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

QUESTION 1

Examine the query and its execution plan:

```
SQL > SELECT cust_last_name, sum (nu12(o.customer_id, 0, 1)) "Count"
          FROM customer c, orders o
          WHERE c.credit_limit > 1000
          AND c.customer_id = o.customer_id(+)
          GROUP By cust_last_name;
```

ID	Operations	Name	Rows	Bytes	Cost (%CPU)
0	SELECT STATEMENT		168	3192	6 (17)
1	HASH GROUP BY		168	3192	6 (17)
* 2	NESTED LOOPS OUTER		260	4940	5 (0)
* 3	TABLE ACCESS FULL	CUSTOMER	260	3900	5 (0)
* 4	INDEX RANGE SCAN	ORD_CUSTOMER_IX	105	420	0 (0)

```
Predicate Information (identified by operation id):
---------------------------------------------------

3 – filter ("C". "CREDIT_LIMIT"> 1000)
4 – access("C". "CUSTOMER_ID" = "0" "CUSTOMER_ID" (+))
    Filter ("O" "CUSTOMER_ID" (+)>0)
```

Which statement is true regarding the execution plan?

A. This query first fetches rows from the CUSTOMERS table that satisfy the conditions, and then the join return NULL from the CUSTOMER_ID column when it does not find any corresponding rows in the ORDERS table.

B. The query fetches rows from CUSTOMERS and ORDERS table simultaneously, and filters the rows that satisfy the conditions from the resultset.

C. The query first fetches rows from the ORDERS table that satisfy the conditions, and then the join returns NULL form CUSTOMER_ID column when it does not find any corresponding rows in the CUSTOMERS table.

D. The query first joins rows from the CUSTOMERS and ORDERS tables and returns NULL for the ORDERS table columns when it does not find any corresponding rows in the ORDERS table, and then fetches the rows that satisfy the conditions from the result set.

Correct Answer: A

QUESTION 2

Your database has the OLTP_SRV service configured for an OLTP application running on a middle tier. This service is used to connect to the database by using

connection pools. The application has three modules. You enabled tracing at the service by executing the following

command:

SQL exec DBMS_MONITOR.SERV_MOD_ACT_TRACE_ENABLE (`OLTP_SRV\\');

What is the correct method of consolidating the trace files generated by the procedure?

A. Use all trace files as input for the tkprof utility to consolidate the trace files for a module.

B. Use one trace file at a time as input for the trcess utility and use tkprof utility to consolidate all the output files for a module.

C. Use the trcess utility to consolidate all trace files into a single output file, which can then be processed by the tkprof utility.

D. Use the tkprof utility to consolidate the trace files and create an output that can directly be used for diagnostic purposes.

Correct Answer: C

Note:

* Oracle provides the trcsess command-line utility that consolidates tracing information based on specific criteria.

The SQL Trace facility and TKPROF are two basic performance diagnostic tools that can help you monitor applications running against the Oracle Server.

Note: SERV_MOD_ACT_TRACE_ENABLE Procedure

Enables SQL tracing for a given combination of Service Name, MODULE and ACTION globally unless an instance_name is specified

Reference: Oracle Database Performance Tuning Guide

QUESTION 3

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 4

Examine the Exhibit1 to view the structure of an indexes for the EMPLOYEES table.

```
SQL> desc employees

Name                          Null?          Type
----------------------------------------------------------------------------------
EMPLOYEE_ID                   NOT NULL       NUMBER (6)
FIRST_NAME                                   VARCHAR 2(20)
LAST_NAME                     NOT NULL       VARCHAR 2(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)

SQL> select index_name, index_type from user_indexes where table_name = 'EMPLOYEES'

INDEX_NAME                    INDEX_TYPE
-----------------------------------------------------------
EMP_NAME_IX                   NORMAL
EMP_MANAGER_IX                NORMAL
EMP_JOB_IX                    NORMAL
EMP_DEPARTMENT_IX             NORMAL
EMP_EMP_ID_PK                 NORMAL
EMP_EMAIL_UK                  NORMAL
```

Examine the query:

SQL> SELECT * FROM employees WHERE employees_id IN (7876, 7900, 7902);

EMPLOYEE_ID is a primary key in the EMPLOYEES table that has 50000 rows.

Which statement is true regarding the execution of the query?

A. The query uses an index skip scan on the EMP_EMP_ID_PK index to fetch the rows.

B. The query uses the INLIST ITERATOR operator to iterate over the enumerated value list, and values are evaluated using an index range scan on the EMP_EMP_ID_PK index.

C. The query uses the INLIST ITERATOR operator to iterate over the enumerated value list, and values are evaluated using a fast full index scan on the EMP_EMP_ID_PK index.

D. The query uses the INLIST ITERATOR operator to iterate over the enumerated value list, and values are evaluated using an index unique scan on the EMP_EMP_ID_PK index.

E. The query uses a fast full index scan on the EMP_EMP_ID_PK index fetch the rows.

Correct Answer: B

How the CBO Evaluates IN-List Iterators

The IN-list iterator is used when a query contains an IN clause with values. The execution plan is identical to what would result for a statement with an equality clause instead of IN except for one additional step. That extra step occurs when the IN-list iterator feeds the equality clause with unique values from the IN-list.

Both of the statements in Example 2-1 and Example 2-1 are equivalent and produce the same plan.

Example 2-1 IN-List Iterators Initial Statement

SELECT header_id, line_id, revenue_amount FROM so_lines_all WHERE header_id IN (1011,1012,1013);

SELECT header_id, line_id, revenue_amount FROM so_lines_all WHERE header_id = 1011 OR header_id = 1012 OR header_id = 1013;

Plan

SELECT STATEMENT INLIST ITERATOR TABLE ACCESS BY INDEX ROWID SO_LINES_ALL INDEX RANGE SCAN SO_LINES_N1

Reference: Database Performance Tuning Guide and Reference

QUESTION 5

Which two statements are true about index full scans?

A. An index fast full scan multi block I/O to read the index structure in its entirely.

B. Index nodes are not retrieved in the index order, and there fore the nodes are not in sequence.

C. An index fast full scan reads the index block by block.

D. An index fast full scan reads the whole index from the lowest value to the higher value.

Correct Answer: AB

A: To speed table and index block access, Oracle uses the db_file_multiblock_read_count parameter (which defaults to 8) to aid in getting full-table scan and full-index scan data blocks into the data buffer cache as fast as possible.

B: The index nodes are not retrieved in index order, the rows will not be sequenced.

Note:

there are some requirements for Oracle to invoke the fast full-index scan.

All of the columns required must be specified in the index. That is, all columns in the select and where clauses must exist in the index.

The query returns more than 10 percent of the rows within the index. This 10 percent figure depends on the degree of multi-block reads and the degree of

parallelism.

You are counting the number of rows in a table that meet a specific criterion. The fast full-index scan is almost always used for count(*) operations.

Reference: index fast full scan tips