

➤ **Vendor: Amazon**

➤ **Exam Code: SAA-C02**

➤ **Exam Name: AWS Certified Solutions Architect - Associate (SAA-C02) Exam**

➤ **New Updated Questions from [Braindump2go](#) (Updated in [May/2022](#))**

Visit Braindump2go and Download Full Version SAA-C02 Exam Dumps

QUESTION 948

A company is building an ecommerce application and needs to store sensitive customer information. The company needs to give customers the ability to complete purchase transactions on the website. The company also needs to ensure that sensitive customer data is protected, even from database administrators.

Which solution meets these requirements?

- A. Store sensitive data in an Amazon Elastic Block Store (Amazon EBS) volume.
Use EBS encryption to encrypt the data.
Use an IAM instance role to restrict access.
- B. Store sensitive data in Amazon RDS for MySQL.
Use AWS Key Management Service (AWS KMS) client-side encryption to encrypt the data.
- C. Store sensitive data in Amazon S3.
Use AWS Key Management Service (AWS KMS) service-side encryption the data.
Use S3 bucket policies to restrict access.
- D. Store sensitive data in Amazon FSx for Windows Server.
Mount the file share on application servers.
Use Windows file permissions to restrict access.

Answer: C

QUESTION 949

A company is launching a new application and will display application metrics on an Amazon CloudWatch dashboard. The company's product manager needs to access this dashboard periodically. The product manager does not have an AWS account. A solution architect must provide access to the product manager by following the principle of least privilege.

Which solution will meet these requirements?

- A. Share the dashboard from the CloudWatch console.
Enter the product manager's email address, and complete the sharing steps.
Provide a shareable link for the dashboard to the product manager.
- B. Create an IAM user specifically for the product manager.
Attach the CloudWatch Read Only Access managed policy to the user.
Share the new login credential with the product manager.
Share the browser URL of the correct dashboard with the product manager.
- C. Create an IAM user for the company's employees.
Attach the View Only Access AWS managed policy to the IAM user.
Share the new login credentials with the product manager.
Ask the product manager to navigate to the CloudWatch console and locate the dashboard by name in the Dashboards section.
- D. Deploy a bastion server in a public subnet. When the product manager requires access to the

[SAA-C02 Exam Dumps](#) [SAA-C02 Exam Questions](#) [SAA-C02 PDF Dumps](#) [SAA-C02 VCE Dumps](#)

<https://www.braindump2go.com/saa-c02.html>

dashboard, start the server and share the RDP credentials. On the bastion server, ensure that the browser is configured to open the dashboard URL with cached AWS credentials that have appropriate permissions to view the dashboard.

Answer: A

QUESTION 950

A company runs a latency-sensitive gaming service in the AWS Cloud. The gaming service runs on a fleet of Amazon EC2 instances behind an Application Load Balancer (ALB). An Amazon DynamoDB table stores the gaming data. All the infrastructure is in a single AWS Region. The main user base is in that same Region.

A solutions architect needs to update the architect to support a global expansion of the gaming service must operate with the least possible latency.

Which solution will meet these requirements?

- A. Create an Amazon CloudFront distribution in front of the ALB.
- B. Deploy an Amazon API Gateway regional API endpoint. Integrate the API endpoint with the ALB.
- C. Create an accelerator in AWS Global Accelerator. Add a listener. Configure the endpoint to point to the ALB.
- D. Deploy the ALB and the fleet of EC2 instances to another Region. Use Amazon Route 53 geolocation routing.

Answer: C

QUESTION 951

A company wants to migrate a Windows-based application from on premises to the AWS Cloud. The application has three tiers, a business tier, and a database tier with Microsoft SQL Server. The company wants to use specific features of SQL Server such as native backups and Data Quality Services. The company also needs to share files for process between the tiers.

How should a solution architect design the architecture to meet these requirements?

- A. Host all three on Amazon instances. Use Amazon FSx File Gateway for file sharing between tiers.
- B. Host all three on Amazon EC2 instances. Use Amazon FSx for Windows file sharing between the tiers.
- C. Host the application tier and the business tier on Amazon EC2 instances. Host the database tier on Amazon RDS. Use Amazon Elastic File system (Amazon EFS) for file sharing between the tiers.
- D. Host the application tier and the business tier on Amazon EC2 instances. Host the database tier on Amazon RDS. Use a Provisioned IOPS SSD (io2) Amazon Elastic Block Store (Amazon EBS) volume for file sharing between the tiers.

Answer: B

QUESTION 952

A solution architect is using an AWS CloudFormation template to deploy a three-tier web application. The web application consists of a web tier and an application that stores and retrieves user data in Amazon DynamoDB tables. The web and application tiers are hosted on Amazon EC2 instances, and the database tier is not publicly accessible. The application EC2 instances need to access the Dynamo tables without exposing API credentials in the template. What should the solution architect do to meet the requirements?

- A. Create an IAM role to read the DynamoDB tables. Associate the role with the application instances by referencing an instance profile.
- B. Create an IAM role that has the required permissions to read and write from the DynamoDB tables. Add the role to the EC2 instance profile, and associate the instance profile with the application instances.

- C. Use the parameter section in the AWS CloudFormation template to have the user input access and secret keys from an already-created IAM user that has the required permissions to read and write from the DynamoDB tables.
- D. Create an IAM user in the AWS CloudFormation template that has the required permissions to read and write from the DynamoDB tables. Use the GetAtt function to retrieve the access secret keys, and pass them to the application instances through the user data.

Answer: B

QUESTION 953

A company that recently started using AWS establishes a Site-to-Site VPN between its on-premises data center and AWS. The company's security mandate states that traffic originating from on premises should stay within the company's private IP space when communicating with an Amazon Elastic Container Service (Amazon ECS) cluster that is hosting a sample web application.

Which solution meets this requirement?

- A. Configure a gateway endpoint for Amazon ECS. Modify the route table to include an entry pointing to the ECS cluster.
- B. Create a Network Load Balancer and AWS PrivateLink endpoint for Amazon ECS in the same VPC that is hosting the ECS cluster.
- C. Create a Network Load Balancer in one VPC and an AWS PrivateLink endpoint for Amazon ECS in another VPC. Connect the two by using VPC peering.
- D. Configure an Amazon Route record with Amazon ECS as the target. Apply a server certificate to Route 53 from AWS Certificate Manager (ACM) for SSL offloading.

Answer: A

QUESTION 954

A company needs to migrate a legacy application from an on-premises data center to the AWS Cloud because of hardware capacity constraints. The application runs 24 hours a day. & days a week,. The application database storage continues to grow over time.

What should a solution architect do to meet these requirements MOST cost-effectively?

- A. Migrate the application layer to Amazon EC2 Spot Instances Migrate the data storage layer to Amazon S3.
- B. Migrate the application layer to Amazon EC2 Reserved Instances Migrate the data storage layer to Amazon RDS On-Demand Instances.
- C. Migrate the application layer to Amazon EC2 Reserved instances Migrate the data storage layer to Amazon Aurora Reserved Instances.
- D. Migrate the application layer to Amazon EC2 On Demand Amazon Migrate the data storage layer to Amazon RDS Reserved instances.

Answer: C

QUESTION 955

A company needs to retain application logs files for a critical application for 10 years. The application team regularly accesses logs from the past month for troubleshooting, but logs older than 1 month are rarely accessed. The application generates more than 10 TB of logs per month.

Which storage option meets these requirements MOST cost-effectively?

- A. Store the logs in Amazon S3. Use AWS Backup to move logs more than 1 month old to S3 Glacier Deep Archive.
- B. Store the logs in Amazon S3. Use S3 Lifecycle policies to move logs more than 1 month old to S3 Glacier Deep Archive.
- C. Store the logs in Amazon CloudWatch Logs. Use AWS Backup to move logs more than 1 month

old to S3 Glacier Deep Archive.

- D. Store the logs in Amazon CloudWatch Logs.
Use Amazon S3 Lifecycle policies to move logs more than 1 month old to S3 Glacier Deep Archive.

Answer: B

QUESTION 956

A company has five organizational units (OUS) as part of its organization in AWS Organization. Each OU correlate to the five business that the company owns. The company research and development R&D business is separating from the company and will need its own organization.

A solutions architect creates a separate new management account for this purpose.

What should a solution architect recommend to meet these requirements?

- A. Have the R&D AWS account be part of both organizations during the transition.
- B. Invite the R&D AWS account to be part of the new organization after the R&D AWS account has left the prior organization.
- C. Create a new R&D AWS account in the new organization.
Migrate resources from the period R&D AWS account to thee new R&D AWS account
- D. Have the R&D AWS account into the now organisation.
Make the now management account a member of the prior organisation

Answer: B

QUESTION 957

A company is hosting a website from an Amazon S3 bucket that is configured for public hosting. The company's security team mandates the usage of secure connections for access to the website. However; HTTP-based URLS and HTTPS-based URLS mist be functional.

What should a solution architect recommend to meet these requirements?

- A. Create an S3 bucket policy to explicitly deny non-HTTPS traffic.
- B. Enable S3 Transfer Acceleration. Select the HTTPS Only bucket property.
- C. Place thee website behind an Elastic Load Balancer that is configured to redirect HTTP traffic to HTTPPS.
- D. Serve the website through an Amazon CloudFront distribution that is configured to redirect HTTP traffic to HTTPS.

Answer: D

QUESTION 958

A company maintains a searchable repository of items on its website. The data is stored in an Amazon RDS for MySQL database table that contains more than 10 million rows. The database has 2 TB of General Purpose SSD storage.

There are millions of updates against this data every day through the company's website.

The company has noticed that some insert operations are taking 10 seconds or longer.

The company has determined that the database storage performance is the problem.

Which solution addresses this performance issue?

- A. Change the storage type to Provisioned IOPS SSD
- B. Change the DB instance to a memory optimized instance class
- C. Change the DB instance to a burstable performance instance class
- D. Enable Multi-AZ RDS read replicas with MySQL native asynchronous replication.

Answer: A

QUESTION 959

A company deploys Amazon EC2 instances that run in a VPC. The EC2 instances load source data into Amazon S3

[SAA-C02 Exam Dumps](#) **[SAA-C02 Exam Questions](#)** **[SAA-C02 PDF Dumps](#)** **[SAA-C02 VCE Dumps](#)**

<https://www.braindump2go.com/saa-c02.html>

buckets so that the data can be processed in the future. According to compliance laws, the data must not be transmitted over the public internet. Servers in the company's on-premises data center will consume the output from an application that runs on the EC2 instances. Which solution will meet these requirements?

- A. Deploy an interface VPC endpoint for Amazon EC2. Create an AWS Site-to-Site VPN connection between the company and the VPC
- B. Deploy a gateway VPC endpoint for Amazon S3. Set up an AWS Direct Connect connection between the on-premises network and the VPC
- C. Set up an AWS Transit Gateway connection from the VPC to the S3 buckets. Create an AWS Site-to-Site VPN connection between the company and the VPC
- D. Set up proxy EC2 instances that have routes to NAT gateways. Configure the proxy EC2 instances to fetch S3 data and feed the application instances

Answer: B

QUESTION 960

A company runs an application on Amazon EC2 instances, that are part of an Auto Scaling group. Traffic to the application increases substantially during business hours. A solutions architect needs to implement an Auto Scaling policy that addresses user latency concerns during periods of high traffic. The company does not want to provision more compute than is necessary.

What should the solutions architect do to meet these requirements?

- A. Configure a predictive scaling policy with the appropriate scaling metric.
- B. Configure a dynamic target tracking scaling policy with the appropriate scaling metric
- C. Configure a scheduled scaling policy that launches additional EC2 instances during business hours
- D. Configure dynamic step or simple scaling policies with Amazon CloudWatch alarms to add and remove EC2 instances based on alarm status

Answer: C

QUESTION 961

A company has a business-critical application that runs on Amazon EC2 instances. The application stores data in an Amazon DynamoDB table. The company must be able to revert the table to any point within the last 24 hours.

Which solution meets these requirements with the LEAST operational overhead?

- A. Configure point-in-time recovery for the table
- B. Use AWS Backup for the table
- C. Use an AWS Lambda function to make an on-demand backup of the table every hour
- D. Turn on streams on the table to capture a log of all changes to the table in the last 24 hours. Store a copy of the stream in an Amazon S3 bucket

Answer: A

QUESTION 962

A web application must send order data to Amazon S3 to support near-time processing.

A solutions architect needs to create an architecture that is scalable and fault tolerant.

Which solutions meet these requirements? (Select TWO.)

- A. Write the order event to an Amazon DynamoDB table. Use Amazon DynamoDB Streams to invoke an AWS Lambda function that parses the payload and writes the data to Amazon S3.
- B. Write the order event to an Amazon Simple Queue Service (Amazon SQS) queue. Use the queue to invoke an AWS Lambda function that parses the payload and writes the data to Amazon S3.
- C. Write the order event to an Amazon Simple Queue Service (Amazon SQS) queue. Use an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function that
- D. Write the order event to an Amazon Simple Queue Service (Amazon SQS) queue. Use an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function that

parses the payload and writes the data to Amazon S3.

- E. Write the order event to an Amazon Simple Notification Service (Amazon SNS) topic. Use an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function that parses the payload and writes the data to Amazon S3.

Answer: AB

QUESTION 963

A company's web application consists of multiple Amazon EC2 instances that run behind an Application Load Balancer in a VPC. An Amazon ROS for MySQL DB instance contains the data. The company needs the ability to automatically detect and respond to suspicious or unexpected behaviour in its AWS environment the company already has added AWS WAF to its architecture.

What should a solutions architect do next to protect against threats?

- A. Use Amazon GuardDuty to perform threat detection.
Configure Amazon EventBridge (Amazon CloudWatch Events) to filter for GuardDuty findings and to invoke an AWS Lambda function to adjust the AWS WAF rules
- B. Use AWS Firewall Manager to perform threat detection.
Configure Amazon EventBridge (Amazon CloudWatch Events) to filter for Firewall Manager findings and to invoke an AWS Lambda function to adjust the AWS WAF web ACL
- C. Use Amazon Inspector to perform threat detection and to update the AWS WAF rules.
Create a VPC network ACL to limit access to the web application
- D. Use Amazon Macie to perform threat detection and to update the AWS WAF rules.
Create a VPC network ACL to limit access to the web application

Answer: A

QUESTION 964

A company needs to ingest and handle large amounts of streaming data that its application generates. The application runs on Amazon EC2 instances and sends data to Amazon Kinesis Data Streams, which is configured with default settings. Every other day the application consumes the data and writes the data to an Amazon S3 bucket for business intelligence (BI) processing the company observes that Amazon S3 is not receiving all the data that the application sends to Kinesis Data Streams.

What should a solutions architect do to resolve this issue?

- A. Update the Kinesis Data Streams default settings by modifying the data retention period.
- B. Update the application to use the Kinesis Producer Library (KPL) to send the data to Kinesis Data Streams.
- C. Update the number of Kinesis shards to handle the throughput of the data that is sent to Kinesis Data Streams.
- D. Turn on S3 Versioning within the S3 bucket to preserve every version of every object that is ingested in the S3 bucket.

Answer: A

QUESTION 965

A company has a web-based map application that provides status information about ongoing repairs. The application sometimes has millions of users. Repair teams have a mobile app that sends current location and status in a JSON message to a REST-based endpoint. Few repairs occur on most days. The company wants the application to be highly available and to scale when large numbers of repairs occur after natural disasters. Customers use the application most often during these times. The company does not want to pay for idle capacity.

What should a solution architect recommend to meet these requirements?

- A. Create a webpage that is based on Amazon S3 to display information.
Use Amazon API Gateway and AWS Lambda to receive the JSON status data. Store the JSON data in Amazon S3.

- B. Use Amazon EC2 instances as web servers across multiple Availability Zones.
Run the EC2 instances in an Auto Scaling group.
Use Amazon API Gateway and AWS Lambda to receive the JSON status data. Store the JSON data in Amazon S3.
- C. Use Amazon EC2 instances as web servers across multiple Availability Zones.
Run the EC2 instances in an Auto Scaling group.
Use a REST endpoint on the EC2 instances to receive the JSON status data.
Store the JSON data in an Amazon RDS Multi-AZ DB instance.
- D. Use Amazon EC2 instances as web servers across multiple Availability zones.
Run the EC2 instances in an Auto Scaling group.
Use a REST endpoint on the EC2 instances to receive the JSON status data.
Store the JSON data in an Amazon DynamoDB table.

Answer: D

QUESTION 966

A company is designing an application that will run on an AWS Lambda function within a VPC. Gateway API will invoke the Lambda function.

A solution architect needs to recommend an Amazon CloudWatch solution that developers can use to identify the users who are generating the most network traffic.

Which solution will meet these requirements?

- A. Configure CloudWatch Lambda Insights. Examine the network usage graph by using the multi-function view in the performance dashboard.
- B. Create a canary in CloudWatch Synthetics. Turn on active tracing. Review the network usage graph in the Monitoring tab of the canary.
- C. Configure VPC Flow Logs to stream to CloudWatch Logs.
Create a CloudWatch Contributor Insights rule from the sample blueprint.
- D. Add the application to CloudWatch Application Insights.
View the graph for top network users in the dashboard that Application Insights creates automatically.

Answer: C

QUESTION 967

A company recently released a new type of internet-connected sensor. The company is expecting to sell thousands of sensors, which are designed to stream high volumes of data each second to a central location. A solutions architect must design a solution that ingests and stores data so that engineering teams can analyze it in near-real time with millisecond responsiveness.

Which solution should the solution architect recommend?

- A. Use an Amazon SQS queue to ingest the data.
Consume the data with an AWS Lambda function which then stores the data in Amazon Redshift.
- B. Use an Amazon SQS queue to ingest the data.
Consume the data with an AWS Lambda function which then stores the data in Amazon DynamoDB.
- C. Use Amazon Kinesis Data Streams to ingest the data.
Consume the data with an AWS Lambda function, which then stores the data in Amazon Redshift.
- D. Use Amazon Kinesis Data Streams to ingest the data.
Consume the data with an AWS Lambda function, which then stores the data in Amazon DynamoDB.

Answer: C

QUESTION 968

A startup company is hosting a website for its customers on an Amazon EC2 instance. The website consists of a stateless Python application and a MySQL database. The website serves only a small amount of traffic. The company is

[SAA-C02 Exam Dumps](#) **[SAA-C02 Exam Questions](#)** **[SAA-C02 PDF Dumps](#)** **[SAA-C02 VCE Dumps](#)**

<https://www.braindump2go.com/saa-c02.html>

concerned about the reliability of the instance and needs to migrate to a highly available architecture. The company cannot modify the application code.

Which combination of actions should a solution architect take to achieve high availability for the website? (Select TWO.)

- A. Provision an internet gateway in each Availability Zone in use.
- B. Migrate the database to on Amazon RDS for MySQL Multi-AZ DB instance
- C. Migrate the database to Amazon DynamoDB, and enable DynamoDB auto scaling.
- D. Use AWS DataSync to synchronize the database data across multiple EC2 instances
- E. Create an Application Load Balancer to distribute traffic to an Auto Scaling group or EC2 instances that are distributed across two Availability Zones.

Answer: BE

QUESTION 969

A company hosts its web applications in the AWS Cloud. The company configures Elastic Load Balancers to use certificate that are imported into AWS Certificate Manager (ACM). The company's security team must be notified 30 days before the expiration of each certificate.

What should a solutions architect recommend to meet the requirement?

- A. Add a rule in ACM to publish a custom message to an Amazon Simple Notification Service (Amazon SNS) topic every day beginning 30 days before any certificate will expire.
- B. Create an AWS Config rule that checks for certificates that will expire within 30 days. Configure Amazon EventBridge (Amazon CloudWatch Events) to invoke a custom alert by way of Amazon Simple Notification Service (Amazon SNS) when AWS Config reports a noncompliant resource
- C. Use AWS Trusted Advisor to check for certificates that will expire within 30 days. Create an Amazon CloudWatch alarm that is based on Trusted Advisor metrics for check status changes. Configure the alarm to send a custom alert by way of Amazon Simple Notification Service (Amazon SNS)
- D. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to detect any certificates that will expire within 30 days. Configure the rule to invoke an AWS Lambda function. Configure the Lambda function to send a custom alert by way of Amazon Simple Notification Service (Amazon SNS).

Answer: A

QUESTION 970

A company wants to reduce the cost of its existing three-tier web architect. The web, application, and database servers are running on Amazon EC2 instance EC2 instance for the development, test and production environments. The EC2 instances average 30% CPU utilization during peak hours and 10% CPU utilization during non-peak hours. The production EC2 instance purchasing solution will meet the company's requirements MOST cost-effectively?

- A. Use Spot Instances for the production EC2 instances. Use Reserved Instances for the development and test EC2 instances
- B. Use Reserved Instances for the production EC2 instances. Use On-Demand Instances for the development and test EC2 instances
- C. Use blocks for the production EC2 instances. Use Reserved instances for the development and test EC2 instances
- D. Use On-Demand Instances for the production EC2 instances. Use Spot blocks for the development and test EC2 instances

Answer: B

QUESTION 971

A company is running a database on an Amazon RDS for MySQL DB instance. The company must maintain a near-real-time replica of the database on premises. The company needs to encrypt the data in transit and is using a 1 Gbps AWS Direct Connect connection.

Which solution will meet these requirements?

- A. Use AWS Data Pipeline to replicate from AWS to on premises over an IPsec VPN on top of the Direct Conned connection.
- B. Use MySQL replication to replicate from AWS to on premises over an IPsec VPN on top of the Direct Connect connection
- C. Use the RDS Multi-AZ feature. Choose on premises as the failover Availability Zone over an IPsec VPN on top of the Direct Connect connection.
- D. Use AWS Database Migration Service (AWS DMS) and Direct Connect with MACsec encryption to continuously replicate the data from AWS to on premises.

Answer: B

QUESTION 972

A company wants to run applications in container in the AWS Cloud. Those applications arc stateless and can tolerate disruptions.

What should a solutions architect do to meet those requirements?

- A. Use Spot Instances in an Amazon EC2 Auto Scaling group to run the application containers
- B. Use Spot Instances in an Amazon Elastic Kubernetes Service (Amazon EKS) managed node group
- C. Use On-Demand Instances in an Amazon EC2 Auto Scaling group to run the application containers
- D. Use On-Demand Instances in an Amazon Elastic Kubernetes Service (Amazon EKS) managed node group.

Answer: A

QUESTION 973

A company runs multiple Windows workloads on AWS. The company's employees use Windows the file shares that are hosted on two Amazon EC2 instances. The file shares synchronize data between themselves and maintain duplicate copies. The company wants a highly available and durable storage solution that preserves how users currently access the files.

What should a solutions architect do to meet those requirements?

- A. Migrate all the data to Amazon S3 Set up IAM authentication for users to access files
- B. Set up an Amazon S3 File Gateway. Mount the S3 File Gateway on the existing EC2 Instances.
- C. Extend the file share environment to Amazon FSx for Windows File Server with a Multi-AZ configuration.
Migrate all the data to FSx for Windows File Server.
- D. Extend the file share environment to Amazon Elastic File System (Amazon EFS) with a Multi-AZ configuration.
Migrate all the data to Amazon EFS.

Answer: C

QUESTION 974

A Company has an application that provides marketing services to stores. The services are based on previous purchased by store customers. The stores upload transaction data to the company through SFTP, and the data is processed an analysed to generate new marketing offers. Some of the files can exceed 200 GB in size. Recently, the company discovered that some of the stores have uploaded file that contains personality identifiable information (PII) that should not have included. The company wants administrators to be alerted if PII is shared again. The company also wants to automate remediation.

What should a solutions architect do to meet those requirements?

- A. Use an Amazon S3 bucket as a secure transfer point. Use Amazon Inspector to scan me objects

in the bucket. If objects contain PII, trigger an S3 Lifecycle policy to remove the objects that contain PII.

- B. Use an Amazon S3 bucket as a secure transfer point. Use Amazon Macie to scan the objects in the bucket. If objects contain PII, Use Amazon Simple Notification Service (Amazon SNS) to trigger a notification to the administrators to remove the objects that contain PII.
- C. Implement custom scanning algorithms in an AWS Lambda function. Trigger the function when objects are loaded into the bucket. If objects contain PII, Use Amazon Simple Notification Service (Amazon SNS) to trigger a notification to the administrators to remove the objects that contain PII.
- D. Implement custom scanning algorithms in an AWS Lambda function. Trigger the function when objects are loaded into the bucket. If objects contain PII, Use Amazon Simple Email Service (Amazon SES) to trigger a notification to the administrators and trigger an S3 Lifecycle policy to remove the objects that contain PII.

Answer: B

QUESTION 975

A company has a web application that runs on Amazon EC2 instances. The company wants end users to authenticate themselves before they use the web application. The web application accesses AWS resources, such as Amazon S3 buckets, on behalf of users who are logged on.

Which combination of actions must a solutions architect take to meet these requirements? (Select TWO).

- A. Configure AWS App Mesh to log on users.
- B. Enable and configure AWS Single Sign-On in AWS Identity and Access Management (IAM).
- C. Define a default IAM role for authenticated users.
- D. Use AWS Identity and Access Management (IAM) for user authentication.
- E. Use Amazon Cognito for user authentication.

Answer: BE

QUESTION 976

A company uses 50 TB of data for reporting. The company wants to move this data from on premises to AWS. A custom application in the company's data center runs a weekly data transformation job. The company plans to pause the application until the data transfer is complete and needs to begin the transfer process as soon as possible.

The data center does not have any available network bandwidth for additional workloads.

A solutions architect must transfer the data and must configure the transformation job to continue to run in the AWS Cloud.

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

- A. Use AWS DataSync to move the data. Create a custom transformation job by using AWS Glue.
- B. Order an AWS Snowcone device to move the data. Deploy the transformation application to the device.
- C. Order an AWS Snowball Edge Storage Optimized device. Copy the data to the device. Create a custom transformation job by using AWS Glue.
- D. Order an AWS Snowball Edge Storage Optimized device that includes Amazon EC2 compute. Copy the data to the device. Create a new EC2 instance on AWS to run the transformation application.

Answer: D

QUESTION 977

A company has an Amazon S3 data lake that is governed by AWS Lake Formation. The company wants to create a visualization in Amazon QuickSight by joining the data in the data lake with operational data that is stored in an Amazon Aurora MySQL database. The company wants to enforce column-level authorization so that the company's marketing team can access only a subset of columns in the database.

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

- A. Use Amazon EMR to ingest the data directly from the database to the QuickSight SPICE engine Include only the required columns
- B. Use AWS Glue Studio to ingest the data from the database to the S3 data lake Attach an IAM policy to the QuickSight users to enforce column-level access control. Use Amazon S3 as the data source in QuickSight
- C. Use AWS Glue Elastic Views to create a materialized view for the database in Amazon S3 Create an S3 bucket policy to enforce column-level access control for the QuickSight users. Use Amazon S3 as the data source in QuickSight.
- D. Use a Lake Formation blueprint to ingest the data from the database to the S3 data lake. Use Lake Formation to enforce column-level access control for the QuickSight users. Use Amazon Athena as the data source in QuickSight

Answer: C

QUESTION 978

A company's ecommerce website has unpredictable traffic and uses AWS Lambda functions to directly access a private Amazon RDS for PostgreSQL DB instance. The company wants to maintain predictable database performance and ensure that the Lambda invocations do not overload the database with too many connections. What should a solutions architect do to meet these requirements?

- A. Point the client driver at an RDS custom endpoint. Deploy the Lambda functions inside a VPC
- B. Point the client driver at an RDS proxy endpoint. Deploy the Lambda functions inside a VPC
- C. Point the client driver at an RDS custom endpoint. Deploy the Lambda functions outside a VPC
- D. Point the client driver at an RDS proxy endpoint. Deploy the Lambda functions outside a VPC

Answer: B

QUESTION 979

A company runs an application that receives data from thousands of geographically dispersed remote devices that use UDP The application processes the data immediately and sends a message back to the device if necessary No data is stored.

The company needs a solution that minimizes latency for the data transmission from the devices. The solution also must provide rapid failover to another AWS Region.

Which solution will meet these requirements?

- A. Configure an Amazon Route 53 failover routing policy. Create a Network Load Balancer (NLB) in each of the two Regions. Configure the NLB to invoke an AWS Lambda function to process the data
- B. Use AWS Global Accelerator. Create a Network Load Balancer (NLB) in each of the two Regions as an endpoint. Create an Amazon Elastic Container Service (Amazon ECS) cluster with the Fargate launch type. Create an ECS service on the cluster Set the ECS service as the target for the NLB Process the data in Amazon ECS.
- C. Use AWS Global Accelerator. Create an Application Load Balancer (ALB) in each of the two Regions as an endpoint. Create an Amazon Elastic Container Service (Amazon ECS) cluster with the Fargate launch type. Create an ECS service on the cluster. Set the ECS service as the target for the ALB Process the data in Amazon ECS.
- D. Configure an Amazon Route 53 failover routing policy. Create an Application Load Balancer (ALB) in each of the two Regions. Create an Amazon Elastic Container Service (Amazon ECS) cluster with the Fargate launch type. Create an ECS service on the cluster Set the ECS service as the target for the ALB Process the data in Amazon ECS.

Answer: C

QUESTION 980

The company wants the application to be highly available with minimum downtime and minimum loss of data. Which solution will meet these requirements with the LEAST operational effort?

- A. Place the EC2 instances in different AWS Regions.
Use Amazon Route 53 health checks to redirect traffic.
Use Aurora PostgreSQL Cross-Region Replication.
- B. Configure the Auto Scaling group to use multiple Availability Zones.
Configure the database as Multi-AZ.
Configure an Amazon RDS Proxy instance for the database.
- C. Configure the Auto Scaling group to use one Availability Zone.
Generate hourly snapshots of the database.
Recover the database from the snapshots in the event of a failure.
- D. Configure the Auto Scaling group to use multiple AWS Regions.
Write the data from the application to Amazon S3.
Use S3 Event Notifications to launch an AWS Lambda function to write the data to the database.

Answer: B

QUESTION 981

A company is developing a new machine learning (ML) model solution on AWS. The models are developed as independent microservices that fetch approximately 1GB of model data from Amazon S3 at startup and load the data into memory. Users access the models through an asynchronous API. Users can send a request or a batch of requests and specify where the results should be sent. The company provides models to hundreds of users. The usage patterns for the models are irregular. Some models could be unused for days or weeks. Other models could receive batches of thousands of requests at a time. Which design should a solutions architect recommend to meet these requirements?

- A. Direct the requests from the API to a Network Load Balancer (NLB).
Deploy the models as AWS Lambda functions that are invoked by the NLB.
- B. Direct the requests from the API to an Application Load Balancer (ALB).
Deploy the models as Amazon Elastic Container Service (Amazon ECS) services that read from an Amazon Simple Queue Service (Amazon SQS) queue.
Use AWS App Mesh to scale the instances of the ECS cluster based on the SQS queue size.
- C. Direct the requests from the API into an Amazon Simple Queue Service (Amazon SQS) queue.
Deploy the models as AWS Lambda functions that are invoked by SQS events.
Use AWS Auto Scaling to increase the number of vCPUs for the Lambda functions based on the SQS queue size.
- D. Direct the requests from the API into an Amazon Simple Queue Service (Amazon SQS) queue.
Deploy the models as Amazon Elastic Container Service (Amazon ECS) services that read from the queue.
Enable AWS Auto Scaling on Amazon ECS for both the cluster and copies of the service based on the queue size.

Answer: C

QUESTION 982

A company has an application with a REST-based interface that allows data to be received in near-real time from a third-party vendor. Once received the application processes and stores the data for further analysis. The application is running on Amazon EC2 instances. The third-party vendor has received many 503 Service Unavailable Errors when sending data to the application. When the data volume spikes, the compute capacity reaches its maximum limit and the application is unable to process all requests. Which design should a solutions architect recommend to provide a more scalable solution?

- A. Use Amazon Kinesis Data Streams to ingest the data.
Process the data using AWS Lambda function.
- B. Use Amazon API Gateway on top of the existing application.
Create a usage plan with a quota limit for the third-party vendor.
- C. Use Amazon Simple Notification Service (Amazon SNS) to ingest the data.
Put the EC2 instances in an Auto Scaling group behind an Application Load Balancer.

[SAA-C02 Exam Dumps](#) [SAA-C02 Exam Questions](#) [SAA-C02 PDF Dumps](#) [SAA-C02 VCE Dumps](#)

<https://www.braindump2go.com/saa-c02.html>

- D. Repackage the application as a container.
Deploy the application using Amazon Elastic Container Service (Amazon ECS) using the EC2 launch type with an Auto Scaling group.

Answer: A

QUESTION 983

A company stores millions of objects in Amazon S3. The data is in JSON format and Apache Parquet format. The data is partitioned and new objects are added daily. A solutions architect needs to create a solution so that employees can use SQL to perform one-time queries against all the data. The solution must avoid code changes and must minimize operational overhead.

Which solution will meet these requirements?

- A. Use S3 Select to perform queries against all the S3 objects.
- B. Create an AWS Glue table and an AWS Glue crawler.
Schedule the crawler to run daily.
Perform queries with Amazon Athena.
- C. Create an Amazon EMR cluster.
Set up an EMR File System (EMRFS) to access the S3 bucket.
Perform queries with Apache Spark.
- D. Create an Amazon Redshift cluster.
Schedule an AWS Lambda function to perform the COPY command on the Redshift cluster to load the S3 data.
Perform queries on the Redshift cluster.

Answer: D