



Technical report:

IBM System Storage N series with SnapManager for Oracle

Best-practice database backup, recovery, and clone management

• • • • • • • • •

Document NS3554-0

May 22, 2008



Table of contents

Abstract	3
Introduction	3
Background on technical issues	3
Purpose and scope	3
Intended audience	4
SnapManager simplifies Oracle database management	5
SMO concepts	5
Centralized administration.....	5
Host and protocol agnostic.....	6
Repository and database schema.....	6
SMO profiles.....	6
Policy-based management.....	8
Architecture overview.....	9
Understanding SMO components.....	9
IBM N series Data ONTAP	9
Database host	9
SnapDrive for UNIX.....	10
SMO	10
Oracle database configuration	10
Oracle database ASM configuration	11
Oracle RMAN itegration	12
SMO GUI vs. CLI	12
N series and SMO host installation and configuration	13
Backup, recovery, and cloning with SMO	13
Backups with SMO.....	13
SMO best-practice backup recommendations	15
Recovery with SMO	16
SMO best-practice recovery recommendations	18
Database cloning with SMO	19
SMO best-practices cloning recommendations	19
Summary	20
Appendix: Performing a block-level restore, using RMAN tasks	21
Trademarks and special notices	24



Abstract

IT enterprise operations continually strive for better solutions to solve critical availability, manageability, and total cost of ownership issues. Backup, recovery, and cloning are all terms that are synonymous with effective database management. IBM System Storage N series with SnapManager for Oracle alleviates Oracle database management complexity by managing the Oracle Database, backup, recovery, and cloning process with an easy to use GUI and a simplified CLI. SnapManager eases management of multiple Oracle databases residing on IBM N series storage via additional IBM System Storage N series features and solution options. Plus, SnapManager seamlessly integrates with additional Oracle database management technologies.

Introduction

The demand for business-critical data availability is driving the need for fast backup, recovery, and cloning of enterprise data. As databases continue to grow in size, tools required to manage this growth need to be increasingly efficient, reliable, and scalable. The database administrator (DBA) and storage administrator require the capability to provision storage as needed while providing reliable methods for database data protection.

IBM® System Storage™ N series with SnapManager® for Oracle (SMO) simplifies Oracle Database management by providing efficient and automated methods for database backup, recovery, and cloning. SMO leverages IBM System Storage N series with Snapshot™ to speed up database backups and hence facilitates increased backup frequency. In addition to simplified backups, the restore process is also faster and more efficient using SMO.

In the enterprise world today, database cloning for test and development purposes also poses significant challenges for DBAs. SMO not only simplifies database cloning but also offers significant time and space savings. The database cloning mechanism utilized in SMO uses IBM System Storage N series with FlexClone™ technology for space-efficient and time-efficient full database clones.

Background on technical issues

This report provides technical background on the features and values that IBM N series with SnapManager solution brings when used in an Oracle database environment. The report further highlights how SMO increases Oracle environment flexibility and reduces TCO as it relates to managing Oracle database backup, recovery, and cloning.

Purpose and scope

Oracle DBAs face daily challenges in database installation, configuration, upgrades, maintenance, growth planning, and establishing sound backup/recovery policies and procedures. These tasks include in great detail the ability to do the following:

- Maintain database installs and configurations in all development, staging, quality assurance (QA), and production environments
- Monitor and manage database performance for database applications to meet and/or exceed service-level agreements (SLAs)
- Perform database management tasks required for loading, editing, removing, and restoring data requested by customer



- Assist with performance and capacity planning
- Perform backup and recovery of all corporate databases
- Meet companies' SLAs as defined by the recovery time objective (RTO) and recovery point objective (RPO)
- Ensure that all tasks are completed within deadlines
- Perform database structural changes without affecting the production database
- Facilitate database copies for testing, performance, and application tuning
- Plan, test, coordinate, and implement strategies for Oracle, operating system (OS), and storage patches, as well as upgrades.

Each one of these tasks requires knowledge and skills while consuming time and resources. Each one of these tasks is dependent upon being able to access data or recover from failed attempts at a patch application. They are also dependent upon storage availability, access, and resources to properly manage, backup, copy, clone, and restore the database. At the core of any DBA's existence is the ability to plan and perform consistent backups of data.

The goal of this report is to provide recommendations for deploying IBM N series SMO.

Intended audience

The intended audience for this report is an Oracle DBA, storage administrator, or someone in charge of the architecting and implementation of the backup, recovery, and cloning process to be used in the Oracle environment. Readers should have a solid understanding of both Oracle and IBM N series technologies.

Although the topic of Oracle database administration is highly complex in nature, this report will help the reader to have a better understanding of what steps are required in order to fully utilize the IBM N series SMO solution.



SnapManager simplifies Oracle database management

One of the core components of N series storage system is the IBM System Storage N series with Data ONTAP® OS. Starting with N series Data ONTAP 7.1, IBM System Storage N series with FlexVol™ technology has allowed users to create highly flexible logical partitions on N series storage systems. In addition to the flexible volumes, Data ONTAP 7.1 or later also offers the ability to use flexible clones. N series with FlexClone allows a writable point-in-time copy of a particular flexible volume. This technology integrates seamlessly with the fast and efficient IBM N series Snapshot technology and IBM System Storage N series with SnapRestore® capability inherent in Data ONTAP 7.1 and later, delivering adding recovery and restore capabilities in Oracle database environments.

Founded in such Data ONTAP advantages, SMO is a host-based application that simplifies Oracle database management tasks by surpassing these very same management complexities and reducing the human error involved. SMO, therefore, provides efficient disk-based backups of Oracle databases while utilizing N series Snapshot technology. In addition to fast backups, N series with SMO supports rapid restore and recovery of a failed Oracle database instance within minutes. Over and above simplifying the management of backups is the ability to create clone copies of database instance within minutes.

SMO concepts

Centralized administration

SMO provides centralized administration of multiple Oracle databases via a graphical user interface (GUI). The GUI provides simple to use database task wizards to perform backup, restore, and clone operations.

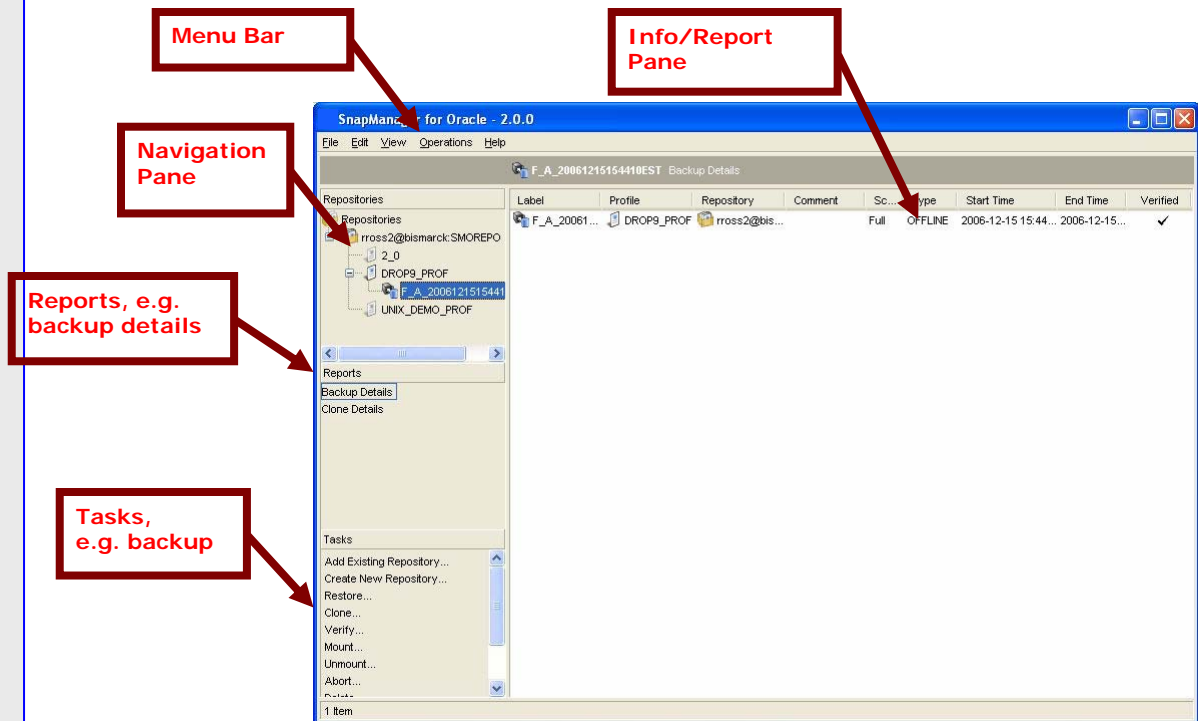


Figure 1) SMO GUI.

Figure 1 highlights the easy-to-use SMO GUI. Centralized administration allows DBAs to administer multiple databases across host systems, which may include managing policies and performing database backup, restore, and clone operations. In addition, reporting and processing information is available for the DBA or storage administrator via a single console.

Host and protocol agnostic

SMO supports multiple protocols and is thus protocol agnostic. SMO 2.0 works seamlessly with Oracle databases using network file system (NFS), internet SCSI (iSCSI), or fibre channel (FC) protocols.

Repository and database schema

N series SMO stores database and backup metadata information in a repository that resides in the Oracle database. This repository consists of a schema that is created upon the initial installation of the product. It is not required that the repository reside on the same host that will be managed by SMO.

SMO profiles

After SMO install, one of the first tasks to complete is a profile creation for the database to be managed. Figure 2 highlights a sample profile.

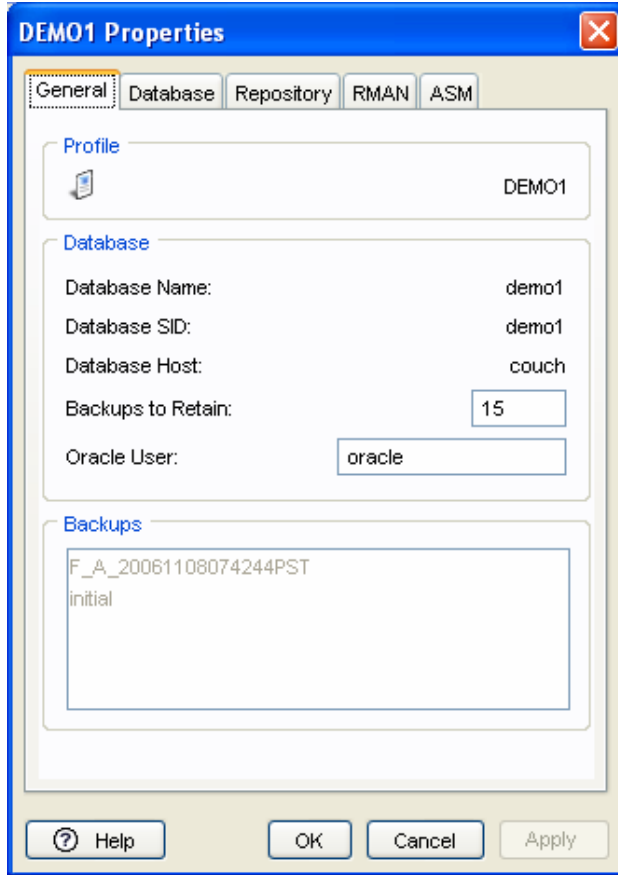


Figure 2) Sample SMO profile.

An SMO profile consists of the following components and is able to leverage both Oracle Recovery Manager (RMAN) and Oracle Automatic Storage Management (ASM) options:

Database credentials	Stored in an encrypted file using DES
Database username/password	Required to access the database
Repository username/password	Required to access the repository
Hostname/username/password	Required to access the host machine with the OS
Number of backups to retain	Backup policy for a profile
RMAN information (optional)	Information necessary to access the RMAN repository connection information if used by database
ASM information (optional)	Information necessary to access an ASM disk group if used by database

Table 1) SMO profile components.

Once the SMO profile is created, information specific to that database will be stored in the repository. Afterward, it is no longer required to specify credentials during subsequent database management operations. Each profile manages one specific database. As an example, if a profile using an RMAN backup is created, then RMAN will be used every time when a database backup/restore operation is initiated via that profile.

Policy-based management

IBM N series SnapManager – and, thus, N series SMO – maintains a sliding window that identifies how many backups are to be kept online via policy. The picture below highlights an example of N series SMO policy-based management.




Figure 3) SMO policy-based management.

As indicated in Figure 3, the number of backups to retain is three. Once the fourth successful backup is executed and completed, the oldest backup will be pruned and the oldest snapshot copy will be removed. The entries in the SMO repository will be updated to reflect the deleted backup as well. In addition, if the backup was registered with RMAN, appropriate changes are also made in RMAN catalog if necessary. Unsuccessful or failed backups are not included in the policy-based sliding window. The backups used to create database clones are also not part of the sliding window.

Architecture overview

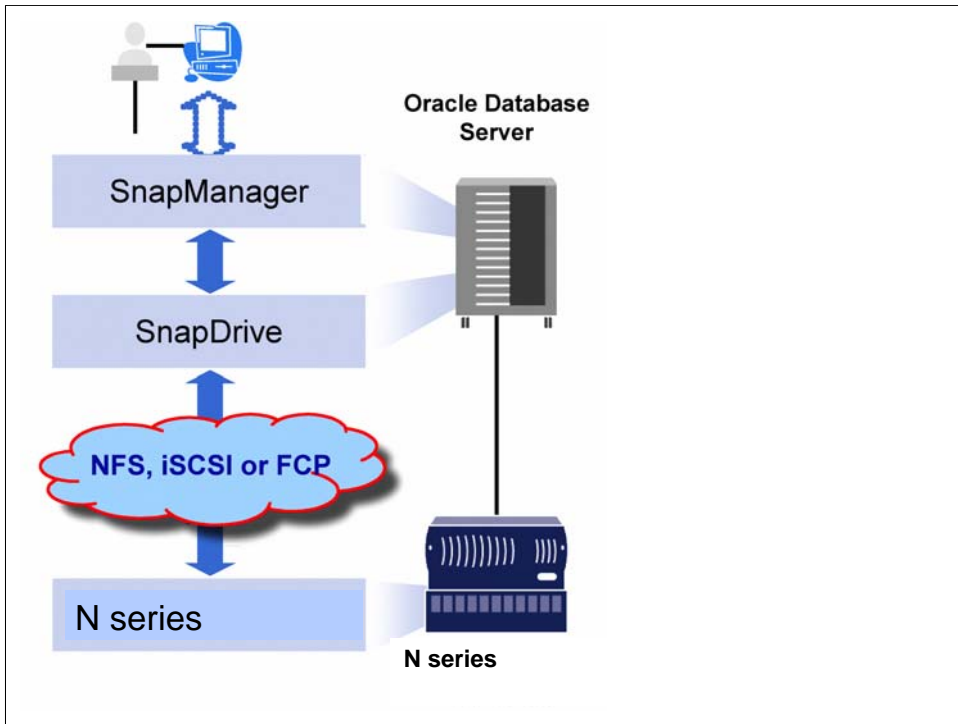


Figure 4) SMO architecture overview.

Figure 4 illustrates the SMO architecture and the components that work together to provide this comprehensive and powerful backup, restore, recovery, and cloning solution for Oracle databases. Along with the easy to use GUI, SMO provides a command-line interface (CLI) that can be executed from the host system where IBM N series with SMO is installed. In other words, all the database management tasks also can be performed via the CLI option. This provides DBAs the flexibility and power to develop scripted solutions if required.

Understanding SMO components

IBM N series Data ONTAP

N series SMO supports Data ONTAP 7.1 and later. It is recommended that the Oracle databases residing on N series storage use flexible volumes. As a result, make sure that following licenses are enabled on the N series storage system:

- A license for at least one protocols (NFS, iSCSI, or FC)
- A license for IBM N series FlexClone (mandatory if using NFS, but optional in FC and iSCSI environments)
- A license for IBM N series SnapRestore.

Database host

IBM N series SMO supports a variety of OS options for database management. Refer to an SMO compatibility matrix for more details.



SnapDrive for UNIX

IBM N series SMO requires that IBM System Storage N series with SnapDrive® for UNIX® (SDU) is installed on the database host before the SMO install (SDU, of course, applies to a UNIX environment; SnapDrive also is available for Windows). SDU simplifies storage management, reduces operational costs, and improves storage efficiency. It automates storage provisioning tasks and simplifies the process of creating error-free, host-consistent data snapshot copies. It provides the capability to traverse the host storage stack and perform the right tasks required for streamlined storage management. On the database host, all communication happens through SDU, thus eliminating all user intervention for database management. At the time of this report writing, SDU 2.2 or later is an appropriate version for use.

SMO

SMO has two components:

- The SMO client (GUI client)
- The SMO server, which resides on the database server.

For SMO 2.0 or later, an installer is available for the SMO client and SMO server. By executing a single command, the user can install and setup the SMO client or server. Note that the SMO installer will automatically install a Java™ Runtime Environment (JRE) for SMO to use. However, this JRE is not made available to the entire system. For instance, on Microsoft® Windows the JRE for SMO will *not* show up in the list of installed programs. It is required that an SMO server be installed on every database server that is to be managed via SMO.

Oracle database configuration

After SMO is installed, the SMO repository will need to be created. When SMO client is started (for the first time only) after install, an existing Oracle database must be available for SMO repository creation. This repository will serve as a recordkeeping for all database credentials, backup, restore, and cloning activity details.

It is not mandatory but a recommended practice is to store the SMO repository database on an N series storage system to allow quick backup and restore capability. At the time of this report writing, SMO cannot create backups/restores of its own repository. However, cross backups of multiple SMO instances are fully supported. (That is, dbX has a repository of dbY, and vice versa.) If the repository database is hosted on N series storage, it is not difficult to backup/restore repository using traditional N series Snapshot and SnapRestore technologies. The database that will be used for the repository can be of minimal size, and it is a best practice to create regular repository backups as well.

Finally, the main database to be managed by SMO should be configured where each of its data files, redo logs, flashback recovery area, and archive logs are stored on the N series storage system. It is also a best practice to not use symbolic links to the location of the data files if linked from local storage to NFS.



Oracle database ASM configuration

SMO provides a seamless integration with Oracle ASM and supports backup, restore, and cloning of ASM disk groups. Keep in mind that if the ASM disk group contains logical unit numbers (LUNs) residing on multiple volumes, consistency group support is required. Consistency group support is provided in N series Data ONTAP 7.2 and later and allows creation of consistent snapshot copies of multiple volumes.

Oracle RMAN integration

N series SMO provides transparent integration with Oracle RMAN. Figure 5 highlights the SMO and RMAN integration process.

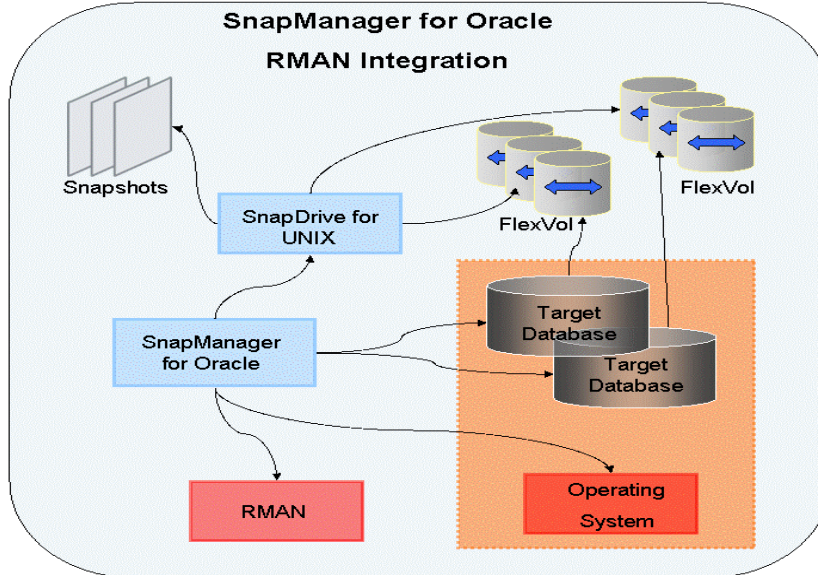


Figure 5) SMO and RMAN integration.

If RMAN is used, N series SMO can catalog the backups it makes with RMAN. Since the backups are registered with RMAN, a DBA can make use of the SnapManager backups and preserve the value of all RMAN functions, such as block-level restore. Refer to the report appendix for an example of block-level restore using RMAN. SMO lets RMAN use the snapshot copies when it performs recovery or restore. For example, you can use RMAN to restore a table within a tablespace, as well as perform full database and tablespace restores and recoveries from snapshot copies created by SMO.

SMO GUI vs. CLI

N series SMO can be managed either from a GUI or CLI. The CLI option provides the additional benefit of allowing scripting capability. All SMO commands can be executed either through the GUI or CLI. The only caveat is that the SMO CLI-based commands need to be executed on the database host where the SMO server product is installed. The SMO CLI can also be executed as a cron job or scheduled through Oracle Enterprise Manager (EM). Utilizing the Oracle EM Grid Control function, SMO can be scheduled and executed as a job.

N series and SMO host installation and configuration

N series SMO is easy to install and configure. The following is a high-level list of the tasks involved in preparing for and installing SMO.

1. Verify SMO install requirements (for example, SDU, which is host-side, and Data ONTAP 7.1 or higher). Also verify appropriate license requirements (again, FlexClone if using NFS and SnapRestore). For more details, refer to the SMO installation and administration guide.
2. Download a copy of the IBM N series SMO software package, as well as a copy of the GUI, if using it.
3. Follow the directions to install the SMO server component on the database host. If using GUI, install the SMO GUI client on a Windows machine.
4. On the database host, start the N series SMO server component using the `"smo_server start"` command.
5. If installing N series SMO for the first time on this host environment, create a new repository for the profiles to use for the databases. If upgrading SMO, update existing repositories.
6. Once SMO is installed, you can verify that all the components are set up correctly by running `"smo system verify"` command. You can also verify the environment setup from the SMO GUI client.

Backup, recovery, and cloning with SMO

N series SMO simplifies database backup, recovery, or cloning tasks by hiding complexity and eliminating user errors. Database backup, recovery, and cloning are discussed in great detail in the relevant SMO installation and administration guide. This section will highlight some details about how SMO addresses day-to-day DBA challenges by providing database consistent backups, guaranteed restores, and faster clones.

Backups with SMO

With SMO, Oracle database backup operation can be initiated from either a GUI client or a CLI on the database host server. N series SMO lets you create a backup of an entire Oracle database or portions of it. The backup is a virtual copy of the database instance and is stored on the same physical medium as the database. The SMO backup consists of the snapshot copies created at the point in time when the database instance was backed up. It uses snapshot copies to create the database backup. You can then use a particular backup to restore and recover the database information at a later time. SMO offers the following choices for Oracle database backup:

- The entire database instance (full backup)
- A group of tablespaces (partial backup)
- A group of data files (partial backup).

Both the CLI command and the GUI wizard require similar information if you want to backup a database. The basic steps for backing up data are as follows:

1. Specify the type of backup you want to perform:

- **Online:** The database is mounted and running (online), the Oracle database status is OPEN, and archive log mode is enabled.
 - **Offline:** The database is mounted, but not available to users (offline), and the Oracle status is either MOUNTED, STARTED, or SHUTDOWN.
 - **Auto:** N series SMO determines whether the database is online or offline and then performs the appropriate backup.
2. (Optional) Specify the verify option if you want N series SMO to verify the backup using the Oracle utility dbv after the backup completes. When you select this option, SMO checks the block structure of the backup to make sure it is in a valid format for Oracle.

Note: The DBV verification might take several minutes. If you prefer to perform this operation at a later time, you can use the SMO verify later operation.

3. Specify name and labeling optional details for the backup:
 - The name of the profile associated with the database.
 - Whether this is a full or partial backup.
 - (Optional) Label for the backup (i.e., an easily identified name for the backup). If you do not supply a label, SMO creates a default label.
 - (Optional) Enter a comment that describes the backup (for example, you might enter the description “production test” so that you can quickly identify the contents of the backup).
4. (Optional) Enter the -force option if you want SMO to force the database to change its state if it is not in the state you requested for the backup operation.

Below is an example of an SMO online full-backup command for a database:

```
$ smo backup create -online -full -profile SMO_TEST -label GOLD_ONLINE_FULL -force
```

Figure 6 highlights how N series SMO significantly improves the efficiency of database backups from a time and space perspective. The time and space savings are realized by SMO use of N series Snapshot technology.

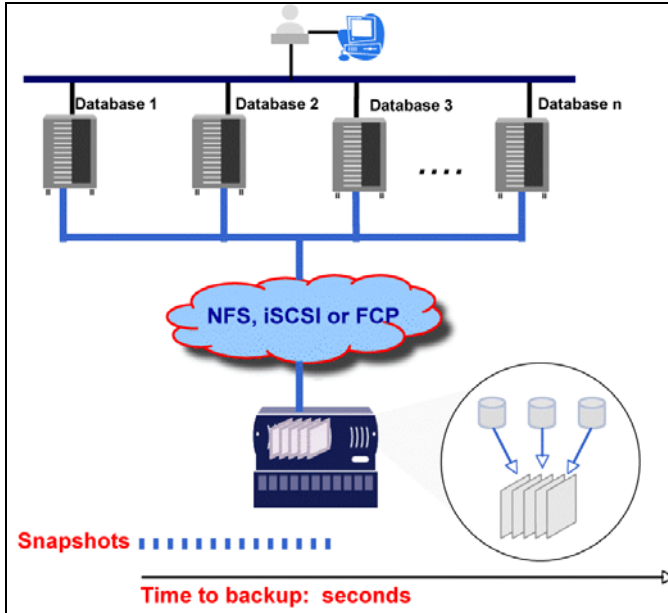


Figure 6) Database backup with SMO.

N series SMO facilitates frequent database backups. The high-performance N series Snapshot also makes it highly scalable. A snapshot copy takes only a few seconds to create — typically less than one second, regardless of the size of the volume or the level of activity on the N series storage system. So, the database is only going to be in hot-backup mode for a very short duration while N series SMO creates snapshot copies of database volumes. Snapshot copies are also very space efficient as snapshot copies created in sequence differ one from the next by the blocks added or changed in the time interval between the two. This block-incremental behavior limits associated storage capacity consumption and offers huge space savings.

As snapshot copies are extremely fast and space efficient, DBAs can afford to create frequent database backups with N series SMO. Below is the summary of recommendations one should consider while using SMO to create database backups:

- It is highly recommended to create frequent database backups as SMO facilitates faster and space efficient backups.
- It is required that each label is unique, including labels used for failed attempts.
- The -force option is required if the database is not mounted (database is in a shutdown state).
- It should be noted that SMO does not back up external tables.
- It is required that database archive logs be enabled in order to complete an online backup.
- All database-related files to back up will need to be stored on the N series storage system.
- The tnsnames and listener should be available and listening on the proper port for SMO to execute.

SMO best-practice backup recommendations

The following are a few SMO best-practices for backup:

- **It is recommended that backups are archived or mirrored off the mainline storage.** This can be accomplished by using IBM System Storage N series with SnapVault® or IBM System Storage N series with SnapMirror® technologies. The classic Oracle backup and archiving problems are solved by adding the capability to back up to an IBM N series ATA-based storage solution – IBM System Storage N series with NearStore® – or a local IBM N series SATA solution – such as IBM System Storage N series N5000. Additional details regarding secondary storage are outside the scope of this report.
- **It is recommended that verification be completed on a regular basis against the completed backups during low activity.** However, it is best to avoid running verifications while snapshot copies are being created.
- **It is recommended that snapshot copies be actively monitored to not overprovision available storage.** At the time of this report writing, the maximum number of snapshot copies that are supported is 255 per volume.

Recovery with SMO

N series SMO provides an Oracle database restore operation that includes a recover option. As a result, SMO can be used to both restore or recover the backup, or SMO can be used to restore the backup and another tool, such as RMAN, can be used to recover the data. The SMO recover option is part of the restore operation, which means you must perform the restore and recover at the same time when using SMO (i.e., you cannot perform a restore operation and then at a later time perform an SMO recover operation).

The SMO backup-restore operation also gives the option of deciding how much of the backup to restore. One can:

- Restore the entire backup.
- Restore only the tablespaces or data files you specify.
- Restore only the control files.

If the recover option is included on the command line, one can recover the database to:

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

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

Both the CLI command and the GUI wizard require similar information if you want to restore and recover a backup. This section highlights the basic steps.

The basic steps for restoring a backup are as follows.

1. Specify the name of the backup. This can be either the label you supplied when you created the backup or the ID N series SMO generated when you created the backup.
2. Specify the name of the profile from which the backup was created.

3. Specify which components of the backup you want to restore:
 - The entire backup
 - Specific tablespace or specific data file
 - Control files.
4. (Optional) Specify that you want to recover the database to:
 - The last transaction
 - An SCN (use only if you are doing a point-in-time recovery of the backup)
 - A specific date and time (use only if you are doing a point-in-time recovery of the backup).
5. (Optional) Decide if you want N series SMO to change the state of the database. If you include the `-force` flag, SMO changes the database state, if necessary. For example, when you include this option, if the database is offline and it needs to be online, SMO forces it online.

Using the CLI, below is an example of SMO recovery to restore and recover a full backup:

```
$ smo restore -profile SMO_PROFILE -label myGoldBackup -complete -recover -alllogs
$ smo restore -profile SMO_PROFILE -label myGoldBackup -controlfiles
```

Figure 7 highlights how N series SMO significantly speeds up recovery time. As highlighted in the previous sections, using the unique Snapshot feature of the N series Data ONTAP OS, N series SnapRestore can quickly restore lost files or volumes from stored snapshot copies. From a single file to a huge production database, SnapRestore does the job in seconds regardless of file or volume size.

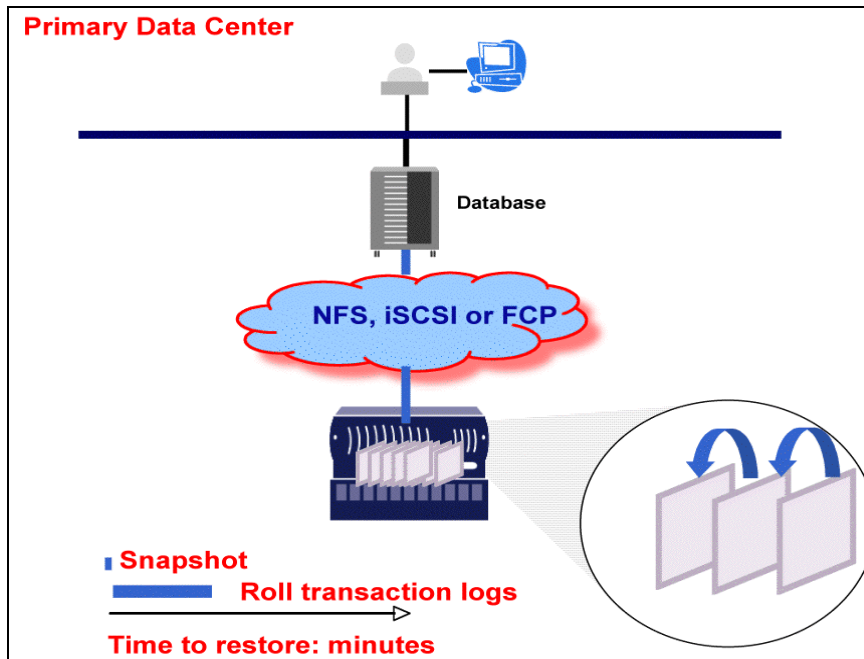


Figure 7) SMO database restore with SnapRestore.

The table below compares the traditional database recovery process with the N series recovery process using SnapRestore. Advantages are the same in an SMO environment.



Traditional database recovery	SnapRestore database recovery
<ul style="list-style-type: none"> • 1 TB database, and all of it requires recovery • Tape recovery at 100 GB/hour • Total recovery: 10 hours + log replay time 	<ul style="list-style-type: none"> • 1 TB database, and all of it requires recovery • SnapRestore recovery duration — seconds • Total recovery: seconds + log replay time

Table 2) Comparison of traditional and SnapRestore database recovery.

N series SMO simplifies and significantly speeds up the database recovery process. SMO hides complexity from users by providing various recovery options – full recovery, partial recovery, point-in-time recovery, or SCN-based recovery. SMO completes all the steps required for database recovery from the end user depending upon the type of recovery selected. By facilitating fast and frequent backups, SMO significantly reduces time required for log replay.

SMO best-practice recovery recommendations

The following best-practices are recommended when using SMO for recovery:

- Avoid using partial database recovery when possible.
- The -force option is required if the database is not mounted (database is in a shutdown state).
- When requiring access to a backup through RMAN, you will need to mount the backup. This is handled by executing the `smo manage mount/unmount ...` command. The appropriate N series SMO installation and administration guide provides the syntax in using the mount and unmount command options.
- It is recommended that any database backups that are no longer in use be removed.

The following table summarizes the options available with the N series SMO backup, restore, and recover operations.

Backup	Restore	Recover Using SnapManager
Entire database instance (Full Backup)	Entire database backup	To a specific date To an SCN To the time of the backup To the last transaction made on the database instance
Entire database instance (Full Backup)	One of the following: Tablespaces Data files Control files	To the last transaction made on the database instance
One of the following: Tablespaces Data files	Entire backup (this is either a set of tablespaces, a set of data files, or the control files)	To the last transaction made on the database instance
One of the following: Tablespaces Data files	One of the following: Tablespaces Data files Control files	To the last transaction made on the database instance

Table 3) Options with N series SMO backup, restore, and recover operations.

Database cloning with SMO

The N series SMO clone feature lets you quickly create a copy of a database instance that is an exact replica of the original database. You create the clone from a full backup or from the current state of the database. SMO offers the following advantages around database cloning:

1. **Speed.** SMO can instantaneously clone data volumes.
2. **Space efficiency.** Creating a clone using SMO requires essentially no overhead. The clone is a writable snapshot copy of the original Oracle database instance, so, like all snapshot copies, it requires space only for the changes between the backup and the clone. It can also grow as needed. In contrast, a physical clone of the database requires that you have sufficient space available to copy the entire database.
3. **Virtual copy.** The Oracle database clone is a virtual copy of the original database. For example, you can use a clone for such operations as testing and bug fixing, platform and update checks, multiple simulations against a large data set, and remote office testing and staging. People can work with the cloned database instance as if it were the original database; however, their changes to the clone do not affect the original.
4. **Simplicity.** SMO simplifies the database cloning process by hiding traditional Oracle database cloning management complexity.

The SMO clone operation uses the N series FlexClone feature available with Data ONTAP 7.1 or later (see Data ONTAP documentation for more information on FlexClone). The basic steps for creating a clone are as follows:

1. Create a full backup of the database instance while it is offline or online.
2. Create a clone specification file for the database instance. N series SMO creates the clone specification file from the information in the backup, or you can manually create the specification file. SMO uses this file to correctly create the clone, based on the information in the original database instance and clone specification.
3. Create the clone.

SMO supports creating a clone of an Oracle database from a backup or from the current state of the database. In other words, you can use an existing database backup to create a clone or perform the backup and clone in one step. SMO only supports cloning a database on the same host. So, you must give the clone a new Oracle System Identifier (SID). As part of the cloning process, SMO creates the required Oracle files and parameter file (`init<SID>.ora`) for the cloned database. When SMO clones the storage for a database, it creates a new file system mountpoint, but does not change the directory structure under the mountpoint. In the clone process, SMO creates `init<clone_SID>.ora` (or `spfile` if clone for clone) and `bdump`, `cdump` and `udump` directories for clone. SMO allows you to clone an Oracle Real Application Cluster (RAC) database instance as well as a nonclustered Oracle database instance. After cloning, the RAC clone starts as a single database instance. For more information on Oracle database cloning, see the appropriate SMO installation and administration guide.

SMO best-practices cloning recommendations

Below are some of the best-practice recommendations for cloning using N series SMO:



- N series SMO only supports cloning a database on the same host.
- It is highly recommended to tone down the system global area and other resource-consuming parameters for the clone database. You can modify these parameters in the clone specification file before clone creation is initiated.
- You can also use IBM System Storage N series with FlexShare™ to assign higher priorities to the production database.

Summary

N series SMO provides a rich feature set that allows the IT organization to take advantage of fast and efficient backup, recovery, and cloning for Oracle databases. This flexibility further reduces the total cost of ownership of the Oracle database and storage infrastructure. N series assists in leading the way to bring DBAs closer to the storage, offers iSCSI, FC, and NFS implementations, while creating flexible tools such as N series SMO to maintain multiple database environments from a single interface. Whether the database environment is in multiple volumes or a single volume, SMO can flawlessly manage database backup, recovery, and cloning processes. This flexibility increases the quality of work while allowing DBAs to meet SLAs and become more efficient in delivering goals.



Appendix: Performing a block-level restore, using RMAN tasks

The following are the basic tasks involved in backing up a database using N series SMO and recovering that backup using RMAN:

- Perform a full offline backup of the database using SMO.
- Verify the backup is catalogued with RMAN.
- Verify the blocks have been backed up correctly.
- Mount the backup using SMO.
- Using RMAN, perform the block restore.
- Verify the blocks have been repaired.

Here are the details, including commands.

1. Full offline backup with N series SMO.

Example:

```
smo backup create -offline -full -profile e -label full_backup -verbose
```

Output:

```
SMO-07109 [INFO ]: Cataloguing all files in backup set with RMAN,  
TAG=SMO_full_backup_1158773581857, RMAN=ES0/controlfile.  
...  
SMO-13037 [INFO ]: Successfully completed operation: Backup  
SMO-13048 [INFO ]: Operation Status: SUCCESS  
SMO-13049 [INFO ]: Elapsed Time: 0:02:20.506  
Operation Id [ff8080810dcc47e3010dcc47eb7a0001] succeeded.
```

2. Verify that backup is cataloged with RMAN.

Example: At the RMAN prompt, do the following:

```
RMAN> list datafilecopy tag SMO_full_backup_1158773581857;
```

Output:

```
using target database control file instead of recovery catalog  
List of Datafile Copies  
Key File S Completion Time Ckp SCN Ckp Time Name  
-----  
335 1 A 20-SEP-06 1347825 20-SEP-06  
/opt/IBMN/smo/mnt/_mnt_weever_revenge_ES0_SMO_E_ES0_F_C_0_ff8080810dcc47e3010dc  
c47eb7a0001/system01.dbf  
336 2 A 20-SEP-06 1347825 20-SEP-06  
/opt/IBMN/smo/mnt/_mnt_weever_revenge_ES0_SMO_E_ES0_F_C_0_ff8080810dcc47e3010dc  
c47eb7a0001/undotbs01.dbf  
334 3 A 20-SEP-06 1347825 20-SEP-06  
/opt/IBMN/smo/mnt/_mnt_weever_revenge_ES0_SMO_E_ES0_F_C_0_ff8080810dcc47e3010dc  
c47eb7a0001/sysaux01.dbf  
333 4 A 20-SEP-06 1347825 20-SEP-06  
/opt/IBMN/smo/mnt/_mnt_weever_revenge_ES0_SMO_E_ES0_F_C_0_ff8080810dcc47e3010dc  
c47eb7a0001/user01.dbf  
337 5 A 20-SEP-06 1347825 20-SEP-06
```



3. Assuming users01.dbf is corrupted, run dbverify as shown below.

Example:

```
dbv FILE=user01.dbf
```

Output:

```
DBVERIFY: Release 10.2.0.1.0 - Production on Wed Sep 20 13:35:44 2006  
Copyright (c) 1982, 2005, Oracle. All rights reserved.  
DBVERIFY - Verification starting : FILE = user01.dbf
```

Page 625 is marked corrupt

```
Corrupt block relative dba: 0x01400271 (file 5, block 625)  
Bad header found during dbv:  
Data in bad block:  
type: 240 format: 6 rdba: 0xed323b81  
last change scn: 0x6f07.faa74628 seq: 0x87 flg: 0x02  
spare1: 0x60 spare2: 0x5 spare3: 0xef7d  
Draft  
112 Performing a block level restore using RMAN  
consistency value in tail: 0xa210fe71  
check value in block header: 0x13c7  
block checksum disabled
```

Page 626 is marked corrupt

```
Corrupt block relative dba: 0x01400272 (file 5, block 626)  
Bad header found during dbv:  
Data in bad block:  
type: 169 format: 2 rdba: 0x5b886118  
last change scn: 0x9bfb.a2ffb25c seq: 0xd1 flg: 0xb4  
spare1: 0x38 spare2: 0xed spare3: 0xf9cb  
consistency value in tail: 0xfc3469f2  
check value in block header: 0xf570  
computed block checksum: 0x98be
```

Page 627 is marked corrupt

```
Corrupt block relative dba: 0x01400273 (file 5, block 627)  
Bad header found during dbv:  
Data in bad block:  
type: 158 format: 7 rdba: 0x2101e16d  
last change scn: 0xe828.42414628 seq: 0xb4 flg: 0xff  
spare1: 0xcc spare2: 0x81 spare3: 0x8665  
consistency value in tail: 0x46d20601  
check value in block header: 0x1a84  
computed block checksum: 0x6c30
```

```
DBVERIFY - Verification complete  
Total Pages Examined : 1280  
Total Pages Processed (Data) : 1123  
Total Pages Failing (Data) : 0  
Total Pages Processed (Index): 0  
Total Pages Failing (Index): 0  
Total Pages Processed (Other): 34  
Total Pages Processed (Seg) : 0  
Total Pages Failing (Seg) : 0  
Total Pages Empty : 120  
Total Pages Marked Corrupt: 3  
Total Pages Influx : 0  
Highest block SCN : 1337349 (0.1337349)  
...
```



4. Mount the backup using N series SMO.

Example:

```
smo backup mount -profile e -label full_backup -verbose
```

Output:

```
SMO-13046 [INFO ]: Operation GUID: ff8080810dcc4ce6010dcc4cf0f00001
SMO-08052 [INFO ]: Beginning to connect mount(s) [/mnt/weever/revenge_ES0] from
logical snapshot SMO_E_ES0_F_C_0_ff8080810dcc47e3010dcc47eb7a0001.
SMO-08025 [INFO ]: Beginning to connect mount /mnt/weever/revenge_ES0 from snapshot
SMO_E_ES0_F_C_0_ff8080810dcc47e3010dcc47eb7a0001_0 of volume revenge_ES0.
SMO-08027 [INFO ]: Finished connecting mount /mnt/weever/revenge_ES0 from snapshot
SMO_E_ES0_F_C_0_ff8080810dcc47e3010dcc47eb7a0001_0 of volume revenge_ES0.
SMO-08053 [INFO ]: Finished connecting mount(s) [/mnt/weever/revenge_ES0] from
logical snapshot SMO_E_ES0_F_C_0_ff8080810dcc47e3010dcc47eb7a0001.
SMO-13037 [INFO ]: Successfully completed operation: Backup Mount
SMO-13048 [INFO ]: Operation Status: SUCCESS
SMO-13049 [INFO ]: Elapsed Time: 0:00:14.258
Operation Id [ff8080810dcc4ce6010dcc4cf0f00001] succeeded.
```

5. Perform block recovery using RMAN.

Example:

```
RMAN> blockrecover datafile '/mnt/weever/revenge_ES0/file01.dbf' block 625, 626,
627 from tag SMO_full_backup_1158773581857;
```

Output:

```
Starting blockrecover at 20-SEP-06
using target database control file instead of recovery catalog
allocated channel: ORA_DISK_1
channel ORA_DISK_1: sid=153 devtype=DISK
channel ORA_DISK_1: restoring block(s) from datafile copy
/opt/IBMN/smo/mnt/_mnt_weever_revenge_ES0_SMO_E_ES0_F_C_0_ff8080810dcc47e3010dcc47e
b7a0001/user01.dbf
starting media recovery
media recovery complete, elapsed time: 00:00:01
Finished blockrecover at 20-SEP-06
```

6. Verify the blocks have been repaired by running dbverify.

Example:

```
dbv FILE=user01.dbf
```

Output:

```
DBVERIFY: Release 10.2.0.1.0 - Production on Wed Sep 20 13:40:01 2006
Copyright (c) 1982, 2005, Oracle. All rights reserved.
DBVERIFY - Verification starting : FILE = user01.dbf
DBVERIFY - Verification complete
Total Pages Examined : 1280
Total Pages Processed (Data) : 1126
Total Pages Failing (Data) : 0
Total Pages Processed (Index): 0
Total Pages Failing (Index): 0
Total Pages Processed (Other): 34
Total Pages Processed (Seg) : 0
Total Pages Failing (Seg) : 0
Total Pages Empty : 120
Total Pages Marked Corrupt : 0
Total Pages Influx : 0
Highest block SCN : 1337349 (0.1337349)
```



Trademarks and special notices

© International Business Machines 1994-2008. IBM, the IBM logo, System Storage, and other referenced IBM products and services are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. All rights reserved.

Data ONTAP, FlexClone, FlexShare, FlexVol, NearStore, Network Appliance, the Network Appliance logo, SnapDrive, SnapManager, SnapMirror, SnapRestore, Snapshot and SnapVault are trademarks or registered trademarks of Network Appliance, Inc., in the U.S. and other countries.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

Information is provided "AS IS" without warranty of any kind.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.