

# Exam Questions DVA-C01

AWS Certified Developer Associate Exam

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



### NEW QUESTION 1

- (Exam Topic 4)

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

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

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

**Answer:** BD

### NEW QUESTION 2

- (Exam Topic 4)

A developer is working on a Python application that runs on Amazon EC2 instances. The developer wants to enable tracing of application requests to debug performance issues in the code.

Which combination of actions should the developer take to achieve this goal? (Select TWO.)

- A. Install the Amazon CloudWatch agent on the EC2 instances.
- B. Install the AWS X-Ray daemon on the EC2 instances
- C. Configure the application to write JSON-formatted logs to `/var/log/cloudwatch`.
- D. Configure the application to write trace data to `/var/log/xray`.
- E. Install and configure the AWS X-Ray SDK for Python in the application.

**Answer:** BE

### NEW QUESTION 3

- (Exam Topic 4)

A company is hosting a workshop for external users and wants to share the reference documents with the external users for 7 days. The company stores the reference documents in an Amazon S3 bucket that the company owns.

What is the MOST secure way to share the documents with the external users?

- A. Use S3 presigned URLs to share the documents with the external user
- B. Set an expiration time of 7 days.
- C. Move the documents to an Amazon WorkDocs folder
- D. Share the links of the WorkDocs folder with the external users.
- E. Create temporary IAM users that have read-only access to the S3 bucket
- F. Share the access keys with the external user
- G. Expire the credentials after 7 days.
- H. Create a role that has read-only access to the S3 bucket
- I. Share the Amazon Resource Name (ARN) of this role with the external users.

**Answer:** A

### NEW QUESTION 4

- (Exam Topic 4)

A distributed application includes an AWS Lambda function that runs successfully in the DEV environment with 128 MB of memory assigned. The same function is failing in the TEST environment. The developer is monitoring the application using AWS X-Ray, but the Lambda function cannot be seen on the X-Ray service graph. The Lambda execution role has AWS X-Ray permissions.

What is the MOST LIKELY cause for AWS X-Ray not showing any data for the Lambda function?

- A. The AWS SDK needs to be included in the AWS Lambda deployment package.
- B. VPC Flow Logs are not enabled for the application VPC.
- C. Active tracing needs to be enabled for the Lambda function
- D. The memory needs to be increased to 2 GB for the TEST environments.

**Answer:** C

#### Explanation:

<https://stackoverflow.com/questions/43728674/enabling-x-ray-support-in-aws-lambda> You need to check the "Enable Active Tracing" checkbox in the Lambda console.

### NEW QUESTION 5

- (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 CodeCommit 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 6

- (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 7

- (Exam Topic 4)

A company is running an application on AWS Elastic Beanstalk in a single-instance environment. The company's deployments must avoid any downtime. Which deployment option will meet these requirements?

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

**Answer:** D

#### NEW QUESTION 8

- (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 Manager
- 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 Store
- 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 9

- (Exam Topic 4)

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

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

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

**Answer:** C

#### NEW QUESTION 10

- (Exam Topic 4)

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

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

- A. Create a new action with AWS Lambda as the action provider. Create a Lambda function that makes an AWS SDK API call to create the CloudFormation stack.
- B. Create a new action with CloudFormation as the action provider. Set the action mode to CREATE\_UPDATE. Target the CloudFormation stack to be launched.
- C. Create a new action with Jenkins as the action provider.

- D. Create and configure a Jenkins job to make an API call by using the AWS CLI to create the CloudFormation stack.
- E. Create a new action with AWS CodeBuild as the action provider Configure the buildspec to make an API call by using the AWS CLI to create the CloudFormation stack

**Answer:** D

#### NEW QUESTION 10

- (Exam Topic 4)

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

Which solution will meet these requirements MOST cost-effectively?

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

**Answer:** D

#### NEW QUESTION 11

- (Exam Topic 4)

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

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

A)

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

B)

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

C)

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

D)

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

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

**Answer:** A

#### Explanation:

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

#### NEW QUESTION 16

- (Exam Topic 4)

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 21

- (Exam Topic 4)

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

Which solution will meet these requirements?



- A. Create a beta stage for the new version of the AP
- B. Send the updated endpoint to the users.
- C. Create a development stage for the new version of the AP
- D. Use a canary deployment.
- E. Create a development stage for the new version of the AP
- F. Promote a canary release.
- G. Create a deployment stag
- H. Enable mutual TLS for the new version of the API.

**Answer: B**

#### NEW QUESTION 25

- (Exam Topic 4)

An application adds a processing date to each transaction that it receives. The application writes each transaction to an Amazon DynamoDB table by using the PutItem operation. Each transaction has a unique ID (transactionID). Sometimes the application receives transactions more than once.

A developer notices that duplicate transactions in DynamoDB have the latest processing date instead of the date when the transaction was first received. Duplicate records happen infrequently, and most of the transactions are unique.

What is the MOST cost-effective solution that the developer can implement to ensure that PutItem does not update an existing record?

- A. Call the GetItem operation first to confirm that the record does not exist
- B. Then call PutItem
- C. Enable the TTL attribute on the DynamoDB table.
- D. Implement a conditional put by using the attribute\_exists(transactionID) condition expression.
- E. Implement a conditional put by using the attribute\_not\_exists(transactionID) condition expression.

**Answer: C**

#### NEW QUESTION 26

- (Exam Topic 4)

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

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

Which solution will meet these requirements MOST cost-effectively?

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

**Answer: C**

#### Explanation:

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

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

#### NEW QUESTION 31

- (Exam Topic 4)

A data-processing application includes an AWS Lambda function that processes data in several steps. Recently, the function has been reaching the Lambda timeout. A developer wants to use AWS X-Ray to find out how long each step is taking so that the developer can determine which step is causing the timeout.

Which combination of actions should the developer take to accomplish this goal? (Select TWO.)

- A. Modify the application to call the PutMetricData API operation after each processing step
- B. Include the time taken in milliseconds.
- C. Use the aws lambda update-function-configuration AWS CLI command to enable active tracing on the Lambda function.
- D. Modify the application to record each processing step in an X-Ray subsegment by using the X-Ray software development kit (SDK).
- E. Add the xray:PutTraceSegments permission and the xray:PutTelemetryRecords permission to the Lambda function's execution role.
- F. Modify the application to put each processing step in a separate Lambda layer
- G. Include all the layers in the Lambda function.

**Answer: B**

#### NEW QUESTION 35

- (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-redis/>

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 37**

- (Exam Topic 4)

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

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

**Answer:** AC

**NEW QUESTION 39**

- (Exam Topic 4)

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

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

**Answer:** A

**NEW QUESTION 40**

- (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 44**

- (Exam Topic 4)

A company has an ecommerce application. To track product reviews, the company's development team uses an Amazon DynamoDB table.

Every record includes the following:

- A Review ID, a 16-digit universally unique identifier (UUID)
- A Product ID and User ID. 16-digit UUIDs that reference other tables
- A Product Rating on a scale of 1—5 - An optional comment from the user

The table partition key is the Review ID. The most performed query against the table is to find the 10 reviews with the highest rating for a given product.

Which index will provide the FASTEST response for this query?

- A. A global secondary index (GSI) with Product ID as the partition key and Product Rating as the sort key
- B. A global secondary index (GSI) with Product ID as the partition key and Review ID as the sort key
- C. A local secondary index (LSI) with Product ID as the partition key and Product Rating as the sort key
- D. A local secondary index (LSI) with Review ID as the partition key and Product ID as the sort key

**Answer:** A

**NEW QUESTION 47**

- (Exam Topic 3)

A company is building a compute-intensive application that will run on a fleet of Amazon EC2 instances. The application uses attached Amazon EBS disks for storing data. The application will process sensitive information and all the data must be encrypted. What should a developer do to ensure the data is encrypted on disk without impacting performance?

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

**Answer:** A

#### NEW QUESTION 49

- (Exam Topic 3)

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

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

**Answer:** B

#### Explanation:

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

[https://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_roles\\_terms-and-concepts.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_terms-and-concepts.html)

#### NEW QUESTION 50

- (Exam Topic 3)

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

Which AWS service should the developer use?

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

**Answer:** C

#### NEW QUESTION 51

- (Exam Topic 3)

A developer has written an AWS Lambda function using Java as the runtime environment. The developer wants to isolate a performance bottleneck in the code.

Which steps should be taken to reveal the bottleneck?

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

**Answer:** C

#### NEW QUESTION 55

- (Exam Topic 3)

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

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

**Answer:** A

#### NEW QUESTION 59

- (Exam Topic 3)

A developer is building a new application that uses an Amazon DynamoDB table. The specification states that all items that are older than 48 hours must be removed.

Which solution will meet this requirement?

- A. Create a new attribute that has the Number data type. Add a local secondary index (LSI) for this attribute and enable TTL with an expiration of 48 hours. In the application code, set the value of this attribute to the current timestamp for each new item that is being inserted.
- B. Create a new attribute that has the String data type. Add a local secondary index (LSI) for this attribute and enable TTL with an expiration of 48 hours. In the application code, set the value of this attribute to the current timestamp for each new item that is being inserted.
- C. Create a new attribute that has the Number data type. Enable TTL on the DynamoDB table for this attribute in the application code. Set the value of this attribute to the current timestamp plus 48 hours for each new item that is being inserted.

D. Create a new attribute that has the String data type Enable TTL on the DynamoDB table for this attribute In the application code set the value of this attribute to the current timestamp plus 48 hours for each new item that is being inserted

**Answer: C**

#### NEW QUESTION 63

- (Exam Topic 3)

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

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

Which of the following solution supports these requirements?

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

**Answer: C**

#### NEW QUESTION 67

- (Exam Topic 3)

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

```
boto.exception.BotoServerError: BotoServerError: 503 Service Unavailable
<?xml version="1.0" encoding="UTF-8"?>
<Response><Errors><Error><Code>RequestLimitExceeded</Code>
<Message>Request limit exceeded.</Message></Error></Errors><RequestID>bfddec84-53b3-4701-b729-dceefb696ced</RequestID>
</Response>
```

What should the developer do to correct this error message?

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

**Answer: D**

#### NEW QUESTION 69

- (Exam Topic 3)

An application uses Amazon DynamoDB as its backend database The application experiences sudden spikes in traffic over the weekend and variable but predictable spikes during weekdays The capacity needs to be set to avoid throttling errors at all times

How can this be accomplished cost-effectively?

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

**Answer: A**

#### NEW QUESTION 74

- (Exam Topic 3)

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

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

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

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

**Answer: A**

#### NEW QUESTION 79

- (Exam Topic 3)

A developer Is designing an AWS Lambda function that create temporary files that are less than 10 MB during execution. The temporary files will be accessed and modified multiple times during execution. The developer has no need to save or retrieve these files in the future.

Where should the temporary file be stored?

- A. the /tmp directory
- B. Amazon EFS
- C. Amazon EBS



D. Amazon S3

**Answer:** A

#### NEW QUESTION 84

- (Exam Topic 3)

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

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

What is the MOST likely cause of this issue?

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

**Answer:** B

#### NEW QUESTION 85

- (Exam Topic 3)

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

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

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

**Answer:** C

#### NEW QUESTION 87

- (Exam Topic 3)

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

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

**Answer:** C

#### NEW QUESTION 92

- (Exam Topic 3)

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

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

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

**Answer:** D

#### NEW QUESTION 97

- (Exam Topic 3)

A developer is creating an application to process a large number of requests. Requests must be processed in order, and each request should be processed only once. How should Amazon SQS be deployed to achieve this?

- A. Configure First in First out (FIFO) delivery in a standard Amazon SQS queue to process requests.
- B. Use an SQS FIFO queue to process requests.
- C. Use the SetOrder attribute to ensure sequential request processing.
- D. Convert the standard queue to a FIFO queue by renaming the queue to use the fifo suffix.

**Answer:** B

#### NEW QUESTION 100

- (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 101

- (Exam Topic 3)

A gaming application stores scores for players in an Amazon DynamoDB table that has four attributes user\_id, user\_name, user\_score, and user\_rank. The users are allowed to update their names only A user is authenticated by web identity federation.

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

A)

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

B)

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

C)

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

D)

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

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

**Answer:** A

#### NEW QUESTION 103

- (Exam Topic 3)

A developer wants to modify the following AWS Cloud Formation template to embed another CloudFormation stack:

```
{
  "AWSTemplateFormatVersion" : "2010-09-09",
  "Resources" : {
    "cfStack" : {
      "Properties" : {
        "TemplateURL" : "https://s3.amazonaws.com/cloudformation-templates/cf.template",
        "Parameters" : {
          "InstanceType" : "t3.small"
        }
      }
    }
  }
}
```

Which syntax should the developer add to the blank line of the CloudFormation template to meet this requirement?

- A. "Mapping" : "AWS::CloudFormation::Stack",
- B. "Type" : "AWS::CloudFormation::NestedStack",
- C. "Type" : "AWS::CloudFormation::Stack",
- D. "Mapping" : "AWS::CloudFormation::NestedStack",

**Answer:** A

#### NEW QUESTION 108

- (Exam Topic 3)

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

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

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

**Answer:** B

#### NEW QUESTION 109

- (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 110

- (Exam Topic 3)

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

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

**Answer:** D

#### NEW QUESTION 115

- (Exam Topic 3)

A developer has created an AWS Lambda function that is written in Python The Lambda function reads data from objects in Amazon S3 and writes data to an Amazon DynamoDB table

The function is successfully invoked from an S3 event notification when an object is created However, the function fails when it attempts to write to the DynamoDB table What is the MOST likely cause of this issue?

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

**Answer:** C

#### NEW QUESTION 118

- (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 121

- (Exam Topic 3)

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

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

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

**Answer:** D

#### NEW QUESTION 122

- (Exam Topic 3)

A developer is building a WebSocket API using Amazon API Gateway. The payload sent to this API is JSON that includes an action key This key can have three different values create, update, and remove The developer must integrate with different routes based on the value of the action key of the incoming JSON payload. How can the developer accomplish this task with the LEAST amount of configuration?

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

**Answer:** D

#### NEW QUESTION 123

- (Exam Topic 3)

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

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

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

**Answer:** D

#### NEW QUESTION 127

- (Exam Topic 3)

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

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

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

**Answer:** AD

#### NEW QUESTION 132



- (Exam Topic 3)

A company is providing read access to objects in an Amazon S3 bucket for different customers. The company uses IAM permissions to restrict access to the S3 bucket. The customers can access only their own files.

Due to a regulation requirement, the company needs to enforce encryption in transit for interactions with Amazon S3.

Which solution will meet these requirements?

- A. Add a bucket policy to the S3 bucket to deny S3 actions when the `aws:SecureTransport` condition is equal to `false`.
- B. Add a bucket policy to the S3 bucket to deny S3 actions when the `s3:x-amz-acl` condition is equal to `public-read`.
- C. Add an IAM policy to the IAM users to enforce the usage of the AWS SDK.
- D. Add an IAM policy to the IAM users that allows S3 actions when the `s3:x-arnz-acl` condition is equal to `bucket-owner-read`.

**Answer: C**

#### NEW QUESTION 136

- (Exam Topic 3)

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

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

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

**Answer: C**

#### NEW QUESTION 138

- (Exam Topic 3)

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

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

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

**Answer: D**

#### NEW QUESTION 143

- (Exam Topic 3)

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

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

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

**Answer: B**

#### NEW QUESTION 147

- (Exam Topic 2)

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

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

**Answer: D**

#### Explanation:

[https://docs.aws.amazon.com/kms/latest/APIReference/API\\_GenerateDataKey.html](https://docs.aws.amazon.com/kms/latest/APIReference/API_GenerateDataKey.html)

#### NEW QUESTION 149

- (Exam Topic 2)

A Developer is publishing critical log data to a log group in Amazon CloudWatch Logs, which was created 2 months ago. The Developer must encrypt the log data using an AWS KMS customer master key (CMK) so future data can be encrypted to comply with the company's security policy.

How can the Developer meet this requirement?

- A. Use the Cloud Watch Logs console and enable the encrypt feature on the log group.
- B. Use the AWS CLI create-log-group command and specify the key Amazon Resource Name (ARN)
- C. Use the KMs console and associate the CMK with the log group
- D. Use the AWS CLI associate-Kms-key command and specify the key Amazon Resource Name (ARN)

**Answer:** C

#### NEW QUESTION 153

- (Exam Topic 2)

n 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 158

- (Exam Topic 2)

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

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

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

**Answer:** C

#### Explanation:

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

#### NEW QUESTION 160

- (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 162

- (Exam Topic 2)

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

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

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

**Answer:** B

#### Explanation:

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

#### NEW QUESTION 165

- (Exam Topic 2)

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

How can the Developer address this issue?

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 167**

- (Exam Topic 2)

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

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

**Answer:** D

**NEW QUESTION 171**

- (Exam Topic 2)

An application uses Amazon Kinesis Data Streams to ingest and process large streams of data records in real time. Amazon EC2 instances consume and process the data from the shards of the Kinesis data stream by using Amazon Kinesis Client Library (KCL). The application handles the failure scenarios and does not require standby workers. The application reports that a specific shard is receiving more data than expected. To adapt to the changes in the rate of data flow, the “hot” shard is resharded.

Assuming that the initial number of shards in the Kinesis data stream is 4, and after resharding the number of shards increased to 6, what is the maximum number of EC2 instances that can be deployed to process data from all the shards?

- A. 12
- B. 6
- C. 4
- D. 1

**Answer:** B

**Explanation:**

Typically, when you use the KCL, you should ensure that the number of instances does not exceed the number of shards (except for failure standby purposes). Each shard is processed by exactly one KCL worker and has exactly one corresponding record processor, so you never need multiple instances to process one shard. However, one worker can process any number of shards, so it's fine if the number of shards exceeds the number of instances.

<https://docs.aws.amazon.com/streams/latest/dev/kinesis-record-processor-scaling.html>

**NEW QUESTION 174**

- (Exam Topic 2)

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

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

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 175**

- (Exam Topic 2)

The upload of a 15 GB object to Amazon S3 fails. The error message reads: “Your proposed upload exceeds the maximum allowed object size.”

What technique will allow the Developer to upload this object?

- A. Upload the object using the multi-part upload API.
- B. Upload the object over an AWS Direct Connect connection.
- C. Contact AWS Support to increase the object size limit.
- D. Upload the object to another AWS region.

**Answer:** A

**Explanation:**

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

**NEW QUESTION 178**

- (Exam Topic 2)

An ecommerce startup is preparing for an annual sales event As the traffic to the company's application increases, the development team wants to be notified when the Amazon EC2 instance's CPU utilization exceeds 80%.

Which solution will meet this requirement?

- A. Create a custom Amazon CloudWatch alarm that sends a notification to an Amazon SNS topic when the CPU utilization exceeds 80%.
- B. Create a custom AWS CloudTrail alarm that sends a notification to an Amazon SNS topic when the CPU utilization exceeds 80%
- C. Create a cron job on the EC2 instance that executes the --describe-instance-information command on the host instance every 15 minutes and sends the results to an Amazon SNS topic
- D. Create an AWS Lambda function that queries the AWS CloudTrail logs for the CPUUtilization metric every 15 minutes and sends a notification to an Amazon SNS topic when the CPU utilization exceeds 80%

**Answer:** C

#### NEW QUESTION 179

- (Exam Topic 2)

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

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

**Answer:** C

#### NEW QUESTION 183

- (Exam Topic 2)

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

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

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

**Answer:** A

#### Explanation:

[https://aws.amazon.com/blogs/networking-and-content-delivery/amazon-s3-amazon-cloudfront-a-match-made-i](https://aws.amazon.com/blogs/networking-and-content-delivery/amazon-s3-amazon-cloudfront-a-match-made-in-heaven/)

#### NEW QUESTION 184

- (Exam Topic 2)

An application displays a status dashboard. The status is updated by 1 KB messages from an SQS queue. Although the status changes infrequently, the Developer must minimize the time between the message arrival in the queue and the dashboard update.

What technique provides the shortest delay in updating the dashboard?

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

**Answer:** A

#### Explanation:

[https://docs.aws.amazon.com/ko\\_kr/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-short-and-long-p](https://docs.aws.amazon.com/ko_kr/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-short-and-long-polling.html)

#### NEW QUESTION 185

- (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 status code, status description, and headers in the Lambda function
- D. Enable the multi-value headers on the ALB

**Answer:** D

#### NEW QUESTION 186

- (Exam Topic 2)

When developing an AWS Lambda function that processes Amazon Kinesis Data Streams, Administrators within the company must receive a notice that includes the processed data.

How should the Developer write the function to send processed data to the Administrators?

- A. Separate the Lambda handler from the core logic
- B. Use Amazon CloudWatch Events to send the processed data
- C. Publish the processed data to an Amazon SNS topic
- D. Push the processed data to Amazon SQS

**Answer:** C



**Explanation:**

<https://stackoverflow.com/questions/13681213/what-is-the-difference-between-amazon-sns-and-amazon-sqs> <https://stackoverflow.com/questions/31484868/can-you-publish-a-message-to-an-sns-topic-using-an-aws-lambda>

**NEW QUESTION 189**

- (Exam Topic 2)

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

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

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

**Answer:** BD

**Explanation:**

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

**NEW QUESTION 190**

- (Exam Topic 2)

A developer is updating an application deployed on AWS Elastic Beanstalk. The new version is incompatible with the old version. To successfully deploy the update, a full cutover to the new updated version must be performed on all instances at one time, with the ability to roll back changes in case of a deployment failure in the new version.

How can this be performed with the LEAST amount of downtime?

- A. Use the Elastic Beanstalk All at once deployment policy to update all instances simultaneously.
- B. Perform an Elastic Beanstalk Rolling update with additional batch deployment.
- C. Deploy the new version in a new Elastic Beanstalk environment and swap environment URLs.
- D. Perform an Elastic Beanstalk Rolling deployment.

**Answer:** C

**NEW QUESTION 194**

- (Exam Topic 2)

A developer is refactoring a monolithic application. The application takes a POST request and performs several operations. Some of the operations are in parallel while others run sequentially. These operations have been refactored into individual AWS Lambda functions. The POST request will be processed by Amazon API Gateway.

How should the developer invoke the Lambda functions in the same sequence using API Gateway\*?

- A. Use Amazon SQS to invoke the Lambda functions.
- B. Use an AWS Step Functions activity to run the Lambda functions.
- C. Use Amazon SNS to trigger the Lambda functions.
- D. Use an AWS Step Functions state machine to orchestrate the Lambda functions.

**Answer:** A

**NEW QUESTION 197**

- (Exam Topic 2)

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

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

**Answer:** C

**Explanation:**

Reference: [https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial\\_cross-account-with-roles.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html)

**NEW QUESTION 201**

- (Exam Topic 2)

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

What is the simplest way to do this?

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

**Answer:** B

**Explanation:**

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

#### NEW QUESTION 206

- (Exam Topic 2)

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

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

- A. Change the StreamViewType parameter value to NEW\_AND\_OLD\_IMAGES for the DynamoDB table
- B. Configure event source mapping for the Lambda function
- C. Map an Amazon SNS topic to the DynamoDB streams
- D. increase the maximum execution time (timeout) setting of the Lambda function

**Answer:** B

#### Explanation:

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

#### NEW QUESTION 208

- (Exam Topic 2)

An application is expected to process many files. Each file takes four minutes to process each AWS Lambda invocation. The Lambda function does not return any important data.

What is the fastest way to process all the files?

- A. First split the files to make them smaller, then process with synchronous RequestResponse Lambda invocations.
- B. Make synchronous RequestResponse Lambda invocations and process the files one by one.
- C. Make asynchronous Event Lambda invocations and process the files in parallel.
- D. First join all the files, then process it all at once with an asynchronous Event Lambda invocation.

**Answer:** C

#### NEW QUESTION 211

- (Exam Topic 2)

An application deployed on AWS Elastic Beanstalk experiences increased error rates during deployments of new application versions, resulting in service degradation for users. The Development team believes that this is because of the reduction in capacity during the deployment steps. The team would like to change the deployment policy configuration of the environment to an option that maintains full capacity during deployment while using the existing instances. Which deployment policy will meet these requirements while using the existing instances?

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

**Answer:** D

#### Explanation:

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

#### NEW QUESTION 216

- (Exam Topic 2)

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

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

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

**Answer:** B

#### Explanation:

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

#### NEW QUESTION 221

- (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 225

- (Exam Topic 2)

A company is adding stored value for gift card) capability to its highly popular casual gaming website. Users need to be able to trade this value for other users' items on the platform. This would require both users' records be updated as a single transaction, or both users' records to be completely rolled back. Which AWS database options can provide the transactional capability required for this new feature? (Select TWO )

- A. Amazon DynamoDB with operations made with the ConsistentRead parameter set to true
- B. Amazon ElastiCache for Memcached with operations made within a transaction block
- C. Amazon Aurora MySQL with operations made within a transaction block
- D. Amazon DynamoDB with reads and writes made using Transact" operations
- E. Amazon Redshift with operations made within a transaction block

**Answer:** CD

#### NEW QUESTION 230

- (Exam Topic 2)

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

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

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

**Answer:** B

#### Explanation:

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

#### NEW QUESTION 231

- (Exam Topic 2)

An application running on Amazon EC2 opens connections to an Amazon RDS SQL Server database. The developer does not want to store the user name and password for the database in the code. The developer would also like to automatically rotate the credentials.

What is the MOST secure way to store and access the database credentials?

- A. Create an IAM role that has permissions to access the database. Attach the role to the EC2 instance.
- B. Use AWS Secrets Manager to store the credential.
- C. Retrieve the credentials from Secrets Manager as needed.
- D. Store the credentials in an encrypted text file in an Amazon S3 bucket. Configure the EC2 instance's user data to download the credentials from Amazon S3 as the instance boots.
- E. Store the user name and password credentials directly in the source code.
- F. No further action is needed because the source code is stored in a private repository.

**Answer:** B

#### NEW QUESTION 234

- (Exam Topic 2)

An application needs to encrypt data that is written to Amazon S3 where the keys are managed in an on-premises data center and the encryption is handled by S3. Which type of encryption should be used?

- A. Use server-side encryption with Amazon S3-managed keys.
- B. Use server-side encryption with AWS KMS-managed keys.
- C. Use client-side encryption with customer master keys.
- D. Use server-side encryption with customer-provided keys.

**Answer:** A

#### NEW QUESTION 236

- (Exam Topic 2)

A Developer created a new AWS account and must create a scalable AWS Lambda function that meets the following requirements for concurrent execution:

- Average execution time of 100 seconds
- 50 requests per second

Which step must be taken prior to deployment to prevent errors?

- A. Implement dead-letter queues to capture invocation errors.
- B. Add an event source from Amazon API Gateway to the Lambda function.
- C. Implement error handling within the application code.
- D. Contact AWS Support to increase the concurrent execution limits.

**Answer:** D

#### Explanation:

<https://aws.amazon.com/about-aws/whats-new/2017/05/aws-lambda-raises-default-concurrent-execution-limit/>

**NEW QUESTION 238**

- (Exam Topic 2)

A Developer accesses AWS CodeCommit over SSH. The SSH keys configured to access AWS CodeCommit are tied to a user with the following permissions:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "codecommit:BatchGetRepositories",
        "codecommit:Get*",
        "codecommit:List*",
        "codecommit:GitPull"
      ],
      "Resource": "*"
    }
  ]
}
```

The Developer needs to create/delete branches.

Which specific IAM permissions need to be added, based on the principle of least privilege?

- A. "codecommit:CreateBranch" "codecommit>DeleteBranch"
- B. "codecommit:Put"
- C. "codecommit:Update"
- D. "codecommit:\*"

**Answer:** A

**Explanation:**

<https://docs.aws.amazon.com/codecommit/latest/userguide/auth-and-access-control-permissions-reference.html#>

<https://docs.aws.amazon.com/codecommit/latest/userguide/auth-and-access-control-iam-identity-based-access-c>

**NEW QUESTION 243**

- (Exam Topic 2)

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

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

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 248**

- (Exam Topic 2)

A company is managing a NoSQL database on-premises to host a critical component of an application, which is starting to have scaling issues. The company wants to migrate the application to Amazon DynamoDB with the following considerations:

- Optimize frequent queries
- Reduce read latencies
- Plan for frequent queries on certain key attributes of the table Which solution would help achieve these objectives?

- A. Create global secondary indexes on keys that are frequently queried Add the necessary attributes into the indexes.
- B. Create local secondary indexes on keys that are frequently queried DynamoDB will fetch needed attributes from the table .
- C. Create DynamoDB global tables to speed up query responses Use a scan to fetch data from the table.
- D. Create an AWS Auto Scaling policy for the DynamoDB table

**Answer:** A

**Explanation:**

"Global secondary index—An index with a partition key and a sort key that can be different from those on the base table. A global secondary index is considered "global" because queries on the index can span all of the data in the base table, across all partitions.

Local secondary index—An index that has the same partition key as the base table, but a different sort key. A local secondary index is "local" in the sense that every partition of a local secondary index is scoped to a base table partition that has the same partition key value. "



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

#### NEW QUESTION 250

- (Exam Topic 2)

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

How can the developer accomplish this?

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

**Answer: B**

#### Explanation:

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

#### NEW QUESTION 252

- (Exam Topic 2)

A social media company is using Amazon Cognito in order to synchronize profiles across different mobile devices, to enable end users to have a seamless experience.

Which of the following configurations can be used to silently notify users whenever an update is available on all other devices?

- A. Modify the user pool to include all the devices which keep them in sync.
- B. Use the SyncCallback interface to receive notifications on the application.
- C. Use an Amazon Cognito stream to analyze the data and push the notifications.
- D. Use the push synchronization feature with the appropriate IAM role.

**Answer: D**

#### Explanation:

<https://docs.aws.amazon.com/cognito/latest/developerguide/push-sync.html>

#### NEW QUESTION 257

- (Exam Topic 1)

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

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

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

**Answer: B**

#### NEW QUESTION 260

- (Exam Topic 1)

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

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

**Answer: B**

#### NEW QUESTION 264

- (Exam Topic 1)

You have an environment that consists of a public subnet using Amazon VPC and 3 instances that are running in this subnet. These three instances can successfully communicate with other hosts on the Internet. You launch a fourth instance in the same subnet, using the same AMI and security group configuration you used for the others, but find that this instance cannot be accessed from the Internet.

What should you do to enable internet access?

- A. Deploy a NAT instance into the public subnet.
- B. Modify the routing table for the public subnet
- C. Configure a publically routable IP Address In the host OS of the fourth instance.
- D. Assign an Elastic IP address to the fourth instance.

**Answer: C**

#### NEW QUESTION 265

- (Exam Topic 1)

A development team consists of 10 team members. Similar to a home directory for each team member the manager wants to grant access to user-specific folders in an Amazon S3 bucket. For the team member with the username "TeamMemberX", the snippet of the IAM policy looks like this:

```
{"Sid": "AllowS3ActionToFolders","Effect": "Allow", "Action":  
["s3:*"], "Resource":  
["arn:aws:s3:::companyname/home/TeamMemberX/*"] }
```

Instead of creating distinct policies for each team member, what approach can be used to make this policy snippet generic for all team members?

- A. Use IAM policy condition
- B. Use IAM policy principal
- C. Use IAM policy variables
- D. Use IAM policy resource

**Answer: C**

**Explanation:**

> [https://docs.aws.amazon.com/IAM/latest/UserGuide/reference\\_policies\\_variables.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_variables.html) > Use AWS Identity and Access Management (IAM) policy variables as placeholders when you don't know the exact value of a resource or condition key when you write the policy.

**NEW QUESTION 266**

- (Exam Topic 1)

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

How can this objective be met?

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

**Answer: C**

**NEW QUESTION 270**

- (Exam Topic 1)

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

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

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

**Answer: C**

**Explanation:**

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

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

**NEW QUESTION 271**

- (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 272**

- (Exam Topic 1)

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

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

What is the MOST likely cause of the problem?

- A. The Lambda function has insufficient memory and needs to be increased to 1 GB to match the Amazon EC2 instance
- B. Files need to be copied to the same Amazon S3 bucket for processing, so the second bucket needs to be deleted.

- C. Lambda functions have a maximum execution limit of 300 seconds, therefore the function is not completing.
- D. There is a problem with the Java runtime for Lambda, and the function needs to be converted to node.js.

**Answer:** C

#### NEW QUESTION 274

- (Exam Topic 1)

A Developer executed a AWS CLI command and received the error shown below:

```
A client error (UnauthorizedOperation) occurred when calling the RunInstances operation:
You are not authorized to perform this operation. Encoded authorization failure message:
oGsbAaIV7wlfj8zUqebHUANHzFbmKzILlxyj__y9xwhIHk99U_cUq1FIeZnskWDjQ1wSHStVfdCEyZILGoccGpCiC
IhORceWF9rRwFTnEcRJ3N9iTrPAE1WHveC5Z54ALPaWlEjHlLg8CaB8d8lCKmxQuylCm0r1Bf2fHJRujAYopMVmga
8olFmKAl9yn_Z5rI120Q9p5ZIMX28zYM4dTulcJQUQjosgrEejfiIMYDda8l7Ooko9H6VmGJX62KfkRa5l7yE6hhh
2bIwA6tpyCJy2LWFRTe4bafqAyoqkarhPA4mGiZyWn4gSqbo8o-
uqSIvWYPweaKGkampa0arcFR4gBD7Ph097WYBkzX9hVjGppLMy4jpXRv
```

What action should the Developer perform to make this error human-readable?

- A. Make a call to AWS KMS to decode the message.
- B. Use the AWS STS decode-authorization-message API to decode the message.
- C. Use an open source decoding library to decode the message.
- D. Use the AWS IAM decode-authorization-message API to decode this message.

**Answer:** B

#### Explanation:

<https://docs.aws.amazon.com/cli/latest/reference/sts/decode-authorization-message.html>

The message is encoded because the details of the authorization status can constitute privileged information that the user who requested the operation should not see. To decode an authorization status message, a user must be granted permissions via an IAM policy to request the `DecodeAuthorizationMessage` (sts:DecodeAuthorizationMessage ) action.

#### NEW QUESTION 278

- (Exam Topic 1)

Which DynamoDB limits can be raised by contacting AWS support? Choose 2 answers

- A. The number of hash keys per account
- B. The maximum storage used per account
- C. The number of tables per account
- D. The number of local secondary indexes per account
- E. The number of provisioned throughput units per account

**Answer:** CE

#### Explanation:

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

#### NEW QUESTION 282

- (Exam Topic 1)

A company is developing an application that will run on several Amazon EC2 instances in an Auto Scaling group and can access a database running on Amazon EC2. The application needs to store secrets required to connect to the database. The application must allow for periodic secret rotation, and there should be no changes to the application when a secret changes.

What is the SAFEST way to meet these requirements?

- A. Associate an IAM role to the EC2 instance where the application is running with permission to access the database.
- B. Use AWS Systems Manager Parameter Store with the SecureString data type to store secrets.
- C. Configure the application to store secrets in Amazon S3 object metadata.
- D. Hard code the database secrets in the application code itself.

**Answer:** B

#### NEW QUESTION 286

- (Exam Topic 1)

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

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

**Answer:** AE

#### NEW QUESTION 288

- (Exam Topic 1)

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

Assuming the SQS queue is configured with the default VisibilityTimeout value, what is the BEST way, upon receiving a message, to ensure that no other

instances can retrieve a message that has already been processed or is currently being processed?

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

**Answer:** A

**Explanation:**

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

**NEW QUESTION 290**

- (Exam Topic 1)

What item operation allows the retrieval of multiple items from a DynamoDB table in a single API call?

- A. GetItem
- B. BatchGetItem
- C. GetMultipleItems
- D. GetItemRange

**Answer:** B

**Explanation:**

[https://docs.aws.amazon.com/amazondynamodb/latest/APIReference/API\\_BatchGetItem.html](https://docs.aws.amazon.com/amazondynamodb/latest/APIReference/API_BatchGetItem.html)

**NEW QUESTION 293**

- (Exam Topic 1)

What is the format of structured notification messages sent by Amazon SNS?

- A. An XML object containing MessageId, UnsubscribeURL, Subject, Message and other values
- B. An JSON object containing MessageId, DuplicateFlag, Message and other values
- C. An XML object containing MessageId, DuplicateFlag, Message and other values
- D. An JSON object containing MessageId, unsubscribeURL, Subject, Message and other values

**Answer:** D

**Explanation:**

<https://docs.aws.amazon.com/sns/latest/dg/sns-message-and-json-formats.html#http-notification-json>

**NEW QUESTION 298**

- (Exam Topic 1)

A Developer created configuration specifications for an AWS Elastic Beanstalk application in a file named healthcheckurl.yaml in the .ebextensions/ directory of their application source bundle. The file contains the following:

```
option_settings:
  - namespace: aws:elasticbeanstalk:application
    option_name: Application Healthcheck URL
    value: /health_check
```

After the application launches, the health check is not being run on the correct path, even though it is valid. What can be done to correct this configuration file?

- A. Convert the file to JSON format.
- B. Rename the file to a .config extension.
- C. Change the configuration section from options\_settings to resources.
- D. Change the namespace of the option settings to a custom namespace.

**Answer:** B

**Explanation:**

Reference: <https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/ebextensions.html>

You can add AWS Elastic Beanstalk configuration files (.ebextensions) to your web application's source code to configure your environment and customize the AWS resources that it contains. Configuration files are YAML- or JSON-formatted documents with a .config file extension that you place in a folder named .ebextensions and deploy in your application source bundle. <https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/ebextensions.html>

**NEW QUESTION 300**

- (Exam Topic 1)

A company has multiple Developers located across the globe who are updating code incrementally for a development project. When Developers upload code concurrently, internet connectivity is slow, and it is taking a long time to upload code for deployment in AWS Elastic Beanstalk. Which step will result in minimized upload and deployment time with the LEAST amount of administrative effort?

- A. Allow the Developers to upload the code to an Amazon S3 bucket, and deploy it directly to Elastic Beanstalk.



- B. Allow the Developers to upload the code to a central FTP server to deploy the application to Elastic Beanstalk.
- C. Create an AWS CodeCommit repository, allow the Developers to commit code to it, and then directly deploy the code to Elastic Beanstalk.
- D. Create a code repository on an Amazon EC2 instance so that all Developers can update the code, and deploy the application from the instance to Elastic Beanstalk.

**Answer:** B

**Explanation:**

<https://aws.amazon.com/premiumsupport/knowledge-center/deploy-codecommit-elastic-beanstalk/>

#### NEW QUESTION 304

- (Exam Topic 1)

An application is designed to use Amazon SQS to manage messages from many independent senders. Each sender's messages must be processed in the order they are received.

Which SQS feature should be implemented by the Developer?

- A. Configure each sender with a unique MessageGroupId
- B. Enable MessageDeduplicationIds on the SQS queue
- C. Configure each message with unique MessageGroupIds.
- D. Enable ContentBasedDeduplication on the SQS queue

**Answer:** A

#### NEW QUESTION 305

- (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 310

- (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 312

- (Exam Topic 1)

Which of the following statements about SQS is true?

- A. Messages will be delivered exactly once and messages will be delivered in First in, First out order
- B. Messages will be delivered exactly once and message delivery order is indeterminate
- C. Messages will be delivered one or more times and messages will be delivered in First in, First out order
- D. Messages will be delivered one or more times and message delivery order is indeterminate

**Answer:** D

**Explanation:**

<https://aws.amazon.com/sqs/features/>

#### NEW QUESTION 317

- (Exam Topic 1)

How should custom libraries be utilized in AWS Lambda?

- A. Host the library on Amazon S3 and reference to it from the Lambda function.
- B. Install the library locally and upload a ZIP file of the Lambda function.
- C. Import the necessary Lambda blueprint when creating the function.
- D. Modify the function runtime to include the necessary library.

**Answer:** D

**Explanation:**

Reference: [https://docs.aws.amazon.com/lambda/latest/dg/env\\_variables.html](https://docs.aws.amazon.com/lambda/latest/dg/env_variables.html)

**NEW QUESTION 322**

- (Exam Topic 1)

If an application is storing hourly log files from thousands of instances from a high traffic web site, which naming scheme would give optimal performance on S3?

- A. Sequential
- B. instanceID\_log-HH-DD-MM-YYYY
- C. instanceID\_log-YYYY-MM-DD-HH
- D. HH-DD-MM-YYYY-log\_instanceID
- E. YYYY-MM-DD-HH-log\_instanceID

**Answer:** B

**Explanation:**

Reference:

[https://acloud.guru/forums/aws-certified-developer-associate/discussion/-KU2dEtJb-LI5ISbH\\_S4/if-an-applicati](https://acloud.guru/forums/aws-certified-developer-associate/discussion/-KU2dEtJb-LI5ISbH_S4/if-an-applicati)

**NEW QUESTION 323**

- (Exam Topic 1)

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

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

**Answer:** AC

**Explanation:**

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

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

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

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

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

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

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

**NEW QUESTION 324**

- (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 329**

- (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 334**

- (Exam Topic 1)

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

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

**Answer:** C

**Explanation:**

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

#### NEW QUESTION 339

- (Exam Topic 1)

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

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

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

**Answer:** BE

#### NEW QUESTION 344

- (Exam Topic 1)

An application is real-time processing millions of events that are received through an API.

What service could be used to allow multiple consumers to process the data concurrently and MOST cost-effectively?

- A. Amazon SNS with fanout to an SQS queue for each application
- B. Amazon SNS with fanout to an SQS FIFO (first-in, first-out) queue for each application
- C. Amazon Kinesis Firehouse
- D. Amazon Kinesis Streams

**Answer:** D

#### NEW QUESTION 345

- (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 350

- (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 353

- (Exam Topic 1)

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

How can the Developer increase the Lambda compute resources?

- A. Run on a larger instance size with more compute capacity.
- B. Increase the maximum execution time.

- C. Specify a larger compute capacity when calling the Lambda function.
- D. Increase the allocated memory for the Lambda function.

**Answer:** D

#### NEW QUESTION 356

- (Exam Topic 1)

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

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

**Answer:** D

#### NEW QUESTION 360

- (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 361

- (Exam Topic 1)

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

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

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

**Answer:** D

#### Explanation:

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

#### NEW QUESTION 366

- (Exam Topic 1)

You attempt to store an object in the US-STANDARD region in Amazon S3, and receive a confirmation that it has been successfully stored. You then immediately make another API call and attempt to read this object. S3 tells you that the object does not exist. What could explain this behavior?

- A. US-STANDARD uses eventual consistency and it can take time for an object to be readable in a bucket
- B. Objects in Amazon S3 do not become visible until they are replicated to a second region.
- C. US-STANDARD imposes a 1 second delay before new objects are readable.
- D. You exceeded the bucket object limit, and once this limit is raised the object will be visible.

**Answer:** A

#### Explanation:

<https://acloud.guru/forums/aws-certified-developer-associate/discussion/-KGngHzVQ03OpeAA9jSP/i-cant-ans> [https://acloud.guru/forums/aws-certified-developer-associate/discussion/-K5WKXRAIJdOu58GREF\\_/s3-questio](https://acloud.guru/forums/aws-certified-developer-associate/discussion/-K5WKXRAIJdOu58GREF_/s3-questio)

#### NEW QUESTION 369



- (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 374

- (Exam Topic 1)

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

What step is required to update the function with the changes?

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

**Answer:** B

**Explanation:**

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

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

#### NEW QUESTION 375

.....

## 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