

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 company has an application that runs on AWS Elastic Beanstalk in a load-balanced environment. The company needs to update the instance types in the environment to a more recent generation of instance types. The company must minimize downtime during the deployment of this configuration change.

Which deployment options will meet these requirements? (Choose two.)

- A. Disabled
- B. Rolling based on Health
- C. Immutable
- D. All at once
- E. Canary

Answer: BC

Explanation:

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

NEW QUESTION 3

- (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 4

- (Exam Topic 4)

A developer runs an application that uses an Amazon API Gateway REST API. The developer needs to implement a solution to proactively monitor the health of both API responses and latencies in case a deployment causes a service disruption despite passing deployment pipeline tests. The solution also must check for endpoint vulnerability and unauthorized changes to APIs, URLs, and website content.

Which solution will meet these requirements?

- A. Use the Amazon CloudWatch Synthetics canary functionality to call the API and check the responses and duration of the request.
- B. Use a custom health check in the API that queries hosts to check the duration of the request.
- C. Implement a custom AWS Lambda function with an Amazon EventBridge event to periodically call the API and check the responses and duration of the request.
- D. Use the built-in API Gateway metrics to monitor the average duration of the API response.

Answer: A

NEW QUESTION 5

- (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 6

- (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 7

- (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 8

- (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 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)

A developer is creating a serverless application that uses an AWS Lambda function. The developer will use AWS CloudFormation to deploy the application. The application will write logs to Amazon CloudWatch Logs. The developer has created a log group in a CloudFormation template for the application to use. The developer needs to modify the CloudFormation template to make the name of the log group available to the application at runtime.

Which solution will meet this requirement?

- A. Use the AWS:: Include transform in CloudFormation to provide the log group's name to the application.
- B. Pass the log group's name to the application in the user data section of the CloudFormation template.
- C. Use the CloudFormation template's Mappings section to specify the log group's name for the application.
- D. Pass the log group's Amazon Resource Name (ARN) as an environment variable to the Lambda function.

Answer: C

NEW QUESTION 10

- (Exam Topic 4)

A company has a serverless application that uses AWS Lambda functions and AWS Systems Manager parameters to store configuration data. The company moves the Lambda functions inside the VPC and into private subnets. The Lambda functions are now producing errors in their attempts to access Systems Manager parameters.

Which solution will allow the Lambda functions to access Systems Manager parameters inside the VPC?

- A. Configure security groups to allow access to Systems Manager.
- B. Create an interface VPC endpoint for Systems Manager.
- C. Use an internet gateway from inside the VPC.
- D. Create a gateway VPC endpoint for Systems Manager.

Answer: B

NEW QUESTION 12

- (Exam Topic 4)

An open-source map application gathers data from several geolocation APIs. The application's source code repository is public and can be used by anyone, but the geolocation APIs must not be directly accessible.

A developer must implement a solution to prevent the credentials that are used to access the APIs from becoming public. The solution also must ensure that the application still functions properly.

Which solution will meet these requirements MOST cost-effectively?

- A. Store the credentials in AWS Secrets Manager
- B. Retrieve the credentials by using the GetSecretValue API operation.
- C. Store the credentials in AWS Key Management Service (AWS KMS). Retrieve the credentials by using the GetPublicKey API operation.
- D. Store the credentials in AWS Security Token Service (AWS STS). Retrieve the credentials by using the GetCallerIdentity API operation.
- E. Store the credentials in AWS Systems Manager Parameter Store
- F. Retrieve the credentials by using the GetParameter API operation.

Answer: D

Explanation:

Secrets Manager: It is paid. The storage cost is \$0.40 per secret per month and API interactions cost is \$0.05 per 10,000 API calls. Parameter Store: For Standard parameters, No additional charge for storage and standard throughput. For higher throughput, API interactions cost is \$0.05 per 10,000 API calls. For Advanced parameters, storage cost is \$0.05 per advanced parameter per month and API interactions cost is \$0.05 per 10,000 API calls.

<https://aws.amazon.com/systems-manager/pricing/>

NEW QUESTION 13

- (Exam Topic 4)

A developer is receiving an intermittent ProvisionedThroughputExceededException error from an application that is based on Amazon DynamoDB. According to the Amazon CloudWatch metrics for the table, the application is not exceeding the provisioned throughput

What could be the cause of the issue?

- A. The DynamoDB table storage size is larger than the provisioned size.
- B. The application is exceeding capacity on a particular hash key.
- C. The DynamoDB table is exceeding the provisioned scaling operations.
- D. The application is exceeding capacity on a particular sort key.

Answer: B

NEW QUESTION 15

- (Exam Topic 4)

A developer is writing a new serverless application for a company. Several other developers must collaborate on the code for this application, and the company expects frequent changes to the code. The developer needs to deploy the code from source control to AWS Lambda with the fewest number of manual steps.

Which strategy for the build and deployment should the developer use to meet these requirements?

- A. Build the code locally, and then upload the code into the source control system
- B. When a release is needed, run AWS CodePipeline to extract the uploaded build and deploy the resources.
- C. Use the AWS Serverless Application Model (AWS SAM) CLI to build and deploy the application from the developer's local machine with the latest version checked out locally.
- D. Use AWS CodeBuild and AWS CodePipeline to invoke builds and corresponding deployments when configured source controlled branches have pull requests merged into them
- E. Use the Lambda console to upload a .zip file of the application that is created by the AWS ServerlessApplication Model (AWS SAM) CLI build command.

Answer: C

NEW QUESTION 19

- (Exam Topic 4)

A media company wants to test its web application more frequently. The company deploys the application by using a separate AWS CloudFormation stack for each environment. The same CloudFormation template is deployed to each stack as the application progresses through the development lifecycle.

A developer needs to build an automated alert for the quality assurance (QA) team. The developer wants the alert to occur for new deployments in the final pre-

production environment.

Which solution will moot these requirements?

- A. Create an Amazon Simple Notification Service (Amazon SNS) topic
- B. Add a subscription to notify the QA team
- C. Update the CloudFormation stack options to point to the SNS topic in the production environment
- D. Most Voted
- E. Create an AWS Lambda function that notifies the QA team
- F. Create an Amazon EventBridge rule to invoke the Lambda function on the default event bus
- G. Filter the events on the CloudFormation service and the CloudFormation stack Amazon Resource Name (ARN).
- H. Create an Amazon CloudWatch alarm that monitors the metrics from CloudFormation
- I. Filter the metrics on the stack name and the stack status
- J. Configure the alarm to notify the QA team.
- K. Create an AWS Lambda function that notifies the QA team
- L. Configure the event source mapping to receive events from CloudFormation
- M. Specify the filtering values to limit invocations to the desired CloudFormation stack.

Answer: A

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/cloudformation-rollback-email/>

<https://aws.amazon.com/premiumsupport/knowledge-center/cloudformation-rollback-email/> <https://www.trendmicro.com/cloudoneconformity/knowledge-base/aws/CloudFormation/cloudformation-stack-n>

NEW QUESTION 20

- (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 implements 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 22

- (Exam Topic 4)

A company uses Amazon DynamoDB as a data store for its order management system. The company frontend application stores orders in a DynamoDB table. The DynamoDB table is configured to send change events to a DynamoDB stream. The company uses an AWS Lambda function to log and process the incoming orders based on data from the DynamoDB stream.

An operational review reveals that the order quantity of incoming orders is sometimes set to 0. A developer needs to create a dashboard that will show how many unique customers this problem affects each day.

What should the developer do to implement the dashboard?

- A. Grant the Lambda function's execution role permissions to upload logs to Amazon CloudWatch Logs. Implement a CloudWatch Logs Insights query that selects the number of unique customers for orders with order quantity equal to 0 and groups the results in 1-day period
- B. Add the CloudWatch Logs Insights query to a CloudWatch dashboard.
- C. Use Amazon Athena to query AWS CloudTrail API logs for API call
- D. Implement an Athena query that selects the number of unique customers for orders with order quantity equal to 0 and groups the results in 1 -day period
- E. Add the Athena query to an Amazon CloudWatch dashboard.
- F. Configure the Lambda function to send events to Amazon EventBridge
- G. Create an EventBridge rule that groups the number of unique customers for orders with order quantity equal to 0 in 1 -day period
- H. Add a CloudWatch dashboard as the target of the rule.
- I. Turn on custom Amazon CloudWatch metrics for the DynamoDB stream of the DynamoDB table. Create a CloudWatch alarm that groups the number of unique customers for orders with order quantity equal to 0 in 1-day period
- J. Add the CloudWatch alarm to a CloudWatch dashboard.

Answer: D

NEW QUESTION 24

- (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 28

- (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.ht](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 32

- (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 35

- (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 39

- (Exam Topic 4)

A developer has an application container, an AWS Lambda function, and an Amazon Simple Queue Service (Amazon SQS) queue. The Lambda function uses the SQS queue as an event source. The Lambda function makes a call to a third-party machine learning API when the function is invoked. The response from the third-party API can take up to 60 seconds to return.

The Lambda function's timeout value is currently 65 seconds. The developer has noticed that the Lambda function sometimes processes duplicate messages from the SQS queue.

What should the developer do to ensure that the Lambda function does not process duplicate messages?

- A. Configure the Lambda function with a larger amount of memory.
- B. Configure an increase in the Lambda function's timeout value.
- C. Configure the SQS queue's delivery delay value to be greater than the maximum time it takes to call the third-party API.
- D. Configure the SQS queue's visibility timeout value to be greater than the maximum time it takes to call the third-party API.

Answer: A

NEW QUESTION 44

- (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 47

- (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 49

- (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 53

- (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 58

- (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 60

- (Exam Topic 4)

A developer creates a web service that performs many critical activities. The web service code uses an AWS SDK to publish noncritical metrics to Amazon CloudWatch by using the PutMetricData API. The web service must return results to the caller as quickly as possible. The response data from the PutMetricData API is not necessary to create the web service response.

Which solution will MOST improve the response time of the web service?

- A. Upgrade to the latest version of the AWS SDK.
- B. Call the PutMetricData API in a background thread.
- C. Use the AWS SDK to perform a synchronous call to an AWS Lambda function.
- D. Call the PutMetricData API within the Lambda function.
- E. Send metric data to an Amazon Simple Queue Service (Amazon SQS) queue.
- F. Configure an AWS Lambda function with the queue as the event source.
- G. Call the PutMetricData API within the Lambda function.

Answer: D

Explanation:

<https://docs.aws.amazon.com/lambda/latest/dg/invoke-async.html#invoke-async-api>

NEW QUESTION 64

- (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 66

- (Exam Topic 4)

A developer is creating an AWS CloudFormation template to deploy Amazon EC2 instances across multiple AWS accounts. The developer must choose the EC2 instances from a list of approved instance types.

How can the developer incorporate the list of approved instance types in the CloudFormation template?

- A. Create a separate CloudFormation template for each EC2 instance type in the list.
- B. In the Resources section of the CloudFormation template, create resources for each EC2 instance type in the list.
- C. In the CloudFormation template, create a separate parameter for each EC2 instance type in the list.
- D. In the CloudFormation template, create a parameter with the list of EC2 instance types as AllowedValues.

Answer: B

NEW QUESTION 69

- (Exam Topic 4)

A company uses AWS CloudFormation to deploy an application that uses an Amazon API Gateway REST API with AWS Lambda function integration. The application uses Amazon DynamoDB for data persistence. The application has three stages, development, testing, and production. Each stage uses its own DynamoDB table.

The company has encountered unexpected issues when promoting changes to the production stage. The changes were successful in the development and testing stages. A developer needs to route 20% of the traffic to the new production stage API with the next production release. The developer needs to route the remaining 80% of the traffic to the existing production stage. The solution must minimize the number of errors that any single customer experiences.

Which approach should the developer take to meet these requirements?

- A. Update 20% of the planned changes to the production stage.
- B. Deploy the new production stage.
- C. Monitor the result.
- D. Repeat this process five times to test all planned changes.
- E. Update the Amazon Route 53 DNS record entry for the production stage API to use a weighted routing policy. Set the weight to a value of 80. Add a second record for the production domain name. Change the second routing policy to a weighted routing policy.
- F. Set the weight of the second policy to a value of 20. Change the alias of the second policy to use the testing stage API.
- G. Deploy an Application Load Balancer (ALB) in front of the REST API. Change the production API Amazon Route 53 record to point traffic to the ALB. Register the production and testing stages as targets of the ALB with weights of 80% and 20%, respectively.
- H. Configure canary settings for the production stage API.
- I. Change the percentage of traffic directed to canary deployment to 20%. Make the planned updates to the production stage. Deploy the changes.

Answer: B

NEW QUESTION 71

- (Exam Topic 4)

A developer is exposing an API by using Amazon API Gateway and AWS Lambda as the backend for an application. The developer wants to add validation rules for a POST method to ensure that the data (from the frontend web form) is valid. The validation rules must include mandatory fields, data type, length, and regular expressions.

Which solution will meet these requirements?

- A. Create an API Gateway model with schema for data validation.
- B. Create API Gateway HTTP request headers for data validation.
- C. Create API Gateway URL query string parameters for data validation.
- D. Create API Gateway URL path parameters for data validation.

Answer: D

Explanation:

<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-method-request-validation.html>

NEW QUESTION 74

- (Exam Topic 4)

A company's security policies require all database passwords to be rotated every 30 days. The company uses different database platforms, including Amazon Aurora databases and proprietary NoSQL document databases, for different applications. A developer needs to implement a solution for password rotation. Which solution will meet these requirements?

- A. Create an AWS Lambda rotation function that has appropriate IAM permissions. Store the password in AWS Secrets Manager. Configure Secrets Manager to rotate the password by using the Lambda function.
- B. Encrypt the existing password with AWS Key Management Service (AWS KMS). Export the existing password. Generate a random password with AWS KMS. Use the AWS KMS password renewal feature to replace the existing password with the new password.
- C. Create an AWS Lambda rotation function that has appropriate IAM permissions. Store the password in AWS Systems Manager Parameter Store. Configure Parameter Store to rotate the password by using the Lambda function.
- D. Integrate AWS Systems Manager Parameter Store with a Key Management Interoperability Protocol (KMIP)-compliant third-party secret manager to enable third-party database password rotation on AWS.

Answer: C

NEW QUESTION 78

- (Exam Topic 4)

A developer is troubleshooting a new AWS Lambda function. The function should run automatically each time a new object is uploaded to an Amazon S3 bucket. However, the developer finds that all calls failed before they reached the application code inside the function. Which of the following is a possible reason for this failure?

- A. The function resource policy does not allow access from Amazon S3.
- B. The function execution role does not allow access from Amazon S3.
- C. The function execution role does not allow access to Amazon S3.
- D. The IAM user does not have access to Amazon S3.

Answer: C

NEW QUESTION 79

- (Exam Topic 4)

A company is using continuous integration/continuous delivery (CI/CD) systems that run on premises. Which AWS service should the developer use to meet these requirements?

- A. AWS Cloud9
- B. AWS CodeBuild
- C. AWS Elastic Beanstalk
- D. AWS CodeDeploy

Answer: D

NEW QUESTION 83

- (Exam Topic 4)

A developer deploys an e-commerce 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 88

- (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 92

- (Exam Topic 4)

A developer has created a Java application that makes HTTP requests directly to AWS services. Application logging shows 5xx HTTP response codes that occur at irregular intervals. The errors are affecting users.

How should the developer update the application to improve the application's resiliency?

- A. Revise the request content in the application code.
- B. Use the AWS SDK for Java to interact with AWS APIs.
- C. Scale out the application so that more instances of the application are running.
- D. Add additional logging to the application code.

Answer: B

NEW QUESTION 93

- (Exam Topic 4)

A developer created a web API that receives requests by using an internet-facing Application Load Balancer (ALB) with an HTTPS listener. The developer configures an Amazon Cognito user pool and wants to ensure that every request to the API is authenticated through Amazon Cognito.

What should the developer do to meet this requirement?

- A. Add a listener rule to the listener to return a fixed response if the Authorization header is missing.
- B. Set the fixed response to 401 Unauthorized.
- C. Create an authentication action for the listener rules of the ALB.
- D. Set the rule action type to authenticate-cognito. Set the OnUnauthenticatedRequest field to "deny."
- E. Create an Amazon API Gateway API.
- F. Configure all API methods to be forwarded to the ALB endpoint. Create an authorizer of the COGNITO_USER_POOLS type.
- G. Configure every API method to use that authorizer.
- H. Create a new target group that includes an AWS Lambda function target that validates the Authorization header by using Amazon Cognito.
- I. Associate the target group with the listener.

Answer: C

NEW QUESTION 95

- (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 96

- (Exam Topic 4)

A developer is debugging an AWS Lambda function behind an Amazon API Gateway. Whenever the API Gateway endpoint is called, HTTP status code 200 is returned even though AWS Lambda is recording a 4xx error.

What change needs to be made to return a proper error code through the API Gateway?

- A. Enable CORS in the API Gateway method settings.
- B. Use a Lambda proxy integration to return HTTP codes and headers.
- C. Enable API Gateway error pass-through.
- D. Return the value in the header X-Amzn-ErrorType.

Answer: B

NEW QUESTION 100

- (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 105

- (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 106

- (Exam Topic 3)

An application is experiencing performance issues based on increased demand. This increased demand is on read-only historical records pulled from an Amazon RDS-hosted database with custom views and queries. A developer improve performance without changing the database structure.

Which approach will improve performance and MINIMIZE management overhead?

- A. Deploy Amazon DynamoDB, move all the data, and point to DynamoDB.
- B. Deploy Amazon ElastiCache for Redis and cache the data for the application.
- C. Deploy Memcached on Amazon EC2 and cache the data for the application.
- D. Deploy Amazon DynamoDB Accelerator (DAX) on Amazon RDS to improve cache performance

Answer: B

NEW QUESTION 110

- (Exam Topic 3)

A developer has written an application that runs on Amazon EC2 instances. The developer is adding

functionality for the application to write objects to an Amazon S3 bucket Which policy must the developer modify to allow the instances to write these objects?

- A. The 1AM policy that is attached to the EC2 instance profile role
- B. The session policy that is applied to the EC2 instance role session
- C. The AWS Key Management Service (AWS KMS) key policy that is attached to the EC2 instance profile role.
- D. The Amazon VPC endpoint policy

Answer: A

NEW QUESTION 112

- (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 me 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 116

- (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 formulti-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 even
- 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 121

- (Exam Topic 3)

A developer is building a backend system for the long-term storage of information from an inventory management system. The information needs to be stored so that other teams can build tools to report and analyze the data
How should the developer implement this solution to achieve the FASTEST running time?

- A. Create an AWS Lambda function that writes to Amazon S3 synchronously Increase the function's concurrency to match the highest expected value of concurrent scans and requests.
- B. Create an AWS Lambda function that writes to Amazon S3 asynchronously Configure a dead-letter queue to collect unsuccessful invocations
- C. Create an AWS Lambda function that writes to Amazon S3 synchronously Set the inventory system to retry failed requests.
- D. Create an AWS Lambda function that writes to an Amazon ElastiCache for Redis cluster asynchronously Configure a dead-letter queue to collect unsuccessful invocations.

Answer: A

NEW QUESTION 125

- (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 approved 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 delete Enable EC2 termination protection on production instances.
- D. Enable EC2 termination protection on production instance
- E. Deny the developer group's permissions policy access to terminate instance
- F. Create a new role that developer can assume when termination is necessary.

Answer: A

NEW QUESTION 128

- (Exam Topic 3)

A developer has written an AWS Lambda function using Java as the runtime environment. The developer wants to isolate a performance bottleneck in the code.
Which steps should be taken to reveal the bottleneck?

- A. Use the Amazon CloudWatch API to write timestamps to a custom CloudWatch metric Use the CloudWatch console to analyze the resulting data
- B. Use the AWS X-Ray API to write trace data into X-Ray from strategic places within the code Use the Amazon CloudWatch console to analyze the resulting data
- C. Use the AWS X-Ray API to write trace data into X-Ray from strategic places within the code
- D. Use the X-Ray console to analyze the resulting data
- E. Use the Amazon CloudWatch API to write timestamps to a custom CloudWatch metric Use the AWS X-Ray console to analyze the resulting data

Answer: C

NEW QUESTION 129

- (Exam Topic 3)

A company runs its APIs using Amazon API Gateway in front of AWS Lambda functions The company wants to add logging at the API level Each API must have production and development environments The developer wants to enable different logging levels in both environments.
How can these requirements be met?

- A. Set up a stage for each environment In each stage, point to different Lambda functions that implement the logging logic in the code Access the logs in Amazon CloudWatch Logs
- B. Set up a stage for each environment In each stage, define a different logging level according to the logging requirements Access the logs in Amazon CloudWatch Logs
- C. Set up a stage and use the same Lambda functions In Amazon CloudWatch Logs set up a filter based on the log level according to the logging requirements
- D. Set up a stage for each environment In each stage, define a variable for the log level Set the value according to the logging requirements.

Answer: A

NEW QUESTION 130

- (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 132

- (Exam Topic 3)

A company has a large number of documents that are stored securely in Amazon S3 The company is creating an application that occasionally will read these documents The application will be deployed on Amazon EC2 instances.
The company's security requirements mandate that no long-term credentials can be stored on the EC2 instances and that only the needed documents can be accessed Only authorized users and applications can access the documents access must be logged by Amazon S3, and each document must follow S3 Lifecycle policies for archival and destruction

What should a developer do to meet these requirements?

- A. Create an event to invoke an AWS Lambda function when a document is uploaded Configure the function to write the documents to an Amazon Elastic File System (Amazon EFS) file system Configure the EC2 instances to mount the EFS file system Configure the application to access the documents that are stored in the file system as needed
- B. Create a user that has programmatic credentials, and attach a policy that allows read access to the S3 bucket Use the AWS CLI to configure those credentials for the EC2 instances to use Create an Amazon Machine Image (AMI), and add the access key and secret access key to the user data section to create environment variables Use the AMI to launch each EC2 instance that runs the application Add application code to use the keys that are stored in the environment variables to access the S3 bucket objects as needed.
- C. Modify the S3 bucket, make the bucket public, and make each object public Add application code to make REST calls to access the objects in the S3 bucket as needed
- D. Create an IAM role with permissions to read objects from Amazon S3 Attach the role to the EC2 instances as an instance profile Add application code to access the objects in the S3 bucket as needed.

Answer: D

NEW QUESTION 135

- (Exam Topic 3)

A developer is building a new application that uses an Amazon DynamoDB table. The specification states that all items that are older than 48 hours must be removed

Which solution will meet this requirement?

- A. Create a new attribute that has the Number data type Add a local secondary index (LSI) for this attribute and enable TTL with an expiration of 48 hours In the application code, set the value of this attribute to the current timestamp for each new item that is being inserted.
- B. Create a new attribute that has the String data type Add a local secondary index (LSI) for this attribute and enable TTL with an expiration of 48 hours In the application code, set the value of this attribute to the current timestamp for each new item that is being inserted.
- C. Create a new attribute that has the Number data type Enable TTL on the DynamoDB table for this attribute in the application code set the value of this attribute to the current timestamp plus 48 hours for each new item that is being inserted
- D. Create a new attribute that has the String data type Enable TTL on the DynamoDB table for this attribute In the application code set the value of this attribute to the current timestamp plus 48 hours for each new item that is being inserted

Answer: C

NEW QUESTION 136

- (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 139

- (Exam Topic 3)

A large company has its application components distributed across.. company needs to collect and visualize trace data across these accounts. What should be used to meet these requirements?

- A. AWS X-Ray
- B. Amazon CloudWatch
- C. Amazon VPC flow logs
- D. Amazon Elasticsearch Service

Answer: A

NEW QUESTION 140

- (Exam Topic 3)

A developer is building a static, client-side rendered website that is powered by ReactJS The code has no server-side generated components and does not need to run any programming languages on the server However the code serves static HTML, CSS, and JavaScript to the client on each request The developer's solution to host the website must maximize performance and cost-effectiveness

Which combination of AWS services or resources should the developer use to meet these requirements?

- A. Application Load Balancer and Amazon EC2
- B. Amazon API Gateway and AWS Lambda
- C. Amazon CloudFront and Amazon S3
- D. Amazon CloudFront and AWS Elastic Beanstalk

Answer: C

NEW QUESTION 144

- (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 148

- (Exam Topic 3)

A developer is creating an event handling system. To handle messages asynchronously, the developer created a standard Amazon SQS queue. Quality assurance testing reveals that some events were processed multiple times.

What is the recommended way to ensure the events are not processed more than once?

- A. Change long polling to short polling.
- B. Use a FIFO queue and configure deduplication
- C. Convert the standard SQS queue into a FIFO queue
- D. Send the messages with message timers

Answer: C

NEW QUESTION 151

- (Exam Topic 3)

A developer receives the following error message when trying to launch or terminate an Amazon EC2 instance using a boto3 script.

```
boto.exception.BotoServerError: BotoServerError: 503 Service Unavailable
<?xml version="1.0" encoding="UTF-8"?>
<Response><Errors><Error><Code>RequestLimitExceeded</Code>
<Message>Request limit exceeded.</Message></Error></Errors><RequestID>bfddec84-53b3-4701-b728-dceefb696ced</RequestID>
</Response>
```

What should the developer do to correct this error message?

- A. Assign an IAM role to the EC2 instance to allow necessary API calls on behalf of the client.
- B. Implement an exponential backoff algorithm for optimizing the number of API requests made to Amazon EC2.
- C. Increase the overall network bandwidth to handle higher API request rates.
- D. Upgrade to the latest AWS CLI version so that boto3 can handle higher request rates

Answer: D

NEW QUESTION 156

- (Exam Topic 3)

A developer implemented a static website hosted in Amazon S3 that makes web service requests hosted in Amazon API Gateway 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-Request-Headers header to the request

Answer: B

Explanation:

<https://docs.aws.amazon.com/apigateway/latest/developerguide/how-to-cors-console.html>

NEW QUESTION 161

- (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 166

- (Exam Topic 3)

A developer is building an application using an Amazon API Gateway REST API backed by an AWS Lambda function that interacts with an Amazon DynamoDB table. During testing, the developer observes high latency when making requests to the API.

How can the developer evaluate the end-to-end latency and identify performance bottlenecks?

- A. Enable AWS CloudTrail logging and use the logs to map each latency and bottleneck
- B. Enable and configure AWS X-Ray tracing on API Gateway and the Lambda function. Use X-Ray to trace and analyze user requests

- C. Enable Amazon CloudWatch Logs for the Lambda function Enable execution logs for API Gateway to view and analyze user request logs.
- D. Enable VPC Flow Logs to capture and analyze network traffic within the VPC

Answer: B

NEW QUESTION 171

- (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 175

- (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 addressked?

- 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 179

- (Exam Topic 3)

A development team is building a new application that will run on Amazon EC2 and use Amazon DynamoDB as a storage layer The developers all have assigned IAM user accounts in the same IAM group The developers currently can launch EC2 instances and they need to be able to launch EC2 instances with an instance role allowing access to Amazon DynamoDB.

Which AWS IAM changes are needed when creating an instance role to provide this functionality?

- A. Create an IAM permission policy attached to the role that allows access to DynamoDB Add a trust policy to the role that allows DynamoDB to assume the role Attach a permissions policy to the development group in AWS IAM that allows developers to use the IAM GetRole and IAM PassRole permissions for the role
- B. Create an IAM permissions policy attached to the role that allows access to DynamoDB Add a trust policy to the role that allows Amazon EC2 to assume the role Attach a permissions policy to the development group in AWS IAM that allows developers to use the IAM PassRole permission for the role
- C. Create an IAM permission policy attached to the role that allows access to Amazon EC2 Add a trust policy to the role that allows DynamoDB to assume the role Attach a permissions policy to the development group in AWS IAM that allows developers to use the IAM PassRole permission for the role
- D. Create an IAM permissions policy attached to the role that allows access to DynamoDB Add a trust policy to the role that allows Amazon EC2 to assume the role Attach a permissions policy to the development group in AWS IAM that allows developers to use the iam GetRole permission for the role.

Answer: C

NEW QUESTION 184

- (Exam Topic 3)

An organization is using Amazon API Gateway to provide a public API called "Survey" for collecting user feedback posts about its products The survey API has "DEV" and "PROD" stages and consists of one resource "/feedback" which allows users to retrieve/create/update single feedback posts.

A version-controlled Swagger file is used to define a new API that retrieves multiple feedback posts To add the new API resource "/listFeedbackForProduct" the developer makes changes to the Swagger file defining an API uploads the file to the organization's version control system, and uses the API Gateway Import API feature to apply the changes to the Survey API After successful import the developer runs the tests against the DEV stage and finds that resource "/listFeedbackForProduct" is not available.

What is MOST likely the reason for resource not being available?

- A. Even though the Swagger import was successful, resource creation failed afterwards
- B. There is a propagation delay of several minutes in creating API Gateway resources after import
- C. The developer needs to restart the API Gateway stage after import in order to apply the changes.
- D. The developer needs to create a new deployment after import in order to deploy the changes

Answer: A

NEW QUESTION 189

- (Exam Topic 3)

A developer is writing an application to analyze the traffic to a fleet of Amazon EC2 instances The EC2 instances run behind a public Application Load Balancer (ALB). An HTTP server runs on each of the EC2 instances, logging all requests to a log file.

The developer wants to capture the client public IP addresses. The developer analyzes the log files and notices only the IP address of the ALB

What must the developer do to capture the client public IP addresses in the log file?

- A. Add a Host header to the HTTP server log configuration file
- B. Install the Amazon CloudWatch Logs agent on each EC2 instance
- C. Configure the agent to write to the log file.
- D. Install the AWS X-Ray daemon on each EC2 instance Configure the daemon to write to the log file.
- E. Add an X-Forwarded-For header to the HTTP server log configuration file.

Answer: C

NEW QUESTION 193

- (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 requirements?

- 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 195

- (Exam Topic 3)

A developer is designing an AWS Lambda function that creates temporary files that are less than 10 MB during execution. The temporary files will be accessed and modified multiple times during execution. The developer has no need to save or retrieve these files in the future.

Where should the temporary file be stored?

- A. the /tmp directory
- B. Amazon EFS
- C. Amazon EBS
- D. Amazon S3

Answer: A

NEW QUESTION 197

- (Exam Topic 3)

A developer used the BatchWriteItem API operation to insert items in an Amazon DynamoDB table. DynamoDB returned a few items as unprocessed due to throttling. The developer decides to retry the records on the unprocessed items.

What should the developer do to reprocess the records with the LEAST number of API calls?

- A. Retry the BatchWriteItem operation immediately
- B. Perform the PutItem operation on the unprocessed items individually instead of using the BatchWriteItem operation
- C. Delay the BatchWriteItem operation by using progressively longer wait times between retries, or exponential backoff
- D. Delete the items that were successfully processed, and reissue a new BatchWriteItem operation

Answer: D

NEW QUESTION 201

- (Exam Topic 3)

A developer has written an application hosted on Amazon EC2 instances. The application generates and uploads thousands of new objects to an Amazon S3 bucket located in the same AWS region. The size of each object is less than 1 MB. The application is taking too long to run.

How can the performance of the application be improved?

- A. Use the S3 Multipart Upload API
- B. Use S3 Transfer Acceleration
- C. Upload the objects in parallel to Amazon S3
- D. Add a random prefix to the object keys

Answer: D

NEW QUESTION 205

- (Exam Topic 3)

A company's fleet of Amazon EC2 instances receives data from millions of users through an API. The servers batch the data, add an object for each user, and upload the objects to an S3 bucket to ensure high access rates. The object attributes are Customer ID, Server ID, TS-Server (TimeStamp and Server ID), the size of the object, and a timestamp. A developer wants to find all the objects for a given user collected during a specified time range.

After creating an S3 object creation event, how can the developer achieve this requirement?

- A. Execute an AWS Lambda function in response to the S3 object creation events that creates an Amazon DynamoDB record for every object with the Customer ID as the partition key and the Server ID as the sort key. Retrieve all the records using the Customer ID and Server ID attributes.
- B. Execute an AWS Lambda function in response to the S3 object creation events that creates an Amazon Redshift record for every object with the Customer ID as the partition key and TS-Server as the sort key. Retrieve all the records using the Customer ID and TS-Server attributes.
- C. Execute an AWS Lambda function in response to the S3 object creation events that creates an Amazon DynamoDB record for every object with the Customer ID as the partition key and TS-Server as the sort key. Retrieve all the records using the Customer ID and TS-Server attributes.
- D. Execute an AWS Lambda function in response to the S3 object creation events that creates an Amazon Redshift record for every object with the Customer ID as the partition key and the Server ID as the sort key.
- E. Retrieve all the records using the Customer ID and Server ID attributes.

Answer: C

NEW QUESTION 209

- (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 211

- (Exam Topic 3)

A developer must cache dependent artifacts from Maven Central, a public package repository, as part of an application's build pipeline. The build pipeline has an AWS CodeArtifact repository where artifacts of the build are published. The developer needs a solution that requires minimum changes to the build pipeline.

Which solution meets these requirements?

- A. Modify the existing CodeArtifact repository to associate an upstream repository with the public package repository
- B. Create a new CodeArtifact repository that has an external connection to the public package repository
- C. Create a new CodeArtifact domain that contains a new repository that has an external connection to the public package repository
- D. Modify the CodeArtifact repository resource policy to allow artifacts to be fetched from the public package repository

Answer: D

NEW QUESTION 213

- (Exam Topic 3)

A gaming application stores scores for players in an Amazon DynamoDB table that has four attributes: user_id, user_name, user_score, and user_rank. The users are allowed to update their names only. A user is authenticated by web identity federation.

Which set of conditions should be added in the policy attached to the role for the dynamodb:PutItem API call?

A)

```
"Condition": {
  "ForAllValues:StringEquals": {
    "dynamodb:LeadingKeys": [
      "${www.amazon.com:user_id}"
    ],
    "dynamodb:Attributes": [
      "user_name"
    ]
  }
}
```

B)

```
"Condition": {
  "ForAllValues:StringEquals": {
    "dynamodb:LeadingKeys": [
      "${www.amazon.com:user_name}"
    ],
    "dynamodb:Attributes": [
      "user_id"
    ]
  }
}
```

C)

```
"Condition": {
  "ForAllValues:StringEquals": {
    "dynamodb:LeadingKeys": [
      "${www.amazon.com:user_id}"
    ],
    "dynamodb:Attributes": [
      "user_name", "user_id"
    ]
  }
}
```

D)

```
"Condition": {
  "ForAllValues:StringEquals": {
    "dynamodb:LeadingKeys": [
      "${www.amazon.com:user_name}"
    ],
    "dynamodb:Attributes": [
      "user_name", "user_id"
    ]
  }
}
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: A

NEW QUESTION 214

- (Exam Topic 3)

A development team decides to adopt a continuous integration/continuous delivery (CI/CD) process using AWS CodePipeline and AWS CodeCommit for a new application. However, management wants a person to review and approve the code before it is deployed to production. How can the development team add a manual approver to the CI/CD pipeline?

- A. Use AWS SES to send an email to approvers when their action is required. Develop a simple application that allows approvers to accept or reject a build. Invoke an AWS Lambda function to advance the pipeline when a build is accepted.
- B. If approved, add an approved tag when pushing changes to the CodeCommit repository.
- C. CodePipeline will proceed to build and deploy approved commits without interruption.
- D. Add an approval step to CodeCommit. Commits will not be saved until approved.
- E. Add an approval action to the pipeline.
- F. Configure the approval action to publish to an Amazon SNS topic when approval is required.
- G. The pipeline execution will stop and wait for an approval.

Answer: D

NEW QUESTION 215

- (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 216

- (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 218

- (Exam Topic 3)

A developer has launched an application that calls an API by way of Amazon API Gateway. It offers information that changes several times a day, but is not updated in real time. The application has become so popular that the API endpoint is overloaded and that traffic to the endpoint must be reduced.

What can the developer do to address the performance issues?

- A. Enable API caching in Amazon ElastiCache.
- B. Enable an Auto Scaling group on the endpoint service and database.
- C. Create an additional API Gateway and use an Application Load Balancer.

Answer: A

NEW QUESTION 223

- (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 227

- (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 232

- (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 237

- (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 240

- (Exam Topic 3)

A developer wants to insert a record into an Amazon DynamoDB table as soon as a new file is added to an Amazon S3 bucket. Which set of steps would be necessary to achieve this?

- A. Create an event with Amazon CloudWatch Events that will monitor the S3 bucket and then insert the records into DynamoDB.
- B. Configure an S3 event to invoke a Lambda function that inserts records into DynamoDB.
- C. Create a Lambda function that will poll the S3 bucket and then insert the records into DynamoDB.
- D. Create a cron job that will run at a scheduled time and insert the records into DynamoDB.

Answer: B

NEW QUESTION 245

- (Exam Topic 3)

A company has a web application that uses an Amazon Cognito user pool for authentication. The company wants to create a login page with the company logo. What should a developer do to meet these requirements?

- A. Create a hosted user interface in Amazon Cognito and customize it with the company logo.
- B. Create a login page with the company logo and upload it to Amazon Cognito.
- C. Create a login page in Amazon API Gateway with the logo and save the link in Amazon Cognito.
- D. Upload the logo to the Amazon Cognito app settings and point to the logo on a custom login page.

Answer: A

NEW QUESTION 250

- (Exam Topic 3)

A developer is designing a distributed application built using a microservices architect spanning multiple AWS accounts. The company's operations team wants to analyze and debug application issues from a centralized account.

How can the developer meet these requirements?

- A. Use an Amazon X-Ray agent with role assumption on to publish data into the centralized account.
- B. Use Amazon X-Ray and create a new IAM user to publish the access keys into the centralized account.
- C. Use VPC Flow Logs to collect application logs across different accounts.
- D. Enable AWS CloudTrail to publish the trails in an Amazon S3 bucket in the centralized account.

Answer: A

NEW QUESTION 254

- (Exam Topic 3)

A company has a three-tier application that is deployed in Amazon Elastic Container Service (Amazon ECS). The application is using an Amazon RDS for MySQL DB Instance. The application performs more database reads than writes.

During times of peak usage, the application's performance degrades. When this performance degradation occurs, the DB instance's ReadLatency metric in Amazon CloudWatch increases suddenly.

How should a developer modify the application to improve performance?

- A. Use Amazon ElastiCache to cache query results.
- B. Scale the ECS cluster to contain more ECS instances.
- C. Add read capacity units (RCUs) to the DB instance.
- D. Modify the ECS task definition to increase the task memory.

Answer: A

NEW QUESTION 256

- (Exam Topic 3)

A developer needs to modify an application architecture to meet new functional requirements. Application data is stored in Amazon DynamoDB and processed for analysis in a nightly batch. The system analysts do not want to wait until the next day to view the processed data and have asked to have it available in near-real time.

Which application architect pattern would enable the data to be processed as it is received?

- A. Event driven
- B. Client served driven
- C. Fan-out driven
- D. Schedule driven

Answer: A

NEW QUESTION 258

- (Exam Topic 3)

A video-hosting website has two types of members: those who pay a fee, and those who do not. Each video upload places a message in Amazon SQS. A fleet of Amazon EC2 instances polls Amazon SQS and processes each video.

The developer needs to ensure that the videos uploaded by the paying members are processed first. How can the developer meet this requirement?

- A. Create two SQS queues: one for paying members, and one for non-paying members. Poll the paying member queue first and then poll the non-paying member queue.
- B. Use SQS to set priorities on individual items within a single queue: give the paying members' videos the highest priority.
- C. Use SQS to set priorities on individual items within a single queue and use Amazon SNS to encode the videos.
- D. Create two Amazon SNS topics: one for paying members and one for non-paying members. Use SNS topic subscription priorities to differentiate between the two types of members.

Answer: B

NEW QUESTION 262

- (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 265

- (Exam Topic 3)

A developer is building an application integrating an Amazon API Gateway with an AWS Lambda function. When calling the API, the developer receives the following error: Wed Nov 03 01:13:00 UTC 2017 : Method completed with status: 502. What should the developer do to resolve the error?

- A. Change the HTTP endpoint of the API to an HTTPS endpoint.
- B. Change the format of the payload sent to the API Gateway.
- C. Change the format of the Lambda function response to the API call.
- D. Change the authorization header in the API call to access the Lambda function.

Answer: C

NEW QUESTION 267

- (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 270

- (Exam Topic 3)

A developer creates an Amazon S3 bucket to store project status files that are uploaded hourly. The developer also creates an AWS Lambda function that will be used to process the project status files.

What should the developer do to invoke the function with the LEAST amount of AWS infrastructure?

- A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to invoke the function every 5 minutes and scan for new objects.
- B. Create an S3 event notification to invoke the function when a new object is created in the S3 bucket.
- C. Create an S3 event notification that publishes a message to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the function to the SNS topic.
- D. Create an S3 event notification that adds a message to an Amazon Simple Queue Service (Amazon SQS) queue. Configure the function to poll the queue.

Answer: B

NEW QUESTION 275

- (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 276

- (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 280

- (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 281

- (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 286

- (Exam Topic 3)

A developer converted an existing program to an AWS Lambda function in the console. The program runs properly on a local laptop, but shows an "Unable to import module" error when tested in the Lambda console

Which of the following can fix the error?

- A. Install the missing module and specify the current directory as the target Create a ZIP file to include all files under the current directory, and upload the ZIP file.
- B. Install the missing module in a lib directory Create a ZIP file to include all files under the lib directory, and upload the ZIP file as a dependency file
- C. In the Lambda code invoke a Linux command to install the missing modules under the /usr/lib directory
- D. In the Lambda console, create a LD_LIBRARY_PATH environment and specify the value for the system library path.

Answer: C

NEW QUESTION 287

- (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 291

- (Exam Topic 3)

A developer is automating a new application deployment with AWS Serverless Application Model (AWS SAM) The new application has one AWS Lambda function and one Amazon S3 bucket The Lambda function must access the S3 bucket to only read objects

How should the developer configure AWS SAM to grant the necessary read privilege to the S3 bucket?

- A. Reference a second Lambda authorizer function
- B. Add a custom S3 bucket policy to the Lambda function
- C. Create an Amazon Simple Queue Service (SQS) topic for only S3 object reads Reference the topic in the template
- D. Add the S3ReadPolicy template to the Lambda function's execution role

Answer: D

NEW QUESTION 292

- (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 293

- (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 298

- (Exam Topic 3)

A Lambda function processes data before sending it to a downstream service Each piece of data is approximately 1 MB in size After a security audit, the function t]is now required to encrypt the data before sending it downstream Which API call is required to perform the encryption?

- A. Pass the data to the KMS ReEncrypt API for encryption
- B. Use the KMS GenerateDataKey API to get an encryption key
- C. Use the KMS GenerateDataKeyWithoutPlainText API to get an encryption key
- D. Pass the data to KMS as part of the Encrypt API for encryption

Answer: D

NEW QUESTION 301

- (Exam Topic 3)

A developer has written the following IAM policy to provide access to an Amazon S3 bucket:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:GetObject",
        "s3:PutObject"
      ],
      "Resource": "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"
    },
    {
      "Effect": "Deny",
      "Action": "s3:*",
      "Resource": "arn:aws:s3:::DOC-EXAMPLE-BUCKET/secrets*"
    }
  ]
}
```

Which access does the policy allow regarding the s3:GetObject and s3:PutObject actions?

- A. Access on all buckets except the "DOC-EXAMPLE-BUCKET" bucket
- B. Access on all buckets that start with "DOC-EXAMPLE-BUCKET" except the "DOC-EXAMPLE-BUCKET/secrets" bucket
- C. Access on all objects in the "DOC-EXAMPLE-BUCKET" bucket along with access to all S3 actions for objects in the "DOC-EXAMPLE-BUCKET" bucket that start with "secrets"
- D. Access on all objects in the "DOC-EXAMPLE-BUCKET" bucket except on objects that start with "secrets"

Answer: D

Explanation:

Meaning:

DOC-EXAMPLE-BUCKET ==> bucket

DOC-EXAMPLE-BUCKET/* ==> contents in the bucket In this example,

ALLOW all "Objects" ==> DOC-EXAMPLE-BUCKET/*

DENY objects starting with secrets ==> DOC-EXAMPLE-BUCKET/secrets* <https://aws.amazon.com/blogs/security/iam-policies-and-bucket-policies-and-acls-oh-my-controlling-access-to-s>

NEW QUESTION 303

- (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 306

- (Exam Topic 2)

A developer must ensure that the IAM credentials used by an application in Amazon EC2 are not misused or compromised What should the developer use to keep user credentials secure?

- A. Environment variables
- B. AWS credentials file
- C. Instance profile credentials
- D. Command line options

Answer: C

NEW QUESTION 310

- (Exam Topic 2)

A company has a two-tier application running on an Amazon EC2 server that handles all of its AWS based

e-commerce activity During peak times, the backend servers that process orders are overloaded with requests. This results in some orders failing to process. A developer needs to create a solution that will re-factor the application.

Which steps will allow for more flexibility during peak times, while still remaining cost-effective? (Select TWO.)

- A. Increase the backend T2 EC2 instance size to x1 to handle the largest possible load throughout the year
- B. Implement an Amazon SQS queue to decouple the front-end and backend servers
- C. Use an Amazon SNS queue to decouple the front-end and backend servers.
- D. Migrate the backend servers to on-premises and pull from an Amazon SNS queue
- E. Modify the backend servers to pull from an Amazon SQS queue.

Answer: CD

NEW QUESTION 315

- (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 316

- (Exam Topic 2)

A Developer is creating a template that uses AWS CloudFormation to deploy an application. This application is serverless and uses Amazon API Gateway, Amazon DynamoDB, and AWS Lambda.

Which tool should the Developer use to define simplified syntax for expressing serverless resources?

- A. CloudFormation serverless intrinsic functions
- B. AWS Serverless Express
- C. An AWS serverless application model
- D. A CloudFormation serverless plugin

Answer: A

NEW QUESTION 318

- (Exam Topic 2)

A developer has written an application that runs on Amazon EC2 instances and generates a value every minute. The Developer wants to monitor and graph the values generated over time without logging in to the instance each time.

Which approach should the Developer use to achieve this goal?

- A. Use the Amazon CloudWatch metrics reported by default for all EC2 instances. View each value from the CloudWatch console.
- B. Develop the application to store each value in a file on Amazon S3 every minute with the Unix timestamp as the name.
- C. Publish each generated value as a custom metric to Amazon CloudWatch using available AWS SDKs.
- D. Store each value as a variable and add the variable to the list of EC2 metrics that should be reported to the Amazon CloudWatch console.

Answer: C

NEW QUESTION 321

- (Exam Topic 2)

A Developer is building a web application that uses Amazon API Gateway to expose an AWS Lambda function to process requests from clients. During testing, the Developer notices that the API Gateway times out even though the Lambda function finishes under the set time limit.

Which of the following API Gateway metrics in Amazon CloudWatch can help the Developer troubleshoot the issue? (Choose two.)

- A. CacheHitCount
- B. IntegrationLatency
- C. CacheMissCount
- D. Latency
- E. Count

Answer: BC

Explanation:

<https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-metrics-and-dimensions.html>

NEW QUESTION 322

- (Exam Topic 2)

While developing an application that runs on Amazon EC2 in an Amazon VPC, a Developer identifies the need for centralized storage of application-level logs.

Which AWS service can be used to securely store these logs?

- A. Amazon EC2 VPC Flow Logs
- B. Amazon CloudWatch Logs
- C. Amazon CloudSearch
- D. AWS CloudTrail

Answer: B

NEW QUESTION 324

- (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 326

- (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 327

- (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 332

- (Exam Topic 2)

A developer is writing an application in AWS Lambda. To simplify testing and deployments, the developer needs the database connection string to be easily changed without modifying the Lambda code. How can this requirement be met?

- A. Store the connection string as a secret in AWS Secrets Manager
- B. Store the connection string in an IAM user account.
- C. Store the connection string in AWS KMS
- D. Store the connection string as a Lambda layer.

Answer: A

NEW QUESTION 335

- (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

NEW QUESTION 339

- (Exam Topic 2)

A company needs a version control system for collaborative software development. Features of the system must include the following:

- Support for batches of changes across multiple files
- Parallel branching
- Version tracking

Which AWS service will meet these requirements?

- A. AWS CodePipeline
- B. Amazon S3
- C. AWS Code Build
- D. AWS CodeCommit

Answer: D

Explanation:

<https://docs.aws.amazon.com/codecommit/latest/userguide/welcome.html>

NEW QUESTION 340

- (Exam Topic 2)

A Developer is investigating an issue whereby certain requests are passing through an Amazon API Gateway endpoint /MyAPI, but the requests do not reach the AWS Lambda function backing /MyAPI. The Developer found that a second Lambda function sometimes runs at maximum concurrency allowed for the given AWS account.

How can the Developer address this issue?

- A. Manually reduce the concurrent execution limit at the account level
- B. Add another API Gateway stage for /MyAPI, and shard the requests
- C. Configure the second Lambda function's concurrency execution limit
- D. Reduce the throttling limits in the API Gateway /MyAPI endpoint

Answer: C

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2017/11/set-concurrency-limits-on-individual-aws-lambda-functi> You can now set a concurrency limit on individual AWS Lambda functions. The concurrency limit you set will reserve a portion of your account level concurrency limit for a given function. This feature allows you to throttle a given function if it reaches a maximum number of concurrent executions allowed, which you can choose to set.

NEW QUESTION 344

- (Exam Topic 2)

A developer wants to ensure the Amazon EC2 instances in AWS Elastic Beanstalk execute a certain set of commands before the application is ready to use Which Elastic Beanstalk feature will allow the developer to accomplish this?

- A. Rolling update
- B. Immutable update
- C. User data
- D. ebextensions

Answer: D

NEW QUESTION 349

- (Exam Topic 2)

An application uses Amazon Kinesis Data Streams to ingest and process large streams of data records in real time. Amazon EC2 instances consume and process the data from the shards of the Kinesis data stream by using Amazon Kinesis Client Library (KCL). The application handles the failure scenarios and does not require standby workers. The application reports that a specific shard is receiving more data than expected. To adapt to the changes in the rate of data flow, the "hot" shard is resharded.

Assuming that the initial number of shards in the Kinesis data stream is 4, and after resharding the number of shards increased to 6, what is the maximum number of EC2 instances that can be deployed to process data from all the shards?

- A. 12
- B. 6
- C. 4
- D. 1

Answer: B

Explanation:

Typically, when you use the KCL, you should ensure that the number of instances does not exceed the number of shards (except for failure standby purposes). Each shard is processed by exactly one KCL worker and has exactly one corresponding record processor, so you never need multiple instances to process one shard. However, one worker can process any number of shards, so it's fine if the number of shards exceeds the number of instances.

<https://docs.aws.amazon.com/streams/latest/dev/kinesis-record-processor-scaling.html>

NEW QUESTION 354

- (Exam Topic 2)

A Developer must analyze performance issues with production-distributed applications written as AWS Lambda functions. These distributed Lambda applications invoke other components that make up the applications.

How should the Developer identify and troubleshoot the root cause of the performance issues in production?

- A. Add logging statements to the Lambda functions, then use Amazon CloudWatch to view the logs.
- B. Use AWS Cloud Trail and then examine the logs

- C. Use AWS X-Ray, then examine the segments and errors
- D. Run Amazon Inspector agents and then analyze performance

Answer: C

Explanation:

<https://aws.amazon.com/blogs/developer/new-analyze-and-debug-distributed-applications-interactively-using-aw>

NEW QUESTION 356

- (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 361

- (Exam Topic 2)

An ecommerce startup is preparing for an annual sales event As the traffic to the company's application increases, the development team wants to be notified when the Amazon EC2 instance's CPU utilization exceeds 80%.

Which solution will meet this requirement?

- A. Create a custom Amazon CloudWatch alarm that sends a notification to an Amazon SNS topic when the CPU utilization exceeds 80%.
- B. Create a custom AWS CloudTrail alarm that sends a notification to an Amazon SNS topic when the CPU utilization exceeds 80%
- C. Create a cron job on the EC2 instance that executes the --describe-instance-information command on the host instance every 15 minutes and sends the results to an Amazon SNS topic
- D. Create an AWS Lambda function that queries the AWS CloudTrail logs for the CPUUtihzation metric every 15 minutes and sends a notification to an Amazon SNS topic when the CPU utilization exceeds 80%

Answer: C

NEW QUESTION 366

- (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 370

- (Exam Topic 2)

A Developer decides lo store highly secure data in Amazon S3 and wants to implement server-side encryption (SSF) with granular control of who can access the master key Company policy requires that the master key be created, rotated, and disabled easily when needed, all for security reasons.

Which solution should be used to moot these requirements?

- A. SSE with Amazon S3 managed keys (SSE-S3)
- B. SSFE with AWS KMS managed keys (SSE KMS)
- C. SSE with AWS Secrets Manager
- D. SSE with customer provided encryption keys

Answer: B

NEW QUESTION 375

- (Exam Topic 2)

A company has a legacy application that was migrated to a fleet of Amazon EC2 instances. The application stores data in a MySQL database that is currently installed on a single EC2 instance The company has decided to migrate the database from the EC2 instance to MySQL on Amazon RDS.

What should the developer do to update the application to support data storage in Amazon RDS?

- A. Update the database connection parameters in the application to point to the new RDS instance
- B. Add a script to the EC2 instance that implements an AWS SDK for requesting database credentials.
- C. Create a new EC2 instance with an IAM role that allows access to the new RDS database
- D. Create an AWS Lambda function that will route traffic from the EC2 instance to the RDS database.

Answer: A

NEW QUESTION 376

- (Exam Topic 2)

An application is using single -node Amazon ElastiCache for Redis instance to improve read performance. Over time, demand for the application has increased exponentially, which has increased the load on the ElastiCache instance. It is critical that this cache layer handles the load and is resilient in case of node failures. What can the Developer do to address the load and resiliency requirements?

- A. Add a read replica instance.
- B. Migrate to a Memcached cluster.
- C. Migrate to an Amazon ElastiCache service cluster.
- D. Vertically scale the ElastiCache instance.

Answer: A

Explanation:

Reference: <https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/Replication.Redis.Groups.html>

NEW QUESTION 380

- (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 385

- (Exam Topic 2)

An application is using Amazon DynamoDB as its data store, and should be able to read 100 items per second as strongly consistent reads. Each item is 5 KB in size.

To what value should the table's provisioned read throughput be set?

- A. 50 read capacity units
- B. 100 read capacity units
- C. 200 read capacity units
- D. 500 read capacity unitsc

Answer: C

NEW QUESTION 390

- (Exam Topic 2)

A Developer has a stateful web server on-premises that is being migrated to AWS. The Developer must have greater elasticity in the new design.

How should the Developer re-factor the application to make it more elastic? (Choose two.)

- A. Use pessimistic concurrency on Amazon DynamoDB
- B. Use Amazon CloudFront with an Auto Scaling group
- C. Use Amazon CloudFront with an AWS Web Application Firewall
- D. Store session state data in an Amazon DynamoDB table
- E. Use an ELB with an Auto Scaling group

Answer: DE

NEW QUESTION 392

- (Exam Topic 2)

A company has a REST application comprised of an Amazon API Gateway and several AWS Lambda functions. A developer is responding to an alert that the API Gateway's HTTP response error rate has unexpectedly increased. The developer must determine must which Lambda function is malfunctioning.

Which method would help the developer make this determination while minimizing delays?

- A. Execute an Amazon Athena query against the API Gateway and Lambda execution logs.
- B. Execute an Amazon CloudWatch Logs Insights query against the API Gateway and Lambda execution logs.
- C. Download the API Gateway and Lambda execution logs from Amazon S3, and perform a line-by-line search against them.
- D. Download the API Gateway and Lambda execution logs from Amazon CloudWatch Events, and perform line-by-line search against them.

Answer: D

NEW QUESTION 396

- (Exam Topic 2)

A developer wants to send multi-value headers to an AWS Lambda function that is registered as a target with an Application Load Balancer (ALB).

What should the developer do to achieve this?

- A. Place the Lambda function and target group in the same account
- B. Send the request body to the Lambda function with a size less than 1 MB 0
- C. Include the Base64 encoding status status code, status description, and headers in the Lambda function

D. Enable the multi-value headers on the ALB

Answer: D

NEW QUESTION 398

- (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 403

- (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 408

- (Exam Topic 2)

A Developer wants to upload data to Amazon S3 and must encrypt the data in transit. Which of the following solutions will accomplish this task? (Choose two.)

- A. Set up hardware VPN tunnels to a VPC and access S3 through a VPC endpoint
- B. Set up Client-Side Encryption with an AWS KMS-Managed Customer Master Key
- C. Set up Server-Side Encryption with AWS KMS-Managed Keys
- D. Transfer the data over an SSL connection
- E. Set up Server-Side Encryption with S3-Managed Keys

Answer: BD

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/dev/UsingEncryption.html>

NEW QUESTION 409

- (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 413

- (Exam Topic 2)

A software company needs to make sure user-uploaded documents are securely stored in Amazon S3. The documents must be encrypted at rest in Amazon S3. The company does not want to manage the security infrastructure in-house, but the company still needs extra protection to ensure it has control over its encryption keys due to industry regulations.

Which encryption strategy should a developer use to meet these requirements?

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

Answer: D

NEW QUESTION 415

- (Exam Topic 2)

A Developer is trying to make API calls using SDK. The IAM user credentials used by the application require multi-factor authentication for all API calls. Which method the Developer use to access the multi-factor authentication protected API?

- A. GetFederationToken
- B. GetCallerIdentity
- C. GetSessionToken
- D. DecodeAuthorizationMessage

Answer: B

NEW QUESTION 417

- (Exam Topic 2)

A Developer needs to deploy an application running on AWS Fargate using Amazon ECS. The application has environment variables that must be passed to a container for the application to initialize

How should the environment variables be passed to the container?

- A. Define an array that includes the environment variables under the environment parameter within the service definition
- B. Define an array that includes the environment variables under the environment parameter within the task definition
- C. Define an array that includes the environment variables under the entrypoint parameter within the task definition
- D. Define in array that includes the environment variables under the entryPoint parameter within the service definition

Answer: B

NEW QUESTION 422

- (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 invoke 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 426

- (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:

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

NEW QUESTION 429

.....

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](#)