



Restore and recover Oracle databases

SnapCenter Software 5.0

NetApp
April 04, 2024

Table of Contents


- Restore and recover Oracle databases 1
 - Restore workflow 1
 - Define a restore and recovery strategy for Oracle databases 1
 - Predefined environment variables for restore specific prescript and postscript. 6
 - Requirements for restoring an Oracle database 7
 - Restore and recover Oracle database 8
 - Restore and recover tablespaces using point-in-time recovery 12
 - Restore and recover pluggable database using point-in-time recovery 14
 - Restore and recover Oracle databases using UNIX commands 17
 - Monitor Oracle database restore operations 17
 - Cancel Oracle database restore operations. 18

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 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 Snapshots 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.

Before you begin

If you have installed the plug-in as a non-root user, you should manually assign the execute permissions to the prescript and postscript directories.

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 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 `/var/opt/snapcenter/sco/lock` 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.

- For SnapLock enabled policies, for ONTAP 9.12.1 and below version, if you specify a Snapshot locking period, the clones created from the tamper proof Snapshots as part of restore will inherit the SnapLock expiry time. Storage admin should manually cleanup the clones post the SnapLock expiry 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 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.

Before you begin

- The backups that are required to perform PITR of tablespaces should be cataloged and mounted.
- If you have installed the plug-in as a non-root user, you should manually assign the execute permissions to the prescript and postscript directories.

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.

Before you begin

- 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.

- If you have installed the plug-in as a non-root user, you should manually assign the execute permissions to the prescript and postscript directories.

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	<div><div><div>i. Select Pluggable databases (PDBs).</div><div>ii. Specify the PDB you want to restore.</div></div><div><div></div><div>You cannot perform PITR on PDB\$SEED database.</div></div></div>

Restore tablespaces in a PDB	<p>i. Select Pluggable database (PDB) tablespaces.</p> <p>ii. Specify the PDB.</p> <p>iii. Specify either a single tablespace or multiple tablespaces you want to restore.</p> <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 Command Reference Guide](#).

- For SnapMirror Business Continuity (SM-BC) restore operation, you must select the backup from the primary location.

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.

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.

From the...	Action
Activity pane	<ul 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.

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.