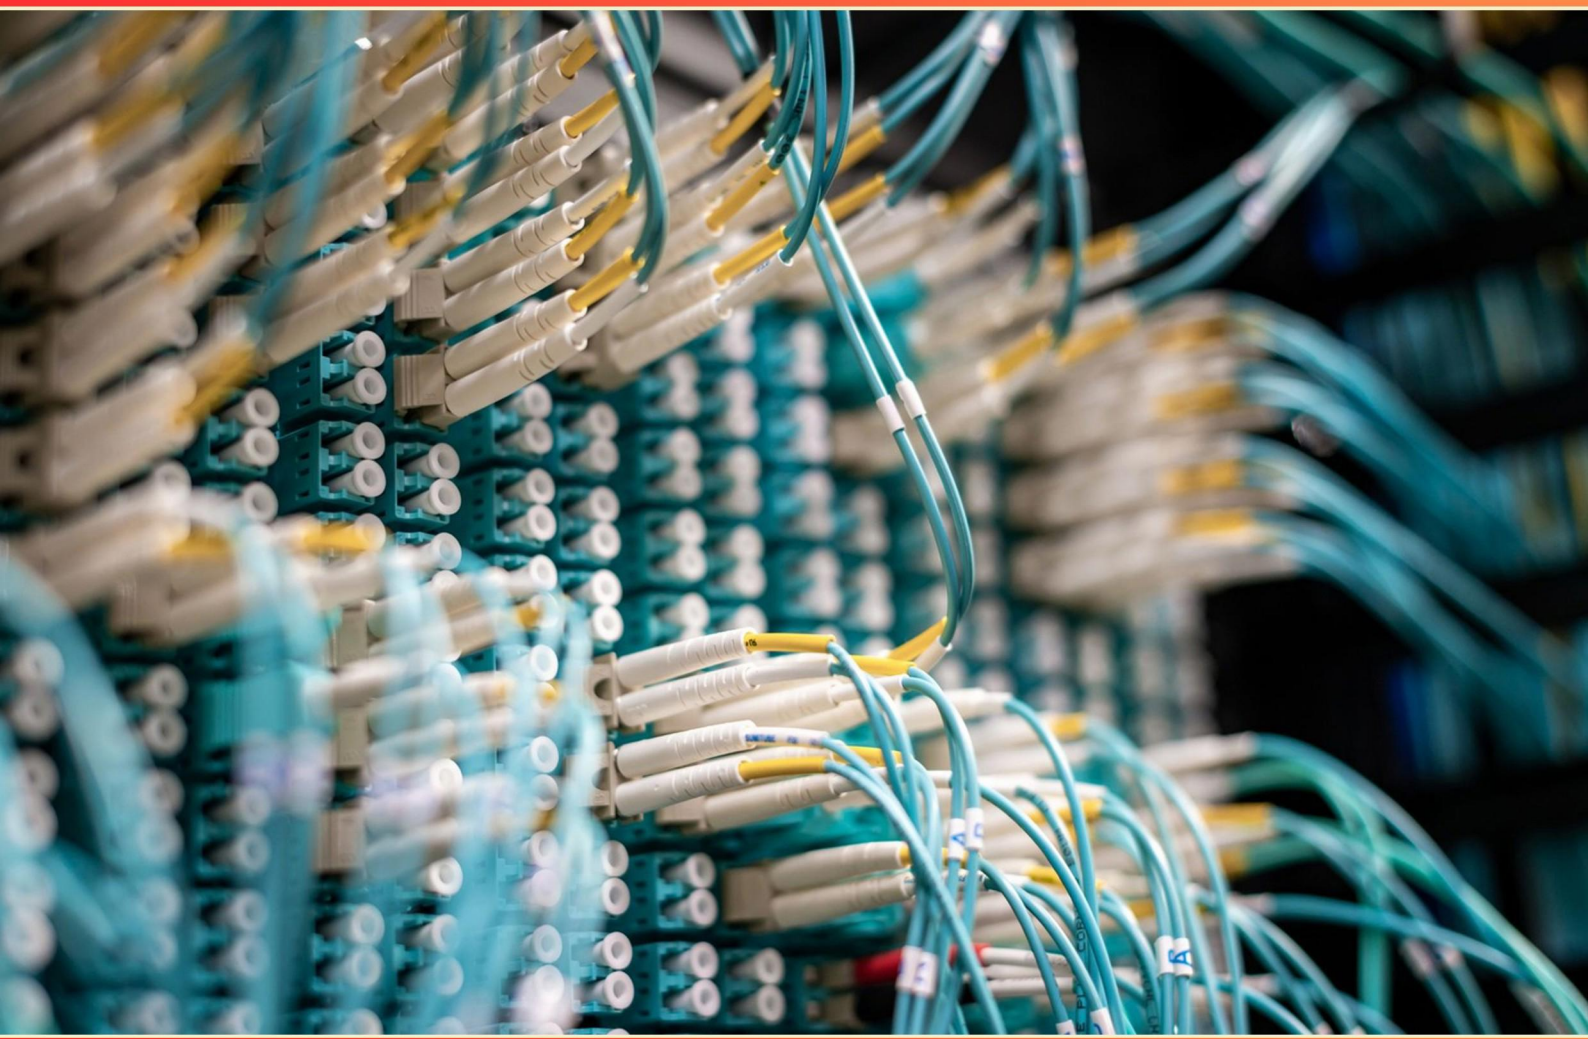
# FABRIC ANALYTICS ENGINEER ASSOCIATE PRACTICE TEST (SAMPLE)

## STUDY GUIDE



## SAMPLE STUDY GUIDE WITH LIMITED QUESTIONS

**VISIT US ONLINE FOR THE FULL VERSION STUDY GUIDE**

https://fabricanalyticsengineerassoc.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. **In the described dataflow with two lakehouses, non-folded steps are executed by which engine?**

   A. Each lakehouse's query engine

   B. The Microsoft Power Query engine

   C. The source lakehouse query engine

   D. The VertiPaq engine

2. **How can you customize the appearance of a Power BI report?**

   A. By embedding external images into the report visuals

   B. By creating a custom theme file with specified colors, fonts and formatting

   C. By applying pre-defined color schemes and fonts

   D. By adjusting the data sources within the report

3. **For ResearchReviewersGroup1, which workspace role should you assign?**

   A. Contributor

   B. Member

   C. Viewer

   D. Administrator

4. **What is the simplest way to analyze data stored as CSV using Spark in Fabric?**

   A. Import the data into a table in a warehouse.

   B. Load the file into a dataframe.

   C. Convert the data to Parquet format.

   D. Use Spark SQL to register a temporary view.

5. **When you create a clone using CREATE TABLE test.FactSales AS CLONE OF Dbo.FactSales, which statement is true?**

   A. A replica of dbo.FactSales is created in the test schema by copying the metadata only

   B. Additional schema changes to dbo.FactSales will also apply to test.factSales

   C. Additional data changes to dbo.FactSales will also apply to test.factSales

   D. The data in test.factSales is a full copy of dbo.FactSales at clone time

SAMPLE

**6. A model includes a products table and a table with monthly category targets; each product can belong to multiple categories and each category can have multiple products. Which cardinality should be set?**

A. One-to-one

B. One-to-many

C. Many-to-many

D. None of the above

**7. Where in the Fabric Admin portal can you enable Fabric environments when migrating from Azure Synapse Analytics due to permission issues?**

A. Protection Metrics

B. Workspaces

C. Tenant settings

D. Capacity settings

**8. What is the primary language used for querying a data warehouse?**

A. Python

B. SQL

C. C#

D. R

**9. To ensure each customer's sales data is written to its own Parquet file, which data pipeline configuration should you implement?**

A. Add a SecureString parameter for Customer ID.

B. Increase the Concurrency count.

C. Partition by customer ID on the customer dimension table.

D. Partition by customer ID on the fact table.

**10. For determining whether queries are mapped to aggregated data via user-defined aggregations, which single tool is most suitable for profiling queries in a semantic model?**

A. DAX Studio

B. Performance Analyzer

C. SQL Server Profiler

D. Tabular Editor

SAMPLE

# Answers

SAMPLE

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

SAMPLE

# Explanations

SAMPLE

## 1. In the described dataflow with two lakehouses, non-folded steps are executed by which engine?

A. Each lakehouse's query engine

B. The Microsoft Power Query engine

C. The source lakehouse query engine

D. The VertiPaq engine

*Non-folded steps are transformations that can't be translated into the lakehouse's native query language and therefore can't be pushed down to the data source. In a dataflow with two lakehouses, these steps must be executed by the Microsoft Power Query engine, which runs the M-based transformations as data moves through the flow. This engine handles the remaining work after pulling data from the sources, ensuring the requested operations are completed even when the source engines can't fold them. The lakehouse query engines would perform folding for steps that they can translate into their native queries, and the VertiPaq engine is the in-memory execution engine used for Power BI data models—not the place where non-folded transformations in a dataflow are executed.*

## 2. How can you customize the appearance of a Power BI report?

A. By embedding external images into the report visuals

B. By creating a custom theme file with specified colors, fonts and formatting

C. By applying pre-defined color schemes and fonts

D. By adjusting the data sources within the report

*Customizing a Power BI report's appearance is best done with a theme file. A custom theme is a JSON file that defines the color palette, fonts, and formatting rules Power BI applies across visuals and pages, so you get a consistent look that matches branding. By importing this theme, you set default colors for charts, fonts for titles and labels, and other formatting so every element follows the same style without tweaking each one individually. This approach is more comprehensive than embedding external images, which only affect specific visuals, and more flexible than using only the built-in pre-defined schemes that may not align with your branding. Changing data sources affects what the report shows, not how it looks.*

## 3. For ResearchReviewersGroup1, which workspace role should you assign?

A. Contributor

B. Member

C. Viewer

D. Administrator

*The main idea here is to apply least-privilege access: give the group only what they need to do their job. For a ResearchReviewersGroup, the task is to review materials without making changes. A read-only view right allows members to see everything they need to review while preventing edits or accidental edits to content or settings. That's why a Viewer is the best fit: it provides visibility without modification rights. Other roles grant additional powers—Contributor can edit or add content, Member typically includes collaboration capabilities, and Administrator has full control over settings and users—so they aren't appropriate for a group whose job is strictly to review. If feedback or comments are required, some systems offer a separate commenting capability, which can be granted without full editing rights.*

SAMPLE

**4. What is the simplest way to analyze data stored as CSV using Spark in Fabric?**

A. Import the data into a table in a warehouse.

B. Load the file into a dataframe.

C. Convert the data to Parquet format.

D. Use Spark SQL to register a temporary view.

*The main idea is that Spark's DataFrame provides the most straightforward and powerful entry point for analyzing structured data like CSV. Loading the CSV into a DataFrame gives you a tabular, column-named representation with optional schema inference, so you can immediately perform filters, aggregations, joins, and transformations using concise, expressive APIs. In Fabric, you'd typically read the CSV with spark.read.csv (with header and, if desired, inferSchema options), and you gain immediate access to the data for analysis. If you later want to use SQL, you can register a temporary view from that DataFrame, but the essential first step is bringing the data into a DataFrame. Other approaches add extra steps or constraints. Importing into a warehouse means moving the data into a separate storage system with its own table and schema management, which is heavier for quick analysis. Converting to Parquet changes the data format and requires a conversion step before you can analyze it, adding overhead. Using Spark SQL to register a temporary view is a useful technique, but it's an optional layer on top of having the data in a DataFrame, not the simplest starting point for analysis.*

**5. When you create a clone using CREATE TABLE test.FactSales AS CLONE OF Dbo.FactSales, which statement is true?**

A. A replica of dbo.FactSales is created in the test schema by copying the metadata only

B. Additional schema changes to dbo.FactSales will also apply to test.factSales

C. Additional data changes to dbo.FactSales will also apply to test.factSales

D. The data in test.factSales is a full copy of dbo.FactSales at clone time

*Zero-copy cloning creates a new table that mirrors the source's structure at the moment of the clone, but does not duplicate the actual data. The clone and the source share the same underlying storage initially, so no full data copy is made; only the metadata (table shape, columns, types) is captured for the new table. Because the clone's metadata is established at clone time, subsequent schema changes to the source don't automatically propagate to the clone unless you apply them there as well. Likewise, data changes to the source don't automatically appear in the clone—each table can diverge once either is modified, via copy-on-write behavior. Therefore, the statement that this creates a replica by copying metadata only is the right description: you get a metadata-level replica without copying the data. The other statements are not true because schema changes don't auto-propagate, data changes don't automatically apply, and there isn't a full data copy made at clone time.*

SAMPLE

## 6. A model includes a products table and a table with monthly category targets; each product can belong to multiple categories and each category can have multiple products. Which cardinality should be set?

A. One-to-one

B. One-to-many

C. Many-to-many

D. None of the above

*This question tests how to model relationships between entities in a relational database. When an entity in one table can be linked to many records in another, and the reverse is also true, you have a many-to-many relationship. Here, a product can belong to multiple categories and a category can contain multiple products, so the link between products and categories is many-to-many. Implementing that typically uses a junction table that stores pairs of product_id and category_id (and possibly extra attributes). The monthly targets table adds another dimension, but it doesn't change the fact that the product–category link itself is many-to-many; it can be modeled alongside with its own tables as needed. Therefore, the appropriate cardinality is many-to-many.*

## 7. Where in the Fabric Admin portal can you enable Fabric environments when migrating from Azure Synapse Analytics due to permission issues?

A. Protection Metrics

B. Workspaces

C. Tenant settings

D. Capacity settings

*Global admin controls at the tenant level govern features that affect the entire organization, including enabling Fabric environments. When you're migrating from Azure Synapse Analytics and run into permission issues, the setting you need to adjust is a tenant-level toggle or policy that enables Fabric environments for the whole tenant. This is why the correct place is Tenant settings—the option sits above individual workspaces or capacity controls and changes how environments can be provisioned across the entire organization. Protection Metrics, Workspaces, and Capacity settings are focused on security metrics, per-workspace resource management, and compute capacity, respectively, and don't control the organization-wide ability to enable new environments for migrations. If you don't see the toggle, you likely need admin access to the tenant settings to make this change.*

## 8. What is the primary language used for querying a data warehouse?

A. Python

B. SQL

C. C#

D. R

*SQL is the language used to query a data warehouse because it's the standard, expressive way to retrieve and analyze structured data stored in tables. It lets you declaratively specify what data you want—filters, joins across tables, groupings, and aggregations—while the database engine handles the best way to execute those operations efficiently. Data warehouses are built around relational structures like facts and dimensions, and SQL provides the powerful constructs (SELECT, JOIN, GROUP BY, HAVING, window functions) needed for analytics across large datasets. Other languages like Python, R, or C# are great for scripting, data processing, and application code, but the actual querying work in a data warehouse is fundamentally done with SQL.*

SAMPLE

**9. To ensure each customer's sales data is written to its own Parquet file, which data pipeline configuration should you implement?**

A. Add a SecureString parameter for Customer ID.

B. Increase the Concurrency count.

C. Partition by customer ID on the customer dimension table.

D. Partition by customer ID on the fact table.

*Partitioning by the customer ID on the fact table ensures each customer's sales records are written into their own Parquet files. The fact table holds the transactional data and the customer ID ties every row to a specific customer, so configuring the pipeline to partition on that key creates separate storage partitions (and thus separate Parquet files) per customer as the data is written. This setup directly controls how the write output is organized, making per-customer files, which is exactly the goal. Partitioning on the dimension table would affect only the dimension data, not the sales facts being written, so it wouldn't guarantee separate Parquet files per customer. Increasing the concurrency count changes how many files can be written in parallel but doesn't ensure per-customer file separation. Adding a SecureString parameter changes configuration or security controls, not how the write output is partitioned.*

**10. For determining whether queries are mapped to aggregated data via user-defined aggregations, which single tool is most suitable for profiling queries in a semantic model?**

A. DAX Studio

B. Performance Analyzer

C. SQL Server Profiler

D. Tabular Editor

*The key idea is to observe exactly how a semantic model processes queries at the server level to see if they are routed to aggregated data via user-defined aggregations. SQL Server Profiler is the most suitable single tool for this because it hooks into the server-side engine (Analysis Services or SQL Server) and captures real-time traces of the executed queries, their texts (DAX/MDX for semantic models), timings, and the components involved. With these traces, you can identify whether a given query is being fulfilled by the aggregated structures defined in the model, showing the path the engine takes when UDAs come into play. DAX Studio specializes in profiling DAX queries and is excellent for detailed client-side query analysis, but UDAs operate at the storage/engine level. Performance Analyzer in Power BI gives you in-report timings and visuals-level queries, which is useful for diagnosing slow visuals but doesn't provide the full server-side trace needed to confirm aggregation routing. Tabular Editor focuses on editing the model's metadata and definitions rather than profiling query execution.*

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://fabricanalyticsengineerassoc.examzify.com

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

SAMPLE