# SAA-C03 Q&As

AWS Certified Solutions Architect - Associate (SAA-C03)

## Pass Amazon SAA-C03 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

**https://www.passapply.com/saa-c03.html**

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon Official Exam Center

- **Instant Download** After Purchase
- **100% Money Back** Guarantee
- **365 Days** Free Update
- **800,000+** Satisfied Customers

**QUESTION 1**

A solutions architect is developing a multiple-subnet VPC architecture. The solution will consist of six subnets in two Availability Zones. The subnets are defined as public, private and dedicated for databases. Only the Amazon EC2 instances running in the private subnets should be able to access a database.

Which solution meets these requirements?

A. Create a now route table that excludes the route to the public subnets\\' CIDR blocks. Associate the route table to the database subnets.

B. Create a security group that denies ingress from the security group used by instances in the public subnets. Attach the security group to an Amazon RDS DB instance.

C. Create a security group that allows ingress from the security group used by instances in the private subnets. Attach the security group to an Amazon RDS DB instance.

D. Create a new peering connection between the public subnets and the private subnets. Create a different peering connection between the private subnets and the database subnets.

Correct Answer: C

Security groups are stateful. All inbound traffic is blocked by default. If you create an inbound rule allowing traffic in, that traffic is automatically allowed back out again. You cannot block specific IP address using Security groups (instead use Network Access Control Lists).

"You can specify allow rules, but not deny rules." "When you first create a security group, it has no inbound rules. Therefore, no inbound traffic originating from another host to your instance is allowed until you add inbound rules to the security group." Source: https://docs.aws.amazon.com/vpc/latest/userguide/VPC_SecurityGroups.html#VPCSecurityGroups

**QUESTION 2**

An application runs on Amazon EC2 instances across multiple Availability Zones The instances run in an Amazon EC2 Auto Scaling group behind an Application Load Balancer The application performs best when the CPU utilization of the EC2 instances is at or near 40%.

What should a solutions architect do to maintain the desired performance across all instances in the group?

A. Use a simple scaling policy to dynamically scale the Auto Scaling group

B. Use a target tracking policy to dynamically scale the Auto Scaling group

C. Use an AWS Lambda function to update the desired Auto Scaling group capacity.

D. Use scheduled scaling actions to scale up and scale down the Auto Scaling group

Correct Answer: B

https://docs.aws.amazon.com/autoscaling/application/userguide/application-auto-scaling-target-tracking.html

**QUESTION 3**

A company has data collection sensors at different locations. The data collection sensors stream a high volume of data to the company. The company wants to design a platform on AWS to ingest and process high-volume streaming data. The solution must be scalable and support data collection in near real time. The company must store the data in Amazon S3 for future reporting.

Which solution will meet these requirements with the LEAST operational overhead?

A. Use Amazon Kinesis Data Firehose to deliver streaming data to Amazon S3.

B. Use AWS Glue to deliver streaming data to Amazon S3.

C. Use AWS Lambda to deliver streaming data and store the data to Amazon S3.

D. Use AWS Database Migration Service (AWS DMS) to deliver streaming data to Amazon S3.

Correct Answer: A

Data collection in near real time = Amazon Kinesis Data Firehose

**QUESTION 4**

A company has a mulli-tier application that runs six front-end web servers in an Amazon EC2 Auto Scaling group in a single Availability Zone behind an Application Load Balancer (ALB). A solutions architect needs lo modify the infrastructure to be highly available without modifying the application.

Which architecture should the solutions architect choose that provides high availability?

A. Create an Auto Scaling group that uses three Instances across each of tv/o Regions.

B. Modify the Auto Scaling group to use three instances across each of two Availability Zones.

C. Create an Auto Scaling template that can be used to quickly create more instances in another Region.

D. Change the ALB in front of the Amazon EC2 instances in a round-robin configuration to balance traffic to the web tier.

Correct Answer: B

High availability can be enabled for this architecture quite simply by modifying the existing Auto Scaling group to use multiple availability zones. The ASG will automatically balance the load so you don\\'t actually need to specify the instances per AZ.

**QUESTION 5**

A company is using a fleet of Amazon EC2 instances to ingest data from on-premises data sources. The data is in JSON format and Ingestion rates can be as high as 1 MB/s. When an EC2 instance is rebooted, the data in-flight is lost. The company\\'s data science team wants to query Ingested data In near-real time.

Which solution provides near-real-time data querying that is scalable with minimal data loss?

A. Publish data to Amazon Kinesis Data Streams Use Kinesis data Analytics to query the data.

B. Publish data to Amazon Kinesis Data Firehose with Amazon Redshift as the destination Use Amazon Redshift to query the data

C. Store ingested data m an EC2 Instance store Publish data to Amazon Kinesis Data Firehose with Amazon S3 as the destination. Use Amazon Athena to query the data.

D. Store ingested data m an Amazon Elastic Block Store (Amazon EBS) volume Publish data to Amazon ElastiCache tor Red Subscribe to the Redis channel to query the data

Correct Answer: A

Publishing data to Amazon Kinesis Data Streams can support ingestion rates as high as 1 MB/s and provide real-time data processing. Kinesis Data Analytics can query the ingested data in real-time with low latency, and the solution can scale as needed to accommodate increases in ingestion rates or querying needs. This solution also ensures minimal data loss in the event of an EC2 instance reboot since Kinesis Data Streams has a persistent data store for up to 7 days by default.

Latest SAA-C03 Dumps SAA-C03 Study Guide SAA-C03 Exam Questions