DOP-C02$^{Q\&As}$

AWS Certified DevOps Engineer - Professional

Pass Amazon DOP-C02 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.leads4pass.com/dop-c02.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon Official Exam Center

- **Instant Download** After Purchase
- **100% Money Back** Guarantee
- **365 Days** Free Update
- **800,000+** Satisfied Customers

QUESTION 1

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 cluster. Create a Chef recipe that populates the content of the \\'etc./cluster/nodes config file and restarts the service by using the current members of the layer. Assign that recipe to the Configure lifecycle event.

B. Put the file nodes config in version control. Create an AWS CodeDeploy deployment configuration and deployment group based on an Amazon EC2 tag value for thecluster nodes. When adding a new node to the cluster update the file with all tagged instances and make a commit in version control. Deploy the new file and restart the services.

C. Create an Amazon S3 bucket and upload a version of the /etc./cluster/nodes config file Create a crontab script that will poll for that S3 file and download it frequently. Use a process manager such as Monit or system, to restart the cluster services when it detects that the new file was modified. When adding a node to the cluster edit the file\\'s most recent members Upload the new file to the S3 bucket.

D. Create a user data script that lists all members of the current security group of the cluster and automatically updates the /etc/cluster/. nodes config. Tile whenever a new instance is added to the cluster.

Correct Answer: A

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

QUESTION 2

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 R TO.

Which solution will meet these requirements?

A. Create a second CloudFront distribution that has the secondary ALB as the default origin. Create Amazon Route 53 alias records that have a failover policy and Evaluate Target Health set to Yes for both CloudFront distributions. Update the application to use the new record set.

B. Create a new origin on the distribution for the secondary ALB. Create a new origin group. Set the original ALB as the

primary origin. Configure the origin group to fail over for HTTP 5xx status codes. Update the default behavior to use the origin group.

C. Create Amazon Route 53 alias records that have a failover policy and Evaluate Target Health set to Yes for both ALBs. Set the TTL of both records to O. Update the distribution\\'s origin to use the new record set.

D. Create a CloudFront function that detects HTTP 5xx status codes. Configure the function to return a 307 Temporary Redirect error response to the secondary ALB if the function detects 5xx status codes. Update the distribution\\'s default behavior to send origin responses to the function.

Correct Answer: B

To implement failover for the application to the secondary Region so that HTTP GET requests meet the desired RTO, the DevOps engineer should use the following solution: Create a new origin on the distribution for the secondary ALB. A CloudFront origin is the source of the content that CloudFront delivers to viewers. By creating a new origin for the secondary ALB, the DevOps engineer can configure CloudFront to route traffic to the secondary Region when the primary Region is unavailable1 Create a new origin group. Set the original ALB as the primary origin. Configure the origin group to fail over for HTTP 5xx status codes. An origin group is a logical grouping of two origins: a primary origin and a secondary origin. By creating an origin group, the DevOps engineer can specify which origin CloudFront should use as a fallback when the primary origin fails. The DevOps engineer can also define which HTTP status codes should trigger a failover from the primary origin to the secondary origin. By setting the original ALB as the primary origin and configuring the origin group to fail over for HTTP 5xx status codes, the DevOps engineer can ensure that CloudFront will switch to the secondary ALB when the primary ALB returns server errors2 Update the default behavior to use the origin group. A behavior is a set of rules that CloudFront applies when it receives requests for specific URLs or file types. The default behavior applies to all requests that do not match any other behaviors. By updating the default behavior to use the origin group, the DevOps engineer can enable failover routing for all requests that are sent to the distribution3 This solution will meet the requirements because it will automate the failover of the application to the secondary Region with zero-second RTO. When CloudFront receives an HTTP GET request, it will first try to route it to the primary ALB in the primary Region. If the primary ALB is healthy and returns a successful response, CloudFront will deliver it to the viewer. If the primary ALB is unhealthy or returns an HTTP 5xx status code, CloudFront will automatically route the request to the secondary ALB in the secondary Region and deliver its response to the viewer. The other options are not correct because they either do not provide zero-second RTO or do not work as expected. Creating a second CloudFront distribution that has the secondary ALB as the default origin and creating Amazon Route 53 alias records that have a failover policy is not a good option because it will introduce additional latency and complexity to the solution. Route 53 health checks and DNS propagation can take several minutes or longer, which means that viewers might experience delays or errors when accessing the application during a failover event. Creating Amazon Route 53 alias records that have a failover policy and Evaluate Target Health set to Yes for both ALBs and setting the TTL of both records to O is not a valid option because it will not work with CloudFront distributions. Route 53 does not support health checks for alias records that point to CloudFront distributions, so it cannot detect if an ALB behind a distribution is healthy or not. Creating a CloudFront function that detects HTTP 5xx status codes and returns a 307 Temporary Redirect error response to the secondary ALB is not a valid option because it will not provide zero-second RTO. A 307 Temporary Redirect error response tells viewers to retry their requests with a different URL, which means that viewers will have to make an additional request and wait for another response from CloudFront before reaching the secondary ALB.

References:

1: Adding, Editing, and Deleting Origins -Amazon CloudFront

2: Configuring Origin Failover -Amazon CloudFront

3: Creating or Updating a Cache Behavior -Amazon CloudFront

QUESTION 3

A company\\'s application is running on Amazon EC2 instances in an Auto Scaling group. A DevOps engineer needs to

ensure there are at least four application servers running at all times. Whenever an update has to be made to the application, the engineer creates a new AMI with the updated configuration and updates the AWS CloudFormation template with the new AMI ID. After the stack finishes, the engineer manually terminates the old instances one by one, verifying that the new instance is operational before proceeding. The engineer needs to automate this process.

Which action will allow for the LEAST number of manual steps moving forward?

A. Update the CloudFormation template to include the UpdatePolicy attribute with the AutoScalingRollingUpdate policy.

B. Update the CloudFormation template to include the UpdatePolicy attribute with the AutoScalingReplacingUpdate policy.

C. Use an Auto Scaling lifecycle hook to verify that the previous instance is operational before allowing the DevOps engineer\\'s selected instance to terminate.

D. Use an Auto Scaling lifecycle hook to confirm there are at least four running instances before allowing the DevOps engineer\\'s selected instance to terminate.

Correct Answer: B

QUESTION 4

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

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

Which solution Will meet these requirements With the MOST operational efficiency?

A. Set the deployment configuration in CodeDeploy to LambdaAllAtOnce Configure automatic rollbacks on the deployment group Create an Amazon CloudWatch alarm that detects HTTP Bad Gateway errors on API Gateway Configure the deployment group to roll back when the number of alarms meets the alarm threshold

B. Set the deployment configuration in CodeDeploy to LambdaCanary10Percent10Minutes. Configure automatic rollbacks on the deployment group Create an Amazon CloudWatch alarm that detects HTTP Bad Gateway errors on API Gateway Configure the deployment group to roll back when the number of alarms meets the

alarm threshold

C. Set the deployment configuration in CodeDeploy to LambdaAllAtOnce Configure manual rollbacks on the deployment group. Create an Amazon Simple Notification Service (Amazon SNS) topc to send notifications every time a deployrnent fads. Configure the SNS topc to Invoke a new Lambda function that stops the current deployment and starts the most recent successful deployment

D. Set the deployment configuration in CodeDeploy to LambdaCanaryIOPercentIOMinutes Configure manual rollbacks on the deployment group Create a metric filter on an Amazon CloudWatch log group for API Gateway to monitor HTTP Bad Gateway errors. Configure the metric filter to Invoke a new Lambda function that stops the current eployment and starts the most recent successful deployment

Correct Answer: B

Option A is incorrect because setting the deployment configuration to LambdaAllAtOnce means that the new version of the application will be deployed to all Lambda functions at once, affecting all customers. This does not meet the

requirement of affecting the fewest customers possible. Moreover, configuring automatic rollbacks on the deployment group is not operationally efficient, as it requires manual intervention to fix the errors and redeploy the application. Option B is correct because setting the deployment configuration to LambdaCanary10Percent10Minutes means that the new version of the application will be deployed to 10 percent of the Lambda functions first, and then to the remaining 90 percent after 10 minutes. This minimizes the impact of errors on customers, as only 10 percent of them will be affected by a faulty deployment. Configuring automatic rollbacks on the deployment group also meets the requirement of reverting to the most recent stable version of the application when an error is detected. Creating a CloudWatch alarm that detects HTTP Bad Gateway errors on API Gateway is a valid way to monitor the health of the application and trigger a rollback if needed. Option C is incorrect because setting the deployment configuration to LambdaAllAtOnce means that the new version of the application will be deployed to all Lambda functions at once, affecting all customers. This does not meet the requirement of affecting the fewest customers possible. Moreover, configuring manual rollbacks on the deployment group is not operationally efficient, as it requires human intervention to stop the current deployment and start a new one. Creating an SNS topic to send notifications every time a deployment fails is not sufficient to detect errors in the application, as it does not monitor the API Gateway responses. Option D is incorrect because configuring manual rollbacks on the deployment group is not operationally efficient, as it requires human intervention to stop the current deployment and start a new one. Creating a metric filter on a CloudWatch log group for API Gateway to monitor HTTP Bad Gateway errors is a valid way to monitor the health of the application, but invoking a new Lambda function to perform a rollback is unnecessary and complex, as CodeDeploy already provides automatic rollback functionality. References: AWS CodeDeploy Deployment Configurations [AWS CodeDeploy Rollbacks] Amazon CloudWatch Alarms

---

QUESTION 5

A company requires an RPO of 2 hours and an RTO of 10 minutes for its data and application at all times. An application uses a MySQL database and Amazon EC2 web servers. The development team needs a strategy for failover and disaster recovery.

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

A. Create an Amazon Aurora cluster in one Availability Zone across multiple Regions as the data store Use Aurora\\'s automatic recovery capabilities in the event of a disaster

B. Create an Amazon Aurora global database in two Regions as the data store. In the event of a failure promote the secondary Region as the primary for the application.

C. Create an Amazon Aurora multi-master cluster across multiple Regions as the data store. Use a Network Load Balancer to balance the database traffic in different Regions.

D. Set up the application in two Regions and use Amazon Route 53 failover-based routing that points to the Application Load Balancers in both Regions. Use hearth checks to determine the availability in a given Region. Use Auto Scaling groups in each Region to adjust capacity based on demand.

E. Set up the application m two Regions and use a multi-Region Auto Scaling group behind Application Load Balancers to manage the capacity based on demand. In the event of a disaster adjust the Auto Scaling group\\'s desired instance count to increase baseline capacity in the failover Region.

Correct Answer: BD

QUESTION 6

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

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

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

A. Update the default KMS key for Secrets Manager to allow only the Lambda function\\'s execution role to decrypt.

B. Create a KMS customer managed key that trusts Secrets Manager and allows the Lambda function\\'s execution role to decrypt. Update Secrets Manager to use the new customer managed key.

C. Create a KMS customer managed key that trusts Secrets Manager and allows the account\\'s :root principal to decrypt. Update Secrets Manager to use the new customer managed key.

D. Ensure that the Lambda function\\'s execution role has the KMS permissions scoped on the resource level. Configure the permissions so that the KMS key can encrypt the Secrets Manager secret.

E. Remove all KMS permissions from the Lambda function\\'s execution role.

Correct Answer: BD

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.

QUESTION 7

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 stream. Subscribe the log group to the data stream. Configure Amazon Kinesis Data Analytics to detect log anomalies in the data stream. Create an AWS Lambda function to use as the output of the data stream. Configure the Lambda function to write to the default Amazon EventBridge event bus in the event of an anomaly finding.

B. Create an Amazon Kinesis Data Firehose delivery stream that delivers events to an Amazon S3 bucket. Subscribe the log group to the delivery stream. Configure Amazon Lookout for Metrics to monitor the data in the S3 bucket for anomalies. Create an AWS Lambda function to run in response to Lookout for Metrics anomaly findings. Configure the Lambda function to publish to the default Amazon EventBridge event bus.

C. Create an AWS Lambda function to detect anomalies. Configure the Lambda function to publish an event to the default Amazon EventBridge event bus if the Lambda function detects an anomaly. Subscribe the Lambda function to the log group.

D. Create an Amazon Kinesis data stream. Subscribe the log group to the data stream. Create an AWS Lambda function to detect log anomalies. Configure the Lambda function to write to the default Amazon EventBridge event bus if the Lambda function detects an anomaly. Set the Lambda function as the processor for the data stream.

Correct Answer: D

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.

QUESTION 8

A root owner is trying to create an IAM user of the various departments. The owner has created groups for each department, but wants to still delineate the user based on the sub division level. E.g. The two users from different sub departments should be identified separately and have separate permissions. How can the root owner configure this?

A. Create a hierarchy of the IAM users which are separated based on the department

B. Create a nested group

C. Use the paths to separate the users of the same group

D. It is not possible to delineate within a group

Correct Answer: C

The path functionality within an IAM group and user allows them to delineate by further levels. In this case the user needs to use the path with each user or group so that the ARN of the user will look similar to:
arn:aws:iam::123456789012:user/division_abc/subdivision_xyz/user1
arn:aws:iam::123456789012:user/division_abc/subdivision_xyz/user2

Reference: http://docs.aws.amazon.com/IAM/latest/UserGuide/Using_Identifiers.html#Identifiers_ARNs

QUESTION 9

An ecommerce company is receiving reports that its order history page is experiencing delays in reflecting the processing status of orders. The order processing system consists of an AWS Lambda function that uses reserved concurrency. The Lambda function processes order messages from an Amazon Simple Queue Service (Amazon SQS)

queue and inserts processed orders into an Amazon DynamoDB table. The DynamoDB table has auto scaling enabled for read and write capacity.

Which actions should a DevOps engineer take to resolve this delay? (Choose two.)

A. Check the ApproximateAgeOfOldestMessage metric for the SQS queue. Increase the Lambda function concurrency limit.

B. Check the ApproximateAgeOfOldestMessage metnc for the SQS queue Configure a redrive policy on the SQS queue.

C. Check the NumberOfMessagesSent metric for the SQS queue. Increase the SQS queue visibility timeout.

D. Check the WriteThrottleEvents metric for the DynamoDB table. Increase the maximum write capacity units (WCUs) for the table\\'s scaling policy.

E. Check the Throttles metric for the Lambda function. Increase the Lambda function timeout.

Correct Answer: AD

A: If the ApproximateAgeOfOldestMessages indicate that orders are remaining in the SQS queue for longer than expected, the reserved concurrency limit may be set too small to keep up with the number of orders entering the queue and is being throttled. D: The DynamoDB table is using Auto Scaling. With Auto Scaling, you create a scaling policy that specifies whether you want to scale read capacity or write capacity (or both), and the minimum and maximum provisioned capacity unit settings for the table. The ThottledWriteRequests metric will indicate if there is a throttling issue on the DynamoDB table, which can be resolved by increasing the maximum write capacity units for the table\\'s Auto Scaling policy. https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/AutoScaling.html

---

QUESTION 10

A DevOps engineer is building a multistage pipeline with AWS CodePipeline to build, verify, stage, test, and deploy an application. A manual approval stage is required between the test stage and the deploy stage. The development team uses a custom chat tool with webhook support that requires near-real-time notifications.

How should the DevOps engineer configure status updates for pipeline activity and approval requests to post to the chat tool?

A. Create an Amazon CloudWatch Logs subscription that filters on CodePipeline Pipeline Execution State Change. Publish subscription events to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the chat webhook URL to the SNS topic, and complete the subscription validation.

B. Create an AWS Lambda function that is invoked by AWS CloudTrail events. When a CodePipeline Pipeline Execution State Change event is detected, send the event details to the chat webhook URL.

C. Create an Amazon EventBridge rule that filters on CodePipeline Pipeline Execution State Change. Publish the events to an Amazon Simple Notification Service (Amazon SNS) topic. Create an AWS Lambda function that sends event details to the chat webhook URL. Subscribe the function to the SNS topic.

D. Modify the pipeline code to send the event details to the chat webhook URL at the end of each stage. Parameterize the URL so that each pipeline can send to a different URL based on the pipeline environment.

Correct Answer: C

https://aws.amazon.com/premiumsupport/knowledge-center/sns-lambda-webhooks-chime-slack-teams/

QUESTION 11

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 notifications. 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.

Correct Answer: C

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

QUESTION 12

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 Logs. Configure the VPC flow log to capture accepted traffic and to send the data to the log group. Create an Amazon CloudWatch metric filter for IP addresses on the deny list. Create a CloudWatch alarm with the metric filter as input. 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.

B. Create an Amazon S3 bucket for log files. Configure the VPC flow log to capture all traffic and to send the data to the S3 bucket. Configure Amazon Athena to return all log files in the S3 bucket for IP addresses on the deny list. Configure Amazon QuickSight to accept data from Athena and to publish the data as a dashboard that the security team can access. Create a threshold alert of 1 for successful access. Configure the alert to automatically notify the security team as frequently as possible when the alert threshold is met.

C. Create an Amazon S3 bucket for log files. Configure the VPC flow log to capture accepted traffic and to send the data to the S3 bucket. Configure an Amazon OpenSearch Service cluster and domain for the log files. Create an AWS Lambda function to retrieve the logs from the S3 bucket, format the logs, and load the logs into the OpenSearch Service cluster. Schedule the Lambda function to run every 5 minutes. 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.

D. Create a log group in Amazon CloudWatch Logs. Create an Amazon S3 bucket to hold query results. Configure the VPC flow log to capture all traffic and to send the data to the log group. Deploy an Amazon Athena CloudWatch connector in AWS Lambda. Connect the connector to the log group. 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 bucket. 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.

Correct Answer: A

QUESTION 13

A company manages AWS accounts for application teams in AWS Control Tower. Individual application teams are responsible for securing their respective AWS accounts.

A DevOps engineer needs to enable Amazon GuardDuty for all AWS accounts in which the application teams have not already enabled GuardDuty. The DevOps engineer is using AWS CloudFormation StackSets from the AWS Control Tower management account.

How should the DevOps engineer configure the CloudFormation template to prevent failure during the StackSets deployment?

A. Create a CloudFormation custom resource that invokes an AWS Lambda function. Configure the Lambda function to conditionally enable GuardDuty if GuardDuty is not already enabled in the accounts.

B. Use the Conditions section of the CloudFormation template to enable GuardDuty in accounts where GuardDuty is not already enabled.

C. Use the CloudFormation Fn. GetAtt intrinsic function to check whether GuardDuty is already enabled If GuardDuty is not already enabled use the Resources section of the CloudFormation template to enable GuardDuty.

D. Manually discover the list of AWS account IDs where GuardDuty is not enabled Use the CloudFormation Fn: ImportValue intrinsic function to import the list of account IDs into the CloudFormation template to skip deployment for the listed AWS accounts.

Correct Answer: A

This solution will meet the requirements because it will use a CloudFormation custom resource to execute custom logic during the stack set operation. A custom resource is a resource that you define in your template and that is associated with an AWS Lambda function. The Lambda function runs whenever the custom resource is created, updated, or deleted, and can perform any actions that are supported by the AWS SDK. In this case, the Lambda function can use the GuardDuty API to check whether GuardDuty is already enabled in each target account, and if not, enable it. This way, the DevOps engineer can avoid deploying the stack set to accounts that already have GuardDuty enabled, and

prevent failure during the deployment.

QUESTION 14

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 changes. Send the event to Amazon CloudWatch Logs. Use CloudWatch Logs Insights to investigate stopped tasks.

B. Configure tasks to write log data in the embedded metric format. Store the logs in Amazon CloudWatch Logs. Monitor the ContainerInstanceCount metric for changes.

C. Configure the EC2 instances to store logs in Amazon CloudWatch Logs. Create a CloudWatch Contributor Insights rule that uses the EC2 instance log data. Use the Contributor Insights rule to investigate stopped tasks.

D. Configure an EC2 Auto Scaling lifecycle hook for the EC2_INSTANCE_TERMINATING scale-in event. Write the SystemEventLog file to Amazon S3. Use Amazon Athena to query the log file for errors.

Correct Answer: A

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

QUESTION 15

A company is using AWS Organizations and wants to implement a governance strategy with the following requirements:

AWS resource access is restricted to the same two Regions for all accounts.

AWS services are limited to a specific group of authorized services for all accounts.

Authentication is provided by Active Directory.

Access permissions are organized by job function and are identical in each account.

Which solution will meet these requirements?

A. Establish an organizational unit (OU) with group policies in the master account to restrict Regions and authorized services. Use AWS CloudFormation StackSets to provision roles with permissions for each job function, including an IAM trust policy for IAM identity provider authentication in each account.

B. Establish a permission boundary in the master account to restrict Regions and authorized services. Use AWS CloudFormation StackSet to provision roles with permissions for each job function, including an IAM trust policy for IAM identity provider authentication in each account.

C. Establish a service control in the master account to restrict Regions and authorized services. Use AWS Resource Access Manager to share master account roles with permissions for each job function, including AWS SSO for authentication in each account.

D. Establish a service control in the master account to restrict Regions and authorized services. Use CloudFormation StackSet to provision roles with permissions for each job function, including an IAM trust policy for IAM identity provider authentication in each account.

Correct Answer: A