



Clone Oracle database

SnapCenter Software 4.8

NetApp
February 12, 2024

Table of Contents

- Clone Oracle database 1
 - Clone workflow 1
 - Define a clone strategy for Oracle databases 1
 - Predefined environment variables for clone specific prescript and postscript 2
 - Requirements for cloning an Oracle database 4
 - Clone an Oracle database backup 6
 - Clone a pluggable database 13
 - Clone Oracle database backups using UNIX commands 17
 - Split an Oracle Database Clone 18
 - Split clone of a pluggable database 19
 - Monitor Oracle database clone operations 20
 - Refresh a clone 20
 - Delete clone of a pluggable database 21

Clone Oracle database


Clone workflow

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

You might clone databases for the following reasons:

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

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



Define a clone strategy for Oracle databases

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

Types of backups supported for cloning

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

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



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



Cloning of archive log backups is not supported.

Types of cloning supported for Oracle databases

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

The SnapCenter Server uses NetApp FlexClone technology to clone backups.

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



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

Clone naming conventions for Oracle databases

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

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

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

Limitations of cloning Oracle databases

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

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

Predefined environment variables for clone specific prescript and postscript

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

Supported predefined environment variables for cloning a database

- **SC_ORIGINAL_SID** specifies the SID of the source database.

This parameter will be populated for application volumes.

Example: NFSB32

- **SC_ORIGINAL_HOST** specifies the name of the source host.

This parameter will be populated for application volumes.

Example: asmrac1.gdl.englab.netapp.com

- **SC_ORACLE_HOME** specifies the path of the target database's Oracle home directory.

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

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

This parameter will be populated for application volumes.

Examples:

- If the database is not running in ARCHIVELOG mode: DATA@RG2_scspr2417819002_07-20-2021_12.16.48.9267_0|LOG@RG2_scspr2417819002_07-20-2021_12.16.48.9267_1
- If the database is running in ARCHIVELOG mode: DATA@RG2_scspr2417819002_07-20-2021_12.16.48.9267_0|LOG:RG2_scspr2417819002_07-20-2021_12.16.48.9267_1,RG2_scspr2417819002_07-21-2021_12.16.48.9267_1,RG2_scspr2417819002_07-22-2021_12.16.48.9267_1

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

Example: AV1|AV2

- **SC_ORIGINAL_OS_USER** specifies the operating system owner of the source database.

Example: oracle

- **SC_ORIGINAL_OS_GROUP** specifies the operating system group of the source database.

Example: oinstall

- **SC_TARGET_SID** specifies the SID of the cloned database.

For PDB clone workflow, the value of this parameter will not be predefined.

This parameter will be populated for application volumes.

Example: clonedb

- **SC_TARGET_HOST** specifies the name of the host where the database will be cloned.

This parameter will be populated for application volumes.

Example: asmrac1.gdl.englab.netapp.com

- **SC_TARGET_OS_USER** specifies the operating system owner of the cloned database.

For PDB clone workflow, the value of this parameter will not be predefined.

Example: oracle

- **SC_TARGET_OS_GROUP** specifies the operating system group of the cloned database.

For PDB clone workflow, the value of this parameter will not be predefined.

Example: oinstall

- **SC_TARGET_DB_PORT** specifies the database port of the cloned database.

For PDB clone workflow, the value of this parameter will not be predefined.

Example: 1521

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

Requirements for cloning an Oracle database

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

- You should have created a backup of the database using SnapCenter.

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

- If you want to customize the control file or redo log file paths, you should have preprovisioned the required file system or Automatic Storage Management (ASM) disk group.

By default, redo log and control files of the cloned database are created on the ASM disk group or the file system provisioned by SnapCenter for the data files of the clone database.

- If you are using ASM over NFS, you should add `/var/opt/snapcenter/scu/clones/*/*` to the existing path defined in the `asm_diskstring` parameter.
- In the `asm_diskstring` parameter, you should configure `AFD:*` if you are using ASMFD or configure `ORCL:*` if you are using ASMLIB.

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

- If you are creating the clone on an alternate host, the alternate host should meet the following requirements:
 - SnapCenter Plug-in for Oracle Database should be installed on the alternate host.
 - The clone host should be able to discover LUNs from primary or secondary storage.
 - If you are cloning from primary storage or secondary (Vault or Mirror) storage to an alternate host, then make sure that an iSCSI session is either established between the secondary storage and the alternate host, or zoned properly for FC.
 - If you are cloning from Vault or Mirror storage to the same host, then make sure that an iSCSI session is either established between the Vault or Mirror storage and the host, or zoned properly for FC.
 - If you are cloning in a virtualized environment, ensure that an iSCSI session is either established

between the primary or secondary storage and the ESX server hosting the alternate host, or zoned properly for FC.

For information, refer to [host utilities documentation](#).

- If the source database is an ASM database:

- The ASM instance should be up and running on the host where the clone will be performed.
- The ASM disk group should be provisioned prior to the clone operation if you want to place archive log files of the cloned database in a dedicated ASM disk group.
- The name of the data disk group can be configured, but ensure that the name is not used by any other ASM disk group on the host where the clone will be performed.

Data files residing on the ASM disk group are provisioned as part of the SnapCenter clone workflow.

- For NVMe, NVMe util should be installed

- The protection type for the data LUN and the log LUN, such as mirror, vault, or mirror-vault, should be the same to discover secondary locators during cloning to an alternate host using log backups.
- You should set the value of `exclude_seed_cdb_view` to `FALSE` in the source database parameter file to retrieve seed PDB related information for cloning a backup of `12_c_database`.

The seed PDB is a system-supplied template that the CDB can use to create PDBs. The seed PDB is named `PDB$SEED`. For information about `PDB$SEED`, see the Oracle Doc ID 1940806.1.



You should set the value before backing up `12_c_database`.

- SnapCenter supports backup of file systems that are managed by the autofs subsystem. If you are cloning the database, ensure that data mount points are not under the root of the autofs mount point because the root user of the plug-in host does not have permission to create directories under the root of the autofs mount point.

If control and redo log files are under data mount point, you should modify the control file path, and then redo log file path accordingly.



You can manually register the new cloned mount points with the autofs subsystem. The new cloned mount points will not be registered automatically.

- If you have a TDE (auto login) and want to clone the database on the same or alternate host, you should copy wallet (key files) under `/etc/ORACLE/WALLET/$ORACLE_SID` from the source database to the cloned database.
- You should set the value of `use_lvm2 = 0` in `/etc/lvm/lvm.conf` and stop the `lvm2-lvmetad` service to successfully perform cloning in storage area network (SAN) environments on Oracle Linux 7 or later or Red Hat Enterprise Linux (RHEL) 7 or later.
- You should install the 13366202 Oracle patch if you are using Oracle database 11.2.0.3 or later and the database ID for the auxiliary instance is changed using an NID script.
- You should ensure that the aggregates hosting the volumes should be in the assigned aggregates list of the storage virtual machine (SVM).
- For NVMe, if any target port has to be excluded from connecting, you should add the target node name and port name in the `/var/opt/snapcenter/scu/etc/nvme.conf` file.

If the file does not exist, you should create the file as shown in the example below:

```
blacklist {
nn-0x<target_node_name_1>:pn-0x<target_port_name_1>
nn-0x<target_node_name_2>:pn-0x<target_port_name_2>
}
```

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

Clone an Oracle database backup

You can use SnapCenter to clone an Oracle database using the backup of the database.

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

The cloning operation creates a copy of the database data files, and creates new online redo log files and control files. The database can be optionally recovered to a specified time, based on the specified recovery options.



Cloning fails if you try to clone a backup that was created on a Linux host to an AIX host or vice-versa.

SnapCenter creates a stand-alone database when cloned from an Oracle RAC database backup. SnapCenter supports creating clone from the backup of a Data Guard standby and Active Data Guard standby databases.

During cloning, SnapCenter mounts the optimal number of log backups based on SCN or dat and time for recovery operations. After recovery, the log backup is unmounted. All such clones are mounted under */var/opt/snapcenter/scu/clones/*. If you are using ASM over NFS, you should add */var/opt/snapcenter/scu/clones/*/** to the existing path defined in the *asm_diskstring* parameter.

While cloning a backup of an ASM database in a SAN environment, udev rules for the cloned host devices are created at */etc/udev/rules.d/999-scu-netapp.rules*. These udev rules associated with the cloned host devices are deleted when you delete the clone.



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


Steps


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

The database topology page is displayed.

4. From the Manage Copies view, select the backups either from Local copies (primary), Mirror copies

(secondary), or Vault copies (secondary).

5. Select the Data backup from the table, and then click .
6. In the Name page, perform one of the following actions:

If you want to...	Steps...
Clone a database (CDB or non CDB)	<p>a. Specify the SID of the clone.</p> <p>The clone SID is not available by default, and the maximum length of the SID is 8 characters.</p> <div style="border: 1px solid #ccc; padding: 5px;"> You should ensure that no database with the same SID exists on the host where the clone will be created.</div>
Clone a pluggable database (PDB)	<p>a. Select PDB Clone.</p> <p>b. Specify the PDB that you want to clone.</p> <p>c. Specify the name of cloned PDB. For the detailed steps to clone a PDB, see Clone a pluggable database.</p>

When you select a mirrored or vault data:



- if there are no log backup at mirror or vault, nothing is selected and the locators are empty.
- if log backups exist in mirror or vault, the latest log backup is selected and corresponding locator is displayed.






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

7. In the Locations page, perform the following actions:

For this field...	Do this...
Clone host	<p>By default, the source database host is populated.</p> <p>If you want to create the clone on an alternate host, select the host having the same version of Oracle and OS as that of the source database host.</p>

For this field...	Do this...
Datafile locations	<p>By default, the datafile location is populated.</p> <p>The SnapCenter default naming convention for SAN or NFS file systems is FileSystemNameofsourcedatabase_CLONESID.</p> <p>The SnapCenter default naming convention for ASM disk groups is SC_HASHCODEofDISKGROUP_CLONESID. The HASHCODEofDISKGROUP is an automatically generated number (2 to 10 digits) that is unique for each ASM disk group.</p> <div style="border-left: 1px solid #ccc; padding-left: 10px; margin: 10px 0;">  <p>If you are customizing the ASM disk group name, ensure that the name length adheres to the maximum length supported by Oracle.</p> </div> <p>If you want to specify a different path, you must enter the datafile mount points or ASM disk group names for clone database. When you customize the datafile path, you must also change the control file and redo log file ASM disk group names or file system either to the same name used for data files or to an existing ASM disk groups or file system.</p>
Control files	<p>By default, the control file path is populated.</p> <p>The control files are placed in the same ASM disk group or file system as that of the data files. If you want to override the control file path, you can provide a different control file path.</p> <div style="border-left: 1px solid #ccc; padding-left: 10px; margin: 10px 0;">  <p>The file system or the ASM disk group should exist on the host.</p> </div> <p>By default, the number of control files will be same as that of the source database. You can modify the number of control files but a minimum of one control file is required to clone the database.</p> <p>You can customize the control file path to a different file system (existing) than that of the source database.</p>

For this field...	Do this...
Redo logs	<p>By default, the redo log file group, path, and their sizes are populated.</p> <p>The redo logs are placed in the same ASM disk group or file system as that of the data files of the cloned database. If you want to override the redo log file path, you can customize the redo log file path to a different file system than that of the source database..</p> <p> The new file system or the ASM disk group should exist on the host.</p> <p>By default, the number of redo log groups, redo log files, and their sizes will be same as that of the source database. You can modify the following parameters:</p> <ul style="list-style-type: none"> • Number of redo log groups <p> A minimum of two redo log groups are required to clone the database.</p> <ul style="list-style-type: none"> • Redo log files in each group and their path <p>You can customize the redo log file path to a different file system (existing) than that of the source database.</p> <p> A minimum of one redo log file is required in the redo log group to clone the database.</p> <ul style="list-style-type: none"> • Sizes of the redo log file

8. On the Credentials page, perform the following actions:

For this field...	Do this...
Credential name for sys user	<p>Select the Credential to be used for defining the sys user password of the clone database.</p> <p>If SQLNET.AUTHENTICATION_SERVICES is set to NONE in sqlnet.ora file on the target host, you should not select None as the Credential in the SnapCenter GUI.</p>

For this field...	Do this...
ASM Instance Credential name	<p>Select None if OS authentication is enabled for connecting to the ASM instance on the clone host.</p> <p>Otherwise, select the Oracle ASM credential configured with either “sys” user or an user having “sysasm” privilege applicable to the clone host.</p>

The Oracle home, user name, and group details are automatically populated from the source database. You can change the values based on the Oracle environment of the host where the clone will be created.

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

- a. Enter the path and the arguments of the prescript that you want to run before the clone operation.

You must store the prescript either in `/var/opt/snapcenter/spl/scripts` or in any folder inside this path. By default, the `/var/opt/snapcenter/spl/scripts` path is populated. If you have placed the script in any folder inside this path, you need to provide the complete path up to the folder where the script is placed.

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

- b. In the Database Parameter settings section, modify the values of prepopulated database parameters that are used to initialize the database.

You can add additional parameters by clicking  .

If you are using Oracle Standard Edition and the database is running in Archive log mode or you want restore a database from archive redo log, add the parameters and specify the path.

- LOG_ARCHIVE_DEST
- LOG_ARCHIVE_DUPLEX_DEST



Fast recovery area (FRA) is not defined in the prepopulated database parameters. You can configure FRA by adding the related parameters.



The default value of `log_archive_dest_1` is `$ORACLE_HOME/clone_sid` and the archive logs of the cloned database will be created in this location. If you have deleted the `log_archive_dest_1` parameter, the archive log location is determined by Oracle. You can define a new location for archive log by editing `log_archive_dest_1` but ensure that the file system or disk group should be existing and made available on the host.

- c. Click **Reset** to get the default database parameter settings.

10. In the PostOps page, **Recover database** and **Until Cancel** are selected by default to perform recovery of the cloned database.

SnapCenter performs recovery by mounting the latest log backup that have the unbroken sequence of archive logs after the data backup that was selected for cloning. The log and data backup should be on primary storage to perform the clone on primary storage and log and data backup should be on secondary storage to perform the clone on secondary storage.


The **Recover database** and **Until Cancel** options are not selected if SnapCenter fails to find the appropriate log backups. You can provide the external archive log location if log backup is not available in **Specify external archive log locations**. You can specify multiple log locations.



If you want to clone a source database that is configured to support flash recovery area (FRA) and Oracle Managed Files (OMF), the log destination for recovery must also adhere to OMF directory structure.

The PostOps page is not displayed if the source database is a Data Guard standby or an Active Data Guard standby database. For Data Guard standby or an Active Data Guard standby database, SnapCenter does not provide an option to select the type of recovery in the SnapCenter GUI but the database is recovered using Until Cancel recovery type without applying any logs.

Field name	Description
Until Cancel	SnapCenter performs recovery by mounting the latest log backup having the unbroken sequence of archive logs after that data backup that was selected for cloning. The cloned database is recovered till the missing or corrupt log file.
Date and time	<p>SnapCenter recovers the database up to a specified date and time. The accepted format is mm/dd/yyyy hh:mm:ss.</p> <div style="display: flex; align-items: center;"> <p>The time can be specified in 24 hour format.</p> </div>
Until SCN (System Change Number)	SnapCenter recovers the database up to a specified system change number (SCN).
Specify external archive log locations	<p>If the database is running in ARCHIVELOG mode, SnapCenter identifies and mounts optimal number of log backups based on the specified SCN or the selected date and time.</p> <p>You can also specify the external archive log location.</p> <div style="display: flex; align-items: center;"> <p>SnapCenter will not automatically identify and mount the log backups if you have selected Until Cancel.</p> </div>

Field name	Description
Create new DBID	<p>By default Create new DBID check box is selected to generate a unique number (DBID) for the cloned database differentiating it from the source database.</p> <p>Clear the check box if you want to assign the DBID of the source database to the cloned database. In this scenario, if you want to register the cloned database with the external RMAN catalog where the source database is already registered, the operation fails.</p>
Create tempfile for temporary tablespace	<p>Select the check box if you want to create a tempfile for the default temporary tablespace of the cloned database.</p> <p>If the check box is not selected, the database clone will be created without the tempfile.</p>
Enter sql entries to apply when clone is created	<p>Add the sql entries that you want to apply when the clone is created.</p>
Enter scripts to run after clone operation	<p>Specify the path and the arguments of the postscript that you want to run after the clone operation.</p> <p>You should store the postscript either in <code>/var/opt/snapcenter/spl/scripts</code> or in any folder inside this path. By default, the <code>/var/opt/snapcenter/spl/scripts</code> path is populated.</p> <p>If you have placed the script in any folder inside this path, you need to provide the complete path up to the folder where the script is placed.</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;"> <p> If the clone operation fails, postscripts will not be executed and cleanup activities will be triggered directly.</p> </div>

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

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



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

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



While performing recovery as part of clone create operation, even if recovery fails, the clone is created with a warning. You can perform manual recovery on this clone to bring the clone database to consistent state.

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

Result

After cloning the database you can refresh the resources page to list the cloned database as one of the resource available for backup. The cloned database can be protected like any other database using the standard backup workflow or can be included in a resource group (either newly created or existing). The cloned database can be further cloned (clone of clones).

After cloning, you should never rename the cloned database.



If you have not performed recovery while cloning, the backing up of the cloned database might fail due to improper recovery and you might have to perform manual recovery. The log backup can also fail if default location which was populated for archive logs is on a non-NetApp storage or if the storage system is not configured with SnapCenter.

In AIX setup, you can use the `lkdev` command to lock and the `rendev` command to rename the disks on which the cloned database resided.

Locking or renaming of devices will not affect the clone deletion operation. For AIX LVM layouts built on SAN devices, renaming of devices will not be supported for the cloned SAN devices.

Find more information

- [Restore or cloning fails with ORA-00308 error message](#)
- [Failed to recover a cloned database](#)
- [Customizable parameters for backup, restore and clone operations on AIX systems](#)

Clone a pluggable database

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


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 `postsript` directories.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select either **Database** or **Resource Group** from the **View** list.
3. Select the database of type single instance (multitenant) either from the database details view or from the resource group details view.

The database topology page is displayed.


4. From the Manage Copies view, select the backups either from Local copies (primary), Mirror copies (secondary), or Vault copies (secondary).
5. Select the backup from the table, and then click .
6. In the Name page, perform the following actions:
 - a. Select **PDB Clone**.
 - b. Specify the PDB that you want to clone.



You can clone only one PDB at a time.

- c. Specify the name of the clone PDB.
7. In the Locations page, perform the following actions:

For this field...	Do this...
Clone host	By default, the source database host is populated. If you want to create the clone on an alternate host, select the host having the same version of Oracle and OS as that of the source database host.
Target CDB	Select the CDB where you want to include the cloned PDB. You should ensure that the target CDB is running.
Database State	Select the Open the cloned PDB in READ-WRITE mode checkbox if you want to open the PDB in READ-WRITE mode.

<p>Datafile locations</p>	<p>By default, the datafile location is populated.</p> <p>The SnapCenter default naming convention for SAN or NFS file systems is FileSystemNameofsourcedatabase_SCJOBID.</p> <p>The SnapCenter default naming convention for ASM disk groups is SC_HASHCODEofDISKGROUP_SCJOBID. The HASHCODEofDISKGROUP is an automatically generated number (2 to 10 digits) that is unique for each ASM disk group.</p> <div style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;">  <p>If you are customizing the ASM disk group name, ensure that the name length adheres to the maximum length supported by Oracle.</p> </div> <p>If you want to specify a different path, you must enter the datafile mount points or ASM disk group names for clone database.</p>
---------------------------	--

The Oracle home, user name, and group details are automatically populated from the source database. You can change the values based on the Oracle environment of the host where the clone will be created.

8. In the PreOps page, perform the following steps:

- a. Enter the path and the arguments of the prescript that you want to run before the clone operation.

You should store the prescript either in `/var/opt/snapcenter/spl/scripts` or in any folder inside this path. By default, the `/var/opt/snapcenter/spl/scripts` path is populated. If you have placed the script in any folder inside this path, you need to provide the complete path up to the folder where the script is placed.

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

- b. In the Auxiliary CDB clone database parameter settings section, modify the values of prepopulated database parameters that are used to initialize the database.


9. Click **Reset** to get the default database parameter settings.


10. In the PostOps page, **Until Cancel** is selected by default to perform recovery of the cloned database.

The **Until Cancel** option is not selected if SnapCenter fails to find the appropriate log backups. You can provide the external archive log location if log backup is not available in **Specify external archive log locations**. You can specify multiple log locations.



If you want to clone a source database that is configured to support flash recovery area (FRA) and Oracle Managed Files (OMF), the log destination for recovery must also adhere to OMF directory structure.

Field name	Description
Until Cancel	<p>SnapCenter performs recovery by mounting the latest log backup having the unbroken sequence of archive logs after that data backup that was selected for cloning.</p> <p>The log and data backup should be on primary storage to perform the clone on primary storage and log and data backup should be on secondary storage to perform the clone on secondary storage. The cloned database is recovered till the missing or corrupt log file.</p>
Date and time	<p>SnapCenter recovers the database up to a specified date and time.</p> <div style="display: flex; align-items: center; margin-top: 10px;">  <p>The time can be specified in 24 hour format.</p> </div>
Until SCN (System Change Number)	<p>SnapCenter recovers the database up to a specified system change number (SCN).</p>
Specify external archive log locations	<p>Specify the external archive log location.</p>
Create new DBID	<p>By default Create new DBID check box is not selected for the auxiliary clone database.</p> <p>Select the check box if you want to generate a unique number (DBID) for the auxiliary cloned database differentiating it from the source database.</p>
Create tempfile for temporary tablespace	<p>Select the check box if you want to create a tempfile for the default temporary tablespace of the cloned database.</p> <p>If the check box is not selected, the database clone will be created without the tempfile.</p>
Enter sql entries to apply when clone is created	<p>Add the sql entries that you want to apply when the clone is created.</p>

Field name	Description
Enter scripts to run after clone operation	<p>Specify the path and the arguments of the postscript that you want to run after the clone operation.</p> <p>You should store the postscript either in <code>/var/opt/snapcenter/spl/scripts</code> or in any folder inside this path.</p> <p>By default, the <code>/var/opt/snapcenter/spl/scripts</code> path is populated. If you have placed the script in any folder inside this path, you need to provide the complete path up to the folder where the script is placed.</p> <div style="display: flex; align-items: center; margin-top: 10px;">  <p>If the clone operation fails, postscripts will not be executed and cleanup activities will be triggered directly.</p> </div>

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

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



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

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

After you finish

If you want to create a backup of the cloned PDB, you should backup the target CDB where the PDB is cloned because backing up only the cloned PDB is not possible. You should create a secondary relationship for the target CDB if you want to create the backup with secondary relationship.

In a RAC setup the storage for cloned PDB is attached only to the node where the PDB clone was performed. The PDBs on the other nodes of the RAC are in MOUNT state. If you want the cloned PDB to be accessible from the other nodes, you should manually attach the storage to the other nodes.

Find more information

- [Restore or cloning fails with ORA-00308 error message](#)
- [Customizable parameters for backup, restore and clone operations on AIX systems](#)

Clone Oracle database backups using UNIX commands

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

About this task

You should execute the following commands to create the Oracle database clone specification file and initiate the clone operation.

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

Steps

1. Create an Oracle database clone specification from a specified backup: *New-SmOracleCloneSpecification*



If secondary data protection policy is unified mirror-vault, then specify only `-IncludeSecondaryDetails`. You do not have to specify `-SecondaryStorageType`.

This command automatically creates an Oracle database clone specification file for the specified source database and its backup. You must also provide a clone database SID so that the specification file created has the automatically generated values for the clone database which you will be creating.



The clone specification file is created at `/var/opt/snapcenter/sco/clone_specs`.

2. Initiate a clone operation from a clone resource group or an existing backup: *New-SmClone*

This command initiates a clone operation. You must also provide an Oracle clone specification file path for the clone operation. You can also specify the recovery options, host where the clone operation to be performed, prescripts, postscripts, and other details.

By default, the archive log destination file for the clone database is automatically populated at `$ORACLE_HOME/CLONE_SIDs`.

Split an Oracle Database Clone

You can use SnapCenter to split a cloned resource from the parent resource. The clone that is split becomes independent of the parent resource.

About this task


- You cannot perform the clone split operation on an intermediate clone.

For example, after you create clone1 from a database backup, you can create a backup of clone1, and then clone this backup (clone2). After you create clone2, clone1 is an intermediate clone, and you cannot perform the clone split operation on clone1. However, you can perform the clone split operation on clone2.

After splitting clone2, you can perform the clone split operation on clone1 because clone1 is no longer the intermediate clone.

- When you split a clone, the backup copies of the clone are deleted.
- For information about clone split operation limitations, see the [ONTAP 9 Logical Storage Management Guide](#).
- Ensure that the volume or aggregate on the storage system is online.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. In the Resources page, select **Database** from the **View** list.
3. Select the cloned resource, (for example, the database or LUN) and then click .
4. Review the estimated size of the clone that is to be split and the required space available on the aggregate, and then click **Start**.
5. Monitor the operation progress by clicking **Monitor > Jobs**.

The clone split operation stops responding if the SMCore service restarts and the databases on which the clone split operation was performed are listed as clones in the Resources page. You should run the *Stop-SmJob* cmdlet to stop the clone split operation, and then retry the clone split operation.

If you want a longer poll time or shorter poll time to check whether the clone is split or not, you can change the value of CloneSplitStatusCheckPollTime parameter in SMCoreServiceHost.exe.config file to set the time interval for SMCore to poll for the status of the clone split operation. The value is in milliseconds and the default value is 5 minutes.

For example,

```
<add key="CloneSplitStatusCheckPollTime" value="300000" />
```



The clone split start operation fails if backup, restore, or another clone split is in progress. You should restart the clone split operation only after the running operations are complete.

Split clone of a pluggable database

You can use SnapCenter to split a cloned pluggable database (PDB).


About this task

If you created a backup of the target CDB where the PDB is cloned, when you split the PDB clone, the cloned PDB is also removed from all the backups of the target CDB containing the cloned PDB.



The PDB clones are not displayed in the inventory or resources view.

Steps







1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. Select the source container database (CDB) from the resource or resource group view.
3. From the Manage Copies view, select **Clones** either from the primary or secondary (mirrored or replicated) storage systems.
4. Select the PDB clone (targetCDB:PDBClone) and then click .
5. Review the estimated size of the clone that is to be split and the required space available on the aggregate, and then click **Start**.
6. Monitor the operation progress by clicking **Monitor > Jobs**.

Monitor Oracle database clone operations


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

About this task

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

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

Steps

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

Refresh a clone

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



You cannot refresh a PDB clone.

What you will need

- Create an online full backup or an offline data backup policy with no scheduled backups enabled.
- Configure the email notification in the policy for backup failures only.
- Define the retention count for the on-demand backups appropriately to ensure that there are no unwanted

backups.

- Ensure that only an online full backup or an offline data backup policy is associated with resource group which is identified for refresh clone operation.
- Create a resource group with only one database.
- If a cron job is created for the clone refresh command, ensure that the SnapCenter schedules and the cron schedules are not overlapping for the database resource group.

For a cron job created for the clone refresh command, ensure that you run `Open-SmConnection` after every 24hrs.

- Ensure that the clone SID is unique for a host.

If multiple refresh clone operations use the same clone specification file or use the clone specification file with same clone SID, existing clone with the SID on the host will be deleted and then the clone will be created.

- Ensure that the backup policy is enabled with secondary protection and the clone specification file is created with “-IncludeSecondaryDetails” to create the clones using secondary backups.
 - If the primary clone specification file is specified but the policy has secondary update option selected, the backup will be created, and update will get transferred to secondary. However, the clone will be created from the primary backup.
 - If the primary clone specification file is specified and the policy does not have secondary update option selected, the backup will be created on primary and clone will be created from primary.

Steps

1. Initiate a connection session with the SnapCenter Server for a specified user: *Open-SmConnection*
2. Create an Oracle database clone specification from a specified backup: *New-SmOracleCloneSpecification*



If secondary data protection policy is unified mirror-vault, then specify only `-IncludeSecondaryDetails`. You do not have to specify `-SecondaryStorageType`.

This command automatically creates an Oracle database clone specification file for the specified source database and its backup. You must also provide a clone database SID so that the specification file created has the automatically generated values for the clone database which you will be creating.



The clone specification file is created at `/var/opt/snapcenter/sco/clone_specs`.

3. Run *Refresh-SmClone*.

If the operation fails with the "PL-SCO-20032: canExecute operation failed with error: PL-SCO-30031: Redo log file +SC_2959770772_clmdb/clmdb/redolog/redo01_01.log exists" error messages, specify a higher value for `-WaitToTriggerClone`.

For detailed information on UNIX commands, see the [SnapCenter Software Command Reference Guide](#).

Delete clone of a pluggable database


You can delete the clone of a pluggable database (PDB) if you no longer require.

If you created a backup of the target CDB where the PDB is cloned, when you delete the PDB clone, the cloned PDB is also removed from the backup of the target CDB.



The PDB clones are not displayed in the inventory or resources view.

Steps

1. In the left navigation pane, click **Resources**, and then select the appropriate plug-in from the list.
2. Select the source container database (CDB) from the resource or resource group view.
3. From the Manage Copies view, select **Clones** either from the primary or secondary (mirrored or replicated) storage systems.
4. Select the PDB clone (targetCDB:PDBClone) and then click .
5. Click **OK**.

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.