



Technical report:
**Oracle E-Business Suite and IBM System
Storage N series with SnapMirror**
Cloning for Data Protection

• • • • • • • • •

Document NS3300-0

January 2, 2008



Table of contents

Abstract	3
Introduction	3
E-Business Suite Cloning Overview	3
System Configuration	5
Hardware Configuration.....	5
Software Configuration	5
Storage Configuration.....	5
Implementing SnapMirror for E-Business Suite Cloning	7
SnapMirror Setup and Operation.....	7
Breaking the Mirror in Preparation for Use	8
Clone Resynchronization Procedures	10
Conclusion	10
Trademarks and Special Notices	11



Abstract

This technical report covers cloning Oracle E-Business Suite 11i using IBM System Storage N series with SnapMirror and Oracle Rapid Clone. It describes the configuration and steps to clone an E-Business Suite 11i in a simple, fast, accurate, and cost-effective method. In today's competitive world, any advantage provides an edge, and any efficiency gained by simplifying the day-to-day operating environment is a huge advantage, and this is what this paper provides.

Introduction

IBM® and Oracle® have mutual customers that require a solution for Oracle E-Business Suite 11i cloning, which is needed to maintain various company activities such as reporting, testing, custom development, and data fixes. The following is a simple, fast, and cost-effective joint IBM and Oracle solution for mutual customers. It uses technology from both companies, IBM System Storage™ N series with SnapMirror® and Oracle Rapid Clone.

There are a myriad of reasons why companies that have implemented E-Business Suite need to clone the apps environment:

- Custom development
- Testing and integration activities with periodic refresh from production
- Data reporting from an instance other than production with reasonably current data
- Preserving month-end and year-end instances for reporting and other purposes
- Testing a data fix before doing it on the production system
- Testing migration or upgrade procedures.

Being able to accomplish these tasks without impacting production is done by cloning the production environment and using the clone in place of production. Being able to clone quickly, reliably, and accurately while maintaining cost effectiveness provides for operational efficiency and boosts competitive advantages by having the production system continuously available for the business while simultaneously being able to work on the cloned system. The fact that cloning is done every day and maybe even multiple times daily magnifies the need to streamline the cloning process.

E-Business Suite Cloning Overview

Cloning is simply making an identical copy of an existing E-Business Suite system, usually the current production system. With Oracle E-Business Suite 11i, 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 need to be modified before the clone system will work as desired.

Moreover, depending on the system configuration available, as depicted in Figure 1, cloning can be any of the following:

- From a multinode system to a similarly configured multinode system
- From a single-node system to a multinode system
- From a single-node system to a single-node system.

Each configuration represents a different challenge to the cloning process. However, with Oracle's Rapid Clone technology, all these configurations are handled with accuracy and ease. The 11.5.8 Rapid Install creates an E-Business Suite system that is Rapid Clone enabled. Pre-release 11.5.8 systems can be configured to use Rapid Clone by migrating to AutoConfig, installing Rapid Clone, and applying other prerequisite patches on the source system. Read Oracle's white paper "Cloning Oracle Applications Release 11*i* with Rapid Clone" for detailed information about Oracle's Rapid Clone technology. Also read Oracle MetaLink note 230672.1 with the same title, which describes the prerequisites and steps of the Rapid Clone process. The white paper and MetaLink note are available from the Oracle website and are requirements for successfully cloning Oracle E-Business Suite systems.

Another useful document is the "Oracle Applications Maintenance Procedures Release 11*i* (11.5.9)," part no. B10641-01, a manual that comes with the distribution CDs. Chapter 4 covers application cloning and includes cloning procedures and requirements.

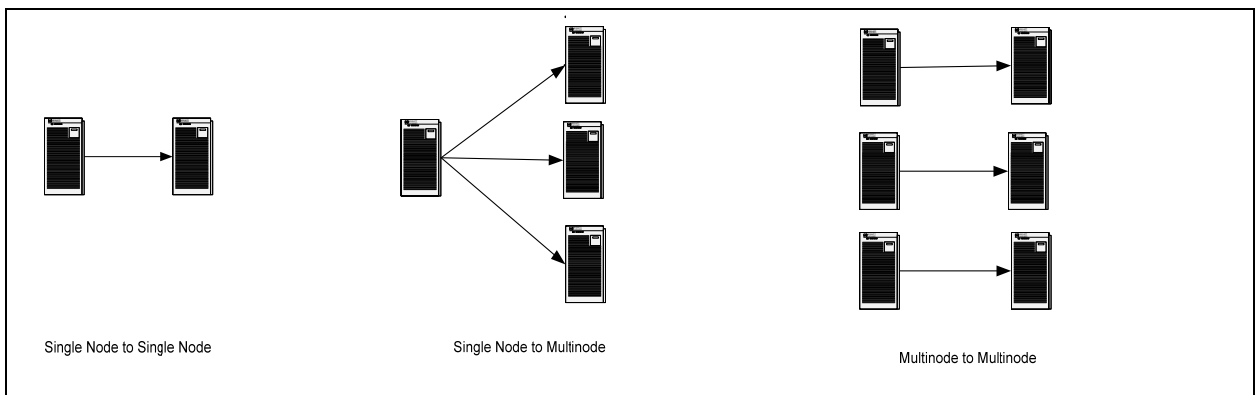


Figure 1) Different Server Geometry Configuration Options for E-Business Suite Cloning.

If you are using an IBM N series as your storage for Oracle E-Business Suite, then use SnapMirror in place of copying files in the steps of the Rapid Clone process. Using SnapMirror allows you to create copies of files at the storage level, quickly, efficiently, and independently of the server. Using SnapMirror eliminates the impact of the copying on the source host and thus maximizes the resources available for the production/source system. Also, the mirroring process can be started well in advance of the actual cloning and thus shortens the window of the cloning process, since only the last incremental updates need to be transferred to complete the copy.

Cloning the Oracle E-Business Suite 11*i* using IBM System Storage N series with Snapshot™ is another possible solution if IBM N series storage is already in use. However this solution is slower since it involves copying the snapshot to a different volume within the same storage or to a separate storage. If the same storage used by production is used for the clone, then there will be a performance hit to the production system due to the sharing of the storage. If it is a different storage then the transfer is either done manually or automated using scripts which is more complex than the way IBM N series with SnapMirror handles the data transfer which is done automatically. Also, SnapMirror gives you the capability to resynchronize so subsequent cloning are completed faster.



System Configuration

The following are the hardware and software configurations of the system used for this paper:

Hardware Configuration

Source server	Sun™ Fire 280R Server running Solaris™ 8
Target server	Sun Fire 280R Server running Solaris 8
Storage system	IBM N series filer

Software Configuration

Oracle	E-Business Suite 11 <i>i</i> (version 11.5.9) with AutoConfig
IBM	IBM System Storage N series with Data ONTAP™

Storage Configuration

- The same mount points are used on both source and target servers.

/vol/apps	-	source volume for Oracle Apps 11 <i>i</i>
/vol/orabin	-	source volume for Oracle 9 <i>i</i> ™ binaries
/vol/oradata	-	source volume for database files
/vol/oralogs	-	source volume for online redo logs
/vol/archlogs	-	source volume for archived logs
- The cloned volumes are the same size and geometry as the source volumes.

/vol/apps_mirr	-	clone of source /vol/apps
/vol/orabin_mirr	-	clone of source /vol/orabin
/vol/oradata_mirr	-	clone of source /vol/oradata
/vol/oralogs_mirr	-	clone of source /vol/oralogs
/vol/archlogs_mirr	-	clone of source /vol/archlogs_mirr
- The source server stores all the database and Oracle E-Business Suite files in the IBM N series filer storage system, and the target server does the same, but with a separate IBM N series filer. This ensures that when the cloned Oracle Apps instance is brought online, the performance of the production system, which is the source, is not affected at all, given the use of separate servers and separate storage. See Figure 2.

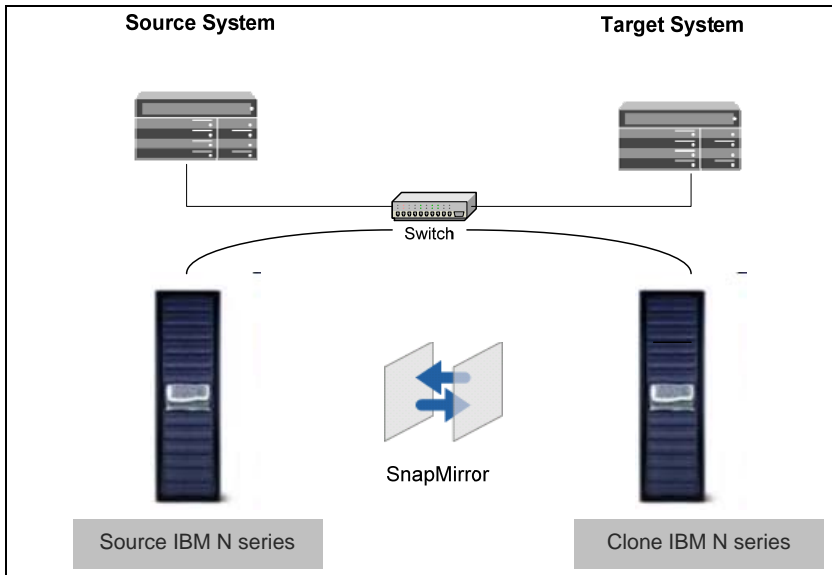


Figure 2) Server and Storage Configuration.

- The source volumes have been created and mounted by the source server.
- Oracle E-Business Suite is in working condition on the source server.
- The source and target servers are configured to mount the volumes from the correct storage systems to the correct mount points and with the proper permissions.

Usage	Source Volume Name	Mount Point on Both Source and Target Servers	Clone Volume Name
E-Business Suite files (APPL_TOP, COMMON_TOP, iAS)	/vol/apps	/oracle/apps	/vol/apps_mirr
ORACLE_HOME	/vol/orabin	/oracle/server	/vol/orabin_mirr
Database files	/vol/oradata	/oracle/data	/vol/oradata_mirr
Online redo logs	/vol/oralogs	/oracle/logs	/vol/oralogs_mirr
Archive logs	/vol/archlogs	/oracle/arch	/vol/archlogs_mirr

Table 1) Storage Configuration of Source and Target Servers.



Implementing SnapMirror for E-Business Suite Cloning

IBM N series with SnapMirror technology provides asynchronous mirroring of data between filer volumes. Data on the source volume is periodically replicated to the target at a user-definable time interval, with the range being from one minute to one month. At the end of each replication event, the mirror target volume becomes an exact block-for-block copy of the mirror source volume. At that point, the two volumes share identical data content and characteristics. The mirror is initialized by effectively copying the entire source volume to the target volume. Once this initial copy is complete, replication events thereafter copy only changed blocks from the source volume to the target volume. This provides a highly efficient data replication mechanism.

SnapMirror Setup and Operation

Setting up SnapMirror is relatively simple and quick. Basic requirements are:

- Identify the source and target volumes.
 - Create a pair of configuration files.
 - Start the mirror.
1. On the source and clone IBM N series storage, make sure that SnapMirror is licensed.
 2. Edit the configuration files on both source and clone storage, and, for simplicity, the respective storage systems are called source and clone.
 - a. `/etc/snapmirror.allow`

The content of this file should be the name of the storage that will participate in the SnapMirror replication, in this case:

```
source
clone
```

- b. `/etc/snapmirror.conf`

The content of this file should be the source storage and source volume; target storage and target volume; SnapMirror options; and schedule expressed in minutes, hours, day of the month, and day of the week. For simplicity, we chose default options and update the mirror once a minute.

```
source: apps clone: apps_mirr - * * * *
source: orabin clone: orabin_mirr - * * * *
source: oradata clone: oradata_mirr - * * * *
source: archlogs clone: archlogs_mirr - * * * *
```

3. On the clone storage, restrict the target volumes so the mirror can be initialized. Log in to the clone storage and use the **vol restrict volname** command to restrict all the target volumes.

```
vol restrict apps_mirr
vol restrict orabin_mirr
vol restrict oradata_mirr
vol restrict archlogs_mirr
```

4. Initialize the mirror, log in to the clone storage, and use the **snapmirror initialize** command.

```
snapmirror initialize -S source:apps clone:apps_mirr
snapmirror initialize -S source:orabin clone:orabin_mirr
snapmirror initialize -S source:oradata clone:oradata_mirr
snapmirror initialize -S source:archlogs clone:archlogs_mirr
```

To monitor the status of the SnapMirror operation, use the command **snapmirror status**. Make sure that the initializations have completed and the mirror (or mirrors) is established. Once the mirror is established, it is updated incrementally every minute (this is what was specified in `snapmirror.conf`).

At this point the mirrors are established, and the next step is to break the mirror and use the clones.

Breaking the Mirror in Preparation for Use

1. For the volumes containing the E-Business Suite files and the Oracle binaries (ORACLE_HOME), break the mirror by issuing the **snapmirror quiesce** and **snapmirror break** commands for each target volume from the clone storage.

```
snapmirror quiesce volume-name
snapmirror break volume-name
```

where *volume-name* is the target volume name.

For example, for the apps volume, break the mirror by doing the following from the clone storage:

```
snapmirror quiesce apps_mirr
snapmirror break apps_mirr
```

Also break the mirror for the target ORACLE_HOME volume.

```
snapmirror quiesce orabin_mirr
snapmirror break orabin_mirr
```


2. For the database and archive log volumes, you'll need to do a hot backup of the database and issue the **snapmirror update**, **snapmirror quiesce**, and **snapmirror break** commands.
 - Force a switch and archive the current redo logs.
SQL> alter system archive log current;
 - Put all the tablespaces in backup mode (part of hot backup).
SQL> alter tablespace *tablespace-name* begin backup;
 - From the clone storage, issue the **snapmirror update**, **snapmirror quiesce**, and **snapmirror break** commands to split the mirror for the volume that contains the database files.
snapmirror update oradata_mirr
snapmirror quiesce oradata_mirr
snapmirror break oradata_mirr
 - Take the tablespaces out of backup mode.
SQL> alter tablespace *tablespace-name* end backup;
 - Again archive the current log to ensure that the archive log is set for database recovery.
 - From the clone storage, split the mirror for the archive logs volume.
snapmirror update archivelogs_mirr
snapmirror quiesce archivelogs_mirr
snapmirror break archivelogs_mirr
 - Mount the cloned volumes from the target server and recreate the control file with the new clone SID. Proceed to recover the database by applying the archived logs from the hot backup.
3. At this point you have a working clone database. Proceed with the Rapid Clone steps to complete the cloning process.



Clone Resynchronization Procedures

To refresh the clone from the source once the SnapMirror source/target relationships have been established, follow the steps outlined below. Resynchronization operations are much faster than initialization, since only the changed blocks are transferred from the source storage.

1. Shut down Oracle E-Business Suite 11*i*, including the database on the target server.

Use the `adstpall.sh` script to shut down Oracle E-Business Suite 11*i* and SQL*Plus to shut down the database.

2. Resynchronize the SnapMirror target volumes; from the clone storage use the **snapmirror resync** command.

```
snapmirror resync -f apps_mirr
snapmirror resync -f oradata_mirr
snapmirror resync -f archivelogs_mirr
```

Monitor the SnapMirror operation to make sure that the resynchronization has completed. From the clone storage use **snapmirror status** to monitor.

3. Once the target storage is in sync, proceed to break the mirror relationships; use the **snapmirror update**, **snapmirror quiesce**, and **snapmirror break** commands. For the database volumes, make sure to put the database in hot backup mode, as previously discussed in the “Breaking the Mirror in Preparation for Use” section.

All the SnapMirror commands can be saved in scripts. The only additional requirement is to establish RSH access to the storage from the server where the scripts (in this case the target server) will be executed. Scripting the whole process allows for automation, which lessens the possibility of user errors and further simplifies the cloning process.

Conclusion

Using SnapMirror simplifies the E-Business Suite cloning process; the use of storage-level mirroring allows the copies to be done quickly, efficiently, and independently of the server. This maximizes the resources on the source server available for production/online use. The mirroring can also be started way ahead of time if possible so that only the last incremental changes need to be transferred during cloning, thus shortening the whole process.

This solution provides for an optimal process for E-Business Suite cloning. This in turn gives you flexibility in the frequency of when cloning is done to satisfy the cloning requirements of the enterprise, be it for development, testing, reporting, or whatever the case may be.

SnapMirror is easy to set up, configure and maintain and most important cost-effective as a mirroring solution. Using IBM N series storage with SnapMirror in conjunction with Oracle’s Rapid Clone greatly simplifies and speeds up the E-Business Suite cloning process. This provides users with the maximum benefit out of their investment in the overall system.



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.

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, SnapMirror and Snapshot are trademarks or registered trademarks of Network Appliance, Inc., in the U.S. and other countries.

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.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

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.