AWS DEVOPS ENGINEER PROFESSIONAL PRACTICE TEST (SAMPLE)

STUDY GUIDE



SAMPLE STUDY GUIDE WITH LIMITED QUESTIONS

VISIT US ONLINE FOR THE FULL VERSION STUDY GUIDE

https://aws-devops.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 the function of AWS IAM?

A. To create virtual machines

B. To manage user access and permissions

C. To provide database services

D. To monitor application performance

2. How can you pass a sensitive value to a CodeBuild project securely?

A. Configure the value in environment variables

B. Store the value in a text file in the repository

C. Configure the value in SSM Parameter Store or Secrets Manager

D. Hardcode the value in the buildspec file

3. What feature should be used for AWS instances that require close physical proximity for optimal performance?

A. Availability Zones

B. Placement Groups

C. VPC Peering

D. Auto Scaling Groups

4. What is Amazon EC2 primarily used for?

A. Running virtual servers in the cloud

B. Data analytics and reporting

C. Securing network traffic

D. Machine learning modeling

5. When adding on-premise instances to a CodeDeploy deployment group, how many IAM users should you create for 5 instances?

A. 1

B. 3

C. 5

D. 10

SAMPLE

6. How long will AWS OpsWorks wait before considering an instance unhealthy?

A. 1 minute

B. 3 minutes

C. 5 minutes

D. 10 minutes

7. What does AWS CodeCommit provide?

A. Secure Git-based repository for source control

B. Continuous integration services

C. API management tools

D. Database management system

8. For load response in OpsWorks, what should you create?

A. Static instances

B. Load-based instances

C. Resource-specific instances

D. Time-based instances

9. When are commands in the "commands" section of an .ebextensions file executed in Elastic Beanstalk?

A. After the application is deployed

B. Before the application version file is extracted

C. After the web server is configured

D. Before the application and web server are set up

10. What AWS service helps in reducing costs for data transfer?

A. Amazon S3

B. Amazon CloudFront

C. Amazon RDS

D. AWS Lambda

SAMPLE

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. What is the function of AWS IAM?

A. To create virtual machines

B. To manage user access and permissions

C. To provide database services

D. To monitor application performance

AWS Identity and Access Management (IAM) is a critical service that focuses on managing user identities and permissions within the AWS cloud environment. Its primary function is to control who has access to AWS resources and what actions they can perform. With IAM, administrators can create and manage AWS users and groups, define permissions for those users and groups, and enable secure delegation of access. This capability allows organizations to ensure that only authorized personnel can access specific resources, aligning with security best practices. IAM enables fine-grained access control, meaning permissions can be assigned at a very specific level—down to individual resources like S3 buckets or EC2 instances. Moreover, IAM supports multi-factor authentication (MFA) and integrates with AWS CloudTrail for logging and tracking access, amplifying security measures. In contrast, the other options describe functions outside the scope of IAM. For example, creating virtual machines is handled by services like Amazon EC2, database services are provided by AWS RDS or DynamoDB, and application performance monitoring falls under tools such as Amazon CloudWatch. Therefore, the focus of AWS IAM squarely lies in user access and permission management within AWS environments.

2. How can you pass a sensitive value to a CodeBuild project securely?

A. Configure the value in environment variables

B. Store the value in a text file in the repository

C. Configure the value in SSM Parameter Store or Secrets Manager

D. Hardcode the value in the buildspec file

Using SSM Parameter Store or Secrets Manager to pass a sensitive value to a CodeBuild project securely is the most appropriate and secure method. Both of these services are specifically designed to store sensitive information such as passwords, API keys, and configuration settings in a secure manner. SSM Parameter Store allows you to securely store and manage parameters, while Secrets Manager provides more advance features for managing secrets, including automatic rotation and cryptographic storage. By integrating these services with your CodeBuild project, you can reference these secure values in your build environment without exposing them in your source code or build configurations. This significantly reduces the risk of accidental leakage or exposure of sensitive information. The alternatives provided can expose sensitive information: Configuring the value in environment variables may seem like a straightforward approach, but it can lead to potential exposure in build logs, making it less secure for sensitive data. Storing the value in a text file in the repository poses a security risk, as it can be accessed by anyone with access to the repository. Hardcoding the value in the buildspec file is not advisable for sensitive values, as it makes them visible in the source code, potentially compromising them.

SAMPLE

3. What feature should be used for AWS instances that require close physical proximity for optimal performance?

A. Availability Zones

B. Placement Groups

C. VPC Peering

D. Auto Scaling Groups

For AWS instances that require close physical proximity to optimize performance, the appropriate feature is Placement Groups. Placement Groups are designed to control the placement of instances within a specific physical location in the data center, which can help benefit workloads that require high throughput or low latency, such as high-performance computing (HPC) or big data applications. By using Placement Groups, instances are launched in close proximity to one another within the same availability zone, thereby minimizing network latency and maximizing network bandwidth. This is particularly crucial for applications that need rapid communication between instances, as it enhances the overall performance and efficiency of the application. The other options, while important for different aspects of cloud architecture, do not offer the same level of control over physical proximity as Placement Groups. For example, Availability Zones provide redundancy and fault tolerance across different data center locations, but they do not ensure that instances are closely located within the same zone. VPC Peering facilitates networking between different virtual private clouds but does not impact the physical locations of instances. Auto Scaling Groups are primarily used for adjusting the number of instances based on demand but do not directly relate to the physical proximity of instances for performance optimization.

4. What is Amazon EC2 primarily used for?

A. Running virtual servers in the cloud

B. Data analytics and reporting

C. Securing network traffic

D. Machine learning modeling

Amazon EC2 (Elastic Compute Cloud) is primarily used for running virtual servers in the cloud. It provides scalable compute capacity, allowing users to deploy and manage virtual machines using a variety of operating systems. This service offers the flexibility to launch and terminate instances as needed, enabling users to adapt to varying workloads and optimize costs. With EC2, users can choose from a range of instance types that are tailored to different application needs, such as memory optimization, storage capacity, or processing power. This versatility makes it possible to run everything from simple websites to complex applications while benefiting from the underlying infrastructure of AWS. The other options, although relevant to cloud computing, do not specifically describe the primary function of Amazon EC2. Data analytics and reporting, securing network traffic, and machine learning modeling may involve other AWS services such as Amazon RDS, AWS Security services, or Amazon SageMaker, but they are not the core purpose of EC2.

SAMPLE

5. When adding on-premise instances to a CodeDeploy deployment group, how many IAM users should you create for 5 instances?

A. 1

B. 3

C. 5

D. 10

When integrating on-premise instances with CodeDeploy, you typically need to create an IAM user for each instance that requires deployment. This is because each on-premise instance must have its own set of credentials (IAM user) for CodeDeploy to authenticate and manage deployments individually. In this scenario, if you have five instances, it is necessary to create five IAM users. Each instance will use a distinct IAM user to ensure the deployment process is secure and manageable. This approach facilitates better access control and monitoring, as each instance's activities can be tracked independently. The other options would not align with the best practices for managing deployments in this context. For example, having just one IAM user for multiple instances would pose security risks, as it would not provide the granularity needed for access management. Similarly, creating three or ten IAM users would either be insufficient or excessive based on the number of instances, leading to unnecessary complexity or inadequate access control.

6. How long will AWS OpsWorks wait before considering an instance unhealthy?

A. 1 minute

B. 3 minutes

C. 5 minutes

D. 10 minutes

AWS OpsWorks designates an instance as unhealthy if it fails to respond within a specific timeframe. The correct timeline for this evaluation is set at 5 minutes. This duration is valuable as it gives the application a sufficient opportunity to recover from transient issues that may occur during operations, such as temporary loss of connectivity or resource constraints that could affect performance. The 5-minute check aligns with AWS's focus on ensuring that instances are resilient and can potentially recover from temporary glitches without prematurely removing them from service. The other options suggest shorter wait times that might not adequately account for brief issues that could resolve themselves, leading to unnecessary restarts or failures. In the context of maintaining application robustness and minimizing disruption, the 5-minute threshold strikes an appropriate balance.

SAMPLE

7. What does AWS CodeCommit provide?

A. Secure Git-based repository for source control

B. Continuous integration services

C. API management tools

D. Database management system

AWS CodeCommit is a fully managed source control service that offers a secure Git-based repository for storing and managing source code. By using CodeCommit, teams can collaborate on code, track changes, and implement version control in a secure environment. The service provides several advantages that make it suitable for managing source code, including encryption at rest and in transit, which ensures that code is protected from unauthorized access. It also facilitates integration with other AWS services, enabling seamless workflows in a DevOps environment. While continuous integration services, API management tools, and database management systems are essential components in a DevOps lifecycle or broader software development practice, they fall outside the specific functionality provided by AWS CodeCommit. CodeCommit strictly focuses on source control, positioning it as a foundational element for developers working with Git repositories in the AWS ecosystem.

8. For load response in OpsWorks, what should you create?

A. Static instances

B. Load-based instances

C. Resource-specific instances

D. Time-based instances

Creating load-based instances in OpsWorks is essential for effective load response management. Load-based instances automatically adjust their number of instances based on the load that the application is experiencing. This means that as user traffic increases, the system can spin up additional instances to handle the increased demand, and conversely, it can scale down when the demand drops. This dynamic approach allows for efficient resource utilization, ensuring that you are only using the instances necessary to meet current needs. It also enhances application performance and availability by preventing bottlenecks during traffic spikes, providing a better experience for users. In contrast, other choices do not accurately reflect this need for scalability based on application load. Static instances would not provide any flexibility in response to changing demands. Resource-specific instances would not inherently scale according to load, and time-based instances focus on predetermined schedules rather than real-time load conditions. Thus, load-based instances are the optimal solution for managing load in OpsWorks effectively.

SAMPLE

9. When are commands in the "commands" section of an .ebextensions file executed in Elastic Beanstalk?

A. After the application is deployed

B. Before the application version file is extracted

C. After the web server is configured

D. Before the application and web server are set up

In the context of AWS Elastic Beanstalk and the execution of commands defined in the "commands" section of an .ebextensions file, these commands are executed before the application and web server are set up. This allows users to run scripts or commands that may be needed to prepare the environment for the application and ensure that any necessary prerequisites are in place before the application itself is deployed. Running commands at this early stage is particularly useful for installing dependencies, configuring the environment, or making any necessary adjustments to the Elastic Beanstalk instance that are critical for the deployment process. This phase occurs early in the lifecycle of the Elastic Beanstalk deployment process, allowing any modifications needed before the application code is extracted and run. Commands in the "commands" section therefore facilitate a customized setup environment that aligns with specific application needs, ensuring that the necessary infrastructure is ready before the main deployment occurs.

10. What AWS service helps in reducing costs for data transfer?

A. Amazon S3

B. Amazon CloudFront

C. Amazon RDS

D. AWS Lambda

Amazon CloudFront is the correct answer because it is a content delivery network (CDN) service that reduces costs associated with data transfer by caching content closer to users. When end-users request content, CloudFront retrieves it from the nearest edge location, which minimizes the distance the data needs to travel. This not only speeds up content delivery but also reduces the amount of data transferred from the origin server, leading to cost savings. Using CloudFront can also help leverage AWS's cost-effective data transfer rates compared to fetching data directly over the internet. The consolidation of data transfer through CloudFront allows users to take advantage of lower bandwidth costs and improve performance for their applications. In contrast, while other services like Amazon S3 facilitate data storage and can serve content directly, they do not provide the same level of cost savings through efficient data transfer optimization as CloudFront. Services like Amazon RDS and AWS Lambda are primarily focused on database management and serverless computing, respectively, and do not provide significant benefits related to reducing data transfer costs directly.

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://aws-devops.examzify.com

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

SAMPLE