



Amazon-Web-Services

Exam Questions DVA-C01

AWS Certified Developer Associate Exam

NEW QUESTION 1

- (Exam Topic 4)

A developer needs to secure the static assets in a company's Amazon S3 bucket that is named DOC-EXAMPLE-BUCKET. The company has an Amazon CloudFront distribution that serves the S3 bucket's assets to the public. The developer has already created the origin access identity (OAI) and has associated the OAI with the distribution. The developer must write a bucket policy that allows only the CloudFront distribution to access the S3 bucket

Which policy will meet this requirement MOST securely?

A)

```
{
  "Effect": "Allow",
  "Principal": {
    "AWS": "arn:aws:iam::cloudfront:user/CloudFront Origin Access Identity E12345OAI12"
  },
  "Action": [
    "s3:GetObject"
  ],
  "Resource": [
    "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"
  ]
}
```

B)

```
{
  "Effect": "Allow",
  "Principal": {
    "AWS": "arn:aws:iam::cloudfront:user/CloudFront Origin Access Identity E12345OAI12"
  },
  "Action": [
    "s3:GetObject",
    "s3:PutObject"
  ],
  "Resource": [
    "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"
  ]
}
```

C)

```
{
  "Effect": "Allow",
  "Principal": {
    "AWS": "arn:aws:iam::cloudfront:user/CloudFront Origin Access Identity E12345OAI12"
  },
  "Action": [
    "s3:GetObject"
  ],
  "Resource": [
    "arn:aws:s3:::DOC-EXAMPLE-BUCKET"
  ]
}
```

D)

```
{
  "Effect": "Allow",
  "Principal": {
    "AWS": "arn:aws:iam::cloudfront:user/CloudFront Origin Access Identity E12345OAI12"
  },
  "Action": [
    "s3:GetObject",
    "s3:PutObject"
  ],
  "Resource": [
    "arn:aws:s3:::DOC-EXAMPLE-BUCKET"
  ]
}
```

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

Answer: C

NEW QUESTION 2

- (Exam Topic 4)

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

Which solution will meet these requirements?

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

Answer: D

Explanation:

<https://docs.aws.amazon.com/cloudhsm/latest/userguide/best-practices.html> Q: Can other AWS services use CloudHSM to store and manage keys? AWS services integrate with AWS Key Management Service, which in turn is integrated with AWS CloudHSM through the KMS custom key store feature. If you want to use the server-side encryption offered by many AWS services (such as EBS, S3, or Amazon RDS), you can do so by configuring a custom key store in AWS KMS.

NEW QUESTION 3

- (Exam Topic 4)

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

(Select TWO.)

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

Answer: BD

NEW QUESTION 4

- (Exam Topic 4)

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

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

Which solution will meet these requirements MOST cost-effectively?

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

Answer: D

Explanation:

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

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

NEW QUESTION 5

- (Exam Topic 4)

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

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

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

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

Answer: B

NEW QUESTION 6

- (Exam Topic 4)

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

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

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

Answer: C

NEW QUESTION 7

- (Exam Topic 4)

A company requires objects that are stored in Amazon S3 to be encrypted. The company is currently using server-side encryption with AWS KMS managed encryption keys (SSE-KMS). A developer needs to optimize the cost-effectiveness of the encryption mechanism without negatively affecting performance. What should the developer do to meet these requirements?

- A. Change the encryption type to customer-provided keys.
- B. Configure the S3 bucket to use an S3 Bucket Key for SSE-KMS.
- C. Use S3 bucket policies to limit the principals who can create objects.
- D. Use a custom policy to limit the number of AWS KMS calls that are allowed.

Answer: B

NEW QUESTION 8

- (Exam Topic 4)

A media company wants to test its web application more frequently. The company deploys the application by using a separate AWS CloudFormation stack for each environment. The same CloudFormation template is deployed to each stack as the application progresses through the development lifecycle. A developer needs to build an automated alert for the quality assurance (QA) team. The developer wants the alert to occur for new deployments in the final pre-production environment. Which solution will meet these requirements?

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

Answer: A

Explanation:

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

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

NEW QUESTION 9

- (Exam Topic 4)

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

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

Answer: A

NEW QUESTION 10

- (Exam Topic 4)

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

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

Which solution should the developer implement to meet these requirements?

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

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

Answer: C

NEW QUESTION 10

- (Exam Topic 4)

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

Which solution will meet these requirements MOST cost-effectively?

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

Answer: D

NEW QUESTION 14

- (Exam Topic 4)

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

A)

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

B)

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

C)

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

D)

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

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

Answer: A

Explanation:

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

NEW QUESTION 16

- (Exam Topic 4)

A developer is migrating a Windows-based legacy application from on premises to AWS. The application will run on Amazon EC2 instances that run Amazon Linux. The application stores a large number of files in an NFS drive. The migration solution must minimize downtime and application code changes.

Which solution should the developer use to migrate the application data?

- A. Create an Amazon S3 bucket
- B. Use the s3 sync command to upload the files to the S3 bucket.
- C. Create an Amazon Elastic Block Store (Amazon EBS) volume
- D. Upload the files to the volume

- E. Attach the volume to the EC2 instances.
- F. Create an Amazon Elastic File System (Amazon EFS) file system.
- G. Use AWS DataSync to transfer the files to Amazon EFS.
- H. Create an Amazon Elastic File System (Amazon EFS) file system.
- I. Mount the EFS file system from the legacy application.
- J. Copy the files to the EFS mount.

Answer: C

NEW QUESTION 18

- (Exam Topic 4)

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

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

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

Answer: AC

Explanation:

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

NEW QUESTION 23

- (Exam Topic 4)

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

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

Which solution will meet these requirements?

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

Answer: A

NEW QUESTION 24

- (Exam Topic 4)

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

Which HTTP header should the developer use for this analysis?

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

Answer: C

Explanation:

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

NEW QUESTION 28

- (Exam Topic 4)

A company is using an Amazon API Gateway REST API endpoint as a webhook to publish events from an on-premises source control management (SCM) system to Amazon EventBridge. The company has configured an EventBridge rule to listen for the events and to control application deployment in a central AWS account. The company needs to receive the same events across multiple receiver AWS accounts. How can a developer meet these requirements without changing the configuration of the SCM system?

- A. Deploy the API Gateway REST API to all the required AWS accounts.
- B. Use the same custom domain name for all the gateway endpoints so that a single SCM webhook can be used for all events from all accounts.
- C. Deploy the API Gateway REST API to all the receiver AWS accounts. Create as many SCM webhooks as the number of AWS accounts.
- D. Grant permission to the central AWS account for EventBridge to access the receiver AWS accounts. Add an EventBridge event bus on the receiver AWS accounts as the targets to the existing EventBridge rule.
- E. Convert the API Gateway type from REST API to HTTP API.

Answer: A

NEW QUESTION 31

- (Exam Topic 4)

A business intelligence application runs on Amazon Elastic Container Service (Amazon ECS) on AWS Fargate. Application-level audits require a searchable log of

all API calls from users to the application. The application's developers must store the logs centrally on AWS. Which solution will meet these requirements?

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

Answer: B

Explanation:

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

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

NEW QUESTION 36

- (Exam Topic 4)

A company is running its website on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances run in an Amazon EC2 Auto Scaling group. A developer needs to secure the internet-facing connection with HTTPS. The developer uses AWS Certificate Manager (ACM) to issue an X.509 certificate. What should the developer do to secure the connection?

- A. Configure the ALB to use the X.509 certificate by using the AWS Management Console.
- B. Configure each EC2 instance to use the same X.509 certificate by using the AWS Management Console.
- C. Export the root key of the X.509 certificate to an Amazon S3 bucket
- D. Configure each EC2 instance to use the same X.509 certificate from the S3 bucket.
- E. Export the root key of the X.509 certificate to an Amazon S3 bucket
- F. Configure the ALB to use the X.509 certificate from the S3 bucket.

Answer: A

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/configure-acm-certificates-ec2/> <https://aws.amazon.com/premiumsupport/knowledge-center/associate-acm-certificate-alb-nlb/>

Configuring an Amazon Issued ACM public certificate for a website that's hosted on an EC2 instance requires exporting the certificate. However, you can't export the certificate because ACM manages the private key that signs and creates the certificate.

Instead, you can associate an ACM certificate with a load balancer or an ACM SSL/TLS certificate with a CloudFront distribution.

Associate an ACM SSL certificate with an Application Load Balancer

Open the Amazon EC2 console.

In the navigation pane, choose Load Balancers, and then choose your Application Load Balancer. Choose Add listener.

For Protocol, choose HTTPS. For port, choose 443.

For Default action(s), choose Forward to, and then select your ALB target group from the dropdown list. For Default SSL certificate, choose From ACM (recommended) and then choose the ACM certificate. Choose Save.

NEW QUESTION 40

- (Exam Topic 4)

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

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

Which approach should the developer take to meet these requirements?

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

Answer: B

NEW QUESTION 44

- (Exam Topic 4)

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

Which solution will meet these requirements?

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

D. Create API Gateway URL path parameters for data validation

Answer: D

Explanation:

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

NEW QUESTION 48

- (Exam Topic 4)

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

A developer needs to implement a solution that optimizes the photos that are served to each device to reduce load time and increase photo quality.

Which solution will meet these requirements MOST cost-effectively?

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

Answer: C

Explanation:

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

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Expiration.html>

NEW QUESTION 51

- (Exam Topic 4)

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

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

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

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

Answer: B

NEW QUESTION 52

- (Exam Topic 4)

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

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

Answer: B

NEW QUESTION 54

- (Exam Topic 4)

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

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

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

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

Answer: C

NEW QUESTION 58

- (Exam Topic 4)

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

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

Answer: A

NEW QUESTION 63

- (Exam Topic 4)

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

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

Answer: AC

NEW QUESTION 65

- (Exam Topic 4)

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

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

Answer: B

NEW QUESTION 68

- (Exam Topic 3)

A company is building a compute-intensive application that will run on a fleet of Amazon EC2 instances. The application uses attached Amazon EBS disks for storing data. The application will process sensitive information and all the data must be encrypted.

What should a developer do to ensure the data is encrypted on disk without impacting performance?

- A. Configure the Amazon EC2 instance fleet to use encrypted EBS volumes for storing data
- B. Add logic to write all data to an encrypted Amazon S3 bucket
- C. Add a custom encryption algorithm to the application that will encrypt and decrypt all data
- D. Create a new Amazon Machine Image (AMI) with an encrypted root volume and store the data on ephemeral disks.

Answer: A

NEW QUESTION 73

- (Exam Topic 3)

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

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

Answer: B

NEW QUESTION 77

- (Exam Topic 3)

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

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

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

Answer: B

NEW QUESTION 81

- (Exam Topic 3)

A developer is creating a role to access Amazon S3 buckets To create the role, the developer uses the AWS CLI create-role command. Which policy should be added to allow the Amazon EC2 service to assume the role?

- A. Managed policy
- B. Trust policy
- C. Inline policy
- D. Service control policy (SCP)

Answer: B

Explanation:

A JSON policy document in which you define the principals that you trust to assume the role. A role trust policy is a required resource-based policy that is attached to a role in IAM. The principals that you can specify in the trust policy include users, roles, accounts, and services.

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_terms-and-concepts.html

NEW QUESTION 82

- (Exam Topic 3)

A developer is building a backend system for the long-term storage of information from an inventory management system. The information needs to be stored so that other teams can build tools to report and analyze the data

How should the developer implement this solution to achieve the FASTEST running time?

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

Answer: A

NEW QUESTION 87

- (Exam Topic 3)

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

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

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

Answer: A

NEW QUESTION 88

- (Exam Topic 3)

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

What does the team need to do to begin using X Ray? (Select TWO)

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

Answer: DE

NEW QUESTION 90

- (Exam Topic 3)

A developer has created a Node js web application on a local development machine. The developer wants to use AWS technology to host the website. The developer needs a solution that requires the least possible operational overhead and no code changes.

Which AWS service should the developer use to meet these requirements?

- A. AWS Elastic Beanstalk
- B. Amazon EC2
- C. AWS Lambda
- D. Amazon Elastic Kubernetes Service (Amazon EKS)

Answer: A

NEW QUESTION 95

- (Exam Topic 3)

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

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

Answer: A

NEW QUESTION 98

- (Exam Topic 3)

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

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

Answer: A

NEW QUESTION 100

- (Exam Topic 3)

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

Which solution below is the most secure?

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

Answer: C

NEW QUESTION 104

- (Exam Topic 3)

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

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

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

Answer: C

NEW QUESTION 107

- (Exam Topic 3)

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

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

Answer: A

NEW QUESTION 112

- (Exam Topic 3)

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

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

Answer: C

NEW QUESTION 116

- (Exam Topic 3)

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

```
boto.exception.BotoServerError: BotoServerError: 503 Service Unavailable
<?xml version="1.0" encoding="UTF-8"?>
<Response><Errors><Error><Code>RequestLimitExceeded</Code>
<Message>Request limit exceeded.</Message></Error></Errors><RequestID>bfddec94-53b3-4701-b729-dceefb696ced</RequestID>
</Response>
```

What should the developer do to correct this error message?

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

Answer: D

NEW QUESTION 117

- (Exam Topic 3)

A company is migrating the content delivery network for its dynamic PHP website to AWS. An Amazon CloudFront web distribution is part of the new infrastructure. The distribution has the following cache behavior settings:

- Allowed HTTP Methods is set to GET, HEAD
- Viewer Protocol Policy is set to HTTP and HTTPS

Developers test the solution and can reach the company's website over HTTP and HTTPS. However, the developers are unable to log in to the previously working administration panel of the website.

Which action will resolve this login issue?

- A. Set Allowed HTTP Methods to GET, HEAD, OPTIONS
- B. Set Viewer Protocol Policy to HTTPS Only
- C. Set Allowed HTTP Methods to GET, HEAD, OPTIONS, PUT, POST, PATCH, DELETE
- D. Set Viewer Protocol Policy to Redirect HTTP to HTTPS

Answer: A

NEW QUESTION 120

- (Exam Topic 3)

Given the following AWS CloudFormation template:

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

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

Answer: D

NEW QUESTION 124

- (Exam Topic 3)

A developer implemented a static website hosted in Amazon S3 that makes web service requests hosted in Amazon API Gateway AWS Lambda. The site is showing an error that reads:

"No 'Access-Control-Allow-Origin' header is present on the requested resource. Origin 'null' is therefore not allowed access." What should the developer do to resolve this issue?

- A. Enable cross-origin resource sharing (CORS) on the S3 bucket.
- B. Enable cross-origin resource sharing (CORS) for the method in API Gateway.
- C. Add the Access-Control-Request-Method header to the request.
- D. Add the Access-Control-Request-Headers header to the request.

Answer: B

Explanation:

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

NEW QUESTION 126

- (Exam Topic 3)

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

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

Answer: A

NEW QUESTION 129

- (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 132

- (Exam Topic 3)

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

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

Answer: C

NEW QUESTION 137

- (Exam Topic 3)

A developer is writing an application to analyze the traffic to a fleet of Amazon EC2 instances. The EC2 instances run behind a public Application Load Balancer (ALB). An HTTP server runs on each of the EC2 instances, logging all requests to a log file. The developer wants to capture the client public IP addresses. The developer analyzes the log files and notices only the IP address of the ALB. What must the developer do to capture the client public IP addresses in the log file?

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

Answer: C

NEW QUESTION 141

- (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 146

- (Exam Topic 3)

A developer used the `BatchWriteItem` API operation to insert items in an Amazon DynamoDB table. DynamoDB returned a few items as unprocessed due to throttling. The developer decides to retry the records on the unprocessed items. What should the developer do to reprocess the records with the LEAST number of API calls?

- A. Retry the `BatchWriteItem` operation immediately
- B. Perform the `PutItem` operation on the unprocessed items individually instead of using the `BatchWriteItem` operation
- C. Delay the `BatchWriteItem` operation by using progressively longer wait times between retries, or exponential backoff

D. Delete the items that were successfully processed, and reissue a new BatchWriteItem operation

Answer: D

NEW QUESTION 151

- (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 154

- (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 159

- (Exam Topic 3)

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

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

A)

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

B)

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

C)

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

D)

```

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

```

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

Answer: A

NEW QUESTION 161

- (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 165

- (Exam Topic 3)

A company is developing a new web application in Python. A developer must deploy the application using AWS Elastic Beanstalk from the AWS Management Console. The developer creates an Elastic Beanstalk source bundle to upload using the console. Which of the following are requirements when creating the source bundle? (Select TWO.)

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

Answer: BD

NEW QUESTION 167

- (Exam Topic 3)

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

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

Answer: D

NEW QUESTION 170

- (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 172

- (Exam Topic 3)

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

What will be the result of this mistake?

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

Answer: A

NEW QUESTION 176

- (Exam Topic 3)

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

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

Answer: A

NEW QUESTION 178

- (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 183

- (Exam Topic 3)

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

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

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

Answer: B

NEW QUESTION 184

- (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 185

- (Exam Topic 3)

A developer is building a WebSocket API using Amazon API Gateway. The payload sent to this API is JSON that includes an action key. This key can have three

different values create, update, and remove The developer must integrate with different routes based on the value of the action key of the incoming JSON payload. How can the developer accomplish this task with the LEAST amount of configuration?

- A. Deploy the WebSocket API to three stages for the respective routes create, update, and remove
- B. Create a new route key and set the name as action
- C. Set the value of the route selection expression to action
- D. Set the value of the route selection expression to `$request.body.action`

Answer: D

NEW QUESTION 186

- (Exam Topic 3)

A company process incoming documents from an Amazon S3 bucket. Users upload documents to an S3 bucket using a web user interface. Upon receiving files in S3, and AWS Lambda function is invoked to process the files, but the Lambda function times out intermittently. If the Lambda function is configured with the default settings, what will happen to the S3 event when there is a timeout exception?

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

Answer: D

NEW QUESTION 190

- (Exam Topic 3)

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

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

Answer: AD

NEW QUESTION 192

- (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 197

- (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 202

- (Exam Topic 3)

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

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

Answer: C

NEW QUESTION 207

- (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 212

- (Exam Topic 3)

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

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

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

Answer: B

NEW QUESTION 213

- (Exam Topic 3)

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

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

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

Answer: A

NEW QUESTION 214

- (Exam Topic 3)

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

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

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

Answer: B

NEW QUESTION 218

- (Exam Topic 3)

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

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

Answer: D

NEW QUESTION 223

- (Exam Topic 3)

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

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

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

Answer: B

NEW QUESTION 226

- (Exam Topic 3)

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

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

Answer: B

NEW QUESTION 230

- (Exam Topic 2)

A company has a two-tier application running on an Amazon EC2 server that handles all of its AWS based e-commerce activity During peak times, the backend servers that process orders are overloaded with requests. This results in some orders failing to process. A developer needs to create a solution that will re-factor the application.

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

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

Answer: CD

NEW QUESTION 234

- (Exam Topic 2)

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

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

Answer: A

NEW QUESTION 236

- (Exam Topic 2)

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

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

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

Answer: A

NEW QUESTION 240

- (Exam Topic 2)

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

Which approach should the Developer use to achieve this goal?

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

Answer: C

NEW QUESTION 244

- (Exam Topic 2)

While developing an application that runs on Amazon EC2 in an Amazon VPC, a Developer identifies the need for centralized storage of application-level logs. Which AWS service can be used to securely store these logs?

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

Answer: B

NEW QUESTION 245

- (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 247

- (Exam Topic 2)

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

How can DynamoDB costs be minimized while maximizing application performance?

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

Answer: C

Explanation:

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

NEW QUESTION 250

- (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 255

- (Exam Topic 2)

A company wants to migrate an imaging service to Amazon EC2 while following security best practices. The images are sourced and read from a non-public Amazon S3 bucket.

What should a developer do to meet these requirements?

- A. Create an IAM user with read-only permissions for the S3 bucket Temporarily store the user credentials in the Amazon EBS volume of the EC2 instance
- B. Create an IAM user with read-only permissions for the S3 bucket
- C. Temporarily store the user credentials in the user data of the EC2 instance.
- D. Create an EC2 service role with read-only permissions for the S3 bucket Attach the role to the EC2 instance
- E. Create an S3 service role with read-only permissions for the S3 bucket Attach the role to the EC2 instance

Answer: A

NEW QUESTION 256

- (Exam Topic 2)

A development team is creating a new application designed to run on AWS. While the test and production environments will run on Amazon EC2 instances, developers will each run their own environment on their laptops.

Which of the following is the simplest and MOST secure way to access AWS services from the local development machines?

- A. Use an IAM role to assume a role and execute API calls using the role.
- B. Create an IAM user to be shared with the entire development team, provide the development team with the access key.
- C. Create an IAM user for each developer on the team: provide each developer with a unique access key
- D. Set up a federation through an Amazon Cognito user pool.

Answer: A

NEW QUESTION 261

- (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 265

- (Exam Topic 2)

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

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

Answer: B

Explanation:

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

NEW QUESTION 266

- (Exam Topic 2)

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

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

Answer: A

Explanation:

https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

NEW QUESTION 270

- (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 271

- (Exam Topic 2)

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

How can the Developer address this issue?

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

Answer: C

Explanation:

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

NEW QUESTION 275

- (Exam Topic 2)

A company requires that AWS Lambda functions written by developers log errors so system administrators can more effectively troubleshoot issues. What should the developers implement to meet this need?

- A. Publish errors to a dedicated Amazon SQS queue
- B. Create an Amazon CloudWatch Events event to trigger based on certain Lambda events.
- C. Report errors through logging statements in Lambda function code.
- D. Set up an Amazon SNS topic that sends logging statements upon failure

Answer: B

NEW QUESTION 279

- (Exam Topic 2)

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

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

- A. Add logging statements to the Lambda functions, then use Amazon CloudWatch to view the logs.
- B. Use AWS Cloud Trail and then examine the logs
- C. Use AWS X-Ray, then examine the segments and errors
- D. Run Amazon Inspector agents and then analyze performance

Answer: C

Explanation:

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

NEW QUESTION 281

- (Exam Topic 2)

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

How can the Developer create the solution?

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

Answer: C

Explanation:

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

NEW QUESTION 286

- (Exam Topic 2)

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

How can the Developer MOST efficiently handle the temporary files?

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

Answer: C

NEW QUESTION 291

- (Exam Topic 2)

A company has developed a new serverless application using AWS Lambda functions that will be deployed using the AWS Serverless Application Model (AWS SAM) CLI. Which step should the developer complete prior to deploying the application?

- A. Compress the application to a .zip file and upload it into AWS Lambda
- B. Test the new AWS Lambda function by first tracing it in AWS X-Ray.
- C. Bundle the serverless application using a SAM package
- D. Create the application environment using the `eb create my-env` command.

Answer: B

NEW QUESTION 295

- (Exam Topic 2)

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 296

- (Exam Topic 2)

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

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

Answer: B

NEW QUESTION 299

- (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 300

- (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 303

- (Exam Topic 2)

A company has implemented AWS CodePipeline to automate its release pipelines. The development team is writing an AWS Lambda function that will send notifications for state changes of each of the actions in the stages. Which steps must be taken to associate the Lambda function with the event source?

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

Answer: B

NEW QUESTION 305

- (Exam Topic 2)

A development team wants to immediately build and deploy an application whenever there is a change to the source code. Which approaches could be used to trigger the deployment? (Select TWO.)

- A. Store the source code in an Amazon S3 bucket. Configure AWS CodePipeline to start whenever a file in the bucket changes.
- B. Store the source code in an encrypted Amazon EBS volume. Configure AWS CodePipeline to start whenever a file in the volume changes.
- C. Store the source code in an AWS CodeCommit repository. Configure AWS CodePipeline to start whenever a change is committed to the repository.
- D. Store the source code in an Amazon S3 bucket. Configure AWS CodePipeline to start every 15 minutes.
- E. Store the source code in an Amazon EC2 instance's ephemeral storage.
- F. Configure the instance to start AWS CodePipeline whenever there are changes to the source code.

Answer: BC

Explanation:

Reference: <https://docs.aws.amazon.com/codepipeline/latest/userguide/tutorials-ecs-ecr-codedeploy.html>

NEW QUESTION 308

- (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 310

- (Exam Topic 2)

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

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

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

Answer: D

NEW QUESTION 313

- (Exam Topic 2)

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

How should the environment variables be passed to the container?

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

Answer: B

NEW QUESTION 314

- (Exam Topic 2)

An organization is using Amazon CloudFront to ensure that its users experience low-latency access to its web application. The organization has identified a need to encrypt all traffic between users and CloudFront, and all traffic between CloudFront and the web application.

How can these requirements be met? (Choose two.)

- A. Use AWS KMS to encrypt traffic between CloudFront and the web application.
- B. Set the Origin Protocol Policy to "HTTPS Only".
- C. Set the Origin's HTTP Port to 443.
- D. Set the Viewer Protocol Policy to "HTTPS Only" or "Redirect HTTP to HTTPS".
- E. Enable the CloudFront option Restrict Viewer Access.

Answer: AB

Explanation:

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/using-https-viewers-to-cloudfront.htm>

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/using-https-cloudfront-to-custom-origi>

NEW QUESTION 317

- (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 318

- (Exam Topic 2)

A Developer wants to debug an application by searching and filtering log data. The application logs are stored in Amazon CloudWatch Logs. The Developer creates a new metric filter to count exceptions in the application logs. However, no results are returned from the logs. What is the reason that no filtered results are being returned?

- A. A setup of the Amazon CloudWatch interface VPC endpoint is required for filtering the CloudWatch Logs in the VPC
- B. CloudWatch Logs only publishes metric data for events that happen after the filter is created
- C. The log group for CloudWatch Logs should be first streamed to Amazon Elasticsearch Service before metric filtering returns the results
- D. Metric data points for logs groups can be filtered only after they are exported to an Amazon S3 bucket

Answer: B

Explanation:

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/MonitoringLogData.html>

NEW QUESTION 322

- (Exam Topic 2)

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

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

Answer: C

Explanation:

Reference: https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

NEW QUESTION 324

- (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 326

- (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 327

- (Exam Topic 2)

A company has implemented AWS CodeDeploy as part of its cloud native CI/CD stack The company enables automatic rollbacks while deploying a new version of a popular web application from in place to Amazon EC2.

What occurs if the deployment of the new version fails due to code regression?

- A. The last known good deployment is automatically restored using the snapshot stored in Amazon S3
- B. CodeDeploy switches the Amazon Route 53 alias records back to the known good green deployment and terminates the failed blue deployment
- C. A new deployment of the last known good version of the application is deployed with a new deployment ID
- D. AWS CodePipeline promotes the most recent deployment with a SUCCEEDED status to production

Answer: B

NEW QUESTION 328

- (Exam Topic 2)

A developer is writing an application that will process data delivered into an Amazon S3 bucket. The data is delivered approximately 10 times a day, and the developer expects the data will be processed in less than 1 minute, on average.

How can the developer deploy and invoke the application with the lowest cost and lowest latency?

- A. Deploy the application as an AWS Lambda function and invoke it with an Amazon CloudWatch alarm triggered by an S3 object upload
- B. Deploy the application as an AWS Lambda function and invoke it with an S3 event notification
- C. Deploy the application as an AWS Lambda function and invoke it with an Amazon CloudWatch scheduled event
- D. Deploy the application onto an Amazon EC2 instance and have it poll the S3 bucket for new objects.

Answer: A

Explanation:

Reference: <https://docs.aws.amazon.com/lambda/latest/dg/with-s3.html>

NEW QUESTION 329

- (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 334

- (Exam Topic 2)

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

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

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

Answer: D

Explanation:

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

NEW QUESTION 336

- (Exam Topic 2)

A developer has written an Amazon Kinesis Data Streams application. As usage grows and traffic increases 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: BC

Explanation:

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

NEW QUESTION 339

- (Exam Topic 2)

A developer needs to create an application that supports Security Assertion Markup Language (SAML) and Facebook authentication. It must also allow access to AWS services, such as Amazon DynamoDB.

Which AWS service or feature will meet these requirements with the LEAST amount of additional coding?

- A. AWSAppSync
- B. Amazon Cognito identity pools
- C. Amazon Cognito user pools
- D. Amazon Lambda@Edge

Answer: B

Explanation:

Reference: <https://aws.amazon.com/blogs/mobile/amazon-cognito-user-pools-supports-federation-with-saml/>

NEW QUESTION 340

- (Exam Topic 2)

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

NEW QUESTION 341

- (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 344

- (Exam Topic 2)

An organization is storing large files in Amazon S3, and is writing a web application to display meta-data about the files to end-users. Based on the metadata a user selects an object to download. The organization needs a mechanism to index the files and provide single-digit millisecond latency retrieval for the metadata. What AWS service should be used to accomplish this?

- A. Amazon DynamoDB
- B. Amazon EC2
- C. AWS Lambda
- D. Amazon RDS

Answer: A

Explanation:

Amazon DynamoDB is a fast and flexible NoSQL database service for all applications that need consistent, single-digit millisecond latency at any scale. It is a fully managed database and supports both document and key-value data models. Its flexible data model and reliable performance make it a great fit for mobile, web, gaming, ad-tech, Internet of Things (IoT), and many other applications.

References:

NEW QUESTION 345

- (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 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 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 GetSession token API in the application code logic to obtain credentials to access the PII table

Answer: BE

NEW QUESTION 349

- (Exam Topic 2)

An AWS Lambda function must access an external site by using a regularly rotated user name and password. These items must be kept securely and cannot be stored in the function code.

What combination of AWS services can be used to accomplish this? (Choose two.)

- A. AWS Certificate Manager (ACM)
- B. AWS Systems Manager Parameter Store
- C. AWS Trusted Advisor
- D. AWS KMS
- E. Amazon GuardDuty

Answer: BD

Explanation:

<https://docs.aws.amazon.com/kms/latest/developerguide/services-parameter-store.html>

NEW QUESTION 350

- (Exam Topic 2)

According to best practice, how should access keys be managed in AWS? (Choose two.)

- A. Use the same access key in all applications for consistency.
- B. Delete all access keys for the account root user.
- C. Leave unused access keys in the account for tracking purposes.
- D. Embed and encrypt access keys in code for continuous deployment.
- E. Use Amazon IAM roles instead of access keys where possible.

Answer: BE

NEW QUESTION 354

- (Exam Topic 2)

A company is creating a REST service using an Amazon API Gateway with AWS Lambda integration. The service must run different versions for testing purposes. What would be the BEST way to accomplish this?

- A. Use an X-Version header to denote which version is being called and pass that header to the Lambda function(s)
- B. Create an API Gateway Lambda authorizer to route API clients to the correct API version
- C. Create an API Gateway resource policy to isolate versions and provide context to the Lambda function(s)
- D. Deploy the API versions as unique stages with unique endpoints and use stage variables to provide further context

Answer: A

Explanation:

The service run different versions for testing purposes. means different stages, stage variable is the way to go.

NEW QUESTION 359

- (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 363

- (Exam Topic 2)

A development team is using AWS Elastic Beanstalk to deploy a two-tier application that consists of a load-balanced web tier and an Amazon RDS database tier in production. The team would like to separate the RDS instance from the Elastic Beanstalk. How can this be accomplished?

- A. Use the Elastic Beanstalk CLI to disassociate the database.
- B. Use the AWS CLI to disassociate the database.
- C. Change the deployment policy to disassociate the database.
- D. Recreate a new Elastic Beanstalk environment without Amazon RDS.

Answer: D

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/decouple-rds-from-beanstalk/>

NEW QUESTION 367

- (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 368

- (Exam Topic 2)

A Developer is going to deploy an AWS Lambda function that requires significant CPU utilization. Which approach will MINIMIZE the average runtime of the function?

- A. Deploy the function into multiple AWS Regions

- C. Use AWS root user to make requests to the application.
- D. Store and retrieve credentials from AWS CodeCommit.

Answer: B

Explanation:

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_use_switch-role-ec2.html

NEW QUESTION 384

- (Exam Topic 1)

A supplier is writing a new RESTful API for customers to query the status of orders. The customers requested the following API endpoint.

<http://www.supplierdomain.com/status/customerID>

Which of the following application designs meet the requirements? (Select two.)

- A. Amazon SQS; Amazon SNS
- B. Elastic Load Balancing; Amazon EC2
- C. Amazon ElastiCache; Amazon Elasticsearch Service
- D. Amazon API Gateway; AWS Lambda
- E. Amazon S3; Amazon CloudFront

Answer: DE

NEW QUESTION 386

- (Exam Topic 1)

A large e-commerce site is being designed to deliver static objects from Amazon S3. The Amazon S3 bucket will server more than 300 GET requests per second.

What should be done to optimize performance? (Select TWO.)

- A. Integrate Amazon CloudFront with Amazon S3.
- B. Enable Amazon S3 cross-region replication.
- C. Delete expired Amazon S3 server log files.
- D. Configure Amazon S3 lifecycle rules.
- E. Randomize Amazon S3 key name prefixes.

Answer: AE

Explanation:

CloudWatch definitely. Random key prefixes is still a valid method of improving performance by using parallel reads. It doesn't mention prefix hashing. For instance prefixes 1/,2/,3/,4/,5/ could provide 5 x parallel streams for S3 as opposed to all objects being in a single folder/prefix e.g. dev/

<https://docs.aws.amazon.com/AmazonS3/latest/dev/optimizing-performance.html>

"There are no limits to the number of prefixes in a bucket. You can increase your read or write performance by parallelizing reads. For example, if you create 10 prefixes in an Amazon S3 bucket to parallelize reads, you could scale your read performance to 55,000 read requests per second." The assumption that prefixes don't matter is incorrect, as described by "Amazon S3 performance guidelines recommended randomizing prefix naming with **hashed characters** to optimize performance for frequent data retrievals. You no longer have to randomize prefix naming for performance, and can use sequential date-based naming for your prefixes"

NEW QUESTION 391

- (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 394

- (Exam Topic 1)

A Developer has implemented a Lambda function that needs to add new customers to an RDS database that is expected to run hundreds of times per hour. The Lambda function is configured to use 512MB of RAM and is based on the following pseudo code:

```
def lambda_handler(event, context):
```

```
    db = database.connect()
```

```
    db.statement('INSERT INTO Customers (CustomerName) VALUES  
                (context.name)')
```

```
    db.close()
```

After testing the Lambda function, the Developer notices that the Lambda execution time is much longer than expected. What should the Developer do to improve

performance?

- A. Increase the amount of RAM allocated to the Lambda function, which will increase the number of threads the Lambda can use.
- B. Increase the size of the RDS database to allow for an increased number of database connections each hour.
- C. Move the database connection and close statement out of the handle
- D. Place the connection in the global space.
- E. Replace RDS with Amazon DynamoDB to implement control over the number of writes per second.

Answer: C

Explanation:

Refer AWS documentation - Lambda Best Practices

Take advantage of Execution Context reuse to improve the performance of your function. Make sure any externalized configuration or dependencies that your code retrieves are stored and referenced locally after initial execution. Limit the re-initialization of variables/objects on every invocation. Instead use static initialization/constructor, global/static variables and singletons. Keep alive and reuse connections (HTTP, database, etc.) that were established during a previous invocation.

NEW QUESTION 396

- (Exam Topic 1)

A set of APIs are exposed to customers using the Amazon API Gateway. These APIs have caching enabled on the API Gateway. Customers have asked for an option to invalidate this cache for each of the APIs.

What action can be taken to allow API customers to invalidate the API Cache?

- A. Ask customers to use AWS credentials to call the InvalidateCache API.
- B. Ask customers to invoke an AWS API endpoint which invalidates the cache.
- C. Ask customers to pass an HTTP header called Cache-Control:max-age=0.
- D. Ask customers to add a query string parameter called "INVALIDATE_CACHE" when making an API call.

Answer: C

Explanation:

<https://developer.mozilla.org/en-US/docs/Web/HTTP/Headers/Cache-Control> <https://docs.aws.amazon.com/apigateway/latest/developerguide/api-gateway-caching.html>

Invalidate an API Gateway Cache Entry A client of your API can invalidate an existing cache entry and reload it from the integration endpoint for individual requests. The client must send a request that contains the Cache-Control: max-age=0 header. The client receives the response directly from the integration endpoint instead of the cache, provided that the client is authorized to do so. This replaces the existing cache entry with the new response, which is fetched from the integration endpoint.

NEW QUESTION 400

- (Exam Topic 1)

In a multi-container Docker environment in AWS Elastic Beanstalk, what is required to configure container instances in the environment?

- A. An Amazon ECS task definition
- B. An Amazon ECS cluster
- C. A Docker in an application package
- D. A CLI for Elastic Beanstalk

Answer: A

Explanation:

Reference: https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker_ecs.html

NEW QUESTION 404

- (Exam Topic 1)

Which of the following services are key/value stores? Choose 3 answers

- A. Amazon ElastiCache
- B. Simple Notification Service
- C. DynamoDB
- D. Simple Workflow Service
- E. Simple Storage Service

Answer: ACE

NEW QUESTION 407

- (Exam Topic 1)

A company has written a Java AWS Lambda function to be triggered whenever a user uploads an image to an Amazon S3 bucket. The function converts the original image to several different formats and then copies the resulting images to another Amazon S3 bucket.

The Developers find that no images are being copied to the second Amazon S3 bucket. They have tested the code on an Amazon EC2 instance with 1GB of RAM, and it takes an average of 500 seconds to complete.

What is the MOST likely cause of the problem?

- A. The Lambda function has insufficient memory and needs to be increased to 1 GB to match the Amazon EC2 instance
- B. Files need to be copied to the same Amazon S3 bucket for processing, so the second bucket needs to be deleted.
- C. Lambda functions have a maximum execution limit of 300 seconds, therefore the function is not completing.
- D. There is a problem with the Java runtime for Lambda, and the function needs to be converted to node.js.

Answer: C

NEW QUESTION 409

- (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 414

- (Exam Topic 1)

A Developer has created a software package to be deployed on multiple EC2 instances using IAM roles. What actions could be performed to verify IAM access to get records from Amazon Kinesis Streams? (Select TWO.)

- A. Use the AWS CLI to retrieve the IAM group.
- B. Query Amazon EC2 metadata for in-line IAM policies.
- C. Request a token from AWS STS, and perform a describe action.
- D. Perform a get action using the `--dry-run` argument.
- E. Validate the IAM role policy with the IAM policy simulator.

Answer: AE

NEW QUESTION 417

- (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 420

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

DVA-C01 Practice Exam Features:

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

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