

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

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

Answer: B

Explanation:

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

NEW QUESTION 2

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

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

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

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

Answer: BE

NEW QUESTION 3

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 4

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

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

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

Answer: BD

Explanation:

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

NEW QUESTION 5

A company uses an organization in AWS Organizations to manage its AWS accounts. The company recently acquired another company that has standalone AWS accounts. The acquiring company's DevOps team needs to consolidate the administration of the AWS accounts for both companies and retain full administrative control of the accounts. The DevOps team also needs to collect and group findings across all the accounts to implement and maintain a security posture.

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

- A. Invite the acquired company's AWS accounts to join the organization
- B. Create an SCP that has full administrative privilege
- C. Attach the SCP to the management account.
- D. Invite the acquired company's AWS accounts to join the organization
- E. Create the OrganizationAccountAccessRole IAM role in the invited account
- F. Grant permission to the management account to assume the role.
- G. Use AWS Security Hub to collect and group findings across all accounts
- H. Use Security Hub to automatically detect new accounts as the accounts are added to the organization.
- I. Use AWS Firewall Manager to collect and group findings across all accounts
- J. Enable all features for the organization
- K. Designate an account in the organization as the delegated administrator account for Firewall Manager.
- L. Use Amazon Inspector to collect and group findings across all accounts
- M. Designate an account in the organization as the delegated administrator account for Amazon Inspector.

Answer: BC

Explanation:

The correct answer is B and C. Option B is correct because inviting the acquired company's AWS accounts to join the organization and creating the OrganizationAccountAccessRole IAM role in the invited accounts allows the management account to assume the role and gain full administrative access to the member accounts. Option C is correct because using AWS Security Hub to collect and group findings across all accounts enables the DevOps team to monitor and improve the security posture of the organization. Security Hub can automatically detect new accounts as the accounts are added to the organization and enable Security Hub for them. Option A is incorrect because creating an SCP that has full administrative privileges and attaching it to the management account does not grant the management account access to the member accounts. SCPs are used to restrict the permissions of the member accounts, not to grant permissions to the management account. Option D is incorrect because using AWS Firewall Manager to collect and group findings across all accounts is not a valid use case for Firewall Manager. Firewall Manager is used to centrally configure and manage firewall rules across the organization, not to collect and group security findings. Option E is incorrect because using Amazon Inspector to collect and group findings across all accounts is not a valid use case for Amazon Inspector. Amazon Inspector is used to assess the security and compliance of applications running on Amazon EC2 instances, not to collect and group security findings across accounts. References:

- ? Inviting an AWS account to join your organization
- ? Enabling and disabling AWS Security Hub
- ? Service control policies
- ? AWS Firewall Manager
- ? Amazon Inspector

NEW QUESTION 6

A DevOps engineer used an AWS CloudFormation custom resource to set up AD Connector. The AWS Lambda function ran and created AD Connector, but CloudFormation is not transitioning from CREATE_IN_PROGRESS to CREATE_COMPLETE.

Which action should the engineer take to resolve this issue?

- A. Ensure the Lambda function code has exited successfully.
- B. Ensure the Lambda function code returns a response to the pre-signed URL.
- C. Ensure the Lambda function IAM role has cloudformation UpdateStack permissions for the stack ARN.
- D. Ensure the Lambda function IAM role has ds ConnectDirectory permissions for the AWS account.

Answer: B

Explanation:

Reference: <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/crpg-ref-responses.html>

NEW QUESTION 7

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

What should the company do to accomplish these goals?

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

Answer: D

Explanation:

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

NEW QUESTION 8

A company has deployed an application in a production VPC in a single AWS account. The application is popular and is experiencing heavy usage. The company's security team wants to add additional security, such as AWS WAF, to the application deployment. However, the application's product manager is concerned about cost and does not want to approve the change unless the security team can prove that additional security is necessary. The security team believes that some of the application's demand might come from users that have IP addresses that are on a deny list. The security team provides the deny list to a DevOps engineer. If any of the IP addresses on the deny list access the application, the security team wants to receive automated notification in near real time so that the security team can document that the application needs additional security. The DevOps engineer creates a VPC flow log for the production VPC.

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

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

Answer: A

NEW QUESTION 9

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

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

Which solution will meet these requirements?

- A. Configure AWS Lambda to poll the SQS queue and invoke a Lambda function to check whether the queue messages are valid
- B. If validation fails, send a copy of the data that is not valid to an Amazon S3 bucket so that the scientists can review and correct the data
- C. When the data is corrected, amend the message in the SQS queue by using a replay Lambda function with the corrected data.
- D. Convert the SQS standard queue to an SQS FIFO queue
- E. Configure AWS Lambda to poll the SQS queue every 10 minutes by using an Amazon EventBridge schedule
- F. Invoke the Lambda function to identify any messages with a SentTimestamp value that is older than 5 minutes, push the data to the same location as the application's output location, and remove the messages from the queue.
- G. Create an SQS dead-letter queue
- H. Modify the existing queue by including a redrive policy that sets the Maximum Receives setting to 1 and sets the dead-letter queue ARN to the ARN of the newly created queue

- I. Instruct the scientists to use the dead-letter queue to review the data that is not valid
- J. Reprocess this data at a later time.
- K. Configure API Gateway to send messages to different SQS virtual queues that are named for each of the satellite
- L. Update the application to use a new virtual queue for any data that it cannot transform, and send the message to the new virtual queue
- M. Instruct the scientists to use the virtual queue to review the data that is not valid
- N. Reprocess this data at a later time.

Answer: C

Explanation:

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

NEW QUESTION 10

A company has an on-premises application that is written in Go. A DevOps engineer must move the application to AWS. The company's development team wants to enable blue/green deployments and perform A/B testing. Which solution will meet these requirements?

- A. Deploy the application on an Amazon EC2 instance, and create an AMI of the instance
- B. Use the AMI to create an automatic scaling launch configuration that is used in an Auto Scaling group
- C. Use Elastic Load Balancing to distribute traffic
- D. When changes are made to the application, a new AMI will be created, which will initiate an EC2 instance refresh.
- E. Use Amazon Lightsail to deploy the application
- F. Store the application in a zipped format in an Amazon S3 bucket
- G. Use this zipped version to deploy new versions of the application to Lightsail
- H. Use Lightsail deployment options to manage the deployment.
- I. Use AWS CodeArtifact to store the application code
- J. Use AWS CodeDeploy to deploy the application to a fleet of Amazon EC2 instances
- K. Use Elastic Load Balancing to distribute the traffic to the EC2 instance
- L. When making changes to the application, upload a new version to CodeArtifact and create a new CodeDeploy deployment.
- M. Use AWS Elastic Beanstalk to host the application
- N. Store a zipped version of the application in Amazon S3. Use that location to deploy new versions of the application
- O. Use Elastic Beanstalk to manage the deployment options.

Answer: D

Explanation:

<https://aws.amazon.com/quickstart/architecture/blue-green-deployment/>

NEW QUESTION 10

A company is implementing an Amazon Elastic Container Service (Amazon ECS) cluster to run its workload. The company architecture will run multiple ECS services on the cluster. The architecture includes an Application Load Balancer on the front end and uses multiple target groups to route traffic. A DevOps engineer must collect application and access logs. The DevOps engineer then needs to send the logs to an Amazon S3 bucket for near-real-time analysis.

Which combination of steps must the DevOps engineer take to meet these requirements? (Choose three.)

- A. Download the Amazon CloudWatch Logs container instance from AWS
- B. Configure this instance as a task
- C. Update the application service definitions to include the logging task.
- D. Install the Amazon CloudWatch Logs agent on the ECS instance
- E. Change the logging driver in the ECS task definition to awslogs.
- F. Use Amazon EventBridge to schedule an AWS Lambda function that will run every 60 seconds and will run the Amazon CloudWatch Logs create-export-task command
- G. Then point the output to the logging S3 bucket.
- H. Activate access logging on the ALB
- I. Then point the ALB directly to the logging S3 bucket.
- J. Activate access logging on the target groups that the ECS services use
- K. Then send the logs directly to the logging S3 bucket.
- L. Create an Amazon Kinesis Data Firehose delivery stream that has a destination of the logging S3 bucket
- M. Then create an Amazon CloudWatch Logs subscription filter for Kinesis Data Firehose.

Answer: BDF

Explanation:

<https://docs.aws.amazon.com/AmazonECS/latest/developerguide/ecs-logging-monitoring.html>

NEW QUESTION 13

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

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

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

- A. In the CloudFormation template, create an AWS::SSM::Document resource that joins the EC2 instance to the AWS Managed Microsoft AD domain by using the parameters for the existing director
- B. Update the launch template to include the SSMAssociation property to use the new SSM document
- C. Attach the AmazonSSMManagedInstanceCore and AmazonSSMDirectoryServiceAccess AWS managed policies to the IAM role that the EC2 instances use.
- D. In the CloudFormation template, update the launch template to include specific tags that propagate on launch

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

Answer: B

Explanation:

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

NEW QUESTION 14

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

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

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

Answer: A

Explanation:

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

NEW QUESTION 17

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

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

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

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

Answer: BD

Explanation:

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

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

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

NEW QUESTION 20

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

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

Which solution will meet these requirements?

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

Answer: B

Explanation:

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

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

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

NEW QUESTION 23

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

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

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

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

Answer: BD

Explanation:

The requirement is to update the infrastructure to ensure that only the Lambda function's execution role can access the values in Secrets Manager. The solution

must apply the principle of least privilege, which means granting the minimum permissions necessary to perform a task.

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

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

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

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

NEW QUESTION 26

An online retail company based in the United States plans to expand its operations to Europe and Asia in the next six months. Its product currently runs on Amazon EC2 instances behind an Application Load Balancer. The instances run in an Amazon EC2 Auto Scaling group across multiple Availability Zones. All data is stored in an Amazon Aurora database instance.

When the product is deployed in multiple regions, the company wants a single product catalog across all regions, but for compliance purposes, its customer information and purchases must be kept in each region.

How should the company meet these requirements with the LEAST amount of application changes?

- A. Use Amazon Redshift for the product catalog and Amazon DynamoDB tables for the customer information and purchases.
- B. Use Amazon DynamoDB global tables for the product catalog and regional tables for the customer information and purchases.
- C. Use Aurora with read replicas for the product catalog and additional local Aurora instances in each region for the customer information and purchases.
- D. Use Aurora for the product catalog and Amazon DynamoDB global tables for the customer information and purchases.

Answer: C

NEW QUESTION 27

AnyCompany is using AWS Organizations to create and manage multiple AWS accounts. AnyCompany recently acquired a smaller company, Example Corp.

During the acquisition process, Example Corp's single AWS account joined AnyCompany's management account through an Organizations invitation.

AnyCompany moved the new member account under an OU that is dedicated to Example Corp.

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

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

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

Answer: C

Explanation:

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

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

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

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

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

NEW QUESTION 30

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

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

- A. Use AWS OpsWorks Stacks to layer the server nodes of that cluster.
- B. Create a Chef recipe that populates the content of the `/etc./cluster/nodes` config file and restarts the service by using the current members of the layer.
- C. Assign that recipe to the Configure lifecycle event.

- D. Put the file nodes config in version contro
- 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 contro
- G. Deploy the new file and restart the services.
- H. Create an Amazon S3 bucket and upload a version of the /etc./cluster/nodes config file Create a crontab script that will poll for that S3 file and download it frequentl
- I. Use a process manager such as Monit or system, to restart the cluster services when it detects that the new file was modifie
- 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 confi
- L. Tile whenever a new instance is added to the cluster.

Answer: A

Explanation:

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

NEW QUESTION 34

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 38

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

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

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

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

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

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

Answer: AD

Explanation:

"AWS CodeArtifact operates in multiple Availability Zones and stores artifact data and metadata in Amazon S3 and Amazon DynamoDB. Your encrypted data is

redundantly stored across multiple facilities and multiple devices in each facility, making it highly available and highly durable."
https://aws.amazon.com/codeartifact/features/ With no internet connectivity, a gateway endpoint becomes necessary to access S3.

NEW QUESTION 42

A company uses an organization in AWS Organizations that has all features enabled. The company uses AWS Backup in a primary account and uses an AWS Key Management Service (AWS KMS) key to encrypt the backups.

The company needs to automate a cross-account backup of the resources that AWS Backup backs up in the primary account. The company configures cross-account backup in the Organizations management account. The company creates a new AWS account in the organization and configures an AWS Backup backup vault in the new account. The company creates a KMS key in the new account to encrypt the backups. Finally, the company configures a new backup plan in the primary account. The destination for the new backup plan is the backup vault in the new account.

When the AWS Backup job in the primary account is invoked, the job creates backups in the primary account. However, the backups are not copied to the new account's backup vault.

Which combination of steps must the company take so that backups can be copied to the new account's backup vault? (Select TWO.)

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

Answer: AE

Explanation:

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

References:

? 1: Creating backup copies across AWS accounts - AWS Backup

? 2: Using AWS Backup with AWS Organizations - AWS Backup

NEW QUESTION 43

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 48

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 53

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

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

Answer: C

Explanation:

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

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

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

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

References:

- ? 1: `AWS::EC2::FlowLog` - AWS CloudFormation
- ? 2: Amazon VPC Flow Logs extends CloudFormation Support to custom format subscriptions, 1-minute aggregation intervals and tagging
- ? 3: Logging IP traffic using VPC Flow Logs - Amazon Virtual Private Cloud
- ? : About AWS Config - AWS Config
- ? : `vpc-flow-logs-enabled` - AWS Config
- ? : Remediate Noncompliant Resources with AWS Config Rules - AWS Config

NEW QUESTION 55

A company has an application that is using a MySQL-compatible Amazon Aurora Multi-AZ DB cluster as the database. A cross-Region read replica has been created for disaster recovery purposes. A DevOps engineer wants to automate the promotion of the replica so it becomes the primary database instance in the event of a failure.

Which solution will accomplish this?

- A. Configure a latency-based Amazon Route 53 CNAME with health checks so it points to both the primary and replica endpoint
- B. Subscribe an Amazon SNS topic to Amazon RDS failure notifications from AWS CloudTrail and use that topic to invoke an AWS Lambda function that will promote the replica instance as the primary.
- C. Create an Aurora custom endpoint to point to the primary database instance
- D. Configure the application to use this endpoint
- E. Configure AWS CloudTrail to run an AWS Lambda function to promote the replica instance and modify the custom endpoint to point to the newly promoted instance.
- F. Create an AWS Lambda function to modify the application's AWS CloudFormation template to promote the replica, apply the template to update the stack, and point the application to the newly promoted instance
- G. Create an Amazon CloudWatch alarm to invoke this Lambda function after the failure event occurs.
- H. Store the Aurora endpoint in AWS Systems Manager Parameter Store
- I. Create an Amazon EventBridge event that detects the database failure and runs an AWS Lambda function to promote the replica instance and update the endpoint URL stored in AWS Systems Manager Parameter Store
- J. Code the application to reload the endpoint from Parameter Store if a database connection fails.

Answer: D

Explanation:

EventBridge is needed to detect the database failure. Lambda is needed to promote the replica as it's in another Region (manual promotion, otherwise). Storing and updating the endpoint in Parameter Store is important in updating the application. Look at High Availability section of Aurora FAQ:
<https://aws.amazon.com/rds/aurora/faqs/>

NEW QUESTION 57

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

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

Which solution will meet these requirements?

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

Answer: A

Explanation:

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

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

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

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

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

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

NEW QUESTION 58

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

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

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

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

Answer: C

Explanation:

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

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

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

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

NEW QUESTION 62

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

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

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

Answer: C

Explanation:

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

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

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

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

? [AWS Command Line Interface]

NEW QUESTION 65

A DevOps engineer needs to configure a blue green deployment for an existing three-tier application. The application runs on Amazon EC2 instances and uses an Amazon RDS database. The EC2 instances run behind an Application Load Balancer (ALB) and are in an Auto Scaling group.

The DevOps engineer has created a launch template and an Auto Scaling group for the blue environment. The DevOps engineer also has created a launch template and an Auto Scaling group for the green environment. Each Auto Scaling group deploys to a matching blue or green target group. The target group also specifies which software blue or green gets loaded on the EC2 instances. The ALB can be configured to send traffic to the blue environment's target group or the green environment's target group. An Amazon Route 53 record for www.example.com points to the ALB.

The deployment must move traffic all at once between the software on the blue environment's EC2 instances to the newly deployed software on the green environment's EC2 instances.

What should the DevOps engineer do to meet these requirements?

- A. Start a rolling restart of the Auto Scaling group for the green environment to deploy the new software on the green environment's EC2 instances. When the rolling restart is complete, use an AWS CLI command to update the ALB to send traffic to the green environment's target group.
- B. Use an AWS CLI command to update the ALB to send traffic to the green environment's target group.
- C. Then start a rolling restart of the Auto Scaling group for the green environment to deploy the new software on the green environment's EC2 instances.
- D. Update the launch template to deploy the green environment's software on the blue environment's EC2 instances. Keep the target groups and Auto Scaling groups unchanged in both environments. Perform a rolling restart of the blue environment's EC2 instances.
- E. Start a rolling restart of the Auto Scaling group for the green environment to deploy the new software on the green environment's EC2 instances. When the rolling restart is complete, update the Route 53 DNS to point to the green environment's endpoint on the ALB.

Answer: A

Explanation:

This solution will meet the requirements because it will use a rolling restart to gradually replace the EC2 instances in the green environment with new instances that have the new software version installed. A rolling restart is a process that terminates and launches instances in batches, ensuring that there is always a minimum number of healthy instances in service. This way, the green environment can be updated without affecting the availability or performance of the application. When the rolling restart is complete, the DevOps engineer can use an AWS CLI command to modify the listener rules of the ALB and change the default action to forward traffic to the green environment's target group. This will switch the traffic from the blue environment to the green environment all at once, as required by the question.

NEW QUESTION 67

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 69

A company has 20 service teams. Each service team is responsible for its own microservice. Each service team uses a separate AWS account for its microservice and a VPC with the 192.168.0.0/22 CIDR block. The company manages the AWS accounts with AWS Organizations.

Each service team hosts its microservice on multiple Amazon EC2 instances behind an Application Load Balancer. The microservices communicate with each other across the public internet. The company's security team has issued a new guideline that all communication between microservices must use HTTPS over private network connections and cannot traverse the public internet.

A DevOps engineer must implement a solution that fulfills these obligations and minimizes the number of changes for each service team. Which solution will meet these requirements?

- A. Create a new AWS account in AWS Organizations Create a VPC in this account and use AWS Resource Access Manager to share the private subnets of this VPC with the organization Instruct the service teams to launch a ne
- B. Network Load Balancer (NLB) and EC2 instances that use the shared private subnets Use the NLB DNS names for communication between microservices.
- C. Create a Network Load Balancer (NLB) in each of the microservice VPCs Use AWS PrivateLink to create VPC endpoints in each AWS account for the NLBs Create subscriptions to each VPC endpoint in each of the other AWS accounts Use the VPC endpoint DNS names for communication between microservices.
- D. Create a Network Load Balancer (NLB) in each of the microservice VPCs Create VPC peering connections between each of the microservice VPCs Update the route tables for each VPC to use the peering links Use the NLB DNS names for communication between microservices.
- E. Create a new AWS account in AWS Organizations Create a transit gateway in this account and use AWS Resource Access Manager to share the transit gateway with the organizatio
- F. In each of the microservice VPC
- G. create a transit gateway attachment to the shared transit gateway Update the route tables of each VPC to use the transit gateway Create a Network Load Balancer (NLB) in each of the microservice VPCs Use the NLB DNS names for communication between microservices.

Answer: B

Explanation:

<https://aws.amazon.com/blogs/networking-and-content-delivery/connecting-networks-with-overlapping-ip-ranges/> Private link is the best option because Transit Gateway doesn't support overlapping CIDR ranges.

NEW QUESTION 70

A company has many applications. Different teams in the company developed the applications by using multiple languages and frameworks. The applications run on premises and on different servers with different operating systems. Each team has its own release protocol and process. The company wants to reduce the complexity of the release and maintenance of these applications.

The company is migrating its technology stacks, including these applications, to AWS. The company wants centralized control of source code, a consistent and automatic delivery pipeline, and as few maintenance tasks as possible on the underlying infrastructure.

What should a DevOps engineer do to meet these requirements?

- A. Create one AWS CodeCommit repository for all application
- B. Put each application's code in a different branc
- C. Merge the branches, and use AWS CodeBuild to build the application
- D. Use AWS CodeDeploy to deploy the applications to one centralized application server.
- E. Create one AWS CodeCommit repository for each of the application
- F. Use AWS CodeBuild to build the applications one at a tim
- G. Use AWS CodeDeploy to deploy the applications to one centralized application server.
- H. Create one AWS CodeCommit repository for each of the application
- I. Use AWS CodeBuild to build the applications one at a time and to create one AMI for each serve
- J. Use AWS CloudFormation StackSets to automatically provision and decommission Amazon EC2 fleets by using these AMIs.
- K. Create one AWS CodeCommit repository for each of the application
- L. Use AWS CodeBuild to build one Docker image for each application in Amazon Elastic Container Registry (Amazon ECR). Use AWS CodeDeploy to deploy the applications to Amazon Elastic Container Service (Amazon ECS) on infrastructure that AWS Fargate manages.

Answer: D

Explanation:

because of "as few maintenance tasks as possible on the underlying infrastructure". Fargate does that better than "one centralized application server"

NEW QUESTION 71

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 72

A company requires its internal business teams to launch resources through pre-approved AWS CloudFormation templates only. The security team requires automated monitoring when resources drift from their expected state.

Which strategy should be used to meet these requirements?

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

Answer: C

Explanation:

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

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

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

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

NEW QUESTION 73

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/singlelogin/latest/userguide/configure-abac.html>

NEW QUESTION 77

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 80

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 81

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 85

A DevOps engineer is deploying a new version of a company's application in an AWS CodeDeploy deployment group associated with its Amazon EC2 instances. After some time, the deployment fails. The engineer realizes that all the events associated with the specific deployment ID are in a Skipped status and code was not deployed in the instances associated with the deployment group. What are valid reasons for this failure? (Select TWO.)

- A. The networking configuration does not allow the EC2 instances to reach the internet via a NAT gateway or internet gateway and the CodeDeploy endpoint cannot be reached.
- B. The IAM user who triggered the application deployment does not have permission to interact with the CodeDeploy endpoint.
- C. The target EC2 instances were not properly registered with the CodeDeploy endpoint.
- D. An instance profile with proper permissions was not attached to the target EC2 instances.
- E. The appspec file was not included in the application revision.
- F. The appspec file was not included in the application revision.

Answer: AD

Explanation:

<https://docs.aws.amazon.com/codedeploy/latest/userguide/troubleshooting-deployments.html#troubleshooting-skipped-lifecycle-events>

NEW QUESTION 87

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

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

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

Answer: C

Explanation:

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

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

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

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

NEW QUESTION 91

A company manages multiple AWS accounts by using AWS Organizations with OUS for the different business divisions. The company is updating their corporate network to use new IP address ranges. The company has 10 Amazon S3 buckets in different AWS accounts. The S3 buckets store reports for the different divisions. The S3 bucket configurations allow only private corporate network IP addresses to access the S3 buckets.

A DevOps engineer needs to change the range of IP addresses that have permission to access the contents of the S3 buckets. The DevOps engineer also needs to revoke the permissions of two OUS in the company.

Which solution will meet these requirements?

- A. Create a new SCP that has two statements, one that allows access to the new range of IP addresses for all the S3 buckets and one that denies access to the old range of IP addresses for all the S3 bucket
- B. Set a permissions boundary for the OrganizationAccountAccessRole role in the two OUS to deny access to the S3 buckets.
- C. Create a new SCP that has a statement that allows only the new range of IP addresses to access the S3 bucket
- D. Create another SCP that denies access to the S3 bucket
- E. Attach the second SCP to the two OUS
- F. On all the S3 buckets, configure resource-based policies that allow only the new range of IP addresses to access the S3 bucket
- G. Create a new SCP that denies access to the S3 bucket
- H. Attach the SCP to the two OUs.
- I. On all the S3 buckets, configure resource-based policies that allow only the new range of IP addresses to access the S3 bucket
- J. Set a permissions boundary for the OrganizationAccountAccessRole role in the two OUS to deny access to the S3 buckets.

Answer: C

Explanation:

The correct answer is C.

A comprehensive and detailed explanation is:

? Option A is incorrect because creating a new SCP that has two statements, one that allows access to the new range of IP addresses for all the S3 buckets and one that denies access to the old range of IP addresses for all the S3 buckets, is not a valid solution. SCPs are not resource-based policies, and they cannot specify the S3 buckets or the IP addresses as resources or conditions. SCPs can only control the actions that can be performed by the principals in the organization, not the access to specific resources. Moreover, setting a permissions boundary for the OrganizationAccountAccessRole role in the two OUs to deny access to the S3 buckets is not sufficient to revoke the permissions of the two OUs, as there might be other roles or users in those OUs that can still access the S3 buckets.

? Option B is incorrect because creating a new SCP that has a statement that allows only the new range of IP addresses to access the S3 buckets is not a valid solution, for the same reason as option A. SCPs are not resource-based policies, and they cannot specify the S3 buckets or the IP addresses as resources or conditions. Creating another SCP that denies access to the S3 buckets and attaching it to the two OUs is also not a valid solution, as SCPs cannot specify the S3 buckets as resources either.

? Option C is correct because it meets both requirements of changing the range of IP addresses that have permission to access the contents of the S3 buckets and revoking the permissions of two OUs in the company. On all the S3 buckets, configuring resource-based policies that allow only the new range of IP addresses to access the S3 buckets is a valid way to update the IP address ranges, as resource-based policies can specify both resources and conditions. Creating a new SCP that denies access to the S3 buckets and attaching it to the two OUs is also a valid way to revoke the permissions of those OUs, as SCPs can deny actions such as s3:PutObject or s3:GetObject on any resource.

? Option D is incorrect because setting a permissions boundary for the OrganizationAccountAccessRole role in the two OUs to deny access to the S3 buckets is not sufficient to revoke the permissions of the two OUs, as there might be other roles or users in those OUs that can still access the S3 buckets. A permissions boundary is a policy that defines the maximum permissions that an IAM entity can have. However, it does not revoke any existing permissions that are granted by other policies.

References:

- ? AWS Organizations
- ? S3 Bucket Policies
- ? Service Control Policies
- ? Permissions Boundaries

NEW QUESTION 94

A company has an application and a CI/CD pipeline. The CI/CD pipeline consists of an AWS CodePipeline pipeline and an AWS CodeBuild project. The CodeBuild project runs tests against the application as part of the build process and outputs a test report. The company must keep the test reports for 90 days.

Which solution will meet these requirements?

- A. Add a new stage in the CodePipeline pipeline after the stage that contains the CodeBuild project
- B. Create an Amazon S3 bucket to store the report
- C. Configure an S3 deploy action type in the new CodePipeline stage with the appropriate path and format for the reports.
- D. Add a report group in the CodeBuild project buildspec file with the appropriate path and format for the report
- E. Create an Amazon S3 bucket to store the report
- F. Configure an Amazon EventBridge rule that invokes an AWS Lambda function to copy the reports to the S3 bucket when a build is complete
- G. Create an S3 Lifecycle rule to expire the objects after 90 days.
- H. Add a new stage in the CodePipeline pipeline
- I. Configure a test action type with the appropriate path and format for the report
- J. Configure the report expiration time to be 90 days in the CodeBuild project buildspec file.
- K. Add a report group in the CodeBuild project buildspec file with the appropriate path and format for the report
- L. Create an Amazon S3 bucket to store the report
- M. Configure the report group as an artifact in the CodeBuild project buildspec file
- N. Configure the S3 bucket as the artifact destination
- O. Set the object expiration to 90 days.

Answer: B

Explanation:

The correct solution is to add a report group in the AWS CodeBuild project buildspec file with the appropriate path and format for the reports. Then, create an

Amazon S3 bucket to store the reports. You should configure an Amazon EventBridge rule that invokes an AWS Lambda function to copy the reports to the S3 bucket when a build is completed. Finally, create an S3 Lifecycle rule to expire the objects after 90 days. This approach allows for the automated transfer of reports to long-term storage and ensures they are retained for the required duration without manual intervention¹. References:

- ? AWS CodeBuild User Guide on test reporting¹.
- ? AWS CodeBuild User Guide on working with report groups².
- ? AWS Documentation on using AWS CodePipeline with AWS CodeBuild³.

NEW QUESTION 99

A company has a new AWS account that teams will use to deploy various applications. The teams will create many Amazon S3 buckets for application-specific purposes and to store AWS CloudTrail logs. The company has enabled Amazon Macie for the account. A DevOps engineer needs to optimize the Macie costs for the account without compromising the account's functionality. Which solutions will meet these requirements? (Select TWO.)

- A. Exclude S3 buckets that contain CloudTrail logs from automated discovery.
- B. Exclude S3 buckets that have public read access from automated discovery.
- C. Configure scheduled daily discovery jobs for all S3 buckets in the account.
- D. Configure discovery jobs to include S3 objects based on the last modified criterion.
- E. Configure discovery jobs to include S3 objects that are tagged as production only.

Answer: AD

Explanation:

To optimize the Macie costs for the account without compromising the account's functionality, the DevOps engineer needs to exclude S3 buckets that do not contain sensitive data from automated discovery. S3 buckets that contain CloudTrail logs are unlikely to have sensitive data, and Macie charges for scanning and monitoring data in S3 buckets. Therefore, excluding S3 buckets that contain CloudTrail logs from automated discovery can reduce Macie costs. Similarly, configuring discovery jobs to include S3 objects based on the last modified criterion can also reduce Macie costs, as it will only scan and monitor new or updated objects, rather than all objects in the bucket.

NEW QUESTION 104

A DevOps engineer is designing an application that integrates with a legacy REST API. The application has an AWS Lambda function that reads records from an Amazon Kinesis data stream. The Lambda function sends the records to the legacy REST API.

Approximately 10% of the records that the Lambda function sends from the Kinesis data stream have data errors and must be processed manually. The Lambda function event source configuration has an Amazon Simple Queue Service (Amazon SQS) dead-letter queue as an on-failure destination. The DevOps engineer has configured the Lambda function to process records in batches and has implemented retries in case of failure.

During testing the DevOps engineer notices that the dead-letter queue contains many records that have no data errors and that already have been processed by the legacy REST API. The DevOps engineer needs to configure the Lambda function's event source options to reduce the number of errorless records that are sent to the dead-letter queue.

Which solution will meet these requirements?

- A. Increase the retry attempts
- B. Configure the setting to split the batch when an error occurs
- C. Increase the concurrent batches per shard
- D. Decrease the maximum age of record

Answer: B

Explanation:

This solution will meet the requirements because it will reduce the number of errorless records that are sent to the dead-letter queue. When you configure the setting to split the batch when an error occurs, Lambda will retry only the records that caused the error, instead of retrying the entire batch. This way, the records that have no data errors and have already been processed by the legacy REST API will not be retried and sent to the dead-letter queue unnecessarily.

<https://docs.aws.amazon.com/lambda/latest/dg/with-kinesis.html>

NEW QUESTION 106

A company requires an RPO of 2 hours and an RTO of 10 minutes for its data and application at all times. An application uses a MySQL database and Amazon EC2 web servers. The development team needs a strategy for failover and disaster recovery.

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

- A. Create an Amazon Aurora cluster in one Availability Zone across multiple Regions as the data store. Use Aurora's automatic recovery capabilities in the event of a disaster.
- B. Create an Amazon Aurora global database in two Regions as the data store.
- C. In the event of a failure, promote the secondary Region as the primary for the application.
- D. Create an Amazon Aurora multi-master cluster across multiple Regions as the data store.
- E. Use a Network Load Balancer to balance the database traffic in different Regions.
- F. Set up the application in two Regions and use Amazon Route 53 failover-based routing that points to the Application Load Balancers in both Regions.
- G. Use health checks to determine the availability in a given Region.
- H. Use Auto Scaling groups in each Region to adjust capacity based on demand.
- I. Set up the application in two Regions and use a multi-Region Auto Scaling group behind Application Load Balancers to manage the capacity based on demand.
- J. In the event of a disaster, adjust the Auto Scaling group's desired instance count to increase baseline capacity in the failover Region.

Answer: BD

NEW QUESTION 108

A DevOps engineer is implementing governance controls for a company that requires its infrastructure to be housed within the United States. The engineer must restrict which AWS Regions can be used, and ensure an alert is sent as soon as possible if any activity outside the governance policy takes place. The controls should be automatically enabled on any new Region outside the United States (US).

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

- A. Create an AWS Organizations SCP that denies access to all non-global services in non-US Region

- B. Attach the policy to the root of the organization.
- C. Configure AWS CloudTrail to send logs to Amazon CloudWatch Logs and enable it for all Region
- D. Use a CloudWatch Logs metric filter to send an alert on any service activity in non-US Regions.
- E. Use an AWS Lambda function that checks for AWS service activity and deploy it to all Region
- F. Write an Amazon EventBridge rule that runs the Lambda function every hour, sending an alert if activity is found in a non-US Region.
- G. Use an AWS Lambda function to query Amazon Inspector to look for service activity in non-US Regions and send alerts if any activity is found.
- H. Write an SCP using the aws: RequestedRegion condition key limiting access to US Region
- I. Apply the policy to all users, groups, and roles

Answer: AB

Explanation:

To implement governance controls that restrict AWS service usage to within the United States and ensure alerts for any activity outside the governance policy, the following actions will meet the requirements:

? A. Create an AWS Organizations SCP that denies access to all non-global services in non-US Regions. Attach the policy to the root of the organization. This action will effectively prevent users and roles in all accounts within the organization from accessing services in non-US Regions¹².

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

References:

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

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

NEW QUESTION 112

A company uses AWS CloudFormation stacks to deploy updates to its application. The stacks consist of different resources. The resources include AWS Auto Scaling groups, Amazon EC2 instances, Application Load Balancers (ALBs), and other resources that are necessary to launch and maintain independent stacks. Changes to application resources outside of CloudFormation stack updates are not allowed.

The company recently attempted to update the application stack by using the AWS CLI. The stack failed to update and produced the following error message:

"ERROR: both the deployment and the CloudFormation stack rollback failed. The deployment failed because the following resource(s) failed to update: [AutoScalingGroup]."

The stack remains in a status of UPDATE_ROLLBACK_FAILED. * Which solution will resolve this issue?

- A. Update the subnet mappings that are configured for the ALB
- B. Run the aws cloudformation update-stack-set AWS CLI command.
- C. Update the IAM role by providing the necessary permissions to update the stack
- D. Run the aws cloudformation continue-update-rollback AWS CLI command.
- E. Submit a request for a quota increase for the number of EC2 instances for the account
- F. Run the aws cloudformation cancel-update-stack AWS CLI command.
- G. Delete the Auto Scaling group resource
- H. Run the aws cloudformation rollback-stack AWS CLI command.

Answer: B

Explanation:

<https://repost.aws/knowledge-center/cloudformation-update-rollback-failed> If your stack is stuck in the UPDATE_ROLLBACK_FAILED state after a failed update, then the only actions that you can perform on the stack are the ContinueUpdateRollback or DeleteStack operations.

NEW QUESTION 114

A company is using AWS CodePipeline to deploy an application. According to a new guideline, a member of the company's security team must sign off on any application changes before the changes are deployed into production. The approval must be recorded and retained.

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

- A. Configure CodePipeline to write actions to Amazon CloudWatch Logs.
- B. Configure CodePipeline to write actions to an Amazon S3 bucket at the end of each pipeline stage.
- C. Create an AWS CloudTrail trail to deliver logs to Amazon S3.
- D. Create a CodePipeline custom action to invoke an AWS Lambda function for approval
- E. Create a policy that gives the security team access to manage CodePipeline custom actions.
- F. Create a CodePipeline manual approval action before the deployment stage
- G. Create a policy that grants the security team access to approve manual approval stages.

Answer: CE

Explanation:

To meet the new guideline for application deployment, the company can use a combination of AWS CodePipeline and AWS CloudTrail. A manual approval action in CodePipeline allows the security team to review and approve changes before they are deployed. This action can be configured to pause the pipeline until approval is granted, ensuring that no changes move to production without the necessary sign-off. Additionally, by creating an AWS CloudTrail trail, all actions taken within CodePipeline, including approvals, are recorded and delivered to an Amazon S3 bucket. This provides an audit trail that can be retained for compliance and review purposes.

References:

? AWS CodePipeline's manual approval action provides a way to ensure that a member of the security team can review and approve changes before they are deployed¹.

? AWS CloudTrail integration with CodePipeline allows for the recording and retention of all pipeline actions, including approvals, which can be stored in Amazon S3 for record-keeping².

NEW QUESTION 116

A DevOps engineer is researching the least expensive way to implement an image batch processing cluster on AWS. The application cannot run in Docker containers and must run on Amazon EC2. The batch job stores checkpoint data on an NFS volume and can tolerate interruptions. Configuring the cluster software from a generic EC2 Linux image takes 30 minutes.

What is the MOST cost-effective solution?

- A. Use Amazon EFS (or checkpoint data)
- B. To complete the job, use an EC2 Auto Scaling group and an On-Demand pricing model to provision EC2 instances temporarily.
- C. Use GlusterFS on EC2 instances for checkpoint data
- D. To run the batch job configure EC2 instances manually When the job completes shut down the instances manually.
- E. Use Amazon EFS for checkpoint data Use EC2 Fleet to launch EC2 Spot Instances and utilize user data to configure the EC2 Linux instance on startup.
- F. Use Amazon EFS for checkpoint data Use EC2 Fleet to launch EC2 Spot Instances Create a custom AMI for the cluster and use the latest AMI when creating instances.

Answer: D

NEW QUESTION 118

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

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

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

Answer: C

Explanation:

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

NEW QUESTION 123

.....

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