



Free Questions for **DVA-C02**  
Shared by **Farrell** on **14-03-2023**

For More Free Questions and Preparation Resources

[Check the Links on Last Page](#)



# Question 1

---

Question Type: MultipleChoice

---

A company has developed a new serverless application using AWS Lambda functions that will be deployed using the AWS Serverless Application Model (AWS SAM) CLI.

Which step should the developer complete prior to deploying the application?

Options:

- A- Compress the application to a zip file and upload it into AWS Lambda.
- B- Test the new AWS Lambda function by first tracing it in AWS X-Ray.
- C- Bundle the serverless application using a SAM package.
- D- Create the application environment using the `eb create my-env` command.

Answer:

---

C

Explanation:

---

This step should be completed prior to deploying the application because it prepares the application artifacts for deployment. The AWS Serverless Application Model (AWS SAM) is a framework that simplifies building and deploying serverless applications on AWS. The AWS SAM CLI is a command-line tool that helps you create, test, and deploy serverless applications using AWS SAM templates. The `sam package` command bundles the application artifacts, such as Lambda function code and API definitions, and uploads them to an Amazon S3 bucket. The command also returns a CloudFormation template that is ready to be deployed with the `sam deploy` command. Compressing the application to a zip file and uploading it to AWS Lambda will not work because it does not use AWS SAM templates or CloudFormation. Testing the new Lambda function by first tracing it in AWS X-Ray will not prepare the application for deployment, but only monitor its performance and errors. Creating the application environment using the `eb create my-env` command will not work because it is a command for AWS Elastic Beanstalk, not AWS SAM.

# Question 2

---

Question Type: MultipleChoice

---

A developer is creating a new REST API by using Amazon API Gateway and AWS Lambda

a. The development team tests the API and validates responses for the known use cases before deploying the API to the production environment.

The developer wants to make the REST API available for testing by using API Gateway locally.

Which AWS Serverless Application Model Command Line Interface (AWS SAM CLI) subcommand will meet these requirements?

Options:

- A- Sam local invoke
- B- Sam local generate-event
- C- Sam local start-lambda
- D- Sam local start-api



Answer:

D

Explanation:

The AWS Serverless Application Model Command Line Interface (AWS SAM CLI) is a command-line tool for local development and testing of Serverless applications<sup>2</sup>. The `sam local start-api` subcommand of AWS SAM CLI is used to simulate a REST API by starting a new local endpoint<sup>3</sup>. Therefore, option D is correct.

## Question 3

Question Type: MultipleChoice

A company built an online event platform. For each event, the company organizes quizzes and generates leaderboards that are based on the quiz scores. The company stores the leaderboard data in Amazon DynamoDB and retains the data for 30 days after an event is complete. The company then uses a scheduled job to delete the old leaderboard data.

The DynamoDB table is configured with a fixed write capacity. During the months when many events occur, the DynamoDB write API requests are throttled when the scheduled delete job runs.

A developer must create a long-term solution that deletes the old leaderboard data and optimizes write throughput.

Which solution meets these requirements?

### Options:

---

- A- Configure a TTL attribute for the leaderboard data
- B- Use DynamoDB Streams to schedule and delete the leaderboard data
- C- Use AWS Step Functions to schedule and delete the leaderboard data.
- D- Set a higher write capacity when the scheduled delete job runs

### Answer:

---

A

### Explanation:

---

DynamoDB TTL (Time-to-Live): A native feature that automatically deletes items after a specified expiration time.

Efficiency: Eliminates the need for scheduled deletion jobs, optimizing write throughput by avoiding potential throttling conflicts.

Seamless Integration: TTL works directly within DynamoDB, requiring minimal development overhead.

[DynamoDB TTL](#)

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

## Question 4

---

Question Type: MultipleChoice

---

An AWS Lambda function requires read access to an Amazon S3 bucket and requires read/write access to an Amazon DynamoDB table. The correct IAM policy already exists.

What is the MOST secure way to grant the Lambda function access to the S3 bucket and the DynamoDB table?

### Options:

---

- A- Attach the existing IAM policy to the Lambda function.
- B- Create an IAM role for the Lambda function. Attach the existing IAM policy to the role. Attach the role to the Lambda function.

- C- Create an IAM user with programmatic access Attach the existing IAM policy to the user. Add the user access key ID and secret access key as environment variables in the Lambda function.
- D- Add the AWS account root user access key ID and secret access key as encrypted environment variables in the Lambda function

Answer:

---

B

Explanation:

---

Principle of Least Privilege: Granting specific permissions through an IAM role is more secure than directly attaching policies to a function or using root user credentials.

IAM Roles for Lambda: Designed to provide temporary credentials to Lambda functions, enhancing security.

Reusability: The existing IAM policy ensures the correct S3 and DynamoDB access is granted.

[IAM Roles for Lambda](#)

Documentation: <https://docs.aws.amazon.com/lambda/latest/dg/lambda-intro-execution-role.html>

[IAM Best Practices: https://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/best-practices.html)

## Question 5

---

Question Type: MultipleChoice

---

A company has an application that runs as a series of AWS Lambda functions. Each Lambda function receives data from an Amazon Simple Notification Service (Amazon SNS) topic and writes the data to an Amazon Aurora DB instance.

To comply with an information security policy, the company must ensure that the Lambda functions all use a single securely encrypted database connection string to access Aurora.

Which solution will meet these requirements'?

Options:

---

- A- Use IAM database authentication for Aurora to enable secure database connections for all the Lambda functions.
- B- Store the credentials and read the credentials from an encrypted Amazon RDS DB instance.
- C- Store the credentials in AWS Systems Manager Parameter Store as a secure string parameter.

D- Use Lambda environment variables with a shared AWS Key Management Service (AWS KMS) key for encryption.

Answer:

---

A

Explanation:

---

This solution will meet the requirements by using IAM database authentication for Aurora, which enables using IAM roles or users to authenticate with Aurora databases instead of using passwords or other secrets. The developer can use IAM database authentication for Aurora to enable secure database connections for all the Lambda functions that access Aurora DB instance. The developer can create an IAM role with permission to connect to Aurora DB instance and attach it to each Lambda function. The developer can also configure Aurora DB instance to use IAM database authentication and enable encryption in transit using SSL certificates. This way, the Lambda functions can use a single securely encrypted database connection string to access Aurora without needing any secrets or passwords. Option B is not optimal because it will store the credentials and read them from an encrypted Amazon RDS DB instance, which may introduce additional costs and complexity for managing and accessing another RDS DB instance. Option C is not optimal because it will store the credentials in AWS Systems Manager Parameter Store as a secure string parameter, which may require additional steps or permissions to retrieve and decrypt the credentials from Parameter Store. Option D is not optimal because it will use Lambda environment variables with a shared AWS Key Management Service (AWS KMS) key for encryption, which may not be secure or scalable as environment variables are stored as plain text unless encrypted with AWS KMS.

## Question 6

---

Question Type: MultipleChoice

---

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"?

### Options:

---

- 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

### Explanation:

---

This solution allows the fastest response for the query because it enables the query to use a single partition key value (the Product ID) and a range of sort key values (the Product Rating) to find the matching items. A global secondary index (GSI) is an index that has a partition key and an optional sort key that are different from those on the base table. A GSI can be created at any time and can be queried or scanned independently of the base table. A local secondary index (LSI) is an index that has the same partition key as the base table, but a different sort key. An LSI can only be created when the base table is created and must be queried together with the base table partition key. Using a GSI with Product ID as the partition key and Review ID as the sort key will not allow the query to use a range of sort key values to find the highest ratings. Using an LSI with Product ID as the partition key and Product Rating as the sort key will not work because Product ID is not the partition key of the base table. Using an LSI with Review ID as the partition key and Product ID as the sort key will not allow the query to use a single partition key value to find the matching items.

## Question 7

---

**Question Type:** MultipleChoice

---

A developer needs to deploy an application running on AWS Fargate using Amazon ECS The application has environment variables that must be passed to a container for the application to

initialize.

How should the environment variables be passed to the container?

### Options:

---

- A- Define an array that includes the environment variables under the environment parameter within the service definition.
- B- Define an array that includes the environment variables under the environment parameter within the task definition.
- C- Define an array that includes the environment variables under the entryPoint parameter within the task definition.
- D- Define an array that includes the environment variables under the entryPoint parameter within the service definition.

### Answer:

---

B

### Explanation:

---

This solution allows the environment variables to be passed to the container when it is launched by AWS Fargate using Amazon ECS. The task definition is a text file that describes one or more containers that form an application. It contains various parameters for configuring the containers, such as CPU and memory requirements, network mode, and environment variables. The environment parameter is an array of key-value pairs that specify environment variables to pass to a container. Defining an array that includes the environment variables under the entryPoint parameter within the task definition will not pass them to the container, but use them as command-line arguments for overriding the default entry point of a container. Defining an array that includes the environment variables under the environment or entryPoint parameter within the service definition will not pass them to the container, but cause an error because these parameters are not valid for a service definition.

## Question 8

---

Question Type: MultipleChoice

---

A company is migrating its PostgreSQL database into the AWS Cloud. The company wants to use a database that will secure and regularly rotate database credentials. The company wants a solution that does not require additional programming overhead.



Which solution will meet these requirements?

### Options:

---

- A- Use Amazon Aurora PostgreSQL for the database. Store the database credentials in AWS Systems Manager Parameter Store Turn on rotation.
- B- Use Amazon Aurora PostgreSQL for the database. Store the database credentials in AWS Secrets Manager Turn on rotation.
- C- Use Amazon DynamoDB for the database. Store the database credentials in AWS Systems Manager Parameter Store Turn on rotation.
- D- Use Amazon DynamoDB for the database. Store the database credentials in AWS Secrets Manager Turn on rotation.

### Answer:

---

B

### Explanation:

---

This solution meets the requirements because it uses a PostgreSQL-compatible database that can secure and regularly rotate database credentials without requiring additional programming overhead. Amazon Aurora PostgreSQL is a relational database service that is compatible with PostgreSQL and offers high performance, availability, and scalability. AWS Secrets Manager is a service that helps you protect secrets needed to access your applications, services, and IT resources. You can store database credentials in AWS Secrets Manager and use them to access your Aurora PostgreSQL database. You can also enable automatic rotation of your secrets according to a schedule or an event. AWS Secrets Manager handles the complexity of rotating secrets for you, such as generating new passwords and updating your database with the new credentials. Using Amazon DynamoDB for the database will not meet the requirements because it is a NoSQL database that is not compatible with PostgreSQL. Using AWS Systems Manager Parameter Store for storing and rotating database credentials will require additional programming overhead to integrate with your database.

## Question 9

---

**Question Type:** MultipleChoice

---

A developer is testing a RESTful application that is deployed by using Amazon API Gateway and AWS Lambda. When the developer tests the user login by using credentials that are not valid, the developer receives an HTTP 405 METHOD\_NOT\_ALLOWED error. The developer has verified that the test is sending the correct request for the resource.

Which HTTP error should the application return in response to the request?

Options:

- A- HTTP 401
- B- HTTP 404
- C- HTTP 503
- D- HTTP 505

Answer:

A



Explanation:

HTTP Status Codes: Each HTTP status code has a specific meaning in RESTful APIs.

HTTP 405 (Method Not Allowed): Indicates that the request method (e.g., POST) is not supported for the specified resource.

HTTP 401 (Unauthorized): Represents a failure to authenticate, which is the appropriate response for invalid login credentials.

HTTP Status Codes: <https://developer.mozilla.org/en-US/docs/Web/HTTP/Status>

## Question 10

Question Type: MultipleChoice

A developer has observed an increase in bugs in the AWS Lambda functions that a development team has deployed in its Node.js application.

To minimize these bugs, the developer wants to implement automated testing of Lambda functions in an environment that closely simulates the Lambda environment.

The developer needs to give other developers the ability to run the tests locally. The developer also needs to integrate the tests into the team's continuous integration and continuous delivery (CI/CD) pipeline before the AWS Cloud Development Kit (AWS CDK) deployment.

Which solution will meet these requirements?

## Options:

---

- A-** Create sample events based on the Lambda documentation. Create automated test scripts that use the `cdk local invoke` command to invoke the Lambda functions. Check the response. Document the test scripts for the other developers on the team. Update the CI/CD pipeline to run the test scripts.
- B-** Install a unit testing framework that reproduces the Lambda execution environment. Create sample events based on the Lambda documentation. Invoke the handler function by using a unit testing framework. Check the response. Document how to run the unit testing framework for the other developers on the team. Update the CI/CD pipeline to run the unit testing framework.
- C-** Install the AWS Serverless Application Model (AWS SAM) CLI tool. Use the `sam local generate-event` command to generate sample events for the automated tests. Create automated test scripts that use the `sam local invoke` command to invoke the Lambda functions. Check the response. Document the test scripts for the other developers on the team. Update the CI/CD pipeline to run the test scripts.
- D-** Create sample events based on the Lambda documentation. Create a Docker container from the Node.js base image to invoke the Lambda functions. Check the response. Document how to run the Docker container for the other developers on the team. Update the CI/CD pipeline to run the Docker container.

## Answer:

---

C

## Explanation:

---

The AWS Serverless Application Model Command Line Interface (AWS SAM CLI) is a command-line tool for local development and testing of Serverless applications<sup>3</sup>. The `sam local generate-event` command of AWS SAM CLI generates sample events for automated tests<sup>3</sup>. The `sam local invoke` command is used to invoke Lambda functions<sup>3</sup>. Therefore, option C is correct.



To Get Premium Files for DVA-C02 Visit

<https://www.p2pexams.com/products/dva-c02>

For More Free Questions Visit

<https://www.p2pexams.com/amazon/pdf/dva-c02>

**20%**  
**DISCOUNT**

**P2P**  
exams