



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.passapply.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

In which three situations must you collect optimizer statistics manually for database objects in addition to automatic statistics collection?

- A. When substantial DML activity occurs between the nightly automatic stats gathering maintenance job
- B. When substantial activity occurs on a partition of the partitioned table.
- C. When a table is used for bulk loads that add 10% or more to the total size of the table
- D. When an index is created or dropped for a column
- E. When the degree of parallelism is explicitly defined for a table

Correct Answer: ABC

http://docs.oracle.com/cd/E11882_01/server.112/e23633/changes.htm#UPGRD12509

QUESTION 2

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 3

You ran a high load SQL statement that used an index through the SQL Tuning Advisor and accepted its recommendation for SQL profile creation. Subsequently you noticed that there has been a 2% growth in number of rows in the tables used by the SQL statement and database statistics have also been refreshed. How does this impact the created SQL profile?

- A. It becomes invalid and no longer used the optimizer.
- B. It remains valid and ensures that the optimizer always use the execution plan that was created before the changes happened.
- C. It remains and allows the optimizer to pick a different plan required.
- D. It becomes invalid and a new SQL profile is created for the statement by the auto tuning task.

Correct Answer: C

QUESTION 4

You plan to bulk load data INSERT INTO . . . SELECT FROM statements.

Which two situations benefit from parallel INSERT operations on tables that have no materialized views defined on them?



- A. Direct path insert of a million rows into a partitioned, index-organized table containing one million rows and a conventional B*tree secondary index.
- B. Direct path insert of a million rows into a partitioned, index-organized table containing 10 rows and a bitmapped secondary index.
- C. Direct path insert of 10 rows into a partitioned, index-organized table containing one million rows and conventional B* tree secondary index.
- D. Direct path insert of 10 rows into a partitioned, index-organized table containing 10 rows and a bitmapped secondary index
- E. Conventional path insert of a million rows into a nonpartitioned, heap-organized containing 10 rows and having a conventional B* tree index.
- F. Conventional path insert of 10 rows into a nonpartitioned, heap-organized table one million rows and a bitmapped index.

Correct Answer: AB

Note:

*

A materialized view is a database object that contains the results of a query.

*

You can use the INSERT statement to insert data into a table, partition, or view in two ways: conventional INSERT and direct-path INSERT.

*

With direct-path INSERT, the database appends the inserted data after existing data in the table. Data is written directly into datafiles, bypassing the buffer cache. Free space in the existing data is not reused. This alternative enhances performance during insert operations and is similar to the functionality of the Oracle direct-path loader utility, SQL*Loader. When you insert into a table that has been created in parallel mode, direct-path INSERT is the default.

*

Direct-path INSERT is not supported for an index-organized table (IOT) if it is not partitioned, if it has a mapping table, or if it is reference by a materialized view.

*

When you issue a conventional INSERT statement, Oracle Database reuses free space in the table into which you are inserting and maintains referential integrity constraints

*

Conventional INSERT always generates maximal redo and undo for changes to both data and metadata, regardless of the logging setting of the table and the archivelog and force logging settings of the database

QUESTION 5



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.

[1Z0-117 VCE Dumps](#)

[1Z0-117 Practice Test](#)

[1Z0-117 Braindumps](#)