



ITDUMPS

QUESTION & ANSWER

Accurate study guides, High passing rate!
IT dumps provides update free of charge
in one year!

[HTTP://WWW.ITDUMPS.COM](http://www.itdumps.com)

Exam : **1Z0-117**

Title : Oracle Database 11g
Release 2: SQL Tuning

Version : Demo

1.Examine the query and its execution plan:

```
SQL > SELECT cust_last_name, sum (nu12(o.customer_id, 0, 1)) "Count"
      FROM customer c, orders o
      WHERE c.credit_limit > 1000
      AND c.customer_id = o.customer_id(+)
      GROUP By cust_last_name;
```

ID	Operations	Name	Rows	Bytes	Cost (%CPU)
0	SELECT STATEMENT		168	3192	6 (17)
1	HASH GROUP BY		168	3192	6 (17)
* 2	NESTED LOOPS OUTER		260	4940	5 (0)
* 3	TABLE ACCESS FULL	CUSTOMER	260	3900	5 (0)
* 4	INDEX RANGE SCAN	ORD_CUSTOMER_IX	105	420	0 (0)

Predicate Information (identified by operation id):

```
-----
3 - filter ("C". "CREDIT_LIMIT"> 1000)
4 - access("C". "CUSTOMER_ID" = "0" "CUSTOMER_ID" (+))
   Filter ("O" "CUSTOMER_ID" (+)>0)
```

Which statement is true regarding the execution plan?

- A. This query first fetches rows from the CUSTOMERS table that satisfy the conditions, and then the join return NULL from the CUSTOMER_ID column when it does not find any corresponding rows in the ORDERS table.
- B. The query fetches rows from CUSTOMERS and ORDERS table simultaneously, and filters the rows that satisfy the conditions from the resultset.
- C. The query first fetches rows from the ORDERS table that satisfy the conditions, and then the join returns NULL form CUSTOMER_ID column when it does not find any corresponding rows in the CUSTOMERS table.
- D. The query first joins rows from the CUSTOMERS and ORDERS tables and returns NULL for the ORDERS table columns when it does not find any corresponding rows in the ORDERS table, and then fetches the rows that satisfy the conditions from the result set.

Answer: A

2.Which three statements are true about histograms?

- A. They capture the distribution of different values in an index for better selectivity estimates.
- B. They can be used only with indexed columns.
- C. They provide metadata about distribution of and occurrences of values in a table column.
- D. They provide improved selectivity estimates in the presence of data skew, resulting in execution plans with uniform distribution.
- E. They help the optimizer in deciding whether to use an index or a full table scan.
- F. They help the optimizer to determine the fastest table join order.

Answer: C,E,F

3.View the exhibit and examine the query and its execution plan from the PLAN_TABLE.

```
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

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.

Answer: C

4.Which four statements are correct about communication between parallel execution process?

- A. The number of logical pathways between parallel execution producers and consumers depends on the degree parallelism.
- B. The shared pool can be used for parallel execution messages buffers.
- C. The large pool can be used for parallel execution messages buffers.
- D. The buffer cache can be used for parallel execution message buffers.
- E. Communication between parallel execution processes is never required if a query uses full partition-wise joins.
- F. Each parallel execution process has an additional connection to the parallel execution coordinator.

Answer: A,B,E,F

5.You have enabled DML by issuing: ALTER session ENABLE PARALLEL DML;

The PARELLEL_DEGREE_POLICY initialization parameter is set to AUTO.

Which two options true about DML statements for which parallel execution is requested?

- A. Statements for which PDML is requested will execute serially estimated time is less than the time

specified by the PARALLEL_MIN_THRESHOLD parameter.

B. Statements for which PDML is requested will be queued if the number of busy parallel execution servers greater than PARALLEL_MIN_SERVERS parameter.

C. Statements for which PDML is requested will always execute in parallel if estimated execution in parallel if estimated execution time is greater than the time specified by the PARALLEL_MIN_TIME_THRESHOLD parameter.

D. Statements for which PDML is requested will be queued if the number of busy parallel execution servers is greater than PARALLEL_SERVERS_TARGET parameter.

E. Statement for which PDML is requested will be queued if the number of busy parallel execution servers is greater than PARALLEL_DEGREE_LIMIT parameter.

Answer: C,D

6.Examine Exhibit1 to view the query and its AUTOTRACE output.

```
SQL> SET AUTOTRACE TRACEONLY
SQL> SELECT prod_category, AVG(amount_sold)
      FROM sales s, products P
      WHERE P.prod_id = S.prod_id
      GROUP BY prod_category;
```

Execution Plan

Plan hash value: 1197568639							
id	Operation	Name	Rows	Bytes	Cost (%CPU)	Time	Pstart
0	SELECT STATEMENT		5	255	639 (11)	00:00:07	
1	HASH GROUP BY		5	255	539 (11)	00:00:07	
2	HASH JOIN		72	3672	538 (11)	00:00:07	
3	VIEW	VW_GBC_5	72	2160	535 (11)	00:00:07	
4	HASH GROUP BY		72	648	535 (11)	00:00:07	
5	PARTITION TANGE ALL		918K	8075K	494 (3)	00:00:06	1
6	TABLE ACCESS FULL	SALES	918K	8075K	684 (3)	00:00:06	1
7	VIEW	INDEX\$_JOIN_002	72	1512	3 (34)	00:00:01	
8	HASH JOIN						
9	INDEX FAST FULL SCAN	PRODUCT_PK	72	1512	1 (0)	00:00:01	
10	INDEX FAST FULL SCAN	PRODUCT_PROD_CAT_IX	72	1512	1 (0)	00:00:01	

Predicate information (identified by operation id)"

```
2- access ("P". "PROD_ID" = "ITEM_1")
8- access (ROEID=ROWID)
```

Statistics

```
-----
0      recursive calls
0      db block gets
1726   consistent gets
0      physical reads
0      Redo size
778   bytes sent via SQL "Net to client"
434   bytes received via SQL "Net from client"
2     SQL "Net roundtrips to/from client"
0     sorts (memory)
```

Which two statements are true about tracing?

A. The displayed plan will be stored in PLAN_TABLE.

B. Subsequent execution of this statement will use the displayed plan that is stored in v\$SQL.

C. The displayed plan may not necessarily be used by the optimizer.

D. The query will not fetch any rows; it will display only the execution plan and statistics.

E. The execution plan generated can be viewed from v\$SQLAREA.

Answer: A,C

7.Which two types of column filtering may benefit from partition pruning?

- A. Equally operates on range-partitioned tables.
- B. In-list operators on system-partitioned tables
- C. Equality operators on system-partitioned tables
- D. Operators on range-partitioned tables
- E. Greater than operators on hash-partitioned tables

Answer: A,D

8.Which two statements about In-Memory Parallel Execution are true?

- A. It can be configured using the Database Resource Manager.
- B. It increases the number of duplicate block images in the global buffer cache.
- C. It requires setting PARALLEL_DEGREE_POLICY to LIMITED.
- D. Objects selected for In-Memory Parallel Execution have blocks mapped to specific RAC instances.
- E. It requires setting PARALLEL_DEGREE_POLICY to AUTO
- F. Objects selected for In-Memory Parallel Execution must be partitioned tables or indexes.

Answer: D,E

9.Which two 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

Answer: A,C

10.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.

Answer: A,B

11. Which are the two prerequisites for enabling star transformation on queries?

- A. The STAR_TRANSFORMATION_ENABLED parameter should be set to TRUE or TEMP_DISABLE.
- B. A B-tree index should be built on each of the foreign key columns of the fact table(s).
- C. A bitmap index should be built on each of the primary key columns of the fact table(s).
- D. A bitmap index should be built on each of the foreign key columns of the fact table(s).
- E. A bitmap index must exist on all the columns that are used in the filter predicates of the query.

Answer: A,E

12. 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.

Answer: A,C,D,G

13. Examine the Exhibit 1 to view the structure of and indexes for EMPLOYEES and DEPARTMENTS tables.

SQL> desc employees

Name	Null?	Type
EMPLOYEE_ID	NOT NULL	NUMBER (6)
FIRST_NAME		VARCHAR2(20)
LAST_NAME	NOT NULL	VARCHAR2(25)
EMAIL	NOT NULL	VARCHAR2 (20)
PHONE-NUMBER		VARCHAR2 (20)
HIRE_DATE	NOT NULL	VARCHAR2(10)
JOB_ID	NOT NULL	VARCHAR2 (10)
SALARY		NUMBER (8, 2)
COMMISSION_PCT		NUMBER (2, 2)
MANAGER_ID		NUMBER (6)
DEPARTMENT_ID		NUMBER (4)

SQL> select index_name, index_type, from user_indexes where table_name = 'EMPLOYEES';

INDEX_NAME	INDEX_TYPE
EMP_NAME_IX	NORMAL
EMP_MANAGER_IX	NORMAL
EMP_JOB_IX	NORMAL
EMP_DEPARTMENT_IX	NORMAL
EMP_EMP_ID_PK	NORMAL
EMP_EMAIL_UK	NORMAL

SQL> desc departments

Name	Null	Type
DEPARTMENT_ID	NOT NULL	NUMBER (4)
DEPARTMENT_NAME	NOT NULL	VARCHAR (30)
MANAGER_ID		NUMBER (6)
LOCATION_ID		NUMBER (4)

SQL> select index_name, index_type from user_indexes where table_name = 'departments';

INDEX_NAME	INDEX_TYPE
DEPT_LOCATION_IX	NORMAL
DEPT_ID_PK	NORMAL

Examine the query and its auto trace output:

```
SQL> SELECT v.*d.department_name
        FROM (SELECT department_id, sum(salary) SUM-SAL
FROM employees GROUP BY department_id) v, department d
        WHERE v.department_id=d.department_id;
```

Execution Plan

Plan hash value: 5235477400

Id	Operation	Name	Rows	Bytes	Cost (%CPU)
[0]	SELECT STATEMENT		11	462	7 (29)
00:00:01					
[1]	Merge Join		11	462	7 (29)
00:00:01					
[2]	TABLE ACCESS BY INDEX ROWID	DEPARTMENTS	27	432	2 (0)
00:00:01					
[3]	INDEX FULL JOIN	DEPT_ID_PK	27	1 (0)	
00:00:01					
[4]	SORT JOIN		11	286	5 (40)
00:00:01					
[5]	VIEW		11	286	4 (25)
00:00:01					
[7]	TABLE ACCESS FULL	EMPLOYEES	107	749	3 (0)

Predicate Information (Identified by operation id):

```
4- access ("V". DEPARTMENT_ID = "D". "DEPARTMENT_ID")
   filter ("V". "DEPARTMENT_ID" = "D" "DEPARTMENT_ID")
```

Statistics

```
152   recursive calls
0     db block gets
36   consistent gets
8     physical reads
0     redo size
963   bytes sent via SQL*NET to client
524   bytes received via SQL*NET from client
2     SQL*NET roundtrips to/from client
2     sorts (memory)
0     sorts (disk)
11    rows processed
```

Which three statements are true regarding the execution plan?

- A. The view operator collects all rows from a query block before they can be processed but higher operations in the plan.
- B. The in-line query in the select list is processed as a view and then joined.
- C. The optimizer pushes the equality predicate into the view to satisfy the join condition.
- D. The optimizer chooses sort-merge join because sorting is required for the join equality predicate.
- E. The optimizer chooses sort-merge join as a join method because an equality predicate is used for joining the tables.

Answer: B,C,E

14. In Your Database, The Cursor_Sharing Parameter is set to EXACT. In the Employees table, the data

is significantly skewed in the DEPTNO column. The value 10 is found in 97% of rows.

Examine the following command and out put.

```
SQL> VARIABLE dno NUMBER
```

```
SQL> EXEC : dno := 10
```

```
SQL> SELECT /*ACS_1*/ count (*), max (empno)
        FROM employees
        WHERE deptno = :dno;
```

```
SQL> SELECT CHILD_NUMBER, EXECUTIONS, BUFFER_GETS, IS_BIND_SENSITIVE, IS_BIND_AWARE, IS_SHAREABLE
        FROM v$sql
        WHERE SQL_TEXT LIKE 'SELECT /*ACES_%';
```

CHILD_NUMBER	EXECUTIONS	BUFFER_GETS	IS_BIND_SEN	IS_BIND_AWA	IS_SHAREABLE
0	2	1010	Y	N	Y

Which three statements are correct?

- A. The DEPTNO column will become bind aware once histogram statistics are collected.
- B. The value for the bind variable will be considered by the optimizer to determine the execution plan.
- C. The same execution plan will always be used irrespective of the bind variable value.
- D. The instance collects statistics and based on the pattern of executions creates a histogram on the column containing the bind value.
- E. Bind peeking will take place only for the first execution of the statement and subsequent execution will use the same plan.

Answer: A,B,D

15. You created a SQL Tuning Set (STS) containing resource-intensive SQL statements. You plan to run the SQL Tuning Advisor.

Which two types of recommendations can be provided by the SQL Tuning Advisor?

- A. Semantic restructuring for each SQL statement
- B. Gathering missing or stale statistics at the schema level for the entire workload
- C. Creating a materialized view to benefit from query rewrite for the entire workload
- D. Gathering missing or stale statistics for objects used by the statements.
- E. Creating a partition table to benefit from partition pruning for each statement

Answer: A,D

16. When would bind peeking be done for queries that vary only in values used in the WHERE clause?

- A. When the column used in the WHERE clause has evenly distributed data and histogram exists on that column.
- B. When the column used in the WHERE clause has evenly distributed data and index exists on that column.
- C. When the column used in the WHERE clause has non uniform distribution of data, uses a bind variable, and no histogram exists for the column.
- D. When the column used in the WHERE clause has non uniform distribution of data and histogram exists for the column.

Answer: B

17.Which type of SQL statement would be selected for tuning by the automatic SQL framework?

- A. Serial queries that are among the costliest in any or all of the four categories: the past week, any day in the past week, any hour in the past week, or single response, and have the potential for improvement
- B. Serial queries that have been tuned within the last 30days and have been SQL profiled by the SQL tuning Advisor.
- C. Serial and parallel queries that top the AWR Top SQL in the past week only and have been SQL profiled by the SQL Tuning Advisor.
- D. Serial queries that top the AWR Top SQL in the past week only and whose poor performance can be traced to concurrency issues.
- E. Serial and parallel queries that are among the costliest in any or all of the four categories: the past week, and day in the past week, any hour in the past week, or a single response, and that can benefit from access method changes.

Answer: E

18.You instance has these parameter settings:

```
PARALLEL_DEGREE_POLICY=AUTO
PARALLEL_SERVERS_TARGET=64
PARALLEL_MIN_MINPERCENT=25
PARALLEL_MAX_SERVERS=128
PARALLEL_MIN_SERVERS=0
PARALLEL_MIN_TIME_THRESHOLD=10
PARALLEL_DEGREE_LIMIT=8
```

Which three statements are true about these settings if no hints are used in a SQL statement?

- A. A statement estimated for more than 10 seconds always has its degree of parallelism computed automatically.
- B. A statement with a computed degree of parallelism greater than 8 will be queued for a maximum of 10 seconds.
- C. A statement that executes for more than 10 seconds always has its degree of parallelism computed automatically.
- D. A statement with a computed degree of parallelism greater than 8 will raise an error.
- E. A statement with any computed degree of parallelism will be queued if the number of busy parallel execution processes exceeds 64.
- F. A statement with a computed degree of parallelism of 20 will be queued if the number of available parallel execution processes is less 5.

Answer: C,E,F

19.Exhibit

```

SQL SELECT id "id", parent_id, position "pos"
lpad(' ', "level")||operation||decode (id, 0, cost||POSITION" operation),
Operations "option" object_name "object", object_node "table_queue",
Other_tag parallel oper type, distribution "row dist", other "slave SQL"
FROM plan_table
Connect by prior id=parent_id START WITH id=0
ORDER By id;

```

Id	par	pos	operations	option	object
---	---	---	-----	-----	-----
0		4	SELECT STATEMENT cost=4		
1	0	1	HASH	GROUP BY	
2	1	1	NESTED LOOPS		
3	2	1	TABLE ACCESS	FULL	DEPARTMENTS
4	2	2	INDEX	RANGE SCAN	EMP_DEPARTMENT_IX

Examine the following SQL statement:

```

SQL> EXPLAIN PLAN FOR
SELECT department_name, count (*)
FROM hr. employees e, hr.departments d
WHERE e.department_id=d.department_id
Group by.ddepartment_name;

```

Examine the exhibit to view the execution plan.

Which statement is true about the execution plan?

- A. The EXPLAIN PLAN generates the execution plan and stores it in c\$SQL_PLAN after executing the query. Subsequent executions will use the same plan.
- B. The EXPLAIN PLAN generates the execution plan and stores it in PLAN_TABLE without executing the query. Subsequent executions will always use the same plan.
- C. The row with the ID 3 is the first step executed in the execution plan.
- D. The row with the ID 0 is the first step executed in the execution plan.
- E. The rows with the ID 3 and 4 are executed simultaneously.

Answer: E

20.Which two types of SQL statements will benefit from dynamic sampling?

- A. SQL statements that are executed parallel
- B. SQL statement that use a complex predicate expression when extended statistics are not available.
- C. SQL statements that are resource-intensive and have the current statistics
- D. SQL statements with highly selective filters on column that has missing index statistics
- E. Short-running SQL statements

Answer: B,D