

1Z0-117^{Q&As}

Oracle Database 11g Release 2: SQL Tuning Exam

Pass home 1Z0-117 Exam with 100% Guarantee

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

<https://www.leads4pass.com/1z0-117.html>

100% Passing Guarantee
100% Money Back Assurance

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

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



QUESTION 1

Examine the parallelism parameters for your instance:

NAME	TYPE	VALUE
parallel_adaptive_multi_user	boolean	TRUE
Parallel_automatic_tuning	boolean	FALSE
parallel_degree_limit	string	CPU
parallel_degree_Opolicy	string	LIMITED
parallel_execution_local	integer	16384
parallel_force_local	boolean	FALSE
parallel_io_cap_enabled	boolean	FALSE
parallel_min_percent	integer	80
parallel_min_servers	integer	20
parallel_min_time_threshold	string	AUTO
parallel_server	boolean	FALSE
parallel_server_instances	integer	1
parallel_servers_target	integer	8
parallel_threads_per_cpu	integer	2

You executed the following query:

```
SQL> SELECT /*+ PARALLEL (AUTO)*/SUM(1_extendedprice) total_rev
      FROM lineitem;
```

What are true about the execution of the query?

- A. It will execute in parallel only if the LINEITEM table has a dictionary DOP defined.
- B. DOP for the statement is determined by the dictionary DOP of the accessed objects.
- C. It is generated to execute in parallel.
- D. It will execute in parallel only if the estimated execution time is 10 or more seconds.
- E. DOP for the statement is calculated automatically.
- F. It may execute serially.

Correct Answer: EF

E:

F (not C): It may execute serially. See note below.

Incorrect:

A, B: Dictionary DOP not used with PARALLEL (AUTO) hint.

D: The default value of parallel_min_time_threshold is 30 (not 10) seconds.

Note:

* parallel_min_percent PARALLEL_MIN_PERCENT operates in conjunction with PARALLEL_MAX_SERVERS and PARALLEL_MIN_SERVERS. It lets you specify the minimum percentage of parallel execution processes (of the value of PARALLEL_MAX_SERVERS) required for parallel execution. Setting this parameter ensures that parallel operations will not execute sequentially unless adequate resources are available. The default value of 0 means that no minimum percentage of processes has been set.

Consider the following settings:

PARALLEL_MIN_PERCENT = 50 PARALLEL_MIN_SERVERS = 5 PARALLEL_MAX_SERVERS = 10

If 8 of the 10 parallel execution processes are busy, only 2 processes are available. If you then request a query with a degree of parallelism of 8, the minimum 50% will not be met.

QUESTION 2

Examine the exhibit to view the query and its execution plan.

```
SQL> SELECT /*+ PRDERED */ E.EMPNO, E.ENAME.D.DNAME  
        FROM emp e, dept d  
        WHERE e.deptno=d. deptno  
        Order by e.deptno, D.DEPTNO;
```

EXECUTION PLAN

Plan hash value: 3232458624

Id	Operation	Name	Rows	Bytes	Cost	(%CPU)	Time
0	SELECT STATEMENT		14	364	8	(25)	00:00:01
1	SORT ORDER BY		14	364	8	(25)	00:00:01
*2	HASH JOIN		14	364	7	(15)	00:00:01
3	TABLE ACCESS FULL	EMP	14	182	3	(0)	00:00:01
4	TABLE ACCESS FULL	DEPT	4	52	(3)	(0)	00:00:01

Predicate Information (Identified by operation id):

2 – access (“E”. “DEPTNO” = “D”. “DEPTNO”)

id	operation
0	SELECT STATEMENT
1	SORT ORDER BY
*2	HASH JOIN
3	TABLE ACCESS FULL
4	TABLE ACCESS FULL

Identify the two correct interpretations that can be made from the execution plan.

- A. The DEPT table is driving table and the EMP table join is the driven table.
- B. Rows from the DEPT table are first hashed by the join key into memory and then joined to the EMP table on the join key.
- C. The EMP table is the driving table and the DEPT table us the driven table.
- D. The rows from the DEPT table are sorted first by the join key and then hashed into memory.
- E. Rows from both the tables are sorted by the join key, but only rows from the DEPT table are hashed into memory.

Correct Answer: CD

If two lines are indented equally, then the top line is normally executed first. Here the line 3 and line 4 are indented equally, so line 3 (TABLE ACCESS FULL EMP) with EMP being the driving table as the ordered hint requests that the tables listed in the FROM clause of a SQL state- ment be joined in the order specified, with the first table in the FROM clause specifying the driving table.

QUESTION 3

Examine the following query and execution plan: Which query transformation technique is used in this scenario?

Examine the following query and execution plan:

```
SQL> SELECT prod_id prod_name
      FROM producers p
      Where prod_list_prices > (SELECT MAX(unit_cost)
      FROM costs c
      WHERE p.prod_id=c.prod_id);
```

Execution Plan

Plan hash value: 2297807746

Pstart	id Pstop	Operation	Name	Rows	Bytes	Cost	(%CPU)	Time
	0	SELECT STATEMENT		4	140	82	(8)	00:00:01
*	1	HASH JOIN		4	140	82	(8)	00:00:01
	2	VIEW	VW_SQL_1	72	1872	78	(7)	00:00:01
	3	HASH GROUP BY		72	648	78	(7)	00:00:01
	4	PARTITION RANGE ALL		82112	721K	74	(2)	00:00:01
1	28							
	5	TABLE ACCESS FULL	COSTS	82112	721K	74	(2)	00:00:01
1		28						
	6	TABLE ACCESS FULL	PRODUCTS	72	648	3	(0)	00:00:01

Predicate Information (Identified by operation id);

```
1- access ("P", "PROD_ID"=by operation id):
  Filter ("PROD_LIST_PRICE"> "MAX(UNIT_COST)")
```

- A. Join predicate push-down
- B. Subquery factoring
- C. Subquery unnesting
- D. Join conversion

Correct Answer: A

*

Normally, a view cannot be joined with an index-based nested loop (i.e., index access) join, since a view, in contrast with a base table, does not have an index defined on it. A view can only be joined with other tables using three methods: hash, nested loop, and sort-merge joins.

*

The following shows the types of views on which join predicate pushdown is currently supported. UNION ALL/UNION view Outer-joined view Anti-joined view Semi-joined view DISTINCT view GROUP-BY view

QUESTION 4

```
SQL> EXPLAIN PLAN SET Statement_id = 'test' for
      SELECT prod_category, avg(amount_sold)
      FROM sales s, products p
      WHERE p.prod_id = s.prod_id
      GROUP BY prod_Category;
```

Explained.

```
SQL> SELECT id "id", parent_id, position "pos"
      lpad(' ', 2 level) || operations || decode(id, 0, 'cost=' || POSITION) "operations"
Options "option" object_name "object"
FROM plan_table
Connect by prior id_parent_id START WITH id = 0
ORDER BY id;
```

id	PARENT_ID	POS	Operation	Option	Object
0		539	SELECT STATEMENT Cost = 539		
1	0	1	HASH	Group By	
2	1	1	HASH JOIN		
3	2	1	VIEW		
4	3	1	HASH	GROUP BY	
5	4	1	PARTITION RANGE	ALL	
6	5	1	TABLE ACCESS	FULL	SALES
7	2	2	VIEW		indes\$_joins\$_002
8	7	1	VIEW RANGE		
9	8	1	INDEX	FAST FULL SCAN	PRODUCTS_PK
10	8	2	INDEX	FAST FULL SCAN	PRODUCTS_PROD_CAT_IX

11 rows are selected

View the exhibit and examine the query and its execution plan from the PLAN_TABLE. Which statement is true about the execution?

- A. The row with the ID column having the value 0 is the first step execution plan.
- B. Rows are fetched from the indexes on the PRODUCTS table and from the SALES table using full table scan simultaneously, and then hashed into memory.
- C. Rows are fetched from the SALES table, and then a hash join operator joins with rows fetched from indexes on the PRODUCTS table.
- D. All the partitions of the SALES table are read in parallel.

Correct Answer: C

QUESTION 5

Examine the query:

```
SQL> SELECT /*+ RESULT_CACHE */ dept, AVG (sal)
FROM emp
GROUP BY deptno;
```

The RESULT_CACHE_MODE parameter is set to MANUAL for the database.

Which two statements are true about the usage of the result cache?

- A. The SQL runtime environment checks whether the query result is cached in the result cache; if the result exists, the optimizer fetches the result from it.
- B. The SQL runtime environment does check for the query result in the result cache because the RESULT_CACHE_MODE parameter is set to MANUAL.
- C. The SQL runtime environment checks for the query result in the result cache only when the query is executed for the second time.
- D. If the query result does not exist in the cache and the query is executed, the result is generated as output, and also sorted in the result cache.

Correct Answer: AD

Note:

*

result_cache_mode: the result cache can be enabled in three ways: via hint, alter session or alter system. Default is MANUAL which means that we need to explicitly request caching via the RESULT_CACHE hint;

*

As its name suggests, the query result cache is used to store the results of SQL queries for re- use in subsequent executions. By caching the results of queries, Oracle can avoid having to repeat the potentially time-consuming and intensive operations that generated the resultset in the first place (for example, sorting/ aggregation, physical I/O, joins etc). The cache results themselves are available across the instance (i.e. for use by sessions other than the one that first executed the query) and are maintained by Oracle in a dedicated area of memory. Unlike our homegrown solutions using associative arrays or global temporary tables, the query result cache is completely transparent to our applications. It is also maintained for consistency automatically, unlike our own caching programs.

*

RESULT_CACHE_MODE specifies when a ResultCache operator is spliced into a query's execution plan.

Values:

/ MANUAL

The ResultCache operator is added only when the query is annotated (that is, hints).

/FORCE

The ResultCache operator is added to the root of all SELECT statements (provided that it is valid to do so).

For the FORCE setting, if the statement contains a NO_RESULT_CACHE hint, then the hint takes precedence over the parameter setting.