

➤ **Vendor: Amazon**

➤ **Exam Code: SAA-C02**

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

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

**[Visit Braindump2go and Download Full Version SAA-C02 Exam Dumps](#)**

**QUESTION 999**

A company is planning to move its data to an Amazon S3 bucket. The data must be encrypted when it is stored in the S3 bucket. Additionally, the encryption key must be automatically rotated every year. Which solution will meet these requirements with the LEAST operational overhead?

- A. Move the data to the S3 bucket.  
Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3).  
Use the built-in key rotation behavior of SSE-S3 encryption keys.
- B. Create an AWS Key Management Service (AWS KMS) customer managed key Enable automatic key rotation.  
Set the S3 bucket's default encryption behavior to use the customer managed KMS key.  
Move the data to the S3 bucket.
- C. Create an AWS Key Management Service (AWS KMS) customer managed key.  
Set the S3 bucket's default encryption behavior to use the customer managed KMS key Move the data to the S3 bucket Manually rotate the KMS key every year.
- D. Encrypt the data with customer key material before moving the data to the S3 bucket.  
Create an AWS Key Management Service (AWS KMS) key without key material. Import the customer key material into the KMS key.  
Enable automatic key rotation,

**Answer: A**

**QUESTION 1000**

A company stores data in an Amazon Aurora PostgreSQL DB cluster. The company must store all the data for 5 years and must delete all the data after 5 years. The company also must indefinitely keep audit logs of actions that are performed within the database. Currently, the company has automated backups configured for Aurora. Which combination of steps should a solutions architect take to meet these requirements? (Select TWO.)

- A. Take a manual snapshot of the DB cluster.
- B. Create a lifecycle policy for the automated backups.
- C. Configure automated backup retention for 5 years.
- D. Configure an Amazon CloudWatch Logs export for the DB cluster.
- E. Use AWS Backup to take the backups and to keep the backups for 5 years.

**Answer: AD**

**QUESTION 1001**

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

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

A gaming company is moving its public scoreboard from a data center to the AWS Cloud. The company uses Amazon EC2 Windows Server instances behind an Application Load Balancer to host its dynamic application. The company needs a highly available storage solution for the application. The application consists of static files and dynamic server-side code.

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

- A. Store the static files on Amazon S3. CloudFront to cache objects at the edge.
- B. Store the static files on Amazon S3. Use Amazon ElastiCache to cache objects at the edge.
- C. Store the server-side code on Amazon Elastic File System (Amazon EFS). Mount the EFS volume on each EC2 instance to share the files.
- D. Store the server-side code on Amazon FSx for Windows File Server. Mount the FSx for Windows File Server volume on each EC2 instance to share the files.
- E. Store the server-side code on a General Purpose SSD (gp2) Amazon Elastic Block Store (Amazon EBS) volume. Mount the EBS volume on each EC2 instance to share the files.

**Answer:** AE

### **QUESTION 1002**

A company hosts a data lake on AWS. The data lake consists of data in Amazon S3 and Amazon RDS for PostgreSQL. The company needs a reporting solution that provides data visualization and includes all the data sources within the data lake. Only the company's management team should have full access to all the visualizations. The rest of the company should have only limited access.

Which solution will meet these requirements?

- A. Create an analysis in Amazon QuickSight.  
Connect all the data sources and create new datasets.  
Publish dashboards to visualize the data.  
Share the dashboards with the appropriate IAM roles.
- B. Create an analysis in Amazon QuickSight.  
Connect all the data sources and create new datasets.  
Publish dashboards to visualize the data.  
Share the dashboards with the appropriate users and groups.
- C. Create an AWS Glue table and crawler for the data in Amazon S3.  
Create an AWS Glue extract, transform, and load (ETL) job to produce reports.  
Publish the reports to Amazon S3.  
Use S3 bucket policies to limit access to the reports.
- D. Create an AWS Glue table and crawler for the data in Amazon S3.  
Use Amazon Athena Federated Query to access data within Amazon RDS for PostgreSQL.  
Generate reports by using Amazon Athena.  
Publish the reports to Amazon S3.  
Use S3 bucket policies to limit access to the reports.

**Answer:** D

### **QUESTION 1003**

A company that primarily runs its application servers on premises has decided to migrate to AWS. The company wants to minimize its need to scale its Internet Small Computer Systems Interface (iSCSI) storage on premises.

The company wants only its recently accessed data to remain stored locally.

Which AWS solution should the company use to meet these requirements?

- A. Amazon S3 File Gateway
- B. AWS Storage Gateway Tape Gateway
- C. AWS Storage Gateway Volume Gateway stored volumes
- D. AWS Storage Gateway Volume Gateway cache volumes

**Answer:** C

**QUESTION 1004**

A payment processing company records all voice communication with its customers and stores the audio files in an Amazon S3 bucket.

The company needs to capture the text from the audio files.

The company must remove from the text any personally identifiable information (PII) that belongs to customers.

What should a solutions architect do to meet these requirements?

- A. Process the audio files by using Amazon Kinesis Video Streams.  
Use an AWS Lambda function to scan for known PII patterns.
- B. When an audio file is uploaded to the S3 bucket, invoke an AWS Lambda function to start an Amazon Textract task to analyze the call recordings.
- C. Configure an Amazon Transcribe transcription job with PII redaction turned on.  
When an audio file is uploaded to the S3 bucket, invoke an AWS Lambda function to start the transcription job.  
Store the output in a separate S3 bucket.
- D. Create an Amazon Connect contact flow that ingests the audio files with transcription turned on.  
Embed an AWS Lambda function to scan for known PII patterns.  
Use Amazon EventBridge (Amazon CloudWatch Events) to start the contact flow when an audio file is uploaded to the S3 bucket.

**Answer: A**

**QUESTION 1005**

A company runs its ecommerce application on AWS. Every new order is published as a message in a RabbitMQ queue that runs on an Amazon EC2 instance in a single Availability Zone.

These messages are processed by a different application that runs on a separate EC2 instance.

This application stores the details in a PostgreSQL database on another EC2 instance. All the EC2 instances are in the same Availability Zone.

The company needs to redesign its architecture to provide the highest availability with the least operational overhead.

What should a solutions architect do to meet these requirements?

- A. Migrate the queue to a redundant pair (active/standby) of RabbitMQ instances on Amazon MQ.  
Create a Multi-AZ Auto Scaling group (or EC2 instances that host the application).  
Create another Multi-AZ Auto Scaling group for EC2 instances that host the PostgreSQL database.
- B. Migrate the queue to a redundant pair (active/standby) of RabbitMQ instances on Amazon MQ.  
Create a Multi-AZ Auto Scaling group for EC2 instances that host the application.  
Migrate the database to run on a Multi-AZ deployment of Amazon RDS for PostgreSQL.
- C. Create a Multi-AZ Auto Scaling group for EC2 instances that host the RabbitMQ queue.  
Create another Multi-AZ Auto Scaling group for EC2 instances that host the application.  
Migrate the database to run on a Multi-AZ deployment of Amazon RDS for PostgreSQL.
- D. Create a Multi-AZ Auto Scaling group for EC2 instances that host the RabbitMQ queue.  
Create another Multi-AZ Auto Scaling group for EC2 instances that host the application.  
Create a third Multi-AZ Auto Scaling group for EC2 instances that host the PostgreSQL database.

**Answer: C**

**QUESTION 1006**

A company collects data from thousands of remote devices by using a RESTful web services application that runs on an Amazon EC2 instance.

The EC2 instance receives the raw data, transforms the raw data, and stores all the data in an Amazon S3 bucket. The number of remote devices will increase into the millions soon.

The company needs a highly scalable solution that minimizes operational overhead.

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

- A. Use AWS Glue to process the raw data in Amazon S3.
- B. Use Amazon Route 53 to route traffic to different EC2 instances.
- C. Add more EC2 instances to accommodate the increasing amount of incoming data.
- D. Send the raw data to Amazon Simple Queue Service (Amazon SQS).  
Use EC2 instances to process the data.
- E. Use Amazon API Gateway to send the raw data to an Amazon Kinesis data stream.  
Configure Amazon Kinesis Data Firehose to use the data stream as a source to deliver the data to Amazon S3.

**Answer:** BE

**QUESTION 1007**

A company needs to keep user transaction data in an Amazon DynamoDB table.

The company must retain the data for 7 years.

What is the MOST operationally efficient solution that meets these requirements?

- A. Use DynamoDB point-in-time recovery to back up the table continuously.
- B. Use AWS Backup to create backup schedules and retention policies for the table.
- C. Create an on-demand backup of the table by using the DynamoDB console.  
Store the backup in an Amazon S3 bucket. Set an S3 Lifecycle configuration for the S3 bucket.
- D. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke an AWS Lambda function.  
Configure the Lambda function to back up the table and to store the backup in an Amazon S3 bucket.  
Set an S3 Lifecycle configuration for the S3 bucket.

**Answer:** C

**QUESTION 1008**

A company has an application that runs on Amazon EC2 instances and uses an Amazon Aurora database.

The EC2 instances connect to the database by using user names and passwords that are stored locally in a file.

The company wants to minimize the operational overhead of credential management.

What should a solutions architect do to accomplish this goal?

- A. Use AWS Secrets Manager.  
Turn on automatic rotation.
- B. Use AWS Systems Manager Parameter Store.  
Turn on automatic rotation.  
Create an Amazon S3 bucket to store objects that are encrypted with an AWS Key
- C. Management Service (AWS KMS) encryption key.  
Migrate the credential file to the S3 bucket.  
Point the application to the S3 bucket.
- D. Create an encrypted Amazon Elastic Block Store (Amazon EBS) volume (on each EC2 instance).  
Attach the new EBS volume to each EC2 instance.  
Migrate the credential file to the new EBS volume.  
Point the application to the new EBS volume.

**Answer:** C

**QUESTION 1009**

An image-processing company has a web application that users use to upload images. The application uploads the images into an Amazon S3 bucket.

The company has set up S3 event notifications to publish the object creation events to an Amazon Simple Queue Service (Amazon SQS) standard queue.

The SQS queue serves as the event source for an AWS Lambda function that processes the images and sends the results to users through email.

Users report that they are receiving multiple email messages for every uploaded image. A solutions architect determines that SQS messages are invoking the Lambda function more than once, resulting in multiple email messages. What should the solutions architect do to resolve this issue with the LEAST operational overhead?

- A. Set up long polling in the SQS queue by increasing the ReceiveMessage wait time to 30 seconds.
- B. Change the SQS standard queue to an SQS FIFO queue.  
Use the message deduplication ID to discard duplicate messages.
- C. Increase the visibility timeout in the SQS queue to a value that is greater than the total of the function timeout and the batch window timeout.
- D. Modify the Lambda function to delete each message from the SQS queue immediately after the message is read before processing.

**Answer: B**

#### **QUESTION 1010**

A company needs to create an Amazon Elastic Kubernetes Service (Amazon EKS) cluster to host a digital media streaming application. The EKS cluster will use a managed node group that is backed by Amazon Elastic Block Store (Amazon EBS) volumes for storage. The company must encrypt all data at rest by using a customer managed key that is stored in AWS Key Management Service (AWS KMS). Which combination of actions will meet this requirement with the LEAST operational overhead? (Select TWO.)

- A. Use a Kubernetes plugin that uses the customer managed key to perform data encryption.
- B. After creation of the EKS cluster, locate the EBS volumes.  
Enable encryption by using the customer managed key.
- C. Enable EBS encryption by default in the AWS Region where the EKS cluster will be created.  
Select the customer managed key as the default key.
- D. Create the EKS cluster. Create an IAM role that has a policy that grants permission to the customer managed key.  
Associate the role with the EKS cluster.
- E. Store the customer managed key as a Kubernetes secret in the EKS cluster.  
Use the customer managed key to encrypt the EBS volumes.

**Answer: AD**

#### **QUESTION 1011**

A company hosts an application on AWS. The application uses AWS Lambda functions and stores data in Amazon DynamoDB tables. The Lambda functions are connected to a VPC that does not have internet access. The traffic to access DynamoDB must not travel across the internet. The application must have write access to only specific DynamoDB tables. Which combination of steps should a solutions architect take to meet these requirements? (Select TWO.)

- A. Attach a VPC endpoint policy for DynamoDB to allow write access to only the specific DynamoDB tables.
- B. Attach a security group to the interface VPC endpoint to allow write access to only the specific DynamoDB tables.
- C. Create a resource-based IAM policy to grant write access to only the specific DynamoDB tables.  
Attach the policy to the DynamoDB tables.
- D. Create a gateway VPC endpoint for DynamoDB that is associated with the Lambda VPC.  
Ensure that the Lambda execution role can access the gateway VPC endpoint.
- E. Create an interface VPC endpoint for DynamoDB that is associated with the Lambda VPC.  
Ensure that the Lambda execution role can access the interface VPC endpoint.

**Answer: AE**

**QUESTION 1012**

An online retail company has more than 50 million active customers and receives more than 25,000 orders each day. The company collects purchase data for customers and stores this data in Amazon S3. Additional customer data is stored in Amazon RDS.

The company wants to make all the data available to various teams so that the teams can perform analytics.

The solution must provide the ability to manage fine-grained permissions for the data and must minimize operational overhead.

Which solution will meet these requirements?

- A. Migrate the purchase data to write directly to Amazon RDS.  
Use RDS access controls to limit access.
- B. Schedule an AWS Lambda function to periodically copy data from Amazon RDS to Amazon S3.  
Create an AWS Glue crawler.  
Use Amazon Athena to query the data.  
Use S3 policies to limit access.
- C. Create a data lake by using AWS Lake Formation.  
Create an AWS Glue JOBC connection to Amazon RDS.  
Register the S3 bucket in Lake Formation.
- D. Formation access controls to limit access.  
Create an Amazon Redshift cluster.  
Schedule an AWS Lambda function to periodically copy data from Amazon S3 and Amazon RDS to Amazon Redshift.  
Use Amazon Redshift access controls to limit access.

**Answer: C**

**QUESTION 1013**

A company runs a web application on Amazon EC2 instances in multiple Availability Zones. The EC2 instances are in private subnets.

A solutions architect implements an internet-facing Application Load Balancer (ALB) and specifies the EC2 instances as the target group.

However, the internet traffic is not reaching the EC2 instances.

How should the solutions architect reconfigure the architecture to resolve this issue?

- A. Replace the ALB with a Network Load Balancer.  
Configure a NAT gateway in a public subnet to allow internet traffic.
- B. Move the EC2 instances to public subnets.  
Add a rule to the EC2 instances' security groups to allow outbound traffic to 0.0.0.0/0.
- C. Update the route tables for the EC2 instances' subnets to send 0.0.0.0/0 traffic through the internet gateway route.  
Add a rule to the EC2 instances' security groups to allow outbound traffic to 0.0.0.0/0.
- D. Create public subnets in each Availability Zone.  
Associate the public subnets with the ALB. Update the route tables for the public subnets with a route to the private subnets.

**Answer: C**

**QUESTION 1014**

A company stores its data objects in Amazon S3 Standard storage. A solutions architect has found that 75% of the data is rarely accessed after 30 days. The company needs all the data to remain immediately accessible with the same high availability and resiliency, but the company wants to minimize storage costs.

Which storage solution will meet these requirements?

- A. Move the data objects to S3 Glacier Deep Archive after 30 days.
- B. Move the data objects to S3 Standard-Infrequent Access (S3 Standard-IA) after 30 days.
- C. Move the data objects to S3 One Zone-Infrequent Access (S3 One Zone-IA) after 30 days.

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

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

D. Move the data objects to S3 One Zone-Infrequent Access (S3 One Zone-IA) immediately.

**Answer: B**

**QUESTION 1015**

A company is building an application in the AWS Cloud. The application will store data in Amazon S3 buckets in two AWS Regions.

The company must use an AWS Key Management Service (AWS KMS) customer managed key to encrypt all data that is stored in the S3 buckets.

The data in both S3 buckets must be encrypted and decrypted with the same KMS key. The data and the key must be stored in each of the two Regions.

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

- A. Create an S3 bucket in each Region.  
Configure the S3 buckets to use server-side encryption with Amazon S3 managed encryption keys (SSE-S3).  
Configure replication between the S3 buckets.
- B. Create a customer managed multi-Region KMS key.  
Create an S3 bucket in each Region.  
Configure replication between the S3 buckets.  
Configure the application to use the KMS key with client-side encryption.
- C. Create a customer managed KMS key and an S3 bucket in each Region.  
Configure the S3 buckets to use server-side encryption with Amazon S3 managed encryption keys (SSE-S3).  
Configure replication between the S3 buckets.
- D. Create a customer managed KMS key and an S3 bucket in each Region.  
Configure the S3 buckets to use server-side encryption with AWS KMS keys (SSE-KMS).  
Configure replication between the S3 buckets.

**Answer: C**

**QUESTION 1016**

A company has an AWS Glue extract, transform, and load (ETL) job that runs every day at the same time. The job processes XML data that is in an Amazon S3 bucket. New data is added to the S3 bucket every day.

A solutions architect notices that AWS Glue is processing all the data during each run.

What should the solutions architect do to prevent AWS Glue from reprocessing old data?

- A. Edit the job to use job bookmarks.
- B. Edit the job to delete data after the data is processed.
- C. Edit the job by setting the NumberOfWorkers field to 1.
- D. Use a FindMatches machine learning (ML) transform.

**Answer: B**

**QUESTION 1017**

A company has an ordering application that stores customer information in Amazon RDS for MySQL. During regular business hours, employees run one-time queries for reporting purposes.

Timeouts are occurring during order processing because the reporting queries are taking a long time to run.

The company needs to eliminate the timeouts without preventing employees from performing queries.

What should a solutions architect do to meet those requirements?

- A. Create a read replica. Move reporting queries to the read replica.
- B. Create a read replica. Distribute the ordering application to the primary DB instance and the read replica.
- C. Migrate the ordering application to Amazon DynamoDB with on-demand capacity.

D. Schedule the reporting queries for non-peak hours.

**Answer: B**

**QUESTION 1018**

A company hosts a serverless application on AWS. The application uses Amazon API Gateway, AWS Lambda, and an Amazon RDS for PostgreSQL database.

The company notices an increase in application errors that result from database connection timeouts during times of peak traffic or unpredictable traffic.

The company needs a solution that reduces the application failures with the least amount of change to the code. What should a solutions architect do to meet these requirements?

- A. Reduce the Lambda concurrency rate.
- B. Enable RDS Proxy on the RDS DB instance.
- C. Resize the RDS DB instance class to accept more connections.
- D. Migrate the database to Amazon DynamoDB with on-demand scaling

**Answer: B**

**QUESTION 1019**

A company stores its application logs in an Amazon CloudWatch Logs log group.

A new policy requires the company to store all application logs in Amazon OpenSearch Service (Amazon Elasticsearch Service) in near-real time.

Which solution will meet this requirement with the LEAST operational overhead?

- A. Configure a CloudWatch Logs subscription to stream the logs to Amazon OpenSearch Service (Amazon Elasticsearch Service).
- B. Create an AWS Lambda function.  
Use the log group to invoke the function to write the logs to Amazon OpenSearch Service (Amazon Elasticsearch Service).
- C. Create an Amazon Kinesis Data Firehose delivery stream.  
Configure the log group as the delivery stream's source.  
Configure Amazon OpenSearch Service (Amazon Elasticsearch Service) as the delivery stream's destination.
- D. Install and configure Amazon Kinesis Agent on each application server to deliver the logs to Amazon Kinesis Data Streams.  
Configure Kinesis Data Streams to deliver the logs to Amazon OpenSearch Service (Amazon Elasticsearch Service)

**Answer: C**

**QUESTION 1020**

A global company hosts its web application on Amazon EC2 instances behind an Application Load Balancer (ALB).

The web application has static data and dynamic data. The company stores its static data in an Amazon S3 bucket.

The company wants to improve performance and reduce latency for the static data and dynamic data.

The company is using its own domain name registered with Amazon Route 53.

What should a solutions architect do to meet these requirements?

- A. Create an Amazon CloudFront distribution that has the S3 bucket and the ALB as origins.  
Configure Route 53 to route traffic to the CloudFront distribution.
- B. Create an Amazon CloudFront distribution that has the ALB as an origin.  
Create an AWS Global Accelerator standard accelerator that has the S3 bucket as an endpoint.  
Configure Route 53 to route traffic to the CloudFront distribution.
- C. Create an Amazon CloudFront distribution that has the S3 bucket as an origin.  
Create an AWS Global Accelerator standard accelerator that has the ALB and the CloudFront distribution as endpoints.

Create a custom domain name that points to the accelerator DNS name.

Use the custom domain name as an endpoint for the web application.

- D. Create an Amazon CloudFront distribution that has the ALB as an origin.  
Create an AWS Global Accelerator standard accelerator that has the S3 bucket as an endpoint.  
Create two domain names.  
Point one domain name to the CloudFront DNS name for dynamic content.  
Point the other domain name to the accelerator DNS name for static content.  
Use the domain names as endpoints for the web application.

**Answer: B**

#### **QUESTION 1021**

A rapidly growing ecommerce company is running its workloads in a single AWS Region.

A solutions architect must create a disaster recovery (DR) strategy that includes a different AWS Region.

The company wants its database to be up to date in the DR Region with the least possible latency.

The remaining infrastructure in the DR Region needs to run at reduced capacity and must be able to scale up if necessary.

Which solution will meet these requirements with the LOWEST recovery time objective (RTO)?

- A. Use an Amazon Aurora global database with a pilot light deployment.
- B. Use an Amazon Aurora global database with a warm standby deployment.
- C. Use an Amazon RDS Multi-AZ DB instance with a pilot light deployment.
- D. Use an Amazon RDS Multi-AZ DB instance with a warm standby deployment.

**Answer: B**

#### **QUESTION 1022**

A company needs to develop a repeatable solution to process time-ordered information from websites around the world.

The company collects the data from the websites by using Amazon Kinesis Data Streams and stores the data in Amazon S3.

The processing logic needs to collect events and handle data from the last 5 years.

The processing logic also must generate results in an S3 bucket so that a business intelligence application can analyze and compare the results.

The processing must be repeated multiple times.

What should a solutions architect do to meet these requirements?

- A. Use Amazon S3 to collect events.  
Create an AWS Lambda function to process the events.  
Create different Lambda functions to handle repeated processing.
- B. Use Amazon EventBridge (Amazon CloudWatch Events) to collect events Set AWS Lambda as an event target.  
Use EventBridge (CloudWatch Events) to create an archive for the events and to replay the events.
- C. Use an Amazon Simple Queue Service (Amazon SQS) FIFO queue to collect events.  
Process the events by using Amazon EC2.  
Use AWS Step Function to create an archive for the events and to replay the events
- D. Use Amazon Managed Streaming for Apache Kafka (Amazon MSK) to collect events.  
Process the events by using Amazon Elastic Kubernetes Service (Amazon EKS).  
Use Amazon MSK to create an archive for the events and to replay the events.

**Answer: B**

#### **QUESTION 1023**

A company's web application consists of an Amazon API Gateway API in front of an AWS Lambda function and an Amazon DynamoDB database.

The Lambda function handles the business logic, and the DynamoDB table hosts the data. The application uses

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

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

Amazon Cognito user pools to identify the individual users of the application.

A solutions architect needs to update the application so that only users who have a subscription can access premium content.

- A. Enable API caching and throttling on the API Gateway API
- B. Set up AWS WAF on the API Gateway API Create a rule to filter users who have a subscription
- C. Apply fine-grained IAM permissions to the premium content in the DynamoDB table
- D. Implement API usage plans and API keys to limit the access of users who do not have a subscription.

**Answer: A**

#### **QUESTION 1024**

A company wants to migrate its existing on-premises monolithic application to AWS. The company wants to keep as much of the front- end code and the backend code as possible.

However, the company wants to break the application into smaller applications. A different team will manage each application.

The company needs a highly scalable solution that minimizes operational overhead.

Which solution will meet these requirements?

- A. Host the application on AWS Lambda Integrate the application with Amazon API Gateway.
- B. Host the application with AWS Amplify.  
Connect the application to an Amazon API Gateway API that is integrated with AWS Lambda.
- C. Host the application on Amazon EC2 instances.  
Set up an Application Load Balancer with EC2 instances in an Auto Scaling group as targets.
- D. Host the application on Amazon Elastic Container Service (Amazon ECS).  
Set up an Application Load Balancer with Amazon ECS as the target.

**Answer: D**

#### **QUESTION 1025**

A company uses a legacy application to produce data in CSV format .The legacy application stores the output data In Amazon S3.

The company is deploying a new commercial off-the-shelf (COTS) application that can perform complex SQL queries to analyze data that is stored Amazon Redshift and Amazon S3 only.

However the COTS application cannot process the csv files that the legacy application produces.

The company cannot update the legacy application to produce data in another format.

The company needs to implement a solution so that the COTS application can use the data that the legacy applicator produces.

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

- A. Create a AWS Glue extract, transform, and load (ETL) job that runs on a schedule.  
Configure the ETL job to process the .csv files and store the processed data in Amazon Redshift.
- B. Develop a Python script that runs on Amazon EC2 instances to convert the .csv files to sql files  
invoke the Python script on cron schedule to store the output files in Amazon S3.
- C. Create an AWS Lambda function and an Amazon DynamoDB table.  
Use an S3 event to invoke the Lambda function.  
Configure the Lambda function to perform an extract transform, and load (ETL) job to process the .csv files and store the processed data in the DynamoDB table.
- D. Use Amazon EventBridge (Amazon CloudWatch Events) to launch an Amazon EMR cluster on a weekly schedule.  
Configure the EMR cluster to perform an extract, transform, and load (ETL) job to process the .csv files and store the processed data in an Amazon Redshift table.

**Answer: C**

**QUESTION 1026**

A company has a business system that generates hundreds of reports each day. The business system saves the reports to a network share in CSV format. The company needs to store this data in the AWS Cloud in near-real time for analysis. Which solution will meet these requirements with the LEAST administrative overhead?

- A. Use AWS DataSync to transfer the files to Amazon S3.  
Create a scheduled task that runs at the end of each day.
- B. Create an Amazon S3 File Gateway.  
Update the business system to use a new network share from the S3 File Gateway.
- C. Use AWS DataSync to transfer the files to Amazon S3.  
Create an application that uses the DataSync API in the automation workflow.
- D. Deploy an AWS Transfer for SFTP endpoint.  
Create a script that checks for new files on the network share and uploads the new files by using SFTP.

**Answer: B**

**QUESTION 1027**

A company produces batch data that comes from different databases. The company also produces live stream data from network sensors and application APIs. The company needs to consolidate all the data into one place for business analytics. The company needs to process the incoming data and then stage the data in different Amazon S3 buckets. Teams will later run onetime queries and import the data into a business intelligence tool to show key performance indicators (KPIs). Which combination of steps will meet these requirements with the LEAST operational overhead? (Select TWO.)

- A. Use Amazon Athena for one-time queries.  
Use Amazon QuickSight to create dashboards for KPIs
- B. Use Amazon Kinesis Data Analytics for one-time queries.  
Use Amazon QuickSight to create dashboards for KPIs
- C. Create custom AWS Lambda functions to move the individual records from the databases to an Amazon Redshift cluster
- D. Use an AWS Glue extract transform, and load (ETL) job to convert the data into JSON format.  
Load the data into multiple Amazon OpenSearch Service (Amazon Elasticsearch Service) clusters.
- E. Use blueprints in AWS Lake Formation to identify the data that can be ingested into a data lake.  
Use AWS Glue to crawl the source, extract the data, and load the data into Amazon S3 in Apache Parquet format

**Answer: CD**

**QUESTION 1028**

A company is using a centralized AWS account to store log data in various Amazon S3 buckets. A solutions architect needs to ensure that the data is encrypted at rest before the data is uploaded to the S3 buckets. The data also must be encrypted in transit. Which solution meets these requirements?

- A. Use client-side encryption to encrypt the data that is being uploaded to the S3 buckets.
- B. Use server-side encryption to encrypt the data that is being uploaded to the S3 buckets.
- C. Create bucket policies that require the use of server-side encryption with S3 managed encryption keys (SSE-S3) for S3 uploads.
- D. Enable the security option to encrypt the S3 buckets through the use of a default AWS Key Management Service (AWS KMS) key.

**Answer: C**

**QUESTION 1029**

A company is running a popular social media website.  
The website gives users the ability to upload images to share with other users.  
The company wants to make sure that the images do not contain inappropriate content.  
The company needs a solution that minimizes development effort.  
What should a solutions architect do to meet these requirements?

- A. Use Amazon Comprehend to detect inappropriate content.  
Use human review for low-confidence predictions.
- B. Use Amazon Rekognition to detect inappropriate content.  
Use human review for low-confidence predictions.
- C. Use Amazon SageMaker to detect inappropriate content.  
Use ground truth to label low-confidence predictions.
- D. Use AWS Fargate to deploy a custom machine learning model to detect inappropriate content.  
Use ground truth to label low-confidence predictions.

**Answer: B**

**QUESTION 1030**

A company has migrated an application to Amazon EC2 Linux instances.  
One of these EC2 instances runs several 1-hour tasks on a schedule.  
These tasks were written by different teams and have no common programming language.  
The company is concerned about performance and scalability while these tasks run on a single instance.  
A solutions architect needs to implement a solution to resolve these concerns.  
Which solution will meet these requirements with the LEAST operational overhead?

- A. Use AWS Batch to run the tasks as jobs.  
Schedule the jobs by using Amazon EventBridge (Amazon CloudWatch Events).
- B. Convert the EC2 instance to a container.  
Use AWS App Runner to create the container on demand to run the tasks as jobs.
- C. Copy the tasks into AWS Lambda functions.  
Schedule the Lambda functions by using Amazon EventBridge (Amazon CloudWatch Events).
- D. Create an Amazon Machine Image (AMI) of the EC2 instance that runs the tasks.  
Create an Auto Scaling group with the AMI to run multiple copies of the instance.

**Answer: C**

**QUESTION 1031**

A company wants to build a data lake on AWS from data that is stored in an on-premises Oracle relational database.  
The data lake must receive ongoing updates from the on-premises database.  
Which solution will meet these requirements with the LEAST operational overhead?

- A. Use AWS DataSync to transfer the data to Amazon S3.  
Use AWS Glue to transform the data and integrate the data into a data lake.
- B. Use AWS Snowball to transfer the data to Amazon S3.  
Use AWS Batch to transform the data and integrate the data into a data lake.
- C. Use AWS Database Migration Service (AWS DMS) to transfer the data to Amazon S3.  
Use AWS Glue to transform the data and integrate the data into a data lake.
- D. Use an Amazon EC2 instance to transfer the data to Amazon S3.  
Configure the EC2 instance to transform the data and integrate the data into a data lake.

**Answer: C**

**QUESTION 1032**

A media company collects and analyzes user activity data on premises.

The company wants to migrate this capability to AWS.

The user activity data store will continue to grow and will be petabytes in size.

The company needs to build a highly available data ingestion solution that facilitates on-demand analytics of existing data and new data with SQL.

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

- A. Send activity data to an Amazon Kinesis data stream.  
Configure the stream to deliver the data to an Amazon S3 bucket.
- B. Send activity data to an Amazon Kinesis Data Firehose delivery stream.  
Configure the stream to deliver the data to an Amazon Redshift cluster
- C. Place activity data in an Amazon S3 bucket.  
Configure Amazon S3 to run an AWS Lambda function on the data as the data arrives in the S3 bucket.
- D. Create an ingestion service on Amazon EC2 instances that are spread across multiple Availability Zones.  
Configure the service to forward data to an Amazon RDS Multi-AZ database.

**Answer: B**

#### **QUESTION 1033**

A company recently launched a variety of new workloads on Amazon EC2 instances in its AWS account.

The company needs to create a strategy to access and administer the instances remotely and securely.

The company needs to implement a repeatable process that works with native AWS services and follows the AWS WellArchitected Framework.

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

- A. Use the EC2 serial console to directly access the terminal interface of each instance for administration.
- B. Attach the appropriate IAM role to each existing instance and new instance.  
Use AWS Systems Manager Session Manager to establish a remote SSH session.
- C. Create an administrative SSH key pair.  
Load the public key into each EC2 instance.  
Deploy a bastion host in a public subnet to provide a tunnel for administration of each instance.
- D. Establish an AWS Site-to-Site VPN connection.  
Instruct administrators to use their local on-premises machines to connect directly to the instances by using SSH keys across the VPN tunnel.

**Answer: B**

#### **QUESTION 1034**

A company hosts a multiplayer gaming application on AWS.

The company wants the application to read data with sub-millisecond latency and run one-time queries on historical data.

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

- A. Use Amazon RDS for data that is frequently accessed.  
Run a periodic custom script to export the data to an Amazon S3 bucket.
- B. Store the data directly in an Amazon S3 bucket.  
Implement an S3 Lifecycle policy to move older data to S3 Glacier Deep Archive for long-term storage.  
Run one-time queries on the data in Amazon S3 by using Amazon Athena
- C. Use Amazon DynamoDB with DynamoDB Accelerator (DAX) for data that is frequently accessed.  
Export the data to an Amazon S3 bucket by using DynamoDB table export.  
Run one-time queries on the data in Amazon S3 by using Amazon Athena.
- D. Use Amazon DynamoDB for data that is frequently accessed.  
Turn on streaming to Amazon Kinesis Data Streams.

Use Amazon Kinesis Data Firehose to read the data from Kinesis Data Streams.  
Store the records in an Amazon S3 bucket.

**Answer: C**