NetApp

Installation and Administration for Windows

SnapManager for SAP

NetApp
November 04, 2025

Table of Contents

Installation and Administration for Windows

What SnapManagerfor SAP is

SnapManager provides the tools required to perform policy-driven data management, schedule and create regular database backups, restore data from these backups in the event of data loss or disaster, and create database clones. You can create backups on primary storage and create protected backups on secondary storage by using post-processing scripts.

SnapManager leverages NetApp technologies when integrating with the latest database releases. SnapManager is integrated with the following NetApp applications and technologies:

- SnapDrive automates storage provisioning tasks and simplifies the process of creating error-free, host-consistent Snapshot copies of the storage.
- Snapshot (a feature of Data ONTAP) creates point-in-time copies of the database.
- SnapVault (a licensed feature of Data ONTAP) leverages disk-based backups for reliable, low-overhead backup and recovery of databases.
- SnapMirror (a licensed feature of Data ONTAP) replicates database data across a global network at high speeds in a simple, reliable, and cost-effective manner.
- SnapRestore (a licensed feature of Data ONTAP) recovers an entire database in seconds, regardless of the capacity or the number of files.
- FlexClone (a licensed feature of Data ONTAP) helps to create fast, space-efficient clones of databases from the Snapshot backups.

SnapManager operates across SAN (FC and iSCSI) protocols.

What SnapManager for SAP does

SnapManager for SAP simplifies and automates database backup, recovery, and cloning by leveraging the Snapshot copies, SnapRestore, and FlexClone technologies.

SnapManager provides the following benefits to database administrators (DBAs):

- Working with Database profiles
  - You can organize and retain host and database information in profiles.

    When you initiate a backup based on a profile, you can reuse the information rather than having to reenter it for every backup. SnapManager also enables you to monitor operations quickly by using profiles.

  - In the profile, you can define the Snapshot copies naming patterns and enter custom (prefix or suffix) text, so that all the Snapshot copies can use the same naming convention that meets business policies.
  - You do not need to know the storage system name because database files are automatically mapped to the associated storage.
  - When you create a new profile, you can specify the option to separate the archive log backup from the data file backup.

You can also update the existing profile to separate the archive log backup from the data file backup.

- Performing the database backup operation
    - Backup of full and partial databases
        - You can create a full or partial backup quickly in a space-efficient way, which allows you to perform backups more frequently.

            The full database backup contains all the data files, control files, and archive log files in a single backup.

            The partial database backup contains specified data files or tablespaces, all the control files, and all the archive log files.

        - You can protect backups to secondary storage by using post-processing scripts.
        - You can schedule backups on an hourly, weekly, daily, monthly, or unlimited basis.
    - Separate back up of data files and archive log files
        - SnapManager (3.2 or later) enables you to back up the data files and archive log files separately. To perform this operation, you must specify the option to separate the archive log files while creating or updating the profile.
        - You can specify the count and duration for which the data file backups will be retained, in the retention policy.
        - You can specify the duration for the archive log file backups to be retained in archive log retention duration.
        - SnapManager (3.2 or later) also consolidates the archive log backups to a minimum number of backups by freeing the archive log backups with duplicate archive log files and retaining only the archive log backups with unique archive log files. However, this consolidation can be optionally disabled.
- Managing the archive log files
    - SnapManager (3.2 or later) enables you to prune the archive log files from the archive log destinations.

        The space occupied by the pruned archive log files is freed when the archive log backups containing these archive log files are purged.

    - SnapManager ensures that the archive log files are backed up before pruning them from the archive log destinations.

        The archive log files, which are not backed up are not pruned.

    - SnapManager ensures that the archive log files are shipped to the Data Guard Standby database while pruning archive log files from a Data Guard Primary database.
    - SnapManager ensures that the archive log files are captured by Oracle's Streams Capture process, if any.
    - Recommendation
        - To manage archive log destination space effectively, you must create the archive log backups, and prune the archive log files along with it.
    - SnapManager consolidates the archive log backups to contain minimum number of backups by freeing the archive log backups with duplicate archive log files and retaining only the archive log backups with unique archive log files.

However, this consolidation can be optionally disabled. The archive log backups, which contain duplicate archive log files are freed and a single backup with unique archive logs is retained.

- Performing the database restore operation
  - You can perform file-based restore operations.

    You can also preview restore operations and obtain a file-by-file analysis of restore operations before the operation is performed.

  - You can reduce the mean time to restore a database by using SnapRestore.
  - SnapManager (3.2 or later) enables you to recover the database automatically by using the archive log files from the backup even if the archive log files are not available in the archive log destination.

    SnapManager (3.2 or later) also provides a way to recover the database by using the archive log files from the external location to a certain extent.

- Performing database cloning for testing and development
  - You can create a clone of a database so that the database can be set up outside the production environment.

    For example, you can clone in the development and test environments for testing upgrades to vital systems.

  - You can clone a database on a primary storage system.
  - SnapManager (3.2 or later) enables you to clone the data file backups with the archive log files available in the backup.
    - You can clone the data file backups only when the archive log backup is taken along with it.
    - You can also clone the data file backups if the archive log files are available in the archive log backups made separately to a certain extent.
    - You can also clone the data file backups of a standalone database to a certain extent with archive log files from any external location accessible by Oracle.
    - If the backups are available from an external location, you can specify the external location during cloning for recovering the cloned database to a consistent state.
  - Cloning of the archive log-only backups is not supported.
- General
  - Integrate with SAP's BR*Tools.

    The BR*Tools package provides SAP tools such as BRARCHIVE, BRBACKUP, BRCONNECT, BRRECOVER, BRRESTORE, BRSPACE, and BRTOOLS.

SnapManager provides the following benefits to storage administrators:

- Supports different SAN protocols.
- Enables you to optimize backups based on the type of backup (full or partial) that works best in your environment.
- Creates space-efficient database backups.
- Creates space-efficient clones.

SnapManager also works with the following Oracle features:

- SnapManager can catalog its backups with Oracle's RMAN.

  If using RMAN, a DBA can make use of SnapManager backups and preserve the value of all RMAN functions, such as block-level restore. SnapManager lets RMAN use the Snapshot copies when it performs recovery or restore. For example, you can use RMAN to restore a table within a tablespace and to perform full database and tablespace restores and recoveries from Snapshot copies made by SnapManager. The RMAN recovery catalog should not be in the database that is being backed up.

Integration with other NetApp applications and technologies

SnapManager for SAP is a stand-alone product that integrates the features from other NetApp products to enable fast backups that require only a small amount of space.

SnapManager integrates with the following NetApp applications and technologies:

Applications and technologies	Description
SnapDrive	SnapManager uses SnapDrive to create Snapshot copies of the storage. Snapshot copies ensure that backups are space-efficient and faster to create than the disk-to-disk backups.
FlexClone (a licensed feature of Data ONTAP)	SnapManager uses the FlexClone feature to create fast, space-efficient clones of backups.
Snapshot (a feature of Data ONTAP)	Snapshot technology creates point-in-time copies of the database.
SnapRestore (a licensed feature of Data ONTAP)	SnapManager reduces the mean time to recover a database by using SnapRestore. SnapRestore can recover individual files to a multi-terabyte volume so that operations can resume quickly.
SnapVault (a licensed feature of Data ONTAP)	SnapVault leverages disk-based backups for reliable, low-overhead backup and recovery of databases.
SnapMirror (a licensed feature of Data ONTAP)	SnapMirror replicates database data across a global network at high speeds in a simple, reliable, and cost-effective manner.

Advantages of using SnapManager

You can use SnapManager for SAP to perform different tasks on the databases and manage data efficiently.

SnapManager for SAP works with storage systems and enables you to perform the following tasks:

- Create space-efficient backups to the primary or secondary storage and schedule backups.

You can create full and partial database backups and apply retention duration policies.SnapManager (3.2 or later) enables you to create backups of only the data files and archive logs.

- ◦ SnapManager (3.2 or later) enables you to perform preprocessing or post-processing before or after the backup and restore operations.
- ◦ SnapManager (3.2 or later) enables you to protect backups by using the postprocessing scripts.
- Restore full or partial databases by using the file-based restore operation.
- Restore and recover database backups automatically.

SnapManager (3.2 or later) enables the restoration and recovery of database backups automatically. SnapManager automatically recovers the restored database by discovering, mounting, and applying the archive log files from the backups.

- Prune archive log files from the archive log destinations when creating backups for only the archive logs.
- Retain the minimum number of archive log backups automatically by retaining only the backups with unique archive log files.
- Track operation details and produce reports by host, profile, backup, or clone.
- Verify the backup status.
- Maintain the history of SnapManager operations associated with a profile.
- Create space-efficient clones of backups on the primary storage.

Create backups using Snapshot copies

SnapManager enables you to create backups on the primary (local) storage and also on the secondary (remote) storage using postprocessing scripts.

Backups created as Snapshot copies are virtual copies of the database and are stored in the same physical medium as the database. Therefore, the backup operation takes less time and requires significantly less space than full, disk-to-disk backups. SnapManager enables you to back up the following:

- All the data files, archive log files, and control files
- Selected data files or tablespaces, all the archive log files, and control files

SnapManager 3.2 or later enables you to optionally back up the following:

- All the data files and the control files
- Selected data files or tablespaces along with the control files
- Archive log files

> The data files, archive log files, and control files can be located on different storage systems, storage system volumes, or logical unit numbers (LUNs). You can also use SnapManager to back up a database when there are multiple databases on the same volume or LUN.

Why you should prune archive log files

SnapManager for SAP enables you to delete archive log files from the active file system that are already backed up.

Pruning enables SnapManager to create backups of distinct archive log files. Pruning, along with the backup retention policy, frees archive log space when backups are purged.

> ℹ️ 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.

Archive log consolidation

SnapManager (3.2 or later) for SAP consolidates the archive log backups to maintain a minimum number of backups for archive log files. SnapManager for SAP identifies and frees the backups that contain archive logs files that are subsets of other backups.

Full or partial restoration of databases

SnapManager provides the flexibility to restore full databases, specific tablespaces, files, control files, or a combination of these entities. SnapManager enables you to restore data by using a file-based restore process.

SnapManager enables database administrators (DBAs) to preview restore operations. The preview feature enables DBAs to view each restore operation on a file-by-file basis.

DBAs can specify the level to which SnapManager restores and recovers information when performing restore operations. For example, DBAs can restore and recover data to specific points in time. The restore point can be a date and time or an Oracle System Change Number (SCN).

SnapManager (3.2 or later) enables you to restore and recover database backups automatically without DBA intervention. You can use SnapManager to create archive log backups, and then use those archive log backups to restore and recover the database backups. Even if the backup's archive log files are managed in an external archive log location, you can specify that external location so those archive logs can help recover the restored database.

Verify backup status

SnapManager can confirm the integrity of the backup using standard Oracle backup verification operations.

Database administrators (DBAs) can perform the verification as part of the backup operation, or at another time. DBAs can set the verify operation to occur during an off-peak time when the load on the host servers is less, or during a scheduled maintenance window.

Database backup clones

SnapManager uses the FlexClone technology to create a writable, space-efficient clone of a database backup. You can modify a clone without changing the backup source.

You might want to clone databases to enable testing or upgrades in nonproduction environments. You can clone a database residing on primary. A clone can be located on the same host or on a different host as the database.

FlexClone technology enables SnapManager to use Snapshot copies of the database to avoid creating an entire physical, disk-to-disk copy. Snapshot copies require less creation time and take up significantly less

space than physical copies.

See the Data ONTAP documentation for more information about FlexClone technology.

Related information

Data ONTAP documentation

Track details and produce reports

SnapManager reduces the level of detail database administrators need to track the status of different operations by offering methods to monitor operations from a single interface.

After administrators specify which databases should be backed up, SnapManager automatically identifies the database files for backup. SnapManager displays information about repositories, hosts, profiles, backups, and clones. You can monitor the operations on specific hosts or databases.

What the SnapManager for SAP architecture is

The SnapManager for SAP architecture includes many components, such as the SnapManager for SAP host, client, and repository. Other components include the primary and secondary storage systems and other NetApp products.

The SnapManager for SAP architecture includes the following architectural components:

- SnapManager host
- SnapManager graphical user interface or command-line interface
- SnapManager repository
- SnapManager for SAP BACKINT interface
- Primary storage system
- Secondary storage systems
- SnapDrive for Windows

The following image shows the architecture of SnapManager for SAP and related components:

SnapManager for SAP Architecture
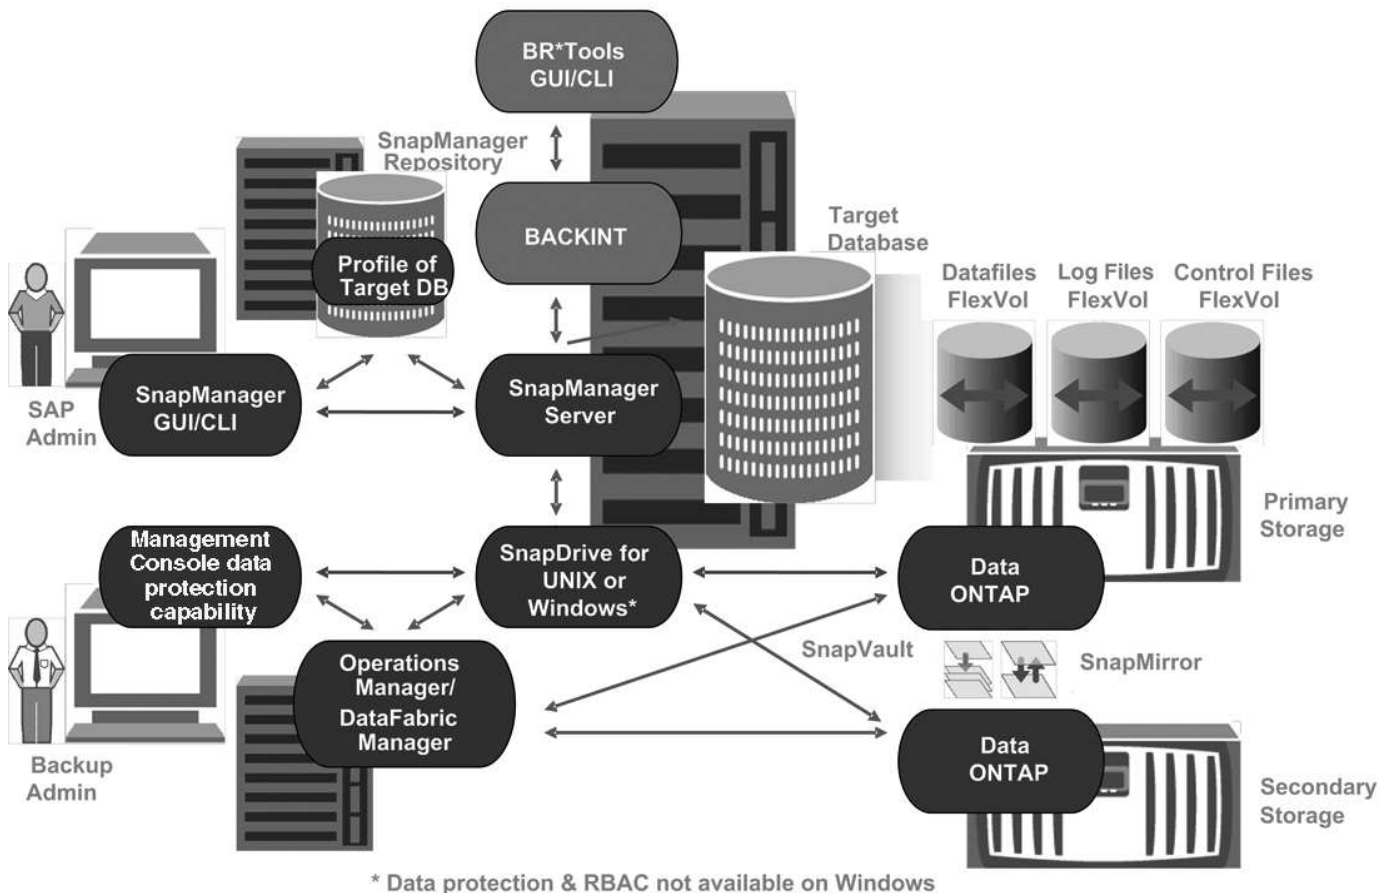
SnapManager host

The SnapManager host is a Windows server, which also runs other NetApp products.

The SnapManager host is installed with the following products:

- SnapDrive for Windows
- Host Utilities

The SnapManager host runs as a service.

The SnapManager host also supports the BACKINT interface, which is used for SAP BR*Tools.

SnapManager graphical user and command-line interfaces

The SnapManager client includes both a graphical user interface (GUI) and a command-line interface (CLI).

SnapManager repository

The repository stores information related to different SnapManager operations, for example, the time of backups, tablespaces and data files backed up, storage systems used, clones made, and Snapshot copies created.

The repository database cannot exist in the same database and also cannot be a part of the same database that SnapManager is backing up. This is because the repository stores the names of the database Snapshot copies created during the backup operations. The repository must be created in a different database than the database that is being backed up. This means that you must have at least two databases: the SnapManager repository database and the target database managed by SnapManager. When you run the SnapManager services, both the databases must be up and running.

> (i) You must not perform any SnapManager operations by using the GUI or CLI when the
> repository database is down.

SnapDrive on SnapManager server

SnapManager uses SnapDrive for Windows to create Snapshot copies of the storage system. SnapDrive resides on the same server as SnapManager.

What repositories are

SnapManager organizes information into profiles, which are then associated with repositories. Profiles contain information about the database that is being managed, while the repository contains data about the operations that are performed on profiles.

The repository records when a backup took place, which files were backed up, and whether a clone was created from the backup. When database administrators restore a database or recover a portion of it, SnapManager queries the repository to determine what was backed up.

Because the repository stores the names of the database Snapshot copies created during backup operations, the repository database cannot exist in the same database and also cannot be a part of the same database that SnapManager is backing up. You must have at least two databases (the SnapManager repository database and the target database being managed by SnapManager) up and running when you execute SnapManager operations.

If you try to open the graphical user interface (GUI) when the repository database is down, the following error message is logged in the `sm_gui.log` file: [WARN]: `SMSAP-01106: Error occurred while querying the repository: No more data to read from socket`. Also, SnapManager operations fail when the repository database is down. For more information about the different error messages, see *Troubleshooting known issues*.

You can use any valid host name, service name, or user name to perform operations. For a repository to support SnapManager operations, the repository user name and service name must consist of only the following characters: alphabetic characters (A-Z), digits (0-9), minus sign (-), underscore (_), and period (.).

The repository port can be any valid port number and the repository host name can be any valid host name. The host name must consist of alphabetic characters (A-Z), digits (0-9), minus sign (-), and period (.), but not an underscore (_).

The repository must be created in an Oracle database. The database that SnapManager uses should be set up in accordance with Oracle procedures for database configuration.

A single repository can contain information about multiple profiles; however, each database is normally associated with only one profile. You can have multiple repositories, with each repository containing multiple profiles.

What profiles are

SnapManager uses profiles to store the information necessary to perform operations on a given database. A profile contains the information about the database including its credentials, backups, and clones. By creating a profile, you do not have to specify database details each time you perform an operation on that database.

A profile can reference only one database. The same database can be referenced by more than one profile. Backups created using one profile cannot be accessed from a different profile, even if both the profiles reference the same database.

Profile information is stored in a repository. The repository contains both the profile information for the database and information about the Snapshot copies that serve as the database backup. The actual Snapshot copies are stored on the storage system. The Snapshot copy names are stored in the repository containing the profile for that database. When you perform an operation on a database, you must select the profile from the repository.

The following figure illustrates how repositories can hold multiple profiles, but also that each profile can define only one database:
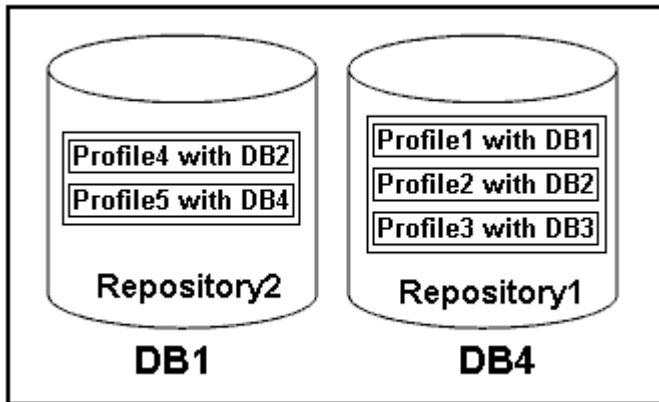
In the preceding example, Repository2 is on database DB1 and Repository1 is on the database DB4.

Each profile contains the credentials for the database associated with the profile. The credentials enable SnapManager to connect to and work with the database. The stored credentials include the user name and password pairs for accessing the host, the repository, the database, and the required connection information if you are using Oracle Recovery Manager (RMAN).

You cannot access a backup that was created using one profile from a different profile, even if both the profiles are associated with the same database. SnapManager places a lock on the database to prevent two incompatible operations from being performed simultaneously.

Profile for creating full and partial backups

You can create profiles to take full backups or partial backups.

The profiles that you specify to create the full and partial backups contain both the data files and archive log files. SnapManager does not allow such profiles to separate the archive log backups from the data file backups. The full and partial backups are retained based on the existing backup retention policies. You can schedule full and partial backups based on the time and frequency that suits you.

Profiles for creating data files-only backups and archive log-only backups

SnapManager (3.2 or later) allows you to create profiles that take backups of the archive log files separately from the data files. After you use the profile to separate the backup types, you can create either data files-only backups or archive log-only backups of the database. You can also create a backup containing both the data files and archive log files together.

The retention policy applies to all the database backups when the archive log backups are not separated. After you separate the archive log backups, SnapManager allows you to specify different retention durations .

Retention policy

SnapManager determines whether a backup should be retained by considering both the retention count (for example, 15 backups) and the retention duration (for example, 10 days of daily backups). A backup expires when its age exceeds the retention duration set for its retention class and the number of backups exceeds the retention count. For example, if the backup count is 15 (meaning that SnapManager has taken 15 successful backups) and the duration requirement is set for 10 days of daily backups, the five oldest, successful, and eligible backups expire.

Archive log retention duration

After the archive log backups are separated, they are retained based on the archive log retention duration. Archive log backups taken with data file backups are always retained along with those data file backups, regardless of the archive log retention duration.

What SnapManager operation states are

SnapManager operations (backup, restore, and clone) can be in different states, with each state indicating the progress of the operation.

Operation state	Description
Succeeded	The operation completed successfully.
Running	The operation has started, but is not finished. For instance, a backup, which takes two minutes, is scheduled to occur at 11:00 a.m.. When you view the **Schedule** tab at 11:01 a.m., the operation appears as Running.
No operation found	The schedule has not run or the last run backup was deleted.
Failed	The operation failed. SnapManager has automatically executed the abort process and cleaned the operation.

Recoverable and unrecoverable events

A recoverable SnapManager event has the following problems:

• The database is not stored on a storage system that runs Data ONTAP.

• SnapDrive for Windows is not installed or cannot access the storage system.

- SnapManager fails to create a Snapshot copy or provision storage if the volume is out of space, the maximum number of Snapshot copies has been reached, or an unanticipated exception occurs.

When a recoverable event occurs, SnapManager performs an abort process and attempts to return the host, database, and storage system to the starting state. If the abort process fails, SnapManager treats the incident as an unrecoverable event.

An unrecoverable (out-of-band) event occurs when any of the following happens:

- A system issue occurs, such as when a host fails.
- The SnapManager process is stopped.
- An in-band abort operation fails when the storage system fails, the logical unit number (LUN) or storage volume is offline, or the network fails.

When an unrecoverable event occurs, SnapManager performs an abort process immediately. The host, database, and storage system might not have returned to the initial states. If this is the case, you must perform a cleanup after the SnapManager operation fails by deleting the orphaned Snapshot copy and removing the SnapManager lock file.

If you want to delete the SnapManager lock file, navigate to `$ORACLE_HOME` on the target machine and delete the `sm_lock_TargetDBName` file. After deleting the file, you must restart the SnapManager for SAP server.

How SnapManager maintains security

You can perform SnapManager operations only if you have the appropriate credentials. Security in SnapManager is governed by user authentication.

SnapManager maintains security by requesting user authentication through password prompts or by setting user credentials. An effective user is authenticated and authorized with the SnapManager server.

SnapManager credentials and user authentication differ significantly from SnapManager 3.0:

- In SnapManager versions earlier than 3.0, you would set an arbitrary server password when you install SnapManager. Anyone who wants to use the SnapManager server would need the SnapManager server password. The SnapManager server password would need to be added to the user credentials by using the `smsap credential set -host` command.
- In SnapManager (3.0 and later), the SnapManager server password has been replaced by individual user operating system (OS) authentication. If you are not running the client from the same server as the host, the SnapManager server performs the authentication by using your OS user names and passwords. If you do not want to be prompted for your OS passwords, you can save the data to your SnapManager user credentials cache by using the `smsap credential set -host` command.

> (i) The `smsap credential set -host` command remembers your credentials when the `host.credentials.persist` property in the `smsap.config` file is set to **true**.

Example

User1 and User2 share a profile called Prof2. User2 cannot perform a backup of Database1 in Host1 without permission to access Host1. User1 cannot clone a database to Host3 without permission to access Host3.

The following table describes different permissions assigned to the users:

Permission type	User1	User2
Host Password	Host1, Host2	Host2, Host3
Repository Password	Repo1	Repo1
Profile Password	Prof1, Prof2	Prof2

In the case where User1 and User2 do not have any shared profiles, assume User1 has permissions for the hosts named Host1 and Host2, and User2 has permissions for the host named Host2. User2 cannot run even the nonprofile commands such as `dump` and `system verify` on Host1.

Access and print online Help

The online Help provides instructions for the tasks that you can perform using the SnapManager graphical user interface. The online Help also provides descriptions of fields on the windows and wizards.

Steps

1. Perform one of the following actions:
   - In the main window, click **Help** > **Help Contents**.
   - In any window or wizard, click **Help** to display help specific to that window.
2. Use the **Table of Contents** in the left pane to navigate through the topics.
3. Click the Printer icon at the top of the help window to print individual topics.

SnapManager for SAP deployment considerations

Before deploying SnapManager in your environment, you should know the other applications and technologies required for different operations.

The following table lists the different applications and technologies:

Applications and technologies	Details
Data ONTAP	SnapManager leverages NetApp tools and technologies, including Snapshot copies.
SnapDrive for Windows	SnapManager uses the SnapDrive features. You must install SnapDrive before running the SnapManager services. SnapManager handles all the interactions with SnapDrive. SnapDrive for Windows must be configured correctly for your storage system and protocol choices.
SnapRestore	SnapManager reduces the mean time to recover a database by using SnapRestore. Each storage system should have a SnapRestore license.

Applications and technologies	Details
FlexClone	FlexClone is a licensed feature in Data ONTAP.
FC and iSCSI protocols	You need to have the licensed versions of the appropriate protocols.

To use BR*Tools commands, the appropriate version of SAP BR*Tools must be installed for your version of Oracle.

Related information

SnapManager for SAP Best Practices: media.netapp.com/documents/tr-3823.pdf

Requirements for running SnapManager

Before deploying SnapManager in your environment, you should know the different requirements.

Before using SnapManager, you must review the compatibility matrices for all the required products. You must also review the following:

- SnapManager and SnapDrive compatibility matrix in the interoperability section for the latest version and patch information for all hosts, storage systems, and other components listed in this section.

- The Configuration Guide for NetApp FCP and iSCSI products.

> ⓘ SnapManager requires specific Oracle versions on some platforms.

See the documentation kit for more information about the recommended configurations for the host and storage systems.

> ⓘ Contact your sales representative if you need a SnapManager configuration that is not mentioned in the documentation kit.

Related information

Interoperability Matrix: support.netapp.com/NOW/products/interoperability

Supported host hardware

Consider the memory, disk space, and CPU requirements.

SnapManager requires the following configuration:

Hardware function	Hardware requirements
Memory	The SnapManager server requires 128 MB of memory. The graphical user interfaces requires a minimum 512 MB RAM to run. Each operation run by the SnapManager server requires 48 MB of additional memory while it is running.
Disk space	128 MB available disk space for the graphical user application (minimum).
CPU speed	1.0 GHz processor speed (minimum).

Supported general configurations

Before installing SnapManager, you must know the general configuration requirements.

SnapManager supports the following general configurations:

- A non-clustered configuration where a single host is connected to a single storage system
- One SnapManager server instance per host

For information about all storage types and versions supported by SnapManager, see the SnapManager and SnapDrive Compatibility Matrix.

Clustered configurations

SnapManager operates in cluster configurations.

SnapManager supports the same host cluster and configurations that the SnapDrive product and Host Utilities Kit support.

SnapManager also supports non-clustered configurations where a single host is connected to a single storage system, supported host clusters, and storage systems running Data ONTAP controller failover.

Database version support and configuration overview

You must know the different database versions and configurations supported with SnapManager. You must perform basic database layout and configuration setup to ensure successful operations.

SnapManager for SAP integrates with Oracle versions 10gR2 (10.2.0.5), 11gR2 (11.2.0.1 and 11.2.0.2), and 12*c*; with native Oracle technology such as Recovery Manager (RMAN); and across fibre channel (FC) and Internet Small Computer System Interface (iSCSI) protocols.

> (i) The Oracle Database 9i is not supported from SnapManager 3.2 and 10g R2 (earlier than 10.2.0.5) is not supported from SnapManager 3.3.1.

General layout and configuration

You can find information about the recommended general database layouts and storage configurations to avoid issues related to disk groups, file types, and tablespaces.

- Do not include files from more than one type of SAN file system in your database.

  All files making up a database must reside on the same type of file system.

- SnapManager requires a multiple of 4K block size.

The following are some guidelines for volume separation:

- Data files for only one database must be in the volume.
- You must use separate volumes for each of the following file classifications: database binaries, data files, online redo log files, archived redo log files, and control files.
- You do not need to create a separate volume for temporary database files, because SnapManager does not back up temporary database files.

SAP uses a standard layout for Oracle database installations. In this layout, SAP places copies of the Oracle control file in `E:\oracle\SID\origlogA, E:\oracle\SID\origlogB, and E:\oracle\SID\sapdata1 file systems`.

> ⓘ Because BR*Tools backups include the Oracle and SAP profiles located in the database subdirectory of the Oracle installation, the Oracle installation must be located on storage.

In the case of a new installation, you can modify the location of the control files by using SAPINST and move the control file normally placed in the sapdata1 file system to a file system that does not reside in the same volume as the data files. (SAPINST is the tool that SAP provides for installing SAP systems.)

For more information, see the *SnapManager for SAP Best Practices*.

Related information

[SnapManager for SAP Best Practices: media.netapp.com/documents/tr-3823.pdf](media.netapp.com/documents/tr-3823.pdf)

Sample database volume layouts

You can refer to sample database volume layouts for help in configuring your database.

Single-instance databases

File types	Volume names	Dedicated volume for file types	Automatic Snapshot copies
Oracle binaries	orabin_ *host name*	Yes	On
Data files	oradata_ *sid*	Yes	Off
Temporary data files	oratemp_ *sid*	Yes	Off

File types	Volume names	Dedicated volume for file types	Automatic Snapshot copies
Control files	oracntrl01_ *sid* (Multiplexed) oracntrl02_ *sid* (Multiplexed)	Yes	Off
Redo logs	oralog01_ *sid* (Multiplexed) oralog02_ *sid* (Multiplexed)	Yes	Off
Archive logs	oraarch_ *sid*	Yes	Off

Limitations when working with SnapManager

You must be aware of the scenarios and limitations that might affect your environment.

Limitations related to database layouts and platforms

- SnapManager supports control files on a file system and does not support control files on raw devices.
- SnapManager operates in a Microsoft clustering (MSCS) environment but does not recognize the state of the MSCS configuration (active or passive) and does not transfer active management of a repository to a standby server in an MSCS cluster.
- The repository database might exist on a host that can be accessed by using more than one IP address.

  If the repository is accessed by using more than one IP address, then the schedule file is created for each of the IP addresses. If the schedule backup is created for a profile (for example, profile A) under one of the IP addresses (for example, IP1), then the schedule file for only that IP address gets updated. If profile A is accessed from another IP address (for example, IP2), the scheduled backup is not listed because the schedule file of IP2 does not have an entry for the schedule that was created under IP1.

  You can wait for the schedule to be triggered from that IP address and the schedule file to be updated, or you can restart the server.

Limitations related to SnapManager configuration

- SnapManager supports databases on MultiStore storage systems with the following requirements:
  - You must configure SnapDrive to set passwords for MultiStore storage systems.
  - SnapDrive cannot create a Snapshot copy of a LUN or file residing in a qtree in a MultiStore storage system if the underlying volume is not in the same MultiStore storage system.
- SnapManager does not support accessing two SnapManager servers running on different ports from a single client (both from the CLI or GUI).

  The port numbers should be the same on the target and remote hosts.

- SnapManager operations fail and you cannot access the GUI when the repository database is down.

  You must verify that the repository database is running when you perform any SnapManager operations.

- SnapManager does not support Live Partition Mobility (LPM) and Live Application Mobility (LAM).
- SnapManager does not support Oracle Wallet Manager and Transparent Data Encryption (TDE).
- SnapManager does not support MetroCluster configurations in raw device mapping (RDM) environments because MetroCluster configurations are yet to be supported by Virtual Storage Console (VSC).

Limitations related to profile management

- If you update the profile to separate the archive log backups, then you cannot perform a rollback operation on the host.
- If you enable a profile from the GUI to create archive log backups, and later try to update the profile by using the Multi Profile Update window or Profile Update window, then you cannot modify that profile to create a full backup.
- If you update multiple profiles in the Multi Profile Update window and some profiles have the **Backup Archivelogs separately** option enabled and other profiles have the option disabled, then the **Backup Archivelogs separately** option is disabled.
- If you update multiple profiles and some profiles have the **Backup Archivelogs separately** option enabled and other profiles have the option disabled, then the **Backup Archivelogs separately** option in the Multi Profile Update window is disabled.
- If you rename the profile, then you cannot roll back the host.

Limitations related to rolling upgrade or rollback operations

- If you try to install an earlier version of SnapManager for a host without performing the rollback operation on the host in the repository, you might not be able to do the following:
  - View the profiles that were created in earlier or later versions of SnapManager for the host.
  - Access backups or clones that were created in earlier or later versions of SnapManager.
  - Perform rolling upgrade or rollback operations on the host.
- After you separate the profiles to create archive log backups, you cannot perform a rollback operation on the related host repository.

Limitations related to backup operations

- During recovery, if the backup is already mounted, SnapManager does not mount the backup again and uses the already mounted backup.

  If the backup is mounted by a different user and you do not have access to the previously mounted backup, then the other user must provide you the permission.

  All archive log files have read permission for users assigned to a group; you might not have the access permission to the archive log file, if the backup is mounted by a different user group. Users can give permission to the mounted archive log files manually, and then retry the restore or recovery operation.

- SnapManager sets the backup state as "PROTECTED", even when one of the Snapshot copies of the database backup is transferred to the secondary storage system.
- You can use the task specification file for scheduled backup only from SnapManager 3.2 or later.

- If the repository database is pointing to more than one IP address and each IP address has a different host name, then the backup scheduling operation is successful for one IP address but fails for the other IP address.

- In ONTAP environments, SnapManager does not support multiple secondary destinations for a source volume.

Limitations related to restore operations

- SnapManager does not support restoring database backups by using fast restore or volume-based restore operations in Windows.

  SnapManager 3.3 displays an error message when you try to perform a fast restore operation from the SnapManager CLI.

Limitations related to clone operations

- SnapManager 3.3 or later does not support using the clone specification XML file created in the releases before SnapManager 3.2.

- If temporary tablespaces are located in a different location from the datafiles location, a clone operation creates the tablespaces in the datafiles location.

  However, if temporary tablespaces are Oracle Managed Files (OMFs) that are located in a different location from the datafiles location, the clone operation does not create the tablespaces in the datafiles location. The OMFs are not managed by SnapManager.

- SnapManager fails to clone a RAC database if you select the `-resetlogs` option.

Limitations related to archive log files and backups

- SnapManager does not support pruning of archive log files from the flash recovery area destination.

- SnapManager does not support pruning of archive log files from the standby destination.

- The archive log backups are retained based on the retention duration and default hourly retention class.

  When the archive log backup retention class is modified by using the SnapManager CLI or GUI, the modified retention class is not considered for backup because archive log backups are retained based on retention duration.

- If you delete the archive log files from the archive log destinations, the archive log backup does not include archive log files older than the missing archive log file.

  If the latest archive log file is missing, then the archive log backup operation fails.

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

- SnapManager consolidates the archive log backups even when you delete the archive log files from the archive log destinations or when the archive log files are corrupted.

Limitations related to changing of target database host name

The following SnapManager operations are not supported when you change the target database host name:

- Changing the target database host name from the SnapManager GUI.

- Rolling back of the repository database after updating the target database host name of the profile.

- Simultaneously updating multiple profiles for a new target database host name.
- Changing the target database host name when any SnapManager operation is running.

Limitations related to the SnapManager CLI or GUI

- The SnapManager CLI commands for the `profile create` operation that are generated from the SnapManager GUI do not have history configuration options.

  You cannot use the `profile create` command to configure history retention settings from the SnapManager CLI.

- SnapManager does not display the GUI in Mozilla Firefox when there is no Java Runtime Environment (JRE) available on the Windows client.

- SnapManager 3.3 does not display the SnapManager GUI in Microsoft Internet Explorer 6 on Windows Server 2008 and Windows 7.

- While updating the target database host name using the SnapManager CLI, if there are one or more open SnapManager GUI sessions, then all of the open SnapManager GUI sessions fail to respond.

- When you install SnapManager on Windows and start the CLI in UNIX, the features that are not supported on Windows are displayed.

Limitations related to SnapMirror and SnapVault

- In some scenarios, you cannot delete the last backup associated with the first Snapshot copy when the volume has a SnapVault relationship established.

  You can delete the backup only when you break the relationship. This issue is because of an ONTAP restriction with base Snapshot copies. In a SnapMirror relationship the base Snapshot copy is created by the SnapMirror engine, and in a SnapVault relationship the base Snapshot copy is the backup created by using SnapManager. For each update, the base Snapshot copy points to the latest backup created by using SnapManager.

Limitations related to Data Guard Standby databases

- SnapManager does not support Logical Data Guard Standby databases.
- SnapManager does not support Active Data Guard Standby databases.
- SnapManager does not allow online backups of Data Guard Standby databases.
- SnapManager does not allow partial backups of Data Guard Standby databases.
- SnapManager does not allow restoring of Data Guard Standby databases.
- SnapManager does not allow pruning of archive log files for Data Guard Standby databases.
- SnapManager does not support Data Guard Broker.

Related information

Documentation on the NetApp Support Site

SnapManager features and Oracle technologies not supported on Windows

SnapManager does not support some SnapManager features, platforms, and Oracle technologies on Windows.

SnapManager does not support the following features, platforms, and Oracle technologies:

- Policy-based data protection available with the integration of Protection Manager
- Role-based access control (RBAC) available with Operations Manager
- Fast restore or volume-based SnapRestore (VBSR)
- Clone split operation
- Oracle Real Application Clusters (RAC) using any protocol
- Oracle Automatic Storage Management (ASM) using any protocol
- Oracle Direct NFS (dNFS)
- Itanium-64 platforms

> (i) The supported hardware platforms for the Windows Operating System are 32-bit and 64-bit (Windows x86 and Windows x86_64).

SnapManager limitations for clustered Data ONTAP

You must know the limitations for some functionalities and SnapManager operations if you are using clustered Data ONTAP.

The following functionalities are not supported if you are using SnapManager on clustered Data ONTAP:

- Raw device mapping (RDM) logical unit number (LUN) for storage virtual machines (SVM)
- A database in which one LUN belongs to a system running Data ONTAP operating in 7-Mode and the other LUN belongs to a system running clustered Data ONTAP
- SnapManager for SAP does not support migration of a Vserver, which is not supported by clustered Data ONTAP
- SnapManager for SAP does not support the clustered Data ONTAP 8.2.1 functionality to specify different export policies for volumes and qtrees

Limitations related to Oracle Database

Before you start working with SnapManager, you must know the limitations related to Oracle Database.

The limitations are as follows:

- SnapManager supports Oracle version 10gR2 and does not support Oracle 10gR1 as the repository or target database.
- SnapManager does not support Oracle Cluster File System (OCFS).
- Support for Oracle Database 9i is deprecated from SnapManager 3.2.
- Support for Oracle Database 10gR2 (earlier than 10.2.0.5) is deprecated from SnapManager 3.3.1.

> (i) Identify the different versions of Oracle databases supported by referring to the Interoperability Matrix.

Related information

Deprecated versions of Oracle database

Oracle database 9i is not supported by SnapManager 3.2 or later, and Oracle database 10gR2 (earlier than 10.2.0.4) is not supported by SnapManager 3.3.1 or later.

If you are using Oracle 9i or 10gR2 (earlier than 10.2.0.4) databases and want to upgrade to SnapManager 3.2 or later, you cannot create new profiles; a warning message is displayed.

If you are using Oracle 9i or 10gR2 (earlier than 10.2.0.4) databases and want to upgrade to SnapManager 3.2 or later, you must perform one of the following:

- Upgrade Oracle 9i or 10gR2 (earlier than 10.2.0.4) databases to either Oracle 10gR2 (10.2.0.5), 11gR1, or 11gR2 databases, and then upgrade to SnapManager 3.2 or 3.3.

  If you are upgrading to Oracle 12*c*, then you must upgrade to SnapManager 3.3.1 or later.

  > ℹ️ Oracle database 12*c* is supported only from SnapManager 3.3.1.

- Manage the Oracle 9i databases using a patch version of SnapManager 3.1.

  You can use SnapManager 3.2 or 3.3 if you want to manage Oracle 10gR2, 11gR1, or 11gR2 databases and use SnapManager 3.3.1 or later if you want to manage Oracle 12*c* databases along with the other supported databases.

Installing SnapManager for SAP

You can download and install SnapManager for SAP in your environment and perform operations such as database backup, restore, recovery, and cloning.

The SnapManager for SAP installation package includes the host server software and the graphical user interface (GUI) client software.

Preparing to install SnapManager for SAP

The environment in which you are installing SnapManager for SAP must meet certain software, hardware, browser, database, and operating system requirements. For the latest information about the requirements, see the Interoperability Matrix.

Preinstallation tasks

Before installing SnapManager for SAP, you must perform some additional tasks to set up your environment. The tasks that you have to perform depend on the operating system and the database version you want to use.

- Install licensed operating system with the appropriate patches.
- Set the operating system's and the Oracle database's languages to English.

For example, to set the language of the Oracle database to English assign `NLS_LANG = AMERICAN_AMERICA.WE8MSWIN1252`. For more information about how to set the language, see the *Troubleshooting SnapManager for SAP* section.

- Install Data ONTAP with licenses enabled for SnapRestore and the supported protocols such as Fibre Channel (FC) and Internet Small Computer System Interface (iSCSI) on all the storage systems.
- Install the following Oracle patches, if you are using Oracle databases 11.2.0.2 and 11.2.0.3:
  - 13413167 for Windows 32-bit
  - 13555974 for Windows 64-bit
- Install SAP BR* Tools.

Related information

[NetApp Interoperability Matrix](#)

[SnapManager for SAP Best Practices](#)

Download the SnapManager for SAP installation package

You can download the SnapManager for SAP installation package from the NetApp Support Site.

Steps

1. Log in to the NetApp Support Site.
2. On the the NetApp Support Site page, click **Downloads** > **Software**.
3. In the Software Download table, go to the SnapManager product row and select **SAP (Windows)** from the **Select Platform** drop-down list.
4. Click **Go**.

   The SnapManager releases are listed.

5. Click **View & Download** for the SnapManager release that you want to install.

   The description page is displayed.

   > ⓘ You should read the information available on this page.

6. At the bottom of this page, click **Continue**.

   The license agreement page is displayed.

   > ⓘ You should read the information available on this page.

7. Click **Accept**.

   The download page is displayed.

8. Download the installation package for your host.

   You can download the installation file to any location on the machine on which you want to install

SnapManager for SAP.

Related information

Install SnapManager for SAP

You can install SnapManager on the host that has one or more databases to be managed. You can install only one SnapManager instance per host.

What you'll need

- You must have completed the required preinstallation tasks.

- The latest latest SnapManager installation package must be downloaded.

- The appropriate version of SnapDrive for Windows must be installed and configured on all the target hosts.

  For information about installing and configuring SnapDrive for Windows, see *SnapDrive for Windows Installation and Administration Guide*.

Steps

1. Double-click the SnapManager installation file:

If the operating system is…	Then use…
Windows x86	`netapp.smsap.windows-x86-version.exe`
Windows x64	`netapp.smsap.windows-x64-version.exe`

   The following message is displayed: `The publisher could not be verified. Are you sure you want to run this software?`

2. Click **OK**.

3. In the **Introduction** window, click **Next**.

4. In the **Choose Install Folder** window, either click **Next** to accept the default installation location or select a new location.

   The default location is: `C:\Program Files\NetApp\SnapManager for SAP`.

5. In the **Menu Availability** window, click **Next**.

6. In the **Specify Service Properties** window, enter the account and password information for Windows service.

   The specified account must be a member of the following groups:

   - The storage system's local administration group

   - The local administrator's group

   - The ORA_DBA group You can specify whether the service must start automatically after reboot or

should be started manually.

7. In the **Pre-Installation Summary** window, click **Install**.

8. In the **Install Complete** window, click **Next**.

9. In the **Important Information** window, click **Done** to exit the installer.

After you finish

After completing the installation, you can verify that the installation was successful:

1. Start the SnapManager server:

   a. In the Windows Services window, select **NetApp SnapManager** *version* **for SAP**.

   b. In the left panel, click **Start**.

2. Verify that the SnapManager system is running correctly:

   a. Click **Start** > **Programs** > **NetApp** > **SnapManager for SAP** > **Start SMSAP Command Line Interface (CLI)**.

   b. In the command-line interface (CLI), enter the following command:
      `smsap system verify`

      The following message is displayed: `Operation Id number succeeded.`

      number is the operation ID number.

Related information

[Documentation on the NetApp Support Site](#)

Integrate with SAP BR* Tools

The SAP BR* Tools that contains SAP tools for Oracle database administration, for example, BRARCHIVE, BRBACKUP, BRCONNECT, BRRECOVER, BRRESTORE, BRSPACE, and BRTOOLS uses the BACKINT interface provided by SnapManager for SAP. To integrate SAP BR* Tools, you must create a link from the BR* Tools directory to `<SnapManager for SAP installation directory>\NetApp\SnapManager for SAP\bin`, where the BACKINT file is installed.

What you'll need

- You must ensure that you have installed SAP BR* Tools.

Steps

1. Create a link from the BR*Tools directory to the `C:\Program Files\NetApp\SnapManager for SAP\bin\backint` file for each SAP instance.

   > (i) You must use the link instead of copying the file so that when you install a new version of SnapManager, the link will point to the new BACKINT interface version.

2. Set credentials for the user that runs the BR*Tools commands.

The operating system user needs the credentials of SnapManager for SAP repository, profile, and server to support the backup and restore of the SAP instance.

3. Specify a different profile name.

   By default, SnapManager uses the profile with the same name as the SAP system identifier when processing commands from BR*Tools. If this system identifier is not unique in your environment, modify the `initSID.utl` SAP initialization file, and create a parameter to specify the correct profile. The `initSID.utl` file is located at `%ORACLE_HOME%\database`.

   Example

   A sample `initSID.utl` file is as follows:

   ```
   # Backup Retention policy.
   # Specifies the retention / lifecycle of backups on the filer.
   #
   ------------------------------------------------------------------------
   ----------
   # Default Value: daily
   # Valid Values: unlimited/hourly/daily/weekly/monthly
   # retain = daily
   # Enabling Fast Restore.
   #
   ------------------------------------------------------------------------
   ----------
   # Default Value: fallback
   # Valid Values: fallback/off
   #
   # fast = fallback
   # profile_name = SID_BRTOOLS
   ```

   (i) The parameter name is always in lowercase and the comments must have a number sign (#).

4. Edit the `initSID.sap` BR*Tools configuration file by performing the following steps:

   a. Open the `initSID.sap` file.

   b. Locate the section containing the backup utility parameter file information.

      Example

      ```
      # backup utility parameter file
      # default: no parameter file
      # util_par_file =
      ```

   c. Edit the last line to include the `initSID.utl` file.

```
# backup utility parameter file
# default: no parameter file
# util_par_file = initSID.utl
```

5. Perform the following steps to complete the linking:

   The steps might vary depending on the operating system (OS) version.

   a. Log in to the host system as an administrator.

   b. Click **Start** > **Control Panel** > **System** > **Advanced** > **Environment**.

   c. Edit the **PATH** variable to add the following path: `Installation directory\NetApp\SnapManager for SAP\bin`

After you finish

Register the BACKINT interface in your System Landscape Directory (SLD) by running the backint register-sld command.

Upgrading SnapManager

You can upgrade to the latest version of SnapManager for SAP from any of the earlier versions. You can either upgrade all the SnapManager hosts simultaneously or perform a rolling upgrade, which allows you to upgrade the hosts in a staggered, host-by-host manner.

Preparing to upgrade SnapManager

The environment in which you want to upgrade SnapManager must meet the specific software, hardware, browser, database, and operating system requirements. For the latest information about the requirements, see the Interoperability Matrix.

You must ensure that you perform the following tasks before upgrading:

- Complete the required preinstallation tasks.
- Download the latest SnapManager for SAP installation package.
- Install and configure the appropriate version of SnapDrive for Windows on all the target hosts.
- Create a backup of the existing SnapManager for SAP repository database.

Related information

Interoperability Matrix

Upgrade SnapManager hosts

You can upgrade all the existing SnapManager hosts to use the latest version of

SnapManager. All the hosts are upgraded simultaneously. However, this might result in downtime of all the SnapManager hosts and the scheduled operations during that time.

Steps

1. Stop the SnapManager server by performing the following steps:

    a. In the **Windows Services** window, select **NetApp SnapManager for SAP**.

    b. In the left panel, click **Stop**.

2. Double-click the SnapManager installation file.

If the operating system is…	Then use…
Windows x86	`netapp.smsap.windows-x86-version_number.exe`
Windows x64	`netapp.smsap.windows-x64-version_number.exe`

The following message is displayed: `The publisher could not be verified. Are you sure you want to run this software?`

3. Click **OK**.

    The Introduction window is displayed.

4. Click **Next**.

    The Choose Install Folder window is displayed.

5. Click **Next** to accept the default installation location or select a new location.

    The default location is: `C:\Program Files\NetApp\SnapManager for SAP`.

6. In the **Menu Availability** window, click **Next**.

7. In the **Specify Service Properties** window, enter the account and password information for Windows service.

    The specified account must be a member of the following groups:

    ◦ The storage system's local administration group

    ◦ The local administrator's group

    ◦ The ORA_DBA group You can specify if the service must start automatically after reboot or the service must be started manually.

8. In the **Pre-Installation Summary** window, click **Install**.

9. In the **Install Complete** window, click **Next**.

10. In the **Important Information** window, click **Done** to exit the installer.

Post-upgrade tasks

After upgrading to a later version of SnapManager, you must update the existing repository. You might also want to modify the backup retention class assigned to the existing backups.

> (i) After upgrading to SnapManager 3.3 or later, you need to set
> *sqlnet.authentication_services* to **NONE** if you want to use database (DB) authentication as the only authentication method. This feature is not supported for RAC databases.

Update the existing repository

You do need not to update the existing repository if you are upgrading from SnapManager 3.3.x to SnapManager 3.4 or later, but for all other upgrade paths you must update the existing repository so that you can access it after the upgrade.

What you'll need

- The upgraded SnapManager server must have been started and verified.

- A backup of the existing repository must exist.

About this task

- If you are upgrading from any version earlier than SnapManager 3.1 to SnapManager 3.3 or later, you must first upgrade to SnapManager 3.2.

  After upgrading to SnapManager 3.2, you can then upgrade to SnapManager 3.3 or later.

- After you update the repository, you cannot use the repository with an earlier version of SnapManager.

Step

1. Update the existing repository:

   smsap repository update -repository -dbname *repository_service_name* -host *repository_host_name* -login -username *repository_user_name* -port *repository_port* The repository user name, repository service name, and repository host name can consist of alphanumeric characters, a minus sign, an underscore, and a period. The repository port can be any valid port number. The other options used while updating the existing repository are as follows:

   - The `force` option

   - The `noprompt` option

   - The `quiet` option

   - The `verbose` option

   Example

```
smsap repository update -repository -dbname HR1
-host server1 -login -username admin -port 1521
```

After you finish

Restart the SnapManager server to restart any associated schedules.

Modify the backup retention class

After upgrading, SnapManager assigns the default backup retention class to the existing backups. You can modify the default retention class values to meet your backup requirements.

About this task

The default backup retention class assigned to the existing backups are as follows:

Backup type	Retention class assignment after upgrade
Backups to be retained forever	Unlimited
Other backups	Daily

Note: You can delete the backups that are retained forever without changing the retention class.

If you upgrade to SnapManager 3.0 or later, the greater of the following two values are assigned to the existing profiles:

• Previous retention count for the profile

• Default values for the retention count and duration of daily backups as specified in the `smsap.config` file

Step

1. Modify the values assigned to `retain.hourly.count` and `retain.hourly.duration` in the `smsap.config` file.

   The `smsap.config` file is located at `default installation location/properties/smsap.config`.

   Example

   You can enter the following values:

   ◦ `retain.hourly.count` = **12**

   ◦ `retain.hourly.duration` = **2**

Upgrading SnapManager hosts by using rolling upgrade

The rolling upgrade approach that enables you to upgrade the hosts in a staggered, host-

by-host manner is supported from SnapManager 3.1.

SnapManager 3.0 or earlier only enabled you to upgrade all the hosts simultaneously. This resulted in downtime of all the SnapManager hosts and the scheduled operations during upgrade operation.

Rolling upgrade provides the following benefits:

- Improved SnapManager performance because only one host is upgraded at one time.

- Ability to test the new features in one SnapManager server host before upgrading the other hosts.

(i)   You can perform rolling upgrade only by using the command-line interface (CLI).
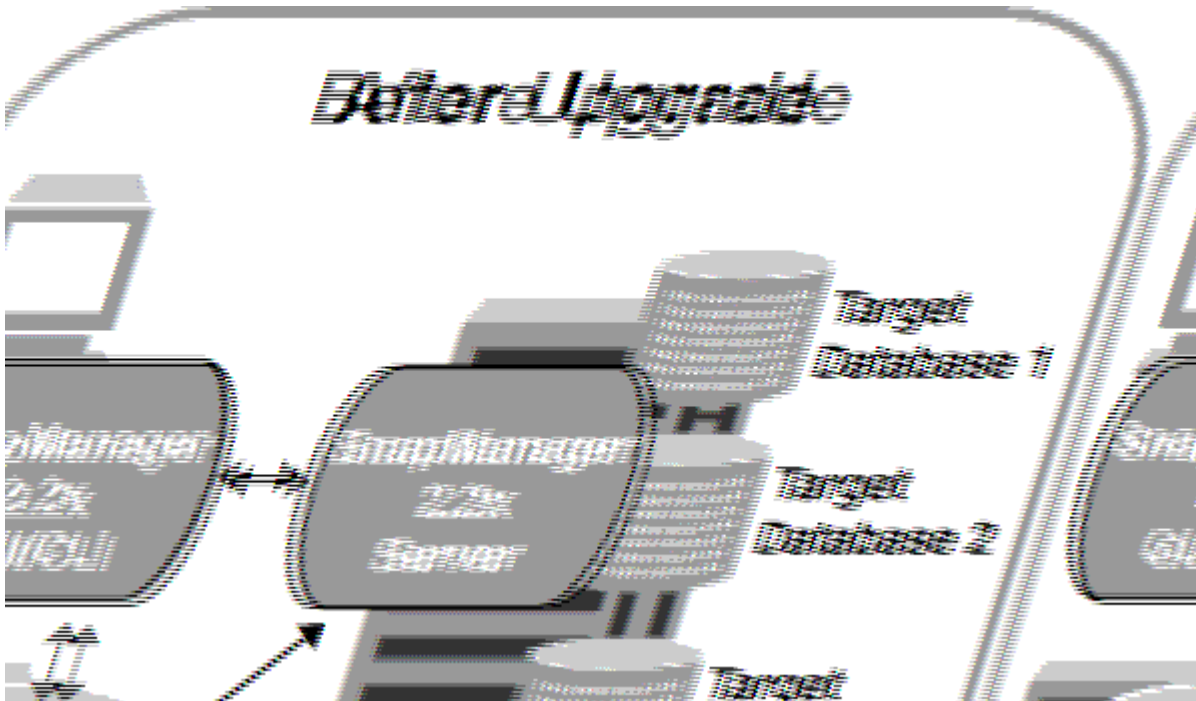
After successful completion of rolling upgrade, the SnapManager hosts, profiles, schedules, backups, and clones associated with the profiles of the target databases are migrated from the repository database of the earlier SnapManager version to the repository database of the new version. The details about the operations performed by using the profiles, schedules, backups, and clones that were created in the earlier SnapManager version are now available in the repository database of the new version. You can start the GUI by using the default configuration values of the user.config file. The values configured in the user.config file of the earlier version of SnapManager are not considered.

The upgraded SnapManager server can now communicate with the upgraded repository database. The hosts that were not upgraded can manage their target databases by using the repository of the earlier version of SnapManager and thereby can use the features available in the earlier version.

(i)   Before performing rolling upgrade, you must ensure that all the hosts under the repository database can be resolved. For information about how to resolve the hosts, see *Troubleshooting SnapManager for SAP*.

Prerequisites for performing rolling upgrades

Before performing a rolling upgrade, you must ensure that your environment meets certain requirements.

- If you are using any version earlier than SnapManager 3.1 and want to perform a rolling upgrade to SnapManager 3.3 or later, you need to first upgrade to 3.2 and then to the latest version.

  You can directly upgrade from SnapManager 3.2 to SnapManager 3.3 or later.

- External scripts that are used to perform any external data protection or data retention must be backed up.

- The SnapManager version to which you want to upgrade must be installed.

  > ⓘ If you are upgrading from any version earlier than SnapManager 3.1 to SnapManager 3.3 or later, you must first install SnapManager 3.2 and perform a rolling upgrade. After upgrading to 3.2, you can then install SnapManager 3.3 or later and perform another rolling upgrade to SnapManager 3.3 or later.

- The SnapDrive for Windows version supported with the SnapManager version to which you want to upgrade must be installed.

  The SnapDrive documentation contains details about installing SnapDrive.

- The repository database must be backed up.

- The amount of SnapManager repository utilization should be minimum.

- If the host to be upgraded is using a repository, SnapManager operations must not be performed on the other hosts that are using the same repository.

  The operations that are scheduled or running on the other hosts wait for the rolling upgrade to finish.

  > ⓘ It is recommended that you perform a rolling upgrade when the repository is least busy, such as over the weekend or when operations are not scheduled.

- Profiles that point to the same repository database must be created with different names in the SnapManager server hosts.

  If you use profiles with the same name, the rolling upgrade involving that repository database fails without warning.

- SnapManager operations must not be performed on the host that is being upgraded.

  > ⓘ The rolling upgrade runs for longer as the number of backups of the hosts being upgraded together increases. The duration of the upgrade can vary depending on the number of profiles and backups associated with a given host.

Related information

[Documentation on the NetApp Support Site: mysupport.netapp.com](mysupport.netapp.com)

Perform roll upgrade on a single host or multiple hosts

You can perform rolling upgrade on a single or multiple SnapManager server hosts by using the command-line interface (CLI). The upgraded SnapManager server host is then managed only with the later version of SnapManager.

What you'll need

You must ensure that all the prerequisites for performing rolling upgrade are completed.

Steps

1. To perform a rolling upgrade on a single host, enter the following command:

   smsap repository rollingupgrade -repository-dbname *repo_service_name* **-host**
   repo_host- _login **-username** *_repo_username* **-port** *_repo_port* **upgrade host**
   host_with_target_database **-force [-quiet | -verbose]**

   Example

   The following command performs the rolling upgrade of all target databases mounted on hostA and a repository database named repoA located on repo_host:

   ```
   smsap repository rollingupgrade
       -repository
         -dbname repoA
         -host repo_host
         -login
          -username repouser
          -port 1521
        -upgradehost hostA
   ```

2. To perform a rolling upgrade on multiple hosts, enter the following command:

   smsap repository rollingupgrade -repository-dbname *repo_service_name* **-host**
   repo_host **-login -username** *repo_username* **-port** *repo_port* **-upgradehost**
   host_with_target_database1,host_with_target_database2 **-force [-quiet |**
   -verbose]

   > ⓘ For multiple hosts, enter the host names separated by a comma and ensure that you do not include any space between the comma and the next host name.Also, ensure that you enter all the host names within double quotation marks.

   Example

   The following command performs the rolling upgrade of all the target databases mounted on the hosts, hostA and hostB and a repository database named repoA located on repo_host:

```
smsap repository rollingupgrade
    -repository
      -dbname repoA
      -host repo_host
      -login
       -username repouser
       -port 1521
     -upgradehost hostA,hostB
```

3. To perform a rolling upgrade on all the hosts on a repository database, enter the following command:

   smsap repository rollingupgrade -repository -dbname *repo_service_name* -host *repo_host* -login -username *repo_username* -port *repo_port* -allhosts -force [- quiet | -verbose]

   After successfully upgrading the repository database, you can perform all the SnapManager operations on the target database.

   Example

   The following command performs the rolling upgrade of all the target databases available on a repository database named repoA located on repo_host:

   ```
   smsap repository rollingupgrade
       -repository
         -dbname repoA
         -host repo_host
         -login
          -username repouser
          -port 1521
         -allhosts
   ```

After you finish

* If the SnapManager server starts automatically, you must restart the server to ensure that you can view the schedules.

* If you upgrade one of the two related hosts, you must upgrade the second host after upgrading the first.

  For example, if you have created a clone from host A to host B or mounted a backup from host A to host B, the hosts A and B are related to each other. When you upgrade host A, a warning message is displayed asking you to upgrade the host B soon after upgrading host A.

  > ⓘ The warning messages are displayed even though the clone is deleted or the backup is unmounted from host B during the rolling upgrade of host A. This is because metadata exists in the repository for the operations performed on the remote host.

What a rollback is

The rollback operation enables you to revert to an earlier version of SnapManager after you perform a rolling upgrade.

> (i) Before performing a rollback, you must ensure that all the hosts under the repository database can be resolved.

When you perform a rollback, the following are rolled back:

- Backups that were created, freed, and deleted by using the SnapManager version from which you are rolling back
- Clones created from a backup that was created by using the SnapManager version from which you are rolling back
- Profile credentials modified by using the SnapManager version from which you are rolling back

The features that were available in the SnapManager version that you were using but are not available in the version to which you are rolling back, are not supported. For example, when you perform a rollback from SnapManager 3.3 or later to SnapManager 3.1, the history configuration set for profiles in SnapManager 3.3 or later is not rolled back to the profiles in SnapManager 3.1. This is because the history configuration feature was not available in SnapManager 3.1.

Limitations for performing a rollback

You must be aware of the scenarios in which you cannot perform a rollback. However, for some of these scenarios you can perform some additional tasks before performing rollback.

The scenarios in which you cannot perform rollback or have to perform the additional tasks are as follows:

- If you perform one of the following operations after performing a rolling upgrade:
  - Create a new profile.
  - Change the mount status of the backup.

    In this scenario, you must first change the mount status to its original state and then perform a rollback.

  - Restore a backup.
  - Change the authentication mode from database authentication to operating system (OS) authentication.

    In this scenario, after performing a rollback, you must manually change the authentication mode from OS to database.

- If the host name for the profile is changed
- If profiles are separated to create archive log backups

  In this scenario, you cannot rollback to a version that is earlier than SnapManager 3.2.

Prerequisites for performing a rollback

Before performing a rollback, you must ensure that your environment meets certain requirements.

- If you are using SnapManager 3.3 or later and want to roll back to a version earlier than SnapManager 3.1, you need to roll back to 3.2 and then to the desired version.
- External scripts that are used to perform any external data protection or data retention must be backed up.
- The SnapManager version to which you want to roll back must be installed.

> ⓘ If you want to perform a rollback from SnapManager 3.3 or later to a version earlier than SnapManager 3.1, you must first install SnapManager 3.2 and perform a rollback. After rolling back to 3.2, you can then install SnapManager 3.1 or earlier and perform another rollback to that version.

- The SnapDrive for Windows version supported with the SnapManager version to which you want to roll back must be installed.

  For information about installing SnapDrive, see SnapDrive documentation set.

- The repository database must be backed up.
- If the host to be rolled back is using a repository, SnapManager operations must not be performed on the other hosts that are using the same repository.

  The operations that are scheduled or running on the other hosts wait for the rollback to complete.

- Profiles that point to the same repository database, must be created with different names in the SnapManager server hosts.

  If you use profiles with the same name, the rollback operation involving that repository database fails without warning.

- SnapManager operations must not be performed on the host which you want to rollback.

  If there is an operation running, you must wait until that operation completes and before proceeding with the rollback.

> ⓘ The rollback operation runs for a longer time as the cumulative number of backups of the hosts that are being rolled back together increases. The duration of the rollback can vary depending on the number of profiles and backups associated with a given host.

Related information

[Documentation on the NetApp Support Site: mysupport.netapp.com](https://mysupport.netapp.com)

Perform a rollback on a single host or multiple hosts

You can perform a rollback on a single or multiple SnapManager server hosts by using the command-line interface (CLI).

What you'll need

You must ensure that all the prerequisites for performing a rollback are complete.

Steps

1. To perform a rollback on a single host, enter the following command:

   smsaprepository rollback -repository -dbname *repo_service_name* **-host** *repo_host* **-login -username** *repo_username* **-port** *repo_port* **-rollbackhost** *host_with_target_database*

   Example

   The following example shows the command to roll back all the target databases that are mounted on hostA and a repository database named repoA located on the repository host, repo_host:

   ```
   smsap repository rollback
       -repository
         -dbname repoA
         -host repo_host
         -login
          -username repouser
          -port 1521
       -rollbackhost hostA
   ```

2. To perform a rollback on multiple hosts, enter the following command:

   smsaprepository rollback -repository-dbname *repo_service_name* **-host** *repo_host* **-login -username** *repo_username* **-port** *repo_port* **-rollbackhost** *host_with_target_database1,host_with_target_database2*

   > (i) For multiple hosts, enter the host names separated by a comma and ensure that there is no space between the comma and the next host name.Also, ensure that you enter the whole set of multiple host names within double quotation marks.

   Example

   The following example shows the command to roll back all the target databases that are mounted on the hosts, hostA, hostB, and a repository database named repoA located on the repository host, repo_host:

   ```
   smsap repository rollback
       -repository
         -dbname repoA
         -host repo_host
         -login
          -username repouser
          -port 1521
       -rollbackhost hostA,hostB
   ```

The hosts, profiles, schedules, backups, and clones that are associated with the profiles of the target databases for the host are reverted to the earlier repository.

Post rollback tasks

You must perform some additional steps after you rollback a repository database and downgrade the SnapManager host from SnapManager 3.2 to SnapManager 3.0, to view the schedules created in the earlier version of the repository database.

Steps

1. Navigate to `C:\Program Files\NetApp\SnapManager for Oracle\repositories`.

   The `repositories` directory might contain two files for each repository. The file name with the number sign (#) is created using SnapManager 3.1 or later and the file name with the hyphen (-) is created using SnapManager 3.0.

   Example

   The file names might be as follows:

   ◦ `repository#SMSAP300a#SMSAPREPO1#10.72.197.141#1521`

   ◦ `repository-smsap300a-smsaprepo1-10.72.197.141-1521`

2. Replace the number sign (#) in the file name with the hyphen (-).

   Example

   The file name that had the number sign (#), now contains hyphen (-): `repository-SMSAP300a-SMSAPREPO1-10.72.197.141-1521`.

Configuring SnapManager

After installing SnapManager, you must perform some additional configuration tasks depending on the environment that you are using.

SnapManager configuration parameters

SnapManager provides a list of configuration parameters that you can edit depending on your requirement. The configuration parameters are stored in the `smsap.config` file. However, the `smsap.config` file might not contain all the supported configuration parameters. You can add the configuration parameters, depending on your requirement.

The following table lists all the supported SnapManager configuration parameters and also explains when to use these parameters:

Parameters	Description

• `retain.hourly.count`
 	
• `retain.hourly.duration`
 	
• `retain.monthly.count`
 	
• `retain.monthly.duration`
	These parameters set the retention policy when you create a profile. For example, you can assign the following values:
	
	`retain.hourly.count = 12`
	
	`retain.hourly.duration = 2`
	
	`retain.monthly.count = 2`
	
	`retain.monthly.duration = 6`

`restore.temporaryVolumeName`	This parameter assigns a name to the temporary volume. When SnapManager uses the indirect method for restoring data from secondary storage, it requires a scratch volume on the primary storage to hold a temporary copy of data until it is copied into the database files and the database is recovered. There is no default value. If you do not specify a value, you must enter a name in the restore command that uses the indirect method. For example, you can assign the following values: `restore.temporaryVolumeName =` **`smsap_temp_volume`**
`host.credentials.persist`	This parameter allows SnapManager to store host credentials. By default, the host credentials are not stored. However, host credentials need to be stored if you have a custom script that runs on a remote clone and requires access to a remote server.You can enable storing of host credentials by assigning **true** to `host.credentials.persist`. SnapManager encrypts and saves the host credentials.
`restorePlanMaxFilesDisplayed`	This parameter enables you to define the maximum number of files to be displayed in the restore preview.By default, SnapManager displays a maximum of 20 files in the restore preview. However, you can change to a value greater than 0. For example, you can assign the following value: • `restorePlanMaxFilesDisplayed = 30` ⓘ If you specify an invalid value, the default number of files are displayed.

Parameters	Description
snapshot.list.timeout.min	This parameter enables you to define the time in minutes for which SnapManager must wait for the `snap list` command to execute when you are performing any SnapManager operations.By default, SnapManager waits for 30 minutes. However, you can change to a value greater than 0. For example, you can assign the following value: • `snapshot.list.timeout.min = 40` ⓘ If you specify an invalid value, the default value is used. For any SnapManager operation, if the `snap list` command execution time exceeds the value assigned to `snapshot.list.timeout.min`, the operation fails with a timeout error message.
pruneIfFileExistsInOtherDestination	This pruning parameter enables you to define the destination of the archive logs files. The archive log files are stored in multiple destinations. While pruning archive log files, SnapManager needs to know the destination of the archive log files. The possible values that you can assign are as follows: • When you want to prune the archive log files from a specified destination, you must assign **false** to `pruneIfFileExistsInOtherDestination`. • When you want to prune the archive log files from an external destination, you must assign **true** to `pruneIfFileExistsInOtherDestination`.

Parameters	Description
`prune.archivelogs.backedup.from.otherdestination`	This pruning parameter enables you to prune the archive log files backed up from the specified archive log destinations or backed up from external archive log destinations. The possible values that you can assign are as follows: • When you want to prune the archive log files from the specified destinations and if the archive log files are backed up from the specified destinations by using `-prune-dest`, you must assign **false** to `prune.archivelogs.backedup.from.other destination`. • When you want to prune the archive log files from specified destinations and if the archive log files are backed up at least once from any one of the other destinations, you must assign **true** to `prune.archivelogs.backedup.from.other destination`.
`maximum.archivelog.files.toprune.atATime`	This pruning parameter enables you to define the maximum number of archive log files that you can prune at a given time. For example, you can assign the following `value:maximum.archivelog.files.toprune.atATime` = **998** (i) The value that can be assigned to `maximum.archivelog.files.toprune.atATime` must be less than 1000.
`archivelogs.consolidate`	This parameter allows SnapManager to free the duplicate archive log backups if you assign **true** to `archivelogs.consolidate`.
`suffix.backup.label.with.logs`	This parameter enables you to specify the suffix that you want to add to differentiate the label names of the data backup and the archive log backup.For example, when you assign logs to `suffix.backup.label.with.logs`, `_logs` is added as a suffix to the archive log backup label. The archive log backup label would then be `arch_logs`.

Parameters	Description
`backup.archivelogs.beyond.missingfiles`	This parameter allows SnapManager to include the missing archive log files in the backup. The archive log files that do not exist in the active file system are not included in the backup. If you want to include all of the archive log files, even those that do not exist in the active file system, you must assign **true** to `backup.archivelogs.beyond.missingfiles`.

You can assign false to ignore the missing archive log files. |
| `srvctl.timeout` | This parameter enables you to define the timeout value for the `srvctl` command.

ⓘ The Server Control (SRVCTL) is a utility to manage RAC instances.

When SnapManager takes more time to execute the srvctl command than the timeout value, the SnapManager operation fails with this error message: `Error: Timeout occurred while executing command: srvctl status.` |
| `snapshot.restore.storageNameCheck` | This parameter allows SnapManager to perform the restore operation with Snapshot copies that were created before migrating from Data ONTAP operating in 7-Mode to clustered Data ONTAP.The default value assigned to the parameter is false. If you have migrated from Data ONTAP operating in 7-Mode to clustered Data ONTAP but want to use the Snapshot copies created before migration, set `snapshot.restore.storageNameCheck`=**true**. |
| `services.common.disableAbort` | This parameter disables cleanup upon failure of long-running operations. You can set `services.common.disableAbort`=**true**.For example, if you are performing a clone operation that runs long and then fails because of an Oracle error, you might not want to clean up the clone. If you set `services.common.disableAbort`=**true**, the clone will not be deleted. You can fix the Oracle issue and restart the clone operation from the point where it failed. |

Parameters	Description
• `backup.sleep.dnfs.layout` • `backup.sleep.dnfs.secs`	These parameters activate the sleep mechanism in the Direct NFS (dNFS) layout. After you create the backup of control files using dNFS or a Network File System (NFS), SnapManager tries to read the control files, but the files might not be found.To enable the sleep mechanism, ensure that `backup.sleep.dnfs.layout`=**true**. The default value is **true**. When you enable the sleep mechanism, you must assign the sleep time to `backup.sleep.dnfs.secs`. The sleep time assigned is in seconds and the value depends upon your environment. The default value is 5 seconds. For example: • `backup.sleep.dnfs.layout`=**true** • `backup.sleep.dnfs.secs`=**2**
• `override.default.backup.pattern` • `new.default.backup.pattern`	When you do not specify the backup label, SnapManager creates a default backup label. These SnapManager parameters allows you to customize the default backup label.To enable customization of the backup label, ensure that the value of `override.default.backup.pattern` is set to **true**. The default value is **false**. To assign the new pattern of the backup label, you can assign keywords such as database name, profile name, scope, mode, and host name to `new.default.backup.pattern`. The keywords should be separated using an underscore. For example, `new.default.backup.pattern`=**dbname_profile_hostname_scope_mode**. ⓘ The timestamp is included automatically at the end of the generated label.

Parameters	Description
`allow.underscore.in.clone.sid`	Oracle supports usage of the underscore in clone SID from Oracle 11gR2. This SnapManager parameter enables you to include an underscore in the clone SID name.To include an underscore in the clone SID name, ensure that the value of `allow.underscore.in.clone.sid` is set to **true**. The default value is **true**. If you are using an Oracle version earlier than Oracle 11gR2 or if you do not want to include an underscore in the clone SID name, set the value to **false**.
`oracle.parameters.with.comma`	This parameter enables you to specify all the Oracle parameters that have comma (,) as the value.While performing any operation SnapManager uses `oracle.parameters.with.comma` to check all the Oracle parameters and skip the splitting of the values. For example, if the value of *nls_numeric_characters*=,, then specify `oracle.parameters.with.comma=`*nls_numeric_characters*. If there are multiple Oracle parameters with comma as the value, you must specify all the parameters in `oracle.parameters.with.comma`.

Parameters	Description
• `archivedLogs.exclude` • `archivedLogs.exclude.fileslike` • `<db-unique-name>.archivedLogs.exclude.fileslike`	These parameters allow SnapManager to exclude the archive log files from the profiles and backups if the database is not on a Snapshot copy-enabled storage system and you want to perform SnapManager operations on that storage system. ⓘ You must include the exclude parameters in the configuration file before creating a profile. The values assigned to these parameters can either be a top-level directory or a mount point where the archive log files are present or a subdirectory. To exclude archive log files from being included in the profile and being backed up, you must include one of the following parameters: • `archivedLogs.exclude` to specify a regular expression for excluding archive log files from all profiles or backups. The archive log files matching the regular expression are excluded from all the profiles and backups. For example, you can set `archivedLogs.exclude = J:\\ARCH\\.*`. ⓘ If the destination has a file separator, then an additional slash symbol (\) must be added to the pattern and the pattern must end with a double-slash pattern (\\.*). • `archivedLogs.exclude.fileslike` to specify an SQL expression for excluding archive log files from all profiles or backups. The archive log files matching the SQL expression are excluded from all the profiles and backups. For example, you can set `archivedLogs.exclude.fileslike = J:\\ARCH2\\%`. ⓘ If the destination has a file separator, then an additional slash symbol (\) must be added to the pattern and the pattern must end with a double-slash pattern (\\%).

Parameters	Description
	• `<db-unique-name>.archivedLogs.exclude.fileslike` to specify an SQL expression for excluding archive log files only from the profile or the backup created for the database with the specified *db-unique-name*. The archive log files matching the SQL expression are excluded from the profile and backups. For example, you can set `mydb.archivedLogs.exclude.fileslike = J:\\ARCH2\\%`. ⓘ If the destination has a file separator, then an additional slash symbol (\\) must be added to the pattern and the pattern must end with a double-slash pattern (\\%). ⓘ The BR*Tools do not support the following parameters even if these parameters are configured to exclude archive log files: • `archivedLogs.exclude.fileslike` • `<db-unique-name>.archivedLogs.exclude.fileslike`

Edit the configuration parameters

Depending on your environment, you can change the default values assigned to the configuration parameter.

Steps

1. Open the configuration file from the following default location:

   `default installation location\properties\smsap.config`

2. Change the default values of the configuration parameters.

   ⓘ You can also add supported configuration parameters that are not included in the configuration file, and assign values to them.

3. Restart the SnapManager for SAP server.

Start SnapManager for SAP

The SnapManager startup section lists the tasks that you perform when you start SnapManager. Use this section also if you are just learning about SnapManager.

What you'll need

Before using SnapManager, you should have performed the following actions:

- Downloaded and installed the SnapManager software.
- Determined whether you will use the graphical user interface or the command-line interface.

Identify an existing database to backup

You can identify the system identifier (SID) of the SnapManager database which is used in creating a profile.

About this task

The standard Oracle user ID for SAP systems is ora `sid` where `sid` is a three-character, alphanumeric value, for example, PRD, PR1, DEV, or QA3.

Steps

1. Click **Start** > **Control Panel** > **Administrative Tools** > **Services**.
2. Verify the Oracle service, `OracleServiceSID`.

   If the service is called OracleServiceFASDB, then the database SID is FASDB.

Verify the Oracle listener status

You can verify the Oracle listener status by using the `lsnrctl status` command.

What you'll need

- You must connect to the database.

About this task

The listener port details are as follows:

- A standard SAP Oracle installation sets the listener port to 1527.
- Anything that references a managed database listener port should have the value as 1527.
- Anything that references the repository database listener port should have the value as 1521.

The Oracle listener's log is located at `%ORACLE_HOME%\network\log` where `ORACLE_HOME is \oracle\SID\oracle version`.

Step

1. At the command prompt, enter the command:

```
lsnrctl status
```

Create Oracle users for the repository database

You can create an Oracle user for the repository database and assign specific privileges to perform different operations on the repository database.

About this task

You must assign the connect and resource privileges to the Oracle user. You do not have to create a user for the repository database with sysdba privileges.

ℹ️ | However, you must create an Oracle user with the sysdba role for the target database.

Steps

1. Log in to SQL *Plus.

   At the command prompt, enter the following command:

   sqlplus '/ as sysdba'

   ```
   SQL*Plus: Release 11.2.0.1.0 Production on Wed Jun 1 06:01:26 2011
   Copyright (c) 1982, 2009, Oracle.  All rights reserved.
   Connected to:
   Oracle Database 11g Enterprise Edition Release 11.2.0.1.0 - Production
   With the Partitioning, Automatic Storage Management, OLAP, Data Mining
   and Real Application Testing options
   ```

2. To create a user, for example *repo1_user*, for the repository with the administrator password, for example, *adminpw1*, enter the following command at the SQL prompt:

   SQL> create user *repo1_user* identified by *adminpw1*;

3. To grant connect and resource privileges to the user, enter the following command:

   grant connect, resource to *repo1_user*;

Create an Oracle user for the target database

You need to create an Oracle user with the sysdba role that connects to the database and performs database operations.

About this task

SnapManager can use any Oracle user with sysdba privileges that exists in the target database, for example, the default "sys" user. You can also create a user in the target database to be used exclusively by SnapManager.

Steps

1. Log in to SQL *Plus.

   At the command prompt, enter the following command:

   sqlplus '/ as sysdba'

2. To create a user, for example *smsap_oper* with the administrator password, for example, *adminpw1*, enter the following command at the SQL prompt:

   SQL> create user *smsap_oper* identified by *adminpw1*;

3. Grant sysdba privileges to the Oracle user by entering the following command:

   SQL> grant sysdba to *smsap_oper*;

Accessing SnapManager

You can access SnapManager by using either the command-line interface (CLI) or graphical user interface (GUI).

You can perform different SnapManager operations in the following ways:

- By entering commands in the CLI on a host that is in the same network as the database host.

  For a list of all the commands and an explanation of their options and arguments, see the Command Reference chapter.

  To access the CLI, click **Start** > **All Programs** > **NetApp** > **SnapManager for SAP** > **Start SMSAP command-line interface (CLI)**.

- By accessing the GUI on a host in the same network as the database host.

  The GUI provides simple and easy-to-use wizards to help you perform different operations.

- By using SAP BR*Tools commands.

Start the SnapManager host server

You can start the SnapManager server by using the Windows services.

Steps
1. Click **Start** > **Control Panel** > **Administrative Tools** > **Services**.
2. In the **Services** window, select NetAppSnapManager 3.3 forSAP.
3. You can start the server in one of three ways:
   - In the left panel, click **Start**.
   - Right-click NetAppSnapManager 3.3 forSAP and select **Start** from the drop-down menu.
   - Double-click NetAppSnapManager 3.3 forSAP and then in the Properties window, click **Start**.

Verify the SnapManager host server status

The server must be running for you to execute commands or initiate SnapManager operations. You must verify the status of the server before performing any operations.

Steps

1. In the Services window, select SnapManager 3.3 forSAP.

2. View the status in the Status column.

Use SnapManager commands

After you start the SnapManager host server, you can use SnapManager by entering commands at the prompt on your host.

Step

1. To perform an operation:

   ○ In case of a Windows host, go to **Start** > **All Programs** > **NetApp** > **SnapManager for SAP** > **Start SMSAP Command Line Interface (CLI)**

Start the SnapManager GUI

If SnapManager is installed on the host, start the graphical user interface (GUI) for SnapManager by selecting the program from a list of programs.

What you'll need

• Ensure that the SnapManager server is started.

About this task

You can start the SnapManager GUI in one of the following ways:

• In the SnapManager host, click **Start** > **All Programs** > **NetApp** > **SnapManager for SAP** > **Start SMSAP GUI**.

• If SnapManager is not installed on the host, use Java Web Start, which downloads SnapManager components and starts the GUI.

Download and start the graphical user interface using Java Web Start

You can use Java Web Start if SnapManager is not installed on the host. Java Web Start downloads SnapManager components and starts the graphical user interface (GUI). The supported JRE versions are 1.5, 1.6, and 1.7.

What you'll need

You must ensure that the following conditions are met:

• The SnapManager server is running.

• A web browser window is open.

Steps

1. In the Microsoft Internet Explorer Web browser window, enter `https://smsap-server.domain.com:port`.

   `smsap-server.domain.com` is the fully qualified host name and domain on which you installed SnapManager and `port` is the listening port for the SnapManager server (27314, by default).

   > ⓘ You must enter `https` in the browser window.

   A dialog box with the following message is displayed.

   > ⓘ A dialog box with the message There is a problem with the site's security certificate…Do you want to proceed?

1. Click **Yes** or **Continue**.

2. Click on the link labeled **Click here to download and install JRE 6.0 and the application**.

   A link labeled Download Java Web Start with the following message is displayed.

   > ⓘ This site might require the following ActiveX control: Java Plug-in 1.6"… "Click here to install

3. In the **Install** window, perform the following steps:

   a. Click the message labeled `Click here to install…`.

      An **Install ActiveX Control** menu is displayed.

   b. Select **Install ActiveX Control…**.

      The following message is displayed.

      > ⓘ Internet Explorer - Security Warning" containing the following text: "Do you want to install this software? Name: Java Plug-in 1.6

   a. Click **Install**.

      A "Java Plug-in 1.6." window for the installer for J2SE Runtime Environment 1.6 is displayed.

   b. Click **Install**.

      A window requesting you to install J2SE Runtime Environment 1.6 is displayed. . In the Install window, perform the following steps:

   c. On the **License Agreement** page, select **I accept the terms in the license agreement** and click **Next**.

   d. On the **Setup Type** page, select **Typical** and click **Next**.

   e. On the **Installation Completed** window, click **Finish**.

      SnapManager begins to download.

      A File Download dialog box with the message `Do you want to save this file?`

`application.jnlp` is displayed. . In the **file download** window, perform the following steps:

f. Install the latest version of JRE 1.6 on the Windows client.

g. Verify that Java is installed by running the following command: `java -version`

   The output should indicate Java version 1.6.0_24 (which is Java 1.6) or later.

h. Change your Windows configuration settings to always open files with extension jnlp with the program Java Web Start Launcher.

   The steps to change the Windows configuration settings varies based on the version of Windows you are using.

i. Enter the SnapManager URL that you have specified in the step 1.

   The SnapManager download starts on the windows client and the following warning is displayed

> ⚠️ `security dialog box is displayed.`
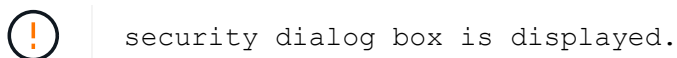
4. Perform the following steps.

   The message contents and button labels vary based on platform.

   a. In the **Warning - Security** dialog box, click **Yes**.

      A dialog box is displayed.

   b. In the host name mismatch dialog box, click **Run**.

      The Warning - Security dialog box with a message about the signature of the SnapManager application is displayed.

   c. Click **Run**.

      A dialog box with the title Java Installer - Security Warning and the following message is displayed.

> ⚠️ `Warning Security - the application's digital signature has an error. Do you want run the application`
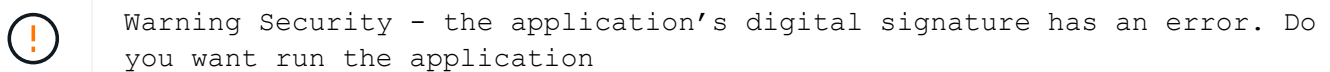
a. Click **Run**.

   The browser downloads and starts the SnapManager for SAP GUI.

Verify the environment

You can verify the environment to make sure SnapDrive and SnapManager are set up correctly.

What you'll need

Download, install, and set up the required prerequisites. Make sure SnapManager is installed and the host

server is running.

Step

1. To verify that SnapDrive is installed and can be run from the root account, run the following command:

```
smsap system verify
```

Verify SnapDrive for Windows

If you have installed SnapDrive for Windows, verify that you can create a Snapshot copy before using SnapManager.

Steps

1. From the Start menu, right-click **My Computer** and select **Manage**.

2. In the **Computer Management** window, click **Storage** > **SnapDrive**.

3. Select a disk.

   See the *SnapDrive for Windows Installation and Administration Guide* for more information about using SnapDrive.

   If you have successfully found disk information for the SnapDrive product, SnapDrive is working correctly.

Related information

SnapDrive for Windows Installation and Administration Guide

Create repositories

SnapManager requires a repository on a host to hold data about the operations you perform.

What you'll need

Ensure that the following tasks are completed:

Steps

1. Create an Oracle user and password in the repository database.

2. Authorize user access to the repository.

For a repository, SnapManager for SAP requires a minimum 4K block size for the tablespace into which it is installed. You can check the block size using the following SQL command:

```
select a.username, a.default_tablespace, b.block_size
from dba_users a, dba_tablespaces b
a.username = repo_user
```

where

- `a.default_tablespace = b.tablespace_name`
- `a.username` = the user name on the repository

About this task

If you are upgrading repositories, you must reboot the SnapManager server to restart any associated schedules.

Step

1. To create the repository, enter the repository create command, using the following general format:

   `smsap repository create -repository -dbname repo_service_name -host repo_host -login -username repo_username -port repo_port [-force] [-noprompt] [-quiet | -verbose]`

   Where:

   ° `-repository -dbname` is the name of the repository database.

   ° `-host` is the name of the host for the repository.

   ° `-username` is the name of the database user who has access to the repository.

   ° `-port` is the port for the host. Other options for this command are as follows:

   `[-force] [-noprompt]`

   `[quiet | -verbose]`

   > ⓘ If you have an existing repository with the same name and you use the `-force` option, all data within an existing repository schema will be overwritten.

Creating a repository

The following command line creates a repository.

```
smsap repository create -repository -dbname HR1
-host server1 -login -username admin -port 1521
```

How to organize repositories

You can organize the SnapManager repositories to meet your business needs. You can organize them in several ways, including by application type and usage.

You can organize repositories in several ways. Two such ways are as follows:

Type	Characteristics
By application	If you have multiple Oracle databases running different applications, you can create a SnapManager repository for every application type. Each SnapManager repository would have profiles for the databases of a particular application type. All production, development, and testing databases of that application type would be managed by the same SnapManager repository. This option would help group similar databases and ease cloning. However, if you have several application types, then you might have to manage several SnapManager repositories, and if you choose to implement another application type, you will need to create another SnapManager repository. Because these SnapManager repositories will be managing production databases, each of these repositories must be on a server with high availability, which can be expensive. Also, having to manage production databases along with development and test databases of the same type in the same SnapManager repository can be a security issue.
By usage	You can distribute the databases among the SnapManager repositories based on their usage (for example, production, development, testing, and training). This option limits the number of repositories to the different types of databases that you have. Because all production databases would be managed by a single SnapManager repository, only production database administrators can be given access to this repository. Also, if you choose to deploy another database for a new application type, then you only need to register the database in the corresponding SnapManager repository instead of creating a new repository. High availability can be provided only for the SnapManager repository that holds the profiles of all the production databases.

SnapManager for Oracle and SnapManager for SAP should not share the same repository. For SnapManager for Oracle and SnapManager for SAP, you must use a different repository (a different Oracle database user) for each product if you have both in your environment. Using a different repository, either in the same or different databases, provides more flexibility by allowing independent upgrade cycles for each product.

Order of performing operations

SnapManager enables you to perform various operations such as creating profiles, performing backups, and cloning backups. These operations must be performed in a specific order.

Steps

1. Create a profile on an existing repository by using the `smsap profile create` command.

(i) The Oracle user specified for the target database must have sysdba privileges.

Example

The following example shows the command to create a profile:

```
smsap profile create -profile prof1 -profile-password prof1cred
-repository -dbname HR1 -login -username admin -host server1 -port 1527
-database -dbname SID -host db_server1 -port 1527
```

2. Create a backup on an existing profile by using the `smsap backup create` command.

Example

The following example shows the command to create a backup:

```
smsap backup create -profile prof1 -full -offline -label
full_backup_prof1 -force
```

3. Restore and recover a database backup on primary storage by using the `smsap backup restore` command.

Example

The following example shows the command to restore a backup:

```
smsap backup restore -profile prof1 -label full_backup_prof1
-complete -recover -alllogs
```

4. Create a clone specification by using the `smsap clone template` command.

You can use the Clone wizard in the graphical user interface (GUI) to create a template clone specification. You can also create the clone specification file by using a text editor.

5. Clone a database with an existing backup by using the `smsap clone create` command.

You must have an existing clone specification or create a clone specification to specify the storage and database specifications for the clone.

Example

The following example shows the command to create a clone:

```
smsap clone create -profile prof1 -backup-label full_backup_prof1
-newsid clone1 -label prof1_clone -clonespec
C:\\clone_spec\\prof1_clonespec.xml
```

Security and credential management

You can manage security in SnapManager by applying user authentication. The user authentication method allows you to access resources such as repositories, hosts, and profiles.

When you perform an operation using either the command-line interface (CLI) or graphical user interface (GUI), SnapManager retrieves the credentials set for repositories and profiles. SnapManager saves credentials from previous installations.

The repository and profiles can be secured with a password. A credential is the password configured for the user for an object, and the password is not configured on the object itself.

You can manage authentication and credentials by performing the following tasks:

- Manage user authentication either through password prompts on operations or by using the `smsap credential set` command.

  Set credentials for a repository, host, or profile.

- View the credentials that govern the resources to which you have access.
- Clear a user's credentials for all resources (hosts, repositories, and profiles).
- Delete a user's credentials for individual resources (hosts, repositories, and profiles).

> (i) If the repository database is on a Windows host, then both local or administrator user and the domain user must have the same credentials.

What user authentication is

SnapManager authenticates the user by using an operating system (OS) login on the host where the SnapManager server is running. You can enable user authentication either through password prompts on operations or by using the smo credential You can enable user authentication either through password prompts on operations or by using the `smsap credential set`.

User authentication requirements depend on where the operation is performed.

- If the SnapManager client is on the same server as the SnapManager host, you are authenticated by the OS credentials.

  You are not prompted for a password because you are already logged in to the host where the SnapManager server is running.

- If the SnapManager client and the SnapManager server are on different hosts, SnapManager needs to authenticate you with both OS credentials.

  SnapManager prompts you for passwords for any operation, if you have not saved your OS credentials in your SnapManager user credential cache. If you enter the `smsap credential set -host` command, you save the OS credentials in your SnapManager credential cache file and so SnapManager does not prompt for the password for any operation.

If you are authenticated with the SnapManager server, you are considered the effective user. The effective user for any operation must be a valid user account on the host on which the operation is executed. For example, if you execute a clone operation, you should be able to log in to the destination host for the clone.

> ⓘ SnapManager for SAP might fail in authorizing users created in Central Active Directory Services, such as LDAP and ADS. To ensure the authentication does not fail, you must set configurable `auth.disableServerAuthorization` to **true**.

As an effective user you can manage credentials in the following ways:

- Optionally, you can configure SnapManager to store user credentials in the SnapManager user credentials file.

  By default, SnapManager does not store host credentials. You might want to change this, for example, if you have custom scripts that require access on a remote host. The remote clone operation is an example of a SnapManager operation that needs the login credentials of a user for a remote host. To have SnapManager remember user host login credentials in the SnapManager user credentials cache, set the `host.credentials.persist` property to **true** in the `smsap.config` file.

- You can authorize user access to the repository.
- You can authorize user access to profiles.
- You can view all user credentials.
- You can clear a user's credentials for all resources (hosts, repositories, and profiles).
- You can delete credentials for individual resources (hosts, repositories, and profiles).

Store encrypted passwords for custom scripts

By default, SnapManager does not store host credentials in the user credentials cache. However, you can change this. You can edit the `smsap.config` file to allow storing of host credentials.

About this task

The `smsap.config` file is located at `<default installation location>\properties\smsap.config`

Steps

1. Edit the `smsap.config` file.

2. Set *host.credentials.persist* to **true**.

Authorize access to the repository

SnapManager enables you to set credentials for database users to access the repository. Using credentials, you can restrict or prevent access to the SnapManager hosts, repositories, profiles, and databases.

About this task

If you set credentials by using the `credential set` command, SnapManager does not prompt you for a password.

You can set user credentials when you install SnapManager or later.

Step

1. Enter the following command:

   smsap credential set -repository -dbname *repo_service_name* **-host** *repo_host*
   -login -username *repo_username* **[-password** *repo_password*] **-port** *repo_port*

Authorize access to profiles

SnapManager enables you to set a password for a profile to prevent unauthorized access.

Step

1. Enter the following command:

   smsap credential set -profile -name *profile_name* **[-password** *password*]

View user credentials

You can list the hosts, profiles, and repositories to which you have access.

Step

1. To list the resources to which you have access, enter this command:

   smsap credential list

Example of viewing user credentials

This example displays the resources to which you have access.

```
smsap credential list
```

```
Credential cache for OS user "user1":
Repositories:
Host1_test_user@SMSAPREPO/hotspur:1521
Host2_test_user@SMSAPREPO/hotspur:1521
user1_1@SMSAPREPO/hotspur:1521
Profiles:
HSDBR (Repository: user1_2_1@SMSAPREPO/hotspur:1521)
PBCASM (Repository: user1_2_1@SMSAPREPO/hotspur:1521)
HSDB (Repository: Host1_test_user@SMSAPREPO/hotspur:1521) [PASSWORD NOT
SET]
Hosts:
Host2
Host5
```

Clear user credentials for all hosts, repositories, and profiles

You can clear the cache of your credentials for resources (hosts, repositories, and profiles). This deletes all of the resource credentials for the user running the command. After clearing the cache, you must authenticate your credentials again to gain access to these secured resources.

Steps

1. To clear your credentials, enter the `smsap credential clear` command from the SnapManager CLI or select **Admin** > **Credentials** > **Clear Cache** from the SnapManager GUI.

2. Exit the SnapManager GUI.

> (i)
> - If you have cleared the credential cache from the SnapManager GUI, you do not need to exit the SnapManager GUI.
>
> - If you have cleared the credential cache from the SnapManager CLI, you must restart SnapManager GUI.
>
> - If you have deleted the encrypted credential file manually, you must restart the SnapManager GUI again.

3. To set the credentials again, repeat the process to set credentials for the repository, profile host, and profile. For additional information on setting the user credentials again, refer to "Setting credentials after clearing credential cache."

Set credentials after clearing the credential cache

After clearing the cache to remove the stored user credentials, you can set the credentials for the hosts, repositories, and profiles.

About this task

You must ensure that you set the same user credentials for the repository, profile host, and profile that you had given earlier. An encrypted credentials file is created while setting the user credentials.

The credentials file is located at `C:\Documents and Settings\Administrator\Application Data\NetApp\smsap\3.3.0`.

From the SnapManager graphical user interface (GUI), if there is no repository under Repositories, perform the following steps:

Steps

1. Click **Tasks** > **Add Existing Repository** to add an existing repository.

2. Perform the following steps to set the credentials for repository:

    a. Right-click the repository and select **Open**.

    b. In the `Repository Credentials Authentication` window, enter the user credentials.

3. Perform the following steps to set the credentials for host:

    a. Right-click the host under the repository and select **Open**.

    b. In the `Host Credentials Authentication` window, enter the user credentials.

4. Perform the following steps to set the credentials for profile:

    a. Right-click the profile under the host and select **Open**.

    b. In the `Profile Credentials Authentication` window, enter the user credentials.

Delete credentials for individual resources

You can delete the credentials for any one of the secured resources, such as a profile, repository, or host. This enables you to remove the credentials for just one resource, rather than clearing the user's credentials for all resources.

Delete user credentials for repositories

You can delete the credentials so a user can no longer access a particular repository. This command enables you to remove the credentials for just one resource, rather than clearing the user's credentials for all resources.

Step

1. To delete repository credentials for a user, enter this command:

   smsap credential delete -repository -dbname *repo_service_name* -host *repo_host* -login -username *repo_username* -port *repo_port*

Delete user credentials for hosts

You can delete the credentials for a host so a user can no longer access it. This command enables you to remove the credentials for just one resource, rather than clearing all the user's credentials for all resources.

Step

1. To delete host credentials for a user, enter this command:

   smsap credential delete -host -name_host_name_-username_-username_

Delete user credentials for profiles

You can delete the user credentials for a profile so a user can no longer access it.

Step

1. To delete profile credentials for a user, enter this command:

```
smsap credential delete -profile -name profile_name
```

Managing profiles for efficient backups

You must create a profile in SnapManager for the database on which you want to perform an operation. You must select the profile and then select the operation that you want to perform.

Tasks related to profiles

You can perform the following tasks:

- Create profiles to enable full or partial backups and backups to primary or secondary storage.

  You can also create profiles to separate the archive log backups from the data file backups.

- Verify profiles.
- Update profiles.
- Delete profiles.

About profiles and authentication

When you create a profile, you can specify a database and choose one of the following methods to connect to the database:

- Oracle authentication with a user name, password, and port
- Operating system (OS) authentication with no user name, password, or port.

  For OS authentication, you must enter the OS account user and group information.

- Database authentication when `sqlnet.authentication_services` is set to **NONE**. SnapManager then uses the database user name and password for all the connections to the target database. SnapManager can also use the Windows native authentication when `sqlnet.authentication_services` is set to **NTS**.

You can set `sqlnet.authentication_services` to **NONE** only in the following environments:

Database layout	Oracle version	Is database authentication supported for the target database
Any non-ASM and non-RAC database	Oracle 10g and Oracle 11g (lesser than 11.2.0.3)	Yes

After you disable `sqlnet.authentication_services` and change the authentication method to database authentication, you must set `sqlnet.authentication_services` to **NONE**.

If you are accessing a profile for the first time, you must enter your profile password. After you enter your credentials, you can view the database backups within the profile.

Create profiles

When creating profiles, you can assign a particular Oracle database user account to the profile. You can set the retention policy for the profile, and set the retention count and duration for each retention class.

About this task

If you do not provide the values of the `-login`, `-password`, and `-port` parameters of the database, the operating system (OS) authentication mode uses the default credentials.

While creating a profile, SnapManager performs a restore eligibility check to determine the restore mechanism that can be used to restore the database.

SnapManager (3.2 or later) enables you to separate archive log files from the data files while creating a new profile or updating an existing profile. After you have separated the backup using the profile, you can either create only the data files-only backup or archive log-only backup of the database. You can use the new profile or the updated profile to create the backup containing both the data files and archive log files. However, you cannot use the profile to create the full backup or revert the settings.

Profiles for creating full and partial backups

You can create profiles to create the full database backup containing the data files, control files, and archive log files and partial database backup containing specified data files or tablespaces, all the control files, and all the archive log files. SnapManager does not allow you to create separate archive log backups using the profiles created for full and partial backups.

Profiles for creating data files-only backups and archivelogs-only backups

When you create a new profile, you can include `-separate-archivelog-backups` to separate the archive log backup from the data file backup. You can also update the existing profile to separate the archive log backup from the data file backup.

By using the new profile options to separate the archive log backups, 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

While creating the profile to separate archive log backups from the data files backup, if the archive log files do not exist in the database for which the profile is created, then a warning message is displayed.

> ⚠️ Archived log file does not exist in the active file system. The
> archived log file versions earlier than the <archive log thread
> version> log file will not be included in the backup

Even if you create backups for this database, the archive log files are not available in the database backups.

> ℹ️ If you encounter an error while creating a profile, use the `smsap system dump` command.
> After you create a profile, if you encounter an error, use the `smsap operation dump` and
> `smsap profile dump` commands.

Step

1. To create a profile with a user name, password, and port (Oracle authentication), enter the following command:

   **smsap profile create -profileprofile [-profile-passwordprofile_password]
   -repository-dbnamerepo_dbname-hostrepo_host-portrepo_port-login
   -usernamerepo_username-database-dbnamedb_dbname-hostdb_host [-siddb_sid] [-
   login-usernamedb_username-passworddb_password-portdb_port] [-rman {-
   controlfile | {-login-usernamerman_username-passwordrman_password-
   tnsnamerman_tnsname} }] [-retain [-hourly [-countn] [-durationm]] [-daily [-
   countn] [-durationm]] [-weekly [-countn] [-durationm]] [-monthly [-countn] [-
   durationm]]] [-commentcomment][-snapname-patternpattern][-summary-
   notification] [-notification [-success-emailemail_address1, email_address2-
   subjectsubject_pattern] [-failure-emailemail_address1, email_address2-
   subjectsubject_pattern]][-separate-archivelog-backups-retain-archivelog-
   backups-hourshours | -daysdays | -weeksweeks| -monthsmonths [-include-with-
   online-backups | -no-include-with-online-backups]] [-dump]**

   Other options for this command are as follows:

   `[-force] [-noprompt]`

   `[quiet | verbose]`

   You can also include other options when creating profiles, depending on how you want to access the database.

If…	Then…
You want to use database authentication to create a profile	Specify the database login details.

If…	Then…
You want to specify a backup retention policy for backups	Specify either the retention count or duration for a retention class, or both. The duration is in units of the class (for example, hours for hourly, days for daily). • `-hourly` is the hourly retention class, for which `[-count n]` `[-duration m]` are the retention count and retention duration, respectively. • `-daily` is the daily retention class, for which `[-count n]` `[-duration m]` are the retention count and retention duration, respectively. • `-weekly` is the weekly retention class, for which `[-count n]` `[-duration m]` are the retention count and retention duration, respectively. • `-monthly` is the monthly retention class, for which `[-count n]` `[-duration m]` are the retention count and retention duration, respectively.

If…	Then…
You want to enable email notification for the completion status of the database operations	Specify the following options and variables: • `-summary-notification` enables you to configure a summary email notification for multiple profiles under a repository database. • `-notification` enables you to receive an email notification for the completion status of the database operation for a profile. • `-success-email``email_address2` enables you to receive an email notification on the successful database operation performed by using a new or existing profile.  • `-failure-email``email_address2` enables you to receive an email notification on the failed database operation performed by using a new or existing profile.  • `-subject``subject_text` specifies the subject text for the email notification while creating a new profile or an existing profile. If the notification settings are not configured for the repository and you try to configure profile or summary notifications by using the CLI, the following message is logged in the console log: `SMSAP-14577: Notification Settings not configured.` If you have configured the notification settings and you try to configure summary notification by using the CLI without enabling summary notification for the repository, the following message is shown in the console log: `SMSAP-14575: Summary notification configuration not available for this repository`

If…	Then…
You want to backup archive log files separately from data files	Specify the following options and variables: • `-separate-archivelog-backups` enables you to separate the archive log backup from the datafile backup. • `-retain-archivelog-backups` sets the retention duration for archive log backups. You must specify a positive retention duration. The archive log backups are retained based on the archive log retention duration. The data files backups are retained based on the existing retention policies. • `-include-with-online-backups` includes the archive log backup along with the online database backup. This option enables you to create an online data files backup and archive logs backup together for cloning. When this option is set, whenever you create an online data files backup, the archive logs backups are created along with the data files immediately. • `-no-include-with-online-backups` does not include the archive log backup along with database backup.
You can collect the dump files after the successful profile create operation	Specify the `-dump` option at the end of the `profile create` command.

Snapshot copy naming

You can specify a naming convention or pattern to describe the Snapshot copies related to the profile you create or update. You can also include custom text in all Snapshot copy names.

You can change the Snapshot copy naming pattern when you create a profile or after the profile has been created. The updated pattern applies only to Snapshot copies that have not yet occurred; Snapshot copies that exist retain the previous snapname pattern.

The following examples show the two Snapshot copy names taken for a volume. The second Snapshot copy listed has *F_H_1* in the middle of its name. The "1" indicates that it is the first Snapshot copy taken in the backup set. The first Snapshot copy listed is the most recent and has a "2," which means it is the second Snapshot copy taken. The "1" Snapshot copy includes the datafiles; the "2" Snapshot copy includes the control files. Because the control file Snapshot copies must be taken after the data file Snapshot copy, two Snapshot copies are required.

```
smsap_profile_sid_f_h_2_8ae482831ad14311011ad14328b80001_0
smsap_profile_sid_f_h_1_8ae482831ad14311011ad14328b80001_0
```

The default pattern includes the required smid, as shown in the following:

- Default pattern: smsap_{profile}_{db-sid}_{scope}_{mode}_{smid}
- Example: smsap_my_profile_rac51_f_h_2_8abc01e915a55ac50115a55acc8d0001_0

You can use the following variables in the Snapshot copy name:

Variable name	Description	Example value
smid (Required)	The SnapManager unique ID is the only required element when creating a name for the Snapshot copy. This ID ensures that you create a unique Snapshot name.	8abc01e915a55ac50115a55acc8d0001_0
class (Optional)	Retention class associated with the backup for the profile and indicated by hourly (h), daily (d), weekly (w), monthly (m), or unlimited (u).	d
comment (Optional)	Comment associated with the backup for the profile. Spaces in this field will be converted to underscores when the Snapshot copy name is complete.	sample_comment_spaces_replaced
date (Optional)	Date that the backup occurs for the profile. Date values are padded with zeros if necessary. (yyyymmdd)	20070218
db-host (Optional)	Database host name associated with the profile being created or updated.	my_host
db-name (Optional)	Database name associated with the Snapshot copy you create.	rac5
db-sid (Optional)	Database sid associated with the Snapshot copy you create.	rac51
label (Optional)	Label associated with the backup for the profile.	sample_label

Variable name	Description	Example value
mode (Optional)	Specifies whether the backup is completed online (h) or offline (c).	h
profile (Optional)	Profile name associated with the backup you create.	my_profile
scope (Optional)	Specifies whether the backup is either full (f) or partial (p).	f
time (Optional)	Time that the backup occurs for the profile. Time values for this variable use the 24-hour clock and are padded with zeros if necessary. For example, 5:32 and 8 seconds appears as 053208 (hhmmss).	170530
time-zone (Optional)	Time zone specified for the target database host.	EST
usertext (Optional)	Custom text that you can enter.	prod

> ⓘ SnapManager for SAP does not support the colon (:) symbol in the long forms of the names for Snapshot copies.

Rename profiles

SnapManager enables you to rename the profile when you update the profile. The SnapManager capabilities that are set on the profile and the operations that can be performed before renaming are retained for the renamed profile.

What you'll need

- You must ensure that there are no SnapManager operations running on the profile while renaming the profile.

About this task

You can rename the profile from both the SnapManager command-line interface (CLI) and graphical user interface (GUI). While updating the profile, SnapManager verifies and updates the profile name in the repository.

> ⓘ SnapManager does not support renaming the profile in the Multi-profile update window.

When you provide a new profile name, the new profile name is added in the client-side credential cache and the earlier profile name is removed. When you rename the profile from a client, the credential cache of only that client is updated. You need to execute the `smsap profile sync` command from each of the clients to update the new credential cache with the new profile name.

You can set the password for the profile by using the `smsap credential set` command.

If the profile name was included in a Snapshot copy naming pattern, when you rename a profile, the new name for the profile gets updated. All the SnapManager operations that are performed on the profile use the new profile name. The backups created with earlier profile continue to have the earlier profile name and are used to perform other SnapManager operations.

If you are performing rolling upgrade of SnapManager server hosts, you must ensure that you perform the complete upgrade before renaming the profile.

The new name for the profile is updated only from the SnapManager client from which the request is made. The SnapManager clients that are connected to the SnapManager server are not notified about the change in profile name. You can check the operation log to know about the change in the profile name.

> ⓘ If a scheduled backup operation begins at the time of renaming the profile, then the scheduled operation fails.

Step

1. Enter the following command:

   smsap profile update -profile *profile* [-new-profile *new_profile_name*]

Change profile passwords

To protect the existing profiles in the repository, you should update the passwords for the profiles. You can apply this updated password when creating a backup using this profile.

Step

1. To update the profile password for an existing profile, enter this command:

   smsap profile update -profile profile_name -profile-password password

Reset the profile password

You can reset the profile password if you do not remember the password that you had provided while creating the profile.

What you'll need

- You must ensure that the SnapManager server is running on the repository database.
- You must have the local administrator credentials of the host on which the repository database is residing.
- You must ensure that the profile is not in use for any operation when the password is being reset for that profile.

About this task

You can reset the password from either the SnapManager CLI or GUI. While resetting the password, SnapManager queries the SnapManager server on the repository host to identify the operating system for the repository host. You must enter the authorized user credentials for connecting to the repository host. The SnapManager server authenticates users with their local administrator credentials on the repository database.

When the authentication is successful, SnapManager resets the profile password on the SnapManager server with the new password.

ⓘ   SnapManager does not maintain the history of the password reset operations.

Step

1. Reset the profile password by entering the following command:

   ```
   smsap password reset -profile profile [-profile-password profile_password] [-repository-hostadmin-password admin_password]
   ```

Authorize access to profiles

SnapManager enables you to set a password for a profile to prevent unauthorized access.

Step

1. Enter the following command:

   ```
   smsap credential set -profile -name profile_name [-password password]
   ```

Verify profiles

You can verify that an existing profile is set up correctly. When you verify a profile, SnapManager checks the environment for the profile you specify and verifies that the profile is set up and the database in this profile is accessible.

Step

1. To verify if the profile is set up correctly, enter this command:

   ```
   smsap profile verify -profile profile_name
   ```

Update profiles

You can update the profiles to modify the profile password, the number of backups to retain, access to the database, the operating system (OS) authentication to database authentication and vice versa, and information about the host. If the Oracle database password information changes, you must also change that information in the profile.

About this task

SnapManager (3.2 or later) enables you to update the profile to separate archive log backups from the data file backups by using the -separate-archivelog-backups option. You can specify separate retention duration for the archive log backup. SnapManager enables you to include the archive log backup along with online database backup. You can also create an online datafile backup and archive log backup together for cloning. When you create an online data files backup, the archive logs backups are immediately created along with the data files.

Step

1. Enter the following command:

   smsap profile update -profile *profile* **[-new-profile** *new_profile_name***] [-profile-password** *profile_password***][-database -*dbnamedb_dbname* -host** *db_host* **[-sid** *db_sid***] [-login -usernamedb_username-password** *db_password***-port** *db_port***]] [{-rman{-controlfile | {-login -username** *rman_username* **-password rman_password-tnsname** *rman_tnsname***}}} | -remove-rman] [-retain [-hourly [-countn] [-durationm]] [-daily [-countn] [-durationm]] [-weekly [-countn] [-durationm]] [-monthly [-countn] [-durationm]]] [-commentcomment][-snapname-patternpattern] [-summary-notification] [-notification [-success-email** *email_address1***,** *email_address2* **-subjects** *ubject_pattern***] [-failure-email** *email_address1***,** *email_address2* **-subject** *subject_pattern***]] [-separate-archivelog-backups-retain-archivelog-backups-hours** *hours* **| -days** *days* **| -weeks** *weeks***| -months** *months* **[-include-with-online-backups | -no-include-with-online-backups]] [-dump]**

   Other options for this command are as follows:

   [`-force`] [`-noprompt`]

   [`quiet` | `verbose`]

If you want to…	Then…
Change the backup retention policy for backups of the database in the profile	Specify either the retention count or retention duration for a retention class, or both to change the retention policy. The duration is in units of the class (for example, hours for hourly, days for daily). • `-hourly` is the hourly retention class, for which [`-countn`] [`-durationm`]] are the retention count and retention duration, respectively. • `-daily` is the daily retention class, for which [`-countn`] [`-durationm`]] are the retention count and retention duration, respectively. • `-weekly` is the weekly retention class, for which [`-countn`] [`-durationm`]] are the retention count and retention duration, respectively. • -monthly is the monthly retention class, for which [`-countn`] [`-durationm`]] are the retention count and retention duration, respectively.

If you want to…	Then…
Enable email notifications for the completion status of the database operations	Specify the following options and variables: • `-summary-notification` enables you to configure a summary email notification for multiple profiles under a repository database. • `-notification` enables you to receive an email notification on the completion status of the database operation for a profile. • `-success-emailemail_address2` enables you to receive an email notification following the completion of a successful database operation performed by using a new or an existing profile. • `-failure-emailemail_address2` enables you to receive an email notification on a failed database operation performed by using a new or an existing profile. • `-subjectsubject_text` specifies subject text for the email notification while creating a new profile or an existing profile. If the notification settings are not configured for the repository and you are trying to configure profile or summary notifications by using the command-line interface (CLI), the following message is logged in the console log: `SMSAP-14577: Notification Settings not configured.` If you have configured the notification settings and you are trying to configure summary notification by using the CLI without enabling summary notification for the repository, the following message is logged in the console log: `SMSAP-14575: Summary notification configuration not available for this repository`

If you want to…	Then…
Update the profile to create backup of the archive log files separately	Specify the following options and variables: • `-separate-archivelog-backups` enables you to create a backup of the archive log files separately from the database files. After you specify this option, you can either create a data files-only backup or an archivelogs-only backup. You cannot create a full backup. Also, you cannot revert the profile settings by separating the backup. SnapManager retains the backups based on the retention policy for the backups that were created before taking archivelogs-only backup. • `-retain-archivelog-backups` sets the retention duration for archive log backups. (i) If you are updating the profile for the first time, you can separate the archive log backups from the data files backup by using the `-separate-archivelog-backups` option; you must provide the retention duration for the archive log backups by using the `-retain-archivelog-backups` option. Setting the retention duration is optional when you later update the profile. • `-include-with-online-backups` specifies that the archive log backup is included along with the database backup. • `-no-include-with-online-backups` specifies the archive log file backup is not included along with the database backup.
Change the host name of the target database	Specify `-hostnew_db_host` to change the host name of the profile.
Collect the dump files after the profile update operation	Specify the `-dump` option.

2. To view the updated profile, enter the following command: `smsap profile show`

Delete profiles

You can delete a profile anytime, as long as it does not contain successful or incomplete

backups. You can delete profiles that contain freed or deleted backups.

Step

1. To delete a profile, enter this command:

   ```
   smsap profile delete -profile profile_name
   ```

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 archivelog-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.

> ℹ️ 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.

> (i) 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:
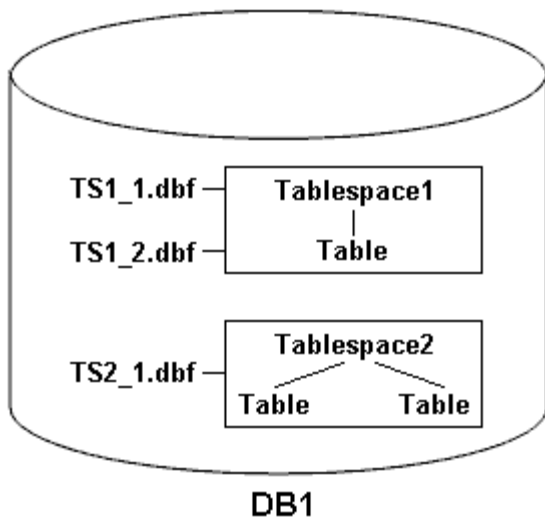
In the illustration, Tablespace1 has one table and two database files associated with it. Tablespace2 has two tables and one database file associated with it.

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
SnapManager makes a backup of everything in database DB1, including the data files, archive logs, and control files.	Complete restore with control files SnapManager restores all data files, tablespaces, and control files in the backup.	You can specify one of the following: • 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.
	Complete restore without control files SnapManager restores all tablespaces and data files, without the control files.	
	Restore either data files or tablespaces with control files Specify one of the following: • Tablespaces • Data files	SnapManager recovers the data to the last transaction made to the database.
	Restore either data files or tablespaces without control filesSnapManager restores one of the following: • Tablespaces • Data files	
	Restore control files only	

Examples of partial backup, restore, and recover operations

Partial backup	Restore	Recover

You can choose one of the following options:
 	
• Tablespaces
 	
  You can specify Tablespace1 and Tablespace2 or only one of them.
 	
• Data files
 	
  You can specify all three database files (TS1_1.dbf, TS1_2.dbf, and TS2_1.dbf), two files, or one file.
 	
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.
	Complete restore SnapManager restores all data files, tablespaces, and control files specified in the partial backup.	SnapManager recovers the data to the last transaction made to the database instance.

Restore either data files or tablespaces with control files SnapManager restores one of the following:
 	
• All the data files specified
• All the tablespaces specified
	Restore either data files or tablespaces without control files SnapManager restores one of the following:
	
	• Tablespaces
	
	  Specify any of the tablespaces. SnapManager restores only the tablespaces specified. If the backup contains Tablespace1, SnapManager restores only that tablespace.
	
	• Data files
	
	  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.
		Restore control files only

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
What happens when two backups are scheduled to occur on two databases at the same time?	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. 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. Before starting a backup operation, SnapManager first determines the following: • 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? If the answer to either question is yes, SnapManager skips the backup. 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. The time window of 30 minutes can be changed in a SnapManager properties file.
Under which user does the backup operation run?	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.
How does the SnapManager scheduler interact with the native operating system scheduler?	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.

Question	Answer
What happens if the clocks in the graphical user interface and the server are not in sync?	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. For recurring backups, the server still fulfills the request. For instance, if the server receives a request to perform hourly backups starting an 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. However, for one-time only backups, the server handles the request as follows: • 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. For instance, consider the following scenario: • 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. 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. 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. You can change the time of five minutes in a SnapManager properties file.
What happens to the scheduled backups for a profile when the profile is deleted?	When a database profile is deleted, the SnapManager server deletes scheduled backups defined for that profile.

Question	Answer

How do scheduled backups behave during Daylight Savings Time or when you change the SnapManager server time?	SnapManager backup schedules get affected due to Daylight Savings Time or when you change the SnapManager server time.
	
	 Consider the following implications when the SnapManager server time is changed:
	
	 • 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:
	
	  ◦ 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 `-offline` to take a backup of the offline database. Specify `-online` to take a backup of the online database. If you use these options, you cannot use the `-auto` option.

If you want to…	Then…
Specify whether you want to let SnapManager handle backing up a database regardless of whether it is online or offline	Specify the `-auto` option. If you use this option, you cannot use the `--offline` or `-online` option.
Specify whether you want to perform a partial backup of specific files	Specify the `-data-files` option and then list the files, separated by commas. For example, list the file names f1, f2, and f3 after the option. Example for creating a partial datafile backup on Windows `smsap backup create -profile nosep -data -files "J:\\mnt\\user\\user.dbf" -online -label partial_datafile_backup -verbose`
Specify whether you want to perform a partial backup of specific tablespaces	Specify the `-data -tablespaces` option and then list the tablespaces, separated by commas. For example, use ts1, ts2, and ts3 after the option. 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. Example for creating a partial tablespace backup `smsap backup create -profile nosep -data -tablespaces tb2 -online -label partial_tablespace_bkup -verbose`

If you want to…	Then…
Specify whether you want to create a unique label for each backup in the following format: full_hot_mybackup_label	For Windows, you might enter this example: ``` smsap backup create -online -full -profile targetdb1_prof1 -label full_hot_my_backup_label -verbose ```
Specify whether you want to create backup of the archive log files separately from the data files	Specify the following options and variables: • `-archivelogs` creates a backup of the archive log files. • `-backup-dest` specifies the archive log file destinations to be backed up. • `-exclude-dest` specifies the archive log destinations to be excluded. • `-label` specifies the label for the archive log file backup. (i) You must provide either the `-backup -dest` option or the `-exclude -dest` option. Providing both these options together along with the backup displays error message `You have specified an invalid backup option. Specify any one of the options: -backup-dest, or exclude-dest.` Example for creating archive log file backups separately on Windows ``` smsap backup create -profile nosep -archivelogs -backup-dest "J:\\mnt\\archive_dest_2\\" -label archivelog_backup -verbose ```

If you want to…	Then…

Specify whether you want to create backup of data files and archive log files together	Specify the following options and variables:
	
	• `-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>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…

Specify whether you want to prune the archive log files while creating a backup	Specify the following options and variables:
	
	• `-prunelogs` specifies to delete the archive log files from the archive log destinations.
	
	  ° `-all` specifies to delete all the archive log files from the archive log destinations.
	
	  ° `-until-scnuntil-scn` specifies to delete the archive log files until a specified SCN.
	
	  ° `-until-dateyyyy-MM-dd:HH:mm:ss` specifies to delete the archive log files until the specified time period.
	
	  ° `-before` option specifies to delete the archive log files before the specified time period (days, months, weeks, hours).
	
	  ° `-prune -destprune_dest1,[prune_dest2` specifies to delete the archive log files from the archive log destinations while creating the backup.
	
	  ⓘ You cannot prune the archive log files when Flash Recovery Area (FRA) is enabled for archive log files.
	
	Example for pruning all archive log files while creating a backup on Windows
	
	<pre>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>

Specify whether you want to add a comment about the backup	Specify `-comment` followed by the description string.
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	Specify the `-force` option.

If you want to…	Then…
Specify whether you want to verify the backup at the same time you create it	Specify the `-verify` option.
Specify whether you want to collect the dump files after the database backup operation	Specify `-dump` 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.

  > (i) 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.archivelog.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.

> (i) 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:
	
	• `-prunelogs` specifies deleting the archive log files while creating a backup.
	
	   ◦ `-all` specifies deleting all the archive log files.
	
	   ◦ `-untilSCN` specifies deleting the archive log files until the specified SCN.
	
	   ◦ `-until-date` specifies deleting the archive logs including the specified date and time.
	
	   ◦ `-before {-months -days -weeks -hours}` 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 `-prune-dest` 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.

2. 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

1. 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: • `-prunelogs` to schedule pruning of the archive log files • `-prune-dest` to prune archive log files from the archive log destinations
Include a name for the schedule	Specify the `-schedule-name` 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: • `-hourly` • `-daily` • `-weekly` • `-monthly` • `-onetimeonly`
Add a comment about the schedule operation	Specify the `-schedule-comment` option followed by the description string.
Specify the start time of the schedule operation	Specify the `-start-time` option in the yyyy-mm-dd hh:mm 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.

> (i)   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.

> (i)   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.

Scheduling database backups

SnapManager (3.2 or later) for SAP enables you to schedule database backups to occur on a regular basis during off-peak hours to maintain high performance. To schedule a backup, you can create a profile, which includes the database information and retention policy, and then set schedules for the backup.

> ⓘ You must schedule the backups as an administrator. If you try to schedule the backups as a non-existing user, SnapManager displays an error message: Invalid user: username: Cannot create schedule backup for a given user

The following are some of the schedule-related tasks:

- Schedule a database backup to occur on an hourly, daily, weekly, monthly, or one-time basis.
- View a list of scheduled backups associated with a profile.
- Update a scheduled backup.

- Suspend a schedule temporarily.
- Resume the suspended schedule.
- Delete a schedule.

> ⓘ The **Run Now Menu Operation** check box is disabled when a scheduled backup is running for that schedule.

Create backup schedules

You can schedule a backup to occur at the time and frequency that are suited for your data and environment.

About this task

From SnapManager 3.2 for SAP, you can schedule the backups of the archive log files separately. However, you must use the profile that you created to separate the archive log files.

If you have scheduled the backups of the data files and archive log files at the same time, then SnapManager creates the data files backup first.

If you select the schedule interval as `-onetimeonly`, then all the pruning options are available. If you select a schedule interval other than `-onetimeonly`, then the pruning options `-until-SCN` and `-until-date` are not supported and the following error message is displayed: `The archive log pruning option you have specified, -until-scn or -until-date for the schedule interval hourly is invalid. Specify either the -onetimeonly option for the schedule interval, or prune the archive logs using any one of the option all, or -before {-months | -days | -weeks| -hours}.`

When a failover happens in Microsoft Windows Server Failover Cluster (WSFC) and Microsoft Cluster Server (MSCS) environments, you must restart the SnapManager for SAP server so that the service (virtual) address is mapped to the active host and the SnapManager schedules are adjusted to the active SnapManager host.

> ⓘ If the same profile and schedule name exists in another repository, the backup scheduling operation is not initiated in that repository. The operation will exit with the following message: `operation is already running.`

Step

1. 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* **| -until-date** *yyyy-MM-dd HH:mm:ss* **| -before {-months | -weeks | -days | -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] [-taskspec -***taskspec***] [-quiet | -verbose]**

If you want to…	Then…
Schedule a backup of an online or offline database	Specify `-offline` or `-online` to schedule a backup of the offline or online database. If you specify these, you cannot use `-auto`.
Let SnapManager handle scheduling of a database regardless of whether it is online or offline	Specify `-auto`. If you specify `-auto`, you cannot use `--offline` or `-online`.
Schedule a backup of data files	Specify `-data -files` to list the files separated by commas. For example, use file names f1,f2,f3.
Schedule a partial backup of specific tablespaces	Specify `-tablespaces` to list the tablespaces separated by commas. For example, use ts1,ts2,ts3.
Schedule backup of archive log files	Specify the following: • `-archivelogs` to schedule backup of the archive log files • `-backup-dest` to schedule archive log file destinations to be included in the backup • `-exclude-dest` to schedule the archive log destinations to be excluded from the backup
Specify the retention class values	Specify `-retain` and indicate whether the backup should be retained according to one of the following retention classes: • `-hourly` • `-daily` • `-weekly` • `-monthly` • `-unlimited` SnapManager defaults to hourly.
Schedule pruning of archive log files	Specify the following: **`-prunelogs` to prune the archive log files while scheduling a backup** `-prune-dest` to specify the archive log destination from which the archive log files are pruned
Include a name for the schedule	Specify `-schedule-name`.

If you want to…	Then…
Schedule backup of the database at a specific time interval	Specify the `interval` option and select the time interval from the following, by which the backups should be created: • `-hourly` • `-daily` • `-weekly` • `-monthly` • `-onetimeonly`
Configure a schedule	Specify `-cronstring` and include the following seven subexpressions that describe the individual option: • 1 refers to seconds. • 2 refers to minutes. • 3 refers to hours. • 4 refers to a day in a month. • 5 refers to the month. • 6 refers to a day in a week. • (Optional) 7 refers to the year. ⓘ If you scheduled your backup with different times in `-cronstring` and `-start-time`, then the schedule of the backup is overwritten and triggered by the `-start-time`.
Add a comment about the backup schedule	Specify `-schedule-comment` followed by the description string.
Specify the start time of the schedule operation	Specify `-start-time` in the yyyy-mm-dd hh:mm format.
Change the user of the scheduled backup operation while scheduling the backup	Specify `-runasuser`. The operation runs as the user (root user or Oracle user) who created the schedule. However, you can use your own user ID, if you have valid credentials for both the database profile and host.

If you want to…	Then…
Enable a pretask or post-task activity of the backup schedule operation by using the pretask and post-task specification XML file	Specify the `-taskspec` option and provide the absolute path of the task specification XML file for performing a preprocessing or a post-processing activity to occur before or after the backup schedule operation.

Update a backup schedule

You can view a list of scheduled operations and update them if necessary. You can update the scheduling frequency, the start time of the schedule, cronstring expression, and the user who scheduled the backup.

Step

1. To update the schedule for a backup, enter this command:

   ```
   smsap schedule update -profile profile_name -schedule-name schedulename [-
   schedule-comment schedule comment] -interval {-hourly | -daily | -weekly |
   -monthly | -onetimeonly} -start-time starttime -cronstring cronstring
   -runasuser runasuser [-quiet | -verbose]
   ```

View a list of scheduled operations

You can view a list of scheduled operations for a profile.

Step

1. To display information about scheduled operation, enter this command:

   ```
   smsap schedule list -profile profile_name [-quiet | -verbose]
   ```

Suspend backup schedules

SnapManager enables you to suspend a backup schedule until the backup schedule is resumed.

About this task

You can suspend the active schedules. If you try to suspend the backup schedule that is already suspended, you might encounter error message "`Cannot suspend: schedule <schedulename> already in suspend state`".

Step

1. To suspend the backup schedule temporarily, enter this command:

   ```
   smsap schedule suspend -profile profile_name -schedule-name schedulename [-
   quiet | -verbose]
   ```

Resume backup schedules

Administrators have the option to resume the suspended backup schedule.

About this task

If you try to resume the active schedules, you might encounter the error message: `"Cannot resume: schedule <schedulename> already in resume state"`.

Step

1. To resume the suspended backup schedule, enter this command:

   ```
   smsap schedule resume -profile profile_name -schedule-name schedulename [-
   quiet | -verbose]
   ```

Delete backup schedules

You can delete backup schedules when they are no longer necessary.

Step

1. To delete the backup schedule, enter this command:

   ```
   smsap schedule delete -profile profile_name -schedule-name schedulename [-
   quiet | -verbose]
   ```

Restoring database backups

SnapManager for SAP enables you to restore a database to the state it was when a Snapshot copy was taken. Because backups are created more frequently, the number of logs that need to be applied is reduced, thus reducing the mean-time-to-recovery (MTTR) for a database.

The following are some of the tasks that you can perform related to restoring and recovering data in databases:

- Perform a file-based restore.
- Restore the entire backup or a portion of it.

  If you restore a portion of it, you specify a group of tablespaces or a group of data files. You can also restore the control files along with the data or just the control files themselves.

- Recover the data based on either a point in time or on all of the available logs, which stores the last transaction committed to the database.

  The point in time can be an Oracle System Change Number (SCN) or a date and time (yyyy-mm-dd:hh:mm:ss). SnapManager uses the 24-hour clock.

- Restore from backups on primary storage (local backups).
- Restore and recover the backup by using SnapManager, or use SnapManager to restore the backup and use another tool, such as Recovery Manager (RMAN), to recover the data.

- Restore backups from alternate locations.

For more information, see the .

You can restore a backup made by a previous version of SnapManager by using SnapManager 3.0 and later versions.

Administrators can perform restore or recovery operations by using the SnapManager graphical user interface (GUI) or by using the command-line interface (CLI).

What database restore is

SnapManager enables you to perform file-based backup and restore operations.

The following table describes the restore methods:

Restore process	Details
File-based restores	Storage-side full file system restore (from primary or secondary): SnapManager performs a full logical unit number (LUN) restore.

Storage-side full file system restore

A storage-side full file system restore is performed when a volume restore cannot be performed, but the entire files system can be restored on the storage system.

When a storage-side file system restore is performed, the following occurs:

- In a SAN environment, all the LUNs used by the file system (and underlying volume group if any) are restored on the storage system.

When a storage-side file system restore is performed, the following occurs, depending on the storage location:

- When SnapManager restores from primary storage systems, the LUNs (SAN) are restored in place via SFSR.
- When SnapManager restores from secondary storage systems, the LUNs (SAN) are copied from secondary storage systems back to the primary storage system over the network.

Because the file system is fully restored, files that are not part of the backup are reverted as well. An override is required if files, which are not part of the restore, exist in the file system that is being restored.

Host-side file restore

A host-side file copy restore is used as a last resort in SAN environments when storage-side file system restore and storage-side file restore cannot be performed.

A host-side file copy restore involves the following tasks:

- Cloning the storage
- Connecting the cloned storage to the host
- Copying files out of the clone file systems back into the active file systems

- Disconnecting the clone storage from the host
- Deleting the clone storage

Backup recovery

In SnapManager, you must perform the restore and recover operations at the same time. You cannot perform a restore operation and then perform a SnapManager recover operation later.

In SnapManager 3.2 or earlier, you can either use SnapManager to restore and recover the backup or use SnapManager to restore the backup and use another tool, such as Oracle Recovery Manager (RMAN), to recover the data. Because SnapManager can register its backups with RMAN, you can use RMAN to restore and recover the database at finer granularities such as blocks. This integration combines the benefits of speed and space efficiency of Snapshot copies with the fine level of control for restoring using RMAN.

> (i) You must recover a database before you can use it. You can use any tool or script to recover a database.

Starting from SnapManager 3.2 for SAP, SnapManager enables the restore of database backups automatically by using the archive log backups. Even when the archive log backups are available in the external location, SnapManager uses the archive log backups from the external location to restore the database backups.

If new data files are added to the database, Oracle recommends that you take 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 Oracle recovery process might fail, because it is unable to create data files. See the Oracle documentation for the process for recovering data files added after a backup.

Database state needed for the restore process

The state of the database that is to be restored depends on the type of restore process that you want to perform and the type of files that are to be included.

The following table lists the state in which the database should be depending on the restore option selected and the type of files you want to include in the restore:

Type of restore	Files included	Database state for this instance
Restore only	Control files	Shutdown
	System files	Mount or Shutdown
	No system files	Any state
Restore and recovery	Control files	Shutdown
	System files	Mount
	No system files	Mount or Open

The database state required by SnapManager for a restore operation depends on the type of restore being

performed (complete, partial, or control files). SnapManager does not transition the database to a lower state (for example, from Open to Mount) unless the `force` option is specified.

SnapManager for SAP does not verify whether SAP is running. SnapManager for SAP waits for a timeout to elapse and then shuts down the database, which can add an hour of additional restore time.

What restore preview plans are

SnapManager provides restore plans before and after a restore operation is completed. The restore plans are used to preview, review, and analyze regarding different restore methods.

Structure of the restore plan

The restore plan consists of the following two sections:

- Preview/Review: This section describes how SnapManager will restore (or has restored) each file.
- Analysis: This section describes why some restore mechanisms were not used during the restore operation.

The Preview/Review section

This section shows how each file will be or has been restored. When you view the restore plan before a restore operation, it is called a preview. When you view it after a restore operation is completed, it is called a review.

The following preview example shows that the files are restored by using storage-side file system restore and storage-side system restore methods. To determine why all the files would not be restored by using the same restore method, see the Analysis section.

```
Preview:
The following files will be restored completely via: storage side full
file system restore
E:\rac6\sysaux.dbf
E:\rac6\system.dbf
```

Each restore method has one subsection that contains information about the files that can be restored using that restore method. The subsections are ordered according to decreasing levels of storage method efficiency.

It is possible for one file to be restored by multiple restore methods. Multiple restore methods are used when the underlying logical unit numbers (LUNs) used for a file system are spread among different storage system volumes and some volumes are eligible for volume restore, while others are not. If multiple restore methods are used to restore the same file, the preview section will be similar to the following:

```
The following files will be restored via a combination of:
[storage side file system restore and storage side system restore]
```

The Analysis section

The Analysis section presents the reasons why some restore mechanisms will not be or were not used. You

can use this information to determine what is required to enable more efficient restore mechanisms.

The following example shows an Analysis section:

```
Analysis:

The following reasons prevent certain files from being
restored completely via: storage side full file system restore
    * LUNs present in snapshot of volume fas960:
         \vol\disks may not be consistent when reverted:
         [fas960:\vol\disks\DG4D1.lun]
      Mapped LUNs in volume fas960:\vol\disks
         not part of the restore scope will be reverted: [DG4D1.lun]

Files to restore:
      E:\disks\sysaux.dbf
      E:\disks\system.dbf
      E:\disks\undotbs1.dbf
      E:\disks\undotbs2.dbf

* Reasons denoted with an asterisk (*) are overridable.
```

In the example, you can override the first failure either from the command-line interface (CLI), or by selecting **Override** in the graphical user interface (GUI). The second failure about mapped LUNs in the volume is mandatory and not overridable.

You can resolve checks by doing the following:

- To resolve a mandatory check failure, change the environment so that the check will pass.

- To resolve an overridable check failure, you can change the environment, or override the check.

  However, you must be careful because overriding the check can result in undesired consequences.

Preview backup restore information

You can preview information about a backup restore process before it occurs to see information about restore eligibility that SnapManager for SAP found on your backup. SnapManager analyzes data on your backup to determine whether the restore process can be completed successfully.

About this task

The restore preview provides the following information:

- Which restore mechanism (storage-side file system restore, storage-side file restore, or host-side file copy restore) can be used to restore each file.

- Why more efficient mechanisms were not used to restore each file, when you specify the `-verbose` option.

If you specify the `-preview` option in the `backup restore` command, SnapManager does not restore anything, but lists the files to be restored and indicates how they will be restored.

> ℹ️ You can preview all types of restore mechanisms. The preview shows information about up to 20 files.

Steps

1. Enter the following command:

   smsap backup restore -profile *profile_name* -label *label* -complete -preview -verbose

   Example

   For example, enter:

   ```
   smsap backup restore -profile targetdb1_prof1
         -label full_bkup_sales_nov_08  -complete -preview -verbose
   ```

   The following example shows the files ready to be restored and lists the different methods used for each:

   ```
   The following files will be restored via storage side full file system
   restore:
   E:\disks\sysaux.dbf
   E:\disks\system.dbf

   The following files will be restored via host side file copy restore:
   E:\disks\undotbs1.dbf
   E:\disks\undotbs2.dbf
   ```

2. Review any reasons why other restore processes cannot be used.

3. Begin the restore operation without the `-preview` option, if only reasons that are overridable are displayed.

   You can still override non-mandatory checks.

Restore backups on primary storage

You can use the `backup restore` command to restore a database backup on primary storage.

About this task

You can use the `backup restore` command options to specify whether SnapManager should restore all or part of the backup. SnapManager also allows you to restore control files along with the data files or tablespaces from the backups in a single user operation. You can include `-controlfiles` with `-complete` to restore control files along with tablespaces and data files.

You can select one of the following options to restore the backup:

If you want to restore…	Use…
The entire backup with all tablespaces and data files	`-complete`
The list of specific tablespaces	`-tablespaces`
Specific data files	`-files`
The control files only	`-controlfiles`
Tablespaces, data files, and control files	`-complete -controlfiles`

You can also restore the backup from an alternate location by specifying `-restorespec`.

If you include `-recover`, you can recover the database to:

- The last transaction that occurred in the database (all logs)
- A specific date and time
- A specific Oracle System Change Number (SCN)
- The time of the backup (no logs)
- Restore only

ⓘ Both date and time recovery and the SCN recovery are point-in-time recoveries.

SnapManager (3.2 or later) provides the ability to recover the restored database backups automatically by using the archive log files. Even if the archive log files are available in the external location, if you specify the `-recover-from-location` option, SnapManager uses the archive log files from the external location to recover the restored database backups.

When you specify the external archive log locations for the recovery of the restored backups, you must ensure that you include the external location names in uppercase. In the file system, all the folders and subfolders names must be in uppercase, because Oracle translates the destination path to uppercase and expects the external destination paths, folder names, and subfolder names to be in uppercase. If you specify the external archive log destination paths in lowercase, Oracle might not be able to identify the specified path, and fails to restore the database.

SnapManager provides the external location to Oracle. But, Oracle does not identify the files from the external destination. This behavior is noticed in flash recovery area destination. These are issues with Oracle and the workaround is to always have backup of archive log files in such database layouts.

If any inconsistent SCN or date is provided, then recovery will stop at the last consistent point recovered with the error message `Recovery succeeded, but insufficient`. You have to manually perform recovery to a consistent state.

For recovery when no logs are applied, SnapManager recovers until the last SCN of the last archive log file created during the backup. If the database is consistent until this SCN, then the database will be opened successfully. If the database is not consistent at this point, SnapManager still attempts to open the database,

which will be opened successfully, if the database is already consistent.

> (i)   SnapManager does not support recovering the archive log-only backups.

If the archive log destination is not a Snapshot-capable storage, SnapManager enables you to recover the restored database backups using the profile. Before performing SnapManager operations on non-Snapshot-capable storage, you should add the destinations for `archivedLogs.exclude` in `smsap.config`.

You must ensure that you set the exclude parameter before creating a profile. Only after setting the exclude parameter in the SnapManager configuration file, the profile creation is successful.

If the backup is already mounted, SnapManager does not mount the backup again and uses the already mounted backup. If the backup is mounted by a different user, and if the current user does not have access to the previously mounted backup, other users have to provide the permissions. All the archive log files have read permissions for the groups owners; the current user might not get the permissions, if the backup is mounted by a different user group. The users can give permissions to the mounted archive log files manually and then retry the restore or recovery.

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

Steps

1. Enter the following command:

    smsap backup restore -profile *profile_name* -label *label* -complete -recover -alllogs [-recover-from-location *path* [,*path2*]]-dump-verbose

    Example

    smsap backup restore -profile targetdb1_prof1 -label full_bkup_sales_nov_08 -complete -recover -alllogs -verbose

2. To restore data for different scenarios, complete one of the following:

If you want to restore…	Command Example
Complete database without control files and recover to a particular SCN number (3794392). In this case, the current control files exist, but all the data files are damaged or lost. Restore and recover the database from an existing full online backup to a point immediately before that SCN.	**smsap backup restore -profile targetdb1_prof1 -label full_bkup_sales_nov_08 -complete -recover -until 3794392 -verbose**
Complete database without control files and recover up to a date and time.	**smsap backup restore -profile targetdb1_prof1 -label full_bkup_sales_nov_08 -complete -recover -until 2008-09-15:15:29:23 -verbose**

If you want to restore…	Command Example
Complete database without control files and recover up to a data and time. In this case, the current control files exist, but all of the data files are damaged or lost or a logical error occurred after a specific time. Restore and recover the database from an existing full online backup to a date and time immediately before the point of failure.	`smsap backup restore -profile targetdb1_prof1 -label full_bkup_sales_nov_08 -complete -recover -until "2008-09-15:15:29:23" -verbose`
Partial database (one or more data files) without control files and recover using all available logs. In this case, the current control files exist, but one or more data files are damaged or lost. Restore those data files and recover the database from an existing full online backup using all available logs.	`smsap backup restore -profile targetdb1_prof1 -label full_bkup_sales_nov_08 -files E:\disks\s02.dbf E:\disks\sales03.dbf E:\disks\sales04.dbf -recover -alllogs -verbose`
Partial database (one or more tablespaces) without control files and recover using all available logs. In this case, the current control files exist, but one or more tablespaces are dropped or one of more data files belonging to the tablespace are damaged or lost. Restore those tablespaces and recover the database from an existing full online backup using all available logs.	`smsap backup restore -profile targetdb1_prof1 -label full_bkup_sales_nov_08 -tablespaces users -recover -alllogs -verbose`
Only control files and recover using all available logs. In this case, the data files exist, but all control files are damaged or lost. Restore just the control files and recover the database from an existing full online backup using all available logs.	`smsap backup restore -profile targetdb1_prof1 -label full_bkup_sales_nov_08 -controlfiles -recover -alllogs -verbose`
Complete database without control files and recover using the backup control files and all available logs. In this case, all data files are damaged or lost. Restore just the control files and recover the database from an existing full online backup using all available logs.	`smsap backup restore -profile targetdb1_prof1 -label full_bkup_sales_nov_08 -complete -using-backup-controlfile -recover -alllogs -verbose`
Recover the restored database using the archive log files from the external archive log location.	`smsap backup restore -profile targetdb1_prof1 -label full_bkup_sales_nov_08 -complete -using-backup-controlfile -recover -alllogs -recover-from-location E:\\archive -verbose`

3. Specify external archive log locations by using the `-recover-from-location` option.

Restore files from an alternate location

SnapManager enables you to restore data files and control files from a location other than that of the Snapshot copies in the original volume.

The original location is the location of the file on the active file system at the time of the backup. The alternate location is the location from which a file will be restored.

You can restore from an alternate location to restore the data files from an intermediate file system to an active file system.

Recovery is automated by SnapManager. When recovering files from external locations, SnapManager uses the `recovery automatic from location` command.

SnapManager also uses Oracle Recovery Manager (RMAN) to recover files. The files to be recovered should be recognizable by Oracle. The file names should be in the default format. When recovering from flash recovery area, SnapManager provides the translated path to Oracle. Oracle though, does not recover from the flash recovery area because it cannot generate the correct file name. Ideally, flash recovery area is a destination that is intended to work with RMAN.

Restore backups from an alternate location overview

To restore a database backup from an alternate location, use the following major steps, each of which is further described in this section.

- Do one of the following, depending on your database layout and what needs to be restored:
  - Restore the required data files from tape, SnapVault, SnapMirror, or any other media to any file system mounted on the database host.
  - Restore the required file system and mount it on the database host.
  - Connect to the required raw devices that exist in the local host.
- Create a restore specification Extensible Markup Language (XML) file that includes the mappings that SnapManager requires to restore from the alternate location to the original location. Save the file in a location that SnapManager can access.
- Use SnapManager to restore and recover the data using the restore specification XML file.

Restoration of the data from files

Before you restore from an alternate location, you need to restore the necessary files from any storage media and restore the files from applications like SnapVault or SnapMirror to a file system mounted on the local host.

You can use the restore from an alternate location operation to copy the files from an alternate file system to an active file system.

You need to specify the alternate locations from which to restore the original files by creating a restore specification.

Restoration of data from the file system

Before you restore data from an alternate location, you must restore the necessary file system and mount it on the local host.

You can invoke the restore operation from an alternate location to copy the files from alternate file systems to active file systems.

To perform this operation, you must specify the alternate mount points from which to restore the original mount points and the original Snapshot copy names by creating a restore specification file.

> ⓘ The Snapshot copy name is a necessary component because the same file system might be snapped multiple times in a single backup operation (for example, once for the data files and once for the log file).

Create restore specifications

The restore specification file is an XML file that contains the original and alternate locations from which the file can be restored. SnapManager uses this specification file to restore files from the specified location.

About this task

You can create the restore specification file by using any text editor. You must use a .xml extension for the file.

Steps

1. Open a text file.

2. Enter the following:

   `<restore-specification xmlns="http://www.netapp.com">`

3. Enter any file mapping information using the format shown in the following example:

   ```
   <file-mapping>
       <original-location>E:\disks\sysaux.dbf</original-location>
       <alternate-location>E:\disks\sysaux.dbf</alternate-location>
   </file-mapping>
   ```

   File mapping specifies where a file is restored from. The original location is the location of the file on the active file system at the time of backup. The alternate location is the location from where the file is restored.

4. Enter any mounted file system mapping information using the format shown in the example:

   ```
   <mountpoint-mapping>
       <original-location>E:\disks\sysaux.dbf</original-location>
       <snapname>snapname</snapname>
       <alternate-location>E:\disks\sysaux.dbf</alternate-location>
   </mountpoint-mapping>
   ```

   Mountpoint refers to directory path `C:\myfs`. The mountpoint mapping specifies the mountpoint from which the files are restored. The original location is the location of the mountpoint in the active file system at the time of backup. The alternate location is the mountpoint from which the files in the original location

are restored. The `snapname` is the name of the Snapshot copy in which the original files were backed up.

> ℹ️ The Snapshot copy name is a necessary component because the same file system can be used multiple times in a single backup operation (for example, once for the data files and once for the logs).

5. Enter the following:

   `</restore-specification>`

6. Save the file as a .xml file and close the specification.

Restore backups from an alternate location

You can restore from an alternate location to restore the data files from an intermediate file system to an active file system.

What you'll need

- Create a restore specification XML file and specify the type of restore method you want to use.

About this task

You can use the `smsap backup restore` command and specify the restore specification XML file you created to restore the backup from an alternate location.

Step

1. Enter the following command:

   `smsap backup restore -profile profile -label label -complete-alllogs -restorespec restorespec`

Cloning database backup

If you clone a database, you can perform tasks such as test an upgrade to a database without affecting the database in production, duplicate a master installation to several training systems, or duplicate a master installation as a base installation to other servers, which have similar requirements.

You can perform the following tasks related to cloning:

- Clone a database from an existing backup.
- Clone a database in its current state, which enables you to create the backup and the clone in one procedure.
- Clone a database and use custom plug-in scripts, which run before or after the clone operation.
- Clone a database to the same host on which the database resides.
- Clone a database by using archive log files from the external archive log location.
- Clone a database to an alternate host.

- View a list of clones.

- View detailed clone information.

- Delete clones.

What Cloning is

You can clone a database to create an exact replica of the original database. You can create the clone from a full backup or from the current state of the database.

Some of the advantages of creating a clone by using SnapManager are as follows:

Advantages	Details
Speed	The SnapManager clone operation uses the FlexClone feature available with Data ONTAP. This enables you to quickly clone large data volumes.
Space efficiency	When you create a clone by using SnapManager, space is needed only for the changes between the backup and the clone. A SnapManager clone is a writable Snapshot copy of the original database and can grow as needed. In contrast, a physical clone of the database requires that you have enough space available to copy the entire database.
Virtual copy	You can use the cloned database as if it were the original database. For example, you can use a clone for testing, platform and update checks, multiple simulations against a large data set, and remote office testing and staging. Changes to the clone do not affect the original database. After the database is cloned, the cloned database is fully operational.
Simplicity	You can clone a database to any host by using SnapManager commands.

You must ensure that the following prerequisites are met before a database can be cloned:

- Delete the `spfile<SID>.ora` file from `$ORACLE_HOME\database`.

- Delete the `init<SID>.ora` file from `$ORACLE_HOME\database`.

- Delete Oracle dump destinations that are specified in the clone specification file.

- Delete the Oracle control files that are specified in the clone specification file.

- Delete the Oracle redo log files that are specified in the clone specification file.

You must give the clone a new system identifier. You cannot simultaneously run two databases with the same system identifier on the same host. You can have a clone on a different host using the same system identifier. You can either give the clone a label or let SnapManager create a label by using the system identifier, date, and time the clone was created.

When you enter a label, you should not include spaces or special characters.

As part of the cloning process, SnapManager creates the necessary Oracle files and parameters for the cloned database. An example of a necessary Oracle file is `init<SID>.ora`.

When you clone a database, SnapManager creates a new `init<SID>.ora` file for the database in the `$ORACLE_HOME\database` directory.

You can clone a database backup to the host in which the database resides or to an alternate host.

If the database you cloned was using a `spfile`, SnapManager creates an `spfile` for the clone. It places this file in the `$ORACLE_HOME\database` directory and creates the directory structure for the diagnostic files. The file name is `spfile <SID>.ora`.

Cloning methods

You can clone a database using one of two methods. The method you choose affects the clone create operation.

The following table describes the cloning methods and their effect on the `clone create` operation and its `-reserve` option. A LUN can be cloned using either method.

Cloning method	Description	clone create -reserve
LUN cloning	A new clone LUN is created within the same volume.	When `-reserve` for a LUN is set to `yes`, space is reserved for the full LUN size within the volume.
Volume cloning	A new FlexClone is created and the clone LUN exists within the new clone volume. Uses FlexClone technology.	When `-reserve` for a volume is set to `yes`, space is reserved for the full volume size within the aggregate.

Create clone specifications

SnapManager for SAP uses a clone specification XML file, which includes the mappings, options, and parameters for use in the clone operation. SnapManager uses this information to determine where to place the files it clones and how to handle diagnostic information, control files, parameters, and so on.

About this task

You can create the clone specification file by using the SnapManager graphical user interface (GUI), command-line interface (CLI), or a text editor.

When you create the clone specification file by using a text editor, you must save it as a `.xml` file. You can use this XML file for other clone operations.

You can also create a clone specification template and then customize it. You can use the `smsap clone template` command or in the GUI, use the Clone wizard.

SnapManager for SAP adds a version string to any clone specification template that it generates. SnapManager for SAP assumes the latest version for any clone specification file that lacks a version string.

If you want to perform remote cloning, do not change the default locations of the data files, redo log files, and control files in the clone specification file. If you change the default location, SnapManager fails to create the clone or creates the clone on a database that does not support Snapshot capability. Therefore, the automatic creation of profile fails.

> ⓘ Though mount point and ASM disk group information are editable from the GUI, you can only change the file name and not the file locations.

You can execute a task multiple times, either with the same or different parameter and value combinations.

SAP uses specific Oracle settings for database configuration. You can find these settings in `init<SID>.ora` located at `$ORACLE_HOME\database`. You should include these in your clone specification.

Steps

1. Open a text file and enter text as shown in the following example:

   Example

   ```
   <clone-specification xmlns="http://www.example.com>
        <storage-specification/>
        <database-specification/>
   </clone-specification>
   ```

2. In the storage specification component, enter the mount points for the data files.

   The storage specification lists the locations for the new storage created for the clone such as data file mount points and raw devices. These items must be mapped from the source to the destination.

   The following example shows a single NFS mount point on an NFS-attached storage:

   Example

   ```
   <mountpoint>
        <source>\oracle\<SOURCE SID>_sapdata</source>
        <destination>\oracle\<TARGET SID>_sapdata</destination>
   </mountpoint>
   ```

3. In the database specification component, identify the control file information as a list of the control files that you want created for the clone.

   The database specification specifies the database options for the clone such as control files, redo logs, archive logs, and Oracle parameters.

   The following example displays the control file syntax that you use in clone specification:

   Example

```
<controlfiles>
     <file>\oracle\<TARGET SID>\origlogA\cntrl\cntrl<TARGET
SID>.dbf</file>
     <file>\oracle\<TARGET SID>\origlogB\cntrl\cntrl<TARGET
SID>.dbf</file>
     <file>\oracle\<TARGET SID>\sapdata1\cntrl\cntrl<TARGET
SID>.dbf</file>
</controlfiles>
```

4. Specify the redo log structure for the clone.

   The following example displays the redo log directory structure for cloning:

   Example

```
<redologs>
     <redogroup>
        <file>E:\oracle\<TARGET SID>\origlogA\log_g11m1.dbf</file>
        <file>E:\oracle\<TARGET SID>\mirrlogA\log_g11m2.dbf</file>
        <number>1</number>
        <size unit="M">100</size>
     </redogroup>
     <redogroup>
        <file>E:\oracle\<TARGET SID>\origlogB\log_g12m1.dbf</file>
        <file>E:\oracle\<TARGET SID>\mirrlogB\log_g12m2.dbf</file>
        <number>2</number>
        <size unit="M">100</size>
     </redogroup>
     <redogroup>
        <file>E:\oracle\<TARGET SID>\origlogA\log_g13m1.dbf</file>
        <file>E:\oracle\<TARGET SID>\mirrlogA\log_g13m2.dbf</file>
        <number>3</number>
        <size unit="M">100</size>
     </redogroup>
     <redogroup>
        <file>E:\oracle\<TARGET SID>\origlogB\log_g14m1.dbf</file>
        <file>E:\oracle\<TARGET SID>\mirrlogB\log_g14m2.dbf</file>
        <number>4</number>
        <size unit="M">100</size>
     </redogroup>
</redologs>
```

5. Specify the Oracle parameters that should be set to different values in the cloned database. If you are using Oracle 10, you must specify the following parameters:

- Background dump
- Core dump
- User dump
- **Optional:** Archive logs

> ⓘ If the parameter values are not set correctly, the clone operation is stopped and you receive an error message.

If you do not specify the location where archive logs are stored, SnapManager creates the clone in `noarchivelog` mode. SnapManager copies this parameter information into the `init.ora` file of the clone.

The following example displays the parameter syntax that you use in clone specification:

Example

```
<parameters>
     <parameter>
          <name>log_archive_dest</name>
          <value>LOCATION=>E:\oracle\<TARGET SID>\oraarch</value>
     </parameter>
     <parameter>
          <name>background_dump_dest</name>
          <value>E:\oracle\<TARGET SID>\saptrace\background</value>
     </parameter>
     <parameter>
          <name>core_dump_dest</name>
          <value>E:\oracle\<TARGET SID>\saptrace\background</value>
     </parameter>
     <parameter>
     <name>user_dump_dest</name>
     <value>E:\oracle\<TARGET SID>\saptrace\usertrace</value>
     </parameter>
</parameters>
```

You can use a default value by using a default element within the parameter element. In the following example, the `os_authentication_prefix` parameter will take the default value because the default element is specified:

Example

```
<parameters>
    <parameter>
        <name>os_authent_prefix</name>
        <default></default>
    </parameter>
</parameters>
```

You can specify an empty string as the value for a parameter by using an empty element. In the following example, the `os_authentication_prefix` will be set to an empty string:

Example

```
<parameters>
    <parameter>
        <name>os_authent_prefix</name>
        <value></value>
    </parameter>
</parameters>
```

> (i) You can use the value from the source database's `init.ora` file for the parameter by not specifying any element.

Example

If a parameter has multiple values, then you can provide the parameter values separated by commas. For example, if you want to move the data files from one location to another, then you can use the `db_file_name_convert` parameter and specify the data file paths separated by commas as seen in the following example:

Example

If you want to move the log files from one location to another, then you can use the `log_file_name_convert` parameter and specify the log file paths separated by commas, as seen in the followowing example:

1. Optional: Specify arbitrary SQL statements to execute against the clone when it is online.

You can use the SQL statements to perform tasks such as re-creating the `temp files` in the cloned database.

> (i) You must ensure that a semicolon is not included at the end of the SQL statement.

The following is a sample SQL statement that you execute as part of the clone operation:

```
<sql-statements>
    <sql-statement>
      ALTER TABLESPACE TEMP ADD
      TEMPFILE 'E:\path\clonename\temp_user01.dbf'
      SIZE 41943040 REUSE AUTOEXTEND ON NEXT 655360
      MAXSIZE 32767M
    </sql-statement>
</sql-statements>
```

Clone specification example

The following example displays the clone specification structure, including both the storage and database specification components, for a Windows environment:

```
<clone-specification xmlns="http://www.example.com>

<storage-specification>
    <storage-mapping>
        <mountpoint>
            <source>D:\oracle\<SOURCE SID>_sapdata</source>
            <destination>D:\oracle\<TARGET SID>_sapdata</destination>
        </mountpoint>
    </storage-mapping>
</storage-specification>

<database-specification>
    <controlfiles>
        <file>D:\oracle\<TARGET SID>\origlogA\cntrl\cntrl<TARGET
SID>.dbf</file>
        <file>D:\oracle\<TARGET SID>\origlogB\cntrl\cntrl<TARGET
SID>.dbf</file>
        <file>D:\oracle\<TARGET SID>\sapdata1\cntrl\cntrl<TARGET
SID>.dbf</file>
    </controlfiles>

    <redologs>
        <redogroup>
            <file>D:\oracle\<TARGET SID>\origlogA\log_g11m1.dbf</file>
            <file>D:\oracle\<TARGET SID>\mirrlogA\log_g11m2.dbf</file>
            <number>1</number>
            <size unit="M">100</size>
        </redogroup>
        <redogroup>
            <file>D:\oracle\<TARGET SID>\origlogB\log_g12m1.dbf</file>
            <file>D:\oracle\<TARGET SID>\mirrlogB\log_g12m2.dbf</file>
```

```
            <number>2</number>
            <size unit="M">100</size>
        </redogroup>
        <redogroup>
            <file>D:\oracle\<TARGET SID>\origlogA\log_g13m1.dbf</file>
            <file>D:\oracle\<TARGET SID>\mirrlogA\log_g13m2.dbf</file>
            <number>3</number>
            <size unit="M">100</size>
        </redogroup>
        <redogroup>
            <file>D:\oracle\<TARGET SID>\origlogB\log_g14m1.dbf</file>
            <file>D:\oracle\<TARGET SID>\mirrlogB\log_g14m2.dbf</file>
            <number>4</number>
            <size unit="M">100</size>
        </redogroup>
    </redologs>

    <parameters>
        <parameter>
            <name>log_archive_dest</name>
            <value>LOCATION=>D:\oracle\<TARGET SID>\oraarch</value>
        </parameter>
        <parameter>
            <name>background_dump_dest</name>
            <value>D:\oracle\<TARGET SID>\saptrace\background</value>
        </parameter>
        <parameter>
            <name>core_dump_dest</name>
            <value>D:\oracle\<TARGET SID>\saptrace\background</value>
        </parameter>
        <parameter>
            <name>user_dump_dest</name>
            <value>D:\oracle\<TARGET SID>\saptrace\usertrace</value>
        </parameter>
    </parameters>
  </database-specification>
</clone-specification>
```

Clone databases and using custom plug-in scripts

SnapManager provides a method for using your custom scripts before and after a clone
operation occurs. For example, you might have created a custom script that validates a
clone database SID and ensures the SID is allowed by your naming policy. Using the
SnapManager clone plug-in, you can include your custom scripts and have them run
automatically before or after a SnapManager clone operation.

Steps

1. View sample plug-in scripts.

2. Create a script from scratch or modify one of the sample plug-in scripts.

   Create your custom script according to SnapManager plug-in script guidelines.

3. Place your custom script in a specified directory location.

4. Update the clone specification XML file and include information about your custom script that should be used during the cloning process.

5. Using a SnapManager command, verify that the custom scripts are operational.

6. When you initiate the clone operation, include the script name and optional parameters.

Clone databases from backups

You can clone a database from a backup by using the `clone create` command.

About this task

You must first create a clone specification file for the database. SnapManager creates the clone based on the information in this specification file.

After you clone a database, you might want to update your `tnsnames.ora` files on your client machines with the new cloned database connection information. The `tnsnames.ora` files are used to connect to an Oracle instance without having to specify the full database information. SnapManager does not update the `tnsnames.ora` files.

SnapManager always creates a backup including archive log files, if you are using the profile created with `-include-with-online-backups`. SnapManager allows you to clone only the full database backups.

SnapManager (3.2 or later) allows you to clone the backups containing the data files and archive log files.

If the archive log is available from an external location, you can specify the external location during cloning for recovering the cloned database to a consistent state. You must ensure that the external location is accessible by Oracle. Cloning of the archive log-only backups is not supported.

Though the archive log backup is created along with the online partial backup, you cannot create a database clone by using this backup.

When you specify the external archive log locations for recovering the cloned database to a consistent state, you must ensure that you include the external location names completely in uppercase. In the file system, the names of all the folders and subfolders must be in uppercase because the Oracle database translates the destination path to uppercase and expects the external destination paths, folder names, and subfolder names to be in uppercase. If you specify the external archive log destination paths in lowercase, database might not be able to identify the specified path, and fails to recover the cloned database.

You can clone the database backup from the external archive log file location only for a stand-alone database.

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

Cloning datafile backup without archive log backup

When the data files backup does not include the archive log backup, SnapManager for SAP clones the database based on the System Change Number (SCN) recorded during the backup. If the cloned database cannot be recovered, the `Archived log file for thread <number> and change <SCN> required to complete recovery` error message is displayed, even though SnapManager for SAP continues to clone the database, and finally succeeds in creating the clone.

When cloning using the data files backup without including the archive log backup, SnapManager recovers the cloned database until the last archive log SCN, which is recorded during the backup.

Steps

1. Create a clone specification file.

2. To create a clone, enter the following command:

   ```
   smsap clone create-backup-labelbackup_name-newsidnew_sid-labelclone_label-
   profileprofile_name-clonespecfull_path_to_clonespecfile [-taskspec _taskspec_]
   [-recover-from-location] path1 [,<path2>...]] [-dump]
   ```

Clone databases in the current state

You can create a backup and a clone of the database from the current state of the database by using a single command.

About this task

When you specify the profile with the `-current` option, SnapManager first creates a backup and then a clone from the current state of the database.

In the profile setting, if you have enabled the backup of data files and archive logs together for cloning, whenever you back up the online data files, the archive logs are also backed up. While cloning the database, SnapManager creates the data files backup along with the archive log backup and creates the database clone. If the archive log backup is not included, SnapManager does not create the archive log backup and therefore cannot create the clone of the database.

Step

1. To clone the database in its current state, enter the following command:

   **smsap clone create -profileprofile_name-current -labelclone_name
   -clonespecclonespec.xml**

   This command takes a full automatic backup (generating the backup label) and immediately makes a clone from that backup, using an existing clone specification that you want to use.

   > ⓘ You can specify the `-dump` option as an optional parameter to collect the dump files after the successful or failed operations. The dump is collected for both the backup and clone operations.

Clone database backups without resetlogs

SnapManager enables you to perform flexible cloning so that you can recover the cloned database manually to a desired point in time without opening the database by using

resetlogs. You can also manually configure the cloned database as a Data Guard Standby database.

About this task

When you can select the `-no-resetlogs` option while creating the clone, SnapManager performs the following activities to create the cloned database:

Steps

1. Executes the preprocessing task activity, if specified, before starting the clone operation

2. Creates the cloned database with the user-specified SID

3. Executes the SQL statements issued against the cloned database.

   Only the SQL statements that can be executed in mount state are successfully executed.

4. Executes the post-processing task activity, if specified.

What tasks you need to do to recover the cloned database manually

Steps

1. Mount the archive log backups and recover the cloned database manually by using the archive log files from the mounted path.

2. After performing manual recovery, open the recovered cloned database with `-resetlogs` option.

3. Create temporary tablespaces, if required.

4. Run the DBNEWID utility.

5. Grant sysdba privilege to the credentials of the cloned database.

While cloning the database backups using the `-no-resetlogs` option, SnapManager leaves the cloned database in the mounted state for manual recovery.

> (i) The cloned database created with the `-no-resetlogs` option is not a complete database. Therefore you must not perform any SnapManager operations on this database, though SnapManager does not restrict you from performing any operations.

If you do not specify the `-no-resetlogs` option, SnapManager applies the archive log files, and opens the database with resetlogs.

Step

1. Enter the following command:

   smsap clone create -profile *profile_name* [-backup-label *backup_name* | -backup -id *backup_id* | current] -newsid *new_sid* -clonespec *full_path_to_clonespecfile* -no-resetlogs

   If you try to specify both `-no-resetlogs` and `recover-from-location` options, SnapManager does not allow you to specify both these options together, and displays the error message: `SMSAP-04084: You must specify either one of the options: -no-resetlogs or -recover-from -location`.

Example

```
smsap clone create -profile product -backup-label full_offline -newsid
PROD_CLONE -clonespec prod_clonespec.xml -label prod_clone-reserve -no
-reset-logs
```

Considerations for cloning a database to an alternate host

Before you can clone to a host other than the one on which the database resides, there are some requirements that must be met.

The following table shows the source and target host setup requirements:

Prerequisite set up	Requirement
Architecture	Must be the same on both the source and target hosts
Operating system and version	Must be the same on both the source and target hosts
SnapManager for SAP	Must be installed and running on both the source and target hosts
Credentials	Must be set for the user to access the target host
Oracle	The same software version must be installed on both the source and target hosts. The Oracle Listener must be running on the target host.
Compatible storage stack	Must be the same on both the source and target hosts
Protocol used to access data files	Must be the same on both the source and target hosts
Domain	Both remote host and the host on which the database resides must be in the domain and not in the workgroup

Clone a database to an alternate host

You can use the `clone create` command to clone a database backup on an alternate host.

What you'll need

- Create a profile or have an existing profile.
- Create a full backup or have an existing database backup.

* Create a clone specification or have an existing clone specification.

Step

1. To clone a database to an alternate host, enter the following command:

   ```
   smsap clone create -backup-label backup_label_name -newsid _new_sid -host
   target_host -label clone_label -comment comment_text -profile profile_name
   -clonespec full_path_to_clonespecfile
   ```

   Oracle does not let you run two databases with the same SID simultaneously on the same host. Because of this, you must supply a new SID for each clone. However, you can have a database on another host with the same SID.

View a list of clones

You can view a list of clones associated with a particular profile.

About this task

The list includes the following information about the clones in a profile:

* The ID for the clone
* Status of the clone operation
* Oracle SID for the clone
* Host on which the clone resides
* Label for the clone

If you specify the `-verbose` option, the output also shows the comments entered for the clone.

Step

1. To display a list of all clones for a profile, enter the following command:

   ```
   smsap clone list -profile profile_name [-quiet | -verbose]
   ```

View detailed clone information

You can view detailed information about a specific clone by using the `clone show` command.

About this task

The `clone show` command displays the following information:

* Clone system identifier and clone ID
* Clone operation status
* Clone create start and end date or time
* Clone label

- Clone comment
- Backup label and ID
- Source database
- Backup start and end time
- Database name, tablespaces, and data files
- Host name and file systems containing data files
- Storage system volumes and Snapshot copies backing the clone

Step

1. Enter the following command:

   smsap clone show -profile *profile_name* [-label *label* | -id *guid*]

Delete clones

You can delete the clones when the size of the Snapshot copy reaches between 10% and 20% of the backup. This also guarantees that the clone has the most current data.

About this task

The label is the unique identifier for each clone in a profile. You can use the clone label or ID, but not the system identifier (SID) to delete the clone.

(i) The clone SID and the clone label are not the same.

When you are deleting a clone, the database must be running. Otherwise, many files and directories for the existing clone will not be deleted, resulting in more work before another clone can be created.

The directories specified for certain Oracle parameters in the clone are destroyed when the clone is deleted, and should only contain data for the cloned database: Archive Log Destinations, Background, Core, and User Dump Destinations. The audit files are not deleted.

(i) You cannot delete a clone when the clone is used in other operations.

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

Step

1. Enter the following command:

   smsap clone delete -profile *profile_name* [-label *label* | -id *guid*] [-force][-dump][-quiet]|[-verbose]

Example

```
smsap clone delete -profile targetdb1_prof1 -label sales0908_clone1
```

Introduction to data protection in SnapManager

SnapManager supports data protection to protect the backups on the secondary or tertiary storage systems. You must set up SnapMirror and SnapVault relationships between the source and the destination volumes.

> ⓘ   Backups created using BR*Tools cannot be protected using SnapManager for SAP.

You can use the post-backup scripts from both command-line interface (CLI) and graphical user interface (GUI) to protect the backups on a primary storage system.

How SnapManager retains backups on the local storage

SnapManager enables you to create backups that meet retention policies, which specify how many successful backups on local storage should be retained. You can specify the number of successful backups that should be retained in the profile for a given database.

You can create backups for the following:

- 10 days of daily backups on primary storage
- 2 months of monthly backups on primary storage
- 7 days of daily backups on secondary storage
- 4 weeks of weekly backups on secondary storage
- 6 months of monthly backups on secondary storage

For each profile in SnapManager, you can change the values for the following nonlimited retention classes:

- Hourly
- Daily
- Weekly
- Monthly

SnapManager determines whether a backup should be retained by considering both the retention count (for example, 15 backups) and the retention duration (for example, 10 days of daily backups). A backup expires when its age exceeds the retention duration set for its retention class or the number of backups exceeds the retention count. For example, if the backup count is 15 (SnapManager has taken 15 successful backups) and the duration requirement is set for 10 days of daily backups, the five oldest successful eligible backups expire.

After a backup expires, SnapManager either frees or deletes the expired backup. SnapManager always retains the last backup taken.

SnapManager counts only the number of successful backups for the retention count and does not consider the following:

Backups not included in the retention count	Additional details
Failed backups	SnapManager retains the information about successful and unsuccessful backups. Although unsuccessful backups require only minimal space in the repository, you might want to delete them. Unsuccessful backups remain in the repository until you delete them.
Backups designated to be retained on an unlimited basis or backups for a different retention class	SnapManager does not delete backups designated to be retained on an unlimited basis. Additionally, SnapManager considers only those backups in the same retention class (for example, SnapManager considers only the hourly backups for the hourly retention count).
Backups mounted from local storage	When Snapshot copies are mounted, they are also cloned and so are not considered eligible for retention. SnapManager cannot delete the Snapshot copies if they are cloned.
Backups that are used to create a clone on local storage	SnapManager retains all the backups that are used to create clones, but does not consider them for the backup retention count.

SnapManager provides a default retention count and duration for each retention class. For example, for the hourly retention class count, SnapManager, by default, retains four hourly backups. You can override these defaults and set the values when creating or updating the profile or change the default values for retention count and duration in the `smsap.config` file.

When local backups expire based on their retention policy, the backups are deleted.

In an archivelog-only backup operation, SnapManager does not archive the redo log files, unlike in the online database backup process. You must add a pretask script to archive the redo log files before performing the archivelog-only backup operation. The pretask script must run the `alter system switch logfile` command.

The following example shows the actions that SnapManager takes on various types of backups, based on a three-daily-backups retention policy (with the count set to retain 3):

Backup date	Status	Retention policy action taken	Explanation
5/10	Successful	Keep	This is the most recent successful backup, so it will be kept.

Backup date	Status	Retention policy action taken	Explanation
5/9	Successful, cloned	Skip	SnapManager does not consider backups used for cloning in the retention policy count. This backup is omitted from the count of successful backups.
5/8	Successful, mounted	Skip	SnapManager does not consider mounted backups in the retention policy count. This backup is omitted from the count of successful backups.
5/7	Failed	Skip	Failed backups are not counted.
5/5	Successful	Keep	SnapManager keeps this second successful daily backup.
5/3	Successful	Keep	SnapManager keeps this third successful daily backup.
5/2	Successful	Delete	SnapManager counts this successful backup, but after SnapManager reaches three successful daily backups, this backup is deleted.

Related information

Documentation on the NetApp Support Site

Considerations for performing data protection

You must be aware of certain considerations for performing data protection.

- To perform clone or restore operations from secondary systems, you must mount the destination volume in the namespace and export it properly.

- You must disable the SnapDrive configuration parameter `check-export-permission-nfs-clone` by setting the value to **off**.

  The SnapDrive for UNIX documentation on the NetApp Support Site contains additional information about the `check-export-permission-nfs-clone parameter`.

- You must configure the SnapMirror relationship for the requested secondary storage volumes in the secondary storage system.
- You must configure the SnapVault relationship for the requested secondary storage qtrees in the secondary storage system for Data ONTAP operating in 7-Mode.
- You must define a policy and rules for the user-defined SnapMirror label if you are using SnapVault post-script for clustered Data ONTAP.

    SnapVault post-script supports clustered Data ONTAP volumes and the SnapMirror relation types DP and XDP. The ONTAP documentation on the NetApp Support Site contains information about configuring SnapMirror and SnapVault.

Documentation on the NetApp Support Site

Licences required for data protection in SnapManager

You must ensure that licenses required for data protection are installed and enabled on the primary and secondary storage systems.

Primary storage systems receive the latest transaction updates for the Oracle database, store the data, and provide local backup protection of the database. The primary storage system also maintains database data files, log files, and control files. Secondary storage systems act as remote storage for the protected backups.

For data protection, the following licenses must be installed and enabled on primary storage systems:

> ⓘ If you want to enable data protection on the secondary storage systems, you must also install and enable these licenses on the secondary storage systems.

- Either Data ONTAP operating in 7-Mode (7.3.1 or later) or clustered Data ONTAP (8.2 or later)
- SnapVault
- SnapRestore
- SnapMirror
- FlexClone is required for cloning.

    FlexClone is also, required for Storage Area Network (SAN) only if SnapDrive is configured to use FlexClone in SAN environments.

- The appropriate protocol, for example, Internet Small Computer System Interface (iSCSI), or Fibre Channel (FC)

Protecting database backups by using postscripts

SnapManager enables you to protect database backups by using postscripts when SnapMirror or SnapVault relationship is established between the primary and secondary storage systems. You can use the default postscripts for post-processing activity of the backup operation from both the SnapManager CLI and GUI.

You can use the following default postscripts available at `default_install_directory\plugins\backup\create\post`:

- The SnapMirror postscript, `Mirror_the_backup.cmd` if you are using Data ONTAP operating in 7-Mode

- The SnapVault postscript, `Vault_the_backup.cmd` if you are using Data ONTAP operating in 7-Mode

- The SnapMirror postscript, `Mirror_the_backup_cDOT.cmd` if you are using clustered Data ONTAP

- The SnapVault postscript, `Vault_the_backup_cDOT.cmd` if you are using clustered Data ONTAP

For more information, see the `readme.txt` available at `default_install_directory\plugins`.

SnapManager 3.1 or earlier provides preprocessing or post-processing scripts only for clone operations. SnapManager 3.2 or later provides preprocessing and post-processing scripts for backup and restore operations. You can use these scripts to run before or after the backup or restore operations.

> (i) The scripts are provided for reference only. They have been tested with SnapDrive 7.0 for Windows or later or later but may not work in all environments. You should customize the scripts based on your secondary protection requirements. The scripts do not work with SnapDrive versions earlier than 6.2.

Sample post-scripts

You can refer the sample scripts and create customized scripts based on your environment to perform mirroring and vaulting. The sample scripts are available at default_install_directory\plugins\backup\create\post.

Mirror_the_backup.cmd

If you are using Data ONTAP operating in 7-Mode, you can use this sample script to mirror the backup. It includes three operations (check, describe, and execute) and executes them at the end of the script. The script also includes error message handling with codes of 0 to 4 and > 4:

```
@echo off
REM $Id:
//depot/prod/capstan/main/src/plugins/windows/examples/backup/create/post/
Mirror_the_backup.cmd#1 $
REM
REM Copyright \(c\) 2011 NetApp, Inc.
REM All rights reserved.
REM
REM
REM This is a sample post-task script to mirror the volumes to the
secondary storage after successful backup operation.
REM|------------------------------------------------------------------
-------------------------------------------|
REM| Pre-requisite/Assumption:
|
REM|   SnapMirror relationship for the requested secondary storage volumes
must be configured in Secondary storage. |
REM|------------------------------------------------------------------
-------------------------------------------|
```

```
REM
REM
REM This script can be used from the SnapManager graphical user interface
(GUI) and command line interface (CLI).
REM
REM To execute the post-task script for the backup operation from
SnapManager GUI, follow these steps:
REM  1. From the Backup wizard > Task Specification page > Post-Tasks tab
> select the post-task scripts from the Available Scripts section.
REM
REM
REM To execute the post-task script for the backup operation from
SnapManager CLI, follow these steps:
REM  1. create a task specification XML file.
REM  2. Add the post-script name in the <post-tasks> tag of the XML file.
REM
REM Example:
REM              <post-tasks>
REM                 <task>
REM                     <name>Mirror the backup</name>
REM                      <description>Mirror the backup</description>
REM                 </task>
REM              <post-tasks>
REM
REM
REM
REM IMPORTANT NOTE: This script is provided for reference only. It has
been tested with SnapDrive 6.3.0 for Windows but may not work in all
environments.  Please review and then customize based on your secondary
protection requirements.
REM
set /a EXIT=0
set name="Mirror the backup"
set description="Mirror the backup"
set parameter=()


if /i "%1" == "-check" goto :check
if /i "%1" == "-execute" goto :execute
if /i "%1" == "-describe" goto :describe

:usage
       echo usage: %0 ^{ -check ^| -describe ^| -execute ^}
       set /a EXIT=99
       goto :exit
```

```
:check
        set /a EXIT=0
        goto :exit

:describe
        echo SM_PI_NAME:%name%
        echo SM_PI_DESCRIPTION:%description%
        echo SM_PRIMARY_MOUNT_POINTS : %SM_PRIMARY_MOUNT_POINTS%
        set /a EXIT=0
        goto :exit

REM - Split the comma-separated PRIMARY_MOUNT_POINTS and Mirror the
PRIMARY_MOUNT_POINTS one-by-one.

:execute
        set /a EXIT=0

        echo "execution started"

        REM FOR %%G IN (%SM_PRIMARY_MOUNT_POINTS%) DO echo %%G

        FOR %%V IN (%SM_PRIMARY_MOUNT_POINTS%) DO sdcli snap update_mirror
-d %%V

        if "%ERRORLEVEL%" NEQ "0" (
            set /a EXIT=4
            exit /b %EXIT%
        )

        echo "execution ended"

        goto :exit

:exit
        echo Command complete.
        exit /b %EXIT%
```

Vault_the_backup.cmd

If you are using Data ONTAP operating in 7-Mode, you can use this sample script to vault the backup. It includes three operations (check, describe, and execute) and executes them at the end of the script. The script also includes error message handling with codes of 0 to 4 and > 4:

```
@echo off
REM $Id:
//depot/prod/capstan/main/src/plugins/windows/examples/backup/create/post/
```

```
Vault_the_backup.cmd#1 $
REM
REM Copyright \(c\) 2011 NetApp, Inc.
REM All rights reserved.
REM
REM
REM This is a sample post-task script to vault the qtrees to the secondary
storage after successful backup operation.
REM|--------------------------------------------------------------------
----------------------------------------|
REM| Pre-requisite/Assumption:
|
REM|   SnapVault relationship for the requested secondary storage qtrees
must be configured in Secondary storage. |
REM|--------------------------------------------------------------------
----------------------------------------|
REM
REM
REM This script can be used from the SnapManager graphical user interface
(GUI) and command line interface (CLI).
REM
REM To execute the post-task script for the backup operation from
SnapManager GUI, follow these steps:
REM  1. From the Backup wizard > Task Specification page > Post-Tasks tab
> select the post-task scripts from the Available Scripts section.
REM
REM
REM To execute the post-task script for the backup operation from
SnapManager CLI, follow these steps:
REM  1. create a task specification XML file.
REM  2. Add the post-script name in the <post-tasks> tag of the XML file.
REM
REM Example:
REM              <post-tasks>
REM                <task>
REM                   <name>Vault the backup</name>
REM                   <description>Vault the backup</description>
REM                </task>
REM              <post-tasks>
REM
REM IMPORTANT NOTE: This script is provided for reference only. It has
been tested with SnapDrive 6.3.0 for Windows but may not work in all
environments.  Please review and then customize based on your secondary
protection requirements.
REM
REM
```

```
REM
REM
REM
set /a EXIT=0
set name="Vault the backup"
set description="Vault the backup"
set parameter=()



if /i "%1" == "-check" goto :check
if /i "%1" == "-execute" goto :execute
if /i "%1" == "-describe" goto :describe

:usage
        echo usage: %0 ^{ -check ^| -describe ^| -execute ^}
        set /a EXIT=99
        goto :exit

:check
        set /a EXIT=0
        goto :exit

:describe
        echo SM_PI_NAME:%name%
        echo SM_PI_DESCRIPTION:%description%
        echo SM_PRIMARY_SNAPSHOTS_AND_MOUNT_POINTS :
%SM_PRIMARY_SNAPSHOTS_AND_MOUNT_POINTS%
        set /a EXIT=0
        goto :exit

REM Split the colon-separated SM_PRIMARY_SNAPSHOTS_AND_MOUNT_POINTS And
SnapVault the mountpoints one-by-one

:execute
        set /a EXIT=0

        echo "execution started"

        FOR %%A IN (%SM_PRIMARY_SNAPSHOTS_AND_MOUNT_POINTS%) DO FOR /F
"tokens=1,2 delims=:" %%B IN ("%%A") DO sdcli snapvault archive  -a %%B
-DS %%C %%B

        if "%ERRORLEVEL%" NEQ "0" (
            set /a EXIT=4
            exit /b %EXIT%
        )
```

```
        echo "execution ended"

        goto :exit

:exit
        echo Command complete.
        exit /b %EXIT%
```

Mirror_the_backup_cDOT.cmd

If you are using clustered Data ONTAP, you can use this sample script to mirror the backup. It includes three operations (check, describe, and execute) and executes them at the end of the script. The script also includes error message handling with codes of 0 to 4 and > 4:

```
@echo off
REM $Id:
//depot/prod/capstan/main/src/plugins/windows/examples/backup/create/post/
Mirror_the_backup_cDOT.cmd#1 $
REM
REM  Copyright \(c\) 2011 NetApp, Inc.
REM All rights reserved.
REM
REM
REM This is a sample post-task script to mirror the volumes to the
secondary storage after successful backup operation.
REM|--------------------------------------------------------------------
----------------------------------------|
REM| Pre-requisite/Assumption:
|
REM|   SnapMirror relationship should be set for the primary volumes and
secondary volumes                            |
REM|--------------------------------------------------------------------
----------------------------------------|
REM
REM
REM This script can be used from the SnapManager graphical user interface
(GUI) and command line interface (CLI).
REM
REM To execute the post-task script for the backup operation from
SnapManager GUI, follow these steps:
REM  1. From the Backup wizard > Task Specification page > Post-Tasks tab
> select the post-task scripts from the Available Scripts section.
REM
REM
REM To execute the post-task script for the backup operation from
SnapManager CLI, follow these steps:
```

```
REM  1. create a task specification XML file.
REM  2. Add the post-script name in the <post-tasks> tag of the XML file.
REM
REM Example:
REM                 <preposttask-specification xmlns="http://www.netapp.com">
REM                   <task-specification>
REM                      <post-tasks>
REM                        <task>
REM                           <name>"Mirror the backup for cDOT"</name>
REM                        </task>
REM                      </post-tasks>
REM                   </task-specification>
REM                 </preposttask-specification>
REM
REM
REM
REM IMPORTANT NOTE: This script is provided for reference only. It has
been tested with SnapDrive 7.0 for Windows but may not work in all
environments.  Please review and then customize based on your secondary
protection requirements.
REM
set /a EXIT=0
set name="Mirror the backup cDOT"
set description="Mirror the backup cDOT"
set parameter=()



if /i "%1" == "-check" goto :check
if /i "%1" == "-execute" goto :execute
if /i "%1" == "-describe" goto :describe

:usage
        echo usage: %0 ^{ -check ^| -describe ^| -execute ^}
        set /a EXIT=99
        goto :exit

:check
        set /a EXIT=0
        goto :exit

:describe
        echo SM_PI_NAME:%name%
        echo SM_PI_DESCRIPTION:%description%
        set /a EXIT=0
        goto :exit
```

```
REM - Split the comma-separated SM_PRIMARY_MOUNT_POINTS then Mirror the
PRIMARY_MOUNT_POINTS one-by-one.

:execute
        set /a EXIT=0

        echo "execution started"

        REM FOR %%G IN (%SM_PRIMARY_MOUNT_POINTS%) DO powershell.exe -file
"c:\snapmirror.ps1" %%G < CON

        powershell.exe -file "c:\snapmirror.ps1"
%SM_PRIMARY_FULL_SNAPSHOT_NAME_FOR_TAG% < CON


        if "%ERRORLEVEL%" NEQ "0" (
            set /a EXIT=4
            exit /b %EXIT%
        )

        echo "execution ended"

        goto :exit

:exit
        echo Command complete.
        exit /b %EXIT%
```

Vault_the_backup_cDOT.cmd

If you are using clustered Data ONTAP, you can use this sample script to vault the backup. It includes three operations (check, describe, and execute) and executes them at the end of the script. The script also includes error message handling with codes of 0 to 4 and > 4:

```
@echo off
REM $Id:
//depot/prod/capstan/main/src/plugins/windows/examples/backup/create/post/
Vault_the_backup_cDOT.cmd#1 $
REM
REM  Copyright \(c\) 2011 NetApp, Inc.
REM All rights reserved.
REM
REM
REM This is a sample post-task script to do vault update to the secondary
storage after successful backup operation.
REM|--------------------------------------------------------------------
```

```
                --------------------------------------------|
REM|  Pre-requsite/Assumption:
 |
REM|    Vaulting relationship with policy and rule needs to be established
between primary and secondary storage volumes |
REM|---------------------------------------------------------------
--------------------------------------------|
REM
REM
REM This script can be used from the SnapManager graphical user interface
(GUI) and command line interface (CLI).
REM
REM To execute the post-task script for the backup operation from
SnapManager GUI, follow these steps:
REM
REM  1. From the Backup wizard > Task Specification page > Post-Tasks tab
> select the post-task scripts from the Available Scripts section.
REM  2. You can view the parameters available in the post-task script in
the Parameter section of the Task Specification page.
REM  3. Provide values to the following parameters:
REM     SNAPSHOT_LABEL     - Label Name to be set for snapshots before
doing the vault update
REM
REM FOR WINDOWS ITS ADVISED TO USE THE post-task script FROM THE GUI BY
SAVING THE BELOW SPEC XML AND GIVING THIS IN THE GUI LOAD XML FILE .
REM
REM To execute the post-task script for the backup operation from
SnapManager CLI, follow these steps:
REM  1. create a task specification XML file.
REM  2. Add the post-script name in the <post-tasks> tag of the XML file .
REM Example:
REM            <preposttask-specification xmlns="http://www.netapp.com">
REM              <task-specification>
REM                <post-tasks>
REM                  <task>
REM                    <name>"Vault the backup for cDOT"</name>
REM                    <parameter>
REM                    <name>SNAPSHOT_LABEL</name>
REM                    <value>TST</value>
REM                    </parameter>
REM                  </task>
REM                </post-tasks>
REM              </task-specification>
REM            </preposttask-specification>
REM
REM
```

```
REM IMPORTANT NOTE: This script is provided for reference only. It has
been tested with SnapDrive 7.0.0 for Windows but may not work in all
environments.
Please review and then customize based on your secondary protection
requirements.
REM
REM
REM Need to take care of the parameter variable, its not like shell script
array handling, so declare a new variable
REM for one more argument and set that variable SM_PI_PARAMETER in the
describe method. Then only that variable will be
REM Visible in the GUI task specification wizard else it wont list.

set /a EXIT=0
set name="Vault the backup for cDOT"
set description="Vault the backup For cDOT volumes"
set parameter=SNAPSHOT_LABEL :

if /i "%1" == "-check" goto :check
if /i "%1" == "-execute" goto :execute
if /i "%1" == "-describe" goto :describe

:usage
        echo usage: %0 ^{ -check ^| -describe ^| -execute ^}
        set /a EXIT=99
        goto :exit

:check
        set /a EXIT=0
        goto :exit

:describe
        echo SM_PI_NAME:%name%
        echo SM_PI_DESCRIPTION:%description%
    echo SM_PI_PARAMETER:%parameter%

        set /a EXIT=0
        goto :exit

REM Split the colon-separated SM_PRIMARY_SNAPSHOTS_AND_MOUNT_POINTS And
SnapVault the mountpoints one-by-one

:execute
        set /a EXIT=0

    echo "execution started"
```

```
    powershell.exe -file "c:\snapvault.ps1"
%SM_PRIMARY_FULL_SNAPSHOT_NAME_FOR_TAG%  %SNAPSHOT_LABEL% < CON


        if "%ERRORLEVEL%" NEQ "0" (
            set /a EXIT=4
            exit /b %EXIT%
        )
    echo "execution ended"


        goto :exit


:exit
        echo Command complete.
        exit /b %EXIT%
```

Create or update the post scripts

You can either create new post scripts or use the scripts available at
`default_install_directory\plugins\backup\create\post`.

About this task

You must structure the script in a particular manner so that it can be executed within the context of a
SnapManager operation. Create the script based on the expected operations, available input parameters, and
return code conventions.

Steps

1. Create a new script or use the available sample scripts.

2. Modify or include the functions, variables, and parameters as required.

3. Save the custom script.

4. Add the script name and required inputs in the post-processing task specification XML file or provide the
   inputs from the GUI by selecting the proper script and input parameter.

> (i) You do not have to provide secondary storage details in the task specification XML file.

Create post-process, task-specification files

SnapManager enables you to create post-processing, task-specification XML files for the
backup operation that include the SnapMirror or SnapVault post-scripts. The scripts allow
you to mirror or vault the backup to secondary storage.

Steps

1. Open a new XML file.

   You can see the sample task-specification XML file available at
   `default_install_directory\plugins\examples`.

2. Add the script name as an input parameter.

3. Save the task-specification XML file.

Use post-processing task specification to mirror volumes

SnapManager for SAP enables you to use the script to mirror the volumes after the backup operation occurs in a Windows environment.

Steps

1. Create a task specification XML file.

2. In the XML file, enter the script name as an input parameter.

3. Save the task specification XML file.

4. Create a protected backup of the database to secondary storage using the following command.

   While creating the protected backup, you must provide the complete path of the saved task specification XML file after the `-taskspec` option.

   Example: `smsapbackup create -profile test_profile -full -online -taskspec "C:\\mirror\\snapmirror.xml"`

The following example shows a post-processing task specification structure if you are using Data ONTAP operating in 7-Mode:

```
#      <post-tasks>
#              <task>
#                   <name>Mirror the backup</name>
#                    <description>Mirror the backup</description>
#              </task>
#           <post-tasks>
```

The following example shows a post-processing task specification structure if you are using clustered Data ONTAP:

```
# <task-specification>
#              <post-tasks>
#                 <task>
#                  <name>"Vault the backup for cDOT"</name>
#                   <parameter>
#                   <name>SNAPSHOT_LABEL</name>
#                 <value>TST</value>
#                   </parameter>
#                 </task>
#              </post-tasks>
#           </task-specification>
#          </preposttask-specification>
```

Use post-processing task specification to vault qtrees

SnapManager for SAP enables you to use the script to vault the qtrees after the backup operation takes place in a Windows environment.

Steps

1. Create a task specification XML file.

2. In the XML file, enter the script name as an input parameter.

3. Save the task specification XML file.

4. Create a protected backup of the database to secondary storage using the following command.

   While creating the protected backup, you must provide the complete path of the saved task specification XML file after the `-taskspec` option.

   Example: `smsapbackup create -profile test_profile -full -online -taskspec "C:\\mirror\\snapvault.xml"`

The following example shows the post-processing task specification structure if you are using Data ONTAP operating in 7-Mode:

```
# <post-tasks>
#                  <task>
#                     <name>Vault the backup</name>
#                      <description>Vault the backup</description>
#               </task>
#           <post-tasks>
```

The following example shows the post-processing task specification structure if you are using clustered Data ONTAP:

```
# <task-specification>
#                  <post-tasks>
#                    <task>
#                      <name>"Vault the backup for cDOT"</name>
#                   <parameter>
#                        <name>SNAPSHOT_LABEL</name>
#                        <value>TST</value>
#                        </parameter>
#                     </task>
#                  </post-tasks>
#               </task-specification>
#             </preposttask-specification>
```

Performing management operations

You can perform management tasks after you have set up and configured SnapManager. These tasks enable you to manage normal operations beyond backing up, restoring, and cloning.

Administrators can perform operations either by using the graphical user interface or command-line interface.

View a list of operations

You can view a summary listing of all the operations performed against a profile.

About this task

You can view the following information when you list operations associated with a particular profile:

- Start and end date when the operation ran
- Operation status
- Operation ID
- Type of operation
- Host that it ran upon

Step

1. To list the summary information of all the operations, use the following command:

   smsap operation list profile -profile *profile_name* -delimiter *character* [-quiet | -verbose]

   When the `-delimiter` option is specified, the command lists each row on a separate line and the attributes in that row are separated by the character specified.

View operation details

You can view detailed information about a particular profile to verify the success or failure of an operation. It can also help you determine the storage resources in use for a particular operation.

About this task

You can view the following details about a particular operation:

- Operation ID
- Type of operation
- Whether the operation was forced
- Runtime information, including status, start and end date of the operation
- The host on which the operation ran, including the Process ID and SnapManager version
- Repository information

- Storage resources in use

Step

1. To view the detailed information for a specific operation ID, enter the following command:

```
smsap operation show -profile profile_name [-label label | -id id] [-quiet |
-verbose]
```

Issue commands from an alternate host

You can issue CLI commands from a host other than the database host and SnapManager will route the commands you enter to the appropriate host.

About this task

For the system to dispatch an operation to the correct host, it must first know where to find the profile for the operation. In this procedure the system keeps the profile to repository mapping information for a file in the user's home directory on the local host.

Step

1. To make the local user's home directory aware of the profile-to-repository mappings so it can route the operation request, enter the following command:

```
smsap profile sync -repository -dbname repo_dbname -host repo_host -port
repo_port -login -username repo_username [-quiet | -verbose]
```

Check the SnapManager software version

You can determine which version of the product you are running on your local host by running the `version` command.

Step

1. To check the SnapManager version, enter this command: `smsap version`

Stop the SnapManager host server

When you have finished using SnapManager, you might want to stop the server.

Step

1. To stop the server, enter the following command, as a root user:

```
smsap_server stop
```

Restart the SnapManager host server

You can restart the SnapManager server through the Services window.

Steps

1. Click **Start** > **Control Panel** > **Administrative Tools** > **Services**.

2. In the Services window, select NetAppSnapManager 3.3 for SAP.

3. You can restart the server in one of the following ways:

   a. In the left panel, click **Restart**.

   b. Right-click NetAppSnapManager 3.3 for SAP and select **Restart** from the drop-down menu.

   c. Double-click NetAppSnapManager 3.3 for SAP and in the properties window that opens, click **Restart**.

Uninstall SnapManager

You can uninstall SnapManager from the host server.

What you'll need

* Stop the host server before uninstalling SnapManager.

Steps

1. Click **Start** > **Control Panel** > **Add or Remove Programs**.

2. Select **SnapManager for SAP**.

3. Click **Uninstall**.

Configuring an email notification

SnapManager enables you to receive an email notification about the completion status of the profile-executed database operations. SnapManager generates the email and helps you to take appropriate action based on the database operation completion status. Configuring email notification is an optional parameter.

You can configure an email notification for an individual profile as a profile notification and for multiple profiles on a repository database as a summary notification.

Profile notification

For an individual profile, you can receive an email for either or both the successful and failed database operations.

> ⓘ By default, email notification is enabled for failed database operations.

Summary notification

Summary notification enables you to receive a summary email about database operations performed using multiple profiles. You can enable hourly, daily, weekly, or monthly notifications.

> ⓘ From SnapManager 3.3, summary notifications are sent only if you specify the host server that has to send the notification. If you upgrade SnapManager from a version earlier than 3.3, the summary notifications might not be sent if you had not specified the host server in the summary notifications configuration.

ⓘ If you create a repository in one node of a database that is on a Real Application Clusters (RAC) environment and enable summary notification, later when you add the same repository to another node of the database, the summary notification email is sent twice.

You can use either profile-level notification or summary notification at a time.

SnapManager enables email notification for the following profile-executed database operations:

- Create backup on primary storage
- Restore backups
- Create clones
- Verify backups

After you create or update profiles with the email notification enabled, you can disable it. If you disable the email notification, you no longer receive email alerts for those profile-executed database operations.

The email that you receive contains the following details:

- Name of the database operation, for example, backup, restore, or clone
- Profile name used for the database operation
- Name of the host server
- System identifier of the database
- Start and end time of the database operation
- Status of the database operation
- Error message, if any
- Warning messages, if any

You can configure the following:

- Mail server for a repository
- Email notification for a new profile
- Email notification for an existing profile
- Summary email notification for multiple profiles under a repository

ⓘ You can configure email notification from both the command-line interface (CLI) and the graphical user interface (GUI).

Configure a mail server for a repository

SnapManager enables you to specify the mail server details from which the email alerts are sent.

About this task

SnapManager enables you to specify the sender's email server host name or IP address, and the email server port number for a repository database name that requires email notification. You can configure the mail server port number in a range from 0 through 65535; the default value is 25. If you require authentication for the email

address, you can specify the user name and password.

You must specify name or IP address of the host server that handles the email notification.

Step

1. To configure the mail server to send email alerts, enter the following command:

   smsap notification set -sender-email *email_address* **-mailhost** *mailhost*
   -mailport *mailport* **[-authentication -username** *username* **-password** *password*]
   -repository -port *repo_port* **-dbname** *repo_service_name* **-host** *repo_host* **-login**
   -username *repo_username*

   Other options for this command are as follows:

   [-force]

   [quiet|-verbose]

To do the following…	Then…
To specify the sender's email address.	Specify the -sender-email option. From SnapManager 3.2 for SAP, you can include hyphen (-) while specifying the domain name of the email address. For example, you can specify the sender email address as -sender-emailuser@org-corp.com.
To specify the sender's email server host name or IP address.	Specify the -mailhost option.
To specify the email server port number for a repository database name that requires email notification. You can configure the mail server port number in a range from zero through 65535; the default value is 25.	Specify the -mailport option.
Specify the user name and password if you require authentication for the email address.	Specify -authentication option followed by the user name and password.

The following example configures the mail server.

```
smsap notification set -sender-email admin1@org.com -mailhost
hostname.org.com -mailport 25 authentication -username admin1 -password
admin1 -repository -port 1521 -dbname SMSAPREPO -host hotspur -login
-username grabal21 -verbose
```

Configure email notification for a new profile

When you are creating a new profile, you can configure to receive an email notification on completion of the database operation.

What you'll need

- You must configure the email address from which the alerts are sent.

- You must use a comma-separated list for multiple email addresses.

    You must ensure that there is no space between the comma and the next email address.

    > ℹ️ You must ensure you enter the set of email addresses within double quotation marks.

Step

1. Enter the following command:

smsap profile create -profile *profile* **[-profile-password** *profile_password*] **-repository -dbname** *repo_service_name* **-host** *repo_host* **-port** *repo_port* **-login -username** *repo_username* **-database -dbname** *db_dbname* **-host** *db_host* **[-sid** *db_sid*] **[-login -username** *db_username* **-password** *db_password* **-port** *db_port*] **[-rman {-controlfile | {-login -username** *rman_username* **-password** *rman_password* **-tnsname** *rman_tnsname*} }] **[-retain [-hourly [-count** *n*] **[-duration** *m*]] **[-daily [-count** *n*] **[-duration** *m*]] **[-weekly [-count** *n*] **[-duration** *m*]] **[-monthly [-count** *n*] **[-duration** *m*]]] **[-comment** *comment*][**-snapname-pattern** *pattern*][**-protect [ ]][-notification [-success-email** *email_address1,email_address2* **-subject** *subject_pattern*] **[-failure -email** *email_address1,email_address2* **-subject** *subject_pattern*]]

Other options for this command are as follows:

[-force]

[quiet|-verbose]

> ℹ️ SnapManager supports up to 1000 characters for email addresses.

When you create a backup of data files and archive log files together using the profile (for creating separate archive log backups), and the data file backup creation fails, the email notification is sent with the data backup as the operation name instead of data backup and archive logs backup. When the data file and archive log file backup operation is successful, you see the output as follows:

```
Profile Name     : PROF_31
Operation Name   : Data Backup and Archive Logs Backup
Database SID     : TENDB1
Database Host    : repo1.rtp.org.com
Start Date       : Fri Sep 23 13:37:21 EDT 2011
End Date         : Fri Sep 23 13:45:24 EDT 2011
Status           : SUCCESS
Error messages   :
```

Customize the email subject for a new profile

You can customize the email subject for the new profile when you create it.

About this task

You can customize the email subject by using the {profile}_{operation-name}_{db-sid}_{db-host}_{start-date}_{end-date}_{status} pattern or enter your own text.

Variable name	Description	Example value
`profile`	Profile name used for the database operation	PROF1
`operation-name`	Database operation name	Backup, Data Backup, Data and Archive Logs Backup
`db-sid`	SID of the database	DB1
`db-host`	Name of the host server	hostA
`start-date`	Start time of the database operation in the mmdd:hh:ss yyyy format	April 27 21:00:45 PST 2012
`end-date`	End time of the database operation in the mmdd:hh:ss yyyy format	April 27 21:10:45 PST 2012
`status`	Database operation status	Success

If you do not provide any value for the variables, then SnapManager displays the following error message: `Missing value(s) -subject.`

Step

1. Enter the following command:

   smsap profile create -profile *profile* **[-profile-password** *profile_password*]
   -repository -dbname *repo_service_name* **-host** *repo_host* **-port** *repo_port* **-login**

```
-username repo_username -database -dbname db_dbname -host db_host [-sid
db_sid] [-login-username db_username -password db_password -port db_port] [-
rman {-controlfile | {-login -username rman_username -password rman_password
-tnsname rman_tnsname} } ] [-retain [-hourly [-count n] [-duration m]] [-daily
[-count n] [-duration m]] [-weekly [-count n] [-duration m]] [-monthly [-count
n] [-duration m]]] [-comment comment][-snapname -pattern pattern][ [ ]] [-
notification [-success -email email_address1, email_address2 -subject
subject_pattern] [-failure -email email_address1, email_address2 -subject
subject_pattern]]
```

Configure email notification for an existing profile

When you are updating an profile, you can configure to receive an email notification on
completion of the database operation.

What you'll need

- You must configure the email address from which the alerts are sent.
- You must enter a single email address or multiple email addresses to which alerts will be sent.

  You can use a comma-separated list for multiple addresses. You must ensure that there is no space
  between the comma and the next email address. Optionally, you can add a subject to the email as well.

  > (i) You must ensure that you enter the set of multiple email addresses within double quotation
  > marks.

Step

1. Enter the following command:

```
smsap profile update -profile profile [-profile-password profile_password][-
database -dbname db_dbname -host db_host [-sid db_sid] [-login -username
db_username -password db_password -port db_port]] [{-rman{-controlfile | {-
login -username rman_username -password rman_password -tnsname rman_tnsname}}}
| -remove-rman] [-retain [-hourly [-count n] [-duration m]] [-daily [-count n]
[-duration m]] [-weekly [-count n] [-duration m]] [-monthly [-count n] [-
duration m]]] [-comment comment][-snapname-pattern pattern][[ [ ]]| [[]] [-
notification [-success -email email_address1,email_address2 -subject
subject_pattern] [-failure -email email_address1,email_address2 -subject
subject_pattern]]
```

You can use the `success` option to receive a notification only for successful database operations and the
`failure` option to receive a notification only for failed database operations.

Customize the email subject for an existing profile

SnapManager enables you to customize the email subject pattern for an existing profile
by updating that profile. This customized subject pattern is applicable only for the updated
profile.

Step

1. Enter the following command:

```
smsap profile update -profile profile [-profile -password profile_password][-
database -dbname db_dbname -host db_host [-sid db_sid] [-login -username
db_username -password db_password -port db_port]] [{-rman{-controlfile | {-
login -username rman_username -password rman_password -tnsname rman_tnsname}}}
| -remove-rman] [-retain [-hourly [-count n] [-duration m]] [-daily [-count n]
[-duration m]] [-weekly [-count n] [-duration m]] [-monthly [-count n] [-
duration m]]] [-comment comment][-snapname -pattern pattern][[-protect [-
protection-policy policy_name]]| [[-noprotect]] [-notification [-success
-email email_address1, email_address2 -subject subject_pattern] [-failure
-email email_address1, email_address2 -subject subject_pattern]]
```

Configure summary email notification for multiple profiles

SnapManager enables you to configure a summary email notification for multiple profiles under a repository database.

About this task

You can set the SnapManager server host as a notification host from which the summary notification email is sent to the recipients. If the SnapManager server host name or IP address is changed, then the notification host can also be updated.

You can select any one of the schedule times at which you require an email notification:

- Hourly: To receive an email notification every hour
- Daily: To receive an email notification daily
- Weekly: To receive an email notification weekly
- Monthly: To receive an email notification monthly

You need to enter a single email address or a comma-separated list of email addresses to receive notifications for operations performed using those profiles. You must ensure that there is no space between the comma and the next email address when you enter multiple email addresses.

> (i) You must ensure that you enter the set of multiple email addresses within double quotation marks.

SnapManager allows you to add a customized email subject using the following variables:

- Profile name used for the database operation.
- Database name
- SID of the database
- Name of the host server
- Start time of the database operation in the yyyymmdd:hh:ss format
- End time of the database operation in the yyyymmdd:hh:ss format
- Database operation status

If you select not to add a customized subject, SnapManager displays an error message: `Missing value -subject.`

Step

1. Enter the following command:

   smsap notification update-summary-notification -repository -port *repo_port*
   -dbname *repo_service_name* **-host** *repo_host* **-login -username** *repo_username*
   -email *email_address1,email_address2* **-subject** *subject-pattern* **-frequency {-**
   daily -time *daily_time* **| -hourly -time** *hourly_time* **| -monthly -time**
   monthly_time **-date {1|2…|31} | -weekly -time** *weekly_time* **-day {1|2|3|4|5|6|7}}**
   -profiles *profile1 profile2* **-notification-host** *notification-host*

   Other options for this command are as follows:

   [-force][-noprompt]

   [quiet|-verbose]

   ```
   smsap notification update-summary-notification –repository –port 1521
   –dbname repo2 –host 10.72.197.133 –login –username oba5 –email-address
   admin@org.com –subject success –frequency -daily -time 19:30:45 –profiles
   sales1 -notification-host wales
   ```

Add a new profile to summary email notifications

After you configure a summary email notification for the repository database, you can add a new profile to summary notification by using the `summary notification` command.

Step

1. Enter the following command:

   smsap profile create -profile *profile_name* **[-profile-password**
   profile_password] **-repository -dbname** *repo_service_name* **-host** *repo_host* **-port**
   repo_port **-login- username** *repo_username* **-database -dbname** *db_dbname* **-host**
   db_host **[-sid** *db_sid*] **[-login -username** *db_username* **-password** *db_password*
   -port *db_port*] **[-rman {-controlfile | {-login -username** *rman_username*
   -password *rman_password* **-tnsname** *rman_tnsname*} **}] [-retain [-hourly -count** *n*]
   [-duration *m*]] **[-daily -count** *n*] **[-duration** *m*]] **[-weekly -count** *n*] **[-duration**
   m]] **[-monthly -count** *n*] **[-duration** *m*]]] **[-comment** *comment*][**-snapname-pattern**
   pattern]**[[ ]] [-summary-notification]**

   Other options for this command are as follows:

   [-force]

   [quiet|-verbose]

Add an existing profile to summary email notifications

SnapManager enables you to add an existing profile to a summary email notification while updating that profile.

Step

1. Enter the following command:

   smsap profile update -profile *profile_name* **[-profile-password**
   profile_password] **-repository -dbname** *repo_service_name* **-host** *repo_host* **-port**
   repo_port **-login -username** *repo_username* **-database -dbname** *db_dbname* **-host**
   db_host **[-sid** *db_sid*] **[-login -username** *db_username* **-password** *db_password*
   -port *db_port*] **[-rman {-controlfile |** {**-login -username** *rman_username*
   -password *rman_password* **-tnsname** *rman_tnsname*} } **] [-retain [-hourly -count** *n*]
   [-duration *m*]] **[-daily -count** *n*] **[-duration** *m*]] **[-weekly -count** *n*] **[-duration**
   m]] **[-monthly -count** *n*] **[-duration** *m*]]] **[-comment** *comment*][**-snapname-pattern**
   pattern]**[[ ]] [-summary-notification]**

Disable email notification for multiple profiles

After you enable the summary email notification for multiple profiles, you can disable them to no longer receive email alerts.

About this task

SnapManager enables you to disable the summary email notification for those profile-executed database operations. From the SnapManager CLI, enter the `notification remove-summary-notification` command to disable the summary email notification for multiple profiles and the name of the repository database for which you do not require email notification.

Step

1. To disable summary notification for multiple profiles on a repository database, enter the following command:

   smsap notification remove-summary-notification -repository -port *repo_port*
   -dbname *repo_service_name* **-host** *repo_host* **-login -username** *repo_username*

The following example shows summary notification being disabled for multiple profiles on a repository database:

```
smsap notification remove-summary-notification –repository –port 1521
-dbname repo2 –host 10.72.197.133 –login –username oba5
```

Creating task specification file and scripts for SnapManager operations

SnapManager for SAP uses a task specification Extensible Markup Language (XML) file that indicates the pretasks and post-tasks for the backup, restore, and clone operations.

You can add the pretask and post-task script names in the XML file for the tasks to be performed before or after the backup, restore, and clone operations.

In SnapManager (3.1 or earlier), you can run the pretask and post-task scripts only for the clone operation. In SnapManager (3.2 or later) for SAP, you can run the pretask and post-task scripts for the backup, restore, and clone operations.

In SnapManager (3.1 or earlier), the task specification section is part of the clone specification XML file. From SnapManager 3.2 for SAP, the task specification section is a separate XML file.

> (i) SnapManager 3.3 or later does not support the use of the clone specification XML file created in the releases before SnapManager 3.2.

In SnapManager (3.2 or later) for SAP, you must ensure that the following conditions are met for successful SnapManager operations:

- For backup and restore operations, use the task specification XML file.
- For the clone operation, provide two specification files: a clone specification XML file and a task specification XML file.

  If you want to enable pretask or post-task activity, you can optionally add the task specification XML file.

You can create the task specification file by using the SnapManager graphical user interface (GUI), command-line interface (CLI), or a text editor. You must use a .xml extension for the file to enable appropriate editing features. You might want to save this file so that you can use it for future backup, restore, and clone operations.

The task specification XML file includes two sections:

- The pretasks section includes scripts that could be run before the backup, restore, and clone operations.
- The post-tasks section includes scripts that could be run after the backup, restore, and clone operations.

The values included in the pretasks and post-tasks sections must adhere to the following guidelines:

- Task name: The name of the task must match the name of the script, which is displayed when you run the `plugin.sh -describe` command.

  > (i) If there is a mismatch, then you might receive the following error message: `the file not found.`

- Parameter name: The name of the parameter must be a string that can be used as an environment variable setting.

  The string must match the parameter name in the custom script, which is displayed when you run the `plugin.sh -describe` command.

You can create the specification file based on the structure of the following sample task specification file:

```
<task-specification>
  <pre-tasks>
<task>
      <name>name</name>
      <parameter>
           <name>name</name>
           <value>value</value>
       </parameter>
    </task>
  </pre-tasks>
  <post-tasks>
    <task>
      <name>name</name>
      <parameter>
           <name>name</name>
           <value>value</value>
       </parameter>
    </task>
  </post-tasks>
</task-specification>
```

ⓘ   The task specification XML file should not contain any policy.

From the SnapManager GUI, you can set the parameter value and save the XML file. You can use the Task Enabling page of the Backup Create wizard, Restore or Recovery wizard, and Clone Create wizard, to load the existing task specification XML file, and use the selected file for the pretask or post-task activity.

A task can be executed multiple times, either with the same or different parameter and value combinations. For example, you could use a Save task to save multiple files.

ⓘ   SnapManager uses the XML tags provided in the task specification file for the preprocessing or post-processing activity for the backup, restore, and clone operations irrespective of the file extension of the task specification file.

Create pretask, post-task, and policy scripts

SnapManager enables you to create the scripts for the preprocessing activity, the post-processing activity, and policy tasks of the backup, restore, and clone operations. You must place the scripts in the correct installation directory to execute the preprocessing activity, post-processing activity, and policy tasks of the SnapManager operation.

About this task

Pretask and post-task script content

All scripts must include the following:

- Specific operations (check, describe, and execute)
- (Optional) Predefined environment variables
- Specific error handling code (return code (rc))

> ℹ️ You must include correct error handling code to validate the script.

You can use the pretask scripts for many purposes, for example, cleaning up a disk space before the SnapManager operation starts. You can also use the post-task scripts, for instance, to estimate whether SnapManager has enough disk space to complete the operation.

Policy task script content

You can execute the policy script without using specific operations such as check, describe, and execute. The script includes the predefined environmental variables (optional) and specific error handling code.

The policy script is executed before the backup, restore, and clone operations.

Supported format

A command file with a .cmd extension can be used as the prescript and post-script.

> ℹ️ If you select the shell script file, the SnapManager operation does not respond. To resolve this, you must provide the command file in the plug-in directory, and perform the SnapManager operation again.

Script installation directory

The directory in which you install the script affects how it is used. You can place the scripts in the directory and execute the script before or after the backup, restore, or clone operation takes place. You must place the script in the directory specified in the table and use it on an optional basis when you specify the backup, restore, or clone operation.

> ℹ️ You must ensure that the plugins directory has the executable permission before using the scripts for the SnapManager operation.

Activity	Backup	Restore	Clone
Preprocessing	<default_installation_directory>\plugins\backup\create\pre	<default_installation_directory>\plugins\restore\create\pre	<default_installation_directory>\plugins\clone\create\pre
Post-processing	<default_installation_directory>\plugins\backup\create\post	<default_installation_directory>\plugins\restore\create\post	<default_installation_directory>\plugins\clone\create\post
Policy-based	<default_installation_directory>\plugins\backup\create\policy	<default_installation_directory>\plugins\restore\create\policy	<default_installation_directory>\plugins\clone\create\policy

Sample scripts locations

The following are some samples of the pretask and post-task scripts for the backup and clone operations available in the installation directory path:

- <default_installation_directory>\plugins\examples\backup\create\pre
- <default_installation_directory>\plugins\examples\backup\create\post
- <default_installation_directory>\plugins\examples\clone\create\pre
- <default_installation_directory>\plugins\examples\clone\create\post

What you can change in the script

If you are creating a new script, you can change only the describe and execute operations. Each script must contain the following variables: `context`, `timeout`, and `parameter`.

The variables you have described in the describe function of the script must be declared at the start of the script. You can add new parameter values in `parameter=()` and then use the parameters in the execute function.

Sample script

The following is a sample script with a user-specified return code for estimating the space in the SnapManager host:

```
@echo off
REM $Id:
//depot/prod/capstan/Rcapstan_ganges/src/plugins/windows/examples/clone/cr
eate/policy/validate_sid.cmd#1 $
REM $Revision: #1 $ $Date: 2011/12/06 $
REM
REM

set /a EXIT=0

set name="Validate SID"
set description="Validate SID used on the target system"
set parameter=()

rem reserved system IDs
set INVALID_SIDS=("ADD" "ALL" "AND" "ANY" "ASC" "COM" "DBA" "END" "EPS"
"FOR" "GID" "IBM" "INT" "KEY" "LOG" "MON" "NIX" "NOT" "OFF" "OMS" "RAW"
"ROW" "SAP" "SET" "SGA" "SHG" "SID" "SQL" "SYS" "TMP" "UID" "USR" "VAR")

if /i "%1" == "-check" goto :check
if /i "%1" == "-execute" goto :execute
if /i "%1" == "-describe" goto :describe

:usage:
    echo usage: %0 "{ -check | -describe | -execute }"
    set /a EXIT=99
```

```
    goto :exit

:check
    set /a EXIT=0
    goto :exit

:describe
    echo SM_PI_NAME:%name%
    echo SM_PI_DESCRIPTION:%description%
    set /a EXIT=0
    goto :exit

:execute
    set /a EXIT=0

    rem SM_TARGET_SID must be set
    if "%SM_TARGET_SID%" == "" (
        set /a EXIT=4
        echo SM_TARGET_SID not set
        goto :exit
    )

    rem exactly three alphanumeric characters, with starting with a letter
    echo %SM_TARGET_SID% | findstr "\<[a-zA-Z][a-zA-Z0-9][a-zA-Z0-9]\>"
>nul
    if %ERRORLEVEL% == 1 (
        set /a EXIT=4
        echo SID is defined as a 3 digit value starting with a letter.
[%SM_TARGET_SID%] is not valid.
        goto :exit
    )

    rem not a SAP reserved SID
    echo %INVALID_SIDS% | findstr /i \"%SM_TARGET_SID%\" >nul
    if %ERRORLEVEL% == 0 (
        set /a EXIT=4
        echo SID [%SM_TARGET_SID%] is reserved by SAP
        goto :exit
    )

    goto :exit



:exit
    echo Command complete.
    exit /b %EXIT%
```

Operations in task scripts

The pretask or post-task scripts that you create must follow a standard SnapManager for SAP plug-in structure.

The pretask and post-task scripts must include the following operations:

- check
- describe
- execute

If any one of these operations is not specified in the pretask or post-task script, then the script becomes invalid.

When you run the `smsap plugin check` command for the pretask or post-task scripts, the returned status of the scripts display error (because the returned status value is not zero).

Operation	Description
check	The SnapManager server runs the `plugin.sh -check` command to ensure that the system has execution permission on the plug-in scripts. You might also include file permission checking on the remote system.

Operation	Description

describe	The SnapManager server runs the `plugin.sh -describe` command to obtain information about your script and match the elements provided by the specification file. Your plug-in script must contain the following description information:
	
	 • `SM_PI_NAME`: Script name. You must provide a value for this parameter.
	
	 • `SM_PI_DESCRIPTION`: Description of the script's purpose. You must provide a value for this parameter.
	
	 • `SM_PI_CONTEXT`: Context in which the script should run-for example, root or orasid. You must provide a value for this parameter.
	
	 • `SM_PI_TIMEOUT`: The maximum time (in milliseconds) that SnapManager should wait for the script to complete processing and terminate execution. You must provide a value for this parameter.
	
	 • `SM_PI_PARAMETER`: One or more custom parameters necessary for your plug-in script to perform processing. Each parameter should be listed in a new output line and include the name of the parameter and a description. When the script completes processing, the parameter value will be provided to your script by an environment variable.
	
	 The following is the sample output of the Followup_activities script.
	
	 <pre>plugin.sh - describe
	
	SM_PI_NAME:Followup_activities
	SM_PI_DESCRIPTION:this script
	contains follow-up activities to
	be executed after the clone create
	operation.
	SM_PI_CONTEXT:root
	SM_PI_TIMEOUT:60000
	SM_PI_PARAMETER:SCHEMAOWNER:Name
	of the database schema owner.
	Command complete.</pre>

Operation	Description
execute	The SnapManager server runs the `plugin.sh -execute` command to start your script to execute the script.

Variables available in the task scripts for the backup operation

SnapManager provides context information in the form of environment variables related to the backup operation that is being performed. For example, your script can retrieve the name of the original host, the name of the retention policy, and the label of the backup.

The following table lists the environment variables that you can use in your scripts:

Variables	Description	Format
`SM_OPERATION_ID`	Specifies the ID of the current operation	string
`SM_PROFILE_NAME`	Specifies the name of the profile used	string
`SM_SID`	Specifies the system identifier of the database	string
`SM_HOST`	Specifies the host name of the database	string
`SM_OS_USER`	Specifies the operating system (OS) owner of the database	string
`SM_OS_GROUP`	Specifies the OS group of the database	string
`SM_BACKUP_TYPE`	Specifies the type of the backup (online, offline, or auto)	string
`SM_BACKUP_LABEL`	Specifies the label of the backup	string
`SM_BACKUP_ID`	Specifies the ID of the backup	string
`SM_BACKUP_RETENTION`	Specifies the retention period	string
`SM_BACKUP_PROFILE`	Specifies the profile used for this backup	string

Variables	Description	Format
SM_ALLOW_DATABASE_SHUTDOWN	Specifies if you want to start up or shut down the databaseIf required you can use the -force option from the command-line interface.	boolean
SM_BACKUP_SCOPE	Specifies the scope of the backup (full or partial)	string
SM_TARGET_FILER_NAME	Specifies the target storage system name ⓘ If more than one storage system is used, then the storage system names must be separated by commas.	string
SM_TARGET_VOLUME_NAME	Specifies the target volume name ⓘ The target volume name must be prefixed with storage device name, for example, SM_TARGET_FILER_NAME/SM_TARGET_VOLUME_NAME.	string
SM_HOST_FILE_SYSTEM	Specifies the host file system	string
SM_SNAPSHOT_NAMES	Specifies the Snapshot list ⓘ Name of the Snapshot copies must be prefixed with the storage system name and volume name. Names of the Snapshot copies are separated by commas.	string array

Variables	Description	Format
SM_ARCHIVE_LOGS_DIRECTORY	Specifies the archive logs directory ⓘ If the archive logs are located in more than one directory, then the names of those directories are separated by commas.	string array
SM_REDO_LOGS_DIRECTORY	Specifies the redo logs directory ⓘ If the redo logs are located in more than one directory, then the names of those directories are separated by commas.	string array
SM_CONTROL_FILES_DIRECTORY	Specifies the control files directory ⓘ If the control files are located in more than one directory, then the names of those directories are separated by commas.	string array
SM_DATA_FILES_DIRECTORY	Specifies the data files directory ⓘ If the data files are located in more than one directory, then the names of those directories are separated by commas.	string array
user_defined	Specifies additional parameters defined by the user. User-defined parameters are not available for plug-ins that are used as policies.	user-defined

Variables available in the task scripts for the restore operation

SnapManager provides context information in the form of environment variables related to the restore operation that is being performed. For example, your script can retrieve the name of the original host and the label of the backup that is restored.

The following table lists the environment variables that you can use in your scripts:

Variables	Description	Format
SM_OPERATION_ID	Specifies the ID of the current operation	string
SM_PROFILE_NAME	Specifies the name of the profile used	string
SM_HOST	Specifies the host name of the database	string
SM_OS_USER	Specifies the operating system (OS) owner of the database	string
SM_OS_GROUP	Specifies the OS group of the database	string
SM_BACKUP_TYPE	Specifies the type of the backup (online, offline, or auto)	string
SM_BACKUP_LABEL	Specifies the backup label	string
SM_BACKUP_ID	Specifies the backup ID	string
SM_BACKUP_PROFILE	Specifies the profile used for the backup	string
SM_RECOVERY_TYPE	Specifies the recovery configuration information	string
SM_VOLUME_RESTORE_MODE	Specifies the volume restore configuration	string

Variables	Description	Format
SM_TARGET_FILER_NAME	Specifies the target storage system name ⓘ If more than one storage system is used, then the storage system names must be separated by commas.	string
SM_TARGET_VOLUME_NAME	Specifies the target volume name ⓘ The target volume name must be prefixed with storage device name, for example, SM_TARGET_FILER_NAME/SM_TARGET_VOLUME_NAME.	string
SM_HOST_FILE_SYSTEM	Specifies the host file system	string
SM_SNAPSHOT_NAMES	Specifies the Snapshot list ⓘ Name of the Snapshot copies must be prefixed with the storage system name and volume name. Names of the Snapshot copies are separated by commas.	string array
SM_ARCHIVE_LOGS_DIRECTORY	Specifies the archive logs directory ⓘ If the archive logs are located in more than one directory, then the names of those directories are separated by commas.	string array

Variables	Description	Format
SM_REDO_LOGS_DIRECTORY	Specifies the redo logs directory ⓘ If the redo logs are located in more than one directory, then the names of those directories are separated by commas.	string array
SM_CONTROL_FILES_DIRECTORY	Specifies the control files directory ⓘ If the control files are located in more than one directory, then the names of those directories are separated by commas.	string array
SM_DATA_FILES_DIRECTORY	Specifies the data files directory ⓘ If the data files are located in more than one directory, then the names of those directories are separated by commas.	string array

Variables available in the task scripts for clone operation

SnapManager provides context information in the form of environment variables related to the clone operation being performed. For example, your script can retrieve the name of the original host, the name of the clone database, and the label of the backup.

The following table lists the environment variables that you can use in your scripts:

Variables	Description	Format
SM_ORIGINAL_SID	SID of the original database	string
SM_ORIGINAL_HOST	Host name associated with the original database	string
SM_ORIGINAL_OS_USER	OS owner of the original database	string

Variables	Description	Format
SM_ORIGINAL_OS_GROUP	OS group of the original database	string
SM_TARGET_SID	SID of the clone database	string
SM_TARGET_HOST	Host name associated with the clone database	string
SM_TARGET_OS_USER	OS owner of the clone database	string
SM_TARGET_OS_GROUP	OS group of the clone database	string
SM_TARGET_DB_PORT	Port of the target database	integer
SM_TARGET_GLOBAL_DB_NAME	Global database name of the target database	string
SM_BACKUP_LABEL	Label of the backup used for the clone	string

Error handling in custom scripts

SnapManager processes the custom script based on the specific return codes. For example, if your custom script returns a value of 0, 1, 2, or 3, SnapManager continues with the clone process. The return code also influences how SnapManager processes and returns the standard output of your script execution.

Return code	Description	Continue processing the operation
0	The script completed successfully.	Yes
1	The script completed successfully, with informational messages.	Yes
2	The script completed, but includes warnings	Yes
3	The script fails, but the operation continues.	Yes
4 or >4	The script fails and the operation stops.	No

View sample plug-in scripts

SnapManager includes scripts that you can use as examples for how to make your own scripts or as a basis for your custom scripts.

About this task

You can find the sample plug-in scripts at the following location:

- `<default_install_directory>\plugins\examples\backup\create`
- `<default_install_directory>\plugins\examples\clone\create`
- `<default_install_directory>\plugins\windows\examples\backup\create\post`

The directory that contains the sample plug-in scripts includes the following subdirectories:

- `policy`: Contains scripts that, when configured, always run on the clone operation.
- `pre`: Contains scripts that, when configured, run before the clone operation.
- `post`: Contains scripts that, when configured, run after the clone operation.

The following table describes the sample scripts:

Script name	Description	Type of script
`validate_sid.sh`	Contains additional checks to the SID used on the target system. The script checks that the SID has the following characteristics: • Contains three alphanumeric characters • Begins with a letter • Does not include reserved SAP SIDs	Policy
`cleanup.sh`	Cleans the target system so that it is ready to store the newly created clone. Preserves or deletes files and directories depending on the need.	Pretask
`sap_follow_up_activities.sh`	Performs follow-up activity tasks as described in *SAP System Copy Guide* and TR-3442, SAP with Oracle on UNIX and NFS and NetApp Storage. For example, this script deletes or modifies table entries in the SAP schema.	Post-task

Script name	Description	Type of script
`os_db_authentication.sh`	Adapts the operating system authentication for the OPS$ user, as recommended in SAP Note 316641. This is a sample of how to process external SQL files.	Post-task
`Mirror_the_backup.cmd`	Mirrors the volumes after the backup operation occurs in a Windows environment when you are using either Data ONTAP operating in 7-Mode.	Post-task
`Vault_the_backup.cmd`	Vaults the qtrees after the backup operation occurs in a Windows environment when you are using either Data ONTAP operating in 7-Mode.	Post-task
`Mirror_the_backup_cDOT.cmd`	Mirrors the volumes after the backup operation occurs in a Windows environment when you are using clustered Data ONTAP.	Post-task
`Vault_the_backup_cDOT.cmd`	Vaults the qtrees after the backup operation occurs in a Windows environment when you are using clustered Data ONTAP.	Post-task

Scripts delivered with SnapManager use the BASH shell by default. You must ensure that support for the BASH shell is installed on your operating system before attempting to run any of the sample scripts.

Steps

1. To verify that you are using the BASH shell, enter the following command at the command prompt:

   bash

   If you do not see an error, the BASH shell is operating properly.

   Alternately, you can enter the `which-bash` command at the command prompt.

2. Locate the script in the following directory:

   `<installdir>\plugins\examples\clone\create`

3. Open the script in a script editor such as vi.

Sample script

The following sample custom script validates database SID names and prevents invalid names from being used in the cloned database. It includes three operations (check, describe, and execute), which are called after

you run the script. The script also includes error message handling with codes 0, 4 and >4.

```
@echo off
REM $Id:
//depot/prod/capstan/Rcapstan_ganges/src/plugins/windows/examples/clone/cr
eate/policy/validate_sid.cmd#1 $
REM $Revision: #1 $ $Date: 2011/12/06 $
REM
REM

set /a EXIT=0

set name="Validate SID"
set description="Validate SID used on the target system"
set parameter=()

rem reserved system IDs
set INVALID_SIDS=("ADD" "ALL" "AND" "ANY" "ASC" "COM" "DBA" "END" "EPS"
"FOR" "GID" "IBM" "INT" "KEY" "LOG" "MON" "NIX" "NOT" "OFF" "OMS" "RAW"
"ROW" "SAP" "SET" "SGA" "SHG" "SID" "SQL" "SYS" "TMP" "UID" "USR" "VAR")

if /i "%1" == "-check" goto :check
if /i "%1" == "-execute" goto :execute
if /i "%1" == "-describe" goto :describe

:usage:
    echo usage: %0 "{ -check | -describe | -execute }"
    set /a EXIT=99
    goto :exit

:check
    set /a EXIT=0
    goto :exit

:describe
    echo SM_PI_NAME:%name%
    echo SM_PI_DESCRIPTION:%description%
    set /a EXIT=0
    goto :exit

:execute
    set /a EXIT=0

    rem SM_TARGET_SID must be set
    if "%SM_TARGET_SID%" == "" (
        set /a EXIT=4
```

```
        echo SM_TARGET_SID not set
        goto :exit
    )


    rem exactly three alphanumeric characters, with starting with a letter
    echo %SM_TARGET_SID% | findstr "\<[a-zA-Z][a-zA-Z0-9][a-zA-Z0-9]\>"
>nul
    if %ERRORLEVEL% == 1 (
        set /a EXIT=4
        echo SID is defined as a 3 digit value starting with a letter.
[%SM_TARGET_SID%] is not valid.
        goto :exit
    )


    rem not a SAP reserved SID
    echo %INVALID_SIDS% | findstr /i \"%SM_TARGET_SID%\" >nul
    if %ERRORLEVEL% == 0 (
        set /a EXIT=4
        echo SID [%SM_TARGET_SID%] is reserved by SAP
        goto :exit
    )


    goto :exit



:exit
    echo Command complete.
    exit /b %EXIT%
```

SAP with Oracle on UNIX and NFS and NetApp Storage: TR-3442

Create task scripts

You can create the pretask, post-task, and policy task scripts for backup, restore, and clone operations, write your script, and include the predefined environment variables in your parameters. You can either create a new script or modify one of the SnapManager sample scripts.

What you'll need

Before you start creating the script, ensure that:

- You must structure the script in a particular manner for it to be run in the context of a SnapManager operation.
- You must create the script based on the expected operations, available input parameters, and return code conventions.

- You must include log messages and redirect the messages to user-defined log files.

Steps

1. Create the task script by customizing the sample script.

   Perform the following:

   a. Locate a sample script in the following installation directory:

      `<default_install_directory>\plugins\examples\backup\create`

      `<default_install_directory>\plugins\examples\clone\create`

   b. Open the script in your script editor.

   c. Save the script with a different name.

2. Modify the functions, variables, and parameters as needed.

3. Save the script in one of the following directories:

   Backup operations scripts

   - `<default_install_directory>\plugins\backup\create\pre`: Executes the script before the backup operation occurs. Use it optionally when you specify the backup creation.
   - `<default_install_directory>\plugins\backup\create\post`: Executes the script after the backup operation occurs. Use it optionally when you specify the backup creation.
   - `<default_install_directory>\plugins\backup\create\policy`: Always executes the script before the backup operation occurs. SnapManager always uses this script for all the backups in the repository.

   Restore operation scripts

   - `<default_install_directory>\plugins\restore\create\pre`: Executes the script before the backup operation occurs. Use it optionally when you specify the backup creation.
   - `<default_install_directory>\plugins\restore\create\post`: Executes the script after the backup operation occurs. Use it optionally when you specify the backup creation.
   - `<default_install_directory>\plugins\restore\create\policy`: Always executes the script before the backup operation occurs. SnapManager always uses this script for all the backups in the repository.

   Clone operation scripts

   - `<default_install_directory>\plugins\clone\create\pre`: Executes the script before the backup operation occurs. Use it optionally when you specify the backup creation.
   - `<default_install_directory>\plugins\clone\create\post`: Executes the script after the backup operation occurs. Use it optionally when you specify the backup creation.
   - `<default_install_directory>\plugins\clone\create\policy`: Always executes the script before the backup operation occurs. SnapManager always uses this script for all the backups in the repository.

Store the task scripts

You must store the pretask, post-task, and policy task scripts in a specified directory on the target server where the backups or clones will be created. For the restore operation, the scripts must be placed in the specified directory on the target server where you want

to restore the backup.

Steps

1. Create your script.
2. Save the script in one of the following locations:

   For the backup operation

Directory	Description
`<default_install_directory >\plugins\backup\create\policy`	The policy scripts run before the backup operations.
`<default_install_directory >\plugins\backup\create\pre`	The preprocessing scripts run the before backup operations.
`<default_install_directory >\plugins\backup\create\pre`	The post-processing scripts run after the backup operations.

   For the restore operation

Directory	Description
`<default_install_directory >\plugins\restore\create\policy`	The policy scripts run before the restore operations.
`<default_install_directory >\plugins\restore\create\pre`	The preprocessing scripts run before the restore operations.
`<default_install_directory >\plugins\restore\create\post`	The post-processing scripts run after the restore operations.

   For the clone operation

Directory	Description
`<default_install_directory >\plugins\clone\create\policy`	The policy scripts run before the clone operations.
`<default_install_directory >\plugins\clone\create\pre`	The preprocessing scripts run before the clone operations.
`<default_install_directory >\plugins\clone\create\post`	The post-processing scripts run after the clone operations.

Verify the installation of plug-in scripts

SnapManager enables you to install and use custom scripts to perform various operations. SnapManager provides plugins for the backup, restore, and clone operations, which you can use to automate your custom scripts before and after the backup, restore, and clone operations.

Step

1. Enter the following command:

   smsap plugin check -osaccount *os db user name*

   If you do not provide the -osaccount option, verification of the plug-in script installation happens for the administrator rather than for a specified user.

   Example

   The following output indicates that the policy1, pre-plugin1, and pre-plugin2 scripts have been installed successfully. However, the post-plugin1 script is not operational.

   ```
           smsap plugin check
    Checking plugin directory structure ...
    <installdir>\plugins\clone\policy
       OK: 'policy1' is executable

    <installdir>\plugins\clone\pre
       OK: 'pre-plugin1' is executable and returned status 0
       OK: 'pre-plugin2' is executable and returned status 0

    <installdir>\plugins\clone\post
       ERROR: 'post-plugin1' is executable and returned status 3
    Command complete.
   ```

Create a task specification file

You can create the task specification files by using graphical user interface (GUI), command-line interface (CLI), or a text editor. These files are used for performing preprocessing or post-processing activity of the backup, restore, or clone operations.

Steps

1. Create a task specification file by using GUI, CLI, or a text editor.

   Example

   You can create the specification file based on the structure of the following sample task specification file:

```
<task-specification>
  <pre-tasks>
                <task>
        <name>name</name>
        <parameter>
            <name>name</name>
            <value>value</value>
         </parameter>
      </task>
  </pre-tasks>
  <post-tasks>
    <task>
        <name>name</name>
        <parameter>
            <name>name</name>
            <value>value</value>
         </parameter>
      </task>
  </post-tasks>
</task-specification>
```

2. Enter the script name.
3. Enter the parameter name and the value assigned to the parameter.
4. Save the XML file in the correct installation directory.

Task specification example

```
   <task-specification>
     <pre-tasks>
       <task>
         <name>clone cleanup</name>
         <description>pre tasks for cleaning up the target
system</description>
         </task>
     </pre-tasks>
     <post-tasks>
       <task>
         <name>SystemCopy follow-up activities</name>
         <description>SystemCopy follow-up activities</description>
         <parameter>
           <name>SCHEMAOWNER</name>
           <value>SAMSR3</value>
         </parameter>
       </task>
       <task>
         <name>Oracle Users for OS based DB authentication</name>
         <description>Oracle Users for OS based DB
authentication</description>
         <parameter>
           <name>SCHEMAOWNER</name>
           <value>SAMSR3</value>
         </parameter>
         <parameter>
           <name>ORADBUSR_FILE</name>


<value\>E:\\mnt\\sam\\oradbusr.sql</value\>
         </parameter>
       </task>
       </post-tasks>
   </task-specification>
```

Perform backup, restore, and clone operations using prescript and post-scripts

You can use your own script while initiating a backup, restore, or clone operation. SnapManager displays a Task-enabling page in the Backup Create wizard, Restore or Recover wizard, or Clone Create wizard, where you can select the script and provide values for any parameters required by the script.

What you'll need

• Install the plug-in scripts in the correct SnapManager installation location.

- Verify that the plug-ins are installed correctly by using the `smsap plugin check` command.
- Ensure that you are using the BASH shell.

About this task

In the command-line interface (CLI), list the script name, select the parameters, and set the values.

Steps

1. To verify that you are using the BASH shell, enter the following command at the command prompt:

   bash

   Alternately, you can enter the `which-bash` command at the prompt, and use the command output as the start parameter of the script.

   The BASH shell is operating properly if you do not see an error.

2. For the backup operation, enter the `-taskspec` option and provide the absolute path of the task specification XML file for performing a preprocessing or a post-processing activity to occur before or after the backup operation:

   ```
   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 | -before {-date yyyy-MM-dd
   HH:mm:ss | -months | -days | -weeks | -hours}} -prune-dest prune_dest1
   ,[prune_dest2]] [-taskspec taskspec] [-include-with-online-backups | -no
   -include-with-online-backups]} -dump [-force] [-quiet | -verbose]
   ```

   If the backup plug-in operation failed, only the plug-in name and return code are displayed. Your plug-in script must include log messages and redirect the messages to the user-defined log files.

3. For the backup restore operation, enter the `-taskspec` option and provide the absolute path of the task specification XML file for performing a preprocessing or a post-processing activity to occur before or after the restore operation:

   ```
   smsap backup restore -profile profile_name {-label <label> | -id <id>} {-files
   <files>|-tablespaces <tablespaces> | -complete | -controlfiles} [-recover {
   -alllogs | -nologs | -until <until>}][-restorespec <restorespec>] [-taskspec
   <taskspec>] [-verify][-force] backup restore -fast [require | override |
   fallback | off] [-preview] -dump [-quiet | -verbose]
   ```

   If the restore plug-in operation failed, only the plug-in name and return code are displayed. Your plug-in script must include log messages and redirect the messages to the user-defined log files.

4. For the clone create operation, enter the `-taskspec` option and provide the absolute path of the task specification XML file for performing a preprocessing or a post-processing activity to occur before or after the clone operation:

   ```
   smsap clone create -profile profile_name {-backup-label backup_name | -backup
   ```

```
-id <backup-id>| -current} -newsid new_sid- clonespec
full_path_to_clonespecfile [-reserve <yes, no, inherit>] [-host <host>] [-
label <label>] [-comment <comment>] {-taskspec <taskspec>} -dump [-quiet |
-verbose]
```

If the clone plug-in operation failed, only the plug-in name and return code are displayed. Your plug-in script must include log messages and redirect the messages to the user-defined log files.

Example of creating a backup using the task specification XML file

```
smsap backup create -profile SALES1 -full -online -taskspec
sales1_taskspec.xml -force -verify
```

Updating storage system name and target database host name associated with a profile

SnapManager 3.3 or later allows you to update the storage system host name or storage system address, and the target database host name associated with a SnapManager profile.

Update the storage system name associated with a profile

SnapManager 3.3 or later provides the ability to update the host name or IP address of a storage system associated with a profile.

What you'll need

You must ensure the following:

- The profile has at least one backup.

  If the profile does not have any backup, then there is no necessity to update the storage system name for that profile.

- No operation is running for the profile.

About this task

You can update the storage system name or IP address by using the SnapManager command-line interface (CLI). While updating the storage system name, the metadata stored in the repository database alone is updated. After renaming the storage system name, you can perform all the SnapManager operations as earlier.

> ⓘ You cannot change the storage system name by using the SnapManager graphical user interface (GUI).

You must ensure that Snapshot copies are available in the new storage system. SnapManager does not verify the existence of the Snapshot copies in the storage system.

However, you must remember the following while performing rolling upgrade and rollback of the host after

renaming the storage system name:

- If you are performing rolling upgrade of the host after renaming the storage system name, you must update the profile with the new storage system name.
- If you are rolling back the host after renaming the storage system, you must ensure that you change the storage system name back to the earlier storage system name so that you can use the profiles, backups, and clones of the earlier storage system for performing SnapManager operations.

Step

1. Enter the following command:

smsap storage rename -profile *profile* -oldname *old_storage_name* -newname *new_storage_name* [quiet | -verbose]

If you want to…	Then…
Update the storage system name associated with a profile	Specify the `-profile` option.
Update the storage system name or IP address associated with a profile	Specify the following options and variables: • `-oldnameold_storage_name` is the host name or IP address of the storage system. • `-newnamenew_storage_name` is the host name or IP address of the storage system.

The following example shows the storage system name being updated:

```
smsap storage rename -profile mjullian -oldname lech -newname hudson
-verbose
```

View a list of storage systems associated with a profile

You can view a list of the storage systems associated with a particular profile.

About this task

The list displays the storage system names associated with the particular profile.

> (i) If there are no backups available for the profile, then you cannot view the storage system name associated with the profile.

Step

1. To display information about storage systems associated with a particular profile, enter this command:

smsap storage list -profile *profile* [-quiet | -verbose]

Example

```
smsap storage list -profile mjubllian
```

```
Sample Output:
Storage Controllers
-------------------
STCO1110-RTP07OLD
```

Update the target database host name associated with a profile

SnapManager (3.2 or later) for SAP provides the ability to update the host name of the target database in the SnapManager profile.

What you'll need

- The local user's home directory must be aware of the profile-to-repository mappings.
- The SnapManager graphical user interface (GUI) sessions must be closed.

About this task

You can update the profile with the new host name by using only the CLI.

Scenarios not supported for changing the target database host name in profile

The following scenarios not supported for changing the target database host name in the profile:

- Changing the target database host name by using the SnapManager GUI
- Rolling back of the repository database after updating the target database host name of the profile
- Updating multiple profiles for a new target database host name by running a single command
- Changing the target database host name when any SnapManager operation is running

> (i) After you update the target database host name in the profile, only the target database host name is changed. All the other configuration parameters set on the profile are retained.

After you update the target database host name, you cannot delete or unmount the backup if the clone or mounted backup is not available in the new host. In such scenarios, running the SnapManager operations from the new host lead to failure as well as stale entries in the earlier host. To perform SnapManager operations, you must revert to the earlier host name by using `profile update`.

Steps

1. Enter the following command:

   smsap profile update -profile *profile* **[-profile-password** *profile_password***][-
   database -dbname** *db_dbname* **-host** *db_host* **[-sid** *db_sid***] [-login-username**
   db_username **-password** *db_password* **-port** *db_port***]] [{-rman{-controlfile | {-
   login-username** *rman_username* **-password** *rman_password* **-tnsname** *rman_tnsname***}}}**

```
| -remove-rman] [-retain [-hourly [-count n] [-duration m]] [-daily [-count n]
[-duration m]] [-weekly [-count n] [-duration m]] [-monthly [-count n] [-
duration m]]] [-comment comment][-snapname-pattern pattern][-summary-
notification] [-notification [-success -email email_address1, email_address2
-subject subject_pattern] [-failure -email email_address1, email_address2
-subject subject_pattern]] [-separate-archivelog-backups-retain-archivelog-
backups -hours hours | -days days | -weeks weeks| -months months [-protect [-
include-with-online-backups | -no-include-with-online-backups]] [-dump]
```

Other options for this command are as follows:

`[-force][-noprompt]`

`[quiet|-verbose]`

If you want to…	Then…
Change the target database host name	Specify `-host new_db_host`

2. To view the target database host name of the profile, enter the following command:

`smsap profile show`

Maintaining history of SnapManager operations

SnapManager for SAP enables you to maintain the history of SnapManager operations associated with a single profile or multiple profiles. You can maintain the history either from the SnapManager command-line interface (CLI) or graphical user interface (GUI). You can view the history of the operations as a report, and use the report for audit compliance purposes.

You can maintain history for the following SnapManager operations:

* Backup create
* Backup verify
* Backup restore
* Clone create

The history information for the SnapManager operations is maintained based on the retention. You can configure different retention classes for each of the supported SnapManager operations.

The following are some retention classes that you can assign:

* Number of days
* Number of weeks
* Number of months
* Number of operations

Based on the retention, SnapManager purges the history automatically. You can also manually purge the history of the SnapManager operations. If you delete or destroy the profile, all the history information associated with the profile is deleted.

> (i) After rollback of the host, you cannot view the history details or perform any history-related operations associated with the profile that has been configured for history maintenance.

Configure history for SnapManager operation

SnapManager for SAP enables you to maintain the history of SnapManager operation from the SnapManager CLI or GUI. You can view the history of the SnapManager operation as a report.

Step

1. To configure the history of SnapManager operation, enter the following command:

   smsap history set -profile {-name, *profile_name* [*profile_name1*, *profile_name2*] | -*all* -repository -login [-password *repo_password*] -username *repo_username* -dbname *repo_dbname* -host *repo_host* -port *repo_port*} -operation {-operations *operation_name* [*operation_name1*, *operation_name2*] | -all} -retain {-count *retain_count* | -daily *retain_daily* | -weekly *retain_weekly* | -monthly *retain_monthly*} [-quiet | -verbose]

   ```
        smsap
   history set -profile -name PROFILE1 -operation -operations backup -retain
   -daily 6 -verbose
   ```

   ```
        smsap
   history set -profile -name PROFILE1 -operation -all -retain -weekly 3
   -verbose
   ```

View a list of SnapManager operation history

You can view the history of a specific or all SnapManager operations as a report based on the retention settings.

Step

1. To view a list of SnapManager history operations, enter the following command:

   smsap history list -profile {-*name, profile_name* [*profile_name1,profile_name2*] | -*all* -repository -login [-password *repo_password*] -username *repo_username* -dbname *repo_dbname* -host *repo_host* -port *repo_port*} -operation {-operations *operation_name* [*operation_name1, operation_name2*] | -all} [-delimiter *delimiter*] [-quiet | -verbose]

View the detailed history of a specific operation associated with a profile

You can view the detailed history of a specific SnapManager operation associated with a profile.

Step

1. To display detailed history information about a specific SnapManager operation associated with a profile, enter the following command:

   smsap history operation-show -profile *profile_name* {-label *label* | -id *id*} [-quiet | -verbose]

Delete history of SnapManager operation

You can delete the history of the SnapManager operation, if you no longer require the history details.

Step

1. To delete the history of the SnapManager operation, enter the following command:

   smsap history purge -profile {-*name*, *profile_name profile_name1*, *profile_name2*] | *all* -repository -login [-password *repo_password*] -username *repo_username* -dbname *repo_dbname* -host *repo_host* -port *repo_port*} -operation {-operations *operation_name* [*operation_name1*, *operation_name2*] | -all} [-quiet | -verbose]

Remove history settings associated with a single profile or multiple profiles

SnapManager enables you to remove the history settings of a SnapManager operation. This operation purges all the history information associated with a single profile or multiple profiles.

Step

1. To remove the history of SnapManager operations associated with a single profile or multiple profiles, enter the following command:

   smsap history remove -profile {-*name*, *profile_name* [*profile_name1*, *profile_name2*] | all -repository -login [-password *repo_password*] -username *repo_username* -dbname *repo_dbname* -host *repo_host* -port *repo_port*} -operation {-operations *operation_name* [*operation_name1*, *operation_name2*] | -all} [-quiet | -verbose]

View SnapManager history configuration details

You can view the history settings for a single profile.

About this task

The SnapManager history operation displays the following information for each profile:

- Operation name

- Retention class

- Retention count

Step

1. To display information about the SnapManager history operation for a specific profile, enter the following command:

   ```
   smsap history show -profile profile_name
   ```

Using BR*Tools with SnapManager for SAP

You can use SAP BR*Tools commands with SnapManager for SAP. BR*Tools is a SAP program package that contains SAP tools for Oracle database administration, for example, BRARCHIVE, BRBACKUP, BRCONNECT, BRRECOVER, BRRESTORE, BRSPACE, and BRTOOLS.

You can complete the following tasks related to using BR*Tools and SnapManager for SAP:

- Disable client access to Snapshot copies

- Specify the profile for BR*Tools backups

- Create database backups using BRBACKUP and BRARCHIVE

- Schedule backups with SAP transaction DB13

- Restore a database using BRRESTORE and BRRECOVER

- Back up and restore files using BR*Tools

- Restore a backup to a different host

What BR*Tools are

You can find information that you would need to use SAP with the storage systems.

For general information about using BR*Tools, including command syntax, see the SAP documentation such as online BR*Tools for Oracle Database Administration.

Profile requirement

The SnapManager for SAP profile must be named correctly to work with BR*Tools. By default, BACKINT uses the profile, which has the same name as the SAP SID, from the repository determined by the user ID that issues the BR*Tools commands.

You need to use a different profile name, if the SAP SID is not unique in your environment. For more information, see *Profile usage for BR* Tools backups*.

You must have a link from the BR*Tools directory to the `C:\Program Files\NetApp\SnapManager for SAP\bin\backint` file installed by SnapManager for SAP. For information about creating the link, see *Integrating with SAP BR* Tools*.

Backups created with versions before BR*Tools 7.00 cannot be verified. The block size of the tablespace or

data file is required to complete the verification; however, versions before BR*Tools 7.00 do not provide this.

About combining BR*Tools and SnapManager for SAP interfaces

You can use a combination of the BR*Tools and the SnapManager for SAP graphical user interface (GUI) or command-line interface (CLI) to perfom following operations:

Operation	Interfaces that you can use
List, restore, recover, and delete database backups (backups that include data files, control files, or online redo log files) created using BRBACKUP	SnapManager for SAP CLI and GUI ⓘ BR*Tools can list and restore only backups created by using BR*Tools.
List and delete backups of other file sets created with BRBACKUP	SnapManager for SAP CLI and GUI
Clone backups of other file sets created with BRBACKUP	A full online or offline backup taken using the BRBACKUP command can be cloned by using the SnapManager for SAP CLI or GUI.

Options available under BR*Tools CLI with SnapManager for SAP

The following options are available under the BR*Tools CLI:

Option	Tasks that can be performed
Instance management	Log all actions in the file system and database tables and also save backup logs and profiles to backup media.
Space management	Allows comprehensive volume management. To use the functions provided in the space management, you must initialize the volumes with BRBACKUP or BRARCHIVE to ensure that they include an SAP-specific label.
Backup and database copy	Create a backup of a database, perform offline redo log (archive log), and verify backups.
Restore and recovery	Restore and recover backups.
Database statistics	Maintain the statistics of tables and indexes.

Similar functions available under BR*Tools GUI

The following operations are available under BR*Tools GUI:

Type of wizard	Tasks that can be performed
Create Repository wizard	Create a repository in a database.
Profile wizard	Create a profile in a repository.
Backup wizard	Create a backup of the profile.
Restore and Recovery wizard	Restore and recover a backup of a profile.
SnapManager Clone Delete wizard	Delete the clone of a backup.
SnapManager Backup Delete wizard	Delete the backup of a profile.

About cloning backups created with BR*Tools

A full online or offline backup taken using the **BRBACKUP** command can be cloned using the SnapManager for SAP CLI or GUI.

To create a clone of the database backup created using the BRBACKUP command, ensure that you set the configuration parameter `brbackup.enable.cloneable.backups` to **true** in the `smsap.config` file and then restart the SnapManager for SAP server. Enter the following command: `brbackup.enable.cloneable.backups=true`

For online backups created using the **BRBACKUP** command, SnapManager for SAP enables you to clone the backup without making any changes from the SnapManager for SAP CLI or GUI.

For offline backups taken created the **BRBACKUP** command, SnapManager for SAP enables you to clone the backup from the SnapManager for SAP CLI or GUI after you perform the following steps:

Steps

1. Add the following configuration variables in the smsap.config file.

   The values for the variables shown below are default values. You can retain these values if you want to set them to default.

   - `brbackup.oracle.maxdatafiles=254`

   - `brbackup.oracle.maxloghistory=1168`

   - `brbackup.oracle.maxinstances=50`

   - `brbackup.oracle.maxlogfiles=255`

   - `brbackup.oracle.maxlogmembers=3`

   - `brbackup.oracle.character_set=UTF8`

2. To change any of these default configuration parameter values in the `smsap.config` file, execute the following command from the SQL command prompt:

   `alter database backup controlfile to trace as file`

   The file can be any user-defined name and is created in the same directory from where SQL*Plus is

invoked.

3. Open the *file* to set these values to the corresponding configuration variables in the `smsap.config` file.

4. Restart the SnapManager server.

About deleting backups created with BR*Tools

BR*Tools does not delete backups. Because SnapManager for SAP backups are based on Snapshot copies, there is a limit to the number of backups you can retain. You need to ensure backups are deleted when they are no longer needed.

On the NetApp storage system, each volume can have a maximum of 255 Snapshot copies. If a volume reaches the limit, backups will fail. A backup made with BRBACKUP typically creates two Snapshot copies of each affected volume.

To prevent reaching the 255 Snapshot copy maximum, you can manage backups in the following ways:

- You can set the retain options in the profile used for BR*Tools operations.

  SnapManager for SAP then automatically deletes older backups as needed.

- You can manually delete backups that are no longer needed by using the SnapManager for SAP CLI or GUI.

Set the path for BR*Tools

You must set the path for BR*Tools to use the BR*Tools command.

About this task

If the path is not set and you try to perform the system dump operation by using BR*Tools, the following error message is logged in `config_summary.txt`: `'brtools' is not recognized as an internal or external command, operable program or batch file`

Steps

1. Click **Start** > **My Computer** and then right-click **Properties** > **Advanced** > **Environmental Variables**.

2. In the **Environmental Variables** window, click **New** to add new BR*tools path or **Edit** to edit the existing BR*Tools path for both the system variables and user variables.

3. Click **OK** to save your changes.

Disabling client access to Snapshot copies

For any storage system volume that uses the NFS protocol and contains SAP data backed up using BR*Tools, you must disable client access to the Snapshot copies of that volume. If client access is enabled, BR*Tools will attempt to create backups of the hidden .snapshot directories that contain previous backups.

To disable client access, use one of the following methods.

- Using Data ONTAP: Using the Data ONTAP command line interface, enter the following command: `vol options volume_name nosnapdir on` wherein `volume_name` is the name of the volume containing

SAP data, for example, `/vol/falls_sap_cer9i_data1`

- Using FilerView: To disable access using the FilerView graphical user interface, you must already have a Snapshot copy of the volume.

  a. In the FilerView left pane, select **Volumes** > **Snapshots** > **Manage**.

  b. On the Manage Snapshots page, click the volume name in the Volume column.

  c. On the Configure Snapshots page, clear the Snapshot Directory Visible check box, and click **Apply.**

For more information about disabling client access to Snapshot copies, see the chapter about Snapshot Management of the *Data ONTAP Data Protection Online Backup and Recovery Guide* for your version of Data ONTAP.

Profile usage for BR*Tools backups

When you run a BR*Tools command that uses the BACKINT interface, SnapManager uses a profile from the repository. The repository is determined by the SnapManager credentials of the user running the BR*Tools command.

By default, SnapManager uses the profile with the same name as the SAP database system identifier.

About creating credentials to access the repository

You can use the `smsap credential set` command to set the repository credentials for the BR*Tools user.

About specifying a different SnapManager profile name

As long as the system identifier is unique for all hosts whose SnapManager profiles are in a given repository, the default profile name is sufficient. You can create the SnapManager profile and name it by using the database system identifier value.

However, if the same system identifier is used on different hosts, or if you want to specify more than one SnapManager profile for use with BR*Tools on a given SAP instance, then you need to define the profile name for BR*Tools commands.

Database operations scheduled within the SAP application run as user. BR*Tools operations scheduled within the SAP applications run as `SAPServiceSID`. These users must have access to the repository and the profile.

The SAP service user SAPServiceSID does not allow the user to log in locally. To change this, you must modify the Local Security Policy by performing the following steps:

Steps

1. Click **Start** > **Programs** > **Administrative Tools** > **Local Security Policy**.

2. Expand **Local Policies** and select**User Rights Assignment**

3. Search for the **Deny log on locally** policy.

4. Right-click and select **Properties**.

5. Delete **SAP Service user** from this list.

About creating the backup utility parameter file

The BR*Tools command can optionally pass a backup utility parameter (`.utl` parameter) file to the BACKINT

interface. By default, the name of this file is `initSID.utl`, where SID is the system identifier of the database.

By default, BR*Tools uses the parameter file specified by the `util_par_file` parameter in the `initSID.sap` file. The backup utility parameter file is normally stored in the same directory as the `initSID.sap` file.

Add `profile_name = <profile>` to the `.utl` file and save. The profile is the name of the SnapManager *profile* you want to use for the BR*Tools commands.

The following table lists additional backup utility parameters included in the `.utl` file for operations like backup retention, fast restore, and data protection:

Operations	Backup utility parameters
Backup retention	• Key: retain • Value: unlimited
hourly	daily
weekly	monthly

The following table describes the different backup utility parameters for these operations:

Operation	Key	Value
Backup retention	fast=	fallback

Adding the backup utility parameter to the BR*Tools initialization file

SAP creates a backup profile file named `initSID.sap` for each SAP database instance in `$ORACLE_HOME\database\`. You can use this file to specify the default backup utility parameter (`.utl`) file used for BR*Tools commands.

Steps

1. Edit the initSID.sap file and locate the line starting with `util_par_file =`.

2. Uncomment this line and add the path to the backup utility parameter file that contains the name of the profile, for example, `util_par_file = initSA1.utl`.

3. If you specify a value for the `util_par_file`, be sure that the file exists.

   If the file is not found, the BRBACKUP commands fail because they attempt to include that file in a backup.

If any of the following error conditions are found, then you need to use the `-u` option as the profile name:

• `Parameter files does not exist`

• `No profile_name entry`

Both `orasid` and `sidadm` need access to the profile that would be used for creating or managing backups created with BR*Tools.

Specifying the backup utility parameter file name on the BR*Tools command

You can optionally specify the backup utility parameter (.utl parameter) file on the BR*Tools commands using the `-r` option. The value on the command line overrides the value specified in the SAP initialization file.

BR*Tools searches for the parameter file in the `$ORACLE_HOME\database\` directory. If you store the file in another location, you need to supply the full path with the `-r` option. For example:

brbackup -r
C:\Oracle\NetApp_fcp_price_10g_enterprise_inst_vol1\database\initCER.utl …

Database backups created using BRBACKUP and BRARCHIVE

The `BRBACKUP` command creates a backup of the SAP database using Snapshot copy on the storage system. You can run the `BRBACKUP` command from the SAP host's command-line interface (CLI), or BR*Tools CLI or graphical user interface (GUI). To back up the offline redo log files, you can use the `BRARCHIVE` command.

`BRBACKUP` backs up the data files, control files, and online redo log files of the SAP database. You should back up the other SAP configuration files, for example, SAP log files, kernel files, and transports requests using `BRBACKUP` with the `SAP_DIR` option, and restore using `BRRESTORE`.

If the value of the *SQLNET.AUTHENTICATION_SERVICES* parameter in the `sqlnet.ora` is set to **NONE**, ensure that the Oracle database user (system) has sysoper privileges. The system is the default user created during the creation of the Oracle database. To enable SYSOPER privileges, run the following command: `grant sysoper to system;`

You should use either the BR*Tools commands or the SnapManager commands to manage the archive log files.

You should not use the combination of SnapManager profiles and the BR*Tools commands for managing the archive log backups with respect to the following operations:

- Creating separate backups of data files and archive log files using the `BRBACKUP` and `BRARCHIVE` commands with the SnapManager profile (created using the option to separate archive log backups)
- Pruning archive log files while creating the backup using the SnapManager profile

> ⓘ SnapManager does not display any warning or error message when you use a combination of SnapManager for SAP profiles and the BR*Tools commands for managing the archive log files.

You should create a profile without using the option to separate archive log backups and use this profile for regular BR*Tools operations.

You can restore the backup by using `BRRESTORE`. For database backups (backups include data files, control files, or online redo log files) created by using BRBACKUP with the `-m all` or `-m full` option, you can also restore the backup by using the SnapManager CLI or GUI.

For more information about how SnapManager handles backing up, see *Backing up databases*.

See the SAP documentation for the specific instructions and syntax for the `BRBACKUP` and `BRARCHIVE` commands. Before using the BRBACKUP and BRARCHIVE command with a storage system, ensure that the following conditions are met:

- If the SnapManager profile name is different from the SAP database system identifier, specify the name of the parameter file that contains the SnapManager profile name.

  You can do so in one of the following ways:

  ◦ Specify the backup utility parameter file (`initSID.utl`) using the `-r` option with the `BRBACKUP` command.

  ◦ Specify the parameter file in the initialization file (`initSID.sap`). For more information, see *Specifying the profile for BR*Tools backups*.

- If it is appropriate for your environment, create additional backups of your SAP data using a different medium, such as tape. Snapshot copies are intended for rapid backup and restoration. They are stored on the same physical media as the data being backed up and are not intended for disaster recovery unless copied to other storage devices.

Schedule backups with SAP transaction DB13

There are several steps required to enable backups by SnapManager for SAP to work from the DBA Planning Calendar transaction DB13.

Steps

1. Change the permissions of the `/oracle/SID/sapbackup` directory to allow write access by the sidadm user identifier.

   Use the following command as `orasid`:

   chmod 775 /oracle/SID/sapbackup

   This will allow members of the dbagroup to write to that directory. The user identifier sidadm is a member of the dba group.

2. Change the permissions of the SnapManager for SAP credentials file for `orasid` to allow `sidadm` access. Use the following command executed by `orasid`:

   Use the following command executed by `orasid`:

   chmod 660 /oracle/SID/.netapp/smsap/credentials

   Follow the directions in section 3.3, "Registering Systems in the Repository" of TR-3582 "SnapManager for SAP Best Practices" to register the sidadm user identifier in the SnapManager for SAP repository.

Database restore by using BRRESTORE and BRRECOVER

The `BRRESTORE` and `BRRECOVER` commands work only with backups created using `BRBACKUP`.

See the SAP documentation for specific instructions and syntax for the `BRRESTORE` and `BRRECOVER` commands.

If the SnapManager profile name is different from the SAP database system identifier, you must specify the name of the parameter file that contains the SnapManager profile name. You can do this in one of the following ways:

- Specify the backup utility parameter file (`initSID.utl`) using the `-r` option in the `BRBACKUP` command.
- Specify the parameter file in the BR*Tools initialization file (`initSID.sap`).

> ⓘ Backups (data files, control files, or online redo log files) that are created by using `BRBACKUP` can be restored from secondary or tertiary storage systems by using the `smsap restore` command. However, it is recommended to back up the other SAP configuration files, for example, SAP log files, kernel files, and transports requests by using `BRBACKUP` with `SAP_DIR`, and restore by using `BRRESTORE`.

Backing up and restoring files using BR*Tools

In addition to backing up database files, you can use BR*Tools to back up and restore any files, such as the SAP system files stored on the storage system.

Follow the SAP BR*Tools documentation for the BRBACKUP and BRRESTORE commands. The following additional information applies when you are using SnapManager for SAP:

- The files must be stored on the storage system.
- Be sure you have permissions for the files you are backing up. For example, backing up SAP system files (`BRBACKUP backup_mode=sap_dir`) or Oracle system files (`BRBACKUP backup_mode=ora_dir`) requires root permission on the host. Use the `su root` command before running BRBACKUP.

Restoring a backup to a different host

You can use the BRRESTORE or `smsap restore` command to restore a backup that you made with BRBACKUP on a different host. The new host must also be running SnapManager for SAP.

Make SnapManager profile available

Before you can run BRRESTORE, you must make the SnapManager profile of the original host available to the new host. Complete the following steps:

- Set credential for repository: Use the `smsap credential set` command to enable the new host to access the SnapManager repository used for the original backup.
- Set credential for profile: Use the `smsap credential set` command to enable the new host to access the SnapManager profile used for the original backup.
- Load the profile on the new host: Use the `smsap profile sync` command to load the SnapManager profile information onto the new host.

Run BRRESTORE

Configure new storage on the new host with the same paths as the original storage on the original host.

Use the `BRRESTORE` command to restore the backup to the new host. By default, the files are restored to the original path locations. You can specify an alternate path for the restored files using the `-m` option of the `BRRESTORE` command. See the SAP documentation for more information on `BRRESTORE`.

SnapManager for SAP command reference

The SnapManager command reference includes the valid usage syntax, options, parameters, and arguments you should supply with the commands, along with examples.

The following issues apply to command usage:

- Commands are case-sensitive.
- SnapManager accepts up to 200 characters and labels up to 80 characters.
- If the shell on your host limits the number of characters that can appear on a command line, you can use the `cmdfile` command.
- Do not use spaces in profile names or label names.
- In the clone specification, do not use spaces in the clone location.

SnapManager can display three levels of messages to the console:

- Error messages
- Warning messages
- Informational messages

You can specify how you want messages displayed. If you specify nothing, SnapManager displays only error messages and warnings to the console. To control the amount of output that SnapManager displays on the console, use one of the following command line options:

- `-quiet`: Displays only error messages to the console.
- `-verbose`: Displays error, warning, and informational messages to the console.

> Regardless of the default behavior, or the level of detail you specify for the display, SnapManager always writes all message types to the log files.

The backint register-sld command

If you want to use SAP BR*Tools, you can run the `backint register-sld` command with SnapManager for SAP to register the BACKINT interface at your System Landscape Directory (SLD). The BACKINT interface is provided by the storage vendor to make the storage system work with the BR*Tools commands. SnapManager for SAP installs its BACKINT interface file at `C:\Program Files\NetApp\SnapManager for SAP\bin\`.

Syntax

```
backint register-sld
-host  host_name
-port  port_id
-username username
-password password
-template template_ID
```

Parameters

- **-host** *host_name*

  Specifies the name of the host on which the SAP SLD is running.

- **-port** *port_id*

  Specifies the ID of the port on which the SAP SLD accepts HTTP requests. The ID must be numeric and less than or equal to nine digits.

- **-username** *username*

  Specifies a valid and authorized SAP SLD user name.

- **-password** *password*

  Specifies a valid and authorized SAP SLD user password. This is optional. If you do not set the password by using `-password`, you will be prompted to enter a password in the following format: `username@http://host:port`

  If you fail to enter the correct password, after three attempts the command fails and exits.

- **-template** *template_id*

  Specifies the name of the master template XML file, which can be used as a base for your custom file.

Example command

The following example shows a successful completion of the command:

```
backint register-sld -host jack12 -port 50100
-username j2ee_admin -password user123 -template E:\template.xml
Operation Id [N96f4142a1442b31ee4636841babbc1d7] succeeded.
```

The smsap_server restart command

This command restarts the SnapManager host server and is entered as root.

Syntax

```
smsap_server restart
[-quiet | -verbose]
```

Parameters

* **-quiet**

  Specifies that only error messages are displayed on the console. The default is to display error and warning messages.

* **-verbose**

  Specifies that error, warning, and informational messages are displayed on the console.

Example command

The following example restarts the host server.

```
smsap_server restart
```

The **smsap_server start** command

This command starts the host server running the SnapManager for SAP software.

Syntax

```
smsap_server start
[-quiet | -verbose]
```

Parameters

* **-quiet**

  Specifies that only error messages are displayed on the console. The default is to display error and warning messages.

* **-verbose**

  Specifies that error, warning, and informational messages are displayed on the console.

Example command

The following example starts the host server.

```
smsap_server start
SMSAP-17100: SnapManager Server started on secure port 25204 with PID
11250
```

The **smsap_server status** command

You can run the `smsap_server status` command to view the status of the SnapManager host server.

Syntax

```
smsap_server status
[-quiet | -verbose]
```

Parameters

- **-quiet**

  Specifies that only error messages are displayed in the console. The default is to display error and warning messages.

- **-verbose**

  Specifies that error, warning, and informational messages are displayed in the console.

Example

The following example displays the status of the host server:

```
smsap_server status
SMSAP-17104: SnapManager Server version 3.3.1 is running on secure port
25204 with PID 11250
and has 0 operations in progress.
```

The **smsap_server stop** command

This command stops the SnapManager host server and is entered at the root.

Syntax

```
smsap_server stop
[-quiet | -verbose]
```

Parameters

* **-quiet**

    Specifies that only error messages are displayed on the console. The default is to display error and warning messages.

* **-verbose**

    Specifies that error, warning, and informational messages are displayed on the console.

Example command

The following example uses the `smsap_server stop` command.

```
smsap_server stop
```

The smsap backup create command

You can run the `backup create` command to create database backups on one or more storage systems.

Syntax

> ℹ️ Before you run this command, you must create a database profile by using the `profile create` command.

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

[-backup-dest path1 [ , path2]]
[-exclude-dest path1 [ , path2]]
[-prunelogs {-all | -until-scn until-scn | -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]
```

Parameters

* **-profile *profile_name***

  Specifies the name of the profile related to the database you want to back up. The profile contains the identifier of the database and other database information.

* **-auto option**

  If the database is in the mounted or offline state, SnapManager performs an offline backup. If the database is in the open or online state, SnapManager performs an online backup. If you use the `-force` option with the `-offline` option, SnapManager forces an offline backup even if the database is currently online.

* **-online option**

  Specifies an online database backup.

  ◦ If the local instance is in the shutdown state and at least one instance is in the open state, you can use the `-force` option to change the local instance to the mounted state.

  ◦ If no instance is in open state, you can use the `-force` option to change the local instance to open state.

* **-offline option**

  Specifies an offline backup while the database is in the shut down state. If the database is in the open or mounted state, the backup fails. If the `-force` option is used, SnapManager attempts to alter the database state to shut down the database for an offline backup.

* **-full option**

Backs up the entire database. This includes all of the data, archived log, and control files. The archived redo logs and control files are backed up no matter what type of backup you perform. If you want to back up only a portion of the database, use the `-files` option or `-tablespaces` option.

* **`-data option`**

  Specifies the data files.

* **`-files list`**

  Backs up only the specified data files plus the archived log and control files. Separate the list of file names with spaces. If the database is in the open state, SnapManager verifies that the appropriate tablespaces are in online backup mode.

* **`-tablespaces tablespaces`**

  Backs up only the specified database tablespaces plus the archived log and control files. Separate the tablespace names with spaces. If the database is in the open state, SnapManager verifies that the appropriate tablespaces are in online backup mode.

* **`-label label`**

  Specifies an optional name for this backup. This name must be unique within the profile. The name can contain letters, numbers, underscores (_), and hyphens (-). It cannot start with a hyphen. If you do not specify a label, SnapManager creates a default label in the scope_type_date format:

  ◦ Scope is either F to indicate a full backup or P to indicate a partial backup.

  ◦ Type is C to indicate an offline (cold) backup, H to indicate an online (hot) backup, or A to indicate auto backup, for example, P_A_20081010060037IST.

  ◦ Date is the year, month, day, and time of the backup.

    SnapManager uses a 24-hour clock.

  For example, if you performed a full backup with the database offline on 16th January 2007, at 5:45:16 p.m. Eastern Standard Time, SnapManager would create the label F_C_20070116174516EST.

* **`-comment string`**

  Specifies an optional comment to describe this backup. Enclose the string in single quotation marks (').

  > ⓘ Some shells strip the quotation marks off. In this case, you must include the quotation mark with a backslash (\). For example, you might need to enter the following: `\' this is a comment\'`.

* **`-verify option`**

  Verifies that the files in the backup are not corrupt by running the Oracle dbv utility.

  > ⓘ If you specify the `-verify` option, the backup operation is not completed until the verify operation is complete.

* **`-force option`**

Forces a state change if the database is not in the correct state. For example, SnapManager might change the state of the database from online to offline, based on the type of backup you specify and the state that the database is in.

- If the local instance is in the shutdown state and at least one instance is in the open state, then using the `-force` option changes the local instance to the mounted state.

- If no instance is in the open state, using the `-force` option changes the local instance to the open state.

- **`-quiet`**

  Displays only error messages in the console. The default is to display error and warning messages.

- **`-verbose`**

  Displays error, warning, and informational messages in the console.

- **`-retain { -hourly | -daily | -weekly | -monthly | -unlimited}`**

  Specifies whether the backup should be retained on an hourly, daily, weekly, monthly, or unlimited basis. If the `-retain` option is not specified, the retention class defaults to `-hourly` option. To retain backups forever, use the `-unlimited` option. The `-unlimited` option makes the backup ineligible for deletion by the retention policy.

- **`-archivelogs option`**

  Creates archive log backup.

- **`-backup-dest path1, [, [path2]]`**

  Specifies the archive log destinations to be backed up for archive log backup.

- **`-exclude-dest path1, [, [path2]]`**

  Specifies the archive log destinations to be excluded from the backup.

- **`-prunelogs {-all | -until-scnuntil-scn | -until-date yyyy-MM-dd:HH:mm:ss | -before {-months | -days | -weeks | -hours}`**

  Deletes the archive log files from the archive log destinations based on options provided while creating a backup. The `-all` option deletes all of the archive log files from the archive log destinations. The `-until-scn` option deletes the archive log files until a specified System Change Number (SCN). The `-until-date` option deletes the archive log files until the specified time period. The `-before` option deletes the archive log files before the specified time period (days, months, weeks, hours).

- **`-prune-dest prune_dest1,prune_dest2`**

  Deletes the archive log files from the archive log destinations while creating the backup.

- **`-taskspec taskspec`**

  Specifies the task specification XML file that can be used for preprocessing activity or post-processing activity of the backup operation. The complete path of the XML file should be provided while giving the `-taskspec` option.

- **-dump option**

  Collects the dump files after a successful or failed database backup operation.

Example command

The following command creates a full online backup, creates a backup to secondary storage, and sets the retention policy to daily:

```
smsap backup create -profile SALES1 -full -online
-label full_backup_sales_May -profile SALESDB -force -retain -daily
Operation Id [8abc01ec0e79356d010e793581f70001] succeeded.
```

The smsap backup delete command

You can run the `backup delete` command to remove backups that are not automatically removed, such as backups that were used to create a clone or backups that failed. You can delete backups retained on an unlimited basis without changing the retention class.

Syntax

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

Parameters

- **-profile** *profile_name*

  Specifies the database associated with the backup you want to remove. The profile contains the identifier of the database and other database information.

- **-id** *guid*

  Specifies the backup with the specified GUID. The GUID is generated by SnapManager when you create a backup. You can use the `smsap backup list` command to display the GUID for each backup.

- **-label** *label*

  Specifies the backup with the specified label. Optionally, specify the scope of the backup as data file or archive log.

- **-data**

Specifies the data files.

* **-archivelogs**

  Specifies the archive log files.

* **-all**

  Specifies all backups. To delete only specified backups instead, use the `-id` or `-label` option.

* **-dump**

  Collects the dump files after a successful or failed backup delete operation.

* **-force**

  Forces the removal of the backup. SnapManager removes the backup even if there are problems in freeing the resources associated with the backup. For example, if the backup was cataloged with Oracle Recovery Manager (RMAN), but the RMAN database no longer exists, including `-force` deletes the backup even though it cannot connect with RMAN.

* **-quiet**

  Displays only error messages in the console. The default is to display error and warning messages.

* **-verbose**

  Displays error, warning, and informational messages in the console.

Example

The following example deletes the backup:

```
smsap backup delete -profile SALES1 -label full_backup_sales_May
Operation Id [8abc01ec0e79004b010e79006da60001] succeeded.
```

The smsap backup free command

You can run the `backup free` command to free the Snapshot copies of the backups without removing the backup metadata from the repository.

Syntax

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

Parameters

* **-profile** *profile_name*

   Specifies the profile associated with the backup you want to free. The profile contains the identifier of the database and other database information.

* **-id** *guid*

   Specifies the resources of the backup with the specified GUID. The GUID is generated by SnapManager when you create a backup. You can use the `smsap backup list` command to display the GUID for each backup. Include the `-verbose` option to display the backup IDs.

* **-label** *label*

   Specifies the backup with the specified label.

* **-data**

   Specifies the data files.

* **-archivelogs**

   Specifies the archive log files.

* **-all**

   Specifies all backups. To delete specified backups instead, use the `-id` or `-label` option.

* **-force**

   Forces the removal of the Snapshot copies.

* **-quiet**

   Displays only error messages in the console. The default is to display error and warning messages.

* **-verbose**

   Displays error, warning, and informational messages in the console.

Example

The following example frees the backup:

```
smsap backup free -profile SALES1 -label full_backup_sales_May
Operation Id [8abc01ec0e79004b010e79006da60001] succeeded.
```

The smsap backup list command

You can run the `backup list` command to display information about the backups in a profile, including information about the retention class and protection status.

Syntax

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

Parameters

* **-profile** *profile_name*

  Specifies the profile you want to list backups for. The profile contains the identifier of the database and other database information.

* **-delimiter** *character*

  Displays each row on a separate line. The attributes in the row are separated by the character specified.

* **-data**

  Specifies the data files.

* **-archivelogs**

  Specifies the archive log files.

* **-quiet**

  Displays only error messages in the console. The default is to display error and warning messages.

* **-verbose**

  Displays error, warning, and informational messages in the console. Include the -verbose option to display the backup IDs.

Example

The following example lists the backups for the SALES1 profile:

```
smsap backup list -profile SALES1 -verbose
Start Date           Status  Scope  Mode     Primary  Label      Retention
Protection
------------------- ------- ------ ------- -------- -----      --------
----------
2007-08-10 14:12:31 SUCCESS FULL   ONLINE  EXISTS   backup2    HOURLY
NOT REQUESTED
2007-08-05 12:08:37 SUCCESS FULL   ONLINE  EXISTS   backup4    UNLIMITED
NOT REQUESTED
2007-08-04 22:03:09 SUCCESS FULL   ONLINE  EXISTS   backup6    UNLIMITED
NOT REQUESTED
```

The smsap backup mount command

You can run the `backup mount` command to mount a backup in order to perform a recover operation by using an external tool.

Syntax

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

[-dump]
[-quiet | -verbose]
```

Parameters

* **-profile** *profile_name*

  Specifies the profile associated with the backup that you want to mount. The profile contains the identifier of the database and other database information.

* **-id** *guid*

  Mounts the backup with the specified GUID. The GUID is generated by SnapManager when you create a backup. You can use the `smsap backup list` command to display the GUID for each backup.

* **-label** *label*

  Mounts the backup with the specified label.

- **-data**

  Specifies the data files.

- **-archivelogs**

  Specifies the archive log files.

- **-host** *host*

  Specifies the host on which you want to mount the backup.

- **-dump**

  Collects the dump files after the successful or failed mount operation.

- **-quiet**

  Displays only error messages in the console. The default setting is to display error and warning messages.

- **-verbose**

  Displays error, warning, and informational messages in the console.

> ℹ️ You must use this command only if you are using an external tool such as Oracle Recovery Manager (RMAN). SnapManager automatically handles the mounting of backups if you use the `smsap backup restore` command to restore the backup. This command displays a list, which shows the paths where the Snapshot copies have been mounted. This list is displayed only when the `-verbose` option is specified.

Example

The following example mounts the backup:

```
smsap backup mount -profile S10_BACKUP -label full_monthly_10 -verbose
[INFO ]: SMSAP-13051: Process PID=6852
[INFO ]: SMSAP-13036: Starting operation Backup Mount on host
hadley.domain.private
[INFO ]: SMSAP-13036: Starting operation Backup Mount on host
hadley.domain.private
[INFO ]: SMSAP-13046: Operation GUID 8abc01573883daf0013883daf5ac0001
starting on Profile FAS_P1
[INFO ]: SD-00025: Beginning to connect filesystem(s) [I:\] from snapshot
smsap_fas_p1_fasdb_d_h_2_8abc0157388344bc01388344c2d50001_0.
[INFO ]: SD-00016: Discovering storage resources for
C:\SnapManager_auto_mounts\I-2012071400592328_0.
[INFO ]: SD-00017: Finished storage discovery for
C:\SnapManager_auto_mounts\I-2012071400592328_0
[INFO ]: SD-00026: Finished connecting filesystem(s) [I:\] from snapshot
smsap_fas_p1_fasdb_d_h_2_8abc0157388344bc01388344c2d50001_0.
[INFO ]: SD-00025: Beginning to connect filesystem(s) [H:\] from snapshot
smsap_fas_p1_fasdb_d_h_1_8abc0157388344bc01388344c2d50001_0.
[INFO ]: SD-00016: Discovering storage resources for
C:\SnapManager_auto_mounts\H-2012071400592312_0.
[INFO ]: SD-00017: Finished storage discovery for
C:\SnapManager_auto_mounts\H-2012071400592312_0.
[INFO ]: SD-00026: Finished connecting filesystem(s) [H:\] from snapshot
smsap_fas_p1_fasdb_d_h_1_8abc0157388344bc01388344c2d50001_0.
[INFO ]: SMSAP-13048: Backup Mount Operation Status: SUCCESS
[INFO ]: SMSAP-13049: Elapsed Time: 0:19:05.620
```

The smsap backup restore command

You can run the `backup restore` command to restore backups of a database or a
portion of a database, and then optionally recover the database information.

Syntax

```
smsap backup restore
-profile profile_name
[-label label | -id id]
[-files files [files...] |
-tablespaces tablespaces [tablespaces...]] |
-complete | -controlfiles]
[-recover {-alllogs | -nologs | -until until} [-using-backup-controlfile]]
[-restorespec restorespec | ]]
[-preview]

[-recover-from-location path1 [, path2]]
[-taskspec taskspec]
[-dump]
[-force]
[-quiet | -verbose]
```

Parameters

* **-profile *profile_name***

  Specifies the database that you want to restore. The profile contains the identifier of the database and other database information.

* **-label *name***

  Restores the backup with the specified label.

* **-id *guid***

  Restores the backup with the specified GUID. The GUID is generated by SnapManager when you create a backup. You can use the `smsap backup list` command to display the GUID for each backup.

* **Choose all or specified files**

  Optionally, you can use one of the following options:

  ○ `-complete`: Restores all the data files in the backup.

  ○ `-tablespaceslist`: Restores only the specified tablespaces from the backup.

    You must use spaces to separate the names in the list.

  ○ `-fileslist`: Restores only the specified data files from the backup.

    You must use spaces to separate the names in the list. If the database is running, SnapManager ensures that the tablespace containing the files is offline.

* **-controlfiles**

  Restores the control files. SnapManager allows you to restore control files along with the data files from the

backups in a single operation. The `-controlfiles` option is independent of other restore scope parameters such as `-complete`, `-tablespaces`, and `-files`.

- **`-recover`**

  Recovers the database after restoring it. You must also specify the point to which you want SnapManager to recover the database by using one of the following options:

  - `-nologs`: Recovers the database to the time of the backup and applies no logs.

    You can use this parameter for online or offline backups.

  - `-alllogs`: Recovers the database to the last transaction or commit, and applies all required logs.

  - `-until date`: Recovers the database up to the date and time specified.

    You must use the year-month-date: hour: minute: second (*yyyy-mm-dd:hh:mm:ss*) format. For hours, use either 12-hour or 24-hour format, depending on the database setting.

  - `-until scn`: Rolls forward the data files until it reaches the specified system change number (SCN).

  - `-using-backup-controlfile`: Recovers the database using the backup control file.

- **`-restorespec`**

  Enables you to restore the data to an active file system and restore from the specified data by providing a mapping of each original Snapshot copy to its active file system. If you do not specify an option, SnapManager restores the data from the Snapshot copies on primary storage. You can specify one of the following options:

  - `-restorespec`: Specifies the data to restore and the restore format.

- **`-preview`**

  Displays the following information:

  - Which restore mechanism (storage-side file system restore, storage-side file restore, or host-side file copy restore) will be used to restore each file

  - Why more efficient mechanisms were not used to restore each file, when you specify the `-verbose` option If you are using the `-preview` option, you must know the following:

  - The `-force` option has no impact on the command.

  - The `-recover` option has no impact on the command. To preview the restore operation, the database must be mounted. If you want to preview a restore plan, and the database currently is not mounted, then SnapManager mounts the database. If the database cannot be mounted, then the command will fail, and SnapManager returns the database to its original state.

  The `-preview` option displays up to 20 files. You can configure the maximum number of files to be displayed in the `smsap.config file`.

- **`-recover-from-location`**

  Specifies the external archive log location of the archive log files. SnapManager takes the archive log files from the external location and uses them for the recovery process.

- **-taskspec**

  Specifies the task specification XML file for preprocessing activity or post-processing activity of the restore operation. You must provide the complete path of the task specification XML file.

- **-dump**

  Specifies to collect the dump files after the restore operation.

- **-force**

  Changes the database state to a lower state than its current state, if necessary.

  By default, SnapManager can change the database state to a higher state during an operation. This option is not required for SnapManager to change the database to a higher state.

- **-quiet**

  Displays only error messages in the console. The default setting is to display error and warning messages.

- **-verbose**

  Displays error, warning, and informational messages in the console. You can use this option to see why more efficient restore processes could not be used to restore the file.

Example

The following example restores a database along with the control files:

```
smsap backup restore -profile SALES1 -label full_backup_sales_May
-complete -controlfiles -force
```

The smsap backup show command

You can use the `backup show` command to display detailed information about a backup, including its protection status, backup retention class, and backups on primary and secondary storage.

Syntax

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

Parameters

- **-profile** *profile_name*

Specifies the profile for which to show backups. The profile contains the identifier of the database and other database information.

- **-label** *label*

  Specifies the label of the backup.

- **-data**

  Specifies the data files.

- **-archivelogs**

  Specifies the archive log files.

- **-id** *id*

  Specifies the backup ID.

- **-quiet**

  Displays only error messages in the console. The default is to display error and warning messages.

- **-verbose**

  Displays error, warning, and informational messages in the console, as well as any clone and verification information.

Example

The following example shows detailed information about the backup:

```
smsap backup show -profile SALES1 -label BTNFS -verbose
Backup id: 8abc013111a450480111a45066210001
Backup status: SUCCESS
Primary storage resources: EXISTS
Protection sate: NOT REQUESTED
Retention class: DAILY
Backup scope: FULL
Backup mode: OFFLINE
Mount status: NOT MOUNTED
Backup label: BTNFS
Backup comment:

Backup start time: 2007-03-30 15:26:30
Backup end time: 2007-03-30 15:34:13
Verification status: OK
Backup Retention Policy: NORMAL
Backup database: hsdb1
Checkpoint: 2700620
Tablespace: SYSAUX
Datafile: E:\disks\data\hsdb\sysaux01.dbf [ONLINE]
...
Control Files:
File: E:\disks\data\control03.ctl
...
Archive Logs:
File: E:\disks\data\archive_logs\2_131_626174106.dbf
...
Volume: hs_data
Snapshot: SMSAP_HSDBR_hsdb1_F_C_1_
8abc013111a450480111a45066210001_0
File: E:\disks\data\hsdb\SMSAPBakCtl_1175283005231_0
...
```

The smsap backup unmount command

You can run the `backup unmount` command to unmount a backup.

Syntax

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

Parameters

* **-profile *profile_name***

  Specifies the profile for which you want to unmount a backup. The profile contains the identifier of the database and other database information.

* **-id *id***

  Unmounts the backup with the specified GUID. The GUID is generated by SnapManager when you create a backup. You can use the `smsap backup list` command to display the GUID for each backup.

* **-label *label***

  Unmounts the backup with the specified label.

* **-data**

  Specifies the data files.

* **-archivelogs**

  Specifies the archive log files.

* **-dump**

  Collects the dump files after a successful or failed unmount operation.

* **-force**

  Unmounts the backup even if there are problems in freeing the resources associated with the backup. SnapManager tries to unmount the backup and clean up any associated resources. The log shows the unmount operation as successful, but you may have to manually clean up resources if there are errors in the log.

* **-quiet**

  Displays only error messages in the console. The default is to display error and warning messages.

* **-verbose**

  Displays error, warning, and informational messages in the console.

Example

The following is an example of an unmount operation:

```
# smsap backup unmount -label test -profile SALES1 -verbose
```

```
[INFO ]: SMSAP-13051: Process PID=9788
[INFO ]: SMSAP-13036: Starting operation Backup Unmount on host
hadley.domain.private
[INFO ]: SMSAP-13036: Starting operation Backup Unmount on host
hadley.domain.private
[INFO ]: SMSAP-13046: Operation GUID 8abc015738849a3d0138849a43900001
starting on Profile FAS_P1
[INFO ]: SD-00031: Beginning to disconnect filesystem(s)
[C:\SnapManager_auto_mounts\H-2012071400592312_0,
C:\SnapManager_auto_mounts\I-2012071400592328_0].
[INFO ]: SD-00032: Finished disconnecting filesystem(s)
[C:\SnapManager_auto_mounts\H-2012071400592312_0,
C:\SnapManager_auto_mounts\I-2012071400592328_0].
[INFO ]: SMSAP-13048: Backup Unmount Operation Status: SUCCESS
[INFO ]: SMSAP-13049: Elapsed Time: 0:07:26.754
```

The smsap backup update command

You can run the `backup update` command to update the backup retention policy.

Syntax

```
smsap backup update
-profile profile_name
[-label label [-data | -archivelogs] | [-id guid]
[-retain {-hourly | -daily | -weekly | -monthly | -unlimited}]
[-comment comment_text]
[-quiet | -verbose]
```

Parameters

* **-profile** *profile_name*

  Specifies the profile for which to update backups. The profile contains the identifier of the database and other database information.

* **-id** *guid*

  Verifies the backup with the specified GUID. The GUID is generated by SnapManager when you create a

backup. You can use the `smsap backup list` command to display the GUID for each backup.

- **-label** *label*

  Specifies the backup label and scope of the backup as data file or archive log.

- **-data**

  Specifies the data files.

- **-archivelogs**

  Specifies the archive log files.

- **-comment** *comment_text*

  Enter text (up to 200 characters) about the backup update. You can include spaces.

- **-quiet**

  Displays only error messages in the console. The default is to display error and warning messages.

- **-verbose**

  Displays error, warning, and informational messages in the console.

- **-retain {-hourly | -daily | -weekly | -monthly | -unlimited}**

  Specifies whether the backup should be retained on an hourly, daily, weekly, monthly, or unlimited basis. If `-retain` is not specified, the retention class defaults to `-hourly`. To retain backups forever, use the `-unlimited` option. The `-unlimited` option makes the backup ineligible for deletion.

Example

The following example updates the backup to be set the retention policy to unlimited:

```
smsap backup update -profile SALES1 -label full_backup_sales_May
-retain -unlimited -comment save_forever_monthly_backup
```

The smsap backup verify command

You can run the `backup verify` command to see if the backup is in a valid format for Oracle.

Syntax

```
smsap backup verify
-profile profile_name
[-label backup_name | [-id guid]
[-retain {-hourly | -daily | -weekly | -monthly | -unlimited}]
[-force]
[-dump]
[-quiet | -verbose]
```

Parameters

- **-profile** *profile_name*

  Specifies the profile for which you want to verify a backup. The profile contains the identifier of the database and other database information.

- **-id** *guid*

  Verifies the backup with the specified GUID. The GUID is generated by SnapManager when you create a backup. You can use the `smsap backup list` command to display the GUID for each backup.

- **-label** *label_name*

  Verifies the backup with the specified label.

- **-dump**

  Collects the dump files after the successful or failed backup verify operation.

- **-force**

  Forces the database into the necessary state to perform the verify operation.

- **-quiet**

  Displays only error messages in the console. The default is to display error and warning messages.

- **-verbose**

  Displays error, warning, and informational messages in the console.

Example

The following is an example of verifying the backup:

```
smsap backup verify -profile SALES1 -label full_backup_sales_May -quiet
```

```
DBVERIFY - Verification starting : FILE = C:\SnapManager_auto_mounts\H-
2012071400592312_0\smsap\datafile\data
```

The smsap clone create command

You can run the `clone create` command to create a clone of a backed-up database. You can clone a backup from primary or secondary storage.

Syntax

```
smsap clone create
-profile profile_name
[-backup-id backup_guid | -backup-label backup_label_name | -current]
-newsid new_sid
[-host target_host]
[-label clone_label]
[-comment string]
-clonespec full_path_to_clonespec_file
]
[-syspassword syspassword]
[-reserve {yes | no | inherit}]

[-no-resetlogs | -recover-from-location path1 [, path2]][-taskspec
taskspec]
[-dump]
[-quiet | -verbose]
```

Parameters

* **-profile** *name*

  Specifies the database that you want to clone. The profile contains the identifier of the database and other database information.

* **-backup-id** *guid*

  Clones the backup with the specified GUID. The GUID is generated by SnapManager when you create a backup. You can use the `smsap backup list-verbose` command to display the GUID for each backup.

* **-backup-label** *backup_label_name*

  Specifies to clone the backup with the specified label name.

* **-current**

  Specifies to create backup and clone from the current state of the database.

> **ⓘ** If the database is in the `noarchivelog` mode, SnapManager will create an offline backup.

* **-newsid** *new_sid*

  Specifies a new, unique Oracle system identifier for the cloned database. The system identifier value is a maximum of eight characters. Oracle does not allow running two databases with the same system identifier on the same host simultaneously.

* **-host** *target_host*

  Specifies the host on which the clone should be created.

* **-label** *clone_label*

  Specifies a label for the clone.

* **-comment** *string*

  Specifies an optional comment to describe this clone. You must enclose the string within single quotation marks.

  > **ⓘ** Some shells delete the quotation marks. If that is true for your shell, you must escape the quotation with a backslash (\). For example, you might need to enter: **\' this is a comment\'**

* **-clonespec** *full_path_to_clonespec_file*

  Specifies the path to the clone specification XML file. This can be a relative or an absolute path name.

* **-syspassword** *syspassword*

  Specifies the password for the sys privileged user.

  > **ⓘ** You must provide the password for the sys privileged user if the database credentials that are provided are not the same for the sys privileged user.

* **-reserve**

  Setting the `-reserve` option to yes ensures that the volume guarantee space reserve is turned on for the new clone volumes. Setting the `-reserve` option to no ensures that the volume guarantee space reserve is turned off for the new clone volumes. Setting the `-reserve` option to inherit ensures that the new clone inherits the space reservation characteristics of the parent Snapshot copy. The default setting is no.

  The following table describes the cloning methods and their effect on the clone create operation and its `-reserve` option. A LUN can be cloned by using either method.

Cloning method	Description	Result
LUN cloning	A new clone LUN is created within the same volume.	When the -reserve option for a LUN is set to yes, space is reserved for the full LUN size within the volume.

Volume cloning	A new FlexClone is created, and the clone LUN exists within the new clone volume. Uses the FlexClone technology.	When the -reserve option for a volume is set to yes, space is reserved for the full volume size within the aggregate.

- **-no-resetlogs**

  Specifies to skip recovering the database, executing the DBNEWID utility, and not opening the database with the resetlogs while creating the clone.

- **-recover-from-location**

  Specifies the external archive log location of the archive log backups where SnapManager takes the archive log files from the external location and uses them for cloning.

- **-taskspec**

  Specifies the task specification XML file for preprocessing activity or post-processing activity of the clone operation. You must provide the complete path of the task specification XML file.

- **-dump**

  Specifies to collect the dump files after the clone create operation.

- **-quiet**

  Displays only error messages in the console. The default setting is to display error and warning messages.

- **-verbose**

  Displays error, warning, and informational messages in the console.

Example

The following example clones the backup by using a clone specification that is created for this clone:

```
smsap clone create -profile SALES1 -backup-label full_backup_sales_May
-newsid
CLONE -label sales1_clone -clonespec E:\\spec\\clonespec.xml
```

```
Operation Id [8abc01ec0e794e3f010e794e6e9b0001] succeeded.
```

The smsap clone delete command

You can run the `clone delete` command to delete a clone. You cannot delete a clone if the clone is use by any operation.

Syntax

```
smsap clone delete
-profile profile_name
[-id guid | -label clone_name]
[-login
[-username db_username -password db_password -port db_port]
]
[-syspassword syspassword]
-force
[-dump]
[-quiet | -verbose]
```

Parameters

* **-profile** *profile_name*

  Specifies the name of the profile containing the clone being deleted. The profile contains the identifier of the database and other database information.

* **-force**

  Deletes the clone even if there are resources associated with the clone.

* **-id** *guid*

  Specifies the GUID for the clone being deleted. The GUID is generated by SnapManager when you create a clone. You can use the `smsap clone list` command to display the GUID for each clone.

* **-label** *name*

  Specifies the label for the clone being deleted.

* **-syspassword** *syspassword*

  Specifies the password for the sys privileged user.

  > ⓘ You must provide the password for the sys privileged user if the database credentials provided are not the same for sys privileged user.

* **-login**

  Allows you to enter the database login details.

* **-username** *db_username*

  Specifies the user name required to access the database.

* **-password** *db_password*

Specifies the password required to access the database.

- **-port** *db_port*

  Specifies the TCP port number used to access the database that the profile describes.

- **-dump**

  Specifies to collect the dump files after the clone delete operation.

- **-quiet**

  Displays only error messages in the console. The default is to display error and warning messages.

- **-verbose**

  Displays error, warning, and informational messages in the console.

Example

The following example deletes the clone:

```
smsap clone delete -profile SALES1 -label SALES_May
Operation Id [8abc01ec0e79004b010e79006da60001] succeeded.
```

The smsap clone list command

This command lists the clones of the database for a given profile.

Syntax

```
smsap clone list
-profile profile_name
-delimiter character
[-quiet | -verbose]
```

Parameters

- **-profile** *profile_name*

  Specifies the list of clones associated with the profile. The profile contains the identifier of the database and other database information.

- **-delimiter** *character*

  When this parameter is specified, the command lists the attributes in each row separated by the character specified.

- **-quiet**

Displays only error messages on the console. The default is to display error and warning messages.

* **-verbose**

  Displays error, warning, and informational messages on the console.

Example command

The following example lists the database clones in the SALES1 profile.

```
smsap clone list -profile SALES1 -verbose
```

```
ID Status SID Host Label Comment
-------- ------- ----- ------- ---------- -------
8ab...01 SUCCESS hsdbc server1 back1clone test comment
```

The smsap clone show command

You can run the `clone show` command to display information about the database clones for the specified profile.

Syntax

```
smsap clone show
-profile profile_name
[-id  guid | -label clone_name]
[-quiet | -verbose]
```

Parameters

* **-profile** *profile_name*

  Specifies the list of clones associated with the profile. The profile contains the identifier of the database and other database information.

* **-id** *guid*

  Shows information about the clone with the specified GUID. The GUID is generated by SnapManager when you create a clone. You can use the `smsap clone show` command to display the GUID for each clone.

* **-label** *label_name*

  Shows information about the clone with the specified label.

* **-quiet**

Displays only error messages in the console. The default is to display error and warning messages.

* **-verbose**

    Displays error, warning, and informational messages in the console.

Example

The following example displays information about the clone:

```
smsap clone show -profile SALES1 -label full_backup_sales_May -verbose
```

The following output shows information about a clone of a backup on primary storage:

```
Clone id: 8abc013111b916e30111b916ffb40001
Clone status: SUCCESS
Clone SID: hsdbc
Clone label: hsdbc
Clone comment: null
Clone start time: 2007-04-03 16:15:50
Clone end time: 2007-04-03 16:18:17
Clone host: Host1
Filesystem: E:\ssys1\data_clone\
File: E:\ssys1\data_clone\hsdb\sysaux01.dbf
File: E:\ssys1\data_clone\hsdb\undotbs01.dbf
File: E:\ssys1\data_clone\hsdb\users01.dbf
File: E:\ssys1\data_clone\hsdb\system01.dbf
File: E:\ssys1\data_clone\hsdb\undotbs02.dbf
Backup id: 8abc013111a450480111a45066210001
Backup label: full_backup_sales_May
Backup SID: hsdb1
Backup comment:
Backup start time: 2007-03-30 15:26:30
Backup end time: 2007-03-30 15:34:13
Backup host: server1
```

The smsap clone template command

This command lets you create a clone specification template.

Syntax

```
smsap clone template
-profile name
[-backup-id guid  | -backup-label backup_name]
[-quiet | -verbose]
```

Parameters

- **-profile** *name*

  Specifies the database you want to create a clone specification of. The profile contains the identifier of the database and other database information.

- **-backup-id** *guid*

  Creates a clone specification from the backup with the specified GUID. The GUID is generated by SnapManager when you create a backup. Use the `smsap backup list` command to display the GUID for each backup.

- **-backup-label** *backup_label_name*

  Creates a clone specification from the backup with the specified backup label.

- **-quiet**

  Displays only error messages on the console. The default is to display error and warning messages.

- **-verbose**

  Displays error, warning, and informational messages on the console.

Example command

The following example creates a clone specification template from the backup with the label full_backup_sales_May. Once the smsap clone template command completes, the clone specification template is complete.

```
smsap clone template -profile SALES1 -backup-label full_backup_sales_May
Operation Id [8abc01ec0e79004b010e79006da60001] succeeded.
```

The smsap clone update command

This command updates information about the clone. You can update the comment.

Syntax

```
smsap clone update
-profile profile_name
[-label label | -id id]
-comment  commment_text
[-quiet | -verbose]
```

Parameters

* **-profile** *profile_name*

  Specifies the name of the profile containing the clone you want to update. The profile contains the identifier of the database and other database information.

* **-id** *id*

  Specifies the ID for the clone. The ID is generated by SnapManager when you create a clone. Use the `smsap clone list` command to display the ID for each clone.

* **-label** *label*

  Specifies the label for the clone.

* **-comment**

  Shows the comment entered in the clone creation. This is an optional parameter.

* **-quiet**

  Displays only error messages on the console. The default is to display error and warning messages.

* **-verbose**

  Displays error, warning, and informational messages on the console.

Example command

The following example updates the clone comment.

```
smsap clone update -profile anson.pcrac5
-label clone_pcrac51_20080820141624EDT  -comment See updated clone
```

The smsap clone detach command

After splitting a cloned volume from its parent volume in Data ONTAP, you can run the `clone detach` command from SnapManager to let SnapManager know that the volume is no longer a clone.

Syntax

```
smsap clone detach -profile profile_name -label clone_label
```

Parameters

* **-profile *profile_name***

  Specifies the profile name from which the clone is created.

* **-label *clone_label***

  Specifies the name generated by the clone operation.

Example

The following command detaches the clone:

```
smsap clone detach -profile SALES1 -label sales1_clone
```

The smsap cmdfile command

You can use the `cmdfile` command to run any command if the shell on your host limits the number of characters that can appear on a command line.

Syntax

```
smsap cmdfile
-file file_name
[-quiet | -verbose]
```

You can include the command in a text file and use the `smsap cmdfile` command to execute the command. You can add only one command in a text file. You must not include smsap in the command syntax.

> (i) The `smsap cmdfile` command replaces the `smsap pfile` command. The `smsap cmdfile` is not compatible with the `smsap pfile` command.

Parameters

* **-file *file_name***

  Specifies the path to text file containing the command you want to execute.

* **-quiet**

  Specifies that only error messages are displayed in the console. The default is to display error and warning messages.

* **-verbose**

  Specifies that error, warning, and informational messages are displayed in the console.

The smsap credential clear command

This command clears the cache of the user credentials for all secured resources.

Syntax

```
smsap credential clear
[-quiet | -verbose]
```

Parameters

* **-quiet**

  Displays only error messages on the console. The default is to display error and warning messages.

* **-verbose**

  Displays error, warning, and informational messages on the console.

Example command

This example clears all of the credentials for the user running the command.

```
smsap credential clear -verbose
```

```
SMSAP-20024 [INFO ]: Cleared credentials for user "user1".
```

The smsap credential delete command

This command deletes the user credentials for a particular secured resource.

Syntax

```
smsap credential delete
[-host -name host_name
-username username] |
[-repository
-dbname repo_service_name
-host repo_host
-login -username repo_username
-port repo_port] |
[-profile
-name profile_name]
[-quiet | -verbose]
```

Parameters

* **-host** *hostname*

  Specifies the name of the host server on which SnapManager is running.

  The `-host` parameter includes the following options:

  ◦ `-name host_name`: Specifies the name of the host for which you will delete the password.

  ◦ `-username user_name`: Specifies the user name on the host.

* **-repository -dbname**

  Specifies the name of the database that stores the profile. Use either the global name or the SID.

  The `-repository` parameter includes the following options:

  ◦ `-dbnamerepo_service_name`: Specifies the name of the database that stores the profile. Use either the global name or the SID.

  ◦ `-host repo_host`: Specifies the name or IP address of the host server the repository database runs on.

  ◦ `-login-username repo_username`: Specifies the user name needed to access the database that stores the repository.

  ◦ `-port repo_port`: Specifies the TCP port number used to access the database that stores the repository.

* **-profile -name** *profile_name*

  Specifies the profile with which the database is associated.

  The `-profile` parameter includes the following option:

  ◦ `-name` *profilename*: Specifies the name of the profile for which you will delete the password.

* **-quiet**

  Displays only error messages on the console. The default is to display error and warning messages.

- **-verbose**

  Displays error, warning, and informational messages on the console.

Example command

This example deletes the credentials of the profile.

```
smsap credential delete -profile -name user1 -verbose
```

```
SMSAP-20022 [INFO ]: Deleted credentials and repository mapping
 for profile "user1" in user credentials for "user1".
```

This example deletes the credentials of the repository.

```
smsap credential delete -repository -dbname SMSAPREPO -host Host2
-login -username user1 -port 1521
```

```
SMSAP-20023 [INFO ]: Deleted repository credentials for
"user1@SMSAPREPO/wasp:1521"
and associated profile mappings in user credentials for "user1".
```

This example deletes the credentials of the host.

```
smsap credential delete -host -name Host2
```

```
SMSAP-20033 [INFO ]: Deleted host credentials for "Host2" in user
credentials for "user1".
```

The smsap credential list command

This command lists all credentials of a user.

Syntax

```
smsap credential list
[-quiet | -verbose]
```

Parameters

- **-quiet**

  Displays only error messages on the console. The default is to display error and warning messages.

- **-verbose**

  Displays error, warning, and informational messages on the console.

Example command

This example displays all of the credentials for the user running the command.

```
smsap credential list
```

```
Credential cache for OS user "user1":
Repositories:
Host1_test_user@SMSAPREPO/hotspur:1521
Host2_test_user@SMSAPREPO/hotspur:1521
user1_1@SMSAPREPO/hotspur:1521
Profiles:
HSDBR (Repository: user1_2_1@SMSAPREPO/hotspur:1521)
PBCASM (Repository: user1_2_1@SMSAPREPO/hotspur:1521)
HSDB (Repository: Host1_test_user@SMSAPREPO/hotspur:1521) [PASSWORD NOT
SET]
Hosts:
Host2
Host5
Host4
Host1
```

The smsap credential set command

This command lets you set the credentials for users to access secure resources, such as hosts, repositories, and database profiles. The host password is the user's password on the host on which SnapManager is running. The repository password is the password of the Oracle user that contains the SnapManager repository schema. The profile password is a password that is made up by the person who creates the profile. For the host and repository options, if the optional `-password` option is not included, you will be prompted to enter a password of the type specified in the command arguments.

Syntax

```
smsap credential set
[-host
-name host_name
-username username]
[-password password]] |
[-repository
-dbname repo_service_name
-host repo_host
-login -username repo_username] [-password repo_password]]
-port repo_port |
[-profile
-name profile_name]
[-password password]]
[-quiet | -verbose]
```

Parameters

- **-host** *hostname*

  Specifies the name or IP address of the host server on which SnapManager is running.

  The -host parameter includes the following options:

  - -name host_name: Specifies the name of the host for which you will set the password.

  - -username user_name: Specifies the user name on the host.

  - -password password: Specifies the password of the user on the host.

- **-repository -dbname**

  Specifies the name of the database that stores the profile. Use either the global name or the SID.

  The -repository parameter includes the following options:

  - -dbnamerepo_service_name: Specifies the name of the database that stores the profile. Use either the global name or the SID.

  - -host repo_host: Specifies the name or IP address of the host server the repository database runs on.

  - -login-username repo_username: Specifies the user name needed to access the database that stores the repository.

  - -password password: Specifies the password needed to access the database that stores the repository.

  - -port repo_port: Specifies the TCP port number used to access the database that stores the repository.

- **-profile -name** *profile_name*

  Specifies the profile with which the database is associated.

The `-profile` parameter includes the following option:

- `-name profilename`: Specifies the name of the profile for which you will set the password.
- `-password password`: Specifies the password needed to access the profile.

- **`-quiet`**

  Displays only error messages on the console. The default is to display error and warning messages.

- **`-verbose`**

  Displays error, warning, and informational messages on the console.

Example command for setting repository credentials

The following example sets credentials for a repository.

```
smsap credential set -repository -dbname SMSAPREPO -host hotspur -port
1527 -login -username chris
Password for chris@hotspur:1527/SMSAPREPO : ********
Confirm password for chris@hotspur:1527/SMSAPREPO : ********
```

```
SMSAP-12345 [INFO ]: Updating credential cache for OS user "admin1"
SMSAP-12345 [INFO ]: Set repository credential for user "user1" on
repo1@Host2.
Operation Id [Nff8080810da9018f010da901a0170001] succeeded.
```

Example command for setting host credentials

Because a host credential represents an actual operating system credential, it must include the username in addition to the password.

```
smsap credential set -host -name bismarck -username avida
Password for avida@bismarck : ********
Confirm password for avida@bismarck : ********
```

The smsap history list command

This command enables you to view a list of history details of the SnapManager operation.

Syntax

```
smsap history list
-profile {-name profile_name [profile_name1, profile_name2] | -all
-repository
-login [-password repo_password]
-username repo_username
-host repo_host
-dbname repo_dbname
-port repo_port}
-operation {-operations operation_name [operation_name1, operation_name2]
| -all}
[-delimiter character]
[-quiet | -verbose]
```

Parameters

- **-profile** *profile*

  Specifies the name of the profile. This name can be up to 30 characters long and must be unique within the host.

- **-repository**

  The options that follow -repository specify the details of the database that stores the profile.

- **-dbname** *repo_dbname*

  Specifies the name of the database that stores the profile. Use either the global name or the SID.

- **-host** *repo_host*

  Specifies the name or IP address of the host computer the repository database runs on.

- **-login**

  Starts the repository login details.

- **-username** *repo_username*

  Specifies the user name needed to access the database that stores the repository.

- **-port** *repo_port*

  Specifies the TCP port number used to access the database that stores the repository.

- **-operation {-operationsoperation_name [operation_name1, operation_name2] |
  -all**

  Specifies the SnapManager operation for which you configure the history.

- **-quiet**

Displays only error messages on the console. The default is to display error and warning messages.

- **-verbose**

  Displays error, warning, and informational messages on the console.

Example command

```
smsap history list -profile -name PROFILE1 -operation -operations backup
-verbose
```

The smsap history operation-show command

This command enables you to view the history of a specific SnapManager operation associated with a profile.

Syntax

```
smsap history operation-show
-profile profile
{-label label | -id id}
[-quiet | -verbose]
```

Parameters

- **-profile** *profile*

  Specifies the name of the profile. This name can be up to 30 characters long and must be unique within the host.

- **-label** *label* | *-idid*

  Specifies the SnapManager operation ID or label for which you want to view the history.

- **-quiet**

  Displays only error messages on the console. The default is to display error and warning messages.

- **-verbose**

  Displays error, warning, and informational messages on the console.

Example command

```
smsap history operation-show -profile PROFILE1 -label backup1 -verbose
```

The smsap history purge command

This command enables you to delete the history of SnapManager operation.

Syntax

```
smsap history purge
-profile {-name profile_name [profile_name1, profile_name2] | -all
-repository
-login [-password repo_password]
-username repo_username
-host repo_host
-dbname repo_dbname
-port repo_port}
-operation {-operations operation_name [operation_name1, operation_name2]
| -all}
[-quiet | -verbose]
```

Parameters

* **-profile** *profile*

  Specifies the name of the profile. This name can be up to 30 characters long and must be unique within the host.

* **-repository**

  The options that follow -repository specify the details of the database that stores the profile.

* **-dbname** *repo_dbname*

  Specifies the name of the database that stores the profile. Use either the global name or the SID.

* **-host** *repo_host*

  Specifies the name or IP address of the host computer the repository database runs on.

* **-login**

  Starts the repository login details.

* **-username** *repo_username*

  Specifies the user name needed to access the database that stores the repository.

* **-port** *repo_port*

  Specifies the TCP port number used to access the database that stores the repository.

* **-operation {-operationsoperation_name** [*operation_name1, operation_name2*] |

-all

Specifies the SnapManager operation for which you configure the history.

- **-quiet**

  Displays only error messages on the console. The default is to display error and warning messages.

- **-verbose**

  Displays error, warning, and informational messages on the console.

Example command

```
smsap history purge -profile -name PROFILE1 -operation -operations backup
-verbose
```

The smsap history remove command

This command enables you to remove the history of SnapManager operations associated with a single profile, multiple profiles, or all profiles under a repository.

Syntax

```
smsap history remove
-profile {-name profile_name [profile_name1, profile_name2] | -all
-repository
-login [-password repo_password]
-username repo_username
-host repo_host
-dbname repo_dbname
-port repo_port}
-operation {-operations operation_name [operation_name, operation_name2] |
-all}
[-quiet | -verbose]
```

Parameters

- **-profile** *profile*

  Specifies the name of the profile. This name can be up to 30 characters long and must be unique within the host.

- **-repository**

  The options that follow -repository specify the details of the database that stores the profile.

- **-dbname** *repo_dbname*

  Specifies the name of the database that stores the profile. Use either the global name or the SID.

- **-host** *repo_host*

  Specifies the name or IP address of the host computer the repository database runs on.

- **-login**

  Starts the repository login details.

- **-username** *repo_username*

  Specifies the user name needed to access the database that stores the repository.

- **-port** *repo_port*

  Specifies the TCP port number used to access the database that stores the repository.

- **-operation {-operations operation_name [operation_name1, operation_name2] | -all**

  Specifies the SnapManager operation for which you configure the history.

- **-quiet**

  Displays only error messages on the console. The default is to display error and warning messages.

- **-verbose**

  Displays error, warning, and informational messages on the console.

Example command

```
smsap history purge -profile -name PROFILE1 -operation -operations backup
-verbose
```

The smsap history set command

You can run the `history set` command to configure the operations for which you want to view the history.

Syntax

```
smsap history set
-profile {-name profile_name [profile_name1, profile_name2] | -all
-repository
-login [password repo_password]
-username repo_username
-host repo_host
-dbname repo_dbname
-port repo_port}
-operation {-operations operation_name [operation_name1, operation_name2]
| -all}
-retain
{-count retain_count | -daily daily_count | -monthly monthly_count |
-weekly weekly_count}
[-quiet | -verbose]
```

Parameters

* **-profile** *profile*

  Specifies the name of the profile. The name can be up to 30 characters long and must be unique within the host.

* **-repository**

  Specifies the details of the database that stores the profile.

* **-dbname** *repo_dbname*

  Specifies the name of the database that stores the profile. You can use either the global name or the system identifier.

* **-host** *repo_host*

  Specifies the name or IP address of the host where the repository database resides.

* **-login**

  Specifies the repository login details.

* **-username** *repo_username*

  Specifies the user name required to access the repository database.

* **-port** *repo_port*

  Specifies the Transmission Control Protocol (TCP) port number used to access the repository database.

* **-operation {-operations** *operation_name* **[operation_name1, operation_name2] |
  -all**

Specifies the SnapManager operations for which you want to configure the history.

- **`-retain {-countretain_count | -dailydaily_count | -monthly-monthly_count | -weeklyweekly_count}`**

  Specifies the retention class of the create backup, verify backup, restore and recover, and create clone operations. The retention class is set based on the operation count number, number of days, weeks, or months.

- **`-quiet`**

  Displays only error messages in the console. The default is to display error and warning messages.

- **`-verbose`**

  Displays error, warning, and informational messages in the console.

Example command

The following example displays information about the backup operation:

```
smsap history set -profile -name PROFILE1 -operation -operations backup
-retain -daily 6
-verbose
```

The smsap history show command

This command enables you to view a detailed history information for a specific profile.

Syntax

```
smsap history show
-profile profile
```

Parameters

- **`-profile profile`**

  Specifies the name of the profile. This name can be up to 30 characters long and must be unique within the host.

- **`-quiet`**

  Displays only error messages on the console. The default is to display error and warning messages.

- **`-verbose`**

  Displays error, warning, and informational messages on the console.

Example command

```
smsap history show -profile -name PROFILE1
-verbose
```

The smsap help command

You can run the `help` command to display information about the SnapManager commands and their options. If you do not supply a command name, it displays a list of valid commands. If you supply a command name, it displays the syntax for that command.

Syntax

```
smsap help
[][backup|cmdfile|clone|credential|help|operation|profile|repository|syste
m|version|plugin|diag|history|schedule|notification|storage|get]
[-quiet | -verbose]
```

Parameters

The following are some command names you can use with this command:

* backup
* clone
* cmdfile
* credential
* diag
* get
* notification
* help
* history
* operation
* plugin
* profile
* repository
* schedule
* storage
* system

* `version`

The smsap notification remove-summary-notification command

This command disables summary notification for multiple profiles on a repository database.

Syntax

```
smsap notification remove-summary-notification
-repository
-dbname repo_service_name
-port repo_port
-host repo_host
-login -username repo_username
[-quiet | -verbose]
```

Parameters

* **-repository**

  The options that follow `-repository` specify the details of the database for the repository.

* **-port** *repo_port*

  Specifies the TCP port number used to access the database that stores the repository.

* **-dbname** *repo_service_name*

  Specifies the name of the database that stores the repository. Use either the global name or the SID.

* **-host** *repo_host*

  Specifies the name or IP address of the host computer the repository database runs on.

* **-login** *repo_username*

  Specifies the login name needed to access the database that stores the repository.

* **-quiet**

  Displays only error messages on the console. The default is to display error and warning messages.

* **-verbose**

  Displays error, warning, and informational messages on the console.

The following example disables summary notification for multiple profiles on a repository database.

```
smsap notification remove-summary-notification -repository -port 1521
-dbname repo2 -host 10.72.197.133 -login -username oba5
```

The smsap notification update-summary-notification command

You can run the `notification update-summary-notification` command to
enable summary notification for a repository database.

Syntax

```
smsap notification update-summary-notification
-repository
-port repo_port
-dbname  repo_service_name
-host repo_host
-login -username repo_username
-email email-address1,email-address2
-subject subject-pattern
-frequency
[-daily -time daily_time |
-hourly -time hourly_time |
-monthly -time monthly_time -date [1|2|3|...|31] |
-weekly  -time weekly_time -day [1|2|3|4|5|6|7]]
-profiles profile1,profile2
-notification-host notification-host
[-quiet | -verbose]
```

Parameters

* **-repository**

  Specifies the details of the repository database.

* **-port** *repo_port*

  Specifies the TCP port number used to access the repository database.

* **-dbname** *repo_service_name*

  Specifies the name of the repository database. You can use either the global name or the system identifier.

* **-host** *repo_host*

  Specifies the name or IP address of the host on which the repository database resides.

* **-login**

Specifies the repository login details. This is optional. If not specified, SnapManager defaults to OS Authentication Connection Mode.

- **`-username`** *`repo_username`*

  Specifies the user name required to access the repository database.

- **`-email`** *`email-address1,e-mail-address2`*

  Specifies email addresses of the recipients.

- **`-subject`** *`subject-pattern`*

  Specifies the email subject pattern.

- **`-frequency { -daily --time`** *`daily_time`* **`| -hourly --time`** *`hourly_time`* **`| -monthly --time`** *`monthly_time`* **`-date {1|2|3…|31 } | -weekly --time`** *`weekly_time`* **`-day {1|2|3|4|5|6|7 } }`**

  Specifies schedule type and schedule time when you want the email notification.

- **`-profiles`** *`profile1, profile2`*

  Specifies profile names that require email notification.

- **`-notification-host`** *`notification-host`*

  Specifies SnapManager server host from which the summary notification email is sent to the recipients. You can provide host name, or IP address for the notification host. You can also update the host IP or host name.

- **`-quiet`**

  Displays only error messages in the console. The default is to display error and warning messages.

- **`-verbose`**

  Displays error, warning, and informational messages in the console.

Example

The following example enables summary notification for a repository database:

```
smsap notification update-summary-notification –repository –port 1521
–dbname repo2 –host 10.72.197.133 –login –username oba5 –email
admin@org.com –subject success –frequency -daily -time 19:30:45 –profiles
sales1
```

The smsap notification set command

You can use the `notification set` command to configure the mail server.

Syntax

```
smsap notification set
-sender-email email_address
-mailhost mailhost
-mailport mailport
[-authentication
-username username
-password password]
-repository
-dbname repo_service_name
-port repo_port]
-host repo_host
-login -username repo_username
[-quiet | -verbose]
```

Parameters

* **-sender-email** *email_address*

  Specifies the sender's email address from which the email alerts are sent. From SnapManager 3.2 for SAP, you can include a hyphen (-) while specifying the domain name of the email address. For example, you can specify the sender email address as `-sender-email07lbfmdatacenter@continental -corporation.com`.

* **-mailhost** *mailhost*

  Specifies the name or IP address of the host server that handles email notifications.

* **-mailport** *mailport*

  Specifies the mail server port number.

* **-authentication -username** *username* *-password password*

  Specifies authentication details for the email address. You must specify the user name and password.

* **-repository**

  Specifies the details of the repository database.

* **-port** *repo_port*

  Specifies the Transmission Control Protocol (TCP) port number used to access the repository database.

* **-dbname** *repo_service_name*

  Specifies the name of the repository database. You can use either the global name or the system identifier.

* **-host** *repo_host*

Specifies the name or IP address of the host where the repository database resides.

- **-login**

  Specifies the repository login details. This is optional. If not specified, SnapManager defaults to OS Authentication Connection Mode.

- **-username** *repo_username*

  Specifies the user name required to access the repository database.

- **-quiet**

  Displays only error messages in the console. The default is to display error and warning messages.

- **-verbose**

  Displays error, warning, and informational messages in the console.

Example

The following example configures the mail server:

```
smsap notification set -sender-email admin@org.com -mailhost
hostname.org.com -mailport 25 authentication -username davis -password
davis -repository -port 1521 -dbname SMSAPREPO -host hotspur
-login -username grabal21 -verbose
```

The smsap operation dump command

You can run the `operation dump` command to create a JAR file that contains diagnostic information about an operation.

Syntax

```
smsap operation dump
-profile profile_name
[-label label_name | -id guid]
[-quiet | -verbose]
```

Parameters

- **-profile** *profile_name*

  Specifies the profile for which you want to create the dump files. The profile contains the identifier of the database and other database information.

- **-label** *label_name*

Creates dump files for the operation and assigns the specified label.

- **-id** *guid*

  Creates dump files for the operation with the specified GUID. The GUID is generated by SnapManager when the operation begins.

- **-quiet**

  Displays only error messages in the console. The default is to display error and warning messages.

- **-verbose**

  Displays error, warning, and informational messages in the console.

Example

The following example creates the dump file for the backup:

```
smsap operation dump -profile SALES1
-id 8abc01ec0e78f3e2010e78f3fdd00001
```

```
Dump file created Path:
C:\userhomedirectory\netapp\smsap\3.3\smsap_dump_8abc01ec0e78f3e2010e78f3f
dd00001.jar
```

The smsap operation list command

This command lists the summary information of all operations recorded against a specified profile.

Syntax

```
smsap operation list
-profile profile_name
[-delimiter character]
[-quiet | -verbose]
```

Parameters

- **-profile** *profile_name*

  Specifies the name of the profile. This name can be up to 30 characters long and must be unique within the host.

- **-delimiter** *character*

(Optional) When this parameter is specified, the command lists each row on a separate line and the attributes in that row are separated by the character specified.

- **-quiet**

  (Optional) Displays only error messages on the console. The default is to display error and warning messages.

- **-verbose**

  (Optional) Displays error, warning, and informational messages on the console.

Example command

The following example lists the summary information of all the operations logged against the specified profile.

```
smsap operation list -profile myprofile
```

```
Start Date Status Operation ID Type Host
-------------------- ---------- --------------------------------
----------------- --
2007-07-16 16:03:57 SUCCESS 8abc01c813d0a1530113d0a15c5f0005 Profile
Create Host3
2007-07-16 16:04:55 FAILED 8abc01c813d0a2370113d0a241230001 Backup Host3
2007-07-16 16:50:56 SUCCESS 8abc01c813d0cc580113d0cc60ad0001 Profile
Update Host3
2007-07-30 15:44:30 SUCCESS 8abc01c81418a88e011418a8973e0001 Remove Backup
Host3
2007-08-10 14:31:27 SUCCESS 8abc01c814510ba20114510bac320001 Backup Host3
2007-08-10 14:34:43 SUCCESS 8abc01c814510e9f0114510ea98f0001 Mount Host3
2007-08-10 14:51:59 SUCCESS 8abc01c814511e6e0114511e78d40001 Unmount Host3
```

The smsap operation show command

You can run the `operation show` command to list the summary information of all the operations performed against the specified profile. The output lists the client user (the user for the client PC) and the effective user (the user in SnapManager who is valid on the selected host).

Syntax

```
smsap operation show
-profile profile_name
[-label label | -id id]
[-quiet | -verbose]
```

Parameters

* **-profile** *profile_name*

  Specifies the name of the profile. This name can be up to 30 characters long and must be unique within the host.

* **-label** *label*

  Specifies the label for the operation.

* **-id** *id*

  Specifies the identifier for the operation.

* **-quiet**

  Optional: Displays only error messages in the console. The default is to display error and warning messages.

* **-verbose**

  Optional: Displays error, warning, and informational messages in the console.

Example

The following command line shows detailed information about an operation:

```
smsap operation show -id 8ac861781d0ac992011d0ac999680001  -profile CER
```

The smsap password reset command

You can run the password `reset command` to reset the password of a profile.

Syntax

```
smsap password reset
-profile profile
[-profile-password profile_password]
[-repository-hostadmin-password repository_hostadmin_password]
[-quiet | -verbose]
```

Parameters

* **-profile** *profile*

  Specifies the name of the profile for which you want to reset the password.

* **-profile-password** *profile_password*

  Specifies the new password for the profile.

* **-repository-hostadmin-password** *admin_password*

  Specifies the authorized user credential with local administrator privilege for the repository database.

* **-quiet**

  Displays only error messages in the console. The default is to display error and warning messages.

* **-verbose**

  Displays error, warning, and informational messages in the console.

The smsap profile create command

You can run the `profile create` command to create a profile of a database in a repository. You must mount the database before you run this command.

Syntax

```
smsap profile create
-profile profile
[-profile-password profile_password]
-repository
-dbname repo_service_name
-host repo_host
-port repo_port
-login -username repo_username
-database
-dbname db_dbname
-host db_host
[-sid  db_sid]
[-login
[-username db_username -password db_password -port db_port]
]
[-rman {-controlfile | {-login
-username  rman_username -password  rman_password}
-tnsname  rman_tnsname}}]
```

```
[-retain
[-hourly [-count n] [-duration m]]
[-daily [-count n] [-duration m]]
[-weekly [-count n] [-duration m]]
[-monthly [-count n] [-duration m]]]]
-comment comment
-snapname-pattern pattern
[]
[-summary-notification]
[-notification
[-success
-email email_address1,email_address2
-subject subject_pattern]
[-failure
-email email_address1,email_address2
-subject subject_pattern]
[-separate-archivelog-backups
-retain-archivelog-backups
-hours hours |
-days days |
-weeks weeks |
-months months
[]
[-include-with-online-backups | -no-include-with-online-backups]]
[-dump]
[-quiet | -verbose]
```

Parameters

* **-profile** *profile*

  Specifies the name of the profile. This name can be up to 30 characters long and must be unique within the host.

* **-profile-password** *profile_password*

  Specify the password for the profile.

* **-repository**

  The options that follow `-repository` specify the details of the database that stores the profile.

* **-dbname** *repo_service_name*

  Specifies the name of the database that stores the profile. Use either the global name or the SID.

* **-host** *repo_host*

  Specifies the name or IP address of the host computer the repository database runs on.

- **-sid** *db_sid*

  Specifies the system identifier of the database that the profile describes. By default, SnapManager uses the database name as the system identifier. If the system identifier is different from the database name, you must specify it with the -sid option.

- **-login**

  Specifies the repository login details.

- **-username** *repo_username*

  Specifies the user name needed to access the repository database.

- **-port** *repo_port*

  Specifies the TCP port number used to access the repository database.

- **-database**

  Specifies the details of the database that the profile describes. This is the database that will be backed up, restored, or cloned.

- **-dbname** *db_dbname*

  Specifies the name of the database that the profile describes. You can use either the global name or the system identifier.

- **-host** *db_host db_host*

  Specifies the name or IP address of the host computer on which the database runs.

- **-login**

  Specifies the database login details.

- **-username** *db_username*

  Specifies the user name needed to access the database that the profile describes.

- **-password** *db_password*

  Specifies the password needed to access the database that the profile describes.

- **-port** *db_port*

  Specifies the TCP port number used to access the database that the profile describes.

- **-retain [-hourly [-count n] [-duration m]] [-daily [-count n] [-duration m]] [-weekly [-count n] [-duration m]] [-monthly [-count n] [-duration m]]**

  Specifies the retention policy for a backup where either or both of a retention count along with a retention duration for a retention class (hourly, daily, weekly, monthly).

For each retention class, either or both of a retention count or a retention duration may be specified. The duration is in units of the class (for example, hours for hourly, days for daily). For instance, if the user specifies only a retention duration of 7 for daily backups, then SnapManager will not limit the number of daily backups for the profile (because the retention count is 0), but SnapManager will automatically delete daily backups created over 7 days ago.

- **`-comment` *`comment`***

  Specifies the comment for a profile describing the profile domain.

- **`-snapname-pattern` *`pattern`***

  Specifies the naming pattern for Snapshot copies. You can also include custom text, for example, HAOPS for highly available operations, in all Snapshot copy names. You can change the Snapshot copy naming pattern when you create a profile or after the profile has been created. The updated pattern applies only to Snapshot copies that have not yet been created. Snapshot copies that exist retain the previous Snapname pattern. You can use several variables in the pattern text.

- **`-summary-notification`**

  Specifies that summary email notification is enabled for the new profile.

- **`-notification -success-email` *`e-mail_address1,e-mail address2`* `-subject` *`subject_pattern`***

  Specifies that email notification is enabled for the new profile so that emails are received by recipients when the SnapManager operation succeeds. You must enter a single email address or multiple email addresses to which email alerts will be sent and an email subject pattern for the new profile.

  You can also include custom subject text for the new profile. You can change the subject text when you create a profile or after the profile has been created. The updated subject applies only to the emails that are not sent. You can use several variables for the email subject.

- **`-notification -failure -email` *`e-mail_address1,e-mail address2`* `-subject` *`subject_pattern`***

  Specifies that enable email notification is enabled for the new profile so that emails are received by recipients when the SnapManager operation fails. You must enter a single email address or multiple email addresses to which email alerts will be sent and an email subject pattern for the new profile.

  You can also include custom subject text for the new profile. You can change the subject text when you create a profile or after the profile has been created. The updated subject applies only to the emails that are not sent. You can use several variables for the email subject.

- **`-separate-archivelog-backups`**

  Specifies that the archive log backup is separated from datafile backup. This is an optional parameter you can provide while creating the profile. After you separate the backup using this option, you can either take data files-only backup or archive logs-only backup.

- **`-retain-archivelog-backups -hours` *`hours`* `| -days` *`days`* `| -weeks` *`weeks`* `| -months` *`months`***

  Specifies that the archive log backups are retained based on the archive log retention duration (hourly,

daily, weekly, monthly).

- **-quiet**

  Displays only error messages in the console. The default is to display error and warning messages.

- **-verbose**

  Displays error, warning, and informational messages in the console.

- **-include-with-online-backups**

  Specifies that the archive log backup is included along with the online database backup.

- **-no-include-with-online-backups**

  Specifies that the archive log backups are not included along with the online database backup.

- **-dump**

  Specifies that the dump files are collected after the successful profile create operation.

Example

The following example shows the creation of a profile with hourly retention policy and email notification:

```
smsap profile create -profile test_rbac -profile-password netapp
-repository -dbname SMSAPREP -host hostname.org.com -port 1521 -login
-username smsaprep -database -dbname RACB -host saal -sid racb1 -login
-username sys -password netapp -port 1521 -rman -controlfile -retain
-hourly -count 30 -verbose
Operation Id [8abc01ec0e78ebda010e78ebe6a40005] succeeded.
```

The smsap profile delete command

You can run the `profile delete` command to delete a profile of the database.

Syntax

```
smsap profile delete
-profile profile
[-quiet | -verbose]
```

Parameters

- **-profile** *profile*

  Specifies the profile to be deleted.

- **-quiet**

    Displays only error messages in the console. The default is to display error and warning messages.

- **-verbose**

    Displays error, warning, and informational messages in the console.

Example

The following example deletes the profile:

```
smsap profile delete -profile SALES1
Operation Id [Ncaf00af0242b3e8dba5c68a57a5ae932] succeeded.
```

The smsap profile dump command

You can run the `profile dump` command to create the `.jar` file that contains diagnostic information about a profile.

Syntax

```
smsap profile dump
-profile profile_name
[-quiet | -verbose]
```

Parameters

- **-profile** *profile_name*

    Specifies the profile for which you want to create the dump files. The profile contains the identifier of the database and other database information.

- **-quiet**

    Displays only error messages in the console. The default is to display error and warning messages.

- **-verbose**

    Displays error, warning, and informational messages in the console.

Example

The following example creates a dump for the profile SALES1:

```
smsap profile dump -profile SALES1
Dump file created
Path:
C:\\userhomedirectory\\netapp\\smsap\\3.3.0\\smsap_dump_SALES1_hostname.ja
r
```

The smsap profile list command

This command displays a list of the current profiles.

Syntax

```
smsap profile list
[-quiet | -verbose]
```

Parameters

* **-quiet**

  Displays only error messages on the console. The default is to display error and warning messages.

* **-verbose**

  Displays error, warning, and informational messages on the console.

Example command

The following example displays existing profiles with their details.

```
smsap profile list -verbose
Profile name: FGTER
Repository:
  Database name: SMSAPREPO
  SID: SMSAPREPO
  Host: hotspur
  Port: 1521
  Username: swagrahn
  Password: ********
Profile name: TEST_RBAC
Repository:
  Database name: smsaprep
  SID: smsaprep
  Host: elbe.rtp.org.com
  Port: 1521
  Username: smsapsaal
```

```
  Password: ********
Profile name: TEST_RBAC_DP_PROTECT
Repository:
  Database name: smsaprep
  SID: smsaprep
  Host: elbe.rtp.org.com
  Port: 1521
  Username: smsapsaal
  Password: ********
Profile name: TEST_HOSTCREDEN_OFF
Repository:
  Database name: smsaprep
  SID: smsaprep
  Host: elbe.rtp.org.com
  Port: 1521
  Username: smsapsaal
  Password: ********
Profile name: SMK_PRF
Repository:
  Database name: smsaprep
  SID: smsaprep
  Host: elbe.rtp.org.com
  Port: 1521
  Username: smsapsaal
  Password: ********
Profile name: FGLEX
Repository:
  Database name: SMSAPREPO
  SID: SMSAPREPO
  Host: hotspur
  Port: 1521
  Username: swagrahn
  Password: ********
```

The smsap profile show command

You can run the `profile show` command to display the information about a profile.

Syntax

```
smsap profile show
-profile profile_name
[-quiet | -verbose]
```

Parameters

- **-profile** *profile_name*

  Specifies the name of the profile. This name can be up to 30 characters long and must be unique within the host.

- **-quiet**

  Displays only error messages in the console. The default is to display error and warning messages.

- **-verbose**

  Displays error, warning, and informational messages in the console.

The smsap profile sync command

This command loads the profile-to-repository mappings for that repository to a file in your home directory on the local host.

Syntax

```
smsap profile sync
-repository
-dbname repo_service_name
-host repo_host
-port repo_port
-login
-username repo_username
[-quiet | -verbose]
```

Parameters

- **-repository**

  The options that follow `-repository` specify the details of the database for the repository.

- **-dbname** *repo_service_name*

  Specifies the repository database for the profile to synchronize.

- **-host**

  Specifies the database host.

- **-port**

  Specifies the port for the host.

- **-login**

Specifies the log in process for the host user.

- **-username**

  Specifies the username for the host.

- **-quiet**

  Displays only error messages on the console. The default is to display error and warning messages.

- **-verbose**

  Displays error, warning, and informational messages on the console.

Example command

The following example shows the result of the command to synchronize the profile-to-repository mappings for the database.

```
smsap profile sync -repository -dbname smrepo -host Host2 -port 1521
-login -username user2
SMSAP-12345 [INFO ]: Loading profile mappings for repository
"user2@Host2:smrepo" into cache for OS User "admin".
Operation Id [Nff8080810da9018f010da901a0170001] succeeded.
```

The smsap profile update command

You can run the `profile update` command to update the information for an existing profile.

Syntax

```
smsap profile update
-profile profile
[-new-profile new_profile_name]
[-profile-password profile_password]
[-database
-dbname db_dbname
-host db_host
[-sid db_sid]
[-login
[-username db_username -password db_password -port db_port]
]
[{-rman  {-controlfile | {{-login
-username  rman_username
-password  rman_password }
[-tnsname  tnsname]}}} |
-remove-rman]


[-retain
[-hourly [-count n] [-duration m]]
[-daily [-count n] [-duration m]]
[-weekly [-count n] [-duration m]]
[-monthly [-count n] [-duration m]]]]
-comment comment
-snapname-patternpattern
[]
[-summary-notification]
[-notification
[-success
-email email_address1,email_address2
-subject subject_pattern]
[-failure
-email email_address1,email_address2
-subject subject_pattern]
[-separate-archivelog-backups
-retain-archivelog-backups
-hours  hours |
-days days |
-weeks weeks |
-months months
[]
[-include-with-online-backups | -no-include-with-online-backups]]
[-dump]
[-quiet | -verbose]
```

Parameters

- **-profile** *profile*

  Specifies the name of the profile. This name can be up to 30 characters long and must be unique within the host.

- **-profile-password** *profile_password*

  Specifies the password for the profile.

- **-new-profile** *new_profile_name*

  Specifies the new name that you can provide for a profile.

- **-database**

  Specifies the details of the database that the profile describes. This is the database that will be backed up, restored, and so on.

- **-dbname** *db_dbname*

  Specifies the name of the database that the profile describes. You can use either the global name or the system identifier.

- **-host** *db_host*

  Specifies the name or IP address of the host computer on which the database runs.

- **-sid** *db_sid*

  Specifies the system identifier of the database that the profile describes. By default, SnapManager uses the database name as the system identifier. If the system identifier is different from the database name, you must specify it using the `-sid` option.

- **-login**

  Specifies the repository login details.

- **-username** *repo_username*

  Specifies the user name required to access the repository database.

- **-port** *repo_port*

  Specifies the TCP port number required to access the repository database.

- **-database**

  Specifies the details of the database that the profile describes. This is the database that will be backed up, restored, or cloned.

- **-dbname** *db_dbname*

Specifies the name of the database that the profile describes. You can use either the global name or the system identifier.

- **-host** *db_host*

  Specifies the name or IP address of the host computer on which the database runs.

- **-login**

  Specifies the database login details.

- **-username** *db_username*

  Specifies the user name required to access the database that the profile describes.

- **-password** *db_password*

  Specifies the password required to access the database that the profile describes.

- **-port** *db_port*

  Specifies the TCP port number required to access the database that the profile describes.

- **-retain [-hourly [-countn] [-duration m]] [-daily [-count n] [-duration m]] [-weekly [-count n][-duration m]] [-monthly [-count n][-duration m]]**

  Specifies the retention class (hourly, daily, weekly, monthly) for a backup.

  For each retention class, a retention count or a retention duration or both can be specified. The duration is in units of the class (for example, hours for hourly or days for daily). For instance, if the user specifies only a retention duration of 7 for daily backups, then SnapManager will not limit the number of daily backups for the profile (because the retention count is 0), but SnapManager will automatically delete daily backups created over 7 days ago.

- **-comment** *comment*

  Specifies the comment for a profile.

- **-snapname-pattern** *pattern*

  Specifies the naming pattern for Snapshot copies. You can also include custom text, for example, HAOPS for highly available operations, in all Snapshot copy names. You can change the Snapshot copy naming pattern when you create a profile or after the profile has been created. The updated pattern applies only to Snapshot copies that have not yet occurred. Snapshot copies that exist retain the previous Snapname pattern. You can use several variables in the pattern text.

- **-summary-notification**

  Specifies that summary email notification is enabled for the existing profile.

- **-notification [-success-email** *e-mail_address1,e-mail address2* **-subject** *subject_pattern*]

  Enables email notification for the existing profile so that emails are received by recipients when the

SnapManager operation succeeds. You must enter a single email address or multiple email addresses to which email alerts will be sent and an email subject pattern for the existing profile.

You can change the subject text while updating the profile or include custom subject text. The updated subject applies only to the emails that are not sent. You can use several variables for the email subject.

* **`-notification [-failure -email e-mail_address1,e-mail address2 -subject subject_pattern]`**

  Enables email notification for the existing profile so that emails are received by recipients when the SnapManager operation fails. You must enter a single email address or multiple email addresses to which email alerts will be sent and an email subject pattern for the existing profile.

  You can change the subject text while updating the profile or include custom subject text. The updated subject applies only to the emails that are not sent. You can use several variables for the email subject.

* **`-separate-archivelog-backups`**

  Separates the archive log backup from datafile backup. This is an optional parameter you can provide while creating the profile. After you separate the backups are separated using this option, you can create either data files-only backup or archive logs-only backup.

* **`-retain-archivelog-backups -hours hours | -days days | -weeks weeks| -months months`**

  Specifies that the archive log backups are retained based on the archive log retention duration (hourly, daily, weekly, monthly).

* **`-include-with-online-backups | -no-include-with-online-backups`**

  Specifies that the archive log backup is included along with the online database backup.

  Specifies that the archive log backups are not included along with the online database backup.

* **`-dump`**

  Specifies that the dump files are collected after the successful profile create operation.

* **`-quiet`**

  Displays only error messages in the console. The default is to display error and warning messages.

* **`-verbose`**

  Displays error, warning, and informational messages in the console.

Example

The following example changes the login information for the database described by the profile and the email notification is configured for this profile:

```
smsap profile update -profile SALES1 -database -dbname SALESDB
-sid SALESDB -login -username admin2 -password d4jPe7bw -port 1521
-host server1 -profile-notification -success -e-mail Preston.Davis@org.com
-subject success
Operation Id [8abc01ec0e78ec33010e78ec3b410001] succeeded.
```

The smsap profile verify command

You can run the `profile verify` command to verify the profile set up. You must mount the database before running this command.

Syntax

```
smsap profile verify
-profile profile_name
[-quiet | -verbose]
```

Parameters

* **-profile**

  Specifies the profile to verify. The profile contains the identifier of the database and other database information.

* **-quiet**

  Displays only error messages in the console. The default is to display error and warning messages.

* **-verbose**

  Displays error, warning, and informational messages in the console.

Example

The following example verifies the profile:

```
smsap profile verify -profile profileA -verbose
  [ INFO] SMSAP-13505: SnapDrive environment verification passed.
[ INFO] SMSAP-13507: JDBC verification for "OS authenticated:
NEWDB/hostA.rtp.com" passed.
[ INFO] SMSAP-13506: SQLPlus verification for database SID "NEWDB" passed.
Environment: [ORACLE_HOME=E:\app\Administrator\product\11.2.0\dbhome_1]
[ INFO] SMSAP-07431: Saving starting state of the database:
Database[NEWDB(OPEN)], Service[RUNNING].
[ INFO] SMSAP-07431: Saving starting state of the database:
Database[NEWDB(OPEN)], Service[RUNNING].
```

```
[ INFO] SD-00016: Discovering storage resources for F:\.
[ INFO] SD-00017: Finished storage discovery for F:\.
[ INFO] SD-00016: Discovering storage resources for F:\.
[ INFO] SD-00017: Finished storage discovery for F:\.
[ INFO] SD-00016: Discovering storage resources for H:\.
[ INFO] SD-00017: Finished storage discovery for H:\.
[ INFO] SD-00016: Discovering storage resources for G:\.
[ INFO] SD-00017: Finished storage discovery for G:\.
[ INFO] SD-00016: Discovering storage resources for I:\.
[ INFO] SD-00017: Finished storage discovery for I:\.
[ WARN] SMSAP-05071: Database profile HADLEY is not eligible for fast
restore:  Restore Plan:
  Preview:

    The following components will be restored completely via: host side
file copy restore
      F:\NEWDB\SYSAUX01.DBF
      F:\NEWDB\SYSTEM01.DBF
      F:\NEWDB\UNDOTBS01.DBF
      F:\NEWDB\USERS01.DBF


  Analysis:

    The following reasons prevent certain components from being restored
completely via: storage side file system restore
      * Files in file system F:\ not part of the restore scope will be
reverted.
    Components not in restore scope:
      F:\_TESTCLN\CONTROL01.CTL
      F:\_TESTCLN\REDO_1.LOG
      F:\_TESTCLN\REDO_2.LOG
      F:\_TESTCLN\REDO_3.LOG
    Components to restore:
      F:\NEWDB\SYSAUX01.DBF
      F:\NEWDB\SYSTEM01.DBF
      F:\NEWDB\UNDOTBS01.DBF
      F:\NEWDB\USERS01.DBF

  * Reasons denoted with an asterisk (*) are overridable.

[ INFO] SMSAP-07433: Returning the database to its initial state: Database
[NEWDB(OPEN)], Service[RUNNING].
[ INFO] SMSAP-13048: Profile Verify Operation Status: SUCCESS
[ INFO] SMSAP-13049: Elapsed Time: 0:19:06.949
Operation Id [N5bc18bd5c3be27a795ce3857093a926a] succeeded.
```

The smsap repository create command

Syntax

This command creates a repository in which to store database profiles and associated credentials. This command also checks to see that the block size is adequate.

```
smsap repository create
-repository
-port repo_port
-dbname repo_service_name
-host repo_host
-login -username repo_username
[-force] [-noprompt]
[-quiet | -verbose]
```

Parameters

* **-repository**

  The options that follow -repository specify the details of the database for the repository

* **-port** *repo_port*

  Specifies the TCP port number used to access the database that stores the repository.

* **-dbname** *repo_service_name*

  Specifies the name of the database that stores the repository. Use either the global name or the SID.

* **-host** *repo_host*

  Specifies the name or IP address of the host computer the repository database runs on.

* **-login**

  Starts the repository login details.

* **-username** *repo_username*

  Specifies the user name needed to access the database that stores the repository.

* **-force**

  Attempts to force the creation of the repository. Using this option results in SnapManager prompting you to backup the repository before creating the repository.

* **-noprompt**

  Does not display the prompt to backup the repository before creating it if you use the `-force` option. Using the `-noprompt` option ensures the prompt does not appear, making it easier to create repositories using a

script.

- **-quiet**

  Displays only error messages on the console. The default is to display error and warning messages.

- **-verbose**

  Displays error, warning, and informational messages on the console.

Command example

The following example creates a repository in the database SMSAPREPO on the host hotspur.

```
smsap repository create -repository -port 1521 -dbname SMSAPREPO -host
hotspur -login -username grabal21 -verbose
SMSAP-09202 [INFO ]: Creating new schema as grabal21 on
jdbc:oracle:thin:@//hotspur:1521/SMSAPREPO.
SMSAP-09205 [INFO ]: Schema generation complete.
SMSAP-09209 [INFO ]: Performing repository version INSERT.
SMSAP-09210 [INFO ]: Repository created with version: 30
SMSAP-13037 [INFO ]: Successfully completed operation: Repository Create
SMSAP-13049 [INFO ]: Elapsed Time: 0:00:08.844
```

The smsap repository delete command

This command deletes a repository used to store database profiles and associated credentials. You can delete a repository only if there are no profiles in the repository.

Syntax

```
smsap repository delete
-repository
-port repo_port
-dbname repo_service_name
-host repo_host
-login -username repo_username
[-force] [-noprompt]
[-quiet | -verbose]
```

Parameters

- **-repository**

  The options that follow -repository specify the details of the database for the repository.

- **-port** *repo_port*

  Specifies the TCP port number used to access the database that stores the repository.

- **-dbname** *repo_service_name*

  Specifies the name of the database that stores the repository. Use either the global name or the SID.

- **-host** *repo_host*

  Specifies the name or IP address of the host computer the repository database runs on.

- **-login**

  Starts the repository login details.

- **-username** *repo_username*

  Specifies the user name needed to access the database that stores the repository.

- **-force**

  Attempts to force the deletion of the repository, even if there are incomplete operations. SnapManager issues a prompt if there are incomplete operations, asking if you are sure you want to delete the repository.

- **-noprompt**

  Does not prompt you before deleting the repository. Using the `-noprompt` option ensures the prompt does not appear, making it easier to delete repositories using a script.

- **-quiet**

  Displays only error messages on the console. The default is to display error and warning messages.

- **-verbose**

  Displays error, warning, and informational messages on the console.

Command example

The following example deletes the repository in the SALESDB database.

```
smsap repository delete -repository -dbname SALESDB
-host server1 -login -username admin -port 1527 -force -verbose
```

The smsap repository rollback command

This command enables you to roll back or revert from a higher version of SnapManager to the original version from which you upgraded.

Syntax

```
smsap repository rollback
-repository
-dbname repo_service_name
-host repo_host
-login -username repo_username
-port repo_port
-rollbackhost host_with_target_database
[-force]
[-quiet | -verbose]
```

Parameters

* **-repository**

  The options that follow -repository specify the details of the database for the repository.

* **-dbname** *repo_service_name*

  Specifies the name of the database that stores the repository. Use either the global name or the SID.

* **-host** *repo_host*

  Specifies the name or IP address of the host computer the repository database runs on.

* **-login**

  Starts the repository login details.

* **-username** *repo_username*

  Specifies the user name needed to access the database that stores the repository.

* **-rollbackhost** *host_with_target_database*

  Specifies the name of the host which will be rolled back from a higher version of SnapManager to the original lower version.

* **-port** *repo_port*

  Specifies the TCP port number used to access the database that stores the repository.

* **-force**

  Attempts to force the update of the repository. SnapManager prompts you to make a backup of the current repository before updating.

* **-noprompt**

  Does not display the prompt before updating the repository database. Using the -noprompt option

ensures the prompt does not appear, making it easier to update repositories using a script.

- **-quiet**

  Displays only error messages on the console. The default is to display error and warning messages.

- **-verbose**

  Displays error, warning, and informational messages on the console.

Example command

The following example updates the repository in the SALESDB database.

```
smsap repository rollback -repository -dbname SALESDB
-host server1 -login -username admin -port 1521 -rollbackhost hostA
```

The smsap repository rollingupgrade command

This command performs rolling upgrade on a single host or multiple hosts and their associated target databases from a lower version of SnapManager to a higher version. The upgraded host is managed only with the higher version of SnapManager.

Syntax

```
smsap repository rollingupgrade
-repository
-dbname repo_service_name
-host repo_host
-login -username repo_username
-port repo_port
-upgradehost host_with_target_database
[-force] [-noprompt]
[-quiet | -verbose]
```

Parameters

- **-repository**

  The options that follow -repository specify the details of the database for the repository.

- **-dbname** *repo_service_name*

  Specifies the name of the database that stores the repository. Use either the global name or the SID.

- **-host** *repo_host*

Specifies the name or IP address of the host computer the repository database runs on.

* **-login**

  Starts the repository login details.

* **-username** *repo_username*

  Specifies the user name needed to access the database that stores the repository.

* **-upgradehost** *host_with_target_database*

  Specifies the name of the host which will be rolling upgraded from a lower version of SnapManager to a higher version.

* **-port** *repo_port*

  Specifies the TCP port number used to access the database that stores the repository.

* **-force**

  Attempts to force the update of the repository. SnapManager prompts you to make a backup of the current repository before updating.

* **-noprompt**

  Does not display the prompt before updating the repository database. Using the `-noprompt` option ensures the prompt does not appear, making it easier to update repositories using a script.

* **-quiet**

  Displays only error messages on the console. The default is to display error and warning messages.

* **-verbose**

  Displays error, warning, and informational messages on the console.

Example command

The following example updates the repository in the SALESDB database.

```
smsap repository rollingupgrade -repository -dbname SALESDB
-host server1 -login -username admin -port 1521 -upgradehost hostA
```

The smsap repository show command

This command displays information about the repository.

Syntax

```
smsap repository show
-repository
-dbname repo_service_name
-host repo_host
-port repo_port
-login -username repo_username
[-quiet | -verbose]
```

Parameters

* **-repository**

  The options that follow -repository specify the details of the database for the repository.

* **-dbname** *repo_service_name*

  Specifies the name of the database that stores the repository. Use either the global name or the SID.

* **-host** *repo_host*

  Specifies the name or IP address of the host computer the repository database runs on.

* **-login**

  Starts the repository login details.

* **-username** *repo_username*

  Specifies the user name needed to access the database that stores the repository.

* **-port** *repo_port*

  Specifies the TCP port number used to access the database that stores the repository.

* **-quiet**

  Displays only error messages on the console. The default is to display error and warning messages.

* **-verbose**

  Displays error, warning, and informational messages on the console.

Command example

The following example shows details about the repository in the SALESDB database.

```
smsap repository show -repository -dbname SALESDB  -host server1
-port 1521 -login -username admin
Repository Definition:
User Name: admin
Host Name: server1
Database Name: SALESDB
Database Port: 1521
Version: 28
Hosts that have run operations using this repository: 2
server2
server3
Profiles defined in this repository: 2
GSF5A
GSF3A
Incomplete Operations: 0
```

The smsap repository update command

This command updates the repository that stores database profiles and associated credentials when you upgrade SnapManager. Any time you install a new version of SnapManager, you must run the repository update command before you can use the new version. You are able to use this command only if there are no incomplete commands in the repository.

Syntax

```
smsap repository update
-repository
-dbname repo_service_name
-host repo_host
-login -username repo_username
-port repo_port
[-force] [-noprompt]
[-quiet | -verbose]
```

Parameters

* **-repository**

  The options that follow -repository specify the details of the database for the repository.

* **-dbname** *repo_service_name*

  Specifies the name of the database that stores the repository. Use either the global name or the SID.

- **-host** *repo_host*

  Specifies the name or IP address of the host computer the repository database runs on.

- **-login**

  Starts the repository login details.

- **-username** *repo_username*

  Specifies the user name needed to access the database that stores the repository.

- **-port** *repo_port*

  Specifies the TCP port number used to access the database that stores the repository.

- **-force**

  Attempts to force the update of the repository. SnapManager prompts you to make a backup of the current repository before updating.

- **-noprompt**

  Does not display the prompt before updating the repository database. Using the `-noprompt` option ensures the prompt does not appear, making it easier to update repositories using a script.

- **-quiet**

  Displays only error messages on the console. The default is to display error and warning messages.

- **-verbose**

  Displays error, warning, and informational messages on the console.

Example command

The following example updates the repository in the SALESDB database.

```
smsap repository update -repository -dbname SALESDB
-host server1 -login -username admin -port 1521
```

The smsap schedule create command

You can use the `schedule create` command to schedule a backup to be created at a specific time.

Syntax

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

[-backup-dest path1 [ , path2]]
[-exclude-dest path1 [ , path2]]
[-prunelogs {-all | -until-scn until-scn | -until -date  yyyy-MM-
dd:HH:mm:ss] | -before {-months | -days | -weeks | -hours}}
-prune-dest prune_dest1,[prune_dest2]]
-schedule-name schedule_name
[-schedule-comment schedule_comment]
-interval {-hourly | -daily | -weekly | -monthly | -onetimeonly}
-cronstring cron_string
-start-time {start_time <yyyy-MM-dd HH:mm>}
-runasuser runasuser
[-taskspec taskspec]
-force
[-quiet | -verbose]
```

Parameters

- **-profile *profile_name***

  Specifies the name of the profile related to the database that you want to schedule the backup for. The profile contains the identifier of the database and other database information.

- **-auto option**

  If the database is in the mounted or offline state, SnapManager performs an offline backup. If the database is in the open or online state, SnapManager performs an online backup. If you use the `-force` option with the `-offline` option, SnapManager forces an offline backup even if the database is currently online.

- **-online option**

  Specifies an online database backup.

- **-offline option**

  Specifies an offline backup while the database is in the shutdown state. If the database is in the open or

mounted state, the backup fails. If the `-force` option is used, SnapManager attempts to alter the database state to shut down the database for an offline backup.

* **`-full option`**

  Backs up the entire database. This includes all of the data, archived log, and control files. The archived redo logs and control files are backed up no matter what type of backup you perform. If you want to back up only a portion of the database, use the `-files` option or `-tablespaces` option.

* **`-files` *list***

  Backs up only the specified data files plus the archived log and control files. Separate the list of file names with spaces. If the database is in open state, SnapManager verifies that the appropriate tablespaces are in online backup mode.

* **`-tablespaces` *tablespaces***

  Backs up only the specified database tablespaces plus the archived log and control files. Separate the tablespace names with spaces. If the database is in open state, SnapManager verifies that the appropriate tablespaces are in online backup mode.

* **`-label` *name***

  Specifies an optional name for this backup. This name must be unique within the profile. The name can contain letters, numbers, underscore (_), and hyphen (-). It cannot start with a hyphen.

  If you do not specify a label, SnapManager creates a default label in the scope_type_date format:

  ◦ Scope is either F to indicate a full backup or P to indicate a partial backup.

  ◦ Type is C to indicate an offline (cold) backup, H to indicate an online (hot) backup, or A to indicate auto backup, for example, P_A_20081010060037IST.

  ◦ Date is the year, month, day, and time of the backup.

    SnapManager uses a 24-hour clock.

  For example, if you performed a full backup with the database offline on 16th January 2007, at 5:45:16 p.m. Eastern Standard Time, SnapManager would create the label F_C_20070116174516EST.

* **`-comment` *string***

  Specifies an optional comment to describe this backup. Enclose the string within single quotation marks (').

  > ⓘ Some shells strip quotation marks off. If that is true for your shell, you must include the quotation mark with a backslash (\). For example, you might need to enter: `\' this is a comment\'`.

* **`-verify option`**

  Verifies that the files in the backup are not corrupt by running the Oracle dbv utility.

  > ⓘ If you specify the `-verify` option, the backup operation is not completed until the verify operation is complete.

* **-force option**

  Forces a state change if the database is not in the correct state. For example, SnapManager might change the state of the database from online to offline, based on the type of backup you specify and the state that the database is in.

  ○ If the local instance is in shutdown state and at least one instance is open, you can change the local instance to mounted by using `-force` option.

  ○ If no instance is open, you can change the local instance to open by using `-force` option.

* **-retain { -hourly | -daily | -weekly | -monthly | -unlimited}**

  Specifies whether the backup should be retained on an hourly, daily, weekly, monthly, or unlimited basis. If `-retain` option is not specified, the retention class defaults to `-hourly`. To retain backups forever, use the `-unlimited` option. The `-unlimited` option makes the backup ineligible for deletion by the retention policy.

* **-archivelogs**

  Specifies creation of an archive log backup.

* **-backup-dest** *path1*, **[, [***path2***]]**

  Specifies the archive log destinations for archive log backup.

* **-exclude-dest** *path1*, **[, [***path2***]]**

  Specifies the archive log destinations to be excluded from the backup.

* **-prunelogs {-all | -until-scnuntil-scn | -until-date** *yyyy-MM-dd:HH:mm:ss* **| -before {-months | -days | -weeks | -hours}**

  Specifies whether to delete the archive log files from the archive log destinations based on options provided while creating a backup. The `-all` option deletes all of the archive log files from the archive log destinations. The `-until-scn` option deletes the archive log files until a specified system change number (SCN). The `-until-date` option deletes the archive log files until the specified time period. The `-before` option deletes the archive log files before the specified time period (days, months, weeks, hours).

* **-schedule-name** *schedule_name*

  Specifies the name that you provide for the schedule.

* **-schedule-comment** *schedule_comment*

  Specifies an optional comment to describe about scheduling the backup.

* **-interval { -hourly | -daily | -weekly | -monthly | -onetimeonly}**

  Specifies the time interval by which the backups are created. You can schedule the backup on an hourly, daily, weekly, monthly, or one time only basis.

* **-cronstring** *cron_string*

  Specifies scheduling the backup using cronstring. Cron expressions are used to configure instances of

CronTrigger. Cron expressions are strings that are made up of the following subexpressions:

- 1 refers to seconds.
- 2 refers to minutes.
- 3 refers to hours.
- 4 refers to a day in a month.
- 5 refers to the month.
- 6 refers to a day in a week.
- 7 refers to the year (optional).

- **-start-time** *yyyy-MM-dd HH:mm*

  Specifies the start time of the scheduled operation. The schedule start time should be included in the yyyy-MM-dd HH:mm format.

- **-runasuser** *runasuser*

  Specifies changing the user (root user or Oracle user) of the scheduled backup operation while scheduling the backup.

- **-taskspec** *taskspec*

  Specifies the task specification XML file that can be used for preprocessing activity or post-processing activity of the backup operation. The complete path of the XML file must be provided with the `-taskspec` option.

- **-quiet**

  Displays only error messages in the console. The default is to display error and warning messages.

- **-verbose**

  Displays error, warning, and informational messages in the console.

The smsap schedule delete command

This command deletes a backup schedule when it is are no longer necessary.

Syntax

```
smsap schedule delete
-profile profile_name
-schedule-name schedule_name
[-quiet | -verbose]
```

Parameters

- **-profile** *profile_name*

Specifies the name of the profile related to the database you want to delete a backup schedule. The profile contains the identifier of the database and other database information.

* **-schedule-name** *schedule_name*

Specifies the schedule name you provided while creating a backup schedule.

The smsap schedule list command

This command lists the scheduled operations associated with a profile.

Syntax

```
smsap schedule list
-profile profile_name
[-quiet | -verbose]
```

Parameters

* **-profile** *profile_name*

Specifies the name of the profile related to the database, using which you can view a list of scheduled operations. The profile contains the identifier of the database and other database information.

The smsap schedule resume command

This command resumes the suspended backup schedule.

Syntax

```
smsap schedule resume
-profile profile_name
-schedule-name schedule_name
[-quiet | -verbose]
```

Parameters

* **-profile** *profile_name*

Specifies the name of the profile related to the database you want to resume the suspended backup schedule. The profile contains the identifier of the database and other database information.

* **-schedule-name** *schedule_name*

Specifies the schedule name you provided while creating a backup schedule.

The smsap schedule suspend command

This command suspends a backup schedule until the backup schedule is resumed.

Syntax

```
smsap schedule suspend
-profile profile_name
-schedule-name schedule_name
[-quiet | -verbose]
```

Parameters

* **-profile** *profile_name*

  Specifies the name of the profile related to the database you want to suspend a backup schedule. The profile contains the identifier of the database and other database information.

* **-schedule-name** **schedule_name**

  Specifies the schedule name you provided while creating a backup schedule.

The smsap schedule update command

This command updates the schedule for a backup.

Syntax

```
smsap schedule update
-profile profile_name
-schedule-name schedule_name
[-schedule-comment schedule_comment]
-interval {-hourly | -daily | -weekly | -monthly | -onetimeonly}
-cronstring cron_string
-start-time {start_time <yyyy-MM-dd HH:mm>}
-runasuser runasuser
[-taskspec taskspec]
-force
[-quiet | -verbose]
```

Parameters

* **-profile** *profile_name*

  Specifies the name of the profile related to the database you want to schedule the back up. The profile contains the identifier of the database and other database information.

- **-schedule-name** *schedule_name*

  Specifies the name that you provide for the schedule.

- **-schedule-comment** *schedule_comment*

  Specifies an optional comment to describe about scheduling the backup.

- **-interval { -hourly | -daily | -weekly | -monthly | -onetimeonly}**

  Indicates the time interval by which the backups are created. You can schedule the backup on an hourly, daily, weekly, monthly, or one time only.

- **-cronstring** *cron_string*

  Specifies to schedule the backup using cronstring. Cron expressions are used to configure instances of CronTrigger. Cron expressions are strings that are actually made up of seven sub-expressions:

  - 1 refers to seconds
  - 2 refers to minutes
  - 3 refers to hours
  - 4 refers to a day in a month
  - 5 refers to the month
  - 6 refers to a day in a week
  - 7 refers to the year (optional)

- **-start-time** *yyyy-MM-dd HH:mm*

  Specifies the start time of the schedule operation. The schedule start time should be included in the format of yyyy-MM-dd HH:mm.

- **-runasuser** *runasuser*

  Specifies to change the user of the scheduled backup operation while scheduling the backup.

- **-taskspec** *taskspec*

  Specifies the task specification XML file that can be used for pre-processing activity or post-processing activity of the backup operation. The complete path of the XML file should be provided which give the `-taskspec` option.

The smsap storage list command

You can run the `storage list` command to display the list of storage systems associated with a particular profile.

Syntax

```
smsap storage list
-profile profile
```

Parameters

* **-profile** *profile*

  Specifies the name of the profile. The name can be up to 30 characters long and must be unique within the host.

Example

The following example displays the storage systems associated with the profile mjullian:

```
smsap storage list -profile mjullian
```

```
Sample Output:
Storage Controllers
--------------------
FAS3020-RTP07OLD
```

The smsap storage rename command

This command updates the name or IP address of the storage system.

Syntax

```
smsap storage rename
-profile profile
-oldname old_storage_name
-newname new_storage_name
[-quiet | -verbose]
```

Parameters

* **-profile** *profile*

  Specifies the name of the profile. This name can be up to 30 characters long and must be unique within the host.

* **-oldname** *old_storage_name*

  Specifies the IP address or name of the storage system before the storage system is renamed. You must enter the IP address or name of the storage system that is displayed when you run the `smsap storage`

list command.

- **-newname** *new_storage_name*

  Specifies the IP address or name of the storage system after the storage system is renamed.

- **-quiet**

  Displays only error messages in the console. The default is to display error and warning messages.

- **-verbose**

  Displays error, warning, and informational messages in the console.

Example

The following example uses the `smsap storage rename` command to rename the storage system:

```
smsap storage rename -profile mjullian -oldname lech -newname hudson
-verbose
```

The smsap system dump command

You can run the `system dump` command to create a JAR file that contains diagnostic information about the server environment.

Syntax

```
smsap system dump
[-quiet | -verbose]
```

Parameters

- **-quiet**

  Displays only error messages in the console. The default is to display error and warning messages.

- **-verbose**

  Displays error, warning, and informational messages in the console.

Example of the system dump command

The following example uses the `smsap system dump` command to create a JAR file:

```
smsap system dump
Path: C:\\userhomedirectory\\netapp\\smsap\\3.3.0\\smsap_dump_hostname.jar
```

The smsap system verify command

This command confirms that all the components of the environment required to run SnapManager are set up correctly.

Syntax

```
smsap system verify
[-quiet | -verbose]
```

Parameters

* **-quiet**

  Displays only error messages on the console. The default is to display error and warning messages.

* **-verbose**

  Displays error, warning, and informational messages on the console.

Example of the system verify command

The following example uses the smsap system verify command.

```
smsap system verify
SMSAP-13505 [INFO ]: Snapdrive verify passed.
SMSAP-13037 [INFO ]: Successfully completed operation: System Verify
SMSAP-13049 [INFO ]: Elapsed Time: 0:00:00.559
Operation Id [N4f4e910004b36cfecee74c710de02e44] succeeded.
```

The smsap version command

You can run the `version` command to determine which version of SnapManager you are running on your local host.

Syntax

```
smsap version
[-quiet | -verbose]
```

Parameters

- **-quiet**

  Displays only error messages in the console. The default is to display error and warning messages.

- **-verbose**

  Displays the build date and contents of each profile. Also displays error, warning, and informational messages in the console.

Example of the version command

The following example displays the version of the SnapManager:

```
smsap version
SnapManager for SAP Version: 3.3.1
```

Troubleshooting SnapManager

You can find information about some of the most common issues that might occur and how you can resolve them.

The following table describes common issues and possible solutions:

Issue-driven question	Possible solution
Are the target database and listener running?	Run the `lsnrctl status` command. Ensure that the database instance is registered with the listener.
Is the storage visible?	Perform the following steps: 1. Right-click My Computer, and then select **Manage**. 2. Click **storage** > **snapdrive** > **Hostname** > **Disks**.

Issue-driven question	Possible solution
Is the SnapManager server running?	Check the status, and then start the server by using Service Configuration. Before you can use the graphical user interface (GUI) or the command-line interface (CLI) to initiate SnapManager commands related to profiles, the server must be running. You can create or update repositories without starting the server, but to execute all other SnapManager operations, the server must be running. To start the SnapManager server, enter the following command: `smsap_server start`
Are all the components required to run SnapManager set up correctly?	Run the `smsap system verify` command to verify that SnapDrive is set up correctly.
Do you have the correct version of SnapManager?	Use the `smsap version` command to check the SnapManager version.

Have you looked at the SnapManager log files to determine if the error messages can help isolate the issue?	SnapManager records all log entries into one set of rotating log files. The log files are found at `C:\program_files\NetApp\SnapManager for SAP\logs`.
	
	If you are using Windows 2008, the logs are at the following locations:
	
	• Operation logs:
	
	   ◦ `C:\Program Files\NetApp\SnapManager for SAP\var\log\smsap`
	
	• Client logs:
	
	   ◦ `C:\Users\Administrator\AppData\Roaming\NetApp\smsap\3.3.0\`
	
	It might also be helpful to look at the logs in the following location: `C:\Documents and Settings\hostname\Application Data\NetApp\smsap\3.3.0\log`
	
	Each operation log is written to its own log file of the form `smsap_of_date_time.log`.

Issue-driven question	Possible solution
If you have archive logs stored on a storage system that is not running Data ONTAP, have you excluded them from consideration for backup with SnapManager?	The `smsap.config` file enables you to exclude certain archive log files. For Windows, the file is at the following location: `C:\program_files\NetApp\smsap\propertie s\smsap.config` Use the format mentioned in the file to exclude the local archive logs. For additional information, see the "Setting configuration properties" topic. You can also exclude the archive log destinations while creating a backup from the SnapManager CLI. For additional information, see the "Creating database backups" topic. You can also exclude the archive log destinations while creating a backup from the SnapManager GUI.
Do you have an MS-DOS window open in the directory in which you are attempting to install or upgrade SnapManager on Windows?	You will see an error message similar to the following: ``` Directory C:\Program Files\NetApp\SnapManager for SAP\bin is currently in use by another program. Any window, opened by you or another user, that is currently referencing this directory must be closed before installation can proceed. ``` Close the window and attempt the installation or upgrade again.

Issue-driven question	Possible solution
Were you unable to connect to the repository?	If connecting to a repository fails, run the `lsnrctl status` command on the repository database and check the active service names. When SnapManager connects to the repository database, it uses the service name of the database. Depending on how the listener is setup, this might be the short service name or the fully qualified service name. When SnapManager connects to a database for a backup, restore, or other operation, it uses the host name and the SID. If the repository does not initialize correctly because it is currently unreachable, you receive an error message asking whether you want to remove the repository. You can remove the repository from your current view so that you can perform operations on other repositories. Also, check whether the corresponding service is running.
Can system resolve the host name?	Check whether the specified host name is on a different subnet. If you receive an error message that SnapManager cannot resolve the host name, then add the host name in the host file.Add the host name to the file located at `C:\WINDOWS\system32\drivers\etc\hosts:` `xxx.xxx.xxx.xxx hostname IP address`
Is SnapDrive running?	To view the status of SnapDrive, go to Services, and then select the SnapDrive service.
Which storage systems are configured to be accessed with SnapDrive?	To find the storage systems configured for SnapDrive, perform the following steps: 1. Right-click **My Computer**, and then select **Manage**. 2. Click **Storage** > **SnapDrive**. 3. Right-click the host name, and then select **transport protocol settings**.

Issue-driven question	Possible solution
How can SnapManager GUI performance be improved?	• Ensure that you have valid user credentials for the repository, profile host, and profile. If your credential is invalid, then clear the user credentials for the repository, profile host, and profile. Reset the same user credentials that you set before for the repository, profile host, and profile. For additional information about setting the user credentials again, see "Setting credentials after clearing credential cache". • Close the unused profiles. If the number of profiles that you have opened is more, the SnapManager GUI performance slows down. • Check whether you enabled **Open On Startup** in the User Preferences window under the **Admin** menu, from the SnapManager GUI. If this is enabled, then the user configuration (`user.config`) file available at `C:\Documents and Settings\Administrator\Application Data\NetApp\smsap\3.3.0\gui\state` is displayed as `openOnStartup=PROFILE`. Because **Open On Startup** is enabled, you must check for recently opened profiles from the SnapManager GUI, using `lastOpenProfiles` in the user configuration (`user.config`) file: `lastOpenProfiles=PROFILE1,PROFILE2,PROFILE3,…` You can delete the profile names listed and always keep a minimum number of profiles as open. • Before installing the new version of SnapManager on the Windows-based environment, delete the SnapManager client-side entries available at the following location: C:\Documents and Settings\Administrator\Application Data\NetApp

Issue-driven question	Possible solution
SnapManager GUI takes more time to refresh when there are multiple SnapManager operations started and running simultaneously in the background. When you right-click the backup (that is already deleted but still gets displayed in the SnapManager GUI), the backup options for that backup are not enabled in the Backup or Clone window.	You need to wait until the SnapManager GUI gets refreshed, and then check for the backup status.

What would you do when the Oracle database is not set in English?	SnapManager operations might fail if the language for an Oracle database is not set to English.Set the language of the Oracle database to English:
	
	1. Verify that the NLS_LANG environment variable is not set: echo%NLS_LANG%
	
	2. Add the following line to the `wrapper.conf` file located at `C:\SnapManager_install_directory\service:`
	
	`set.NLS_LANG=AMERICAN_AMERICA.WE8MSWIN1252`
	
	3. Restart the SnapManager server:
	
	`smsap_server restart`
	
	ⓘ If the system environment variable is set to NLS_LANG, you must edit the script to not overwrite NLS_LANG.

Issue-driven question	Possible solution

What would you do when the backup scheduling operation fails if the repository database points to more than one IP and each IP has a different host name?	1. Stop the SnapManager server.
	
	2. Delete the schedule files in the repository directory from the hosts where you want to trigger the backup schedule.
	
	   The schedule file names can be in the following formats:
	
	      ○ repository#repo_username#repository_database_name#repository_host#repo_port
	
	      ○ repository-repo_usernamerepository_database_name-repository_host-repo_port
	
	   (i) You must ensure that you delete the schedule file in the format that matches the repository details.
	
	1. Restart the SnapManager server.
	
	2. Open other profiles under the same repository from the SnapManager GUI to ensure that you do not miss any schedule information of those profiles.

What would you do when the SnapManager operation fails with credential file lock error?	SnapManager locks the credential file before updating, and unlocks it after updating.When multiple operations run simultaneously, one of the operations might lock the credential file to update it. If another operation tries to access the locked credential file at the same time, the operation fails with the file lock error. Configure the following parameters in the smsap.config file depending on the frequency of simultaneous operations: • fileLock.retryInterval = 100 milliseconds • fileLock.timeout = 5000 milliseconds (i) The values assigned to the parameters must be in milliseconds.

Issue-driven question	Possible solution

What would you do when the backup verify operation's intermediate status shows failed in the Monitor tab even though the backup verify operation is still running?	The error message is logged in the sm_gui.log file. You must look in the log file to determine the new values for the `operation.heartbeatInterval` and `operation.heartbeatThreshold` parameters which will resolve this issue.
	
	1. Add the following parameters in the `smsap.config` file:
	
	   ◦ `operation.heartbeatInterval` = 5000
	
	   ◦ `operation.heartbeatThreshold` = 5000 The default value assigned by SnapManager is 5000.
	
	2. Assign the new values to these parameters.
	
	   (i) The values assigned to the parameters must be in milliseconds.
	
	3. Restart the SnapManager server and perform the operation again.

What to do when you encounter a heap-space issue?	When you encounter a heap-space issue during SnapManager for SAP operations, you must perform the following steps:
	
	1. Navigate to the SnapManager for SAP installation directory.
	
	2. Open the `launchjava` file from the `installationdirectory\bin\launchjava` path.
	
	3. Increase the value of the java -Xmx160m Java heap-space parameter.
	
	   For example, you can increase the default value of 160m to 200m.
	
	   (i) If you have increased the value of the Java heap-space parameter in the earlier versions of SnapManager for SAP, you should retain that value.

Issue-driven question	Possible solution
What would you do when the SnapManager services do not start in a Windows environment and the following error message is displayed: Windows could not start Snap Manager on Local Computer. For more information, review the System Event log. If this is a non-Microsoft service, contact service vendor, and refer to service-specific error code 1?	Configure the following parameters in the wrapper.conf file located at `Installation_directory\service`. • The wrapper startup timeout parameter defines the maximum permissible time between the wrapper starting the Java Virtual Machine (JVM) and response from the JVM that the application has started. The default value is set to 90 seconds. However, you can change a value greater than 0. If you specify an invalid value, the default is used instead. • The `wrapper.ping.timeout` parameter defines the maximum permissible time between the wrapper pinging the JVM and the response from the JVM. The default value is set to 90 seconds. However, you can change to a value greater than 0. If you specify an invalid value, the default is used instead.

Dump files

The dump files are compressed log files containing information about SnapManager and its environment. The different types of log files created are operation, profile, and system dump file.

You can use the dump command or the **Create Diagnostics** tab in the graphical user interface (GUI) to collect information about an operation, a profile, or the environment. A system dump does not require a profile; however, the profile and operation dumps require profiles.

SnapManager includes the following diagnostic information in the dump file:

- The steps performed
- The length of time for each step to complete
- The outcome of each step
- Error, if any, that occurred during the operation

> ⓘ SnapManager log files or dump files enable read and write permissions only for the root users and the other users who belong to root user group.

SnapManager also includes the following information in the file:

- Operating system version and architecture

- Environment variables

- Java version

- SnapManager version and architecture

- SnapManager preferences

- SnapManager messages

- log4j properties

- SnapDrive version and architecture

- SnapDrive log files

- Oracle version

- Oracle OPatch local inventory details

- Repository database Oracle version

- Target database type (stand alone)

- Target database role (primary, physical standby, or logical standby)

- Target database Oracle Recovery Manager (RMAN) setup (no RMAN integration, RMAN with control files, or RMAN with catalog file)

- Target database Oracle version

- System identifier (SID) of the target database

- Repository database service name

- Database instances installed on the host

- Profile descriptor

- Shared memory maximum

- Swap space information

- Memory information

- Multipath environment

- Host utilities version

- Microsoft Internet Small Computer System Interface (iSCSI) software initiator version for Windows

- BACKINT interface versions

- BR tool version

- Patch level

- Output of the `system verify` command

The dump file also lists the SnapManager limitations on Windows.

SnapManager dump files also contain the SnapDrive data collector file and the Oracle alert log file. You can collect the Oracle alert log file by using the `smsap operation dump` and `smsap profile dump` commands.

> ⓘ System dump does not contain Oracle alert logs; however, the profile and operation dumps contain the alert logs.

Even if the SnapManager host server is not running, you can access the dump information by using the

command-line interface (CLI) or the GUI.

If you encounter a problem that you cannot resolve, you can send these files to NetApp Global Services.

Create operation-level dump files

You can use the `smsap operation dump` command with the name or ID of the failed operation to get log information about a particular operation. You can specify different log levels to gather information about a specific operation, profile, host, or environment.

Step

1. Enter the following command:

   `smsap operation dump -id` _guid_

   (i)   The `smsap operation dump` command provides a super set of the information provided by the `smsap profile dump` command, which in turn provides a super set of the information provided by the `smsap system dump` command.

   Dump file location:

   ```
   Path:\<user-home>\Application
   Data\NetApp\smsap\3.3.0\smsap_dump_8abc01c814649ebd0114649ec69d0001.jar
   ```

Create profile-level dump files

You can find log information about a particular profile by using the `smsap profile dump` command with the name of the profile.

Step

1. Enter the following command:

   `smsap profile dump -profile` _profile_name_

   Dump file location:

   ```
   Path:\<user-home>\Application
   Data\NetApp\smsap\3.3.0\smsap_dump_8abc01c814649ebd0114649ec69d0001.jar
   ```

   (i)   If you encounter an error while creating a profile, use the `smsap system dump` command. After you have successfully created a profile, use the `smsap operation dump` and `smsap profile dump` commands.

Create system-level dump files

You can use the `smsap system dump` command to get log information about the

SnapManager host and environment. You can specify different log levels to collect information about a specific operation, profile, or host and environment.

Step

1. Enter the following command:

```
smsap system dump
```

Resulting dump

```
Path:\<user-home>\Application
Data\NetApp\smsap\3.3.0\smsap_dump_server_host.jar
```

How to locate dump files

The dump file is located at the client system for easy access. These files are helpful if you need to troubleshoot a problem related to profile, system, or any operation.

The dump file is located in the user's home directory on the client system.

- If you are using the graphical user interface (GUI), the dump file is located at:

```
user_home\Application Data\NetApp\smsap\3.3.0\smsap_dump
dump_file_type_name
server_host.jar
```

- If you are using the command-line interface (CLI), the dump file is located at:

```
user_home\.netapp\smsap\3.3.0\smsap_dump_dump_file_type_name
server_host.jar
```

The dump file contains the output of the dump command. The name of the file depends on the information supplied. The following table shows the types of dump operations and the resulting file names:

Type of dump operation	Resulting file name
Operation dump command with operation ID	`smsap_dump_operation-id.jar`

Type of dump operation	Resulting file name
Operation dump command with no operation ID	`smsap operation dump -profile VH1 -verbose` The following output is displayed: ``` smsap operation dump -profile VH1 -verbose [INFO] SMSAP-13048: Dump Operation Status: SUCCESS [INFO] SMSAP-13049: Elapsed Time: 0:00:01.404 Dump file created. Path: user_home\Application Data\ontap\smsap\3.3.0\smsap_dump_ VH1_kaw.rtp.foo.com.jar ```
System dump command	`smsap_dump_host-name.jar`
Profile dump command	`smsap_dump_profile-name_host-name.jar`

How to collect dump files

You can include `-dump` in the SnapManager command to collect the dump files after a successful or failed SnapManager operation.

You can collect dump files for the following SnapManager operations:

- Creating profiles
- Updating profiles
- Creating backups
- Verifying backups
- Deleting backups
- Freeing backups
- Mounting backups
- Unmounting backups
- Restoring backups
- Creating clones
- Deleting clones

> ⓘ When you create a profile, you can collect dump files only if the operation is successful. If you encounter an error while creating a profile, you must use the `smsap system dump` command. For successful profiles, you can use the `smsap operation dump` and `smsap profile dump` commands to collect the dump files.

Example

```
smsap backup create -profile targetdb1_prof1 -auto -full -online -dump
```

Collecting additional log information for easier debugging

If you need additional logs to debug a failed SnapManager operation, you must set an external environment variable `server.log.level`. This variable overrides the default log level and dumps all the log messages in the log file. For example, you can change the log level to DEBUG, which logs additional messages and can assist in debugging issues.

The SnapManager logs can be found at the following locations:

* `SnapManager_install_directory\log`

To override the default log level, you must perform the following steps:

1. Create a `platform.override` text file in the SnapManager installation directory.

2. Add the `server.log.level` parameter in the `platform.override` text file.

3. Assign a value (***TRACE, DEBUG, INFO, WARN, ERROR, FATAL, or PROGRESS***) to the `server.log.level` parameter.

   For example, to change the log level to *ERROR*, set the value of `server.log.level` to *ERROR*.

   `server.log.level=ERROR`

4. Restart the SnapManager server.

   (i) If the additional log information is not required, you can delete the `server.log.level` parameter from the `platform.override` text file.

SnapManager manages the volume of server log files based on the user-defined values of the following parameters in the `smsap.config` file:

* `log.max_log_files`
* `log.max_log_file_size`
* `log.max_rolling_operation_factory_logs`

Troubleshooting clone issues

You can find information about that might occur during clone operations and how you can resolve them.

Symptom	Explanation	Workaround
The clone operation fails when the archive destination is set to `USE_DB_RECOVERY_FILE_DEST.`	When the archive destination is referring to `USE_DB_RECOVERY_FILE_DEST`, Flash recovery area (FRA) actively manages the archive log. SnapManager does not use the FRA location during clone or restore operations and thus the operations fail.	Change the archive destination to actual archive log location instead of the FRA location.

Symptom	Explanation	Workaround
The clone operation fails with the following error message: `Cannot perform operation: Clone Create. Root cause: ORACLE-00001: Error executing SQL: [ALTER DATABASE OPEN RESETLOGS;]. The command returned: ORA-01195: online backup of file 1 needs more recovery to be consistent.`	This issue occurs if Oracle listener fails to connect to the database.	If you are using SnapManager GUI to clone a backup, perform the following actions: 1. From the Repository tree, click **Repository** > **Host** > **Profile** to display the backups. 2. Right-click the backup that you want to clone and select **Clone**. 3. On the Clone Initialization page, enter the mandatory values and select the clone specification method. 4. On the Clone Specification page, select **Parameters**. 5. Click the **Parameter** tab. 6. In the **Parameter Name** field, enter the name as `local_listener` and click **OK**. 7. Select the **Override Default** check box for the local_listener row. 8. Click any parameter, then double-click the local_listener parameter, and enter the following value: `(ADDRESS=(PROTOCOL=TCP)(HOST=<your_host_name>)(PORT=<port#>))` 9. Click **Save To File**. 10. Click **Next** and continue with the clone create wizard. If you are using CLI to clone a backup, you must include the following information in the **<parameters>** tag of the clone specification file:

Symptom	Explanation	Workaround
The clone operation fails with an error message saying that the mountpoint you are using is already in use.	SnapManager does not let you mount a clone over an existing mountpoint. So an incomplete clone did not remove the mountpoint.	Specify a different mountpoint to be used by the clone, or unmount the problematic mountpoint.
The clone operation fails with an error message about data files not having a .dbf extension.	Some versions of the Oracle NID utility do not work with data files unless the files use a .dbf extension.	• Rename the data file to give it a .dbf extension. • Repeat the backup operation. • Clone the new backup.
The clone operation fails due to unmet requirements.	You are attempting to create a clone; however, some of the prerequisites have not been met.	Proceed as described in *Creating a clone* to meet the prerequisites.
SnapManager for SAP fails to clone Oracle 10gR2 (10.2.0.5) physical Oracle Data Guard Standby databases.	SnapManager for SAP does not disable the managed recovery mode while performing an offline backup of the Oracle 10gR2 (10.2.0.5) physical standby databases created using Oracle Data Guard services. Due to this issue, the offline backup taken is inconsistent. When SnapManager for SAP tries to clone the offline backup, it does not even try to perform any recovery on the cloned database. Because the backup is inconsistent, the cloned database requires recovery, and thus SAP fails to create the clone successfully.	Upgrade the Oracle database to the Oracle 11gR1 (11.1.0.7 patch).

Troubleshooting graphical user interface issues

You can find information about some common known graphical user interface (GUI) issues that might help you resolve them.

Issue	Explanation	Workaround

While accessing the SnapManager GUI to perform an operation, the following error message might be displayed: `SMSAP-20111 : Authentication failed for user on host.`	This issue occurs if the password of the user is changed in the host on which the SnapManager server is running.After the password is changed, the credential cache that is created for the user who launched the GUI becomes invalid. SnapManager GUI still uses the credentials in the cache to authenticate and thus the authentication fails.	You must perform one of following tasks:
		
		• Delete the credentials of the user whose password was changed and then add the new credentials in the cache by running the following commands:
		
		  a. `smsap credential delete`
		
		  b. `smsap credential set`
		
		• Clear the entire cache by running the `smsap credential clear` command. Re-open the GUI and set the credentials, if prompted.

Security warning is displayed while using Java Web Start to access SnapManager GUI.	While accessing SnapManager GUI using Java Web Start, a security warning is displayed.This issue occurs because JNLP jars are self-signed and Java version used by SnapManager does not allow self-signed jars at high security level.	Either change the security settings to medium in the java control panel or add the SnapManager GUI URL to the exception list.
The SnapManager web start GUI displays the incorrect version.	After downgrading SnapManager from a later version to an earlier version when you launch the web start GUI, the later version of the SnapManager web start GUI is launched.	You must also clear the cache by performing the following steps: 1. Click **Start** and select **Run**. 2. Enter the following: `javaws -viewer` 3. On the Java cache viewer screen, right-click the SnapManager application and select **Delete**.

Issue	Explanation	Workaround

When you restart the GUI and try to check the backups for a certain profile, you see only the names of the profiles.	SnapManager does not display any information about a profile until you open it.	Perform the following steps:
		
		1. Right-click the profile and select **Open** from the menu.
		
		   SnapManager displays the Profile Authentication dialog box.
		
		2. Enter the host user name and password.
		
		   SnapManager displays the backup list.
		
		   ⓘ You only need to authenticate the profile once as long as the credentials are valid and remain in the cache.

The GUI installation on Windows succeeds, but with errors.	The user account used to install the GUI does not have enough permissions to set up the icons and shortcuts for all the users. The user account does not have permissions to modify the `C:\Documents and Settings\All Users` directory.	• Reinstall the GUI with different settings. In Choose Icon Availability, clear the **Make these shortcuts available to all users of this PC** check box. • Log in using a user account that is not restricted and reinstall the GUI.
When you open the first repository in the GUI, an error message similar to the following is displayed: `The Profile name XXXX clashes with previously loaded repository.`	Identically named profiles cannot exist in a repository. Also, you can open only one repository at a time.	Reference the conflicting profiles from two different operating system (OS) users or rename the profile by issuing an SQL statement for the repository: **UPDATE SMSAP_33_PROFILE SET NAME = 'NEW_NAME' WHERE NAME = 'OLD_NAME'**

Issue	Explanation	Workaround
An error message similar to the following is displayed: `SMSAP-01092: Unable to initialize repository repo1@ does not exist:repo1SMSAP-11006: Cannot resolve host does not exist`	The repository is inaccessible, perhaps because it no longer exists. The GUI initializes the list of repositories from the credentials file.	The error message asks if you would like to remove this repository so that no attempt is made to load it in the future. If you do not need to access this repository, click **Delete** to remove it from the GUI view. This removes the reference to the repository in the credentials file and the GUI does not attempt to load the repository again.
SnapManager takes a longer time to load the database tree structure and results in a timeout error message being displayed on the SnapManager GUI.	When you try to perform a partial backup operation from the SnapManager GUI, SnapManager tries to load the credentials for all the profiles, and if there are any invalid entries, SnapManager tries to validate the entry and this results in a timeout error message being displayed.	Delete the credentials of the unused host, repository, and profile by using the credential delete command from the SnapManager command-line interface (CLI).
The custom scripts for the preprocessing or postprocessing activity to occur before or after the backup, restore, or clone operations, are not visible from the SnapManager GUI.	When you add custom scripts in the custom backup, restore, or clone script location after you start the respective wizard, the custom scripts are not displayed under the Available Scripts list.	Restart the SnapManager host server and then open the SnapManager GUI.
You cannot use the clone specification XML file created in SnapManager (3.1 or earlier) for the clone operation.	From SnapManager 3.2 for SAP, the task specification section (task-specification) is provided as a separate task specification XML file.	If you are using SnapManager 3.2 for SAP, you must remove the task specification section from the clone specification XML or create a new clone specification XML file.SnapManager 3.3 or later does not support the clone specification XML file created in SnapManager 3.2 or earlier releases.

Issue	Explanation	Workaround
SnapManager operation on the GUI does not proceed after you have cleared user credentials by using the smsap credential clear command from the SnapManager CLI or by clicking **Admin** > **Credentials** > **Clear** > **Cache** from the SnapManager GUI.	The credentials set for the repositories, hosts, and profiles are cleared. SnapManager verifies user credentials before starting any operation.When user credentials are invalid, SnapManager fails to authenticate. When a host or a profile is deleted from the repository, the user credentials are still available in the cache. These unnecessary credential entries slow down the SnapManager operations from the GUI.	Restart the SnapManager GUI depending on how the cache is cleared. ⓘ • If you have cleared the credential cache from the SnapManager GUI, you do not need to exit the SnapManager GUI. • If you have cleared the credential cache from the SnapManager CLI, you must restart the SnapManager GUI. • If you have deleted the encrypted credential file manually, you must restart the SnapManager GUI. Set the credentials that you have given for the repository, profile host, and profile. From the SnapManager GUI, if there is no repository mapped under the Repositories tree, perform the following steps: 1. Click **Tasks** > **Add Exisiting repository** 2. Right-click the repository, click **Open**, and enter the user credentials in the **Repository Credentials Authentication** window. 3. Right-click the host under the repository, click **Open**, and enter the user credentials in **Host Credentials Authentication**. Right-click the profile under the host, click **Open**, and enter the user credentials in **Profile**

314

Issue	Explanation	Workaround
You cannot open the SnapManager GUI by using Java Web Start GUI due to weaker Secure Sockets Layer (SSL) cipher strength of the browser.	SnapManager does not support SSL ciphers weaker than 128 bits.	Upgrade the browser version and check the cipher strength.

Troubleshooting known issues

You should be aware of some known issues that might occur when you use SnapManager, and how to work around them.

SnapManager for SAP fails to identify the profile as a Cluster-Mode profile

If the Cluster-Mode profile name is not present in the `cmode_profiles.config` file in the SnapManager for SAP installation directory, the following error message might trigger:

```
Please configure DFM server using snapdrive config set -dfm user_name
appliance_name.
```

Also, while upgrading the SnapManager for SAP, if you delete the `/opt/NetApp/smsap/*` folder, then the `cmode_profiles.config` file that has the Cluster-Mode profile names also get deleted. This issue also triggers the same error message.

Workaround

Update the profile: `smsap profile update-profile <profile_name\>`

> ⓘ If SnapManager for SAP is installed in the `/opt/NetApp/smsap/` path, then the file location will be `/opt/NetApp/smsap/cmode_profile/cmode_profiles.config`.

The server fails to start

When starting the server, you might see an error message similar to the following:

```
SMSAP-01104: Error invoking command: SMSAP-17107: SnapManager Server failed to
start on port 8074 because of the following errors: java.net.BindException:
Address already in use
```

This might be because the SnapManager listening ports (27314 and 27315, by default) are currently in use by another application.

This error can also occur if the `smsap_server` command is already running, but SnapManager does not detect the existing process.

Workaround

You can reconfigure either SnapManager or the other application to use different ports.

To reconfigure SnapManager, edit the following file: `C:\Program Files\NetApp\SnapManager for SAP\properties\smsap.config`

You assign the following values:

- SMSAP Server.port=27314
- SMSAP Server.rmiRegistry.port=27315
- remote.registry.ocijdbc.port= 27315

  The remote.registry.ocijdbc.port must be the same as Server.rmiRegistry.port.

To start the SnapManager server, perform the following steps:

Steps

1. Click **Start** > **Control Panel** > **Administrative Tools** > **Services**.
2. In the Services window open, select NetApp SnapManager 3.3 for SAP.
3. You can start the server in one of three ways:
   - In the left panel, click **Start**.
   - Right-click NetApp SnapManager 3.3 for SAP and select **Start** from the drop-down menu.
   - Double-click NetApp SnapManager 3.3 for SAP and in the Properties window that opens, click **Start**.

Unable to manage archive log file destination names if the destination names are part of other destination names

While creating an archive log backup, if the user excludes a destination that is part of other destination names, then the other destination names are also excluded.

For example, assume that there are three destinations available to be excluded: `E:\\arch`, `G:\\arch`, and `H:\\arch`. While creating the archive log file backup, if you exclude `E:\\arch` by using the command

```
smsap backup create -profile almsamp1 -data -online -archivelogs  -exclude
-dest E:\\arch
```

, SnapManager for SAP excludes all the destinations starting with `E:\\arch`.

Workaround

- Add a path separator after destinations are configured in `v$archive_dest`. For example, change the `E:\\arch` to `E:\\arch\`.
- While creating a backup, include destinations instead of excluding any destination.

Repository database size grows with time and not with the number of backups

The repository database size grows with time because SnapManager operations insert or delete data within the schema in the repository database tables, which results in high index space usage.

Workaround

You must monitor and rebuild the indexes according to the Oracle guidelines to control the space consumed by the repository schema.

The SnapManager GUI cannot be accessed and SnapManager operations fail when the repository database is down

SnapManager operations fail and you cannot access the GUI when the repository database is down.

The following table lists the different actions you might want to perform, and their exceptions:

Operations	Exceptions
Opening a closed repository	The following error message is logged in `sm_gui.log: [WARN ]: SMSAP-01106: Error occurred while querying the repository: Closed Connection java.sql.SQLException: Closed Connection.`
Refreshing an opened repository by pressing F5	A repository exception is displayed in the GUI and also logs a NullPointerException in the `sm_gui.log` file.
Refreshing the host server	A NullPointerException is logged in the `sumo_gui.log` file.
Creating a new profile	A NullPointerException is displayed in the Profile Configuration window.
Refreshing a profile	The following SQL exception is logged in `sm_gui.log: [WARN ]: SMSAP-01106: Error occurred while querying the repository: Closed Connection.`
Accessing a backup	The following error message is logged in `sm_gui.log: Failed to lazily initialize a collection.`
Viewing clone properties	The following error message is logged in `sm_gui.log` and `sumo_gui.log: Failed to lazily initialize a collection.`

Workaround

You must ensure that the repository database is running when you want to access the GUI or want to perform any SnapManager operations.

Unable to create temporary files for the cloned database

When temporary tablespace files of the target database are placed in mount points different from the mount point of the data files, the clone create operation is successful but SnapManager fails to create temporary files for the cloned database.

Workaround

You must perform either of the following:

- Ensure that the target database is laid out so that temporary files are placed in the same mount point as that of the data files.

- Manually create or add temporary files in the cloned database.

Back up of Data Guard Standby database fails

If any archive log location is configured with the service name of the primary database, the back up of Data Guard Standby database fails.

Workaround

In the GUI, you must clear **Specify External Archive Log location** corresponding to the service name of the primary database.

Running multiple parallel operations fails in SnapManager

When you run multiple parallel operations on separate databases that reside on the same storage system, the igroup for LUNs associated with both the databases might get deleted because of one of the operations. Later, if the other operation attempts to use the deleted igroup, SnapManager displays an error message.

For example, if you are running the backup delete and backup create operations on different databases almost at the same time, the backup create operation fails. The following sequential steps show what happens when you run backup delete and backup create operations on different databases almost at the same time.

Steps

1. Run the `backup delete` command.

2. Run the `backup create` command.

3. The `backup create` command identifies the already existing igroup and uses the same igroup for mapping the LUN.

4. The `backup delete` command deletes the backup LUN, which was mapped to the same igroup.

5. The `backup delete` command then deletes the igroup because there are no LUNs associated with the igroup.

6. The `backup create` command creates the backup and tries to map to the igroup that does not exist, and therefore the operation fails.

What to do

You must create igroup for each storage system used by the database by using the following command: **sdcli igroup create**

Unable to restore RAC database from one of the RAC nodes where the profile was not created

In an Oracle RAC environment where both nodes belong to the same cluster, if you attempt a restore operation from a node which is different from the node where the backup was created, the restore operation fails.

For example, if you create a backup in Node A and try to restore from Node B, the restore operation fails.

What to do

Before performing restore operation from node B, perform the following in node B:

Steps

1. Add the repository.

2. Sync the profile by running the command `smo profile sync`.

3. Set the credential for the profile to be used for restore operation by running the command `smo credential set`.

4. Update the profile to add the new host name and the corresponding SID by running the command `smo profile update`.

Where to go for more information

You can find information about the basic tasks involved in installing and using SnapManager.

Document	Description
SnapManager description page	This page provides information about SnapManager, pointers to online documentation, and a link to the SnapManager download page, from which you can download the software.
Data ONTAP SAN Configuration Guide for 7-Mode	This document is available at NetApp Support. It is a dynamic, online document that contains the most up-to-date information about the requirements for setting up a system in a SAN environment. It provides the current details about storage systems and host platforms, cabling issues, switch issues, and configurations.
SnapManager and SnapDrive Compatibility Matrix	This document is available in the Interoperability section at Interoperability Matrix Tool. It is a dynamic, online document that contains the most up-to-date information specific to SnapManager and its platform requirements.
SnapManager Release Notes	This document comes with SnapManager. You can also download a copy from NetApp Support. It contains any last-minute information that you need to get the configuration up and running smoothly.
NetApp host attach and support kits documentation	NetApp Support.

Document	Description
Host operating system and database information	These documents provide information about your host operating system and database software.

Error message classifications

You can determine the cause of an error if you know the message classifications.

The following table provides information about the numerical ranges for the different types of messages you might see with SnapManager:

Group	Range	Usage
ENVIRONMENT	1000-1999	Used to log the state or issues with the operating environment of SnapManager. This group includes messages about the systems that SnapManager interacts with, such as the host, storage system, database, and so on.
BACKUP	2000-2999	Associated with the database backup process.
RESTORE	3000-3999	Associated with the database restore process.
CLONE	4000-4999	Associated with the database clone process.
PROFILE	5000-5999	Used for managing profiles.
MANAGE	6000-6999	Used for managing backups.
VIRTUAL DATABASE INTERFACE	7000-7999	Associated with the virtual database interface.
VIRTUAL STORAGE INTERFACE	8000-8999	Associated with the virtual storage interface.
REPOSITORY	9000-9999	Associated with the Repository interface.

Group	Range	Usage
METRICS	10000-10999	Associated with the size of the database backup, elapsed time to perform the backup, time to restore the database, the number of times a database has been cloned, and so on.
VIRTUAL HOST INTERFACE	11000-11999	Associated with the virtual host interface. This is the interface to the host operating system.
EXECUTION	12000-12999	Associated with the execution package, including spawning and processing operating system calls.
PROCESS	13000-13999	Associated with the process component of SnapManager.
UTILITIES	14000-14999	Associated with SnapManager utilities, global context, and so on.
DUMP/DIAGNOSTICS	15000-15999	Associated with dump or diagnostic operations.
HELP	16000-16999	Associated with help.
SERVER	17000-17999	Used in the SnapManager server administration.
API	18000-18999	Associated with the API.
BACKINT	19000-19999	Associated with BACKINT.
AUTH	20000-20999	Associated with the authorization of credentials.

Error messages

You can find information about the error messages associated with different SnapManager operations.

Most common error messages

The following table lists some of the most common and important errors associated with SnapManager for SAP:

Error message	Explanation	Resolution
`SD-10038: File system is not writable.`	SnapManager process does not have write access to the file system.	You must ensure that the SnapManager process has write access to the file system. After correcting this, you may need to take another snapshot.
`SMSAP-05075: Unable to create Profile. You must configure the DP/XDP relationship properly or choose the correct protection policy per the underlying relationship.`	The underlying volumes are not in a SnapVault or SnapMirror relationship.	You must configure a data protection relationship between the source and destination volumes and initialize the relationship.
`[SMSAP-05503] You have specified the same name to the profile. Specify a different name to rename the profile.`	Profiles with identical names cannot exist in a repository.	Provide a profile name that is not in use.
`SMSAP-05505: Unable to update dataset metadata.`	Dataset might have been deleted or does not exist.	Before updating the dataset metadata, verify that the dataset exists by using the NetApp Management Console.
`SMSAP-05506: You cannot update the profile since there are operation(s) running on the profile. You must wait until the operation(s) complete and then update the profile`	Profile cannot be updated when backup, restore, and cloning operations are in progress.	Update the profile after completion of the current operation.
`SMSAP-05509: Invalid archivelog primary retention duration - Specify a positive integer value.`	Retention duration of archive log backups cannot be negative.	Specify a positive value for the retention duration of archive log backups.
`SMSAP-07463: This backup restore requires the database to be in required state. Failed to bring the database to the required state.`	Database is not in the required state for a backup operation.	Check that the database is in a relevant state before creating a backup copy. The state of the database that is to be restored depends on the type of restore process that you want to perform and the type of files that are to be included.

Error message	Explanation	Resolution
`SMSAP-09315: After performing repository upgrade or update operation, you might not receive the summary notification for notifications set in previous version unless you update the summary notification with the notification host details.`	Notification settings are not configured for the repository after a rolling upgrade.	After a rolling upgrade, update summary notification settings to receive notifications.
`SMSAP-02076: Label name should not contain any special characters other than underscore.`	Label name contains special characters other than the underscore.	The label name must be unique within the profile. The name can contain letters, numbers, an underscore (_), and a hyphen (-) (but cannot start with a hyphen). Ensure that labels do not contain any special characters except the underscore.
`SMSAP-06308: Exception when attempting to start schedule: java.lang.NullPointerException`	The fully qualified domain name (FQDN) of the profile host is configured instead of the system's host name and the FQDN of the profile host cannot be resolved.	Ensure that you use the system's host name and not the FQDN.
`Failure in ExecuteRestoreSteps: ORACLE-10003: Error executing SQL "DROP DISKGROUP;control diskgroup name; INCLUDING CONTENTS" against Oracle database +ASM1: ORA-15039: diskgroup not dropped ORA-15027: active use of diskgroup; "control diskgroup name;" precludes its dismount`	The operation to restore a backup with control files fails to drop the control disk group. This issue occurs if there are stale backed up control files in the control disk group.	Identify the stale backed up control files and manually delete them.

Error message	Explanation	Resolution
`RMAN-06004: ORACLE error from recovery catalog database: ORA-01424: missing or illegal character following the escape character`	Backup create operation failed to remove the backup copy from the catalog when SnapManager is integrated with RMAN.	Check if there are any external scripts used for removing the backups from RMAN. Execute the `CROSSCHECK BACKUP` command in RMAN to update the RMAN repository and the `resync catalog` command to synchronize the control file of the target database with the recovery catalog.
`[DEBUG]: Exception while pruning backup. java.lang.IllegalStateException: [Assertion failed] - this state invariant must be true`	Multiple Snapshot copies are created for a single operation ID.	Delete the Snapshot copies manually and use scripts to delete the entries from the repository.
System time and the time displayed by SnapManager in the log files do not match or not synchronized.	A time zone change is not yet supported by Java 7.	Apply the `tzupdater` patch provided by Oracle.
`DISC-00001: Unable to discover storage: The following identifier does not exist or is not of the expected type: ASM File`	Data or control files or redo logs are multiplexed in an ASM database.	Remove the Oracle multiplexing.
`0001-SMSAP-02016: There may have been external tables in the database not backed up as part of this backup operation (since the database was not OPEN during this backup ALL_EXTERNAL_LOCATIONS could not be queried to determine whether or not external tables exist).`	SnapManager does not backup external tables (for example, tables that are not stored in .dbf files). This issue occurs because the database was not open during the backup, SnapManager cannot determine if any external tables are being used.	There might have been external tables in the database that are not backed up as part of this operation (because the database was not open during the backup).

Error message	Explanation	Resolution
`0002-332 Admin error: Could not check SD.SnapShot.Clone access on volume "volume_name" for user username on Operations Manager server(s) "dfm_server". Reason: Invalid resource specified. Unable to find its ID on Operations Manager server "dfm_server"`	Proper access privileges and roles are not set.	Set access privileges or roles for the users who are trying to execute the command.
`[WARN] FLOW-11011: Operation aborted [ERROR] FLOW-11008: Operation failed: Java heap space.`	There are more number of archive log files in the database than the maximum allowed.	1. Navigate to the SnapManager installation directory. 2. Open the `launch-java` file. 3. Increase the value of the `java -Xmx160m` Java heap space parameter . For example, you can modify the value from the default value of 160m to 200m as `java -Xmx200m`.
`SMSAP-21019: The archive log pruning failed for the destination: "E:\dest" with the reason: "ORACLE-00101: Error executing RMAN command: [DELETE NOPROMPT ARCHIVELOG 'E:\dest']`	Archive log pruning fails in one of the destinations. In such a scenario, SnapManager continues to prune the archive log files from the other destinations. If any files are manually deleted from the active file system, the RMAN fails to prune the archive log files from that destination.	Connect to RMAN from the SnapManager host. Run the RMAN `CROSSCHECK ARCHIVELOG ALL` command and perform the pruning operation on the archive log files again.
`SMSAP-13032: Cannot perform operation: Archive log Prune. Root cause: RMAN Exception: ORACLE-00101: Error executing RMAN command.`	The archive log files are manually deleted from the archive log destinations.	Connect to RMAN from the SnapManager host. Run the RMAN `CROSSCHECK ARCHIVELOG ALL` command and perform the pruning operation on the archive log files again.

Error message	Explanation	Resolution

`Unable to parse shell output: (java.util.regex.Matcher[pattern=Command complete. region=0,18 lastmatch=]) does not match (name:backup_script)`
 	
`Unable to parse shell output: (java.util.regex.Matcher[pattern=Command complete. region=0,25 lastmatch=]) does not match (description:backup script)`
 	
`Unable to parse shell output: (java.util.regex.Matcher[pattern=Command complete. region=0,9 lastmatch=]) does not match (timeout:0)`
	Environment variables are set not set correctly in the pre-task or post-task scripts.	Check if the pre-task or post-task scripts follow the standard SnapManager plug-in structure. For additional information about using the environmental variables in the script, see Operations in task scripts.

`ORA-01450: maximum key length (6398) exceeded.`	When you perform an upgrade from SnapManager 3.2 for SAP to SnapManager 3.3 for SAP, the upgrade operation fails with this error message. This issue might occur because of one of the following reasons: • The block size of the tablespace in which the repository exists is less than 8k. • The `nls_length_semantics` parameter is set to `char`.	You must assign the values to the following parameters: • `block_size`=**8192** • `nls_length`=**byte** After modifying the parameter values, you must restart the database. For more information, see the Knowledge Base article 2017632.

Error messages associated with the database backup process (2000 series)

The following table lists the common errors associated with the database backup process:

Error message	Explanation	Resolution
`SMSAP-02066: You cannot delete or free the archive log backup "data-logs" as the backup is associated with data backup "data-logs".`	The archive log backup is taken along with the data files backup, and you tried to delete the archive log backup.	Use the `-force` option to delete or free the backup.
`SMSAP-02067: You cannot delete, or free the archive log backup "data-logs" as the backup is associated with data backup "data-logs" and is within the assigned retention duration.`	The archive log backup is associated with the database backup and is within the retention period, and you tried to delete the archive log backup.	Use the `-force` option to delete or free the backup.
`SMSAP-07142: Archived Logs excluded due to exclusion pattern <exclusion> pattern.`	You exclude some archive log files during the profile create or backup create operation.	No action is required.
`SMSAP-07155: <count> archived log files do not exist in the active file system. These archived log files will not be included in the backup.`	The archive log files do not exist in the active file system during the profile create or backup create operation. These archived log files are not included in the backup.	No action is required.
`SMSAP-07148: Archived log files are not available.`	No archive log files are created for the current database during the profile create or backup create operation.	No action is required.
`SMSAP-07150: Archived log files are not found.`	All the archive log files are missing from the file system or excluded during the profile create or backup create operation.	No action is required.

Error message	Explanation	Resolution
`SMSAP-13032: Cannot perform operation: Backup Create. Root cause: ORACLE-20001: Error trying to change state to OPEN for database instance dfcln1: ORACLE-20004: Expecting to be able to open the database without the RESETLOGS option, but oracle is reporting that the database needs to be opened with the RESETLOGS option. To keep from unexpectedly resetting the logs, the process will not continue. Please ensure that the database can be opened without the RESETLOGS option and try again.`	You try to back up the cloned database that was created with the `-no-resetlogs` option. The cloned database is not a complete database.However, you can perform SnapManager operations such as creating profiles and backups, and so on with the cloned database, but the SnapManager operations fail because the cloned database is not configured as a complete database.	Recover the cloned database or convert the database into a Data Guard Standby database.

Error messages associated with the restore process (3000 series)

The following table shows the common errors associated with the restore process:

Error message	Explanation	Resolution
`SMSAP-03031:Restore specification is required to restore backup <variable> because the storage resources for the backup has already been freed.`	You attempted to restore a backup that has its storage resources freed without specifying a restore specification.	Specify a restore specification.
`SMSAP-03032:Restore specification must contain mappings for the files to restore because the storage resources for the backup has already been freed. The files that need mappings are: <variable> from Snapshots: <variable>`	You attempted to restore a backup that has its storage resources freed along with a restore specification that does not contain mapping for all the files to be restored.	Correct the restore specification file so that the mappings match the files to be restored.

Error message	Explanation	Resolution
`ORACLE-30028: Unable to dump log file <filename>. The file may be missing/inaccessible/corrupted. This log file will not be used for recovery.`	The online redo log files or archive log files cannot be used for recovery.This error occurs due to following reasons: • The online redo log files or archived log files mentioned in the error message do not have sufficient change numbers to apply for recovery. This occurs when the database is online without any transactions. The redo log or archived log files do not have any valid change numbers that can be applied for recovery. • The online redo log file or archived log file mentioned in the error message does not have sufficient access privileges for Oracle. • The online redo log file or archived log file mentioned in the error message is corrupted and cannot be read by Oracle. • The online redo log file or archived log file mentioned in the error message is not found in the path mentioned.	If the file mentioned in the error message is an archived log file and if you have manually provided for recovery, ensure that the file has full access permissions to Oracle.Even if the file has full permissions, and the message continues, the archive log file does not have any change numbers to be applied for recovery, and this message can be ignored.

Error messages associated with the clone process (4000 series)

The following table shows the common errors associated with the clone process:

Error message	Explanation	Resolution
`SMSAP-04133: Dump destination must not exist`	You are using SnapManager to create new clones; however, the dump destinations to be used by the new clone already exist. SnapManager cannot create a clone if the dump destinations exist.	Remove or rename the old dump destinations before you create a clone.

Error message	Explanation	Resolution
`SMSAP-13032:Cannot perform operation: Clone Create. Root cause: ORACLE-00001: Error executing SQL: [ALTER DATABASE OPEN RESETLOGS;]. The command returned: ORA-38856: cannot mark instance UNNAMED_INSTANCE_2 (redo thread 2) as enabled.`	The clone creation fails when you create the clone from the standby database using the following setup: • The standby is created by using RMAN for taking the data files backup.	Add the `_no_recovery_through_reset logs=TRUE` parameter in the clone specification file before creating the clone. See Oracle documentation (ID 334899.1) for additional information. Ensure that you have your Oracle metalink user name and password.
	You did not provide a value for a parameter in the clone specification file.	You must either provide a value for the parameter or delete that parameter if it is not required from the clone specification file.

Error messages associated with the managing profile process (5000 series)

The following table shows the common errors associated with the clone process:

Error message	Explanation	Resolution
`SMSAP-20600: Profile "profile1" not found in repository "repo_name". Please run "profile sync" to update your profile-to-repository mappings.`	The dump operation cannot be performed when profile creation fails.	Use `smsap system dump`.

Error messages associated with freeing backup resources (backups 6000 series)

The following table shows the common errors associated with backup tasks:

Error message	Explanation	Resolution
`SMSAP-06030: Cannot remove backup because it is in use: <variable>`	You attempted to perform the backup free operation using commands, when the backup is mounted or is marked to be retained on an unlimited basis.	Unmount the backup or change the unlimited retention policy. If clones exist, delete them.
`SMSAP-06045: Cannot free backup <variable> because the storage resources for the backup have already been freed`	You attempted to perform the backup free operation using commands, when the backup has been already freed.	You cannot free the backup if it is already freed.

Error message	Explanation	Resolution
`SMSAP-06047: Only successful backups can be freed. The status of backup <ID> is <status>.`	You attempted to perform the backup free operation using commands, when the backup status is not successful.	Try again after a successful backup.
`SMSAP-13082: Cannot perform operation <variable> on backup <ID> because the storage resources for the backup have been freed.`	Using commands, you attempted to mount a backup that has its storage resources freed.	You cannot mount, clone, verify, or perform backint restore on a backup that has its storage resources freed.

Error messages associated with the rolling upgrade process (9000 series)

The following table shows the common errors associated with the rolling upgrade process:

Error message	Explanation	Resolution
`SMSAP-09234:Following hosts does not exist in the old repository. <hostnames>.`	You tried to perform rolling upgrade of a host, which does not exist in the previous repository version.	Check whether the host exists in the previous repository using the `repository show-repository` command from the earlier version of the SnapManager CLI.
`SMSAP-09255:Following hosts does not exist in the new repository. <hostnames>.`	You tried to perform roll back of a host, which does not exist in the new repository version.	Check whether the host exists in the new repository using the `repository show-repository` command from the later version of the SnapManager CLI.
`SMSAP-09256:Rollback not supported, since there exists new profiles <profilenames>.for the specified hosts <hostnames>.`	You tried to roll back a host that contains new profiles existing in the repository. However, these profiles did not exist in the host of the earlier SnapManager version.	Delete new profiles in the later or upgraded version of SnapManager before the rollback.
`SMSAP-09257:Rollback not supported, since the backups <backupid> are mounted in the new hosts.`	You tried to roll back a later version of the SnapManager host that has mounted backups. These backups are not mounted in the earlier version of the SnapManager host.	Unmount the backups in the later version of the SnapManager host, and then perform the rollback.

Error message	Explanation	Resolution
`SMSAP-09258:Rollback not supported, since the backups <backupid> are unmounted in the new hosts.`	You tried to roll back a later version of the SnapManager host that has backups that are being unmounted.	Mount the backups in the later version of the SnapManager host, and then perform the rollback.
`SMSAP-09298:Cannot update this repository since it already has other hosts in the higher version. Please perform rollingupgrade for all hosts instead.`	You performed a rolling upgrade on a single host and then updated the repository for that host.	Perform a rolling upgrade on all the hosts.
`SMSAP-09297: Error occurred while enabling constraints. The repository might be in inconsistent state. It is recommended to restore the backup of repository you have taken before the current operation.`	You attempted to perform a rolling upgrade or rollback operation if the repository database is left in an inconsistent state.	Restore the repository that you backed up earlier.

Execution of operations (12,000 series)

The following table shows the common errors associated with operations:

Error message	Explanation	Resolution
`SMSAP-12347 [ERROR]: SnapManager server not running on host <host> and port <port>. Please run this command on a host running the SnapManager server.`	While setting up the profile, you entered information about the host and port. However, SnapManager cannot perform these operations because the SnapManager server is not running on the specified host and port.	Enter the command on a host running the SnapManager server. You can check the port with the `lsnrctl status` command and see the port on which the database is running. Change the port in the backup command, if needed.

Execution of process components (13,000 series)

The following table shows the common errors associated with the process component of SnapManager:

Error message	Explanation	Resolution
`SMSAP-13083: Snapname pattern with value "x" contains characters other than letters, numbers, underscore, dash, and curly braces.`	When creating a profile, you customized the Snapname pattern; however, you included special characters that are not allowed.	Remove special characters other than letters, numbers, underscore, dash, and braces.
`SMSAP-13084: Snapname pattern with value "x" does not contain the same number of left and right braces.`	When you were creating a profile, you customized the Snapname pattern; however, the left and right braces do not match.	Enter matching opening and closing brackets in the Snapname pattern.
`SMSAP-13085: Snapname pattern with value "x" contains an invalid variable name of "y".`	When you were creating a profile, you customized the Snapname pattern; however, you included a variable that is not allowed.	Remove the offending variable. To see a list of acceptable variables, see Snapshot copy naming.
`SMSAP-13086 Snapname pattern with value "x" must contain variable "smid".`	When you were creating a profile, you customized the Snapname pattern; however, you omitted the required `smid` variable.	Insert the required `smid` variable.

Error messages associated with SnapManager Utilities (14,000 series)

The following table shows the common errors associated with SnapManager utilities:

Error message	Explanation	Resolution
`SMSAP-14501: Mail ID cannot be blank.`	You did not enter the email address.	Enter a valid email address.
`SMSAP-14502: Mail subject cannot be blank.`	You did not enter the email subject.	Enter the appropriate email subject.
`SMSAP-14506: Mail server field cannot be blank.`	You did not enter the email server host name or IP address.	Enter the valid mail server host name or IP address.
`SMSAP-14507: Mail Port field cannot be blank.`	You did not enter the email port number.	Enter the email server port number.
`SMSAP-14508: From Mail ID cannot be blank.`	You did not enter the sender's email address.	Enter a valid sender's email address.

Error message	Explanation	Resolution
`SMSAP-14509: Username cannot be blank.`	You enabled authentication and did not provide the user name.	Enter the email authentication user name.
`SMSAP-14510: Password cannot be blank. Please enter the password.`	You enabled authentication and did not provide the password.	Enter the email authentication password.
`SMSAP-14550: Email status <success/failure>.`	The port number, mail server, or receiver's email address is invalid.	Provide proper values during email configuration.
`SMSAP-14559: Sending email notification failed: <error>.`	This could be due to invalid port number, invalid mail server, or invalid receiver's mail address.	Provide proper values during email configuration.
`SMSAP-14560: Notification failed: Notification configuration is not available.`	Notification sending failed, because notification configuration is not available.	Add notification configuration.
`SMSAP-14565: Invalid time format. Please enter time format in HH:mm.`	You have entered time in an incorrect format.	Enter the time in the format: hh:mm.
`SMSAP-14566: Invalid date value. Valid date range is 1-31.`	The date configured is incorrect.	Date should be in the range from 1 through 31.
`SMSAP-14567: Invalid day value. Valid day range is 1-7.`	The day configured is incorrect.	Enter the day range from 1 through 7.
`SMSAP-14569: Server failed to start Summary Notification schedule.`	The SnapManager server got shut down due to unknown reasons.	Start the SnapManager server.
`SMSAP-14570: Summary Notification not available.`	You have not configured summary notification.	Configure the summary notification.
`SMSAP-14571: Both profile and summary notification cannot be enable.`	You have selected both the profile and summary notification options.	Enable either the profile notification or summary notification.

Error message	Explanation	Resolution
`SMSAP-14572: Provide success or failure option for notification.`	You have not enabled the success or failure options.	You must select either success or failure option or both.