



1Z0-117^{Q&As}

Oracle Database 11g Release 2: SQL Tuning Exam





Pass Oracle 1Z0-117 Exam with 100% Guarantee

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

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

100% Passing Guarantee
100% Money Back Assurance

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

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





QUESTION 1

Which statement is true about the usage of the STAR_TRANSFORMATION hint in a query?

- A. The optimizer always uses a plan in which the transformation is used.
- B. The optimizer uses transformation only if the cost is less than a query executing without transformation.
- C. The optimizer always generates subqueries to transform a query.
- D. The optimizer always uses bitmap indexes on the primary key column for any dimension table to transform a query.

Correct Answer: B

(http://docs.oracle.com/cd/E11882_01/server.112/e41084/sql_elements006.htm#SQLRF50508) says that "Even if the hint is specified, there is no guarantee that the transformation will take place".

QUESTION 2

Which three factors does the estimator depend on for overall cost estimation of a given execution plan?

- A. Cardinality
- B. Sort area size
- C. OPTIMIZER_FEATURE_ENABLE parameter
- D. NOT NULL_FEATURE_ENABLE parameter
- E. NOT NULL constraint on a unique key column
- F. Library cache size
- G. The units of work such as disk input/output, CPU usage, and memory used in an operation

Correct Answer: ACG

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

Note: The estimator determines the overall cost of a given execution plan. The estimator generates three different types of measures to achieve this goal:

*

Selectivity

This measure represents a fraction of rows from a row set. The selectivity is tied to a query predicate, such as last_name='Smith', or a combination of predicates.

*

Cardinality



This measure represents the number of rows in a row set. 1,

*

Cost

This measure represents units of work or resource used. The query optimizer uses disk I/O, CPU usage, and memory usage as units of work.

If statistics are available, then the estimator uses them to compute the measures. The statistics improve the degree of accuracy of the measures.

QUESTION 3

An application accessing your database got the following error in response to SQL query:

ORA-12827: insufficient parallel query slaves available

View the parallel parameters for your instance:

NAME	TYPE	VALUE
fast_start_parallel_rollback	string	LOW
parallel_adaptive_multi_tuning	boolean	TRUE
parallel_automatic_tuning	boolean	FALSE
parallel_degree_limit	string	32
parallel_degree_policy	string	LIMITED
parallel_execution_message_size	integer	16384
parallel_force_local	boolean	FALSE
parallel_io_cap_enabled	boolean	FALSE
parallel_max_servers	integer	128
parallel_min_servers	integer	50
parallel__server	Integer	0
parallel_server_instances	string	AUTO
parallel_server	boolean	1
parallel_servers_target	integer	8
parallel_threads_servers_per_cpu	integer	2

No hints are used and the session use default parallel settings.

What four changes could you make to help avoid the error and ensure that the query executes in parallel?

- A. Set PARELLEL_DEGREE_POLICY to AUTO.
- B. Increase the value of PARELLEL_MAX_SERVERS.
- C. Increase PARELLEL_SERVERS_TARGET.



- D. Decrease PARELLEL_MIN_PERCENT.
- E. Increase PARELLEL_MIN_SERVERS.
- F. Decrease PARELLEL_MIN_TIME_THRESHOLD.
- G. Increase PARELLEL__MIN_TIME_THRESHOLD.

Correct Answer: ACDG

C: 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.

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

D: Note: ORA-12827: insufficient parallel query slaves available Cause: PARALLEL_MIN_PERCENT parameter was specified and fewer than minimum slaves were acquired Action: either re-execute query with lower PARALLEL_MIN_PERCENT or wait until some running queries are completed, thus freeing up slaves

A, G: 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.

QUESTION 4

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 5

Examine the exhibit.



Id	Operations	Name	Rows	Bytes	Cost	(%CPU)	Time	IQ	IN-OUT	PQ Disturb
0	SELECT STATEMENT		14	588	5	(20)	00:00:01			
1	PX COORDINATOR									
2	PS SEND QC (RANDOM)	:lq10001	14	588	5	(20)	00:00:01	Q1, 01,	P -> S	QC(RAND)
3	HASH JOIN		14	588	5	(20)	00:00:01	Q1, 01	PCWP	
4	SUFFER SORT							Q1, 01	PCWC	
5	PX RECEIVE		4	88	2	(0)	00:00:01	Q1, 01	PCWP	
6	PX SEND BROADCAST	:lq10000	4	88	2	(0)	00:00:01		S -> P	BROADCAST
7	TABLE ACCESS FULL	DEPARTMENTS	4	88	2	(0)	00:00:01			
8	PX BLOCK ITERATOR		14	280	2	(0)	00:00:01	Q1, 01	PCWC	
9	TABLE ACCESS FUL	EMPLOYEES	14	280	2	(0)	00:00:01	Q1, 01	PCWP	

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

- A. A full partition-wise join performed between the EMPLOYEES and DEPARTMENTS tables.
- B. A full table scan on the DEPARTMENTS table performed serially by the query coordinator.
- C. A full table scan on the DEPARTMENTS table is performed serially by a single parallel execution server process.
- D. A partial partition-wise join performed between the EMPLOYEES and DEPARTMENTS tables.
- E. A full table scan on the EMPLOYEES table is done in parallel.

Correct Answer: E

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

each of the parallel servers involved.

Incorrect:

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

Note:

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

*

Example with same structure of execution plan:



Id	Operation	Name	Rows (Estim)	Cost	Time Active(s)	Start Active	Execs	Rows (Actual)	Read Reqs	Read Bytes	Mem (Max)	Activity (%)	Activity Detail (# samples)
0	SELECT STATEMENT				1	+3	1	1					
1	SORT AGGREGATE		1		1	+3	1	1					
2	PX COORDINATOR				1	+3	9	8					
3	PX SEND QC (RANDOM)	:TQ10001	1		6	+2	8	8					
4	SORT AGGREGATE		1		6	+2	8	8					
5	HASH JOIN		46K	291	6	+2	8	452K			38M		
6	BUFFER SORT				6	+2	8	603K			19M		
7	PX RECEIVE		75272	76	6	+2	8	603K					
8	PX SEND BROADCAST	:TQ10000	75272	76	1	+3	1	603K					
9	INDEX FAST FULL SCAN	I_OBJ1	75272	76	1	+3	1	75322					
10	PX BLOCK ITERATOR		46K	214	6	+2	8	452K					
11	TABLE ACCESS FULL	MYOBJ	46K	214	6	+2	104	452K	1623	44MB			

Here's how to read the plan:

1.

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

2.

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

3.

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

4.

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

5.

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

6.

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

7.

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

8.

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

9.

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

10.



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

[1Z0-117 Practice Test](#)

[1Z0-117 Study Guide](#)

[1Z0-117 Braindumps](#)