

# Exam Questions DVA-C01

AWS Certified Developer Associate Exam

<https://www.2passeasy.com/dumps/DVA-C01/>



#### NEW QUESTION 1

- (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 2

- (Exam Topic 4)

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

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

**Answer:** C

#### Explanation:

Client-side and server-side encryption:

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

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

#### NEW QUESTION 3

- (Exam Topic 4)

A 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 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 5

- (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 6

- (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 7

- (Exam Topic 4)

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

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

Which solution will meet these requirements MOST cost-effectively?

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

**Answer:** D

#### Explanation:

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

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

#### NEW QUESTION 8

- (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 locall
- 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 pipelin
- L. Use AWS CodeBuild as the provide
- M. Add the new stage before the stage that deploys code revisions to the test environmen
- N. Write a buildspec that fails the CodeBuild stage if any test does not pas
- O. Use the test reports feature of CodeBuild to integrate the report with the CodeBuild consol
- P. View the test results in CodeBuil
- Q. Resolve any issues.
- R. Add a new stage to the pipelin
- S. Use Jenkins as the provide
- T. Configure CodePipeline to use Jenkins to run the unit test
- . Write a Jenkinsfile that fails the stage if any test does not pas
- . Use the test report plugin for Jenkins to integrate the report with the Jenkins dashboar
- . View the test results in Jenkin
- . Resolve any issues.

**Answer:** C

#### Explanation:

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

#### NEW QUESTION 9

- (Exam Topic 4)

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

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

production environment.

Which solution will moot these requirements?

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

**Answer:** A

**Explanation:**

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

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

**NEW QUESTION 10**

- (Exam Topic 4)

Which solution will meet these requirements?

- A. Build the container from the amazon/aws-xray-daemon base image
- B. Use the AWS X-Ray SDK to instrument the application.
- C. Install the Amazon CloudWatch agent on the container image
- D. Use the CloudWatch SDK to publish custom metrics from each of the microservices.
- E. Install the AWS X-Ray daemon on each of the ECS instances.
- F. Configure AWS CloudTrail data events to capture the traffic between the microservices.

**Answer:** C

**NEW QUESTION 10**

- (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 13**

- (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 14**

- (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 16

- (Exam Topic 4)

A developer uses server-side encryption with Amazon S3 managed encryption keys (SSE-S3) to store data in Amazon S3. The developer needs to decrypt and download the encrypted objects by using the GetObject API call.

What is the LEAST amount of information that the developer must provide in the API call to meet this requirement?

- A. The S3 object key only
- B. The S3 object key and the encryption key
- C. The S3 object key and the Amazon Resource Name (ARN) of the AWS Key Management Service (AWS KMS) key
- D. The S3 object key and a randomly salted Hash-based Message Authentication Code (HMAC) value of the encryption key

**Answer:** A

#### Explanation:

[https://docs.aws.amazon.com/AmazonS3/latest/API/API\\_GetObject.html](https://docs.aws.amazon.com/AmazonS3/latest/API/API_GetObject.html)

Encryption request headers, like x-amz-server-side-encryption, should not be sent for GET requests if your object uses server-side encryption with KMS keys (SSE-KMS) or server-side encryption with Amazon S3–managed encryption keys (SSE-S3). If your object does use these types of keys, you'll get an HTTP 400 BadRequest error.

[https://docs.aws.amazon.com/AmazonS3/latest/API/API\\_GetObject.html#API\\_GetObject\\_Examples](https://docs.aws.amazon.com/AmazonS3/latest/API/API_GetObject.html#API_GetObject_Examples)

<https://awscli.amazonaws.com/v2/documentation/api/latest/reference/s3api/get-object.html>

Sample Request:

The following request returns the object my-image.jpg. GET /my-image.jpg HTTP/1.1

Host: bucket.s3.<Region>.amazonaws.com Date: Mon, 3 Oct 2016 22:32:00 GMT

Authorization: authorization string

The following example uses the get-object command to download an object from Amazon S3: aws s3api get-object --bucket text-content --key dir/my\_images.tar.bz2 my\_images.tar.bz2

#### NEW QUESTION 20

- (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 22

- (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](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 23

- (Exam Topic 4)

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

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

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

**Answer: B**

#### NEW QUESTION 26

- (Exam Topic 4)

A developer is testing a new file storage application that uses an Amazon CloudFront distribution to serve content from an Amazon S3 bucket. The distribution accesses the S3 bucket by using an origin access identity (OAI). The S3 bucket's permissions explicitly deny access to all other users.

The application prompts users to authenticate on a login page and then uses signed cookies to allow users to access their personal storage directories. The developer has configured the distribution to use its default cache behavior with restricted viewer access and has set the origin to point to the S3 bucket. However, when the developer tries to navigate to the login page, the developer receives a 403 Forbidden error.

The developer needs to implement a solution to allow unauthenticated access to the login page. The solution also must keep all private content secure.

Which solution will meet these requirements?

- A. Add a second cache behavior to the distribution with the same origin as the default cache behavior
- B. Set the path pattern for the second cache behavior to the path of the login page, and make viewer access unrestricted
- C. Keep the default cache behavior's settings unchanged.
- D. Add a second cache behavior to the distribution with the same origin as the default cache behavior
- E. Set the path pattern for the second cache behavior to \*, and make viewer access restricted
- F. Change the default cache behavior's path pattern to the path of the login page, and make viewer access unrestricted.
- G. Add a second origin as a failover origin to the default cache behavior
- H. Point the failover origin to the S3 bucket
- I. Set the path pattern for the primary origin to \* and make viewer access restricted
- J. Set the path pattern for the failover origin to the path of the login page, and make viewer access unrestricted.
- K. Add a bucket policy to the S3 bucket to allow read access
- L. Set the resource on the policy to the Amazon Resource Name (ARN) of the login page object in the S3 bucket
- M. Add a CloudFront function to the default cache behavior to redirect unauthorized requests to the login page's S3 URI.

**Answer: B**

#### Explanation:

Adding a second cache behavior to the distribution with the same origin as the default cache behavior and setting the path pattern to \* will allow access to all files in the S3 bucket. Changing the default cache behavior's path pattern to the path of the login page and making viewer access unrestricted will allow unauthenticated users to access the login page, while keeping all other private content secure.

#### NEW QUESTION 27

- (Exam Topic 4)

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

Which solution will meet these requirements?

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

**Answer: D**

#### Explanation:

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

#### NEW QUESTION 29

- (Exam Topic 4)

A company uses the AWS SDK for JavaScript in the Browser to build a web application and then hosts the application on Amazon S3. The company wants the application to support 10,000 users concurrently. The company selects Amazon DynamoDB to store user preferences in a table. There is a requirement to uniquely identify users at any scale.

Which solution will meet these requirements?

- A. Create a user cookie
- B. Attach an IAM role to the S3 bucket that hosts the application.
- C. Deploy an Amazon CloudFront distribution with an origin access identity (OAI) to access the S3 bucket.
- D. Configure and use Amazon Cognito
- E. Access DynamoDB with the authenticated users.
- F. Create an IAM user for each user
- G. Use fine-grained access control on the DynamoDB table to control access.

**Answer: C**

#### Explanation:

This will allow the application to support 10,000 users concurrently and will provide a unique identifier for each user. By using Amazon Cognito, the company can authenticate users and then access DynamoDB with the authenticated users to store their preferences in a table. This approach will allow the company to control access to the DynamoDB table and to scale to any number of users. Creating a user cookie or deploying an Amazon CloudFront distribution with an OAI would not

solve the problem because these solutions do not provide a way to uniquely identify users or control access to DynamoDB. Creating an IAM user for each user and using fine-grained access control on the DynamoDB table would not be practical or scalable because it would require the company to manage and maintain a large number of IAM users.

When dealing with user profiles in serverless applications we often turn to Cognito for managing their credentials while the app itself will store user entities.  
<https://www.sorenandersen.com/manage-user-profile-data-between-cognito-and-dynamodb/>

### NEW QUESTION 30

- (Exam Topic 4)

A company is migrating a web application from on premises to AWS. The company needs to move session storage from the application code to a shared service as part of the migration. The session storage data must be encrypted at rest.

Which AWS services meet these requirements? (Choose two.)

- A. Amazon ElastiCache for Redis
- B. Amazon ElastiCache for Memcached
- C. Amazon CloudWatch
- D. AWS CloudTrail
- E. Amazon DynamoDB

**Answer:** AE

#### Explanation:

<https://aws.amazon.com/blogs/security/amazon-elasticache-now-supports-encryption-for-elasticache-for-redis/>

A) ElastiCache for Redis is always a good option as a distributed cache for session management - <https://aws.amazon.com/getting-started/hands-on/building-fast-session-caching-with-amazon-elasticache-for-red>

It also supports encrypt at rest - <https://docs.aws.amazon.com/AmazonElastiCache/latest/red-ug/at-rest-encryption.html>

E) DynamoDB is also common to store session state with TTL support. And all user data stored in Amazon DynamoDB is fully encrypted at rest - <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/EncryptionAtRest.html>

### NEW QUESTION 34

- (Exam Topic 4)

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

What should the developer do to meet this requirement?

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

**Answer:** C

### NEW QUESTION 39

- (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 bundl
- B. Grant the application access by using IAM policies.
- C. Store the configuration data in AWS Systems Manager Parameter Stor
- 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 Manage
- 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 43

- (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. Miich 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 44

- (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 46

- (Exam Topic 3)

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

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

**Answer:** B

#### NEW QUESTION 48

- (Exam Topic 3)

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

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

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

**Answer:** A

#### NEW QUESTION 49

- (Exam Topic 3)

A developer is building a web and mobile application for two types of users regular users and guest users Regular users are required to log in, but guest users do not log in Users should see only their data regardless of whether they authenticate Users need AWS credentials before they can access AWS resources What is the MOST secure solution that the developer can implement to allow access for guest users?

- A. Use an Amazon Cognito credentials provider to issue temporary credentials that are linked to an unauthenticated role that has access to the required resources.
- B. Set up an IAM user that has permissions to the required resource
- C. Hardcode the IAM credentials in the web and mobile application
- D. Generate temporary keys that are stored in AWS Key Management Service (AWS KMS) Use the temporary keys to access the required resources
- E. Generate temporary credential
- F. Store the temporary credentials in AWS Secrets Manager Use the temporary credentials to access the required resources

**Answer:** D

#### NEW QUESTION 51

- (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 53

- (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 56

- (Exam Topic 3)

A developer is creating a script to automate the deployment process for a serverless application. The developer wants to use an existing AWS Serverless Application Model (AWS SAM) template for the application. What should the developer use for the project? (Select TWO)

- A. Call `aws cloudformation package` to create the deployment package. Call `aws cloudformation deploy` to deploy the package afterward.
- B. Call `sam package` to create the deployment package. Call `sam deploy` to deploy the package afterward.
- C. Call `aws s3 cp` to upload the AWS SAM template to Amazon S3. Call `aws lambda update-function-code` to create the application.
- D. Create a ZIP package locally and call `aws serverlessrepo create-application` to create the application.
- E. Create a ZIP package and upload it to Amazon S3. Call `aws cloudformation create-stack` to create the application.

**Answer: BC**

#### NEW QUESTION 59

- (Exam Topic 3)

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

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

**Answer: C**

#### NEW QUESTION 63

- (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 that includes the company logo. What should a developer do to meet these requirements?

- A. Create a hosted user interface (UI) in Amazon Cognito. Customize the hosted UI with the company logo.
- B. Create a login page that includes the company logo.
- C. Upload the login page to Amazon Cognito.
- D. Create a login page that includes the company logo in Amazon API Gateway. Save the link in Amazon Cognito.
- E. Upload the company logo to an Amazon S3 bucket. Specify the S3 object path in app client settings in Amazon Cognito.

**Answer: B**

#### NEW QUESTION 65

- (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 67

- (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 69

- (Exam Topic 3)

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

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

**Answer:** D

#### NEW QUESTION 71

- (Exam Topic 3)

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

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

**Answer:** D

#### NEW QUESTION 75

- (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 79

- (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 83

- (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 88

- (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 90

- (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 91

- (Exam Topic 3)

A development team uses AWS Elastic Beanstalk to deploy a Java-based web application. The team wants to ensure that the changes to the source code and the configuration are always deployed on new instances. The team configures the Elastic Beanstalk environment to use immutable updates. However an error occurs the first time a change is deployed with the new update policy. What is the MOST likely cause of this issue?

- A. Immutable updates are not supported for Java-based applications
- B. The account has reached its on-demand instance limit
- C. Immutable updates are only supported for m4 large and larger instance types.
- D. The developer must also modify the ebextensions/immutable-updates config file to enable immutable updates

**Answer:** A

#### NEW QUESTION 92

- (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 96

- (Exam Topic 3)

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

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

**Answer:** D

#### NEW QUESTION 98

- (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 102

- (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 104

- (Exam Topic 3)

A developer needs to deploy a new version to an AWS Elastic Beanstalk application. How can the developer accomplish this task?

- A. Upload and deploy the new application version in the Elastic Beanstalk console
- B. Use the `eb init` CLI command to deploy a new version
- C. Terminate the current Elastic Beanstalk environment and create a new one
- D. Modify the `ebextensions` folder to add a source option to services

**Answer:** A

#### NEW QUESTION 105

- (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 107

- (Exam Topic 3)

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

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

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

**Answer:** B

#### NEW QUESTION 111

- (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 112

- (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 117

- (Exam Topic 3)

A developer is troubleshooting a three-tier application, which is deployed on Amazon EC2 instances. There is a connectivity problem between the application servers and the database servers.

Which AWS services or tools should be used to identify the faulty component? (Select TWO.)

- A. AWS CloudTrail.
- B. AWS Trusted Advisor
- C. Amazon VPC Flow Logs
- D. Network access control lists
- E. AWS Config rules

**Answer:** CD

#### NEW QUESTION 122

- (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 123

- (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 126

- (Exam Topic 3)

A developer wants to use React to build a web and mobile application. The application will be hosted on AWS. The application must authenticate users and then allow users to store and retrieve files that they own. The developer wants to use Facebook for authentication.

Which CLI will MOST accelerate the development and deployment of this application on AWS?

- A. AWS CLI
- B. AWS Amplify CLI
- C. AWS Serverless Application Model (AWS SAM) CLI
- D. Amazon Elastic Container Service (Amazon ECS) CLI

**Answer:** B

#### NEW QUESTION 128

- (Exam Topic 3)

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

Which solution will meet these requirements?

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

**Answer:** A

#### NEW QUESTION 133

- (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 135**

- (Exam Topic 3)

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

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

**Answer:** D

**NEW QUESTION 136**

- (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-s3/>

**NEW QUESTION 137**

- (Exam Topic 3)

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

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

**Answer:** B

**NEW QUESTION 142**

- (Exam Topic 3)

A developer is working on a serverless project based in Java. Initial testing shows a cold start takes about 8 seconds on average for AWS Lambda functions. What should the developer do to reduce the cold start time? (Select TWO)

- A. Add the Spring Framework to the project and enable dependency injection.
- B. Reduce the deployment package by including only the needed modules from the AWS SDK for Java.
- C. Increase the memory allocation setting for the Lambda function.
- D. Increase the timeout setting for the Lambda function.
- E. Change the Lambda invocation mode from synchronous to asynchronous.

**Answer:** BC

**NEW QUESTION 144**

- (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 147

- (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 152

- (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 153

- (Exam Topic 2)

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

What should a developer do to meet these requirements?

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

**Answer:** A

#### NEW QUESTION 154

- (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 156

- (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 158

- (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 159

- (Exam Topic 2)

A Developer is migrating existing applications to AWS. These applications use MongoDB as their primary data store, and they will be deployed to Amazon EC2 instances. Management requires that the Developer minimize changes to applications while using AWS services  
Which solution should the Developer use to host MongoDB in AWS?

- A. Install MongoDB on the same instance where the application is running
- B. Deploy Amazon DocumentDB in MongoDB compatibility mode
- C. Use Amazon API Gateway to translate API calls from MongoDB to Amazon DynamoDB.
- D. Replicate the existing MongoDB workload to Amazon DynamoDB

**Answer:** D

#### NEW QUESTION 160

- (Exam Topic 2)

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

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

**Answer:** C

**Explanation:**

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

#### NEW QUESTION 164

- (Exam Topic 2)

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

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

**Answer:** B

#### NEW QUESTION 167

- (Exam Topic 2)

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

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

**Answer:** C

**Explanation:**

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

#### NEW QUESTION 169

- (Exam Topic 2)

A Developer must trigger an AWS Lambda function based on the item lifecycle activity in an Amazon DynamoDB table.  
How can the Developer create the solution?

- A. Enable a DynamoDB stream that publishes an Amazon SNS message
- B. Trigger the Lambda function synchronously from the SNS message.
- C. Enable a DynamoDB stream that publishes an SNS message



- D. Trigger the Lambda function asynchronously from the SNS message.
- E. Enable a DynamoDB stream, and trigger the Lambda function synchronously from the stream.
- F. Enable a DynamoDB stream, and trigger the Lambda function asynchronously from the stream.

**Answer:** C

**Explanation:**

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

#### NEW QUESTION 174

- (Exam Topic 2)

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

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

**Answer:** B

#### NEW QUESTION 177

- (Exam Topic 2)

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

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

**Answer:** D

#### NEW QUESTION 181

- (Exam Topic 2)

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

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

**Answer:** B

#### NEW QUESTION 185

- (Exam Topic 2)

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

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

**Answer:** A

#### NEW QUESTION 187

- (Exam Topic 2)

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

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

**Answer:** D

#### NEW QUESTION 190

- (Exam Topic 2)

An organization is using Amazon CloudFront to ensure that its users experience low-latency access to its web application. The organization has identified a need to encrypt all traffic between users and CloudFront, and all traffic between CloudFront and the web application. How can these requirements be met? (Choose two.)

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

**Answer:** AB

#### Explanation:

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

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

#### NEW QUESTION 194

- (Exam Topic 2)

A company is developing an application that will be accessed through the Amazon API Gateway REST API. Registered users should be the only ones who can access certain resources of this API. The token being used should expire automatically and needs to be refreshed periodically. How can a developer meet these requirements'?

- A. Create an Amazon Cognito identity pool, configure the Amazon Cognito Authorizer in API Gateway, and use the temporary credentials generated by the identity pool
- B. Create and maintain a database record for each user with a corresponding token and use an AWS Lambda authorizer in API Gateway
- C. Create an Amazon Cognito user pool, configure the Cognito Authorizer in API Gateway, and use the identity or access token
- D. Create an IAM user for each API user, attach an invoke permissions policy to the AP
- E. and use an IAM authorizer in API Gateway.

**Answer:** C

#### Explanation:

Reference: <https://aws.amazon.com/premiumsupport/knowledge-center/cognito-custom-scopes-api-gateway/>

#### NEW QUESTION 197

- (Exam Topic 2)

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

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

**Answer:** B

#### Explanation:

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

#### NEW QUESTION 199

- (Exam Topic 2)

An advertising company has a dynamic website with heavy traffic. The company wants to migrate the website infrastructure to AWS to handle everything except website development. Which solution BEST meets these requirements?

- A. Use AWS VM Import to migrate a web server image to AWS. Launch the image on a compute-optimized Amazon EC2 instance.
- B. Launch multiple Amazon Lightsail instances behind a load balance.
- C. Set up the website on those instances.
- D. Deploy the website code in an AWS Elastic Beanstalk environment.
- E. Use Auto Scaling to scale the number of instances.
- F. Use Amazon S3 to host the website.
- G. Use Amazon CloudFront to deliver the content at scale.

**Answer:** C

#### NEW QUESTION 201

- (Exam Topic 2)

A company has 25,000 employees and is growing. The company is creating an application that will be accessible to its employees only. A developer is using Amazon S3 to store images and Amazon RDS to store application data. The company requires that all employee information remain in the legacy Security Assertion Markup Language (SAML) employee directory only and is not interested in mirroring any employee information on AWS. How can the developer provide authorized access for the employees who will be using this application so each employee can access their own application data only?

- A. Use Amazon VPC and keep all resources inside the VP
- B. and use a VPC link for the S3 bucket with the bucket policy.
- C. Use Amazon Cognito user pools, federate with the SAML provider and use user pool groups with an IAM policy
- D. Use an Amazon Cognito identity pool, federate with the SAML provider, and use an IAM condition key with a value for the cognito-identity.amazonaws.com sub variable to grant access to the employees.
- E. Create a unique IAM role for each employee and have each employee assume the role to access the application so they can access their personal data only.

**Answer:** B

#### NEW QUESTION 204

- (Exam Topic 2)

A developer has written an Amazon Kinesis Data Streams application. As usage grows and traffic increases over time, the application is regularly receiving ProvisionedThroughputExceededException error messages  
Which steps should the developer take to resolve the error? (Select TWO.)

- A. Use Auto Scaling to scale the stream for better performance
- B. Increase the delay between the GetRecords call and the PutRecords call.
- C. Increase the number of shards in the data stream
- D. Specify a shard iterator using the ShardIterator parameter.
- E. Implement exponential backoff on the GetRecords call and the PutRecords call.

**Answer:** BC

#### Explanation:

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

#### NEW QUESTION 205

- (Exam Topic 2)

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

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

**Answer:** B

#### Explanation:

Reference:

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

#### NEW QUESTION 208

- (Exam Topic 2)

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

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

**Answer:** C

#### Explanation:

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

#### NEW QUESTION 212

- (Exam Topic 2)

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

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

**Answer:** A

#### Explanation:

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

#### NEW QUESTION 216

- (Exam Topic 2)

An on-premises application is implemented using a Linux, Apache, MySQL and PHP (LAMP) stack. The Developer wants to run this application in AWS. Which of the following sets of AWS services can be used to run this stack?

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

**Answer:** C

#### NEW QUESTION 219

- (Exam Topic 2)

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

What should the Developer implement to achieve the design goals?

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

**Answer:** A

#### NEW QUESTION 222

- (Exam Topic 2)

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

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

**Answer:** D

#### Explanation:

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

#### NEW QUESTION 227

- (Exam Topic 2)

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

What is the best option to store the session state?

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

**Answer:** A

#### Explanation:

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

#### NEW QUESTION 230

- (Exam Topic 2)

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

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

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

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

**Answer:** AD

#### NEW QUESTION 231

- (Exam Topic 2)

A developer is preparing a deployment package using AWS Cloud Formation. The package consists of two separate templates: one for the infrastructure and one for the application. The application has to be inside the VPC that is created from the infrastructure template

How can the application stack refer to the VPC created from the infrastructure template?



- A. Use the Ret function to import the VPC into the application stack from the infrastructure template
- B. Use the export flag in the infrastructure template, and then use the Fn::ImportValue function in the application template
- C. Use the DependsOn attribute to specify that the application instance depends on the VPC in the application template
- D. Use the Fn::GetAtt function to include the attribute of the VPC in the application template.

**Answer:** A

#### NEW QUESTION 234

- (Exam Topic 2)

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

How can the developer accomplish this?

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

**Answer:** B

#### Explanation:

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

#### NEW QUESTION 235

- (Exam Topic 1)

A Developer is designing a new application that uses Amazon S3. To satisfy compliance requirements, the Developer must encrypt the data at rest.

How can the Developer accomplish this?

- A. Use s3:x-amz-acl as a condition in the S3 bucket policy.
- B. Use Amazon RDS with default encryption.
- C. Use aws:SecureTransport as a condition in the S3 bucket policy.
- D. Turn on S3 default encryption for the S3 bucket.

**Answer:** D

#### NEW QUESTION 238

- (Exam Topic 1)

When a Developer tries to run an AWS CodeBuild project, it raises an error because the length of all environment variables exceeds the limit for the combined maximum of characters.

What is the recommended solution?

- A. Add the export LC\_ALL="en\_US.utf8" command to the pre\_build section to ensure POSIX localization.
- B. Use Amazon Cognito to store key-value pairs for large numbers of environment variables.
- C. Update the settings for the build project to use an Amazon S3 bucket for large numbers of environment variables.
- D. Use AWS Systems Manager Parameter Store to store large numbers of environment variables.

**Answer:** D

#### NEW QUESTION 243

- (Exam Topic 1)

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

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

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

**Answer:** AE

#### Explanation:

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

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

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

#### NEW QUESTION 245

- (Exam Topic 1)

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

```
def lambda_handler(event, context):  
  
    db = database.connect()  
  
    db.statement('INSERT INTO Customers (CustomerName) VALUES  
    (context.name)')  
  
    db.close()
```

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

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

**Answer:** C

**Explanation:**

Refer AWS documentation - Lambda Best Practices

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

**NEW QUESTION 246**

- (Exam Topic 1)

A serverless application uses an API Gateway and AWS Lambda.

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

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

**Answer:** A

**NEW QUESTION 248**

- (Exam Topic 1)

A Developer is creating a web application that requires authentication, but also needs to support guest access to provide users limited access without having to authenticate. What service can provide support for the application to allow guest access?

- A. IAM temporary credentials using AWS STS.
- B. Amazon Directory Service
- C. Amazon Cognito with unauthenticated access enabled
- D. IAM with SAML integration

**Answer:** C

**Explanation:**

<https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/serverless-getting-started-hello> <https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/sam-cli-command-reference-sa>

<https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/sam-cli-command-reference-sa>

**NEW QUESTION 251**

- (Exam Topic 1)

An application overwrites an object in Amazon S3, and then immediately reads the same object. Why would the application sometimes retrieve the old version of the object?

- A. S3 overwrite PUTS are eventually consistent, so the application may read the old object.
- B. The application needs to add extra metadata to label the latest version when uploading to Amazon S3.
- C. All S3 PUTS are eventually consistent, so the application may read the old object.
- D. The application needs to explicitly specify latest version when retrieving the object.

**Answer:** A

**NEW QUESTION 253**

- (Exam Topic 1)

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

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

- A. Use Amazon Cognito user pools and a custom authorizer to authenticate and authorize users based on JSON Web Tokens.
- B. Build a OpenID token broker with Amazon and Facebook

- C. Users will authenticate with these identifyproviders and pass the JSON Web Token to the API to authenticate each API call.
- D. Store user credentials in Amazon DynamoDB and have the application retrieve temporary credentials from AWS ST
- E. Make API calls by passing user credentials to the APIs for authentication and authorization.
- F. Use Amazon RDS to store user credentials and pass them to the APIs for authentications and authorization.

**Answer:** A

#### NEW QUESTION 255

- (Exam Topic 1)

An application stores payroll information nightly in DynamoDB for a large number of employees across hundreds of offices. Item attributes consist of individual name, office identifier, and cumulative daily hours. Managers run reports for ranges of names working in their office. One query is. "Return all Items in this office for names starting with A through E".

Which table configuration will result in the lowest impact on provisioned throughput for this query?

- A. Configure the table to have a hash index on the name attribute, and a range index on the office identifier
- B. Configure the table to have a range index on the name attribute, and a hash index on the office identifier
- C. Configure a hash index on the name attribute and no range index
- D. Configure a hash index on the office Identifier attribute and no range index

**Answer:** B

#### Explanation:

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

Partition key and sort key – Referred to as a composite primary key, this type of key is composed of two attributes. The first attribute is the partition key, and the second attribute is the sort key. DynamoDB uses the partition key value as input to an internal hash function. The output from the hash function determines the partition (physical storage internal to DynamoDB) in which the item will be stored. All items with the same partition key value are stored together, in sorted order by sort key value.

#### NEW QUESTION 260

- (Exam Topic 1)

A company is building a stock trading application that requires sub-millisecond latency in processing trading requests. Amazon DynamoDB is used to store all the trading data that is used to process each request. After load testing the application, the development team found that due to data retrieval times, the latency requirement is not satisfied. Because of sudden high spikes in the number of requests, DynamoDB read capacity has to be significantly over-provisioned to avoid throttling.

What steps should be taken to meet latency requirements and reduce the cost of running the application?

- A. Add Global Secondary Indexes for trading data.
- B. Store trading data in Amazon S3 and use Transfer Acceleration.
- C. Add retries with exponential back-off for DynamoDB queries
- D. Use DynamoDB Accelerator to cache trading data.

**Answer:** D

#### Explanation:

Refer AWS documentation - DynamoDB Accelerator

Amazon DynamoDB Accelerator (DAX) is a fully managed, highly available, in-memory cache foDr ynamoDB that delivers up to a 10x performance improevment – from milliseconds to microseconds – even at millions of requests per second. DAX does all the heavy lifting required to add in-memory acceleration to your DynamoDB tables, without requiring developers to manage cache invalidation, data population, or cluster management. Now you can focus on building great applications for your customers without worrying about performance at scale.

#### NEW QUESTION 261

- (Exam Topic 1)

A nightly batch job loads 1 million new records into a DynamoDB table. The records are only needed for one hour, and the table needs to be empty by the next night's batch job.

Which is the MOST efficient and cost-effective method to provide an empty table?

- A. Use DeleteItem using a ConditionExpression.
- B. Use BatchWriteItem to empty all of the rows.
- C. With a recursive function that scans and calls out DeleteItem.
- D. Create and then delete the table after the task has completed.

**Answer:** D

#### Explanation:

"Deleting an entire table is significantly more efficient than removing items one-by-one, which essentially doubles the write throughput as you do as many delete operations as put operations"

#### NEW QUESTION 264

- (Exam Topic 1)

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

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

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

**Answer:** C

#### NEW QUESTION 265

- (Exam Topic 1)

A company is creating an application that will require users to access AWS services and allow them to reset their own passwords.

Which of the following would allow the company to manage users and authorization while allowing users to reset their own passwords?

- A. Amazon Cognito identify pools and AWS STS
- B. Amazon Cognito identity pools and AWS IAM
- C. Amazon Cognito user pools and AWS KMS
- D. Amazon Cognito user pools and identity pools

**Answer:** D

#### Explanation:

<https://serverless-stack.com/chapters/cognito-user-pool-vs-identity-pool.html>

#### NEW QUESTION 267

- (Exam Topic 1)

Which of the following programming languages have an officially supported AWS SDK? Choose 2 answers

- A. Perl
- B. PHP
- C. Pascal
- D. Java
- E. SQL

**Answer:** BD

#### NEW QUESTION 272

- (Exam Topic 1)

A Developer has an application that can upload tens of thousands of objects per second to Amazon S3 in parallel within a single AWS account. As part of new requirements, data stored in S3 must use server side encryption with AWS KMS (SSE-KMS). After creating this change, performance of the application is slower. Which of the following is MOST likely the cause of the application latency?

- A. Amazon S3 throttles the rate at which uploaded objects can be encrypted using Customer Master Keys.
- B. The AWS KMS API calls limit is less than needed to achieve the desired performance.
- C. The client encryption of the objects is using a poor algorithm.
- D. KMS requires that an alias be used to create an independent display name that can be mapped to a CMK.

**Answer:** B

#### Explanation:

<https://aws.amazon.com/about-aws/whats-new/2018/08/aws-key-management-service-increases-api-requests-pe> KMS API access limit is 10k/sec in us-east and some others and 5.5k/sec for the rest of the regions. Client can request this limit to be changed.

#### NEW QUESTION 277

- (Exam Topic 1)

A Developer is building a mobile application and needs any update to user profile data to be pushed to all devices accessing the specific identity. The Developer does not want to manage a back end to maintain the user profile data.

What is the MOST efficient way for the Developer to achieve these requirements using Amazon Cognito?

- A. Use Cognito federated identities.
- B. Use a Cognito user pool.
- C. Use Cognito Sync.
- D. Use Cognito events.

**Answer:** C

#### Explanation:

Amazon Cognito Sync is an AWS service and client library that enables cross-device syncing of application-related user data. You can use it to synchronize user profile data across mobile devices and the web without requiring your own backend. <https://docs.aws.amazon.com/cognito/latest/developerguide/cognito-sync.html>

#### NEW QUESTION 282

- (Exam Topic 1)

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

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

**Answer:** CE

#### NEW QUESTION 283

- (Exam Topic 1)



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

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

**Answer:** A

**Explanation:**

[https://docs.aws.amazon.com/AmazonECS/latest/developerguide/task\\_IAM\\_role.html#:~:targetText=Amazon%](https://docs.aws.amazon.com/AmazonECS/latest/developerguide/task_IAM_role.html#:~:targetText=Amazon%20task,IAM,role,task,IAM,role)

#### NEW QUESTION 288

- (Exam Topic 1)

A Developer has been asked to build a real-time dashboard web application to visualize the key prefixes and storage size of objects in Amazon S3 buckets. Amazon DynamoDB will be used to store the Amazon S3 metadata.

What is the optimal and MOST cost-effective design to ensure that the real-time dashboard is kept up to date with the state of the objects in the Amazon S3 buckets?

- A. Use an Amazon CloudWatch event backed by an AWS Lambda function
- B. Issue an Amazon S3 API call to get a list of all Amazon S3 objects and persist the metadata within DynamoDB
- C. Have the web application poll the DynamoDB table to reflect this change.
- D. Use Amazon S3 Event Notification backed by a Lambda function to persist the metadata into DynamoDB
- E. Have the web application poll the DynamoDB table to reflect this change.
- F. Run a cron job within an Amazon EC2 instance to list all objects within Amazon S3 and persist the metadata into DynamoDB
- G. Have the web application poll the DynamoDB table to reflect this change.
- H. Create a new Amazon EMR cluster to get all the metadata about Amazon S3 objects; persist the metadata into DynamoDB
- I. Have the web application poll the DynamoDB table to reflect this change.

**Answer:** A

#### NEW QUESTION 293

- (Exam Topic 1)

A Developer is writing an imaging micro service on AWS Lambda. The service is dependent on several libraries that are not available in the Lambda runtime environment.

Which strategy should the Developer follow to create the Lambda deployment package?

- A. Create a ZIP file with the source code and all dependent libraries.
- B. Create a ZIP file with the source code and a script that installs the dependent libraries at runtime.
- C. Create a ZIP file with the source code
- D. Stage the dependent libraries on an Amazon S3 bucket indicated by the Lambda environment variable LD\_LIBRARY\_PATH
- E. Create a ZIP file with the source code and a buildspec.yaml file that installs the dependent libraries on AWS Lambda.

**Answer:** B

#### NEW QUESTION 295

- (Exam Topic 1)

Games-R-Us is launching a new game app for mobile devices. Users will log into the game using their existing Facebook account and the game will record player data and scoring information directly to a DynamoDB table.

What is the most secure approach for signing requests to the DynamoDB API?

- A. Create an IAM user with access credentials that are distributed with the mobile app to sign the requests
- B. Distribute the AWS root account access credentials with the mobile app to sign the requests
- C. Request temporary security credentials using web identity federation to sign the requests
- D. Establish cross account access between the mobile app and the DynamoDB table to sign the requests

**Answer:** C

#### NEW QUESTION 299

- (Exam Topic 1)

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

What deployment policy would satisfy these requirements?

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

**Answer:** A

#### NEW QUESTION 304

- (Exam Topic 1)

A company maintains a REST service using Amazon API Gateway and the API Gateway native API key validation. The company recently launched a new registration page, which allows users to sign up for the service. The registration page creates a new API key using CreateApiKey and sends the new key to the user. When the user attempts to call the API using this key, the user receives a 403 Forbidden error. Existing users are unaffected and can still call the API.

What code updates will grant these new users access to the API?

- A. The createDeployment method must be called so the API can be redeployed to include the newly created API key.
- B. The updateAuthorizer method must be called to update the API's authorizer to include the newly created API key.
- C. The importApiKeys method must be called to import all newly created API keys into the current stage of the API.
- D. The createUsagePlanKey method must be called to associate the newly created API key with the correct usage plan.

**Answer:** D

**Explanation:**

<https://stackoverflow.com/questions/39061041/using-an-api-key-in-amazon-api-gateway>

#### NEW QUESTION 308

- (Exam Topic 1)

A Development team has pushed out 10 applications running on several Amazon EC2 instances. The Operations team is asking for a graphical representation of one key performance metric for each application. These metrics should be available on one screen for easy monitoring. Which steps should the Developer take to accomplish this using Amazon CloudWatch?

- A. Create a custom namespace with a unique metric name for each application.
- B. Create a custom dimension with a unique metric name for each application.
- C. Create a custom event with a unique metric name for each application.
- D. Create a custom alarm with a unique metric name for each application.

**Answer:** A

**Explanation:**

<https://aws.amazon.com/premiumsupport/knowledge-center/cloudwatch-custom-metrics/>

#### NEW QUESTION 311

- (Exam Topic 1)

A company recently migrated its web, application and NoSQL database tiers to AWS. The company is using Auto Scaling to scale the web and application tiers. More than 95 percent of the Amazon DynamoDB requests are repeated read-requests. How can the DynamoDB NoSQL tier be scaled up to cache these repeated requests?

- A. Amazon EMR
- B. Amazon DynamoDB Accelerator
- C. Amazon SQS
- D. Amazon CloudFront

**Answer:** B

**Explanation:**

Reference: <https://aws.amazon.com/dynamodb/dax/>

#### NEW QUESTION 315

- (Exam Topic 1)

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

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

**Answer:** A

#### NEW QUESTION 320

- (Exam Topic 1)

A Developer must build an application that uses Amazon DynamoDB. The requirements state that items being stored in the DynamoDB table will be 7KB in size and that reads must be strongly consistent. The maximum read rate is 3 items per second, and the maximum write rate is 10 items per second. How should the Developer size the DynamoDB table to meet these requirements?

- A. Read: 3 read capacity units Write: 70 write capacity units
- B. Read: 6 read capacity units Write: 70 write capacity units
- C. Read: 6 read capacity units Write: 10 write capacity units
- D. Read: 3 read capacity units Write: 10 write capacity units

**Answer:** B

**Explanation:**

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

#### NEW QUESTION 324

- (Exam Topic 1)

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

- A. Save the secrets in Amazon S3 with AWS KMS server-side encryption, and use a signed URL to access them by using the IAM role from Amazon EC2 instances.
- B. Use the instance metadata to store the secrets and to programmatically access the secrets from EC2 instances.
- C. Use the Amazon DynamoDB client-side encryption library to save the secrets in DynamoDB and to programmatically access the secrets from EC2 instances.
- D. Use AWS SSM Parameter Store to store the secrets and to programmatically access them by using the IAM role from EC2 instances.

**Answer:** D

**Explanation:**

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

**NEW QUESTION 329**

- (Exam Topic 1)

A Developer is developing an application that manages financial transactions. To improve security, multi-factor authentication (MFA) will be required as part of the login protocol.

What services can the Developer use to meet these requirements?

- A. Amazon DynamoDB to store MFA session data, and Amazon SNS to send MFA codes
- B. Amazon Cognito with MFA
- C. AWS Directory Service
- D. AWS IAM with MFA enabled

**Answer:** B

**Explanation:**

AWS documentation - Cognito MFA Managing Security

You can add multi-factor authentication (MFA) to a user pool to protect the identity of your users. MFA adds a second authentication method that doesn't rely solely on user name and password. You can choose to use SMS text messages, or time-based one-time (TOTP) passwords as second factors in signing in your users.

You can also use adaptive authentication with its risk-based model to predict when you might need another authentication factor. It's part of the user pool advanced security features, which also include protections against compromised credentials.

**NEW QUESTION 334**

- (Exam Topic 1)

A company uses Amazon DynamoDB for managing and tracking orders. The DynamoDB table is partitioned based on the order date. The company receives a huge increase in orders during a sales event, causing DynamoDB writes to throttle, and the consumed throughput is far below the provisioned throughput.

According to AWS best practices, how can this issue be resolved with MINIMAL costs?

- A. Create a new DynamoDB table for every order date.
- B. Increase the read and write capacity units of the DynamoDB table.
- C. Add a random number suffix to the partition key values.
- D. Add a global secondary index to the DynamoDB table.

**Answer:** C

**Explanation:**

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

**NEW QUESTION 336**

- (Exam Topic 1)

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

Which combination of steps will grant the application access? (Select TWO.)

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

**Answer:** BC

**Explanation:**

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

**NEW QUESTION 338**

- (Exam Topic 1)

During non-peak hours, a Developer wants to minimize the execution time of a full Amazon DynamoDB table scan without affecting normal workloads. The workloads average half of the strongly consistent read capacity units during non-peak hours.

How would the Developer optimize this scan?

- A. Use parallel scans while limiting the rate
- B. Use sequential scans
- C. Increase read capacity units during the scan operation
- D. Change consistency to eventually consistent during the scan operation

**Answer:** A

**Explanation:**

<https://aws.amazon.com/blogs/developer/rate-limited-scans-in-amazon-dynamodb/>

#### NEW QUESTION 343

- (Exam Topic 1)

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

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

**Answer:** E

**Explanation:**

<https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-visibility-timeout.html> Visibility timeout:  
default value = 30 seconds, minimum = 0 seconds, maximum = 12 hours

#### NEW QUESTION 344

- (Exam Topic 1)

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

How can the Developer provide this authentication functionality?

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

**Answer:** A

**Explanation:**

Reference:

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

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

#### NEW QUESTION 347

- (Exam Topic 1)

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

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

**Answer:** ACE

#### NEW QUESTION 350

- (Exam Topic 1)

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

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

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

- A. Launch and run the load-tester EC2 instance from us-east-1 instead.
- B. Re-configure the load-testing software to re-resolve DNS for each web request.
- C. Use a 3rd-party load-testing service which offers globally-distributed test clients.
- D. Configure ELB and Auto Scaling to distribute across us-west-2a and us-west-2c.
- E. Configure ELB session stickiness to use the app-specific session cookie.

**Answer:** CE

#### NEW QUESTION 355

- (Exam Topic 1)

For a deployment using AWS CodeDeploy, what is the run order of the hooks for in-place deployments?

- A. Before Install -> Application Stop -> Application Start -> After Install
- B. Application Stop -> Before Install -> After Install -> Application Start
- C. Before Install -> Application Stop -> Validate Service -> Application Start
- D. Application Stop -> Before Install -> Validate Service -> Application Start

**Answer:**



B

### NEW QUESTION 358

- (Exam Topic 1)

An e-commerce site allows returning users to log in to display customized web pages. The workflow is shown in the image below:



An application is running on EC2 instances. Amazon RDS is used for the database that stores user accounts and preferences. The website freezes or is slow to load while waiting for the login step to complete. The remaining components of the site are well-optimized.

Which of the following techniques will resolve this issue? (Select Two.)

- A. Implement the user login page as an asynchronous Lambda function.
- B. Use Amazon ElastiCache for MemCached to cache user data.
- C. Use Amazon Application Load Balancer to load balance the traffic to the website.
- D. Call the database asynchronously so the code can continue executing.
- E. Batch login requests from hundreds of users together as a single read request to the database.

**Answer:** BD

#### Explanation:

<https://docs.aws.amazon.com/sdk-for-javascript/v2/developer-guide/making-asynchronous-calls.html>

### NEW QUESTION 360

- (Exam Topic 1)

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

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

**Answer:** C

### NEW QUESTION 361

.....

## THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual DVA-C01 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the DVA-C01 Product From:

<https://www.2passeasy.com/dumps/DVA-C01/>

## Money Back Guarantee

### **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