



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

Your DevOps team is responsible for a multi-tier, Windows-based web application consisting of web servers, Amazon RDS database instances, and a load balancer behind Amazon Route53. You have been asked by your manager to build a cost-effective rolling deployment solution for this web application. What method should you use?

- A. Re-deploy your application on an AWS OpsWorks stack. Use the AWS OpsWorks done stack feature to allow updates between duplicate stacks.
- B. Re-deploy your application on Elastic Beanstalk and take advantage of Elastic BeanStalk rolling updates.
- C. Re-deploy your application using an AWS CloudFormation template, launch a new AWS CloudFormation stack during each deployment, and then tear down the old stack.
- D. Re-deploy your application using an AWS CloudFormation template. Use AWS CloudFormation rolling deployment policies, create a new policy for your AWS CloudFormation stack, and initiate an update stack operation to deploy new code.

Correct Answer: D

QUESTION 2

You are administering a continuous integration application that polls version control for changes and then launches new Amazon EC2 instances for a full suite of build tests. What should you do to ensure the lowest overall cost while being able to run as many tests in parallel as possible?

- A. Perform syntax checking on the continuous integration system before launching a new Amazon EC2 instance for build test, unit and integration tests.
- B. Perform syntax and build tests on the continuous integration system before launching the new Amazon EC2 instance unit and integration tests.
- C. Perform all tests on the continuous integration system, using AWS OpsWorks for unit, integration, and build tests.
- D. Perform syntax checking on the continuous integration system before launching a new AWS Data Pipeline for coordinating the output of unit, integration, and build tests.

Correct Answer: B

QUESTION 3

A company wants to use Amazon ECS to provide a Docker container runtime environment. For compliance reasons, all Amazon EBS volumes used in the ECS cluster must be encrypted. Rolling updates will be made to the cluster instances

and the company wants the instances drained of all tasks before being terminated.

How can these requirements be met? (Choose two.)

- A. Modify the default ECS AMI user data to create a script that executes docker rm

Correct Answer: CD

QUESTION 4

If designing a single playbook to run across multiple Linux distributions that have distribution specific commands, what would be the best method to allow a successful run?

- A. Enable fact gathering and use the ``when`` conditional to match the distribution to the task.
- B. This is not possible, a separate playbook for each target Linux distribution is required.
- C. Use ``ignore_errors: true`` in the tasks.
- D. Use the ``shell`` module to write your own checks for each command that is ran.

Correct Answer: A

Ansible provides a method to only run a task when a condition is met using the ``when`` declarative. With gather facts enabled, the play has access to the distribution name of the Linux system, thus, tasks can be tailored to a specific distribution and ran only when the condition is met, e.g.: ``- when: ansible_os_family == "Debian"``.

Reference: http://docs.ansible.com/ansible/playbooks_conditionals.html

QUESTION 5

Your current log analysis application takes more than four hours to generate a report of the top 10 users of your web application. You have been asked to implement a system that can report this information in real time, ensure that the report is always up to date, and handle increases in the number of requests to your web application. Choose the option that is cost-effective and can fulfill the requirements.

- A. Publish your data to CloudWatch Logs, and configure your application to autoscale to handle the load on demand.
- B. Publish your log data to an Amazon S3 bucket. Use AWS CloudFormation to create an Auto Scaling group to scale your post-processing application which is configured to pull down your log files stored on Amazon S3.
- C. Post your log data to an Amazon Kinesis data stream, and subscribe your log-processing application so that is configured to process your logging data.
- D. Configure an Auto Scaling group to increase the size of your Amazon EMR cluster.
- E. Create a multi-AZ Amazon RDS MySQL cluster, post the logging data to MySQL, and run a map reduce job to retrieve the required information on user counts.

Correct Answer: C

QUESTION 6

You need to know when you spend \$1000 or more on AWS. What's the easy way for you to see that notification?

- A. AWS CloudWatch Events tied to API calls, when certain thresholds are exceeded, publish to SNS.



- B. Scrape the billing page periodically and pump into Kinesis.
- C. AWS CloudWatch Metrics + Billing Alarm + Lambda event subscription. When a threshold is exceeded, email the manager.
- D. Scrape the billing page periodically and publish to SNS.

Correct Answer: C

Explanation: Even if you're careful to stay within the free tier, it's a good idea to create a billing alarm to notify you if you exceed the limits of the free tier. Billing alarms can help to protect you against unknowingly accruing charges if you inadvertently use a service outside of the free tier or if traffic exceeds your expectations.

Reference: <http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/freetier-alarms.html>

QUESTION 7

A company is using an AWS CloudFormation template to deploy web applications. The template requires that manual changes be made for each of the three major environments: production, staging, and development. The current sprint includes the new implementation and configuration of AWS CodePipeline for automated deployments.

What changes should the DevOps Engineer make to ensure that the CloudFormation template is reusable across multiple pipelines?

- A. Use a CloudFormation custom resource to query the status of the CodePipeline to determine which environment is launched. Dynamically alter the launch configuration of the Amazon EC2 instances.
- B. Set up a CodePipeline pipeline for each environment to use input parameters. Use CloudFormation mappings to switch associated UserData for the Amazon EC2 instances to match the environment being launched.
- C. Set up a CodePipeline pipeline that has multiple stages, one for each development environment. Use AWS Lambda functions to trigger CloudFormation deployments to dynamically alter the UserData of the Amazon EC2 instances launched in each environment.
- D. Use CloudFormation input parameters to dynamically alter the LaunchConfiguration and UserData sections of each Amazon EC2 instance every time the CloudFormation stack is updated.

Correct Answer: C

QUESTION 8

Which of these is not a CloudFormation Helper Script?

- A. cfn-signal
- B. cfn-hup
- C. cfn-request
- D. cfn-get-metadata

Correct Answer: C



This is the complete list of CloudFormation Helper Scripts: cfn-init, cfn-signal, cfn-get-metadata, cfn-hup

Reference: <http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/cfn-helper-scriptsreference.html>

QUESTION 9

A DevOps engineer has automated a web service deployment using AWS CodePipeline with the following steps:

1.

An AWS CodeBuild project compiles the deployment artifact and runs unit tests.

2.

An AWS CodeDeploy deployment group deploys the web service to Amazon EC2 instances in the staging environment.

3.

A CodeDeploy deployment group deploys the web service to EC2 instances in the production environment.

The quality assurance (QA) team has asked for permission to inspect the build artifact before the deployment to the production environment occurs. The QA team wants to run an internal automated penetration testing tool (invoked using a REST API call) to run some manual tests.

Which combination of actions will fulfill this request? (Choose two.)

- A. Insert a manual approval action between the test and deployment actions of the pipeline.
- B. Modify the buildspec.yml file for the compilation stage to require manual approval before completion.
- C. Update the CodeDeploy deployment group so it requires manual approval to proceed.
- D. Update the pipeline to directly trigger the REST API for the automated penetration testing tool.
- E. Update the pipeline to invoke a Lambda function that triggers the REST API for the automated penetration testing tool.

Correct Answer: BC

Reference: <https://docs.aws.amazon.com/codebuild/latest/userguide/sample-codedeploy.html>

QUESTION 10

When thinking of DynamoDB, what are true of Local Secondary Key properties?

- A. Either the partition key or the sort key can be different from the table, but not both.
- B. Only the sort key can be different from the table.
- C. The partition key and sort key can be different from the table.
- D. Only the partition key can be different from the table.

Correct Answer: B

Global secondary index ?an index with a partition key and a sort key that can be different from those on the table. A global secondary index is considered "global" because queries on the index can span all of the data in a table, across all partitions.

Reference: <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/SecondaryIndexes.html>

QUESTION 11

An Engineering team manages a Node.js e-commerce application. The current environment consists of the following components:

1.
Amazon S3 buckets for storing content
2.
Amazon EC2 for the front-end web servers
3.
AWS Lambda for executing image processing
4.
Amazon DynamoDB for storing session-related data

The team expects a significant increase in traffic to the site. The application should handle the additional load without interruption. The team ran initial tests by adding new servers to the EC2 front-end to handle the larger load, but the

instances took up to 20 minutes to become fully configured. The team wants to reduce this configuration time.

What changes will the Engineering team need to implement to make the solution the MOST resilient and highly available while meeting the expected increase in demand?

- A. Use AWS OpsWorks to automatically configure each new EC2 instance as it is launched. Configure the EC2 instances by using an Auto Scaling group behind an Application Load Balancer across multiple Availability Zones. Implement Amazon DynamoDB Auto Scaling. Use Amazon Route 53 to point the application DNS record to the Application Load Balancer.
- B. Deploy a fleet of EC2 instances, doubling the current capacity, and place them behind an Application Load Balancer. Increase the Amazon DynamoDB read and write capacity units. Add an alias record that contains the Application Load Balancer endpoint to the existing Amazon Route 53 DNS record that points to the application.
- C. Configure Amazon CloudFront and have its origin point to Amazon S3 to host the web application. Implement Amazon DynamoDB Auto Scaling. Use Amazon Route 53 to point the application DNS record to the CloudFront DNS name.
- D. Use AWS Elastic Beanstalk with a custom AMI including all web components. Deploy the platform by using an Auto Scaling group behind an Application Load Balancer across multiple Availability Zones. Implement Amazon DynamoDB Auto Scaling. Use Amazon Route 53 to point the application DNS record to the Elastic Beanstalk load balancer.

Correct Answer: D



QUESTION 12

Your application uses Amazon SQS and Auto Scaling to process background jobs. The Auto Scaling policy is based on the number of messages in the queue, with a maximum Instance count of 100. Since the application was launched, the group has never scaled above 50. The Auto Scaling group has now scaled to 100, the queue size is increasing, and very few Jobs are being completed. The number of messages being sent to the queue is at normal levels.

What should you do to identify why the queue size is unusually high, and to reduce it?

- A. Temporarily increase the Auto Scaling group's desired value to 200. When the queue size has been reduced, reduce it to 50.
- B. Analyze the application logs to identify possible reasons for message processing failure and resolve the cause for failures.
- C. Create additional Auto Scaling groups, enabling the processing of the queue to be performed in parallel.
- D. Analyze CloudTrail logs for Amazon SQS to ensure that the instances' Amazon EC2 role has permission to receive messages from the queue.

Correct Answer: B

QUESTION 13

A company is migrating its public-facing software to AWS. The company plans to use Amazon EC2 to run application code and Amazon RDS to store all application data. The company wants to primarily use one Region with failover capabilities to a secondary Region and Amazon Route 53 to route traffic. The RPO is 2 hours and the RTO is 4 hours.

Which combination of steps should be used to meet these requirements while MINIMIZING cost? (Choose three.)

- A. Create an AWS CloudFormation template to provision the application server and database instance in a single Region.
- B. Create an AWS CloudFormation template to provision the application tier of the application and a multi-Region database instance.
- C. Configure Amazon CloudWatch Events rules to run every hour. Trigger AWS Lambda functions to create an RDS snapshot and copy it to the secondary Region.
- D. Configure Amazon CloudWatch Events rules to run every 3 hours. Trigger AWS Lambda functions to create an RDS snapshot and copy it to the secondary Region.
- E. In the event of a failure, deploy a new AWS CloudFormation stack in a secondary region to provision the application resources and a new RDS instance using the copied snapshot and a Route 53 failover routing policy.
- F. In the event of a failure, deploy a new AWS CloudFormation stack in a secondary region to provision the application resources and a replica of the RDS database using the copied snapshot and a Route 53 latency-based routing policy.

Correct Answer: BDE



QUESTION 14

A company runs a three-tier web application in its production environment, which is built on a single AWS CloudFormation template made up of Amazon EC2 instances behind an ELB Application Load Balancer. The instances run in an EC2

Auto Scaling group across multiple Availability Zones. Data is stored in an Amazon RDS Multi-AZ DB instance with read replicas. Amazon Route 53 manages the application's public DNS record.

A DevOps Engineer must create a workflow to mitigate a failed software deployment by rolling back changes in the production environment when a software cutover occurs for new application software. What steps should the Engineer perform to meet these requirements with the LEAST amount of downtime?

- A. Use CloudFormation to deploy an additional staging environment and configure the Route 53 DNS with weighted records. During cutover, change the Route 53 A record weights to achieve an even traffic distribution between the two environments. Validate the traffic in the new environment and immediately terminate the old environment if tests are successful.
- B. Use a single AWS Elastic Beanstalk environment to deploy the staging and production environments. Update the environment by uploading the ZIP file with the new application code. Swap the Elastic Beanstalk environment CNAME. Validate the traffic in the new environment and immediately terminate the old environment if tests are successful.
- C. Use a single AWS Elastic Beanstalk environment and an AWS OpsWorks environment to deploy the staging and production environments. Update the environment by uploading the ZIP file with the new application code into the Elastic Beanstalk environment deployed with the OpsWorks stack. Validate the traffic in the new environment and immediately terminate the old environment if tests are successful.
- D. Use AWS CloudFormation to deploy an additional staging environment, and configure the Route 53 DNS with weighted records. During cutover, increase the weight distribution to have more traffic directed to the new staging environment as workloads are successfully validated. Keep the old production environment in place until the new staging environment handles all traffic.

Correct Answer: D

QUESTION 15

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

[DOP-C01 VCE Dumps](#)

[DOP-C01 Study Guide](#)

[DOP-C01 Braindumps](#)