Exam Questions 1Z0-064

Oracle Database 12c: Performance Management and Tuning

https://www.2passeasy.com/dumps/1Z0-064/

NEW QUESTION 1
Examine the Load Profile section of an AWR report:

	Per Second	Per Transaction	Per Exec	Per Call
DB Time(s):	2.0	0.9	0.02	0.02
DB CPU(s):	0.5	0.2	0.01	0.01
Redo size(bytes):	25,972.2	12,131.8		
Logical reads (blocks):	9,444.6	4,411.6		
Block changes:	144.7	67.6		
Physical reads (blocks):	8,671.9	4,050.7		
Physical writes (blocks):	2,641.5	1,233.9		
User calls:	83.9	39.2		
Parses (SQL):	30.7	14.3		
Hard parses(SQL):	0.4	0.2		
SQL Work Area (MB)	4.6	2.1		
Logons:	2.5	1.2		
Executes (SQL):	88.6	41.4		
Rollbacks:	0.0	0.0		
Transactions:	2.1			

Which two inferences can you derive from the details in this section? (Choose two.)

A. The values for Redo size and Block changes imply that only updates were performed by transactions.
B. The values for Parses (SQL) and Hard parses (SQL) imply that cursor sharing occurred quite often.
C. The values for DB Time and DB CPU imply that the database had a high proportion of idle time during the specified snapshot interval.
D. The values for SQL Work Area and User calls imply that only sort-based operations were performed.
E. The values for Logical reads and Physical reads imply that the number of disk reads per second was less than the total number of DB block reads and consistent gets per second.

Answer: BD

NEW QUESTION 2
Examine the partial TOP 10 Foreground Events by Total Wait Time section of an AWR report:

Top 10 Foreground Events by Total Wait Time

Event	Waits	Time (s)	Avg wait (ms)	%Total Call Time	Wait Class
enq: TX - allocate ITL entry	9,799	28,698	2929	32.9	Configurat
db file sequential read	4,827,509	25,964	5	29.7	User I/O
read by other session	2,998,307	18,118	6	20.7	User I/O
CPU time		6,872		7.9	
direct path read	222,425	4,782	21	5.5	User I/O

What should you examine to diagnose the cause of the top three wait events? (Choose the best answer.)

A. the V$ACTIVE_SESSION_HISTORY view
B. the Time Model Statistics section of the AWR report
C. the SQL statements based on elapsed time from the AWR report
D. the Latch Activity section
E. the Segment Statistics section of the AWR report

Answer: B

NEW QUESTION 3
You are administering a database that supports an OLTP workload. Users complain about the degraded response time of a query. You want to gather new statistics for objects accessed by the query and test query performance with the new statistics without affecting other sessions connected to the instance.
The STALE_PERCENT statistic preference is set to a default value and the STATISTICS_LEVEL parameter is set to TYPICAL.
Which two actions would you take to accomplish the task? (Choose two.)

A. Set the STALE_PERCENT statistic preference to a higher value than the default, and then gather statistics.
B. Set the STATISTICS_LEVEL parameter to ALL for the instance.
C. Set the INCREMENTAL preference to TRUE, and then gather statistics.
D. Set the OPTIMIZER_USE_PENDING_STATISTICS parameter to TRUE for the session in which you want to test the query.
E. Set the PUBLISH statistic preference to FALSE, and then gather statistics.
F. Set the NO_INVALIDATE statistic preference to TRUE, and then gather statistics.

Answer: BE

NEW QUESTION 4
Which two are prerequisites for running the I/O calibration tool? (Choose two.)

A. The database must be in MOUNT state.
B. The database should be opened in restricted mode.
C. For determining latency time, the STATISTICS_LEVEL parameter must be set to TYPICAL or ALL.
D. The disks to be tested must be configured to use asynchronous I/O for data files.
E. The database instance must be started using an SPFILE.

Answer: CD

NEW QUESTION 5
In the CUSTOMERS table, the values in the CUST_STATE column are dependent on the values in the COUNTRY_ID column. You want to make the optimizer aware of this dependency when these columns are used together in WHERE clause predicates that contain equalities or in-lists.
Which two methods achieve this? (Choose two.)

A. gathering statistics on the CUSTOMERS table and its dependent objects, and then locking the statistics
B. using SQL plan directives to generate an optimal plan
C. setting the dynamic statistics level to 4 and setting the OPTIMIZER_USE_PENDING_STATISTICS initialization parameter to true
D. creating column group statistics, regathering statistics, and ensuring that histograms exist on both these columns

Answer: AD

NEW QUESTION 6
Your database supports an OLTP workload during the day and batch processing at night. You want to monitor performance metrics to detect any degradation of performance in both types of workloads over a time period of 30 days.
Examine this list of possible steps:
1. Create a fixed baseline.
2. Create a baseline template.
3. Create a new moving window baseline.
4. Increase the retention period default value to 30 days.
5. Increase the size of the existing moving window baseline to 30 days.
6. Create warning and critical alerts for the relevant metrics.
7. Enable adaptive thresholds to detect the workload patterns and specify a high- significance-level threshold type.
8. Enable adaptive thresholds to detect the workload patterns and set different threshold values as a percentage of the maximum value.
Which option represents the required steps in the correct order? (Choose the best answer.)

A. 5, 7
B. 2, 4, 3
C. 3, 4, 8
D. 4, 5, 7
E. 5, 1, 6, 8

Answer: E

NEW QUESTION 7
You have been asked to assess if using column store compression (previously known as hybrid columnar compression or HCC) would help improve the performance of queries on some large tables.
Which three aspects should you consider before you choose this compression method? (Choose three.)

A. Check whether direct path load operations are used to insert rows in the table.
B. Check whether the table is frequently queried using full table scans as column store compression only minimizes I/O during full table scans.
C. Check whether the table is frequently updated because it will have overhead for insert and update operations.
D. Check whether the table has LOB columns as it will minimize I/O for the queries.
E. Check whether the table blocks are sparsely populated as this will defragment the blocks.

Answer: ABD

NEW QUESTION 8
Examine the parameters:

NAME	TYPE	VALUE
parallel_degree_policy	string	MANUAL
workarea_size_policy	string	AUTO
sort_area_size	integer	65536
memory_max_target	big integer	0
memory_target	big integer	0
pga_aggregate_target	big integer	256M
sga_target	big integer	1G

Your database supports a mixed workload and users have dedicated server connections. Users complain about the increased response time of a few queries that are performing large sort operations. On investigation, you notice an increase in the number of multipass work area executions and high number of direct path write wait events.
Which two actions could improve the performance? (Choose two.)

A. increasing the value of the SORT_AREA_SIZE parameter
B. increasing the value of the PGA_AGGREGATE_TARGET parameter
C. enabling Automatic Memory Management for the instance
D. increasing the size of the default temporary tablespace
E. using parallel hint in queries performing large sort operations
F. enabling Automatic Shared Memory Management for the instance

Answer: AF

NEW QUESTION 9
You recently joined a new team administering a database.
You notice that full table scans are performing poorly compared with full table scans on the databases you administered in a previous job.
You decide that performance problems are caused by a misconfiguration of factors affecting full table scans.
Which three factors should you investigate to determine the cause of the poorly performing Full Table Scans (FTS)? (Choose three.)

A. value of DB_FILE_MULTIBLOCK_READ_COUNT
B. storing query results in the result cache
C. setting of the DISK_ASYNC_IO parameter to TRUE
D. setting of the OPTIMIZER_MODE parameter to ALL_ROWS
E. use of parallel queries
F. block size of the tablespaces in which the tables being scanned are stored
G. value of the OPTIMIZER_DYNAMIC_SAMPLING parameter

Answer: ABC

NEW QUESTION 10
Examine the structure of the EMPLOYEES table.

```
SQL> desc employees

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

EMPLOYEE_ID is the primary key. No other indexes exist on this table. View the Exhibit to examine the commands and their output.

```
SQL> select department_id, count(department_id) from employees group by
department_id order by 2;

DEPARTMENT_ID  COUNT(DEPARTMENT_ID)
-------------  --------------------
           40                     1
           10                     1
           70                     1
           20                     2
          110                     2
           90                     3
           60                     5
           30                     6
          100                     6
           80                    34
           50                    45
11 rows selected.

SQL> var dept_id number
SQL> exec :dept_id := 50
SQL> select count(*) from employees where department_id= :dept_id;
   COUNT(*)
----------
        45
SQL> /
   COUNT(*)
----------
        45

SQL> SELECT CHILD_NUMBER, IS_BIND_SENSITIVE AS "BIND_SENSI", IS_BIND_AWARE AS
"BIND_AWARE", IS_SHAREABLE AS "BIND_SHARE" FROM V$SQL
WHERE SQL_TEXT LIKE 'select count(*) from emp%';

CHILD_NUMBER   BIND_SENSI   BIND_AWARE   BIND_SHARE
------------   ----------   ----------   ----------
           0   N            N            Y
```

Which two actions should you perform to make the cursor bind aware? (Choose two.)

A. Create a histogram on the DEPARTMENT_ID column.
B. Change the default CURSOR_SHARING value to FORCE.
C. Execute the query with the same DEPARTMENT_ID value multiple times.
D. Create an index on the DEPARTMENT_ID column.
E. Gather statistics for the index.
F. Regather statistics on the table.

Answer: CD

NEW QUESTION 10
For your database some users complain about not being able to execute transactions. Upon investigation, you find that the problem is caused by some users performing long- running transactions that consume huge amounts of space in the UNDO tablespace.
You want to control the usage of the UNDO tablespace only for these user sessions. How would you avoid the issue from repeating in future? (Choose the best answer.)

A. Create a profile for the users with the LOGICAL_READS_PER_SESSION and LOGICAL_READS_PER_CALL limits defined.
B. Create external roles to restrict the usage of the UNDO tablespace and assign them to the users.
C. Set the threshold for UNDO tablespace usage for the users.
D. Implement a Database Resource Manager plan by mapping the users to a resource consumer group with limits defined for UNDO tablespace usage.

Answer: D

NEW QUESTION 14
Which three methods can you use to create a pre-change SQL trial to capture performance data by using SQL Performance Analyzer? (Choose three.)

A. executing SQL statements in an SQL Tuning Set (STS) on a test database by using database links to the production database.
B. generating only execution plans on a test database without actually running SQL statements.
C. generating an execution plan and statistics for selective SQL statements captured in an STS
D. loading performance data and execution plans from an STS.
E. generating both execution plans and statistics for each SQL statement in an STS by actually running the SQL statements on a test database.

Answer: BDE

Explanation: Reference: https://docs.oracle.com/cd/E11882_01/server.112/e41481/spa_pre_change.htm#RATUG1 81

NEW QUESTION 16

Which two statements are true about DB time in V$SYS_TIME_MODEL? (Choose two.)

A. DB time cannot exceed the total elapsed time (walk clock time) since the database instance started.
B. DB time cannot exceed the maximum number of concurrent sessions multiplied by the actual elapsed time for each session.
C. DB time includes the time spent on client processes and background processes.
D. Reducing DB time allows a database instance to support more user requests by using the same resources.
E. DB time is always greater than or equal to the DB CPU time.

Answer: DE

NEW QUESTION 21
You execute this query twice in a session:

```
SQL>select product_name
from order_items o, product_information p
where o.unit_price = 15 and quantity > 1
and p.product_id = o.product_id;
```

Then you query V$SQL_SHARED_CURSOR for details about child cursors as shown.

```
SQL>select c.child_number, c.use_feedback_stats from v$sql_shared_cursor c
where c.sql_id = 'an4zdfz0h7513';

CHILD_NUMBER   USE_FEEDBACK_STATS
------------   ------------------
           0   Y
           1   N
```

Which two statements are true? (Choose two.)

A. No statistics were collected during the first execution of the query.
B. A subsequent execution of the query in this session is likely to undergo a soft parse.
C. The second execution of the query was hard parsed because the estimated cardinality was inaccurate.
D. A subsequent execution of the query in this session will undergo a hard parse.
E. The second execution of the query was hard parsed because extended statistics were collected after the first execution of the query.

Answer: BC

NEW QUESTION 25
You are administering a database that supports an OLTP workload. The CURSOR_SHARING parameter is set to EXACT for the instance. The performance of queries issued by one of the modules has degraded. The queries executed by the module are almost identical in syntax. To investigate, you analyze the latest AWR report and find a large number of latch:shared pool wait events and also a high percentage of the hard parse elapsed time.
Which two can be reasons for this? (Choose two.)

A. The I/O performance is slow.
B. Bind variables are not used for similar queries, causing hard parses.
C. Repeated access to a small number of blocks.
D. Excessive time is spent on finding cached cursors in the library cache.
E. The CURSOR_SHARING parameter is set to EXACT, which does not allow similar queries to share a cursor.

Answer: BC

NEW QUESTION 30
Which three statements are true about using Real-Time Database Operations? (Choose three.)

A. The STATISTICS_LEVEL initialization parameter must be set to ALL to enable automatic SQL monitoring for all long-running queries.
B. The CONTROL_MANAGEMENT_PACK_ACCESS initialization parameter must be set to DIAGNOSTIC+TUNUNG to use Real-Time Database Operations.
C. The STATISTICS_LEVEL initialization parameter can be set to TYPICAL or ALL to enable Real-Time Database Operations.
D. Real-Time Database Operations can be enabled only at the system level.
E. Real-Time Database Operations can be created by using the DBMS_MONITOR or DBMS_SESSION packages.
F. Database operation monitoring starts automatically when a database operation consumes at least five seconds of the CPU or I/O time in a single execution.

Answer: BCF

NEW QUESTION 33
Which two statements are true about Compare Period ADDM? (Choose two.)

A. It is automatically invoked whenever the AWR Compare Period report is invoked.
B. It is automatically invoked whenever ADDM is run by default.
C. It verifies if there is any change in the workload or average resource consumption by the SQL executed during the two specified time periods, to ensure 100% accuracy.
D. It can be used to create a comparison report between the Database Replay workload capture report and the replay report.

Answer: CD

NEW QUESTION 37
Identify two effects of the DB_FILE_MULTIBLOCK_READ_COUNT parameter on the optimizer. (Choose two.)

A. Decreasing the value of DB_FILE_MULTIBLOCK_READ_COUNT from the default increases the cost of index probes for DSS workloads.
B. A full table scan can become cheaper than index scans if the database instance has a high enough DB_FILE MULTIBLOCK_READ_COUNT for both OLTP and DSS workloads.
C. Increasing the value of DB_FILE_MULTIBLOCK_READ_COUNT within OS limits lowers the costing of an index probe that is done in conjunction with a nested loop for OLTP workloads.
D. In DSS workloads where full table scans may run in parallel and bypass the buffer cache, decreasing the value of DB_FILE_MULTIBLOCK_READ_COUNT from the default increases the cost of full table scans.
E. Increasing the value of DB_FILE_MULTIBLOCK_READ_COUNT within OS limits lowers the cost of full table scans and can result in the optimizer choosing a full table scan over an index scan for both OLTP and DSS workloads.

Answer: BE

NEW QUESTION 39
Examine this list of possible tasks:
1. Ensure that STATISTICS_LEVEL is set to TYPICAL or ALL.
2. Ensure that TIMED_STATISTICS is set to TRUE.
3. Set MAX_DUMP_FILE_SIZE to UNLIMITED and DIAGNOSTIC_DEST to an appropriate destination.
4. Ensure that SQL_TRACE is set to TRUE.
5. Enable tracing at the database instance level by using the DBMS_MONITOR.DATABASE_TRACE_ENABLE procedure.
6. Enable tracing in the required session by using the DBMS_SESSION.SET_SQL_TRACE procedure.
7. Run TKPROF with the EXPLAIN parameter on the output trace file.
8. Run the trcsess utility on the output trace files, and then run TKPROF on the output of the trcsess utility.
Select the minimum tasks to perform, in the correct order, to generate both a formatted trace file with timing information and an explain plan for each SQL statement for all sessions. (Choose the best answer.)

A. 1, 2, 5, 8
B. 1, 3, 6, 7
C. 2, 4, 5, 8
D. 1, 3, 4, 5, 6, 7
E. 1, 2, 4, 8

Answer: C

NEW QUESTION 40
Examine the partial TKPROF output for an SQL statement:

```
SQL> SELECT city_id
     FROM city_names
     WHERE code = 'DLR'?

call       count     cpu    elapsed       disk  query    current   rows
-------  ------  ------    -------     -----  -----    -------  -----
Parse          1    0.06       0.10         0      0          0      0
Execute        1    0.02       0.02         0      0          0      0
Fetch          1    0.23       0.30        31     31          3      1

Misses in library cache during parse: 0
Parsing user id: 02 (USER2)

Rows      Execution Plan
------    ---------------------------------------------------
     0    SELECT STATEMENT
  2340      TABLE ACCESS (BY ROWID) OF 'CITY_NAMES'
     0        INDEX (RANGE SCAN) OF 'CITY_NAMES_NAME' (NON-UNIQUE)
```

Which two inferences can definitely be made from this output? (Choose two.)

A. Array fetch operations were not performed for this query.
B. No hard parse was performed for this query.
C. The number of logical I/Os is almost equal to the number of physical I/Os.
D. Another transaction held a shared lock on the table, thereby causing a significant delay.

Answer: BD

NEW QUESTION 41
Your database supports multiple applications. The applications run on the middle tier and use connection pooling for connecting to the database.
You notice that the sessions created by the applications are competing for resources. You want to statistically measure the workload and set priorities.
What action must you perform to achieve this? (Choose the best answer.)

A. Create services for the applications and set a relative priority by assigning them to application users and using the DBMS_MONITOR.SERV_MOD_ACT_TRACE_ENABLE procedure to trace the services.
B. Create services for the applications and set a relative priority by assigning them to application users and using the DBMS_MONITOR.SESSION_TRACE_ENABLE procedure to trace the services.
C. Create services for the applications and set the relative priority of services within an instance by mapping the services directly to consumer groups.
D. Create services for the applications and set a relative priority by assigning them to application users.

Answer: A

NEW QUESTION 42
You want to capture AWR data to monitor performance variation every Monday between 9:00 AM and 12:00 PM for three months and automatically remove the older AWR data every fortnight.
How would you achieve this? (Choose the best answer.)

A. Create AWR baselines.
B. Create SQL plan baselines.
C. Create repeating baseline templates.
D. Create database services and make sure that user connections use them to connect to the database instance.
E. Create a single baseline template.

Answer: D

NEW QUESTION 43
Examine the query and its output:

```
SQL>select sid,state,wait_time/100 "WAIT TIME IN SECONDS", event from v$session where
username='HR';
```

Output:

```
SID         STATE               WAIT TIME IN SECONDS     EVENT
-------     -----------         --------------------     -------------
2832        WAITED KNOWN TIME      2029                  rdbms ipc message
3346        WAITING                   0                  enq: TX - row lock
contention
4208        WAITING                   0                  SQL*Net message from client
```

Which two statements are true? (Choose two.)

A. Session 2832 had to wait 2029 seconds for a message to arrive because of a network bottleneck.
B. Session 4208 is either idle or experiencing poor response time due to a network or resource bottleneck on the client process.
C. Session 3346 is in wait state because it wants to lock a row in a block in which other sessions have already locked rows, and there is no free ITL slot available in this block.
D. Session 3346 is in wait state because either it is waiting to update a row that is locked by another session or another session is trying to insert the same key value in a UNIQUE index.
E. Session 4208 is definitely idle and should be killed to free network resources.

Answer: AD

NEW QUESTION 47
Examine the query and its output:

```
SQL> SELECT sid, seq#, event, p1text, p1, p2text, p2, p3text, p3, wait_time,
seconds_in_wait, state FROM v$session_wait WHERE sid = 24;

SID  SEQ#  EVENT                    P1TEXT   P1   P2TEXT   P2    P3TEXT    P3   WAIT_TIME
---  ----  ---------                ------ ----  ------   ---   --------  ---  ----------
24    104  db file scattered read   file#    12  block#   1221  blocks     8        -1
```

Which two inferences can be definitely derived from this output? (Choose two.)

A. The db file scattered read event has occurred 104 times in this session for file# 12.
B. The session has completed performing a full table scan.
C. The SQL statements in this session are performing excessive disk reads.
D. The multiblock factor is 8 for this I/O but it could vary for the other I/O events.

Answer: AC

NEW QUESTION 48
You are administering a database that supports a DSS workload. Automatic Shared Memory Management is enabled for the database instance. Users issue queries to perform large soft operations and complain about degraded performance of the queries. On investigation, you notice that the queries are performing multipass work area executions and the I/O contention on one of the temporary tablespaces is very high.
Which two can be possible resolutions for this issue? (Choose two.)

A. Increase the size of the large pool.
B. Increase the value of the PGA_AGGREGATE_TARGET parameter.
C. Create a temporary tablespace group and assign it to users.
D. Increase the value of the PGA_AGGREGATE_LIMIT parameter.
E. Create another temporary tablespace and assign it to users.
F. Enable temporary undo.

Answer: CD

NEW QUESTION 53
You are administering a database that supports a mixed workload. Many applications are running on the middle tier that use connection pools to connect to the database instance. Application users perform OLTP operations during the day and another application performs batch job operations at night. You want to measure and prioritize the two workloads.
Which action would you take to achieve this? (Choose the best answer.)

A. Create database services for the applications, assign individual sessions created by the applications to consumer groups, and then set a priority.
B. Assign profiles to users running the batch operations and make sure that a priority is set for resource limits in profiles.
C. Create database services for the applications and assign different profiles to the sessions to set a relative priority for resource usage.
D. Create database services for the applications, create a job class associated with the service, batch the jobs, and then create jobs by using the job class.

Answer: C

NEW QUESTION 55
Which two statements are true about ADDM? (Choose two.)

A. It analyzes the performance of a database instance based on the time period covered by the most recent AWR snapshot, and generates recommendations based on hard-coded criteria.
B. It can analyze performance issues that occurred in past events provided they fall within the AWR retention period.
C. ADDM resource utilization and cost of analysis depends on the actual load on the database and the number of performance problems analyzed.
D. It first identifies the performance symptoms, and then refines them to reach the root cause with the singular aim of reducing the DB CPU metric.
E. It documents only those components and wait classes that are significantly impacting the performance of the database.

Answer: AB

NEW QUESTION 57
Examine the initialization parameters set for a database instance:

NAME	TYPE	VALUE
dbwr_io_slaves	integer	0
db_writer_processes	integer	1
filesystemio_options	string	NONE
disk_asynch_io	boolean	TRUE

The database supports an OLTP workload. Applications connect to the instance using shared server connections and perform small, random I/Os. All the data files are on the same disk. You notice free buffer wait events for sessions in the database instance.
To solve the problem, you increase the size of the buffer cache. But after some time, you notice sessions waiting again on free buffer waits.
What will you recommend to alleviate the issue? (Choose the best answer.)

A. Run the I/O calibration tool.
B. Configure the database instance to make asynchronous I/O available to DBWR.
C. Spread the data files over multiple disks, controllers, and I/O buses to ensure that there are no hotspots in the I/O subsystem.
D. Configure dedicated server connections for the applications.

Answer: B

NEW QUESTION 61
For which three problem categories does Automatic Database Diagnostic Monitor (ADDM) provide analysis and recommendations by default? (Choose three.)

A. for network stack-related bandwidth contention
B. for concurrency issues because of buffer busy problems
C. for high-load PL/SQL execution and compilation, and high-load Java usage
D. for application-level lock contention.

Answer: BCD

NEW QUESTION 63
Your database supports an OLTP system.
Examine the parameter values configured in your database:
sga_max_size = 480M sga_target = 480M pga_aggregate_target = 160M
The CUSTOMERS table contains 8,000 rows. The CUST_ID column is the primary key and the COUNTRY_ID column contains only three possible values: 1111, 2222, and 3333.
You execute the commands:
SQL> EXECUTE DBMS_STATS.GATHER_TABLE_STATS('SH','CUSTOMERS');
PL/SQL procedure successfully completed.
SQL> CREATE INDEX COUNTRY_IDX ON CUSTOMERS (COUNTRY_ID);
Index created.
You then perform a series of INSERT, UPDATE, and DELETE operations on the table. View the Exhibit to examine the query and its execution plan.

```
SQL> SELECT COUNT(*)
 FROM CUSTOMERS
 WHERE COUNTRY_ID = 2222;

 COUNT(*)
----------
       150

SQL> select * from table(dbms_xplan.display_cursor(null,null,'basic rows'));

PLAN_TABLE_OUTPUT
------------------------------------------------------------------------------
EXPLAINED SQL STATEMENT:
------------------------
SELECT COUNT(*) FROM CUSTOMERS WHERE COUNTRY_ID = 2222;

Plan hash value: 568322376
-------------------------------------------------
|  ID   | Operation             | Name      | Rows  |
-------------------------------------------------
|    0  | SELECT STATEMENT      |           |       |
|    1  |  SORT AGGREGATE       |           |     1 |
|    2  |   TABLE ACCESS FULL   | CUSTOMERS |  8000 |
-------------------------------------------------
```

Which two options can improve the performance of the query without significantly slowing down the DML operations? (Choose two.)

A. creating a bitmap index on the COUNTRY_ID column
B. regathering statistics on the CUSTOMERS table
C. gathering statistics on the COUNTRY_IDX index
D. creating a histogram on the COUNTRY_ID column
E. increasing the size of the PGA
F. creating an SQL profile
G. creating a KEEP cache

Answer: AD

NEW QUESTION 67
Which two actions should you take to monitor the throughput generated by the modules of an application? (Choose two.)

A. Use the Resource Manager.
B. Enable SQL Trace at the session level.
C. Create a service.
D. Use a dedicated server configuration.
E. Use the DBMS_APPLICATION_INFO package to define the current module and action so that they appear in V$SESSION.

Answer: BE

NEW QUESTION 70
Examine the parameter values configured in your database:
sga_max_size = 480M sga_target = 480M pga_aggregate_target = 160M
The CUSTOMERS table contains 8,000 rows. The CUST_ID column is the primary key and the COUNTRY_ID column contains only three possible values: 1111, 2222, and 3333.
You execute the commands:
SQL> EXECUTE DBMS_STATS.GATHER_TABLE_STATS('SH','CUSTOMERS');
PL/SQL procedure successfully completed.
SQL> CREATE INDEX COUNTRY_IDX ON CUSTOMERS (COUNTRY_ID);
Index created.
You then perform several INSERT, UPDATE, and DELETE operations, significantly altering the data in the table.
View the Exhibit to examine the query and its execution plan.

```
SQL> SELECT COUNT(*)
   FROM CUSTOMERS
   WHERE COUNTRY_ID = 2222;

   COUNT(*)
------------
          150

SQL> select * from table(dbms_xplan.display_cursor(null,null,'basic rows'));

PLAN_TABLE_OUTPUT
------------------------------------------------------------------------------
EXPLAINED SQL STATEMENT:
------------------------
SELECT COUNT(*) FROM CUSTOMERS WHERE COUNTRY_ID = 2222;

Plan hash value: 568322376
-----------------------------------------------------
|  Id   | Operation           |    Name      | Rows   |
-----------------------------------------------------
|    0  | SELECT STATEMENT    |              |        |
|    1  |  SORT AGGREGATE     |              |      1 |
|    2  |   TABLE ACCESS FULL | CUSTOMERS    |   8000 |
-----------------------------------------------------
```

Which three options would improve the performance of the query? (Choose three.)

A. creating a bitmap index on the COUNTRY_ID column
B. regathering statistics on the CUSTOMERS table
C. creating a histogram on the COUNTRY_ID column
D. increasing the size of the PGA
E. creating an SQL profile
F. creating a KEEP cache

Answer: ABF

NEW QUESTION 71
You are administering a database that supports multiple applications, which make dedicated connections to the database instance by using different services.
You execute the command to enable tracing of the ORCL1 service:

```
SQL> EXECUTE DBMS_MONITOR.SERV_MOD_ACT_TRACE_ENABLE (service_name => 'ORCL1', WAITS =>
TRUE, BINDS => NULL, instance_name => 'ORCL',plan_stat => NULL);
```

Which two statements are true? (Choose two.)

A. A single trace file is generated for all sessions mapped to the ORCL1 service.
B. SQL trace is enabled for all modules and actions for sessions mapped to the ORCL1 service.
C. An SQL trace file is generated for each session that maps to the ORCL1 service.
D. An SQL trace file is generated for each of the modules using the ORCL1 service.
E. SQL trace is not enabled for the service because a module name is not specified.

Answer: AC

NEW QUESTION 75
Which two actions can reduce library cache latch contention for an OLTP application that repeatedly executes queries containing a mix of literals and bind variables? (Choose two.)

A. setting the OPEN_CURSORS parameter to hold a sufficient number of concurrently open cursors
B. coding the application such that a frequently executed statement is parsed only onceand executed repeatedly as required
C. setting the CURSOR_SHARING parameter to EXACT
D. avoiding the granting of privileges on objects associated with cursors during peak load
E. enabling Automatic Memory Management and allocating at least 50% of the available memory for SHARED_POOL_SIZE
F. configuring shared server connections

Answer: BE

Explanation: Reference: http://docs.oracle.com/cd/B28359_01/server.111/b28274/memory.htm

NEW QUESTION 77
Which two situations can lead to sparsely populated index blocks? (Choose two.)

A. Data is frequently inserted using direct path load into a table with an index.
B. Indexed columns in a table are frequently updated.
C. Values in an indexed column are inserted using monotonically incrementing sequences.
D. Bulk delete operations are performed on a table with indexes.

E. Online table move operations are performed frequently on a table with indexes.

Answer: BD

NEW QUESTION 78
.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual 1Z0-064 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the 1Z0-064 Product From:

https://www.2passeasy.com/dumps/1Z0-064/

Money Back Guarantee

1Z0-064 Practice Exam Features:

* 1Z0-064 Questions and Answers Updated Frequently

* 1Z0-064 Practice Questions Verified by Expert Senior Certified Staff

* 1Z0-064 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* 1Z0-064 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year