



Product overview

SnapManager Oracle

NetApp
April 15, 2021

This PDF was generated from https://docs.netapp.com/us-en/snapmanager-oracle/unix-administration/concept_create_backups_using_snapshot_copies.html on April 15, 2021. Always check docs.netapp.com for the latest.

Table of Contents

- Product overview 1
 - SnapManager highlights 1
 - Create backups using Snapshot copies 1
 - 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
 - Accessing and printing online Help 8
 - Recommended general database layouts and storage configurations 8
 - Limitations when working with SnapManager 20

Product overview

SnapManager for Oracle automates and simplifies the complex, manual, and time-consuming processes associated with the backup, recovery, and cloning of Oracle 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 integrates with native Oracle technologies such as Oracle Real Application Clusters (Oracle RAC), Automatic Storage Management (ASM), and Direct NFS across FC, iSCSI, and NFS protocols. Optionally, backups created using SnapManager can be cataloged with Oracle Recovery Manager (RMAN) to preserve the 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 Oracle databases on the UNIX host and with NetApp Snapshot, SnapRestore, and FlexClone technologies on the back end. It offers an easy-to-use user interface (UI) as well as a 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

You can back up the data files and archive log files separately.

- Scheduling backups
- Restoring full or partial databases 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 systems

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 Oracle 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 Oracle consolidates the archive log backups to maintain a minimum number of backups for archive log files. SnapManager for Oracle 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).

DBAs can use SnapManager to restore the database and use another tool to recover the information. DBAs are not required to use SnapManager for both operations.

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:](https://mysupport.netapp.com/documentation/productsatoz/index.html)
mysupport.netapp.com/documentation/productsatoz/index.html

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]: SMO-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.

Related information

[Managing profiles for efficient backups](#)

What SnapManager operation states are

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

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

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.
- An Automatic Storage Management (ASM) database is configured, but the ASM instance is not running.
- 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 Oracle 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 sm 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 sm credential set -host command.



The sm credential set -host command remembers your credentials when the host.credentials.persist property in the sm.config file is set to true.

Example

User1 and User2 share a profile called Prof2. User2 cannot perform a backup of Database1 in Host1 without

permission to access Host1. User1 cannot clone a database to Host3 without permission to access Host3.

The following table describes different permissions assigned to the users:

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

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

Accessing and printing online Help

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

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

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 register SnapManager backups with Oracle Recovery Manager (RMAN), you must create RMAN-enabled profiles.

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:

- Multiple databases cannot share the same Automatic Storage Management (ASM) disk group.
- 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.

Defining the database home with the oratab file

SnapManager uses the oratab file during operations to determine the Oracle database home directory. An entry for your Oracle database must be in the oratab file for SnapManager to work correctly. The oratab file is created during the Oracle software installation.

The oratab file resides in different locations based on the host operating system as shown in the following table:

Host operating system	File location
Linux	/etc/oratab
Solaris	/var/opt/oracle/oratab
IBM AIX	/etc/oratab

The sample oratab file contains the following information:

```
+ASM1:/u01/app/11.2.0/grid:N # line added by Agent
oelpro:/u01/app/11.2.0/oracle:N # line added by Agent
# SnapManager generated entry (DO NOT REMOVE THIS LINE)
smoclone:/u01/app/11.2.0/oracle:N
```



After Oracle is installed, you must ensure that the oratab file resides in the location specified in the previous table. If the oratab file does not reside in the correct location per your operating system, you must contact technical support for assistance.

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 or an Automatic Storage Management (ASM) environment, 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.

Requirements for using ASM databases with SnapManager

You must know the requirements for using Automatic Storage Management (ASM) databases with SnapManager. Knowing these requirements can help you avoid issues with the ASMLib, partitions, and clone specifications, among other things.

- SnapManager (3.0.3 or later) uses the new sysasm privilege available with Oracle 11gR2 instead of the sysdba privilege to administer an Oracle ASM instance.

If you use the sysdba privilege to run administrative commands on the ASM instance, an error message is displayed. The database uses the sysdba privilege to access disk groups. If you connect to the ASM instance using the sysasm privilege, you have complete access to all the available Oracle ASM disk groups and management functions.



If you are using Oracle 10gR2 and 11gR1, you must continue to use the sysdba privilege.

- SnapManager (3.0.3 or later) supports backing up databases that are stored directly on ASM disk groups when the disk group also contains an Automatic Cluster File System (ACFS) volume.

These files are indirectly protected by SnapManager and might be restored with the remaining contents of an ASM diskgroup, but SnapManager (3.0.3 or later) does not support ACFS.



ACFS is a multiplatform, scalable file-system storage management technology available with Oracle 11gR2. ACFS extends ASM functionality to support customer files maintained outside the Oracle database.

- SnapManager (3.0.3 or later) supports the backup of files that are stored on ASM disk groups when the disk group also contains Oracle Cluster Registry (OCR) files or voting disk files; however, restore operations require slower, host-based or partial-file snap restore (PFSR) method.

It is best to have OCR and voting disks on disk groups that do not contain database files.

- Each disk used for ASM must contain only one partition.
- The partition hosting the ASM data must be properly aligned to avoid severe performance problems.

This implies that the LUN must be of the correct type and the partition must have an offset that is a multiple of 4K bytes.



For details about how to create partitions that are aligned to 4K, see the Knowledge Base article 1010717.

- ASM configuration is not specified as part of the clone specification.

You must manually remove the ASM configuration information in clone specifications that were created using SnapManager 2.1 before upgrading the host to SnapManager (2.2 or later).

- SnapManager 3.1, 3.1p1, and 3.2 or later support ASMLib 2.1.4.
- SnapManager 3.1p4 or later support ASMLib 2.1.4, 2.1.7, and 2.1.8.

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***
Red Hat Enterprise Linux 5x or later or Oracle Enterprise Linux 5x or later	Yes	No	Yes	ASM with ASMLib**
SUSE Linux Enterprise Server 10 SP4 or SUSE Linux Enterprise Server 11	Yes	No	Yes	ASM with ASMLib**

Operating system	Single partition	Multiple partition	Non-partition devices	File system or RAW devices
SUSE Linux Enterprise Server 10 SP4 or later	Yes	No	No	ASM without ASMLib**
SUSE Linux Enterprise Server 11				

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

Support for ASMLib

SnapManager supports different versions of ASMLib, although there are several factors you must consider when using SnapManager with ASMLib.

SnapManager supports ASMLib 2.1.4, 2.1.7, and 2.1.8. All SnapManager operations can be performed with ASMLib 2.1.4, 2.1.7, and 2.1.8.

If you have upgraded from ASMLib 2.1.4 to ASM 2.1.7, you can use the same profiles and backups created with ASMLib 2.1.4 to restore the backups and create the clones.

You must consider the following when using SnapManager with ASMLib:

- SnapManager 3.1 does not support ASMLib 2.1.7.

SnapManager 3.1p4 or later support ASMLib 2.1.4, 2.1.7, and 2.1.8.

- After performing a rolling upgrade from SnapManager 3.1 to 3.2, the backups created by using ASMLib 2.1.7 work only if the repository is rolled back to SnapManager 3.1 and ASMLib 2.1.7 is downgraded to ASMLib 2.1.4.
- After performing a rolling upgrade from SnapManager 3.1 to 3.2, backups created by using ASMLib 2.1.7 do not work if the repository is rolled back to SnapManager 3.1 with ASMLib 2.1.7.

The rollback succeeds, but the profiles and backups cannot be used.

Support for ASM databases without ASMLib

SnapManager supports ASM without ASMLib, by default. The basic requirement is that the devices that are used for ASM disk groups must be partitioned.

When ASMLib is not installed, the device permissions related to ASM disk groups are changed to root:disk when you perform the following operations:

- Restart the host
- Restore a database from the primary storage by using volume-based SnapRestore (VBSR)
- Restore a database from the secondary storage

You can set the proper device permissions by assigning true to the `oracleasm.support.without.asmlib` configuration variable in `smo.conf`. The devices related to the ASM disk groups are added or removed from the

initasmdisks file whenever new devices are added or removed from the host. The initasmdisks file is located at /etc/initasmdisks.

For example, if you set `oracleasm.support.without.asmlib=true` and then perform a backup mount, new devices are added to initasmdisks. When the host is restarted, the device permissions and ownership are maintained by the startup scripts.



The default value for `oracleasm.support.without.asmlib` is false.

Related information

[Supported partition devices](#)

Supported scripts

The `asmmain.sh` and `asmquerydisk.sh` scripts allow you to change the grid user, group, and the user, all of which are used to query the ASM disks. The scripts must always be executed from the root.

The `asmmain.sh` is the main script file called from any operation that adds or deletes devices. The `asmmain.sh` script calls another script internally, which needs to be executed from the root that has oracle grid credentials. This script queries the ASM disk group's devices, then adds those entries in the `initasmdisk` file with the permission and the ownership of the devices. You can change the permissions and ownership of this file based on your environment and the `regex` pattern that is used for matching only the `/dev/mapper/*p1`.

The `asmquerydisk.sh` script is used to query the disk list, which is used to create the ASM disk group. You must assign values to `ORACLE_BASE`, `ORACLE_HOME`, and `ORACLE_SID`, depending on your configuration.

The scripts are located at `/opt/NetApp/smo/plugins/examples/noasmlib`. However, these scripts must be moved to `/opt/NetApp/smo/plugins/noasmlib` before starting the SnapManager for Oracle server on the host.

Limitations of using scripts to support an ASM database without ASMLib

You must be aware of certain limitations to using scripts to support an ASM database without ASMLib.

- The scripts provide an alternative solution for any kernel version, but only if ASMLib is not installed.
- The permissions for the scripts must be set in such a way that the scripts can be accessed by root, grid, oracle, or equivalent users.
- The scripts do not support restoration from a secondary location.

Deploying and running the scripts

You can deploy and run the `asmmain.sh` and `asmquerydisk.sh` scripts to support ASM databases without ASMLib.

These scripts do not follow the pre-scripts or post-scripts syntax and workflow is called when `initasmdisks` is enabled. You can change anything related to your configuration settings in the scripts. It is recommended to verify if everything in the scripts are working as expected by performing a quick dry run.



These scripts do not harm your system on failures nor will they impact your system. These scripts are executed to update the ASM-related disks to have proper permissions and ownership, so that the disks will always be under ASM instance control.

1. Create the ASM disk groups with the partitioned disks.
2. Create the Oracle database on the DISK GROUPS.
3. Stop the SnapManager for Oracle server.



In an RAC environment, you need perform this step on all the RAC nodes.

4. Modify the smo.conf to include the following parameters:
 - a. oracleasm.support.without.asmlib = true
 - b. oracleasm.support.without.asmlib.ownership = true
 - c. oracleasm.support.without.asmlib.username = user name of your ASM instance environment
 - d. oracleasm.support.without.asmlib.groupname = group name of your ASM instance environmentThese modifications set the permissions for the absolute path only, which means instead of partition device, permissions will be set only for dm-* device.
5. Modify the plugins scripts available in /opt/NetApp/smo/plugins/examples/noasmlib to include your configuration settings in the scripts.
6. Copy the scripts to /opt/NetApp/smo/plugins/noasmlib before starting the SnapManager for Oracle server on the host.
7. Navigate to the /opt/NetApp/smo directory and perform a dry run by running the following script: sh plugins/noasmlib/asmmain.sh

The etc/initasmdisks file is created, which is the main file that is used.

You can confirm that the etc/initasmdisks file contains all the devices related to configured the ASM database, such as:

```
chown -R grid:oinstall /dev/mapper/360a98000316b61396c3f394645776863p1
  chmod 777 /dev/mapper/360a98000316b61396c3f394645776863p1
  chown -R grid:oinstall
/dev/mapper/360a980003754322f7a2b433469714239p1
  chmod 777 /dev/mapper/360a980003754322f7a2b433469714239p1
  chown -R grid:oinstall
/dev/mapper/360a980003754322f7a2b433469714241p1
  chmod 777 /dev/mapper/360a980003754322f7a2b433469714241p1
  chown -R grid:oinstall
/dev/mapper/360a980003754322f7a2b433469714243p1
  chmod 777 /dev/mapper/360a980003754322f7a2b433469714243p1
```

8. Start the SnapManager for Oracle server.
9. Configure SnapDrive for UNIX by adding the following to snapdrive.conf file.disconnect-luns-before-vbsr=on

10. Restart the SnapDrive for UNIX server.



In an RAC environment, you need perform the Step 3 through Step 10 for all the RAC nodes.

The `/etc/initasmdisks` file created, must be executed from either one of the startup scripts or from a script that is newly defined in the `rc3.d`. The `/etc/initasmdisks` file should always be executed before the `oracleha` service starts.

Example

```
# ls -ltr *ohasd*
      lrwxrwxrwx 1 root root 17 Aug  7 02:34 S96ohasd ->
/etc/init.d/ohasd
      lrwxrwxrwx 1 root root 17 Aug  7 02:34 K15ohasd ->
/etc/init.d/ohasd
```

In the following example, `sh -x/etc/initasmdisks` will not be available by default, and you need to append it as the first line in the function `start_stack()` in an `ohasd` script:

```
start_stack()
{
sh -x /etc/initasmdisks
# see init.ohasd.sbs for a full rationale case $PLATFORM in Linux
}
```

Support for Oracle RAC ASM databases without ASMLib

If you are using Oracle RAC databases, the RAC nodes must be updated with the `initasmdisks` file whenever an operation is performed in the master RAC node.

If no authentication is required to log in into the RAC nodes from the master node, the `asmmain.sh` performs a secure copy (SCP) of `initasmdisks` to all the RAC nodes. The master node's `initasmdisks` file will be called whenever restore happens, and the `asmmain.sh` script can be updated to invoke the same script in all the RAC nodes.

The `/etc/initasmdisks` file created that must be executed from either one of the startup scripts or from a newly defined script in the `rc3.d`. The `/etc/initasmdisks` file should always be executed before the `oracleha` service starts.

Support for Oracle 10g ASM databases without ASMLib

If you are using Oracle 10g, the `asmcmd` command is not available for listing disks. You can use the `sql` query to obtain the disks list.

The `disk_list.sql` script is included in the existing scripts provided in the `examples` directory to support `sql` queries. When you execute `theasmquerydisk.sh` script, the `disk_list.sql` script must be executed manually. The example script lines are added with comments in the `asmquerydisk.sh` file. This file can either be placed in the

/home/grid location or another location of your choice.

Sample scripts to support ASM databases without ASMLib

The sample scripts are available in the plugins/examples/noasmlib directory of the SnapManager for Oracle installation directory.

asmmain.sh

```
#!/bin/bash
griduser=grid
gridgroup=oinstall

# Run the script which takes the disklist from the asmcmd
# use appropriate user , here grid user is being used to run
# asmcmd command.
su -c "plugins/noasmlib/asmdiskquery.sh" -s /bin/sh grid
cat /home/grid/disklist

# Construct the final file as .bak file with propre inputs
awk -v guser=$griduser -v gggroup=$gridgroup '/^\s*/dev\/mapper/ { print
"chown -R "guser":"gggroup" "$1; print "chmod 777 " $1; }'
/home/grid/disklist > /etc/initasmdisks.bak

# move the bak file to the actual file.
mv /etc/initasmdisks.bak /etc/initasmdisks

# Set full full permission for this file to be called while rebooting and
restore
chmod 777 /etc/initasmdisks

# If the /etc/initasmdisks needs to be updated in all the RAC nodes
# or /etc/initasmdisks script has to be executed in the RAC nodes, then
the following
# section needs to be uncommented and used.
#
# Note: To do scp or running scripts in remote RAC node via ssh, it needs
password less login
# for root user with ssh keys shared between the two nodes.
#
# The following 2 lines are used for updating the file in the RAC nodes:
# scp /etc/initasmdisks root@racnode1:/etc/initasmdisks
# scp /etc/initasmdisks root@racnode2:/etc/initasmdisks
#
# In order to execute the /etc/initasmdisks in other RAC nodes
# The following must be added to the master RAC node /etc/initasmdisks
file
```

```

# from the asmmain.sh script itself. The above scp transfer will make sure
# the permissions and mode for the disk list contents are transferred to
the other RAC nodes
# so now appending any command in the /etc/initasmdisks will be retained
only in the master RAC node.
# The following lines will add entries to the /etc/initasmdisks file in
master RAC node only. When this script is executed
# master RAC node, /etc/initasmdisks in all the RAC nodes will be
executed.
# echo 'ssh racnode1 /etc/initasmdisks' >> /etc/initasmdisks
# echo 'ssh racnode2 /etc/initasmdisks' >> /etc/initasmdisks

```

asmquerydisk.sh

```

#!/bin/bash
export ORACLE_BASE=/u01/app/oracle
export ORACLE_HOME=/u01/app/grid/product/11.2.0.3/grid
export ORACLE_SID=+ASM
export PATH=$ORACLE_HOME/bin:$PATH

# Get the Disk List and save this in a file called dglist.
asmcmd lsdsk > /home/grid/disklist

# In oracle 10g the above used command 'asmcmd' is not available so use
SQL
# query can be used to take the disk list. Need to uncomment the following
# line and comment the above incase oracle 10g is being in use.
# The disk_list.sql script is availbe in this noasm lib examples folder
itself
# which can be modified as per customer needs.
# sqlplus "/as sysdba" @/home/grid/disk_list.sql > /home/grid/disklist

```

disk_list.sql

```

# su - oracle
-bash-4.1$ cat disk_list.sql
select path from v$asm_disk;
exit
-bash-4.1$

```

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_host name	Yes	On
Data files	oradata_sid	Yes	Off
Temporary data files	oratemp_sid	Yes	Off
Control files	oracntrl01_sid (Multiplexed) oracntrl02_sid (Multiplexed)	Yes	Off
Redo logs	oralog01_sid (Multiplexed) oralog02_sid (Multiplexed)	Yes	Off
Archive logs	oraarch_sid	Yes	Off

Real Application Clusters (RAC) databases

File types	Volume names	Dedicated volume for file types	Automatic Snapshot copies
Oracle binaries	orabin_host name	Yes	On
Data files	oradata_dbname	Yes	Off
Temporary data files	oratemp_dbname	Yes	Off
Control files	oracntrl01_dbname (Multiplexed) oracntrl02_dbname (Multiplexed)	Yes	Off
Redo logs	oralog01_dbname (Multiplexed) oralog02_dbname (Multiplexed)	Yes	Off
Archive logs	oraarch_dbname	Yes	Off
Cluster files	oracrs_clustername	Yes	On

Single instance of an Automatic Storage Management (ASM) database

File types	Volume names	LUN names	Dedicated volume for file types	Automatic Snapshot copies
Oracle binaries	orabin_host name	orabin_host namelun	Yes	On
Data files	oradata_sid	oradata_sidlun	Yes	Off
Temporary data files	oratemp_sid	Oratemp_sidlun	Yes	Off
Control files	oracntrl01_sid (Multiplexed) oracntrl02_sid (Multiplexed)	oracntrl01_sidlun (Multiplexed) oracntrl02_sidlun (Multiplexed)	Yes	Off
Redo logs	oralog01_dbname (Multiplexed) oralog02_dbname (Multiplexed)	oralog01_dbnamelu n (Multiplexed) oralog02_dbnamelu n (Multiplexed)	Yes	Off

File types	Volume names	LUN names	Dedicated volume for file types	Automatic Snapshot copies
Archive logs	oraarch_sid	Oraarch_sidlun	Yes	Off

ASM RAC databases

File types	Volume names	LUN names	Dedicated volume for file types	Automatic Snapshot copies
Oracle binaries	orabin_host name	orabin_host namelun	Yes	On
Data files	oradata_sid	oradata_sidlun	Yes	Off
Temporary data files	oratemp_sid	Oratemp_sidlun	Yes	Off
Control files	oracntrl01_sid (Multiplexed) oracntrl02_sid (Multiplexed)	oracntrl01_sidlun (Multiplexed) oracntrl02_sidlun (Multiplexed)	Yes	Off
Redo logs	oralog01_dbname (Multiplexed) oralog02_dbname (Multiplexed)	oralog01_dbnamelu n (Multiplexed) oralog02_dbnamelu n (Multiplexed)	Yes	Off
Archive logs	oraarch_sid	Oraarch_sidlun	Yes	Off
Cluster files	oracrs_clustername	oracrs_clusternamelun	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 or in an ASM disk group 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

- SnapManager can be configured to catalog database backups with RMAN.

If an RMAN recovery catalog is used, the recovery catalog must be in a different database than the database that is backed up.

- 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

- Backup creation might fail if you run SnapManager operations concurrently on the same host against a different ASM database.
- 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.

- When a backup or clone operation is executed simultaneously on the 10gR2 and 11gR2 RAC databases over ASM, then one of the backup or clone creation operations fails.

This failure is because of a known Oracle limitation.

- 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.

- In an ASM environment, if the Oracle Cluster Registry (OCR) and voting disk files coexist on a disk group that has data files, then the fast restore preview operation displays the wrong directory structure for the OCR and voting disk.

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: mysupport.netapp.com](http://mysupport.netapp.com)

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 Oracle does not support migration of a Vserver, which is not supported by clustered Data ONTAP
- SnapManager for Oracle 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 versions 10gR2, 11gR1, 11gR2 and 12c, but does not support Oracle 10gR1 as the repository or target database.
- SnapManager will not support the use of a SCAN IP address in place of a host name.

SCAN IP is a new feature in Oracle 11gR2.

- SnapManager does not support Oracle Cluster File System (OCFS).
- Oracle 11g in a Direct NFS (dNFS) environment allows additional mount point configurations in the `oranfstab` file, such as multiple paths for load balancing.

SnapManager does not modify the `oranfstab` file. You must manually add any additional properties that you want the cloned database to use, in the `oranfstab` file.

- 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: support.netapp.com/NOW/products/interoperability](http://support.netapp.com/NOW/products/interoperability)

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.
- Raw devices backed by a logical volume manager are not supported for protection of Oracle data.

Raw device storage and Automatic Storage Management (ASM) disk groups must be provisioned directly on physical devices. In some cases, partitioning is required.

- A Linux environment without logical volume management requires a partition.

Copyright Information

Copyright © 2021 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

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.