

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.pass2lead.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

The management team at a company with a large on-premises OpenStack environment wants to move non-production workloads to AWS. An AWS Direct Connect connection has been provisioned and configured to connect the environments. Due to contractual obligations, the production workloads must remain on-premises, and will be moved to AWS after the next contract negotiation. The company follows Center for Internet Security (CIS) standards for hardening

images; this configuration was developed using the company's configuration management system.

Which solution will automatically create an identical image in the AWS environment without significant overhead?

- A. Write an AWS CloudFormation template that will create an Amazon EC2 instance. Use cloud-init to install the configuration management agent, use cfn-wait to wait for configuration management to successfully apply, and use an AWS Lambda-backed custom resource to create the AMI.
- B. Log in to the console, launch an Amazon EC2 instance, and install the configuration management agent. When changes are applied through the configuration management system, log in to the console and create a new AMI from the instance.
- C. Create a new AWS OpsWorks layer and mirror the image hardening standards. Use this layer as the baseline for all AWS workloads.
- D. When a change is made in the configuration management system, a job in Jenkins is triggered to use the VM Import command to create an Amazon EC2 instance in the Amazon VPC. Use lifecycle hooks to launch an AWS Lambda function to create the AMI.

Correct Answer: D

QUESTION 2

A company is hosting a web application in an AWS Region. For disaster recovery purposes, a second region is being used as a standby. Disaster recovery requirements state that session data must be replicated between regions in near-real

time and 1% of requests should route to the secondary region to continuously verify system functionality. Additionally, if there is a disruption in service in the main region, traffic should be automatically routed to the secondary region, and the

secondary region must be able to scale up to handle all traffic.

How should a DevOps Engineer meet these requirements?

- A. In both regions, deploy the application on AWS Elastic Beanstalk and use Amazon DynamoDB global tables for session data. Use an Amazon Route 53 weighted routing policy with health checks to distribute the traffic across the regions.
- B. In both regions, launch the application in Auto Scaling groups and use DynamoDB for session data. Use a Route 53 failover routing policy with health checks to distribute the traffic across the regions.
- C. In both regions, deploy the application in AWS Lambda, exposed by Amazon API Gateway, and use Amazon RDS PostgreSQL with cross-region replication for session data. Deploy the web application with client-side logic to call the

API Gateway directly.

D. In both regions, launch the application in Auto Scaling groups and use DynamoDB global tables for session data. Enable an Amazon CloudFront weighted distribution across regions. Point the Amazon Route 53 DNS record at the CloudFront distribution.

Correct Answer: A

QUESTION 3

A company wants to use Amazon DynamoDB for maintaining metadata on its forums. See the sample data set in the image below.

Thread

ForumName	Subject	LastPostDateTime	Thread
"S3"	"aaa"	"2015-03-15:17:24:31"	12
"S3"	"bbb"	"2015-01-22:23:18:01"	3
"S3"	"ccc"	"2015-02-31:13:14:21"	4
"S3"	"ddd"	"2015-01-03:09:21:11"	9
"EC2"	"yyy"	"2015-02-12:11:07:56"	18
"EC2"	"zzz"	"2015-01-18:07:33:42"	0
"RDS"	"ttt"	"2015-01-19:01:13:24"	3
"RDS"	"sss"	"2015-03-11:06:53:00"	11
"RDS"	"ttt"	"2015-10-22:12:19:44"	5

A DevOps Engineer is required to define the table schema with the partition key, the sort key, the local secondary index, projected attributes, and fetch operations. The schema should support the following example searches using the least provisioned read capacity units to minimize cost. -Search within ForumName for items where the subject starts with `a\\`. -Search forums within the given LastPostDateTime time frame. -Return the thread value where LastPostDateTime is within the last three months.

Which schema meets the requirements?

A. Use Subject as the primary key and ForumName as the sort key. Have LSI with LastPostDateTime as the sort key

and fetch operations for thread.

B. Use ForumName as the primary key and Subject as the sort key. Have LSI with LastPostDateTime as the sort key and the projected attribute thread.

C. Use ForumName as the primary key and Subject as the sort key. Have LSI with Thread as the sort key and the projected attribute LastPostDateTime.

D. Use Subject as the primary key and ForumName as the sort key. Have LSI with Thread as the sort key and fetch operations for LastPostDateTime.

Correct Answer: A

QUESTION 4

A company is using AWS CodeCommit as its source code repository. After an internal audit, the compliance team mandates that any code change that go into the master branch must be committed by senior developers.

Which solution will meet these requirements?

A. Create two repositories in CodeCommit: one for working and another for the master. Create separate IAM groups for senior developers and developers. Assign the resource-level permissions on the repositories tied to the IAM groups. After the code changes are reviewed, sync the approved files to the master code commit repository.

B. Create a repository in CodeCommit. Create separate IAM groups for senior developers and developers. Assign code commit permissions for both groups, with code merge permissions for the senior developers group. Create a trigger to notify senior developers with a URL link to approve or deny commit requests delivered through Amazon SNS. Once a senior developer approves the code, the code gets merged to the master branch.

C. Create a repository in CodeCommit with a working and master branch. Create separate IAM groups for senior developers and developers. Use an IAM policy to assign each IAM group their corresponding branches. Once the code is merged to the working branch, senior developers can pull the changes from the working branch to the master branch.

D. Create a repository in CodeCommit. Create separate IAM groups for senior developers and developers. Use AWS Lambda triggers on the master branch and get the user name of the developer at the event object of the Lambda function. Validate the user name with the IAM group to approve or deny the commit.

Correct Answer: A

QUESTION 5

You have a web application that is currently running on a collection of micro instance types in a single AZ behind a single load balancer. You have an Auto Scaling group configured to scale from 2 to 64 instances. When reviewing your CloudWatch metrics, you see that sometimes your Auto Scaling group is running 64 micro instances. The web application is reading and writing to a DynamoDB-configured backend and configured with 800 Write Capacity Units and 800 Read Capacity Units. Your customers are complaining that they are experiencing long load times when viewing your website. You have investigated the DynamoDB CloudWatch metrics; you are under the provisioned Read and write Capacity Units and there is no throttling. How do you scale your service to improve the load times and ensure the principles of high availability?

A. Change your Auto Scaling group configuration to include multiple AZs.

B. Change your Auto Scaling group configuration to include multiple AZs, and increase the number of Read Capacity

Units in your DynamoDB table by a factor of three, because you will need to be calling DynarnODB from three AZs.

C. Add a second load balancer to your Auto Scaling group so that you can support more inbound connections per second.

D. Change your Auto Scaling group configuration to use larger instances and include multiple AZ's instead of one.

Correct Answer: D

[DOP-C01 Practice Test](#)

[DOP-C01 Exam Questions](#)

[DOP-C01 Braindumps](#)