DOP-C01^{Q&As}

AWS Certified DevOps Engineer - Professional (DOP-C01)

Pass Amazon DOP-C01 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.passapply.com/aws-devops-engineer-professional.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon
Official Exam Center

🛠️ **Instant Download** After Purchase

🛠️ **100% Money Back** Guarantee

🛠️ **365 Days** Free Update

🛠️ **800,000+** Satisfied Customers

QUESTION 1

You have an application running a specific process that is critical to the application\'s functionality, and have added the health check process to your Auto Scaling Group. The instances are showing healthy but the application itself is not working as it should. What could be the issue with the health check, since it is still showing the instances as healthy.

A. You do not have the time range in the health check properly configured

B. It is not possible for a health check to monitor a process that involves the application

C. The health check is not configured properly

D. The health check is not checking the application process

Correct Answer: D

If you have custom health checks, you can send the information from your health checks to Auto Scaling so that Auto Scaling can use this information. For example, if you determine that an instance is not functioning as expected, you can set the health status of the instance to Unhealthy. The next time that Auto Scaling performs a health check on the instance, it will determine that the instance is unhealthy and then launch a replacement instance.

QUESTION 2

A root owner is trying to create an IAM user of the various departments. The owner has created groups for each department, but wants to still delineate the user based on the sub division level. E.g. The two users from different sub departments should be identified separately and have separate permissions. How can the root owner configure this?

A. Create a hierarchy of the IAM users which are separated based on the department

B. Create a nested group

C. Use the paths to separate the users of the same group

D. It is not possible to delineate within a group

Correct Answer: C

Explanation: The path functionality within an IAM group and user allows them to delineate by further levels. In this case the user needs to use the path with each user or group so that the ARN of the user will look similar to:
arn:aws:iam::123456789012:user/division_abc/subdivision_xyz/user1
arn:aws:iam::123456789012:user/division_abc/subdivision_xyz/user2

Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/Using_Identifiers.html#Identifiers_ARNs

QUESTION 3

A company has developed a static website hosted on an Amazon S3 bucket. The website is deployed using AWS CloudFormation. The CloudFormation template defines an S3 bucket and a custom resource that copies content into the

bucket from a source location. The company has decided that it needs to move the website to a new location, so the

existing CloudFormation stack must be deleted and re-created. However, CloudFormation reports that the stack could not be

deleted cleanly.

What is the MOST likely cause and how can the DevOps Engineer mitigate this problem for this and future versions of the website?

A. Deletion has failed because the S3 bucket has an active website configuration. Modify the CloudFormation template to remove the WebsiteConfiguration property from the S3 bucket resource.

B. Deletion has failed because the S3 bucket is not empty. Modify the custom resource\\'s AWS Lambda function code to recursively empty the bucket when RequestType is Delete.

C. Deletion has failed because the custom resource does not define a deletion policy. Add a DeletionPolicy property to the custom resource definition with a value of RemoveOnDeletion.

D. Deletion has failed because the S3 bucket is not empty. Modify the S3 bucket resource in the CloudFormation template to add a DeletionPolicy property with a value of Empty.

Correct Answer: B

QUESTION 4

A company runs a database on a single Amazon EC2 instance in a development environment. The data is stored on separate Amazon EBS volumes that are attached to the EC2 instance. An Amazon Route 53 A record has been created and configured to point to the EC2 instance. The company would like to automate the recovery of the database instance when an instance or Availability Zone (AZ) fails. The company also wants to keep its costs low. The RTO is 4 hours and the RPO is 12 hours.

Which solution should a DevOps Engineer implement to meet these requirements?

A. Run the database in an Auto Scaling group with a minimum and maximum instance count of 1 in multiple AZs. Add a lifecycle hook to the Auto Scaling group and define an Amazon CloudWatch Events rule that is triggered when a lifecycle event occurs. Have the CloudWatch Events rule invoke an AWS Lambda function to detach or attach the Amazon EBS data volumes from the EC2 instance based on the event. Configure the EC2 instance UserData to mount the data volumes (retry on failure with a short delay), then start the database and update the Route 53 record.

B. Run the database on two separate EC2 instances in different AZs with one active and the other as a standby. Attach the data volumes to the active instance. Configure an Amazon CloudWatch Events rule to invoke an AWS Lambda function on EC2 instance termination. The Lambda function launches a replacement EC2 instance. If the terminated instance was the active node, then the function attaches the data volumes to the standby node. Start the database and update the Route 53 record.

C. Run the database in an Auto Scaling group with a minimum and maximum instance count of 1 in multiple AZs. Create an AWS Lambda function that is triggered by a scheduled Amazon CloudWatch Events rule every 4 hours to take a snapshot of the data volume and apply a tag. Have the instance UserData get the latest snapshot, create a new volume from it, and attach and mount the volume. Then start the database and update the Route 53 record.

D. Run the database on two separate EC2 instances in different AZs. Configure one of the instances as a master and the other as a standby. Set up replication between the master and standby instances. Point the Route 53 record to the master. Configure an Amazon CloudWatch Events rule to invoke an AWS Lambda function upon the EC2 instance termination. The Lambda function launches a replacement EC2 instance. If the terminated instance was the active node, the function promotes the standby to master and points the Route 53 record to it.

Correct Answer: D

QUESTION 5

You want to pass queue messages that are 1GB each. How should you achieve this?

A. Use Kinesis as a buffer stream for message bodies. Store the checkpoint id for the placement in the Kinesis Stream in SQS.

B. Use the Amazon SQS Extended Client Library for Java and Amazon S3 as a storage mechanism for message bodies.

C. Use SQS\\'s support for message partitioning and multi-part uploads on Amazon S3.

D. Use AWS EFS as a shared pool storage medium. Store filesystem pointers to the files on disk in the SQS message bodies.

Correct Answer: B

You can manage Amazon SQS messages with Amazon S3. This is especially useful for storing and retrieving messages with a message size of up to 2 GB. To manage Amazon SQS messages with Amazon S3, use the Amazon SQS

Extended Client Library for Java.

Reference:

http://docs.aws.amazon.com/AWSSimpleQueueService/latest/SQSDeveloperGuide/s3-messages.html

[DOP-C01 Study Guide](#)   [DOP-C01 Exam Questions](#)   [DOP-C01 Braindumps](#)