# 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.pass4itsure.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**

What is the correct syntax for the AWS command to create a single region trail?

A. aws create-trail --name trailname --s3-object objectname

B. aws cloudtrail --s3-regionname IPaddress create-trail --name trailname

C. aws cloudtrail create-trail --name trailname --s3-bucket-name bucketname

D. aws cloudtrail create-trail --name trailname --s3-portnumber IPaddress

Correct Answer: C

The command aws cloudtrail create-trail --name trailname --s3-bucket-name bucketname will create a single region trail. You must create a S3 bucket before you execute the command, with proper CloudTrail permissions applied to it (and you must have the AWS command line tools (CLI) on your system).

Reference: http://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-create-and-update-a-trailby-using-the-aws-cli.html

**QUESTION 2**

You need your API backed by DynamoDB to stay online during a total regional AWS failure. You can tolerate a couple minutes of lag or slowness during a large failure event, but the system should recover with normal operation after those

few minutes.

What is a good approach?

A. Set up DynamoDB cross-region replication in a master-standby configuration, with a single standby in another region. Create an Auto Scaling Group behind an ELB in each of the two regions DynamoDB is running in. Add a Route53 Latency DNS Record with DNS Failover, using the ELBs in the two regions as the resource records.

B. Set up a DynamoDB Multi-Region table. Create an Auto Scaling Group behind an ELB in each of the two regions DynamoDB is running in. Add a Route53 Latency DNS Record with DNS Failover, using the ELBs in the two regions as the resource records.

C. Set up a DynamoDB Multi-Region table. Create a cross-region ELB pointing to a cross-region Auto Scaling Group, and direct a Route53 Latency DNS Record with DNS Failover to the cross-region ELB.

D. Set up DynamoDB cross-region replication in a master-standby configuration, with a single standby in another region. Create a cross-region ELB pointing to a cross-region Auto Scaling Group, and direct a Route53 Latency DNS Record with DNS Failover to the cross-region ELB.

Correct Answer: A

There is no such thing as a cross-region ELB, nor such thing as a cross-region Auto Scaling Group, nor such thing as a DynamoDB Multi-Region Table. The only option that makes sense is the cross-regional replication version with two ELBs

and ASGs with Route53 Failover and Latency DNS.

Reference:

http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Streams.CrossRegionRepl.html

**QUESTION 3**

A DevOps Engineer must track the health of a stateless RESTful service sitting behind a Classic ILoad Balancer. The deployment of new application revisions is through a CI/CD pipeline. If the service\\'s latency increases beyond a defined threshold, deployment should be stopped until the service has recovered.

Which of the following methods allow for the QUICKEST detection time?

A. Use Amazon CloudWatch metrics provided by Elastic Load Balancing to calculate average latency. Alarm and stop deployment when latency increases beyond the defined threshold.

B. Use AWS Lambda and Elastic Load Balancing access logs to detect average latency. Alarm and stop deployment when latency increases beyond the defined threshold.

C. Use AWS CodeDeploy\\'s MinimumHealthyHosts setting to define thresholds for rolling back deployments. If these thresholds are breached, roll back the deployment.

D. Use Metric Filters to parse application logs in Amazon CloudWatch Logs. Create a filter for latency. Alarm and stop deployment when latency increases beyond the defined threshold.

Correct Answer: C

**QUESTION 4**

What are the default memory limit policies for a Docker container?

A. Limited memory, limited kernel memory

B. Unlimited memory, limited kernel memory

C. Limited memory, unlimited kernel memory

D. Unlimited memory, unlimited kernel memory

Correct Answer: D

Kernel memory limits are expressed in terms of the overall memory allocated to a container. Consider the following scenarios: Unlimited memory, unlimited kernel memory: This is the default behavior. Unlimited memory, limited kernel memory: This is appropriate when the amount of memory needed by all cgroups is greater than the amount of memory that actually exists on the host machine. You can configure the kernel memory to never go over what is available on the host machine, and containers which need more memory need to wait for it. Limited memory, umlimited kernel memory: The overall memory is limited, but the kernel memory is not. Limited memory, limited kernel memory: Limiting both user and kernel memory can be useful for debugging memory-related problems. If a container is using an unexpected amount of either type of memory, it will run out of memory without affecting other containers or the host machine. Within this setting, if the kernel memory limit is lower than the user memory limit, running out of kernel memory will cause the container to experience an OOM error. If the kernel memory limit is higher than the user memory limit, the kernel limit will not cause the container to experience an OOM.

Reference: https://docs.docker.com/engine/admin/resource_constraints/#--kernel-memory-details

**QUESTION 5**

A company is building a solution for storing files containing Personally Identifiable Information (PII) on AWS.

Requirements state:

1.

All data must be encrypted at rest and in transit.

2.

All data must be replicated in at least two locations that are at least 500 miles apart.

Which solution meets these requirements?

A. Create primary and secondary Amazon S3 buckets in two separate Availability Zones that are at least 500 miles apart. Use a bucket policy to enforce access to the buckets only through HTTPS. Use a bucket policy to enforce Amazon S3 SSE-C on all objects uploaded to the bucket. Configure cross- region replication between the two buckets.

B. Create primary and secondary Amazon S3 buckets in two separate AWS Regions that are at least 500 miles apart. Use a bucket policy to enforce access to the buckets only through HTTPS. Use a bucket policy to enforce S3-Managed Keys (SSE-S3) on all objects uploaded to the bucket. Configure cross- region replication between the two buckets.

C. Create primary and secondary Amazon S3 buckets in two separate AWS Regions that are at least 500 miles apart. Use an IAM role to enforce access to the buckets only through HTTPS. Use a bucket policy to enforce Amazon S3Managed Keys (SSE-S3) on all objects uploaded to the bucket. Configure cross-region replication between the two buckets.

D. Create primary and secondary Amazon S3 buckets in two separate Availability Zones that are at least 500 miles apart. Use a bucket policy to enforce access to the buckets only through HTTPS. Use a bucket policy to enforce AWS KMS encryption on all objects uploaded to the bucket. Configure cross- region replication between the two buckets. Create a KMS Customer Master Key (CMK) in the primary region for encrypting objects.

Correct Answer: B