



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

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.geekcert.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

Refer to the Exhibit.

SQL> DESC stored

Name	Null?	Type
STORE_ID	NOT NULL	NUMBER (4)
STORE_NAME		VARCHAR2 (12)
STORE_ADDRESS		VARCHAR2(20)
START_DATE		DATE

SQL> DESC Sales

NAME	NULL?	TYPE
SALES_ID	NOT NULL	NUMBER(4)
ITEM_ID		NUMBER(4)
UQANTITY		NUMBER (10)
SALES_DATE		DATE
STORE_ID		NUMBER (4)

Execution plan: What must be the correct order of steps that the optimizer executes based on the ID column the execution plan?




Plan hash value: 29632623819

Id	Operation	Name	Rows	Bytes	Cost	(%CPU)
0	SELECT STATEMENT		3	189	10	(10)
1	NESTED LOOPS		3	189	10	(10)
2	NESTED LOOPS		3	141	7	(15)
*3	TABLE ACCESS FULL	EMPLOYEES	3	60	4	(25)
4	TABLE ACCESS BY INDEX ROWID	JOBS	19	513	2	(50)
*5	INDEX UNIQUE SCAN	JOB_ID_PK	1			
6	TABLE ACCESS BY INDEX ROWID	DEPARTMENTS	27	432	2	(50)
7	INDEX UNIQUE SCAN	DEPT_ID_PK	1			

PREDICATE Information (identified by operation id):

- 3 – filter (“E”. “EMPLOYEE\_ID” < 103)
- 5 – access (“E”. “JOB\_ID” = “J\_ID”)
- 7 – access (“E”. “DEPARTMENT\_ID” = “D”. “DEPARTMENT\_ID”)



- A. 3, 5, 4, 6, 7
- B. 3, 5, 4, 7, 6
- C. 3, 4, 5, 7, 6
- D. 4, 5, 3, 7, 6

Correct Answer: D

### QUESTION 2

Examine the Exhibit.



```
CREATE TABLE dept AS SELECT * FROM departments;
ALTER TABLE dept PARALLEL 2;
```

```
CREATE TABLE emp_range_did PARTITION BY RANGE (department_id)
(PARTITION emp_p1 VALUES LESS THAN (150),
PARTITION emp_p5 VALUES LESS THAN (MAXVALUE) )
AS SELECT * FROM employees;
```

```
ALTER TABLE emp_range_did PARALLEL 2;
```

```
EXPLAIN PLAN FOR
SELECT /*+ PQ_DISTRIBUTE (d NONE PARTITION) ORDERED */ e.last_name, d.department_name
FROM emp_range_did e, dept d
WHERE e.department_id = d.department_id;
```

Id	Operations	Name	Rows	Bytes	Cost	Pstart	Pstop
0	SELECT STATEMENT		284	16188	6		
1	PX COORDINATOR						
2	PX SEND QC (RANDOM) :	TQ10001	284	16188	6		
3	HASH JOIN		284	16188	6		
4	PX PARTITION RANGE ALL		284	7668	2	1	2
5	TABLE ACCESS FULL	EMP_RANGE_DID	284	7668	2	1	2
6	BUFFER SORT						
7	PX RECEIVE		21	630	2		
8	PX SEND PARTITION (KEY) :	TQ10000	21	630	2		
9	TABLE ACCESS FULL	DEPT	21	630	2		

Which two options are true about the execution plan and the set of statements?

- A. The query uses a partial partition-wise join.
- B. The degree of parallelism is limited to the number of partitions in the EMP\_RANGE\_DID table.
- C. The DEPT table id dynamically distributed based on the partition keys of the EMP\_RANGE\_DID table.
- D. The server process serially scans the entire DEPT table for each range partition on the EMP\_RANGE\_DID table.
- E. The query uses a full partition-wise join.

Correct Answer: AD

### QUESTION 3



Tracing has been enabled for the HR user. You execute the following command to check the contents of the orcl\_25052.trc trace file, which was generated during tracing:

```
Tkprof orcl_2505.trc output_tk.prf  
EXPLAIN = hr/hr TABLE = hr.temp_plan_table INSERT = script.sql SYS = NO  
SORT = (EXECUTE, FCHCPU)
```

Which two statements are correct about the execution of the command?

- A. SCRIPT.SQL stores the statistics for all traced SQL statements.
- B. Execution plans for SQL statements are stored in TEMP\_PLAN\_TABLE and can be queried by the user.
- C. SQL statements in the output files are stored in the order of elapsed time.
- D. TKPROF use TEMP\_PLAN\_TABLE in the HR schema as a temporary plan table.
- E. Recursive SQL statements are included in the output file.

Correct Answer: AD

INSERT Creates a SQL script that stores the trace file statistics in the database. TKPROF creates this script with the name filename3. This script creates a table and inserts a row of statistics for each traced SQL statement into the table.

#### QUESTION 4

How can you reduce fragmentation of an index without affecting the current transactions that are using the index?

- A. Use the ANALYZE INDEX . . . command
- B. Use the ALTER INDEX . . . VALIDATE STRUCTURE command
- C. Use the ALTER INDEX . . . REBUILD ONLINE command
- D. Use the ALTER INDEX . . . DEALLOCATE UNUSED command

Correct Answer: D

Use the deallocate\_unused\_clause to explicitly deallocate unused space at the end of the index and make the freed space available for other segments in the tablespace.

If index is range-partitioned or hash-partitioned, then Oracle Database deallocates unused space from each index partition. If index is a local index on a composite-partitioned table, then Oracle Database deallocates unused space from each index subpartition.

Reference: Oracle Database SQL Language Reference 11g, alter index

#### QUESTION 5

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.

Correct Answer: CEF

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

[Latest 1Z0-117 Dumps](#)

[1Z0-117 PDF Dumps](#)

[1Z0-117 Study Guide](#)