



Clone an Oracle database backup

SnapCenter Software

Soumik Das, Archana
August 05, 2021

This PDF was generated from https://docs.netapp.com/us-en/snapcenter/protect-sco/task_clone_an_oracle_database_backup.html on September 16, 2021. Always check docs.netapp.com for the latest.

Table of Contents

Clone an Oracle database backup 1

Clone an Oracle database backup

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

About this task

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



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

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

During cloning, SnapCenter mounts the log backup 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. On 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-left: 1px solid #ccc; padding-left: 10px; margin-left: 20px;">  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. On 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 <code>SQLNET.AUTHENTICATION_SERVICES</code> is set to <code>NONE</code> in <code>sqlnet.ora</code> 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. On 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.

- 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. On 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> <p> The time can be specified in 24 hour format.</p>
Until SCN (System Change Number)	SnapCenter recovers the database up to a specified system change number (SCN).
Specify external archive log locations	Specify the external archive log location.
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>

Field name	Description
Enter sql entries to apply when clone is created	Add the sql entries that you want to apply when the clone is created.
Enter scripts to run after clone operation	<p>Specify the path and the arguments of the postscript that you want to run after the clone operation.</p> <p>You should store the postscript either in <code>/var/opt/snapcenter/spl/scripts</code> or in any folder inside this path. By default, the <code>/var/opt/snapcenter/spl/scripts</code> path is populated.</p> <p>If you have placed the script in any folder inside this path, you need to provide the complete path up to the folder where the script is placed.</p>

- 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 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.

- 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](#)

Copyright Information

Copyright © 2021 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

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.