

1Z0-064^{Q&As}

Oracle Database 12c: Performance Management and Tuning

Pass Oracle 1Z0-064 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.certbus.com/1z0-064.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

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

Correct Answer: AD

QUESTION 2

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.

Correct Answer: C

QUESTION 3

Examine the partial AWR report taken for a time period of 60 minutes: Which two inferences can you draw from this report? (Choose two.)

Top 10 Foreground Events by Total Wait Time

Event	Waits	Time (s)	Avg wait (ms)	%Total Call Time	Wait Class
resmgr: cpu quantum	475,956	152,859	320	75.2	Scheduler
CPU time		47,880		23.5	
db file sequential read	3,374,890	16,868	5	7.8	User I/O
db file scattered read	196,265	4,278	22	2.1	User I/O
log file sync	177,735	4,579	29	5.4	Commit
.....					
.....					
.....					

Operating System Statistics DB/Inst: ****/**** Snaps: 56708/56709

Statistic	Total
.....	
BUSY_TIME	5,707,832
IDLE_TIME	2
.....	
NUM_CPUS	32

- A. The database user calls are issuing frequent explicit commits.
- B. The CPUs are busy executing server processes and background processes for a considerable amount of CPU time.
- C. The database user calls are spending most of their time in I/O for single block reads.
- D. The database user calls are spending most of their time waiting for sessions that are in more important consumer groups.

Correct Answer: BC

QUESTION 4

You want to enable the Database Smart Flash Cache feature on your database. For this purpose, you configure two flash devices: /dev/sdj and /dev/sdk, with sizes of 64 GB and 32 GB, respectively.

Examine the parameters set for your database instance:

NAME	TYPE	VALUE
db_flash_cache_file	string	/dev/sdj, /dev/sdk
db_flash_cache_size	big integer	0
memory_max_target	big integer	16G
memory_target	big integer	16G
sga_target	big integer	0

How would you make sure that the database uses these devices for Database Smart Flash Cache?

- A. by disabling Automatic Memory Management, setting SGA_TARGET to 96G, and setting the required memory that you want to set for SGA
- B. by setting the value of the DB_FLASH_CACHE_SIZE parameter to 8G
- C. by configuring the DB_FLASH_CACHE_SIZE parameter and making sure that all flash devices are of the same size
- D. by setting the DB_FLASH_CACHE_SIZE parameter to 86G and the MEMORY_TARGET parameter to 112G
- E. by setting the value of the DB_FLASH_CACHE_SIZE parameter to 64G, 32G

Correct Answer: E

QUESTION 5

Your database supports an OLTP workload where a large number of syntactically similar queries are executed. Examine the Instance Efficiency Percentages in the Automatic Workload Repository (AWR) report of the last hour: Which two inferences can be drawn from the report?

Instance Efficiency Percentages (Target 100%)

Buffer Nowait %:	100.00	Redo NoWait %:	100.00
Buffer Hit %:	79.49	In-memory Sort %:	100.00
Library Hit %:	63.12	Soft Parse %:	32.72
Execute to Parse %:	4.12	Latch Hit %:	99.95
Parse CPU to Parses Elapsed %:	1.98	%Non-Parse CPU:	36.94

- A. The optimizer is waiting for resources during parsing of the queries.
- B. The CPU is spending more time on pinning cursors in the library cache.
- C. The database buffer cache is undersized and is causing contention.
- D. Cursors are not getting shared in the library cache.

Correct Answer: AD

QUESTION 6

Examine the Load Profile and partial Top 10 Foreground Events by Total Wait Time sections from an AWR report.

Load Profile	Per Second	Per Transaction
DDB Time(s) :	0.3	0.15
DB CPU(s)	0.2	0.4
Redo size (bytes) :	18,680.98	4,365.06
Logical reads (blocks) :	106,671.46	24,488.48
Block changes:	109.86	25.47
Physical reads (blocks)	2.99	0.7
Physical writes (blocks)	7.97	1.86
Read IO requests:	2.9	32.3
Write IO requests:	0.4	8.2
Read IO (MB) :	0.1	0.5
Write IO (MB) :	0.0	0.2
User calls:	497.39	105.37
Parses (SQL) :	39.68	9.27
Hard parses (SQL) :	0.12	0.04
SQL Work Area (MB) :	14.56	3.66
Logons:	0.13	0.03
Executes (SQL) :	55.94	12.15
Rollbacks:	0.0	0.1
Transactions:	4.15	

Top 10 Foreground Events by Total Wait Time

Event	Waits	Total Wait Time (sec)	Wait Avg (ms)	% DB Time
CPU time		6,581		38.1
db file sequential read	19,870	185	9	3.6
SQL*Net more data from client	229,931	104	0	.8
log file sync	58,341	103	2	.7
log switch/archive	10	98	9,791	.6

Which two areas should you examine next to identify possible bottlenecks?

- A. the application code because of CPU-intensive activities
- B. the application code because user calls are performing several queries that require sorting
- C. the "SQL ordered by Gets" section of the AWR report to check for excessive logical I/O
- D. the "SQL ordered by Reads" section of the AWR report to check for excessive physical reads

Correct Answer: AC

QUESTION 7

Examine the parameters set for your database instance:

NAME	TYPE	VALUE
db_block_size	integer	8192
db_2k_cache_size	big integer	0
db_4k_cache_size	big integer	0
db_8k_cache_size	big integer	0
db_16k_cache_size	big integer	0
db_32k_cache_size	big integer	0

To investigate the slow response time of queries on the TRANS table, you gather table and execute the query:

```
SQL> SELECT chain_cnt, round(chain_cnt/num_rows*100,2) pct_chained, avg_row_len, pct_free, pct_used
FROM user_tables
WHERE table_name = 'TRANS';
```

CHAIN_CNT	PCT_CHAINED	AVG_ROW_LEN	PCT_FREE	PCT_USED
4789	100	3691	10	40

The table is stored in a tablespace with Automatic Segment Space Management (ASSM), and some rows of the TRANS table are migrated and chained.

Which two actions would you recommend to improve query response time?

- A. Reorganize the TRANS table online by using the DBMS_REDEFINITION package.
- B. Create a bigger non-standard blocksize tablespace and move the TRANS table to that tablespace.
- C. Move the TRANS table to a tablespace with manual segment space management with a lower value set for the PCTUSED attribute.
- D. Move the TRANS table to a tablespace with manual segment space management with a higher value set for the PCTFREE attribute.

Correct Answer: AD

QUESTION 8

Examine the partial Activity Over Time section of an Active Session History (ASH) report:

Slot Time (Duration)	Slot Count	Event	Event Count	% Event
14:10:50 (1.2 min)	5	control file sequential read	4	0.11
		CPU + Wait for CPU	1	0.03
14:12:00 (3.0 min)	9	CPU + Wait for CPU	5	0.14
		control file parallel write	2	0.05
		null event	1	0.03

Which two inferences are correct? (Choose two.)

- A. In the first time slot, five different sampled sessions were connected to the database instance.
- B. In the second time slot, out of the nine sampled sessions connected to the database instance, only one sampled session was idle at the time of report generation.
- C. In the first time slot, only one sampled session was using the CPU.
- D. In the second time slot, five different sampled sessions were using the CPU.
- E. In the second time slot, 0.14% of the time was spent on the CPU.

Correct Answer: AE

QUESTION 9

Examine the Load Profile section of an AWR report: Which two inferences can you derive from the details in this section? (Choose two.)

	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			

- 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.

Correct Answer: BD

QUESTION 10

You have been asked to use table compression for two large tables. Given are the details of the tables:

The TRANS_DET table:

The table is used by an OLTP application.

High volume insert and update operations are performed on the table.

The table is frequently queried using index range scans.

The TRANS_HISTORY table:

The table is used by a DSS application.

High volume bulk loads are performed on the table.

The table is used to store archival data on which large table full-table scans (FTS) are performed.

Which row store compression would you recommend for these tables with minimal overhead on performance? (Choose the best answer.)

- A. basic table compression for both the tables
- B. advanced row compression for both the tables
- C. basic table compression for the TRANS_HISTORY table and advanced row compression for the TRANS_DET table
- D. basic table compression for the TRANS_DET table and advanced row compression for the TRANS_HISTORY table
- E. warehouse compression for the TRANS_DET table and archive compression for the TRANS_HISTORY table

Correct Answer: A

QUESTION 11

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

- A. DB time cannot exceed the total elapsed time (wall 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.

Correct Answer: DE

QUESTION 12

You are administering a database that supports a mixed workload. The CURSOR_SHARING parameter is set to the default value. While analyzing the latest Automatic Workload Repository (AWR) report, you find a large number of cursor: pin S wait on X, cursor: pin X wait on S, and library cache mutex waits in the Top 10 foreground events section. Examine the Instance Efficiency Percentages section in the AWR report:

Instance Efficiency Percentages (Target 100%)

Buffer Nowait %:	100.00	Redo NoWait %:	100.00
Buffer Hit %:	99.95	In-memory Sort %:	100.00
Library Hit %:	62.17	Soft Parse %:	52.72
Execute to Parse %:	47.12	Latch Hit %:	97.95
Parse CPU to Parse Elapsed %:	53.98	% Non-Parse CPU:	70.94

Which three statements are true in this scenario? (Choose three.)

- A. Sessions are waiting for mutexes in share mode on cursors but other sessions are holding the mutexes in exclusive mode.
- B. The CPU is spending more time in finding cursors in the library cache.
- C. Cursors are not getting shared, resulting in a large number of hard parses.
- D. Sessions are waiting for mutexes in exclusive mode on cursors but other sessions are holding the mutexes in share mode.
- E. The buffers required by queries are not found in the buffer cache, thereby increasing expensive disk I/O.

Correct Answer: BDE

QUESTION 13

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 a SQL profile
- G. creating a KEEP cache

Correct Answer: AD

QUESTION 14

In which three situations does DB time always increase? (Choose three.)

- A. when the host is CPU bound for foreground processes
- B. when I/O wait time increases for foreground processes

- C. when more connections are made to a database instance
- D. when CPU consumption by background processes increases
- E. when wait time for data to be sent over a network increases

Correct Answer: BCD

QUESTION 15

Examine the Time Model Statistics section of an AWR report:

Statistic Name	Time (s)	% of DB Time
sql execute elapsed time	12,416.14	86.45
DB CPU	9,223.70	64.22
parse time elapsed	935.61	6.51
hard parse elapsed time	884.73	6.16
failed parse elapsed time	21.39	.72
PL/SQL execution elapsed time	153.51	1.07
hard parse (sharing criteria) elapsed time	25.96	0.18
connection management call elapsed time	14.00	0.10
hard parse (bind mismatch) elapsed time	4.74	0.03
PL/SQL compilation elapsed time	1.20	0.01
repeated bind elapsed time	0.22	0.00
sequence load elapsed time	0.11	0.00
DB time	14,352.96	
background elapsed time	731.00	
background cpu time	72.00	

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

- A. The available CPU resources were not utilized to their maximum capacity.
- B. All sequence numbers used during this AWR time interval were cached.
- C. A large number of connected user sessions were idle.
- D. New child cursors were created because of new bind values or usage of literal values as well as different bind types or sizes.
- E. The DB CPU time was not spent exclusively for processing SQL statements.

Correct Answer: DE