

ORACLE

Oracle Database In-Memory Advisor

Usage Examples

March, 2022, Version 2.0

Copyright © 2022, Oracle and/or its affiliates

Public

Purpose statement

This document provides usage examples for the use of the Oracle Database In-Memory Advisor. It is intended solely to give you some general guidelines so that you can determine how best to use Database In-Memory in your environment.

Disclaimer

This document in any form, software or printed matter, contains proprietary information that is the exclusive property of Oracle. Your access to and use of this confidential material is subject to the terms and conditions of your Oracle software license and service agreement, which has been executed and with which you agree to comply. This document and information contained herein may not be disclosed, copied, reproduced or distributed to anyone outside Oracle without prior written consent of Oracle. This document is not part of your license agreement nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

This document is for informational purposes only and is intended solely to assist you in planning for the implementation and upgrade of the product features described. 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 in this document remains at the sole discretion of Oracle. Due to the nature of the product architecture, it may not be possible to safely include all features described in this document without risking significant destabilization of the code.

Table of contents

Purpose statement	2
Disclaimer	2
Intended Audience	4
Introduction	4
How to Use This Document	4
Database In-Memory Advisor	5
Scenarios for Running the In-Memory Advisor	5
Running the In-Memory Advisor on a Live Workload	5
Generating Recommendations for Different In-Memory Sizes	9
Running the In-Memory Advisor On AWR Data from a Different Database	12
Export the data	12
Copy the data	17
Import the data	17
Run the Advisor on the imported data	20
Running the In-Memory Advisor on a Batch Workload	24
Running the In-Memory Advisor on a Pluggable Database	25
Running the In-Memory Advisor with a Customized Script	29
Conclusion	30

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 Database In-Memory, see the [Oracle Database In-Memory technical brief](#) 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 In-Memory 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 In-Memory Advisor can be downloaded from My Oracle Support (MOS) note 1965343.1. MOS note 1965343.1 also contains a technical brief describing how to install and run the In-Memory 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 In-Memory Advisor in that manner.

Database In-Memory Advisor

The goal of Oracle Database In-Memory is to optimize analytical processing in the database. The In-Memory Advisor analyzes the analytical processing workload present in your database to determine an estimated benefit for the database as a whole if that analytical workload is optimized.

The In-Memory Advisor differentiates analytics processing from other database activity based upon SQL plan cardinality, Active Session History (ASH), use of parallel query, and other statistics.

The In-Memory Advisor estimates the In-Memory size of objects based upon statistics and heuristic compression factors and, optionally, the DBMS_COMPRESSION package (in Oracle Database 12.1.0.2 and above).

The In-Memory Advisor estimates analytic processing performance improvement factors based upon the following:

- Elimination of user I/O waits, cluster transfer waits, buffer cache latch waits, etc.
- Certain query processing advantages related to specific compression types.
- Decompression cost heuristics per specific compression types.
- SQL plan selectivity, number of columns in the result set, etc.

Scenarios for Running the In-Memory Advisor

Running the In-Memory Advisor on a Live Workload

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

These steps can also be used to run the In-Memory 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 In-Memory 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 In-Memory 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 19.0.0.0.0 - Production on Wed Mar 23 10:53:06 2022  
Version 19.14.0.0.0
```

Copyright (c) 1982, 2021, Oracle. All rights reserved.

Connected to:

Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production
Version 19.14.0.0.0

```
SQL> @imadvisor_recommendations.sql
```

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_20220323105312
Enter value for task_name:

Advisor task name specified: im_advisor_task_20220323105312 (default)

New Advisor task will be named: im_advisor_task_20220323105312...

Analyzing and reporting on a live workload on this database (DBID=524914937)...

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: -120

Report begin time specified: -120

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

Enter value for duration:

Report duration specified:

Using 2022-MAR-23 08:54:09.000000000 as report begin time

Using 2022-MAR-23 10:54:10.000000000 as report end time

You may optionally specify a comma separated list of object owner
and name patterns to be considered for In Memory Placement.

Example:

```
GEEK_SUMMARY.%,%.GEEK_%
```

Press ENTER to consider all objects.

Enter value for consider_objects_like:

Considering all objects for In Memory placement.

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*
1.610GB	37	258	6.0X
1.530GB	35	15	1.1X
1.449GB	33	15	1.1X
1.369GB	31	15	1.1X
1.288GB	29	15	1.1X
1.208GB	27	15	1.1X
1.127GB	26	15	1.1X
1.047GB	24	15	1.1X
989.3MB	22	15	1.1X
906.9MB	20	15	1.1X
824.4MB	18	15	1.1X
742.0MB	16	15	1.1X
659.5MB	15	15	1.1X
577.1MB	13	15	1.1X
494.7MB	11	15	1.1X
412.2MB	9	15	1.1X
329.8MB	7	15	1.1X
247.3MB	5	15	1.1X
164.9MB	4	15	1.1X
82.44MB	2	15	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=1.610GB):

```
The Advisor is optimizing for an In-Memory size of 1.610GB...
Fetching recommendation files for task: im_advisor_task_20220323105312
Placing recommendation files in: the current working directory
```

```
Fetches file: imadvisor_im_advisor_task_20220323105312.html
Purpose:      recommendation report primary html page
```

```
Fetches file: imadvisor_im_advisor_task_20220323105312.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 In-Memory 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 In-Memory Advisor to report on.
 - Review the reports

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

```
SQL> @imadvisor_recommendations.sql
```

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
dbim_pdb_20220309111559	2022-MAR-09 11:17:54
dbim_pdb_20220321153010	2022-MAR-21 15:31:10

Default task_name (new task): im_advisor_task_20220322162636

Enter value for task_name: dbim_pdb_20220321153010

Advisor task name specified: dbim_pdb_20220321153010

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: 3GB

The In-Memory Advisor will optimize for this In-Memory size: 3GB

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: 03/21/22 15:31:10

Report begin time specified: 03/21/22 15:31:10

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

Enter value for duration: 30

Report duration specified: 30

Using 2022-MAR-21 15:31:10.000000000 as report begin time

Using 2022-MAR-21 16:01:10.000000000 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 is optimizing for an In-Memory size of 3GB...

Fetching recommendation files for task: dbim_pdb_20220321153010

Placing recommendation files in: the current working directory

Fetches file: imadvisor_dbim_pdb_20220321153010.html

Purpose: recommendation report primary html page

Fetches file: imadvisor_dbim_pdb_20220321153010.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 creating and running a new task from scratch.

SQL>

Running the In-Memory Advisor On AWR Data from a Different Database

You can run the In-Memory 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 In-Memory 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 In-Memory Advisor
- Install the In-Memory 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 In-Memory 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 19.0.0.0.0 - Production on Wed Mar 23 14:03:41 2022  
Version 19.14.0.0.0
```

```
Copyright (c) 1982, 2021, Oracle. All rights reserved.
```

```
Connected to:
```

```
Oracle Database 19c Enterprise Edition Release 19.0.0.0.0 - Production  
Version 19.14.0.0.0
```

```
SQL> set pages 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                      |
|-------------|---------|---------------------------|
| * 524914937 | DBIM    | localhost.lo<br>caldomain |

The default database id is the local one: '524914937'. To use this database id, press <return> to continue, otherwise enter an alternative.

Enter value for dbid:

Using 524914937 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: 2

Listing the last 2 days of Completed Snapshots

DB Name	Snap Id	Snap Started
DBIM	21	22 Mar 2022 17:17
	22	22 Mar 2022 17:21
	23	22 Mar 2022 18:00
	24	22 Mar 2022 19:00
	25	22 Mar 2022 20:00
	26	22 Mar 2022 21:00
	27	22 Mar 2022 22:00

```

28 22 Mar 2022 23:00
29 23 Mar 2022 00:00
30 23 Mar 2022 01:00
31 23 Mar 2022 02:00
32 23 Mar 2022 03:00
33 23 Mar 2022 04:00
34 23 Mar 2022 05:01
35 23 Mar 2022 06:00
36 23 Mar 2022 07:00
37 23 Mar 2022 08:01
38 23 Mar 2022 09:00
39 23 Mar 2022 09:42
40 23 Mar 2022 09:43
41 23 Mar 2022 11:00
42 23 Mar 2022 12:00
43 23 Mar 2022 13:00
44 23 Mar 2022 14:00

```

Specify the Begin and End Snapshot Ids

~~~~~

Enter value for begin\_snap: 37

Begin Snapshot Id specified: 37

Enter value for end\_snap: 42

End Snapshot Id specified: 42

Specify the Directory Name

~~~~~

Directory Name	Directory Path
ADVISORDIR	/home/oracle/InMemory/imadvisor
DATA_PUMP_DIR	/u01/app/oracle/admin/dbim/dpdump/
DBMS_OPTIM_ADMINDIR	/u01/app/oracle/product/db19/db_1/rdbms/admin
DBMS_OPTIM_LOGDIR	/u01/app/oracle/product/db19/db_1/cfgtoollogs
DPDIR	/u02/expdp
JAVA\$JOX\$CUJS\$DIRECTORY\$	/u01/app/oracle/product/db19/db_1/javavm/admin/
OPATCH_INST_DIR	/u01/app/oracle/product/db19/db_1/OPatch
OPATCH_LOG_DIR	/u01/app/oracle/product/db19/db_1/rdbms/log
OPATCH_SCRIPT_DIR	/u01/app/oracle/product/db19/db_1/QOpatch
ORACLE_BASE	/u01/app/oracle
ORACLE_HOME	/u01/app/oracle/product/db19/db_1
ORACLE_OCM_CONFIG_DIR	/u01/app/oracle/product/db19/db_1/ccr/state
ORACLE_OCM_CONFIG_DIR2	/u01/app/oracle/product/db19/db_1/ccr/state
SDO_DIR_ADMIN	/u01/app/oracle/product/db19/db_1/md/admin
SDO_DIR_WORK	

```
XMLDIR          /u01/app/oracle/product/db19/db_1/rdbms/xml
XSDDIR          /u01/app/oracle/product/db19/db_1/rdbms/xml/schem
a
```

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\_37\_42.  
To use this name, press <return> to continue, otherwise enter  
an alternative.

Enter value for file\_name:

Using the dump file prefix: awrdat\_37\_42

```
|
| ~~~~~
| The AWR extract dump file will be located
| in the following directory/file:
| /home/oracle/InMemory/imadvisor
| awrdat_37_42.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:
| /home/oracle/InMemory/imadvisor
| awrdat_37_42.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: The Automatic Workload Repository (AWR) augment you are about to create  
\* can only be used with an AWR export from this local database  
\* (DBID=524914937).  
\*

\* Also note: It is best to capture an AWR augment in the same timeframe  
\* (preferably after) the corresponding AWR export.  
\*\*\*\*\*

You may optionally include one the following SQL Tuning Set(s) as part of the  
AWR augment.

Enter the sqlset\_owner and sqlset\_name in the prompts below to include a SQL  
Tuning Set.

Otherwise, press ENTER for both sqlset\_owner and sqlset\_name if you do not wish  
to include any SQL Sets.

SYS SYS\_AUTO\_STS

Enter value for sqlset\_owner:

Enter value for sqlset\_name:

|                               |                                                    |
|-------------------------------|----------------------------------------------------|
| ADVISORDIR                    | /home/oracle/InMemory/imadvisor                    |
| DATA_PUMP_DIR                 | /u01/app/oracle/admin/dbim/dpdump/                 |
| DBMS_OPTIM_ADMINDIR           | /u01/app/oracle/product/db19/db_1/rdbms/admin      |
| DBMS_OPTIM_LOGDIR             | /u01/app/oracle/product/db19/db_1/cfgtoollogs      |
| DPDIR                         | /u02/expdp                                         |
| JAVA\$JOX\$CUJSS\$DIRECTORY\$ | /u01/app/oracle/product/db19/db_1/javavm/admin/    |
| OPATCH_INST_DIR               | /u01/app/oracle/product/db19/db_1/OPatch           |
| OPATCH_LOG_DIR                | /u01/app/oracle/product/db19/db_1/rdbms/log        |
| OPATCH_SCRIPT_DIR             | /u01/app/oracle/product/db19/db_1/QOPatch          |
| ORACLE_BASE                   | /u01/app/oracle                                    |
| ORACLE_HOME                   | /u01/app/oracle/product/db19/db_1                  |
| ORACLE_OCM_CONFIG_DIR         | /u01/app/oracle/product/db19/db_1/ccr/state        |
| ORACLE_OCM_CONFIG_DIR2        | /u01/app/oracle/product/db19/db_1/ccr/state        |
| SDO_DIR_ADMIN                 | /u01/app/oracle/product/db19/db_1/md/admin         |
| SDO_DIR_WORK                  |                                                    |
| XMLDIR                        | /u01/app/oracle/product/db19/db_1/rdbms/xml        |
| XSDDIR                        | /u01/app/oracle/product/db19/db_1/rdbms/xml/schema |

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?

imadvisor\_awr\_augment\_37\_42

Using "imadvisor\_awr\_augment\_37\_42" as the AWR augment dump file name prefix...

Setting up the AWR agument staging schema...

No errors.

old 942: dmp\_name := NVL('&&dump\_file\_name\_prefix', :dmp\_name\_default);

new 942: dmp\_name := NVL('imadvisor\_awr\_augment\_37\_42', :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: /home/oracle/InMemory/imadvisor

Data Pump dump file: imadvisor\_awr\_augment\_37\_42.dmp

Data Pump log file: imadvisor\_awr\_augment\_37\_42\_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 21.0.0.0.0 - Production on Wed Mar 23 15:52:26 2022

Version 21.4.0.0.0

Copyright (c) 1982, 2021, Oracle. All rights reserved.

Connected to:

Oracle Database 21c Enterprise Edition Release 21.0.0.0.0 - Production  
Version 21.4.0.0.0

```
SQL> create directory advisordir as '/home/oracle/InMemory/imadvisor';
```

Directory created.

```
SQL> set pages 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 ~
```

```
~~~~~
```

Specify the Directory Name

```
~~~~~
```

| Directory Name               | Directory Path                                           |
|------------------------------|----------------------------------------------------------|
| ADVISORDIR                   | /home/oracle/InMemory/imadvisor                          |
| DATA_PUMP_DIR                | /u01/app/oracle/admin/dbim/dpdump/                       |
| DBMS_OPTIM_ADMINDIR          | /u01/app/oracle/product/21.0.0/db_1/rdbms/admin          |
| DBMS_OPTIM_LOGDIR            | /u01/app/oracle/product/21.0.0/db_1/cfgtoollogs          |
| JAVA\$JOX\$CUJS\$DIRECTORY\$ | /u01/app/oracle/product/21.4.0/db_1/javavm/admin/        |
| OPATCH_INST_DIR              | /u01/app/oracle/product/21.4.0/db_1/OPatch               |
| OPATCH_LOG_DIR               | /u01/app/oracle/homes/OraDB21Home1/rdbms/log             |
| OPATCH_SCRIPT_DIR            | /u01/app/oracle/product/21.4.0/db_1/QOpatch              |
| ORACLE_BASE                  | /u01/app/oracle                                          |
| ORACLE_HOME                  | /u01/app/oracle/product/21.0.0/db_1                      |
| ORACLE_OCM_CONFIG_DIR        | /u01/app/oracle/homes/OraDB21Home1/ccr/state             |
| ORACLE_OCM_CONFIG_DIR2       | /u01/app/oracle/homes/OraDB21Home1/ccr/state             |
| SDO_DIR_ADMIN                | /u01/app/oracle/product/21.0.0/db_1/md/admin             |
| XMLDIR                       | /u01/app/oracle/product/21.0.0/db_1/rdbms/xml            |
| XSDDIR                       | /u01/app/oracle/product/21.0.0/db_1/rdbms/xml/sch<br>ema |

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

```
Enter value for directory_name: ADVISORDIR
```

```
Using the dump directory: ADVISORDIR
```

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\_37\_42

Loading from the file name: awrdat\_37\_42.dmp

```
|
| ~~~~~
| Loading the AWR data from the following
| directory/file:
| /home/oracle/InMemory/imadvisor
| awrdat_37_42.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:
| /home/oracle/InMemory/imadvisor
| awrdat_37_42.log
|
```

End of AWR Load

SQL>

SQL>

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

```
ADVISORDIR /home/oracle/InMemory/imadvisor
DATA_PUMP_DIR /u01/app/oracle/admin/dbim/dpdump/
DBMS_OPTIM_ADMINDIR /u01/app/oracle/product/21.0.0/db_1/rdbms/admin
DBMS_OPTIM_LOGDIR /u01/app/oracle/product/21.0.0/db_1/cfgtoollogs
JAVAJOXCUJS$DIRECTORY$ /u01/app/oracle/product/21.4.0/db_1/javavm/admin/
OPATCH_INST_DIR /u01/app/oracle/product/21.4.0/db_1/OPatch
OPATCH_LOG_DIR /u01/app/oracle/homes/OraDB21Home1/rdbms/log
OPATCH_SCRIPT_DIR /u01/app/oracle/product/21.4.0/db_1/QOpatch
ORACLE_BASE /u01/app/oracle
ORACLE_HOME /u01/app/oracle/product/21.0.0/db_1
ORACLE_OCM_CONFIG_DIR /u01/app/oracle/homes/OraDB21Home1/ccr/state
ORACLE_OCM_CONFIG_DIR2 /u01/app/oracle/homes/OraDB21Home1/ccr/state
SDO_DIR_ADMIN /u01/app/oracle/product/21.0.0/db_1/md/admin
XMLDIR /u01/app/oracle/product/21.0.0/db_1/rdbms/xml
XSDDIR /u01/app/oracle/product/21.0.0/db_1/rdbms/xml/schema
```

Please enter the Oracle directory object to use for import (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, please enter an alternative name prefix?

imadvisor\_awr\_augment\_37\_42

Using "imadvisor\_awr\_augment\_37\_42" as the AWR augment dump file name prefix...

Gathering information about the AWR agument...

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>

## Run the Advisor on the imported data

```
$ sqlplus / as sysdba
```

```
SQL*Plus: Release 21.0.0.0.0 - Production on Wed Mar 23 15:59:31 2022
Version 21.4.0.0.0
```

```
Copyright (c) 1982, 2021, Oracle. All rights reserved.
```

```
Connected to:
```

```
Oracle Database 21c Enterprise Edition Release 21.0.0.0.0 - Production
Version 21.4.0.0.0
```

```
SQL> @imadvisor_recommendations.sql
```

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         |
|--------------------------------|----------------------|
| dbim_pdb_20220309111559        | 2022-MAR-09 11:17:54 |
| dbim_pdb_20220309112209        | 2022-MAR-09 11:22:54 |
| im_advisor_task_20220321152708 | 2022-MAR-21 15:28:31 |
| dbim_pdb_20220321153010        | 2022-MAR-21 15:31:10 |
| dbim_pdb_20220321153937        | 2022-MAR-21 15:41:04 |
| dbim_pdb_20220321154151        | 2022-MAR-21 15:42:34 |
| im_advisor_task_20220322080610 | 2022-MAR-22 08:06:55 |
| im_advisor_task_20220322080734 | 2022-MAR-22 08:08:11 |
| im_advisor_task_20220322080900 | 2022-MAR-22 08:10:25 |

Default task\_name (new task): im\_advisor\_task\_20220323155943

Enter value for task\_name: task\_imp\_awr

Advisor task name specified: task\_imp\_awr

New Advisor task will be named: task\_imp\_awr...

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  
-----  
524914937

Enter value for dbid: 524914937

Analyzing and reporting on an augmented AWR workload with DBID=524914937...

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: 03/23/22 07:00:39

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 2022-MAR-23 07:00:39.000000000 as report begin time

Using 2022-MAR-23 12:00:20.000000000 as report end time

You may optionally specify a comma separated list of object owner and name patterns to be considered for In Memory Placement.

Example:

GEEK\_SUMMARY.%,%.GEEK\_%

Press ENTER to consider all objects.

Enter value for consider\_objects\_like:

Considering all objects for In Memory placement.

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* |
|--|-------------------|--------------------------------------|--------------------------------------------------------------------------|-------------------------------------------------------------------------------|
|  | 1.557GB           | 35                                   | 256                                                                      | 5.7X                                                                          |
|  | 1.479GB           | 34                                   | 27                                                                       | 1.1X                                                                          |
|  | 1.401GB           | 32                                   | 27                                                                       | 1.1X                                                                          |
|  | 1.323GB           | 30                                   | 27                                                                       | 1.1X                                                                          |
|  | 1.245GB           | 28                                   | 27                                                                       | 1.1X                                                                          |
|  | 1.167GB           | 27                                   | 27                                                                       | 1.1X                                                                          |

|         |    |    |      |
|---------|----|----|------|
| 1.090GB | 25 | 27 | 1.1X |
| 1.012GB | 23 | 27 | 1.1X |
| 956.3MB | 21 | 27 | 1.1X |
| 876.6MB | 19 | 27 | 1.1X |
| 796.9MB | 18 | 27 | 1.1X |
| 717.2MB | 16 | 27 | 1.1X |
| 637.6MB | 14 | 27 | 1.1X |
| 557.9MB | 12 | 27 | 1.1X |
| 478.2MB | 11 | 27 | 1.1X |
| 398.5MB | 9  | 27 | 1.1X |
| 318.8MB | 7  | 27 | 1.1X |
| 239.1MB | 5  | 27 | 1.1X |
| 159.4MB | 4  | 27 | 1.1X |
| 79.69MB | 2  | 27 | 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=1.557GB):

The Advisor is optimizing for an In-Memory size of 1.557GB...  
 Fetching recommendation files for task: task\_imp\_awr  
 Placing recommendation files in: the current working directory

Fetches file: imadvisor\_task\_imp\_awr.html  
 Purpose: recommendation report primary html page

Fetches file: imadvisor\_task\_imp\_awr.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 In-Memory 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 In-Memory Advisor for a batch reporting workload is the same as running the In-Memory Advisor on a live (currently executing workload). You just need to specify the time period that you wish the In-Memory 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 In-Memory 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 In-Memory Advisor on a Pluggable Database

The 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 In-Memory Advisor in the root container, the CDB. Then you can run the In-Memory 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 In-Memory 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.

```
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\_20220309111559

Enter value for task\_name: dbim\_pdb\_20220309111559

Advisor task name specified: dbim\_pdb\_20220309111559

New Advisor task will be named: dbim\_pdb\_20220309111559...

Analyzing and reporting on a live workload on this database (DBID=543831963)...

Enter value for pdb\_name: dbimpdb

dbimpdb

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: <ENTER>

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 2022-MAR-09 10:00:00.000000000 as report begin time
Using 2022-MAR-09 11:00:00.000000000 as report end time
```

```
You may optionally specify a comma separated list of object owner
and name patterns to be considered for In Memory Placement.
Example:
```

```
GEEK_SUMMARY.%,%.GEEK_%
```

```
Press ENTER to consider all objects.
```

```
Enter value for consider_objects_like:
```

```
Considering all objects for In Memory placement.
```

```
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* |
|-------------------|--------------------------------------|--------------------------------------------------------------------------|-------------------------------------------------------------------------------|
| 2.662GB           | 46                                   | 26                                                                       | 3.4X                                                                          |
| 2.529GB           | 44                                   | 20                                                                       | 2.2X                                                                          |
| 2.396GB           | 42                                   | 20                                                                       | 2.2X                                                                          |
| 2.263GB           | 39                                   | 20                                                                       | 2.2X                                                                          |
| 2.129GB           | 37                                   | 20                                                                       | 2.2X                                                                          |
| 1.996GB           | 35                                   | 20                                                                       | 2.2X                                                                          |
| 1.863GB           | 32                                   | 20                                                                       | 2.2X                                                                          |
| 1.730GB           | 30                                   | 20                                                                       | 2.2X                                                                          |
| 1.597GB           | 28                                   | 20                                                                       | 2.2X                                                                          |
| 1.464GB           | 25                                   | 20                                                                       | 2.2X                                                                          |
| 1.331GB           | 23                                   | 20                                                                       | 2.2X                                                                          |
| 1.198GB           | 21                                   | 20                                                                       | 2.2X                                                                          |
| 1.065GB           | 18                                   | 0                                                                        |                                                                               |
| 954.0MB           | 16                                   | 0                                                                        |                                                                               |
| 817.7MB           | 14                                   | 0                                                                        |                                                                               |
| 681.4MB           | 12                                   | 0                                                                        |                                                                               |
| 545.1MB           | 9                                    | 0                                                                        |                                                                               |
| 408.9MB           | 7                                    | 0                                                                        |                                                                               |
| 272.6MB           | 5                                    | 0                                                                        |                                                                               |
| 136.3MB           | 2                                    | 0                                                                        |                                                                               |

\*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=2.662GB):

The Advisor is optimizing for an In-Memory size of 2.662GB...  
 Fetching recommendation files for task: dbim\_pdb\_20220309111559  
 Placing recommendation files in: the current working directory

Fetches file: imadvisor\_dbim\_pdb\_20220309111559.html  
 Purpose: recommendation report primary html page

Fetches file: imadvisor\_dbim\_pdb\_20220309111559.sql  
Purpose: recommendation DDL sqlplus script

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

SQL>

### Running the In-Memory Advisor with a Customized Script

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

There are many additional options using the PL/SQL interface. In this example, a script will be used to run the In-Memory 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 In-Memory Advisor with a customized script.

The high level steps are:

- Create the In-Memory 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
```

```
Fetches file: imadvisor_TASK03.html
```

```
Purpose: recommendation report primary html page
```

```
Fetches file: imadvisor_sql_TASK03.html
```

```
Purpose: SQL detail secondary html page with link from primary html page
```

```
Fetches 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 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.

---

## Connect with us

Call **+1.800.ORACLE1** or visit **oracle.com**. Outside North America, find your local office at: **oracle.com/contact**.

 [blogs.oracle.com](https://blogs.oracle.com)

 [facebook.com/oracle](https://facebook.com/oracle)

 [twitter.com/oracle](https://twitter.com/oracle)

---

Copyright © 2022, 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.

This device has not been authorized as required by the rules of the Federal Communications Commission. This device is not, and may not be, offered for sale or lease, or sold or leased, until authorization is obtained.

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. 0120

Disclaimer: If you are unsure whether your data sheet needs a disclaimer, read the revenue recognition policy. If you have further questions about your content and the disclaimer requirements, e-mail [REVREC\\_US@oracle.com](mailto:REVREC_US@oracle.com).

---