



Clone cloud-native Oracle databases

BlueXP backup and recovery

NetApp
April 18, 2024

Table of Contents

- Clone cloud-native Oracle databases 1
 - Clone concepts and requirements 1
 - Clone cloud-native Oracle databases 7

Clone cloud-native Oracle databases

Clone concepts and requirements

You can clone an Oracle database residing on Amazon FSx for NetApp ONTAP or Cloud Volumes ONTAP using the backup of the database either to the source database host or to an alternate host. You can clone the backup from primary storage systems.

Before cloning the database, you should understand the clone concepts and ensure that all the requirements are met.

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.
You should have successfully created online data and log backup for the cloning operation to succeed.
- In the `asm_diskstring` parameter, you should configure:
 - `AFD:*` if you are using ASMFD
 - `ORCL:*` if you are using ASMLIB
 - `/dev/<exact_device_location>` if you are using ASMUDEV
- If you are creating the clone on an alternate host, the alternate host should meet the following requirements:
 - The plug-in should be installed on the alternate host.
 - Oracle software should be installed on the alternate host.
 - The clone host should be able to discover LUNs from storage if you are cloning a database residing on iSCSI SAN storage.
If you are cloning to an alternate host, then make sure that an iSCSI session is established between the storage and the alternate host.
 - 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 diskgroup should be provisioned prior to the clone operation if you want to place archive log files of the cloned database in a dedicated ASM diskgroup.
 - The name of the data diskgroup can be configured but ensure that the name is not used by any other ASM diskgroup on the host where the clone will be performed.
 - Data files residing on the ASM diskgroup are provisioned as part of the clone workflow.

Limitations

- Cloning of databases residing on Azure NetApp Files is not supported.
- Cloning of databases residing on Qtree is not supported.
- Backing up a cloned database is not supported.
- If daily automatic backups are enabled on Amazon FSx for NetApp ONTAP, the cloned volumes on Amazon FSx for NetApp ONTAP cannot be deleted from BlueXP UI because FSx would have created

backups on the cloned volumes.

You should delete the cloned volumes after deleting all the backups for the volume from FSx UI and then delete the clones from the BlueXP UI using force option.

Clone methods

You can create clone either using the basic method or using the clone specification file.

Clone using basic method

You can create the clone with the default configurations based on the source database and the selected backup.

- The database parameters, home, and OS user are defaulted to the source database.
- The data file paths are named based on the naming scheme selected.
- The pre-script, post-script, and SQL statements cannot be specified.
- The recovery option is by default **until cancel** and it uses the log backup associated with the data backup for recovery

Clone using specification file

You can define the configurations in the clone specification file and use it to clone the database. You can download the specification file, modify it to your requirement, and then upload the file. [Learn more](#).

The different parameters defined in the specification file and that can be modified are as follows:

| Parameter | Description |
|-------------------------|--|
| control_files | Location of control files for the clone database. The number of control files will be same as source database. If you want to override the control file path, you can provide a different control file path. The file system or the ASM diskgroup should exist on the host. |
| redo_logs | Location, size, redo group number of redo logs. A minimum of two redo log groups are required to clone the 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. The file system or the ASM diskgroup should exist on the host. |
| oracle_version | Version of Oracle on the target host. |
| oracle_home | Oracle home on the target host. |
| enable_archive_log_mode | Controls the archive log mode for the clone database |

| Parameter | Description |
|---------------------|--|
| database_parameters | Database parameters for the cloned database |
| sql_statements | The SQL statements to be executed on the database after cloning |
| os_user_detail | Oracle OS user on the target clone database |
| database_port | Port used for communicating with the database if OS authentication is disabled on the host. |
| asm_port | Port used for communicating with ASM database if credentials are provided in the create clone input. |
| skip_recovery | Does not perform recovery operation. |
| until_scn | Recovers the database up to the specified system change number (scn). |
| until_time | Recovers the database up to the specified date and time. The accepted format is <i>mm/dd/yyyy hh:mm:ss</i> . |
| until_cancel | Recovers by mounting the log backup associated with the data backup that was selected for cloning. The cloned database is recovered till the missing or corrupt log file. |
| log_paths | Additional locations of archive log paths to be used for recovering the cloned database. |
| source_location | Location of the diskgroup or mount point on the source database host. |
| clone_location | Location of the diskgroup or mount point that needs to be created on the target host corresponding to the source location. |
| location_type | It can be either ASM_Diskgroup Or mountpoint. The values are auto-populated at the time of downloading the file. You should not edit this parameter. |

| Parameter | Description |
|-------------|--|
| pre_script | Script to be executed on the target host before creating the clone. |
| post_script | Script to be executed on the target host after creating the clone. |
| path | Absolute Path of the script on the clone host. You should store the script either in /var/opt/snapcenter/spl/scripts or in any folder inside this path. |
| timeout | The timeout time specified for the script running on the target host. |
| arguments | Arguments specified for the scripts. |

Clone naming scheme

Clone naming scheme defines what will be the location of the mount points and name of the diskgroups of the cloned database. You can either select **Identical** or **Auto-generated**.

Identical naming scheme

If you select the clone naming scheme as **Identical**, the location of mount points and the name of the diskgroups of the cloned database will be same as the source database.

For example, if the mount point of the source database is `/netapp_sourcedb/data_1 , +DATA1_DG`, for the cloned database the mount point remains the same for both NFS and ASM on SAN.

- Configurations like number and path of control files and redo files will be same as source.



If the redo logs or control file paths are located on the non-data volumes, then the user should have provisioned the ASM diskgroup or mountpoint in the target host.

- Oracle OS user and Oracle version will be same as source database.
- Clone storage volume name will be in the following format
`sourceVolNameSCS_Clone_CurrentTimeStampNumber`.

For example, if the volume name on the source database is `sourceVolName`, the cloned volume name will be `sourceVolNameSCS_Clone_1661420020304608825`.



The *CurrentTimeStampNumber* provides the uniqueness in volume name.

Auto-generated naming scheme

If you select the cloning scheme as **Auto-generated**, the location of mount points and the name of the diskgroups of the cloned database will be appended with a suffix.

- If you have selected the basic clone method, the suffixed will be the **Clone SID**.
- If you have selected the specification file method, the suffix will be the **Suffix** that was specified while downloading the clone specification file.

For example, if the mount point of the source database is `/netapp_sourcedb/data_1` and the **Clone SID** or the **Suffix** is `HR`, then the mount point of the cloned database will be `/netapp_sourcedb/data_1_HR`.

- Number of control files and redo log files will be same as the source.
- All redo log files and control files will be located on one of the cloned data mount points or data ASM diskgroups.
- Clone storage volume name will be in the following format
`sourceVolNameSCS_Clone_CurrentTimeStampNumber`.

For example, if the volume name on the source database is `sourceVolName`, the cloned volume name will be `sourceVolNameSCS_Clone_1661420020304608825`.



The *CurrentTimeStampNumber* provides the uniqueness in volume name.

- The format of the NAS mount point will be *SourceNASMountPoint_suffix*.
- The format of the ASM diskgroup will be *SourceDiskgroup_suffix*.



If the number of characters in the clone diskgroup is greater than 25 then it will have *SC_HashCode_suffix*.

Database parameters

The value of the following database parameters will be same as that of the source database irrespective of the clone naming scheme.

- `log_archive_format`
- `audit_trail`
- `processes`
- `pga_aggregate_target`
- `remote_login_passwordfile`
- `undo_tablespace`
- `open_cursors`
- `sga_target`
- `db_block_size`

The value of the following database parameters will be appended with a suffix based on the clone SID.

- `audit_file_dest = {sourcedatabase_parametervalue}_suffix`
- `log_archive_dest_1 = {sourcedatabase_oraclehome}_suffix`

Supported predefined environment variables for clone specific prescript and postscript

You can use the supported predefined environment variables when you execute the prescript and postscript while 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_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

Supported delimiters

- **@** is used to separate data from its database name and to separate the value from its key.
Example: DATA@RG2_scspr2417819002_07-20-2021_12.16.48.9267_0|LOG@RG2_scspr2417819002_07-20-2021_12.16.48.9267_1
- **|** is used to separate the data between two different entities for SC_BACKUP_NAME parameter.
Example: DATA@RG2_scspr2417819002_07-20-2021_12.16.48.9267_0|LOG@RG2_scspr2417819002_07-20-2021_12.16.48.9267_1
- **,** is used to separate set of variables for the same key.

Example: 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

Clone cloud-native Oracle databases

You can clone an Oracle database residing on Amazon FSx for NetApp ONTAP or Cloud Volumes ONTAP using the backup of the database either to the source database host or to an alternate host.



You might clone databases for the following reasons:


- To test functionality that must 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.


Before you begin

You should understand the clone concepts and ensure that all the requirements are met. [Learn more](#).

Steps

1. Click  corresponding to the database that you want to clone and click **View Details**.
2. Click  corresponding to the data backup and click **Clone**.
3. In the Clone Details page, select one of the clone options.
4. Depending on the option selected, perform the following actions:

| If you have selected... | Do this... |
|-------------------------|---|
| Basic | <ol style="list-style-type: none"> 1. Select the clone host. 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. 2. Specify the SID of the clone. 3. Select the clone naming scheme. If the database is cloned to the source host, the clone naming scheme will be auto-generated. If the database is cloned to an alternate host, clone naming scheme will be identical. 4. Specify the Oracle home path. 5. (Optional) Specify the database credentials. <ul style="list-style-type: none"> ◦ Database credential: If the OS User authentication is disabled, you should provide a password for the sys user to connect to the cloned database on the same or target host. ◦ ASM credential: If the OS user authentication is disabled on the target host, you should provide a sysasm privileged user's credentials to connect to the ASM instance on the target host. <div data-bbox="966 1199 1023 1255"></div> <div data-bbox="1073 1178 1450 1276">Ensure that the listener is up and running on the target host.</div> 6. Click Next. 7. Click Clone. |

| If you have selected... | Do this... |
|---------------------------|---|
| Specification file | <ol style="list-style-type: none"> 1. Click Download File to download the specification file. 2. Select the clone naming scheme. If you select, Auto-generated, you should specify the suffix. 3. Edit the specification file as per the requirement and upload it by clicking the Browse button. 4. Select the clone host. 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. 5. Specify the SID of the clone. 6. (Optional) Specify the database credentials. <ul style="list-style-type: none"> ◦ Database credential: If the OS User authentication is disabled, you should provide a password for the sys user to connect to the cloned database on the same or target host. ◦ ASM credential: If the OS user authentication is disabled on the target host, you should provide a sysasm privileged user's credentials to connect to the ASM instance on the target host. <div data-bbox="971 1247 1024 1304"></div> <div data-bbox="1084 1228 1442 1325">Ensure that the listener is up and running on the target host.</div> 7. Click Next. 8. Click Clone. |

5. Click  adjacent to **Filter By** and select **Clone options > Clones** to view the clones.

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.