



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 three options are true about parallel queries when PARALLEL_DEGREE_POLICY is set to MANUAL and the session is using the default settings for parallel query, DDL, and DML?

- A. A subquery in a parallel DML is parallelized only if it includes a parallel hint.
- B. The number of parallel execution servers requested for a cursor is based on the greatest degree of parallelism associated with any object accessed by the cursor.
- C. A SELECT statement can be executed in parallel only if no scalar subqueries are contained in the SELECT list.
- D. In a CREATE TABLE . . . AS SELECT (CTAS) statement, SELECT is parallelized only if create TABLE is parallelized.
- E. In an INSERT INTO . . . SELECT FROM statement, INSERT is parallelized if select is parallelized.
- F. Single row inserts are never executed in parallel.

Correct Answer: CEF

* Decision to Parallelize

A SELECT statement can be parallelized only if the following conditions are satisfied:

/ The query includes a parallel hint specification (PARALLEL or PARALLEL_INDEX) or the schema objects referred to in the query have a PARALLEL declaration associated with them. / At least one of the tables specified in the query requires one of the following:

A full table scan

An index range scan spanning multiple partitions

/ (C) No scalar subqueries are in the SELECT list.

*

By default, the system only uses parallel execution when a parallel degree has been explicitly set on an object or if a parallel hint is specified in the SQL statement.

*

CREATE TABLE ... AS SELECT in Parallel

Parallel execution lets you parallelize the query and create operations of creating a table as a subquery from another table or set of tables. This can be extremely useful in the creation of summary or rollup tables.

Clustered tables cannot be created and populated in parallel.

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

MANUAL

Disables automatic degree of parallelism, statement queuing, and in-memory parallel execution. This reverts the



behavior of parallel execution to what it was prior to Oracle Database 11g Release 2 (11.2). This is the default.

Incorrect:

A:

*

For parallel DML (INSERT, UPDATE, MERGE, and DELETE), the reference object that determines the DOP (degree of parallelism) is the table being modified by an insert, update, or delete operation. Parallel DML also adds some limits to the DOP to prevent deadlock. If the parallel DML statement includes a subquery, the subquery's DOP is the same as the DML operation.

*

For parallel DDL, the reference object that determines the DOP is the table, index, or partition being created, rebuilt, split, or moved. If the parallel DDL statement includes a subquery, the subquery's DOP is the same as the DDL operation.

D: The CREATE TABLE ... AS SELECT statement contains two parts: a CREATE part (DDL) and a SELECT part (query). Oracle Database can parallelize both parts of the statement. The query part of a CREATE TABLE ... AS SELECT statement can be parallelized only if the following conditions are satisfied:

The query includes a parallel hint specification (PARALLEL or PARALLEL_INDEX) or the CREATE part of the statement has a PARALLEL clause specification or the schema objects referred to in the query have a PARALLEL declaration associated with them.

At least one of the tables specified in the query requires one of the following: a full table scan or an index range scan spanning multiple partitions.

Reference: Oracle Database VLDB and Partitioning Guide, Using Parallel Execution

QUESTION 2

While tuning a SQL statement, the SQL Tuning Advisor finds an existing SQL profile for a statement that has stale statistics. Automatic optimizer statistics is enabled for the database.

What does the optimizer do in this situation?

- A. Updates the existing SQL profiles for which the statistics are stale.
- B. Makes the statistics information available to GATHER_DATABASE_STATS_JOB_PROC
- C. Starts the statistics collection process by running GATHER_STATS_JOB
- D. Writes a warning message in the alert log file

Correct Answer: B

Automatic optimizer statistics collection calls the DBMS_STATS.GATHER_DATABASE_STATS_JOB_PROC procedure. This internal procedure operates similarly to the DBMS_STATS.GATHER_DATABASE_STATS procedure using the GATHER AUTO option. The main difference is that GATHER_DATABASE_STATS_JOB_PROCPrioritizes database objects that require statistics, so that objects that most need updated statistics are processed first, before the maintenance window closes.



Note:

* The optimizer relies on object statistics to generate execution plans. If these statistics are stale or missing, then the optimizer does not have the necessary information it needs and can generate poor execution plans. The Automatic Tuning Optimizer checks each query object for missing or stale statistics, and produces two types of output:

/ Recommendations to gather relevant statistics for objects with stale or no statistics

Because optimizer statistics are automatically collected and refreshed, this problem occurs only when automatic optimizer statistics collection is disabled. See "Managing Automatic Optimizer Statistics Collection".

/ Auxiliary statistics for objects with no statistics, and statistic adjustment factor for objects with stale statistics

The database stores this auxiliary information in an object called a SQL profile.

* Oracle recommends that you enable automatic optimizer statistics collection. In this case, the database automatically collects optimizer statistics for tables with absent or stale statistics. If fresh statistics are required for a table, then the database collects them both for the table and associated indexes.

Automatic collection eliminates many manual tasks associated with managing the optimizer. It also significantly reduces the risks of generating poor execution plans because of missing or stale statistics.

Automatic optimizer statistics collection calls the DBMS_STATS.GATHER_DATABASE_STATS_JOB_PROC procedure. This internal procedure operates similarly to the DBMS_STATS.GATHER_DATABASE_STATS procedure using the GATHER AUTO option. The main difference is that GATHER_DATABASE_STATS_JOB_PROC prioritizes database objects that require statistics, so that objects that most need updated statistics are processed first, before the maintenance window closes. Reference: Oracle Database Performance Tuning Guide, Managing Automatic Optimizer Statistics Collection

QUESTION 3

Examine the query:

```
SQL> SELECT /*+ RESULT_CACHE */ dept, AVG (sal)
FROM emp
GROUP BY deptno;
```

The RESULT_CACHE_MODE parameter is set to MANUAL for the database.

Which two statements are true about the usage of the result cache?

- A. The SQL runtime environment checks whether the query result is cached in the result cache; if the result exists, the optimizer fetches the result from it.
- B. The SQL runtime environment does check for the query result in the result cache because the RESULT_CACHE_MODE parameter is set to MANUAL.
- C. The SQL runtime environment checks for the query result in the result cache only when the query is executed for the second time.
- D. If the query result does not exist in the cache and the query is executed, the result is generated as output, and also sorted in the result cache.



Correct Answer: AD

Note:

*

result_cache_mode: the result cache can be enabled in three ways: via hint, alter session or alter system. Default is MANUAL which means that we need to explicitly request caching via the RESULT_CACHE hint;

*

As its name suggests, the query result cache is used to store the results of SQL queries for re- use in subsequent executions. By caching the results of queries, Oracle can avoid having to repeat the potentially time-consuming and intensive operations that generated the resultset in the first place (for example, sorting/ aggregation, physical I/O, joins etc). The cache results themselves are available across the instance (i.e. for use by sessions other than the one that first executed the query) and are maintained by Oracle in a dedicated area of memory. Unlike our homegrown solutions using associative arrays or global temporary tables, the query result cache is completely transparent to our applications. It is also maintained for consistency automatically, unlike our own caching programs.

*

RESULT_CACHE_MODE specifies when a ResultCache operator is spliced into a query's execution plan.

Values:

/ MANUAL

The ResultCache operator is added only when the query is annotated (that is, hints).

/FORCE

The ResultCache operator is added to the root of all SELECT statements (provided that it is valid to do so).

For the FORCE setting, if the statement contains a NO_RESULT_CACHE hint, then the hint takes precedence over the parameter setting.

QUESTION 4

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 OPTIMIZER_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 determine 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: OPTIMIZER_USE_PENDING_STATISTICS specifies whether or not the optimizer uses pending statistics when compiling SQL statements.

QUESTION 5

You are administering a database supporting an OLTP workload. A new module was added to one of the applications recently in which you notice that the SQL statements are highly resource intensive in terms of CPU, I/O and temporary space. You created a SQL Tuning Set (STS) containing all resource-intensive SQL statements. You want to analyze the entire workload captured in the STS. You plan to run the STS through the SQL Advisor.

Which two recommendations can you get?

- A. Combing similar indexes into a single index
- B. Implementing SQL profiles for the statements
- C. Syntactic and semantic restructuring of SQL statements
- D. Dropping unused or invalid index.
- E. Creating invisible indexes for the workload



F. Creating composite indexes for the workload

Correct Answer: CF

The output of the SQL Tuning Advisor is in the form of an advice or recommendations, along with a rationale for each recommendation and its expected benefit. The recommendation relates to collection of statistics on objects , creation of new indexes (F), restructuring of the SQL statement (C), or creation of a SQL profile. You can choose to accept the recommendation to complete the tuning of the SQL statements.

Reference: Oracle Database Performance Tuning Guide 11g , SQL Tuning Advisor

[Latest 1Z0-117 Dumps](#)

[1Z0-117 PDF Dumps](#)

[1Z0-117 VCE Dumps](#)