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.certbus.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

One of your databases supports a mixed workload.

When monitoring SQL performance, you detect many direct paths reads full table scans.

What are the two possible causes?

A. Histograms statistics not available

B. Highly selective filter on indexed columns

C. Too many sort operations performed by queries

D. Indexes not built on filter columns

E. Too many similar type of queries getting executed with cursor sharing disabled

Correct Answer: BD

Note:

* The direct path read Oracle metric occurs during Direct Path operations when the data is asynchronously read from the database files into the PGA instead of

into the SGA data buffer.

Direct reads occur under these conditions:

-

When reading from the TEMP tablespace (a sort operation)

-

When reading a parallel full-table scan (parallel query factotum (slave) processes)

-Reading a LOB segment

* The optimizer uses a full table scan in any of the following cases:

-Lack of Index

-Large Amount of Data

-Small Table

-High Degree of Parallelism

QUESTION 2

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.

QUESTION 3

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" )
```

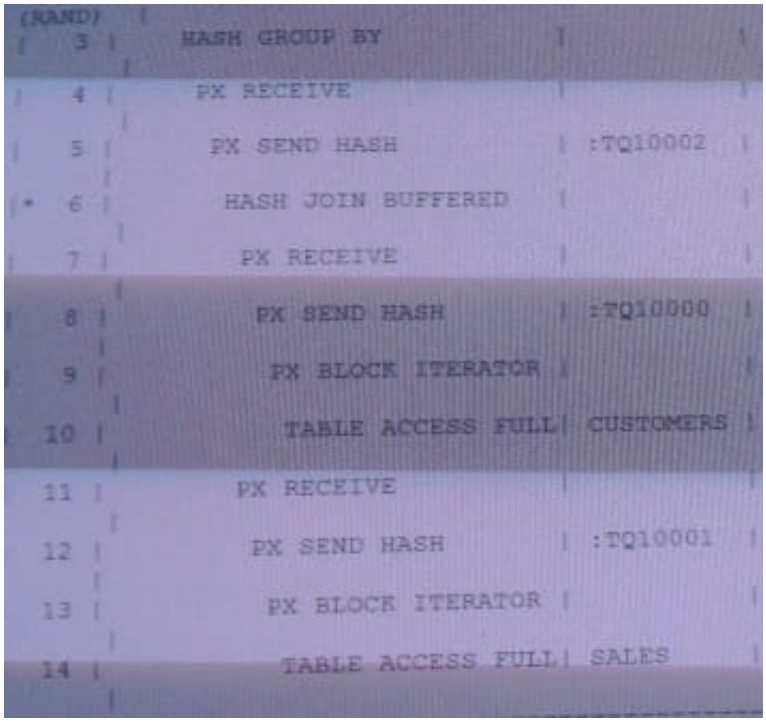

Which statement is correct about the parallel executions plan?

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.

QUESTION 4

Which four types of column filtering may benefit from partition pruning when accessing tables via partitioned indexes?

A. Equality operates on List-Partitioned Indexes

B. Not Equal operates on a Global Hash-Partitioned Indexes

C. Equality operates on System-Partitioned Tables

D. In-List operates on Range-Partitioned Indexes

E. Not Equal operates on a local Hash-Partitioned Indexes

F. Equality operates on Range-Partitioned Indexes

G. Equality operates on Hash-Partitioned Indexes

Correct Answer: ADFG

Oracle Database prunes partitions when you use range, LIKE, equality (A, F), and IN-list (D) predicates on the range or list partitioning columns, and when you use equality (G) and IN-list predicates on the hash partitioning columns.

Reference: Oracle Database VLDB and Partitioning Guide 11g, Information that can be Used for Partition Pruning

QUESTION 5

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 6

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 7

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 8

You are administering a database supporting an OLTP workload where the users perform frequent queries for fetching a new rows as possible, involving join operations on recently inserted data. In addition at night, a few DSS queries are also performed. Examine the initialization parameters for the instance:

NAME	TYPE	VALUE
optimizer_dynamic_sampling	integer	2
optimizer_index_caching	integer	90
optimizer_index_cost_adj	integer	25
optimizer_mode	string	ALL_ROWS

Which two options would you use for the optimizer?

A. Set the OPTIMIZER_MODE initialization parameter to FIRST_ROWS_n.

B. Add the hint ALL_ROWS in the DOS queries.

C. Set the OPTIMIZER_INDEX_CACHING initialization parameter to 0.

D. Add a hint INDEX_COMBINE in all DSS queries.

E. Set the OPTIMIZER_INDEX_COST_ADJ initialization parameter to 100.

Correct Answer: AE

The last appended rows are more likely to be found quickly with FIRST_ROWS_n.

E: Make it not to prioritize index instead if table scan. 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 9

Examine the exhibit to view the query and its execution plan.

```
SQL> SELECT /*+ PRDERED */ E.EMPNO, E.ENAME.D.DNAME
          FROM emp e, dept d
          WHERE e.deptno=d. deptno
          Order by e.deptno, D.DEPTNO;
```

EXECITION PLAN

Plan hash value: 3232458624

Id	Operation	Name	Rows	Bytes	Cost	(%CPU)	Time
0	SELECT STATEMENT		14	364	8	(25)	00:00:01
1	SORT ORDER BY		14	364	8	(25)	00:00:01
*2	HASH JOIN		14	364	7	(15)	00:00:01
3	TABLE ACCESS FULL	EMP	14	182	3	(0)	00:00:01
4	TABLE ACCESS FULL	DEPT	4	52	(3)	(0)	00:00:01

Predicate Information (Identified by operation id):

---

2 – access (“E” . “DEPTNO” = “D”. “DEPTNO”)

Identify the two correct interpretations that can be made from the execution plan.

A. The DEPT table is driving table and the EMP table join is the driven table.

B. Rows from the DEPT table are first hashed by the join key into memory and then joined to the EMP table on the join key.

C. The EMP table is the driving table and the DEPT table us the driven table.

D. The rows from the DEPT table are sorted first by the join key and then hashed into memory.

E. Rows from both the tables are sorted by the join key, but only rows from the DEPT table are hashed into memory.

Correct Answer: CD

If two lines are indented equally, then the top line is normally executed first. Here the line 3 and line 4 are indented equally, so line 3 (TABLE ACCESS FULL EMP) with EMP being the driving table as the ordered hint requests that the tables listed in the FROM clause of a SQL state- ment be joined in the order specified, with the first table in the FROM clause specifying the driving table.

QUESTION 10

Auto DOP is enabled for your instance. You execute the following statements:

```
SQL > ALTER TABLE employees PARALLEL 2;
SQL> ALTER TABLE departments NOPARALLEL;

SQL SELECT /*+ PARALLEL (3) */ last_name, d.department_name
          FROM employees e, departments_id=d.department_id
          WHERE e.department_id=d.department_id;
```

Which three are true about the execution of the join?

A. Dictionary DOP is used to calculate statements DOP.

B. Hinted DOP is used to calculate statement DOP.

C. The EMPLOYEES table is accessed in parallel.

D. The DEPARTMENTS table is accessed in parallel.

E. The hint operates at the level of each table accessed by the statement.

Correct Answer: BCE

C: As per ALTER TABLE employees PARALLEL 2;

Incorrect:

not D: As per ALTER TABLE departments NOPARALLEL;

QUESTION 11

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 12

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 13

You instance has these parameter settings:

```
PARALLEL_DEGREE_POLICY=AUTO
PARALLEL_SERVERS_TARGET=64
PARALLEL_MIN_MINPERCENT=25
PARALLEL_MAX_SERVERS=128
PARALLEL_MIN_SERVERS=0
PARALLEL_MIN_TIME_THRESHOLD=10
PARALLEL_DEGREE_LIMIT=8
```

Which three statements are true about these settings if no hints are used in a SQL statement?

A. A statement estimated for more than 10 seconds always has its degree of parallelism computed automatically.

B. A statement with a computed degree of parallelism greater than 8 will be queued for a maximum of 10 seconds.

C. A statement that executes for more than 10 seconds always has its degree of parallelism computed automatically.

D. A statement with a computed degree of parallelism greater than 8 will raise an error.

E. A statement with any computed degree of parallelism will be queued if the number of busy parallel execution processes exceeds 64.

F. A statement with a computed degree of parallelism of 20 will be queued if the number of available parallel execution processes is less 5.

Correct Answer: CEF

C (not A): PARALLEL_MIN_TIME_THRESHOLD specifies the minimum execution time a statement should have before the statement is considered for automatic degree of parallelism. By default, this is set to 30 seconds. Automatic degree of parallelism is only enabled if PARALLEL_DEGREE_POLICY is set to AUTO or LIMITED.

: PARALLEL_DEGREE_LIMIT integer

A numeric value for this parameter specifies the maximum degree of parallelism the optimizer can choose for a SQL statement when automatic degree of parallelism is active. Automatic degree of parallelism is only enabled if PARALLEL_DEGREE_POLICY is set to AUTO or LIMITED.

E: PARALLEL_SERVERS_TARGET specifies the number of parallel server processes allowed to run parallel statements before statement queuing will be used. When the parameter PARALLEL_DEGREE_POLICY is set to AUTO, Oracle will queue SQL statements that require parallel 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.

F: PARALELL_MIN_MINPERCENT

PARALLEL_MIN_PERCENT operates in conjunction with PARALLEL_MAX_SERVERS and PARALLEL_MIN_SERVERS. It lets you specify the minimum percentage of parallel execution processes (of the value of PARALLEL_MAX_SERVERS) required for parallel execution. Setting this parameter ensures that parallel operations will not execute sequentially unless adequate resources are available. The default value of 0 means that no minimum percentage of processes has been set.

Consider the following settings: PARALLEL_MIN_PERCENT = 50 PARALLEL_MIN_SERVERS = 5 PARALLEL_MAX_SERVERS = 10

If 8 of the 10 parallel execution processes are busy, only 2 processes are available. If you then request a query with a degree of parallelism of 8, the minimum 50% will not be met.

Note: With automatic degree of parallelism, Oracle automatically decides whether or not a statement should execute in parallel and what degree of parallelism the statement should use. The optimizer automatically determines the degree of parallelism for a statement based on the resource requirements of the statement.

However, the optimizer will limit the degree of parallelism used to ensure parallel server processes do not flood the system. This limit is enforced by PARALLEL_DEGREE_LIMIT.

Values:

CPU

IO

integer

A numeric value for this parameter specifies the maximum degree of parallelism the optimizer can choose for a SQL statement when automatic degree of parallelism is active. Automatic degree of parallelism is only enabled if PARALLEL_DEGREE_POLICY is set to AUTO or LIMITED.

Reference: PARALLEL_MIN_TIME_THRESHOLD PARALLEL_DEGREE_LIMIT PARALELL_MIN_MINPERCENT PARALELL_SERVERS_TARGET

QUESTION 14

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 15

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