aws training and certification

AWS Certified Solutions Architect – Associate (SAA-C02) Exam Guide

Introduction

The AWS Certified Solutions Architect – Associate (SAA-C02) exam is intended for individuals who perform in a solutions architect role. The exam validates a candidate's ability to design secure and robust solutions by using AWS technologies.

The exam also validates a candidate's ability to complete the following tasks:

- Design a solution by using appropriate AWS services and by following architectural principles based on requirements
- Provide implementation guidance based on best practices to the organization throughout the workload lifecycle

Target candidate description

The target candidate should have at least 1 year of hands-on experience designing secure, high-performing, cost-effective, highly available, and scalable systems by using AWS services.

Recommended AWS knowledge

The target candidate should have the following knowledge:

- Hands-on experience using compute, networking, storage, management, and database AWS services
- The ability to identify and define technical requirements for a solution that involves AWS technology
- The ability to identify which AWS services meet a given technical requirement
- An understanding of best practices for building well-architected solutions on AWS
- An understanding of the AWS global infrastructure
- An understanding of AWS security services and features in relation to traditional services

What is considered out of scope for the target candidate?

The following is a non-exhaustive list of related job tasks that the target candidate is not expected to be able to perform. These items are out of scope for the exam:

- Design a complex, hybrid network architecture
- Design identity federation within multiple accounts
- Design an architecture that meets compliance requirements
- Incorporate specialized services in a design
- Develop deployment strategies
- Create a migration strategy for complex multi-tier applications

For a detailed list of specific tools and technologies that might be covered on the exam, as well as a list of in-scope AWS services, refer to the Appendix.

Exam content

Response types

There are two types of questions on the exam:

- **Multiple choice:** Has one correct response and three incorrect responses (distractors)
- **Multiple response:** Has two or more correct responses out of five or more response options

Select one or more responses that best complete the statement or answer the question. Distractors, or incorrect answers, are response options that a candidate with incomplete knowledge or skill might choose. Distractors are generally plausible responses that match the content area.

Unanswered questions are scored as incorrect; there is no penalty for guessing. The exam includes 50 questions that will affect your score.

Unscored content

The exam includes 15 unscored questions that do not affect your score. AWS collects information about candidate performance on these unscored questions to evaluate these questions for future use as scored questions. These unscored questions are not identified on the exam.

Exam results

The AWS Certified Solutions Architect – Associate exam is a pass or fail exam. The exam is scored against a minimum standard established by AWS professionals who follow certification industry best practices and guidelines.

Your results for the exam are reported as a scaled score of 100–1,000. The minimum passing score is 720. Your score shows how you performed on the exam as a whole and whether or not you passed. Scaled scoring models help equate scores across multiple exam forms that might have slightly different difficulty levels.

Your score report could contain a table of classifications of your performance at each section level. This information provides general feedback about your exam performance. The exam uses a compensatory scoring model, which means that you do not need to achieve a passing score in each section. You need to pass only the overall exam.

Each section of the exam has a specific weighting, so some sections have more questions than other sections have. The table contains general information that highlights your strengths and weaknesses. Use caution when interpreting section-level feedback.

Content outline

This exam guide includes weightings, test domains, and objectives for the exam. It is not a comprehensive listing of the content on the exam. However, additional context for each of the objectives is available to help guide your preparation for the exam. The following table lists the main content domains and their weightings. The table precedes the complete exam content outline, which includes the additional context. The percentage in each domain represents only scored content.

Domain	% of Exam
Domain 1: Design Resilient Architectures	30%
Domain 2: Design High-Performing Architectures	28%
Domain 3: Design Secure Applications and Architectures	24%
Domain 4: Design Cost-Optimized Architectures	18%
TOTAL	**100%**

Domain 1: Design Resilient Architectures

1.1 Design a multi-tier architecture solution

- Determine a solution design based on access patterns.
- Determine a scaling strategy for components used in a design.
- Select an appropriate database based on requirements.
- Select an appropriate compute and storage service based on requirements.

1.2 Design highly available and/or fault-tolerant architectures

- Determine the amount of resources needed to provide a fault-tolerant architecture across Availability Zones.
- Select a highly available configuration to mitigate single points of failure.
- Apply AWS services to improve the reliability of legacy applications when application changes are not possible.
- Select an appropriate disaster recovery strategy to meet business requirements.
- Identify key performance indicators to ensure the high availability of the solution.

1.3 Design decoupling mechanisms using AWS services

- Determine which AWS services can be leveraged to achieve loose coupling of components.
- Determine when to leverage serverless technologies to enable decoupling.

1.4 Choose appropriate resilient storage

- Define a strategy to ensure the durability of data.
- Identify how data service consistency will affect the operation of the application.
- Select data services that will meet the access requirements of the application.
- Identify storage services that can be used with hybrid or non-cloud-native applications.

Domain 2: Design High-Performing Architectures

2.1 Identify elastic and scalable compute solutions for a workload

- Select the appropriate instance(s) based on compute, storage, and networking requirements.
- Choose the appropriate architecture and services that scale to meet performance requirements.
- Identify metrics to monitor the performance of the solution.

2.2 Select high-performing and scalable storage solutions for a workload

- Select a storage service and configuration that meets performance demands.
- Determine storage services that can scale to accommodate future needs.

2.3 Select high-performing networking solutions for a workload

- Select appropriate AWS connectivity options to meet performance demands.
- Select appropriate features to optimize connectivity to AWS public services.
- Determine an edge caching strategy to provide performance benefits.
- Select appropriate data transfer service for migration and/or ingestion.

2.4 Choose high-performing database solutions for a workload

- Select an appropriate database scaling strategy.
- Determine when database caching is required for performance improvement.
- Choose a suitable database service to meet performance needs.

Domain 3: Design Secure Applications and Architectures

3.1 Design secure access to AWS resources

- Determine when to choose between users, groups, and roles.
- Interpret the net effect of a given access policy.
- Select appropriate techniques to secure a root account.
- Determine ways to secure credentials using features of AWS IAM.
- Determine the secure method for an application to access AWS APIs.
- Select appropriate services to create traceability for access to AWS resources.

3.2 Design secure application tiers

- Given traffic control requirements, determine when and how to use security groups and network ACLs.
- Determine a network segmentation strategy using public and private subnets.
- Select the appropriate routing mechanism to securely access AWS service endpoints or internet-based resources from Amazon VPC.
- Select appropriate AWS services to protect applications from external threats.

3.3 Select appropriate data security options

- Determine the policies that need to be applied to objects based on access patterns.
- Select appropriate encryption options for data at rest and in transit for AWS services.
- Select appropriate key management options based on requirements.

Domain 4: Design Cost-Optimized Architectures

4.1 Identify cost-effective storage solutions

- Determine the most cost-effective data storage options based on requirements.
- Apply automated processes to ensure that data over time is stored on storage tiers that minimize costs.

4.2 Identify cost-effective compute and database services

- Determine the most cost-effective Amazon EC2 billing options for each aspect of the workload.
- Determine the most cost-effective database options based on requirements.
- Select appropriate scaling strategies from a cost perspective.
- Select and size compute resources that are optimally suited for the workload.
- Determine options to minimize total cost of ownership (TCO) through managed services and serverless architectures.

4.3 Design cost-optimized network architectures

- Identify when content delivery can be used to reduce costs.
- Determine strategies to reduce data transfer costs within AWS.
- Determine the most cost-effective connectivity options between AWS and on-premises environments.

Appendix

Which key tools, technologies, and concepts might be covered on the exam?

The following is a non-exhaustive list of the tools and technologies that could appear on the exam. This list is subject to change and is provided to help you understand the general scope of services, features, or technologies on the exam. The general tools and technologies in this list appear in no particular order. AWS services are grouped according to their primary functions. While some of these technologies will likely be covered more than others on the exam, the order and placement of them in this list is no indication of relative weight or importance:

- Compute
- Cost management
- Database
- Disaster recovery
- High availability
- Management and governance
- Microservices and component decoupling
- Migration and data transfer
- Networking, connectivity, and content delivery
- Security
- Serverless design principles
- Storage

AWS services and features

Analytics:

- Amazon Athena
- Amazon Elasticsearch Service (Amazon ES)
- Amazon EMR
- AWS Glue
- Amazon Kinesis
- Amazon QuickSight

AWS Billing and Cost Management:

- AWS Budgets
- Cost Explorer

Application Integration:

- Amazon Simple Notification Service (Amazon SNS)
- Amazon Simple Queue Service (Amazon SQS)

Compute:

- Amazon EC2
- AWS Elastic Beanstalk
- Amazon Elastic Container Service (Amazon ECS)
- Amazon Elastic Kubernetes Service (Amazon EKS)
- Elastic Load Balancing
- AWS Fargate
- AWS Lambda

Database:

- Amazon Aurora
- Amazon DynamoDB
- Amazon ElastiCache
- Amazon RDS
- Amazon Redshift

Management and Governance:

- AWS Auto Scaling
- AWS Backup
- AWS CloudFormation
- AWS CloudTrail
- Amazon CloudWatch
- AWS Config
- Amazon EventBridge (Amazon CloudWatch Events)
- AWS Organizations
- AWS Resource Access Manager
- AWS Systems Manager
- AWS Trusted Advisor

Migration and Transfer:

- AWS Database Migration Service (AWS DMS)
- AWS DataSync
- AWS Migration Hub
- AWS Server Migration Service (AWS SMS)
- AWS Snowball
- AWS Transfer Family

Networking and Content Delivery:

- Amazon API Gateway
- Amazon CloudFront
- AWS Direct Connect
- AWS Global Accelerator
- Amazon Route 53
- AWS Transit Gateway
- Amazon VPC (and associated features)

Security, Identity, and Compliance:

- AWS Certificate Manager (ACM)
- AWS Directory Service
- Amazon GuardDuty
- AWS Identity and Access Management (IAM)
- Amazon Inspector
- AWS Key Management Service (AWS KMS)
- Amazon Macie
- AWS Secrets Manager
- AWS Shield
- AWS Single Sign-On
- AWS WAF

Storage:

- Amazon Elastic Block Store (Amazon EBS)
- Amazon Elastic File System (Amazon EFS)
- Amazon FSx
- Amazon S3
- Amazon S3 Glacier
- AWS Storage Gateway

AWS Certified Solutions Architect – Associate (SAA-C02)
Sample Exam Questions

1) A customer relationship management (CRM) application runs on Amazon EC2 instances in multiple Availability Zones behind an Application Load Balancer.

If one of these instances fails, what occurs?

A) The load balancer will stop sending requests to the failed instance.
B) The load balancer will terminate the failed instance.
C) The load balancer will automatically replace the failed instance.
D) The load balancer will return 504 Gateway Timeout errors until the instance is replaced.

2) A company needs to perform asynchronous processing, and has Amazon SQS as part of a decoupled architecture. The company wants to ensure that the number of empty responses from polling requests are kept to a minimum.

What should a solutions architect do to ensure that empty responses are reduced?

A) Increase the maximum message retention period for the queue.
B) Increase the maximum receives for the redrive policy for the queue.
C) Increase the default visibility timeout for the queue.
D) Increase the receive message wait time for the queue.

3) A company currently stores data for on-premises applications on local drives. The chief technology officer wants to reduce hardware costs by storing the data in Amazon S3 but does not want to make modifications to the applications. To minimize latency, frequently accessed data should be available locally.

What is a reliable and durable solution for a solutions architect to implement that will reduce the cost of local storage?

A) Deploy an SFTP client on a local server and transfer data to Amazon S3 using AWS Transfer for SFTP.
B) Deploy an AWS Storage Gateway volume gateway configured in cached volume mode.
C) Deploy an AWS DataSync agent on a local server and configure an S3 bucket as the destination.
D) Deploy an AWS Storage Gateway volume gateway configured in stored volume mode.

4) A company runs a public-facing three-tier web application in a VPC across multiple Availability Zones. Amazon EC2 instances for the application tier running in private subnets need to download software patches from the internet. However, the instances cannot be directly accessible from the internet.

Which actions should be taken to allow the instances to download the needed patches? (Select TWO.)

A) Configure a NAT gateway in a public subnet.
B) Define a custom route table with a route to the NAT gateway for internet traffic and associate it with the private subnets for the application tier.
C) Assign Elastic IP addresses to the application instances.
D) Define a custom route table with a route to the internet gateway for internet traffic and associate it with the private subnets for the application tier.
E) Configure a NAT instance in a private subnet.

5) A solutions architect wants to design a solution to save costs for Amazon EC2 instances that do not need to run during a 2-week company shutdown. The applications running on the instances store data in instance memory (RAM) that must be present when the instances resume operation.

Which approach should the solutions architect recommend to shut down and resume the instances?

A) Modify the application to store the data on instance store volumes. Reattach the volumes while restarting them.
B) Snapshot the instances before stopping them. Restore the snapshot after restarting the instances.
C) Run the applications on instances enabled for hibernation. Hibernate the instances before the shutdown.
D) Note the Availability Zone for each instance before stopping it. Restart the instances in the same Availability Zones after the shutdown.

6) A company plans to run a monitoring application on an Amazon EC2 instance in a VPC. Connections are made to the instance using its private IPv4 address. A solutions architect needs to design a solution that will allow traffic to be quickly directed to a standby instance if the application fails and becomes unreachable.

Which approach will meet these requirements?

A) Deploy an Application Load Balancer configured with a listener for the private IP address and register the primary instance with the load balancer. Upon failure, de-register the instance and register the secondary instance.
B) Configure a custom DHCP option set. Configure DHCP to assign the same private IP address to the secondary instance when the primary instance fails.
C) Attach a secondary elastic network interface (ENI) to the instance configured with the private IP address. Move the ENI to the standby instance if the primary instance becomes unreachable.
D) Associate an Elastic IP address with the network interface of the primary instance. Disassociate the Elastic IP from the primary instance upon failure and associate it with a secondary instance.

7) An analytics company is planning to offer a site analytics service to its users. The service will require that the users' webpages include a JavaScript script that makes authenticated GET requests to the company's Amazon S3 bucket.

What must a solutions architect do to ensure that the script will successfully execute?

A) Enable cross-origin resource sharing (CORS) on the S3 bucket.
B) Enable S3 versioning on the S3 bucket.
C) Provide the users with a signed URL for the script.
D) Configure a bucket policy to allow public execute privileges.

8) A company's security team requires that all data stored in the cloud be encrypted at rest at all times using encryption keys stored on-premises.

Which encryption options meet these requirements? (Select TWO.)

A) Use Server-Side Encryption with Amazon S3 Managed Keys (SSE-S3).
B) Use Server-Side Encryption with AWS KMS Managed Keys (SSE-KMS).
C) Use Server-Side Encryption with Customer Provided Keys (SSE-C).
D) Use client-side encryption to provide at-rest encryption.
E) Use an AWS Lambda function triggered by Amazon S3 events to encrypt the data using the customer's keys.

9) A company needs to maintain access logs for a minimum of 5 years due to regulatory requirements. The data is rarely accessed once stored, but must be accessible with one day's notice if it is needed.

What is the MOST cost-effective data storage solution that meets these requirements?

A) Store the data in Amazon S3 Glacier Deep Archive storage and delete the objects after 5 years using a lifecycle rule.
B) Store the data in Amazon S3 Standard storage and transition to Amazon S3 Glacier after 30 days using a lifecycle rule.
C) Store the data in logs using Amazon CloudWatch Logs and set the retention period to 5 years.
D) Store the data in Amazon S3 Standard-Infrequent Access (S3 Standard-IA) storage and delete the objects after 5 years using a lifecycle rule.

10) A company uses Reserved Instances to run its data-processing workload. The nightly job typically takes 7 hours to run and must finish within a 10-hour time window. The company anticipates temporary increases in demand at the end of each month that will cause the job to run over the time limit with the capacity of the current resources. Once started, the processing job cannot be interrupted before completion. The company wants to implement a solution that would allow it to provide increased capacity as cost-effectively as possible.

What should a solutions architect do to accomplish this?

A) Deploy On-Demand Instances during periods of high demand.
B) Create a second Amazon EC2 reservation for additional instances.
C) Deploy Spot Instances during periods of high demand.
D) Increase the instance size of the instances in the Amazon EC2 reservation to support the increased workload.

Answers

1) A – An Application Load Balancer (ALB) sends requests to healthy instances only. An ALB performs periodic health checks on targets in a target group. An instance that fails health checks for a configurable number of consecutive times is considered unhealthy. The load balancer will no longer send requests to the instance until it passes another health check.

2) D – When the ReceiveMessageWaitTimeSeconds property of a queue is set to a value greater than zero, long polling is in effect. Long polling reduces the number of empty responses by allowing Amazon SQS to wait until a message is available before sending a response to a ReceiveMessage request.

3) B – An AWS Storage Gateway volume gateway connects an on-premises software application with cloud-backed storage volumes that can be mounted as Internet Small Computer System Interface (iSCSI) devices from on-premises application servers. In cached volumes mode, all the data is stored in Amazon S3 and a copy of frequently accessed data is stored locally.

4) A, B – A NAT gateway forwards traffic from the instances in the private subnet to the internet or other AWS services, and then sends the response back to the instances. After a NAT gateway is created, the route tables for private subnets must be updated to point internet traffic to the NAT gateway.

5) C – Hibernating an instance saves the contents of RAM to the Amazon EBS root volume. When the instance restarts, the RAM contents are reloaded.

6) C – A secondary ENI can be added to an instance. While primary ENIs cannot be detached from an instance, secondary ENIs can be detached and attached to a different instance.

7) A – Web browsers will block the execution of a script that originates from a server with a different domain name than the webpage. Amazon S3 can be configured with CORS to send HTTP headers that allow the script execution.

8) C, D – Server-Side Encryption with Customer-Provided Keys (SSE-C) enables Amazon S3 to encrypt objects server side using an encryption key provided in the PUT request. The same key must be provided in GET requests for Amazon S3 to decrypt the object. Customers also have the option to encrypt data client side before uploading it to Amazon S3 and decrypting it after downloading it. AWS SDKs provide an S3 encryption client that streamlines the process.

9) A – Data can be stored directly in Amazon S3 Glacier Deep Archive. This is the cheapest S3 storage class.

10) A – While Spot Instances would be the least costly option, they are not suitable for jobs that cannot be interrupted or must complete within a certain time period. On-Demand Instances would be billed for the number of seconds they are running.