1Z0-117Q&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.pass2lead.com/1z0-117.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by home Official Exam Center

QUESTION 1

Which three are benefits of In-Memory Parallel Execution?

A. Reduction in the duplication of block images across multiple buffer caches

B. Reduction in CPU utilization

C. Reduction in the number of blocks accessed

D. Reduction in physical I/O for parallel queries

E. Ability to exploit parallel execution servers on remote instance

Correct Answer: ACD

Note: 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.

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

QUESTION 2

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 3

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 4

Examine the initialization parameters for a database an OLTP overload.

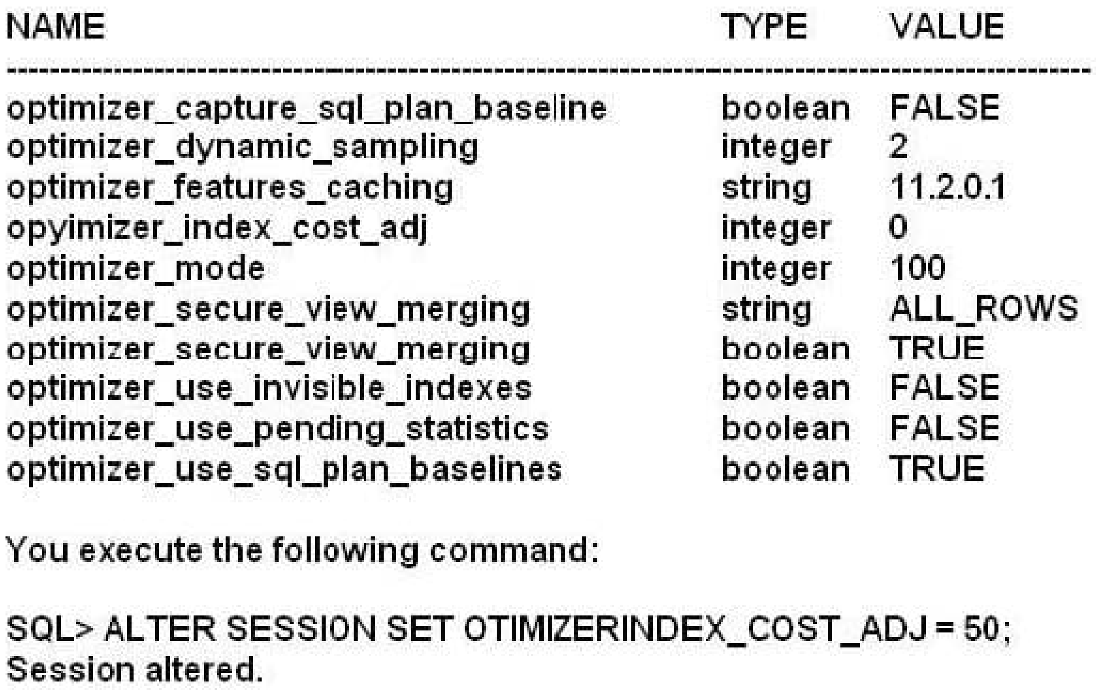

```
NAME                                    TYPE      VALUE
------------------------------------------------------------------
optimizer_capture_sql_plan_baseline     boolean   FALSE
optimizer_dynamic_sampling              integer   2
optimizer_features_caching              string    11.2.0.1
opyimizer_index_cost_adj                integer   0
optimizer_mode                          integer   100
optimizer_secure_view_merging           string    ALL_ROWS
optimizer_secure_view_merging           boolean   TRUE
optimizer_use_invisible_indexes         boolean   FALSE
optimizer_use_pending_statistics        boolean   FALSE
optimizer_use_sql_plan_baselines        boolean   TRUE
```

You execute the following command:

```
SQL> ALTER SESSION SET OTIMIZERINDEX_COST_ADJ = 50;
Session altered.
```

What is the effect of changing the value of the parameter?

A. It influences the optimizer to always use the value of the parameter?

B. It influences the optimizer to use indexes instead of full table scans as the estimated cost of the using index is reduced.

C. It influences the optimizer to use full table scans instead of index scans as the estimated cost of full table scan is reduced.

D. It influenced the optimizer to use bitmap indexes as the estimated cost conversion from bimap is rowed us reduced.

Correct Answer: B

OPTIMIZER_INDEX_COST_ADJ

OPTIMIZER_INDEX_COST_ADJ lets you tune optimizer behavior for access path selection to be more or less index friendly--that is, to make the optimizer more or less prone to selecting an index access path over a full table scan.

The default for this parameter is 100 percent, at which the optimizer evaluates index access paths at the regular cost. Any other value makes the optimizer evaluate the access path at that percentage of the regular cost. For example, a setting of 50 makes the index access path look half as expensive as normal.

QUESTION 5

You are administering a database that supports an OLTP workload in which one of the applications inserts rows in a table until 12 noon every, after which multiple years perform frequent queries on the table. You want the statistics to be more representative of the table population.

What must be done to ensure that an optimizer uses the latest statistics of the table?

A. Set the STALE_PERCENT preference to 0.

B. Set the OPTIMIZER_MODE parameter to ALL_ROWS.

C. Set the OPTIMIZER_DYNAMIC_SAMPLING parameter to 0.

D. Use the FIRST_ROWS_n hint in the queries.

E. Unlock and gather statistics for the table after inserts are done and lock them again.

Correct Answer: E

*

For tables that are substantially modified in batch operations, such as with bulk loads, gather statistics on these tables as part of the batch operation. Call the DBMS_STATS procedure as soon as the load operation completes.

* Statistics for a table or schema can be locked. After statistics are locked, you can make no modifications to the statistics until the statistics have been unlocked. Locking procedures are useful in a static environment in which you want to guarantee that the statistics never change.

The DBMS_STATS package provides two procedures for locking (LOCK_SCHEMA_STATS and LOCK_TABLE_STATS) and two procedures for unlocking statistics (UNLOCK_SCHEMA_STATS andUNLOCK_TABLE_STATS).

Incorrect:

A: STALE_PERCENT cannot be set to 0.

* With the DBMS_STATS package you can view and modify optimizer statistics gathered for database objects.

STALE_PERCENT - This value determines the percentage of rows in a table that have to change before the statistics on that table are deemed stale and should be regathered. The default value is 10%.

B: Optimizer_mode applies to the database, not to a specific table.

*

Possible values for optimizer_mode = choose/ all_rows/ first_rows/ first_rows[n]

Important facts about ALL_ROWS

1.

ALL_ROWS considers both index scan and full scan and based on their contribution to the overall query, it uses them. If Selectivity of a column is low, optimizer may use index to fetch the data (for example `where employee_code=7712\\'), but if selectivity of column is quite high (\\'where deptno=10\\'), optimizer may consider doing Full table scan. With ALL_ROWS, optimizer has more freedom to its job at its best.

2.

Good for OLAP system, where work happens in batches/procedures. (While some of the report may still use FIRST_ROWS depending upon the anxiety level of report reviewers)

3.

Likes hash joins over nested loop for larger data sets.

C: Optimizer dynamic sampling refers to the ability of the SQL optimizer to take a sample of rows from a table to calculate missing statistics. Dynamic sampling can be controlled with the OPTIMIZER_DYNAMIC_SAMPLING parameter or the DYNAMIC_SAMPLING hint.

level 0 - do not use dynamic sampling

D: First_row_n cannot be used in this way.

QUESTION 6

You plan to bulk load data using INSERT /*+PARALLEL*/ INTO . . . . SELECT FROM statements.

Which four types of operations can execute in parallel on tables that have no bitmapped indexes or materialized views defined on term?

A. Direct path insert of a million rows into a partitioned, index-organized table containing one million rows.

B. Direct path insert of a million rows into a partitioned, index-organized table containing 10 million rows.

C. Direct path insert of a million rows into a nonpartitioned, index-organized table containing one million rows.

D. Direct path insert of a million rows into a nonpartitioned, heap-organized table containing 10 million rows.

E. Direct path insert of a million rows into a nonpartitioned, heap-organized table containing one million rows.

Correct Answer: ABDE

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.

QUESTION 7

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

Examine the query:

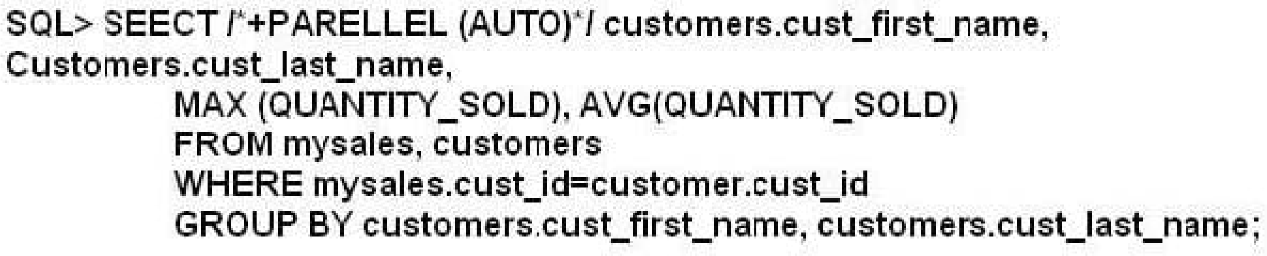

```
SQL> SEECT /*+PARELLEL (AUTO)*/ customers.cust_first_name,
Customers.cust_last_name,
            MAX (QUANTITY_SOLD), AVG(QUANTITY_SOLD)
            FROM mysales, customers
            WHERE mysales.cust_id=customer.cust_id
            GROUP BY customers.cust_first_name, customers.cust_last_name;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

Which two are true about the execution of this query?

A. Dictionary DOP will be used, if present, on the tables referred in the query.

B. DOP is calculated if the calculated DOP is 1.

C. DOP is calculated automatically.

D. Calculated DOP will always by 2 or more.

E. The statement will execute with auto DOP only when PARALLEL_DEGREE_POLICY is set to AUTO.

Correct Answer: AC

*

PARALLEL (AUTO): The database computes the degree of parallelism (C), which can be 1 or greater (not D). If the computed degree of parallelism is 1, then the statement runs serially.

*

You can use the PARALLEL hint to force parallelism. It takes an optional parameter: the DOP at which the statement should run. In addition, theNO_PARALLEL hint overrides a PARALLEL parameter in the DDL that created or altered the table.

The following example illustrates computing the DOP the statement should use: SELECT /*+ parallel(auto) */ ename, dname FROM emp e, dept d

WHERE e.deptno=d.deptno;

* When the parameter PARALLEL_DEGREE_POLICY is set to AUTO, Oracle Database auto- matically decides if a statement should execute in parallel or not and what DOP it should use. Oracle Database also determines if the statement can be executed immediately or if it is queued until more system resources are available. Finally, Oracle Database decides if the statement can take advantage of the aggregated cluster memory or not.

QUESTION 8

You are working on a database that supports an OLTP workload. You see a large number of hard parses occurring and several almost identical SQL statements in the library cache that vary only in the literal values in the WHERE clause conditions.

Which two methods can you use to reduce hard parsing?

A. Replace literals with bind variables and evolve a baseline for the statement.

B. Use the RESULT_CACHE hint in the queries.

C. Create baselines for the almost identical SQL statement by manually loading them from the cursor cache.

D. Set the CURSOR_SHARING parameter to SIMILAR.

Correct Answer: AD

A: We can reduce this Hard parsing by using bindvariables

D: SIMILAR

Causes statements that may differ in some literals, but are otherwise identical, to share a cursor, unless the literals affect either the meaning of the statement or the degree to which the plan is optimized.

Note:

A hard parse is when your SQL must be re-loaded into the shared pool. A hard parse is worse than a soft parse because of the overhead involved in shared pool RAM allocation and memory management. Once loaded, the SQL must then be completely re-checked for syntax and semantics and an executable generated.

Excessive hard parsing can occur when your shared_pool_size is too small (and reentrant SQL is paged out), or when you have non-reusable SQL statements without host variables.

See the cursor_sharing parameter for a easy way to make SQL reentrant and remember that you should always use host variables in you SQL so that they can be reentrant.

Reference: Oracle Database Reference, CURSOR_SHARING

QUESTION 9

Which two types of column filtering may benefit from partition pruning?

A. Equally operates on range-partitioned tables.

B. In-list operators on system-partitioned tables

C. Equality operators on system-partitioned tables

D. Operators on range-partitioned tables

E. Greater than operators on hash-partitioned tables

Correct Answer: AD

The query optimizer can perform pruning whenever a WHERE condition can be reduced to either one of the following two cases:

partition_column = constant

partition_column IN (constant1, constant2, ..., constantN)

In the first case, the optimizer simply evaluates the partitioning expression for the value given, determines which partition contains that value, and scans only this partition. In many cases, the equal sign can be replaced with another arithmetic comparison, including , =, and . Some queries using BETWEEN in the WHERE clause can also take advantage of partition pruning.

Note:

*

The core concept behind partition pruning is relatively simple, and can be described as "Do not scan partitions where there can be no matching values".

When the optimizer can make use of partition pruning in performing a query, execution of the query can be an order of magnitude faster than the same query against a nonpartitioned table containing the same column definitions and data.

*

Example:

Suppose that you have a partitioned table t1 defined by this statement:

CREATE TABLE t1 (

fname VARCHAR(50) NOT NULL,

lname VARCHAR(50) NOT NULL,

region_code TINYINT UNSIGNED NOT NULL,

dob DATE NOT NULL

)

PARTITION BY RANGE(region_code) (

PARTITION p0 VALUES LESS THAN (64),

PARTITION p1 VALUES LESS THAN (128),

PARTITION p2 VALUES LESS THAN (192),

PARTITION p3 VALUES LESS THAN MAXVALUE

);

Consider the case where you wish to obtain results from a query such as this one:

SELECT fname, lname, region_code, dob

FROM t1

WHERE region_code > 125 AND region_code

p0 or

It is easy to see that none of the rows which ought to be returned will be in either of the partitions p3; that is, we need to search only in partitions p1 and p2 to find

matching rows. By doing so, it is possible to expend much less time and effort in finding matching rows than would be required to scan all partitions "cutting away"

of unneeded partitions is known as pruning.

in the table. This

QUESTION 10

Examine the execution plan:

```
----------------------------------------------------------------------------------------------------------------------
Id      Operation                          Name        Pstart     Pstop      IN-OUT   PQ      Distrib
----------------------------------------------------------------------------------------------------------------------
0       SELECT STATEMENT
1       PX COORDINATOR
2       PX SEND QC (RANDOM)                :TQ10002                          P->S     QC      RAND
*3      FILTER                                                               PCWC
4       HASH GROUP BY                                                        PCWP
5       PX RECEIVE                                                           PCWP
6       PX SEND HASH                       :TQ10001                          P->P             HASH
7       HASH GROUP BY                                                        PCWP
*8      HASH JOIN                                                            PCWP
9       PART JOIN FILTER CREATE            :BF0000                           PCWP
10      PX RECEIVE                                                           PCWP
11      PX SEND PARTITION (KEY)            :TQ10000                          P->P             PART(KEY)
12      PX BLOCK ITERATOR                                                    PCWC
13      TABLE ACCESS FULL              CUSTOMERS                             PCWP
14      PX PARTITION HASH JOIN-FILTER                  :BF0000    :BF0000    PCWC
*15     TABLE ACCESS FULL                  SALES       :BF0000    :BF0000    PCWP
----------------------------------------------------------------------------------------------------------------------
```

```
Predicate Information (identified by operation id):
--------------------------------------------------------
3 – filter (COUNT(SYS_OP_CSR(SYS_OP_MSR(COUNT(0)), 0))>100)
8 – access (“S”. “CUST_ID” = “C”. CUST_ID)
15 – filter (“S”. “TIME_ID” <=TO_DATE (‘ 1999 – 10 – 01 00:00:00’, ‘syyy-mm-dd ff24:mi:ss’ )
AND
“S”. “TIME_ID”>=TO_DATE(‘1999-07-01’)
            00:00:00’, ‘syyy-mm-dd hh24:mi:ss’))
```

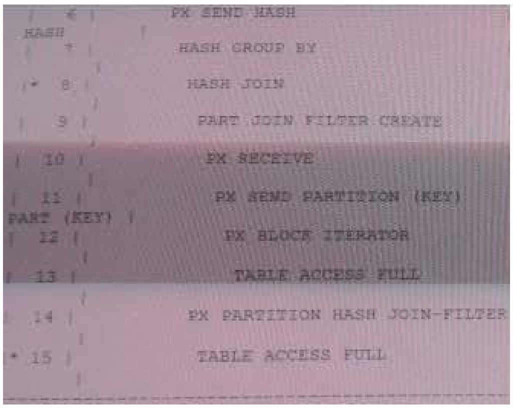

```
|   6 |        PX SEND HASH
 HASH       |
|   7 |         HASH GROUP BY
|*  8 |          HASH JOIN
|   9 |           PART JOIN FILTER CREATE
|  10 |            PX RECEIVE
|  11 |             PX SEND PARTITION (KEY)
PART (KEY) |
|  12 |              PX BLOCK ITERATOR
|  13 |               TABLE ACCESS FULL
|  14 |           PX PARTITION HASH JOIN-FILTER
|* 15 |            TABLE ACCESS FULL
------------------------------------------------------------
```

Which two are true regarding the execution plan?

A. The CUSTOMERS table is hash partitioned.

B. The SALES table is hash partitioned.

C. The CUSTOMERS table is scanned first and selected partitions from the SALES table are scanned based on the BLOOM Filter created during the scan of the CUSTOMERS table.

D. The SALES table is scanned first and selected partitions from the CUSTOMERS table are scanned based on the Bloom Filter created during the scan of the SALES table.

E. Both the CUSTOMERS and SALES tables are scanned simultaneously and rows from the CUSTOMERS table are joined to row of the SALES table.

F. The CUSTOMERS table is range partitioned.

Correct Answer: BC

B: As per line 14 and 15.

C: As per exhibit line 13 is execute before line 15.

Incorrect:

Not E: As per line 13 and 15 they are not executed simultaneously.

QUESTION 11

Which two tasks are performed during the optimization stage of a SQL statement?

A. Evaluating the expressions and conditions in the query

B. Checking the syntax and analyzing the semantics of the statement

C. Separating the clauses of the SQL statement into structures that can be processed

D. Inspecting the integrity constraints and optimizing the query based on this metadata

E. Gathering the statistics before creating the execution plan for the statement

Correct Answer: DE

Note:

Oracle SQL is parsed before execution, and a hard parse includes these steps: *

1.

Loading into shared pool - The SQL source code is loaded into RAM for parsing. (the "hard" parse step)

2.

Syntax parse - Oracle parses the syntax to check for misspelled SQL keywords.

3.

Semantic parse - Oracle verifies all table and column names from the dictionary and checks to see if you are authorized to see the data.

4.

Query Transformation - If enabled (query_rewrite=true), Oracle will transform complex SQL into simpler, equivalent forms and replace aggregations with

materialized views, as appropriate.

5.

Optimization - Oracle then creates an execution plan, based on your schema statistics (or maybe with statistics from dynamic sampling in 10g).

6.

Create executable - Oracle builds an executable file with native file calls to service the SQL query.

*

The parsing process performs two main functions:

o Syntax Check: is the statement a valid one. Does it make sense given the SQL grammar documented in the SQL Reference Manual. Does it follow all of the

rules for SQL.

o Semantic Analysis: Going beyond the syntax ? is the statement valid in light of the objects in the database (do the tables and columns referenced exist). Do you

have access to the objects ? are the proper privileges in place? Are there ambiguities in the statement ? for example if there are two tables T1 and T2 and both

have a column X, the query ?select X from T1, T2 where ?? is ambiguous, we don?t know which table to get X from. And so on.

So, you can think of parsing as basically a two step process, that of a syntax check to check the validity of the statement and that of a semantic check ? to ensure

the statement can execute properly.

Reference: Oracle hard-parse vs. soft parse

QUESTION 12

You execute the following query for the first time:

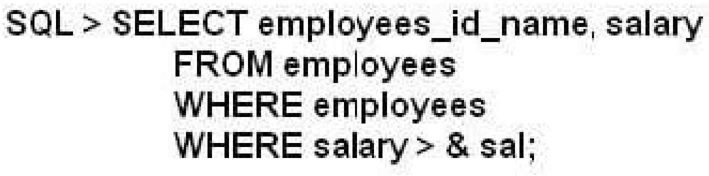

```
SQL > SELECT employees_id_name, salary
             FROM employees
             WHERE employees
             WHERE salary > & sal;
```

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:

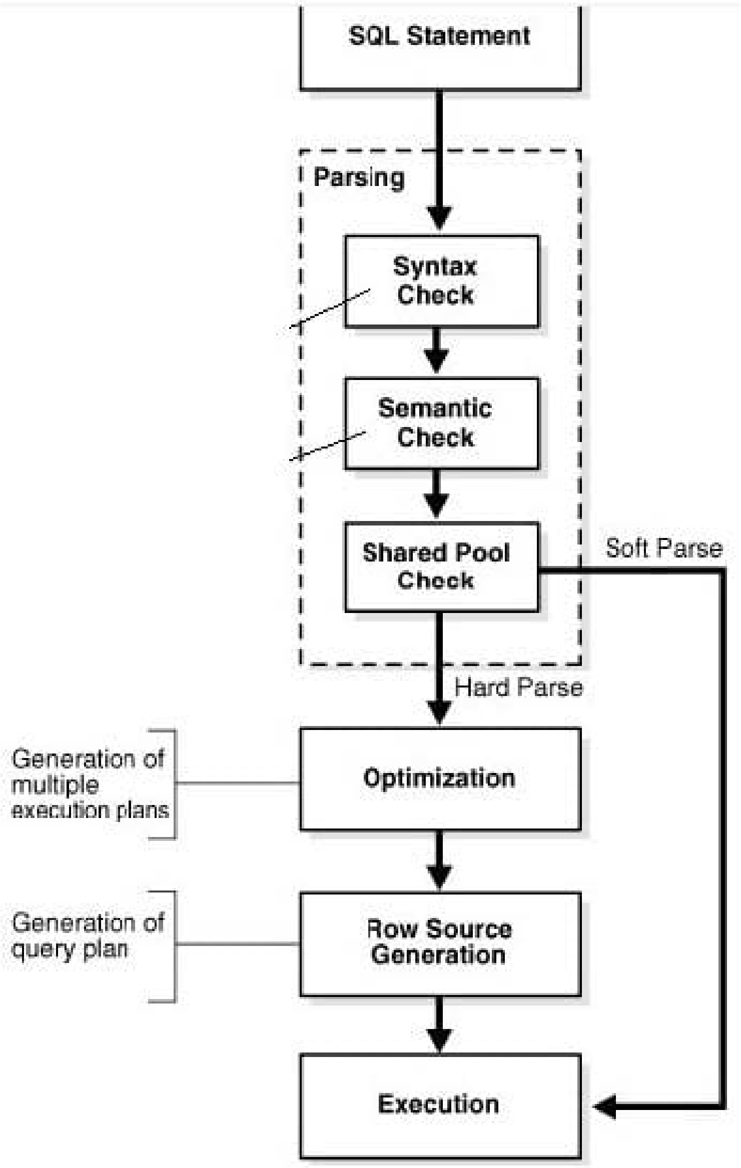

Incorrect:

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 13

Examine the parallelism parameter for your instance:

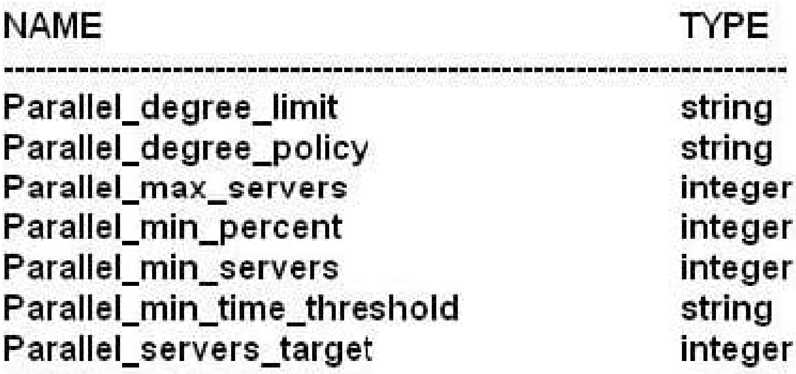

```
NAME                                  TYPE      VALUE
-----------------------------------------------------
Parallel_degree_limit                 string    CPU
Parallel_degree_policy                string    AUTO
Parallel_max_servers                  integer   128
Parallel_min_percent                  integer   0
Parallel_min_servers                  integer   0
Parallel_min_time_threshold           string    AUTO
Parallel_servers_target               integer   64
```

parallel_servers_target

Now examine the resource plan containing parallel statement directives:

Consumer Group resource plan containing parallel statement directives:

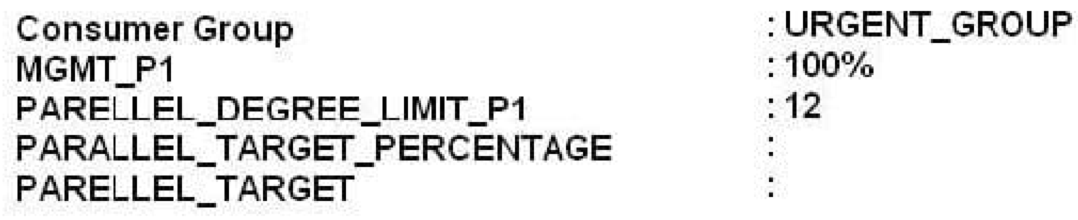

```
Consumer Group                      : URGENT_GROUP
MGMT_P1                             : 100%
PARELLEL_DEGREE_LIMIT_P1            : 12
PARALLEL_TARGET_PERCENTAGE          :
PARELLEL_TARGET                     :
```

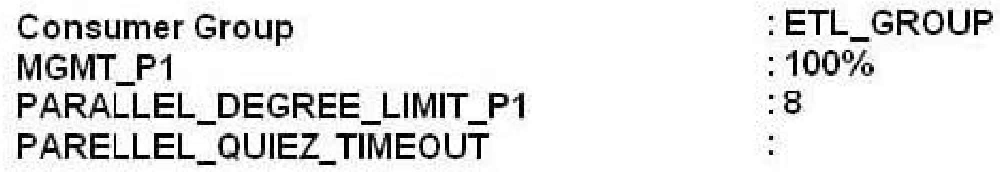

```
Consumer Group                      : ETL_GROUP
MGMT_P1                             : 100%
PARALLEL_DEGREE_LIMIT_P1            : 8
PARELLEL_QUIEZ_TIMEOUT              :
```

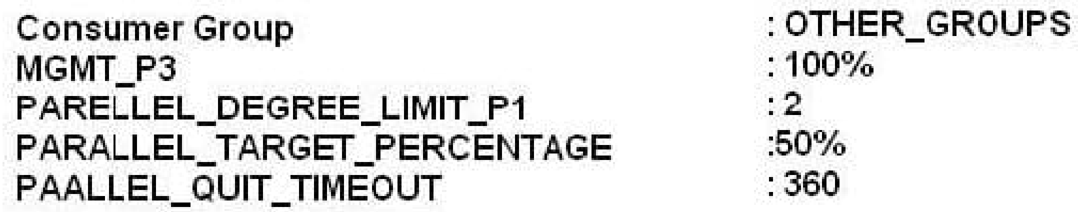

```
Consumer Group                      : OTHER_GROUPS
MGMT_P3                             : 100%
PARELLEL_DEGREE_LIMIT_P1            : 2
PARALLEL_TARGET_PERCENTAGE          :50%
PAALLEL_QUIT_TIMEOUT                : 360
```

Which two are true about parallel statement queuing when this plan is active?

A. Urgent_group sessions collectively can consume up to 64 parallel execution servers before queuing starts for this consumer group.

B. ETL_GROUP sessions can collectively consume up to 64 parallel execution servers before the queuing starts for this consumer.

C. A single OTHER_GROUPS session will execute serially once it is queued for six minutes.

D. A single ETL_GROUP session can consume up to eight parallel execution servers.

E. A single ETL_GROUP session can consume up to 32 parallel execution servers.

F. A single OTHER_GROUPS session will execute in parallel once it is queued for six minutes.

Correct Answer: AD

(http://docs.oracle.com/cd/E11882_01/server.112/e25494/dbrm.htm#ADMIN13466)

QUESTION 14

You need to migrate database from oracle Database 10g to 11g. You want the SQL workload to start the 10g plans in the 11g database instance and evolve better plans.

Examine the following steps:

1.

Capture the pre-Oracle Database 11g plans in a SQL Tuning Set (STS)

2.

Export the STS from the 10g system.

3.

Import the STS into Oracle Database 11g.

4.

Set the OPTIMIZER_FEATURES_ENABLE parameter to 10.2.0.

5.

Run SQL Performance Analyzer for the STS.

6.

Set the OPTIMIZER_FEATURES_ENABLE parameter to 11.2.0.

7.

Rerun the SQL Performance Analyzer for the STS.

8.

Set OPTIMIZER_CAPTURE_SQL_PLAN_BASELINE to TRUE.

9.

Use DBMS_SPM.EVOLVE_SQL_BASELINE function to evolve the plans.

10.

Set the OPTIMIZER_USE_SQL_PLAN_BASELINE to TRUE.

Identify the required steps in the correct order.

A. 1, 2, 3, 4, 5, 6, 7,

B. 4, 8, 10

C. 1, 2, 3, 4, 8, 10

D. 1, 2, 3, 6, 9, 5

E. 1, 2, 3, 5, 9, 10

Correct Answer: C

Step 1: (1)

Step 2: (2)

Step 3: (3)

Step 4: (4)

By setting the parameter OPTIMIZER_FEATURES_ENABLE to the 10g version used before the upgrade, you should be able to revert back to the same execution

plans you had prior to the upgrade.

Step 5: (8)

OPTIMIZER_CAPTURE_SQL_PLAN_BASELINES

In Oracle Database 11g a new feature called SQL Plan Management (SPM) has been introduced to guarantees any plan changes that do occur lead to better

performance. When OPTIMIZER_CAPTURE_SQL_PLAN_BASELINES is set to TRUE (default FALSE) Oracle will

automatically capture a SQL plan baseline for every repeatable SQL statement on the system. The execution plan found at parse time will be added to the SQL

plan baseline as an accepted plan.

Step 6: (10)

OPTIMIZER_USE_SQL_PLAN_BASELINES enables or disables the use of SQL plan baselines stored in SQL Management Base. When enabled, the optimizer

looks for a SQL plan baseline for the SQL statement being compiled. If one is found in SQL Management Base, then the optimizer will cost each of the baseline

plans and pick one with the lowest cost.

QUESTION 15

Examine the Exhibit and view the query and its execution plan.

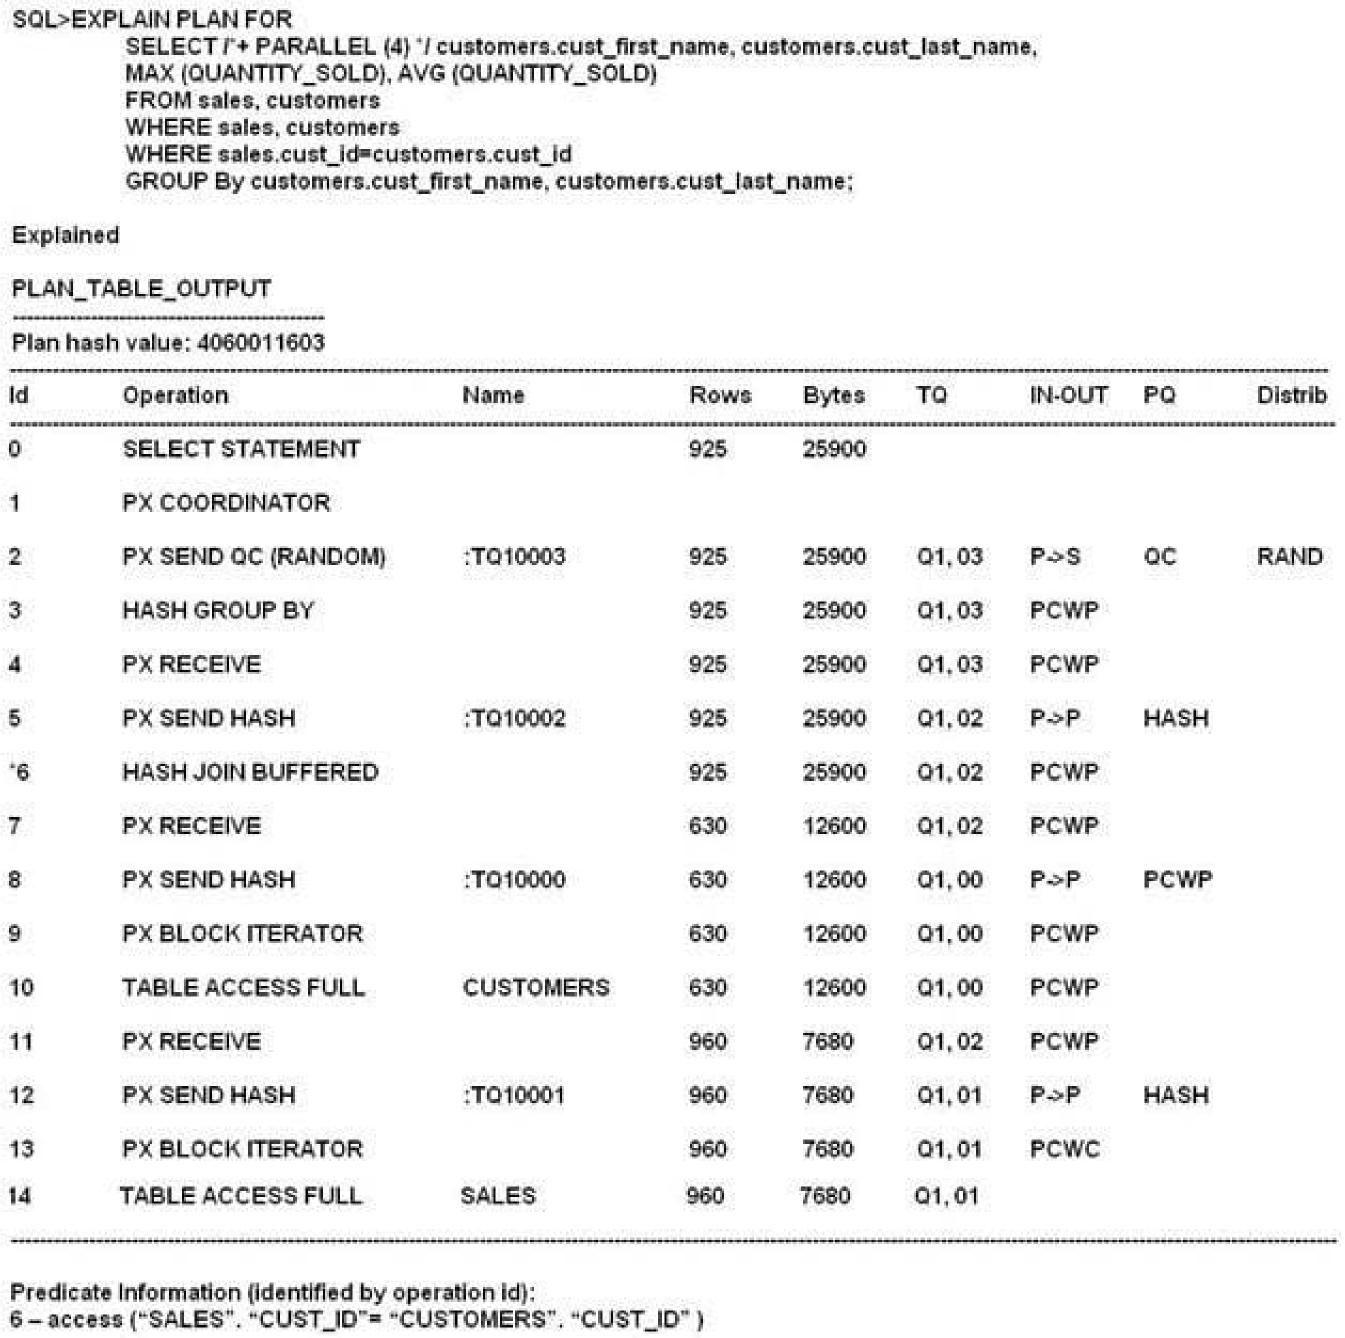

```
SQL>EXPLAIN PLAN FOR
          SELECT /*+ PARALLEL (4) */ customers.cust_first_name, customers.cust_last_name,
          MAX (QUANTITY_SOLD), AVG (QUANTITY_SOLD)
          FROM sales, customers
          WHERE sales, customers
          WHERE sales.cust_id=customers.cust_id
          GROUP By customers.cust_first_name, customers.cust_last_name;

Explained

PLAN_TABLE_OUTPUT
------------------------------------
Plan hash value: 4060011603
------------------------------------------------------------------------------------------------------------
Id    Operation               Name          Rows    Bytes    TQ        IN-OUT    PQ      Distrib
------------------------------------------------------------------------------------------------------------
0     SELECT STATEMENT                      925     25900

1     PX COORDINATOR

2     PX SEND QC (RANDOM)     :TQ10003      925     25900    Q1, 03    P->S      QC      RAND

3     HASH GROUP BY                         925     25900    Q1, 03    PCWP

4     PX RECEIVE                            925     25900    Q1, 03    PCWP

5     PX SEND HASH            :TQ10002      925     25900    Q1, 02    P->P      HASH

*6    HASH JOIN BUFFERED                    925     25900    Q1, 02    PCWP

7     PX RECEIVE                            630     12600    Q1, 02    PCWP

8     PX SEND HASH            :TQ10000      630     12600    Q1, 00    P->P      PCWP

9     PX BLOCK ITERATOR                     630     12600    Q1, 00    PCWP

10    TABLE ACCESS FULL       CUSTOMERS     630     12600    Q1, 00    PCWP

11    PX RECEIVE                            960     7680     Q1, 02    PCWP

12    PX SEND HASH            :TQ10001      960     7680     Q1, 01    P->P      HASH

13    PX BLOCK ITERATOR                     960     7680     Q1, 01    PCWC
14    TABLE ACCESS FULL       SALES         960     7680     Q1, 01
------------------------------------------------------------------------------------------------------------

Predicate Information (identified by operation id):
6 – access (“SALES”. “CUST_ID”= “CUSTOMERS”. “CUST_ID” )
```

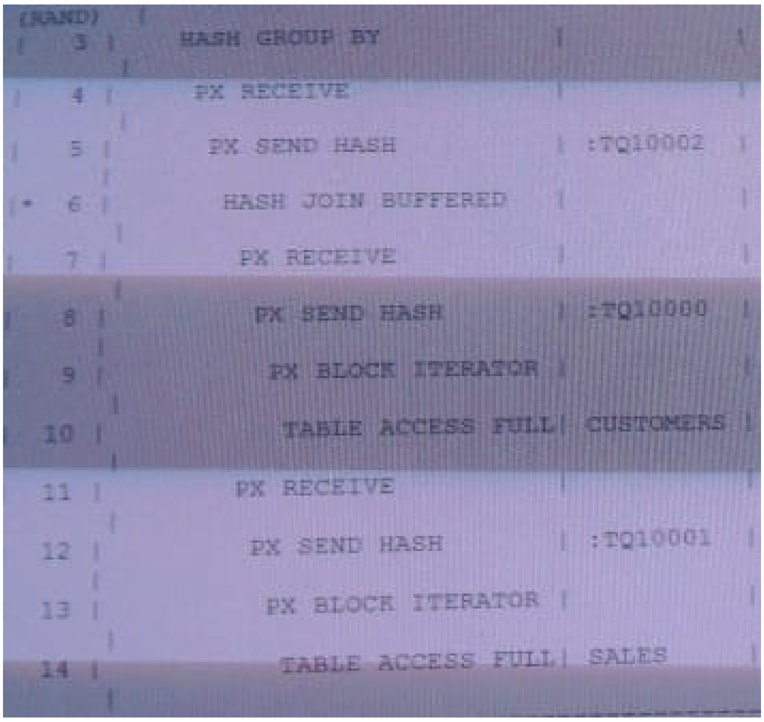

```
(RAND)  |
|   3 |       HASH GROUP BY             |           |
|   4 |        PX RECEIVE               |           |
|   5 |         PX SEND HASH            | :TQ10002  |
|*  6 |          HASH JOIN BUFFERED     |           |
|   7 |           PX RECEIVE            |           |
|   8 |            PX SEND HASH         | :TQ10000  |
|   9 |             PX BLOCK ITERATOR   |           |
|  10 |              TABLE ACCESS FULL  | CUSTOMERS |
|  11 |           PX RECEIVE            |           |
|  12 |            PX SEND HASH         | :TQ10001  |
|  13 |             PX BLOCK ITERATOR   |           |
|  14 |              TABLE ACCESS FULL  | SALES     |
```

Which statement is correct about the parallel executions plan?

A. The CUSTOMERS and SALES tables are scanned simultaneously in parallel and then joined in parallel.

B. First, the CUSTOMERS table is scanned in parallel, then the SALES table is scanned in parallel, and then they are joined serially.

C. First, the SALES table is scanned in parallel, then the CUSTOMERS table us scanned in parallel, and then they are joined in parallel.

D. The CUSTOMERS and SALES tables are scanned simultaneously in parallel and then joined serially.

E. First, the CUSTOMERS table is scanned in parallel, then the SALES table us scanned in parallel, and then they are joined in parallel.

Correct Answer: A

As per exhibit:

Line 7 and line 11 are run in parallel.

Line 8 and line 12 are run in parallel.

Line 9 and line 13 are run in parallel.

Line 10 and line 14 are run in parallel.

Line 6 is a PCWP (parallel combined with parent) and the parent is a P-> P (Parallel to parallel) operation.

1Z0-117 Practice Test 1Z0-117 Study Guide 1Z0-117 Braindumps