


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.

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. If a data source is only available in narrow time windows, which dataflow design best ensures the source is copied as-is and usable for later transformations?**
 - A. Create a linked table referencing the data from another dataflow.
 - B. Create a shared dataset that can be reused by multiple Power BI reports.
 - C. Create a staging dataflow that will only copy the data from the source as-is.
 - D. Create a transformation dataflow that will apply all necessary transformations.
- 2. Which requirement describes how Research division workspaces should be organized to support OneLake data hub filtering based on department name?**
 - A. Integrate with Sales division
 - B. Group logically for OneLake data hub filtering
 - C. Use per-minute billing for all workspaces
 - D. Enable Lakehouse Explorer access for all users
- 3. When choosing between filtering on a dimension table versus a fact table, which yields faster performance?**
 - A. Filtering on a dimension table
 - B. Filtering on a fact table
 - C. Filtering on a calculated column
 - D. Filtering on a measure
- 4. In numeric data summaries, which option correctly states an additional function included beyond COUNT, MEAN, and STD?**
 - A. MIN and MAX
 - B. TOP and UNIQUE
 - C. MEDIAN and MODE
 - D. SUM and PRODUCT
- 5. To enforce RLS with a dynamic rule that restricts access to sales regions per user, what should you create?**
 - A. Create an OLS role and use a dynamic rule.
 - B. Create an RLS role and use a dynamic rule.
 - C. Create an RLS role and use a static rule.
 - D. Create an OLS role and use a static rule.

6. Which action would you take to reference external Gen2 data in the lakehouse without copying it?
- A. Create a shortcut.
 - B. Create a Dataflow that queries it directly.
 - C. Copy the data into a staging folder.
 - D. Create a notebook to read it.
7. To group workspaces for certain company divisions that should support OneLake data hub filtering, which pairing of grouping method and tool is correct?
- A. Capacity / The Fabric Admin portal
 - B. Domain / The Fabric Admin portal
 - C. Tenant / OneLake data hub
 - D. Domain / The Microsoft Entra admin center
8. In OneLake security, which permission should you assign to User2 to meet the requirement 'Read all SQL endpoint data'?
- A. Read
 - B. ReadAll
 - C. ReadSpark
 - D. ReadData
9. What is the benefit of storing different layers of your lakehouse in separate workspaces?
- A. It can enhance security, manage capacity use, and optimize cost-effectiveness.
 - B. It makes it easier to share data with colleagues.
 - C. There's no benefit of storing different layers of your lakehouse in separate workspaces.
 - D. It reduces data redundancy.
10. In the SQL example, what are the Ranking values for the first two records?
- A. 0 and 1
 - B. 1 and 1
 - C. 1 and 2
 - D. 10 and 9

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. If a data source is only available in narrow time windows, which dataflow design best ensures the source is copied as-is and usable for later transformations?
- A. Create a linked table referencing the data from another dataflow.
 - B. Create a shared dataset that can be reused by multiple Power BI reports.
 - C. Create a staging dataflow that will only copy the data from the source as-is.**
 - D. Create a transformation dataflow that will apply all necessary transformations.

When a data source is only available in narrow time windows, you want a precise, untouched copy of what the source produced at that moment so you can use it later for any needed transformations. A staging dataflow is built for that purpose: it ingests the source and copies it as-is into a staging area, creating a persistent, raw dataset that can be transformed later without re-accessing the live source. This decouples data capture from processing, allowing you to validate and re-process the data as needed, even if the source window is closed. It also gives you a reproducible baseline to apply transformations when you're ready, without risking loss or drift from subsequent source changes. Other designs don't fit as well. A linked table depends on another dataflow and may not provide a true, standalone snapshot. A shared dataset is about consumption by reports, not about preserving a raw ingestion. A transformation dataflow applies changes during load, which means you don't keep the original as-is copy for later use.

2. Which requirement describes how Research division workspaces should be organized to support OneLake data hub filtering based on department name?
- A. Integrate with Sales division
 - B. Group logically for OneLake data hub filtering**
 - C. Use per-minute billing for all workspaces
 - D. Enable Lakehouse Explorer access for all users

Grouping workspaces logically by department and related functions is the best way to support OneLake data hub filtering based on department name. When the Research division's workspaces are organized in a way that mirrors the department structure and are consistently labeled with the department name, the data hub can apply filters cleanly and predictably. This alignment makes it straightforward to present or restrict data to the appropriate department, enhances governance and policy enforcement, and simplifies discovery for users who need department-specific data. Other options don't directly enable department-based filtering. Integrating with another division doesn't create the department-aligned structure the filter relies on. Per-minute billing is about cost rather than data organization or access control. Enabling access for all users affects visibility but doesn't establish the department-based filtering the hub uses to segment data.

3. When choosing between filtering on a dimension table versus a fact table, which yields faster performance?

- A. Filtering on a dimension table**
- B. Filtering on a fact table
- C. Filtering on a calculated column
- D. Filtering on a measure

Pruning data early by filtering on dimension tables tends to be faster because dimension tables are smaller and offer selective predicates that quickly narrow down the relevant rows. When you filter on a dimension, the engine can identify the subset of dimension rows that meet the predicate and then use that reduced set to join to the large fact table, cutting down the number of fact rows that must be read and processed. This early reduction often enables the use of indexes and efficient joins, which speeds up the query. Filtering on the large fact table would typically require scanning many more rows unless there are highly selective indexes or partitions on the filter column. Filtering on a calculated column adds extra per-row computation, which can prevent index usage and slow down the plan. Filtering on a measure is usually applied after aggregation, so it doesn't prune the rowset efficiently.

4. In numeric data summaries, which option correctly states an additional function included beyond COUNT, MEAN, and STD?

- A. MIN and MAX**
- B. TOP and UNIQUE
- C. MEDIAN and MODE
- D. SUM and PRODUCT

When summarizing numeric data, you want to convey not just the center and spread but also the range of values. COUNT gives how many observations, MEAN gives the average, and STD shows variability around that average. The natural next piece is the extreme values: the smallest and largest observations, i.e., the minimum and maximum. Including these two bounds completes a simple, informative snapshot of the data, showing where the data sit on the scale and how wide the spread can be. Other options introduce metrics that aren't as universally part of a basic descriptive summary—MEDIAN and MODE can be useful but aren't as universally paired with COUNT, MEAN, and STD, and TOP/UNIQUE or SUM/PRODUCT aren't standard additions to that particular trio.

5. To enforce RLS with a dynamic rule that restricts access to sales regions per user, what should you create?

- A. Create an OLS role and use a dynamic rule.
- B. Create an RLS role and use a dynamic rule.**
- C. Create an RLS role and use a static rule.
- D. Create an OLS role and use a static rule.

Row-Level Security lets you filter data rows based on who is querying. To enforce access that varies by user (restricting sales regions per user), you need a mechanism that adapts to the identity of the person asking. An RLS role provides the per-row filtering capability, and a dynamic rule makes that filter depend on the current user's identity. In practice, the dynamic rule uses the user's login (or a mapped attribute) to determine which regions are allowed and then only returns rows matching those regions. This per-user mapping is what enables true regional access control, rather than a one-size-fits-all filter. Using an Object-Level Security role wouldn't filter individual rows, only access to the object as a whole, so it wouldn't enforce per-row regional restrictions. A static rule would apply the same filter to everyone, which fails to differentiate access by user. So combining a Row-Level Security role with a dynamic rule is the correct approach to achieve per-user region restrictions.

6. Which action would you take to reference external Gen2 data in the lakehouse without copying it?

- A. Create a shortcut.**
- B. Create a Dataflow that queries it directly.
- C. Copy the data into a staging folder.
- D. Create a notebook to read it.

Creating a shortcut is the way to reference external Gen2 data in the lakehouse without copying it. A shortcut acts as a metadata pointer that links to the data's real location in Gen2, letting the lakehouse surface and query the data as if it were stored nearby while the actual data stays in place. Because you aren't duplicating the data, you keep a single source of truth, avoid extra storage, and ensure the data remains up to date with its source. When queries run, the engine reads directly from the external Gen2 location through the pointer, enabling seamless analysis without moving data. The other options either copy the data, which violates the requirement, or don't establish a persistent lakehouse-level reference that supports governance and reuse.

7. To group workspaces for certain company divisions that should support OneLake data hub filtering, which pairing of grouping method and tool is correct?

- A. Capacity / The Fabric Admin portal
- B. Domain / The Fabric Admin portal**
- C. Tenant / OneLake data hub
- D. Domain / The Microsoft Entra admin center

Grouping by domain matches how the company is structured, which lets OneLake data hub filtering apply per division. The Fabric Admin portal is the right place to configure this—it's where you manage workspaces, assign them to domains, and enforce domain-based filtering for OneLake. Capacity and tenant are resource or broader-scope boundaries that don't align with division-level filtering, and the Microsoft Entra admin center handles identities rather than organizing workspaces for OneLake filtering. So the domain-based grouping implemented through the Fabric Admin portal is the correct pairing.

8. In OneLake security, which permission should you assign to User2 to meet the requirement 'Read all SQL endpoint data'?

- A. Read
- B. ReadAll
- C. ReadSpark
- D. ReadData**

This tests how permissions map to specific data endpoints in OneLake. To meet the requirement of reading all SQL endpoint data, you need the permission that specifically enables access to data exposed by SQL endpoints. ReadData is designed for reading data via SQL endpoints, providing exactly that scope without granting broader access to other data types or endpoints. Read would be too generic, ReadAll could grant broader access beyond SQL endpoints, and ReadSpark is for Spark endpoints. So ReadData best satisfies the need to read all SQL endpoint data.

9. What is the benefit of storing different layers of your lakehouse in separate workspaces?

A. It can enhance security, manage capacity use, and optimize cost-effectiveness.

B. It makes it easier to share data with colleagues.

C. There's no benefit of storing different layers of your lakehouse in separate workspaces.

D. It reduces data redundancy.

Separating different layers of your lakehouse into separate workspaces lets you control access, manage resource usage, and optimize costs. By isolating layers, you can grant permissions narrowly to sensitive data and prevent unnecessary access, boosting security. It also makes capacity management clearer because compute, storage, and quotas can be allocated per workspace, so heavy workloads on one layer don't derail others, leading to more predictable performance. Finally, cost-effectiveness comes from being able to track and constrain spending per workspace, apply appropriate storage tiers, and avoid cross-layer budget spillover. While sharing data can be easier to arrange when needed, separation is not primarily about easier sharing. And storing layers separately does not inherently reduce data redundancy; duplication is more about data design and workflows than workspace boundaries.

10. In the SQL example, what are the Ranking values for the first two records?

A. 0 and 1

B. 1 and 1

C. 1 and 2

D. 10 and 9

Ranking with window functions assigns a position to each row based on the ORDER BY clause inside OVER. The first row gets 1, the second gets 2, and so on when there aren't ties that alter the sequence. In the given example, the top two records have distinct sort values, so their ranking appears as 1 for the first row and 2 for the second. If there were ties for the top value, a true RANK would show 1 and 1 for the first two, with a gap before the next rank, whereas ROW_NUMBER would still yield 1 and 2 because it's a strict sequential numbering. The 1 and 2 outcome thus reflects a non-tied, consecutive ranking for the first two records.

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