

# Amazon-Web-Services

## Exam Questions DOP-C02

AWS Certified DevOps Engineer - Professional



### NEW QUESTION 1

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

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

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

**Answer: C**

#### Explanation:

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

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

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

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

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

### NEW QUESTION 2

A company is examining its disaster recovery capability and wants the ability to switch over its daily operations to a secondary AWS Region. The company uses AWS CodeCommit as a source control tool in the primary Region.

A DevOps engineer must provide the capability for the company to develop code in the secondary Region. If the company needs to use the secondary Region, developers can add an additional remote URL to their local Git configuration.

Which solution will meet these requirements?

- A. Create a CodeCommit repository in the secondary Region
- B. Create an AWS CodeBuild project to perform a Git mirror operation of the primary Region's CodeCommit repository to the secondary Region's CodeCommit repository
- C. Create an AWS Lambda function that invokes the CodeBuild project
- D. Create an Amazon EventBridge rule that reacts to merge events in the primary Region's CodeCommit repository
- E. Configure the EventBridge rule to invoke the Lambda function.
- F. Create an Amazon S3 bucket in the secondary Region
- G. Create an AWS Fargate task to perform a Git mirror operation of the primary Region's CodeCommit repository and copy the result to the S3 bucket
- H. Create an AWS Lambda function that initiates the Fargate task
- I. Create an Amazon EventBridge rule that reacts to merge events in the CodeCommit repository
- J. Configure the EventBridge rule to invoke the Lambda function.
- K. Create an AWS CodeArtifact repository in the secondary Region
- L. Create an AWS CodePipeline pipeline that uses the primary Region's CodeCommit repository for the source action
- M. Create a Cross-Region stage in the pipeline that packages the CodeCommit repository contents and stores the contents in the CodeArtifact repository when a pull request is merged into the CodeCommit repository.
- N. Create an AWS Cloud9 environment and a CodeCommit repository in the secondary Region
- O. Configure the primary Region's CodeCommit repository as a remote repository in the AWS Cloud9 environment
- P. Connect the secondary Region's CodeCommit repository to the AWS Cloud9 environment.

**Answer: A**

#### Explanation:

The best solution to meet the disaster recovery capability and allow developers to switch over to a secondary AWS Region for code development is option A. This involves creating a CodeCommit repository in the secondary Region and setting up an AWS CodeBuild project to perform a Git mirror operation of the primary Region's CodeCommit repository to the secondary Region's repository. An AWS Lambda function is then created to invoke the CodeBuild project. Additionally, an Amazon EventBridge rule is configured to react to merge events in the primary Region's CodeCommit repository and invoke the Lambda function. This setup ensures that the secondary Region's repository is always up-to-date with the primary repository, allowing for a seamless transition in case of a disaster recovery event.

References:

? AWS CodeCommit User Guide on resilience and disaster recovery<sup>1</sup>.

? AWS Documentation on monitoring CodeCommit events in Amazon EventBridge and Amazon CloudWatch Events<sup>2</sup>.

### NEW QUESTION 3

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

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

- A. Create an AWS Organizations SCP that denies access to all non-global services in non-US Region
- B. Attach the policy to the root of the organization.
- C. Configure AWS CloudTrail to send logs to Amazon CloudWatch Logs and enable it for all Region
- D. Use a CloudWatch Logs metric filter to send an alert on any service activity in non-US Regions.
- E. Use an AWS Lambda function that checks for AWS service activity and deploy it to all Region

- F. Write an Amazon EventBridge rule that runs the Lambda function every hour, sending an alert if activity is found in a non-US Region.
- G. Use an AWS Lambda function to query Amazon Inspector to look for service activity in non-US Regions and send alerts if any activity is found.
- H. Write an SCP using the aws: RequestedRegion condition key limiting access to US Region
- I. Apply the policy to all users, groups, and roles

**Answer:** AB

**Explanation:**

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

? A. Create an AWS Organizations SCP that denies access to all non-global services in non-US Regions. Attach the policy to the root of the organization. This action will effectively prevent users and roles in all accounts within the organization from accessing services in non-US Regions<sup>12</sup>.

? B. Configure AWS CloudTrail to send logs to Amazon CloudWatch Logs and enable it for all Regions. Use a CloudWatch Logs metric filter to send an alert on any service activity in non-US Regions. This action will allow monitoring of all AWS Regions and will trigger alerts if any activity is detected in non-US Regions, ensuring that the governance team is notified as soon as possible<sup>3</sup>.

References:

? AWS Documentation on Service Control Policies (SCPs) and how they can be used to manage permissions and restrict access based on Regions<sup>12</sup>.

? AWS Documentation on monitoring CloudTrail log files with Amazon CloudWatch Logs to set up alerts for specific activities<sup>3</sup>.

**NEW QUESTION 4**

A company has a data ingestion application that runs across multiple AWS accounts. The accounts are in an organization in AWS Organizations. The company needs to monitor the application and consolidate access to the application. Currently the company is running the application on Amazon EC2 instances from several Auto Scaling groups. The EC2 instances have no access to the internet because the data is sensitive Engineers have deployed the necessary VPC endpoints. The EC2 instances run a custom AMI that is built specifically for the application.

To maintain and troubleshoot the application, system administrators need the ability to log in to the EC2 instances. This access must be automated and controlled centrally. The company's security team must receive a notification whenever the instances are accessed.

Which solution will meet these requirements?

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

**Answer:** C

**Explanation:**

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

**NEW QUESTION 5**

A company has a mobile application that makes HTTP API calls to an Application Load Balancer (ALB). The ALB routes requests to an AWS Lambda function. Many different versions of the application are in use at any given time, including versions that are in testing by a subset of users. The version of the application is defined in the user-agent header that is sent with all requests to the API.

After a series of recent changes to the API, the company has observed issues with the application. The company needs to gather a metric for each API operation by response code for each version of the application that is in use. A DevOps engineer has modified the Lambda function to extract the API operation name, version information from the user-agent header and response code.

Which additional set of actions should the DevOps engineer take to gather the required metrics?

- A. Modify the Lambda function to write the API operation name, response code, and version number as a log line to an Amazon CloudWatch Logs log group
- B. Configure a CloudWatch Logs metric filter that increments a metric for each API operation name
- C. Specify response code and application version as dimensions for the metric.
- D. Modify the Lambda function to write the API operation name, response code, and version number as a log line to an Amazon CloudWatch Logs log group
- E. Configure a CloudWatch Logs Insights query to populate CloudWatch metrics from the log line
- F. Specify response code and application version as dimensions for the metric.
- G. Configure the ALB access logs to write to an Amazon CloudWatch Logs log group
- H. Modify the Lambda function to respond to the ALB with the API operation name, response code, and version number as response metadata
- I. Configure a CloudWatch Logs metric filter that increments a metric for each API operation name
- J. Specify response code and application version as dimensions for the metric.
- K. Configure AWS X-Ray integration on the Lambda function
- L. Modify the Lambda function to create an X-Ray subsegment with the API operation name, response code, and version number
- M. Configure X-Ray insights to extract an aggregated metric for each API operation name and to publish the metric to Amazon CloudWatch
- N. Specify response code and application version as dimensions for the metric.

**Answer:** A

**Explanation:**

"Note that the metric filter is different from a log insights query, where the experience is interactive and provides immediate search results for the user to investigate.

No automatic action can be invoked from an insights query. Metric filters, on the other hand, will generate metric data in the form of a time series. This lets you

create alarms that integrate into your ITSM processes, execute AWS Lambda functions, or even create anomaly detection models."  
<https://aws.amazon.com/blogs/mt/quantify-custom-application-metrics-with-amazon-cloudwatch-logs-and-metric-filters/>

#### NEW QUESTION 6

A growing company manages more than 50 accounts in an organization in AWS Organizations. The company has configured its applications to send logs to Amazon CloudWatch Logs.

A DevOps engineer needs to aggregate logs so that the company can quickly search the logs to respond to future security incidents. The DevOps engineer has created a new AWS account for centralized monitoring.

Which combination of steps should the DevOps engineer take to make the application logs searchable from the monitoring account? (Select THREE.)

- A. In the monitoring account, download an AWS CloudFormation template from CloudWatch to use in Organization
- B. Use CloudFormation StackSets in the organization's management account to deploy the CloudFormation template to the entire organization.
- C. Create an AWS CloudFormation template that defines an IAM role
- D. Configure the role to allow logs.amazonaws.com to perform the logs:Link action if the aws:ResourceAccount property is equal to the monitoring account ID
- E. Use CloudFormation StackSets in the organization's management account to deploy the CloudFormation template to the entire organization.
- F. Create an IAM role in the monitoring account
- G. Attach a trust policy that allows logs.amazonaws.com to perform the iam:CreateSink action if the aws:PrincipalOrgId property is equal to the organization ID.
- H. In the organization's management account, enable the logging policies for the organization.
- I. Use CloudWatch Observability Access Manager in the monitoring account to create a sink
- J. Allow logs to be shared with the monitoring account
- K. Configure the monitoring account data selection to view the Observability data from the organization ID.
- L. In the monitoring account, attach the CloudWatchLogsReadOnlyAccess AWS managed policy to an IAM role that can be assumed to search the logs.

**Answer:** BCF

#### Explanation:

? To aggregate logs from multiple accounts in an organization, the DevOps engineer needs to create a cross-account subscription<sup>1</sup> that allows the monitoring account to receive log events from the sharing accounts.

? To enable cross-account subscription, the DevOps engineer needs to create an IAM role in each sharing account that grants permission to CloudWatch Logs to link the log groups to the destination in the monitoring account<sup>2</sup>. This can be done using a CloudFormation template and StackSets<sup>3</sup> to deploy the role to all accounts in the organization.

? The DevOps engineer also needs to create an IAM role in the monitoring account that allows CloudWatch Logs to create a sink for receiving log events from other accounts<sup>4</sup>. The role must have a trust policy that specifies the organization ID as a condition.

? Finally, the DevOps engineer needs to attach the

CloudWatchLogsReadOnlyAccess policy<sup>5</sup> to an IAM role in the monitoring account that can be used to search the logs from the cross-account subscription.

References: 1: Cross-account log data sharing with subscriptions 2: Create an IAM role for CloudWatch Logs in each sharing account 3: AWS CloudFormation StackSets 4: Create an IAM role for CloudWatch Logs in your monitoring account 5: CloudWatchLogsReadOnlyAccess policy

#### NEW QUESTION 7

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 application
- B. Configure an AWS Lambda function to add a tag to the EC2 instance that produced the login event and mark the instance to be decommissioned
- C. Create an Amazon EventBridge rule to invoke a second Lambda function once a day that will terminate all instances with this tag.
- D. Create an Amazon CloudWatch alarm that will be invoked by the login event
- E. Send the notification to an Amazon Simple Notification Service (Amazon SNS) topic that the operations team is subscribed to, and have them terminate the EC2 instance within 24 hours.
- F. Create an Amazon CloudWatch alarm that will be invoked by the login event
- G. Configure the alarm to send to an Amazon Simple Queue Service (Amazon SQS) queue
- H. Use a group of worker instances to process messages from the queue, which then schedules an Amazon EventBridge rule to be invoked.
- I. Create a CloudWatch Logs subscription to an AWS Lambda function
- J. Configure the function to add a tag to the EC2 instance that produced the login event and mark the instance to be decommissioned
- 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 8

A large enterprise is deploying a web application on AWS. The application runs on Amazon

EC2 instances behind an Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. The application stores data in an Amazon RDS for Oracle DB instance and Amazon DynamoDB. There are separate environments for development testing and production.

What is the MOST secure and flexible way to obtain password credentials during deployment?

- A. Retrieve an access key from an AWS Systems Manager securestring parameter to access AWS service
- B. Retrieve the database credentials from a Systems Manager SecureString parameter.
- C. Launch the EC2 instances with an EC2 IAM role to access AWS services Retrieve the database credentials from AWS Secrets Manager.
- D. Retrieve an access key from an AWS Systems Manager plaintext parameter to access AWS service
- E. Retrieve the database credentials from a Systems Manager SecureString parameter.
- F. Launch the EC2 instances with an EC2 IAM role to access AWS services Store the database passwords in an encrypted config file with the application artifacts.

**Answer:** B

#### Explanation:

AWS Secrets Manager is a secrets management service that helps you protect access to your applications, services, and IT resources. This service enables you to easily rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle. Using Secrets Manager, you can secure and manage secrets used to access resources in the AWS Cloud, on third-party services, and on-premises. SSM parameter store and AWS Secret manager are both a secure option. However, Secrets manager is more flexible and has more options like password generation. Reference: <https://www.1strategy.com/blog/2019/02/28/aws-parameter-store-vs-aws-secrets-manager/>

#### NEW QUESTION 9

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

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

Which solution will meet these requirements?

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

**Answer: C**

#### Explanation:

The correct answer is C.

A comprehensive and detailed explanation is:

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

? Option B is incorrect because creating a new SCP that has a statement that allows

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

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

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

References:

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

#### NEW QUESTION 10

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

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

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

**Answer: BDE**

#### Explanation:

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

#### NEW QUESTION 10

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

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

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

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

**Answer: C**

#### Explanation:

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

#### NEW QUESTION 15

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

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

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

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

**Answer: AE**

#### Explanation:

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

References:

- ? Introducing mutual TLS authentication for Amazon API Gateway<sup>1</sup>.
- ? Configuring mutual TLS authentication for a REST API<sup>2</sup>.
- ? AWS Private Certificate Authority details<sup>3</sup>.
- ? AWS Certificate Manager Private Certificate Authority updates<sup>4</sup>.

#### NEW QUESTION 20

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

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

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

**Answer: A**

#### Explanation:

The correct answer is A.

A comprehensive and detailed explanation is:

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

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

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

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

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

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

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

References:

? CloudWatch Logs Insights

? Amazon Athena

? AWS X-Ray

? AWS Lambda

? Amazon DynamoDB

### NEW QUESTION 21

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

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

**Answer:** D

#### Explanation:

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

### NEW QUESTION 25

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 26

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

The team uses a developer IAM role to access the environment The role is configured with the Administratoraccess managed policy. The company has created a

new CloudFormationDeployment IAM role that has the following policy.

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

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

- A. Remove the AdministratorAccess polic
- B. Assign the ReadOnlyAccess managed IAM policy to the developer rol
- C. Instruct the developers to use the CloudFormationDeployment role as a CloudFormation service role when the developers deploy new stacks.
- D. Update the trust of CloudFormationDeployment role to allow the developer IAM role to assume the CloudFormationDepoyment role.
- E. Configure the IAM to be to get and pass the CloudFormationDeployment role if cloudformation actions for resources,
- F. Update the trust Of the CloudFormationDepoyment role to anow the cloudformation.amazonaws.com AWS principal to perform the iam:AssumeR01e action
- G. Remove me Administratoraccess polic
- H. Assign the ReadOnly/Access managed IAM policy to the developer role Instruct the developers to assume the CloudFormatondeployment role when the developers new stacks
- I. Add an IAM policy to CloudFormationDeploment to allow cloudformation \* on an Add a policy that allows the iam.PassR01e action for ARN of if iam PassedT0Service equal cloudformation.amazonaws.com

**Answer:** ADF

**Explanation:**

A comprehensive and detailed explanation is:

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

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

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

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

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

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

References:

- ? 1: AWS CloudFormation service roles
- ? 2: How to use trust policies with IAM roles
- ? 3: AWS::IAM::Policy
- ? 4: IAM: Pass an IAM role to a specific AWS service

**NEW QUESTION 27**

A company has 20 service learns Each service team is responsible for its own microservice. Each service team uses a separate AWS account for its microservice and a VPC with the 192 168 0 0/22 CIDR block. The company manages the AWS accounts with AWS Organizations. Each service team hosts its microservice on multiple Amazon EC2 instances behind an Application Load Balancer. The microservices communicate with each

other across the public internet. The company's security team has issued a new guideline that all communication between microservices must use HTTPS over private network connections and cannot traverse the public internet.

A DevOps engineer must implement a solution that fulfills these obligations and minimizes the number of changes for each service team. Which solution will meet these requirements?

- A. Create a new AWS account in AWS Organizations Create a VPC in this account and use AWS Resource Access Manager to share the private subnets of this VPC with the organization Instruct the service teams to launch a ne
- B. Network Load Balancer (NLB) and EC2 instances that use the shared private subnets Use the NLB DNS names for communication between microservices.
- C. Create a Network Load Balancer (NLB) in each of the microservice VPCs Use AWS PrivateLink to create VPC endpoints in each AWS account for the NLBs Create subscriptions to each VPC endpoint in each of the other AWS accounts Use the VPC endpoint DNS names for communication between microservices.
- D. Create a Network Load Balancer (NLB) in each of the microservice VPCs Create VPC peering connections between each of the microservice VPCs Update the route tables for each VPC to use the peering links Use the NLB DNS names for communication between microservices.
- E. Create a new AWS account in AWS Organizations Create a transit gateway in this account and use AWS Resource Access Manager to share the transit gateway with the organizatio
- F. In each of the microservice VPC
- G. create a transit gateway attachment to the shared transit gateway Update the route tables of each VPC to use the transit gateway Create a Network Load Balancer (NLB) in each of the microservice VPCs Use the NLB DNS names for communication between microservices.

**Answer: B**

**Explanation:**

<https://aws.amazon.com/blogs/networking-and-content-delivery/connecting-networks-with-overlapping-ip-ranges/> Private link is the best option because Transit Gateway doesn't support overlapping CIDR ranges.

**NEW QUESTION 28**

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

Which action will accomplish this?

- A. Add a health check to the Auto Scaling group to invoke an AWS Lambda function whenever an instance status is impaired.
- B. Configure the Auto Scaling group to send a notification to an Amazon SNS topic whenever a failed instance launch occurs.
- C. Create an Amazon CloudWatch alarm that invokes an AWS Lambda function when a failed Attachinstances Auto Scaling API call is made.
- D. Create a status check alarm on Amazon EC2 to send a notification to an Amazon SNS topic whenever a status check fail occurs.

**Answer: B**

**Explanation:**

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

**NEW QUESTION 33**

A security review has identified that an AWS CodeBuild project is downloading a database population script from an Amazon S3 bucket using an unauthenticated request. The security team does not allow unauthenticated requests to S3 buckets for this project.

How can this issue be corrected in the MOST secure manner?

- A. Add the bucket name to the AllowedBuckets section of the CodeBuild project setting
- B. Update the build spec to use the AWS CLI to download the database population script.
- C. Modify the S3 bucket settings to enable HTTPS basic authentication and specify a token
- D. Update the build spec to use cURL to pass the token and download the database population script.
- E. Remove unauthenticated access from the S3 bucket with a bucket 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 38**

A company deploys updates to its Amazon API Gateway API several times a week by using an AWS CodePipeline pipeline. As part of the update process the company exports the JavaScript SDK for the API from the API. Gateway console and uploads the SDK to an Amazon S3 bucket

The company has configured an Amazon CloudFront distribution that uses the S3 bucket as an origin Web client then download the SDK by using the CloudFront distribution's endpoint. A DevOps engineer needs to implement a solution to make the new SDK available automatically during new API deployments.

Which solution will meet these requirements?

- A. Create a CodePipeline action immediately after the deployment stage of the AP
- B. Configure the action to invoke an AWS Lambda functio
- C. Configure the Lambda function to download the SDK from API Gateway, upload the SDK to the S3 bucket and create a CloudFront invalidation for the SDK path.

- D. Create a CodePipeline action immediately after the deployment stage of the API Configure the action to use the CodePipeline integration with AP
- E. Gateway to export the SDK to Amazon S3 Create another action that uses the CodePipeline integration with Amazon S3 to invalidate the cache for the SDK path.
- F. Create an Amazon EventBridge rule that reacts to UpdateStage events from aws apigateway Configure the rule to invoke an AWS Lambda function to download the SDK from API Gateway upload the SDK to the S3 bucket and call the CloudFront API to create an invalidation for the SDK path.
- G. Create an Amazon EventBridge rule that reacts to Creat
- H. Deployment events from aws apigatewa
- I. Configure the rule to invoke an AWS Lambda function to download the SDK from AP
- J. Gateway upload the SDK to the S3 bucket and call the S3 API to invalidate the cache for the SDK path.

**Answer: A**

**Explanation:**

This solution would allow the company to automate the process of updating the SDK and making it available to web clients. By adding a CodePipeline action immediately after the deployment stage of the API, the Lambda function will be invoked automatically each time the API is updated. The Lambda function should be able to download the new SDK from API Gateway, upload it to the S3 bucket and also create a CloudFront invalidation for the SDK path so that the latest version of the SDK is available for the web clients. This is the most straight forward solution and it will meet the requirements.

**NEW QUESTION 39**

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

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

Which type of events will match this event pattern?

- A. Failed deploy and build actions across all the pipelines
- B. All rejected or failed approval actions across all the pipelines
- C. All the events across all pipelines
- D. Approval actions across all the pipelines

**Answer: B**

**Explanation:**

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

**NEW QUESTION 43**

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 47

A company has a guideline that every Amazon EC2 instance must be launched from an AMI that the company's security team produces Every month the security team sends an email message with the latest approved AMIs to all the development teams.

The development teams use AWS CloudFormation to deploy their applications. When developers launch a new service they have to search their email for the latest AMIs that the security department sent. A DevOps engineer wants to automate the process that the security team uses to provide the AMI IDs to the development teams.

What is the MOST scalable solution that meets these requirements?

A. Direct the security team to use CloudFormation to create new versions of the AMIs and to list! the AMI ARNs in an encrypted Amazon S3 object as part of the stack's Outputs Section Instruct the developers to use a cross-stack reference to load the encrypted S3 object and obtain the most recent AMI ARNs.

B. Direct the security team to use a CloudFormation stack to create an AWS CodePipeline pipeline that builds new AMIs and places the latest AMI ARNs in an encrypted Amazon S3 object as part of the pipeline output Instruct the developers to use a cross-stack reference within their own CloudFormation template to obtain the S3 object location and the most recent AMI ARNs.

C. Direct the security team to use Amazon EC2 Image Builder to create new AMIs and to place the AMI ARNs as parameters in AWS Systems Manager Parameter Store Instruct the developers to specify a parameter of type SSM in their CloudFormation stack to obtain the most recent AMI ARNs from Parameter Store.

D. Direct the security team to use Amazon EC2 Image Builder to create new AMIs and to create an Amazon Simple Notification Service (Amazon SNS) topic so that every development team can receive notification

E. When the development teams receive a notification instruct them to write an AWS Lambda function that will update their CloudFormation stack with the most recent AMI ARNs.

**Answer: C**

**Explanation:**

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/dynamic-references.html>

#### NEW QUESTION 48

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

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

A. Use AWS OpsWorks Stacks to layer the server nodes of that cluste

B. Create a Chef recipe that populates the content of the 'etc./cluster/nodes config file and restarts the service by using the current members of the laye

C. Assign that recipe to the Configure lifecycle event.

D. Put the file nodes config in version contro

E. Create an AWS CodeDeploy deployment configuration and deployment group based on an Amazon EC2 tag value for thecluster node

F. When adding a new node to the cluster update the file with all tagged instances and make a commit in version contro

G. Deploy the new file and restart the services.

H. Create an Amazon S3 bucket and upload a version of the /etc./cluster/nodes config file Create a crontab script that will poll for that S3 file and download it frequentl

I. Use a process manager such as Monit or system, to restart the cluster services when it detects that the new file was modifie

J. When adding a node to the cluster edit the file's most recent members Upload the new file to the S3 bucket.

K. Create a user data script that lists all members of the current security group of the cluster and automatically updates the /etc/cluster/. nodes confi

L. Tile whenever a new instance is added to the cluster.

**Answer: A**

**Explanation:**

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

OpsWorks Stacks triggers a Setup event. This event runs the associated layer's Setup recipes, which typically handle tasks such as installing and configuring packages

#### NEW QUESTION 51

A company has a single AWS account that runs hundreds of Amazon EC2 instances in a single AWS Region. New EC2 instances are launched and terminated each hour in the account. The account also includes existing EC2 instances that have been running for longer than a week. The company's security policy requires all running EC2 instances to use an EC2 instance profile. If an EC2 instance does not have an instance profile attached, the EC2 instance must use a default instance profile that has no IAM permissions assigned. A DevOps engineer reviews the account and discovers EC2 instances that are running without an instance profile. During the review, the DevOps engineer also observes that new EC2 instances are being launched without an instance profile. Which solution will ensure that an instance profile is attached to all existing and future EC2 instances in the Region?

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

**Answer: B**

#### Explanation:

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

#### NEW QUESTION 53

A company 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 grant
- 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 54

A developer is maintaining a fleet of 50 Amazon EC2 Linux servers. The servers are part of an Amazon EC2 Auto Scaling group, and also use Elastic Load Balancing for load balancing. Occasionally, some application servers are being terminated after failing ELB HTTP health checks. The developer would like to perform a root cause analysis on the issue, but before being able to access application logs, the server is terminated. How can log collection be automated?

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

**Answer: D**

#### Explanation:

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

#### NEW QUESTION 59

A company wants to use AWS CloudFormation for infrastructure deployment. The company has strict tagging and resource requirements and wants to limit the deployment to two Regions. Developers will need to deploy multiple versions of the same application. Which solution ensures resources are deployed in accordance with company policy?

- A. Create AWS Trusted Advisor checks to find and remediate unapproved CloudFormation StackSets.
- B. Create a CloudFormation drift detection operation to find and remediate unapproved CloudFormation StackSets.
- C. Create CloudFormation StackSets with approved CloudFormation templates.
- D. Create AWS Service Catalog products with approved CloudFormation templates.

**Answer: D**

#### Explanation:

service catalog uses stacksets and can enforce tag and restrict resources AWS Customer case with tag enforcement  
<https://aws.amazon.com/ko/blogs/apn/enforce-centralized-tag-compliance-using-aws-service-catalog-amazon-dynamodb-aws-lambda-and-amazon-cloudwatch-events/> And Youtube video showing how to restrict resources per user with portfolio <https://www.youtube.com/watch?v=LzvhTcqyog>

#### NEW QUESTION 60

A business has an application that consists of five independent AWS Lambda functions. The DevOps engineer has built a CI/CD pipeline using AWS CodePipeline and AWS CodeBuild that builds tests packages and deploys each Lambda function in sequence. The pipeline uses an Amazon EventBridge rule to ensure the pipeline starts as quickly as possible after a change is made to the application source code.

After working with the pipeline for a few months the DevOps engineer has noticed the pipeline takes too long to complete. What should the DevOps engineer implement to BEST improve the speed of the pipeline?

- A. Modify the CodeBuild projects within the pipeline to use a compute type with more available network throughput.
- B. Create a custom CodeBuild execution environment that includes a symmetric multiprocessing configuration to run the builds in parallel.
- C. Modify the CodePipeline configuration to run actions for each Lambda function in parallel by specifying the same runorder.
- D. Modify each CodeBuild project to run within a VPC and use dedicated instances to increase throughput.

**Answer: C**

#### Explanation:

<https://docs.aws.amazon.com/codepipeline/latest/userguide/reference-pipeline-structure.html>  
AWS doc: "To specify parallel actions, use the same integer for each action you want to run in parallel. For example, if you want three actions to run in sequence in a stage, you would give the first action the runOrder value of 1, the second action the runOrder value of 2, and the third the runOrder value of 3. However, if you want the second and third actions to run in parallel, you would give the first action the runOrder value of 1 and both the second and third actions the runOrder value of 2."

#### NEW QUESTION 61

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

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

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

**Answer: A**

#### Explanation:

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

#### NEW QUESTION 65

A company is hosting a static website from an Amazon S3 bucket. The website is available to customers at example.com. The company uses an Amazon Route 53 weighted routing policy with a TTL of 1 day. The company has decided to replace the existing static website with a dynamic web application. The dynamic web application uses an Application Load Balancer (ALB) in front of a fleet of Amazon EC2 instances.

On the day of production launch to customers, the company creates an additional Route 53 weighted DNS record entry that points to the ALB with a weight of 255 and a TTL of 1 hour. Two days later, a DevOps engineer notices that the previous static website is displayed sometimes when customers navigate to example.com.

How can the DevOps engineer ensure that the company serves only dynamic content for example.com?

- A. Delete all objects, including previous versions, from the S3 bucket that contains the static website content.
- B. Update the weighted DNS record entry that points to the S3 bucket.
- C. Apply a weight of 0. Specify the domain reset option to propagate changes immediately.
- D. Configure webpage redirect requests on the S3 bucket with a hostname that redirects to the ALB.
- E. Remove the weighted DNS record entry that points to the S3 bucket from the example.com hosted zone.
- F. Wait for DNS propagation to become complete.

**Answer: D**

#### NEW QUESTION 68

A company needs to implement failover for its application. The application includes an Amazon CloudFront distribution and a public Application Load Balancer

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

- A. Create a second CloudFront distribution that has the secondary ALB as the default origin
- B. Create Amazon Route 53 alias records that have a failover policy and Evaluate Target Health set to Yes for both CloudFront distribution
- C. Update the application to use the new record set.
- D. Create a new origin on the distribution for the secondary AL
- E. Create a new origin group
- F. Set the original ALB as the primary origin
- G. Configure the origin group to fail over for HTTP 5xx status code
- H. Update the default behavior to use the origin group.
- I. Create Amazon Route 53 alias records that have a failover policy and Evaluate TargetHealth set to Yes for both ALB
- J. Set the TTL of both records to 0. Update the distribution's origin to use the new record set.
- K. Create a CloudFront function that detects HTTP 5xx status code
- L. Configure the function to return a 307 Temporary Redirect error response to the secondary ALB if the function detects 5xx status code
- M. Update the distribution's default behavior to send origin responses to the function.

**Answer: B**

**Explanation:**

The best solution to implement failover for the application is to use CloudFront origin groups. Origin groups allow CloudFront to automatically switch to a secondary origin when the primary origin is unavailable or returns specific HTTP status codes that indicate a failure<sup>1</sup>. This way, CloudFront can serve the requests from the secondary ALB in the secondary Region without any delay or redirection. To set up origin groups, the DevOps engineer needs to create a new origin on the distribution for the secondary ALB, create a new origin group with the original ALB as the primary origin and the secondary ALB as the secondary origin, and configure the origin group to fail over for HTTP 5xx status codes. Then, the DevOps engineer needs to update the default behavior to use the origin group instead of the single origin<sup>2</sup>.

The other options are not as effective or efficient as the solution in option B. Option A is not suitable because creating a second CloudFront distribution will increase the complexity and cost of the application. Moreover, using Route 53 alias records with a failover policy will introduce some delay in detecting and switching to the secondary CloudFront distribution, which may not meet the zero-second RTO requirement. Option C is not feasible because CloudFront does not support using Route 53 alias records as origins<sup>3</sup>. Option D is not advisable because using a CloudFront function to redirect the requests to the secondary ALB will add an extra round-trip and latency to the failover process, which may also not meet the zero-second RTO requirement.

References:

- ? 1: Optimizing high availability with CloudFront origin failover - Amazon CloudFront
- ? 2: Creating an origin group - Amazon CloudFront
- ? 3: Values That You Specify When You Create or Update a Web Distribution - Amazon CloudFront

**NEW QUESTION 72**

A company is implementing an Amazon Elastic Container Service (Amazon ECS) cluster to run its workload. The company architecture will run multiple ECS services on the cluster. The architecture includes an Application Load Balancer on the front end and uses multiple target groups to route traffic. A DevOps engineer must collect application and access logs. The DevOps engineer then needs to send the logs to an Amazon S3 bucket for near-real-time analysis.

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

- A. Download the Amazon CloudWatch Logs container instance from AWS
- B. Configure this instance as a task
- C. Update the application service definitions to include the logging task.
- D. Install the Amazon CloudWatch Logs agent on the ECS instance
- E. Change the logging driver in the ECS task definition to awslogs.
- F. Use Amazon EventBridge to schedule an AWS Lambda function that will run every 60 seconds and will run the Amazon CloudWatch Logs create-export-task command
- G. Then point the output to the logging S3 bucket.
- H. Activate access logging on the ALB
- I. Then point the ALB directly to the logging S3 bucket.
- J. Activate access logging on the target groups that the ECS services use
- K. Then send the logs directly to the logging S3 bucket.
- L. Create an Amazon Kinesis Data Firehose delivery stream that has a destination of the logging S3 bucket
- M. Then create an Amazon CloudWatch Logs subscription filter for Kinesis Data Firehose.

**Answer: BDF**

**Explanation:**

<https://docs.aws.amazon.com/AmazonECS/latest/developerguide/ecs-logging-monitoring.html>

**NEW QUESTION 77**

A company's security team requires that all external Application Load Balancers (ALBs) and Amazon API Gateway APIs are associated with AWS WAF web ACLs. The company has hundreds of AWS accounts, all of which are included in a single organization in AWS Organizations. The company has configured AWS Config for the organization. During an audit, the company finds some externally facing ALBs that are not associated with AWS WAF web ACLs. Which combination of steps should a DevOps engineer take to prevent future violations? (Choose two.)

- A. Delegate AWS Firewall Manager to a security account.
- B. Delegate Amazon GuardDuty to a security account.
- C. Create an AWS Firewall Manager policy to attach AWS WAF web ACLs to any newly created ALBs and API Gateway APIs.
- D. Create an Amazon GuardDuty policy to attach AWS WAF web ACLs to any newly created ALBs and API Gateway APIs.
- E. Configure an AWS Config managed rule to attach AWS WAF web ACLs to any newly created ALBs and API Gateway APIs.

**Answer: AC**

**Explanation:**

If instead you want to automatically apply the policy to existing in-scope resources, choose Auto remediate any noncompliant resources. This option creates a web ACL in each applicable account within the AWS organization and associates the web ACL with the resources in the accounts. When you choose Auto remediate any noncompliant resources, you can also choose to remove existing web ACL associations from in-scope resources, for the web ACLs that aren't managed by another active Firewall Manager policy. If you choose this option, Firewall Manager first associates the policy's web ACL with the resources, and then removes the prior associations. If a resource has an association with another web ACL that's managed by a different active Firewall Manager policy, this choice doesn't affect that association.

#### NEW QUESTION 79

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 threshold
- 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 threshold
- 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 threshold
- 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 threshold
- 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 83

A company has multiple development teams in different business units that work in a shared single AWS account. All Amazon EC2 resources that are created in the account must include tags that specify who created the resources. The tagging must occur within the first hour of resource creation.

A DevOps engineer needs to add tags to the created resources that include the user ID that created the resource and the cost center ID. The DevOps engineer configures an AWS Lambda function with the cost center mappings to tag the resources. The DevOps engineer also sets up AWS CloudTrail in the AWS account.

An Amazon S3 bucket stores the CloudTrail event logs.

Which solution will meet the tagging requirements?

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

**Answer: D**

#### Explanation:

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

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

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

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

References:

? S3 event notifications

? Server access logging

? Amazon EventBridge rules

? AWS CloudTrail

#### NEW QUESTION 88

A company has configured an Amazon S3 event source on an AWS Lambda function. The company needs the Lambda function to run when a new object is created or an existing object is modified in a particular S3 bucket. The Lambda function will use the S3 bucket name and the S3 object key of the incoming event to read the contents of the created or modified S3 object. The Lambda function will parse the contents and save the parsed contents to an Amazon DynamoDB table.

The Lambda function's execution role has permissions to read from the S3 bucket and to write to the DynamoDB table. During testing, a DevOps engineer discovers that the Lambda

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

Which solution will resolve this problem?

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

**Answer: B**

**Explanation:**

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

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

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

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

References:

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

**NEW QUESTION 89**

A company runs an application on Amazon EC2 instances. The company uses a series of AWS CloudFormation stacks to define the application resources. A developer performs updates by building and testing the application on a laptop and then uploading the build output and CloudFormation stack templates to Amazon S3. The developer's peers review the changes before the developer performs the CloudFormation stack update and installs a new version of the application onto the EC2 instances.

The deployment process is prone to errors and is time-consuming when the developer updates each EC2 instance with the new application. The company wants to automate as much of the application deployment process as possible while retaining a final manual approval step before the modification of the application or resources.

The company already has moved the source code for the application and the CloudFormation templates to AWS CodeCommit. The company also has created an AWS CodeBuild project to build and test the application.

Which combination of steps will meet the company's requirements? (Choose two.)

- A. Create an application group and a deployment group in AWS CodeDeploy
- B. Install the CodeDeploy agent on the EC2 instances.
- C. Create an application revision and a deployment group in AWS CodeDeploy
- D. Create an environment in CodeDeploy
- E. Register the EC2 instances to the CodeDeploy environment.
- F. Use AWS CodePipeline to invoke the CodeBuild job, run the CloudFormation update, and pause for a manual approval step
- G. After approval, start the AWS CodeDeploy deployment.
- H. Use AWS CodePipeline to invoke the CodeBuild job, create CloudFormation change sets for each of the application stacks, and pause for a manual approval step
- I. After approval, run the CloudFormation change sets and start the AWS CodeDeploy deployment.
- J. Use AWS CodePipeline to invoke the CodeBuild job, create CloudFormation change sets for each of the application stacks, and pause for a manual approval step
- K. After approval, start the AWS CodeDeploy deployment.

**Answer: AD**

**Explanation:**

A- <https://docs.aws.amazon.com/codedeploy/latest/userguide/codedeploy-agent.html> D - This option correctly utilizes AWS CodePipeline to invoke the CodeBuild job and create CloudFormation change sets. It adds a manual approval step before executing the change sets and starting the AWS CodeDeploy deployment. This ensures that the deployment process is automated while retaining the final manual approval step.

**NEW QUESTION 92**

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

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

- A. Deploy the application on AWS Elastic Beanstalk
- B. Deploy an Amazon RDS for MySQL DB instance as part of the Elastic Beanstalk configuration.
- C. Deploy the application on AWS Elastic Beanstalk
- D. Deploy a separate Amazon RDS for MySQL DB instance outside of Elastic Beanstalk.
- E. Configure a notification email address that alerts the application team in the AWS Elastic Beanstalk configuration.
- F. Configure an Amazon EventBridge rule to monitor AWS Health event
- G. Use an Amazon Simple Notification Service (Amazon SNS) topic as a target to alert the application team.
- H. Use the immutable deployment method to deploy new application versions.
- I. Use the rolling deployment method to deploy new application versions.

**Answer:** BDE

**Explanation:**

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

? B. Deploy the application on AWS Elastic Beanstalk. Deploy a separate Amazon

RDS for MySQL DB instance outside of Elastic Beanstalk. This approach ensures that the database persists independently of the Elastic Beanstalk environment, which can be torn down and recreated without affecting the database<sup>123</sup>.

? E. Use the immutable deployment method to deploy new application

versions. Immutable deployments provide a zero-downtime deployment method that ensures that if any part of the deployment process fails, the environment is rolled back to the original state automatically<sup>4</sup>.

? D. Configure an Amazon EventBridge rule to monitor AWS Health events. Use an

Amazon Simple Notification Service (Amazon SNS) topic as a target to alert the application team. This setup allows for automated monitoring and alerting of the application team in case of deployment failures or other health events<sup>56</sup>.

References:

? AWS Elastic Beanstalk documentation on deploying Ruby applications<sup>1</sup>.

? AWS documentation on application auto-scaling<sup>7</sup>.

? AWS documentation on automated deployment strategies with automatic rollbacks and alerts<sup>456</sup>.

**NEW QUESTION 93**

A company has deployed an application in a production VPC in a single AWS account. The application is popular and is experiencing heavy usage. The company's security team wants to add additional security, such as AWS WAF, to the application deployment. However, the application's product manager is concerned about cost and does not want to approve the change unless the security team can prove that additional security is necessary.

The security team believes that some of the application's demand might come from users that have IP addresses that are on a deny list. The security team provides the deny list to a DevOps engineer. If any of the IP addresses on the deny list access the application, the security team wants to receive automated notification in near real time so that the security team can document that the application needs additional security. The DevOps engineer creates a VPC flow log for the production VPC.

Which set of additional steps should the DevOps engineer take to meet these requirements MOST cost-effectively?

- A. Create a log group in Amazon CloudWatch Log
- B. Configure the VPC flow log to capture accepted traffic and to send the data to the log grou
- C. Create an Amazon CloudWatch metric filter for IP addresses on the deny lis
- D. Create a CloudWatch alarm with the metric filter as inpu
- E. Set the period to 5 minutes and the datapoints to alarm to 1. Use an Amazon Simple Notification Service (Amazon SNS) topic to send alarm notices to the security team.
- F. Create an Amazon S3 bucket for log file
- G. Configure the VPC flow log to capture all traffic and to send the data to the S3 bucke
- H. Configure Amazon Athena to return all log files in the S3 bucket for IP addresses on the deny lis
- I. Configure Amazon QuickSight to accept data from Athena and to publish the data as a dashboard that the security team can acces
- J. Create a threshold alert of 1 for successful acces
- K. Configure the alert to automatically notify the security team as frequently as possible when the alert threshold is met.
- L. Create an Amazon S3 bucket for log file
- M. Configure the VPC flow log to capture accepted traffic and to send the data to the S3 bucke
- N. Configure an Amazon OpenSearch Service cluster and domain for the log file
- O. Create an AWS Lambda function to retrieve the logs from the S3 bucket, format the logs, and load the logs into the OpenSearch Service cluste
- P. Schedule the Lambda function to run every 5 minute
- Q. Configure an alert and condition in OpenSearch Service to send alerts to the security team through an Amazon Simple Notification Service (Amazon SNS) topic when access from the IP addresses on the deny list is detected.
- R. Create a log group in Amazon CloudWatch Log
- S. Create an Amazon S3 bucket to hold query result
- T. Configure the VPC flow log to capture all traffic and to send the data to the log grou
- . Deploy an Amazon Athena CloudWatch connector in AWS Lambd
- . Connect the connector to the log grou
- . Configure Athena to periodically query for all accepted traffic from the IP addresses on the deny list and to store the results in the S3 bucke
- . Configure an S3 event notification to automatically notify the security team through an Amazon Simple Notification Service (Amazon SNS) topic when new objects are added to the S3 bucket.

**Answer:** A

**NEW QUESTION 97**

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 99

A company needs a strategy for failover and disaster recovery of its data and application. The application uses a MySQL database and Amazon EC2 instances. The company requires a maximum RPO of 2 hours and a maximum RTO of 10 minutes for its data and application at all times. Which combination of deployment strategies will meet these requirements? (Select TWO.)

- A. Create an Amazon Aurora Single-AZ cluster in multiple AWS Regions as the data stor
- B. Use Aurora's automatic recovery capabilities in the event of a disaster.
- C. Create an Amazon Aurora global database in two AWS Regions as the data stor
- D. In the event of a failure, promote the secondary Region to the primary for the applicatio
- E. Update the application to use the Aurora cluster endpoint in the secondary Region.
- F. Create an Amazon Aurora cluster in multiple AWS Regions as the data stor
- G. Use a Network Load Balancer to balance the database traffic in different Regions.
- H. Set up the application in two AWS Region
- I. Use Amazon Route 53 failover routing that points to Application Load Balancers in both Region
- J. Use health checks and Auto Scaling groups in each Region.
- K. Set up the application in two AWS Region
- L. Configure AWS Global Accelerator to point to Application Load Balancers (ALBs) in both Region
- M. Add both ALBs to a single endpoint grou
- N. Use health checks and Auto Scaling groups in each Region.

**Answer: BE**

#### Explanation:

To meet the requirements of failover and disaster recovery, the company should use the following deployment strategies:

? Create an Amazon Aurora global database in two AWS Regions as the data store.

In the event of a failure, promote the secondary Region to the primary for the application. Update the application to use the Aurora cluster endpoint in the secondary Region. This strategy can provide a low RPO and RTO for the data, as Aurora global database replicates data with minimal latency across Regions and allows fast and easy failover<sup>12</sup>. The company can use the Amazon Aurora cluster endpoint to connect to the current primary DB cluster without needing to change any application code<sup>1</sup>.

? Set up the application in two AWS Regions. Configure AWS Global Accelerator to

point to Application Load Balancers (ALBs) in both Regions. Add both ALBs to a single endpoint group. Use health checks and Auto Scaling groups in each Region. This strategy can provide high availability and performance for the application, as AWS Global Accelerator uses the AWS global network to route traffic to the closest healthy endpoint<sup>3</sup>. The company can also use static IP addresses that are assigned by Global Accelerator as a fixed entry point for their application<sup>1</sup>. By using health checks and Auto Scaling groups, the company can ensure that their application can scale up or down based on demand and handle any instance failures<sup>4</sup>.

The other options are incorrect because:

? Creating an Amazon Aurora Single-AZ cluster in multiple AWS Regions as the data store would not provide a fast failover or disaster recovery solution, as the company would need to manually restore data from backups or snapshots in another Region in case of a failure.

? Creating an Amazon Aurora cluster in multiple AWS Regions as the data store and using a Network Load Balancer to balance the database traffic in different Regions would not work, as Network Load Balancers do not support cross-Region routing. Moreover, this strategy would not provide a consistent view of the data across Regions, as Aurora clusters do not replicate data automatically between Regions unless they are part of a global database.

? Setting up the application in two AWS Regions and using Amazon Route 53 failover routing that points to Application Load Balancers in both Regions would not provide a low RTO, as Route 53 failover routing relies on DNS resolution, which can take time to propagate changes across different DNS servers and clients. Moreover, this strategy would not provide deterministic routing, as Route 53 failover routing depends on DNS caching behavior, which can vary depending on different factors.

### NEW QUESTION 101

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 IAM 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 105

A company requires that its internally facing web application be highly available. The architecture is made up of one Amazon EC2 web server instance and one NAT instance that provides outbound internet access for updates and accessing public data.

Which combination of architecture adjustments should the company implement to achieve high availability? (Choose two.)

- A. Add the NAT instance to an EC2 Auto Scaling group that spans multiple Availability Zone
- B. Update the route tables.
- C. Create additional EC2 instances spanning multiple Availability Zone
- D. Add an Application Load Balancer to split the load between them.
- E. Configure an Application Load Balancer in front of the EC2 instance
- F. Configure Amazon CloudWatch alarms to recover the EC2 instance upon host failure.
- G. Replace the NAT instance with a NAT gateway in each Availability Zone
- H. Update the route tables.
- I. Replace the NAT instance with a NAT gateway that spans multiple Availability Zone
- J. Update the route tables.

**Answer:** BD

#### Explanation:

<https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-gateway.html>

#### NEW QUESTION 108

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 performance<sup>1</sup>.

? 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 performance<sup>1</sup>.

? 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 processing<sup>2</sup>.

? 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 processing<sup>2</sup>.

? 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 store<sup>3</sup>.

? 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 needs<sup>3</sup>.

References:

? 1: Batch writes

? 2: Multi-measure records vs. single-measure records

? 3: Storage

#### NEW QUESTION 109

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 main 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 for 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 111

A company wants to migrate its content sharing web application hosted on Amazon EC2 to a serverless architecture. The company currently deploys changes to its application by creating a new Auto Scaling group of EC2 instances and a new Elastic Load Balancer, and then shifting the traffic away using an Amazon Route 53 weighted routing policy.

For its new serverless application, the company is planning to use Amazon API Gateway and AWS Lambda. The company will need to update its deployment processes to work with the new application. It will also need to retain the ability to test new features on a small number of users before rolling the features out to the entire user base.

Which deployment strategy will meet these requirements?

- A. Use AWS CDK to deploy API Gateway and Lambda function
- B. When code needs to be changed, update the AWS CloudFormation stack and deploy the new version of the APIs and Lambda function
- C. Use a Route 53 failover routing policy for the canary release strategy.
- D. Use AWS CloudFormation to deploy API Gateway and Lambda functions using Lambda function version
- E. When code needs to be changed, update the CloudFormation stack with the new Lambda code and update the API versions using a canary release strateg
- F. Promote the new version when testing is complete.
- G. Use AWS Elastic Beanstalk to deploy API Gateway and Lambda function
- H. When code needs to be changed, deploy a new version of the API and Lambda function
- I. Shift traffic gradually using an Elastic Beanstalk blue/green deployment.
- J. Use AWS OpsWorks to deploy API Gateway in the service layer and Lambda functions in a custom laye
- K. When code needs to be changed, use OpsWorks to perform a blue/green deployment and shift traffic gradually.

**Answer:** B

#### Explanation:

<https://docs.aws.amazon.com/serverless-application-model/latest/developerguide/automating-updates-to-serverless-apps.html>

#### NEW QUESTION 112

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

- A. Create an Amazon EventBridge rule to capture task state change
- B. Send the event to Amazon CloudWatch Log
- C. Use CloudWatch Logs Insights to investigate stopped tasks.
- D. Configure tasks to write log data in the embedded metric forma
- E. Store the logs in Amazon CloudWatch Log
- F. Monitor the ContainerInstanceCount metric for changes.
- G. Configure the EC2 instances to store logs in Amazon CloudWatch Log
- H. Create a CloudWatch Contributor Insights rule that uses the EC2 instance log dat
- I. Use the Contributor Insights rule to investigate stopped tasks.
- J. Configure an EC2 Auto Scaling lifecycle hook for the EC2\_INSTANCE\_TERMINATING scale-in even
- K. Write the SystemEventLog file to Amazon S3. Use Amazon Athena to query the log file for errors.

**Answer:** A

#### Explanation:

The best solution to log and review all stopped tasks for errors is to use Amazon EventBridge and Amazon CloudWatch Logs. Amazon EventBridge allows the DevOps engineer to create a rule that matches task state change events from Amazon ECS. The rule can then send the event data to Amazon CloudWatch Logs as the target. Amazon CloudWatch Logs can store and monitor the log data, and also provide CloudWatch Logs Insights, a feature that enables the DevOps engineer to interactively search and analyze the log data. Using CloudWatch Logs Insights, the DevOps engineer can filter and aggregate the log data based on various fields, such as cluster, task, container, and reason. This way, the DevOps engineer can easily identify and investigate the stopped tasks and their errors. The other options are not as effective or efficient as the solution in option A. Option B is not suitable because the embedded metric format is designed for custom metrics, not for logging task state changes. Option C is not feasible because the EC2 instances do not store the task state change events in their logs. Option D is not relevant because the EC2\_INSTANCE\_TERMINATING lifecycle hook is triggered when an EC2 instance is terminated by the Auto Scaling group, not when a task is stopped by Amazon ECS. References:

? : [Creating a CloudWatch Events Rule That Triggers on an Event - Amazon Elastic Container Service](#)

? : [Sending and Receiving Events Between AWS Accounts - Amazon EventBridge](#)

? : [Working with Log Data - Amazon CloudWatch Logs](#)

? : [Analyzing Log Data with CloudWatch Logs Insights - Amazon CloudWatch Logs](#)

? : [Embedded Metric Format - Amazon CloudWatch](#)

? : [Amazon EC2 Auto Scaling Lifecycle Hooks - Amazon EC2 Auto Scaling](#)

#### NEW QUESTION 113

A company is using AWS CodePipeline to deploy an application. According to a new guideline, a member of the company's security team must sign off on any application changes before the changes are deployed into production. The approval must be recorded and retained.

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

- A. Configure CodePipeline to write actions to Amazon CloudWatch Logs.
- B. Configure CodePipeline to write actions to an Amazon S3 bucket at the end of each pipeline stage.
- C. Create an AWS CloudTrail trail to deliver logs to Amazon S3.
- D. Create a CodePipeline custom action to invoke an AWS Lambda function for approva
- E. Create a policy that gives the security team access to manage CodePipeline custom actions.
- F. Create a CodePipeline manual approval action before the deployment ste

G. Create a policy that grants the security team access to approve manual approval stages.

**Answer:** CE

**Explanation:**

To meet the new guideline for application deployment, the company can use a combination of AWS CodePipeline and AWS CloudTrail. A manual approval action in CodePipeline allows the security team to review and approve changes before they are deployed. This action can be configured to pause the pipeline until approval is granted, ensuring that no changes move to production without the necessary sign-off. Additionally, by creating an AWS CloudTrail trail, all actions taken within CodePipeline, including approvals, are recorded and delivered to an Amazon S3 bucket. This provides an audit trail that can be retained for compliance and review purposes.

References:

? AWS CodePipeline's manual approval action provides a way to ensure that a member of the security team can review and approve changes before they are deployed<sup>1</sup>.

? AWS CloudTrail integration with CodePipeline allows for the recording and retention of all pipeline actions, including approvals, which can be stored in Amazon S3 for record-keeping<sup>2</sup>.

**NEW QUESTION 118**

A company recently created a new AWS Control Tower landing zone in a new organization in AWS Organizations. The landing zone must be able to demonstrate compliance with the Center for Internet Security (CIS) Benchmarks for AWS Foundations.

The company's security team wants to use AWS Security Hub to view compliance across all accounts. Only the security team can be allowed to view aggregated Security Hub Findings. In addition, specific users must be able to view findings from their own accounts within the organization. All accounts must be enrolled in Security Hub after the accounts are created.

Which combination of steps will meet these requirements in the MOST automated way? (Select THREE.)

- A. Turn on trusted access for Security Hub in the organization's management account
- B. Create a new security account by using AWS Control Tower. Configure the new security account as the delegated administrator account for Security Hub
- C. In the new security account, provide Security Hub with the CIS Benchmarks for AWS Foundations standards.
- D. Security Hub with the CIS Benchmarks for AWS Foundations standards.
- E. Turn on trusted access for Security Hub in the organization's management account
- F. From the management account, provide Security Hub with the CIS Benchmarks for AWS Foundations standards.
- G. Create an AWS IAM identity Center (AWS Single Sign-On) permission set that includes the required permissions. Use the CreateAccountAssignment API operation to associate the security team users with the permission set and with the delegated security account.
- H. Create an SCP that explicitly denies any user who is not on the security team from accessing Security Hub.
- I. In Security Hub, turn on automatic enablement.
- J. In the organization's management account, create an Amazon EventBridge rule that reacts to the CreateManagedAccount event. Create an AWS Lambda function that uses the Security Hub CreateMembers API operation to add new accounts to Security Hub
- K. Configure the EventBridge rule to invoke the Lambda function.

**Answer:** ACE

**Explanation:**

<https://docs.aws.amazon.com/securityhub/latest/userguide/accounts-orgs-auto-enable.html>

**NEW QUESTION 119**

.....

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

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

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

2nd - Questions and Answers in PDF Format

### **DOP-C02 Practice Exam Features:**

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

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