



Back up Oracle databases

SnapCenter Software 4.5

NetApp
July 18, 2024

Table of Contents


- Back up Oracle databases 1
 - Backup workflow 1
 - Define a backup strategy for Oracle databases 2
 - Determine whether Oracle databases are available for backup 8
 - Create backup policies for Oracle databases 10
 - Create resource groups and attach policies for Oracle databases 13
 - Requirements for backing up an Oracle database 16
- Back up Oracle resources 16
 - Back up Oracle database resource groups 19
 - Back up Oracle databases using UNIX commands 20
- Monitor Oracle database backup operations 21
- Cancel backup operations of Oracle databases 22
- View Oracle database backups and clones in the Topology page 23

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 (ASMFDD)



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>

For Oracle versions 11g to 12cR1	For Oracle versions 12cR2 to 18c
<p>Standalone: The standalone databases are discovered only on the basis of <code>/etc/oratab</code> entries.</p> <p>You should have the database entries in the <code>/etc/oratab</code> file.</p>	<p>Standalone: The standalone databases are discovered based on the entries in the <code>/etc/oratab</code> file and the output of the <code>srvctl config</code> command.</p>
<p>ASM: The ASM instance entry should be available in the <code>/etc/oratab</code> file.</p>	<p>ASM: The ASM instance entry need not be in the <code>/etc/oratab</code> file.</p>
<p>RAC One Node: The RAC One Node databases are discovered only on the basis of <code>/etc/oratab</code> 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 <code>/etc/oratab</code> 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 <code>/etc/oratab</code> 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 style="border-left: 1px solid #ccc; padding-left: 10px; margin-top: 10px;"> <p> If you have relocated a database from one node to another node and if the <code>oratab</code> entry in the earlier node is not deleted, you should manually delete the <code>oratab</code> entry to avoid the same database being displayed twice.</p> </div>	<p>RAC One Node: The RAC One Node databases are discovered using the <code>srvctl config</code> 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 `cdbrac` has three instances: `cdbrac1` on `node1`, `cdbrac2` on `node2`, and `cdbrac3` 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:

```
dts1_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.

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.

- 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.

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 style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;"> <p> 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.</p> </div> <div style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;"> <p> 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.</p> </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 style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;"> <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.

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. On the Resources page, click **New Resource Group**.
3. On the Name page, perform the following actions:

For this field...	Do this...
Name	Enter a name for the resource group.  The resource group name should not exceed 250 characters.
Tags	Enter one or more labels that will help you later search for the resource group. 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.
Use custom name format for Snapshot copy	Select this check box, and enter a custom name format that you want to use for the Snapshot copy name. For example, customtext_resource group_policy_hostname or resource group_hostname. By default, a timestamp is appended to the Snapshot copy name.

For this field...	Do this...
Exclude archive log destinations from backup	Specify the destinations of the archive log files that you do not want to back up.

4. On 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. On 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. On 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 seconds 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. On 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. On 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. 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 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	Select Run scheduled verification , and then select the schedule type from the drop-down list. <div style="border-left: 1px solid #ccc; padding-left: 10px; margin-top: 10px;">  <p>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.</p> </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. On 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. On 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. On 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. On 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. On 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. On 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	<ol style="list-style-type: none"> In the left navigation pane, click Monitor > Jobs. Select the operation and click Cancel Job.
Activity pane	<ol style="list-style-type: none"> After initiating the backup job, click  on the Activity pane to view the five most recent operations. Select the operation. 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](#).

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.