

DP-700 MICROSOFT FABRIC DATA ENGINEER PRACTICE EXAM (SAMPLE)

STUDY GUIDE



SAMPLE STUDY GUIDE WITH LIMITED QUESTIONS

VISIT US ONLINE FOR
THE FULL VERSION STUDY GUIDE

<https://dp700msfabricdataengineer.examzify.com>



Copyright © 2026 by Examzify - A Kaluba Technologies Inc. product.

ALL RIGHTS RESERVED.

No part of this book may be reproduced or transferred in any form or by any means, graphic, electronic, or mechanical, including photocopying, recording, web distribution, taping, or by any information storage retrieval system, without the written permission of the author.

Notice: Examzify makes every reasonable effort to obtain accurate, complete, and timely information about this product from reliable sources.

SAMPLE

Table of Contents

Copyright	1
Table of Contents	2
Introduction	3
How to Use This Guide	4
Questions	5
Answers	8
Explanations	10
Next Steps	15

SAMPLE

Introduction

Preparing for a certification exam can feel overwhelming, but with the right tools, it becomes an opportunity to build confidence, sharpen your skills, and move one step closer to your goals. At Examzify, we believe that effective exam preparation isn't just about memorization, it's about understanding the material, identifying knowledge gaps, and building the test-taking strategies that lead to success.

This guide was designed to help you do exactly that.

Whether you're preparing for a licensing exam, professional certification, or entry-level qualification, this book offers structured practice to reinforce key concepts. You'll find a wide range of multiple-choice questions, each followed by clear explanations to help you understand not just the right answer, but why it's correct.

The content in this guide is based on real-world exam objectives and aligned with the types of questions and topics commonly found on official tests. It's ideal for learners who want to:

- Practice answering questions under realistic conditions,
- Improve accuracy and speed,
- Review explanations to strengthen weak areas, and
- Approach the exam with greater confidence.

We recommend using this book not as a stand-alone study tool, but alongside other resources like flashcards, textbooks, or hands-on training. For best results, we recommend working through each question, reflecting on the explanation provided, and revisiting the topics that challenge you most.

Remember: successful test preparation isn't about getting every question right the first time, it's about learning from your mistakes and improving over time. Stay focused, trust the process, and know that every page you turn brings you closer to success.

Let's begin.

How to Use This Guide

This guide is designed to help you study more effectively and approach your exam with confidence. Whether you're reviewing for the first time or doing a final refresh, here's how to get the most out of your Examzify study guide:

1. Start with a Diagnostic Review

Skim through the questions to get a sense of what you know and what you need to focus on. Your goal is to identify knowledge gaps early.

2. Study in Short, Focused Sessions

Break your study time into manageable blocks (e.g. 30 – 45 minutes). Review a handful of questions, reflect on the explanations.

3. Learn from the Explanations

After answering a question, always read the explanation, even if you got it right. It reinforces key points, corrects misunderstandings, and teaches subtle distinctions between similar answers.

4. Track Your Progress

Use bookmarks or notes (if reading digitally) to mark difficult questions. Revisit these regularly and track improvements over time.

5. Simulate the Real Exam

Once you're comfortable, try taking a full set of questions without pausing. Set a timer and simulate test-day conditions to build confidence and time management skills.

6. Repeat and Review

Don't just study once, repetition builds retention. Re-attempt questions after a few days and revisit explanations to reinforce learning. Pair this guide with other Examzify tools like flashcards, and digital practice tests to strengthen your preparation across formats.

There's no single right way to study, but consistent, thoughtful effort always wins. Use this guide flexibly, adapt the tips above to fit your pace and learning style. You've got this!

Questions

SAMPLE

- 1. In the Medallion Architecture, what is the typical purpose of the Silver layer?**
 - A. To store raw, unprocessed data.
 - B. To store curated, business-ready aggregates.
 - C. To store cleaned, integrated data ready for analysis.
 - D. To store model training data.
- 2. Which option best describes the Lakehouse data store in Fabric?**
 - A. Optimized for unstructured data only
 - B. Combines data warehousing and data lake capabilities
 - C. Stores only raw CSV
 - D. Is the same as OneLake
- 3. What is a lakehouse table in Fabric, and how is it typically stored?**
 - A. A lakehouse table is a managed Delta table stored in OneLake
 - B. A lakehouse table is a Delta Lake format table stored on OneLake
 - C. A lakehouse table is a JSON document store
 - D. A lakehouse table is a CSV file on OneLake
- 4. What is order of evaluation of access in Fabric?**
 - A. Microsoft Entra ID authentication, Fabric access, Data security
 - B. Data security, Fabric access, Microsoft Entra ID authentication
 - C. Fabric access, Microsoft Entra ID authentication, Data security
 - D. Data security, Microsoft Entra ID authentication, Fabric access
- 5. How do you manage access control lists (ACLs) in OneLake?**
 - A. ACLs are managed automatically and cannot be edited by users.
 - B. ACLs can only grant read access, not write access.
 - C. Set ACLs on folders or datasets to grant read/write permissions to users/groups; use least privilege and auditing to track changes and ensure compliance.
 - D. ACLs are deprecated in favor of role-based access without specifics.

- 6. How would you implement time-based retention in a lakehouse to comply with policy?**
- A. Manual purge data on a quarterly basis
 - B. Disable retention to preserve data indefinitely
 - C. Retention is not supported in Fabric lakehouse
 - D. Use retention policies by dataset with automatic purging or archiving after defined retention periods; ensure backups or point-in-time restore.
- 7. In Fabric Real-Time Intelligence, where can ingested data reside for subsequent analysis?**
- A. Data lake only
 - B. On-premises data store
 - C. Lakehouse or Eventhouse
 - D. SQL data warehouse only
- 8. Which dynamic management view provides details about SQL commands running in the data warehouse?**
- A. sys.dm_exec_connections
 - B. sys.dm_exec_sessions
 - C. sys.dm_exec_requests
 - D. sys.dm_exec_query_stats
- 9. Which language is optimized for querying real-time data in an eventhouse?**
- A. Python
 - B. Java
 - C. KQL
 - D. SQL
- 10. You want to include data in an external Azure Data Lake Store Gen2 location in your lakehouse, without the requirement to copy the data. What should you do?**
- A. A Data pipeline that uses a Copy Data activity to load the external data into a file
 - B. Create a Dataflow Gen2 that extracts the data and loads it into a table
 - C. Create an external table
 - D. Create a shortcut

Answers

SAMPLE

1. C
2. B
3. B
4. A
5. C
6. D
7. C
8. B
9. C
10. D

SAMPLE

Explanations

SAMPLE

1. In the Medallion Architecture, what is the typical purpose of the Silver layer?

- A. To store raw, unprocessed data.
- B. To store curated, business-ready aggregates.
- C. To store cleaned, integrated data ready for analysis.**
- D. To store model training data.

In the Medallion Architecture, the Silver layer is where data is cleaned, standardized, and integrated so it's ready for analysis. After raw data lands in the Bronze layer, the Silver layer applies data quality checks, resolves inconsistencies, and joins related datasets from multiple Bronze sources to produce consistent, analytics-ready tables or views. This makes it easy for analysts and BI tools to run queries without dealing with raw, unstructured, or disjointed data. This is why the purpose is to store cleaned, integrated data ready for analysis. Raw data belongs in the Bronze layer, curated business-ready aggregates live in the Gold layer, and model training data isn't the primary function of the Silver layer—even though ML can use curated data, the Silver layer's primary role is to provide refined, analysis-ready datasets.

2. Which option best describes the Lakehouse data store in Fabric?

- A. Optimized for unstructured data only
- B. Combines data warehousing and data lake capabilities**
- C. Stores only raw CSV
- D. Is the same as OneLake

The Lakehouse in Fabric is designed to merge the strengths of data lakes and data warehouses. It provides scalable storage for diverse data types—structured, semi-structured, and unstructured—while also offering SQL-based querying, metadata management, schema enforcement, and ACID transactions for reliable analytics. This combination lets you store raw data and curated, optimized data in one place and analyze it efficiently with governance features. So, it best describes a system that combines data warehousing and data lake capabilities. It's not limited to unstructured data, not restricted to raw CSV files, and it isn't just the same as OneLake—the Lakehouse sits on top of storage like OneLake to enable analytics.

3. What is a lakehouse table in Fabric, and how is it typically stored?

- A. A lakehouse table is a managed Delta table stored in OneLake
- B. A lakehouse table is a Delta Lake format table stored on OneLake**
- C. A lakehouse table is a JSON document store
- D. A lakehouse table is a CSV file on OneLake

In Fabric, a lakehouse table is a Delta Lake format table stored on OneLake. Delta Lake provides a transaction log and ACID guarantees on top of the data lake, giving you reliable schema enforcement, atomic commits, and features like time travel and concurrent reads/writes. Storing the table on OneLake means the Delta files and metadata live in Fabric's unified storage layer, making the data accessible and governable across tools and pipelines. Data stored as JSON or CSV lacks the Delta Lake transaction log and schema enforcement, so those formats don't offer the same lakehouse capabilities.

4. What is order of evaluation of access in Fabric?

- A. Microsoft Entra ID authentication, Fabric access, Data security
- B. Data security, Fabric access, Microsoft Entra ID authentication
- C. Fabric access, Microsoft Entra ID authentication, Data security
- D. Data security, Microsoft Entra ID authentication, Fabric access

Access evaluation starts with proving who you are, then deciding what you're allowed to do, and finally enforcing protections on the data you access. First, Microsoft Entra ID authentication establishes the user's identity. Without a verified identity, you can't reliably determine permissions. Next, Fabric access checks evaluate whether that authenticated user has the required rights to access the particular Fabric resources (workspaces, datasets, pipelines) they're requesting. Only after access is allowed do data security policies come into play, applying data-level protections such as masking, encryption, or other safeguards to ensure the data remains protected according to policy. Approaches that try to apply access or security before identifying the user don't make sense, so the correct sequence is: authenticate, then verify Fabric access, then enforce data security.

5. How do you manage access control lists (ACLs) in OneLake?

- A. ACLs are managed automatically and cannot be edited by users.
- B. ACLs can only grant read access, not write access.
- C. Set ACLs on folders or datasets to grant read/write permissions to users/groups; use least privilege and auditing to track changes and ensure compliance.
- D. ACLs are deprecated in favor of role-based access without specifics.

OneLake uses access control lists attached to folders and datasets to manage who can view or modify data. By assigning read and write permissions to specific users or groups on those resources, you implement precise access rights. Following the least-privilege principle means granting only the minimum permissions each user needs to do their job, reducing risk. Keeping an audit trail of permission changes and data-access events helps you verify compliance and detect unusual activity. This approach is preferable because ACLs are explicit and flexible, allowing both read and write rights as needed, and they can be adjusted by authorized admins. They aren't automatically fixed, they aren't limited to read-only, and they aren't deprecated in favor of a single, blanket role model—RBAC and ACLs can work together to provide both broad and fine-grained access control.

6. How would you implement time-based retention in a lakehouse to comply with policy?

- A. Manual purge data on a quarterly basis
- B. Disable retention to preserve data indefinitely
- C. Retention is not supported in Fabric lakehouse
- D. Use retention policies by dataset with automatic purging or archiving after defined retention periods; ensure backups or point-in-time restore.

Time-based retention in a lakehouse is implemented with dataset-level retention policies that automatically purge or archive data after defined retention periods. This approach enforces the data lifecycle consistently, reduces manual tasks, and helps meet policy requirements. You can configure each dataset with its own retention window, and after that window the system automatically purges or moves data to an archive. To protect against accidental deletion or to support audits, maintain backups or enable point-in-time restore so you can recover data if needed. Manual purges are error-prone and inconsistent, disabling retention misses policy requirements, and saying retention isn't supported is incorrect in a typical lakehouse setup.

7. In Fabric Real-Time Intelligence, where can ingested data reside for subsequent analysis?

- A. Data lake only
- B. On-premises data store
- C. Lakehouse or Eventhouse**
- D. SQL data warehouse only

In Fabric Real-Time Intelligence, ingested data isn't limited to a single storage type—you can land it in either a Lakehouse or an Eventhouse for analysis. A Lakehouse combines the scalability and breadth of a data lake with the structured query and governance features of a data warehouse, making it suitable for both raw and curated data used in analytics. An Eventhouse is tailored for streaming, low-latency event data, enabling real-time processing and immediate insights. This flexibility lets you choose the storage target that best fits your workflow: long-term analytics on diverse data in a lakehouse, or rapid real-time analysis on streaming data in an eventhouse. The other options are too narrow: a data lake alone doesn't capture the combined analytics capabilities of a lakehouse, on-premises storage isn't the cloud-native model for Fabric, and a SQL data warehouse only doesn't cover the streaming/event use case.

8. Which dynamic management view provides details about SQL commands running in the data warehouse?

- A. sys.dm_exec_connections
- B. sys.dm_exec_sessions**
- C. sys.dm_exec_requests
- D. sys.dm_exec_query_stats

To see what SQL commands are running in the data warehouse, you look at runtime state provided by dynamic management views that expose active sessions. The session-focused view gives you one row per active session, including who is connected and when their activity started, and you can retrieve the exact statement being executed for that session by querying the associated SQL text (via joining with the text view). This makes it the most direct way to map a live command to the user or session that issued it. Other DMVs serve different angles: the connections view shows connection details, not the running commands; the requests view reports current requests but without the broader session context; and the query stats view aggregates historical performance rather than showing the current, in-flight SQL text. Thus, the per-session view is the best starting point for details about SQL commands currently running in the data warehouse, especially when you pair it with the SQL text retrieval.

9. Which language is optimized for querying real-time data in an eventhouse?

- A. Python
- B. Java
- C. KQL**
- D. SQL

Real-time event data querying benefits from a language that handles fast time-based filtering, parsing semi-structured logs, and quick aggregations over time windows. KQL is built for these workloads in Microsoft's telemetry and log analytics platforms, including eventstores and dashboards. It supports concise time range filtering, field extraction, and aggregations like summarizing counts or metrics by time bins, enabling interactive, near-instantaneous exploration of streaming data. While Python and Java are general-purpose languages and SQL works well for structured relational data, they don't provide the same native, time-series-oriented operators and performance optimizations that KQL offers for real-time event data.

10. You want to include data in an external Azure Data Lake Store Gen2 location in your lakehouse, without the requirement to copy the data. What should you do?

- A. A Data pipeline that uses a Copy Data activity to load the external data into a file
- B. Create a Dataflow Gen2 that extracts the data and loads it into a table
- C. Create an external table
- D. Create a shortcut**

The concept being tested is bringing external data into a lakehouse without duplicating it by using a lightweight pointer. Creating a shortcut to the Azure Data Lake Storage Gen2 location serves exactly this need: it provides a metadata reference to data that already lives in ADLS Gen2, so the lakehouse can access and query it directly without copying the data into the lakehouse storage. The data remains in its original location, preserving storage costs and avoiding duplication, while still making it available for analysis through the lakehouse interface and governance surface. Other options involve moving or duplicating data: a data pipeline or dataflow would copy data into a table or file in the lakehouse, which contradicts the requirement to avoid copying. An external table could reference external data, but shortcuts are the intended, simpler mechanism in this scenario to integrate external ADLS Gen2 data into the lakehouse without copying.

SAMPLE

Next Steps

Congratulations on reaching the final section of this guide. You've taken a meaningful step toward passing your certification exam and advancing your career.

As you continue preparing, remember that consistent practice, review, and self-reflection are key to success. Make time to revisit difficult topics, simulate exam conditions, and track your progress along the way.

If you need help, have suggestions, or want to share feedback, we'd love to hear from you. Reach out to our team at hello@examzify.com.

Or visit your dedicated course page for more study tools and resources:

<https://dp700msfabricdataengineer.examzify.com>

We wish you the very best on your exam journey. You've got this!

SAMPLE