



Amazon

Exam Questions AWS-Certified-DevOps-Engineer-Professional

Amazon AWS Certified DevOps Engineer Professional

About ExamBible

Your Partner of IT Exam

Found in 1998

ExamBible is a company specialized on providing high quality IT exam practice study materials, especially Cisco CCNA, CCDA, CCNP, CCIE, Checkpoint CCSE, CompTIA A+, Network+ certification practice exams and so on. We guarantee that the candidates will not only pass any IT exam at the first attempt but also get profound understanding about the certificates they have got. There are so many alike companies in this industry, however, ExamBible has its unique advantages that other companies could not achieve.

Our Advances

* 99.9% Uptime

All examinations will be up to date.

* 24/7 Quality Support

We will provide service round the clock.

* 100% Pass Rate

Our guarantee that you will pass the exam.

* Unique Gurantee

If you do not pass the exam at the first time, we will not only arrange FULL REFUND for you, but also provide you another exam of your claim, ABSOLUTELY FREE!

NEW QUESTION 1

A company runs an application on one Amazon EC2 instance. Application metadata is stored in Amazon S3 and must be retrieved if the instance is restarted. The instance must restart or relaunch automatically if the instance becomes unresponsive.

Which solution will meet these requirements?

- A. Create an Amazon CloudWatch alarm for the StatusCheckFailed metri
- B. Use the recover action to stop and start the instanc
- C. Use an S3 event notification to push the metadata to the instance when the instance is back up and running.
- D. Configure AWS OpsWorks, and use the auto healing feature to stop and start the instanc
- E. Use a lifecycle event in OpsWorks to pull the metadata from Amazon S3 and update it on the instance.
- F. Use EC2 Auto Recovery to automatically stop and start the instance in case of a failur
- G. Use an S3 event notification to push the metadata to the instance when the instance is back up and running.
- H. Use AWS CloudFormation to create an EC2 instance that includes the UserData property for the EC2 resourc
- I. Add a command in UserData to retrieve the application metadata from Amazon S3.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/mt/how-to-set-up-aws-opsworks-stacks-auto-healing-notifications-in-amazon-cloudwatch-events/>

NEW QUESTION 2

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 Conflg 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 3

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

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

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

Answer: C

Explanation:

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

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

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

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

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

NEW QUESTION 4

A company is using an organization in AWS Organizations to manage multiple AWS accounts. The company's development team wants to use AWS Lambda functions to meet resiliency requirements and is rewriting all applications to work with Lambda functions that are deployed in a VPC. The development team is using Amazon Elastic File System (Amazon EFS) as shared storage in Account A in the organization.

The company wants to continue to use Amazon EPS with Lambda Company policy requires all serverless projects to be deployed in Account B.

A DevOps engineer needs to reconfigure an existing EFS file system to allow Lambda functions to access the data through an existing EPS access point.

Which combination of steps should the DevOps engineer take to meet these requirements? (Select THREE.)

- A. Update the EFS file system policy to provide Account B with access to mount and write to the EFS file system in Account A.
- B. Create SCPs to set permission guardrails with fine-grained control for Amazon EFS.
- C. Create a new EFS file system in Account B Use AWS Database Migration Service (AWS DMS) to keep data from Account A and Account B synchronized.
- D. Update the Lambda execution roles with permission to access the VPC and the EFS file system.

- E. Create a VPC peering connection to connect Account A to Account B.
- F. Configure the Lambda functions in Account B to assume an existing IAM role in Account A.

Answer: AEF

Explanation:

A Lambda function in one account can mount a file system in a different account. For this scenario, you configure VPC peering between the function VPC and the file system VPC. <https://docs.aws.amazon.com/lambda/latest/dg/services-efs.html> <https://aws.amazon.com/ru/blogs/storage/mount-amazon-efs-file-systems-cross-account-from-amazon-eks/>

* 1. Need to update the file system policy on EFS to allow mounting the file system into Account B.

File System Policy

\$ cat file-system-policy.json

```
{
  "Statement": [
    {
      "Effect": "Allow", "Action": [
        "elasticfilesystem:ClientMount", "elasticfilesystem:ClientWrite"
      ],
      "Principal": {
        "AWS": "arn:aws:iam::<aws-account-id-A>:root" # Replace with AWS account ID of EKS cluster
      }
    }
  ]
}
```

* 2. Need VPC peering between Account A and Account B as the pre-requisite

* 3. Need to assume cross-account IAM role to describe the mounts so that a specific mount can be chosen.

NEW QUESTION 5

A development team manually builds an artifact locally and then places it in an Amazon S3 bucket. The application has a local cache that must be cleared when a deployment occurs. The team runs a command to do this downloads the artifact from Amazon S3 and unzips the artifact to complete the deployment.

A DevOps team wants to migrate to a CI/CD process and build in checks to stop and roll back the deployment when a failure occurs. This requires the team to track the progression of the deployment.

Which combination of actions will accomplish this? (Select THREE)

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

Answer: BDE

NEW QUESTION 6

A company is developing a new application. The application uses AWS Lambda functions for its compute tier. The company must use a canary deployment for any changes to the Lambda functions. Automated rollback must occur if any failures are reported.

The company's DevOps team needs to create the infrastructure as code (IaC) and the CI/CD pipeline for this solution.

Which combination of steps will meet these requirements? (Choose three.)

- A. Create an AWS CloudFormation template for the applicatio
- B. Define each Lambda function in the template by using the AWS::Lambda::Function resource typ
- C. In the template, include a version for the Lambda function by using the AWS::Lambda::Version resource typ
- D. Declare the CodeSha256 propert
- E. Configure an AWS::Lambda::Alias resource that references the latest version of the Lambda function.
- F. Create an AWS Serverless Application Model (AWS SAM) template for the applicatio
- G. Define each Lambda function in the template by using the AWS::Serverless::Function resource typ
- H. For each function, include configurations for the AutoPublishAlias property and the DeploymentPreference propert
- I. Configure the deployment configuration type to LambdaCanary10Percent10Minutes.
- J. Create an AWS CodeCommit repositor
- K. Create an AWS CodePipeline pipelin
- L. Use the CodeCommit repository in a new source stage that starts the pipelin
- M. Create an AWS CodeBuild project to deploy the AWS Serverless Application Model (AWS SAM) templat
- N. Upload the template and source code to the CodeCommit repositor
- O. In the CodeCommit repository, create a buildspec.yml file that includes the commands to build and deploy the SAM application.
- P. Create an AWS CodeCommit repositor
- Q. Create an AWS CodePipeline pipelin
- R. Use the CodeCommit repository in a new source stage that starts the pipelin
- S. Create an AWS CodeDeploy deployment group that is configured for canary deployments with a DeploymentPreference type of Canary10Percent10Minute
- T. Upload the AWS CloudFormation template and source code to the CodeCommit repositor
- . In the CodeCommit repository, create an appspec.yml file that includes the commands to deploy the CloudFormation template.
- . Create an Amazon CloudWatch composite alarm for all the Lambda function
- . Configure an evaluation period and dimensions for Lambd
- . Configure the alarm to enter the ALARMstate if any errors are detected or if there is insufficient data.
- . Create an Amazon CloudWatch alarm for each Lambda functio
- . Configure the alarms to enter the ALARM state if any errors are detecte
- . Configure an evaluation period, dimensions for each Lambda function and version, and the namespace as AWS/Lambda on the Errors metric.

Answer: BCF

Explanation:

The requirement is to create the infrastructure as code (IaC) and the CI/CD pipeline for the Lambda application that uses canary deployment and automated rollback. To do this, the DevOps team needs to use the following steps:

? Create an AWS Serverless Application Model (AWS SAM) template for the application. AWS SAM is a framework that simplifies the development and deployment of serverless applications on AWS. AWS SAM allows customers to define Lambda functions and other resources in a template by using a simplified syntax. For each Lambda function, the DevOps team can include configurations for the `AutoPublishAlias` property and the `DeploymentPreference` property. The `AutoPublishAlias` property specifies the name of the alias that points to the latest version of the function. The `DeploymentPreference` property specifies how CodeDeploy deploys new versions of the function. By configuring the deployment configuration type to `LambdaCanary10Percent10Minutes`, the DevOps team can enable canary deployment with 10% of traffic shifted to the new version every 10 minutes.

? Create an AWS CodeCommit repository. Create an AWS CodePipeline pipeline.

Use the CodeCommit repository in a new source stage that starts the pipeline. Create an AWS CodeBuild project to deploy the AWS SAM template. CodeCommit is a fully managed source control service that hosts Git repositories. CodePipeline is a fully managed continuous delivery service that automates the release process of software applications. CodeBuild is a fully managed continuous integration service that compiles source code and runs tests. By using these services, the DevOps team can create a CI/CD pipeline for the Lambda application. The pipeline should use the CodeCommit repository as the source stage, where the DevOps team can upload the SAM template and source code. The pipeline should also use a CodeBuild project as the build stage, where the SAM template can be built and deployed.

? Create an Amazon CloudWatch alarm for each Lambda function. Configure the alarms to enter the ALARM state if any errors are detected. Configure an evaluation period, dimensions for each Lambda function and version, and the namespace as `AWS/Lambda` on the Errors metric. CloudWatch is a service that monitors and collects metrics from AWS resources and applications. CloudWatch alarms are actions that are triggered when a metric crosses a specified threshold. By creating CloudWatch alarms for each Lambda function, the DevOps team can monitor the health and performance of each function version during deployment. By configuring the alarms to enter the ALARM state if any errors are detected, the DevOps team can enable automated rollback if any failures are reported.

NEW QUESTION 7

A company has an application that runs on a fleet of Amazon EC2 instances. The application requires frequent restarts. The application logs contain error messages when a restart is required. The application logs are published to a log group in Amazon CloudWatch Logs.

An Amazon CloudWatch alarm notifies an application engineer through an Amazon Simple Notification Service (Amazon SNS) topic when the logs contain a large number of restart-related error messages. The application engineer manually restarts the application on the instances after the application engineer receives a notification from the SNS topic.

A DevOps engineer needs to implement a solution to automate the application restart on the instances without restarting the instances.

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

- A. Configure an AWS Systems Manager Automation runbook that runs a script to restart the application on the instance
- B. Configure the SNS topic to invoke the runbook.
- C. Create an AWS Lambda function that restarts the application on the instance
- D. Configure the Lambda function as an event destination of the SNS topic.
- E. Configure an AWS Systems Manager Automation runbook that runs a script to restart the application on the instance
- F. Create an AWS Lambda function to invoke the runbook
- G. Configure the Lambda function as an event destination of the SNS topic.
- H. Configure an AWS Systems Manager Automation runbook that runs a script to restart the application on the instance
- I. Configure an Amazon EventBridge rule that reacts when the CloudWatch alarm enters ALARM state
- J. Specify the runbook as a target of the rule.

Answer: D

Explanation:

This solution meets the requirements in the most operationally efficient manner by automating the application restart process on the instances without restarting them. When the CloudWatch alarm enters the ALARM state, the EventBridge rule is triggered, which in turn invokes the Systems Manager Automation runbook that contains the script to restart the application on the instances.

NEW QUESTION 8

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 9

A company is launching an application that stores raw data in an Amazon S3 bucket. Three applications need to access the data to generate reports. The data must be redacted differently for each application before the applications can access the data.

Which solution will meet these requirements?

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

Answer: D

Explanation:

? The best solution is to use S3 Object Lambda¹, which allows you to add your own code to S3 GET, LIST, and HEAD requests to modify and process data as it is returned to an application². This way, you can redact the data differently for each application without creating and storing multiple copies of the data or running proxies.

? The other solutions are less efficient or scalable because they require replicating the data to multiple buckets, streaming the data through Kinesis, or storing the data in S3 access points.

References: 1: Amazon S3 Features | Object Lambda | AWS 2: Transforming objects with S3 Object Lambda - Amazon Simple Storage Service

NEW QUESTION 10

A company is testing a web application that runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Auto Scaling group across multiple Availability Zones. The company uses a blue green deployment process with immutable instances when deploying new software.

During testing users are being automatically logged out of the application at random times. Testers also report that when a new version of the application is deployed all users are logged out. The development team needs a solution to ensure users remain logged m across scaling events and application deployments.

What is the MOST operationally efficient way to ensure users remain logged in?

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

Answer: D

Explanation:

<https://aws.amazon.com/caching/session-management/>

NEW QUESTION 10

A company needs a strategy for failover and disaster recovery of its data and application. The application uses a MySQL database and Amazon EC2 instances. The company requires a maximum RPO of 2 hours and a maximum RTO of 10 minutes for its data and application at all times.

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

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

- M. Add both ALBs to a single endpoint 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 15

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

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

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

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

Answer: A

NEW QUESTION 16

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 21

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

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

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

Answer: BDE

Explanation:

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

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

RDS for MySQL DB instance outside of Elastic Beanstalk. This approach ensures that the database persists independently of the Elastic Beanstalk environment, which can be torn down and recreated without affecting the database¹²³.

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

versions. Immutable deployments provide a zero-downtime deployment method that ensures that if any part of the deployment process fails, the environment is rolled back to the original state automatically⁴.

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

Amazon Simple Notification Service (Amazon SNS) topic as a target to alert the application team. This setup allows for automated monitoring and alerting of the application team in case of deployment failures or other health events⁵⁶.

References:

? AWS Elastic Beanstalk documentation on deploying Ruby applications¹.

? AWS documentation on application auto-scaling⁷.

? AWS documentation on automated deployment strategies with automatic rollbacks and alerts⁴⁵⁶.

NEW QUESTION 24

A space exploration company receives telemetry data from multiple satellites. Small packets of data are received through Amazon API Gateway and are placed directly into an Amazon Simple Queue Service (Amazon SQS) standard queue. A custom application is subscribed to the queue and transforms the data into a standard format.

Because of inconsistencies in the data that the satellites produce, the application is occasionally unable to transform the data. In these cases, the messages remain in the SQS queue. A DevOps engineer must develop a solution that retains the failed messages and makes them available to scientists for review and future processing.

Which solution will meet these requirements?

- A. Configure AWS Lambda to poll the SQS queue and invoke a Lambda function to check whether the queue messages are valid
- B. If validation fails, send a copy of the data that is not valid to an Amazon S3 bucket so that the scientists can review and correct the data
- C. When the data is corrected, amend the message in the SQS queue by using a replay Lambda function with the corrected data.
- D. Convert the SQS standard queue to an SQS FIFO queue
- E. Configure AWS Lambda to poll the SQS queue every 10 minutes by using an Amazon EventBridge schedule
- F. Invoke the Lambda function to identify any messages with a SentTimestamp value that is older than 5 minutes, push the data to the same location as the application's output location, and remove the messages from the queue.
- G. Create an SQS dead-letter queue
- H. Modify the existing queue by including a receive policy that sets the Maximum Receives setting to 1 and sets the dead-letter queue ARN to the ARN of the newly created queue
- I. Instruct the scientists to use the dead-letter queue to review the data that is not valid
- J. Reprocess this data at a later time.
- K. Configure API Gateway to send messages to different SQS virtual queues that are named for each of the satellites
- L. Update the application to use a new virtual queue for any data that it cannot transform, and send the message to the new virtual queue
- M. Instruct the scientists to use the virtual queue to review the data that is not valid
- N. Reprocess this data at a later time.

Answer: C

Explanation:

Create an SQS dead-letter queue. Modify the existing queue by including a redrive policy that sets the Maximum Receives setting to 1 and sets the dead-letter queue ARN to the ARN of the newly created queue. Instruct the scientists to use the dead-letter queue to review the data that is not valid. Reprocess this data at a later time.

NEW QUESTION 26

A company's production environment uses an AWS CodeDeploy blue/green deployment to deploy an application. The deployment includes Amazon EC2 Auto Scaling groups that launch instances that run Amazon Linux 2.

A working appspec. yml file exists in the code repository and contains the following text.

```
version: 0.0
os: linux
files:
  - source: /
    destination: /var/www/html/application
```

A DevOps engineer needs to ensure that a script downloads and installs a license file onto the instances before the replacement instances start to handle request traffic. The DevOps engineer adds a hooks section to the appspec. yml file.

Which hook should the DevOps engineer use to run the script that downloads and installs the license file?

- A. AfterBlockTraffic
- B. BeforeBlockTraffic
- C. BeforeInstall
- D. Download Bundle

Answer: C

Explanation:

This hook runs before the new application version is installed on the replacement instances. This is the best place to run the script because it ensures that the license file is downloaded and installed before the replacement instances start to handle request traffic. If you use any other hook, you may encounter errors or inconsistencies in your application.

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 29

A DevOps engineer manages a large commercial website that runs on Amazon EC2. The website uses Amazon Kinesis Data Streams to collect and process web logs. The DevOps engineer manages the Kinesis consumer application, which also runs on Amazon EC2.

Sudden increases of data cause the Kinesis consumer application to fall behind and the Kinesis data streams drop records before the records can be processed. The DevOps engineer must implement a solution to improve stream handling.

Which solution meets these requirements with the MOST operational efficiency?

- A. Modify the Kinesis consumer application to store the logs durably in Amazon S3. Use Amazon EMR to process the data directly on Amazon S3 to derive customer insights. Store the results in Amazon S3.
- B. Horizontally scale the Kinesis consumer application by adding more EC2 instances based on the Amazon CloudWatch GetRecords IteratorAgeMilliseconds metric. Increase the retention period of the Kinesis data streams.
- C. Convert the Kinesis consumer application to run as an AWS Lambda function.
- D. Configure the Kinesis data streams as the event source for the Lambda function to process the data streams.
- E. Increase the number of shards in the Kinesis data streams to increase the overall throughput so that the consumer application processes the data faster.

Answer: B

Explanation:

<https://docs.aws.amazon.com/streams/latest/dev/monitoring-with-cloudwatch.html>

GetRecords.IteratorAgeMilliseconds - The age of the last record in all GetRecords calls made against a Kinesis stream, measured over the specified time period. Age is the difference between the current time and when the last record of the GetRecords call was written to the stream. The Minimum and Maximum statistics can be used to track the progress of Kinesis consumer applications. A value of zero indicates that the records being read are completely caught up.

NEW QUESTION 31

A company is using an AWS CodeBuild project to build and package an application. The packages are copied to a shared Amazon S3 bucket before being deployed across multiple AWS accounts.

The buildspec.yml file contains the following:

```
version: 0.2
phases:
  build:
    commands:
      - go build -o myapp
  post_build:
    commands:
      - aws s3 cp --acl authenticated-read myapp s3://artifacts/
```

The DevOps engineer has noticed that anybody with an AWS account is able to download the artifacts.

What steps should the DevOps engineer take to stop this?

- A. Modify the post_build command to use --acl public-read and configure a bucket policy that grants read access to the relevant AWS accounts only.
- B. Configure a default ACL for the S3 bucket that defines the set of authenticated users as the relevant AWS accounts only and grants read-only access.
- C. Create an S3 bucket policy that grants read access to the relevant AWS accounts and denies read access to the principal “*”.
- D. Modify the post_build command to remove --acl authenticated-read and configure a bucket policy that allows read access to the relevant AWS accounts only.

Answer: D

Explanation:

When setting the flag authenticated-read in the command line, the owner gets FULL_CONTROL. The AuthenticatedUsers group (Anyone with an AWS account) gets READ access. Reference: <https://docs.aws.amazon.com/AmazonS3/latest/userguide/acl-overview.html>

NEW QUESTION 34

A company is using AWS to run digital workloads. Each application team in the company has its own AWS account for application hosting. The accounts are consolidated in an organization in AWS Organizations.

The company wants to enforce security standards across the entire organization. To avoid noncompliance because of security misconfiguration, the company has enforced the use of AWS CloudFormation. A production support team can modify resources in the production environment by using the AWS Management Console to troubleshoot and resolve application-related issues.

A DevOps engineer must implement a solution to identify in near real time any AWS service misconfiguration that results in noncompliance. The solution must automatically remediate the issue within 15 minutes of identification. The solution also must track noncompliant resources and events in a centralized dashboard with accurate timestamps.

Which solution will meet these requirements with the LEAST development overhead?

- A. Use CloudFormation drift detection to identify noncompliant resource
- B. Use drift detection events from CloudFormation to invoke an AWS Lambda function for remediation
- C. Configure the Lambda function to publish logs to an Amazon CloudWatch Logs log group
- D. Configure an Amazon CloudWatch dashboard to use the log group for tracking.
- E. Turn on AWS CloudTrail in the AWS account
- F. Analyze CloudTrail logs by using Amazon Athena to identify noncompliant resource
- G. Use AWS Step Functions to track query results on Athena for drift detection and to invoke an AWS Lambda function for remediation
- H. For tracking, set up an Amazon QuickSight dashboard that uses Athena as the data source.
- I. Turn on the configuration recorder in AWS Config in all the AWS accounts to identify noncompliant resource
- J. Enable AWS Security Hub with the ~no-enable-default-standards option in all the AWS account
- K. Set up AWS Config managed rules and custom rule
- L. Set up automatic remediation by using AWS Config conformance pack
- M. For tracking, set up a dashboard on Security Hub in a designated Security Hub administrator account.
- N. Turn on AWS CloudTrail in the AWS account
- O. Analyze CloudTrail logs by using Amazon CloudWatch Logs to identify noncompliant resource
- P. Use CloudWatch Logs filters for drift detection
- Q. Use Amazon EventBridge to invoke the Lambda function for remediation
- R. Stream filtered CloudWatch logs to Amazon OpenSearch Service
- S. Set up a dashboard on OpenSearch Service for tracking.

Answer: C

Explanation:

The best solution is to use AWS Config and AWS Security Hub to identify and remediate noncompliant resources across multiple AWS accounts. AWS Config enables continuous monitoring of the configuration of AWS resources and evaluates them against desired configurations. AWS Config can also automatically remediate noncompliant resources by using conformance packs, which are a collection of AWS Config rules and remediation actions that can be deployed as a single entity. AWS Security Hub provides a comprehensive view of the security posture of AWS accounts and resources. AWS Security Hub can aggregate and normalize the findings from AWS Config and other AWS services, as well as from partner solutions. AWS Security Hub can also be used to create a dashboard for tracking noncompliant resources and events in a centralized location.

The other options are not optimal because they either require more development overhead, do not provide near real time detection and remediation, or do not provide a centralized dashboard for tracking.

Option A is not optimal because CloudFormation drift detection is not a near real time solution. Drift detection has to be manually initiated on each stack or resource, or scheduled using a cron expression. Drift detection also does not provide remediation

actions, so a custom Lambda function has to be developed and invoked. CloudWatch Logs and dashboard can be used for tracking, but they do not provide a comprehensive view of the security posture of the AWS accounts and resources.

Option B is not optimal because CloudTrail logs analysis using Athena is not a near real time solution. Athena queries have to be manually run or scheduled using

a cron expression. Athena also does not provide remediation actions, so a custom Lambda function has to be developed and invoked. Step Functions can be used to orchestrate the query and remediation workflow, but it adds more complexity and cost. QuickSight dashboard can be used for tracking, but it does not provide a comprehensive view of the security posture of the AWS accounts and resources.

Option D is not optimal because CloudTrail logs analysis using CloudWatch Logs is not a near real time solution. CloudWatch Logs filters have to be manually created or updated for each resource type and configuration change. CloudWatch Logs also does not provide remediation actions, so a custom Lambda function has to be developed and invoked. EventBridge can be used to trigger the Lambda function, but it adds more complexity and cost. OpenSearch Service dashboard can be used for tracking, but it does not provide a comprehensive view of the security posture of the AWS accounts and resources. References:

- ? AWS Config conformance packs
- ? Introducing AWS Config conformance packs
- ? Managing conformance packs across all accounts in your organization

NEW QUESTION 36

A company builds an application that uses an Application Load Balancer in front of Amazon EC2 instances that are in an Auto Scaling group. The application is stateless. The Auto Scaling group uses a custom AMI that is fully prebuilt. The EC2 instances do not have a custom bootstrapping process. The AMI that the Auto Scaling group uses was recently deleted. The Auto Scaling group's scaling activities show failures because the AMI ID does not exist. Which combination of steps should a DevOps engineer take to meet these requirements? (Select THREE.)

- A. Create a new launch template that uses the new AMI.
- B. Update the Auto Scaling group to use the new launch template.
- C. Reduce the Auto Scaling group's desired capacity to 0.
- D. Increase the Auto Scaling group's desired capacity by 1.
- E. Create a new AMI from a running EC2 instance in the Auto Scaling group.
- F. Create a new AMI by copying the most recent public AMI of the operating system that the EC2 instances use.

Answer: ABF

Explanation:

To restore the functionality of the Auto Scaling group after the AMI was deleted, the DevOps engineer needs to create a new AMI and update the Auto Scaling group to use it. The DevOps engineer can create a new AMI by copying the most recent public AMI of the operating system that the EC2 instances use. This will ensure that the new AMI has the same operating system as the custom AMI that was deleted. The DevOps engineer can then create a new launch template that uses the new AMI and update the Auto Scaling group to use the new launch template. This will allow the Auto Scaling group to launch new instances with the new AMI.

NEW QUESTION 39

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

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

Answer: B

Explanation:

The best solution to implement failover for the application is to use CloudFront origin groups. Origin groups allow CloudFront to automatically switch to a secondary origin when the primary origin is unavailable or returns specific HTTP status codes that indicate a failure¹. This way, CloudFront can serve the requests from the secondary ALB in the secondary Region without any delay or redirection. To set up origin groups, the DevOps engineer needs to create a new origin on the distribution for the secondary ALB, create a new origin group with the original ALB as the primary origin and the secondary ALB as the secondary origin, and configure the origin group to fail over for HTTP 5xx status codes. Then, the DevOps engineer needs to update the default behavior to use the origin group instead of the single origin².

The other options are not as effective or efficient as the solution in option B. Option A is not suitable because creating a second CloudFront distribution will increase the complexity and cost of the application. Moreover, using Route 53 alias records with a failover policy will introduce some delay in detecting and switching to the secondary CloudFront distribution, which may not meet the zero-second RTO requirement. Option C is not feasible because CloudFront does not support using Route 53 alias records as origins³. Option D is not advisable because using a CloudFront function to redirect the requests to the secondary ALB will add an extra round-trip and latency to the failover process, which may also not meet the zero-second RTO requirement.

References:

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

NEW QUESTION 41

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 46

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 auto-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 even
- 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 50

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 .ebextensions 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 54

A company updated the AWS CloudFormation template for a critical business application. The stack update process failed due to an error in the updated template and AWS CloudFormation automatically began the stack rollback process. Later a DevOps engineer discovered that the application was still unavailable and that the stack was in the UPDATE_ROLLBACK_FAILED state.

Which combination of actions should the DevOps engineer perform so that the stack rollback can complete successfully? (Select TWO.)

- A. Attach the AWS CloudFormation FullAccess IAM policy to the AWS CloudFormation role.
- B. Automatically recover the stack resources by using AWS CloudFormation drift detection.
- C. Issue a ContinueUpdateRollback command from the AWS CloudFormation console or the AWS CLI.
- D. Manually adjust the resources to match the expectations of the stack.
- E. Update the existing AWS CloudFormation stack by using the original template.

Answer: CD

Explanation:

<https://docs.aws.amazon.com/cli/latest/reference/cloudformation/continue-update-rollback.html> For a specified stack that is in the UPDATE_ROLLBACK_FAILED state, continues rolling it back to the UPDATE_ROLLBACK_COMPLETE state. Depending on the cause of the failure, you can manually fix the error and continue the rollback. By continuing the rollback, you can return your stack to a working state (the UPDATE_ROLLBACK_COMPLETE state), and then try to update the stack again.

NEW QUESTION 58

A global company manages multiple AWS accounts by using AWS Control Tower. The company hosts internal applications and public applications.

Each application team in the company has its own AWS account for application hosting. The accounts are consolidated in an organization in AWS Organizations.

One of the AWS Control Tower member accounts serves as a centralized DevOps account with CI/CD pipelines that application teams use to deploy applications to their respective target AWS accounts. An IAM role for deployment exists in the centralized DevOps account.

An application team is attempting to deploy its application to an Amazon Elastic Kubernetes Service (Amazon EKS) cluster in an application AWS account. An IAM role for deployment exists in the application AWS account. The deployment is through an AWS CodeBuild project that is set up in the centralized DevOps account. The CodeBuild project uses an IAM service role for CodeBuild. The deployment is failing with an Unauthorized error during attempts to connect to the cross-account EKS cluster from CodeBuild.

Which solution will resolve this error?

- A. Configure the application account's deployment IAM role to have a trust relationship with the centralized DevOps account
- B. Configure the trust relationship to allow the sts:AssumeRole action
- C. Configure the application account's deployment IAM role to have the required access to the EKS cluster
- D. Configure the EKS cluster aws-auth ConfigMap to map the role to the appropriate system permissions.
- E. Configure the centralized DevOps account's deployment IAM role to have a trust relationship with the application account
- F. Configure the trust relationship to allow the sts:AssumeRole action
- G. Configure the centralized DevOps account's deployment IAM role to allow the required access to CodeBuild.
- H. Configure the centralized DevOps account's deployment IAM role to have a trust relationship with the application account
- I. Configure the trust relationship to allow the sts:AssumeRoleWithSAML action
- J. Configure the centralized DevOps account's deployment IAM role to allow the required access to CodeBuild.
- K. Configure the application account's deployment IAM role to have a trust relationship with the AWS Control Tower management account
- L. Configure the trust relationship to allow the sts:AssumeRole action
- M. Configure the application account's deployment IAM role to have the required access to the EKS cluster
- N. Configure the EKS cluster aws-auth ConfigMap to map the role to the appropriate system permissions.

Answer: A

Explanation:

In the source AWS account, the IAM role used by the CI/CD pipeline should have permissions to access the source code repository, build artifacts, and any other resources required for the build process. In the destination AWS accounts, the IAM role used for deployment should have permissions to access the AWS resources required for deploying the application, such as EC2 instances, RDS databases, S3 buckets, etc. The exact permissions required will depend on the specific resources being used by the application. The IAM role used for deployment in the destination accounts should also have permissions to assume the IAM role for deployment in the centralized DevOps account. This is typically done using an IAM role trust policy that allows the destination account to assume the DevOps account role.

NEW QUESTION 60

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

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

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

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

Answer: BD

Explanation:

The requirement is to update the infrastructure to ensure that only the Lambda function's execution role can access the values in Secrets Manager. The solution must apply the principle of least privilege, which means granting the minimum permissions necessary to perform a task.

To do this, the DevOps engineer needs to use the following steps:

? Create a KMS customer managed key that trusts Secrets Manager and allows the Lambda function's execution role to decrypt. A customer managed key is a symmetric encryption key that is fully managed by the customer. The customer can define the key policy, which specifies who can use and manage the key. By creating a customer managed key, the DevOps engineer can restrict the decryption permission to only the Lambda function's execution role, and prevent other principals from accessing the secret values. The customer managed key also needs to trust Secrets Manager, which means allowing Secrets Manager to use the key to encrypt and decrypt secrets on behalf of the customer.

? Update Secrets Manager to use the new customer managed key. Secrets Manager allows customers to choose which KMS key to use for encrypting each secret. By default, Secrets Manager uses the default KMS key for Secrets Manager, which is a service-managed key that is shared by all customers in the same AWS Region. By updating Secrets Manager to use the new customer managed key, the DevOps engineer can ensure that only the Lambda function's execution role can decrypt the secret values using that key.

? Ensure that the Lambda function's execution role has the KMS permissions scoped on the resource level. The Lambda function's execution role is an IAM role that grants permissions to the Lambda function to access AWS services and resources. The role needs to have KMS permissions to use the customer managed key for decryption. However, to apply the principle of least privilege, the role should have the permissions scoped on the resource level, which means specifying the ARN of the customer managed key as a condition in the IAM policy statement. This way, the role can only use that specific key and not any other KMS keys in the account.

NEW QUESTION 63

A DevOps engineer is architecting a continuous development strategy for a company's software as a service (SaaS) web application running on AWS. For application and security reasons users subscribing to this application are distributed across multiple Application Load Balancers (ALBs) each of which has a dedicated Auto Scaling group and fleet of Amazon EC2 instances. The application does not require a build stage and when it is committed to AWS CodeCommit, the application must trigger a simultaneous deployment to all ALBs Auto Scaling groups and EC2 fleets.

Which architecture will meet these requirements with the LEAST amount of configuration?

- A. Create a single AWS CodePipeline pipeline that deploys the application in parallel using unique AWS CodeDeploy applications and deployment groups created for each ALB-Auto Scaling group pair.
- B. Create a single AWS CodePipeline pipeline that deploys the application using a single AWS CodeDeploy application and single deployment group.

- C. Create a single AWS CodePipeline pipeline that deploys the application in parallel using a single AWS CodeDeploy application and unique deployment group for each ALB-Auto Scaling group pair.
- D. Create an AWS CodePipeline pipeline for each ALB-Auto Scaling group pair that deploys the application using an AWS CodeDeploy application and deployment group created for the same ALB-Auto Scaling group pair.

Answer: C

Explanation:

<https://docs.aws.amazon.com/codedeploy/latest/userguide/deployment-groups.html>

NEW QUESTION 64

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 69

A company manages an application that stores logs in Amazon CloudWatch Logs. The company wants to archive the logs to an Amazon S3 bucket. Logs are rarely accessed after 90 days and must be retained for 10 years. Which combination of steps should a DevOps engineer take to meet these requirements? (Select TWO.)

- A. Configure a CloudWatch Logs subscription filter to use AWS Glue to transfer all logs to an S3 bucket.
- B. Configure a CloudWatch Logs subscription filter to use Amazon Kinesis Data Firehose to stream all logs to an S3 bucket.
- C. Configure a CloudWatch Logs subscription filter to stream all logs to an S3 bucket.
- D. Configure the S3 bucket lifecycle policy to transition logs to S3 Glacier after 90 days and to expire logs after 3,650 days.
- E. Configure the S3 bucket lifecycle policy to transition logs to Reduced Redundancy after 90 days and to expire logs after 3,650 days.

Answer: BD

Explanation:

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

NEW QUESTION 70

An application running on a set of Amazon EC2 instances in an Auto Scaling group requires a configuration file to operate. The instances are created and maintained with AWS CloudFormation. A DevOps engineer wants the instances to have the latest configuration file when launched and wants changes to the configuration file to be reflected on all the instances with a minimal delay when the CloudFormation template is updated. Company policy requires that application configuration files be maintained along with AWS infrastructure configuration files in source control. Which solution will accomplish this?

- A. In the CloudFormation template add an AWS Config rule.
- B. Place the configuration file content in the rule's `InputParameters` property and set the `Scope` property to the EC2 Auto Scaling group.
- C. Add an AWS Systems Manager Resource Data Sync resource to the template to poll for updates to the configuration.
- D. In the CloudFormation template add an EC2 launch template resource.
- E. Place the configuration file content in the launch template.
- F. Configure the `cf-n-init` script to run when the instance is launched and configure the `cf-n-hup` script to poll for updates to the configuration.
- G. In the CloudFormation template add an EC2 launch template resource.
- H. Place the configuration file content in the launch template.
- I. Add an AWS Systems Manager Resource Data Sync resource to the template to poll for updates to the configuration.
- J. In the CloudFormation template add CloudFormation `Metadata`.
- K. Place the configuration file content in the `Metadata`.
- L. Configure the `cf-n-init` script to run when the instance is launched and configure the `cf-n-hup` script to poll for updates to the configuration.

Answer: D

Explanation:

Use the `AWS::CloudFormation::Init` type to include metadata on an Amazon EC2 instance for the `cf-n-init` helper script. If your template calls the `cf-n-init` script, the script looks for resource metadata rooted in the `AWS::CloudFormation::Init` metadata key. Reference:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-resource-init.html>

NEW QUESTION 75

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

Which combination of steps will meet these requirements? (Choose three.)

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

Answer: BCF

Explanation:

Using the `principalTag` in the Permission Set inline policy a logged in user belonging to a specific AD group in the IDP can be permitted access to perform operations on certain resources if their group matches the group used in the `PrincipalTag`. Basically you are narrowing the scope of privileges assigned via Permission policies conditionally based on whether the logged in user belongs to a specific AD Group in IDP. The mapping of the AD group to the request attributes can be done using SSO attributes where we can pass other attributes like the SAML token as well.

<https://docs.aws.amazon.com/singlesignon/latest/userguide/abac.html>

NEW QUESTION 77

A DevOps engineer has implemented a CI/CO pipeline to deploy an AWS CloudFormation template that provisions a web application. The web application consists of an Application Load Balancer (ALB) a target group, a launch template that uses an Amazon Linux 2 AMI an Auto Scaling group of Amazon EC2 instances, a security group and an Amazon RDS for MySQL database. The launch template includes user data that specifies a script to install and start the application.

The initial deployment of the application was successful. The DevOps engineer made changes to update the version of the application with the user data. The CI/CD pipeline has deployed a new version of the template. However, the health checks on the ALB are now failing. The health checks have marked all targets as unhealthy.

During investigation the DevOps engineer notices that the CloudFormation stack has a status of `UPDATE_COMPLETE`. However, when the DevOps engineer connects to one of the EC2 instances and checks `/var/log` messages, the DevOps engineer notices that the Apache web server failed to start successfully because of a configuration error.

How can the DevOps engineer ensure that the CloudFormation deployment will fail if the user data fails to successfully finish running?

- A. Use the `cfn-signal` helper script to signal success or failure to CloudFormation. Use the `WaitOnResourceSignals` update policy within the CloudFormation template. Set an appropriate timeout for the update policy.
- B. Create an Amazon CloudWatch alarm for the `UnhealthyHostCount` metric.
- C. Include an appropriate alarm threshold for the target group. Create an Amazon Simple Notification Service (Amazon SNS) topic as the target to signal success or failure to CloudFormation.
- D. Create a lifecycle hook on the Auto Scaling group by using the `AWS::AutoScaling::LifecycleHook` resource. Create an Amazon Simple Notification Service (Amazon SNS) topic as the target to signal success or failure to CloudFormation. Set an appropriate timeout on the lifecycle hook.
- E. Use the Amazon CloudWatch agent to stream the `cloud-init` logs. Create a subscription filter that includes an AWS Lambda function with an appropriate invocation timeout. Configure the Lambda function to use the `SignalResource` API operation to signal success or failure to CloudFormation.

Answer: A

Explanation:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/aws-attribute-updatepolicy.html>

NEW QUESTION 81

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

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

Answer: D

Explanation:

This solution will meet the requirements because it will use Amazon S3 as a common data lake for both the application logs and the API logs. Amazon S3 is a service that provides scalable, durable, and secure object storage for any type of data. You can use the Amazon CloudWatch agent to send logs from your EC2

instances to S3 buckets, and use AWS CloudTrail to deliver the API logs to S3 buckets as well. You can also use Amazon Athena to query both sets of logs in S3 using standard SQL, without loading or transforming them. Athena is a serverless interactive query service that allows you to analyze data in S3 using a variety of data formats, such as JSON, CSV, Parquet, and ORC.

NEW QUESTION 83

A DevOps engineer is building a continuous deployment pipeline for a serverless application that uses AWS Lambda functions. The company wants to reduce the customer impact of an unsuccessful deployment. The company also wants to monitor for issues.

Which deploy stage configuration will meet these requirements?

- A. Use an AWS Serverless Application Model (AWS SAM) template to define the serverless applicatio
- B. Use AWS CodeDeploy to deploy the Lambda functions with the Canary10Percent15Minutes Deployment Preference Typ
- C. Use Amazon CloudWatch alarms to monitor the health of the functions.
- D. Use AWS CloudFormation to publish a new stack update, and include Amazon CloudWatch alarms on all resource
- E. Set up an AWS CodePipeline approval action for a developer to verify and approve the AWS CloudFormation change set.
- F. Use AWS CloudFormation to publish a new version on every stack update, and include Amazon CloudWatch alarms on all resource
- G. Use the RoutingConfig property of the AWS::Lambda::Alias resource to update the traffic routing during the stack update.
- H. Use AWS CodeBuild to add sample event payloads for testing to the Lambda function
- I. Publish a new version of the functions, and include Amazon CloudWatch alarm
- J. Update the production alias to point to the new versio
- K. Configure rollbacks to occur when an alarm is in the ALARM state.

Answer: D

Explanation:

Use routing configuration on an alias to send a portion of traffic to a second function version. For example, you can reduce the risk of deploying a new version by configuring the alias to send most of the traffic to the existing version, and only a small percentage of traffic to the new version.

<https://docs.aws.amazon.com/lambda/latest/dg/configuration-aliases.html>

The following are the steps involved in the deploy stage configuration that will meet the requirements:

? Use AWS CodeBuild to add sample event payloads for testing to the Lambda functions.

? Publish a new version of the functions, and include Amazon CloudWatch alarms.

? Update the production alias to point to the new version.

? Configure rollbacks to occur when an alarm is in the ALARM state.

This configuration will help to reduce the customer impact of an unsuccessful deployment

by deploying the new version of the functions to a staging environment first. This will allow the DevOps engineer to test the new version of the functions before deploying it to production.

The configuration will also help to monitor for issues by including Amazon CloudWatch alarms. These alarms will alert the DevOps engineer if there are any problems with the new version of the functions.

NEW QUESTION 87

A company needs to ensure that flow logs remain configured for all existing and new VPCs in its AWS account. The company uses an AWS CloudFormation stack to manage its VPCs. The company needs a solution that will work for any VPCs that any IAM user creates.

Which solution will meet these requirements?

- A. Add the resource to the CloudFormation stack that creates the VPCs.
- B. Create an organization in AWS Organization
- C. Add the company's AWS account to the organizatio
- D. Create an SCP to prevent users from modifying VPC flow logs.
- E. Turn on AWS Confi
- F. Create an AWS Config rule to check whether VPC flow logs are turned o
- G. Configure automatic remediation to turn on VPC flow logs.
- H. Create an IAM policy to deny the use of API calls for VPC flow log
- I. Attach the IAM policy to all IAM users.

Answer: C

Explanation:

To meet the requirements of ensuring that flow logs remain configured for all existing and new VPCs in the AWS account, the company should use AWS Config and automatic remediation. AWS Config is a service that enables customers to assess, audit, and evaluate the configurations of their AWS resources. AWS Config continuously monitors and records the configuration changes of the AWS resources and evaluates them against desired configurations. Customers can use AWS Config rules to define the desired configuration state of their AWS resources and trigger actions when a resource configuration violates a rule.

One of the AWS Config rules that customers can use is vpc-flow-logs-enabled, which checks whether VPC flow logs are enabled for all VPCs in an AWS account. Customers can also configure automatic remediation for this rule, which means that AWS Config will automatically enable VPC flow logs for any VPCs that do not have them enabled. Customers can specify the destination (CloudWatch Logs or S3) and the traffic type (all, accept, or reject) for the flow logs as remediation parameters. By using AWS Config and automatic remediation, the company can ensure that flow logs remain configured for all existing and new VPCs in its AWS account, regardless of who creates them or how they are created.

The other options are not correct because they do not meet the requirements or follow best practices. Adding the resource to the CloudFormation stack that creates the VPCs is not a sufficient solution because it will only work for VPCs that are created by using the CloudFormation stack. It will not work for VPCs that are created by using other methods, such as the console or the API. Creating an organization in AWS Organizations and creating an SCP to prevent users from modifying VPC flow logs is not a good solution because it will not ensure that flow logs are enabled for all VPCs in the first place. It will only prevent users from disabling or changing flow logs after they are enabled. Creating an IAM policy to deny the use of API calls for VPC flow logs and attaching it to all IAM users is not a valid solution because it will prevent users from enabling or disabling flow logs at all.

It will also not work for VPCs that are created by using other methods, such as the console or CloudFormation.

References:

? 1: AWS::EC2::FlowLog - AWS CloudFormation

? 2: Amazon VPC Flow Logs extends CloudFormation Support to custom format subscriptions, 1-minute aggregation intervals and tagging

? 3: Logging IP traffic using VPC Flow Logs - Amazon Virtual Private Cloud

? : About AWS Config - AWS Config

? : vpc-flow-logs-enabled - AWS Config

? : Remediate Noncompliant Resources with AWS Config Rules - AWS Config

NEW QUESTION 89

A company runs its container workloads in AWS App Runner. A DevOps engineer manages the company's container repository in Amazon Elastic Container Registry (Amazon ECR).

The DevOps engineer must implement a solution that continuously monitors the container repository. The solution must create a new container image when the solution detects an operating system vulnerability or language package vulnerability.

Which solution will meet these requirements?

- A. Use EC2 Image Builder to create a container image pipeline
- B. Use Amazon ECR as the target repository
- C. Turn on enhanced scanning on the ECR repository
- D. Create an Amazon EventBridge rule to capture an Inspector2 finding event
- E. Use the event to invoke the image pipeline
- F. Re-upload the container to the repository.
- G. Use EC2 Image Builder to create a container image pipeline
- H. Use Amazon ECR as the target repository
- I. Enable Amazon GuardDuty Malware Protection on the container workload
- J. Create an Amazon EventBridge rule to capture a GuardDuty finding event
- K. Use the event to invoke the image pipeline.
- L. Create an AWS CodeBuild project to create a container image
- M. Use Amazon ECR as the target repository
- N. Turn on basic scanning on the repository
- O. Create an Amazon EventBridge rule to capture an ECR image action event
- P. Use the event to invoke the CodeBuild project
- Q. Re-upload the container to the repository.
- R. Create an AWS CodeBuild project to create a container image
- S. Use Amazon ECR as the target repository
- T. Configure AWS Systems Manager Compliance to scan all managed nodes
- . Create an Amazon EventBridge rule to capture a configuration compliance state change event
- . Use the event to invoke the CodeBuild project.

Answer: A

Explanation:

The solution that meets the requirements is to use EC2 Image Builder to create a container image pipeline, use Amazon ECR as the target repository, turn on enhanced scanning on the ECR repository, create an Amazon EventBridge rule to capture an Inspector2 finding event, and use the event to invoke the image pipeline. Re-upload the container to the repository.

This solution will continuously monitor the container repository for vulnerabilities using enhanced scanning, which is a feature of Amazon ECR that provides detailed information and guidance on how to fix security issues found in your container images. Enhanced scanning uses Inspector2, a security assessment service that integrates with Amazon ECR and generates findings for any vulnerabilities detected in your images. You can use Amazon EventBridge to create a rule that triggers an action when an Inspector2 finding event occurs. The action can be to invoke an EC2 Image Builder pipeline, which is a service that automates the creation of container images. The pipeline can use the latest patches and updates to build a new container image and upload it to the same ECR repository, replacing the vulnerable image.

The other options are not correct because they do not meet all the requirements or use services that are not relevant for the scenario.

Option B is not correct because it uses Amazon GuardDuty Malware Protection, which is a feature of GuardDuty that detects malicious activity and unauthorized behavior on your AWS accounts and resources. GuardDuty does not scan container images for vulnerabilities, nor does it integrate with Amazon ECR or EC2 Image Builder.

Option C is not correct because it uses basic scanning on the ECR repository, which only provides a summary of the vulnerabilities found in your container images. Basic scanning does not use Inspector2 or generate findings that can be captured by Amazon EventBridge. Moreover, basic scanning does not provide guidance on how to fix the vulnerabilities.

Option D is not correct because it uses AWS Systems Manager Compliance, which is a feature of Systems Manager that helps you monitor and manage the compliance status of your AWS resources based on AWS Config rules and AWS Security Hub standards. Systems Manager Compliance does not scan container images for vulnerabilities, nor does it integrate with Amazon ECR or EC2 Image Builder.

NEW QUESTION 90

A company uses AWS Organizations to manage its AWS accounts. The company has a root OU that has a child OU. The root OU has an SCP that allows all actions on all resources. The child OU has an SCP that allows all actions for Amazon DynamoDB and AWS Lambda, and denies all other actions.

The company has an AWS account that is named vendor-data in the child OU. A DevOps engineer has an IAM user that is attached to the AdministratorAccess IAM policy in the vendor-data account. The DevOps engineer attempts to launch an Amazon EC2 instance in the vendor-data account but receives an access denied error.

Which change should the DevOps engineer make to launch the EC2 instance in the vendor-data account?

- A. Attach the AmazonEC2FullAccess IAM policy to the IAM user.
- B. Create a new SCP that allows all actions for Amazon EC2. Attach the SCP to the vendor-data account.
- C. Update the SCP in the child OU to allow all actions for Amazon EC2.
- D. Create a new SCP that allows all actions for Amazon EC2. Attach the SCP to the root OU.

Answer: C

Explanation:

The correct answer is C. Updating the SCP in the child OU to allow all actions for Amazon EC2 will enable the DevOps engineer to launch the EC2 instance in the vendor-data account. SCPs are applied to OUs and accounts in a hierarchical manner, meaning that the SCPs attached to the parent OU are inherited by the child OU and accounts. Therefore, the SCP in the child OU overrides the SCP in the root OU and denies all actions except for DynamoDB and Lambda. By adding EC2 to the allowed actions in the child OU's SCP, the DevOps engineer can access EC2 resources in the vendor-data account.

Option A is incorrect because attaching the AmazonEC2FullAccess IAM policy to the IAM user will not grant the user access to EC2 resources. IAM policies are evaluated after SCPs, so even if the IAM policy allows EC2 actions, the SCP will still deny them.

Option B is incorrect because creating a new SCP that allows all actions for EC2 and attaching it to the vendor-data account will not work. SCPs are not cumulative, meaning that only one SCP is applied to an account at a time. The SCP attached to the account will be the SCP attached to the OU that contains the account. Therefore, option B will not change the SCP that is applied to the vendor-data account.

Option D is incorrect because creating a new SCP that allows all actions for EC2 and attaching it to the root OU will not work. As explained earlier, the SCP in the child OU overrides the SCP in the root OU and denies all actions except for DynamoDB and Lambda. Therefore, option D will not affect the SCP that is applied to the vendor-data account.

NEW QUESTION 92

A company's security policies require the use of security hardened AMIs in production environments. A DevOps engineer has used EC2 Image Builder to create a pipeline that builds the AMIs on a recurring schedule.

The DevOps engineer needs to update the launch templates of the company's Auto Scaling groups. The Auto Scaling groups must use the newest AMIs during the launch of Amazon EC2 instances.

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

- A. Configure an Amazon EventBridge rule to receive new AMI events from Image Builder
- B. Target an AWS Systems Manager Run Command document that updates the launch templates of the Auto Scaling groups with the newest AMI ID.
- C. Configure an Amazon EventBridge rule to receive new AMI events from Image Builder
- D. Target an AWS Lambda function that updates the launch templates of the Auto Scaling groups with the newest AMI ID.
- E. Configure the launch template to use a value from AWS Systems Manager Parameter Store for the AMI ID
- F. Configure the Image Builder pipeline to update the Parameter Store value with the newest AMI ID.
- G. Configure the Image Builder distribution settings to update the launch templates with the newest AMI ID
- H. Configure the Auto Scaling groups to use the newest version of the launch template.

Answer: C

Explanation:

? The most operationally efficient solution is to use AWS Systems Manager Parameter Store¹ to store the AMI ID and reference it in the launch template². This way, the launch template does not need to be updated every time a new AMI is created by Image Builder. Instead, the Image Builder pipeline can update the Parameter Store value with the newest AMI ID³, and the Auto Scaling group can launch instances using the latest value from Parameter Store.

? The other solutions require updating the launch template or creating a new version of it every time a new AMI is created, which adds complexity and overhead. Additionally, using EventBridge rules and Lambda functions or Run Command documents introduces additional dependencies and potential points of failure.

References: 1: AWS Systems Manager Parameter Store 2: Using AWS Systems Manager parameters instead of AMI IDs in launch templates 3: Update an SSM parameter with Image Builder

NEW QUESTION 97

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

Which action will accomplish this?

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

Answer: B

Explanation:

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

NEW QUESTION 100

A media company has several thousand Amazon EC2 instances in an AWS account. The company is using Slack and a shared email inbox for team communications and important updates. A DevOps engineer needs to send all AWS-scheduled EC2 maintenance notifications to the Slack channel and the shared inbox. The solution must include the instances' Name and Owner tags.

Which solution will meet these requirements?

- A. Integrate AWS Trusted Advisor with AWS Config Configure a custom AWS Config rule to invoke an AWS Lambda function to publish notifications to an Amazon Simple Notification Service (Amazon SNS) topic Subscribe a Slack channel endpoint and the shared inbox to the topic.
- B. Use Amazon EventBridge to monitor for AWS Health Events Configure the maintenance events to target an Amazon Simple Notification Service (Amazon SNS) topic Subscribe an AWS Lambda function to the SNS topic to send notifications to the Slack channel and the shared inbox.
- C. Create an AWS Lambda function that sends EC2 maintenance notifications to the Slack channel and the shared inbox Monitor EC2 health events by using Amazon CloudWatch metrics Configure a CloudWatch alarm that invokes the Lambda function when a maintenance notification is received.
- D. Configure AWS Support integration with AWS CloudTrail Create a CloudTrail lookup event to invoke an AWS Lambda function to pass EC2 maintenance notifications to Amazon Simple Notification Service (Amazon SNS) Configure Amazon SNS to target the Slack channel and the shared inbox.

Answer: B

Explanation:

<https://docs.aws.amazon.com/health/latest/ug/cloudwatch-events-health.html>

NEW QUESTION 104

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 organization
- 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 108

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

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

Answer: B

Explanation:

To implement failover for the application to the secondary Region so that HTTP GET requests meet the desired RTO, the DevOps engineer should use the following solution:

? Create a new origin on the distribution for the secondary ALB. A CloudFront origin is the source of the content that CloudFront delivers to viewers. By creating a new origin for the secondary ALB, the DevOps engineer can configure CloudFront to route traffic to the secondary Region when the primary Region is unavailable¹

? Create a new origin group. Set the original ALB as the primary origin. Configure the origin group to fail over for HTTP 5xx status codes. An origin group is a logical grouping of two origins: a primary origin and a secondary origin. By creating an origin group, the DevOps engineer can specify which origin CloudFront should use as a fallback when the primary origin fails. The DevOps engineer can also define which HTTP status codes should trigger a failover from the primary origin to the secondary origin. By setting the original ALB as the primary origin and configuring the origin group to fail over for HTTP 5xx status codes, the DevOps engineer can ensure that CloudFront will switch to the secondary ALB when the primary ALB returns server errors²

? Update the default behavior to use the origin group. A behavior is a set of rules that CloudFront applies when it receives requests for specific URLs or file types. The default behavior applies to all requests that do not match any other behaviors. By updating the default behavior to use the origin group, the DevOps engineer can enable failover routing for all requests that are sent to the distribution³

This solution will meet the requirements because it will automate the failover of the application to the secondary Region with zero-second RTO. When CloudFront receives an HTTP GET request, it will first try to route it to the primary ALB in the primary Region. If the primary ALB is healthy and returns a successful response, CloudFront will deliver it to the viewer. If the primary ALB is unhealthy or returns an HTTP 5xx status code, CloudFront will automatically route the request to the secondary ALB in the secondary Region and deliver its response to the viewer. The other options are not correct because they either do not provide zero-second RTO or do not work as expected. Creating a second CloudFront distribution that has the secondary ALB as the default origin and creating Amazon Route 53 alias records that have a failover policy is not a good option because it will introduce additional latency and complexity to the solution. Route 53 health checks and DNS propagation can take several minutes or longer, which means that viewers might experience delays or errors when accessing the application during a failover event. Creating Amazon Route 53 alias records that have a failover policy and Evaluate Target Health set to Yes for both ALBs and setting the TTL of both records to 0 is not a valid option because it will not work with CloudFront distributions. Route 53 does not support health checks for alias records that point to CloudFront distributions, so it cannot detect if an ALB behind a distribution is healthy or not. Creating a CloudFront function that detects HTTP 5xx status codes and returns a 307 Temporary Redirect error response to the secondary ALB is not a valid option because it will not provide zero-second RTO. A 307 Temporary Redirect error response tells viewers to retry their requests with a different URL, which means that viewers will have to make an additional request and wait for another response from CloudFront before reaching the secondary ALB.

References:

- ? 1: Adding, Editing, and Deleting Origins - Amazon CloudFront
- ? 2: Configuring Origin Failover - Amazon CloudFront
- ? 3: Creating or Updating a Cache Behavior - Amazon CloudFront

NEW QUESTION 112

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

The team uses a developer IAM role to access the environment. The role is configured with the AdministratorAccess managed policy. The company has created a new CloudFormationDeployment IAM role that has the following policy.

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

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

- A. Remove the AdministratorAccess polic
- B. Assign the ReadOnlyAccess managed IAM policy to the developer rol
- C. Instruct the developers to use the CloudFormationDeployment role as a CloudFormation service role when the developers deploy new stacks.
- D. Update the trust of CloudFormationDeployment role to allow the developer IAM role to assume the CloudFormationDepoyment role.
- E. Configure the IAM to be to get and pass the CloudFormationDeployment role if cloudformation actions for resources,
- F. Update the trust Of the CloudFormationDepoyment role to anow the cloudformation.amazonaws.com AWS principal to perform the iam:AssumeR01e action
- G. Remove me Administratoraccess polic
- H. Assign the ReadOnly/Access managed IAM policy to the developer role Instruct the developers to assume the CloudFormatondeployment role when the developers new stacks
- I. Add an IAM policy to CloudFormationDeplyment 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 114

A company recently migrated its legacy application from on-premises to AWS. The application is hosted on Amazon EC2 instances behind an Application Load Balancer which is behind Amazon API Gateway. The company wants to ensure users experience minimal disruptions during any deployment of a new version of the application. The company also wants to ensure it can quickly roll back updates if there is an issue.

Which solution will meet these requirements with MINIMAL changes to the application?

- A. Introduce changes as a separate environment parallel to the existing one Configure API Gateway to use a canary release deployment to send a small subset of user traffic to the new environment.
- B. Introduce changes as a separate environment parallel to the existing one Update the application's DNS alias records to point to the new environment.
- C. Introduce changes as a separate target group behind the existing Application Load Balancer Configure API Gateway to route user traffic to the new target group in steps.
- D. Introduce changes as a separate target group behind the existing Application Load Balancer Configure API Gateway to route all traffic to the Application Load Balancer which then sends the traffic to the new target group.

Answer: A

Explanation:

API Gateway supports canary deployment on a deployment stage before you direct all traffic to that stage. A parallel environment means we will create a new ALB and a target group that will target a new set of EC2 instances on which the newer version of the app will be deployed. So the canary setting associated to the new version of the API will connect with the new ALB instance which in turn will direct the traffic to the new EC2 instances on which the newer version of the application is deployed.

NEW QUESTION 115

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

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

Answer: C

Explanation:

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

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

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

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

NEW QUESTION 118

A DevOps team manages an API running on-premises that serves as a backend for an Amazon API Gateway endpoint. Customers have been complaining about high response latencies, which the development team has verified using the API Gateway latency metrics in Amazon CloudWatch. To identify the cause, the team needs to collect relevant data without introducing additional latency. Which actions should be taken to accomplish this? (Choose two.)

- A. Install the CloudWatch agent server side and configure the agent to upload relevant logs to CloudWatch.
- B. Enable AWS X-Ray tracing in API Gateway, modify the application to capture request segments, and upload those segments to X-Ray during each request.
- C. Enable AWS X-Ray tracing in API Gateway, modify the application to capture request segments, and use the X-Ray daemon to upload segments to X-Ray.
- D. Modify the on-premises application to send log information back to API Gateway with each request.
- E. Modify the on-premises application to calculate and upload statistical data relevant to the API service requests to CloudWatch metrics.

Answer: AC

Explanation:

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/install-CloudWatch-Agent-on-premise.html>
<https://docs.aws.amazon.com/xray/latest/devguide/xray-api-sendingdata.html>

NEW QUESTION 120

A company has chosen AWS to host a new application. The company needs to implement a multi-account strategy. A DevOps engineer creates a new AWS account and an organization in AWS Organizations. The DevOps engineer also creates the OU structure for the organization and sets up a landing zone by using AWS Control Tower.

The DevOps engineer must implement a solution that automatically deploys resources for new accounts that users create through AWS Control Tower Account Factory. When a user creates a new account, the solution must apply AWS CloudFormation templates and SCPs that are customized for the OU or the account to automatically deploy all the resources that are attached to the account. All the OUs are enrolled in AWS Control Tower.

Which solution will meet these requirements in the MOST automated way?

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

Answer: D

Explanation:

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

NEW QUESTION 125

.....

Relate Links

100% Pass Your AWS-Certified-DevOps-Engineer-Professional Exam with ExamBible Prep Materials

<https://www.exambible.com/AWS-Certified-DevOps-Engineer-Professional-exam/>

Contact us

We are proud of our high-quality customer service, which serves you around the clock 24/7.

Viste - <https://www.exambible.com/>