

# 1Z0-117<sup>Q&As</sup>

Oracle Database 11g Release 2: SQL Tuning Exam

## Pass Oracle 1Z0-117 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.pass2lead.com/1z0-117.html>

100% Passing Guarantee  
100% Money Back Assurance

Following Questions and Answers are all new published by Oracle  
Official Exam Center

- ⚙️ **Instant Download** After Purchase
- ⚙️ **100% Money Back** Guarantee
- ⚙️ **365 Days** Free Update
- ⚙️ **800,000+** Satisfied Customers



## QUESTION 1

You execute the following query:

```
SQL> SELECT /*+NO_MERGE (v) PUSH_PRED (v) */ *  
        FROM employees e,  
        (SELECT manager_id  
         FROM employees) v  
        WHERE e.manager_id = v.manager_id(+)  
        AND e.employee_id = 100;
```

Which statement is true about the usage of these hints in the query?

- A. The optimizer pushes the join predicate into the inline view.
- B. The optimizer evaluates the subquery as a first and then filters out rows.
- C. The optimizer performs a join operation first and then filters out the rows.
- D. The hint will have no effect because one of the join resultsets is an inline view.

Correct Answer: A

The PUSH\_PRED hint forces pushing of a join predicate into the view.

For example:

```
SELECT /*+ NO_MERGE(v) PUSH_PRED(v) */ *  
        FROM employees e,  
        (SELECT manager_id  
         FROM employees  
        ) v  
        WHERE e.manager_id = v.manager_id(+)  
        AND e.employee_id = 100;
```

When the PUSH\_PRED hint is used without an argument, it should be placed in the view query block. When PUSH\_PRED is used with the view name as an

argument, it should be placed in the surrounding query.

Reference: Oracle Database Performance Tuning Guide, PUSH\_PRED

## QUESTION 2

You are administering a database that supports an OLTP application. To set statistics preferences, you issued the

following command:

```
SQL > DBMS_STATS.SET_GLOBAL_PREFS ('ESTIMATE_PERCENT\','9');
```

What will be the effect of executing this procedure?

- A. It will influence the gathering of statistics for a table based on the value specified for ESTIMATE\_PERCENT provided on table preferences for the same table exist.
- B. It will influence dynamic sampling for a query to estimate the statistics based on ESTIMATE\_PERCENT.
- C. The automatic statistics gathering job running in the maintenance window will use global preferences unless table preferences for the same table exist.
- D. New objects created will use global preference even if table preferences are specified.

Correct Answer: C

Note:

\*

With the DBMS\_STATS package you can view and modify optimizer statistics gathered for database objects.

\*

The SET\_GLOBAL\_PREFS procedure is used to set the global statistics preferences.

\*

ESTIMATE\_PERCENT - The value determines the percentage of rows to estimate. The valid range is [0.000001,100]. Use the constant DBMS\_STATS.AUTO\_SAMPLE\_SIZE to have Oracle determine the appropriate sample size for good statistics. This is the default.

---

### QUESTION 3

Which two statements are true about index full scans?

- A. An index fast full scan multi block I/O to read the index structure in its entirety.
- B. Index nodes are not retrieved in the index order, and therefore the nodes are not in sequence.
- C. An index fast full scan reads the index block by block.
- D. An index fast full scan reads the whole index from the lowest value to the higher value.

Correct Answer: AB

A: To speed table and index block access, Oracle uses the db\_file\_multiblock\_read\_count parameter (which defaults to 8) to aid in getting full-table scan and full-index scan data blocks into the data buffer cache as fast as possible.

B: The index nodes are not retrieved in index order, the rows will not be sequenced.

Note:

there are some requirements for Oracle to invoke the fast full-index scan.

All of the columns required must be specified in the index. That is, all columns in the select and where clauses must exist in the index.

The query returns more than 10 percent of the rows within the index. This 10 percent figure depends on the degree of multi-block reads and the degree of

parallelism.

You are counting the number of rows in a table that meet a specific criterion. The fast full-index scan is almost always used for count(\*) operations.

Reference: index fast full scan tips

#### QUESTION 4

You executed the following statement:

```
SQL> EXPLAIN PLAN SET STATEMENT_ID = 'emp_dept' FOR
      SELECT e.ename e.sal, d.dname
      FROM emp e, dept d
      WHERE e.dept_id = d.dept_id;
```

Which three statements are true about EXPLAIN PLAN?

- A. The execution plan is saved in PLAN\_TABLE without executing the query.
- B. The execution plan for the query is generated and displayed immediately as the output.
- C. The execution plan generated may not necessarily be the execution plan used during query execution.
- D. The execution plan is saved in DBA\_HIST\_SQL\_PLAN without executing the query.
- E. The execution plan generated can be viewed using the DBMS\_XPLAIN.DISPLAY function.
- F. The execution plan generated can be fetched from the library cache by using the DBMS\_XPLAIN.DISPLAY function.

Correct Answer: ACE

\*

(A, not D): The explain plan process stores data in the PLAN\_TABLE.

\*

EXPLAIN PLAN

The EXPLAIN PLAN method doesn't require the query to be run (A), greatly reducing the time it takes to get an execution plan for long-running queries compared to AUTOTRACE.

E: Use the DBMS\_XPLAIN.DISPLAY function to display the execution plan.

\* The DBMS\_XPLAN package provides an easy way to display the output of the EXPLAIN PLAN command in several, predefined formats. You can also use the DBMS\_XPLAN package to display the plan of a statement stored in the Automatic Workload Repository (AWR) or stored in a SQL tuning set. It further provides a way to display the SQL execution plan and SQL execution runtime statistics for cached SQL cursors based on the information stored in the V\$SQL\_PLAN and V\$SQL\_PLAN\_STATISTICS\_ALL fixed views.

Note:

\*

First the query must be explained.

```
SQL> EXPLAIN PLAN FOR
```

```
2 SELECT *
```

```
3 FROM emp e, dept d
```

```
4 WHERE e.deptno = d.deptno
```

```
5 AND e.ename = \'SMITH\';
```

Explained.

```
SQL>
```

Then the execution plan displayed. (not B)

```
SQL> @$ORACLE_HOME/rdbms/admin/utlxpls.sql
```

Plan Table

```
----- | Operation | Name | Rows | Bytes| Cost | Pstart| Pstop |
----- | SELECT STATEMENT |      |      |      |      |      |      |
| NESTED LOOPS |      |      |      |      |      |      |
| TABLE ACCESS FULL |EMP |      |      |      |      |      |
| TABLE ACCESS BY INDEX ROWID |DEPT |      |      |      |      |      |
| INDEX UNIQUE SCAN |PK_DEPT |      |      |      |      |      |
```

8 rows selected.

```
SQL>
```

For parallel queries use the "utlxplp.sql" script instead of "utlxpls.sql".

---

## QUESTION 5

Which four types of column filtering may benefit from partition pruning when accessing tables via partitioned indexes?

A. Equality operates on List-Partitioned Indexes

- B. Not Equal operates on a Global Hash-Partitioned Indexes
- C. Equality operates on System-Partitioned Tables
- D. In-List operates on Range-Partitioned Indexes
- E. Not Equal operates on a local Hash-Partitioned Indexes
- F. Equality operates on Range-Partitioned Indexes
- G. Equality operates on Hash-Partitioned Indexes

Correct Answer: ADFG

Oracle Database prunes partitions when you use range, LIKE, equality (A, F), and IN-list (D) predicates on the range or list partitioning columns, and when you use equality (G) and IN-list predicates on the hash partitioning columns.

Reference: Oracle Database VLDB and Partitioning Guide 11g, Information that can be Used for Partition Pruning

[Latest 1Z0-117 Dumps](#)

[1Z0-117 Study Guide](#)

[1Z0-117 Braindumps](#)