

1Z0-515^{Q&As}

Data Warehousing 11g Essentials

Pass Oracle 1Z0-515 Exam with 100% Guarantee

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

<https://www.leads4pass.com/1z0-515.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

Flash in the Oracle Exadata Database Machine is of lower latency (and enables faster access to data) than:

- A. Memory
- B. Memory and locally attached disk
- C. Memory and network-attached disk
- D. Locally attached and network-attached disk

Correct Answer: D

Explanation:

The Exadata Smart Flash Cache is also used to reduce the latency of log write I/O eliminating performance bottlenecks that might occur due to database logging.

Flash memory has very good average write latency. It is faster than locally attached and network- attached disks. However, it is not faster than Memory (RAM).

References:

QUESTION 2

How can you implement near real time data integration with Oracle Data Integrator?

- A. By accessing Change Data Capture records from logs
- B. By using Exchange Partition
- C. By mining Oracle UNDO segments
- D. By reading operating system logs

Correct Answer: A

Explanation: Conventional "Extract, Transform, Load" (ETL) tools closely intermix data transformation rules with integration process procedures, requiring the development of both data transformations and data flow. Oracle Data Integrator (ODI) takes a different approach to integration by clearly separating the declarative rules (the "what") from the actual implementation (the "how"). With ODI, declarative rules describing mappings and transformations are defined graphically, through a drag-and-drop interface, and stored independently from the implementation. ODI automatically generates the data flow, which can be fine-tuned if required. This innovative approach for declarative design has also been applied to ODI's framework for Changed Data Capture. ODI's CDC moves only changed data to the target systems and can be integrated with Oracle GoldenGate, thereby enabling the kind of real time integration that businesses require.

References:

QUESTION 3

Which is NOT an advantage provided by partitioning?

- A. Reduces storage requirements for tables
- B. Can add to the benefits of parallelism through parallel partition-wise joins
- C. Can improve performance by reducing I/O
- D. Provides added flexibility for maintenance operations

Correct Answer: A

Explanation:

Table storage requirements would increase, but the benefits are huge.

Oracle partitioning is a divide-and-conquer approach to improving Oracle maintenance and SQL performance. Anyone with un-partitioned databases over 500 gigabytes is courting disaster. Databases become unmanageable, and serious problems occur:

*

SQL may perform poorly - Without Oracle partitioning, SQL queries with full-table scans take hours to complete. In a full scan, the smaller the Oracle partition, the faster the performance. Also, index range scans become inefficient.

*

Recovery - Files recovery takes days, not minutes

*

Maintenance - Rebuilding indexes (important to re-claim space and improve performance)

Oracle partitioning has many benefits to improve performance and manageability:

*

Stable

*

Robust

*

Faster backups

*

Less overhead

*

Easier management Maintenance of Oracle partitioned tables is improved because maintenance can be focused on particular portions of tables. For maintenance operations across an entire database object, it is possible to perform

these operations on a per-partition basis, thus dividing the maintenance process into more manageable chunks. (not D)

*

Faster SQL Oracle is partition-aware, and some SQL may improve its speed by several orders of magnitude (over 100x faster).

-Index range scans Oracle partitioning physically sequences rows in index-order causing a dramatic improvement (over 10x faster) in the speed of partition-key scans.

-Full-table scans Oracle partition pruning only accesses those data blocks required by the query.

-Table joins Oracle partition-wise joins take the specific sub-set of the query partitions, causing huge speed improvements on nested loop and hash joins. (not C)

-You can also improve the performance of massive join operations when large amounts of data (for example, several million rows) are joined together by using partition-wise joins. (not B)

-Updates Oracle parallel query for partitions improves batch load speed.

References:

QUESTION 4

Identify the true statement about adaptive parallelism.

- A. It is turned on by default.
- B. It is turned off by default.
- C. You should always leave the default setting
- D. There is no such thing.

Correct Answer: A

Explanation: Adaptive Parallelism: The adaptive multiuser algorithm, which is enabled by default, reduces the degree of parallelism as the load on the system increases. When using the Oracle Database adaptive parallelism capabilities, the database uses an algorithm at SQL execution time to determine whether a parallel operation should receive the requested DOP or have its DOP lower to ensure the system is not overloaded. In a system that makes aggressive use of parallel execution by using a high DOP, the adaptive algorithm adjusts the DOP down when only a few operations are running in parallel. While the algorithm still ensures optimal resource utilization, users may experience inconsistent response times. Using solely the adaptive parallelism capabilities in an environment that requires deterministic response times is not advised. Adaptive parallelism is controlled through the database initialization parameter `PARALLEL_ADAPTIVE_MULTI_USER`.

References:

QUESTION 5

You are looking to size a data warehouse configuration. If the I/O throughput for the CPUs is 25 GB/s, the I/O throughput for the HBA is 18 GB/s, and the I/O throughput for the disk subsystem is 6 GB/s, what is the overall throughput of the data warehouse?

- A. 25 GB/s
- B. 18 GB/s
- C. 6 GB/s
- D. It depends on how many processors are in the servers.

Correct Answer: C

Explanation:

In this scenario the disk subsystem is the bottleneck. It determines the throughput.

Note: Each of the components must provide sufficient I/O bandwidth to ensure a well-balanced I/O system.

The end-to-end I/O system consists of more components than just the CPUs and disks. A well-balanced I/O system must provide approximately the same bandwidth across all components in the I/O system.

These components include:

*

Host bus adapters (HBAs), the connectors between the server and the storage.

*

Switches, in between the servers and a storage area network (SAN) or network attached storage (NAS).

*

Ethernet adapters for network connectivity (GigE NIC or Infiniband). In an Oracle Real Application Clusters (Oracle RAC) environment, you need an additional private port for the interconnect between the nodes that you should not include when sizing the system for I/O throughput. The interconnect must be sized separately, taking into account factors such as internode parallel execution.

*

Wires that connect the individual components.

References:

QUESTION 6

Your customer is looking to implement ad-hoc analysis in a data warehouse. Which approach is least likely to be used assuming that the customer does not want the expense of managing view?

- A. Star schema
- B. Snowflake schema
- C. Third normal form schema
- D. OLAP

Correct Answer: C

Explanation: Data warehouses often use denormalized or partially denormalized schemas (such as a star schema) to optimize query performance. On the other hand OLTP (Online Transaction Processing) systems often use fully normalized schemas to optimize update/insert/delete performance, and to guarantee data consistency.

References:

QUESTION 7

You want partitions to be automatically created when data that does not fit into current date range loaded. Which type of partitioning would you implement?

- A. Hash
- B. List
- C. Invisible
- D. Interval

Correct Answer: D

Explanation: Interval Partitioning was introduced in 11g, interval partitions are extensions to range partitioning. These provide automation for equi-sized range partitions. Partitions are created as metadata and only the start partition is made persistent. The additional segments are allocated as the data arrives. The additional partitions and local indexes are automatically created.

Note: Partitioning is one of the most sought after options for data warehousing. Almost all Oracle data warehouses use partitioning to improve the performance of queries and also to ease the day-to-day maintenance complexities. Starting with 11G, more partitioning options have been provided and these should reduce the burden of the DBA to a great extent.

References:

QUESTION 8

Which feature would enable higher availability during maintenance operations while also improving query response performance?

- A. Partitioning
- B. Materialized views
- C. Bitmap Indexing
- D. OLAP

Correct Answer: A

Explanation: Partitioning enhances the performance, manageability, and availability of a wide variety of applications and helps reduce the total cost of ownership for storing large amounts of data. Partitioning allows tables, indexes, and index-organized tables to be subdivided into smaller pieces, enabling these database objects to be managed and accessed at

a finer level of granularity. Oracle provides a rich variety of partitioning strategies and extensions to address every business requirement. Moreover, since it is entirely transparent, partitioning can be applied to almost any application without the need for potentially expensive and time consuming application changes.

References:

QUESTION 9

Data Guard compresses data:

- A. Always
- B. When using logical standby
- C. When using physical standby
- D. When catching up after a network failure

Correct Answer: C

Explanation:

A Physical standby database replicates the exact contents of its primary database across the Oracle Net network layer. While the physical storage locations can be different, the data in the database will be exactly the same as the primary database.

Incorrect answer:

A, B: Logical standby databases convert the redo generated at the Primary database into data and SQL and then re-apply those SQL transactions on the Logical standby, thus physical structures and organization will be different from the Primary database. Users can read from logical standby databases while the changes are being applied and, if the GUARD is set to STANDBY (ALTER DATABASE GUARD STANDBY;), write to tables in the Logical standby database that are not being maintained by SQL Apply. Unfortunately there are a number of unsupported objects (ie: tables or sequences owned by SYS, tables that use table compression, tables that underlie a materialized view or Global temporary tables (GTTs)) and unsupported data types (ie: Datatypes BFILE, ROWID, and UROWID, user-defined TYPEs, Multimedia data types like Oracle Spatial, ORDDICOM, and Oracle Text Collections (e.g. nested tables, VARRAYs), SecureFile LOBs, OBJECT RELATIONAL XMLTypes and BINARY XML).[2] Physical standby may be appropriate in such a case.

QUESTION 10

You can use Oracle Data Mining unstructured data.

- A. TRUE
- B. FALSE

Correct Answer: A

Explanation: Data that cannot be meaningfully interpreted as numerical or categorical is considered unstructured for purposes of data mining. It has been estimated that as much as 85% of enterprise data falls into this category. Extracting meaningful information from this unstructured data can be critical to the success of a business.

Unstructured data may be binary objects, such as image or audio files, or text objects, which are language-based. Oracle Data Mining supports text objects. Text must undergo a transformation process before it can be mined. Once the data has been properly transformed, the case table can be used for building, testing, or scoring data mining models. Most Oracle Data Mining algorithms support text

References:

QUESTION 11

One goal of your Information Lifecycle Management strategy using Oracle's ILM capabilities is to reduce cost or online storage. Identify two database options that would help in enabling such a strategy.

- A. RAC and Advanced Compression
- B. RAC and Partitioning
- C. Partitioning and Advanced Compression
- D. RAC One and Advanced Compression

Correct Answer: B

Explanation: Advanced compression: Advanced Compression, an option introduced in Oracle Database 11 g Enterprise Edition, offers a comprehensive set of compression capabilities to help organizations reduce costs, while maintaining or improving performance. It significantly reduces the storage footprint of databases through compression of structured data (numbers, characters) as well as unstructured data (documents, spreadsheets, XML and other files). It provides enhanced compression for database backups and also includes network compression capabilities for faster synchronization of standby databases. Archival Compression:

*

Built on HCC technology

*

Compression algorithm optimized for maximum storage savings

*

Benefits any application with data retention requirements

*

Best approach for ILM and data archival

Partitioning: There are a number of benefits to partitioning data. Partitioning provides an easy way to distribute the data across appropriate storage devices depending on its usage, while still keeping the data online and stored on the most cost-effective device. Since partitioning is completely transparent to anyone accessing the data, no application changes are required, thus partitioning can be implemented at any time.

Note There is a wide variety of information held in an organization today, for example it could be an email message, a picture, or an order in an Online Transaction Processing System. Therefore, once the type of data being retained has been identified, you already have an understanding of what its evolution and final destiny is likely to be. The challenge now before all organizations, is to understand how their data evolves and grows, monitor how its usage changes over time, and decide how long it should survive. In addition, the evolving rules and regulations such as Sarbanes-Oxley, place additional constraints on the data that is being retained and some organizations now require that data is deleted when there is no longer a legal requirement to keep it, to avoid expensive e-discovery when the data is requested for a legal matter.

Implementing ILM using Oracle Database 11g Page 4 Information Lifecycle Management (ILM) is designed to address these issues, with a combination of processes, policies, software and hardware so that the appropriate technology can be used for each phase of the lifecycle of the data.1

References:

QUESTION 12

Which is NOT an available composite partition in Oracle Database 11g?

- A. range-list
- B. list-list
- C. list-range
- D. interval-hash

Correct Answer: D

Explanation:

Extended Composite Partitioning

In previous releases of Oracle, composite partitioning was limited to Range-Hash and Range-List partitioning. Oracle 11g Release 1 extends this to allow the following composite partitioning schemes:

Range-Hash (available since 8i)

Range-List (available since 9i)

Range-Range

List-Range

List-Hash

List-List

Note: interval-hash is a valid Interval partitioning.

References:

QUESTION 13

You want to enable result set caching to quickly see if this feature will help the performance of your application. Which is the quickest way to make this determination?

- A. Set RESULT_CACHE_MODE = FORCE in the initialization file.
- B. Set RESULT_CACHE = ENABLED in the initialization file.
- C. Set RESULT_CACHE_MAX_SIZE = 0.
- D. Set RESULT_CACHE = ENABLED in the initialization file and use a RESULT_CACHE hint in queries.

Correct Answer: A

Explanation: The RESULT_CACHE_MODE initialization parameter determines the SQL query result cache mode. The parameter specifies when a ResultCache operator is spliced into a query's execution plan. The parameter accepts the following values: FORCE The ResultCache operator is added to the root of all SELECT statements, if that is possible. However, if the statement contains a NO_RESULT_CACHE hint, then the hint takes precedence over the parameter setting. MANUAL The ResultCache operator is added, only if you use the RESULT_CACHE hint in the SQL query.

References:

QUESTION 14

What data can you compress using Advanced Compression in Oracle Database 11g?

- A. Read only data
- B. Data that can be updated, inserted and/or deleted (DML)
- C. Only data being archived
- D. Data warehousing data

Correct Answer: B

Explanation:

Oracle Database 11g has new option named as Oracle Advanced Table Compression option which aims at reducing space occupied by data for both OLTP and warehouse databases. This option provides the following types of compression:

*

Compression of data tables even for OLTP environment. (Previous versions had compression option for tables that are mostly read only).

- *
Compression of unstructured data in SecureFiles.
 - *
Compression of RMAN backups.
 - *
Compression in Data Pump Export files.
 - *
Compression of redo data transmitted to a standby database during redo gap resolution (when data guard is configured).
-

QUESTION 15

Which unique method of improving performance is NOT used by the Oracle Exadata Database Machine?

- A. Flash to improve query performance
- B. Reduces the amount of data required to flow through I/O
- C. Increases the I/O using InfiniBand
- D. Performs analysis in a special in-memory database

Correct Answer: D

References:

[1Z0-515 VCE Dumps](#)

[1Z0-515 Study Guide](#)

[1Z0-515 Braindumps](#)