



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.geekcert.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:

```
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)
559	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 2

```
SQL> EXPLAIN PLAN SET Statement_id = 'test' for
      SELECT prod_category, avg(amount_sold)
      FROM sales s, products p
      WHERE p.prod_id = s.prod_id
      GROUP BY prod_Category;
```

Explained.

```
SQL> SELECT id "id", parent_id, position "pos"
      lpad(' ', 2 level) || operations || decode(id, 0, 'cost=' || POSITION) "operations"
Options "option" object_name "object"
FROM plan_table
Connect by prior id_parent_id START WITH id = 0
ORDER BY id;
```

id	PARENT_ID	POS	Operation	Option	Object
0		539	SELECT STATEMENT Cost = 539		
1	0	1	HASH	Group By	
2	1	1	HASH JOIN		
3	2	1	VIEW		
4	3	1	HASH	GROUP BY	
5	4	1	PARTITION RANGE	ALL	
6	5	1	TABLE ACCESS	FULL	SALES
7	2	2	VIEW		indes\$_joins\$_002
8	7	1	VIEW RANGE		
9	8	1	INDEX	FAST FULL SCAN	PRODUCTS_PK
10	8	2	INDEX	FAST FULL SCAN	PRODUCTS_PROD_CAT_IX

11 rows are selected

View the exhibit and examine the query and its execution plan from the PLAN_TABLE. Which statement is true about the execution?

- A. The row with the ID column having the value 0 is the first step execution plan.
- B. Rows are fetched from the indexes on the PRODUCTS table and from the SALES table using full table scan simultaneously, and then hashed into memory.
- C. Rows are fetched from the SALES table, and then a hash join operator joins with rows fetched from indexes on the PRODUCTS table.
- D. All the partitions of the SALES table are read in parallel.

Correct Answer: C



QUESTION 3

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

Examine the query:

```
SQL> SELECT /*+PARALLEL (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 be 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, the NO_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 automatically 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 4



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

QUESTION 5

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)

[1Z0-117 Study Guide](#)

[1Z0-117 Exam Questions](#)

[1Z0-117 Braindumps](#)