

DVA-C02 dumps PDF & DVA-C02 exam guide & DVA-C02 test simulate



DVA-C02
AWS Certified
Developer
Associate

Certification Questions
& Exams Dumps

www.edurely.com

P.S. Free & New DVA-C02 dumps are available on Google Drive shared by ActualtestPDF: <https://drive.google.com/open?id=1XGNQg8aLIQz78IYsxo427NPGsy4vyPyw>

We regard the customer as king so we put a high emphasis on the trust of every users, therefore our security system can protect you both in payment of DVA-C02 guide braindumps and promise that your computer will not be infected during the process of payment on our DVA-C02 Study Materials. Moreover, if you end up the cooperation between us, we have the responsibility to delete your personal information on DVA-C02 exam prep. In a word, We have data protection act for you to avoid information leakage!

As we all know it is not easy and smooth for everyone to obtain the DVA-C02 certification, and especially for those people who cannot make full use of their sporadic time and are not able to study in a productive way. But you are lucky, we can provide you with well-rounded services on DVA-C02 practice DVA-C02 test materials to help you improve ability and come over difficulties when you have trouble studying. We would be very pleased and thankful if you can spare your valuable time to have a look about features of our DVA-C02 study materials.

>> [DVA-C02 Exam Tutorials](#) <<

DVA-C02 Reliable Practice Questions | 100% DVA-C02 Exam Coverage

In order to meet the requirements of our customers, Our DVA-C02 test questions carefully designed the automatic correcting system for customers. It is known to us that practicing the incorrect questions is very important for everyone, so our DVA-C02 exam question provide the automatic correcting system to help customers understand and correct the errors. If you are finding a study material in order to get away from your exam, you can spend little time to know about our DVA-C02 Test Torrent, it must suit for you.

Amazon AWS Certified Developer - Associate Sample Questions (Q411-Q416):

NEW QUESTION # 411

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?

- A. Test the new AWS Lambda function by first tracing it in AWS X-Ray.
- B. Compress the application to a zip file and upload it into AWS Lambda.
- **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.

NEW QUESTION # 412

A company is planning to use AWS CodeDeploy to deploy an application to Amazon Elastic Container Service (Amazon ECS). During the deployment of a new version of the application, the company initially must expose only 10% of live traffic to the new version of the deployed application. Then, after 15 minutes elapse, the company must route all the remaining live traffic to the new version of the deployed application.

Which CodeDeploy predefined configuration will meet these requirements?

- A. `CodeDeployDefault_LambdaCanary10Percent5Minutes`
- B. `CodeDeployDefault_LambdaCanary10Percent15Minutes`
- C. `CodeDeployDefault_ECSECSLinear10PercentEvery1_Minutes`
- **D. `CodeDeployDefault_ECSCanary10Percent15Minutes`**

Answer: D

Explanation:

* CodeDeploy Predefined Configurations: CodeDeploy offers built-in deployment configurations for common scenarios.

* Canary Deployment: Canary deployments gradually shift traffic to a new version, ideal for controlled rollouts like this requirement.

* `CodeDeployDefault_ECSCanary10Percent15Minutes`: This configuration matches the company's requirements, shifting 10% of traffic initially and then completing the rollout after 15 minutes.

References:

* AWS CodeDeploy Deployment Configurations: <https://docs.aws.amazon.com/codedeploy/latest/userguide/deployment-configurations-create.html>

NEW QUESTION # 413

A developer deployed an application to an Amazon EC2 instance. The application needs to know the public IPv4 address of the instance. How can the application find this information?

- **A. Query the instance metadata from `http://169.254.169.254/latest/meta-data/`.**
- B. Check the hosts file of the operating system.
- C. Query the Amazon Machine Image (AMI) information from `http://169.254.169.254/latest/meta-data/ami/`.
- D. Query the instance user data from `http://169.254.169.254/latest/user-data/`.

Answer: A

Explanation:

* Instance Metadata Service: EC2 instances have access to an internal metadata service. It provides instance-specific information like instance ID, security groups, and public IP address.

* Accessing Metadata:

* Make an HTTP GET request to the base URL: <http://169.254.169.254/latest/meta-data/>

* You'll get a list of available categories. The public IPv4 address is under public-ipv4.

References:

* Instance Metadata and User Data: <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/instancedata-data-retrieval.html>

NEW QUESTION # 414

A developer is creating a serverless application that uses an AWS Lambda function. The developer will use AWS CloudFormation to deploy the application. The application will write logs to Amazon CloudWatch Logs. The developer has created a log group in a CloudFormation template for the application to use. The developer needs to modify the CloudFormation template to make the name of the log group available to the application at runtime. Which solution will meet this requirement?

- A. Pass the log group's name to the application in the user data section of the CloudFormation template.
- B. Use the `AWS::Include` transform in CloudFormation to provide the log group's name to the application.
- C. Use the CloudFormation template's Mappings section to specify the log group's name for the application.
- **D. Pass the log group's Amazon Resource Name (ARN) as an environment variable to the Lambda function.**

Answer: D

Explanation:

* CloudFormation and Lambda Environment Variables:

* CloudFormation is an excellent tool to manage infrastructure as code, including the log group resource.

* Lambda functions can access environment variables at runtime, making them a suitable way to pass configuration information like the log group ARN.

* CloudFormation Template Modification:

* In your CloudFormation template, define the log group resource.

* In the Lambda function resource, add an Environment section:

YAML

Environment:

Variables:

```
LOG_GROUP_ARN: !Ref LogGroupResourceName
```

Use code with caution.

content_copy

* The `!Ref` intrinsic function retrieves the log group's ARN, which CloudFormation generates during stack creation.

* Using the ARN in Your Lambda Function:

* Within your Lambda code, access the `LOG_GROUP_ARN` environment variable.

* Configure your logging library (e.g., Python's logging module) to send logs to the specified log group.

NEW QUESTION # 415

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

There is a requirement to uniquely identify users at any scale.

Which solution will meet these requirements?

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

Answer: A

Explanation:

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

P.S. Free 2025 Amazon DVA-C02 dumps are available on Google Drive shared by ActualtestPDF: <https://drive.google.com/open?id=1XGNQg8aLIQz78IYsxo427NPGsy4vyPyw>