


AWS CLOUD ARCHITECTING PRACTICE EXAM (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	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. Which type of connections use a virtual private gateway in an Amazon VPC?**
 - A. Direct Connect and VPN
 - B. PrivateLink and VPC Peering
 - C. Internet Gateway and NAT Gateway
 - D. Route 53 and CloudFront
- 2. True or False? When choosing a Region for Amazon S3, geographical proximity to your users is important.**
 - A. True
 - B. False
 - C. Only for latency-sensitive applications
 - D. Only if using public S3 buckets
- 3. Which Amazon service is best for delivering content with low latency?**
 - A. Amazon S3
 - B. Amazon CloudFront
 - C. Amazon RDS
 - D. Amazon Lambda
- 4. 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
- 5. Which service is best suited for real-time data processing over large distributed data streams?**
 - A. Amazon EMR
 - B. Amazon Kinesis
 - C. Amazon S3
 - D. Amazon RDS

- 6. What defines how to launch Amazon EC2 instances in Amazon Auto Scaling?**
- A. Launch Configuration
 - B. Instance Profile
 - C. Security Group
 - D. Auto Scaling Policy
- 7. What is an essential benefit of using Amazon CloudFront?**
- A. Data backup
 - B. Content delivery optimization
 - C. Database scaling
 - D. Server management
- 8. Which service helps you see an inventory of your AWS resources and configurations?**
- A. AWS CloudTrail
 - B. Amazon CloudWatch
 - C. AWS Config
 - D. Amazon S3
- 9. What action should be taken if an EC2 instance is unreachable?**
- A. Stop and start the instance
 - B. Modify the security group rules
 - C. Check the instance status checks
 - D. All of the above
- 10. What is the maximum size of a single message that can be sent through Amazon SNS?**
- A. 128 KB
 - B. 256 KB
 - C. 512 KB
 - D. 1 MB

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. Which type of connections use a virtual private gateway in an Amazon VPC?

- A. Direct Connect and VPN**
- B. PrivateLink and VPC Peering
- C. Internet Gateway and NAT Gateway
- D. Route 53 and CloudFront

A virtual private gateway in an Amazon VPC is designed to facilitate secure connections to an on-premises network or to other networks from within a VPC. The two primary connection types that utilize a virtual private gateway are AWS Direct Connect and VPN (Virtual Private Network) connections. AWS Direct Connect provides a dedicated network connection from your premises to AWS. This connection allows you to bypass the internet, providing more reliability, faster speeds, and lower data transfer costs. When you set up a Direct Connect, it connects to the virtual private gateway, allowing you to access your VPC resources securely. Similarly, a VPN connection uses the virtual private gateway to establish an encrypted tunnel between your on-premises network and your VPC over the public Internet. This ensures that data is securely transmitted between the two locations, utilizing the secure gateway as a point of termination for the VPN connection. Both Direct Connect and VPN rely on the virtual private gateway to manage and secure the data flow in and out of the VPC, making this the correct answer. Other connection types listed do not utilize a virtual private gateway, as they serve different purposes or connect VPCs internally rather than to external networks.

2. True or False? When choosing a Region for Amazon S3, geographical proximity to your users is important.

- A. True**
- B. False
- C. Only for latency-sensitive applications
- D. Only if using public S3 buckets

Choosing a region for Amazon S3 based on geographical proximity to your users is important because it can significantly impact the performance of your application. When users are closer to the region where your data is stored, they experience lower latency when accessing that data. This is especially crucial for applications that require faster data retrieval or have real-time requirements, as geographical distance can introduce delays in data transfer. Additionally, selecting an appropriate region helps ensure compliance with local data residency laws and regulations, which may dictate where data can be stored or processed. Proximity also plays a critical role in reducing the cost of data transfer, as cross-region data requests can incur additional charges compared to accessing data within the same region. While some might argue that this factor is only relevant for latency-sensitive applications or when using specific S3 configurations, the principle of geographical proximity remains a fundamental consideration for any application relying on data stored in S3. Thus, the statement is indeed true.

3. Which Amazon service is best for delivering content with low latency?

- A. Amazon S3
- B. Amazon CloudFront**
- C. Amazon RDS
- D. Amazon Lambda

Amazon CloudFront is the most suitable service for delivering content with low latency due to its design as a content delivery network (CDN). CloudFront caches copies of your content at various edge locations around the world, which are strategically placed closer to end-users. This geographic distribution reduces the distance data must travel, significantly improving response times and minimizing latency. When a user requests content, CloudFront can serve it from the nearest edge location. This retrieval process is much faster than retrieving content from a central server or a data center that may be located far away from the user, thereby achieving low-latency delivery. In addition, CloudFront optimizes file delivery through various mechanisms such as dynamic content delivery, forming multiple connections, and effectively managing high volumes of requests. In contrast, while Amazon S3 is a highly durable storage solution, it does not inherently provide the low-latency delivery advantages that CloudFront does when serving content. Amazon RDS is a relational database service that manages and scales databases, but it is primarily focused on data storage and retrieval rather than content delivery. Amazon Lambda, a serverless compute service, processes and executes code in response to events but does not specifically address the needs for low-latency content delivery.

4. 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

5. Which service is best suited for real-time data processing over large distributed data streams?

- A. Amazon EMR
- B. Amazon Kinesis**
- C. Amazon S3
- D. Amazon RDS

Amazon Kinesis is the most suitable service for real-time data processing over large distributed data streams. It is specifically designed to handle streaming data, allowing you to collect, process, and analyze real-time data at scale. Kinesis supports various use cases, including log and event data collection, real-time analytics, and machine learning applications, making it ideal for scenarios that require rapid ingestion and processing of data. The service enables applications to process thousands of records per second from multiple sources, and it seamlessly integrates with analytics tools such as Amazon Redshift and Amazon Elasticsearch Service for further processing. Key features like data shards and the ability to scale throughput also contribute to its capability to manage large volumes of real-time data effectively. This makes Kinesis a robust choice for applications that require immediate insights and reactions to incoming data, differentiating it from other AWS services that may be better suited for batch processing or storage rather than real-time stream processing.

6. What defines how to launch Amazon EC2 instances in Amazon Auto Scaling?

- A. Launch Configuration**
- B. Instance Profile
- C. Security Group
- D. Auto Scaling Policy

The determination of how to launch Amazon EC2 instances within Auto Scaling groups is primarily handled by a Launch Configuration. A Launch Configuration is a template that specifies the instance type, Amazon Machine Image (AMI), key pair, security groups, and other parameters required for launching EC2 instances within the Auto Scaling group. This template allows the Auto Scaling service to create and configure instances consistently based on the defined settings when scaling out to meet demand or scaling in to reduce capacity. In contrast, the other options serve different purposes in the AWS ecosystem. An Instance Profile is used to associate permissions with the instances, allowing them to access AWS resources securely but does not define the launch configurations for new instances. A Security Group acts as a virtual firewall to control inbound and outbound traffic for your instances, but it is not responsible for determining the instance launch parameters. An Auto Scaling Policy defines the conditions under which scaling actions occur, such as when to add or remove instances, but it does not dictate how those instances should be launched.

7. What is an essential benefit of using Amazon CloudFront?

- A. Data backup
- B. Content delivery optimization**
- C. Database scaling
- D. Server management

Using Amazon CloudFront provides significant benefits in content delivery optimization. CloudFront is a Content Delivery Network (CDN) that accelerates the distribution of static and dynamic web content, such as HTML, CSS, JavaScript, images, and videos, to users based on their geographic location. By caching content at edge locations close to the end-users, CloudFront reduces latency and improves loading speed, which enhances overall user experience and satisfaction. Additionally, CloudFront utilizes various optimizations, such as using HTTP/2, allowing for multiplexing and prioritization of requests, thereby further enhancing the delivery speed of content. It also integrates seamlessly with other AWS services, ensuring that content can be delivered efficiently while benefiting from AWS's scale and infrastructure. The other options listed do not align with the primary purpose of CloudFront. Data backup pertains to storage services like Amazon S3 or AWS Backup. Database scaling is relevant to services like Amazon RDS or DynamoDB. Server management includes handling compute resources through services like Amazon EC2, rather than content distribution. Thus, content delivery optimization is clearly the defining benefit of using Amazon CloudFront.

8. Which service helps you see an inventory of your AWS resources and configurations?

- A. AWS CloudTrail
- B. Amazon CloudWatch
- C. AWS Config**
- D. Amazon S3

The correct answer is AWS Config. This service is specifically designed to provide visibility into the inventory of AWS resources and their configurations. AWS Config continuously monitors and records your AWS resource configurations and allows you to automate the evaluation of recorded configurations against desired configurations. By using AWS Config, you can assess compliance with your organization's policies and standards, track changes over time, and understand how resources relate to one another. This capability is crucial for effective governance, risk management, and auditing within the AWS environment. AWS CloudTrail, while useful for logging API calls for resource management and monitoring user activity across your AWS account, does not provide a detailed inventory of existing resources or their configurations. Instead, it focuses on tracking API usage and changes made at the operational level. Amazon CloudWatch serves a different purpose, primarily concerned with monitoring AWS cloud resources and applications in real time. It provides metrics and log data for AWS resources but does not offer a comprehensive view of resource configuration inventories. Amazon S3 is a scalable storage solution that allows you to store and retrieve data. While you can store configuration data or logs in S3, it does not inherently manage or provide visibility into your AWS resources' inventory or configurations. Thus, AWS Config is the only service among the options provided that

9. What action should be taken if an EC2 instance is unreachable?

- A. Stop and start the instance
- B. Modify the security group rules
- C. Check the instance status checks
- D. All of the above**

When an EC2 instance is unreachable, it is important to consider a multifaceted approach to diagnosing and resolving the issue. Stopping and starting the instance can often resolve transient issues. This action effectively reboots the instance, which can help clear up any underlying problems that may be preventing the instance from responding. Modifying the security group rules is also crucial because it addresses potential networking issues. If the rules inadvertently block access, the instance may seem unreachable. Adjusting these rules can restore the necessary traffic flow to and from the instance. Checking the instance status checks is an essential step in troubleshooting. AWS performs automated system checks on the instance, and reviewing the results can indicate whether the instance has passed all checks or if there are underlying hardware or software issues. Thus, taking all of these actions combined provides a comprehensive strategy for diagnosing and potentially resolving the issue of an unreachable EC2 instance. Each step deals with different aspects of what could be causing the accessibility problem, making it a logical collective choice to take all the listed actions.

10. What is the maximum size of a single message that can be sent through Amazon SNS?

- A. 128 KB
- B. 256 KB**
- C. 512 KB
- D. 1 MB

Amazon Simple Notification Service (SNS) has a maximum message size limit of 256 KB for a single message. This means that when you publish a message to an SNS topic, the content of that message cannot exceed this size. It's important to note this limit when designing applications that utilize SNS to ensure that your messages are structured within this constraint. If you need to send larger messages, you can consider breaking the message into smaller parts or using alternative methods such as Amazon S3 for sending payloads that exceed this limit while only including a reference to the location of the data in the SNS message. Understanding these limits is crucial for effective message handling and ensuring that your system operates without errors related to message size.

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