

MICROSOFT CERTIFIED: AZURE DATA FUNDAMENTALS (DP-900) PRACTICE TEST (SAMPLE) STUDY GUIDE



Prepare with structure. Pass with confidence



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. When is it advisable to use Azure Analysis Services?**
 - A. For very high volumes of data
 - B. For smaller volumes of data and correlation
 - C. For streaming data processing
 - D. For data archiving
- 2. What are cognitive analytics based on?**
 - A. Simple statistical models
 - B. Machine learning predictions based on previous models
 - C. Descriptive historical facts
 - D. An intuitive understanding of data
- 3. What type of data does a data lake primarily consist of?**
 - A. Structured data
 - B. Unstructured data
 - C. Organized data
 - D. Processed data
- 4. What distinguishes a clustered index from a regular index?**
 - A. It requires a copy of the indexed data
 - B. It sorts the actual table data based on the index
 - C. It improves data security and encryption
 - D. It allows for unlimited data storage capacity
- 5. Which API is NOT supported by Azure Cosmos DB?**
 - A. SQL API
 - B. GraphQL API
 - C. Table API
 - D. MongoDB API
- 6. What is the data storage capacity limit for Azure Table Storage?**
 - A. 1 TB
 - B. 5 PB
 - C. 10 TB
 - D. 100 GB

- 7. What does Azure Databricks provide on the Azure platform?**
- A. A comprehensive data warehouse
 - B. A machine learning environment for big data processing
 - C. A relational database service
 - D. A file storage solution
- 8. What is a key benefit of using SQL Server in a VM?**
- A. Limited data storage options
 - B. Always managed by Azure
 - C. No compatibility issues with on-premise SQL Server
 - D. Requires automated updates
- 9. What capability does Azure Synapse Analytics provide concerning SQL tables?**
- A. It can transform unstructured data into SQL tables
 - B. It can only create views and not tables
 - C. It does not support SQL tables
 - D. It helps to convert read data into SQL tables
- 10. Which of the following would most likely NOT be considered streaming data?**
- A. Sensor data
 - B. Transaction logs
 - C. Batch CSV files
 - D. IoT Hub data

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. When is it advisable to use Azure Analysis Services?

- A. For very high volumes of data
- B. For smaller volumes of data and correlation**
- C. For streaming data processing
- D. For data archiving

Using Azure Analysis Services is particularly advisable for scenarios that involve smaller volumes of data and where the goal is to perform complex analytics and correlation across that data. This service excels in analytical scenarios where you need a semantic layer to facilitate easy access to data across different sources. It allows businesses to create rich analytical models, which can result in meaningful insights derived from the structured data stored in databases. This is particularly beneficial for organizations that need to perform detailed analysis and reporting, using data to drive business intelligence decisions. By enabling users to explore data in a familiar manner, Azure Analysis Services simplifies data interaction and enhances the efficiency of querying against analytical models. In contrast, options that mention very high volumes of data, streaming data processing, or data archiving do not align with the core capabilities of Azure Analysis Services. While it can handle significant datasets, it is often not the best choice compared to other services optimized for big data processing or real-time streaming analytics. Similarly, it is not designed for archiving data, which is better suited to other Azure storage solutions.

2. What are cognitive analytics based on?

- A. Simple statistical models
- B. Machine learning predictions based on previous models**
- C. Descriptive historical facts
- D. An intuitive understanding of data

Cognitive analytics leverage machine learning techniques to analyze and interpret complex data patterns. This approach goes beyond simple analysis by employing algorithms that can learn from previous data and improve over time. Machine learning, a subset of artificial intelligence, enables systems to make predictions and identify trends based on vast amounts of data without being explicitly programmed for each task. In the context of cognitive analytics, the learning processes allow systems to adapt and refine their predictions based on historical data, making them particularly effective in processing unstructured data and deriving deeper insights that traditional statistical models may overlook. This capability is essential for making data-driven decisions in complex environments, where traditional analytical methods may fall short in managing the intricacies of large datasets. The other options focus on more simplistic or conventional approaches to data analysis, which do not capture the advanced, adaptive nature of cognitive analytics driven by machine learning technologies.

3. What type of data does a data lake primarily consist of?

- A. Structured data
- B. Unstructured data**
- C. Organized data
- D. Processed data

A data lake primarily consists of unstructured data, which is a key characteristic that differentiates it from more traditional data storage options like relational databases. Unstructured data can include a wide variety of formats such as text files, images, videos, audio files, and more, which do not have a predefined structure or schema when they are stored. This flexibility allows data lakes to accommodate massive volumes of diverse data types without the constraints imposed by structured formats. The ability to ingest unstructured data makes data lakes particularly valuable for big data analytics, machine learning, and other advanced processing tasks, where insights may be derived from vast amounts of raw data. As organizations increasingly rely on various forms of data beyond traditional structured datasets, data lakes serve as a central hub for storing and processing this unstructured information, enabling more comprehensive data analysis and utilization.

4. What distinguishes a clustered index from a regular index?

- A. It requires a copy of the indexed data
- B. It sorts the actual table data based on the index**
- C. It improves data security and encryption
- D. It allows for unlimited data storage capacity

A clustered index is distinct from a regular index primarily because it dictates the physical order of the data within the table itself. When a clustered index is created on a table, the rows of the table are sorted and stored on disk according to the indexed columns. This means that the data is organized in the same order as defined by the clustered index, which allows for efficient retrieval of data, especially in queries that sort or filter by the indexed columns. Because the clustered index determines how the data is stored, there can only be one clustered index per table, as the data can only be physically arranged in one order at a time. This feature of a clustered index contrasts with a non-clustered index, which maintains a separate structure that references the actual table rows and does not affect their physical order. The other options address characteristics that aren't core distinctions of clustered versus regular indexes. For instance, a copy of the indexed data might be relevant in specific contexts, but it does not define the fundamental difference. Data security and encryption relate to database protection measures, while unlimited data storage capacity is not a feature of either index type. Thus, the correct understanding of a clustered index is focused on its role in sorting and storing the actual data in the table in a

5. Which API is NOT supported by Azure Cosmos DB?

- A. SQL API
- B. GraphQL API**
- C. Table API
- D. MongoDB API

Azure Cosmos DB supports a variety of APIs to cater to different application needs, one of which includes the SQL API for querying JSON documents using SQL-like syntax. Additionally, it offers a Table API for key-value data structures, and a MongoDB API that allows applications using the MongoDB protocol to interact seamlessly with Cosmos DB. However, the GraphQL API is not natively supported by Azure Cosmos DB. While GraphQL is a popular query language for APIs, Cosmos DB has not integrated it into its supported API offerings at this time. So, it is essential to understand that while many APIs can be utilized with Azure Cosmos DB, GraphQL is not among them. This highlights the importance of checking the official Azure documentation for the latest supported features and capabilities of services like Azure Cosmos DB.

6. What is the data storage capacity limit for Azure Table Storage?

- A. 1 TB
- B. 5 PB**
- C. 10 TB
- D. 100 GB

Azure Table Storage is designed for massive scalability and is optimized for storing large amounts of structured data. The platform allows users to store petabytes of data, making it suitable for applications that require high availability and high throughput. The data storage capacity limit for Azure Table Storage is indeed 5 PB (petabytes). This means that users can leverage Azure Table Storage to store vast amounts of data without the concerns of running into capacity limits that other storage solutions might impose. This scalability is a significant advantage for organizations that deal with big data applications and need to manage large datasets efficiently. In contrast, the other options do not reflect the capabilities of Azure Table Storage. The limits mentioned in those choices fall significantly below the true storage capacity, which would not meet the needs of most enterprise data scenarios.

7. What does Azure Databricks provide on the Azure platform?

- A. A comprehensive data warehouse
- B. A machine learning environment for big data processing**
- C. A relational database service
- D. A file storage solution

Azure Databricks is an analytics service powered by Apache Spark that provides a collaborative environment primarily focused on big data processing and machine learning. It combines the capabilities of data engineering, data science, and machine learning workflows into a single platform, enabling users to process large volumes of data quickly and efficiently. The platform allows users to build machine learning models, conduct data analysis, and visualize data, leveraging the scalability of Apache Spark and the ease of integration with other Azure services. This makes it particularly suitable for scenarios where large datasets need to be processed or where complex machine learning algorithms need to be executed. Other options, while relevant in their own contexts, do not accurately capture the core function of Azure Databricks. A comprehensive data warehouse, a relational database service, and a file storage solution each serve different purposes within the Azure ecosystem. They provide specialized capabilities for data storage, management, and retrieval rather than focusing on the collaborative and integrated experience for big data processing and machine learning that Azure Databricks is designed to offer.

8. What is a key benefit of using SQL Server in a VM?

- A. Limited data storage options
- B. Always managed by Azure
- C. No compatibility issues with on-premise SQL Server**
- D. Requires automated updates

Utilizing SQL Server in a virtual machine (VM) offers the key benefit of no compatibility issues with on-premise SQL Server. Organizations that are already using SQL Server on their internal servers can seamlessly transition their databases and applications to a VM in Azure without worrying about significant changes to their existing setup. This ensures that the same SQL Server features, tools, and configurations remain available, which simplifies migration and reduces the learning curve for administrators and developers. Maintaining this compatibility supports continuity in operations and minimizes downtime or disruptions that might occur during such a transition. Given that many enterprises rely heavily on SQL Server, this aspect is crucial for ensuring that their applications work as expected when moving to a cloud environment.

9. What capability does Azure Synapse Analytics provide concerning SQL tables?

- A. It can transform unstructured data into SQL tables
- B. It can only create views and not tables
- C. It does not support SQL tables
- D. It helps to convert read data into SQL tables**

Azure Synapse Analytics offers robust capabilities regarding SQL tables, especially in the context of analytics and data integration. It allows users to ingest and process data from various sources, including files and databases, and enables the conversion of structured, semi-structured, and unstructured data into a structured format that can be represented as SQL tables. The platform's integration with data lakes and other data storage options also facilitates this transformation, allowing analysts to work with data in a more organized manner. By using Azure Synapse Analytics, users can create SQL tables that reflect data loaded into the system, allowing for efficient querying and analysis. This capability is crucial for organizations that need to make sense of vast amounts of data for reporting, analytics, or decision-making purposes. Moreover, Synapse provides seamless execution of complex queries against these tables and integration across different data processing engines, enhancing performance and accessibility. This correct choice captures the essence of how Azure Synapse Analytics excels at converting incoming data into SQL table structures for better analytical capabilities.

10. Which of the following would most likely NOT be considered streaming data?

- A. Sensor data
- B. Transaction logs
- C. Batch CSV files**
- D. IoT Hub data

Streaming data is typically characterized by its ability to be generated continuously and in real-time, allowing for immediate processing and analysis. When considering the options presented, batch CSV files stand out as the most likely candidate that does not fit into the streaming data category. Batch CSV files are collections of data that are gathered and stored at one time, rather than being generated and processed in real-time. These files represent a snapshot of data at a specific moment rather than an ongoing flow of information. This contrasts sharply with other choices like sensor data, transaction logs, and IoT Hub data, which are continuously produced and ideally suited for streaming analytics, where insights can be derived from the data as it arrives. By understanding the distinction between batch processing and streaming data, it's clear why batch CSV files do not fit the definition of streaming data, which is inherently focused on continuous data flow and real-time processing.

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