Free PDF Quiz 2026 Efficient Amazon DOP-C02: Valid AWS Certified DevOps Engineer - Professional Exam Duration

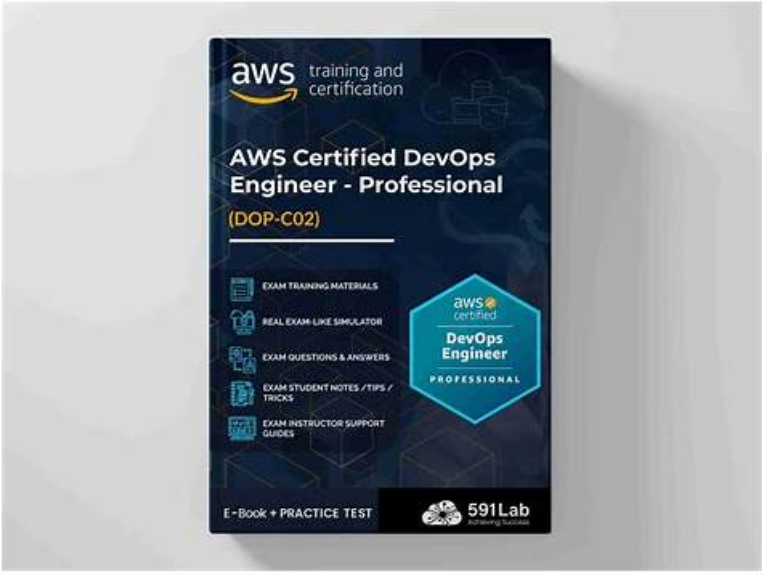

BTW, DOWNLOAD part of ExamDiscuss DOP-C02 dumps from Cloud Storage: https://drive.google.com/open?id=1xlJW-3h3nkU7cZj9kVseDz0tgqNmIP4M

The ExamDiscuss is a revolutionary platform for professionals and students looking to pass the Prepare for your AWS Certified DevOps Engineer - Professional (DOP-C02) exam and advance their careers. Our mission is to provide a comprehensive, convenient, and cost-effective preparation material for individuals to prepare for the DOP-C02 Certification Exam.

Experts at ExamDiscuss have also prepared Amazon DOP-C02 practice exam software for your self-assessment. This is especially handy for preparation and revision. You will be provided with an examination environment and you will be presented with actual exam Amazon DOP-C02 Exam Questions. This sort of preparation method enhances your knowledge which is crucial to excelling in the actual certification exam.

>> Valid DOP-C02 Exam Duration <<

Amazon DOP-C02 Accurate Answers - Online DOP-C02 Version

Owning the ExamDiscuss DOP-C02 exam certification training materials is equal to have a bright future, and equal to own the key to success. After you purchase ExamDiscuss's DOP-C02 certification exam training materials, we will provide one year free renewal service. If there's any quality problem in DOP-C02 Exam Dumps or you fail DOP-C02 exam certification, we will give a full refund unconditionally.

Amazon DOP-C02 Exam measures a candidate's ability to design and manage highly available and scalable systems on AWS. It covers a wide range of topics, including continuous integration and delivery (CI/CD) pipelines, monitoring and logging, security and compliance, infrastructure as code (IaC), and automation tools.

Amazon AWS Certified DevOps Engineer - Professional Sample Questions (Q155-Q160):

NEW QUESTION # 155
A company is using AWS CodeDeploy to automate software deployment. The deployment must meet these requirements:
* A number of instances must be available to serve traffic during the deployment Traffic must be balanced across those instances, and the instances must automatically heal in the event of failure.
* A new fleet of instances must be launched for deploying a new revision automatically, with no manual provisioning.

* Traffic must be rerouted to the new environment to half of the new instances at a time. The deployment should succeed if traffic is rerouted to at least half of the instances; otherwise, it should fail.
* Before routing traffic to the new fleet of instances, the temporary files generated during the deployment process must be deleted.
* At the end of a successful deployment, the original instances in the deployment group must be deleted immediately to reduce costs.
How can a DevOps engineer meet these requirements?

- A. Use an Application Load Balancer and a blue/green deployment. Associate the Auto Scaling group and the Application Load Balancer target group with the deployment group. Use the Automatically copy Auto scaling group option, and use CodeDeployDefault.HalfAtAtime as the deployment configuration. Instruct AWSCodeDeploy to terminate the original instances in the deployment group, and use the BeforeAllowTraffic hook within appspec.yml to delete the temporary tiles.
- B. Use an Application Load Balancer and an in-place deployment. Associate the Auto Scaling group and Application Load Balancer target group with the deployment group. Use the Automatically copy Auto Scaling group option, and use CodeDeployDefaulLAIIatOnce as a deployment configuration. Instruct AWS CodeDeploy to terminate the original instances in the deployment group, and use the BlockTraffic hook within appspec.yml to delete the temporary files.
- C. Use an Application Load Balancer and an in-place deployment. Associate the Auto Scaling group with the deployment group. Use the Automatically copy Auto Scaling group option. and use CodeDeployDefault.OneAtAtime as the deployment configuration. Instruct AWS CodeDeploy to terminate the original instances in the deployment group, and use the AllowTraffic hook within appspec.yml to delete the temporary files.
- D. Use an Application Load Balancer and a blue/green deployment. Associate the Auto Scaling group and Application Load Balancer target group with the deployment group. Use the Automatically copy Auto Scaling group option, create a custom deployment configuration with minimum healthy hosts defined as 50%. and assign the configuration to the deployment group. Instruct AWS CodeDeploy to terminate the original instances in the deployment group, and use the BeforeBlockTraffic hook within appspec.yml to delete the temporary files.

Answer: A

Explanation:
Step 1: Use a Blue/Green Deployment Strategy
A blue/green deployment strategy is necessary to meet the requirement of launching a new fleet of instances for each deployment and ensuring availability. In a blue/green deployment, the new version (green environment) is deployed to a separate set of instances, while the old version (blue environment) remains active. After testing the new version, traffic can be gradually shifted.
Action: Use AWS CodeDeploy's blue/green deployment configuration.
Why: Blue/green deployment minimizes downtime and ensures that traffic is shifted only to healthy instances.
Reference:
Step 2: Use an Application Load Balancer and Auto Scaling Group
The Application Load Balancer (ALB) is essential to balance traffic across multiple instances, and Auto Scaling ensures the deployment scales automatically to meet demand.
Action: Associate the Auto Scaling group and Application Load Balancer target group with the deployment group.
Why: This configuration ensures that traffic is evenly distributed and that instances automatically scale based on traffic load.
Step 3: Use Custom Deployment Configuration
The company requires that traffic be rerouted to at least half of the instances to succeed. AWS CodeDeploy allows you to configure custom deployment settings with specific thresholds for healthy hosts.
Action: Create a custom deployment configuration where 50% of the instances must be healthy.
Why: This ensures that the deployment continues only if at least 50% of the new instances are healthy.
Step 4: Clean Temporary Files Using Hooks
Before routing traffic to the new environment, the temporary files generated during the deployment must be deleted. This can be achieved using the BeforeAllowTraffic hook in the appspec.yml file.
Action: Use the BeforeAllowTraffic lifecycle event hook to clean up temporary files before routing traffic to the new environment.
Why: This ensures that the environment is clean before the new instances start serving traffic.
Step 5: Terminate Original Instances After Deployment
After a successful deployment, AWS CodeDeploy can automatically terminate the original instances (blue environment) to save costs.
Action: Instruct AWS CodeDeploy to terminate the original instances after the new instances are healthy.
Why: This helps in cost reduction by removing unused instances after the deployment.
This corresponds to Option C: Use an Application Load Balancer and a blue/green deployment. Associate the Auto Scaling group and the Application Load Balancer target group with the deployment group. Use the Automatically copy Auto Scaling group option, and use CodeDeployDefault.HalfAtATime as the deployment configuration. Instruct AWS CodeDeploy to terminate the original instances in the deployment group, and use the BeforeAllowTraffic hook within appspec.yml to delete the temporary files.

NEW QUESTION # 156

A company has deployed an application in a production VPC in a single AWS account. The application is popular and is experiencing heavy usage. The company's security team wants to add additional security, such as AWS WAF, to the application deployment. However, the application's product manager is concerned about cost and does not want to approve the change unless the security team can prove that additional security is necessary.
The security team believes that some of the application's demand might come from users that have IP addresses that are on a deny list. The security team provides the deny list to a DevOps engineer. If any of the IP addresses on the deny list access the application, the security team wants to receive automated notification in near real time so that the security team can document that the application needs additional security. The DevOps engineer creates a VPC flow log for the production VPC.
Which set of additional steps should the DevOps engineer take to meet these requirements MOST cost-effectively?

- A. Create a log group in Amazon CloudWatch Logs. Create an Amazon S3 bucket to hold query results.Configure the VPC flow log to capture all traffic and to send the data to the log group. Deploy an Amazon Athena CloudWatch connector in AWS Lambda. Connect the connector to the log group.Configure Athena to periodically query for all accepted traffic from the IP addresses on the deny list and to store the results in the S3 bucket. Configure an S3 event notification to automatically notify the security team through an Amazon Simple Notification Service (Amazon SNS) topic when new objects are added to the S3 bucket.
- B. Create an Amazon S3 bucket for log files. Configure the VPC flow log to capture accepted traffic and to send the data to the S3 bucket. Configure an Amazon OpenSearch Service cluster and domain for the log files. Create an AWS Lambda function to retrieve the logs from the S3 bucket, format the logs, and load the logs into the OpenSearch Service cluster. Schedule the Lambda function to run every 5 minutes.
  Configure an alert and condition in OpenSearch Service to send alerts to the security team through an Amazon Simple Notification Service (Amazon SNS) topic when access from the IP addresses on the deny list is detected.
- C. Create an Amazon S3 bucket for log files. Configure the VPC flow log to capture all traffic and to send the data to the S3 bucket. Configure Amazon Athena to return all log files in the S3 bucket for IP addresses on the deny list. Configure Amazon QuickSight to accept data from Athena and to publish the data as a dashboard that the security team can access. Create a threshold alert of 1 for successful access.
  Configure the alert to automatically notify the security team as frequently as possible when the alert threshold is met.
- D. Create a log group in Amazon CloudWatch Logs. Configure the VPC flow log to capture accepted traffic and to send the data to the log group. Create an Amazon CloudWatch metric filter for IP addresses on the deny list. Create a CloudWatch alarm with the metric filter as input. Set the period to 5 minutes and the datapoints to alarm to 1. Use an Amazon Simple Notification Service (Amazon SNS) topic to send alarm notices to the security team.

Answer: D

NEW QUESTION # 157
A company is launching an application that stores raw data in an Amazon S3 bucket. Three applications need to access the data to generate reports. The data must be redacted differently for each application before the applications can access the data.
Which solution will meet these requirements?

- A. For each application, create an S3 access point that uses the raw data's S3 bucket as the destination. Create an AWS Lambda function that is invoked by object creation events in the raw data's S3 bucket. Program the Lambda function to redact data for each application. Store the data in each application's S3 access point. Configure each application to consume data from its own S3 access point.
- B. Create an S3 access point that uses the raw data's S3 bucket as the destination. For each application, create an S3 Object Lambda access point that uses the S3 access point. Configure the AWS Lambda function for each S3 Object Lambda access point to redact data when objects are retrieved. Configure each application to consume data from its own S3 Object Lambda access point.
- C. Create an S3 bucket for each application. Configure S3 Same-Region Replication (SRR) from the raw data's S3 bucket to each application's S3 bucket. Configure each application to consume data from its own S3 bucket.
- D. Create an Amazon Kinesis data stream. Create an AWS Lambda function that is invoked by object creation events in the raw data's S3 bucket. Program the Lambda function to redact data for each application. Publish the data on the Kinesis data stream. Configure each application to consume data from the Kinesis data stream.

Answer: B

Explanation:
The best solution is to use S3 Object Lambda1, which allows you to add your own code to S3 GET, LIST, and HEAD requests to modify and process data as it is returned to an application2. This way, you can redact the data differently for each application without creating and storing multiple copies of the data or running proxies.
The other solutions are less efficient or scalable because they require replicating the data to multiple buckets, streaming the data

through Kinesis, or storing the data in S3 access points.

NEW QUESTION # 158
A company deploys an application on on-premises devices in the company's on-premises data center. The company uses an AWS Direct Connect connection between the data center and the company's AWS account.
During initial setup of the on-premises devices and during application updates, the application needs to retrieve configuration files from an Amazon Elastic File System (Amazon EFS) file system. All traffic from the on-premises devices to Amazon EFS must remain private and encrypted. The on-premises devices must follow the principle of least privilege for AWS access. The company's DevOps team needs the ability to revoke access from a single device without affecting the access of the other devices. Which combination of steps will meet these requirements? (Select TWO.)

- A. Use the native Linux NFS client to mount the EFS file system.
- B. Generate certificates for each on-premises device in AWS Private Certificate Authority. Create a trust anchor in IAM Roles Anywhere that references an AWS Private CA. Create an IAM role that trusts IAM Roles Anywhere. Attach the AmazonElasticFileSystemClientReadWriteAccess to the role. Create an IAM Roles Anywhere profile for the IAM role. Configure the AWS CLI on the on-premises devices to use the aws_signing_helper command to obtain credentials.
- C. Create an IAM user that has an access key and a secret key for all devices. Attach the AmazonElasticFileSystemClientReadWriteAccess policy to the IAM user. Configure the AWS CLI on the on-premises devices to use the IAM user's access key and secret key.
- D. Use the amazon-efs-utils package to mount the EFS file system.
- E. Create an IAM user that has an access key and a secret key for each device. Attach the AmazonElasticFileSystemFullAccess policy to all IAM users. Configure the AWS CLI on the on- premises devices to use the IAM user's access key and secret key.

Answer: B,D

NEW QUESTION # 159
A company runs a web application that extends across multiple Availability Zones. The company uses an Application Load Balancer (ALB) for routing. AWS Fargate (or the application and Amazon Aurora for the application data The company uses AWS CloudFormation templates to deploy the application The company stores all Docker images in an Amazon Elastic Container Registry (Amazon ECR) repository in the same AWS account and AWS Region.
A DevOps engineer needs to establish a disaster recovery (DR) process in another Region. The solution must meet an RPO of 8 hours and an RTO of 2 hours The company sometimes needs more than 2 hours to build the Docker images from the Dockerfile Which solution will meet the RTO and RPO requirements MOST cost-effectively?

- A. Copy the CloudFormation templates to an Amazon S3 bucket in the DR Region. Use Amazon EventBridge to schedule an AWS Lambda function to take an hourly snapshot of the Aurora database and of the most recent Docker image in the ECR repository. Copy the snapshot and the Docker image to the DR Region in case of DR, use the CloudFormation template with the most recent Aurora snapshot and the Docker image from the local ECR repository to launch a new CloudFormation stack in the DR Region
- B. Copy the CloudFormation templates to an Amazon S3 bucket in the DR Region. Deploy a second application CloudFormation stack in the DR Region. Reconfigure Aurora to be a global database Update both CloudFormation stacks when a new application release in the current Region is needed. In case of DR. update, the application DNS records to point to the new ALB.
- C. Copy the CloudFormation templates and the Dockerfile to an Amazon S3 bucket in the DR Region Use AWS Backup to configure automated Aurora cross-Region hourly snapshots In case of DR, build the most recent Docker image and upload the Docker image to an ECR repository in the DR Region Use the CloudFormation template that has the most recent Aurora snapshot and the Docker image from the ECR repository to launch a new CloudFormation stack in the DR Region Update the application DNS records to point to the new ALB
- D. Copy the CloudFormation templates to an Amazon S3 bucket in the DR Region Configure Aurora automated backup Cross-Region Replication Configure ECR Cross-Region Replication. In case of DR use the CloudFormation template with the most recent Aurora snapshot and the Docker image from the local ECR repository to launch a new CloudFormation stack in the DR Region Update the application DNS records to point to the new ALB

Answer: D

Explanation:
The most cost-effective solution to meet the RTO and RPO requirements is option B. This option involves copying the CloudFormation templates to an Amazon S3 bucket in the DR Region, configuring Aurora automated backup Cross-Region

Replication, and configuring ECR Cross-Region Replication. In the event of a disaster, the CloudFormation template with the most recent Aurora snapshot and the Docker image from the local ECR repository can be used to launch a new CloudFormation stack in the DR Region. This approach avoids the need to build Docker images from the Dockerfile, which can sometimes take more than 2 hours, thus meeting the RTO requirement. Additionally, the use of automated backups and replication ensures that the RPO of 8 hours is met.
Reference:
AWS Documentation on Disaster Recovery: Plan for Disaster Recovery (DR) - Reliability Pillar AWS Blog on Establishing RPO and RTO Targets: Establishing RPO and RTO Targets for Cloud Applications AWS Documentation on ECR Cross-Region Replication: Amazon ECR Cross-Region Replication AWS Documentation on Aurora Cross-Region Replication: Replicating Amazon Aurora DB Clusters Across AWS Regions

NEW QUESTION # 160
.....