



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.pass4itsure.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**

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. Combining 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 2

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 d  
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 3

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;
```

EXECUTION 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")

id	operation
0	SELECT STATEMENT
1	SORT ORDER BY
*2	HASH JOIN
3	TABLE ACCESS FULL
4	TABLE ACCESS FULL

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 statement be joined in the order specified, with the first table in the FROM clause specifying the driving table.

QUESTION 4

Which two statements are true about index full scans?

A. An index fast full scan multi block I/O to read the index structure in its entirety.

B. Index nodes are not retrieved in the index order, and therefore the nodes are not in sequence.

C. An index fast full scan reads the index block by block.

D. An index fast full scan reads the whole index from the lowest value to the higher value.

Correct Answer: AB

A: To speed table and index block access, Oracle uses the `db_file_multiblock_read_count` parameter (which defaults to 8) to aid in getting full-table scan and full-index scan data blocks into the data buffer cache as fast as possible.

B: The index nodes are not retrieved in index order, the rows will not be sequenced.

Note:

there are some requirements for Oracle to invoke the fast full-index scan.

All of the columns required must be specified in the index. That is, all columns in the select and where clauses must exist in the index.

The query returns more than 10 percent of the rows within the index. This 10 percent figure depends on the degree of multi-block reads and the degree of

parallelism.

You are counting the number of rows in a table that meet a specific criterion. The fast full-index scan is almost always used for `count(*)` operations.

Reference: index fast full scan tips

QUESTION 5

Examine the exhibit.



Id	Operation	Name	Pstprt	Pstop	IN-OUT	PQ	Distrib
0	SELECT STATEMENT						
1	PX COORDINATOR						
2	PX SEND QC (RANDOM)	:TQ10001			P->S	QC	(RAND)
*3	FILTER				PCWC		
4	HASH GROUP BY				PCWP		
5	PX RECEIVE				PCWP		
6	PX SEND HASH	:TQ10000			P->P	HASH	
7	HASH GRIYP BY		1	16	PCWC		
8	PX PARTITION HASH ALL		1	16	PCWP		
*9	HASH JOIN				PCWP		
10	TABLE ACCESS FULL	CUSTOMERS	1	16	PCWP		
11	PX PARTITION RANGE ITERATOR		8	9	PCWC		
*12	TABLE ACCESS FULL	SALES	113	144	PCWP		

Predicate information (identified by operation id):

3 – filter (COUNT (SYS_OP_CSR(SYS_OP_MSR(COUNT(''), 0))>100)

9 – access ("S". "CUST_ID"= "C". "CUST_ID")

12 – filter ("S". "TIME_ID"<= TO_DATE ('1999-10-01 00:00:00', 'syyy-mm-dd hh2:mi:ss') AND
"S". "TIME_ID">=TO_DATE('1999-07-01
00:00:00', 'syyy-mm-dd hh24:mi:ss')

Which two are true concerning the execution plan?

- A. No partition-wise join is used
- B. A full partition-wise join is used
- C. A partial partition-wise join is used
- D. The SALES table is composite partitioned

Correct Answer: BD

* The following example shows the execution plan for the full partition-wise join with the sales table range partitioned by time_id, and subpartitioned by hash on

cust_id. ----- | Id | Operation | Name | Pstart| Pstop
|IN-OUT| PQ Distrib |

----- | 0 | SELECT STATEMENT | | | |

| 1 | PX COORDINATOR | | | |

| 2 | PX SEND QC (RANDOM) | :TQ10001 | | P->S | QC (RAND) |

|* 3 | FILTER | | | PCWC |

| 4 | HASH GROUP BY | | | PCWP |

| 5 | PX RECEIVE | | | PCWP |

| 6 | PX SEND HASH | :TQ10000 | | P->P | HASH |

| 7 | HASH GROUP BY | | | PCWP |



| 8 | PX PARTITION HASH ALL | | 1 | 16 | PCWC | |

|* 9 | HASH JOIN | | | PCWP | |

| 10 | TABLE ACCESS FULL | CUSTOMERS | 1 | 16 | PCWP | |

| 11 | PX PARTITION RANGE ITERATOR | | 8 | 9 | PCWC | |

|* 12 | TABLE ACCESS FULL | SALES | 113 | 144 | PCWP | |

Predicate Information (identified by operation id):

3 - filter(COUNT(SYS_OP_CSR(SYS_OP_MSR(COUNT(*)),0))>100)

9 - access("S"."CUST_ID"="C"."CUST_ID")

12 - filter("S"."TIME_ID"=TO_DATE(\\ 1999-07-01

00:00:00\\', \\syddd-mm-dd hh24:mi:ss\\'))

* Full partition-wise joins can occur if two tables that are co-partitioned on the same key are joined in a query. The tables can be co-partitioned at the partition level, or at the subpartition level, or at a combination of partition and subpartition levels. Reference partitioning is an easy way to guarantee co-partitioning. Full partition-wise joins can be executed in serial and in parallel.

Reference: Oracle Database VLDB and Partitioning Guide, Full Partition-Wise Joins: Composite

-Single-Level

[Latest 1Z0-117 Dumps](#)

[1Z0-117 PDF Dumps](#)

[1Z0-117 Study Guide](#)