

Exam Questions AWS-Certified-Developer-Associate

Amazon AWS Certified Developer - Associate

<https://www.2passeasy.com/dumps/AWS-Certified-Developer-Associate/>



NEW QUESTION 1

- (Exam Topic 1)

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

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

Answer: C

NEW QUESTION 2

- (Exam Topic 1)

A Developer has been asked to make changes to the source code of an AWS Lambda function. The function is managed using an AWS CloudFormation template. The template is configured to load the source code from an Amazon S3 bucket. The Developer manually created a .ZIP file deployment package containing the changes and put the file into the correct location on Amazon S3. When the function is invoked, the code changes have not been applied. What step is required to update the function with the changes?

- A. Delete the .ZIP file on S3, and re-upload by using a different object key name.
- B. Update the CloudFormation stack with the correct values for the function code properties S3Bucket, S3Key, or S3ObjectVersion.
- C. Ensure that the function source code is base64-encoded before uploading the deployment package to S3.
- D. Modify the execution role of the Lambda function to allow S3 access permission to the deployment package .ZIP file.

Answer: B

Explanation:

Changes to a deployment package in Amazon S3 are not detected automatically during stack updates. To update the function code, change the object key or version in the template.

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-properties-lambda-function-code.htm>

NEW QUESTION 3

- (Exam Topic 1)

A Developer must re-implement the business logic for an order fulfilment system. The business logic has to make requests to multiple vendors to decide where to purchase an item. The whole process can take up to a week to complete.

What is the MOST efficient and SIMPLEST way to implement a system that meets these requirements?

- A. Use AWS Step Functions to execute parallel Lambda functions, and join the results.
- B. Create an AWS SQS for each vendor, poll the queue from a worker instance, and join the results.
- C. Use AWS Lambda to asynchronously call a Lambda function for each vendor, and join the results.
- D. Use Amazon CloudWatch Events to orchestrate the Lambda functions.

Answer: A

Explanation:

<https://aws.amazon.com/step-functions/>

NEW QUESTION 4

- (Exam Topic 1)

A company has three different environments: Development, QA, and Production. The company wants to deploy its code first in the Development environment, then QA, and then Production.

Which AWS service can be used to meet this requirement?

- A. Use AWS CodeCommit to create multiple repositories to deploy the application.
- B. Use AWS CodeBuild to create, configure, and deploy multiple build application projects.
- C. Use AWS Data Pipeline to create multiple data pipeline provisions to deploy the application.
- D. Use AWS CodeDeploy to create multiple deployment groups.

Answer: D

Explanation:

<https://docs.aws.amazon.com/codedeploy/latest/userguide/deployment-groups.html>

"You can associate more than one deployment group with an application in CodeDeploy. This makes it possible to deploy an application revision to different sets of instances at different times. For example, you might use one deployment group to deploy an application revision to a set of instances tagged Test where you ensure the quality of the code. Next, you deploy the same application revision to a deployment group with instances tagged Staging for additional verification. Finally, when you are ready to release the latest application to customers, you deploy to a deployment group that includes instances tagged Production. "

NEW QUESTION 5

- (Exam Topic 1)

A startup's photo-sharing site is deployed in a VPC. An ELB distributes web traffic across two subnets. ELB session stickiness is configured to use the AWS-generated session cookie, with a session TTL of 5 minutes. The webserver Auto Scaling Group is configured as: min-size=4, max-size=4.

The startup is preparing for a public launch, by running load-testing software installed on a single EC2 instance running in us-west-2a. After 60 minutes of load-testing, the webserver logs show:

Which recommendations can help ensure load-testing HTTP requests are evenly distributed across the four web servers? Choose 2 answers

- A. Launch and run the load-tester EC2 instance from us-east-1 instead.
- B. Re-configure the load-testing software to re-resolve DNS for each web request.
- C. Use a 3rd-party load-testing service which offers globally-distributed test clients.

- D. Configure ELB and Auto Scaling to distribute across us-west-2a and us-west-2c.
- E. Configure ELB session stickiness to use the app-specific session cookie.

Answer: CE

NEW QUESTION 6

- (Exam Topic 1)

An application reads data from an Amazon DynamoDB table. Several times a day, for a period of 15 seconds, the application receives multiple ProvisionedThroughputExceeded errors.
How should this exception be handled?

- A. Create a new global secondary index for the table to help with the additional requests.
- B. Retry the failed read requests with exponential backoff.
- C. Immediately retry the failed read requests.
- D. Use the DynamoDB "UpdateItem" API to increase the provisioned throughput capacity of the table.

Answer: B

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Programming.Errors.html>

NEW QUESTION 7

- (Exam Topic 1)

How can software determine the public and private IP addresses of the Amazon EC2 instance that it is running on?

- A. Query the appropriate Amazon CloudWatch metric.
- B. Use ipconfig or ifconfig command.
- C. Query the local instance userdata.
- D. Query the local instance metadata.

Answer: D

Explanation:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/instancedata-data-categories.html>

Data: local-hostname Descriptions: The private IPv4 DNS hostname of the instance. In cases where multiple network interfaces are present, this refers to the eth0 device (the device for which the device number is 0). Data: public-ipv4 Descriptions: The public IPv4 address. If an Elastic IP address is associated with the instance, the value returned is the Elastic IP address.

NEW QUESTION 8

- (Exam Topic 1)

Which of the following statements about SWF are true? Choose 3 answers

- A. SWF tasks are assigned once and never duplicated
- B. SWF requires an S3 bucket for workflow storage
- C. SWF workflow executions can last up to a year
- D. SWF triggers SNS notifications on task assignment
- E. SWF uses deciders and workers to complete tasks
- F. SWF requires at least 1 EC2 instance per domain

Answer: ACE

NEW QUESTION 9

- (Exam Topic 1)

A Developer is writing a Linux-based application to run on AWS Elastic Beanstalk. Application requirements state that the application must maintain full capacity during updates while minimizing cost.

Which type of Elastic Beanstalk deployment policy should the Developer specify for the environment?

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

Answer: D

Explanation:

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

NEW QUESTION 10

- (Exam Topic 1)

What is the maximum number of S3 Buckets available per AWS account?

- A. 100 per region
- B. there is no limit
- C. 100 per account
- D. 500 per account
- E. 100 per IAM user

Answer:

C

Explanation:

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

NEW QUESTION 10

- (Exam Topic 1)

A Developer is writing a mobile application that allows users to view images from an S3 bucket. The users must be able to log in with their Amazon login, as well as Facebook® and/or Google® accounts.

How can the Developer provide this authentication functionality?

- A. Use Amazon Cognito with web identity federation.
- B. Use Amazon Cognito with SAML-based identity federation.
- C. Use AWS IAM Access/Secret keys in the application code to allow Get* on the S3 bucket.
- D. Use AWS STS AssumeRole in the application code and assume a role with Get* permissions on the S3 bucket.

Answer: A

Explanation:

Reference:

<http://jayendrapatil.com/tag/iam-role/>

<https://docs.aws.amazon.com/sdk-for-javascript/v2/developer-guide/loading-browser-credentials-federated-id.ht>

NEW QUESTION 14

- (Exam Topic 1)

If a message is retrieved from a queue in Amazon SQS, how long is the message inaccessible to other users by default?

- A. 0 seconds
- B. 1 hour
- C. 1 day
- D. forever
- E. 30 seconds

Answer: E

Explanation:

<https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-visibility-timeout.html> Visibility timeout:

default value = 30 seconds, minimum = 0 seconds, maximum = 12 hours

NEW QUESTION 16

- (Exam Topic 1)

A Developer has created a Lambda function and is finding that the function is taking longer to complete than expected. After some debugging, the Developer has discovered that increasing compute capacity would improve performance.

How can the Developer increase the Lambda compute resources?

- A. Run on a larger instance size with more compute capacity.
- B. Increase the maximum execution time.
- C. Specify a larger compute capacity when calling the Lambda function.
- D. Increase the allocated memory for the Lambda function.

Answer: D

NEW QUESTION 21

- (Exam Topic 1)

A company wants to implement a continuous integration for its workloads on AWS. The company wants to trigger unit test in its pipeline for commits-on its code repository, and wants to be notified of failure events in the pipeline.

How can these requirements be met?

- A. Store the source code in AWS CodeCommi
- B. Create a CodePipeline to automate unit testin
- C. Use Amazon SNS to trigger notifications of failure events.
- D. Store the source code in GitHu
- E. Create a CodePipeline to automate unit testin
- F. Use Amazon SES to trigger notifications of failure events.
- G. Store the source code on GitHu
- H. Create a CodePipeline to automate unit testin
- I. Use Amazon CloudWatch to trigger notifications of failure events.
- J. Store the source code in AWS CodeCommi
- K. Create a CodePipeline to automate unit testin
- L. Use Amazon CloudWatch to trigger notification of failure events.

Answer: D

NEW QUESTION 25

- (Exam Topic 1)

A legacy service has an XML-based SOAP interface. The Developer wants to expose the functionality of the service to external clients with the Amazon API Gateway. Which technique will accomplish this?

- A. Create a RESTful API with the API Gateway; transform the incoming JSON into a valid XML message for the SOAP interface using mapping templates.
- B. Create a RESTful API with the API Gateway; pass the incoming JSON to the SOAP interface through an Application Load Balancer.
- C. Create a RESTful API with the API Gateway; pass the incoming XML to the SOAP interface through an Application Load Balancer.
- D. Create a RESTful API with the API Gateway; transform the incoming XML into a valid message for the SOAP interface using mapping templates.

Answer: A

Explanation:

<https://blog.codecentric.de/en/2016/12/serverless-soap-legacy-api-integration-java-aws-lambda-aws-api-gateway>

NEW QUESTION 28

- (Exam Topic 1)

An application running on Amazon EC2 instances must access objects within an Amazon S3 bucket that are encrypted using server-side encryption using AWS KMS encryption keys (SSE-KMS). The application must have access to the customer master key (CMK) to decrypt the objects. Which combination of steps will grant the application access? (Select TWO.)

- A. Write an S3 bucket policy that grants the bucket access to the key.
- B. Grant access to the key in the IAM EC2 role attached to the application's EC2 instances.
- C. Write a key policy that enables IAM policies to grant access to the key.
- D. Grant access to the key in the S3 bucket's ACL
- E. Create a Systems Manager parameter that exposes the KMS key to the EC2 instances.

Answer: BC

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/decrypt-kms-encrypted-objects-s3/> IAM role needs access to the keys to decrypt the object and key policies must allow role access to the key. Key policies are the primary way to control access to customer master keys (CMKs) in AWS KMS. You need the permission to decrypt the AWS KMS key. When a user sends a GET request, Amazon S3 checks if the AWS Identity and Access Management (IAM) user or role that sent the request is authorized to decrypt the key associated with the object. If the IAM user or role belongs to the same AWS account as the key, then the permission to decrypt must be granted on the AWS KMS key's policy.

NEW QUESTION 33

- (Exam Topic 1)

You have written an application that uses the Elastic Load Balancing service to spread traffic to several web servers. Your users complain that they are sometimes forced to login again in the middle of using your application, after they have already logged in. This is not behavior you have designed. What is a possible solution to prevent this happening?

- A. Use instance memory to save session state.
- B. Use instance storage to save session state.
- C. Use EBS to save session state
- D. Use ElastiCache to save session state.
- E. Use Glacier to save session state.

Answer: D

Explanation:

<https://aws.amazon.com/caching/session-management/>

NEW QUESTION 37

- (Exam Topic 1)

A corporate web application is deployed within an Amazon VPC, and is connected to the corporate data center via IPsec VPN. The application must authenticate against the on-premise LDAP server. Once authenticated, logged-in users can only access an S3 keyspace specific to the user. Which two approaches can satisfy the objectives? Choose 2 answers

- A. The application authenticates against LDAP
- B. The application then calls the IAM Security Service to login to IAM using the LDAP credential
- C. The application can use the IAM temporary credentials to access the appropriate S3 bucket.
- D. The application authenticates against LDAP, and retrieves the name of an IAM role associated with the user
- E. The application then calls the IAM Security Token Service to assume that IAM Role
- F. The application can use the temporary credentials to access the appropriate S3 bucket.
- G. The application authenticates against IAM Security Token Service using the LDAP credential
- H. The application uses those temporary AWS security credentials to access the appropriate S3 bucket.
- I. Develop an identity broker which authenticates against LDAP, and then calls IAM Security Token Service to get IAM federated user credential
- J. The application calls the identity broker to get IAM federated user credentials with access to the appropriate S3 bucket.
- K. Develop an identity broker which authenticates against IAM Security Token Service to assume an IAM Role to get temporary AWS security credential
- L. The application calls the identity broker to get AWS temporary security credentials with access to the appropriate S3 bucket.

Answer: BD

Explanation:

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_temp_request.html

NEW QUESTION 38

- (Exam Topic 1)

A Developer uses AWS CodeDeploy to automate application deployment that connects to an external MySQL database. The Developer wants to securely access the encrypted secrets, such as API keys and database passwords. Which of the following solutions would involve the LEAST administrative effort?

- A. Save the secrets in Amazon S3 with AWS KMS server-side encryption, and use a signed URL to access them by using the IAM role from Amazon EC2

instances.

- B. Use the instance metadata to store the secrets and to programmatically access the secrets from EC2 instances.
- C. Use the Amazon DynamoDB client-side encryption library to save the secrets in DynamoDB and to programmatically access the secrets from EC2 instances.
- D. Use AWS SSM Parameter Store to store the secrets and to programmatically access them by using the IAM role from EC2 instances.

Answer: D

Explanation:

<https://docs.aws.amazon.com/systems-manager/latest/userguide/systems-manager-parameter-store.html>

NEW QUESTION 39

- (Exam Topic 1)

A Developer wants to use AWS X-Ray to trace a user request end-to-end throughout the software stack. The Developer made the necessary changes in the application tested it, and found that the application is able to send the traces to AWS X-Ray. However, when the application is deployed to an EC2 instance, the traces are not available.

Which of the following could create this situation? (Select two.)

- A. The traces are reaching X-Ray, but the Developer does not have access to view the records.
- B. The X-Ray daemon is not installed on the EC2 instance.
- C. The X-Ray endpoint specified in the application configuration is incorrect.
- D. The instance role does not have "xray:BatchGetTraces" and "xray:GetTraceGraph" permissions.
- E. The instance role does not have "xray:PutTraceSegments" and "xray:PutTelemetryRecords" permissions.

Answer: BE

NEW QUESTION 44

- (Exam Topic 1)

An application will ingest data at a very high throughput from many sources and must store the data in an Amazon S3 bucket. Which service would BEST accomplish this task?

- A. Amazon Kinesis Firehose
- B. Amazon S3 Acceleration Transfer
- C. Amazon SQS
- D. Amazon SNS

Answer: A

NEW QUESTION 47

- (Exam Topic 1)

A Developer wants to find a list of items in a global secondary index from an Amazon DynamoDB table. Which DynamoDB API call can the Developer use in order to consume the LEAST number of read capacity units?

- A. Scan operation using eventually-consistent reads
- B. Query operation using strongly-consistent reads
- C. Query operation using eventually-consistent reads
- D. Scan operation using strongly-consistent reads

Answer: C

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/bp-query-scan.html>

NEW QUESTION 51

- (Exam Topic 1)

A Developer is creating a serverless website with content that includes HTML files, images, videos, and JavaScript (client-side scripts).

Which combination of services should the Developer use to create the website?

- A. Amazon S3 and Amazon CloudFront
- B. Amazon EC2 and Amazon ElastiCache
- C. Amazon ECS and Redis
- D. AWS Lambda and Amazon API Gateway

Answer: A

Explanation:

Reference: <https://d1.awsstatic.com/whitepapers/Building%20Static%20Websites%20on%20AWS.pdf>

NEW QUESTION 56

- (Exam Topic 1)

A Developer is using AWS CLI, but when running list commands on a large number of resources, it is timing out.

What can be done to avoid this time-out?

- A. Use pagination
- B. Use shorthand syntax
- C. Use parameter values
- D. Use quoting strings

Answer: A

Explanation:

Reference: <https://docs.aws.amazon.com/cli/latest/userguide/cli-usage-pagination.html>

NEW QUESTION 57

- (Exam Topic 1)

A customer wants to deploy its source code on an AWS Elastic Beanstalk environment. The customer needs to perform deployment with minimal outage and should only use existing instances to retain application access log. What deployment policy would satisfy these requirements?

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

Answer: A

NEW QUESTION 61

- (Exam Topic 1)

Which of the following is an example of a good DynamoDB hash key schema for provisioned throughput efficiency?

- A. User ID, where the application has many different users.
- B. Status Code where most status codes are the same
- C. Device ID, where one is by far more popular than all the others.
- D. Game Type, where there are three possible game types

Answer: A

NEW QUESTION 64

- (Exam Topic 1)

What type of block cipher does Amazon S3 offer for server side encryption?

- A. Triple DES
- B. Advanced Encryption Standard
- C. Blowfish
- D. RC5

Answer: B

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/dev/serv-side-encryption.html>

NEW QUESTION 66

- (Exam Topic 1)

A Developer is working on an application that tracks hundreds of millions of product reviews in an Amazon DynamoDB table. The records include the data elements shown in the table:

Name	Type	Description
reviewID	Number	16 digit UUID
starRating	Number	Integer 1-5 of user rating
comment	String	User comment string
productID	Number	Product ID being reviewed

Which field, when used as the partition key, would result in the MOST consistent performance using DynamoDB?

- A. starRating
- B. reviewID
- C. comment
- D. productID

Answer: B

NEW QUESTION 68

- (Exam Topic 1)

Which of the following are valid SNS delivery transports? Choose 2 answers

- A. HTTP
- B. UDP
- C. SMS
- D. DynamoDB
- E. Named Pipes

Answer: AC

Explanation:

<https://aws.amazon.com/sns/faqs/>

Q: What are the different delivery formats/transport for receiving notifications?

In order for customers to have broad flexibility of delivery mechanisms, Amazon SNS supports notifications over multiple transport protocols. Customers can select one of the following transports as part of the subscription requests:

“HTTP”, “HTTPS” – Subscribers specify a URL as part of the subscription registration; notifications will be delivered through an HTTP POST to the specified URL.

“Email”, “Email-JSON” – Messages are sent to registered addresses as email. Email-JSON sends notifications as a JSON object, while Email sends text-based email.

“SQS” – Users can specify an SQS standard queue as the endpoint; Amazon SNS will enqueue a notification message to the specified queue (which subscribers can then process using SQS APIs such as ReceiveMessage, DeleteMessage, etc.). Note that FIFO queues are not currently supported.

“SMS” – Messages are sent to registered phone numbers as SMS text messages.

NEW QUESTION 70

- (Exam Topic 1)

Company B provides an online image recognition service and utilizes SQS to decouple system components for scalability. The SQS consumers poll the imaging queue as often as possible to keep end-to-end throughput as high as possible. However, Company B is realizing that polling in tight loops is burning CPU cycles and increasing costs with empty responses.

How can Company B reduce the number of empty responses?

- A. Set the imaging queue visibility Timeout attribute to 20 seconds
- B. Set the Imaging queue ReceiveMessageWaitTimeSeconds attribute to 20 seconds
- C. Set the imaging queue MessageRetentionPeriod attribute to 20 seconds
- D. Set the DelaySeconds parameter of a message to 20 seconds

Answer: B

Explanation:

<https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-short-and-long-polling>.

NEW QUESTION 74

- (Exam Topic 1)

A static website is hosted in an Amazon S3 bucket. Several HTML pages on the site use JavaScript to download images from another Amazon S3 bucket. These images are not displayed when users browse the site.

What is the possible cause for the issue?

- A. The referenced Amazon S3 bucket is in another region.
- B. The images must be stored in the same Amazon S3 bucket.
- C. Port 80 must be opened on the security group in which the Amazon S3 bucket is located.
- D. Cross Origin Resource Sharing must be enabled on the Amazon S3 bucket.

Answer: D

NEW QUESTION 79

- (Exam Topic 1)

A company needs to encrypt data at rest, but it wants to leverage an AWS managed service using its own master key.

Which of the following AWS services can be used to meet these requirements?

- A. SSE with Amazon S3
- B. SSE with AWS KMS
- C. Client-side encryption
- D. AWS IAM roles and policies

Answer: B

NEW QUESTION 82

- (Exam Topic 1)

When uploading an object, what request header can be explicitly specified in a request to Amazon S3 to encrypt object data when saved on the server side?

- A. x-amz-storage-class
- B. Content-MD5
- C. x-amz-security-token
- D. x-amz-server-side-encryption

Answer: D

Explanation:

<https://docs.aws.amazon.com/AmazonS3/latest/dev/UsingServerSideEncryption.html#APISupportforServer-Sid>

NEW QUESTION 83

- (Exam Topic 1)

Which of the following are valid arguments for an SNS Publish request? Choose 3 answers

- A. TopicArn
- B. Subject
- C. Destination
- D. Format
- E. Message
- F. Language

Answer: ABE

Explanation:

https://docs.aws.amazon.com/sns/latest/api/API_Publish.html Message
MessageAttributes MessageStructure PhoneNumber Subject
TargetArn TopicArn

NEW QUESTION 84

- (Exam Topic 1)

A Developer is testing a Docker-based application that uses the AWS SDK to interact with Amazon DynamoDB. In the local development environment, the application has used IAM access keys. The application is now ready for deployment onto an ECS cluster. How should the application authenticate with AWS services in production?

- A. Configure an ECS task IAM role for the application to use
- B. Refactor the application to call AWS STS AssumeRole based on an instance role
- C. Configure AWS access key/secret access key environment variables with new credentials
- D. Configure the credentials file with a new access key/secret access key

Answer: A

Explanation:

[https://docs.aws.amazon.com/AmazonECS/latest/developerguide/task_IAM_role.html#:~:targetText=Amazon%](https://docs.aws.amazon.com/AmazonECS/latest/developerguide/task_IAM_role.html#:~:targetText=Amazon%20)

NEW QUESTION 89

- (Exam Topic 1)

Which of the following services are included at no additional cost with the use of the AWS platform? Choose 2 answers

- A. Simple Storage Service
- B. Elastic Compute Cloud
- C. Auto Scaling
- D. Elastic Load Balancing
- E. CloudFormation
- F. Simple Workflow Service

Answer: CE

NEW QUESTION 93

- (Exam Topic 1)

A web application is using Amazon Kinesis Streams for clickstream data that may not be consumed for up to 12 hours. How can the Developer implement encryption at rest for data within the Kinesis Streams?

- A. Enable SSL connections to Kinesis
- B. Use Amazon Kinesis Consumer Library
- C. Encrypt the data once it is at rest with a Lambda function
- D. Enable server-side encryption in Kinesis Streams

Answer: D

Explanation:

<https://docs.aws.amazon.com/streams/latest/dev/what-is-sse.html>

<https://aws.amazon.com/about-aws/whats-new/2017/07/amazon-kinesis-streams-introduces-server-side-encrypti>

NEW QUESTION 98

- (Exam Topic 1)

A Development team currently supports an application that uses an in-memory store to save accumulated game results. Individual results are stored in a database. As part of migrating to AWS, the team needs to use automatic scaling. The team knows this will yield inconsistent results. Where should the team store these accumulated game results to BEST allow for consistent results without impacting performance?

- A. Amazon S3
- B. Amazon RDS
- C. Amazon ElastiCache
- D. Amazon Kinesis

Answer: C

NEW QUESTION 102

- (Exam Topic 1)

A Developer writes an AWS Lambda function and uploads the code in a .ZIP file to Amazon S3. The Developer makes changes to the code and uploads a new .ZIP file to Amazon S3. However, Lambda executes the earlier code. How can the Developer fix this in the LEAST disruptive way?

- A. Create another Lambda function and specify the new .ZIP file.
- B. Call the update-function-code API.
- C. Remove the earlier .ZIP file first, then add the new .ZIP file.
- D. Call the create-alias API.

Answer: B

Explanation:

<https://docs.aws.amazon.com/cli/latest/reference/lambda/update-function-code.html>

NEW QUESTION 107

- (Exam Topic 1)

An application on AWS is using third-party APIs. The Developer needs to monitor API errors in the code, and wants to receive notifications if failures go above a set threshold value.

How can the Developer achieve these requirements?

- A. Publish a custom metric on Amazon CloudWatch and use Amazon SES for notification.
- B. Use an Amazon CloudWatch API-error metric and use Amazon SNS for notification.
- C. Use an Amazon CloudWatch API-error metric and use Amazon SES for notification.
- D. Publish a custom metric on Amazon CloudWatch and use Amazon SNS for notification.

Answer: D

NEW QUESTION 110

- (Exam Topic 1)

You are inserting 1000 new items every second in a DynamoDB table. Once an hour these items are analyzed and then are no longer needed. You need to minimize provisioned throughput, storage, and API calls.

Given these requirements, what is the most efficient way to manage these Items after the analysis?

- A. Retain the items in a single table
- B. Delete items individually over a 24 hour period
- C. Delete the table and create a new table per hour
- D. Create a new table per hour

Answer: C

NEW QUESTION 111

- (Exam Topic 1)

An application has hundreds of users. Each user may use multiple devices to access the application. The Developer wants to assign unique identifiers to these users regardless of the device they use.

Which of the following methods should be used to obtain unique identifiers?

- A. Create a user table in Amazon DynamoDB as key-value pairs of users and their device
- B. Use these keys as unique identifiers.
- C. Use IAM-generated access key IDs for the users as the unique identifier, but do not store secret keys.
- D. Implement developer-authenticated identities by using Amazon Cognito, and get credentials for these identities.
- E. Assign IAM users and roles to the user
- F. Use the unique IAM resource ID as the unique identifier.

Answer: C

NEW QUESTION 116

- (Exam Topic 1)

An application under development is required to store hundreds of video files. The data must be encrypted within the application prior to storage, with a unique key for each video file.

How should the Developer code the application?

- A. Use the KMS Encrypt API to encrypt the data
- B. Store the encrypted data key and data.
- C. Use a cryptography library to generate an encryption key for the application
- D. Use the encryption key to encrypt the data
- E. Store the encrypted data.
- F. Use the KMS GenerateDataKey API to get a data key
- G. Encrypt the data with the data key
- H. Store the encrypted data key and data.
- I. Upload the data to an S3 bucket using server side-encryption with an AWS KMS key.

Answer: C

Explanation:

https://docs.aws.amazon.com/kms/latest/APIReference/API_GenerateDataKey.html

NEW QUESTION 121

- (Exam Topic 1)

A Developer is creating a mobile application with a limited budget. The solution requires a scalable service that will enable customers to sign up and authenticate into the mobile application while using the organization's current SAML 2.0 identity provider.

Which AWS service should be used to meet these requirements?

- A. AWS Lambda
- B. Amazon Cognito
- C. AWS IAM
- D. Amazon EC2

Answer: B

NEW QUESTION 125

- (Exam Topic 1)

A Developer is asked to implement a caching layer in front of Amazon RDS. Cached content is expensive to regenerate in case of service failure. Which implementation below would work while maintaining maximum uptime?

- A. Implement Amazon ElastiCache Redis in Cluster Mode
- B. Install Redis on an Amazon EC2 instance.
- C. Implement Amazon ElastiCache Memcached.
- D. Migrate the database to Amazon Redshift.

Answer: A

Explanation:

<https://docs.aws.amazon.com/AmazonElastiCache/latest/mem-ug/SelectEngine.html>

NEW QUESTION 128

- (Exam Topic 1)

A Lambda function is packaged for deployment to multiple environments, including development, test, production, etc. Each environment has unique set of resources such as databases, etc.

How can the Lambda function use the resources for the current environment?

- A. Apply tags to the Lambda functions.
- B. Hardcore resources in the source code.
- C. Use environment variables for the Lambda functions.
- D. Use separate function for development and production.

Answer: C

NEW QUESTION 131

- (Exam Topic 1)

A Developer is receiving HTTP 400: ThrottlingException errors intermittently when calling the Amazon CloudWatch API. When a call fails, no data is retrieved. What best practice should first be applied to address this issue?

- A. Contact AWS Support for a limit increase.
- B. Use the AWS CLI to get the metrics
- C. Analyze the applications and remove the API call
- D. Retry the call with exponential backoff

Answer: D

Explanation:

https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/cloudwatch_limits.html

NEW QUESTION 132

- (Exam Topic 1)

A Developer is creating a mobile application that will not require users to log in. What is the MOST efficient method to grant users access to AWS resources?

- A. Use an identity provider to securely authenticate with the application.
- B. Create an AWS Lambda function to create an IAM user when a user accesses the application.
- C. Create credentials using AWS KMS and apply these credentials to users when using the application.
- D. Use Amazon Cognito to associate unauthenticated users with an IAM role that has limited access to resources.

Answer: D

Explanation:

<https://docs.aws.amazon.com/cognito/latest/developerguide/iam-roles.html>

NEW QUESTION 134

- (Exam Topic 1)

A Developer needs to use AWS X-Ray to monitor an application that is deployed on EC2 instances. What steps have to be executed to perform the monitoring?

- A. Deploy the X-Ray SDK with the application and use X-Ray annotation.
- B. Install the X-Ray daemon and instrument the application code.
- C. Install the X-Ray daemon and configure it to forward data to Amazon CloudWatch Events.
- D. Deploy the X-Ray SDK with the application and instrument the application code.

Answer: C

NEW QUESTION 135

- (Exam Topic 1)

In a move toward using microservices, a company's Management team has asked all Development teams to build their services so that API requests depend only on that service's data store. One team is building a Payments service which has its own database; the service needs data that originates in the Accounts database. Both are using Amazon DynamoDB.

What approach will result in the simplest, decoupled, and reliable method to get near-real time updates from the Accounts database?

- A. Use Amazon Glue to perform frequent ETL updates from the Accounts database to the Payments database.
- B. Use Amazon ElastiCache in Payments, with the cache updated by triggers in the Accounts database.

- C. Use Amazon Kinesis Data Firehouse to deliver all changes from the Accounts database to the Payments database.
- D. Use Amazon DynamoDB Streams to deliver all changes from the Accounts database to the Payments database.

Answer: D

Explanation:

Reference:

<https://aws.amazon.com/blogs/database/how-to-perform-ordered-data-replication-between-applications-by-using>

NEW QUESTION 137

- (Exam Topic 1)

You are writing to a DynamoDB table and receive the following exception: "ProvisionedThroughputExceededException". though according to your Cloudwatch metrics for the table, you are not exceeding your provisioned throughput.

What could be an explanation for this?

- A. You haven't provisioned enough DynamoDB storage instances
- B. You're exceeding your capacity on a particular Range Key
- C. You're exceeding your capacity on a particular Hash Key
- D. You're exceeding your capacity on a particular Sort Key
- E. You haven't configured DynamoDB Auto Scaling triggers

Answer: C

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.CoreComponents.html#Ho>

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.Partitions.html>

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/bp-partition-key-design.html>

NEW QUESTION 139

- (Exam Topic 1)

An application takes 40 seconds to process instructions received in an Amazon SQS message.

Assuming the SQS queue is configured with the default VisibilityTimeout value, what is the BEST way, upon receiving a message, to ensure that no other instances can retrieve a message that has already been processed or is currently being processed?

- A. Use the ChangeMessageVisibility API to increase the VisibilityTimeout, then use the DeleteMessage API to delete the message.
- B. Use the DeleteMessage API call to delete the message from the queue, then call DeleteQueue API to remove the queue.
- C. Use the ChangeMessageVisibility API to decrease the timeout value, then use the DeleteMessage API to delete the message.
- D. Use the DeleteMessageVisibility API to cancel the VisibilityTimeout, then use the DeleteMessage API to delete the message.

Answer: A

Explanation:

<https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-visibility-timeout.html> In SQS, messages remain there. It is the consumer's responsibility to delete it, once consumed and processed.

NEW QUESTION 144

- (Exam Topic 1)

A Developer created a dashboard for an application using Amazon API Gateway, Amazon S3, AWS Lambda, and Amazon RDS. The Developer needs an authentication mechanism allowing a user to sign in and view the dashboard. It must be accessible from mobile applications, desktops, and tablets, and must remember user preferences across platforms.

Which AWS service should the Developer use to support this authentication scenario?

- A. AWS KMS
- B. Amazon Cognito
- C. AWS Directory Service
- D. Amazon IAM

Answer: B

Explanation:

Cognito user pool provides sign up and sign in functionality along with identity pool which provides temp credentials for using aws services.

NEW QUESTION 148

- (Exam Topic 1)

Which of the following is chosen as the default region when making an API call with an AWS SDK?

- A. ap-northeast-1
- B. us-west-2
- C. us-east-1
- D. eu-west-1
- E. us-central-1

Answer: C

Explanation:

<https://docs.aws.amazon.com/sdk-for-java/v1/developer-guide/java-dg-region-selection.html>. This section applies only when using a client builder to access AWS services. AWS clients created by using the client constructor will not automatically determine region from the environment and will, instead, use the default SDK region (USEast1).

NEW QUESTION 149

- (Exam Topic 1)

A company developed a set of APIs that are being served through the Amazon API Gateway. The API calls need to be authenticated based on OpenID identity providers such as Amazon or Facebook. The APIs should allow access based on a custom authorization model.

Which is the simplest and MOST secure design to use to build an authentication and authorization model for the APIs?

- A. Use Amazon Cognito user pools and a custom authorizer to authenticate and authorize users based on JSON Web Tokens.
- B. Build a OpenID token broker with Amazon and Facebook
- C. Users will authenticate with these identity providers and pass the JSON Web Token to the API to authenticate each API call.
- D. Store user credentials in Amazon DynamoDB and have the application retrieve temporary credentials from AWS STS
- E. Make API calls by passing user credentials to the APIs for authentication and authorization.
- F. Use Amazon RDS to store user credentials and pass them to the APIs for authentications and authorization.

Answer: A

NEW QUESTION 152

- (Exam Topic 1)

You run an ad-supported photo sharing website using S3 to serve photos to visitors of your site. At some point you find out that other sites have been linking to the photos on your site, causing loss to your business.

What is an effective method to mitigate this?

- A. Store photos on an EBS volume of the web server
- B. Remove public read access and use signed URLs with expiry dates.
- C. Use CloudFront distributions for static content.
- D. Block the IPs of the offending websites in Security Groups.

Answer: B

Explanation:

<https://aws.amazon.com/getting-started/projects/building-fast-session-caching-with-amazon-elasticache-for-redis/>

NEW QUESTION 157

- (Exam Topic 1)

A company is migrating its on-premises database to Amazon RDS for MySQL. The company has read-heavy workloads, and wants to make sure it re-factors its code to achieve optimum read performance for its queries.

How can this objective be met?

- A. Add database retries to effectively use RDS with vertical scaling
- B. Use RDS with multi-AZ deployment
- C. Add a connection string to use an RDS read replica for read queries
- D. Add a connection string to use a read replica on an EC2 instance.

Answer: C

NEW QUESTION 159

- (Exam Topic 1)

A Developer is creating an Auto Scaling group whose instances need to publish a custom metric to Amazon CloudWatch.

Which method would be the MOST secure way to authenticate a CloudWatch PUT request?

- A. Create an IAM user with PutMetricData permission and put the user credentials in a private repository; have applications pull the credentials as needed.
- B. Create an IAM user with PutMetricData permission, and modify the Auto Scaling launch configuration to inject the user credentials into the instance user data.
- C. Modify the CloudWatch metric policies to allow the PutMetricData permission to instances from the Auto Scaling group.
- D. Create an IAM role with PutMetricData permission and modify the Auto Scaling launching configuration to launch instances using that role.

Answer: D

NEW QUESTION 160

- (Exam Topic 1)

A serverless application uses an API Gateway and AWS Lambda.

Where should the Lambda function store its session information across function calls?

- A. In an Amazon DynamoDB table
- B. In an Amazon SQS queue
- C. In the local filesystem
- D. In an SQLite session table using `–DSQLITE_ENABLE_SESSION`

Answer: A

NEW QUESTION 163

- (Exam Topic 1)

A company is migrating a single-server, on-premises web application to AWS. The company intends to use multiple servers behind an Elastic Load Balancer (ELB) to balance the load, and will also store session data in memory on the web server. The company does not want to lose that session data if a server fails or goes offline, and it wants to minimize user's downtime.

Where should the company move session data to MOST effectively reduce downtime and make users' session data more fault tolerant?

- A. An Amazon ElastiCache for Redis cluster
- B. A second Amazon EBS volume

- C. The web server's primary disk
- D. An Amazon EC2 instance dedicated to session data

Answer: A

NEW QUESTION 167

- (Exam Topic 1)

An AWS Lambda function must read data from an Amazon RDS MySQL database in a VPC and also reach a public endpoint over the internet to get additional data.

Which steps must be taken to allow the function to access both the RDS resource and the public endpoint? (Select TWO.)

- A. Modify the default configuration for the Lambda function to associate it with an Amazon VPC private subnet.
- B. Modify the default network access control list to allow outbound traffic.
- C. Add a NAT Gateway to the VPC.
- D. Modify the default configuration of the Lambda function to associate it with a VPC public subnet.
- E. Add an environmental variable to the Lambda function to allow outbound internet access.

Answer: AC

Explanation:

Reference: <https://docs.aws.amazon.com/lambda/latest/dg/vpc.html>

NEW QUESTION 169

- (Exam Topic 1)

A meteorological system monitors 600 temperature gauges, obtaining temperature samples every minute and saving each sample to a DynamoDB table. Each sample involves writing 1K of data and the writes are evenly distributed over time.

How much write throughput is required for the target table?

- A. 1 write capacity unit
- B. 10 write capacity units
- C. 60 write capacity units
- D. 600 write capacity units
- E. 3600 write capacity units

Answer: B

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ReadWriteCapacityMode.h>

NEW QUESTION 170

- (Exam Topic 1)

How is provisioned throughput affected by the chosen consistency model when reading data from a DynamoDB table?

- A. Strongly consistent reads use the same amount of throughput as eventually consistent reads
- B. Strongly consistent reads use more throughput than eventually consistent reads.
- C. Strongly consistent reads use less throughput than eventually consistent reads
- D. Strongly consistent reads use variable throughput depending on read activity

Answer: B

NEW QUESTION 172

- (Exam Topic 1)

Company A has an S3 bucket containing premier content that they intend to make available to only paid subscribers of their website. The S3 bucket currently has default permissions of all objects being private to prevent inadvertent exposure of the premier content to non-paying website visitors.

How can Company A provide only paid subscribers the ability to download a premier content file in the S3 bucket?

- A. Apply a bucket policy that grants anonymous users to download the content from the S3 bucket
- B. Generate a pre-signed object URL for the premier content file when a paid subscriber requests a download
- C. Add a bucket policy that requires Multi-Factor Authentication for requests to access the S3 bucket objects
- D. Enable server-side encryption on the S3 bucket for data protection against the non-paying website visitors

Answer: B

NEW QUESTION 173

- (Exam Topic 2)

A Developer is writing an application that runs on Amazon EC2 instances in an Auto scaling group. The application data is stored in an Amazon DynamoDB table and records are constantly updated by all instances. An instance sometimes retrieves old data. The Developer wants to correct this by making sure the reads are strongly consistent.

How can the developer accomplish this?

- A. Set consistentRead to true when calling Getitem.
- B. Create a new DynamoDB Accelerator (DAX) table.
- C. Set consistency to strong when calling Update Table.
- D. Use the GetshardIterator command.

Answer: B

Explanation:

Reference: <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/HowItWorks.ReadConsistency.html>

NEW QUESTION 174

- (Exam Topic 2)

A company is developing a report executed by AWS Step Functions Amazon CloudWatch shows errors in the Step Functions task state machine To troubleshoot each task, the state input needs to be included along with the error message in the state output.

Which coding practice can preserve both the original input and the error for the state?

- A. Use ResultPath in a Catch statement to include the error with the original input
- B. Use inputPath in a Catch statement and set the value to null.
- C. Use ErrorEquals in a Retry statement to include the error with the original input
- D. Use OutputPath in a Retry statement and set the value to \$.

Answer: A

Explanation:

Use ResultPath in a Catch to include the error with the original input.

Reference: <https://docs.aws.amazon.com/step-functions/latest/dg/input-output-resultpath.html>

NEW QUESTION 176

- (Exam Topic 2)

A development team is designing a mobile app that requires multi-factor authentication Which steps should be taken to achieve this? (Select TWO)

- A. Use Amazon Cognito to create a user pool and create users in the user pool
- B. Send multi-factor authentication text codes to users with the Amazon SNS Publish API call in the app code
- C. Enable multi-factor authentication for the Amazon Cognito user pool
- D. Use AWS IAM to create IAM users
- E. Enable multi-factor authentication for the users created in AWS IAM.

Answer: AC

NEW QUESTION 181

- (Exam Topic 2)

A developer added a new feature to an application running on an Amazon EC2 instance that uses Amazon SQS After deployment, the developer noticed a significant increase in Amazon SQS costs. When monitoring the Amazon SQS metrics on Amazon CloudWatch. the developer found that on average one message per minute is posted on this queue.

What can be done to reduce Amazon SQS costs for this application?

- A. Increase the Amazon SQS queue polling timeout
- B. Scale down the Amazon SQS queue to the appropriate size for low traffic demand.
- C. Configure push delivery via Amazon SNS instead of polling the Amazon SQS queue
- D. Use an Amazon SQS first-in, first-out (FIFO) queue instead of a standard queue.

Answer: A

NEW QUESTION 184

- (Exam Topic 2)

A company has an application where reading objects from Amazon S3 is based on the type of user The user types are registered user and guest user The company has 25.000 users and is growing Information is pulled from an S3 bucket depending on the user type.

Which approaches are recommended to provide access to both user types? (Select TWO.)

- A. Provide a different access key and secret access key in the application code for registered users and guest users to provide read access to the objects
- B. Use S3 bucket policies to restrict read access to specific IAM users
- C. Use Amazon Cognito to provide access using authenticated and unauthenticated roles
- D. Create a new IAM user for each user and grant read access.
- E. Use the AWS IAM service and let the application assume the different roles using the AWS Security Token Service (AWS STS) AssumeRole action depending on the type of user and provide read access to Amazon S3 using the assumed role

Answer: BC

NEW QUESTION 188

- (Exam Topic 2)

A Developer is building a three-tier web application that should be able to handle a minimum of 5000 requests per minute. Requirements state that the web tier should be completely stateless while the application maintains session state for the users.

How can session data be externalized, keeping latency at the LOWEST possible value?

- A. Create an Amazon RDS instance, then implement session handling at the application level to leverage a database inside the RDS database instance for session data storage
- B. Implement a shared file system solution across the underlying Amazon EC2 instances, then implement session handling at the application level to leverage the shared file system for session data storage
- C. Create an Amazon ElastiCache Memcached cluster, then implement session handling at the application level to leverage the cluster for session data storage
- D. Create an Amazon DynamoDB table, then implement session handling at the application level to leverage the table for session data storage

Answer: A

Explanation:

<https://aws.amazon.com/caching/session-management/>

NEW QUESTION 190

- (Exam Topic 2)

An e-commerce web application that shares session state on-premises is being migrated to AWS. The application must be fault tolerant, natively highly scalable, and any service interruption should not affect the user experience.

What is the best option to store the session state?

- A. Store the session state in Amazon ElastiCache
- B. Store the session state in Amazon CloudFront
- C. Store the session state in Amazon S3
- D. Enable session stickiness using elastic load balancers

Answer: A

Explanation:

<https://aws.amazon.com/caching/session-management/>

NEW QUESTION 192

- (Exam Topic 2)

An application is running on a cluster of Amazon EC2 instance. While trying to read objects stored within a single Amazon S3 bucket that are encrypted with server-side encryption with AWS KMS managed keys (SSE-KMS), the application receives the following error:

Service : AWSKMS: Status Code: 400: Code : ThrottlingException

Which combination of steps should be taken to prevent this failure? (Select TWO.)

- A. Contact AWS Support to request an AWS KMS rate limit increase.
- B. Perform error retries with exponential backoff in the application code.
- C. Contact AWS Support to request a S3 rate limit increase.
- D. Import a customer master key (CMK) with a larger key size.
- E. Use more than one customer master key (CMK) to encrypt S3 data

Answer: AD

NEW QUESTION 195

- (Exam Topic 2)

An application writes items to an Amazon DynamoDB table. As the application scales to thousands of instances, calls to the DynamoDB API generate occasional ThrottlingException errors. The application is coded in a language incompatible with the AWS SDK.

How should the error be handled?

- A. Add exponential backoff to the application logic
- B. Use Amazon SQS as an API message bus
- C. Pass API calls through Amazon API Gateway
- D. Send the items to DynamoDB through Amazon Kinesis Data Firehose

Answer: A

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/throttled-ddb/>

SDKs automatically add exponential backoff. If not using the AWS SDKs, add your own backoff logic to the application code.

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Programming.Errors.html#Programming>.

NEW QUESTION 198

- (Exam Topic 2)

A company runs an e-commerce website that uses Amazon DynamoDB where pricing for items is dynamically updated in real time. At any given time, multiple updates may occur simultaneously for pricing information on a particular product. This is causing the original editor's changes to be overwritten without a proper review process.

Which DynamoDB write option should be selected to prevent this overwriting?

- A. Concurrent writes
- B. Conditional writes
- C. Atomic writes
- D. Batch writes

Answer: B

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/WorkingWithItems.html#WorkingWithIt>

NEW QUESTION 203

- (Exam Topic 2)

An on-premises application is implemented using a Linux, Apache, MySQL and PHP (LAMP) stack. The Developer wants to run this application in AWS.

Which of the following sets of AWS services can be used to run this stack?

- A. Amazon API Gateway, Amazon S3
- B. AWS Lambda, Amazon DynamoDB
- C. Amazon EC2, Amazon Aurora
- D. Amazon Cognito, Amazon RDS
- E. Amazon ECS, Amazon EBS

Answer: C

NEW QUESTION 207

- (Exam Topic 2)

A web application is designed to allow new users to create accounts using their email addresses. The application will store attributes for each user, and is expecting millions of user to sign up.

What should the Developer implement to achieve the design goals?

- A. Amazon Cognito user pools
- B. AWS Mobile Hub user data storage
- C. Amazon Cognito Sync
- D. AWS Mobile Hub cloud logic

Answer: A

NEW QUESTION 208

- (Exam Topic 2)

An Amazon RDS database instance is used by many applications to look up historical data. The query rate is relatively constant. When the historical data is updated each day, the resulting write traffic slows the read query performance and affects all application users.

What can be done to eliminate the performance impact on application users?

- A. Make sure Amazon RDS is Multi-AZ so it can better absorb increased traffic.
- B. Create an RDS Read Replica and direct all read traffic to the replica.
- C. Implement Amazon ElastiCache in front of Amazon RDS to buffer the write traffic.
- D. Use Amazon DynamoDB instead of Amazon RDS to buffer the read traffic.

Answer: B

Explanation:

<https://aws.amazon.com/rds/details/read-replicas/>

NEW QUESTION 213

- (Exam Topic 2)

A developer is provided with an HTTPS clone URL for an AWS CodeCommit repository. What needs to be configured before cloning this repository?

- A. Use AWS KMS to set up public and private keys for use with CodeCommit.
- B. Set up the Git credential helper to use an AWS credential profile, and enable the helper to send the path to the repositories.
- C. Generate encryption keys using AWS CloudHSM, then export the key for use with AWS CodeCommit.
- D. Use AWS certificate manager to provision public and private SSL/TLS certificates.

Answer: B

Explanation:

AWS credential profile, and enabling the Git credential helper to send the path to repositories: Reference:

<https://docs.aws.amazon.com/codecommit/latest/userguide/setting-up-https-unixes.html>

NEW QUESTION 218

- (Exam Topic 2)

A company stores all personally identifiable information (PII) in an Amazon DynamoDB table named PII in Account A. An application running on Amazon EC2 instances in Account B requires access to the PII table. An administrators in Account A created an IAM role named AccessPII with privileges to access the PII table, and made account B a trusted entity.

Which combination of actional steps should Developers take to access the table? (Select TWO)

- A. Ask an Administrator in Account B to allow the EC2 IAM role permission to assume the AccessPII role
- B. Ask an Administrator in Account B to allow the EC2 IAM role permission to assume the AccessPII role with predefined service control policies
- C. Ask an Administrator in Account A to allow the EG2 IAM role permission to assume the AccessPII role with predefined service control policies.
- D. Include the AssumeRole API in the application code logic to obtain credentials to access the PII table
- E. Include the GetSession token API in the application code logic to obtain credentials to access the PII table

Answer: BE

NEW QUESTION 222

- (Exam Topic 2)

A Development team wants to instrument their code to provide more detailed information to AWS X-Ray than simple outgoing and incoming requests. This will generate large amounts of data, so the Development team wants to implement indexing so they can filter the data.

What should the Development team do to achieve this?

- A. Add annotations to the segment document and the code
- B. Add metadata to the segment document and the code
- C. Configure the necessary X-Ray environment variables
- D. Install required plugins for the appropriate AWS SDK

Answer: A

Explanation:

<https://docs.aws.amazon.com/xray/latest/devguide/xray-sdk-python-segment.html> <https://docs.aws.amazon.com/xray/latest/devguide/xray-concepts.html#xray-concepts-annotations>

NEW QUESTION 224

- (Exam Topic 2)

A developer is setting up Amazon API Gateway for their company's products. The API will be used by registered developers to query and update their environments. The company wants to limit the amount of requests end users can send for both cost and security reasons. Management wants to offer registered developers the option of buying larger packages that allow for more requests.

How can the developer accomplish this with the LEAST amount of overhead management?

- A. Enable throttling for the API Gateway stage
- B. Set a value for both the rate and burst capacity
- C. If a registered user chooses a larger package, create a stage for them, adjust the values, and share the new URL with them.
- D. Set up Amazon CloudWatch API logging in API Gateway. Create a filter based on the user and requestTime fields and create an alarm on this filter. Write an AWS Lambda function to analyze the values and requester information, and respond accordingly. Set up the function as the target for the alarm. If a registered user chooses a larger package, update the Lambda code with the values.
- E. Enable Amazon CloudWatch metrics for the API Gateway stage. Set up CloudWatch alarms based off the Count metric and the ApiName, Method, Resource, and Stage dimensions to alert when request rates pass the threshold. Set the alarm action to Deny. If a registered user chooses a larger package, create a user-specific alarm and adjust the values.
- F. Set up a default usage plan, specify values for the rate and burst capacity, and associate it with a stage. If a registered user chooses a larger package, create a custom plan with the appropriate values and associate the plan with the user.

Answer: D

NEW QUESTION 229

- (Exam Topic 2)

A Developer is trying to deploy a serverless application using AWS CodeDeploy. The application was updated and needs to be redeployed. What file does the Developer need to update to push that change through CodeDeploy?

- A. dockerrun.aws.json
- B. buildspec.yml
- C. appspec.yml
- D. ebextensions.config

Answer: C

Explanation:

<https://docs.aws.amazon.com/codedeploy/latest/userguide/application-revisions-push.html>

NEW QUESTION 232

- (Exam Topic 2)

A company is developing a web application that allows its employees to upload a profile picture to a private Amazon S3 bucket. There is no size limit for the profile pictures, which should be displayed every time an employee logs in. For security reasons, the pictures cannot be publicly accessible.

What is a viable long-term solution for this scenario?

- A. Generate a presigned URL when a picture is uploaded. Save the URL in an Amazon DynamoDB table. Return the URL to the browser when the employee logs in.
- B. Save the picture's S3 key in an Amazon DynamoDB table. Create an Amazon S3 VPC endpoint to allow the employees to download pictures once they log in.
- C. Encode a picture using base64. Save the base64 string in an Amazon DynamoDB table. Allow the browser to retrieve the string and convert it to a picture.
- D. Save the picture's S3 key in an Amazon DynamoDB table.
- E. Use a function to generate a presigned URL every time an employee logs in.
- F. Return the URL to the browser.

Answer: B

Explanation:

Reference:

<https://aws.amazon.com/premiumsupport/knowledge-center/s3-private-connection-noauthentication/>

NEW QUESTION 235

- (Exam Topic 2)

A website's page load times are gradually increasing as more users access the system at the same time. Analysis indicates that a user profile is being loaded from a database in all the web pages being visited by each user and this is increasing the database load and the page load latency. To address this issue the Developer decides to cache the user profile data.

Which caching strategy will address this situation MOST efficiently?

- A. Create a new Amazon EC2 Instance and run a NoSQL database on it.
- B. Cache the profile data within this database using the write-through caching strategy.
- C. Create an Amazon ElastiCache cluster to cache the user profile data.
- D. Use a cache-aside caching strategy.
- E. Use a dedicated Amazon RDS instance for caching profile data.
- F. Use a write-through caching strategy.
- G. Create an ElastiCache cluster to cache the user profile data.
- H. Use a write-through caching strategy.

Answer: B

Explanation:

<https://docs.aws.amazon.com/AmazonElastiCache/latest/mem-ug/Strategies.html>

NEW QUESTION 240

- (Exam Topic 2)

A software engineer developed an AWS Lambda function in Node.js to do some CPU-intensive data

processing. With the default settings, the Lambda function takes about 5 minutes to complete. Which approach should a developer take to increase the speed of

completion"

- A. Instead of using Node.js
- B. rewrite the Lambda function using Python
- C. Instead of packaging the libraries in the ZIP file with the function move them to a Lambda layer and use the layer with the function.
- D. Allocate the maximum available CPU units to the function
- E. Increase the available memory to the function.

Answer: D

NEW QUESTION 244

- (Exam Topic 2)

An application needs to use the IP address of the client in its processing. The application has been moved into AWS and has been placed behind an Application Load Balancer (ALB). However, all the client IP addresses now appear to be the same. The application must maintain the ability to scale horizontally. Based on this scenario, what is the MOST cost-effective solution to this problem?

- A. Remove the application from the ALB
- B. Delete the ALB and change Amazon Route 53 to direct traffic to the instance running the application.
- C. Remove the application from the ALB
- D. Create a Classic Load Balancer in its place
- E. Direct traffic to the application using the HTTP protocol.
- F. Alter the application code to inspect the X-Forwarded-For header
- G. Ensure that the code can work properly if a list of IP addresses is passed in the header.
- H. Alter the application code to inspect a custom header
- I. Alter the client code to pass the IP address in the custom header.

Answer: C

NEW QUESTION 247

- (Exam Topic 2)

An application runs on multiple EC2 instances behind an ELB.

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

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

Answer: C

Explanation:

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

NEW QUESTION 252

- (Exam Topic 2)

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

How can the developer troubleshoot the failure?

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

Answer: A

NEW QUESTION 256

- (Exam Topic 2)

A developer is implementing authentication and authorization for an application. The developer needs to ensure that the user credentials are never exposed. Which approach should the developer take to meet this requirement?

- A. Store the user credentials in Amazon DynamoDB. Build an AWS Lambda function to validate the credentials and authorize users.
- B. Deploy a custom authentication and authorization API on an Amazon EC2 instance.
- C. Store the user credentials in Amazon S3 and encrypt the credentials using Amazon S3 server-side encryption.
- D. Use Amazon Cognito to configure a user pool, and use the Cognito API to authenticate and authorize the users.
- E. Store the user credentials in Amazon RDS. Enable the encryption option for the Amazon RDS DB instances. Build an API using AWS Lambda to validate the credentials and authorize users.

Answer: C

NEW QUESTION 261

- (Exam Topic 2)

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

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

- A. Change the StreamViewType parameter value to NEW_AND_OLD_IMAGES for the DynamoDB table

- B. Configure event source mapping for the Lambda function
- C. Map an Amazon SNS topic to the DynamoDB streams
- D. increase the maximum execution time (timeout) setting of the Lambda function

Answer: B

Explanation:

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

NEW QUESTION 262

- (Exam Topic 2)

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

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

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

Answer: B

NEW QUESTION 264

- (Exam Topic 2)

A company needs to ingest terabytes of data each hour from thousands of sources that are delivered almost continually throughout the day. The volume of messages generated varies over the course of the day. Messages must be delivered in real time for fraud detection and live operational dashboards.

Which approach will meet these requirements?

- A. Send the messages to an Amazon SQS queue, then process the messages by using a fleet of Amazon EC2 instances
- B. Use the Amazon S3 API to write messages to an S3 bucket, then process the messages by using Amazon Redshift
- C. Use AWS Data Pipeline to automate the movement and transformation of data
- D. Use Amazon Kinesis Data Streams with Kinesis Client Library to ingest and deliver messages

Answer: D

Explanation:

<https://aws.amazon.com/streaming-data/>

NEW QUESTION 269

- (Exam Topic 2)

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

What is the simplest way to do this?

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

Answer: B

Explanation:

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

NEW QUESTION 274

- (Exam Topic 2)

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

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

Answer: C

NEW QUESTION 278

- (Exam Topic 2)

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

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

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

E. Implement exponential backoff on the GetRecords call and the PutRecords call.

Answer: BD

Explanation:

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

NEW QUESTION 282

- (Exam Topic 2)

A developer has discovered that an application responsible for processing messages in an Amazon SQS queue is routinely falling behind. The application is capable of processing multiple messages in one execution, but is only receiving one message at a time. What should the developer do to increase the number of messages the application receives?

- A. Call the ChangeMessageVisibility API for the queue and set MaxNumberOfMessages to a value greater than the default of 1.
- B. Call the AddPermission API to set MaxNumberOfMessages for the ReceiveMessage action to a value greater than the default of 1.
- C. Call the ReceiveMessage API to set MaxNumberOfMessages to a value greater than the default of 1
- D. Call the SetQueueAttributes API for the queue and set MaxNumberOfMessages to a value greater than the default of 1.

Answer: A

NEW QUESTION 287

- (Exam Topic 2)

A company is launching an ecommerce website and will host the static data in Amazon S3. The company expects approximately 1 000 transactions per second (TPS) for GET and PUT requests in total. Logging must be enabled to track all requests and must be retained for auditing purposes. What is the MOST cost-effective solution?

- A. Enable AWS CloudTrail logging for the S3 bucket-level action and create a lifecycle policy to move the data from the log bucket to Amazon S3 Glacier in 90 days
- B. Enable S3 server access logging and create a lifecycle policy to expire the data in 90 days
- C. Enable AWS CloudTrail logging for the S3 bucket-level action and create a lifecycle policy to expire the data in 90 days
- D. Enable S3 server access logging and create a lifecycle policy to move the data to Amazon S3 Glacier in 90 days.

Answer: C

Explanation:

Reference: <https://docs.aws.amazon.com/AmazonS3/latest/dev/cloudtrail-request-identification.html>

NEW QUESTION 290

- (Exam Topic 2)

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

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

Answer: B

NEW QUESTION 292

- (Exam Topic 2)

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

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

Answer: BD

Explanation:

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

NEW QUESTION 295

- (Exam Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 298

- (Exam Topic 2)

A company is using AWS CodePipeline to deliver one of its applications. The delivery pipeline is triggered by changes to the master branch of an AWS CodeCommit repository and uses AWS CodeBuild to implement the test and build stages of the process and AWS CodeDeploy to deploy the application. The pipeline has been operating successfully for several months and there have been no modifications. Following a recent change to the application's source code, AWS CodeDeploy has not deployed the updates application as expected. What are the possible causes? (Choose two.)

- A. The change was not made in the master branch of the AWS CodeCommit repository.
- B. One of the earlier stages in the pipeline failed and the pipeline has terminated.
- C. One of the Amazon EC2 instances in the company's AWS CodePipeline cluster is inactive.
- D. The AWS CodePipeline is incorrectly configured and is not executing AWS CodeDeploy.
- E. AWS CodePipeline does not have permissions to access AWS CodeCommit.

Answer: AB

NEW QUESTION 299

- (Exam Topic 2)

A company is using continuous integration and continuous delivery systems. A Developer now needs to automate a software package deployment to both Amazon EC2 instances and virtual servers running on-premises. Which AWS service should be used to accomplish this?

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

Answer: D

NEW QUESTION 300

- (Exam Topic 2)

A developer wants to send multi-value headers to an AWS Lambda function that is registered as a target with an Application Load Balancer (ALB). What should the developer do to achieve this?

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

Answer: D

NEW QUESTION 305

- (Exam Topic 2)

A company has a REST application comprised of an Amazon API Gateway and several AWS Lambda functions. A developer is responding to an alert that the API Gateway's HTTP response error rate has unexpectedly increased. The developer must determine which Lambda function is malfunctioning. Which method would help the developer make this determination while minimizing delays?

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

Answer: D

NEW QUESTION 307

- (Exam Topic 2)

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

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

Answer: A

Explanation:

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

NEW QUESTION 312

- (Exam Topic 2)

An application displays a status dashboard. The status is updated by 1 KB messages from an SQS queue. Although the status changes infrequently, the

Developer must minimize the time between the message arrival in the queue and the dashboard update.
What technique provides the shortest delay in updating the dashboard?

- A. Retrieve the messages from the queue using long polling every 20 seconds.
- B. Reduce the size of the messages by compressing them before sending.
- C. Retrieve the messages from the queue using short polling every 10 seconds.
- D. Reduce the size of each message payload by sending it in two parts.

Answer: A

Explanation:

https://docs.aws.amazon.com/ko_kr/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-short-and-long-p

NEW QUESTION 314

- (Exam Topic 2)

A developer is building an application that needs to store data in Amazon S3. Management requires that the data be encrypted before it is sent to Amazon S3 for storage. The encryption keys need to be managed by the security team.
Which approach should the developer take to meet these requirements?

- A. Implement server-side encryption using customer-provided encryption keys (SSE-C).
- B. Implement server-side encryption by using client-side master key.
- C. Implement client-side encryption using an AWS KMS managed customer master key (CMK).
- D. Implement Client-side encryption using Amazon S3 managed keys.

Answer: C

Explanation:

Reference: <https://aws.amazon.com/s3/faqs/>

NEW QUESTION 315

- (Exam Topic 2)

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

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

Answer: D

NEW QUESTION 317

- (Exam Topic 2)

To include objects defined by the AWS Serverless Application Model (SAM) in an AWS CloudFormation template, in addition to Resources, what section MUST be included in the document root?

- A. Conditions
- B. Globals
- C. Transform
- D. Properties

Answer: C

Explanation:

<https://github.com/awslabs/serverless-application-model/blob/master/versions/2016-10-31.md> <https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/sam-specification-template-an>

NEW QUESTION 321

- (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 324

- (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 328

- (Exam Topic 2)

An on-premises application makes repeated calls to store files to Amazon S3. As usage of the application has increased, "LimitExceeded" errors are being logged. What should be changed to fix this error?

- A. Implement exponential backoffs in the application.
- B. Load balance the application to multiple servers.
- C. Move the application to Amazon EC2.
- D. Add a one second delay to each API call.

Answer: A

NEW QUESTION 330

- (Exam Topic 2)

A developer is creating an AWS Lambda function that generates a new file each time it runs. Each new file must be checked into an AWS CodeCommit repository hosted in the same AWS account.

How should the developer accomplish this?

- A. When the Lambda function starts, use the Git CLI to clone the repository.
- B. Check the new file into the cloned repository and push the change.
- C. After the new file is created in Lambda, use cURL to invoke the CodeCommit API.
- D. Send the file to the repository.
- E. Use an AWS SDK to instantiate a CodeCommit client.
- F. Invoke the put_file method to add the file to the repository.
- G. Upload the new file to an Amazon S3 bucket.
- H. Create an AWS step function to accept S3 events.
- I. In the step function, add the new file to the repository.

Answer: D

NEW QUESTION 335

- (Exam Topic 2)

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

Which encryption option will meet these requirements?

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

Answer: B

NEW QUESTION 340

- (Exam Topic 2)

A company stores all personally identifiable information (PII) in an Amazon DynamoDB table named PII in Account A. An application running on Amazon EC2 instances in Account B requires access to the PII table. An administrator in Account A created an IAM role named AccessPII with privileges to access the PII table and made Account B a trusted entity.

Which combination of additional steps should developers take to access the table? (Select TWO)

- A. Ask an administrator in Account B to allow the EC2 IAM role permission to assume the AccessPII role.
- B. Ask an administrator in Account B to allow the EC2 IAM role permission to assume the AccessPII role with predefined service control policies.
- C. Ask an administrator in Account A to allow the EC2 IAM role permission to assume the AccessPII role with predefined service control policies.
- D. Include the AssumeRole API in the application code logic to obtain credentials to access the PII table.
- E. Include the GetSessionToken API in the application code logic to obtain credentials to access the PII table.

Answer: AD

NEW QUESTION 343

- (Exam Topic 2)

A company experienced partial downtime during the last deployment of a new application. AWS Elastic Beanstalk split the environment's Amazon EC2 instances into batches and deployed a new version one batch at a time after taking them out of service. Therefore, full capacity was not maintained during deployment. The developer plans to release a new version of the application, and is looking for a policy that will maintain full capacity and minimize the impact of the failed deployment.

Which deployment policy should the developer use?

- A. Immutable
- B. All at Once
- C. Rolling

D. Rolling with an Additional Batch

Answer: A

Explanation:

Immutable infrastructure has become a new norm in IT operations. Immutable Deployment is one of those approaches, and it simply means: Immutable: the “staging” environment, once ready to become production, doesn't change. If we need to change something, we then deploy new code on completely new infrastructure. The benefits of an immutable infrastructure include more consistency and reliability in your infrastructure and a simpler, more predictable deployment process

NEW QUESTION 345

- (Exam Topic 2)

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

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

Answer: C

Explanation:

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

NEW QUESTION 348

- (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 349

- (Exam Topic 2)

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

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

Answer: C

NEW QUESTION 353

- (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 358

- (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 360

- (Exam Topic 3)

A developer is working with a Docker application that needs to be quickly deployed using AWS without changing the infrastructure or configuring health checks. The application should be configured so that changes and updates can be made automatically without any downtime. Which solution will meet these requirements?

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

Answer: A

NEW QUESTION 362

- (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 366

- (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 370

- (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 374

- (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 376

- (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 380

- (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 382

- (Exam Topic 3)

A developer at a company writes an AWS CloudFormation template. The template refers to subnets that were created by a separate AWS CloudFormation template that the company's network team wrote. When the developer attempts to launch the stack for the first time, the launch fails. Which template coding mistakes could have caused this failure? (Select TWO.)

- A. The developer's template does not use the Ref intrinsic function to refer to the subnets
- B. The developer's template does not use the ImportValue intrinsic function to refer to the subnets
- C. The Mappings section of the developer's template does not refer to the subnets.
- D. The network team's template does not export the subnets in the Outputs section
- E. The network team's template does not export the subnets in the Mappings section

Answer: BD

NEW QUESTION 386

- (Exam Topic 3)

A developer has written a multi-threaded application that is running on a fleet of Amazon EC2 instances. The operations team has requested a graphical method to monitor the number of running threads over time. What is the MOST efficient way to fulfill this request?

- A. Periodically send the thread count to AWS X-Ray segments, then generate a service graph on demand
- B. Create a custom Amazon CloudWatch metric and periodically perform a PutMetricData call with the current thread count.
- C. Periodically log thread count data to Amazon S3. Use Amazon Kinesis to process the data into a graph.
- D. Periodically write the current thread count to a table using Amazon DynamoDB and use Amazon CloudFront to create a graph

Answer: D

NEW QUESTION 388

- (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 391

- (Exam Topic 3)

A company's ecommerce website is experiencing massive traffic spikes, which are causing performance problems in the company database. Users are reporting that accessing the website takes a long time

A developer wants to implement a caching layer using Amazon ElastiCache. The website is required to be responsive no matter which product a user views, and the updates to product information and prices must be strongly consistent

- A. Which cache writing policy will satisfy these requirements?
- B. Write to the cache directly and sync the backend at a later time.
- C. Write to the backend first and wait for the cache to expire.
- D. Write to the cache and the backend at the same time
- E. Write to the backend first and invalidate the cache

Answer: E

NEW QUESTION 396

- (Exam Topic 3)

A developer is building a highly secure healthcare application using .. 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 .. 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 400

- (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 404

- (Exam Topic 3)

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

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

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

Answer: A

NEW QUESTION 408

- (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 412

- (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 414

- (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 418

- (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 419

- (Exam Topic 3)

A developer has created an AWS Lambda function that is written in Python The Lambda function reads data from objects in Amazon S3 and writes data to an Amazon DynamoDB table The function is successfully invoked from an S3 event notification when an object is created However, the function fails when it attempts to write to the DynamoDB table What is the MOST likely cause of this issue?

- A. The Lambda function's concurrency limit has been exceeded
- B. The DynamoDB table requires a global secondary index (GSI) to support writes.
- C. The Lambda function does not have IAM permissions to write to DynamoDB ID.
- D. The DynamoDB table is not running in the same Availability Zone as the Lambda function.

Answer: C

NEW QUESTION 424

- (Exam Topic 3)

A company has an application that is based on Amazon EC2 The company provides API access to the application through Amazon API Gateway and uses Amazon DynamoDB to store the application's data A developer is investigating performance issues that are affecting the application. During peak usage, the application is overwhelmed by a large number of identical data read requests that come through APIs What is the MOST operationally efficient way for the developer to improve the application's performance"

- A. Use DynamoDB Accelerator (DAX) to cache database responses
- B. Configure Amazon EC2 Auto Scaling policies to meet fluctuating demand
- C. Enable API Gateway caching to cache API responses
- D. Use Amazon ElastiCache to cache application responses.

Answer: D

NEW QUESTION 429

- (Exam Topic 3)

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

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

Answer: D

NEW QUESTION 432

- (Exam Topic 3)

A developer is designing a web application in which new users will use their email addresses to create accounts Millions of users are expected to sign up. The application will store attributes for each user Which AWS service or feature should the developer implement to meet these requirements?

- A. Amazon Cognito user pools
- B. AWS Mobile Hub User File Storage
- C. AWS AppSync
- D. AWS Mobile Hub Cloud Logic

Answer: A

NEW QUESTION 433

- (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 434

- (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 437

- (Exam Topic 3)

A company wants to migrate an existing web application to AWS. The application consists of two web servers and a MySQL database.

The company wants the application to automatically scale in response to demand. The company also wants to reduce its operational overhead for database backups and maintenance. The company needs the ability to deploy multiple versions of the application concurrently.

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

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

Answer: B

NEW QUESTION 438

- (Exam Topic 3)

How does Envelope Encryption work in AWS KMS?

- A. The Customer Master Key is used to encrypt/decrypt a data key. The Plaintext Data Key is used to encrypt customer data.
- B. Two encryption keys are used. The Customer Master Key encrypts customer data.
- C. The Data Key is used to re-encrypt the encrypted data.
- D. Two encryption keys are used. The Data Key encrypts customer data. The Customer Master Key is used to re-encrypt the encrypted data.
- E. The Customer Master Key is used to encrypt/decrypt a data key.
- F. The Encrypted Data Key is used to encrypt customer data.

Answer: A

NEW QUESTION 441

- (Exam Topic 3)

A company has a web application in an Amazon Elastic Container Service (Amazon ECS) cluster running hundreds of secure services in AWS Fargate containers.

The services are in target groups routed by an Application Load Balancer (ALB). Application users log in to the website anonymously, but they must be authenticated using any OpenID Connect protocol-compatible identity provider (IdP) to access the secure services.

Which authentication approach would meet these requirements with the LEAST amount of effort?

- A. Configure the services to use Amazon Cognito.
- B. Configure the ALB to use Amazon Cognito.
- C. Configure the services to use AWS Security Token Service (AWS STS) with the OpenID Connect IdP.
- D. Configure the Amazon ECS cluster to use AWS Security Token Service (AWS STS) with the OpenID Connect IdP.

Answer: A

NEW QUESTION 443

- (Exam Topic 3)

A developer is building an application. The application's front end is developed in JavaScript, and the data is stored in an Amazon DynamoDB table. During testing, the application returns an HTTP 5xx error from the strongly consistent reads to the DynamoDB table; "Internal server error (Service: AmazonDynamoDBv2, Status Code: 500; Error Code: InternalServerError)."

Which actions should the developer take to mitigate this error? (Select TWO)

- A. Avoid strongly consistent reads
- B. Use DynamoDB Accelerator (DAX)
- C. Increase read/write capacity of DynamoDB to meet the peak load.
- D. Retry the failed read requests with exponential backoff
- E. Configure DynamoDB auto scaling

Answer: AD

NEW QUESTION 446

- (Exam Topic 3)

A developer is working on an ecommerce website. The developer wants to review server logs without logging in to each of the application servers individually. The website runs on multiple Amazon EC2 instances, is written in Python, and needs to be highly available.

How can the developer update the application to meet these requirements with MINIMUM changes?

- A. Rewrite the application to be cloud native and to run on AWS Lambda where the logs can be reviewed in Amazon CloudWatch.
- B. Set up centralized logging by using Amazon Elasticsearch Service (Amazon ES), Logstash, and Kibana
- C. Scale down the application to one larger EC2 instance where only one instance is recording logs.
- D. Install the unified Amazon CloudWatch agent on the EC2 instance
- E. Configure the agent to push the application logs to CloudWatch.

Answer: D

NEW QUESTION 449

- (Exam Topic 3)

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

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

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

Answer: C

NEW QUESTION 450

- (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 451

- (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 452

- (Exam Topic 3)

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

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

- A. Retry the BatchWriteItem operation immediately
- B. Perform the PutItem operation on the unprocessed items individually instead of using the BatchWriteItem operation

- C. Delay the BatchWriteln operation by using progressively longer wait times between retries, or exponential backoff
- D. Delete the items that were successfully processed, and reissue a new BatchWriteln operation

Answer: D

NEW QUESTION 454

- (Exam Topic 3)

A development team is migrating a monolithic application to Amazon API Gateway with AWS Lambda integrations using the AWS CDK. The zip deployment package exceeds the Lambda direct upload deployment package size limit. How should the Lambda function be deployed?

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

Answer: D

NEW QUESTION 456

- (Exam Topic 3)

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

Which DynamoDB feature should be used to meet these requirements?

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

Answer: B

Explanation:

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

NEW QUESTION 460

- (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 463

- (Exam Topic 3)

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

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

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

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

Answer: A

NEW QUESTION 468

- (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 472

- (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 476

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual AWS-Certified-Developer-Associate Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the AWS-Certified-Developer-Associate Product From:

<https://www.2passeasy.com/dumps/AWS-Certified-Developer-Associate/>

Money Back Guarantee

AWS-Certified-Developer-Associate Practice Exam Features:

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