

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 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 3

A DevOps engineer is implementing governance controls for a company that requires its infrastructure to be housed within the United States. The engineer must restrict which AWS Regions can be used, and ensure an alert is sent as soon as possible if any activity outside the governance policy takes place. The controls should be automatically enabled on any new Region outside the United States (US).

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

- A. Create an AWS Organizations SCP that denies access to all non-global services in non-US Region.
- B. Attach the policy to the root of the organization.
- C. Configure AWS CloudTrail to send logs to Amazon CloudWatch Logs and enable it for all Region.
- D. Use a CloudWatch Logs metric filter to send an alert on any service activity in non-US Regions.
- E. Use an AWS Lambda function that checks for AWS service activity and deploy it to all Region.
- F. Write an Amazon EventBridge rule that runs the Lambda function every hour, sending an alert if activity is found in a non-US Region.
- G. Use an AWS Lambda function to query Amazon Inspector to look for service activity in non-US Regions and send alerts if any activity is found.
- H. Write an SCP using the aws:RequestedRegion condition key limiting access to US Region.
- I. Apply the policy to all users, groups, and roles.

Answer: AB

Explanation:

To implement governance controls that restrict AWS service usage to within the United States and ensure alerts for any activity outside the governance policy, the following actions will meet the requirements:

? A. Create an AWS Organizations SCP that denies access to all non-global services in non-US Regions. Attach the policy to the root of the organization. This action will effectively prevent users and roles in all accounts within the organization from accessing services in non-US Regions.

? B. Configure AWS CloudTrail to send logs to Amazon CloudWatch Logs and enable it for all Regions. Use a CloudWatch Logs metric filter to send an alert on any service activity in non-US Regions. This action will allow monitoring of all AWS Regions and will trigger alerts if any activity is detected in non-US Regions, ensuring that the governance team is notified as soon as possible³.

References:

? AWS Documentation on Service Control Policies (SCPs) and how they can be used to manage permissions and restrict access based on Regions¹².

? AWS Documentation on monitoring CloudTrail log files with Amazon CloudWatch Logs to set up alerts for specific activities³.

NEW QUESTION 4

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

What would cause this?

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

Answer: C

Explanation:

This failure is typically due to incorrectly configured health checks in Elastic Load Balancing for the Classic Load Balancer, Application Load Balancer, or Network Load Balancer used to manage traffic for the deployment group. To resolve the issue, review and correct any errors in the health check configuration for the load balancer. <https://docs.aws.amazon.com/codedeploy/latest/userguide/troubleshooting-deployments.html#troubleshooting-deployments-allowtraffic-no-logs>

NEW QUESTION 5

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 metric
- B. Use the recover action to stop and start the instance
- 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 instance
- 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 failure
- 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 resource
- 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 6

A company has an application that runs on AWS Lambda and sends logs to Amazon CloudWatch Logs. An Amazon Kinesis data stream is subscribed to the log groups in CloudWatch Logs. A single consumer Lambda function processes the logs from the data stream and stores the logs in an Amazon S3 bucket.

The company's DevOps team has noticed high latency during the processing and ingestion of some logs.

Which combination of steps will reduce the latency? (Select THREE.)

- A. Create a data stream consumer with enhanced fan-out
- B. Set the Lambda function that processes the logs as the consumer.
- C. Increase the ParallelizationFactor setting in the Lambda event source mapping.
- D. Configure reserved concurrency for the Lambda function that processes the logs.
- E. Increase the batch size in the Kinesis data stream.
- F. Turn off the ReportBatchItemFailures setting in the Lambda event source mapping.
- G. Increase the number of shards in the Kinesis data stream.

Answer: ABC

Explanation:

The latency in processing and ingesting logs can be caused by several factors, such as the throughput of the Kinesis data stream, the concurrency of the Lambda function, and the configuration of the event source mapping. To reduce the latency, the following steps can be taken:

? Create a data stream consumer with enhanced fan-out. Set the Lambda function that processes the logs as the consumer. This will allow the Lambda function to receive records from the data stream with dedicated throughput of up to 2 MB per second per shard, independent of other consumers¹. This will reduce the contention and delay in accessing the data stream.

? Increase the ParallelizationFactor setting in the Lambda event source mapping. This will allow the Lambda service to invoke more instances of the function concurrently to process the records from the data stream². This will increase the processing capacity and reduce the backlog of records in the data stream.

? Configure reserved concurrency for the Lambda function that processes the logs. This will ensure that the function has enough concurrency available to handle the increased load from the data stream³. This will prevent the function from being throttled by the account-level concurrency limit.

The other options are not effective or may have negative impacts on the latency. Option D is not suitable because increasing the batch size in the Kinesis data stream will increase the amount of data that the Lambda function has to process in each invocation, which may increase the execution time and latency⁴. Option E is not advisable because turning off the ReportBatchItemFailures setting in the Lambda event source mapping will prevent the Lambda service from retrying the failed records, which may result in data loss. Option F is not necessary because increasing the number of shards in the Kinesis data stream will increase the throughput of the data stream, but it will not affect the processing speed of the Lambda function, which is the bottleneck in this scenario.

References:

? 1: Using AWS Lambda with Amazon Kinesis Data Streams - AWS Lambda

? 2: AWS Lambda event source mappings - AWS Lambda

? 3: Managing concurrency for a Lambda function - AWS Lambda

- ? 4: AWS Lambda function scaling - AWS Lambda
- ? : AWS Lambda event source mappings - AWS Lambda
- ? : Scaling Amazon Kinesis Data Streams with AWS CloudFormation - Amazon Kinesis Data Streams

NEW QUESTION 7

A company has a new AWS account that teams will use to deploy various applications. The teams will create many Amazon S3 buckets for application- specific purposes and to store AWS CloudTrail logs. The company has enabled Amazon Macie for the account.

A DevOps engineer needs to optimize the Macie costs for the account without compromising the account's functionality.

Which solutions will meet these requirements? (Select TWO.)

- A. Exclude S3 buckets that contain CloudTrail logs from automated discovery.
- B. Exclude S3 buckets that have public read access from automated discovery.
- C. Configure scheduled daily discovery jobs for all S3 buckets in the account.
- D. Configure discovery jobs to include S3 objects based on the last modified criterion.
- E. Configure discovery jobs to include S3 objects that are tagged as production only.

Answer: AD

Explanation:

To optimize the Macie costs for the account without compromising the account's functionality, the DevOps engineer needs to exclude S3 buckets that do not contain sensitive data from automated discovery. S3 buckets that contain CloudTrail logs are unlikely to have sensitive data, and Macie charges for scanning and monitoring data in S3 buckets. Therefore, excluding S3 buckets that contain CloudTrail logs from automated discovery can reduce Macie costs. Similarly, configuring discovery jobs to include S3 objects based on the last modified criterion can also reduce Macie costs, as it will only scan and monitor new or updated objects, rather than all objects in the bucket.

NEW QUESTION 8

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 projec
- 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 pipelin
- 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 fil
- N. Configure the S3 bucket as the artifact destinatio
- 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 intervention¹. References:

- ? AWS CodeBuild User Guide on test reporting¹.
- ? AWS CodeBuild User Guide on working with report groups².
- ? AWS Documentation on using AWS CodePipeline with AWS CodeBuild³.

NEW QUESTION 9

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 installle
- 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 installle
- D. Create an AppSpec file that contains application deployment scripts and grants access to CodeDeploy.
- E. Create an application in CodeDeplo
- F. Configure an in-place deployment typ
- 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 AM
- I. Configure CodeDeploy to deploy the newly created AMI.
- J. Create an application in CodeDeplo
- K. Configure an in-place deployment typ
- 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 CodeDeplo

- O. Configure an in-place deployment typ
- 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 10

A company manages multiple AWS accounts by using AWS Organizations with OUS for the different business divisions. The company is updating their corporate network to use new IP address ranges. The company has 10 Amazon S3 buckets in different AWS accounts. The S3 buckets store reports for the different divisions. The S3 bucket configurations allow only private corporate network IP addresses to access the S3 buckets.

A DevOps engineer needs to change the range of IP addresses that have permission to access the contents of the S3 buckets. The DevOps engineer also needs to revoke the permissions of two OUS in the company.

Which solution will meet these requirements?

- A. Create a new SCP that has two statements, one that allows access to the new range of IP addresses for all the S3 buckets and one that denies access to the old range of IP addresses for all the S3 bucket
- B. Set a permissions boundary for the OrganizationAccountAccessRole role in the two OUS to deny access to the S3 buckets.
- C. Create a new SCP that has a statement that allows only the new range of IP addresses to access the S3 bucket
- D. Create another SCP that denies access to the S3 bucket
- E. Attach the second SCP to the two OUS
- F. On all the S3 buckets, configure resource-based policies that allow only the new range of IP addresses to access the S3 bucket
- G. Create a new SCP that denies access to the S3 bucket
- H. Attach the SCP to the two OUs.
- I. On all the S3 buckets, configure resource-based policies that allow only the new range of IP addresses to access the S3 bucket
- J. Set a permissions boundary for the OrganizationAccountAccessRole role in the two OUS to deny access to the S3 buckets.

Answer: C

Explanation:

The correct answer is C.

A comprehensive and detailed explanation is:

? Option A is incorrect because creating a new SCP that has two statements, one that allows access to the new range of IP addresses for all the S3 buckets and one that denies access to the old range of IP addresses for all the S3 buckets, is not a valid solution. SCPs are not resource-based policies, and they cannot specify the S3 buckets or the IP addresses as resources or conditions. SCPs can only control the actions that can be performed by the principals in the organization, not the access to specific resources. Moreover, setting a permissions boundary for the OrganizationAccountAccessRole role in the two OUs to deny access to the S3 buckets is not sufficient to revoke the permissions of the two OUs, as there might be other roles or users in those OUs that can still access the S3 buckets.

? Option B is incorrect because creating a new SCP that has a statement that allows only the new range of IP addresses to access the S3 buckets is not a valid solution, for the same reason as option A. SCPs are not resource-based policies, and they cannot specify the S3 buckets or the IP addresses as resources or conditions. Creating another SCP that denies access to the S3 buckets and attaching it to the two OUs is also not a valid solution, as SCPs cannot specify the S3 buckets as resources either.

? Option C is correct because it meets both requirements of changing the range of IP addresses that have permission to access the contents of the S3 buckets and revoking the permissions of two OUs in the company. On all the S3 buckets, configuring resource-based policies that allow only the new range of IP addresses to access the S3 buckets is a valid way to update the IP address ranges, as resource-based policies can specify both resources and conditions. Creating a new SCP that denies access to the S3 buckets and attaching it to the two OUs is also a valid way to revoke the permissions of those OUs, as SCPs can deny actions such as s3:PutObject or s3:GetObject on any resource.

? Option D is incorrect because setting a permissions boundary for the OrganizationAccountAccessRole role in the two OUs to deny access to the S3 buckets is not sufficient to revoke the permissions of the two OUs, as there might be other roles or users in those OUs that can still access the S3 buckets. A permissions boundary is a policy that defines the maximum permissions that an IAM entity can have. However, it does not revoke any existing permissions that are granted by other policies.

References:

- ? AWS Organizations
- ? S3 Bucket Policies
- ? Service Control Policies
- ? Permissions Boundaries

NEW QUESTION 10

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

The company has enabled AWS Config in each existing AWS account in the organization.

A DevOps engineer must implement a solution that enables AWS Config automatically for all future AWS accounts that are created in the organization.

Which solution will meet this requirement?

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

Answer: B

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2020/02/aws-cloudformation-stacksets-introduces-automatic-deployments-across-accounts-and-regions-through-aws-organizations/>

NEW QUESTION 12

A company's developers use Amazon EC2 instances as remote workstations. The company is concerned that users can create or modify EC2 security groups to allow unrestricted inbound access.

A DevOps engineer needs to develop a solution to detect when users create unrestricted security group rules. The solution must detect changes to security group rules in near real time, remove unrestricted rules, and send email notifications to the security team. The DevOps engineer has created an AWS Lambda function that checks for security group ID from input, removes rules that grant unrestricted access, and sends notifications through Amazon Simple Notification Service (Amazon SNS).

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

- A. Configure the Lambda function to be invoked by the SNS topic
- B. Create an AWS CloudTrail subscription for the SNS topic
- C. Configure a subscription filter for security group modification events.
- D. Create an Amazon EventBridge scheduled rule to invoke the Lambda function
- E. Define a schedule pattern that runs the Lambda function every hour.
- F. Create an Amazon EventBridge event rule that has the default event bus as the source
- G. Define the rule's event pattern to match EC2 security group creation and modification event
- H. Configure the rule to invoke the Lambda function.
- I. Create an Amazon EventBridge custom event bus that subscribes to events from all AWS services
- J. Configure the Lambda function to be invoked by the custom event bus.

Answer: C

Explanation:

To meet the requirements, the DevOps engineer should create an Amazon EventBridge event rule that has the default event bus as the source. The rule's event pattern should match EC2 security group creation and modification events, and it should be configured to invoke the Lambda function. This solution will allow for near real-time detection of security group rule changes and will trigger the Lambda function to remove any unrestricted rules and send email notifications to the security team. <https://repost.aws/knowledge-center/monitor-security-group-changes-ec2>

NEW QUESTION 16

A company uses an Amazon API Gateway regional REST API to host its application API. The REST API has a custom domain. The REST API's default endpoint is deactivated.

The company's internal teams consume the API. The company wants to use mutual TLS between the API and the internal teams as an additional layer of authentication.

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

- A. Use AWS Certificate Manager (ACM) to create a private certificate authority (CA). Provision a client certificate that is signed by the private CA.
- B. Provision a client certificate that is signed by a public certificate authority (CA). Import the certificate into AWS Certificate Manager (ACM).
- C. Upload the provisioned client certificate to an Amazon S3 bucket
- D. Configure the API Gateway mutual TLS to use the client certificate that is stored in the S3 bucket as the trust store.
- E. Upload the provisioned client certificate private key to an Amazon S3 bucket
- F. Configure the API Gateway mutual TLS to use the private key that is stored in the S3 bucket as the trust store.
- G. Upload the root private certificate authority (CA) certificate to an Amazon S3 bucket
- H. Configure the API Gateway mutual TLS to use the private CA certificate that is stored in the S3 bucket as the trust store.

Answer: AE

Explanation:

Mutual TLS (mTLS) authentication requires two-way authentication between the client and the server. For Amazon API Gateway, you can enable mTLS for a custom domain name, which requires clients to present X.509 certificates to verify their identity to access your API. To set up mTLS, you would typically use AWS Certificate Manager (ACM) to create a private certificate authority (CA) and provision a client certificate signed by this private CA. The root CA certificate is then uploaded to an Amazon S3 bucket and configured in API Gateway as the trust store¹².

References:

- ? Introducing mutual TLS authentication for Amazon API Gateway¹.
- ? Configuring mutual TLS authentication for a REST API².
- ? AWS Private Certificate Authority details³.
- ? AWS Certificate Manager Private Certificate Authority updates⁴.

NEW QUESTION 17

A company sells products through an ecommerce web application. The company wants a dashboard that shows a pie chart of product transaction details. The company wants to integrate the dashboard with the company's existing Amazon CloudWatch dashboards.

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

- A. Update the ecommerce application to emit a JSON object to a CloudWatch log group for each processed transaction
- B. Use CloudWatch Logs Insights to query the log group and to visualize the results in a pie chart format. Attach the results to the desired CloudWatch dashboard.
- C. Update the ecommerce application to emit a JSON object to an Amazon S3 bucket for each processed transaction
- D. Use Amazon Athena to query the S3 bucket and to visualize the results in a pie chart format
- E. Export the results from Athena. Attach the results to the desired CloudWatch dashboard
- F. Update the ecommerce application to use AWS X-Ray for instrumentation
- G. Create a new X-Ray subsegment. Add an annotation for each processed transaction
- H. Use X-Ray traces to query the data and to visualize the results in a pie chart format. Attach the results to the desired CloudWatch dashboard
- I. Update the ecommerce application to emit a JSON object to a CloudWatch log group for each processed transaction. Create an AWS Lambda function to aggregate and write the results to Amazon DynamoDB
- J. Create a Lambda subscription filter for the log filter
- K. Attach the results to the desired CloudWatch dashboard.

Answer: A

Explanation:

The correct answer is A.

A comprehensive and detailed explanation is:

? Option A is correct because it meets the requirements with the most operational efficiency. Updating the ecommerce application to emit a JSON object to a CloudWatch log group for each processed transaction is a simple and cost-effective way to collect the data needed for the dashboard. Using CloudWatch Logs Insights to query the log group and to visualize the results in a pie chart format is also a convenient and integrated solution that leverages the existing CloudWatch dashboards. Attaching the results to the desired CloudWatch dashboard is straightforward and does not require any additional steps or services.

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

Updating the ecommerce application to emit a JSON object to an Amazon S3 bucket for each processed transaction is a valid way to store the data, but it requires creating and managing an S3 bucket and its permissions. Using Amazon Athena to query the S3 bucket and to visualize the results in a pie chart format is also a valid way to analyze the data, but it incurs charges based on the amount of data scanned by each query. Exporting the results from Athena and attaching them to the desired CloudWatch dashboard is also an extra step that adds more overhead and latency.

? Option C is incorrect because it uses AWS X-Ray for an inappropriate purpose.

Updating the ecommerce application to use AWS X-Ray for instrumentation is a good practice for monitoring and tracing distributed applications, but it is not designed for aggregating product transaction details. Creating a new X-Ray subsegment and adding an annotation for each processed transaction is possible, but it would clutter the X-Ray service map and make it harder to debug performance issues. Using X-Ray traces to query the data and to visualize the results in a pie chart format is also possible, but it would require custom code and logic that are not supported by X-Ray natively. Attaching the results to the desired CloudWatch dashboard is also not supported by X-Ray directly, and would require additional steps or services.

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

Updating the ecommerce application to emit a JSON object to a CloudWatch log group for each processed transaction is a simple and cost-effective way to collect the data needed for the dashboard, as in option A. However, creating an AWS Lambda function to aggregate and write the results to Amazon DynamoDB is redundant, as CloudWatch Logs Insights can already perform aggregation queries on log data. Creating a Lambda subscription filter for the log file is also redundant, as CloudWatch Logs Insights can already access log data directly. Attaching the results to the desired CloudWatch dashboard would also require additional steps or services, as DynamoDB does not support native integration with CloudWatch dashboards.

References:

? CloudWatch Logs Insights

? Amazon Athena

? AWS X-Ray

? AWS Lambda

? Amazon DynamoDB

NEW QUESTION 22

A company has multiple accounts in an organization in AWS Organizations. The company's SecOps team needs to receive an Amazon Simple Notification Service (Amazon SNS) notification if any account in the organization turns off the Block Public Access feature on an Amazon S3 bucket. A DevOps engineer must implement this change without affecting the operation of any AWS accounts. The implementation must ensure that individual member accounts in the organization cannot turn off the notification.

Which solution will meet these requirements?

- A. Designate an account to be the delegated Amazon GuardDuty administrator account
- B. Turn on GuardDuty for all accounts across the organization
- C. In the GuardDuty administrator account, create an SNS topic
- D. Subscribe the SecOps team's email address to the SNS topic
- E. In the same account, create an Amazon EventBridge rule that uses an event pattern for GuardDuty findings and a target of the SNS topic.
- F. Create an AWS CloudFormation template that creates an SNS topic and subscribes the SecOps team's email address to the SNS topic
- G. In the template, include an Amazon EventBridge rule that uses an event pattern of CloudTrail activity for s3:PutBucketPublicAccessBlock and a target of the SNS topic
- H. Deploy the stack to every account in the organization by using CloudFormation StackSets.
- I. Turn on AWS Config across the organization
- J. In the delegated administrator account, create an SNS topic
- K. Subscribe the SecOps team's email address to the SNS topic
- L. Deploy a conformance pack that uses the s3-bucket-level-public-access-prohibited AWS Config managed rule in each account and uses an AWS Systems Manager document to publish an event to the SNS topic to notify the SecOps team.
- M. Turn on Amazon Inspector across the organization
- N. In the Amazon Inspector delegated administrator account, create an SNS topic
- O. Subscribe the SecOps team's email address to the SNS topic
- P. In the same account, create an Amazon EventBridge rule that uses an event pattern for public network exposure of the S3 bucket and publishes an event to the SNS topic to notify the SecOps team.

Answer: C

Explanation:

Amazon GuardDuty is primarily on threat detection and response, not configuration monitoring. A conformance pack is a collection of AWS Config rules and remediation actions that can be easily deployed as a single entity in an account and a Region or across an organization in AWS Organizations.

<https://docs.aws.amazon.com/config/latest/developerguide/conformance-packs.html> <https://docs.aws.amazon.com/config/latest/developerguide/s3-account-level-public-access-blocks.html>

NEW QUESTION 24

A DevOps engineer is creating an AWS CloudFormation template to deploy a web service. The web service will run on Amazon EC2 instances in a private subnet behind an Application Load Balancer (ALB). The DevOps engineer must ensure that the service can accept requests from clients that have IPv6 addresses.

What should the DevOps engineer do with the CloudFormation template so that IPv6 clients can access the web service?

- A. Add an IPv6 CIDR block to the VPC and the private subnet for the EC2 instance
- B. Create route table entries for the IPv6 network, use EC2 instance types that support IPv6, and assign IPv6 addresses to each EC2 instance.
- C. Assign each EC2 instance an IPv6 Elastic IP address
- D. Create a target group, and add the EC2 instances as target
- E. Create a listener on port 443 of the ALB, and associate the target group with the ALB.
- F. Replace the ALB with a Network Load Balancer (NLB). Add an IPv6 CIDR block to the VPC and subnets for the NLB, and assign the NLB an IPv6 Elastic IP address.
- G. Add an IPv6 CIDR block to the VPC and subnets for the AL
- H. Create a listener on port 443. and specify the dualstack IP address type on the AL
- I. Create a target group, and add the EC2 instances as target
- J. Associate the target group with the ALB.

Answer: D

Explanation:

it involves adding an IPv6 CIDR block to the VPC and subnets for the ALB and specifying the dualstack IP address type on the ALB listener. This allows the ALB to listen on both IPv4 and IPv6 addresses, and forward requests to the EC2 instances that are added as targets to the target group associated with the ALB.

NEW QUESTION 28

A company's application is currently deployed to a single AWS Region. Recently, the company opened a new office on a different continent. The users in the new office are experiencing high latency. The company's application runs on Amazon EC2 instances behind an Application Load Balancer (ALB) and uses Amazon DynamoDB as the database layer. The instances run in an EC2 Auto Scaling group across multiple Availability Zones. A DevOps engineer is tasked with minimizing application response times and improving availability for users in both Regions.

Which combination of actions should be taken to address the latency issues? (Choose three.)

- A. Create a new DynamoDB table in the new Region with cross-Region replication enabled.
- B. Create new ALB and Auto Scaling group global resources and configure the new ALB to direct traffic to the new Auto Scaling group.
- C. Create new ALB and Auto Scaling group resources in the new Region and configure the new ALB to direct traffic to the new Auto Scaling group.
- D. Create Amazon Route 53 records, health checks, and latency-based routing policies to route to the ALB.
- E. Create Amazon Route 53 aliases, health checks, and failover routing policies to route to the ALB.
- F. Convert the DynamoDB table to a global table.

Answer: CDF

Explanation:

C. Create new ALB and Auto Scaling group resources in the new Region and configure the new ALB to direct traffic to the new Auto Scaling group. This will allow users in the new Region to access the application with lower latency by reducing the network hops between the user and the application servers.

* D. Create Amazon Route 53 records, health checks, and latency-based routing policies to route to the ALB. This will enable Route 53 to route user traffic to the nearest healthy ALB, based on the latency between the user and the ALBs.

* F. Convert the DynamoDB table to a global table. This will enable reads and writes to the table in both Regions with low latency, improving the overall response time of the application

NEW QUESTION 33

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 m Systems Manager also is 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 instance
- 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 IoT 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

NEW QUESTION 35

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 error
- 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 39

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 42

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

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

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

The company wants ensure that only CloudFormation can use the new role. The development team cannot make any manual changes to the deployed resources. Which combination of steps meet these requirements? (Select THREE.)

- A. Remove the AdministratorAccess polic
- B. Assign the ReadOnlyAccess managed IAM policy to the developer rol
- C. Instruct the developers to use the CloudFormationDeployment role as a CloudFormation service role when the developers deploy new stacks.
- D. Update the trust of CloudFormationDeployment role to allow the developer IAM role to assume the CloudFormationDeployment role.
- E. Configure the IAM to be to get and pass the CloudFormationDeployment role if cloudformation actions for resources,

- F. Update the trust Of the CloudFormationDeployment role to allow the cloudformation.amazonaws.com AWS principal to perform the iam:AssumeRole action
- G. Remove the AdministratorAccess policy
- H. Assign the ReadOnlyAccess managed IAM policy to the developer role Instruct the developers to assume the CloudFormationDeployment role when they deploy new stacks
- I. Add an IAM policy to CloudFormationDeployment to allow cloudformation: * on an Add a policy that allows the iam:PassRole action for ARN of CloudFormationDeployment if iam:PassedToService equals cloudformation.amazonaws.com

Answer: ADF

Explanation:

A comprehensive and detailed explanation is:

? Option A is correct because removing the AdministratorAccess policy and assigning the ReadOnlyAccess managed IAM policy to the developer role is a valid way to prevent the developers from making any manual changes to the deployed resources. The AdministratorAccess policy grants full access to all AWS resources and actions, which is not necessary for the developers. The ReadOnlyAccess policy grants read-only access to most AWS resources and actions, which is sufficient for the developers to view the status of their stacks. Instructing the developers to use the CloudFormationDeployment role as a CloudFormation service role when they deploy new stacks is also a valid way to ensure that only CloudFormation can use the new role. A CloudFormation service role is an IAM role that allows CloudFormation to make calls to resources in a stack on behalf of the user1. The user can specify a service role when they create or update a stack, and CloudFormation will use that role's credentials for all operations that are performed on that stack1.

? Option B is incorrect because updating the trust of CloudFormationDeployment role to allow the developer IAM role to assume the CloudFormationDeployment role is not a valid solution. This would allow the developers to manually assume the CloudFormationDeployment role and perform actions on the deployed resources, which is not what the company wants. The trust of CloudFormationDeployment role should only allow the cloudformation.amazonaws.com AWS principal to assume the role, as in option D.

? Option C is incorrect because configuring the IAM user to be able to get and pass the CloudFormationDeployment role if cloudformation actions for resources is not a valid solution. This would allow the developers to manually pass the CloudFormationDeployment role to other services or resources, which is not what the company wants. The IAM user should only be able to pass the CloudFormationDeployment role as a service role when they create or update a stack with CloudFormation, as in option A.

? Option D is correct because updating the trust of CloudFormationDeployment role to allow the cloudformation.amazonaws.com AWS principal to perform the iam:AssumeRole action is a valid solution. This allows CloudFormation to assume the CloudFormationDeployment role and access resources in other services on behalf of the user2. The trust policy of an IAM role defines which entities can assume the role2. By specifying cloudformation.amazonaws.com as the principal, you grant permission only to CloudFormation to assume this role.

? Option E is incorrect because instructing the developers to assume the CloudFormationDeployment role when they deploy new stacks is not a valid solution. This would allow the developers to manually assume the CloudFormationDeployment role and perform actions on the deployed resources, which is not what the company wants. The developers should only use the CloudFormationDeployment role as a service role when they deploy new stacks with CloudFormation, as in option A.

? Option F is correct because adding an IAM policy to CloudFormationDeployment that allows cloudformation: * on all resources and adding a policy that allows the iam:PassRole action for ARN of CloudFormationDeployment if iam:PassedToService equals cloudformation.amazonaws.com are valid solutions. The first policy grants permission for CloudFormationDeployment to perform any action with any resource using cloudformation.amazonaws.com as a service principal3. The second policy grants permission for passing this role only if it is passed by cloudformation.amazonaws.com as a service principal4. This ensures that only CloudFormation can use this role.

References:

? 1: AWS CloudFormation service roles

? 2: How to use trust policies with IAM roles

? 3: AWS::IAM::Policy

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

NEW QUESTION 46

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 47

A company's security policies require the use of security hardened AMIs in production environments. A DevOps engineer has used EC2 Image Builder to create a pipeline that builds the AMIs on a recurring schedule.

The DevOps engineer needs to update the launch templates of the company's Auto Scaling groups. The Auto Scaling groups must use the newest AMIs during the launch of Amazon EC2 instances.

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

- A. Configure an Amazon EventBridge rule to receive new AMI events from Image Builder
- B. Target an AWS Systems Manager Run Command document that updates the launch templates of the Auto Scaling groups with the newest AMI ID.
- C. Configure an Amazon EventBridge rule to receive new AMI events from Image Builder

- D. Target an AWS Lambda function that updates the launch templates of the Auto Scaling groups with the newest AMI ID.
- E. Configure the launch template to use a value from AWS Systems Manager Parameter Store for the AMI ID.
- F. Configure the Image Builder pipeline to update the Parameter Store value with the newest AMI ID.
- G. Configure the Image Builder distribution settings to update the launch templates with the newest AMI ID.
- H. Configure the Auto Scaling groups to use the newest version of the launch template.

Answer: C

Explanation:

? The most operationally efficient solution is to use AWS Systems Manager Parameter Store¹ to store the AMI ID and reference it in the launch template². This way, the launch template does not need to be updated every time a new AMI is created by Image Builder. Instead, the Image Builder pipeline can update the Parameter Store value with the newest AMI ID³, and the Auto Scaling group can launch instances using the latest value from Parameter Store.

? The other solutions require updating the launch template or creating a new version of it every time a new AMI is created, which adds complexity and overhead. Additionally, using EventBridge rules and Lambda functions or Run Command documents introduces additional dependencies and potential points of failure.

References: 1: AWS Systems Manager Parameter Store 2: Using AWS Systems Manager parameters instead of AMI IDs in launch templates 3: Update an SSM parameter with Image Builder

NEW QUESTION 49

To run an application, a DevOps engineer launches an Amazon EC2 instance with public IP addresses in a public subnet. A user data script obtains the application artifacts and installs them on the instances upon launch. A change to the security classification of the application now requires the instances to run with no access to the internet. While the instances launch successfully and show as healthy, the application does not seem to be installed. Which of the following should successfully install the application while complying with the new rule?

- A. Launch the instances in a public subnet with Elastic IP addresses attached
- B. Once the application is installed and running, run a script to disassociate the Elastic IP addresses afterwards.
- C. Set up a NAT gateway
- D. Deploy the EC2 instances to a private subnet
- E. Update the private subnet's route table to use the NAT gateway as the default route.
- F. Publish the application artifacts to an Amazon S3 bucket and create a VPC endpoint for S3. Assign an IAM instance profile to the EC2 instances so they can read the application artifacts from the S3 bucket.
- G. Create a security group for the application instances and allow only outbound traffic to the artifact repository
- H. Remove the security group rule once the install is complete.

Answer: C

Explanation:

EC2 instances running in private subnets of a VPC can now have controlled access to S3 buckets, objects, and API functions that are in the same region as the VPC. You can use an S3 bucket policy to indicate which VPCs and which VPC Endpoints have access to your S3 buckets 1-
<https://aws.amazon.com/pt/blogs/aws/new-vpc-endpoint-for-amazon-s3/>

NEW QUESTION 54

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

During a code review the DevOps engineer notices that the script does not verify whether the data was successfully copied to Amazon S3. The DevOps engineer must update the script to ensure that data is not corrupted during transmission. The script must use MD5 checksums to verify data integrity before the on-premises data is deleted.

Which solutions for the script will meet these requirements? (Select TWO.)

- A. Check the returned response for the Versioned Compare the returned Versioned against the MD5 checksum.
- B. Include the MD5 checksum within the Content-MD5 parameter
- C. Check the operation's return status to find out if an error was returned.
- D. Include the checksum digest within the tagging parameter as a URL query parameter.
- E. Check the returned response for the ETag
- F. Compare the returned ETag against the MD5 checksum.
- G. Include the checksum digest within the Metadata parameter as a name-value pair. After upload use the S3 HeadObject operation to retrieve metadata from the object.

Answer: BD

Explanation:

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

NEW QUESTION 57

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 59

A company is performing vulnerability scanning for all Amazon EC2 instances across many accounts. The accounts are in an organization in AWS Organizations. Each account's VPCs are attached to a shared transit gateway. The VPCs send traffic to the internet through a central egress VPC. The company has enabled Amazon Inspector in a delegated administrator account and has enabled scanning for all member accounts.

A DevOps engineer discovers that some EC2 instances are listed in the "not scanning" tab in Amazon Inspector.

Which combination of actions should the DevOps engineer take to resolve this issue? (Choose three.)

- A. Verify that AWS Systems Manager Agent is installed and is running on the EC2 instances that Amazon Inspector is not scanning.
- B. Associate the target EC2 instances with security groups that allow outbound communication on port 443 to the AWS Systems Manager service endpoint.
- C. Grant inspector: StartAssessmentRun permissions to the IAM role that the DevOps engineer is using.
- D. Configure EC2 Instance Connect for the EC2 instances that Amazon Inspector is not scanning.
- E. Associate the target EC2 instances with instance profiles that grant permissions to communicate with AWS Systems Manager.
- F. Create a managed-instance activation
- G. Use the Activation Code and the Activation ID to register the EC2 instances.

Answer: ABE

Explanation:

<https://docs.aws.amazon.com/inspector/latest/user/scanning-ec2.html>

NEW QUESTION 63

A rapidly growing company wants to scale for developer demand for AWS development environments. Development environments are created manually in the AWS Management Console. The networking team uses AWS CloudFormation to manage the networking infrastructure, exporting stack output values for the Amazon VPC and all subnets. The development environments have common standards, such as Application Load Balancers, Amazon EC2 Auto Scaling groups, security groups, and Amazon DynamoDB tables.

To keep up with demand, the DevOps engineer wants to automate the creation of development environments. Because the infrastructure required to support the application is expected to grow, there must be a way to easily update the deployed infrastructure. CloudFormation will be used to create a template for the development environments.

Which approach will meet these requirements and quickly provide consistent AWS environments for developers?

- A. Use Fn::ImportValue intrinsic functions in the Resources section of the template to retrieve Virtual Private Cloud (VPC) and subnet value
- B. Use CloudFormation StackSets for the development environments, using the Count input parameter to indicate the number of environments needed
- C. Use the UpdateStackSet command to update existing development environments.
- D. Use nested stacks to define common infrastructure component
- E. To access the exported values, use TemplateURL to reference the networking team's template
- F. To retrieve Virtual Private Cloud (VPC) and subnet values, use Fn::ImportValue intrinsic functions in the Parameters section of the root template
- G. Use the CreateChangeSet and ExecuteChangeSet commands to update existing development environments.
- H. Use nested stacks to define common infrastructure component
- I. Use Fn::ImportValue intrinsic functions with the resources of the nested stack to retrieve Virtual Private Cloud (VPC) and subnet value
- J. Use the CreateChangeSet and ExecuteChangeSet commands to update existing development environments.
- K. Use Fn::ImportValue intrinsic functions in the Parameters section of the root template to retrieve Virtual Private Cloud (VPC) and subnet value
- L. Define the development resources in the order they need to be created in the CloudFormation nested stack
- M. Use the CreateChangeSet
- N. and ExecuteChangeSet commands to update existing development environments.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference-importvalue.html>

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference-importvalue.html> CF of network exports the VPC, subnet or needed information CF of application imports the above information to its stack and UpdateChangeSet/ ExecuteChangeSet

NEW QUESTION 65

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 documentation¹². 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 70

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 yaml 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 75

A company manages multiple AWS accounts in AWS Organizations. The company's security policy states that AWS account root user credentials for member accounts must not be used. The company monitors access to the root user credentials.

A recent alert shows that the root user in a member account launched an Amazon EC2 instance. A DevOps engineer must create an SCP at the organization's root level that will prevent the root user in member accounts from making any AWS service API calls.

Which SCP will meet these requirements?

A)

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "*",
      "Resource": "*",
      "Condition": {
        "StringNotLike": { "aws:PrincipalArn": "arn:aws:iam::*:root" }
      }
    }
  ]
}
```

B)

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": "*",
      "Resource": "*",
      "Principal": { "AWS": "arn:aws:iam::*:root" }
    }
  ]
}
```

C)

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": "*",
      "Resource": "*",
      "Condition": {
        "StringLike": { "aws:PrincipalArn": "arn:aws:iam::*:root" }
      }
    }
  ]
}
```

D)


```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "*",
      "Resource": "*",
      "Principal": "root"
    }
  ]
}
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: D

NEW QUESTION 78

A DevOps engineer has implemented a CI/CO pipeline to deploy an AWS Cloud Formation 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 Cloud Formation 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 metri
- 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 83

A company wants to use a grid system for a proprietary enterprise m-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 cluste
- 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 laye
- C. Assign that recipe to the Configure lifecycle event.
- D. Put the file nodes config in version contro
- E. Create an AWS CodeDeploy deployment configuration and deployment group based on an Amazon EC2 tag value for thecluster node
- F. When adding a new node to the cluster update the file with all tagged instances and make a commit in version contro
- 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 frequentl
- I. Use a process manager such as Monit or system, to restart the cluster services when it detects that the new file was modifie
- 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 confi
- L. Tile 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 87

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 90

A company builds a container image in an AWS CodeBuild project by running Docker commands. After the container image is built, the CodeBuild project uploads the container image to an Amazon S3 bucket. The CodeBuild project has an IAM service role that has permissions to access the S3 bucket.

A DevOps engineer needs to replace the S3 bucket with an Amazon Elastic Container Registry (Amazon ECR) repository to store the container images. The DevOps engineer creates an ECR private image repository in the same AWS Region of the CodeBuild project. The DevOps engineer adjusts the IAM service role with the permissions that are necessary to work with the new ECR repository. The DevOps engineer also places new repository information into the docker build command and the docker push command that are used in the buildspec.yml file.

When the CodeBuild project runs a build job, the job fails when the job tries to access the ECR repository.

Which solution will resolve the issue of failed access to the ECR repository?

- A. Update the buildspec.yml file to log in to the ECR repository by using the aws ecr get-login-password AWS CLI command to obtain an authentication token
- B. Update the docker login command to use the authentication token to access the ECR repository.
- C. Add an environment variable of type SECRETS_MANAGER to the CodeBuild project
- D. In the environment variable, include the ARN of the CodeBuild project's IAM service role
- E. Update the buildspec.yml file to use the new environment variable to log in with the docker login command to access the ECR repository.
- F. Update the ECR repository to be a public image repository
- G. Add an ECR repository policy that allows the IAM service role to have access.
- H. Update the buildspec.yml file to use the AWS CLI to assume the IAM service role for ECR operation
- I. Add an ECR repository policy that allows the IAM service role to have access.

Answer: A

Explanation:

Update the buildspec.yml file to log in to the ECR repository by using the aws ecr get-login-password AWS CLI command to obtain an authentication token.

Update the docker login command to use the authentication token to access the ECR repository.

This is the correct solution. The aws ecr get-login-password AWS CLI command retrieves and displays an authentication token that can be used to log in to an ECR repository. The docker login command can use this token as a password to authenticate with the ECR repository. This way, the CodeBuild project can push and pull images from the ECR repository without any errors. For more information, see Using Amazon ECR with the AWS CLI and get-login-password.

NEW QUESTION 91

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 93

A company must encrypt all AMIs that the company shares across accounts. A DevOps engineer has access to a source account where an unencrypted custom AMI has been built. The DevOps engineer also has access to a target account where an Amazon EC2 Auto Scaling group will launch EC2 instances from the AMI. The DevOps engineer must share the AMI with the target account.

The company has created an AWS Key Management Service (AWS KMS) key in the source account.

Which additional steps should the DevOps engineer perform to meet the requirements? (Choose three.)

- A. In the source account, copy the unencrypted AMI to an encrypted AM
- B. Specify the KMS key in the copy action.
- C. In the source account, copy the unencrypted AMI to an encrypted AM
- D. Specify the default Amazon Elastic Block Store (Amazon EBS) encryption key in the copy action.
- E. In the source account, create a KMS grant that delegates permissions to the Auto Scaling group service-linked role in the target account.
- F. In the source account, modify the key policy to give the target account permissions to create a gran
- G. In the target account, create a KMS grant that delegates permissions to the Auto Scaling group service-linked role.
- H. In the source account, share the unencrypted AMI with the target account.
- I. In the source account, share the encrypted AMI with the target account.

Answer: ADF

Explanation:

The Auto Scaling group service-linked role must have a specific grant in the source account in order to decrypt the encrypted AMI. This is because the service-linked role does not have permissions to assume the default IAM role in the source account. The following steps are required to meet the requirements:

- ? In the source account, copy the unencrypted AMI to an encrypted AMI. Specify the KMS key in the copy action.
- ? In the source account, create a KMS grant that delegates permissions to the Auto Scaling group service-linked role in the target account.
- ? In the source account, share the encrypted AMI with the target account.
- ? In the target account, attach the KMS grant to the Auto Scaling group service- linked role.

The first three steps are the same as the steps that I described earlier. The fourth step is required to grant the Auto Scaling group service-linked role permissions to decrypt the AMI in the target account.

NEW QUESTION 94

A company uses AWS Secrets Manager to store a set of sensitive API keys that an AWS Lambda function uses. When the Lambda function is invoked, the Lambda function retrieves the API keys and makes an API call to an external service. The Secrets Manager secret is encrypted with the default AWS Key Management Service (AWS KMS) key.

A DevOps engineer needs to update the infrastructure to ensure that only the Lambda function's execution role can access the values in Secrets Manager. The solution must apply the principle of least privilege.

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

- A. Update the default KMS key for Secrets Manager to allow only the Lambda function's execution role to decrypt.
- B. Create a KMS customer managed key that trusts Secrets Manager and allows the Lambda function's execution role to decrypt
- C. Update Secrets Manager to use the new customer managed key.
- D. Create a KMS customer managed key that trusts Secrets Manager and allows the account's :root principal to decrypt
- E. Update Secrets Manager to use the new customer managed key.
- F. Ensure that the Lambda function's execution role has the KMS permissions scoped on the resource level
- G. Configure the permissions so that the KMS key can encrypt the Secrets Manager secret.
- H. Remove all KMS permissions from the Lambda function's execution role.

Answer: BD

Explanation:

The requirement is to update the infrastructure to ensure that only the Lambda function's execution role can access the values in Secrets Manager. The solution must apply the principle of least privilege, which means granting the minimum permissions necessary to perform a task.

To do this, the DevOps engineer needs to use the following steps:

- ? Create a KMS customer managed key that trusts Secrets Manager and allows the Lambda function's execution role to decrypt. A customer managed key is a symmetric encryption key that is fully managed by the customer. The customer can define the key policy, which specifies who can use and manage the key. By creating a customer managed key, the DevOps engineer can restrict the decryption permission to only the Lambda function's execution role, and prevent other principals from accessing the secret values. The customer managed key also needs to trust Secrets Manager, which means allowing Secrets Manager to use the key to encrypt and decrypt secrets on behalf of the customer.
- ? Update Secrets Manager to use the new customer managed key. Secrets Manager allows customers to choose which KMS key to use for encrypting each secret. By default, Secrets Manager uses the default KMS key for Secrets Manager, which is a service-managed key that is shared by all customers in the same AWS Region. By updating Secrets Manager to use the new customer managed key, the DevOps engineer can ensure that only the Lambda function's execution role can decrypt the secret values using that key.
- ? Ensure that the Lambda function's execution role has the KMS permissions scoped on the resource level. The Lambda function's execution role is an IAM role that grants permissions to the Lambda function to access AWS services and resources. The role needs to have KMS permissions to use the customer managed key for decryption. However, to apply the principle of least privilege, the role should have the permissions scoped on the resource level, which means specifying the ARN of the customer managed key as a condition in the IAM policy statement. This way, the role can only use that specific key and not any other KMS keys in the account.

NEW QUESTION 97

A developer is maintaining a fleet of 50 Amazon EC2 Linux servers. The servers are part of an Amazon EC2 Auto Scaling group, and also use Elastic Load Balancing for load balancing.

Occasionally, some application servers are being terminated after failing ELB HTTP health checks. The developer would like to perform a root cause analysis on the issue, but before being able to access application logs, the server is terminated.

How can log collection be automated?

- A. Use Auto Scaling lifecycle hooks to put instances in a Pending:Wait stat
- B. Create an Amazon CloudWatch alarm for EC2 Instance Terminate Successful and trigger an AWS Lambda function that invokes an SSM Run Command script to collect logs, push them to Amazon S3, and complete the lifecycle action once logs are collected.
- C. Use Auto Scaling lifecycle hooks to put instances in a Terminating:Wait stat
- D. Create an AWS Config rule for EC2 Instance-terminate Lifecycle Action and trigger a step function that invokes a script to collect logs, push them to Amazon S3, and complete the lifecycle action once logs are collected.
- E. Use Auto Scaling lifecycle hooks to put instances in a Terminating:Wait stat
- F. Create an Amazon CloudWatch subscription filter for EC2 Instance Terminate Successful and trigger a CloudWatch agent that invokes a script to collect logs, push them to Amazon S3, and complete the lifecycle action once logs are collected.
- G. Use Auto Scaling lifecycle hooks to put instances in a Terminating:Wait stat
- H. Create an Amazon EventBridge rule for EC2 Instance-terminate Lifecycle Action and trigger an AWS Lambda function that invokes an SSM Run Command script to collect logs, push them to Amazon S3, and complete the lifecycle action once logs are collected.

Answer: D

Explanation:

<https://blog.fourninecloud.com/auto-scaling-lifecycle-hooks-to-export-server-logs-when-instance-terminating-58e06d7c0d6a>

NEW QUESTION 101

A DevOps engineer manages a web application that runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances run in an EC2 Auto Scaling group across multiple Availability Zones. The engineer needs to implement a deployment strategy that:

Launches a second fleet of instances with the same capacity as the original fleet. Maintains the original fleet unchanged while the second fleet is launched. Transitions traffic to the second fleet when the second fleet is fully deployed. Terminates the original fleet automatically 1 hour after transition.

Which solution will satisfy these requirements?

- A. Use an AWS CloudFormation template with a retention policy for the ALB set to 1 hour
- B. Update the Amazon Route 53 record to reflect the new ALB.
- C. Use two AWS Elastic Beanstalk environments to perform a blue/green deployment from the original environment to the new one
- D. Create an application version lifecycle policy to terminate the original environment in 1 hour.
- E. Use AWS CodeDeploy with a deployment group configured with a blue/green deployment configuration. Select the option Terminate the original instances in the deployment group with a waiting period of 1 hour.
- F. Use AWS Elastic Beanstalk with the configuration set to Immutable
- G. Create an extension using the Resources key that sets the deletion policy of the ALB to 1 hour, and deploy the application.

Answer: C

Explanation:

https://docs.aws.amazon.com/codedeploy/latest/APIReference/API_BlueInstanceTerminationOption.html

The original revision termination settings are configured to wait 1 hour after traffic has been rerouted before terminating the blue task set.

<https://docs.aws.amazon.com/AmazonECS/latest/developerguide/deployment-type-bluegreen.html>

NEW QUESTION 103

A company uses AWS Organizations to manage multiple accounts. Information security policies require that all unencrypted Amazon EBS volumes be marked as non-compliant. A DevOps engineer needs to automatically deploy the solution and ensure that this compliance check is always present.

Which solution will accomplish this?

- A. Create an AWS CloudFormation template that defines an AWS Inspector rule to check whether EBS encryption is enabled
- B. Save the template to an Amazon S3 bucket that has been shared with all accounts within the company
- C. Update the account creation script pointing to the CloudFormation template in Amazon S3.
- D. Create an AWS Config organizational rule to check whether EBS encryption is enabled and deploy the rule using the AWS CLI
- E. Create and apply an SCP to prohibit stopping and deleting AWS Config across the organization.
- F. Create an SCP in Organization
- G. Set the policy to prevent the launch of Amazon EC2 instances without encryption on the EBS volumes using a conditional expression
- H. Apply the SCP to all AWS accounts
- I. Use Amazon Athena to analyze the AWS CloudTrail output, looking for events that deny an ec2: RunInstances action.
- J. Deploy an IAM role to all accounts from a single trusted account
- K. Build a pipeline with AWS CodePipeline with a stage in AWS Lambda to assume the IAM role, and list all EBS volumes in the account
- L. Publish a report to Amazon S3.

Answer: B

Explanation:

<https://docs.aws.amazon.com/config/latest/developerguide/ec2-ebs-encryption-by-default.html>

NEW QUESTION 107

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 109

A company is using AWS to run digital workloads. Each application team in the company has its own AWS account for application hosting. The accounts are consolidated in an organization in AWS Organizations.

The company wants to enforce security standards across the entire organization. To avoid noncompliance because of security misconfiguration, the company has enforced the use of AWS CloudFormation. A production support team can modify resources in the production environment by using the AWS Management

Console to troubleshoot and resolve application-related issues.

A DevOps engineer must implement a solution to identify in near real time any AWS

service misconfiguration that results in noncompliance. The solution must automatically remediate the issue within 15 minutes of identification. The solution also must track noncompliant resources and events in a centralized dashboard with accurate timestamps.

Which solution will meet these requirements with the LEAST development overhead?

- A. Use CloudFormation drift detection to identify noncompliant resource
- B. Use drift detection events from CloudFormation to invoke an AWS Lambda function for remediation
- C. Configure the Lambda function to publish logs to an Amazon CloudWatch Logs log group
- D. Configure an Amazon CloudWatch dashboard to use the log group for tracking.
- E. Turn on AWS CloudTrail in the AWS account
- F. Analyze CloudTrail logs by using Amazon Athena to identify noncompliant resource
- G. Use AWS Step Functions to track query results on Athena for drift detection and to invoke an AWS Lambda function for remediation
- H. For tracking, set up an Amazon QuickSight dashboard that uses Athena as the data source.
- I. Turn on the configuration recorder in AWS Config in all the AWS accounts to identify noncompliant resource
- J. Enable AWS Security Hub with the `--no-enable-default-standards` option in all the AWS account
- K. Set up AWS Config managed rules and custom rule
- L. Set up automatic remediation by using AWS Config conformance pack
- M. For tracking, set up a dashboard on Security Hub in a designated Security Hub administrator account.
- N. Turn on AWS CloudTrail in the AWS account
- O. Analyze CloudTrail logs by using Amazon CloudWatch Logs to identify noncompliant resource
- P. Use CloudWatch Logs filters for drift detection
- Q. Use Amazon EventBridge to invoke the Lambda function for remediation
- R. Stream filtered CloudWatch logs to Amazon OpenSearch Service
- S. Set up a dashboard on OpenSearch Service for tracking.

Answer: C

Explanation:

The best solution is to use AWS Config and AWS Security Hub to identify and remediate noncompliant resources across multiple AWS accounts. AWS Config enables continuous monitoring of the configuration of AWS resources and evaluates them against desired configurations. AWS Config can also automatically remediate noncompliant resources by using conformance packs, which are a collection of AWS Config rules and remediation actions that can be deployed as a single entity. AWS Security Hub provides a comprehensive view of the security posture of AWS accounts and resources. AWS Security Hub can aggregate and normalize the findings from AWS Config and other AWS services, as well as from partner solutions. AWS Security Hub can also be used to create a dashboard for tracking noncompliant resources and events in a centralized location.

The other options are not optimal because they either require more development overhead, do not provide near real time detection and remediation, or do not provide a centralized dashboard for tracking.

Option A is not optimal because CloudFormation drift detection is not a near real time solution. Drift detection has to be manually initiated on each stack or resource, or scheduled using a cron expression. Drift detection also does not provide remediation

actions, so a custom Lambda function has to be developed and invoked. CloudWatch Logs and dashboard can be used for tracking, but they do not provide a comprehensive view of the security posture of the AWS accounts and resources.

Option B is not optimal because CloudTrail logs analysis using Athena is not a near real time solution. Athena queries have to be manually run or scheduled using a cron expression. Athena also does not provide remediation actions, so a custom Lambda function has to be developed and invoked. Step Functions can be used to orchestrate the query and remediation workflow, but it adds more complexity and cost. QuickSight dashboard can be used for tracking, but it does not provide a comprehensive view of the security posture of the AWS accounts and resources.

Option D is not optimal because CloudTrail logs analysis using CloudWatch Logs is not a near real time solution. CloudWatch Logs filters have to be manually created or updated for each resource type and configuration change. CloudWatch Logs also does not provide remediation actions, so a custom Lambda function has to be developed and invoked. EventBridge can be used to trigger the Lambda function, but it adds more complexity and cost. OpenSearch Service dashboard can be used for tracking, but it does not provide a comprehensive view of the security posture of the AWS accounts and resources. References:

? AWS Config conformance packs

? Introducing AWS Config conformance packs

? Managing conformance packs across all accounts in your organization

NEW QUESTION 111

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

The buildspec.yml file contains the following:

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

The DevOps engineer has noticed that anybody with an AWS account is able to download the artifacts.

What steps should the DevOps engineer take to stop this?

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

Answer: D

Explanation:

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

NEW QUESTION 114

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?

- A. Modify the Kinesis consumer application to store the logs durably in Amazon S3. Use Amazon EMR to process the data directly on Amazon S3 to derive customer insights. Store the results in Amazon S3.
- B. Horizontally scale the Kinesis consumer application by adding more EC2 instances based on the Amazon CloudWatch GetRecords.IteratorAge.Milliseconds metric. Increase the retention period of the Kinesis data streams.
- C. Convert the Kinesis consumer application to run as an AWS Lambda function.
- D. Configure the Kinesis data streams as the event source for the Lambda function to process the data streams.
- E. Increase the number of shards in the Kinesis data streams to increase the overall throughput so that the consumer application processes the data faster.

Answer: B

Explanation:

<https://docs.aws.amazon.com/streams/latest/dev/monitoring-with-cloudwatch.html>
GetRecords.IteratorAge.Milliseconds - The age of the last record in all GetRecords calls made against a Kinesis stream, measured over the specified time period. Age is the difference between the current time and when the last record of the GetRecords call was written to the stream. The Minimum and Maximum statistics can be used to track the progress of Kinesis consumer applications. A value of zero indicates that the records being read are completely caught up.

NEW QUESTION 119

A company that uses electronic health records is running a fleet of Amazon EC2 instances with an Amazon Linux operating system. As part of patient privacy requirements, the company must ensure continuous compliance for patches for operating system and applications running on the EC2 instances. How can the deployments of the operating system and application patches be automated using a default and custom repository?

- A. Use AWS Systems Manager to create a new patch baseline including the custom repository.
- B. Run the AWS-RunPatchBaseline document using the run command to verify and install patches.
- C. Use AWS Direct Connect to integrate the corporate repository and deploy the patches using Amazon CloudWatch scheduled events, then use the CloudWatch dashboard to create reports.
- D. Use yum-config-manager to add the custom repository under /etc/yum.repos.d and run yum-config-manager-enable to activate the repository.
- E. Use AWS Systems Manager to create a new patch baseline including the corporate repository.
- F. Run the AWS-AmazonLinuxDefaultPatchBaseline document using the run command to verify and install patches.

Answer: A

Explanation:

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

NEW QUESTION 120

A company has developed a serverless web application that is hosted on AWS. The application consists of Amazon S3, Amazon API Gateway, several AWS Lambda functions, and an Amazon RDS for MySQL database. The company is using AWS CodeCommit to store the source code. The source code is a combination of AWS Serverless Application Model (AWS SAM) templates and Python code. A security audit and penetration test reveal that user names and passwords for authentication to the database are hardcoded within CodeCommit repositories. A DevOps engineer must implement a solution to automatically detect and prevent hardcoded secrets. What is the MOST secure solution that meets these requirements?

- A. Enable Amazon CodeGuru Profile
- B. Decorate the handler function with `@with_lambda_profiler()`. Manually review the recommendation report.
- C. Write the secret to AWS Systems Manager Parameter Store as a secure string.
- D. Update the SAM templates and the Python code to pull the secret from Parameter Store.
- E. Associate the CodeCommit repository with Amazon CodeGuru Reviewer.
- F. Manually check the code review for any recommendation.
- G. Choose the option to protect the secret.
- H. Update the SAM templates and the Python code to pull the secret from AWS Secrets Manager.
- I. Enable Amazon CodeGuru Profile
- J. Decorate the handler function with `@with_lambda_profiler()`. Manually review the recommendation report.
- K. Choose the option to protect the secret.
- L. Update the SAM templates and the Python code to pull the secret from AWS Secrets Manager.
- M. Associate the CodeCommit repository with Amazon CodeGuru Reviewer.
- N. Manually check the code review for any recommendation.
- O. Write the secret to AWS Systems Manager Parameter Store as a string.
- P. Update the SAM templates and the Python code to pull the secret from Parameter Store.

Answer: B

Explanation:

<https://docs.aws.amazon.com/codecommit/latest/userguide/how-to-amazon-codeguru-reviewer.html>

NEW QUESTION 125

A company has multiple development teams in different business units that work in a shared single AWS account. All Amazon EC2 resources that are created in the account must include tags that specify who created the resources. The tagging must occur within the first hour of resource creation. A DevOps engineer needs to add tags to the created resources that include the user ID that created the resource and the cost center ID. The DevOps engineer configures an AWS Lambda function with the cost center mappings to tag the resources. The DevOps engineer also sets up AWS CloudTrail in the AWS account. An Amazon S3 bucket stores the CloudTrail event logs. Which solution will meet the tagging requirements?

- A. Create an S3 event notification on the S3 bucket to invoke the Lambda function for s3.ObjectTagging:Put event
- B. Enable bucket versioning on the S3 bucket.
- C. Enable server access logging on the S3 bucket
- D. Create an S3 event notification on the S3 bucket for s3.ObjectTagging.* events
- E. Create a recurring hourly Amazon EventBridge scheduled rule that invokes the Lambda function
- F. Modify the Lambda function to read the logs from the S3 bucket
- G. Create an Amazon EventBridge rule that uses Amazon EC2 as the event source
- H. Configure the rule to match events delivered by CloudTrail
- I. Configure the rule to target the Lambda function

Answer: D

Explanation:

? Option A is incorrect because S3 event notifications do not support s3.ObjectTagging:Put events. S3 event notifications only support events related to object creation, deletion, replication, and restore. Moreover, enabling bucket versioning on the S3 bucket is not relevant to the tagging requirements, as it only keeps multiple versions of objects in the bucket.

? Option B is incorrect because enabling server access logging on the S3 bucket does not help with tagging the resources. Server access logging only records requests for access to the bucket or its objects. It does not capture the user ID or the cost center ID of the resources. Furthermore, creating an S3 event notification on the S3 bucket for s3.ObjectTagging:Put events is not possible, as explained in option A.

? Option C is incorrect because creating a recurring hourly Amazon EventBridge scheduled rule that invokes the Lambda function is not efficient or timely. The Lambda function would have to read the logs from the S3 bucket every hour and tag the resources accordingly, which could incur unnecessary costs and delays. A better solution would be to trigger the Lambda function as soon as a resource is created, rather than waiting for an hourly schedule.

? Option D is correct because creating an Amazon EventBridge rule that uses Amazon EC2 as the event source and matches events delivered by CloudTrail is a valid way to tag the resources. CloudTrail records all API calls made to AWS services, including EC2, and delivers them as events to EventBridge. The EventBridge rule can filter the events based on the user ID and the resource type, and then target the Lambda function to tag the resources with the cost center ID. This solution meets the tagging requirements in a timely and efficient manner.

References:

- ? S3 event notifications
- ? Server access logging
- ? Amazon EventBridge rules
- ? AWS CloudTrail

NEW QUESTION 127

A company has configured an Amazon S3 event source on an AWS Lambda function. The company needs the Lambda function to run when a new object is created or an existing object is modified in a particular S3 bucket. The Lambda function will use the S3 bucket name and the S3 object key of the incoming event to read the contents of the created or modified S3 object. The Lambda function will parse the contents and save the parsed contents to an Amazon DynamoDB table. The Lambda function's execution role has permissions to read from the S3 bucket and to write to the DynamoDB table. During testing, a DevOps engineer discovers that the Lambda

function does not run when objects are added to the S3 bucket or when existing objects are modified.

Which solution will resolve this problem?

- A. Increase the memory of the Lambda function to give the function the ability to process large files from the S3 bucket.
- B. Create a resource policy on the Lambda function to grant Amazon S3 the permission to invoke the Lambda function for the S3 bucket
- C. Configure an Amazon Simple Queue Service (Amazon SQS) queue as an OnFailure destination for the Lambda function
- D. Provision space in the /tmp folder of the Lambda function to give the function the ability to process large files from the S3 bucket

Answer: B

Explanation:

? Option A is incorrect because increasing the memory of the Lambda function does not address the root cause of the problem, which is that the Lambda function is not triggered by the S3 event source. Increasing the memory of the Lambda function might improve its performance or reduce its execution time, but it does not affect its invocation. Moreover, increasing the memory of the Lambda function might incur higher costs, as Lambda charges based on the amount of memory allocated to the function.

? Option B is correct because creating a resource policy on the Lambda function to grant Amazon S3 the permission to invoke the Lambda function for the S3 bucket is a necessary step to configure an S3 event source. A resource policy is a JSON document that defines who can access a Lambda resource and under what conditions. By granting Amazon S3 permission to invoke the Lambda function, the company ensures that the Lambda function runs when a new object is created or an existing object is modified in the S3 bucket.

? Option C is incorrect because configuring an Amazon Simple Queue Service (Amazon SQS) queue as an On-Failure destination for the Lambda function does not help with triggering the Lambda function. An On-Failure destination is a feature that allows Lambda to send events to another service, such as SQS or Amazon Simple Notification Service (Amazon SNS), when a function invocation fails. However, this feature only applies to asynchronous invocations, and S3 event sources use synchronous invocations. Therefore, configuring an SQS queue as an On-Failure destination would have no effect on the problem.

? Option D is incorrect because provisioning space in the /tmp folder of the Lambda function does not address the root cause of the problem, which is that the Lambda function is not triggered by the S3 event source. Provisioning space in the /tmp folder of the Lambda function might help with processing large files from the S3 bucket, as it provides temporary storage for up to 512 MB of data. However, it does not affect the invocation of the Lambda function.

References:

- ? Using AWS Lambda with Amazon S3
- ? Lambda resource access permissions
- ? AWS Lambda destinations
- ? [AWS Lambda file system]

NEW QUESTION 130

A DevOps engineer is planning to deploy a Ruby-based application to production. The application needs to interact with an Amazon RDS for MySQL database and should have automatic scaling and high availability. The stored data in the database is critical and should persist regardless of the state of the application stack. The DevOps engineer needs to set up an automated deployment strategy for the application with automatic rollbacks. The solution also must alert the application team when a deployment fails.

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

- A. Deploy the application on AWS Elastic Beanstalk
- B. Deploy an Amazon RDS for MySQL DB instance as part of the Elastic Beanstalk configuration.
- C. Deploy the application on AWS Elastic Beanstalk
- D. Deploy a separate Amazon RDS for MySQL DB instance outside of Elastic Beanstalk.

- E. Configure a notification email address that alerts the application team in the AWS Elastic Beanstalk configuration.
- F. Configure an Amazon EventBridge rule to monitor AWS Health event
- G. Use an Amazon Simple Notification Service (Amazon SNS) topic as a target to alert the application team.
- H. Use the immutable deployment method to deploy new application versions.
- I. Use the rolling deployment method to deploy new application versions.

Answer: BDE

Explanation:

For deploying a Ruby-based application with requirements for interaction with an Amazon RDS for MySQL database, automatic scaling, high availability, and data persistence, the following steps will meet the requirements:

? B. Deploy the application on AWS Elastic Beanstalk. Deploy a separate Amazon RDS for MySQL DB instance outside of Elastic Beanstalk. This approach ensures that the database persists independently of the Elastic Beanstalk environment, which can be torn down and recreated without affecting the database¹²³.

? E. Use the immutable deployment method to deploy new application versions. Immutable deployments provide a zero-downtime deployment method that ensures that if any part of the deployment process fails, the environment is rolled back to the original state automatically⁴.

? D. Configure an Amazon EventBridge rule to monitor AWS Health events. Use an Amazon Simple Notification Service (Amazon SNS) topic as a target to alert the application team. This setup allows for automated monitoring and alerting of the application team in case of deployment failures or other health events⁵⁶.

References:

? AWS Elastic Beanstalk documentation on deploying Ruby applications¹.

? AWS documentation on application auto-scaling⁷.

? AWS documentation on automated deployment strategies with automatic rollbacks and alerts⁴⁵⁶.

NEW QUESTION 133

A development team uses AWS CodeCommit for version control for applications. The development team uses AWS CodePipeline, AWS CodeBuild, and AWS CodeDeploy for CI/CD infrastructure. In CodeCommit, the development team recently merged pull requests that did not pass long-running tests in the code base. The development team needed to perform rollbacks to branches in the codebase, resulting in lost time and wasted effort.

A DevOps engineer must automate testing of pull requests in CodeCommit to ensure that reviewers more easily see the results of automated tests as part of the pull request review.

What should the DevOps engineer do to meet this requirement?

- A. Create an Amazon EventBridge rule that reacts to the pullRequestStatusChanged event
- B. Create an AWS Lambda function that invokes a CodePipeline pipeline with a CodeBuild action that runs the tests for the applicatio
- C. Program the Lambda function to post the CodeBuild badge as a comment on the pull request so that developers will see the badge in their code review.
- D. Create an Amazon EventBridge rule that reacts to the pullRequestCreated event
- E. Create an AWS Lambda function that invokes a CodePipeline pipeline with a CodeBuild action that runs the tests for the applicatio
- F. Program the Lambda function to post the CodeBuild test results as a comment on the pull request when the test results are complete.
- G. Create an Amazon EventBridge rule that reacts to pullRequestCreated and pullRequestSourceBranchUpdated event
- H. Create an AWS Lambda function that invokes a CodePipeline pipeline with a CodeBuild action that runs the tests for the applicatio
- I. Program the Lambda function to post the CodeBuild badge as a comment on the pull request so that developers will see the badge in their code review.
- J. Create an Amazon EventBridge rule that reacts to the pullRequestStatusChanged event
- K. Create an AWS Lambda function that invokes a CodePipeline pipeline with a CodeBuild action that runs the tests for the applicatio
- L. Program the Lambda function to post the CodeBuild test results as a comment on the pull request when the test results are complete.

Answer: C

Explanation:

<https://aws.amazon.com/es/blogs/devops/complete-ci-cd-with-aws-codecommit-aws-codebuild-aws-codedeploy-and-aws-codepipeline/>

NEW QUESTION 138

A DevOps engineer used an AWS Cloud Formation custom resource to set up AD Connector. The AWS Lambda function ran and created AD Connector, but Cloud Formation is not transitioning from CREATE_IN_PROGRESS to CREATE_COMPLETE.

Which action should the engineer take to resolve this issue?

- A. Ensure the Lambda function code has exited successfully.
- B. Ensure the Lambda function code returns a response to the pre-signed URL.
- C. Ensure the Lambda function IAM role has cloudformation UpdateStack permissions for the stack ARN.
- D. Ensure the Lambda function IAM role has ds ConnectDirectory permissions for the AWS account.

Answer: B

Explanation:

Reference: <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/crpg-ref-responses.html>

NEW QUESTION 140

.....

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 FirstTry
- * DOP-C02 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

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