# Free Questions for 1Z0-084 by vceexamstest

## Shared by Delaney on 15-04-2024

**For More Free Questions and Preparation Resources**

**Check the Links on Last Page**

# Question 1

**Question Type:** MultipleChoice

You need to collect and aggregate statistics for the ACCTG service and PAYROLL module, and execute:

```
BEGIN
  DBMS_MONITOR.SERV_MOD_ACT_STAT_ENABLE(
    service_name => 'ACCTG' ,
    module_name  => 'PAYROLL' );
END;
```

Where do you find the output of this command?

## Options:

**A-** By viewing V$SERV_MOD_ACT_STATS

**B-** In $ORACLE_BASE/diag/rdbms/<db unique name>/<instance name>/trace

**C-** By viewing V$SERVICE_STATS

**D-** In the current working directory

## Answer:

A

## Explanation:

When you enable statistics gathering for a specific service and module using DBMS_MONITOR.SERV_MOD_ACT_STAT_ENABLE, the output is aggregated and can be viewed using the V$SERV_MOD_ACT_STATS dynamic performance view. This view contains the cumulative statistics of database activity broken down by service and module, which is exactly what you collect when executing the provided command.

B (Incorrect): While many types of trace files are located in the Diagnostic Destination directory ($ORACLE_BASE/diag), the aggregated statistics for services and modules are not written to trace files but are instead viewable through dynamic performance views.

C (Incorrect): The V$SERVICE_STATS view provides service-level statistics but does not provide the combined service/module-level breakdown.

D (Incorrect): The output of the PL/SQL block is not written to a file in the current working directory; it is stored in the data dictionary and accessible via dynamic performance views.

Oracle Database PL/SQL Packages and Types Reference: DBMS_MONITOR

Oracle Database Reference: V$SERV_MOD_ACT_STATS

# Question 2

**Question Type:** MultipleChoice

Which two Oracle Database features use database services?

## Options:

**A-** Oracle Automatic Reoptimization

**B-** Database Resource Manager

**C-** Oracle SQL Performance Management

**D-** Oracle SQL Tuning Advisor

**E-** Oracle Scheduler

## Answer:

B, E

## Explanation:

Database services in Oracle are used to manage how resources are allocated and how workloads are managed within the database.
The features that use database services are:

B (Correct): Database Resource Manager (DBRM) uses services to control resource allocation to different workloads. It ensures that resources are assigned to the most critical tasks first, based on the service associated with the workload.

E (Correct): Oracle Scheduler can also utilize database services. Jobs in Oracle Scheduler can be assigned to different services to control resource allocation and prioritization.

The other features mentioned are related to SQL performance but do not directly utilize database services in the way Resource Manager and Scheduler do:

A: Oracle Automatic Reoptimization is a feature that allows the database to automatically improve the execution plan of a SQL statement after it is executed, based on the actual performance metrics, but it does not directly use database services.

C: Oracle SQL Performance Management involves various components of SQL tuning and monitoring, but it does not use database services to operate.

D: Oracle SQL Tuning Advisor provides advice on how to tune SQL queries for better performance. While it can be used in conjunction with services for managing and analyzing workloads, it doesn't use services in its core functionality.

Oracle Database Administrator's Guide: Administering Services

Oracle Database Administrator's Guide: Managing Resources with Oracle Database Resource Manager

Oracle Database Scheduler Developer's Guide: Using the Scheduler

## Question 3

**Question Type:** **MultipleChoice**

Examine this AWR report excerpt:

```
Top 10 Foreground Events by Total Wait Time

Event                              Waits Total Wait Time (sec) Avg Wait % DB time Wait Class
DB CPU                                                 14.7K               42.3
db file sequential read             3,467,279            14K    4.04ms      40.4 User I/O
gc current grant busy                 386,743          224.4  539.50us       4.1 Cluster
log file sync                          50,598          162.3    2.02ms       2.8 Commit
gc current block 2-way                427,553           96.3  237.36us       1.7 Cluster


SQL ordered by Reads

Physical Reads Executions Reads per Exec Elapsed Time (s)  %CPU   %IO SQL Id        SQL
Module       SQL Text
       511,342     28,519          17.93         6,766.64 57.91 28.60 al7r9r9jz63k2 JDBC
Thin Client UPDATE SALES SET...
       154,496     28,530           5.42         4,556.02 68.74 20.89 4d3ffsuuh8f2q JDBC
Thin Client UPDATE SALES SET...
        95,933  1,695,105           0.06           213.16 16.18 66.99 wqsjwz37p93ug JDBC
Thin Client select COUNT(*) FROM...
```

You must reduce the impact of database I/O, without increasing buffer cache size and without modifying the SQL statements.

Which compression option satisfies this requirement?

## Options:

**A-** MN STORE COMPRESS FOR QUERY LOW

**B-** STORE COMPRESS

**C-** ROW STORE COMPRESS ADVANCED

**D-** COLUMN STORE COMPRESS FOR QUERY HIGH

## Answer:

C

## Explanation:

To reduce the impact of database I/O without increasing the size of the buffer cache and without modifying SQL statements, you can use table compression. Among the given options, ROW STORE COMPRESS ADVANCED is the most suitable form of table compression to satisfy this requirement.

Advanced row compression (ROW STORE COMPRESS ADVANCED) is designed to work well with all supported types of data, whether it's OLTP or data warehouse environments. It offers a higher level of compression than basic table compression (ROW STORE COMPRESS BASIC) without significant overhead during DML operations. This feature can help reduce the amount of I/O required to retrieve data by storing it more efficiently on disk.

A, B, D: While COLUMN STORE COMPRESS FOR QUERY HIGH and ROW STORE COMPRESS are both valid compression types, COLUMN STORE COMPRESS FOR QUERY HIGH applies to the In-Memory column store and is not available in all versions and

editions, and ROW STORE COMPRESS is less advanced than ROW STORE COMPRESS ADVANCED.

Oracle Database Concepts Guide: Table Compression

Oracle Database Performance Tuning Guide: Row Compression

# Question 4

**Question Type:** MultipleChoice

Which three statements are true about using the in Memory (IM) column store?

## Options:

**A-** It does not improve performance for queries using cached results of function evaluations on columns from the same table.

**B-** It does not improve performance for queries that use join groups on columns from different tables.

**C-** It can improve OLTP workload performance by avoiding the use of indexes.

**D-** It does not improve performance for queries using user-defined virtual column results.

**E-** It does not require all database data to fit in memory to improve query performance.

**F-** It improves performance for queries joining several tables using bloom filter joins.

## Answer:

C, E, F

## Explanation:

The Oracle In-Memory (IM) column store feature enhances the performance of databases by providing a fast columnar storage format for analytical workloads while also potentially benefiting OLTP workloads.

C (True): It can improve OLTP workload performance by providing a faster access path for full table scans and reducing the need for indexes in certain scenarios, as the In-Memory store allows for efficient in-memory scans.

E (True): The In-Memory column store does not require all database data to fit in memory. It can be used selectively for performance-critical tables or partitions, and Oracle Database will manage the population and eviction of data as needed.

F (True): In-Memory column store can significantly improve performance for queries joining several tables, especially when bloom filters are used, as they are highly efficient with the columnar format for large scans and join processing.

The other options provided are not correct in the context of the In-Memory column store:

A (False): While In-Memory column store is designed for analytical queries rather than caching results of function evaluations, it does not specifically avoid improving performance for queries using cached results of function evaluations.

B (False): In-Memory column store can improve the performance of queries that use join groups, which can be used to optimize joins on columns from different tables.

D (False): In-Memory column store can improve the performance of queries using expressions, including user-defined virtual columns, because it supports expression statistics which help in optimizing such queries.

Oracle Database In-Memory Guide: In-Memory Column Store in Oracle Database

Oracle Database In-Memory Guide: In-Memory Joins

Oracle Database In-Memory Guide: In-Memory Aggregation

# Question 5

**Question Type:** **MultipleChoice**

You must write a statement that returns the ten most recent sales. Examine this statement:

```
SELECT s.prod_id, p.prod_name, s.amount_sold, c.cust_name, c.cust_city
FROM   sales s, customers c, products p
WHERE  s.cust_id = c.cust_id AND
       s.prod_id = p.prod_id
ORDER BY s.time_id
FETCH FIRST 10 ROWS ONLY;
```

Users complain that the query executes too slowly. Examine the statement's current execution plan:

```
------------------------------------
-------------------------------------------------------------------------------------------------------------------------------------------
| Id | Operation             | Name      | Starts | E-Rows | A-Rows | A-Time      | Buffers | Reads | OMem | 1Mem | Used-Mem
-------------------------------------------------------------------------------------------------------------------------------------------
|  0 | SELECT STATEMENT      |           |      1 |        |     10 |00:00:00.32 |    3130 |     3 |      |      |
|* 1 | COUNT STOPKEY         |           |      1 |        |     10 |00:00:00.32 |    3130 |     3 |      |      |
|  2 | VIEW                  |           |      1 |   918K|     10 |00:00:00.32 |    3130 |     3 |      |      |
|* 3 | SORT ORDER BY STOPKEY |           |      1 |   918K|     10 |00:00:00.32 |    3130 |     3 | 2048 | 2048 |  2048 (0
|* 4 | HASH JOIN             |           |      1 |   918K|    918K|00:00:00.22 |    3130 |     3 | 1250K| 1250K| 1579K (0
|  5 | TABLE ACCESS FULL     | PRODUCTS  |      1 |    72 |     72 |00:00:00.01 |       3 |     3 |      |      |
|* 6 | HASH JOIN             |           |      1 |   918K|    918K|00:00:00.15 |    3126 |     0 | 4696K| 1834K| 4597K (0
|  7 | TABLE ACCESS FULL     | CUSTOMERS |      1 |  55500 |  55500 |00:00:00.01 |    1521 |     0 |      |      |
|  8 | PARTITION RANGE ALL   |           |      1 |   918K|    918K|00:00:00.10 |    1604 |     0 |      |      |
|  9 | TABLE ACCESS FULL     | SALES     |     28 |   918K|    918K|00:00:00.10 |    1604 |     0 |      |      |
-------------------------------------------------------------------------------------------------------------------------------------------
Predicate Information (identified by operation id):
---------------------------------------------------

1 - filter(ROWNUM<11)
3 - filter(ROWNUM<11)
4 - access("S"."PROD_ID"="P"."PROD_ID")
6 - access("S"."CUST_ID"="C"."CUST_ID")
```

What must you do to reduce the execution time and why?

## Options:

**A-** Create an index on SALES.TIME_ID to force the return of rows in the order specified by the ORDER BY clause.

**B-** Replace the FETCH FIRST clause with ROWNUM to enable the use of an index on SALES.

**C-** Collect a new set of statistics on PRODUCT, CUSTOMERS, and SALES because the current stats are inaccurate.

**D-** Enable Adaptive Plans so that Oracle can change the Join method as well as the Join order for this query.

**E-** Create an index on SALES.CUST_ID to force an INDEX RANGE SCAN on this index followed by a NESTED LOOP join between CUSTOMERS and SALES.

## Answer:

A

## Explanation:

The execution plan shows a full table access for the SALES table. To reduce the execution time, creating an index on SALES.TIME_ID would be beneficial as it would allow the database to quickly sort and retrieve the most recent sales without the need to perform a full table scan, which is I/O intensive and slower. By indexing TIME_ID, which is used in the ORDER BY clause, the optimizer can take advantage of the index to efficiently sort and limit the result set to the ten most recent sales.

B (Incorrect): Replacing FETCH FIRST with ROWNUM would not necessarily improve the performance unless there is an appropriate index that the optimizer can use to avoid sorting the entire result set.

C (Incorrect): There is no indication that the current statistics are inaccurate; hence, collecting new statistics may not lead to performance improvement.

D (Incorrect): While adaptive plans can provide performance benefits by allowing the optimizer to adapt the execution strategy, the main issue here is the lack of an index on the ORDER BY column.

E (Incorrect): Creating an index on SALES.CUST_ID could improve join performance but would not address the performance issue caused by the lack of an index on the ORDER BY column.

Oracle Database SQL Tuning Guide: Managing Indexes

Oracle Database SQL Tuning Guide: Using Indexes and Clusters

# Question 6

**Question Type:** MultipleChoice

A database instance is suffering poor I/O performance on two frequently accessed large tables.

No Big Table caching occurs in the database.

Examine these parameter settings:

```
PARALLEL_DEGREE_POLICY=LIMITED
DB_BIG_TABLE_CACHE_PERCENT_TARGET=10
DB_CACHE_SIZE=500M
DB_KEEP_CACHE_SIZE=0M
```

Which are two actions either one of which will allow Big Table caching to occur?

## Options:

**A-** Increasing DB_BIG_TABLE_CACHE_PERCENT_TARGET to at least 50

**B-** Setting DB_KEEP_CACHE_SIZE to at least 50M

**C-** Increasing DB_BIG_TABLE_CACHE_PERCENT_TARGET to at least 25

**D-** Increasing DB_CACHESIZE to 1 G

**E-** Setting PARALLEL_DEGREE_POLICYADAPTIVE

**F-** Setting PARALLEL_DEGREE_POLICYAUTO

## Answer:

C, D

## Explanation:

Big Table caching is a feature that allows frequently accessed large tables to be cached in memory to improve I/O performance. From the parameter settings provided, Big Table caching is not occurring because DB_BIG_TABLE_CACHE_PERCENT_TARGET is set to 10, which is the minimum threshold for enabling the feature, but the size of the cache is too small for the big tables to be effectively cached.

To enable Big Table caching, one of the following actions could be taken:

C (Correct): Increasing DB_BIG_TABLE_CACHE_PERCENT_TARGET to at least 25. This action would allocate a larger percentage of the buffer cache for storing big tables, which could allow for caching large tables and thus improve I/O performance.

D (Correct): Increasing DB_CACHE_SIZE to 1G. Since the size of the buffer cache is a determining factor for how much data can be cached, increasing this parameter would provide more memory space for big tables to be cached.

Options A, B, E, and F will not enable Big Table caching because:

A: Increasing DB_BIG_TABLE_CACHE_PERCENT_TARGET to 50 without adjusting the overall size of the cache might still not be sufficient if the DB_CACHE_SIZE is not large enough to hold the big tables.

B: Setting DB_KEEP_CACHE_SIZE to at least 50M only specifies a separate buffer pool for objects with the KEEP cache attribute and does not affect Big Table caching.

E: and F: Changing the PARALLEL_DEGREE_POLICY to ADAPTIVE or AUTO influences the behavior of parallel execution but does not directly enable or influence Big Table caching.

Oracle Database Performance Tuning Guide: Big Table Caching

Oracle Database Reference: DB_BIG_TABLE_CACHE_PERCENT_TARGET

Oracle Database Reference: DB_CACHE_SIZE

## Question 7

**Question Type:** MultipleChoice

What are the least elevated values of statistics_level and C0NTR0LJ4ANAGEMENT_PACK_ACCESS that allow the usage of Monitoring of Database Operations?

## Options:

**A-** STATISTICS_LEVEL=ALL and CONTROL_MANAGEMENT_PACK_ACCESS=DIAGOSTIC+TUNING

**B-** STATISTICS_LEVEL=BASIC and CONTROL_MANAGEMENT_PACK ACCESS=DIAGOSTIC

**C-** STATISTICS_LEVEL=TYPICAL and CONTROL_MANAGEMENT_PACK_ACCESS-DIAGOSTIC*TUNING

**D-** STATISTICS_LEVEL=TYPICAL and CONTROL_MANAGEMENT_PACK_ACCESS=DIAGOSTIC

## Answer:

A

## Explanation:

Monitoring of Database Operations requires that the STATISTICS_LEVEL parameter be set to ALL and CONTROL_MANAGEMENT_PACK_ACCESS be set to DIAGNOSTIC+TUNING. These settings enable all the advisory features and automatic tuning features within the Oracle Database, including the Automatic Workload Repository (AWR), Automatic Database Diagnostic Monitor (ADDM), and the full functionality of the SQL Tuning Advisor and SQL Access Advisor, which are components of the

Diagnostic and Tuning packs.

STATISTICS_LEVEL=ALL: This setting enables the collection of all system statistics for problem detection and self-tuning purposes.

CONTROL_MANAGEMENT_PACK_ACCESS=DIAGNOSTIC+TUNING: This grants access to both the Diagnostic Pack and the Tuning Pack, which are essential for detailed performance monitoring and tuning capabilities.

Oracle Database Reference: STATISTICS_LEVEL

Oracle Database Licensing Information User Manual: Oracle Database Management Packs

# Question 8

**Question Type:** **MultipleChoice**

This error occurred more than four hours ago in the database:

ORA-04036 PGA memory used by the instance exceeds PGA_AGGREGATE_LIMIT

You want to know which process and query were at fault.

Which two views should you use for this purpose?

## Options:

**A-** DBA_HIST_ACTIVE_SESS_HISTORY

**B-** DBA_HIST_SQLSTAT

**C-** DBA_HIST_SQLTEXT

**D-** DBA_HIST_PGASTAT

**E-** DBA_HIST_PROCESS_MEM_SUMMARY

## Answer:

A, E

## Explanation:

To investigate the cause of the ORA-04036 error, which indicates that PGA memory usage exceeds the PGA_AGGREGATE_LIMIT, the appropriate views to consult are DBA_HIST_ACTIVE_SESS_HISTORY and DBA_HIST_PROCESS_MEM_SUMMARY.

DBA_HIST_ACTIVE_SESS_HISTORY: This view provides historical information about active sessions in the database. It includes details about the SQL executed, the execution context, and the resources consumed by each session. By examining this view, you can identify the specific sessions and SQL queries that were active and potentially consuming excessive PGA memory around the time the ORA-04036 error occurred.

DBA_HIST_PROCESS_MEM_SUMMARY: This view contains historical summaries of memory usage by processes. It can help in identifying the processes that were consuming a significant amount of PGA memory, leading to the ORA-04036 error. This view provides

aggregated memory usage information over time, making it easier to pinpoint the processes responsible for high PGA memory consumption.

Together, these views offer a comprehensive overview of the memory usage patterns and specific queries or processes that might have contributed to exceeding the PGA_AGGREGATE_LIMIT, resulting in the ORA-04036 error.

Oracle Database Reference: DBA_HIST_ACTIVE_SESS_HISTORY

Oracle Database Reference: DBA_HIST_PROCESS_MEM_SUMMARY

Oracle Database Performance Tuning Guide: Managing Memory

# Question 9

**Question Type:** MultipleChoice

You are informed that the RMAN session that is performing the database duplication is much slower than usual. You want to know the approximate time when the rman operation will be completed.

Which view has this information?

## Options:

**A-** V$SESSION

**B-** V$SESSTAT

**C-** V$RMAN_BACKUP_JOB_DETAILS

**D-** V$SESSION_LONGOPS

## Answer:

D

## Explanation:

In Oracle Database, the V$SESSION_LONGOPS view provides insights into various operations within the database that are expected to take more than six seconds to complete. These include operations related to RMAN (Recovery Manager), such as database duplication tasks. This view displays information about the progress of these long-running operations, including the start time, elapsed time, and estimated time to completion.

When an RMAN session is performing a database duplication and is observed to be slower than usual, checking the V$SESSION_LONGOPS view can give an approximation of when the RMAN operation might complete. This view includes fields like TIME_REMAINING and ELAPSED_SECONDS that help in estimating the completion time of the operation based on its current progress.

Oracle Database Reference: V$SESSION_LONGOPS

Oracle Database Backup and Recovery User's Guide: Monitoring RMAN Jobs