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

STUDY GUIDE



SAMPLE STUDY GUIDE WITH LIMITED QUESTIONS

VISIT US ONLINE FOR THE FULL VERSION STUDY GUIDE

https://
mscertifiedazuredatafundamentals.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 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.

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 a primary key in a relational database?

A. A unique identifier for each row in a table

B. A column that does not contain null values

C. A relationship between two tables

D. A column that can contain duplicate values

2. Which of the following best describes IaaS?

A. A cloud service model that manages the entire infrastructure for users

B. A model providing virtual infrastructure that mimics on-premises data centers

C. A platform that automatically handles database services

D. A service for managing software applications in the cloud

3. What is the default access tier in Azure Blob Storage?

A. Cool Tier

B. Archive Tier

C. Hot Tier

D. Premium Tier

4. Which approach enhances Azure database security by defining which IPs can connect?

A. Network Access Control Lists

B. IP whitelisting in the firewall

C. Role-based access control

D. Application Gateway

5. What is the availability of Azure Data Services?

A. 99.0

B. 99.5

C. 99.9

D. 99.99

6. Data processed by Azure Synapse Analytics is distributed how?

A. Randomly across nodes

B. Evenly across compute nodes

C. Based on user demand

D. Concentrated in the control node

SAMPLE

7. Which of the following elements is NOT part of an RBAC role?

A. Security Principal

B. Role Definition

C. Access Control List

D. Scope

8. How is data generally defined?

A. A collection of theories and hypotheses

B. A set of facts used for decision making

C. A statistical model for predictions

D. An abstract representation of information

9. What constitutes a key in Azure Table Storage?

A. Unique identifier across all storage accounts

B. Combination of partition identifier and row key

C. Randomly generated string

D. Solely the row key

10. Which of the following formats is considered semi-structured data?

A. SQL databases

B. JSON

C. CSV files

D. Binary data

SAMPLE

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. What is a primary key in a relational database?

A. A unique identifier for each row in a table

B. A column that does not contain null values

C. A relationship between two tables

D. A column that can contain duplicate values

A primary key serves as a unique identifier for each row in a relational database table, ensuring that no two rows can have the same key value. This unique identification is critical for maintaining data integrity and enabling easy retrieval of records. The primary key must always contain a value, meaning it cannot be null, which reinforces its role as a reliable reference point for each entry. Being established as the unique footprint of a row, the primary key guarantees that each record can be distinctly accessed and manipulated without confusion or ambiguity. This allows for efficient data operations, such as updates or deletions, enhancing the overall organization and structure of the data within the database. While there are various attributes associated with primary keys—such as the prohibition of null values—what fundamentally defines them is their role in uniquely identifying records in a table, which is pivotal for relational database management.

2. Which of the following best describes IaaS?

A. A cloud service model that manages the entire infrastructure for users

B. A model providing virtual infrastructure that mimics on-premises data centers

C. A platform that automatically handles database services

D. A service for managing software applications in the cloud

The correct choice effectively describes Infrastructure as a Service (IaaS) by emphasizing its provision of virtual infrastructure that closely resembles traditional on-premises data centers. In the IaaS model, users are able to rent virtualized computing resources over the internet, such as servers, networking features, storage, and the virtualization layer itself. This model enables users to have control over the operating systems and applications they choose to deploy, allowing them to manage their own software stack while the underlying physical infrastructure is maintained by the cloud provider. The emphasis on mimicking on-premises data centers is significant because it highlights IaaS's flexibility and scalability, allowing businesses to expand their resources as needed without the overhead associated with maintaining physical hardware. This facilitates easier resource provisioning, management of workloads, and scaling operations up or down according to business needs. In contrast, other service models do not align as well with the characteristics of IaaS. For instance, services that manage the entire infrastructure for users pertain more closely to Platform as a Service (PaaS) or Software as a Service (SaaS), where the user's control is limited or entirely abstracted away. Similarly, models that automatically handle database services also fall under PaaS, and managing software applications

SAMPLE

3. What is the default access tier in Azure Blob Storage?

A. Cool Tier

B. Archive Tier

C. Hot Tier

D. Premium Tier

In Azure Blob Storage, the default access tier is the Hot Tier. This is designed for data that is accessed frequently. Choosing the Hot Tier allows organizations to store blobs that require rapid access for applications such as web applications, mobile applications, and data analytics. The Hot Tier balances performance with cost, making it suitable for scenarios where low-latency access to data is a priority. The other tiers serve specific purposes based on data access needs. The Cool Tier is intended for infrequently accessed data that is stored for at least 30 days, which may not be suitable for scenarios where quick access is required. The Archive Tier is designed for rarely accessed data that can tolerate several hours of retrieval time, making it unsuitable for immediate access needs. The Premium Tier provides high-performance storage for workloads that demand low latency, but it is not the default tier and is typically used for specific performance-intensive scenarios. Overall, the Hot Tier is aligned with the default settings aimed at providing versatility in data management for most common use cases.

4. Which approach enhances Azure database security by defining which IPs can connect?

A. Network Access Control Lists

B. IP whitelisting in the firewall

C. Role-based access control

D. Application Gateway

The approach of IP whitelisting in the firewall effectively enhances Azure database security by controlling which IP addresses are permitted to connect to the database. This method involves specifying a list of allowed IP addresses, meaning that only these addresses can establish a connection to the database. By limiting access in this way, organizations can reduce the risk of unauthorized access, as any connection attempt from an IP address not included in the whitelist is automatically blocked. This is particularly useful in protecting sensitive data and maintaining compliance with security best practices. Other options like Network Access Control Lists and Role-based access control pertain to network configurations and user permissions rather than directly managing IP connections to the database. Application Gateway primarily serves as a web traffic load balancer and security layer for applications rather than focusing on IP connection management for databases. Therefore, IP whitelisting in the firewall stands out as the most direct method for enhancing security by regulating the specific IPs allowed access to an Azure database.

SAMPLE

5. What is the availability of Azure Data Services?

A. 99.0

B. 99.5

C. 99.9

D. 99.99

The availability of Azure Data Services is often represented by the Service Level Agreement (SLA) that guarantees a certain percentage of uptime. Microsoft Azure Data Services, including services like Azure SQL Database, Azure Cosmos DB, and others, typically offer a high level of availability, ensuring that users have reliable access to their data. In this case, the option reflecting 99.99% availability indicates a commitment to a high level of service, meaning the service will be operational and accessible 99.99% of the time over a given period. This high availability is crucial for enterprises that rely on these services for critical operations and data storage, as it minimizes downtime and enhances reliability for users and applications. Understanding the context of service availability is essential, as businesses often prioritize platforms that can ensure minimal interruptions and provide consistently strong performance. Therefore, this commitment to 99.99% uptime aligns directly with the needs of users who require dependable data services for their operations.

6. Data processed by Azure Synapse Analytics is distributed how?

A. Randomly across nodes

B. Evenly across compute nodes

C. Based on user demand

D. Concentrated in the control node

Data processed by Azure Synapse Analytics is distributed evenly across compute nodes to optimize performance and resource utilization. This approach allows for parallel processing, where multiple nodes can handle different segments of the data simultaneously. By evenly distributing the workload, Azure Synapse can leverage its scalable architecture to ensure efficient query performance and faster data processing. The even distribution also minimizes bottlenecks that can occur when one node is overwhelmed while others are under-utilized. This architecture is particularly beneficial for large datasets, allowing for a smooth and balanced workflow during data ingestion, transformation, and analysis. The efficiency gained from this method is critical for analytics tasks where speed and responsiveness are key factors. Other options may suggest alternative methods of distribution that do not align with the operational design of Azure Synapse Analytics. For instance, random distribution could lead to uneven workloads, while distribution based on user demand might not be efficient for all scenarios, and concentrating data in a control node could negate the benefits of having a distributed system in the first place.

SAMPLE

7. Which of the following elements is NOT part of an RBAC role?

A. Security Principal

B. Role Definition

C. Access Control List

D. Scope

In Role-Based Access Control (RBAC) within Azure, the elements that define an RBAC role include a security principal, role definition, and scope. A security principal is an identity that can be assigned a role, which could be a user, a group, or an application. This identity is crucial because it represents who is being granted the permissions. A role definition outlines the permissions that are granted. It specifies what actions are allowed on resources, integrating the principle of least privilege by ensuring that users obtain only the permissions necessary for their tasks. Scope refers to the boundaries of where the permissions apply. This could include specific resources, resource groups, or the entire subscription. Defining the scope is vital for controlling the extent of access and managing permissions effectively across different environments. The correct answer indicates that the Access Control List (ACL) is not an element of an RBAC role. While ACLs are used in other contexts for managing access permissions at a resource level, RBAC employs a more abstracted approach, managing permissions through roles rather than explicit lists tied to individual resources. This distinction clarifies that RBAC focuses on roles rather than granular ACLs to handle permissions efficiently.

8. How is data generally defined?

A. A collection of theories and hypotheses

B. A set of facts used for decision making

C. A statistical model for predictions

D. An abstract representation of information

Data is generally defined as a set of facts used for decision making, capturing various forms of quantitative and qualitative information that can be analyzed to derive insights. This definition emphasizes the utility of data in supporting informed decisions across various contexts, including business, science, healthcare, and more. The importance of data lies in its ability to represent measurable aspects of reality, such as sales figures, customer behaviors, survey results, or performance metrics. This factual basis provides a foundation for analysis, interpretation, and ultimately, informed decision-making processes. By harnessing data in decision making, organizations can identify trends, make predictions, and develop strategies based on empirical evidence rather than assumptions. The other options present ideas that align with certain aspects of data but do not encapsulate its general definition as comprehensively. For example, a collection of theories and hypotheses pertains more to the realm of scientific inquiry rather than raw data itself. Similarly, a statistical model for predictions involves the application of data to generate insights but does not define what data is. An abstract representation of information captures certain details about how data might be displayed but is too vague to define the essence of what data itself constitutes. Thus, the definition as a set of facts stands as the most accurate and comprehensive depiction of data in this

SAMPLE

9. What constitutes a key in Azure Table Storage?

A. Unique identifier across all storage accounts

B. Combination of partition identifier and row key

C. Randomly generated string

D. Solely the row key

In Azure Table Storage, a key is defined by a combination of the partition key and the row key. This unique combination helps in efficiently organizing and accessing data within the storage. The partition key determines the partition that the entity belongs to, which plays a crucial role in the performance and scalability of the table. By distributing entities into different partitions, Azure Table Storage can optimize data operations and increase performance. The row key, on the other hand, is unique within the defined partition. Together, the partition key and row key form a composite key that uniquely identifies each entity within the table. This structured approach allows for efficient querying and retrieval of data, as well as enhanced performance during read and write operations. Understanding this structure is fundamental for effectively leveraging Azure Table Storage in application development and data management scenarios.

10. Which of the following formats is considered semi-structured data?

A. SQL databases

B. JSON

C. CSV files

D. Binary data

Semi-structured data is characterized by a lack of a fixed schema, meaning it does not conform to the strict structure of traditional databases but still contains some organizational properties that make it easier to analyze. JSON (JavaScript Object Notation) is a prime example of semi-structured data because it uses a flexible format with key-value pairs, allowing for hierarchical organization. In JSON, data objects can vary in structure, meaning that different instances of a data entry can have different attributes or nested structures, while still maintaining the overall conventions of JSON syntax. This adaptability makes JSON ideal for representing complex data relationships and is commonly used in web applications and APIs. Other formats, such as SQL databases, possess a rigid schema that enforces structure and type for each piece of data, making them structured. CSV files, while they do allow for a simple representation of tabular data, are generally considered structured as well since they follow a clear and consistent format of rows and columns. Binary data lacks a defined structure altogether and is not human-readable, placing it outside the scope of semi-structured data. Thus, JSON stands out as the correct answer for this question.

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://mscertifiedazuredatafundamentals.examzify.com

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

SAMPLE