

HOW A RAT WILL SAVE YOUR JOB

Dan Norris, Piocon

What Is It?

Oracle Real Application Testing is the marketing name given to two new features introduced with Oracle Database 11g: Database Replay and SQL Performance Analyzer. These features provide a way to accurately simulate a database workload to protect from the unpredictable impact of routine database changes. Routine changes like creation of a new index or an application upgrade can have unpredictable side effects. The only way to ensure that this change won't have any unintended consequences is to test the normal SQL workload on the database after making the changes. Historically, the most difficult part of this testing process is to create something that approaches a realistic workload. Database Replay provides the ability to capture actual workload and "play" it back on the same or a different environment.

At the highest levels, using Real Application Testing involves four major steps or phases:

1. **Workload capture:** Capture the workload from a running production database as normal processing continues.
2. **Workload preprocessing:** Following capture, a processing step is necessary to transform the captured information into a format suitable for replay.
3. **Workload replay:** Typically on a separate system that has changes implemented, the workload is replayed while capturing performance information.
4. **Analysis and reporting:** Once the workload is replayed, Oracle provides tools to compare the performance of the modified environment with the original baseline environment to determine positive or negative impact. This is where the SQL Performance Analyzer shines and shows its usefulness as it compares before and after performance of SQL statements.

Using these four phases, you can guide yourself through the process. In the current release, the manual processes are labor-intensive, but the Oracle Enterprise Manager makes the process relatively easy. Usability is almost certain to be on the enhancement list for future releases.

In this whitepaper, we'll review the mechanics of using Database Replay and how its use may benefit you as you make changes to applications and databases.

Capture Processes

Capturing a workload employs a PL/SQL package interface to configure, start, and stop the capture process. The capture can be filtered to capture only the actions of some of the database users instead of its default behavior to capture all database activity. The capture process operates at a database level, not an instance level, so RAC environments need no special configuration.

Where they run

Capture processes run on the database server where activity is being captured. These processes are controlled by the PL/SQL package interface (DBMS_WORKLOAD_CAPTURE). On UNIX-based systems, you can see the processes running after you start capture. The capture process is handled by dedicated server processes, no additional background processes are spawned for capture. A directory object is created in the database prior to capture starting and that directory is where the capture files, generated by the capture process, will be placed. These are placed on a filesystem directory since putting them in the database would create a recursive activity that the capture would also be capturing.

What they take as input, output

Starting a capture process takes relatively few arguments. However, there are several inputs to consider. First, let's review the list of items that are not able to be captured currently:

- Direct path loads from external files using tools like SQL*Loader
- Shared server requests
- Oracle Streams
- Advanced Replication streams

- Non-PL/SQL-based Advanced Queuing (AQ)
- Flashback queries
- Oracle Call Interface (OCI) based object navigations
- Non SQL-based object access
- Distributed transactions (any distributed transactions that are captured will be replayed as local transactions)
- Remote DESCRIBE and COMMIT operations

These items will not be captured and seeing things like "Flashback queries" on the list reminds us that all SQL query activity, not just changes, will be captured and available for replay. The think time in between queries as well as the ordering of events will be replayed just as it happened on the production server where it was captured. This provides the most realistic representation of workload available.

Besides the exclusion list above, you can also filter the data that is captured at capture time by creating filters. When defining a filter, you enter criteria that will be used by the capture to either include or exclude items matched by the filter(s). Filters can be defined on any of the following items: `INSTANCE_NUMBER` (useful for RAC environments), `USER`, `MODULE`, `ACTION`, `PROGRAM`, or `SERVICE`. Multiple filters can be defined to either filter for multiple users or for multiple different types of criteria. To add a filter, use the `DBMS_WORKLOAD_CAPTURE.ADD_FILTER` procedure.

The capture process requires two parameters and three additional parameters are optional. The two required parameters are the name (helpful for monitoring and stopping the capture) and the `dir` (database directory object name) where capture output will be written. Additionally, three optional parameters will likely be used frequently. The optional parameters are `duration`, `default_action`, and `auto_unrestrict`. `Duration` is self-explanatory and specified in sections. If not specified, capture will continue until the user stops the capture with `DBMS_WORKLOAD_CAPTURE.FINISH_CAPTURE`. The next optional parameter, `default_action`, determines the default behavior for the capture process and allows values of `INCLUDE` or `EXCLUDE`. If `INCLUDE` is specified (which is the default), all activity is captured and the filters are considered exclusion filters to remove items from the capture. If `EXCLUDE` is specified, then only the activity that matches a filter will be captured. The last parameter, `auto_unrestrict`, is a Boolean that when set to `TRUE` will take instances out of restricted mode once the capture is started successfully.

The output of the capture process is the capture files that reside in the directory given to the `DBMS_WORKLOAD_CAPTURE.START_CAPTURE` procedure as the second argument. These files can be moved to another server for processing in the next phase of the database replay process.

Overhead

The capture process does introduce some overhead to the database server. Current testing shows that the overhead amounts to approximately 4.5%. Compared with the cost of not performing the testing, this is a minor impact and is very similar to the overhead introduced by today's monitoring tools.

Version requirements

Currently capture is supported on version 10.2.0.4 and higher databases. Capturing a workload on Oracle Database 10g can be especially useful when preparing for an upgrade to Oracle Database 11g. While workloads can be captured on Oracle Database 10g, they cannot be replayed on Oracle Database 10g.

Workload Preprocessing

Once the capture is complete, it must be processed in order to make it ready for playback. The processing step is not difficult to complete. It only requires a single command. Note that it is this step that makes it clear that there can be only one set of capture files in a given capture directory. That is because the only parameter given to the `DBMS_WORKLOAD_REPLAY.PROCESS_CAPTURE` procedure is the capture directory where capture files were directed.

Note that the preprocessing does not have to be done on the capture source machine. It is typically done on the same server where the replay is going to take place so that the replay source machine isn't burdened with the additional overhead of preprocessing.

Replay Clients

To replay your processed workload, you must coordinate a few commands, some in the database and one on the operating system. When running database capture and replay, AWR automatically collects snapshots at the start and end of the capture and replay as well as at its normal regular intervals during the capture and replay. The replay clients are processes that run on the operating system, outside the database, to execute the workload that was captured and processed.

Another important consideration when replaying the workload is to properly configure the test database that will be used to execute the replay. The target for the replay should have the same data as the production database did when the capture was started. When these are exactly the same, you ensure realistic results. One popular way of preparing the replay target database is to restore a production backup and recover it until the proper SCN (obtained from the capture's logfile).

Where they run

Following the capture process, no other processing takes place on the capture source machine. The replay clients will usually be executed from other machines (not the replay database server) so that the overhead of replaying the workload isn't taxed on the database server. That way, the database server can focus on the single task of running the database, providing a realistic view of most production environments (where application servers are employed to do non-database application tasks).

To begin a replay session, you first need to execute a few PL/SQL procedures to set up the database for replay. First, run the `DBMS_WORKLOAD_REPLAY.INITIALIZE_REPLAY` to place the database in `INIT FOR REPLAY` mode. You must supply the name of the replay and the replay directory object name (where the preprocessed capture data resides) as arguments to `INITIALIZE_REPLAY`. After that's complete, then prepare the database for replay by running `DBMS_WORKLOAD_REPLAY.PREPARE_REPLAY` (no arguments required).

Now that the database is prepared, one or more replay clients will need to be started on one or more replay client hosts (you can also start the replay client(s) on the database server). The replay client is the binary `$ORACLE_HOME/bin/wrc`. It requires a few parameters to connect to the database and locate the replay directory. It also optionally allows several others (see the documentation, section 22.4.2.5 in the Oracle Database Performance Tuning Guide 11g Release 1 (11.1) for more information). Once one or more replay clients are connected and running, you can start the replay with `DBMS_WORKLOAD_REPLAY.START_REPLAY`. This will begin the workload replay by distributing instructions to the replay clients per the parameters you provide.

What they take as input, output

There are several procedures required to start and manage replay in addition to the workload replay client (`wrc`) utility. Let's examine arguments to each of these.

Procedure or Utility	Argument	Required or Optional	Description
<code>DBMS_WORKLOAD_REPLAY.INITIALIZE_REPLAY</code>	<code>replay_name</code>	Required	An arbitrary name for the replay
<code>DBMS_WORKLOAD_REPLAY.INITIALIZE_REPLAY</code>	<code>replay_dir</code>	Required	The database directory object where the capture files are located
<code>DBMS_WORKLOAD_REPLAY.PREPARE_REPLAY</code>	<code>synchronization</code>	Optional	A Boolean parameter indicating if COMMIT order is preserved from the capture, default TRUE
<code>DBMS_WORKLOAD_REPLAY.PREPARE_REPLAY</code>	<code>connect_time_scale</code>	Optional	Scaling factor to determine the time elapsed between the capture start and when sessions connected to the database, expressed as a percentage, default 100
<code>DBMS_WORKLOAD_REPLAY.PREPARE_REPLAY</code>	<code>think_time_scale</code>	Optional	Scaling factor to determine the time elapsed between two successive calls from the same session, expressed as a percentage, default 100
<code>DBMS_WORKLOAD_REPLAY.PREPARE_REPLAY</code>	<code>think_time_auto_correct</code>	Optional	A Boolean parameter to indicate if replay should adjust think time between calls such that the elapsed replay time matches elapsed capture time, default TRUE
<code>DBMS_WORKLOAD_REPLAY.REMAP_CONNECTION</code>	<code>connection_id</code>	Required	A connection id from <code>DBA_WORKLOAD_CONNECTION_MAP</code>

Procedure or Utility	Argument	Required or Optional	Description
DBMS_WORKLOAD_REPLAY.REMAP_CONNECTION	replay_connection	Required	The new connection string that should be used during replay
DBMS_WORKLOAD_REPLAY.START_REPLAY	none		
wrc	mode	Required	The replay client mode, one of <code>calibrate</code> , <code>replay</code> , or <code>list_hosts</code>
wrc	replaydir	Optional	The directory containing replay files, default is the current directory
wrc	process_per_cpu, threads_per_cpu	Optional	<code>process_per_cpu</code> and <code>threads_per_process</code> are optional parameters (defaults 4 and 50, respectively) that may be provided in <code>calibrate</code> mode
wrc	userid, password, server, workdir, debug, connection_override	Optional	These parameters may be provided in <code>replay</code> mode. <code>userid</code> and <code>password</code> are self-explanatory. <code>server</code> is the TNS connection string to the replay database. <code>workdir</code> specifies a directory path where debug information may be written if <code>debug</code> is set to <code>files</code> . <code>debug</code> may be set to <code>files</code> , <code>stdout</code> , <code>both</code> , or <code>none</code> (default is <code>none</code>). <code>connection_override</code> is a Boolean when <code>TRUE</code> ignores remapping and uses the value of the <code>server</code> parameter for all connection remapping.

The output for the procedures should not include much other than the usual "PL/SQL Procedure executed". The `wrc` command may have some diagnostic output that will be helpful for troubleshooting, but it is not required as input to any other downstream process (so no need to capture it). All data about the replay and performance of the replay is captured in the database by AWR and other processes.

Overhead

Due to the replay architecture, overhead on the database server should be minimized since the replay clients can be run on separate servers. When the `wrc` clients are run on the replay database server, there is some overhead to the database server, so that should be a consideration when attempting to perform a realistic comparison.

It's also important to realize that the `wrc` executable provides all the necessary infrastructure to perform the replay, so client platform.

Analyzing Replay Results

Once replay is complete, you will want to compare the replay system results with the results from the capture system. The SQL Performance Analyzer, the second feature under the name Real Application Testing, provides an interface to make just such a comparison.

Using the SQL Performance Analyzer, Workload Report, and AWR Compare Period Report, you can quickly see the impact of your changes—for better or for worse. When attempting to identify how a single change may impact a particular SQL statement, the SQL Performance Analyzer is the best tool. For reviewing overall impact to the whole database, an AWR Compare Period Report will likely be the most insightful.

Employing RAT

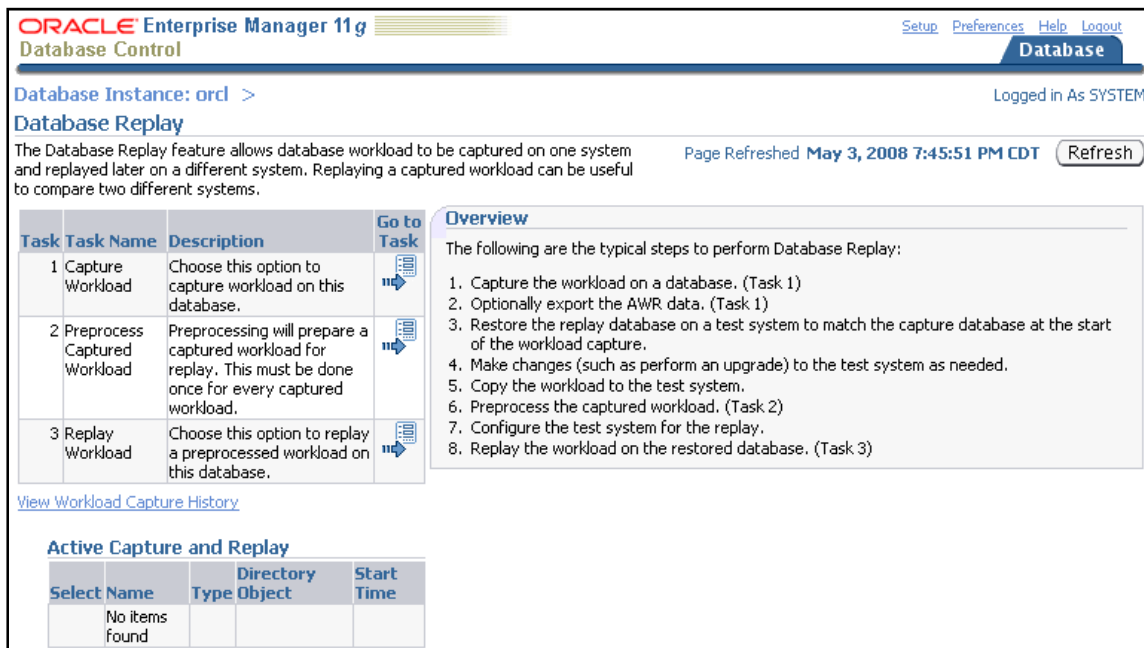
Putting these features to work is made easier by employing Oracle Enterprise Manager (OEM). There are many procedures to run and settings to configure that must be done properly and in sequence to achieve usable test results. Navigating the entire process is much easier when guided by OEM. In this section, we'll demonstrate an example capture and replay using OEM.

Oracle Enterprise Manager Interface

The OEM interfaces for Database Replay and SQL Performance Analyzer are easy to find under the Software and Support tab > Real Application Testing.



Going to the Database Replay link leads you to a page where you can launch the capture, preprocessing and replay operations.



Capture screens

To create a capture, OEM leads you through all the steps. First, you have to agree to all the conditions of the capture. This is OEM's way of telling you that it can't check the prerequisites very well and also warn you that in order to get a good capture, you'll want to restart the database. You have to acknowledge all three prerequisites before you can continue with the capture.

ORACLE Enterprise Manager 11g Database Control

Plan Environment Options Parameters Schedule Review

Capture Workload: Plan Environment

Database orcl
Logged In As SYSTEM

Cancel Step 1 of 5 Next

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.

Prerequisite	Acknowledge
Restarting the database prior to workload capture is recommended for the best workload replay result. Consider scheduling the workload capture at a time when the database can be restarted.	<input type="checkbox"/>
Make sure there is enough disk space to hold the captured workload. Consider doing a short duration workload capture and use it for estimating disk space requirement of a full workload capture.	<input type="checkbox"/>
Make sure you can restore the replay database to match the capture database at the start of the workload capture. A successful workload replay depends on application transactions accessing application data identical to that on a capture system. Common ways to restore application data state include point-in-time recovery, flashback, and import/export.	<input type="checkbox"/>

Cancel Step 1 of 5 Next

On the next screen in the process, we define the capture options. For my test capture, I am not going to restart the database and I'm going to use an inclusion filter to only include one user's actions (DAN) in the capture.

ORACLE Enterprise Manager 11g Database Control

Plan Environment Options Parameters Schedule Review

Capture Workload: Options

Database orcl
Logged In As SYSTEM

Cancel Back Step 2 of 5 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.

Workload Filters

Workload filters are used to customize what workload will be captured. By default, most external client requests made to the database are captured. Refer to system documentation for more information.

Filter Mode: Inclusion

Included Sessions

Only those sessions specified below will be captured.

Filter Name	Type	Session Attribute	Value	Remove
only DAN	Included	User	DAN	

Add Another Row

TIP You may use % for wildcard in a filter value.

Cancel Back Step 2 of 5 Next

With options defined, we need to name the capture and define where the capture files will be output. Note that the directory you choose is a directory object defined in the database. If you need to create a new destination, the Create Directory Object will allow you to create one without leaving this wizard.

The screenshot shows the 'Capture Workload: Parameters' step of the Oracle Enterprise Manager 11g Database Control wizard. The breadcrumb trail at the top indicates the current step is 'Parameters', with previous steps being 'Plan Environment', 'Options', 'Schedule', and 'Review'. The page title is 'Capture Workload: Parameters'. Below the title, it shows 'Database orcl' and 'Logged In As SYSTEM'. There are 'Cancel', 'Back', 'Step 3 of 5', and 'Next' buttons. The main section is titled 'Workload Capture Parameters' and contains two fields: 'Capture Name' with the value 'only-dan-odtug-capture' and 'Directory Object' with a dropdown menu showing 'DAN_CAPTURE_DEST' and a 'Create Directory Object' button. A note below the dropdown states: 'Select a directory object to hold the captured workload. The selected directory object cannot already contain a captured workload.' At the bottom, there are 'Cancel', 'Back', 'Step 3 of 5', and 'Next' buttons.

Now that the capture is defined, we need to start it by defining a job to begin the capture using the next screen in the wizard. I'm going to define this to start immediately and run until I stop it. You may find that helpful as well to ensure that you capture exactly what you want and then stop the capture.

The screenshot shows the 'Capture Workload: Schedule' step of the Oracle Enterprise Manager 11g Database Control wizard. The breadcrumb trail at the top indicates the current step is 'Schedule', with previous steps being 'Plan Environment', 'Options', 'Parameters', and 'Review'. The page title is 'Capture Workload: Schedule'. Below the title, it shows 'Database orcl' and 'Logged In As SYSTEM'. There are 'Cancel', 'Back', 'Step 4 of 5', and 'Next' buttons. The main section is titled 'Job Parameters' and contains a 'Job Name' field with the value 'dan-capture-job-odtug' and an empty 'Description' field. Below this is the 'Job Schedule' section with the instruction: 'Choose a start time and a capture duration so that the workload you are interested in replaying at a later time can be captured.' It has two sub-sections: 'Start' and 'Capture Duration'. Under 'Start', 'Immediately' is selected with a radio button, and 'Later' is unselected. The 'Date' field is set to 'May 3, 2008' and the 'Time' field is set to '7:45:00 PM'. Under 'Capture Duration', 'Not Specified' is selected with a radio button, and 'Duration' is unselected. The 'Hours' and 'Minutes' fields are both set to '0'. Below this is the 'Job Credentials' section with 'Host Credentials' fields for 'Username' (oracle), 'Password' (masked with asterisks), and 'Confirm Password' (masked with asterisks). There is a checkbox for 'Save as Preferred Credential' which is unchecked. At the bottom, there are 'Cancel', 'Back', 'Step 4 of 5', and 'Next' buttons.

The final screen in the wizard provides a summary of the capture task and gives you one last chance to make sure everything's correct before starting.

ORACLE Enterprise Manager 11g Database Control

Plan Environment Options Parameters Schedule **Review**

Capture Workload: Review

Database **orcl** Logged In As **SYSTEM** [Cancel] [Back] Step 5 of 5 [Submit]

Review the following settings for capturing the workload.

Job Name **DAN-CAPTURE-JOB-ODTUG**
 Capture Name **only-dan-odtug-capture**
 Directory Object **DAN_CAPTURE_DEST**
 Start Time **Immediately**
 Capture Duration **Not Specified**

Database Restart

Restart Database **No**

Workload Filters: Included Sessions

Filter Name	Type	Session Attribute	Value
only DAN	Included	User	DAN

[Cancel] [Back] Step 5 of 5 [Submit]

Once the capture is started, you'll be directed to the workload capture monitoring screen.

ORACLE Enterprise Manager 11g Database Control

Database Instance: orcl > Database Replay > Logged In As SYSTEM

Confirmation

The workload capture has started.

Once the capture is complete you will need to do the following prior to replaying the workload on a different system:

1. Optionally export the AWR data.
2. Restore the replay database on a test system to match the capture database at the start of the workload capture.
3. Make changes (such as perform an upgrade) to the test system as needed.
4. Copy the workload to the test system.
5. Preprocess the captured workload.

View Data Real Time: 60 Second Refresh

View Workload Capture: only-dan-odtug-capture

Page Refreshed May 3, 2008 8:17:44 PM CDT [Refresh] [OK]

Status **In Progress** [Stop Capture]

Summary

Name	only-dan-odtug-capture	Captured Data Size (MB)	0.00
Directory Object	DAN_CAPTURE_DEST	Duration (hh:mm:ss)	00:00:24
Database Name	ORCL	Start Time	May 3, 2008 8:17:20 PM CDT
Capture Database Version	11.1.0.6.0	End Time	N/A
DBID	1180446881	Start SCN	790636
Capture Error Code	None	End SCN	N/A
Capture Error Message	None		

At this stage, you should begin the workload that you wish to capture. Once the workload is complete or you've captured as much as you need, click the Stop Capture button. Once the capture is stopped, OEM will ask if you want to export AWR data as shown here:

ORACLE Enterprise Manager 11g Database Control

Export AWR Data

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

Exporting the AWR data from this database enables in-depth capture and replay analysis. A database scheduler job will be created to perform the export immediately.

TIP If you choose not to export the AWR data now, you may perform the export at a later time from the page that lists the capture history on this database.

[No] [Yes]

Having the AWR data along with the capture files is convenient, so I choose Yes. A job is submitted to export the AWR data which takes a few minutes in my case.

With the capture complete, we can move on to the preprocessing step. We go back to the Database Replay screen and click on the link for task 2, "Preprocess Captured Workload".

Workload capture processing screens

In the preprocessing wizard, the first prompt is for the directory object that references the location of the capture files. Once you choose a directory object, the wizard locates the capture files, reads them, and presents a summary of the capture.

ORACLE Enterprise Manager 11g Database Control

Database Instance: orcl > Database Replay > Preprocess Captured Workload

Directory: DAN_CAPTURE_DEST

Capture Summary	
Name	only-dan-odtug-capture
Status	Completed
Directory Object	DAN_CAPTURE_DEST
Database Name	ORCL
Capture Database Version	11.1.0.6.0
DBID	1180446881
Capture Error Code	None
Capture Error Message	None
Captured Data Size (MB)	0.02
Duration (hh:mm:ss)	00:11:06
Start Time	May 3, 2008 8:17:20 PM CDT
End Time	May 3, 2008 8:28:26 PM CDT
Start SCN	790636
End SCN	791573
Preprocessed Database Version	N/A

There are no options to specify for preprocessing, so the next screen summarizes the preprocessing information.

ORACLE Enterprise Manager 11g Database Control

Preprocess Captured Workload: Database Version

Database: orcl, Version: 11.1.0.6, Capture Name: only-dan-odtug-capture, Logged In As: SYSTEM

The current database version is 11.1.0.6.

Continue only if you intend to replay the captured workload on a database of the same version.

Next, the wizard moves on to scheduling the preprocessing information.

The screenshot shows the 'Preprocess Captured Workload: Schedule' screen in Oracle Enterprise Manager 11g. At the top, the breadcrumb trail is 'Database Version' > 'Schedule' > 'Review'. The main title is 'Preprocess Captured Workload: Schedule'. Below the title, the following information is displayed: Database: orcl, Version: 11.1.0.6, Capture Name: only-dan-odtug-capture, and Logged In As: SYSTEM. There are 'Cancel', 'Back', 'Step 2 of 3', and 'Next' buttons. The instruction reads: 'Specify the following to schedule the preprocessing job.' The 'Job Parameters' section includes a 'Job Name' field with the value 'dan-preprocess-odtug' and an empty 'Description' field. The 'Start' section has two radio buttons: 'Immediately' (selected) and 'Later'. Below this, a 'Date' field shows 'May 3, 2008' and a 'Time' field shows '8:45:00 AM'. The 'Host Credentials' section includes fields for 'Username' (oracle), 'Password' (masked with asterisks), and 'Confirm Password' (masked with asterisks), along with a 'Save as Preferred Credential' checkbox. At the bottom right, there are 'Cancel', 'Back', 'Step 2 of 3', and 'Next' buttons.

I chose to run this preprocessing immediately, but it could be scheduled for a future time (possibly overnight) if you need to put it in a maintenance window. The final screen is a summary of all preprocessing settings before submitting the job request.

The screenshot shows the 'Preprocess Captured Workload: Review' screen. The breadcrumb trail is 'Database Version' > 'Schedule' > 'Review'. The main title is 'Preprocess Captured Workload: Review'. Below the title, it shows 'Logged In As: SYSTEM' and buttons for 'Cancel', 'Back', 'Step 3 of 3', and 'Submit'. The text states: 'Workload only-dan-odtug-capture will be preprocessed on database 'orcl''. A summary table is displayed:

Job Name	DAN-PREPROCESS-ODTUG
Database	orcl
Preprocessed Database Version	11.1.0.6
Directory Object	DAN_CAPTURE_DEST
Capture Name	only-dan-odtug-capture
Captured Data Size (MB)	0.02
Start Time	Immediately

After submission, you'll see a confirmation with a link to the job status screen.

The screenshot shows the 'Confirmation' screen. The breadcrumb trail is 'Database Instance: orcl >'. The main title is 'Confirmation'. Below the title, it says: 'Job 'DAN-PREPROCESS-ODTUG' to prepare the workload has been created successfully.' There is a 'View Job' link. Below this, a list of instructions is provided: 'Once the job is complete you will need to do the following, if you have not already, prior to replaying the workload on the replay system: 1. Restore the replay database on a test system to match the capture database at the start of the workload capture. 2. Make changes (such as perform an upgrade) to the test system as needed. 3. Copy the workload to the test system.'

Clicking on the View Job link shows the job status screen.

ORACLE Enterprise Manager 11g Database Control

Jobs > Job Run: DAN-PREPROCESS-ODTUG > Execution: orcl

Page Refreshed May 3, 2008 8:51:20 PM CDT [Delete Run] [Create Like] [Edit] [View Definition]

Summary
 The Stop and Suspend operations will wait for the current step to complete. A suspended job can be resumed later, at the next step. [Stop] [Suspend]

Status: **Running**
 Scheduled: **May 3, 2008 8:51:13 PM (UTC-05:00)**
 Started: **May 3, 2008 8:51:13 PM (UTC-05:00)**
 Ended:
 Elapsed Time:
 Notification: **No**

Type: **SQL Script**
 Owner: **SYSTEM**
 Description:
 Parameters: **-S**
 DB Username: **SYSTEM**
 DB Role:
 Username: **oracle**

Targets: [Text Box]
 Status: All [Dropdown]
 [Go]

Expand All | Collapse All

Name	Targets	Status	Started	Ended	Elapsed Time (seconds)
Execution: orcl	orcl	Running	May 3, 2008 8:51:13 PM (UTC-05:00)		6
Step: SQL Script	orcl	Running	May 3, 2008 8:51:19 PM (UTC-05:00)		0

Once the preprocessing job is complete, you are now ready to replay the workload. Remember that you may need to reset the database state or recover it to an earlier point in time. With everything in place, you're ready to navigate to the Database Replay screen and choose task 3, "Replay Workload".

Replay setup screens

The first screen in the replay wizard prompts for the location of the preprocessed capture files by choosing the correct directory object. When a directory object is chosen, the capture summary is shown.

ORACLE Enterprise Manager 11g Database Control

Database Instance: orcl > Database Replay > Replay Workload

Logged in As SYSTEM

The captured workload must have been preprocessed and copied to the replay system. A directory object for the directory with the copied workload must exist in the replay system. [Cancel] [Set Up Replay]

Directory
 Select a directory object that contains a preprocessed workload.
 * Directory Object: DAN_CAPTURE_DEST [Dropdown] [Create Directory Object]

Capture Summary

Name	only-dan-odtug-capture2	Captured Data Size (MB)	0.01
Status	Completed	Duration (hh:mm:ss)	00:08:38
Directory Object	DAN_CAPTURE_DEST ⓘ	Start Time	May 3, 2008 8:54:40 PM CDT
Database Name	ORCL	End Time	May 3, 2008 9:03:18 PM CDT
Capture Database Version	11.1.0.6.0	Start SCN	804716
DBID	1180446881	End SCN	805907
Capture Error Code	None	Preprocessed Database Version	11.1.0.6.0
Capture Error Message	None		

▶ Capture Details

Replay History

Select Name	Status	Duration (hh:mm:ss)	Start Time	End Time	AWR Data Exported
No items found					

Click the "Set Up Replay" and you'll be prompted again with a list of prerequisites that you must acknowledge several tasks that need to be completed to make the replay successful.

ORACLE Enterprise Manager 11g Database Control

Setup Preferences Help Logout Database

Database Instance: orcl > Database Replay > Replay Workload > **Replay Workload: Prerequisites** Logged in As SYSTEM

The following items should be completed before setting up a replay: Cancel Continue

Restore Database
Restore the replay database to match the capture database at the start of the workload capture. A successful workload replay depends on application transactions accessing application data identical to that on the capture system. Common ways to restore application data state include point-in-time recovery, flashback, and import/export.

Perform System Changes
You should make any desired changes to the replay system, including any database or system upgrade, prior to replay. The primary purpose of Database Replay is to test the effect of system changes on a real captured application workload. Therefore, the changes you make, combined with the captured workload, define the test you are conducting.

Resolve References to External Systems
A captured workload may contain references to external systems that may only be meaningful in the capture environment. For example, database links in a captured production environment may reference external production databases that should not be referenced during replay. In such a case, you should modify any external references that could jeopardize the production environment during replay.

Set Up Replay Clients
The Replay Client is a multithreaded program (an executable named wrcl) where each thread submits a captured session's workload. This program is made available as part of the standard Oracle Client as well as the Oracle Instant Client. You should install one or more Replay Clients preferably on systems other than the database host. In addition, each Replay Client must be able to access the directory that holds the preprocessed workload. Refer to system documentation for more information.

OEM shows its additional value by offering additional warning as well as opportunity to update references to external systems.

ORACLE Enterprise Manager 11g Database Control

Setup Preferences Help Logout Database

Database Instance: orcl > Database Replay > Replay Workload > **Replay Workload: Prerequisites** > Logged in As SYSTEM

ATTENTION: Replay should be performed in a COMPLETELY ISOLATED test environment

Warning
A captured workload may contain references to external systems that may only be meaningful in the capture environment. Replaying workload with unresolved references to external systems may cause unexpected problems in the production environment.
A replay should be performed in a COMPLETELY ISOLATED test environment which may include hosts, networks, e-mail servers, storage systems, and other devices. Make sure that all references to external systems have been resolved in the replay environment and replaying a workload will cause no harm to your production environment.

Replay Workload: References to External Systems
References to external systems may cause problems during the replay. Cancel Continue

Use the links below to verify potential references to external systems and modify those that are invalid.

- Database Links - This link takes you outside of the Database Replay process.
- Directory Objects - This link takes you outside of the Database Replay process.
- Streams - This link takes you outside of the Database Replay process.

TIP There may be more references to external systems than those found via the above links.

After updating the external references, continue to the first step in executing the replay, entering a name for the replay.

ORACLE Enterprise Manager 11g Database Control

Setup Preferences Help Logout Database

Choose Initial Options Customize Options Prepare Replay Clients Wait for Client Connections Review

Replay Workload: Choose Initial Options Cancel Step 1 of 5 Next

Database orcl
Capture Name only-dan-odtug-capture2
Logged In As SYSTEM

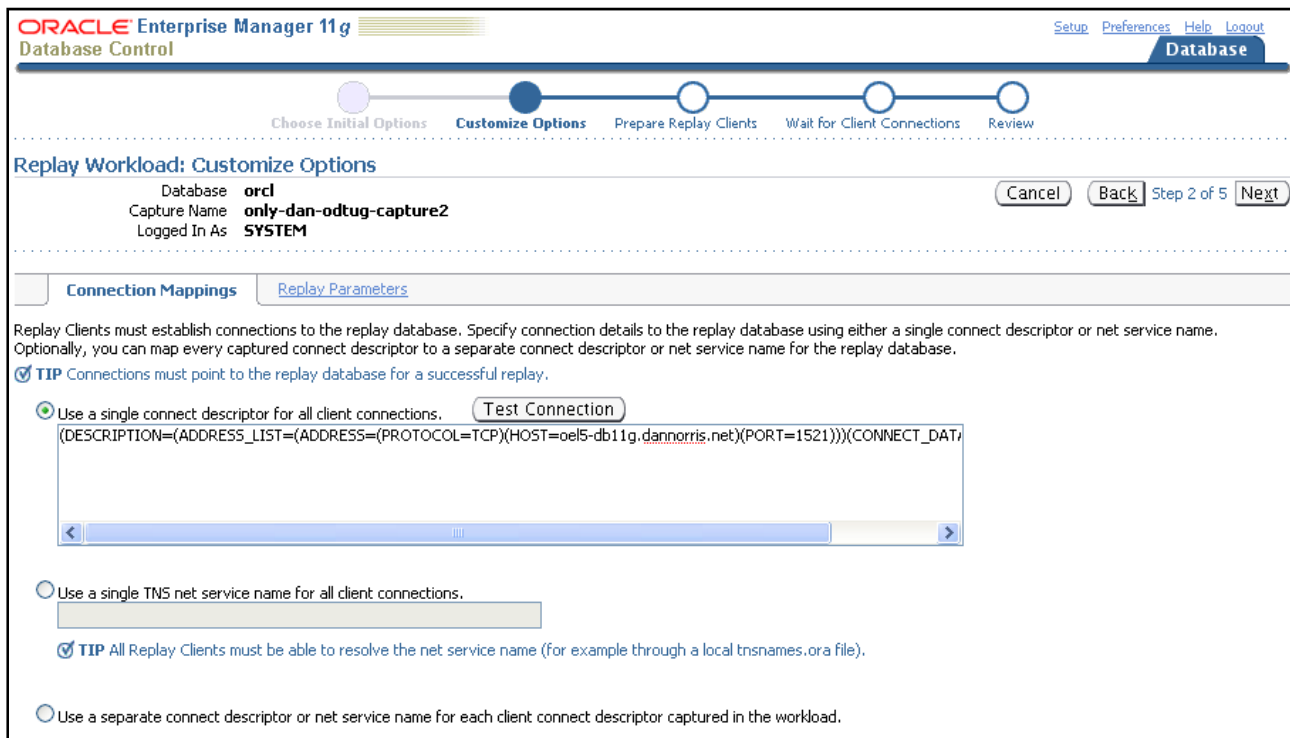
* Replay Name dan-odtug-replay

Initial Options
Choose the initial replay options.

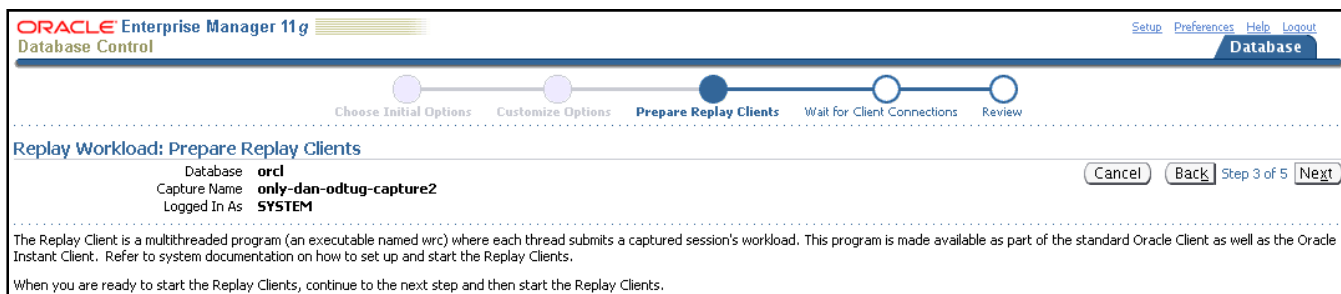
Use the default replay options
 Use replay options from a previous replay

Replay Name

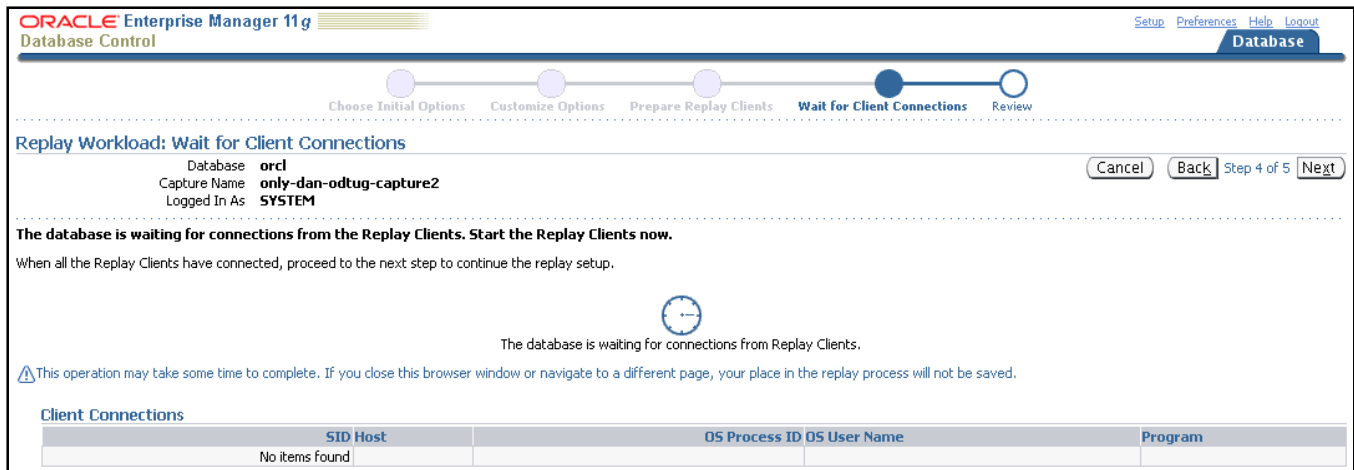
Following the replay name screen, you can customize the connect descriptor used to connect the replay sessions. In my case, I'm keeping the default since I'm using the same database for capture and replay.



After confirming that the connection mapping is correct, the next step is to prepare the replay clients. This screen gives you an opportunity to confirm that you're ready to start replay.



Click next and you'll see the screen where replay waits for clients to connect.

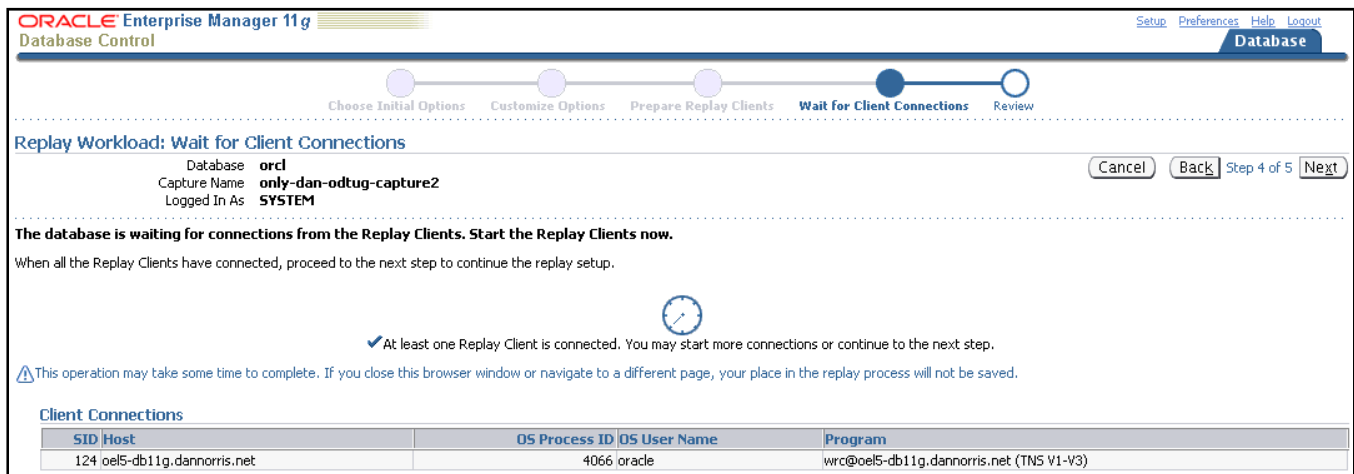


At this point, you need to go to the OS and start a replay client. I did this by running the following command (shown with the output it returns too):

```
[oracle@oel5-db11g ~]$ wrc userid=system password=***** replaydir=/tmp/dan-capture
Workload Replay Client: Release 11.1.0.6.0 - Production on Sun May 4 21:17:42 2008
Copyright (c) 1982, 2007, Oracle. All rights reserved.
```

Wait for the replay to start (21:17:42)

Once you start the replay client, it connects to the database and the OEM screen shows the connection has been made.



Click next to proceed to the summary screen where you can start the replay.

ORACLE Enterprise Manager 11g Database Control

Setup Preferences Help Logout Database

Choose Initial Options Customize Options Prepare Replay Clients Wait for Client Connections **Review**

Replay Workload: Review

Logged In As SYSTEM Cancel Back Step 5 of 5 Submit

Information

Time for resetting the clock: May 3, 2008 8:54:40 PM CDT.

It is recommended that the system time on the database host platform be changed to a value that is close to the capture start time. This must be done just before replay is started. Not doing so might present an invalid data set to the replayed time-sensitive workload, thus causing data divergence. Examples include statements that use the SYSDATE and SYSTIMESTAMP functions. Resetting the time will also minimize job scheduling inconsistencies between replay and capture.

Workload only-dan-odtug-capture2 will be replayed on database 'orcl'.

Database	orcl
Capture Name	only-dan-odtug-capture2
Replay Name	dan-odtug-replay
Directory Object	DAN_CAPTURE_DEST
Connected Replay Clients	1

Click Submit and the replay is started. You're directed to the replay status screen to monitor the replay activity.

ORACLE Enterprise Manager 11g Database Control

Setup Preferences Help Logout Database

Database Instance: orcl > Database Replay >

Logged In As SYSTEM

View Data Real Time: 15 Second Refresh

View Workload Replay: dan-odtug-replay

Page Refreshed May 4, 2008 9:24:03 PM CDT Refresh OK

Status **In Progress** Stop Replay

Summary

Replay Name	dan-odtug-replay	Capture Name	only-dan-odtug-capture2
Directory Object	DAN_CAPTURE_DEST	Duration (hh:mm:ss)	00:02:16
Database Name	ORCL	Prepare Time	May 4, 2008 9:15:25 PM CDT
DBID	1180446881	Start Time	May 4, 2008 9:21:46 PM CDT
Replay Error Code	N/A	End Time	N/A
Replay Error Message	None		

Workload Profile Connection Mappings Replay Parameters Report

Network Time (hh:mm:ss) 00:00:00 Clients 1
Think Time (hh:mm:ss) 00:00:00 Clients Finished 0

Elapsed Time Comparison

Category	Elapsed Time (Minutes)
Capture	8:16
Replay	2:16

Assessing the Replay

The Elapsed Time Comparison chart shows how much time the replayed workload has taken to accomplish the same amount of work as captured.

When the Replay bar is shorter than the Capture bar then the replay environment is processing the workload faster than the capture environment.

The divergence table gives information about both the data and error discrepancies between the replay and capture environments, which can be used as a measure of the replay quality.

View Workload Replay Report

Once you're done with replay, you can proceed to the result analysis utilities. You'll notice that the Report subtab on the screen above will become available once the replay has stopped.

The screenshot shows a web-based interface with a navigation bar at the top containing four tabs: 'Workload Profile', 'Connection Mappings', 'Replay Parameters', and 'Report'. The 'Report' tab is active. Below the navigation bar, there are four distinct report configuration sections, each with a 'Run Report' button.

- Workload Replay Report:** A single 'Run Report' button.
- AWR Compare Period Report:** Two dropdown menus for 'First Workload Capture or Replay' (selected: 'only-dan-odtug-capture2 (May 3, 2008 8:54:40 PM)') and 'Second Workload Capture or Replay' (selected: 'only-dan-odtug-replay (May 4, 2008 9:53:18 PM)'). A 'Run Report' button is below.
- AWR Report:** A dropdown menu for 'Workload Capture or Replay' (selected: 'only-dan-odtug-replay (May 4, 2008 9:53:18 PM)'). A 'Run Report' button is below.
- ASH Report:** A dropdown menu for 'Workload Capture or Replay' (selected: 'only-dan-odtug-replay (May 4, 2008 9:53:18 PM)'). Below this are date and time pickers: 'Start Date' (May 4, 2008), 'End Date' (May 4, 2008), 'Start Time' (9:53 AM), and 'End Time' (9:57 AM). A 'Filter' dropdown is set to 'SID'. A 'Run Report' button is at the bottom.

Replay result analysis screens

The quickest way to get feedback on the affect of the replay is to review the AWR Compare Period report. This will compare an AWR report that covered the capture with an AWR report that covers the replay. You can obtain these AWR reports and compare them yourself at any time, but OEM provides all the comparison information in a single, easy-to-read format. To obtain the comparison, on the Workload Replay screen, once replay has completed, go to the Report subtab (shown above).

When you click the Run Report button, a popup window will be launched to display the comparison report (note that you may need to disable popup blockers to allow this report to display).

WORKLOAD REPOSITORY COMPARE PERIOD REPORT

Snapshot Set	DB Name	DB Id	Instance	Inst num	Release	Cluster	Host	Std Block Size
First (1st)	ORCL	1180446881	orcl	1	11.1.0.6.0	NO	oel5-db11g.dannorris.net	8192
Second (2nd)	ORCL	1180446881	orcl	1	11.1.0.6.0	NO	oel5-db11g.dannorris.net	8192

Snapshot Set	Begin Snap Id	Begin Snap Time	End Snap Id	End Snap Time	Avg Active Users	Elapsed Time (min)	DB time (min)
1st	37	03-May-08 20:54:37 (Sat)	38	03-May-08 21:03:49 (Sat)	0.03	9.19	0.30
2nd	65	04-May-08 21:53:18 (Sun)	66	04-May-08 21:57:56 (Sun)	0.07	4.64	0.33
%Diff					133.33	-49.51	8.87

Host Configuration Comparison

	1st	2nd	Diff	%Diff
Number of CPUs:	1	1	0	0.00
Physical Memory:	1273M	1273M	0M	0.00
Load at Start Snapshot:	.29	.05	-.24	-82.76
Load at End Snapshot:	.3	.27	-.03	-10.00
%User Time:	2.04	4.25	2.22	108.33
%System Time:	6.92	9.84	2.92	42.20
%Idle Time:	90.56	85.21	-5.35	-5.91
%IO Wait Time:	.39	.86	.47	120.51

System Configuration Comparison

	1st	2nd	Diff	%Diff
SGA Target:			0M	0.00
Buffer Cache:	128M	124M	-4M	-3.13
Shared Pool Size:	228M	232M	4M	1.75
Large Pool Size:	4M	4M	0M	0.00
Java Pool Size:	12M	12M	0M	0.00
Streams Pool Size:	8M	8M	0M	0.08
Log Buffer:	5,748K	5,748K	0K	0.00
PGA Aggregate Target:			0M	0.00
Undo Management:	AUTO	AUTO		

Load Profile

	1st per sec	2nd per sec	%Diff	1st per txn	2nd per txn	%Diff
DB time:	0.03	0.07	133.33	0.07	0.16	128.57
CPU time:	0.03	0.06	100.00	0.06	0.15	150.00
Redo size:	7,738.55	6,280.29	-18.84	15,795.96	14,434.71	-8.62

Reviewing the contents of this report will likely offer some good insights on whether the changes made between the capture and replay resulted in a positive or negative overall effect. One quick way to see the high-level effects is to review the comparison of the top wait events:

Top Timed Events

- ◆ Events with a "-" did not make the Top list in this set of snapshots, but are displayed for comparison purposes

1st						2nd					
Event	Wait Class	Waits	Time(s)	Avg Time(ms)	%DB time	Event	Wait Class	Waits	Time(s)	Avg Time(ms)	%DB time
CPU time			15.99		87.48	CPU time			17.57		88.31
os thread startup	Concurrency	9	1.94	215.21	10.60	SGA: allocation forcing compon	Other	4	2.86	714.23	14.36
log file parallel write	System I/O	297	1.93	6.51	10.59	os thread startup	Concurrency	6	1.51	252.10	7.60
control file parallel write	System I/O	195	1.72	8.84	9.43	log file parallel write	System I/O	163	1.35	8.27	6.78
clb file sequential read	User I/O	373	0.97	2.60	5.30	control file parallel write	System I/O	135	1.15	8.54	5.79
-						-clb file sequential read	User I/O	354	1.12	3.18	5.65

There are many other reporting options such as ASH reports, individual AWR reports, and SQL Performance Analyzer tasks as well. With so many reporting options, it's the subject of another session to go into detail of each method. Note that some of these methods may require additional licensing.

Practical uses, save your job

Most of this document has been about the mechanics of how to use database replay as it is a relatively complex set of tasks that are very powerful when executed properly. As replaying a workload also involves many steps, missing part of the process can yield a frustrating round of troubleshooting.

In this section, we will review some of the most common and useful practical applications for Database Replay.

Normal application change rollout

While many discussions about Database Replay focus on the "large tasks," the technology is just as useful for ongoing change management. When new application components, like additional reports, need to be introduced, they can sometimes have negative side effects. By simply adding the new report to a database clone and replaying the production workload, the effect can be measured and interpreted prior to a potential production issue being created.

Tuning Efforts

When performing application tuning, you first need to reproduce the issue. With Database Replay, reproducing the problem is exactly what you will do, easily and reliably. Once you've cloned the production environment and replayed the workload to reproduce the problem, you can then make the changes necessary to address the performance issue, such as creating a new index. New indexes sometimes have negative side effects, so by replaying the entire workload, you'll hopefully see the benefit of the new index. Additionally, you should also be able to identify any negative side effects that might have popped up as well.

Upgrade testing

I think the most obvious use for Database Replay is when applying patches or performing upgrades. These activities are well-known to cause unintended side effects and most sites have a lot of difficulties reproducing realistic workloads to sufficiently test their applications. Employing Database Replay to capture workload on the previous version, a replay on a test upgraded environment will show whether the upgrade will benefit the database without additional effort, or if some additional changes must be made first.

With capture available on Oracle Database 10g Release 2, the upgrade to Oracle Database 11g can be much safer than it might have been otherwise.

Conclusion

Oracle Real Application Testing's Database Replay feature is a great new addition that is almost certain to be useful to environments of all sizes and types. Database Replay's architecture can be scaled to any workload and can be used as an important part of upgrade planning, patching, application releases, and tuning efforts.

Of course, there is always a catch, so before going too much further, let's cover the cost for this great new feature.

Licensing the option

Oracle Real Application Testing is an option on top of the Oracle Database Enterprise Edition license. So, that means that Standard Edition customers are not eligible to use this option. For Oracle Database Enterprise Edition licensees, the Oracle Real Application Testing option has a list price of \$10,000 per CPU or \$200 per Named User Plus. While only half the cost of the Real Application Clusters (RAC) option, it is still expensive and will certainly slow the uptake of this new feature. However, it seems like a small price to pay for the additional stability and change assurance that Real Application Testing provides.

In the wild: is anyone using it yet?

There are customers using Real Application Testing for their production needs today. However, there is a common pattern of many customers waiting to upgrade to a new release until a first patch set is released. For some, the wait is even longer—until the Release 2 of the major release is available. So, we haven't yet seen a lot of Oracle Database 11g upgrades yet, certainly not enough to gauge the overall deployment success for Real Application Testing.

Comparison of RAT vs. Traditional Load Testing

Finally, one must consider the effectiveness of Real Application Testing versus the traditional load testing methods employing a synthetic load via the application. Application servers are not always reliable tools to reproduce a workload the same way each time and even if they can be programmed to reproduce the same workload in the same way each time, it is almost always a synthetic workload that makes no guarantee that it replicates real production workload.

In conclusion, if you're trying to test the application servers' capabilities, the using a load testing tool that exercises the application server is appropriate. However, if you wish to test the database to measure the effect of a change to the database, the Database Replay is the only way to ensure that you've tested your database's real workload from your real production system. Combining Database Replay with flashback database gives you the ability to replay multiple times without having to contemplate a lengthy and laborious restore and recovery process between each replay.

References

- "A look at Real Application Testing from a customer's perspective", Jeremiah Wilton, ORA-600 Consulting
- "Oracle Database 11g: Real Application Testing Overview", Jagan R. Athreya, Oracle
- "Database Replay", "Oracle Database 11g: The Top New Features For DBAs and Developers", Arup Nanda, OTN
- Chapter 22, "Database Replay", in Oracle Database Performance Tuning Guide 11g Release 1 (11.1)
- Oracle Technology Global Price List, April 1, 2008, <http://www.oracle.com/corporate/pricing/technology-price-list.pdf>

From the Lawyers

The information contained herein should be deemed reliable but not guaranteed. The author has made every attempt to provide current and accurate information. If you have any comments or suggestions, please contact the author at [dnorris\(at\)piocon.com](mailto:dnorris(at)piocon.com).