



Backing up databases

SnapManager for SAP

NetApp
February 20, 2023

This PDF was generated from <https://docs.netapp.com/us-en/snapmanager-sap/windows/concept-what-snapmanager-database-backups-are.html> on February 20, 2023. Always check docs.netapp.com for the latest.

Table of Contents

- Backing up databases 1
 - What SnapManager database backups are 2
 - What full and partial backups are 2
 - About control file and archive log file handling 8
 - What database backup scheduling is 9
 - Create database backups 12
 - What AutoSupport is 24
 - Verify database backups 26
 - Change the backup retention policy 26
 - View a list of backups 28
 - View backup details 28
 - Mount backups 29
 - Unmount backups 30
 - Free backups 30
 - Delete backups 31

Backing up databases

SnapManager enables the backing up of data on local storage resources by using post-processing scripts.

SnapManager provides the following options to back up, restore, and recover the data in your database:

- Back up the entire database or a portion of it.

If you back up a portion of it, specify a group of tablespaces or a group of data files.

- Back up the data files and archive log files separately.
- Back up databases to primary storage (also called local storage) and protect them by backing them up to secondary by using postprocessing scripts.
- Schedule routine backups.

How SnapManager (3.2 or later) differs from earlier SnapManager versions

SnapManager (3.1 or earlier) enables you to create full database backups that contain data files, control files, and archive log files.

SnapManager (3.1 or earlier) manages only the data files. The archive log files are maintained by using solutions outside SnapManager.

SnapManager (3.1 or earlier) imposes the following constraints in managing database backups:

- Performance impact

When you perform a full, online database backup (when the database is in the backup mode), the performance of the database reduces for the period of time until the backup is created. In SnapManager (3.2 or later), limited database backups and frequent archive log backups can be taken. Taking frequent archive log backups helps in preventing the database from being placed in backup mode.

- Manual restore and recovery

When the required archive log files do not exist in the active file system, database administrators have to identify which backup contains the archive log files, mount the database backups, and recover the restored database. This process is time consuming.

- Space constraints

When a database backup is created, the archive log destinations become full causing the database not to respond until sufficient space is created on the storage. In SnapManager (3.2 or later), the archive log files can be pruned from the active file system to free space periodically.

Why archive log backups are important

Archive log files are required to roll the database forward after a restore operation is performed. Every transaction on an Oracle database is captured in the archive log files (if the database is in the archive log mode). Database administrators can restore the database backups by using the archive log files.

Advantages of archive log-only backups

- Provides separate retention duration for archivelog-only backups

You can have less retention duration for the archivelog-only backups that are required for recovery.

- Protects the archivelog-only backups by using post-processing scripts
- Improves the performance of the database
- Consolidates archive log backups

SnapManager consolidates the archive log backups every time you take a backup by freeing the duplicate archive log backups.

What SnapManager database backups are

SnapManager enables you to perform different backup tasks. You can assign retention classes to specify how long the backup can be retained; once that time limit is reached, the backup is deleted.

- Create backups on the primary storage
- Create protected backups on the secondary storage resources by using postprocessing scripts
- Verify that the backups completed successfully
- View a list of backups
- Schedule backups by using the graphical user interface
- Manage the number of backups retained
- Free backup resources
- Mount and unmount backups
- Delete backups

SnapManager creates backups by using one of the following retention classes:

- Hourly
- Daily
- Weekly
- Monthly
- Unlimited

If new data files are added to the database, you should create a new backup immediately. Also, if you restore a backup taken before the new data files were added and attempt to recover to a point after the new data files were added, the automatic recovery process might fail. See the Oracle documentation to learn more about the process for recovering the data files added after a backup.

What full and partial backups are

You can choose to back up the entire database or just a portion of it. If you choose to back up a portion of the database, you can choose to back up a group of tablespaces or data files. You can choose to take a separate backup of both tablespaces and data files.

The following table lists the benefits and consequences of each type of backup:

Backup type	Advantages	Disadvantages
Full	Minimizes the number of Snapshot copies. For online backups, each tablespace is in backup mode for the entire time of the backup operation. SnapManager takes one Snapshot copy for each volume that the database uses, plus one Snapshot copy for each volume that the log files occupy.	For online backups, each tablespace is in backup mode for the entire time of the backup operation.
Partial	Minimizes the amount of time each tablespace spends in backup mode. SnapManager groups the Snapshot copies it takes by tablespace. Each tablespace is in backup mode only long enough to create the Snapshot copies. This method of grouping the Snapshot copies minimizes the physical block writes in the log files during an online backup.	The backup can require creating Snapshot copies of multiple tablespaces in the same volume. This method can cause SnapManager to create multiple Snapshot copies of a single volume during the backup operation.



Although you can perform a partial backup, you must always perform a full backup of the entire database.

Backup types and the number of Snapshot copies

The backup type (full or partial) affects the number of Snapshot copies that SnapManager creates. For a full backup, SnapManager creates a Snapshot copy of each volume, while for a partial backup, SnapManager creates a Snapshot copy of each tablespace file.



Data ONTAP limits the maximum number of Snapshot copies to 255 per volume. You might reach this maximum only if you configure SnapManager to retain a large number of backups where each backup consists of numerous Snapshot copies.

To keep an adequate pool of backups available while ensuring that the maximum limit of Snapshot copies per volume is not reached, you must remove backups when they are no longer needed. You can configure the SnapManager retention policy to remove successful backups after reaching a specific threshold for a specific backup frequency. For example, after SnapManager creates four successful daily backups, SnapManager removes the daily backups created on the previous day.

The following tables show how SnapManager creates Snapshot copies based on the backup type. The example in the tables assumes that database Z includes two volumes, each volume includes two tablespaces (TS1 and TS2), and each tablespace includes two database files (`ts1.data1`, `ts1.data2`, `ts2.data1`, and `ts2.data2`).

These tables show how the two types of backups produce different numbers of Snapshot copies.

SnapManager creates Snapshot copies at the volume level instead of the tablespace level, which usually reduces the number of Snapshot copies it must create.



Both backups also create Snapshot copies of the log files.

Volumes in database	Tablespace TS1 (includes 2 database files)	Tablespace TS2 (includes 2 database files)	Snapshot copies created	Total number of Snapshot copies
E:\data	TS1.data1	TS2.data1	1 per volume	2

Volumes in database	Tablespace TS1 (includes 2 database files)	Tablespace TS2 (includes 2 database files)	Snapshot copies created	Total number of Snapshot copies
E:\data	TS1.data1	TS2.data1	2 per file	4

Full online backups

During a full online backup, SnapManager backs up the entire database and creates Snapshot copies at the volume level (not at the tablespace level).

SnapManager creates two Snapshot copies for each backup. If all the files needed by the database are in a single volume, then both Snapshot copies appear in that volume.

When you specify a full backup, SnapManager performs the following actions:

Steps

1. Places the entire database in the online backup mode
2. Creates Snapshot copies of all the volumes containing database files
3. Takes the database out of the online backup mode
4. Forces a log switch and then archives the log files

This also flushes the redo information to disk.

5. Generates backup control files
6. Creates a Snapshot copy of the log files and the backup control files

When performing a full backup, SnapManager places the entire database in the online backup mode. An individual tablespace (for example, E:\data\system.data1) is in the online backup mode longer than certain tablespaces or data files that were specified.

When a database goes into backup mode, Oracle writes entire blocks to the logs and does not merely write the delta between backups. Because databases do more work in online backup mode, choosing a full backup places a greater load on the host.

Although performing full backups places a greater load on the host, full backups require fewer Snapshot copies, resulting in fewer storage requirements.

Partial online backups

Instead of a full backup, you can choose to perform a partial backup of the tablespaces in a database. While SnapManager takes a Snapshot copy of volumes for *full* backups, SnapManager takes a Snapshot copy of each specified tablespace for *partial* backups.

Because the tablespace level is the lowest level that Oracle allows into backup mode, SnapManager processes backups at the tablespace level, even if you specify a data file in a tablespace.

With a partial backup, each tablespace exists in backup mode for a shorter amount of time compared to a full backup. During an online backup, the database is always available to users; however, the database must perform more work and the host must perform more physical I/O. In addition, because it is taking Snapshot copies of each tablespace specified or each tablespace containing a specified data file instead of the entire volume, SnapManager takes more Snapshot copies.

SnapManager takes Snapshot copies of specific tablespaces or data files. The partial backup algorithm is a loop that SnapManager repeats until it has taken a Snapshot copy of each specified tablespace or data file.



Although you can perform a partial backup, it is recommended that you always perform a full backup of the entire database.

During a partial backup, SnapManager performs these actions:


Steps

1. Places the tablespace containing the data files into backup mode.
2. Takes a Snapshot copy of all the volumes used by the tablespace.
3. Takes the tablespace out of backup mode.
4. Continues this process, until it has taken a Snapshot copy of all the tablespaces or files.
5. Forces a log switch and then archives the log files.
6. Generates backup control files.
7. Takes a Snapshot copy of the log files and the backup control files.

Examples of backup, restore, and recover operations

You can find information about some of the backup, restore, and recover scenarios that you can use to accomplish your data protection goals.

The following illustration shows the contents of the tablespace:



In the illustration, Tablespace1 has one table and two database files associated with it. Tablespace2 has two tables and one database file associated with it.

The following tables describe some full and partial backup, restore, and recover scenarios:

Examples of full backup, restore, and recover operations

Full backup	Restore	Recover
<p>SnapManager makes a backup of everything in database DB1, including the data files, archive logs, and control files.</p>	<p>Complete restore with control files SnapManager restores all data files, tablespaces, and control files in the backup.</p>	<p>You can specify one of the following:</p> <ul style="list-style-type: none"> • SCN - Enter an SCN, such as 384641. • Date/Time - Enter a date and time of the backup, such as 2005-11-25:19:06:22. • The last transaction made to the database.
	<p>Complete restore without control files SnapManager restores all tablespaces and data files, without the control files.</p>	
	<p>Restore either data files or tablespaces with control files Specify one of the following:</p> <ul style="list-style-type: none"> • Tablespaces • Data files 	<p>SnapManager recovers the data to the last transaction made to the database.</p>
	<p>Restore either data files or tablespaces without control files SnapManager restores one of the following:</p> <ul style="list-style-type: none"> • Tablespaces • Data files 	
	<p>Restore control files only</p>	

Examples of partial backup, restore, and recover operations

Partial backup	Restore	Recover
<p>You can choose one of the following options:</p> <ul style="list-style-type: none"> • Tablespaces <p>You can specify Tablespace1 and Tablespace2 or only one of them.</p> <ul style="list-style-type: none"> • Data files <p>You can specify all three database files (TS1_1.dbf, TS1_2.dbf, and TS2_1.dbf), two files, or one file.</p> <p>Regardless of which option you select, the backup includes all the control files. Archive log files are included in the partial backup if the profile is not enabled to create the archive log backups separately.</p>	<p>Complete restore SnapManager restores all data files, tablespaces, and control files specified in the partial backup.</p>	<p>SnapManager recovers the data to the last transaction made to the database instance.</p>
<p>Restore either data files or tablespaces with control files SnapManager restores one of the following:</p> <ul style="list-style-type: none"> • All the data files specified • All the tablespaces specified 	<p>Restore either data files or tablespaces without control files SnapManager restores one of the following:</p> <ul style="list-style-type: none"> • Tablespaces <p>Specify any of the tablespaces. SnapManager restores only the tablespaces specified. If the backup contains Tablespace1, SnapManager restores only that tablespace.</p> <ul style="list-style-type: none"> • Data files <p>Specify any of the database files. SnapManager restores only the data files specified. If the backup contains database files (TS1_1.dbf and TS1_2.dbf), SnapManager restores only those files.</p>	<p>Restore control files only</p>

About control file and archive log file handling

SnapManager includes the control files and optionally includes archive log files with each

backup. Archive log files are used for recovery operations.

The database uses control files to identify names, locations, and sizes of the database files. SnapManager includes control files in each backup because control files are used in the restore process.

The changes to a database are tracked by using the online redo logs, which are eventually archived and known as archived redo logs (or archive logs). SnapManager (3.2 or later) enables you to backup data files and archive log files separately with different retentions and frequencies. SnapManager can take backups of only the archive logs or combined backups of data files and archive logs. SnapManager provides complete automated management of archive logs, and does not require any manual intervention for database recovery and also allows pruning of archive logs from one or more archive log destinations after the backup is taken.



To see which tablespaces and data files are included in a backup, use the `backup show` command or the Backup Properties window.

The following table illustrates how SnapManager handles control and archive log files during each operation:

Type of operation	Control files	Archive log files
Backup	Included with each backup	Can be included with each backup
Restore	Can be restored either alone or along with the tablespaces or data files	Can be used for the recovery process

What database backup scheduling is

You can schedule, update, and monitor backups for databases by using the Schedule tab of the graphical user interface.

The following table addresses some common scheduling questions:

Question	Answer
What happens to the scheduled backups when the SnapManager server restarts?	When the SnapManager server restarts, it automatically restarts all the schedules. However, SnapManager does not follow-up on any missed occurrences.

Question	Answer
<p>What happens when two backups are scheduled to occur on two databases at the same time?</p>	<p>SnapManager starts backup operations one at a time and then allows the backups to run in parallel. For example, if a database administrator creates six daily backup schedules for six different database profiles to occur at 1:00 a.m., all six backups run in parallel.</p> <p>If multiple backups are scheduled to occur on a single database profile in a short period of time, the SnapManager server runs only the backup operation with the longest retention duration.</p> <p>Before starting a backup operation, SnapManager first determines the following:</p> <ul style="list-style-type: none"> • Within the last 30 minutes, has another schedule successfully created a backup, with greater retention, for the same profile? • Within the next 30 minutes, will another schedule attempt to create a backup, with greater retention, for the same profile? <p>If the answer to either question is yes, SnapManager skips the backup.</p> <p>For example, a database administrator might create a daily, weekly, and monthly schedule for a database profile, all of which are scheduled to take backups at 1:00 a.m. On that one day of the month when three backups are scheduled to occur simultaneously at 1:00 a.m., SnapManager runs only the backup operation based on the monthly schedule.</p> <p>The time window of 30 minutes can be changed in a SnapManager properties file.</p>
<p>Under which user does the backup operation run?</p>	<p>The operation runs under the user who created the schedule. However, you can change this to your own user ID, if you have valid credentials for both the database profile and host. For instance, by launching Scheduled Backup Properties for the backup schedule created by Avida Davis, Stella Morrow can select her user ID in Perform this operation as user to run the scheduled backup.</p>
<p>How does the SnapManager scheduler interact with the native operating system scheduler?</p>	<p>On the SnapManager server, you cannot view the scheduled backups via the operating system's native scheduler. For instance, after creating a scheduled backup, you do not see a new entry in the Scheduled Tasks window.</p>

Question	Answer
<p>What happens if the clocks in the graphical user interface and the server are not in sync?</p>	<p>The clocks on the client and server are not synchronized. Therefore, you can schedule a backup in which the start time is in the future on the client but in the past on the server.</p> <p>For recurring backups, the server still fulfills the request. For instance, if the server receives a request to perform hourly backups starting on 01/30/08 at 3:00 p.m. but the current time is 3:30 p.m. on that day, the server performs its first backup at 4:00 p.m. and continues to perform backups every hour.</p> <p>However, for one-time only backups, the server handles the request as follows:</p> <ul style="list-style-type: none"> • If the start time is within the last five minutes of the current server time, SnapManager immediately begins the backup. • If the start time is greater than five minutes, SnapManager does not initiate the backup. <p>For instance, consider the following scenario:</p> <ul style="list-style-type: none"> • The clock in the graphical interface host is three minutes behind the actual time. • The current time on the client is 8:58 a.m. • You schedule a one-time backup to occur at 9:00 a.m. • You schedule another one-time backup to occur at 8:30 a.m. <p>When the server receives the first request, the time on the server is 9:01 a.m. Although the start time of the backup is in the past, SnapManager immediately performs the backup.</p> <p>When the server receives the second request, the start time of the backup is more than five minutes in the past. You will receive a message that the schedule request failed because the start time is in the past.</p> <p>You can change the time of five minutes in a SnapManager properties file.</p>
<p>What happens to the scheduled backups for a profile when the profile is deleted?</p>	<p>When a database profile is deleted, the SnapManager server deletes scheduled backups defined for that profile.</p>

Question	Answer
<p>How do scheduled backups behave during Daylight Savings Time or when you change the SnapManager server time?</p>	<p>SnapManager backup schedules get affected due to Daylight Savings Time or when you change the SnapManager server time.</p> <p>Consider the following implications when the SnapManager server time is changed:</p> <ul style="list-style-type: none"> • After the backup schedule is triggered, if the SnapManager server time falls back, then the backup schedule does not trigger again. • If Daylight Savings Time begins before the scheduled start time, the backup schedules are triggered automatically. • For example, if you are in the United States and you schedule hourly backups at 4 a.m. that should occur every 4 hours, backups will occur at 4 a.m., 8 a.m., 12 a.m., 4 a.m., 8 p.m., and midnight on the days before and after Daylight Savings Time adjustments in March and November. • Note the following if backups are scheduled for 2:30 a.m. every night: <ul style="list-style-type: none"> ◦ When the clocks fall back an hour, as the backup is already triggered, the backup does not trigger again. ◦ When the clocks spring forward an hour, the backup triggers immediately. If you are in the United States and want to avoid this issue, you must schedule your backups to start outside the 2:00 a.m. to 3:00 a.m. interval.

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...
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	Specify <code>-offline</code> to take a backup of the offline database. Specify <code>-online</code> to take a backup of the online database. If you use these options, you cannot use the <code>-auto</code> option.

If you want to...	Then...
<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="846 310 1485 569" style="border: 1px solid #ccc; padding: 10px; background-color: #f9f9f9;"> <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="846 701 1485 959" style="border: 1px solid #ccc; padding: 10px; background-color: #f9f9f9;"> <pre>smsap backup create -profile nosep -data -files "J:\\mnt\\user\\user.dbf" -online -label partial_datafile_backup -verbose</pre> </div>
<p>Specify whether you want to perform a partial backup of specific tablespaces</p>	<p>Specify the <code>-data -tablespaces</code> option and then list the tablespaces, separated by commas. For example, use <code>ts1</code>, <code>ts2</code>, and <code>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> <div data-bbox="846 1415 1485 1631" style="border: 1px solid #ccc; padding: 10px; background-color: #f9f9f9;"> <pre>smsap backup create -profile nosep -data -tablespaces tb2 -online -label partial_tablespace_bkup -verbose</pre> </div>

If you want to...	Then...
<p>Specify whether you want to create a unique label for each backup in the following format: full_hot_mybackup_label</p>	<p>For Windows, you might enter this example:</p> <pre data-bbox="846 226 1481 443"> smsap backup create -online -full -profile targetdbl_prof1 -label full_hot_my_backup_label -verbose </pre>
<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 data-bbox="867 562 1489 890" 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 data-bbox="873 940 1463 1052" style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  <p>You must provide either the <code>-backup-dest</code> option or the <code>-exclude-dest</code> option.</p> </div> <p>Providing both these options together along with the backup displays error message <code>You have specified an invalid backup option. Specify any one of the options: -backup-dest, or exclude-dest.</code></p> <p>Example for creating archive log file backups separately on Windows</p> <pre data-bbox="846 1398 1481 1614"> smsap backup create -profile nosep -archivelogs -backup-dest "J:\mnt\archive_dest_2\" -label archivelog_backup -verbose </pre>

If you want to...	Then...
Specify whether you want to create backup of data files and archive log files together	<p>Specify the following options and variables:</p> <ul style="list-style-type: none">• -data option to specify the data files.• -archivelogs option to specify the archive log files. Example for backing up data files and archive log files together on Windows <pre data-bbox="889 415 1489 716">smsap backup create -profile nosep -data -online -archivelogs -backup-dest "J:\\mnt\\archive_dest_2\\" -label data_arch_backup -verbose</pre>

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 #ccc; padding: 5px; margin-top: 10px;">  <p>You cannot prune the archive log files when Flash Recovery Area (FRA) is enabled for archive log files.</p> </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.toprun.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...
Prune archive log files	Specify the following options: <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.
Include the destination from where the archive log files are to be pruned	Specify the <code>-prune-dest</code> option.

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 <code>yyyy-mm-dd hh:mm</code> format.

What AutoSupport is

The AutoSupport feature enables SnapManager server to send AutoSupport messages to the storage system after the backup operation is complete.



SnapManager sends AutoSupport messages only for the successful backup operations.

You can enable or disable AutoSupport by assigning the following values to the `auto_support.on` configuration parameter in the `smsap.config` configuration file:

- **TRUE** - Enables AutoSupport
- **FALSE** - Disables AutoSupport



By default, AutoSupport is enabled in SnapManager.

Add storage systems operating in clustered Data ONTAP to the SnapManager server host

You must add the storage systems operating in clustered Data ONTAP to the SnapManager server host to enable AutoSupport. In SnapManager 3.3 and earlier, AutoSupport was supported only on storage systems operating in 7-Mode.

Steps

1. Add an Admin Storage Virtual Machine (SVM, formerly known as Vserver) and a SVM operating in clustered Data ONTAP to the SnapManager server host: **`sdcli transport_protocol set -f AdminVserver_name or Vserver_name -type HTTP -user admin`**

The enter the following command: message is displayed.

2. Enter the password that you provided while creating SVM.

After you run the command successfully, the New transport protocol has been set. message is displayed.

Enable AutoSupport in SnapManager

You must enable AutoSupport, so that storage systems receive messages from the SnapManager server for every successful backup operation.

About this task

AutoSupport can be enabled in two ways:

- By default, the new installation of SnapManager does not contain the `auto_support.on` parameter in the `smsap.config` configuration file. This implies that autosupport is enabled.
- You can manually configure the `auto_support.on` parameter.

Steps

1. Stop the SnapManager server.
2. In the `smsap.config` configuration file, set the value of the `auto_support.on` parameter to **TRUE**.

Example

```
auto_support.on=TRUE
```

3. Restart the SnapManager server.

Disable AutoSupport in SnapManager

You must disable AutoSupport if you do not want the storage system to receive messages from the SnapManager server for every successful backup operation.

About this task

By default, AutoSupport is enabled if the configuration file does not contain the `auto_support.on` parameter. In this scenario, you must add the `auto_support.on` parameter in the configuration file and set the value to **FALSE**.

Steps

1. Stop the SnapManager server.
2. In the `smsap.config` configuration file, set the value of the `auto_support.on` parameter to **FALSE**.

Example

```
auto_support.on=FALSE
```

3. Restart the SnapManager server.

Verify database backups

You can use the `backup verify` command to verify that the blocks in the database backup are not corrupted. The verify operation invokes the Oracle Database Verify utility for each data file in the backup.

About this task

SnapManager enables you to perform the verify operation at any time that is convenient for you and the users on your system. You can perform the verification immediately after creating the backup. You must specify the profile containing the backup and either the label or the ID of the backup you created.



The backup verify operation fails in a Windows environment if you are using SnapManager 3.0 and Oracle database 11.1.0.7. You must use Oracle database 11.2.0.1 or later.



You can specify `-dump` to collect the dump files after the backup verify operation.

Step

1. Enter the following command:

```
smsap backup verify -profile profile_name [-label label | -id id] [-force] [-dump] [-quiet | -verbose]
```

Change the backup retention policy

You can change properties of a backup so it is eligible or ineligible for deletion according to the retention policy.

About this task

When you create a backup, you can set its retention policy. You can later choose to either keep that backup for a longer period than the retention policy allows or specify that you no longer need the backup and want the retention policy to manage it.

Retain backups forever

You can specify that a backup should be ineligible for deletion by the retention policy to keep the backup indefinitely.

Step

1. To specify that a backup be retained on an unlimited basis, enter this command:

```
smsap backup update -profile profile_name {-label label [data | -archivelogs] | -id id} -retain -unlimited
```

Assign backups with a specific retention class

DBAs can assign a specific retention class of hourly, daily, weekly, or monthly to backups.

Assigning a specific retention class makes the backups performed under this change eligible for deletion.

Step

1. To assign a specific backup retention class, enter this command:

```
smsap backup update -profile profile_name {-label label [data | -archivelogs] | -id id} -retain [-hourly | -daily | -weekly | -monthly ]
```

Change the retention policy default behavior

When a backup expires based on the retention policy, SnapManager determines whether to delete the backup based on the retention settings. Deletion of backups is the default behavior. You can change this default behavior and choose to free the unprotected backups instead.

About this task

By default, Snap Manager deletes the backup when they expire.

Steps

1. Access the following default location:

```
default smsap installation location\properties\smsap.config
```

2. Edit the `smsap.config` file.
3. Set the `retain.alwaysFreeExpiredBackups` property in the `smsap.config` file to **true**.

For example, `retain.alwaysFreeExpiredBackups = true`

Free or delete retention policy exempt backups

Backups with the retention class of "unlimited" cannot be deleted or freed directly. To delete or free these backups, you must first assign another retention class, such as hourly, daily, weekly, or monthly. To delete or free a backup that is exempt from the retention policy, you must first update the backup to make it eligible for deletion or free it.

Steps

1. To update the backup to make it eligible for deletion by the retention policy, enter this command:

```
smsap backup update -profile profile_name {-label label [data | -archivelogs] | -id id} -retain [-hourly | -daily | -weekly | -monthly]
```

2. After updating the backup so it is eligible for deletion, you can either delete the backup or free backup resources.

- To delete the backup, enter this command:

```
smsap backup delete -profile profile_name {-label label [data | -archivelogs] | -id id | -all}
```

- To free the backup resources, rather than delete the backup, enter this command:

```
smsap backup free -profile profile_name {-label label [data | -archivelogs] | -id id | -all} [-force] [ -dump] [-quiet | -verbose]
```

View a list of backups

You can check which backups were created for a profile and the backup state by using the `smsap backup list` command. For each profile, the command displays the information about the most recent backup first and then continues until the information for all the backups is displayed.

Step

1. Enter the following command:

```
smsap backup list -profile profile_name [-delimiter character] [data | -archivelogs] [-quiet | -verbose]
```

View backup details

You can view the detailed information about a particular backup in a profile by using the `smsap backup show` command.

About this task

The `smsap backup show` command displays the following information for each backup:

- The backup ID
- Whether the backup succeeded or failed
- Backup scope (full, partial, online, or offline)
- Backup mode
- Mount status
- The backup label
- Comment
- The date and time when the operation started and ended
- Information about whether the backup was verified
- The backup retention class
- The database and host name
- The checkpoint System Change Number (SCN)
- The end backup SCN (for online backups only)
- The tablespaces and data files from the database backed up
- The control files from the database backed up
- The archive logs from the database backed up

- The storage system and volumes where the files are located
- The Snapshot copies made and their location
- The status of the primary storage resources
- The backup protection status
- Backup mode

If you specify the `-verbose` option, the following additional information is displayed:

- The clones made from the backup, if there are any
- Verification information
- If the backup is mounted, SnapManager displays the mount points in use

For the archive log file backup, the same information is displayed as that of the other database backup except for the following information:

- Checkpoint SCN
- End Backup SCN
- Tablespace
- Control files

However, archive log file backup contains the following additional information:

- The first change number of the backup
- The next change number of the backup
- Thread number
- Reset logs ID
- Incarnation
- Log file name

Step

1. Enter the following command:

```
smsap backup show -profile profile_name [-label label [data | -archivelogs] |  
-id id [-quiet | -verbose]
```

Mount backups

SnapManager automatically handles the mounting of a backup to make it available to the host. You can also mount backups in scenarios where you use an external tool, such as Oracle Recovery Manager (RMAN), to access the files in the backup.

About this task

The `smsap backup mount` command displays a list of paths where the Snapshot copies consisting of the backup have been mounted.



You can optionally collect the dump files after a successful or failed backup mount operation.

Step

1. To Mount a backup enter the following command:

```
smsap backup mount -profile profile_name label [data | -archivelogs] | -id id
[-host host] [-dump] [-quiet | -verbose]
```

Unmount backups

SnapManager automatically unmounts the backup to make it unavailable to the host server. SnapManager also allows you to unmount if you are using an external tool, such as Oracle Recovery Manager (RMAN), to access the files in the backup, and to change the state of the backup to remove access.

About this task

You can optionally collect the dump files after a successful or failed unmount backup operation.

Step

1. Enter the following command:

```
smsap backup unmount -profile profile_name {label label [data | -archivelogs]
| -id id} [-quiet | -verbose]
```

Free backups

You can free backups, which deletes the Snapshot copies without deleting the backup metadata. This function frees the space occupied by the backup. You can use the `smsap backup free` command to free the backups.

What you'll need

For a backup to be eligible for freeing, you must ensure the following:

- Backup was successful
- Backup is not to be mounted
- Backup does not have clones
- Backup is not to be retained by using an unlimited retention policy
- Backup is not already freed

About this task

You can specify the `-dump` option as an optional parameter to collect the dump files after the successful or failed backup free operation.

Step

1. Enter the following command:

```
smsap backup free -profile profile_name {-label label [data | -archivelogs] |  
-id id | -all} -force[-dump] [-quiet] [-force]
```

Delete backups

You must delete backups when you no longer need them, which frees the space those backups occupy. If you remove backups, you reduce the chance of reaching the limit of 255 Snapshot copies per volume.

What you'll need

- You must ensure that the backup was not used to create a clone.

About this task

You can delete backups retained on an unlimited basis without changing the retention class.

You can optionally collect the dump files after the successful or failed backup delete operation.

If you want to delete the archive log backups, you need to check for the retention duration set for the archive log backup. If the archive log backup is within the retention duration and the archive log files are required for recovery of a restored database, you cannot delete the archive log backup.

Steps

1. Verify that the operations are complete by entering the following command:

```
smsap operation list -profile profile_name-dump -quiet -verbose
```

2. To delete a backup, enter the following command:

```
smsap backup delete -profile profile_name [-label label [data | -archivelogs]  
| -id id | -all] [-force] [ -dump] [-quiet | -verbose]
```

Use the `-force` option to force the removal of the backup. Forcing the removal of a backup that has incomplete operations might leave the backup in an inconsistent state.

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.