ORACLE®

Your. Open. World.

Beat-up Your Oracle Data Guard Standby with Oracle Real Application Testing – It’s Payback Time!

Larry M. Carpenter
Distinguished Product Manager, Oracle USA
Sreekanth Chintala
Senior Database Engineer, Dell Inc.

Agenda

- Snapshot Standby Databases
  - What are they?
  - What are their requirements?
  - What can they be used for?
- Real Application Testing
  - What is it?
  - How does it work with Snapshot Standby Databases?
- Dell Success Story.
- Using Database Replay on a Snapshot Standby
  - Capture a workload on the Primary database
  - Replay that workload on a Snapshot Standby Database

Oracle Data Guard

Best Protection at Lowest Cost

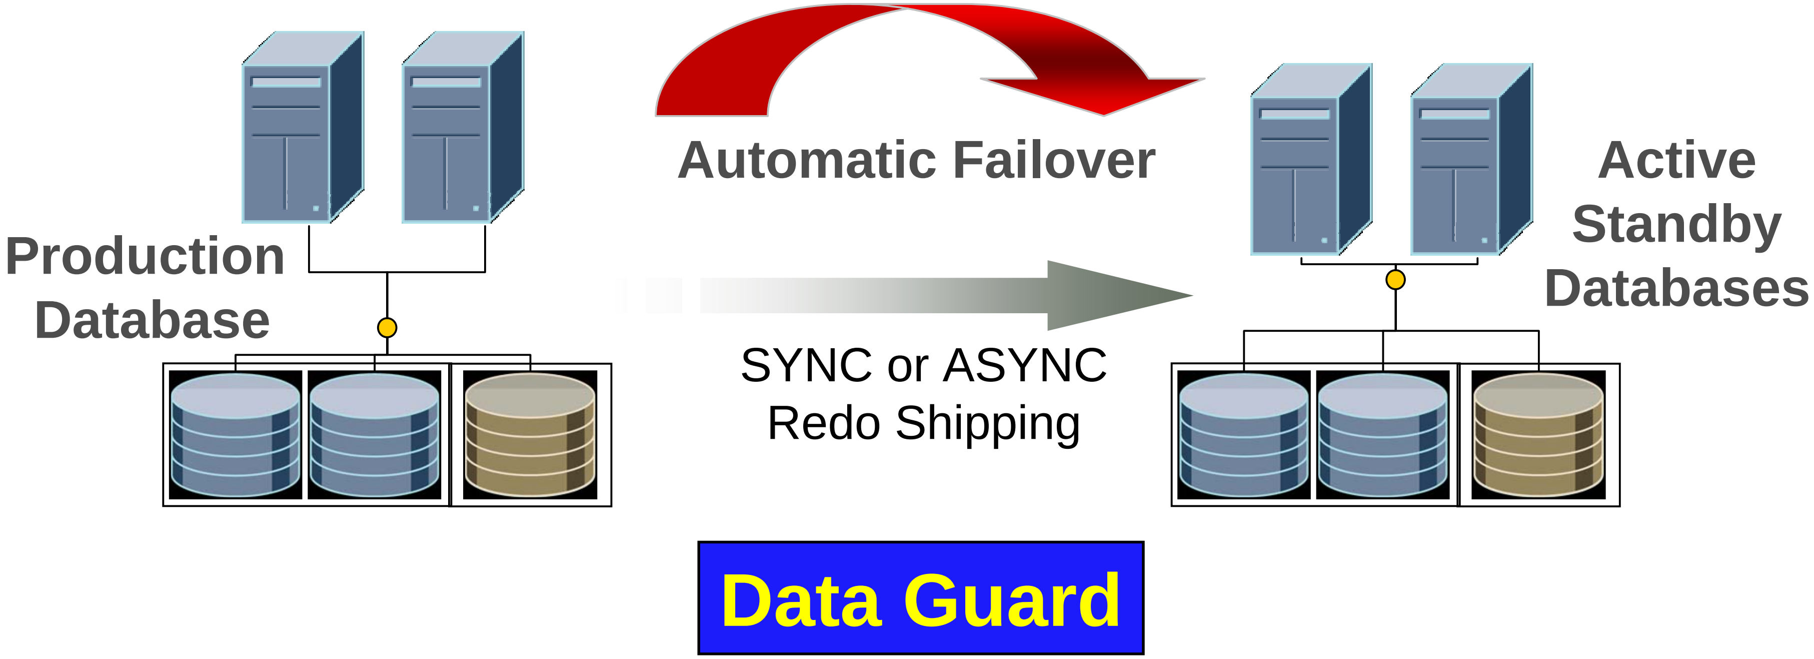

Unlocking the Value of Standby Databases

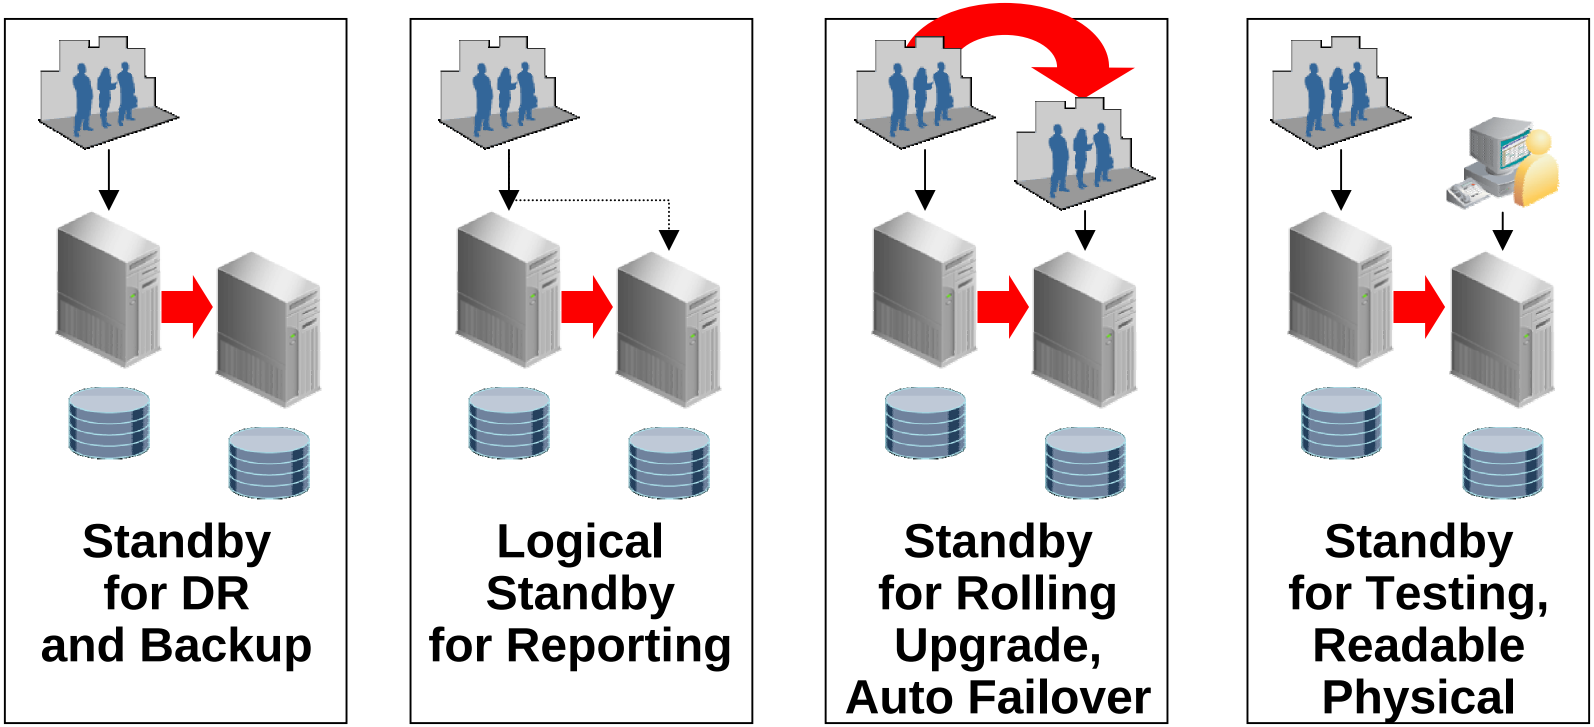

Standby
for DR
and Backup

Logical
Standby
for Reporting

Standby
for Rolling
Upgrade,
Auto Failover

Standby
for Testing,
Readable
Physical

ORACLE 8i INTERNET

ORACLE 9i INTERNET

ORACLE DATABASE 10g

ORACLE DATABASE 11g

Why Snapshot Standby Databases?

- Combined with Real Application Testing
  - Provides a simple way to test and maintain protection
- Eliminate guesswork.
  - Performance test results are realistic and reliable
- Know the impact.
  - Realistic system resource consumption on production data during performance testing
- Query tuning
  - No longer dependent on best execution plan from old data
- Change assurance.
  - Introduce changes with confidence

Ideal Complement for Real Application Testing

- Physical standby databases can take a beating and come back for more!
  - Snapshot Standby databases allow real world testing without compromising data protection.
  - Real Application Testing with Database Replay provides that real world testing
  - Since Physical standby databases are an exact replica of your Production database the workload replay uses the same data set.
  - No need for extra storage
  - Maintain Disaster Protection while testing

Snapshot Standby
Oracle Database 11*g*

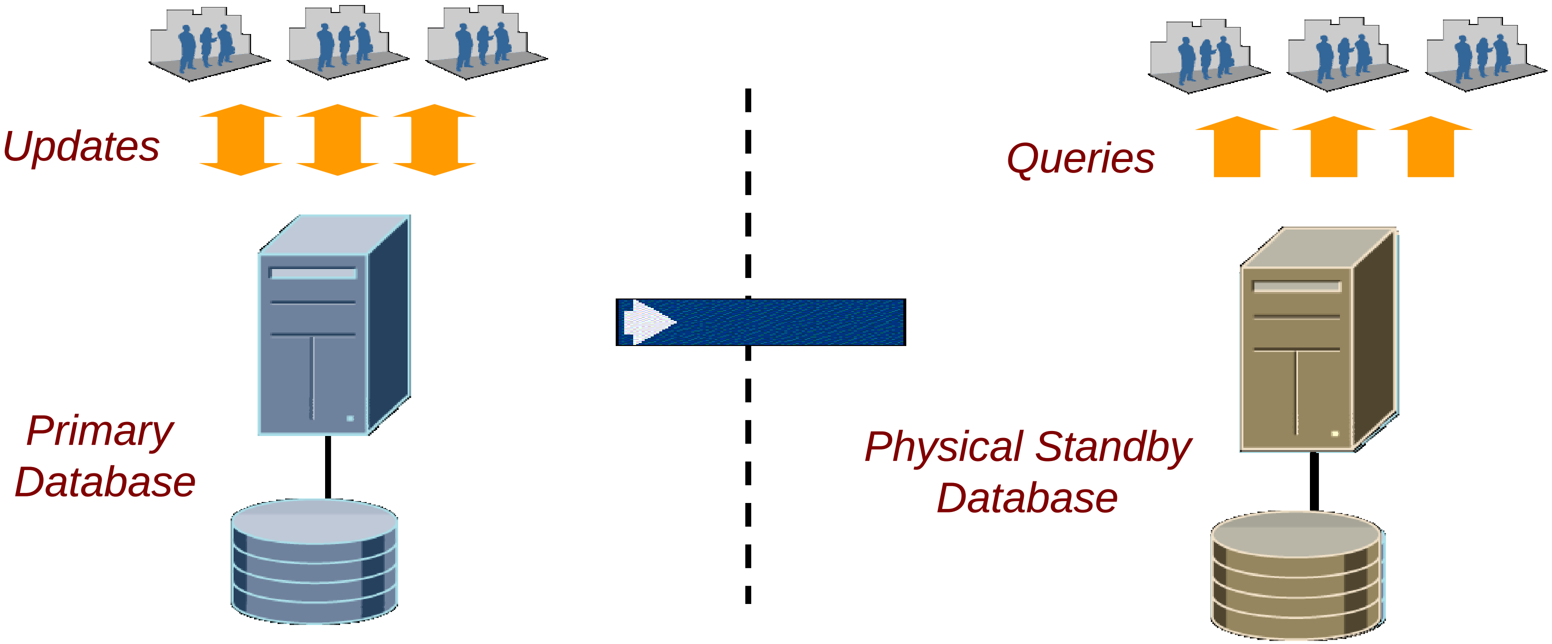

Snapshot Standby – Simple

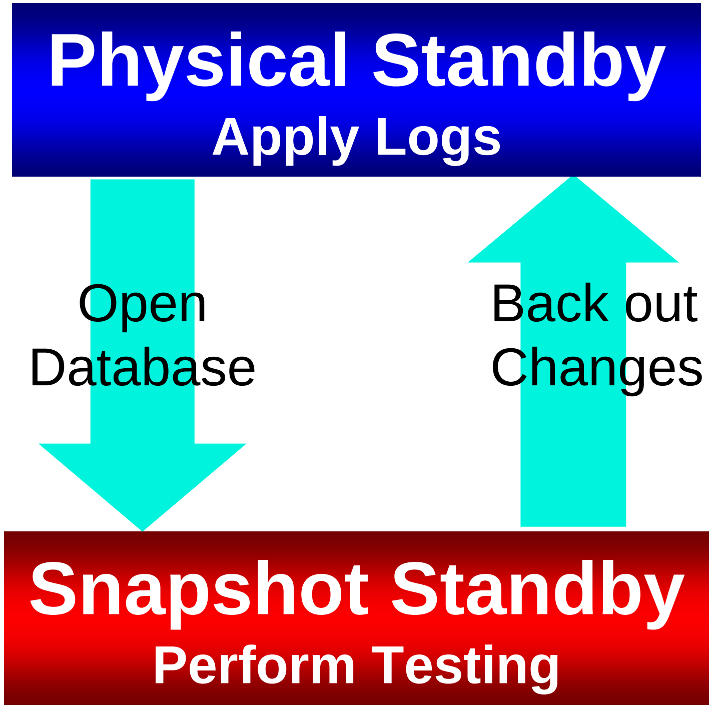

Continuous Redo Shipping while in Snapshot mode

- Convert Physical Standby to Snapshot Standby and open for writes by testing applications
  - 2 Steps vs 15 in 10.2
- Discard testing writes and catch-up to primary by applying logs
- Preserves zero data loss
  - But no real time query or automatic failover
- No idle resources
- Similar to storage snapshots, but:
  - Provides DR at the same time
  - Uses single copy of storage

Snapshot Standby Requirements

- Physical Standby
  - One or more Physical standby databases can be Snapshot standby databases at any time
- Flashback Database
  - Necessary to undo all of your testing
  - See Note 565535.1 Flashback Database Best Practices & Performance
- Flash Recovery Area
  - Required for Flashback Database
- More details on setting up these requirements with Data Guard in the Appendix of this talk.
  - Download the talk!

Real Application Testing

- Database Replay
  - Database and System administrators can test
    - Database upgrades, patches, parameter changes
    - Configuration changes
      - Conversion from a single instance to RAC, ASM, etc.
    - Storage, network, interconnect changes
    - Operating system, hardware migrations, patches, upgrades, parameter changes
- SQL Performance Analyzer
  - Test changes as with Replay and in addition:
    - Schema changes such as adding new indexes, partitioning or materialized views
    - Gathering optimizer statistics. SQL tuning actions, for example, creating SQL profiles

Real Application Testing and Data Guard

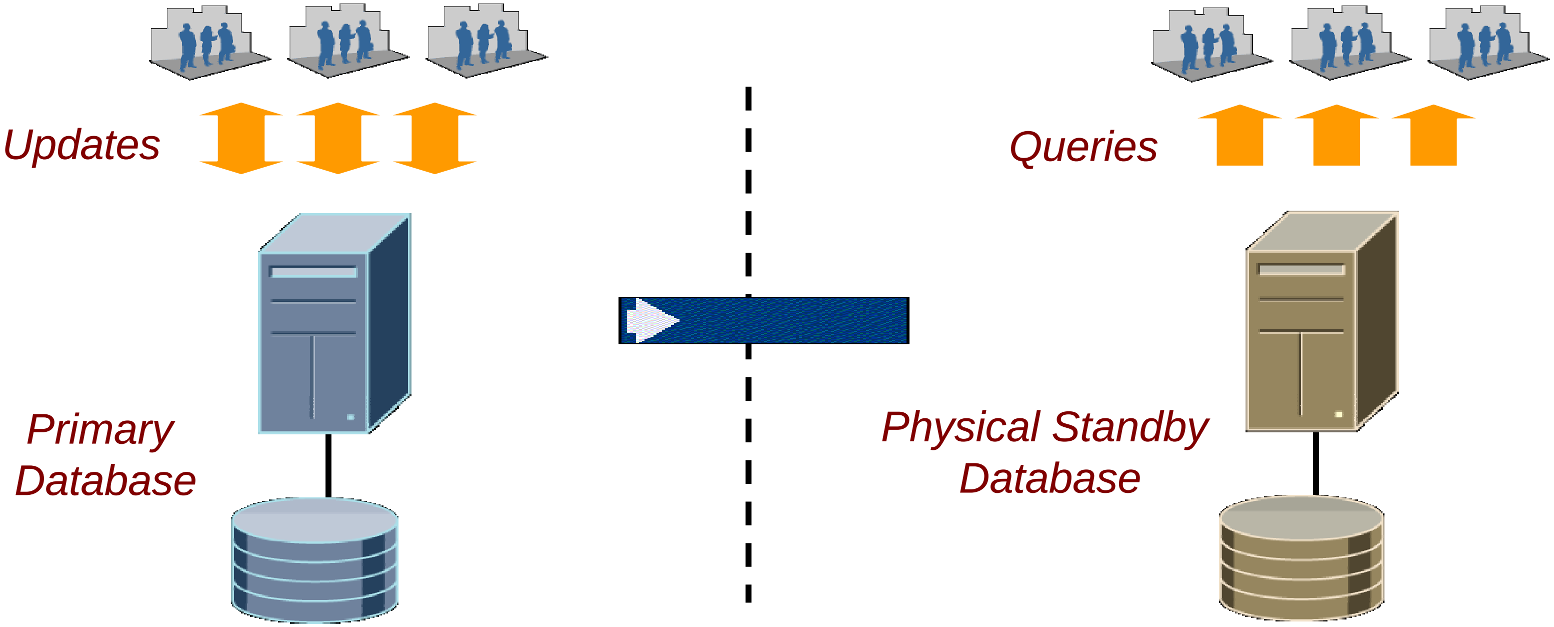

Real Application Testing and Data Guard

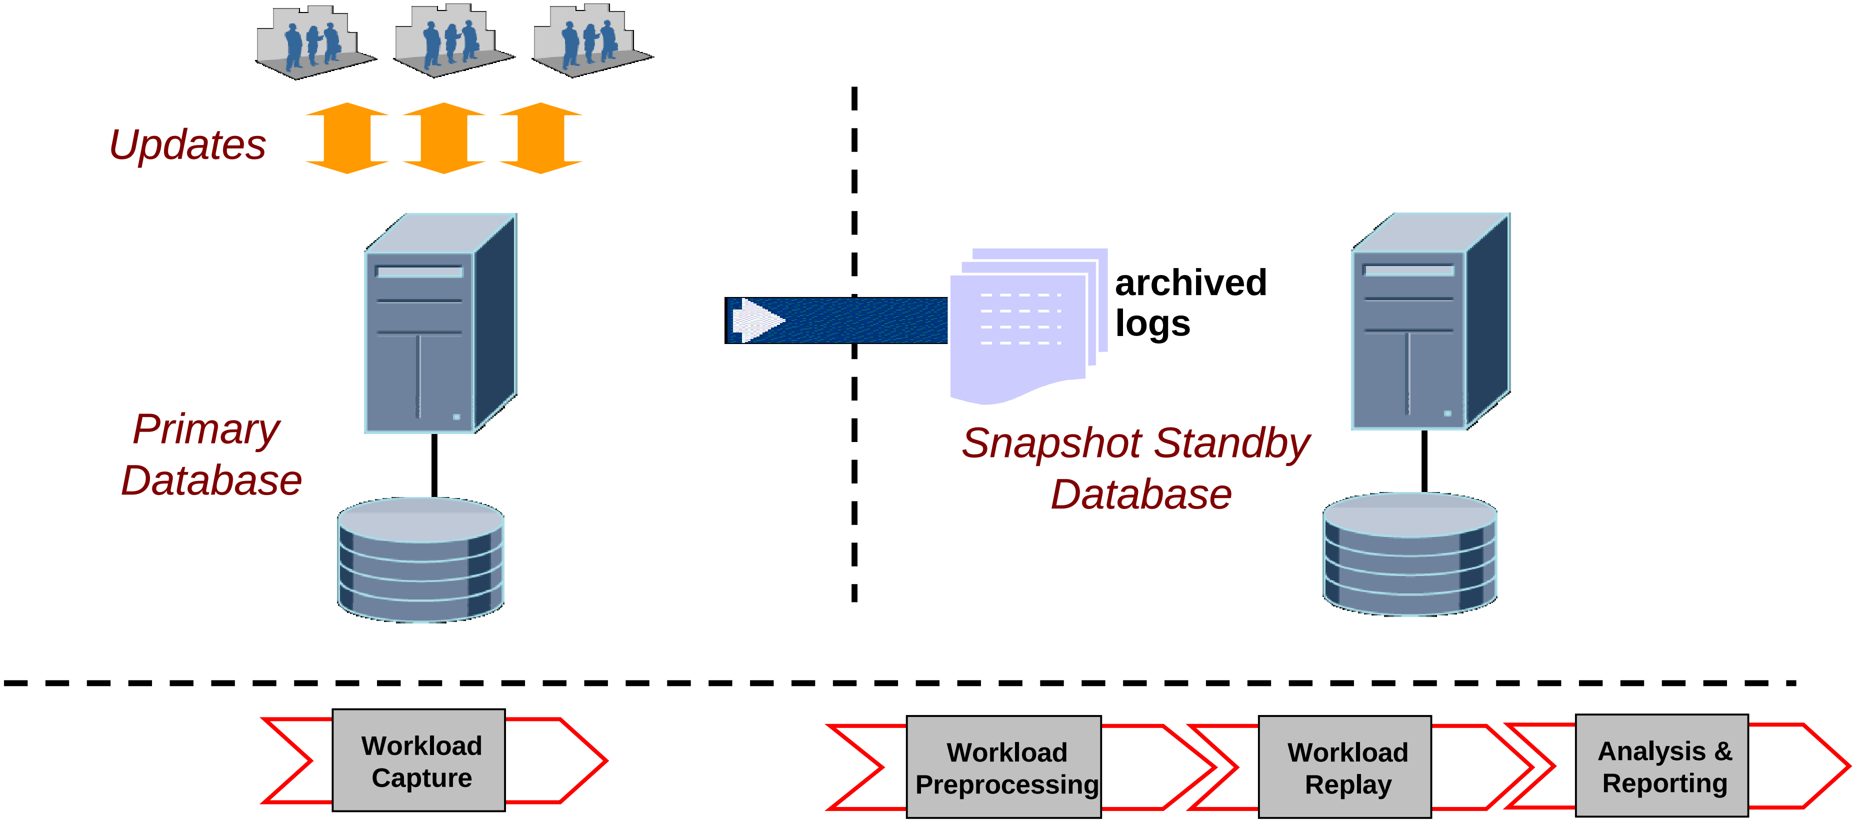

Real Application Testing and Data Guard

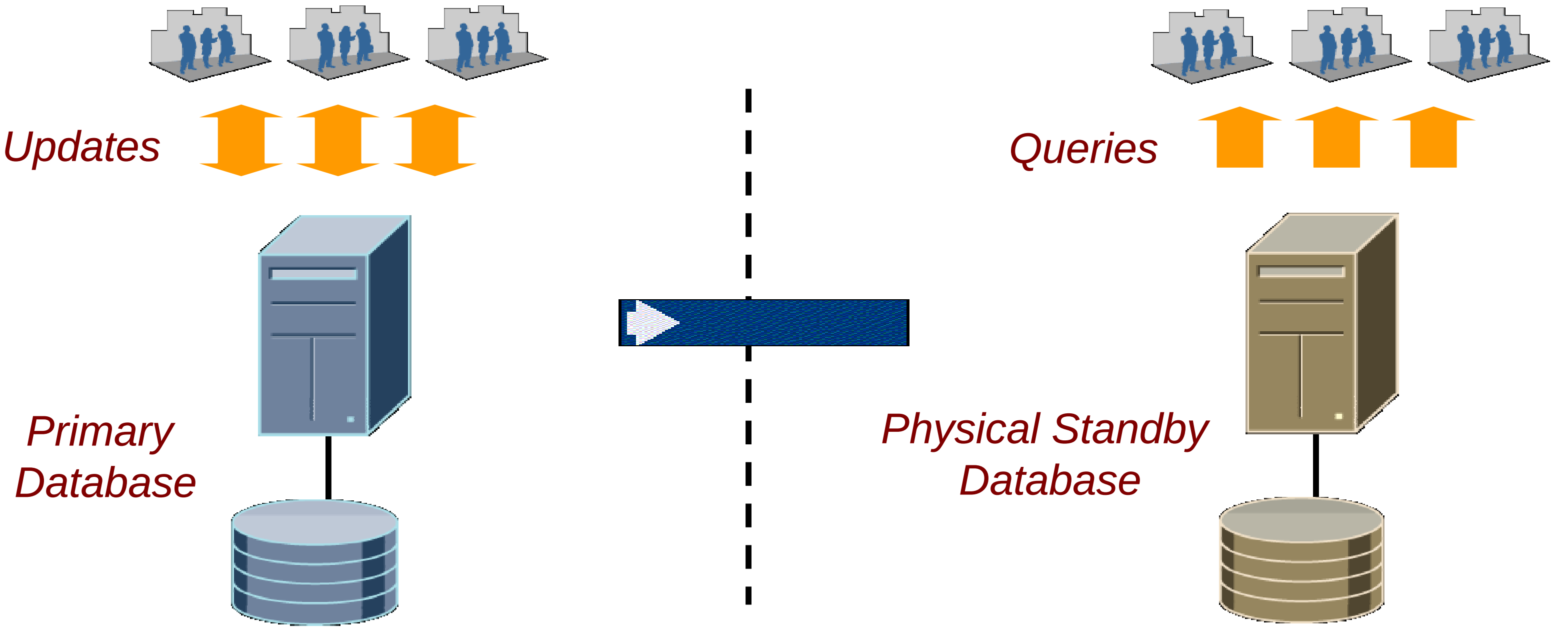

Introducing Sreekanth Cintala

Senior Database Engineer, Dell Inc.
And a long time Data Guard user

Dell Inc. slides can be found in S298770_2.pdf

Testing

- Everybody performs testing
  - At least they had better be testing!
- Requires time and effort in populating test environments frequently
- Needs extra storage
- Workloads do not mimic production
  - Rarely catch production problems
- Even simple tests like checking a new index can be very time consuming!

Using Real Application Testing and Data Guard Snapshot Standby

Production Setup	Test on Snapshot Standby
Non-ASM storage	ASM storage
No Flash Recovery Area	Uses Flash Recovery Area
Flashback Database Not Enabled	Flashback Database Enabled
Lower O/S Version	Higher O/S Version
Old Hardware (Disks/System)	New Hardware (Disks/System)
Simple Indexes	Complex/More Indexes
Parameters	Changed Parameters
No Partitioning	New Partitioning
No Compression	Compression Enabled
9.2.0.8, 10.2.0.4, 11.1.0.6	11.1.0.6+ (*Requires that Compatibility is not changed*)

Configuring Database Control on the Primary Database

- If necessary

```
[Chicago] emca -config dbcontrol db -repos recreate

STARTED EMCA at Sep 2, 2008 4:16:55 PM
EM Configuration Assistant, Version 11.1.0.5.0 Production
Copyright (c) 2003, 2005, Oracle.  All rights reserved.

Enter the following information:
.
.
Do you wish to continue? [yes(Y)/no(N)]: Y
.
.
INFO: >>>>>>>>>>> The Database Control URL is
  https://stadu67.us.oracle.com:1158/em <<<<<<<<<<<

Enterprise Manager configuration completed successfully
FINISHED EMCA at Sep 2, 2008 4:23:32 PM
```

Configuring for Database Replay

- Create a capture directory

```
[Chicago] pwd
/scratch/oracle/Replay/
[Chicago] mkdir Capture
```

- Setup your workload files and scripts
  - We used the Oracle By Example Scripts for Database Replay
- Startup Database Control if necessary

```
[Chicago] emctl start dbconsole
Starting Oracle Enterprise Manager 11g
  Database Control ............ started.
Logs are generated in directory
  /scratch/OracleHomes/OraHome111/stadu67.us.
  oracle.com_Chicago/sysman/log
```

Convert Your Physical Standby to a Snapshot Standby

- Convert before you start your workload generation
  - Allows you to maintain the same starting point.

```
DGMGRL> convert database boston to snapshot standby;
Converting database "boston" to a Snapshot Standby
  database, please wait...
Database "boston" converted successfully
DGMGRL> show configuration
...
  Databases:
    chicago - Primary database
    boston  - Snapshot standby database
...
```

- Requires extra Flash Recovery Area space as Flashback logs are not recycled when a Guaranteed Restore Point is in effect

Redo is still being received

- On the Standby database

```
SQL> select client_process,process,sequence#,
           status from v$managed_standby
     where client_process='LGWR'
     or process like 'MRP%';

CLIENT_P PROCESS    SEQUENCE# STATUS
-------- --------- ---------- ------------
LGWR     RFS              323 IDLE

1 rows selected.
```

Configure Database Control on the Snapshot Standby

- Can be done in parallel with the capture

```
[Boston] emca -config dbcontrol db -repos recreate
STARTED EMCA at Sep 2, 2008 5:11:25 PM
EM Configuration Assistant, Version 11.1.0.5.0
  Production
Copyright (c) 2003, 2005, Oracle.  All rights
  reserved.

Enter the following information:
...
You have specified the following settings
...
Do you wish to continue? [yes(Y)/no(N)]: Y
INFO: >>>>>>>>>>> The Database Control URL is
  https://stadu67.us.oracle.com:5501/em <<<<<<<<<<<
Enterprise Manager configuration completed
  successfully
FINISHED EMCA at Sep 2, 2008 5:26:34 PM
```

Capturing a Workload

- Log in to Database Control of your Primary database
  - https://stadu67.us.oracle.com:**1158**/em/

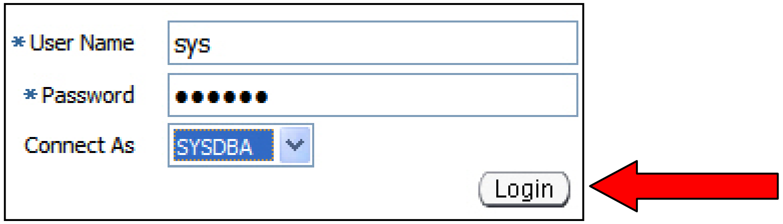

* User Name: sys

* Password: ●●●●●●

Connect As: SYSDBA

- Select the “Software and Support” link at the top

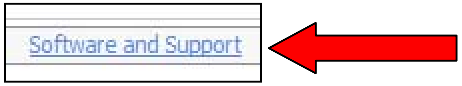

- Under “Real Application Testing”
  - Select “Database Replay”

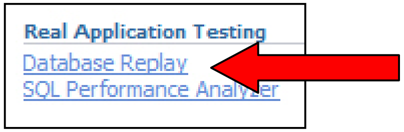

Capturing a Workload

- On the “Capture Workload task” line
  - Click the “Go to Task” icon.

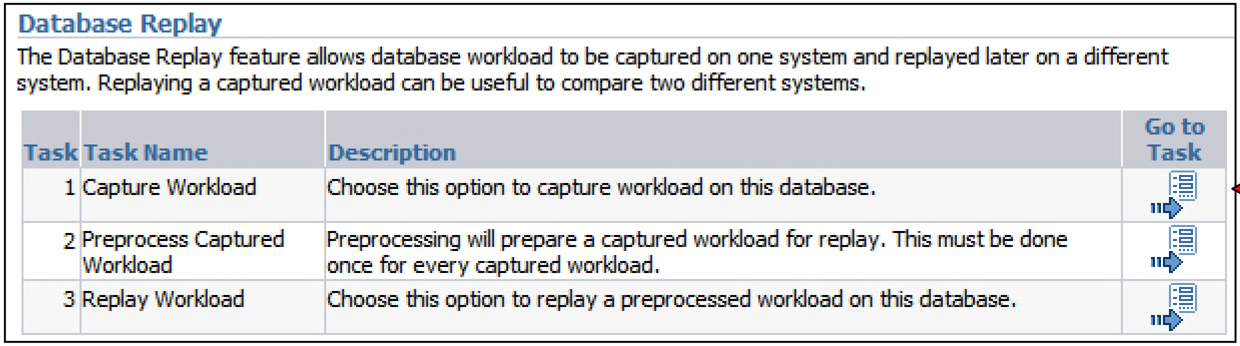

Database Replay

The Database Replay feature allows database workload to be captured on one system and replayed later on a different system. Replaying a captured workload can be useful to compare two different systems.

Task	Task Name	Description	Go to Task
1	Capture Workload	Choose this option to capture workload on this database.	➡
2	Preprocess Captured Workload	Preprocessing will prepare a captured workload for replay. This must be done once for every captured workload.	➡
3	Replay Workload	Choose this option to replay a preprocessed workload on this database.	➡

- Read and acknowledge the prerequisites.

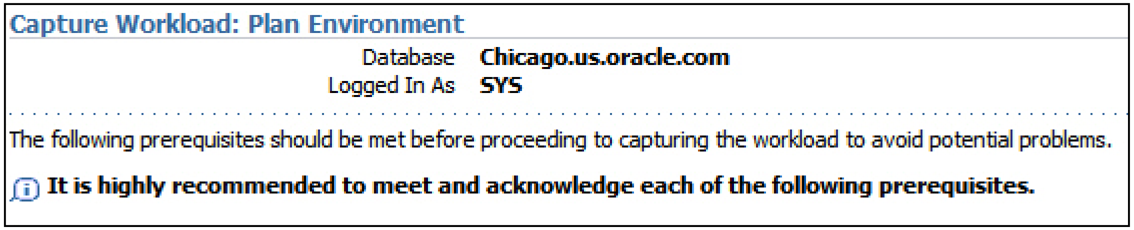

Capture Workload: Plan Environment

Database **Chicago.us.oracle.com**
Logged In As **SYS**

The following prerequisites should be met before proceeding to capturing the workload to avoid potential problems.

ⓘ **It is highly recommended to meet and acknowledge each of the following prerequisites.**

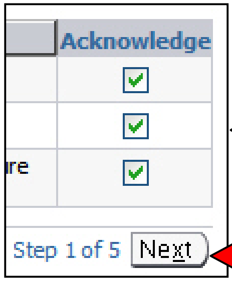

	Acknowledge
	☑
	☑
re	☑

Capturing a Workload

- Choose to restart the Primary or not and click next.

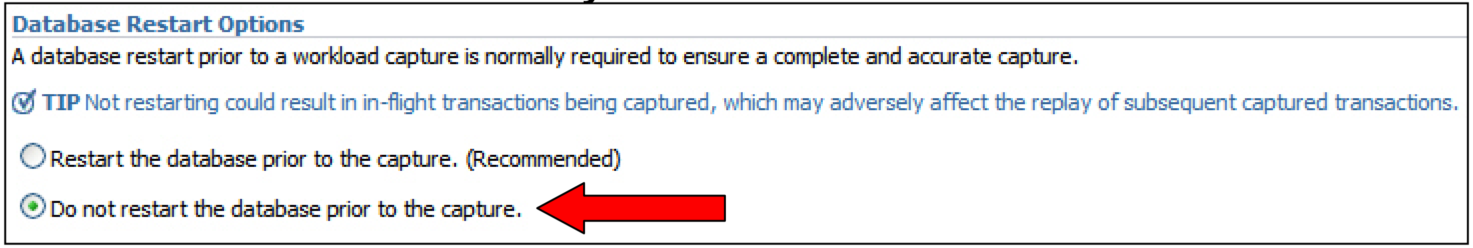

Database Restart Options

A database restart prior to a workload capture is normally required to ensure a complete and accurate capture.

TIP Not restarting could result in in-flight transactions being captured, which may adversely affect the replay of subsequent captured transactions.

○ Restart the database prior to the capture. (Recommended)

● Do not restart the database prior to the capture.

- Create the directory object for your Replay directory

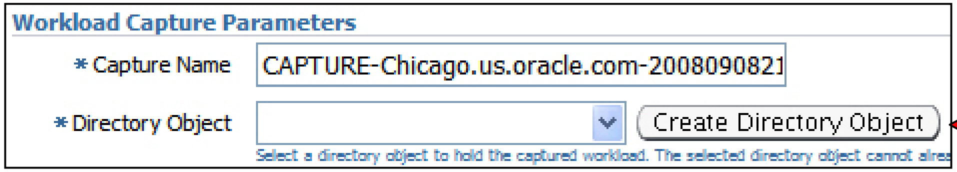

Workload Capture Parameters

* Capture Name: CAPTURE-Chicago.us.oracle.com-200809082

* Directory Object:

Select a directory object to hold the captured workload. The selected directory object cannot alre

  - Test the file system if required using the button provided

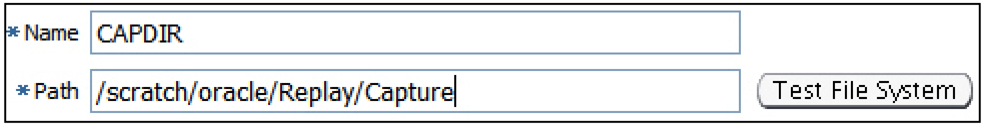

* Name: CAPDIR

* Path: /scratch/oracle/Replay/Capture

  - Each Workload Capture requires its own directory.
  - Click the OK Button

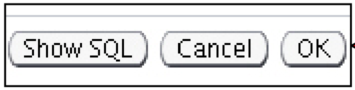

Capturing a Workload

- Once the directory object is created click ‘Next’

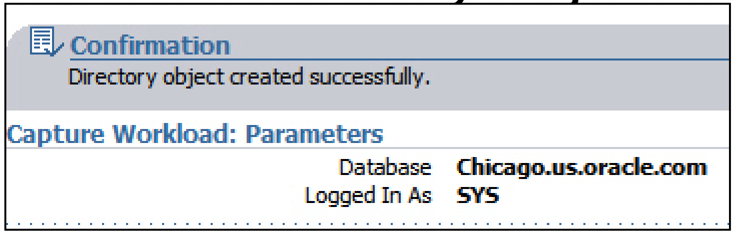

Confirmation
Directory object created successfully.

Capture Workload: Parameters

Database	**Chicago.us.oracle.com**
Logged In As	**SYS**

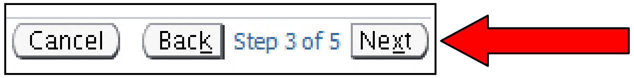

- Set a schedule for your capture (Start and Stop)

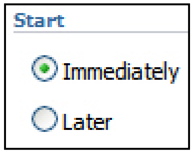

Start

● Immediately
○ Later

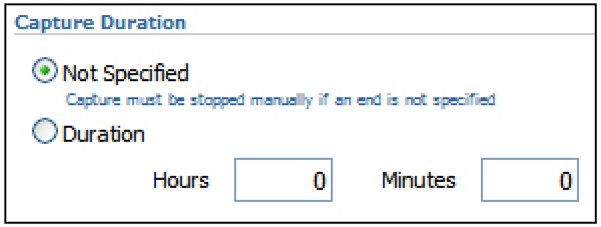

Capture Duration

● Not Specified
Capture must be stopped manually if an end is not specified
○ Duration
Hours 0 Minutes 0

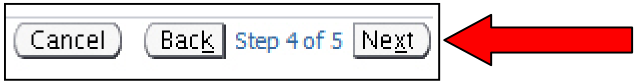

- Submit your Capture job

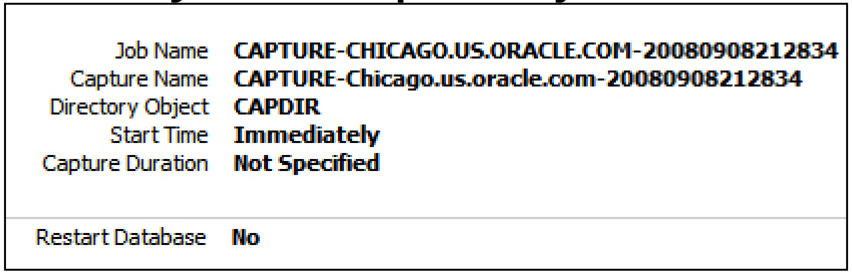

Job Name	**CAPTURE-CHICAGO.US.ORACLE.COM-20080908212834**
Capture Name	**CAPTURE-Chicago.us.oracle.com-20080908212834**
Directory Object	**CAPDIR**
Start Time	**Immediately**
Capture Duration	**Not Specified**
Restart Database	**No**

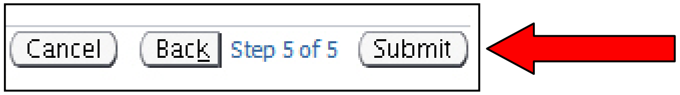

Capturing a Workload

- Execute your workload
  - In our case we ran a set of programs

```
[Chicago] wcr_demo_workload.csh
...
[7]  + Done wcr_demo -d 40 -c salmon -t 0.00001 -u 2000
[6]  + Done wcr_demo -d 30 -c orange -t 0.00001 -u 2000
[5]  + Done wcr_demo -c yellow -d 10 -t 0.00001 -u 2000
[4]  + Done wcr_demo -u 3000 -t 0.01 -c black
[3]  + Done wcr_demo -u 3000 -t 0.01 -c blue
[2]  + Done wcr_demo -u 3000 -t 0.01 -c green
[1]  + Done wcr_demo -u 3000 -t 0.01 -c red
```

Capturing a Workload

- When the programs complete
  - Stop the capture

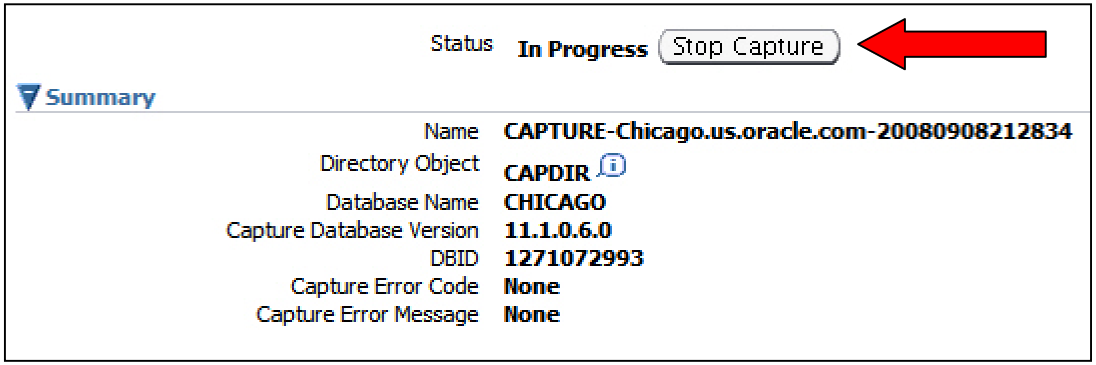

Status **In Progress**

Summary

Name	**CAPTURE-Chicago.us.oracle.com-20080908212834**
Directory Object	**CAPDIR**
Database Name	**CHICAGO**
Capture Database Version	**11.1.0.6.0**
DBID	**1271072993**
Capture Error Code	**None**
Capture Error Message	**None**

  - Export the AWR data

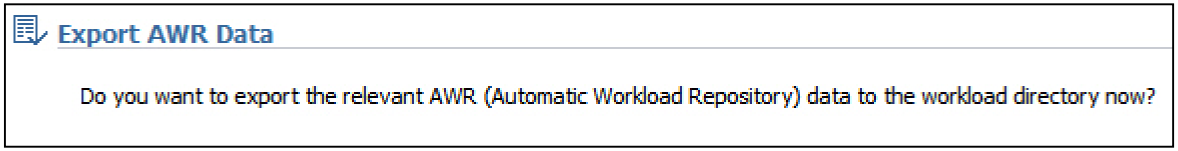

Export AWR Data

Do you want to export the relevant AWR (Automatic Workload Repository) data to the workload directory now?

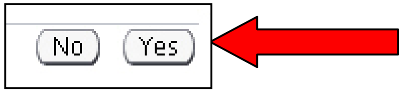

Workload Profile

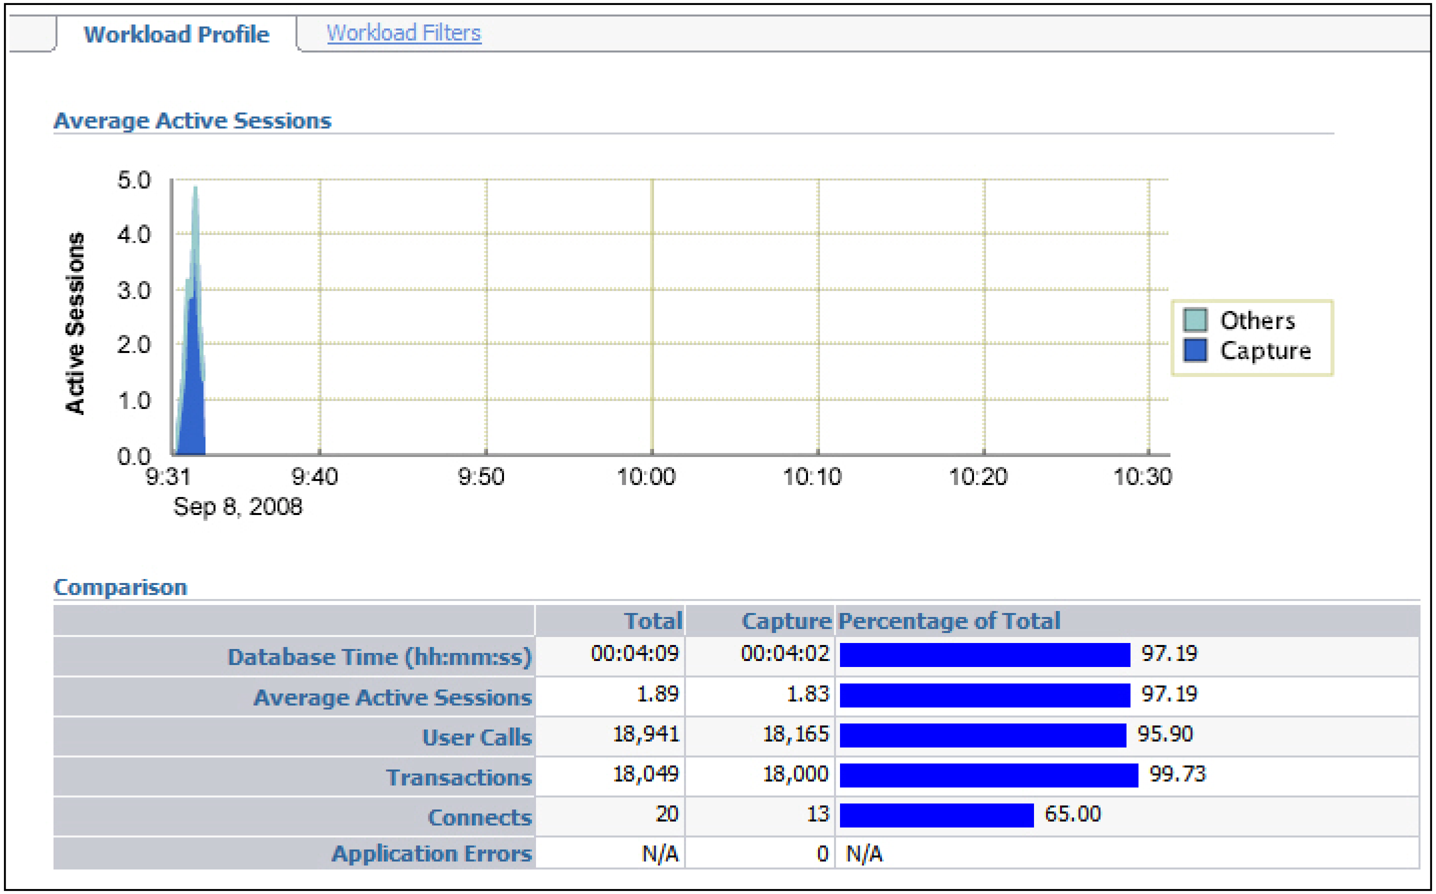

Comparison

	Total	Capture	Percentage of Total
Database Time (hh:mm:ss)	00:04:09	00:04:02	97.19
Average Active Sessions	1.89	1.83	97.19
User Calls	18,941	18,165	95.90
Transactions	18,049	18,000	99.73
Connects	20	13	65.00
Application Errors	N/A	0	N/A

Processing the Captured Workload

- Copy the captured workload files to the standby.
  - Unless you want to process the workload on the Primary
- Log in to Database Control of your Standby database
  - `https://stadu67.us.oracle.com:5501/em/`
- Navigate to the Database Replay page and select the Preprocess Captured Workload task

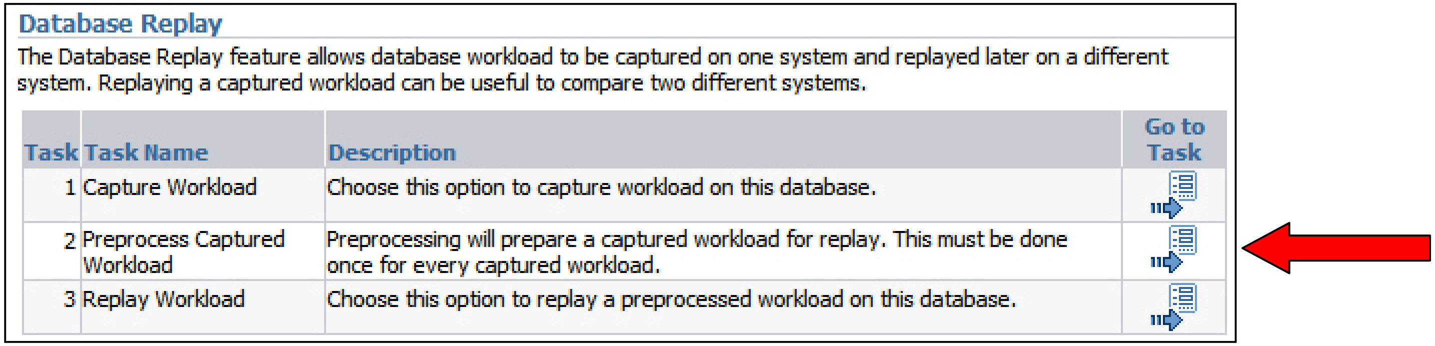

Database Replay

The Database Replay feature allows database workload to be captured on one system and replayed later on a different system. Replaying a captured workload can be useful to compare two different systems.

Task	Task Name	Description	Go to Task
1	Capture Workload	Choose this option to capture workload on this database.	
2	Preprocess Captured Workload	Preprocessing will prepare a captured workload for replay. This must be done once for every captured workload.	
3	Replay Workload	Choose this option to replay a preprocessed workload on this database.	

Processing the Captured Workload

- On the next pages, create the Capture directory object to point to the copied Capture files

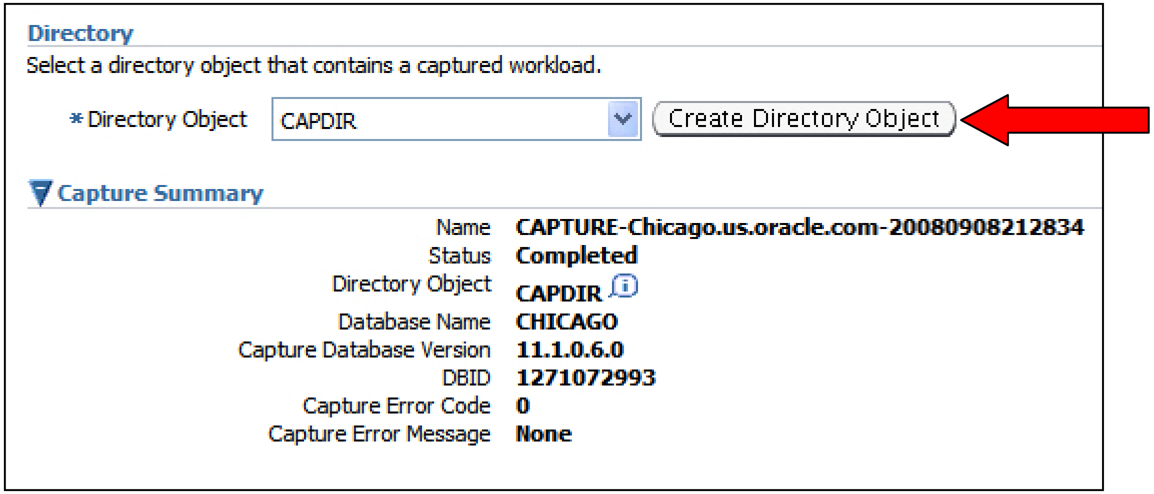

Directory

Select a directory object that contains a captured workload.

* Directory Object	CAPDIR

Capture Summary

Name	CAPTURE-Chicago.us.oracle.com-20080908212834
Status	Completed
Directory Object	CAPDIR
Database Name	CHICAGO
Capture Database Version	11.1.0.6.0
DBID	1271072993
Capture Error Code	0
Capture Error Message	None

- Click the ‘PreProcess Workload” button.

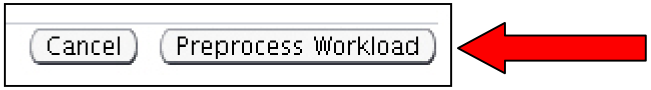

Processing the Captured Workload

- Configure the job
  - Verify the Replay Version and click the Next button.

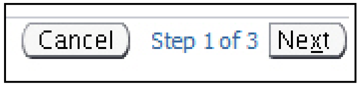

  - Set the Schedule and click Next

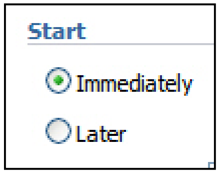
Start

● Immediately

○ Later

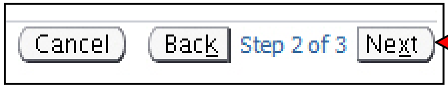

- Submit the Job

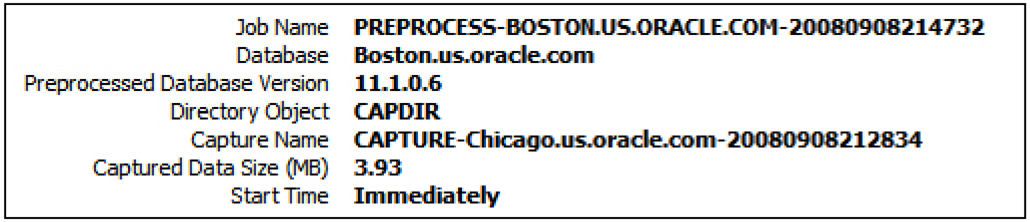

Job Name	**PREPROCESS-BOSTON.US.ORACLE.COM-20080908214732**
Database	**Boston.us.oracle.com**
Preprocessed Database Version	**11.1.0.6**
Directory Object	**CAPDIR**
Capture Name	**CAPTURE-Chicago.us.oracle.com-20080908212834**
Captured Data Size (MB)	**3.93**
Start Time	**Immediately**

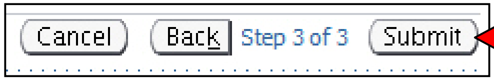

Replaying the Workload

- Run the Replay once without any changes first
  - To obtain your baseline set of metrics
- Then make runs of the Replay with your changes.
- Compare the results of the various runs

- We will run the baseline first and then several indexes to the Snapshot standby to change the workload execution statistics.

Setting the Stage

- Create a guaranteed restore point in the snapshot standby and restart the standby.

```
SQL> create restore point beforereplay
      guarantee flashback database;
SQL> shutdown immediate
SQL> startup
```

  - Allows for multiple replays with the same starting point.
- A Guaranteed Restore Point will require more Flash Recovery Area space as the Flashback logs are not recycled when a GRP is in effect.
- This applies to standard Snapshot standby database as well since they also have their own GRP.

Replaying the Workload

- Navigate to the Database Replay page and select the Replay Workload task and Click the “Go to Task” icon.

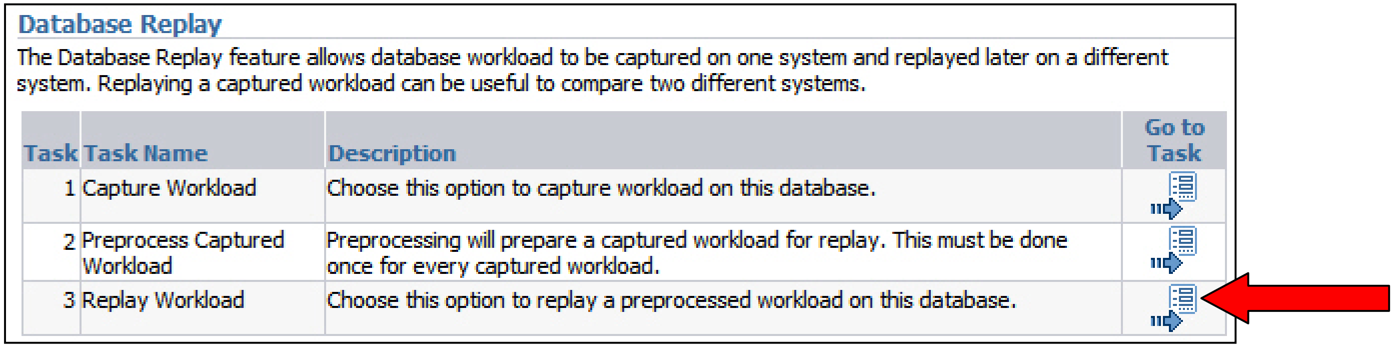

Database Replay

The Database Replay feature allows database workload to be captured on one system and replayed later on a different system. Replaying a captured workload can be useful to compare two different systems.

Task	Task Name	Description	Go to Task
1	Capture Workload	Choose this option to capture workload on this database.	
2	Preprocess Captured Workload	Preprocessing will prepare a captured workload for replay. This must be done once for every captured workload.	
3	Replay Workload	Choose this option to replay a preprocessed workload on this database.	

Configure for Replay

- Create a directory object to point to the capture files.
- Click the ‘Set Up Replay’ button

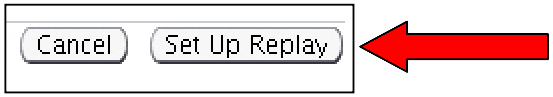

- Ensure the Prerequisites have been met
  - Restore Database Requirement
    - We are using a Snapshot Standby and Flashback Database so we do not need to restore the database
  - Resolve References to External Systems
- Click the ‘Continue’ button

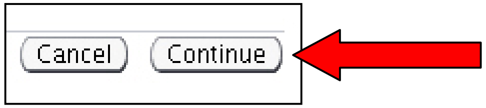

Determine how many Clients you need

- Replay Client is a multithreaded program called ‘wrc’
- When run with the ‘mode=calibrate’ option it will tell you the minimum number of clients to run so that it can start the appropriate number of ‘user’ threads.

```
[Boston] wrc mode=calibrate REPLAYDIR=/scratch/oracle/Replay/Capture

Workload Replay Client: Release 11.1.0.6.0 - Production
Copyright (c) 1982, 2007, Oracle.  All rights reserved.

Report for Workload in: /scratch/oracle/Replay/Capture
-----------------------
Recommendation:
Consider using at least 1 clients divided among 1 CPU(s).

Workload Characteristics:
- max concurrency: 7 sessions
- total number of sessions: 11
```

Configuring Replay

- Choose your Replay options
- Start up the clients

```
[Chicago] wrc REPLAYDIR=/path USERID=xxx PASSWORD=xxx
```

- When Replay sees the clients you can continue

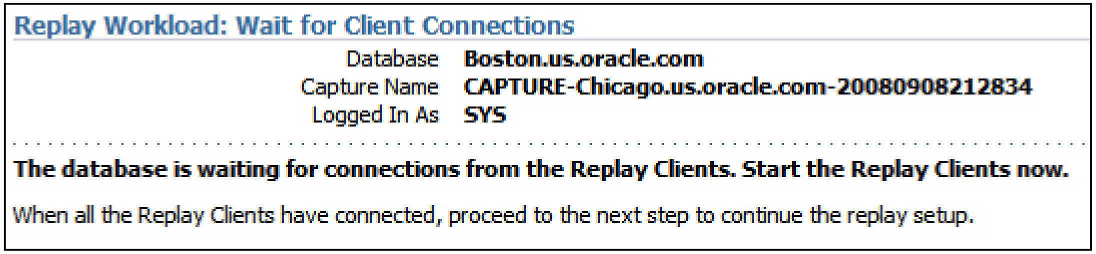

Replay Workload: Wait for Client Connections

Database	**Boston.us.oracle.com**
Capture Name	**CAPTURE-Chicago.us.oracle.com-20080908212834**
Logged In As	**SYS**

The database is waiting for connections from the Replay Clients. Start the Replay Clients now.

When all the Replay Clients have connected, proceed to the next step to continue the replay setup.

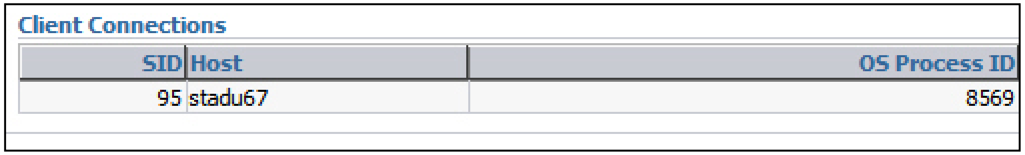

Client Connections

SID	Host	OS Process ID
95	stadu67	8569

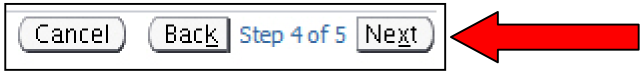

Running the Replay

- Click the ‘Submit’ button to start the replay

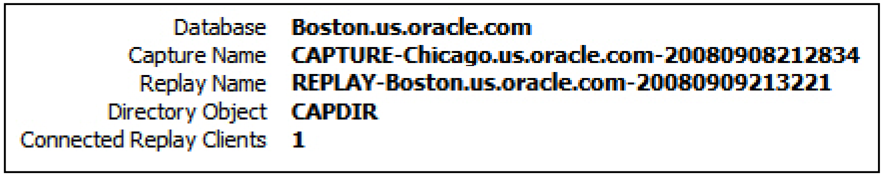

Database	Boston.us.oracle.com
Capture Name	CAPTURE-Chicago.us.oracle.com-20080908212834
Replay Name	REPLAY-Boston.us.oracle.com-20080909213221
Directory Object	CAPDIR
Connected Replay Clients	1

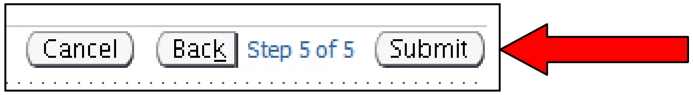

- And wait until it is completed

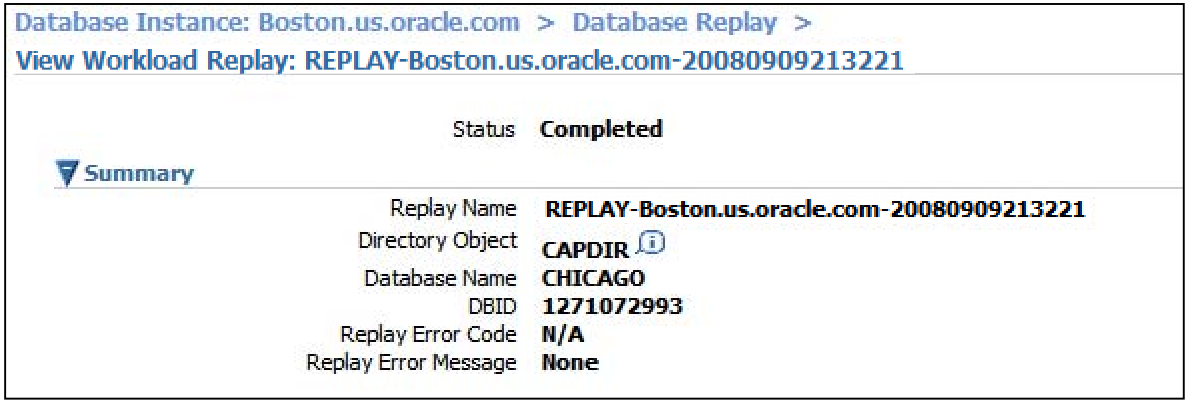

Database Instance: Boston.us.oracle.com > Database Replay >

View Workload Replay: REPLAY-Boston.us.oracle.com-20080909213221

Status	Completed

▼ Summary

Replay Name	REPLAY-Boston.us.oracle.com-20080909213221
Directory Object	CAPDIR ⓘ
Database Name	CHICAGO
DBID	1271072993
Replay Error Code	N/A
Replay Error Message	None

Replay Complete

- Baseline run
  - Create and save the Workload Report

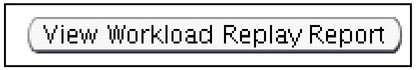

  - Create and save the AWR report

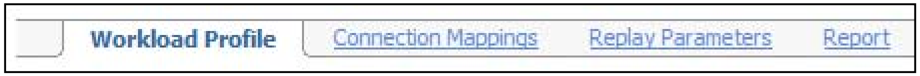

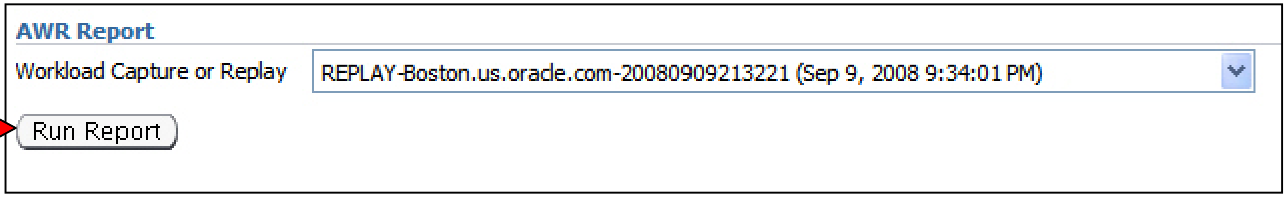

AWR Report

Workload Capture or Replay: REPLAY-Boston.us.oracle.com-20080909213221 (Sep 9, 2008 9:34:01 PM)

Rewind the Snapshot Standby

- Flashback the Snapshot standby to the restore point

```
SQL> shutdown immediate
SQL> startup mount
SQL> flashback database to restore point beforereplay;
SQL> alter database open resetlogs;
SQL> shutdown immediate
SQL> startup
```

- The second restart is not technically necessary
  - It is to ensure that we start with the database in the same state as the Baseline run.

Prepare for the next test and Replay the Workload

- We’ll set up some new indexes.

```
SQL> create index wcr_grid_xcoor_ycoor on wcr_grid(xcoor, ycoor);
SQL> create index wcr_grid_pixid_xcoor_ycoor on
       wcr_grid(pixid,xcoor, ycoor);
SQL> create index wcr_grid_pixid_xcoor on wcr_grid(pixid, xcoor);
SQL> create index wcr_grid_pixid_ycoor on wcr_grid(pixid, ycoor);
SQL> create index wcr_grid_pixid_xcoor_color on
       wcr_grid(pixid, xcoor, color);
SQL> create index wcr_grid_pixid_ycoor_color on
       wcr_grid(pixid, ycoor, color);
SQL> create index wcr_grid_pixid_color on wcr_grid(pixid, color);
SQL> create index wcr_grid_xcoor_color on wcr_grid(xcoor, color);
SQL> create index wcr_grid_ycoor_color on wcr_grid(ycoor, color);
```

- Re-run the entire Replay and compare the results.
  - Same steps as with the Baseline run.

Replay Complete

- New Indexes run
  - Create and save the Workload Report

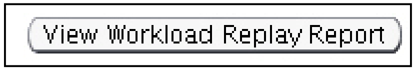

  - Create and save the AWR report

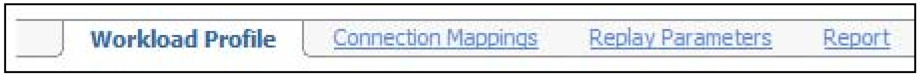

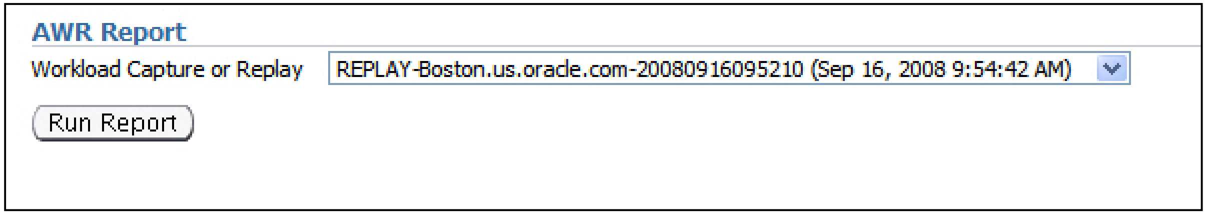
AWR Report

Workload Capture or Replay: REPLAY-Boston.us.oracle.com-20080916095210 (Sep 16, 2008 9:54:42 AM)

Generate the Differences Report

- Create the AWR Differences Report

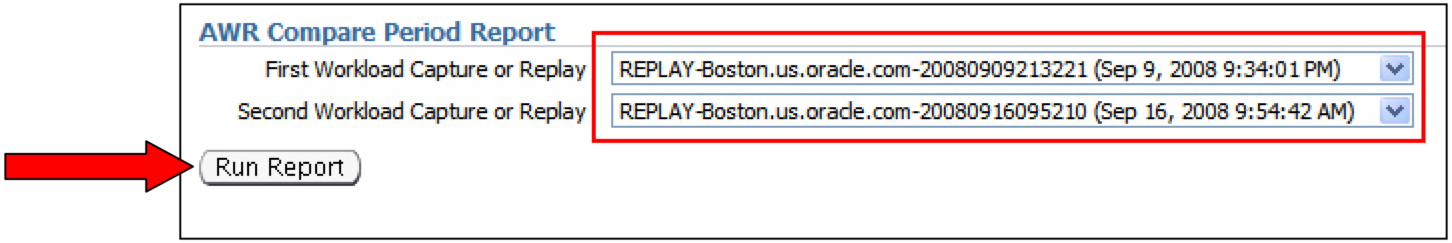

AWR Compare Period Report

First Workload Capture or Replay: REPLAY-Boston.us.oracle.com-20080909213221 (Sep 9, 2008 9:34:01 PM)

Second Workload Capture or Replay: REPLAY-Boston.us.oracle.com-20080916095210 (Sep 16, 2008 9:54:42 AM)

- Requires an Import of the Baseline run AWR data

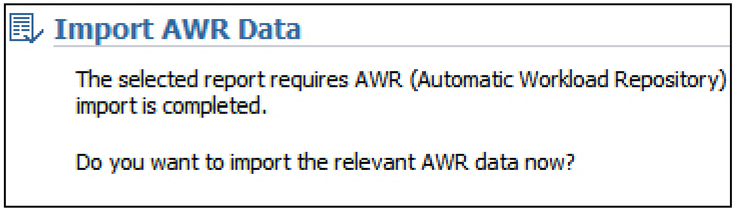

Import AWR Data

The selected report requires AWR (Automatic Workload Repository) import is completed.

Do you want to import the relevant AWR data now?

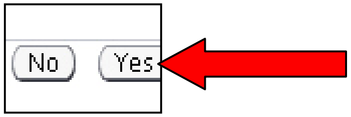

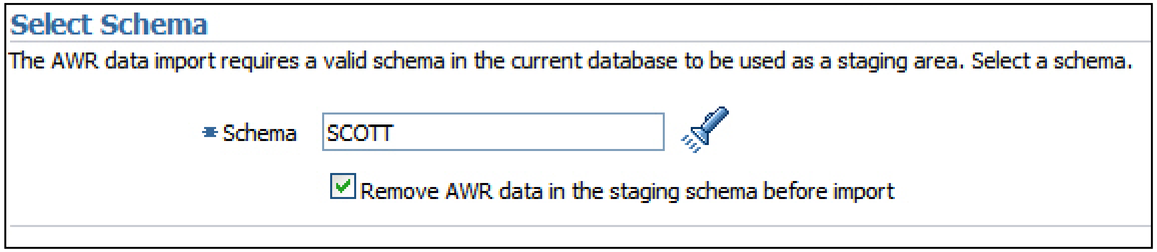

Select Schema

The AWR data import requires a valid schema in the current database to be used as a staging area. Select a schema.

* Schema: SCOTT

☑ Remove AWR data in the staging schema before import

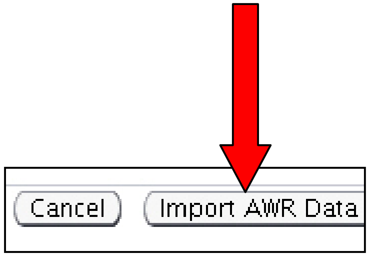

Importing the AWR Data

- Wait for the job to complete

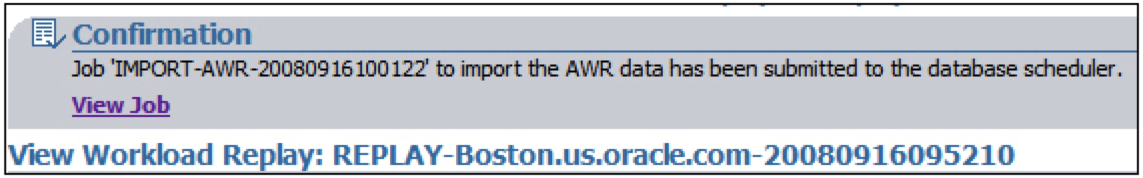

Confirmation

Job 'IMPORT-AWR-20080916100122' to import the AWR data has been submitted to the database scheduler.

View Job

View Workload Replay: REPLAY-Boston.us.oracle.com-20080916095210

- And then execute the run again

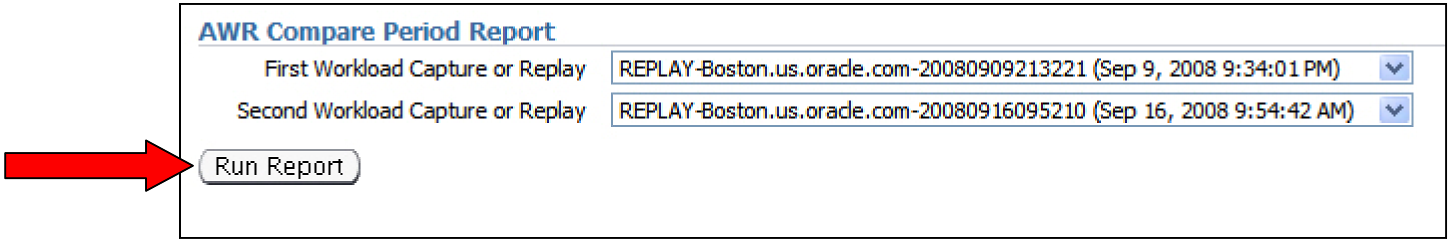

AWR Compare Period Report

First Workload Capture or Replay	REPLAY-Boston.us.oracle.com-20080909213221 (Sep 9, 2008 9:34:01 PM)
Second Workload Capture or Replay	REPLAY-Boston.us.oracle.com-20080916095210 (Sep 16, 2008 9:54:42 AM)

- And save the report

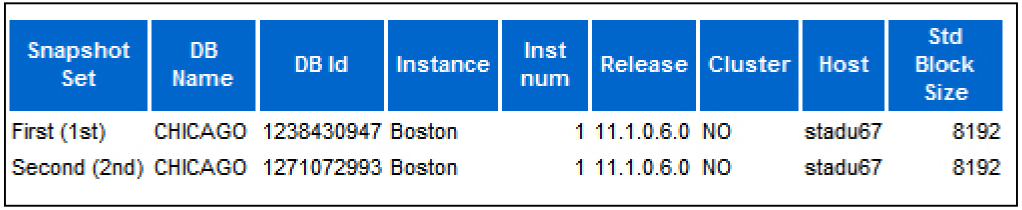

Snapshot Set	DB Name	DB Id	Instance	Inst num	Release	Cluster	Host	Std Block Size
First (1st)	CHICAGO	1238430947	Boston	1	11.1.0.6.0	NO	stadu67	8192
Second (2nd)	CHICAGO	1271072993	Boston	1	11.1.0.6.0	NO	stadu67	8192

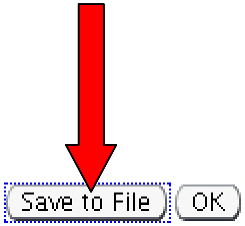

Comparing the Run times

- Baseline Run

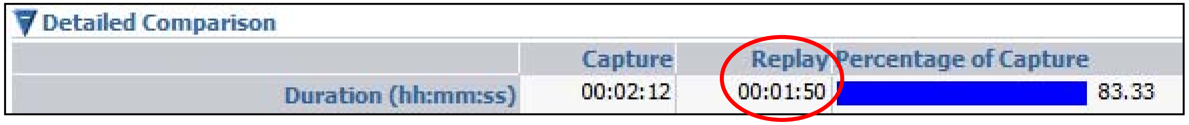

Detailed Comparison

	Capture	Replay	Percentage of Capture
Duration (hh:mm:ss)	00:02:12	00:01:50	83.33

- New Indexes Run

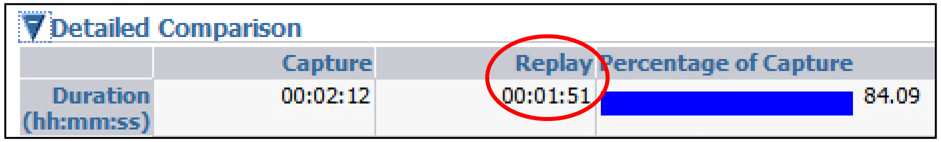

Detailed Comparison

	Capture	Replay	Percentage of Capture
Duration (hh:mm:ss)	00:02:12	00:01:51	84.09

- Small workload but not much difference?
  - Well, let’s look a little closer.

AWR Baseline versus New Indexes

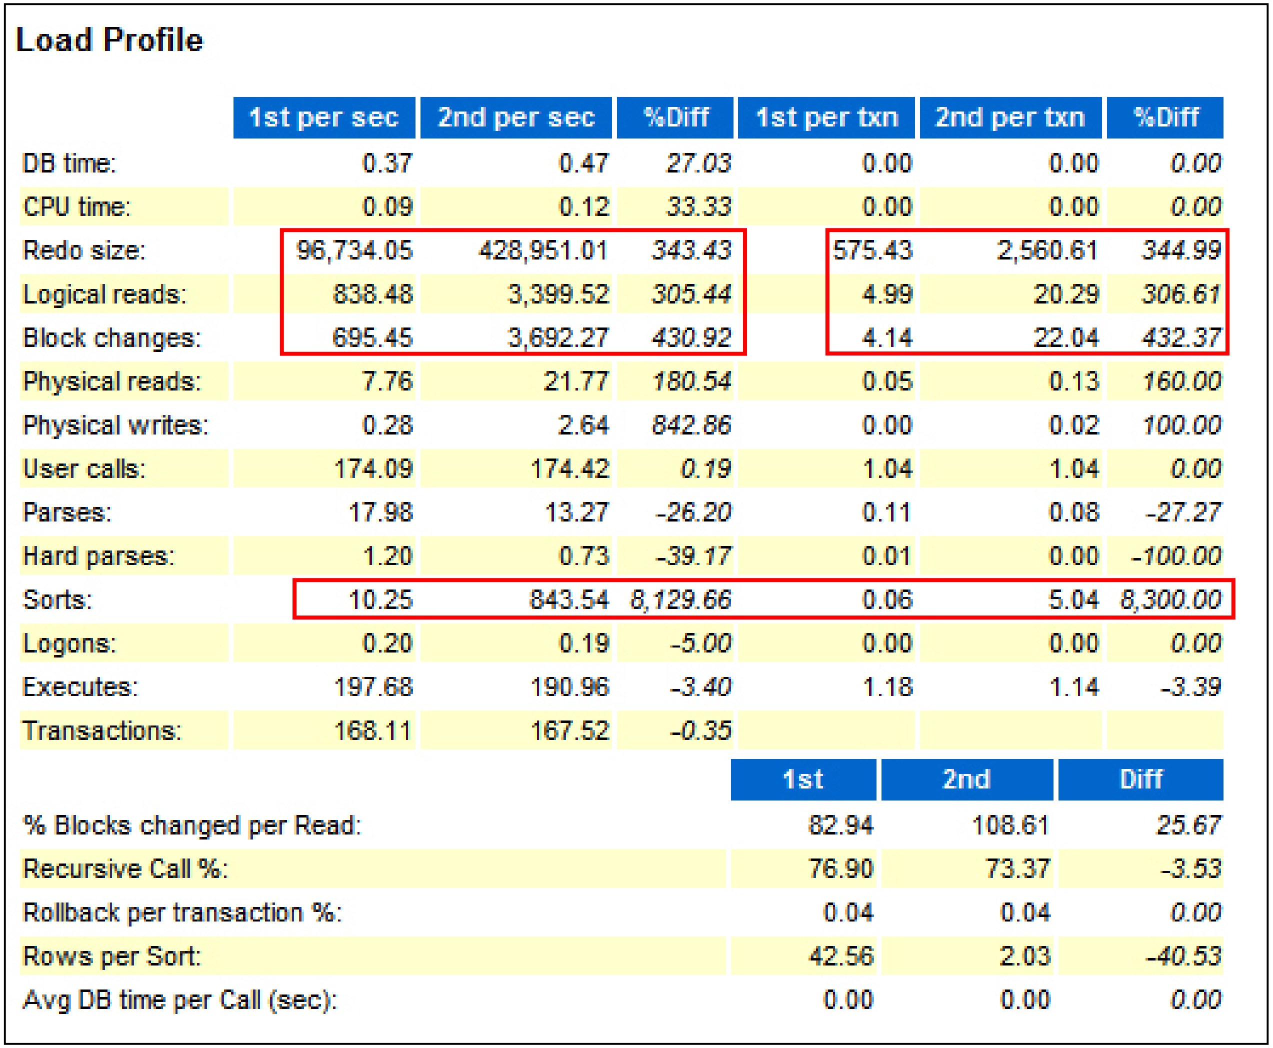

Load Profile

	1st per sec	2nd per sec	%Diff	1st per txn	2nd per txn	%Diff
DB time:	0.37	0.47	27.03	0.00	0.00	0.00
CPU time:	0.09	0.12	33.33	0.00	0.00	0.00
Redo size:	96,734.05	428,951.01	343.43	575.43	2,560.61	344.99
Logical reads:	838.48	3,399.52	305.44	4.99	20.29	306.61
Block changes:	695.45	3,692.27	430.92	4.14	22.04	432.37
Physical reads:	7.76	21.77	180.54	0.05	0.13	160.00
Physical writes:	0.28	2.64	842.86	0.00	0.02	100.00
User calls:	174.09	174.42	0.19	1.04	1.04	0.00
Parses:	17.98	13.27	-26.20	0.11	0.08	-27.27
Hard parses:	1.20	0.73	-39.17	0.01	0.00	-100.00
Sorts:	10.25	843.54	8,129.66	0.06	5.04	8,300.00
Logons:	0.20	0.19	-5.00	0.00	0.00	0.00
Executes:	197.68	190.96	-3.40	1.18	1.14	-3.39
Transactions:	168.11	167.52	-0.35			

	1st	2nd	Diff
% Blocks changed per Read:	82.94	108.61	25.67
Recursive Call %:	76.90	73.37	-3.53
Rollback per transaction %:	0.04	0.04	0.00
Rows per Sort:	42.56	2.03	-40.53
Avg DB time per Call (sec):	0.00	0.00	0.00

Production is still Protected

RPO Unaffected

- While the standby was open Read Write in Snapshot mode the Primary was still shipping the redo

```
SQL> select client_process,process,sequence#, status
     from v$managed_standby
     where client_process='LGWR'
     or process like 'MRP%';

 CLIENT_P PROCESS    SEQUENCE# STATUS
 -------- --------- ---------- ------------
 LGWR     RFS              351 WRITING
```

- At the beginning Production and Standby databases were at Sequence number 321

Converting back to Physical Standby

- Drop the Replay restore point

```
   SQL> drop restore point beforereplay;
```

- Convert the Snapshot to a Physical standby database

```
DGMGRL> convert database boston to physical standby;
Converting database "boston" to a Physical Standby
  database, please wait...
...
Database "boston" converted successfully
```

Automatic Resynchronization

- Redo Apply begins applying the redo that came in while it was in Snapshot standby mode.
- The Primary continues to send the current redo.

```
SQL> select client_process,process,sequence#,status
     from v$managed_standby
     where client_process='LGWR'
     or process like 'MRP%';

CLIENT_P PROCESS    SEQUENCE# STATUS
-------- --------- ---------- ------------
N/A      MRP0             325 APPLYING_LOG
LGWR     RFS              353 IDLE
```

But my Production is still on 10.2!

- Everything we've talked about so far has been 11.1
    - Test on a standby with upgraded ASM and CRS?
        - No changes to the standby. Run Replay as before.
    - Test on a release higher than 11.1?
        - Upgrade the Snapshot standby to a higher release.
        - Requires that you keep COMPATIBILITY at the lower level.
- But what about 10.2?
    - Can I capture a workload on my 10.2 Production database?
        - Yes!
    - And Replay it on my Physical Standby?
    - Yes!
        - Provided that the standby is at 11.1 or higher.
            - It just requires a little more effort.
            - But the results will be worth it!

Using the Capture feature on 10.2

- Follow Metalink note 560977.1
  - “Real Application Testing Now Available for Earlier Releases ”
    - Enables Capture on 10.2
- Convert your physical standby to Read Write
  - More details on the next slide
- Run and Capture your workload on the Primary
- Move the Capture files and Process them.

Configuring a Read Write Standby Oracle Database 10g

- Multiple steps
  - Refer to the Data Guard Concepts and Administration manual, section “12.6 Using a Physical Standby Database for Read/Write Testing and Reporting”
- Manual solution available to maintain RPO.
  - Since Redo is not shipped during this period you can create an Archive Log Repository.
  - Note 434164.1 Data Guard Redo Log Repository Example

Converting to a Read Write Standby Oracle Database 10g

- 10.2 – Basic Steps Required (See 12.6 of the manual!)

```
Standby
   SQL> alter database recover managed standby database cancel;
   SQL> create restore point myrw guarantee flashback database;
Primary
   SQL> alter system archive log current;
   SQL> alter system set log_archive_dest_state_2=defer;
Standby
   SQL> alter database activate standby database;
   SQL> shutdown immediate
   SQL> startup mount;
   SQL> alter database set standby database
        to maximize performance;
   SQL> alter system set log_archive_dest_state_2=defer;
   SQL> alter database open;
```

Performing the Database Replay

- Upgrade your Read Write Standby to 11.1
  - Leaving COMPATIBILITY at 10.2
- Reconfigure Database Control as before.
- Create your ‘beforereplay’ Restore Point.
- Run your Replay
- Analyze the results.
- Flashback to your ‘beforereplay’ restore point
- Rerun the Replay as required.
- Convert your standby back to a Physical standby.
  - Requires you mount the standby back in the 10.2 home after it is converted back into a physical standby.

Converting back to Physical Standby Oracle Database 10g

- 10.2 – Steps Required

Standby

```
SQL> shutdown immediate
SQL> startup mount;
SQL> drop restore point beforereplay;
SQL> flashback database to restore point myrw;
SQL> alter database convert to physical standby;
SQL> shutdown immediate
SQL> startup mount;
SQL> recover managed standby database
     using current logfile disconnect;
```

Primary

```
SQL> alter system set log_archive_dest_state_2=enable;
```

Conclusion

- Real Application Testing
  - Provides an easy method to test your workloads against various changes in your environment without impacting the Primary database
- Data Guard Snapshot Standby database
  - Allows you to replay those workloads with the changes without extra databases and extra storage while maintaining your data protection all at the same time!
- Beat up your standby, not your Production!

Oracle Data Guard

- It can take a beating and come back for more!

HA Sessions, Labs, Demos From Oracle Development

Mon, Sep 22

- 2:30 pm - Database 11*g*: Next-Gen HA, Moscone South 103

Tue, Sep 23

- 9:00 am - Active-Active Data Centers, Moscone South 103
- 11:30 am - Sharding with Oracle, Moscone South 302
- 11:30 am - HA with Oracle VM, Moscone West 3024
- 1:00 pm - Active Data Guard, Moscone South 104

Wed, Sep 24

- 9:00 am - Fusion Middleware Grid HA, Marriott Nob Hill AB
- 11:30 am - RMAN Best Practices, Moscone South 103
- 5:00 pm - Data Guard & Real Application Testing, Moscone 102
- 5:00 pm - EM in Secure MAA, Moscone West 2001

Wed, Sep 24 (contd.)

- 5:00 pm - E-Business Suite HA, Moscone West 2002/04

Thu, Sep 25

- 9:00 am - Oracle Secure Backup, Moscone South 102
- 10:30 am - Streams Replication, Moscone South 102
- 12:00 pm - Rolling Database Upgrades, Moscone South 103
- 1:30 pm - Streams Performance, Moscone South 102
- 3:00 pm - Oracle Grid Computing, Moscone South 303
- 3:00 pm - E-Business Suite R12 MAA, Moscone West 2007
- 3:00 pm - Siebel MAA, Moscone South 308
- 3:00 pm - Fusion SOA HA & Scalability, Marriott Salon 14/15

Hands On Labs - Thu, Sep 25

- 10:30 - 11:30 am, 12:00 - 1:00 pm - Active Data Guard, Marriott Golden Gate A3

DEMOgrounds, Mon-Thu

- Active Data Guard, Streams, Oracle Secure Backup, RMAN/Flashback, MAA

For More Information

search.oracle.com

Maximum Availability Architecture 🔍

or

oracle.com

ORACLE®

Appendix

- Snapshot Standby Reference
- Flash Recovery Area Reference
- Flashback Database Reference

Snapshot Standby Reference

What is a Snapshot Standby?

- An exact Read Write copy of the Primary database
- Preserves zero data loss
  - Continuous redo transport while open read-write
- Truly leverages your standby database and DR hardware for multiple purposes
- Similar to storage snapshots
  - But it provides DR at the same time
  - And uses a single copy of the storage

Flash Recovery Area Reference

The Flash Recovery Area

- Centralized disk location for all recovery files.
- Ensures that the database is completely recoverable by using files stored in the recovery area
  - Archived redo logs
  - Flashback logs
  - Control file autobackups
  - Backup pieces
  - Full backup copy of Production
  - And, If the disks are fast then you can also keep copies of the:
    - Current control file
    - Online redo logs

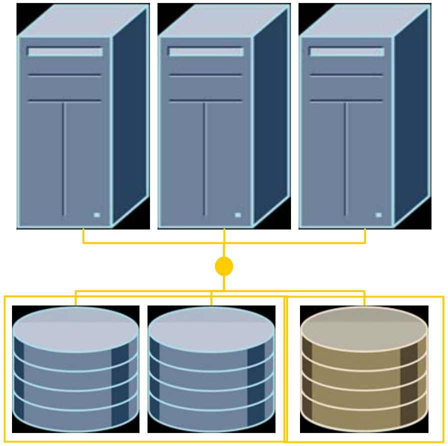

Flash Recovery Area and Data Guard?

- Simplifies the management of archived redo log files
  - You don’t have to worry about cleaning up log files that are no longer necessary to your standby databases.
- Required for Flashback Database
  - Which can be used for many things with Data Guard.

Enabling the Flash Recovery Area

- Set `DB_RECOVERY_FILE_DEST_SIZE` parameter
  - Specifies the maximum total bytes to be used.
  - How big?
    - Total size of archive logs between backups multiplied by 2.
    - Amount of Redo generated during Flashback Retention period.
    - Typical size of your incremental backups.
      - Incremental backup size is very dependent on workload.
  - All of the above AND an on-disk backup?
    - Add in the size of the database minus the temp files.
- Set `DB_RECOVERY_FILE_DEST` parameter
  - Location of the flash recovery area.
    - Can be a directory, file system, or an ASM disk group.
      - It cannot be a raw file system.

What about my old archiving?

- Change it to use the Flash Recovery Area

```
    LOG_ARCHIVE_DEST_1 =
      'LOCATION = /disk2/archive
      VALID_FOR=(ALL_LOGFILES,ALL_ROLES)‘
```

To

```
    LOG_ARCHIVE_DEST_1 =
      'LOCATION=USE_DB_RECOVERY_FILE_DEST
      VALID_FOR=(ALL_LOGFILES,ALL_ROLES)‘
```

Automate Archive Log Management

- Move backups to a Physical standby database
- Use RMAN to setup the Archive log retention policy
  - Primary database retention policy

    ```
    CONFIGURE ARCHIVELOG DELETION POLICY TO
    [APPLIED ON | SHIPPED TO] ALL STANDBY;
    ```
  - On the Physical Standby where the backups are being done.

    ```
    CONFIGURE ARCHIVELOG DELETION POLICY TO NONE;
    ```
  - On all other standby databases

    ```
    CONFIGURE ARCHIVELOG DELETION POLICY TO
    APPLIED ON STANDBY;
    ```
- After Switchover or Failover, re-execute the proper configure commands again as appropriate

Maintaining Protection (RPO) in 10.2 while in Read Write mode

- Create an Archive Log Repository (ALR)
  - Note 434164.1 Data Guard Redo Log Repository Example
  - Redo is shipped to reduce RPO in case of a disaster at Production
- Put the Archive Log Repository in place before you begin the Snapshot process.
- When the Physical standby has been returned
  - Use these logs to catch it up with the Primary
  - Or use an incremental backup of the Primary
- This is not required in Oracle Database 11g

Flashback Database Reference

Flashback Database

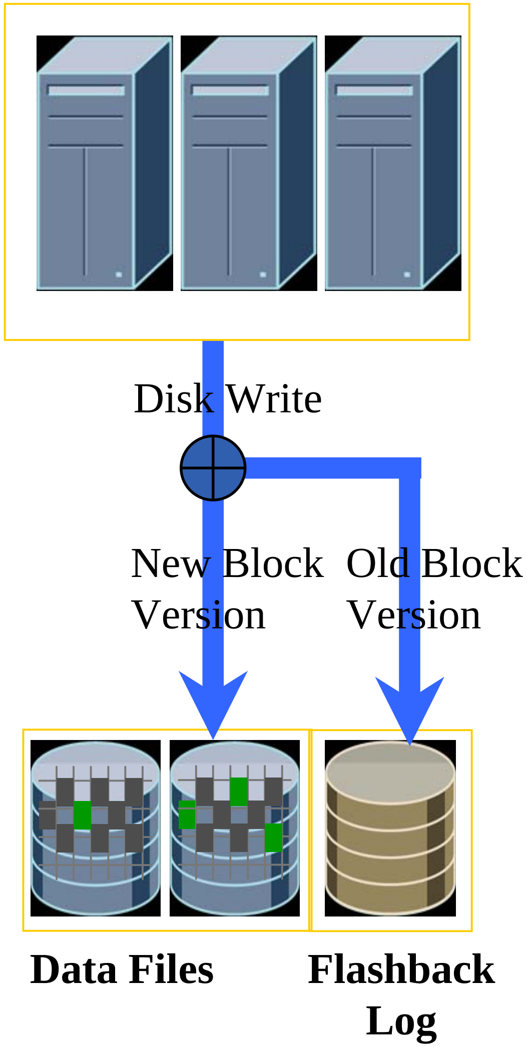

- A strategy for point in time recovery
- Flashback Log captures old versions of changed blocks
  - A continuous backup
  - Replay log to restore DB to time
  - Restores just changed blocks

"Rewind" button for the Database

Why use Flashback Database?

- Easy correction of User errors without impacting production or having the standby lag behind.
  - Rewind the standby to fix user errors
  - Failovers can happen without affecting RTO
- Enables a failed Primary to become a standby after a failover without re-copying all the files.
- Required for Snapshot standby
  - Opening your Physical standby for read and write
  - Without it this talk would be meaningless
- See Note 565535.1 Flashback Database Best Practices & Performance

Enabling Flashback Database

- Configure Flashback parameters on both your Primary and Physical standby databases

```
DB_FLASHBACK_RETENTION_TARGET
```

- On the Primary

```
SQL> SHUTDOWN IMMEDIATE
SQL> STARTUP MOUNT
SQL> ALTER DATABASE FLASHBACK ON;
SQL> ALTER DATABASE OPEN;
```

- On the Standby

```
SQL> ALTER DATABASE FLASHBACK ON;
```

  - Physical standby databases are usually mounted anyway
  - Cancel Managed Recovery before enabling Flashback Database.
  - If you were using Active Data Guard follow the same procedure as the Primary Database and restart the apply after the standby database is open again.

Monitoring Flashback Database

- Check the Flashback settings

```
SQL> SELECT flashback_on, current_scn
  2  FROM v$database;
```

- Monitor the Flashback retention target

```
SQL> SELECT * FROM v$flashback_database_log;
```

- Adjust Flash Recovery Area disk quota

```
SQL> SELECT estimated_flashback_size
  2  FROM v$flashback_database_log;
```

- Monitor logging in the Flashback logs

```
SQL> SELECT * FROM v$flashback_database_stat;
```

ORACLE IS THE INFORMATION COMPANY