



1Z0-117^{Q&As}

Oracle Database 11g Release 2: SQL Tuning Exam

Pass home 1Z0-117 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.passapply.com/1z0-117.html>

100% Passing Guarantee
100% Money Back Assurance

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

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





QUESTION 1

Examine the following query and execution plan: Which query transformation technique is used in this scenario?

Examine the following query and execution plan:

```
SQL> SELECT prod_id prod_name
      FROM producers p
      Where prod_list_prices > (SELECT MAX(unit_cost)
      FROM costs c
      WHERE p.prod_id=c.prod_id);
```

Execution Plan

Plan hash value: 2297807746

Pstart	id Pstop	Operation	Name	Rows	Bytes	Cost	(%CPU)	Time
	0	SELECT STATEMENT		4	140	82	(8)	00:00:01
	1	HASH JOIN		4	140	82	(8)	00:00:01
	2	VIEW	VW_SQL_1	72	1872	78	(7)	00:00:01
	3	HASH GROUP BY		72	648	78	(7)	00:00:01
	4	PARTITION RANGE ALL		82112	721K	74	(2)	00:00:01
1	28							
	5	TABLE ACCESS FULL	COSTS	82112	721K	74	(2)	00:00:01
1		28						
	6	TABLE ACCESS FULL	PRODUCTS	72	648	3	(0)	00:00:01

Predicate Information (Identified by operation id);

1- access ("P", "PROD_ID"=by operation id):
Filter ("PROD_LIST_PRICE"> "MAX(UNIT_COST)")

- A. Join predicate push-down
- B. Subquery factoring
- C. Subquery unnesting
- D. Join conversion

Correct Answer: A

*

Normally, a view cannot be joined with an index-based nested loop (i.e., index access) join, since a view, in contrast with a base table, does not have an index defined on it. A view can only be joined with other tables using three methods: hash, nested loop, and sort-merge joins.



*

The following shows the types of views on which join predicate pushdown is currently supported. UNION ALL/UNION view Outer-joined view Anti-joined view Semi-joined view DISTINCT view GROUP-BY view

QUESTION 2

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 3

You are administering a database that supports an OLTP application. To set statistics preferences, you issued the following command:

```
SQL > DBMS_STATS.SET_GLOBAL_PREFS ('ESTIMATE_PERCENT', '9');
```

What will be the effect of executing this procedure?

- A. It will influence the gathering of statistics for a table based on the value specified for `ESTIMATE_PERCENT` provided on table preferences for the same table exist.
- B. It will influence dynamic sampling for a query to estimate the statistics based on `ESTIMATE_PERCENT`.
- C. The automatic statistics gathering job running in the maintenance window will use global preferences unless table preferences for the same table exist.



D. New objects created will use global preference even if table preferences are specified.

Correct Answer: C

Note:

*

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

*

The SET_GLOBAL_PREFS procedure is used to set the global statistics preferences.

*

ESTIMATE_PERCENT - The value determines the percentage of rows to estimate. The valid range is [0.000001,100]. Use the constant DBMS_STATS.AUTO_SAMPLE_SIZE to have Oracle determine the appropriate sample size for good statistics. This is the default.

QUESTION 4

Examine the Following Query and execution plan:



```
SQL> SELECT C.cust_last_name, c.cust_city
FROM customers C,
     (SELECT DISTINCT S.S.cust_id)
     FROM SALES S, costs CT
     WHERE S.Prod_id and CT.unit_price > 70 v
     WHERE S.prod_id= CT.prod_id and CT.unit_price > 70)v
WHERE C.cust_state_province = 'CA' and C.cust_id = V.cust_id;
```

695 rows selected.

Execution Plan

Plan hash value: 3834618923

id	Operation	Name	Rows	Bytes	Cots	(%CPU)	TIME	Pstart	Pstop
0	SELECT STATEMENT		49	1764	50194	(2)	00:10:03		
1	NESTED LOOPS SEMT		49	1764	50194	(2)	00:10:03		
2	TABLE ACCESS FULL	CUSTOMERS	383	13022	406	(1)	00:00:06		
3	VIEW PUSHED PREDICATE		18206	36410	130	3	00:00:02		
4	HASH JOIN		143K	2615K	130	(3)	00:00:02		
5	PARTITION RANGE ALL		130	1170	54	(0)	00:00:01	1	30
6	TABLE ACCESS BY LOCAL INDEX ROWID	SALES	130	1170	54	(0)	00:00:01		
7	BTMAP CONVERSION TO ROWIDS								
8	BITMAP INDEX SINGLE VALUE	SALES_CUST_BIX						1	8
9	PARITION RANGE ALL		79173	695K	74	(2)	00:00:01	1	28
10	TABLE ACCESS FULL	COSTS	79172	695K	74	(2)	00:00:01	1	28

Predicate Information (identified by operation id)

```
2 - filter ("C", "CUST_STATE_PROVINCE"='CA')
4 - access ("S", "PROD_ID"="CUST.PROD_ID")
8 - access ("S", "CUST_ID"="CUST_ID")
10 - filter ("CT", "UNIT_PRICE"> 70)
```

Statistics

```
16986 recursive calls
0 db block gets
231975 consistent gets
4030 physical reads
0 redo size
19001 bytes sent via SQL*NET to client
389 bytes received via SQL*NET from client
45 SQL*Net roundtrips to/from client
113 sorts (memory)
0 sorts (disk)
659 rows processed
```

Which query transformation technique is used by the optimizer?

- A. Filter push down
- B. Subquery factoring
- C. Subquery unnesting
- D. Predicate pushing

Correct Answer: D



Note:

* In the execution plan BX, note the keyword `VIEW PUSHED PREDICATE` indicates that the view has undergone the join predicate pushdown transformation.

QUESTION 5

You are administering a database supporting an OLTP workload. A new module was added to one of the applications recently in which you notice that the SQL statements are highly resource intensive in terms of CPU, I/O and temporary space. You created a SQL Tuning Set (STS) containing all resource-intensive SQL statements. You want to analyze the entire workload captured in the STS. You plan to run the STS through the SQL Advisor.

Which two recommendations can you get?

- A. Combing similar indexes into a single index
- B. Implementing SQL profiles for the statements
- C. Syntactic and semantic restructuring of SQL statements
- D. Dropping unused or invalid index.
- E. Creating invisible indexes for the workload
- F. Creating composite indexes for the workload

Correct Answer: CF

The output of the SQL Tuning Advisor is in the form of an advice or recommendations, along with a rationale for each recommendation and its expected benefit. The recommendation relates to collection of statistics on objects , creation of new indexes (F), restructuring of the SQL statement (C), or creation of a SQL profile. You can choose to accept the recommendation to complete the tuning of the SQL statements.

Reference: Oracle Database Performance Tuning Guide 11g , SQL Tuning Advisor

QUESTION 6

Examine the parallelism parameters for your instance:



NAME	TYPE	VALUE
parallel_adaptive_multi_user	boolean	TRUE
Parallel_automatic_tuning	boolean	FALSE
parallel_degree_limit	string	CPU
parallel_degree_Opolicy	string	LIMITED
parallel_execution_local	integer	16384
parallel_force_local	boolean	FALSE
parallel_io_cap_enabled	boolean	FALSE
parallel_min_percent	integer	80
parallel_min_servers	integer	20
parallel_min_time_threshold	string	AUTO
parallel_server	boolean	FALSE
parallel_server_instances	integer	1
parallel_servers_target	integer	8
parallel_threads_per_cpu	integer	2

You executed the following query:

```
SQL> SELECT /*+ PARALLEL (AUTO)*/SUM(1_extendedprice) total_rev
FROM lineitem;
```

What are true about the execution of the query?

- A. It will execute in parallel only if the LINEITEM table has a dictionary DOP defined.
- B. DOP for the statement is determined by the dictionary DOP of the accessed objects.
- C. It is generated to execute in parallel.
- D. It will execute in parallel only if the estimated execution time is 10 or more seconds.
- E. DOP for the statement is calculated automatically.
- F. It may execute serially.

Correct Answer: EF

E:

F (not C): It may execute serially. See note below.

Incorrect:

A, B: Dictionary DOP not used with PARALLEL (AUTO) hint.

D: The default value of parallel_min_time_threshold is 30 (not 10) seconds.

Note:

* parallel_min_percent 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.

QUESTION 7

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 these 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 8

The following parameter values are set for the instance:

OPTIMIZER_CAPTURE_SQL_BASELINE = FALSE OPTIMIZER_USESQL_PLAN_BASELINE = TRUE

The SQL plan baseline for a SQL statement contains an accepted plan.

You want to add a new plan automatically as an accepted plan to the existing SQL plan baseline.

Examine the following tasks.

1.
Set the OPTIMIZER_CAPTURE_SQL_PLAN_BASELINE parameter to TRUE.
- 2.



Evolve the new plan.

3.

Fix the existing accepted plan.

4.

Manually load the new plan.

Identify the task(s) that must be performed to accomplish this.

A. 1, 2, and 3

B. 4 and 3

C. 1, 4, and 3

D. Only 4

E. 1, 2, 4, and 3

F. 1 and 2

Correct Answer: D

Manual Plan Loading Manual plan loading can be used in conjunction with, or as an alternative to automatic plan capture. The load operations are performed using the DBMS_SPM package, which allows SQL plan baselines to be loaded from SQL tuning sets or from specific SQL statements in the cursor cache. Manually loaded statements are flagged as accepted by default. If a SQL plan baseline is present for a SQL statement, the plan is added to the baseline, otherwise a new baseline is created.

Note:

*

The value of the OPTIMIZER_CAPTURE_SQL_PLAN_BASELINES parameter, whose default value is FALSE, determines if the system should automatically capture SQL plan baselines. When set to TRUE, the system records a plan history for SQL statements. The first plan for a specific statement is automatically flagged as accepted. Alternative plans generated after this point are not used until it is verified they do not cause performance degradations. Plans with acceptable performance are added to the SQL plan baseline during the evolution phase.

*

Managing SQL plan baselines involves three phases:

Capturing SQL Plan Baselines Selecting SQL Plan Baselines Evolving SQL Plan Baselines

*

Evolving SQL Plan Baselines

Evolving a SQL plan baseline is the process by which the optimizer determines if non-accepted plans in the baseline should be accepted. As mentioned

previously, manually loaded plans are automatically marked as accepted, so manual loading forces the evolving process. When plans are loaded automatically,



the baselines are evolved using the

EVOLVE_SQL_PLAN_BASELINE function, which returns a CLOB reporting its results.

Reference: SQL Plan Management in Oracle Database 11g Release 1

QUESTION 9

You are administering a database that supports an OLTP workload. Automatic optimizer statistics collection is scheduled in the night maintenance window. Some of the tables get updated frequently during day time and you notice a performance degradation for queries using those tables due to stale statistics.

Which two options might help to avoid the performance degradation of the queries?

- A. Set the global statistics preference STALE_PERCENT to 0.
- B. Use dynamically sampling hint for the queries on frequently updated tables.
- C. Create histogram for the columns used frequently in the WHERE clause.
- D. Gather statistics with global statistics preference NO_VALIDATE to TRUE.
- E. Set the OPTIMIZER_USE_PENDING_STATISTICS parameter to TRUE.

Correct Answer: BC

B: Dynamic sampling first became available in Oracle9i Database Release 2. It is the ability of the cost-based optimizer (CBO) to sample the tables a query references during a hard parse, to determine better default statistics for unanalyzed segments, and to verify its "guesses." This sampling takes place only at hard parse time and is used to dynamically generate better statistics for the optimizer to use, hence the name dynamic sampling.

There are two ways to use dynamic sampling:

The OPTIMIZER_DYNAMIC_SAMPLING parameter can be set at the database instance level and can also be overridden at the session level with the ALTER

SESSION command.

The DYNAMIC_SAMPLING query hint can be added to specific queries.

C: A histogram is a collection of information about the distribution of values within a column.

In some cases, the distribution of values within a column of a table will affect the optimizer's decision to use an index vs. perform a full-table scan. This scenario occurs when the value with a where clause has a disproportional amount of values, making a full-table scan cheaper than index access.

Histograms are also important for determine the optimal table join order.

Incorrect:

A: Too much statistics would be gathered.

Note: 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%.

D: In Oracle PL/SQL, the VALIDATE keyword defines the state of a constraint on a column in a table.

E: OPTIMIZER_USE_PENDING_STATISTICS specifies whether or not the optimizer uses pending statistics when compiling SQL statements.

QUESTION 10

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 11

Partial details of an execution plan.

```

SELECT STATEMENT
SORT GROUP BY
HASH JOIN
TABLE ACCESS FULL CHANNELS
HASH JOIN
TABLE ACCESS FULL CUSTOMERS
HASH JOIN
TABLE ACCESS FULL TIMES
PARTITION RANGE ITERATOR
TABLE ACCESS BY LOCAL INDEX ROWID SALES
BITMAP CONVERSION TO ROWIDS
BITMAP AND
BITMAP MERGE
BITMAP KEY ITERATOR
BUFFER SORT
TABLE ACCESS FULL CHANNELS
BITMAP INDEX RANGE SCAN SALES_CHANNEL_BIX
BITMAP MERGE
BITMAP KEY ITERATION
BUFFER SORT
TABLE ACCESS FULL TIMES
BITMAP INDEX RANGE SCAN SALES_TIME_BIX
BITMAP MERGE
BITMAP KEY ITERATION
BUFFER SORT
TABLE ACCESS FULL TIMES
BITMAP INDEX RANGE SCAN SALES_TIME_BIX

```

Which statement correctly describes the BITMAP AND operation?

- A. It produces a bitmap, representing dimension table rows from all dimension tables that join with qualified fact table rows.
- B. It produces a concentration of the bitmaps for all dimension tables.
- C. It produces a bitmap, representing fact table rows that do not join with qualified dimension table rows from all dimension tables.
- D. It produces a bitmap, representing fact table rows that join with qualified dimension table rows from all dimension tables.

Correct Answer: D

Example:



Additional set operations will be done for the customer dimension and the product dimension. At this point in the star query processing, there are three bitmaps. Each bitmap corresponds to a separate dimension table, and each bitmap represents the set of rows of the fact table that satisfy that individual dimension's constraints.

These three bitmaps are combined into a single bitmap using the bitmap AND operation. This final bitmap represents the set of rows in the fact table that satisfy all of the constraints on the dimension table.

Reference: Oracle Database Data Warehousing Guide, Star Transformation with a Bitmap Index

QUESTION 12

Examine the parallelism parameters for you 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

Now examine the DSS_PLAN with parallel statement directives:

```
Consumer Group : URGENT_GROUP
MGMT_P1 : 100%
PARALLEL_DEGREE_LIMIT_91 : 12
PARALLEL_TARGET_PERCENTAGE :
PARALLEL_QUEUE_TIMEOUT :
```

```
Consumer Group : ETL_GROUP
MGMT_P2 : 100%
PARALLEL_DEGREE_LIMIT_P1 : 8
PARALLEL_TARGET_PERCENTAGE : 50%
PARALLEL_QUEUE_TIMEOUT :
```

```
Consumer Group : OTHER_GROUPS
MGMT_P3 : 100%
PARALLEL_DEGREE_LIMIT_P1 : 2
PARALLEL_TARGET_PERCENTAGE : 50%
PARALLEL_QUEUE_TIMEOUT : 360
```

Which two are true about the DSS_PLAN resource manager plan?



- A. URGENT_GROUPS sessions will always be dequeued before sessions from other groups.
- B. OTHER_GROUPS sessions are queued for maximum of six minutes.
- C. ETL_GROUP sessions can collectively consume 64 parallel execution servers before queuing starts for this consumer group.
- D. An ETL_GRP sessions will be switched to URGENT_GROUPS if the session requests more than eight parallel executions servers.
- E. URGENT_GROUP sessions will not be queued if 64 parallel execution servers are busy because their PARALLEL_TARGET_PERCENTAGE is not specified.

Correct Answer: AB

B: PARALLEL_QUEUE_TIMEOUT Parallel Queue Timeout

When you use parallel statement queuing, if the database does not have sufficient resources to execute a parallel statement, the statement is queued until the

required resources become available. However, there is a chance that a parallel statement may be waiting in the parallel statement queue for longer than is

desired. You can prevent such scenarios by specifying the maximum time a parallel statement can wait in the parallel statement queue.

The PARALLEL_QUEUE_TIMEOUT directive attribute enables you to specify the maximum time, in seconds, that a parallel statement can wait in the parallel

statement queue before it is timed out. The PARALLEL_QUEUE_TIMEOUT attribute can be set for each consumer group.

Incorrect:

Not C: ETL_GROUP PARALLEL_TARGET_PERCENTAGE is 50%. So ETL_GROUP can only consume 32 servers.

Note:

* If you want more per-workload management, you must use the following directive attributes:

/ MGMT_Pn

Management attributes control how a parallel statement is selected from the parallel statement queue for execution. You can prioritize the parallel statements of one consumer group over another by setting a higher value for the management attributes of that group.

/ PARALLEL_TARGET_PERCENTAGE

/ PARALLEL_QUEUE_TIMEOUT

/ PARALLEL_DEGREE_LIMIT_P1

* PARALLEL_DEGREE_LIMIT_P1 Degree of Parallelism Limit

You can limit the maximum degree of parallelism for any operation within a consumer group. The degree of parallelism is the number of parallel execution servers that are associated with a single operation. Use the PARALLEL_DEGREE_LIMIT_P1 directive attribute to specify the degree of parallelism for a consumer group.



The degree of parallelism limit applies to one operation within a consumer group; it does not limit the total degree of parallelism across all operations within the consumer group. However, you can combine both the `PARALLEL_DEGREE_LIMIT_P1` and the `PARALLEL_TARGET_PERCENTAGE` directive attributes to achieve the desired control.

* `PARALLEL_TARGET_PERCENTAGE` Parallel Target Percentage

It is possible for a single consumer group to launch enough parallel statements to use all the available parallel servers. If this happens, when a high-priority parallel statement from a different consumer group is run, no parallel servers are available to allocate to this group. You can avoid such a scenario by limiting the number of parallel servers that can be used by a particular consumer group.

Use the `PARALLEL_TARGET_PERCENTAGE` directive attribute to specify the maximum percentage of the parallel server pool that a particular consumer group can use. The number of parallel servers used by a particular consumer group is counted as the sum of the parallel servers used by all sessions in that consumer group.

QUESTION 13

Examine the Exhibit.



SQL> EXPLAIN FRO SELECT /*+parallel*/cust_city.sum(amount_sold)

FROM sales.s.customers c WHERE s.cust_id=c.cust_id
GROUP BY cust_city;

SQL> SELECT *FROM TABLE (dbms_xplain display)

PLAN_TABLE_OUTPUT

Plan hash value: 24252389986

Id	Operation	Name	Rows	Bytes	Cost	(%CPU)	Time	Pstart	Pstop	TQ	IN-OUT	PC Distrib
0	SELECT STATEMENT		620	20460	512	(4)	00:00:07					
1	PX COORDINATOR											
2	PX SEND QC (RANDOM)	:TQ10003	620	20460	512	(4)	00:00:07			Q1,03	P->S	QC (RAND)
3	HASH GROUP BY		620	20460	512	(4)	00:00:07			Q1,03	PCWP	
4	PX RECEIVE		620	20460	512	(4)	00:00:07			Q1,03	PCWP	
5	PX SEND HASH	:TQ10002	620	20460	512	(4)	00:00:07			Q1,02	P->P	HASH
6	HASH GROUP BY		620	20460	512	(4)	00:00:07			Q1,02	PCWP	
7	HASH JOIN		7059	227K	551	(4)	00:00:07			Q1,02	PCWP	
8	PX SERVICE		7059	124K	286	(7)	00:00:04			Q1,01	P->P	BROADCAST
9	PX SEND BROADCAST	:TQ10001	7059	124K	286	(7)	00:00:04			Q1,01	PCWP	
10	VIEW	VW_GBC_5	7059	124K	286	(7)	00:00:04			Q1,01	PCWP	
11	HASH GROUP BY		7059	70590	286	(7)	00:00:04			Q1,01	PCWP	
12	PX RECEIVE		7059	70590	286	(7)	00:00:04			Q1,01	PCWP	
13	PX SEND HASH	:TQ7059	7059	70590	286	(7)	00:00:04			Q1,00	P->P	HASH
14	HASH GROUP BY		7059	70590	286	(7)	00:00:04			Q1,00	PCWP	
15	PX BLOCK ITERATOR		918K	6973K	271	(2)	00:00:04	1	28	Q1,00	PCWE	
16	TABLE ACCESS FULL	SALES	918K	7973K	271	(2)	00:00:04	1	28	Q1,00	PCWP	
17	PX BLOCK ITERATOR		55500	612K	255	(1)	00:00:03			Q1,02	PCEC	
18	TABLE ACCESS FULL	CUSTOMERS	55500	612K	255	(1)	00:00:03			Q1,02	PCWP	

Predicate Information (Identify by operation id)


7 - access ("ITEM-1" = "C", "CUST_ID")

Note

PLAN_TABLE_OUTPUT

- automatic DOP Computed Degree of Parallelism is 2

34 rows selected





Given two sets of parallel execution processes, SS1 and SS2, which is true?

- A. Each process SS1 reads some of the rows from the CUSTOMERS table and sends all the rows it reads to each process in SS2.
- B. Each process in SS1 reads all the rows from the CUSTOMERS table and distributes the rows evenly among the processes in SS2.
- C. Each process in SS1 reads some of the rows from the SALES table and sends all the rows it reads to each process in SS2.
- D. Each process in SS1 reads all the rows from the SALES table and distributes the rows evenly among the processes in SS2.
- E. Each process in SS1 reads some of the rows from the SALES table and distributes the rows evenly among the processes in SS2.
- F. Each process in the SS1 reads some of the rows from the CUSTOMERS table and distributes the rows evenly among the processes in SS2.

Correct Answer: D

Note:

*

The execution starts with line 16 (accessing the SALES table), followed by line 15.

*

PX BLOCKITERATOR The PX BLOCK ITERATOR row source represents the splitting up of the table EMP2 into pieces so as to divide the scan workload between the parallel scan slaves. The PX SEND and PX RECEIVE row sources represent the pipe that connects the two slave sets as rows flow up from the parallel scan, get repartitioned through the HASHtable queue, and then read by and aggregated on the top slave set.

QUESTION 14

A database instance is configured in the shared server mode and it supports multiple applications running on a middle tier. These applications connect to the database using different services. You enabled the statistics gathering for the service by using the following command:

```
SQL > EXECUTE DBMS_MONITOR.SERV_MOD_ACT_STAT_ENABLE ('APPS1\\', NULL, NULL);
```

Which two statements are true regarding statistics gathered for APPS1 service?

- A. The statistics are gathered for all the modules and actions within the service.
- B. The statistics are collected at the session level for all sessions connected using the service.
- C. The statistics are aggregated and stored in the V\$SERV_MOD_ACT_STATS view.
- D. The statistics are gathered for all the modules using the service only when DBMS_APPLICATION_INFO.SET_MODULE is executed to register with the service.
- E. Statistics gathering is enabled only for the subsequent sessions using the service.



F. The statistics are gathered for all the applications using the service only when DBMS_APPLICATION_INFO.SET_ACTION is executed to register with the service.

Correct Answer: AC

SERV_MOD_ACT_STAT_ENABLE Procedure

This procedure enables statistic gathering for a given combination of Service Name, MODULE and ACTION. Calling this procedure enables statistic gathering for a hierarchical combination of Service name, MODULE name, and ACTION name on all instances for the same database. Statistics are accessible by means of the V\$SERV_MOD_ACT_STATS view.

Note:

* Syntax

```
DBMS_MONITOR.SERV_MOD_ACT_STAT_ENABLE(  
service_name IN VARCHAR2,  
module_name IN VARCHAR2,  
action_name IN VARCHAR2 DEFAULT ALL_ACTIONS);
```

Parameter, Description

service_name

Name of the service for which statistic aggregation is enabled

module_name

Name of the MODULE. An additional qualifier for the service. It is a required parameter.

action_name

Name of the ACTION. An additional qualifier for the Service and MODULE name. Omitting the parameter (or supplying ALL_ACTIONS constant) means enabling

aggregation for all Actions for a given Server/Module combination. In this case, statistics are aggregated on the module level.

Reference: Oracle Database PL/SQL Packages and Types Reference 11g, SERV_MOD_ACT_STAT_ENABLE Procedure

QUESTION 15

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

[1Z0-117 VCE Dumps](#)

[1Z0-117 Practice Test](#)

[1Z0-117 Braindumps](#)