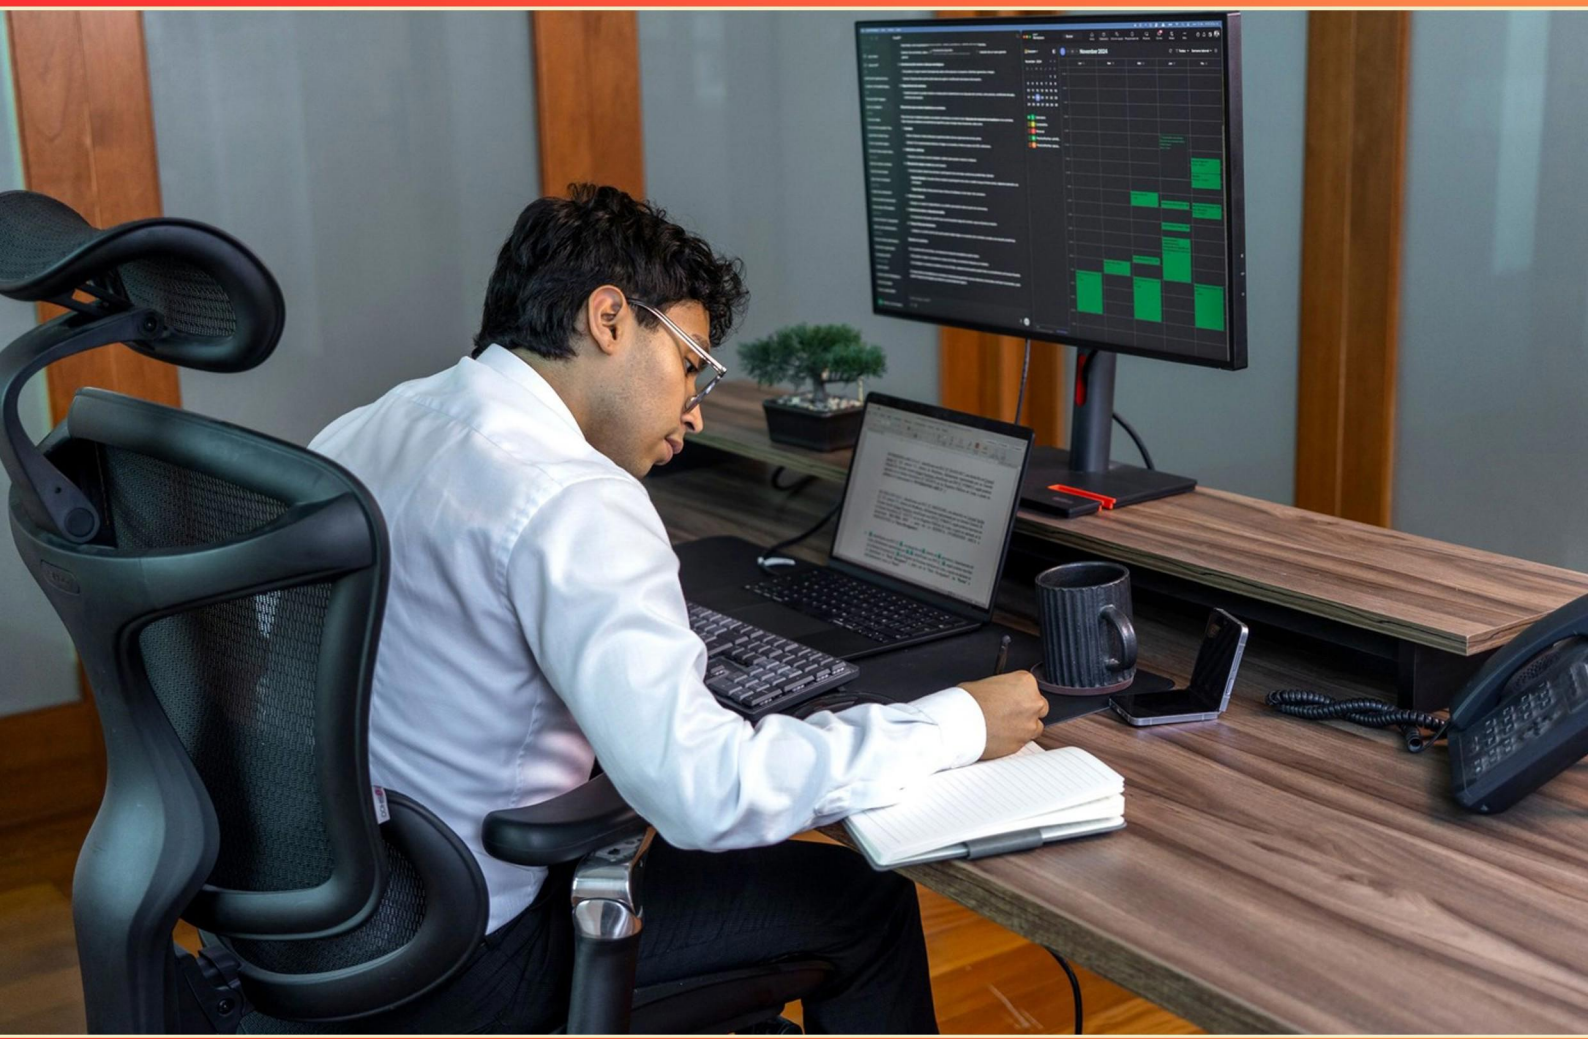
DP-700 MICROSOFT FABRIC DATA ENGINEER PRACTICE EXAM (SAMPLE)

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.

SAMPLE

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!

SAMPLE

Questions

SAMPLE

1. 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

2. Time travel in a dataset allows access to historical versions. What is the primary benefit?

A. It speeds up current query performance.

B. It eliminates the need for data governance.

C. It creates duplicate copies for every change.

D. It allows access to historical versions of datasets through time travel to reproduce or rollback analytics.

3. How would you describe a lakehouse table stored on OneLake enabling BI-friendly SQL with ACID transactions?

A. A NoSQL key-value store

B. A relational database managed separately

C. A CSV file

D. A Delta Lake table on OneLake enabling BI-friendly SQL queries with ACID transactions

4. Which of the following is NOT a part of the Dataflow Gen2 UI in Microsoft Fabric?

A. Data Preview pane

B. Query Settings pane

C. Query Development canvas

D. Power Query ribbon

5. Which action enables interactive Spark exploration of a file in the lakehouse?

A. Open a notebook to run Spark interactively.

B. Switch to the SQL analytics endpoint mode.

C. Create a Dataflow Gen2.

D. Publish a dashboard.

SAMPLE

6. What data retention strategies are typical in Fabric?

A. Define retention policies per dataset, archive or purge stale data, support point-in-time restore, and implement lifecycle management to balance cost and compliance.

B. Retain all data forever with no policy.

C. Purge everything after a fixed period regardless of usage.

D. Retention is not configurable in Fabric.

7. Where are lakehouse Delta tables stored?

A. On OneLake storage

B. In a separate cloud provider's storage

C. In local on-premises storage only

D. In memory

8. Which component in Microsoft Fabric is used to trigger downstream processes in response to events?

A. Activator

B. Pipelines

C. Notebooks

D. Real-Time Intelligence

9. What is the role of Git in the CI/CD process in Fabric?

A. Git is used to deploy code changes to different environments.

B. Git is used for automated testing of code changes.

C. Git automatically consolidates pipelines across projects.

D. Git lets your team collaborate using branches, and provides version control. It helps manage incremental code changes, and see code history.

10. What is the primary advantage of Dynamic Data Masking (DDM)?

A. It changes the actual data in the database.

B. It improves query performance.

C. It limits data exposure by obscuring sensitive information in real time.

D. It requires complex coding to implement.

SAMPLE

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. 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.

2. Time travel in a dataset allows access to historical versions. What is the primary benefit?

A. It speeds up current query performance.

B. It eliminates the need for data governance.

C. It creates duplicate copies for every change.

D. It allows access to historical versions of datasets through time travel to reproduce or rollback analytics.

Access to historical versions through time travel provides a reliable way to reproduce analyses and revert to earlier data states. This capability allows you to pull the exact dataset as it existed at a specific point in time, so you can rerun analytics, verify results, or compare how data has changed over time. It also supports undoing mistakes by rolling back to a known good state, which is essential for debugging, auditing, and maintaining data integrity. The other options don't capture this core benefit: time travel isn't primarily about speeding up current queries, it doesn't replace data governance, and it doesn't automatically create duplicates for every change.

3. How would you describe a lakehouse table stored on OneLake enabling BI-friendly SQL with ACID transactions?

A. A NoSQL key-value store

B. A relational database managed separately

C. A CSV file

D. A Delta Lake table on OneLake enabling BI-friendly SQL queries with ACID transactions

Using a Delta Lake table on OneLake gives you a lakehouse that supports BI-friendly SQL with ACID transactions. Delta Lake adds ACID guarantees, scalable metadata management, and data versioning on top of cloud storage, so you can perform reliable, concurrent reads and writes and run SQL-based queries for business intelligence without worrying about partial writes. On OneLake, this means you query the data with familiar SQL tools while preserving transactional integrity, combining the openness of a data lake with the reliability and queryability you expect from a database. The other options don't fit because a NoSQL key-value store typically lacks full SQL access and strong ACID behavior at scale; a plain CSV file has no transactional guarantees or optimized BI querying; and a relational database managed separately isn't leveraging the lake storage pattern and the Delta Lake capabilities.

SAMPLE

4. Which of the following is NOT a part of the Dataflow Gen2 UI in Microsoft Fabric?

A. Data Preview pane

B. Query Settings pane

C. Query Development canvas

D. Power Query ribbon

Dataflow Gen2's UI is organized around focused panels and a command ribbon, not a freeform development canvas. The Data Preview pane is where you view sample data as you shape your transformations, giving you immediate feedback on how changes will look in the dataset. The Query Settings pane exposes configuration options and the current query properties so you can control source details, load behavior, privacy, and the sequence of applied steps. The Power Query ribbon provides the standard set of commands you'd expect from Power Query—tools to apply transforms, manage steps, and access the M language editor. The element that isn't part of the Dataflow Gen2 UI is a dedicated Query Development canvas. In Gen2, you don't develop the query on a visual canvas; instead you use the steps list and, if needed, the Advanced Editor to modify the M code. That's why the development canvas isn't part of the interface.

5. Which action enables interactive Spark exploration of a file in the lakehouse?

A. Open a notebook to run Spark interactively.

B. Switch to the SQL analytics endpoint mode.

C. Create a Dataflow Gen2.

D. Publish a dashboard.

Interactive Spark exploration is enabled by opening a notebook to run Spark interactively. In a lakehouse environment, a notebook starts a Spark session you can use with PySpark, Scala, or Spark SQL to read the file, transform it, and inspect results step by step. This immediate feedback loop lets you experiment with filters, joins, and aggregations directly on the data. Switching to the SQL analytics endpoint mode provides a SQL-based interface for querying the lakehouse, which is great for BI tools, but it doesn't execute Spark code for interactive exploration. Creating a Dataflow Gen2 is about building data pipelines for movement and transformation, not interactive Spark experimentation. Publishing a dashboard is for visualization and sharing results, not for live, code-driven exploration of the data.

SAMPLE

6. What data retention strategies are typical in Fabric?

A. Define retention policies per dataset, archive or purge stale data, support point-in-time restore, and implement lifecycle management to balance cost and compliance.

B. Retain all data forever with no policy.

C. Purge everything after a fixed period regardless of usage.

D. Retention is not configurable in Fabric.

In Fabric, data retention is handled through policy-driven controls that apply to each dataset, balancing cost, compliance, and recoverability. Defining retention policies per dataset lets different datasets follow the specific regulatory or business needs they have. Archiving older data moves it to cheaper storage while keeping it accessible if needed, and purging stale data helps control costs and meet privacy rules. Point-in-time restore provides a way to recover data from a specific moment, which is important after mistakes or ransomware within the retention window. Lifecycle management automates moving and deleting data across storage tiers to maintain ongoing cost and compliance balance. The other options don't fit because keeping everything forever has no policy, purging everything after a fixed period overlooks value and usage, and retention is indeed configurable in Fabric.

7. Where are lakehouse Delta tables stored?

A. On OneLake storage

B. In a separate cloud provider's storage

C. In local on-premises storage only

D. In memory

Delta tables for a lakehouse are stored in OneLake storage, Microsoft Fabric's unified data lake. The actual data files and the Delta transaction log reside in OneLake, giving a durable, centrally managed storage layer that Fabric compute engines access. They aren't kept in memory, on-premises only, or in a separate cloud provider's storage—the OneLake storage location is the single home for lakehouse data across Fabric workloads.

8. Which component in Microsoft Fabric is used to trigger downstream processes in response to events?

A. Activator

B. Pipelines

C. Notebooks

D. Real-Time Intelligence

When you want something to start automatically because something happened, you're dealing with event-driven orchestration. Activator is the component in Fabric that does exactly that: it listens for events and triggers downstream work in response, such as starting pipelines or other processes when data arrives, a status changes, or a message is received. This makes workflows reactive and loosely coupled, so upstream events drive the next steps without manual intervention. Pipelines, while they orchestrate sequences of steps, aren't primarily the event trigger—they're the workflows that do the work once triggered. Notebooks are for interactive analysis and ad-hoc computation, not for automating production triggers. Real-Time Intelligence focuses on streaming insights and analytics, not on initiating downstream processes in response to events.

SAMPLE

9. What is the role of Git in the CI/CD process in Fabric?

A. Git is used to deploy code changes to different environments.

B. Git is used for automated testing of code changes.

C. Git automatically consolidates pipelines across projects.

D. Git lets your team collaborate using branches, and provides version control. It helps manage incremental code changes, and see code history.

Git in a Fabric CI/CD setup is about version control and collaboration. It serves as the single source of truth for code, allowing teams to work in parallel through branches, commits, and a complete history of changes. Each change is captured as a commit, branches let you isolate features or fixes, and pull requests enable code review and traceability. In the Fabric CI/CD workflow, pipelines respond to commits or PRs, build and test the code, and deploy, but the versioning and collaboration core comes from Git. This is why the described role fits best: it emphasizes collaboration using branches, robust version control, the management of incremental changes, and the ability to see code history. Deployment, automated testing performed within pipelines, or cross-project pipeline consolidation are handled by the CI/CD tooling rather than Git itself.

10. What is the primary advantage of Dynamic Data Masking (DDM)?

A. It changes the actual data in the database.

B. It improves query performance.

C. It limits data exposure by obscuring sensitive information in real time.

D. It requires complex coding to implement.

Dynamic Data Masking protects sensitive information by masking data at query time for non-privileged users. It obscures the actual values in real time when results are returned, without changing the data stored in the database. This lets legitimate users see usable results while helping prevent accidental exposure for others. For example, a credit card or SSN might appear masked unless the user has explicit permission to unmask. This approach preserves data integrity and avoids any need to rewrite data or deploy complex logic.

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