

Free PDF Quiz Amazon - Useful DOP-C02 - AWS Certified DevOps Engineer - Professional Official Study Guide



BONUS!!! Download part of TestSimulate DOP-C02 dumps for free: https://drive.google.com/open?id=1UAgVXDxwFbe1EjtA-dlE2h_1pfd3BJg2

The price of our DOP-C02 learning guide is among the range which you can afford and after you use our DOP-C02 study materials you will certainly feel that the value of the DOP-C02 exam questions far exceed the amount of the money you pay for the pass rate of our practice quiz is 98% to 100% which is unmatched in the market. Choosing our DOP-C02 Study Guide equals choosing the success and the perfect service.

The DOP-C02 prep guide adopt diversified such as text, images, graphics memory method, have to distinguish the markup to learn information, through comparing different color font, as well as the entire logical framework architecture, let users of the DOP-C02 training dump on the premise of grasping the overall layout, better clues to the formation of targeted long-term memory, and through the cycle of practice, let the knowledge more deeply printed in my mind. The DOP-C02 Exam Questions are so scientific and reasonable that you can easily remember everything of the DOP-C02 exam.

[**>> DOP-C02 Official Study Guide <<**](#)

DOP-C02 Real Test Practice Materials - DOP-C02 Study Guide - TestSimulate

We hope that you can use your time as much as possible for learning on the DOP-C02 practice questions. So we have considered every detail of the DOP-C02 study guide to remove all unnecessary programs. If you try to download our DOP-C02 study materials, you will find that they are so efficient! And even you free download the demos on the website, you can feel the convenience and efficiency. It is simple and easy to study with our DOP-C02 learning braindumps.

Amazon AWS Certified DevOps Engineer - Professional Sample Questions (Q368-Q373):

NEW QUESTION # 368

A DevOps engineer is implementing governance controls for a company that requires its infrastructure to be housed within the United States. The engineer must restrict which AWS Regions can be used, and ensure an alert is sent as soon as possible if any activity outside the governance policy takes place. The controls should be automatically enabled on any new Region outside the United States (US).

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

- A. Create an AWS Organizations SCP that denies access to all non-global services in non-US Regions. Attach the policy to the root of the organization.
- B. Use an AWS Lambda function that checks for AWS service activity and deploy it to all Regions. Write an Amazon EventBridge rule that runs the Lambda function every hour, sending an alert if activity is found in a non-US Region.
- C. Write an SCP using the aws:RequestedRegion condition key limiting access to US Regions. Apply the policy to all users, groups, and roles
- D. Use an AWS Lambda function to query Amazon Inspector to look for service activity in non-US Regions and send alerts if any activity is found.
- E. Configure AWS CloudTrail to send logs to Amazon CloudWatch Logs and enable it for all Regions. Use a CloudWatch Logs metric filter to send an alert on any service activity in non-US Regions.

Answer: A,E

Explanation:

To implement governance controls that restrict AWS service usage to within the United States and ensure alerts for any activity outside the governance policy, the following actions will meet the requirements:

A). Create an AWS Organizations SCP that denies access to all non-global services in non-US Regions. Attach the policy to the root of the organization. This action will effectively prevent users and roles in all accounts within the organization from accessing services in non-US Regions¹².

B). Configure AWS CloudTrail to send logs to Amazon CloudWatch Logs and enable it for all Regions. Use a CloudWatch Logs metric filter to send an alert on any service activity in non-US Regions. This action will allow monitoring of all AWS Regions and will trigger alerts if any activity is detected in non-US Regions, ensuring that the governance team is notified as soon as possible³.

AWS Documentation on Service Control Policies (SCPs) and how they can be used to manage permissions and restrict access based on Regions¹².

AWS Documentation on monitoring CloudTrail log files with Amazon CloudWatch Logs to set up alerts for specific activities³.

NEW QUESTION # 369

A DevOps team uses AWS CodePipeline, AWS CodeBuild, and AWS CodeDeploy to deploy an application.

The application is a REST API that uses AWS Lambda functions and Amazon API Gateway. Recent deployments have introduced errors that have affected many customers.

The DevOps team needs a solution that reverts to the most recent stable version of the application when an error is detected. The solution must affect the fewest customers possible.

Which solution Will meet these requirements With the MOST operational efficiency?

- A. Set the deployment configuration in CodeDeploy to LambdaAllAtOnce. Configure automatic rollbacks on the deployment group. Create an Amazon CloudWatch alarm that detects HTTP Bad Gateway errors on API Gateway. Configure the deployment group to roll back when the number of alarms meets the alarm threshold.
- B. Set the deployment configuration in CodeDeploy to LambdaAllAtOnce. Configure manual rollbacks on the deployment group. Create an Amazon Simple Notification Service (Amazon SNS) topic to send notifications every time a deployment fails. Configure the SNS topic to invoke a new Lambda function that stops the current deployment and starts the most recent successful deployment.
- C. Set the deployment configuration in CodeDeploy to LambdaCanary10Percent10Minutes. Configure manual rollbacks on the deployment group. Create a metric filter on an Amazon CloudWatch log group for API Gateway to monitor HTTP Bad Gateway errors. Configure the metric filter to invoke a new Lambda function that stops the current deployment and starts the most recent successful deployment.
- D. Set the deployment configuration in CodeDeploy to LambdaCanary10Percent10Minutes. Configure automatic rollbacks on the deployment group. Create an Amazon CloudWatch alarm that detects HTTP Bad Gateway errors on API Gateway. Configure the deployment group to roll back when the number of alarms meets the alarm threshold.

Answer: D

Explanation:

Explanation

Option A is incorrect because setting the deployment configuration to LambdaAllAtOnce means that the new version of the application will be deployed to all Lambda functions at once, affecting all customers.

This does not meet the requirement of affecting the fewest customers possible. Moreover, configuring automatic rollbacks on the deployment group is not operationally efficient, as it requires manual intervention to fix the errors and redeploy the application.

Option B is correct because setting the deployment configuration to LambdaCanary10Percent10Minutes means that the new version of the application will be deployed to 10 percent of the Lambda functions first, and then to the remaining 90 percent after 10 minutes. This minimizes the impact of errors on customers, as only 10 percent of them will be affected by a faulty deployment.

Configuring automatic rollbacks on the deployment group also meets the requirement of reverting to the most recent stable version of the application when an error is detected. Creating a CloudWatch alarm that detects HTTP Bad Gateway errors on API Gateway is a valid way to monitor the health of the application and trigger a rollback if needed.

Option C is incorrect because setting the deployment configuration to LambdaAllAtOnce means that the new version of the application will be deployed to all Lambda functions at once, affecting all customers.

This does not meet the requirement of affecting the fewest customers possible. Moreover, configuring manual rollbacks on the deployment group is not operationally efficient, as it requires human intervention to stop the current deployment and start a new one. Creating an SNS topic to send notifications every time a deployment fails is not sufficient to detect errors in the application, as it does not monitor the API Gateway responses.

Option D is incorrect because configuring manual rollbacks on the deployment group is not operationally efficient, as it requires human intervention to stop the current deployment and start a new one. Creating a metric filter on a CloudWatch log group for API Gateway to monitor HTTP Bad Gateway errors is a valid way to monitor the health of the application, but invoking a new Lambda function to perform a rollback is unnecessary and complex, as CodeDeploy already provides automatic rollback functionality.

References:

[AWS CodeDeploy Deployment Configurations](#)

[\[AWS CodeDeploy Rollbacks\]](#)

[Amazon CloudWatch Alarms](#)

NEW QUESTION # 370

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 average response time becomes greater than the longest response time that the application supports. Configure the CloudWatch alarm to use the SNS topic.
- B. Create an Amazon CloudWatch Synthetics 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.
- C. 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.
- D. 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.

Answer: B

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.

* 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. 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 # 371

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 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.
- B. 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.
- C. 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.
- D. 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.
- E. 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.

Answer: B,E

Explanation:

For the critical data, using a warm pool1 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 EC23. Using a launch template with the CloudWatch agent4 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.

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 # 372

A company needs to ensure that flow logs remain configured for all existing and new VPCs in its AWS account. The company uses an AWS CloudFormation stack to manage its VPCs. The company needs a solution that will work for any VPCs that any IAM user creates.

Which solution will meet these requirements?

- A. Turn on AWS Config. Create an AWS Config rule to check whether VPC flow logs are turned on. Configure automatic remediation to turn on VPC flow logs.
- B. Create an organization in AWS Organizations. Add the company's AWS account to the organization. Create an SCP to prevent users from modifying VPC flow logs.
- C. Add the resource to the CloudFormation stack that creates the VPCs.
- D. Create an IAM policy to deny the use of API calls for VPC flow logs. Attach the IAM policy to all IAM users.

Answer: A

Explanation:

To meet the requirements of ensuring that flow logs remain configured for all existing and new VPCs in the AWS account, the company should use AWS Config and automatic remediation. AWS Config is a service that enables customers to assess, audit, and evaluate the configurations of their AWS resources. AWS Config continuously monitors and records the configuration changes of the AWS resources and evaluates them against desired configurations. Customers can use AWS Config rules to define the desired configuration state of their AWS resources and trigger actions when a resource configuration violates a rule.

One of the AWS Config rules that customers can use is `vpc-flow-logs-enabled`, which checks whether VPC flow logs are enabled for all VPCs in an AWS account. Customers can also configure automatic remediation for this rule, which means that AWS Config will automatically enable VPC flow logs for any VPCs that do not have them enabled. Customers can specify the destination (CloudWatch Logs or S3) and the traffic type (all, accept, or reject) for the flow logs as remediation parameters. By using AWS Config and automatic remediation, the company can ensure that flow logs remain configured for all existing and new VPCs in its AWS account, regardless of who creates them or how they are created.

The other options are not correct because they do not meet the requirements or follow best practices. Adding the resource to the CloudFormation stack that creates the VPCs is not a sufficient solution because it will only work for VPCs that are created by using the CloudFormation stack. It will not work for VPCs that are created by using other methods, such as the console or the API.

Creating an organization in AWS Organizations and creating an SCP to prevent users from modifying VPC flow logs is not a good solution because it will not ensure that flow logs are enabled for all VPCs in the first place. It will only prevent users from disabling or changing flow logs after they are enabled. Creating an IAM policy to deny the use of API calls for VPC flow logs and attaching it to all IAM users is not a valid solution because it will prevent users from enabling or disabling flow logs at all. It will also not work for VPCs that are created by using other methods, such as the console or CloudFormation.

1: AWS::EC2::FlowLog - AWS CloudFormation

2: Amazon VPC Flow Logs extends CloudFormation Support to custom format subscriptions, 1-minute aggregation intervals and tagging

3: Logging IP traffic using VPC Flow Logs - Amazon Virtual Private Cloud About AWS Config - AWS Config `vpc-flow-logs-enabled` - AWS Config Remediate Noncompliant Resources with AWS Config Rules - AWS Config

NEW QUESTION # 373

.....

After using our DOP-C02 study materials, you will feel your changes. These changes will increase your confidence in continuing your studies on DOP-C02 real exam. Believe me, as long as you work hard enough, you can certainly pass the exam in the shortest possible time. The rest of the time, you can use to seize more opportunities. As long as you choose DOP-C02 simulating exam, we will be responsible to you.

Exam DOP-C02 Fees: <https://www.testsimulate.com/DOP-C02-study-materials.html>

Amazon DOP-C02 Official Study Guide The whole content is perfect and updated, Amazon DOP-C02 Official Study Guide All in all, large corporation appreciates people who have many certificates, In fact, you can enjoy the first-class services of our DOP-C02 exam guide, which in turn protects you from any unnecessary troubles, Amazon DOP-C02 Official Study Guide There are a bunch of people around the world who are worrying about their condition at present: want to improve their competitiveness above the average people and live without enough proof, eager to stand out to become an outstanding people with well-paid salary, now, it is time to realize their dreams and reject to be a daydreamer any more.

It's important to understand the difference between risk and a vulnerability, We also offer free Amazon DOP-C02 exam questions updates for 1 year after purchase, as well as a free DOP-C02 Practice Exam Questions demo before purchase.

2026 Professional DOP-C02 – 100% Free Official Study Guide | Exam AWS Certified DevOps Engineer - Professional Fees

The whole content is perfect and updated, All DOP-C02 Official Study Guide in all, large corporation appreciates people who have many certificates, In fact, you can enjoy the first-class services of our DOP-C02 exam guide, which in turn protects you from any unnecessary troubles.

There are a bunch of people around the world who are worrying DOP-C02 about their condition at present: want to improve their competitiveness above the average people and live without enough proof, eager to stand out to become an outstanding DOP-C02 Online Bootcamps people with well-paid salary, now, it is time to realize their dreams and reject to be a daydreamer any more.

Above all, DOP-C02 valid exam dumps will live up to your expectations and help you achieve your goals.

- DOP-C02 Latest Exam Pattern DOP-C02 Reliable Test Voucher DOP-C02 Test Sample Questions Easily obtain free download of **【 DOP-C02 】** by searching on \Rightarrow www.verifieddumps.com \Leftarrow DOP-C02 Reliable Source

P.S. Free & New DOP-C02 dumps are available on Google Drive shared by TestSimulate: https://drive.google.com/open?id=1UAqVXDxwFbe1EjtA-dle2h_1pfld3BJg2