

- **Vendor: Amazon**
- **Exam Code: DVA-C02**
- **Exam Name: AWS Certified Developer - Associate (DVA-C02)**
- **New Updated Questions from [Braindump2go](#)**
- **(Updated in [November/2023](#))**

### **Visit Braindump2go and Download Full Version DVA-C02 Exam Dumps**

#### **QUESTION 226**

A developer is creating an AWS Lambda function that will generate and export a file. The function requires 100 MB of temporary storage for temporary files while running. 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 Amazon Elastic Block Store (Amazon EBS) and delete the files at the end of the Lambda function.
- B. Copy the files to Amazon Elastic File System (Amazon 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 Amazon S3 bucket with a lifecycle policy to delete the files.

**Answer: C**

#### **Explanation:**

The /tmp directory is a dedicated temporary storage location provided by AWS Lambda for storing temporary files during the execution of the function.

It's cost-effective and efficient because it doesn't involve additional AWS services or storage costs.

AWS Lambda automatically manages the /tmp directory for you, including clearing its contents after the function execution is complete. You don't need to explicitly delete the files; Lambda takes care of it.

#### **QUESTION 227**

A company uses Amazon DynamoDB as a data store for its order management system. The company frontend application stores orders in a DynamoDB table. The DynamoDB table is configured to send change events to a DynamoDB stream. The company uses an AWS Lambda function to log and process the incoming orders based on data from the DynamoDB stream.

An operational review reveals that the order quantity of incoming orders is sometimes set to 0. A developer needs to create a dashboard that will show how many unique customers this problem affects each day.

What should the developer do to implement the dashboard?

- A. Grant the Lambda function's execution role permissions to upload logs to Amazon CloudWatch Logs. Implement a CloudWatch Logs Insights query that selects the number of unique customers for orders with order quantity equal to 0 and groups the results in 1-day periods. Add the CloudWatch Logs Insights query to a CloudWatch dashboard.
- B. Use Amazon Athena to query AWS CloudTrail API logs for API calls. Implement an Athena query that selects the number of unique customers for orders with order quantity equal to 0 and groups the results in 1-day periods. Add the Athena query to an Amazon CloudWatch dashboard.
- C. Configure the Lambda function to send events to Amazon EventBridge. Create an EventBridge rule that groups the number of unique customers for orders with order quantity equal to 0 in 1-day periods. Add a CloudWatch dashboard as the target of the rule.

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

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

- D. Turn on custom Amazon CloudWatch metrics for the DynamoDB stream of the DynamoDB table. Create a CloudWatch alarm that groups the number of unique customers for orders with order quantity equal to 0 in 1-day periods. Add the CloudWatch alarm to a CloudWatch dashboard.

**Answer:** D

#### **QUESTION 228**

A developer needs to troubleshoot an AWS Lambda function in a development environment. The Lambda function is configured in VPC mode and needs to connect to an existing Amazon RDS for SQL Server DB instance. The DB instance is deployed in a private subnet and accepts connections by using port 1433.

When the developer tests the function, the function reports an error when it tries to connect to the database. Which combination of steps should the developer take to diagnose this issue? (Choose two.)

- A. Check that the function's security group has outbound access on port 1433 to the DB instance's security group. Check that the DB instance's security group has inbound access on port 1433 from the function's security group.
- B. Check that the function's security group has inbound access on port 1433 from the DB instance's security group. Check that the DB instance's security group has outbound access on port 1433 to the function's security group.
- C. Check that the VPC is set up for a NAT gateway. Check that the DB instance has the public access option turned on.
- D. Check that the function's execution role permissions include `rds:DescribeDBInstances`, `rds:ModifyDBInstance`, and `rds:DescribeDBSecurityGroups` for the DB instance.
- E. Check that the function's execution role permissions include `ec2:CreateNetworkInterface`, `ec2:DescribeNetworkInterfaces`, and `ec2:DeleteNetworkInterface`.

**Answer:** AD

**Explanation:**

<https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/lambda-rds-connect.html>

#### **QUESTION 229**

A developer needs to launch a new Amazon EC2 instance by using the AWS CLI.

Which AWS CLI command should the developer use to meet this requirement?

- A. `aws ec2 bundle-instance`
- B. `aws ec2 start-instances`
- C. `aws ec2 confirm-product-instance`
- D. `aws ec2 run-instances`

**Answer:** D

**Explanation:**

So, to create a new EC2 instance using the AWS CLI, you would typically use the `aws ec2 run-instances` command, providing the necessary parameters such as the AMI ID, instance type, security groups, and key pair, among others.

#### **QUESTION 230**

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

Which approach addresses these requirements?

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

**Answer:** D

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

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

**Explanation:**

AWS CloudFormation: It allows you to define your infrastructure as code using templates, which can be version-controlled. You can create, update, and delete stacks of AWS resources in a controlled and predictable manner. This aligns with the requirement to deploy multiple identical copies of the infrastructure, stage changes, and revert to previous versions.

AWS CodeCommit: It provides a fully managed source control service, allowing you to store and version-control your CloudFormation templates. This ensures that you can manage and track changes to your infrastructure configurations.

**QUESTION 231**

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: D**

**Explanation:**

dynamodb:GetItem: This permission allows the Lambda function to retrieve an item from DynamoDB.

dynamodb:UpdateItem: This permission allows the Lambda function to update the attributes of an item in DynamoDB.

dynamodb:PutItem: This permission allows the Lambda function to create a new item if it doesn't already exist in the DynamoDB table.

**QUESTION 232**

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

What could be causing this issue?

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

**Answer: A**

**Explanation:**

In a caching setup using Amazon ElastiCache in front of Amazon DynamoDB, if the cache is not being invalidated or updated when data in DynamoDB is changed, it can result in stale data being served from the cache, leading to the observed behavior.

To resolve this issue, you should implement a mechanism to invalidate or update the cache whenever the price of an item is changed in DynamoDB to ensure that the most up-to-date data is retrieved from the cache or DynamoDB.

**QUESTION 233**

A company requires that all applications running on Amazon EC2 use IAM roles to gain access to AWS services. A developer is modifying an application that currently relies on IAM user access keys stored in environment variables to access Amazon DynamoDB tables using boto, the AWS SDK for Python.

The developer associated a role with the same permissions as the IAM user to the EC2 instance, then deleted the IAM user. When the application was restarted, the AWS AccessDeniedException messages started appearing in the application logs. The developer was able to use their personal account on the server to run DynamoDB API commands using the AWS CLI.

What is the MOST likely cause of the exception?

- A. IAM policies might take a few minutes to propagate to resources.
- B. Disabled environment variable credentials are still being used by the application.
- C. The AWS SDK does not support credentials obtained using an instance role.
- D. The instance's security group does not allow access to http://169.254.169.254.

**Answer: B**

#### **QUESTION 234**

A company has an existing application that has hardcoded database credentials. A developer needs to modify the existing application. The application is deployed in two AWS Regions with an active-passive failover configuration to meet company's disaster recovery strategy.

The developer needs a solution to store the credentials outside the code. The solution must comply with the company's disaster recovery strategy.

Which solution will meet these requirements in the MOST secure way?

- A. Store the credentials in AWS Secrets Manager in the primary Region. Enable secret replication to the secondary Region. Update the application to use the Amazon Resource Name (ARN) based on the Region.
- B. Store credentials in AWS Systems Manager Parameter Store in the primary Region. Enable parameter replication to the secondary Region. Update the application to use the Amazon Resource Name (ARN) based on the Region.
- C. Store credentials in a config file. Upload the config file to an S3 bucket in the primary Region. Enable Cross-Region Replication (CRR) to an S3 bucket in the secondary region. Update the application to access the config file from the S3 bucket, based on the Region.
- D. Store credentials in a config file. Upload the config file to an Amazon Elastic File System (Amazon EFS) file system. Update the application to use the Amazon EFS file system Regional endpoints to access the config file in the primary and secondary Regions.

**Answer: A**

**Explanation:**

<https://docs.aws.amazon.com/secretsmanager/latest/userguide/create-manage-multi-region-secrets.html>

#### **QUESTION 235**

A developer is receiving HTTP 400: ThrottlingException errors intermittently when calling the Amazon CloudWatch API. When a call fails, no data is retrieved.

What best practice should first be applied to address this issue?

- A. Contact AWS Support for a limit increase.
- B. Use the AWS CLI to get the metrics.
- C. Analyze the applications and remove the API call.
- D. Retry the call with exponential backoff.

**Answer: D**

#### **QUESTION 236**

An application needs to use the IP address of the client in its processing. The application has been moved into AWS and has been placed behind an Application Load Balancer (ALB). However, all the client IP addresses now appear to be the same. The application must maintain the ability to scale horizontally.

Based on this scenario, what is the MOST cost-effective solution to this problem?

- A. Remove the application from the ALB. Delete the ALB and change Amazon Route 53 to direct traffic to the instance running the application.
- B. Remove the application from the ALB. Create a Classic Load Balancer in its place. Direct traffic to the application using the HTTP protocol.
- C. Alter the application code to inspect the X-Forwarded-For header. Ensure that the code can work properly if a list of IP addresses is passed in the header.
- D. Alter the application code to inspect a custom header. Alter the client code to pass the IP address in the custom header.

**Answer: C**

**Explanation:**

If you need to see external IP address and your app is behind ALB, always use x-forwarded-for.  
<https://docs.aws.amazon.com/elasticloadbalancing/latest/application/x-forwarded-headers.html>

#### **QUESTION 237**

A developer is designing a serverless application that customers use to select seats for a concert venue. Customers send the ticket requests to an Amazon API Gateway API with an AWS Lambda function that acknowledges the order and generates an order ID. The application includes two additional Lambda functions: one for inventory management and one for payment processing. These two Lambda functions run in parallel and write the order to an Amazon Dynamo DB table.

The application must provide seats to customers according to the following requirements. If a seat is accidentally sold more than once, the first order that the application received must get the seat. In these cases, the application must process the payment for only the first order. However, if the first order is rejected during payment processing, the second order must get the seat. In these cases, the application must process the payment for the second order. Which solution will meet these requirements?

- A. Send the order ID to an Amazon Simple Notification Service (Amazon SNS) FIFO topic that fans out to one Amazon Simple Queue Service (Amazon SQS) FIFO queue for inventory management and another SQS FIFO queue for payment processing.
- B. Change the Lambda function that generates the order ID to initiate the Lambda function for inventory management. Then initiate the Lambda function for payment processing.
- C. Send the order ID to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the Lambda functions for inventory management and payment processing to the topic.
- D. Deliver the order ID to an Amazon Simple Queue Service (Amazon SQS) queue. Configure the Lambda functions for inventory management and payment processing to poll the queue.

**Answer: A**

#### **QUESTION 238**

An application uses AWS X-Ray to generate a large amount of trace data on an hourly basis. A developer wants to use filter expressions to limit the returned results through user-specified custom attributes.

How should the developer use filter expressions to filter the results in X-Ray?

- A. Add custom attributes as annotations in the segment document.
- B. Add custom attributes as metadata in the segment document.
- C. Add custom attributes as new segment fields in the segment document.
- D. Create new sampling rules that are based on custom attributes.

**Answer: A**

**Explanation:**

If you add annotations with the X-Ray SDK, you can also filter based on the presence of an annotation key or the value of a key.

<https://docs.aws.amazon.com/xray/latest/devguide/xray-console-filters.html>

#### **QUESTION 239**

A web application is using Amazon Kinesis Data Streams for clickstream data that may not be consumed for up to 12



hours.

How can the developer implement encryption at rest for data within the Kinesis Data Streams?

- A. Enable SSL connections to Kinesis.
- B. Use Amazon Kinesis Consumer Library.
- C. Encrypt the data once it is at rest with a Lambda function.
- D. Enable server-side encryption in Kinesis Data Streams.

**Answer: D**

**Explanation:**

Amazon Kinesis Data Streams is designed for real-time data streaming and allows multiple consumers to process data concurrently and in real-time. It can handle millions of events and provides a scalable and cost-effective solution for handling high-throughput data streams.

#### **QUESTION 240**

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 Firehose
- D. Amazon Kinesis Data Streams

**Answer: D**

**Explanation:**

Amazon Kinesis Data Streams is designed for real-time data streaming and allows multiple consumers to process data concurrently and in real-time. It can handle millions of events and provides a scalable and cost-effective solution for handling high-throughput data streams.

#### **QUESTION 241**

Given the following AWS CloudFormation template:

```
Description: Creates a new Amazon S3 bucket for shared content. Uses a random bucket name to avoid conflicts.

Resources:

  ContentBucket:
    Type: AWS::S3::Bucket
    Outputs:

      ContentBucketName:
        Value: !Ref ContentBucket
```

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

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

**Answer: A**

**Explanation:**

By adding an Export declaration to the Outputs section of the original CloudFormation template, you can make the bucket name available for other templates to import and use. This allows you to reference the bucket name directly in other templates without the need for additional resources or custom logic.

**QUESTION 242**

A developer has built an application that inserts data into an Amazon DynamoDB table. The table is configured to use provisioned capacity. The application is deployed on a burstable nano Amazon EC2 instance. The application logs show that the application has been failing because of a `ProvisionedThroughputExceededException` error. Which actions should the developer take to resolve this issue? (Choose two.)

- A. Move the application to a larger EC2 instance.
- B. Increase the number of read capacity units (RCUs) that are provisioned for the DynamoDB table.
- C. Reduce the frequency of requests to DynamoDB by implementing exponential backoff.
- D. Increase the frequency of requests to DynamoDB by decreasing the retry delay.
- E. Change the capacity mode of the DynamoDB table from provisioned to on-demand.

**Answer:** CE

**QUESTION 243**

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

**Answer:** A

**QUESTION 244**

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

**QUESTION 245**

A developer registered an AWS Lambda function as a target for an Application Load Balancer (ALB) using a CLI command. However, the Lambda function is not being invoked when the client sends requests through the ALB. Why is the Lambda function not being invoked?

- A. A Lambda function cannot be registered as a target for an ALB.
- B. A Lambda function can be registered with an ALB using AWS Management Console only.
- C. The permissions to invoke the Lambda function are missing.
- D. Cross-zone is not enabled on the ALB.

**Answer:** C

**QUESTION 246**

A developer is creating an AWS Lambda function that will connect to an Amazon RDS for MySQL instance. The developer wants to store the database credentials. The database credentials need to be encrypted and the database password needs to be automatically rotated.

Which solution will meet these requirements?

- A. Store the database credentials as environment variables for the Lambda function. Set the environment variables to rotate automatically.
- B. Store the database credentials in AWS Secrets Manager. Set up managed rotation on the database credentials.
- C. Store the database credentials in AWS Systems Manager Parameter Store as secure string parameters. Set up managed rotation on the parameters.
- D. Store the database credentials in the X-Amz-Security-Token parameter. Set up managed rotation on the parameter.

**Answer: B**

**QUESTION 247**

A developer wants to reduce risk when deploying a new version of an existing AWS Lambda function. To test the Lambda function, the developer needs to split the traffic between the existing version and the new version of the Lambda function.

Which solution will meet these requirements?

- A. Configure a weighted routing policy in Amazon Route 53. Associate the versions of the Lambda function with the weighted routing policy.
- B. Create a function alias. Configure the alias to split the traffic between the two versions of the Lambda function.
- C. Create an Application Load Balancer (ALB) that uses the Lambda function as a target. Configure the ALB to split the traffic between the two versions of the Lambda function.
- D. Create the new version of the Lambda function as a Lambda layer on the existing version. Configure the function to split the traffic between the two layers.

**Answer: B**