



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.passapply.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

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 2

Identify the type of refresh that is NOT supported by materialized views.



- A. Deferred
- B. Incremental
- C. Full
- D. Heuristic

Correct Answer: D

Explanation:

Use the CREATE MATERIALIZED VIEW statement to create a materialized view. A materialized view is a database object that contains the results of a query.

Incorrect answer:

A: Specify DEFERRED to indicate that the materialized view is to be populated by the next REFRESH operation.

B: Oracle Database uses the default index to speed up incremental (FAST) refresh of the materialized view.

C: By default, Oracle Database creates a primary key materialized view with refresh on demand only. If a materialized view log exists on the table, then the column can be altered to be capable of fast refresh. If no such log exists, then only full refresh of the column is possible.

References:

QUESTION 3

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 4

What is the difference between an ETL (Extraction Transformation Load) approach and an ELT (Extraction Load Transformation) approach to data integration? Select one.

- A. ETL can operate between heterogeneous data sources.
- B. ELT requires a separate transformation server.
- C. ELT transforms data on the target server.
- D. ELT cannot be used for incremental data loading.

Correct Answer: C

Explanation:

There are two approaches to consider for data integration: ELT and ETL. The difference between ETL and ELT lies in the environment in which the data transformations are applied. In traditional ETL, the transformation takes place when the data is en route from the source to the target system. In ELT, the data is loaded into the target system, and then transformed within the target system environment.

References:

QUESTION 5

How many Exadata Storage Server cells are there in a Full Rack Exadata database machine configuration that has 8 Database Server nodes?

- A. 2
- B. 14
- C. 16
- D. 24

Correct Answer: B

Explanation:

QUESTION 6

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 7

Which questions CANNOT be addressed by Oracle Data Mining?

- A. Fraud detection
- B. Prediction of customer behavior
- C. Root cause de
- D. Identify factors associated with a business problem

Correct Answer: C

Explanation:

Data Mining can provide valuable results:

*Predict customer behavior (Classification) (not B)

*Predict or estimate a value (Regression)

*Segment a population (Clustering)

*Identify factors more associated with a business

problem (Attribute Importance) (not D)

*

Find profiles of targeted people or items (Decision Trees)

*

Determine important relationships and market baskets within the population (Associations)

*



Find fraudulent or rare events (Anomaly Detection) (not A)

References:

QUESTION 8

What would you do to compress data in partitions that are frequently updated in Oracle Database 11g?

- A. Use Hybrid Columnar Compression.
- B. Use Advanced Compression Option.
- C. Use Hybrid Partitions.
- D. Avoid compressing any data.

Correct Answer: B

Explanation:

Advanced Compression features in Oracle Database 11g include:

* Online Transaction Processing (OLTP) Table Compression: This breakthrough compression feature compresses table data during all types of data manipulation operations, including conventional INSERT or UPDATE. OLTP Table Compression leverages a sophisticated and intelligent algorithm that minimizes compression overhead during write operations, thereby making it viable for highly transactional workloads

References:

QUESTION 9

Identify the dimension that appears most often in queries in a data warehouse.

- A. Product dimension
- B. Time dimension
- C. Cost dimension
- D. Location dimension

Correct Answer: B

Explanation: In a data warehouse, a dimension is a data element that categorizes each item in a data set into non-overlapping regions. A data warehouse dimension provides the means to "slice and dice" data in a data warehouse. Dimensions provide structured labeling information to otherwise unordered numeric measures. For example, "Customer", "Date", and "Product" are all dimensions that could be applied meaningfully to a sales receipt. A dimensional data element is similar to a categorical variable in statistics. The primary function of dimensions is threefold: to provide filtering, grouping and labeling. For example, in a data warehouse where each person is categorized as having a gender of male, female or unknown, a user of the data warehouse would then be able to filter or categorize each presentation or report by either filtering based on the gender dimension or displaying results broken out by the gender.



QUESTION 10

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 11

You are looking for some general design principles that could be used in designing every large scale data warehouse you create. Identify the principle that would have the widest applicability.

- A. Partition your tables appropriately to produce partition-wise joins.
- B. Always use a star schema or snowflake schema design.
- C. Do as much analytics as possible in your BI tools.
- D. Always use Oracle OLAP.

Correct Answer: A

Explanation:

Partition-wise joins can be full or partial. Oracle decides which type of join to use. A full partition-wise join divides a large join into smaller joins between a pair of partitions from the two joined tables. To use this feature, you must equipartition both tables on their join keys, or use reference partitioning.

Oracle Database can perform partial partition-wise joins only in parallel. Unlike full partition-wise joins,



partial partition-wise joins require you to partition only one table on the join key, not both tables.

Note: Partition-wise joins reduce query response time by minimizing the amount of data exchanged among parallel execution servers when joins execute in parallel. This significantly reduces response time and improves the use of both CPU and memory resources. In Oracle Real Application Clusters (RAC) environments, partition-wise joins also avoid or at least limit the data traffic over the interconnect, which is the key to achieving good scalability for massive join operations.

References:

QUESTION 12

For data warehousing, identify the benefits that would NOT be provided by the use of RAC.

- A. Distribute workload across all the nodes.
- B. Distribute workload to some of the nodes.
- C. Provide parallel query servers.
- D. Provide high availability for all the operations.

Correct Answer: B

Explanation:

With Oracle RAC the workload can be distributed across all cluster nodes, parallel query servers can be provided through the Parallel Query tool, and high availability can be obtained through, for example, Oracle Clusterware.

Note: Oracle RAC (Real Application Clusters) is a cluster database with a shared cache architecture that overcomes the limitations of traditional shared-nothing and shared-disk approaches to provide highly scalable and available database solutions for all your business applications. Oracle RAC is a key component of Oracle's private cloud architecture. Oracle RAC support is included in the Oracle Database Standard Edition for higher levels of system uptime.

References:

QUESTION 13

Which feature of Oracle Warehouse Builder can be used to help ensure data quality?

- A. Data extraction
- B. Data profiling
- C. Logical mapping
- D. Exception reporting

Correct Answer: B



Explanation: After you connect to your data sources in Oracle Warehouse Builder (including Oracle databases, sources accessed through gateways, and flat file sources) you can apply full-featured data profiling to generate statistics about data quality, and to discover complex patterns, foreign key relationships, and functional dependencies. You can then design complex data rules and create data auditors to monitor compliance with those rules in any source or target system in your landscape, regardless of whether those sources are loaded using Oracle Warehouse Builder or other ETL tools.

For customers who have selected solutions other than Oracle Warehouse Builder for data profiling and data quality, these can be applied independently of Oracle Warehouse Builder ETL and design features.

Note: Oracle Warehouse Builder is a full-featured data integration, data warehousing, data quality and metadata management solution designed for the Oracle database. Oracle Warehouse Builder is an integral part of Oracle Database 11g Release 2 (11.2) and is installed as part of every database installation (other than Oracle Database XE).

The major feature areas of Oracle Warehouse Builder include: *Data modeling *Extraction, Transformation, and Load (ETL) *Data profiling and data quality *Metadata management *Business-level integration of ERP application data *Integration with Oracle business intelligence tools for reporting purposes *Advanced data lineage and impact analysis Oracle Warehouse Builder is also an extensible data integration and data quality solutions platform. Oracle Warehouse Builder can be extended to manage metadata specific to any application, and can integrate with new data source and target types, and implement support for new data access mechanisms and platforms, enforce your organization's best practices, and foster the reuse of components across solutions.

References:

QUESTION 14

You have analyzed your client's workload and the SQL Access Advisor in Enterprise Manager recommends that you create some materialized views to improve performance. What should you do in order to most simply implement this change?

- A. Rewrite all the queries in the application to identify materialized view
- B. Rewrite existing queries. New queries will automatically use the views.
- C. Respond positively to the Advisor to create the materialized views.
- D. Build virtual views on a third normal form schema.

Correct Answer: C

Explanation: Enterprise Manager provides a very simple interface for the SQL Access Advisor (Advisor Central > SQL Access Advisor > SQL Access Advisor). The first page allows you to create tasks to test existing indexes, materialized view and partitions, or create tasks to suggest new structures. The "Workload Source" page allows you to define the workload to associate with the task. The basic options allow the workload to be gathered from the cursor cache, an existing SQL tuning set, or a hypothetical workload based on specific schema objects. The "Recommendation Options" page allows you to define which type of recommendations you are interested in (Indexes, Materialized Views and Partitioning). After reviewing the result of the analysis you can decide if you should accept or ignore the suggested recommendations.

Note: The SQL Access Advisor was introduced in Oracle 10g to make suggestions about additional indexes and materialized views which might improve system performance.

References:



QUESTION 15

What two types of results can be cached in the Result Set Cache?

- A. Results of an SQL query
- B. Results from a PL/SQL function
- C. Sequence object results
- D. Result sets derived from data dictionary tables

Correct Answer: AB

Explanation:

Your applications sometime send repetitive queries to the database. To improve the response time of repetitive queries, results of queries, query fragments, and PL/SQL functions can be cached in memory. A result cache stores the results of queries shared across all sessions. When these queries are executed repeatedly, the results are retrieved directly from the cache memory.

You must annotate a query or query fragment with a result cache hint to indicate that results are to be stored in the query result cache.

The query result set can be cached in the following ways:

*

Server-side Cache

*

Client Result Cache

Oracle Database 11g Release 1 (11.1) provides support for server-side Result Set caching for both JDBC types. The server-side result cache is used to cache the results of the current queries, query fragments, and PL/SQL functions in memory and then to use the cached results in future executions of the query, query fragment, or PL/SQL function. The cached results reside in the result cache memory portion of the SGA. A cached result is automatically invalidated whenever a database object used in its creation is successfully modified. The server-side caching can be of the following two types:

*

SQL Query Result Cache (A)

*

PL/SQL Function Result Cache (B)

[Latest 1Z0-515 Dumps](#)

[1Z0-515 Practice Test](#)

[1Z0-515 Study Guide](#)