

Microsoft Certified: Azure Data Fundamentals (DP-900) Practice Test (Sample)

Study Guide



Everything you need from our exam experts!

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	16

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. What best defines batch data?**
 - A. A load of data that is continuous**
 - B. A load of data with a beginning and end**
 - C. A type of data primarily focused on real-time analysis**
 - D. A method of handling data with no defined structure**
- 2. What is the role of a Security Principal in Azure RBAC?**
 - A. An object representing available resources**
 - B. An object representing a user or group requesting access**
 - C. An object defining the level of access**
 - D. An object that manages access permissions**
- 3. Which of the following best describes what data includes?**
 - A. Only numerical information used in analysis**
 - B. Text descriptions and qualitative observations**
 - C. A combination of facts and figures that assist decision making**
 - D. Only structured data formatted for database entry**
- 4. What is a common use case for non-relational data?**
 - A. Handling high-volume transactional data**
 - B. Managing financial records with strict schemas**
 - C. Modeling social interactions in web and mobile applications**
 - D. Creating traditional relational database applications**
- 5. What is considered the cheapest method for storing data in Azure?**
 - A. File Storage**
 - B. Blob Storage**
 - C. SQL Database**
 - D. Queue Storage**
- 6. How does Azure address security for its databases?**
 - A. By requiring a public IP address for all connections**
 - B. By allowing public network access only**
 - C. By managing permissions at the database level**
 - D. By using encryption exclusively**

7. Which of the following is NOT a building block of Power BI?

- A. Reports**
- B. Apps**
- C. Dashboards**
- D. Tiles**

8. If you want to keep earlier versions of a blob, what feature should you enable?

- A. Snapshots**
- B. Versioning**
- C. Soft Delete**
- D. Access Control**

9. Which of the following describes a characteristic of Blob Storage?

- A. It serves structured data only**
- B. It is suitable for backup and recovery**
- C. It has to be accessed through SQL queries**
- D. It is only accessible through Azure Portal**

10. In a relational database, what enforces data integrity?

- A. Data types and relationships between tables**
- B. Flexible schemas and unstructured data**
- C. The absence of primary and foreign keys**
- D. Manual data entry by users**

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. What best defines batch data?

- A. A load of data that is continuous
- B. A load of data with a beginning and end**
- C. A type of data primarily focused on real-time analysis
- D. A method of handling data with no defined structure

Batch data is characterized as a load of data that has a clear beginning and end, meaning it is collected, processed, and stored over a defined period before being analyzed or utilized. This can involve processing a large volume of data at once, typically at scheduled intervals, such as daily or weekly, rather than instantaneously. The nature of batch processing allows for the accumulation of data that can then be processed all at once, making it suitable for tasks such as reporting, data transformations, or aggregating information from various sources. This definition aligns with the concept of batch processing in data systems, where data is often gathered in bulk before any analytical or operational activities take place, making it distinct from continuous or real-time data processing, which lacks a defined endpoint and operates on an ongoing basis. Understanding this distinction is crucial in contexts where selecting the appropriate data processing strategy can significantly impact performance and outcomes.

2. What is the role of a Security Principal in Azure RBAC?

- A. An object representing available resources
- B. An object representing a user or group requesting access**
- C. An object defining the level of access
- D. An object that manages access permissions

In Azure Role-Based Access Control (RBAC), a Security Principal is fundamentally an object that represents a user, group, service principal, or managed identity that requests access to Azure resources. This means that the Security Principal is the entity that can be granted permissions to perform actions on resources within Azure. By understanding that the Security Principal symbolizes the individual or group seeking entry to a resource, it is also important to recognize how this fits within the broader Azure RBAC framework. For instance, after identifying who the Security Principal is, Azure policies can dictate what resources they can interact with and what operations they can perform based on assigned roles. The other choices, while related to Azure RBAC concepts, refer to distinct components. Objects representing resources or defining levels of access pertain to what a Security Principal interacts with, while the object managing access permissions relates more to how those permissions are structured and enforced within Azure's RBAC model. Hence, the correct understanding of a Security Principal focuses on it being the entity making access requests within this system.

3. Which of the following best describes what data includes?

- A. Only numerical information used in analysis
- B. Text descriptions and qualitative observations
- C. A combination of facts and figures that assist decision making**
- D. Only structured data formatted for database entry

The correct answer highlights that data encompasses a combination of facts and figures, which can be quantitative and qualitative, that assist in decision-making processes. This definition captures the essence of what data is meant to represent in various contexts, as it goes beyond numeric values to include various types of information that can inform and support decisions across different industries. Understanding data in this comprehensive way is essential because it can take many forms including structured data, which is organized in a predefined manner, but also unstructured data, such as text or multimedia content, which are increasingly important in modern data analysis. The emphasis on assisting decision-making underscores data's role in driving insights and outcomes in business, research, and more. On the other hand, focusing solely on numerical information, as suggested in one of the choices, limits the view of data to a narrow scope and neglects qualitative data, which can be equally valuable. Similarly, restricting the definition to only structured formats does not recognize the broader landscape of data, including unstructured data that can be generated from various sources like social media, customer feedback, or text documents. Thus, the correct answer provides a holistic view of data, reflecting its multi-faceted nature and critical role in decision-making.

4. What is a common use case for non-relational data?

- A. Handling high-volume transactional data
- B. Managing financial records with strict schemas
- C. Modeling social interactions in web and mobile applications**
- D. Creating traditional relational database applications

The selection of modeling social interactions in web and mobile applications as a common use case for non-relational data is accurate due to several factors inherent to the nature of non-relational databases, often characterized by their flexibility and ability to handle diverse data types. Non-relational databases, commonly known as NoSQL databases, excel in scenarios where data is unstructured or semi-structured and does not fit neatly into tables with fixed schemas, as is typical in relational databases. In the context of social interactions, applications often need to store and retrieve vast amounts of data that can vary widely in format—such as user profiles, posts, likes, comments, and interactions—on a scale that can grow exponentially as user engagement increases. Moreover, non-relational databases typically provide advantages in horizontal scalability. This means they can efficiently handle large volumes of concurrent reads and writes, which is crucial for applications expecting to handle a fast-paced influx of user-generated content. Additionally, non-relational databases offer various data models, such as document, key-value, graph, or wide-column stores, which can be optimized for specific queries and data retrieval patterns that arise in social applications. Overall, the capability to manage diverse data formats, combined with the need for scalability and flexibility, makes non

5. What is considered the cheapest method for storing data in Azure?

- A. File Storage
- B. Blob Storage**
- C. SQL Database
- D. Queue Storage

Blob storage is recognized as the most economical option for storing large amounts of unstructured data in Azure. This storage solution is optimized for scenarios such as backup, archival, and serving files directly to webpages or applications. The pricing model of Blob Storage is specifically designed for high capacity and low storage costs, particularly when dealing with large datasets. Blob storage offers different tiers, including hot, cool, and archive storage, allowing users to optimize costs based on access frequency. The hot tier is ideal for frequently accessed data, while the cool and archive tiers are meant for less frequently accessed data, offering even lower prices for storage. When compared to other Azure storage options, Blob Storage is particularly beneficial for applications involving large files, media storage, and big data analytics due to its scalability and cost-effectiveness. The other storage options, while useful in their own contexts—for instance, SQL Database for structured data and transactions or File Storage for shared file access—do not match Blob Storage's affordability for large-scale data storage needs.

6. How does Azure address security for its databases?

- A. By requiring a public IP address for all connections
- B. By allowing public network access only
- C. By managing permissions at the database level**
- D. By using encryption exclusively

Azure addresses security for its databases primarily by managing permissions at the database level. This means that access control is implemented through authentication and authorization mechanisms, ensuring that only specified users and applications can access the database and perform certain actions based on their assigned roles. This role-based access control helps in maintaining the principle of least privilege, reducing the risk of unauthorized data access. Effective permission management allows database administrators to set fine-tuned policies that dictate which users can read, write, update, or delete data. This layer of security is crucial for protecting sensitive data, as it allows organizations to customize their security posture according to their specific requirements and compliance standards. In contrast, requiring a public IP address for all connections, allowing public network access only, and relying exclusively on encryption do not comprehensively address the multifaceted nature of database security within Azure. Public network requirements can expose databases to potential attacks, while encryption is an important security measure but does not control access or permissions on its own. Therefore, managing permissions at the database level offers a more robust and effective approach to database security in Azure.

7. Which of the following is NOT a building block of Power BI?

- A. Reports**
- B. Apps**
- C. Dashboards**
- D. Tiles**

Power BI is a powerful tool for data visualization and business intelligence, built upon several key components that help users analyze and represent data effectively. Reports, dashboards, and tiles are all essential aspects of the Power BI environment. Reports are comprehensive collections of visualizations and data analyses, often spread across multiple pages, allowing users to dive deep into their datasets. Dashboards offer a high-level overview by bringing together multiple visualizations, known as tiles, from either one report or multiple reports into a single page for quick insights. Tiles represent the individual visualizations or metrics displayed on a report or dashboard, offering a visual summary that can be quickly interpreted. While apps are related to Power BI, they are not considered a foundational building block like the others mentioned. Instead, apps serve as a way for organizations to package and deliver Power BI content (including dashboards and reports) to a wider audience. Therefore, the correct answer highlights that apps, while useful in the Power BI ecosystem, do not constitute a foundational building block in the same context as reports, dashboards, and tiles.

8. If you want to keep earlier versions of a blob, what feature should you enable?

- A. Snapshots**
- B. Versioning**
- C. Soft Delete**
- D. Access Control**

To retain earlier versions of a blob in Azure, the appropriate feature to enable is versioning. Blob versioning allows you to automatically save and maintain multiple versions of a blob as it changes over time. Whenever a blob is modified or overwritten, a new version is created, which means that you can access and restore previous versions whenever necessary. This is particularly useful for scenarios where you might need to roll back to an earlier state of your data due to accidental deletions, overwrites, or any other form of data loss. In contrast, other features serve different purposes: Snapshots provide a point-in-time view of a blob, allowing you to create a read-only copy of the blob as it exists at that moment, but they do not automatically enable the retention of multiple versions as versioning does. Soft delete protects blobs from accidental deletion by preserving them for a configurable retention period, allowing recovery, but it does not enable version management or multiple versions of the same blob. Access control governs permissions for who can perform actions on the blobs but does not affect the storage of multiple versions. By utilizing versioning, users can efficiently manage historical data changes, ensuring that earlier versions are preserved and can be accessed when needed, supporting data integrity and recovery strategies in Azure Storage

9. Which of the following describes a characteristic of Blob Storage?

- A. It serves structured data only**
- B. It is suitable for backup and recovery**
- C. It has to be accessed through SQL queries**
- D. It is only accessible through Azure Portal**

Blob Storage is indeed suitable for backup and recovery, as it is designed to store large amounts of unstructured data, including documents, images, videos, and backups. This storage solution is highly durable and easily scalable, making it ideal for storing data that needs to be retained and retrieved later, such as backups of critical systems or application data. It supports various redundancy options to ensure data durability, which further enhances its use for backup and disaster recovery scenarios. In contrast to the other responses, Blob Storage does not exclusively handle structured data; it is specifically built for unstructured data, making it versatile for a wide range of data types. Additionally, it does not require access through SQL queries, as it supports REST APIs and client libraries, enabling access through various programming languages instead of just structured query language. Lastly, while Azure Portal provides one way to access Blob Storage, it can also be accessed programmatically through SDKs, tools like Azure Storage Explorer, and other interfaces, making it more flexible than being limited to the Azure Portal alone.

10. In a relational database, what enforces data integrity?

- A. Data types and relationships between tables**
- B. Flexible schemas and unstructured data**
- C. The absence of primary and foreign keys**
- D. Manual data entry by users**

In a relational database, data integrity is primarily enforced through the use of data types and the establishment of relationships between tables. Data types ensure that the data entered into a field conforms to specific standards (such as integer, string, date, etc.), which helps prevent errors and inconsistencies in the data. Relationships between tables are established through the use of primary keys and foreign keys. A primary key uniquely identifies each record in a table, ensuring that there are no duplicate entries. Foreign keys create links between tables, enforcing the rules that govern how data in one table relates to data in another. This helps maintain consistency across the database and ensures that relationships between entities are preserved, which is critical for data integrity. In contrast, options that mention flexible schemas and unstructured data do not contribute to enforcing integrity within a relational database, as such databases are designed to work with structured data and predefined schemas. Additionally, the absence of primary and foreign keys would lead to a lack of relationships and potential inconsistencies, undermining data integrity. Lastly, manual data entry can introduce human error, which can also detract from data integrity. Thus, the establishment of data types and relationships is fundamental in maintaining accurate and reliable data within relational databases.

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://mscertifiedazuredatafundamentals.examzify.com>

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

SAMPLE