

AWS Academy Cloud Operations Practice Exam (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. How can you make an EC2 instance resilient to failure?**
 - A. By using Amazon RDS**
 - B. By using an Auto Scaling group across multiple Availability Zones**
 - C. By deploying a single instance**
 - D. By using a static IP address**

- 2. What data structure in Amazon DynamoDB is used to store name-value pair collections?**
 - A. String type**
 - B. List type**
 - C. Map type**
 - D. Binary type**

- 3. What is the key characteristic of a VPC?**
 - A. It automatically scales applications based on demand**
 - B. It provides a private, isolated section of the AWS Cloud**
 - C. It enables real-time data processing**
 - D. It supports containerized applications**

- 4. What features does the AWS Well-Architected Tool provide?**
 - A. API integration capabilities**
 - B. Code deployment automation**
 - C. Best practices reviews and improvement recommendations**
 - D. Cost estimation for resources**

- 5. Which database service supports petabytes of data and provides columnar storage on high-performance disks?**
 - A. Amazon RDS**
 - B. Amazon Aurora**
 - C. Amazon Redshift**
 - D. Amazon DynamoDB**

6. Which AWS service allows real-time monitoring and management of AWS resources?

- A. Amazon Aurora**
- B. Amazon CloudWatch**
- C. Amazon VPC**
- D. Amazon SQS**

7. Which AWS service simplifies the database migration process to AWS?

- A. AWS Database Migration Service (DMS)**
- B. AWS Snowball**
- C. AWS Glue**
- D. AWS Data Pipeline**

8. What is a key benefit of creating reusable infrastructure templates?

- A. Increased manual intervention**
- B. Consistency in deployments**
- C. Higher cost of maintenance**
- D. Reduced scalability**

9. What type of storage is Amazon EBS?

- A. File storage**
- B. Block storage**
- C. Object storage**
- D. Hybrid storage**

10. Which AWS service provides a scalable object storage solution for data archiving?

- A. Amazon Glacier**
- B. Amazon S3**
- C. Amazon EBS**
- D. Amazon RDS**

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. How can you make an EC2 instance resilient to failure?

- A. By using Amazon RDS
- B. By using an Auto Scaling group across multiple Availability Zones**
- C. By deploying a single instance
- D. By using a static IP address

Using an Auto Scaling group across multiple Availability Zones is an effective strategy to make an EC2 instance resilient to failure. This approach allows for the automatic scaling of instances based on demand and ensures high availability by distributing instances across different Availability Zones within a region. If one Availability Zone experiences an outage, the Auto Scaling group can launch new instances in another Availability Zone to maintain application availability. This capability provides not only redundancy but also the flexibility to handle varying loads efficiently. Other options, while related to Amazon services, do not provide the same level of resilience. For example, Amazon RDS is a managed database service and, while it can be part of an application's architecture, it does not enhance the resilience of EC2 instances directly. Deploying a single instance does not provide redundancy; if that instance fails, your application becomes unavailable. Finally, using a static IP address is primarily a networking configuration and does not inherently include features to handle instance failures or ensure continuity of service. Thus, leveraging an Auto Scaling group is the most effective method for ensuring resiliency against instance failures.

2. What data structure in Amazon DynamoDB is used to store name-value pair collections?

- A. String type
- B. List type
- C. Map type**
- D. Binary type

In Amazon DynamoDB, the Map type is designed specifically to store name-value pair collections. This data structure allows users to create an object that contains one or more key-value pairs. Each key within the Map must be a unique string, while the associated value can be of any DynamoDB data type, including strings, numbers, lists, or even other maps. This flexibility enables developers to store complex data structures easily, making it suitable for representing structured data in NoSQL databases. For example, when storing user profiles, a Map could be used to hold attributes like first name, last name, and email, each associated with their respective values. This facilitates a more organized way of handling data, especially when dealing with nested or hierarchical information. The other data types mentioned have different purposes. The String type is used for single text values, the List type is for ordered collections of items, which might not be key-value pairs, and the Binary type is for data that is not represented as text, such as images or files. Thus, the Map type is the correct and most appropriate choice for storing collections of name-value pairs in DynamoDB.

3. What is the key characteristic of a VPC?

- A. It automatically scales applications based on demand
- B. It provides a private, isolated section of the AWS Cloud**
- C. It enables real-time data processing
- D. It supports containerized applications

A Virtual Private Cloud (VPC) is primarily characterized by its ability to provide a private and isolated section of the AWS Cloud. This feature is essential because it allows users to define a virtual network that is logically separated from other networks in the AWS Cloud. Within a VPC, users can control various aspects of their networking environment, including the selection of IP address ranges, subnets, routing tables, and security settings. This isolation is crucial for scenarios requiring heightened security, allowing organizations to host sensitive applications and data while maintaining control over their infrastructure. The other options describe functionalities that may apply to various services within AWS but do not directly define the core characteristic of a VPC. For instance, automatic scaling, real-time data processing, and support for containerized applications are features of different AWS services or architectures but are not intrinsic to the VPC itself. Therefore, the defining trait of a VPC is its role as a secure and isolated network environment within the broader AWS infrastructure.

4. What features does the AWS Well-Architected Tool provide?

- A. API integration capabilities
- B. Code deployment automation
- C. Best practices reviews and improvement recommendations**
- D. Cost estimation for resources

The AWS Well-Architected Tool is designed to help architects build secure, high-performing, resilient, and efficient infrastructure for their applications. One of its primary features is providing best practices reviews and improvement recommendations. Using this tool, users can assess their architectures against AWS's best practices, which cover various pillars such as operational excellence, security, reliability, performance efficiency, and cost optimization. After an assessment, the tool offers specific recommendations that can help users improve their architecture to better align with AWS's well-architected framework. This feedback is invaluable for enhancing overall system design, compliance, and operational efficiency. In contrast, while API integration capabilities, code deployment automation, and cost estimation for resources are all important aspects of cloud operations, they highlight other functionalities not directly provided by the Well-Architected Tool. Instead, these features are part of different AWS services and tools that cater to specific operational needs in the cloud environment.

5. Which database service supports petabytes of data and provides columnar storage on high-performance disks?

- A. Amazon RDS**
- B. Amazon Aurora**
- C. Amazon Redshift**
- D. Amazon DynamoDB**

Amazon Redshift is designed specifically for large-scale data storage and analytics, supporting petabytes of data effectively. It utilizes a columnar storage architecture, which is optimal for analytical queries because it allows for efficient reading of only the required data columns rather than entire rows. This capability leads to faster query performance and reduces the amount of I/O required. Additionally, Amazon Redshift operates on high-performance disks known as SSDs, which further enhances the speed of data retrieval and processing. This combination of columnar storage and high-performance disk usage makes Redshift particularly well-suited for data warehousing and complex analytical workloads, enabling organizations to run large-scale analytics and retrieve insights from massive datasets efficiently. In contrast, services like Amazon RDS and Amazon Aurora are primarily intended for transactional databases and do not utilize the same columnar storage approach best suited for analytic tasks. Amazon DynamoDB, while allowing for scalability, is a NoSQL database primarily designed for key-value and document data models, rather than the structured data queries typically associated with petabyte-scale analytics. Therefore, Redshift stands out as the appropriate choice for handling petabytes of structured data with high efficiency.

6. Which AWS service allows real-time monitoring and management of AWS resources?

- A. Amazon Aurora**
- B. Amazon CloudWatch**
- C. Amazon VPC**
- D. Amazon SQS**

Amazon CloudWatch is the service designed specifically for real-time monitoring and management of AWS resources. It provides key functionalities to collect, analyze, and visualize metrics, which can include resource utilization, application performance, and operational health. By offering dashboards that display real-time information about resource performance, along with alerts based on predefined thresholds, CloudWatch enables administrators to respond to operational changes promptly, making it an essential tool for maintaining optimal performance and resource management in a cloud environment. Other AWS services, such as Amazon Aurora, focus primarily on managed database solutions, while Amazon VPC is related to networking and the definition of virtual networks, and Amazon SQS is a message queuing service. None of these services are inherently designed for monitoring and managing AWS resources in real time like CloudWatch does.

7. Which AWS service simplifies the database migration process to AWS?

- A. AWS Database Migration Service (DMS)**
- B. AWS Snowball**
- C. AWS Glue**
- D. AWS Data Pipeline**

AWS Database Migration Service (DMS) is specifically designed to facilitate the migration of databases to AWS. It supports various source database engines and allows users to migrate their data seamlessly with minimal downtime. DMS can handle the complexities involved in a database migration, such as schema conversion and data transformation, making it straightforward for users to move their on-premises databases or other database environments into AWS services. The service allows for continuous data replication, which means that source databases can remain operational during the migration process. This capability significantly reduces the risk of data loss and system downtime. Other services mentioned, such as AWS Snowball, AWS Glue, and AWS Data Pipeline, serve different purposes. For instance, Snowball is designed for transferring large amounts of data into and out of AWS quickly and securely, while Glue is primarily an ETL (Extract, Transform, Load) service for data preparation and integration. Data Pipeline offers a way to process and move data between AWS services but is not specifically aimed at migrating databases. Thus, while these services have their functions, they do not simplify the database migration process in the same targeted manner as AWS DMS.

8. What is a key benefit of creating reusable infrastructure templates?

- A. Increased manual intervention**
- B. Consistency in deployments**
- C. Higher cost of maintenance**
- D. Reduced scalability**

Creating reusable infrastructure templates offers significant advantages, particularly in achieving consistency in deployments. By using templates, teams can ensure that every deployment follows the same configuration and settings, minimizing the risk of human error that can occur when configurations are manually input. This uniformity makes it easier to maintain, update, and troubleshoot environments, as all deployments adhere to the same standards. Moreover, the use of reusable templates can streamline the deployment process across different environments—development, testing, and production. This standardization not only enhances reliability but also accelerates the deployment cycle, allowing teams to deploy applications faster and more efficiently. In contrast, options that suggest increased manual intervention, higher maintenance costs, or reduced scalability are contrary to the benefits associated with infrastructure as code and the use of templates. Templates are designed specifically to automate and simplify processes, thereby reducing the need for ongoing manual oversight and promoting overall efficiency.

9. What type of storage is Amazon EBS?

- A. File storage
- B. Block storage**
- C. Object storage
- D. Hybrid storage

Amazon Elastic Block Store (EBS) is classified as block storage, which is designed for scenarios requiring fine-grained control over the data in storage. Block storage divides data into fixed-sized chunks, or blocks, and each block can be treated independently, allowing for efficient read and write operations. This structure makes EBS ideal for use with Amazon EC2 instances, where it functions like a traditional hard drive, providing persistent storage that can be attached to instances as needed. Block storage is particularly suited for applications that require low-latency and high-performance data access, such as databases or enterprise applications, which further emphasizes the suitability of Amazon EBS for these use cases. With EBS, users can create volumes that can be attached or detached from EC2 instances, enabling them to scale their storage requirements dynamically and manage their data effectively.

10. Which AWS service provides a scalable object storage solution for data archiving?

- A. Amazon Glacier**
- B. Amazon S3
- C. Amazon EBS
- D. Amazon RDS

Amazon Glacier is specifically designed as a scalable object storage service that focuses on data archiving and long-term storage. It is optimized for infrequently accessed data and offers a low-cost solution for storing massive amounts of data that do not need to be accessed quickly. The primary use case for Glacier is to store data in a secure, durable manner while minimizing costs for storage that is rarely retrieved, making it ideal for compliance, backup, or disaster recovery scenarios. Unlike other options, such as Amazon S3, which is a more general-purpose storage solution that supports frequent access and various storage classes including archival solutions, Glacier specifically targets long-term data archiving with features like retrieval policies that control how quickly archived data can be accessed. Amazon EBS provides block storage primarily for use with Amazon EC2 instances, which is more suited for data that requires low-latency access and high performance, and is not tailored for archiving purposes. Amazon RDS, on the other hand, is a managed database service meant for relational databases, completely unrelated to object storage or archiving functions. Therefore, when considering the specific need for scalable object storage geared towards data archiving, Amazon Glacier stands out as the appropriate solution.

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://awsacademycloudops.examzify.com>

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

SAMPLE