schedules column.

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. If you want to attach the report of the backup operation performed on the resource, and then select **Attach Job Report**.



For email notification, you must have specified the SMTP server details using the either the GUI or the PowerShell command Set-SmSmtServer.

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

The application volume topology page is displayed.

9. Click **Back up Now**.

10. In the Backup page, perform the following steps:
 - a. If you have applied multiple policies to the resource, from the **Policy** drop-down list, select the policy that you want to use for backup.
 - b. Click **Backup**.
11. Monitor the operation progress by clicking **Monitor > Jobs**.

Back up the application volumes resource group

You can back up the resource group containing only application volumes or a mix of application volumes and database. A backup operation on the resource group is performed on all resources defined in the resource group.



If the resource group has multiple application volumes, all the application volumes should either have SnapMirror or SnapVault replication policy.

About this task

By default, consistency group (CG) backups are created. If you want to create volume based backups, you should set the value of **EnableOracleNdvVolumeBasedBackup** to true in the *web.config* file.

Steps

1. In the left navigation pane, click **Resources**, and then select the Oracle Database plug-in from the list.
2. In the Resources page, select **Resource Group** from the **View** list.

You can search the resource group either by entering the resource group name in the search box, or by clicking , and then selecting the tag. You can then click  to close the filter pane.
3. In the Resource Groups page, select the resource group that you want to back up, and then click **Back up Now**.
4. In the Backup page, perform the following steps:
 - a. If you have associated multiple policies 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. Click **Backup**.
5. Monitor the operation progress by clicking **Monitor > Jobs**.



Verification operation will be performed only for the databases and not for the application volumes.


Clone application volume backup

You can use SnapCenter to clone the application volume backups.

Steps

1. In the left navigation pane, click **Resources**, and then select the Oracle Database plug-in from the list.
2. In the Resources page, select **Application Volume** from the **View** list.
3. Select the application volume either from the application volume details view or from the resource group details view.

The application volume topology page is displayed.

4. From the Manage Copies view, select the backups either from Local copies (primary), Mirror copies (secondary), or Vault copies (secondary).
5. Select the backup from the table, and then click .
6. In the Location page, perform the following actions:

For this field...	Do this...
Plug-in host	Select the host where you want to create the clone.
Target Resource Name	Specify the resource name.

7. In the Scripts page, specify the names of the scripts to be executed before cloning, commands to mount a file system, and names of the scripts to be executed after cloning.
8. 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. If you want to attach the report of the clone operation performed, select **Attach Job Report**.




For email notification, you must have specified the SMTP server details using either the GUI or the PowerShell command Set-SmSmtServer.

9. Review the summary, and then click **Finish**.

Split an application volume clone

You can use SnapCenter to split a cloned resource from the parent resource. The clone that is split becomes independent of the parent resource.

Steps

1. In the left navigation pane, click **Resources**, and then select the Oracle Database plug-in from the list.
2. In the Resources page, select **Application Volume** from the **View** list.
3. Select the cloned resource and click .
4. Review the estimated size of the clone that is to be split and the required space available on the aggregate, and then click **Start**.
5. Monitor the operation progress by clicking **Monitor > Jobs**.

Delete an application volume clone


You can delete clones if you find them no longer necessary. You cannot delete clones that acts like source for

other clones.

Steps

1. In the left navigation pane, click **Resources**, and then select the Oracle Database plug-in from the list.
2. In the Resources page, select **Application Volume** from the **View** list.
3. Select the resource or resource group from the list.

The resource or the resource group topology page is displayed.

4. From the Manage Copies view, select **Clones** either from the primary or secondary (mirrored or replicated) storage systems.
5. Select the clone, and then click .
6. In the Delete Clone page, perform the following actions:
 - a. In the **Pre clone delete** field, enter the names of the scripts to be executed before deleting the clone.
 - b. In the **Unmount** field, enter the commands to unmount the clone before deleting the clone.
7. Click **OK**.

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.