

AWS Cloud Architecting Practice Exam (Sample)

Study Guide



Everything you need from our exam experts!

This is a sample study guide. To access the full version with hundreds of questions,

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 from reliable sources accurate, complete, and timely information about this product.

SAMPLE

Table of Contents

Copyright	1
Table of Contents	2
Introduction	3
How to Use This Guide	4
Questions	6
Answers	9
Explanations	11
Next Steps	17

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. Don't worry about getting everything right, 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, and take breaks to retain information better.

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.

7. Use Other Tools

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. Which Amazon Elastic Compute Cloud (EC2) pricing model should you choose for applications with short-term, spiky, or unpredictable workloads that cannot be interrupted?**
 - A. On-Demand Instances**
 - B. Spot Instances**
 - C. Reserved Instances**
 - D. Dedicated Instances**
- 2. True or False: Amazon EC2 instances are immediately available for use once the OS has booted and status checks are passed.**
 - A. True**
 - B. False**
 - C. Not Determined**
 - D. Depends on Instance Type**
- 3. What role does AWS Identity and Access Management (IAM) play in AWS Lambda?**
 - A. Defines billing options for Lambda usage**
 - B. Manages permissions for Lambda to access necessary AWS resources**
 - C. Controls the geographical location of Lambda functions**
 - D. Determines the version of the Lambda runtime**
- 4. Which statements describe a service-oriented architecture (SOA)?**
 - A. It relies heavily on a monolithic design structure**
 - B. Implementation details become irrelevant**
 - C. It requires a single programming language throughout**
 - D. Services are tightly coupled with other components**
- 5. True or False: It is a best practice to use custom route tables for each subnet for granular routing.**
 - A. True**
 - B. False**
 - C. Only for specific subnets**
 - D. Depends on the architecture**

6. A customer needs per-second billing as an option. Which types of purchase options do they have regarding instance type?

- A. On-Demand Instances and Spot Instances**
- B. Reserved Instances only**
- C. Dedicated Hosts only**
- D. Spot Instances only**

7. Which of the following responsibilities of the customer is part of the shared responsibility model?

- A. Evaluate soft limits and request increases**
- B. Manage physical security of data centers**
- C. Perform patch management**
- D. Control network segmentation**

8. What is the main advantage of using encryption for data at rest?

- A. Faster data retrieval**
- B. Protection against unauthorized access**
- C. Reduced storage costs**
- D. Improved application performance**

9. Which design principle for achieving performance efficiency uses an AWS service to remove the need to run and maintain servers?

- A. Use server-based architectures**
- B. Use serverless architectures**
- C. Use on-premises solutions**
- D. Use hybrid cloud environments**

10. Which statement accurately describes a private subnet?

- A. Accessible from the public internet**
- B. Has a direct route to an internet gateway**
- C. Not directly accessible from the public internet**
- D. Uses a public IP for all instances**

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. Which Amazon Elastic Compute Cloud (EC2) pricing model should you choose for applications with short-term, spiky, or unpredictable workloads that cannot be interrupted?

- A. On-Demand Instances**
- B. Spot Instances**
- C. Reserved Instances**
- D. Dedicated Instances**

Choosing On-Demand Instances for applications with short-term, spiky, or unpredictable workloads that require high availability aligns perfectly with their pricing model and use case characteristics. On-Demand Instances allow you to pay for computing capacity by the hour or second, depending on the instance type, without having to commit to a long-term contract. This flexibility ensures that your applications can scale up during peak times and reduce usage when demand decreases, making it ideal for workloads that have unpredictable usage patterns. Moreover, On-Demand Instances guarantee that your resources are always available when needed, addressing the crucial requirement that these workloads should not be interrupted. In scenarios where uptime and immediate availability are critical, choosing On-Demand pricing offers peace of mind and reliability. In contrast, other pricing models, such as Spot Instances, are designed for applications that can tolerate interruptions and are ideal for cost savings when the demand for instances is low. Reserved Instances require a commitment for a set term, making them less suitable for unpredictable workloads. Dedicated Instances, while providing isolated hardware for compliance or regulatory requirements, do not inherently accommodate fluctuations in demand like On-Demand Instances do. Thus, On-Demand Instances serve as the best choice for the specified workload characteristics.

2. True or False: Amazon EC2 instances are immediately available for use once the OS has booted and status checks are passed.

- A. True**
- B. False**
- C. Not Determined**
- D. Depends on Instance Type**

The statement is true: Amazon EC2 instances are indeed available for use once the operating system has successfully booted and all associated status checks have passed. When an EC2 instance is launched, it goes through a series of initialization processes, which culminate in the booting of the operating system. Once this process is complete and the instance is running without errors, the instance is considered ready for use. At this point, users can access the instance via SSH for Linux or RDP for Windows and start deploying applications or services. This is fundamental to understanding AWS's on-demand cloud infrastructure, where the speed of instance provisioning and availability is a significant advantage. Instances that pass their health checks are thus reliable for immediate workloads. Other possible choices do not apply here since the availability of an instance does not depend on any unspecified conditions, nor is it ambiguous or reliant on different instance types in this context. The operational readiness of EC2 instances follows a clear and consistent pattern across all types once the OS has booted and checks are satisfactory.

3. What role does AWS Identity and Access Management (IAM) play in AWS Lambda?

- A. Defines billing options for Lambda usage
- B. Manages permissions for Lambda to access necessary AWS resources**
- C. Controls the geographical location of Lambda functions
- D. Determines the version of the Lambda runtime

AWS Identity and Access Management (IAM) plays a crucial role in managing permissions for AWS Lambda functions to access other AWS resources. This capability allows you to define what actions a Lambda function can perform and which resources it can interact with. For example, when a Lambda function needs to read from an S3 bucket, IAM policies can be set to grant that specific permission. By using IAM roles, you can securely control access and maintain the principle of least privilege, ensuring that functions have only the permissions necessary to perform their tasks. The other options are relevant aspects of AWS but do not accurately describe the function of IAM in relation to Lambda. Billing options pertain to cost management rather than access control, geographical locations are managed through the deployment settings of Lambda and the use of AWS Regions, and versions of the Lambda runtime are determined by the AWS environment setup rather than IAM.

4. Which statements describe a service-oriented architecture (SOA)?

- A. It relies heavily on a monolithic design structure
- B. Implementation details become irrelevant**
- C. It requires a single programming language throughout
- D. Services are tightly coupled with other components

A service-oriented architecture (SOA) primarily revolves around the design of software systems as a collection of services that communicate with one another. One of the defining characteristics of SOA is that the implementation details of one service are abstracted away from the users of that service. This means that consumers of the service do not need to know how the service is implemented; they only need to know how to interact with it through well-defined interfaces or protocols. As a result, this degree of abstraction allows for greater flexibility and the ability to change the underlying implementations without impacting the consumers, thereby facilitating maintenance, scaling, and integration of different services. The other options suggest aspects that are contrary to the principles of SOA. A monolithic design structure undermines the modularity that SOA promotes, while the requirement for a single programming language contradicts the flexibility SOA offers in terms of technology choices. Tightly coupled components are also at odds with the loose coupling that SOA aims to achieve, which enhances service scalability and reusability.

5. True or False: It is a best practice to use custom route tables for each subnet for granular routing.

- A. True**
- B. False**
- C. Only for specific subnets**
- D. Depends on the architecture**

Using custom route tables for each subnet is considered a best practice in AWS for achieving granular routing, making the statement true. By employing custom route tables, you can control the routing of traffic to and from resources more effectively within your Virtual Private Cloud (VPC). This flexibility allows you to specify routes that are tailored to the specific needs of each subnet, which can be critical for optimizing performance, security, and connectivity. For example, a public subnet may require a route to an internet gateway to allow external traffic, whereas a private subnet might direct traffic through a NAT gateway for outbound internet access without exposing its resources directly. Additionally, custom route tables can facilitate integration with different services, such as peering connections, VPNs, or Direct Connect, depending on the architecture's requirements. The option stating that this practice is only for specific subnets does not capture the full scope of the best practice, which applies broadly across different subnet configurations. The notion that it depends on the architecture also doesn't wholly embrace the acknowledged benefits of segregating route management through custom route tables, as the best practice is relevant in various architectural scenarios for improved network control. Overall, utilizing custom route tables is vital for creating a well-optimized and secure network environment in AWS.

6. A customer needs per-second billing as an option. Which types of purchase options do they have regarding instance type?

- A. On-Demand Instances and Spot Instances**
- B. Reserved Instances only**
- C. Dedicated Hosts only**
- D. Spot Instances only**

Per-second billing is an important feature for customers who want precise control over their costs based on actual usage, especially for workloads that are variable or unpredictable. On-Demand Instances allow users to pay for compute capacity by the second, with no long-term commitments. This provides flexibility, as users can start and stop instances as needed, making it ideal for applications with short-term or unpredictable workloads. Spot Instances also offer per-second billing, as they allow users to bid on spare AWS capacity, providing a cost-effective solution for flexible applications that can tolerate interruptions. This combination of On-Demand and Spot Instances ensures customers can benefit from per-second billing based on their specific usage scenarios. The other options do not provide per-second billing. Reserved Instances require a commitment for a one- or three-year term and are primarily billed on an hourly basis, although they can provide lower rates for predictable workloads. Dedicated Hosts are dedicated physical servers to run your instances and follow the same hourly billing structure. Thus, when considering the need for per-second billing, On-Demand Instances and Spot Instances are the appropriate choices.

7. Which of the following responsibilities of the customer is part of the shared responsibility model?

- A. Evaluate soft limits and request increases**
- B. Manage physical security of data centers**
- C. Perform patch management**
- D. Control network segmentation**

The shared responsibility model in AWS clearly delineates the division of responsibilities between AWS and its customers. In this context, customers are responsible for managing specific elements of their applications and services that run on the AWS platform. Evaluating soft limits and requesting increases is a task that falls under the customer's purview. This means that customers must monitor their usage of AWS resources and understand their limits, including service quotas, and take proactive steps to request increases if they anticipate needing more resources. This responsibility is essential for effectively managing performance and ensuring applications run smoothly. Physical security of data centers, on the other hand, is managed entirely by AWS. AWS is responsible for securing the infrastructure that houses the data centers, including hardware and physical facilities. Performing patch management is also a responsibility that falls on AWS when it pertains to the underlying infrastructure. Customers, however, must manage patching for their applications and the operating systems they deploy on AWS services. Lastly, control of network segmentation is a task that customers must implement within their own configuration of AWS services, while AWS provides the tools and infrastructure to enable this. Therefore, managing the limits and understanding the request process directly correlates to the customer's responsibilities within the shared responsibility model, highlighting essential operational management for their deployment on

8. What is the main advantage of using encryption for data at rest?

- A. Faster data retrieval**
- B. Protection against unauthorized access**
- C. Reduced storage costs**
- D. Improved application performance**

Using encryption for data at rest primarily provides protection against unauthorized access. This means that even if an unauthorized user gains physical access to the storage medium where the data is stored, they cannot read or use the data without the proper decryption key. This adds a critical layer of security, ensuring that sensitive information, such as personal data or proprietary business data, remains confidential and secure against threats such as data breaches, theft, or insider threats. The implementation of encryption at rest is especially important in cloud environments where data multitude could be stored across various locations and platforms. Compliance with legal and regulatory standards regarding data protection often mandates the use of strong encryption to safeguard this data. Therefore, encryption serves as a robust defense mechanism against potential vulnerabilities associated with stored data, providing peace of mind to businesses and individuals alike regarding the confidentiality and integrity of their information.

9. Which design principle for achieving performance efficiency uses an AWS service to remove the need to run and maintain servers?

- A. Use server-based architectures**
- B. Use serverless architectures**
- C. Use on-premises solutions**
- D. Use hybrid cloud environments**

The design principle for achieving performance efficiency that involves using an AWS service to eliminate the need to run and maintain servers is centered around serverless architectures. Serverless architectures allow developers to focus on writing code and building applications without the overhead of managing the underlying infrastructure. Services like AWS Lambda, API Gateway, and DynamoDB facilitate this by automatically handling scaling and server management, allowing developers to execute code in response to events without provisioning or managing servers. This leads to cost savings, automatic scaling based on demand, and a simplified operational experience, aligning perfectly with the goals of performance efficiency and agility in application development. In contrast, server-based architectures typically require ongoing management of server instances, which can introduce more complexity and operational overhead. On-premises solutions involve maintaining physical servers, which can be less efficient in terms of scaling and performance management compared to a cloud-based serverless approach. Hybrid cloud environments mix on-premises resources with cloud services, but they still involve server management and do not leverage the full benefits of removing server maintenance requirements entirely. By utilizing a serverless architecture, organizations can enhance performance efficiency while minimizing operational burdens.

10. Which statement accurately describes a private subnet?

- A. Accessible from the public internet**
- B. Has a direct route to an internet gateway**
- C. Not directly accessible from the public internet**
- D. Uses a public IP for all instances**

A private subnet is defined by its inability to be accessed directly from the public internet. This characteristic is essential for scenarios where resources need to be safeguarded from external access while still allowing them to communicate with other instances within the same virtual private cloud (VPC) or with the public internet through specific configurations like a virtual private network (VPN) or a bastion host. In a private subnet, instances do not receive public IP addresses, which further enhances their security by preventing unsolicited traffic from the internet. This design is typically used for resources such as databases and application servers that do not need to be exposed to the public and should only be accessed from within the private network or via controlled means. The other statements describe characteristics of public subnets or misrepresent the configuration of a private subnet, making them less applicable in this context.

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://awscloudarchitecting.examzify.com>

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

SAMPLE