

AWS Certified Solutions Architect Professional Practice (SAP-C02) Practice Exam (Sample)

Study Guide



Everything you need from our exam experts!

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	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. 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 service is the fully managed offering for Apache Kafka on AWS?**
 - A. Amazon Kinesis Data Streams**
 - B. AWS Glue**
 - C. Amazon Athena**
 - D. Amazon MSK**

- 2. Which ETL service reliably captures, transforms, and delivers streaming data to data lakes, data stores, and analytics services?**
 - A. Amazon Kinesis Data Streams**
 - B. Amazon MSK**
 - C. Amazon Kinesis Data Firehose**
 - D. AWS Data Pipeline**

- 3. Which AWS service can be used to update the firmware of IoT devices?**
 - A. AWS IoT Core**
 - B. AWS IoT Device Management**
 - C. Amazon API Gateway**
 - D. Amazon AppStream 2.0**

- 4. In migration strategy, what is the concept of a "pilot light"?**
 - A. Run a full production environment in the cloud before migrating on-prem systems.**
 - B. Run a small, always-on portion of the environment in the cloud to speed up full production deployment later.**
 - C. Migrate only the data and not the applications.**
 - D. Use a one-time data snapshot for cutover.**

- 5. Which service provides a hosted notebook environment to run data science and machine learning experiments without managing servers?**
 - A. Amazon Polly**
 - B. Amazon Textract**
 - C. Amazon Translate**
 - D. Amazon SageMaker**

- 6. For migrating legacy on-prem storage with large data volumes, which pattern helps minimize downtime?**
- A. Use AWS DMS for ongoing replication for initial load, with a cutover plan; optionally Snowball/Storage Gateway for initial data transfer.**
 - B. Migrate all data and services in a single weekend without replication.**
 - C. Use only network optimization to speed transfers.**
 - D. Use manual data validation after cutover.**
- 7. Which AWS service is based on the same deep learning technology used by Amazon's computer vision scientists to analyze billions of images and videos daily?**
- A. Amazon Translate**
 - B. AWS CLI**
 - C. Amazon Polly**
 - D. Amazon Rekognition**
- 8. Which service is used to create customized licensing rules that emulate the terms of licensing agreements, and then enforce these rules?**
- A. AWS Compute Optimizer**
 - B. AWS License Manager**
 - C. Amazon Managed Grafana**
 - D. Amazon Managed Service for Prometheus**
- 9. Which service helps you securely host code repositories in AWS and supports private access by default?**
- A. AWS CodeArtifact**
 - B. AWS CodeCommit**
 - C. AWS X-Ray**
 - D. AWS Cloud9**

10. Which fully managed service enables you to build and run applications that use Apache Kafka to process streaming data?

- A. Amazon MSK**
- B. Amazon Kinesis Data Streams**
- C. AWS Glue**
- D. Amazon Athena**

SAMPLE

Answers

SAMPLE

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

SAMPLE

Explanations

SAMPLE

1. Which service is the fully managed offering for Apache Kafka on AWS?

- A. Amazon Kinesis Data Streams**
- B. AWS Glue**
- C. Amazon Athena**
- D. Amazon MSK**

Amazon MSK is the fully managed service for Apache Kafka on AWS. It provisions and manages the Kafka brokers, handles software patching and maintenance, and takes care of scaling and automatic failover, so you don't have to operate the Kafka cluster yourself. You still use the Kafka client APIs to produce and consume data, manage topics and partitions, and set data retention, but the underlying infrastructure and reliability are managed by AWS. This differs from the other options because they are not Apache Kafka. Kinesis Data Streams is AWS's own streaming service with a different API and semantics. AWS Glue is an ETL and data integration service, not a streaming platform. Amazon Athena is a serverless query service for data in S3, not a streaming system. If you need Kafka compatibility and the ecosystem around it while offloading cluster maintenance, MSK is the right choice.

2. Which ETL service reliably captures, transforms, and delivers streaming data to data lakes, data stores, and analytics services?

- A. Amazon Kinesis Data Streams**
- B. Amazon MSK**
- C. Amazon Kinesis Data Firehose**
- D. AWS Data Pipeline**

The essential idea here is a service that handles streaming data end-to-end for ETL: it captures the stream, can transform the data on the fly, and reliably delivers it to data lakes, data stores, or analytics services. Amazon Kinesis Data Firehose is built for this pattern. It is a fully managed, serverless service that ingests streaming records and delivers them directly to destinations such as Amazon S3 for data lakes, Redshift for data warehousing, Elasticsearch Service for analytics, or Splunk. It can perform transformations by integrating with a Lambda function, allowing you to modify records as they pass through, before they're stored or analyzed. Firehose also handles buffering, retries, and automatic scaling, ensuring reliable delivery without you having to manage infrastructure. In contrast, a pure streaming ingestion service like Kinesis Data Streams requires separate processing to transform and load data, a managed Kafka service like MSK focuses on broker-level streaming, and a workflow tool like AWS Data Pipeline is geared toward batch processing rather than continuous streaming ETL.

3. Which AWS service can be used to update the firmware of IoT devices?

- A. AWS IoT Core**
- B. AWS IoT Device Management**
- C. Amazon API Gateway**
- D. Amazon AppStream 2.0**

Updating firmware across a fleet of IoT devices is a use case for over-the-air (OTA) deployments, where you orchestrate software updates remotely. AWS IoT Device Management provides this capability through OTA updates (Jobs). You store the firmware image in S3, create an OTA update, and target devices by thing or group. The service handles distributing the new firmware, monitoring progress, and collecting status from devices as they install, with options to roll back if needed. This fleet-wide orchestration is exactly what you need for firmware updates. AWS IoT Core handles secure connectivity, messaging, and device shadows, but it doesn't offer built-in OTA fleet update orchestration. API Gateway is for exposing APIs, and AppStream 2.0 is for streaming applications to users, not managing firmware on IoT devices.

4. In migration strategy, what is the concept of a "pilot light"?

- A. Run a full production environment in the cloud before migrating on-prem systems.**
- B. Run a small, always-on portion of the environment in the cloud to speed up full production deployment later.**
- C. Migrate only the data and not the applications.**
- D. Use a one-time data snapshot for cutover.**

In a pilot light approach, you keep a small, always-on portion of the environment in the cloud that runs the essential components and data. This baseline is continuously kept up to date and ready, so when you're ready to cut over, you can quickly scale up to full production in the cloud. The idea is to minimize upfront infrastructure while still having a ready-to-go footprint that speeds deployment and reduces downtime during migration. This works well because keeping a minimal cloud footprint with automated replication of critical data means you don't start from scratch at cutover, and you can ramp up to full capacity more quickly and with less risk. Rationale for the other options: starting with a full production cloud environment defeats the cost-saving and speed goals of pilot light, because you've already moved everything before the migration. Migrating only data ignores the need to run and test the application and services in the cloud. A one-time data snapshot for cutover lacks the continuously running, minimal cloud environment that facilitates rapid and reliable transition.

5. Which service provides a hosted notebook environment to run data science and machine learning experiments without managing servers?

- A. Amazon Polly**
- B. Amazon Textract**
- C. Amazon Translate**
- D. Amazon SageMaker**

Running data science and machine learning experiments without managing servers is enabled by a fully managed ML development environment. Amazon SageMaker offers hosted notebooks (SageMaker Notebooks and SageMaker Studio) that run in the cloud with the infrastructure managed by the service. You can create and run Jupyter notebooks, prepare data, train models, tune hyperparameters, and deploy endpoints—all without provisioning or maintaining servers. The service handles provisioning, scaling, and security so you focus on experiments and models. The other options are not about hosted notebook environments: Polly is a text-to-speech service, Textract extracts text from documents, and Translate translates text between languages. They don't provide the end-to-end notebook-based workflow for data science experiments, hence SageMaker is the correct choice.

6. For migrating legacy on-prem storage with large data volumes, which pattern helps minimize downtime?

- A. Use AWS DMS for ongoing replication for initial load, with a cutover plan; optionally Snowball/Storage Gateway for initial data transfer.**
- B. Migrate all data and services in a single weekend without replication.**
- C. Use only network optimization to speed transfers.**
- D. Use manual data validation after cutover.**

Minimizing downtime during a large on-prem to AWS migration hinges on keeping the source and target in sync until the moment you cut over. The best pattern uses ongoing replication to load the initial data while the on-prem system remains active, plus a well-planned cutover to redirect traffic to the cloud. AWS DMS supports change data capture, so it can continuously replicate changes from on-prem to the target during the migration. This means the destination stays up to date, and when you flip the switch, only a small delta needs to be applied, keeping the cutover brief. For very large datasets, you can transfer the initial bulk offline—using Snowball or Storage Gateway—to reduce network impact, then resume replication to catch any remaining changes until cutover. A plan that tries to move everything in one weekend without replication would require long downtime and risks data drift. Relying only on network speed won't guarantee synchronization, and performing manual data validation after cutover tends to extend the downtime rather than shorten it.

7. Which AWS service is based on the same deep learning technology used by Amazon's computer vision scientists to analyze billions of images and videos daily?

- A. Amazon Translate**
- B. AWS CLI**
- C. Amazon Polly**
- D. Amazon Rekognition**

Amazon Rekognition is the AWS service built on the same deep learning vision technology used by Amazon's computer vision scientists to analyze billions of images and videos daily. It exposes APIs that let you detect objects, scenes, and activities in images; perform facial analysis and comparison; recognize text within visuals (OCR); and even analyze video frames for motion and events. This family of deep learning models is specifically tuned for image and video understanding, so Rekognition directly embodies that technology in a managed service that you can call from your apps. Translate focuses on language translation, Polly handles text-to-speech, and the AWS CLI is just a command-line interface for interacting with AWS services, not a vision capability. So the service tied to image and video analysis using that deep learning foundation is Amazon Rekognition.

8. Which service is used to create customized licensing rules that emulate the terms of licensing agreements, and then enforce these rules?

- A. AWS Compute Optimizer**
- B. AWS License Manager**
- C. Amazon Managed Grafana**
- D. Amazon Managed Service for Prometheus**

Defining and enforcing software licensing terms across your AWS environment is handled by AWS License Manager. It lets you create license configurations with custom rules that mirror the terms of your licensing agreements—such as the number of licenses, permitted instance types, regions, or usage limits—and then enforce those rules by comparing actual usage against your entitlements. This centralized approach helps you track licenses, stay compliant, and automatically prevent overuse. The other services don't address licensing management or enforcement: one focuses on optimizing cost and performance, another provides a managed visualization platform, and the last is a managed service for collecting and querying metrics. So AWS License Manager is the appropriate choice for creating and enforcing customized licensing rules.

9. Which service helps you securely host code repositories in AWS and supports private access by default?

- A. AWS CodeArtifact
- B. AWS CodeCommit**
- C. AWS X-Ray
- D. AWS Cloud9

Securely hosting code repositories with private access by default is provided by a fully managed Git-based code hosting service in AWS. This service stores your Git repositories and uses IAM for authentication and authorization, making repositories private by default so only authorized users and roles can access them. It also offers encryption at rest and in transit and integrates smoothly with other AWS developer tools, so you don't have to manage your own git server. The other services mentioned serve different purposes: a package repository for software libraries, a tracing tool for distributed applications, and a cloud-based integrated development environment, none of which are designed primarily to host private code repositories with IAM-driven access controls.

10. Which fully managed service enables you to build and run applications that use Apache Kafka to process streaming data?

- A. Amazon MSK**
- B. Amazon Kinesis Data Streams
- C. AWS Glue
- D. Amazon Athena

The main idea is a fully managed service that provides Apache Kafka as the underlying streaming platform. Amazon MSK fits this exactly because it runs Apache Kafka for you as a managed service. With MSK, you don't have to worry about provisioning Kafka brokers, configuring Zookeeper (in older Kafka setups) or handling patching, failover, scaling, or monitoring. AWS handles the operational heavy lifting, while you run your applications that produce and consume streaming data using the familiar Kafka APIs, clients, and ecosystem (Kafka Connect, Kafka Streams, etc.). This makes it easy to reuse existing Kafka-based applications with minimal changes, and you can still apply Kafka security and VPC controls, encryption, and monitoring through AWS. The other options don't provide Apache Kafka as a managed service. Kinesis Data Streams is AWS's own streaming service with a different API, not Kafka-compatible. AWS Glue is an ETL/service integration tool, and Amazon Athena is a query service for data in S3.

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://awssapc02.examzify.com>

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

SAMPLE