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.leads4pass.com/dop-c02.html

100% Passing Guarantee
100% Money Back Assurance

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

- **Instant Download** After Purchase
- **100% Money Back** Guarantee
- **365 Days** Free Update
- **800,000+** Satisfied Customers

QUESTION 1

A company is using Amazon S3 buckets to store important documents. The company discovers that some S3 buckets are not encrypted. Currently, the company\\'s IAM users can create new S3 buckets without encryption. The company is implementing a new requirement that all S3 buckets must be encrypted.

A DevOps engineer must implement a solution to ensure that server-side encryption is enabled on all existing S3 buckets and all new S3 buckets. The encryption must be enabled on new S3 buckets as soon as the S3 buckets are created. The default encryption type must be 256-bit Advanced Encryption Standard (AES-256).

Which solution will meet these requirements?

A. Create an AWS Lambda function that is invoked periodically by an Amazon EventBridge scheduled rule. Program the Lambda function to scan all current S3 buckets for encryption status and to set AES-256 as the default encryption for any S3 bucket that does not have an encryption configuration.

B. Set up and activate the s3-bucket-server-side-encryption-enabled AWS Config managed rule. Configure the rule to use the AWS-EnableS3BucketEncryption AWS Systems Manager Automation runbook as the remediation action. Manually run the re-evaluation process to ensure that existing S3 buckets are compliant.

C. Create an AWS Lambda function that is invoked by an Amazon EventBridge event rule. Define the rule with an event pattern that matches the creation of new S3 buckets. Program the Lambda function to parse the EventBridge event, check the configuration of the S3 buckets from the event, and set AES-256 as the default encryption.

D. Configure an IAM policy that denies the s3:CreateBucket action if the s3:x-amz-server-side-encryption condition key has a value that is not AES-256. Create an IAM group for all the company\\'s IAM users. Associate the IAM policy with the IAM group.

Correct Answer: B

QUESTION 2

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 3

A DevOps engineer manages a company\\'s Amazon Elastic Container Service (Amazon ECS) cluster. The cluster runs on several Amazon EC2 instances that are in an Auto Scaling group. The DevOps engineer must implement a solution that logs and reviews all stopped tasks for errors.

Which solution will meet these requirements?

A. Create an Amazon EventBridge rule to capture task state changes. Send the event to Amazon CloudWatch Logs. Use CloudWatch Logs Insights to investigate stopped tasks.

B. Configure tasks to write log data in the embedded metric format. Store the logs in Amazon CloudWatch Logs. Monitor the ContainerInstanceCount metric for changes.

C. Configure the EC2 instances to store logs in Amazon CloudWatch Logs. Create a CloudWatch Contributor Insights rule that uses the EC2 instance log data. Use the Contributor Insights rule to investigate stopped tasks.

D. Configure an EC2 Auto Scaling lifecycle hook for the EC2_INSTANCE_TERMINATING scale-in event. Write the SystemEventLog file to Amazon S3. Use Amazon Athena to query the log file for errors.

Correct Answer: A

The best solution to log and review all stopped tasks for errors is to use Amazon EventBridge and Amazon CloudWatch Logs. Amazon EventBridge allows the DevOps engineer to create a rule that matches task state change events from Amazon ECS. The rule can then send the event data to Amazon CloudWatch Logs as the target. Amazon CloudWatch Logs can store and monitor the log data, and also provide CloudWatch Logs Insights, a feature that enables the DevOps engineer to interactively search and analyze the log data. Using CloudWatch Logs Insights, the DevOps engineer can filter and aggregate the log data based on various fields, such as cluster, task, container, and reason. This way, the DevOps engineer can easily identify and investigate the stopped tasks and their errors. The other options are not as effective or efficient as the solution in option A. Option B is not suitable because the embedded metric format is designed for custom metrics, not for logging task state changes. Option C is not feasible because the EC2 instances do not store the task state change events in their logs. Option D is not relevant because the EC2_INSTANCE_TERMINATING lifecycle hook is triggered when an EC2 instance is terminated by the Auto Scaling group, not when a task is stopped by Amazon ECS. References: : Creating a CloudWatch Events Rule That Triggers on an Event -Amazon Elastic Container Service : Sending and Receiving Events Between AWS Accounts -Amazon EventBridge : Working with Log Data -Amazon CloudWatch Logs : Analyzing Log Data with CloudWatch Logs Insights -Amazon CloudWatch Logs : Embedded Metric Format -Amazon CloudWatch : Amazon EC2 Auto Scaling Lifecycle Hooks -Amazon EC2 Auto Scaling

---

QUESTION 4

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 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.

B. 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.

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.

D. 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 CodeDeployDefault.AllatOnce 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.

Correct Answer: C

QUESTION 5

A company uses AWS Control Tower and AWS CloudFormation to manage its AWS accounts and to create AWS resources. The company requires all Amazon S3 buckets to be encrypted with AWS Key Management Service (AWS KMS) when the S3 buckets are created in a CloudFormation stack.

Which solution will meet this requirement?

A. Use AWS Organizations. Attach an SCP that denies the s3:PutObject permission if the request does not include an x-amz-server-side-encryption header that requests server-side encryption with AWS KMS keys (SSE-KMS).

B. Use AWS Control Tower with a multi-account environment. Configure and enable proactive AWS Control Tower controls on all OUs with CloudFormation hooks.

C. Use AWS Control Tower with a multi-account environment. Configure and enable detective AWS Control Tower controls on all OUs with CloudFormation hooks.

D. Use AWS Organizations. Create an AWS Config organizational rule to check whether a KMS encryption key is enabled for all S3 buckets. Deploy the rule. Create and apply an SCP to prevent users from stopping and deleting AWS Config across all AWS accounts,

Correct Answer: B

Proactive controls: Proactive controls are preventative measures that block actions violating defined policies before they occur. This ensures encryption gets applied automatically during S3 bucket creation within CloudFormation stacks. CloudFormation hooks: Hooks enable Control Tower to intercept and enforce policies on CloudFormation stack

operations, making it ideal for enforcing encryption during resource creation. Multi-account environment: Since the requirement applies across all accounts, Control Tower\\'s multi-account capabilities ensure consistent policy enforcement throughout the organization.

QUESTION 6

What is the proper (best practice) way to begin a playbook?

A. - hosts: all

B. ...

C. ###

D. --

Correct Answer: D

All YAML files can begin with `---\\' and end with `...\\' to indicate where YAML starts and ends. While this is optional, it is considered best practice.

Reference: http://docs.ansible.com/ansible/YAMLSyntax.html

QUESTION 7

A DevOps engineer manages a large commercial website that runs on Amazon EC2. The website uses Amazon Kinesis Data Streams to collect and process web togs. The DevOps engineer manages the Kinesis consumer application, which also runs on Amazon EC2.

Sudden increases of data cause the Kinesis consumer application to (all behind and the Kinesis data streams drop records before the records can be processed. The DevOps engineer must implement a solution to improve stream handling.

Which solution meets these requirements with the MOST operational efficiency?

A. Modify the Kinesis consumer application to store the logs durably in Amazon S3 Use Amazon EMR to process the data directly on Amazon S3 to derive customer insights Store the results in Amazon S3.

B. Horizontally scale the Kinesis consumer application by adding more EC2 instances based on the Amazon CloudWatch GetRecords IteratorAgeMilliseconds metric Increase the retention period of the Kinesis data streams.

C. Convert the Kinesis consumer application to run as an AWS Lambda function. Configure the Kinesis data streams as the event source for the Lambda function to process the data streams

D. Increase the number of shards in the Kinesis data streams to increase the overall throughput so that the consumer application processes the data faster.

Correct Answer: B

https://docs.aws.amazon.com/streams/latest/dev/monitoring-with-cloudwatch.html GetRecords.IteratorAgeMilliseconds -The age of the last record in all GetRecords calls made against a Kinesis stream, measured over the specified time period. Age is the difference between the current time and when the last record of the GetRecords call was written to the stream. The Minimum and Maximum statistics can be used to track the progress of Kinesis consumer applications.

A value of zero indicates that the records being read are completely caught up.

QUESTION 8

A company is running an application on Amazon EC2 instances in an Auto Scaling group. Recently an issue occurred that prevented EC2 instances from launching successfully and it took several hours for the support team to discover the issue. The support team wants to be notified by email whenever an EC2 instance does not start successfully.

Which action will accomplish this?

A. Add a health check to the Auto Scaling group to invoke an AWS Lambda function whenever an instance status is impaired.

B. Configure the Auto Scaling group to send a notification to an Amazon SNS topic whenever a failed instance launch occurs.

C. Create an Amazon CloudWatch alarm that invokes an AWS Lambda function when a failed Attachinstances Auto Scaling API call is made.

D. Create a status check alarm on Amazon EC2 to send a notification to an Amazon SNS topic whenever a status check fail occurs.

Correct Answer: B

https://docs.aws.amazon.com/autoscaling/ec2/userguide/ASGettingNotifications.html#auto-scaling-sns-notifications

QUESTION 9

Which statement is true about configuring proxy support for Amazon Inspector agent on Linuxbased systems?

A. Amazon Inspector proxy support on Linux-based systems is achieved through installing proxyenabled version of the agent which comes with pre-configured files that you need to edit to match your environment.

B. Amazon Inspector agent does NOT support the use of proxy on Linux-based systems.

C. Amazon Inspector proxy configuration on Linux-based system is included in awsagent.env file under /etc/init.d/

D. Amazon Inspector agent proxy settings on Linux-based systems are configured through WinHTTP proxy.

Correct Answer: C

To install an AWS agent on an EC2 instance that uses a proxy server Create a file called awsagent.env and save it in the /etc/init.d/ directory. Edit awsagent.env to include these environment variables in the following format: export https_proxy=https://hostname:port export http_proxy=http://hostname:port export no_proxy= 123.456.789.111

Reference: https://docs.aws.amazon.com/inspector/latest/userguide/inspector_agents-on-win.html#inspectoragent-proxy

QUESTION 10

A development team is building an ecommerce application and is using Amazon Simple Notification Service (Amazon

SNS) to send order messages to multiple endpoints. One of the endpoints is an external HTTP endpoint that is not always available. The development team needs to receive a notification if an order message is not delivered to the HTTP endpoint.

What should a DevOps engineer do to meet these requirements?

A. Create an Amazon Simple Queue Service (Amazon SQS) queue. On the SNS topic, configure a redrive policy that sends undelivered messages to the SQS queue. Create an Amazon CloudWatch alarm for the new SQS queue to notify the development team when messages are delivered to the queue.

B. Create an Amazon Simple Queue Service (Amazon SQS) queue. On the HTTP endpoint subscription of the SNS topic, configure a redrive policy that sends undelivered messages to the SQS queue. Create an Amazon CloudWatch alarm for the new SQS queue to notify the development team when messages are delivered to the queue.

C. On the SNS topic, configure an HTTPS delivery policy that will retry delivery until the order message is delivered successfully. Configure the backoffFunction parameter in the policy to notify the development team when a message cannot be delivered within the set constraints.

D. On the HTTP endpoint subscription of the SNS topic, configure an HTTPS delivery policy that will retry delivery until the order message is delivered successfully. Configure the backoffFunction parameter in the policy to notify the development team when a message cannot be delivered within the set constraints.

Correct Answer: C

QUESTION 11

A DevOps engineer is working on a data archival project that requires the migration of on-premises data to an Amazon S3 bucket. The DevOps engineer develops a script that incrementally archives on-premises data that is older than 1 month to Amazon S3. Data that is transferred to Amazon S3 is deleted from the on-premises location The script uses the S3 PutObject operation.

During a code review the DevOps engineer notices that the script does not verity whether the data was successfully copied to Amazon S3. The DevOps engineer must update the script to ensure that data is not corrupted during transmission.

The script must use MD5 checksums to verify data integrity before the on-premises data is deleted. Which solutions for the script will meet these requirements\\'? (Select TWO.)

A. Check the returned response for the Versioned Compare the returned Versioned against the MD5 checksum.

B. Include the MD5 checksum within the Content-MD5 parameter. Check the operation call\\'s return status to find out if an error was returned.

C. Include the checksum digest within the tagging parameter as a URL query parameter.

D. Check the returned response for the ETag. Compare the returned ETag against the MD5 checksum.

E. Include the checksum digest within the Metadata parameter as a name-value pair After upload use the S3 HeadObject operation to retrieve metadata from the object.

Correct Answer: BD

https://docs.aws.amazon.com/AmazonS3/latest/userguide/checking-object-integrity.html

QUESTION 12

A company\\'s development team uses AVMS Cloud Formation to deploy its application resources The team must use for an changes to the environment The team cannot use AWS Management Console or the AWS CLI to make manual changes directly.

The team uses a developer IAM role to access the environment The role is configured with the Admnistratoraccess managed policy. The company has created a new Cloudformationdeployment IAM role that has the following policy.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "elasticloadbalancing:*",
                "lambda:*",
                "dynamodb:*"
            ],
            "Resource": "*"
        }
    ]
}
```

The company wants ensure that only CloudFormation can use the new role. The development team cannot make any manual changes to the deployed resources.

Which combination of steps meet these requirements? (Select THREE.)

A. Remove the AdministratorAccess policy. Assign the ReadOnlyAccess managed IAM policy to the developer role. Instruct the developers to use the CloudFormationDeployment role as a CloudFormation service role when the developers deploy new stacks.

B. Update the trust of CloudFormationDeployment role to allow the developer IAM role to assume the CloudFormationDepoyment role.

C. Configure the IAM to be to get and pass the CloudFormationDeployment role if cloudformation actions for resources,

D. Update the trust Of the CloudFormationDepoyment role to anow the cloudformation.amazonaws.com AWS principal to perform the iam:AssumeR01e action

E. Remove me Administratoraccess policy. Assign the ReadOnly/Access managed IAM policy to the developer role Instruct the developers to assume the CloudFormatondeployment role when the developers new stacks

F. Add an IAM policy to CloudFormationDeplyment to allow cloudformation * on an Add a policy that allows the iam.PassR01e action for ARN of if iam PassedT0Service equal cloudformation.amazonaws.com

Correct Answer: ADF

A comprehensive and detailed explanation is: Option A is correct because removing the AdministratorAccess policy and assigning the ReadOnlyAccess managed IAM policy to the developer role is a valid way to prevent the developers from making any manual changes to the deployed resources. The AdministratorAccess policy grants full access to all AWS resources and actions, which is not necessary for the developers. The ReadOnlyAccess policy grants read-only access to most AWS resources and actions, which is sufficient for the developers to view the status of their stacks. Instructing the developers to use the CloudFormationDeployment role as a CloudFormation service role when they deploy new stacks is also a valid way to ensure that only CloudFormation can use the new role. A CloudFormation service role is an IAM role that allows CloudFormation to make calls to resources in a stack on behalf of the user1. The user can specify a service role when they create or update a stack, and CloudFormation will use that role\\'s credentials for all operations that are performed on that stack1. Option B is incorrect because updating the trust of CloudFormationDeployment role to allow the developer IAM role to assume the CloudFormationDeployment role is not a valid solution. This would allow the developers to manually assume the CloudFormationDeployment role and perform actions on the deployed resources, which is not what the company wants. The trust of CloudFormationDeployment role should only allow the cloudformation.amazonaws.com AWS principal to assume the role, as in option D. Option C is incorrect because configuring the IAM user to be able to get and pass the CloudFormationDeployment role if cloudformation actions for resources is not a valid solution. This would allow the developers to manually pass the CloudFormationDeployment role to other services or resources, which is not what the company wants. The IAM user should only be able to pass the CloudFormationDeployment role as a service role when they create or update a stack with CloudFormation, as in option A. Option D is correct because updating the trust of CloudFormationDeployment role to allow the cloudformation.amazonaws.com AWS principal to perform the iam:AssumeRole action is a valid solution. This allows CloudFormation to assume the CloudFormationDeployment role and access resources in other services on behalf of the user2. The trust policy of an IAM role defines which entities can assume the role2. By specifying cloudformation.amazonaws.com as the principal, you grant permission only to CloudFormation to assume this role. Option E is incorrect because instructing the developers to assume the CloudFormationDeployment role when they deploy new stacks is not a valid solution. This would allow the developers to manually assume the CloudFormationDeployment role and perform actions on the deployed resources, which is not what the company wants. The developers should only use the CloudFormationDeployment role as a service role when they deploy new stacks with CloudFormation, as in option A. Option F is correct because adding an IAM policy to CloudFormationDeployment that allows cloudformation:* on all resources and adding a policy that allows the iam:PassRole action for ARN of CloudFormationDeployment if iam:PassedToService equals cloudformation.amazonaws.com are valid solutions. The first policy grants permission for CloudFormationDeployment to perform any action with any resource using cloudformation.amazonaws.com as a service principal3. The second policy grants permission for passing this role only if it is passed by cloudformation.amazonaws.com as a service principal4. This ensures that only CloudFormation can use this role. References:

1: AWS CloudFormation service roles

2: How to use trust policies with IAM roles

3: AWS::IAM::Policy

4: IAM: Pass an IAM role to a specific AWS service

QUESTION 13

A company runs its container workloads in AWS App Runner. A DevOps engineer manages the company\\'s container repository in Amazon Elastic Container Registry (Amazon ECR).

The DevOps engineer must implement a solution that continuously monitors the container repository. The solution must create a new container image when the solution detects an operating system vulnerability or language package vulnerability.

Which solution will meet these requirements?

A. Use EC2 Image Builder to create a container image pipeline. Use Amazon ECR as the target repository. Turn on

enhanced scanning on the ECR repository. Create an Amazon EventBridge rule to capture an Inspector2 finding event. Use the event to invoke the image pipeline. Re-upload the container to the repository.

B. Use EC2 Image Builder to create a container image pipeline. Use Amazon ECR as the target repository. Enable Amazon GuardDuty Malware Protection on the container workload. Create an Amazon EventBridge rule to capture a GuardDuty finding event. Use the event to invoke the image pipeline.

C. Create an AWS CodeBuild project to create a container image. Use Amazon ECR as the target repository. Turn on basic scanning on the repository. Create an Amazon EventBridge rule to capture an ECR image action event. Use the event to invoke the CodeBuild project. Re-upload the container to the repository.

D. Create an AWS CodeBuild project to create a container image. Use Amazon ECR as the target repository. Configure AWS Systems Manager Compliance to scan all managed nodes. Create an Amazon EventBridge rule to capture a configuration compliance state change event. Use the event to invoke the CodeBuild project.

Correct Answer: A

The solution that meets the requirements is to use EC2 Image Builder to create a container image pipeline, use Amazon ECR as the target repository, turn on enhanced scanning on the ECR repository, create an Amazon EventBridge rule to capture an Inspector2 finding event, and use the event to invoke the image pipeline. Re-upload the container to the repository. This solution will continuously monitor the container repository for vulnerabilities using enhanced scanning, which is a feature of Amazon ECR that provides detailed information and guidance on how to fix security issues found in your container images. Enhanced scanning uses Inspector2, a security assessment service that integrates with Amazon ECR and generates findings for any vulnerabilities detected in your images. You can use Amazon EventBridge to create a rule that triggers an action when an Inspector2 finding event occurs. The action can be to invoke an EC2 Image Builder pipeline, which is a service that automates the creation of container images. The pipeline can use the latest patches and updates to build a new container image and upload it to the same ECR repository, replacing the vulnerable image. The other options are not correct because they do not meet all the requirements or use services that are not relevant for the scenario. Option B is not correct because it uses Amazon GuardDuty Malware Protection, which is a feature of GuardDuty that detects malicious activity and unauthorized behavior on your AWS accounts and resources. GuardDuty does not scan container images for vulnerabilities, nor does it integrate with Amazon ECR or EC2 Image Builder. Option C is not correct because it uses basic scanning on the ECR repository, which only provides a summary of the vulnerabilities found in your container images. Basic scanning does not use Inspector2 or generate findings that can be captured by Amazon EventBridge. Moreover, basic scanning does not provide guidance on how to fix the vulnerabilities. Option D is not correct because it uses AWS Systems Manager Compliance, which is a feature of Systems Manager that helps you monitor and manage the compliance status of your AWS resources based on AWS Config rules and AWS Security Hub standards. Systems Manager Compliance does not scan container images for vulnerabilities, nor does it integrate with Amazon ECR or EC2 Image Builder.

QUESTION 14

You have just recently deployed an application on EC2 instances behind an ELB. After a couple of weeks, customers are complaining on receiving errors from the application. You want to diagnose the errors and are trying to get errors from the ELB access logs. But the ELB access logs are empty. What is the reason for this.

A. You do not have the appropriate permissions to access the logs

B. You do not have your CloudWatch metrics correctly configured

C. ELB Access logs are only available for a maximum of one week

D. Access logging is an optional feature of Elastic Load Balancing that is disabled by default

Correct Answer: D

Clastic Load Balancing provides access logs that capture detailed information about requests sent to your load balancer. Cach log contains information such as the time the request was received, the client\\'s IP address, latencies, request paths, and server responses. You can use these access logs to analyze traffic patterns and to troubleshoot issues. Access logging is an optional feature of Elastic Load Balancing that is disabled by default. After you enable access logging for your load balancer. Clastic Load Balancing captures the logs and stores them in the Amazon S3 bucket that you specify. You can disable access logging at any time.

QUESTION 15

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 a new DynamoDB table in the new Region with cross-Region replication enabled.

B. Create new ALB and Auto Scaling group global resources and configure the new ALB to direct traffic to the new Auto Scaling group.

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 Amazon Route 53 records, health checks, and latency-based routing policies to route to the ALB.

E. Create Amazon Route 53 aliases, health checks, and failover routing policies to route to the ALB.

F. Convert the DynamoDB table to a global table.

Correct Answer: CDF

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. This will allow users in the new Region to access the application with lower latency by reducing the network hops between the user and the application servers.

D. Create Amazon Route 53 records, health checks, and latency-based routing policies to route to the ALB. This will enable Route 53 to route user traffic to the nearest healthy ALB, based on the latency between the user and the ALBs.
F. Convert the DynamoDB table to a global table. This will enable reads and writes to the table in both Regions with low latency, improving the overall response time of the application