

Amazon

Exam Questions AWS-Certified-DevOps-Engineer-Professional

Amazon AWS Certified DevOps Engineer Professional



NEW QUESTION 1

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 2

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

services.

Which solution will meet these requirements?

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

Answer: D

Explanation:

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

NEW QUESTION 3

A company runs applications in AWS accounts that are in an organization in AWS Organizations. The applications use Amazon EC2 instances and Amazon S3. The company wants to detect potentially compromised EC2 instances, suspicious network activity, and unusual API activity in its existing AWS accounts and in any AWS accounts that the company creates in the future. When the company detects one of these events, the company wants to use an existing Amazon Simple Notification Service (Amazon SNS) topic to send a notification to its operational support team for investigation and remediation.

Which solution will meet these requirements in accordance with AWS best practices?

- A. In the organization's management account, configure an AWS account as the AmazonGuardDuty administrator account

- B. In the GuardDuty administrator account add the company's existing AWS accounts to GuardDuty as members In the GuardDuty administrator account create an Amazon EventBridge rule with an event pattern to match GuardDuty events and to forward matching events to the SNS topic.
- C. In the organization's management account configure Amazon GuardDuty to add newly created AWS accounts by invitation and to send invitations to the existing AWS accounts Create an AWS Cloud Formation stack set that accepts the GuardDuty invitation and creates an Amazon EventBridge rule Configure the rule with an event pattern to match GuardDuty events and to forward matching events to the SNS topic
- D. GuardDuty events and to forward matching events to the SNS topic
- E. Configure the Cloud Formation stack set to deploy into all AWS accounts in the organization.
- F. In the organization's management account
- G. create an AWS CloudTrail organization trail Activate the organization trail in all AWS accounts in the organization
- H. Create an SCP that enables VPC Flow Logs in each account in the organization
- I. Configure AWS Security Hub for the organization Create an Amazon EventBridge rule with an event pattern to match Security Hub events and to forward matching events to the SNS topic.
- J. In the organization's management account configure an AWS account as the AWS CloudTrail administrator account in the CloudTrail administrator account create a CloudTrail organization trail
- K. Add the company's existing AWS accounts to the organization trail Create an SCP that enables VPC Flow Logs in each account in the organization
- L. Configure AWS Security Hub for the organization
- M. Create an Amazon EventBridge rule with an event pattern to match Security Hub events and to forward matching events to the SNS topic.

Answer: B

Explanation:

It allows the company to detect potentially compromised EC2 instances, suspicious network activity, and unusual API activity in its existing AWS accounts and in any AWS accounts that the company creates in the future using Amazon GuardDuty. It also provides a solution for automatically adding future AWS accounts to GuardDuty by configuring GuardDuty to add newly created AWS accounts by invitation and to send invitations to the existing AWS accounts.

NEW QUESTION 4

A company uses AWS Key Management Service (AWS KMS) keys and manual key rotation to meet regulatory compliance requirements. The security team wants to be notified when any keys have not been rotated after 90 days. Which solution will accomplish this?

- A. Configure AWS KMS to publish to an Amazon Simple Notification Service (Amazon SNS) topic when keys are more than 90 days old.
- B. Configure an Amazon EventBridge event to launch an AWS Lambda function to call the AWS Trusted Advisor API and publish to an Amazon Simple Notification Service (Amazon SNS) topic.
- C. Develop an AWS Config custom rule that publishes to an Amazon Simple Notification Service (Amazon SNS) topic when keys are more than 90 days old.
- D. Configure AWS Security Hub to publish to an Amazon Simple Notification Service (Amazon SNS) topic when keys are more than 90 days old.

Answer: C

Explanation:

<https://aws.amazon.com/blogs/security/how-to-use-aws-config-to-determine-compliance-of-aws-kms-key-policies-to-your-specifications/>

NEW QUESTION 5

A company is building a new pipeline by using AWS CodePipeline and AWS CodeBuild in a build account. The pipeline consists of two stages. The first stage is a CodeBuild job to build and package an AWS Lambda function. The second stage consists of deployment actions that operate on two different AWS accounts a development environment account and a production environment account. The deployment stages use the AWS CloudFormation action that CodePipeline invokes to deploy the infrastructure that the Lambda function requires.

A DevOps engineer creates the CodePipeline pipeline and configures the pipeline to encrypt build artifacts by using the AWS Key Management Service (AWS KMS) AWS managed key for Amazon S3 (the aws/s3 key). The artifacts are stored in an S3 bucket When the pipeline runs, the CloudFormation actions fail with an access denied error.

Which combination of actions must the DevOps engineer perform to resolve this error? (Select TWO.)

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

Answer: BE

NEW QUESTION 6

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

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

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

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

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

F. Configure the Lambda functions in Account B to assume an existing IAM role in Account A.

Answer: AEF

Explanation:

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

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

File System Policy

\$ cat file-system-policy.json

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

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

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

NEW QUESTION 7

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

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

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

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

Answer: BDE

NEW QUESTION 8

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

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

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

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

Answer: BCF

Explanation:

The requirement is to create the infrastructure as code (IaC) and the CI/CD pipeline for the Lambda application that uses canary deployment and automated

rollback. To do this, the DevOps team needs to use the following steps:

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

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

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

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

NEW QUESTION 9

A company has containerized all of its in-house quality control applications. The company is running Jenkins on Amazon EC2 instances, which require patching and upgrading. The compliance officer has requested a DevOps engineer begin encrypting build artifacts since they contain company intellectual property. What should the DevOps engineer do to accomplish this in the MOST maintainable manner?

- A. Automate patching and upgrading using AWS Systems Manager on EC2 instances and encrypt Amazon EBS volumes by default.
- B. Deploy Jenkins to an Amazon ECS cluster and copy build artifacts to an Amazon S3 bucket with default encryption enabled.
- C. Leverage AWS CodePipeline with a build action and encrypt the artifacts using AWS Secrets Manager.
- D. Use AWS CodeBuild with artifact encryption to replace the Jenkins instance running on EC2 instances.

Answer: D

Explanation:

The following are the steps involved in accomplishing this in the most maintainable manner:

? Use AWS CodeBuild with artifact encryption to replace the Jenkins instance running on EC2 instances.

? Configure CodeBuild to encrypt the build artifacts using AWS Secrets Manager.

? Deploy the containerized quality control applications to CodeBuild.

This approach is the most maintainable because it eliminates the need to manage Jenkins on EC2 instances. CodeBuild is a managed service, so the DevOps engineer does not need to worry about patching or upgrading the service. <https://docs.aws.amazon.com/codebuild/latest/userguide/security-encryption.html> Build artifact encryption - CodeBuild requires access to an AWS KMS CMK in order to encrypt its build output artifacts. By default, CodeBuild uses an AWS Key Management Service CMK for Amazon S3 in your AWS account. If you do not want to use this CMK, you must create and configure a customer-managed CMK. For more information Creating keys.

NEW QUESTION 10

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

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

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

Answer: ADF

Explanation:

A comprehensive and detailed explanation is:

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

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

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

? Option D is correct because treating each log as a multi-measure record is a recommended practice for optimizing the query performance in Timestream. Treating each log as a multi-measure record would result in creating a single record for each timestamp, which would reduce the storage size and the query latency. Moreover, treating each log as a multi-measure record would allow querying multiple measures for the same timestamp without using joins, which would simplify and speed up the query processing².

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

? Option F is correct because configuring the memory store retention period to be shorter than the magnetic store retention period is a valid option in Timestream.

The memory store retention period determines how long data is kept in the memory store, which is optimized for fast point-in-time queries. The magnetic store retention period determines how long data is kept in the magnetic store, which is optimized for fast analytical queries. By configuring these retention periods appropriately, you can balance your storage costs and query performance according to your application needs.

References:

? 1: Batch writes

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

? 3: Storage

NEW QUESTION 10

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

Which solution will provide the notification with the LEAST operational effort?

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

Answer: C

Explanation:

The correct answer is C. Configuring AWS CloudTrail to send log data events to an Amazon CloudWatch Logs log group and creating a CloudWatch logs metric filter to match failed ConsoleLogin events is the simplest and most efficient way to monitor and alert on failed login attempts. Creating a CloudWatch alarm that is based on the metric filter and configuring an alarm action to send messages to the SNS topic will ensure that the security team is notified when multiple failed login attempts occur. This solution requires the least operational effort compared to the other options.

Option A is incorrect because it involves configuring AWS CloudTrail to send log management events instead of log data events. Log management events are used to track changes to CloudTrail configuration, such as creating, updating, or deleting a trail. Log data events are used to track API activity in AWS accounts, such as login attempts. Therefore, option A will not capture the failed ConsoleLogin events.

Option B is incorrect because it involves creating an Amazon Athena query and two Amazon EventBridge rules to monitor and alert on failed login attempts. This is a more complex and costly solution than using CloudWatch logs and alarms. Moreover, option B relies on the query returning a failure, which may not happen if the query is executed successfully but does not find any failed logins.

Option D is incorrect because it involves configuring AWS CloudTrail to send log data events to an Amazon S3 bucket and configuring an Amazon S3 event notification for the s3:ObjectCreated event type. This solution will not work because the s3:ObjectCreated event type does not allow filtering by ConsoleLogin failed events. The event notification will be triggered for any object created in the S3 bucket, regardless of the event type. Therefore, option D will generate a lot of false positives and unnecessary notifications. References:

? AWS CloudTrail Log File Examples

? Creating CloudWatch Alarms for CloudTrail Events: Examples

? Monitoring CloudTrail Log Files with Amazon CloudWatch Logs

NEW QUESTION 14

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 15

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

Which solution will meet these requirements?

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

Answer: D

Explanation:

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

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

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

NEW QUESTION 16

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

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

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

Answer: BE

Explanation:

To meet the requirements of failover and disaster recovery, the company should use the following deployment strategies:

? Create an Amazon Aurora global database in two AWS Regions as the data store.

In the event of a failure, promote the secondary Region to the primary for the application. Update the application to use the Aurora cluster endpoint in the secondary Region. This strategy can provide a low RPO and RTO for the data, as Aurora global database replicates data with minimal latency across Regions and allows fast and easy failover¹². The company can use the Amazon Aurora cluster endpoint to connect to the current primary DB cluster without needing to change any application code¹.

? Set up the application in two AWS Regions. Configure AWS Global Accelerator to

point to Application Load Balancers (ALBs) in both Regions. Add both ALBs to a single endpoint group. Use health checks and Auto Scaling groups in each Region. This strategy can provide high availability and performance for the application, as AWS Global Accelerator uses the AWS global network to route traffic to the closest healthy endpoint³. The company can also use static IP addresses that are assigned by Global Accelerator as a fixed entry point for their application¹. By using health checks and Auto Scaling groups, the company can ensure that their application can scale up or down based on demand and handle any instance failures⁴.

The other options are incorrect because:

? Creating an Amazon Aurora Single-AZ cluster in multiple AWS Regions as the data store would not provide a fast failover or disaster recovery solution, as the company would need to manually restore data from backups or snapshots in another Region in case of a failure.

? Creating an Amazon Aurora cluster in multiple AWS Regions as the data store and using a Network Load Balancer to balance the database traffic in different Regions would not work, as Network Load Balancers do not support cross-Region routing. Moreover, this strategy would not provide a consistent view of the data across Regions, as Aurora clusters do not replicate data automatically between Regions unless they are part of a global database.

? Setting up the application in two AWS Regions and using Amazon Route 53 failover routing that points to Application Load Balancers in both Regions would not provide a low RTO, as Route 53 failover routing relies on DNS resolution, which can take time to propagate changes across different DNS servers and clients. Moreover, this strategy would not provide deterministic routing, as Route 53 failover routing depends on DNS caching behavior, which can vary depending on different factors.

NEW QUESTION 18

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 21

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

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

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

Answer: BDE

Explanation:

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

? B. Deploy the application on AWS Elastic Beanstalk. Deploy a separate Amazon RDS for MySQL DB instance outside of Elastic Beanstalk. This approach ensures that the database persists independently of the Elastic Beanstalk environment, which can be torn down and recreated without affecting the database¹²³.

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

? D. Configure an Amazon EventBridge rule to monitor AWS Health events. Use an Amazon Simple Notification Service (Amazon SNS) topic as a target to alert the application team. This setup allows for automated monitoring and alerting of the application team in case of deployment failures or other health events⁵⁶.

References:

? AWS Elastic Beanstalk documentation on deploying Ruby applications¹.

? AWS documentation on application auto-scaling⁷.

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

NEW QUESTION 24

A company runs an application on Amazon EC2 instances. The company uses a series of AWS CloudFormation stacks to define the application resources. A developer performs updates by building and testing the application on a laptop and then uploading the build output and CloudFormation stack templates to Amazon S3. The developer's peers review the changes before the developer performs the CloudFormation stack update and installs a new version of the application onto the EC2 instances.

The deployment process is prone to errors and is time-consuming when the developer updates each EC2 instance with the new application. The company wants to automate as much of the application deployment process as possible while retaining a final manual approval step before the modification of the application or resources.

The company already has moved the source code for the application and the CloudFormation templates to AWS CodeCommit. The company also has created an AWS CodeBuild project to build and test the application.

Which combination of steps will meet the company's requirements? (Choose two.)

- A. Create an application group and a deployment group in AWS CodeDeploy
- B. Install the CodeDeploy agent on the EC2 instances.
- C. Create an application revision and a deployment group in AWS CodeDeploy
- D. Create an environment in CodeDeploy
- E. Register the EC2 instances to the CodeDeploy environment.
- F. Use AWS CodePipeline to invoke the CodeBuild job, run the CloudFormation update, and pause for a manual approval step
- G. After approval, start the AWS CodeDeploy deployment.

- H. Use AWS CodePipeline to invoke the CodeBuild job, create CloudFormation change sets for each of the application stacks, and pause for a manual approval step
- I. After approval, run the CloudFormation change sets and start the AWS CodeDeploy deployment.
- J. Use AWS CodePipeline to invoke the CodeBuild job, create CloudFormation change sets for each of the application stacks, and pause for a manual approval step
- K. After approval, start the AWS CodeDeploy deployment.

Answer: AD

Explanation:

A- <https://docs.aws.amazon.com/codedeploy/latest/userguide/codedeploy-agent.html> D - This option correctly utilizes AWS CodePipeline to invoke the CodeBuild job and create CloudFormation change sets. It adds a manual approval step before executing the change sets and starting the AWS CodeDeploy deployment. This ensures that the deployment process is automated while retaining the final manual approval step.

NEW QUESTION 26

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 satellites
- L. Update the application to use a new virtual queue for any data that it cannot transform, and send the message to the new virtual queue
- M. Instruct the scientists to use the virtual queue to review the data that is not valid
- N. Reprocess this data at a later time.

Answer: C

Explanation:

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

NEW QUESTION 30

A company has configured an Amazon S3 event source on an AWS Lambda function. The company needs the Lambda function to run when a new object is created or an existing object is modified in a particular S3 bucket. The Lambda function will use the S3 bucket name and the S3 object key of the incoming event to read the contents of the created or modified S3 object. The Lambda function will parse the contents and save the parsed contents to an Amazon DynamoDB table. The Lambda function's execution role has permissions to read from the S3 bucket and to write to the DynamoDB table. During testing, a DevOps engineer discovers that the Lambda

function does not run when objects are added to the S3 bucket or when existing objects are modified.

Which solution will resolve this problem?

- A. Increase the memory of the Lambda function to give the function the ability to process large files from the S3 bucket.
- B. Create a resource policy on the Lambda function to grant Amazon S3 the permission to invoke the Lambda function for the S3 bucket
- C. Configure an Amazon Simple Queue Service (Amazon SQS) queue as an OnFailure destination for the Lambda function
- D. Provision space in the /tmp folder of the Lambda function to give the function the ability to process large files from the S3 bucket

Answer: B

Explanation:

? Option A is incorrect because increasing the memory of the Lambda function does not address the root cause of the problem, which is that the Lambda function is not triggered by the S3 event source. Increasing the memory of the Lambda function might improve its performance or reduce its execution time, but it does not affect its invocation. Moreover, increasing the memory of the Lambda function might incur higher costs, as Lambda charges based on the amount of memory allocated to the function.

? Option B is correct because creating a resource policy on the Lambda function to grant Amazon S3 the permission to invoke the Lambda function for the S3 bucket is a necessary step to configure an S3 event source. A resource policy is a JSON document that defines who can access a Lambda resource and under what conditions. By granting Amazon S3 permission to invoke the Lambda function, the company ensures that the Lambda function runs when a new object is created or an existing object is modified in the S3 bucket.

? Option C is incorrect because configuring an Amazon Simple Queue Service (Amazon SQS) queue as an On-Failure destination for the Lambda function does not help with triggering the Lambda function. An On-Failure destination is a feature that allows Lambda to send events to another service, such as SQS or Amazon Simple Notification Service (Amazon SNS), when a function invocation fails. However, this feature only applies to asynchronous invocations, and S3 event sources use synchronous invocations. Therefore, configuring an SQS queue as an On-Failure destination would have no effect on the problem.

? Option D is incorrect because provisioning space in the /tmp folder of the Lambda function does not address the root cause of the problem, which is that the Lambda function is not triggered by the S3 event source. Provisioning space in the /tmp folder of the Lambda function might help with processing large files from the S3 bucket, as it provides temporary storage for up to 512 MB of data. However, it does not affect the invocation of the Lambda function.

References:

? Using AWS Lambda with Amazon S3

- ? Lambda resource access permissions
- ? AWS Lambda destinations
- ? [AWS Lambda file system]

NEW QUESTION 31

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 32

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 37

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 40

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

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

Answer: A

Explanation:

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

NEW QUESTION 45

A DevOps engineer is setting up a container-based architecture. The engineer has decided to use AWS CloudFormation to automatically provision an Amazon ECS cluster and an Amazon EC2 Auto Scaling group to launch the EC2 container instances. After successfully creating the CloudFormation stack, the engineer noticed that, even though the ECS cluster and the EC2 instances were created successfully and the stack finished the creation, the EC2 instances were associating with a different cluster.

How should the DevOps engineer update the CloudFormation template to resolve this issue?

- A. Reference the EC2 instances in the AWS: ECS: Cluster resource and reference the ECS cluster in the AWS: ECS: Service resource.
- B. Reference the ECS cluster in the AWS: AutoScaling: LaunchConfiguration resource of the UserData property.
- C. Reference the ECS cluster in the AWS:EC2: Instance resource of the UserData property.
- D. Reference the ECS cluster in the AWS: CloudFormation: CustomResource resource to trigger an AWS Lambda function that registers the EC2 instances with the appropriate ECS cluster.

Answer: B

Explanation:

The UserData property of the AWS: AutoScaling: LaunchConfiguration resource can be used to specify a script that runs when the EC2 instances are launched. This script can include the ECS cluster name as an environment variable for the ECS agent running on the EC2 instances. This way, the EC2 instances will register with the correct ECS cluster. Option A is incorrect because the AWS: ECS: Cluster resource does not have a property to reference the EC2 instances. Option C is incorrect because the EC2 instances are launched by the Auto Scaling group, not by the AWS: EC2: Instance resource. Option D is incorrect because using a custom resource and a Lambda function is unnecessary and overly complex for this scenario. References: AWS::AutoScaling::LaunchConfiguration, Amazon ECS Container Agent Configuration

NEW QUESTION 46

A company plans to use Amazon CloudWatch to monitor its Amazon EC2 instances. The company needs to stop EC2 instances when the average of the NetworkPacketsIn metric is less than 5 for at least 3 hours in a 12-hour time window. The company must evaluate the metric every hour. The EC2 instances must continue to run if there is missing data for the NetworkPacketsIn metric during the evaluation period.

A DevOps engineer creates a CloudWatch alarm for the NetworkPacketsIn metric. The DevOps engineer configures a threshold value of 5 and an evaluation period of 1 hour.

Which set of additional actions should the DevOps engineer take to meet these requirements?

- A. Configure the Datapoints to Alarm value to be 3 out of 12. Configure the alarm to treat missing data as breaching the threshold
- B. Add an AWS Systems Manager action to stop the instance when the alarm enters the ALARM state.
- C. Configure the Datapoints to Alarm value to be 3 out of 12. Configure the alarm to treat missing data as not breaching the threshold
- D. Add an EC2 action to stop the instance when the alarm enters the ALARM state.
- E. Configure the Datapoints to Alarm value to be 9 out of 12. Configure the alarm to treat missing data as breaching the threshold
- F. Add an EC2 action to stop the instance when the alarm enters the ALARM state.
- G. Configure the Datapoints to Alarm value to be 9 out of 12. Configure the alarm to treat missing data as not breaching the threshold
- H. Add an AWS Systems Manager action to stop the instance when the alarm enters the ALARM state.

Answer: B

Explanation:

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

NEW QUESTION 47

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

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

Which solution meets these requirements with the MOST operational efficiency?

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

Answer: B

Explanation:

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

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

NEW QUESTION 50

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

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

A DevOps engineer must implement a solution to identify in near real time any AWS

service misconfiguration that results in noncompliance. The solution must automatically remediate the issue within 15 minutes of identification. The solution also must track noncompliant resources and events in a centralized dashboard with accurate timestamps.

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

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

Answer: C

Explanation:

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

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

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

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

Option B is not optimal because CloudTrail logs analysis using Athena is not a near real time solution. Athena queries have to be manually run or scheduled using a cron expression. Athena also does not provide remediation actions, so a custom Lambda function has to be developed and invoked. Step Functions can be used to orchestrate the query and remediation workflow, but it adds more complexity and cost. QuickSight dashboard can be used for tracking, but it does not provide a

comprehensive view of the security posture of the AWS accounts and resources.

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

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

NEW QUESTION 52

A company's security team requires that all external Application Load Balancers (ALBs) and Amazon API Gateway APIs are associated with AWS WAF web ACLs. The company has hundreds of AWS accounts, all of which are included in a single organization in AWS Organizations. The company has configured AWS Config for the organization. During an audit, the company finds some externally facing ALBs that are not associated with AWS WAF web ACLs. Which combination of steps should a DevOps engineer take to prevent future violations? (Choose two.)

- A. Delegate AWS Firewall Manager to a security account.
- B. Delegate Amazon GuardDuty to a security account.
- C. Create an AWS Firewall Manager policy to attach AWS WAF web ACLs to any newly created ALBs and API Gateway APIs.
- D. Create an Amazon GuardDuty policy to attach AWS WAF web ACLs to any newly created ALBs and API Gateway APIs.
- E. Configure an AWS Config managed rule to attach AWS WAF web ACLs to any newly created ALBs and API Gateway APIs.

Answer: AC

Explanation:

If instead you want to automatically apply the policy to existing in-scope resources, choose Auto remediate any noncompliant resources. This option creates a web ACL in each applicable account within the AWS organization and associates the web ACL with the resources in the accounts. When you choose Auto remediate any noncompliant resources, you can also choose to remove existing web ACL associations from in-scope resources, for the web ACLs that aren't managed by another active Firewall Manager policy. If you choose this option, Firewall Manager first associates the policy's web ACL with the resources, and then removes the prior associations. If a resource has an association with another web ACL that's managed by a different active Firewall Manager policy, this choice doesn't affect that association.

NEW QUESTION 55

A company uses AWS Directory Service for Microsoft Active Directory as its identity provider (IdP). The company requires all infrastructure to be defined and deployed by AWS CloudFormation.

A DevOps engineer needs to create a fleet of Windows-based Amazon EC2 instances to host an application. The DevOps engineer has created a CloudFormation template that contains an EC2 launch template, IAM role, EC2 security group, and EC2 Auto Scaling group. The DevOps engineer must implement a solution that joins all EC2 instances to the domain of the AWS Managed Microsoft AD directory.

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

- A. In the CloudFormation template, create an AWS::SSM::Document resource that joins the EC2 instance to the AWS Managed Microsoft AD domain by using the parameters for the existing director
- B. Update the launch template to include the SSMAssociation property to use the new SSM document
- C. Attach the AmazonSSMManagedInstanceCore and AmazonSSMDirectoryServiceAccess AWS managed policies to the IAM role that the EC2 instances use.
- D. In the CloudFormation template, update the launch template to include specific tags that propagate on launch
- E. Create an AWS::SSM::Association resource to associate the AWS- JoinDirectoryServiceDomain Automation runbook with the EC2 instances that have the specified tag
- F. Define the required parameters to join the AWS Managed Microsoft AD director
- G. Attach the AmazonSSMManagedInstanceCore and AmazonSSMDirectoryServiceAccess AWS managed policies to the IAM role that the EC2 instances use.
- H. Store the existing AWS Managed Microsoft AD domain connection details in AWS Secrets Manager
- I. In the CloudFormation template, create an AWS::SSM::Association resource to associate the AWS-CreateManagedWindowsInstanceWithApproval Automation runbook with the EC2 Auto Scaling group
- J. Pass the ARNs for the parameters from Secrets Manager to join the domain
- K. Attach the AmazonSSMDirectoryServiceAccess and SecretsManagerReadWrite AWS managed policies to the IAM role that the EC2 instances use.
- L. Store the existing AWS Managed Microsoft AD domain administrator credentials in AWS Secrets Manager
- M. In the CloudFormation template, update the EC2 launch template to include user data
- N. Configure the user data to pull the administrator credentials from Secrets Manager and to join the AWS Managed Microsoft AD domain
- O. Attach the AmazonSSMManagedInstanceCore and SecretsManagerReadWrite AWS managed policies to the IAM role that the EC2 instances use.

Answer: B

Explanation:

To meet the requirements, the DevOps engineer needs to create a solution that joins all EC2 instances to the domain of the AWS Managed Microsoft AD directory with the most operational efficiency. The DevOps engineer can use AWS Systems Manager Automation to automate the domain join process using an existing runbook called AWS- JoinDirectoryServiceDomain. This runbook can join Windows instances to an AWS Managed Microsoft AD or Simple AD directory by using PowerShell commands. The DevOps engineer can create an AWS::SSM::Association resource in the CloudFormation template to associate the runbook with the EC2 instances that have specific tags. The tags can be defined in the launch template and propagated on launch to the EC2 instances. The DevOps engineer can also define the required parameters for the runbook, such as the directory ID, directory name, and organizational unit. The DevOps engineer can attach the AmazonSSMManagedInstanceCore and AmazonSSMDirectoryServiceAccess AWS managed policies to the IAM role that the EC2 instances use. These policies grant the necessary permissions for Systems Manager and Directory Service operations.

NEW QUESTION 56

A company manages a web application that runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The EC2 instances run in an Auto Scaling group across multiple Availability Zones. The application uses an Amazon RDS for MySQL DB instance to store the data. The company has configured Amazon Route 53 with an alias record that points to the ALB.

A new company guideline requires a geographically isolated disaster recovery (DR) site with an RTO of 4 hours and an RPO of 15 minutes.

Which DR strategy will meet these requirements with the LEAST change to the application stack?

- A. Launch a replica environment of everything except Amazon RDS in a different Availability Zone Create an RDS read replica in the new Availability Zone: and configure the new stack to point to the local RDS DB instance

- B. Add the new stack to the Route 53 record set by using a health check to configure a failover routing policy.
- C. Launch a replica environment of everything except Amazon RDS in a different AW
- D. Region Create an RDS read replica in the new Region and configure the new stack to point to the local RDS DB instanc
- E. Add the new stack to the Route 53 record set by using a health check to configure a latency routing policy.
- F. Launch a replica environment of everything except Amazon RDS ma different AWS Regio
- G. In the event of an outage copy and restore the latest RDS snapshot from the primar
- H. Region to the DR Region Adjust the Route 53 record set to point to the ALB in the DR Region.
- I. Launch a replica environment of everything except Amazon RDS in a different AWS Regio
- J. Create an RDS read replica in the new Region and configure the new environment to point to the local RDS DB instanc
- K. Add the new stack to the Route 53 record set by using a health check to configure a failover routing polic
- L. In the event of an outage promote the read replica to primary.

Answer: D

NEW QUESTION 58

A company has developed an AWS Lambda function that handles orders received through an API. The company is using AWS CodeDeploy to deploy the Lambda function as the final stage of a CI/CD pipeline.

A DevOps engineer has noticed there are intermittent failures of the ordering API for a few seconds after deployment. After some investigation the DevOps engineer believes the failures are due to database changes not having fully propagated before the Lambda function is invoked. How should the DevOps engineer overcome this?

- A. Add a BeforeAllowTraffic hook to the AppSpec file that tests and waits for any necessary database changes before traffic can flow to the new version of the Lambda function.
- B. Add an AfterAllowTraffic hook to the AppSpec file that forces traffic to wait for any pending database changes before allowing the new version of the Lambda function to respond.
- C. Add a BeforeAllowTraffic hook to the AppSpec file that tests and waits for any necessary database changes before deploying the new version of the Lambda function.
- D. Add a validateService hook to the AppSpec file that inspects incoming traffic and rejects the payload if dependent services such as the database are not yet ready.

Answer: A

Explanation:

<https://docs.aws.amazon.com/codedeploy/latest/userguide/reference-appspec-file-structure-hooks.html#appspec-hooks-lambda>

NEW QUESTION 62

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 grou
- B. Create a warm pool instance in the stopped stat
- C. Define the warm pool siz
- D. Create a new version of the launch template that has detailed monitoring enable
- E. use Spot Instances.
- F. For the critical data, modify the existing Auto Scaling grou
- G. Create a warm pool instance in the stopped stat
- H. Define the warm pool siz
- I. Create a new version of the launch template that has detailed monitoring enable
- J. Use On-Demand Instances.
- K. For the critical dat
- L. modify the existing Auto Scaling grou
- M. Create a lifecycle hook to ensure that bootstrap scripts are completed successfull
- N. Ensure that the application on the instances is ready to accept traffic before the instances are registere
- 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 templat
- Q. Configure the launch template to install the unified Amazon CloudWatch agent and to configure the CloudWatch agent with a custom memory utilization metri
- R. Use Spot Instance
- S. Add the new Auto Scaling group as the target group for the AL
- T. Modify the application to use two target groups for critical data and noncritical data.
- . For the noncritical data, create a second Auto Scaling grou
- . Choose the predefined memory utilization metric type for the target tracking scaling polic
- . Use Spot Instance
- . Add the new Auto Scaling group as the target group for the AL
- . 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 66

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

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

Answer: B

Explanation:

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

codes. Then, the DevOps engineer needs to update the default behavior to use the origin group instead of the single origin².

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

References:

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

NEW QUESTION 67

A company is hosting a static website from an Amazon S3 bucket. The website is available to customers at example.com. The company uses an Amazon Route 53 weighted routing policy with a TTL of 1 day. The company has decided to replace the existing static website with a dynamic web application. The dynamic web application uses an Application Load Balancer (ALB) in front of a fleet of Amazon EC2 instances.

On the day of production launch to customers, the company creates an additional Route 53 weighted DNS record entry that points to the ALB with a weight of 255 and a TTL of 1 hour. Two days later, a DevOps engineer notices that the previous static website is displayed sometimes when customers navigate to example.com.

How can the DevOps engineer ensure that the company serves only dynamic content for example.com?

- A. Delete all objects, including previous versions, from the S3 bucket that contains the static website content.
- B. Update the weighted DNS record entry that points to the S3 bucket
- C. Apply a weight of 0. Specify the domain reset option to propagate changes immediately.
- D. Configure webpage redirect requests on the S3 bucket with a hostname that redirects to the ALB.
- E. Remove the weighted DNS record entry that points to the S3 bucket from the example.com hosted zone
- F. Wait for DNS propagation to become complete.

Answer: D

NEW QUESTION 69

A company uses AWS Storage Gateway in file gateway mode in front of an Amazon S3 bucket that is used by multiple resources. In the morning when business begins, users do not see the objects processed by a third party the previous evening. When a DevOps engineer looks directly at the S3 bucket, the data is there, but it is missing in Storage Gateway.

Which solution ensures that all the updated third-party files are available in the morning?

- A. Configure a nightly Amazon EventBridge event to invoke an AWS Lambda function to run the RefreshCache command for Storage Gateway.
- B. Instruct the third party to put data into the S3 bucket using AWS Transfer for SFTP.
- C. Modify Storage Gateway to run in volume gateway mode.
- D. Use S3 Same-Region Replication to replicate any changes made directly in the S3 bucket to Storage Gateway.

Answer: A

Explanation:

https://docs.aws.amazon.com/storagegateway/latest/APIReference/API_RefreshCache.html "It only updates the cached inventory to reflect changes in the inventory of the objects in the S3 bucket. This operation is only supported in the S3 File Gateway types."

NEW QUESTION 70

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

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

- A. Create an Amazon CloudWatch Logs subscription that filters on CodePipeline Pipeline Execution State Change
- B. Publish subscription events to an Amazon Simple Notification Service (Amazon SNS) topic
- C. Subscribe the chat webhook URL to the SNS topic, and complete the subscription validation.
- D. Create an AWS Lambda function that is invoked by AWS CloudTrail event
- E. When a CodePipeline Pipeline Execution State Change event is detected, send the event details to the chat webhook URL.
- F. Create an Amazon EventBridge rule that filters on CodePipeline Pipeline Execution State Change
- G. Publish the events to an Amazon Simple Notification Service (Amazon SNS) topic
- H. Create an AWS Lambda function that sends event details to the chat webhook URL
- I. Subscribe the function to the SNS topic.
- J. Modify the pipeline code to send the event details to the chat webhook URL at the end of each stage
- K. Parameterize the URL so that each pipeline can send to a different URL based on the pipeline environment.

Answer: C

Explanation:

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

NEW QUESTION 73

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

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

- A. Create an Amazon EventBridge rule for the CloudTrail StopLogging event
- B. Create an AWS Lambda function that uses the AWS SDK to call StartLogging on the ARN of the resource in which StopLogging was called
- C. Add the Lambda function ARN as a target to the EventBridge rule.
- D. Deploy the AWS-managed CloudTrail-enabled AWS Config rule set with a periodic interval to 1 hour
- E. Create an Amazon EventBridge rule for AWS Config rules compliance change
- F. Create an AWS Lambda function that uses the AWS SDK to call StartLogging on the ARN of the resource in which StopLogging was called
- G. Add the Lambda function ARN as a target to the EventBridge rule.
- H. Create an Amazon EventBridge rule for a scheduled event every 5 minutes
- I. Create an AWS Lambda function that uses the AWS SDK to call StartLogging on a CloudTrail trail in the AWS account
- J. Add the Lambda function ARN as a target to the EventBridge rule.
- K. Launch a t2.nano instance with a script running every 5 minutes that uses the AWS SDK to query CloudTrail in the current account
- L. If the CloudTrail trail is disabled have the script re-enable the trail.

Answer: A

Explanation:

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

NEW QUESTION 75

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 76

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

Which solution will accomplish this?

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

Answer: B

Explanation:

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

NEW QUESTION 79

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 stat
- 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 stat
- 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 stat
- 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 stat
- 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 82

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 83

A company must encrypt all AMIs that the company shares across accounts. A DevOps engineer has access to a source account where an unencrypted custom AMI has been built. The DevOps engineer also has access to a target account where an Amazon EC2 Auto Scaling group will launch EC2 instances from the AMI. The DevOps engineer must share the AMI with the target account.

The company has created an AWS Key Management Service (AWS KMS) key in the source account.

Which additional steps should the DevOps engineer perform to meet the requirements? (Choose three.)

- A. In the source account, copy the unencrypted AMI to an encrypted AM
- B. Specify the KMS key in the copy action.
- C. In the source account, copy the unencrypted AMI to an encrypted AM
- D. Specify the default Amazon Elastic Block Store (Amazon EBS) encryption key in the copy action.
- E. In the source account, create a KMS grant that delegates permissions to the Auto Scaling group service-linked role in the target account.
- F. In the source account, modify the key policy to give the target account permissions to create a gran
- G. In the target account, create a KMS grant that delegates permissions to the Auto Scaling group service-linked role.
- H. In the source account, share the unencrypted AMI with the target account.
- I. In the source account, share the encrypted AMI with the target account.

Answer: ADF

Explanation:

The Auto Scaling group service-linked role must have a specific grant in the source account in order to decrypt the encrypted AMI. This is because the service-linked role does not have permissions to assume the default IAM role in the source account. The following steps are required to meet the requirements:

- ? In the source account, copy the unencrypted AMI to an encrypted AMI. Specify the KMS key in the copy action.
- ? In the source account, create a KMS grant that delegates permissions to the Auto Scaling group service-linked role in the target account.
- ? In the source account, share the encrypted AMI with the target account.
- ? In the target account, attach the KMS grant to the Auto Scaling group service-linked role.

The first three steps are the same as the steps that I described earlier. The fourth step is required to grant the Auto Scaling group service-linked role permissions to decrypt the AMI in the target account.

NEW QUESTION 86

A company builds a container image in an AWS CodeBuild project by running Docker commands. After the container image is built, the CodeBuild project uploads the container image to an Amazon S3 bucket. The CodeBuild project has an IAM service role that has permissions to access the S3 bucket. A DevOps engineer needs to replace the S3 bucket with an Amazon Elastic Container Registry (Amazon ECR) repository to store the container images. The DevOps engineer creates an ECR private image repository in the same AWS Region of the CodeBuild project. The DevOps engineer adjusts the IAM service role with the permissions that are necessary to work with the new ECR repository. The DevOps engineer also places new repository information into the docker build command and the docker push command that are used in the buildspec.yml file. When the CodeBuild project runs a build job, the job fails when the job tries to access the ECR repository. Which solution will resolve the issue of failed access to the ECR repository?

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

Answer: A

Explanation:

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

NEW QUESTION 88

A development team wants to use AWS CloudFormation stacks to deploy an application. However, the developer IAM role does not have the required permissions to provision the resources that are specified in the AWS CloudFormation template. A DevOps engineer needs to implement a solution that allows the developers to deploy the stacks. The solution must follow the principle of least privilege. Which solution will meet these requirements?

- A. Create an IAM policy that allows the developers to provision the required resource
- B. Attach the policy to the developer IAM role.
- C. Create an IAM policy that allows full access to AWS CloudFormation
- D. Attach the policy to the developer IAM role.
- E. Create an AWS CloudFormation service role that has the required permission
- F. Grant the developer IAM role a `cloudformation:*` action
- G. Use the new service role during stack deployments.
- H. Create an AWS CloudFormation service role that has the required permission
- I. Grant the developer IAM role the `iam:PassRole` permission
- J. Use the new service role during stack deployments.

Answer: D

Explanation:

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-iam-servicerole.html>

NEW QUESTION 91

A company has a single AWS account that runs hundreds of Amazon EC2 instances in a single AWS Region. New EC2 instances are launched and terminated each hour in the account. The account also includes existing EC2 instances that have been running for longer than a week. The company's security policy requires all running EC2 instances to use an EC2 instance profile. If an EC2 instance does not have an instance profile attached, the EC2 instance must use a default instance profile that has no IAM permissions assigned. A DevOps engineer reviews the account and discovers EC2 instances that are running without an instance profile. During the review, the DevOps engineer also observes that new EC2 instances are being launched without an instance profile. Which solution will ensure that an instance profile is attached to all existing and future EC2 instances in the Region?

- A. Configure an Amazon EventBridge rule that reacts to EC2 RunInstances API call
- B. Configure the rule to invoke an AWS Lambda function to attach the default instance profile to the EC2 instances.
- C. Configure the `ec2-instance-profile-attached` AWS Config managed rule with a trigger type of configuration change
- D. Configure an automatic remediation action that invokes an AWS Systems Manager Automation runbook to attach the default instance profile to the EC2 instances.
- E. Configure an Amazon EventBridge rule that reacts to EC2 StartInstances API call
- F. Configure the rule to invoke an AWS Systems Manager Automation runbook to attach the default instance profile to the EC2 instances.
- G. Configure the `iam-role-managed-policy-check` AWS Config managed rule with a trigger type of configuration change
- H. Configure an automatic remediation action that invokes an AWS Lambda function to attach the default instance profile to the EC2 instances.

Answer: B

Explanation:

<https://docs.aws.amazon.com/config/latest/developerguide/ec2-instance-profile-attached.html>

NEW QUESTION 96

A company manages an application that stores logs in Amazon CloudWatch Logs. The company wants to archive the logs to an Amazon S3 bucket. Logs are rarely accessed after 90 days and must be retained for 10 years.

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

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

Answer: BD

Explanation:

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

NEW QUESTION 100

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

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

- A. Use AWS OpsWorks Stacks to layer the server nodes of that cluster.
- 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 systemd, 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 file.
- L. Trigger a new instance to be 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 102

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

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

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

Answer: BCF

Explanation:

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

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

NEW QUESTION 106

A company uses AWS CodeArtifact to centrally store Python packages. The CodeArtifact repository is configured with the following repository policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "codeartifact:DescribePackageVersion",
        "codeartifact:DescribeRepository",
        "codeartifact:GetPackageVersionReadme",
        "codeartifact:GetRepositoryEndpoint",
        "codeartifact:ListPackageVersionAssets",
        "codeartifact:ListPackageVersionDependencies",
        "codeartifact:ListPackageVersions",
        "codeartifact:ListPackages",
        "codeartifact:ReadFromRepository"
      ],
      "Effect": "Allow",
      "Resource": "*",
      "Principal": "*",
      "Condition": {
        "StringEquals": {
          "aws:PrincipalOrgID": [
            "o-xxxxxxxxxxxx"
          ]
        }
      }
    }
  ]
}
```

A development team is building a new project in an account that is in an organization in AWS Organizations. The development team wants to use a Python library that has already been stored in the CodeArtifact repository in the organization. The development team uses AWS CodePipeline and AWS CodeBuild to build the new application. The CodeBuild job that the development team uses to build the application is configured to run in a VPC. Because of compliance requirements the VPC has no internet connectivity.

The development team creates the VPC endpoints for CodeArtifact and updates the CodeBuild buildspec yml file. However, the development team cannot download the Python library from the repository.

Which combination of steps should a DevOps engineer take so that the development team can use Code Artifact? (Select TWO.)

- A. Create an Amazon S3 gateway endpoint. Update the route tables for the subnets that are running the CodeBuild job.
- B. Update the repository policy's Principal statement to include the ARN of the role that the CodeBuild project uses.
- C. Share the CodeArtifact repository with the organization by using AWS Resource Access Manager (AWS RAM).
- D. Update the role that the CodeBuild project uses so that the role has sufficient permissions to use the CodeArtifact repository.
- E. Specify the account that hosts the repository as the delegated administrator for CodeArtifact in the organization.

Answer: AD

Explanation:

"AWS CodeArtifact operates in multiple Availability Zones and stores artifact data and metadata in Amazon S3 and Amazon DynamoDB. Your encrypted data is redundantly stored across multiple facilities and multiple devices in each facility, making it highly available and highly durable."

<https://aws.amazon.com/codeartifact/features/> With no internet connectivity, a gateway endpoint becomes necessary to access S3.

NEW QUESTION 109

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 111

A company uses an organization in AWS Organizations that has all features enabled. The company uses AWS Backup in a primary account and uses an AWS Key Management Service (AWS KMS) key to encrypt the backups. The company needs to automate a cross-account backup of the resources that AWS Backup backs up in the primary account. The company configures cross-account backup in the Organizations management account. The company creates a new AWS account in the organization and configures an AWS Backup backup vault in the new account. The company creates a KMS key in the new account to encrypt the backups. Finally, the company configures a new backup plan in the primary account. The destination for the new backup plan is the backup vault in the new account. When the AWS Backup job in the primary account is invoked, the job creates backups in the primary account. However, the backups are not copied to the new account's backup vault. Which combination of steps must the company take so that backups can be copied to the new account's backup vault? (Select TWO.)

- A. Edit the backup vault access policy in the new account to allow access to the primary account.
- B. Edit the backup vault access policy in the primary account to allow access to the new account.
- C. Edit the backup vault access policy in the primary account to allow access to the KMS key in the new account.
- D. Edit the key policy of the KMS key in the primary account to share the key with the new account.
- E. Edit the key policy of the KMS key in the new account to share the key with the primary account.

Answer: AE

Explanation:

To enable cross-account backup, the company needs to grant permissions to both the backup vault and the KMS key in the destination account. The backup vault access policy in the destination account must allow the primary account to copy backups into the vault. The key policy of the KMS key in the destination account must allow the primary account to use the key to encrypt and decrypt the backups. These steps are described in the AWS documentation¹². Therefore, the correct answer is A and E.

References:

- ? 1: Creating backup copies across AWS accounts - AWS Backup
- ? 2: Using AWS Backup with AWS Organizations - AWS Backup

NEW QUESTION 114

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 value
- 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 component
- 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 component
- I. Use Fn::ImportValue intrinsic functions with the resources of the nested stack to retrieve Virtual Private Cloud (VPC) and subnet value
- 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 value
- 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 119

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

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

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

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