

Exam Questions DOP-C02

AWS Certified DevOps Engineer - Professional

<https://www.2passeasy.com/dumps/DOP-C02/>



NEW QUESTION 1

A company is using an organization in AWS Organizations to manage multiple AWS accounts. The company's development team wants to use AWS Lambda functions to meet resiliency requirements and is rewriting all applications to work with Lambda functions that are deployed in a VPC. The development team is using Amazon Elastic File System (Amazon EFS) as shared storage in Account A in the organization.

The company wants to continue to use Amazon EFS with Lambda. Company policy requires all serverless projects to be deployed in Account B.

A DevOps engineer needs to reconfigure an existing EFS file system to allow Lambda functions to access the data through an existing EFS access point.

Which combination of steps should the DevOps engineer take to meet these requirements? (Select THREE.)

- A. Update the EFS file system policy to provide Account B with access to mount and write to the EFS file system in Account A.
- B. Create SCPs to set permission guardrails with fine-grained control for Amazon EFS.
- C. Create a new EFS file system in Account B Use AWS Database Migration Service (AWS DMS) to keep data from Account A and Account B synchronized.
- D. Update the Lambda execution roles with permission to access the VPC and the EFS file system.
- E. Create a VPC peering connection to connect Account A to Account B.
- F. Configure the Lambda functions in Account B to assume an existing IAM role in Account A.

Answer: AEF

Explanation:

A Lambda function in one account can mount a file system in a different account. For this scenario, you configure VPC peering between the function VPC and the file system VPC. <https://docs.aws.amazon.com/lambda/latest/dg/services-efs.html>

<https://aws.amazon.com/ru/blogs/storage/mount-amazon-efs-file-systems-cross-account-from-amazon-eks/>

* 1. Need to update the file system policy on EFS to allow mounting the file system into Account B.

File System Policy

\$ cat file-system-policy.json

```
{
  "Statement": [
    {
      "Effect": "Allow", "Action": [
        "elasticfilesystem:ClientMount", "elasticfilesystem:ClientWrite"
      ],
      "Principal": {
        "AWS": "arn:aws:iam::<aws-account-id-A>:root" # Replace with AWS account ID of EKS cluster
      }
    }
  ]
}
```

* 2. Need VPC peering between Account A and Account B as the pre-requisite

* 3. Need to assume cross-account IAM role to describe the mounts so that a specific mount can be chosen.

NEW QUESTION 2

A company uses a single AWS account to test applications on Amazon EC2 instances. The company has turned on AWS Config in the AWS account and has activated the restricted-ssh AWS Config managed rule.

The company needs an automated monitoring solution that will provide a customized notification in real time if any security group in the account is not compliant with the restricted-ssh rule. The customized notification must contain the name and ID of the noncompliant security group.

A DevOps engineer creates an Amazon Simple Notification Service (Amazon SNS) topic in the account and subscribes the appropriate personnel to the topic.

What should the DevOps engineer do next to meet these requirements?

- A. Create an Amazon EventBridge rule that matches an AWS Config evaluation result of NON_COMPLIANT for the restricted-ssh rule
- B. Configure an input transformer for the EventBridge rule Configure the EventBridge rule to publish a notification to the SNS topic.
- C. Configure AWS Config to send all evaluation results for the restricted-ssh rule to the SNS topic. Configure a filter policy on the SNS topic to send only notifications that contain the text of NON_COMPLIANT in the notification to subscribers.
- D. Create an Amazon EventBridge rule that matches an AWS Config evaluation result of NON_COMPLIANT for the restricted-ssh rule Configure the EventBridge rule to invoke AWS Systems Manager Run Command on the SNS topic to customize a notification and to publish the notification to the SNS topic
- E. Create an Amazon EventBridge rule that matches all AWS Config evaluation results of NON_COMPLIANT Configure an input transformer for the restricted-ssh rule Configure the EventBridge rule to publish a notification to the SNS topic.

Answer: A

Explanation:

Create an Amazon EventBridge (Amazon CloudWatch Events) rule that matches an AWS Config evaluation result of NON_COMPLIANT for the restricted-ssh rule. Configure an input transformer for the EventBridge (CloudWatch Events) rule. Configure the EventBridge (CloudWatch Events) rule to publish a notification to the SNS topic. This approach uses Amazon EventBridge (previously known as Amazon CloudWatch Events) to filter AWS Config evaluation results based on the restricted-ssh rule and its compliance status (NON_COMPLIANT). An input transformer can be used to customize the information contained in the notification, such as the name and ID of the noncompliant security group. The EventBridge (CloudWatch Events) rule can then be configured to publish a notification to the SNS topic, which will notify the appropriate personnel in real-time.

NEW QUESTION 3

A company has a legacy application A DevOps engineer needs to automate the process of building the deployable artifact for the legacy application. The solution must store the deployable artifact in an existing Amazon S3 bucket for future deployments to reference

Which solution will meet these requirements in the MOST operationally efficient way?

- A. Create a custom Docker image that contains all the dependencies for the legacy application Store the custom Docker image in a new Amazon Elastic Container Registry (Amazon ECR) repository Configure a new AWS CodeBuild project to use the custom Docker image to build the deployable artifact and to save the artifact to the S3 bucket.
- B. Launch a new Amazon EC2 instance Install all the dependencies (or the legacy application) on the EC2 instance Use the EC2 instance to build the deployable artifact and to save the artifact to the S3 bucket.
- C. Create a custom EC2 Image Builder image Install all the dependencies for the legacy application on the image Launch a new Amazon EC2 instance from the image Use the new EC2 instance to build the deployable artifact and to save the artifact to the S3 bucket.
- D. Create an Amazon Elastic Kubernetes Service (Amazon EKS) cluster with an AWS Fargate profile that runs in multiple Availability Zones Create a custom

Docker image that contains all the dependencies for the legacy application Store the custom Docker image in a new Amazon Elastic Container Registry (Amazon ECR) repository Use the custom Docker image inside the EKS cluster to build the deployable artifact and to save the artifact to the S3 bucket.

Answer: A

Explanation:

This approach is the most operationally efficient because it leverages the benefits of containerization, such as isolation and reproducibility, as well as AWS managed services. AWS CodeBuild is a fully managed build service that can compile your source code, run tests, and produce deployable software packages. By using a custom Docker image that includes all dependencies, you can ensure that the environment in which your code is built is consistent. Using Amazon ECR to store Docker images lets you easily deploy the images to any environment. Also, you can directly upload the build artifacts to Amazon S3 from AWS CodeBuild, which is beneficial for version control and archival purposes.

NEW QUESTION 4

A company is using an AWS CodeBuild project to build and package an application. The packages are copied to a shared Amazon S3 bucket before being deployed across multiple AWS accounts.

The buildspec.yml file contains the following:

```
version: 0.2
phases:
  build:
    commands:
      - go build -o myapp
  post_build:
    commands:
      - aws s3 cp --acl authenticated-read myapp s3://artifacts/
```

The DevOps engineer has noticed that anybody with an AWS account is able to download the artifacts. What steps should the DevOps engineer take to stop this?

- A. Modify the post_build command to use --acl public-read and configure a bucket policy that grants read access to the relevant AWS accounts only.
- B. Configure a default ACL for the S3 bucket that defines the set of authenticated users as the relevant AWS accounts only and grants read-only access.
- C. Create an S3 bucket policy that grants read access to the relevant AWS accounts and denies read access to the principal ""*"".
- D. Modify the post_build command to remove --acl authenticated-read and configure a bucket policy that allows read access to the relevant AWS accounts only.

Answer: D

Explanation:

When setting the flag authenticated-read in the command line, the owner gets FULL_CONTROL. The AuthenticatedUsers group (Anyone with an AWS account) gets READ access. Reference: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/acl-overview.html>

NEW QUESTION 5

An application runs on Amazon EC2 instances behind an Application Load Balancer (ALB). A DevOps engineer is using AWS CodeDeploy to release a new version. The deployment fails during the AllowTraffic lifecycle event, but a cause for the failure is not indicated in the deployment logs.

What would cause this?

- A. The appspec
- B. yml file contains an invalid script that runs in the AllowTraffic lifecycle hook.
- C. The user who initiated the deployment does not have the necessary permissions to interact with the ALB.
- D. The health checks specified for the ALB target group are misconfigured.
- E. The CodeDeploy agent was not installed in the EC2 instances that are part of the ALB target group.

Answer: C

Explanation:

This failure is typically due to incorrectly configured health checks in Elastic Load Balancing for the Classic Load Balancer, Application Load Balancer, or Network Load Balancer used to manage traffic for the deployment group. To resolve the issue, review and correct any errors in the health check configuration for the load balancer.

<https://docs.aws.amazon.com/codedeploy/latest/userguide/troubleshooting-deployments.html#troubleshooting-d>

NEW QUESTION 6

A company has an application that is using a MySQL-compatible Amazon Aurora Multi-AZ DB cluster as the database. A cross-Region read replica has been created for disaster recovery purposes. A DevOps engineer wants to automate the promotion of the replica so it becomes the primary database instance in the event of a failure.

Which solution will accomplish this?

- A. Configure a latency-based Amazon Route 53 CNAME with health checks so it points to both the primary and replica endpoint
- B. Subscribe an Amazon SNS topic to Amazon RDS failure notifications from AWS CloudTrail and use that topic to invoke an AWS Lambda function that will promote the replica instance as the primary.
- C. Create an Aurora custom endpoint to point to the primary database instance
- D. Configure the application to use this endpoint
- E. Configure AWS CloudTrail to run an AWS Lambda function to promote the replica instance and modify the custom endpoint to point to the newly promoted instance.
- F. Create an AWS Lambda function to modify the application's AWS CloudFormation template to promote the replica, apply the template to update the stack, and point the application to the newly promoted instance
- G. Create an Amazon CloudWatch alarm to invoke this Lambda function after the failure event occurs.
- H. Store the Aurora endpoint in AWS Systems Manager Parameter Store
- I. Create an Amazon EventBridge event that detects the database failure and runs an AWS Lambda function to promote the replica instance and update the endpoint URL stored in AWS Systems Manager Parameter Store
- J. Code the application to reload the endpoint from Parameter Store if a database connection fails.

Answer: D

Explanation:

EventBridge is needed to detect the database failure. Lambda is needed to promote the replica as it's in another Region (manual promotion, otherwise). Storing and updating the endpoint in Parameter store is important in updating the application. Look at High Availability section of Aurora FAQ:
<https://aws.amazon.com/rds/aurora/faqs/>

NEW QUESTION 7

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 Log
- B. Configure the VPC flow log to capture accepted traffic and to send the data to the log group
- C. Create an Amazon CloudWatch metric filter for IP addresses on the deny list
- D. Create a CloudWatch alarm with the metric filter as input
- E. 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.
- F. Create an Amazon S3 bucket for log file
- G. Configure the VPC flow log to capture all traffic and to send the data to the S3 bucket
- H. Configure Amazon Athena to return all log files in the S3 bucket for IP addresses on the deny list
- I. Configure Amazon QuickSight to accept data from Athena and to publish the data as a dashboard that the security team can access
- J. Create a threshold alert of 1 for successful access
- K. Configure the alert to automatically notify the security team as frequently as possible when the alert threshold is met.
- L. Create an Amazon S3 bucket for log file
- M. Configure the VPC flow log to capture accepted traffic and to send the data to the S3 bucket
- N. Configure an Amazon OpenSearch Service cluster and domain for the log file
- O. 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
- P. Schedule the Lambda function to run every 5 minutes
- Q. 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.
- R. Create a log group in Amazon CloudWatch Log
- S. 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
- T. Deploy an Amazon Athena CloudWatch connector in AWS Lambda
- U. Connect the connector to the log group
- V. 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
- W. 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.

Answer: A

NEW QUESTION 8

A company requires an RPO of 2 hours and an RTO of 10 minutes for its data and application at all times. An application uses a MySQL database and Amazon EC2 web servers. The development team needs a strategy for failover and disaster recovery.

Which combination of deployment strategies will meet these requirements? (Select TWO.)

- A. Create an Amazon Aurora cluster in one Availability Zone across multiple Regions as the data store. Use Aurora's automatic recovery capabilities in the event of a disaster.
- B. Create an Amazon Aurora global database in two Regions as the data store.
- C. In the event of a failure, promote the secondary Region as the primary for the application.
- D. Create an Amazon Aurora multi-master cluster across multiple Regions as the data store.
- E. Use a Network Load Balancer to balance the database traffic in different Regions.
- F. Set up the application in two Regions and use Amazon Route 53 failover-based routing that points to the Application Load Balancers in both Regions.
- G. Use health checks to determine the availability in a given Region.
- H. Use Auto Scaling groups in each Region to adjust capacity based on demand.
- I. Set up the application in two Regions and use a multi-Region Auto Scaling group behind Application Load Balancers to manage the capacity based on demand.
- J. In the event of a disaster, adjust the Auto Scaling group's desired instance count to increase baseline capacity in the failover Region.

Answer: BD

NEW QUESTION 9

A company is adopting AWS CodeDeploy to automate its application deployments for a Java-Apache Tomcat application with an Apache Webserver. The development team started with a proof of concept, created a deployment group for a developer environment, and performed functional tests within the application. After completion, the team will create additional deployment groups for staging and production.

The current log level is configured within the Apache settings, but the team wants to change this configuration dynamically when the deployment occurs, so that they can set different log level configurations depending on the deployment group without having a different application revision for each group.

How can these requirements be met with the LEAST management overhead and without requiring different script versions for each deployment group?

- A. Tag the Amazon EC2 instances depending on the deployment group
- B. Then place a script into the application revision that calls the metadata service and the EC2 API to identify which deployment group the instance is part of
- C. Use this information to configure the log level setting
- D. Reference the script as part of the AfterInstall lifecycle hook in the appspec.yml file.
- E. Create a script that uses the CodeDeploy environment variable DEPLOYMENT_GROUP_NAME to identify which deployment group the instance is part of
- F. Use this information to configure the log level setting
- G. Reference this script as part of the BeforeInstall lifecycle hook in the appspec.yml file.

- H. Create a CodeDeploy custom environment variable for each environmen
- I. Then place a script into the application revision that checks this environment variable to identify which deployment group the instance is part o
- J. Use this information to configure the log level setting
- K. Reference this script as part of the ValidateService lifecycle hook in the appspec.yml file.
- L. Create a script that uses the CodeDeploy environment variable DEPLOYMENT_GROUP_ID to identify which deployment group the instance is part of to configure the log level setting
- M. Reference this script as part of the Install lifecycle hook in the appspec.yml file.

Answer: B

Explanation:

The following are the steps that the company can take to change the log level dynamically when the deployment occurs:

- > Create a script that uses the CodeDeploy environment variable DEPLOYMENT_GROUP_NAME to identify which deployment group the instance is part of.
- > Use this information to configure the log level settings.
- > Reference this script as part of the BeforeInstall lifecycle hook in the appspec.yml file.

The DEPLOYMENT_GROUP_NAME environment variable is automatically set by CodeDeploy when the deployment is triggered. This means that the script does not need to call the metadata service or the EC2 API to identify the deployment group.

This solution is the least complex and requires the least management overhead. It also does not require different script versions for each deployment group.

The following are the reasons why the other options are not correct:

- > Option A is incorrect because it would require tagging the Amazon EC2 instances, which would be a manual and time-consuming process.
- > Option C is incorrect because it would require creating a custom environment variable for each environment. This would be a complex and error-prone process.
- > Option D is incorrect because it would use the DEPLOYMENT_GROUP_ID environment variable.

However, this variable is not automatically set by CodeDeploy, so the script would need to call the metadata service or the EC2 API to get the deployment group ID. This would add complexity and overhead to the solution.

NEW QUESTION 10

A DevOps team is merging code revisions for an application that uses an Amazon RDS Multi-AZ DB cluster for its production database. The DevOps team uses continuous integration to periodically verify that the application works. The DevOps team needs to test the changes before the changes are deployed to the production database.

Which solution will meet these requirements'?

- A. Use a buildspec file in AWS CodeBuild to restore the DB cluster from a snapshot of the production database run integration tests, and drop the restored database after verification.
- B. Deploy the application to productio
- C. Configure an audit log of data control language (DCL) operations to capture database activities to perform if verification fails.
- D. Create a snapshot of the DB duster before deploying the application Use the Update requires Replacement property on the DB instance in AWS CloudFormation to deploy the application and apply the changes.
- E. Ensure that the DB cluster is a Multi-AZ deploymen
- F. Deploy the application with the update
- G. Fail over to the standby instance if verification fails.

Answer: A

Explanation:

This solution will meet the requirements because it will create a temporary copy of the production database using a snapshot, run the integration tests on the copy, and delete the copy after the tests are done. This way, the production database will not be affected by the code revisions, and the DevOps team can test the changes before deploying them to production. A buildspec file is a YAML file that contains the commands and settings that CodeBuild uses to run a build1. The buildspec file can specify the steps to restore the DB cluster from a snapshot, run the integration tests, and drop the restored database2

NEW QUESTION 10

A DevOps engineer at a company is supporting an AWS environment in which all users use AWS IAM Identity Center (AWS Single Sign-On). The company wants to immediately disable credentials of any new IAM user and wants the security team to receive a notification.

Which combination of steps should the DevOps engineer take to meet these requirements? (Choose three.)

- A. Create an Amazon EventBridge rule that reacts to an IAM CreateUser API call in AWS CloudTrail.
- B. Create an Amazon EventBridge rule that reacts to an IAM GetLoginProfile API call in AWS CloudTrail.
- C. Create an AWS Lambda function that is a target of the EventBridge rul
- D. Configure the Lambda function to disable any access keys and delete the login profiles that are associated with the IAM user.
- E. Create an AWS Lambda function that is a target of the EventBridge rul
- F. Configure the Lambda function to delete the login profiles that are associated with the IAM user.
- G. Create an Amazon Simple Notification Service (Amazon SNS) topic that is a target of the EventBridge rul
- H. Subscribe the security team's group email address to the topic.
- I. Create an Amazon Simple Queue Service (Amazon SQS) queue that is a target of the Lambda function.Subscribe the security team's group email address to the queue.

Answer: ACE

NEW QUESTION 12

A company uses AWS Storage Gateway in file gateway mode in front of an Amazon S3 bucket that is used by multiple resources. In the morning when business begins, users do not see the objects processed by a third party the previous evening. When a DevOps engineer looks directly at the S3 bucket, the data is there, but it is missing in Storage Gateway.

Which solution ensures that all the updated third-party files are available in the morning?

- A. Configure a nightly Amazon EventBridge event to invoke an AWS Lambda function to run the RefreshCache command for Storage Gateway.
- B. Instruct the third party to put data into the S3 bucket using AWS Transfer for SFTP.
- C. Modify Storage Gateway to run in volume gateway mode.
- D. Use S3 Same-Region Replication to replicate any changes made directly in the S3 bucket to Storage Gateway.

Answer: A

Explanation:

https://docs.aws.amazon.com/storagegateway/latest/APIReference/API_RefreshCache.html " It only updates the cached inventory to reflect changes in the inventory of the objects in the S3 bucket. This operation is only supported in the S3 File Gateway types."

NEW QUESTION 16

An ecommerce company is receiving reports that its order history page is experiencing delays in reflecting the processing status of orders. The order processing system consists of an AWS Lambda function that uses reserved concurrency. The Lambda function processes order messages from an Amazon Simple Queue Service (Amazon SQS) queue and inserts processed orders into an Amazon DynamoDB table. The DynamoDB table has auto scaling enabled for read and write capacity.

Which actions should a DevOps engineer take to resolve this delay? (Choose two.)

- A. Check the ApproximateAgeOfOldestMessage metric for the SQS queue
- B. Increase the Lambda function concurrency limit.
- C. Check the ApproximateAgeOfOldestMessage metric for the SQS queue. Configure a redrive policy on the SQS queue.
- D. Check the NumberOfMessagesSent metric for the SQS queue
- E. Increase the SQS queue visibility timeout.
- F. Check the WriteThrottleEvents metric for the DynamoDB table
- G. Increase the maximum write capacity units (WCUs) for the table's scaling policy.
- H. Check the Throttles metric for the Lambda function
- I. Increase the Lambda function timeout.

Answer: AD

Explanation:

A: If the ApproximateAgeOfOldestMessages indicate that orders are remaining in the SQS queue for longer than expected, the reserved concurrency limit may be set too small to keep up with the number of orders entering the queue and is being throttled. D: The DynamoDB table is using Auto Scaling. With Auto Scaling, you create a scaling policy that specifies whether you want to scale read capacity or write capacity (or both), and the minimum and maximum provisioned capacity unit settings for the table. The ThrottledWriteRequests metric will indicate if there is a throttling issue on the DynamoDB table, which can be resolved by increasing the maximum write capacity units for the table's Auto Scaling policy. <https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/AutoScaling.html>

NEW QUESTION 18

A company has migrated its container-based applications to Amazon EKS and want to establish automated email notifications. The notifications sent to each email address are for specific activities related to EKS components. The solution will include Amazon SNS topics and an AWS Lambda function to evaluate incoming log events and publish messages to the correct SNS topic.

Which logging solution will support these requirements?

- A. Enable Amazon CloudWatch Logs to log the EKS component
- B. Create a CloudWatch subscription filter for each component with Lambda as the subscription feed destination.
- C. Enable Amazon CloudWatch Logs to log the EKS component
- D. Create CloudWatch Logs Insights queries linked to Amazon EventBridge events that invoke Lambda.
- E. Enable Amazon S3 logging for the EKS component
- F. Configure an Amazon CloudWatch subscription filter for each component with Lambda as the subscription feed destination.
- G. Enable Amazon S3 logging for the EKS component
- H. Configure S3 PUT Object event notifications with AWS Lambda as the destination.

Answer: A

Explanation:

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/SubscriptionFilters.html#LambdaFunctionExample>
<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/SubscriptionFilters.html>

NEW QUESTION 20

A company manages a web application that runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The EC2 instances run in an Auto Scaling group across multiple Availability Zones. The application uses an Amazon RDS for MySQL DB instance to store the data. The company has configured Amazon Route 53 with an alias record that points to the ALB.

A new company guideline requires a geographically isolated disaster recovery (DR) site with an RTO of 4 hours and an RPO of 15 minutes.

Which DR strategy will meet these requirements with the LEAST change to the application stack?

- A. Launch a replica environment of everything except Amazon RDS in a different Availability Zone. Create an RDS read replica in the new Availability Zone: and configure the new stack to point to the local RDS DB instance
- B. Add the new stack to the Route 53 record set by using a health check to configure a failover routing policy.
- C. Launch a replica environment of everything except Amazon RDS in a different AW
- D. Region Create an RDS read replica in the new Region and configure the new stack to point to the local RDS DB instance
- E. Add the new stack to the Route 53 record set by using a health check to configure a latency routing policy.
- F. Launch a replica environment of everything except Amazon RDS in a different AWS Region
- G. In the event of an outage copy and restore the latest RDS snapshot from the primary
- H. Region to the DR Region Adjust the Route 53 record set to point to the ALB in the DR Region.
- I. Launch a replica environment of everything except Amazon RDS in a different AWS Region
- J. Create an RDS read replica in the new Region and configure the new environment to point to the local RDS DB instance
- K. Add the new stack to the Route 53 record set by using a health check to configure a failover routing policy
- L. In the event of an outage promote the read replica to primary.

Answer: D

NEW QUESTION 22

A company requires its developers to tag all Amazon Elastic Block Store (Amazon EBS) volumes in an account to indicate a desired backup frequency. This requirement includes EBS volumes that do not require backups. The company uses custom tags named Backup_Frequency that have values of none, daily, or

weekly that correspond to the desired backup frequency. An audit finds that developers are occasionally not tagging the EBS volumes. A DevOps engineer needs to ensure that all EBS volumes always have the Backup_Frequency tag so that the company can perform backups at least weekly unless a different value is specified. Which solution will meet these requirements?

- A. Set up AWS Config in the account
- B. Create a custom rule that returns a compliance failure for all Amazon EC2 resources that do not have a Backup Frequency tag applied
- C. Configure a remediation action that uses a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly.
- D. Set up AWS Config in the account
- E. Use a managed rule that returns a compliance failure for EC2::Volume resources that do not have a Backup Frequency tag applied
- F. Configure a remediation action that uses a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly.
- G. Turn on AWS CloudTrail in the account
- H. Create an Amazon EventBridge rule that reacts to EBS CreateVolume event
- I. Configure a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly
- J. Specify the runbook as the target of the rule.
- K. Turn on AWS CloudTrail in the account
- L. Create an Amazon EventBridge rule that reacts to EBS CreateVolume events or EBS ModifyVolume event
- M. Configure a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly
- N. Specify the runbook as the target of the rule.

Answer: B

Explanation:

The following are the steps that the DevOps engineer should take to ensure that all EBS volumes always have the Backup_Frequency tag so that the company can perform backups at least weekly unless a different value is specified:

- > Set up AWS Config in the account.
- > Use a managed rule that returns a compliance failure for EC2::Volume resources that do not have a Backup Frequency tag applied.
- > Configure a remediation action that uses a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly.

The managed rule AWS::Config::EBSVolumesWithoutBackupTag will return a compliance failure for any EBS volume that does not have the Backup_Frequency tag applied. The remediation action will then use the Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly to the EBS volume.

NEW QUESTION 26

A company manages AWS accounts for application teams in AWS Control Tower. Individual application teams are responsible for securing their respective AWS accounts.

A DevOps engineer needs to enable Amazon GuardDuty for all AWS accounts in which the application teams have not already enabled GuardDuty. The DevOps engineer is using AWS CloudFormation StackSets from the AWS Control Tower management account.

How should the DevOps engineer configure the CloudFormation template to prevent failure during the StackSets deployment?

- A. Create a CloudFormation custom resource that invokes an AWS Lambda function
- B. Configure the Lambda function to conditionally enable GuardDuty if GuardDuty is not already enabled in the accounts.
- C. Use the Conditions section of the CloudFormation template to enable GuardDuty in accounts where GuardDuty is not already enabled.
- D. Use the CloudFormation Fn::GetAtt intrinsic function to check whether GuardDuty is already enabled. If GuardDuty is not already enabled use the Resources section of the CloudFormation template to enable GuardDuty.
- E. Use the CloudFormation Fn::GetAtt intrinsic function to check whether GuardDuty is already enabled. If GuardDuty is not already enabled use the Resources section of the CloudFormation template to enable GuardDuty.
- F. Manually discover the list of AWS account IDs where GuardDuty is not enabled. Use the CloudFormation Fn::ImportValue intrinsic function to import the list of account IDs into the CloudFormation template to skip deployment for the listed AWS accounts.

Answer: A

Explanation:

This solution will meet the requirements because it will use a CloudFormation custom resource to execute custom logic during the stack set operation. A custom resource is a resource that you define in your template and that is associated with an AWS Lambda function. The Lambda function runs whenever the custom resource is created, updated, or deleted, and can perform any actions that are supported by the AWS SDK. In this case, the Lambda function can use the GuardDuty API to check whether GuardDuty is already enabled in each target account, and if not, enable it. This way, the DevOps engineer can avoid deploying the stack set to accounts that already have GuardDuty enabled, and prevent failure during the deployment.

NEW QUESTION 28

A company has enabled all features for its organization in AWS Organizations. The organization contains 10 AWS accounts. The company has turned on AWS CloudTrail in all the accounts. The company expects the number of AWS accounts in the organization to increase to 500 during the next year. The company plans to use multiple OUs for these accounts.

The company has enabled AWS Config in each existing AWS account in the organization. A DevOps engineer must implement a solution that enables AWS Config automatically for all future AWS accounts that are created in the organization.

Which solution will meet this requirement?

- A. In the organization's management account, create an Amazon EventBridge rule that reacts to a CreateAccount API call
- B. Configure the rule to invoke an AWS Lambda function that enables trusted access to AWS Config for the organization.
- C. In the organization's management account, create an AWS CloudFormation stack set to enable AWS Config
- D. Configure the stack set to deploy automatically when an account is created through Organizations.
- E. In the organization's management account, create an SCP that allows the appropriate AWS Config API calls to enable AWS Config
- F. Apply the SCP to the root-level OU.
- G. In the organization's management account, create an Amazon EventBridge rule that reacts to a CreateAccount API call
- H. Configure the rule to invoke an AWS Systems Manager Automation runbook to enable AWS Config for the account.

Answer: B

NEW QUESTION 33

A company has containerized all of its in-house quality control applications. The company is running Jenkins on Amazon EC2 instances, which require patching and upgrading. The compliance officer has requested a DevOps engineer begin encrypting build artifacts since they contain company intellectual property. What

should the DevOps engineer do to accomplish this in the MOST maintainable manner?

- A. Automate patching and upgrading using AWS Systems Manager on EC2 instances and encrypt Amazon EBS volumes by default.
- B. Deploy Jenkins to an Amazon ECS cluster and copy build artifacts to an Amazon S3 bucket with default encryption enabled.
- C. Leverage AWS CodePipeline with a build action and encrypt the artifacts using AWS Secrets Manager.
- D. Use AWS CodeBuild with artifact encryption to replace the Jenkins instance running on EC2 instances.

Answer: D

Explanation:

The following are the steps involved in accomplishing this in the most maintainable manner:

- > Configure CodeBuild to encrypt the build artifacts using AWS Secrets Manager.
- > Deploy the containerized quality control applications to CodeBuild.

This approach is the most maintainable because it eliminates the need to manage Jenkins on EC2 instances. CodeBuild is a managed service, so the DevOps engineer does not need to worry about patching or upgrading the service.

<https://docs.aws.amazon.com/codebuild/latest/userguide/security-encryption.html> Build artifact encryption - CodeBuild requires access to an AWS KMS CMK in order to encrypt its build output artifacts. By default, CodeBuild uses an AWS Key Management Service CMK for Amazon S3 in your AWS account. If you do not want to use this CMK, you must create and configure a customer-managed CMK. For more information Creating keys.

NEW QUESTION 38

A production account has a requirement that any Amazon EC2 instance that has been logged in to manually must be terminated within 24 hours. All applications in the production account are using Auto Scaling groups with the Amazon CloudWatch Logs agent configured.

How can this process be automated?

- A. Create a CloudWatch Logs subscription to an AWS Step Functions applicatio
- B. Configure an AWS Lambda function to add a tag to the EC2 instance that produced the login event and mark the instance to be decommissioned
- C. Create an Amazon EventBridge rule to invoke a second Lambda function once a day that will terminate all instances with this tag.
- D. Create an Amazon CloudWatch alarm that will be invoked by the login event
- E. Send the notification to an Amazon Simple Notification Service (Amazon SNS) topic that the operations team is subscribed to, and have them terminate the EC2 instance within 24 hours.
- F. Create an Amazon CloudWatch alarm that will be invoked by the login event
- G. Configure the alarm to send to an Amazon Simple Queue Service (Amazon SQS) queue
- H. Use a group of worker instances to process messages from the queue, which then schedules an Amazon EventBridge rule to be invoked.
- I. Create a CloudWatch Logs subscription to an AWS Lambda function
- J. Configure the function to add a tag to the EC2 instance that produced the login event and mark the instance to be decommissioned. Create an Amazon EventBridge rule to invoke a daily Lambda function that terminates all instances with this tag.

Answer: D

Explanation:

"You can use subscriptions to get access to a real-time feed of log events from CloudWatch Logs and have it delivered to other services such as an Amazon Kinesis stream, an Amazon Kinesis Data Firehose stream, or AWS Lambda for custom processing, analysis, or loading to other systems. When log events are sent to the receiving service, they are Base64 encoded and compressed with the gzip format." See <https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/Subscriptions.html>

NEW QUESTION 39

An IT team has built an AWS CloudFormation template so others in the company can quickly and reliably deploy and terminate an application. The template creates an Amazon EC2 instance with a user data script to install the application and an Amazon S3 bucket that the application uses to serve static webpages while it is running.

All resources should be removed when the CloudFormation stack is deleted. However, the team observes that CloudFormation reports an error during stack deletion, and the S3 bucket created by the stack is not deleted.

How can the team resolve the error in the MOST efficient manner to ensure that all resources are deleted without errors?

- A. Add a DeletionPolicy attribute to the S3 bucket resource, with the value Delete forcing the bucket to be removed when the stack is deleted.
- B. Add a custom resource with an AWS Lambda function with the DependsOn attribute specifying the S3 bucket, and an IAM role
- C. Write the Lambda function to delete all objects from the bucket when RequestType is Delete.
- D. Identify the resource that was not deleted
- E. Manually empty the S3 bucket and then delete it.
- F. Replace the EC2 and S3 bucket resources with a single AWS OpsWorks Stacks resource
- G. Define a custom recipe for the stack to create and delete the EC2 instance and the S3 bucket.

Answer: B

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/cloudformation-s3-custom-resources/>

NEW QUESTION 40

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 (or checkpoint data)
- B. To complete the job, use an EC2 Auto Scaling group and an On-Demand pricing model to provision EC2 instances temporarily.
- C. Use GlusterFS on EC2 instances for checkpoint data
- D. To run the batch job configure EC2 instances manually When the job completes shut down the instances manually.
- E. 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.
- F. Use Amazon EFS for checkpoint data Use EC2 Fleet to launch EC2 Spot Instances Create a customAMI for the cluster and use the latest AMI when creating instances.

Answer: D

NEW QUESTION 45

The security team depends on AWS CloudTrail to detect sensitive security issues in the company's AWS account. The DevOps engineer needs a solution to auto-remediate CloudTrail being turned off in an AWS account.

What solution ensures the LEAST amount of downtime for the CloudTrail log deliveries?

- A. Create an Amazon EventBridge rule for the CloudTrail StopLogging even
- B. Create an AWS Lambda function that uses the AWS SDK to call StartLogging on the ARN of the resource in which StopLogging was called
- C. Add the Lambda function ARN as a target to the EventBridge rule.
- D. Deploy the AWS-managed CloudTrail-enabled AWS Config rule set with a periodic interval to 1 hour. Create an Amazon EventBridge rule for AWS Config rules compliance change
- E. Create an AWS Lambda function that uses the AWS SDK to call StartLogging on the ARN of the resource in which StopLogging was called
- F. Add the Lambda function ARN as a target to the EventBridge rule.
- G. Create an Amazon EventBridge rule for a scheduled event every 5 minutes
- H. Create an AWS Lambda function that uses the AWS SDK to call StartLogging on a CloudTrail trail in the AWS account
- I. Add the Lambda function ARN as a target to the EventBridge rule.
- J. Launch a t2 nano instance with a script running every 5 minutes that uses the AWS SDK to query CloudTrail in the current account
- K. If the CloudTrail trail is disabled have the script re-enable the trail.

Answer: A

Explanation:

<https://aws.amazon.com/blogs/mt/monitor-changes-and-auto-enable-logging-in-aws-cloudtrail/>

NEW QUESTION 49

A company has a single AWS account that runs hundreds of Amazon EC2 instances in a single AWS Region. New EC2 instances are launched and terminated each hour in the account. The account also includes existing EC2 instances that have been running for longer than a week.

The company's security policy requires all running EC2 instances to use an EC2 instance profile. If an EC2 instance does not have an instance profile attached, the EC2 instance must use a default instance profile that has no IAM permissions assigned.

A DevOps engineer reviews the account and discovers EC2 instances that are running without an instance profile. During the review, the DevOps engineer also observes that new EC2 instances are being launched without an instance profile.

Which solution will ensure that an instance profile is attached to all existing and future EC2 instances in the Region?

- A. Configure an Amazon EventBridge rule that reacts to EC2 RunInstances API call
- B. Configure the rule to invoke an AWS Lambda function to attach the default instance profile to the EC2 instances.
- C. Configure the ec2-instance-profile-attached AWS Config managed rule with a trigger type of configuration change
- D. Configure an automatic remediation action that invokes an AWS Systems Manager Automation runbook to attach the default instance profile to the EC2 instances.
- E. Configure an Amazon EventBridge rule that reacts to EC2 StartInstances API call
- F. Configure the rule to invoke an AWS Systems Manager Automation runbook to attach the default instance profile to the EC2 instances.
- G. Configure the iam-role-managed-policy-check AWS Config managed rule with a trigger type of configuration change
- H. Configure an automatic remediation action that invokes an AWS Lambda function to attach the default instance profile to the EC2 instances.

Answer: B

Explanation:

<https://docs.aws.amazon.com/config/latest/developerguide/ec2-instance-profile-attached.html>

NEW QUESTION 53

A company has many AWS accounts. During AWS account creation the company uses automation to create an Amazon CloudWatch Logs log group in every AWS Region that the company operates in. The automation configures new resources in the accounts to publish logs to the provisioned log groups in their Region. The company has created a logging account to centralize the logging from all the other accounts. A DevOps engineer needs to aggregate the log groups from all the accounts to an existing Amazon S3 bucket in the logging account.

Which solution will meet these requirements in the MOST operationally efficient manner?

- A. In the logging account create a CloudWatch Logs destination with a destination policy
- B. For each new account subscribe the CloudWatch Logs log groups to the destination
- C. Destination Configure a single Amazon Kinesis data stream and a single Amazon Kinesis Data Firehose delivery stream to deliver the logs from the CloudWatch Logs destination to the S3 bucket.
- D. In the logging account create a CloudWatch Logs destination with a destination policy for each Region. For each new account subscribe the CloudWatch Logs log groups to the destination
- E. Configure a single Amazon Kinesis data stream and a single Amazon Kinesis Data Firehose delivery stream to deliver the logs from all the CloudWatch Logs destinations to the S3 bucket.
- F. In the logging account create a CloudWatch Logs destination with a destination policy for each Region. For each new account subscribe the CloudWatch Logs log groups to the destination. Configure an Amazon Kinesis data stream and an Amazon Kinesis Data Firehose delivery stream for each Region to deliver the logs from the CloudWatch Logs destinations to the S3 bucket.
- G. In the logging account create a CloudWatch Logs destination with a destination policy
- H. For each new account subscribe the CloudWatch Logs log groups to the destination
- I. Configure a single Amazon Kinesis data stream to deliver the logs from the CloudWatch Logs destination to the S3 bucket.

Answer: C

Explanation:

This solution will meet the requirements in the most operationally efficient manner because it will use CloudWatch Logs destination to aggregate the log groups from all the accounts to a single S3 bucket in the logging account. However, unlike option A, this solution will create a CloudWatch Logs destination for each region, instead of a single destination for all regions. This will improve the performance and reliability of the log delivery, as it will avoid cross-region data transfer and latency issues. Moreover, this solution will use an Amazon Kinesis data stream and an Amazon Kinesis Data Firehose delivery stream for each region, instead of a single stream for all regions. This will also improve the scalability and throughput of the log delivery, as it will avoid bottlenecks and throttling issues that may occur with a single stream.

NEW QUESTION 57

A company has an application that runs on a fleet of Amazon EC2 instances. The application requires frequent restarts. The application logs contain error messages when a restart is required. The application logs are published to a log group in Amazon CloudWatch Logs. An Amazon CloudWatch alarm notifies an application engineer through an Amazon Simple Notification Service (Amazon SNS) topic when the logs contain a large number of restart-related error messages. The application engineer manually restarts the application on the instances after the application engineer receives a notification from the SNS topic. A DevOps engineer needs to implement a solution to automate the application restart on the instances without restarting the instances. Which solution will meet these requirements in the MOST operationally efficient manner?

- A. Configure an AWS Systems Manager Automation runbook that runs a script to restart the application on the instance
- B. Configure the SNS topic to invoke the runbook.
- C. Create an AWS Lambda function that restarts the application on the instance
- D. Configure the Lambda function as an event destination of the SNS topic.
- E. Configure an AWS Systems Manager Automation runbook that runs a script to restart the application on the instance
- F. Create an AWS Lambda function to invoke the runbook
- G. Configure the Lambda function as an event destination of the SNS topic.
- H. Configure an AWS Systems Manager Automation runbook that runs a script to restart the application on the instance
- I. Configure an Amazon EventBridge rule that reacts when the CloudWatch alarm enters ALARM state
- J. Specify the runbook as a target of the rule.

Answer: D

Explanation:

This solution meets the requirements in the most operationally efficient manner by automating the application restart process on the instances without restarting them. When the CloudWatch alarm enters the ALARM state, the EventBridge rule is triggered, which in turn invokes the Systems Manager Automation runbook that contains the script to restart the application on the instances.

NEW QUESTION 62

A company is using an Amazon Aurora cluster as the data store for its application. The Aurora cluster is configured with a single DB instance. The application performs read and write operations on the database by using the cluster's instance endpoint. The company has scheduled an update to be applied to the cluster during an upcoming maintenance window. The cluster must remain available with the least possible interruption during the maintenance window. What should a DevOps engineer do to meet these requirements?

- A. Add a reader instance to the Aurora cluster
- B. Update the application to use the Aurora cluster endpoint for write operation
- C. Update the Aurora cluster's reader endpoint for reads.
- D. Add a reader instance to the Aurora cluster
- E. Create a custom ANY endpoint for the cluster
- F. Update the application to use the Aurora cluster's custom ANY endpoint for read and write operations.
- G. Turn on the Multi-AZ option on the Aurora cluster
- H. Update the application to use the Aurora cluster endpoint for write operation
- I. Update the Aurora cluster's reader endpoint for reads.
- J. Turn on the Multi-AZ option on the Aurora cluster
- K. Create a custom ANY endpoint for the cluster. Update the application to use the Aurora cluster's custom ANY endpoint for read and write operations.

Answer: C

Explanation:

To meet the requirements, the DevOps engineer should do the following:

- > Turn on the Multi-AZ option on the Aurora cluster.
- > Update the application to use the Aurora cluster endpoint for write operations.
- > Update the Aurora cluster's reader endpoint for reads.

Turning on the Multi-AZ option will create a replica of the database in a different Availability Zone. This will ensure that the database remains available even if one of the Availability Zones is unavailable.

Updating the application to use the Aurora cluster endpoint for write operations will ensure that all writes are sent to both the primary and replica databases. This will ensure that the data is always consistent.

Updating the Aurora cluster's reader endpoint for reads will allow the application to read data from the replica database. This will improve the performance of the application during the maintenance window.

NEW QUESTION 64

A DevOps team manages an API running on-premises that serves as a backend for an Amazon API Gateway endpoint. Customers have been complaining about high response latencies, which the development team has verified using the API Gateway latency metrics in Amazon CloudWatch. To identify the cause, the team needs to collect relevant data without introducing additional latency. Which actions should be taken to accomplish this? (Choose two.)

- A. Install the CloudWatch agent server side and configure the agent to upload relevant logs to CloudWatch.
- B. Enable AWS X-Ray tracing in API Gateway, modify the application to capture request segments, and upload those segments to X-Ray during each request.
- C. Enable AWS X-Ray tracing in API Gateway, modify the application to capture request segments, and use the X-Ray daemon to upload segments to X-Ray.
- D. Modify the on-premises application to send log information back to API Gateway with each request.
- E. Modify the on-premises application to calculate and upload statistical data relevant to the API service requests to CloudWatch metrics.

Answer: AC

Explanation:

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/install-CloudWatch-Agent-on-premise.htm>

<https://docs.aws.amazon.com/xray/latest/devguide/xray-api-sendingdata.html>

NEW QUESTION 66

A company uses AWS CodePipeline pipelines to automate releases of its application. A typical pipeline consists of three stages: build, test, and deployment. The company has been using a separate AWS CodeBuild project to run scripts for each stage. However, the company now wants to use AWS CodeDeploy to handle the deployment stage of the pipelines.

The company has packaged the application as an RPM package and must deploy the application to a fleet of Amazon EC2 instances. The EC2 instances are in an EC2 Auto Scaling group and are launched from a common AMI.

Which combination of steps should a DevOps engineer perform to meet these requirements? (Choose two.)

- A. Create a new version of the common AMI with the CodeDeploy agent installed
- B. Update the IAM role of the EC2 instances to allow access to CodeDeploy.
- C. Create a new version of the common AMI with the CodeDeploy agent installed
- D. Create an AppSpec file that contains application deployment scripts and grants access to CodeDeploy.
- E. Create an application in CodeDeploy
- F. Configure an in-place deployment type
- G. Specify the Auto Scaling group as the deployment target
- H. Add a step to the CodePipeline pipeline to use EC2 Image Builder to create a new AMI
- I. Configure CodeDeploy to deploy the newly created AMI.
- J. Create an application in CodeDeploy
- K. Configure an in-place deployment type
- L. Specify the Auto Scaling group as the deployment target
- M. Update the CodePipeline pipeline to use the CodeDeploy action to deploy the application.
- N. Create an application in CodeDeploy
- O. Configure an in-place deployment type
- P. Specify the EC2 instances that are launched from the common AMI as the deployment target
- Q. Update the CodePipeline pipeline to use the CodeDeploy action to deploy the application.

Answer: AD

Explanation:

<https://docs.aws.amazon.com/codedeploy/latest/userguide/integrations-aws-auto-scaling.html>

NEW QUESTION 69

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual DOP-C02 Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the DOP-C02 Product From:

<https://www.2passeasy.com/dumps/DOP-C02/>

Money Back Guarantee

DOP-C02 Practice Exam Features:

- * DOP-C02 Questions and Answers Updated Frequently
- * DOP-C02 Practice Questions Verified by Expert Senior Certified Staff
- * DOP-C02 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * DOP-C02 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year