

AWS-Certified-DevOps-Engineer-Professional Dumps

Amazon AWS Certified DevOps Engineer Professional

<https://www.certleader.com/AWS-Certified-DevOps-Engineer-Professional-dumps.html>



NEW QUESTION 1

A company has an application that runs on AWS Lambda and sends logs to Amazon CloudWatch Logs. An Amazon Kinesis data stream is subscribed to the log groups in CloudWatch Logs. A single consumer Lambda function processes the logs from the data stream and stores the logs in an Amazon S3 bucket. The company's DevOps team has noticed high latency during the processing and ingestion of some logs. Which combination of steps will reduce the latency? (Select THREE.)

- A. Create a data stream consumer with enhanced fan-out
- B. Set the Lambda function that processes the logs as the consumer.
- C. Increase the ParallelizationFactor setting in the Lambda event source mapping.
- D. Configure reserved concurrency for the Lambda function that processes the logs.
- E. Increase the batch size in the Kinesis data stream.
- F. Turn off the ReportBatchItemFailures setting in the Lambda event source mapping.
- G. Increase the number of shards in the Kinesis data stream.

Answer: ABC

Explanation:

The latency in processing and ingesting logs can be caused by several factors, such as the throughput of the Kinesis data stream, the concurrency of the Lambda function, and the configuration of the event source mapping. To reduce the latency, the following steps can be taken:

? Create a data stream consumer with enhanced fan-out. Set the Lambda function that processes the logs as the consumer. This will allow the Lambda function to receive records from the data stream with dedicated throughput of up to 2 MB per second per shard, independent of other consumers¹. This will reduce the contention and delay in accessing the data stream.

? Increase the ParallelizationFactor setting in the Lambda event source mapping. This will allow the Lambda service to invoke more instances of the function concurrently to process the records from the data stream². This will increase the processing capacity and reduce the backlog of records in the data stream.

? Configure reserved concurrency for the Lambda function that processes the logs. This will ensure that the function has enough concurrency available to handle the increased load from the data stream³. This will prevent the function from being throttled by the account-level concurrency limit.

The other options are not effective or may have negative impacts on the latency. Option D is not suitable because increasing the batch size in the Kinesis data stream will increase the amount of data that the Lambda function has to process in each invocation, which may increase the execution time and latency⁴. Option E is not advisable because turning off the ReportBatchItemFailures setting in the Lambda event source mapping will prevent the Lambda service from retrying the failed records, which may result in data loss. Option F is not necessary because increasing the number of shards in the Kinesis data stream will increase the throughput of the data stream, but it will not affect the processing speed of the Lambda function, which is the bottleneck in this scenario.

References:

? 1: Using AWS Lambda with Amazon Kinesis Data Streams - AWS Lambda

? 2: AWS Lambda event source mappings - AWS Lambda

? 3: Managing concurrency for a Lambda function - AWS Lambda

? 4: AWS Lambda function scaling - AWS Lambda

? : AWS Lambda event source mappings - AWS Lambda

? : Scaling Amazon Kinesis Data Streams with AWS CloudFormation - Amazon Kinesis Data Streams

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

Answer: B

Explanation:

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

NEW QUESTION 3

A company is adopting AWS CodeDeploy to automate its application deployments for a Java-Apache Tomcat application with an Apache Webserver. The development team started with a proof of concept, created a deployment group for a developer environment, and performed functional tests within the application. After completion, the team will create additional deployment groups for staging and production.

The current log level is configured within the Apache settings, but the team wants to change this configuration dynamically when the deployment occurs, so that they can set different log level configurations depending on the deployment group without having a different application revision for each group.

How can these requirements be met with the LEAST management overhead and without requiring different script versions for each deployment group?

- A. Tag the Amazon EC2 instances depending on the deployment group
- B. Then place a script into the application revision that calls the metadata service and the EC2 API to identify which deployment group the instance is part of
- C. Use this information to configure the log level setting
- D. Reference the script as part of the AfterInstall lifecycle hook in the appspec.yml file.
- E. Create a script that uses the CodeDeploy environment variable DEPLOYMENT_GROUP_NAME to identify which deployment group the instance is part of
- F. Use this information to configure the log level setting
- G. Reference this script as part of the BeforeInstall lifecycle hook in the appspec.yml file.
- H. Create a CodeDeploy custom environment variable for each environment
- I. Then place a script into the application revision that checks this environment variable to identify which deployment group the instance is part of
- J. Use this information to configure the log level setting
- K. Reference this script as part of the ValidateService lifecycle hook in the appspec.yml file.
- L. Create a script that uses the CodeDeploy environment variable DEPLOYMENT_GROUP_ID to identify which deployment group the instance is part of to configure the log level setting
- M. Reference this script as part of the Install lifecycle hook in the appspec.yml file.

Answer: B

Explanation:

The following are the steps that the company can take to change the log level dynamically when the deployment occurs:

? Create a script that uses the CodeDeploy environment variable `DEPLOYMENT_GROUP_NAME` to identify which deployment group the instance is part of.

? Use this information to configure the log level settings.

? Reference this script as part of the `BeforeInstall` lifecycle hook in the `appspec.yml` file.

The `DEPLOYMENT_GROUP_NAME` environment variable is automatically set by CodeDeploy when the deployment is triggered. This means that the script does not need to call the metadata service or the EC2 API to identify the deployment group.

This solution is the least complex and requires the least management overhead. It also does not require different script versions for each deployment group.

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

? Option A is incorrect because it would require tagging the Amazon EC2 instances, which would be a manual and time-consuming process.

? Option C is incorrect because it would require creating a custom environment variable for each environment. This would be a complex and error-prone process.

? Option D is incorrect because it would use

the `DEPLOYMENT_GROUP_ID` environment variable. However, this variable is not automatically set by CodeDeploy, so the script would need to call the metadata service or the EC2 API to get the deployment group ID. This would add complexity and overhead to the solution.

NEW QUESTION 4

A company uses a single AWS account to test applications on Amazon EC2 instances. The company has turned on AWS Config in the AWS account and has activated the restricted-ssh AWS Config managed rule.

The company needs an automated monitoring solution that will provide a customized notification in real time if any security group in the account is not compliant with the restricted-ssh rule. The customized notification must contain the name and ID of the noncompliant security group.

A DevOps engineer creates an Amazon Simple Notification Service (Amazon SNS) topic in the account and subscribes the appropriate personnel to the topic.

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

- A. Create an Amazon EventBridge rule that matches an AWS Config evaluation result of `NON_COMPLIANT` for the restricted-ssh rule
- B. Configure an input transformer for the EventBridge rule Configure the EventBridge rule to publish a notification to the SNS topic.
- C. Configure AWS Config to send all evaluation results for the restricted-ssh rule to the SNS topic
- D. Configure a filter policy on the SNS topic to send only notifications that contain the text of `NON_COMPLIANT` in the notification to subscribers.
- E. Create an Amazon EventBridge rule that matches an AWS Config evaluation result of `NON_COMPLIANT` for the restricted-ssh rule Configure the EventBridge rule to invoke AWS Systems Manager Run Command on the SNS topic to customize a notification and to publish the notification to the SNS topic
- F. Create an Amazon EventBridge rule that matches all AWS Config evaluation results of `NON_COMPLIANT` Configure an input transformer for the restricted-ssh rule Configure the EventBridge rule to publish a notification to the SNS topic.

Answer: A

Explanation:

Create an Amazon EventBridge (Amazon CloudWatch Events) rule that matches an AWS Config evaluation result of `NON_COMPLIANT` for the restricted-ssh rule. Configure an input transformer for the EventBridge (CloudWatch Events) rule. Configure the EventBridge (CloudWatch Events) rule to publish a notification to the SNS topic. This approach uses Amazon EventBridge (previously known as Amazon CloudWatch Events) to filter AWS Config evaluation results based on the restricted-ssh rule and its compliance status (`NON_COMPLIANT`). An input transformer can be used to customize the information contained in the notification, such as the name and ID of the noncompliant security group. The EventBridge (CloudWatch Events) rule can then be configured to publish a notification to the SNS topic, which will notify the appropriate personnel in real-time.

NEW QUESTION 5

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

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

Which solution will meet these requirements?

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

Answer: A

Explanation:

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

References:

? AWS CodeCommit User Guide on resilience and disaster recovery.

? AWS Documentation on monitoring CodeCommit events in Amazon EventBridge and Amazon CloudWatch Events2.

NEW QUESTION 6

A company runs an application with an Amazon EC2 and on-premises configuration. A DevOps engineer needs to standardize patching across both environments. Company policy dictates that patching only happens during non-business hours. Which combination of actions will meet these requirements? (Choose three.)

- A. Add the physical machines into AWS Systems Manager using Systems Manager Hybrid Activations.
- B. Attach an IAM role to the EC2 instances, allowing them to be managed by AWS Systems Manager.
- C. Create IAM access keys for the on-premises machines to interact with AWS Systems Manager.
- D. Run an AWS Systems Manager Automation document to patch the systems every hour.
- E. Use Amazon EventBridge scheduled events to schedule a patch window.
- F. Use AWS Systems Manager Maintenance Windows to schedule a patch window.

Answer: ABF

Explanation:

<https://docs.aws.amazon.com/systems-manager/latest/userguide/sysman-managed-instance-activation.html>

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 is using an Amazon Aurora cluster as the data store for its application. The Aurora cluster is configured with a single DB instance. The application performs read and write operations on the database by using the cluster's instance endpoint. The company has scheduled an update to be applied to the cluster during an upcoming maintenance window. The cluster must remain available with the least possible interruption during the maintenance window. What should a DevOps engineer do to meet these requirements?

- A. Add a reader instance to the Aurora cluster
- B. Update the application to use the Aurora cluster endpoint for write operation
- C. Update the Aurora cluster's reader endpoint for reads.
- D. Add a reader instance to the Aurora cluster
- E. Create a custom ANY endpoint for the cluster
- F. Update the application to use the Aurora cluster's custom ANY endpoint for read and write operations.
- G. Turn on the Multi-AZ option on the Aurora cluster
- H. Update the application to use the Aurora cluster endpoint for write operation
- I. Update the Aurora cluster's reader endpoint for reads.
- J. Turn on the Multi-AZ option on the Aurora cluster
- K. Create a custom ANY endpoint for the cluster
- L. Update the application to use the Aurora cluster's custom ANY endpoint for read and write operations.

Answer: C

Explanation:

To meet the requirements, the DevOps engineer should do the following:

? Turn on the Multi-AZ option on the Aurora cluster.

? Update the application to use the Aurora cluster endpoint for write operations.

? Update the Aurora cluster's reader endpoint for reads.

Turning on the Multi-AZ option will create a replica of the database in a different Availability Zone. This will ensure that the database remains available even if one of the Availability Zones is unavailable.

Updating the application to use the Aurora cluster endpoint for write operations will ensure that all writes are sent to both the primary and replica databases. This will ensure that the data is always consistent.

Updating the Aurora cluster's reader endpoint for reads will allow the application to read data from the replica database. This will improve the performance of the application during the maintenance window.

NEW QUESTION 9

A development team uses AWS CodeCommit, AWS CodePipeline, and AWS CodeBuild to develop and deploy an application. Changes to the code are submitted by pull requests. The development team reviews and merges the pull requests, and then the pipeline builds and tests the application.

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

Which solution will meet these requirements?

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

Answer: B

Explanation:

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

NEW QUESTION 10

A company is storing 100 GB of log data in csv format in an Amazon S3 bucket. SQL developers want to query this data and generate graphs to visualize it. The SQL developers also need an efficient automated way to store metadata from the csv file.

Which combination of steps will meet these requirements with the LEAST amount of effort? (Select THREE.)

- A. Filter the data through AWS X-Ray to visualize the data.
- B. Filter the data through Amazon QuickSight to visualize the data.
- C. Query the data with Amazon Athena.
- D. Query the data with Amazon Redshift.
- E. Use the AWS Glue Data Catalog as the persistent metadata store.
- F. Use Amazon DynamoDB as the persistent metadata store.

Answer: BCE

Explanation:

<https://docs.aws.amazon.com/glue/latest/dg/components-overview.html>

NEW QUESTION 10

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

The firewall appliance sends logs to Amazon CloudWatch Logs and includes event severities of CRITICAL, HIGH, MEDIUM, LOW, and INFO. The security team wants to receive an alert if any CRITICAL events occur.

What should the DevOps engineer do to meet these requirements?

- A. Create an Amazon CloudWatch Synthetics canary to monitor the firewall status
- B. If the firewall reaches a CRITICAL state or logs a CRITICAL event use a CloudWatch alarm to publish a notification to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the security team's email address to the topic.
- C. Create an Amazon CloudWatch metric filter by using a search for CRITICAL events. Publish a custom metric for the findings.
- D. Use a CloudWatch alarm based on the custom metric to publish a notification to an Amazon Simple Notification Service (Amazon SNS) topic.
- E. Subscribe the security team's email address to the topic.
- F. Enable Amazon GuardDuty in the network operations account.
- G. Configure GuardDuty to monitor flow logs. Create an Amazon EventBridge event rule that is invoked by GuardDuty events that are CRITICAL. Define an Amazon Simple Notification Service (Amazon SNS) topic as a target. Subscribe the security team's email address to the topic.
- H. Use AWS Firewall Manager to apply consistent policies across all accounts.
- I. Create an Amazon EventBridge event rule that is invoked by Firewall Manager events that are CRITICAL. Define an Amazon Simple Notification Service (Amazon SNS) topic as a target. Subscribe the security team's email address to the topic.
- J. EventBridge event rule that is invoked by Firewall Manager events that are CRITICAL. Define an Amazon Simple Notification Service (Amazon SNS) topic as a target. Subscribe the security team's email address to the topic.

Answer: B

Explanation:

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

NEW QUESTION 11

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 in 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 load 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 12

A company uses a series of individual Amazon CloudFormation templates to deploy its multi-Region Applications. These templates must be deployed in a specific order. The company is making more changes to the templates than previously expected and wants to deploy new templates more efficiently. Additionally, the data engineering team must be notified of all changes to the templates.

What should the company do to accomplish these goals?

- A. Create an AWS Lambda function to deploy the CloudFormation templates in the required order. Use stack policies to alert the data engineering team.
- B. Host the CloudFormation templates in Amazon S3. Use Amazon S3 events to directly trigger CloudFormation updates and Amazon SNS notifications.
- C. Implement CloudFormation StackSets and use drift detection to trigger update alerts to the data engineering team.
- D. Leverage CloudFormation nested stacks and stack sets (or deployments). Use Amazon SNS to notify the data engineering team.

Answer: D

Explanation:

This solution will meet the requirements because it will use CloudFormation nested stacks and stack sets to deploy the templates more efficiently and consistently across multiple regions. Nested stacks allow the company to separate out common components and reuse templates, while stack sets allow the company to create stacks in multiple accounts and regions with a single template. The company can also use Amazon SNS to send notifications to the data engineering team whenever a change is made to the templates or the stacks. Amazon SNS is a service that allows you to publish messages to subscribers, such as email addresses, phone numbers, or other AWS services. By using Amazon SNS, the company can ensure that the data engineering team is aware of all changes to the templates and can take appropriate actions if needed. What is Amazon SNS? - Amazon Simple Notification Service

NEW QUESTION 17

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 19

A development team uses AWS CodeCommit for version control for applications. The development team uses AWS CodePipeline, AWS CodeBuild, and AWS CodeDeploy for CI/CD infrastructure. In CodeCommit, the development team recently merged pull requests that did not pass long-running tests in the code base. The development team needed to perform rollbacks to branches in the codebase, resulting in lost time and wasted effort.

A DevOps engineer must automate testing of pull requests in CodeCommit to ensure that reviewers more easily see the results of automated tests as part of the pull request review.

What should the DevOps engineer do to meet this requirement?

- A. Create an Amazon EventBridge rule that reacts to the pullRequestStatusChanged event

- B. Create an AWS Lambda function that invokes a CodePipeline pipeline with a CodeBuild action that runs the tests for the applicatio
- C. Program the Lambda function to post the CodeBuild badge as a comment on the pull request so that developers will see the badge in their code review.
- D. Create an Amazon EventBridge rule that reacts to the pullRequestCreated event
- E. Create an AWS Lambda function that invokes a CodePipeline pipeline with a CodeBuild action that runs the tests for the applicatio
- F. Program the Lambda function to post the CodeBuild test results as a comment on the pull request when the test results are complete.
- G. Create an Amazon EventBridge rule that reacts to pullRequestCreated and pullRequestSourceBranchUpdated event
- H. Create an AWS Lambda function that invokes a CodePipeline pipeline with a CodeBuild action that runs the tests for the applicatio
- I. Program the Lambda function to post the CodeBuild badge as a comment on the pull request so that developers will see the badge in their code review.
- J. Create an Amazon EventBridge rule that reacts to the pullRequestStatusChanged event
- K. Create an AWS Lambda function that invokes a CodePipeline pipeline with a CodeBuild action that runs the tests for the applicatio
- L. Program the Lambda function to post the CodeBuild test results as a comment on the pull request when the test results are complete.

Answer: C

Explanation:

<https://aws.amazon.com/es/blogs/devops/complete-ci-cd-with-aws-codecommit-aws-codebuild-aws-codedeploy-and-aws-codepipeline/>

NEW QUESTION 22

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 projec
- D. In the environment variable, include the ARN of the CodeBuild project's IAM service rol
- 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 repositor
- 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 27

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 Beanstal
- B. Deploy an Amazon RDS for MySQL DB instance as part of the Elastic Beanstalk configuration.
- C. Deploy the application on AWS Elastic Beanstal
- 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 32

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 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
- 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 satellite
- 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 35

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

A DevOps engineer needs to add tags to the created resources that include the user ID that created the resource and the cost center ID. The DevOps engineer configures an AWS Lambda function with the cost center mappings to tag the resources. The DevOps engineer also sets up AWS CloudTrail in the AWS account. An Amazon S3 bucket stores the CloudTrail event logs.

Which solution will meet the tagging requirements?

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

Answer: D

Explanation:

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

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

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

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

References:

- ? S3 event notifications
- ? Server access logging
- ? Amazon EventBridge rules
- ? AWS CloudTrail

NEW QUESTION 37

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

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

What is the MOST secure solution that meets these requirements?

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

Answer: B

Explanation:

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

NEW QUESTION 39

A company has an application that runs on Amazon EC2 instances that are in an Auto Scaling group. When the application starts up, the application needs to process data from an Amazon S3 bucket before the application can start to serve requests. The size of the data that is stored in the S3 bucket is growing. When the Auto Scaling group adds new instances, the application now takes several minutes to download and process the data before the application can serve requests. The company must reduce the time that elapses before new EC2 instances are ready to serve requests.

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

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

Answer: A

Explanation:

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

NEW QUESTION 42

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 46

A company's application uses a fleet of Amazon EC2 On-Demand Instances to analyze and process data. The EC2 instances are in an Auto Scaling group. The Auto Scaling group is a target group for an Application Load Balancer (ALB). The application analyzes critical data that cannot tolerate interruption. The application also analyzes noncritical data that can withstand interruption.

The critical data analysis requires quick scalability in response to real-time application demand. The noncritical data analysis involves memory consumption. A DevOps engineer must implement a solution that reduces scale-out latency for the critical data. The solution also must process the noncritical data.

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

- A. For the critical data, modify the existing Auto Scaling group
- B. Create a warm pool instance in the stopped state
- C. Define the warm pool size
- D. Create a new version of the launch template that has detailed monitoring enabled
- E. Use Spot Instances.
- F. For the critical data, modify the existing Auto Scaling group
- G. Create a warm pool instance in the stopped state
- H. Define the warm pool size
- I. Create a new version of the launch template that has detailed monitoring enabled
- J. Use On-Demand Instances.
- K. For the critical data
- L. Modify the existing Auto Scaling group
- M. Create a lifecycle hook to ensure that bootstrap scripts are completed successfully
- N. Ensure that the application on the instances is ready to accept traffic before the instances are registered
- O. Create a new version of the launch template that has detailed monitoring enabled.
- P. For the noncritical data, create a second Auto Scaling group that uses a launch template
- Q. Configure the launch template to install the unified Amazon CloudWatch agent and to configure the CloudWatch agent with a custom memory utilization metric
- R. Use Spot Instance
- S. Add the new Auto Scaling group as the target group for the ALB
- T. Modify the application to use two target groups for critical data and noncritical data.
- . For the noncritical data, create a second Auto Scaling group
- . Choose the predefined memory utilization metric type for the target tracking scaling policy
- . Use Spot Instance
- . Add the new Auto Scaling group as the target group for the ALB
- . Modify the application to use two target groups for critical data and noncritical data.

Answer: BD

Explanation:

? For the critical data, using a warm pool¹ can reduce the scale-out latency by having pre-initialized EC2 instances ready to serve the application traffic. Using On-Demand Instances can ensure that the instances are always available and not interrupted by Spot interruptions².

? For the noncritical data, using a second Auto Scaling group with Spot Instances can reduce the cost and leverage the unused capacity of EC2³. Using a launch template with the CloudWatch agent⁴ can enable the collection of memory utilization metrics, which can be used to scale the group based on the memory demand. Adding the second group as a target group for the ALB and modifying the application to use two target groups can enable routing the traffic based on the data type.

References: 1: Warm pools for Amazon EC2 Auto Scaling 2: Amazon EC2 On-Demand Capacity Reservations 3: Amazon EC2 Spot Instances 4: Metrics collected by the CloudWatch agent

NEW QUESTION 50

A company is divided into teams. Each team has an AWS account and all the accounts are in an organization in AWS Organizations. Each team must retain full administrative rights to its AWS account. Each team also must be allowed to access only AWS services that the company approves for use. AWS services must gain approval through a request and approval process.

How should a DevOps engineer configure the accounts to meet these requirements?

- A. Use AWS CloudFormation StackSets to provision IAM policies in each account to deny access to restricted AWS services
- B. In each account, configure AWS Config rules that ensure that the policies are attached to IAM principals in the account.
- C. Use AWS Control Tower to provision the accounts into OUs within the organization. Configure AWS Control Tower to enable AWS IAM Identity Center (AWS Single Sign-On). Configure IAM Identity Center to provide administrative access. Include deny policies on user roles for restricted AWS services.
- D. Place all the accounts under a new top-level OU within the organization. Create an SCP that denies access to restricted AWS services. Attach the SCP to the OU.
- E. Create an SCP that allows access to only approved AWS services
- F. Attach the SCP to the root OU of the organization
- G. Remove the FullAWSAccess SCP from the root OU of the organization.

Answer: C

Explanation:

<https://docs.aws.amazon.com/vpc/latest/userguide/managed-prefix-lists.html> A managed prefix list is a set of one or more CIDR blocks. You can use prefix lists to make it easier to configure and maintain your security groups and route tables. <https://docs.aws.amazon.com/vpc/latest/userguide/sharing-managed-prefix-lists.html> With AWS Resource Access Manager (AWS RAM), the owner of a prefix list can share a prefix list with the following: Specific AWS accounts inside or outside of its organization in AWS Organizations An organizational unit inside its organization in AWS Organizations An entire organization in AWS Organizations

NEW QUESTION 53

A business has an application that consists of five independent AWS Lambda functions.

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

After working with the pipeline for a few months, the DevOps engineer has noticed the pipeline takes too long to complete.

What should the DevOps engineer implement to BEST improve the speed of the pipeline?

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

Answer: C

Explanation:

<https://docs.aws.amazon.com/codepipeline/latest/userguide/reference-pipeline-structure.html>

AWS doc: "To specify parallel actions, use the same integer for each action you want to run in parallel. For example, if you want three actions to run in sequence in a stage, you would give the first action the runOrder value of 1, the second action the runOrder value of 2, and the third the runOrder value of 3. However, if you

want the second and third actions to run in parallel, you would give the first action the runOrder value of 1 and both the second and third actions the runOrder value of 2."

NEW QUESTION 54

A company wants to use AWS CloudFormation for infrastructure deployment. The company has strict tagging and resource requirements and wants to limit the deployment to two Regions. Developers will need to deploy multiple versions of the same application.

Which solution ensures resources are deployed in accordance with company policy?

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

Answer: D

Explanation:

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

NEW QUESTION 55

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 EB extension 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 56

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 1AM role to allow the required access to CodeBuild.
- H. Configure the centralized DevOps account's deployment 1AM 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 1AM role to allow the required access to CodeBuild.
- K. Configure the application account's deployment 1AM 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 1AM 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 63

A developer is maintaining a fleet of 50 Amazon EC2 Linux servers. The servers are part of an Amazon EC2 Auto Scaling group, and also use Elastic Load Balancing for load balancing.

Occasionally, some application servers are being terminated after failing ELB HTTP health checks. The developer would like to perform a root cause analysis on the issue, but before being able to access application logs, the server is terminated.

How can log collection be automated?

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

Answer: D

Explanation:

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

NEW QUESTION 64

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 69

AnyCompany is using AWS Organizations to create and manage multiple AWS accounts. AnyCompany recently acquired a smaller company, Example Corp. During the acquisition process, Example Corp's single AWS account joined AnyCompany's management account through an Organizations invitation. AnyCompany moved the new member account under an OU that is dedicated to Example Corp.

AnyCompany's DevOps engineer has an IAM user that assumes a role that is named OrganizationAccountAccessRole to access member accounts. This role is configured with a full access policy. When the DevOps engineer tries to use the AWS Management Console to assume the role in Example Corp's new member account, the DevOps engineer receives the following error message: "Invalid information in one or more fields. Check your information or contact your administrator."

Which solution will give the DevOps engineer access to the new member account?

- A. In the management account, grant the DevOps engineer's IAM user permission to assume the OrganizationAccountAccessRole IAM role in the new member account.
- B. In the management account, create a new SCP. In the SCP, grant the DevOps engineer's IAM user full access to all resources in the new member account.
- C. Attach the SCP to the OU that contains the new member account.
- D. In the new member account, create a new IAM role that is named OrganizationAccountAccessRole.
- E. Attach the AdministratorAccess AWS managed policy to the role.
- F. In the role's trust policy, grant the management account permission to assume the role.
- G. In the new member account, edit the trust policy for the OrganizationAccountAccessRole IAM role.
- H. Grant the management account permission to assume the role.

Answer: C

Explanation:

The problem is that the DevOps engineer cannot assume the OrganizationAccountAccessRole IAM role in the new member account that joined AnyCompany's management account through an Organizations invitation. The solution is to create a new IAM role with the same name and trust policy in the new member account.

? Option A is incorrect, as it does not address the root cause of the error. The DevOps engineer's IAM user already has permission to assume the OrganizationAccountAccessRole IAM role in any member account, as this is the default role name that AWS Organizations creates when a new account joins an organization. The error occurs because the new member account does not have this role, as it was not created by AWS Organizations.

? Option B is incorrect, as it does not address the root cause of the error. An SCP is a policy that defines the maximum permissions for account members of an organization or organizational unit (OU). An SCP does not grant permissions to IAM users or roles, but rather limits the permissions that identity-based policies or resource-based policies grant to them. An SCP also does not affect how IAM roles are assumed by other principals.

? Option C is correct, as it addresses the root cause of the error. By creating a new IAM role with the same name and trust policy as the OrganizationAccountAccessRole IAM role in the new member account, the DevOps engineer can assume this role and access the account. The new role should have the AdministratorAccess AWS managed policy attached, which grants full access to all AWS resources in the account. The trust policy should allow the management account to assume the role, which can be done by specifying the management account ID as a principal in the policy statement.

? Option D is incorrect, as it assumes that the new member account already has the OrganizationAccountAccessRole IAM role, which is not true. The new member account does not have this role, as it was not created by AWS Organizations. Editing the trust policy of a non-existent role will not solve the problem.

NEW QUESTION 74

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 cfn-init script to run when the instance is launched and configure the cfn-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 Init metadata.
- K. Place the configuration file content in the metadata.
- L. Configure the cfn-init script to run when the instance is launched and configure the cfn-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 cfn-init helper script. If your template calls the cfn-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 76

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

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

- A. Use AWS OpsWorks Stacks to layer the server nodes of that cluster.
- B. Create a Chef recipe that populates the content of the /etc./cluster/nodes config file and restarts the service by using the current members of the layer.
- C. Assign that recipe to the Configure lifecycle event.
- D. Put the file nodes config in version control.
- E. Create an AWS CodeDeploy deployment configuration and deployment group based on an Amazon EC2 tag value for the cluster node.
- F. When adding a new node to the cluster, update the file with all tagged instances and make a commit in version control.

- G. Deploy the new file and restart the services.
- H. Create an Amazon S3 bucket and upload a version of the /etc./cluster/nodes config file Create a crontab script that will poll for that S3 file and download it frequently
- I. Use a process manager such as Monit or system, to restart the cluster services when it detects that the new file was modified
- J. When adding a node to the cluster edit the file's most recent members Upload the new file to the S3 bucket.
- K. Create a user data script that lists all members of the current security group of the cluster and automatically updates the /etc/cluster/. nodes config
- L. Tile whenever a new instance is added to the cluster.

Answer: A

Explanation:

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

NEW QUESTION 78

A company is hosting a web application in an AWS Region. For disaster recovery purposes, a second region is being used as a standby. Disaster recovery requirements state that session data must be replicated between regions in near-real time and 1% of requests should route to the secondary region to continuously verify system functionality. Additionally, if there is a disruption in service in the main region, traffic should be automatically routed to the secondary region, and the secondary region must be able to

scale up to handle all traffic.

How should a DevOps engineer meet these requirements?

- A. In both regions, deploy the application on AWS Elastic Beanstalk and use Amazon DynamoDB global tables for session data
- B. Use an Amazon Route 53 weighted routing policy with health checks to distribute the traffic across the regions.
- C. In both regions, launch the application in Auto Scaling groups and use DynamoDB for session data
- D. Use a Route 53 failover routing policy with health checks to distribute the traffic across the regions.
- E. In both regions, deploy the application in AWS Lambda, exposed by Amazon API Gateway, and use Amazon RDS for PostgreSQL with cross-region replication for session data
- F. Deploy the web application with client-side logic to call the API Gateway directly.
- G. In both regions, launch the application in Auto Scaling groups and use DynamoDB global tables for session data
- H. Enable an Amazon CloudFront weighted distribution across region
- I. Point the Amazon Route 53 DNS record at the CloudFront distribution.

Answer: D

NEW QUESTION 82

A rapidly growing company wants to scale for developer demand for AWS development environments. Development environments are created manually in the AWS Management Console. The networking team uses AWS CloudFormation to manage the networking infrastructure, exporting stack output values for the Amazon VPC and all subnets. The development environments have common standards, such as Application Load Balancers, Amazon EC2 Auto Scaling groups, security groups, and Amazon DynamoDB tables.

To keep up with demand, the DevOps engineer wants to automate the creation of development environments. Because the infrastructure required to support the application is expected to grow, there must be a way to easily update the deployed infrastructure. CloudFormation will be used to create a template for the development environments.

Which approach will meet these requirements and quickly provide consistent AWS environments for developers?

- A. Use Fn::ImportValue intrinsic functions in the Resources section of the template to retrieve Virtual Private Cloud (VPC) and subnet values
- B. Use CloudFormation StackSets for the development environments, using the Count input parameter to indicate the number of environments needed
- C. Use the UpdateStackSet command to update existing development environments.
- D. Use nested stacks to define common infrastructure components
- E. To access the exported values, use TemplateURL to reference the networking team's template
- F. To retrieve Virtual Private Cloud (VPC) and subnet values, use Fn::ImportValue intrinsic functions in the Parameters section of the root template
- G. Use the CreateChangeSet and ExecuteChangeSet commands to update existing development environments.
- H. Use nested stacks to define common infrastructure components
- I. Use Fn::ImportValue intrinsic functions with the resources of the nested stack to retrieve Virtual Private Cloud (VPC) and subnet values
- J. Use the CreateChangeSet and ExecuteChangeSet commands to update existing development environments.
- K. Use Fn::ImportValue intrinsic functions in the Parameters section of the root template to retrieve Virtual Private Cloud (VPC) and subnet values
- L. Define the development resources in the order they need to be created in the CloudFormation nested stack
- M. Use the CreateChangeSet
- N. and ExecuteChangeSet commands to update existing development environments.

Answer: C

Explanation:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference-importvalue.html>

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/intrinsic-function-reference-importvalue.html> CF of network exports the VPC, subnet or needed information CF of application imports the above information to its stack and UpdateChangeSet/ ExecuteChangeSet

NEW QUESTION 86

A company is performing vulnerability scanning for all Amazon EC2 instances across many accounts. The accounts are in an organization in AWS Organizations. Each account's VPCs are attached to a shared transit gateway. The VPCs send traffic to the internet through a central egress VPC. The company has enabled Amazon Inspector in a delegated administrator account and has enabled scanning for all member accounts.

A DevOps engineer discovers that some EC2 instances are listed in the "not scanning" tab in Amazon Inspector.

Which combination of actions should the DevOps engineer take to resolve this issue? (Choose three.)

- A. Verify that AWS Systems Manager Agent is installed and is running on the EC2 instances that Amazon Inspector is not scanning.
- B. Associate the target EC2 instances with security groups that allow outbound communication on port 443 to the AWS Systems Manager service endpoint.
- C. Grant inspector: StartAssessmentRun permissions to the IAM role that the DevOps engineer is using.

- D. Configure EC2 Instance Connect for the EC2 instances that Amazon Inspector is not scanning.
- E. Associate the target EC2 instances with instance profiles that grant permissions to communicate with AWS Systems Manager.
- F. Create a managed-instance activation
- G. Use the Activation Code and the Activation ID to register the EC2 instances.

Answer: ABE

Explanation:

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

NEW QUESTION 90

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 application
- B. Use AWS CodeDeploy to deploy the Lambda functions with the Canary10Percent15Minutes Deployment Preference Type
- 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 resources
- 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 resources
- 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 version
- 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 95

A company wants to deploy a workload on several hundred Amazon EC2 instances. The company will provision the EC2 instances in an Auto Scaling group by using a launch template.

The workload will pull files from an Amazon S3 bucket, process the data, and put the results into a different S3 bucket. The EC2 instances must have least-privilege permissions and must use temporary security credentials.

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

- A. Create an IAM role that has the appropriate permissions for S3 bucket
- B. Add the IAM role to an instance profile.
- C. Update the launch template to include the IAM instance profile.
- D. Create an IAM user that has the appropriate permissions for Amazon S3. Generate a secret key and token.
- E. Create a trust anchor and profile
- F. Attach the IAM role to the profile.
- G. Update the launch template
- H. Modify the user data to use the new secret key and token.

Answer: AB

Explanation:

To meet the requirements of deploying a workload on several hundred EC2 instances with least-privilege permissions and temporary security credentials, the company should use an IAM role and an instance profile. An IAM role is a way to grant permissions to an entity that you trust, such as an EC2 instance. An instance profile is a container for an IAM role that you can use to pass role information to an EC2 instance when the instance starts. By using an IAM role and an instance profile, the EC2 instances can automatically receive temporary security credentials from the AWS Security Token Service (STS) and use them to access the S3 buckets. This way, the company does not need to manage or rotate any long-term credentials, such as IAM users or access keys.

To use an IAM role and an instance profile, the company should create an IAM role that has the appropriate permissions for S3 buckets. The permissions should allow the EC2 instances to read from the source S3 bucket and write to the destination S3 bucket. The company should also create a trust policy for the IAM role that specifies that EC2 is allowed to assume the role. Then, the company should add the IAM role to an instance profile. An instance profile can have only one IAM role, so the company does not need to create multiple roles or profiles for this scenario.

Next, the company should update the launch template to include the IAM instance profile. A launch template is a way to save launch parameters for EC2 instances, such as the instance type, security group, user data, and IAM instance profile. By using a launch template, the company can ensure that all EC2 instances in the Auto Scaling group have consistent configuration and permissions. The company should specify the name or ARN of the IAM instance profile in the launch template. This way, when the Auto Scaling group launches new EC2 instances based on the launch template, they will automatically receive the IAM role and its permissions through the instance profile.

The other options are not correct because they do not meet the requirements or follow best practices. Creating an IAM user and generating a secret key and token is not a good option because it involves managing long-term credentials that need to be rotated regularly. Moreover, embedding credentials in user data is not

secure because user data is visible to anyone who can describe the EC2 instance. Creating a trust anchor and profile is not a valid option because trust anchors are used for certificate-based authentication, not for IAM roles or instance profiles. Modifying user data to use a new secret key and token is also not a good option because it requires updating user data every time the credentials change, which is not scalable or efficient.

References:

? 1: AWS Certified DevOps Engineer - Professional Certification | AWS Certification

| AWS

? 2: DevOps Resources - Amazon Web Services (AWS)

? 3: Exam Readiness: AWS Certified DevOps Engineer - Professional

? : IAM Roles for Amazon EC2 - AWS Identity and Access Management

? : Working with Instance Profiles - AWS Identity and Access Management

? : Launching an Instance Using a Launch Template - Amazon Elastic Compute Cloud

? : Temporary Security Credentials - AWS Identity and Access Management

NEW QUESTION 98

A company has its AWS accounts in an organization in AWS Organizations. AWS Config is manually configured in each AWS account. The company needs to implement a solution to centrally configure AWS Config for all accounts in the organization. The solution also must record resource changes to a central account. Which combination of actions should a DevOps engineer perform to meet these requirements? (Choose two.)

- A. Configure a delegated administrator account for AWS Config
- B. Enable trusted access for AWS Config in the organization.
- C. Configure a delegated administrator account for AWS Config
- D. Create a service-linked role for AWS Config in the organization's management account.
- E. Create an AWS CloudFormation template to create an AWS Config aggregator
- F. Configure a CloudFormation stack set to deploy the template to all accounts in the organization.
- G. Create an AWS Config organization aggregator in the organization's management account
- H. Configure data collection from all AWS accounts in the organization and from all AWS Regions.
- I. Create an AWS Config organization aggregator in the delegated administrator account
- J. Configure data collection from all AWS accounts in the organization and from all AWS Regions.

Answer: AE

Explanation:

<https://aws.amazon.com/blogs/mt/org-aggregator-delegated-admin/> <https://docs.aws.amazon.com/organizations/latest/userguide/services-that-can-integrate-config.html>

NEW QUESTION 102

A DevOps engineer is working on a data archival project that requires the migration of on-premises data to an Amazon S3 bucket. The DevOps engineer develops a script that incrementally archives on-premises data that is older than 1 month to Amazon S3. Data that is transferred to Amazon S3 is deleted from the on-premises location. The script uses the S3 PutObject operation.

During a code review the DevOps engineer notices that the script does not verify whether the data was successfully copied to Amazon S3. The DevOps engineer must update the script to ensure that data is not corrupted during transmission. The script must use MD5 checksums to verify data integrity before the on-premises data is deleted.

Which solutions for the script will meet these requirements? (Select TWO.)

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

Answer: BD

Explanation:

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

NEW QUESTION 103

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

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

Which type of events will match this event pattern?

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

Answer: B

Explanation:

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

NEW QUESTION 104

To run an application, a DevOps engineer launches an Amazon EC2 instance with public IP addresses in a public subnet. A user data script obtains the application artifacts and installs them on the instances upon launch. A change to the security classification of the application now requires the instances to run with no access to the internet. While the instances launch successfully and show as healthy, the application does not seem to be installed. Which of the following should successfully install the application while complying with the new rule?

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

Answer: C

Explanation:

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

NEW QUESTION 108

A company hosts applications in its AWS account. Each application logs to an individual Amazon CloudWatch log group. The company's CloudWatch costs for ingestion are increasing. A DevOps engineer needs to identify which applications are the source of the increased logging costs. Which solution will meet these requirements?

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

Answer: C

Explanation:

The correct answer is C.
 A comprehensive and detailed explanation is:
 ? Option A is incorrect because using CloudWatch metrics to create a custom expression that identifies the CloudWatch log groups that have the most data being written to them is not a valid solution. CloudWatch metrics do not provide information about the size or volume of data being ingested by CloudWatch logs.

CloudWatch metrics only provide information about the number of events, bytes, and errors that occur within a log group or stream. Moreover, creating a custom expression with CloudWatch metrics would require using the search_web tool, which is not necessary for this use case.

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

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

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

References:

? CloudWatch Metrics

? CloudWatch Logs Insights

? AWS Cost Explorer

? AWS CloudTrail

NEW QUESTION 109

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

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

- A. Add the bucket name to the AllowedBuckets section of the CodeBuild project setting
- B. Update the build spec to use the AWS CLI to download the database population script.
- C. Modify the S3 bucket settings to enable HTTPS basic authentication and specify a token
- D. Update the build spec to use cURL to pass the token and download the database population script.
- E. Remove unauthenticated access from the S3 bucket with a bucket policy
- F. Modify the service role for the CodeBuild project to include Amazon S3 access
- G. Use the AWS CLI to download the database population script.
- H. Remove unauthenticated access from the S3 bucket with a bucket policy
- I. Use the AWS CLI to download the database population script using an IAM access key and a secret access key.

Answer: C

Explanation:

A bucket policy is a resource-based policy that defines who can access a specific S3 bucket and what actions they can perform on it. By removing unauthenticated access from the bucket policy, you can prevent anyone without valid credentials from accessing the bucket. A service role is an IAM role that allows an AWS service, such as CodeBuild, to perform actions on your behalf. By modifying the service role for the CodeBuild project to include Amazon S3 access, you can grant the project permission to read and write objects in the S3 bucket. The AWS CLI is a command-line tool that allows you to interact with AWS services, such as S3, using commands in your terminal. By using the AWS CLI to download the database population script, you can leverage the service role credentials and encryption to secure the data transfer.

For more information, you can refer to these web pages:

? [Using bucket policies and user policies - Amazon Simple Storage Service]

? [Create a service role for CodeBuild - AWS CodeBuild]

? [AWS Command Line Interface]

NEW QUESTION 113

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 Store to store the AMI ID and reference it in the launch template. 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 ID, 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 114

An IT team has built an AWS CloudFormation template so others in the company can quickly and reliably deploy and terminate an application. The template creates an Amazon EC2 instance with a user data script to install the application and an Amazon S3 bucket that the application uses to serve static webpages

while it is running.

All resources should be removed when the CloudFormation stack is deleted. However, the team observes that CloudFormation reports an error during stack deletion, and the S3 bucket created by the stack is not deleted.

How can the team resolve the error in the MOST efficient manner to ensure that all resources are deleted without errors?

- A. Add a DeletionPolicy attribute to the S3 bucket resource, with the value Delete forcing the bucket to be removed when the stack is deleted.
- B. Add a custom resource with an AWS Lambda function with the DependsOn attribute specifying the S3 bucket, and an IAM role
- C. Write the Lambda function to delete all objects from the bucket when RequestType is Delete.
- D. Identify the resource that was not deleted
- E. Manually empty the S3 bucket and then delete it.
- F. Replace the EC2 and S3 bucket resources with a single AWS OpsWorks Stacks resource
- G. Define a custom recipe for the stack to create and delete the EC2 instance and the S3 bucket.

Answer: B

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/cloudformation-s3-custom-resources/>

NEW QUESTION 118

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

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

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

Answer: AD

Explanation:

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

NEW QUESTION 121

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

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

Which solution will meet these requirements?

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

Answer: B

Explanation:

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

? Set up AWS Config in the account.

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

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

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

NEW QUESTION 123

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 new
- 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 VPCs
- 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 125

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 to 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 meets these requirements? (Select THREE.)

- A. Remove the AdministratorAccess policy
- B. Assign the ReadOnlyAccess managed IAM policy to the developer role
- C. Instruct the developers to use the CloudFormationDeployment role as a CloudFormation service role when the developers deploy new stacks.
- D. Update the trust of CloudFormationDeployment role to allow the developer IAM role to assume the CloudFormationDeployment role.
- E. Configure the IAM to be able to get and pass the CloudFormationDeployment role if cloudformation actions for resources,
- F. Update the trust of the CloudFormationDeployment role to allow the cloudformation.amazonaws.com AWS principal to perform the iam:AssumeRole action
- G. Remove the AdministratorAccess policy
- H. Assign the ReadOnlyAccess managed IAM policy to the developer role. Instruct the developers to assume the CloudFormationDeployment role when they deploy new stacks
- I. Add an IAM policy to CloudFormationDeployment to allow cloudformation:* on an IAM role. Add a policy that allows the iam:PassRole action for ARN of iam:PassedToService 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 user. 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 stack.

? 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 130

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

Which solution will meet these requirements?

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

Answer: A

Explanation:

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

NEW QUESTION 132

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

Which logging solution will support these requirements?

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

Answer: A

Explanation:

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

NEW QUESTION 133

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 only
- 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 only

- 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 only
- 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 only
- 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 136

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your AWS-Certified-DevOps-Engineer-Professional Exam with Our Prep Materials Via below:

<https://www.certleader.com/AWS-Certified-DevOps-Engineer-Professional-dumps.html>