

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.leads4pass.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

Refer to the Exhibit.

SQL> DESC stored

Name	Null?	Type
STORE_ID	NOT NULL	NUMBER (4)
STORE_NAME		VARCHAR2 (12)
STORE_ADDRESS		VARCHAR2(20)
START_DATE		DATE

SQL> DESC Sales

NAME	NULL?	TYPE
SALES_ID	NOT NULL	NUMBER(4)
ITEM_ID		NUMBER(4)
UQANTITY		NUMBER (10)
SALES_DATE		DATE
STORE_ID		NUMBER (4)


Execution plan: What must be the correct order of steps that the optimizer executes based on the ID column the execution plan?

Plan hash value: 29632623819

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	2	(50)
*5	INDEX UNIQUE SCAN	JOB_ID_PK	1			
6	TABLE ACCESS BY INDEX ROWID	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_ID”)
- 7 – access (“E”. “DEPARTMENT_ID” = “D”. “DEPARTMENT_ID”)



- A. 3, 5, 4, 6, 7
- B. 3, 5, 4, 7, 6
- C. 3, 4, 5, 7, 6
- D. 4, 5, 3, 7, 6

Correct Answer: D

QUESTION 2

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.

QUESTION 3

```
SQL> EXPLAIN PLAN SET Statement_id = 'test' for
      SELECT prod_category, avg(amount_sold)
      FROM sales s, products p
      WHERE p.prod_id = s.prod_id
      GROUP BY prod_Category;
```

Explained.

```
SQL> SELECT id "id", parent_id, position "pos"
      lpad(' ', 2 level) || operations || decode (id, 0, 'cost=' || POSITION) "operations"
Options "option" object_name "object"
FROM plan_table
Connect by prior id_parent_id START WITH id = 0
ORDER BY id;
```

id	PARENT_ID	POS	Operation	Option	Object
0		539	SELECT STATEMENT Cost = 539		
1	0	1	HASH	Group By	
2	1	1	HASH JOIN		
3	2	1	VIEW		
4	3	1	HASH	GROUP BY	
5	4	1	PARTITION RANGE	ALL	
6	5	1	TABLE ACCESS	FULL	SALES
7	2	2	VIEW		indes\$_joins\$_002
8	7	1	VIEW RANGE		
9	8	1	INDEX	FAST FULL SCAN	PRODUCTS_PK
10	8	2	INDEX	FAST FULL SCAN	PRODUCTS_PROD_CAT_IX

11 rows are selected

View the exhibit and examine the query and its execution plan from the PLAN_TABLE. Which statement is true about the execution?

- A. The row with the ID column having the value 0 is the first step execution plan.
- B. Rows are fetched from the indexes on the PRODUCTS table and from the SALES table using full table scan simultaneously, and then hashed into memory.
- C. Rows are fetched from the SALES table, and then a hash join operator joins with rows fetched from indexes on the PRODUCTS table.
- D. All the partitions of the SALES table are read in parallel.

Correct Answer: C

QUESTION 4

Examine the Exhibit and view the structure of an indexes for the EMPLOYEES table.

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

INDEX_NAME	INDEX_TYPE	COLUMN_NAME	UNIQUENES	TABLE_NAME
EMP_EMAIL_UK	NORMAL	EMAIL	UNIQUE	EMPLOYEES
EMP_EMP_ID_PK	NORMAL	EMPLOYEE_ID	UNIQUE	EMPLOYEES
EMP_DEPARTMENT_IX	NORMAL	DEPARTMENT_ID	NONUNIQUE	EMPLOYEES
EMP_JOB_IX	NORMAL	JOB_ID	NONUNIQUE	EMPLOYEES
EMP_MANAGER_IX	NORMAL	MANAGER	NONUNIQUE	EMPLOYEES

Examine the output rkprof:

```
SQL > SELECT employees_id, last_name, salary, department_id
FROM employees
WHERE employees_id = 126;
```

Calls	Count	CPU	Elapsed	Disk	Query	Current	Rows
Parse	1	0.01	0.01	0	0	0	0
Execute	1	0.00	0.00	0	3	0	0
Fetch	2	0.020	.55	453	797	0	1
Total	1	0.21	0.62	453	797	0	1

Which two actions might improve the performance of the query?

- A. Use the ALL_ROWS hint in the query.
- B. Collect the histogram statistics for the EMPLOYEE_ID column.
- C. Decrease the value for the DB_FILE_MULTIBLOCK_READ_COUNT initialization parameter.
- D. Decrease the index on the EMPLOYEE_ID if not being used.
- E. Set the OPTIMIZER_MODE parameter to ALL_ROWS.

Correct Answer: AE

A: The ALL_ROWS hint instructs the optimizer to optimize a statement block with a goal of best throughput, which is minimum total resource consumption.

E: optimizer_mode=all_rows - This optimizer mode favors full-table scans (especially parallel full-table-scans) in cases

where the server resources will be minimized. The all_rows mode is generally used during batch-oriented processing and for data warehouses where the goal is to minimize server resource consumption.


QUESTION 5

Examine the Exhibit.

Id	Operation	Name	TQ	IN-OUT	PQ Distrib
0	SELECT STATEMENT				
1	PX COORDINATOR				
2	PX SEND QC (RANDOM)	:TQ10002	Q1, 02	P->S	QC (RAND)
*3	HASH JOIN BUFFERED		Q1, 02	PCWP	
4	PX JOIN FILTER CREATE	:BF0000	Q1, 02	PCWP	
5	PX RECEIVE		Q1, 02	PCWP	
6	PX SEND HASH	:TQ10000	Q1, 00	P->P	HASH
7	PX BLOCK ITERATOR		Q1, 00	PCWP	
*8	TABLE ACCESS FULL	T1	Q1, 00	PCWP	
9	PX RECEIVE		Q1, 02	PCWP	
10	PX SEND HASH	:TQ10001	Q1, 01	P->P	HASH
11	PX JOIN FILTER USE	:BF0000	Q1, 01	PCWP	
12	PX BLOCK ITERATOR		Q1, 01	PCWC	
13	TABLE ACCESS FULL	T2	Q1, 01	PCWP	

3 – access (“T1”. “ID” = “T2”. “ID”)

8 – filter (“T1”. “MOD” = 42)



Which two statements are true about the bloom filter in the execution plan?

- A. The bloom filter prevents all rows from table T1 that do not join T2 from being needlessly distributed.
- B. The bloom filter prevents all rows from table T2 that do not join table T1 from being needlessly distributed.
- C. The bloom filter prevents some rows from table T2 that do not join table T1 from being needlessly distributed.
- D. The bloom filter is created in parallel by the set of parallel execution processes that scanned table T2.
- E. The bloom filter is created in parallel by the set of parallel execution processes that later perform join.
- F. The bloom filter is created in parallel by the set of parallel execution processes that scanned table T1.

Correct Answer: BF

*

PX JOIN FILTER CREATE The bloom filter is created in line 4.

*

PX JOIN FILTER USE The bloom filter is used in line 11.

Note:

*

You can identify a bloom pruning in a plan when you see :BF0000 in the Pstart and Pstop columns of the execution plan and PART JOIN FILTER CREATE in the operations column.

*

A Bloom filter is a probabilistic algorithm for doing existence tests in less memory than a full list of keys would require. In other words, a Bloom filter is a method for representing a set of n elements (also called keys) to support membership queries.

*

The Oracle database makes use of Bloom filters in the following 4 situations:

-To reduce data communication between slave processes in parallel joins: mostly in RAC

-

To implement join-filter pruning: in partition pruning, the optimizer analyzes FROM and WHERE clauses in SQL statements to eliminate unneeded partitions when building the partition access list

-

To support result caches: when you run a query, Oracle will first see if the results of that query have already been computed and cached by some session or

user, and if so, it will retrieve the answer from the server result cache instead of gathering all of the database blocks

-

To filter members in different cells in Exadata: Exadata performs joins between large tables and small lookup tables, a very common scenario for data warehouses with star schemas. This is implemented using Bloom filters as to determine

whether a row is a member of the desired result set.

[Latest 1Z0-117 Dumps](#)

[1Z0-117 PDF Dumps](#)

[1Z0-117 Study Guide](#)