

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 developer is migrating to Amazon Cognito from a custom user management solution that stores user information in a database. The developer has created a..... Amazon Cognito user pool. The developer needs to migrate the existing user information to the user pool without forcing users to change their passwords. Which solution will meet these requirements?

- A. Import users from a .csv file.
- B. Add an OpenID Connect (OIDC) identity provider to the user pool.
- C. Import users from a .json file.
- D. Import users with a user migration AWS Lambda trigger.

Answer: B

NEW QUESTION 3

- (Exam Topic 4)

A company has an application that runs on AWS Elastic Beanstalk in a load-balanced environment. The company needs to update the instance types in the environment to a more recent generation of instance types. The company must minimize downtime during the deployment of this configuration change. Which deployment options will meet these requirements? (Choose two.)

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

Answer: BC

Explanation:

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

NEW QUESTION 4

- (Exam Topic 4)

A developer needs to use Amazon DynamoDB to store customer orders. The developer's company requires all customer data to be encrypted at rest with a key the company generates.

What should the developer do to meet these requirements?

- A. Create the DynamoDB table with encryption set to Non
- B. Code the application to use the key to decrypt the data when the application reads from the table
- C. Code the application to use the key to encrypt the data when the application writes to the table.
- D. Store the key by using AWS Key Management Service (AWS KMS). Choose an AWS KMS customer managed key during creation of the DynamoDB table
- E. Provide the Amazon Resource Name (ARN) of the AWS KMS key.
- F. Store the key by using AWS Key Management Service (AWS KMS). Create the DynamoDB table with default encryption
- G. Include the kms:Encrypt parameter with the Amazon Resource Name (ARN) of the AWS KMS key when using the DynamoDB software development kit (SDK).
- H. Store the key by using AWS Key Management Service (AWS KMS). Choose an AWS KMS AWS managed key during creation of the DynamoDB table Provide the Amazon Resource Name (ARN) of the AWS KMS key.

Answer: C

NEW QUESTION 5

- (Exam Topic 4)

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

Which solution will meet these requirements?

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

Answer: A

NEW QUESTION 6

- (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 7

- (Exam Topic 4)

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

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

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

Answer: D

NEW QUESTION 8

- (Exam Topic 4)

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

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

Answer: CD

NEW QUESTION 9

- (Exam Topic 4)

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

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

Answer: B

NEW QUESTION 10

- (Exam Topic 4)

A developer is integrating Amazon ElastiCache in an application. The cache will store data from a database. The cached data must populate real-time dashboards. Which caching strategy will meet these requirements?

- A. A read-through cache
- B. A write-behind cache
- C. A lazy-loading cache
- D. A write-through cache

Answer: D

Explanation:

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

NEW QUESTION 10

- (Exam Topic 4)

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

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

Which solution will meet these requirements?

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

Answer: B

NEW QUESTION 13

- (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 14

- (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 17

- (Exam Topic 4)

A developer is managing an application that uploads user files to an Amazon S3 bucket named companybucket. The company wants to maintain copies of all the files uploaded by users for compliance purposes, while ensuring users still have access to the data through the application. Which IAM permissions should be applied to users to ensure they can create but not remove files from the bucket?

A.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "statement1",
      "Effect": "Allow",
      "Action": [
        "s3:GetObject",
        "s3:PutObject",
        "s3:DeleteObject"
      ],
      "Resource": [
        "arn:aws:s3:::companybucket"
      ]
    }
  ]
}
```

B. Text Description automatically generated

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "statement1",
      "Effect": "Allow",
      "Action": [
        "s3:CreateBucket",
        "s3:GetBucketLocation"
      ],
      "Resource": [
        "arn:aws:s3:::companybucket"
      ]
    }
  ]
}
```

C. Text, letter Description automatically generated

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "statement1",
      "Effect": "Allow",
      "Action": [
        "s3:GetObject",
        "s3:PutObject",
        "s3:DeleteObject",
        "s3:PutObjectRetention"
      ],
      "Resource": [
        "arn:aws:s3:::companybucket"
      ]
    }
  ]
}
```

D. Text Description automatically generated {

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "statement1",
      "Effect": "Allow",
      "Action": [
        "s3:GetObject",
        "s3:PutObject"
      ],
      "Resource": [
        "arn:aws:s3:::companybucket"
      ]
    }
  ]
}
```

Answer: D

Explanation:

To accomplish: "can create but not remove files"

-- Need: "Put Object"

-- Don't need: "Delete Object" <https://docs.aws.amazon.com/cli/latest/reference/s3api/put-object.html>

NEW QUESTION 20

- (Exam Topic 4)

A developer has an application that is composed of many different AWS Lambda functions. The Lambda functions all use some of the same dependencies. To avoid security issues, the developer is constantly updating the dependencies of all of the Lambda functions. The result is duplicated effort for each function. Now can the developer keep the dependencies of the Lambda functions up to date with the LEAST additional complexity?

- A. Define a maintenance window for the Lambda functions to ensure that the functions get updated copies of the dependencies.
- B. Upgrade the Lambda functions to the most recent runtime version.
- C. Define a Lambda layer that contains all of the shared dependencies.
- D. Use an AWS CodeCommit repository to host the dependencies in a centralized location.

Answer: C

NEW QUESTION 23

- (Exam Topic 4)

A team of developers is using an AWS CodePipeline pipeline as a continuous integration and continuous delivery (CI/CD) mechanism for a web application. A developer has written unit tests to programmatically test the functionality of the application code. The unit tests produce a test report that shows the results of each individual check. The developer now wants to run these tests automatically during the CI/CD process. Which solution will meet this requirement with the LEAST operational effort?

- A. Write a Git pre-commit hook that runs the tests before every commi
- B. Ensure that each developer who is working on the project has the pre-commit hook installed local
- C. Review the test report and resolve any issues before pushing changes to AWS CodeCommit.
- D. Add a new stage to the pipelin
- E. Use AWS CodeBuild as the provide
- F. Add the new stage after the stage that deploys code revisions to the test environmen
- G. Write a buildspec that fails the CodeBuild stage if any test does not pas
- H. Use the test reports feature of CodeBuild to integrate the report with the CodeBuild consol
- I. View the test results in CodeBuil

- J. Resolve any issues.
- K. Add a new stage to the pipeline
- L. Use AWS CodeBuild as the provider
- M. Add the new stage before the stage that deploys code revisions to the test environment
- N. Write a buildspec that fails the CodeBuild stage if any test does not pass
- O. Use the test reports feature of CodeBuild to integrate the report with the CodeBuild console
- P. View the test results in CodeBuild
- Q. Resolve any issues.
- R. Add a new stage to the pipeline
- S. Use Jenkins as the provider
- T. Configure CodePipeline to use Jenkins to run the unit test
- . Write a Jenkinsfile that fails the stage if any test does not pass
- . Use the test report plugin for Jenkins to integrate the report with the Jenkins dashboard
- . View the test results in Jenkins
- . Resolve any issues.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/devops/test-reports-with-aws-codebuild/>

NEW QUESTION 24

- (Exam Topic 4)

A company is creating a continuous integration and continuous delivery (CI/CD) process by using AWS CodePipeline for its application on AWS. The CI/CD process will pull code from an AWS CodeCommit repository, create the application infrastructure by using AWS CloudFormation, deploy the frontend code to an Amazon S3 bucket that is configured for static website hosting, and deploy the application backend on an Amazon Elastic Container Service (Amazon ECS) cluster.

A developer needs to create a new CodePipeline stage that creates the application infrastructure. Which solution will meet these requirements with the LEAST operational overhead?

- A. Create a new action with AWS Lambda as the action provider. Create a Lambda function that makes an AWS SDK API call to create the CloudFormation stack.
- B. Create a new action with CloudFormation as the action provider. Set the action mode to CREATE_UPDATE. Target the CloudFormation stack to be launched.
- C. Create a new action with Jenkins as the action provider.
- D. Create and configure a Jenkins job to make an API call by using the AWS CLI to create the CloudFormation stack.
- E. Create a new action with AWS CodeBuild as the action provider. Configure the buildspec to make an API call by using the AWS CLI to create the CloudFormation stack.

Answer: D

NEW QUESTION 29

- (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 32

- (Exam Topic 4)

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

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

Which solution should the developer implement to meet these requirements?

- A. After successful login, submit a Cognito request for user token.
- B. When redirecting to the welcome screen, provide the identity token in the Authorization header of the request.

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

Answer: C

NEW QUESTION 34

- (Exam Topic 4)

A company has an application that uses Amazon Cognito user pools as an identity provider. The company must secure access to user records. The company implements multi-factor authentication (MFA). The company also wants to send a login activity notification by email every time a user logs in. What is the MOST operationally efficient solution that meets this requirement?

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

Answer: C

NEW QUESTION 37

- (Exam Topic 4)

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

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

Answer: B

NEW QUESTION 38

- (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 42

- (Exam Topic 4)

A developer is implementing an AWS Lambda function that will be invoked when an object is uploaded to Amazon S3. The developer wants to test the Lambda function in a local development machine before publishing the function to a production AWS account. Which solution will meet these requirements with the LEAST operational overhead?

- A. Upload an object to Amazon S3 by using the aws s3api put-object CLI command
- B. Wait for the local Lambda invocation from the S3 event.

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

Answer: D

NEW QUESTION 44

- (Exam Topic 4)

A developer designed an application on an Amazon EC2 instance. The application makes API requests to objects in an Amazon S3 bucket. Which combination of steps will ensure that the application makes the API requests in the MOST secure manner? (Select TWO.)

- A. Create an IAM user that has permissions to the S3 bucket
- B. Add the user to an IAM group.
- C. Create an IAM role that has permissions to the S3 bucket.
- D. Add the IAM role to an instance profile
- E. Attach the instance profile to the EC2 instance.
- F. Create an IAM role that has permissions to the S3 bucket
- G. Assign the role to an IAM group.
- H. Store the credentials of the IAM user in the environment variables on the EC2 instance.

Answer: BC

NEW QUESTION 47

- (Exam Topic 4)

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

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

Answer: A

NEW QUESTION 51

- (Exam Topic 4)

An application that is running on Amazon EC2 instances stores data in an Amazon S3 bucket. All the data must be encrypted in transit. How can a developer ensure that all traffic to the S3 bucket is encrypted?

- A. Install certificates on the EC2 instances.
- B. Create a private VPC endpoint.
- C. Configure the S3 bucket with server-side encryption with AWS KMS managed encryption keys (SSE-KMS).
- D. Create an S3 bucket policy that denies traffic when the value for the aws:SecureTransport condition key is false.

Answer: C

NEW QUESTION 53

- (Exam Topic 4)

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

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

Answer: D

Explanation:

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

NEW QUESTION 54

- (Exam Topic 4)

A company is developing a microservice that will manage customer account data in an Amazon DynamoDB table. Insert, update, and delete requests will be rare.

Read traffic will be heavy. The company must have the ability to access customer data quickly by using a customer ID. The microservice can tolerate stale data. Which solution will meet these requirements with the FEWEST possible read capacity units (RCUs)?

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

Answer: A

Explanation:

Key points: "Read heavy", "access data quickly", "can tolerate stale data" To achieve: "FEWEST" possible (RCUs)
For items up to 4 KB in size, one RCU can perform one strongly consistent read request per second. For items up to 4 KB in size, one RCU can perform two eventually consistent read requests per second. Transactional read requests require two RCUs to perform one read per second for items up to 4 KB. For example, a strongly consistent read of an 8 KB item would require two RCUs, an eventually consistent read of an 8 KB item would require one RCU, and a transactional read of an 8 KB item would require four RCUs. <https://aws.amazon.com/dynamodb/pricing/provisioned/>

NEW QUESTION 55

- (Exam Topic 4)

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

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

Answer: D

NEW QUESTION 58

- (Exam Topic 4)

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

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

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

Answer: A

Explanation:

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

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

NEW QUESTION 63

- (Exam Topic 4)

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

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

Answer: A

NEW QUESTION 65

- (Exam Topic 4)

A company is migrating a legacy application to a serverless application on AWS. The legacy application consists of a set of web services that are exposed by an Amazon API Gateway API. A developer needs to replace the existing implementation of web services with AWS Lambda functions. The developer needs to test a new version of the API that uses the functions in production. The developer must minimize the impact of the testing on the application's users.

Which solution will meet these requirements?

- A. Create a beta stage for the new version of the API.
- B. Send the updated endpoint to the users.
- C. Create a development stage for the new version of the API.
- D. Use a canary deployment.
- E. Create a development stage for the new version of the API.

- F. Promote a canary release.
- G. Create a deployment stag
- H. Enable mutual TLS for the new version of the API.

Answer: B

NEW QUESTION 70

- (Exam Topic 4)

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

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

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

Answer: C

Explanation:

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

Note:

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

NEW QUESTION 74

- (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 79

- (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 81

- (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 stag
- B. Deploy the new production stag
- 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 polic
- 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 AP
- 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 84

- (Exam Topic 4)

A company is using AWS Elastic Beanstalk to deploy a three-tier application. The application uses an Amazon RDS DB instance as the database tier. The com wants to decouple the DB instance from the Elastic Beanstalk environment.

Which combination of steps should a developer take to meet this requirement? (Select TWO.)

- A. Create a new Elastic Beanstalk environment that connects to the DB instance.
- B. Create a new DB instance from a snapshot of the previous DB instance.
- C. Use the Elastic Beanstalk CLI to decouple the DB instance.
- D. Use the AWS CLI to decouple the DB instance.
- E. Modify the current Elastic Beanstalk environment to connect to the DB instance.

Answer: BE

NEW QUESTION 89

- (Exam Topic 4)

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

Which of the following is a possible reason for this failure?

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

Answer: C

NEW QUESTION 94

- (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 97

- (Exam Topic 4)

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

What should the developer do to accomplish these tasks?

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

Answer: C

Explanation:

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

NEW QUESTION 101

- (Exam Topic 4)

A company has a new application. The company needs to secure sensitive configuration data such as database connection strings, application license codes, and API keys that the application uses to access external resources. The company must track access to the configuration data for auditing purposes. The resources are managed outside the application.

The company is not required to manage rotation of the connection strings, license codes, and API keys in the application. The company must implement a solution to securely store the configuration data and to give the application access to the configuration data. The solution must comply with security best practices.

Which solution will meet these requirements MOST cost-effectively?

- A. Store the configuration data in an encrypted file on the source code bundle
- B. Grant the application access by using IAM policies.
- C. Store the configuration data in AWS Systems Manager Parameter Store
- D. Grant the application access by using IAM policies.
- E. Store the configuration data on an Amazon Elastic Block Store (Amazon EBS) encrypted volume. Attach the EBS volume to an Amazon EC2 instance to provide the application with access to the data.
- F. Store the configuration data in AWS Secrets Manager
- G. Grant the application access by using IAM policies.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/mt/the-right-way-to-store-secrets-using-parameter-store/> <https://docs.aws.amazon.com/managedservices/latest/userguide/sys-man-param-store.html> AWS Systems Manager Parameter Store (AMS SSPS):

AWS Systems Manager Parameter Store provides secure, hierarchical storage for configuration data management and secrets management. You can store data such as passwords, database strings, and license codes as parameter values.

NEW QUESTION 102

- (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 106

- (Exam Topic 4)

A developer needs to deploy an application to AWS Elastic Beanstalk for a company. The application consists of a single Docker image. The company's automated continuous integration and continuous delivery (CI/CD) process builds the Docker image and pushes the image to a public Docker registry.

How should the developer deploy the application to Elastic Beanstalk?

- A. Create a Dockerfile
- B. Configure Elastic Beanstalk to build the application as a Docker image.
- C. Create a docker-compose.yml file
- D. Use the Elastic Beanstalk CLI to deploy the application.
- E. Create a .zip file that contains the Docker image
- F. Upload the .zip file to Elastic Beanstalk.
- G. Create a Dockerfile
- H. Run the Elastic Beanstalk CLI `eb local run` command in the same directory.

Answer: B

Explanation:

<https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/docker.html#single-container-docker.deploy-remote> Deploy a remote Docker image to Elastic Beanstalk After testing your container locally, deploy it to an Elastic Beanstalk environment. Elastic Beanstalk uses the docker-compose.yml file to pull and run your image if you are using Docker Compose. Otherwise, Elastic Beanstalk uses the Dockerrun.aws.json instead.

Use the EB CLI to create an environment and deploy your image.

```
~/remote-docker$ eb create environment-name
```

NEW QUESTION 110

- (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 115

- (Exam Topic 4)

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

- Average item size: 10 KB
 - Item reads each second: 10 strongly consistent
 - Item writes each second: 2 transactional
- Which read and write capacity cost-effectively meets these requirements?

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

Answer: B

NEW QUESTION 119

- (Exam Topic 3)

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

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

Answer: B

NEW QUESTION 120

- (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 124

- (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 126

- (Exam Topic 3)

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

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

Answer: C

NEW QUESTION 128

- (Exam Topic 3)

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

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

E. Use the Amazon CloudWatch API to write timestamps to a custom CloudWatch metric Use the AWS X-Ray console to analyze the resulting data

Answer: C

NEW QUESTION 131

- (Exam Topic 3)

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

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

Answer: A

NEW QUESTION 133

- (Exam Topic 3)

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

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

Answer: A

NEW QUESTION 138

- (Exam Topic 3)

A developer wants to secure sensitive configuration data such as passwords, database strings, and application license codes. Access to this sensitive information must be tracked for future audit purposes.

- A. In an encrypted file on the source code bundle; grant the application access with Amazon IAM
- B. In the Amazon EC2 Systems Manager Parameter Store; grant the application access with IAM
- C. On an Amazon EBS encrypted volume attach the volume to an Amazon EC2 instance to access the data
- D. As an object in an Amazon S3 bucket, grant an Amazon EC2 instance access with an IAM role.

Answer: B

NEW QUESTION 139

- (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 143

- (Exam Topic 3)

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

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

Answer: D

NEW QUESTION 148

- (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 149

- (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 150

- (Exam Topic 3)

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

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

Answer: A

NEW QUESTION 152

- (Exam Topic 3)

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

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

Answer: D

NEW QUESTION 153

- (Exam Topic 3)

A company has an AWS Lambda function that runs hourly, reads log files that are stored in Amazon S3, and forwards alerts to Amazon Simple Notification Service (Amazon SNS) topics based on content. A developer wants to add a custom metric to the Lambda function to track the number of alerts of each type for each run. The developer needs to log this information in Amazon CloudWatch in a metric that is named Lambda/AlertCounts. How should the developer modify the Lambda function to meet this requirement with the LEAST operational overhead?

- A. Add a print statement to standard out for each alert and the number of occurrences
- B. Add a call to the PutMetricData API operation. Pass an array for alerts and the number of occurrences in the Values and Counts fields with a namespace of "Lambda/AlertCounts"
- C. Add a call to the PutMetricAlarm API operation. Pass an array of alerts in the metrics member with the namespace of "Lambda/AlertCounts"
- D. Add a call to the PutDashboard API operation. Pass an array of alerts in the metrics member with the namespace of "Lambda/AlertCounts"

Answer: B

NEW QUESTION 157

- (Exam Topic 3)

A developer is creating an event handling system. To handle messages asynchronously, the developer created a standard Amazon SQS queue. Quality assurance testing reveals that some events were processed multiple times. What is the recommended way to ensure the events are not processed more than once?

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

Answer: C

NEW QUESTION 159

- (Exam Topic 3)

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

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

Answer: A

NEW QUESTION 164

- (Exam Topic 3)

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

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

What should the developer do to correct this error message?

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

Answer: D

NEW QUESTION 169

- (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 174

- (Exam Topic 3)

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

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

Answer: A

NEW QUESTION 176

- (Exam Topic 3)

A developer is building an application using an Amazon API Gateway REST API backed by an AWS Lambda function that interacts with an Amazon DynamoDB table. During testing, the developer observes high latency when making requests to the API. How can the developer evaluate the end-to-end latency and identify performance bottlenecks?

- A. Enable AWS CloudTrail logging and use the logs to map each latency and bottleneck.
- B. Enable and configure AWS X-Ray tracing on API Gateway and the Lambda function. Use X-Ray to trace and analyze user requests.
- C. Enable Amazon CloudWatch Logs for the Lambda function. Enable execution logs for API Gateway to view and analyze user request logs.
- D. Enable VPC Flow Logs to capture and analyze network traffic within the VPC.

Answer: B

NEW QUESTION 180

- (Exam Topic 3)

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

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

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

Answer: A

NEW QUESTION 184

- (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 189

- (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 191

- (Exam Topic 3)

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

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

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

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

Answer: A

NEW QUESTION 195

- (Exam Topic 3)

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

Which DynamoDB feature should be used to meet these requirements?

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

Answer: B

Explanation:

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

NEW QUESTION 199

- (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 201

- (Exam Topic 3)

A developer used the BatchWriteItem API operation to insert items in an Amazon DynamoDB table. The operation 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 206

- (Exam Topic 3)

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

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

Answer: DE

NEW QUESTION 207

- (Exam Topic 3)

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

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

Answer: D

NEW QUESTION 209

- (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 211

- (Exam Topic 3)

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

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

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

Answer: C

NEW QUESTION 212

- (Exam Topic 3)

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

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

- A. Rewrite the application to be cloud native and to run on AWS Lambda where the logs can be reviewed in Amazon CloudWatch.

- B. Set up centralized logging by using Amazon Elasticsearch Service (Amazon ES), Logstash, and Kibana
- C. Scale down the application to one larger EC2 instance where only one instance is recording logs.
- D. Install the unified Amazon CloudWatch agent on the EC2 instance
- E. Configure the agent to push the application logs to CloudWatch.

Answer: D

NEW QUESTION 215

- (Exam Topic 3)

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

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

Answer: C

NEW QUESTION 219

- (Exam Topic 3)

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

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

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

Answer: AD

NEW QUESTION 220

- (Exam Topic 3)

An application uploads photos to an Amazon S3 bucket. Each photo that is uploaded to the S3 bucket must be resized to a thumbnail image by the application. Each thumbnail image is uploaded with a new name in the same S3 bucket.

Which AWS service can a developer configure to directly process each single S3 event (or each S3 object upload)?

- A. Amazon EC2
- B. Amazon Elastic Container Service (Amazon ECS)
- C. AWS Elastic Beanstalk
- D. AWS Lambda

Answer: D

NEW QUESTION 221

- (Exam Topic 3)

A company has a web application in an Amazon Elastic Container Service (Amazon ECS) cluster running hundreds of secure services in AWS Fargate containers. The services are in target groups routed by an Application Load Balancer (ALB). Application users log in to the website anonymously, but they must be authenticated using any OpenID Connect protocol-compatible identity provider (IdP) to access the secure services. Which authentication approach would meet these requirements with the LEAST amount of effort?

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

Answer: A

NEW QUESTION 224

- (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 225

- (Exam Topic 3)

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

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

Answer: D

NEW QUESTION 230

- (Exam Topic 3)

A company is launching a new web application in the AWS Cloud. The company's development team is using AWS Elastic Beanstalk for deployment and maintenance. According to the company's change management process, the development team must evaluate changes for a specific time period before completing the rollout.

Which deployment policy meets this requirement?

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

Answer: A

NEW QUESTION 233

- (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 236

- (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 241

- (Exam Topic 3)

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

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

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

Answer: A

NEW QUESTION 246

- (Exam Topic 3)

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

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

Answer: A

NEW QUESTION 251

- (Exam Topic 3)

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

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

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

Answer: B

NEW QUESTION 254

- (Exam Topic 3)

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

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

Answer: C

NEW QUESTION 256

- (Exam Topic 3)

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

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

Answer: D

NEW QUESTION 260

- (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 263

- (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 266

- (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 267

- (Exam Topic 3)

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

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

Answer: AD

NEW QUESTION 268

- (Exam Topic 3)

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

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

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

Answer: A

NEW QUESTION 272

- (Exam Topic 3)

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

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

Answer: D

NEW QUESTION 277

- (Exam Topic 3)

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

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

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

Answer: A

NEW QUESTION 279

- (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 280

- (Exam Topic 3)

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

Which approach addresses these requirements?

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

Answer: D

NEW QUESTION 283

- (Exam Topic 3)

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

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

Answer: D

NEW QUESTION 285

- (Exam Topic 3)

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

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

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

Answer: D

NEW QUESTION 286

- (Exam Topic 3)

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

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

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

Answer: D

NEW QUESTION 288

- (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 292

- (Exam Topic 3)

A company's ecommerce website is experiencing massive traffic spikes, which are causing performance problems in the company database. Users are reporting

that accessing the website takes a long time

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

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

Answer: E

NEW QUESTION 295

- (Exam Topic 3)

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

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

Answer: BE

NEW QUESTION 300

- (Exam Topic 3)

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

What is the MOST likely cause of this error?

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

Answer: D

NEW QUESTION 303

- (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 305

- (Exam Topic 3)

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

What is the MOST efficient way to fulfill this request?

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

Answer: D

NEW QUESTION 306

- (Exam Topic 3)

A developer at a company writes an AWS CloudFormation template. The template refers to subnets that were created by a separate AWS CloudFormation template that the company's network team wrote. When the developer attempts to launch the stack for the first time, the launch fails.

Which template coding mistakes could have caused this failure? (Select TWO.)

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

Answer: BD

NEW QUESTION 307

- (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 311

- (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 312

- (Exam Topic 3)

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

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

Answer: D

NEW QUESTION 317

- (Exam Topic 3)

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

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

Answer: A

NEW QUESTION 320

- (Exam Topic 3)

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

```
aws dynamodb get-item table-name demoman-table --key '{"id": <"N"; "1993"}' ' The command returned errors and no rows were returned
```

What is the MOST likely cause of these issues?

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

Answer: D

NEW QUESTION 325

- (Exam Topic 3)

A developer is working on a serverless application that needs to process any changes to an Amazon DynamoDB table with an AWS Lambda function How should the developer configure the Lambda function to detect changes to the DynamoDB table?

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

Answer: C

NEW QUESTION 330

- (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 334

- (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 336

- (Exam Topic 3)

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

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

Answer: CE

NEW QUESTION 338

- (Exam Topic 3)

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

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

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

Answer: D

NEW QUESTION 341

- (Exam Topic 3)

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

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

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

Answer: C

NEW QUESTION 342

- (Exam Topic 3)

A developer has designed a customer, facing application that is running on an Amazon EC2 instance. The application logs every request made to it. The application usually runs seamlessly, but a spike in traffic generates several logs that cause the disk to fill up and eventually run out of memory Company policy requires old to be centralized for analysis.

Which long-term solution should the developer employ to prevent the issue from reoccurring?

- A. Install the Amazon CloudWatch agent on the instance to send the logs to CloudWatc
- B. Delete the logs from the instance once they are sent to CloudWatch.
- C. Enable AWS Auto Scaling on Amazon Elastic Block Store (Amazon EBS) to automatically add volumes to the instance when it reaches a specified threshold.
- D. Enable AWS Auto Scaling on Amazon Elastic Block Store (Amazon EBS) a outomatically add volumeto the instance when it reaches a specified threshold.
- E. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to pull the logs from the instance.Configure the rule to delete the logs they have been pulled.

Answer: D

NEW QUESTION 346

- (Exam Topic 3)

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

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

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

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

Answer: D

Explanation:

Meaning:

DOC-EXAMPLE-BUCKET ==> bucket

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

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

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

NEW QUESTION 348

- (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 349

- (Exam Topic 2)

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

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

Answer: C

NEW QUESTION 351

- (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 352

- (Exam Topic 2)

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

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

Answer: D

NEW QUESTION 355

- (Exam Topic 2)

A Developer must encrypt a 100-GB object using AWS KMS. What is the BEST approach?

- A. Make an Encrypt API call to encrypt the plaintext data as ciphertext using a customer master key (CMK)
- B. Make an Encrypt API call to encrypt the plaintext data as ciphertext using a customer master key (CMK) with imported key material
- C. Make a GenerateDataKey API call that returns a plaintext key and an encrypted copy of a data ke
- D. Use a plaintext key to encrypt the data
- E. Make a GenerateDataKeyWithoutPlaintext API call that returns an encrypted copy of a data ke
- F. Use an encrypted key to encrypt the data

Answer: D

Explanation:

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

NEW QUESTION 360

- (Exam Topic 2)

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

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

- A. Run the following:`docker pull REPOSITORY_URI : TAG`
- B. Run the output of the following:`aws ecr get-login`and then run:`docker pull REPOSITORY_URI : TAG`
- C. Run the following:`aws ecr get-login`and then run:`docker pull REPOSITORY_URI : TAG`
- D. Run the output of the following:`aws ecr get-download-url-for-layer`and then run:`docker pull REPOSITORY_URI : TAG`

Answer: B

Explanation:

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

NEW QUESTION 361

- (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 362

- (Exam Topic 2)

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

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

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

Answer: BC

Explanation:

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

NEW QUESTION 366

- (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 367

- (Exam Topic 2)

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

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

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

Answer: AD

NEW QUESTION 369

- (Exam Topic 2)

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

How can DynamoDB costs be minimized while maximizing application performance?

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

Answer: C

Explanation:

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

NEW QUESTION 370

- (Exam Topic 2)

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 371

- (Exam Topic 2)

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

How should the developer accomplish this?

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

Answer: D

NEW QUESTION 373

- (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 374

- (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 378

- (Exam Topic 2)

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

How can this requirement be met?

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

Answer: A

NEW QUESTION 382

- (Exam Topic 2)

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

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

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

Answer: C

NEW QUESTION 385

- (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 388

- (Exam Topic 2)

A developer is working on an AWS Lambda function that accesses Amazon DynamoDB The Lambda function must retrieve an item and update some of its attributes. or create the item if it does not exist The Lambda function has access to the primary key.

Which IAM permissions should the developer request for the Lambda function to achieve this functionality?

- A. dynamodb:DeleteItem dynamodb:GetItem dynamodb:PutItem
- B. dynamodb:UpdateItem dynamodb:GetItem dynamodb:DescribeTable
- C. dynamodb:GetRecords dynamodb:PutItem dynamodb:updateTable
- D. dynamodb:UpdateItem dynamodb:GetItem dynamodb:PutItem

Answer: C

Explanation:

Reference: <https://docs.aws.amazon.com/AWSJavaScriptSDK/latest/AWS/DynamoDB.html>

NEW QUESTION 393

- (Exam Topic 2)

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

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

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

Answer: B

Explanation:

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

NEW QUESTION 396

.....

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](#)