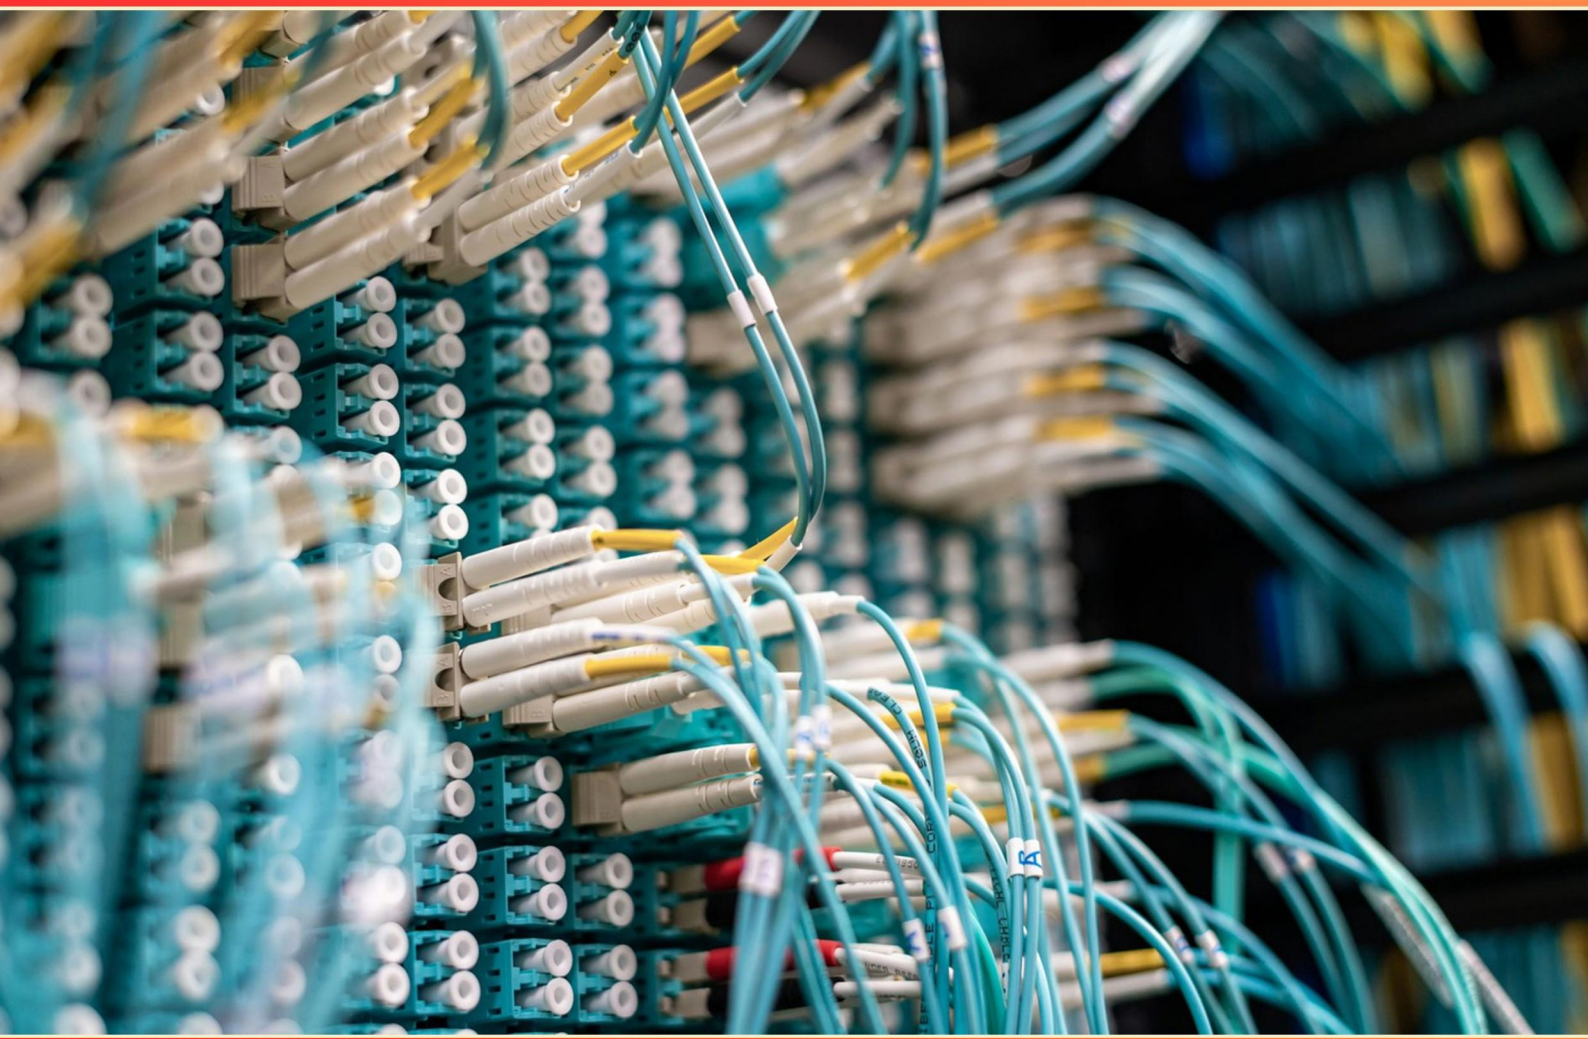
AWS SOLUTIONS ARCHITECT ASSOCIATE PRACTICE TEST (SAMPLE)

STUDY GUIDE



SAMPLE STUDY GUIDE WITH LIMITED QUESTIONS

VISIT US ONLINE FOR THE FULL VERSION STUDY GUIDE

https://aws-solutionsarchitect.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 benefit does AWS Data Pipeline offer for data processing?

A. It allows for real-time analytics

B. It tracks user access patterns

C. It enables scheduled data workflows

D. It enhances cloud storage capacity

2. What is the purpose of Amazon Inspector?

A. To monitor cloud resource costs

B. To assess applications for vulnerabilities

C. To streamline resource creation

D. To analyze machine learning predictions

3. Which AWS service is designed to automate database backups?

A. AWS Backup

B. AWS CloudTrail

C. AWS Lambda

D. AWS Snowball

4. Which AWS service is best suited for building a data warehouse?

A. Amazon RDS

B. Amazon Redshift

C. AWS Lambda

D. Amazon DynamoDB

5. If I want to run a database on an EC2 instance, which storage option is most recommended?

A. RDS

B. S3

C. Glacier

D. EBS

SAMPLE

6. Which service allows for automated backups and recovery for AWS resources?

A. AWS Backup

B. AWS Storage Gateway

C. AWS Data Pipeline

D. AWS Snow Storage

7. Are automated backups enabled by default for new DB Instances in RDS?

A. True

B. False

C. Only for specific databases

D. It depends on the configuration

8. What protocol does AWS S3 use for secure communications?

A. FTP

B. HTTPS

C. HTTP

D. SSH

9. What is the maximum size of an individual object that can be stored in Amazon S3?

A. 1 terabyte

B. 5 terabytes

C. 10 terabytes

D. 15 terabytes

10. Is the AWS platform certified PCI DSS 1.0 compliant?

A. True

B. False

SAMPLE

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. What benefit does AWS Data Pipeline offer for data processing?

A. It allows for real-time analytics

B. It tracks user access patterns

C. It enables scheduled data workflows

D. It enhances cloud storage capacity

AWS Data Pipeline primarily provides the capability to design and manage data workflows that can be scheduled to run at specific intervals or times. This means that users can automate the movement and transformation of data between different AWS services and on-premises data sources according to a predefined schedule. By allowing users to define how and when data is processed, AWS Data Pipeline supports efficient data management and ensures that data is consistently available for analysis or reporting tasks. Scheduled data workflows are particularly beneficial for scenarios where data must be processed regularly, such as daily, hourly, or weekly data ingestion and ETL (Extract, Transform, Load) operations. This capability allows organizations to maintain up-to-date datasets without the need for manual intervention, thereby reducing errors and saving time. Consequently, the other options presented do not align with the primary functions of AWS Data Pipeline. While real-time analytics is essential in many applications, it is more closely associated with services such as Amazon Kinesis. Tracking user access patterns relates to security and monitoring services like AWS CloudTrail. Enhancing cloud storage capacity is linked to services such as Amazon S3 or EBS but is not a direct function of AWS Data Pipeline.

2. What is the purpose of Amazon Inspector?

A. To monitor cloud resource costs

B. To assess applications for vulnerabilities

C. To streamline resource creation

D. To analyze machine learning predictions

Amazon Inspector is designed specifically to assess applications for vulnerabilities and provide security assessments. This service helps developers and security teams ensure that their applications are secure by scanning for common vulnerabilities and deviations from best practices. By automatically evaluating and reporting on the security state of applications deployed in AWS, Amazon Inspector facilitates proactive identification of potential security issues before they can be exploited by malicious actors. It leverages a set of predefined rules and compliance templates, offering insights into various aspects of application security, including network configurations, software dependencies, and layered defenses. The information provided by Amazon Inspector enables teams to take corrective action to bolster security, thus enhancing the overall security posture of their applications. This focus on identifying vulnerabilities sets Amazon Inspector apart from other services that deal with resource monitoring, resource creation, or machine learning analytics, which serve distinct and separate purposes in the AWS ecosystem.

SAMPLE

3. Which AWS service is designed to automate database backups?

A. AWS Backup

B. AWS CloudTrail

C. AWS Lambda

D. AWS Snowball

AWS Backup is specifically designed to automate the backup processes for various AWS services, including databases. It provides a centralized way to manage backups across services such as Amazon RDS, Amazon DynamoDB, and EFS, allowing you to configure backup plans that automate the scheduling and retention of backups. This service simplifies the management of backups, ensuring that they are performed consistently and reliably according to the policies you set. In contrast, AWS CloudTrail focuses on logging and monitoring account activity for security and compliance rather than on backup automation. AWS Lambda is a serverless compute service that allows you to run code in response to events, but it is not inherently related to database backup management. AWS Snowball is a physical data transport solution used for transferring large amounts of data into and out of AWS and does not provide backup capabilities. Therefore, AWS Backup is the most suitable choice for automating database backups.

4. Which AWS service is best suited for building a data warehouse?

A. Amazon RDS

B. Amazon Redshift

C. AWS Lambda

D. Amazon DynamoDB

Amazon Redshift is specifically designed for building data warehouses and is optimized for analyzing large volumes of structured data. It is a fully managed, petabyte-scale data warehouse service that supports complex queries and data analytics. Redshift leverages columnar storage, which allows for efficient compression and more effective querying of large datasets, making it highly suitable for business intelligence and reporting tasks. With Redshift, you can quickly run complex queries across vast amounts of data and integrate with various data visualization tools, further enhancing its utility as a data warehousing solution. The architecture of Redshift also supports scalable analytics, meaning organizations can start with minimal resources and scale up as data requirements grow. Other options, while valuable AWS services, do not specialize in the requirements of data warehousing. For instance, Amazon RDS is used for relational databases and may not handle large-scale analytics tasks efficiently. AWS Lambda is a serverless compute service designed for running code in response to events, rather than managing large volumes of structured data. Amazon DynamoDB is a NoSQL database service suited for fast and flexible data storage but lacks the specific querying capabilities and optimizations that data warehousing requires. Thus, among these choices, Amazon Redshift stands out as the most fitting service for building a data warehouse

SAMPLE

5. If I want to run a database on an EC2 instance, which storage option is most recommended?

A. RDS

B. S3

C. Glacier

D. EBS

Running a database on an EC2 instance requires a storage solution that provides low-latency access, consistent performance, and the capability to perform frequent I/O operations. EBS (Elastic Block Store) is specifically designed for this kind of use case. It offers block-level storage that can be attached to an EC2 instance, making it ideal for databases where you need to read and write data quickly. EBS volumes are persistent, which means that the data stored on them is retained even if the EC2 instance is stopped or terminated. Additionally, EBS supports features like snapshots for data backup and restoration, which is crucial for maintaining database integrity and availability. In comparison, while RDS is a managed database service, it is not the storage option you would directly choose for running a database on EC2 since RDS itself abstracts the underlying storage. S3 is designed for object storage, which is better suited for storing unstructured data rather than for databases requiring fast access. Glacier is a long-term archival storage service, making it unsuitable for the needs of a running database that requires quick access to its data. These characteristics make EBS the most recommended storage option for databases running on EC2 instances, as it meets the performance and reliability requirements for such workloads

6. Which service allows for automated backups and recovery for AWS resources?

A. AWS Backup

B. AWS Storage Gateway

C. AWS Data Pipeline

D. AWS Snow Storage

AWS Backup is designed specifically for creating and managing backups across various AWS services in a centralized manner. It provides automated backup capabilities for services like Amazon EBS (Elastic Block Store), Amazon RDS (Relational Database Service), Amazon DynamoDB, Amazon EFS (Elastic File System), and AWS Storage Gateway, among others. This service allows users to set up backup policies, manage retention periods, and define backup schedules, ensuring that critical data is regularly backed up and can be restored efficiently when needed. In contrast, AWS Storage Gateway is a hybrid cloud storage service that connects on-premises environments with cloud storage, primarily used for integrating on-premises applications with cloud storage but does not specifically provide a comprehensive backup solution. AWS Data Pipeline is focused on moving and processing data between different AWS services and does not address backup and recovery directly. AWS Snow Storage refers to AWS Snow Family, which is designed for data transfer and edge computing rather than for managing backup and recovery processes. Thus, AWS Backup is the most suitable choice for automated backups and recovery of AWS resources.

SAMPLE

7. Are automated backups enabled by default for new DB Instances in RDS?

A. True

B. False

C. Only for specific databases

D. It depends on the configuration

Automated backups are indeed enabled by default for new DB Instances in Amazon RDS. When you launch a new RDS instance, AWS automatically enables automated backups unless you explicitly disable this feature during the creation process. The automated backup system allows for point-in-time recovery, which means you can restore your database instance to any specific time within the retention period, which is typically set to 7 days by default but can be extended to a maximum of 35 days. This feature is part of RDS's managed service offerings, which aim to simplify database management by handling routine tasks such as backups and patching. The fact that automated backups are enabled by default ensures that users can quickly restore their databases without manual intervention, providing a high level of data durability and availability in the event of data loss or corruption. The other available choices do not accurately reflect the standard behavior of RDS when creating new instances, therefore reinforcing the rationale that automated backups are a default setting designed to enhance database reliability from the start.

8. What protocol does AWS S3 use for secure communications?

A. FTP

B. HTTPS

C. HTTP

D. SSH

Amazon S3 (Simple Storage Service) employs HTTPS (Hypertext Transfer Protocol Secure) for secure communications. HTTPS is an extension of HTTP and incorporates SSL/TLS protocols to encrypt data in transit between clients and the S3 service. This encryption ensures that sensitive data, such as files stored in S3, is protected from eavesdropping and man-in-the-middle attacks, maintaining the confidentiality and integrity of the information being transmitted. Using HTTPS is essential for secure cloud storage operations, as it safeguards the data not only while being uploaded to S3 but also during requests to retrieve that data. This is particularly important for businesses that handle sensitive information, as they must comply with regulations regarding data protection. Other protocols listed, such as FTP (File Transfer Protocol) and SSH (Secure Shell), do not apply to S3 communications. FTP, for instance, is often used for transferring files but lacks built-in encryption, making it unsuitable for transmitting sensitive data. SSH is primarily a protocol for secure remote server access rather than for data storage communications. HTTP, while functional, does not provide the security benefits that HTTPS does since it transmits data in plain text. Thus, the correct choice demonstrates a fundamental aspect of securing cloud storage interactions with AWS S3.

SAMPLE

9. What is the maximum size of an individual object that can be stored in Amazon S3?

A. 1 terabyte

B. 5 terabytes

C. 10 terabytes

D. 15 terabytes

The maximum size of an individual object that can be stored in Amazon S3 is indeed 5 terabytes. This limit is significant as it allows users to store very large files, such as high-resolution videos, large datasets, or backups of entire databases, in a highly durable and available object storage service. To upload objects larger than 5 gigabytes, AWS provides a multipart upload capability, which allows you to upload a single object as a set of parts instead of a single upload. This feature not only helps in managing large uploads more efficiently but also improves the resilience and performance of the upload process. The capacity of Amazon S3 to handle individual objects up to 5 terabytes aligns with its design as a scalable object storage system capable of managing vast amounts of unstructured data, making it suitable for a variety of use cases across different industries.

10. Is the AWS platform certified PCI DSS 1.0 compliant?

A. True

B. False

The statement about whether the AWS platform is certified PCI DSS 1.0 compliant highlights an aspect of AWS's commitment to security and compliance in the cloud. PCI DSS, which stands for Payment Card Industry Data Security Standard, outlines requirements for organizations that handle credit card information to ensure that data is protected. The correct assertion regarding AWS's compliance with PCI DSS indicates that AWS has pursued and achieved compliance with various versions of the PCI DSS. However, it should be noted that compliance with a version as old as 1.0 is not typical for most service providers today. In fact, the PCI DSS standards evolve over time, and organizations are generally expected to comply with the most current version applicable at the time. AWS is compliant with PCI DSS, but the specifics may pertain more to ensuring adherence to the latest standards rather than older iterations like 1.0. Therefore, while AWS does comply with the PCI DSS framework, stating compliance with PCI DSS 1.0 specifically may not represent the most accurate current practice. This context shows the importance of understanding the versioning in compliance standards and the continuous updates required in security practices within cloud services.

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-solutionsarchitect.examzify.com

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

SAMPLE