



An Oracle White Paper  
June, 2016

# Oracle Database In-Memory Advisor Usage Examples

## Disclaimer

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Intended Audience .....	2
Introduction .....	2
How to Use This Document .....	2
Scenarios for Running the Advisor .....	3
Running the Advisor on a Live Workload .....	3
Generating Recommendations for Different In-Memory Sizes .....	8
Running the Advisor On AWR Data from a Different Database....	11
Running the Advisor on a Batch Workload.....	25
Running the Advisor on a Pluggable Database.....	26
Running the Advisor with a Customized Script .....	31
Conclusion .....	33

## Intended Audience

Readers are assumed to have hands-on experience with Oracle Database technologies from the perspective of a DBA or performance specialist.

## Introduction

Oracle Database 12.1.0.2 introduced Oracle Database In-Memory allowing a single database to efficiently support mixed analytic and transactional workloads. An Oracle Database configured with Database In-Memory delivers optimal performance for transactions while simultaneously supporting real-time analytics and reporting.

For complete details about Oracle Database In-Memory, see the [Oracle Database In-Memory whitepaper](#) and the [Oracle Database In-Memory Page](#) on oracle.com.

This paper contains examples of how to use the Oracle Database In-Memory Advisor. The Advisor analyzes your workload and makes specific recommendations regarding how to size Oracle Database In-Memory and which objects would render the greatest benefit to your system when placed In-Memory. The Oracle Database In-Memory Advisor can be downloaded from My Oracle Support (MOS) note 1965343.1. MOS note 1965343.1 also contains a whitepaper describing how to install and run the Advisor.

## How to Use This Document

This document is designed to give detailed examples of using the In-Memory Advisor in a variety of scenarios. The easiest way to use this document is to review the table of contents for the scenario that best matches your intended use and then go directly to that section to see an example of using the Advisor in that manner.

## Scenarios for Running the Advisor

### Running the Advisor on a Live Workload

The Oracle Database In-Memory Advisor was designed to be easily executed on a database. Below are the steps to run the Advisor on a workload currently executing on your system.

These steps can also be used to run the Advisor on any time frame within the Automatic Workload Repository (AWR) retention period on your system. By default, AWR data is retained for eight days. Therefore, if you have the default AWR configuration on your system, you can use these steps to run the Advisor for any time period in the last eight days.

The high level steps are:

- Run SQLPLUS as an appropriately privileged user such as SYSTEM or a user who has been granted ADVISOR privilege
- Execute the imadvisor\_recommendations.sql script
  - Supply a task name
  - If you know the amount of memory you'd like to use, enter it at the prompt, or hit <ENTER> to let the Advisor give a list of sizes and estimated benefits
  - Enter the time range for analysis
  - If not specified at the earlier prompt, select a size from the table of estimated sizes and benefits
- Review the reports

Below is an example that session would look like. User input is highlighted in yellow.

```
$ sqlplus / as sysdba
```

```
SQL*Plus: Release 12.1.0.2.0 Production on Tue May 31 17:14:15 2016  
Copyright (c) 1982, 2014, Oracle. All rights reserved.
```

```
Connected to:
```

```
Oracle Database 12c Enterprise Edition Release 12.1.0.2.0 - 64bit Production  
With the Partitioning, OLAP, Advanced Analytics and Real Application Testing  
options
```

```
SQL> @imadvisor_recommendations
```

This script creates and runs an In-Memory Advisor task that analyzes your workload to determine an optimal In-Memory configuration.

This script then generates an HTML recommendation report file in the current working directory: imadvisor\_<task\_name>.html

This script also generates a sqlplus DDL script to implement the recommendations: imadvisor\_<task\_name>.sql

NOTE: Once you have existing tasks, you can use this script again with a task that has already gathered and analyzed statistics to optimize for a different In-Memory size.

Default task\_name (new task): im\_advisor\_task\_20160531171426  
Enter value for task\_name: TASK01

Advisor task name specified: TASK01

New Advisor task will be named: TASK01...

Analyzing and reporting on a live workload on this database  
(DBID=2026721262)...

The In-Memory Advisor optimizes the In-Memory configuration for a specific In-Memory size that you choose.

After analysis, the In-Memory Advisor can provide you a list of performance benefit estimates for a range of In-Memory sizes. You may then choose the In-Memory size for which you wish to optimize.

If you already know the specific In-Memory size you wish, please enter the value now. Format: nnnnnn[KB|MB|GB|TB]

Or press <ENTER> to get performance estimates first.  
Enter value for inmemory\_size:

The In-Memory Advisor will display performance benefit estimates after analysis.

Enter begin time for report:

```
-- Valid input formats:
--   To specify absolute begin time:
--     [MM/DD[/YY]] HH24:MI[:SS]
--     Examples: 02/23/03 14:30:15
```

```
--          02/23 14:30:15
--          14:30:15
--          14:30
--      To specify relative begin time: (start with '-' sign)
--      -[HH24:]MI
--      Examples: -1:15  (SYSDATE - 1 Hr 15 Mins)
--              -25    (SYSDATE - 25 Mins)
```

Default begin time: -60

Enter value for begin\_time: -60

Report begin time specified: -60

Enter duration in minutes starting from begin time:  
(defaults to SYSDATE - begin\_time)

Enter value for duration: 60

Report duration specified: 60

Using 2016-MAY-31 16:15:08.000000000 as report begin time

Using 2016-MAY-31 17:15:08.000000000 as report end time

In-Memory Advisor: Adding statistics...

In-Memory Advisor: Finished adding statistics.

In-Memory Advisor: Analyzing statistics...

In-Memory Advisor: Finished analyzing statistics.

The Advisor estimates the following performance benefits:

IN-MEMORY SIZE	PERCENTAGE OF MAXIMUM SGA SIZE	ESTIMATED ANALYTICS PROCESSING TIME REDUCTION (SECONDS) *	ESTIMATED ANALYTICS PROCESSING PERFORMANCE IMPROVEMENT FACTOR*
1.074GB	97	837	2.8X
1.021GB	92	778	2.5X
990.3MB	87	778	2.5X
935.2MB	82	347	1.4X
880.2MB	77	347	1.4X
825.2MB	73	347	1.4X
770.2MB	68	347	1.4X
715.2MB	63	347	1.4X
660.2MB	58	347	1.4X
605.2MB	53	347	1.4X
550.1MB	48	347	1.4X
495.1MB	44	347	1.4X
440.1MB	39	347	1.4X
385.1MB	34	347	1.4X
330.1MB	29	347	1.4X
275.1MB	24	347	1.4X
220.1MB	19	154	1.1X
165.0MB	15	154	1.1X
110.0MB	10	154	1.1X
55.01MB	5	105	1.1X

\*Estimates: The In-Memory Advisor's estimates are useful for making In-Memory decisions. But they are not precise. Due to performance variations caused by workload diversity, the Advisor's performance estimates are conservatively limited to no more than 10.0X faster.

Choose the In-Memory size you wish for optimization (default=1.074GB):

1.074GB

The Advisor is optimizing for an In-Memory size of 1.074GB...

Fetching recommendation files for task: TASK01

Placing recommendation files in: the current working directory

Fetches file: imadvisor\_TASK01.html

Purpose: recommendation report primary html page

Fetches file: imadvisor\_TASK01.sql

Purpose: recommendation DDL sqlplus script

You can re-run this task with this script and specify a different an In-Memory size. Re-running a task to optimize for a different In-Memory size is faster than creatng and running a new task from scratch.

SQL>

## Generating Recommendations for Different In-Memory Sizes

When you run the In-Memory Advisor on a specific time range, the Advisor analyzes and stores recommendations for a range of possible In-Memory sizes. You can produce a report of recommendations for a different In-Memory size for the same workload by specifying an existing analysis task. This requires significantly fewer resources than rerunning the full analysis:

The high level steps are:

- Run SQLPLUS as an appropriately privileged user such as SYSTEM or a user who has been granted ADVISOR privilege
- Execute the `imadvisor_recommendations.sql` script
  - Choose an existing task name
  - Enter the new amount of memory for Database In-Memory for the Advisor to report on.
  - Review the reports

Below is what that session would look like. User input is highlighted in yellow.

```
$ sqlplus / as sysdba
```

```
SQL*Plus: Release 12.1.0.2.0 Production on Tue May 31 17:31:34 2016  
Copyright (c) 1982, 2014, Oracle. All rights reserved.
```

```
Connected to:
```

```
Oracle Database 12c Enterprise Edition Release 12.1.0.2.0 - 64bit Production  
With the Partitioning, OLAP, Advanced Analytics and Real Application Testing  
options
```

```
SQL> @imadvisor_recommendations
```

This script creates and runs an In-Memory Advisor task that analyzes your workload to determine an optimal In-Memory configuration.

This script then generates an HTML recommendation report file in the current working directory: `imadvisor_<task_name>.html`

This script also generates a sqlplus DDL script to implement the recommendations: `imadvisor_<task_name>.sql`

NOTE: You may specify one of your existing tasks if you wish to optimize for a different In-Memory size.

Using an existing, executed task is faster than a new task since a new task requires statistics gathering and analysis.

But if you wish to analyze a different workload or use a different statistics capture window or add a SQLSET, you must specify a new task.

The following is a list of your existing tasks:

```
TASK_NAME                DATE_CREATED
-----
TASK01                   2016-MAY-31 17:15:13
```

```
Default task_name (new task): im_advisor_task_20160531173140
Enter value for task_name: TASK01
```

Advisor task name specified: TASK01

Analyzing and reporting on a live workload on this database  
(DBID=2026721262)...

The Advisor estimates the following performance benefits:

		ESTIMATED	ESTIMATED
		ANALYTICS	ANALYTICS
		PROCESSING	PROCESSING
	PERCENTAGE	TIME	PERFORMANCE
IN-MEMORY	OF MAXIMUM	REDUCTION	IMPROVEMENT
SIZE	SGA SIZE	(SECONDS) *	FACTOR*
1.074GB	97	837	2.8X
1.021GB	92	778	2.5X
990.3MB	87	778	2.5X
935.2MB	82	347	1.4X
880.2MB	77	347	1.4X
825.2MB	73	347	1.4X
770.2MB	68	347	1.4X
715.2MB	63	347	1.4X
660.2MB	58	347	1.4X

605.2MB	53	347	1.4X
550.1MB	48	347	1.4X
495.1MB	44	347	1.4X
440.1MB	39	347	1.4X
385.1MB	34	347	1.4X
330.1MB	29	347	1.4X
275.1MB	24	347	1.4X
220.1MB	19	154	1.1X
165.0MB	15	154	1.1X
110.0MB	10	154	1.1X
55.01MB	5	105	1.1X
*Estimates: The In-Memory Advisor's estimates are useful for making In-Memory decisions. But they are not precise. Due to performance variations caused by workload diversity, the Advisor's performance estimates are conservatively limited to no more than 10.0X faster.			

Choose the In-Memory size you wish for optimization (default=1.074GB):

990.3MB

The Advisor is optimizing for an In-Memory size of 990.3MB...

Fetching recommendation files for task: TASK01

Placing recommendation files in: the current working directory

Fetches file: imadvisor\_TASK01.html

Purpose: recommendation report primary html page

Fetches file: imadvisor\_TASK01.sql

Purpose: recommendation DDL sqlplus script

You can re-run this task with this script and specify a different an In-Memory

size. Re-running a task to optimize for a different In-Memory size is faster than creatng and running a new task from scratch.

SQL>

## Running the Advisor On AWR Data from a Different Database

You can run the Advisor on a workload from another database by exporting the Automatic Workload Data (AWR) along with some supplemental information from the production database and importing it into another database. You might want to do this to lessen the impact on the production database or to preserve the AWR data for longer than the default.

In addition to the AWR data, the Advisor needs some data dictionary information to complete its analysis. A script is provided to export and then import the additional required information. The additional export script should be run immediately after running the AWR export script.

The high level steps to do this are:

- Run SQLPLUS as an appropriately privileged user such as SYSTEM or a user who has been granted ADVISOR privilege
- Execute the AWR extract script: \$ORACLE\_HOME/rdbms/admin/awrextr.sql
- Execute the additional data 'augment' export script : imadvisor\_awr\_augment\_export.sql
- Copy the AWR export and augment export dump files to the system with the database that will be used to run the Advisor
- Install the Advisor on the target database
- Import the AWR data using: \$ORACLE\_HOME/rdbms/admin/awrload.sql (Note: you can only import AWR data into the same or higher version of the Oracle database.)
- Import the AWR augment data using the import script: imadvisor\_augment\_import.sql
- Run the Advisor as described in prior use cases

Below is what that session would look like. User input is highlighted in yellow.

### Export the data

```
$ sqlplus / as sysdba
SQL*Plus: Release 12.1.0.2.0 Production on Wed Jun 1 16:04:08 2016
Copyright (c) 1982, 2014, Oracle. All rights reserved.
```

Connected to:

```
Oracle Database 12c Enterprise Edition Release 12.1.0.2.0 - 64bit Production
With the Partitioning, OLAP, Advanced Analytics and Real Application Testing
options
```

```
SQL> set pagesize 100
SQL> @?/rdbms/admin/awrextr.sql
~~~~~
AWR EXTRACT
```

```

~~~~~
~~~~~
~ This script will extract the AWR data for a range of snapshots ~
~ into a dump file. The script will prompt users for the ~
~ following information: ~
~ (1) database id ~
~ (2) snapshot range to extract ~
~ (3) name of directory object ~
~ (4) name of dump file ~
~~~~~

```

Databases in this Workload Repository schema

```
~~~~~
```

DB Id	DB Name	Host
* 2026721262	O120250	node50
957820695	O120250	node50

The default database id is the local one: '2026721262'. To use this database id, press <return> to continue, otherwise enter an alternative.

Enter value for dbid:

Using 2026721262 for Database ID

Specify the number of days of snapshots to choose from

```
~~~~~
```

Entering the number of days (n) will result in the most recent (n) days of snapshots being listed. Pressing <return> without specifying a number lists all completed snapshots.

Enter value for num\_days: 1

Listing the last day's Completed Snapshots

DB Name	Snap Id	Snap Started
O120250	1785	01 Jun 2016 11:00
	1786	01 Jun 2016 12:00
	1787	01 Jun 2016 13:00

```

1788 01 Jun 2016 14:00
1789 01 Jun 2016 15:00
1790 01 Jun 2016 16:00

```

Specify the Begin and End Snapshot Ids

~~~~~

Enter value for begin\_snap: 1785

Begin Snapshot Id specified: 1785

Enter value for end\_snap: 1790

End Snapshot Id specified: 1790

Specify the Directory Name

~~~~~

Directory Name	Directory Path
ADVISORDIR	/tmp
DATA_PUMP_DIR	/u01/oracle/admin/o120250/dpdump/
OPATCH_INST_DIR	/u01/oracle/product/12.1.0/dbhome_1/OPatch
OPATCH_LOG_DIR	/u01/oracle/product/12.1.0/dbhome_1/QOpatch
OPATCH_SCRIPT_DIR	/u01/oracle/product/12.1.0/dbhome_1/QOpatch
ORACLE_BASE	/
ORACLE_HOME	/
ORACLE_OCM_CONFIG_DIR2	/u01/oracle/product/12.1.0/dbhome_1/ccr/state
XMLDIR	/u01/oracle/product/12.1.0/dbhome_1/rdbms/xml

Choose a Directory Name from the above list (case-sensitive).

Enter value for directory\_name: ADVISORDIR

Using the dump directory: ADVISORDIR

Specify the Name of the Extract Dump File

~~~~~

The prefix for the default dump file name is awrdat\_1785\_1790.

To use this name, press <return> to continue, otherwise enter an alternative.

Enter value for file\_name:

Using the dump file prefix: awrdat\_1785\_1790

```
|
| ~~~~~
| The AWR extract dump file will be located
| in the following directory/file:
|   /tmp
|   awrdat_1785_1790.dmp
| ~~~~~
|
| *** AWR Extract Started ...
|
| This operation will take a few moments. The
| progress of the AWR extract operation can be
| monitored in the following directory/file:
|   /tmp
|   awrdat_1785_1790.log
|
```

End of AWR Extract

SQL>

SQL>

SQL> @imadvisor\_awr\_augment\_export.sql

```
*****
* This script will create for you an Automatic Workload Repository (AWR)
* augment that will supply additional data required by the Oracle Database
* In-Memory Advisor.  Without this data augment, you cannot use an AWR export
* with the In-Memory Advisor.  (But with it, you can!)
*
* Note: While the AWR on this database holds data imported from other
* databases, only AWR data from this local database (DBID=2026721262) can
* be used with the AWR augment you are about to create.
*
* Please make sure your AWR export is for DBID=2026721262.
*
* Also note: It is best to capture an AWR augment in the same timeframe
* (preferably after) the corresponding AWR export.
*****
```

```

ADVISORDIR                /tmp
DATA_PUMP_DIR              /u01/oracle/admin/o120250/dpdump/
OPATCH_INST_DIR           /u01/oracle/product/12.1.0/dbhome_1/OPatch
OPATCH_LOG_DIR            /u01/oracle/product/12.1.0/dbhome_1/QOpatch
OPATCH_SCRIPT_DIR         /u01/oracle/product/12.1.0/dbhome_1/QOpatch
ORACLE_BASE                /
ORACLE_HOME                /
ORACLE_OCM_CONFIG_DIR2    /u01/oracle/product/12.1.0/dbhome_1/ccr/state
XMLDIR                     /u01/oracle/product/12.1.0/dbhome_1/rdbms/xml

```

Please enter the Oracle directory object to use for export  
(default=DATA\_PUMP\_DIR)?

**ADVISORDIR**

Using directory ADVISORDIR...

The default IM Advisor AWR augment dump file name prefix is  
imadvisor\_awr\_augment.

Please press <return> to use this name prefix; otherwise, enter an  
alternative name prefix?

Using "imadvisor\_awr\_augment" as the AWR augment dump file name prefix...

Setting up the AWR agument staging schema...

No errors.

```

old 931:  dmp_name := NVL('&&dump_file_name_prefix', :dmp_name_default);
new 931:  dmp_name := NVL('imadvisor_awr_augment', :dmp_name_default);

```

PL/SQL procedure successfully completed.

Exporting AWR agument data...

IMADVISOR\_AWR\_AUGMENT\_EXPORT Data Pump status: SUCCESS

PL/SQL procedure successfully completed.

Dropping the AWR agument staging schema...

Directory path for Data Pump dump and log files: /tmp

Data Pump dump file: imadvisor\_awr\_augment.dmp

Data Pump log file: imadvisor\_awr\_augment\_export.log

PL/SQL procedure successfully completed.

All done.

SQL>

### Copy the data

```
$ scp /location_of_dump_files/*.dmp username@differenthost:/targetdir
```

### Import the data

```
$ sqlplus / as sysdba
```

```
SQL*Plus: Release 12.1.0.2.0 Production on Thu Jun 2 07:18:24 2016
```

```
Copyright (c) 1982, 2014, Oracle. All rights reserved.
```

```
Connected to:
```

```
Oracle Database 12c Enterprise Edition Release 12.1.0.2.0 - 64bit Production
With the Partitioning, OLAP, Advanced Analytics and Real Application Testing
options
```

```
SQL> set pagesize 100
```

```
SQL> @?/rdbms/admin/awrload.sql
```

```
~~~~~
```

```
AWR LOAD
```

```
~~~~~
```

```
~~~~~
```

```
~ This script will load the AWR data from a dump file. The ~
```

```
~ script will prompt users for the following information: ~
```

```
~ (1) name of directory object ~
```

```
~ (2) name of dump file ~
```

```
~ (3) staging schema name to load AWR data into ~
```

```
~~~~~
```

```
Specify the Directory Name
```

```
~~~~~
```

| Directory Name    | Directory Path                              |
|-------------------|---------------------------------------------|
| DATA_PUMP_DIR     | /u01/oracle/admin/o120249/dpdump/           |
| DUMPDIR           | /scratch/oracle/imadvisor                   |
| OPATCH_INST_DIR   | /u01/oracle/product/12.1.0/dbhome_1/OPatch  |
| OPATCH_LOG_DIR    | /u01/oracle/product/12.1.0/dbhome_1/QOpatch |
| OPATCH_SCRIPT_DIR | /u01/oracle/product/12.1.0/dbhome_1/QOpatch |

```
ORACLE_BASE           /
ORACLE_HOME           /

SUBDIR                /u01/oracle/product/12.1.0/dbhome_1/demo/schema

XMLDIR                /u01/oracle/product/12.1.0/dbhome_1/rdbms/xml
XSDDIR                /u01/product/12.1.0/dbhome_1/xml/schema
```

Choose a Directory Name from the list above (case-sensitive).

Enter value for directory\_name: DATA\_PUMP\_DIR

Using the dump directory: DATA\_PUMP\_DIR

Specify the Name of the Dump File to Load

~~~~~

Please specify the prefix of the dump file (.dmp) to load:

Enter value for file\_name: awrdat\_1785\_1790

Loading from the file name: awrdat\_1785\_1790.dmp

Staging Schema to Load AWR Snapshot Data

~~~~~

The next step is to create the staging schema where the AWR snapshot data will be loaded. After loading the data into the staging schema, the data will be transferred into the AWR tables in the SYS schema.

The default staging schema name is AWR\_STAGE.

To use this name, press <return> to continue, otherwise enter an alternative.

Enter value for schema\_name:

Using the staging schema name: AWR\_STAGE

Choose the Default tablespace for the AWR\_STAGE user

~~~~~

Choose the AWR\_STAGE users's default tablespace. This is the

tablespace in which the AWR data will be staged.

| TABLESPACE_NAME | CONTENTS  | DEFAULT TABLESPACE |
|-----------------|-----------|--------------------|
| EXAMPLE         | PERMANENT |                    |
| SYSAUX          | PERMANENT | *                  |
| USERS           | PERMANENT |                    |

Pressing <return> will result in the recommended default tablespace (identified by \*) being used.

Enter value for default\_tablespace: **USERS**

Using tablespace USERS as the default tablespace for the AWR\_STAGE

Choose the Temporary tablespace for the AWR\_STAGE user  
 ~~~~~  
 Choose the AWR\_STAGE user's temporary tablespace.

| TABLESPACE_NAME | CONTENTS  | DEFAULT TEMP TABLESPACE |
|-----------------|-----------|-------------------------|
| TEMP            | TEMPORARY | *                       |

Pressing <return> will result in the database's default temporary tablespace (identified by \*) being used.

Enter value for temporary\_tablespace:

Using tablespace TEMP as the temporary tablespace for AWR\_STAGE

... Creating AWR\_STAGE user

```
|
| ~~~~~
| Loading the AWR data from the following
| directory/file:
| /scratch/oracle/imadvisor
| awrdat_1785_1790.dmp
| ~~~~~
|
| *** AWR Load Started ...
```

```
|
| This operation will take a few moments. The
| progress of the AWR load operation can be
| monitored in the following directory/file:
| /scratch/oracle/imadvisor
| awrdat_1785_1790.log
|
... Dropping AWR_STAGE user
```

End of AWR Load

SQL>

SQL>

SQL>

SQL> @imadvisor\_awr\_augment\_import.sql

```
DATA_PUMP_DIR          /u01/oracle/admin/o120249/dpdump/
DUMPDIR                /scratch/oracle/imadvisor
LOG_FILE_DIR           /u01/product/12.1.0/dbhome_1/demo/schema/log/
OPATCH_INST_DIR       /u01/oracle/product/12.1.0/dbhome_1/OPatch
OPATCH_LOG_DIR        /u01/oracle/product/12.1.0/dbhome_1/QOpatch
OPATCH_SCRIPT_DIR     /u01/oracle/product/12.1.0/dbhome_1/QOpatch
ORACLE_BASE            /
ORACLE_HOME            /
SUBDIR                 /u01/product/12.1.0/dbhome_1/demo/schema
XMLDIR                 /u01/oracle/product/12.1.0/dbhome_1/rdbms/xml
```

Please enter the Oracle directory object to use for import  
(default=DATA\_PUMP\_DIR)?

DATA\_PUMP\_DIR

Using directory DATA\_PUMP\_DIR...

The default IM Advisor AWR augment dump file name prefix is  
imadvisor\_awr\_augment.

Please press <return> to use this name prefix; otherwise, please enter an  
alternative name prefix?

Using "imadvisor\_awr\_augment" as the AWR augment dump file name prefix...

Gathering information about the AWR argument...

IMADVISOR\_AWR\_AUGMENT\_MASTER Data Pump status: SUCCESS

Setting up the AWR agument schema...

Importing AWR agument data...

IMADVISOR\_AWR\_AUGMENT\_IMPORT Data Pump status: SUCCESS

Granting access on the AWR augment schema to DBMS\_INMEMORY\_ADVISOR...

All done.

SQL> exit

### Run the Advisor on the imported data

```
$ sqlplus / as sysdba
```

```
SQL*Plus: Release 12.1.0.2.0 Production on Thu Jun 2 07:23:45 2016
```

```
Copyright (c) 1982, 2014, Oracle. All rights reserved.
```

```
Connected to:
```

```
Oracle Database 12c Enterprise Edition Release 12.1.0.2.0 - 64bit Production  
With the Partitioning, OLAP, Advanced Analytics and Real Application Testing  
options
```

```
SQL> @imadvisor_recommendations
```

This script creates and runs an In-Memory Advisor task that analyzes your workload to determine an optimal In-Memory configuration.

This script then generates an HTML recommendation report file in the current working directory: imadvisor\_<task\_name>.html

This script also generates a sqlplus DDL script to implement the recommendations: imadvisor\_<task\_name>.sql

NOTE: Once you have existing tasks, you can use this script again with a task that has already gathered and analyzed statistics to optimize for a different In-Memory size.

```
Default task_name (new task): im_advisor_task_20160602072358
```

```
Enter value for task_name: TASK02
```

```
Advisor task name specified: TASK02
```

New Advisor task will be named: TASK02...

By default, the Advisor runs against a live workload on this database.  
This database also has imported, augmented AWR workloads.

Press ENTER or respond NO to run against a live workload.  
Respond YES to run against an augmented AWR workload.

Enter value for run\_against\_augmented\_awr: YES

The Advisor can use the following augmented AWR imports:

Augmented AWR Import DBID

-----

2026721262

Enter value for dbid: 2026721262

Analyzing and reporting on an augmented AWR workload with DBID=2026721262...

The In-Memory Advisor optimizes the In-Memory configuration for a specific  
In-Memory size that you choose.

After analysis, the In-Memory Advisor can provide you a list of performance  
benefit estimates for a range of In-Memory sizes. You may then choose the  
In-Memory size for which you wish to optimize.

If you already know the specific In-Memory size you wish, please enter  
the value now. Format: nnnnnnn[KB|MB|GB|TB]

Or press <ENTER> to get performance estimates first.

Enter value for inmemory\_size:

The In-Memory Advisor will display performance benefit estimates after  
analysis.

Enter begin time for report:

```
-- Valid input formats:
--   To specify absolute begin time:
--     [MM/DD[/YY]] HH24:MI[:SS]
--     Examples: 02/23/03 14:30:15
--                02/23 14:30:15
--                14:30:15
--                14:30
--   To specify relative begin time: (start with '-' sign)
--     -[HH24:]MI
--     Examples: -1:15 (SYSDATE - 1 Hr 15 Mins)
--                -25 (SYSDATE - 25 Mins)
```

Default begin time: 06/01/16 10:00:25

Enter value for begin\_time:

Report begin time specified:

Enter duration in minutes starting from begin time:  
(defaults to <latest-snapshot-end-time> - begin\_time)

Enter value for duration:

Report duration specified:

Using 2016-JUN-01 10:00:25.000000000 as report begin time

Using 2016-JUN-01 16:00:16.000000000 as report end time

In-Memory Advisor: Adding statistics...

In-Memory Advisor: Finished adding statistics.

In-Memory Advisor: Analyzing statistics...

In-Memory Advisor: Finished analyzing statistics.

The Advisor estimates the following performance benefits:

| IN-MEMORY<br>SIZE | PERCENTAGE<br>OF MAXIMUM<br>SGA SIZE | ESTIMATED<br>ANALYTICS<br>PROCESSING<br>TIME<br>REDUCTION<br>(SECONDS) * | ESTIMATED<br>ANALYTICS<br>PROCESSING<br>PERFORMANCE<br>IMPROVEMENT<br>FACTOR* |
|-------------------|--------------------------------------|--------------------------------------------------------------------------|-------------------------------------------------------------------------------|
| 1.074GB           | 97                                   | 6164                                                                     | 2.4X                                                                          |
| 1.021GB           | 92                                   | 5946                                                                     | 2.3X                                                                          |
| 990.3MB           | 87                                   | 5946                                                                     | 2.3X                                                                          |
| 935.2MB           | 82                                   | 2970                                                                     | 1.4X                                                                          |
| 880.2MB           | 77                                   | 2970                                                                     | 1.4X                                                                          |
| 825.2MB           | 73                                   | 2970                                                                     | 1.4X                                                                          |
| 770.2MB           | 68                                   | 2970                                                                     | 1.4X                                                                          |
| 715.2MB           | 63                                   | 2970                                                                     | 1.4X                                                                          |
| 660.2MB           | 58                                   | 2970                                                                     | 1.4X                                                                          |
| 605.2MB           | 53                                   | 2970                                                                     | 1.4X                                                                          |
| 550.1MB           | 48                                   | 2970                                                                     | 1.4X                                                                          |
| 495.1MB           | 44                                   | 2970                                                                     | 1.4X                                                                          |
| 440.1MB           | 39                                   | 2970                                                                     | 1.4X                                                                          |
| 385.1MB           | 34                                   | 2970                                                                     | 1.4X                                                                          |
| 330.1MB           | 29                                   | 2970                                                                     | 1.4X                                                                          |
| 275.1MB           | 24                                   | 2970                                                                     | 1.4X                                                                          |
| 220.1MB           | 19                                   | 1317                                                                     | 1.1X                                                                          |
| 165.0MB           | 15                                   | 1317                                                                     | 1.1X                                                                          |
| 110.0MB           | 10                                   | 1317                                                                     | 1.1X                                                                          |
| 55.01MB           | 5                                    | 1010                                                                     | 1.1X                                                                          |

\*Estimates: The In-Memory Advisor's estimates are useful for making In-Memory decisions. But they are not precise. Due to performance variations caused by workload diversity, the Advisor's performance estimates are conservatively limited to no more than 10.0X faster.

Choose the In-Memory size you wish for optimization (default=1.074GB):  
1.074GB

The Advisor is optimizing for an In-Memory size of 1.074GB...  
Fetching recommendation files for task: TASK02  
Placing recommendation files in: the current working directory

Fetches file: imadvisor\_TASK02.html  
Purpose: recommendation report primary html page

Fetches file: imadvisor\_TASK02.sql  
Purpose: recommendation DDL sqlplus script

You can re-run this task with this script and specify a different an In-Memory size. Re-running a task to optimize for a different In-Memory size is faster than creatng and running a new task from scratch.

SQL>

## Running the Advisor on a Batch Workload

Oracle Database In-Memory optimizes analytical workload. Analytical workload tends to be long running SQL which scans a large amount of data. It's common to have batch operations that run to produce reports regularly. Oracle Database In-Memory is well suited for this type of workload.

The In-Memory Advisor can run efficiently if its input time period is constrained to batch reporting workload executions. Running the Advisor for a batch reporting workload is the same as running the Advisor on a live (currently executing workload). You just need to specify the time period that you wish the Advisor to analyze when prompted for the begin time and duration.

The high level steps are:

- Run SQLPLUS as an appropriately privileged user such as SYSTEM or a user who has been granted ADVISOR privilege
- Execute the imadvisor\_recommendations.sql script
  - Supply a task name
  - If you know the amount of memory you'd like to use, enter it at the prompt, or hit RETURN to let the Advisor give a list of sizes and estimated benefits
  - Enter the time range to run the analysis for the batch reporting workload
  - Select a size from the table of estimated sizes and benefits
- Review the reports

## Running the Advisor on a Pluggable Database

The Oracle Database In-Memory Advisor supports multitenant databases which were introduced with Oracle Database 12.1.

One way to run the In-Memory Advisor on multitenant databases is to install the Advisor in the root container, the CDB. Then you can run the Advisor on the root container and specify one of the pluggable databases – PDBs.

The high level steps are:

- Install the In-Memory Advisor in the CDB\$ROOT of the database

Run SQLPLUS with a connection to the root container as an appropriately privileged user such as SYSTEM or a user who has been granted ADVISOR privilege

- Execute the imadvisor\_recommendations.sql script
  - Supply a task name
  - Supply a pluggable database name, PDB, to be analyzed
  - If you know the amount of memory you'd like to use, enter it at the prompt, or hit RETURN to let the Advisor give a list of sizes and estimated benefits
  - Enter the time range for analysis
  - Select a size from the table of estimated sizes and benefits
- Review the reports

Below is an example that session would look like. User input is highlighted in yellow.

```
$ sqlplus / as sysdba
```

```
SQL*Plus: Release 12.2.0.0.3 Production on Sat Jun 4 14:57:05 2016  
Copyright (c) 1982, 2016, Oracle. All rights reserved.
```

```
Connected to:
```

```
Oracle Database 12c Enterprise Edition Release 12.2.0.1.0 - 64bit Production
```

```
SQL> @imadvisor_recommendations
```

This script creates and runs an In-Memory Advisor task that analyzes your workload to determine an optimal In-Memory configuration.

This script then generates an HTML recommendation report file in the current working directory: imadvisor\_<task\_name>.html

This script also generates a sqlplus DDL script to implement the recommendations: imadvisor\_<task\_name>.sql

NOTE: Once you have existing tasks, you can use this script again with a task that has already gathered and analyzed statistics to optimize for a different In-Memory size.

Default task\_name (new task): im\_advisor\_task\_20160604145711

Enter value for task\_name: TASK05

Advisor task name specified: TASK05

New Advisor task will be named: TASK05...

Analyzing and reporting on a live workload on this database (DBID=1049723161)...

Enter value for pdb\_name: SALES\_NA

SALES\_NA

The In-Memory Advisor optimizes the In-Memory configuration for a specific In-Memory size that you choose.

After analysis, the In-Memory Advisor can provide you a list of performance benefit estimates for a range of In-Memory sizes. You may then choose the In-Memory size for which you wish to optimize.

If you already know the specific In-Memory size you wish, please enter the value now. Format: nnnnnnn[KB|MB|GB|TB]

Or press <ENTER> to get performance estimates first.

Enter value for inmemory\_size:

The In-Memory Advisor will display performance benefit estimates after analysis.

Enter begin time for report:

-- Valid input formats:

-- To specify absolute begin time:

```
--      [MM/DD[/YY]] HH24:MI[:SS]
--      Examples: 02/23/03 14:30:15
--                02/23 14:30:15
--                14:30:15
--                14:30
--      To specify relative begin time: (start with '-' sign)
--      -[HH24:]MI
--      Examples: -1:15  (SYSDATE - 1 Hr 15 Mins)
--                -25   (SYSDATE - 25 Mins)
```

Default begin time: -60

Enter value for begin\_time:

Report begin time specified:

Enter duration in minutes starting from begin time:  
(defaults to SYSDATE - begin\_time)

Enter value for duration:

Report duration specified:

Using 2016-JUN-04 13:57:46.000000000 as report begin time

Using 2016-JUN-04 14:57:47.000000000 as report end time

In-Memory Advisor: Adding statistics...

In-Memory Advisor: Finished adding statistics.

In-Memory Advisor: Analyzing statistics...

In-Memory Advisor: Finished analyzing statistics.

The Advisor estimates the following performance benefits:

| IN-MEMORY<br>SIZE | PERCENTAGE<br>OF MAXIMUM<br>SGA SIZE | ESTIMATED<br>ANALYTICS<br>PROCESSING<br>TIME<br>REDUCTION<br>(SECONDS) * | ESTIMATED<br>ANALYTICS<br>PROCESSING<br>PERFORMANCE<br>IMPROVEMENT<br>FACTOR* |
|-------------------|--------------------------------------|--------------------------------------------------------------------------|-------------------------------------------------------------------------------|
| 1.074GB           | 37                                   | 227                                                                      | 3.1X                                                                          |
| 1.021GB           | 35                                   | 207                                                                      | 2.6X                                                                          |
| 990.3MB           | 33                                   | 207                                                                      | 2.6X                                                                          |
| 935.2MB           | 31                                   | 76                                                                       | 1.3X                                                                          |
| 880.2MB           | 29                                   | 76                                                                       | 1.3X                                                                          |
| 825.2MB           | 27                                   | 76                                                                       | 1.3X                                                                          |
| 770.2MB           | 26                                   | 76                                                                       | 1.3X                                                                          |
| 715.2MB           | 24                                   | 76                                                                       | 1.3X                                                                          |
| 660.2MB           | 22                                   | 76                                                                       | 1.3X                                                                          |
| 605.2MB           | 20                                   | 76                                                                       | 1.3X                                                                          |
| 550.1MB           | 18                                   | 76                                                                       | 1.3X                                                                          |
| 495.1MB           | 16                                   | 76                                                                       | 1.3X                                                                          |
| 440.1MB           | 15                                   | 76                                                                       | 1.3X                                                                          |
| 385.1MB           | 13                                   | 76                                                                       | 1.3X                                                                          |
| 330.1MB           | 11                                   | 76                                                                       | 1.3X                                                                          |
| 275.1MB           | 9                                    | 76                                                                       | 1.3X                                                                          |
| 220.1MB           | 7                                    | 35                                                                       | 1.1X                                                                          |
| 165.0MB           | 5                                    | 35                                                                       | 1.1X                                                                          |
| 110.0MB           | 4                                    | 35                                                                       | 1.1X                                                                          |
| 55.01MB           | 2                                    | 23                                                                       | 1.1X                                                                          |

\*Estimates: The In-Memory Advisor's estimates are useful for making In-Memory decisions. But they are not precise. Due to performance variations caused by workload diversity, the Advisor's performance estimates are conservatively limited to no more than 10.0X faster.

Choose the In-Memory size you wish for optimization (default=1.074GB):  
1.074GB

The Advisor is optimizing for an In-Memory size of 1.074GB...  
Fetching recommendation files for task: TASK05  
Placing recommendation files in: the current working directory

Fetches file: imadvisor\_TASK05.html  
Purpose: recommendation report primary html page

Fetches file: imadvisor\_TASK05.sql  
Purpose: recommendation DDL sqlplus script

You can re-run this task with this script and specify a different an In-Memory size. Re-running a task to optimize for a different In-Memory size is faster than creatng and running a new task from scratch.

SQL>

## Running the Advisor with a Customized Script

The Oracle Database In-Memory Advisor ships with a predefined script you can use to run the Advisor. There also is a PL/SQL interface to the Advisor which is described in the whitepaper which documents the installation and usage of the Advisor.

There are many additional options using the PL/SQL interface. In this example, a script will be used to run the Advisor with data from the ADD\_HIST\_STATISTICS PL/SQL procedure, which takes as input AWR snapshot ids, rather than the ADD\_STATISTICS procedure in the default script which takes a time range. Here is an example of running the Advisor with a customized script.

The high level steps are:

- Create the Advisor script
- Run SQLPLUS as an appropriately privileged user
- Execute script
- Review the reports

Here is a sample script, using the ADD\_HIST\_STATISTICS procedure:

```
SET SERVEROUTPUT ON;
BEGIN
  BEGIN
    dbms_inmemory_advisor.drop_task ('TASK03', force=>TRUE);
  EXCEPTION
    WHEN OTHERS THEN NULL;
  END;
  dbms_inmemory_advisor.create_task ('TASK03');
  dbms_inmemory_advisor.add_hist_statistics ('TASK03',1814,1820);
  dbms_inmemory_advisor.execute_task ('TASK03');
  dbms_inmemory_advisor.generate_recommendations('TASK03',
    directory_name=>'ADVISORDIR');
END;
/

DEFINE task_name='TASK03';
@imadvisor_fetch_recommendations.sql
```

Here is the output from the session where this script was executed:

```
$ sqlplus / as sysdba
SQL*Plus: Release 12.1.0.2.0 Production on Thu Jun 2 17:55:07 2016
Copyright (c) 1982, 2014, Oracle. All rights reserved.

Connected to:
Oracle Database 12c Enterprise Edition Release 12.1.0.2.0 - 64bit Production
With the Partitioning, OLAP, Advanced Analytics and Real Application Testing
options

SQL> SET SERVEROUTPUT ON;
SQL> SQL> BEGIN
 2  BEGIN
 3      dbms_inmemory_advisor.drop_task ('TASK03', force=>TRUE);
 4  EXCEPTION
 5      WHEN OTHERS THEN NULL;
 6  END;
 7  dbms_inmemory_advisor.create_task ('TASK03');
 8  dbms_inmemory_advisor.add_hist_statistics ('TASK03',1814,1820);
 9
10  dbms_inmemory_advisor.execute_task ('TASK03');
11
12  dbms_inmemory_advisor.generate_recommendations ('TASK03');
13  END;
14  /

DEFINE task_name='TASK03';
@imadvisor_fetch_recommendations.sql

PL/SQL procedure successfully completed.

SQL> SQL> SQL> Fetching recommendation files for task: TASK03
Enter value for client_directory_path: .
Placing recommendation files in: the current working directory

Fetched file: imadvisor_TASK03.html
Purpose:      recommendation report primary html page

Fetched file: imadvisor_sql_TASK03.html
Purpose:      SQL detail secondary html page with link from primary html page

Fetched file: imadvisor_object_TASK03.html
```

Purpose: object detail secondary html page with link from primary html

Fetches file: imadvisor\_auxiliary\_TASK03.html

Purpose: rationale secondary html page with link from primary html page

Fetches file: imadvisor\_TASK03.sql

Purpose: recommendation DDL sqlplus script

SQL>

## Conclusion

The Oracle Database In-Memory Advisor is a flexible tool that can be used to help you identify the workload and objects in your database that will benefit the most with Oracle Database In-Memory.



White Paper Title  
June 2016

Author: [OPTIONAL]  
Contributing Authors: [OPTIONAL]

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

Worldwide Inquiries:  
Phone: +1.650.506.7000  
Fax: +1.650.506.7200

oracle.com



Copyright © 2016, Oracle and/or its affiliates. All rights reserved.

This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0113

**Hardware and Software, Engineered to Work Together**