



Technical Report

Cloning Oracle E-Business Suite on NetApp Using NetApp and Oracle Technologies

Padmanabhan Sadagopan, Habib Rangoonwala, Sunil Kumar, NetApp
April 2010 | TR-3840

TABLE OF CONTENTS

1	INTRODUCTION	3
2	ORACLE E-BUSINESS SUITE	3
3	ARCHITECTURE	3
3.1	ORACLE APPS TIER	4
3.2	ORACLE DB TIER (HOT BACKUP AND COLD BACKUP)	4
4	FILE SYSTEM: DIRECTORY STRUCTURE	4
5	ORACLE E-BUSINESS SUITE CLONING	5
5.1	SINGLE-NODE CLONING	5
5.2	MULTINODE CLONING	5
6	ORACLE CLONING UTILITIES	5
6.1	RAPID CLONE	5
6.2	AUTOCONFIG	5
7	NETAPP TECHNOLOGIES	5
7.1	FLEXCLONE	5
7.2	SNAPMIRROR	6
8	SYSTEM CONFIGURATION	6
8.1	HARDWARE CONFIGURATION	6
8.2	SOFTWARE CONFIGURATION	6
8.3	NETAPP VOLUME TO DIRECTORY MAPPING	6
8.4	FILE SYSTEM LAYOUT	7
8.5	SNAPMIRROR FREQUENCY	8
8.6	SNAPSHOT BACKUP	8
9	CLONING STEPS	9
9.1	STANDARD RAPID CLONE	9
9.2	CLONE USING AUTOCONFIG	10
10	CREATE TEST ENVIRONMENT FROM PRODUCTION	14
11	ADVANCED CLONING STEPS	14
11.1	SINGLE-NODE TO MULTINODE	14
11.2	MULTINODE TO MULTINODE	14
12	BENEFITS	15
13	REFERENCES	15
14	SUMMARY	15
15	ACKNOWLEDGEMENTS	16

1 INTRODUCTION

This technical report describes the cloning of Oracle® E-Business Suite 11*i* using NetApp® FlexClone® and Oracle's Rapid Clone technologies.

It describes the configuration and steps to clone an E-Business Release 11*i* in a simple, fast, accurate, and cost-effective method. This solution uses NetApp SnapMirror® technology to replicate the file system between storage systems.

2 ORACLE E-BUSINESS SUITE

The Oracle E-Business Suite Architecture is a framework for multitiered, distributed computing that supports Oracle E-Business Suite products. In this model, various servers or services are distributed among three levels, or tiers.

A server (or services) is a process or group of processes that runs on a single machine and provides a particular functionality. For example, Web services process HTTP requests, forms services process requests for activities related to Oracle forms, and the concurrent processing server supports data-intensive programs that run in the background.

A tier is a logical grouping of services, potentially spread across more than one physical machine. The three-tier architecture that composes an Oracle E-Business Suite installation consists of the following:

- **Database tier.** Supports and manages the Oracle Database
- **Application tier.** Supports and manages the various Oracle E-Business Suite components
- **Desktop tier.** Provides the user interface by an add-on component to a standard Web browser

A machine may be referred to as a **node**, particularly in the context of a group of computers that work closely together in a cluster. Each tier may consist of one or more nodes, and each node can potentially accommodate more than one tier. For example, the database can reside on the same node as one or more application tier components.

However, note that a node is also a software concept, referring to a logical grouping of servers.

3 ARCHITECTURE

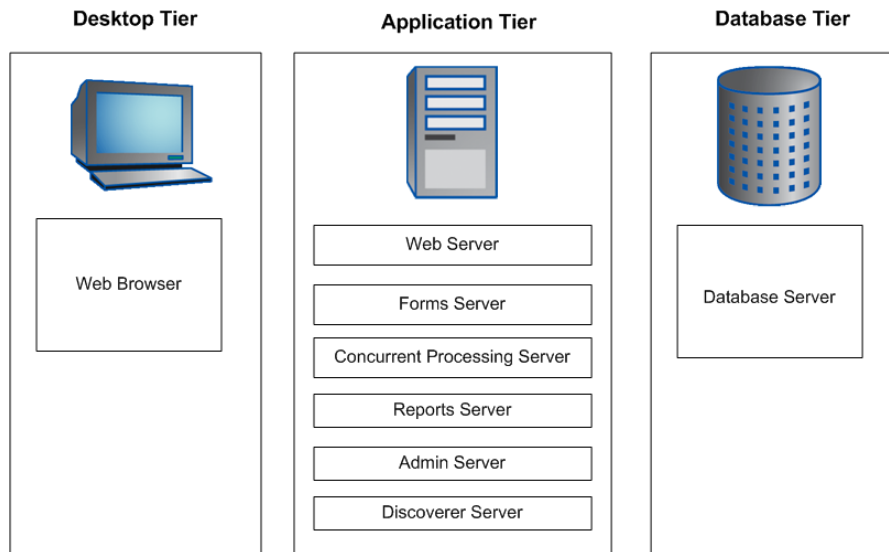


Figure 1) Architecture details.

3.1 ORACLE APPS TIER

The application tier, or “middle tier,” has a dual role: hosting the various servers and service groups that process the business logic, and managing communication between the desktop tier and the database tier.

Three servers or service groups compose the basic application tier for Oracle E-Business Suite:

- Web services
- Forms services
- Concurrent processing server

3.2 ORACLE DB TIER (HOT BACKUP AND COLD BACKUP)

The database tier contains the Oracle Database server that stores and manages all the data maintained by Oracle E-Business Suite. This includes the various types of files where the tables, indexes, and other database objects for your system physically reside, as well as the database executables. The database also stores the Oracle E-Business Suite online help information.

The database server communicates with the services and servers on the application tier, which mediate the communications between the database and the clients: there is no direct communication between the database and clients.

4 FILE SYSTEM: DIRECTORY STRUCTURE

Figure 2 depicts the Oracle Apps directory structure.

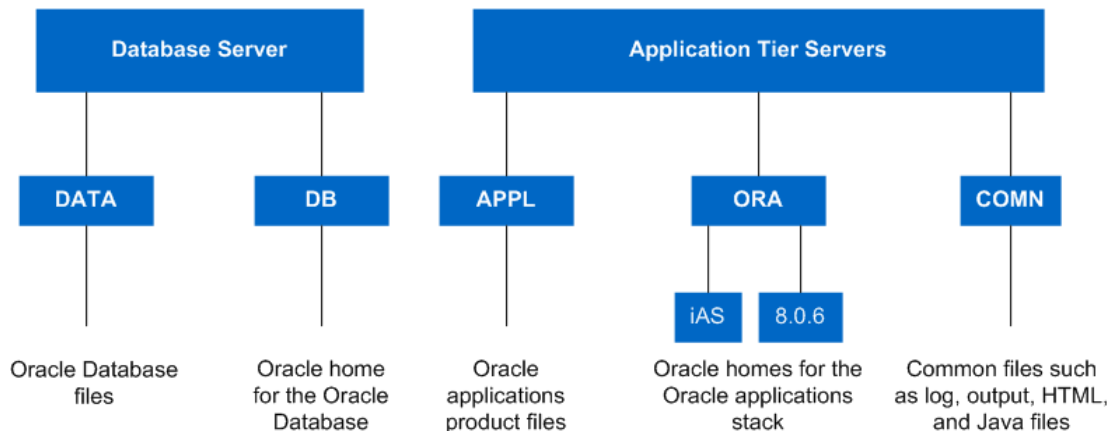


Figure 2) Oracle Apps directory.

- The <dbname>DATA or DATA_TOP directory is located on the database server machine and contains the system tablespaces, redo log files, data tablespaces, index tablespaces, and database files.
- The <dbname>DB directory is located on the database server machine and contains the ORACLE_HOME for the database.
- The <dbname>APPL or APPL_TOP directory contains the product directories and files for Oracle Applications.
- The <dbname>ORA directory contains the ORACLE_HOME for the Applications technology stack components.
- The <dbname>COMN or COMN_TOP (or COMMON_TOP) directory contains directories and files used across products.
- The <dbname>DATA file system contains the data (.dbf) files of the Oracle Database. Rapid Install installs the system, data, and index files in directories below several mount points on the database

server. You can specify the directory names of the mount points on the database server during installation.

- The DB `ORACLE_HOME` (Applications database home) is located in the `<dbname>DB` directory. It contains the files needed for running and maintaining the Oracle Applications database.
- The 8.0.6 directory contains the `ORACLE_HOME` for the developer 6*i* products (forms, reports, and graphics). The product libraries in the 8.0.6 `ORACLE_HOME` are used to relink Oracle Applications executables.
- The iAS directory, also under the `<dbname>ORA`, contains the `ORACLE_HOME` for Oracle9*i* Application Server.

5 ORACLE E-BUSINESS SUITE CLONING

Cloning is creating an identical copy of an existing E-Business Suite system, usually the current production system. With Oracle E-Business Suite 11*i*, simply copying all of the components does not yield a working clone. There are a number of configuration files in the applications layer and within the database that should be modified before the clone system will work as desired.

5.1 SINGLE-NODE CLONING

Single-node cloning is the simplest and most basic configuration where all services run on a single node.

5.2 MULTINODE CLONING

In multinode cloning, there are multiple nodes where the database runs on a dedicated node, and application servers, depending on requirement and configuration, may vary from 1 to *n*.

We will explain the steps needed to perform a clone between single-node system and a single-node system.

6 ORACLE CLONING UTILITIES

6.1 RAPID CLONE

Rapid Clone is the utility that uses the configuration technology used by Rapid Install. Rapid Clone helps in creating a replica of an 11*i* Instance.

6.2 AUTOCONFIG

AutoConfig is a configuration tool that supports automated configuration of an Oracle Applications instance. All of the information required for configuring an Applications instance is collected into a central repository called the Applications Context. When the AutoConfig tool runs, it uses information from the Applications Context file to generate configuration files and update database profiles.

7 NETAPP TECHNOLOGIES

7.1 FLEXCLONE

FlexClone is a powerful feature introduced in Data ONTAP[®] 7G that adds a new level of agility and efficiency to storage operations by allowing an individual to create an instant writable clone of a flexible volume (FlexVol[®] volume).

A FlexClone volume is a writable point-in-time image of a FlexVol volume or another FlexClone volume. With FlexClone, it takes only a few seconds to create a clone of a FlexVol volume, and such a volume can be created without interrupting access to the parent volume on which the clone is based.

The clone volume uses space very efficiently, allowing both the original FlexVol volume and the FlexClone volume to share common data, storing only the data that changes between the original volume and the

clone. This provides a huge potential saving in storage space, resources, and cost. In addition, a FlexClone volume has all the features and capabilities of a regular FlexVol volume, including the ability to be grown or shrunk and the ability to be the source of another FlexClone volume.

FlexClone volumes also enable administrators to access the destination mirror created through the NetApp SnapMirror product. Previously, it was necessary to break the mirror in order to make any changes to the destination copy. With FlexClone, an administrator can now clone a Snapshot™ copy held in the mirror and make it available for both reading and writing at the remote site while allowing the mirror facility to continue running unaffected.

7.2 SNAPMIRROR

The Data ONTAP SnapMirror feature allows an administrator to mirror Snapshot copies of volumes or qtrees from a source volume or qtree to a destination volume or qtree. Replication of data can be performed at regular intervals to make the information available at the destination volume or qtree. The result of this process is an online, read-only volume, or qtree, that contains the same data as the source at the time of the most recent update.

SnapMirror requires a license code. The basic deployment of SnapMirror consists of the following basic components: source volumes or qtrees, destination volumes or qtrees.

SnapMirror can operate in three different types of modes: asynchronous, synchronous, and semisynchronous, namely, SnapMirror Async, SnapMirror Sync, and SnapMirror Semi-Sync.

8 SYSTEM CONFIGURATION

The following are the hardware and software configurations of the system used for developing this technical report.

8.1 HARDWARE CONFIGURATION

- Fujitsu PrimePower 1500
- Solaris™ 10
- Memory 64GB
- NetApp FAS 980
- Disk shelves DS14MK2 SHLF, 18x144GB, FC, 15K RPM, ESH2

8.2 SOFTWARE CONFIGURATION

- NetApp Data ONTAP 7.2 or greater
- NFS, FlexClone, and SnapMirror software licenses
- Oracle E-Business Suite 11.5.9
- Oracle Database: 10.2.0.4

8.3 NETAPP VOLUME TO DIRECTORY MAPPING

1. Create a storage volume for the Database and Apps tier.

```
PRODFLR1> vol create prod_oradata1 aggr01 1024g
PRODFLR2> vol create prod_oradata2 aggr01 1024g
PRODFLR2> vol create prod_dbbin aggr01 500g
PRODFLR2> vol create prod_applmgr aggr01 500g
PRODFLR1> vol create prod_applcsf aggr01 500g
PRODFLR1> vol create prod_arch aggr01 1024g
```

2. Create a storage volume for the Database and Apps tier on a CLONE/STAGE storage system.

```
STGFLR1> vol create vsm_clone_oradata1 aggr01 1024g
STGFLR2> vol create vsm_clone_oradata2 aggr01 1024g
STGFLR2> vol create vsm_clone_dbbin aggr01 500g
STGFLR2> vol create vsm_clone_applmgr aggr01 500g
STGFLR1> vol create vsm_clone_applcsf aggr01 500g
STGFLR1> vol create vsm_clone_arch aggr01 1024g
```

3. Restrict the volume:

```
STGFLR1> vol restrict vsm_clone_oradata1
STGFLR2> vol restrict vsm_clone_oradata2
STGFLR2> vol restrict vsm_clone_dbbin
STGFLR2> vol restrict vsm_clone_applmgr
STGFLR1> vol restrict vsm_clone_applcsf
STGFLR1> vol restrict vsm_clone_arch
```

4. Initialize SnapMirror:

```
PRODFLR1> snapmirror initialize -S prodflr1:prod_oradata1
stgflr1:vsm_clone_oradata1

PRODFLR2> snapmirror initialize -S prodflr2:prod_oradata2
stgflr2:vsm_clone_oradata2

PRODFLR2> snapmirror initialize -S prodflr2:prod_dbbin
stgflr2:vsm_clone_dbbin

PRODFLR2> snapmirror initialize -S prodflr2:prod_applmgr
stgflr2:vsm_clone_applmgr

PRODFLR1> snapmirror initialize -S prodflr1:prod_applcsf
stgflr1:vsm_clone_applcsf

PRODFLR1> snapmirror initialize -S prodflr1:prod_arch stgflr1:vsm_clone_arch
```

5. After the baseline copy is complete, you can edit `/etc/snapmirror.conf` and schedule the automatic transfer based on your setup/requirements. A sample entry in `/etc/snapmirror.conf` is as follows:

```
prodflr02:prod_dbbin stgflr02:vsm_clone_dbbin - 0 6 * *
## above entry will sync vsm_clone_dbbin everyday at 6am
## complete this for all the volumes
```

8.4 FILE SYSTEM LAYOUT

Following is the file system layout:

```
<sid>/applmgr
<sid>/applcsf
<sid>/oracle
<sid>/oradata/data01
<sid>/oradata/data02
<sid>/oradata/redo01
<sid>/oradata/redo02
```

<sid>/oradata/arch

This file system structure is available both at the source and the target. You may design your layout according to your need.

8.5 SNAPMIRROR FREQUENCY

Table 1 describes the SnapMirror frequency based on the file system usage. You may choose a different frequency based on your requirement and design.

Table 1) SnapMirror frequency.

File System	Frequency
<sid>/applmgr	Once every 24 hours
<sid>/applcsf	
<sid>/oracle	
<sid>/oradata/data01	Every hour
<sid>/oradata/data02	
<sid>/oradata/arch	
<sid>/oradata/redo01/02	Not required

8.6 SNAPSHOT BACKUP

The first step in the cloning process is to create a Snapshot copy. These Snapshot copies will be considered as a backup for data on the NetApp volumes.

To back up a running Oracle E-Business Suite system, create a Snapshot copy of the database and application volumes using the following commands:

```
SQL> ALTER DATABASE BEGIN BACKUP;
PRODFLR1> snap create prod_oradata1 hotbackup.0
PRODFLR2> snap create prod_oradata2 hotbackup.0
PRODFLR2> snap create prod_dbbin hotbackup.0
PRODFLR2> snap create prod_applmgr hotbackup.0
PRODFLR1> snap create prod_applcsf hotbackup.0
SQL> ALTER DATABASE END BACKUP;
SQL> ALTER SYSTEM SWITCH LOGFILE;
SQL> ARCHIVE LOG ALL
PRODFLR1> snap create prod_arch hotbackup.0
```

Note: This is also called a Snapshot backup.

9 CLONING STEPS

Following are two types of cloning method you can use:

- Standard rapid clone
- Cloning use AutoConfig

9.1 STANDARD RAPID CLONE

The following procedure illustrates the most preferred way of cloning an E-Business Environment that is supported by Oracle. There are several prerequisite procedures required before starting the clone. See **Note: 230672.1** for the detailed steps and prerequisites. Steps at a high level are listed below (considering it's a shared APPL_TOP):

1. Run `adpreclone.pl dbTier` on the Database Node.
2. Run `adpreclone.pl appstier` on the Application Node.
3. Create a hotbackup Snapshot copy for application and database volumes:

```
SQL> ALTER DATABASE BEGIN BACKUP;
PRODFLR1> snap create prod_oradata1 hotbackup.0
PRODFLR2> snap create prod_oradata2 hotbackup.0
PRODFLR2> snap create prod_dbbin hotbackup.0
PRODFLR2> snap create prod_applmgr hotbackup.0
PRODFLR1> snap create prod_applcsf hotbackup.0
SQL> ALTER DATABASE END BACKUP;
SQL> ALTER SYSTEM SWITCH LOGFILE;
SQL> ARCHIVE LOG ALL
PRODFLR1> snap create prod_arch hotbackup.0
```

4. Run `snapmirror update` to sync the volume on remote/target storage system:

```
STGFLR1> snapmirror update stgflr1:vsm_clone_oradata1
STGFLR2> snapmirror update stgflr2:vsm_clone_oradata2
STGFLR2> snapmirror update stgflr2:vsm_clone_dbbin
STGFLR2> snapmirror update stgflr2:vsm_clone_applmgr
STGFLR1> snapmirror update stgflr2:vsm_clone_applcsf
STGFLR1> snapmirror update stgflr1:vsm_clone_arch
```

Note: SnapMirror is not required if the storage system is the same for the destination instance.

5. Create a FlexClone volume from a mirrored volume:

```
STGFLR1> vol clone create dev1_oradata1 -b vsm_clone_oradata1 hotbackup.0
STGFLR2> vol clone create dev1_oradata2 -b vsm_clone_oradata2 hotbackup.0
STGFLR2> vol clone create dev1_dbbin -b vsm_clone_dbbin hotbackup.0
STGFLR2> vol clone create dev1_applmgr -b vsm_clone_applmgr hotbackup.0
STGFLR1> vol clone create dev1_applcsf -b vsm_clone_applcsf hotbackup.0
STGFLR1> vol clone create dev1_arch -b vsm_clone_arch hotbackup.0
```

6. Split the FlexClone volume.

```
STGFLR1> vol clone split dev1_oradata1
STGFLR2> vol clone split dev1_oradata2
STGFLR2> vol clone split dev1_dbbin
STGFLR2> vol clone split dev1_applmgr
STGFLR1> vol clone split dev1_applcsf
STGFLR1> vol clone split dev1_arch
```

Note: Make sure that sufficient space is available for a SPLIT operation and that the operation completes successfully.

To monitor the split operation, you can use the `vol clone split status` command.

IMPORTANT:

This is an optional step. The split command creates an exact replica on the target system. The drawback is that it occupies the same amount of space as source (your space requirement becomes twice that of the source).

7. Complete the database recovery

8. Run SETUP_CLEAN as APPS user after the database is up on target/ cloned environment:

```
SQL> EXEC FND_CONC_CLONE.SETUP_CLEAN
```

9. Run the `adcfgclone.pl` command on the application tier.

10. Complete postclone steps specific to your environment.

11. Bring up the services.

9.2 CLONE USING AUTOCONFIG

This is a nonstandard way of cloning an E-Business Environment that may not be supported by Oracle. There are several steps required before actually starting the clone.

For detailed steps and prerequisites, see **Note: 230672.1**. The figures shown here are samples only.

The following procedure outlines how to create a clone (considering that it's a shared `APPL_TOP`).

1. Create a hot backup Snapshot copy for application and database volumes.

```
SQL> ALTER DATABASE BEGIN BACKUP;
PRODFLR1> snap create prod_oradata1 hotbackup.0
PRODFLR2> snap create prod_oradata2 hotbackup.0
PRODFLR2> snap create prod_dbbin hotbackup.0
PRODFLR2> snap create prod_applmgr hotbackup.0
PRODFLR1> snap create prod_applcsf hotbackup.0
SQL> ALTER DATABASE END BACKUP;
SQL> ALTER SYSTEM SWITCH LOGFILE;
SQL> ARCHIVE LOG ALL
PRODFLR1> snap create prod_arch hotbackup.0
```

2. Run `snapmirror update` to sync the volume on the remote/target storage system.

```
STGFLR1> snapmirror update stgflr1:vsm_clone_oradata1
STGFLR2> snapmirror update stgflr2:vsm_clone_oradata2
STGFLR2> snapmirror update stgflr2:vsm_clone_dbbin
```

```

STGFLR2> snapmirror update stgflr2:vsm_clone_applmgr
STGFLR1> snapmirror update stgflr2:vsm_clone_applcsf
STGFLR1> snapmirror update stgflr1:vsm_clone_arch

```

Note: SnapMirror is not required if the storage system is the same for the destination instance.

3. Create a FlexClone volume from a mirrored volume.

```

STGFLR1> vol clone create dev1_oradata1 -b vsm_clone_oradata1 hotbackup.0
STGFLR2> vol clone create dev1_oradata2 -b vsm_clone_oradata2 hotbackup.0
STGFLR2> vol clone create dev1_dbbin -b vsm_clone_dbbin hotbackup.0
STGFLR2> vol clone create dev1_applmgr -b vsm_clone_applmgr hotbackup.0
STGFLR1> vol clone create dev1_applcsf -b vsm_clone_applcsf hotbackup.0
STGFLR1> vol clone create dev1_arch -b vsm_clone_arch hotbackup.0

```

```

Sun Microsystems Inc.   SunOS 5.10   Generic January 2005
Volume "dev1_oradata" is now offline.
Volume "dev1_oradata" destroyed.
Volume "dev1_oradata2" is now offline.
Volume "dev1_oradata2" destroyed.
Volume "dev1_interface" is now offline.
Volume "dev1_interface" destroyed.
Volume "dev1_bin01" is now offline.
Volume "dev1_bin01" destroyed.
Volume "dev1_app_shared" is now offline.
Volume "dev1_app_shared" destroyed.
Volume "dev1_app_shared_appltop" is now offline.
Volume "dev1_app_shared_appltop" destroyed.
Creation of clone volume "dev1_oradata" has completed.
Creation of clone volume "dev1_oradata2" has completed.
Creation of clone volume "dev1_bin01" has completed.
Creation of clone volume "dev1_app_shared" has completed.
Creation of clone volume "dev1_app_shared_appltop" has completed.
Creation of clone volume "dev1_interface" has completed.

```

Figure 3) An example of an offline/destroy and recreate FlexVol volumes.

4. Split the FlexClone volume:

```

STGFLR1> vol clone split dev1_oradata1
STGFLR2> vol clone split dev1_oradata2
STGFLR2> vol clone split dev1_dbbin
STGFLR2> vol clone split dev1_applmgr
STGFLR1> vol clone split dev1_applcsf
STGFLR1> vol clone split dev1_arch

```

Note: Make sure that sufficient space is available for SPLIT operation, and the operation completes successfully.

Use the `vol clone split status` command to monitor the SPLIT operation.

IMPORTANT:

This is an optional step. The split command creates an exact replica on the target system. The drawback is that it occupies the same amount of space as source (your space requirement becomes twice that of the source).

5. Complete the database recovery.

```
SQL*Plus: Release 10.2.0.4.0 - Production on Thu Mar 18 21:53:49 2010
Copyright (c) 1982, 2007, Oracle. All Rights Reserved.

Connected to an idle instance.

ORACLE instance started.

Total System Global Area 4294967296 bytes
Fixed Size                2238584 bytes
Variable Size             926472072 bytes
Database Buffers          3355443200 bytes
Redo Buffers              10813440 bytes

Control file created.

ORA-00279: change 6699416920207 generated at 03/18/2010 17:05:12 needed for
thread 1
ORA-00289: suggestion : /dev1/arch/DEV1321290_1459126785.arc
ORA-00280: change 6699416920207 for thread 1 is in sequence #321290

ORA-00279: change 6699416930142 generated at 03/18/2010 17:06:45 needed for
thread 1
ORA-00289: suggestion : /dev1/arch/DEV1321291_1459126785.arc
ORA-00280: change 6699416930142 for thread 1 is in sequence #321291
ORA-00278: log file '/dev1/oracle/hCLONE/DEV1321290_1459126785.arc' no
longer needed for this recovery

ORA-00279: change 6699416930146 generated at 03/18/2010 17:06:45 needed for
thread 1
ORA-00289: suggestion : /dev1/arch/DEV1321292_1459126785.arc
ORA-00280: change 6699416930146 for thread 1 is in sequence #321292
ORA-00278: log file '/dev1/oracle/hCLONE/DEV1321291_1459126785.arc' no
longer needed for this recovery

Media recovery cancelled.

Database altered.

Tablespace altered.

Tablespace altered.
```

Figure 4) Steps to complete the media recovery on the target instance (output is hypothetical only).

6. Run SETUP_CLEAN as APPS user after the database is up on the target/cloned environment:

```
SQL> EXEC FND_CONC_CLONE.SETUP_CLEAN
```

7. Run AutoConfig on each application tier node.

```
AutoConfig is configuring the Applications environment...
AutoConfig will consider the custom templates if present.
Using APPPL_TOP location      : /dev1/applmgr/1159
Classpath                     : /dev1/applmgr/common/util/j2sdk1.4.2_13/jre/lib/rt.jar:/dev1/applmgr/common/util/j2sdk1.4.2_13/lib/dt.jar:/dev1/applmgr/common/util/j2sdk1.4.2_13/lib/tools.jar:/dev1/applmgr/common/java/appsborg2.zip:/dev1/applmgr/common/java
Using Context file           : /dev1/applmgr/1159/admin/DEV1_hostname1.xml

Context Value Management will now update the Context file

Updating Context file...COMPLETED

Attempting upload of Context file and templates to database...COMPLETED

Configuring templates from all of the product tops...
Configuring AD_TOP.....COMPLETED
Configuring FND_TOP.....COMPLETED
Configuring ICX_TOP.....COMPLETED
Configuring IEO_TOP.....COMPLETED
Configuring ECK_TOP.....COMPLETED
Configuring BIS_TOP.....COMPLETED
Configuring GL_TOP.....COMPLETED
Configuring AMS_TOP.....COMPLETED
Configuring WSH_TOP.....COMPLETED
Configuring OKE_TOP.....COMPLETED
Configuring OKL_TOP.....COMPLETED
Configuring OKS_TOP.....COMPLETED
Configuring CSF_TOP.....COMPLETED
Configuring XNC_TOP.....COMPLETED
Configuring IGS_TOP.....COMPLETED
Configuring IBY_TOP.....COMPLETED
Configuring PA_TOP.....COMPLETED
Configuring JTF_TOP.....COMPLETED
Configuring MWF_TOP.....COMPLETED
Configuring CN_TOP.....COMPLETED
Configuring CSI_TOP.....COMPLETED
Configuring WIP_TOP.....COMPLETED
Configuring CSE_TOP.....COMPLETED
Configuring EAM_TOP.....COMPLETED
Configuring IMT_TOP.....COMPLETED
Configuring FTE_TOP.....COMPLETED
Configuring ONT_TOP.....COMPLETED
Configuring AR_TOP.....COMPLETED
Configuring AHL_TOP.....COMPLETED
Configuring CSD_TOP.....COMPLETED
Configuring IGC_TOP.....COMPLETED

AutoConfig completed successfully.
```

Figure 5) Running AutoConfig to complete the clone (output is hypothetical only).

8. Complete the post clone steps specific to your environment.

9. Bring up the services.

```
adapcctl.sh version 115.54

Apache Web Server Listener is not running.
Starting Apache Web Server Listener (dedicated HTTP) ...
Removing gantt cache directory
Successfully removed the gantt cache directory:/dev1/applmgr/common/html/cabo/images/cache/gantt
Apache Web Server Listener (PLSQL) is not running.
Starting Apache Web Server Listener (dedicated PLSQL) ...

adapcctl.sh: exiting with status 0
```

Figure 6) Bringing up the Apache services (output is hypothetical only).

```
adalnctl.sh version

Checking for FNDFS executable.
Starting listener process APPS_DEV1.

adalnctl.sh: exiting with status 0
```

Figure 7) Bringing up the APPS Listener (output is hypothetical only).

```
You are running adcmctl.sh version 115.28

Starting concurrent manager for DEV1 ...
Starting DEV1_0318@DEV1 Internal Concurrent Manager
Default printer is noprint

adcmctl.sh: exiting with status 0
```

Figure 8) Bringing up the concurrent managers (output is hypothetical only).

Note: You can also use the `adstrta1.sh` script.

10 CREATE TEST ENVIRONMENT FROM PRODUCTION

Follow the steps in sections 9.1 or 9.2 based on your requirements.

11 ADVANCED CLONING STEPS

11.1 SINGLE-NODE TO MULTINODE

Single-node to multinode cloning differs only if you do not have shared `APPL_TOP`. If you have shared `APPL_TOP`, then there is no change. The only difference would be to complete the FlexClone operation as many times as the number of nodes you have instead of just one in the case of shared `APPL_TOP`. Other steps specific to Oracle are documented in **Note: 230672.1**.

11.2 MULTINODE TO MULTINODE

Multinode to multinode cloning requires no change in the case of shared `APPL_TOP`. For nonshared `APPL_TOP`, you must set up the SnapMirror relationship based on your requirement, and you can follow cloning steps specific to Oracle as documented in **Note: 230672.1**.

12 BENEFITS

The unique NetApp software features add substantial value for E-Business Suite deployments, such as:

- **Enhanced database backup.** Protection of the Oracle Database at the core of the E-Business Suite is critically important.
- **Rapid database recovery.** If a problem does occur, then the database must be brought back online quickly so that business operations can continue.
- **Flexible options for business continuance.** Disasters happen. Every business needs a strategy to protect critical business data from site-wide disasters.
- **Simplified storage operations.** Keeping ahead of the storage required by an active database can be no small task. NetApp storage appliances make it easy to provide additional storage for Oracle tablespaces and roll back segments without taking the database offline.
- **Shared code.** The ability to share E-Business code—`APPL_TOPs`—is a new capability of 11*i*. NetApp storage appliances make it simple to share a single `APPL_TOP` across many nodes.

Table 2) Benefits.

Task	NetApp FlexClone Benefit
Application testing	<ul style="list-style-type: none">• Make necessary changes to infrastructure without worrying about the production system.• Avoid making untested changes on production systems to reduce downtime.• Provides low risk, low stress, and higher confidence when applying changes to production after testing on FlexClone volumes.
System deployment	<ul style="list-style-type: none">• Maintain a template environment and use FlexClone volumes to build and deploy identical environments.• Create clones faster, simpler, and available as needed for reliable and predictable testing.
IT operations	<ul style="list-style-type: none">• Maintain multiple copies of production systems, saving disk space.• Instantly refresh clones to create a working set that is as close to a live production system as is practical.

13 REFERENCES

Oracle Metalink **Note: 230672.1**.

14 SUMMARY

Using NetApp FlexClone and SnapMirror technologies with Oracle Rapid Clone greatly simplifies the Oracle E-Business Suite cloning process. It allows the copies to be created quickly, efficiently, and independently of the server. This maximizes the resources on the source server available for production/online use.

15 ACKNOWLEDGEMENTS

Bill L Heffelfinger, NetApp

Lynne Thieme, NetApp

Daniel Morgan, NetApp

Neil Gerren, NetApp

Uday Shet, NetApp

Esther Smitha, NetApp

NetApp provides no representations or warranties regarding the accuracy, reliability or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.