

Oracle

1Z0-064 Exam

**Oracle Database 12c: Performance Management and Tuning
Exam**

**Questions & Answers
Demo**

Version: 10.0

Question: 1

Your database supports a mixed workload. The ERP application creates short sessions and performs small, random I/Os; the REPORTING application executes long-running DSS queries.

You want to set a priority for the workload generated by the ERP application and optimize resource usage for them.

Which three objectives can be achieved by the Resource Manager? (Choose three.)

- A. limiting the amount of time that a session is idle and blocking other sessions of the ERP application
- B. limiting the amount of undo generated by operations performed by sessions created by the ERP application
- C. creating two resource plans with resource limits defined for the workload generated by the applications and automatically changing resource plans based on the workload
- D. allocating a lower percentage of CPU to sessions used by the REPORTING application than to those used by the ERP application
- E. limiting the physical I/O performed by the sessions or users of the ERP application that are connected to the database

Answer: B,D,E

Question: 2

Examine the partial TKPROF output for an SQL statement:

```
SQL> SELECT city_id
      FROM city_names
      WHERE code = 'DLR'?
```

call	count	cpu	elapsed	disk	query	current	rows
Parse	1	0.06	0.10	0	0	0	0
Execute	1	0.02	0.02	0	0	0	0
Fetch	1	0.23	0.30	31	31	3	1

```
Misses in library cache during parse: 0
Parsing user id: 02 (USER2)
```

Rows	Execution Plan
0	SELECT STATEMENT
2340	TABLE ACCESS (BY ROWID) OF 'CITY_NAMES'
0	INDEX (RANGE SCAN) OF 'CITY_NAMES_NAME' (NON-UNIQUE)

Which two inferences can definitely be made from this output? (Choose two.)

- A. Array fetch operations were not performed for this query.
- B. No hard parse was performed for this query.
- C. The number of logical I/Os is almost equal to the number of physical I/Os.
- D. Another transaction held a shared lock on the table, thereby causing a significant delay.
- E. An uncommitted transaction made a series of updates to the NAME_ID column just before the execution of this query.

Answer: B,D

Question: 3

In your database, the measured 99th percentile value is used as the maximum value. You set a warning threshold level of 110% of maximum trigger as an alert.

What is the outcome? (Choose the best answer.)

- A. It generates an error because the warning threshold cannot exceed 100%.
- B. It generates an error because the percentage of maximum threshold cannot be set with a significance-level threshold value.
- C. It generates an alert when an observed metric is 99% of the 99th percentile value as measured over the moving window baseline.
- D. It generates an alert when an observed metric is 110% of the 99th percentile value as measured over the moving window baseline.
- E. It generates an alert when 1 in 100 observations for an observed metric exceeds the 99th percentile value as measured over the fixed baseline.

Answer: A

Question: 4

You want to capture AWR data to monitor performance variation every Monday between 9:00 AM and 12:00 PM for three months and automatically remove the older AWR data every fortnight. How would you achieve this? (Choose the best answer.)

- A. Create AWR baselines.
- B. Create SQL plan baselines.
- C. Create repeating baseline templates.
- D. Create database services and make sure that user connections use them to connect to the database instance.
- E. Create a single baseline template.

Answer: D

Question: 5

Which two actions can reduce library cache latch contention for an OLTP application that repeatedly executes queries containing a mix of literals and bind variables? (Choose two.)

- A. setting the OPEN_CURSORS parameter to hold a sufficient number of concurrently open cursors
- B. coding the application such that a frequently executed statement is parsed only once and executed repeatedly as required
- C. setting the CURSOR_SHARING parameter to EXACT
- D. avoiding the granting of privileges on objects associated with cursors during peak load
- E. enabling Automatic Memory Management and allocating at least 50% of the available memory for SHARED_POOL_SIZE
- F. configuring shared server connections

Answer: C,E

Question: 6

You are administering a database that supports multiple applications, which make dedicated connections to the database instance by using different services.

You execute the command to enable tracing of the ORCL1 service:

```
SQL> EXECUTE DBMS_MONITOR.SERV_MOD_ACT_TRACE_ENABLE (service_name => 'ORCL1', WAITS => TRUE, BINDS => NULL, instance_name => 'ORCL', plan_stat => NULL);
```

Which two statements are true? (Choose two.)

- A. A single trace file is generated for all sessions mapped to the ORCL1 service.
- B. SQL trace is enabled for all modules and actions for sessions mapped to the ORCL1 service.
- C. A SQL trace file is generated for each session that maps to the ORCL1 service.

- D. A SQL trace file is generated for each of the modules using the ORCL1 service.
- E. SQL trace is not enabled for the service because a module name is not specified.

Answer: A,C

Question: 7

In which three situations can dynamic statistics be used? (Choose three.)

- A. when the sampling time is a small fraction of the total time for a query
- B. when an execution plan is suboptimal because of complex predicates
- C. when extended statistics are not available for SQL statements that use complex predicates
- D. when a query is on a partitioned table with a global index
- E. when index statistics are missing on a column that is used in SQL statements with highly selective filters

Answer: B,C,D

Question: 8

You want to generate statistics for new objects added to an OLTP application without affecting old statistics and the sessions that currently use them.

Which two tasks would you perform to test the new statistics? (Choose two.)

- A. Set the OPTIMIZER_USE_PENDING_STATISTICS initialization parameter to TRUE for the session.
- B. Set the STALE_PERCENT statistics preference to zero and then gather statistics.
- C. Set the PUBLISH statistics preference to FALSE and then gather statistics.
- D. Use the DBMS_STATS.PUBLISH_PENDING_STATS procedure to make pending statistics the current statistics.
- E. Set the NO_INVALIDATE statistics preference to FALSE and gather statistics without affecting old statistics.

Answer: A,B

Question: 9

You plan to upgrade your production database from Oracle Database 11g to 12c. As part of the upgrade, you want to introduce new indexes and materialized views. You have already created a test system with Oracle Database 12c, having the same structure and data as the production database, along with new schema objects to be added to the production database.

You want to identify regressed SQL statements, if any, which may have been caused by schema changes and the change in the optimizer version.

Which two methods would you use to achieve this? (Choose two.)

- A. Create an SQL Tuning Set (STS) for the SQL statements on the production database and submit as input to the SQL Tuning Advisor on the test database.
- B. Create an STS for the SQL statements on the production database and submit as input to the SQL Performance Analyzer with the OPTIMIZER_FEATURES_ENABLE parameter first set to 11.2.0.1, and then to 12.1.0.1 on the test database.
- C. Generate an Automatic Workload Repository (AWR) compare periods report with snapshots taken before and after schema changes on the test database.
- D. Capture the production database workload, replay it on the test system by using Database Replay, and analyze by using the workload replay compare period report.
- E. Create an STS for the SQL statements on the production database and submit as input to the SQL Access Advisor on the test database.
- F. Create an STS for the SQL statements on the production database before and after changes and submit as input to the SQL Performance Analyzer on the test database.

Answer: A,D

Question: 10

Examine the Time Model Statistics section of an AWR report:

Statistic Name	Time (s)	% of DB Time
sql execute elapsed time	12,416.14	86.45
DB CPU	9,223.70	64.22
parse time elapsed	935.61	6.51
hard parse elapsed time	884.73	6.16
failed parse elapsed time	21.39	.72
PL/SQL execution elapsed time	153.51	1.07
hard parse (sharing criteria) elapsed time	25.96	0.18
connection management call elapsed time	14.00	0.10
hard parse (bind mismatch) elapsed time	4.74	0.03
PL/SQL compilation elapsed time	1.20	0.01
repeated bind elapsed time	0.22	0.00
sequence load elapsed time	0.11	0.00
DB time	14,362.96	
background elapsed time	731.00	
background cpu time	72.00	

Which two inferences can be definitely derived from this section? (Choose two.)

- A. The available CPU resources were not utilized to their maximum capacity.
- B. All sequence numbers used during this AWR time interval were cached.
- C. A large number of connected user sessions were idle.
- D. New child cursors were created because of new bind values or usage of literal values as well as different bind types or sizes.
- E. The DB CPU time was not spent exclusively for processing SQL statements.

Answer: D,E
