1Z0-117 Q&As

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.geekcert.com/1z0-117.html

100% Passing Guarantee
100% Money Back Assurance

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

QUESTION 1

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 FRO
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	TQ	IN-OUT	PQ DISTRIB
0	SELECT STATEMENT		284	16188	6					
1	PX COORDINATOR									
2	PX SEND QC (RANDOM) :	TQ10001	284	16188	6			Q1,01	P->S	QC (RAND)
* 3	HASH JOIN		284	16188	6			Q1,01	PCWP	
4	PX PARTITION RANGE ALL		284	7668	2	1	2	Q1,01	PCWC	
5	TABLE ACCESS FULL	EMP_RANGE_DID	284	7668	2	1	2	Q1,0	PCWP	
6	BUFFER SORT							Q1,01	PCWC	
7	PX RECEIVE		21	630	2			Q1,01	PCWP	
8	PX SEND PARTITION (KEY) :	TQ10000	21	630	2				S->P	PART (KEY)
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 2

Examine Exhibit1 to view the query and its AUTOTRACE output.

```
SQL> SET AUTOTRACE TRACEONLY
SQL> SELECT prod_category, AVG(amount_sold)
          FROM sales s, products P
WHERE P.prod_id = S.prod_id
GROUP BY prod_category;

Execution Plan
```

Plan hash volume: 1197568639

id	Operation	Name	Rows	Bytes	Cost (%CPU)	Time	Pstart
0	SELECT STATEMENT		5	255	639 (11)	00:00:07	
1	HASH GROUP BY		5	255	539 (11)	00:00:07	
2	HASH JOIN		72	3672	538 (11)	00:00:07	
3	VIEW	VW_GBC_5	72	2160	535 (11)	00:00:07	
4	HASH GROUP BY		72	648	535 (11)	00:00:07	
5	PARTITION TANGE ALL		918K	8075K	494 (3)	00:00:06	1
6	TABLE ACCESS FULL	SALES	918K	8075K	684 (3)	00:00:06	1
7	VIEW	INDEX$_JOIN_002	72	1512	3 (34)	00:00:01	
8	HASH JOIN						
9	INDEX FAST FULL SCAN	PRODUCT_PK	72	1512	1 (0)	00:00:01	
10	INDEX FAST FULL SCAN	PRODUCT_PROD_CAT_IX	72	1512	1 (0)	00:00:01	

```
Predicate information (identified by operation id)"
---------------------------------------------------
2- access ("P". "PROD_ID" = "ITEM_1")
8- access (ROEID=ROWID)

Statistics
----------------------------------------------------------
          0  recursive calls
          0  db block gets
       1726  consistent gets
          0  physical reads
          0  Redo size
        778  bytes sent via SQL "Net to client
        434  bytes received via SQL "Net from client
          2  SQL "Net roundtrips to/from client
          0  sorts (memory)
```

Which two statements are true about tracing?

A. The displayed plan will be stored in PLAN_TABLE.

B. Subsequent execution of this statement will use the displayed plan that is stored in v$SQL.

C. The displayed plan may not necessarily be used by the optimizer.

D. The query will not fetch any rows; it will display only the execution plan and statistics.

E. The execution plan generated can be viewed from v$SQLAREA.

Correct Answer: AD

The PLAN_TABLE is automatically created as a public synonym to a global temporary table. This temporary table holds

the output of EXPLAIN PLAN statements for all users. PLAN_TABLE is the default sample output table into which the EXPLAIN PLAN statement inserts rows describing execution plans

QUESTION 3

You plan to bulk load data using INSERT /*+PARALLEL*/ INTO . . . . SELECT FROM statements.

Which four types of operations can execute in parallel on tables that have no bitmapped indexes or materialized views defined on term?

A. Direct path insert of a million rows into a partitioned, index-organized table containing one million rows.

B. Direct path insert of a million rows into a partitioned, index-organized table containing 10 million rows.

C. Direct path insert of a million rows into a nonpartitioned, index-organized table containing one million rows.

D. Direct path insert of a million rows into a nonpartitioned, heap-organized table containing 10 million rows.

E. Direct path insert of a million rows into a nonpartitioned, heap-organized table containing one million rows.

Correct Answer: ABDE

Direct-path INSERT is not supported for an index-organized table (IOT) if it is not partitioned, if it has a mapping table, or if it is reference by a materialized view.

QUESTION 4

Which two tasks are performed during the optimization stage of a SQL statement?

A. Evaluating the expressions and conditions in the query

B. Checking the syntax and analyzing the semantics of the statement

C. Separating the clauses of the SQL statement into structures that can be processed

D. Inspecting the integrity constraints and optimizing the query based on this metadata

E. Gathering the statistics before creating the execution plan for the statement

Correct Answer: DE

Note:

Oracle SQL is parsed before execution, and a hard parse includes these steps: *

1.

Loading into shared pool - The SQL source code is loaded into RAM for parsing. (the "hard" parse step)

2.

Syntax parse - Oracle parses the syntax to check for misspelled SQL keywords.

3.

Semantic parse - Oracle verifies all table and column names from the dictionary and checks to see if you are authorized to see the data.

4.

Query Transformation - If enabled (query_rewrite=true), Oracle will transform complex SQL into simpler, equivalent forms and replace aggregations with

materialized views, as appropriate.

5.

Optimization - Oracle then creates an execution plan, based on your schema statistics (or maybe with statistics from dynamic sampling in 10g).

6.

Create executable - Oracle builds an executable file with native file calls to service the SQL query.

*

The parsing process performs two main functions:

o Syntax Check: is the statement a valid one. Does it make sense given the SQL grammar documented in the SQL Reference Manual. Does it follow all of the

rules for SQL.

o Semantic Analysis: Going beyond the syntax ? is the statement valid in light of the objects in the database (do the tables and columns referenced exist). Do you

have access to the objects ? are the proper privileges in place? Are there ambiguities in the statement ? for example if there are two tables T1 and T2 and both

have a column X, the query ?select X from T1, T2 where ?? is ambiguous, we don?t know which table to get X from. And so on.

So, you can think of parsing as basically a two step process, that of a syntax check to check the validity of the statement and that of a semantic check ? to ensure

the statement can execute properly.

Reference: Oracle hard-parse vs. soft parse

QUESTION 5

You need to migrate database from oracle Database 10g to 11g. You want the SQL workload to start the 10g plans in the 11g database instance and evolve better plans.

Examine the following steps:

1.

Capture the pre-Oracle Database 11g plans in a SQL Tuning Set (STS)

2.

Export the STS from the 10g system.

3.

Import the STS into Oracle Database 11g.

4.

Set the OPTIMIZER_FEATURES_ENABLE parameter to 10.2.0.

5.

Run SQL Performance Analyzer for the STS.

6.

Set the OPTIMIZER_FEATURES_ENABLE parameter to 11.2.0.

7.

Rerun the SQL Performance Analyzer for the STS.

8.

Set OPTIMIZER_CAPTURE_SQL_PLAN_BASELINE to TRUE.

9.

Use DBMS_SPM.EVOLVE_SQL_BASELINE function to evolve the plans.

10.

Set the OPTIMIZER_USE_SQL_PLAN_BASELINE to TRUE.

Identify the required steps in the correct order.

A. 1, 2, 3, 4, 5, 6, 7,

B. 4, 8, 10

C. 1, 2, 3, 4, 8, 10

D. 1, 2, 3, 6, 9, 5

E. 1, 2, 3, 5, 9, 10

Correct Answer: C

Step 1: (1)

Step 2: (2)

Step 3: (3)

Step 4: (4)

By setting the parameter OPTIMIZER_FEATURES_ENABLE to the 10g version used before the upgrade, you should be able to revert back to the same execution

plans you had prior to the upgrade.

Step 5: (8)

OPTIMIZER_CAPTURE_SQL_PLAN_BASELINES

In Oracle Database 11g a new feature called SQL Plan Management (SPM) has been introduced to guarantees any plan changes that do occur lead to better

performance. When OPTIMIZER_CAPTURE_SQL_PLAN_BASELINES is set to TRUE (default FALSE) Oracle will

automatically capture a SQL plan baseline for every repeatable SQL statement on the system. The execution plan found at parse time will be added to the SQL

plan baseline as an accepted plan.

Step 6: (10)

OPTIMIZER_USE_SQL_PLAN_BASELINES enables or disables the use of SQL plan baselines stored in SQL Management Base. When enabled, the optimizer

looks for a SQL plan baseline for the SQL statement being compiled. If one is found in SQL Management Base, then the optimizer will cost each of the baseline

plans and pick one with the lowest cost.