

# Amazon-Web-Services

## Exam Questions DVA-C01

AWS Certified Developer Associate Exam



### NEW QUESTION 1

- (Exam Topic 4)

A company is building a serverless application that uses AWS Lambda. The application includes Lambda functions that are exposed by Amazon API Gateway. The functions will use several large third-party libraries, and the build artifacts will exceed 50 MB in size.

Which combination of steps should a developer take to prepare and perform the deployment? (Select TWO.)

- A. Issue the `aws lambda update-function-code` CLI command with the `-zip-file fileb://my-function.zip` parameter
- B. Upload the build artifact to Amazon S3.
- C. Issue the `aws cloudformation package` CLI command.
- D. Issue the `aws lambda update-function-code` CLI command with the `-s3-bucket` and `-s3-key` parameters.
- E. Issue the `aws lambda update-function-code` CLI command with a parameter that points to the source code in AWS CodeCommit.

**Answer:** BD

### NEW QUESTION 2

- (Exam Topic 4)

A developer at a company recently created a serverless application to process and show data from business reports. The application's user interface (UI) allows users to select and start processing the files. The UI displays a message when the result is available to view. The application uses AWS Step Functions with AWS Lambda functions to process the files. The developer used Amazon API Gateway and Lambda functions to create an API to support the UI.

The company's UI team reports that the request to process a file is often returning timeout errors because of the size or complexity of the files. The UI team wants the API to provide an immediate response so that the UI can display a message while the files are being processed. The backend process that is invoked by the API needs to send an email message when the report processing is complete.

What should the developer do to configure the API to meet these requirements?

- A. Change the API Gateway route to add an `X-Amz-Invocation-Type` header with a static value of 'Event' in the integration request.
- B. Deploy the API Gateway stage to apply the changes.
- C. Change the configuration of the Lambda function that implements the request to process a file. Configure the maximum age of the event so that the Lambda function will run asynchronously.
- D. Change the API Gateway timeout value to match the Lambda function timeout value.
- E. Deploy the API Gateway stage to apply the changes.
- F. Change the API Gateway route to add an `X-Amz-Target` header with a static value of 'Async' in the integration request.
- G. Deploy the API Gateway stage to apply the changes.

**Answer:** A

### NEW QUESTION 3

- (Exam Topic 4)

A company is hosting a workshop for external users and wants to share the reference documents with the external users for 7 days. The company stores the reference documents in an Amazon S3 bucket that the company owns.

What is the MOST secure way to share the documents with the external users?

- A. Use S3 presigned URLs to share the documents with the external user.
- B. Set an expiration time of 7 days.
- C. Move the documents to an Amazon WorkDocs folder.
- D. Share the links of the WorkDocs folder with the external users.
- E. Create temporary IAM users that have read-only access to the S3 bucket.
- F. Share the access keys with the external user.
- G. Expire the credentials after 7 days.
- H. Create a role that has read-only access to the S3 bucket.
- I. Share the Amazon Resource Name (ARN) of this role with the external users.

**Answer:** A

### NEW QUESTION 4

- (Exam Topic 4)

A distributed application includes an AWS Lambda function that runs successfully in the DEV environment with 128 MB of memory assigned. The same function is failing in the TEST environment. The developer is monitoring the application using AWS X-Ray, but the Lambda function cannot be seen on the X-Ray service graph. The Lambda execution role has AWS X-Ray permissions.

What is the MOST LIKELY cause for AWS X-Ray not showing any data for the Lambda function?

- A. The AWS SDK needs to be included in the AWS Lambda deployment package.
- B. VPC Flow Logs are not enabled for the application VPC.
- C. Active tracing needs to be enabled for the Lambda function.
- D. The memory needs to be increased to 2 GB for the TEST environments.

**Answer:** C

#### Explanation:

<https://stackoverflow.com/questions/43728674/enabling-x-ray-support-in-aws-lambda> You need to check the "Enable Active Tracing" checkbox in the Lambda console.

### NEW QUESTION 5

- (Exam Topic 4)

A company is building an application for stock trading. The application needs sub-millisecond latency for processing trade requests. The company uses Amazon DynamoDB to store all the trading data that is used to process each trading request.

A development team performs load testing on the application and finds that the data retrieval time is higher than expected. The development team needs a solution that reduces the data retrieval time with the least possible effort.

Which solution meets these requirements?

- A. Add local secondary indexes (LSIs) for the trading data
- B. Store the trading data in Amazon S3, and use S3 Transfer Acceleration.
- C. Add retries with exponential backoff for DynamoDB queries.
- D. Use DynamoDB Accelerator (DAX) to cache the trading data

**Answer: D**

#### **NEW QUESTION 6**

- (Exam Topic 4)

A company is running an application on AWS Elastic Beanstalk in a single-instance environment. The company's deployments must avoid any downtime. Which deployment option will meet these requirements?

- A. All at once
- B. Rolling
- C. Rolling with additional batch
- D. Immutable

**Answer: D**

#### **NEW QUESTION 7**

- (Exam Topic 4)

A developer is writing a new web application that will be deployed and managed with AWS Elastic Beanstalk. The application will include an Amazon RDS DB instance. What steps should the developer take to access the RDS DB instance from the code? (Select TWO.)

- A. Modify the endpoint name using either the AWS Management Console or AWS CLI
- B. Upload the driver to Amazon S3 and reference it in the code
- C. Download the appropriate database driver and include it with the application.
- D. Construct a connection string using the Elastic Beanstalk environment variables
- E. Create a CNAME record referencing database instances ALIAS.

**Answer: CD**

#### **NEW QUESTION 8**

- (Exam Topic 4)

A developer has built an application running on AWS Lambda using AWS Serverless Application Model (AWS SAM). What is the correct sequence of steps to successfully deploy the application?

- A. \* 1. Build the SAM template in Amazon EC2.\* 2. Package the SAM template to Amazon EBS storage.\* 3. Deploy the SAM template from Amazon EBS.
- B. \* 1. Build the SAM template locally.\* 2. Package the SAM template onto Amazon S3.\* 3. Deploy the SAM template from Amazon S3.
- C. \* 1. Build the SAM template locally.\* 2. Deploy the SAM template from Amazon S3.\* 3. Package the SAM template for use.
- D. \* 1. Build the SAM template locally.\* 2 Package the SAM template from AWS CodeCommit.

**Answer: B**

#### **NEW QUESTION 9**

- (Exam Topic 4)

A movie fan club hosts a serverless web application in an Amazon S3 bucket. The application uses an AWS Lambda function that is exposed by an Amazon API Gateway API. The function queries an Amazon DynamoDB table to list actors sorted by movie. In the DynamoDB table, Actor is the primary key, Movie is the sort key, and Role and Year are attributes.

In the web application, a developer wants to add a page that is named Phase 1 that lists only the movies that were released between 2008 and 2012. The developer needs to fetch the Phase 1 items in a way that minimizes the impact on the DynamoDB table.

Which solution will meet these requirements?

- A. Create a global secondary index (GSI) with the Year attribute as the sort key
- B. Create a Lambda function to return the results from a new method in the API.
- C. Design a Lambda function that scans the DynamoDB table and filters the results for the Phase 1 items. Invoke the function from a new method in the API.
- D. Use a DynamoDB stream to send items that are filtered by Year to a new DynamoDB table
- E. Invoke a Lambda function from a new method in the API.
- F. Set up an Amazon CloudFront distribution
- G. Create a Lambda@Edge function to filter the items that are returned from the API request.

**Answer: B**

#### **NEW QUESTION 10**

- (Exam Topic 4)

Which solution will meet these requirements?

- A. Build the container from the amazon/aws-xray-daemon base image
- B. Use the AWS X-Ray SDK to instrument the application.
- C. Install the Amazon CloudWatch agent on the container image
- D. Use the CloudWatch SDK to publish custom metrics from each of the microservices.
- E. Install the AWS X-Ray daemon on each of the ECS instances.
- F. Configure AWS CloudTrail data events to capture the traffic between the microservices.

**Answer: C**

### NEW QUESTION 10

- (Exam Topic 4)

A company is using Amazon Cognito user pools for sign-up and login functionality for a web application. The company is using Amazon RDS for the application's data persistence and is using Amazon API Gateway and AWS Lambda for the application's API functionality. Users must provide their first name, last name, email address, and phone number to sign up. All API endpoints have a Cognito user pool authorizer to guard against unauthenticated requests.

A developer wants to show a personalized welcome screen to users after they log in. The welcome screen needs to show the user's first name and the user's previous login date. According to company policy, developers who work on the web application cannot store any personally identifiable information in RDS instances.

Which solution should the developer implement to meet these requirements?

- A. After successful login, submit a Cognito request for user token
- B. When redirecting to the welcome screen, provide the identity token in the Authorization header of the request
- C. Extract the user name from the given\_name claim and the user's universally unique identifier (UUID) from the sub claim inside the identity token
- D. Use the UUID as the key to store and retrieve the previous login information from the database.
- E. After successful login, submit a Cognito request for user token
- F. When redirecting to the welcome screen, provide the access token in the Authorization header of the request
- G. Extract the user name from the given\_name claim and the user's universally unique identifier (UUID) from the sub claim inside the access token
- H. Use the UUID as the key to store and retrieve the previous login information from the database.
- I. After successful login, submit a Cognito request for user token
- J. When redirecting to the welcome screen, provide the identity token in the Authorization header of the request
- K. Extract the user name from the given\_name claim and the user's universally unique identifier (UUID) from the iss claim inside the identity token
- L. Use the UUID as the key to store and retrieve the previous login information from the database.
- M. After successful login, submit a Cognito request for user token
- N. When redirecting to the welcome screen, provide the access token in the Authorization header of the request
- O. Extract the user name from the given name claim and the user's universally unique identifier (UUID) from the iss claim inside the access token
- P. Use the UUID as the key to store and retrieve the previous login information from the database.

**Answer: C**

### NEW QUESTION 12

- (Exam Topic 4)

A company has an application that uses Amazon Cognito user pools as an identity provider. The company must secure access to user records. The company is implementing multi-factor authentication (MFA). The company also wants to send a login activity notification by email every time a user logs in.

What is the MOST operationally efficient solution that meets this requirement?

- A. Create an AWS Lambda function that uses Amazon Simple Email Service (Amazon SES) to send the email notification
- B. Add an Amazon API Gateway to invoke the function
- C. Call the API from the client side when login confirmation is received.
- D. Create an AWS Lambda function that uses Amazon Simple Email Service (Amazon SES) to send the email notification
- E. Add an Amazon Cognito post authentication Lambda trigger for the function.
- F. Create an AWS Lambda function that uses Amazon Simple Email Service (Amazon SES) to send the email notification
- G. Create an Amazon CloudWatch Logs log subscription filter to invoke the function based on the login status.
- H. Configure Amazon Cognito to stream all logs to Amazon Kinesis Data Firehose
- I. Create an AWS Lambda function to process the streamed logs and it send the email notification based on the login status of each user.

**Answer: C**

### NEW QUESTION 13

- (Exam Topic 4)

A developer is creating a serverless orchestrator that performs a series of steps to process incoming IoT data. The orchestrator transforms the data, performs a series of calculations, and stores the results in Amazon DynamoDB. The entire process is completed in less than a minute.

The orchestrator must process tens of thousands of transactions each second. The orchestrator must not miss a transaction and must process each transaction at least once.

Which solution will meet these requirements MOST cost-effectively?

- A. Use Amazon Simple Notification Service (Amazon SNS) to process the data through an HTTP or HTTPS endpoint.
- B. Use AWS Step Functions to process the data as Standard Workflows.
- C. Use AWS Step Functions to process the data as Synchronous Express Workflows.
- D. Use AWS Step Functions to process the data as Asynchronous Express Workflows.

**Answer: D**

### NEW QUESTION 14

- (Exam Topic 4)

A developer is storing JSON files in an Amazon S3 bucket. The developer wants to securely share an object with a specific group of people. How can the developer securely provide temporary access to the objects that are stored in the S3 bucket?

- A. Set object retention on the file
- B. Use the AWS software development kit (SDK) to restore the object before subsequent request
- C. Provide the bucket's URL.
- D. Use the AWS software development kit (SDK) to generate a presigned URL
- E. Provide the presigned URL.
- F. Set a bucket policy that restricts access after a period of time
- G. Provide the bucket's S3 URL.
- H. Configure static web hosting on the S3 bucket
- I. Provide the bucket's web URL.

**Answer: B**

### NEW QUESTION 15

- (Exam Topic 4)

A company that manages movie reviews wants to make its movie review data available to its customers by calling a set of REST web service endpoints. The company will develop the retrieval functionality as AWS Lambda functions and will expose the functionality to customers as an Amazon API Gateway REST API. The company needs to ensure that no consumer exceeds 100 requests a day to the API during the initial deployment. The company decides to use API Gateway API keys to restrict access. The company creates and issues API keys for each customer.

What should the company do next to meet these requirements with the LEAST administrative effort?

- A. Create a usage plan that applies throttling at 100 requests a da
- B. Associate the usage plan with the API keys of all customers.
- C. Create an Amazon DynamoDB table to track all the requests that use a particular API ke
- D. For each request to the API, count the number of records in the DynamoDB table for that day for the API ke
- E. If the number of requests is 100 or greater, generate an exception.
- F. Create a usage plan that applies a quota of 100 requests a da
- G. Associate the usage plan with the API keys of all customers.
- H. Create an Amazon Aurora table to track all the requests that use a particular API ke
- I. For each request to the API, count the number of records in the Aurora table for that day for the API ke
- J. If the number of requests is 100 or greater, generate an exception.

**Answer: C**

#### Explanation:

What is quota in API gateway?

API quotas usually describe a certain amount of calls for longer intervals. For example, your API quota might be 5,000 calls per month. Remember that this could be combined with a rate limit or throttling setup e.g. 20 TPS (Transactions per Second).

What is throttling in AWS API gateway?

AWS throttling limits are applied across all accounts and clients in a region. These limit settings exist to prevent your API—and your account—from being overwhelmed by too many requests. These limits are set by AWS and can't be changed by a customer.

<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-create-usage-plans-with-console.html> Choose Enable quota, and set specify a value (for example, 5000) for a selected time interval (for example, Month).

### NEW QUESTION 19

- (Exam Topic 4)

A developer deployed an application to an Amazon EC2 instance. The application needs to know the public IPv4 address of the instance. How can the application find this information?

- A. Query the instance metadata from <http://169.254.169.254/latest/meta-data/>.
- B. Query the instance user data from <http://169.254.169.254/latest/user-data/>.
- C. Query the Amazon Machine Image (AMI) information from <http://169.254.169.254/latest/meta-data/ami/>.
- D. Check the hosts file of the operating system.

**Answer: A**

### NEW QUESTION 20

- (Exam Topic 4)

A company manages a financial services application that stores a large volume of data in an Amazon DynamoDB table. A developer needs to improve the performance of the DynamoDB read queries without increasing the cost.

Which solution meets these requirements?

- A. Use parallel scans
- B. Add a local secondary index (LSI).
- C. Create a DynamoDB Accelerator (DAX) cluster.
- D. Query with the Projection Expression parameter

**Answer: C**

### NEW QUESTION 23

- (Exam Topic 4)

A developer is implementing an AWS Lambda function that will be invoked when an object is uploaded to Amazon S3. The developer wants to test the Lambda function in a local development machine before publishing the function to a production AWS account.

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

- A. Upload an object to Amazon S3 by using the `aws s3api put-object` CLI command
- B. Wait for the local Lambda invocation from the S3 event.
- C. Create a sample JSON text file for a put object S3 event
- D. Invoke the Lambda function locally
- E. Use the `aws lambda invoke` CLI command with the JSON file and Lambda function name as arguments.
- F. Use the `sam local start-lambda` CLI command to start Lambda
- G. Use the `sam local generate-event s3 put` CLI command to create the Lambda test JSON file
- H. Use the `sam local invoke` CLI command with the JSON file as the argument to invoke the Lambda function.
- I. Create a JSON string for the put object S3 event
- J. In the AWS Management Console, use the JSON string to create a test event for the local Lambda function
- K. Perform the test.

**Answer: D**

### NEW QUESTION 25

- (Exam Topic 4)

A developer is writing an AWS Lambda function. The Lambda function needs to access items that are stored in an Amazon DynamoDB table. What is the MOST secure way to configure this access for the Lambda function?

- A. Create an IAM user that has permissions to access the DynamoDB tabl
- B. Create an access key for this use
- C. Store the access key ID and secret..... key in the Lambda function environment variables.
- D. Add a resource-based policy to the DynamoDB table to allow access from the Lambda function's IAM role.
- E. Create an IAM policy that allows access to the DynamoDB tabl
- F. Attach this policy to the Lambda function's 1AM role.
- G. Create a DynamoDB Accelerator (DAX) cluste
- H. Configure the Lambda function to use the DAX cluster to access the DynamoDB table.

**Answer:** A

#### NEW QUESTION 29

- (Exam Topic 4)

A developer is migrating a legacy monolithic application to AWS and wants to convert the application's internal processes to microservices. The application's internal processes communicate through internal asynchronous messaging. Occasionally messages need to be reprocessed by multiple microservices. How should the developer migrate the application's internal messaging to AWS to meet these requirements?

- A. Use Amazon Simple Queue Service (Amazon SQS) queues to communicate messages between the microservices.
- B. Use Amazon API Gateway to provide REST interfaces between the microservices.
- C. Use Amazon Kinesis Data Streams to communicate messages between the microservices.
- D. Use Amazon API Gateway to provide WebSocket APIs between the microservices.

**Answer:** A

#### NEW QUESTION 31

- (Exam Topic 4)

A banking application processes thousands of transactions each second. Each transaction payload must have end-to-end encryption. The application encrypts each transaction locally by using the AWS Key Management Service (AWS KMS) GenerateDataKey operation. A developer is testing the application and receives a Throttling Exception error.

Which actions are best practices to resolve this error? (Select TWO.)

- A. Use the LocalCryptoMaterialsCache feature of the AWS Encryption SDK encryption library.
- B. Call the AWS KMS Encrypt operation directly to allow AWS KMS to encrypt the data.
- C. Create a case in the AWS Support Center to increase the quota for the account.
- D. Use Amazon Simple Queue Service (Amazon SQS) to queue the requests to AWS KMS.
- E. Switch to an AWS KMS custom key store.

**Answer:** AC

#### Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/kms-throttlingexception-error/>

#### NEW QUESTION 32

- (Exam Topic 4)

A company that has multiple offices uses an Amazon DynamoDB table to store employee payroll information. Item attributes consist of employee names, office identifiers, and cumulative daily hours worked. The most frequently used query extracts a report of an alphabetical subset of employees for a specific office. Which design of the DynamoDB table primary key will have the MINIMUM performance impact?

- A. Partition key on the office identifier and sort key on the employee name
- B. Partition key on the employee name and sort key on the office identifier
- C. Partition key on the employee name
- D. Partition key on the office identifier

**Answer:** B

#### NEW QUESTION 33

- (Exam Topic 4)

A developer is designing a serverless application for an ecommerce website. An Amazon API Gateway API exposes..... user operations. The website features shopping carts for the users. The shopping carts must be stored for extended..... the front-end application. The load on the application will vary significantly based on the time of day and the promotional sales that are offered..... scale automatically to meet these changing demands.

Which solution will meet these requirements?

- A. Store the data objects on an Amazon RDS DB instanc
- B. Cache the data objects in memory by using Amazon ElastiCache.
- C. Store the data objects on Amazon EC2 instances behind an Application Load Balance
- D. Use session affinity (sticky sessions) for each user's shopping cart.
- E. Store the data objects in Amazon S3 bucket
- F. Cache the data objects by using Amazon CloudFront with the maximum TTL.
- G. Store the data objects in Amazon DynamoDB table
- H. Cache the data objects by using DynamoDB Accelerator (DAX).

**Answer:** D

#### Explanation:

A developer is designing a serverless application for an ecommerce website." DynamoDB is the best option to keep the serverless design.

#### NEW QUESTION 35

- (Exam Topic 4)

A developer is monitoring an application that runs on an Amazon EC2 instance. The developer has configured a custom Amazon CloudWatch metric with a granularity of 1 second. If any issues occur, the developer wants to be notified within 30 seconds by Amazon Simple Notification Service (Amazon SNS). What should the developer do to meet this requirement?

- A. Configure a high-resolution CloudWatch alarm.
- B. Set up a custom CloudWatch dashboard.
- C. Use Amazon CloudWatch Logs Insights.
- D. Change to a default CloudWatch metric.

**Answer: D**

#### NEW QUESTION 40

- (Exam Topic 4)

A developer is building a three-tier application with an Application Load Balancer (ALB), Amazon EC2 instances, and Amazon RDS. There is an alias record in Amazon Route 53 that points to the ALB. When the developer tries to access the ALB from a laptop, the request times out. Which logs should the developer investigate to verify that the request is reaching the AWS network?

- A. VPC Flow Logs
- B. Amazon Route 53 logs
- C. AWS Systems Manager Agent logs
- D. Amazon CloudWatch agent logs

**Answer: A**

#### NEW QUESTION 44

- (Exam Topic 4)

A development team set up a pipeline to launch a test environment. The developers want to automate tests for their application. The team created an AWS CodePipeline stage to deploy the application to a test environment in batches using AWS Elastic Beanstalk. A later CodePipeline stage contains a single action that uses AWS CodeBuild to run numerous automated Selenium-based tests on the deployed application. The team must speed up the pipeline without removing any of the individual tests.

Which set of actions will MOST effectively speed up application deployment and testing?

- A. Set up an all-at-once deployment in Elastic Beanstalk.
- B. Run tests in parallel with multiple CodeBuild actions.
- C. Set up a rolling update in Elastic Beanstalk.
- D. Run tests in serial with a single CodeBuild action.
- E. Set up an immutable update in Elastic Beanstalk.
- F. Run tests in serial with a single CodeBuild action.
- G. Set up a traffic-splitting deployment in Elastic Beanstalk.
- H. Run tests in parallel with multiple CodeBuild actions.

**Answer: A**

#### Explanation:

<https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.deploy-existing-version.html>

All at once – The quickest deployment method. Suitable if you can accept a short loss of service, and if quick deployments are important to you. With this method, Elastic Beanstalk deploys the new application version to each instance. Then, the web proxy or application server might need to restart. As a result, your application might be unavailable to users (or have low availability) for a short time.

#### NEW QUESTION 45

- (Exam Topic 4)

A developer deployed an application to an Amazon EC2 instance. The application needs to know the public IPv4 address of the instance. How can the application find this information?

- A. Query the instance metadata from <http://169.254.169.254/latest/meta-data/>.
- B. Query the instance user data from <http://169.254.169.254/latest/user-data/>.
- C. Query the Amazon Machine Image (AMI) information from <http://169.254.169.254/latest/meta-data/ami/>.
- D. Check the hosts file of the operating system.

**Answer: A**

#### NEW QUESTION 47

- (Exam Topic 4)

A developer is designing a serverless application for a game in which users register and log in through a web browser. The application makes requests on behalf of users to a set of AWS Lambda functions that run behind an Amazon API Gateway HTTP API.

The developer needs to implement a solution to register and log in users on the application's sign-in page. The solution must minimize operational overhead and must minimize ongoing management of user identities.

Which solution will meet these requirements?

- A. Create Amazon Cognito user pools for external social identity provider.
- B. Configure IAM roles for the identity pools.
- C. Program the sign-in page to create users' IAM groups with the IAM roles attached to the groups.
- D. Create an Amazon RDS for SQL Server DB instance to store the users and manage the permissions to the backend resources in AWS.
- E. Configure the sign-in page to register and store the users and their passwords in an Amazon DynamoDB table with an attached IAM policy.

**Answer: A**

#### NEW QUESTION 48

- (Exam Topic 4)

An ecommerce application is running behind an Application Load Balancer. A developer observes some unexpected load on the application during non-peak hours. The developer wants to analyze patterns for the client IP addresses that use the application. Which HTTP header should the developer use for this analysis?

- A. The X-Forwarded-Proto header
- B. The X-Forwarded-Host header
- C. The X-Forwarded-For header
- D. The X-Forwarded-Port header

**Answer: C**

#### Explanation:

<https://developer.mozilla.org/en-US/docs/Web/HTTP/Headers/X-Forwarded-Proto>

#### NEW QUESTION 51

- (Exam Topic 4)

A developer is designing a serverless application that customers use to select seats for a concert venue. Customers send the ticket requests to an Amazon API Gateway API with an AWS Lambda function that acknowledges the order and generates an order ID. The application includes two additional Lambda functions: one for inventory management and one for payment processing. These two Lambda functions run in parallel and write the order to an Amazon Dynamo DB table. The application must provide seats to customers according to the following requirements. If a seat is accidentally sold more than once, the first order that the application received must get the seat. In these cases, the application must process the payment for only the first order. However, if the first order is rejected during payment processing, the second order must get the seat. In these cases, the application must process the payment for the second order. Which solution will meet these requirements?

- A. Send the order ID to an Amazon Simple Notification Service (Amazon SNS) FIFO topic that fans out to one Amazon Simple Queue Service (Amazon SQS) FIFO queue for inventory management and another SQS FIFO queue for payment processing.
- B. Change the Lambda function that generates the order ID to initiate the Lambda function for inventory management.
- C. Then initiate the Lambda function for payment processing.
- D. Send the order ID to an Amazon Simple Notification Service (Amazon SNS) topic.
- E. Subscribe the Lambda functions for inventory management and payment processing to the topic.
- F. Deliver the order ID to an Amazon Simple Queue Service (Amazon SQS) queue.
- G. Configure the Lambda functions for inventory management and payment processing to poll the queue.

**Answer: A**

#### Explanation:

Inventory & Payment functions are running in parallel. So going with Fanout option. <https://docs.aws.amazon.com/sns/latest/dg/sns-common-scenarios.html>

#### NEW QUESTION 54

- (Exam Topic 4)

An application adds a processing date to each transaction that it receives. The application writes each transaction to an Amazon DynamoDB table by using the PutItem operation. Each transaction has a unique ID (transactionID). Sometimes the application receives transactions more than once. A developer notices that duplicate transactions in DynamoDB have the latest processing date instead of the date when the transaction was first received. Duplicate records happen infrequently, and most of the transactions are unique. What is the MOST cost-effective solution that the developer can implement to ensure that PutItem does not update an existing record?

- A. Call the GetItem operation first to confirm that the record does not exist.
- B. Then call PutItem.
- C. Enable the TTL attribute on the DynamoDB table.
- D. Implement a conditional put by using the attribute\_exists(transactionID) condition expression.
- E. Implement a conditional put by using the attribute\_not\_exists(transactionID) condition expression.

**Answer: C**

#### NEW QUESTION 58

- (Exam Topic 4)

A developer deploys a custom application to three Amazon EC2 instances. The application processes messages from an Amazon Simple Queue Service (Amazon SQS) standard queue with default settings. When the developer runs a load test on the Amazon SQS queue, the developer discovers that the application processes many messages multiple times. How can the developer ensure that the application processes each message exactly once?

- A. Modify the SQS standard queue to an SQS FIFO queue.
- B. Process the messages on one EC2 instance instead of three instances.
- C. Create a new SQS FIFO queue.
- D. Point the application to the new queue.
- E. Increase the DelaySeconds value on the current SQS queue.

**Answer: C**

#### Explanation:

<https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/FIFO-queues-moving.html> Moving from a standard queue to a FIFO queue: If you have an existing application that uses standard queues and you want to take advantage of the ordering or exactly-once processing features of FIFO queues, you need to configure the queue and your application correctly.

Note:

You can't convert an existing standard queue into a FIFO queue. To make the move, you must either create a new FIFO queue for your application or delete your existing standard queue and recreate it as a FIFO queue.

### NEW QUESTION 60

- (Exam Topic 4)

A developer is creating a web application that collects highly regulated and confidential user data through a POST request. The web application is served through Amazon CloudFront. User names and phone numbers must be encrypted at the edge and must remain encrypted throughout the entire application stack. What is the MOST secure way to meet these requirements?

- A. Enforce Match Viewer with HTTPS Only on CloudFront.
- B. Use only the newest TLS security policy on CloudFront.
- C. Enforce a signed URL on CloudFront on the front end.
- D. Use field-level encryption on CloudFront.

**Answer: D**

#### Explanation:

Field-level encryption allows you to enable your users to securely upload sensitive information to your web servers. The sensitive information provided by your users is encrypted at the edge, close to the user, and remains encrypted throughout your entire application stack.  
<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/field-level-encryption.html>

### NEW QUESTION 61

- (Exam Topic 4)

A company is expanding the compatibility of its photo-sharing mobile app to hundreds of additional devices with unique screen dimensions and resolutions. Photos are stored in Amazon S3 in their original format and resolution. The company uses an Amazon CloudFront distribution to serve the photos. The app includes the dimension and resolution of the display as GET parameters with every request.

A developer needs to implement a solution that optimizes the photos that are served to each device to reduce load time and increase photo quality. Which solution will meet these requirements MOST cost-effectively?

- A. Use S3 Batch Operations to invoke an AWS Lambda function to create new variants of the photos with the required dimensions and resolution.
- B. Create a dynamic CloudFront origin that automatically maps the request of each device to the corresponding photo variant.
- C. Use S3 Batch Operations to invoke an AWS Lambda function to create new variants of the photos with the required dimensions and resolution.
- D. Create a Lambda@Edge function to route requests to the corresponding photo variant by using request headers.
- E. Create a Lambda@Edge function that optimizes the photos upon request and returns the photos as a response.
- F. Change the CloudFront TTL cache policy to the maximum value possible.
- G. Create a Lambda@Edge function that optimizes the photos upon request and returns the photos as a response.
- H. In the same function, store a copy of processed photos on Amazon S3 for subsequent requests.

**Answer: C**

#### Explanation:

This solution will meet the requirements most cost-effectively because it allows the developer to use a Lambda@Edge function to optimize the photos on the fly, without the need to pre-generate multiple variants of the photos for different devices. This approach can reduce the overall storage and compute costs associated with generating and storing multiple photo variants. Additionally, changing the CloudFront TTL cache policy to the maximum value possible can help reduce the number of times the Lambda@Edge function needs to be executed, further reducing the cost.  
<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Expiration.html>

### NEW QUESTION 63

- (Exam Topic 4)

A company has an application that provides blog hosting services to its customers. The application includes an Amazon DynamoDB table with a primary key. The primary key consists of the customers' Username as a partition key and the NumberOfBlogs as a sort key. The application stores the TotalReactionsOnBlogs as an attribute on the same DynamoDB table.

A developer needs to implement an operation to retrieve the top 10 customers based on the greatest number of reactions on their blogs. This operation must not consume the DynamoDB table's existing read capacity.

What should the developer do to meet these requirements in the MOST operationally efficient manner?

- A. For the existing DynamoDB table, create a new global secondary index (GSI) that has the Username as a partition key and the TotalReactionsOnBlogs as a sort key.
- B. For the existing DynamoDB table, create a new local secondary index (LSI) that has the Username as a partition key and the TotalReactionsOnBlogs as a sort key.
- C. Back up and restore the DynamoDB table to a new DynamoDB table. Create a new global secondary index (GSI) that has the Username as a partition key and the TotalReactionsOnBlogs as a sort key. Delete the old DynamoDB table.
- D. Back up and restore the DynamoDB table to a new DynamoDB table.
- E. Create a new local secondary index (LSI) that has the Username as a partition key and the TotalReactionsOnBlogs as a sort key.
- F. Delete the old DynamoDB table.

**Answer: B**

### NEW QUESTION 66

- (Exam Topic 4)

A developer is working on an application that is deployed on an Amazon EC2 instance. The application needs to transfer a file to an Amazon S3 bucket. What should the developer do to authenticate the application's access to the S3 bucket in the MOST secure way?

- A. Create an access key for an IAM user.
- B. Store the access key in the application's environment variables.
- C. Create an IAM role.
- D. Create an access key for the role.
- E. Store the access key in the application's environment variables.
- F. Associate an IAM role with the EC2 instance.
- G. Use the instance metadata service to retrieve the credentials.
- H. Configure a bucket policy for the S3 bucket.
- I. Allow access from the EC2 instance ID in the bucket policy.

**Answer: B**

**NEW QUESTION 67**

- (Exam Topic 4)

A data-processing application includes an AWS Lambda function that processes data in several steps. Recently, the function has been reaching the Lambda timeout. A developer wants to use AWS X-Ray to find out how long each step is taking so that the developer can determine which step is causing the timeout. Which combination of actions should the developer take to accomplish this goal? (Select TWO.)

- A. Modify the application to call the PutMetricData API operation after each processing step.
- B. Include the time taken in milliseconds.
- C. Use the aws lambda update-function-configuration AWS CLI command to enable active tracing on the Lambda function.
- D. Modify the application to record each processing step in an X-Ray subsegment by using the X-Ray software development kit (SDK).
- E. Add the xray:PutTraceSegments permission and the xray:PutTelemetryRecords permission to the Lambda function's execution role.
- F. Modify the application to put each processing step in a separate Lambda layer.
- G. Include all the layers in the Lambda function.

**Answer: B**

**NEW QUESTION 72**

- (Exam Topic 4)

A developer deploys an ecommerce application on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances run in an Amazon EC2 Auto Scaling group. The EC2 instances are based on an Amazon Machine Image (AMI) that uses an Amazon Elastic Block Store (Amazon EBS) root volume. After deployment, the developer notices that a third of the instances seem to be idle. These instances are not receiving requests from the load balancer. The developer verifies that all the instances are registered with the load balancer. The developer must implement a solution to allow the EC2 instances to receive requests from the load balancer.

Which action will meet this requirement?

- A. Reregister the failed instances with the ALB.
- B. Enable all Availability Zones for the ALB.
- C. Use the instance refresh feature to redeploy the EC2 Auto Scaling group.
- D. Restart the EC2 instances that are not receiving traffic.

**Answer: C**

**Explanation:**

<https://aws.amazon.com/blogs/compute/introducing-instance-refresh-for-ec2-auto-scaling/>

**NEW QUESTION 77**

- (Exam Topic 4)

An AWS Lambda function requires read access to an Amazon S3 bucket and requires read/write access to an Amazon DynamoDB table. The correct IAM policy already exists. What is the MOST secure way to grant the Lambda function access to the S3 bucket and the DynamoDB table?

- A. Attach the existing IAM policy to the Lambda function.
- B. Create an IAM role for the Lambda function.
- C. Attach the existing IAM policy to the role.
- D. Attach the role to the Lambda function.
- E. Create an IAM user with programmatic access.
- F. Attach the existing IAM policy to the user.
- G. Add the user access key ID and secret access key as environment variables in the Lambda function.
- H. Add the AWS account root user access key ID and secret access key as encrypted environment variables in the Lambda function.

**Answer: B**

**NEW QUESTION 79**

- (Exam Topic 4)

A developer is writing a web application that allows users to sign in. The application will run on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances will run in an Auto Scaling group across multiple Availability Zones.

How can the developer ensure that users stay signed in when the Auto Scaling group is scaled down?

- A. Enable sticky sessions on the ALB target group.
- B. Create an Amazon DynamoDB table.
- C. Configure the application to use the DynamoDB table to store session state such as login status.
- D. Create an Amazon Elastic Block Store (Amazon EBS) volume.
- E. Use EBS Multi-Attach to attach the volume to all instances in the Auto Scaling group.
- F. Configure the application to use the volume to store session state such as login status.
- G. Enable deregistration delay on the ALB target group.

**Answer: B**

**NEW QUESTION 83**

- (Exam Topic 4)

A developer is writing an application that will run on Amazon EC2 instances in an Auto Scaling group. The developer wants to externalize the session state to support the application. Which AWS services or resources can the developer use to meet these requirements? (Select TWO.)

- A. Amazon DynamoDB
- B. Amazon Cognito
- C. Amazon ElastiCache
- D. Application Load Balancer

E. Amazon Simple Queue Service (Amazon SQS)

**Answer:** AC

#### **NEW QUESTION 86**

- (Exam Topic 4)

A company is concerned that a malicious user could deploy unauthorized changes to the code for an AWS Lambda function. What can a developer do to ensure that only trusted code is deployed to Lambda?

- A. Turn on the trusted code option in AWS CodeDeploy
- B. Add the CodeDeploy digital certificate to the Lambda package before deploying the package to Lambda
- C. Define the code signing configuration in the Lambda console Use AWS Signer to digitally sign the Lambda package before deploying the package to Lambda
- D. Link Lambda to AWS Key Management Service (AWS KMS) in the Lambda console
- E. Use AWS KMS to digitally sign the Lambda package before deploying the package to Lambda.
- F. Set the KmsKeyArn property of the Lambda function to the Amazon Resource Name (ARN) of a trusted key before deploying the package to Lambda.

**Answer:** B

#### **NEW QUESTION 91**

- (Exam Topic 4)

A developer has an Amazon DynamoDB table that must be in provisioned mode to comply with user requirements. The application needs to support the following:

- Average item size: 10 KB
- Item reads each second: 10 strongly consistent
- Item writes each second: 2 transactional

Which read and write capacity cost-effectively meets these requirements?

- A. Read 10; write 2
- B. Read 30; write 40
- C. Use on-demand scaling
- D. Read 300; write 400

**Answer:** B

#### **NEW QUESTION 93**

- (Exam Topic 4)

A company has an ecommerce application. To track product reviews, the company's development team uses an Amazon DynamoDB table.

Every record includes the following:

- A Review ID, a 16-digit universally unique identifier (UUID)
- A Product ID and User ID. 16-digit UUIDs that reference other tables
- A Product Rating on a scale of 1—5 - An optional comment from the user

The table partition key is the Review ID. The most performed query against the table is to find the 10 reviews with the highest rating for a given product.

Which index will provide the FASTEST response for this query?

- A. A global secondary index (GSI) with Product ID as the partition key and Product Rating as the sort key
- B. A global secondary index (GSI) with Product ID as the partition key and Review ID as the sort key
- C. A local secondary index (LSI) with Product ID as the partition key and Product Rating as the sort key
- D. A local secondary index (LSI) with Review ID as the partition key and Product ID as the sort key

**Answer:** A

#### **NEW QUESTION 95**

- (Exam Topic 3)

A developer has an AWS CodePipeline pipeline that invokes AWS CodeBuild in the build stage. The developer wants to pass in a variable from CodePipeline so that the variable can be read in the CodeBuild buildspec.yml file.

How can the developer accomplish this goal?

- A. Configure a unique CodePipeline variable namespace and variables as key-value pairs that define each of the variables required in CodeBuild
- B. Configure a CodePipeline environment variable that contains a JSON document that defines each of the variables required in CodeBuild
- C. Configure an AWS CloudFormation stack set that contains a JSON document that defines each of the variables required in CodeBuild. Reference the stack set from CodePipeline
- D. Configure an AWS CodeArtifact repository to store each environment variable. Reference CodeArtifact from CodePipeline and CodeBuild

**Answer:** B

#### **NEW QUESTION 96**

- (Exam Topic 3)

An application runs on multiple EC2 instances behind an ELB.

Where is the session data best written so that it can be served reliably across multiple requests?

- A. Write data to Amazon ElasticCache.
- B. Write data to Amazon Elastic Block Store.
- C. Write data to Amazon EC2 instance Block Store.
- D. Write data to the root filesystem.

**Answer:** A

#### **NEW QUESTION 99**

- (Exam Topic 3)

A company wants to make sure that only one user from its Admin group has the permanent right to delete an Amazon EC2 resource. There should be no changes in the existing policy under the Admin group. What should a developer use to meet these requirements?

- A. AWS managed policy
- B. Inline policy
- C. IAM trust relationship
- D. AWS Security Token Service (AWS STS)

**Answer: B**

#### NEW QUESTION 103

- (Exam Topic 3)

A developer is creating a serverless web application and maintains different branches of code. The developer wants to avoid updating the Amazon API Gateway target endpoint each time a new code push is performed. What solution would allow the developer to perform a code push efficiently, without the need to update the API Gateway?

- A. Associate different AWS Lambda functions to an API Gateway target endpoint
- B. Create different stages in API Gateway, then associate API Gateway with AWS Lambda
- C. Create aliases and versions in AWS Lambda
- D. Tag the AWS Lambda functions with different names

**Answer: B**

#### NEW QUESTION 106

- (Exam Topic 3)

A developer works in an environment with multiple AWS accounts that have AWS Lambda functions processing the same 100 KB payloads. The developer wants to centralize the point of origin of the payloads to one account and have all the Lambda functions be invoked whenever the initiating event occurs in the parent account.

How can the developer design the workflow in the MOST efficient way, so all the multi-account Lambda functions get invoked when the event occurs?

- A. Create a Lambda function in the parent account and use cross-account IAM roles with the AWS Security Token Service (AWS STS) AssumeRole API call to make AWS Lambda invoke the API call to invoke all the cross-account Lambda functions.
- B. Subscribe all the multi-account Lambda functions to an Amazon SNS topic and make a SNS Publish API call with the payload to the SNS topic.
- C. Set up an Amazon SQS queue with the queue policy permitting the ReceiveMessage action for multi-account Lambda function
- D. Then send the payload to the SQS queue using the sqs:SendMessage permission and poll the queue using multi-account Lambda functions.
- E. Use a worker on an Amazon EC2 instance to poll for the payload event
- F. Invoke all Lambda functions using the Lambda Invoke API after using cross-account IAM roles with the AWS Security Token Service (AWS STS) AssumeRole API call.

**Answer: B**

#### Explanation:

<https://aws.amazon.com/es/blogs/compute/cross-account-integration-with-amazon-sns/>

#### NEW QUESTION 107

- (Exam Topic 3)

A company recently experienced some unexpected downtime. After investigating, the company determines that a developer mistakenly terminated several production Amazon EC2 instances.

What should the company do to BEST protect against accidental terminations in the future?

- A. Enable EC2 termination protection on all production instances unless approval has been given through AWS Resource Access Manager.
- B. Modify the developer group's permissions policy to deny them access to delete production instances unless approval has been given through AWS Resource Access Manager.
- C. Modify the developer group's permission policy to require multi-factor authentication (MFA) only production instances are being deleted. Enable EC2 termination protection on production instances.
- D. Enable EC2 termination protection on production instances
- E. Deny the developer group's permissions policy access to terminate instances
- F. Create a new role that the developer can assume when termination is necessary.

**Answer: A**

#### NEW QUESTION 108

- (Exam Topic 3)

A developer has written code for an application and wants to share it with other developers on the team to receive feedback. The shared application code needs to be stored long-term with multiple versions and batch change tracking.

Which AWS service should the developer use?

- A. AWS CodeBuild
- B. Amazon S3
- C. AWS CodeCommit
- D. AWS Cloud9

**Answer: C**

#### NEW QUESTION 112

- (Exam Topic 3)

A developer is building a web and mobile application for two types of users: regular users and guest users. Regular users are required to log in, but guest users do

not log in Users should see only their data regardless of whether they authenticate Users need AWS credentials before they can access AWS resources What is the MOST secure solution that the developer can implement to allow access for guest users?

- A. Use an Amazon Cognito credentials provider to issue temporary credentials that are linked to an unauthenticated role that has access to the required resources.
- B. Set up an IAM user that has permissions to the required resource
- C. Hardcode the IAM credentials in the web and mobile application
- D. Generate temporary keys that are stored in AWS Key Management Service (AWS KMS) Use the temporary keys to access the required resources
- E. Generate temporary credential
- F. Store the temporary credentials in AWS Secrets Manager Use the temporary credentials to access the required resources

**Answer: D**

#### **NEW QUESTION 114**

- (Exam Topic 3)

An application development team decides to use AWS X Ray to monitor application code to analyze performance and perform root cause analysis What does the team need to do to begin using X Ray? (Select TWO )

- A. Log instrumentation output into an Amazon SQS queue
- B. Use a visualization tool to view application traces
- C. Instrument application code using the AWS SDK
- D. Install the X-Ray agent on the application servers
- E. Create an Amazon DynamoDB table to store the trace logs

**Answer: DE**

#### **NEW QUESTION 116**

- (Exam Topic 3)

A developer must increase read performance from an unencrypted Amazon S3 bucket. The application requires 100,000 read requests each second Cost-effectiveness is a priority. What would be the SIMPLEST approach to implement these requirements?

- A. Create 20 or more prefixes in Amazon S3 Place files by prefix
- B. Read in parallel by prefixes
- C. Create 20 or more AWS accounts Create a bucket in each account Read in parallel by bucket
- D. Deploy Memcached on Amazon EC2 Cache the files in memory Retrieve from the Memcached cache
- E. Copy all files to Amazon DynamoDB Index the files with S3 metadata Retrieve from DynamoDB

**Answer: A**

#### **NEW QUESTION 119**

- (Exam Topic 3)

A developer must modify an Alexa skill backed by an AWS Lambda function to access an Amazon DynamoDB table in a second account A role in the second account has been created with permissions to access the table How should the table be accessed?

- A. Modify the Lambda function execution role's permissions to include the new role
- B. Change the Lambda function execution role to be the new role
- C. Assume the new role in the Lambda function when accessing the table
- D. Store the access key and the secret key for the new role and use them when accessing the table

**Answer: A**

#### **NEW QUESTION 123**

- (Exam Topic 3)

A developer must build a mobile application that allows users to read and write data from an Amazon DynamoDB table to store user state for each unique user. The solution needs to limit data access to allow users access only to their own data Which solution below is the most secure?

- A. Embed AWS access credentials into the application and create DynamoDB queries that limit user access.
- B. Use Amazon Cognito identity pools to assign unique identifiers and provide user access
- C. Modify the DynamoDB table to allow public read and writes, then add client-side filtering
- D. Create a web portal for users to create an account on AWS Directory Service

**Answer: C**

#### **NEW QUESTION 128**

- (Exam Topic 3)

A developer is building an application on Amazon EC2 The developer encountered an "Access Denied" error on some of the API calls to AWS services while testing The developer needs to modify permissions that have been already given to the instance How can these requirements be met with minimal changes and minimum downtime?

- A. Make a new IAM role with the needed permissions Stop the instance
- B. Attach the new IAM role to the instance Start the instance.
- C. Delete the existing IAM role Attach a new IAM role with the needed permissions
- D. Stop the instance Update the attached IAM role adding the needed permission
- E. Start the instance
- F. Update the attached IAM role adding the needed permissions

**Answer: D**

#### NEW QUESTION 130

- (Exam Topic 3)

A developer has built a market application that stores pricing data in Amazon DynamoDB with Amazon ElastiCache in front. The prices of items in the market change frequently. Sellers have begun complaining that, after they update the price of an item, the price does not actually change in the product listing. What could be causing this issue?

- A. The cache is not being invalidated when the price of the item is changed
- B. The price of the item is being retrieved using a write-through ElastiCache cluster
- C. The DynamoDB table was provisioned with insufficient read capacity
- D. The DynamoDB table was provisioned with insufficient write capacity.

**Answer: A**

#### NEW QUESTION 135

- (Exam Topic 3)

An application running on multiple Amazon EC2 instances pulls messages from an SQS queue. A requirement for the application is that all messages must be encrypted at rest.

Developers are instructed to use methods that allow for centralized support requirements whenever possible.

Which of the following solutions supports these requirements?

- A. Encrypt individual messages by using client-side encryption with customer managed keys, then write to the SQS queue.
- B. Encrypt individual messages by using SQS Extended Client and the Amazon S3 encryption client.
- C. Create an SQS queue, and encrypt the queue by using server-side encryption with AWS KMS
- D. Create an SQS queue and encrypt the queue by using client-side encryption

**Answer: C**

#### NEW QUESTION 139

- (Exam Topic 3)

Multiple development teams are working on a project to migrate a monolithic application to a microservices-based application running on AWS Lambda. The teams need a way to centrally manage code that is shared across multiple functions.

Which approach requires the LEAST maintenance?

- A. Each team maintains the code for the common components in their own code repository
- B. They build and deploy the components with their Lambda functions together.
- C. One team builds a Lambda layer to include the common components and shares the layer with the other teams
- D. Each team builds and publishes the component they want to share to an Amazon S3 bucket. The Lambda functions will download the components from the bucket
- E. One team builds a Docker container for the common components and shares the container with the other teams

**Answer: C**

#### NEW QUESTION 140

- (Exam Topic 3)

A developer is building an AWS Lambda function that will dynamically generate and send a weekly newsletter to 100,000 users. This newsletter contains both static text and images. The developer needs a fast and highly scalable place to store the images that will be hyperlinked in the newsletter.

Where should the developer store these images?

- A. Use an Amazon DynamoDB table with DynamoDB Streams and read capacity auto scaling enabled
- B. Use an Amazon S3 bucket and S3 Transfer Acceleration to speed up the image download
- C. Use an Amazon Aurora database with a public DNS endpoint and auto scaling enabled
- D. Use an Amazon S3 backed Amazon CloudFront distribution with a high Time-to-Live (TTL) to maximize caching

**Answer: D**

#### NEW QUESTION 141

- (Exam Topic 3)

An e-commerce application is using Amazon Simple Notification Service (Amazon SNS) with an AWS Lambda subscription to save all new orders into an Amazon DynamoDB table. The company wants to record all the orders that are more than a certain amount of money in a separate table. The company wants to avoid changes to the processes that post orders to Amazon SNS or the current Lambda function that saves the orders to the DynamoDB table.

How can a developer implement this feature with the LEAST change to the existing application?

- A. Create another Lambda subscription with the SNS message attribute value matching a filter option to save the appropriate orders to a separate table
- B. Create another SNS topic, and also send orders in that topic. Create a Lambda subscription with a numeric value filter option to save the appropriate orders to a separate table
- C. Create another Lambda subscription with the SNS message numeric value matching a filter option to save the appropriate orders to a separate table
- D. Modify the Lambda code to filter the orders and save the appropriate orders to a separate table

**Answer: D**

#### NEW QUESTION 146

- (Exam Topic 3)

A company uses a third-party tool to build, bundle, and package its applications on-premises, and store them locally. The company uses Amazon EC2 instances to run its front-end applications. How can an application be deployed from the source control system onto the EC2 instances?

- A. Use AWS CodeDeploy and point it to the local storage to directly deploy a bundle to a zipped file
- B. Use AWS CodeDeploy and point it to the local storage to directly deploy a bundle to a zip file

- C. or tar.gz format
- D. Upload the bundle to an Amazon S3 bucket and specify the S3 location when doing a deployment using AWS CodeDeploy
- E. Create a repository using AWS CodeCommit to automatically trigger a deployment to the EC2 instances
- F. Use AWS CodeBuild to automatically deploy the latest build to the latest EC2 instances

**Answer: B**

#### **NEW QUESTION 149**

- (Exam Topic 3)

A developer has code stored in an Amazon S3 bucket. The code must be deployed as an AWS Lambda function across multiple accounts in the same Region as the S3 bucket. The Lambda function will be deployed using an AWS CloudFormation template that is run for each account. What is the MOST secure approach to allow access to the Lambda code in the S3 bucket?

- A. Grant the CloudFormation execution role S3 list and get permissions. Add a bucket policy to Amazon S3 with the Principal of "AWS": [account numbers].
- B. Grant the CloudFormation execution role S3 get permissions. Add a bucket policy to Amazon S3 with the Principal of "".
- C. Use a service-based link to grant the Lambda function S3 list and get permissions by explicitly adding the S3 bucket's account number in the resource.
- D. Use a service-based link to grant the Lambda function S3 get permissions and add a Resource of "\*" to allow access to the S3 bucket.

**Answer: A**

#### **NEW QUESTION 152**

- (Exam Topic 3)

Given the following AWS CloudFormation template:

What is the MOST efficient way to reference the new Amazon S3 bucket from another AWS CloudFormation template?

- A. Add an Export declaration to the outputs section of the original template and use ImportValue in other templates.
- B. Add Exported: True to the ContentBucket in the original template and use ImportResource in other templates.
- C. Create a custom AWS CloudFormation resource that gets the bucket name from the ContentBucket resource of the first stack.
- D. Use Fn: : Include to include the existing template in other template and use the ContentBucket resource directly.

**Answer: D**

#### **NEW QUESTION 156**

- (Exam Topic 3)

An application contains two components: one component to handle HTTP requests, and another component to handle background processing tasks. Each component must scale independently. The developer wants to deploy this application using AWS Elastic Beanstalk. How should this application be deployed, based on these requirements?

- A. Deploy the application in a single Elastic Beanstalk environment.
- B. Deploy each component in a separate Elastic Beanstalk environment.
- C. Use multiple Elastic Beanstalk environments for the HTTP component but one environment for the background task component.
- D. Use multiple Elastic Beanstalk environments for the background task component but one environment for the HTTP component.

**Answer: A**

#### **NEW QUESTION 161**

- (Exam Topic 3)

An application uses Amazon DynamoDB as its backend database. The application experiences sudden spikes in traffic over the weekend and variable but predictable spikes during weekdays. The capacity needs to be set to avoid throttling errors at all times. How can this be accomplished cost-effectively?

- A. Use provisioned capacity with AWS Auto Scaling throughout the week.
- B. Use on-demand capacity for the weekend and provisioned capacity with AWS Auto Scaling during the weekdays.
- C. Use on-demand capacity throughout the week.
- D. Use provisioned capacity with AWS Auto Scaling enabled during the weekend and reserved capacity enabled during the weekdays.

**Answer: A**

#### **NEW QUESTION 163**

- (Exam Topic 3)

An application is processing clickstream data using Amazon Kinesis. The clickstream data feed into Kinesis experiences periodic spikes. The PutRecords API call occasionally fails and the logs show that the failed call returns the response shown below.

```
{
  "FailedRecordCount": 1,
  "Records": [
    {
      "SequenceNumber": "21269319989900637946712965403778482371",
      "ShardId": "shardId-0000000000001"
    },
    {
      "ErrorCode": "ProvisionedThroughputExceededException",
      "ErrorMessage": "Rate exceeded for shard shardId-0000000000001 in
        stream exampleStreamName under account 123456789."
    },
    {
      "SequenceNumber": "21269319989999637946712965403778482985",
      "ShardId": "shardId-0000000000002"
    }
  ]
}
```

Which techniques will help mitigate this exception? (Select TWO.)

- A. Implement retries with exponential backoff
- B. Use a PutRecord API instead of PutRecords
- C. Reduce the frequency and/or size of the requests
- D. Use Amazon SNS instead of Kinesis.
- E. Reduce the number of KCL consumers.

**Answer:** AC

**NEW QUESTION 165**

- (Exam Topic 3)

A developer tested an application locally and then deployed it to AWS Lambda While testing the application remotely the Lambda function fails with an access denied message. How can this issue be addresssed?

- A. Update the Lambda function's execution role to include the missing permissions
- B. Update the Lambda function's resource policy to include the missing permissions
- C. Include an IAM policy document at the root of the deployment package and redeploy the Lambda function.
- D. Redeploy the Lambda function using an account with access to the AdministratorAccess policy

**Answer:** A

**NEW QUESTION 168**

- (Exam Topic 3)

A developer is building an application that reads 90 Items of data each second from an Amazon DynamoDB table. Each item is 3 KB in size. The table is configured to use eventually consistent reads  
 How many read capacity units should the developer provision for the table?

- A. 25
- B. 35
- C. 45
- D. 85

**Answer:** C

**NEW QUESTION 170**

- (Exam Topic 3)

A developer is using Amazon DynamoDB to store application data . The developer wants to further improve application performance by reducing response times for read and write operations.

Which DynamoDB feature should be used to meet these requirement?

- A. Amazon DynamoDB Streams
- B. Amazon DynamoDB Accelerator
- C. Amazon DynamoDB global tables
- D. Amazon DynamoDB transactions

**Answer:** B

**Explanation:**

<https://aws.amazon.com/ko/blogs/database/amazon-dynamodb-accelerator-dax-a-read-throughwrite-through-cac>

**NEW QUESTION 171**

- (Exam Topic 3)

A development team is migrating a monolithic application to Amazon API Gateway with AWS Lambda integrations using the AWS CD The zip deployment package

exceeds the Lambda direct upload deployment package size limit. How should the Lambda function be deployed?

- A. Use the zip file to create a Lambda layer and reference it using the `-code` CLI parameter
- B. Create a Docker image and reference the image using the `--docker-image` CLI parameter
- C. Upload a deployment package using the `--zip-file` CLI parameter
- D. Upload a deployment package to Amazon S3 and reference Amazon S3 using the `--code` CLI parameter

**Answer:** D

#### NEW QUESTION 176

- (Exam Topic 3)

A developer is developing an application that uses signed requests (Signature Version 4) to call other AWS services. The developer has created a canonical request, has created the string to sign, and has calculated signing information. Which methods could the developer use to complete a signed request? (Select TWO)

- A. Add the signature to an HTTP header that is named Authorization
- B. Add the signature to a session cookie
- C. Add the signature to an HTTP header that is named Authentication
- D. Add the signature to a query string parameter that is named X-Amz-Signature
- E. Add the signature to an HTTP header that is named WWW-Authenticate

**Answer:** DE

#### NEW QUESTION 177

- (Exam Topic 3)

A company is launching a polling application. The application will store the results of each poll in an Amazon DynamoDB table. Management wants to remove poll data after a few days and store an archive of those records in Amazon S3.

Which approach would allow the application to archive each poll's data while keeping complexity to a MINIMUM?

- A. Enable Time to Live (TTL) on the DynamoDB table
- B. Enable DynamoDB Streams on the table and store the records removed from the stream in Amazon S3.
- C. Schedule an AWS Lambda function to periodically scan the DynamoDB table
- D. Use the BatchWrite operation to delete the results of a scan. Enable DynamoDB Stream on the table and store the records removed from the stream in Amazon S3.
- E. Enable DynamoDB Streams on the table
- F. Configure the stream as trigger for AWS Lambda
- G. Save records to Amazon S3 when records on the stream are modified.
- H. Enable cross-Region replication on the S3 bucket to achieve the poll data.

**Answer:** C

#### NEW QUESTION 179

- (Exam Topic 3)

A developer has written an application that writes data to Amazon DynamoDB. The DynamoDB table has been configured to use conditional writes. During peak usage times, writes are failing due to a `ConditionalCheckFailedException` error. How can the developer increase the application's reliability when multiple clients are attempting to write to the same record?

- A. Write the data to an Amazon SNS topic.
- B. Increase the amount of write capacity for the table to anticipate short-term spikes or bursts in write operations.
- C. Implement a caching solution, such as DynamoDB Accelerator or Amazon ElastiCache.
- D. Implement error retries and exponential backoff with jitter.

**Answer:** C

#### NEW QUESTION 182

- (Exam Topic 3)

A company wants to migrate an existing web application to AWS. The application consists of two web servers and a MySQL database. The company wants the application to automatically scale in response to demand. The company also wants to reduce its operational overhead for database backups and maintenance. The company needs the ability to deploy multiple versions of the application concurrently. What is the MOST operationally efficient solution that meets these requirements?

- A. Deploy the application to AWS Elastic Beanstalk
- B. Migrate the database to an Amazon RDS Multi-AZ DB instance
- C. Create an Amazon Machine Image (AMI) that contains the application code
- D. Create an Auto Scaling group that is based on the AMI. Integrate the Auto Scaling group with an Application Load Balancer for the web server
- E. Migrate the database to a MySQL instance that runs on an Amazon EC2 instance.
- F. Deploy the application to AWS Elastic Beanstalk
- G. Migrate the database to a MySQL instance that runs on an Amazon EC2 instance.
- H. Create an Amazon Machine Image (AMI) that contains the application code
- I. Create an Auto Scaling group that is based on the AMI
- J. Integrate the Auto Scaling group with an Application Load Balancer for the web server
- K. Migrate the database to an Amazon RDS Multi-AZ DB Instance

**Answer:** B

#### NEW QUESTION 184

- (Exam Topic 3)

A company is launching a new web application in the AWS Cloud. The company's development team is using AWS Elastic Beanstalk for deployment and

maintenance. According to the company's change management process, the development team must evaluate changes for a specific time period before completing the rollout.

Which deployment policy meets this requirement?

- A. Immutable
- B. Rolling
- C. Rolling with additional batch
- D. Traffic splitting

**Answer: A**

**NEW QUESTION 185**

- (Exam Topic 3)

A developer wants to modify the following AWS Cloud Formation template to embed another CloudFormation stack:

```
{
  "AWSTemplateFormatVersion" : "2010-09-09",
  "Resources" : {
    "cfStack" : {
      "Properties" : {
        "TemplateURL" : "https://s3.amazonaws.com/cloudformation-templates/cf.template",
        "Parameters" : {
          "InstanceType" : "t3.small"
        }
      }
    }
  }
}
```

Which syntax should the developer add to the blank line of the CloudFormation template to meet this requirement?

- A. "Mapping" : "AWS::CloudFormation::Stack",
- B. "Type" : "AWS::CloudFormation::NestedStack",
- C. "Type" : "AWS::CloudFormation::Stack",
- D. "Mapping" : "AWS::CloudFormation::NestedStack",

**Answer: A**

**NEW QUESTION 189**

- (Exam Topic 3)

An application is using a custom library to make HTTP calls directly to AWS service endpoints. The application is experiencing transient errors that are causing processes to stop when each error is first encountered. A request has been made to make the application more resilient by adding error retries and exponential backoff.

How should a developer implement the changes with MINIMAL custom code?

- A. Add a Retry-After HTTP header to API requests.
- B. Use the AWS CLI to configure the retry settings in a named profile
- C. Change the custom library to retry on 5xx errors only
- D. Use an AWS SDK and set retry-specific configurations.

**Answer: D**

**NEW QUESTION 190**

- (Exam Topic 3)

A company is developing a new web application in Python. A developer must deploy the application using AWS Elastic Beanstalk from the AWS Management Console. The developer creates an Elastic Beanstalk source bundle to upload using the console.

Which of the following are requirements when creating the source bundle? (Select TWO.)

- A. The source bundle must include the ebextensions.yaml file.
- B. The source bundle must not include a top-level directory.
- C. The source bundle must be compressed with any required dependencies in a top-level parent folder.
- D. The source bundle must be created as a single zip or war file.
- E. The source bundle must be uploaded into Amazon EFS.

**Answer: BD**

**NEW QUESTION 194**

- (Exam Topic 3)

A developer is designing a web application in which new users will use their email addresses to create accounts. Millions of users are expected to sign up. The application will store attributes for each user.

Which AWS service or feature should the developer implement to meet these requirements?

- A. Amazon Cognito user pools
- B. AWS Mobile Hub User File Storage

- C. AWS AppSync
- D. AWS Mobile Hub Cloud Logic

**Answer:** A

#### NEW QUESTION 199

- (Exam Topic 3)

A development team uses AWS Elastic Beanstalk to deploy a Java-based web application. The team wants to ensure that the changes to the source code and the configuration are always deployed on new instances. The team configures the Elastic Beanstalk environment to use immutable updates. However, an error occurs the first time a change is deployed with the new update policy.

What is the MOST likely cause of this issue?

- A. Immutable updates are not supported for Java-based applications.
- B. The account has reached its on-demand instance limit.
- C. Immutable updates are only supported for m4 large and larger instance types.
- D. The developer must also modify the `ebextensions/immutable-updates` config file to enable immutable updates.

**Answer:** A

#### NEW QUESTION 204

- (Exam Topic 3)

A developer is creating AWS CloudFormation templates to manage an application's deployment in Amazon Elastic Container Service (Amazon ECS) through AWS CodeDeploy. The developer wants to automatically deploy new versions of the application to a percentage of users before the new version becomes available for all users.

How should the developer manage the deployment of the new version?

- A. Modify the CloudFormation template to include a Transform section and the AWS `"CodeDeploy::BlueGreen` hook.
- B. Deploy the new version in a new CloudFormation stack. After testing is complete, update the application's DNS records for the new stack.
- C. Run CloudFormation stack updates on the application stack to deploy new application versions when they are available.
- D. Create a nested stack for the new version.
- E. Include a Transform section and the AWS `:CodeDeploy BlueGreen` hook.

**Answer:** B

#### NEW QUESTION 205

- (Exam Topic 3)

A developer is planning to use an Amazon API Gateway and AWS Lambda to provide a REST API. The developer will have three distinct environments to manage: development, test, and production. How should the application be deployed while minimizing the number of resources to manage?

- A. Create a separate API Gateway and separate Lambda function for each environment in the same Region.
- B. Assign a Region for each environment and deploy API Gateway and Lambda to each Region.
- C. Create one API Gateway with multiple stages with one Lambda function with multiple aliases.
- D. Create one API Gateway and one Lambda function, and use a REST parameter to identify the environment.

**Answer:** C

#### NEW QUESTION 209

- (Exam Topic 3)

What is required to trace Lambda-based applications with AWS X-Ray?

- A. Send logs from the Lambda application to an S3 bucket; trigger a Lambda function from that bucket to send data to AWS X-Ray.
- B. Trigger a Lambda function from the application logs in Amazon CloudWatch to submit tracing data to AWS X-Ray.
- C. Use an IAM execution role to give the Lambda function permissions and enable tracing.
- D. Update and add AWS X-Ray daemon code to relevant parts of the Lambda function to set up the trace.

**Answer:** D

#### NEW QUESTION 210

- (Exam Topic 3)

A developer needs to deploy a new version to an AWS Elastic Beanstalk application. How can the developer accomplish this task?

- A. Upload and deploy the new application version in the Elastic Beanstalk console.
- B. Use the `eb init` CLI command to deploy a new version.
- C. Terminate the current Elastic Beanstalk environment and create a new one.
- D. Modify the `ebextensions` folder to add a source option to services.

**Answer:** A

#### NEW QUESTION 215

- (Exam Topic 3)

A developer is writing an AWS Lambda function. The developer wants to log key events that occur during the Lambda function and include a unique identifier to associate the events with a specific function invocation.

Which of the following will help the developer accomplish this objective?

- A. Obtain the request identifier from the Lambda context object. Architect the application to write logs to the console.
- B. Obtain the request identifier from the Lambda event object. Architect the application to write logs to a file.
- C. Obtain the request identifier from the Lambda event object. Architect the application to write logs to the console.

D. Obtain the request identifier from the Lambda context object Architect the application to write logs to a file.

**Answer:** A

**NEW QUESTION 220**

- (Exam Topic 3)

A developer needs to manage AWS infrastructure as code and must be able to deploy multiple identical copies of the infrastructure, stage changes, and revert to previous versions.

Which approach addresses these requirements?

- A. Use cost allocation reports and AWS OpsWorks to deploy and manage the infrastructure.
- B. Use Amazon CloudWatch metrics and alerts along with resource tagging to deploy and manage the infrastructure.
- C. Use AWS Elastic Beanstalk and AWS CodeCommit to deploy and manage the infrastructure.
- D. Use AWS CloudFormation and AWS CodeCommit to deploy and manage the infrastructure.

**Answer:** D

**NEW QUESTION 222**

- (Exam Topic 3)

A developer is creating an application that is based on an AWS Lambda function. The function uses the AWS SDK to read product price data from an Amazon S3 bucket and to write user information to an Amazon Aurora DB instance The Lambda function runs often, up to a few times each minute To meet performance requirements, the developer must minimize the run duration of the Lambda function

Which actions can help the developer increase the performance? (Select TWO)

- A. Initialize SDK clients and database connections outside of the function handler
- B. Read the S3 product price data initially and cache it locally in the /tmp directory
- C. Use environment variables to pass operational parameters to the function.
- D. Use most-restrictive permissions when setting the IAM policies for the Lambda IAM role
- E. Split the code into different Lambda functions to keep the functions smaller

**Answer:** AC

**NEW QUESTION 225**

- (Exam Topic 3)

A developer is building a serverless application using AWS Lambda and must create a REST API using an HTTP GET method What needs to be defined to meet this requirement? (Select TWO )

- A. A Lambda@Edge function
- B. An Amazon API Gateway with a Lambda function
- C. An exposed GET method in an Amazon API Gateway ID.
- D. An exposed GET method in the Lambda function
- E. An exposed GET method in Amazon Route 53

**Answer:** BE

**NEW QUESTION 228**

- (Exam Topic 3)

A developer is attempting to use the Amazon S3 PutObject API operation to upload an object to an S3 bucket that has default encryption enabled. The developer receives a 400 Bad Request error.

What is the MOST likely cause of this error?

- A. The API operation cannot access the encryption key
- B. The HTTP Content-Length header is missing.
- C. The object exceeds the maximum object size that is allowed.
- D. The S3 bucket exceeds the maximum storage capacity that is allowed

**Answer:** D

**NEW QUESTION 232**

- (Exam Topic 3)

A physician's office management application requires that all data in transit between an EC2 instance and an Amazon EBS volume be encrypted Which of the following techniques fulfills this requirement? (Select TWO )

- A. Create encrypted snapshots into Amazon S3
- B. Use Amazon RDS with encryption
- C. Use IAM roles to limit access to the Amazon EBS volume
- D. Enable EBS encryption
- E. Leverage OS-level encryption

**Answer:** AD

**NEW QUESTION 236**

- (Exam Topic 3)

A developer is deploying an application in the AWS Cloud by using AWS Cloud Formation The application will connect to an existing Amazon RDS database The hostname of the RDS database is stored in AWS Systems Manager Parameter Store as a plaintext value The developer needs to incorporate the database hostname into the Cloud Formation template to initialize the application when the stack is created

How should the developer reference the parameter that contains the database hostname?

- A. Use the ssm dynamic reference
- B. Use the Ref intrinsic function
- C. Use the Fn: ImportVatue intrinsic function
- D. Use the ssm-secure dynamic reference.

**Answer: C**

#### NEW QUESTION 240

- (Exam Topic 3)

A company is providing read access to objects in an Amazon S3 bucket for different customers. The company uses IAM permissions to restrict access to the S3 bucket. The customers can access only their own files.

Due to a regulation requirement, the company needs to enforce encryption in transit for interactions with Amazon S3.

Which solution will meet these requirements?

- A. Add a bucket policy to the S3 bucket to deny S3 actions when the aws:SecureTransport condition is equal to false.
- B. Add a bucket policy to the S3 bucket to deny S3 actions when the s3:x-amz-aci condition is equal to public-read.
- C. Add an IAM policy to the IAM users to enforce the usage of the AWS SDK.
- D. Add an IAM policy to the IAM users that allows S3 actions when the s3:x-arnz-acl condition is equal to bucket-owner-read.

**Answer: C**

#### NEW QUESTION 245

- (Exam Topic 3)

A developer is trying to get data from an Amazon DynamoDB table called demoman-table. The developer configured the AWS CLI to use a specific IAM user's credentials and executed the following command:

```
aws dynamodb get-item table-name demoman-table --key '{"id": <"N"; "1993"}'
```

 The command returned errors and no rows were returned.

What is the MOST likely cause of these issues?

- A. The command is incorrect; it should be rewritten to use `aws dynamodb get-item` with a string argument.
- B. The developer needs to log a ticket with AWS Support to enable access to the demoman-table.
- C. Amazon DynamoDB cannot be accessed from the AWS CLI and needs to be called via the REST API.
- D. The IAM user needs an associated policy with read access to demoman-table.

**Answer: D**

#### NEW QUESTION 249

- (Exam Topic 3)

A developer is working on a serverless application that needs to process any changes to an Amazon DynamoDB table with an AWS Lambda function.

How should the developer configure the Lambda function to detect changes to the DynamoDB table?

- A. Create an Amazon Kinesis data stream, and attach it to the DynamoDB table. Create a trigger to connect the data stream to the Lambda function.
- B. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke the Lambda function on a regular schedule. Connect to the DynamoDB table from the Lambda function to detect changes.
- C. Enable DynamoDB Streams on the table. Create a trigger to connect the DynamoDB stream to the Lambda function.
- D. Create an Amazon Kinesis Data Firehose delivery stream, and attach it to the DynamoDB table. Configure the delivery stream destination as the Lambda function.

**Answer: C**

#### NEW QUESTION 253

- (Exam Topic 3)

A developer wants to use React to build a web and mobile application. The application will be hosted on AWS. The application must authenticate users and then allow users to store and retrieve files that they own. The developer wants to use Facebook for authentication.

Which CLI will MOST accelerate the development and deployment of this application on AWS?

- A. AWS CLI
- B. AWS Amplify CLI
- C. AWS Serverless Application Model (AWS SAM) CLI
- D. Amazon Elastic Container Service (Amazon ECS) CLI

**Answer: B**

#### NEW QUESTION 255

- (Exam Topic 3)

A developer is working with a Docker application that needs to be quickly deployed using AWS without changing the infrastructure or configuring health checks.

The application should be configured so that changes and updates can be made automatically without any downtime.

Which solution will meet these requirements?

- A. Use AWS Elastic Beanstalk for application deployment and select an all-at-once update policy.
- B. Use AWS Elastic Beanstalk for application deployment and select a rolling deployment policy.
- C. Deploy the Docker container on an Amazon EC2 instance in an Auto Scaling group and configure a health check on the EC2 instance.
- D. Deploy the Docker container using AWS Lambda and enable Amazon CloudWatch monitoring.

**Answer: A**

#### NEW QUESTION 258

- (Exam Topic 3)

An AWS Lambda function accesses two Amazon DynamoDB tables. A developer wants to improve the performance of the Lambda function by identifying bottlenecks in the function. How can the developer inspect the timing of the DynamoDB API calls?

- A. Add DynamoDB as an event source to the Lambda function
- B. View the performance with Amazon CloudWatch metrics
- C. Place an Application Load Balancer (ALB) in front of the two DynamoDB tables
- D. Inspect the ALB logs
- E. Limit Lambda to no more than five concurrent invocations Monitor from the Lambda console
- F. Enable AWS X-Ray tracing for the function
- G. View the traces from the X-Ray service.

**Answer: D**

#### NEW QUESTION 261

- (Exam Topic 3)

A three-tier application hosted on AWS uses Amazon RDS for MySQL as its database. A developer must ensure the database credentials are stored and accessed securely.

What is the MOST secure way for the developer to achieve this?

- A. Store the credentials in a configuration file and commit it to the GIT repository.
- B. Store the credentials in AWS Secrets Manager and enable automatic secret rotation.
- C. Store the credentials using Amazon RDS and enable automatic rotation
- D. Store the credentials in code and handle credentials rotation within the application.

**Answer: A**

#### NEW QUESTION 263

- (Exam Topic 3)

A developer is changing the configuration for a CPU-intensive AWS Lambda function that runs once an hour. The function usually takes 45 seconds to run, but sometimes the run time is up to 1 minute. The timeout parameter is set to 3 minutes, and all other parameters are set to default.

The developer needs to optimize the run time of this function. Which solution will meet this requirement?

- A. Redeploy the function within the default VPC
- B. Increase the function's memory.
- C. Redeploy the function by using Lambda layers
- D. Increase the function's reserved concurrency

**Answer: B**

#### NEW QUESTION 265

- (Exam Topic 3)

Which of the following are good use cases for how Amazon ElastiCache can help an application? (Select TWO.)

- A. Improve the performance of S3 PUT operations
- B. Improve the latency of deployments performed by AWS CodeDeploy
- C. Improve latency and throughput for read-heavy application workloads.
- D. Reduce the time required to merge AWS CodeCommit branches
- E. Improve performance of compute-intensive applications.

**Answer: CE**

#### NEW QUESTION 269

- (Exam Topic 3)

A development team uses AWS Elastic Beanstalk for application deployment. The team has configured the application version lifecycle policy to limit the number of application versions to 25. However, even with the lifecycle policy, the source bundle is deleted from the Amazon S3 source bucket.

What should a developer do in the Elastic Beanstalk application version lifecycle settings to retain the source code in the S3 bucket?

- A. Change the Set the application versions limit by total count setting to zero.
- B. Disable the Lifecycle policy setting
- C. Change the Set the application version limit by age setting to zero.
- D. Set Retention to Retain source bundle in S3.

**Answer: C**

#### NEW QUESTION 274

- (Exam Topic 3)

A developer is building an application that processes a stream of user-supplied data. The data stream must be consumed by multiple Amazon EC2 based processing applications in parallel and in real time. Each processor must be able to resume without losing data if there is a service interruption. The Application Architect plans to add other processors in the near future, and wants to minimize the amount of data duplication involved.

Which solution will satisfy these requirements?

- A. Publish the data to Amazon SQS
- B. Publish the data to Amazon Kinesis Data Firehose
- C. Publish the data to Amazon CloudWatch Events.
- D. Publish the data to Amazon Kinesis Data Streams.

**Answer: A**

#### NEW QUESTION 275

- (Exam Topic 3)

A company is running a custom application on a set of on-premises Linux servers that are accessed using Amazon API Gateway. AWS X-Ray tracing has been enabled on the API test stage

How can a developer enable X-Ray tracing on the on-premises servers with the LEAST amount of configuration"

- A. Install and run the X-Ray SDK on the on-premises servers to capture and relay the data to the X-Ray service.
- B. Install and run the X-Ray daemon on the on-premises servers to capture and relay the data to the X-Ray service
- C. Capture incoming requests on-premises and configure an AWS Lambda function to pull, process, and relay relevant data to X-Ray using the PutTraceSegments API call
- D. Capture incoming requests on-premises and configure an AWS Lambda function to pull, process, and relay relevant data to X-Ray using the PutTelemetryRecords API call.

**Answer: B**

#### NEW QUESTION 276

- (Exam Topic 3)

When using the AWS Encryption SDK now does the developer keep track of the data encryption keys used to encrypt data?

- A. The developer must manually Keep track of the data encryption keys used for each data object
- B. The SDK encrypts the data encryption key and stores it (encrypted) as part of the returned ciphertext
- C. The SDK stores the data encryption keys automatically in Amazon S3
- D. The data encryption key is stored in the userdata for the EC2 instance

**Answer: B**

#### NEW QUESTION 281

- (Exam Topic 3)

A developer is working on a serverless project based in Java. Initial testing shows a cold start takes about 8 seconds on average for AWS Lambda functions.

What should the developer do to reduce the cold start time" (Select TWO)

- A. Add the Spring Framework to the project and enable dependency injection
- B. Reduce the deployment package by including only the needed modules from the AWS SDK for Java.
- C. Increase the memory allocation setting for the Lambda function.
- D. Increase the timeout setting for the Lambda function.
- E. Change the Lambda invocation mode from synchronous to asynchronous.

**Answer: BC**

#### NEW QUESTION 284

- (Exam Topic 2)

A Developer is migrating an on-premises application to AWS. The application currently takes user uploads and saves them to a local directory on the server. All uploads must be saved and made immediately available to all instances in an Auto scaling group.

Which approach will meet these requirements?

- A. Use Amazon EBS and configure the application AMI to use a snapshot of the same EBS instance on boot.
- B. Use Amazon S3 and rearchitect the application so all uploads are placed in S3.
- C. Use instance storage and share it between instances launched from the same Amazon machine image (AMI).
- D. Use Amazon EBS and file synchronization software to achieve eventual consistency among the auto scaling group.

**Answer: B**

#### Explanation:

Use Amazon S3 and rearchitect the application so all uploads are placed in S3. Even though you could do EBS attachment to ASG launch config userdata for ec2 instances going to serve , But you need to select the ASG in single AZ where your EBS is located otherwise it will not work since EBS is AZ locked.

#### NEW QUESTION 287

- (Exam Topic 2)

A company has an AWS CloudFormation template that is stored as a single file. The template is able to launch and create a full infrastructure stack.

Which best practice would increase the maintainability of the template?

- A. Use nested stacks for common template patterns.
- B. Embed credentials to prevent typos.
- C. Remove mappings to decrease the number of variables.
- D. Use AWS::Include to reference publicly-hosted template files.

**Answer: A**

#### NEW QUESTION 291

- (Exam Topic 2)

AWS CodeBuild builds code for an application, creates the Docker image, pushes the image to Amazon Elastic Container Registry (Amazon ECR), and tags the image with a unique identifier.

If the Developers already have AWS CLI configured on their workstations, how can the Docker images be pulled to the workstations?

- A. Run the following:docker pull REPOSITORY URI : TAG
- B. Run the output of the following:aws ecr get-loginand then run:docker pull REPOSITORY URI : TAG
- C. Run the following:aws ecr get-loginand then run:docker pull REPOSITORY URI : TAG
- D. Run the output of the following:aws ecr get-download-url-for-layerand then run:docker pull REPOSITORY URI : TAG

**Answer:** B

**Explanation:**

<https://docs.aws.amazon.com/cli/latest/reference/ecr/get-login.html>

**NEW QUESTION 294**

- (Exam Topic 2)

A Development team would like to migrate their existing application code from a GitHub repository to AWS CodeCommit. What needs to be created before they can migrate a cloned repository to CodeCommit over HTTPS?

- A. A GitHub secure authentication token
- B. A public and private SSH key file
- C. A set of Git credentials generated from IAM
- D. An Amazon EC2 IAM role with CodeCommit permissions

**Answer:** C

**Explanation:**

<https://docs.aws.amazon.com/codecommit/latest/userguide/how-to-migrate-repository-existing.html>

**NEW QUESTION 295**

- (Exam Topic 2)

A Developer is storing sensitive documents in Amazon S3 that will require encryption at rest. The encryption keys must be rotated annually, at least. What is the easiest way to achieve this?

- A. Encrypt the data before sending it to Amazon S3
- B. Import a custom key into AWS KMS with annual rotation enabled
- C. Use AWS KMS with automatic key rotation
- D. Export a key from AWS KMS to encrypt the data

**Answer:** C

**Explanation:**

<https://docs.aws.amazon.com/kms/latest/developerguide/rotate-keys.html> <https://docs.aws.amazon.com/kms/latest/developerguide/custom-key-store-overview.html>

You can use the same techniques to view and manage the CMKs in your custom key store that you use for CMKs in the AWS KMS key store. You can control access with IAM and key policies, create tags and aliases, enable and disable the CMKs, and schedule key deletion. You can use the CMKs for cryptographic operations and use them with AWS services that integrate with AWS KMS. However, you cannot enable automatic key rotation and you cannot import key material into a CMK in a custom key store.

Q: Can I rotate my keys? Yes. You can choose to have AWS KMS automatically rotate CMKs every year, provided that those keys were generated within AWS KMS HSMs. Automatic key rotation is not supported for imported keys, asymmetric keys, or keys generated in an AWS CloudHSM cluster using the AWS KMS custom key store feature. If you choose to import keys to AWS KMS or asymmetric keys or use a custom key store, you can manually rotate them by creating a new CMK and mapping an existing key alias from the old CMK to the new CMK. <https://aws.amazon.com/kms/faqs/>

**NEW QUESTION 298**

- (Exam Topic 2)

A Developer is publishing critical log data to a log group in Amazon CloudWatch Logs, which was created 2 months ago. The Developer must encrypt the log data using an AWS KMS customer master key (CMK) so future data can be encrypted to comply with the company's security policy. How can the Developer meet this requirement?

- A. Use the Cloud Watch Logs console and enable the encrypt feature on the log group.
- B. Use the AWS CLI create-log-group command and specify the key Amazon Resource Name (ARN)
- C. Use the KMS console and associate the CMK with the log group
- D. Use the AWS CLI associate-kms-key command and specify the key Amazon Resource Name (ARN)

**Answer:** C

**NEW QUESTION 302**

- (Exam Topic 2)

Queries to an Amazon DynamoDB table are consuming a large amount of read capacity. The table has a significant number of large attributes. The application does not need all of the attribute data.

How can DynamoDB costs be minimized while maximizing application performance?

- A. Batch all the writes, and perform the write operations when no or few reads are being performed.
- B. Create a global secondary index with a minimum set of projected attributes.
- C. Implement exponential backoffs in the application.
- D. Load balance the reads to the table using an Application Load Balancer.

**Answer:** C

**Explanation:**

<https://docs.aws.amazon.com/AWSEC2/latest/APIReference/query-api-troubleshooting.html>

**NEW QUESTION 307**

- (Exam Topic 2)

A developer is storing sensitive data generated by an application in Amazon S3. The developer wants to encrypt the data at rest. A company policy requires an audit trail of when the master key was used and by whom.

Which encryption option will meet these requirements?

- A. Server-side encryption with Amazon S3 managed keys (SSE-S3)
- B. Server-side encryption with AWS KMS managed keys (SSE-KMS)
- C. Server-side encryption with customer-provided keys (SSE-C)
- D. Server-side encryption with self-managed keys

**Answer: B**

#### NEW QUESTION 311

- (Exam Topic 2)

A developer has built an application running on AWS Lambda using AWS Serverless Application Model (AWS SAM). What is the correct order of execution to successfully deploy the application?

- A. \* 1 Build the SAM template in Amazon EC2\* 2 Package the SAM template to Amazon EBS storage\* 3. Deploy the SAM template from Amazon EBS.
- B. \* 1 Build the SAM template locally\* 2 Package the SAM template onto Amazon S3\* 3. Deploy the SAM template from Amazon S3.
- C. \* 1 Build the SAM template locally\* 2. Deploy the SAM template from Amazon S3. \* 3 Package the SAM template for use
- D. \* 1 Build the SAM template locally\* 2 Package the SAM template from AWS CodeCommit
- E. \* 3 Deploy the SAM template to CodeCommit

**Answer: B**

#### Explanation:

Reference:

<https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/serverlessdeploying.html>

#### NEW QUESTION 315

- (Exam Topic 2)

A company is running a Docker application on Amazon ECS. The application must scale based on user load in the last 15 seconds. How should a Developer instrument the code so that the requirement can be met?

- A. Create a high-resolution custom Amazon CloudWatch metric for user activity data, then publish data every 30 seconds
- B. Create a high-resolution custom Amazon CloudWatch metric for user activity data, then publish data every 5 seconds
- C. Create a standard-resolution custom Amazon CloudWatch metric for user activity data, then publish data every 30 seconds
- D. Create a standard-resolution custom Amazon CloudWatch metric for user activity data, then publish data every 5 seconds

**Answer: B**

#### Explanation:

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/publishingMetrics.html#high-resolution-m>

#### NEW QUESTION 317

- (Exam Topic 2)

An application is being developed to audit several AWS accounts. The application will run in Account A and must access AWS services in Accounts B and C. What is the MOST secure way to allow the application to call AWS services in each audited account?

- A. Configure cross-account roles in each audited account
- B. Write code in Account A that assumes those roles
- C. Use S3 cross-region replication to communicate among accounts, with Amazon S3 event notifications to trigger Lambda functions
- D. Deploy an application in each audited account with its own role
- E. Have Account A authenticate with the application
- F. Create an IAM user with an access key in each audited account
- G. Write code in Account A that uses those access keys

**Answer: A**

#### Explanation:

[https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial\\_cross-account-with-roles.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html)

#### NEW QUESTION 322

- (Exam Topic 2)

A development team wants to run their container workloads on Amazon ECS. Each application container needs to share data with another container to collect logs and metrics.

What should the development team do to meet these requirements?

- A. Create two pod specifications. Make one to include the application container and the other to include the other container. Link the two pods together.
- B. Create two task definitions. Make one to include the application container and the other to include the other container.
- C. Mount a shared volume between the two tasks.
- D. Create one task definition. Specify both containers in the definition. Mount a shared volume between those two containers.
- E. Create a single pod specification. Include both containers in the specification. Mount a persistent volume to both containers.

**Answer: C**

#### NEW QUESTION 325

- (Exam Topic 2)

A Developer must trigger an AWS Lambda function based on the item lifecycle activity in an Amazon DynamoDB table.

How can the Developer create the solution?

- A. Enable a DynamoDB stream that publishes an Amazon SNS message.
- B. Trigger the Lambda function synchronously from the SNS message.
- C. Enable a DynamoDB stream that publishes an SNS message.
- D. Trigger the Lambda function asynchronously from the SNS message.
- E. Enable a DynamoDB stream, and trigger the Lambda function synchronously from the stream.
- F. Enable a DynamoDB stream, and trigger the Lambda function asynchronously from the stream.

**Answer:** C

**Explanation:**

<https://docs.aws.amazon.com/lambda/latest/dg/with-ddb.html>

**NEW QUESTION 327**

- (Exam Topic 2)

A Developer is creating a Lambda function that will generate and export a file. The function requires 100 MB of temporary storage for temporary files while executing. These files will not be needed after the function is complete.

How can the Developer MOST efficiently handle the temporary files?

- A. Store the files in EBS and delete the files at the end of the Lambda function.
- B. Copy the files to EFS and delete the files at the end of the Lambda function.
- C. Store the files in the /tmp directory and delete the files at the end of the Lambda function.
- D. Copy the files to an S3 bucket with a lifecycle policy to delete the files.

**Answer:** C

**NEW QUESTION 331**

- (Exam Topic 2)

After installing the AWS CLI, a Developer tries to run the command `aws configure` but receives the following error:

Error: aws: command not found

What is the most likely cause of this error?

- A. The aws executable is not in the PATH environment variable.
- B. Access to the aws executable has been denied to the installer.
- C. Incorrect AWS credentials were provided.
- D. The aws script does not have an executable file mode.

**Answer:** A

**Explanation:**

<https://docs.aws.amazon.com/cli/latest/userguide/cli-chap-troubleshooting.html>

**NEW QUESTION 333**

- (Exam Topic 2)

A company needs to distribute firmware updates to its customers around the world.

Which service will allow easy and secure control of the access to the downloads at the lowest cost?

- A. Use Amazon CloudFront with signed URLs for Amazon S3
- B. Create a dedicated Amazon CloudFront Distribution for each customer
- C. Use Amazon CloudFront with AWS Lambda@Edge
- D. Use Amazon API Gateway and AWS Lambda to control access to an S3 bucket

**Answer:** A

**Explanation:**

<https://aws.amazon.com/blogs/networking-and-content-delivery/amazon-s3-amazon-cloudfront-a-match-made-i>

**NEW QUESTION 335**

- (Exam Topic 2)

A developer has created a new AWS IAM user that has `s3:putobject` permission to write to a specific Amazon bucket. This S3 bucket uses server-side encryption with AWS KMS managed keys (SEE-KMS) as the encryption. Using the access key and secret key of the IAM user, the application received an access denied error when calling the `PutObject` API.

How can this issue be resolved?

- A. Update the policy of the IAM user to allow the `s3:Encrypt` action.
- B. Update the bucket policy of the S3 bucket to allow the IAM user to upload objects
- C. Update the policy of the IAM user to allow the `kms:GenerateDatakey` action
- D. Update the ACL of the bucket to allow the IAM user to upload objects

**Answer:** C

**NEW QUESTION 336**

- (Exam Topic 2)

A developer is writing a web application that must share secure documents with end users. The documents are stored in a private Amazon S3 bucket. The application must allow only authenticated users to download specific documents when requested, and only for a duration of 15 minutes.

How can the developer meet these requirements?

- A. Copy the documents to a separate S3 bucket that has a lifecycle policy for deletion after 15 minutes
- B. Create a presigned S3 URL using the AWS SDK with an expiration time of 15 minutes

- C. Use server-side encryption with AWS KMS managed keys (SSE-KMS) and download the documents using HTTPS
- D. Modify the S3 bucket policy to only allow specific users to download the documents Revert the change after 15 minutes.

**Answer:** B

#### NEW QUESTION 339

- (Exam Topic 2)

A company has implemented AWS CodePipeline to automate its release pipelines The development team is writing an AWS Lambda function that will send notifications for state changes of each of the actions in the stages.

Which steps must be taken to associate the Lambda function with the event source?

- A. Create a trigger that invokes the Lambda function from the Lambda console by selecting CodePipeline as the event source
- B. Create an event trigger and specify the Lambda function from the CodePipeline console.
- C. Create an Amazon CloudWatch alarm that monitors status changes in CodePipeline and triggers the Lambda function
- D. Create an Amazon CloudWatch Events rule that uses CodePipeline as an event source.

**Answer:** B

#### NEW QUESTION 341

- (Exam Topic 2)

A company is running an application built on AWS Lambda functions. One Lambda function has performance issues when it has to download a 50MB file from the Internet in every execution. This function is called multiple times a second.

What solution would give the BEST performance increase?

- A. Cache the file in the /tmp directory
- B. Increase the Lambda maximum execution time
- C. Put an Elastic Load Balancer in front of the Lambda function
- D. Cache the file in Amazon S3

**Answer:** A

#### Explanation:

<https://docs.aws.amazon.com/lambda/latest/dg/runtimes-context.html>

#### NEW QUESTION 343

- (Exam Topic 2)

A developer implemented a static website hosted in amazon s3 that makes web service requests in amazon api gateway and aws lambda. The site is showing an error that reads.

"No 'access-control-allow-origin' header is present on the requested resource. Origin 'null is therefore not allowed access "

What should the developer do to resolve this issue?

- A. Enable cross-origin resource sharing (cors) on the s3 bucket
- B. Enable cross-origin resource sharing (cors) for the method in api gateway
- C. Add the access-control-request-method header to the request
- D. Add the access-control-inquest headers header to the request

**Answer:** A

#### Explanation:

Reference: <https://forums.aws.amazon.com/thread.jspa?threadid=252972>

#### NEW QUESTION 347

- (Exam Topic 2)

A developer is creating a new application that will be accessed by users through an API created using Amazon API Gateway The users need to be authenticated by a third-party Security Assertion Markup Language (SAML) identity provider Once authenticated, users will need access to other AWS services such as Amazon S3 and Amazon DynamoDB

How can these requirements be met?

- A. Use an Amazon Cognito user pool with SAML as the resource server
- B. Use Amazon Cognito Identity pools with a SAML identity provider as one of the authentication providers
- C. Use the AWS IAM service to provide the sign-up and sign-in functionality.
- D. Use Amazon CloudFront signed URLs to connect with the SAML identity provider

**Answer:** B

#### NEW QUESTION 349

- (Exam Topic 2)

A developer is migrating code to an AWS Lambda function that will access an Amazon Aurora MySQL database.

What is the MOST secure way to authenticate the function to the database?

- A. Store the database credentials as encrypted parameters in AWS Systems Manager Parameter Store Obtain the credentials from Systems Manager when the Lambda function needs to connect to the database
- B. Store the database credentials in AWS Secrets Manager Let Secrets Manager handle the rotation of the credentials, as required
- C. Store the database credentials in an Amazon S3 bucket that has a restrictive bucket policy for the Lambda role only when accessing the credentials Use AWS KMS to encrypt the data
- D. Create a policy with rds-db connect access to the database and attach it to the role assigned to the Lambda function

**Answer:** B

#### NEW QUESTION 351

- (Exam Topic 2)

A developer has written an Amazon Kinesis Data Streams application. As usage grows and traffic over time, the application is regularly receiving `ProvisionedThroughputExceededException` error messages.

Which steps should the Developer take to resolve the error? (Select Two.)

- A. Use Auto scaling to scale the stream for better performance.
- B. Increase the delay between the `GetRecords` call and the `PutRecords` call.
- C. Increase the number of shards in the data stream.
- D. Specify a shard iterator using the `shardIterator` parameter.
- E. Implement exponential backoff on the `GetRecords` call and the `PutRecords` call.

**Answer:** BD

#### Explanation:

Reference: <https://docs.aws.amazon.com/streams/latest/dev/troubleshooting-consumers.html>

#### NEW QUESTION 353

- (Exam Topic 2)

A company is developing an application that will be accessed through the Amazon API Gateway REST API. Registered users should be the only ones who can access certain resources of this API. The token being used should expire automatically and needs to be refreshed periodically.

How can a developer meet these requirements?

- A. Create an Amazon Cognito identity pool, configure the Amazon Cognito Authorizer in API Gateway, and use the temporary credentials generated by the identity pool.
- B. Create and maintain a database record for each user with a corresponding token and use an AWS Lambda authorizer in API Gateway.
- C. Create an Amazon Cognito user pool, configure the Cognito Authorizer in API Gateway, and use the identity or access token.
- D. Create an IAM user for each API user, attach an inline permissions policy to the AP
- E. and use an IAM authorizer in API Gateway.

**Answer:** C

#### Explanation:

Reference: <https://aws.amazon.com/premiumsupport/knowledge-center/cognito-custom-scopes-api-gateway/>

#### NEW QUESTION 355

- (Exam Topic 2)

An application running on an Amazon Linux EC2 instance needs to manage the AWS infrastructure. How can the EC2 instance be configured to make AWS API calls securely?

- A. Sign the AWS CLI command using the signature version 4 process.
- B. Run the `aws configure` AWS CLI command and specify the access key id and secret access key.
- C. Specify a role for the EC2 instance with the necessary privileges.
- D. Pass the access key id and secret access key as parameters for each AWS CLI command.

**Answer:** C

#### NEW QUESTION 357

- (Exam Topic 2)

A developer needs temporary access to resources in a second account. What is the MOST secure way to achieve this?

- A. Use the Amazon Cognito user pools to get short-lived credentials for the second account.
- B. Create a dedicated IAM access key for the second account, and send it by mail.
- C. Create a cross-account access role, and use `sts:AssumeRole` API to get short-lived credentials.
- D. Establish trust, and add an SSH key for the second account to the IAM user.

**Answer:** C

#### Explanation:

Reference: [https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial\\_cross-account-with-roles.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html)

#### NEW QUESTION 358

- (Exam Topic 2)

A company caches session information for a web application in an Amazon DynamoDB table. The company wants an automated way to delete old items from the table.

What is the simplest way to do this?

- A. Write a script that deletes old records; schedule the scripts as a cron job on an Amazon EC2 instance.
- B. Add an attribute with the expiration time; enable the Time To Live feature based on that attribute.
- C. Each day, create a new table to hold session data; delete the previous day's table.
- D. Add an attribute with the expiration time; name the attribute `ItemExpiration`.

**Answer:** B

#### Explanation:

Reference: <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/time-to-live-ttl-how-to.html>

### NEW QUESTION 361

- (Exam Topic 2)

A company runs continuous integration/continuous delivery (CI/CD) pipeline for its application on AWS CodePipeline. A developer must write unit tests and run them as part of the pipelines before staging the artifacts for testing.

How should the Developer incorporate unit tests as part of CI/CD pipeline?

- A. Create a separate codePipeline pipeline to run unit tests.
- B. Update the AWS codeBuild build specification to include a phase for running unit tests.
- C. Install the AWS CodeDeploy agent on an Amazon EC2 instance to run unit tests.
- D. Create a testing branch in AWS CodeCommit to run unit tests.

**Answer: B**

### NEW QUESTION 366

- (Exam Topic 2)

A Developer has been asked to create an AWS Lambda function that is triggered any time updates are made to items in an Amazon DynamoDB table. The function has been created, and appropriate permissions have been added to the Lambda execution role. Amazon DynamoDB streams have been enabled for the table, but the function is still not being triggered.

Which option would enable DynamoDB table updates to trigger the Lambda function?

- A. Change the StreamViewType parameter value to NEW\_AND\_OLD\_IMAGES for the DynamoDB table
- B. Configure event source mapping for the Lambda function
- C. Map an Amazon SNS topic to the DynamoDB streams
- D. increase the maximum execution time (timeout) setting of the Lambda function

**Answer: B**

#### Explanation:

[https://docs.aws.amazon.com/en\\_us/amazondynamodb/latest/developerguide/Streams.Lambda.Tutorial.html](https://docs.aws.amazon.com/en_us/amazondynamodb/latest/developerguide/Streams.Lambda.Tutorial.html) Create an event source mapping to tell Lambda to send records from your stream to a Lambda function. You can create multiple event source mappings to process the same data with multiple Lambda functions, or process items from multiple streams with a single function.

### NEW QUESTION 371

- (Exam Topic 2)

A team of Developers must migrate an application running inside an AWS Elastic Beanstalk environment from a Classic Load Balancer to an Application Load Balancer.

Which steps should be taken to accomplish the task using the AWS Management Console?

- A. \*1. Update the application code in the existing deployment.\* 2. Select a new load balancer type before running the deployment.\* 3. Deploy the new version of the application code to the environment.
- B. \*1. Create a new environment with the same configurations except for the load balancer type.\* 2. Deploy the same application version as used in the original environment.\* 3. Run the swap-environment-cnames action.
- C. \*1. Clone the existing environment, changing the associated load balancer type.\*2. Deploy the same application version as used in the original environment.\*3. Run the swap-environment-cnames action.
- D. \*1. Edit the environment definitions in the existing deployment.\*2. Change the associated load balancer type according to the requirements.\*3. Rebuild the environment with the new load balancer type.

**Answer: B**

#### Explanation:

<https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/using-features.managing.elb.html>

By default, Elastic Beanstalk creates an Application Load Balancer for your environment when you enable load balancing with the Elastic Beanstalk console or the EB CLI. It configures the load balancer to listen for HTTP traffic on port 80 and forward this traffic to instances on the same port. You can choose the type of load balancer that your environment uses only during environment creation. Later, you can change settings to manage the behavior of your running environment's load balancer, but you can't change its type.

### NEW QUESTION 375

- (Exam Topic 2)

A developer uses Amazon S3 buckets for static website hosting. The developer creates one S3 bucket for the code and another S3 bucket for the assets, such as image and video files. Access is denied when a user attempts to access the assets bucket from the code bucket, with the website application showing a 403 error. How should the developer solve this issue?

- A. Create an IAM role and apply it to the assets bucket for the code bucket to be granted access
- B. Edit the bucket policy of the assets bucket to open access to all principals
- C. Edit the cross-origin resource sharing (CORS) configuration of the assets bucket to allow any origin to access the assets
- D. Change the code bucket to use AWS Lambda functions instead of static website hosting.

**Answer: C**

### NEW QUESTION 378

- (Exam Topic 2)

An application ingests a large number of small messages and stores them in a database. The application uses AWS Lambda. A development team is making changes to the application's processing logic. In testing, it is taking more than 15 minutes to process each message. The team is concerned the current backend may time out.

Which changes should be made to the backend system to ensure each message is processed in the MOST scalable way?

- A. Add the messages to an Amazon SQS queue Set up an Amazon EC2 instance to poll the queue and process messages as they arrive.
- B. Add the messages to an Amazon SQS queue
- C. Set up Amazon EC2 instances in an Auto Scaling group to poll the queue and process the messages as they arrive.

- D. Create a support ticket to increase the Lambda timeout to 60 minutes to allow for increased processing time
- E. Change the application to directly insert the body of the message into an Amazon RDS database.

**Answer:** A

#### NEW QUESTION 380

- (Exam Topic 2)

An application runs on multiple EC2 instances behind an ELB.

Where is the session data best written so that it can be served reliably across multiple requests?

- A. Write data to Amazon ElastiCache
- B. Write data to Amazon Elastic Block Store.
- C. Write data to Amazon EC2 Instance Store.
- D. Write data to the root filesystem.

**Answer:** C

#### Explanation:

Reference:<https://docs.aws.amazon.com/aws-technical-content/latest/microservices-on-aws/microservices-on-aw>

#### NEW QUESTION 385

- (Exam Topic 2)

A developer is testing an application that invokes an AWS Lambda function asynchronously. During the testing phase, the Lambda function fails to process after two retries.

How can the developer troubleshoot the failure?

- A. Configure AWS CloudTrail logging to investigate the invocation failures.
- B. Configure Dead Letter Queues by sending events to Amazon SQS for investigation.
- C. Configure Amazon Simple Workflow Service to process any direct unprocessed events.
- D. Configure AWS Config to process any direct unprocessed events.

**Answer:** A

#### NEW QUESTION 388

- (Exam Topic 2)

A Developer is writing a REST service that will add items to a shopping list. The service is built on Amazon API Gateway with AWS Lambda integrations. The shopping list items are sent as query string parameters in the method request.

How should the Developer convert the query string parameters to arguments for the Lambda function?

- A. Enable request validation
- B. Include the Amazon Resource Name (ARN) of the Lambda function
- C. Change the integration type
- D. Create a mapping template

**Answer:** D

#### Explanation:

<https://docs.aws.amazon.com/apigateway/latest/developerguide/integrating-api-with-aws-services-lambda.html#>

#### NEW QUESTION 389

- (Exam Topic 2)

A development team is working on a mobile app that allows users to upload pictures to Amazon S3. The team expects the app will be used by hundreds of thousands of users during a single event simultaneously. Once the pictures are uploaded, the backend service will scan and parse the pictures for inappropriate content.

Which approach is the MOST resilient way to achieve this goal which also smooths out temporary volume spikes for the backend service?

- A. Develop an AWS Lambda function to check the upload folder in the S3 bucket.
- B. If new uploaded pictures are detected, the Lambda function will scan and parse them.
- C. Once a picture is uploaded to Amazon S3, publish the event to an Amazon SQS queue.
- D. Use the queue as an event source to trigger an AWS Lambda function. In the Lambda function, scan and parse the picture.
- E. When the user uploads a picture, invoke an API hosted in Amazon API Gateway.
- F. The API will invoke an AWS Lambda function to scan and parse the picture.
- G. Create a state machine in AWS Step Functions to check the upload folder in the S3 bucket.
- H. If a new picture is detected, invoke an AWS Lambda function to scan and parse it.

**Answer:** B

#### NEW QUESTION 390

.....

## **Thank You for Trying Our Product**

### **We offer two products:**

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

### **DVA-C01 Practice Exam Features:**

- \* DVA-C01 Questions and Answers Updated Frequently
- \* DVA-C01 Practice Questions Verified by Expert Senior Certified Staff
- \* DVA-C01 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- \* DVA-C01 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The DVA-C01 Practice Test Here](#)**