



Administration for UNIX

SnapManager for SAP

NetApp

November 04, 2025

This PDF was generated from <https://docs.netapp.com/us-en/snapmanager-sap/unix-administration/concept-create-backups-using-snapshot-copies.html> on November 04, 2025. Always check docs.netapp.com for the latest.

Table of Contents

Administration for UNIX	1
Product overview	1
SnapManager highlights	1
Create backups using Snapshot copies	2
Why you should prune archive log files	2
Archive log consolidation	2
Full or partial restoration of databases	2
Verify backup status	3
Database backup clones	3
Track details and produce reports	3
What repositories are	4
What profiles are	4
What SnapManager operation states are	6
How SnapManager maintains security	7
Recommended general database layouts and storage configurations	8
Limitations when working with SnapManager	12
Configuring SnapManager	18
SnapManager configuration parameters	18
Security and credential management	26
What user authentication is	26
About role-based access control	27
Clear user credentials for all hosts, repositories, and profiles	28
Delete credentials for individual resources	28
Managing profiles for efficient backups	28
Tasks related to profiles	28
About profiles and authentication	29
Snapshot copy naming	30
Backing up databases	31
What SnapManager database backups are	33
What full and partial backups are	33
About control file and archive log file handling	39
What database backup scheduling is	40
Create database backups	43
What AutoSupport is	43
Change the backup retention policy	44
Scheduling database backups	44
Restoring database backups	44
What database restore is	45
Restore files from an alternate location	62
Cloning database backup	64
What Cloning is	64
Cloning methods	66
Create clone specifications	66

Considerations for cloning a database to an alternate host	66
Splitting a clone	67
Introduction to data protection in SnapManager	67
What protection policies are	68
What protection states are	69
What resource pools are	69
About different protection policies	70
Configuring and enabling policy-driven data protection	71
How SnapManager retains backups on the local storage	72
Considerations for performing data protection	75
Restoring protected backups from secondary storage	76
SnapManager for SAP uses Protection Manager to protect a database backup	78
Details of the target database	78
Primary and secondary storage configuration and topology	78
Backup schedule and retention strategy	82
Workflow summary for local and secondary database backup	83
Protected backup configuration and execution	84
Database restore from backup	84
Performing management operations	85
Configuring an email notification	85
Configure email notification for a new profile	87
Configure email notification for an existing profile	87
Creating task specification file and scripts for SnapManager operations	87
Create pretask, post-task, and policy scripts	88
Updating storage system name and target database host name associated with a profile	96
Maintaining history of SnapManager operations	96
Using BR*Tools with SnapManager for SAP	97
What BR*Tools are	97
Disabling client access to Snapshot copies	100
Profile usage for BR*Tools backups	100
Database backups created using BRBACKUP and BRARCHIVE	103
Database restore by using BRRESTORE and BRRECOVER	104
Backing up and restoring files using BR*Tools	105
Restoring a backup to a different host	105
Restoring a backup from a secondary storage to a different location	105
SnapManager for SAP command reference	106
The backint register-sld command	106
The smsap_server restart command	107
The smsap_server start command	108
The smsap_server status command	109
The smsap_server stop command	110
The smsap backup create command	110
The smsap backup delete command	114
The smsap backup free command	116
The smsap backup list command	117

The smsap backup mount command	119
The smsap backup restore command	121
The smsap backup show command	126
The smsap backup unmount command	127
The smsap backup update command	129
The smsap backup verify command	130
The smsap clone create command	132
The smsap clone delete command	135
The smsap clone list command	137
The smsap clone show command	138
The smsap clone template command	140
The smsap clone update command	141
The smsap clone split-delete command	142
The smsap clone split-estimate command	143
The smsap clone split command	144
The smsap clone split-result command	149
The smsap clone split-stop command	150
The smsap clone split-status command	151
The smsap clone detach command	151
The smsap cmdfile command	152
The smsap credential clear command	153
The smsap credential delete command	154
The smsap credential list command	156
The smsap credential set command	157
The smsap history list command	159
The smsap history operation-show command	161
The smsap history purge command	161
The smsap history remove command	163
The smsap history set command	164
The smsap history show command	166
The smsap help command	167
The smsap notification remove-summary-notification command	168
The smsap notification update-summary-notification command	169
The smsap notification set command	171
The smsap operation dump command	173
The smsap operation list command	174
The smsap operation show command	175
The smsap password reset command	177
The smsap plugin check command	178
The smsap profile create command	179
The smsap profile delete command	184
The smsap profile destroy command	185
The smsap profile dump command	186
The smsap profile list command	187
The smsap profile show command	188

The smsap profile sync command	190
The smsap profile update command	192
The smsap profile verify command	197
The smsap protection-policy command	200
The smsap repository create command	201
The smsap repository delete command	202
The smsap repository rollback command	204
The smsap repository rollingupgrade command	205
The smsap repository show command	207
The smsap repository update command	208
The smsap schedule create command	210
The smsap schedule delete command	215
The smsap schedule list command	215
The smsap schedule resume command	215
The smsap schedule suspend command	216
The smsap schedule update command	216
The smsap storage list command	218
The smsap storage rename command	218
The smsap system dump command	219
The smsap system verify command	220
The smsap version command	221
Troubleshooting SnapManager	221
Dump files	227
Troubleshooting clone issues	232
Troubleshooting graphical user interface issues	235
Troubleshooting SnapDrive issues	240
Troubleshooting storage system renaming issue	241
Troubleshooting known issues	243
Mounting a FlexClone volume fails in NFS environment	249
Running multiple parallel operations fails in SnapManager	249
Where to go for more information	250
Error message classifications	251
Error messages	252
Most common error messages	252
Error messages associated with the database backup process (2000 series)	259
Data protection errors	260
Error messages associated with the restore process (3000 series)	264
Error messages associated with the clone process (4000 series)	266
Error messages associated with the managing profile process (5000 series)	267
Error messages associated with freeing backup resources (backups 6000 series)	267
Virtual storage interface errors (virtual storage interface 8000 series)	268
Error messages associated with the rolling upgrade process (9000 series)	269
Execution of operations (12,000 series)	270
Execution of process components (13,000 series)	270
Error messages associated with SnapManager Utilities (14,000 series)	272

Administration for UNIX

Product overview

SnapManager for SAP automates and simplifies the complex, manual, and time-consuming processes associated with the backup, recovery, and cloning of databases. You can use SnapManager with ONTAP SnapMirror technology to create copies of backups on another volume and with ONTAP SnapVault technology to archive backups efficiently to disk.

SnapManager provides the tools required, such as OnCommand Unified Manager and integration with SAP's BR* Tools, to perform policy-driven data management, schedule and create regular database backups, and restore data from these backups in the event of data loss or disaster.

SnapManager also integrates with native Oracle technologies, such as Oracle Real Application Clusters (Oracle RAC) and Oracle Recovery Manager (RMAN) to preserve backup information. These backups can be used later in block-level restore or tablespace point-in-time recovery operations.

SnapManager highlights

SnapManager features seamless integration with databases on the UNIX host and with Snapshot, SnapRestore, and FlexClone technologies on the back end. It offers an easy-to-use user interface (UI) as well as command-line interface (CLI) for administrative functions.

SnapManager enables you to perform the following database operations and manage data efficiently:

- Creating space-efficient backups on primary or secondary storage
SnapManager enables you to back up the data files and archive log files separately.
- Scheduling backups
- Restoring full or partial databases by using a file-based or volume-based restore operation
- Recovering databases by discovering, mounting, and applying archive log files from backups
- Pruning archive log files from archive log destinations when creating backups of only the archive logs
- Retaining a minimum number of archive log backups automatically by retaining only the backups that contain unique archive log files
- Tracking operation details and generating reports
- Verifying backups to ensure that backups are in a valid block format and that none of the backed-up files are corrupted
- Maintaining a history of operations performed on the database profile

A profile contains information about the database to be managed by SnapManager.

- Creating space-efficient clones of backups on primary or secondary storage

SnapManager enables you to split a clone of a database.

Create backups using Snapshot copies

SnapManager enables you to create backups on the primary (local) storage and also on the secondary (remote) storage using protection policies or 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.



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 processor a faster, volume-based restore process. Database administrators can select the process they want to use or let SnapManager decide which process is appropriate.

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 or secondary storage. 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. You can also identify the protected backups and determine whether backups are in process or scheduled to occur.

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.

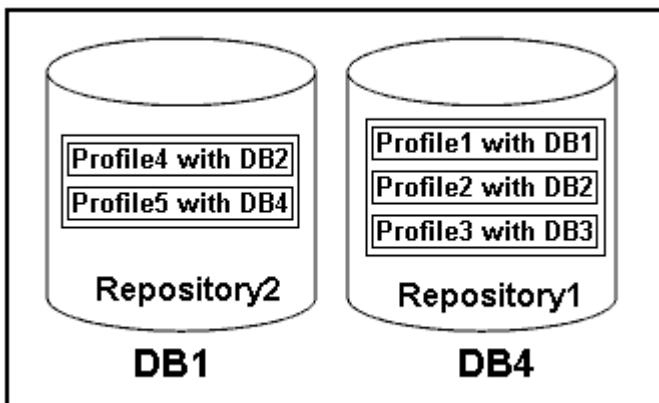
What profiles are

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 and protected based on the existing protection 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 and protection policies for the archive log backups.

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	<p>The operation failed. SnapManager has automatically executed the abort process and cleaned the operation.</p> <p> You can split the clone that is created. When you stop the clone split operation you started and the operation is stopped successfully, the clone split operation state displays as failed.</p>

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 UNIX 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 and role-based access control (RBAC). RBAC enables database administrators to restrict the operations that SnapManager can perform against the volumes and LUNs that hold the data files in a database.

Database administrators enable RBAC for SnapManager by using SnapDrive. Database administrators then assign permissions to SnapManager roles and assign these roles to the users in the Operations Manager graphical user interface (GUI) or command-line interface (CLI). RBAC permission checks happen in the DataFabric Manager server.

In addition to role-based access, 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.



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

Permission type	User1	User2
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.

Recommended general database layouts and storage configurations

Knowing the recommended general database layouts and storage configurations can help you avoid issues related to disk groups, file types, and tablespaces.

- Do not include files from more than one type of SAN file system or volume manager 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.
- Include the database system identifier in the `oratab` file.

Include an entry in the `oratab` file for each database to be managed. SnapManager relies on the `oratab` file to determine which Oracle home to use.

If you want to leverage the new volume-based restore or full disk group restore, consider the following guidelines related to file systems and disk groups:

- A disk group containing data files cannot contain other types of files.
- The logical unit number (LUN) for the data file disk group must be the only object in the storage volume.

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.

The control file in the `sapdata1` file system conflicts with the SnapManager requirements for separating the control files and data files into separate volumes and must be adjusted to allow fast restore capability.



Because BR*Tools backups include the Oracle and SAP profiles located in the `dfs` 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.)

However, in the case of a system that has already been installed, you must move the control file from the file system to allow fast restore using SnapManager. You can do this by creating a new file system in a volume that does not contain data files, moving the control file to that file system, and then creating a symbolic link from the previous file system to the directory for the new file system. It is necessary to stop SAP and the Oracle database when moving the control file to avoid database errors.

Before making any modifications, listing of the files in the sapdata1 directory containing the control file might look like this:

```
hostname:/
# ls -l /oracle/SID/sapdata1/cntrl
-rw-r----- 1 orasid  dba  9388032 Jun 19 01:51 cntrlsID.dbf
```

After the modification, the listing might look like this:

```
hostname:/
# ls -sl /oracle/SID/sapdata1
0 lrwxrwxrwx 1 root root 19 2008-08-06 14:55 cntrl -> /oracle/SID/control
0 -rw-r--r-- 1 root root 0 2008-08-06 14:57 data01.dbf

# ls -sl /oracle/SID/control
0 -rw-r--r-- 1 root root 0 2008-08-06 14:56 cntrlsID.dbf
```

Requirements for using RAC databases with SnapManager

You must know the recommendations for using Real Application Clusters (RAC) databases with SnapManager. The recommendations include port numbers, passwords, and authentication mode.

- In database authentication mode, the listener on each node that interacts with an instance of the RAC database must be configured to use the same port number.

The listener that interacts with the primary database instance must be started prior to initiating a backup.

- In operating system authentication mode, the SnapManager server must be installed and running on each node in the RAC environment.
- The database user password (for example, for a system administrator or a user with the sysdba privilege) must be same for all the Oracle database instances in a RAC environment.

Supported partition devices

You must know the different partition devices that are supported in SnapManager.

The following table provides partition information and how it can be enabled for different operating systems:

Operating system	Single partition	Multiple partition	Non-partition devices	File system or RAW devices
Red Hat Enterprise Linux 5x or Oracle Enterprise Linux 5x	Yes	No	No	ext3*
Red Hat Enterprise Linux 6x or Oracle Enterprise Linux 6x	Yes	No	No	ext3 or ext4*
SUSE Linux Enterprise Server 11	Yes	No	No	ext3*
SUSE Linux Enterprise Server 10	No	No	Yes	ext3***

*

For a non-MPIO environment, enter the following command:

```
sfdisk -uS -f -L -q /dev/ device_name
```

For an MPIO environment, enter the following commands:

- sfdisk -uS -f -L -q /dev/ device_name
- kpartx -a -p p /dev/mapper/ device_name

Not applicable.

For more information on the operating system versions supported, refer to the Interoperability Matrix.

Requirements for using databases with NFS and SnapManager

You must know the requirements for using databases with Network File System (NFS) and SnapManager. The recommendations include running as root, attribute caching, and symbolic links.

- You must run SnapManager as root; SnapManager must be able to access the file systems that contain data files, control files, online redo logs, archive logs, and the database home.

Set either of the following NFS export options to ensure that root can access the file systems:

- `root=host name`
- `rw=host name, anon=0`

- You must disable attribute caching for all the volumes that contain database data files, control files, redo and archive logs, and the database home.

Export the volumes by using the noac (for Solaris and AIX) or actimeo=0 (for Linux) options.

- You must link the database data files from local storage to NFS to support symbolic links at the mount point-level only.

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_ <i>host name</i>	Yes	On
Data files	oradata_ <i>sid</i>	Yes	Off
Temporary data files	oratemp_ <i>sid</i>	Yes	Off
Control files	oracntrl01_ <i>sid</i> (Multiplexed) oracntrl02_ <i>sid</i> (Multiplexed)	Yes	Off
Redo logs	oralog01_ <i>sid</i> (Multiplexed) oralog02_ <i>sid</i> (Multiplexed)	Yes	Off
Archive logs	oraarch_ <i>sid</i>	Yes	Off

Real Application Clusters (RAC) databases

File types	Volume names	Dedicated volume for file types	Automatic Snapshot copies
Oracle binaries	orabin_ <i>host name</i>	Yes	On
Data files	oradata_ <i>dbname</i>	Yes	Off
Temporary data files	oratemp_ <i>dbname</i>	Yes	Off

File types	Volume names	Dedicated volume for file types	Automatic Snapshot copies
Control files	oracntrl01_ <i>dbname</i> (Multiplexed) oracntrl02_ <i>dbname</i> (Multiplexed)	Yes	Off
Redo logs	oralog01_ <i>dbname</i> (Multiplexed) oralog02_ <i>dbname</i> (Multiplexed)	Yes	Off
Archive logs	oraarch_ <i>dbname</i>	Yes	Off
Cluster files	oracrs_ <i>clustername</i>	Yes	On

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.
- In Red Hat Enterprise Linux (RHEL) and Oracle Enterprise Linux 4.7, 5.0, 5.1, 5.2, and 5.3, the ext3 file system is not supported when deploying Oracle over raw devices by using dynamic multipathing (DMP) in a multipath network I/O (MPIO) environment.

This issue is noticed in SnapManager only when using SnapDrive 4.1 for UNIX or earlier versions.

- SnapManager on RHEL does not support partitioning of disks using the **parted** utility.

This is an issue with the RHEL **parted** utility.

- In a RAC configuration, when a profile name is updated from RAC node A, the schedule file for the profile is updated only for RAC node A.

The schedule file for the same profile on RAC node B is not updated and contains the earlier schedule information. When a scheduled backup is triggered from node B, the scheduled backup operation fails because node B contains the earlier schedule file. However, the scheduled backup operation succeeds from node A, on which the profile is renamed. You can restart the SnapManager server so that you receive the latest schedule file for the profile on node B.

- 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

- SnapDrive for UNIX supports more than one type of file system and volume manager on certain platforms.

The file system and volume manager used for database files must be specified in the SnapDrive configuration file as the default file system and volume manager.

- 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.

- All LUNs within a volume should reside at the volume level or inside qtrees, but not both.

This is because if the data is residing on the qtrees and you mount the volume, then the data inside the qtrees is not protected.

- 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.
- SnapManager integrated with Protection Manager supports the backup of multiple volumes in primary storage to a single volume in secondary storage for SnapVault and qtree SnapMirror.

Dynamic secondary volume sizing is not supported. The Provisioning Manager and Protection Manager Administration Guide For Use with DataFabric Manager Server 3.8 has for more information about this.

- SnapManager does not support vaulting of backups using the post-processing script.
- 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.
- After upgrading to SnapManager 3.4 or later, any backups scheduled with post-processing scripts using SnapManager 3.3.1 cannot be updated.

You must delete the existing schedule and create a new schedule.

Limitations related to restore operations

- When you use an indirect method for performing a restore operation and the archive log files required for recovery are available only in backups from the secondary storage system, SnapManager fails to recover the database.

This is because SnapManager cannot mount the backup of archive log files from the secondary storage system.

- When SnapManager performs a volume restore operation, the archive log backup copies that are made after the corresponding backup is restored are not purged.

When the data files and archive log file destination exist on the same volume, the data files can be restored

through a volume restore operation if there are no archive log files available in the archive log file destination. In such a scenario, the archive log Snapshot copies that are created after the backup of the data files are lost.

You should not delete all of the archive log files from the archive log destination.

Limitations related to clone operations

- You cannot view any numerical values between 0 and 100 for the progress of the clone split operation because of the speed with which the inodes are discovered and processed by the storage system containing the flexible volume.
- SnapManager does not support receiving emails only for the successful clone split operations.
- SnapManager only supports splitting a FlexClone.
- The cloning of online database backup of the RAC database that uses external archive log file location fails because of failure in recovery.

The cloning fails because Oracle fails to find and apply the archive log files for recovery from the external archive log location. This is an Oracle limitation. For more information, see the Oracle Bug ID: 13528007. Oracle does not apply archive log from the non-default location at the [Oracle support site](#). You must have a valid Oracle metalink user name and password.

- 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 UNIX client.
- 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.

Limitations related to SnapMirror and SnapVault

- The SnapVault post-processing script is not supported if you are using Data ONTAP operating in 7-Mode.
- If you are using ONTAP, you cannot perform volume-based SnapRestore (VBSR) on the backups that were created in the volumes that have SnapMirror relationships established.

This is because of an ONTAP limitation, which does not allow you to break the relationship when doing a VBSR. However, you can perform a VBSR on the last or most recently created backup only when the volumes have SnapVault relationships established.

- If you are using Data ONTAP operating in 7-Mode and want to perform a VBSR on the backups that were created in the volumes that have SnapMirror relationships established, you can set the `override-vbsr-snapmirror-check` option to `ON` in SnapDrive for UNIX.

The SnapDrive documentation has more information about this.

- 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 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:

- Data protection capabilities if SnapManager is integrated with OnCommand Unified Manager
- 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.



Identify the different versions of Oracle databases supported by referring to the Interoperability Matrix.

Related information

[Interoperability Matrix](#)

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 12c, then you must upgrade to SnapManager 3.3.1 or later.



Oracle database 12c 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 12c databases along with the other supported databases.

Volume management restrictions

SnapManager has certain volume management restrictions that might affect your environment.

You can have multiple disk groups for a database; however, the following limitations apply to all disk groups for a given database:

- Disk groups for the database can be managed by only one volume manager.
- A Linux environment without logical volume management requires a partition.

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
<ul style="list-style-type: none"> • <code>retain.hourly.count</code> • <code>retain.hourly.duration</code> • <code>retain.monthly.count</code> • <code>retain.monthly.duration</code> 	<p>These parameters set the retention policy when you create a profile. For example, you can assign the following values:</p> <pre>retain.hourly.count = 12 retain.hourly.duration = 2 retain.monthly.count = 2 retain.monthly.duration = 6</pre>
<code>restore.secondaryAccessPolicy</code>	<p>This parameter defines how SnapManager can access data on secondary storage when it cannot be restored directly by using Protection Manager. The different ways to access the data on secondary storage are as follows:</p> <ul style="list-style-type: none"> • Direct (default) <p>When <code>restore.secondaryAccessPolicy</code> is set to <code>direct</code>, SnapManager clones the data on secondary storage, mounts the cloned data from the secondary storage to the host, and then copies data out of the clone into the active environment.</p> <ul style="list-style-type: none"> • Indirect <p>If you assign <code>indirect</code> to <code>restore.secondaryAccessPolicy</code>, SnapManager copies data to a temporary volume on primary storage, mounts data from the temporary volume to the host, and then copies data out of the temporary volume into the active environment.</p> <p>The indirect method must be used only if the host does not have direct access to the secondary storage system. This method takes twice as long as the direct method because two copies of the data are made.</p> <p> In a Storage Area Network (SAN) with Network File System (NFS) as the protocol, SnapManager does not need to connect directly to secondary storage to perform a restore.</p>
<code>restore.temporaryVolumeName</code>	<p>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:</p> <pre>restore.temporaryVolumeName = smsap_temp_volume</pre>

Parameters	Description
retain.alwaysFreeExpiredBackups	<p>This parameter allows SnapManager to free backups when they expire and when a fast restore is performed, even if data protection is not configured. This parameter frees the protected backups that expire and deletes the unprotected backups that expire. The possible values that you can assign are as follows:</p> <ul style="list-style-type: none"> • True <p>If you assign true to <code>retain.alwaysFreeExpiredBackups</code>, SnapManager frees the expired backups regardless of whether the backups are protected.</p> <p>The backups are deleted either when they are not protected or if the protected copies on secondary storage have also expired.</p> <ul style="list-style-type: none"> • False <p>If you assign false to <code>retain.alwaysFreeExpiredBackups</code>, SnapManager frees the expired backups that are protected.</p>
host.credentials.persist	<p>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 <code>host.credentials.persist</code>. SnapManager encrypts and saves the host credentials.</p>
restorePlanMaxFilesDisplayed	<p>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:</p> <ul style="list-style-type: none"> • <code>restorePlanMaxFilesDisplayed = 30</code> <p> If you specify an invalid value, the default number of files are displayed.</p>
snapshot.list.timeout.min	<p>This parameter enables you to define the time in minutes for which SnapManager must wait for the <code>snap list</code> 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:</p> <ul style="list-style-type: none"> • <code>snapshot.list.timeout.min = 40</code> <p> If you specify an invalid value, the default value is used.</p> <p>For any SnapManager operation, if the <code>snap list</code> command execution time exceeds the value assigned to <code>snapshot.list.timeout.min</code>, the operation fails with a timeout error message.</p>

Parameters	Description
pruneIfFileExistsInOtherDestination	<p>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:</p> <ul style="list-style-type: none"> When you want to prune the archive log files from a specified destination, you must assign false to <code>pruneIfFileExistsInOtherDestination</code>. When you want to prune the archive log files from an external destination, you must assign true to <code>pruneIfFileExistsInOtherDestination</code>.
prune.archivelogs.backedup.from.otherdestination	<p>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:</p> <ul style="list-style-type: none"> 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 <code>-prune-dest</code>, you must assign false to <code>prune.archivelogs.backedup.from.otherdestination</code>. 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 <code>prune.archivelogs.backedup.from.otherdestination</code>.
maximum.archivelog.files.toprune.atATime	<p>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:</p> <pre>maximum.archivelog.files.toprune.atATime = 998</pre> <p> The value that can be assigned to <code>maximum.archivelog.files.toprune.atATime</code> must be less than 1000.</p>
archivelogs.consolidate	<p>This parameter allows SnapManager to free the duplicate archive log backups if you assign true to <code>archivelogs.consolidate</code>.</p>
suffix.backup.label.with.logs	<p>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.</p> <p>For example, when you assign logs to <code>suffix.backup.label.with.logs</code>, <code>_logs</code> is added as a suffix to the archive log backup label. The archive log backup label would then be <code>arch_logs</code>.</p>

Parameters	Description
backup.archiveLogs.beyond.missingfiles	<p>This parameter allows SnapManager to include the missing archive log files in the backup.</p> <p>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 <code>backup.archiveLogs.beyond.missingfiles</code>.</p> <p>You can assign false to ignore the missing archive log files.</p>
srvctl.timeout	<p>This parameter enables you to define the timeout value for the <code>srvctl</code> command.</p> <p> The Server Control (SRVCTL) is a utility to manage RAC instances.</p> <p>When SnapManager takes more time to execute the <code>srvctl</code> command than the timeout value, the SnapManager operation fails with this error message: <code>Error: Timeout occurred while executing command: srvctl status</code>.</p>
snapshot.restore.storageNameCheck	<p>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 <code>snapshot.restore.storageNameCheck=true</code>.</p>
services.common.disableAbort	<p>This parameter disables cleanup upon failure of long-running operations. You can set <code>services.common.disableAbort=true</code>. 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 <code>services.common.disableAbort=true</code>, the clone will not be deleted. You can fix the Oracle issue and restart the clone operation from the point where it failed.</p>

Parameters	Description
<ul style="list-style-type: none"> • <code>backup.sleep.dnfs.layout</code> • <code>backup.sleep.dnfs.secs</code> 	<p>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.</p> <p>To enable the sleep mechanism, ensure that <code>backup.sleep.dnfs.layout=true</code>. The default value is <code>true</code>.</p> <p>When you enable the sleep mechanism, you must assign the sleep time to <code>backup.sleep.dnfs.secs</code>. The sleep time assigned is in seconds and the value depends upon your environment. The default value is 5 seconds.</p> <p>For example:</p> <ul style="list-style-type: none"> • <code>backup.sleep.dnfs.layout=true</code> • <code>backup.sleep.dnfs.secs=2</code>
<ul style="list-style-type: none"> • <code>override.default.backup.pattern</code> • <code>new.default.backup.pattern</code> 	<p>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.</p> <p>To enable customization of the backup label, ensure that the value of <code>override.default.backup.pattern</code> is set to <code>true</code>. The default value is <code>false</code>.</p> <p>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 <code>new.default.backup.pattern</code>. The keywords should be separated using an underscore. For example, <code>new.default.backup.pattern=dbname_profile_hostname_scope_mod e</code>.</p> <p> The timestamp is included automatically at the end of the generated label.</p>
<code>allow.underscore.in.clone.sid</code>	<p>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.</p> <p>To include an underscore in the clone SID name, ensure that the value of <code>allow.underscore.in.clone.sid</code> is set to <code>true</code>. The default value is true.</p> <p>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 <code>false</code>.</p>

Parameters	Description
oracle.parameters.with.comma	<p>This parameter enables you to specify all the Oracle parameters that have comma (,) as the value. While performing any operation SnapManager uses <code>oracle.parameters.with.comma</code> to check all the Oracle parameters and skip the splitting of the values.</p> <p>For example, if the value of <code>nls_numeric_characters=,,</code> then specify <code>oracle.parameters.with.comma=nls_numeric_characters</code>. If there are multiple Oracle parameters with comma as the value, you must specify all the parameters in <code>oracle.parameters.with.comma</code>.</p>

Parameters	Description
<ul style="list-style-type: none"> archivedLogs.exclude archivedLogs.exclude.fileslike <db-unique-name>.archivedLogs.exclude.fileslike 	<p>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.</p> <p> You must include the exclude parameters in the configuration file before creating a profile.</p> <p>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. If a top-level directory or a mount point is specified and if data protection is enabled for a profile on the host, then that mount point or directory is not included in the dataset that is created in Protection Manager. When there are multiple archive log files to be excluded from the host, you must separate the archive log file paths by using commas.</p> <p>To exclude archive log files from being included in the profile and being backed up, you must include one of the following parameters:</p> <ul style="list-style-type: none"> archivedLogs.exclude to specify a regular expression for excluding archive log files from all profiles or backups. <p>The archive log files matching the regular expression are excluded from all the profiles and backups.</p> <p>For example, you can set <code>archivedLogs.exclude = /arch/logs/on/local/disk1/.*,/arch/logs/on/local/disk2/..</code></p> <p>For ASM databases, you can set <code>archivedLogs.exclude = \\+KHDB_ARCH_DEST/khdb/archivelog/.*,\\+KHDB_NONNARCHTWO/khdb/archivelog/.*.</code></p> <ul style="list-style-type: none"> archivedLogs.exclude.fileslike to specify an SQL expression for excluding archive log files from all profiles or backups. <p>The archive log files matching the SQL expression are excluded from all the profiles and backups.</p> <p>For example, you can set <code>archivedLogs.exclude.fileslike = /arch/logs/on/local/disk1/%,/arch/logs/on/local/disk2/%.</code></p> <ul style="list-style-type: none"> <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 <code>db-unique-name</code>. <p>The archive log files matching the SQL expression are excluded from the profile and backups.</p> <p>For example, you can set <code>mydb.archivedLogs.exclude.fileslike = /arch/logs/on/local/disk1/%,/arch/logs/on/local/disk2/%.</code></p> <p> The BR*Tools do not support the following parameters even if these parameters are configured to exclude archive log files:</p>

Security and credential management

• `<db-unique-name>.archivedLogs.exclude.fileslike`
You can manage security in SnapManager by applying user authentication and role-based access control (RBAC). The user authentication method allows you to access resources such as repositories, hosts, and profiles. RBAC allows you to restrict the operations that SnapManager can perform against the volumes and LUNs containing the data files in your database.

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).

You can manage role-based access by performing the following tasks:

- Enable RBAC for SnapManager by using SnapDrive.
- Assign users to roles and set role capabilities by using the Operations Manager console.
- Optionally, enable SnapManager to store encrypted passwords by editing the `smsap.config` file.

If Protection Manager is installed, access to the features is affected in the following ways:

- If Protection Manager is installed, when you create a database profile, SnapManager creates a dataset and populates the dataset with the volumes that contain the database files.

After a backup operation, SnapManager keeps the dataset contents synchronized with the database files.

- If Protection Manager is not installed, SnapManager cannot create a dataset and you cannot set protection on profiles.

What user authentication is

In addition to using role-based access control (RBAC), 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 `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).

About role-based access control

Role-based access control (RBAC) lets you control who has access to SnapManager operations. RBAC allows administrators to manage groups of users by defining roles and assigning users to those roles. You might want to use SnapManager RBAC in environments where RBAC is already in place.

RBAC includes the following components:

- Resources: Volumes and LUNs that hold the datafiles that make up your database.
- Capabilities: Types of operations that can be performed on a resource.

- Users: People to whom you grant capabilities.
- Roles: A set of resources and capabilities allowed on resources. You assign a specific role to a user who should perform those capabilities.

You enable RBAC in SnapDrive. You can then configure specific capabilities per role in the Operations Manager Web graphical user interface or command-line interface. RBAC checks occur in the DataFabric Manager server.

The following table lists some roles and their typical tasks, as set in Operations Manager.

Role	Typical tasks
SAP basis administrator	<ul style="list-style-type: none"> • Creating, maintaining, and monitoring an Oracle database that resides on a host • Scheduling and creating database backups • Ensuring that backups are valid and can be restored • Cloning databases
Server administrator	<ul style="list-style-type: none"> • Setting up storage systems and aggregates • Monitoring volumes for free space • Provisioning storage on requests from users • Configuring and monitoring disaster recovery mirroring
Storage architect	<ul style="list-style-type: none"> • Making architectural decisions on storage • Planning storage capacity growth • Planning disaster recovery strategies • Delegating capabilities to members of the team

If RBAC is in use (meaning that Operations Manager is installed and RBAC is enabled in SnapDrive), the storage administrator needs to assign RBAC permissions on all of the volumes and storage systems for the database files.

Clear user credentials for all hosts, repositories, and profiles

Delete credentials for individual resources

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, secondary, or even tertiary 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.



To use OS authentication for the Real Application Cluster (RAC) databases, the SnapManager server must be running on each node of the RAC environment and the database password must be the same for all Oracle instances in a RAC environment. SnapManager uses the database user name and password to connect to every RAC instance in the profile.

- 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.



To use database authentication for an Automatic Storage Management (ASM) instance, you must enter the user name and password that you use to log in to the ASM instance.

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	Is database authentication supported for the ASM instance
Any non-ASM and non-RAC database	Oracle 10g and Oracle 11g (lesser than 11.2.0.3)	Yes	No
Stand-alone ASM database on UNIX	Oracle 11.2.0.3 and later	Yes	Yes
ASM instance on RAC database on UNIX	Oracle 11.2.0.3	No	No
RAC database on NFS	Oracle 11.2.0.3	Yes	No



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.

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

Variable name	Description	Example value
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
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.

Backing up databases

SnapManager enables the backing up of data on local storage resources by using post-processing scripts or by protecting backups on secondary or tertiary storage resources. The choice to back up to secondary storage provides an additional layer that preserves data in the case of a disaster.

SnapManager also enables storage administrators to configure their backups based on policy plans. By using

SnapManager, administrators can identify backups that do not conform to policy requirements and rectify those immediately.

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 or tertiary storage (also called remote storage).
- 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 based on archive log protection policies

You can select different protection policies for archivelog-only backups based on their requirement.

- 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
- 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

The Protection Manager must be installed to use protection policies for protecting backups. A backup can have one of these protection states: not requested, not protected, or protected.

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

What full and partial backups are

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

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

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



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

Backup types and the number of Snapshot copies

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



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

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

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

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

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



Both backups also create Snapshot copies of the log files.

Volumes in database	Tablespace TS1 (includes 2 database files)	Tablespace TS2 (includes 2 database files)	Snapshot copies created	Total number of Snapshot copies
/vol/volA	TS1.data1	TS2.data1	1 per volume	2
/vol/volB	TS1.data2	TS2.data2	1 per volume	

Volumes in database	Tablespace TS1 (includes 2 database files)	Tablespace TS2 (includes 2 database files)	Snapshot copies created	Total number of Snapshot copies
/vol/volA	TS1.data1	TS2.data1	2 per file	4
/vol/volB	TS1.data2	TS2.data2	2 per file	

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:

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, /oracle/CER/sapdata1/system_1/system.data1) is in the online backup mode longer than certain tablespaces or data files that were specified.

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

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

Partial online backups

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

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

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

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



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

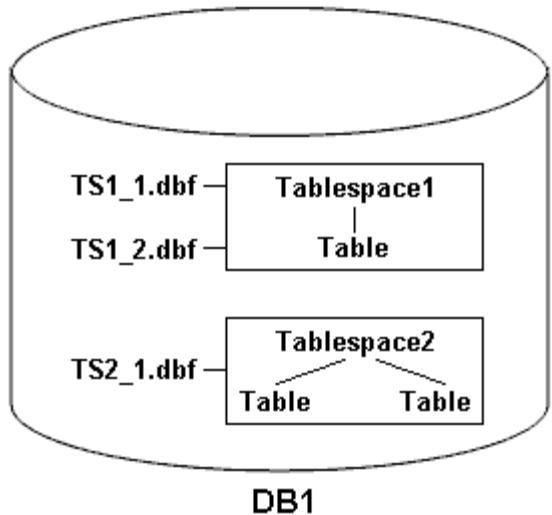
During a partial backup, SnapManager performs these actions:

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: <ul style="list-style-type: none"> • SCN - Enter an SCN, such as 384641. • Date/Time - Enter a date and time of the backup, such as 2005-11-25:19:06:22. • The last transaction made to the database.
	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: <ul style="list-style-type: none"> • Tablespaces • Data files 	SnapManager recovers the data to the last transaction made to the database.
	Restore either data files or tablespaces without control files SnapManager restores one of the following: <ul style="list-style-type: none"> • Tablespaces • Data files 	
	Restore control files only	

Examples of partial backup, restore, and recover operations

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

About control file and archive log file handling

SnapManager includes the control files and optionally includes archive log files with each backup. Archive log files are used for recovery operations.

The database uses control files to identify names, locations, and sizes of the database files. SnapManager

includes control files in each backup because control files are used in the restore process.

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



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

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

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

What database backup scheduling is

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

The following table addresses some common scheduling questions:

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

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

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

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

Create database backups

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.

Change the backup retention policy

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 either a root user or an Oracle user. 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.

Restoring database backups

SnapManager for SAP enables you to restore a database to the state it was when a Snapshot copy was taken. In addition to the file-based restore process, SnapManager supports volume-based fast restore technology, which reduces the restore time significantly compared to other recovery methods. 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 or a volume-based restore, which is the fastest method of restoring database backups and is the default that SnapManager uses.
- 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.
- Restore protected backups from secondary storage (remote backups) and from an alternate location by using the restore specification file.

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 volume-based or file-based backup and restore operations.

The following table describes the restore methods:

Restore process	Details
Volume-based fast restores (from primary storage)	SnapManager restores the data files of a database by restoring a full volume. This default process is the fastest method for restoring your database.
File-based restores	<p>Storage-side full file system restore (from primary or secondary): SnapManager performs a full logical unit number (LUN) restore.</p> <p>Storage-side file restore: SnapManager performs a single file snap restore (SFSR) in a NAS environment. In an SFSR, each of the files or LUNs that represents the protected objects are restored.</p>
	Host-side file copy restore (from primary or secondary): SnapManager clones the local backup using either a LUN or a FlexClone. The clone is mounted and then SnapManager copies the host files from the clone into the active file system.



You cannot restore a backup from the secondary storage, if the backup also exists on the primary storage.

When the fast restore operation is completed, SnapManager performs the following tasks:

- Freed more recent backups (taken after the backup was restored) in the same profile, because their Snapshot copies no longer exist on the primary storage.
- Deletes all Snapshot copies for backups in the same profile that had any Snapshot copies automatically deleted by the fast restore process.

This prevents backups from being partially freed. For example, Backup_A was created first and then Backup_B was created. Each has a Snapshot copy for the data files and one for the archive logs. After SnapManager restores Backup_A by using the fast restore process, SnapManager automatically deletes the data file Snapshot copy from Backup_B. Because the archive log is not restored in the fast restore process, SnapManager must delete Backup_B's Snapshot copy of the archive logs after the fast restore process completes.

Fast restore

Fast restore or volume-based restore is so named because it is the fastest possible restore method. The entire storage system volume is reverted to a Snapshot copy. At the storage level, this restore is almost instantaneous. However, performing a volume restore can have the following negative consequences, and therefore must be used with caution:

- The entire storage side volume is reverted, including the following:
 - Files that were not considered as part of the backup
 - Other files, file systems, or LUNs on the volume
- All the Snapshot copies that were created after the Snapshot copy to which the volume is being reverted are deleted.

For example, you can no longer restore Tuesday's backup if you volume restored Monday's backup.

- Relationships to secondary storage systems are broken if the restored Snapshot copy is older than the baseline Snapshot copy in the relationship.

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.
- In a NAS environment, every file in the file system is restored on the storage system.

For NAS environments, this restore mechanism does not provide additional benefit over storage side file restore.

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)or files (NAS) are restored in place via SFSR.
- When SnapManager restores from secondary storage systems, the LUNs (SAN)or files (NAS) 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.

Storage-side file restore

A storage-side file restore is sometimes performed when a storage-side file system restore cannot be performed. In a storage-side file restore, individual files within a file system are restored directly on the storage systems.

This type of restore can be performed only in NFS environments or in certain cases in ASM environments.

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

- When SnapManager restores NFS files from primary storage systems, the individual files are restored in place by using SFSR.
- When SnapManager restores NFS files from secondary storage systems, the individual files are copied back to the primary storage system over the storage network.

Host-side file restore

A host-side file copy restore is used as a last resort in SAN environments when fast restore, 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

When restoring from the secondary storage, SnapManager first attempts to restore data directly from the secondary storage system to the primary storage system (without involving the host). If SnapManager cannot perform this type of restore (for example, if files not part of the restore exist in a file system), then SnapManager will perform host-side file copy restore. SnapManager has two methods of performing a host-side file copy restore from the secondary storage. The method SnapManager selects is configured in the `smsap.config` file.

- Direct: SnapManager clones the data on the secondary storage, mounts the cloned data from the secondary storage system to the host, and then copies data out of the clone into the active environment. This is the default secondary access policy.
- Indirect: SnapManager first copies the data to a temporary volume on the primary storage, then mounts the data from the temporary volume to the host, and then copies data out of the temporary volume into the active environment. This secondary access policy should be used only if the host does not have direct access to the secondary storage system. Restores using this method take twice as long as the direct secondary access policy because two copies of the data are made.

The decision whether to use the direct or indirect method is controlled by the value of the `restore.secondaryAccessPolicy` parameter in the `smsap.config` configuration file. The default is direct.

Guidelines for when you can use fast restore

Specific rules apply for using fast restore to achieve optimal restore performance. In some cases, you cannot use fast restore.

To achieve optimal restore performance (volume restore or full disk group restore), you must adhere to the following rules:

- Only complete restores of full backups are eligible for fast restore.
- Only data files are eligible for fast restore.
- Data files must be the only files in a volume to be eligible for fast restore.

Although temporary data files can reside in the volume, control files, logs, pfiles, or other files must reside on a separate volume from the data files. You must set up an Oracle database with data files on a separate volume from control files, archived logs, and online log files.

- Data files for only one database must be present in the volume.
- Multiple file systems can be used, but the files in those file systems must be data files for only one database.
- SAP requires a slightly different file layout.

The “General layout and configuration” section contains details.

- If `BRRESTORE` is used to restore the database, fast restore is used with the `fast` parameter in the backup utility parameter file.



To check whether a previously created backup is restorable by using fast restore, you can use the `-preview` option of the `smsap backup restore` command.

The fast restore process cannot be used in the following cases:

- On partial backups
- On backups from the secondary storage if the backup also exists on the primary storage

You cannot restore these using the file-based or volume-based restore.

- On backups protected with SnapVault

The fast restore process cannot be used for backups that were created earlier than the last protected backup. However, you can use the fast restore process for backups created after the last protected backup. For example, consider backups A, B, and C. B is the last backup to transfer to secondary storage by using SnapVault. You can fast restore B and C, but you cannot fast restore A because it was created earlier than the last protected backup. SnapVault needs a baseline SnapVault to compute the time difference and send to the secondary storage the next time a backup is transferred to the secondary storage. The last protected backup acts as the baseline Snapshot copy. Therefore, using the fast restore process prevents SnapVault from being able to recognize the baseline.

- FlexClones or LUN clones that use Snapshot copies that were created after the Snapshot copy to which the volume is being reverted

For example, the clones can be the result of a later backup that is being mounted or being cloned by SnapManager.

- LUNs that are not part of the active SnapDrive Snapshot copy

You cannot perform a fast restore along with other types of restores for the same backup. For example, if one data volume can be restored by using the fast restore process but another data volume cannot, neither is

restored by using the fast restore process. You can choose a file-based restore in this case.

Additionally, you should consider the following points about database restores:

- SnapManager never restores archive logs or redo logs but mounts the backup of archive log files and uses them for recovery.
- SnapManager never restores control files by using volume restore.
- If you want to restore control files and data files, SnapManager performs the restore in two steps.

SnapManager restores the control files first and then the data files.

- If SnapManager finds temporary files in the same volume as the standard tablespace files, you do not need to issue an override to perform a volume-level restore.

After a volume restore, the TEMP tablespace is brought back online.

Both SnapManager for SAP and the BACKINT interface use the same logic when determining which restore mechanism can be used. All restore methods can be used whether the backup was taken with SnapManager for SAP or the BACKINT interface, and whether the restore is performed via SnapManager for SAP or the BACKINT interface.

Related information

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

Advantages and disadvantages of using fast restore

DBAs should be aware of the advantages and disadvantages of using volume-based fast restores.

Restoring database backups using fast restores provides the following advantages:

- Volume-based restores reduce the time needed to restore backups.
- SnapManager provides fast restore eligibility checks. SnapManager analyzes the database backup and displays information about whether it can perform the volume-based restore.
- You can preview the restore operation and decide whether to continue with the recommended path or override the recommendation with your selected process.

Restoring database backups using fast restores has the following disadvantages:

- The entire file system is reverted, including files that were not considered part of the backup. Other files, file systems, or LUNs on the volume will also be reverted.
- SnapManager removes all Snapshot copies that were taken after the Snapshot you are reverting to. In effect, you lose the history after the Snapshot copy date. For example, you cannot restore Tuesday's backup if you already restored Monday's backup.

You can avoid the disadvantages by following these recommendations:

- Optimize the database layout according to best practices.
- Protect backups to secondary storage. However, if you delete Snapshot copies from primary storage, you cannot use fast restores to restore them from secondary storage.

Fast restore eligibility checks

When you choose to perform a fast restore of a backup, SnapManager first performs an eligibility check to determine whether the fast restore process can be used.

SnapManager provides the following types of checks:

- **Mandatory checks:** SnapManager can perform the fast restore process only if all the conditions under this check pass.
- **Overridable checks:** If the conditions under this check fail, administrators can override the check to force a fast restore process. However, you must override these checks with caution.

The following table lists issues that you might encounter and indicates whether the fast restore eligibility check can be overridden:

Issue	Pass required	Details
ACFS, Voting Disk, or OCR is present on ASM Disk group in 11gR2	Yes	Fast restore cannot be performed. Resolution: None Cannot override.
Only backups created with SnapManager 3.0 or later can be fast restored	Yes	Cannot override.
Only Snapshot copies created with SnapDrive for UNIX 4.0 or later can be fast restored	Yes	Cannot override.
Volume is a root volume	Yes	Volume being restored is a root volume on the storage system. Resolution: Do not use the root volume on the storage system. Cannot override.
Volume restore is not available on Windows	Yes	Volume being restored is a root volume on the storage system. Resolution: None Cannot override.
Volume restore is disabled	Yes	Volume restore has been disabled. Resolution: Enable volume restore by selecting different options when starting the restore. In the command-line interface, do not use <code>-fast -off</code> . Cannot override.

Issue	Pass required	Details
Control files and data files on the same volume	Yes	<p>For online backups, control files and data files cannot be on the same volume because SnapManager takes two Snapshot copies of the volume (one in which the data files are consistent in hot backup mode, and one in which the backup control files are consistent after hot backup mode is complete). The volume restore will revert to the first Snapshot copy, which deletes the second Snapshot copy containing the backup control files. When a data file-only restore occurs, the control files are reverted to an inconsistent state, and SnapManager restores the backup control file and then opens the database with the resetlogs option, which is not desired behavior.</p> <p>Resolution: Migrate control files and data files onto separate file systems that do not share the same underlying volume. This does not help the restore in which the check failed, but will help future backup restore operations.</p> <p>Cannot override.</p>

Issue	Pass required	Details
Archive logs and data files must not exist on the same volume	Yes	<p>Database archive logs and data files reside in file systems backed by the same storage system volume. If a volume restore was performed, SnapManager cannot open the database after a restore of an online backup because the archived log file that is written after the database is taken out of hot backup mode is not available. Also, you would not be able to roll forward through later transactions that might have been in the archive log files.</p> <p>Resolution: Migrate archive logs and data files onto separate file systems that do not share the same underlying storage system volume. This does not help the restore in which the check failed, but will help future backup restore operations.</p> <p>Cannot override.</p>
Online logs and data files must not exist on the same volume	Yes	<p>Database online redo logs and data files reside in file systems backed by the same storage system volume. If a volume restore was performed, recovery cannot use the online redo logs because they would have been reverted.</p> <p>Resolution: Migrate online redo logs and data files onto separate file systems that do not share the same underlying storage system volume. This does not help the restore in which the check failed, but will help future backup restore operations.</p> <p>Cannot override.</p>

Issue	Pass required	Details
Files in the file system not part of the restore scope are reverted	Yes	<p>Files visible on the host, other than the files being restored, exist in a file system on the volume. If a fast restore or a storage side file system restore was performed, the files visible on the host would be reverted to their original content when the Snapshot copy is created. If SnapManager discovers 20 or less files, they are listed in the eligibility check. Otherwise, SnapManager displays a message that you should investigate the file system.</p> <p>Resolution: Migrate the files not used by the database onto a different file system that uses a different volume. Alternatively, delete the files.</p> <p>If SnapManager cannot determine the file purpose, you can override the check failure. If you override the check, the files not in the restore scope are reverted. Override this check only if you are certain that reverting the files will not adversely affect anything.</p>

Issue	Pass required	Details
File systems in the specified volume group not part of the restore scope are reverted	No	<p>Multiple file systems are in the same volume group, but not all file systems are requested to be restored. Storage side file system restore and fast restore cannot be used to restore individual file systems within a volume group because the LUNs used by the volume group contain data from all file systems. All file systems within a volume group must be restored at the same time to use fast restore or storage side file system restore. If SnapManager discovers 20 or less files, SnapManager lists them in the eligibility check. Otherwise, SnapManager provides a message that you should investigate the file system.</p> <p>Resolution: Migrate the files not used by the database onto a different volume group. Alternatively, delete the file systems in the volume group.</p> <p>Can override.</p>
Host volumes in specified volume group not part of the restore scope are reverted	No	<p>Multiple host volumes (logical volumes) are in the same volume group, but not all host volumes are requested to be restored. This check is similar to File systems in volume group not part of the restore scope will be reverted except that the other host volumes in the volume group are not mounted as file systems on the host. Resolution: Migrate host volumes used by the database onto a different volume group. Or, delete the other host volumes in the volume group.</p> <p>If you override the check, all the host volumes in the volume group are restored. Override this check only if you are certain that reverting the other host volumes does not adversely affect anything.</p>

Issue	Pass required	Details
File extents have changed since the last backup	Yes	Cannot override.
Mapped LUNs in volume not part of restore scope are reverted	Yes	<p>LUNs other than those requested to be restored in the volume are currently mapped to a host. A volume restore cannot be performed because other hosts or applications using these LUNs will become unstable. If the LUN names end with an underscore and an integer index (for example, _0 or _1), these LUNs are typically clones of other LUNs within the same volume. It is possible that another backup of the database is mounted, or a clone of another backup exists.</p> <p>Resolution: Migrate LUNs not used by the database onto a different volume. If the mapped LUNs are clones, look for mounted backups of the same database or clones of the database, and unmount the backup or remove the clone.</p> <p>Cannot override.</p>
Unmapped LUNs in volume not part of the restore scope are reverted	No	<p>LUNs other than those requested to be restored in the volume exist. These LUNs are not currently mapped to any host, so restoring them does not disrupt any active processes. However, the LUNs may be temporarily unmapped.</p> <p>Resolution: Migrate LUNs not used by the database onto a different volume, or delete the LUNs.</p> <p>If you override this check, the volume restore will revert these LUNs to the state at which the Snapshot copy was made. If the LUN did not exist when the Snapshot copy was made, the LUN will not exist after the volume restore. Override this check only if you are certain that reverting the LUNs does not adversely affect anything.</p>

Issue	Pass required	Details
LUNs present in Snapshot copy of volume might not be consistent when reverted	No	<p>During Snapshot copy creation, LUNs other than those for which the Snapshot copy was requested, existed in the volume. These other LUNs may not be in a consistent state. Resolution: Migrate LUNs not used by the database onto a different volume, or delete the LUNs. This does not help the restore process in which the check failed, but will help restores of future backups taken after the LUNs are moved or deleted.</p> <p>If you override this check, the LUNs reverts to the inconsistent state at which the Snapshot copy was made. Override this check only if you are certain that reverting the LUNs does not adversely affect anything.</p>
New Snapshot copies have volume clone	Yes	<p>Clones have been created of Snapshot copies that were created after the Snapshot copy is requested to be restored. Because a volume restore will delete later Snapshot copies, and a Snapshot copy cannot be deleted if it has a clone, a volume restore cannot be performed. Resolution: Delete clones of later Snapshot copies.</p> <p>Cannot override.</p>
Newer backups are mounted	Yes	<p>Backups taken after the backup is restored are mounted. Because a volume restore deletes later Snapshot copies, a Snapshot copy cannot be deleted if it has a clone, a backup mount operation creates cloned storage, and a volume restore cannot be performed. Resolution: Unmount the later backup, or restore from a backup taken after the mounted backup.</p> <p>Cannot override.</p>

Issue	Pass required	Details
Clones of newer backups exist	Yes	<p>Backups taken after the backup is restored have been cloned. Because a volume restore deletes later Snapshot copies, and a Snapshot copy cannot be deleted if it has a clone, a volume restore cannot be performed. Resolution: Delete the clone of the newer backup, or restore from a backup taken after the backups that have clones.</p> <p>Cannot override.</p>
New Snapshot copies of volume is lost	No	<p>Performing a volume restore deletes all Snapshot copies created after the Snapshot copy to which the volume is being restored. If SnapManager can map a later Snapshot copy back to a SnapManager backup in the same profile, then the "Newer backups will be freed or deleted" message appears. If SnapManager cannot map a later Snapshot copy back to a SnapManager backup in the same profile, this message does not appear. Resolution: Restore from a later backup, or delete the later Snapshot copies.</p> <p>Can override.</p>

Issue	Pass required	Details
Newer backups will be freed or deleted	No	<p>Performing a volume restore deletes all the Snapshot copies created after the Snapshot copy to which the volume is being restored. Therefore, any backups created after the backup that is being restored are either deleted or freed. Later backups are deleted in the following scenarios:</p> <ul style="list-style-type: none"> • The backup state is not PROTECTED • <code>retain.alwaysFreeExpire</code> <code>dBackups</code> is false in <code>smsap.config</code> <p>Later backups are freed in the following scenarios:</p> <ul style="list-style-type: none"> • The backup state is PROTECTED • <code>retain.alwaysFreeExpire</code> <code>dBackups</code> is true false in <code>smsap.config</code> <p>Resolution: Restore from a later backup, or free or delete later backups.</p> <p>If you override this check, backups created after the backup that is being restored are deleted or freed.</p>
SnapMirror relationship for volume is lost	Yes (If RBAC is disabled or you do not have RBAC permission)	<p>Restoring a volume to a Snapshot copy earlier than the baseline Snapshot copy in a SnapMirror relationship destroys the relationship. Resolution: Restore from a backup created after the relationship's baseline Snapshot copy. Alternatively, break the storage relationship manually (and then re-create and re-baseline the relationship after the restore is complete).</p> <p>Can override, if RBAC is enabled and you have RBAC permission.</p>

Issue	Pass required	Details
SnapVault relationship for volume is lost if the fast restore process occurred	Yes (If RBAC is disabled or you do not have RBAC permission)	<p>Restoring a volume to a Snapshot copy earlier than the baseline Snapshot copy in a SnapVault relationship destroys the relationship. Resolution: Restore from a backup created after the relationship's baseline Snapshot copy. Alternatively, break the storage relationship manually (and then re-create and re-baseline the relationship after the restore is complete).</p> <p>Cannot override, if RBAC is enabled and you have RBAC permission.</p>
NFS files in volume not part of the restore scope are reverted	No	<p>Files present in the storage system volume, which are not visible on the host, are reverted if a volume restore is performed. Resolution: Migrate files not used by the database onto a different volume or delete the files.</p> <p>Can override. If you override this check failure, the LUNs are deleted.</p>
CIFS shares exist for volume	No	<p>The volume being restored has CIFS shares. Other hosts might be accessing files in the volume during the volume restore. Resolution: Remove unneeded CIFS shares.</p> <p>Can override.</p>
Restoring from alternate location	Yes	<p>A restore specification was provided for the restore operation that specifies that the files be restored from an alternate location. Only host-side copy utilities can be used to restore from an alternate location.</p> <p>Resolution: None.</p> <p>Cannot override.</p>

Issue	Pass required	Details
Storage side file system restore is not supported in a RAC database	Yes	Cannot override.

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.



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	Database state for other instance (RAC only)
Restore only	Control files	Shutdown	Shutdown
	System files	Mount or Shutdown	Mount or Shutdown
	No system files	Any state	Any state
Restore and recovery	Control files	Shutdown	Shutdown
	System files	Mount	Mount or Shutdown
	No system files	Mount or Open	Any

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 using fast volume-based restore, 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: fast restore  
+DG1/rac6/users.dbf
```

```
The following files will be restored completely via: storage side file  
system restore  
+DG2/rac6/sysaux.dbf  
+DG2/rac6/system.dbf
```

```
The following files will be restored completely via: storage side system  
restore  
+DG2/rac6/undotbs1.dbf  
+DG2/rac6/undotbs2.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. In the example above, the fast restore method is more efficient than the storage file system restore method and so is displayed first.

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:  
[fast restore, storage side file system restore. 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: fast restore

- * LUNs present in snapshot of volume fas960:

- /vol/rac_6_asm_disks may not be consistent when reverted:

- [fas960:/vol/rac6_asm_disks/DG4D1.lun]

- Mapped LUNs in volume fas960:/vol/rac_6_asm_disks

- not part of the restore scope will be reverted: [DG4D1.lun]

```
Files to restore:
```

- +DG2/rac6/sysaux.dbf

- +DG2/rac6/system.dbf

- +DG2/rac6/undotbs1.dbf

- +DG2/rac6/undotbs2.dbf

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

In the example, the first failure is overridable by using `-fast -override` 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.

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 the following data from an alternate location:

- The data files from an intermediate file system to an active file system
- The blocks of data from an intermediate raw device into an active raw device

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).

Restoration of the data from raw devices

Before you restore from an alternate location, you need to connect to the necessary raw devices that exist on the local host.

You can invoke the restore from an alternate location operation to copy the blocks of data from alternate raw devices to active raw devices. To perform this operation, you need to specify the alternate raw device from which to restore the original raw device by creating a restore specification.

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 protected backup on the secondary or even tertiary storage.
- 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.
- Clone a RAC database.
- 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 can clone a backup on the primary (local) storage or a protected backup that is on the secondary (remote) storage. However, you cannot clone a backup if the backup operation is in progress or the backup has been transferred to the secondary storage.

You must ensure that the following prerequisites are met before a database can be cloned:

- Ensure that the [/etc|/var/opt/oracle]/oratab directory does not contain an entry pointing to the target system identifier.
- Delete the spfile<SID>.ora file from \$ORACLE_HOME/dbs.
- Delete the init<SID>.ora file from \$ORACLE_HOME/dbs.
- 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/dbs directory.

You can clone a Real Application Cluster (RAC) database as well as a nonclustered database. A RAC clone starts as a single database.

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/dbs` 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	<code>clone create -reserve</code>
LUN cloning	A new clone LUN is created within the same volume.	When <code>-reserve</code> 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 <code>-reserve</code> for a volume is set to yes, space is reserved for the full volume size within the aggregate.

Create clone specifications

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.

Prerequisite set up	Requirement
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
Volume managers	Must be configured on both the source and target hosts and must be of compatible versions

Splitting a clone

SnapManager enables you to split and manage an existing clone that was created by using the FlexClone technology. In the FlexClone technology, the clone and original database share the same physical data blocks.

Before you perform the clone split operation, you can know that the estimated size of the clone to be split and the required space available on the aggregate.

A new profile is generated by SnapManager if the clone split operation is successful. If SnapManager fails to create the new profile, you can manually create a new profile. You can use the new profile to create database backups, restore data, and create clones. If the clone split operation is successful, irrespective of whether the new profile is created or not, the clone-related metadata is removed from the repository database.

You can perform the following tasks related to splitting clones:

- View the clone split estimate.
- Split a clone on a primary storage.
- Split a clone on a secondary storage.
- View the clone split operation status.
- Stop the clone split operation.
- Destroy the profile along with the underlying storage.
- Delete the profile created for a split clone.

When you split a clone from its parent volume, the Snapshot copies associated with the cloned volume are deleted. The backups created for the cloned database before the clone split process cannot be used because the Snapshot copies of these backups are deleted, and the backups remain as stale entries in the repository.

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.

If you are using Data ONTAP operating in 7-Mode, SnapManager provides policy-driven data protection by integrating with Protection Manager (OnCommand Unified Manager). This automates replicating SnapManager backups on a primary storage system to a secondary storage system or even to a tertiary storage system by using SnapVault or SnapMirror policies created by the storage or backup administrator in Protection Manager. Retention on primary storage is controlled by SnapManager based on the retention defined during profile

creation and the retention class tagged during the backup creation. Secondary storage backup retention is controlled by the policy defined in Protection Manager.

If you are using clustered Data ONTAP, SnapManager 3.4 provides *SnapManager_cDOT_Mirror* and *SnapManager_cDOT_Vault* policies for data protection. While creating a profile, you can select these policies depending on the SnapMirror or SnapVault relationship that was established using clustered Data ONTAP CLI or System Manager. When you create a backup selecting the profile for which you enabled protection, the backups are protected to a secondary storage system.



Backups created using BR*Tools cannot be protected using SnapManager for SAP.

If you were using SnapManager 3.3.1 with clustered Data ONTAP, the backups were protected using post-scripts which were selected while creating profiles. If you want to use those profiles, after upgrading to SnapManager 3.4 you must perform the following operations.

- You must update the profiles to select either *SnapManager_cDOT_Mirror* or *SnapManager_cDOT_Vault* policy and delete the post-script that was used for data protection.
- After updating profile to use *SnapManager_cDOT_Vault* policy, you must delete existing backup schedules and create new schedules to specify the SnapVault label for the backups.
- If the profiles were created in SnapManager 3.3.1 without selecting the post-scripts, you must update the profiles to select either *SnapManager_cDOT_Mirror* or *SnapManager_cDOT_Vault* policy to enable data protection.



If you have backups in the secondary storage system that were mirrored or vaulted using SnapManager 3.3.1 post-scripts, you cannot use those backups for restore or cloning using SnapManager 3.4.

If you are using clustered Data ONTAP, SnapManager 3.4.2 supports multiple protection relationships (SnapMirror and SnapVault) on source volumes. Only one SnapMirror and one SnapVault relationship per volume is supported. You must create separate profiles, each with the *SnapManager_cDOT_Mirror* and the *SnapManager_cDOT_Vault* policy selected.



Snapdrive for Unix 5.3.2 and later is required to use multiple protection policies.

What protection policies are

Protection policies are rules that govern how database backups are protected. You can select the protection policies while creating the profile.

A protection policy defines the following parameters:

- When to transfer copies to secondary storage
- The maximum amount of data that should be transferred at scheduled times
- How long to retain copies for each backup location
- Warning and error thresholds for lag times

When protection is enabled, SnapManager creates a dataset for the database. A dataset consists of a collection of storage sets along with configuration information associated with their data. The storage sets associated with a dataset include a primary storage set used to export data to clients, and the set of replicas and archives that exist on other storage sets. Datasets represent exportable user data. If the administrator

disables protection for a database, SnapManager deletes the dataset.

What protection states are

SnapManager shows the state of each backup. Administrators must know the different states and monitor the state of their backups.

A database backup can have the following protection states:

Status	Definition	Explanation
Protected	Protection was requested and has been enabled.	Protection is enabled for the backup in SnapManager and the Protection Manager successfully copied the backup to another set of physical disks (also referred to as secondary storage). If the Protection Manager removes a backup from secondary storage due to a retention policy, the backup can return to a Not protected state.
Not protected	Protection was requested, but not completed.	Protection is enabled for the backup, but the backup is not copied to another set of physical disks. The backup is not yet protected, or protection failed, or it was protected earlier but is no longer protected. When you create a backup, the initial protection state of the backup is either Not requested or Not protected. If the backup is not protected, it becomes protected when it is transferred to secondary storage.
Not requested	Protection was not requested.	Protection is not enabled for the backup. A logical copy of the data exists on the same physical disks (also referred to as a local backup). If protection is not requested when the backup was created, protection on the backup is always shown as Not requested.

What resource pools are

A resource pool is a collection of unused physical storage (such as storage systems or aggregates) from which new volumes or LUNs can be provisioned to contain data. If you assign a storage system to a resource pool, all the aggregates on that storage system become available for provisioning.

Storage administrators use the Protection Manager's console to assign a resource pool to the backup and mirror copies. The provision application can then automatically provision volumes out of the physical resources in the resource pool to contain backups and mirror copies.

For protected profiles, SnapManager displays information about a profile and indicates whether a storage resource pool has been assigned to that profile. If not, the profile is considered "non-conformant." After a storage resource pool has been assigned to the corresponding profile's dataset, the profile is considered "conformant".

About different protection policies

You can select different policies to protect the backups on the secondary or tertiary storage systems.

If you are using Data ONTAP operating in 7-Mode and SnapManager is integrated with Protection Manager, you must select one of the following protection policies while creating the profile. The Protection Manager's Management Console provides templates to configure protection policies for the datasets. Even though disaster recovery protection policies are listed in the SnapManager user interface, these policies are not supported.

Policy	Description
Back up	A dataset is backed up locally and also from the primary to secondary storage by using SnapVault or SnapMirror.
Back up, then mirror	A dataset is backed up from the primary to secondary storage by using SnapVault or SnapMirror, and then mirrored to a SnapMirror partner.
Local Snapshot copies only	A dataset uses only local Snapshot copies in the primary storage.
Mirror	A dataset is mirrored from the primary to secondary storage by using SnapMirror.
Mirror and back up	A dataset is mirrored from the primary to secondary storage by using SnapMirror, and then backed up to the secondary storage by using SnapVault or SnapMirror.
Mirror and mirror	A dataset is mirrored from the primary to secondary storage on two different SnapMirror partners.
Mirror, then back up	A dataset is mirrored from the primary to secondary storage by using SnapMirror, and then backed up to tertiary storage by using SnapVault or SnapMirror.
Mirror, then mirror	A dataset is mirrored from the primary to secondary storage by using SnapMirror, and then mirrored to an additional SnapMirror partner.
No protection	A dataset has no Snapshot copies, backups, or mirror-copy protection of any kind.
Remote backup only	Data on a storage system is backed up remotely to secondary storage by using SnapVault or SnapMirror. The licensed application carries out no local backup on the primary storage. This protection policy can be applied to third-party systems with Open Systems SnapVault installed.

If you are using clustered Data ONTAP, you must select one of following protection policies while creating the profile.

Policy	Description
SnapManager_cDOT_Mirror	Mirrors the backup.
SnapManager_cDOT_Vault	Vaults the backup.

Configuring and enabling policy-driven data protection

You must configure SnapDrive and the DataFabric Manager server to enable data protection on the profile to protect backups on the secondary storage systems. You can select the protection policies in the Protection Manager's console to specify how database backups will be protected.



You must ensure that OnCommand Unified Manager is installed on a separate server to enable data protection.

Understanding enabling or disabling of data protection in profile

You can enable or disable data protection while creating or updating a database profile.

To create a protected backup of a database on the secondary storage resources, database administrators and storage administrators perform the following actions.

If you want to...	Then...
Create or edit a profile	<p>To create or edit a profile, perform the following:</p> <ul style="list-style-type: none"> Enable backup protection to the secondary storage. If you are using Data ONTAP operating in 7-Mode and have installed Protection Manager, you can select the policies created by the storage or backup administrator in Protection Manager. <p>If you are using Data ONTAP operating in 7-Mode and protection is enabled, SnapManager creates a dataset for the database. A dataset consists of a collection of storage sets along with configuration information associated with their data. The storage sets associated with a dataset include a primary storage set used to export data to clients, and the set of replicas and archives that exist on other storage sets. Datasets represent exportable user data. If the administrator disables protection for a database, SnapManager deletes the dataset.</p> <ul style="list-style-type: none"> If you are using ONTAP, you must select either the <i>SnapManager_cDOT_Mirror</i> or <i>SnapManager_cDOT_Vault</i> policy depending on the SnapMirror or SnapVault relationship created. <p>When you disable backup protection, a warning message is displayed stating that the dataset will be deleted and restoring or cloning backups for this profile will not be possible.</p>

If you want to...	Then...
View the profile	Because the storage administrator has not yet assigned storage resources to implement the protection policy, the profile shows up as nonconformant in both the SnapManager graphical user interface and the <code>profile show</code> command output.
Assign storage resources in the Protection Manager Management Console	In the Protection Manager Management Console, the storage administrator views the unprotected dataset and assigns a resource pool for each node of the dataset that is associated with the profile. The storage administrator then makes sure that secondary volumes are provisioned and protection relationships are initialized.
View the conformant profile in SnapManager	In SnapManager, the database administrator sees that the profile has changed to conformant state in both the graphical user interface and in the <code>profile show</code> command output, indicating that resources were assigned.
Create the backup	<ul style="list-style-type: none"> • Select full backup. • Also, select whether the backup should be protected and select the primary retention class (for example, hourly or daily). • If you are using Data ONTAP operating in 7-Mode and want to protect the backup immediately to secondary storage overriding the Protection Manager protection schedule, specify the <code>-protectnow</code> option. • If you are using ONTAP and want to protect the backup immediately to the secondary storage, specify the <code>protect</code> option. <p> The <code>protectnow</code> option is not applicable in clustered Data ONTAP.</p>
View the backup	The new backup is shown as scheduled for protection, but not yet protected (in the SnapManager interface and in the <code>backup show</code> command output). The Protection state is shown as "Not protected".
View the backup list	After the storage administrator verifies that the backup has been copied to secondary storage, SnapManager changes the backup Protection state from "Not protected" to "Protected".

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.

Backups not included in the retention count	Additional details
Backups that are cloned or mounted on secondary storage and that use the mirror protection policy	If SnapManager deletes the Snapshot copies for the backup on the primary storage resource and the Snapshot copies are mirrored, the next backup to the secondary storage will fail.

When you free a backup from its primary storage resources, the primary resources (Snapshot copies) used by the backup are destroyed, but the backup metadata is still available. SnapManager does not consider freed backups in 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.

Backups on primary storage can be protected by backing up to secondary storage. While SnapManager manages the retention and scheduling of backups on primary storage, the Protection Manager manages the retention and scheduling of backups on secondary storage.

When local backups expire based on their retention policy, they are either deleted or freed, depending on whether they are protected.

- If they are protected, the local backups are freed. Their storage resources or Snapshot copies are deleted, but the backups remain in the SnapManager repository and are available for restoration from the secondary storage. You do not have to free backups (for example, with the `backup free` command). Backups are freed until the backup no longer exists on the secondary storage, and at that point, the backup is deleted.
- If they are not protected, the local 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.
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.

Backup date	Status	Retention policy action taken	Explanation
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.

- In NAS environments, you must configure the primary and secondary NAS data path by using the `snapdrive config set -mgmtpath management_path management_path management_path datapath_path` command.

For example, `snapdrive config set -mgmtpath f3050-197-91 f3050-197-91 f3050-197-91 f3050-220-91`, where f3050-197-91 is the management path and f3050-220-91 is the data path.

[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 (depending on the protection policy)
- SnapRestore
- SnapMirror (depending on the protection policy)
- FlexClone is required for Network File System (NFS) and 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, NFS, Internet Small Computer System Interface (iSCSI), or Fibre Channel (FC)

SnapVault or SnapMirror should be on the primary and secondary storage systems based on the protection policies used. The basic backup protection policies require only SnapVault installed on the supporting systems. The policies that include mirror protection require SnapMirror installed on the supporting systems. The backup and mirror disaster recovery policies require SnapMirror installed on the supporting systems.

Restoring protected backups from secondary storage

You can restore protected backups from secondary storage. However, you cannot restore backups from secondary storage if the backup also exists on primary storage.

Restores of protected backups overview

You can choose the restore method that you want to use to restore the backup data from secondary storage to primary storage.

The following table explains the different scenarios and methods that you can use to restore a backup from secondary storage:

Restore target	Explanation
Directly to primary storage	<p>Returns the data from the secondary storage system directly to the original location on the primary storage system over the same network that was used to protect the data.</p> <p>SnapManager uses the direct storage method whenever possible. This method is not possible if the data is in a file system on storage area network (SAN) and if any of the following conditions apply:</p> <ul style="list-style-type: none">• Other non-database files are not being restored in the same file system.• Snapshot copies of the control files and data files in a file system being restored were taken at different times.• The logical unit number (LUN) is in a volume group, but other LUNs in the same volume group are not being restored.
Directly to host	<p>Clones the data on the secondary storage system and mounts the cloned data on the host. After the data is cloned and mounted, SnapManager copies it into its original location.</p>
Indirectly to storage or host	<p>Returns the data from the secondary storage system to a new location on the primary system over the same network that was used to protect the data and to mount the new storage on the host. After the data is returned and mounted, SnapManager copies it into its original location. The indirect storage method might require a long time to return the data.</p> <p>SnapManager first copies data to a scratch volume on the primary host before SnapManager uses it to restore and recover the database. Whether the scratch data is automatically deleted depends on the protocol used.</p> <ul style="list-style-type: none">• For SAN, SnapManager deletes the returned data.• For network-attached storage (NAS), SnapManager deletes the contents of the returned qtrees, but does not delete the qtrees themselves. To delete the qtrees, administrators should mount the scratch volume and remove the qtrees using the UNIX rmdir command.

If you cannot directly return data to storage, SnapManager can return data either directly to host or indirectly to storage or host. The method depends on the policy governing whether the organization allows connection directly to secondary storage or requires data to be copied over the storage network. You can manage this policy by setting configuration information in the smsap.config file.

SnapManager for SAP uses Protection Manager to protect a database backup

SnapManager for SAP and Protection Manager, when installed on a UNIX host and on the server respectively, give the SnapManager database administrator (DBA) the ability to configure and carry out policy-based Oracle database backups to secondary storage, and to restore, if necessary, the backed up data from secondary to primary storage.

In the following example, a DBA, who is using SnapManager, creates a profile for a local backup on primary storage and another profile for a protected backup to secondary storage. Then this DBA works with his network storage administrator, who is using the Protection Manager's console, to configure a policy-based backup of that database from primary to secondary storage.

Details of the target database

This example of integrated database protection describes the protection of a payroll database. The following data is used in the example.

The database administrator (DBA) at TechCo, a 3000-person company headquartered in Atlanta, must create a consistent backup of the production payroll database, PAYDB. The protection strategy for backing up to primary and secondary storage requires that the DBA and the storage administrator work together to back up the Oracle database both locally on primary storage and also remotely, to secondary storage at a remote location.

- **Profile information**

When creating a profile in SnapManager, you need the following data:

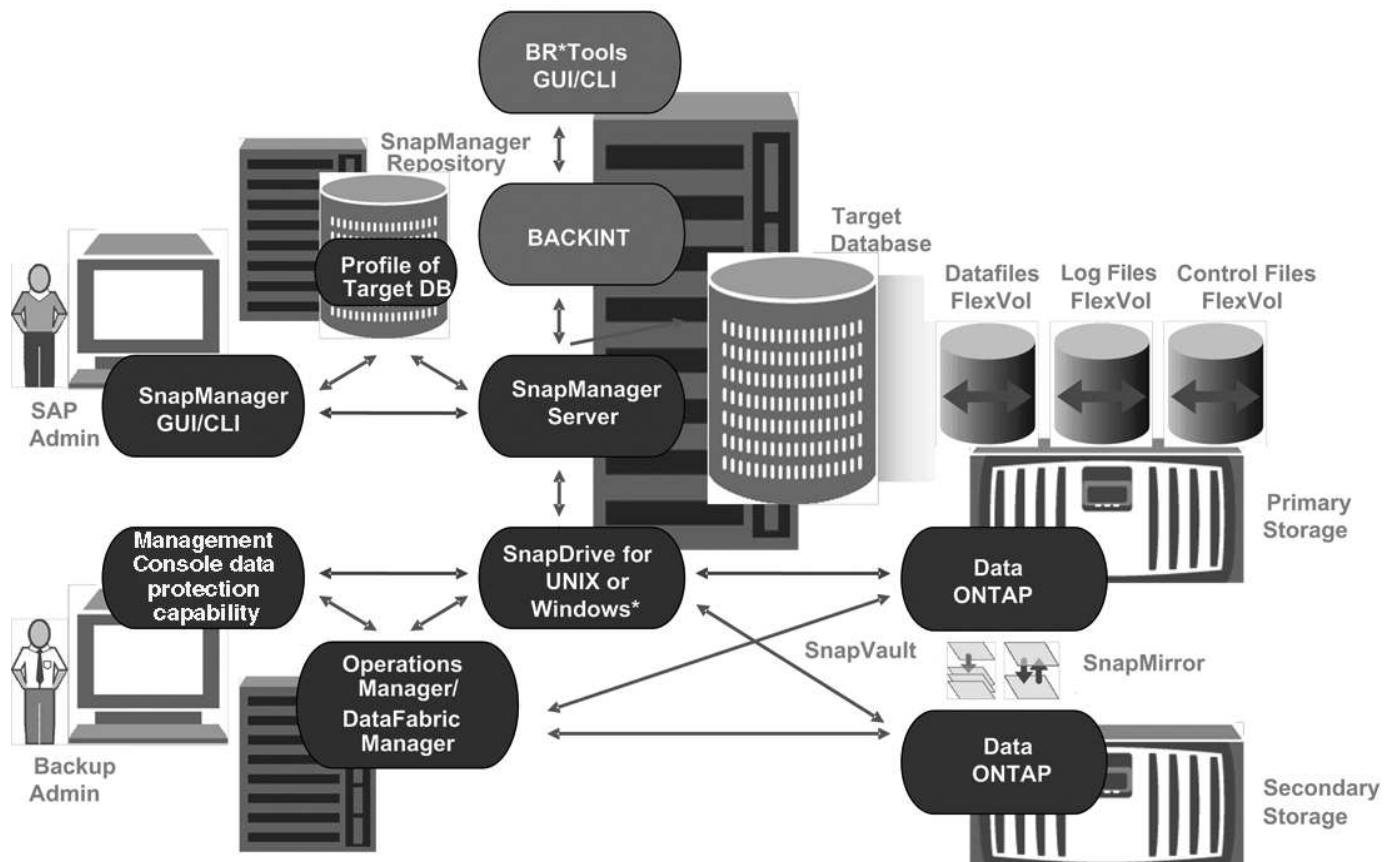
- Database name: P01
- Host name: prod01.sample.com
- Database ID: P01
- Profile name: P01_BACKUP
- Connection mode: Database authentication
- Snapshot naming scheme: *smsap_hostname_dbsid_smsaprofile_scope_mode_smid* (which translates to "smsap_prod01.sample.com_p01_p01_backup_f_h_x")
- Repository user: <sid>rep, which translates to p01rep.

Primary and secondary storage configuration and topology

In this example, the TechCo corporation runs its payroll database on a database server that is also a SnapManager for SAP host and stores its payroll database data and configuration files on primary storage systems at company headquarters. The corporate requirement is to protect that database with daily and weekly backups to local storage as well as backups to storage systems at a secondary storage site fifty miles away.

The following illustration shows the SnapManager for SAP and the NetApp Management Console data protection capability components required to support local and secondary backup protection.

SnapManager for SAP Architecture



To manage the payroll database and support its local and secondary backup protection as illustrated in the previous graphic, the following deployment is used.

- **SnapManager host**

The SnapManager host, payroll.techco.com, is located at company headquarters and runs on a UNIX server, which also runs the database program that generates and maintains the payroll database.

- **Connections**

To support local backup and secondary backup protection, the SnapManager host has network connections to the following components:

- SnapManager for SAP client
- SnapManager repository, which runs the database program, SnapDrive for UNIX, and SnapManager
- Primary storage systems
- Secondary storage systems
- DataFabric Manager server

- **Installed products**

The SnapManager host is installed with the following products for this example:

- SnapManager server
- SnapDrive for UNIX
- Host Utilities

- **TechCo primary storage systems**

The payroll database, including associated data files, log files, and control files, reside on the primary storage systems. These are located at TechCo company headquarters along with the SnapManager host and the network connecting primary storage and the SnapManager host. The latest payroll database transactions and updates are written to the primary storage systems. Snapshot copies, which provide local backup protection of the payroll database, also reside on the primary storage systems.

- **Connections**

To support secondary backup protection, the primary storage systems have network connections to the following components:

- SnapManager host running the database program, SnapDrive for UNIX, and SnapManager
- Secondary storage systems
- DataFabric Manager server

- **Installed products**

The following licenses must be enabled on these systems for this example:

- Data ONTAP 7.3.1 or later
- SnapVaultData ONTAP Primary
- FlexVol (required for NFS)
- SnapRestore
- NFS protocol

- **TechCo secondary storage systems**

The secondary storage systems, located at a network-connected secondary storage site fifty miles away, are used to store secondary backups of the payroll database.

- **Connections**

To support secondary backup protection, the secondary storage systems have network connections to the following components:

- Primary storage systems
- DataFabric Manager server

- **Installed products**

The following licenses must be enabled on the secondary storage systems for this example:

- Data ONTAP
- SnapVaultData ONTAP Secondary
- SnapRestore

- FlexVol (required for NFS)
- NFS protocol

- **DataFabric Manager server**

The DataFabric Manager server, techco_dfm, is located at company headquarters in a location accessible by the storage administrator. The DataFabric Manager server, among other functions, coordinates the backup tasks between primary and secondary storage.

- **Connections**

To support secondary backup protection, the DataFabric Manager server maintains network connections to the following components:

- NetApp Management Console
- Primary storage systems
- Secondary storage systems

- **Installed products**

The DataFabric Manager server is licensed for the following server products for this example:

- DataFabric Manager

- **SnapManager repository**

The SnapManager repository, located on a dedicated server, stores data about operations performed by SnapManager, for example the time of backups, tablespaces and datafiles backed up, storage systems used, clones made, and Snapshot copies created. When a DBA attempts a full or partial restore, SnapManager queries the repository to identify backups that were created by SnapManager for SAP for restoration.

- **Connections**

To support secondary backup protection, the secondary storage systems have network connections to the following components:

- SnapManager host
- SnapManager for SAP client

- **NetApp Management Console**

The NetApp Management Console is the graphical user interface console used by the storage administrator to configure schedules, policies, datasets, and resource pool assignments to enable backup to secondary storage systems, which are accessible to the storage administrator.

- **Connections**

To support secondary backup protection, NetApp Management Console has network connections to the following components:

- Primary storage systems
- Secondary storage systems
- DataFabric Manager server

- **SnapManager for SAP client**

The SnapManager for SAP client is the graphical user interface and command line console used by the DBA for the payroll database in this example to configure and carry out local backup and backup to secondary storage.

- **Connections**

To support local backup and secondary backup protection, SnapManager for SAP client has network connections to the following components:

- SnapManager host
- SnapManager repository, running the database program, SnapDrive for UNIX, and SnapManager
- Database host (if separate from the host running SnapManager)
- DataFabric Manager server

- **Installed products**

To support local backup and secondary backup protection, the SnapManager for SAP client software must be installed on this component.

Backup schedule and retention strategy

The DBA wants to ensure that backups are available in case of a loss of data, in case of a disaster, and for regulatory reasons. This requires a carefully thought out retention policy for the various databases.

For the production payroll database, the DBA adheres to the following TechCo retention strategy:

Backup frequency	Retention duration	Backup time	Type of storage
Once daily	10 days	7 p.m.	Primary (local)
Once daily	10 days	7 p.m.	Secondary (archive)
Once weekly	52 weeks	Saturdays 1 a.m.	Secondary (archive)

- **Local backup advantages**

Daily local backup provides database protection, which is instantaneous, uses zero network bandwidth, uses a minimum of additional storage space, provides instantaneous restore, and provides finely-grained backup and restore capability.

Because the final weekly backups of the payroll database are retained for a minimum 52 weeks at a secondary storage site, there is no need to retain the daily backups any longer than 10 days.

- **Protected backup advantages**

Daily and weekly backups to secondary storage at a remote location guarantee that if the data at the primary storage site is damaged, the target database is still protected and can be restored from secondary storage.

The daily backups to secondary storage are made to protect against primary storage system damage. Because the final weekly backups of the payroll database are retained for a minimum 52 weeks, there is no need to retain the daily backups any longer than 10 days.

Workflow summary for local and secondary database backup

In this example, the DBA (using SnapManager) and the storage administrator (using the NetApp Management Console data protection capability) coordinate actions to configure local backup and secondary backup (also known as a protected backup) of the target database.

The sequence of actions carried out are summarized as follows:

- **Secondary resource pool configuration**

The storage administrator uses the NetApp Management Console data protection capability to configure a resource pool of storage systems at the secondary site that can be used to store the payroll database backup.

- **Secondary backup scheduling**

The storage administrator uses the NetApp Management Console data protection capability to configure secondary backup schedules.

- **Protection policy configuration**

The storage administrator uses the NetApp Management Console data protection capability to configure a secondary backup protection policy for the target database. The protection policy includes the schedules and specifies the base type of protection to implement backup protection (backup, mirror, or a combination of both), and names retention policies for primary data, secondary, and sometimes tertiary storage nodes.

- **Database profile configuration and protection policy assignment**

The DBA uses SnapManager to create or edit a profile of the target database that supports secondary backup. While configuring the profile, the DBA:

- Enables backup protection to secondary storage.
- Assigns the new protection policy, which was created in and retrieved from the NetApp Management Console data protection capability, to this profile.

Assigning the protection policy automatically includes the target database in a partially provisioned, but nonconformant the NetApp Management Console data protection capability dataset. When fully provisioned, the dataset configuration enables backup of the target database to secondary storage.

The dataset name uses this syntax: `smsap_hostname_databasename`, which translates to "smsap_prod01.sample.com_p01".

- **Secondary and tertiary storage provisioning**

The storage administrator uses the NetApp Management Console data protection capability to assign resource pools to provision the secondary and sometimes tertiary storage nodes (if the assigned protection policy specifies tertiary storage nodes).

- **Backup on local storage**

The DBA opens the profile with protection enabled in SnapManager and creates a full backup to local storage. The new backup shows in SnapManager as scheduled for protection, but not yet protected.

- **Secondary backup confirmation**

Because the backup was based on a protection-enabled profile, the backup is transferred to secondary according to the protection policy's schedule. The DBA uses SnapManager to confirm the transferral of the backup to secondary storage. After the backup has been copied to secondary storage, SnapManager changes the backup Protection State from "Not protected" to "Protected."

Protected backup configuration and execution

You must configure SnapManager and Protection Manager to support database backup to secondary storage. The database administrator and the storage administrator must coordinate their actions.

Database restore from backup

Use SnapManager for SAP to restore a local backup on primary storage

You can restore local backups that exist on primary storage. The entire process is performed using SnapManager for SAP.

About this task

You can also preview information about a backup restore process. You might want to do this to see information about restore eligibility of a backup. SnapManager analyzes data on a backup to determine whether the restore process can be completed by using the volume-based restore or the file-based restore method.

The restore preview shows the following information:

- Which restore mechanism (fast restore, 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.

In preview of the restore plan, SnapManager does not restore anything. The preview shows information up to 20 files.

If you want to preview a restore of data files but the database is not mounted, then SnapManager mounts the database. If the database cannot be mounted, then the operation fails and SnapManager returns the database to its original state.

Steps

1. From the **Repository** tree, right-click the backup you want to restore, and select **Restore**.
2. On the Restore and Recovery Wizard Welcome page, click **Next**.
3. On the **Restore Configuration Information** page, select **Complete Datafile/Tablespace Restore with Control Files**.
4. Click **Allow shutdown of database if necessary**.

SnapManager changes the database state, if necessary. For example, if the database is offline and it needs to be online, SnapManager forces it online.

5. On the **Recovery Configuration Information** page, click **All Logs**.

SnapManager restores and recovers the database to the last transaction and applies all required logs.

6. On the **Restore Source Location Configuration** page, view the information about the backup on primary and click **Next**.

If the backup exists only on primary storage, SnapManager restores the backup from the primary storage.

7. On the **Volume Restore Configuration Information** page, select **Attempt volume restore** to attempt volume restore method.

8. Click **Fallback to file-based restore**.

This allows SnapManager to use the file-based restore method if the volume restore method cannot be used.

9. Click **Preview** to see the eligibility checks for fast restore and information about mandatory and overridable checks.

10. On the **Perform Operation** page, verify the information you have entered, and click **Restore**.

11. To view details about the process, click **Operation Details**.

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.

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
- Split 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 email notification for a new profile

Configure email notification for an existing profile

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.



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.



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

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 <code>plugin.sh -check</code> 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	<p>The SnapManager server runs the <code>plugin.sh -describe</code> 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:</p> <ul style="list-style-type: none"> • <code>SM_PI_NAME</code>: Script name. You must provide a value for this parameter. • <code>SM_PI_DESCRIPTION</code>: Description of the script's purpose. You must provide a value for this parameter. • <code>SM_PI_CONTEXT</code>: Context in which the script should run—for example, root or orasid. You must provide a value for this parameter. • <code>SM_PI_TIMEOUT</code>: 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. • <code>SM_PI_PARAMETER</code>: 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. <p>The following is the sample output of the <code>Followup_activities</code> script.</p> <pre data-bbox="437 956 1408 1389"> 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>
execute	The SnapManager server runs the <code>plugin.sh -execute</code> 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
<i>SM_OPERATION_ID</i>	Specifies the ID of the current operation	string
<i>SM_PROFILE_NAME</i>	Specifies the name of the profile used	string
<i>SM_SID</i>	Specifies the system identifier of the database	string
<i>SM_HOST</i>	Specifies the host name of the database	string
<i>SM_OS_USER</i>	Specifies the operating system (OS) owner of the database	string
<i>SM_OS_GROUP</i>	Specifies the OS group of the database	string
<i>SM_BACKUP_TYPE</i>	Specifies the type of the backup (online, offline, or auto)	string
<i>SM_BACKUP_LABEL</i>	Specifies the label of the backup	string
<i>SM_BACKUP_ID</i>	Specifies the ID of the backup	string
<i>SM_BACKUP_RETENTION</i>	Specifies the retention period	string
<i>SM_BACKUP_PROFILE</i>	Specifies the profile used for this backup	string
<i>SM_ALLOW_DATABASE_SHUTDOWN</i>	Specifies if you want to start up or shut down the database if required you can use the -force option from the command-line interface.	boolean
<i>SM_BACKUP_SCOPE</i>	Specifies the scope of the backup (full or partial)	string
<i>SM_BACKUP_PROTECTION</i>	Specifies if backup protection is enabled	boolean
<i>SM_TARGET_FILER_NAME</i>	Specifies the target storage system name	string
	 If more than one storage system is used, then the storage system names must be separated by commas.	

Variables	Description	Format
<i>SM_TARGET_VOLUME_NAME</i>	<p>Specifies the target volume name</p> <p> The target volume name must be prefixed with storage device name, for example, <i>SM_TARGET_FILER_NAME/SM_TARGET_VOLUME_NAME</i>.</p>	string
<i>SM_HOST_FILE_SYSTEM</i>	Specifies the host file system	string
<i>SM_SNAPSHOT_NAMES</i>	<p>Specifies the Snapshot list</p> <p> 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.</p>	string array
<i>SM_ARCHIVE_LOGS_DIRECTORY</i>	<p>Specifies the archive logs directory</p> <p> If the archive logs are located in more than one directory, then the names of those directories are separated by commas.</p>	string array
<i>SM_REDO_LOGS_DIRECTORY</i>	<p>Specifies the redo logs directory</p> <p> If the redo logs are located in more than one directory, then the names of those directories are separated by commas.</p>	string array
<i>SM_CONTROL_FILES_DIRECTORY</i>	<p>Specifies the control files directory</p> <p> If the control files are located in more than one directory, then the names of those directories are separated by commas.</p>	string array
<i>SM_DATA_FILES_DIRECTORY</i>	<p>Specifies the data files directory</p> <p> If the data files are located in more than one directory, then the names of those directories are separated by commas.</p>	string array

Variables	Description	Format
<code>user_defined</code>	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
<code>SM_OPERATION_ID</code>	Specifies the ID of the current operation	string
<code>SM_PROFILE_NAME</code>	Specifies the name of the profile used	string
<code>SM_HOST</code>	Specifies the host name of the database	string
<code>SM_OS_USER</code>	Specifies the operating system (OS) owner of the database	string
<code>SM_OS_GROUP</code>	Specifies the OS group of the database	string
<code>SM_BACKUP_TYPE</code>	Specifies the type of the backup (online, offline, or auto)	string
<code>SM_BACKUP_LABEL</code>	Specifies the backup label	string
<code>SM_BACKUP_ID</code>	Specifies the backup ID	string
<code>SM_BACKUP_PROFILE</code>	Specifies the profile used for the backup	string
<code>SM_RECOVERY_TYPE</code>	Specifies the recovery configuration information	string
<code>SM_VOLUME_RESTORE_MODE</code>	Specifies the volume restore configuration	string
<code>SM_TARGET_FILER_NAME</code>	Specifies the target storage system name	string
	 If more than one storage system is used, then the storage system names must be separated by commas.	

Variables	Description	Format
<i>SM_TARGET_VOLUME_NAME</i>	<p>Specifies the target volume name</p> <p> The target volume name must be prefixed with storage device name, for example, <i>SM_TARGET_FILER_NAME/SM_TARGET_VOLUME_NAME</i>.</p>	string
<i>SM_HOST_FILE_SYSTEM</i>	Specifies the host file system	string
<i>SM_SNAPSHOT_NAMES</i>	<p>Specifies the Snapshot list</p> <p> 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.</p>	string array
<i>SM_ARCHIVE_LOGS_DIRECTORY</i>	<p>Specifies the archive logs directory</p> <p> If the archive logs are located in more than one directory, then the names of those directories are separated by commas.</p>	string array
<i>SM_REDO_LOGS_DIRECTORY</i>	<p>Specifies the redo logs directory</p> <p> If the redo logs are located in more than one directory, then the names of those directories are separated by commas.</p>	string array
<i>SM_CONTROL_FILES_DIRECTORY</i>	<p>Specifies the control files directory</p> <p> If the control files are located in more than one directory, then the names of those directories are separated by commas.</p>	string array
<i>SM_DATA_FILES_DIRECTORY</i>	<p>Specifies the data files directory</p> <p> If the data files are located in more than one directory, then the names of those directories are separated by commas.</p>	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
<i>SM_ORIGINAL_SID</i>	SID of the original database	string
<i>SM_ORIGINAL_HOST</i>	Host name associated with the original database	string
<i>SM_ORIGINAL_OS_USER</i>	OS owner of the original database	string
<i>SM_ORIGINAL_OS_GROUP</i>	OS group of the original database	string
<i>SM_TARGET_SID</i>	SID of the clone database	string
<i>SM_TARGET_HOST</i>	Host name associated with the clone database	string
<i>SM_TARGET_OS_USER</i>	OS owner of the clone database	string
<i>SM_TARGET_OS_GROUP</i>	OS group of the clone database	string
<i>SM_TARGET_DB_PORT</i>	Port of the target database	integer
<i>SM_TARGET_GLOBAL_DB_NAME</i>	Global database name of the target database	string
<i>SM_BACKUP_LABEL</i>	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

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.

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
- Clone split

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.



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.

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 /opt/NetApp/smsap/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 perform 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	<p>SnapManager for SAP CLI and GUI</p> <p> BR*Tools can list and restore only backups created by using BR*Tools.</p>
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.

Type of wizard	Tasks that can be performed
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:

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.

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 `sidadm`. These users must have access to the repository and the profile.

To set the credentials for a user, perform the following steps:

1. Log in as `sidadm`.
2. Set the repository credentials by entering the command:

`smsap credential set command`

3. Synchronize the profiles by entering the following command:

`smsap profile sync`

4. Set the password for the discovered profiles by entering the command:

`smsap credential set`

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	<ul style="list-style-type: none">• Key: retain• Value: unlimited hourly daily weekly monthly

Operations	Backup utility parameters
Fast restore	<ul style="list-style-type: none"> • Key: fast • Value: require override fallback off
Data Protection	<ul style="list-style-type: none"> • Key: protect • Value: yes no left blank <p> The different options mentioned for value parameter are similar to the <code>-protect</code> option used in command-line interface.</p>
Back up data from the secondary storage to a location set as per users preferences	<ul style="list-style-type: none"> • Key: preferred_backup_locations • Value: Data protection node names as a comma-separated list. • Default: No values set. <p> For adding a data protection policy node name, see About protection policy descriptions.</p>
Restore data from the secondary storage to a location to a preferred location	<ul style="list-style-type: none"> • Key: restore_from_nearest_backup_location • Value: yes no • Default: yes

The following table describes the different backup utility parameters for these operations:

Operation	Key	Value
Backup retention	fast=	fallback
Fast restore	retain=	hourly
Data protection	protect=	no

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/dbs/`. You can use this file to specify the default backup utility parameter (`.utl`) file used for BR*Tools commands.

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/dbs/` directory. If you store the file in another location, you need to supply the full path with the `-r` option. For example:

```
brbackup -r /opt/NetApp_fcp_price_10g_enterprise_inst_vol1/dbs/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.ut1`) 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.

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.ut1`) 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.

You can override the checks on a volume-based, fast restore using BRRESTORE by specifying `fast=override` in the `.ut1` file (the same file in which you specify the profile name if it is different from the database system identifier).

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.

Restoring a backup from a secondary storage to a different location

You can use the BRRESTORE or smsap restore command to restore a secondary backup to a different location. The new host must also be running SnapManager.

You must set the preferred_backup_locations and restore_from_nearest_backup_location parameters in the initsID.utl file to restore the secondary backup to a desired location.

1. Edit the initsID.utl file.
2. Add preferred_backup_locations and restore_from_nearest_backup_location in the file.



The PREFERRED_BACKUP_LOCATIONS environment variable overrides the value set for the preferred_backup_locations parameter in the initSID.utl file.

- If the value for preferred_backup_locations does not match any protection policy node names and restore_from_nearest_backup_location is set to **No**, the operation fails.
- If either the value for preferred_backup_locations is not set or restore_from_nearest_backup_location is set to **Yes**, the restore operation can be performed from the nearest backup location.

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 /opt/NetApp/smsap/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 /u/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



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] }
[-protect | -noprotect | -protectnow]
[-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.

You can create an online backup of a Real Application Clusters (RAC) database, as long as the primary is in the open state, or the primary is mounted and an instance is in the open state. You can use the **-force** option for online backups if the local instance is in the shutdown state, or no instance is in the open state. The version of Oracle must be 10.2.0.5; otherwise, the database will hang if any instance in the RAC is mounted.

- 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.

With an online RAC database backup, use the `-force` option if the local instance is in the shutdown state, or no instance is in the open state.



The version of Oracle must be 10.2.0.5; otherwise, the database will hang if any instance in the RAC is mounted.

- 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.

- **-protect | -noprotect | -protectnow**

Indicates whether the backup should be protected to secondary storage. The `-noprotect` option specifies that the backup should not be protected to secondary storage. Only full backups are protected. If neither option is specified, SnapManager protects the backup as the default if the backup is a full backup and the profile specifies a protection policy. The `-protectnow` option is applicable only for Data ONTAP operating in 7-Mode. The option specifies that the backup be protected immediately to secondary storage.

- **-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:

```
smssap backup create -profile SALES1 -full -online  
-label full_backup_sales_May -profile SALESDB -force -retain -daily  
Operation Id [8abc01ec0e79356d010e793581f70001] succeeded.
```

The smssap 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 14:31:27 SUCCESS FULL    ONLINE   EXISTS   backup1   DAILY
PROTECTED
2007-08-10 14:12:31 SUCCESS FULL    ONLINE   EXISTS   backup2   HOURLY
NOT PROTECTED
2007-08-10 10:52:06 SUCCESS FULL    ONLINE   EXISTS   backup3   HOURLY
PROTECTED
2007-08-05 12:08:37 SUCCESS FULL    ONLINE   EXISTS   backup4   UNLIMITED
NOT PROTECTED
2007-08-05 09:22:08 SUCCESS FULL    OFFLINE  EXISTS   backup5   HOURLY
PROTECTED
2007-08-04 22:03:09 SUCCESS FULL    ONLINE   EXISTS   backup6   UNLIMITED
NOT REQUESTED
2007-07-30 18:31:05 SUCCESS FULL    OFFLINE  EXISTS   backup7   HOURLY
PROTECTED
```

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
[-from-secondary {-copy-id id}]
[-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.

- **`-from-secondary -copy-id id`**

Mounts the backup from secondary storage. If this option is not specified, SnapManager mounts the backup from primary storage. You can use this option if the backup is freed.

You must specify the `-copy-id` option whenever you specify the `-from-secondary` option. If there is more than one backup on the secondary storage system, the `-copy-id` option is used to specify which backup copy on the secondary storage should be used to mount the backup.



If you are using Data ONTAP operating in 7-Mode, you must specify a valid value for the `-copy-id` option. However, if you are using clustered Data ONTAP, the `-copy-id` option is not required.

- **`-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.

 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
SMSAP-13046 [INFO ]: Operation GUID 8abc013111b9088e0111b908a7560001
starting on Profile S10_BACKUP
SMSAP-08052 [INFO ]: Beginning to connect mount(s)
[/oracle/S10_mirrlogs, /oracle/S10_sapdata] from logical snapshot
SMSAP_S10_BACKUP_S10_F_C_1_8abc013111a450480111a45066210001.
SMSAP-08025 [INFO ]: Beginning to connect mount
/oracle/S10_mirrlogs from snapshot
SMSAP_S10_BACKUP_S10_F_C_1_8abc013111a450480111a45066210001_0 of
volume saplog_S10.
SMSAP-08027 [INFO ]: Finished connecting mount /oracle/S10_mirrlogs from
snapshot
SMSAP_S10_BACKUP_S10_F_C_1_8abc013111a450480111a45066210001_0 of
volume saplog_S10.
SMSAP-08025 [INFO ]: Beginning to connect mount /oracle/S10_sapdata
from snapshot
SMSAP_S10_BACKUP_S10_F_C_1_8abc013111a450480111a45066210001_0 of
volume sapdata_S10.
SMSAP-08027 [INFO ]: Finished connecting mount /oracle/S10_sapdata
from snapshot
SMSAP_S10_BACKUP_S10_F_C_1_8abc013111a450480111a45066210001_0 of
volume sapdata_S10.
SMSAP-08053 [INFO ]: Finished connecting mount(s)
[/oracle/S10_mirrlogs, /oracle/S10_sapdata] from logical snapshot
SMSAP_S10_BACKUP_S10_F_C_1_8abc013111a450480111a45066210001.
SMSAP-13037 [INFO ]: Successfully completed operation: Backup Mount
SMSAP-13048 [INFO ]: Operation Status: SUCCESS
SMSAP-13049 [INFO ]: Elapsed Time: 0:01:00.981
Operation Id [8abc013111b9088e0111b908a7560001] succeeded.
```

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 | -from-secondary [-temp-volume temp_volume] [-
copy-id id]]
[-preview]
[-fast {-require | -override | -fallback | -off}]
[-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.
- **-tablespaces *list***: Restores only the specified tablespaces from the backup.

You must use spaces to separate the names in the list.

- **-files *list***: 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.

- `-from-secondary`: Restores the data from secondary storage.

You cannot use this option if the backup exists on primary storage; the primary backup must be freed before a backup can be restored from secondary storage. If you use a temporary volume, you must specify the volume by using the `-temp-volume` option.

You must specify the `-copy-id` option whenever you specify the `-from-secondary` option. If there is more than one backup on the secondary storage system, the `-copy-id` option is used to specify which backup copy on the secondary storage should be used for the restore operation.



If you are using Data ONTAP operating in 7-Mode, you must specify a valid value for the `-copy-id` option. However, if you are using clustered Data ONTAP, the `-copy-id` option is not required

When restoring from secondary storage, SnapManager first attempts to restore data directly from the secondary storage system to the primary storage system (without involving the host). If SnapManager cannot perform this type of restore (for example, if the files are not part of the file system), then SnapManager will fall back to a host-side file copy restore. SnapManager has two methods for performing a host-side file copy restore from secondary. The method that SnapManager selects is configured in the `smsap.config` file.

- Direct: SnapManager clones the data on secondary storage, mounts the cloned data from the secondary storage system to the host, and then copies data out of the clone into the active

environment.

This is the default secondary access policy.

- Indirect: SnapManager first copies the data to a temporary volume on primary storage, mounts the data from the temporary volume to the host, and then copies data out of the temporary volume into the active environment.

This policy should be used only if the host does not have direct access to the secondary storage system. Restores using the indirect method will take twice as long as the direct secondary access policy because two copies of the data are made.

The decision whether to use the direct or indirect method is controlled by the value of the `restore.secondaryAccessPolicy` parameter in the `smsap.config` configuration file.

- **-preview**

Displays the following information:

- Which restore mechanism (fast restore, 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.
 - The `-fast` option (`-require`, `-override`, `-fallback`, or `-off`) has significant impact on the output. 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.

- **-fast**

Enables you to choose the process to use in the restore operation. You can force SnapManager to use the volume-based fast restore process rather than other restore processes, if all mandatory restore eligibility conditions are met. If you are aware that a volume restore cannot be performed, you can also use this process to prevent SnapManager from conducting eligibility checks and the restore operation by using the fast restore process.

The `-fast` option includes the following parameters:

- `-require`: Enables you to force SnapManager to perform a volume restore, if all restore eligibility conditions are met.

If you specify the `-fast` option, but do not specify any parameter for `-fast`, SnapManager uses the `-require` parameter as a default.

- `-override`: Enables you to override the non-mandatory eligibility checks and perform the volume-based fast restore process.

◦ **-fallback**: Enables you to restore the database by using any method that SnapManager determines.

If you do not specify the **-fast** option, SnapManager uses the default `backup restore -fast fallback` option.

◦ **-off**: Enables you to avoid the time required to perform eligibility checks.

• **-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. For Real Application Clusters (RAC), you must include the **-force** option if SnapManager has to change the state of any RAC instance to a lower state.

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:

```
smasap 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 state: PROTECTED
  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: /mnt/ssys1/data/hsdb/sysaux01.dbf [ONLINE]
  ...
  Control Files:
  File: /mnt/ssys1/data/control03.ctl
  ...
  Archive Logs:
  File: /mnt/ssys1/data/archive_logs/2_131_626174106.dbf
  ...
  Host: Host1
  Filesystem: /mnt/ssys1/data
  File: /mnt/ssys1/data/hsdb/SMSAPBakCtl_1175283005231_0
  ...
  Volume: hs_data
  Snapshot: SMSAP_HSDBR_hsdb1_F_C_1_
  8abc013111a450480111a45066210001_0
  File: /mnt/ssys1/data/hsdb/SMSAPBakCtl_1175283005231_0
  ...
  Protected copies on Secondary Storage:
    14448939 - manow
    88309228 - graffe
```

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
```

```
SMSAP-13046 [INFO ]: Operation GUID 8abc013111b909eb0111b90a02f50001
starting on Profile SALES1
SMSAP-08028 [INFO ]: Beginning to disconnect connected mount(s)
[/u/user1/mnt/_mnt_ssys1_logs_SMSAP_SALES1_hsdb1_F_C_1_8abc013111a45048011
1a45066210001,
/u/user1/mnt/_mnt_ssys1_data_SMSAP_SALES1_hsdb1_F_C_1_8abc013111a45048011
1a45066210001].
SMSAP-08030 [INFO ]: Done disconnecting connected mount(s)
[/u/user1/mnt/_mnt_ssys1_logs_SMSAP_SALES1_hsdb1_F_C_1_8abc013111a45048011
1a45066210001,
/u/user1//mnt/_mnt_ssys1_data_SMSAP_SALES1_hsdb1_F_C_1_8abc013111a45048011
1a45066210001].
SMSAP-13037 [INFO ]: Successfully completed operation: Backup Unmount
SMSAP-13048 [INFO ]: Operation Status: SUCCESS
SMSAP-13049 [INFO ]: Elapsed Time: 0:00:33.715
Operation Id [8abc013111b909eb0111b90a02f50001] succeeded.
```

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 =
+SMSAP_1_1161675083835/smsmsap/datafile/data.277.582482539 ...
```

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
[-asminstance -asmusername asminstance_username -asmpassword
asminstance_password]
[-syspassword syspassword] [-reserve {yes | no | inherit}]
[-from-secondary {-copy-id id}
[-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.

- **`-asminstance`**

Specifies the credentials that are used to log in to the ASM instance.

- **`-asmusername asminstance_username`**

Specifies the user name that is used to log in to the ASM instance.

- **`-asmpassword asminstance_password`**

Specifies the password that is used to log in to ASM instance.

- **`-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 <code>-reserve</code> 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 <code>-reserve</code> option for a volume is set to yes, space is reserved for the full volume size within the aggregate.

- **`-from-secondary [-copy-id copy_id]`**

Specifies that SnapManager should clone a copy of a backup that has been protected to secondary storage. If this option is not specified, SnapManager clones the copy from primary storage.

You must specify the `-copy-id` option whenever you specify the `-from-secondary` option. If there is more than one backup on the secondary storage system, the `-copy-id` option is used to specify which backup copy on the secondary storage should be used for cloning.



If you are using Data ONTAP operating in 7-Mode, you must specify a valid value for the `-copy-id` option. However, if you are using clustered Data ONTAP, the `-copy-id` option is not required.

- **`-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  
/opt/<path>/smsap/clonespecs/sales1_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 used 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]  
[-asminstance -asmusername asminstance_username -asmpassword  
asminstance_password]  
[-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.

- **-asminstance**

Specifies the credentials that are used to log in to the Automatic Storage Management (ASM) instance.

- **-asmusername *asminstance_username***

Specifies the user name used to log in to the ASM instance.

- **-asmpassword *asminstance_password***

Specifies the password used to log in to ASM instance.

- **-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: /mnt/ssys1/data_clone
File: /mnt/ssys1/data_clone/hsdb/sysaux01.dbf
File: /mnt/ssys1/data_clone/hsdb/undotbs01.dbf
File: /mnt/ssys1/data_clone/hsdb/users01.dbf
File: /mnt/ssys1/data_clone/hsdb/system01.dbf
File: /mnt/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 following output shows information about a clone of a protected backup on secondary storage:

```

clone show -label clone_CLSTEST -profile
TEST_USER_NFSTEST_DIRMAC
Clone id:8abc01ec16514aec0116514af52f0001
Clone status: SUCCESS
Clone SID: CLSTEST
Clone label: clone_CLSTEST
Clone comment:comment_for_clone_CLSTEST
Clone start time: 2007-11-18 00:46:10
Clone end time: 2007-11-18 00:47:54
Clone host: dirmac
Filesystem: /ant/fish/bt_dirmac_nfs_clone
File: /ant/fish/bt_dirmac_nfs_clone/datafiles/sysaux01.dbf
File: /ant/fish/bt_dirmac_nfs_clone/datafiles/system01.dbf
File: /ant/fish/bt_dirmac_nfs_clone/datafiles/undotbs01.dbf
File: /ant/fish/bt_dirmac_nfs_clone/datafiles/users01.dbf
Backup id: 8abc01ec16514883011651488b580001
Backup label:full_backup
Backup SID: NFSTEST
Backup comment:
Backup start time: 2007-11-18 00:43:32
Backup end time: 2007-11-18 00:45:30
Backup host: dirmac
Storage System: fish (Secondary storage)
Volume: bt_dirmac_nfs
Snapshot: smsap_user_nfstest_b_nfstest_f_c_1_8abc01ec16511d6a0116511d735900
01_0
File: /ant/fish/bt_dirmac_nfs/archlogs/1_14_638851420.dbf
File: /ant/fish/bt_dirmac_nfs/datafiles/sysaux01.dbf
File: /ant/fish/bt_dirmac_nfs/datafiles/undotbs01.dbf
File: /ant/fish/bt_dirmac_nfs/archlogs/1_13_638851420.dbf
File: /ant/fish/bt_dirmac_nfs/archlogs_2/1_16_638851420.dbf
File: /ant/fish/bt_dirmac_nfs/datafiles/users01.dbf
File: /ant/fish/bt_dirmac_nfs/controlfiles/SMSAPBakCtl_1195361899651_2
File: /ant/fish/bt_dirmac_nfs/datafiles/system01.dbf

```

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_label_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 comment_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 split-delete` command

This command lets you delete a clone split operation cycle entry from a repository

database.

Syntax

```
smsap clone split-delete
-profile profile [-host hostname]
[-label split-label | -id split-id]
[-quiet | -verbose]
```

Parameters

- **-profile *profile***

Specifies the profile name of the clone.

- **-host *hostname***

Specifies the hostname in which the clone exists.

- **-label *split-label***

Specifies the label name generated by clone split start process.

- **-id *split-id***

Specifies the unique ID generated by clone split start process.

- **-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 smsap clone split-estimate command

This command enables you to view the clone split amount of storage consumed estimate.

Syntax

```
smsap clone split-estimate
-profile profile
[-host hostname]
[-label clone-label | -id clone-id]
[-quiet | -verbose]
```

Parameters

- **-profile *profile***

Specifies the profile name of the clone.

- **-host *hostname***

Specifies the hostname in which the clone exists.

- **-label *clone-label***

Specifies the label name generated by clone process.

- **-id *clone-id***

Specifies the unique ID generated by clone process.

- **-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 smsap clone split command

You can run the `clone split` command to split a clone. The split clone becomes independent of the original clone. SnapManager generates a new profile after the clone split operation and you can use this profile to manage the split clone.

Syntax

```

smsap clone split
-profile clone-profile
[-host hostname]
{-label clone-label | -id clone-id} [-split-label split-operation_label]
[-comment comment]
-new-profile new-profile-name [-profile-password new-profile_password]
getRepository -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} }]
-osaccount osaccount
-osgroup osgroup
[-retain
[-hourly [-count n] [-duration m] ]
[-daily [-count n] [-duration m] ]
[-weekly [-count n] [-duration m] ]
[-monthly [-count n] [-duration m] ]
[-profile-comment profile-comment]
[-snapname-pattern pattern]
[-protect [-protection-policy policy_name] | [-noprotect] ]
[-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 [-protection-policy policy_name | -noprotect]
[-include-with-online-backups | -no-include-with-online-backups] ]
[-dump]
[-quiet | -verbose]

```

Parameters

- **-profile clone-profile**

Specifies the profile name from which the clone is created.

- **-host *hostname***

Specifies the host name in which the clone exists.

- **-label *clone-label***

Specifies the label name generated by the clone operation.

- **-id *clone-id***

Specifies the unique ID generated by the clone operation.

- **-split-label *split-operation_label***

Specifies the label name generated by the clone operation.

- **-new-profile *new-profile_name***

Specifies the new profile name that SnapManager will generate after a successful split operation.

- **-profile-password *new-profile_password***

Specifies the password for the profile.

- **-repository**

Specifies the details of the database for the repository.

- **-dbname *repo_service_name***

Specifies the name of the database that stores the repository. You can use either the global name or system identifier.

- **-host *repo_host***

Specifies the name or IP address of the host computer on which the repository database resides.

- **-port *repo_port***

Specifies the Transmission Control Protocol (TCP) port number used to access 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 host on which the repository database resides.

- **-database**

Specifies the details of 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 system identifier.

- **-host *db_host***

Specifies the name or IP address of the host computer on which the database resides.

- **-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 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.

- **-osaccount *osaccount***

Specifies the name of the Oracle database user account. SnapManager uses this account to perform the Oracle operations such as startup and shutdown. It is typically the user who owns the Oracle software on the host, for example, orasid.

- **-osgroup *osgroup***

Specifies the name of the Oracle database group name associated with the orasid account.



The **-osaccount** and **-osgroup** variables are required for UNIX but not allowed for databases running on Windows.

- **-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.

For each retention class, either or both the retention count or retention duration might be specified. The duration is in units of the class (for example, hours for hourly, days for daily). For instance, if you specify 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.

- **-profile-comment *profile-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.

- **-protect -protection-policy *policy_name***

Specifies whether the backup should be protected to secondary storage.



If `-protect` is specified without `-protection-policy`, then the dataset will not have a protection policy. If `-protect` is specified and `-protection-policy` is not set when the profile is created, then it may be set later by the `smsap profile update` command or set by the storage administrator by using the Protection Manager's console.

- **-summary-notification**

Specifies the details for configuring summary email notification for multiple profiles under a repository database. SnapManager generates this email.

- **-notification**

Specifies the details for configuring email notification for the new profile. SnapManager generates this email. The email notification enables the database administrator to receive emails on the succeeded or failed status of the database operation that is performed by using this profile.

- **-success**

Specifies that email notification is enabled for a profile for when the SnapManager operation succeeds.

- **-email *email address 1* *email address 2***

Specifies the email address of the recipient.

- **-subject *subject-pattern***

Specifies the email subject.

- **-failure**

Specifies that email notification is enabled for a profile for when the SnapManager operation fails.

- **-separate-archivelog-backups**

Specifies that the archive log backup is separated from the datafile backup. This is an optional parameter, which you can provide while creating the profile. After the backups are separated by using this option, you can either create datafiles-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, or monthly).

- **protect [-protection-policy *policy_name*] | -noprotect**

Specifies that the archive log files are protected based on the archive log protection policy.

Specifies that the archive log files are not protected by using the **-noprotect** option.

- **-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 not collected after the successful profile create operation.

- **-quiet**

Displays only error messages in the console. The default setting displays error and warning messages.

- **-verbose**

Displays error, warning, and informational messages in the console.

The smsap clone split-result command

Syntax

This command lets you view the result of the clone split process.

```
smsap clone split-result
-profile profile
[-host hostname]
[-label split-label | -id split-id]
[-quiet | -verbose]
```

Parameters

- **-profile *profile***

Specifies the profile name of the clone.

- **-host *hostname***

Specifies the hostname in which the clone exists.

- **-label *split-label***

Specifies label name generated by clone split start process.

- **-id *split-id***

Specifies unique ID generated by clone split start process.

- **-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 smsap clone split-stop command

This command stops the running clone split process.

Syntax

```
smsap clone split-stop
-profile profile
[-host hostname]
[-label split-label | -id split-id]
[-quiet | -verbose]
```

Parameters

- **-profile *profile***

Specifies the profile name of the clone.

- **-host *hostname***

Specifies the hostname in which the clone exists.

- **-label *split-label***

Specifies the label name generated by clone process.

- **-id *split-id***

Specifies the unique ID generated by clone process.

- **-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 smsap clone split-status command

This command lets you know the progress of running split process.

Syntax

```
smsap clone split-status
-profile profile
[-host hostname]
[-label split-label | -id split-id]
[-quiet | -verbose]
```

Parameters

- **-profile *profile***

Specifies the profile name of the clone.

- **-host *hostname***

Specifies the hostname in which the clone exists.

- **-label *split-label***

Specifies the label name generated by clone process.

- **-id *split-id***

Specifies the unique ID generated by clone process.

- **-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 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.



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.

Example

This example creates a profile by including the `profile create` command in `command.txt` located at `/tmp`. You can then run the `smsap cmdfile` command.

The text file contains the following information:

```
profile create -profile SALES1 -repository -dbname SNAPMGRR
    -login -username server1_user -password ontap -port 1521 -host server1
    -database -dbname SMSMSAP -sid SMSMSAP -login -username sys -password
    oracle -port 1521
    -host Host2 -osaccount oracle -osgroup db2
```

You can now create the profile by running the `smsap cmdfile` command with the `command.txt` file:

```
smsap cmdfile -file /tmp/command.txt
```

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:

- -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 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.

```
smssp 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.

```
smssp 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:

- **-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 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 {-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 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* | -id *id***

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 {-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 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_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 {-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

```
smssap history purge -profile -name PROFILE1 -operation -operations backup  
-verbose
```

The smssap 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
-database 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.

- **-database *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 {-count *retain_count* | -daily *daily_count* | -monthly *monthly_count* | -weekly *weekly_count*}**

Specifies the retention class of the create backup, verify backup, restore and recover, and create and split 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|protection-policy  
|repository|system|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
- protection policy
- 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
-database 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.

- **-database *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:

```
smssap 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:/userhomedirectory/.netapp/smsap/3.3/smsap_dump_8abc01ec0e78f3e2010e7
8f3fdd00001.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:

```
smssap operation show -id 8ac861781d0ac992011d0ac999680001 -profile CER
```

Operation Attempted

Operation ID: 8ac861781d0ac992011d0ac999680001
Type: Backup
For profile: CER
With Force: No
Label: 081017180043
Comments: BRBACKUP

Operation Runtime Information

Status: SUCCESS
Start date: 2008-10-16 18:01:00 IST
End date: 2008-10-17 18:01:26 IST
Client user: oracle
Effective user: oracle
By schedule: none

Host

Host Run upon: lnx225-248.lab.eng.org.com
Process ID: 29096
SnapManager version: 3.3

Repository

Connection: krishna@smsaprep/10.72.225.155:1521
Repository version: 62

Error messages

The operation did not generate any error messages.

Resources in use

Snapshots:
f270-225-
57:/vol/f270_lnx225_248_10gr2_sap_oracle_cer: smsap_cer_cer1_f_c_2_8ac86178
1d0ac992011d0ac999680001_0
Storage components:
/sapbackup/backCER1.log (File)
/sapbackup/bdzbalta.anf (File)
/sapreorg/spaceCER1.log (File)
/mnt/oracle/CER (File System)
...

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 root 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 plugin check command

SnapManager enables you to install and use custom scripts for various operations. SnapManager offers backup, restore, and clone plug-ins to automate your custom scripts before and after the backup, restore, and clone operations. Before you use the backup, restore, and clone plug-in, you can run the plugin check command to verify the installation of plug-in scripts. Custom scripts are stored in three directories: policy (for scripts that should always be run before the backup, restore, or clone operation occurs), pre (for preprocessing scripts), and post (for post-processing scripts).

Syntax

```
smsap plugin check
-osaccount os_db_user_name
```

Parameter

- **-osaccount**

Specifies the operating system (OS) database user name. If you do not enter the -osaccount option,

SnapManager checks the plug-in scripts as root user rather than for a specific user.

Example

The following example shows that the plugin check command found the policy1 custom script stored in the policy directory as an executable. The example also shows that the two other custom scripts stored in the pre directory return no error messages (shown with a status of 0); however, the fourth custom script (post-plugin1), which was found in the post directory, contains errors (shown with a status of 3).

```
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
<installdir>/plugins/backup/policy
OK: 'policy1' is executable
<installdir>/plugins/backup/pre
OK: 'pre-plugin1' is executable and returned status 0
OK: 'pre-plugin2' is executable and returned status 0
<installdir>/plugins/backup/post
ERROR: 'post-plugin1' is executable and returned status 3
<installdir>/plugins/restore/policy
OK: 'policy1' is executable
<installdir>/plugins/restore/pre
OK: 'pre-plugin1' is executable and returned status 0
OK: 'pre-plugin2' is executable and returned status 0
<installdir>/plugins/restore/post
ERROR: 'post-plugin1' is executable and returned status 3
Command complete.
```

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]
[-asminstance -asmusername asminstance_username -asmpassword
asminstance_password]
[-rman {-controlfile | {-login
-username rman_username -password rman_password\}
-tnsname rman_tnsname\}}
[-osaccount osaccount] [-osgroup osgroup]
[-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]
[-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 [-protection-policy policy_name | -noprotect
[-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.

For example, if you are using Oracle Real Application Clusters (RAC), you must specify the system identifier of the RAC instance on the RAC node from which SnapManager is executed.

- **-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.

- **-asminstance**

Specifies the credentials that are used to log in to the Automatic Storage Management (ASM) instance.

- **-asmusername *asminstance_username***

Specifies the user name used to log in to the ASM instance.

- **-asmpassword *asminstance_password***

Specifies the password used to log in to ASM instance.

- **-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.

- **-osaccount *osaccount***

Specifies the name of the Oracle database user account. SnapManager uses this account to perform the Oracle operations such as startup and shutdown. It is typically the user who owns the Oracle software on the host, for example, `orasid`.

- **-osgroup *osgroup***

Specifies the name of the Oracle database group name associated with the `orasid` account.

- **-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.

- **-protect -protection-policy *policy***

Indicates whether the backup should be protected to secondary storage.



If `-protect` is specified without `-protection-policy`, then the dataset will not have a protection policy. If `-protect` is specified and `-protection-policy` is not set when the profile is created, then it may be set later by `smsap profile update` command or set by the storage administrator through Protection Manager's console.

- **-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).

- **protect [-protection-policy *policy_name* | -noprotect]**

Specifies to protect the archive log files based on the archive log protection policy.

The -noprotect option specifies not to protect 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-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:

```
smssap 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 smssap 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 destroy command

This command deletes the split clone (database) along with the profile generated by SnapManager during the clone split process.

Syntax

```
smsap profile destroy
-profile profile
[-host hostname]
[-quiet | -verbose]
```

Parameters

- **-profile *profile***

Specifies the profile that SnapManager generates after a successful clone split process.

- **-host *hostname***

Specifies the hostname in which the split clone exists.

- **-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 deletes the profile named SALES1.

```
  smsap profile destroy -profile SALES1
```

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:/userhomedirectory/.netapp/smsap/3.3.0/smsap_dump_SALES1_hostname.jar
```

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: swagrahm
  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.

Example

The following example shows the details of the profile:

```

smsap profile show -profile TEST_RBAC_DP_PROTECT
Profile name: TEST_RBAC_DP_PROTECT
Comment:
Target database:
  Database name: racb
  SID: racb1
  Host: saal
  Port: 1521
  Username: sys
  Password: *****
Repository:
  Database name: smsaprep
  SID: smsaprep
  Host: elbe.rtp.org.com
  Port: 1521
  Username: smsapsaal
  Password: *****
RMAN:
  Use RMAN via control file
Oracle user account: oracle
Oracle user group: dba
Snapshot Naming:
  Pattern: smsap_{profile}_{db-sid}_{scope}_{mode}_{smid}
  Example:
    smsap_test_rbac_dp_protect_racb1_f_h_1_8abc01e915a55ac50115a55acc8d0001_0
Protection:
  Dataset: smsap_saal_racb
  Protection policy: Back up
  Conformance status: CONFORMANT
Local backups to retain:
  Hourly: 4 copies
  Daily: 7 day(s)
  Weekly: 4 week(s)
  Monthly: 12 month(s)

```

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]
  [-asminstance -asmusername asminstance_username -asmpassword
  asminstance_password]
  [{-rman {-controlfile | {-login
  -username rman_username
  -password rman_password }
  [-tnsname tnsname]}} | {-remove-rman}
  -osaccount osaccount
  -osgroup osgroup
  [-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]
  [-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 [-protection-policy policy_name] | [-noprotect]
  [-include-with-online-backups | -no-include-with-online-backups]]
  [-dump]
  [-quiet | -verbose]]]
```

Parameters

If protection policy is set on the profile, you cannot change the policy using SnapManager. You must change the policy using the Protection Manager's console.

- **-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.

For example, if you are using Oracle Real Application Clusters (RAC), you must specify the SID system identifier of the RAC instance on the RAC node from which SnapManager is executed.

- **-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.

- **-asminstance**

Specifies the credentials that are used to log in to the Automatic Storage Management (ASM) instance.

- **-asmusername *asminstance_username***

Specifies the user name used to log in to the ASM instance.

- **-asmpassword *asminstance_password***

Specifies the password used to log in to ASM instance.

- **-osaccount *osaccount***

Specifies the name of the Oracle database user account. SnapManager uses this account to perform the Oracle operations such as startup and shutdown. It is typically the user who owns the Oracle software on the host, for example, orasid.

- **-osgroup *osgroup***

Specifies the name of the Oracle database group name associated with the orasid account.

- **-retain [-hourly [-count *n*] [-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.

- **-protect [-protection-policy policy_name] | [-noprotect]**

Indicates whether the backup should be protected to secondary storage or not.



If `-protect` is specified without `-protection-policy`, then the dataset will not have a protection policy. If `-protect` is specified and `-protection-policy` is not set when the profile is created, then it may be set later by `smsap profile update` command or set by the storage administrator by using the Protection Manager's console.

The `-noprotect` option specifies not to protect the profile to secondary storage.

- **-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).

- **-protect [-protection-policy policy_name] | -noprotect**

Specifies that the archive log files are protected based on the archive log protection policy.

Specifies that the archive log files are not protected by using the **-noprotect** option.

- **-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:
CER/hostA.rtp.com" passed.
[ INFO] SMSAP-13506: SQLPlus verification for database SID "CER" passed.
Environment: [ORACLE_HOME=/u02/app/oracle/product/11.2.0.2]
[ INFO] SMSAP-07431: Saving starting state of the database: CER(OPEN) .
[ INFO] SMSAP-07431: Saving starting state of the database: CER(OPEN) .
[ INFO] SD-00016: Discovering storage resources for
/vol/hostA_sap_datavol_CER.
[ INFO] SD-00017: Finished storage discovery for /vol/
hostA_sap_datavol_CER.
[ INFO] SD-00016: Discovering storage resources for
/vol/hostA_sap_datavol_CER.
[ INFO] SD-00017: Finished storage discovery for /vol/
hostA_sap_datavol_CER.
[ INFO] SD-00016: Discovering storage resources for
/vol/hostA_sap_cntrlvol_CER.
[ INFO] SD-00017: Finished storage discovery for
/vol/hostA_sap_cntrlvol_CER.
```

```
[ INFO] SD-00016: Discovering storage resources for  
/vol/hostA_sap_redovol_CER.  
[ INFO] SD-00017: Finished storage discovery for  
/vol/hostA_sap_redovol_CER.  
[ INFO] SD-00016: Discovering storage resources for  
/vol/hostA_sap_archivevol_CER.  
[ INFO] SD-00017: Finished storage discovery for  
/vol/hostA_sap_archivevol_CER.  
[ INFO] SD-00040: Beginning to discover filesystem(s) upon host volume  
group hostA-3_SdDg.  
[ INFO] SD-00041: Finished discovering filesystem(s) upon host volume  
group hostA-3_SdDg.  
[ INFO] SD-00040: Beginning to discover filesystem(s) upon host volume  
group hostA-2_SdDg.  
[ INFO] SD-00041: Finished discovering filesystem(s) upon host volume  
group hostA-2_SdDg.  
[ INFO] SD-00040: Beginning to discover filesystem(s) upon host volume  
group hostA_s_SdDg.  
[ INFO] SD-00041: Finished discovering filesystem(s) upon host volume  
group hostA_s_SdDg.  
[ INFO] SD-00040: Beginning to discover filesystem(s) upon host volume  
group hostA-1_SdDg.  
[ INFO] SD-00041: Finished discovering filesystem(s) upon host volume  
group hostA-1_SdDg.  
[ WARN] SMSAP-05071: Database profile profileA is not eligible for fast  
restore: Restore Plan:
```

Preview:

The following components will be restored completely via: host side
file copy restore

```
/vol/hostA_sap_datavol_CER/CER/sapdata1/oradata/CER/sysaux01.dbf  
/vol/hostA_sap_datavol_CER/CER/sapdata1/oradata/CER/system01.dbf  
/vol/hostA_sap_datavol_CER/CER/sapdata1/oradata/CER/undotbs01.dbf  
/vol/hostA_sap_datavol_CER/CER/sapdata1/oradata/CER/users01.dbf
```

Analysis:

The following reasons prevent certain components from being restored
completely via: storage side file system restore

* Files in file system /vol/hostA_sap_datavol_CER not part of the
restore scope will be reverted.

* File systems in volume group hostA-1_SdDg not part of the restore
scope will be reverted: [/vol/hostA_sap_datavol_CER]

Components not in restore scope:

```
/vol/hostA_sap_datavol_CER/CER/sapdata1/cfgtoollogs/catbundle/catbundle_PS
```

```
U_CER_APPLY_2011Dec15_00_52_21.log
```

```
/vol/hostA_sap_datavol_CER/CER/sapdata1/cfgtoollogs/catbundle/catbundle_PS  
U_CER_GENERATE_2011Dec15_00_52_16.log
```

Components to restore:

```
/vol/hostA_sap_datavol_CER/CER/sapdata1/oradata/CER/sysaux01.dbf  
/vol/hostA_sap_datavol_CER/CER/sapdata1/oradata/CER/system01.dbf  
/vol/hostA_sap_datavol_CER/CER/sapdata1/oradata/CER/undotbs01.dbf  
/vol/hostA_sap_datavol_CER/CER/sapdata1/oradata/CER/users01.dbf
```

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

```
[ INFO] SMSAP-07433: Returning the database to its initial state: CER  
(OPEN).  
[ INFO] SMSAP-13048: Profile Verify Operation Status: SUCCESS  
[ INFO] SMSAP-13049: Elapsed Time: 0:01:17.857  
Operation Id [Nab0240e8200dae6f17ecf21060bc6de8] succeeded.
```

The smsap protection-policy command

You can run the **protection-policy** command to list the protection policies that can be applied to a profile. The protection policy can be applied when a new profile is created or an existing profile is updated. You can also set the protection policy for the profile using the Protection Manager console.

Syntax

```
  smsap protection-policy list
```



The Protection Manager and SnapDrive must be installed on the server for you to use this command.

Parameters

- **list**

Displays the list of protection policies that can be set on a profile.

Example

The following example lists the protection policies that can be set to a profile:

```
  smsap protection-policy list
```

Back up
Back up, then mirror
Chain of two mirrors
DR Back up
DR Back up, then mirror
DR Mirror
DR Mirror and back up
DR Mirror and mirror
DR Mirror, then back up
DR Mirror, then mirror
Local backups only
Mirror
Mirror and back up
Mirror to two destinations
Mirror, then back up
No protection
Partial-volume Mirror
Remote backups only

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
-database 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

```
smssap repository update
-repository
-database 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 <em>profile_name</em>
[-full{-auto | -online | -offline}
[-retain <em>-hourly</em> | <em>-daily</em> | <em>-weekly</em> | <em>-
monthly</em> | <em>-unlimited</em> [-verify]] |
-data [[-files <em>files</em> [<em>files</em>] |
-tablespaces <em>tablespaces</em> [<em>tablespaces</em>] {-auto | -online |
-offline}
[-retain <em>-hourly</em> | -daily | <em>-weekly</em> | <em>-monthly</em>
| <em>-unlimited</em> [-verify]] |
[-archivelogs}]
[-label <em>label</em>]
[-comment <em>comment</em>]
[-protect | -noprotect | -protectnow] [-backup-dest <em>path1</em> [ ,
<em>path2</em>]
[-exclude-dest <em>path1</em> [ , <em>path2</em>]] [-prunelogs {-all | -
until-scn <em>until-scn</em> | -until -date <em>yyyy-MM-
dd:HH:mm:ss</em>} | -before {-months | -days | -weeks | -hours}}
-prune-dest <em>prune_dest1</em>, [<em>prune_dest2</em>]] -schedule-name
<em>schedule_name</em>
[-schedule-comment <em>schedule_comment</em>] -interval {<em>-hourly</em>
| <em>-daily</em> | <em>-weekly</em> | <em>-monthly</em> | <em>-
onetimeonly</em>}
-cronstring <em>cron_string</em>-start-time {<em>start_time < yyyy-MM-dd
HH:mm</em>>}
-runasuser <em>runasuser</em>
[-taskspec <em>taskspec</em>]-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.

You can create an online backup of a Real Application Clusters (RAC) database, as long as the primary is in the open or mounted state and an instance is in the open state. You can use the **-force** option for online backups if the local instance is in the shutdown state, or no instance is open.

- If the local instance is in the shutdown state and at least one instance is open, you can use the `-force` option to change the local instance to mounted.
- If no instance is in open state, you can use the `-force` option to change the local instance to open.

- **`-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.

With an online RAC database backup, use the `-force` option if the local instance is in shutdown state, or no instance is open.



The version of Oracle must be 10.2.0.5; otherwise, the database will hang if any instance in the RAC is mounted.

- 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.

- **-protect | -noprotect | -protectnow**

Indicates whether the backup should be protected to secondary storage. The `-noprotect` option specifies that the backup should not be protected to secondary storage. Only full backups are protected. If neither option is specified, SnapManager protects the backup as the default if the backup is a full backup and the profile specifies a protection policy. The `-protectnow` option is applicable only for Data ONTAP operating in 7-Mode. The option specifies that the backup be protected immediately to secondary storage.

- **-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 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 <em>profile_name</em>
-schedule-name <em>schedule_name</em> [-schedule-comment
<em>schedule_comment</em>]
-interval {<em>-hourly</em> | <em>-daily</em> | <em>-weekly</em> | <em>-
monthly</em> | <em>-onetimeonly</em>}
-cronstring <em>cron_string</em> -start-time {<em>start_time < yyyy-MM-dd
HH:mm</em>}
-runasuser <em>runasuser</em> [-taskspec <em>taskspec</em>] -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-RTP070LD

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:/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 [N4f4e910004b36cfecce74c710de02e44] 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 <code>lsnrctl status</code> command. Ensure that the database instance is registered with the listener.
Is the storage visible?	Run the <code>snapdrive storage show -all</code> command.
Is the storage writable?	Edit a file in the mountpoint that you just created. Use the <code>touch filename</code> command. If the file is created, then your storage is writable. You must ensure that the storage is writable by the user that SnapManager runs as (for example, as root on UNIX).
Is the SnapManager server running?	<p>Run <code>smsap_server status</code> and try to start the server by using the <code>smsap_server start</code> command.</p> <p>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.</p> <p>To start the SnapManager server, enter the following command: <code>smsap_server start</code>.</p>
Are all the components required to run SnapManager set up correctly?	Run the <code>smsap system verify</code> command to verify that SnapDrive is set up correctly.
Do you have the correct version of SnapManager?	Use the <code>smsap version</code> 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?	<p>SnapManager records all log entries into one set of rotating log files. The log files are found at <code>/var/log/smsap</code>.</p> <p>The log files are found at <code>C:\program_files\NetApp\SnapManager for SAP\logs</code>.</p> <p>It might also be helpful to look at the logs in the following location:</p> <p><code>/usr_home/.netapp/smsap/3.3.0/log/</code></p> <p>Each operation log is written to its own log file of the form <code>smsap_of_date_time.log</code>.</p>

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?	<p>The <code>smsap.config</code> file enables you to exclude certain archive log files. For UNIX, the files are at the following location: <code>/opt/NetApp/smsap/properties/smsap.config</code></p> <p>Use the format mentioned in the file to exclude the local archive logs. For additional information, see the “Setting configuration properties” topic.</p> <p>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.</p> <p>You can also exclude the archive log destinations while creating a backup from the SnapManager GUI.</p>
Do you have a FlexClone license if you are using SnapManager with NFS databases?	<p>A FlexClone license is required to take full advantage of SnapManager with NFS databases. SnapManager uses the FlexClone feature to accomplish these tasks:</p> <ul style="list-style-type: none"> • Mount backups of NFS databases • Verify backups of NFS databases • Clone NFS databases
Were you unable to connect to the repository?	<p>If connecting to a repository fails, run the <code>lsnrctl status</code> 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.</p> <p>Also, check whether the repository instance is running by running the <code>ps -eaf grep instance - name</code> command.</p>
Can system resolve the host name?	<p>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 <code>/etc/hosts</code>:</p> <p><code>xxx.xxx.xxx.xxx hostname IP address</code></p>
Is SnapDrive running?	<p>Check whether the SnapDrive daemon is running:</p> <p><code>-snapdrived status</code></p> <p>If the daemon is not running, a message appears indicating that there is a connection error.</p>

Issue-driven question	Possible solution
Which storage systems are configured to be accessed with SnapDrive?	<p>Run the command:</p> <pre data-bbox="479 234 850 264">-snapdrive config list</pre>
How can SnapManager GUI performance be improved?	<ul data-bbox="496 354 1498 413" style="list-style-type: none"> • Ensure that you have valid user credentials for the repository, profile host, and profile. <p>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”.</p> <ul data-bbox="496 650 850 680" style="list-style-type: none"> • Close the unused profiles. <p>If the number of profiles that you have opened is more, the SnapManager GUI performance slows down.</p> <ul data-bbox="496 815 1405 874" style="list-style-type: none"> • Check whether you enabled Open On Startup in the User Preferences window under the Admin menu, from the SnapManager GUI. <p>If this is enabled, then the user configuration (<code>user.config</code>) file available at <code>/root/.netapp/smsap/3.3.0/gui/state</code> is displayed as <code>openOnStartup=PROFILE</code>.</p> <p>Because Open On Startup is enabled, you must check for recently opened profiles from the SnapManager GUI, using <code>lastOpenProfiles</code> in the user configuration (<code>user.config</code>) file:</p> <pre data-bbox="528 1170 1286 1199">lastOpenProfiles=PROFILE1,PROFILE2,PROFILE3,...</pre> <p>You can delete the profile names listed and always keep a minimum number of profiles as open.</p> <ul data-bbox="496 1343 1442 1402" style="list-style-type: none"> • The protected profile takes more time to refresh than the profile that is not protected. <p>The protected profile is refreshed at a time interval, based on the value specified in the <code>protectionStatusRefreshRate</code> parameter of the user configuration (<code>user.config</code>) file.</p> <p>You can increase the value from the default value (300 seconds) so that the protected profiles are refreshed only after specified time interval.</p> <ul data-bbox="496 1681 1426 1776" style="list-style-type: none"> • Before installing the new version of SnapManager on the UNIX-based environment, delete the SnapManager client-side entries available at the following location: <pre data-bbox="528 1812 747 1841">/root/.netapp</pre>

Issue-driven question	Possible solution
<p>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.</p>	<p>You need to wait until the SnapManager GUI gets refreshed, and then check for the backup status.</p>
<p>What would you do when the Oracle database is not set in English?</p>	<p>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:</p> <ol style="list-style-type: none"> <li data-bbox="494 830 1106 988">1. Add the following under the initial comments in <code>/etc/init.d/smsap_server</code> <ul style="list-style-type: none"> <li data-bbox="551 914 975 946">◦ <code>NLS_LANG=American_America</code> <li data-bbox="551 963 812 994">◦ <code>export NLS_LANG</code> <li data-bbox="494 1015 1310 1079">2. Restart the SnapManager server using the following command: <code>smsap_server restart</code> <p> If the login scripts such as <code>.bash_profile</code>, <code>.bashrc</code>, and <code>.cshrc</code> for the Oracle user is set to <code>NLS_LANG</code>, you must edit the script to not overwrite <code>NLS_LANG</code>.</p>

Issue-driven question	Possible solution
<p>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?</p>	<ol style="list-style-type: none"> 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. <p>The schedule file names can be in the following formats:</p> <ul style="list-style-type: none"> ◦ repository#repo_username#repository_database_name#repository_host#repo_port ◦ repository-repo_usernamerepository_database_name-repository_host-repo_port <p> You must ensure that you delete the schedule file in the format that matches the repository details.</p> <ol style="list-style-type: none"> 3. Restart the SnapManager server. 4. Open other profiles under the same repository from the SnapManager GUI to ensure that you do not miss any schedule information of those profiles.
<p>What would you do when the SnapManager operation fails with credential file lock error?</p>	<p>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.</p> <p>Configure the following parameters in the smsap.config file depending on the frequency of simultaneous operations:</p> <ul style="list-style-type: none"> • fileLock.retryInterval = 100 milliseconds • fileLock.timeout = 5000 milliseconds <p> The values assigned to the parameters must be in milliseconds.</p>
<p>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?</p>	<p>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.</p> <ol style="list-style-type: none"> 1. Add the following parameters in the smsap.config file: <ul style="list-style-type: none"> ◦ operation.heartbeatInterval = 5000 ◦ operation.heartbeatThreshold = 5000 The default value assigned by SnapManager is 5000. 2. Assign the new values to these parameters. <p> The values assigned to the parameters must be in milliseconds.</p> <ol style="list-style-type: none"> 3. Restart the SnapManager server and perform the operation again.

Issue-driven question	Possible solution
What to do when you encounter a heap-space issue?	<p>When you encounter a heap-space issue during SnapManager for SAP operations, you must perform the following steps:</p> <ol style="list-style-type: none"> 1. Navigate to the SnapManager for SAP installation directory. 2. Open the <code>launchjava</code> file from the <code>installationdirectory/bin/launchjava</code> path. 3. Increase the value of the <code>java -Xmx160m</code> Java heap-space parameter. <p>For example, you can increase the default value of 160m to 200m.</p> <p> 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.</p>
What would you do if you cannot use the protected backups to restore or clone?	<p>This issue is observed if you were using SnapManager 3.3.1 with clustered Data ONTAP and have upgraded to SnapManager 3.4. The backups were protected using post-scripts in SnapManager 3.3.1. From SnapManager 3.4, the backups are protected using either <code>SnapManager_cDOT_Mirror</code> or <code>SnapManager_cDOT_Vault</code> policies which are selected while creating a profile. After upgrading to SnapManager 3.4, you might still be using the old profiles and thus backups are protected using backup scripts, but you cannot use them for restore or cloning using SnapManager.</p> <p>You must update the profile and select either <code>SnapManager_cDOT_Mirror</code> or <code>SnapManager_cDOT_Vault</code> policy and delete the post-script that was used for data protection in SnapManager 3.3.1.</p>
What would you do if scheduled backups are not getting protected (SnapVault)?	<p>After upgrading to SnapManager 3.4 and updating the profile to use <code>SnapManager_cDOT_Vault</code> policy for protection, you must delete old backup schedules and create new schedules to specify the SnapVault label while creating the schedule.</p>

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
- Automatic Storage Management (ASM) instance OPatch local inventory details
- Storage system version
- Oracle oratab file
- Oracle listener status
- Oracle network configuration files (`listener.ora` and `tnsnames.ora`)
- Repository database Oracle version
- Target database type (stand alone or Real Application Clusters (RAC))
- 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 ASM instance version
- 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
- Kernel version
- FSTAB

- Protocol used by Snapdrive
- Multipath environment
- RAC
- Supported volume manager
- Operations Manager version
- Supported file system
- Host utilities version
- BACKINT interface versions
- BR tool version
- Patch level
- Output of the `system verify` command
- Output of the `sdconfcheck` command

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.

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	<code>smsap_dump_ <i>operation-id</i>.jar</code>
Operation dump command with no operation ID	<code>smsap operation dump -profile VH1 -verbose</code> The following output is displayed: <pre>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: /oracle/VH1/<path>/smsap/3.3.0/smsap_dump_VH1_kaw.rtp.foo. com.jar</pre>
System dump command	<code>smsap_dump_ <i>host-name</i>.jar</code>
Profile dump command	<code>smsap_dump_ <i>profile-name_host-name</i>.jar</code>

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

- Splitting 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:

- `/var/log/smsap`

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.



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 <code>USE_DB_RECOVERY_FILE_DEST</code> .	When the archive destination is referring to <code>USE_DB_RECOVERY_FILE_DEST</code> , 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
<p>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.</p>	<p>This issue occurs if Oracle listener fails to connect to the database.</p>	<p>If you are using SnapManager GUI to clone a backup, perform the following actions:</p> <ol style="list-style-type: none"> 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 <i>local_listener</i> and click OK. 7. Select the Override Default check box for the <i>local_listener</i> row. 8. Click any parameter, then double-click the <i>local_listener</i> 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. <p>If you are using CLI to clone a backup, you must include the following information in the <parameters> tag of the clone specification file:</p>

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.	<ul style="list-style-type: none"> • 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 <i>Creating a clone</i> to meet the prerequisites.
SnapManager fails to generate a new profile after the clone split operation and the user does not know if the new profile is created.	SnapManager fails to prompt if a new profile is not created after the clone split operation. Because the prompt is not displayed, you might assume that the profile is created.	From the SnapManager command-line interface (CLI), enter the <code>clone split-result</code> command to view the detailed result of the clone split operation.
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).
Cloning a backup to a remote host fails with the following error message <code>Error: Access is denied.</code>	While mounting, if the IP address of the host is provided to the <code>snap mount</code> command, the cloning operation might fail. This issue occurs if the host on which the database resides is in workgroup while the remote host is in domain, or vice-versa.	You must ensure that both remote host and the host on which the database resides are in the domain and not in the workgroup.

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.	<p>You must perform one of following tasks:</p> <ul style="list-style-type: none">• Delete the credentials of the user whose password was changed and then add the new credentials in the cache by running the following commands:<ol style="list-style-type: none">a. <code>smsap credential delete</code>b. <code>smsap credential set</code>• Clear the entire cache by running the <code>smsap credential clear</code> 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.	<p>You must also clear the cache by performing the following steps:</p> <ol style="list-style-type: none">1. Start the console.2. Enter the following: <code>javaws -viewer</code>3. On the Java cache viewer screen, right-click the SnapManager application and select Delete.

Issue	Explanation	Workaround
<p>When you restart the GUI and try to check the backups for a certain profile, you see only the names of the profiles.</p>	<p>SnapManager does not display any information about a profile until you open it.</p>	<p>Perform the following steps:</p> <ol style="list-style-type: none"> 1. Right-click the profile and select Open from the menu. <p>SnapManager displays the Profile Authentication dialog box.</p> <ol style="list-style-type: none"> 2. Enter the host user name and password. <p>SnapManager displays the backup list.</p> <p> You only need to authenticate the profile once as long as the credentials are valid and remain in the cache.</p>
<p>When you open the first repository in the GUI, an error message similar to the following is displayed:</p> <p>The Profile name XXXX clashes with previously loaded repository.</p>	<p>Identically named profiles cannot exist in a repository. Also, you can open only one repository at a time.</p>	<p>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'</p>
<p>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</p>	<p>The repository is inaccessible, perhaps because it no longer exists. The GUI initializes the list of repositories from the credentials file.</p>	<p>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.</p>

Issue	Explanation	Workaround
<p>Profile creation fails because host credentials fail to authenticate in the SUSE Linux Enterprise Server 10 and SUSE Linux Enterprise Server 11 platforms.</p>	<p>SnapManager uses Pluggable Authentication Module (PAM) to authenticate users. In the SUSE Linux Enterprise Server versions 10 and 11 platforms, there is no <code>snapmanager</code> file by default in the <code>/etc/pam.d</code> directory that provides the required authentication details. Thus the host credentials fail.</p>	<p>To successfully log in to the host in the SUSE Linux Enterprise Server 10 and 11 platforms, perform the following steps:</p> <ol style="list-style-type: none"> 1. Create a <code>snapmanager</code> file in <code>/etc/pam.d/</code>. 2. Add the following content to the <code>snapmanager</code> file located at <code>/etc/pam.d/snapmanager</code>: <div data-bbox="1101 559 1496 982" style="border: 1px solid black; padding: 10px;"> <pre>#%PAM-1.0 auth include common-auth account include common-account password include common-password session include common-session</pre> </div> <ol style="list-style-type: none"> 3. Save the file and retry the profile creation operation.
<p>SnapManager takes a longer time to load the database tree structure and results in a timeout error message being displayed on the SnapManager GUI.</p>	<p>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.</p>	<p>Delete the credentials of the unused host, repository, and profile by using the <code>credential delete</code> command from the SnapManager command-line interface (CLI).</p>
<p>SnapManager fails to generate a new profile after the clone split operation and you do not know if the new profile is created.</p>	<p>SnapManager fails to prompt you if a new profile is not created after the clone split operation. Because no message is displayed for the failed operation, you might assume that the profile is created.</p>	<p>To know if a new profile is created for the clone split operation, perform the following steps:</p> <ol style="list-style-type: none"> 1. Click the Monitor tab, right-click the clone split operation entry and select Properties. 2. In the Profile Properties window, click the Logs tab to view the clone split operation and profile creation logs.

Issue	Explanation	Workaround
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
<p>SnapManager operation on the GUI does not proceed after you have cleared user credentials by using the <code>smsap credential clear</code> command from the SnapManager CLI or by clicking Admin > Credentials > Clear > Cache from the SnapManager GUI.</p>	<p>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.</p>	<p>Restart the SnapManager GUI depending on how the cache is cleared.</p> <ul style="list-style-type: none"> • 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. <p>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:</p> <ol style="list-style-type: none"> 1. Click Tasks > Add Existing 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. <p>Right-click the profile under the host, click Open, and enter the user credentials in Profile.</p>

Issue	Explanation	Workaround
<p>The error message Unable to list the protection policies for the following reason: Protection Manager is temporarily unavailable is displayed when you select None from the Protection Manager Protection Policy drop-down menu of the Profile Properties window and the policy settings page of the Profile create wizard.</p>	<p>The Protection Manager is not configured with SnapManager or the Protection Manager is not running.</p>	<p>No action is necessary.</p>
<p>You cannot open the SnapManager GUI by using Java Web Start GUI due to weaker Secure Sockets Layer (SSL) cipher strength of the browser.</p>	<p>SnapManager does not support SSL ciphers weaker than 128 bits.</p>	<p>Upgrade the browser version and check the cipher strength.</p>

Troubleshooting SnapDrive issues

There are a few common issues you might run into when using SnapManager with SnapDrive products.

First, you must determine if the issue is related to SnapManager for SAP or SnapDrive. If the issue is a SnapDrive error, SnapManager for SAP gives an error message similar to:

```
SMSAP-12111: Error executing snapdrive command "<snapdrive command>":  
<snapdrive error>
```

The following is an example of a SnapDrive error message where SMSAP-12111 is the SnapManager error number. The 0001-770 numbering scheme represents SnapDrive for UNIX errors.

```
SMSAP-12111: Error executing snapdrive command  
"/usr/sbin/snapdrive snap restore -file  
/mnt pathname/ar_anzio_name_10gR2_arrac1/data/undotbs02.dbf -snapname  
pathname.company.com:  
/vol/ar_anzio_name_10gR2_arrac1:  
TEST_ARRAC1_YORKTOW_arrac12_F_C_0_8abc01b20f9ec03d010f9ec06bee0001_0":  
0001-770  
Admin error: Inconsistent number of files returned when listing contents  
of  
/vol/ar_anzio_name_10gR2_arrac1/.snapshot/  
TEST_ARRAC1_YORKTOW_arrac12_F_C_0_8abc01b20f9ec03d010f9ec06bee0001_0/data  
on filer pathname.
```

The following are the most common SnapDrive for UNIX error messages related to LUN discovery, configuration issues, and space. If you receive any of these errors, see the Troubleshooting chapter of the *SnapDrive Installation and Administration Guide*.

Symptom	Explanation
0001-136 Admin error: Unable to log on to filer: <filer> Please set user name and/or password for <filer>	Initial SnapDrive configuration
0001-382 Admin error: Multipathing rescan failed	LUN discovery error
0001-462 Admin error: Failed to unconfigure multipathing for <LUN>: spd5: cannot stop device. Device busy.	LUN discovery error
0001-476 Admin error: Unable to discover the device associated with ... 0001-710 Admin error: OS refresh of LUN failed ...	LUN discovery error
0001-680 Admin error: Host OS requires an update to internal data to allow LUN creation or connection. Use 'snapdrive config prepare luns' or update this information manually...	LUN discovery error
0001-817 Admin error: Failed to create volume clone ... : FlexClone not licensed	Initial SnapDrive configuration
0001-878 Admin error: HBA assistant not found. Commands involving LUNs should fail.	LUN discovery error

Troubleshooting storage system renaming issue

You might face issues when renaming a storage system or after you have successfully renamed the storage system.

When you try to rename the storage system, the operation might fail with the following error message: SMSAP-05085 No storage controller "fas3020-rtp07New" is found to be associated with the profile

You must enter the IP address or name of the storage system that is listed when you run the `smsap storage list` command.

After you rename the storage system, SnapManager operations might fail if SnapManager fails to recognize the storage system. You must perform some additional steps in the DataFabric Manager server host and the SnapManager server host to resolve this issue.

Perform the following steps in the DataFabric Manager server host:

1. Delete the IP address and host of the earlier storage system in the host file located at `/etc/hosts` in the DataFabric Manager server host.

2. Add the new IP address and host of the new storage system in the host file located at /etc/hosts in the DataFabric Manager server host.
3. Change the storage host name by entering the following command:

```
dfm host rename -a old_host_namenew_host_name
```

4. Set the new IP address in the host by entering the following command:

```
dfm host set old_host_name_or_objId hostPrimaryAddress =  
new_storage_controller_ip_address
```



You must perform this step only if you have specified the IP address as the new storage system name.

5. Update the new storage system name in the DataFabric Manager server host by entering the following command:

```
dfm host diag old_storage_name
```

You can verify that the earlier storage controller name is replaced with new storage controller name by entering the following command:

```
dfm host discover new_storage_name
```

Perform the following steps as a root user in the SnapManager server host.



When you enter the new storage controller name, ensure that you use the system alias name and not the fully qualified domain name (FQDN).

1. Delete the earlier storage system name by entering the following command:

```
snapdrive config delete old_storage_name
```



If you do not delete the earlier storage system name, then all the SnapManager operations fail.

2. Delete the IP address and host of the earlier storage system in the host file located at /etc/hosts in the target database host.
3. Add the new IP address and host of the new storage system in the host file located at /etc/hosts in the target database host.
4. Add the new storage system name by entering the following command:

```
snapdrive config set root new_storage_name
```

5. Map the earlier and later storage system names by entering the following command:

```
snapdrive config migrate set old_storage_namenew_storage_name
```

6. Delete the management path of the earlier storage system by entering the following command:

```
snapdrive config delete -mgmtpath old_storage_name
```

7. Add the management path for the new storage system by entering the following command:

```
snapdrive config set -mgmtpath new_storage_name
```

8. Update the dataset for both data files and archive log files with the new storage system name by entering the following command:

```
snapdrive dataset changehostname -dn dataset_name -oldname old_storage_name  
-newname new_storage_name
```

9. Update the profile for the new storage system name by entering the following command:

```
smsap storage rename -profile profile_name -oldname old_storage_name -newname  
new_storage_name
```

10. Verify the storage system associated with the profile by entering the following command:

```
smsap storage list -profile profile_name
```

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: `/opt/NTAP/smsap/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, enter the following command: `smsap_server start`



An error message is displayed if the server is already running.

If the server is already running, perform the following steps:

1. Stop the server by entering the following command: `smsap_server stop`
2. Restart the server by entering the following command: `smsap_server start`

Terminating a currently running SnapManager operation

If SnapManager server hangs and you cannot execute any operations successfully, you can terminate SnapManager and its operations.

Workaround

SnapManager works with both SnapManager and Protection Manager. You must perform the following steps to list the different processes running and stop the last process running.

1. List all SnapDrive processes that are running: `ps`

`ps`

Example:

```
ps | grep snapdrive
```

2. Stop the SnapDrive process or processes: `kill <pid>`

`pid` is the list of processes you found using the `ps` command.



Do not stop all SnapDrive processes. You might want to end only the last process that is running.

3. If one of the operations involves restoring a protected backup from secondary storage, open the Protection Manager console and perform the following:
 - a. From the System menu, select **Jobs**.
 - b. Select **Restore**.
 - c. Check for the name of the dataset that matches the one in the SnapManager profile.
 - d. Right-click and select **Cancel**.
4. List the SnapManager processes:
 - a. Log in as a root user.
 - b. List the processes by using the `ps` command.

Example: `ps | grep java`

5. End the SnapManager process: `kill <pid>`

Unable to delete or free the last protected backup

When you create the first backup for a profile on secondary storage, SnapManager sends all the information about the backup to Protection Manager. For subsequent backups related to this profile, SnapManager sends only the modified information. If you remove the last protected backup, SnapManager loses the ability to identify the differences between backups and must find a way to rebaseline these relationships. Therefore, attempting to delete the last protected backup results in an error message being displayed.

Workaround

You can delete the profile or only the profile backup.

To delete the profile, perform the following steps:

1. Delete the profile's backups.
2. Update the profile and disable protection in the profile.

This deletes the dataset.

3. Delete the last protected backup.
4. Delete the profile.

To delete only the backup, perform the following steps:

1. Create another backup copy of the profile.
2. Transfer that backup copy to secondary storage.
3. Delete the previous backup copy.

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: /dest, /dest1, and /dest2. While creating the archive log file backup, if you exclude /dest by using the command

```
  smsap backup create -profile almsamp1 -data -online -archivelogs -exclude  
  -dest /dest
```

, SnapManager for SAP excludes all the destinations starting with /dest.

Workaround

- Add a path separator after destinations are configured in v\$archive_dest. For example, change the /dest to /dest/.
- While creating a backup, include destinations instead of excluding any destination.

Restoring control files that are multiplexed on Automatic Storage Management (ASM) and non-ASM storage fails

When the control files are multiplexed on ASM and non-ASM storage, the backup operation is successful. However, when you try to restore control files from that successful backup, the restore operation fails.

SnapManager clone operation fails

When you clone a backup in SnapManager, the DataFabric Manager server might fail to discover volumes, and display the following error message:

```
SMSAP-13032: Cannot perform operation: Clone Create. Root cause: SMSAP-11007:  
Error cloning from snapshot: FLOW-11019: Failure in ExecuteConnectionSteps: SD-  
00018: Error discovering storage for /mnt/datafile_clone3: SD-10016: Error  
executing snapdrive command "/usr/sbin/snapdrive storage show -fs  
/mnt/datafile_clone3": 0002-719 Warning: Could not check SD.Storage.Read access  
on volume filer:/vol/SnapManager_20091122235002515_voll for user user-vm5\oracle  
on Operations Manager servers x.x.x.x
```

Reason: Invalid resource specified. Unable to find its Id on Operations Manager server 10.x.x.x

This occurs if the storage system has large number of volumes.

Workaround

You must perform one of the following:

- From the Data Fabric Manager server, run

```
  dfm host discover storage_system
```

You can also add the command in a shell script file and schedule a job in the DataFabric Manager server to run the script at a frequent interval.

- Increase the value of *dfm-rbac-retries* in the *Snapdrive.conf* file.

SnapDrive uses the default refresh interval value and default number of retries. The default value of *dfm-rbac-retry-sleep-secs* is 15 seconds and *dfm-rbac-retries* is 12 iterations.



The Operations Manager refresh interval depends on the number of storage systems, number of storage objects in the storage system, and the load on the DataFabric Manager server.

As a recommendation, perform the following:

1. From the DataFabric Manager server, manually run the following command for all the secondary storage systems associated with the dataset:

```
dfm host discover storage_system
```

2. Double the time taken to perform the host discovery operation and assign that value to *dfm-rbac-retry-sleep-secs*.

For example, if the operation took 11 seconds, you can set the value of *dfm-rbac-retry-sleep-secs* to 22 (11*2).

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 <code>sm_gui.log</code> : [WARN] : SMSAP-01106: Error occurred while querying the repository: Closed Connection <code>java.sql.SQLException: Closed Connection</code> .
Refreshing an opened repository by pressing F5	A repository exception is displayed in the GUI and also logs a <code>NullPointerException</code> in the <code>sm_gui.log</code> file.
Refreshing the host server	A <code>NullPointerException</code> is logged in the <code>sumo_gui.log</code> file.
Creating a new profile	A <code>NullPointerException</code> is displayed in the Profile Configuration window.

Operations	Exceptions
Refreshing a profile	The following SQL exception is logged in <code>sm_gui.log</code> : [WARN] : SMSAP-01106: Error occurred while querying the repository: Closed Connection.
Accessing a backup	The following error message is logged in <code>sm_gui.log</code> : Failed to lazily initialize a collection.
Viewing clone properties	The following error message is logged in <code>sm_gui.log</code> and <code>sumo_gui.log</code> : 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.

Unable to migrate the protocol from NFSv3 to NFSv4

You can migrate the protocol from NFSv3 to NFSv4 by enabling the `enable-migrate-nfs-version` parameter in the `snapdrive.conf` file. During the migration, SnapDrive considers only the protocol version, irrespective of the mount point options such as `rw`, `largefiles`, `nosuid`, and so on.

However, after migrating the protocol to NFSv4, when you restore the backup that was created by using NFSv3, the following occurs:

- If NFSv3 and NFSv4 are enabled at the storage level, the restore operation is successful but is mounted with the mount point options that were available during backup.
- If only NFSv4 is enabled at the storage level, the restore operation is successful and only the protocol version (NFSv4) is retained.

However, the other mount point options such as `rw`, `largefiles`, `nosuid`, and so on are not retained.

Workaround

You must manually shut down the database, unmount the database mount points, and mount with the options

available prior to the restore.

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.

Mounting a FlexClone volume fails in NFS environment

When SnapManager creates a FlexClone of a volume in an NFS environment, an entry is added in the `/etc/exports` file. The clone or backup fails to mount on a SnapManager host with an error message.

The error message is: 0001-034 Command error: mount failed: mount:
filer1:/vol/SnapManager_20090914112850837_voll4 on /opt/NTAPsmsap/mnt/-
ora_data02-20090914112850735_1 - WARNING unknown option "zone=voll4" nfs mount:
filer1:/vol/SnapManager_20090914112850837_voll4: Permission denied.

At the same time, the following message is generated at the storage system console: Mon Sep 14
23:58:37 PDT [filer1: export.auto.update.disabled: warning]: /etc/exports was not
updated for voll4 when the vol clone create command was run. Please either
manually update /etc/exports or copy /etc/exports.new to it.

This message might not be captured in the AutoSupport messages.



You might encounter similar issues while cloning FlexVol volumes on NFS. You can follow the same steps to enable the `nfs.export.auto-update` option.

What to do

1. Set the `nfs.export.auto-update` option *on* so that the `/etc/exports` file is updated automatically.

`options nfs.export.auto-update on`



In the HA pair configuration, ensure that you set the NFS exports option to on for both the storage systems.

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.

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 and use the following command to update SDU with the igroup information:

```
snapshot igrup add
```

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.
<i>Data ONTAP SAN Configuration Guide for 7-Mode</i>	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.

Document	Description
SnapManager Release Notes	<p>This document comes with SnapManager. You can also download a copy from NetApp Support.</p> <p>It contains any last-minute information that you need to get the configuration up and running smoothly.</p>
NetApp host attach and support kits documentation	NetApp Support .
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-DS-10021: Unable to set protection policy of dataset <dataset-name> to <new-protection-policy> because the protection policy is already set to <old-protection-policy>. Please use Protection Manager to change the protection policy	After the protection policy of a dataset is set, SnapManager will not allow you to change the protection policy, because it might require realigning the baseline relationships and result in the loss of existing backups on the secondary storage.	Update the protection policy using Protection Manager's Management Console, which provides more options on migrating from one protection policy to another.
0001-SD-10028: SnapDrive Error (id:2618 code:102) Unable to discover the device associated with "lun_path". If multipathing in use, possible multipathing configuration error. Please verify configuration and retry.	The host is not able to discover LUNs created on the storage systems.	Ensure that the transport protocol is properly installed and configured. Ensure that SnapDrive can create and discover a LUN on the storage system.

Error message	Explanation	Resolution
0001-SD-10028: SnapDrive Error (id:2836 code:110) Failed to acquire dataset lock on volume "storage name":"temp_volume_name"	You tried to restore using the indirect storage method and the temporary volume specified does not exist on the primary storage.	Create a temporary volume on the primary storage. Or, specify the correct volume name, if a temporary volume is already created.
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).
0001-SMSAP-11027: Cannot clone or mount snapshots from secondary storage because the snapshots are busy. Try cloning or mounting from an older backup.	You tried to create a clone or mount Snapshot copies from the secondary storage of the latest protected backup.	Clone or mount from an older backup.
0001-SMSAP-12346: Cannot list protection policies because Protection Manager product is not installed or SnapDrive is not configured to use it. Please install Protection Manager and/or configure SnapDrive...	You tried to list protection policies on a system where SnapDrive is not configured to use Protection Manager.	Install Protection Manager and configure SnapDrive to use Protection Manager.
0001-SMSAP-13032: Cannot perform operation: Backup Delete. Root cause: 0001-SMSAP-02039: Unable to delete backup of dataset: SD-10028: SnapDrive Error (id:2406 code:102) Failed to delete backup id: "backup_id" for dataset, error(23410):Snapshot "snapshot_name" on volume "volume_name" is busy.	You tried to free or delete the latest protected backup or a backup containing Snapshot copies that are baselines in a mirror relationship.	Free or delete the protected 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.	<ol style="list-style-type: none"> 1. Navigate to the SnapManager installation directory. 2. Open the <code>launch-java</code> file. 3. Increase the value of the <code>java -Xmx160m</code> Java heap space parameter. For example, you can modify the value from the default value of 160m to 200m as <code>java -Xmx200m</code>.
SD-10028: SnapDrive Error (id:2868 code:102) Could not locate remote snapshot or remote qtree.	SnapManager displays the backups as protected even if the protection job in Protection Manager is only partially successful. This condition occurs when dataset conformance is in progress (when the baseline Snapshots are getting mirrored).	Take a new backup after the dataset is conformant.
SMSAP-21019: The archive log pruning failed for the destination: "/mnt/destination_name/" with the reason: "ORACLE- 00101: Error executing RMAN command: [DELETE NOPROMPT ARCHIVELOG '/mnt/destination_name/']	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.

Error message	Explanation	Resolution
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.
<pre>Unable to parse shell output: (javax.util.regex.Matcher[pattern=Command complete. region=0,18 lastmatch=]) does not match (name:backup_script) Unable to parse shell output: (javax.util.regex.Matcher[pattern=Command complete. region=0,25 lastmatch=]) does not match (description:backup script) Unable to parse shell output: (javax.util.regex.Matcher[pattern=Command complete. region=0,9 lastmatch=]) does not match (timeout:0)</pre>	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.	<p>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:</p> <ul style="list-style-type: none"> • The block size of the tablespace in which the repository exists is less than 8k. • The <code>nls_length_semantics</code> parameter is set to <code>char</code>. 	<p>You must assign the values to the following parameters:</p> <ul style="list-style-type: none"> • <code>block_size=8192</code> • <code>nls_length=byte</code> <p>After modifying the parameter values, you must restart the database.</p> <p>For more information, see the Knowledge Base article 2017632.</p>

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 <i>-force</i> 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 <i>-force</i> 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, splitting clones, 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.

Data protection errors

The following table shows the common errors associated with data protection:

Error message	Explanation	Resolution
Backup protection is requested but the database profile does not have a protection policy. Please update the protection policy in the database profile or do not use the 'protect' option when creating backups.	You try to create a backup with protection to secondary storage; however, the profile associated with this backup does not have a protection policy specified.	Edit the profile and select a protection policy. Re-create the backup.
Cannot delete profile because data protection is enabled but the Protection Manager is temporarily unavailable. Please try again later.	You try to delete a profile that has protection enabled; however, Protection Manager is unavailable.	Ensure that appropriate backups are stored in either primary or secondary storage. Disable protection in the profile. When Protection Manager is available again, return to the profile and delete it.

Error message	Explanation	Resolution
Cannot list protection policies because Protection Manager is temporarily unavailable. Please try again later.	While setting up the backup profile, you enabled protection on the backup so that the backup would be stored on secondary storage. However, SnapManager cannot retrieve the protection policies from Protection Manager Management Console.	Disable protection in the profile temporarily. Continue creating a new profile or updating an existing profile. When Protection Manager is available again, return to the profile.
Cannot list protection policies because Protection Manager product is not installed or SnapDrive is not configured to use it. Please install Protection Manager and/or configure SnapDrive.	While setting up the backup profile, you enabled protection on the backup so that the backup would be stored on secondary storage. However, SnapManager cannot retrieve the protection policies from Protection Manager's Management Console. The Protection Manager is not installed or SnapDrive is not configured.	Install Protection Manager. Configure SnapDrive. Return to the profile, reenable protection, and select the protection policies available in Protection Manager's Management Console.
Cannot set protection policy because Protection Manager is temporarily unavailable. Please try again later.	While setting up the backup profile, you enabled protection on the backup so that the backup would be stored on secondary storage. However, SnapManager cannot retrieve the protection policies from Protection Manager's Management Console.	Disable protection in the profile temporarily. Continue creating or updating the profile. When Protection Manager's Management Console is available, return to the profile.
Creating new dataset <dataset_name> for database <dbname> on host <host>.	You attempted to create a backup profile. SnapManager creates a dataset for this profile.	No action necessary.
Data protection is not available because Protection Manager is not installed.	While setting up the backup profile, you attempted to enable protection on the backup so that the backup would be stored on secondary storage. However, SnapManager cannot access protection policies from Protection Manager's Management Console. The Protection Manager is not installed.	Install Protection Manager.
Deleted dataset <dataset_name> for this database.	You deleted a profile. SnapManager will delete the associated dataset.	No action is necessary.

Error message	Explanation	Resolution
<p>Deleting profile with protection enabled and Protection Manager is no longer configured.</p> <p>Deleting profile from SnapManager but not cleaning up dataset in Protection Manager.</p>	<p>You attempted to delete a profile that has protection enabled; however, Protection Manager is no longer installed, or no longer configured, or has expired. SnapManager will delete the profile, but not the profile's dataset from Protection Manager's Management Console.</p>	<p>Reinstall or reconfigure Protection Manager. Return to the profile and delete it.</p>
<p>Invalid retention class. Use "smsap help backup" to see a list of available retention classes.</p>	<p>When setting up the retention policy, you attempted to use an invalid retention class.</p>	<p>Create a list of valid retention classes by entering this command: smsap help backup</p> <p>Update the retention policy with one of the available classes.</p>
<p>Specified protection policy is not available. Use "smsap protection-policy list" to see a list of available protection policies.</p>	<p>While setting up the profile, you enabled protection and entered a protection policy that is not available.</p>	<p>Identify available protection policies, by entering the following command: smsap protection-policy list</p>
<p>Using existing dataset <dataset_name> for database <dbname> on host <host> since the dataset already existed.</p>	<p>You attempted to create a profile; however, the dataset for the same database profile already exists.</p>	<p>Check the options from the existing profile and ensure that they match what you need in the new profile.</p>
<p>Using existing dataset <dataset_name> for RAC database <dbname> since profile <profile_name> for the same RAC database already exists for instance <SID> on host <hostname>.</p>	<p>You attempted to create a profile for a RAC database; however, the dataset for the same RAC database profile already exists.</p>	<p>Check the options from the existing profile and ensure that they match what you need in the new profile.</p>

Error message	Explanation	Resolution
<p>The dataset <dataset_name> with protection policy <existing_policy_name> already exists for this database. You have specified protection policy <new_policy_name>. The dataset's protection policy will be changed to <new_policy_name>. You can change the protection policy by updating the profile.</p>	<p>You attempted to create a profile with protection enabled and a protection policy selected. However, the dataset for the same database profile already exists, but has a different protection policy. SnapManager will use the newly specified policy for the existing dataset.</p>	<p>Review this protection policy and determine if this is the policy you want to use for the dataset. If not, edit the profile and change the policy.</p>
<p>Protection Manager deletes the local backups created by SnapManager for SAP</p>	<p>The Protection Manager's Management Console deletes or frees the local backups created by SnapManager based on the retention policy defined in the Protection Manager. The retention class set for the local backups is not considered while deleting or freeing the local backups. When the local backups are transferred to a secondary storage system, the retention class set for the local backups on the primary storage system are not considered. The retention class specified in the transfer schedule is assigned to the remote backup.</p>	<p>Run the <code>dfpm dataset fix_smsap</code> command from the Protection Manager server every time a new dataset is created. Now the backups are not deleted based on the retention policy set in Protection Manager's Management Console.</p>

Error message	Explanation	Resolution
<p>You have selected to disable protection for this profile. This could potentially delete the associated dataset in Protection Manager and destroy the replication relationships created for that dataset. You will also not be able to perform SnapManager operations such as restoring or cloning the secondary or tertiary backups for this profile. Do you wish to continue (Y/N) ?</p>	<p>You tried to disable protection for a protected profile while updating the profile from the SnapManager CLI or GUI. You can disable protection for the profile using the <code>-noprotect</code> option from the SnapManager CLI or clearing the Protection Manager Protection Policy check box in the Policies properties window from the SnapManager GUI. When you disable protection for the profile, SnapManager for SAP deletes the dataset from Protection Manager's Management Console, which unregisters all of the secondary and tertiary backup copies associated with that dataset.</p> <p>After a dataset is deleted, all secondary and tertiary backup copies are orphaned. Neither the Protection Manager nor SnapManager for SAP have the ability to access those backup copies. The backup copies can no longer be restored by using SnapManager for SAP.</p> <p> The same warning message is displayed even when the profile is not protected.</p>	<p>This is a known issue in SnapManager for SAP and expected behavior within Protection Manager when destroying a dataset. There is no workaround. The orphaned backups need to be managed manually.</p>

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
<p>SMSAP-03031:Restore specification is required to restore backup <variable> because the storage resources for the backup has already been freed.</p>	<p>You attempted to restore a backup that has its storage resources freed without specifying a restore specification.</p>	<p>Specify a restore specification.</p>

Error message	Explanation	Resolution
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.
ORACLE-30028: Unable to dump log file <filename>. The file may be missing/inaccessible/corrupted. This log file will not be used for recovery.	<p>The online redo log files or archive log files cannot be used for recovery. This error occurs due to following reasons:</p> <ul style="list-style-type: none"> 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.
SMSAP-03038: Cannot restore from secondary because the storage resources still exist on primary. Please restore from primary instead.	You tried to restore from secondary storage, but Snapshot copies exist on the primary storage.	Always restore from the primary if the backup has not been freed.

Error message	Explanation	Resolution
SMSAP-03054: Mounting backup archbkp1 to feed archivelogs. DS-10001: Connecting mountpoints. [ERROR] FLOW-11019: Failure in ExecuteConnectionSteps: SD-10028: SnapDrive Error (id:2618 code:305). The following files could not be deleted. The corresponding volumes might be read-only. Retry the command with older snapshots.[ERROR] FLOW-11010: Operation transitioning to abort due to prior failure.	During recovery, SnapManager tries to mount the latest backup from secondary to feed the archive log files from secondary. Though, if there are any other backups, the recovery can succeed. But, if there are no other backups, the recovery might fail.	Do not delete the latest backups from primary, so that SnapManager can use the primary backup for recovery.

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.
SMSAP-04908: Not a FlexClone.	The clone is a LUN clone. This applies for Data ONTAP 8.1 7-mode as well as clustered Data ONTAP.	SnapManager supports clone split on the FlexClone technology only.
SMSAP-04904: No clone split operation running with <i>split-idsplit_id</i>	The operation ID is invalid or no clone split operation is in progress.	Provide a valid split ID or split label for the clone split status, result, and stop operations.
SMSAP-04906: Stop clone split operation failed with <i>split-idsplit_id</i>	The split operation is complete.	Check whether the split process is in progress by using the clone split-status or clone split-result command.

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.	<p>The clone creation fails when you create the clone from the standby database using the following setup:</p> <ul style="list-style-type: none"> The primary database is a RAC setup and the standby database is standalone. The standby is created by using RMAN for taking the data files backup. 	Add the <code>_no_recovery_through_reset_logs=TRUE</code> 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.
[INFO] Operation failed. Syntax errors in clone specification: [error: cvc-complex-type.2.4c: Expected elements 'value@http://www.example.com default@http://www.example.com' before the end of the content in element parameter@http://www.example.com]	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 <code>smsapsystem dump</code> .

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.
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.

Virtual storage interface errors (virtual storage interface 8000 series)

The following table shows the common errors associated with virtual storage interface tasks:

Error message	Explanation	Resolution
SMSAP-08017 Error discovering storage for ./.	<p>SnapManager attempted to locate storage resources, but found data files, control files, or logs in the root/ directory. These files should reside in a subdirectory. The root file system might be a hard drive in your local machine. SnapDrive cannot take Snapshot copies at this location and SnapManager cannot perform operations on these files.</p>	<p>Check to see if data files, control files, or redo logs are in the root directory. If so, move them to their correct locations or re-create control files or redo logs in their correct locations. The basic mountpoint is always</p> <ul style="list-style-type: none"> • /oracle/<SID> on UNIX-based environments • [drive:]\\oracle\\<SID> on Windows-based environments <p>SAP uses four redo log groups with two members:</p> <ul style="list-style-type: none"> • One member is stored in origlog{A B}. • The other member is mirrored to mirrlog{A B}.

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.

Error message	Explanation	Resolution
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.
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 <i>smid</i> variable.	Insert the required <i>smid</i> variable.
SMSAP-13902: Clone Split Start failed.	<p>There could be multiple reasons for this error:</p> <ul style="list-style-type: none"> • No space in the volume. • SnapDrive is not running. • Clone could be a LUN clone. • FlexVol volume has restricted Snapshot copies. 	Check for the available space in the volume by using the clone split-estimate command. Confirm that the FlexVol volume has no restricted Snapshot copies.
SMSAP-13904: Clone Split Result failed.	This could be due to failure in the SnapDrive or storage system.	Try working on a new clone.
SMSAP-13906: Split operation already running for clone label <i>clone-label</i> or ID <i>clone-id</i> .	You are trying to split a clone that is already split.	The clone is already split and the clone related metadata will be removed.
SMSAP-13907: Split operation already running for clone label <i>clone-label</i> or ID <i>clone-id</i> .	You are trying to split a clone that is undergoing the split process.	You must wait until the split operation completes.

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.
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.

Error message	Explanation	Resolution
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.
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.

Common SnapDrive for UNIX error messages

The following table shows the common errors related to SnapDrive for UNIX:

Error message	Explanation
0001-136 Admin error: Unable to log on to filer: <filer> Please set user name and/or password for <filer>	Initial configuration error
0001-382 Admin error: Multipathing rescan failed	LUN discovery error
0001-462 Admin error: Failed to unconfigure multipathing for <LUN>: spd5: cannot stop device. Device busy.	LUN discovery error
0001-476 Admin error: Unable to discover the device associated with...	LUN discovery error
0001-680 Admin error: Host OS requires an update to internal data to allow LUN creation or connection. Use 'snapdrive config prepare luns' or update this information manually...	LUN discovery error

Error message	Explanation
0001-710 Admin error: OS refresh of LUN failed...	LUN discovery error
0001-817 Admin error: Failed to create volume clone... : FlexClone not licensed	Initial configuration error
0001-817 Admin error: Failed to create volume clone... : Request failed as space cannot be guaranteed for the clone.	Space issue
0001-878 Admin error: HBA assistant not found. Commands involving LUNs should fail.	LUN discovery error
SMSAP-12111: Error executing snapdrive command "<snapdrive command>": <snapdrive error>	SnapDrive for UNIX generic error

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.