

Exam Questions AWS-Certified-DevOps-Engineer-Professional

Amazon AWS Certified DevOps Engineer Professional

<https://www.2passeasy.com/dumps/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's DevOps engineer uses AWS Systems Manager to perform maintenance tasks during maintenance windows. The company has a few Amazon EC2 instances that require a restart after notifications from AWS Health. The DevOps engineer needs to implement an automated solution to remediate these notifications. The DevOps engineer creates an Amazon EventBridge rule.

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

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

Answer: C

Explanation:

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

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

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

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

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

NEW QUESTION 3

A company is migrating its on-premises Windows applications and Linux applications to AWS. The company will use automation to launch Amazon EC2 instances to mirror the on-premises configurations. The migrated applications require access to shared storage that uses SMB for Windows and NFS for Linux.

The company is also creating a pilot light disaster recovery (DR) environment in another AWS Region. The company will use automation to launch and configure the EC2 instances in the DR Region. The company needs to replicate the storage to the DR Region.

Which storage solution will meet these requirements?

- A. Use Amazon S3 for the application storage
- B. Create an S3 bucket in the primary Region and an S3 bucket in the DR Region
- C. Configure S3 Cross-Region Replication (CRR) from the primary Region to the DR Region.

- D. Use Amazon Elastic Block Store (Amazon EBS) for the application storage
- E. Create a backup plan in AWS Backup that creates snapshots of the EBS volumes that are in the primary Region and replicates the snapshots to the DR Region.
- F. Use a Volume Gateway in AWS Storage Gateway for the application storage
- G. Configure Cross-Region Replication (CRR) of the Volume Gateway from the primary Region to the DR Region.
- H. Use Amazon FSx for NetApp ONTAP for the application storage
- I. Create an FSx for ONTAP instance in the DR Region
- J. Configure NetApp SnapMirror replication from the primary Region to the DR Region.

Answer: D

Explanation:

To meet the requirements of migrating its on-premises Windows and Linux applications to AWS and creating a pilot light DR environment in another AWS Region, the company should use Amazon FSx for NetApp ONTAP for the application storage. Amazon FSx for NetApp ONTAP is a fully managed service that provides highly reliable, scalable, high-performing, and feature-rich file storage built on NetApp's popular ONTAP file system. FSx for ONTAP supports multiple protocols, including SMB for Windows and NFS for Linux, so the company can access the shared storage from both types of applications. FSx for ONTAP also supports NetApp SnapMirror replication, which enables the company to replicate the storage to the DR Region. NetApp SnapMirror replication is efficient, secure, and incremental, and it preserves the data deduplication and compression benefits of FSx for ONTAP. The company can use automation to launch and configure the EC2 instances in the DR Region and then use NetApp SnapMirror to restore the data from the primary Region.

The other options are not correct because they do not meet the requirements or follow best practices. Using Amazon S3 for the application storage is not a good option because S3 is an object storage service that does not support SMB or NFS protocols natively. The company would need to use additional services or software to mount S3 buckets as file systems, which would add complexity and cost. Using Amazon EBS for the application storage is also not a good option because EBS is a block storage service that does not support SMB or NFS protocols natively. The company would need to set up and manage file servers on EC2 instances to provide shared access to the EBS volumes, which would add overhead and maintenance. Using a Volume Gateway in AWS Storage Gateway for the application storage is not a valid option because Volume Gateway does not support SMB protocol. Volume Gateway only supports iSCSI protocol, which means that only Linux applications can access the shared storage.

References:

- ? 1: What is Amazon FSx for NetApp ONTAP? - FSx for ONTAP
- ? 2: Amazon FSx for NetApp ONTAP
- ? 3: Amazon FSx for NetApp ONTAP | NetApp
- ? 4: AWS Announces General Availability of Amazon FSx for NetApp ONTAP
- ? : Replicating Data with NetApp SnapMirror - FSx for ONTAP
- ? : What Is Amazon S3? - Amazon Simple Storage Service
- ? : What Is Amazon Elastic Block Store (Amazon EBS)? - Amazon Elastic Compute Cloud
- ? : What Is AWS Storage Gateway? - AWS Storage Gateway

NEW QUESTION 4

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 application
- B. Define each Lambda function in the template by using the `AWS::Lambda::Function` resource type
- C. In the template, include a version for the Lambda function by using the `AWS::Lambda::Version` resource type
- D. Declare the `CodeSha256` property
- 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 application
- G. Define each Lambda function in the template by using the `AWS::Serverless::Function` resource type
- H. For each function, include configurations for the `AutoPublishAlias` property and the `DeploymentPreference` property
- I. Configure the deployment configuration type to `LambdaCanary10Percent10Minutes`.
- J. Create an AWS CodeCommit repository
- K. Create an AWS CodePipeline pipeline
- L. Use the CodeCommit repository in a new source stage that starts the pipeline
- M. Create an AWS CodeBuild project to deploy the AWS Serverless Application Model (AWS SAM) template
- N. Upload the template and source code to the CodeCommit repository
- 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 repository
- Q. Create an AWS CodePipeline pipeline
- R. Use the CodeCommit repository in a new source stage that starts the pipeline
- 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 repository
- . In the CodeCommit repository, create an `appspect.yml` file that includes the commands to deploy the CloudFormation template.
- . Create an Amazon CloudWatch composite alarm for all the Lambda functions
- . Configure an evaluation period and dimensions for Lambda
- . Configure the alarm to enter the `ALARM` state if any errors are detected or if there is insufficient data.
- . 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.

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 5

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

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

Answer: ABF

Explanation:

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

NEW QUESTION 6

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 7

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

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

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

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

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

Answer: D

Explanation:

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

NEW QUESTION 8

A company requires that its internally facing web application be highly available. The architecture is made up of one Amazon EC2 web server instance and one NAT instance that provides outbound internet access for updates and accessing public data.

Which combination of architecture adjustments should the company implement to achieve high availability? (Choose two.)

- A. Add the NAT instance to an EC2 Auto Scaling group that spans multiple Availability Zone
- B. Update the route tables.
- C. Create additional EC2 instances spanning multiple Availability Zone
- D. Add an Application Load Balancer to split the load between them.
- E. Configure an Application Load Balancer in front of the EC2 instance
- F. Configure Amazon CloudWatch alarms to recover the EC2 instance upon host failure.
- G. Replace the NAT instance with a NAT gateway in each Availability Zone
- H. Update the route tables.
- I. Replace the NAT instance with a NAT gateway that spans multiple Availability Zone
- J. Update the route tables.

Answer: BD

Explanation:

<https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-gateway.html>

NEW QUESTION 9

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 10

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

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

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

Answer: BCE

Explanation:

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

NEW QUESTION 10

A company 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 11

A company has a legacy application A DevOps engineer needs to automate the process of building the deployable artifact for the legacy application. The solution must store the deployable artifact in an existing Amazon S3 bucket for future deployments to reference Which solution will meet these requirements in the MOST operationally efficient way?

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

Answer: A

Explanation:

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

NEW QUESTION 14

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 bucket1.

? 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 19

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 22

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 27

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 grou

J. Pass the ARNs for the parameters from Secrets Manager to join the domai

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 Manage

M. In the CloudFormation template, update the EC2 launch template to include user dat

N. Configure the user data to pull the administrator credentials from Secrets Manager and to join the AWS Managed Microsoft AD domai

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 30

A company builds an application that uses an Application Load Balancer in front of Amazon EC2 instances that are in an Auto Scaling group. The application is stateless. The Auto Scaling group uses a custom AMI that is fully prebuilt. The EC2 instances do not have a custom bootstrapping process.

The AMI that the Auto Scaling group uses was recently deleted. The Auto Scaling group's scaling activities show failures because the AMI ID does not exist.

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

A. Create a new launch template that uses the new AMI.

B. Update the Auto Scaling group to use the new launch template.

C. Reduce the Auto Scaling group's desired capacity to 0.

D. Increase the Auto Scaling group's desired capacity by 1.

E. Create a new AMI from a running EC2 instance in the Auto Scaling group.

F. Create a new AMI by copying the most recent public AMI of the operating system that the EC2 instances use.

Answer: ABF

Explanation:

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

NEW QUESTION 32

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 EC2. 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 35

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 36

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=LzvhTcqyog>

NEW QUESTION 41

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

Launches a second fleet of instances with the same capacity as the original fleet. Maintains the original fleet unchanged while the second fleet is launched.

Transitions traffic to the second fleet when the second fleet is fully deployed. Terminates the original fleet automatically 1 hour after transition.

Which solution will satisfy these requirements?

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

Answer: C

Explanation:

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

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

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

NEW QUESTION 44

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

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

- A. Create a single AWS CodePipeline pipeline that deploys the application in parallel using unique AWS CodeDeploy applications and deployment groups created

for each ALB-Auto Scaling group pair.

B. Create a single AWS CodePipeline pipeline that deploys the application using a single AWS CodeDeploy application and single deployment group.

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

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

Answer: C

Explanation:

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

NEW QUESTION 49

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 51

A company wants to use a grid system for a proprietary enterprise 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 cron 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

L. Trigger whenever a new instance is added to the cluster.

Answer: A

Explanation:

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

NEW QUESTION 55

A company manages multiple AWS accounts in AWS Organizations. The company's security policy states that AWS account root user credentials for member accounts must not be used. The company monitors access to the root user credentials.

A recent alert shows that the root user in a member account launched an Amazon EC2 instance. A DevOps engineer must create an SCP at the organization's root level that will prevent the root user in member accounts from making any AWS service API calls.

Which SCP will meet these requirements?

A)

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "*",
      "Resource": "*",
      "Condition": {
        "StringNotLike": { "aws:PrincipalArn": "arn:aws:iam::*:root" }
      }
    }
  ]
}
```

B)

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": "*",
      "Resource": "*",
      "Principal": { "AWS": "arn:aws:iam::*:root" }
    }
  ]
}
```

C)

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": "*",
      "Resource": "*",
      "Condition": {
        "StringLike": { "aws:PrincipalArn": "arn:aws:iam::*:root" }
      }
    }
  ]
}
```

D)

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "*",
      "Resource": "*",
      "Principal": "root"
    }
  ]
}
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: D

NEW QUESTION 56

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 60

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 61

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

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

Answer: D

Explanation:

This solution will meet the requirements because it will use Amazon S3 as a common data lake for both the application logs and the API logs. Amazon S3 is a service that provides scalable, durable, and secure object storage for any type of data. You can use the Amazon CloudWatch agent to send logs from your EC2 instances to S3 buckets, and use AWS CloudTrail to deliver the API logs to S3 buckets as well. You can also use Amazon Athena to query both sets of logs in S3 using standard SQL, without loading or transforming them. Athena is a serverless interactive query service that allows you to analyze data in S3 using a variety of data formats, such as JSON, CSV, Parquet, and ORC.

NEW QUESTION 64

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

Which of the following should successfully install the application while complying with the new rule?

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

Answer: C

Explanation:

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

NEW QUESTION 68

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

A DevOps engineer needs to identify which applications are the source of the increased logging costs.

Which solution will meet these requirements?

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

Answer: C

Explanation:

The correct answer is C.

A comprehensive and detailed explanation is:

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

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

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

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

References:

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

NEW QUESTION 71

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

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

Which solution will meet these requirements?

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

- T. Configure AWS Systems Manager Compliance to scan all managed node
- . Create an Amazon EventBridge rule to capture a configuration compliance state change even
- . Use the event to invoke the CodeBuild project.

Answer: A

Explanation:

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

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

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

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

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

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

NEW QUESTION 75

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

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

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

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

Answer: C

Explanation:

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

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

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

NEW QUESTION 77

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

Which action will accomplish this?

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

Answer: B

Explanation:

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

NEW QUESTION 81

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

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

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

- A. Add a DeletionPolicy attribute to the S3 bucket resource, with the value Delete forcing the bucket to be removed when the stack is deleted.

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

Answer: B

Explanation:

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

NEW QUESTION 84

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

Which solution will meet these requirements?

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

Answer: B

Explanation:

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

NEW QUESTION 88

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

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

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

Which solution will meet these requirements?

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

Answer: C

Explanation:

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

? Configure provisioned concurrency on the Lambda function.

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

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

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

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

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

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

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

NEW QUESTION 91

A company's DevOps engineer is creating an AWS Lambda function to process notifications from an Amazon Simple Notification Service (Amazon SNS) topic. The Lambda function will process the notification messages and will write the contents of the notification messages to an Amazon RDS Multi-AZ DB instance.

During testing a database administrator accidentally shut down the DB instance. While the database was down the company lost several of the SNS notification messages that were delivered during that time.

The DevOps engineer needs to prevent the loss of notification messages in the future Which solutions will meet this requirement? (Select TWO.)

- A. Replace the RDS Multi-AZ DB instance with an Amazon DynamoDB table.
- B. Configure an Amazon Simple Queue Service (Amazon SQS) queue as a destination of the Lambda function.
- C. Configure an Amazon Simple Queue Service (Amazon SQS) dead-letter queue for the SNS topic.
- D. Subscribe an Amazon Simple Queue Service (Amazon SQS) queue to the SNS topic Configure the Lambda function to process messages from the SQS queue.
- E. Replace the SNS topic with an Amazon EventBridge event bus Configure an EventBridge rule on the new event bus to invoke the Lambda function for each event.

Answer: CD

Explanation:

These solutions will meet the requirement because they will prevent the loss of notification messages in the future. An Amazon SQS queue is a service that provides a reliable, scalable, and secure message queue for asynchronous communication between distributed components. You can use an SQS queue to buffer messages from an SNS topic and ensure that they are delivered and processed by a Lambda function, even if the function or the database is temporarily unavailable.

Option C will configure an SQS dead-letter queue for the SNS topic. A dead-letter queue is a queue that receives messages that could not be delivered to any subscriber after a specified number of retries. You can use a dead-letter queue to store and analyze failed messages, or to reprocess them later. This way, you can avoid losing messages that could not be delivered to the Lambda function due to network errors, throttling, or other issues. Option D will subscribe an SQS queue to the SNS topic and configure the Lambda function to process messages from the SQS queue. This will decouple the SNS topic from the Lambda function and provide more flexibility and control over the message delivery and processing. You can use an SQS queue to store messages from the SNS topic until they are ready to be processed by the Lambda function, and also to retry processing in case of failures. This way, you can avoid losing messages that could not be processed by the Lambda function due to database errors, timeouts, or other issues.

NEW QUESTION 93

A company's application development team uses Linux-based Amazon EC2 instances as bastion hosts. Inbound SSH access to the bastion hosts is restricted to specific IP addresses, as defined in the associated security groups. The company's security team wants to receive a notification if the security group rules are modified to allow SSH access from any IP address.

What should a DevOps engineer do to meet this requirement?

- A. Create an Amazon EventBridge rule with a source of aws.cloudtrail and the event name AuthorizeSecurityGroupIngres
- B. Define an Amazon Simple Notification Service (Amazon SNS) topic as the target.
- C. Enable Amazon GuardDuty and check the findings for security groups in AWS Security Hub
- D. Configure an Amazon EventBridge rule with a custom pattern that matches GuardDuty events with an output of NON_COMPLIAN
- E. Define an Amazon Simple Notification Service (Amazon SNS) topic as the target.
- F. Create an AWS Config rule by using the restricted-ssh managed rule to check whether security groups disallow unrestricted incoming SSH traffic
- G. Configure automatic remediation to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic.
- H. Enable Amazon Inspector
- I. Include the Common Vulnerabilities and Exposures-1.1 rules package to check the security groups that are associated with the bastion host
- J. Configure Amazon Inspector to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic.

Answer: A

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/monitor-security-group-changes-ec2/>

NEW QUESTION 97

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

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

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

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

Answer: B

Explanation:

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

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

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

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

References:

- ? AWS CodeDeploy Deployment Configurations
- ? [AWS CodeDeploy Rollbacks]
- ? Amazon CloudWatch Alarms

NEW QUESTION 100

A company has multiple AWS accounts. The company uses AWS IAM Identity Center (AWS Single Sign-On) that is integrated with AWS Toolkit for Microsoft

Azure DevOps. The attributes for access control feature is enabled in IAM Identity Center.

The attribute mapping list contains two entries. The department key is mapped to

`${path:enterprise.department}`. The costCenter key is mapped to

`${path:enterprise.costCenter}`.

All existing Amazon EC2 instances have a department tag that corresponds to three company departments (d1, d2, d3). A DevOps engineer must create policies based on the matching attributes. The policies must minimize administrative effort and must grant each Azure AD user access to only the EC2 instances that are tagged with the user's respective department name.

Which condition key should the DevOps engineer include in the custom permissions policies to meet these requirements?

A.

```
"Condition": {
  "ForAllValues:StringEquals": {
    "aws:TagKeys": ["department"]
  }
}
```

B.

```
"Condition": {
  "StringEquals": {
    "aws:PrincipalTag/department": "${aws:ResourceTag/department}"
  }
}
```

C.

```
"Condition": {
  "StringEquals": {
    "ec2:ResourceTag/department": "${aws:PrincipalTag/department}"
  }
}
```

D.

```
"Condition": {
  "ForAllValues:StringEquals": {
    "ec2:ResourceTag/department": ["d1", "d2", "d3"]
  }
}
```

A.

Answer: C

Explanation:

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

NEW QUESTION 101

An Amazon EC2 instance is running in a VPC and needs to download an object from a restricted Amazon S3 bucket. When the DevOps engineer tries to download the object, an AccessDenied error is received,

What are the possible causes for this error? (Select TWO,)

- A. The S3 bucket default encryption is enabled.
- B. There is an error in the S3 bucket policy.
- C. The object has been moved to S3 Glacier.
- D. There is an error in the IAM role configuration.
- E. S3 Versioning is enabled.

Answer: BD

Explanation:

These are the possible causes for the AccessDenied error because they affect the permissions to access the S3 object from the EC2 instance. An S3 bucket policy is a resource-based policy that defines who can access the bucket and its objects, and what actions they can perform. An IAM role is an identity that can be assumed by an EC2 instance to grant it permissions to access AWS services and resources. If there is an error in the S3 bucket policy or the IAM role configuration, such as a missing or incorrect statement, condition, or principal, then the EC2 instance may not have the necessary permissions to download the object from the S3 bucket . <https://docs.aws.amazon.com/AmazonS3/latest/userguide/example-bucket-policies.html>
<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html>

NEW QUESTION 105

A DevOps team manages an API running on-premises that serves as a backend for an Amazon API Gateway endpoint. Customers have been complaining about

high response latencies, which the development team has verified using the API Gateway latency metrics in Amazon CloudWatch. To identify the cause, the team needs to collect relevant data without introducing additional latency. Which actions should be taken to accomplish this? (Choose two.)

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

Answer: AC

Explanation:

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

NEW QUESTION 108

A DevOps engineer needs to back up sensitive Amazon S3 objects that are stored within an S3 bucket with a private bucket policy using S3 cross-Region replication functionality. The objects need to be copied to a target bucket in a different AWS Region and account. Which combination of actions should be performed to enable this replication? (Choose three.)

- A. Create a replication IAM role in the source account
- B. Create a replication IAM role in the target account.
- C. Add statements to the source bucket policy allowing the replication IAM role to replicate objects.
- D. Add statements to the target bucket policy allowing the replication IAM role to replicate objects.
- E. Create a replication rule in the source bucket to enable the replication.
- F. Create a replication rule in the target bucket to enable the replication.

Answer: ADE

Explanation:

S3 cross-Region replication (CRR) automatically replicates data between buckets across different AWS Regions. To enable CRR, you need to add a replication configuration to your source bucket that specifies the destination bucket, the IAM role, and the encryption type (optional). You also need to grant permissions to the IAM role to perform replication actions on both the source and destination buckets. Additionally, you can choose the destination storage class and enable additional replication options such as S3 Replication Time Control (S3 RTC) or S3 Batch Replication. <https://medium.com/cloud-techies/s3-same-region-replication-srr-and-cross-region-replication-crr-34d446806bab> <https://aws.amazon.com/getting-started/hands-on/replicate-data-using-amazon-s3-replication/> <https://docs.aws.amazon.com/AmazonS3/latest/userguide/replication.html>

NEW QUESTION 113

A video-sharing company stores its videos in Amazon S3. The company has observed a sudden increase in video access requests, but the company does not know which videos are most popular. The company needs to identify the general access pattern for the video files. This pattern includes the number of users who access a certain file on a given day, as well as the number of times a file is accessed. 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? or of pull requests for certain files. How can the company meet these requirements with the LEAST amount of effort?

- A. Activate S3 server access logging
- B. Import the access logs into an Amazon Aurora database
- C. Use an Aurora SQL query to analyze the access patterns.
- D. Activate S3 server access logging
- E. Use Amazon Athena to create an external table with the log file
- F. Use Athena to create a SQL query to analyze the access patterns.
- G. Invoke an AWS Lambda function for every S3 object access event
- H. Configure the Lambda function to write the file access information, such as user, S3 bucket, and file key, to an Amazon Kinesis Data Analytics for SQL application
- I. S3 bucket, and file key, to an Amazon Aurora database
- J. Use an Aurora SQL query to analyze the access patterns.
- K. Record an Amazon CloudWatch Logs log message for every S3 object access event
- L. Configure a CloudWatch Logs log stream to write the file access information, such as user, S3 bucket, and file key, to an Amazon Kinesis Data Analytics for SQL application
- M. Perform a sliding window analysis.

Answer: B

Explanation:

Activating S3 server access logging and using Amazon Athena to create an external table with the log files is the easiest and most cost-effective way to analyze access patterns. This option requires minimal setup and allows for quick analysis of the access patterns with SQL queries. Additionally, Amazon Athena scales automatically to match the query load, so there is no need for additional infrastructure provisioning or management.

NEW QUESTION 118

A company's application is currently deployed to a single AWS Region. Recently, the company opened a new office on a different continent. The users in the new office are experiencing high latency. The company's application runs on Amazon EC2 instances behind an Application Load Balancer (ALB) and uses Amazon DynamoDB as the database layer. The instances run in an EC2 Auto Scaling group across multiple Availability Zones. A DevOps engineer is tasked with minimizing application response times and improving availability for users in both Regions. Which combination of actions should be taken to address the latency issues? (Choose three.)

- A. Create a new DynamoDB table in the new Region with cross-Region replication enabled.
- B. Create new ALB and Auto Scaling group global resources and configure the new ALB to direct traffic to the new Auto Scaling group.

- C. Create new ALB and Auto Scaling group resources in the new Region and configure the new ALB to direct traffic to the new Auto Scaling group.
- D. Create Amazon Route 53 records, health checks, and latency-based routing policies to route to the ALB.
- E. Create Amazon Route 53 aliases, health checks, and failover routing policies to route to the ALB.
- F. Convert the DynamoDB table to a global table.

Answer: CDF

Explanation:

C. Create new ALB and Auto Scaling group resources in the new Region and configure the new ALB to direct traffic to the new Auto Scaling group. This will allow users in the new Region to access the application with lower latency by reducing the network hops between the user and the application servers.

* D. Create Amazon Route 53 records, health checks, and latency-based routing policies to route to the ALB. This will enable Route 53 to route user traffic to the nearest healthy ALB, based on the latency between the user and the ALBs.

* F. Convert the DynamoDB table to a global table. This will enable reads and writes to the table in both Regions with low latency, improving the overall response time of the application

NEW QUESTION 121

A DevOps engineer is creating an AWS CloudFormation template to deploy a web service. The web service will run on Amazon EC2 instances in a private subnet behind an Application Load Balancer (ALB). The DevOps engineer must ensure that the service can accept requests from clients that have IPv6 addresses. What should the DevOps engineer do with the CloudFormation template so that IPv6 clients can access the web service?

- A. Add an IPv6 CIDR block to the VPC and the private subnet for the EC2 instance
- B. Create route table entries for the IPv6 network, use EC2 instance types that support IPv6, and assign IPv6 addresses to each EC2 instance.
- C. Assign each EC2 instance an IPv6 Elastic IP address
- D. Create a target group, and add the EC2 instances as target
- E. Create a listener on port 443 of the ALB, and associate the target group with the ALB.
- F. Replace the ALB with a Network Load Balancer (NLB). Add an IPv6 CIDR block to the VPC and subnets for the NLB, and assign the NLB an IPv6 Elastic IP address.
- G. Add an IPv6 CIDR block to the VPC and subnets for the AL
- H. Create a listener on port 443. and specify the dualstack IP address type on the AL
- I. Create a target group, and add the EC2 instances as target
- J. Associate the target group with the ALB.

Answer: D

Explanation:

it involves adding an IPv6 CIDR block to the VPC and subnets for the ALB and specifying the dualstack IP address type on the ALB listener. This allows the ALB to listen on both IPv4 and IPv6 addresses, and forward requests to the EC2 instances that are added as targets to the target group associated with the ALB.

NEW QUESTION 123

A company sells products through an ecommerce web application The company wants a dashboard that shows a pie chart of product transaction details. The company wants to integrate the dashboard With the company's existing Amazon CloudWatch dashboards Which solution Will meet these requirements With the MOST operational efficiency?

- A. Update the ecommerce application to emit a JSON object to a CloudWatch log group for each processed transactio
- B. Use CloudWatch Logs Insights to query the log group and to visualize the results in a pie chart format Attach the results to the desired CloudWatch dashboard.
- C. Update the ecommerce application to emit a JSON object to an Amazon S3 bucket for each processed transactio
- D. Use Amazon Athena to query the S3 bucket and to visualize the results In a Pie chart forma
- E. Export the results from Athena Attach the results to the desired CloudWatch dashboard
- F. Update the ecommerce application to use AWS X-Ray for instrumentatio
- G. Create a new X-Ray subsegment Add an annotation for each processed transactio
- H. Use X-Ray traces to query the data and to visualize the results in a pie chart format Attach the results to the desired CloudWatch dashboard
- I. Update the ecommerce application to emit a JSON object to a CloudWatch log group for each processed transaction_ Create an AWS Lambda function to aggregate and write the results to Amazon DynamoD
- J. Create a Lambda subscription filter for the log fil
- K. Attach the results to the desired CloudWatch dashboard.

Answer: A

Explanation:

The correct answer is A.

A comprehensive and detailed explanation is:

? Option A is correct because it meets the requirements with the most operational efficiency. Updating the ecommerce application to emit a JSON object to a CloudWatch log group for each processed transaction is a simple and cost- effective way to collect the data needed for the dashboard. Using CloudWatch Logs Insights to query the log group and to visualize the results in a pie chart format is also a convenient and integrated solution that leverages the existing CloudWatch dashboards. Attaching the results to the desired CloudWatch dashboard is straightforward and does not require any additional steps or services.

? Option B is incorrect because it introduces unnecessary complexity and cost.

Updating the ecommerce application to emit a JSON object to an Amazon S3 bucket for each processed transaction is a valid way to store the data, but it requires creating and managing an S3 bucket and its permissions. Using Amazon Athena to query the S3 bucket and to visualize the results in a pie chart format is also a valid way to analyze the data, but it incurs charges based on the amount of data scanned by each query. Exporting the results from Athena and attaching them to the desired CloudWatch dashboard is also an extra step that adds more overhead and latency.

? Option C is incorrect because it uses AWS X-Ray for an inappropriate purpose.

Updating the ecommerce application to use AWS X-Ray for instrumentation is a good practice for monitoring and tracing distributed applications, but it is not designed for aggregating product transaction details. Creating a new X-Ray subsegment and adding an annotation for each processed transaction is possible, but it would clutter the X-Ray service map and make it harder to debug performance issues. Using X-Ray traces to query the data and to visualize the results in a pie chart format is also possible, but it would require custom code and logic that are not supported by X-Ray natively. Attaching the results to the desired CloudWatch dashboard is also not supported by X-Ray directly, and would require additional steps or services.

? Option D is incorrect because it introduces unnecessary complexity and cost.

Updating the ecommerce application to emit a JSON object to a CloudWatch log group for each processed transaction is a simple and cost-effective way to collect the data needed for the dashboard, as in option A. However, creating an AWS Lambda function to aggregate and write the results to Amazon DynamoDB is redundant, as CloudWatch Logs Insights can already perform aggregation queries on log data. Creating a Lambda subscription filter for the log file is also

redundant, as CloudWatch Logs Insights can already access log data directly. Attaching the results to the desired CloudWatch dashboard would also require additional steps or services, as DynamoDB does not support native integration with CloudWatch dashboards.

References:

- ? CloudWatch Logs Insights
- ? Amazon Athena
- ? AWS X-Ray
- ? AWS Lambda
- ? Amazon DynamoDB

NEW QUESTION 125

A company uses an Amazon API Gateway regional REST API to host its application API. The REST API has a custom domain. The REST API's default endpoint is deactivated.

The company's internal teams consume the API. The company wants to use mutual TLS between the API and the internal teams as an additional layer of authentication.

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

- A. Use AWS Certificate Manager (ACM) to create a private certificate authority (CA). Provision a client certificate that is signed by the private CA.
- B. Provision a client certificate that is signed by a public certificate authority (CA). Import the certificate into AWS Certificate Manager (ACM).
- C. Upload the provisioned client certificate to an Amazon S3 bucket
- D. Configure the API Gateway mutual TLS to use the client certificate that is stored in the S3 bucket as the trust store.
- E. Upload the provisioned client certificate private key to an Amazon S3 bucket
- F. Configure the API Gateway mutual TLS to use the private key that is stored in the S3 bucket as the trust store.
- G. Upload the root private certificate authority (CA) certificate to an Amazon S3 bucket
- H. Configure the API Gateway mutual TLS to use the private CA certificate that is stored in the S3 bucket as the trust store.

Answer: AE

Explanation:

Mutual TLS (mTLS) authentication requires two-way authentication between the client and the server. For Amazon API Gateway, you can enable mTLS for a custom domain name, which requires clients to present X.509 certificates to verify their identity to access your API. To set up mTLS, you would typically use AWS Certificate Manager (ACM) to create a private certificate authority (CA) and provision a client certificate signed by this private CA. The root CA certificate is then uploaded to an Amazon S3 bucket and configured in API Gateway as the trust store¹².

References:

- ? Introducing mutual TLS authentication for Amazon API Gateway¹.
- ? Configuring mutual TLS authentication for a REST API².
- ? AWS Private Certificate Authority details³.
- ? AWS Certificate Manager Private Certificate Authority updates⁴.

NEW QUESTION 126

A company is launching an application. The application must use only approved AWS services. The account that runs the application was created less than 1 year ago and is assigned to an AWS Organizations OU.

The company needs to create a new Organizations account structure. The account structure must have an appropriate SCP that supports the use of only services that are currently active in the AWS account.

The company will use AWS Identity and Access Management (IAM) Access Analyzer in the solution.

Which solution will meet these requirements?

- A. Create an SCP that allows the services that IAM Access Analyzer identifies
- B. Create an OU for the account
- C. Move the account into the new OU
- D. Attach the new SCP to the new OU
- E. Detach the default FullAWSAccess SCP from the new OU.
- F. Create an SCP that denies the services that IAM Access Analyzer identifies
- G. Create an OU for the account
- H. Move the account into the new OU
- I. Attach the new SCP to the new OU.
- J. Create an SCP that allows the services that IAM Access Analyzer identifies
- K. Attach the new SCP to the organization's root.
- L. Create an SCP that allows the services that IAM Access Analyzer identifies
- M. Create an OU for the account
- N. Move the account into the new OU
- O. Attach the new SCP to the management account
- P. Detach the default FullAWSAccess SCP from the new OU.

Answer: A

Explanation:

To meet the requirements of creating a new Organizations account structure with an appropriate SCP that supports the use of only services that are currently active in the AWS account, the company should use the following solution:

? Create an SCP that allows the services that IAM Access Analyzer identifies. IAM Access Analyzer is a service that helps identify potential resource-access risks by analyzing resource-based policies in the AWS environment. IAM Access Analyzer can also generate IAM policies based on access activity in the AWS CloudTrail logs. By using IAM Access Analyzer, the company can create an SCP that grants only the permissions that are required for the application to run, and denies all other services. This way, the company can enforce the use of only approved AWS services and reduce the risk of unauthorized access¹²

? Create an OU for the account. Move the account into the new OU. An OU is a container for accounts within an organization that enables you to group accounts that have similar business or security requirements. By creating an OU for the account, the company can apply policies and manage settings for the account as a group. The company should move the account into the new OU to make it subject to the policies attached to the OU³

? Attach the new SCP to the new OU. Detach the default FullAWSAccess SCP from the new OU. An SCP is a type of policy that specifies the maximum permissions for an organization or organizational unit (OU). By attaching the new SCP to the new OU, the company can restrict the services that are available to all accounts in that OU, including the account that runs the application. The company should also detach the default FullAWSAccess SCP from the new OU, because this policy allows all actions on all AWS services and might override or conflict with the new SCP⁴⁵

The other options are not correct because they do not meet the requirements or follow best practices. Creating an SCP that denies the services that IAM Access Analyzer identifies is not a good option because it might not cover all possible services that are not approved or required for the application. A deny policy is also

more difficult to maintain and update than an allow policy. Creating an SCP that allows the services that IAM Access Analyzer identifies and attaching it to the organization's root is not a good option because it might affect other accounts and OUs in the organization that have different service requirements or approvals. Creating an SCP that allows the services that IAM Access Analyzer identifies and attaching it to the management account is not a valid option because SCPs cannot be attached directly to accounts, only to OUs or roots.

References:

- ? 1: Using AWS Identity and Access Management Access Analyzer - AWS Identity and Access Management
- ? 2: Generate a policy based on access activity - AWS Identity and Access Management
- ? 3: Organizing your accounts into OUs - AWS Organizations
- ? 4: Service control policies - AWS Organizations
- ? 5: How SCPs work - AWS Organizations

NEW QUESTION 129

A DevOps engineer has developed an AWS Lambda function. The Lambda function starts an AWS CloudFormation drift detection operation on all supported resources for a specific CloudFormation stack. The Lambda function then exits its invocation. The DevOps engineer has created an Amazon EventBridge scheduled rule that invokes the Lambda function every hour. An Amazon Simple Notification Service (Amazon SNS) topic already exists in the AWS account. The DevOps engineer has subscribed to the SNS topic to receive notifications.

The DevOps engineer needs to receive a notification as soon as possible when drift is detected in this specific stack configuration.

Which solution will meet these requirements?

- A. Configure the existing EventBridge rule to also target the SNS topic. Configure an SNS subscription filter policy to match the CloudFormation stack.
- B. Attach the subscription filter policy to the SNS topic.
- C. Create a second Lambda function to query the CloudFormation API for the drift detection results for the stack. Configure the second Lambda function to publish a message to the SNS topic. If drift is detected, adjust the existing EventBridge rule to also target the second Lambda function.
- D. Configure Amazon GuardDuty in the account with drift detection for all CloudFormation stacks.
- E. Create a second EventBridge rule that reacts to the GuardDuty drift detection event finding for the specific CloudFormation stack.
- F. Configure the SNS topic as a target of the second EventBridge rule.
- G. Configure AWS Config in the account.
- H. Use the cloudformation-stack-drift-detection-check managed rule.
- I. Create a second EventBridge rule that reacts to a compliance change event for the CloudFormation stack.
- J. Configure the SNS topic as a target of the second EventBridge rule.

Answer: D

Explanation:

A comprehensive and detailed explanation is:

? Option A is incorrect because EventBridge rules cannot filter events based on the message body or attributes of the target service. Therefore, configuring an SNS subscription filter policy to match the CloudFormation stack will not work. The SNS topic will receive all events from the EventBridge rule, regardless of the stack name or drift status.

? Option B is incorrect because it introduces unnecessary complexity and cost.

Creating a second Lambda function to query the CloudFormation API for the drift detection results is redundant, since CloudFormation already publishes drift detection events to EventBridge. Moreover, invoking two Lambda functions every hour will incur more charges than invoking one.

? Option C is incorrect because GuardDuty does not provide drift detection for CloudFormation stacks. GuardDuty is a threat detection service that monitors for malicious activity and unauthorized behavior in AWS accounts and workloads. It does not monitor or report on configuration changes or drifts in CloudFormation stacks.

? Option D is correct because it leverages AWS Config and its managed rule for drift detection. AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. It can detect configuration changes and drifts in CloudFormation stacks using the cloudformation-stack-drift-detection-check managed rule. This rule triggers an AWS Config event when a stack drifts from its expected template configuration. By creating a second EventBridge rule that reacts to this event for the specific stack, the DevOps engineer can configure the SNS topic as a target and receive a notification as soon as possible when drift is detected.

References:

- ? AWS Config
- ? Amazon SNS subscription filter policies
- ? Amazon EventBridge rules

NEW QUESTION 133

A company has enabled all features for its organization in AWS Organizations. The organization contains 10 AWS accounts. The company has turned on AWS CloudTrail in all the accounts. The company expects the number of AWS accounts in the organization to increase to 500 during the next year. The company plans to use multiple OUs for these accounts.

The company has enabled AWS Config in each existing AWS account in the organization.

A DevOps engineer must implement a solution that enables AWS Config automatically for all future AWS accounts that are created in the organization.

Which solution will meet this requirement?

- A. In the organization's management account, create an Amazon EventBridge rule that reacts to a CreateAccount API call.
- B. Configure the rule to invoke an AWS Lambda function that enables trusted access to AWS Config for the organization.
- C. In the organization's management account, create an AWS CloudFormation stack set to enable AWS Config.
- D. Configure the stack set to deploy automatically when an account is created through Organizations.
- E. In the organization's management account, create an SCP that allows the appropriate AWS Config API calls to enable AWS Config.
- F. Apply the SCP to the root-level OU.
- G. In the organization's management account, create an Amazon EventBridge rule that reacts to a CreateAccount API call.
- H. Configure the rule to invoke an AWS Systems Manager Automation runbook to enable AWS Config for the account.

Answer: B

Explanation:

<https://aws.amazon.com/about-aws/whats-new/2020/02/aws-cloudformation-stacksets-introduces-automatic-deployments-across-accounts-and-regions-through-aws-organizations/>

NEW QUESTION 138

A company recently launched multiple applications that use Application Load Balancers. Application response time often slows down when the applications experience problems. A DevOps engineer needs to implement a monitoring solution that alerts the company when the applications begin to perform slowly. The

DevOps engineer creates an Amazon Simple Notification Service (Amazon SNS) topic and subscribe the company's email address to the topic. What should the DevOps engineer do next to meet the requirements?

- A. Create an Amazon EventBridge rule that invokes an AWS Lambda function to query the applications on a 5-minute interval. Configure the Lambda function to publish a notification to the SNS topic when the applications return errors.
- B. Create an Amazon CloudWatch Synthetics canary that runs a custom script to query the applications on a 5-minute interval.
- C. Configure the canary to use the SNS topic when the applications return errors.
- D. Create an Amazon CloudWatch alarm that uses the AWS/ApplicationELB namespace RequestCountPerTarget metric. Configure the CloudWatch alarm to send a notification when the number of connections becomes greater than the configured number of threads that the application supports. Configure the CloudWatch alarm to use the SNS topic.
- E. Create an Amazon CloudWatch alarm that uses the AWS/ApplicationELB namespace RequestCountPerTarget metric. Configure the CloudWatch alarm to send a notification when the average response time becomes greater than the longest response time that the application supports. Configure the CloudWatch alarm to use the SNS topic.

Answer: B

Explanation:

? Option A is incorrect because creating an Amazon EventBridge rule that invokes an AWS Lambda function to query the applications on a 5-minute interval is not a valid solution. EventBridge rules can only trigger Lambda functions based on events, not on time intervals. Moreover, querying the applications on a 5-minute interval might incur unnecessary costs and network overhead, and might not detect performance issues in real time.

? Option B is correct because creating an Amazon CloudWatch Synthetics canary that runs a custom script to query the applications on a 5-minute interval is a valid solution. CloudWatch Synthetics canaries are configurable scripts that monitor endpoints and APIs by simulating customer behavior. Canaries can run as often as once per minute, and can measure the latency and availability of the applications. Canaries can also send notifications to an Amazon SNS topic when they detect errors or performance issues¹.

? Option C is incorrect because creating an Amazon CloudWatch alarm that uses the AWS/ApplicationELB namespace RequestCountPerTarget metric is not a valid solution. The RequestCountPerTarget metric measures the number of requests completed or connections made per target in a target group². This metric does not reflect the application response time, which is the requirement. Moreover, configuring the CloudWatch alarm to send a notification when the number of connections becomes greater than the configured number of threads that the application supports is not a valid way to measure the application performance, as it depends on the application design and implementation.

? Option D is incorrect because creating an Amazon CloudWatch alarm that uses the AWS/ApplicationELB namespace RequestCountPerTarget metric is not a valid solution, for the same reason as option C. The RequestCountPerTarget metric does not reflect the application response time, which is the requirement. Moreover, configuring the CloudWatch alarm to send a notification when the average response time becomes greater than the longest response time that the application supports is not a valid way to measure the application performance, as it does not account for variability or outliers in the response time distribution.

References:

? 1: Using synthetic monitoring

? 2: Application Load Balancer metrics

NEW QUESTION 141

.....

THANKS FOR TRYING THE DEMO OF OUR PRODUCT

Visit Our Site to Purchase the Full Set of Actual AWS-Certified-DevOps-Engineer-Professional Exam Questions With Answers.

We Also Provide Practice Exam Software That Simulates Real Exam Environment And Has Many Self-Assessment Features. Order the AWS-Certified-DevOps-Engineer-Professional Product From:

<https://www.2passeasy.com/dumps/AWS-Certified-DevOps-Engineer-Professional/>

Money Back Guarantee

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