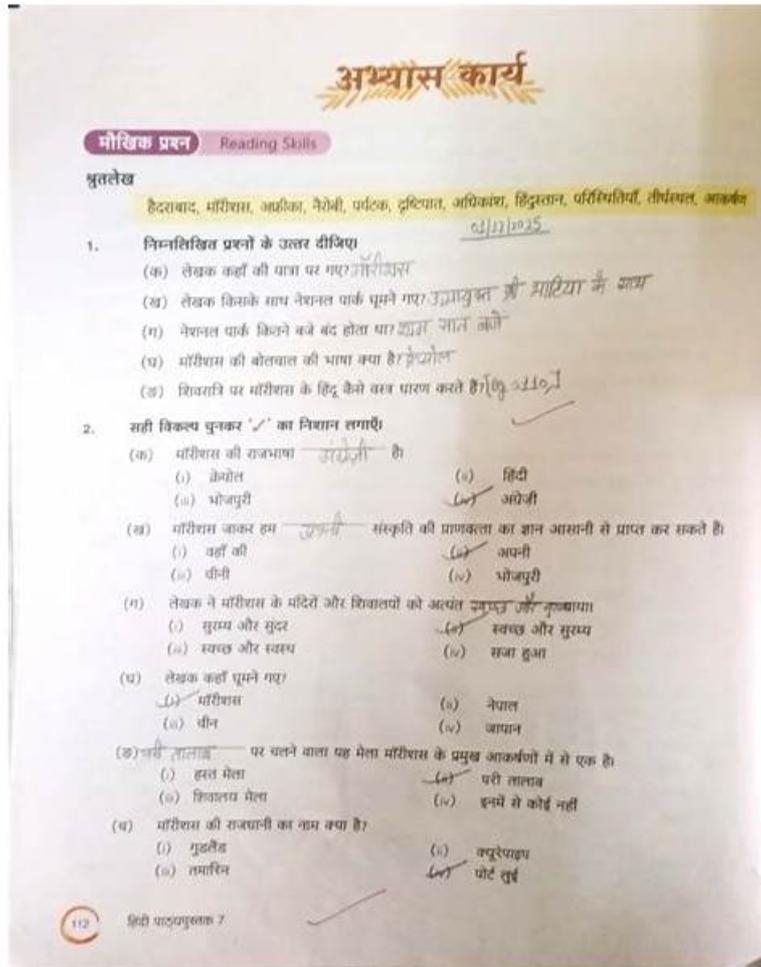


# Latest DOP-C02 Dumps Questions | DOP-C02 Learning Mode



Scanned with OKEN Scanner

P.S. Free 2026 Amazon DOP-C02 dumps are available on Google Drive shared by Actual4test: [https://drive.google.com/open?id=1Pd2F\\_lrJZ5v7XABmTYYwWUVZQFDFKFiq](https://drive.google.com/open?id=1Pd2F_lrJZ5v7XABmTYYwWUVZQFDFKFiq)

Our DOP-C02 exam questions have a lot of advantages. First, our DOP-C02 practice materials are reasonably priced with accessible prices that everyone can afford. Second, they are well-known in this line so their quality and accuracy is unquestionable that everyone trusts with confidence. Third, our DOP-C02 Study Guide is highly efficient that you have great possibility pass the exam within a week based on regular practice attached with the newest information.

The DOP-C02 Exam covers a broad range of topics related to DevOps, including continuous integration and delivery, infrastructure as code, monitoring and logging, security and compliance, and automation and optimization of AWS services. To pass the exam, candidates must demonstrate their ability to design and implement scalable, reliable, and secure DevOps solutions using AWS technologies and best practices. AWS Certified DevOps Engineer - Professional certification is highly valued by employers and can help DevOps professionals advance their careers and increase their earning potential.

>> Latest DOP-C02 Dumps Questions <<

## DOP-C02 Learning Mode, New DOP-C02 Test Pass4sure

A generally accepted view on society is only the professionals engaged in professionally work, and so on, only professional in accordance with professional standards of study materials, as our AWS Certified DevOps Engineer - Professional study questions, to bring more professional quality service for the user. Our study materials can give the user confidence and strongly rely on feeling,

lets the user in the reference appendix not alone on the road, because we are to accompany the examinee on DOP-C02 Exam, candidates need to not only learning content of teaching, but also share his arduous difficult helper, so believe us, we are so professional company.

## Amazon AWS Certified DevOps Engineer - Professional Sample Questions (Q121-Q126):

### NEW QUESTION # 121

A company runs applications in AWS accounts that are in an organization in AWS Organizations. The applications use Amazon EC2 instances and Amazon S3.

The company wants to detect potentially compromised EC2 instances, suspicious network activity, and unusual API activity in its existing AWS accounts and in any AWS accounts that the company creates in the future. When the company detects one of these events, the company wants to use an existing Amazon Simple Notification Service (Amazon SNS) topic to send a notification to its operational support team for investigation and remediation.

Which solution will meet these requirements in accordance with AWS best practices?

- A. In the organization's management account, create an AWS CloudTrail organization trail. Activate the organization trail in all AWS accounts in the organization. Create an SCP that enables VPC Flow Logs in each account in the organization. Configure AWS Security Hub for the organization. Create an Amazon EventBridge rule with an event pattern to match Security Hub events and to forward matching events to the SNS topic.
- B. In the organization's management account, configure an AWS account as the Amazon GuardDuty administrator account. In the GuardDuty administrator account, add the company's existing AWS accounts to GuardDuty as members. In the GuardDuty administrator account, create an Amazon EventBridge rule with an event pattern to match GuardDuty events and to forward matching events to the SNS topic.
- C. In the organization's management account, configure an AWS account as the AWS CloudTrail administrator account. In the CloudTrail administrator account, create a CloudTrail organization trail. Add the company's existing AWS accounts to the organization trail. Create an SCP that enables VPC Flow Logs in each account in the organization. Configure AWS Security Hub for the organization. Create an Amazon EventBridge rule with an event pattern to match Security Hub events and to forward matching events to the SNS topic.
- D. In the organization's management account, configure Amazon GuardDuty to add newly created AWS accounts by invitation and to send invitations to the existing AWS accounts. Create an AWS CloudFormation stack set that accepts the GuardDuty invitation and creates an Amazon EventBridge rule. Configure the rule with an event pattern to match GuardDuty events and to forward matching events to the SNS topic. Configure the CloudFormation stack set to deploy into all AWS accounts in the organization.

**Answer: D**

Explanation:

Explanation:

It allows the company to detect potentially compromised EC2 instances, suspicious network activity, and unusual API activity in its existing AWS accounts and in any AWS accounts that the company creates in the future using Amazon GuardDuty. It also provides a solution for automatically adding future AWS accounts to GuardDuty by configuring GuardDuty to add newly created AWS accounts by invitation and to send invitations to the existing AWS accounts.

### NEW QUESTION # 122

A company recently launched multiple applications that use Application Load Balancers. Application response time often slows down when the applications experience problems. A DevOps engineer needs to implement a monitoring solution that alerts the company when the applications begin to perform slowly. The DevOps engineer creates an Amazon Simple Notification Service (Amazon SNS) topic and subscribes the company's email address to the topic. What should the DevOps engineer do next to meet the requirements?

- A. Create an Amazon CloudWatch alarm that uses the AWS/ApplicationELB namespace RequestCountPerTarget metric. Configure the CloudWatch alarm to send a notification when the number of connections becomes greater than the configured number of threads that the application supports. Configure the CloudWatch alarm to use the SNS topic.
- B. Create an Amazon EventBridge rule that invokes an AWS Lambda function to query the applications on a 5-minute interval. Configure the Lambda function to publish a notification to the SNS topic when the applications return errors.
- C. Create an Amazon CloudWatch Synthetic canary that runs a custom script to query the applications on a 5-minute interval. Configure the canary to use the SNS topic when the applications return errors.
- D. Create an Amazon CloudWatch alarm that uses the AWS/ApplicationELB namespace RequestCountPerTarget metric.

Configure the CloudWatch alarm to send a notification when the average response time becomes greater than the longest response time that the application supports Configure the CloudWatch alarm to use the SNS topic

**Answer: C**

Explanation:

Option A is incorrect because creating an Amazon EventBridge rule that invokes an AWS Lambda function to query the applications on a 5-minute interval is not a valid solution. EventBridge rules can only trigger Lambda functions based on events, not on time intervals. Moreover, querying the applications on a 5-minute interval might incur unnecessary costs and network overhead, and might not detect performance issues in real time.

Option B is correct because creating an Amazon CloudWatch Synthetics canary that runs a custom script to query the applications on a 5-minute interval is a valid solution. CloudWatch Synthetics canaries are configurable scripts that monitor endpoints and APIs by simulating customer behavior. Canaries can run as often as once per minute, and can measure the latency and availability of the applications. Canaries can also send notifications to an Amazon SNS topic when they detect errors or performance issues<sup>1</sup>.

Option C is incorrect because creating an Amazon CloudWatch alarm that uses the AWS/ApplicationELB namespace RequestCountPerTarget metric is not a valid solution. The RequestCountPerTarget metric measures the number of requests completed or connections made per target in a target group<sup>2</sup>. This metric does not reflect the application response time, which is the requirement. Moreover, configuring the CloudWatch alarm to send a notification when the number of connections becomes greater than the configured number of threads that the application supports is not a valid way to measure the application performance, as it depends on the application design and implementation.

Option D is incorrect because creating an Amazon CloudWatch alarm that uses the AWS/ApplicationELB namespace RequestCountPerTarget metric is not a valid solution, for the same reason as option C. The RequestCountPerTarget metric does not reflect the application response time, which is the requirement. Moreover, configuring the CloudWatch alarm to send a notification when the average response time becomes greater than the longest response time that the application supports is not a valid way to measure the application performance, as it does not account for variability or outliers in the response time distribution.

References:

1: Using synthetic monitoring

2: Application Load Balancer metrics

### NEW QUESTION # 123

A company's application uses a fleet of Amazon EC2 On-Demand Instances to analyze and process data. The EC2 instances are in an Auto Scaling group. The Auto Scaling group is a target group for an Application Load Balancer (ALB). The application analyzes critical data that cannot tolerate interruption. The application also analyzes noncritical data that can withstand interruption. The critical data analysis requires quick scalability in response to real-time application demand. The noncritical data analysis involves memory consumption. A DevOps engineer must implement a solution that reduces scale-out latency for the critical data. The solution also must process the noncritical data.

Which combination of steps will meet these requirements? (Select TWO.)

- A. For the critical data, modify the existing Auto Scaling group. Create a warm pool instance in the stopped state. Define the warm pool size. Create a new version of the launch template that has detailed monitoring enabled. Use Spot Instances.
- B. For the noncritical data, create a second Auto Scaling group. Choose the predefined memory utilization metric type for the target tracking scaling policy. Use Spot Instances. Add the new Auto Scaling group as the target group for the ALB. Modify the application to use two target groups for critical data and noncritical data.
- C. For the noncritical data, create a second Auto Scaling group that uses a launch template. Configure the launch template to install the unified Amazon CloudWatch agent and to configure the CloudWatch agent with a custom memory utilization metric. Use Spot Instances. Add the new Auto Scaling group as the target group for the ALB. Modify the application to use two target groups for critical data and noncritical data.
- D. For the critical data, modify the existing Auto Scaling group. Create a warm pool instance in the stopped state. Define the warm pool size. Create a new version of the launch template that has detailed monitoring enabled. Use On-Demand Instances.
- E. For the critical data, modify the existing Auto Scaling group. Create a lifecycle hook to ensure that bootstrap scripts are completed successfully. Ensure that the application on the instances is ready to accept traffic before the instances are registered. Create a new version of the launch template that has detailed monitoring enabled.

**Answer: C,D**

Explanation:

Explanation

For the critical data, using a warm pool can reduce the scale-out latency by having pre-initialized EC2 instances ready to serve the application traffic. Using On-Demand Instances can ensure that the instances are always available and not interrupted by Spot

interruptions2.

For the noncritical data, using a second Auto Scaling group with Spot Instances can reduce the cost and leverage the unused capacity of EC2. Using a launch template with the CloudWatch agent<sup>4</sup> can enable the collection of memory utilization metrics, which can be used to scale the group based on the memory demand. Adding the second group as a target group for the ALB and modifying the application to use two target groups can enable routing the traffic based on the data type.

References: 1: Warm pools for Amazon EC2 Auto Scaling 2: Amazon EC2 On-Demand Capacity Reservations 3: Amazon EC2 Spot Instances 4: Metrics collected by the CloudWatch agent

#### NEW QUESTION # 124

A DevOps engineer uses AWS WAF to manage web ACLs across an AWS account. The DevOps engineer must ensure that AWS WAF is enabled for all Application Load Balancers (ALBs) in the account. The DevOps engineer uses an AWS CloudFormation template to deploy an individual ALB and AWS WAF as part of each application stack's deployment process. If AWS WAF is removed from the ALB after the ALB is deployed, AWS WAF must be added to the ALB automatically. Which solution will meet these requirements with the MOST operational efficiency?

- A. Configure an Amazon EventBridge rule to periodically call an AWS Lambda function that calls the detect-stack-drift API on the CloudFormation template. Configure the Lambda function to modify the ALB attributes with `waf.fail_open.enabled` set to true if the `AWS::WAFv2::WebACLAssociation` resource shows a status of drifted.
- B. Configure an Amazon EventBridge rule to periodically call an AWS Lambda function that calls the detect-stack-drift API on the CloudFormation template. Configure the Lambda function to delete and redeploy the CloudFormation stack if the `AWS::WAFv2::WebACLAssociation` resource shows a status of drifted.
- C. Enable AWS Config. Add the `alb-waf-enabled managed rule`. Create an AWS Systems Manager Automation document to add AWS WAF to an ALB. Edit the rule to automatically remediate. Select the Systems Manager Automation document as the remediation action.
- D. Enable AWS Config. Add the `alb-waf-enabled managed rule`. Create an Amazon EventBridge rule to send all AWS Config ConfigurationItemChangeNotification notification types to an AWS Lambda function. Configure the Lambda function to call the AWS Config `start-resource-evaluation` API in detective mode.

**Answer: C**

Explanation:

Comprehensive and Detailed Explanation From Exact Extract:

AWS Config has a managed rule called `alb-waf-enabled` that checks whether AWS WAF is enabled on ALBs. AWS Config supports automatic remediation actions that can be triggered when noncompliance is detected.

By creating a Systems Manager Automation document that adds AWS WAF to the ALB and associating it as the remediation action for the AWS Config rule, the system can automatically detect and remediate any removal of AWS WAF from ALBs without manual intervention.

This is the most operationally efficient and reliable approach to ensure continuous compliance.

Option B lacks automatic remediation. Options C and D rely on drift detection and Lambda, which add complexity and risk downtime during stack replacement.

Reference:

AWS Config Managed Rules:

"The `alb-waf-enabled` rule checks for AWS WAF association with ALBs and supports automatic remediation using Systems Manager Automation." (AWS Config Managed Rules) AWS Config Remediation:

"AWS Config automatic remediation can invoke Systems Manager Automation documents to remediate noncompliance." (AWS Config Remediation)

#### NEW QUESTION # 125

A DevOps team manages an API running on-premises that serves as a backend for an Amazon API Gateway endpoint. Customers have been complaining about high response latencies, which the development team has verified using the API Gateway latency metrics in Amazon CloudWatch. To identify the cause, the team needs to collect relevant data without introducing additional latency. Which actions should be taken to accomplish this? (Choose two.)

- A. Modify the on-premises application to calculate and upload statistical data relevant to the API service requests to CloudWatch metrics.
- B. Modify the on-premises application to send log information back to API Gateway with each request.
- C. Install the CloudWatch agent server side and configure the agent to upload relevant logs to CloudWatch.
- D. Enable AWS X-Ray tracing in API Gateway, modify the application to capture request segments, and upload those

