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

⚙ **Instant Download** After Purchase

⚙ **100% Money Back** Guarantee

⚙ **365 Days** Free Update

⚙ **800,000+** Satisfied Customers

QUESTION 1

Examine the Exhibit.

Id	Operation	Name	TQ	IN-OUT	PQ Distrib
0	SELECT STATEMENT				
1	PX COORDINATOR				
2	PX SEND QC (RANDOM)	:TQ10002	Q1, 02	P->S	QC (RAND)
*3	HASH JOIN BUFFERED		Q1, 02	PCWP	
4	PX JOIN FILTER CREATE	:BF0000	Q1, 02	PCWP	
5	PX RECEIVE		Q1, 02	PCWP	
6	PX SEND HASH	:TQ10000	Q1, 00	P->P	HASH
7	PX BLOCK ITERATOR		Q1, 00	PCWP	
*8	TABLE ACCESS FULL	T1	Q1, 00	PCWP	
9	PX RECEIVE		Q1, 02	PCWP	
10	PX SEND HASH	:TQ10001	Q1, 01	P->P	HASH
11	PX JOIN FILTER USE	:BF0000	Q1, 01	PCWP	
12	PX BLOCK ITERATOR		Q1, 01	PCWC	
13	TABLE ACCESS FULL	T2	Q1, 01	PCWP	

3 – access ("T1". "ID" = "T2". "ID")
8 – filter ("T1". "MOD" = 42)

Which two statements are true about the bloom filter in the execution plan?

A. The bloom filter prevents all rows from table T1 that do not join T2 from being needlessly distributed.

B. The bloom filter prevents all rows from table T2 that do not join table T1 from being needlessly distributed.

C. The bloom filter prevents some rows from table T2 that do not join table T1 from being needlessly distributed.

D. The bloom filter is created in parallel by the set of parallel execution processes that scanned table T2.

E. The bloom filter is created in parallel by the set of parallel execution processes that later perform join.

F. The bloom filter is created in parallel by the set of parallel execution processes that scanned table T1.

Correct Answer: BF

*

PX JOIN FILTER CREATE The bloom filter is created in line 4.

*

PX JOIN FILTER USE The bloom filter is used in line 11.

Note:

*

You can identify a bloom pruning in a plan when you see :BF0000 in the Pstart and Pstop columns of the execution plan and PART JOIN FILTER CREATE in the operations column.

*

A Bloom filter is a probabilistic algorithm for doing existence tests in less memory than a full list of keys would require. In other words, a Bloom filter is a method for representing a set of n elements (also called keys) to support membership queries.

*

The Oracle database makes use of Bloom filters in the following 4 situations:

-To reduce data communication between slave processes in parallel joins: mostly in RAC

-

 To implement join-filter pruning: in partition pruning, the optimizer analyzes FROM and WHERE clauses in SQL statements to eliminate unneeded partitions when building the partition access list

-

 To support result caches: when you run a query, Oracle will first see if the results of that query have already been computed and cached by some session or

user, and if so, it will retrieve the answer from the server result cache instead of gathering all of the database blocks

-

 To filter members in different cells in Exadata: Exadata performs joins between large tables and small lookup tables, a very common scenario for data warehouses with star schemas. This is implemented using Bloom filters as to determine

whether a row is a member of the desired result set.

QUESTION 2

You are administering a database that supports an OLTP workload. Automatic optimizer statistics collection is scheduled in the night maintenance window. Some of the tables get updated frequently during day time and you notice a performance degradation for queries using those tables due to stale statistics.

Which two options might help to avoid the performance degradation of the queries?

A. Set the global statistics preference STALE_PERCENT to 0.

B. Use dynamically sampling hint for the queries on frequently updated tables.

C. Create histogram for the columns used frequently in the WHERE clause.

D. Gather statistics with global statistics preference NO_VALIDATE to TRUE.

E. Set the OPTIMZER_USE_PENDING_STATISTICS parameter to TRUE.

Correct Answer: BC

B: Dynamic sampling first became available in Oracle9i Database Release 2. It is the ability of the cost-based optimizer (CBO) to sample the tables a query references during a hard parse, to determine better default statistics for unanalyzed segments, and to verify its "guesses." This sampling takes place only at hard parse time and is used to dynamically generate better statistics for the optimizer to use, hence the name dynamic sampling.

There are two ways to use dynamic sampling:

The OPTIMIZER_DYNAMIC_SAMPLING parameter can be set at the database instance level and can also be overridden at the session level with the ALTER

SESSION command.

The DYNAMIC_SAMPLING query hint can be added to specific queries.

C: A histogram is a collection of information about the distribution of values within a column.

In some cases, the distribution of values within a column of a table will affect the optimizer\\'s decision to use an index vs. perform a full-table scan. This scenario occurs when the value with a where clause has a disproportional amount of values, making a full-table scan cheaper than index access.

Histograms are also important for determinine the optimal table join order.

Incorrect:

A: Too much statistics would be gathered.

Note: STALE_PERCENT - This value determines the percentage of rows in a table that have to change before the statistics on that table are deemed stale and

should be regathered. The default value is 10%.

D: In Oracle PL/SQL, the VALIDATE keyword defines the state of a constraint on a column in a table.

E: OPTIMZER_USE_PENDING_STATISTICS specifies whether or not the optimizer uses pending statistics when compiling SQL statements.

QUESTION 3

Examine the Exhibit and view the query and its execution plan.

```
SQL>EXPLAIN PLAN FOR
        SELECT /*+ PARALLEL (4) */ customers.cust_first_name, customers.cust_last_name,
        MAX (QUANTITY_SOLD), AVG (QUANTITY_SOLD)
        FROM sales, customers
        WHERE sales, customers
        WHERE sales.cust_id=customers.cust_id
        GROUP By customers.cust_first_name, customers.cust_last_name;

Explained

PLAN_TABLE_OUTPUT
-----------------------------------------
Plan hash value: 4060011603
```
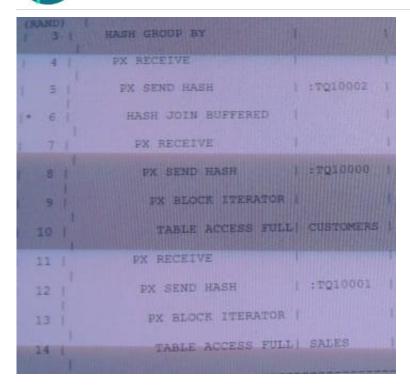
Id	Operation	Name	Rows	Bytes	TQ	IN-OUT	PQ	Distrib
0	SELECT STATEMENT		925	25900				
1	PX COORDINATOR							
2	PX SEND QC (RANDOM)	:TQ10003	925	25900	Q1,03	P->S	QC	RAND
3	HASH GROUP BY		925	25900	Q1,03	PCWP		
4	PX RECEIVE		925	25900	Q1,03	PCWP		
5	PX SEND HASH	:TQ10002	925	25900	Q1,02	P->P	HASH	
*6	HASH JOIN BUFFERED		925	25900	Q1,02	PCWP		
7	PX RECEIVE		630	12600	Q1,02	PCWP		
8	PX SEND HASH	:TQ10000	630	12600	Q1,00	P->P	PCWP	
9	PX BLOCK ITERATOR		630	12600	Q1,00	PCWP		
10	TABLE ACCESS FULL	CUSTOMERS	630	12600	Q1,00	PCWP		
11	PX RECEIVE		960	7680	Q1,02	PCWP		
12	PX SEND HASH	:TQ10001	960	7680	Q1,01	P->P	HASH	
13	PX BLOCK ITERATOR		960	7680	Q1,01	PCWC		
14	TABLE ACCESS FULL	SALES	960	7680	Q1,01			

```
------------------------------------------------------------------------------------------------
Predicate Information (identified by operation id):
6 - access ("SALES". "CUST_ID"= "CUSTOMERS". "CUST_ID" )
```

```
(RAND)
  3 |      HASH GROUP BY                    |         |
  4 |        PX RECEIVE                     |         |
  5 |         PX SEND HASH        | :TQ10002 |
* 6 |        HASH JOIN BUFFERED    |         |
  7 |          PX RECEIVE                   |         |
  8 |           PX SEND HASH      | :TQ10000 |
  9 |            PX BLOCK ITERATOR |         |
 10 |             TABLE ACCESS FULL| CUSTOMERS |
 11 |          PX RECEIVE                   |         |
 12 |           PX SEND HASH      | :TQ10001 |
 13 |            PX BLOCK ITERATOR |         |
 14 |             TABLE ACCESS FULL| SALES |
```

Which statement is correct about the parallel executions plan?

A. The CUSTOMERS and SALES tables are scanned simultaneously in parallel and then joined in parallel.

B. First, the CUSTOMERS table is scanned in parallel, then the SALES table is scanned in parallel, and then they are joined serially.

C. First, the SALES table is scanned in parallel, then the CUSTOMERS table us scanned in parallel, and then they are joined in parallel.

D. The CUSTOMERS and SALES tables are scanned simultaneously in parallel and then joined serially.

E. First, the CUSTOMERS table is scanned in parallel, then the SALES table us scanned in parallel, and then they are joined in parallel.

Correct Answer: A

As per exhibit:

Line 7 and line 11 are run in parallel.

Line 8 and line 12 are run in parallel.

Line 9 and line 13 are run in parallel.

Line 10 and line 14 are run in parallel.

Line 6 is a PCWP (parallel combined with parent) and the parent is a P-> P (Parallel to parallel) operation.

QUESTION 4

You enabled auto degree of parallelism (DOP) for your instance.

Examine the query:

```
SQL> SEECT /*+PARELLEL (AUTO)*/ customers.cust_first_name,
Customers.cust_last_name,
         MAX (QUANTITY_SOLD), AVG(QUANTITY_SOLD)
         FROM mysales, customers
         WHERE mysales.cust_id=customer.cust_id
         GROUP BY customers.cust_first_name, customers.cust_last_name;
```

Which two are true about the execution of this query?

A. Dictionary DOP will be used, if present, on the tables referred in the query.

B. DOP is calculated if the calculated DOP is 1.

C. DOP is calculated automatically.

D. Calculated DOP will always by 2 or more.

E. The statement will execute with auto DOP only when PARALLEL_DEGREE_POLICY is set to AUTO.

Correct Answer: AC

*

 PARALLEL (AUTO): The database computes the degree of parallelism (C), which can be 1 or greater (not D). If the computed degree of parallelism is 1, then the statement runs serially.

*

 You can use the PARALLEL hint to force parallelism. It takes an optional parameter: the DOP at which the statement should run. In addition, theNO_PARALLEL hint overrides a PARALLEL parameter in the DDL that created or altered the table.

The following example illustrates computing the DOP the statement should use: SELECT /*+ parallel(auto) */ ename, dname FROM emp e, dept d

WHERE e.deptno=d.deptno;

* When the parameter PARALLEL_DEGREE_POLICY is set to AUTO, Oracle Database auto- matically decides if a statement should execute in parallel or not and what DOP it should use. Oracle Database also determines if the statement can be executed immediately or if it is queued until more system resources are available. Finally, Oracle Database decides if the statement can take advantage of the aggregated cluster memory or not.

QUESTION 5

An application user complains about statement execution taking longer than usual. You find that the query uses a bind variable in the WHERE clause as follows:

```
SELECT prod_category, AVG(amount_sold)
FROM sales s, products p
WHERE p.prod_id = s.prod_id AND prod_category != :pcat
GROUP BY prod_category
```

You want to view the execution plan of the query that takes into account the value in the bind variable PCAT. Which two methods can you use to view the required execution plan?

A. Use the DBMS_XPLAN.DISPLAY function to view the execution plan.

B. Identify the SQL_ID for the statementsand use DBMS_XPLAN.DISPLAY_CURSOR for that SQL_ID to view the execution plan.

C. Identify the SQL_ID for the statement and fetch the execution plan PLAN_TABLE.

D. View the execution plan for the statement from V$SQL_PLAN.

E. Execute the statement with different bind values and set AUTOTRACE enabled for session.

Correct Answer: BD

D: V$SQL_PLAN contains the execution plan information for each child cursor loaded in the library cache.

B: The DBMS_XPLAN package supplies five table functions:

DISPLAY_SQL_PLAN_BASELINE - to display one or more execution plans for the SQL statement identified by SQL handle

DISPLAY - to format and display the contents of a plan table.

DISPLAY_AWR - to format and display the contents of the execution plan of a stored SQL statement in the AWR.

DISPLAY_CURSOR - to format and display the contents of the execution plan of any loaded cursor.

DISPLAY_SQLSET - to format and display the contents of the execution plan of statements stored in a SQL tuning set.

Latest 1Z0-117 Dumps   1Z0-117 PDF Dumps   1Z0-117 Braindumps