

ISLEVER

# 1Z0-117

Oracle Database 11g Release 2: SQL Tuning  
Exam

DEMO

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

<https://www.islever.com/oracle.html>

For the most up-to-date exam questions and materials, we recommend visiting our website, where you can access the latest content and resources.

---

**QUESTION NO: 1**

Examine the query and its execution plan:

```
SQL > SELECT cust_last_name, sum (nu12(o.customer_id, 0, 1)) "Count"
      FROM customer c, orders o
      WHERE c.credit_limit > 1000
      AND c.customer_id = o.customer_id(+)
      GROUP By cust_last_name;
```

ID	Operations	Name	Rows	Bytes	Cost (%CPU)
0	SELECT STATEMENT		168	3192	6 (17)
1	HASH GROUP BY		168	3192	6 (17)
* 2	NESTED LOOPS OUTER		260	4940	5 (0)
* 3	TABLE ACCESS FULL	CUSTOMER	260	3900	5 (0)
* 4	INDEX RANGE SCAN	ORD_CUSTOMER_IX	105	420	0 (0)

Predicate Information (identified by operation id):

-----  
3 – filter ("C". "CREDIT\_LIMIT"> 1000)  
4 – access("C". "CUSTOMER\_ID" = "0" "CUSTOMER\_ID" (+))  
Filter ("O" "CUSTOMER\_ID" (+)>0)

Which statement is true regarding the execution plan?

- A.** This query first fetches rows from the CUSTOMERS table that satisfy the conditions, and then the join return NULL from the CUSTOMER\_ID column when it does not find any corresponding rows in the ORDERS table.
- B.** The query fetches rows from CUSTOMERS and ORDERS table simultaneously, and filters the rows that satisfy the conditions from the resultset.
- C.** The query first fetches rows from the ORDERS table that satisfy the conditions, and then the join returns NULL form CUSTOMER\_ID column when it does not find any corresponding rows in the CUSTOMERS table.
- D.** The query first joins rows from the CUSTOMERS and ORDERS tables and returns NULL for the ORDERS table columns when it does not find any corresponding rows in the ORDERS table, and then fetches the rows that satisfy the conditions from the result set.

**Answer: A**

**Explanation:**

**QUESTION NO: 2**

Which three statements are true about histograms?

- A.** They capture the distribution of different values in an index for better selectivity estimates.
- B.** They can be used only with indexed columns.

- 
- C. They provide metadata about distribution of and occurrences of values in a table column.
  - D. They provide improved selectivity estimates in the presence of data skew, resulting in execution plans with uniform distribution.
  - E. They help the optimizer in deciding whether to use an index or a full table scan.
  - F. They help the optimizer to determine the fastest table join order.

**Answer: C,E,F**

**Explanation:** C:A histogram is a frequency distribution (metadata) that describes the distribution of data values within a table.

E:It's well established that histograms are very useful for helping the optimizer choose between a full-scan and an index-scan.

F:Histograms may help the Oracle optimizer in deciding whether to use an index vs. a full-table scan (where index values are skewed) or help the optimizer determine the fastest table join order. For determining the best table join order, the WHERE clause of the query can be inspected along with the execution plan for the original query. If the cardinality of the table is too-high, then histograms on the most selective column in the WHERE clause will tip-off the optimizer and change the table join order.

Note:

\*The Oracle Query Optimizer uses histograms to predict better query plans. The ANALYZE command or DBMS\_STATS package can be used to compute these histograms.

Incorrect:

B:Histograms are NOT just for indexed columns.

– Adding a histogram to an un-indexed column that is used in a where clause can improve performance.

D:Histograms Opportunities

Any column used in a where clause with skewed data

Columns that are not queried all the time

Reduced overhead for insert, update, delete

**QUESTION NO: 3**

```
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.

**Answer: C**

**Explanation:**

#### QUESTION NO: 4

Which four statements are correct about communication between parallel execution process?

- A. The number of logical pathways between parallel execution producers and consumers depends on the degree parallelism.
- B. The shared pool can be used for parallel execution messages buffers.
- C. The large pool can be used for parallel execution messages buffers.
- D. The buffer cache can be used for parallel execution message buffers.
- E. Communication between parallel execution processes is never required if a query uses full partition-wise joins.