- **Vendor:** Amazon
- **Exam Code:** SAA-C02
- **Exam Name:** AWS Certified Solutions Architect - Associate (SAA-C02) Exam
- **New Updated Questions from Braindump2go (Updated in May/2020)**

QUESTION 23
A Solutions Architect is creating an application running in an Amazon VPC that needs to access AWS Systems Manager Parameter Store. Network security rules prohibit any route table entry with a 0.0.0.0/0 destination.
What infrastructure addition will allow access to the AWS service while meeting the requirements?

A. VPC peering
B. NAT instance
C. NAT gateway
D. AWS PrivateLink

Answer: D
Explanation:
You can privately access AWS Systems Manager APIs from your VPC (created using Amazon Virtual Private Cloud) by creating VPC Endpoints. With VPC Endpoints, the routing between the VPC and AWS Systems Manager is handled by the AWS network without the need for an internet gateway, NAT gateway, or VPN connection. The latest generation of VPC Endpoints used by AWS Systems Manager are powered by AWS PrivateLink, a technology that enables private connectivity between AWS services using Elastic Network Interfaces (ENIs) with private IP addresses in your VPCs. To learn more about PrivateLink, visit the PrivateLink documentation.
https://docs.aws.amazon.com/vpc/latest/userguide/vpce-interface.html

QUESTION 24
A photo-sharing website running on AWS allows users to generate thumbnail images of photos stored in Amazon S3. An Amazon DynamoDB table maintains the locations of photos, and thumbnails are easily re- created from the originals if they are accidentally deleted.
How should the thumbnail images be stored to ensure the LOWEST cost?

A. Amazon S3 Standard-Infrequent Access (S3 Standard-IA) with cross-region replication
B. Amazon S3
C. Amazon Glacier
D. Amazon S3 with cross-region replication

Answer: B

QUESTION 25
A company is implementing a data lake solution on Amazon S3. Its security policy mandates that the data stored in Amazon S3 should be encrypted at rest.
Which options can achieve this? (Select TWO.)

A. Use S3 server-side encryption with an Amazon EC2 key pair.
B. Use S3 server-side encryption with customer-provided keys (SSE-C).
C. Use S3 bucket policies to restrict access to the data at rest.

D. Use client-side encryption before ingesting the data to Amazon S3 using encryption keys.
E. Use SSL to encrypt the data while in transit to Amazon S3.

Answer: BD

QUESTION 26
A solutions architect has created a new AWS account and must secure AWS account root user access.
Which combination of actions will accomplish this? (Select TWO.)

A. Ensure the root user uses a strong password
B. Enable multi-factor authentication to the root user
C. Store root user access keys in an encrypted Amazon S3 bucket
D. Add the root user to a group containing administrative permissions.
E. Apply the required permissions to the root user with an inline policy document

Answer: BC
Explanation:
https://docs.aws.amazon.com/IAM/latest/UserGuide/id_root-user.html

QUESTION 27
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 ElastiCache to remove some of the workload from the EC2 instances

Answer: C

QUESTION 28
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

Answer: A

QUESTION 29
A company's website is used to sell products to the public.
The site runs on Amazon EC2 instances in an Auto Scaling group behind an Application Load Balancer (ALB).
There is also an Amazon CloudFront distribution and AWS WAF is being used to protect against SQL injection attacks.
The ALB is the origin for the CloudFront distribution.
A recent review of security logs revealed an external malicious IP that needs to be blocked from accessing the website.
What should a solutions architect do to protect the application?

A. Modify the network ACL on the CloudFront distribution to add a deny rule for the malicious IP address
B. Modify the configuration of AWS WAF to add an IP match condition to block the malicious IP address
C. Modify the network ACL for the EC2 instances in the target groups behind the ALB to deny the malicious IP address
D. Modify the security groups for the EC2 instances in the target groups behind the ALB to deny the malicious IP address

Answer: D

QUESTION 30
A marketing company is storing CSV files in an Amazon S3 bucket for statistical analysis.
An application on an Amazon EC2 instance needs permission to efficiently process the CSV data stored in the S3 bucket.
Which action will MOST securely grant the EC2 instance access to the S3 bucket?

A. Attach a resource-based policy to the S3 bucket
B. Create an 1AM user for the application with specific permissions to the S3 bucket
C. Associate an 1AM role with least privilege permissions to the EC2 instance profile
D. Store AWS credentials directly on the EC2 instance for applications on the instance to use for API calls

Answer: C

QUESTION 31
A solutions architect is designing a solution where users will De directed to a backup static error page it 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 health checks fail for the ALB.

Answer: B

QUESTION 32
A solutions architect is designing the cloud architecture for a new application being deployed on AWS.
The process should run in parallel while adding and removing application nodes as needed based on the number of jobs to be processed.
The processor application is stateless.
The solutions architect must ensure that the application is loosely coupled and the job items are durably stored.
Which design should the solutions architect use?

A. Create an Amazon SNS topic to send the jobs that need to be processed.

Create an Amazon Machine Image (AMI) that consists of the processor application.
Create a launch configuration that uses the AMI.
Create an Auto Scaling group using the launch configuration.
Set the scaling policy for the Auto Scaling group to add and remove nodes based on CPU usage

B. Create an Amazon SQS queue to hold the jobs that need to be processed.
Create an Amazon Machine Image (AMI) that consists of the processor application.
Create a launch configuration that uses the AMI.
Create an Auto Scaling group using the launch configuration.
Set the scaling policy for the Auto Scaling group to add and remove nodes based on network usage

C. Create an Amazon SQS queue to hold the jobs that needs to be processed.
Create an Amazon Machine Image (AMI) that consists of the processor application.
Create a launch template that uses the AMI.
Create an Auto Scaling group using the launch template.
Set the scaling policy for the Auto Scaling group to add and remove nodes based on the number of items in the SQS queue

D. Create an Amazon SNS topic to send the jobs that need to be processed.
Create an Amazon Machine Image (AMI) that consists of the processor application.
Create a launch template that uses the AMI.
Create an Auto Scaling group using the launch template.
Set the scaling policy for the Auto Scaling group to add and remove nodes based on the number of messages published to the SNS topic.

Answer: D

QUESTION 33
A company has a legacy application that processes 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.

Answer: A

QUESTION 34
A solutions architect at an ecommerce company wants to back up application log data to Amazon S3.
The solutions architect is unsure how frequently the logs will be accessed or which logs will be accessed the most.
The company wants to keep costs as low as possible by using the appropriate S3 storage class.
Which S3 storage class should be implemented to meet these requirements?

A. S3 Glacier
B. S3 Intelligent-Tiering
C. S3 Standard-Infrequent Access (S3 Standard-IA)
D. S3 One Zone-Infrequent Access (S3 One Zone-IA)

Answer: B