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.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

QUESTION 1

Examine the initializing parameters:

Name	Type	VALUE
Optimizer_dynamic_sampling	integer	2
Optimizer_index_caching	integer	50
Optimizer_index_cost_adj	integer	100
Optimizer_mode	string	ALL_ROWS
Optimizer_use_invisible_indexes	boolean	FALSE
Optimizer_use_pending_statistics	boolean	FALSE

An index exists on the column used in the WHERE of a query. You execute the query for the first time today and notice that the query is not using the index. The CUSTOMERS table has 55000 rows.

View the exhibit and examine the query and its execution plan.

```
SQL> SELECT * FROM customers WHERE cust_city_id = 51166;

Execution plan
----------------------------------------------------------
Plan hash value: 1351338989
-------------------------------------------------------------------------------------------------------------------------
Id  Operation          Name       Rows  Bytes  Cost  (%CPU)  Time
-------------------------------------------------------------------------------------------------------------------------
0   SELECT STATEMENT              437   79097  406   (1)     00:00:05
*1  TABLE ACCESS FULL  CUSTOMERS  437   79097  406   (1)     00:00:05
-------------------------------------------------------------------------------------------------------------------------

Predicate Information (identified by operation id):

1- filter ("CUST_CITY_ID"=51166)
```

What can be the two reasons for full table scan?

A. The value of the OPTIMIZER_INDEX_COST_ADJ parameter is set to a low value.

B. The blocks fetched by the query are greater than the value specified by the DB_FILE_MULTIBLOCK_READ_COUNT parameter.

C. The statistics for the CUSTOMERS table and the indexes stale.

D. The OPTIMIZER_MODE parameter is set to ALL_ROWS.

E. Histogram statistics for CUST_CITY_ID are missing.

F. Average number of rows per block for the CUSTOMERS table is low.

Correct Answer: CE

C: Old statistics could cause this problem. "Histograms are feature in CBO and it helps to optimizer to determine how data are skewed(distributed) with in the column. Histogram is good to create for the column which

are included in the WHERE clause where the column is highly skewed. Histogram helps to

optimizer to decide whether to use an index or full-table scan or help the optimizer

"

determine the fastest table join order.

QUESTION 2

View Exhibit1 and examine the structure and indexes for the MYSALES table.

Name	NULL?	Type
PROD_ID	NOT NULL	NUMBER
CUST_ID	NOT NULL	NUMBER
TIME_ID	NOT NULL	DATE
CHANNEL_ID	NOT NULL	NUMBER
PROMO_ID	NOT NULL	NUMBER
QUANTITY_ID	NOT NULL	NUMBER (10, 2)
AMOUNT_SOLD	NOT NULL	NUMBER (10, .)

TABLE_NAME	UNIQUENES	COLUMN_NAME	INDEX_NAME
MYSALES	NONUNIQUE	PPROD_ID	MYSALES_PRODIS_IDX
MYSALES	NONUNIQUE	PCUST_ID	MYSALES_PRODIS_IDX

The application uses the MYSALES table to insert sales record. But this table is also extensively used for generating sales reports. The PROD_ID and CUST_ID columns are frequently used in the WHERE clause of the queries. These columns have few distinct values relative to the total number of rows in the table. The MYSALES table has 4.5 million rows.

View exhibit 2 and examine one of the queries and its autotrace output.

```
SQL> select sum(AMOUNT_SOLD) from mysales where PROD_ID = 30 and CUST_ID=25939;

SUM(AMOUNT_SOLD)
----------------
          4137.6

Execution Plan
----------------------------------------------------------
Plan hash value: 3710121278
```

Id	Operation	Name	Rows	Bytes	Cost	(%CPU)	Time
0	SELECT STATEMENT		1	14	843	(1)	00:00:11
1	SORT AGGREGATE		1	14			
2	TABLE ACCESS BY INDEX ROWID	MYSALES	57	798	843	(1)	00:00:11
3	BITMAP CONVERSION TO WOWIDS				11	(0)	00"00"01
7	BITMAPCONVSERION FROM ROWIDS						
8	INDEX RANGE SCAN	MYSALES_PROD_IDX	4165		819	(1)	00:00:10

```
Predicate information (identified by operation id):
6- access ("CUST_ID"=25939)
6- access ("PROD_ID"=30)

Statistics
----------------------------------------------------------
1          recursive calls
0          db block gets
2202       consistent gets
104        physical reads
0          redo size
432        bytes sent via SQL*NET to client
420        bytes received via SQL*Net from client
2          SQL*NET roudtrips to/from client
0          sorts (memory)
1          rows processed
```

Which two methods can examine one of the queries and its autotrace output?

A. Drop the current standard balanced B* Tree indexes on the CUST_ID and PROD_ID columns and re-create as bitmapped indexes.

B. Use the INDEX_COMBINE hint in the query.

C. Create a composite index involving the CUST_ID and PROD_ID columns.

D. Rebuild the index to rearrange the index blocks to have more rows per block by decreasing the value for PCTFRE attribute.

E. Collect histogram statistics for the CUST_ID and PROD_ID columns.

Correct Answer: BC

B: The INDEX hint explicitly chooses an index scan for the specified table. You can use the INDEX hint for domain, B-tree, bitmap, and bitmap join indexes. However, Oracle recommends using INDEX_COMBINE rather than INDEX for the combination of multiple indexes, because it is a more versatile hint.

C: Combining the CUST_ID and PROD_ID columns into an composite index would improve performance.

QUESTION 3

Partial details of an execution plan.

```
SELECT STATEMENT
SORT GROUP BY
HASH JOIN
TABLE ACCESS FULL                          CHANNELS
HASH JOIN
TABLE ACCESS FULL                          CUSTOMERS
HASH JOIN
TABLE ACCESS FULL                          TIMES
PARTITION RANGE ITERATOR
TABLE ACCESS BY LOCAL INDEX ROWID          SALES
BITMAP CONVERSION TO ROWIDS
BITMAP AND
BITMAP MERGE
BITMAP KEY ITERATOR
BUFFER SORT
TABLE ACCESS FULL                          CHANNELS
BITMAP INDEX RANGE SCAN                    SALES_CHANNEL_BIX
BITMAP MERGE
BITMAP KEY ITERATION
BUFFER SORT
TABLE ACCESS FULL                          TIMES
BITMAP INDEX RANGE SCAN                    SALES_TIME_BIX
BITMAP MERGE
BITMAP KEY ITERATION
BUFFER SORT
TABLE ACCESS FULL                          TIMES
BITMAP INDEX RANGE SCAN                    SALES_TIME_BIX
```

Which statement correctly describes the BITMAP AND operation?

A. It produces a bitmap, representing dimension table rows from all dimension tables that join with qualified fact table rows.

B. It produces a concentration of the bitmaps for all dimension tables.

C. It produces a bitmap, representing fact table rows that do not join with qualified dimension table rows from all dimension tables.

D. It produces a bitmap, representing fact table rows that join with qualified dimension table rows from all dimension tables.

Correct Answer: D

Example:

Additional set operations will be done for the customer dimension and the product dimension. At this point in the star query processing, there are three bitmaps. Each bitmap corresponds to a separate dimension table, and each bitmap represents the set of rows of the fact table that satisfy that individual dimension\\'s constraints.

These three bitmaps are combined into a single bitmap using the bitmap AND operation. This final bitmap represents the set of rows in the fact table that satisfy all of the constraints on the dimension table.

Reference: Oracle Database Data Warehousing Guide, Star Transformation with a Bitmap Index

QUESTION 4

You notice some performance degradation for a high-load SQL statement in your database. After investigations, you run the SQL Tuning Advisor, which recommends a SQL Profile. You accept the profile recommendation resulting in a new, tuned execution plan for the statement.

Your database uses SQL plan management and a SQL plan baseline exists for this SQL statement.

Which statement is true?

A. The database adds the tuned plan to the SQL plan baseline as a nonfixed plan.

B. The database adds the tuned plan to the SQL plan baseline as a fixed plan.

C. The optimizer uses the new tuned plan only when a reproducible fixed plan is present.

D. The created SQL profile will continuously adapt to all changes made to the database, the object, and to the system statistics over an extended length of time.

Correct Answer: A

Note:

*

When the SQL Tuning Advisor recommends that a SQL Profile be used, you should accept the SQL Profile that is recommended. In cases where the SQL Tuning Advisor recommends that an index and a SQL Profile be used, both should be used. You can use the DBMS_SQLTUNE.ACCEPT_SQL_PROFILE procedure to accept a SQL Profile recommended by the SQL Tuning Advisor. This creates and stores a SQL Profile in the database.

*

When tuning SQL statements with the SQL Tuning Advisor, if the advisor finds a tuned plan and verifies its performance to be better than a plan chosen from the corresponding SQL plan baseline, it makes a recommendation to accept a SQL profile. When the SQL profile is accepted, the tuned plan is added to the corresponding SQL plan baseline.

*

If SQL plan management is used and there is already an existing plan baseline for the SQL statement, a new plan baseline will be added when a SQL profile is created.

*

SQL plan management is a preventative mechanism that records and evaluates the execution plans of SQL statements over time, and builds SQL plan baselines composed of a set of existing plans known to be efficient. The SQL plan

baselines are then used to preserve performance of corresponding SQL statements, regardless of changes occurring in the system.

*

SQL plan baseline is fixed if it contains at least one enabled plan whose FIXED attribute is set to YES.

*

ACCEPT_SQL_PROFILE Procedure and Function

This procedure creates a SQL Profile recommended by the SQL Tuning Advisor. The SQL text is normalized for matching purposes though it is stored in the data dictionary in de-normalized form for readability.

QUESTION 5

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