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

QUESTION 1

You enable auto degree of parallelism (DOP) for your database instance.

Examine the following query:

```
SQL> SELECT /*+ PARELLEL (MANUAL) */ customers.cust_first_name,
Customers.cust_last_name,
        MAX (QUANTITY_SOLD), AVG (QUANTITY_SOLD)
        FROM mysales, customers
        WHERE mysales.cust_id=customers.cust_id
        GROUP BY customer.cust_first_name, customers.cust_last_name;
```

Which two are true about the execution of statement?

A. Dictionary DOP for the objects accessed by the query is used to determine the statement DOP.

B. Auto DOP is used to determine the statement DOP only if the estimated serial execution time exceeds PARALLEL_MIN_TIME_THRESHOLD.

C. Dictionary DOP is used to determine the statement DOP only if the estimated serial execution time exceeds PARALLEL_MIN_TIME_THRESHOLD.

D. The statement will be queued if insufficient parallel execution slaves are available to satisfy the statements DOP.

E. The statement will be queued if the number of busy parallel execution servers exceeds PARALLEL_SERVERS_TARGET.

F. The statements may execute serially.

Correct Answer: AF

*

Parallel (Manual): The optimizer is forced to use the parallel settings of the objects in the state- ment.

*

MANUAL - This is the default. Disables Auto DOP (not B), statement queuing (not D, Not E) and in-memory parallel execution. It reverts the behavior of parallel

execution to what it was pre- vious to Oracle Database 11g, Release 2 (11.2).

*

PARELLEL (MANUAL)

You can use the PARALLEL hint to force parallelism. It takes an optional parameter: the DOP at which the statement should run.

The following example forces the statement to use Oracle Database 11g Release 1 (11.1) behav- ior:

SELECT /*+ parallel(manual) */ ename, dname FROM emp e, dept d WHERE e.deptno=d.deptno;

*

PARALLEL_SERVERS_TARGET specifies the number of parallel server processes allowed to run parallel statements before statement queuing will be used. When the parameter PARAL- LEL_DEGREE_POLICY is set to AUTO, Oracle will queue SQL statements that require paral- lel execution, if the necessary parallel server processes are not available. Statement queuing will begin once the number of parallel server processes active on the system is equal to or greater than PARALLEL_SERVER_TARGET.

By default, PARALLEL_SERVER_TARGET is set lower than the maximum number of parallel server processes allowed on the system (PARALLEL_MAX_SERVERS) to ensure each parallel statement will get all of the parallel server resources required and to prevent overloading the sys- tem with parallel server processes.

Note that all serial (non-parallel) statements will execute immediately even if statement queuing has been activated.

QUESTION 2

A new application module is deployed on middle tier and is connecting to your database. You want to monitor the performance of the SQL statements generated from the application.

To accomplish this, identify the required steps in the correct order from the steps given below:

1.

Use DBNMS_APPLICATION_INFO to set the name of the module

2.

Use DBMS_MONITOR.SERV_MOD_ACT_STAT_ENABLE to enable statistics gathering for the module.

3.

Use DBMS_MONITOR.SERV_MOD_ACT_TRACE_ENABLE to enable tracing for the service

4.

Use the trcsess utility to consolidate the trace files generated.

5.

Use the tkprof utility to convert the trace files into formatted output.

A. 1, 2, 3, 4, 5

B. 2, 3, 1, 4, 5

C. 3, 1, 2, 4, 5

D. 1, 2, 4, 5

E. 1, 3, 4, 5

F. 2, 1, 4, 5

Correct Answer: A

Note:

*

Before tracing can be enabled, the environment must first be configured to enable gathering of statistics.

*

(gather statistics): DBMS_MONITOR.SERV_MOD_ACT_STAT_ENABLE

Enables statistic gathering for a given combination of Service Name, MODULE and ACTION

*

DBMS_MONITOR.SERV_MOD_ACT_TRACE_ENABLE

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

dbms_monitor.serv_mod_act_trace_enable(

service_name IN VARCHAR2,

module_name IN VARCHAR2 DEFAULT ANY_MODULE,

action_name IN VARCHAR2 DEFAULT ANY_ACTION,

waits IN BOOLEAN DEFAULT TRUE,

binds IN BOOLEAN DEFAULT FALSE,

instance_name IN VARCHAR2 DEFAULT NULL,

plan_stat IN VARCHAR2 DEFAULT NULL);

SELECT instance_name

FROM gv$instance;

exec dbms_monitor.serv_mod_act_trace_enable(\\'TESTSERV\\', dbms_monitor.all_modules, dbms_monitor.all_actions, TRUE, TRUE, \\'orabase\\');

exec dbms_monitor.serv_mod_act_trace_disable(\\'TESTSERV\\', dbms_monitor.all_modules, dbms_monitor.all_actions, \\'orabase\\');

*

When solving tuning problems, session traces are very useful and offer vital information. Traces are simple and straightforward for dedicated server sessions,

but for shared server sessions, many processes are involved. 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.

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.

*

Once the trace files have been consolidated (with trcsess), tkprof can be run against the consolidated trace file for reporting purposes.

QUESTION 3

Examine the exhibit.

Id	Operations	Name	Rows	Bytes	Cost	(%CPU)	Time	IQ	IN-OUT	PQ Disturb
0	SELECT STATEMENT		14	588	5	(20)	00:00:01			
1	PX COORDINATOR									
2	PS SEND QC (RANDOM)	:iq10001	14	588	5	(20)	00:00:01	Q1, 01,	P ->S	QC(RAND)
3	HASH JOIN		14	588	5	(20)	00:00:01	Q1, 01	PCWP	
4	SUFFER SORT							Q1, 01	PCWC	
5	PX RECEIVE		4	88	2	(0)	00:00:01	Q1, 01	PCWP	
6	PX SEND BROADCAST	:IQ10000	4	88	2	(0)	00:00:01		S->P	BROADCAST
7	TABLE ACCESS FULL	DEPARTMENTS	4	88	2	(0)	00:00:01			
8	PX BLOCK ITERATOR		14	280	2	(0)	00:00:01	Q1, 01	PCWC	
9	TABLE ACCESS FUL	EMPLOYEES	14	280	2	(0)	00:00:01	Q1, 01	PCWP	

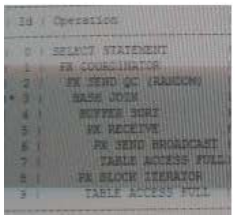

Which is true based on the information obtainable from the execution plan?

A. A full partition-wise join performed between the EMPLOYEES and DEPARTMENTS tables.

B. A full table scan on the DEPARTMENTS table performed serially by the query coordinator.

C. A full table scan on the DEPARTMENTS table is performed serially by a single parallel execution server process.

D. A partial partition-wise join performed between the EMPLOYEES and DEPARTMENTS tables.

E. A full table scan on the EMPLOYEES table is done in parallel.

Correct Answer: E

PX BLOCK ITERATOR This operation is typically the first step in a parallel pipeline. The BLOCK ITERATOR breaks up the table into chunks that are processed by

each of the parallel servers involved.

Incorrect:

B, C: The scan on the Departsments table is done in parallel.

Note:

* As per exhibit: Line 7 is run first, followed by line 6.

*

Example with same structure of execution plan:

Id	Operation	Name	Rows (Estim)	Cost	Time Active(s)	Start Active	Execs	Rows (Actual)	Read Reqs	Read Bytes	Mem (Max)	Activity (%)	Activity Detail (# samples)
0	SELECT STATEMENT				1	+3	1	1					
1	SORT AGGREGATE		1		1	+3	1	1					
2	PX COORDINATOR				1	+3	9	8					
3	PX SEND QC (RANDOM)	:TQ10001	1		6	+-2	8	8					
4	SORT AGGREGATE		1		6	+-2	8	8					
5	HASH JOIN		446K	291	6	+-2	8	452K			38M		
6	BUFFER SORT				6	+-2	8	603K			19M		
7	PX RECEIVE		75272	76	6	+-2	8	603K					
8	PX SEND BROADCAST	:TQ10000	75272	76	1	+3	1	603K					
9	INDEX FAST FULL SCAN	I_OBJ1	75272	76	1	+3	1	75322					
10	PX BLOCK ITERATOR		446K	214	6	+-2	8	452K					
11	TABLE ACCESS FULL	MYOBJ	446K	214	6	+-2	104	452K	1623	44MB			

Here\\'s how to read the plan:

1.

The first thing done is at line 9 an index fast full scan on SYS.OBJ$.I_OBJ1 index. This is done in parallel, as indicated from the "PX SEND" line above.

2.

In line 8, we\\'re doing a "PX SEND BROADCAST" operation. When joining tables in parallel, Oracle can choose to either broadcast results (rows) from one operation to apply to the other table scan, or it can choose PX SEND HASH. In this case, our CBO determined that a BROADCOAST was appropriate because the results from the OBJ$ table were much lower than the MYOBJ table

3.

Line 7, the PX RECEIVE step, is basically the consumer of the broadcasted rows in step 8

4.

Line 6 is an in-memory BUFFER SORT of the rows returned from the index scan on OBJ$

5.

Lines 11 and 10, respectively, indicate the full scan and PX BOCK ITERATOR operation for the granules involved in the 8 PQ servers

6.

In line 5, Oracle is doing a hash join on the resulting rows from the parallel scans on MYOBJ and OBJ$

7.

Line 4 is a per-PQ server sort of data from the joind PQ servers

8.

Line 3 is the consumer QC that holds the result of the each of the PQ servers

9.

Line 2 is the PX Coordinator (QC) collecting, or consuming the rows of the joined data

10.

Line 1 is the final SORT AGGREGATE line that performs the grouping function

QUESTION 4

Which two statements about In-Memory Parallel Execution are true?

A. It can be configured using the Database Resource Manager.

B. It increases the number of duplicate block images in the global buffer cache.

C. It requires setting PARALLEL_DEGREE_POLICY to LIMITED.

D. Objects selected for In-Memory Parallel Execution have blocks mapped to specific RAC instances.

E. It requires setting PARALLEL_DEGREE_POLICY to AUTO

F. Objects selected for In-Memory Parallel Execution must be partitioned tables or indexes.

Correct Answer: DE

D, E: In-Memory Parallel Execution

When the parameter PARALLEL_DEGREE_POLICY is set to AUTO, Oracle Database decides if an object that is accessed using parallel execution would benefit from being cached in the SGA (also called the buffer cache). The decision to cache an object is based on a well-defined set of heuristics including the size of the object and frequency on which it is accessed. In an Oracle RAC environment, Oracle Database maps pieces of the object into each of the buffer caches on the active instances. By creating this mapping, Oracle Database automatically knows which buffer cache to access to find different parts or pieces of the object. Using this information, Oracle Database prevents multiple instances from reading the same information from disk over and over again, thus maximizing the amount of memory that can cache objects. If the size of the object is larger than the size of the buffer cache (single instance) or the size of the buffer cache multiplied by the number of active instances in an Oracle RAC cluster, then the object is read using direct-path reads.

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

AUTO Enables automatic degree of parallelism, statement queuing, and in-memory parallel execution.

Incorrect:

C: LIMITED Enables automatic degree of parallelism for some statements but statement queuing and in- memory Parallel Execution are disabled. Automatic degree of parallelism is only applied to those statements that access tables or indexes decorated explicitly with the PARALLEL clause. Tables and indexes that have a degree of parallelism specified will use that degree of parallelism.

Reference: Oracle Database VLDB and Partitioning Guide 11g, How Parallel Execution Works

QUESTION 5

Examine the query and its execution plan: Which two statements are true regarding the execution plan?

```
SQL > SELECT cust_last_name, sum (nv12(o.customer_id, 0, 1)) "Count"
          FROM customers c, orders o
          WHERE c.credit_limit > 1000
          AND c.customer_id = o.customer_id(+)
          GROUP BY cust_last_name;
```

Id	Operation	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	CUSTOMERS	260	3900	5	(0)
*4	INDEX RANGE SCAN	ORD_CUSTOMERS_IX	105	420	0	(0)

```
Predicate Information (Identified by operation id);
---------------------------------------------------
3 – filter ("C". "CREDIT_LIMIT"> 1000)
4 – access ("C". "CUSTOMERS_ID"= "0". "CUSTOMER_ID"(+))
    Filter ("O". "CUSTOMER_ID"(+)>0)
```

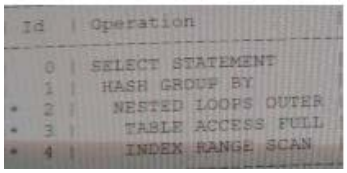

A. For every row of CUSTOMERS table, the row matching the join predicate from the ORDERS table are returned.

B. An outer join returns NULL for the ORDERS table columns along with the CUSTOMERS table rows when it does not find any corresponding rows in the ORDER table.

C. The data is aggregated from the ORDERS table before joining to CUSTOMERS.

D. The NESTED LOOP OUTER join is performed because the OPTIMZER_MODE parameter is set to ALL_ROWS.

Correct Answer: BD

B: An outer join extends the result of a simple join. An outer join returns all rows that satisfy the join condition and also returns some or all of those rows from one table for which no rows from the other satisfy the join condition.

Note:

*

All_rows attempts to optimize the query to get the very last row as fast as possible. This makes sense in a stored procedure for example where the client does

not regain control until the stored procedure completes. You don\\'t care if you have to wait to get the first row if the last

row gets back to you twice as fast. In a

client server/interactive application you may well care about that.

*

The optimizer uses nested loop joins to process an outer join in the following circumstances:

/ It is possible to drive from the outer table to inner table.

/ Data volume is low enough to make the nested loop method efficient.

*

First_rows attempts to optimize the query to get the very first row back to the client as fast as possible. This is good for an interactive client server environment

where the client runs a query and shows the user the first 10 rows or so and waits for them to page down to get more.

QUESTION 6

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 7

Examine the Exhibit.

```
CREATE TABLE dept AS SELECT * FROM departments;
ALTER TABLE dept PARALLEL 2;

CREATE TABLE emp_range_did PARTITION BY RANGE (department_id)
            (PARTITION emp_p1 VALUES LESS THAN (150),
            PARTITION emp_p5 VALUES LESS THAN (MAXVALUE) )
    AS SELECT * FROM employees;

ALTER TABLE emp_range_did PARALLEL 2;

EXPLAIN PLAN FRO
SELECT /*PQ_DISTRIBUTE (d NONE PARTITION) ORDERED */ e.last_name, d.department_name
            FROM emp_range_did e, dept d
            WHERE e.department_id = d.department_id;

-------------------------------------------------------------------------------------------------

|   Id |    Operations                          Name              Rows    Bytes   Cost    Pstart    Pstop
TQ   |IN-OUT | |PQ  DISTRIB|
-------------------------------------------------------------------------------------------

|  0|       SELECT STATEMENT                                     284     16188   6
     |       |
|  1|       PX COORDINATOR
     |       |
|  2|       PX SEND QC (RANDOM) :          TQ10001               284     16188   6
Q1, 01       |P->S| QC (RAND)|
|*  3|       HASH JOIN                                           284     16188   6
Q1, 01       |PCWP|
|  4|       PX PARTITION RANGE ALL                               284     7668    2       1         2
Q1, 01       |PCWC|
|  5|       TABLE ACCESS FULL              EMP_RANGE_DID         284     7668    2       1         2
Q1, 0        |PCWP|
|  6|       BUFFER SORT
Q1, 01       |PCWC|
|  7|       PX RECEIVE                                          21      630     2
Q1, 01       |PCWP|
|  8|       PX SEND PARTITION (KEY) :      TQ10000               21      630     2
             |S->P|PART (KEY)|
|  9|       TABLE ACCESS FULL              DEPT                  21      630     2
             |              |              |
-------------------------------------------------------------------------------------------
```

Which two options are true about the execution plan and the set of statements?

A. The query uses a partial partition-wise join.

B. The degree of parallelism is limited to the number of partitions in the EMP_RANGE_DID table.

C. The DEPT table id dynamically distributed based on the partition keys of the EMP_RANGE_DID table.

D. The server process serially scans the entire DEPT table for each range partition on the EMP_RANGE_DID table.

E. The query uses a full partition-wise join.

Correct Answer: AD

QUESTION 8

Which three are tasks performed in the hard parse stage of a SQL statement executions?

A. Semantics of the SQL statement are checked.

B. The library cache is checked to find whether an existing statement has the same hash value.

C. The syntax of the SQL statement is checked.

D. Information about location, size, and data type is defined, which is required to store fetched values in variables.

E. Locks are acquired on the required objects.

Correct Answer: BDE

Parse operations fall into the following categories, depending on the type of statement submitted and the result of the hash check: A) Hard parse

If Oracle Database cannot reuse existing code, then it must build a new executable version of the application code. This operation is known as a hard parse, or a

library cache miss. The database always perform a hard parse of DDL.

During the hard parse, the database accesses the library cache and data dictionary cache numerous times to check the data dictionary. When the database

accesses these areas, it uses a serialization device called a latch on required objects so that their definition does not change (see "Latches"). Latch contention

increases statement execution time and decreases concurrency.

B) Soft parse

A soft parse is any parse that is not a hard parse. If the submitted statement is the same as a reusable SQL statement in the shared pool, then Oracle Database

reuses the existing code. This reuse of code is also called a library cache hit.

Soft parses can vary in the amount of work they perform. For example, configuring the session cursor cache can sometimes reduce the amount of latching in the

soft parses, making them "softer."

In general, a soft parse is preferable to a hard parse because the database skips the optimization and row source generation steps, proceeding straight to

execution.

Incorrect: A, C: During the parse call, the database performs the following checks: Syntax Check Semantic Check Shared Pool Check The hard parse is within Shared Pool check. Reference: Oracle Database Concepts 11g, SQL Parsing

QUESTION 9

Examine Exhibit1 to view the query and its AUTOTRACE output.

```
SQL> SET AUTOTRACE TRACEONLY
SQL> SELECT prod_category, AVG(amount_sold)
          FROM sales s, products P
WHERE P.prod_id = S.prod_id
GROUP BY prod_category;

Execution Plan
```

Plan hash volume: 1197568639

id	Operation	Name	Rows	Bytes	Cost (%CPU)	Time	Pstart
0	SELECT STATEMENT		5	255	639 (11)	00:00:07	
1	HASH GROUP BY		5	255	539 (11)	00:00:07	
2	HASH JOIN		72	3672	538 (11)	00:00:07	
3	VIEW	VW_GBC_5	72	2160	535 (11)	00:00:07	
4	HASH GROUP BY		72	648	535 (11)	00:00:07	
5	PARTITION TANGE ALL		918K	8075K	494 (3)	00:00:06	1
6	TABLE ACCESS FULL	SALES	918K	8075K	684 (3)	00:00:06	1
7	VIEW	INDEX$_JOIN_002	72	1512	3 (34)	00:00:01	
8	HASH JOIN						
9	INDEX FAST FULL SCAN	PRODUCT_PK	72	1512	1 (0)	00:00:01	
10	INDEX FAST FULL SCAN	PRODUCT_PROD_CAT_IX	72	1512	1 (0)	00:00:01	

```
Predicate information (identified by operation id)"
---------------------------------------------------
2- access ("P". "PROD_ID" = "ITEM_1")
8- access (ROEID=ROWID)

Statistics
---------------------------------------------------
          0          recursive calls
          0          db block gets
          1726       consistent gets
          0          physical reads
          0          Redo size
          778        bytes sent via SQL "Net to client
          434        bytes received via SQL "Net from client
          2          SQL "Net roundtrips to/from client
          0          sorts (memory)
```

Which two statements are true about tracing?

A. The displayed plan will be stored in PLAN_TABLE.

B. Subsequent execution of this statement will use the displayed plan that is stored in v$SQL.

C. The displayed plan may not necessarily be used by the optimizer.

D. The query will not fetch any rows; it will display only the execution plan and statistics.

E. The execution plan generated can be viewed from v$SQLAREA.

Correct Answer: AD

The PLAN_TABLE is automatically created as a public synonym to a global temporary table. This temporary table holds the output of EXPLAIN PLAN statements for all users. PLAN_TABLE is the default sample output table into which the EXPLAIN PLAN statement inserts rows describing execution plans

QUESTION 10

You are administering a database that supports a DSS workload, where in an application a set of queries use the query

rewrite on materialized views. You notice that these queries are performing poorly.

Which two actions can you make to improve the performance of theses queries?

A. Use DBMS_MVIEW.EXPLAIN_REWRITE to analyze whether the queries are rewritten.

B. USE DBMS_ADVISOR.QUICK_TUNE to analyze the query rewrite usage of materialized views for the entire workload.

C. Create an STS for all the queries and use SQL performance Analyzer to generate recommendations for determining the regressed SQL statements.

D. Create an STS for all the queries in the application and use the SQL Tuning Advisor to generate recommendations.

E. Create an STS for all the queries in the application and use the SQL Access Advisor to generate a recommendation for optimizing materialized views for maximum query rewrite usage and fast refresh.

Correct Answer: DE

http://docs.oracle.com/cd/E11882_01/server.112/e41573/advisor.htm#PFGRF94911

QUESTION 11

Which statement is true about an automatic SQL task?

A. It will attempt to tune the currently running SQL statements that are highly resource intensive.

B. It will automatically implement new SQL profiles for the statements that have existing SQL profiles.

C. It will attempt to tune all-long-running queries that have existing SQL profiles.

D. It will automatically implement SQL profiles if a three-fold benefit can be achieved and automatic profile implementation is enabled.

E. It will tune all the top SQL statements from AWR irrespective of the time it takes to complete the task in a maintenance window.

Correct Answer: D

Optionally, implements the SQL profiles provided they meet the criteria of threefold performance improvement

The database considers other factors when deciding whether to implement the SQL profile. For example, the database does not implement a profile when the objects referenced in the statement have stale optimizer statistics. SQL profiles that have been implemented automatically show type is AUTO in the DBA_SQL_PROFILES view. If the database uses SQL plan management, and if a SQL plan baseline exists for the SQL statement, then the database adds a new plan baseline when creating the SQL profile. As a result, the optimizer uses the new plan immediately after profile creation.

Incorrect:

E: Oracle Database automatically runs SQL Tuning Advisor on selected high-load SQL statements from the Automatic Workload Repository (AWR) that qualify as

tuning candidates. This task, called Automatic SQL Tuning, runs in the default maintenance windows on a nightly basis. By default, automatic SQL tuning runs for

at most one hour.

Note:

After automatic SQL tuning begins, the database performs the following steps:

1. Identifies SQL candidates in the AWR for tuning

Oracle Database analyzes statistics in AWR and generates a list of potential SQL statements that are eligible for tuning. These statements include repeating high-load statements that have a significant impact on the database.

The database tunes only SQL statements that have an execution plan with a high potential for improvement. The database ignores recursive SQL and statements that have been tuned recently (in the last month), parallel queries, DML, DDL, and SQL statements with performance problems caused by concurrency issues.

The database orders the SQL statements that are selected as candidates based on their performance impact. The database calculates the impact by summing the CPU time and the I/O times in AWR for the selected statement in the past week.

2.

Tunes each SQL statement individually by calling SQL Tuning Advisor

During the tuning process, the database considers and reports all recommendation types, but it can implement only SQL profiles automatically.

3.

Tests SQL profiles by executing the SQL statement

4.

Optionally, implements the SQL profiles provided they meet the criteria of threefold performance improvement. The database considers other factors when deciding whether to implement the SQL profile. For example, the database does not implement a profile when the objects referenced in the statement have stale optimizer statistics. SQL profiles that have been implemented automatically show type is AUTO in the

DBA_SQL_PROFILES view.If the database uses SQL plan management, and if a SQL plan baseline exists for the SQL statement, then the database adds a new plan baseline when creating the SQL profile. As a result, the optimizer uses the new plan immediately after profile creation.

Reference: Oracle Database Performance Tuning Guide, Automatic SQL Tuning

QUESTION 12

Which statement is true about an SQL plan baselines that are fixed?

A. New plans are added automatically by the optimizer to the baseline and are automatically evolved.

B. New, better plans are added automatically as a fixed plan baseline.

C. New plan can be manually loaded to the baseline from the cursor cache or a SQL tuning set.

D. New plans can be added as fixed plans to the baseline by using the SQL Tuning Advisor to generate a SQL profile and by accepting the SQL profile.

Correct Answer: D

When a SQL statement with a fixed SQL plan baseline is tuned using the SQL Tuning Advisor, a SQL profile recommendation has special meaning. When the

SQL profile is accepted, the tuned plan is added to the fixed SQL plan baseline as a non-fixed plan. However, as described above, the optimizer will not use the

tuned plan as long as a reproducible fixed plan is present. Therefore, the benefit of SQL tuning may not be realized. To enable the use of the tuned plan, manually

alter the tuned plan to a fixed plan by setting its FIXED attribute to YES.

Note:

It is also possible to influence the optimizer\\'s choice of plan when it is selecting a plan from a SQL plan baseline. SQL plan baselines can be marked as fixed.

Fixed SQL plan baselines indicate to the optimizer that they are preferred. If the optimizer is costing SQL plan baselines and one of the plans is fixed, the optimizer

will only cost the fixed plan and go with that if it is reproducible.

If the fixed plan(s) are not reproducible the optimizer will go back and cost the remaining SQL plan baselines and select the one with the lowest cost. Note that

costing a plan is nowhere near as expensive as a hard parse. The optimizer is not looking at all possible access methods but at one specific access path.

Reference: Oracle Database Performance Tuning Guide 11g, Using Fixed SQL Plan Baselines Reference: SQL Plan Management in Oracle Database 11g

QUESTION 13

You execute the following query for the first time:

```
SQL > SELECT employees_id_name, salary
          FROM employees
          WHERE employees
          WHERE salary > & sal;
```

Examine the SQL statement processing steps:

1.

The value of the variable SAL is obtained to run the query.

2.

The syntax of the query is checked

3.

A parse tree for the query is generated

4.

Semantics for the query are checked

5.

The required rows are fetched

6.

The SQL is executed to produce the required result.

Which is the correct order of execution of the above query?

A. 1, 2 3, 4, 5, 6

B. 1, 4, 2, 3, 6, 5

C. 2, 4, 1, 3, 6, 5

D. 2, 3, 1, 4, 6, 5

Correct Answer: C

Step 1: Syntax check (2)

Step 2: Semantic check (4)

Step 3: Shared pool check

Step 4: Optimization

Step 5: Row Source Optimization (generation of query plan) (3)

Etc.

Note:

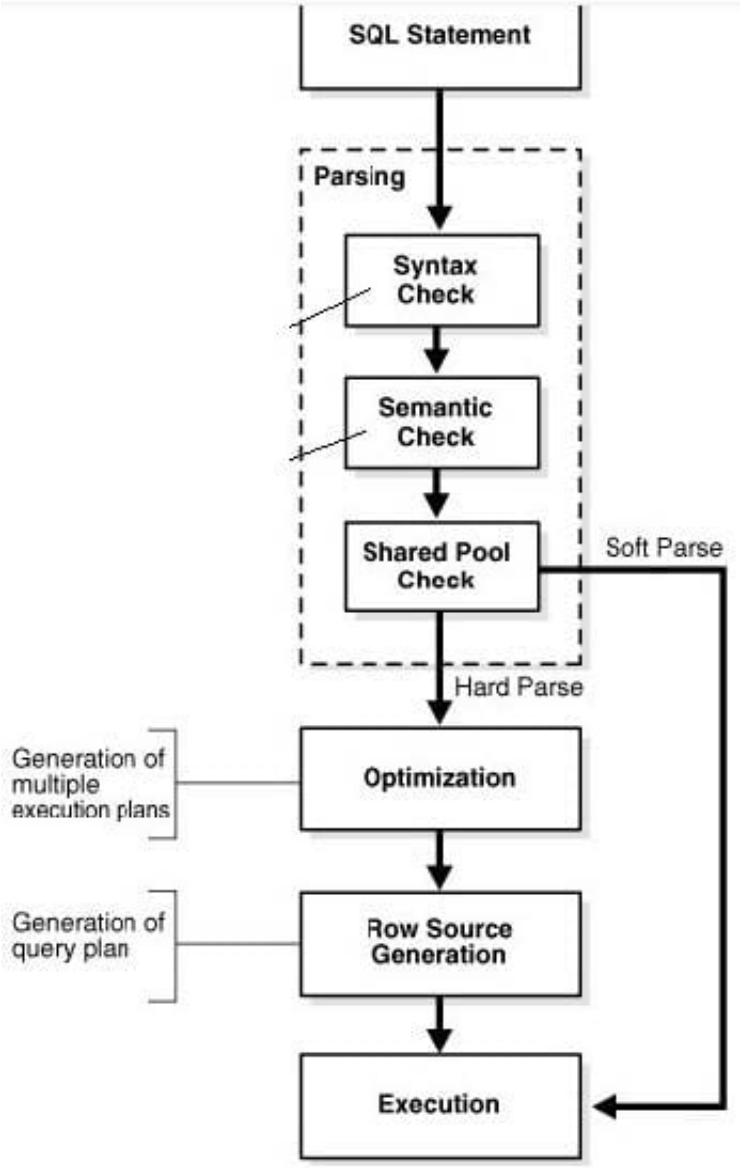

Incorrect:

Not A: First execute then fetch rows.

Not B: Check of syntax is before check of semantics.

Not D: Parse tree is after semantics.

QUESTION 14

Examine the exhibit.

Id	Operation	Name	Pstsrt	Pstop	IN-OUT	PQ	Distrib
0	SELECT STATEMENT						
1	PX COORDINATOR						
2	PX SEND QC (RANDOM)	:TQ10001			P->S	QC	(RAND)
*3	FILTER				PCWC		
4	HASH GROUP BY				PCWP		
5	PX RECEIVE				PCWP		
6	PX SEND HASH	:TQ10000			P->P	HASH	
7	HASH GRIYP BY		1	16	PCWC		
8	PX PARTITION HASH ALL		1	16	PCWP		
*9	HASH JOIN				PCWP		
10	TABLE ACCESS FULL	CUSTOMERS	1	16	PCWP		
11	PX PARTITION RANGE ITERATOR		8	9	PCWC		
*12	TABLE ACCESS FULL	SALES	113	144	PCWP		

Predicate information (identified by operation id):

3 – filter (COUNT (SYS_OP_CSR(SYS_OP_MSR(COUNT(*)), 0)>100)
9 – access ("S". "CUST_ID"= "C". "CUST_ID")
12 – filter ("S". "TIME_ID"<= TO_DATE ('1999-10-01 00:00:00', 'syyyy-mm-dd hh2:mi:ss') AND "S", "TIME_ID"> =TO_DATE('1999-07-01 00:00:00', 'syyyy-mm-dd hh24:mi:ss')

Which two are true concerning the execution plan?

A. No partition-wise join is used

B. A full partition-wise join is used

C. A partial partition-wise join is used

D. The SALES table is composite partitioned

Correct Answer: BD

* The following example shows the execution plan for the full partition-wise join with the sales table range partitioned by time_id, and subpartitioned by hash on

cust_id. ---------- | Id | Operation | Name | Pstart| Pstop |IN-OUT| PQ Distrib |

---------- | 0 | SELECT STATEMENT | | | | | |

| 1 | PX COORDINATOR | | | | | |

| 2 | PX SEND QC (RANDOM) | :TQ10001 | | | P->S | QC (RAND) |

|* 3 | FILTER | | | | PCWC | |

| 4 | HASH GROUP BY | | | | PCWP | |

| 5 | PX RECEIVE | | | | PCWP | |

| 6 | PX SEND HASH | :TQ10000 | | | P->P | HASH |

| 7 | HASH GROUP BY | | | | PCWP | |

| 8 | PX PARTITION HASH ALL | | 1 | 16 | PCWC | |

|* 9 | HASH JOIN | | | | PCWP | |

| 10 | TABLE ACCESS FULL | CUSTOMERS | 1 | 16 | PCWP | |

| 11 | PX PARTITION RANGE ITERATOR| | 8 | 9 | PCWC | |

|* 12 | TABLE ACCESS FULL | SALES | 113 | 144 | PCWP | |

Predicate Information (identified by operation id):

3 - filter(COUNT(SYS_OP_CSR(SYS_OP_MSR(COUNT(*)),0))>100)

9 - access("S"."CUST_ID"="C"."CUST_ID")

12 - filter("S"."TIME_ID"=TO_DATE(\\' 1999-07-01

00:00:00\\', \\'syyyy-mm-dd hh24:mi:ss\\'))

* Full partition-wise joins can occur if two tables that are co-partitioned on the same key are joined in a query. The tables can be co-partitioned at the partition level, or at the subpartition level, or at a combination of partition and subpartition levels. Reference partitioning is an easy way to guarantee co-partitioning. Full partition-wise joins can be executed in serial and in parallel.

Reference: Oracle Database VLDB and Partitioning Guide, Full Partition-Wise Joins: Composite

-Single-Level

QUESTION 15

Which three statements are true about the usage of optimizer hints?

A. Whenever a query uses table aliases, the hints in the query must use the aliases.

B. The OPTIMIZER_FEATURES_ENABLE parameter must be set to a version supports the hints used.

C. The optimizer uses the execution plan with lower cost even if a hint is specified.

D. A schema name for the table must be used in the hint if the table us qualified in the FROM clause.

E. Hints can be used to override the optimization approach specified with the OPTIMIZER_MODE parameter.

F. A statement block can have only one hint, and that hint must be immediately after SELECT, UPDATE, INSERT, MERGE, or DELETE keyword.

Correct Answer: ABE

*

You must specify the table to be accessed exactly as it appears in the statement. If the statement uses an alias for the table, then use the alias rather than the table name in the hint.

*

OPTIMIZER_FEATURES_ENABLE acts as an umbrella parameter for enabling a series of optimizer features based on an Oracle release number.

For example, if you upgrade your database from release 10.1 to release 11.1, but you want to keep the release 10.1 optimizer behavior, you can do so by setting this parameter to 10.1.0. At a later time, you can try the enhancements introduced in releases up to and including release 11.1 by setting the parameter to

11.1.0.6.

* If a SQL statement has a hint specifying an optimization approach and goal, then the optimizer uses the specified approach regardless of the presence or absence of statistics, the value of the OPTIMIZER_MODE initialization parameter, and the OPTIMIZER_MODE parameter of the ALTER SESSION statement.