

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.pass2lead.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 plan to bulk load data INSERT INTO . . . SELECT FROM statements.

Which two situations benefit from parallel INSERT operations on tables that have no materialized views defined on them?

- A. Direct path insert of a million rows into a partitioned, index-organized table containing one million rows and a conventional B*tree secondary index.
- B. Direct path insert of a million rows into a partitioned, index-organized table containing 10 rows and a bitmapped secondary index.
- C. Direct path insert of 10 rows into a partitioned, index-organized table containing one million rows and conventional B* tree secondary index.
- D. Direct path insert of 10 rows into a partitioned, index-organized table containing 10 rows and a bitmapped secondary index
- E. Conventional path insert of a million rows into a nonpartitioned, heap-organized containing 10 rows and having a conventional B* tree index.
- F. Conventional path insert of 10 rows into a nonpartitioned, heap-organized table one million rows and a bitmapped index.

Correct Answer: AB

Note:

*

A materialized view is a database object that contains the results of a query.

*

You can use the INSERT statement to insert data into a table, partition, or view in two ways: conventional INSERT and direct-path INSERT.

*

With direct-path INSERT, the database appends the inserted data after existing data in the table. Data is written directly into datafiles, bypassing the buffer cache. Free space in the existing data is not reused. This alternative enhances performance during insert operations and is similar to the functionality of the Oracle direct-path loader utility, SQL*Loader. When you insert into a table that has been created in parallel mode, direct-path INSERT is the default.

*

Direct-path INSERT is not supported for an index-organized table (IOT) if it is not partitioned, if it has a mapping table, or if it is reference by a materialized view.

*

When you issue a conventional INSERT statement, Oracle Database reuses free space in the table into which you are inserting and maintains referential integrity constraints

*

Conventional INSERT always generates maximal redo and undo for changes to both data and metadata, regardless of the logging setting of the table and the archivelog and force logging settings of the database

QUESTION 2

Examine the statements being executed for the first time:

```
SQL SELECT last_name, email, salary, FROM employees WHERE employees_id = 100;
```

```
SQL> SELECT last_name, email, salary FROM employees WHERE employees_id = 110;
```

```
SQL > SELECT last_name, email, salary FROM employees WHERE employees_id = 130;
```

Steps followed by a SQL statement during parsing:

1.
Search for a similar statement in the shared pool.
2.
Search for an identical statement in the shared pool.
3.
Search the SQL area of identical statement already in the shared pool.
4.
Proceed through the remaining steps of the parse phase to ensure that the execution plan of the existing statements is applicable to the view statement.
5.
Perform hard parsing.
6.
Share the SQL area of the similar statement already in the shared pool.

Identify the required steps in the correct sequence used by the third query.

- A. 5, 1, 3, 4
- B. 2, 4, 3
- C. 5, 2, 3, 4
- D. 1, 4, 3
- E. Only 5

F. 2, 5

Correct Answer: F

2) before 5).

Note:

* When application code is run, Oracle attempts to reuse existing code if it has been executed previously and can be shared. If the parsed representation of the statement does exist in the library cache and it can be shared, then Oracle reuses the existing code. This is known as a soft parse, or a library cache hit. If Oracle is unable to use existing code, then a new executable version of the application code must be built. This is known as a hard parse, or a library cache miss.

Reference: Oracle Database Performance Tuning Guide, SQL Sharing Criteria

QUESTION 3

You are administering a database supporting a DDS workload in which some tables are updated frequently but not queried often. You have SQL plan baseline for these tables and you do not want the automatic maintenance task to gather statistics for these tables regularly.

Which task would you perform to achieve this?

- A. Set the INCREMENTAL statistic preference FALSE for these tables.
- B. Set the STALE_PERCENT static preference to a higher value for these tables.
- C. Set the GRANULARITY statistic preference to AUTO for these tables.
- D. Set the PUBLISH statistic preference to TRUE for these tables.

Correct Answer: B

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

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%. Reference: Oracle Database PL/SQL Packages and Types Reference

QUESTION 4

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:

*

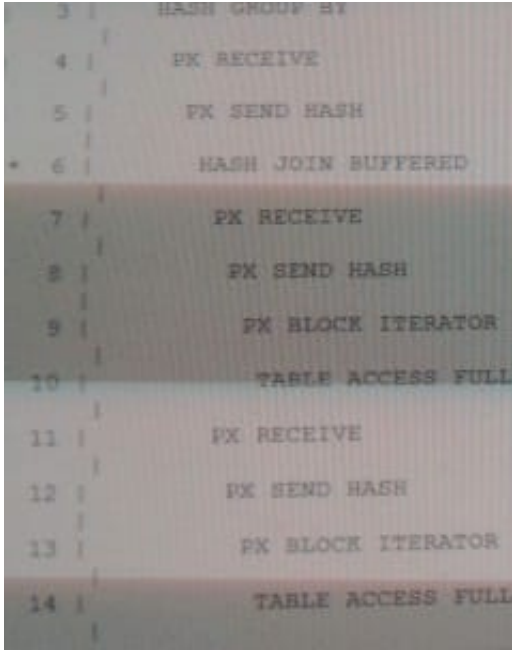
The execution starts with line 16 (accessing the SALES table), followed by line 15.

*

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 5

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



What two statements are true?

- A. The HASH GROUP BY operation is the consumer of the HASH operation.
- B. The HASH operation is the consumer of the HASH GROUP BY operation.
- C. The HASH GROUP BY operation is the consumer of the TABLE ACCESS FULL operation for the CUSTOMER table.
- D. The HASH GROUP BY operation is consumer of the TABLE ACCESS FULL operation for the SALES table.
- E. The SALES table scan is a producer for the HASH JOIN operation.

Correct Answer: AE

A, not C, not D: Line 3, HASH GROUP BY, consumes line 6 (HASH JOIN BUFFERED).

E: Line 14, TABLE ACCESS FULL (Sales), is one of the two producers for line 6 (HASH JOIN).

[Latest 1Z0-117 Dumps](#)

[1Z0-117 PDF Dumps](#)

[1Z0-117 Exam Questions](#)