Amazon Web Services

Exam Questions DOP-C02

AWS Certified DevOps Engineer - Professional

NEW QUESTION 1
A company uses AWS Organizations and AWS Control Tower to manage all the company's AWS accounts. The company uses the Enterprise Support plan.
A DevOps engineer is using Account Factory for Terraform (AFT) to provision new accounts. When new accounts are provisioned, the DevOps engineer notices that the support plan for the new accounts is set to the Basic Support plan. The DevOps engineer needs to implement a solution to provision the new accounts with the Enterprise Support plan.
Which solution will meet these requirements?

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

Answer: D

Explanation:
AWS Organizations is a service that helps to manage multiple AWS accounts. AWS Control Tower is a service that makes it easy to set up and govern secure, compliant multi-account AWS environments. Account Factory for Terraform (AFT) is an AWS Control Tower feature that provisions new accounts using Terraform templates. To provision new accounts with the Enterprise Support plan, the DevOps engineer can set the aft_feature_enterprise_support feature flag to True in the AFT deployment input configuration. This flag enables the Enterprise Support plan for newly provisioned accounts.
https://docs.aws.amazon.com/controltower/latest/userguide/aft-feature-options.html

NEW QUESTION 2
A production account has a requirement that any Amazon EC2 instance that has been logged in to manually must be terminated within 24 hours. All applications in the production account are using Auto Scaling groups with the Amazon CloudWatch Logs agent configured.
How can this process be automated?

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

Answer: D

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

NEW QUESTION 3
A DevOps engineer is deploying a new version of a company's application in an AWS CodeDeploy deployment group associated with its Amazon EC2 instances.
After some time, the deployment fails. The engineer realizes that all the events associated with the specific deployment ID are in a Skipped status and code was not deployed in the instances associated with the deployment group.
What are valid reasons for this failure? (Select TWO.).

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

Answer: AD

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

NEW QUESTION 4
A company runs a workload on Amazon EC2 instances. The company needs a control that requires the use of Instance Metadata Service Version 2 (IMDSv2) on all EC2 instances in the AWS account. If an EC2 instance does not prevent the use of Instance Metadata Service Version 1 (IMDSv1), the EC2 instance must be terminated.
Which solution will meet these requirements?

A. Set up AWS Config in the accoun
B. Use a managed rule to check EC2 instance

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

Answer: B

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

NEW QUESTION 5
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 transactio
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 transactio
D. Use Amazon Athena to query the S3 bucket and to visualize the results In a Pie chart forma
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 instrumentatio
G. Create a new X-Ray subsegment Add an annotation for each processed transactio
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 DynamoD
J. Create a Lambda subscription filter for the log fil
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 6
A company uses AWS Organizations to manage its AWS accounts. The company has a root OU that has a child OU. The root OU has an SCP that allows all actions on all resources. The child OU has an SCP that allows all actions for Amazon DynamoDB and AWS Lambda, and denies all other actions.
The company has an AWS account that is named vendor-data in the child OU. A DevOps engineer has an 1AM user that is attached to the AdministratorAccess 1AM policy in the vendor-data account. The DevOps engineer attempts to launch an Amazon EC2 instance in the vendor-data account but receives an access denied error.
Which change should the DevOps engineer make to launch the EC2 instance in the vendor-data account?

A. Attach the AmazonEC2FullAccess 1AM policy to the 1AM user.
B. Create a new SCP that allows all actions for Amazon EC2. Attach the SCP to the vendor-data account.

C. Update the SCP in the child OU to allow all actions for Amazon EC2.
D. Create a new SCP that allows all actions for Amazon EC2. Attach the SCP to the root OU.

Answer: C

Explanation:
 The correct answer is C. Updating the SCP in the child OU to allow all actions for Amazon EC2 will enable the DevOps engineer to launch the EC2 instance in the vendor-data account. SCPs are applied to OUs and accounts in a hierarchical manner, meaning that the SCPs attached to the parent OU are inherited by the child OU and accounts. Therefore, the SCP in the child OU overrides the SCP in the root OU and denies all actions except for DynamoDB and Lambda. By adding EC2 to the allowed actions in the child OU's SCP, the DevOps engineer can access EC2 resources in the vendor-data account.
Option A is incorrect because attaching the AmazonEC2FullAccess IAM policy to the IAM user will not grant the user access to EC2 resources. IAM policies are evaluated after SCPs, so even if the IAM policy allows EC2 actions, the SCP will still deny them.
Option B is incorrect because creating a new SCP that allows all actions for EC2 and attaching it to the vendor-data account will not work. SCPs are not cumulative, meaning that only one SCP is applied to an account at a time. The SCP attached to the account will be the SCP attached to the OU that contains the account. Therefore, option B will not change the SCP that is applied to the vendor-data account.
Option D is incorrect because creating a new SCP that allows all actions for EC2 and attaching it to the root OU will not work. As explained earlier, the SCP in the child OU overrides the SCP in the root OU and denies all actions except for DynamoDB and Lambda. Therefore, option D will not affect the SCP that is applied to the vendor-data account.

NEW QUESTION 7
A DevOps engineer is working on a data archival project that requires the migration of on- premises data to an Amazon S3 bucket. The DevOps engineer develops a script that incrementally archives on-premises data that is older than 1 month to Amazon S3. Data that is transferred to Amazon S3 is deleted from the on-premises location The script uses the S3 PutObject operation.
During a code review the DevOps engineer notices that the script does not verity whether the data was successfully copied to Amazon S3. The DevOps engineer must update the script to ensure that data is not corrupted during transmission. The script must use MD5 checksums to verify data integrity before the on-premises data is deleted.
Which solutions for the script will meet these requirements'? (Select TWO.)

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

Answer: BD

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

NEW QUESTION 8
A company that runs many workloads on AWS has an Amazon EBS spend that has increased over time. The DevOps team notices there are many unattached EBS volumes. Although there are workloads where volumes are detached, volumes over 14 days old are stale and no longer needed. A DevOps engineer has been tasked with creating automation that deletes unattached EBS volumes that have been unattached for 14 days.
Which solution will accomplish this?

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

Answer: C

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

NEW QUESTION 9
A company is deploying a new application that uses Amazon EC2 instances. The company needs a solution to query application logs and AWS account API activity Which solution will meet these requirements?

A. Use the Amazon CloudWatch agent to send logs from the EC2 instances to Amazon CloudWatch Logs Configure AWS CloudTrail to deliver the API logs to

Amazon S3 Use CloudWatch to query both sets of logs.
B. Use the Amazon CloudWatch agent to send logs from the EC2 instances to Amazon CloudWatch Logs Configure AWS CloudTrail to deliver the API logs to CloudWatch Logs Use CloudWatch Logs Insights to query both sets of logs.
C. Use the Amazon CloudWatch agent to send logs from the EC2 instances to Amazon Kinesis Configure AWS CloudTrail to deliver the API logs to Kinesis Use Kinesis to load the data into Amazon Redshift Use Amazon Redshift to query both sets of logs.
D. Use the Amazon CloudWatch agent to send logs from the EC2 instances to Amazon S3 Use AWS CloudTrail to deliver the API togs to Amazon S3 Use Amazon Athena to query both sets of logs in Amazon S3.

Answer: D

Explanation:
 This solution will meet the requirements because it will use Amazon S3 as a common data lake for both the application logs and the API logs. Amazon S3 is a service that provides scalable, durable, and secure object storage for any type of data. You can use the Amazon CloudWatch agent to send logs from your EC2 instances to S3 buckets, and use AWS CloudTrail to deliver the API logs to S3 buckets as well. You can also use Amazon Athena to query both sets of logs in S3 using standard SQL, without loading or transforming them. Athena is a serverless interactive query service that allows you to analyze data in S3 using a variety of data formats, such as JSON, CSV, Parquet, and ORC.

NEW QUESTION 10
A company uses AWS Secrets Manager to store a set of sensitive API keys that an AWS Lambda function uses. When the Lambda function is invoked, the Lambda function retrieves the API keys and makes an API call to an external service. The Secrets Manager secret is encrypted with the default AWS Key Management Service (AWS KMS) key.
A DevOps engineer needs to update the infrastructure to ensure that only the Lambda function's execution role can access the values in Secrets Manager. The solution must apply the principle of least privilege.
Which combination of steps will meet these requirements? (Select TWO.)

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

Answer: BD

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

NEW QUESTION 11
A company uses AWS Organizations to manage multiple accounts. Information security policies require that all unencrypted Amazon EBS volumes be marked as non-compliant. A DevOps engineer needs to automatically deploy the solution and ensure that this compliance check is always present.
Which solution will accomplish this?

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

Answer: B

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

NEW QUESTION 12
A company is using AWS CodePipeline to automate its release pipeline. AWS CodeDeploy is being used in the pipeline to deploy an application to Amazon Elastic

Container Service (Amazon ECS) using the blue/green deployment model. The company wants to implement scripts to test the green version of the application before shifting traffic. These scripts will complete in 5 minutes or less. If errors are discovered during these tests, the application must be rolled back. Which strategy will meet these requirements?

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

Answer: C

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

NEW QUESTION 13
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 threshol
D. Add an EC2 action to stop the instance when the alarm enters the ALARM state.
E. Configure the Datapoints to Alarm value to be 9 out of 12. Configure the alarm to treat missing data as breaching the 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 threshol
H. Add an AWS Systems Manager action to stop the instance when the alarm enters the ALARM state.

Answer: B

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

NEW QUESTION 14
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 bucket1.
? 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 15
A company uses AWS and has a VPC that contains critical compute infrastructure with predictable traffic patterns. The company has configured VPC flow logs that are published to a log group in Amazon CloudWatch Logs.
The company's DevOps team needs to configure a monitoring solution for the VPC flow logs to identify anomalies in network traffic to the VPC over time. If the monitoring solution detects an anomaly, the company needs the ability to initiate a response to the anomaly.
How should the DevOps team configure the monitoring solution to meet these requirements?

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

Answer: D

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

NEW QUESTION 16
A company uses an organization in AWS Organizations to manage its AWS accounts. The company recently acquired another company that has standalone AWS accounts. The acquiring company's DevOps team needs to consolidate the administration of the AWS accounts for both companies and retain full administrative control of the accounts. The DevOps team also needs to collect and group findings across all the accounts to implement and maintain a security posture.
Which combination of steps should the DevOps team take to meet these requirements? (Select TWO.)

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

Answer: BC

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

NEW QUESTION 17
A company detects unusual login attempts in many of its AWS accounts. A DevOps engineer must implement a solution that sends a notification to the company's security team when multiple failed login attempts occur. The DevOps engineer has already created an Amazon Simple Notification Service (Amazon SNS) topic and has subscribed the security team to the SNS topic.
Which solution will provide the notification with the LEAST operational effort?

A. Configure AWS CloudTrail to send log management events to an Amazon CloudWatch Logs log grou
B. Create a CloudWatch Logs metric filter to match failed ConsoleLogin event
C. Create a CloudWatch alarm that is based on the metric filte
D. Configure an alarm action to send messages to the SNS topic.
E. Configure AWS CloudTrail to send log management events to an Amazon S3 bucke
F. Create an Amazon Athena query that returns a failure if the query finds failed logins in the logs in the S3 bucke
G. Create an Amazon EventBridge rule to periodically run the quer
H. Create a second EventBridge rule to detect when the query fails and to send a message to the SNS topic.
I. Configure AWS CloudTrail to send log data events to an Amazon CloudWatch Logs log grou
J. Create a CloudWatch logs metric filter to match failed Consolel_ogin event
K. Create a CloudWatch alarm that is based on the metric filte
L. Configure an alarm action to send messages to the SNS topic.
M. Configure AWS CloudTrail to send log data events to an Amazon S3 bucke
N. Configure an Amazon S3 event notification for the s3:ObjectCreated event typ
O. Filter the event type by ConsoleLogin failed event
P. Configure the event notification to forward to the SNS topic.

Answer: C

Explanation:
The correct answer is C. Configuring AWS CloudTrail to send log data events to an Amazon CloudWatch Logs log group and creating a CloudWatch logs metric filter to match failed ConsoleLogin events is the simplest and most efficient way to monitor and alert on failed login attempts. Creating a CloudWatch alarm that is based on the metric filter and configuring an alarm action to send messages to the SNS topic will ensure that the security team is notified when multiple failed login attempts occur. This solution requires the least operational effort compared to the other options.
Option A is incorrect because it involves configuring AWS CloudTrail to send log management events instead of log data events. Log management events are used to track changes to CloudTrail configuration, such as creating, updating, or deleting a trail. Log data events are used to track API activity in AWS accounts, such as login attempts. Therefore, option A will not capture the failed ConsoleLogin events.
Option B is incorrect because it involves creating an Amazon Athena query and two Amazon EventBridge rules to monitor and alert on failed login attempts. This is a more complex and costly solution than using CloudWatch logs and alarms. Moreover, option B relies on the query returning a failure, which may not happen if the query is executed successfully but does not find any failed logins.
Option D is incorrect because it involves configuring AWS CloudTrail to send log data events to an Amazon S3 bucket and configuring an Amazon S3 event notification for the s3:ObjectCreated event type. This solution will not work because the s3:ObjectCreated event type does not allow filtering by ConsoleLogin failed events. The event notification will be triggered for any object created in the S3 bucket, regardless of the event type. Therefore, option D will generate a lot of false positives and unnecessary notifications. References:
? AWS CloudTrail Log File Examples
? Creating CloudWatch Alarms for CloudTrail Events: Examples
? Monitoring CloudTrail Log Files with Amazon CloudWatch Logs

NEW QUESTION 18
A company is building a new pipeline by using AWS CodePipeline and AWS CodeBuild in a build account. The pipeline consists of two stages. The first stage is a CodeBuild job to build and package an AWS Lambda function. The second stage consists of deployment actions that operate on two different AWS accounts a development environment account and a production environment account. The deployment stages use the AWS Cloud Format ion action that CodePipeline invokes to deploy the infrastructure that the Lambda function requires.
A DevOps engineer creates the CodePipeline pipeline and configures the pipeline to encrypt build artifacts by using the AWS Key Management Service (AWS KMS) AWS managed key for Amazon S3 (the aws/s3 key). The artifacts are stored in an S3 bucket When the pipeline runs, the Cloud Formation actions fail with an access denied error.
Which combination of actions must the DevOps engineer perform to resolve this error? (Select TWO.)

A. Create an S3 bucket in each AWS account for the artifacts Allow the pipeline to write to the S3 bucket
B. Create a CodePipeline S3 action to copy the artifacts to the S3 bucket in each AWS account Update the CloudFormation actions to reference the artifacts S3 bucket in the production account.
C. Create a customer managed KMS key Configure the KMS key policy to allow the IAM roles used by the CloudFormation action to perform decrypt operations Modify the pipeline to use the customer managed KMS key to encrypt artifacts.
D. Create an AWS managed KMS key Configure the KMS key policy to allow the development account and the production account to perform decrypt operation
E. Modify the pipeline to use the KMS key to encrypt artifacts.
F. In the development account and in the production account create an IAM role for CodePipelin
G. Configure the roles with permissions to perform CloudFormation operations and with permissions to retrieve and decrypt objects from the artifacts S3 bucke
H. In the CodePipeline account configure the CodePipeline CloudFormation action to use the roles.
I. In the development account and in the production account create an IAM role for CodePipeline Configure the roles with permissions to perform CloudFormation operations and with permissions to retrieve and decrypt objects from the artifacts S3bucke
J. In the CodePipelme account modify the artifacts S3 bucket policy to allow the roles access Configure the CodePipeline CloudFormation action to use the roles.

Answer: BE

NEW QUESTION 19
A company uses AWS CloudFormation stacks to deploy updates to its application. The stacks consist of different resources. The resources include AWS Auto Scaling groups, Amazon EC2 instances, Application Load Balancers (ALBs), and other resources that are necessary to launch and maintain independent stacks. Changes to application resources outside of CloudFormation stack updates are not allowed.
The company recently attempted to update the application stack by using the AWS CLI. The stack failed to update and produced the following error message:
"ERROR: both the deployment and the CloudFormation stack rollback failed. The deployment failed because the following resource(s) failed to update: [AutoScalingGroup]."
The stack remains in a status of UPDATE_ROLLBACK_FAILED. * Which solution will resolve this issue?

A. Update the subnet mappings that are configured for the ALB

B. Run the aws cloudformation update-stack-set AWS CLI command.
C. Update the 1AM role by providing the necessary permissions to update the stac
D. Run the aws cloudformation continue-update-rollback AWS CLI command.
E. Submit a request for a quota increase for the number of EC2 instances for the accoun
F. Run the aws cloudformation cancel-update-stack AWS CLI command.
G. Delete the Auto Scaling group resourc
H. Run the aws cloudformation rollback-stack AWS CLI command.

Answer: B

Explanation:
 https://repost.aws/knowledge-center/cloudformation-update-rollback-failed If your stack is stuck in the UPDATE_ROLLBACK_FAILED state after a failed update, then the only actions that you can perform on the stack are the ContinueUpdateRollback or DeleteStack operations.

NEW QUESTION 20
......

Thank You for Trying Our Product

We offer two products:

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

2nd - Questons and Answers in PDF Format

DOP-C02 Practice Exam Features:

* DOP-C02 Questions and Answers Updated Frequently

* DOP-C02 Practice Questions Verified by Expert Senior Certified Staff

* DOP-C02 Most Realistic Questions that Guarantee you a Pass on Your FirstTry

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

100% Actual & Verified — Instant Download, Please Click
Order The DOP-C02 Practice Test Here