

AWS-DevOps-Engineer-Professional Study Guide: AWS Certified DevOps Engineer - Professional & AWS-DevOps-Engineer-Professional Dumps Torrent & AWS-DevOps-Engineer-Professional Latest Dumps



BONUS!!! Download part of CramPDF AWS-DevOps-Engineer-Professional dumps for free: https://drive.google.com/open?id=1mOrS_cim_m43IjrTUJAn96JvXErB7een

Nowadays, using electronic AWS-DevOps-Engineer-Professional exam materials to prepare for the exam has become more and more popular, so now, you really should not be restricted to paper materials any more, our electronic Amazon AWS-DevOps-Engineer-Professional preparation dumps will surprise you with their effectiveness and usefulness. I can assure you that you will pass the exam as well as getting the related AWS-DevOps-Engineer-Professional Certification under the guidance of our AWS-DevOps-Engineer-Professional training materials as easy as pie.

The AWS Certified DevOps Engineer Professional (DOP-C01) exam is a highly sought-after certification for experienced DevOps professionals looking to validate their knowledge and skills in deploying, managing, and operating highly available, fault-tolerant, and scalable systems on the AWS platform. AWS Certified DevOps Engineer - Professional certification is designed to test candidates' proficiency in deploying and managing continuous delivery systems and methodologies on AWS, as well as their ability to implement and manage automation technologies and tools that enable secure and reliable infrastructure deployments.

Amazon AWS-DevOps (AWS Certified DevOps Engineer - Professional (DOP-C01)) Certification Exam is designed for professionals who possess advanced knowledge and skills in the field of DevOps. AWS Certified DevOps Engineer - Professional certification validates the expertise of the professionals in deploying, managing, and operating highly available, scalable, and fault-tolerant systems on the AWS platform. AWS-DevOps-Engineer-Professional Exam Tests the candidate's proficiency in DevOps principles, practices, and tools, including continuous integration, continuous delivery, infrastructure as code, monitoring, and logging.

>> New AWS-DevOps-Engineer-Professional Exam Topics <<

Pass4sure AWS Certified DevOps Engineer - Professional certification - Amazon AWS-DevOps-Engineer-Professional sure exam practice

We have always been made rapid progress on our Amazon AWS-DevOps-Engineer-Professional training materials because of the merits of high-efficiency and perfect after-sales services online for 24 hours. Studying with our AWS-DevOps-Engineer-Professional Actual Exam, you can get the most professional information and achieve your dreaming scores by your first go.

The AWS Certified DevOps Engineer - Professional (DOP-C01) certification is highly valued in the industry and is recognized by employers worldwide. It is a testament to the skills and knowledge of the candidate in AWS DevOps and demonstrates their ability to design, deploy, and manage DevOps systems on the AWS platform. AWS Certified DevOps Engineer - Professional certification opens up new job opportunities and career advancement prospects for professionals in the field.

Amazon AWS Certified DevOps Engineer - Professional Sample Questions (Q214-Q219):

NEW QUESTION # 214

Your system uses a multi-master, multi-region DynamoDB configuration spanning two regions to achieve high availability. For the first time since launching your system, one of the AWS Regions in which you operate went down for 3 hours, and the failover worked correctly. However, after recovery, your users are experiencing strange bugs, in which users on different sides of the globe see different data. What is a likely design issue that was not accounted for when launching?

- A. The system did not include repair logic and request replay buffering logic for post-failure, to re- synchronize data to the Region that was unavailable for a number of hours.
- B. The system did not implement DynamoDB Table Defragmentation for restoring partition performance in the Region that experienced an outage, so data is served stale.
- C. The system did not use DynamoDB Consistent Read requests, so the requests in different areas are not utilizing consensus across Regions at runtime.
- D. The system does not have Lambda Function Repair Automations, to perform table scans and check for corrupted partition blocks inside the Table in the recovered Region.

Answer: A

Explanation:

When using multi-region DynamoDB systems, it is of paramount importance to make sure that all requests made to one Region are replicated to the other. Under normal operation, the system in question would correctly perform write replays into the other Region. If a whole Region went down, the system would be unable to perform these writes for the period of downtime. Without buffering write requests somehow, there would be no way for the system to replay dropped cross- region writes, and the requests would be serviced differently depending on the Region from which they were served after recovery.

<http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/Streams.CrossRegionRepl.html>

NEW QUESTION # 215

For AWS CloudFormation, which is true?

- A. Custom resources using SNS do not need a `<code>ServiceToken</code>` property.
- B. **Custom resources using Lambda and `<code>Code.ZipFile</code>` allow inline nodejs resource composition.**
- C. Custom resources using SNS have a default timeout of 3 minutes.
- D. Custom resources using Lambda do not need a `<code>ServiceToken</code>` property

Answer: B

Explanation:

Code is a property of the AWS::Lambda::Function resource that enables you to specify the source code of an AWS Lambda (Lambda) function.

You can point to a file in an Amazon Simple Storage Service (Amazon S3) bucket or specify your source code as inline text (for nodejs runtime environments only).

<http://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/template-custom-resources.html>

NEW QUESTION # 216

Your CTO has asked you to make sure that you know what all users of your AWS account are doing to change resources at all times. She wants a report of who is doing what over time, reported to her once per week, for as broad a resource type group as

possible. How should you do this?

- A. Use AWS Config with an SNS subscription on a Lambda, and insert these changes over time into a DynamoDB table. Generate reports based on the contents of this table.
- B. Use CloudWatch Events Rules with an SNS topic subscribed to all AWS API calls. Subscribe the CTO to an email type delivery on this SNS Topic.
- C. Use AWS IAM credential reports to deliver a CSV of all uses of IAM User Tokens over time to the CTO.
- D. Create a global AWS CloudTrail Trail. Configure a script to aggregate the log data delivered to S3 once per week and deliver this to the CTO.

Answer: D

Explanation:

This is the ideal use case for AWS CloudTrail.

CloudTrail provides visibility into user activity by recording API calls made on your account.

CloudTrail records important information about each API call, including the name of the API, the identity of the caller, the time of the API call, the request parameters, and the response elements returned by the AWS service. This information helps you to track changes made to your AWS resources and to troubleshoot operational issues. CloudTrail makes it easier to ensure compliance with internal policies and regulatory standards.

<https://aws.amazon.com/cloudtrail/faqs/>

NEW QUESTION # 217

A company hosts parts of a Python-based application using AWS Elastic Beanstalk. An Elastic Beanstalk CLI is being used to create and update the environments. The Operations team detected an increase in requests in one of the Elastic Beanstalk environments that caused downtime overnight. The team noted that the policy used for AWS Auto Scaling is NetworkOut. Based on load testing metrics, the team determined that the application needs to scale CPU utilization to improve the resilience of the environments. The team wants to implement this across all environments automatically. Following AWS recommendations, how should this automation be implemented?

- A. Using ebextensions, place a command within the container_commands key to perform an API call to modify the scaling metric to CPUUtilization for the Auto Scaling configuration. Use leader_only to execute this command in only the first instance launched within the environment.
- B. Using ebextensions, create a custom resource that modifies the AWSEBAutoScalingScaleUpPolicy and AWSEBAutoScalingScaleDownPolicy resources to use CPUUtilization as a metric to scale for the Auto Scaling group.
- C. Using ebextensions, configure the option setting MeasureName to CPUUtilization within the aws:autoscaling:trigger namespace.
- D. Using ebextensions, place a script within the files key and place it in /opt/elasticbeanstalk/hooks/appdeploy/pre to perform an API call to modify the scaling metric to CPUUtilization for the Auto Scaling configuration. Use leader_only to place this script in only the first instance launched within the environment.

Answer: C

NEW QUESTION # 218

One of your engineers has written a web application in the Go Programming language and has asked your DevOps team to deploy it to AWS. The application code is hosted on a Git repository.

What are your options? (Select Two)

- A. Create a new AWS Elastic Beanstalk application and configure a Go environment to host your application, Using Git check out the latest version of the code, once the local repository for Elastic Beanstalk is configured use "eb create" command to create an environment and then use "eb deploy" command to deploy the application.
- B. Write a Dockerfile that installs the Go base image and fetches your application using Git, Create an AWS CloudFormation template that creates and associates an AWS::EC2::Instance resource type with an AWS::EC2::Container resource type.
- C. Write a Dockerfile that installs the Go base image and uses Git to fetch your application. Create a new AWS OpsWorks stack that contains a Docker layer that uses the Dockerfile to deploy your container and then use the Dockerfile to automate the deployment.
- D. Write a Dockerfile that installs the Go base image and fetches your application using Git, Create a new AWS Elastic Beanstalk application and use this Dockerfile to automate the deployment.

Answer: A,D

Explanation:

Explanation

Opsworks works with Chef recipes and not with Docker containers so Option B and C are invalid.

There is no AWS::CC2::Container resource for Cloudformation so Option D is invalid.

Below is the documentation on Clastic beanstalk and Docker

Clastic Beanstalk supports the deployment of web applications from Docker containers. With Docker containers, you can define your own runtime environment. You can choose your own platform, programming language, and any application dependencies (such as package managers or tools), that aren't supported by other platforms. Docker containers are self-contained and include all the configuration information and software your web application requires to run.

For more information on Clastic beanstalk and Docker, please visit the link:

* http://docs.aws.amazon.com/elasticbeanstalk/latest/dg/create_deploy_docker.html

* <https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/eb-cli3-getting-started.html>

* <https://docs.aws.amazon.com/elasticbeanstalk/latest/dg/eb-cli3-getting-started.html>

What's more, part of that CramPDF AWS-DevOps-Engineer-Professional dumps now are free: https://drive.google.com/open?id=1mOrS_cim_m43IjrTUJAn96JvXErB7een