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

For AWS CloudFormation, which stack state refuses UpdateStack calls?

A. UPDATE_ROLLBACK_FAILED

B. UPDATE_ROLLBACK_COMPLETE

C. UPDATE_COMPLETE

D. CREATE_COMPLETE

Correct Answer: A

When a stack is in the UPDATE_ROLLBACK_FAILED state, you can continue rolling it back to return it to a working state (to UPDATE_ROLLBACK_COMPLETE). You cannot update a stack that is in the UPDATE_ROLLBACK_FAILED state. However, if you can continue to roll it back, you can return the stack to its original settings and try to update it again.

Reference: http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-cfn-updating- stackscontinueupdaterollback.html

QUESTION 2

You have just come from your Chief Information Security Officer\\'s (CISO) office with the instructions to provide an audit report of all AWS network rules used by the organization\\'s Amazon EC2 instances. You have discovered that a single Describe-Security-Groups API call will return all of an account\\'s security groups and rules within a region. You create the following pseudo-code to create the required report:

-Parse "aws ec2 describe-security-groups" output

-For each security group

-Create report of ingress and egress rules

Which two additional pieces of logic should you include to meet the CISO\\'s requirements? (Choose two.)

A. Parse security groups in each region.

B. Parse security groups in each Availability Zone and region.

C. Evaluate VPC network access control lists.

D. Evaluate AWS CloudTrail logs.

E. Evaluate Elastic Load Balancing access control lists.

F. Parse CloudFront access control lists.

Correct Answer: AC

QUESTION 3

A company is using AWS CodeDeploy to automate software deployment. The deployment must meet these requirements:

1.

A number of instances must be available to serve traffic during the deployment. Traffic must be balanced across those instances, and the instances must automatically heal in the event of failure.

2.

A new fleet of instances must be launched for deploying a new revision automatically, with no manual provisioning.

3.

Traffic must be rerouted to the new environment to half of the new instances at a time. The deployment should succeed if traffic is rerouted to at least half of the instances; otherwise, it should fail.

4.

Before routing traffic to the new fleet of instances, the temporary files generated during the deployment process must be deleted.

5.

At the end of a successful deployment, the original instances in the deployment group must be deleted immediately to reduce costs.

How can a DevOps Engineer meet these requirements?

A. Use an Application Load Balancer and an in-place deployment. Associate the Auto Scaling group with the deployment group. Use the Automatically copy Auto Scaling group option, and use CodeDeployDefault.OneAtAtime as the deployment configuration. Instruct AWS CodeDeploy to terminate the original instances in the deployment group, and use the AllowTraffic hook within appspec.yml to delete the temporary files.

B. Use an Application Load Balancer and a blue/green deployment. Associate the Auto Scaling group and the Application Load Balancer target group with the deployment group. Use the Automatically copy Auto Scaling group option, create a custom deployment configuration with minimum healthy hosts defined as 50%, and assign the configuration to the deployment group. Instruct AWS CodeDeploy to terminate the original instances in the deployment group, and use the BeforeBlockTraffic hook within appspec.yml to delete the temporary files.

C. Use an Application Load Balancer and a blue/green deployment. Associate the Auto Scaling group and the Application Load Balancer target group with the deployment group. Use the Automatically copy Auto Scaling group option, and use CodeDeployDefault.HalfAtAtime as the deployment configuration. Instruct AWS CodeDeploy to terminate the original instances in the deployment group, and use the BeforeAllowTraffic hook within appspec.yml to delete the temporary files.

D. Use an Application Load Balancer and an in-place deployment. Associate the Auto Scaling group and Application Load Balancer target group with the deployment group. Use the Automatically copy Auto Scaling group option, and use CodeDeployDefault.AllatOnce as a deployment configuration. Instruct AWS CodeDeploy to terminate the original instances in the deployment group, and use the BlockTraffic hook within appspec.yml to delete the temporary files.

Correct Answer: C

Reference:

https://docs.aws.amazon.com/codedeploy/latest/APIReference/API_BlueGreenDeploymentConfig uration.html

QUESTION 4

What is server immutability?

A. Not updating a server after creation.

B. The ability to change server counts.

C. Updating a server after creation.

D. The inability to change server counts.

Correct Answer: A

... disposable upgrades offer a simpler way to know if your application has unknown dependencies. The underlying EC2 instance usage is considered temporary or ephemeral in nature for the period of deployment until the current release is active. During the new release, a new set of EC2 instances are rolled out by terminating older instances. This type of upgrade technique is more common in an immutable infrastructure. Reference: https://d0.awsstatic.com/whitepapers/overview-of-deployment-options-on-aws.pdf

QUESTION 5

Which of the following is NOT an advantage of Docker\\'s content addressable storage model?

A. random UUIDs improve filesystem performance

B. improved security

C. guarantees data integrity after push, pull, load, and save operations

D. avoids content ID collisions

Correct Answer: A

Docker 1.10 introduced a new content addressable storage model. This is a completely new way to address image and layer data on disk. Previously, image and layer data was referenced and stored using a randomly generated UUID. In the new model this is replaced by a secure content hash. The new model improves security, provides a built-in way to avoid ID collisions, and guarantees data integrity after pull, push, load, and save operations. It also enables better sharing of layers by allowing many images to freely share their layers even if they did not come from the same build.

Reference: https://docs.docker.com/engine/userguide/storagedriver/imagesandcontainers/#contentaddressable-storage