Technical White Paper
August 2010

Forward Recovery of an Oracle Database Using Data Replicator Software

ORACLE

Contents

Introduction

Restoring data in a database to a previous point in time often means restoring a previous backup from disk or tape, and then rolling forward to the point in time before a problem occurred. One of the best ways to accomplish this task is to use the Oracle Roll Forward Recovery process. By using the Oracle Roll Forward Recovery process, you can start with a full backup image of the database, and then roll forward through the archive logs to recover to a specific System Change Number (SCN), Date and Time, or until you, as the Oracle Database administrator, cancel the recovery process.

The one restriction of using the Oracle Roll Forward Recovery process is that the entire database is restored to this point in time. If you are required to restore the database without disrupting the existing database, you must surface the data on another host, and then isolate the process of retrieving this data. For example, suppose your company wants to create a standby database environment where it is possible to recover a database from the last good point in time full database backup. Your company wants to be able to roll forward all of the logs necessary to bring up the database to a specific SCN or to a specific date and time before data corruption or data errors took place. Just as likely, suppose that a program deleted many accounts in error. You must be able to retrieve these accounts from a point in time before the error happened. After retrieving this data, you can verify the information, and then insert the data back into the production database.

This document shows you a method for completing the forward recovery of an Oracle Database using Data Replicator Software (DRS) with the Oracle Roll Forward Recovery. You have done similar restorations many times before. This method, however, uses a secondary server, allowing the primary database to remain in production throughout the procedure.

The intended reader for this document is an Oracle Database administrator with experience in the following areas:

- Oracle Database and its related components
- Cloning an Oracle Database
- Storage, including an understanding of data services and the premium features that support those services

Benefits of Using This Method

When you use the Oracle Roll Forward Recovery process along with DRS, you can create a copy of the database, and then clone it onto a different server. Then you can perform certain critical functions on the copy rather than on the production database. Using a copy for critical functions has several advantages:

- No disruption of the source production database
- Fewer staff resources
- Anytime duplication of the database
- Data replication for recovery of the production database

The DRS premium feature is an option for real-time replication of data that creates copies of the source volumes on a second target storage array. The second storage array might reside within the same computing environment, or for disaster recovery purposes, might reside many miles away or even in a different state or country. A fully-synchronized DRS setup provides many ways to use the remote mirrors without disrupting the primary source volumes. DRS provides rapid recovery from a disaster or catastrophic failure of the primary site by using role reversals to promote the secondary volumes to a primary role. Hosts can then read and write to the newly promoted volumes and business operations can continue.

By implementing the DRS premium feature, the issue of restoring the data from disk or tape or restoring the whole source database to a point in time in the past is eliminated. DRS lets you mirror the data to another set of volumes, and then stop the mirroring process at any time. You can use these older mirrored volumes to bring up a copy of the source database, and then restore the database on another server with no impact to the source database. After the recovery process is complete, you can use Oracle Roll Forward Recovery to roll forward the data in either of the following ways:

- Retrieve the data from this cloned database and update the source database.
- Reverse the DRS process and sync the source database drives from the cloned database drives.

Setting Up the System Environment

Use the reference architectures as guidelines for setting up your own system environment. The reference architectures serve as examples only. You can modify the system environment as necessary to suit your needs. There are two reference architectures:

- The 6540 reference architecture consists of one primary sever and one secondary server attached to two 6540 controllers. See Figure 1 on page 5.
- The Sun Storage 6580/6780 reference architecture consists of one primary sever and one secondary server attached to one Sun Storage 6580/6780 controller and one 6540 controller. See Figure 2 on page 6.

The reference architectures were developed for testing seven common scenarios that an Oracle Database administrator would regularly encounter. This document describes only one of those

scenarios. The other six scenarios are described in separate documents. Those documents are listed in "Appendix A: References."

The reference architectures use three premium features—Volume Copy, Snapshot, and Data Replicator Software (DRS)—to test various backup and recovery methods and to test multiple development and upgrade scenarios. The specific Oracle scenario described in this document might not use all three premium features.

Technical Specifications

Here are the specifications for components of the reference architectures.

Operating System Information

Red Hat Enterprise Linux 5.1 Version 2.6.18-53.el5

The RDAC driver is MPP Driver Version 99.03.C000.0005

Oracle Information

Oracle Database EE 10g Release 2 Version 10.2.0.3

Oracle Database EE 11g Release 1 Version 11.1.0.6

Storage Information

6540

Firmware Version 07.60.36.13
NVSRAM Version N6091-760843-004

Sun Storage 6580/6780

Firmware version 07.60.36.13
NVSRAM version N7091-760843-004

Sun Storage Common Array Manager (CAM) 6.7.0

Volume Configuration

For each reference architecture, the primary storage array has the following volumes:

- Oracle distribution volume /u01
- Primary Oracle Database volumes
- Snapshot volumes
- Oracle Volume Copy volumes

For each reference architecture, the secondary storage array has the following volumes:

- Snapshot volumes
- Secondary DRS volumes for the Oracle distribution and database volumes

Reference Architecture

Figure 1 shows that the 6540 reference architecture consists of one primary server and one secondary server attached to two 6540 controllers.

Figure 2 on page 6 shows that the Sun Storage 6580/6780 reference architecture consists of one primary server and one secondary server attached to one Sun Storage 6580/6780 controller and one 6540 controller.

Figure 1 6540 Reference Architecture

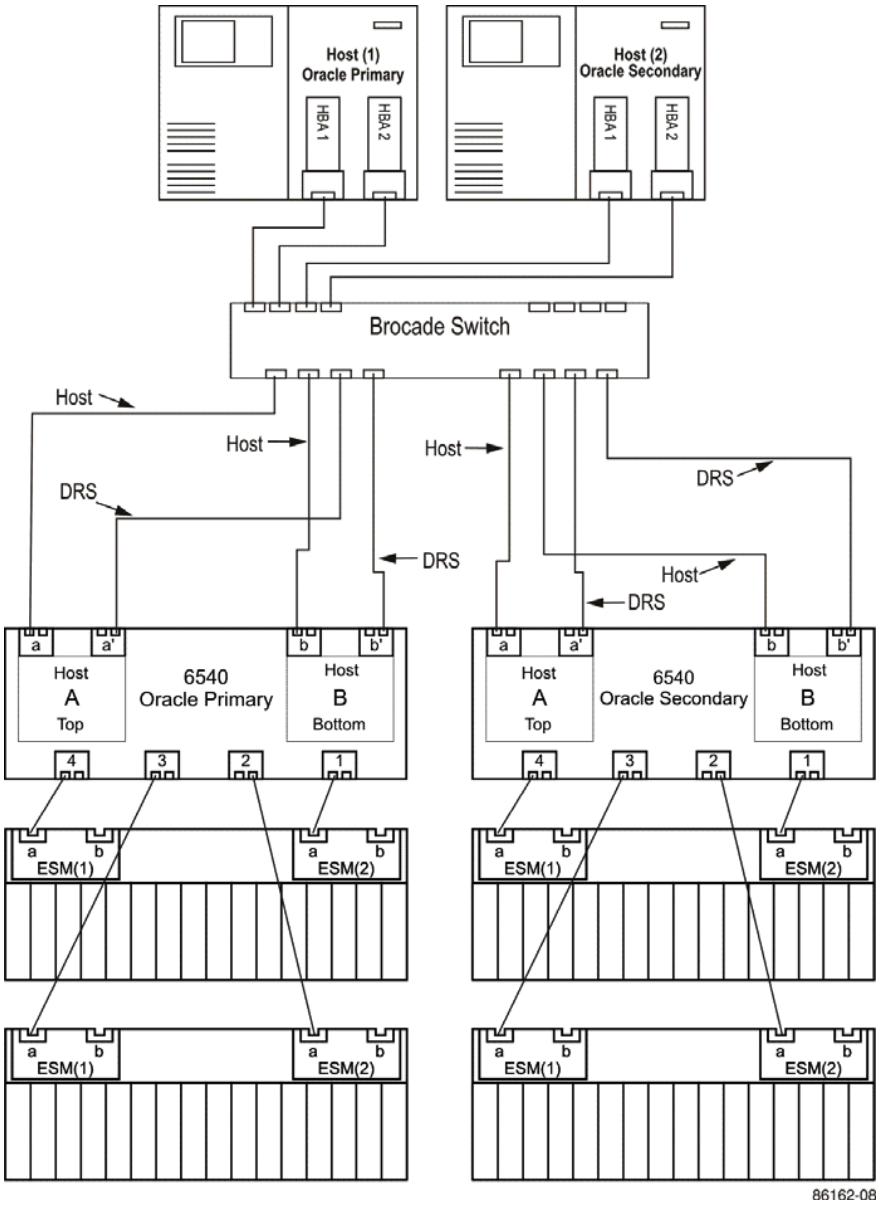

Figure 2 Sun Storage 6580/6780 Reference Architecture

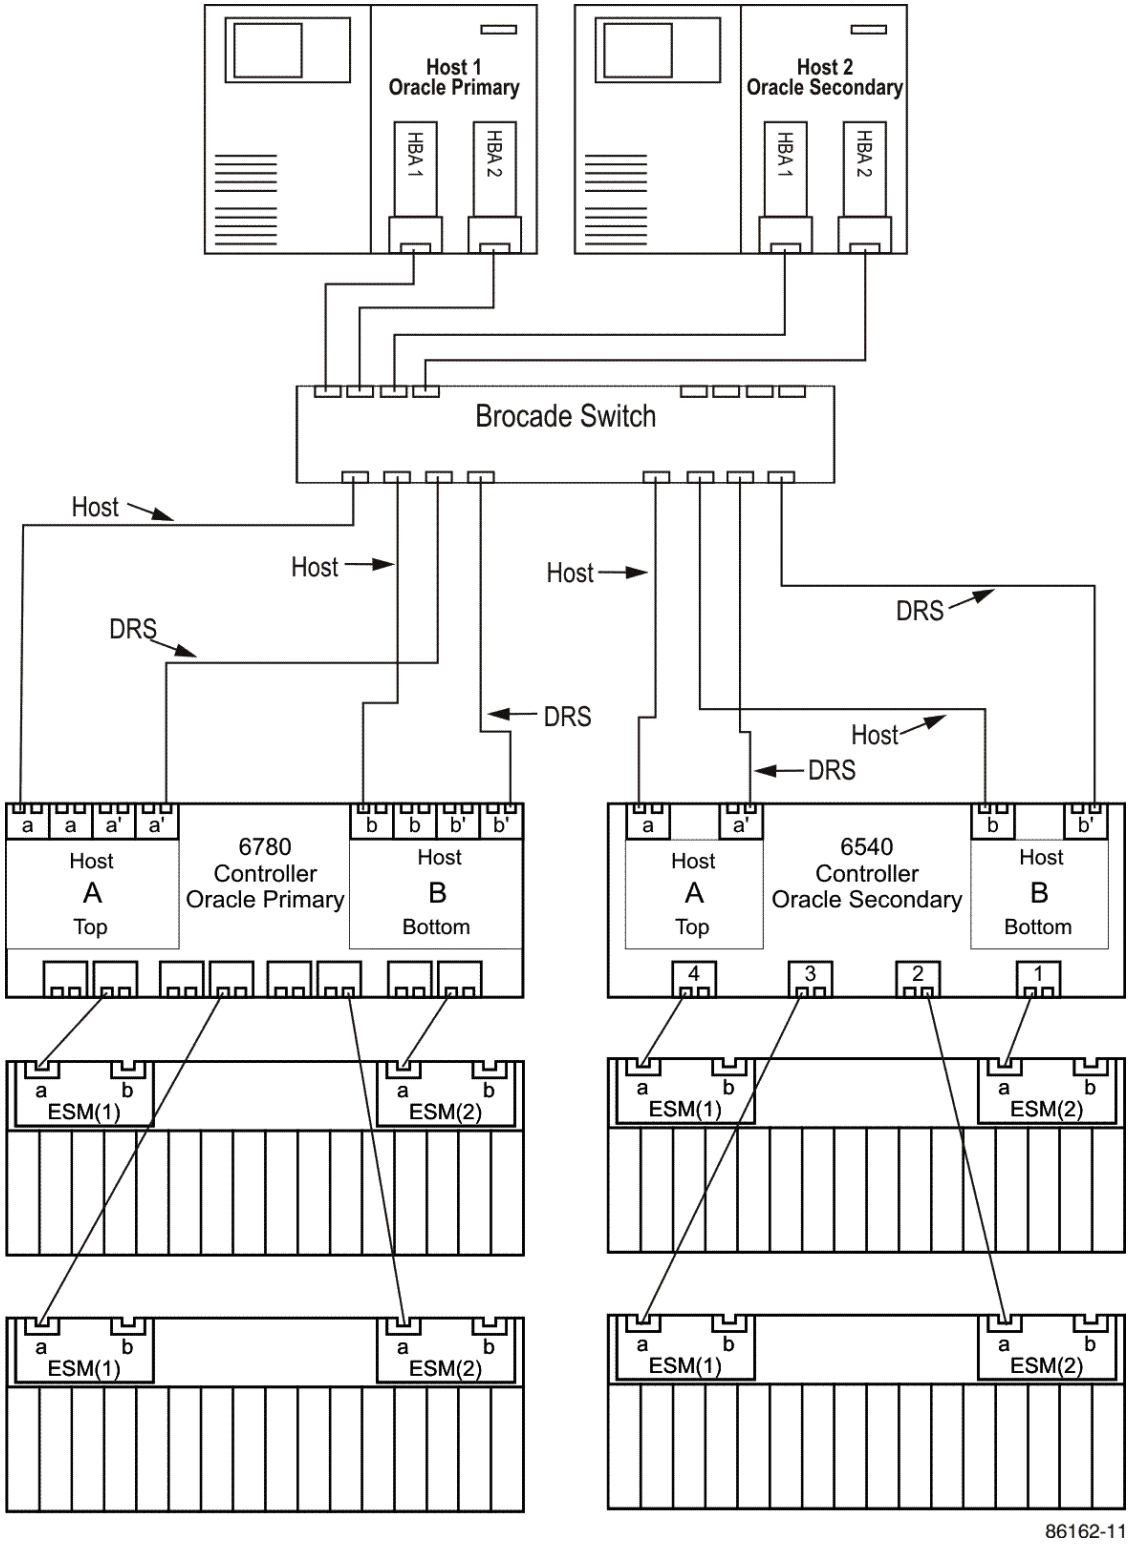

Volume Configuration

Figure 3 shows the volume configuration in the tests of Sun Storage Common Array Manager licensable premium features. Both reference architectures use the same volume configuration.

Figure 2 Sun Storage 6580/6780 Reference Architecture

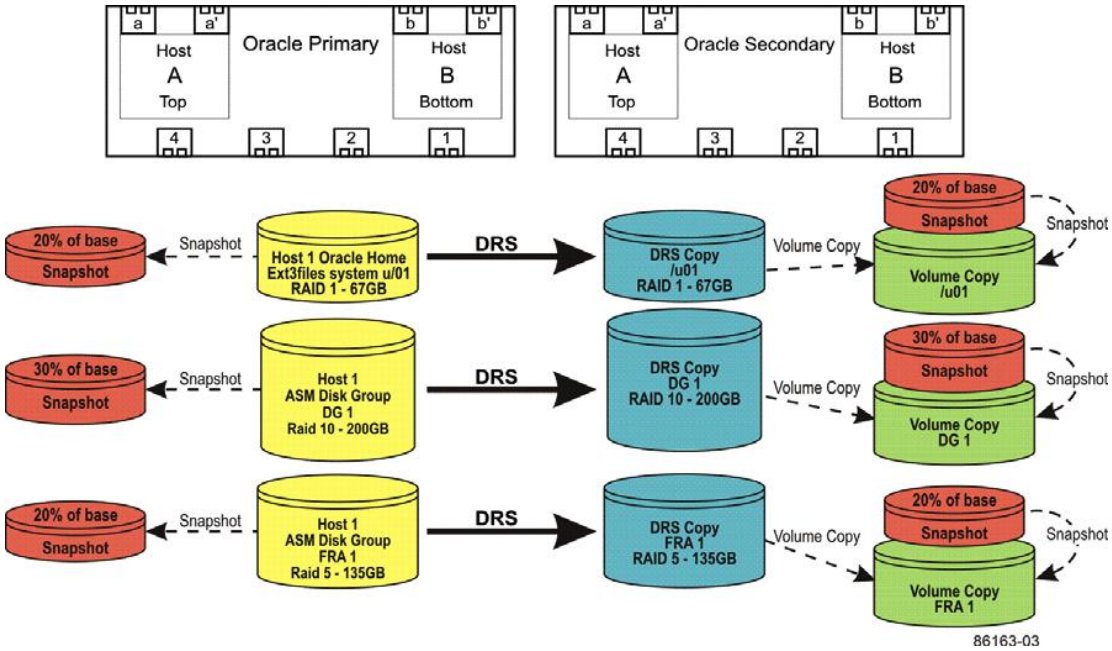

Setting Up the Licensable Premium Features

After your system environment is set up, your next step is to enable the DRS licensable feature. Licensable features must be enabled on all arrays separately. You must enable a licensable feature only once for the feature to become fully enabled throughout each storage array. To enable the licensable feature, obtain the Feature Key file from Oracle for each array.

Enabling a Licensable Premium Feature

1. Using the Sun Storage Common Array Manager browser interface, open the **Licensable Feature Summary** page by expanding the **Administration** folder under each storage array and click **Licensing**.
2. In the **Available Features** table, click the **Add License**... button.
3. In the **Add License** page, click the **Browse...** button and navigate to the *.KEY file that contains the activation key for the licensable feature. When Oracle sent the *.KEY file, you specified its file name and location.
4. Select the file, and then click **Open**.
5. Back in the **Add License** page, click the **Enable** button.

Setting Up Data Replicator Software

DRS lets you create an identical copy of the source volumes on a second storage array or on a remote storage array. DRS requires a physical hardware connection between two storage arrays. Connect the upper host ports to the last host ports on the storage array using a SAN switch infrastructure. Both controllers require the following connections:

- Connect the last host port of controller A on the target storage array to the last host port of controller A on the source storage array.
- Connect the last host port of controller B on the target storage array to the last host port of controller B on the source storage array.

DRS disables any host port connections to the last port when DRS is initialized.

Detailed Example Using This Method

After you set up the DRS premium feature, log on to Oracle to complete the following procedures. This section gives a general overview of the work that you must do in Oracle. The following sections give the detailed steps that you must perform in sequence using Sun Storage Common Array Manager and Oracle. In either reference architecture, the source database, residing on HOST1, consisted of two components:

- An Oracle 11g Automatic Storage Management (ASM) database, named pfdb1. The database was set up to use the Flash Recovery Area (FRA) in the database.
- An Oracle ASM instance running, named +ASM.

All volumes on HOST1 were set up to use DRS. The DRS copies provided all of the files needed to start the database after restoring or surfacing the data on another server. The target server, HOST2, was set up to use the mapped DRS drives to perform the recovery process.

On your system, configure the source database to run with archive logging, and then start the DRS process between the two storage arrays. Set up a process to stop the mirroring process on the ASM diskgroup DG1 (Oracle datafiles) each Sunday night. This weekly stop time provides the base to start the recovery process. The DRS process continues to run on the ASM diskgroup, FRA, which contains all of the archive log files generated by the source database. If data corruption occurs during the week, stop the DRS process on the FRA1 diskgroup, and then map all of the mirrored volumes mapped on HOST2. Recover the Oracle Database 11g to a point in time prior to when the corruption occurred.

Setting Up the Servers

1. Set up the primary server and the storage array. After the primary server and the storage array are running in an optimal state, validate the following components:

    - Databases are optimal and running
    - Primary storage is optimal
    - Secondary storage is optimal
    - DRS premium feature is enabled and the DRS links are connected

2. Set up the secondary server. Let the secondary server remain idle until it is needed to perform the flashback process. Do not map any volumes to the secondary server yet.

3. Using Sun Storage Common Array Manager, initiate an DRS copy by creating a new replication set. This is done using the **Create New Replication Set** wizard. The wizard is launched by clicking the **New...** button in the **Replication Sets** table found in the **Replication Set Summary** page. The **Replication Set Summary** page is shown by clicking on **Replication Sets** under the primary array. The target volume must already exist on the secondary array and be greater than or equal in size to the source volume on the primary array.

   The DRS copy should include the following components:

   - Source volumes /u01 (Oracle home, Oracle binary files, and filesystem database)
   - DG1 (ASM database disk group)
   - FRA1 (Flash Recovery ASM disk group)

NOTE: The source and target volumes must be owned by the same controller on their respective arrays in order to create a replication set relationship between them.
For example, if controller A on the primary array owns the source volume, controller A must also own the target volume on the secondary array.

The test case used the following **Create New Replication Set** wizard properties:

- Mode: **Asynchronous**
- ***Add to write consistency group*** checkbox enabled.
- Replication priority: **Highest**
- Resynchronization method: **Automatic**

4. Before preparing the source database, wait for the DRS volume copy to complete. Using the Common Array Manager browser interface, notice the **Syncronization Progress** field of the **Replication Sets** table in the **Replication Set Summary** page. The DRS volume copy has completed when the value of this field for the new replication set created in step 3 changes from "Synchronization In Progress" to "Replicating" which may take some time depending on the size of the volume.

Preparing the Source Database

The source database, an instance of Oracle Database 11g using Oracle ASM, resides on HOST1 and is named pfdb1. Perform the following steps on HOST1.

1. Set up the source database to run in archive log mode. In the test case, the archive log destination references the ASM diskgroup FRA1.

2. Create two tables, scott.tb1 and scott.tb2. As a test, populate the tables with hundreds of rows of data.

3. Stop the DRS process on the ASM diskgroup DG1. The DRS process continues to run on the ASM diskgroup FRA. Make sure that synchronization is complete before you quiesce the database.

SQL> alter system switch logfile;

SQL> alter database begin backup;

SQL> alter system quiesce restricted;

SQL> select active_state from v$instance;

```
ACTIVE_STATE
----------
QUIESCED
```

4. From the Sun Storage Common Array Manager browser interface, click the replication set name associated with the DRS volume DG1 in the **Replication Sets** table on the **Replication Set Summary** page for either array . This will display the **Replication Set Details** page. Click the **Suspend** button. Back in the **Replication Set Summary** page, confirm that the **Synchronization Progress** field now displays the value "Suspended" for the replication set. If the replication set is not intended to be used again, an alternative to suspending the replication set is to simply delete it. Deleting a replication set will not alter the contents of either the source or target volumes.

5. Bring the source database back online.

SQL> alter system unquiesce;

SQL> alter database end backup;

```
Database altered.
```

A couple of active jobs were creating archive log files.

6. Issue a query to find out the current SCN and the current date and time before data deletion occurred. The query result will become the restore point.

SQL> select current_scn from v$database;

```
CURRENT_SCN
-----------
23907385
```

SQL> ! date

```
Thu Apr 10 16:50:13 CDT 2008
```

7. After discovering the restore point in the previous step, you can drop the scott.tb1 table and the scott.tb2 table.

SQL> drop table scott.tb1;

```
Table dropped.
```

SQL> drop table scott.tb2;

Table dropped.

8. Stop the DRS process on the ASM diskgroup FRA1.

   Stopping DRS on the ASM diskgroup lets the current archived log files roll forward the database from the mirrored DG1 disks. Make sure that synchronization is complete before you quiesce the database.

   SQL> alter system switch logfile;

   SQL> alter database begin backup;

   SQL> alter system quiesce restricted;

   SQL> select active_state from v$instance;

   ACTIVE_STATE

   QUIESCED

9. From the Sun Storage Common Array Manager browser interface, click the replication set name associated with the DRS volume FRA1 in the **Replication Sets** table on the **Replication Set Summary** page for either array. This will display the **Replication Set Details** page. Click the **Suspend** button. Back in the **Replication Set Summary** page, confirm that the **Synchronization Progress** field now displays the value "Suspended" for the replication set. If the replication set is not intended to be used again, an alternative to suspending the replication set is to simply delete it. Deleting a replication set will not alter the contents of either the source or target volumes.

10. Bring the source database back online.

    SQL> alter system unquiesce;

    SQL> alter database end backup;

    Database altered.

Preparing the Target

HOST2 is the target location where the source database, an instance of Oracle Database 11g using Oracle ASM, and named pfdb1, will surface. Perform the following steps on HOST2.

1. Using the Sun Storage Common Array Manager browser interface, map the target DRS volume to HOST2 and make the volume visible to the operating system.

   a. After expanding the secondary array's folder, click **Mappings** to display the **Mapping Summary** page.
   b. Click **New...** to launch the **Create New Mappings** wizard.
   c. Choose a specific LUN and make note of the number chosen.
   d. Run the /usr/sbin/mppBusRescan utility on HOST2 to have the operating system look for the newly mapped target.

e. Run the /opt/mpp/lsvdev command on HOST2 to confirm that the LUN chosen for the target DRS volume in step c above is shown next to the secondary array's name.

f. Make a mount point with mkdir /u01.

g. Mount the file system as /u01 using the device shown next to the LUN in step e.

2. Scan for ASM disks. Type the following command.

```
/etc/init.d/oracleasm scandisks
```

3. Start Oracle Cluster Services (OCS). Type the following command.

```
localconfig reset $ORACLE_HOME
```

4. If you have not already changed the environment settings to start the ASM instance of Oracle Database 11g and the instance of Oracle Database 11g, named pfdb1, on HOST2, change the environment settings now.

5. Because the database being recovered is using Oracle ASM, you must complete these two tasks before you start the recovery process.

   a. Locate the archive log files in Oracle ASM. To locate the archive log files in Oracle ASM, log in to ASMCMD, and then drill down to find the logs. In the following example from the test case, the drill down located three archive log files in Oracle ASM.

```
Export ORACLE_SID=+ASM
        Asmcmd
ASMCMD> cd +FRA1/PFDB1/ONLINELOG
ASMCMD> ls –al
        Type              Redund Striped Time              Sys Name
        ONLINELOG UNPROT FINE     APR1016:00:00 Y none => group_1.257.646402135
        ONLINELOG UNPROT FINE     APR1016:00:00 Y none => group_2.258.646402141
        ONLINELOG UNPROT FINE     APR1016:00:00 Y none => group_3.259.646402147
```

   b. Reference these logs during the roll forward recovery process.

   The process of recovering the database begins.

6. If you have not already done so, make the necessary environment settings to start the database.

In the following example from the test case, an inconsistent with file error appears because the control files were multiplexed in both diskgroup DG1 and diskgroup FRA1. Also, the FRA1 diskgroup has a later version of the control file compared to the control file in DG1.

```
export ORACLE_SID=pfdb1

oracle@HOST2$ sqlplus

SQL*Plus: Release 11.1.0.6.0 - Production on Thu Apr 10 17:42:47 2008

Enter your user name as follows.

    / as sysdba

    Connected to an idle instance.

    SQL> startup mount

        ORACLE instance started.

        Total System Global Area    1686925312 bytes
        Fixed Size                  2145024 bytes
        Variable Size               1207960832 bytes
        Database Buffers            469762048 bytes
        Redo Buffers                7057408 bytes
        ORA-00214: control file '+FRA1/pfdb1/controlfile/current.256.646402127' version
        40686 inconsistent with file '+DG1/pfdb1/controlfile/current.260.646402127' version
        40661
```

7. To continue, you must delete the control file entry in the spfile that points to the DG1 diskgroup, shut down the database, and then start the database in mount mode.

```
    SQL> show parameter control_files
        NAME            TYPE        VALUE

        control_files   string      +DG1/pfdb1/controlfile/current
                                    .260.646402127, +FRA1/pfdb1/
                                    controlfile/current.256.646402127


    SQL> alter system set control_files ='+FRA1/pfdb1/controlfile/current.256.646402127'
    scope=spfile;

        System altered.

    SQL> shutdown abort

        ORACLE instance shut down.

    SQL> startup mount

        ORACLE instance started.
        Total System Global Area    1686925312 bytes
        Fixed Size                  2145024 bytes
        Variable Size               1207960832 bytes
        Database Buffers            469762048 bytes
        Redo Buffers                7057408 bytes
        Database mounted.
```

The recover command begins the roll forward process. In the test case, the SCN supplied the restore point. You also can use a date and time for the restore point, or, as a database administrator, you can CANCEL when you observe that the process is at the right point.

8. When prompted for the log file name, supply the name of the first log file that you located in ASMCMD in step 7.

IMPORTANT You must disable Flashback before you proceed with the recover process.
The Flashback logs impact performance too much to leave Flashback on.

To check to see if Flashback is on, enter the following command.
Sql> Select flashback_on from v$database;
If the returned value is YES, enter the following command.
Sql> alter database flashback off;

9. Recover database to the specified SCN.

```
SQL> recover database until change 23907385 using backup controlfile;

ORA-00279: change 23903106 generated at 04/10/2008 11:13:06 needed for thread 1
ORA-00289: suggestion : +FRA1
ORA-00280: change 23903106 for thread 1 is in sequence #1
Specify log: {<RET>=suggested | filename | AUTO | CANCEL}
+FRA1/PFDB1/ONLINELOG/ group_1.257.646402135

Log applied
Recovery complete
```

Because the first log file contained the recovery point in time, the recovery process applies the data from this archive log and indicates that the recovery process is complete.

10. Open the database with the resetlogs command. The resetlogs command completes the recovery process.

```
SQL> alter database open resetlogs;
```

11. To verify that the recovery process worked, issue a query against the table, scott.tb1, that you dropped from the source database. If the scott.tb1 table is present and has the correct number of rows, the recovery process is a success. In the following example from the test case, scott.tb1 is present and contains the correct number of rows.

```
SQL> select count(1) from scott.tb1;

COUNT(1)
-------
289
```

12. After the recovery process is complete, you can use Oracle Roll Forward Recovery to roll forward the data in either of the following ways:

    - Retrieve the data from this cloned database and update the source database.

- Reverse the DRS process and sync the source database drives from the cloned database drives.

Keep This Document Handy

When a database must be rolled back to a point in time before corruption occurred, new data from the point of corruption to the present is temporarily lost. The restored data is brought current by rolling forwards as many logs as necessary. This procedure differs from standard restorations because it employs a secondary server, allowing the primary database to remain in production throughout the procedure. Letting the primary database remain in production gives you the flexibility to complete a roll forward restore whenever you want. You will want to use this method from now on, because it can be implemented without pre-scheduling. Corruptions can be repaired in real-time and you can maintain your data in good shape at all times. Because you might have to perform a roll forward recovery on short notice, keep this document in a close but safe place for quick access.

If you prefer to have professional consultants service your roll forward recovery process, see "Contact Information."

For documentation on other scenarios that an Oracle Database administrator might encounter, see "Appendix A: References."

Appendix A: References

This document is one of seven detailed examples that explain how to complete common but important tasks often required from an Oracle database administrator. Refer to the other documents as needed.

• Cloning an Oracle Database to the Same Server Using Snapshot and Volume Copy

• Cloning an Oracle Database Using Data Replicator Software

• Forward Recovery of an Oracle Database Using Data Replicator Software

• Migrating to Oracle 11g Using Data Replicator Software with Transportable Tablespaces

• Recovering from Catastrophic Failures Using Data Replicator Software for Data Replication

• Safely Upgrading an Oracle Database Using Data Replicator Software

• Selective Restores Using Data Replicator Software with Oracle Flashback Database

Oracle is committed to developing practices and products that help protect the environment

Forward Recovery of an Oracle Database
Using Data Replicator Software
August 2010

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
oracle.com

SOFTWARE. HARDWARE. COMPLETE.