



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

View Exhibit1 and examine the structure and indexes for the MYSALES table.

Name	NULL?	Type
PROD_ID	NOT NULL	NUMBER
CUST_ID	NOT NULL	NUMBER
TIME_ID	NOT NULL	DATE
CHANNEL_ID	NOT NULL	NUMBER
PROMO_ID	NOT NULL	NUMBER
QUANTITY_ID	NOT NULL	NUMBER (10, 2)
AMOUNT_SOLD	NOT NULL	NUMBER (10, .)

TABLE_NAME	UNIQUENES	COLUMN_NAME	INDEX_NAME
MYSALES	NONUNIQUE	PPROD_ID	MYSALES_PRODIS_IDX
MYSALES	NONUNIQUE	PCUST_ID	MYSALES_PRODIS_IDX

The application uses the MYSALES table to insert sales record. But this table is also extensively used for generating sales reports. The PROD_ID and CUST_ID columns are frequently used in the WHERE clause of the queries. These columns have few distinct values relative to the total number of rows in the table. The MYSALES table has 4.5 million rows.

View exhibit 2 and examine one of the queries and its autotrace output.



SQL> select sum(AMOUNT_SOLD) from mysales where PROD_ID = 30 and CUST_ID=25939;

```
SUM(AMOUNT_SOLD)
-----
                4137.6
```

Execution Plan

Plan hash value: 3710121278

Id	Operation	Name	Rows	Bytes	Cost	(%CPU)	Time
0	SELECT STATEMENT		1	14	843	(1)	00:00:11
1	SORT AGGREGATE		1	14			
2	TABLE ACCESS BY INDEX ROWID	MYSALES	57	798	843	(1)	00:00:11
3	BITMAP CONVERSION TO ROWIDS				11	(0)	00:00:01
7	BITMAP CONVERSION FROM ROWIDS						
8	INDEX RANGE SCAN	MYSALES_PROD_IDX	4165		819	(1)	00:00:10

Predicate information (identified by operation id):

```
6- access ("CUST_ID"=25939)
6- access ("PROD_ID"=30)
```

Statistics

```
1 recursive calls
0 db block gets
2202 consistent gets
104 physical reads
0 redo size
432 bytes sent via SQL*NET to client
420 bytes received via SQL*Net from client
2 SQL*NET roundtrips to/from client
0 sorts (memory)
1 rows processed
```

Which two methods can examine one of the queries and its autotrace output?

- A. Drop the current standard balanced B* Tree indexes on the CUST_ID and PROD_ID columns and re-create as bitmapped indexes.
- B. Use the INDEX_COMBINE hint in the query.
- C. Create a composite index involving the CUST_ID and PROD_ID columns.
- D. Rebuild the index to rearrange the index blocks to have more rows per block by decreasing the value for PCTFRE attribute.
- E. Collect histogram statistics for the CUST_ID and PROD_ID columns.

Correct Answer: BC

B: The INDEX hint explicitly chooses an index scan for the specified table. You can use the INDEX hint for domain, B-tree, bitmap, and bitmap join indexes. However, Oracle recommends using INDEX_COMBINE rather than INDEX for the combination of multiple indexes, because it is a more versatile hint.

C: Combining the CUST_ID and PROD_ID columns into an composite index would improve performance.



QUESTION 2

Which two statements are true about the trcsess utility?

- A. It merges multiple trace files and produces a formatted output file.
- B. It merges multiple trace files from a particular session into one single trace file.
- C. It produces multiple files only for DBA sessions, which can be consolidated into one formatted file using the tkprof utility.
- D. It produces multiple files for a service, which can be consolidated into one formatted file using the tkprof utility.
- E. It merges files pertaining to a user session scattered across different processes in a shared server configuration.

Correct Answer: AB

The trcsess utility consolidates trace output from selected trace files based on several criteria:

Session ID Client ID Service name Action name Module name After trcsess merges the trace information into a single output file, the output file could be processed by TKPROF.

Note:

*

trcsess is useful for consolidating the tracing of a particular session for performance or debugging purposes. Tracing a specific session is usually not a problem in the dedicated server model as a single dedicated process serves a session during its lifetime. You can see the trace information for the session from the trace file belonging to the dedicated server serving it. However, in a shared server configuration a user session is serviced by different processes from time to time. The trace pertaining to the user session is scattered across different trace files belonging to different processes. This makes it difficult to get a complete picture of the life cycle of a session.

Reference: Oracle Database Performance Tuning Guide 11g , Using the trcsess Utility

*

Now there is a new tool, a command line utility called trcsess to help read the trace files. The trcsess command-line utility consolidates trace information from selected trace files, based on specified criteria. The criteria include session id, client id, service name, action name and module name.

QUESTION 3

You are logged in as the HR user and you execute the following procedure:

```
SQL > exec DBMS_STATS.SET_TABLE_PREFS ('HR\','EMPLOYEES\','PUBLISH\','FALSE\'); SQL> exec DBMS_STATS.GATHER_TABLE_STATS ('HR\','EMPLOYEES\');
```

Which statement is true about the newly gathered statistics?

- A. They are temporary and purged when the session exits.
- B. They are used by the optimizer for all sessions.



- C. They are locked and cannot be overwritten.
- D. They are marked as pending and stored in the pending statistics table.

Correct Answer: D

In previous database versions, new optimizer statistics were automatically published when they were gathered. In 11g this is still the default action, but you now have the option of keeping the newly gathered statistics in a pending state until you choose to publish them.

The DBMS_STATS.GET_PREFS function allows you to check the `\\PUBLISH\\` attribute to see if statistics are automatically published. The default value of TRUE means they are automatically published, while FALSE indicates they are held in a pending state.

Reference: Statistics Collection Enhancements in Oracle Database 11g Release 1, Pending Statistics

QUESTION 4

How can you reduce fragmentation of an index without affecting the current transactions that are using the index?

- A. Use the ANALYZE INDEX . . . command
- B. Use the ALTER INDEX . . . VALIDATE STRUCTURE command
- C. Us the ALTER INDEX . . . REBUILD ONLINE command
- D. Use the ALTER INDEX . . . DEALLOCATE UNUSED command

Correct Answer: D

Use the `deallocate_unused_clause` to explicitly deallocate unused space at the end of the index and make the freed space available for other segments in the tablespace.

If index is range-partitioned or hash-partitioned, then Oracle Database deallocates unused space from each index partition. If index is a local index on a composite-partitioned table, then Oracle Database deallocates unused space from each index subpartition.

Reference: Oracle Database SQL Language Reference 11g, alter index

QUESTION 5

You have created some materialized views to improve the performance of several queries.

Which four must be true to enable sessions to benefit from improved response time made possible by these materialized views?

- A. Query rewrite must be enabled for the sessions.
- B. Bitmap indexes must exist on all the columns involved in the join operations for the defining query of the MVIEWs.
- C. All or part of the query results must be obtainable from one or more MVIEWs.
- D. Bitmap join indexes must exist on all the columns involved in the join operations.



- E. Session users must have query rewrite privilege.
- F. The MVIEWS must be enabled for query rewrite.
- G. All or part of the query results must be obtainable from one MVIEW.

Correct Answer: ABCF

A: For a given user's session, ALTER SESSION can be used to disable or enable query rewrite for that session only.

B: Bitmap indexes on the join columns would improve performance.

C (not G) : One of the major benefits of creating and maintaining materialized views is the ability to take advantage of query rewrite, which transforms a SQL statement expressed in terms of tables or views into a statement accessing one or more materialized views that are defined on the detail tables.

F:

*

A materialized view is only eligible for query rewrite if the ENABLE QUERY REWRITE clause has been specified, either initially when the materialized view was first created or subsequently with an ALTER MATERIALIZED VIEW statement.

*

Enabling or disabling query rewrite: by the CREATE or ALTER statement for individual materialized views by the initialization parameter QUERY_REWRITE_ENABLED by the REWRITE and NOREWRITE hints in SQL statements

Note:

*

A materialized view is a replica of a target master from a single point in time. The master can be either a master table at a master site or a master materialized view at a materialized view site. Whereas in multimaster replication tables are continuously updated by other master sites, materialized views are updated from one or more masters through individual batch updates, known as a refreshes, from a single master site or master materialized view site.

[Latest 1Z0-117 Dumps](#)

[1Z0-117 Practice Test](#)

[1Z0-117 Exam Questions](#)