

Amazon-Web-Services

Exam Questions DVA-C01

AWS Certified Developer Associate Exam



NEW QUESTION 1

- (Exam Topic 4)

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

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

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

Answer: BD

NEW QUESTION 2

- (Exam Topic 4)

A developer is creating a Java application that will store sensitive data in an Amazon DynamoDB table. The data must be encrypted at all times. How can the developer meet this requirement?

- A. Enable encryption at rest by using an AWS Key Management Service (AWS KMS) AWS owned key for the DynamoDB table.
- B. Enable encryption at rest by using an AWS Key Management Service (AWS KMS) customer managed key for the DynamoDB table.
- C. Implement client-side encryption in the application code by using the DynamoDB Encryption Client.
- D. Use an HTTPS connection to encrypt data in transit.

Answer: C

Explanation:

Client-side and server-side encryption:

"DynamoDB Encryption Client" supports client-side encryption, where you encrypt your table data before you send it to DynamoDB. However, DynamoDB provides a server-side encryption at rest feature that transparently encrypts your table when it is persisted to disk and decrypts it when you access the table.

<https://docs.aws.amazon.com/dynamodb-encryption-client/latest/devguide/client-server-side.html>

NEW QUESTION 3

- (Exam Topic 4)

A developer is deploying a company's application to Amazon EC2 instances. The application generates gigabytes of data files each day. The files are rarely accessed, but the files must be available to the application's users within minutes of a request during the first year of storage. The company must retain the files for 7 years.

How can the developer implement the application to meet these requirements MOST cost-effectively?

- A. Store the files in an Amazon S3 bucket. Use the S3 Glacier Instant Retrieval storage class.
- B. Create an S3 Lifecycle policy to transition the files to the S3 Glacier Deep Archive storage class after 1 year.
- C. Store the files in an Amazon S3 bucket.
- D. Use the S3 Standard storage class.
- E. Create an S3 Lifecycle policy to transition the files to the S3 Glacier Flexible Retrieval storage class after 1 year.
- F. Store the files on an Amazon Elastic Block Store (Amazon EBS) volume.
- G. Use Amazon Data Lifecycle Manager (Amazon DLM) to create snapshots of the EBS volumes and to store those snapshots in Amazon S3.
- H. Store the files on an Amazon Elastic File System (Amazon EFS) mount.
- I. Configure EFS lifecycle management to transition the files to the EFS Standard-Infrequent Access (Standard-IA) storage class after 1 year.

Answer: A

Explanation:

Amazon S3 Glacier Instant Retrieval is an archive storage class that delivers the lowest-cost storage for long-lived data that is rarely accessed and requires retrieval in milliseconds. With S3 Glacier Instant Retrieval, you can save up to 68% on storage costs compared to using the S3 Standard-Infrequent Access (S3 Standard-IA) storage class, when your data is accessed once per quarter.

<https://aws.amazon.com/s3/storage-classes/glacier/instant-retrieval/>

NEW QUESTION 4

- (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 5

- (Exam Topic 4)

A company must encrypt sensitive data that the company will store in Amazon S3. A developer must retain total control over the company's AWS Key Management Service (AWS KMS) key and the company's data keys. The company currently uses an on-premises hardware security module (HSM) solution. The company wants to move its key management onto AWS.

Which solution will meet these requirements?

- A. Implement server-side encryption with AWS KMS managed keys (SSE-KMS). Use AWS CloudHSM to generate the KMS key and data keys to use with AWS KMS.
- B. Implement server-side encryption with customer-provided encryption keys (SSE-C). Use AWS CloudHSM to generate the KMS key and manage the data keys that the company will use to read and write objects to Amazon S3.
- C. Implement server-side encryption with Amazon S3 managed encryption keys (SSE-S3). Use AWS CloudHSM to generate the KMS key and manage the data keys that the company will use to read and write objects to Amazon S3.
- D. Implement server-side encryption with AWS KMS managed keys (SSE-KMS). Use the AWS KMS custom key store feature to manage the data key
- E. Then read or write objects to Amazon S3 as normal.

Answer: D

Explanation:

<https://docs.aws.amazon.com/cloudhsm/latest/userguide/best-practices.html> Q: Can other AWS services use CloudHSM to store and manage keys?

AWS services integrate with AWS Key Management Service, which in turn is integrated with AWS CloudHSM through the KMS custom key store feature. If you want to use the server-side encryption offered by many AWS services (such as EBS, S3, or Amazon RDS), you can do so by configuring a custom key store in AWS KMS.

NEW QUESTION 6

- (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 7

- (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 8

- (Exam Topic 4)

A developer is using AWS CodeDeploy to automate a company's application deployments to Amazon EC2. Which application specification file properties are required to ensure the software deployments do not fail?

(Select TWO.)

- A. The file must be a JSON-formatted file named appspec.json.
- B. The file must be a YAML-formatted file named appspec.yml
- C. The file must be stored in AWS CodeBuikJ and referenced from the application's source code.
- D. The file must be placed in the root of the directory structure of the application's source code.
- E. The file must be stored in Amazon S3 and referenced from the application's source code.

Answer: BD

NEW QUESTION 9

- (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 10

- (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 10

- (Exam Topic 4)

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

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

Answer: B

NEW QUESTION 11

- (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 16

- (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 18

- (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 22

- (Exam Topic 4)

A company hosts a monolithic application on Amazon EC2 instances. The company starts converting some features of the application to a serverless architecture by using Amazon API Gateway and AWS Lambda. After the migration, some users report problems with payment processing.

Upon inspection, a developer discovers that the Lambda function that calls the external payment API is taking longer than expected. Therefore, the API Gateway requests are timing out.

What should the developer do to resolve this issue in the serverless architecture?

- A. Use the EC2 instances to make the API calls to the payment API.
- B. Use Amazon Simple Queue Service (Amazon SQS) with API Gateway and the Lambda function to asynchronously call the payment API.
- C. Increase the API Gateway timeout duration to match the payment API time.
- D. Increase the Lambda function's memory to increase the network bandwidth and increase the speed of the payment API calls.

Answer: B

NEW QUESTION 23

- (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 meet 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 pre-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 25

- (Exam Topic 4)

A developer notices timeouts from the AWS CLI when the developer runs list commands. What should the developer do to avoid these timeouts?

- A. Use the --page-size parameter to request a smaller number of items.
- B. Use shorthand syntax to separate the list by a single space.
- C. Use the --yaml-stream output for faster viewing of large datasets.
- D. Use quotation marks around strings to enclose data structure.

Answer: A

NEW QUESTION 30

- (Exam Topic 4)

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

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

Which solution should the developer implement to meet these requirements?

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

- O. Extract the user name from the given name claim and the user's universally unique identifier (UUID) from the iss claim inside the access token
- P. Use the UUID as the key to store and retrieve the previous login information from the database.

Answer: C

NEW QUESTION 34

- (Exam Topic 4)

A developer is creating a serverless orchestrator that performs a series of steps to process incoming IoT data. The orchestrator transforms the data, performs a series of calculations, and stores the results in Amazon DynamoDB. The entire process is completed in less than a minute. The orchestrator must process tens of thousands of transactions each second. The orchestrator must not miss a transaction and must process each transaction at least once.

Which solution will meet these requirements MOST cost-effectively?

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

Answer: D

NEW QUESTION 37

- (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 42

- (Exam Topic 4)

A developer is writing an application that stores data in an Amazon DynamoDB table by using the PutItem API operation. The table has a partition key of streamID and has a sort key of seqID. The developer needs to make sure that the PutItem invocation does not overwrite the existing partition key and sort key.

Which condition expression will maintain the uniqueness of the partition key and the sort key?

A)

```
condition = 'attribute_not_exists(streamID)
            AND
            attribute_not_exists(seqID) '
```

B)

```
condition = 'attribute_not_exists(PARTITION)
            AND
            attribute_not_exists(SORT) '
```

C)

```
condition = 'attribute_exists(streamID)
            AND
            attribute_exists(seqID) '
```

D)

```
condition = 'attribute_exists(PARTITION)
            AND
            attribute_exists(SORT) '
```

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

Answer: A

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Expressions.ConditionExpressions.html>

NEW QUESTION 43

- (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 47

- (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 51

- (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 55

- (Exam Topic 4)

A developer is writing an AWS Lambda function. The Lambda function needs to access items that are stored in an Amazon DynamoDB table.

What is the MOST secure way to configure this access for the Lambda function?

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

Answer: A

NEW QUESTION 59

- (Exam Topic 4)

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

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

- A. Use the `LocalCryptoMaterialsCache` feature of the AWS Encryption SDK encryption library.
- B. Call the AWS KMS `Encrypt` operation directly to allow AWS KMS to encrypt the data.
- C. Create a case in the AWS Support Center to increase the quota for the account.

- D. Use Amazon Simple Queue Service (Amazon SQS) to queue the requests to AWS KMS.
- E. Switch to an AWS KMS custom key store.

Answer: AC

Explanation:

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

NEW QUESTION 61

- (Exam Topic 4)

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

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

Answer: D

Explanation:

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

NEW QUESTION 66

- (Exam Topic 4)

A company is developing a microservice that will manage customer account data in an Amazon DynamoDB table. Insert, update, and delete requests will be rare. Read traffic will be heavy. The company must have the ability to access customer data quickly by using a customer ID. The microservice can tolerate stale data. Which solution will meet these requirements with the FEWEST possible read capacity units (RCUs)?

- A. Read the table by using eventually consistent reads.
- B. Read the table by using strongly consistent reads.
- C. Read the table by using transactional reads.
- D. Read the table by using strongly consistent PartiQL queries.

Answer: A

Explanation:

Key points: "Read heavy", "access data quickly", "can tolerate stale data" To achieve: "FEWEST" possible (RCUs)

For items up to 4 KB in size, one RCU can perform one strongly consistent read request per second. For items up to 4 KB in size, one RCU can perform two eventually consistent read requests per second. Transactional read requests require two RCUs to perform one read per second for items up to 4 KB. For example, a strongly consistent read of an 8 KB item would require two RCUs, an eventually consistent read of an 8 KB item would require one RCU, and a transactional read of an 8 KB item would require four RCUs. <https://aws.amazon.com/dynamodb/pricing/provisioned/>

NEW QUESTION 68

- (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 70

- (Exam Topic 4)

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

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 75

- (Exam Topic 4)

A company has a front-end application that runs on four Amazon EC2 instances behind an Elastic Load Balancer (ELB) in a production environment that is provisioned by AWS Elastic Beanstalk. A developer needs to deploy and test new application code while updating the Elastic Beanstalk platform from the current version to a newer version of Node.js. The solution must result in zero downtime for the application.

Which solution meets these requirements?

- A. Clone the production environment to a different platform versio
- B. Deploy the new application code, and test i
- C. Swap the environment URLs upon verification.
- D. Deploy the new application code in an all-at-once deployment to the existing EC2 instance
- E. Test the cod
- F. Redeploy the previous code if verification fails.
- G. Perform an immutable update to deploy the new application code to new EC2 instance
- H. Serve traffic to the new instances after they pass health checks
- I. Use a rolling deployment for the new application cod
- J. Apply the code to a subset of EC2 instances until the tests pas
- K. Redeploy the previous code if the tests fail.

Answer: D

NEW QUESTION 80

- (Exam Topic 4)

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

Which HTTP header should the developer use for this analysis?

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

Answer: C

Explanation:

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

NEW QUESTION 82

- (Exam Topic 4)

A developer is creating a photo website. Amazon Route 53 hosts the website's domain. The developer wants to store the application code and images in an Amazon S3 bucket. The developer also wants to use Amazon CloudFront to deliver the images to users.

The developer has created the S3 bucket and a CloudFront distribution. The developer wants the images to be accessed only through the website's domain. Users must not use the S3 URLs.

Which solution will meet these requirements?

- A. Create a CloudFront origin access identity (OAI). Associate the OAI with the CloudFront distribution. Modify the S3 bucket policy to allow access from only the OA
- B. Create an alias in Route 53 that points the website domain to the S3 bucket.
- C. Create a CloudFront origin access identity (OAI). Associate the OAI with the CloudFront distribution. Modify the S3 bucket policy to allow access from only the OA
- D. Update the Route 53 records to point the website domain to the CloudFront domain name.
- E. Block public access in the S3 bucket polic
- F. Configure CloudFront to use the S3 bucket endpoint. Create an alias in Route 53 that points the website domain to the S3 bucket.
- G. Block public access in the S3 bucket polic
- H. Configure CloudFront to use the S3 bucket endpoint. Create an alias in Route 53 that points the website domain to the CloudFront domain name.

Answer: A

NEW QUESTION 83

- (Exam Topic 4)

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

Which solution will meet these requirements?

- A. Send the order ID to an Amazon Simple Notification Service (Amazon SNS) FIFO topic that fans out to one Amazon Simple Queue Service (Amazon SQS) FIFO queue for inventory management and another SQS FIFO queue for payment processing.
- B. Change the Lambda function that generates the order ID to initiate the Lambda function for inventory managemen

- C. Then initiate the Lambda function for payment processing.
- D. Send the order ID to an Amazon Simple Notification Service (Amazon SNS) topic.
- E. Subscribe the Lambda functions for inventory management and payment processing to the topic.
- F. Deliver the order ID to an Amazon Simple Queue Service (Amazon SQS) queue.
- G. Configure the Lambda functions for inventory management and payment processing to poll the queue.

Answer: A

Explanation:

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

NEW QUESTION 84

- (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 89

- (Exam Topic 4)

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

How can the developer ensure that the application processes each message exactly once?

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

Answer: C

Explanation:

<https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/FIFO-queues-moving.html> Moving from a standard queue to a FIFO queue:

If you have an existing application that uses standard queues and you want to take advantage of the ordering or exactly-once processing features of FIFO queues, you need to configure the queue and your application correctly.

Note:

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

NEW QUESTION 93

- (Exam Topic 4)

A business intelligence application runs on Amazon Elastic Container Service (Amazon ECS) on AWS Fargate. Application-level audits require a searchable log of all API calls from users to the application. The application's developers must store the logs centrally on AWS.

Which solution will meet these requirements?

- A. Install the Amazon CloudWatch agent on the Amazon EC2 host that runs Fargate.
- B. Configure the awslogs log driver in the ECS task definition.
- C. Configure AWS CloudTrail for the ECS containers.
- D. Install the ECS logs collector on the ECS hosts.

Answer: B

Explanation:

https://docs.aws.amazon.com/AmazonECS/latest/developerguide/using_awslogs.html

Configuring the awslogs log driver in the ECS task definition will allow the application to store the logs centrally on AWS. The awslogs log driver sends logs to Amazon CloudWatch Logs, which is a managed service that provides search and analysis of log data. This solution will meet the requirements of storing the logs centrally on AWS and making them searchable. Installing the Amazon CloudWatch agent on the Amazon EC2 host or installing the ECS logs collector on the ECS hosts will not work because the application is running on AWS Fargate and not on Amazon EC2. AWS CloudTrail is not a suitable solution because it is used to record API calls made to AWS services, not application-level API calls.

NEW QUESTION 97

- (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 98

- (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 100

- (Exam Topic 4)

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

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

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

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

Answer: B

NEW QUESTION 101

- (Exam Topic 4)

A company is using continuous integration/continuous delivery (CI/CD) systems. A team that runs 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 103

- (Exam Topic 4)

A company has an online order website that uses Amazon DynamoDB to store item inventory. A sample of the inventory object is as follows:

```
{
  "Id": { "N": "456"},
  "Price": { "N": "650"},
  "ProductCategory": { "S": "Sporting Goods" }
}
```

A developer needs to reduce all inventory prices by 100 as long as the resulting price would not be less than 500. What should the developer do to make this change with the LEAST number of calls to DynamoDB?

- A. Perform a DynamoDB Query operation with the following condition expression: `Price >= 500`.
- B. If the price is ≥ 600 , perform an UpdateItem operation to update the price.
- C. Perform a DynamoDB UpdateItem operation with a condition expression of `Price >= 600`.
- D. Perform a DynamoDB UpdateItem operation with a condition expression of `ProductCategory IN {"S": "Sporting Goods"} and Price = 600`.

Answer: C

NEW QUESTION 107

- (Exam Topic 4)

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

Which action will meet this requirement?

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

Answer: C

Explanation:

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

NEW QUESTION 109

- (Exam Topic 4)

What are the MINIMUM properties required in the resources section of the AppSpace file for CodeDeploy to deploy the ECS service successfully?

- A. name, alias, currentversion, and targetversion
- B. TaskDefinition, ContainerName, and PlatformVersion
- C. TaskDefinitionContainerName, ContainerPort
- D. name, Currentversion, NetworkConfiguration, and Platform Version

Answer: A

NEW QUESTION 111

- (Exam Topic 4)

A developer is building varKHJS microservices for an application that will run on Amazon EC2 instances. The developer needs to monitor the end-to-end view of the requests between the microservices and debug any issues in the various microservices.

What should the developer do to accomplish these tasks?

- A. Use Amazon CloudWatch to aggregate the microservices' logs and metrics, and build the monitoring dashboard
- B. Use AWS CloudTrail to aggregate the microservices' logs and metrics, and build the monitoring dashboard
- C. Use the AWS X-Ray SDK to add instrumentation in all the microservices, and monitor using the X-Ray service map.
- D. Use AWS Health to monitor the health of all the microservices

Answer: C

Explanation:

<https://docs.aws.amazon.com/whitepapers/latest/microservices-on-aws/distributed-monitoring.html#distributed-t> <https://aws.amazon.com/xray/>

NEW QUESTION 113

- (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 118

- (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 121

- (Exam Topic 4)

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

- A. Amazon DynamoDB
- B. Amazon Cognito
- C. Amazon ElastiCache
- D. Application Load Balancer
- E. Amazon Simple Queue Service (Amazon SQS)

Answer: AC

NEW QUESTION 122

- (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 124

- (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..... 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

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/lambda-vpc-parameter-store/>

NEW QUESTION 125

- (Exam Topic 4)

A security review for a software company's application infrastructure shows that there is no test coverage in any of the company's deployment pipelines. A developer must fix this issue as soon as possible. The company has been integrating the AWS Cloud Development Kit (AWS CDK) into the deployment process. However, much of the pipeline still uses AWS CloudFormation templates. The developer needs to add test coverage to all the deployment code.

Which solution will meet these requirements with the LEAST amount of configuration?

- A. Write unit tests by using the AWS CDK assertions modul
- B. Create CloudFormation template instances by using the module's Template class for the existing CloudFormation templates and the module's Capture class for the CDK stacks.
- C. Write unit tests by using the AWS CDK assertions modul
- D. Create CloudFormation template instances by using the module's Template.fromStringO method for the existing CloudFormation templates and the module's Template fromStackQ method for the CDK stacks.
- E. Convert the CloudFormation templates into CDK stacks by using the AWS CDK CfnInclude construct. Write unit tests against the templates by using CloudFormation rule assertions.
- F. Convert the CDK stacks into CloudFormation templates by using the AWS CDK CfnInclude construct Write unit tests against the templates by using CloudFormation rule assertions

Answer: A

NEW QUESTION 128

- (Exam Topic 4)

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

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

Answer: B

NEW QUESTION 133

- (Exam Topic 4)

A developer is creating a new batch application that will run on an Amazon EC2 instance. The application requires read access to an Amazon S3 bucket. The

developer needs to follow security best practices to grant S3 read access to the application.
Which solution meets these requirements?

- A. Add the permissions to an IAM policy
- B. Attach the policy to a role Attach the role to the EC2 instance profile.
- C. Add the permissions inline to an IAM group
- D. Attach the group to the EC2 instance profile.
- E. Add the permissions to an IAM policy
- F. Attach the policy to a user Attach the user to the EC2 instance profile.
- G. Add the permissions to an IAM policy
- H. Use IAM web identity federation to access the S3 bucket with the policy

Answer: A

NEW QUESTION 138

- (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 142

- (Exam Topic 3)

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

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

Answer: B

NEW QUESTION 145

- (Exam Topic 3)

An application runs on multiple EC2 instances behind an ELB. Where is the session data best written so that it can be served reliably across multiple requests?

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

Answer: A

NEW QUESTION 149

- (Exam Topic 3)

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

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

Answer: B

NEW QUESTION 153

- (Exam Topic 3)

A developer from AnyCompany's AWS account needs access to the Example Corp AWS account. AnyCompany uses an identity provider that is compatible with OpenID Connect.

What is the MOST secure way for Example Corp to allow developer access?

- A. Create a cross-account role and call the AssumeRole API operation
- B. Create a user in the Example Corp account and provide the access keys
- C. Create a user in the Example Corp account and provide the credentials

D. Create a cross-account role and call the AssumeRoleWithWebIdentity API operation

Answer: B

NEW QUESTION 158

- (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 instanc
- F. Create a new role that developer can assume when termination is necessary.

Answer: A

NEW QUESTION 160

- (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 165

- (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 cod
- 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 166

- (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 m 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 169

- (Exam Topic 3)

An application development team decides to use AWS X Ray to monitor application code to analyze performance and perform r 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 172

- (Exam Topic 3)

A developer must increase read performance from an unencrypted Amazon S3 bucket. The application requires 100.000 read requests each second Cost-effectiveness is a priority. What would be the SIMPLEST approach to implement these requirements?

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

Answer: A

NEW QUESTION 173

- (Exam Topic 3)

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

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

Answer: A

NEW QUESTION 178

- (Exam Topic 3)

A company is using AWS CloudFormation templates to deploy AWS resources. The company needs to update one of its AWS CloudFormation stacks What can the company do to find out how the changes will impact the resources that are running?

- A. Investigate the change sets
- B. Investigate the stack policies
- C. Investigate the Metadata section.
- D. Investigate the Resources section.

Answer: A

NEW QUESTION 182

- (Exam Topic 3)

An IAM role is attached to an Amazon EC2 instance that explicitly denies access to all Amazon S3 API actions. The EC2 instance credentials file specifies the IAM access key and secret access key, which allow full administrative access.

Given that multiple modes of IAM access are present for this EC2 instance, which of the following is correct?

- A. The EC2 instance will only be able to list the S3 buckets
- B. The EC2 instance will only be able to list the contents of one S3 bucket at a time
- C. The EC2 instance will be able to perform all actions on any S3 bucket
- D. The EC2 instance will not be able to perform any S3 action on any S3 bucket.

Answer: C

NEW QUESTION 186

- (Exam Topic 3)

A photo sharing website gets millions of new images every week The images are stored in Amazon S3 under a formatted date prefix A developer wants to move images to a few S3 buckets for analysis and further processing Images are not required to be moved in real time
What is the MOST efficient method for performing this task?

- A. Use S3 PutObject events to Invoke AWS Lambda Then Lambda will copy the files to the other objects
- B. Create an AWS Lambda function that will pull a day of Images from the origin bucket and copy them to the other buckets.
- C. Use S3 Batch Operations to create jobs for images to be copied to each Individual bucket.
- D. Use Amazon EC2 to batch pull images from multiple days and copy them to the other buckets

Answer: D

NEW QUESTION 189

- (Exam Topic 3)

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

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

Which of the following solution supports these requirements?

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

Answer: C

NEW QUESTION 191

- (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 193

- (Exam Topic 3)

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

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

Answer: D

NEW QUESTION 198

- (Exam Topic 3)

A company has deployed an application on AWS Elastic Beanstalk. The company has configured the Auto Scaling group that is associated with the Elastic Beanstalk environment to have five Amazon EC2 instances. If the capacity is fewer than four EC2 instances during the deployment, application performance degrades. The company is using the all-at-once deployment policy. What is the MOST cost-effective way to solve the deployment issue?

- A. Change the Auto Scaling group to six desired instances
- B. Change the deployment policy to traffic splitting. Specify an evaluation time of 1 hour.
- C. Change the deployment policy to rolling with additional batch. Specify a batch size of 1.
- D. Change the deployment policy to rolling. Specify a batch size of 2.

Answer: C

NEW QUESTION 199

- (Exam Topic 3)

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

- A. Use AWS CodeDeploy and point it to the local storage to directly deploy a bundle in a zip format
- B. Use a tar.gz format
- C. Upload the bundle to an Amazon S3 bucket and specify the S3 location when doing a deployment using AWS CodeDeploy
- D. Create a repository using AWS CodeCommit to automatically trigger a deployment to the EC2 instances
- E. Use AWS CodeBuild to automatically deploy the latest build to the latest EC2 instances

Answer: B

NEW QUESTION 203

- (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 207

- (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 210

- (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 211

- (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 addressed?

- 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 214

- (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 218

- (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 220

- (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 225

- (Exam Topic 3)

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

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

Answer: DE

NEW QUESTION 227

- (Exam Topic 3)

A company has an internal website that gives users the ability to access contract data that is stored in an Amazon RDS DB instance. The number of contracts has increased and several users have reported slow retrieval of the contract data. The company wants to set up a cache to improve the latency. A developer must create a solution that ensures data resiliency. The data must be encrypted and must be partitioned by department. Which solution will meet these requirements?

- A. Amazon ElastiCache for Memcached with cluster mode enabled
- B. Amazon ElastiCache for Redis with cluster mode enabled
- C. Amazon ElastiCache for Redis with cluster mode disabled
- D. Amazon ElastiCache for Memcached with cluster mode disabled

Answer: C

NEW QUESTION 229

- (Exam Topic 3)

A developer is working on a serverless application. The application uses Amazon API Gateway, AWS Lambda functions that are written in Python, and Amazon DynamoDB.

Which combination of steps should the developer take so that the Lambda functions can be debugged in the event of application failures? (Select TWO)

- A. Configure an AWS CloudTrail trail to deliver log files to an Amazon S3 bucket
- B. Ensure that the Lambda functions write log messages to stdout and stderr
- C. Enable an AWS CloudTrail trail for the Lambda function
- D. Ensure that the execution role for the Lambda function has access to write to Amazon CloudWatch Logs.
- E. Use the Amazon CloudWatch metric for Lambda errors to create a CloudWatch alarm.

Answer: DE

NEW QUESTION 233

- (Exam Topic 3)

A company hosts a microservices application that uses Amazon API Gateway, AWS Lambda, Amazon Simple Queue Service (Amazon SQS), and Amazon DynamoDB. One of the Lambda functions adds messages to an SQS FIFO queue.

When a developer checks the application logs, the developer finds a few duplicated items in a DynamoDB table. The items were inserted by another polling function that processes messages from the queue.

What is the MOST likely cause of this issue?

- A. Write operations on the DynamoDB table are being throttled
- B. The SQS queue delivered the message to the function more than once
- C. API Gateway duplicated the message in the SQS queue
- D. The polling function timeout is greater than the queue visibility timeout

Answer: B

NEW QUESTION 238

- (Exam Topic 3)

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

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

Answer: C

NEW QUESTION 243

- (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 245

- (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 246

- (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 250

- (Exam Topic 3)

A developer supports an application that accesses data in an Amazon DynamoDB table. One of the item attributes is `expirationDate` in the timestamp format. The application uses this attribute to find items, archive them, and remove them from the table based on the timestamp value.

The application will be decommissioned soon, and the developer must find another way to implement this functionality. The developer needs a solution that will require the least amount of code to write.

Which solution will meet these requirements?

- A. Enable TTL on the `expirationDate` attribute in the table.
- B. Create a DynamoDB stream.
- C. Create an AWS Lambda function to process the deleted items.
- D. Create a DynamoDB trigger for the Lambda function.
- E. Create two AWS Lambda functions: one to delete the items and one to process the items. Create a DynamoDB stream. Use the `DeleteItem` API operation to delete the items based on the `expirationDate` attribute. Use the `GetRecords` API operation to get the items from the DynamoDB stream and process them.
- F. Create two AWS Lambda functions: one to delete the items and one to process the items. Create an Amazon EventBridge (Amazon CloudWatch Events) scheduled rule to invoke the Lambda functions. Use the `DeleteItem` API operation to delete the items based on the `expirationDate` attribute. Use the `GetRecords` API operation to get the items from the DynamoDB table and process them.
- G. Enable TTL on the `expirationDate` attribute in the table. Specify an Amazon Simple Queue Service (Amazon SQS) dead-letter queue as the target to delete the items. Create an AWS Lambda function to process the items.

Answer: C

NEW QUESTION 252

- (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 257

- (Exam Topic 3)

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

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

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

Answer: BD

NEW QUESTION 261

- (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 266

- (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 267

- (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 270

- (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 271

- (Exam Topic 3)

A team deployed an AWS CloudFormation template to update a stack that already included an Amazon RDS DB instance. However, before the deployment of the update, the team changed the name of the DB instance on the template by mistake. The `DeletionPolicy` attribute for all resources was not changed from the default values.

What will be the result of this mistake?

- A. AWS CloudFormation will create a new database and delete the old one.
- B. AWS CloudFormation will create a new database and keep the old one.
- C. AWS CloudFormation will overwrite the existing database and rename it.
- D. AWS CloudFormation will leave the existing database and will not create a new one.

Answer: A

NEW QUESTION 273

- (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 275

- (Exam Topic 3)

A developer is working on a web application that runs on Amazon Elastic Container Service (Amazon ECS) and uses an Amazon DynamoDB table to store data. The application performs a large number of read requests against a small set of the table data.

How can the developer improve the performance of these requests? (Select TWO)

- A. Create an Amazon ElastiCache cluster. Configure the application to cache data in the cluster.
- B. Create a DynamoDB Accelerator (DAX) cluster. Configure the application to use the DAX cluster for DynamoDB requests.
- C. Configure the application to make strongly consistent read requests against the DynamoDB table.
- D. Increase the read capacity of the DynamoDB table.
- E. Enable DynamoDB adaptive capacity.

Answer: AD

NEW QUESTION 279

- (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 281

- (Exam Topic 3)

A developer has created a REST API using Amazon API Gateway. The developer wants to log who and how each caller accesses the API. The developer also wants to control how long the logs are kept. What should the developer do to meet these requirements?

- A. Enable API Gateway execution logging. Delete old logs using API Gateway retention settings.
- B. Enable API Gateway access logs. Use Amazon CloudWatch retention settings to delete old logs.
- C. Enable detailed Amazon CloudWatch metrics. Delete old logs with a recurring AWS Lambda function.
- D. Create and use API Gateway usage plan.
- E. Delete old logs with a recurring AWS Lambda function.

Answer: A

NEW QUESTION 282

- (Exam Topic 3)

A developer is writing a new AWS Serverless Application Model (AWS SAM) template with a new AWS Lambda function. The Lambda function runs complex code. The developer wants to test the Lambda function with more CPU power.

What should the developer do to meet this requirement?

- A. Increase the runtime engine version.
- B. Increase the timeout.
- C. Increase the number of Lambda layers.
- D. Increase the memory.

Answer: D

NEW QUESTION 283

- (Exam Topic 3)

A developer has built an application using Amazon Cognito for authentication and authorization. After a user is successfully logged in to the application, the application creates a user record in an Amazon DynamoDB table.

What is the correct flow to authenticate the user and create a record in the DynamoDB table?

- A. Authenticate and get a token from an Amazon Cognito user pool.
- B. Use the token to access DynamoDB.
- C. Authenticate and get a token from an Amazon Cognito identity pool.
- D. Use the token to access DynamoDB.
- E. Authenticate and get a token from an Amazon Cognito user pool. Exchange the token for AWS credentials with an Amazon Cognito identity pool.
- F. Use the credential to access DynamoDB.
- G. Authenticate and get a token from an Amazon Cognito identity pool.
- H. Exchange the token for AWS credentials with an Amazon Cognito user pool.
- I. Use the credentials to access DynamoDB.

Answer: B

NEW QUESTION 286

- (Exam Topic 3)

A developer is building a highly secure healthcare application using .NET. The application requires writing temporary data to /tmp storage on an AWS Lambda function. How should the developer encrypt this data?

- A. Enable Amazon EBS volume encryption with an AWS KMS .NET configuration so that all storage attached to the Lambda function is encrypted.
- B. Set up the Lambda function with a role and key policy to access an AWS KMS CMK. Use the CMK to generate a data key used to encrypt all data prior to writing to /tmp storage.
- C. Use OpenSSL to generate a symmetric encryption key on Lambda startup. Use this key to encrypt the data prior to writing to /tmp.
- D. Use an on-premises hardware security module (HSM) to generate keys where the Lambda function requests a data key from the HSM and uses that to encrypt data on all requests to the function.

Answer: D

NEW QUESTION 288

- (Exam Topic 3)

A company has three AWS Lambda functions that are written in Node.js. The Lambda functions include a mix of custom code and open-source modules. When bugs are occasionally detected in the open-source modules, all three Lambda functions must be patched.

What is the MOST operationally efficient solution to deploy a patched open-source library for all three Lambda functions?

- A. Create a custom AWS CloudFormation public registry extension Reference a GitHub repository that hosts the open-source modules in the extension Configure CloudFormation to scan the repository once each day Write an AWS Serverless Application Model (AWS SAM) template to redeploy the three Lambda functions upon a scan notification change.
- B. Create an Amazon CloudFront distribution with an Amazon S3 bucket as the origin Upload the patched modules to Amazon S3 when needed Modify each Lambda function to download the patched modules from the CloudFront distribution during the cold start.
- C. Launch an Amazon EC2 instance Host a private open-source module registry on the EC2 instance Upload the modified open-source modules to the private registry when needed
- D. Modify each Lambda function deployment script to download the modules from the private registry Redeploy the three new Lambda functions.
- E. Create a Lambda layer with the open-source modules Modify all three Lambda functions to depend on the layer Remove the open-source modules from each Lambda function Patch the Lambda layer with the modified open-source modules when needed Update the Lambda functions to reference the new layer version

Answer: D

NEW QUESTION 290

- (Exam Topic 3)

A company is using continuous integration/continuous deliver (CI/CD) system. A developer must automate the deployment of an application software package to Amazon EC2 instances and virtual servers that run on premises.

Which AWS services 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 292

- (Exam Topic 3)

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

What is the MOST likely cause of this error?

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

Answer: D

NEW QUESTION 296

- (Exam Topic 3)

A company process incoming documents from an Amazon S3 bucket. Users upload documents to an S3 bucket using a web user interface. Upon receiving files in S3, an AWS Lambda function is invoked to process the files, but the Lambda function times out intermittently.

If the Lambda function is configured with the default settings, what will happen to the S3 event when there is a timeout exception?

- A. Notification of a failed S3 event is sent as an email through Amazon SNS.
- B. The S3 event is sent to the default Dead Letter Queue.
- C. The S3 event is processed until it is successful.
- D. The S3 event is discarded after the event is retried twice.

Answer: D

NEW QUESTION 301

- (Exam Topic 3)

A physician's office management application requires that all data in transit between an EC2 instance and an Amazon EBS volume be encrypted

Which of the following techniques fulfills this requirement? (Select TWO)

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

Answer: AD

NEW QUESTION 302

- (Exam Topic 3)

A developer must extend an existing application that is based on the AWS Services Application Model (AWS SAM). The developer has used the AWS SAM CLI to create the project. The project contains different AWS Lambda functions.

Which combination of commands must the developer use to redeploy the AWS SAM application (Select TWO.)

- A. Sam init
- B. Sam validate
- C. Sam build
- D. Sam deploy
- E. Sam publish

Answer: AD

NEW QUESTION 303

- (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 307

- (Exam Topic 3)

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

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

Answer: C

NEW QUESTION 309

- (Exam Topic 3)

A developer must allow guest users without logins to access an Amazon Cognito-enabled site to view files stored within an Amazon S3 bucket How should the developer meet these requirements'?

- A. Create a blank user ID in a user pool, add to the user group, and grant access to AWS resources
- B. Create a new identity pool, enable access to unauthenticated identities and grant access to AWS resources
- C. Create a new user pool, enable access to unauthenticated identities, and grant access to AWS resources.
- D. Create a new user pool disable authentication access, and grant access to AWS resources

Answer: C

NEW QUESTION 314

- (Exam Topic 3)

A company is using Amazon API Gateway to manage its public-facing API. The CISO requires that the APIs be used by test account users only. What is the MOST secure way to restrict API access to users of this particular AWS account?

- A. Client-side SSL certificates for authentication
- B. API Gateway resource policies
- C. Cross-origin resource sharing (CORS)
- D. Usage plans

Answer: D

NEW QUESTION 318

- (Exam Topic 3)

A developer has written an application that uses Amazon API Gateway and AWS Lambda The developer needs to configure the application so that the developer can visualize the application's components and Identify performance bottlenecks What should the developer do to meet these requirements?

- A. Enable AWS X-Ray tracing on the API Gateway stage
- B. Enable AWS X-Ray tracing on the API Gateway methods
- C. Enable Amazon CloudWatch Logs for API Gateway
- D. Enable Amazon CloudWatch Logs for Lambda

Answer: A

NEW QUESTION 320

- (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 325

- (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 326

- (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 330

- (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 table
- 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 331

- (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 333

- (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 334

- (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 335

- (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 338

- (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 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 341

- (Exam Topic 3)

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

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

Answer: B

NEW QUESTION 344

- (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 348

- (Exam Topic 2)

A company has an AWS CloudFormation template that is stored as a single file. The template is able to launch and create a full infrastructure stack. Which best practice would increase the maintainability of the template?

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

Answer: A

NEW QUESTION 351

- (Exam Topic 2)

A gaming company is developing a mobile game application for iOS® and Android® platforms. This mobile game securely stores user data locally on the device. The company wants to allow users to use multiple device for the game, which requires user data synchronization across device. Which service should be used to synchronize user data across devices without the need to create a backend application?

- A. AWS Lambda
- B. Amazon S3
- C. Amazon DynamoDB
- D. Amazon Cognito

Answer: D

NEW QUESTION 355

- (Exam Topic 2)

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 360

- (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 364

- (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 366

- (Exam Topic 2)

A Developer is storing sensitive documents in Amazon S3 that will require encryption at rest. The encryption keys must be rotated annually, at least.

What is the easiest way to achieve this?

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

Answer: C

Explanation:

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

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

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

NEW QUESTION 367

- (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 368

- (Exam Topic 2)

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

How can DynamoDB costs be minimized while maximizing application performance?

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

Answer: C

Explanation:

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

NEW QUESTION 370

- (Exam Topic 2)

An application uses Lambda functions to extract metadata from files uploaded to an S3 bucket; the metadata is stored in Amazon DynamoDB. The application starts behaving unexpectedly, and the Developer wants to examine the logs of the Lambda function code for errors.

Based on this system configuration, where would the Developer find the logs?

- A. Amazon S3
- B. AWS CloudTrail
- C. Amazon CloudWatch
- D. Amazon DynamoDB

Answer: C

NEW QUESTION 371

- (Exam Topic 2)

A stock market monitoring application uses Amazon Kinesis for data ingestion. During simulated tests of peak data rates, the Kinesis stream cannot keep up with the incoming data.

What step will allow Kinesis to accommodate the traffic during peak hours?

- A. Install the Kinesis Producer Library (KPL) for ingesting data into the stream.
- B. Reduce the data retention period to allow for more data ingestion using DecreaseStreamRetentionPeriod.
- C. Increase the shard count of the stream using UpdateShardCount.
- D. Ingest multiple records into the stream in a single call using PutRecords.

Answer: C

Explanation:

<https://docs.aws.amazon.com/streams/latest/dev/developing-producers-with-kpl.html>

NEW QUESTION 372

- (Exam Topic 2)

A company needs a new REST API that can return information about the contents of an Amazon S3 bucket, such as a count of the objects stored in it. The company has decided that the new API should be written as a microservice using AWS Lambda and Amazon API Gateway.

How should the Developer ensure that the microservice has the necessary access to the Amazon S3 bucket, while adhering to security best practices?

- A. Create an IAM user that has permissions to access the Amazon S3 bucket, and store the IAM user credentials in the Lambda function source code.
- B. Create an IAM role that has permissions to access the Amazon S3 bucket and assign it to the Lambda function as its execution role.
- C. Create an Amazon S3 bucket policy that specifies the Lambda service as its principal and assign it to the Amazon S3 bucket.
- D. Create an IAM role, attach the AmazonS3FullAccess managed policy to it, and assign the role to the Lambda function as its execution role.

Answer: B

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/lambda-execution-role-s3-bucket/>

NEW QUESTION 377

- (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 378

- (Exam Topic 2)

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

What should the development team do to meet these requirements?

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

Answer: C

NEW QUESTION 383

- (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-functions/> 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 385

.....

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](#)