

ISLEVER

1Z0-515

Data Warehousing 11g Essentials

DEMO

<https://www.islever.com/1z0-515.html>

<https://www.islever.com/oracle.html>

For the most up-to-date exam questions and materials, we recommend visiting our website, where you can access the latest content and resources.

Note: The answer is for reference only, you need to understand all question.

Q1. Identify the true statement about REF partitions.

- A. REF partitions have no impact on partition-wise joins.
- B. Changes to partitioning in the parent table are automatically reflected in the child table.
- C. Changes in the data in a parent table are reflected in a child table.
- D. REF partitions can save storage space in the parent table.

Answer: B

Q2. Identify the control structure that would NOT be defined as part of a data flow with Oracle Data Integrator.

- A. Loops
- B. Conditions
- C. Error handling
- D. GOTOS

Answer: D

Q3. For which type of query is the SQL result cache automatically disabled?

- A. Queries that access data which changes frequently
- B. Queries that return large amounts of data
- C. Queries that use SQL functions such as SYSDATE
- D. Queries that are used infrequently

Answer: C

Q4. Which is NOT among Oracle SQL Analytic functions included in Oracle Database 11g?

- A. Ranking functions
- B. Substring functions
- C. Window aggregate functions
- D. LAG/LEAD functions
- E. All of the above

Answer: E

Q5. How many Exadata Storage Server cells can be used in a grid?

- A. 7
- B. 14
- C. 128
- D. No practical limit

Answer: D

Q6. Identify the action that you CANNOT perform using Database Resource Manager.

- A. Define Consumer Groups.
- B. Create rules to map sessions to Consumer Groups.
- C. Define a Resource Plan.
- D. Allocate individual CPUs to Consumer Groups.

Answer: D

Q7. You will be implementing a data warehouse for one of your customers. In your design process, which index type is most likely to be used to improve the performance of some queries where the data is of low cardinality?

- A. Bitmap indexes
- B. B*-tree indexes
- C. Reverse indexes
- D. Invisible indexes

Answer: A

Q8. You think that result set caching might provide some benefits for your current data warehouse scenario.

You perform some analysis on the composition of the queries used in the scenario. Identify the result of the analysis that would indicate the most potential for improvement with result set caching.

- A. The scenario consists mainly of queries that are used infrequently.
- B. The scenario consists mainly of queries that work on data which changes frequently.
- C. The scenario consists mainly of queries with long run times and small result sets.
- D. All data warehouse scenarios will benefit from result set caching.

Answer: C

Q9.

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.

Answer: B

Q10. 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

Answer: D

Q11. Identify the benefit of using interval partitioning.

- A. Automatic creation of new partitions based on hash values
- B. Automatic creation of new partitions based on the value of data being entered
- C. Improved performance compared to range partitions
- D. Automatic transfer of older partitions lower cost storage

Answer: B

Q12. Your customer wants to implement an ILM strategy. The customer must have which option when deploying Oracle's ILM Assistant to implement this strategy?

- A. RAC
- B. Partitioning
- C. OLAP
- D. OracleClusterware

Answer: B