



# DOP-C02<sup>Q&As</sup>

AWS Certified DevOps Engineer - Professional

## Pass Amazon DOP-C02 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.passapply.com/dop-c02.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 growing company manages more than 50 accounts in an organization in AWS Organizations. The company has configured its applications to send logs to Amazon CloudWatch Logs.

A DevOps engineer needs to aggregate logs so that the company can quickly search the logs to respond to future security incidents. The DevOps engineer has created a new AWS account for centralized monitoring.

Which combination of steps should the DevOps engineer take to make the application logs searchable from the monitoring account? (Select THREE.)

- A. In the monitoring account, download an AWS CloudFormation template from CloudWatch to use in Organizations. Use CloudFormation StackSets in the organization's management account to deploy the CloudFormation template to the entire organization.
- B. Create an AWS CloudFormation template that defines an IAM role. Configure the role to allow logs-amazonaws.com to perform the logs:Link action if the aws:ResourceAccount property is equal to the monitoring account ID. Use CloudFormation StackSets in the organization's management account to deploy the CloudFormation template to the entire organization.
- C. Create an IAM role in the monitoring account. Attach a trust policy that allows logs.amazonaws.com to perform the iam:CreateSink action if the aws:PrincipalOrgId property is equal to the organization ID.
- D. In the organization's management account, enable the logging policies for the organization.
- E. use CloudWatch Observability Access Manager in the monitoring account to create a sink. Allow logs to be shared with the monitoring account. Configure the monitoring account data selection to view the Observability data from the organization ID.
- F. In the monitoring account, attach the CloudWatchLogsReadOnlyAccess AWS managed policy to an IAM role that can be assumed to search the logs.

Correct Answer: BCF

To aggregate logs from multiple accounts in an organization, the DevOps engineer needs to create a cross-account subscription that allows the monitoring account to receive log events from the sharing accounts.

To enable cross-account subscription, the DevOps engineer needs to create an IAM role in each sharing account that grants permission to CloudWatch Logs to link the log groups to the destination in the monitoring account. This can be done

using a CloudFormation template and StackSets to deploy the role to all accounts in the organization.

The DevOps engineer also needs to create an IAM role in the monitoring account that allows CloudWatch Logs to create a sink for receiving log events from other accounts. The role must have a trust policy that specifies the organization ID

as a condition.

Finally, the DevOps engineer needs to attach the CloudWatchLogsReadOnlyAccess policy to an IAM role in the monitoring account that can be used to search the logs from the cross-account subscription.

References:

1: Cross-account log data sharing with subscriptions



- 2: Create an IAM role for CloudWatch Logs in each sharing account
  - 3: AWS CloudFormation StackSets
  - 4: Create an IAM role for CloudWatch Logs in your monitoring account
  - 5: CloudWatchLogsReadOnlyAccess policy
- 

## QUESTION 2

A company's DevOps engineer is creating an AWS Lambda function to process notifications from an Amazon Simple Notification Service (Amazon SNS) topic. The Lambda function will process the notification messages and will write the contents of the notification messages to an Amazon RDS Multi-AZ DB instance.

During testing a database administrator accidentally shut down the DB instance. While the database was down the company lost several of the SNS notification messages that were delivered during that time.

The DevOps engineer needs to prevent the loss of notification messages in the future

Which solutions will meet this requirement? (Select TWO.)

- A. Replace the RDS Multi-AZ DB instance with an Amazon DynamoDB table.
- B. Configure an Amazon Simple Queue Service (Amazon SQS) queue as a destination of the Lambda function.
- C. Configure an Amazon Simple Queue Service (Amazon SQS) dead-letter queue for the SNS topic.
- D. Subscribe an Amazon Simple Queue Service (Amazon SQS) queue to the SNS topic. Configure the Lambda function to process messages from the SQS queue.
- E. Replace the SNS topic with an Amazon EventBridge event bus. Configure an EventBridge rule on the new event bus to invoke the Lambda function for each event.

Correct Answer: CD

These solutions will meet the requirement because they will prevent the loss of notification messages in the future. An Amazon SQS queue is a service that provides a reliable, scalable, and secure message queue for asynchronous communication between distributed components. You can use an SQS queue to buffer messages from an SNS topic and ensure that they are delivered and processed by a Lambda function, even if the function or the database is temporarily unavailable. Option C will configure an SQS dead-letter queue for the SNS topic. A dead-letter queue is a queue that receives messages that could not be delivered to any subscriber after a specified number of retries. You can use a dead-letter queue to store and analyze failed messages, or to reprocess them later. This way, you can avoid losing messages that could not be delivered to the Lambda function due to network errors, throttling, or other issues. Option D will subscribe an SQS queue to the SNS topic and configure the Lambda function to process messages from the SQS queue. This will decouple the SNS topic from the Lambda function and provide more flexibility and control over the message delivery and processing. You can use an SQS queue to store messages from the SNS topic until they are ready to be processed by the Lambda function, and also to retry processing in case of failures. This way, you can avoid losing messages that could not be processed by the Lambda function due to database errors, timeouts, or other issues.

---

## QUESTION 3

A company is implementing a well-architected design for its globally accessible API stack. The design needs to ensure both high reliability and fast response times for users located in North America and Europe.



The API stack contains the following three tiers:

Amazon API Gateway

AWS Lambda

Amazon DynamoDB

Which solution will meet the requirements?

- A. Configure Amazon Route 53 to point to API Gateway APIs in North America and Europe using health checks. Configure the APIs to forward requests to a Lambda function in that Region. Configure the Lambda functions to retrieve and update the data in a DynamoDB table in the same Region as the Lambda function.
- B. Configure Amazon Route 53 to point to API Gateway APIs in North America and Europe using latency-based routing and health checks. Configure the APIs to forward requests to a Lambda function in that Region. Configure the Lambda functions to retrieve and update the data in a DynamoDB global table.
- C. Configure Amazon Route 53 to point to API Gateway in North America, create a disaster recovery API in Europe, and configure both APIs to forward requests to the Lambda functions in that Region. Retrieve the data from a DynamoDB global table. Deploy a Lambda function to check the North America API health every 5 minutes. In the event of a failure, update Route 53 to point to the disaster recovery API.
- D. Configure Amazon Route 53 to point to API Gateway API in North America using latency-based routing. Configure the API to forward requests to the Lambda function in the Region nearest to the user. Configure the Lambda function to retrieve and update the data in a DynamoDB table.

Correct Answer: B

---

#### QUESTION 4

What are the bare minimum requirements for a valid Ansible playbook?

- A. The hosts, connection type, fact gathering, vars and tasks.
- B. The hosts declaration and tasks
- C. A YAML file with a single line containing `---\`.
- D. At least one play with at least a hosts declaration

Correct Answer: D

Ansible Playbooks are a series of plays and must contain at a minimum, one play. A play generally consists of hosts to run on, a list of tasks, variables and roles, and any additional instructions, such as connection type, fact gathering, remote

username, etc. that the tasks will need to complete. The only requirement for a valid play is to declare the hosts.

Reference:

[http://docs.ansible.com/ansible/playbooks\\_intro.html](http://docs.ansible.com/ansible/playbooks_intro.html)

---

**QUESTION 5**

What storage driver does Docker generally recommend that you use if it is available?

- A. zfs
- B. btrfs
- C. aufs
- D. overlay

Correct Answer: C

After you have read the storage driver overview, the next step is to choose the best storage driver for your workloads. In making this decision, there are three high-level factors to consider: If multiple storage drivers are supported in your kernel, Docker has a prioritized list of which storage driver to use if no storage driver is explicitly configured, assuming that the prerequisites for that storage driver are met: If aufs is available, default to it, because it is the oldest storage driver. However, it is not universally available.

Reference: <https://docs.docker.com/engine/userguide/storagedriver/selectadriver/>

[DOP-C02 PDF Dumps](#)

[DOP-C02 Study Guide](#)

[DOP-C02 Exam Questions](#)