



# Technical report: Oracle and IBM N series Storage Systems

*Applications for LUN cloning*

• • • • • • • • •

*Document NS3266-0*

October 3, 2007



## Table of contents

<b>Abstract</b> .....	<b>3</b>
<b>Introduction</b> .....	<b>3</b>
<b>Assumptions</b> .....	<b>3</b>
<b>Infrastructure</b> .....	<b>4</b>
<b>Cloning Databases Using Writeable Snapshots</b> .....	<b>4</b>
Single-host clone creation .....	5
Second-host clone creation .....	6
Considerations for volume management software .....	8
<b>Seeding Test Environments Using LUN Cloning</b> .....	<b>8</b>
<b>Reorganizing Read-Mostly Databases Using Writeable Snapshots and LUN Cloning</b> .....	<b>8</b>
<b>Sample Scripts</b> .....	<b>10</b>
clonelun.ksh: .....	10
clonelun.bat:.....	11
begin_backup.sql:.....	11
end_backup.sql:.....	11
startup.sql: .....	11
<b>Conclusions</b> .....	<b>12</b>
<b>Caveats</b> .....	<b>12</b>
<b>Trademarks and Special Notices</b> .....	<b>13</b>



## Abstract

---

*IBM System Storage N series leverages unique Snapshot capabilities to deliver snapshot-backed LUN cloning to enable space-efficient and expedient creation of cloned Oracle databases for reporting, test environments, and database reorganization.*

## Introduction

---

This document explores the use of writeable IBM® System Storage® N series with Snapshot™-backed logical unit numbers (LUNs) and LUN cloning capabilities in conjunction with an Oracle® database. These features, available in IBM N series storage systems running Data ONTAP® 7.1 or later, enable space-efficient and expedient creation of cloned Oracle databases for reporting, test environments, and database reorganization.

## Assumptions

---

For the information contained in this report to be useful to the reader, several assumptions are made:

- The reader has at least intermediate Windows 2000 administration skills and has access to the administrator login or an account that is a member of the local or domain administrator's group, or has at least intermediate Sun Solaris or other UNIX® system administration skills and has access to the root user account.
- The reader has at least intermediate Oracle 9i™ administration skills and has the Oracle9i installation media or has Oracle 9i already installed on the target system.
- The reader has at least intermediate knowledge of IBM N series administration, and has administrative access to the system via the command-line interface (CLI) as well as via IBM System Storage N series with FilerView®.
- The storage and host have the required licenses necessary to perform the activities outlined in this document.
- The target system must have the required protocol interconnects to perform the activities outlined in this document.
- For the purposes of this report, storage system administration is performed via telnet for clarity. FilerView or scripts utilizing "rsh" may also be used where applicable. Microsoft® Windows® implementations should use the IBM System Storage N series with SnapDrive® console or CLI.



## Infrastructure

---

The specifics of the hardware configuration are not of central importance to this paper, so hardware specifics will not be discussed here. Each scenario as presented will have specific infrastructure requirements, as will the protocol and system used. Software platform differences will affect the specific commands used in each procedure. Infrastructure requirements will be addressed in each section.

## Cloning Databases Using Writeable Snapshots

---

Database cloning is often accomplished by making a complete copy of all database files to an alternate location. In cases where the clone is to be run on the same machine, the database is renamed by recreating the control files.

Writeable IBM N series Snapshot-backed LUNs afford the database administrator the ability to clone a database by taking a snapshot, creating a snapshot-backed LUN, and connecting that LUN to a host. Because of the structure of the write anywhere file layout (WAFL<sup>®</sup>) file system, the clone actually shares unchanged blocks with the original database; no data copy is involved. Changed blocks are written to new areas on disk and the WAFL inode pointers are updated with this location information.

Clone databases created in writeable snapshots should be designed with temporary use in mind. Reporting databases, which normally have a short lifespan, are an excellent application for writeable snapshots.

The procedure for creating clone databases using writeable snapshots will depend on whether the snapshot-backed LUN will be connected to a second host. If the snapshot is connected to the same host, the database name must be changed before the database can be started. If a second host is used, the database name can remain the same.

## Single-host clone creation

Because Oracle cannot allow two databases with identical names to run on the same host, and changing the database's name using datafiles in hot backup mode is not possible without recovering the database first, only cold backups can be used to create single-host clones.

The procedure for creating a clone from a cold backup is detailed in the following steps:

1. Backup the controlfile to the user trace directory using the "alter database backup controlfile to trace" command in SQLPlus.
2. Shut down the target database cleanly. Issuing "shutdown abort" is not an acceptable shutdown method.
3. Create the snapshot of all database LUNs using a distinctive name. At a minimum, the snapshot must contain the redo logs and all database files. If the database runs in "archivelog" mode, the archived logs will be needed.

```
filer> snap create /vol/dbvol snap1
```

4. Create a new parameter file for the new Oracle instance. It is better to copy the original database parameter file and update the parameters to reflect the new drive letters or pathnames to use.
5. Copy the controlfile trace from the user dump directory and edit the file to reflect the new database name. Be sure to insert the "set" command in front of the database name, and verify that all datafile locations point to the clone database drive letter or pathname.
6. Create snapshot-backed LUNs using the "lun clone create" command from the filer command line and map the LUN to the host, as shown in the example. Optionally, use the "-o noreserve" switch to disable space reservation, conserving disk space at the expense of database robustness.

```
filer> lun clone create /vol/dbvol/writeable.lun dbhost
-o noreserve -b /vol/dbvol/.snapshot/snap1/dblun.lun
filer> lun map /vol/dbvol/writeable.lun dbhost
lun map: auto-assigned dbhost=1
```

7. On the host, mount or connect the writable LUNs to the appropriate drive letter or pathname. Ensure that the path used matches the contents of the new parameter file file and controlfile creation script.
8. Run SQLPlus, start the database in NOMOUNT mode and recreate the controlfiles using the modified trace file from the prior step.
9. The database can now be opened.

## Second-host clone creation

Creating a clone database using a second host is a more robust method for creating reporting databases. Since the clone database runs on a second host, the instance does not have to compete for resources with the production instance on the same machine. Also, the second host could be on an entirely different front-end network from the production machine, allowing more flexibility in determining access control. In addition, the database name need not be changed; a different net service name can be configured on the clients, but the Oracle System Identifier (SID) need not be changed.

The two hosts must be connected to the same storage fabric as the filer. The particular protocol used in the fabric is of no consequence; either iSCSI or Fibre Channel (FC) protocols may be employed.

A basic diagram of the architecture is shown in Figure 1:

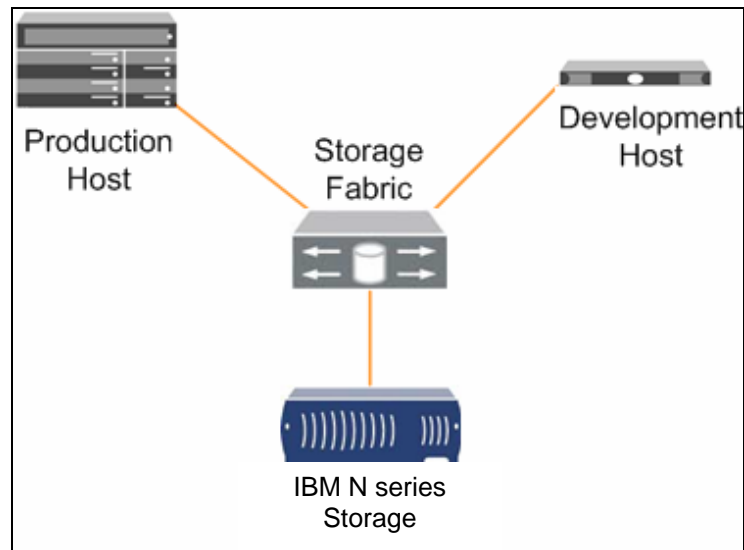


Figure 1. IBM N series architecture.

In this section, it is assumed that the original database is run in “archivelog” mode, as all production database instances should be for recoverability. The production system can remain online during the cloning process; Oracle’s hot backup mode will be used to make the clone. Scripting can be employed for all of these steps to automatically refresh the reporting database.



The creation of a clone database for a second host is detailed in the following steps:

1. Backup the controlfile to the user trace directory using the “alter database backup controlfile to trace” command in SQLPlus.
2. Using the “alter tablespace ... begin backup” command in SQLPlus, put all of the tablespaces in hot backup mode.
3. Create the snapshot of all database LUNs using a distinctive name. At a minimum, the snapshot must contain the redo logs and all database files. Additionally, any archive logs created during the backup will be needed.

```
filer> snap create /vol/dbvol snap1
```

4. Using the “alter tablespace ... end backup” command in SQLPlus, take all of the tablespaces out of hot backup mode.
5. Copy the Oracle parameter file from the production database. Some parameters, such as the database buffer cache, shared pool size, and large pool size can be adjusted for the target system if necessary.
6. Create snapshot-backed LUNs using the “lun clone create” command from the filer command line. Optionally, the “-o noreserve” switch can be used to disable space reservation to conserve disk space at the expense of database robustness.

```
filer> lun clone create /vol/dbvol/writeable.lun -o noreserve  
-b /vol/dbvol/.snapshot/snap1/dblun.lun  
filer> lun map /vol/dbvol/writeable.lun dbhost2  
lun map: auto-assigned dbhost2=1
```

7. On the host, mount or connect the writable LUNs to the appropriate drive letter or pathname. Ensure that the path matches the contents of the new parameter file and controlfile creation script. If this is not possible, you can recreate the controlfiles by using the controlfile trace created previously.
8. Start the database in MOUNT mode.
9. Issue “alter database end backup” to take all tablespaces out of hot backup mode.
10. Issue the “recover automatic database” from the SQLPlus command line. At this point, Oracle may prompt for the archived logs needed to complete the recovery.
11. Issue “alter database open” to make the database available.

At this point, redo log archiving can be disabled if necessary to save disk resources. Any additional setup, such as granting user access, can also be performed before make the database available for use.

## Considerations for volume management software

When using volume management software to manage LUNs, additional procedures and considerations may need to be addressed. Though volume management software is outside of the scope of this report, this section will enumerate some of these issues.

- Most volume management software, including the Symantec Volume Manager, write signatures to the LUN that contain unique host identification information, as well as information about other disks in the group. Additional steps will be required to “import” these groups onto the same host (where possible) or the second host. Importing a snapshot-backed LUN volume onto the same machine as the original should not be attempted; at the time of this report writing, Symantec does not support the procedure required.
- When volume management software is used to stripe or mirror LUNs, all LUNs in a volume set must be created as snapshot-backed LUNs.

## Seeding Test Environments Using LUN Cloning

---

Creating test environments for developer access can be a resource-consuming process. These environments are usually not created with great frequency, and can have a long lifespan.

Any snapshots for which there exist snapshot-backed LUNs will be marked as busy. They cannot be destroyed until that LUN is destroyed. Additionally, the snapshot-backed LUN is created in the active filesystem, which means that it will appear in subsequent snapshots. It is possible to create a dependency on a snapshot such that it cannot be destroyed without destroying all subsequent snapshots. LUN cloning solves this potential problem by making the LUN permanent in the active WAFL filesystem, removing the dependency on the snapshot. Once cloned, the snapshot slot is freed for reuse in the backup rotation.

LUN cloning is performed using the “lun clone” commands from the filer command line. LUNs are cloned to the same pathname as the original. LUN clones inherit their space reservation settings from the snapshot-backed LUN they are cloned from.

The process for creating cloned databases using LUN cloning is the same as previously outlined for reporting databases, with one additional step: the snapshot-backed LUN is cloned after the last step, as in the following command example.

```
filer> lun clone start /vol/dbvol/writeable.lun
filer> lun clone status /vol/dbvol/writeable.lun
lun clone: Done 330112 of 1311298 blocks (25% complete).
Wed Mar 26 10:55:00 EST [vdisk_admin:info]: Cloning completed on lun
/vol/dbvol/writeable.lun
```

After the lun clone command completes, the LUN is permanent and is no longer dependent on the snapshot.

## Reorganizing Read-Mostly Databases Using Writeable Snapshots and LUN Cloning

---

Read-mostly environments generally are data warehouses, where data is loaded in batches and queries are run to extract data. In this case, read-mostly environments could also be considered to be databases where data is inserted only during certain periods of the day.





The LUN manipulation features of Data ONTAP can be employed to allow online reorganization of these types of environments, incurring minimal downtime since the database can be available while the reorganization takes place. The database is then briefly taken offline to complete the switchover to the reorganized copy.

The steps required to implement this type of reorganization are detailed in the following steps:

1. Clone the LUNs containing the database data as shown in the previous sections.
2. If possible, alter all tablespaces in the original database read-only, to avoid data changes that would be lost.
3. Reorganize the database. Reorganization can include dropping or resizing datafiles, migrating data between tablespaces, or any other data manipulation needed.
4. Once reorganization is complete, record the LUN paths and their associated numbers using the "lun show -v" command.
5. Shut down the original database, then unmap and destroy its associated LUNs.
6. Map the cloned LUNs to the original LUNs, preserving the LUN numbering. If reorganization involved combining or removing LUNs, ensure that the host configuration is updated to reflect the change.

## Sample Scripts

---

For reference, sample scripts are included in this report to aid in the creation of automated solutions for database cloning. These scripts may require extensive modification before being used. The scripts will require remote administrative access to both the filer and the source database. In the case of Microsoft Windows, the second host will also need to have administrative privileges on the host where the source database resides. This is normally established through domain credentials.

### clonelun.ksh:

```
#!/bin/ksh
# put oracle in hot backup mode
su - oracle -c "sqlplus /nolog @begin_backup.sql"

# take the snapshot
rsh filer snap create /vol/dbvol snap1

# take the database out of hot backup
su - oracle -c "sqlplus /nolog @end_backup.sql"

# create the snapshot-backed LUN
rsh filer lun clone create /vol/dbvol/writeable.lun -o noreserve -b
/vol/dbvol/.snapshot/snap1/dblun.lun

rsh filer lun map /vol/dbvol/writeable.lun dbhost2

# mount the filesystem on the local machine
fsck -y /dev/dsk/c2t1d0s1
mount /dev/dsk/c2t1d0s1 /u01

# startup Oracle
su - oracle -c "sqlplus /nolog @startup.sql"
```

### clonelun.bat:

```
rem put oracle in hot backup mode
sqlplus /nolog @begin_backup.sql

rem take the snapshot
sdcli snap create -m prodhost -s snap1 -D 0

rem take oracle out of hot backup
sqlplus /nolog @end_backup.sql

rem create the snapshot backed lun and connect to current machine
sdcli snap mount -m devhost -r prodhost -s snap1 -k 0 -d 0

rem start the Oracle instance
sqlplus /nolog @startup.sql
```

### begin\_backup.sql:

```
connect sys/sys@source as sysdba;
alter tablespace system begin backup;
alter tablespace sys_undots begin backup;
alter tablespace users begin backup;
```

### end\_backup.sql:

```
connect sys/sys@source as sysdba;
alter tablespace system end backup;
alter tablespace sys_undots end backup;
alter tablespace users end backup;
```

### startup.sql:

```
connect / as sysdba;
startup mount;
alter database end backup;
recover automatic database;
alter database open;
```



## Conclusions

---

IBM N series with Snapshot technology has always given the database administrator a robust and extremely fast backup mechanism. Storage virtualization provided by Data ONTAP gives the system administrator and database administrator powerful tools to quickly and efficiently create cloned databases for a variety of uses.

## Caveats

---

You can find support on the Oracle website for IBM N series storage. However, this procedure has not been tested with all of the combinations of hardware and software options available. There may be significant differences in your configuration that will alter the procedures necessary to accomplish the objectives outlined in this paper.



## Trademarks and Special Notices

---

© International Business Machines 1994-2007. 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

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

Network Appliance, the Network Appliance logo, Data ONTAP, FilerView, SnapDrive, Snapshot and WAFL are trademarks or registered trademarks of Network Appliance, Inc., in the U.S. and other countries.

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

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

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

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

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.