

KillTest

高品質 より良いサービス



練習問題

<https://www.killtest.jp>

1年間無料アップデート

**Exam : AWS Certified Developer
- Associate 2018**

**Title : AWS Certified Big Data -
Specialty**

Version : DEMO

1.A Developer created a dashboard for an application using Amazon API Gateway, Amazon S3, AWS Lambda, and Amazon RDS. The Developer needs an authentication mechanism allowing a user to sign in and view the dashboard. It must be accessible from mobile applications, desktops, and tablets, and must remember user preferences across platforms.

Which AWS service should the Developer use to support this authentication scenario?

- A. AWS KMS
- B. Amazon Cognito
- C. AWS Directory Service
- D. Amazon IAM

Answer: D

2.A Developer has created an S3 bucket s3://mycoolapp and has enabled server access logging that points to the folder s3://mycoolapp/logs. The Developer moved 100 KB of Cascading Style Sheets (CSS) documents to the folder s3://mycoolapp/css, and then stopped work. When the developer came back a few days later, the bucket was 50 GB.

What is the MOST likely cause of this situation?

- A. The CSS files were not compressed and S3 versioning was enabled.
- B. S3 replication was enabled on the bucket.
- C. Logging into the same bucket caused exponential log growth.
- D. An S3 lifecycle policy has moved the entire CSS file to S3 Infrequent Access.

Answer: B

3.A Developer is creating an Auto Scaling group whose instances need to publish a custom metric to Amazon CloudWatch.

Which method would be the MOST secure way to authenticate a CloudWatch PUT request?

- A. Create an IAM user with PutMetricData permission and put the user credentials in a private repository; have applications pull the credentials as needed.
- B. Create an IAM user with PutMetricData permission, and modify the Auto Scaling launch configuration to inject the user credentials into the instance user data.
- C. Modify the CloudWatch metric policies to allow the PutMetricData permission to instances from the Auto Scaling group.
- D. Create an IAM role with PutMetricData permission and modify the Auto Scaling launching configuration to launch instances using that role.

Answer: D

4.A Developer is working on an application that tracks hundreds of millions of product reviews in an Amazon DynamoDB table.

The records include the data elements shown in the table:

Name	Type	Description
reviewID	Number	16 digit UUID
starRating	Number	Integer 1-5 of user rating
comment	String	User comment string
productID	Number	Product ID being reviewed

Which field, when used as the partition key, would result in the MOST consistent performance using

DynamoDB?

- A. starRating
- B. reviewID
- C. comment
- D. productID

Answer: B

5.A Developer has written a serverless application using multiple AWS services. The business logic is written as a Lambda function which has dependencies on third-party libraries. The Lambda function endpoints will be exposed using Amazon API Gateway. The Lambda function will write the information to Amazon DynamoDB. The Developer is ready to deploy the application but must have the ability to rollback.

How can this deployment be automated, based on these requirements?

- A. Deploy using Amazon Lambda API operations to create the Lambda function by providing a deployment package.
- B. Use an AWS CloudFormation template and use CloudFormation syntax to define the Lambda function resource in the template.
- C. Use syntax conforming to the Serverless Application Model in the AWS CloudFormation template to define the Lambda function resource.
- D. Create a bash script which uses AWS CLI to package and deploy the application.

Answer: A

Explanation:

Reference: <https://docs.aws.amazon.com/lambda/latest/dg/automating-deployment.html>