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.pass4itsure.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

Which two statements about In-Memory Parallel Execution are true?

A. It can be configured using the Database Resource Manager.

B. It increases the number of duplicate block images in the global buffer cache.

C. It requires setting PARALLEL_DEGREE_POLICY to LIMITED.

D. Objects selected for In-Memory Parallel Execution have blocks mapped to specific RAC instances.

E. It requires setting PARALLEL_DEGREE_POLICY to AUTO

F. Objects selected for In-Memory Parallel Execution must be partitioned tables or indexes.

Correct Answer: DE

D, E: In-Memory Parallel Execution

When the parameter PARALLEL_DEGREE_POLICY is set to AUTO, Oracle Database decides if an object that is accessed using parallel execution would benefit from being cached in the SGA (also called the buffer cache). The decision to cache an object is based on a well-defined set of heuristics including the size of the object and frequency on which it is accessed. In an Oracle RAC environment, Oracle Database maps pieces of the object into each of the buffer caches on the active instances. By creating this mapping, Oracle Database automatically knows which buffer cache to access to find different parts or pieces of the object. Using this information, Oracle Database prevents multiple instances from reading the same information from disk over and over again, thus maximizing the amount of memory that can cache objects. If the size of the object is larger than the size of the buffer cache (single instance) or the size of the buffer cache multiplied by the number of active instances in an Oracle RAC cluster, then the object is read using direct-path reads.

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

AUTO Enables automatic degree of parallelism, statement queuing, and in-memory parallel execution.

Incorrect:

C: 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.

Reference: Oracle Database VLDB and Partitioning Guide 11g, How Parallel Execution Works

QUESTION 2

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 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

Examine the Exhibit to view the structure of and indexes for the EMPLOYEES and DEPARTMENTS tables:

Name	Null?	Type
EMPLOYEE_ID	NOT NULL	NUMBER (6)
FIRST_NAME		VARCHAR2 (20)
LAST_NAME	NOT NULL	VARCHAR2 (25)
EMAIL	NOT NULL	VARCHAR2 (25)
PHONE_NUMBER		VARCHAR2 (20)
HIRE_DATE	NOT NULL	DATE
JOB_ID	NOT NULL	VARCHAR2 (10)
SALARY		NUMBER (8, 2)
COMISSION_PCT		NUMBER (2, 2)
MANAGER_ID		NUMBER (6)
DEPARTMENT_ID		NUMBER (4)

INDEX_NAME	INDEX_TYPE	COLUMN_NAME
EMP_NAME_IX	NORMAL	LAST_NAME
EMP_MANAGER_IX	NORMAL	MANAGER_ID
EMP_JOB_IX	NORMAL	JOB_ID
EMP_DEPARTMENT_IX	NORMAL	DEPARTMENT ID
EMP_EMP_PK	NORMAL	EMPLOYEE_ID
EMP_EMAIL_UK	NORMAL	EMAIL

Departments

Name	Null?	Type
DEPARTMENT_ID	NOT NULL	NUMBER (4)
DEPARTMENT_NAME	NOT NULL	VARCHAR (30)
MANAGER_ID		NUMBER (6)
LOCATION_ID		NUMBER (4)

EXAMINE the SQL statement and its execution plan:

```
SQL> EXPLAIN PLAN FOR
SELECT e.employee_id, j.job_title, e.salary, d.department_name
        FROM employees e, jobs j, departments d
        WHERE e.employees_id < 103
        AND e.job_id = j.job_id
        AND e.department_id = d.department_id;
```

EXPLAIN Plan output.

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	4	(50)
*5	INDEX UNIQUE SCAN	JOB_ID_PK	1			
6	TABLE ACCESS BY INDEX ROW ID	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". "JOB_ID")
7 – access ("E". "DEPARTMENT_ID". "DEPARTMENT_ID")
```

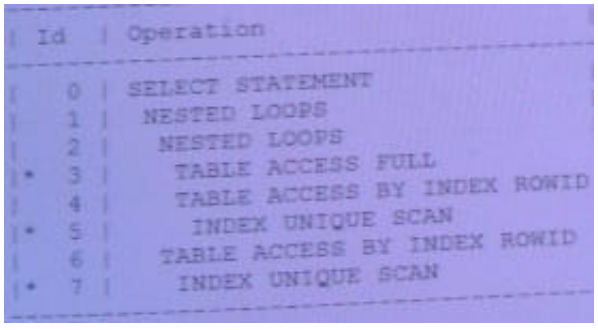

Which two statements are correct regarding the execution plan?

A. Step 2 is performing nested operation on JOB_ID column of the JOBS table, which is the driven table and the EMPLOYEES table is the driven table.

B. In step 2 for every row returned by the JOBS table matching rows from the EMPLOYEES table are accessed.

C. Step 1 is performing nested loop operation on the DEPARTMENT_ID column of the DEPARTMENTS table, which is the driven table and results returned by step 2 in the driving resultset.

D. The performance of the query can be improved by creating bitmap index on the JOB_ID column of the EMPLOYEES table.

E. The performance of the query can be improved by creating bitmapped index on the DEPARTMENT_ID column of the EMPLOYEES table.

Correct Answer: BE

As per exhibit:

B, not A, Not C: First is line 5 executed, followed by line 4, followed by line 3.

Step 2 is line 4.

E: The Department_ID column has lower cardinality compared to the JOB_ID column, so it is better suited for a bitmapped index.

Note:

*

Oracle bitmap indexes are very different from standard b-tree indexes. In bitmap structures, a two-dimensional array is created with one column for every row in the table being indexed. Each column represents a distinct value within the bitmapped index. This two-dimensional array represents each value within the index multiplied by the number of rows in the table.

At row retrieval time, Oracle decompresses the bitmap into the RAM data buffers so it can be rapidly scanned for matching values. These matching values are delivered to Oracle in the form of a Row-ID list, and these Row-ID values may directly access the required information.

*

The real benefit of bitmapped indexing occurs when one table includes multiple bitmapped indexes. Each individual column may have low cardinality. The creation of multiple bitmapped indexes provides a very powerful method for rapidly answering difficult SQL queries.

*

Oracle bitmap indexes are very different from standard b-tree indexes. In bitmap structures, a two-dimensional array is created with one column for every row in the table being indexed. Each column represents a distinct value within the bitmapped index. This two-dimensional array represents each value within the index multiplied by the number of rows in the table.

At row retrieval time, Oracle decompresses the bitmap into the RAM data buffers so it can be rapidly scanned for matching values. These matching values are delivered to Oracle in the form of a Row-ID list, and these Row-ID values may directly access the required information.

Latest 1Z0-117 Dumps 1Z0-117 VCE Dumps 1Z0-117 Braindumps