[2025-November-NewBraindump2go DOP-C02 Dumps Free[Q340-Q370

2025/November Latest Braindump2go DOP-C02 Exam Dumps with PDF and VCE Free Updated Today! Following are some new Braindump2go DOP-C02 Real Exam Questions!QUESTION 340A company uses Amazon Redshift as its data warehouse solution. The company wants to create a dashboard to view changes to the Redshift users and the queries the users perform.Which combination of steps will meet this requirement? (Choose two.)A. Create an Amazon CloudWatch log group. Create an AWS CloudTrail trail that writes to the CloudWatch log group.B. Create a new Amazon S3 bucket. Configure default audit logging on the Redshift cluster. Configure the S3 bucket as the target.C. Configure the Redshift cluster database audit logging to include user activity logs. Configure Amazon CloudWatch as the target.D. Create an Amazon CloudWatch dashboard that has a log widget. Configure the widget to display user details from the Redshift logs.E. Create an AWS Lambda function that uses Amazon Athena to query the Redshift logs. Create an Amazon CloudWatch dashboard that has a custom widget type that uses the Lambda function.Answer: BCExplanation:Amazon Redshift audit logging allows you to capture information about the activities performed on the database, including changes to users and the queries executed. By enabling default audit logging and specifying an S3 bucket as the target, you can store the logs in a centralized location. This step ensures that user activity and database changes are captured. Redshift's database audit logging can include user activity logs, which track the SQL queries performed by users and the changes they make. By configuring these logs and sending them to Amazon CloudWatch, you can monitor user activity in real time, making it easier to integrate with a monitoring and alerting dashboard.By enabling audit logging for Amazon Redshift and sending the logs to S3 and CloudWatch, you can track changes to Redshift users and queries effectively and integrate the data into a dashboard for monitoring purposes.QUESTION 341A company uses an organization in AWS Organizations to manage its 500 AWS accounts. The organization has all features enabled. The AWS accounts are in a single OU. The developers need to use the CostCenter tag key for all resources in the organization's member accounts. Some teams do not use the CostCenter tag key to tag their Amazon EC2 instances.The cloud team wrote a script that scans all EC2 instances in the organization's member accounts. If the EC2 instances do not have a CostCenter tag key, the script will notify AWS account administrators. To avoid this notification, some developers use the CostCenter tag key with an arbitrary string in the tag value.The cloud team needs to ensure that all EC2 instances in the organization use a CostCenter tag key with the appropriate cost center value.Which solution will meet these requirements?A. Create an SCP that prevents the creation of EC2 instances without the CostCenter tag key. Create a tag policy that requires the CostCenter tag to be values from a known list of cost centers for all EC2 instances. Attach the policy to the OU. Update the script to scan the tag keys and tag values.Modify the script to update noncompliant resources with a default approved tag value for the CostCenter tag key.B. Create an SCP that prevents the creation of EC2 instances without the CostCenter tag key. Attach the policy to the OU. Update the script to scan the tag keys and tag values and notify the administrators when the tag values are not valid.C. Create an SCP that prevents the creation of EC2 instances without the CostCenter tag key. Attach the policy to the OU. Create an IAM permission boundary in the organization's member accounts that restricts the CostCenter tag values to a list of valid cost centers.D. Create a tag policy that requires the CostCenter tag to be values from a known list of cost centers for all EC2 instances. Attach the policy to the OU.Configure an AWS Lambda function that adds an empty CostCenter tag key to an EC2 instance. Create an Amazon EventBridge rule that matches events to the RunInstances API action with the Lambda function as the target.Answer: A Explanation:Service Control Policies (SCPs) can be used to prevent the creation of EC2 instances without a required tag, such as the CostCenter tag. This ensures that any new EC2 instance created in the organization must include the CostCenter tag key.Tag policies allow you to define rules for specific tags, such as specifying a known list of valid values for the CostCenter tag. By attaching a tag policy to the OU, you enforce that the CostCenter tag must contain a value from the approved list. This ensures that developers cannot use arbitrary values for the CostCenter tag.The existing script can be enhanced to not only scan for missing CostCenter tags but also check for invalid tag values. It can automatically remediate noncompliant resources by updating them with a default approved tag value or notifying administrators if necessary. This makes the solution proactive in fixing noncompliant resources while reducing manual intervention.By combining an SCP to enforce the presence of the CostCenter tag with a tag policy that validates the tag value against a list of approved cost centers, option A ensures both the presence and correctness of the CostCenter tag across all EC2 instances. This solution provides proactive governance while allowing the existing script to handle exceptions and enforce compliance.QUESTION 342A DevOps engineer uses a pipeline in AWS CodePipeline. The pipeline has a build action and a deploy action for a single-page web application that is delivered to an Amazon S3 bucket. Amazon CloudFront serves the web application. The build action creates an artifact for the web application.The DevOps engineer has created an AWS CloudFormation template that defines the S3 bucket and configures the S3 bucket to host the application. The DevOps engineer has configured a CloudFormation deploy action before the S3 action. The CloudFormation deploy action creates the S3 bucket. The DevOps engineer

needs to configure the S3 deploy action to use the S3 bucket from the CloudFormation template.Which combination of steps will meet these requirements? (Choose two.)A. Add an output named BucketName to the CloudFormation template. Set the output's value to refer to the S3 bucket from the CloudFormation template. Configure the output value to export to an AWS::SSM::Parameter resource named Stackvariables.B. Add an output named BucketName to the CloudFormation template. Set the output's value to refer to the S3 bucket from the CloudFormation template. Set the CloudFormation action's namespace to StackVariables in the pipeline.C. Configure the output artifacts of the CloudFormation action in the pipeline to be an AWS Systems Manager Parameter Store parameter named StackVariables. Name the artifact BucketName.D. Configure the build artifact from the build action as the input to the CodePipeline S3 deploy action. Configure the deploy action to deploy to the S3 bucket by using the StackVariables.BucketName variable.E. Configure the build artifact from the build action and the AWS Systems Manager parameter as the inputs to the deploy action. Configure the deploy action to deploy to the S3 bucket by using the StackVariables.BucketName variable.Answer: BDExplanation:You need to ensure that the S3 bucket name is exposed from the CloudFormation stack as an output. This allows the pipeline to reference the bucket dynamically after the CloudFormation deploy action has created it.By adding an output named BucketName in the CloudFormation template, you can export the S3 bucket's name so that it can be used in subsequent pipeline actions.Setting the CloudFormation action's namespace to StackVariables in the pipeline makes the output available under that namespace, allowing you to reference it later.The build artifact from the build action contains the web application code that needs to be deployed. You can configure the S3 deploy action to use this build artifact as input.The deploy action will use the S3 bucket created by the CloudFormation stack by referencing the BucketName from the StackVariables namespace.This ensures that the application is deployed to the correct bucket without hardcoding the bucket name.QUESTION 343 A company used a lift and shift strategy to migrate a workload to AWS. The company has an Auto Scaling group of Amazon EC2 instances. Each EC2 instance runs a web application, a database, and a Redis cache.Users are experiencing large variations in the web application's response times. Requests to the web application go to a single EC2 instance that is under significant load. The company wants to separate the application components to improve availability and performance.Which solution will meet these requirements?A. Create a Network Load Balancer and an Auto Scaling group for the web application. Migrate the database to an Amazon Aurora Serverless database. Create an Application Load Balancer and an Auto Scaling group for the Redis cache.B. Create an Application Load Balancer and an Auto Scaling group for the web application. Migrate the database to an Amazon Aurora database that has a Multi-AZ deployment. Create a Network Load Balancer and an Auto Scaling group in a single Availability Zone for the Redis cache.C. Create a Network Load Balancer and an Auto Scaling group for the web application. Migrate the database to an Amazon Aurora Serverless database. Create an Amazon ElastiCache (Redis OSS) cluster for the cache. Create a target group that has a DNS target type that contains the ElastiCache (Redis OSS) cluster hostname.D. Create an Application Load Balancer and an Auto Scaling group for the web application. Migrate the database to an Amazon Aurora database that has a Multi-AZ deployment. Create an Amazon ElastiCache (Redis OSS) cluster for the cache.Answer: DExplanation:Application Load Balancer is designed to handle HTTP/HTTPS traffic, making it suitable for routing requests to the web application. By creating an Auto Scaling group, you can distribute the load across multiple instances, improving the performance and availability of the web application, and preventing any single instance from being overwhelmed.Amazon Aurora provides high availability and performance improvements over a traditional database hosted on an EC2 instance. The Multi-AZ deployment ensures that the database remains available even if one Availability Zone experiences an outage, which enhances availability and fault tolerance.Amazon ElastiCache is a fully managed Redis service that provides high availability and performance improvements for caching. This separates the Redis cache from the EC2 instance, improving overall application performance by offloading cache management to a managed service, which is optimized for this use case.By using an Application Load Balancer, Multi-AZ Aurora, and Amazon ElastiCache, the solution improves the separation of the web application, database, and cache components, enhancing both performance and availability while leveraging AWS managed services.QUESTION 344A company is using AWS Organizations and wants to implement a governance strategy with the following requirements:- AWS resource access is restricted to the same two Regions for all accounts.- AWS services are limited to a specific group of authorized services for all accounts.- Authentication is provided by Active Directory.- Access permissions are organized by job function and are identical in each account.Which solution will meet these requirements?A. Establish an organizational unit (OU) with group policies in the management account to restrict Regions and authorized services. Use AWS CloudFormation StackSets to provision roles with permissions for each job function, including an IAM trust policy for IAM identity provider authentication in each account.B. Establish a permission boundary in the management account to restrict Regions and authorized services. Use AWS CloudFormation StackSets to provision roles with permissions for each job function, including an IAM trust policy for IAM identity provider authentication in each account.C. Establish a service control policy in the management account to restrict Regions and authorized services. Use AWS Resource Access Manager (AWS RAM) to share

management account roles with permissions for each job function, including AWS IAM Identity Center for authentication in each account.D. Establish a service control policy in the management account to restrict Regions and authorized services. Use AWS CloudFormation StackSets to provision roles with permissions for each job function, including an IAM trust policy for IAM identity provider authentication in each account.Answer: DExplanation:Service Control Policies (SCPs) are the appropriate tool for enforcing governance in AWS Organizations. SCPs allow you to restrict AWS resources, services, and regions across all accounts in an organization. By setting an SCP, you can ensure that access to AWS resources is limited to the specific two regions and that only the authorized AWS services can be used. SCPs are applied at the organizational level and govern what IAM users and roles within an account can do, providing a strong governance mechanism.CloudFormation StackSets are ideal for deploying IAM roles and permissions consistently across multiple accounts in AWS Organizations. You can define roles that are organized by job function and ensure that these roles are provisioned identically in each account. This simplifies management by creating consistent access policies across all accounts.The solution mentions using an IAM trust policy for IAM identity provider authentication, which allows you to integrate Active Directory with AWS for authentication. By configuring a trust relationship, the accounts can authenticate with Active Directory Federation Services (ADFS) or another compatible identity provider (IdP), ensuring centralized authentication for users across accounts.By using SCPs for governance and CloudFormation StackSets for consistent role provisioning, this solution provides an efficient and scalable approach to meet the company's governance, regional restrictions, service limitations, and centralized authentication requirements.QUESTION 345A company detects unusual login attempts in many of its AWS accounts. A DevOps engineer must implement a solution that sends a notification to the company's security team when multiple failed login attempts occur. The DevOps engineer has already created an Amazon Simple Notification Service (Amazon SNS) topic and has subscribed the security team to the SNS topic.Which solution will provide the notification with the LEAST operational effort?A. Configure AWS CloudTrail to send management events to an Amazon CloudWatch Logs log group. Create a CloudWatch Logs metric filter to match failed ConsoleLogin events. Create a CloudWatch alarm that is based on the metric filter. Configure an alarm action to send messages to the SNS topic.B. Configure AWS CloudTrail to send management events to an Amazon S3 bucket. Create an Amazon Athena query that returns a failure if the query finds failed logins in the logs in the S3 bucket. Create an Amazon EventBridge rule to periodically run the query. Create a second EventBridge rule to detect when the query fails and to send a message to the SNS topic.C. Configure AWS CloudTrail to send data events to an Amazon CloudWatch Logs log group. Create a CloudWatch logs metric filter to match failed ConsoleLogin events. Create a CloudWatch alarm that is based on the metric filter. Configure an alarm action to send messages to the SNS topic.D. Configure AWS CloudTrail to send data events to an Amazon S3 bucket. Configure an Amazon S3 event notification for the s3:ObjectCreated event type. Filter the event type by ConsoleLogin failed events. Configure the event notification to forward to the SNS topic.Answer: AExplanation:AWS CloudTrail management events include login attempts like ConsoleLogin. By sending these events to Amazon CloudWatch Logs, you can track and analyze these logs in real-time, which is ideal for detecting failed login attempts.A CloudWatch Logs metric filter can be configured to detect specific patterns (such as failed ConsoleLogin events) in the logs. This provides an efficient and automated way to detect multiple failed login attempts.Once the metric filter is in place, you can create a CloudWatch alarm that triggers when the number of failed login attempts exceeds a predefined threshold. This alarm can automatically notify the security team by sending a message to the SNS topic.This solution is efficient because it:- Uses existing services like CloudTrail and CloudWatch, which are designed for monitoring and logging.- Provides real-time detection and notification with minimal configuration.- Requires no custom code or manual intervention after setup.It provides the best balance of simplicity, real-time monitoring, and minimal operational effort for detecting and notifying about failed login attempts.QUESTION 346A company has deployed a new REST API by using Amazon API Gateway. The company uses the API to access confidential data. The API must be accessed from only specific VPCs in the company.Which solution will meet these requirements?A. Create and attach a resource policy to the API Gateway API. Configure the resource policy to allow only the specific VPC IDs.B. Add a security group to the API Gateway API. Configure the inbound rules to allow only the specific VPC IP address ranges.C. Create and attach an IAM role to the API Gateway API. Configure the IAM role to allow only the specific VPC IDs.D. Add an ACL to the API Gateway API. Configure the outbound rules to allow only the specific VPC IP address ranges.Answer: AExplanation:API Gateway resource policies are used to control access to your API based on various conditions, such as IP addresses, VPC IDs, or VPC endpoints.In this case, you can create a resource policy that explicitly allows access to the API only from requests originating from the specified VPC IDs. This approach is the most direct and secure way to restrict access to your API.A resource policy provides fine-grained control over which VPCs can access the API, ensuring that only the specified VPCs within the company can communicate with the API. This is a highly effective way to control access to sensitive resources.Since the API is used to access confidential data, restricting access at the API Gateway level with resource policies ensures that unauthorized requests from external VPCs or IP addresses are blocked, reducing the attack surface.

Using resource policies is the correct approach to restrict access to your API to specific VPCs, providing the necessary security for accessing confidential data.QUESTION 347A company runs a website by using an Amazon Elastic Container Service (Amazon ECS) service that is connected to an Application Load Balancer (ALB). The service was in a steady state with tasks responding to requests successfully.A DevOps engineer updated the task definition with a new container image and deployed the new task definition to the service. The DevOps engineer noticed that the service is frequently stopping and starting new tasks because the ALB healtth checks are failing.What should the DevOps engineer do to troubleshoot the failed deployment?A. Ensure that a security group associated with the service allows traffic from the ALB.B. Increase the ALB health check grace period for the service.C. Increase the service minimum healthy percent setting.D. Decrease the ALB health check interval.Answer: A Explanation:When the ALB health checks are failing, one common cause is that the security group associated with the ECS tasks (or service) is not allowing traffic from the ALB. The health checks from the ALB need to reach the ECS tasks to confirm that they are healthy and responsive.If the security group associated with the ECS tasks is too restrictive and does not permit traffic from the ALB, the health checks will fail, causing ECS to stop and restart tasks. Ensuring that the ECS service's security group allows inbound traffic from the ALB's security group or IP range will allow the health checks to function properly.Since the ECS service is stopping and starting new tasks frequently, this indicates that the ALB is marking tasks as unhealthy due to failed health checks. The most common reason for this is that the health check requests from the ALB cannot reach the tasks due to security group restrictions. By ensuring that the security group associated with the ECS tasks allows traffic from the ALB, the health checks can succeed, resolving the issue with failed tasks and frequent restarts.QUESTION 348A company that uses electronic patient health records runs a fleet of Amazon EC2 instances with an Amazon Linux operating system. The company must continuously ensure that the EC2 instances are running operating system patches and application patches that are in compliance with current privacy regulations. The company uses a custom repository to store application patches.A DevOps engineer needs to automate the deployment of operating system patches and application patches. The DevOps engineer wants to use both the default operating system patch repository and the custom patch repository.Which solution will meet these requirements with the LEAST effort?A. Use AWS Systems Manager to create a new custom patch baseline that includes the default operating system repository and the custom repository.Run the AWS-RunPatchBaseline document by using the Run command to verify and install patches. Use the BaselineOverride API to configure the new custom patch baseline.B. Use AWS Direct Connect to integrate the custom repository with the EC2 instances. Use Amazon EventBridge events to deploy the patches.C. Use the yum-config-manager command to add the custom repository to the /etc/yum.repos.d configuration. Run the yum-config-manager-enable command to activate the new repository.D. Use AWS Systems Manager to create a patch baseline for the default operating system repository and a second patch baseline for the custom repository. Run the AWS-RunPatchBaseline document by using the Run command to verify and install patches. Use the BaselineOverride API to configure the default patch baseline and the custom patch baseline.Answer: AExplanation:AWS Systems Manager Patch Manager allows for the creation of custom patch baselines that can include both the default operating system repository and a custom repository. By using the AWS-RunPatchBaseline document, you can automate the verification and installation of patches, ensuring that the EC2 instances remain compliant with the latest updates. Using the BaselineOverride API provides the flexibility to apply the custom baseline as needed without managing multiple baselines separately, reducing effort and complexity. This approach integrates all required patch sources in a streamlined, automated manner.QUESTION 349A company use an organization in AWS Organizations to manage multiple AWS accounts. The company has enabled all features enabled for the organization. The company configured the organization as a hierarchy of OUs under the root OU. The company recently registered all its OUs and enrolled all its AWS accounts in AWS Control Tower.The company needs to customize the AWS Control Tower managed AWS Config configuration recorder in each of the company's AWS accounts. The company needs to apply the customizations to both the existing AWS accounts and to any new AWS accounts that the company enrolls in AWS Control Tower in the future.Which combination of steps will meet these requirements? (Choose three.)A. Create a new AWS account. Create an AWS Lambda function in the new account to apply the customizations to the AWS Config configuration recorder in each AWS account in the organization.B. Create a new AWS account as an AWS Config delegated administrator. Create an AWS Lambda function in the delegated administrator account to apply the customizations to the AWS Config configuration recorder in the delegated administrator account.C. Configure an Amazon EventBridge rule in the AWS Control Tower management account to invoke an AWS Lambda function when the Organizations OU is registered or reregistered. Re-register the root Organizations OU.D. Configure the AWSControlTowerExecution IAM role in each AWS account in the organization to be assumable by an AWS Lambda function.Configure the Lambda function to assume the AWSControlTowerExecution IAM role.E. Create an IAM role in the AWS Control Tower management account that an AWS Lambda function can assume. Grant the IAM role permission to assume the AWSControlTowerExecution IAM role in any account in the organization. Configure the Lambda function to use the new IAM

role.F. Configure an Amazon EventBridge rule in the AWS Control Tower management account to invoke an AWS Lambda function when an AWS account is updated or enrolled in AWS Control Tower or when the landing zone is updated. Re-register each Organizations OU in the organization.Answer: CEDExplanation:C and F: Using Amazon EventBridge rules in the AWS Control Tower management account helps ensure that the Lambda function is triggered whenever there is an OU registration, re-registration, or account update/enrollment event. This setup allows the customizations to apply automatically to both existing and new AWS accounts.E: Creating an IAM role in the AWS Control Tower management account that the Lambda function can assume allows it to manage AWS Config in all accounts across the organization. This role needs permission to assume the AWSControlTowerExecution IAM role, which is already established in each account by AWS Control Tower.This combination of steps ensures that customizations to the AWS Config configuration recorder are applied automatically to both current and future accounts in AWS Control Tower with minimal manual intervention.QUESTION 350A company runs an application in an Auto Scaling group of Amazon EC2 instances behind an Application Load Balancer (ALB). The EC2 instances run Docker containers that make requests to a MySQL database that runs on separate EC2 instances.A DevOps engineer needs to update the application to use a serverless architecture.Which solution will meet this requirement with the FEWEST changes?A. Replace the containers that run on EC2 instances and the ALB with AWS Lambda functions. Replace the MySQL database with an Amazon Aurora Serverless v2 database that is compatible with MySQL.B. Replace the containers that run on EC2 instances with AWS Fargate. Replace the MySQL database with an Amazon Aurora Serverless v2 database that is compatible with MySQL.C. Replace the containers that run on EC2 instances and the ALB with AWS Lambda functions. Replace the MySQL database with Amazon DynamoDB tables.D. Replace the containers that run on EC2 instances with AWS Fargate. Replace the MySQL database with Amazon DynamoDB tables.Answer: BExplanation:This solution achieves a serverless architecture with the fewest changes. By moving the application containers from EC2 instances to AWS Fargate, the DevOps engineer eliminates the need to manage EC2 instances, creating a serverless, containerized solution. Replacing the MySQL database with Amazon Aurora Serverless v2 (which is MySQL-compatible) allows the application to continue using a relational database without requiring extensive modifications to the database layer or schema. This approach preserves compatibility with the existing MySQL database while leveraging serverless benefits.QUESTION 351A company uses an organization in AWS Organizations to manage 10 AWS accounts. All features are enabled, and trusted access for AWS CloudFormation is enabled.A DevOps engineer needs to use CloudFormation to deploy an IAM role to the Organizations management account and all member accounts in the organization.Which solution will meet these requirements with the LEAST operational overhead?A. Create a CloudFormation StackSet that has service-managed permissions. Set the root OU as a deployment target.B. Create a CloudFormation StackSet that has service-managed permissions. Set the root OU as a deployment target. Deploy a separate CloudFormation stack in the Organizations management account.C. Create a CloudFormation StackSet that has self-managed permissions. Set the root OU as a deployment target.D. Create a CloudFormation StackSet that has self-managed permissions. Set the root OU as a deployment target. Deploy a separate CloudFormation stack in the Organizations management account.Answer: AExplanation:Using a CloudFormation StackSet with service-managed permissions and targeting the root OU allows deployment of the IAM role across all accounts in the organization with minimal operational overhead. Service-managed permissions simplify the deployment process by automating the necessary permissions for CloudFormation to operate across all accounts in the organization, without requiring additional setup or manual intervention.This solution ensures that the IAM role is consistently deployed to both the management account and all member accounts in the organization with the least amount of effort.QUESTION 352A company runs an application that stores artifacts in an Amazon S3 bucket. The application has a large user base. The application writes a high volume of objects to the S3 bucket. The company has enabled event notifications for the S3 bucket.When the application writes an object to the S3 bucket, several processing tasks need to be performed simultaneously. The company's DevOps team needs to create an AWS Step Functions workflow to orchestrate the processing tasks.Which combination of steps should the DevOps team take to meet these requirements with the LEAST operational overhead? (Choose two.)A. Create a Standard workflow that contains a parallel state that defines the processing tasks. Create an Asynchronous Express workflow that contains a parallel state that defines the processing tasks.B. Create a Synchronous Express workflow that contains a map state that defines the processing tasks.C. Create an Amazon EventBridge rule to match when a new S3 object is created. Configure the EventBridge rule to invoke an AWS Lambda function. Configure the Lambda function to start the processing workflow.D. Create an Amazon EventBridge rule to match when a new S3 object is created. Configure the EventBridge rule to start the processing workflow.Answer: ADExplanation:Using an AWS Step Functions Standard or Asynchronous Express workflow with a parallel state enables the orchestration of multiple processing tasks concurrently, which is ideal for handling multiple tasks triggered by each new object event with low operational complexity.An Amazon EventBridge rule can be configured to directly start the Step Functions workflow upon the creation of a new S3 object, removing the need for

additional Lambda functions and thus reducing operational overhead.QUESTION 353A DevOps team supports an application that runs in an Amazon Elastic Container Service (Amazon ECS) cluster behind an Application Load Balancer (ALB). Currently, the DevOps team uses AWS CodeDeploy to deploy the application by using a blue/green all-at-once strategy. Recently, the DevOps team had to roll back a deployment when a new version of the application dramatically increased response times for requests.The DevOps team needs use to a deployment strategy that will allow the team to monitor a new version of the application before the team shifts all traffic to the new version. If a new version of the application increases response times, the deployment should be rolled back as quickly as possible.Which combination of steps will meet these requirements? (Choose two.)A. Modify the CodeDeploy deployment to use the CodeDeployDefault.ECSCanary10Percent5Minutes configuration.B. Modify the CodeDeploy deployment to use the CodeDeployDefault.ECSLinear10PercentEvery3Minutes configuration.C. Create an Amazon CloudWatch alarm to monitor the UnHealthyHostCount metric for the ALB. Set the alarm to activate if the metric is higher than the desired value. Associate the alarm with the CodeDeploy deployment group. Modify the deployment group to roll back when a deployment fails.D. Create an Amazon CloudWatch alarm to monitor the TargetResponseTime metric for the ALB. Set the alarm to activate if the metric is higher than the desired value. Associate the alarm with the CodeDeploy deployment group. Modify the deployment group to roll back when alarm thresholds are met.E. Create an Amazon CloudWatch alarm to monitor the TargetConnectionErrorCount metric for the ALB. Set the alarm to activate if the metric is higher than the desired value. Associate the alarm with the CodeDeploy deployment group. Modify the deployment group to roll back when alarm thresholds are met.Answer: ADExplanation:The CodeDeployDefault.ECSCanary10Percent5Minutes configuration uses a canary deployment strategy that shifts 10% of traffic to the new version initially and waits 5 minutes before proceeding. This provides time to monitor the new version's performance on a subset of traffic before shifting all traffic, allowing for quicker rollback if issues arise. Monitoring the TargetResponseTime metric for the ALB ensures that any increase in response times due to the new version is detected promptly.Associating this CloudWatch alarm with the CodeDeploy deployment group allows for automatic rollback if response times exceed the defined threshold, which meets the requirement of quickly rolling back when issues are detected.QUESTION 354A security team must record the configuration of AWS resources, detect issues, and send notifications for findings. The main workload in the AWS account consists of an Amazon EC2 Auto Scaling group that scales in and out several times during the day.The team wants to be notified within 2 days if any Amazon EC2 security group allows traffic on port 22 for 0.0.0.0/0. The team also needs a snapshot of the configuration of the AWS resources to be taken routinely.The security team has already created and subscribed to an Amazon Simple Notification Service (Amazon SNS) topic.Which solution meets these requirements?A. Configure AWS Config to use periodic recording for the AWS account. Deploy the vpc-sg-port-restriction-check AWS Config managed rule. Configure AWS Config to use the SNS topic as the target for notifications.B. Configure AWS Config to use configuration change recording for the AWS account. Deploy the vpc-sg-open-only-to-authorized-ports AWS Config managed rule. Configure AWS Config to use the SNS topic as the target for notifications.C. Configure AWS Config to use configuration change recording for the AWS account. Deploy the ssh-restricted AWS Config managed rule. Configure AWS Config to use the SNS topic as the target for notifications.D. Create an AWS Lambda function to evaluate security groups and publish a message to the SNS topic. Use an Amazon EventBridge rule to schedule the Lambda function to run once a day.Answer: A Explanation:Periodic recording with AWS Config ensures that resource configurations are recorded routinely, capturing snapshots of the AWS resources as requested by the security team.The vpc-sg-port-restriction-check AWS Config managed rule is specifically designed to detect when security groups allow unrestricted access on certain ports, such as port 22 (SSH) for 0.0.0.0/0, meeting the security team's requirement to detect and notify on such configurations.By configuring AWS Config to use the existing SNS topic for notifications, the team will be notified within the required timeframe if any issues are detected.QUESTION 355A company has proprietary data available by using an Amazon CloudFront distribution. The company needs to ensure that the distribution is accessible by only users from the corporate office that have a known set of IP address ranges. An AWS WAF web ACL is associated with the distribution and has a default action set to Count.Which solution will meet these requirements with the LEAST operational overhead?A. Create a new regex pattern set. Add the regex pattern set to a new rule group. Create a new web ACL that has a default action set to Block. Associate the web ACL with the CloudFront distribution. Add a rule that allows traffic based on the new rule group.B. Create an AWS WAF IP address set that matches the corporate office IP address range. Create a new web ACL that has a default action set to Allow. Associate the web ACL with the CloudFront distribution. Add a rule that allows traffic from the IP address set.C. Create a new regex pattern set. Add the regex pattern set to a new rule group. Set the default action on the existing web ACL to Allow. Add a rule that has priority 0 that allows traffic based on the regex pattern set.D. Create a WAF IP address set that matches the corporate office IP address range. Set the default action on the existing web ACL to Block. Add a rule that has priority 0 that allows traffic from the IP address set.Answer: DExplanation:WAF IP address set: By creating a WAF IP address set

that matches the known IP address ranges of the corporate office, you can explicitly allow only those IPs to access the distribution. Set the default action to Block on the existing web ACL and add a high-priority rule (priority 0) to allow traffic from the corporate IP address set. This configuration will block all traffic by default, except traffic coming from the allowed corporate IP range.QUESTION 356A company runs several applications in the same AWS account. The applications send logs to Amazon CloudWatch.A data analytics team needs to collect performance metrics and custom metrics from the applications. The analytics team needs to transform the metrics data before storing the data in an Amazon S3 bucket. The analytics team must automatically collect any new metrics that are added to the CloudWatch namespace.Which solution will meet these requirements with the LEAST operational overhead?A. Configure a CloudWatch metric stream to include metrics from the application and the CloudWatch namespace. Configure the metric stream to deliver the metrics to an Amazon Data Firehose delivery stream. Configure the Firehose delivery stream to invoke an AWS Lambda function to transform the data. Configure the delivery stream to send the transformed data to the S3 bucket.B. Configure a CloudWatch metrics stream to include all the metrics and to deliver the metrics to an Amazon Data Firehose delivery stream. Configure the Firehose delivery stream to invoke an AWS Lambda function to transform the data. Configure the delivery stream to send the transformed data to the S3 bucket.C. Configure metric filters for the CloudWatch logs to create custom metrics. Configure a CloudWatch metric stream to deliver the application metrics to the S3 bucket.D. Configure subscription filters on the application log groups to target an Amazon Data Firehose delivery stream. Configure the Firehose delivery stream to invoke an AWS Lambda function to transform the data. Configure the delivery stream to send the transformed data to the S3 bucket.Answer: BExplanation:CloudWatch metric stream: Setting up a metric stream for all metrics ensures that any new metrics added to the CloudWatch namespace are automatically included, meeting the requirement to capture new metrics without additional configuration.Amazon Kinesis Data Firehose: The metric stream can be configured to deliver data to a Firehose delivery stream, which allows transformation of the data using an AWS Lambda function before sending it to the S3 bucket.Lambda for data transformation: This setup allows the analytics team to transform the data as needed before it reaches S3, maintaining flexibility and scalability with minimal operational overhead.QUESTION 357A company uses an HPC platform to run analysis jobs for data. The company uses AWS CodeBuild to create container images and store the images on Amazon Elastic Container Registry (Amazon ECR). The images are then deployed on Amazon Elastic Kubernetes Service (Amazon EKS).To maintain compliance, the company needs to ensure that the images are signed before the images are deployed on Amazon EKS. The signing keys must be rotated periodically and must be managed automatically. The company needs to track who generates the signatures.Which solution will meet these requirements with the LEAST operational effort?A. Use CodeBuild to retrieve the image that was previously pushed to Amazon ECR. Use AWS Signer to sign the image. Use AWS CloudTrail to track who generates the signatures.B. Use AWS Lambda to retrieve the image that was previously pushed to Amazon ECR. Use a Lambda function to sign the image. Use Amazon CloudWatch to track who generates the signatures.C. Use AWS Lambda to retrieve the image that was previously pushed to Amazon ECR. Use AWS Signer to sign the image. Use Amazon CloudWatch to track who generates the signatures.D. Use CodeBuild to build the image. Sign the image by using AWS Signer before pushing the image to Amazon ECR. Use AWS CloudTrail to track who generates the signatures.Answer: DExplanation:CodeBuild for image building: Using CodeBuild for building the images aligns with the existing workflow, minimizing changes.AWS Signer for signing: AWS Signer can be integrated to sign images before they are pushed to Amazon ECR, ensuring compliance with minimal additional steps.AWS CloudTrail for tracking: CloudTrail can be used to track who generates the signatures, meeting the requirement for auditability.QUESTION 358A company uses an AWS CodeArtifact repository to store Python packages that the company developed internally. A DevOps engineer needs to use AWS CodeDeploy to deploy an application to an Amazon EC2 instance. The application uses a Python package that is stored in the CodeArtifact repository. A BeforeInstall lifecycle event hook will install the package.The DevOps engineer needs to grant the EC2 instance access to the CodeArtifact repository.Which solution will meet this requirement?A. Create a service-linked role for CodeArtifact. Associate the role with the EC2 instance. Use the aws codeartifact get-authorization-token CLI command on the instance.B. Configure a resource-based policy for the CodeArtifact repository that allows the ReadFromRepository action for the EC2 instance principal.C. Configure ACLs on the CodeArtifact repository to allow the EC2 instance to access the Python package.D. Create an instance profile that contains an IAM role that has access to CodeArtifact. Associate the instance profile with the EC2 instance. Use the aws codeartifact login CLI command on the instance.Answer: DExplanation:Instance profile with IAM role: By creating an instance profile that contains an IAM role with permissions to access the CodeArtifact repository, the EC2 instance is granted the necessary access to retrieve packages.aws codeartifact login: The aws codeartifact login CLI command configures the Python package manager (such as pip) on the instance to use the CodeArtifact repository, handling authentication and authorization automatically using the instance profile credentials.QUESTION 359A company has a file-reading application that saves files to a database that runs on Amazon EC2

instances. Regulations require the company to delete files from EC2 instances every day at a specific time. The company must delete database records that are older than 60 days.The database record deletion must occur after the file deletions. The company has created scripts to delete files and database records. The company needs to receive an email notification for any failure of the deletion scripts.Which solution will meet these requirements with the LEAST development effort?A. Use AWS Systems Manager State Manager to automatically invoke a Systems Manager Automation document at the specified time each day.Configure the Automation document to use a run command to run the deletion scripts in sequential order. Create an Amazon EventBridge rule to use Amazon Simple Notification Service (Amazon SNS) to send failure notifications to the company.B. Use AWS Systems Manager State Manager to automatically invoke a Systems Manager Automation document at the specified time each day.Configure the Automation document to use a run command to run the deletion scripts in sequential order. Create a conditional statement inside the Automation document as the last step to check for errors. Use Amazon Simple Email Service (Amazon SES) to send failure notifications as email messages to the company.C. Create an Amazon EventBridge rule that invokes an AWS Lambda function at the specified time. Add the necessary permissions for the invocation to the Lambda function's resource-based policy. Configure the Lambda function to run the deletion scripts in sequential order. Configure the Lambda function to use Amazon Simple Notification Service (Amazon SNS) to send failure notifications to the company.D. Create an Amazon EventBridge rule that invokes an AWS Lambda function at the specified time. Add the necessary permissions for the invocation to the Lambda function's resource-based policy. Configure the Lambda function to run the deletion scripts in sequential order. Configure the Lambda function to use Amazon Simple Email Service (Amazon SES) to send failure notifications as email messages to the company.Answer: AExplanation:AWS Systems Manager State Manager: This service allows for scheduled automation tasks, including invoking a Systems Manager Automation document at a specified time each day.Sequential execution in Automation document: Using the Automation document with a run command enables the scripts to execute in the required order, with the database record deletion occurring after file deletion.Amazon EventBridge rule and SNS for notifications: An EventBridge rule can capture any failures from the Automation document execution and send failure notifications via Amazon SNS. This setup provides reliable error handling and notification with minimal development effort.QUESTION 360A company uses an organization in AWS Organizations that has all features enabled to manage its AWS accounts. Amazon EQ instances run in the AWS accounts.The company requires that all current EC2 instances must use Instance Metadata Service Version 2 (IMDSv2). The company needs to block AWS API calls that originate from EC2 instances that do not use IMDSv2.Which solution will meet these requirements?A. Create a new SCP statement that denies the ec2:RunInstances action when the ec2:MetadataHttpTokens condition key is not equal to the value of required. Attach the SCP to the root of the organization.B. Create a new SCP statement that denies the ec2:RunInstances action when the ec2:MetadataHttpPutResponseHopLimit condition key value is greater than two. Attach the SCP to the root of the organization.C. Create a new SCP statement that denies "*" when the ec2:RoleDelivery condition key value is less than two. Attach the SCP to the root of the organization.D. Create a new SCP statement that denies when the ec2:MetadataHttpTokens condition key value is not equal to required. Attach the SCP to the root of the organization.Answer: AExplanation:SCP with ec2:condition key: The ec2:MetadataHttpTokens condition key can enforce the use of IMDSv2 by requiring the ec2:RunInstances action to use the MetadataHttpTokens parameter set to required. This configuration prevents instances from launching without IMDSv2 enabled. Attaching the SCP to the root of the organization: Applying this SCP at the organization root level ensures the policy applies across all accounts, meeting the requirement for all EC2 instances to use IMDSv2.QUESTION 361A DevOps team supports an application that runs on a large number of Amazon EC2 instances in an Auto Scaling group. The DevOps team uses AWS CloudFormation to deploy the EC2 instances. The application recently experienced an issue. A single instance returned errors to a large percentage of requests. The EC2 instance responded as healthy to both Amazon EC2 and Elastic Load Balancing health checks.The DevOps team collects application logs in Amazon CloudWatch by using the embedded metric format. The DevOps team needs to receive an alert if any EC2 instance is responsible for more than half of all errors.Which combination of steps will meet these requirements with the LEAST operational overhead? (Choose two.)A. Create a CloudWatch Contributor Insights rule that groups logs from the CloudWatch application logs based on instance ID and errors.B. Create a resource group in AWS Resource Groups. Use the CloudFormation stack to group the resources for the application. Add the application to CloudWatch Application Insights. Use the resource group to identify the application.C. Create a metric filter for the application logs to count the occurrence of the term "Error." Create a CloudWatch alarm that uses the METRIC_COUNT function to determine whether errors have occurred. Configure the CloudWatch alarm to send a notification to an Amazon Simple Notification Service (Amazon SNS) topic to notify the DevOps team.D. Create a CloudWatch alarm that uses the INSIGHT_RULE_METRIC function to determine whether a specific instance is responsible for more than half of all errors reported by EC2 instances. Configure the CloudWatch alarm to send a notification to an Amazon Simple Notification Service (Amazon SNS) topic to notify the DevOps team.E. Create a CloudWatch subscription filter

for the application logs that filters for errors and invokes an AWS Lambda function. Configure the Lambda function to send the instance ID and error and in a notification to an Amazon Simple Notification Service (Amazon SNS) topic to notify the DevOps team.Answer: ADExplanation:Using a CloudWatch Contributor Insights rule to group logs based on instance ID and error frequency allows the team to pinpoint which instance is contributing to errors, providing detailed insights without additional custom logic.The INSIGHT_RULE_METRIC function in CloudWatch can be used to evaluate the Contributor Insights data to determine if any instance is responsible for more than half of all errors. This approach enables automated monitoring and alerts based on specific error patterns.QUESTION 362A company is using AWS CloudFormation to perform deployments of its application environment. A deployment failed during a recent update to the existing CloudFormation stack. A DevOps engineer discovered that some resources in the stack were manually modified.The DevOps engineer needs a solution that detects manual modification of resources and sends an alert to the DevOps lead.Which solution will meet these requirements with the LEAST operational effort?A. Create an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the DevOps lead to the topic by using an email address. Create an AWS Config managed rule that has the CLOUDFORMATION_STACK_DRIFT_DETECTION_CHECK identifier. Create an Amazon EventBridge rule that is invoked on the NON_COMPLIANT resources status. Set the SNS topic as the rule target.B. Tag all CloudFormation resources with a specific tag. Create an AWS Config custom rule by using the AWS Config Rules Development Kit Library (RDKlib) that checks all resource changes that have the specific tag. Configure the custom rule to mark all the tagged resource changes as NON_COMPLIANT when the change is not performed by CloudFormation. Create an Amazon EventBridge rule that is invoked on the NON_COMPUANT resources status. Create an AWS Lambda function that sends an email message to the DevOps lead. Set the Lambda function as the rule target.C. Create an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the DevOps lead to the topic by using an email address. Create an AWS Config managed rule that has the CLOUDFORMATION_STACK_DRIFT_DETECTION_CHECK identifier. Create an Amazon EventBridge rule that is invoked on the COMPLIANT resources status. Set the SNS topic as the rule target.D. Create an AWS Config managed rule that has the CLOUDFORMATION_STACK_DRIFT_DETECTION_CHECK identifier. Create an Amazon EventBridge rule that is invoked on the NON_COMPLIANT resources status. Create an AWS Lambda function that sends an email message to the DevOps lead. Set the Lambda function as the rule target.Answer: AExplanation:AWS Config managed rule: The CLOUDFORMATION_STACK_DRIFT_DETECTION_CHECK managed rule automatically detects drift (manual changes) in CloudFormation stacks, which directly addresses the need to detect manual modifications of resources.Amazon SNS for alerts: By creating an SNS topic and subscribing the DevOps lead to it, you can easily send notifications when drift is detected.EventBridge for triggering notifications: Setting up an EventBridge rule that triggers on the NON_COMPLIANT status from the AWS Config rule allows for automated alerting without the need for custom logic or additional resources.QUESTION 363A DevOps engineer deployed multiple AWS accounts by using AWS Control Tower to support different business, technical, and administrative units in a company. A security team needs the DevOps engineer to automate AWS Control Tower guardrails for the company. The guardrails must be applied to all accounts in an OU of the company's organization in AWS Organizations.The security team needs a solution that has version control and can be reviewed and rolled back if necessary. The security team will maintain the management of the solution in its OU. The security team wants to limit the type of guardrails that are allowed and allow only new guardrails that are approved by the security team.Which solution will meet these requirements with the MOST operational efficiency?A. Create individual AWS CloudFormation templates that align to a guardrail. Store the templates in an AWS CodeCommit repository. Create an AWS::ControlTower::EnableControl logical resource in the template for each OU in the organization. Configure an AWS Code Build project that an Amazon EventBridge rule will invoke for the security team's AWS CodeCommit changes.B. Create individual AWS CloudFormation templates that align to a guardrail. Store the templates in an AWS CodeCommit repository. Create an AWS::ControlTower::EnableControl logical resource in the template for each account in the organization. Configure an AWS CodePipeline pipeline in the security team's account. Advise the security team to invoke the pipeline and provide these parameters when starting the pipeline.C. Create individual AWS CloudFormation templates that align to a guardrail. Store the templates in an AWS CodeCommit repository. Create an AWS::ControlTower::EnableControl logical resource in the template for each OU in the organization. Configure an AWS CodePipeline pipeline in the security team's account that an Amazon EventBridge rule will invoke for the security team's CodeCommit changes.D. Configure an AWS CodePipeline pipeline in the security team's account that an Amazon EventBridge rule will invoke for PutObject events to an Amazon S3 bucket. Create individual AWS CloudFormation templates that align to a guardrail. Store the templates in the S3 bucket. Create an AWS::ControlTower::EnableControl logical resource in the template for each OU in the organization.Answer: CExplanation:CloudFormation templates: Creating individual CloudFormation templates for each guardrail allows for version control, making it easy to review, approve, and roll back changes if necessary.AWS CodeCommit for storage: Storing the templates in AWS CodeCommit provides a robust version control system for

managing the guardrails.AWS::ControlTower::EnableControl: Including this resource in the templates enables the speific guardrails for the targeted organizational units (OUs).AWS CodePipeline with EventBridge: Configuring a CodePipeline that triggers on changes in CodeCommit using an Amazon EventBridge rule automates the application of guardrails while allowing the security team to maintain control over the deployment process. This setup enhances operational efficiency by automating the deployment while allowing for oversight and version control.QUESTION 364A company runs a web application on Amazon Elastic Kubernetes Service (Amazon EKS). The company uses Amazon CloudFront to distribute the application. The company recently enabled AWS WAF. The company set up Amazon CloudWatch Logs to send logs to an aws-waf-logs log group.The company wants a DevOps engineer to receive alerts if there are sudden changes in blocked traffic. The company does not want to receive alerts for other changes in AWS WAF log behavior. The company will tune AWS WAF rules over time.The DevOps engineer is currently subscribed to an Amazon Simple Notification Service (Amazon SNS) topic in the environment.Which solution will meet these requirements?A. Create a CloudWatch Logs metrics filter for blocked requests on the AWS WAF log group to create a custom metric. Create a CloudWatch alarm by using CloudWatch anomaly detection and the published custom metric. Configure the alarm to notify the SNS topic to alert the DevOps engineer.B. Create a CloudWatch anomaly detector for the log group. Create a CloudWatch alarm by using metrics that the CloudWatch anomaly detector publishes. Use the high setting for the LogAnomalyPriority metric. Configure the alarm to go into alarm state if a static threshold of one anomaly is detected. Configure the alarm to notify the SNS topic to alert the DevOps engineer.C. Create a CloudWatch metrics filter for counted requests on the AWS WAF log group to create a custom metric. Create a CloudWatch alarm that activates when the sum of blocked requests in the custom metric during a period of 1 hour is greater than a static estimate for the acceptable number of blocked requests in 1 hour. Configure the alarm to notify the SNS topic to alert the DevOps engineer.D. Create a CloudWatch anomaly detector for the log group. Create a CloudWatch alarm by using metrics that the CloudWatch anomaly detector publishes. Use the medium setting for the LogAnomalyPriority metric. Configure the alarm to go into alarm state if a sum of anomalies over 1 hour is greater than an expected value. Configure the alarm to notify the SNS topic to alert the DevOps engineer.Answer: AExplanation:CloudWatch Logs metrics filter: By creating a metrics filter specifically for blocked requests in the AWS WAF log group, you can generate a custom metric that reflects only the relevant traffic being blocked, isolating it from other log behaviors.CloudWatch anomaly detection: Using CloudWatch anomaly detection on this custom metric allows for the identification of sudden changes in blocked traffic patterns without being affected by normal fluctuations or rule tuning, addressing the requirement to alert only on significant changes. SNS notifications: Configuring the alarm to notify the SNS topic ensures that the DevOps engineer receives timely alerts about any unexpected spikes in blocked traffic, aligning with the company's alerting strategy.QUESTION 365A video platform company is migrating its video catalog to AWS. The company will host MP4 videos files in an Amazon S3 bucket. The company will use Amazon CloudFront and Amazon EC2 instances to serve the video files.Users first connect to a frontend application that redirects to a video URL. The video URL contains an authorization token in CloudFront. The cache is activated on the CloudFront distribution. Authorization token check activity needs to be logged in Amazon CloudWatch.The company wants to prevent direct access to video files on CloudFront and Amazon S3 and wants to implement checks of the authorization token that the frontend application provides. The company also wants to perform regular rolling updates of the code that checks the authorization token signature. Which solution will meet these requirements with the LEAST operational effort?A. Implement an authorization token check in Lambda@Edge as a trigger on the CloudFront distribution. Enable CloudWatch logging for the Lambda@Edge function. Attach the Lambda@Edge function to the CloudFront distribution. Implement CloudFront continuous deployment to perform updates.B. Implement an authorization token check in CloudFront Functions. Enable CloudWatch logging for the CloudFront function. Attach the CloudFront function to the CloudFront distribution. Implement CloudFront continuous deployment to perform updates.C. Implement an authorization token check in the application code that is installed on the EC2 instances. Install the CloudWatch agent on the EC2 instances. Configure the application to log to the CloudWatch agent. Implement a second CloudFront distribution. Migrate the traffic from the first CloudFront distribution by using Amazon Route 53 weighted routing.D. Implement an authorization token check in CloudFront Functions. Enable CloudWatch logging for the CloudFront function. Attach the CloudFront function to the CloudFront distribution. Implement a second CloudFront distribution. Migrate the traffic from the first CloudFront distribution by using Amazon Route 53 weighted routing.Answer: BExplanation:CloudFront Functions: Using CloudFront Functions to check the authorization token allows you to perform lightweight, low-latency processing directly at the edge. This is particularly suitable for simple token validation without the overhead of a full Lambda function.CloudWatch Logging: Enabling CloudWatch logging for the CloudFront function provides visibility into authorization token check activities, meeting the requirement for logging without additional operational effort.Continuous Deployment: Implementing continuous deployment for CloudFront functions allows for seamless updates to the token validation logic without significant downtime or operational

complexity.QUESTION 366A company uses an organization in AWS Organizations to manage multiple AWS accounts in a hierarchical structure. An SCP that is associated with the organization root allows IAM users to be created.A DevOps team must be able to create IAM users with any level of permissions. Developers must also be able to create IAM users. However, developers must not be able to grant new IAM users excessive permissions. The developers have the CreateAndManageUsers role in each account.The DevOps team must be able to prevent other users from creating IAM users.Which combination of steps will meet these requirements? (Choose two.)A. Create an SCP in the organization to deny users the ability to create and modify IAM users. Attach the SCP to the root of the organization. Attach the CreateAndManageUsers role to developers.B. Create an SCP in the organization to grant users that have the DeveloperBoundary policy attached the ability to create new IAM users and to modify IAM users. Configure the SCP to require users to attach the PermissionBoundaries policy to any new IAM user. Attach the SCP to the root of the organization.C. Create an IAM permissions policy named PermissionBoundaries within each account. Configure the PermissionBoundaries policy to specify the maximum permissions that a developer can grant to a new IAM user.D. Create an IAM permissions policy named PermissionBoundaries within each account. Configure PermissionsBoundaries to allow users who have the PermissionBoundaries policy to create new IAM users.E. Create an IAM permissions policy named DeveloperBoundary within each account. Configure the DeveloperBoundary policy to allow developers to create IAM users and to assign policies to IAM users of only if the developer includes the PermissionBoundaries policy as the permissions boundary.Attach the DeveloperBoundary policy to the CreateAndManageUsers role within each account.Answer: ACExplanation:By creating a Service Control Policy (SCP) that denies users the ability to create and modify IAM users and attaching it to the root of the organization, you ensure that only users in the DevOps team (who have the necessary permissions) can create IAM users. This effectively prevents developers from creating excessive IAM users while still allowing the DevOps team the required capabilities.Creating an IAM permissions policy named PermissionBoundaries allows for defining the maximum permissions that can be granted to new IAM users by developers. By implementing this policy, you ensure that even if developers have the ability to create IAM users, they cannot assign more permissions than specified in the boundaries policy, thereby preventing excessive permissions.QUESTION 367A company has deployed a landing zone that has a well-defined AWS Organizations structure and an SCP. The company's development team can create their AWS resources only by using AWS CloudFormation and the AWS Cloud Development Kit (AWS CDK).A DevOps engineer notices that Amazon Simple Queue Service (Amazon SQS) queues that are deployed in different CloudFormation stacks have different configurations. The DevOps engineer also notices that the application cost allocation tag is not always set.The DevOps engineer needs a solution that will enforce tagging and promote the reuse of code. The DevOps engineer needs to avoid different configurations for the deployed SQS queues.What should the DevOps engineer do to meet these requirements?A. Create an Organizations tag policy to enforce the cost allocation tag in CloudFormation stacks. Instruct the development team to use CloudFormation to define SQS queues. Instruct the development team to deploy the SQS queues by using CloudFormation StackSets.B. Update the SCP to enforce the cost allocation tag in CloudFormation stacks. Instruct the development team to use CloudFormation modules to define SQS queues. Instruct the development team to deploy the SQS queues by using CloudFormation stacks.C. Use AWS CDK tagging to enforce the cost allocation tag in CloudFormation StackSets. Instruct the development team to use the AWS CDK to define SQS queues. Instruct the development team to deploy the SQS queues by using CDK stacks.D. Use AWS CDK tagging to enforce the cost allocation tag in CloudFormation stacks. Instruct the development team to use the AWS CDK to define SQS queues. Instruct the development team to deploy the SQS queues by using CDK feature flags.Answer: CExplanation: AWS CDK Tagging: By using AWS CDK's tagging capabilities, the DevOps engineer can enforce specific tags (like cost allocation tags) on the SQS queues at the code level. This ensures consistency across all deployments, as the tagging logic is incorporated into the infrastructure as code.Defining SQS Queues with AWS CDK: Instructing the development team to use the AWS CDK to define their SQS queues promotes code reuse and consistency in configuration. The CDK allows for defining resources in a more programmatic way, which can help standardize the SQS configuration across different stacks.CloudFormation StackSets: By deploying the SQS queues using CDK stacks, the development team can easily manage and deploy consistent configurations across multiple accounts or regions through StackSets, ensuring that all resources adhere to the same standards.QUESTION 368A DevOps team manages a company's AWS account. The company wants to ensure that specific AWS resource configuration changes are automatically reverted.Which solution will meet this requirement?A. Use AWS Config rules to detect changes in resource configurations. Configure remediation action that uses AWS Systems Manager Automation documents to revert the configuration changes.B. Use Amazon CloudWatch alarms to monitor resource metrics. When an alarm is activated, use an Amazon Simple Notification Service (Amazon SNS) topic to notify an administrator to manually reverts the configuration changes.C. Use AWS CloudFormation to create a stack that deploys the necessary configuration changes. Update the stack when configuration changes need to be reverted.D. Use AWS Trusted Advisor to check for noncompliant configurations. Manually apply necessary changes

based on Trusted Advisor recommendations.Answer: AExplanation:AWS Config Rules: AWS Config allows you to set up rules that monitor the configuration of AWS resources. You can define rules that detect specific configuration changes, ensuring compliance with your desired configuration standards.Remediation Action: By configuring a remediation action in AWS Config, you can automatically trigger a remediation process (using Systems Manager Automation documents) to revert any unwanted changes as soon as they are detected. This provides an automated and efficient way to maintain the desired state of your AWS resources.QUESTION 369When thinking of DynamoDB, what are true of Local Secondary Key properties?A. Either the partition key or the sort key can be different from the table, but not both.B. Only the sort key can be different from the table.C. The partition key and sort key can be different from the table.D. Only the partition key can be different from the table.Answer: B Explanation:Global secondary index ?an index with a partition key and a sort key that can be different from those on the table. A global secondary index is considered "global" because queries on the index can span all of the data in a table, across all partitions. Reference: http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/SecondaryIndexes.htmlQUESTION 370Which deployment method, when using AWS Auto Scaling Groups and Auto Scaling Launch Configurations, enables the shortest time to live for individual servers?A. Pre-baking AMIs with all code and configuration on deploys.B. Using a Dockerfile bootstrap on instance launch.C. Using UserData bootstrapping scripts.D. Using AWS EC2 Run Commands to dynamically SSH into fleets.Answer: AExplanation:Note that the bootstrapping process can be slower if you have a complex application or multiple applications to install. Managing a fleet of applications with several build tools and dependencies can be a challenging task during rollouts. Furthermore, your deployment service should be designed to do faster rollouts to take advantage of Auto Scaling. Prebaking is a process of embedding a significant portion of your application artifacts within your base AMI. During the deployment process you can customize application installations by using EC2 instance artifacts such as instance tags, instance metadata, and Auto Scaling groups.Reference:https://d0.awsstatic.com/whitepapers/overview-of-deployment-options-on-aws.pdfResources From:
1.2025 Latest Braindump2go DOP-C02 Exam Dumps (PDF & VCE) Free Share:**https://www.braindump2go.com/dop-c02.html**
2.2025 Latest Braindump2go DOP-C02 PDF and DOP-C02 VCE Dumps Free Share:
https://drive.google.com/drive/folders/1FhCZoaDCriYOlfYbMyFXhVN9z4p7HNoX?usp=sharing3.2025 Free Braindump2go DOP-C02 Exam Questions Download:
https://www.braindump2go.com/free-online-pdf/DOP-C02-VCE-Dumps(340-370).pdfFree Resources from Braindump2go,We Devoted to Helping You 100% Pass All Exams!