

# Amazon-Web-Services

## Exam Questions DOP-C02

AWS Certified DevOps Engineer - Professional



### NEW QUESTION 1

A company's DevOps engineer uses AWS Systems Manager to perform maintenance tasks during maintenance windows. The company has a few Amazon EC2 instances that require a restart after notifications from AWS Health. The DevOps engineer needs to implement an automated solution to remediate these notifications. The DevOps engineer creates an Amazon EventBridge rule.

How should the DevOps engineer configure the EventBridge rule to meet these requirements?

- A. Configure an event source of AWS Health, a service of EC2, and an event type that indicates instance maintenance
- B. Target a Systems Manager document to restart the EC2 instance.
- C. Configure an event source of Systems Manager and an event type that indicates a maintenance window
- D. Target a Systems Manager document to restart the EC2 instance.
- E. Configure an event source of AWS Health, a service of EC2, and an event type that indicates instance maintenance
- F. Target a newly created AWS Lambda function that registers an automation task to restart the EC2 instance during a maintenance window.
- G. Configure an event source of EC2 and an event type that indicates instance maintenance
- H. Target a newly created AWS Lambda function that registers an automation task to restart the EC2 instance during a maintenance window.

**Answer: C**

#### Explanation:

AWS Health provides real-time events and information related to your AWS infrastructure. It can be integrated with Amazon EventBridge to act upon the health events automatically. If the maintenance notification from AWS Health indicates that an EC2 instance requires a restart, you can set up an EventBridge rule to respond to such events. In this case, the target of this rule would be a Lambda function that would trigger a Systems Manager automation to restart the EC2 instance during a maintenance window. Remember, AWS Health is the source of the events (not EC2 or Systems Manager), and AWS Lambda can be used to execute complex remediation tasks, such as scheduling maintenance tasks via Systems Manager.

The following are the steps involved in configuring the EventBridge rule to meet these requirements:

? Configure an event source of AWS Health, a service of EC2, and an event type that indicates instance maintenance.

? Target a newly created AWS Lambda function that registers an automation task to restart the EC2 instance during a maintenance window.

The AWS Lambda function will be triggered by the event from AWS Health. The function will then register an automation task to restart the EC2 instance during the next maintenance window.

### NEW QUESTION 2

A company wants to use AWS development tools to replace its current bash deployment scripts. The company currently deploys a LAMP application to a group of Amazon EC2 instances behind an Application Load Balancer (ALB). During the deployments, the company unit tests the committed application, stops and starts services, unregisters and re-registers instances with the load balancer, and updates file permissions. The company wants to maintain the same deployment functionality through the shift to using AWS services.

Which solution will meet these requirements?

- A. Use AWS CodeBuild to test the application
- B. Use bash scripts invoked by AWS CodeDeploy's appspec.yml file to restart services, and deregister and register instances with the ALB
- C. Use the appspec.yml file to update file permissions without a custom script.
- D. Use AWS CodePipeline to move the application from the AWS CodeCommit repository to AWS CodeDeploy
- E. Use CodeDeploy's deployment group to test the application, unregister and re-register instances with the ALB and restart service
- F. and restart service
- G. Use the appspec.yml file to update file permissions without a custom script.
- H. Use AWS CodePipeline to move the application source code from the AWS CodeCommit repository to AWS CodeDeploy
- I. Use CodeDeploy to test the application
- J. Use CodeDeploy's appspec.yml file to restart services and update permissions without a custom script
- K. Use AWS CodeBuild to unregister and re-register instances with the ALB.
- L. Use AWS CodePipeline to trigger AWS CodeBuild to test the application
- M. Use bash scripts invoked by AWS CodeDeploy's appspec.yml file to restart service
- N. Unregister and re-register the instances in the AWS CodeDeploy deployment group with the ALB
- O. Update the appspec.yml file to update file permissions without a custom script.

**Answer: D**

#### Explanation:

<https://aws.amazon.com/blogs/devops/how-to-test-and-debug-aws-codedeploy-locally-before-you-ship-your-code/#:~:text=You%20can%20test%20application%20code,local%20server%20or%20EC2%20instance.>

### NEW QUESTION 3

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
- D. Configure a filter policy on the SNS topic to send only notifications that contain the text of NON\_COMPLIANT in the notification to subscribers.
- E. 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
- F. 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 4**

A DevOps engineer needs to apply a core set of security controls to an existing set of AWS accounts. The accounts are in an organization in AWS Organizations. Individual teams will administer individual accounts by using the AdministratorAccess AWS managed policy. For all accounts, AWS CloudTrail and AWS Config must be turned on in all available AWS Regions. Individual account administrators must not be able to edit or delete any of the baseline resources. However, individual account administrators must be able to edit or delete their own CloudTrail trails and AWS Config rules. Which solution will meet these requirements in the MOST operationally efficient way?

- A. Create an AWS CloudFormation template that defines the standard account resource
- B. Deploy the template to all accounts from the organization's management account by using CloudFormation StackSet
- C. Set the stack policy to deny Update:Delete actions.
- D. Enable AWS Control Towe
- E. Enroll the existing accounts in AWS Control Towe
- F. Grant the individual account administrators access to CloudTrail and AWS Config.
- G. Designate an AWS Config management accoun
- H. Create AWS Config recorders in all accounts by using AWS CloudFormation StackSet
- I. Deploy AWS Config rules to the organization by using the AWS Config management accoun
- J. Create a CloudTrail organization trail in the organization's management accoun
- K. Deny modification or deletion of the AWS Config recorders by using an SCP.
- L. Create an AWS CloudFormation template that defines the standard account resource
- M. Deploy the template to all accounts from the organization's management account by using Cloud Formation StackSets Create an SCP that prevents updates or deletions to CloudTrail resources or AWS Config resources unless the principal is an administrator of the organization's management account.

**Answer: D**

**NEW QUESTION 5**

A DevOps engineer is designing an application that integrates with a legacy REST API. The application has an AWS Lambda function that reads records from an Amazon Kinesis data stream. The Lambda function sends the records to the legacy REST API.

Approximately 10% of the records that the Lambda function sends from the Kinesis data stream have data errors and must be processed manually. The Lambda function event source configuration has an Amazon Simple Queue Service (Amazon SQS) dead-letter queue as an on-failure destination. The DevOps engineer has configured the Lambda function to process records in batches and has implemented retries in case of failure.

During testing the DevOps engineer notices that the dead-letter queue contains many records that have no data errors and that already have been processed by the legacy REST API. The DevOps engineer needs to configure the Lambda function's event source options to reduce the number of errorless records that are sent to the dead-letter queue.

Which solution will meet these requirements?

- A. Increase the retry attempts
- B. Configure the setting to split the batch when an error occurs
- C. Increase the concurrent batches per shard
- D. Decrease the maximum age of record

**Answer: B**

**Explanation:**

This solution will meet the requirements because it will reduce the number of errorless records that are sent to the dead-letter queue. When you configure the setting to split the batch when an error occurs, Lambda will retry only the records that caused the error, instead of retrying the entire batch. This way, the records that have no data errors and have already been processed by the legacy REST API will not be retried and sent to the dead-letter queue unnecessarily.

<https://docs.aws.amazon.com/lambda/latest/dg/with-kinesis.html>

**NEW QUESTION 6**

A company runs an application on one Amazon EC2 instance. Application metadata is stored in Amazon S3 and must be retrieved if the instance is restarted. The instance must restart or relaunch automatically if the instance becomes unresponsive.

Which solution will meet these requirements?

- A. Create an Amazon CloudWatch alarm for the StatusCheckFailed metri
- B. Use the recover action to stop and start the instanc
- C. Use an S3 event notification to push the metadata to the instance when the instance is back up and running.
- D. Configure AWS OpsWorks, and use the auto healing feature to stop and start the instanc
- E. Use a lifecycle event in OpsWorks to pull the metadata from Amazon S3 and update it on the instance.
- F. Use EC2 Auto Recovery to automatically stop and start the instance in case of a failur
- G. Use an S3 event notification to push the metadata to the instance when the instance is back up and running.
- H. Use AWS CloudFormation to create an EC2 instance that includes the UserData property for the EC2 resourc
- I. Add a command in UserData to retrieve the application metadata from Amazon S3.

**Answer: B**

**Explanation:**

<https://aws.amazon.com/blogs/mt/how-to-set-up-aws-opsworks-stacks-auto-healing-notifications-in-amazon-cloudwatch-events/>

**NEW QUESTION 7**

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 environment
- I. Then place a script into the application revision that checks this environment variable to identify which deployment group the instance is part of
- 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 8**

A development team is using AWS CodeCommit to version control application code and AWS CodePipeline to orchestrate software deployments. The team has decided to use a remote main branch as the trigger for the pipeline to integrate code changes. A developer has pushed code changes to the CodeCommit repository, but noticed that the pipeline had no reaction, even after 10 minutes.

Which of the following actions should be taken to troubleshoot this issue?

- A. Check that an Amazon EventBridge rule has been created for the main branch to trigger the pipeline.
- B. Check that the CodePipeline service role has permission to access the CodeCommit repository.
- C. Check that the developer's IAM role has permission to push to the CodeCommit repository.
- D. Check to see if the pipeline failed to start because of CodeCommit errors in Amazon CloudWatch Logs.

**Answer: A**

**Explanation:**

When you create a pipeline from CodePipeline during the step-by-step it creates a CloudWatch Event rule for a given branch and repo like this:

```
{
  "source": [ "aws.codecommit"
],
  "detail-type": [
    "CodeCommit Repository State Change"
],
  "resources": [
    "arn:aws:codecommit:us-east-1:xxxxx:repo-name"
],
  "detail": {
    "event": [ "referenceCreated", "referenceUpdated"
],
    "referenceType": [ "branch"
],
    "referenceName": [ "master"
]
}
}
```

<https://docs.aws.amazon.com/codepipeline/latest/userguide/pipelines-trigger-source-repo-changes-console.html>

**NEW QUESTION 9**

A company is implementing a well-architected design for its globally accessible API stack. The design needs to ensure both high reliability and fast response times for users located in North America and Europe.

The API stack contains the following three tiers: Amazon API Gateway  
 AWS Lambda Amazon DynamoDB

Which solution will meet the requirements?

- A. Configure Amazon Route 53 to point to API Gateway APIs in North America and Europe using health check

- B. Configure the APIs to forward requests to a Lambda function in that Region
- C. Configure the Lambda functions to retrieve and update the data in a DynamoDB table in the same Region as the Lambda function.
- D. Configure Amazon Route 53 to point to API Gateway APIs in North America and Europe using latency-based routing and health check
- E. Configure the APIs to forward requests to a Lambda function in that Region
- F. Configure the Lambda functions to retrieve and update the data in a DynamoDB global table.
- G. Configure Amazon Route 53 to point to API Gateway in North America, create a disaster recovery API in Europe, and configure both APIs to forward requests to the Lambda functions in that Region
- H. Retrieve the data from a DynamoDB global table
- I. Deploy a Lambda function to check the North America API health every 5 minutes
- J. In the event of a failure, update Route 53 to point to the disaster recovery API.
- K. Configure Amazon Route 53 to point to API Gateway API in North America using latency-based routing
- L. Configure the API to forward requests to the Lambda function in the Region nearest to the user
- M. Configure the Lambda function to retrieve and update the data in a DynamoDB table.

**Answer: B**

#### NEW QUESTION 10

A company has a data ingestion application that runs across multiple AWS accounts. The accounts are in an organization in AWS Organizations. The company needs to monitor the application and consolidate access to the application. Currently the company is running the application on Amazon EC2 instances from several Auto Scaling groups. The EC2 instances have no access to the internet because the data is sensitive. Engineers have deployed the necessary VPC endpoints. The EC2 instances run a custom AMI that is built specifically for the application. To maintain and troubleshoot the application, system administrators need the ability to log in to the EC2 instances. This access must be automated and controlled centrally. The company's security team must receive a notification whenever the instances are accessed. Which solution will meet these requirements?

- A. Create an Amazon EventBridge rule to send notifications to the security team whenever a user logs in to an EC2 instance. Use EC2 Instance Connect to log in to the instance.
- B. Deploy Auto Scaling groups by using AWS CloudFormation. Use the cfn-init helper script to deploy appropriate VPC routes for external access. Rebuild the custom AMI so that the custom AMI includes AWS Systems Manager Agent.
- C. Deploy a NAT gateway and a bastion host that has internet access. Create a security group that allows incoming traffic on all the EC2 instances from the bastion host. Install AWS Systems Manager Agent on all the EC2 instances. Use Auto Scaling group lifecycle hooks for monitoring and auditing access. Use Systems Manager Session Manager to log into the instances. Send logs to a log group in Amazon CloudWatch Log.
- D. Export data to Amazon S3 for auditing. Send notifications to the security team by using S3 event notifications.
- E. Use EC2 Image Builder to rebuild the custom AMI. Include the most recent version of AWS Systems Manager Agent in the image. Configure the Auto Scaling group to attach the AmazonSSMManagedInstanceCore role to all the EC2 instances. Use Systems Manager Session Manager to log in to the instances. Enable logging of session details to Amazon S3. Create an S3 event notification for new file uploads to send a message to the security team through an Amazon Simple Notification Service (Amazon SNS) topic.
- F. Use AWS Systems Manager Automation to build Systems Manager Agent into the custom AMI. Configure AWS Config to attach an SCP to the root organization account to allow the EC2 instances to connect to Systems Manager. Use Systems Manager Session Manager to log in to the instances. Enable logging of session details to Amazon S3. Create an S3 event notification for new file uploads to send a message to the security team through an Amazon Simple Notification Service (Amazon SNS) topic.

**Answer: C**

#### Explanation:

Even if AmazonSSMManagedInstanceCore is a managed policy and not an IAM role, I will go with C because this policy is to be attached to an IAM role for EC2 to access Systems Manager.

#### NEW QUESTION 10

A company uses AWS Organizations and AWS Control Tower to manage all the company's AWS accounts. The company uses the Enterprise Support plan. A DevOps engineer is using Account Factory for Terraform (AFT) to provision new accounts. When new accounts are provisioned, the DevOps engineer notices that the support plan for the new accounts is set to the Basic Support plan. The DevOps engineer needs to implement a solution to provision the new accounts with the Enterprise Support plan. Which solution will meet these requirements?

- A. Use an AWS Config conformance pack to deploy the account-part-of-organizations AWS Config rule and to automatically remediate any noncompliant accounts.
- B. Create an AWS Lambda function to create a ticket for AWS Support to add the account to the Enterprise Support plan.
- C. Grant the Lambda function the support:ResolveCase permission.
- D. Add an additional value to the control\_tower\_parameters input to set the AWSEnterpriseSupport parameter as the organization's management account number.
- E. Set the aft\_feature\_enterprise\_support feature flag to True in the AFT deployment input configuration.
- F. Redeploy AFT and apply the changes.

**Answer: D**

#### Explanation:

AWS Organizations is a service that helps to manage multiple AWS accounts. AWS Control Tower is a service that makes it easy to set up and govern secure, compliant multi-account AWS environments. Account Factory for Terraform (AFT) is an AWS Control Tower feature that provisions new accounts using Terraform templates. To provision new accounts with the Enterprise Support plan, the DevOps engineer can set the aft\_feature\_enterprise\_support feature flag to True in the AFT deployment input configuration. This flag enables the Enterprise Support plan for newly provisioned accounts.

<https://docs.aws.amazon.com/controltower/latest/userguide/aft-feature-options.html>

#### NEW QUESTION 11

A company has an application and a CI/CD pipeline. The CI/CD pipeline consists of an AWS CodePipeline pipeline and an AWS CodeBuild project. The CodeBuild project runs tests against the application as part of the build process and outputs a test report. The company must keep the test reports for 90 days. Which solution will meet these requirements?

- A. Add a new stage in the CodePipeline pipeline after the stage that contains the CodeBuild project.
- B. Create an Amazon S3 bucket to store the report.

- C. Configure an S3 deploy action type in the new CodePipeline stage with the appropriate path and format for the reports.
- D. Add a report group in the CodeBuild project buildspec file with the appropriate path and format for the report
- E. Create an Amazon S3 bucket to store the report
- F. Configure an Amazon EventBridge rule that invokes an AWS Lambda function to copy the reports to the S3 bucket when a build is complete
- G. Create an S3 Lifecycle rule to expire the objects after 90 days.
- H. Add a new stage in the CodePipeline pipeline
- I. Configure a test action type with the appropriate path and format for the report
- J. Configure the report expiration time to be 90 days in the CodeBuild project buildspec file.
- K. Add a report group in the CodeBuild project buildspec file with the appropriate path and format for the report
- L. Create an Amazon S3 bucket to store the report
- M. Configure the report group as an artifact in the CodeBuild project buildspec file
- N. Configure the S3 bucket as the artifact destination
- O. Set the object expiration to 90 days.

**Answer: B**

**Explanation:**

The correct solution is to add a report group in the AWS CodeBuild project buildspec file with the appropriate path and format for the reports. Then, create an Amazon S3 bucket to store the reports. You should configure an Amazon EventBridge rule that invokes an AWS Lambda function to copy the reports to the S3 bucket when a build is completed. Finally, create an S3 Lifecycle rule to expire the objects after 90 days. This approach allows for the automated transfer of reports to long-term storage and ensures

they are retained for the required duration without manual intervention<sup>1</sup>. References:

- ? AWS CodeBuild User Guide on test reporting<sup>1</sup>.
- ? AWS CodeBuild User Guide on working with report groups<sup>2</sup>.
- ? AWS Documentation on using AWS CodePipeline with AWS CodeBuild<sup>3</sup>.

**NEW QUESTION 16**

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 17**

A company has an organization in AWS Organizations. The organization includes workload accounts that contain enterprise applications. The company centrally manages users from an operations account. No users can be created in the workload accounts. The company recently added an operations team and must provide the operations team members with administrator access to each workload account.

Which combination of actions will provide this access? (Choose three.)

- A. Create a SysAdmin role in the operations account
- B. Attach the AdministratorAccess policy to the role
- C. Modify the trust relationship to allow the sts:AssumeRole action from the workload accounts.
- D. Create a SysAdmin role in each workload account
- E. Attach the AdministratorAccess policy to the role
- F. Modify the trust relationship to allow the sts:AssumeRole action from the operations account.
- G. Create an Amazon Cognito identity pool in the operations account
- H. Attach the SysAdmin role as an authenticated role.
- I. In the operations account, create an IAM user for each operations team member.
- J. In the operations account, create an IAM user group that is named SysAdmin
- K. Add an IAM policy that allows the sts:AssumeRole action for the SysAdmin role in each workload account
- L. Add all operations team members to the group.
- M. Create an Amazon Cognito user pool in the operations account
- N. Create an Amazon Cognito user for each operations team member.

**Answer: BDE**

**Explanation:**

[https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial\\_cross-account\\_with-roles.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account_with-roles.html)

#### NEW QUESTION 18

A DevOps engineer is deploying a new version of a company's application in an AWS CodeDeploy deployment group associated with its Amazon EC2 instances. After some time, the deployment fails. The engineer realizes that all the events associated with the specific deployment ID are in a Skipped status and code was not deployed in the instances associated with the deployment group.

What are valid reasons for this failure? (Select TWO.).

- A. The networking configuration does not allow the EC2 instances to reach the internet via a NAT gateway or internet gateway and the CodeDeploy endpoint cannot be reached.
- B. The IAM user who triggered the application deployment does not have permission to interact with the CodeDeploy endpoint.
- C. The target EC2 instances were not properly registered with the CodeDeploy endpoint.
- D. An instance profile with proper permissions was not attached to the target EC2 instances.
- E. The appspec
- F. yml file was not included in the application revision.

**Answer:** AD

#### Explanation:

[https://docs.aws.amazon.com/codedeploy/latest/userguide/troubleshooting\\_deployments.html#troubleshooting-skipped-lifecycle-events](https://docs.aws.amazon.com/codedeploy/latest/userguide/troubleshooting_deployments.html#troubleshooting-skipped-lifecycle-events)

#### NEW QUESTION 23

A company runs a workload on Amazon EC2 instances. The company needs a control that requires the use of Instance Metadata Service Version 2 (IMDSv2) on all EC2 instances in the AWS account. If an EC2 instance does not prevent the use of Instance Metadata Service Version 1 (IMDSv1), the EC2 instance must be terminated.

Which solution will meet these requirements?

- A. Set up AWS Config in the account
- B. Use a managed rule to check EC2 instance
- C. Configure the rule to remediate the findings by using AWS Systems Manager Automation to terminate the instance.
- D. Create a permissions boundary that prevents the ec2:RunInstance action if the ec2:MetadataHttpTokens condition key is not set to a value of required.
- E. Attach the permissions boundary to the IAM role that was used to launch the instance.
- F. Set up Amazon Inspector in the account
- G. Configure Amazon Inspector to activate deep inspection for EC2 instance
- H. Create an Amazon EventBridge rule for an Inspector2 finding
- I. Set an AWS Lambda function as the target to terminate the instance.
- J. Create an Amazon EventBridge rule for the EC2 instance launch successful event
- K. Send the event to an AWS Lambda function to inspect the EC2 metadata and to terminate the instance.

**Answer:** B

#### Explanation:

To implement a control that requires the use of IMDSv2 on all EC2 instances in the account, the DevOps engineer can use a permissions boundary. A permissions boundary is a policy that defines the maximum permissions that an IAM entity can have. The DevOps engineer can create a permissions boundary that prevents the ec2:RunInstance action if the ec2:MetadataHttpTokens condition key is not set to a value of required. This condition key enforces the use of IMDSv2 on EC2 instances. The DevOps engineer can attach the permissions boundary to the IAM role that was used to launch the instance. This way, any attempt to launch an EC2 instance without using IMDSv2 will be denied by the permissions boundary.

#### NEW QUESTION 28

A company has multiple member accounts that are part of an organization in AWS Organizations. The security team needs to review every Amazon EC2 security group and their inbound and outbound rules. The security team wants to programmatically retrieve this information from the member accounts using an AWS Lambda function in the management account of the organization.

Which combination of access changes will meet these requirements? (Choose three.)

- A. Create a trust relationship that allows users in the member accounts to assume the management account IAM role.
- B. Create a trust relationship that allows users in the management account to assume the IAM roles of the member accounts.
- C. Create an IAM role in each member account that has access to the AmazonEC2ReadOnlyAccess managed policy.
- D. Create an IAM role in each member account to allow the sts:AssumeRole action against the management account IAM role's ARN.
- E. Create an IAM role in the management account that allows the sts:AssumeRole action against the member account IAM role's ARN.
- F. Create an IAM role in the management account that has access to the AmazonEC2ReadOnlyAccess managed policy.

**Answer:** BCE

#### Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/lambda-function-assume-iam-role/> <https://kreuzwerker.de/post/aws-multi-account-setups-reloaded>

#### NEW QUESTION 31

A DevOps engineer has developed an AWS Lambda function. The Lambda function starts an AWS CloudFormation drift detection operation on all supported resources for a specific CloudFormation stack. The Lambda function then exits its invocation. The DevOps engineer has created an Amazon EventBridge scheduled rule that invokes the Lambda function every hour. An Amazon Simple Notification Service (Amazon SNS) topic already exists in the AWS account. The DevOps engineer has subscribed to the SNS topic to receive notifications.

The DevOps engineer needs to receive a notification as soon as possible when drift is detected in this specific stack configuration.

Which solution will meet these requirements?

- A. Configure the existing EventBridge rule to also target the SNS topic. Configure an SNS subscription filter policy to match the CloudFormation stack.
- B. Attach the subscription filter policy to the SNS topic.
- C. Create a second Lambda function to query the CloudFormation API for the drift detection results for the stack. Configure the second Lambda function to publish a message to the SNS topic. If drift is detected, adjust the existing EventBridge rule to also target the second Lambda function.
- D. Configure Amazon GuardDuty in the account with drift detection for all CloudFormation stacks.

- E. Create a second EventBridge rule that reacts to the GuardDuty drift detection event finding for the specific CloudFormation stack
- F. Configure the SNS topic as a target of the second EventBridge rule.
- G. Configure AWS Config in the account
- H. Use the cloudformation-stack-drift-detection-check-managed-rule
- I. Create a second EventBridge rule that reacts to a compliance change event for the CloudFormation stack
- J. Configure the SNS topic as a target of the second EventBridge rule.

**Answer: D**

**Explanation:**

A comprehensive and detailed explanation is:

? Option A is incorrect because EventBridge rules cannot filter events based on the message body or attributes of the target service. Therefore, configuring an SNS subscription filter policy to match the CloudFormation stack will not work. The SNS topic will receive all events from the EventBridge rule, regardless of the stack name or drift status.

? Option B is incorrect because it introduces unnecessary complexity and cost.

Creating a second Lambda function to query the CloudFormation API for the drift detection results is redundant, since CloudFormation already publishes drift detection events to EventBridge. Moreover, invoking two Lambda functions every hour will incur more charges than invoking one.

? Option C is incorrect because GuardDuty does not provide drift detection for CloudFormation stacks. GuardDuty is a threat detection service that monitors for malicious activity and unauthorized behavior in AWS accounts and workloads. It does not monitor or report on configuration changes or drifts in CloudFormation stacks.

? Option D is correct because it leverages AWS Config and its managed rule for drift detection. AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. It can detect configuration changes and drifts in CloudFormation stacks using the cloudformation-stack-drift-detection-check-managed-rule. This rule triggers an AWS Config event when a stack drifts from its expected template configuration. By creating a second EventBridge rule that reacts to this event for the specific stack, the DevOps engineer can configure the SNS topic as a target and receive a notification as soon as possible when drift is detected.

References:

- ? AWS Config
- ? Amazon SNS subscription filter policies
- ? Amazon EventBridge rules

**NEW QUESTION 35**

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. 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 36**

A company deploys its corporate infrastructure on AWS across multiple AWS Regions and Availability Zones. The infrastructure is deployed on Amazon EC2 instances and connects with AWS IoT Greengrass devices. The company deploys additional resources on on-premises servers that are located in the corporate headquarters.

The company wants to reduce the overhead involved in maintaining and updating its resources. The company's DevOps team plans to use AWS Systems Manager to implement automated management and application of patches. The DevOps team confirms that Systems Manager is available in the Regions that the resources are deployed in. Systems Manager is also available in a Region near the corporate headquarters.

Which combination of steps must the DevOps team take to implement automated patch and configuration management across the company's EC2 instances, IoT devices, and on-premises infrastructure? (Select THREE.)

- A. Apply tags to all the EC2 instances
- B. AWS IoT Greengrass devices, and on-premises server
- C. Use Systems Manager Session Manager to push patches to all the tagged devices.
- D. Use Systems Manager Run Command to schedule patching for the EC2 instances, AWS IoT Greengrass devices, and on-premises servers.
- E. Use Systems Manager Patch Manager to schedule patching for the EC2 instances, AWS IoT Greengrass devices, and on-premises servers as a Systems Manager maintenance window task.
- F. Configure Amazon EventBridge to monitor Systems Manager Patch Manager for updates to patch baseline
- G. Associate Systems Manager Run Command with the event to initiate a patch action for all EC2 instances, AWS IoT Greengrass devices, and on-premises servers.
- H. Create an IAM instance profile for Systems Manager. Attach the instance profile to all the EC2 instances in the AWS account
- I. For the AWS IoT Greengrass devices and on-premises servers, create an IAM service role for Systems Manager.
- J. Generate a managed-instance activation. Use the Activation Code and Activation ID to install Systems Manager Agent (SSM Agent) on each server in the on-premises environment. Update the AWS IoT Greengrass IAM token exchange role. Use the role to deploy SSM Agent on all the IoT devices.

**Answer: CEF**

**Explanation:**

[https://aws.amazon.com/blogs/mt/how-to-centrally-manage-aws-iot-greengrass-devices-using-aws-systems-manager/?force\\_isolation=true](https://aws.amazon.com/blogs/mt/how-to-centrally-manage-aws-iot-greengrass-devices-using-aws-systems-manager/?force_isolation=true)

**NEW QUESTION 40**

A video-sharing company stores its videos in Amazon S3. The company has observed a sudden increase in video access requests, but the company does not know which videos are most popular. The company needs to identify the general access pattern for the video files. This pattern includes the number of users who access a certain file on a given day, as well as the number of times a file is accessed. 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 logs. The DevOps engineer manages the Kinesis consumer application, which also runs on Amazon EC2. Sudden increases of data cause the Kinesis consumer application to fall 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? How can the company meet these requirements with the LEAST amount of effort?

- A. Activate S3 server access logging
- B. Import the access logs into an Amazon Aurora database
- C. Use an Aurora SQL query to analyze the access patterns.
- D. Activate S3 server access logging
- E. Use Amazon Athena to create an external table with the log file
- F. Use Athena to create a SQL query to analyze the access patterns.
- G. Invoke an AWS Lambda function for every S3 object access event
- H. Configure the Lambda function to write the file access information, such as user, S3 bucket, and file key, to an Amazon Kinesis Data Analytics for SQL application
- I. S3 bucket, and file key, to an Amazon Aurora database
- J. Use an Aurora SQL query to analyze the access patterns.
- K. Record an Amazon CloudWatch Logs log message for every S3 object access event
- L. Configure a CloudWatch Logs log stream to write the file access information, such as user, S3 bucket, and file key, to an Amazon Kinesis Data Analytics for SQL application
- M. Perform a sliding window analysis.

**Answer: B**

**Explanation:**

Activating S3 server access logging and using Amazon Athena to create an external table with the log files is the easiest and most cost-effective way to analyze access patterns. This option requires minimal setup and allows for quick analysis of the access patterns with SQL queries. Additionally, Amazon Athena scales automatically to match the query load, so there is no need for additional infrastructure provisioning or management.

**NEW QUESTION 41**

A company has deployed a critical application in two AWS Regions. The application uses an Application Load Balancer (ALB) in both Regions. The company has Amazon Route 53 alias DNS records for both ALBs. The company uses Amazon Route 53 Application Recovery Controller to ensure that the application can fail over between the two Regions. The Route 53 ARC configuration includes a routing control for both Regions. The company uses Route 53 ARC to perform quarterly disaster recovery (DR) tests. During the most recent DR test, a DevOps engineer accidentally turned off both routing controls. The company needs to ensure that at least one routing control is turned on at all times. Which solution will meet these requirements?

- A. In Route 53 ARC, create a new assertion safety rule
- B. create a new assertion safety rule
- C. Apply the assertion safety rule to the two routing controls
- D. Configure the rule with the ATLEAST type with a threshold of 1.
- E. In Route 53 ARC, create a new gating safety rule
- F. Apply the assertion safety rule to the two routing controls
- G. Configure the rule with the OR type with a threshold of 1.
- H. In Route 53 ARC, create a new resource set
- I. Configure the resource set with an AWS: Route53: HealthCheck resource type
- J. Specify the ARNs of the two routing controls as the target resources
- K. Create a new readiness check for the resource set.
- L. In Route 53 ARC, create a new resource set
- M. Configure the resource set with an AWS: Route53RecoveryReadiness: DNSTargetResource resource type
- N. Add the domain names of the two Route 53 alias DNS records as the target resources
- O. Create a new readiness check for the resource set.

**Answer: A**

**Explanation:**

The correct solution is to create a new assertion safety rule in Route 53 ARC and apply it to the two routing controls. An assertion safety rule is a type of safety rule that ensures that a minimum number of routing controls are always enabled. The ATLEAST type of assertion safety rule specifies the minimum number of routing controls that must be enabled for the rule to evaluate as healthy. By setting the threshold to 1, the rule ensures that at least one routing control is always turned on. This prevents the scenario where both routing controls are accidentally turned off and the application becomes unavailable in both Regions. The other solutions are incorrect because they do not use safety rules to prevent both routing controls from being turned off. A gating safety rule is a type of safety rule that prevents routing control state changes that violate the rule logic. The OR type of gating safety rule specifies that one or more routing controls must be enabled for the rule to evaluate as healthy. However, this rule does not prevent a user from turning off both routing controls manually. A resource set is a collection of resources that are tested for readiness by Route 53 ARC. A readiness check is a test that verifies that all the resources in a resource set are operational. However, these concepts are not related to routing control states or safety rules. Therefore, creating a new resource set and a new readiness check will not ensure that at least one routing control is turned on at all times. References:  
 ? Routing control in Amazon Route 53 Application Recovery Controller  
 ? Viewing and updating routing control states in Route 53 ARC  
 ? Creating a control panel in Route 53 ARC  
 ? Creating safety rules in Route 53 ARC

**NEW QUESTION 44**

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

**NEW QUESTION 45**

A DevOps team uses AWS CodePipeline, AWS CodeBuild, and AWS CodeDeploy to deploy an application. The application is a REST API that uses AWS Lambda functions and Amazon API Gateway. Recent deployments have introduced errors that have affected many customers.

The DevOps team needs a solution that reverts to the most recent stable version of the application when an error is detected. The solution must affect the fewest customers possible.

Which solution will meet these requirements with the MOST operational efficiency?

- A. Set the deployment configuration in CodeDeploy to LambdaAllAtOnce. Configure automatic rollbacks on the deployment group. Create an Amazon CloudWatch alarm that detects HTTP Bad Gateway errors on API Gateway. Configure the deployment group to roll back when the number of alarms meets the alarm threshold.
- B. Set the deployment configuration in CodeDeploy to LambdaCanary10Percent10Minute.
- C. Configure automatic rollbacks on the deployment group. Create an Amazon CloudWatch alarm that detects HTTP Bad Gateway errors on API Gateway. Configure the deployment group to roll back when the number of alarms meets the alarm threshold.
- D. Set the deployment configuration in CodeDeploy to LambdaAllAtOnce. Configure manual rollbacks on the deployment group.
- E. Create an Amazon Simple Notification Service (Amazon SNS) topic to send notifications every time a deployment fails.
- F. Configure the SNS topic to invoke a new Lambda function that stops the current deployment and starts the most recent successful deployment.
- G. Set the deployment configuration in CodeDeploy to LambdaCanary10Percent10Minutes. Configure manual rollbacks on the deployment group. Create a metric filter on an Amazon CloudWatch log group for API Gateway to monitor HTTP Bad Gateway errors.
- H. Configure the metric filter to invoke a new Lambda function that stops the current deployment and starts the most recent successful deployment.

**Answer:** B

**Explanation:**

? Option A is incorrect because setting the deployment configuration to LambdaAllAtOnce means that the new version of the application will be deployed to all Lambda functions at once, affecting all customers. This does not meet the requirement of affecting the fewest customers possible. Moreover, configuring automatic rollbacks on the deployment group is not operationally efficient, as it requires manual intervention to fix the errors and redeploy the application.

? Option B is correct because setting the deployment configuration to LambdaCanary10Percent10Minutes means that the new version of the application will be deployed to 10 percent of the Lambda functions first, and then to the remaining 90 percent after 10 minutes. This minimizes the impact of errors on customers, as only 10 percent of them will be affected by a faulty deployment. Configuring automatic rollbacks on the deployment group also meets the requirement of reverting to the most recent stable version of the application when an error is detected. Creating a CloudWatch alarm that detects HTTP Bad Gateway errors on API Gateway is a valid way to monitor the health of the application and trigger a rollback if needed.

? Option C is incorrect because setting the deployment configuration to LambdaAllAtOnce means that the new version of the application will be deployed to all Lambda functions at once, affecting all customers. This does not meet the requirement of affecting the fewest customers possible. Moreover, configuring manual rollbacks on the deployment group is not operationally efficient, as it requires human intervention to stop the current deployment and start a new one. Creating an SNS topic to send notifications every time a deployment fails is not sufficient to detect errors in the application, as it does not monitor the API Gateway responses.

? Option D is incorrect because configuring manual rollbacks on the deployment group is not operationally efficient, as it requires human intervention to stop the current deployment and start a new one. Creating a metric filter on a CloudWatch log group for API Gateway to monitor HTTP Bad Gateway errors is a valid way to monitor the health of the application, but invoking a new Lambda function to perform a rollback is unnecessary and complex, as CodeDeploy already provides automatic rollback functionality.

References:

- ? AWS CodeDeploy Deployment Configurations
- ? [AWS CodeDeploy Rollbacks]
- ? Amazon CloudWatch Alarms

**NEW QUESTION 49**

A company recently migrated its legacy application from on-premises to AWS. The application is hosted on Amazon EC2 instances behind an Application Load Balancer which is behind Amazon API Gateway. The company wants to ensure users experience minimal disruptions during any deployment of a new version of the application. The company also wants to ensure it can quickly roll back updates if there is an issue.

Which solution will meet these requirements with MINIMAL changes to the application?

- A. Introduce changes as a separate environment parallel to the existing one. Configure API Gateway to use a canary release deployment to send a small subset of user traffic to the new environment.
- B. Introduce changes as a separate environment parallel to the existing one. Update the application's DNS alias records to point to the new environment.
- C. Introduce changes as a separate target group behind the existing Application Load Balancer. Configure API Gateway to route user traffic to the new target group in steps.
- D. Introduce changes as a separate target group behind the existing Application Load Balancer. Configure API Gateway to route all traffic to the Application Load Balancer which then sends the traffic to the new target group.

**Answer:** A

**Explanation:**

API Gateway supports canary deployment on a deployment stage before you direct all traffic to that stage. A parallel environment means we will create a new ALB and a target group that will target a new set of EC2 instances on which the newer version of the app will be deployed. So the canary setting associated to the new version of the API will connect with the new ALB instance which in turn will direct the traffic to the new EC2 instances on which the newer version of the application is deployed.

#### NEW QUESTION 50

A company's application development team uses Linux-based Amazon EC2 instances as bastion hosts. Inbound SSH access to the bastion hosts is restricted to specific IP addresses, as defined in the associated security groups. The company's security team wants to receive a notification if the security group rules are modified to allow SSH access from any IP address.

What should a DevOps engineer do to meet this requirement?

- A. Create an Amazon EventBridge rule with a source of aws.cloudtrail and the event name AuthorizeSecurityGroupIngres
- B. Define an Amazon Simple Notification Service (Amazon SNS) topic as the target.
- C. Enable Amazon GuardDuty and check the findings for security groups in AWS Security Hub
- D. Configure an Amazon EventBridge rule with a custom pattern that matches GuardDuty events with an output of NON\_COMPLIANT
- E. Define an Amazon Simple Notification Service (Amazon SNS) topic as the target.
- F. Create an AWS Config rule by using the restricted-ssh managed rule to check whether security groups disallow unrestricted incoming SSH traffic
- G. Configure automatic remediation to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic.
- H. Enable Amazon Inspector
- I. Include the Common Vulnerabilities and Exposures-1.1 rules package to check the security groups that are associated with the bastion host
- J. Configure Amazon Inspector to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic.

**Answer:** A

#### Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/monitor-security-group-changes-ec2/>

#### NEW QUESTION 55

A company has multiple development groups working in a single shared AWS account. The Senior Manager of the groups wants to be alerted via a third-party API call when the creation of resources approaches the service limits for the account.

Which solution will accomplish this with the LEAST amount of development effort?

- A. Create an Amazon CloudWatch Event rule that runs periodically and targets an AWS Lambda function
- B. Within the Lambda function, evaluate the current state of the AWS environment and compare deployed resource values to resource limits on the account
- C. Notify the Senior Manager if the account is approaching a service limit.
- D. Deploy an AWS Lambda function that refreshes AWS Trusted Advisor checks, and configure an Amazon CloudWatch Events rule to run the Lambda function periodically
- E. Create another CloudWatch Events rule with an event pattern matching Trusted Advisor events and a target Lambda function
- F. In the target Lambda function, notify the Senior Manager.
- G. Deploy an AWS Lambda function that refreshes AWS Personal Health Dashboard checks, and configure an Amazon CloudWatch Events rule to run the Lambda function periodically
- H. Create another CloudWatch Events rule with an event pattern matching Personal Health Dashboard events and a target Lambda function
- I. In the target Lambda function, notify the Senior Manager.
- J. Add an AWS Config custom rule that runs periodically, checks the AWS service limit status, and streams notifications to an Amazon SNS topic
- K. Deploy an AWS Lambda function that notifies the Senior Manager, and subscribe the Lambda function to the SNS topic.

**Answer:** B

#### Explanation:

To meet the requirements, the company needs to create a solution that alerts the Senior Manager when the creation of resources approaches the service limits for the account with the least amount of development effort. The company can use AWS Trusted Advisor, which is a service that provides best practice recommendations for cost optimization, performance, security, and service limits. The company can deploy an AWS Lambda function that refreshes Trusted Advisor checks, and configure an Amazon CloudWatch Events rule to run the Lambda function periodically. This will ensure that Trusted Advisor checks are up to date and reflect the current state of the account. The company can then create another CloudWatch Events rule with an event pattern matching Trusted Advisor events and a target Lambda function. The event pattern can filter for events related to service limit checks and their status. The target Lambda function can notify the Senior Manager via a third-party API call if the event indicates that the account is approaching or exceeding a service limit.

#### NEW QUESTION 56

A company provides an application to customers. The application has an Amazon API Gateway REST API that invokes an AWS Lambda function. On initialization, the Lambda function loads a large amount of data from an Amazon DynamoDB table. The data load process results in long cold-start times of 8-10 seconds. The DynamoDB table has DynamoDB Accelerator (DAX) configured.

Customers report that the application intermittently takes a long time to respond to requests. The application receives thousands of requests throughout the day. In the middle of the day, the application experiences 10 times more requests than at any other time of the day. Near the end of the day, the application's request volume decreases to 10% of its normal total.

A DevOps engineer needs to reduce the latency of the Lambda function at all times of the day.

Which solution will meet these requirements?

- A. Configure provisioned concurrency on the Lambda function with a concurrency value of 1. Delete the DAX cluster for the DynamoDB table.
- B. Configure reserved concurrency on the Lambda function with a concurrency value of 0.
- C. Configure provisioned concurrency on the Lambda function
- D. Configure AWS Application Auto Scaling on the Lambda function with provisioned concurrency values set to a minimum of 1 and a maximum of 100.
- E. Configure reserved concurrency on the Lambda function
- F. Configure AWS Application Auto Scaling on the API Gateway API with a reserved concurrency maximum value of 100.

**Answer:** C

#### Explanation:

The following are the steps that the DevOps engineer should take to reduce the latency of the Lambda function at all times of the day:

? Configure provisioned concurrency on the Lambda function.

? Configure AWS Application Auto Scaling on the Lambda function with provisioned concurrency values set to a minimum of 1 and a maximum of 100.

The provisioned concurrency setting ensures that there is always a minimum number of Lambda function instances available to handle requests. The Application Auto Scaling setting will automatically scale the number of Lambda function instances up or down based on the demand for the application.

This solution will ensure that the Lambda function is able to handle the increased load during the middle of the day, while also keeping the cold-start latency low.

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

? Option A is incorrect because it will not reduce the cold-start latency of the Lambda function.

- ? Option B is incorrect because it will not scale the number of Lambda function instances up or down based on demand.
- ? Option D is incorrect because it will only configure reserved concurrency on the API Gateway API, which will not affect the Lambda function.

#### NEW QUESTION 58

A company needs to implement failover for its application. The application includes an Amazon CloudFront distribution and a public Application Load Balancer (ALB) in an AWS Region. The company has configured the ALB as the default origin for the distribution. After some recent application outages, the company wants a zero-second RTO. The company deploys the application to a secondary Region in a warm standby configuration. A DevOps engineer needs to automate the failover of the application to the secondary Region so that HTTP GET requests meet the desired RTO. Which solution will meet these requirements?

- A. Create a second CloudFront distribution that has the secondary ALB as the default origin
- B. Create Amazon Route 53 alias records that have a failover policy and Evaluate Target Health set to Yes for both CloudFront distribution
- C. Update the application to use the new record set.
- D. Create a new origin on the distribution for the secondary AL
- E. Create a new origin group
- F. Set the original ALB as the primary origin
- G. Configure the origin group to fail over for HTTP 5xx status code
- H. Update the default behavior to use the origin group.
- I. Create Amazon Route 53 alias records that have a failover policy and Evaluate Target Health set to Yes for both ALB
- J. Set the TTL of both records to
- K. Update the distribution's origin to use the new record set.
- L. Create a CloudFront function that detects HTTP 5xx status code
- M. Configure the function to return a 307 Temporary Redirect error response to the secondary ALB if the function detects 5xx status code
- N. Update the distribution's default behavior to send origin responses to the function.

**Answer: B**

#### Explanation:

To implement failover for the application to the secondary Region so that HTTP GET requests meet the desired RTO, the DevOps engineer should use the following solution:

? Create a new origin on the distribution for the secondary ALB. A CloudFront origin is the source of the content that CloudFront delivers to viewers. By creating a new origin for the secondary ALB, the DevOps engineer can configure CloudFront to route traffic to the secondary Region when the primary Region is unavailable<sup>1</sup>

? Create a new origin group. Set the original ALB as the primary origin. Configure the origin group to fail over for HTTP 5xx status codes. An origin group is a logical grouping of two origins: a primary origin and a secondary origin. By creating an origin group, the DevOps engineer can specify which origin CloudFront should use as a fallback when the primary origin fails. The DevOps engineer can also define which HTTP status codes should trigger a failover from the primary origin to the secondary origin. By setting the original ALB as the primary origin and configuring the origin group to fail over for HTTP 5xx status codes, the DevOps engineer can ensure that CloudFront will switch to the secondary ALB when the primary ALB returns server errors<sup>2</sup>

? Update the default behavior to use the origin group. A behavior is a set of rules that CloudFront applies when it receives requests for specific URLs or file types. The default behavior applies to all requests that do not match any other behaviors. By updating the default behavior to use the origin group, the DevOps engineer can enable failover routing for all requests that are sent to the distribution<sup>3</sup>

This solution will meet the requirements because it will automate the failover of the application to the secondary Region with zero-second RTO. When CloudFront receives an HTTP GET request, it will first try to route it to the primary ALB in the primary Region. If the primary ALB is healthy and returns a successful response, CloudFront will deliver it to the viewer. If the primary ALB is unhealthy or returns an HTTP 5xx status code, CloudFront will automatically route the request to the secondary ALB in the secondary Region and deliver its response to the viewer. The other options are not correct because they either do not provide zero-second RTO or do not work as expected. Creating a second CloudFront distribution that has the secondary ALB as the default origin and creating Amazon Route 53 alias records that have a failover policy is not a good option because it will introduce additional latency and complexity to the solution. Route 53 health checks and DNS propagation can take several minutes or longer, which means that viewers might experience delays or errors when accessing the application during a failover event. Creating Amazon Route 53 alias records that have a failover policy and Evaluate Target Health set to Yes for both ALBs and setting the TTL of both records to 0 is not a valid option because it will not work with CloudFront distributions. Route 53 does not support health checks for alias records that point to CloudFront distributions, so it cannot detect if an ALB behind a distribution is healthy or not. Creating a CloudFront function that detects HTTP 5xx status codes and returns a 307 Temporary Redirect error response to the secondary ALB is not a valid option because it will not provide zero-second RTO. A 307 Temporary Redirect error response tells viewers to retry their requests with a different URL, which means that viewers will have to make an additional request and wait for another response from CloudFront before reaching the secondary ALB.

References:

- ? 1: Adding, Editing, and Deleting Origins - Amazon CloudFront
- ? 2: Configuring Origin Failover - Amazon CloudFront
- ? 3: Creating or Updating a Cache Behavior - Amazon CloudFront

#### NEW QUESTION 59

A company's application teams use AWS CodeCommit repositories for their applications. The application teams have repositories in multiple AWS accounts. All accounts are in an organization in AWS Organizations. Each application team uses AWS IAM Identity Center (AWS Single Sign-On) configured with an external IdP to assume a developer IAM role. The developer role allows the application teams to use Git to work with the code in the repositories. A security audit reveals that the application teams can modify the main branch in any repository. A DevOps engineer must implement a solution that allows the application teams to modify the main branch of only the repositories that they manage. Which combination of steps will meet these requirements? (Select THREE.)

- A. Update the SAML assertion to pass the user's team name
- B. Update the IAM role's trust policy to add an access-team session tag that has the team name.
- C. Create an approval rule template for each team in the Organizations management account
- D. Associate the template with all the repositories
- E. Add the developer role ARN as an approver.
- F. Create an approval rule template for each account
- G. Associate the template with all repositories
- H. Add the "aws:ResourceTag/access-team": "\$ ;{aws:PrincipalTag/access-team}" condition to the approval rule template.
- I. For each CodeCommit repository, add an access-team tag that has the value set to the name of the associated team.
- J. Attach an SCP to the account
- K. Include the following statement:

```

    {
      "Effect": "Deny",
      "Action": [
        "codecommit:GitPush",
        "codecommit:PutFile",
        "codecommit:Merge*"
      ],
      "Resource": "*",
      "Condition": {
        "StringEqualsIfExists": {
          "codecommit:References": ["refs/heads/main"]
        },
        "StringNotEquals": {
          "aws:ResourceTag/access-team": "$ ;{aws:PrincipalTag/access-team}"
        },
        "Null": {
          "codecommit:References": "false"
        }
      }
    }
  }
}

```

L. Create an IAM permissions boundary in each account

M. Include the following statement: {

```

    "Effect": "Allow",
    "Action": [
      "codecommit:GitPush",
      "codecommit:PutFile",
      "codecommit:Merge*"
    ],
    "Resource": "*",
    "Condition": {
      "StringEqualsIfExists": {
        "codecommit:References": ["refs/heads/main"]
      },
      "StringNotEquals": {
        "aws:ResourceTag/access-team": "$ ;{aws:PrincipalTag/access-team}"
      },
      "Null": {
        "codecommit:References": "false"
      }
    }
  }
}

```

**Answer:** ADF

**Explanation:**

Short Explanation: To meet the requirements, the DevOps engineer should update the SAML assertion to pass the user's team name, update the IAM role's trust policy to add an access-team session tag that has the team name, create an IAM permissions boundary in each account, and for each CodeCommit repository, add an access-team tag that has the value set to the name of the associated team.

References:

? Updating the SAML assertion to pass the user's team name allows the DevOps engineer to use IAM tags to identify which team a user belongs to. This can help enforce fine-grained access control based on the user's team membership<sup>1</sup>.

? Updating the IAM role's trust policy to add an access-team session tag that has the team name allows the DevOps engineer to use IAM condition keys to restrict access based on the session tag value<sup>2</sup>. For example, the DevOps engineer can use the aws:PrincipalTag condition key to match the access-team tag of the user with the access-team tag of the repository<sup>3</sup>.

? Creating an IAM permissions boundary in each account allows the DevOps engineer to set the maximum permissions that an identity-based policy can grant to an IAM entity. An entity's permissions boundary allows it to perform only the actions that are allowed by both its identity-based policies and its permissions boundaries<sup>4</sup>. For example, the DevOps engineer can use a permissions boundary policy to limit the actions that a user can perform on CodeCommit repositories based on their access-team tag<sup>5</sup>.

? For each CodeCommit repository, adding an access-team tag that has the value set to the name of the associated team allows the DevOps engineer to use resource tags to identify which team manages a repository. This can help enforce fine-grained access control based on the resource tag value<sup>6</sup>.

? The other options are incorrect because:

**NEW QUESTION 63**

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. Configure a remediation

- action that uses a custom AWS Systems Manager Automation runbook to apply the Backup\_Frequency tag with a value of weekly.
- C. Set up AWS Config in the account
  - D. Use a managed rule that returns a compliance failure for EC2::Volume resources that do not have a Backup Frequency tag applied
  - E. Configure a remediation action that uses a custom AWS Systems Manager Automation runbook to apply the Backup\_Frequency tag with a value of weekly.
  - F. Turn on AWS CloudTrail in the account
  - G. Create an Amazon EventBridge rule that reacts to EBS CreateVolume event
  - H. Configure a custom AWS Systems Manager Automation runbook to apply the Backup\_Frequency tag with a value of weekly
  - I. Specify the runbook as the target of the rule.
  - J. Turn on AWS CloudTrail in the account
  - K. Create an Amazon EventBridge rule that reacts to EBS CreateVolume events or EBS ModifyVolume event
  - L. Configure a custom AWS Systems Manager Automation runbook to apply the Backup\_Frequency tag with a value of weekly
  - M. 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 68**

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 73**

A DevOps engineer needs to configure a blue green deployment for an existing three-tier application. The application runs on Amazon EC2 instances and uses an Amazon RDS database. The EC2 instances run behind an Application Load Balancer (ALB) and are in an Auto Scaling group.

The DevOps engineer has created a launch template and an Auto Scaling group for the blue environment. The DevOps engineer also has created a launch template and an Auto Scaling group for the green environment. Each Auto Scaling group deploys to a matching blue or green target group. The target group also specifies which software blue or green gets loaded on the EC2 instances. The ALB can be configured to send traffic to the blue environment's target group or the green environment's target group. An Amazon Route 53 record for www.example.com points to the ALB.

The deployment must move traffic all at once between the software on the blue environment's EC2 instances to the newly deployed software on the green environment's EC2 instances.

What should the DevOps engineer do to meet these requirements?

- A. Start a rolling restart to the Auto Scaling group for the green environment to deploy the new software on the green environment's EC2 instances. When the rolling restart is complete, use an AWS CLI command to update the ALB to send traffic to the green environment's target group.
- B. Use an AWS CLI command to update the ALB to send traffic to the green environment's target group.
- C. Then start a rolling restart of the Auto Scaling group for the green environment to deploy the new software on the green environment's EC2 instances.
- D. Update the launch template to deploy the green environment's software on the blue environment's EC2 instances. Keep the target groups and Auto Scaling groups unchanged in both environments. Perform a rolling restart of the blue environment's EC2 instances.
- E. Start a rolling restart of the Auto Scaling group for the green environment to deploy the new software on the green environment's EC2 instances. When the rolling restart is complete, update the Route 53 DNS to point to the green environment's endpoint on the ALB.

**Answer: A**

**Explanation:**

This solution will meet the requirements because it will use a rolling restart to gradually replace the EC2 instances in the green environment with new instances that have the new software version installed. A rolling restart is a process that terminates and launches instances in batches, ensuring that there is always a minimum number of healthy instances in service. This way, the green environment can be updated without affecting the availability or performance of the application. When the rolling restart is complete, the DevOps engineer can use an AWS CLI command to modify the listener rules of the ALB and change the default action to forward traffic to the green environment's target group. This will switch the traffic from the blue environment to the green environment all at once, as required by the question.

#### NEW QUESTION 74

A security review has identified that an AWS CodeBuild project is downloading a database population script from an Amazon S3 bucket using an unauthenticated request. The security team does not allow unauthenticated requests to S3 buckets for this project. How can this issue be corrected in the MOST secure manner?

- A. Add the bucket name to the AllowedBuckets section of the CodeBuild project setting
- B. Update the build spec to use the AWS CLI to download the database population script.
- C. Modify the S3 bucket settings to enable HTTPS basic authentication and specify a token
- D. Update the build spec to use cURL to pass the token and download the database population script.
- E. Remove unauthenticated access from the S3 bucket with a bucket policy
- F. Modify the service role for the CodeBuild project to include Amazon S3 access
- G. Use the AWS CLI to download the database population script.
- H. Remove unauthenticated access from the S3 bucket with a bucket policy
- I. Use the AWS CLI to download the database population script using an IAM access key and a secret access key.

**Answer: C**

#### Explanation:

A bucket policy is a resource-based policy that defines who can access a specific S3 bucket and what actions they can perform on it. By removing unauthenticated access from the bucket policy, you can prevent anyone without valid credentials from accessing the bucket. A service role is an IAM role that allows an AWS service, such as CodeBuild, to perform actions on your behalf. By modifying the service role for the CodeBuild project to include Amazon S3 access, you can grant the project permission to read and write objects in the S3 bucket. The AWS CLI is a command-line tool that allows you to interact with AWS services, such as S3, using commands in your terminal. By using the AWS CLI to download the database population script, you can leverage the service role credentials and encryption to secure the data transfer.

For more information, you can refer to these web pages:

? [Using bucket policies and user policies - Amazon Simple Storage Service]

? [Create a service role for CodeBuild - AWS CodeBuild]

? [AWS Command Line Interface]

#### NEW QUESTION 75

A DevOps engineer is working on a project that is hosted on Amazon Linux and has failed a security review. The DevOps manager has been asked to review the company buildspec.yaml file for an AWS CodeBuild project and provide recommendations. The buildspec. yaml file is configured as follows:

```
env:
  variables:
    AWS_ACCESS_KEY_ID: AKIAJF7BRFWJBA4GHXNA
    AWS_SECRET_ACCESS_KEY: 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
    AWS_DEFAULT_REGION: us-east-1
    DB_PASSWORD: kuj5RptFa3va
  phases:
    build:
      commands:
        - aws s3 cp s3://db-deploy-bucket/my.cnf.template /tmp/my.cnf
        - sed -i '' s/DB_PW/${DB_PASSWORD}/ /tmp/my.cnf
        - aws s3 cp s3://db-deploy-bucket/instance.key /tmp/instance.key
        - chmod 600 /tmp/instance.key
        - scp -i /tmp/instance.key /tmp/my.cnf root@10.25.15.23:/etc/my.cnf
        - ssh -i /tmp/instance.key root@10.25.15.23 /etc/init.d/mysqld restart
```

What changes should be recommended to comply with AWS security best practices? (Select THREE.)

- A. Add a post-build command to remove the temporary files from the container before termination to ensure they cannot be seen by other CodeBuild users.
- B. Update the CodeBuild project role with the necessary permissions and then remove the AWS credentials from the environment variable.
- C. Store the db\_password as a SecureString value in AWS Systems Manager Parameter Store and then remove the db\_password from the environment variables.
- D. Move the environment variables to the 'db.-deploy-bucket' Amazon S3 bucket, add a prebuild stage to download then export the variables.
- E. Use AWS Systems Manager run command versus scp and ssh commands directly to the instance.

**Answer: BCE**

#### Explanation:

B. Update the CodeBuild project role with the necessary permissions and then remove the AWS credentials from the environment variable. C. Store the DB\_PASSWORD as a SecureString value in AWS Systems Manager Parameter Store and then remove the DB\_PASSWORD from the environment variables. E. Use AWS Systems Manager run command versus scp and ssh commands directly to the instance.

#### NEW QUESTION 80

A company uses AWS Organizations to manage its AWS accounts. The company has a root OU that has a child OU. The root OU has an SCP that allows all actions on all resources. The child OU has an SCP that allows all actions for Amazon DynamoDB and AWS Lambda, and denies all other actions. The company has an AWS account that is named vendor-data in the child OU. A DevOps engineer has an IAM user that is attached to the AdministratorAccess IAM policy in the vendor-data account. The DevOps engineer attempts to launch an Amazon EC2 instance in the vendor-data account but receives an access denied error. Which change should the DevOps engineer make to launch the EC2 instance in the vendor-data account?

- A. Attach the AmazonEC2FullAccess IAM policy to the IAM user.
- B. Create a new SCP that allows all actions for Amazon EC2. Attach the SCP to the vendor-data account.
- C. Update the SCP in the child OU to allow all actions for Amazon EC2.
- D. Create a new SCP that allows all actions for Amazon EC2. Attach the SCP to the root OU.

**Answer: C**

#### Explanation:

The correct answer is C. Updating the SCP in the child OU to allow all actions for Amazon EC2 will enable the DevOps engineer to launch the EC2 instance in the

vendor-data account. SCPs are applied to OUs and accounts in a hierarchical manner, meaning that the SCPs attached to the parent OU are inherited by the child OU and accounts. Therefore, the SCP in the child OU overrides the SCP in the root OU and denies all actions except for DynamoDB and Lambda. By adding EC2 to the allowed actions in the child OU's SCP, the DevOps engineer can access EC2 resources in the vendor-data account.

Option A is incorrect because attaching the AmazonEC2FullAccess IAM policy to the IAM user will not grant the user access to EC2 resources. IAM policies are evaluated after SCPs, so even if the IAM policy allows EC2 actions, the SCP will still deny them.

Option B is incorrect because creating a new SCP that allows all actions for EC2 and attaching it to the vendor-data account will not work. SCPs are not cumulative, meaning that only one SCP is applied to an account at a time. The SCP attached to the account will be the SCP attached to the OU that contains the account. Therefore, option B will not change the SCP that is applied to the vendor-data account.

Option D is incorrect because creating a new SCP that allows all actions for EC2 and attaching it to the root OU will not work. As explained earlier, the SCP in the child OU overrides the SCP in the root OU and denies all actions except for DynamoDB and Lambda. Therefore, option D will not affect the SCP that is applied to the vendor-data account.

#### NEW QUESTION 84

A company hosts applications in its AWS account. Each application logs to an individual Amazon CloudWatch log group. The company's CloudWatch costs for ingestion are increasing.

A DevOps engineer needs to identify which applications are the source of the increased logging costs.

Which solution will meet these requirements?

- A. Use CloudWatch metrics to create a custom expression that identifies the CloudWatch log groups that have the most data being written to them.
- B. Use CloudWatch Logs Insights to create a set of queries for the application log groups to identify the number of logs written for a period of time.
- C. Use AWS Cost Explorer to generate a cost report that details the cost for CloudWatch usage.
- D. Use AWS CloudTrail to filter for CreateLogStream events for each application.

**Answer: C**

#### Explanation:

The correct answer is C.

A comprehensive and detailed explanation is:

? Option A is incorrect because using CloudWatch metrics to create a custom expression that identifies the CloudWatch log groups that have the most data being written to them is not a valid solution. CloudWatch metrics do not provide information about the size or volume of data being ingested by CloudWatch logs.

CloudWatch metrics only provide information about the number of events, bytes, and errors that occur within a log group or stream. Moreover, creating a custom expression with CloudWatch metrics would require using the search\_web tool, which is not necessary for this use case.

? Option B is incorrect because using CloudWatch Logs Insights to create a set of queries for the application log groups to identify the number of logs written for a period of time is not a valid solution. CloudWatch Logs Insights can help analyze and filter log events based on patterns and expressions, but it does not provide information about the cost or billing of CloudWatch logs. CloudWatch Logs Insights also charges based on the amount of data scanned by each query, which could increase the logging costs further.

? Option C is correct because using AWS Cost Explorer to generate a cost report that details the cost for CloudWatch usage is a valid solution. AWS Cost Explorer is a tool that helps visualize, understand, and manage AWS costs and usage over time. AWS Cost Explorer can generate custom reports that show the breakdown of costs by service, region, account, tag, or any other dimension. AWS Cost Explorer can also filter and group costs by usage type, which can help identify the specific CloudWatch log groups that are the source of the increased logging costs.

? Option D is incorrect because using AWS CloudTrail to filter for CreateLogStream events for each application is not a valid solution. AWS CloudTrail is a service that records API calls and account activity for AWS services, including CloudWatch logs. However, AWS CloudTrail does not provide information about the cost or billing of CloudWatch logs. Filtering for CreateLogStream events would only show when a new log stream was created within a log group, but not how much data was ingested or stored by that log stream.

References:

? CloudWatch Metrics

? CloudWatch Logs Insights

? AWS Cost Explorer

? AWS CloudTrail

#### NEW QUESTION 86

A company is implementing AWS CodePipeline to automate its testing process. The company wants to be notified when the execution state fails and used the following custom event pattern in Amazon EventBridge:

```
{
  "source": [
    "aws.codepipeline"
  ],
  "detail-type": [
    "CodePipeline Action Execution State Change"
  ],
  "detail": {
    "state": [
      "FAILED"
    ]
  },
  "type": {
    "category": ["Approval"]
  }
}
```

Which type of events will match this event pattern?

- A. Failed deploy and build actions across all the pipelines

- B. All rejected or failed approval actions across all the pipelines
- C. All the events across all pipelines
- D. Approval actions across all the pipelines

**Answer: B**

**Explanation:**

Action-level states in events  
 Action state Description  
 STARTED The action is currently running. SUCCEEDED The action was completed successfully.  
 FAILED For Approval actions, the FAILED state means the action was either rejected by the reviewer or failed due to an incorrect action configuration.  
 CANCELED The action was canceled because the pipeline structure was updated.

**NEW QUESTION 88**

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

- A. Configure the AWS Config ec2-volume-inuse-check managed rule with a configuration changes trigger type and an Amazon EC2 volume resource target
- B. Create a new Amazon CloudWatch Events rule scheduled to execute an AWS Lambda function in 14 days to delete the specified EBS volume.
- C. Use Amazon EC2 and Amazon Data Lifecycle Manager to configure a volume lifecycle policy
- D. Set the interval period for unattached EBS volumes to 14 days and set the retention rule to delete
- E. Set the policy target volumes as \*
- F. Create an Amazon CloudWatch Events rule to execute an AWS Lambda function daily
- G. The Lambda function should find unattached EBS volumes and tag them with the current date, and delete unattached volumes that have tags with dates that are more than 14 days old.
- H. Use AWS Trusted Advisor to detect EBS volumes that have been detached for more than 14 days
- I. Execute an AWS Lambda function that creates a snapshot and then deletes the EBS volume.

**Answer: C**

**Explanation:**

The requirement is to create automation that deletes unattached EBS volumes that have been unattached for 14 days. To do this, the DevOps engineer needs to use the following steps:

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

**NEW QUESTION 92**

A company has its AWS accounts in an organization in AWS Organizations. AWS Config is manually configured in each AWS account. The company needs to implement a solution to centrally configure AWS Config for all accounts in the organization. The solution also must record resource changes to a central account. Which combination of actions should a DevOps engineer perform to meet these requirements? (Choose two.)

- A. Configure a delegated administrator account for AWS Config
- B. Enable trusted access for AWS Config in the organization.
- C. Configure a delegated administrator account for AWS Config
- D. Create a service-linked role for AWS Config in the organization's management account.
- E. Create an AWS CloudFormation template to create an AWS Config aggregator
- F. Configure a CloudFormation stack set to deploy the template to all accounts in the organization.
- G. Create an AWS Config organization aggregator in the organization's management account
- H. Configure data collection from all AWS accounts in the organization and from all AWS Regions.
- I. Create an AWS Config organization aggregator in the delegated administrator account
- J. Configure data collection from all AWS accounts in the organization and from all AWS Regions.

**Answer: AE**

**Explanation:**

<https://aws.amazon.com/blogs/mt/org-aggregator-delegated-admin/> <https://docs.aws.amazon.com/organizations/latest/userguide/services-that-can-integrate-config.html>

**NEW QUESTION 94**

A company wants to deploy a workload on several hundred Amazon EC2 instances. The company will provision the EC2 instances in an Auto Scaling group by using a launch template. The workload will pull files from an Amazon S3 bucket, process the data, and put the results into a different S3 bucket. The EC2 instances must have least-privilege permissions and must use temporary security credentials. Which combination of steps will meet these requirements? (Select TWO.)

- A. Create an IAM role that has the appropriate permissions for S3 bucket
- B. Add the IAM role to an instance profile.
- C. Update the launch template to include the IAM instance profile.
- D. Create an IAM user that has the appropriate permissions for Amazon S3. Generate a secret key and token.
- E. Create a trust anchor and profile
- F. Attach the IAM role to the profile.

- G. Update the launch template
- H. Modify the user data to use the new secret key and token.

**Answer:** AB

**Explanation:**

To meet the requirements of deploying a workload on several hundred EC2 instances with least-privilege permissions and temporary security credentials, the company should use an IAM role and an instance profile. An IAM role is a way to grant permissions to an entity that you trust, such as an EC2 instance. An instance profile is a container for an IAM role that you can use to pass role information to an EC2 instance when the instance starts. By using an IAM role and an instance profile, the EC2 instances can automatically receive temporary security credentials from the AWS Security Token Service (STS) and use them to access the S3 buckets. This way, the company does not need to manage or rotate any long-term credentials, such as IAM users or access keys.

To use an IAM role and an instance profile, the company should create an IAM role that has the appropriate permissions for S3 buckets. The permissions should allow the EC2 instances to read from the source S3 bucket and write to the destination S3 bucket. The company should also create a trust policy for the IAM role that specifies that EC2 is allowed to assume the role. Then, the company should add the IAM role to an instance profile. An instance profile can have only one IAM role, so the company does not need to create multiple roles or profiles for this scenario.

Next, the company should update the launch template to include the IAM instance profile. A launch template is a way to save launch parameters for EC2 instances, such as the instance type, security group, user data, and IAM instance profile. By using a launch template, the company can ensure that all EC2 instances in the Auto Scaling group have consistent configuration and permissions. The company should specify the name or ARN of the IAM instance profile in the launch template. This way, when the Auto Scaling group launches new EC2 instances based on the launch template, they will automatically receive the IAM role and its permissions through the instance profile.

The other options are not correct because they do not meet the requirements or follow best practices. Creating an IAM user and generating a secret key and token is not a good option because it involves managing long-term credentials that need to be rotated regularly. Moreover, embedding credentials in user data is not secure because user data is visible to anyone who can describe the EC2 instance. Creating a trust anchor and profile is not a valid option because trust anchors are used for certificate-based authentication, not for IAM roles or instance profiles. Modifying user data to use a new secret key and token is also not a good option because it requires updating user data every time the credentials change, which is not scalable or efficient.

References:

- ? 1: AWS Certified DevOps Engineer - Professional Certification | AWS Certification | AWS
- ? 2: DevOps Resources - Amazon Web Services (AWS)
- ? 3: Exam Readiness: AWS Certified DevOps Engineer - Professional
- ? : IAM Roles for Amazon EC2 - AWS Identity and Access Management
- ? : Working with Instance Profiles - AWS Identity and Access Management
- ? : Launching an Instance Using a Launch Template - Amazon Elastic Compute Cloud
- ? : Temporary Security Credentials - AWS Identity and Access Management

**NEW QUESTION 97**

A DevOps engineer is building a continuous deployment pipeline for a serverless application that uses AWS Lambda functions. The company wants to reduce the customer impact of an unsuccessful deployment. The company also wants to monitor for issues.

Which deploy stage configuration will meet these requirements?

- A. Use an AWS Serverless Application Model (AWS SAM) template to define the serverless application
- B. Use AWS CodeDeploy to deploy the Lambda functions with the Canary10Percent15Minutes Deployment Preference Type
- C. Use Amazon CloudWatch alarms to monitor the health of the functions.
- D. Use AWS CloudFormation to publish a new stack update, and include Amazon CloudWatch alarms on all resource
- E. Set up an AWS CodePipeline approval action for a developer to verify and approve the AWS CloudFormation change set.
- F. Use AWS CloudFormation to publish a new version on every stack update, and include Amazon CloudWatch alarms on all resource
- G. Use the RoutingConfig property of the AWS::Lambda::Alias resource to update the traffic routing during the stack update.
- H. Use AWS CodeBuild to add sample event payloads for testing to the Lambda function
- I. Publish a new version of the functions, and include Amazon CloudWatch alarm
- J. Update the production alias to point to the new version
- K. Configure rollbacks to occur when an alarm is in the ALARM state.

**Answer:** D

**Explanation:**

Use routing configuration on an alias to send a portion of traffic to a second function version. For example, you can reduce the risk of deploying a new version by configuring the alias to send most of the traffic to the existing version, and only a small percentage of traffic to the new version.

<https://docs.aws.amazon.com/lambda/latest/dg/configuration-aliases.html>

The following are the steps involved in the deploy stage configuration that will meet the requirements:

- ? Use AWS CodeBuild to add sample event payloads for testing to the Lambda functions.
- ? Publish a new version of the functions, and include Amazon CloudWatch alarms.
- ? Update the production alias to point to the new version.
- ? Configure rollbacks to occur when an alarm is in the ALARM state.

This configuration will help to reduce the customer impact of an unsuccessful deployment

by deploying the new version of the functions to a staging environment first. This will allow the DevOps engineer to test the new version of the functions before deploying it to production.

The configuration will also help to monitor for issues by including Amazon CloudWatch alarms. These alarms will alert the DevOps engineer if there are any problems with the new version of the functions.

**NEW QUESTION 102**

A company uses an organization in AWS Organizations that has all features enabled. The company uses AWS Backup in a primary account and uses an AWS Key Management Service (AWS KMS) key to encrypt the backups.

The company needs to automate a cross-account backup of the resources that AWS Backup backs up in the primary account. The company configures cross-account backup in the Organizations management account. The company creates a new AWS account in the organization and configures an AWS Backup backup vault in the new account. The company creates a KMS key in the new account to encrypt the backups. Finally, the company configures a new backup plan in the primary account. The destination for the new backup plan is the backup vault in the new account.

When the AWS Backup job in the primary account is invoked, the job creates backups in the primary account. However, the backups are not copied to the new account's backup vault.

Which combination of steps must the company take so that backups can be copied to the new account's backup vault? (Select TWO.)

- A. Edit the backup vault access policy in the new account to allow access to the primary account.
- B. Edit the backup vault access policy in the primary account to allow access to the new account.
- C. Edit the backup vault access policy in the primary account to allow access to the KMS key in the new account.
- D. Edit the key policy of the KMS key in the primary account to share the key with the new account.
- E. Edit the key policy of the KMS key in the new account to share the key with the primary account.

**Answer:** AE

**Explanation:**

To enable cross-account backup, the company needs to grant permissions to both the backup vault and the KMS key in the destination account. The backup vault access policy in the destination account must allow the primary account to copy backups into the vault. The key policy of the KMS key in the destination account must allow the primary account to use the key to encrypt and decrypt the backups. These steps are described in the AWS documentation<sup>12</sup>. Therefore, the correct answer is A and E.

References:

- ? 1: Creating backup copies across AWS accounts - AWS Backup
- ? 2: Using AWS Backup with AWS Organizations - AWS Backup

**NEW QUESTION 103**

A company uses AWS CodeArtifact to centrally store Python packages. The CodeArtifact repository is configured with the following repository policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "codeartifact:DescribePackageVersion",
        "codeartifact:DescribeRepository",
        "codeartifact:GetPackageVersionReadme",
        "codeartifact:GetRepositoryEndpoint",
        "codeartifact:ListPackageVersionAssets",
        "codeartifact:ListPackageVersionDependencies",
        "codeartifact:ListPackageVersions",
        "codeartifact:ListPackages",
        "codeartifact:ReadFromRepository"
      ],
      "Effect": "Allow",
      "Resource": "*",
      "Principal": "*",
      "Condition": {
        "StringEquals": {
          "aws:PrincipalOrgID": [
            "o-xxxxxxxxxxxx"
          ]
        }
      }
    }
  ]
}
```

A development team is building a new project in an account that is in an organization in AWS Organizations. The development team wants to use a Python library that has already been stored in the CodeArtifact repository in the organization. The development team uses AWS CodePipeline and AWS CodeBuild to build the new application. The CodeBuild job that the development team uses to build the application is configured to run in a VPC. Because of compliance requirements the VPC has no internet connectivity.

The development team creates the VPC endpoints for CodeArtifact and updates the CodeBuild buildspec yml file. However, the development team cannot download the Python library from the repository.

Which combination of steps should a DevOps engineer take so that the development team can use Code Artifact? (Select TWO.)

- A. Create an Amazon S3 gateway endpoint. Update the route tables for the subnets that are running the CodeBuild job.
- B. Update the repository policy's Principal statement to include the ARN of the role that the CodeBuild project uses.
- C. Share the CodeArtifact repository with the organization by using AWS Resource Access Manager (AWS RAM).
- D. Update the role that the CodeBuild project uses so that the role has sufficient permissions to use the CodeArtifact repository.
- E. Specify the account that hosts the repository as the delegated administrator for CodeArtifact in the organization.

**Answer:** AD

**Explanation:**

"AWS CodeArtifact operates in multiple Availability Zones and stores artifact data and metadata in Amazon S3 and Amazon DynamoDB. Your encrypted data is redundantly stored across multiple facilities and multiple devices in each facility, making it highly available and highly durable."

<https://aws.amazon.com/codeartifact/features/> With no internet connectivity, a gateway endpoint becomes necessary to access S3.

**NEW QUESTION 105**

A DevOps engineer has implemented a CI/CO pipeline to deploy an AWS CloudFormation template that provisions a web application. The web application consists of an Application Load Balancer (ALB), a target group, a launch template that uses an Amazon Linux 2 AMI, an Auto Scaling group of Amazon EC2 instances, a security group, and an Amazon RDS for MySQL database. The launch template includes user data that specifies a script to install and start the application.

The initial deployment of the application was successful. The DevOps engineer made changes to update the version of the application with the user data. The

CI/CD pipeline has deployed a new version of the template. However, the health checks on the ALB are now failing. The health checks have marked all targets as unhealthy.

During investigation, the DevOps engineer notices that the CloudFormation stack has a status of UPDATE\_COMPLETE. However, when the DevOps engineer connects to one of the EC2 instances and checks /var/log messages, the DevOps engineer notices that the Apache web server failed to start successfully because of a configuration error.

How can the DevOps engineer ensure that the CloudFormation deployment will fail if the user data fails to successfully finish running?

- A. Use the cfn-signal helper script to signal success or failure to CloudFormation. Use the WaitOnResourceSignals update policy within the CloudFormation template. Set an appropriate timeout for the update policy.
- B. Create an Amazon CloudWatch alarm for the UnhealthyHostCount metric.
- C. Include an appropriate alarm threshold for the target group. Create an Amazon Simple Notification Service (Amazon SNS) topic as the target to signal success or failure to CloudFormation.
- D. Create a lifecycle hook on the Auto Scaling group by using the AWS AutoScaling LifecycleHook resource. Create an Amazon Simple Notification Service (Amazon SNS) topic as the target to signal success or failure to CloudFormation. Set an appropriate timeout on the lifecycle hook.
- E. Use the Amazon CloudWatch agent to stream the cloud-init logs. Create a subscription filter that includes an AWS Lambda function with an appropriate invocation timeout. Configure the Lambda function to use the SignalResource API operation to signal success or failure to CloudFormation.

**Answer:** A

**Explanation:**

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-updatepolicy.html>

**NEW QUESTION 108**

An ecommerce company has chosen AWS to host its new platform. The company's DevOps team has started building an AWS Control Tower landing zone. The DevOps team has set the identity store within AWS IAM Identity Center (AWS Single Sign-On) to an external identity provider (IdP) and has configured SAML 2.0. The DevOps team wants a robust permission model that applies the principle of least privilege. The model must allow the team to build and manage only the team's own resources.

Which combination of steps will meet these requirements? (Choose three.)

- A. Create IAM policies that include the required permission.
- B. Include the aws:PrincipalTag condition key.
- C. Create permission set.
- D. Attach an inline policy that includes the required permissions and uses the aws:PrincipalTag condition key to scope the permissions.
- E. Create a group in the IdP.
- F. Place users in the group.
- G. Assign the group to accounts and the permission sets in IAM Identity Center.
- H. Create a group in the IdP.
- I. Place users in the group.
- J. Assign the group to OUs and IAM policies.
- K. Enable attributes for access control in IAM Identity Center.
- L. Apply tags to user.
- M. Map the tags as key-value pairs.
- N. Enable attributes for access control in IAM Identity Center.
- O. Map attributes from the IdP as key-value pairs.

**Answer:** BCF

**Explanation:**

Using the principalTag in the Permission Set inline policy, a logged-in user belonging to a specific AD group in the IDP can be permitted access to perform operations on certain resources if their group matches the group used in the PrincipalTag. Basically, you are narrowing the scope of privileges assigned via Permission policies conditionally based on whether the logged-in user belongs to a specific AD Group in IDP. The mapping of the AD group to the request attributes can be done using SSO attributes where we can pass other attributes like the SAML token as well.

<https://docs.aws.amazon.com/singlesignon/latest/userguide/abac.html>

**NEW QUESTION 111**

A company wants to use a grid system for a proprietary enterprise in-memory data store on top of AWS. This system can run in multiple server nodes in any Linux-based distribution. The system must be able to reconfigure the entire cluster every time a node is added or removed. When adding or removing nodes, an /etc./cluster/nodes config file must be updated listing the IP addresses of the current node members of that cluster.

The company wants to automate the task of adding new nodes to a cluster. What can a DevOps engineer do to meet these requirements?

- A. Use AWS OpsWorks Stacks to layer the server nodes of that cluster.
- B. Create a Chef recipe that populates the content of the /etc./cluster/nodes config file and restarts the service by using the current members of the layer.
- C. Assign that recipe to the Configure lifecycle event.
- D. Put the file nodes config in version control.
- E. Create an AWS CodeDeploy deployment configuration and deployment group based on an Amazon EC2 tag value for the cluster node.
- F. When adding a new node to the cluster, update the file with all tagged instances and make a commit in version control.
- G. Deploy the new file and restart the services.
- H. Create an Amazon S3 bucket and upload a version of the /etc./cluster/nodes config file. Create a crontab script that will poll for that S3 file and download it frequently.
- I. Use a process manager such as Monit or systemd, to restart the cluster services when it detects that the new file was modified.
- J. When adding a node to the cluster, edit the file's most recent members. Upload the new file to the S3 bucket.
- K. Create a user data script that lists all members of the current security group of the cluster and automatically updates the /etc./cluster/nodes config file whenever a new instance is added to the cluster.
- L. Update the file whenever a new instance is added to the cluster.

**Answer:** A

**Explanation:**

You can run custom recipes manually, but the best approach is usually to have AWS OpsWorks Stacks run them automatically. Every layer has a set of built-in recipes assigned each of five lifecycle events—Setup, Configure, Deploy, Undeploy, and Shutdown. Each time an event occurs for an instance, AWS OpsWorks Stacks runs the associated recipes for each of the instance's layers, which handle the corresponding tasks. For example, when an instance finishes booting, AWS OpsWorks Stacks triggers a Setup event. This event runs the associated layer's Setup recipes, which typically handle tasks such as installing and configuring

packages

#### NEW QUESTION 115

A company manages an application that stores logs in Amazon CloudWatch Logs. The company wants to archive the logs to an Amazon S3 bucket. Logs are rarely accessed after 90 days and must be retained for 10 years.

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

- A. Configure a CloudWatch Logs subscription filter to use AWS Glue to transfer all logs to an S3 bucket.
- B. Configure a CloudWatch Logs subscription filter to use Amazon Kinesis Data Firehose to stream all logs to an S3 bucket.
- C. Configure a CloudWatch Logs subscription filter to stream all logs to an S3 bucket.
- D. Configure the S3 bucket lifecycle policy to transition logs to S3 Glacier after 90 days and to expire logs after 3,650 days.
- E. Configure the S3 bucket lifecycle policy to transition logs to Reduced Redundancy after 90 days and to expire logs after 3,650 days.

**Answer:** BD

#### Explanation:

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

#### NEW QUESTION 119

An application running on a set of Amazon EC2 instances in an Auto Scaling group requires a configuration file to operate. The instances are created and maintained with AWS CloudFormation. A DevOps engineer wants the instances to have the latest configuration file when launched and wants changes to the configuration file to be reflected on all the instances with a minimal delay when the CloudFormation template is updated. Company policy requires that application configuration files be maintained along with AWS infrastructure configuration files in source control.

Which solution will accomplish this?

- A. In the CloudFormation template add an AWS Config rule
- B. Place the configuration file content in the rule's InputParameters property and set the Scope property to the EC2 Auto Scaling group
- C. Add an AWS Systems Manager Resource Data Sync resource to the template to poll for updates to the configuration.
- D. In the CloudFormation template add an EC2 launch template resource
- E. Place the configuration file content in the launch template
- F. Configure the cfn-init script to run when the instance is launched and configure the cfn-hup script to poll for updates to the configuration.
- G. In the CloudFormation template add an EC2 launch template resource
- H. Place the configuration file content in the launch template
- I. Add an AWS Systems Manager Resource Data Sync resource to the template to poll for updates to the configuration.
- J. In the CloudFormation template add CloudFormation metadata
- K. Place the configuration file content in the metadata
- L. Configure the cfn-init script to run when the instance is launched and configure the cfn-hup script to poll for updates to the configuration.

**Answer:** D

#### Explanation:

Use the `AWS::CloudFormation::Init` type to include metadata on an Amazon EC2 instance for the `cfn-init` helper script. If your template calls the `cfn-init` script, the script looks for resource metadata rooted in the `AWS::CloudFormation::Init` metadata key. Reference:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-init.html>

#### NEW QUESTION 122

A development team wants to use AWS CloudFormation stacks to deploy an application. However, the developer IAM role does not have the required permissions to provision the resources that are specified in the AWS CloudFormation template. A DevOps engineer needs to implement a solution that allows the developers to deploy the stacks. The solution must follow the principle of least privilege.

Which solution will meet these requirements?

- A. Create an IAM policy that allows the developers to provision the required resource
- B. Attach the policy to the developer IAM role.
- C. Create an IAM policy that allows full access to AWS CloudFormation
- D. Attach the policy to the developer IAM role.
- E. Create an AWS CloudFormation service role that has the required permission
- F. Grant the developer IAM role a `cloudformation:*` action
- G. Use the new service role during stack deployments.
- H. Create an AWS CloudFormation service role that has the required permission
- I. Grant the developer IAM role the `iam:PassRole` permission
- J. Use the new service role during stack deployments.

**Answer:** D

#### Explanation:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-iam-service-role.html>

#### NEW QUESTION 124

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 129**

AnyCompany is using AWS Organizations to create and manage multiple AWS accounts. AnyCompany recently acquired a smaller company, Example Corp. During the acquisition process, Example Corp's single AWS account joined AnyCompany's management account through an Organizations invitation. AnyCompany moved the new member account under an OU that is dedicated to Example Corp. AnyCompany's DevOps engineer has an IAM user that assumes a role that is named OrganizationAccountAccessRole to access member accounts. This role is configured with a full access policy. When the DevOps engineer tries to use the AWS Management Console to assume the role in Example Corp's new member account, the DevOps engineer receives the following error message: "Invalid information in one or more fields. Check your information or contact your administrator."

Which solution will give the DevOps engineer access to the new member account?

- A. In the management account, grant the DevOps engineer's IAM user permission to assume the OrganizationAccountAccessRole IAM role in the new member account.
- B. In the management account, create a new SCP. In the SCP, grant the DevOps engineer's IAM user full access to all resources in the new member account.
- C. Attach the SCP to the OU that contains the new member account.
- D. In the new member account, create a new IAM role that is named OrganizationAccountAccessRole.
- E. Attach the AdministratorAccess AWS managed policy to the role.
- F. In the role's trust policy, grant the management account permission to assume the role.
- G. In the new member account, edit the trust policy for the OrganizationAccountAccessRole IAM role.
- H. Grant the management account permission to assume the role.

**Answer: C**

**Explanation:**

The problem is that the DevOps engineer cannot assume the OrganizationAccountAccessRole IAM role in the new member account that joined AnyCompany's management account through an Organizations invitation. The solution is to create a new IAM role with the same name and trust policy in the new member account.

? Option A is incorrect, as it does not address the root cause of the error. The DevOps engineer's IAM user already has permission to assume the OrganizationAccountAccessRole IAM role in any member account, as this is the default role name that AWS Organizations creates when a new account joins an organization. The error occurs because the new member account does not have this role, as it was not created by AWS Organizations.

? Option B is incorrect, as it does not address the root cause of the error. An SCP is a policy that defines the maximum permissions for account members of an organization or organizational unit (OU). An SCP does not grant permissions to IAM users or roles, but rather limits the permissions that identity-based policies or resource-based policies grant to them. An SCP also does not affect how IAM roles are assumed by other principals.

? Option C is correct, as it addresses the root cause of the error. By creating a new IAM role with the same name and trust policy as the OrganizationAccountAccessRole IAM role in the new member account, the DevOps engineer can assume this role and access the account. The new role should have the AdministratorAccess AWS managed policy attached, which grants full access to all AWS resources in the account. The trust policy should allow the management account to assume the role, which can be done by specifying the management account ID as a principal in the policy statement.

? Option D is incorrect, as it assumes that the new member account already has the OrganizationAccountAccessRole IAM role, which is not true. The new member account does not have this role, as it was not created by AWS Organizations. Editing the trust policy of a non-existent role will not solve the problem.

**NEW QUESTION 132**

An online retail company based in the United States plans to expand its operations to Europe and Asia in the next six months. Its product currently runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. All data is stored in an Amazon Aurora database instance.

When the product is deployed in multiple regions, the company wants a single product catalog across all regions, but for compliance purposes, its customer information and purchases must be kept in each region.

How should the company meet these requirements with the LEAST amount of application changes?

- A. Use Amazon Redshift for the product catalog and Amazon DynamoDB tables for the customer information and purchases.
- B. Use Amazon DynamoDB global tables for the product catalog and regional tables for the customer information and purchases.
- C. Use Aurora with read replicas for the product catalog and additional local Aurora instances in each region for the customer information and purchases.
- D. Use Aurora for the product catalog and Amazon DynamoDB global tables for the customer information and purchases.

**Answer: C**

**NEW QUESTION 137**

.....

## **Thank You for Trying Our Product**

### **We offer two products:**

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

### **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

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The DOP-C02 Practice Test Here](#)**