DOP-C02$^{Q\&As}$

AWS Certified DevOps Engineer - Professional

Pass Amazon DOP-C02 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.geekcert.com/dop-c02.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon Official Exam Center

QUESTION 1

When logging with Amazon CloudTrail, API call information for services with regional end points is ____.

A. captured and processed in the same region as to which the API call is made and delivered to the region associated with your Amazon S3 bucket

B. captured, processed, and delivered to the region associated with your Amazon S3 bucket

C. captured in the same region as to which the API call is made and processed and delivered to the region associated with your Amazon S3 bucket

D. captured in the region where the end point is located, processed in the region where the CloudTrail trail is configured, and delivered to the region associated with your Amazon S3 bucket

Correct Answer: A

When logging with Amazon CloudTrail, API call information for services with regional end points (EC2, RDS etc.) is captured and processed in the same region as to which the API call is made and delivered to the region associated with your Amazon S3 bucket. API call information for services with single end points (IAM, STS etc.) is captured in the region where the end point is located, processed in the region where the CloudTrail trail is configured, and delivered to the region associated with your Amazon S3 bucket.

Reference: https://aws.amazon.com/cloudtrail/faqs/

QUESTION 2

A company that uses electronic health records is running a fleet of Amazon EC2 instances with an Amazon Linux operating system. As part of patient privacy requirements, the company must ensure continuous compliance for patches for operating system and applications running on the EC2 instances.

How can the deployments of the operating system and application patches be automated using a default and custom repository?

A. Use AWS Systems Manager to create a new patch baseline including the custom repository. Run the AWS-RunPatchBaseline document using the run command to verify and install patches.

B. Use AWS Direct Connect to integrate the corporate repository and deploy the patches using Amazon CloudWatch scheduled events, then use the CloudWatch dashboard to create reports.

C. Use yum-config-manager to add the custom repository under /etc/yum.repos.d and run yum-config-manager-enable to activate the repository.

D. Use AWS Systems Manager to create a new patch baseline including the corporate repository. Run the AWS-AmazonLinuxDefaultPatchBaseline document using the run command to verify and install patches.

Correct Answer: A

https://docs.aws.amazon.com/systems-manager/latest/userguide/patch-manager-how-it-works-alt-source-repository.html

QUESTION 3

A DevOps engineer is researching the least expensive way to implement an image batch processing cluster on AWS. The application cannot run in Docker containers and must run on Amazon EC2. The batch job stores checkpoint data on an NFS volume and can tolerate interruptions. Configuring the cluster software from a generic EC2 Linux image takes 30 minutes.

What is the MOST cost-effective solution?

A. Use Amazon EFS for checkpoint data. To complete the job, use an EC2 Auto Scaling group and an On-Demand pricing model to provision EC2 instances temporarily.

B. Use GlusterFS on EC2 instances for checkpoint data. To run the batch job, configure EC2 instances manually. When the job completes, shut down the instances manually.

C. Use Amazon EFS for checkpoint data. Use EC2 Fleet to launch EC2 Spot Instances, and utilize user data to configure the EC2 Linux instance on startup.

D. Use Amazon EFS for checkpoint data. Use EC2 Fleet to launch EC2 Spot Instances. Create a custom AMI for the cluster and use the latest AMI when creating instances.

Correct Answer: D

QUESTION 4

A DevOps engineer needs to grant several external contractors access to a legacy application that runs on an Amazon Linux Amazon EC2 instance. The application server is available only in a private subnet. The contractors are not authorized for VPN access.

What should the DevOps engineer do to grant the contactors access to the application server?

A. Create an IAM user and SSH keys for each contractor. Add the public SSH key to the application server\\'s SSH authorized_keys file. Instruct the contractors to install the AWS CLI and AWS Systems Manager Session Manager plugin, update their AWS credentials files with their private keys, and use the aws ssm start-session command to gain access to the target application server instance ID.

B. Ask each contractor to securely send their SSH public key. Add this public key to the application server\\'s SSH authorized-keys file. Instruct the contractors to use their private key to connect to the application server through SSH.

C. Ask each contractor to securely send their SSH public key. Use EC2 pairs to import their key. Update the application server\\'s SSH authorized_keys file. Instruct the contractors to use their private key to connect to the application server through SSH.

D. Create an IAM user for each contractor with programmatic access. Add each user to an IAM group that has a policy that allows the ssm:StartSession action. Instruct the contractors to install the AWS CLI and AWS Systems Manager Session Manager plugin, update their AWS credentials files with their access keys, and use the aws ssm start-session to gain access to the target application server instance ID.

Correct Answer: B

QUESTION 5

A company is running a custom-built application that processes records. All the components run on Amazon EC2

instances that run in an Auto Scaling group. Each record\\'s processing is a multistep sequential action that is compute-intensive. Each step is always completed in 5 minutes or less.

A limitation of the current system is that if any steps fail, the application has to reprocess the record from the beginning. The company wants to update the architecture so that the application must reprocess only the failed steps.

What is the MOST operationally efficient solution that meets these requirements?

A. Create a web application to write records to Amazon S3. Use S3 Event Notifications to publish to an Amazon Simple Notification Service (Amazon SNS) topic. Use an EC2 instance to poll Amazon SNS and start processing. Save intermediate results to Amazon S3 to pass on to the next step.

B. Perform the processing steps by using logic in the application. Convert the application code to run in a container. Use AWS Fargate to manage the container instances. Configure the container to invoke itself to pass the state from one step to the next.

C. Create a web application to pass records to an Amazon Kinesis data stream. Decouple the processing by using the Kinesis data stream and AWS Lambda functions.

D. Create a web application to pass records to AWS Step Functions. Decouple the processing into Step Functions tasks and AWS Lambda functions.

Correct Answer: D