

Amazon

Exam Questions AWS-Certified-DevOps-Engineer-Professional

Amazon AWS Certified DevOps Engineer Professional



NEW QUESTION 1

A company uses Amazon S3 to store proprietary information. The development team creates buckets for new projects on a daily basis. The security team wants to ensure that all existing and future buckets have encryption logging and versioning enabled. Additionally, no buckets should ever be publicly read or write accessible.

What should a DevOps engineer do to meet these requirements?

- A. Enable AWS CloudTrail and configure automatic remediation using AWS Lambda.
- B. Enable AWS Config rules and configure automatic remediation using AWS Systems Manager documents.
- C. Enable AWS Trusted Advisor and configure automatic remediation using Amazon EventBridge.
- D. Enable AWS Systems Manager and configure automatic remediation using Systems Manager documents.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/mt/aws-config-auto-remediation-s3-compliance/> <https://aws.amazon.com/blogs/aws/aws-config-rules-dynamic-compliance-checking-for-cloud-resources/>

NEW QUESTION 2

A company wants to use AWS development tools to replace its current bash deployment scripts. The company currently deploys a LAMP application to a group of Amazon EC2 instances behind an Application Load Balancer (ALB). During the deployments, the company unit tests the committed application, stops and starts services, unregisters and re-registers instances with the load balancer, and updates file permissions. The company wants to maintain the same deployment functionality through the shift to using AWS services.

Which solution will meet these requirements?

- A. Use AWS CodeBuild to test the applicatio
- B. Use bash scripts invoked by AWS CodeDeploy's appspec.yml file to restart services, and deregister and register instances with the AL
- C. Use the appspec.yml file to update file permissions without a custom script.
- D. Use AWS CodePipeline to move the application from the AWS CodeCommit repository to AWS CodeDeplo
- E. Use CodeDeploy's deployment group to test the application, unregister and re-register instances with the AL
- F. and restart service
- G. Use the appspec.yml file to update file permissions without a custom script.
- H. Use AWS CodePipeline to move the application source code from the AWS CodeCommit repository to AWS CodeDeplo
- I. Use CodeDeploy to test the applicatio
- J. Use CodeDeploy's appspec.yml file to restart services and update permissions without a custom scrip
- K. Use AWS CodeBuild to unregister and re-register instances with the ALB.
- L. Use AWS CodePipeline to trigger AWS CodeBuild to test the applicatio
- M. Use bash scripts invoked by AWS CodeDeploy's appspec.yml file to restart service
- N. Unregister and re-register the instances in the AWS CodeDeploy deployment group with the AL
- O. Update the appspec.yml file to update file permissions without a custom script.

Answer: D

Explanation:

<https://aws.amazon.com/blogs/devops/how-to-test-and-debug-aws-codedeploy-locally-before-you-ship-your-code/#:~:text=You%20can%20test%20application%20code,local%20server%20or%20EC2%20instance.>

NEW QUESTION 3

A company is developing an application that will generate log events. The log events consist of five distinct metrics every one tenth of a second and produce a large amount of data. The company needs to configure the application to write the logs to Amazon Time stream. The company will configure a daily query against the Timestream table.

Which combination of steps will meet these requirements with the FASTEST query performance? (Select THREE.)

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

Answer: ADF

Explanation:

A comprehensive and detailed explanation is:

? Option A is correct because using batch writes to write multiple log events in a single write operation is a recommended practice for optimizing the performance and cost of data ingestion in Timestream. Batch writes can reduce the number of network round trips and API calls, and can also take advantage of parallel processing by Timestream. Batch writes can also improve the compression ratio of data in the memory store and the magnetic store, which can reduce the storage costs and improve the query performance1.

? Option B is incorrect because writing each log event as a single write operation is not a recommended practice for optimizing the performance and cost of data ingestion in Timestream. Writing each log event as a single write operation would increase the number of network round trips and API calls, and would also reduce the compression ratio of data in the memory store and the magnetic store. This would increase the storage costs and degrade the query performance1.

? Option C is incorrect because treating each log as a single-measure record is not a recommended practice for optimizing the query performance in Timestream. Treating each log as a single-measure record would result in creating multiple records for each timestamp, which would increase the storage size and the query latency. Moreover, treating each log as a single-measure record would require using joins to query multiple measures for the same timestamp, which would add complexity and overhead to the query processing2.

? Option D is correct because treating each log as a multi-measure record is a recommended practice for optimizing the query performance in Timestream. Treating each log as a multi-measure record would result in creating a single record for each timestamp, which would reduce the storage size and the query latency. Moreover, treating each log as a multi-measure record would allow querying multiple measures for the same timestamp without using joins, which would

simplify and speed up the query processing².

? Option E is incorrect because configuring the memory store retention period to be longer than the magnetic store retention period is not a valid option in Timestream. The memory store retention period must always be shorter than or equal to the magnetic store retention period. This ensures that data is moved from the memory store to the magnetic store before it expires out of the memory store³.

? Option F is correct because configuring the memory store retention period to be shorter than the magnetic store retention period is a valid option in Timestream. The memory store retention period determines how long data is kept in the memory store, which is optimized for fast point-in-time queries. The magnetic store retention period determines how long data is kept in the magnetic store, which is optimized for fast analytical queries. By configuring these retention periods appropriately, you can balance your storage costs and query performance according to your application needs³.

References:

? 1: Batch writes

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

? 3: Storage

NEW QUESTION 4

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 5

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 group
- B. Create a CloudWatch Logs metric filter to match failed ConsoleLogin event
- C. Create a CloudWatch alarm that is based on the metric filter
- 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 bucket
- F. Create an Amazon Athena query that returns a failure if the query finds failed logins in the logs in the S3 bucket
- G. Create an Amazon EventBridge rule to periodically run the query
- 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 group
- J. Create a CloudWatch logs metric filter to match failed ConsoleLogin event
- K. Create a CloudWatch alarm that is based on the metric filter
- 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 bucket
- N. Configure an Amazon S3 event notification for the s3:ObjectCreated event type
- 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 6

A development team uses AWS CodeCommit, AWS CodePipeline, and AWS CodeBuild to develop and deploy an application. Changes to the code are submitted

by pull requests. The development team reviews and merges the pull requests, and then the pipeline builds and tests the application.

Over time, the number of pull requests has increased. The pipeline is frequently blocked because of failing tests. To prevent this blockage, the development team wants to run the unit and integration tests on each pull request before it is merged.

Which solution will meet these requirements?

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

Answer: B

Explanation:

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

NEW QUESTION 7

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 organization
- 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 organization
- 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 accounts
- 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 accounts
- J. Enable all features for the organization
- 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 accounts
- 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 8

A company's DevOps engineer is working in a multi-account environment. The company uses AWS Transit Gateway to route all outbound traffic through a network operations account. In the network operations account all account traffic passes through a firewall appliance for inspection before the traffic goes to an internet gateway.

The firewall appliance sends logs to Amazon CloudWatch Logs and includes event

severities of CRITICAL, HIGH, MEDIUM, LOW, and INFO. The security team wants to receive an alert if any CRITICAL events occur.

What should the DevOps engineer do to meet these requirements?

- A. Create an Amazon CloudWatch Synthetic canary to monitor the firewall status
- B. If the firewall reaches a CRITICAL state or logs a CRITICAL event use a CloudWatch alarm to publish a notification to an Amazon Simple Notification Service (Amazon SNS) topic Subscribe the security team's email address to the topic.
- C. Create an Amazon CloudWatch metric filter by using a search for CRITICAL events Publish a custom metric for the findings
- D. Use a CloudWatch alarm based on the custom metric to publish a notification to an Amazon Simple Notification Service (Amazon SNS) topic
- E. Subscribe the security team's email address to the topic.
- F. Enable Amazon GuardDuty in the network operations account

- G. Configure GuardDuty to monitor flow logs Create an Amazon EventBridge event rule that is invoked by GuardDuty events that are CRITICAL Define an Amazon Simple Notification Service (Amazon SNS) topic as a target Subscribe the security team's email address to the topic.
- H. Use AWS Firewall Manager to apply consistent policies across all account
- I. Create an Amazon EventBridge event rule that is invoked by Firewall Manager events that are CRITICAL Define an Amazon Simple Notification Service (Amazon SNS) topic as a target Subscribe the security team's email address to the topic.

Answer: B

Explanation:

"The firewall appliance sends logs to Amazon CloudWatch Logs and includes event severities of CRITICAL, HIGH, MEDIUM, LOW, and INFO"

NEW QUESTION 9

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 store
- 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 store
- D. In the event of a failure, promote the secondary Region to the primary for the application
- 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 store
- G. Use a Network Load Balancer to balance the database traffic in different Regions.
- H. Set up the application in two AWS Regions
- I. Use Amazon Route 53 failover routing that points to Application Load Balancers in both Regions
- J. Use health checks and Auto Scaling groups in each Region.
- K. Set up the application in two AWS Regions
- L. Configure AWS Global Accelerator to point to Application Load Balancers (ALBs) in both Regions
- M. Add both ALBs to a single endpoint group
- 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 failover¹². The company can use the Amazon Aurora cluster endpoint to connect to the current primary DB cluster without needing to change any application code¹.

? 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 endpoint³. The company can also use static IP addresses that are assigned by Global Accelerator as a fixed entry point for their application¹. 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 failures⁴.

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 10

A company has a legacy application A DevOps engineer needs to automate the process of building the deployable artifact for the legacy application. The solution must store the deployable artifact in an existing Amazon S3 bucket for future deployments to reference

Which solution will meet these requirements in the MOST operationally efficient way?

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

Answer: A

Explanation:

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

which is beneficial for version control and archival purposes.

NEW QUESTION 10

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
- B. Subscribe the log group to the data stream
- C. Configure Amazon Kinesis Data Analytics to detect log anomalies in the data stream
- D. Create an AWS Lambda function to use as the output of the data stream
- 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 bucket
- G. Subscribe the log group to the delivery stream
- H. Configure Amazon Lookout for Metrics to monitor the data in the S3 bucket for anomalies
- 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 anomalies
- L. Configure the Lambda function to publish an event to the default Amazon EventBridge event bus if the Lambda function detects an anomaly
- M. Subscribe the Lambda function to the log group.
- N. Create an Amazon Kinesis data stream
- O. Subscribe the log group to the data stream
- P. Create an AWS Lambda function to detect anomalies
- Q. Configure the Lambda function to write to the default Amazon EventBridge event bus if the Lambda function detects an anomaly
- 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 11

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

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

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 12

A company has developed a serverless web application that is hosted on AWS. The application consists of Amazon S3, Amazon API Gateway, several AWS Lambda functions, and an Amazon RDS for MySQL database. The company is using AWS CodeCommit to store the source code. The source code is a combination of AWS Serverless Application Model (AWS SAM) templates and Python code.

A security audit and penetration test reveal that user names and passwords for authentication to the database are hardcoded within CodeCommit repositories. A DevOps engineer must implement a solution to automatically detect and prevent hardcoded secrets.

What is the MOST secure solution that meets these requirements?

- A. Enable Amazon CodeGuru Profile
- B. Decorate the handler function with `@with_lambda_profiler()`. Manually review the recommendation report
- C. Write the secret to AWS Systems Manager Parameter Store as a secure string
- D. Update the SAM templates and the Python code to pull the secret from Parameter Store.
- E. Associate the CodeCommit repository with Amazon CodeGuru Reviewer

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

Answer: B

Explanation:

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

NEW QUESTION 16

A Company uses AWS CodeCommit for source code control. Developers apply their changes to various feature branches and create pull requests to move those changes to the main branch when the changes are ready for production.

The developers should not be able to push changes directly to the main branch. The company applied the `AWSCodeCommitPowerUser` managed policy to the developers' IAM role, and now these developers can push changes to the main branch directly on every repository in the AWS account.

What should the company do to restrict the developers' ability to push changes to the main branch directly?

- A. Create an additional policy to include a Deny rule for the `GitPush` and `PutFile` action
- B. Include a restriction for the specific repositories in the policy repositories in the policy statement with a condition that references the main branch.
- C. Remove the IAM policy, and add an `AWSCodeCommitReadOnly` managed polic
- D. Add an Allow rule for the `GitPush` and `PutFile` actions for the specific repositories in the policy statement with a condition that references the mam branch.
- E. Modify the IAM policy Include a Deny rule for the `GitPush` and `PutFile` actions for the specific repositories in the policy statement with a condition that references the main branch.
- F. Create an additional policy to include an Allow rule for the `GitPush` and `PutFile` action
- G. Include a restriction for the specific repositories in the policy statement with a condition that references the feature branches.

Answer: A

Explanation:

By default, the `AWSCodeCommitPowerUser` managed policy allows users to push changes to any branch in any repository in the AWS account. To restrict the developers' ability to push changes to the main branch directly, an additional policy is needed that explicitly denies these actions for the main branch.

The Deny rule should be included in a policy statement that targets the specific repositories and includes a condition that references the main branch. The policy statement should look something like this:

```
{
  "Effect": "Deny", "Action": [ "codecommit:GitPush", "codecommit:PutFile"
],
  "Resource": "arn:aws:codecommit:<region>:<account-id>:<repository-name>", "Condition": {
    "StringEqualsIfExists": { "codecommit:References": [ "refs/heads/main"
]
}
}
```

NEW QUESTION 18

A company has an application that runs on Amazon EC2 instances that are in an Auto Scaling group. When the application starts up. the application needs to process data from an Amazon S3 bucket before the application can start to serve requests.

The size of the data that is stored in the S3 bucket is growing. When the Auto Scaling group adds new instances, the application now takes several minutes to download and process the data before the application can serve requests. The company must reduce the time that elapses before new EC2 instances are ready to serve requests.

Which solution is the MOST cost-effective way to reduce the application startup time?

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

Answer: A

Explanation:

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

NEW QUESTION 22

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

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

Answer: A

Explanation:

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

NEW QUESTION 27

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 28

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 file
- 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 file
- 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 29

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 31

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 "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 34

A company has an application that includes AWS Lambda functions. The Lambda functions run Python code that is stored in an AWS CodeCommit repository. The company has recently experienced failures in the production environment because of an error in the Python code. An engineer has written unit tests for the Lambda functions to help avoid releasing any future defects into the production environment.

The company's DevOps team needs to implement a solution to integrate the unit tests into an existing AWS CodePipeline pipeline. The solution must produce reports about the unit tests for the company to view.

Which solution will meet these requirements?

- A. Associate the CodeCommit repository with Amazon CodeGuru Reviewer
- B. Create a new AWS CodeBuild project
- C. In the CodePipeline pipeline, configure a test stage that uses the new CodeBuild project
- D. Create a buildspec.yml file in the CodeCommit repository
- E. In the buildspec.yml file, define the actions to run a CodeGuru review.
- F. Create a new AWS CodeBuild project
- G. In the CodePipeline pipeline, configure a test stage that uses the new CodeBuild project
- H. Create a CodeBuild report group
- I. Create a buildspec.yml file in the CodeCommit repository
- J. In the buildspec.yml file, define the actions to run the unit tests with an output of JUNITXML in the build phase section. Configure the test reports to be uploaded to the new CodeBuild report group.
- K. Create a new AWS CodeArtifact repository
- L. Create a new AWS CodeBuild project
- M. In the CodePipeline pipeline, configure a test stage that uses the new CodeBuild project
- N. Create an appspec.yml file in the original CodeCommit repository
- O. In the appspec.yml file, define the actions to run the unit tests with an output of CUCUMBERJSON in the build phase section
- P. Configure the test reports to be sent to the new CodeArtifact repository.
- Q. Create a new AWS CodeBuild project
- R. In the CodePipeline pipeline, configure a test stage that uses the new CodeBuild project
- S. Create a new Amazon S3 bucket
- T. Create a buildspec.yml file in the CodeCommit repository
- . In the buildspec.yml file, define the actions to run the unit tests with an output of HTML in the phases section
- . In the reports section, upload the test reports to the S3 bucket.

Answer: B

Explanation:

The correct answer is B. Creating a new AWS CodeBuild project and configuring a test stage in the AWS CodePipeline pipeline that uses the new CodeBuild project is the best way to integrate the unit tests into the existing pipeline. Creating a CodeBuild report group and uploading the test reports to the new CodeBuild report group will produce reports about the unit tests for the company to view. Using JUNITXML as the output format for the unit tests is supported by CodeBuild and will generate a valid report. Option A is incorrect because Amazon CodeGuru Reviewer is a service that provides automated code reviews and recommendations for improving code quality and performance. It is not a tool for running unit tests or producing test reports. Therefore, option A will not meet the requirements.

Option C is incorrect because AWS CodeArtifact is a service that provides secure, scalable, and cost-effective artifact management for software development. It is not a tool for running unit tests or producing test reports. Moreover, option C uses CUCUMBERJSON as the output format for the unit tests, which is not supported by CodeBuild and will not generate a valid report.

Option D is incorrect because uploading the test reports to an Amazon S3 bucket is not the best way to produce reports about the unit tests for the company to view. CodeBuild has a built-in feature to create and manage test reports, which is more convenient and efficient than using S3. Furthermore, option D uses HTML as the output format for the unit tests, which is not supported by CodeBuild and will not generate a valid report.

NEW QUESTION 37

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

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

- A. Create an Amazon EventBridge rule for the CloudTrail StopLogging event
- B. Create an AWS Lambda function that uses the AWS SDK to call StartLogging on the ARN of the resource in which StopLogging was called
- C. Add the Lambda function ARN as a target to the EventBridge rule.
- D. Deploy the AWS-managed CloudTrail-enabled AWS Config rule set with a periodic interval to 1 hour

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

Answer: A

Explanation:

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

NEW QUESTION 42

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

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

Which solution will satisfy these requirements?

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

Answer: C

Explanation:

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

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

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

NEW QUESTION 46

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

Which solution will accomplish this?

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

Answer: B

Explanation:

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

NEW QUESTION 51

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

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

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 53

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 54

A company is performing vulnerability scanning for all Amazon EC2 instances across many accounts. The accounts are in an organization in AWS Organizations. Each account's VPCs are attached to a shared transit gateway. The VPCs send traffic to the internet through a central egress VPC. The company has enabled Amazon Inspector in a delegated administrator account and has enabled scanning for all member accounts. A DevOps engineer discovers that some EC2 instances are listed in the "not scanning" tab in Amazon Inspector. Which combination of actions should the DevOps engineer take to resolve this issue? (Choose three.)

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

Answer: ABE

Explanation:

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

NEW QUESTION 55

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

- A. Check the returned response for the Versioned Compare the returned Versioned against the MD5 checksum.
- B. Include the MD5 checksum within the Content-MD5 parameter
- C. Check the 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 57

To run an application, a DevOps engineer launches an Amazon EC2 instance with public IP addresses in a public subnet. A user data script obtains the application artifacts and installs them on the instances upon launch. A change to the security classification of the application now requires the instances to run with no access

to the internet. While the instances launch successfully and show as healthy, the application does not seem to be installed. Which of the following should successfully install the application while complying with the new rule?

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

Answer: C

Explanation:

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

NEW QUESTION 59

A company hosts applications in its AWS account. Each application logs to an individual Amazon CloudWatch log group. The company's CloudWatch costs for ingestion are increasing.

A DevOps engineer needs to identify which applications are the source of the increased logging costs. Which solution will meet these requirements?

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

Answer: C

Explanation:

The correct answer is C.

A comprehensive and detailed explanation is:

? Option A is incorrect because using CloudWatch metrics to create a custom expression that identifies the CloudWatch log groups that have the most data being written to them is not a valid solution. CloudWatch metrics do not provide information about the size or volume of data being ingested by CloudWatch logs. CloudWatch metrics only provide information about the number of events, bytes, and errors that occur within a log group or stream. Moreover, creating a custom expression with CloudWatch metrics would require using the search_web tool, which is not necessary for this use case.

? Option B is incorrect because using CloudWatch Logs Insights to create a set of queries for the application log groups to identify the number of logs written for a period of time is not a valid solution. CloudWatch Logs Insights can help analyze and filter log events based on patterns and expressions, but it does not provide information about the cost or billing of CloudWatch logs. CloudWatch Logs Insights also charges based on the amount of data scanned by each query, which could increase the logging costs further.

? Option C is correct because using AWS Cost Explorer to generate a cost report that details the cost for CloudWatch usage is a valid solution. AWS Cost Explorer is a tool that helps visualize, understand, and manage AWS costs and usage over time. AWS Cost Explorer can generate custom reports that show the breakdown of costs by service, region, account, tag, or any other dimension. AWS Cost Explorer can also filter and group costs by usage type, which can help identify the specific CloudWatch log groups that are the source of the increased logging costs.

? Option D is incorrect because using AWS CloudTrail to filter for CreateLogStream events for each application is not a valid solution. AWS CloudTrail is a service that records API calls and account activity for AWS services, including CloudWatch logs. However, AWS CloudTrail does not provide information about the cost or billing of CloudWatch logs. Filtering for CreateLogStream events would only show when a new log stream was created within a log group, but not how much data was ingested or stored by that log stream.

References:

- ? CloudWatch Metrics
- ? CloudWatch Logs Insights
- ? AWS Cost Explorer
- ? AWS CloudTrail

NEW QUESTION 60

A DevOps engineer is working on a project that is hosted on Amazon Linux and has failed a security review. The DevOps manager has been asked to review the company buildspec.yaml file for an AWS CodeBuild project and provide recommendations. The buildspec.yaml file is configured as follows:

```
env:
  variables:
    AWS_ACCESS_KEY_ID: AKIAJF7BRFWJBA4GHXNA
    AWS_SECRET_ACCESS_KEY: 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
    AWS_DEFAULT_REGION: us-east-1
    DB_PASSWORD: cuj5RptFa3va
  phases:
    build:
      commands:
        - aws s3 cp s3://db-deploy-bucket/my.cnf.template /tmp/my.cnf
        - sed -i 's/DB_PW/${DB_PASSWORD}/' /tmp/my.cnf
        - aws s3 cp s3://db-deploy-bucket/instance.key /tmp/instance.key
        - chmod 600 /tmp/instance.key
        - scp -i /tmp/instance.key /tmp/my.cnf root@10.25.15.23:/etc/my.cnf
        - ssh -i /tmp/instance.key root@10.25.15.23 /etc/init.d/mysqld restart
```

What changes should be recommended to comply with AWS security best practices? (Select THREE.)

- A. Add a post-build command to remove the temporary files from the container before termination to ensure they cannot be seen by other CodeBuild users.

- B. Update the CodeBuild project role with the necessary permissions and then remove the AWS credentials from the environment variable.
- C. Store the db_password as a SecureString value in AWS Systems Manager Parameter Store and then remove the db_password from the environment variables.
- D. Move the environment variables to the 'db.-deploy-bucket' Amazon S3 bucket, add a prebuild stage to download then export the variables.
- E. Use AWS Systems Manager run command versus sec and ssh commands directly to the instance.

Answer: BCE

Explanation:

B. Update the CodeBuild project role with the necessary permissions and then remove the AWS credentials from the environment variable. C. Store the DB_PASSWORD as a SecureString value in AWS Systems Manager Parameter Store and then remove the DB_PASSWORD from the environment variables. E. Use AWS Systems Manager run command versus scp and ssh commands directly to the instance.

NEW QUESTION 61

An ecommerce company is receiving reports that its order history page is experiencing delays in reflecting the processing status of orders. The order processing system consists of an AWS Lambda function that uses reserved concurrency. The Lambda function processes order messages from an Amazon Simple Queue Service (Amazon SQS) queue and inserts processed orders into an Amazon DynamoDB table. The DynamoDB table has auto scaling enabled for read and write capacity.

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

- A. Check the ApproximateAgeOfOldestMessage metric for the SQS queue
- B. Increase the Lambda function concurrency limit.
- C. Check the ApproximateAgeOfOldestMessage metric for the SQS queue. Configure a redrive policy on the SQS queue.
- D. Check the NumberOfMessagesSent metric for the SQS queue
- E. Increase the SQS queue visibility timeout.
- F. Check the WriteThrottleEvents metric for the DynamoDB table
- G. Increase the maximum write capacity units (WCUs) for the table's scaling policy.
- H. Check the Throttles metric for the Lambda function
- I. Increase the Lambda function timeout.

Answer: AD

Explanation:

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

NEW QUESTION 65

A company requires its developers to tag all Amazon Elastic Block Store (Amazon EBS) volumes in an account to indicate a desired backup frequency. This requirement includes EBS volumes that do not require backups. The company uses custom tags named Backup_Frequency that have values of none, daily, or weekly that correspond to the desired backup frequency. An audit finds that developers are occasionally not tagging the EBS volumes.

A DevOps engineer needs to ensure that all EBS volumes always have the Backup_Frequency tag so that the company can perform backups at least weekly unless a different value is specified.

Which solution will meet these requirements?

- A. Set up AWS Config in the account
- B. Create a custom rule that returns a compliance failure for all Amazon EC2 resources that do not have a Backup Frequency tag applied. Configure a remediation action that uses a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly.
- C. Set up AWS Config in the account
- D. Use a managed rule that returns a compliance failure for EC2::Volume resources that do not have a Backup Frequency tag applied
- E. Configure a remediation action that uses a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly.
- F. Turn on AWS CloudTrail in the account
- G. Create an Amazon EventBridge rule that reacts to EBS CreateVolume event
- H. Configure a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly
- I. Specify the runbook as the target of the rule.
- J. Turn on AWS CloudTrail in the account
- K. Create an Amazon EventBridge rule that reacts to EBS CreateVolume events or EBS ModifyVolume event
- L. Configure a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly
- M. Specify the runbook as the target of the rule.

Answer: B

Explanation:

The following are the steps that the DevOps engineer should take to ensure that all EBS volumes always have the Backup_Frequency tag so that the company can perform backups at least weekly unless a different value is specified:

? Set up AWS Config in the account.

? Use a managed rule that returns a compliance failure for EC2::Volume resources that do not have a Backup Frequency tag applied.

? Configure a remediation action that uses a custom AWS Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly.

The managed rule AWS::Config::EBSVolumesWithoutBackupTag will return a compliance failure for any EBS volume that does not have the Backup_Frequency tag applied. The remediation action will then use the Systems Manager Automation runbook to apply the Backup_Frequency tag with a value of weekly to the EBS volume.

NEW QUESTION 66

A company has 20 service teams. 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 67

A company provides an application to customers. The application has an Amazon API Gateway REST API that invokes an AWS Lambda function. On initialization, the Lambda function loads a large amount of data from an Amazon DynamoDB table. The data load process results in long cold-start times of 8-10 seconds. The DynamoDB table has DynamoDB Accelerator (DAX) configured.

Customers report that the application intermittently takes a long time to respond to requests. The application receives thousands of requests throughout the day. In the middle of the day, the application experiences 10 times more requests than at any other time of the day. Near the end of the day, the application's request volume decreases to 10% of its normal total.

A DevOps engineer needs to reduce the latency of the Lambda function at all times of the day.

Which solution will meet these requirements?

- A. Configure provisioned concurrency on the Lambda function with a concurrency value of 1. Delete the DAX cluster for the DynamoDB table.
- B. Configure reserved concurrency on the Lambda function with a concurrency value of 0.
- C. Configure provisioned concurrency on the Lambda functio
- D. Configure AWS Application Auto Scaling on the Lambda function with provisioned concurrency values set to a minimum of 1 and a maximum of 100.
- E. Configure reserved concurrency on the Lambda functio
- F. Configure AWS Application Auto Scaling on the API Gateway API with a reserved concurrency maximum value of 100.

Answer: C

Explanation:

The following are the steps that the DevOps engineer should take to reduce the latency of the Lambda function at all times of the day:

? Configure provisioned concurrency on the Lambda function.

? Configure AWS Application Auto Scaling on the Lambda function with provisioned concurrency values set to a minimum of 1 and a maximum of 100.

The provisioned concurrency setting ensures that there is always a minimum number of Lambda function instances available to handle requests. The Application Auto Scaling setting will automatically scale the number of Lambda function instances up or down based on the demand for the application.

This solution will ensure that the Lambda function is able to handle the increased load during the middle of the day, while also keeping the cold-start latency low.

The following are the reasons why the other options are not correct:

? Option A is incorrect because it will not reduce the cold-start latency of the Lambda function.

? Option B is incorrect because it will not scale the number of Lambda function instances up or down based on demand.

? Option D is incorrect because it will only configure reserved concurrency on the API Gateway API, which will not affect the Lambda function.

NEW QUESTION 70

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 CloudFormationDeployment role.
- E. Configure the IAM to be to get and pass the CloudFormationDeployment role if cloudformation actions for resources,
- F. Update the trust Of the CloudFormationDeployment role to 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 CloudFormationDeplymet 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 75

A DevOps team is merging code revisions for an application that uses an Amazon RDS Multi-AZ DB cluster for its production database. The DevOps team uses continuous integration to periodically verify that the application works. The DevOps team needs to test the changes before the changes are deployed to the production database.

Which solution will meet these requirements'?

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

Answer: A

Explanation:

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

NEW QUESTION 77

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

Which logging solution will support these requirements?

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

Answer: A

Explanation:

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/SubscriptionFilters.html#LambdaFunctionExample>
<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/SubscriptionFilters.html>

NEW QUESTION 81

A DevOps team uses AWS CodePipeline, AWS CodeBuild, and AWS CodeDeploy to deploy an application. The application is a REST API that uses AWS Lambda functions and Amazon API Gateway. Recent deployments have introduced errors that have affected many customers. The DevOps team needs a solution that reverts to the most recent stable version of the application when an error is detected. The solution must affect the fewest customers possible.

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

- A. Set the deployment configuration in CodeDeploy to LambdaAllAtOnce. Configure automatic rollbacks on the deployment group. Create an Amazon CloudWatch alarm that detects HTTP Bad Gateway errors on API Gateway. Configure the deployment group to roll back when the number of alarms meets the alarm threshold.
- B. Set the deployment configuration in CodeDeploy to LambdaCanary10Percent10Minute.
- C. Configure automatic rollbacks on the deployment group. Create an Amazon CloudWatch alarm that detects HTTP Bad Gateway errors on API Gateway. Configure the deployment group to roll back when the number of alarms meets the alarm threshold.
- D. Set the deployment configuration in CodeDeploy to LambdaAllAtOnce. Configure manual rollbacks on the deployment group.
- E. Create an Amazon Simple Notification Service (Amazon SNS) topic to send notifications every time a deployment fails.
- F. Configure the SNS topic to invoke a new Lambda function that stops the current deployment and starts the most recent successful deployment.
- G. Set the deployment configuration in CodeDeploy to LambdaCanary10Percent10Minutes. Configure manual rollbacks on the deployment group. Create a metric filter on an Amazon CloudWatch log group for API Gateway to monitor HTTP Bad Gateway errors.
- H. Configure the metric filter to invoke a new Lambda function that stops the current deployment and starts the most recent successful deployment.

Answer: B

Explanation:

? Option A is incorrect because setting the deployment configuration to LambdaAllAtOnce means that the new version of the application will be deployed to all Lambda functions at once, affecting all customers. This does not meet the requirement of affecting the fewest customers possible. Moreover, configuring automatic rollbacks on the deployment group is not operationally efficient, as it requires manual intervention to fix the errors and redeploy the application.

? Option B is correct because setting the deployment configuration to LambdaCanary10Percent10Minutes means that the new version of the application will be deployed to 10 percent of the Lambda functions first, and then to the remaining 90 percent after 10 minutes. This minimizes the impact of errors on customers, as only 10 percent of them will be affected by a faulty deployment. Configuring automatic rollbacks on the deployment group also meets the requirement of reverting to the most recent stable version of the application when an error is detected. Creating a CloudWatch alarm that detects HTTP Bad Gateway errors on API Gateway is a valid way to monitor the health of the application and trigger a rollback if needed.

? Option C is incorrect because setting the deployment configuration to LambdaAllAtOnce means that the new version of the application will be deployed to all Lambda functions at once, affecting all customers. This does not meet the requirement of affecting the fewest customers possible. Moreover, configuring manual rollbacks on the deployment group is not operationally efficient, as it requires human intervention to stop the current deployment and start a new one. Creating an SNS topic to send notifications every time a deployment fails is not sufficient to detect errors in the application, as it does not monitor the API Gateway responses.

? Option D is incorrect because configuring manual rollbacks on the deployment group is not operationally efficient, as it requires human intervention to stop the current deployment and start a new one. Creating a metric filter on a CloudWatch log group for API Gateway to monitor HTTP Bad Gateway errors is a valid way to monitor the health of the application, but invoking a new Lambda function to perform a rollback is unnecessary and complex, as CodeDeploy already provides automatic rollback functionality.

References:

? AWS CodeDeploy Deployment Configurations

? [AWS CodeDeploy Rollbacks]

? Amazon CloudWatch Alarms

NEW QUESTION 82

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: CDE

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 86

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

Which solution will meet these requirements?

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

Answer: C

Explanation:

Amazon GuardDuty is primarily on threat detection and response, not configuration monitoring. A conformance pack is a collection of AWS Config rules and remediation actions that can be easily deployed as a single entity in an account and a Region or across an organization in AWS Organizations.
<https://docs.aws.amazon.com/config/latest/developerguide/conformance-packs.html> <https://docs.aws.amazon.com/config/latest/developerguide/s3-account-level-public-access-blocks.html>

NEW QUESTION 90

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 file
- 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 93

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 account
- 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 `required`
- E. Attach the permissions boundary to the IAM role that was used to launch the instance.
- F. Set up Amazon Inspector in the account
- G. Configure Amazon Inspector to activate deep inspection for EC2 instance
- H. Create an Amazon EventBridge rule for an Inspector2 finding
- 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 event
- 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 94

A healthcare services company is concerned about the growing costs of software licensing for an application for monitoring patient wellness. The company wants to create an audit process to ensure that the application is running exclusively on Amazon EC2 Dedicated Hosts. A DevOps engineer must create a workflow to audit the application to ensure compliance.

What steps should the engineer take to meet this requirement with the LEAST administrative overhead?

- A. Use AWS Systems Manager Configuration Compliance
- B. Use calls to the `put-compliance-items` API action to scan and build a database of noncompliant EC2 instances based on their host placement configuration
- C. Use an Amazon DynamoDB table to store these instance IDs for fast access
- D. Generate a report through Systems Manager by calling the `list-compliance-summaries` API action.
- E. Use custom Java code running on an EC2 instance
- F. Set up EC2 Auto Scaling for the instance depending on the number of instances to be checked
- G. Send the list of noncompliant EC2 instance IDs to an Amazon SQS queue
- H. Set up another worker instance to process instance IDs from the SQS queue and write them to Amazon DynamoDB
- I. Use an AWS Lambda function to terminate noncompliant instance IDs obtained from the queue, and send them to an Amazon SNS email topic for distribution.
- J. Use AWS Config
- K. Identify all EC2 instances to be audited by enabling Config Recording on all Amazon EC2 resources for the region
- L. Create a custom AWS Config rule that triggers an AWS Lambda function by using the `"config-rule-change-triggered"` blueprint. Modify the `LambdaevaluateCompliance()` function to verify host placement to return a `NON_COMPLIANT` result if the instance is not running on an EC2 Dedicated Host
- M. Use the AWS Config report to address noncompliant instances.
- N. Use AWS CloudTrail
- O. Identify all EC2 instances to be audited by analyzing all calls to the EC2 `RunCommand` API action
- P. Invoke a AWS Lambda function that analyzes the host placement of the instance
- Q. Store the EC2 instance ID of noncompliant resources in an Amazon RDS for MySQL DB instance
- R. Generate a report by querying the RDS instance and exporting the query results to a CSV text file.

Answer: C

Explanation:

The correct answer is C. Using AWS Config to identify and audit all EC2 instances based on their host placement configuration is the most efficient and scalable solution to ensure compliance with the software licensing requirement. AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. By creating a custom AWS Config rule that triggers a Lambda function to verify host placement, the DevOps engineer can automate the process of checking whether the instances are running on EC2 Dedicated Hosts or not. The Lambda function can return a `NON_COMPLIANT` result if the instance is not running on an EC2 Dedicated Host, and the AWS Config report can provide a summary of the compliance status of the instances. This solution requires the least administrative overhead compared to the other options.

Option A is incorrect because using AWS Systems Manager Configuration Compliance to scan and build a database of noncompliant EC2 instances based on their host placement configuration is a more complex and costly solution than using AWS Config. AWS Systems Manager Configuration Compliance is a feature of AWS Systems Manager that enables you to scan your managed instances for patch compliance and configuration inconsistencies. To use this feature, the DevOps engineer would need to install the Systems Manager Agent on each EC2 instance, create a State Manager association to run the `put-compliance-items` API action periodically, and use a DynamoDB table to store the instance IDs of noncompliant resources. This solution would also require more API calls and storage costs than using AWS Config.

Option B is incorrect because using custom Java code running on an EC2 instance to check and terminate noncompliant EC2 instances is a more cumbersome and error-prone solution than using AWS Config. This solution would require the DevOps engineer to write and maintain the Java code, set up EC2 Auto Scaling for the instance, use an SQS queue and another worker instance to process the instance IDs, use a Lambda function and an SNS topic to terminate and notify the noncompliant instances, and handle any potential failures or exceptions in the workflow. This solution would also incur more compute,

storage, and messaging costs than using AWS Config.

Option D is incorrect because using AWS CloudTrail to identify and audit EC2 instances by analyzing the EC2 RunCommand API action is a less reliable and accurate solution than using AWS Config. AWS CloudTrail is a service that enables you to monitor and log the API activity in your AWS account. The EC2 RunCommand API action is used to execute commands on one or more EC2 instances. However, this API action does not necessarily indicate the host placement of the instance, and it may not capture all the instances that are running on EC2 Dedicated Hosts or not. Therefore, option D would not provide a comprehensive and consistent audit of the EC2 instances.

NEW QUESTION 99

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 subscription¹ 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 account². This can be done using a CloudFormation template and StackSets³ 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 accounts⁴. The role must have a trust policy that specifies the organization ID as a condition.

? Finally, the DevOps engineer needs to attach the

CloudWatchLogsReadOnlyAccess policy⁵ 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 104

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 plan.
- 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 configuration.
- 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 105

A company hosts its staging website using an Amazon EC2 instance backed with Amazon EBS storage. The company wants to recover quickly with minimal data losses in the event of network connectivity issues or power failures on the EC2 instance.

Which solution will meet these requirements?

- A. Add the instance to an EC2 Auto Scaling group with the minimum, maximum, and desired capacity set to 1.
- B. Add the instance to an EC2 Auto Scaling group with a lifecycle hook to detach the EBS volume when the EC2 instance shuts down or terminates.
- C. Create an Amazon CloudWatch alarm for the StatusCheckFailed System metric and select the EC2 action to recover the instance.
- D. Create an Amazon CloudWatch alarm for the StatusCheckFailed Instance metric and select the EC2 action to reboot the instance.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-recover.html>

NEW QUESTION 107

A development team is using AWS CodeCommit to version control application code and AWS CodePipeline to orchestrate software deployments. The team has decided to use a remote main branch as the trigger for the pipeline to integrate code changes. A developer has pushed code changes to the CodeCommit repository, but noticed that the pipeline had no reaction, even after 10 minutes.

Which of the following actions should be taken to troubleshoot this issue?

- A. Check that an Amazon EventBridge rule has been created for the main branch to trigger the pipeline.
- B. Check that the CodePipeline service role has permission to access the CodeCommit repository.
- C. Check that the developer's IAM role has permission to push to the CodeCommit repository.
- D. Check to see if the pipeline failed to start because of CodeCommit errors in Amazon CloudWatch Logs.

Answer: A

Explanation:

When you create a pipeline from CodePipeline during the step-by-step it creates a CloudWatch Event rule for a given branch and repo like this:

```
{
"source": [ "aws.codecommit"
],
"detail-type": [
"CodeCommit Repository State Change"
],
"resources": [
"arn:aws:codecommit:us-east-1:xxxxx:repo-name"
],
"detail": {
"event": [ "referenceCreated", "referenceUpdated"
],
"referenceType": [ "branch"
],
"referenceName": [ "master"
]
}
}
```

<https://docs.aws.amazon.com/codepipeline/latest/userguide/pipelines-trigger-source-repo-changes-console.html>

NEW QUESTION 108

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 Regions¹².

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

References:

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

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

NEW QUESTION 112

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

What would cause this?

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

Answer: C

Explanation:

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

NEW QUESTION 114

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 IAM role by providing the necessary permissions to update the stack
- 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 account
- F. Run the aws cloudformation cancel-update-stack AWS CLI command.
- G. Delete the Auto Scaling group resource
- 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 118

A company has many AWS accounts. During AWS account creation the company uses automation to create an Amazon CloudWatch Logs log group in every AWS Region that the company operates in. The automation configures new resources in the accounts to publish logs to the provisioned log groups in their Region. The company has created a logging account to centralize the logging from all the other accounts. A DevOps engineer needs to aggregate the log groups from all the accounts to an existing Amazon S3 bucket in the logging account.

Which solution will meet these requirements in the MOST operationally efficient manner?

- A. In the logging account create a CloudWatch Logs destination with a destination policy
- B. For each new account subscribe the CloudWatch Logs log groups to the destination
- C. Destination Configure a single Amazon Kinesis data stream and a single Amazon Kinesis Data Firehose delivery stream to deliver the logs from the CloudWatch Logs destination to the S3 bucket.
- D. In the logging account create a CloudWatch Logs destination with a destination policy for each Region
- E. For each new account subscribe the CloudWatch Logs log groups to the destination
- F. Configure a single Amazon Kinesis data stream and a single Amazon Kinesis Data Firehose delivery stream to deliver the logs from all the CloudWatch Logs destinations to the S3 bucket.
- G. In the logging account create a CloudWatch Logs destination with a destination policy for each Region
- H. For each new account subscribe the CloudWatch Logs log groups to the destination Configure an Amazon Kinesis data stream and an Amazon Kinesis Data Firehose delivery stream for each Region to deliver the logs from the CloudWatch Logs destinations to the S3 bucket.
- I. In the logging account create a CloudWatch Logs destination with a destination policy
- J. For each new account subscribe the CloudWatch Logs log groups to the destination
- K. Configure a single Amazon Kinesis data stream to deliver the logs from the CloudWatch Logs destination to the S3 bucket.

Answer: C

Explanation:

This solution will meet the requirements in the most operationally efficient manner because it will use CloudWatch Logs destination to aggregate the log groups from all the accounts to a single S3 bucket in the logging account. However, unlike option A, this solution will create a CloudWatch Logs destination for each region, instead of a single destination for all regions. This will improve the performance and reliability of the log delivery, as it will avoid cross-region data transfer and latency issues. Moreover, this solution will use an Amazon Kinesis data stream and an Amazon Kinesis Data Firehose delivery stream for each region, instead of a single stream for all regions. This will also improve the scalability and throughput of the log delivery, as it will avoid bottlenecks and throttling issues that may occur with a single stream.

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

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

- * AWS-Certified-DevOps-Engineer-Professional Questions and Answers Updated Frequently
- * AWS-Certified-DevOps-Engineer-Professional Practice Questions Verified by Expert Senior Certified Staff
- * AWS-Certified-DevOps-Engineer-Professional Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * AWS-Certified-DevOps-Engineer-Professional Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The AWS-Certified-DevOps-Engineer-Professional Practice Test Here](#)