Welcome to Cert007 - Your Ultimate IT Certification Partner

- Real Exam Questions
- Free Updates
- Expert Support
- Instant Access
- Money-Back Guarantee

Visit us at **https://www.cert007.com/** for more information

Exam : **DP-750**

Title : Implementing Data Engineering Solutions Using Azure Databricks

Version : DEMO

1.A company has an Azure Databricks job that runs every hour to process landing-zone files and write curated Delta tables. Each run should use fresh isolated compute and terminate when the workload finishes.
Which compute option should be used?
A. All-purpose compute
B. SQL warehouse
C. Job compute
D. Interactive shared cluster
Answer: C
Explanation:
Job compute is designed for scheduled production jobs. It can be created for each job run and terminated automatically after completion, which improves isolation and cost control.

2.A group of business analysts needs to query governed Unity Catalog tables by using SQL. They do not need to run notebooks or manage Spark cluster settings.
Which compute option is most appropriate?
A. SQL warehouse
B. All-purpose compute
C. Single-user cluster
D. Cluster pool
Answer: A
Explanation:
SQL warehouses are designed for SQL analytics workloads and allow analysts to query governed tables without managing notebook clusters.

3.A data engineering team is developing and testing PySpark notebooks interactively. The users need to attach notebooks, inspect intermediate results, and rerun cells during development.
Which compute type should they use?
A. Serverless SQL warehouse
B. All-purpose compute
C. Job compute
D. Delta Sharing compute
Answer: B
Explanation:
All-purpose compute is best suited for interactive notebook development and iterative testing.

4.A nightly ETL workload has unpredictable input volume. Some nights require only a few workers, while month-end processing requires significantly more resources. The company also wants to avoid paying for idle compute.
Which two settings should be configured? Choose two.
A. Autoscaling
B. Fixed worker count
C. Disable auto-termination
D. Auto-termination

Answer: A, D
Explanation:
Autoscaling adjusts worker capacity based on demand. Auto-termination stops idle compute to reduce unnecessary cost.

5.An administrator wants to prevent users from creating clusters with unsupported runtimes, oversized node types, or missing auto-termination settings.
What should the administrator configure?
A. Table ACLs
B. Delta Sharing recipients
C. External locations
D. Compute policies
Answer: D
Explanation:
Compute policies can enforce rules for cluster configuration, including allowed node types, runtimes, autoscaling limits, and auto-termination behavior.