Demo Questions

Amazon SAA-C02 Exam

AWS Certified Solutions Architect - Associate SAA-C02

Thank you for downloading SAA-C02 Exam PDF

Question #1 *Topic 1*

A solutions architect is designing a solution where users will be directed to a backup static error page if the primary website is unavailable. The primary website's DNS records are hosted in Amazon Route 53 where their domain is pointing to an Application Load Balancer (ALB).
Which configuration should the solutions architect use to meet the company's needs while minimizing changes and infrastructure overhead?

- A. Point a Route 53 alias record to an Amazon CloudFront distribution with the ALB as one of its origins. Then, create custom error pages for the distribution.
- B. Set up a Route 53 active-passive failover configuration. Direct traffic to a static error page hosted within an Amazon S3 bucket when Route 53 health checks determine that the ALB endpoint is unhealthy.
- C. Update the Route 53 record to use a latency-based routing policy. Add the backup static error page hosted within an Amazon S3 bucket to the record so the traffic is sent to the most responsive endpoints.
- D. Set up a Route 53 active-active configuration with the ALB and an Amazon EC2 instance hosting a static error page as endpoints. Route 53 will only send requests to the instance if the heal checks fail for the ALB.

Correct Answer: *B*

Question #2 *Topic 1*

A solutions architect is designing a high performance computing (HPC) workload on Amazon EC2. The EC2 instances need to communicate to each other frequently and require network performance with low latency and high throughput.
Which EC2 configuration meets these requirements?

- A. Launch the EC2 instances in a cluster placement group in one Availability Zone.
- B. Launch the EC2 instances in a spread placement group in one Availability Zone.
- C. Launch the EC2 instances in an Auto Scaling group in two Regions and peer the VPCs.
- D. Launch the EC2 instances in an Auto Scaling group spanning multiple Availability Zones.

Correct Answer: *A*

Question #3 *Topic 1*

A company wants to host a scalable web application on AWS. The application will be accessed by users from different geographic regions of the world.
Application users will be able to download and upload unique data up to gigabytes in size. The development team wants a cost-effective solution to minimize upload and download latency and maximize performance.
What should a solutions architect do to accomplish this?

- A. Use Amazon S3 with Transfer Acceleration to host the application.
- B. Use Amazon S3 with CacheControl headers to host the application.
- C. Use Amazon EC2 with Auto Scaling and Amazon CloudFront to host the application.
- D. Use Amazon EC2 with Auto Scaling and Amazon ElastiCache to host the application.

Correct Answer: *C*
Reference:
https://aws.amazon.com/ec2/autoscaling/

Question #4 *Topic 1*

A company is migrating from an on-premises infrastructure to the AWS Cloud. One of the company's applications stores files on a Windows file server farm that uses Distributed File System Replication (DFSR) to keep data in sync. A solutions architect needs to replace the file server farm.
Which service should the solutions architect use?

- A. Amazon EFS

- B. Amazon FSx
- C. Amazon S3
- D. AWS Storage Gateway

Correct Answer: *B*
Reference:
https://docs.aws.amazon.com/fsx/latest/WindowsGuide/migrate-files-to-fsx-datasync.html

Question #5 *Topic 1*

A company has a legacy application that process data in two parts. The second part of the process takes longer than the first, so the company has decided to rewrite the application as two microservices running on Amazon ECS that can scale independently. How should a solutions architect integrate the microservices?

- A. Implement code in microservice 1 to send data to an Amazon S3 bucket. Use S3 event notifications to invoke microservice 2.
- B. Implement code in microservice 1 to publish data to an Amazon SNS topic. Implement code in microservice 2 to subscribe to this topic.
- C. Implement code in microservice 1 to send data to Amazon Kinesis Data Firehose. Implement code in microservice 2 to read from Kinesis Data Firehose.
- D. Implement code in microservice 1 to send data to an Amazon SQS queue. Implement code in microservice 2 to process messages from the queue.

Correct Answer: *C*

Question #6 *Topic 1*

A company captures clickstream data from multiple websites and analyzes it using batch processing. The data is loaded nightly into Amazon Redshift and is consumed by business analysts. The company wants to move towards near-real-time data processing for timely insights. The solution should process the streaming data with minimal effort and operational overhead.
Which combination of AWS services are MOST cost-effective for this solution? (Choose two.)

- A. Amazon EC2
- B. AWS Lambda
- C. Amazon Kinesis Data Streams
- D. Amazon Kinesis Data Firehose

- E. Amazon Kinesis Data Analytics

Correct Answer: *AD*
Reference:
https://d0.awsstatic.com/whitepapers/whitepaper-streaming-data-solutions-on-aws-with-amazon-kinesis.pdf
(9)

Question #7 *Topic 1*

A company's application runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. On the first day of every month at midnight, the application becomes much slower when the month-end financial calculation batch executes. This causes the CPU utilization of the EC2 instances to immediately peak to 100%, which disrupts the application.
What should a solutions architect recommend to ensure the application is able to handle the workload and avoid downtime?

- A. Configure an Amazon CloudFront distribution in front of the ALB.
- B. Configure an EC2 Auto Scaling simple scaling policy based on CPU utilization.
- C. Configure an EC2 Auto Scaling scheduled scaling policy based on the monthly schedule.
- D. Configure Amazon ElastiGache to remove some of the workload from the EC2 instances.

Correct Answer: *C*

Question #8 *Topic 1*

A company runs a multi-tier web application that hosts news content. The application runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an EC2 Auto Scaling group across multiple Availability Zones and use an Amazon Aurora database. A solutions architect needs to make the application more resilient to periodic increases in request rates.
Which architecture should the solutions architect implement? (Choose two.)

- A. Add AWS Shield.
- B. Add Aurora Replica.
- C. Add AWS Direct Connect.
- D. Add AWS Global Accelerator.
- E. Add an Amazon CloudFront distribution in front of the Application Load Balancer.

Correct Answer: *DE*

Question #9 *Topic 1*

An application running on AWS uses an Amazon Aurora Multi-AZ deployment for its database. When evaluating performance metrics, a solutions architect discovered that the database reads are causing high I/O and adding latency to the write requests against the database.

What should the solutions architect do to separate the read requests from the write requests?

- A. Enable read-through caching on the Amazon Aurora database.
- B. Update the application to read from the Multi-AZ standby instance.
- C. Create a read replica and modify the application to use the appropriate endpoint.
- D. Create a second Amazon Aurora database and link it to the primary database as a read replica.

Correct Answer: *C*
Reference:
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_ReadRepl.html

Question #10 *Topic 1*

A recently acquired company is required to build its own infrastructure on AWS and migrate multiple applications to the cloud within a month. Each application has approximately 50 TB of data to be transferred. After the migration is complete, this company and its parent company will both require secure network connectivity with consistent throughput from their data centers to the applications. A solution architect must ensure one-time data migration and ongoing network connectivity.

Which solution will meet these requirements?

- A. AWS Direct Connect for both the initial transfer and ongoing connectivity.
- B. AWS Site-to-Site VPN for both the initial transfer and ongoing connectivity.
- C. AWS Snowball for the initial transfer and AWS Direct Connect for ongoing connectivity.
- D. AWS Snowball for the initial transfer and AWS Site-to-Site VPN for ongoing connectivity.

Correct Answer: *D*
Reference:
https://docs.aws.amazon.com/dms/latest/userguide/CHAP_LargeDBs.html