Amazon

Exam Questions AWS-Certified-DevOps-Engineer-Professional

Amazon AWS Certified DevOps Engineer Professional

NEW QUESTION 1
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-ou
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 2
A company runs applications in AWS accounts that are in an organization in AWS Organizations The applications use Amazon EC2 instances and Amazon S3. The company wants to detect potentially compromised EC2 instances suspicious network activity and unusual API activity in its existing AWS accounts and in any AWS accounts that the company creates in the future When the company detects one to these events the company wants to use an existing Amazon Simple Notification Service (Amazon SNS) topic to send a notification to its operational support team for investigation and remediation.
Which solution will meet these requirements in accordance with AWS best practices?

A. In the organization's management account configure an AWS account as the AmazonGuardDuty administrator accoun
B. In the GuardDuty administrator account add the company's existing AWS accounts to GuardDuty as members In the GuardDuty administrator account create an Amazon EventBridge rule with an event pattern to match GuardDuty events and to forward matching events to the SNS topic.
C. In the organization's management account configure Amazon GuardDuty to add newly created AWS accounts by invitation and to send invitations to the existing AWS accounts Create an AWS Cloud Formation stack set that accepts the GuardDuty invitation and creates an Amazon EventBridge rule Configure the rule with an event pattern to matc
D. GuardDuty events and to forward matching events to the SNS topi
E. Configure the Cloud Formation stack set to deploy into all AWS accounts in the organization.
F. In the organization's management accoun
G. create an AWS CloudTrail organization trail Activate the organization trail in all AWS accounts in the organizatio
H. Create an SCP that enables VPC Flow Logs in each account in the organizatio
I. Configure AWS Security Hub for the organization Create an Amazon EventBridge rule with an even pattern to match Security Hub events and to forward matching events to the SNS topic.
J. In the organization's management account configure an AWS account as the AWS CloudTrail administrator account in the CloudTrail administrator account create a CloudTrail organization trai
K. Add the company's existing AWS accounts to the organization trail Create an SCP that enables VPC Flow Logs in each account in the organizatio
L. Configure AWS Security Hub for the organizatio
M. Create an Amazon EventBridge rule with an event pattern to match Security Hub events and to forward matching events to the SNS topic.

Answer: B

Explanation:
 It allows the company to detect potentially compromised EC2 instances, suspicious network activity, and unusual API activity in its existing AWS accounts and in any AWS accounts that the company creates in the future using Amazon GuardDuty. It also provides a solution for automatically adding future AWS accounts to GuardDuty by configuring GuardDuty to add newly created AWS accounts by invitation and to send invitations to the existing AWS accounts.

NEW QUESTION 3
A development team manually builds an artifact locally and then places it in an Amazon S3 bucket. The application has a local cache that must be cleared when a deployment occurs. The team runs a command to do this downloads the artifact from Amazon S3 and unzips the artifact to complete the deployment.
A DevOps team wants to migrate to a CI/CD process and build in checks to stop and roll back the deployment when a failure occurs. This requires the team to track the progression of the deployment.
Which combination of actions will accomplish this? (Select THREE)

A. Allow developers to check the code into a code repository Using Amazon EventBridge on every pull into the mam branch invoke an AWS Lambda function to build the artifact and store it in Amazon S3.
B. Create a custom script to clear the cache Specify the script in the BeforeInstall lifecycle hook in the AppSpec file.
C. Create user data for each Amazon EC2 instance that contains the clear cache script Once deployed test the application If it is not successful deploy it again.
D. Set up AWS CodePipeline to deploy the application Allow developers to check the code into a code repository as a source tor the pipeline.
E. Use AWS CodeBuild to build the artifact and place it in Amazon S3 Use AWS CodeDeploy to deploy the artifact to Amazon EC2 instances.
F. Use AWS Systems Manager to fetch the artifact from Amazon S3 and deploy it to all the instances.

Answer: BDE

NEW QUESTION 4
A company has an AWS CodePipeline pipeline that is configured with an Amazon S3 bucket in the eu-west-1 Region. The pipeline deploys an AWS Lambda application to the same Region. The pipeline consists of an AWS CodeBuild project build action and an AWS CloudFormation deploy action.
The CodeBuild project uses the aws cloudformation package AWS CLI command to build an artifact that contains the Lambda function code's .zip file and the CloudFormation template. The CloudFormation deploy action references the CloudFormation template from the output artifact of the CodeBuild project's build action.
The company wants to also deploy the Lambda application to the us-east-1 Region by using the pipeline in eu-west-1. A DevOps engineer has already updated the CodeBuild project to use the aws cloudformation package command to produce an additional output artifact for us-east-1.
Which combination of additional steps should the DevOps engineer take to meet these requirements? (Choose two.)

A. Modify the CloudFormation template to include a parameter for the Lambda function code's zip file locatio
B. Create a new CloudFormation deploy action for us-east-1 in the pipelin
C. Configure the new deploy action to pass in the us-east-1 artifact location as a parameter override.
D. Create a new CloudFormation deploy action for us-east-1 in the pipelin
E. Configure the new deploy action to use the CloudFormation template from the us-east-1 output artifact.
F. Create an S3 bucket in us-east-1. Configure the S3 bucket policy to allow CodePipeline to have read and write access.
G. Create an S3 bucket in us-east-1. Configure S3 Cross-Region Replication (CRR) from the S3 bucket in eu-west-1 to the S3 bucket in us-east-1.
H. Modify the pipeline to include the S3 bucket for us-east-1 as an artifact stor
I. Create a new CloudFormation deploy action for us-east-1 in the pipelin
J. Configure the new deploy action to use the CloudFormation template from the us-east-1 output artifact.

Answer: AB

Explanation:
 A. The CloudFormation template should be modified to include a parameter that indicates the location of the .zip file containing the Lambda function's code. This allows the CloudFormation deploy action to use the correct artifact depending on the region. This is critical because Lambda functions need to reference their code artifacts from the same region they are being deployed in. B. You would also need to create a new CloudFormation deploy action for the us-east-1 Region within the pipeline. This action should be configured to use the CloudFormation template from the artifact that was specifically created for us- east-1.

NEW QUESTION 5
A company is using an Amazon Aurora cluster as the data store for its application. The Aurora cluster is configured with a single DB instance. The application performs read and write operations on the database by using the cluster's instance endpoint.
The company has scheduled an update to be applied to the cluster during an upcoming maintenance window. The cluster must remain available with the least possible interruption during the maintenance window.
What should a DevOps engineer do to meet these requirements?

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

Answer: C

Explanation:
 To meet the requirements, the DevOps engineer should do the following:
? Turn on the Multi-AZ option on the Aurora cluster.
? Update the application to use the Aurora cluster endpoint for write operations.
? Update the Aurora cluster's reader endpoint for reads.
Turning on the Multi-AZ option will create a replica of the database in a different Availability Zone. This will ensure that the database remains available even if one of the Availability Zones is unavailable.
Updating the application to use the Aurora cluster endpoint for write operations will ensure that all writes are sent to both the primary and replica databases. This will ensure that the data is always consistent.
Updating the Aurora cluster's reader endpoint for reads will allow the application to read data from the replica database. This will improve the performance of the application during the maintenance window.

NEW QUESTION 6
A company is developing an application that will generate log events. The log events consist of five distinct metrics every one tenth of a second and produce a large amount of data The company needs to configure the application to write the logs to Amazon Time stream The company will configure a daily query against the Timestream table.
Which combination of steps will meet these requirements with the FASTEST query performance? (Select THREE.)

A. Use batch writes to write multiple log events in a Single write operation

B. Write each log event as a single write operation
C. Treat each log as a single-measure record
D. Treat each log as a multi-measure record
E. Configure the memory store retention period to be longer than the magnetic store retention period
F. Configure the memory store retention period to be shorter than the magnetic store retention period

Answer: ADF

Explanation:
A comprehensive and detailed explanation is:
? Option A is correct because using batch writes to write multiple log events in a single write operation is a recommended practice for optimizing the performance and cost of data ingestion in Timestream. Batch writes can reduce the number of network round trips and API calls, and can also take advantage of parallel processing by Timestream. Batch writes can also improve the compression ratio of data in the memory store and the magnetic store, which can reduce the storage costs and improve the query performance1.
? Option B is incorrect because writing each log event as a single write operation is not a recommended practice for optimizing the performance and cost of data ingestion in Timestream. Writing each log event as a single write operation would increase the number of network round trips and API calls, and would also reduce the compression ratio of data in the memory store and the magnetic store. This would increase the storage costs and degrade the query performance1.
? Option C is incorrect because treating each log as a single-measure record is not a recommended practice for optimizing the query performance in Timestream. Treating each log as a single-measure record would result in creating multiple records for each timestamp, which would increase the storage size and the query latency. Moreover, treating each log as a single-measure record would require using joins to query multiple measures for the same timestamp, which would add complexity and overhead to the query processing2.
? Option D is correct because treating each log as a multi-measure record is a recommended practice for optimizing the query performance in Timestream. Treating each log as a multi-measure record would result in creating a single record for each timestamp, which would reduce the storage size and the query latency. Moreover, treating each log as a multi-measure record would allow querying multiple measures for the same timestamp without using joins, which would simplify and speed up the query processing2.
? Option E is incorrect because configuring the memory store retention period to be longer than the magnetic store retention period is not a valid option in Timestream. The memory store retention period must always be shorter than or equal to the magnetic store retention period. This ensures that data is moved from the memory store to the magnetic store before it expires out of the memory store3.
? Option F is correct because configuring the memory store retention period to be shorter than the magnetic store retention period is a valid option in Timestream. The memory store retention period determines how long data is kept in the memory store, which is optimized for fast point-in-time queries. The magnetic store retention period determines how long data is kept in the magnetic store, which is optimized for fast analytical queries. By configuring these retention periods appropriately, you can balance your storage costs and query performance according to your application needs3.
References:
? 1: Batch writes
? 2: Multi-measure records vs. single-measure records
? 3: Storage

NEW QUESTION 7
A development team uses AWS CodeCommit, AWS CodePipeline, and AWS CodeBuild to develop and deploy an application. Changes to the code are submitted by pull requests. The development team reviews and merges the pull requests, and then the pipeline builds
and tests the application.
Over time, the number of pull requests has increased. The pipeline is frequently blocked because of failing tests. To prevent this blockage, the development team wants to run the unit and integration tests on each pull request before it is merged.
Which solution will meet these requirements?

A. Create a CodeBuild project to run the unit and integration test
B. Create a CodeCommit approval rule templat
C. Configure the template to require the successful invocation of the CodeBuild projec
D. Attach the approval rule to the project's CodeCommit repository.
E. Create an Amazon EventBridge rule to match pullRequestCreated events from CodeCommit Create a CodeBuild project to run the unit and integration test
F. Configure the CodeBuild project as a target of the EventBridge rule that includes a custom event payload with the CodeCommit repository and branch information from the event.
G. Create an Amazon EventBridge rule to match pullRequestCreated events from CodeCommi
H. Modify the existing CodePipeline pipeline to not run the deploy steps if the build is started from a pull reques
I. Configure the EventBridge rule to run the pipeline with a custom payload that contains the CodeCommit repository and branch information from the event.
J. Create a CodeBuild project to run the unit and integration test
K. Create a CodeCommit notification rule that matches when a pull request is created or update
L. Configure the notification rule to invoke the CodeBuild project.

Answer: B

Explanation:
CodeCommit generates events in CloudWatch, CloudWatch triggers the CodeBuild https://aws.amazon.com/es/blogs/devops/complete-ci-cd-with-aws-codecommit-aws-codebuild-aws-codedeploy-and-aws-codepipeline/

NEW QUESTION 8
A company uses an organization in AWS Organizations to manage its AWS accounts. The company recently acquired another company that has standalone AWS accounts. The acquiring company's DevOps team needs to consolidate the administration of the AWS accounts for both companies and retain full administrative control of the accounts. The DevOps team also needs to collect and group findings across all the accounts to implement and maintain a security posture.
Which combination of steps should the DevOps team take to meet these requirements? (Select TWO.)

A. Invite the acquired company's AWS accounts to join the organizatio
B. Create an SCP that has full administrative privilege
C. Attach the SCP to the management account.
D. Invite the acquired company's AWS accounts to join the organizatio
E. Create the OrganizationAccountAccessRole 1AM role in the invited account
F. Grant permission to the management account to assume the role.
G. Use AWS Security Hub to collect and group findings across all account
H. Use Security Hub to automatically detect new accounts as the accounts are added to the organization.
I. Use AWS Firewall Manager to collect and group findings across all account
J. Enable all features for the organizatio

K. Designate an account in the organization as the delegated administrator account for Firewall Manager.
L. Use Amazon Inspector to collect and group findings across all account
M. Designate an account in the organization as the delegated administrator account for Amazon Inspector.

Answer: BC

Explanation:
 The correct answer is B and C. Option B is correct because inviting the acquired company's AWS accounts to join the organization and creating the OrganizationAccountAccessRole IAM role in the invited accounts allows the management account to assume the role and gain full administrative access to the member accounts. Option C is correct because using AWS Security Hub to collect and group findings across all accounts enables the DevOps team to monitor and improve the security posture of the organization. Security Hub can automatically detect new accounts as the accounts are added to the organization and enable Security Hub for them. Option A is incorrect because creating an SCP that has full administrative privileges and attaching it to the management account does not grant the management account access to the member accounts. SCPs are used to restrict the permissions of the member accounts, not to grant permissions to the management account. Option D is incorrect because using AWS Firewall Manager to collect and group findings across all accounts is not a valid use case for Firewall Manager. Firewall Manager is used to centrally configure and manage firewall rules across the organization, not to collect and group security findings. Option E is incorrect because using Amazon Inspector to collect and group findings across all accounts is not a valid use case for Amazon Inspector. Amazon Inspector is used to assess the security and compliance of applications running on Amazon EC2 instances, not to collect and group security findings across accounts. References:
? Inviting an AWS account to join your organization
? Enabling and disabling AWS Security Hub
? Service control policies
? AWS Firewall Manager
? Amazon Inspector

NEW QUESTION 9
A company is launching an application that stores raw data in an Amazon S3 bucket. Three applications need to access the data to generate reports. The data must be redacted differently for each application before
the applications can access the data.
Which solution will meet these requirements?

A. Create an S3 bucket for each applicatio
B. Configure S3 Same-Region Replication (SRR) from the raw data's S3 bucket to each application's S3 bucke
C. Configure each application to consume data from its own S3 bucket.
D. Create an Amazon Kinesis data strea
E. Create an AWS Lambda function that isinvoked by object creation events in the raw data's S3 bucke
F. Program the Lambda function to redact data for each applicatio
G. Publish the data on the Kinesis data strea
H. Configure each application to consume data from the Kinesis data stream.
I. For each application, create an S3 access point that uses the raw data's S3 bucket as the destinatio
J. Create an AWS Lambda function that is invoked by object creation events in the raw data's S3 bucke
K. Program the Lambda function to redact data for each applicatio
L. Store the data in each application's S3 access poin
M. Configure each application to consume data from its own S3 access point.
N. Create an S3 access point that uses the raw data's S3 bucket as the destinatio
O. For each application, create an S3 Object Lambda access point that uses the S3 access poin
P. Configure the AWS Lambda function for each S3 Object Lambda access point to redact data when objects are retrieve
Q. Configure each application to consume data from its own S3 Object Lambda access point.

Answer: D

Explanation:
? The best solution is to use S3 Object Lambda1, which allows you to add your own code to S3 GET, LIST, and HEAD requests to modify and process data as it is returned to an application2. This way, you can redact the data differently for each application without creating and storing multiple copies of the data or running proxies.
? The other solutions are less efficient or scalable because they require replicating the data to multiple buckets, streaming the data through Kinesis, or storing the data in S3 access points.
References: 1: Amazon S3 Features | Object Lambda | AWS 2: Transforming objects with S3 Object Lambda - Amazon Simple Storage Service

NEW QUESTION 10
A company is testing a web application that runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. The company uses a blue green deployment process with immutable instances when deploying new software.
During testing users are being automatically logged out of the application at random times. Testers also report that when a new version of the application is deployed all users are logged out. The development team needs a solution to ensure users remain logged m across scaling events and application deployments.
What is the MOST operationally efficient way to ensure users remain logged in?

A. Enable smart sessions on the load balancer and modify the application to check tor an existing session.
B. Enable session sharing on the toad balancer and modify the application to read from the session store.
C. Store user session information in an Amazon S3 bucket and modify the application to read session information from the bucket.
D. Modify the application to store user session information in an Amazon ElastiCache cluster.

Answer: D

Explanation:
 https://aws.amazon.com/caching/session-management/

NEW QUESTION 10
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 11
A company uses a series of individual Amazon Cloud Formation templates to deploy its multi-Region Applications. These templates must be deployed in a specific order. The company is making more changes to the templates than previously expected and wants to deploy new templates more efficiently. Additionally, the data engineering team must be notified of all changes to the templates.
What should the company do to accomplish these goals?

A. Create an AWS Lambda function to deploy the Cloud Formation templates m the required order Use stack policies to alert the data engineering team.
B. Host the Cloud Formation templates in Amazon S3 Use Amazon S3 events to directly trigger CloudFormation updates and Amazon SNS notifications.
C. Implement CloudFormation StackSets and use drift detection to trigger update alerts to the data engineering team.
D. Leverage CloudFormation nested stacks and stack sets (or deployments Use Amazon SNS to notify the data engineering team.

Answer: D

Explanation:
 This solution will meet the requirements because it will use CloudFormation nested stacks and stack sets to deploy the templates more efficiently and consistently across multiple regions. Nested stacks allow the company to separate out common components and reuse templates, while stack sets allow the company to create stacks in multiple accounts and regions with a single template. The company can also use Amazon SNS to send notifications to the data engineering team whenever a change is made to the templates or the stacks. Amazon SNS is a service that allows you to publish messages to subscribers, such as email addresses, phone numbers, or other AWS services. By using Amazon SNS, the company can ensure that the data engineering team is aware of all changes to the templates and can take appropriate actions if needed. What is Amazon SNS? - Amazon Simple Notification Service

NEW QUESTION 16
A company has a legacy application A DevOps engineer needs to automate the process of building the deployable artifact for the legacy application. The solution must store the deployable artifact in an existing Amazon S3 bucket for future deployments to reference
Which solution will meet these requirements in the MOST operationally efficient way?

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

Answer: A

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

NEW QUESTION 18
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 even
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 even
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 even
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 21
A company uses AWS and has a VPC that contains critical compute infrastructure with predictable traffic patterns. The company has configured VPC flow logs that are published to a log group in Amazon CloudWatch Logs.
The company's DevOps team needs to configure a monitoring solution for the VPC flow logs to identify anomalies in network traffic to the VPC over time. If the monitoring solution detects an anomaly, the company needs the ability to initiate a response to the anomaly.
How should the DevOps team configure the monitoring solution to meet these requirements?

A. Create an Amazon Kinesis data strea
B. Subscribe the log group to the data strea
C. Configure Amazon Kinesis Data Analytics to detect log anomalies in the data strea
D. Create anAWS Lambda function to use as the output of the data strea
E. Configure the Lambda function to write to the default Amazon EventBridge event bus in the event of an anomaly finding.
F. Create an Amazon Kinesis Data Firehose delivery stream that delivers events to an Amazon S3 bucke
G. Subscribe the log group to the delivery strea
H. Configure Amazon Lookout for Metrics to monitor the data in the S3 bucket for anomalie
I. Create an AWS Lambda function to run in response to Lookout for Metrics anomaly finding
J. Configure the Lambda function to publish to the default Amazon EventBridge event bus.
K. Create an AWS Lambda function to detect anomalie
L. Configure the Lambda function to publish an event to the default Amazon EventBridge event bus if the Lambda function detects an anomal
M. Subscribe the Lambda function to the log group.
N. Create an Amazon Kinesis data strea
O. Subscribe the log group to the data strea
P. Create an AWS Lambda function to detect log anomalie
Q. Configure the Lambda function to write to the default Amazon EventBridge event bus if the Lambda function detects an anomal
R. Set the Lambda function as the processor for the data stream.

Answer: D

Explanation:
To meet the requirements, the DevOps team needs to configure a monitoring solution for the VPC flow logs that can detect anomalies in network traffic over time and initiate a response to the anomaly. The DevOps team can use Amazon Kinesis Data Streams to ingest and process streaming data from CloudWatch Logs. The DevOps team can subscribe the log group to a Kinesis data stream, which will deliver log events from CloudWatch Logs to Kinesis Data Streams in near real-time. The DevOps team can then create an AWS Lambda function to detect log anomalies using machine learning or statistical methods. The Lambda function can be set as a processor for the data stream, which means that it will process each record from the stream before sending it to downstream applications or destinations. The Lambda function can also write to the default Amazon EventBridge event bus if it detects an anomaly, which will allow other AWS services or custom applications to respond to the anomaly event.

NEW QUESTION 23
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 1AM 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 1AM 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 toke
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 projec
D. In the environment variable, include the ARN of the CodeBuild project's lAM service rol
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 repositor
G. Add an ECR repository policy that allows the 1AM service role to have access.
H. Update the buildspec.yml file to use the AWS CLI to assume the 1AM service role for ECR operation
I. Add an ECR repository policy that allows the 1AM service role to have access.

Answer: A

Explanation:
 (A) When Docker communicates with an Amazon Elastic Container Registry (ECR) repository, it requires authentication. You can authenticate your Docker client to the Amazon ECR registry with the help of the AWS CLI (Command Line Interface). Specifically, you can use the "aws ecr get-login-password" command to get an authorization token and then use Docker's "docker login" command with that token to authenticate to the registry. You would need to perform these steps in your buildspec.yml file before attempting to push or pull images from/to the ECR repository.

NEW QUESTION 24
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 bucke
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 Larnbda functio
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 sourc
H. Configure the rule to match events delivered by CloudTrai
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 28
A company has an application that runs on Amazon EC2 instances that are in an Auto Scaling group. When the application starts up. the application needs to process data from an Amazon S3 bucket before the application can start to serve requests.
The size of the data that is stored in the S3 bucket is growing. When the Auto Scaling group adds new instances, the application now takes several minutes to download and process the data before the application can serve requests. The company must reduce the time that elapses before new EC2 instances are ready to serve requests.
Which solution is the MOST cost-effective way to reduce the application startup time?

A. Configure a warm pool for the Auto Scaling group with warmed EC2 instances in the Stopped stat
B. Configure an autoscaling:EC2_INSTANCE_LAUNCHING lifecycle hook on the Auto Scaling grou
C. Modify the application to complete the lifecycle hook when the application is ready to serve requests.
D. Increase the maximum instance count of the Auto Scaling grou
E. Configure an autoscaling:EC2_INSTANCE_LAUNCHING lifecycle hook on the Auto Scaling grou
F. Modify the application to complete the lifecycle hook when the application is ready to serve requests.
G. Configure a warm pool for the Auto Scaling group with warmed EC2 instances in the Running stat
H. Configure an autoscaling:EC2_INSTANCE_LAUNCHING lifecycle hook on the Auto Scaling grou
I. Modify the application to complete the lifecycle hook when the application is ready to serve requests.
J. Increase the maximum instance count of the Auto Scaling grou
K. Configure an autoscaling:EC2_INSTANCE_LAUNCHING lifecycle hook on the Auto Scaling grou
L. Modify the application to complete the lifecycle hook and to place the new instance in the Standby state when the application is ready to serve requests.

Answer: A

Explanation:
Option A is the most cost-effective solution. By configuring a warm pool of EC2 instances in the Stopped state, the company can reduce the time it takes for new instances to be ready to serve requests. When the Auto Scaling group launches a new instance, it can attach the stopped EC2 instance from the warm pool. The instance can then be started up immediately, rather than having to wait for the data to be downloaded and processed. This reduces the overall startup time for the application.

NEW QUESTION 29
A company plans to use Amazon CloudWatch to monitor its Amazon EC2 instances. The company needs to stop EC2 instances when the average of the NetworkPacketsIn metric is less than 5 for at least 3 hours in a 12-hour time window. The company must evaluate the metric every hour. The EC2 instances must continue to run if there is missing data for the NetworkPacketsIn metric during the evaluation period.
A DevOps engineer creates a CloudWatch alarm for the NetworkPacketsIn metric. The DevOps engineer configures a threshold value of 5 and an evaluation period of 1 hour.
Which set of additional actions should the DevOps engineer take to meet these requirements?

A. Configure the Datapoints to Alarm value to be 3 out of 12. Configure the alarm to treat missing data as breaching the threshol
B. Add an AWS Systems Manager action to stop the instance when the alarm enters the ALARM state.
C. Configure the Datapoints to Alarm value to be 3 out of 12. Configure the alarm to treat missing data as not breaching the threshol
D. Add an EC2 action to stop the instance when the alarm enters the ALARM state.
E. Configure the Datapoints to Alarm value to be 9 out of 12. Configure the alarm to treat missing data as breaching the threshol
F. Add an EC2 action to stop the instance when the alarm enters the ALARM state.
G. Configure the Datapoints to Alarm value to be 9 out of 12. Configure the alarm to treat missing data as not breaching the threshol
H. Add an AWS Systems Manager action to stop the instance when the alarm enters the ALARM state.

Answer: B

Explanation:
To meet the requirements, the DevOps engineer needs to configure the CloudWatch alarm to stop the EC2 instances when the average of the NetworkPacketsIn metric is less than 5 for at least 3 hours in a 12-hour time window. This means that the alarm should trigger when 3 out of 12 datapoints are below the threshold of 5. The alarm should also treat missing data as not breaching the threshold, so that the EC2 instances continue to run if there is no data for the metric during the evaluation period. The DevOps engineer can add an EC2 action to stop the instance when the alarm enters the ALARM state, which is a built-in action type for

CloudWatch alarms.

NEW QUESTION 32
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 37
A company is using AWS CodePipeline to automate its release pipeline. AWS CodeDeploy is being used in the pipeline to deploy an application to Amazon Elastic Container Service (Amazon ECS) using the blue/green deployment model. The company wants to implement scripts to test the green version of the application before shifting traffic. These scripts will complete in 5 minutes or less. If errors are discovered during these tests, the application must be rolled back.
Which strategy will meet these requirements?

A. Add a stage to the CodePipeline pipeline between the source and deploy stage
B. Use AWS CodeBuild to create a runtime environment and build commands in the buildspec file to invoke test script
C. If errors are found, use the aws deploy stop-deployment command to stop the deployment.
D. Add a stage to the CodePipeline pipeline between the source and deploy stage
E. Use this stage to invoke an AWS Lambda function that will run the test script
F. If errors are found, use the aws deploy stop-deployment command to stop the deployment.
G. Add a hooks section to the CodeDeploy AppSpec fil
H. Use the AfterAllowTestTraffic lifecycle event to invoke an AWS Lambda function to run the test script
I. If errors are found, exit the Lambda function with an error to initiate rollback.
J. Add a hooks section to the CodeDeploy AppSpec fil
K. Use the AfterAllowTraffic lifecycle event to invoke the test script
L. If errors are found, use the aws deploy stop-deployment CLI command to stop the deployment.

Answer: C

Explanation:
 https://docs.aws.amazon.com/codedeploy/latest/userguide/reference-appspec-file-structure-hooks.html

NEW QUESTION 39
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 enable
B. Save the template to an Amazon S3 bucket that has been shared with all accounts within the compan
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 CL
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 expressio
H. Apply the SCP to all AWS account
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 accoun
K. Build a pipeline withAWS CodePipeline with a stage in AWS Lambda to assume the IAM role, and list all EBS volumes in the accoun
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 42
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 47
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 toke
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 projec
D. In the environment variable, include the ARN of the CodeBuild project's IAM service rol
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 repositor
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 authentica-tion 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 48
A development team wants to use AWS CloudFormation stacks to deploy an application. However, the developer IAM role does not have the required permissions to provision the resources that are specified in the AWS CloudFormation template. A DevOps engineer needs to implement a solution that allows the developers to deploy the stacks. The solution must follow the principle of least privilege.
Which solution will meet these requirements?

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

Answer: D

Explanation:
https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-iam-servicerole.html

NEW QUESTION 49
An ecommerce company has chosen AWS to host its new platform. The company's DevOps team has started building an AWS Control Tower landing zone. The DevOps team has set the identity store within AWS IAM Identity Center (AWS Single Sign-On) to external identity provider (IdP) and has configured SAML 2.0. The DevOps team wants a robust permission model that applies the principle of least privilege. The model must allow the team to build and manage only the team's own resources.
Which combination of steps will meet these requirements? (Choose three.)

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

Answer: BCF

Explanation:
 Using the principalTag in the Permission Set inline policy a logged in user belonging to a specific AD group in the IDP can be permitted access to perform operations on certain resources if their group matches the group used in the PrincipleTag. Basically you are narrowing the scope of privileges assigned via Permission policies conditionally based on whether the logged in user belongs to a specific AD Group in IDP. The mapping of the AD group to the request attributes can be done using SSO attributes where we can pass other attributes like the SAML token as well.
https://docs.aws.amazon.com/singlesignon/latest/userguide/abac.html

NEW QUESTION 51
A DevOps engineer has implemented a CI/CO pipeline to deploy an AWS Cloud Format ion 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 /varar/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 56
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-xxxxxxxxxx"
                    ]
                }
            }
        }
    ]
}
```

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 thatare 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 59
A company is hosting a web application in an AWS Region. For disaster recovery purposes, a second region is being used as a standby. Disaster recovery requirements state that session data must be replicated between regions in near-real time and 1% of requests should route to the secondary region to continuously verify system functionality. Additionally, if there is a disruption in service in the main region, traffic should be automatically routed to the secondary region, and the secondary region must be able to
scale up to handle all traffic.
How should a DevOps engineer meet these requirements?

A. In both regions, deploy the application on AWS Elastic Beanstalk and use Amazon DynamoDB global tables for session dat
B. Use an Amazon Route 53 weighted routing policy with health checks to distribute the traffic across the regions.
C. In both regions, launch the application in Auto Scaling groups and use DynamoDB for session dat
D. Use a Route 53 failover routing policy with health checks to distribute the traffic across the regions.
E. In both regions, deploy the application in AWS Lambda, exposed by Amazon API Gateway, and use Amazon RDS for PostgreSQL with cross-region replication for session dat
F. Deploy the web application with client-side logic to call the API Gateway directly.
G. In both regions, launch the application in Auto Scaling groups and use DynamoDB global tables for session dat
H. Enable an Amazon CloudFront weighted distribution across region
I. Point the Amazon Route 53 DNS record at the CloudFront distribution.

Answer: D

NEW QUESTION 63
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 66
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 neede
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 templat
F. To retrieve Virtual Private Cloud (VPC) and subnet values, use Fn::ImportValue intrinsic functions in the Parameters section of the root templat
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 CreateChangeSe
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 67
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 activatio
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 72
A company hosts a security auditing application in an AWS account. The auditing application uses an IAM role to access other AWS accounts. All the accounts are in the same organization in AWS Organizations.
A recent security audit revealed that users in the audited AWS accounts could modify or delete the auditing application's IAM role. The company needs to prevent

any modification to the auditing application's IAM role by any entity other than a trusted administrator IAM role.
Which solution will meet these requirements?

A. Create an SCP that includes a Deny statement for changes to the auditing application's IAM rol
B. Include a condition that allows the trusted administrator IAM role to make change
C. Attach the SCP to the root of the organization.
D. Create an SCP that includes an Allow statement for changes to the auditing application's IAM role by the trusted administrator IAM rol
E. Include a Deny statement for changes by all other IAM principal
F. Attach the SCP to the IAM service in each AWS account where the auditing application has an IAM role.
G. Create an IAM permissions boundary that includes a Deny statement for changes to the auditing application's IAM rol
H. Include a condition that allows the trusted administrator IAM role to make change
I. Attach the permissions boundary to the audited AWS accounts.
J. Create an IAM permissions boundary that includes a Deny statement for changes to the auditing application's IAM rol
K. Include a condition that allows the trusted administrator IAM role to make change
L. Attach the permissions boundary to the auditing application's IAM role in the AWS accounts.

Answer: A

Explanation:
 https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps. html?icmpid=docs_orgs_console
SCPs (Service Control Policies) are the best way to restrict permissions at the organizational level, which in this case would be used to restrict modifications to the IAM role used by the auditing application, while still allowing trusted administrators to make changes to it. Options C and D are not as effective because IAM permission boundaries are applied to IAM entities (users, groups, and roles), not the account itself, and must be applied to all IAM entities in the account.

NEW QUESTION 75
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 profil
F. Attach the IAM role to the profile.
G. Update the launch templat
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 78
A company has an application that is using a MySQL-compatible Amazon Aurora Multi-AZ DB cluster as the database. A cross-Region read replica has been created for disaster recovery purposes. A DevOps engineer wants to automate the promotion of the replica so it becomes the primary database instance in the event of a failure.
Which solution will accomplish this?

A. Configure a latency-based Amazon Route 53 CNAME with health checks so it points to both the primary and replica endpoint
B. Subscribe an Amazon SNS topic to Amazon RDS failure notifications from AWS CloudTrail and use that topic to invoke an AWS Lambda function that will

promote the replica instance as the primary.
C. Create an Aurora custom endpoint to point to the primary database instanc
D. Configure the application to use this endpoin
E. Configure AWS CloudTrail to run an AWS Lambda function to promote the replica instance and modify the custom endpoint to point to the newly promoted instance.
F. Create an AWS Lambda function to modify the application's AWS CloudFormation template to promote the replica, apply the template to update the stack, and point the application to the newly promoted instanc
G. Create an Amazon CloudWatch alarm to invoke this Lambda function after the failure event occurs.
H. Store the Aurora endpoint in AWS Systems Manager Parameter Stor
I. Create an Amazon EventBridge event that detects the database failure and runs an AWS Lambda function to promote the replica instance and update the endpoint URL stored in AWS Systems Manager Parameter Stor
J. Code the application to reload the endpoint from Parameter Store if a database connection fails.

Answer: D

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

NEW QUESTION 79
A company hosts applications in its AWS account Each application logs to an individual Amazon CloudWatch log group. The company's CloudWatch costs for ingestion are increasing
A DevOps engineer needs to Identify which applications are the source of the increased logging costs.
Which solution Will meet these requirements?

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

Answer: C

Explanation:
The correct answer is C.
A comprehensive and detailed explanation is:
? Option A is incorrect because using CloudWatch metrics to create a custom expression that identifies the CloudWatch log groups that have the most data being written to them is not a valid solution. CloudWatch metrics do not provide information about the size or volume of data being ingested by CloudWatch logs. CloudWatch metrics only provide information about the number of events, bytes, and errors that occur within a log group or stream. Moreover, creating a custom expression with CloudWatch metrics would require using the search_web tool, which is not necessary for this use case.
? Option B is incorrect because using CloudWatch Logs Insights to create a set of queries for the application log groups to identify the number of logs written for a period of time is not a valid solution. CloudWatch Logs Insights can help analyze and filter log events based on patterns and expressions, but it does not provide information about the cost or billing of CloudWatch logs. CloudWatch Logs Insights also charges based on the amount of data scanned by each query, which could increase the logging costs further.
? Option C is correct because using AWS Cost Explorer to generate a cost report that details the cost for CloudWatch usage is a valid solution. AWS Cost Explorer is a tool that helps visualize, understand, and manage AWS costs and usage over time. AWS Cost Explorer can generate custom reports that show the breakdown of costs by service, region, account, tag, or any other dimension. AWS Cost Explorer can also filter and group costs by usage type, which can help identify the specific CloudWatch log groups that are the source of the increased logging costs.
? Option D is incorrect because using AWS CloudTrail to filter for CreateLogStream events for each application is not a valid solution. AWS CloudTrail is a service that records API calls and account activity for AWS services, including CloudWatch logs. However, AWS CloudTrail does not provide information about the cost or billing of CloudWatch logs. Filtering for CreateLogStream events would only show when a new log stream was created within a log group, but not how much data was ingested or stored by that log stream.
References:
? CloudWatch Metrics
? CloudWatch Logs Insights
? AWS Cost Explorer
? AWS CloudTrail

NEW QUESTION 80
A security review has identified that an AWS CodeBuild project is downloading a database population script from an Amazon S3 bucket using an unauthenticated request. The security team does not allow unauthenticated requests to S3 buckets for this project.
How can this issue be corrected in the MOST secure manner?

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

Answer: C

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

credentials and encryption to secure the data transfer.
For more information, you can refer to these web pages:
? [Using bucket policies and user policies - Amazon Simple Storage Service]
? [Create a service role for CodeBuild - AWS CodeBuild]
? [AWS Command Line Interface]

NEW QUESTION 85

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

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

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

What should the DevOps engineer do to meet these requirements?

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

Answer: A

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

NEW QUESTION 90

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 metnc for the SQS queue Configure a redrive policy on the SQS queue.
D. Check the NumberOfMessagesSent metric for the SQS queu
E. Increase the SQS queue visibility timeout.
F. Check the WriteThrottleEvents metric for the DynamoDB tabl
G. Increase the maximum write capacity units (WCUs) for the table's scaling policy.
H. Check the Throttles metric for the Lambda functio
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 ThottledWriteRequests 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 91

A company's DevOps engineer is creating an AWS Lambda function to process notifications from an Amazon Simple Notification Service (Amazon SNS) topic. The Lambda function will process the notification messages and will write the contents of the notification messages to an Amazon RDS Multi-AZ DB instance.
During testing a database administrator accidentally shut down the DB instance. While the database was down the company lost several of the SNS notification messages that were delivered during that time.
The DevOps engineer needs to prevent the loss of notification messages in the future Which solutions will meet this requirement? (Select TWO.)

A. Replace the RDS Multi-AZ DB instance with an Amazon DynamoDB table.
B. Configure an Amazon Simple Queue Service (Amazon SQS) queue as a destination of the Lambda function.
C. Configure an Amazon Simple Queue Service (Amazon SQS> dead-letter queue for the SNS topic.
D. Subscribe an Amazon Simple Queue Service (Amazon SQS) queue to the SNS topic Configure the Lambda function to process messages from the SQS queue.
E. Replace the SNS topic with an Amazon EventBridge event bus Configure an EventBridge rule on the new event bus to invoke the Lambda function for each event.

Answer: CD

Explanation:

These solutions will meet the requirement because they will prevent the loss of notification messages in the future. An Amazon SQS queue is a service that provides a reliable, scalable, and secure message queue for asynchronous communication between distributed components. You can use an SQS queue to buffer messages from an SNS topic and ensure that they are delivered and processed by a Lambda function, even if the function or the database is temporarily unavailable.

Option C will configure an SQS dead-letter queue for the SNS topic. A dead-letter queue is a queue that receives messages that could not be delivered to any subscriber after a specified number of retries. You can use a dead-letter queue to store and analyze failed messages, or to reprocess them later. This way, you can avoid losing messages that could not be delivered to the Lambda function due to network errors, throttling, or other issues. Option D will subscribe an SQS queue to the SNS topic and configure the Lambda function to process messages from the SQS queue. This will decouple the SNS topic from the Lambda function and provide more flexibility and control over the message delivery and processing. You can use an SQS queue to store messages from the SNS topic until they are ready to be processed by the Lambda function, and also to retry processing in case of failures. This way, you can avoid losing messages that could not be processed by the Lambda function due to database errors, timeouts, or other issues.

NEW QUESTION 96
A DevOps team is merging code revisions for an application that uses an Amazon RDS Multi-AZ DB cluster for its production database. The DevOps team uses continuous integration to periodically verify that the application works. The DevOps team needs to test the changes before the changes are deployed to the production database.
Which solution will meet these requirements'?

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

Answer: A

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

NEW QUESTION 97
A company has migrated its container-based applications to Amazon EKS and want to establish automated email notifications. The notifications sent to each email address are for specific activities related to EKS components. The solution will include Amazon SNS topics and an AWS Lambda function to evaluate incoming log events and publish messages to the correct SNS topic.
Which logging solution will support these requirements?

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

Answer: A

Explanation:
https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/SubscriptionFilters.html#La mbdaFunctionExample
https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/SubscriptionFilters.html

NEW QUESTION 98
A company requires its internal business teams to launch resources through pre-approved AWS CloudFormation templates only. The security team requires automated monitoring when resources drift from their expected state.
Which strategy should be used to meet these requirements?

A. Allow users to deploy CloudFormation stacks using a CloudFormation service role onl
B. Use CloudFormation drift detection to detect when resources have drifted from their expected state.
C. Allow users to deploy CloudFormation stacks using a CloudFormation service role onl
D. Use AWS Config rules to detect when resources have drifted from their expected state.
E. Allow users to deploy CloudFormation stacks using AWS Service Catalog onl
F. Enforce the use of a launch constrain
G. Use AWS Config rules to detect when resources have drifted from their expected state.
H. Allow users to deploy CloudFormation stacks using AWS Service Catalog onl
I. Enforce the use of a template constrain
J. Use Amazon EventBridge notifications to detect when resources have drifted from their expected state.

Answer: C

Explanation:
The correct answer is C. Allowing users to deploy CloudFormation stacks using AWS Service Catalog only and enforcing the use of a launch constraint is the best way to ensure that the internal business teams launch resources through pre-approved CloudFormation templates only. AWS Service Catalog is a service that enables organizations to create and manage catalogs of IT services that are approved for use on AWS. A launch constraint is a rule that specifies the role that AWS Service Catalog assumes when launching a product.

By using a launch constraint, the DevOps engineer can control the permissions that the users have when launching a product. Using AWS Config rules to detect when resources have drifted from their expected state is the best way to automate the monitoring of the resources. AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. AWS Config rules are custom or managed rules that AWS Config uses to evaluate whether your AWS resources comply with your desired configurations. By using AWS Config rules, the DevOps engineer can track the changes in the resources and identify any non-compliant resources.

Option A is incorrect because allowing users to deploy CloudFormation stacks using a CloudFormation service role only is not the best way to ensure that the internal business teams launch resources through pre-approved CloudFormation templates only. A CloudFormation service role is an IAM role that CloudFormation assumes to create, update, or delete the stack resources. By using a CloudFormation service role, the DevOps engineer can control the permissions that CloudFormation has when acting on the resources, but not the permissions that the users have when launching a stack. Therefore, option A does not prevent the users from launching resources that are not approved by the company. Using CloudFormation drift detection to detect when resources have drifted from their expected state is a valid way to monitor the resources, but it is not as automated and scalable as using AWS Config rules. CloudFormation drift detection is a feature that enables you to detect whether a stack's actual configuration differs, or has drifted, from its expected configuration. To use this feature, the DevOps engineer would need to manually initiate a drift detection operation on the stack or the stack resources, and then view the drift status and details in the CloudFormation console or API.

Option B is incorrect because allowing users to deploy CloudFormation stacks using a CloudFormation service role only is not the best way to ensure that the internal business teams launch resources through pre-approved CloudFormation templates only, as explained in option A. Using AWS Config rules to detect when resources have drifted from their expected state is a valid way to monitor the resources, as explained in option C. Option D is incorrect because enforcing the use of a template constraint is not the best way to ensure that the internal business teams launch resources through pre-approved CloudFormation templates only. A template constraint is a rule that defines the values or properties that users can specify when launching a product. By using a template constraint, the DevOps engineer can control the parameters that the users can provide when launching a product, but not the permissions that the users have when launching a product. Therefore, option D does not prevent the users from launching resources that are not approved by the company. Using Amazon EventBridge notifications to detect when resources have drifted from their expected state is a less reliable and consistent solution than using AWS Config rules. Amazon EventBridge is a service that enables you to connect your applications with data from a variety of sources. Amazon EventBridge can deliver a stream of real-time data from event sources, such as AWS services, and route

that data to targets, such as AWS Lambda functions. However, to use this solution, the DevOps engineer would need to configure the event source, the event bus, the event rule, and the event target for each resource type that needs to be monitored, which is more complex and error-prone than using AWS Config rules.

NEW QUESTION 100
A company has multiple AWS accounts. The company uses AWS IAM Identity Center (AWS Single Sign-On) that is integrated with AWS Toolkit for Microsoft Azure DevOps. The attributes for access control feature is enabled in IAM Identity Center.
The attribute mapping list contains two entries. The department key is mapped to
${path:enterprise.department}. The costCenter key is mapped to
${path:enterprise.costCenter}.
All existing Amazon EC2 instances have a department tag that corresponds to three company departments (d1, d2, d3). A DevOps engineer must create policies based on the matching attributes. The policies must minimize administrative effort and must grant each Azure AD user access to only the EC2 instances that are tagged with the user's respective department name.
Which condition key should the DevOps engineer include in the custom permissions policies to meet these requirements?
A.

```
"Condition": {
    "ForAllValues:StringEquals": {
        "aws:TagKeys": ["department"]
    )
}
```

B.

```
"Condition": {
    "StringEquals": {
        "aws:PrincipalTag/department": "$(aws:ResourceTag/department)"
    )
}
```

C.

```
"Condition": {
    "StringEquals": {
        "ec2:ResourceTag/department": "$(aws:PrincipalTag/department)"
    )
}
```

D.

```
"Condition": {
    "ForAllValues:StringEquals": {
        "ec2:ResourceTag/department": ["d1", "d2", "d3"]
    )
}
```

A.

Answer: C

Explanation:
https://docs.aws.amazon.com/singlesignon/latest/userguide/configure- abac.html

NEW QUESTION 102
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 103
A company has chosen AWS to host a new application. The company needs to implement a multi-account strategy. A DevOps engineer creates a new AWS account and an organization in AWS Organizations. The DevOps engineer also creates the OU structure for the organization and sets up a landing zone by using AWS Control Tower.
The DevOps engineer must implement a solution that automatically deploys resources for new accounts that users create through AWS Control Tower Account Factory. When a user creates a new account, the solution must apply AWS CloudFormation templates and SCPs that are customized for the OU or the account to automatically deploy all the resources that are attached to the account. All the OUs are enrolled in AWS Control Tower.
Which solution will meet these requirements in the MOST automated way?

A. Use AWS Service Catalog with AWS Control Towe
B. Create portfolios and products in AWS Service Catalo
C. Grant granular permissions to provision these resource
D. Deploy SCPs by using the AWS CLI and JSON documents.
E. Deploy CloudFormation stack sets by using the required template
F. Enable automatic deploymen
G. Deploy stack instances to the required account
H. Deploy a CloudFormationstack set to the organization's management account to deploy SCPs.
I. Create an Amazon EventBridge rule to detect the CreateManagedAccount even
J. Configure AWS Service Catalog as the target to deploy resources to any new account
K. Deploy SCPs by using the AWS CLI and JSON documents.
L. Deploy the Customizations for AWS Control Tower (CfCT) solutio
M. Use an AWS CodeCommit repository as the sourc
N. In the repository, create a custom package that includes the CloudFormation templates and the SCP JSON documents.

Answer: D

Explanation:
The CfCT solution is designed for the exact purpose stated in the question. It extends the capabilities of AWS Control Tower by providing you with a way to automate resource provisioning and apply custom configurations across all AWS accounts created in the Control Tower environment. This enables the company to implement additional account customizations when new accounts are provisioned via the Control Tower Account Factory. The CloudFormation templates and SCPs can be added to a CodeCommit repository and will be automatically deployed to new accounts when they are created. This provides a highly automated solution that does not require manual intervention to deploy resources and SCPs to new accounts.

NEW QUESTION 105
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 bucke
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 bucke
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 bucke
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 107
A DevOps engineer has developed an AWS Lambda function The Lambda function starts an AWS CloudFormation drift detection operation on all supported resources for a specific CloudFormation stack The Lambda function then exits Its invocation The DevOps engineer has created an Amazon EventBrdge scheduled rule that Invokes the Lambda function every hour. An Amazon Simple Notification Service (Amazon SNS) topic already exists In the AWS account. The DevOps engineer has subscribed to the SNS topic to receive notifications
The DevOps engineer needs to receive a notification as soon as possible when drift is detected in this specific stack configuration.
Which solution Will meet these requirements?

A. Configure the existing EventBridge rule to also target the SNS topic Configure an SNS subscription filter policy to match the Cloud Formation stac
B. Attach the subscription filter policy to the SNS tomc
C. Create a second Lambda function to query the CloudFormation API for the drift detection results for the stack Configure the second Lambda function to publish a message to the SNS topic If drift ts detected Adjust the existing EventBridge rule to also target the second Lambda function
D. Configure Amazon GuardDuty in the account with drift detection for all CloudFormation stack
E. Create a second EventBndge rule that reacts to the GuardDuty drift detection event finding for the specific CloudFormation stac
F. Configure the SNS topic as a target of the second EventBridge rule.
G. Configure AWS Config in the accoun
H. Use the cloudformation-stack-drift-detection- check managed rul
I. Create a second EventBndge rule that reacts to a compliance change event for the CloudFormaUon stac
J. Configure the SNS topc as a target of the second EventBridge rule.

Answer: D

Explanation:
A comprehensive and detailed explanation is:
? Option A is incorrect because EventBridge rules cannot filter events based on the message body or attributes of the target service. Therefore, configuring an SNS subscription filter policy to match the CloudFormation stack will not work. The SNS topic will receive all events from the EventBridge rule, regardless of the stack name or drift status.
? Option B is incorrect because it introduces unnecessary complexity and cost.
Creating a second Lambda function to query the CloudFormation API for the drift detection results is redundant, since CloudFormation already publishes drift detection events to EventBridge. Moreover, invoking two Lambda functions every hour will incur more charges than invoking one.
? Option C is incorrect because GuardDuty does not provide drift detection for CloudFormation stacks. GuardDuty is a threat detection service that monitors for malicious activity and unauthorized behavior in AWS accounts and workloads. It does not monitor or report on configuration changes or drifts in CloudFormation stacks.
? Option D is correct because it leverages AWS Config and its managed rule for drift detection. AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. It can detect configuration changes and drifts in CloudFormation stacks using the cloudformation-stack-drift-detection- check managed rule. This rule triggers an AWS Config event when a stack drifts from its expected template configuration. By creating a second EventBridge rule that reacts to this event for the specific stack, the DevOps engineer can configure the SNS topic as a target and receive a notification as soon as possible when drift is detected.
References:
? AWS Config
? Amazon SNS subscription filter policies
? Amazon EventBridge rules

NEW QUESTION 111
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 topi
B. Create an AWS CloudTrail subscription for the SNS topi
C. Configure a subscription filter for security group modification events.
D. Create an Amazon EventBridge scheduled rule to invoke the Lambda functio
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 sourc
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 service
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 115

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

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

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

Answer: AD

Explanation:
https://docs.aws.amazon.com/codedeploy/latest/userguide/troubleshooting- deployments.html#troubleshooting-skipped-lifecycle-events

NEW QUESTION 120

A company recently launched multiple applications that use Application Load Balancers. Application response time often slows down when the applications experience problems A DevOps engineer needs to Implement a monitoring solution that alerts the company when the applications begin to perform slowly The DevOps engineer creates an Amazon Simple Notification Semce (Amazon SNS) topic and subscribe the company's email address to the topic
What should the DevOps engineer do next to meet the requirements?

A. Create an Amazon EventBridge rule that invokes an AWS Lambda function to query the applications on a 5-minute interval Configure the Lambda function to publish a notification to the SNS topic when the applications return errors.
B. Create an Amazon CloudWatch Synthetics canary that runs a custom script to query the applications on a 5-minute interva
C. Configure the canary to use the SNS topic when the applications return errors.
D. Create an Amazon CloudWatch alarm that uses the AWS/AppljcabonELB namespace RequestCountPerTarget metric Configure the CloudWatch alarm to send a notification when the number of connections becomes greater than the configured number of threads that the application supports Configure the CloudWatch alarm to use the SNS topic.
E. Create an Amazon CloudWatch alarm that uses the AWS/ApplicationELB namespace RequestCountPerTarget metric Configure the CloudWatch alarm to send a notification when the average response time becomes greater than the longest response time that the application supports Configure the CloudWatch alarm to use the SNS topic

Answer: B

Explanation:
? Option A is incorrect because creating an Amazon EventBridge rule that invokes an AWS Lambda function to query the applications on a 5-minute interval is not a valid solution. EventBridge rules can only trigger Lambda functions based on events, not on time intervals. Moreover, querying the applications on a 5-minute interval might incur unnecessary costs and network overhead, and might not detect performance issues in real time.
? Option B is correct because creating an Amazon CloudWatch Synthetics canary that runs a custom script to query the applications on a 5-minute interval is a valid solution. CloudWatch Synthetics canaries are configurable scripts that monitor endpoints and APIs by simulating customer behavior. Canaries can run as often as once per minute, and can measure the latency and availability of the
applications. Canaries can also send notifications to an Amazon SNS topic when they detect errors or performance issues1.
? Option C is incorrect because creating an Amazon CloudWatch alarm that uses the AWS/ApplicationELB namespace RequestCountPerTarget metric is not a valid solution. The RequestCountPerTarget metric measures the number of requests completed or connections made per target in a target group2. This metric does not reflect the application response time, which is the requirement. Moreover, configuring the CloudWatch alarm to send a notification when the number of connections becomes greater than the configured number of threads that the application supports is not a valid way to measure the application performance, as it depends on the application design and implementation.
? Option D is incorrect because creating an Amazon CloudWatch alarm that uses the AWS/ApplicationELB namespace RequestCountPerTarget metric is not a valid solution, for the same reason as option C. The RequestCountPerTarget metric does not reflect the application response time, which is the requirement. Moreover, configuring the CloudWatch alarm to send a notification when the average response time becomes greater than the longest response time that the application supports is not a valid way to measure the application performance, as it does not account for variability or outliers in the response time distribution. References:
? 1: Using synthetic monitoring
? 2: Application Load Balancer metrics

NEW QUESTION 122

A DevOps engineer has automated a web service deployment by using AWS CodePipeline with the following steps:
1) An AWS CodeBuild project compiles the deployment artifact and runs unit tests.
2) An AWS CodeDeploy deployment group deploys the web service to Amazon EC2 instances in the staging environment.
3) A CodeDeploy deployment group deploys the web service to EC2 instances in the production environment.
The quality assurance (QA) team requests permission to inspect the build artifact before the deployment to the production environment occurs. The QA team wants to run an internal penetration testing tool to conduct manual tests. The tool will be invoked by a REST API call.
Which combination of actions should the DevOps engineer take to fulfill this request? (Choose two.)

A. Insert a manual approval action between the test actions and deployment actions of the pipeline.
B. Modify the buildspec.yml file for the compilation stage to require manual approval before completion.
C. Update the CodeDeploy deployment groups so that they require manual approval to proceed.
D. Update the pipeline to directly call the REST API for the penetration testing tool.
E. Update the pipeline to invoke an AWS Lambda function that calls the REST API for thepenetration testing tool.

Answer: AE

NEW QUESTION 124

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

How can this process be automated?

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

Answer: D

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

NEW QUESTION 129
......

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

AWS-Certified-DevOps-Engineer-Professional Practice Exam Features:

* AWS-Certified-DevOps-Engineer-Professional Questions and Answers Updated Frequently

* AWS-Certified-DevOps-Engineer-Professional Practice Questions Verified by Expert Senior Certified Staff

* AWS-Certified-DevOps-Engineer-Professional Most Realistic Questions that Guarantee you a Pass on Your FirstTry

* AWS-Certified-DevOps-Engineer-Professional Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The AWS-Certified-DevOps-Engineer-Professional Practice Test Here](https://www.surepassexam.com/AWS-Certified-DevOps-Engineer-Professional-exam-dumps.html)