

➤ **Vendor: Amazon**

➤ **Exam Code: DVA-C02**

➤ **Exam Name: AWS Certified Developer - Associate (DVA-C02)**

➤ **New Updated Questions from [Braindump2go](#) (Updated in [March/2024](#))**

[Visit Braindump2go and Download Full Version DVA-C02 Exam Dumps](#)

QUESTION 343

A developer is storing many objects in a single Amazon S3 bucket. The developer needs to optimize the S3 bucket for high request rates.

How should the developer store the objects to meet this requirement?

- A. Store the objects by using S3 Intelligent-Tiering.
- B. Store the objects at the root of the S3 bucket.
- C. Store the objects by using object key names distributed across multiple prefixes.
- D. Store each object with an object tag named "prefix" that contains a unique value.

Answer: C

Explanation:

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

QUESTION 344

A company deploys a new application to AWS. The company is streaming application logs to Amazon CloudWatch Logs. The company's development team must receive notification by email when the word "ERROR" appears in any log lines. A developer sets up an Amazon Simple Notification Service (Amazon SNS) topic and subscribes the development team to the topic.

What should the developer do next to meet the requirements?

- A. Select the appropriate log group. Create a CloudWatch metric filter with "ERROR" as the search term. Create an alarm on this metric that notifies the SNS topic when the metric is 1 or higher.
- B. In CloudWatch Logs Insights, select the appropriate log group. Create a metric query to search for the term "ERROR" in the logs. Create an alarm on this metric that notifies the SNS topic when the metric is 1 or higher.
- C. Select the appropriate log group. Create an SNS subscription filter with "ERROR" as the filter pattern. Select the SNS topic as the destination.
- D. Create a CloudWatch alarm that includes "ERROR" as a filter pattern, a log group dimension that defines the appropriate log group, and a destination that notifies the SNS topic.

Answer: A

Explanation:

<https://docs.aws.amazon.com/sns/latest/dg/sns-monitoring-using-cloudwatch.html>

QUESTION 345

A company uses Amazon Simple Queue Service (Amazon SQS) to decouple its microservices architecture. Some messages in an SQS queue contain sensitive information. A developer must implement a solution that encrypts all the data at rest.

Which solution will meet this requirement?

[DP-500 Exam Dumps](#) [DP-500 Exam Questions](#) [DP-500 PDF Dumps](#) [DP-500 VCE Dumps](#)

<https://www.braindump2go.com/dva-c02.html>

- A. Enable server-side encryption for the SQS queue by using an SQS managed encryption key (SSE-SQS).
- B. Use the aws:SecureTransport condition in the queue policy to ensure that only HTTPS (TLS) is used for all requests to the SQS queue.
- C. Use AWS Certificate Manager (ACM) to generate an SSL/TLS certificate. Reference the certificate when messages are sent to the queue.
- D. Set a message attribute in the SQS SendMessage request for messages that are sent to the queue. Set the Name to ENCRYPT. Set the Value to TRUE.

Answer: A

Explanation:

<https://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/sqs-server-side-encryption.html>

QUESTION 346

A company recently deployed a new serverless user portal. Users have reported that part of the portal is slow. The initial analysis found a single Amazon API Gateway endpoint that is responsible for the performance issues. The endpoint integrates with an AWS Lambda function. However, the Lambda function interacts with other APIs and AWS services.

How can a developer find the source of the increased response time by using operational best practices?

- A. Update the Lambda function by adding logging statements with high-precision timestamps before and after each external request. Deploy the updated Lambda function. After accumulating enough usage data, examine the Amazon CloudWatch logs for the Lambda function to determine the likely sources for the increased response time.
- B. Instrument the Lambda function with the AWS X-Ray SDK. Add HTTP and HTTPS interceptors and SDK client handlers. Deploy the updated Lambda function. Turn on X-Ray tracing. After accumulating enough usage data, use the X-Ray service map to examine the average response times to determine the likely sources.
- C. Review the Lambda function's Amazon CloudWatch metrics by using the metrics explorer. Apply anomaly detection to the Duration metric and the Throttles metric. Review the anomalies to determine the likely sources.
- D. Use Amazon CloudWatch Synthetics to create a new canary. Turn on AWS X-Ray tracing on the canary. Configure the canary to scan the user portal. After accumulating enough usage data, use the CloudWatch Synthetics canary dashboard to view the metrics from the canary.

Answer: B

QUESTION 347

A developer is building an event-driven application by using AWS Lambda and Amazon EventBridge. The Lambda function needs to push events to an EventBridge event bus. The developer uses an SDK to run the PutEvents EventBridge action and specifies no credentials in the code. After deploying the Lambda function, the developer notices that the function is failing and there are AccessDeniedException errors in the logs.

How should the developer resolve this issue?

- A. Configure a VPC peering connection between the Lambda function and EventBridge.
- B. Modify their AWS credentials to include permissions for the PutEvents EventBridge action.
- C. Modify the Lambda function execution role to include permissions for the PutEvents EventBridge action.
- D. Add a resource-based policy to the Lambda function to include permissions for the PutEvents EventBridge action.

Answer: C

QUESTION 348

A company's application has an AWS Lambda function that processes messages from IoT devices. The company wants to monitor the Lambda function to ensure that the Lambda function is meeting its required service level

agreement (SLA).

A developer must implement a solution to determine the application's throughput in near real time. The throughput must be based on the number of messages that the Lambda function receives and processes in a given time period. The Lambda function performs initialization and post-processing steps that must not factor into the throughput measurement.

What should the developer do to meet these requirements?

- A. Use the Lambda function's ConcurrentExecutions metric in Amazon CloudWatch to measure the throughput.
- B. Modify the application to log the calculated throughput to Amazon CloudWatch Logs. Use Amazon EventBridge to invoke a separate Lambda function to process the logs on a schedule.
- C. Modify the application to publish custom Amazon CloudWatch metrics when the Lambda function receives and processes each message. Use the metrics to calculate the throughput.
- D. Use the Lambda function's Invocations metric and Duration metric to calculate the throughput in Amazon CloudWatch.

Answer: C

Explanation:

Use the metrics to calculate the throughput. This is because custom metrics can provide a more accurate measure of throughput, as they can be configured to only increment when a message is received and processed by the Lambda function. This would exclude the time spent on initialization and post-processing, which are not part of the throughput measurement.

QUESTION 349

A developer is using an AWS CodePipeline pipeline to provide continuous integration and continuous delivery (CI/CD) support for a Java application. The developer needs to update the pipeline to support the introduction of a new application dependency .jar file. The pipeline must start a build when a new version of the .jar file becomes available. Which solution will meet these requirements?

- A. Create an Amazon S3 bucket to store the dependency .jar file. Publish the dependency .jar file to the S3 bucket. Use an Amazon Simple Notification Service (Amazon SNS) notification to start a CodePipeline pipeline build.
- B. Create an Amazon Elastic Container Registry (Amazon ECR) private repository. Publish the dependency .jar file to the repository. Use an ECR source action to start a CodePipeline pipeline build.
- C. Create an Amazon Elastic Container Registry (Amazon ECR) private repository. Publish the dependency .jar file to the repository. Use an Amazon Simple Notification Service (Amazon SNS) notification to start a CodePipeline pipeline build.
- D. Create an AWS CodeArtifact repository. Publish the dependency .jar file to the repository. Use an Amazon EventBridge rule to start a CodePipeline pipeline build.

Answer: D

Explanation:

<https://docs.aws.amazon.com/codeartifact/latest/ug/configure-service-events-codepipeline.html#configure-service-events-codepipeline-create-rule>

QUESTION 350

A company with multiple branch locations has an analytics and reporting application. Each branch office pushes a sales report to a shared Amazon S3 bucket at a predefined time each day. The company has developed an AWS Lambda function that analyzes the reports from all branch offices in a single pass. The Lambda function stores the results in a database.

The company needs to start the analysis once each day at a specific time.

Which solution will meet these requirements MOST cost-effectively?

- A. Configure an S3 event notification to invoke the Lambda function when a branch office uploads a sales report.
- B. Create an AWS Step Functions state machine that invokes the Lambda function once each day at

[DP-500 Exam Dumps](#) [DP-500 Exam Questions](#) [DP-500 PDF Dumps](#) [DP-500 VCE Dumps](#)

<https://www.braindump2go.com/dva-c02.html>

the predefined time.

- C. Configure the Lambda function to run continuously and to begin analysis only at the predefined time each day.
- D. Create an Amazon EventBridge scheduled rule that invokes the Lambda function once each day at the predefined time.

Answer: D

QUESTION 351

A developer has an application that asynchronously invokes an AWS Lambda function. The developer wants to store messages that resulted in failed invocations of the Lambda function so that the application can retry the call later. What should the developer do to accomplish this goal with the LEAST operational overhead?

- A. Set up Amazon CloudWatch Logs log groups to filter and store the messages in an Amazon S3 bucket. Import the messages in Lambda. Run the Lambda function again.
- B. Configure Amazon EventBridge to send the messages to Amazon Simple Notification Service (Amazon SNS) to initiate the Lambda function again.
- C. Implement a dead-letter queue for discarded messages. Set the dead-letter queue as an event source for the Lambda function.
- D. Send Amazon EventBridge events to an Amazon Simple Queue Service (Amazon SQS) queue. Configure the Lambda function to pull messages from the SQS queue. Run the Lambda function again.

Answer: C

Explanation:

<https://docs.aws.amazon.com/lambda/latest/dg/invoke-async.html#invoke-dlq>

QUESTION 352

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

Explanation:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-cfn-updating-stacks-changesets.html>

QUESTION 353

A company stores all personally identifiable information (PII) in an Amazon DynamoDB table named PII in Account A. Developers are working on an application that is running on Amazon EC2 instances in Account B. The application in Account B requires access to the PII table.

An administrator in Account A creates an IAM role named AccessPII that has permission to access the PII table. The administrator also creates a trust policy that specifies Account B as a principal that can assume the role.

Which combination of steps should the developers take in Account B to allow their application to access the PII table? (Choose two.)

- A. Allow the EC2 IAM role the permission to assume the AccessPII role.
- B. Allow the EC2 IAM role the permission to access the PII table.
- C. Include the AWS API in the application code logic to obtain temporary credentials from the EC2 IAM role to access the PII table.
- D. Include the AssumeRole API operation in the application code logic to obtain temporary credentials to access the PII table.

- E. Include the GetSessionToken API operation in the application code logic to obtain temporary credentials to access the PII table.

Answer: AD

QUESTION 354

A gaming website gives users the ability to trade game items with each other on the platform. The platform requires both users' records to be updated and persisted in one transaction. If any update fails, the transaction must roll back. Which AWS solutions can provide the transactional capability that is required for this feature? (Choose 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 DynamoDB with reads and writes made by using Transact* operations
- D. Amazon Aurora MySQL with operations made within a transaction block
- E. Amazon Athena with operations made within a transaction block

Answer: CD

Explanation:

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/transaction-apis.html>

QUESTION 355

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

How should the developer reference the parameter that contains the database hostname?

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

Answer: A

QUESTION 356

A company uses an AWS Lambda function to call a third-party service. The third-party service has a limit of requests each minute. If the number of requests exceeds the limit, the third-party service returns rate-limiting errors. A developer needs to configure the Lambda function to avoid receiving rate limiting errors from the third-party service. Which solution will meet these requirements?

- A. Set the reserved concurrency on the Lambda function to match the number of concurrent requests that the third-party service allows.
- B. Decrease the memory that is allocated to the Lambda function.
- C. Set the provisioned concurrency on the Lambda function to match the number of concurrent requests that the third-party service allows.
- D. Increase the timeout value that is specified on the Lambda function.

Answer: A

Explanation:

<https://docs.aws.amazon.com/lambda/latest/dg/provisioned-concurrency.html>

QUESTION 357

A developer is building a new containerized application by using AWS Copilot. The developer uses the AWS Copilot command line interface (CLI) to deploy the application during development. The developer committed the application code to a new AWS CodeCommit repository. The developer must create an automated deployment process before releasing the new application to production.

[DP-500 Exam Dumps](#) [DP-500 Exam Questions](#) [DP-500 PDF Dumps](#) [DP-500 VCE Dumps](#)

<https://www.braindump2go.com/dva-c02.html>

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

- A. Create a buildspec file that invokes the AWS Copilot CLI commands to build and deploy the application. Use the AWS Copilot CLI to create an AWS CodePipeline that uses the CodeCommit repository in the source stage and AWS CodeBuild in the build stage.
- B. Use the AWS Serverless Application Model (AWS SAM) CLI to bootstrap and initialize an AWS CodePipeline configuration. Use the CodeCommit repository as the source. Invoke the AWS Copilot CLI to build and deploy the application.
- C. Use the AWS Copilot CLI to define the AWS Copilot pipeline and to deploy the AWS CodePipeline. Select CodeCommit as the source for the AWS CodePipeline.
- D. Define an AWS CloudFormation template for an AWS CodePipeline with CodeCommit as the source. Configure the template as an AWS Copilot CLI add-on. Use the AWS Copilot CLI to deploy the application.

Answer: C

QUESTION 358

A developer is creating a new application for a pet store. The application will manage customer rewards points. The developer will use Amazon DynamoDB to store the data for the application. The developer needs to optimize query performance and limit partition overload before actual performance analysis.

Which option should the developer use for a partition key to meet these requirements?

- A. A randomly generated universally unique identifier (UUID)
- B. The customer's full name
- C. The date when the customer signed up for the rewards program
- D. The name of the customer's pet

Answer: A

QUESTION 359

A developer uses AWS IAM Identity Center (AWS Single Sign-On) to interact with the AWS CLI and AWS SDKs on a local workstation. API calls to AWS services were working when the SSO access was first configured. However, the developer is now receiving Access Denied errors. The developer has not changed any configuration files or scripts that were previously working on the workstation.

What is the MOST likely cause of the developer's access issue?

- A. The access permissions to the developer's AWS CLI binary file have changed.
- B. The permission set that is assumed by IAM Identity Center does not have the necessary permissions to complete the API call.
- C. The credentials from the IAM Identity Center federated role have expired.
- D. The developer is attempting to make API calls to the incorrect AWS account.

Answer: C

Explanation:

<https://docs.aws.amazon.com/IAM/latest/UserGuide/security-creds.html>