


AWS CERTIFIED DEVELOPER ASSOCIATE PRACTICE EXAM (SAMPLE)

STUDY GUIDE



SAMPLE STUDY GUIDE WITH LIMITED QUESTIONS

**VISIT US ONLINE FOR
THE FULL VERSION STUDY GUIDE**

<https://awscertifieddevassociate.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.

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. Was SQS the second service offered by AWS?**
 - A. True
 - B. False
 - C. It is unknown
 - D. It was the first service
- 2. What is the read throughput setting for an application that needs to read 25 items of 13KB per second using eventually consistent reads?**
 - A. 25
 - B. 50
 - C. 75
 - D. 100
- 3. What is a key benefit of using SWF for workflows?**
 - A. It reduces the amount of code needed
 - B. It automatically optimizes task execution
 - C. It reliably tracks the status of processing steps
 - D. It eliminates the need for worker programs
- 4. Which indexes in DynamoDB allow for a choice of sort keys?**
 - A. Global secondary indexes
 - B. Local secondary indexes
 - C. Both Global and Local secondary indexes
 - D. No indexes provide sort keys
- 5. What is true about a scan operation in DynamoDB?**
 - A. A scan reads every item in a table or secondary index
 - B. A scan is more efficient than a query
 - C. A scan cannot read items from multiple tables
 - D. A scan returns items based on a specified key
- 6. Which AWS service is ideal for backing up real-time streaming data?**
 - A. Kinesis Firehose
 - B. Amazon Glacier
 - C. Amazon S3
 - D. Amazon RDS

- 7. How many eventually consistent reads per second does one read capacity unit provide?**
- A. One
 - B. Two
 - C. Three
 - D. Four
- 8. Which statement is true regarding EBS storage and instance store in EC2?**
- A. EBS is ephemeral, and data is lost when the instance stops
 - B. Instance store is persistent, and data remains after stopping
 - C. EBS is persistent, and data remains when an instance stops
 - D. Both EBS and instance store are persistent
- 9. Does Amazon S3 provide unlimited storage capacity?**
- A. Yes
 - B. No
 - C. Only in specific regions
 - D. Only for paying customers
- 10. How can instances in a private subnet initiate a connection to the internet using a NAT instance?**
- A. Forward requests through a public IP
 - B. Send traffic through an Elastic Load Balancer
 - C. Use the NAT instance as a proxy
 - D. Directly assign them public IPs

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. Was SQS the second service offered by AWS?

- A. True
- B. False**
- C. It is unknown
- D. It was the first service

The correct answer is false because Amazon Simple Queue Service (SQS) was not the second service offered by AWS. In fact, the first service introduced by AWS was Amazon S3 (Simple Storage Service), which launched in March 2006. SQS was released later, in January 2004, after the initial batch of services. AWS began with a strong focus on providing infrastructure services, and while SQS is one of the foundational messaging services in cloud architecture, it is important to recognize the chronology of service launches. Understanding the history of AWS service releases is essential for grasping the evolution of cloud computing solutions and how they fit into the overall AWS ecosystem.

2. What is the read throughput setting for an application that needs to read 25 items of 13KB per second using eventually consistent reads?

- A. 25
- B. 50**
- C. 75
- D. 100

When determining the read throughput setting for an application using eventually consistent reads, it is important to understand how AWS DynamoDB calculates read capacity units. Each read capacity unit allows you to read one item up to 4 KB in size per second when using eventually consistent reads. In this case, the application needs to read 25 items that are each 13 KB in size. To calculate the required read capacity units: 1. Since each item is 13 KB, you can determine how many read capacity units are needed by dividing the item size by the maximum size per read capacity unit. 2. Each 13 KB item requires 4 read capacity units because $13 \text{ KB} / 4 \text{ KB} = 3.25$, which rounds up to 4 units. 3. Therefore, for 25 items: $25 \text{ items} * 4 \text{ read capacity units} = 100 \text{ read capacity units}$. Since eventually consistent reads allow reading two items for one read capacity unit, you can divide the total read capacity units by 2: $100 \text{ read capacity units} / 2 = 50$. Thus, the required read throughput setting for this scenario is indeed 50. This number ensures that the application can handle the required read load efficiently using the eventually consistent read strategy.

3. What is a key benefit of using SWF for workflows?

- A. It reduces the amount of code needed
- B. It automatically optimizes task execution
- C. It reliably tracks the status of processing steps**
- D. It eliminates the need for worker programs

Using Amazon Simple Workflow Service (SWF) offers the significant benefit of reliably tracking the status of processing steps throughout the workflow. SWF is designed to manage and coordinate distributed applications, enabling developers to define workflows that can encapsulate complex task executions involving multiple workers. One of the core features of SWF is its ability to maintain the state of each task within the workflow, allowing for better visibility and management of task execution. This capability ensures that developers can determine where in the process a specific task is at any given point. If a failure occurs, the workflow can be resumed from the point of failure rather than needing to restart the entire process, which enhances both reliability and efficiency. This tracking mechanism also aids in handling retries, managing timeouts, and orchestrating parallel processing, which might otherwise become challenging in distributed systems without such oversight. Overall, this leads to improved workflow robustness and operational integrity, making it easier for developers to create resilient applications.

4. Which indexes in DynamoDB allow for a choice of sort keys?

- A. Global secondary indexes
- B. Local secondary indexes**
- C. Both Global and Local secondary indexes
- D. No indexes provide sort keys

Local secondary indexes (LSIs) in DynamoDB allow for a choice of sort keys while sharing the same partition key as the base table. This means that for each item in the main table, you can create multiple local secondary indexes, each with a different sort key. This capability is particularly useful when you want to query data based on different attributes or ranges of values while keeping the same partitioning scheme. By allowing multiple sort keys, LSIs enable more flexible querying options on the same set of items, enhancing the ability to retrieve data efficiently based on additional criteria beyond the primary key. It's important to note that while global secondary indexes (GSIs) can provide additional flexibility—allowing for different partition keys and sort keys across potentially different data distributions—they do not utilize the same partition key, thus making LSIs unique in their ability to share this aspect. This distinction underlines the local secondary index's capability to serve specific use cases where sorting based on multiple attributes under a uniform partitioning strategy is desired.

5. What is true about a scan operation in DynamoDB?

- A. A scan reads every item in a table or secondary index**
- B. A scan is more efficient than a query
- C. A scan cannot read items from multiple tables
- D. A scan returns items based on a specified key

A scan operation in DynamoDB does indeed read every item in a table or secondary index. This means that when performing a scan, the operation examines every single item in the specified table or index, retrieving the data irrespective of the item's keys or any sort of filter. It does not require any specific query criteria and can potentially return a large amount of data, depending on the size of the table. This characteristic of a scan operation is crucial for understanding its impact on performance and cost. Since it processes every item, it tends to be less efficient than a query, which retrieves items based on specific key conditions and can narrow down the search significantly. Additionally, scans are limited to one table or one secondary index at a time, meaning they cannot pull data from multiple tables simultaneously. While scans can utilize filters to reduce the amount of data returned, they do not operate based on keys like queries do; rather, they can return items regardless of their keys unless additional conditions are applied.

6. Which AWS service is ideal for backing up real-time streaming data?

A. Kinesis Firehose

B. Amazon Glacier

C. Amazon S3

D. Amazon RDS

Kinesis Firehose is the ideal service for backing up real-time streaming data. This service is specifically designed to handle large streams of data in real-time, allowing you to reliably capture streaming data, transform it if necessary, and then store it in a variety of destinations such as Amazon S3, Amazon Redshift, or Amazon Elasticsearch Service. The ability of Kinesis Firehose to process data continuously and deliver it with minimal latency makes it particularly suited for scenarios where immediate insights from streaming data are crucial. It is capable of automatically scaling to match the throughput of incoming data, which is essential for applications that generate a continuous flow of information. Other services listed, while useful in various contexts, do not specialize in real-time streaming data backing. Amazon S3, for example, is designed for object storage and can be a destination for the data processed by Kinesis Firehose but does not inherently manage real-time streaming data. Amazon Glacier is optimized for long-term archival storage and is not suited for real-time access. Amazon RDS is meant for relational database management and not for handling continuous streaming data. Hence, Kinesis Firehose stands out as the best option for this purpose.

7. How many eventually consistent reads per second does one read capacity unit provide?

A. One

B. Two

C. Three

D. Four

When using Amazon DynamoDB, one read capacity unit allows for one eventually consistent read per second for items up to 4 KB in size. However, due to the nature of eventually consistent reads, which can return stale data that is not up-to-date with the latest writes, it's important to note that the read can utilize cached data or data from replicas, allowing DynamoDB to potentially serve more reads under the same unit cost. In the context of multiple-choice answers, the notion that one read capacity unit provides two eventually consistent reads per second is inaccurate. Instead, it is designed to facilitate one read operation efficiently within the constraints of the total KB limit. Therefore, a clear understanding of read capacity units helps clarify their behavior regarding read operations and throughput management in DynamoDB. This understanding also informs decisions on how to provision capacity for applications needing high throughput, ensuring that developers allocate resources appropriately based on their anticipated read patterns and data sizes.

8. Which statement is true regarding EBS storage and instance store in EC2?

- A. EBS is ephemeral, and data is lost when the instance stops
- B. Instance store is persistent, and data remains after stopping
- C. EBS is persistent, and data remains when an instance stops**
- D. Both EBS and instance store are persistent

EBS (Elastic Block Store) volumes are designed to provide durable, persistent storage for Amazon EC2 instances. When an EC2 instance that uses EBS storage is stopped, the data stored on EBS persists, allowing you to retain your data even when the instance is not running. This is particularly useful for applications that require reliable storage, as you can stop and start instances without losing your data. In contrast, instance store volumes are ephemeral, meaning that any data stored on them is lost if the instance is stopped or terminated. This characteristic makes instance store suitable for temporary storage needs, but it is not reliable for data that needs to persist beyond the life of the instance. Thus, the statement that EBS is persistent and retains its data when an instance stops accurately reflects the nature of EBS in the context of EC2 storage.

9. Does Amazon S3 provide unlimited storage capacity?

- A. Yes**
- B. No
- C. Only in specific regions
- D. Only for paying customers

Amazon S3 does indeed provide unlimited storage capacity. The service is designed to seamlessly scale to accommodate a virtually limitless amount of data. As a cloud storage solution, S3 allows users to store and retrieve any amount of data at any time. This scalability is one of the primary advantages of using cloud services like S3, as it eliminates the need for organizations to manage and provision hardware for storage separately. Data can be uploaded, stored, and managed through S3 without the constraints that typically come with on-premises storage solutions, such as fixed disk space or hardware failures. Once users start utilizing S3, they can continue to add data without hitting a predefined limit. This makes it a flexible solution suitable for a wide range of applications, from backup and archiving to big data analytics. The other options mentioned do not reflect the inherent capabilities of S3. There aren't any geographical limitations affecting storage capacity, nor is there a restriction based on user payment status regarding the overall ability to store data.

10. How can instances in a private subnet initiate a connection to the internet using a NAT instance?

- A. Forward requests through a public IP
- B. Send traffic through an Elastic Load Balancer
- C. Use the NAT instance as a proxy**
- D. Directly assign them public IPs

Instances in a private subnet can initiate a connection to the internet using a NAT (Network Address Translation) instance by using it as a proxy. When instances in the private subnet send their outgoing traffic to the NAT instance, the NAT instance modifies the source IP address of the traffic to its own public IP address before forwarding it to the internet. In doing this, the NAT instance acts as an intermediary, enabling the private instances to access external resources while still not being directly exposed to the internet. This method ensures that the instances retain their private IP addresses internally while still being able to communicate with external networks. When responses from the internet return to the NAT instance, it translates the destination address back to the original private IP of the requesting instance, allowing the traffic to flow back into the private subnet seamlessly. The other options do not effectively describe how a NAT instance facilitates internet connectivity for private subnet instances. Forwarding requests through a public IP does not integrate with the NAT instance functionality. Using an Elastic Load Balancer is not applicable because load balancers distribute incoming traffic rather than provide outbound internet access. Directly assigning public IPs to instances in a private subnet defeats the purpose of being in a private subnet, as those instances would then be exposed to the internet directly

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://awscertifieddevassociate.examzify.com>

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

SAMPLE