



Create database backups

SnapManager for SAP

NetApp
February 20, 2023

Table of Contents

- Create database backups 1
 - About this task 1
 - Example 8
- Prune archive log files 8
- Consolidate archive log backups 10
- Schedule archive log file pruning 11

Create database backups

You can create backups of entire databases or portions of databases, including tablespaces, data files, or control files.

About this task

SnapManager provides Snapshot copy capabilities for databases across many host-side storage stacks, including NFS, Veritas, and others.

Administrators can optionally register backups with Oracle RMAN, which facilitates the use of RMAN to restore and recover the database at finer granularities such as blocks.

While defining the profile, you can customize the names of the Snapshot copies created by backups of that profile. For example, you might insert a prefix string of **HOPS** to denote High Operations backups.

In addition to defining unique names for Snapshot copies created by backups, you can also create unique labels for the backups themselves. When you create a backup, it is a good practice to supply a name for the backup so you have an easy way to identify it by using the `-label` parameter. This name must be unique for all backups created within a particular profile. The name can contain letters, numbers, underscore (`_`), and hyphen (`-`). It cannot start with a hyphen. Labels are case-sensitive. You might want to append information such as operating system environment variables, system date, and backup type.

If you do not supply a label, SnapManager creates a default label name in the form `scope_mode_datestring`, where `scope` is full or partial and `mode` is offline, online, or automatic (the letter `c` for cold, `h` for hot, or `a` for automatic).

From SnapManager 3.4, you can provide your own backup label by overriding the default backup label created by SnapManager. You must set the value of the `override.default.backup.pattern` parameter to **true** and specify the new backup label in the `new.default.backup.pattern` parameter. The backup label pattern can contain keywords such as database name, profile name, scope, mode and hostname, which has to be separated by underscore. For example,

```
new.default.backup.pattern=dbname_profile_hostname_scope_mode.
```



The timestamp will be included automatically at the end of the generated label.

When you enter a comment, you can include spaces and special characters. In contrast, when you enter a label, do not include spaces or special characters.

For each backup, SnapManager automatically generates a GUID, which is a 32-character HEX string. To determine the GUID, you must run the `backup list` command with the `-verbose` option.

You can create a full backup of a database while it is online or offline. To let SnapManager handle backing up a database regardless of whether it is online or offline, you should use the `-auto` option.

While creating a backup, if you have enabled pruning and the summary notification was enabled in the profile, two separate emails are triggered. One email is for the backup operation and the other for the pruning. You can correlate these emails by comparing the backup name and backup ID contained in these emails.

You can create a cold backup when the database is in the shutdown state. If the database is in a mounted state, change it to a shutdown state and perform the offline backup (cold backup).

SnapManager (3.2 or later) enables you to back up the archive log files separately from the data files, enabling you to manage the archive log files efficiently.

To create the archive log backups separately, you must create a new profile or update the existing profile to separate the archive log backups by using the `-separate-archivelog-backups` option. Using the profile, you can perform the following SnapManager operations:

- Create an archive log backup.
- Delete an archive log backup.
- Mount an archive log backup.
- Free an archive log backup.

The backup options vary depending on the profile settings:

- Using a profile that is not separated to take archive log backups separately allows you to do the following:
 - Create a full backup.
 - Create a partial backup.
 - Specify archive log destinations to be backed up for archive log files.
 - Specify archive log destinations to be excluded from the backup.
 - Specify the pruning options for deleting the archive log files from the archive log destinations.
- Using a profile that is separated to take archive log backups allows you to do the following:
 - Create a data files-only backup.
 - Create an archivelogs-only backup.
 - While creating a data files-only backup, include the archive log backup along with the online data files only backup for cloning.

If you have included archive log backups along with data files in the **Profile Settings** page of the **Profile Create** wizard from the SnapManager GUI, and if you have not selected the **Archivelogs** option in the **Backup Create** wizard, SnapManager always creates the archive log backup along with data files for all online backups.

In such a situation, from the SnapManager CLI, you can consider all the archive log destinations for backup except for the exclude destinations specified in the SnapManager configuration file. But you cannot prune these archive log files. However, you can still use the `-archivelogs` option to specify the archive log file destination and prune the archive log files from the SnapManager CLI.

If you are creating the backup using the `-auto` option and specify the `-archivelogs` option, SnapManager creates either an online or offline backup based on the current status of the backup.

- SnapManager creates an offline backup when the database is offline and does not include the archive log files in the backup.
- SnapManager creates an online backup including archive log files when the database is online.
- While creating the archivelogs-only backup:
 - Specify the archive log destination to be backed up along with the archivelogs-only backup
 - Specify the archive log destinations to be excluded from the archive logs-only backup
 - Specify the pruning options for deleting the archive log files from the archive log destinations

• Scenarios not supported

- You cannot create the archivelog-only backup along with an offline data files-only backup.
- You cannot prune the archive log files when the archive log files are not backed up.
- You cannot prune the archive log files when Flash Recovery Area (FRA) is enabled for archive log files.

If you specify the archive log location in Flash Recovery Area, you must ensure that you also specify the archive log location in the `archive_log_dest` parameter.



While creating archive log backups, you must enter the full archive log destinations paths within double quotation marks and the destination paths separated by commas. The path separator should be given as two backslashes (`\\`) instead of one.

When you specify the label for online data files backup with included archive log backup, the label is applied for data files backup, and the archive log backup will be suffixed with (`_logs`). This suffix can be configured by changing the parameter `suffix.backup.label.with.logs` parameter in the SnapManager configuration file.

For example, you can specify the value as `suffix.backup.label.with.logs=arc` so that the `_logs` default value is changed to `_arc`.

If you have not specified any archive log destinations to be included in the backup, then SnapManager includes all the archive log destinations configured in the database.

If any archive log files are missing in any one of the destinations, SnapManager skips all these archive log files created before the missing archive log files even if these files are available in other archive log destination.

While creating archive log backups, you must specify the archive log file destinations to be included in the backup, and can set the configuration parameter to include the archive log files always beyond the missing files in the backup.



By default, this configuration parameter is set to **true** to include all the archive log files, beyond missing files. If you are using your own archive log pruning scripts or manually deleting archive log files from the archive log destinations, you can disable this parameter, so that SnapManager can skip the archive log files and proceed further with the backup.

SnapManager does not support the following SnapManager operations for archive log backups:

- Clone the archive log backup
- Restore archive log backup
- Verify archive log backup

SnapManager also supports backing up the archive log files from the flash recovery area destinations.

Step

1. Enter the following command:


```
smsap backup create -profile profile_name {[-full {-online | -offline | -auto}
[-retain {-hourly | -daily | -weekly | -monthly | -unlimited}] [-verify] | [-
data [[-files files [files]] | [-tablespaces -tablespaces [-tablespaces]] [-
datalabel label] {-online | -offline | -auto} [-retain {-hourly | [-daily |
```

```
-weekly | -monthly | -unlimited}] [-verify] | [-archivelogs [-label label] [-comment comment] [-backup-dest path1 [, [path2]]] [-exclude-dest path1 [, path2]]] [-prunelogs {-all | -untilSCN untilSCN | -until-date yyyy-MM-dd:HH:mm:ss | -before {-months | -days | -weeks | -hours}}] -prune-dest prune_dest1, [prune_dest2] [-taskspec taskspec] [-dump] [-force] [-quiet | -verbose]
```

If you want to...	Then...
<p>Specify whether you want to take a backup of an online or offline database, rather than allowing SnapManager to handle whether it is online or offline</p>	<p>Specify <code>-offline</code> to take a backup of the offline database. Specify <code>-online</code> to take a backup of the online database.</p> <p>If you use these options, you cannot use the <code>-auto</code> option.</p>
<p>Specify whether you want to let SnapManager handle backing up a database regardless of whether it is online or offline</p>	<p>Specify the <code>-auto</code> option. If you use this option, you cannot use the <code>--offline</code> or <code>-online</code> option.</p>
<p>Specify whether you want to perform a partial backup of specific files</p>	<div data-bbox="846 810 1487 1066" style="border: 1px solid #ccc; padding: 10px; margin-bottom: 10px;"> <p>Specify the <code>-data-files</code> option and then list the files, separated by commas. For example, list the file names <code>f1</code>, <code>f2</code>, and <code>f3</code> after the option.</p> </div> <p>Example for creating a partial datafile backup on Windows</p> <div data-bbox="846 1199 1487 1455" style="border: 1px solid #ccc; padding: 10px;"> <pre>smsap backup create -profile nosep -data -files "J:\mnt\user\user.dbf" -online -label partial_datafile_backup -verbose</pre> </div>

If you want to...	Then...
Specify whether you want to perform a partial backup of specific tablespaces	<p>Specify the <code>-data -tablespaces</code> option and then list the tablespaces, separated by commas. For example, use <code>ts1, ts2, and ts3</code> after the option.</p> <p>SnapManager supports backing up of read-only tablespaces. While creating the backup, SnapManager changes the read-only table spaces to read-write. After creating the backup, the tablespaces are changed to read-only.</p> <p>Example for creating a partial tablespace backup</p> <pre>smsap backup create -profile nosep -data -tablespaces tb2 -online -label partial_tablespace_bkup -verbose</pre>
Specify whether you want to create a unique label for each backup in the following format: full_hot_mybackup_label	<p>For Windows, you might enter this example:</p> <pre>smsap backup create -online -full -profile targetdbl_prof1 -label full_hot_my_backup_label -verbose</pre>

If you want to...	Then...
<p>Specify whether you want to create backup of the archive log files separately from the data files</p>	<p>Specify the following options and variables:</p> <ul style="list-style-type: none"> • <code>-archivelogs</code> creates a backup of the archive log files. • <code>-backup-dest</code> specifies the archive log file destinations to be backed up. • <code>-exclude-dest</code> specifies the archive log destinations to be excluded. • <code>-label</code> specifies the label for the archive log file backup. <div style="border: 1px solid #ccc; padding: 5px; margin: 10px 0;">  You must provide either the <code>-backup-dest</code> option or the <code>-exclude-dest</code> option. </div> <p>Providing both these options together along with the backup displays error message <code>You have specified an invalid backup option.</code> Specify any one of the options: <code>-backup-dest</code>, or <code>exclude-dest</code>.</p> <p>Example for creating archive log file backups separately on Windows</p> <pre style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> smsap backup create -profile nosep -archivelogs -backup-dest "J:\mnt\archive_dest_2\" -label archivelog_backup -verbose </pre>
<p>Specify whether you want to create backup of data files and archive log files together</p>	<p>Specify the following options and variables:</p> <ul style="list-style-type: none"> • <code>-data</code> option to specify the data files. • <code>-archivelogs</code> option to specify the archive log files. Example for backing up data files and archive log files together on Windows <div style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;"> <pre style="background-color: #f0f0f0; padding: 5px;"> smsap backup create -profile nosep -data -online -archivelogs -backup-dest "J:\mnt\archive_dest_2\" -label data_arch_backup -verbose </pre> </div>

If you want to...	Then...
<p>Specify whether you want to prune the archive log files while creating a backup</p>	<p>Specify the following options and variables:</p> <ul style="list-style-type: none"> • <code>-prunelogs</code> specifies to delete the archive log files from the archive log destinations. <ul style="list-style-type: none"> ◦ <code>-all</code> specifies to delete all the archive log files from the archive log destinations. ◦ <code>-until-scnuntil-scn</code> specifies to delete the archive log files until a specified SCN. ◦ <code>-until-dateyyyy-MM-dd:HH:mm:ss</code> specifies to delete the archive log files until the specified time period. ◦ <code>-before</code> option specifies to delete the archive log files before the specified time period (days, months, weeks, hours). ◦ <code>-prune</code> <code>-destprune_dest1,[prune_dest2</code> specifies to delete the archive log files from the archive log destinations while creating the backup. <div style="border: 1px solid gray; padding: 5px; margin-top: 10px;">  You cannot prune the archive log files when Flash Recovery Area (FRA) is enabled for archive log files. </div> <p>Example for pruning all archive log files while creating a backup on Windows</p> <pre style="background-color: #f0f0f0; padding: 10px; border: 1px solid #ccc;"> smsap backup create -profile nosep -archivelogs -label archive_prunebackup1 -backup-dest "E:\\oracle\\MDV\\oraarch\\MDVarc h,J:\\ " -prunelogs -all -prune-dest "E:\\oracle\\MDV\\oraarch\\MDVarc h,J:\\" -verbose </pre>
<p>Specify whether you want to add a comment about the backup</p>	<p>Specify <code>-comment</code> followed by the description string.</p>
<p>Specify whether you want to force the database into the state you have specified to back it up, regardless of the state it is currently in</p>	<p>Specify the <code>-force</code> option.</p>

If you want to...	Then...
Specify whether you want to verify the backup at the same time you create it	Specify the <code>-verify</code> option.
Specify whether you want to collect the dump files after the database backup operation	Specify <code>-dump</code> option at the end of the backup create command.

Example

```
smsap backup create -profile targetdb1_prof1 -full -online -force -verify
```

Prune archive log files

You can prune the archive log files from the archive log locations while creating a backup.

What you'll need

- Archive log files must be backed up by the current backup operation.

If pruning is specified along with other backups that do not contain archive log files, the archive log files are not pruned.

- The database must be in the mounted state.

If the database is not in mounted state, enter the `-force` option along with backup command.

About this task

While performing a backup operation, you can specify the following:

- Scope of pruning:
 - Delete all the archive log files.
 - Delete the archive log files until the specified System Change Number (SCN).
 - Delete the archive log files until the specified time.
 - Delete the archive log files before the specified time period.
- Destination from where the archive log files must be pruned.



Even when the archive log file pruning fails in one destination, SnapManager continues to prune the archive log files from the other destinations.

Before deleting the archive log files, SnapManager verifies the following:

- Archive log files are backed up at least once.

- Archive log files are shipped to Oracle Dataguard Standby database, if any.
- Archive log files are captured by Oracle streams capture process, if any.

If the archive log files are backed up, shipped to standby, and captured by the capture process, SnapManager deletes all the archive log files in a single execution. However, if there are any archive log files that are not backed up, not shipped to standby, or not captured by the capture process, SnapManager deletes the archive log files one-by-one. The deletion of archive logs files in a single execution is faster than deleting archive logs one-by-one.

SnapManager can also group the archive log files and delete them batch-by-batch. Each batch will have a maximum of 998 files. This value can be configured below 998 by using the configuration parameter `maximum.archive.log.files.toprune.atATime` in the `smsap.config` file.

SnapManager uses Oracle Recovery Manager (RMAN) commands to delete the archive log files. However, SnapManager does not integrate with the RMAN retention policies and deletion policies.



If you delete the archive log files from the archive log destinations, the pruning of archive log files fails.

SnapManager does not support pruning of archive log files in the following scenarios:

- Archive log files are located in the flash recovery area.
- Archive log files are located in the Standby database.
- Archive log files are managed by both SnapManager and RMAN.

Step

1. Enter the following command:

```
smsap backup create -profile profile_name {[-full {-online | -offline | -auto}
[-retain {-hourly | [-daily | -weekly | -monthly | -unlimited]} [-verify] | [-
data [[-files files [files]] | [-tablespaces -tablespaces [-tablespaces]] [-
datalabel label] {-online | -offline | -auto} [-retain {-hourly | [-daily |
-weekly | -monthly | -unlimited]} [-verify] | [-archivelogs [-label label] [-
comment comment][-backup-dest path1 [, [path2]]] [-exclude-dest path1 [,
path2]]] [-prunelogs {-all | -untilSCN untilSCN | -until-date yyyy-MM-
dd:HH:mm:ss | -before {-months | -days | -weeks | -hours}} -prune-dest
prune_dest1, [prune_dest2]] [-taskspec taskspec]} -dump [-force] [-quiet |
-verbose]
```

If you want to...	Then...
<p>Prune archive log files</p>	<p>Specify the following options:</p> <ul style="list-style-type: none"> • <code>-prunelogs</code> specifies deleting the archive log files while creating a backup. <ul style="list-style-type: none"> ◦ <code>-all</code> specifies deleting all the archive log files. ◦ <code>-untilSCN</code> specifies deleting the archive log files until the specified SCN. ◦ <code>-until-date</code> specifies deleting the archive logs including the specified date and time. ◦ <code>-before {-months -days -weeks -hours}</code> specifies deleting the archive log files before the specified time period.
<p>Include the destination from where the archive log files are to be pruned</p>	<p>Specify the <code>-prune-dest</code> option.</p>

Consolidate archive log backups

SnapManager consolidates the archivelog-only backups every time you take a backup by freeing up the duplicate archivelog-only backups. By default, consolidation is enabled.

About this task

SnapManager identifies the archivelog-only backups which has archive log files in other backups and frees them to maintain minimum number of archivelog-only backups with unique archive log files.

If the archivelog-only backups are freed by consolidation, then these backups are deleted based on the archive log retention duration.

When the database is in the shutdown or nomount state during archive log consolidation, SnapManager changes the database to the mount state.

If the backup or pruning of archive log files fails, then consolidation will not be done. Consolidation of archivelog-only backups is followed only after successful backups and successful pruning operations.

Steps

1. To enable consolidation of the archivelog-only backups, modify the configuration parameter `consolidation` and set the value as `true` in the SnapManager configuration file (`smsap.config`).

Once the parameter is set, the archivelog-only backups are consolidated.

If the newly-created archivelog-only backup contains the same archive log files in any of the earlier archivelog-only backups, then the earlier archive-log only backups are freed.



SnapManager does not consolidate the archive log backup taken along with the datafiles backup. SnapManager consolidates the archivelog-only backup.



SnapManager consolidates the archive log backups even when user manually deletes the archive log files from the archive log destinations or when the archive log files are corrupted and might be included the backup.

- To disable consolidation of the archive log backups, modify the configuration parameter `consolidation` and set the value as `false` in the SnapManager configuration file (`smsap.config`).

Schedule archive log file pruning

When you create a backup, you can schedule the pruning of archive log files to occur at a specified time.

About this task

SnapManager allows you to prune the archive log files periodically from the active file system.

Step

- Enter the following command:

```
smsap schedule create -profile profile_name {[-full {-online | -offline |
-auto}[-retain [-hourly | -daily | -weekly | -monthly | -unlimited] [-verify]]
| [-data [-files files [files]] | [-tablespaces -tablespaces [-tablespaces]]
{-online | -offline | -auto}[-retain [-hourly | -daily | -weekly | -monthly |
-unlimited] [-verify]] | [-archivelogs]} [-comment comment] [-backup-dest
path1 [, [path2]] [-exclude-dest path1 [, path2]] [-prunelogs{-all | -untilSCN
untilSCN | -before {-date yyyy-MM-dd HH:mm:ss | -months months | -weeks weeks
| -days days | -hours hours}} -prune-dest prune_dest1, prune_dest2] -schedule
-name schedule_name [-schedule-comment schedule_comment] -interval {-hourly |
-daily | -weekly | -monthly | -onetimeonly} -cronstring cronstring -start-time
{start-time start_time <yyyy-MM-dd HH:mm>} -runasuser -runasuser [-force] [-
quiet | -verbose]
```

If you want to...	Then...
Schedule pruning of archive log files	Specify the following options: <ul style="list-style-type: none"> <code>-prunelogs</code> to schedule pruning of the archive log files <code>-prune-dest</code> to prune archive log files from the archive log destinations
Include a name for the schedule	Specify the <code>-schedule-name</code> option.

If you want to...	Then...
Schedule pruning of archive log files at specific time interval	Specify the interval option and indicate whether the archive log files should be pruned based on the following interval classes: <ul style="list-style-type: none"> • -hourly • -daily • -weekly • -monthly • -onetimeonly
Add a comment about the schedule operation	Specify the <code>-schedule-comment</code> option followed by the description string.
Specify the start time of the schedule operation	Specify the <code>-start-time</code> option in the yyyy-mm-dd hh:mm format.

Copyright information

Copyright © 2023 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.