High Hit Rate Reasonable DOP-C02 Exam Price Help You to Get Acquainted with Real DOP-C02 Exam Simulation

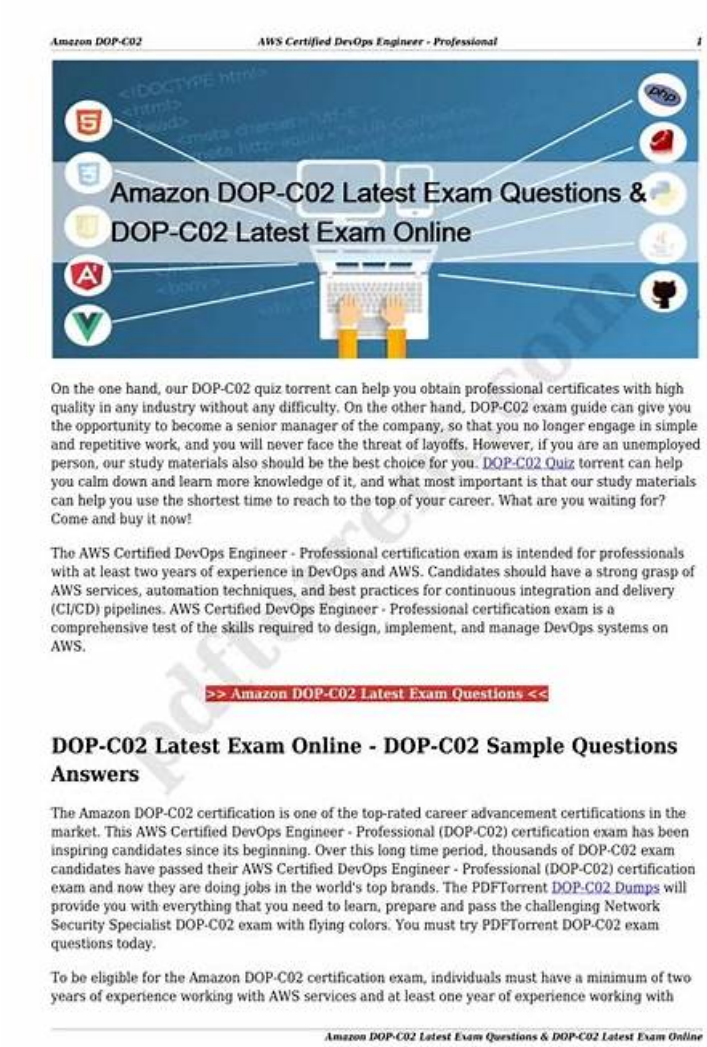

P.S. Free 2025 Amazon DOP-C02 dumps are available on Google Drive shared by TestKingFree: https://drive.google.com/open?id=1KtLKT_XmdeMIU2eu9Ai-oeyFNzeDuGqa

We are sure you can seep great deal of knowledge from our DOP-C02 study prep in preference to other materials obviously. Our DOP-C02 practice materials have variant kinds including PDF, app and software versions. As DOP-C02 Exam Questions with high prestige and esteem in the market, we hold sturdy faith for you. And you will find that our DOP-C02 learning quiz is quite popular among the candidates all over the world.

Amazon DOP-C02 (AWS Certified DevOps Engineer - Professional) Certification Exam is a highly sought-after certification for professionals who are looking to validate their skills and knowledge in the field of DevOps. AWS Certified DevOps Engineer - Professional certification exam is designed to test the candidate's proficiency in deploying, managing, and operating highly available, scalable, and fault-tolerant systems on the AWS cloud platform.

>> Reasonable DOP-C02 Exam Price <<

DOP-C02 Tesking Torrent - DOP-C02 Pdf Questions & DOP-C02 Practice Training

Eliminates confusion while taking the Amazon DOP-C02 certification exam. Prepares you for the format of your DOP-C02 exam dumps, including multiple-choice questions and fill-in-the-blank answers. Comprehensive, up-to-date coverage of the entire AWS Certified DevOps Engineer - Professional (DOP-C02) certification curriculum. Amazon DOP-C02 practice questions are based on recently released DOP-C02 exam objectives.

Amazon AWS Certified DevOps Engineer - Professional Sample Questions (Q283-Q288):

NEW QUESTION # 283
A company deploys an application to Amazon EC2 instances. The application runs Amazon Linux 2 and uses AWS CodeDeploy. The application has the following file structure for its code repository:

The appspec.yml file has the following contents in the files section:

What will the result be for the deployment of the config.txt file?

- A. The config.txt file will be deployed to only /var/www/html/config/config txt
- B. The config.txt file will be deployed to only /usr/local/src/config txt
- C. The config txt file will be deployed to /usr/local/src/config.txt and to /var/www/html/application/web/config txt
- D. The config.txt file will be deployed to /usr/local/src/config.txt and to /var/www/html/config/config txt.

Answer: B

Explanation:
Deployment of config.txt file based on the appspec.yml:
The appspec.yml file specifies that config/config.txt should be copied to /usr/local/src/config.txt.
The source: / directive in the appspec.yml indicates that the entire directory structure starting from the root of the application source should be copied to the specified destination, which is /var/www/html.
Result of the Deployment:
The config.txt file will be specifically deployed to /usr/local/src/config.txt as per the explicit file mapping.
The entire directory structure including application/web will be copied to /var/www/html, but this does not include config/config.txt since it has a specific destination defined.
Thus, the config.txt file will be deployed only to /usr/local/src/config.txt.
Therefore, the correct answer is:
C . The config.txt file will be deployed to only /usr/local/src/config.txt.
Reference:
AWS CodeDeploy AppSpec File Reference
AWS CodeDeploy Deployment Process

NEW QUESTION # 284
A company manages environments for its application in multiple AWS accounts. Each environment account is in a different OU in AWS Organizations.
A DevOps team is responsible for the application deployment process across the environments. The deployment process uses an AWS CodePipeline pipeline in a Shared Services account. The DevOps team members are in the same user group. The team members have administrative access to all accounts through AWS IAM Identity Center.
A recent deployment problem in the development environment required the DevOps team to perform manual steps. The deployment to the production environment then resulted in an incident that caused the pipeline to fail, blocking new deployments for several hours.
A DevOps engineer needs to ensure that only the pipeline can perform deployments in the production environment. The DevOps engineer must have access to the environment in case of an emergency.
Which solution will meet these requirements with the MOST operational efficiency?

- A. Update the DevOps group to have the ReadOnlyAccess permission set for the production accounts. Configure the DevOps engineer user with a new permission set that has AdministratorAccess permissions and that allows the user to assume the pipeline role. Add an SCP that denies modification of resources by any entity other than the pipeline role.
- B. Update the DevOps group to be able to assume the pipeline role for the production accounts. Configure a new user in IAM Identity Center for the DevOps engineer with a new permission set that has AdministratorAccess permissions. Add an SCP that denies modification of resources by any entity other than the DevOps engineer.
- C. Create an SCP that denies all write actions for the DevOps team members on the production OU. Use a specific tag to tag the resources that CodePipeline provisions. Add an SCP that denies modification of tagged resources by any entity other than the DevOps engineer.
- D. Create an SCP that denies all write actions for the DevOps team members on the production OU. Configure a new user in IAM Identity Center for the DevOps engineer with a new permission set that has AdministratorAccess permissions. Add an SCP that denies modification of resources by any entity other than the pipeline role.

Answer: A

Explanation:
Comprehensive and Detailed Explanation From Exact Extract:
The requirement is to restrict production deployments strictly to the pipeline, while still allowing emergency access to a specific engineer.
* The best approach is to restrict the DevOps team toread-only accessin production accounts, minimizing risk of manual changes (Option A).
* The DevOps engineer can have an admin permission set but assume thepipeline IAM rolefor deployment, enforcing strict control.
* Applying an SCP todeny modification by anyone other than the pipeline roleenforces this at the organization level.Option B is similar but unnecessarily creates separate IAM users, increasing management overhead. Option C grants the DevOps engineer broader permissions that may conflict with controls. Option D complicates management with tagging and SCPs, increasing operational overhead.
Reference:
AWS Organizations Service Control Policies (SCPs):"SCPs can restrict what actions identities in member accounts can perform, even for administrators."(AWS Organizations SCP Documentation) IAM Identity Center Role Assumption Best Practices:"Use role assumption for limited elevated permissions instead of broad admin access."(AWS IAM Best Practices)

NEW QUESTION # 285
A company's application is currently deployed to a single AWS Region. Recently, the company opened a new office on a different continent. The users in the new office are experiencing high latency. The company's application runs on Amazon EC2 instances behind an Application Load Balancer (ALB) and uses Amazon DynamoDB as the database layer. The instances run in an EC2 Auto Scaling group across multiple Availability Zones. A DevOps engineer is tasked with minimizing application response times and improving availability for users in both Regions.
Which combination of actions should be taken to address the latency issues? (Choose three.)

- A. Create Amazon Route 53 aliases, health checks, and failover routing policies to route to the ALB.
- B. Create a new DynamoDB table in the new Region with cross-Region replication enabled.
- C. Create new ALB and Auto Scaling group resources in the new Region and configure the new ALB to direct traffic to the new Auto Scaling group.
- D. Create new ALB and Auto Scaling group global resources and configure the new ALB to direct traffic to the new Auto Scaling group.
- E. Create Amazon Route 53 records, health checks, and latency-based routing policies to route to the ALB.
- F. Convert the DynamoDB table to a global table.

Answer: C,E,F

NEW QUESTION # 286
A company that runs many workloads on AWS has an Amazon EBS spend that has increased over time. The DevOps team notices there are many unattached EBS volumes. Although there are workloads where volumes are detached, volumes over 14 days old are stale and no longer needed. A DevOps engineer has been tasked with creating automation that deletes unattached EBS volumes that have been unattached for 14 days.
Which solution will accomplish this?

- A. Configure the AWS Config ec2-volume-inuse-check managed rule with a configuration changes trigger type and an

Amazon EC2 volume resource target. Create a new Amazon CloudWatch Events rule scheduled to execute an AWS Lambda function in 14 days to delete the specified EBS volume.

- B. Use AWS Trusted Advisor to detect EBS volumes that have been detached for more than 14 days.Execute an AWS Lambda function that creates a snapshot and then deletes the EBS volume.
- C. Use Amazon EC2 and Amazon Data Lifecycle Manager to configure a volume lifecycle policy. Set the interval period for unattached EBS volumes to 14 days and set the retention rule to delete. Set the policy target volumes as *.
- D. Create an Amazon CloudWatch Events rule to execute an AWS Lambda function daily. The Lambda function should find unattached EBS volumes and tag them with the current date, and delete unattached volumes that have tags with dates that are more than 14 days old.

Answer: D

Explanation:
Explanation
The requirement is to create automation that deletes unattached EBS volumes that have been unattached for 14 days. To do this, the DevOps engineer needs to use the following steps:
* Create an Amazon CloudWatch Events rule to execute an AWS Lambda function daily. CloudWatch Events is a service that enables event-driven architectures by delivering events from various sources to targets. Lambda is a service that lets you run code without provisioning or managing servers. By creating a CloudWatch Events rule that executes a Lambda function daily, the DevOps engineer can schedule a recurring task to check and delete unattached EBS volumes.
* The Lambda function should find unattached EBS volumes and tag them with the current date, and
* delete unattached volumes that have tags with dates that are more than 14 days old. The Lambda function can use the EC2 API to list and filter unattached EBS volumes based on their state and tags.
The function can then tag each unattached volume with the current date using the create-tags command.
The function can also compare the tag value with the current date and delete any unattached volume that has been tagged more than 14 days ago using the delete-volume command.

NEW QUESTION # 287
A DevOps engineer is building an application that uses an AWS Lambda function to query an Amazon Aurora MySQL DB cluster. The Lambda function performs only read queries. Amazon EventBridge events invoke the Lambda function.
As more events invoke the Lambda function each second, the database's latency increases and the database's throughput decreases. The DevOps engineer needs to improve the performance of the application.
Which combination of steps will meet these requirements? (Select THREE.)

- A. Use Amazon RDS Proxy to create a proxy. Connect the proxy to the Aurora cluster reader endpoint. Set a maximum connections percentage on the proxy.
- B. Connect to the proxy endpoint from the Lambda function.
- C. Connect to the Aurora cluster endpoint from the Lambda function.
- D. Implement the database connection opening and closing inside the Lambda event handler code.
- E. Implement the database connection opening outside the Lambda event handler code.
- F. Implement database connection pooling inside the Lambda code. Set a maximum number of connections on the database connection pool.

Answer: A,B,E

Explanation:
Explanation
Short Explanation: To improve the performance of the application, the DevOps engineer should use Amazon RDS Proxy, implement the database connection opening outside the Lambda event handler code, and connect to the proxy endpoint from the Lambda function.
References:
Amazon RDS Proxy is a fully managed, highly available database proxy for Amazon Relational Database Service (RDS) that makes applications more scalable, more resilient to database failures, and more secure1. By using Amazon RDS Proxy, the DevOps engineer can reduce the overhead of opening and closing connections to the database, which can improve latency and throughput2.
The DevOps engineer should connect the proxy to the Aurora cluster reader endpoint, which allows read-only connections to one of the Aurora Replicas in the DB cluster3. This can help balance the load across multiple read replicas and improve performance for read-intensive workloads4.
The DevOps engineer should implement the database connection opening outside the Lambda event handler code, which means using a global variable to store the database connection object5. This can enable connection reuse across multiple invocations of the Lambda function, which can reduce latency and improve performance.

The DevOps engineer should connect to the proxy endpoint from the Lambda function, which is a unique URL that represents the proxy. This can allow the Lambda function to access the database through the proxy, which can provide benefits such as connection pooling, load balancing, failover handling, and enhanced security.
The other options are incorrect because:
Implementing database connection pooling inside the Lambda code is unnecessary and redundant when using Amazon RDS Proxy, which already provides connection pooling as a service.
Implementing the database connection opening and closing inside the Lambda event handler code is inefficient and costly, as it can increase latency and consume more resources for each invocation of the Lambda function.
Connecting to the Aurora cluster endpoint from the Lambda function is not optimal for read-only queries, as it can direct traffic to either the primary instance or one of the Aurora Replicas in the DB cluster. This can result in inconsistent performance and potential conflicts with write operations on the primary instance.

NEW QUESTION # 288
.....