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.passapply.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

An application frequently executed similar types of queries that vary only in the usage of literals in the WHERE clause. You plan to use bind variable in place of literal values.

The CURSOR_SHARING parameter to set to EXACT.

Which two statements are true about the usage of bind variables?

A. The number of latch gets in the library cache will be reduced.

B. Bind peeking will take place and subsequent execution of queries can have different plans based on the cardinality of the bind value in the column.

C. Bind peeking will take place and subsequent execution of queries can have different plans only when the histograms exists on a column used in WHERE clause.

D. Bind peeking will not happen and the optimizer will use the same plan for all bind values if no histograms exist on a column use in the WHERE clause.

E. Bind peeking will happen and subsequent execution of queries will have the same parent cursor but different child cursors.

Correct Answer: AB

CURSOR_SHARING determines what kind of SQL statements can share the same cursors.

EXACT

Only allows statements with identical text to share the same cursor.

Note:

EXACT-This is the default setting. With this value in place, the query is not rewritten to use bind variables.

With CURSOR_SHARING=EXACT (the default), every unique SQL statement executed will create a new entry in V$SQL, it will be hard-parsed, and an execution plan will be created just for it. There can be hundreds or thousands of very similar queries in the shared pool that differ only in the literals used in the SQL statement itself. This implies that the database is forced to hard-parse virtually every query, which, in turn, not only consumes a lot of CPU cycles but also leads to decreased scalability.

The database just cannot hard-parse hundreds or thousands of SQL statements concurrently-- the application ends up waiting for the shared pool to become available. One of the major scalability inhibitors in the database is not using bind variables. That was the motivation behind adding CURSOR_SHARING=FORCE .

QUESTION 2

See the code fragment:

```
PARALLEL_DEGREE_POLICY=MANUAL
PARALLEL_MIN_PERCENT = 50
PARALLEL_MAX_SERVERS = 128
PARALLEL_MIN_SERVERS = 0
PARALLEL_DEGREE_LIMIT = 8
```

You execute the following query:

```
SELECT /*+ FULL (c, 8) */ c.cust_last_name, s.time_id, s.quantity_id,
s.quality_sold
FROM oe.customers c, sh.sales s
WHERE s.cust_id = c.customer_id;
```

You receive the following error message:

ORA-12827: insufficient parallel query slaves available

Which three parameter settings could you change to avoid this error?

A. Decrease the value of PARALLEL_MIN_PERCENT

B. Increase the value of PARALLEL_MAX_SERVERS

C. Increase the value of PARALLEL_MIN_SERVERS

D. Reduce the value of PARALLEL_MIN_TIME_THRESHOLF

E. Increase the value of PARALLEL_DEGREE_LIMIT

F. Set the PARALLEL_DEGREE_POLICY = AUTO

G. Set the PARALLEL_DEGREE_POLICY = LIMITED

Correct Answer: ABG

A: ORA-12827: insufficient parallel query slaves available Cause: PARALLEL_MIN_PERCENT parameter was specified and fewer than minimum slaves were acquired Action: either re-execute query with lower PARALLEL_MIN_PERCENT or wait until some running queries are completed, thus freeing up slaves

B: Your query doesn\\'t run because you\\'ve told Oracle not to run it unless at least 5% of the parallel execution processes are available for your query.

Set PARALLEL_MIN_PERCENT=0 or increase the number of parallel execution processes by increasing the PARALLEL_MAX_SERVERS parameter.

G: PARALLEL_DEGREE_POLICY

PARALLEL_DEGREE_POLICY specifies whether or not automatic degree of Parallelism, statement queuing, and in-memory parallel execution will be enabled.

LIMITED

Enables automatic degree of parallelism for some statements but statement queuing and in- memory Parallel Execution are disabled. Automatic degree of parallelism is only applied to those statements that access tables or indexes decorated explicitly with the PARALLEL clause. Tables and indexes that have a degree of parallelism specified will use that degree of parallelism.

Note: PARALLEL_MIN_PERCENT operates in conjunction with PARALLEL_MAX_SERVERS and PARALLEL_MIN_SERVERS. It lets you specify the minimum percentage of parallel execution processes (of the value of PARALLEL_MAX_SERVERS) required for parallel execution. Setting this parameter ensures that parallel operations will not execute sequentially unless adequate resources are available. The default value of 0 means that no minimum percentage of processes has been set.

Consider the following settings:

PARALLEL_MIN_PERCENT = 50 PARALLEL_MIN_SERVERS = 5 PARALLEL_MAX_SERVERS = 10

If 8 of the 10 parallel execution processes are busy, only 2 processes are available. If you then request a query with a degree of parallelism of 8, the minimum 50% will not be met.

You can use this parameter in conjunction with PARALLEL_ADAPTIVE_MULTI_USER. In a multi-user environment, an individual user or application can set PARALLEL_MIN_PERCENT to a minimum value until sufficient resources are available on the system and an acceptable degree of parallelism is returned.

QUESTION 3

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 4

Which three statements are true the Automatic Tuning Optimizer (ATO)?

A. It identifies the objects with stale or missing statistics and gathers statistics automatically.

B. It investigates the effect of new or modified indexes on the access paths for a workload and recommends running that statistics through the SQL Access Advisor.

C. It recommends a SQL profile to help create a better execution plan.

D. It picks up resource-intensive SQL statements from the ADDM and recommends the use of materialized views to improve query performance.

E. It identifies the syntactic, semantic, or design problems with structure of SQL statements leading to poor performance and suggests restricting the statements.

F. It identifies resource-intensive SQL statements, runs them through the SQL Tuning Advisor, and implements the recommendations automatically.

Correct Answer: ADF

Under tuning mode, the optimizer can take several minutes to tune a single statement. It is both time and resource intensive to invoke Automatic Tuning Optimizer every time a query must be hard-parsed. Automatic Tuning Optimizer is meant for complex and high-load SQL statements that have nontrivial impact on the database.

Automatic Database Diagnostic Monitor (ADDM) proactively identifies high-load SQL state- ments that are good candidates for SQL tuning. The automatic SQL tuning feature also automati- cally identifies problematic SQL statements

and implements tuning recommendations during system maintenance windows as an automated maintenance task.

The Automatic Tuning Optimizer performs the following types of tuning analysis:

Statistics Analysis SQL Profiling Access Path Analysis SQL Structure Analysis Alternative Plan Analysis

Note:

* Oracle Database uses the optimizer to generate the execution plans for submitted SQL state- ments. The optimizer operates in the following modes:

Normal mode The optimizer compiles the SQL and generates an execution plan. The normal mode generates a reasonable plan for the vast majority of SQL statements. Under normal mode, the optimizer op- erates with very strict time constraints, usually a fraction of a second.

Tuning mode

The optimizer performs additional analysis to check whether it can further improve the plan pro- duced in normal mode. The optimizer output is not an execution plan, but a series of actions, along with their rationale and expected benefit for producing a significantly better plan. When running in tuning mode, the optimizer is known as the Automatic Tuning Optimizer.

QUESTION 5

Exhibit

```
SQL SELECT id "id", parent_id, position "pos"
lpad(' ', "level )||operation||decode (id, 0, cost||POSITION" operation),
Operations "option" object_name "onject", object_node "table_queue",
Other_tag parallel oper type, distribution "row dist", other "slave SQL"
FROM plan_table
Connect by prior id=parent_id START WITH id=0
ORDER By id;
```

Id	par	pos	operations	option	object
0		4	SELECT STATEMENT cost=4		
1	0	1	HASH	GROUP BY	
2	1	1	NESTED LOOPS		
3	2	1	TABLE ACCESS	FULL	DEPARTMENTS
4	2	2	INDEX	RANGE SCAN	EMP_DEPARTMENT_IX

Examine the following SQL statement:

```
SQL> EXPLAIN PLAN FOR
SELECT department_name, count (')
FROM hr. employees e, hr.departments d
WHERE e.department_id=d.department_id
Group by.ddepartment_name;
```

Examine the exhibit to view the execution plan. Which statement is true about the execution plan?

A. The EXPLAIN PLAN generates the execution plan and stores it in c$SQL_PLAN after executing the query. Subsequent executions will use the same plan.

B. The EXPLAIN PLAN generates the execution plan and stores it in PLAN_TABLE without executing the query. Subsequent executions will always use the same plan.

C. The row with the ID 3 is the first step executed in the execution plan.

D. The row with the ID 0 is the first step executed in the execution plan.

E. The rows with the ID 3 and 4 are executed simultaneously.

Correct Answer: E

Note the other_tag parallel in the execution plan.

Note:

Within the Oracle plan_table, we see that Oracle keeps the parallelism in a column called other_tag. The other_tag column will tell you the type of parallel

operation that is being performed within your query.

For parallel queries, it is important to display the contents of the other_tag in the execution.

1Z0-117 VCE Dumps 1Z0-117 Practice Test 1Z0-117 Exam Questions