

Microsoft DP-700

Implementing Data Engineering Solutions Using Microsoft

Thank You for Downloading DP-700 Updated Exam Questions

<https://certs4sale.com/microsoft/dp-700-pdf-exam-dumps>

Microsoft DP-700

Implementing Data Engineering Solutions Using Microsoft Fabric

FABRIC DATA ENGINEER
ASSOCIATE

TARGET CERTIFICATION

Microsoft Certified: Fabric Data Engineer Associate

EXAM CODE

DP-700

DOCUMENT PURPOSE

Practice Questions & Detailed Solutions

FOCUS AREAS

Ingestion, Lakehouses, Warehouses & Governance

Question 1

MULTIPLE CHOICE

Your organization is implementing a Microsoft Fabric workspace for a data engineering team. You need to configure workspace settings to support multiple teams working on different projects while maintaining role-based access control. Users from the marketing team should only see and work with marketing-related items, while the analytics team should have full workspace visibility. Which two actions should you take to achieve this configuration?

- A. Use workspace roles to assign Marketing Contributor and Analytics Admin roles to respective team members, then configure item-level permissions by creating a security group for each team
- B. Create separate workspaces for each team and use workspace roles to manage access
- C. Configure only dataset-level Row-Level Security (RLS) without using workspace roles, as this is sufficient for all access control needs
- D. Use only Azure AD group membership and rely entirely on Fabric capacity settings to manage team isolation
- E. Assign all users as workspace admins and handle data access restrictions only through report-level filters

Correct Answer: A, B

Explanation:

The correct approach combines workspace roles (which control who can access workspace-level items and features) with item-level permissions (which control access to specific items like datasets and reports). Additionally, creating separate workspaces is a valid architectural pattern for team isolation when projects are sufficiently independent. Workspace roles alone do not provide item-level granularity, and RLS is a data-level control that complements but doesn't replace workspace and item permissions. Assigning all users as admins

Question 2

SINGLE CHOICE

You are designing a data ingestion solution for a healthcare organization that receives patient data from multiple hospital systems throughout the day. The data arrives in batches at irregular intervals (every 2-8 hours), and the organization requires that data be available for analysis within 30 minutes of arrival. You need to recommend a loading pattern that balances latency requirements with cost efficiency. Which loading pattern should you implement?

- A. Implement a scheduled batch load with a 30-minute interval using Data Factory pipelines, combined with incremental load logic to move only new or changed records
- B. Use real-time streaming ingestion with Azure Event Hubs to capture all data events, as this is the only pattern that meets the 30-minute SLA
- C. Implement full batch loads every 8 hours without incremental logic, since the data arrives in batches anyway
- D. Use manual data imports through Power BI Desktop with a scheduled task, as this minimizes infrastructure costs
- E. Implement a 5-minute polling interval on all source systems to detect new data, regardless of actual data arrival patterns

Correct Answer: A

Explanation:

A scheduled batch load with incremental logic is the optimal pattern here because it meets the 30-minute latency requirement, handles the irregular arrival intervals efficiently, and minimizes costs compared to true streaming. Since data arrives in batches (not continuous streams), streaming ingestion would be overengineered and more expensive. Full batch loads every 8 hours violate the SLA. Manual imports lack scalability and reliability.

Question 3

MULTIPLE CHOICE

Your team is implementing a Data Factory pipeline in Microsoft Fabric that orchestrates a multi-stage ETL process: raw data ingestion from an API, transformation using a Spark notebook, and loading into a warehouse table. During a recent run, the pipeline failed at the transformation stage with a Spark execution error. You need to diagnose the issue efficiently. Which two actions would provide the most relevant diagnostic information?

- A. Review the Spark notebook execution logs in the monitoring section and examine the pipeline run history to identify the specific error message and line number
- B. Check the Fabric capacity utilization metrics to confirm whether compute resources were insufficient during the pipeline run

Question 4

SINGLE CHOICE

You are implementing a lifecycle management strategy for a Fabric workspace where data evolves through stages: raw ingestion (kept for 6 months), enriched datasets (kept for 24 months), and published reports (kept indefinitely). You need to ensure that old intermediate datasets are removed to control storage costs while preserving historical reports and their underlying data relationships. Which approach should you use to implement this lifecycle policy?

- A.** Use Fabric's dataset retention policies in the workspace settings to configure different TTL (time-to-live) values for different item types, combined with metadata tagging to identify lifecycle stage
- B.** Implement manual deletion scripts that run monthly to remove datasets older than the retention period without tagging or automation
- C.** Archive all data to external cold storage after 6 months and delete it from Fabric immediately, then restore on demand for any historical analysis
- D.** Configure only lakehouse retention policies and ignore intermediate datasets, since reports will continue to function with missing source data
- E.** Store all items indefinitely in Fabric, then use Power BI's data refresh policies to manage memory consumption

Correct Answer: A

Explanation:

Fabric's dataset retention policies with metadata tagging enable automated, granular lifecycle management where different item types can have different retention periods. This approach is automated, auditable, and prevents accidental deletion of important items. Manual deletion scripts are error-prone and lack governance.

Question 5

SINGLE CHOICE

You are monitoring a Fabric warehouse that supports a critical business dashboard used by executives daily. Over the past week, you've noticed that query performance has degraded gradually: queries that previously completed in 2 seconds now take 8-12 seconds. No schema changes have occurred, and data volume has grown only 5%. You need to identify the root cause and optimize performance. Which investigation approach is most appropriate?

- A.** Review warehouse query execution plans and statistics in Fabric's monitoring tools to identify missing indexes or outdated statistics, then rebuild indexes and update statistics if needed
- B.** Immediately increase the warehouse SKU/compute tier, as query slowdown always indicates insufficient compute capacity

End of Practice Assessment

You have completed the sample set of questions for the **Microsoft DP-700: Implementing Data Engineering Solutions Using Microsoft Fabric** exam.

TOTAL QUESTIONS

5 Included

PASSING SCORE

700 / 1000

THANK YOU FOR DOWNLOADING DP-700 UPDATED EXAM QUESTIONS

Note: Thanks For Trying The Demo Of Our DP-700 Exam Product

Visit Our Site to Purchase the Full Set of Actual DP-700 Exam Questions With Answers.

Money Back Guarantee



Click The Link Below

<https://certs4sale.com/microsoft/dp-700-pdf-exam-dumps>