



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

A company uses a complex system that consists of networking, IAM policies, and multiple three-tier applications. Requirements are still being defined for a new system, so the number of AWS components present in the final design is not known. The DevOps Engineer needs to begin defining AWS resources using AWS CloudFormation to automate and version-control the new infrastructure. What is the best practice for using CloudFormation to create new environments?

- A. Manually construct the networking layer using Amazon VPC and then define all other resources using CloudFormation.
- B. Create a single template to encompass all resources that are required for the system so there is only one template to version-control.
- C. Create multiple separate templates for each logical part of the system, use cross-stack references in CloudFormation, and maintain several templates in version control.
- D. Create many separate templates for each logical part of the system, and provide the outputs from one to the next using an Amazon EC2 instance running SDK for granular control.

Correct Answer: C

QUESTION 2

For auditing, analytics, and troubleshooting purposes, a DevOps Engineer for a data analytics application needs to collect all of the application and Linux system logs from the Amazon EC2 instances before termination. The company, on

average, runs 10,000 instances in an Auto Scaling group. The company requires the ability to quickly find logs based on instance IDs and date ranges.

Which is the MOST cost-effective solution?

- A. Create an EC2 Instance-terminate Lifecycle Action on the group, write a termination script for pushing logs into Amazon S3, and trigger an AWS Lambda function based on S3 PUT to create a catalog of log files in an Amazon DynamoDB table with the primary key being Instance ID and sort key being Instance Termination Date.
- B. Create an EC2 Instance-terminate Lifecycle Action on the group, write a termination script for pushing logs into Amazon CloudWatch Logs, create a CloudWatch Events rule to trigger an AWS Lambda function to create a catalog of log files in an Amazon DynamoDB table with the primary key being Instance ID and sort key being Instance Termination Date.
- C. Create an EC2 Instance-terminate Lifecycle Action on the group, create an Amazon CloudWatch Events rule based on it to trigger an AWS Lambda function for storing the logs in Amazon S3, and create a catalog of log files in an Amazon DynamoDB table with the primary key being Instance ID and sort key being Instance Termination Date.
- D. Create an EC2 Instance-terminate Lifecycle Action on the group, push the logs into Amazon Kinesis Data Firehose, and select Amazon ES as the destination for providing storage and search capability.

Correct Answer: C

QUESTION 3



Your Company wants to perform A/B testing on a new website feature for 20 percent of its users. The website uses CloudFront for whole site delivery, with some content cached for up to 24 hours. How do you enable this testing for the required proportion of users while minimizing performance impact?

- A. Configure the web servers to handle two domain names. The feature is switched on or off depending on which domain name is used for a request. Configure a CloudFront origin for each domain name, and configure the CloudFront distribution to use one origin for 20 percent of users and the other origin for the other 80 percent.
- B. Configure the CloudFront distribution to forward a cookie specific to this feature. For requests where the cookie is not set, the web servers set its value to "on" for 20 percent of responses and "off" for 80 percent. For requests where the cookie is set, the web servers use its value to determine whether the feature should be on or off for the response.
- C. Create a second stack of web servers that host the website with the feature on. Using Amazon Route53, create two resource record sets with the same name: one with a weighting of "1" and a value of this new stack; the other a weighting of "4" and a value of the existing stack. Use the resource record set's name as the CloudFront distribution's origin.
- D. Invalidate all of the CloudFront distribution's cache items that the feature affects. On future requests, the web servers create responses with the feature on for 20 percent of users, and off for 80 percent. The web servers set "Cache-Control: no-cache" on all of these responses.

Correct Answer: B

QUESTION 4

A legacy web application stores access logs in a proprietary text format. One of the security requirements is to search application access events and correlate them with access data from many different systems.

These searches should be near-real time.

Which solution offloads the processing load on the application server and provides a mechanism to search the data in near-real time?

- A. Install the Amazon CloudWatch Logs agent on the application server and use CloudWatch Events rules to search logs for access events. Use Amazon CloudSearch as an interface to search for events.
- B. Use the third-party file-input plugin Logstash to monitor the application log file, then use a custom dissect filter on the agent to parse the log entries into the JSON format. Output the events to Amazon ES to be searched. Use the Elasticsearch API for querying the data.
- C. Upload the log files to Amazon S3 by using the S3 sync command. Use Amazon Athena to define the structure of the data as a table, with Athena SQL queries to search for access events.
- D. Install the Amazon Kinesis Agent on the application server, configure it to monitor the log files, and send it to a Kinesis stream. Configure Kinesis to transform the data by using an AWS Lambda function, and forward events to Amazon ES for analysis. Use the Elasticsearch API for querying the data.

Correct Answer: D

QUESTION 5

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

[DOP-C01 Practice Test](#)

[DOP-C01 Study Guide](#)

[DOP-C01 Braindumps](#)