Technical Report

Migration of Oracle Database to Data ONTAP 8.0 64-bit Aggregates

Niranjan Mohapatra and Karthikeyan Nagalingam, NetApp

November 2009 | TR-3792

EXECUTIVE SUMMARY

Migration of large Oracle Databases such as data warehouses and decision support systems, which can easily consume many terabytes, presents one of the biggest technical challenges for any organization. This document discusses various procedures for the nondisruptive upgrade and migration of Oracle® Database to Data ONTAP® 8.0 7-Mode 64-bit aggregates. This nondisruptive process of migration not only allows each node of active-active controllers to be upgraded individually to a newer version of Data ONTAP or firmware, it also provides the ability to transparently perform hardware upgrades and maintenance on the active-active controller nodes.

TABLE OF CONTENTS

1 INTRODUCTION

Typically in a database environment you may need more space in the same aggregate as your data grows. Older releases of Data ONTAP restricted the creation of 64-bit aggregates. With the release of Data ONTAP 8.0 7-Mode, you can create 64-bit aggregates to store *n*-size databases.

This document describes an end-to-end solution for the migration of Data ONTAP from Data ONTAP 7G to Data ONTAP 8.0 7-Mode and the migration of database volumes to 64-bit aggregates. of Data ONTAP 8.0 7-Mode in an Oracle Database environment.

There are several reasons to migrate your database to Data ONTAP 8.0 7-Mode:

- In a business enterprise, there are no limits for database growth. Using a 64-bit aggregate, the entire database can be stored in a single aggregate.
- Managing a single aggregate is easier than managing multiple aggregates where the database is stored.
- Sequential throughput improves performance when using 64-bit aggregates.
- There is a performance issue when using 1TB SATA drives to create a 32-bit aggregate, because it imposes a maximum size limit of 16TB. Because of fewer disk spindles, read/write performance degrades. Using Data ONTAP 8.0 7-Mode, you can create 64-bit aggregates by using a large number of disk spindles, which improves the read/write performance.
- There will not be any new release of Data ONTAP 7G. All new features will be included in Data ONTAP 8.0 7-Mode.

1.1 AUDIENCE

This guide is for NetApp customers using Oracle Database on NetApp® storage, database administrators, data center managers, sales engineers (SEs), consulting sales engineers (CSEs), professional services engineers (PSEs), professional services consultants (PSCs), contracted delivery partners (CDPs), and channel partner engineers.

This document assumes familiarity with NetApp storage solutions. It is helpful to have a working knowledge of upgrading an Oracle Database to a NetApp 64-bit aggregate in Data ONTAP 8.0 7-Mode.

1.2 SCOPE

Various migration processes described in this document have different impacts on the performance of the database. Therefore, NetApp recommends performing this migration at off-peak hours, when traffic to the database is very low.

Before performing the migration, be sure to review the following checklist:

- Review the release notes for this Data ONTAP release.
- Review the requirements for upgrading to the current release from your existing software.
- Create a secondary plan, in the unlikely event that you need to revert to the Data ONTAP release running on your system before the upgrade.
- If you have storage systems in an active-active configuration, select the appropriate upgrade method.
- If your storage system is in a SAN environment, verify that all components of your SAN configuration are compatible with the upgraded Data ONTAP release by consulting the NetApp Interoperability Matrix on the NOW™ (NetApp on the Web) site.
- If you run the SnapMirror® software, identify storage systems with destination and source volumes.
- If necessary, perform any required preliminary procedures before upgrading to the new Data ONTAP release. Required procedures might include the following:
  - Resolving upgrade issues, including performing an intermediate upgrade
  - Updating disk firmware
  - Updating disk shelf firmware
  - Upgrading storage system firmware

2 NONDISRUPTIVE UPGRADE OVERVIEW

For organizations that rely on Oracle Database for their business applications such as Oracle E-Business Suite, PeopleSoft, Siebel, and SAP® in a production environment, the biggest challenge is to nondisruptively migrate the operating system (Data ONTAP) running on the NetApp storage to a newer release or version where the database is stored or alive to achieve the new functionality provided by NetApp. It is a challenge for organizations to properly plan and implement, which puts pressure on the architecture teams and administrators to deploy the upgrade faster.

NetApp provides nondisruptive upgrades so that the process of migrating Data ONTAP software on each of the two storage cluster nodes in an active-active controller configuration can take place without affecting I/O to connected client machines. This procedure supports all the storage protocols (NFS, iSCSI, FCP) except CIFS.

The most important task here is migrating the existing database to 64-bit aggregates after upgrading to Data ONTAP 8.0 7-Mode.

Three types of migration processes are offered:

- **Using Oracle Automatic Storage Management (ASM)**: Add an ASM disk created in a 64-bit aggregate to the existing ASM normal mirrored disk group and remove the old ASM disks of the 32-bit aggregate from the existing ASM disk group one at a time after completing the rebalance. This rebalancing will have some performance impact, but it can be done online, with no downtime required. NetApp recommends performing this operation during off-peak hours.
- **Using NetApp qtree SnapMirror**: This method requires very minimal downtime. Using qtree SnapMirror, first replicate the volumes of the 32-bit aggregate on the 64-bit aggregate online. Bring down the database for some time to mount it on the 64-bit aggregate to maintain data consistency once the base level of replication is done.
- **Using NDMP**: Use NDMP to copy between the two volumes of 32-bit aggregate and 64-bit aggregate if your database has downtime. The downtime varies depending on the size of the database.

Note: Test the whole migration process before applying it in production.

2.1 WHEN NOT TO USE NONDISRUPTIVE UPGRADES

Nondisruptive upgrade methods cannot be used in all circumstances. Upgrades may be disruptive if any of the following conditions are true:

- You have storage systems that are actively serving CIFS to clients. Because CIFS is session oriented, sessions must be terminated before upgrade procedures to prevent data loss.
- You have storage systems that are actively serving FTP or NDMP clients that cannot be postponed. Because these protocols are session oriented, outstanding sessions must complete, and these services must be disabled to use nondisruptive upgrades.
- You need to update firmware for AT-FC-based or AT-FC2-based disk shelves. Client services may encounter delays accessing data when disk shelf firmware is updated to AT-FC or AT-FC2 modules. To prevent data loss, all session-oriented services must be terminated before you begin an update procedure.
- You need to update disk firmware and have RAID 4 aggregates on your system. Standard disk firmware updates automatically take disks in RAID 4 aggregates offline until the update is complete. Services and data are unavailable until they are back online.

  Note: If you upgrade RAID protection to RAID-DP®, disk firmware updates take place in the background and are nondisruptive.
- Your configuration does not meet the nondisruptive requirement for FCP or iSCSI hosts. For these conditions, standard upgrades are recommended.

2.2 PREREQUISITES FOR NONDISRUPTIVE UPGRADES

Before beginning a nondisruptive upgrade process, make sure that your systems meet the configuration and utilization requirements. Use the Upgrade Advisor tool (if available in your environment) to help determine nondisruptive upgrade requirements.

Before beginning the nondisruptive upgrade procedure, do the following:

- Verify that an active-active configuration exists, where a partner controller takes over I/O during the upgrade process.
- Remove or replace all failed disk drives before beginning the nondisruptive upgrade. This is necessary because failed disk drives prevent giveback operations and can introduce loop instability throughout the storage system.
- Verify that there are no old core files in the `/etc/core` directory.
- Verify that the systems run the latest disk and disk shelf firmware.
- Verify that you are able to open a terminal session to the console port of both controllers in an active-active configuration by using one of the following methods:
  - Direct serial connection
  - A console server
  - The system's remote LAN modules (RLM), if available
  - The system's Baseboard Management Controllers (BMC), if available

Note: Telnet, SSH, and FilerView® sessions do not work during the nondisruptive upgrade because network connections to the controllers are lost during the takeover and giveback operations.

Major nondisruptive upgrades (nondisruptive upgrades between release families) to Data ONTAP 7.3 releases are supported in Data ONTAP 7.2.2 and later Data ONTAP 7.2 releases.

- If you are running a release of the Data ONTAP 7.2 family earlier than 7.2.2, to perform a nondisruptive upgrade to a Data ONTAP 8.0 7-Mode release, first perform a nondisruptive upgrade to the latest Data ONTAP 7.3 release.
- If you are running Data ONTAP 7.0 or 7.1, to perform a nondisruptive upgrade to a Data ONTAP 7.3 release, first upgrade to the latest 7.2.*x* release and then upgrade to Data ONTAP 7.3.
- Minor nondisruptive upgrades (nondisruptive upgrades within release families) are supported on all previous Data ONTAP 7.3 releases.

Do not exceed the maximum of 300 FlexVol® volumes for major or minor nondisruptive upgrades.

On all platforms, do not exceed the maximum values for the system elements shown in Table 1.

Table 1) Maximum values for selected system elements.

Element Value	Per Storage Controller
Snapshot® copies	No more than 10 times the number of FlexVol volumes
CPU utilization	No greater than 50%
Disk utilization	No greater than 50%

3 TEST SETUP

Figure 1 shows a high-level overview of the test setup.

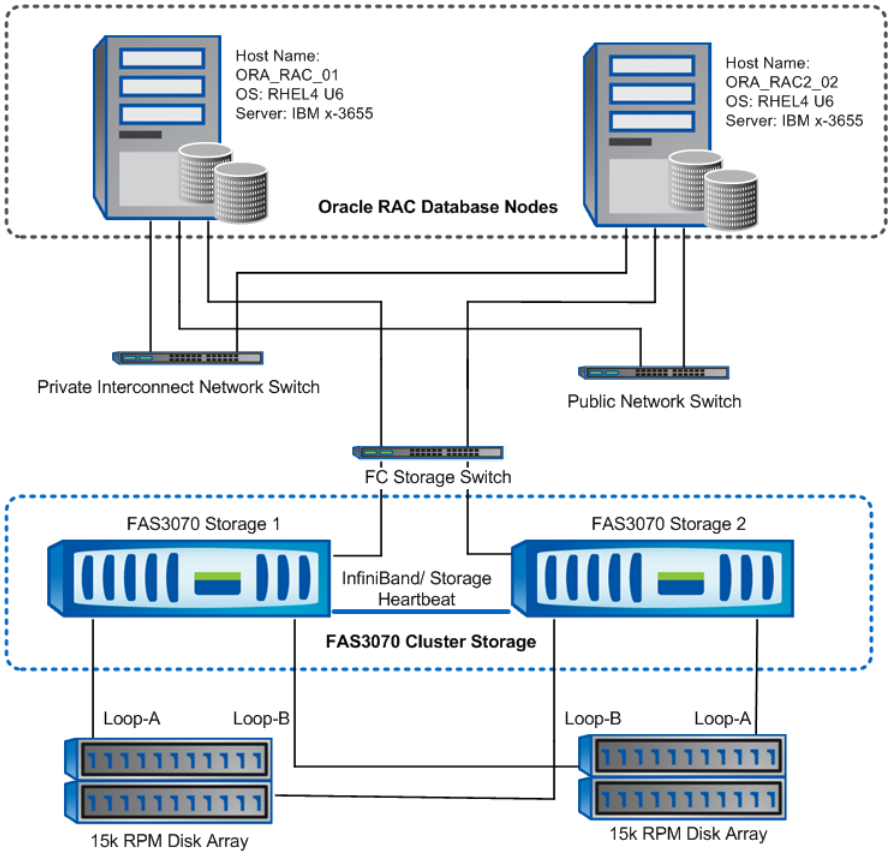

Figure 1) High-level overview of the test setup.

The test setup consists of an Oracle Real Application Cluster (RAC) with two nodes (IBM eServer xSeries with four Intel® Xeon™ CPUs and 4GB RAM) as members in each RAC. The hosts in the RAC are named ORA_RAC_01 and ORA_RAC_02. Two NetApp FAS3070 systems in an active-active cluster configuration are required for the upgrade to Data ONTAP 8.0 7-Mode.

4 INSTALLING AND CONFIGURING STORAGE AND DATABASE

4.1 LOW-LEVEL STORAGE INFRASTRUCTURE ARCHITECTURE

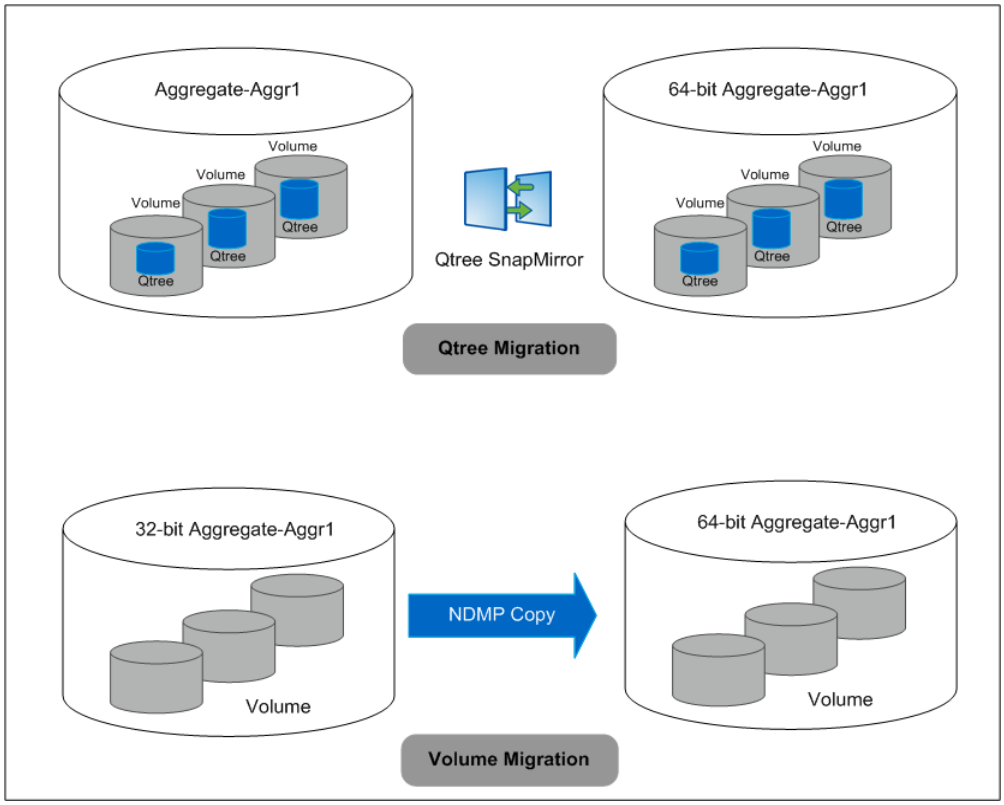

Figure 2) Low-level storage infrastructure.

4.2 NETAPP STORAGE CONFIGURATION

The following tables show the storage layout (aggregate or volumes) for a 1TB Oracle database.

Table 1) Aggregate layout.

Controller	Aggregate Name	Option	# of Disks/Type/Size	Purpose
FAS-3070_1	ORADATA_32	RAID-DP Aggr	10 no.s / 15K RPM FC Drives / 172GB	Oracle data files and control files copy
FAS-3070_1	REDO_FRA_32	RAID-DP Aggr	10 no.s / 15K RPM FC Drives / 172GB	Redo logs, archive logs, FRA, copy of control files
FAS-3070_1	ORADATA_64	RAID-DP Aggr	10 no.s / 15K RPM FC Drives / 172GB	Oracle data files and control files copy
FAS-3070_1	REDO_FRA_64	RAID-DP Aggr	10 no.s / 15K RPM FC Drives / 172GB	Redo logs, archive logs, FRA, copy of control files
FAS-3070_2	ORADATA_32	RAID-DP Aggr	10 no.s / 15K RPM FC Drives / 172GB	Oracle data files and control files copy
FAS-3070_2	REDO_FRA_32	RAID-DP Aggr	10 no.s / 15K RPM FC Drives / 172GB	Redo logs, archive logs, FRA, copy of control files

FAS-3070_2	ORADATA_64	RAID-DP Aggr	10 no.s / 15K RPM FC Drives / 172GB	Oracle data files and control files copy
FAS-3070_2	REDO_FRA_64	RAID-DP Aggr	10 no.s / 15K RPM FC Drives / 172GB	Redo logs, archive logs, FRA, copy of control files

Table 2) Volume/LUN layout.

Controller	Volume/LUN name	Aggregate Name	Size	Purpose
FAS-3070_1	Oradata_32 /Data1_32	ORADATA_32	1200GB	Oracle data files and control files copy
FAS-3070_1	Redo_FRA_32 / FRA1_32	REDO_FRA_32	1200GB	Redo logs, archive logs, FRA, copy of control files
FAS-3070_1	Oradata_64 / Data1_64	ORADATA_64	1200GB	Oracle data files and control files copy
FAS-3070_1	Redo_FRA_64 / FRA1_64	REDO_FRA_64	1200GB	Redo logs, archive logs, FRA, copy of control files
FAS-3070_2	Oradata_32 /Data2_32	ORADATA_32	1200GB	Oracle data files and control files copy
FAS-3070_2	Redo_FRA_32 / FRA2_32	REDO_FRA_32	1200GB	Redo logs, archive logs, FRA, copy of control files
FAS-3070_2	Oradata_64 / Data2_64	ORADATA_64	1200GB	Oracle data files and control files copy
FAS-3070_2	Redo_FRA_64 / FRA2_64	REDO_FRA_64	1200GB	Redo logs, archive logs, FRA, copy of control files

4.3 DATABASE SERVER INSTALLATION AND CONFIGURATION

HOST CONFIGURATION

Oracle Real Application Clusters (RAC) were set up with two nodes (IBM eServer xSeries with four Intel Xeon CPUs and 4GB RAM) as members in each RAC. The hosts in the RAC are named ORA_RAC_01 and ORA_RAC_02.

1. Install the following after the host installation and configuration:
   - Oracle Cluster Ready Service (CRS) 10*g* R2 (10.2.0.4)
   - Oracle Database 10*g* R2 (10.2.0.4)
2. Configure Listener and ASM instances.

Note: For details about Oracle RAC installation and configuration, refer http://media.netapp.com/documents/tr-3594.pdf

DATABASE STORAGE LAYOUT WITH ASM

Table 3 describes the Automatic Storage Management (ASM) layout for the database creation.

Create the database in the ASM disk groups.

Table 3) ASM layout.

ASM Disk Group Name	ASM Disk Name	ASM Mirror/ Redundancy Level	Purpose
DATADG	DATA1 , DATA2	Normal	Store all data files, spfile and control files
FRADG	FRA1, FRA2	Normal	Store Redo log files, Archive log files, FRA, copy of control files.

5 MIGRATING DATA ONTAP 7.2.X TO DATA ONTAP 8.0 7-MODE

There is no supported upgrade from Data ONTAP 7.2.x to Data ONTAP 8.0 7-Mode. To perform the upgrade, first upgrade from Data ONTAP 7.2.x to 7.3.x, and then from 7.3.x to Data ONTAP 8.0 7-Mode.

5.1 UPGRADE DATA ONTAP FROM 7.2.X TO 7.3.X

This upgrade process can be done online without any downtime by using NetApp NDU (nondisruptive upgrade) and cluster storage. For full information, see the *Data ONTAP Upgrade Guide from 7.2.x to 7.3.x*.

Note: NetApp recommends backing up your database before performing the upgrade.

5.2 UPGRADE DATA ONTAP FROM 7.3.X TO DATA ONTAP 8.0 7-MODE

This upgrade process can be done online without any downtime by using NetApp NDU (nondisruptive upgrade) and cluster storage. For full information, see the *Data ONTAP Upgrade Guide from 7.2.x to 8.0 7-Mode*.

Note: NetApp recommends backing up your database before performing the upgrade.

5.3 CREATE 64-BIT AGGREGATES AND FLEXIBLE VOLUMES

After you upgrade Data ONTAP from a previous release to Data ONTAP 8.0 7-Mode, create the 64-bit aggregates and also create volumes or LUNs as per the existing (32-bit) setups.

All the volumes or LUNs can be created online without disrupting the production systems.

6 MIGRATING THE DATABASE TO THE 64-BIT AGGREGATE

This section describes three migration processes. Choose the process that is most suitable for your environment:

- Migrate the database online.
- Migrate the database with minimal downtime.
- Migrate the database offline.

6.1 MIGRATE THE DATABASE ONLINE

This migration process requires that your database runs on an ASM file system with either normal redundancy or high redundancy. NetApp recommends backing up your database before migration.

1. Create ASM disks of the same or larger size than the existing production ASM disks in the 64-bit aggregate's volumes.
2. Add the newly created ASM disk to the existing disk group and remove the old ASM disks one at a time after the rebalancing operation is done.

NetApp recommends performing this migration in off-peak hours because the rebalancing job consumes resources.

The migration process can be done online without disturbing I/O to the database.

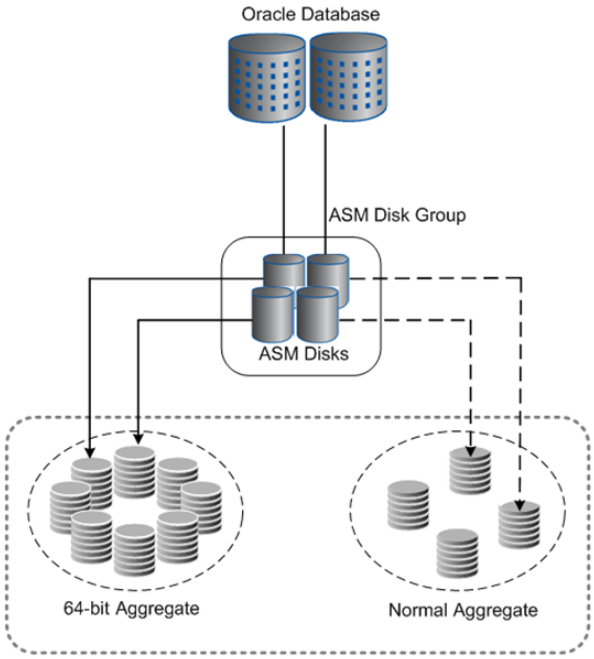

Figure 3) Migrate the database online.

6.2 MIGRATE THE DATABASE WITH MINIMAL DOWNTIME

In this migration process, the database requires minimal downtime during migration from the normal aggregate to the 64-bit aggregate. The database need not be in ASM and can be any other file system. This migration process uses the qtree SnapMirror functionality of NetApp storage.

To migrate your database, follow these steps:

1. Create 64-bit aggregates, volumes, qtrees, LUNs, and so on, as required.
2. Create the SnapMirror relationship between 32-bit aggregate volumes and the 64-bit aggregate volume based on the Snapshot copy created for the 32-bit aggregate volumes by using qtree SnapMirror.

3. When the baseline data transfer has completed successfully (which can be done online without disturbing I/O to the database), bring down the database instance to sync the final delta changes in the database.
4. Synchronize the 32-bit aggregate volume with the 64-bit aggregate volume.
5. Temporarily rename the 32-bit aggregate volume to some other name, because you need to rename the 64-bit aggregate volume with the original name of 32-bit aggregate volume.
6. Rename the 64-bit aggregate volume with the original name of 32-bit aggregate volume.
7. Start the database for production use.

This migration process requires minimal downtime to sync the last delta changes in the normal aggregate volume to the 64-bit aggregate volume.

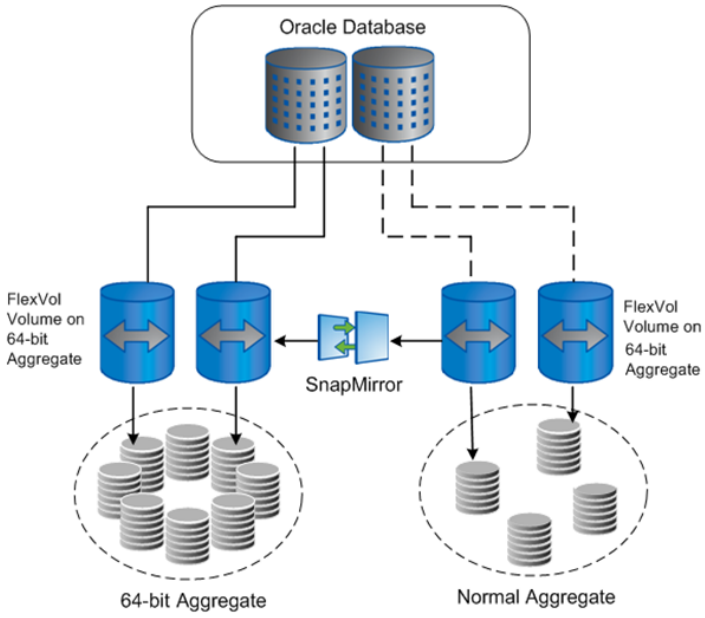

Figure 4) Migrate the database with minimal downtime.

6.3 MIGRATE THE DATABASE OFFLINE

This migration process requires full downtime of the database. To perform offline migration of your database from normal aggregates to a 64-bit aggregate, follow these steps:

1. Create the volume, qtrees, and LUNS (as required) in 64-bit aggregates.
2. Bring down the database stored in the normal aggregate.
3. Use the `NDMP Copy storage` command to copy the files from the normal aggregate volumes to the 64-bit aggregate volumes.
4. Temporarily rename the normal aggregate volumes to some other name, because you need to rename the 64-bit aggregate volume to the original name of normal aggregate volume.
5. Rename the 64-bit aggregate volume with the original name of 32-bit normal aggregate volume.
6. Start the database and continue to use it for production.

Note: Validate this solution in a test environment before applying the changes in a production environment.

6.4 POSTMIGRATION VALIDATION

After completing the migration, you may want to compare the number and sizes of objects, including data and definitions (such as item types and attributes), between the 32-bit and 64-bit aggregates.

You can use your application to connect to the database and run the test script to validate data. Oracle 11*g* has a feature called Real Application Testing, which can be used to capture the database load from the 32-bit aggregate and replay it after migrating the database to the 64-bit aggregate to simulate a realistic load to the database. After the replay you can generate the report to validate the data.

7 SUMMARY

Using Oracle ASM, NetApp qtree SnapMirror, and NDMP copy, databases can be migrated from a 32-bit aggregate to 64-bit aggregate on NetApp storage. Using Oracle ASM, the database can be migrated online without affecting production I/O. Using NetApp QSM minimizes the downtime required to sync last-minute changes occurring in the 32-bit aggregate. Using this process, upgrading Data ONTAP from 7G to Data ONTAP 8.0 7-Mode and baseline migration from 32-bit aggregate to 64-bit aggregate can be done online. When using NDMP copy, the database must be down completely before migrating from a 32-bit aggregate to a 64-bit aggregate. The Data ONTAP upgrade can be done nondisruptively.

Please send any errors, omissions, differences, new discoveries, and comments about this paper to niranjan@netapp.com or karthikeyan.nagalingam@netapp.com .

APPENDIX A: MIGRATION EXAMPLES

EXAMPLE 1: MIGRATE USING ORACLE ASM

This example shows the creation of a disk group called `oradata`, which has two ASM disks (disk1 and disk2) existing in a 32-bit aggregate used to store the database. There is ongoing I/O to the database. We created two more ASM disks in the 64-bit aggregate, called `bigdata1` and `bigdata2`.

To migrate the database from the 32-bit aggregate to the 64-bit aggregate:

32-bit disks	64-bit disks
DATA1	BIGDATA1
DATA2	BIGDATA2

```
SQL> alter diskgroup ORADATA add disk 'ORCL:BIGDATA1' rebalance power 0;

Diskgroup altered.

SQL> alter diskgroup ORADATA add disk 'ORCL:BIGDATA2' rebalance power 0;

Diskgroup altered.

SQL> alter diskgroup ORADATA drop disk 'DATA1' rebalance power 0;

Diskgroup altered.

SQL> alter diskgroup ORADATA drop disk 'DATA2' rebalance power 11;

Diskgroup altered.

SQL> select name, state from v$asm_disk_stat where name like '%DATA%';

NAME                            STATE

------------------------------------

ORADATA1                        DROPPING

ORADATA2                        DROPPING

BIGDATA1                        NORMAL

BIGDATA2                        NORMAL

select name, state from v$asm_disk;

SQL> select * from v$asm_operation;

GROUP_NUMBER OPERA STAT      POWER     ACTUAL      SOFAR   EST_WORK
EST_RATE EST_MINUTES

------------ ----- ---- ---------- ---------- ---------- ---------- --------
-- -----------

1 REBAL RUN           1          1       2518    1517939       1015
1493

SQL> select * from v$asm_operation;

no rows selected

SQL> select name, state from v$asm_disk_stat where name like '%DATA%';

NAME                    STATE

------------------------------ --------

BIGDATA1                NORMAL

BIGDATA2                NORMAL
```

EXAMPLE 2: MIGRATE USING NETAPP QTREE SNAPMIRROR

In this example, we created a volume called `ora10data` in the 32-bit aggregate where the database was created. Also, a Snapshot copy was taken for that volume.

1. Check the Snapshot copy created in `ora10data`:

```
btcppe181*> snap list ora10data

Volume ora10data

working...

%/used       %/total      date          name

----------  ----------  ------------  --------

7% ( 7%)    0% ( 0%)     Aug 07 10:53
smo_b1_newdb1_f_h_2_8ac9c5ef22e50a790122e50a808f0001_0_backup1

  7% ( 0%)    0% ( 0%)   Aug 07 10:53
smo_b1_newdb1_f_h_1_8ac9c5ef22e50a790122e50a808f0001_0_backup1

btcppe181*>
```

2. Migrate the `ora10data` volume by using the Snapshot copy named `smo_b1_newdb1_f_h_2_8ac9c5ef22e50a790122e50a808f0001_0_backup1` to the `ora10data64` volume.

```
btcppe181*> snapmirror initialize -s
smo_b1_newdb1_f_h_2_8ac9c5ef22e50a790122e50a808f0001_0_backup1  -S
/vol/ora10data/ora10dataq /vol/ora10data64/ora10dataq

Transfer started.

Monitor progress with 'snapmirror status' or the snapmirror log.

btcppe181*>
```

3. Migrate the non-qtree `ora10data` volume from the `smo_b1_newdb1_f_h_2_8ac9c5ef22e50a790122e50a808f0001_0_backup1` Snapshot copy.

```
btcppe181*> snapmirror initialize -s
smo_b1_newdb1_f_h_2_8ac9c5ef22e50a790122e50a808f0001_0_backup1 -S
/vol/ora10data/- /vol/ora10data64/root

Transfer started.

Monitor progress with 'snapmirror status' or the snapmirror log.

btcppe181*>
```

4. Check the SnapMirror status:

```
btcppe181*> snapmirror status

Snapmirror is on.

Source    Destination           State          Lag        Status

btcppe181:/vol/ora10archive/ora10archiveq
btcppe181:/vol/ora10archive64/ora10archiveq  Snapmirrored   21:25:56
Idle

btcppe181:/vol/ora10data/ora10dataq
btcppe181:/vol/ora10data64/ora10dataq        Uninitialized  -
Transferring  (2592 MB done)

btcppe181:/vol/ora10data/-
btcppe181:/vol/ora10data64/root              Snapmirrored   21:26:10
Transferring

btcppe181:/vol/ora10archive/-
btcppe181:/vol/ora10archive64/root           Source         21:25:56
Idle
```

```
btcppe181*>

btcppe181*> snapmirror status

Snapmirror is on.

Source                                          Destination
State          Lag        Status

btcppe181:/vol/ora10archive/ora10archiveq
btcppe181:/vol/ora10archive64/ora10archiveq  Snapmirrored   21:25:56
Idle

btcppe181:/vol/ora10data/ora10dataq
btcppe181:/vol/ora10data64/ora10dataq        Uninitialized  -
Transferring  (2592 MB done)

btcppe181:/vol/ora10data/-
btcppe181:/vol/ora10data64/root              Snapmirrored   21:26:10
Transferring

btcppe181:/vol/ora10archive/-
btcppe181:/vol/ora10archive64/root           Source         21:25:56
Idle

btcppe181*>
```

5. Migrate the SMO (Backup) specific files from the `ora10data` qtree to `ora10data64`:

```
btcppe181*> snapmirror quiesce  btcppe181:/vol/ora10data64/root

snapmirror quiesce: in progress

This can be a long-running operation. Use Control - C (^C) to interrupt.

snapmirror quiesce: /vol/ora10data64/root :  Successfully quiesced

btcppe181*> snapmirror break  btcppe181:/vol/ora10data64/root

snapmirror break: Destination /vol/ora10data64/root is now writable.

btcppe181*>


btcppe181*> exportfs -p "rw,anon=0" /vol/ora10data64

btcppe181*> exportfs -r

btcppe181*> exportfs –a


[root@oraserv1 ~]# mount -o rw btcppe181:/vol/ora10data64 /64

[root@oraserv1 ~]# cd /64

[root@oraserv1 64]# ls -ltra

total 24

drwxrwxrwx  2 root root 4096 Jul 27 13:45 root

drwxr-xr-x 41 root root 4096 Jul 27 14:13 ..

drwxrwxrwx  2 root root 4096 Jul 28 11:10 ora10dataq

drwxr-xr-x  5 root root 4096 Jul 28  2009 .

drwxrwxrwx  4 root root 4096 Jul 28  2009 .snapshot

[root@oraserv1 64]# cd root

[root@oraserv1 root]# ls -ltra

total 12

-rw-r--r-- 1 root root   28 Jul 27 13:45 .snapdrive_lock
```

```
drwxrwxrwx 2 root root 4096 Jul 27 13:45 .
-rw-r--r-- 1 root root 2081 Jul 27 13:45 .stoc.xml
drwxr-xr-x 5 root root 4096 Jul 28  2009 ..
[root@oraserv1 root]# cp .stoc.xml ..
[root@oraserv1 root]# cd ..
[root@oraserv1 64]# ls -ltra
total 28
drwxrwxrwx  2 root root 4096 Jul 27 13:45 root
drwxr-xr-x 41 root root 4096 Jul 27 14:13 ..
drwxrwxrwx  2 root root 4096 Jul 28 11:10 ora10dataq
drwxrwxrwx  4 root root 4096 Jul 28  2009 .snapshot
drwxr-xr-x  5 root root 4096 Jul 28  2009 .
-rw-r--r--  1 root root 2081 Jul 28  2009 .stoc.xml
[root@oraserv1 64]# rm -fr root
[root@oraserv1 64]# ls -ltra
total 24
drwxr-xr-x 41 root root 4096 Jul 27 14:13 ..
drwxrwxrwx  2 root root 4096 Jul 28 11:10 ora10dataq
drwxrwxrwx  4 root root 4096 Jul 28  2009 .snapshot
-rw-r--r--  1 root root 2081 Jul 28  2009 .stoc.xml
drwxr-xr-x  4 root root 4096 Jul 28  2009 .
[root@oraserv1 64]# cd
[root@oraserv1 ~]# umount /64
btcppe181*> snapmirror status
Snapmirror is on.
Source                                      Destination
State          Lag        Status
btcppe181:/vol/ora10archive/ora10archiveq
btcppe181:/vol/ora10archive64/ora10archiveq  Snapmirrored   29:07:12
Idle
btcppe181:/vol/ora10data/ora10dataq
btcppe181:/vol/ora10data64/ora10dataq        Snapmirrored   29:07:26
Idle
btcppe181:/vol/ora10archive/-
btcppe181:/vol/ora10archive64/root           Source         29:07:12
Idle
btcppe181:/vol/ora10data/-
btcppe181:/vol/ora10data64/root              Source         29:07:26
Idle
btcppe181*>
```

6. Create the `smo_b1_newdb1_f_h_2_8ac9c5ef22e50a790122e50a808f0001_0_backup1` Snapshot copy in `ora10data64` for the `ora10data` volume:

```
btcppe181*> snap create ora10data64
smo_b1_newdb1_f_h_2_8ac9c5ef22e50a790122e50a808f0001_0_backup1
```

7. Check the Snapshot copy just created:

```
btcppe181*> snap list ora10data64

Volume ora10data64

working...

%/used       %/total  date          name

----------  ----------  ------------  --------

0% ( 0%)    0% ( 0%)  Aug 09 08:23
smo_b1_newdb1_f_h_2_8ac9c5ef22e50a790122e50a808f0001_0_backup1

4% ( 4%)    0% ( 0%)  Aug 09 08:16
btcppe181(0118056873)_ora10data64_root-dst.2

8% ( 4%)    0% ( 0%)  Aug 09 08:15
btcppe181(0118056873)_ora10data64_ora10dataq-dst.2 (busy,snapmirror)

btcppe181*>

btcppe181*> snap list ora10data

Volume ora10data

working...


%/used       %/total  date          name

----------  ----------  ------------  --------

7% ( 7%)    0% ( 0%)  Aug 07 10:53
smo_b1_newdb1_f_h_2_8ac9c5ef22e50a790122e50a808f0001_0_backup1
(snapmirror)

  7% ( 0%)    0% ( 0%)  Aug 07 10:53
smo_b1_newdb1_f_h_1_8ac9c5ef22e50a790122e50a808f0001_0_backup1

btcppe181*>
```

8. Perform steps 1-7 for the `ora10archive`, `ora10fra`, and `ora10log` volumes:
9. Check the LUNs in the 64-bit aggregate volume:

```
btcppe181*> lun show

/vol/demodata2/demolun1       15g (16106127360)   (r/w, online, mapped)

/vol/ora10archive/ora10archiveq/lun1    100g (107374182400)  (r/w,
online, mapped)

/vol/ora10archive64/ora10archiveq/lun1    100g (107374182400)  (r/o,
online)

/vol/ora10data/ora10dataq/lun1    1.0t (1106377703424) (r/w, online,
mapped)

/vol/ora10data64/ora10dataq/lun1    1.0t (1106377703424) (r/o, online)

/vol/ora10fra/ora10fraq/lun1     10g (10737418240)   (r/w, online,
mapped)

/vol/ora10fra64/ora10fraq/lun1      8g (8589934592)    (r/w, online)

/vol/ora10log/ora10logq/lun1     15g (16106127360)   (r/w, online,
mapped)

/vol/ora10log64/ora10logq/lun1   17.0g (18254659584)   (r/w, online)

btcppe181*>
```

10. Check the Snapshot copy created in the `ora10data` and `ora10data64` volumes:

```
btcppe181*> snap list ora10data
```

```
Volume ora10data

working...

%/used       %/total  date          name

----------  ----------  ------------  --------

7% ( 7%)    0% ( 0%)  Aug 07 10:53
smo_b1_newdb1_f_h_2_8ac9c5ef22e50a790122e50a808f0001_0_backup1

7% ( 0%)    0% ( 0%)  Aug 07 10:53
smo_b1_newdb1_f_h_1_8ac9c5ef22e50a790122e50a808f0001_0_backup1
(snapmirror)

btcppe181*> snap list ora10data64

Volume ora10data64

working...

%/used       %/total  date          name

----------  ----------  ------------  --------

0% ( 0%)    0% ( 0%)  Aug 09 08:32
smo_b1_newdb1_f_h_1_8ac9c5ef22e50a790122e50a808f0001_0_backup1

4% ( 4%)    0% ( 0%)  Aug 09 08:28
btcppe181(0118056873)_ora10data64_root-dst.2

8% ( 4%)    0% ( 0%)  Aug 09 08:25
btcppe181(0118056873)_ora10data64_ora10dataq-dst.4 (busy,snapmirror)

10% ( 3%)    0% ( 0%)  Aug 09 08:23
smo_b1_newdb1_f_h_2_8ac9c5ef22e50a790122e50a808f0001_0_backup1

btcppe181*>
```

11. Update the `ora10data` volume with SnapMirror Semi-Sync:

```
btcppe181*> snapmirror update -S /vol/ora10data/ora10dataq
/vol/ora10data64/ora10dataq

Transfer started.

Monitor progress with 'snapmirror status' or the snapmirror log.

btcppe181*>

btcppe181*> snapmirror status

Snapmirror is on.

Source                                    Destination
State          Lag        Status

btcppe181:/vol/src/qt src                 btcppe181:/vol/dst/qt dst
Snapmirrored   238:11:58  Idle

btcppe181:/vol/ora10archive/ora10archiveq
btcppe181:/vol/ora10archive64/ora10archiveq  Snapmirrored   00:03:32
Idle

btcppe181:/vol/ora10data/ora10dataq
btcppe181:/vol/ora10data64/ora10dataq        Snapmirrored   00:00:35
Idle

btcppe181:/vol/ora10archive/-
btcppe181:/vol/ora10archive64/root           Source         45:46:58
Idle

btcppe181:/vol/ora10data/-
btcppe181:/vol/ora10data64/root              Source         45:47:36
Idle
```

12. Update the `ora10archive64`, `ora10log64`, `ora10fra64`, and `ora10data64` volumes using semi-sync and update the `/etc/snapmirror.conf` file:

```
btcppe181*> wrfile /etc/snapmirror.conf

btcppe181:/vol/ora10archive/ora10archiveq
btcppe181:/vol/ora10archive64/ora10archiveq restart=never semi-sync

btcppe181:/vol/ora10data/ora10dataq
btcppe181:/vol/ora10data64/ora10dataq restart=never semi-sync

btcppe181:/vol/ora10fra/ora10fraq  btcppe181:/vol/ora10fra64/ora10fraq
restart=never semi-sync

btcppe181:/vol/ora10log/ora10logq  btcppe181:/vol/ora10log64/ora10logq
restart=never semi-sync

read: error reading standard input: Interrupted system call

btcppe181*>

btcppe181*> snapmirror status

Snapmirror is on.

Source                                      Destination
State          Lag        Status

btcppe181:/vol/src/qt src                   btcppe181:/vol/dst/qt dst
Snapmirrored   239:56:24  Idle

btcppe181:/vol/ora10archive/ora10archiveq
btcppe181:/vol/ora10archive64/ora10archiveq  Snapmirrored   00:00:16
Idle

btcppe181:/vol/ora10data/ora10dataq
btcppe181:/vol/ora10data64/ora10dataq        Snapmirrored   00:00:16
Idle

btcppe181:/vol/ora10fra/ora10fraq
btcppe181:/vol/ora10fra64/ora10fraq          Snapmirrored   00:00:16
Idle

btcppe181:/vol/ora10log/ora10logq
btcppe181:/vol/ora10log64/ora10logq          Snapmirrored   00:00:16
Idle

btcppe181:/vol/ora10archive/-
btcppe181:/vol/ora10archive64/root           Source         47:31:24
Idle

btcppe181:/vol/ora10data/-
btcppe181:/vol/ora10data64/root              Source         47:32:02
Idle

btcppe181*>
```

13. Downtime steps:

a. Stop the CRS:

```
[root@node2 ~]# /ora10home/app/product/10.2.0/crs_1/bin/crsctl stop crs

Stopping resources. This could take several minutes.

Resource or relatives are currently involved in another operation.
Retrying stop resources.

Successfully stopped CRS resources.

Stopping CSSD.

Shutting down CSS daemon.

Shutdown request successfully issued.
```

b. Stop the ASM:

```
[root@node2 ~]# /etc/init.d/oracleasm stop
Dropping Oracle ASMLib disks:                              [  OK  ]
Shutting down the Oracle ASMLib driver:                    [  OK  ]
[root@node2 ~]#
 [root@node2 ~]#
ls -ltr /dev/mapper/* | grep `/etc/init.d/oracleasm querydisk  FRA2 |
awk -F ',' '{ print $2 }' | awk -F ']' '{print $1}'`
```

c. Stop the multipath:

```
[root@node2 ~]# multipath –F
```

d. Unload the Qlogic driver:

```
[root@node2 ~]# modprobe -vr qla2xxx
remove /sbin/modprobe -r --first-time --ignore-remove qla2xxx && {
/sbin/modprobe -r --ignore-remove qla2xxx_conf; }
rmmod /lib/modules/2.6.9-67.ELsmp/kernel/drivers/scsi/qla2xxx/qla2xxx.ko
rmmod /lib/modules/2.6.9-
67.ELsmp/kernel/drivers/scsi/qla2xxx/qla2xxx_conf.ko
[root@node2 ~]#
```

e. Wait until the `snapmirrored` status appears on both storage systems.

f. Perform steps a through f for other Oracle cluster members.

14. Break the SnapMirror relationship in `btcppe181` for `ora10data`:

```
btcppe181*> snapmirror release /vol/ora10data/-
btcppe181:/vol/ora10data64/root
btcppe181*>
btcppe181*> snapmirror status
Snapmirror is on.
Source                                        Destination
State          Lag        Status
btcppe181:/vol/ora10data/ora10dataq
btcppe181:/vol/ora10data64/ora10dataq        Snapmirrored   00:00:06
Idle
btcppe181*>
btcppe181*> snapmirror release /vol/ora10data/-
btcppe181:/vol/ora10data64/root
btcppe181*> snapmirror quiesce  btcppe181:/vol/ora10data64/ora10dataq
snapmirror quiesce: in progress
This can be a long-running operation. Use Control - C (^C) to interrupt.
snapmirror quiesce: /vol/ora10data64/ora10dataq :  Successfully quiesced
btcppe181*> snapmirror Sun Aug  9 13:17:37 GMT [btcppe181:
wafl.scan.ownblocks.done:info]: Completed block ownership calculation on
volume ora10data64.
btcppe181*> snapmirror break   btcppe181:/vol/ora10data64/ora10dataq
snapmirror break: Destination /vol/ora10data64/ora10dataq is now
writable.
btcppe181*>
```

```
btcppe181*> snapmirror status

Snapmirror is on.

Source   Destination                      State          Lag        Status

btcppe181:/vol/ora10data/ora10dataq
btcppe181:/vol/ora10data64/ora10dataq      Broken-off     00:01:19   Idle

btcppe181*>

btcppe181*>
```

For NFS

1. Rename the 32-bit volume to a temporary name (for example, `ora10data` to `ora10data_old`) and rename the 64-bit volume to the 32-bit volume name (for example, `ora10data64` to `ora10data`) in both the btcppe181 and btcppe182 storage systems.

```
btcppe181*> vol rename ora10data ora10data_old

Sun AugRenaming volume ora10data (fsid 5847796) to ora10data_old: start
time 285602757

  9 15:43:51 GMT [btcppe181: wafl.vvol.renamed:info]: Volume 'ora10data'
renamed to 'ora10data_old'.

'ora10data' renamed to 'ora10data_old'

btcppe181*> vol rename ora10data64 ora10data

Sun AugRenaming volume ora10data64 (fsid c605138) to ora10data: start
time 285610518

  9 15:43:58 GMT [btcppe181: wafl.vvol.renamed:info]: Volume
'ora10data64' renamed to 'ora10data'.

'ora10data64' renamed to 'ora10data'

btcppe181*>
```

2. Restart the Oracle Database.

FOR FCP AND ISCSI LUNS

1. Add the 64-bit disk LUNs to an FCP initiator group such as the btcppe181 and btcppe182 storage systems.

```
btcppe181*> lun show

/vol/demodata2/demolun1       15g (16106127360)   (r/w, online, mapped)

/vol/ora10archive/ora10archiveq/lun1    100g (107374182400) (r/w,
online,mapped)

/vol/ora10archive64/ora10archiveq/lun1    100g (107374182400)  (r/w,
online)

/vol/ora10data/ora10dataq/lun1    1.0t (1106377703424) (r/w, online,
mapped)

/vol/ora10data64/ora10dataq/lun1    1.0t (1106377703424) (r/w, online)

/vol/ora10fra/ora10fraq/lun1     10g (10737418240)   (r/w, online,
mapped)

/vol/ora10fra64/ora10fraq/lun1     10g (10737418240)   (r/w, online)

/vol/ora10log/ora10logq/lun1     15g (16106127360)   (r/w, online,
mapped)

/vol/ora10log64/ora10logq/lun1     15g (16106127360)   (r/w, online)

btcppe181*>

btcppe181*> lun show -m
```

```
LUN path                                Mapped to          LUN ID  Protocol
----------------------------------------------------------------------------
/vol/demodata2/demolun1                 crsgroup       1234     FCP
/vol/ora10archive/ora10archiveq/lun1    crsgroup       2        FCP
/vol/ora10archive64/ora10archiveq/lun1  crsgroup       642      FCP
/vol/ora10data/ora10dataq/lun1          crsgroup       0        FCP
/vol/ora10data64/ora10dataq/lun1        crsgroup       640      FCP
/vol/ora10fra/ora10fraq/lun1            crsgroup       4        FCP
/vol/ora10fra64/ora10fraq/lun1          crsgroup       644      FCP
/vol/ora10log/ora10logq/lun1            crsgroup       6        FCP
/vol/ora10log64/ora10logq/lun1          crsgroup       646      FCP
btcppe181*>
```

2. Add and drop the LUNs using the procedure described in “Migrate Using Oracle ASM,” earlier in this appendix.

EXAMPLE 3: MIGRATE USING NETAPP NDMPCOPY

You can perform volume-level migration from 32-bit to 64-bit volumes by using NDMP:

1. Migrate the 32-bit volume data to the 64-bit volume:

```
btcppe10> ndmpcopy -d /vol/vol32 /vol/vol64
Ndmpcopy: Starting copy [ 2 ] ...
Ndmpcopy: btcppe10: Notify: Connection established
Ndmpcopy: btcppe10: Notify: Connection established
Ndmpcopy: btcppe10: Connect: Authentication successful
Ndmpcopy: btcppe10: Connect: Authentication successful
Ndmpcopy: btcppe10: Log: DUMP: creating
"/vol/vol32/../snapshot_for_backup.3" snapshot.
Ndmpcopy: btcppe10: Log: DUMP: Using Full Volume Dump
Ndmpcopy: btcppe10: Log: DUMP: Date of this level 0 dump: Mon May 25
05:46:24 2009.
Ndmpcopy: btcppe10: Log: DUMP: Date of last level 0 dump: the epoch.
Ndmpcopy: btcppe10: Log: DUMP: Dumping /vol/vol32 to NDMP connection
Ndmpcopy: btcppe10: Log: DUMP: mapping (Pass I)[regular files]
Ndmpcopy: btcppe10: Log: DUMP: mapping (Pass II)[directories]
Ndmpcopy: btcppe10: Log: DUMP: estimated 1618148 KB.
Ndmpcopy: btcppe10: Log: DUMP: dumping (Pass III) [directories]
Ndmpcopy: btcppe10: Log: DUMP: dumping (Pass IV) [regular files]
Ndmpcopy: btcppe10: Log: RESTORE: Mon May 25 05:46:27 2009: Begin level 0
restore
Ndmpcopy: btcppe10: Log: RESTORE: Mon May 25 05:46:27 2009: Reading
directories from the backup
Ndmpcopy: btcppe10: Log: RESTORE: Mon May 25 05:46:27 2009: Creating
files and directories.
```

```
Ndmpcopy: btcppe10: Log: RESTORE: Mon May 25 05:46:27 2009: Writing data
to files.

Ndmpcopy: btcppe10: Log: DUMP: dumping (Pass V) [ACLs]

Ndmpcopy: btcppe10: Log: DUMP: 1618171 KB

Ndmpcopy: btcppe10: Log: DUMP: DUMP IS DONE

Ndmpcopy: btcppe10: Log: RESTORE: Mon May 25 05:46:52 2009: Restoring NT
ACLs.

Ndmpcopy: btcppe10: Log: RESTORE: RESTORE IS DONE

Ndmpcopy: btcppe10: Log: RESTORE: The destination path is /vol/vol64/

Ndmpcopy: btcppe10: Log: DUMP: Deleting
"/vol/vol32/../snapshot_for_backup.3" snapshot.

Ndmpcopy: btcppe10: Notify: restore successful

Ndmpcopy: btcppe10: Notify: dump successful

Ndmpcopy: Transfer successful [ 37 seconds ]

Ndmpcopy: Done
```

2. For NFS protocol, perform the volume rename procedure.
3. For SAN, perform the add/drop disk procedure in ASM.

APPENDIX B: REFERENCES

#	Topic	Reference
1	RAID-DP	http://media.netapp.com/documents/wp_3298.pdf http://www.usenix.org/publications/library/proceedings/fast04/tech/corbett/corbett.pdf
2	SATA	http://media.netapp.com/documents/rp-0046.pdf http://media.netapp.com/documents/tr-3574.pdf
3	Thin Provisioning and FlexVol	http://media.netapp.com/documents/tr-3563.pdf http://media.netapp.com/documents/FlexVols.pdf
4	Database Layout	http://media.netapp.com/documents/tr-3411.pdf
5	Snapshot	http://media.netapp.com/documents/tr-3001.pdf http://media.netapp.com/documents/ar1038.pdf
6	FlexClone®	http://media.netapp.com/documents/tr-3347.pdf http://media.netapp.com/documents/tr-3348.pdf http://media.netapp.com/documents/tr-3373.pdf http://media.netapp.com/documents/tr-3373.pdf http://media.netapp.com/documents/tr-3646.pdf
7	Thin Replication (SnapMirror)	http://media.netapp.com/documents/tr-3446.pdf http://media.netapp.com/documents/tr-3326.pdf http://www.netapp.com/us/library/research-papers/rp-0007.html