Free Questions for DOP-C02 by go4braindumps

Shared by Odonnell on 24-05-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

A company's application teams use AWS CodeCommit repositories for their applications. The application teams have repositories in multiple AWS accounts. All accounts are in an organization in AWS Organizations.

Each application team uses AWS IAM Identity Center (AWS Single Sign-On) configured with an external IdP to assume a developer IAM role. The developer role allows the application teams to use Git to work with the code in the repositories.

A security audit reveals that the application teams can modify the main branch in any repository. A DevOps engineer must implement a solution that allows the application teams to modify the main branch of only the repositories that they manage.

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

Options:

A- Update the SAML assertion to pass the user's team name. Update the IAM role's trust policy to add an access-team session tag that has the team name.

B- Create an approval rule template for each team in the Organizations management account. Associate the template with all the repositories. Add the developer role ARN as an approver.

C- Create an approval rule template for each account. Associate the template with all repositories. Add the 'aws:ResourceTag/access-team':'$;{aws:PrincipalTag/access-team}' condition to the approval rule template.

D- For each CodeCommit repository, add an access-team tag that has the value set to the name of the associated team.

E- Attach an SCP to the accounts. Include the following statement:

```
{
    "Effect": "Allow",
    "Action": [
        "codecommit:GitPush",
        "codecommit:PutFile",
        "codecommit:Merge*"
    ],
    "Resource": "*",
    "Condition": {
        "StringEqualsIfExists": {
            "codecommit:References": ["refs/heads/main"]
        },
        "StringNotEquals": {
            "aws:ResourceTag/access-team": "$ ;{aws:PrincipalTag/access-t
        },
        "Null": {
            "codecommit:References": "false"
        }
    }
}
```

F- Create an IAM permissions boundary in each account. Include the following statement:

```
{
    "Effect": "Allow",
    "Action": [
        "codecommit:GitPush",
        "codecommit:PutFile",
        "codecommit:Merge*"
    ],
    "Resource": "*",
    "Condition": {
        "StringEqualsIfExists": {
            "codecommit:References": ["refs/heads/main"]
        },
        "StringNotEquals": {
            "aws:ResourceTag/access-team": "$ ;{aws:PrincipalTag/access-t
        },
        "Null": {
            "codecommit:References": "false"
        }
    }
}
```

Answer:

A, D, F

Explanation:

Short To meet the requirements, the DevOps engineer should update the SAML assertion to pass the user's team name, update the IAM role's trust policy to add an access-team session tag that has the team name, create an IAM permissions boundary in each account, and for each CodeCommit repository, add an access-team tag that has the value set to the name of the associated team.

Updating the SAML assertion to pass the user's team name allows the DevOps engineer to use IAM tags to identify which team a user belongs to. This can help enforce fine-grained access control based on the user's team membership1.

Updating the IAM role's trust policy to add an access-team session tag that has the team name allows the DevOps engineer to use IAM condition keys to restrict access based on the session tag value2. For example, the DevOps engineer can use the aws:PrincipalTag condition key to match the access-team tag of the user with the access-team tag of the repository3.

Creating an IAM permissions boundary in each account allows the DevOps engineer to set the maximum permissions that an identity-based policy can grant to an IAM entity. An entity's permissions boundary allows it to perform only the actions that are allowed by both its identity-based policies and its permissions boundaries4. For example, the DevOps engineer can use a permissions boundary policy to limit the actions that a user can perform on CodeCommit repositories based on their access-team tag5.

For each CodeCommit repository, adding an access-team tag that has the value set to the name of the associated team allows the DevOps engineer to use resource tags to identify which team manages a repository. This can help enforce fine-grained access control based on the resource tag value6.

The other options are incorrect because:

Creating an approval rule template for each team in the Organizations management account is not a valid option, as approval rule templates are not supported by AWS Organizations. Approval rule templates are specific to CodeCommit and can only be associated with one or more repositories in the same AWS Region where they are created7.

Creating an approval rule template for each account is not a valid option, as approval rule templates are not designed to restrict access to modify branches. Approval rule templates are designed to require approvals from specified users or groups before merging pull requests8.

Attaching an SCP to the accounts is not a valid option, as SCPs are not designed to restrict access based on tags. SCPs are designed to restrict access based on service actions and resources across all users and roles in an organization's account9.

Question 2

Question Type: **MultipleChoice**

A DevOps engineer is building an application that uses an AWS Lambda function to query an Amazon Aurora MySQL DB cluster. The Lambda function performs only read queries. Amazon EventBridge events invoke the Lambda function.

As more events invoke the Lambda function each second, the database's latency increases and the database's throughput decreases. The DevOps engineer needs to improve the performance of the application.

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

Options:

A- Use Amazon RDS Proxy to create a proxy. Connect the proxy to the Aurora cluster reader endpoint. Set a maximum connections percentage on the proxy.

B- Implement database connection pooling inside the Lambda code. Set a maximum number of connections on the database connection pool.

C- Implement the database connection opening outside the Lambda event handler code.

D- Implement the database connection opening and closing inside the Lambda event handler code.

E- Connect to the proxy endpoint from the Lambda function.

F- Connect to the Aurora cluster endpoint from the Lambda function.

Answer:

A, C, E

Explanation:

Verified Answer: A, C, and E

Short To improve the performance of the application, the DevOps engineer should use Amazon RDS Proxy, implement the database connection opening outside the Lambda event handler code, and connect to the proxy endpoint from the Lambda function.

Amazon RDS Proxy is a fully managed, highly available database proxy for Amazon Relational Database Service (RDS) that makes applications more scalable, more resilient to database failures, and more secure1. By using Amazon RDS Proxy, the DevOps engineer can reduce the overhead of opening and closing connections to the database, which can improve latency and throughput2.

The DevOps engineer should connect the proxy to the Aurora cluster reader endpoint, which allows read-only connections to one of the Aurora Replicas in the DB cluster3. This can help balance the load across multiple read replicas and improve performance for read-intensive workloads4.

The DevOps engineer should implement the database connection opening outside the Lambda event handler code, which means using a global variable to store the database connection object5. This can enable connection reuse across multiple invocations of the Lambda function, which can reduce latency and improve performance.

The DevOps engineer should connect to the proxy endpoint from the Lambda function, which is a unique URL that represents the proxy. This can allow the Lambda function to access the database through the proxy, which can provide benefits such as connection pooling, load balancing, failover handling, and enhanced security.

The other options are incorrect because:

Implementing database connection pooling inside the Lambda code is unnecessary and redundant when using Amazon RDS Proxy, which already provides connection pooling as a service.

Implementing the database connection opening and closing inside the Lambda event handler code is inefficient and costly, as it can increase latency and consume more resources for each invocation of the Lambda function.

Connecting to the Aurora cluster endpoint from the Lambda function is not optimal for read-only queries, as it can direct traffic to either the primary instance or one of the Aurora Replicas in the DB cluster. This can result in inconsistent performance and potential conflicts with write operations on the primary instance.

Question 3

Question Type: MultipleChoice

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?

Options:

A- 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. Use the event to invoke the image

pipeline. Re-upload the container to the repository.

B- Use EC2 Image Builder to create a container image pipeline. Use Amazon ECR as the target repository. Enable Amazon GuardDuty Malware Protection on the container workload. Create an Amazon EventBridge rule to capture a GuardDuty finding event. Use the event to invoke the image pipeline.

C- Create an AWS CodeBuild project to create a container image. Use Amazon ECR as the target repository. Turn on basic scanning on the repository. Create an Amazon EventBridge rule to capture an ECR image action event. Use the event to invoke the CodeBuild project. Re-upload the container to the repository.

D- Create an AWS CodeBuild project to create a container image. Use Amazon ECR as the target repository. 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.

Question 4

Question Type: **MultipleChoice**

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

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

Options:

A- Configure AWS CloudTrail to send log 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 log 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 log data events to an Amazon CloudWatch Logs log group. Create a CloudWatch logs metric filter to match failed Consolel_ogin 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 log 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:

C

Question 5

Question Type: MultipleChoice

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

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

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

Which solution will meet these requirements?

Options:

A- Create an SCP that allows the services that IAM Access Analyzer identifies. Create an OU for the account. Move the account into the new OU. Attach the new SCP to the new OU. Detach the default FullAWSAccess SCP from the new OU.

B- Create an SCP that denies the services that IAM Access Analyzer identifies. Create an OU for the account. Move the account into the new OIJ. Attach the new SCP to the new OU.

C- Create an SCP that allows the services that IAM Access Analyzer identifies. Attach the new SCP to the organization's root.

D- Create an SCP that allows the services that IAM Access Analyzer identifies. Create an OU for the account. Move the account into the new OU. Attach the new SCP to the management account. Detach the default FullAWSAccess SCP from the new OU.

Answer:

A

Explanation:

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

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

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

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

The other options are not correct because they do not meet the requirements or follow best practices. Creating an SCP that denies the services that IAM Access Analyzer identifies is not a good option because it might not cover all possible services that are not approved or required for the application. A deny policy is also more difficult to maintain and update than an allow policy. Creating an SCP that allows the services that IAM Access Analyzer identifies and attaching it to the organization's root is not a good option because it might affect other accounts and OUs in the organization that have different service requirements or approvals. Creating an SCP that allows the services that IAM Access Analyzer identifies and attaching it to the management account is not a valid option because SCPs cannot be attached directly to accounts, only to OUs or roots.

References:

1: Using AWS Identity and Access Management Access Analyzer - AWS Identity and Access Management

2: Generate a policy based on access activity - AWS Identity and Access Management

3: Organizing your accounts into OUs - AWS Organizations

4: Service control policies - AWS Organizations

5: How SCPs work - AWS Organizations

Question 6

Question Type: MultipleChoice

A company deploys a web application on Amazon EC2 instances that are behind an Application Load Balancer (ALB). The company stores the application code in an AWS CodeCommit repository. When code is merged to the main branch, an AWS Lambda function invokes an AWS CodeBuild project. The CodeBuild project packages the code, stores the packaged code in AWS CodeArtifact, and invokes AWS Systems Manager Run Command to deploy the packaged code to the EC2 instances.

Previous deployments have resulted in defects, EC2 instances that are not running the latest version of the packaged code, and inconsistencies between instances.

Which combination of actions should a DevOps engineer take to implement a more reliable deployment solution? (Select TWO.)

Options:

A- Create a pipeline in AWS CodePipeline that uses the CodeCommit repository as a source provider. Configure pipeline stages that run the CodeBuild project in parallel to build and test the application. In the pipeline, pass the CodeBuild project output artifact to an AWS CodeDeploy action.

B- Create a pipeline in AWS CodePipeline that uses the CodeCommit repository as a source provider. Create separate pipeline stages that run a CodeBuild project to build and then test the application. In the pipeline, pass the CodeBuild project output artifact to an AWS CodeDeploy action.

C- Create an AWS CodeDeploy application and a deployment group to deploy the packaged code to the EC2 instances. Configure the

ALB for the deployment group.

D- Create individual Lambda functions that use AWS CodeDeploy instead of Systems Manager to run build, test, and deploy actions.

E- Create an Amazon S3 bucket. Modify the CodeBuild project to store the packages in the S3 bucket instead of in CodeArtifact. Use deploy actions in CodeDeploy to deploy the artifact to the EC2 instances.

Answer:

A, C

Explanation:

To implement a more reliable deployment solution, a DevOps engineer should take the following actions:

Create a pipeline in AWS CodePipeline that uses the CodeCommit repository as a source provider. Configure pipeline stages that run the CodeBuild project in parallel to build and test the application. In the pipeline, pass the CodeBuild project output artifact to an AWS CodeDeploy action. This action will improve the deployment reliability by automating the entire process from code commit to deployment, reducing human errors and inconsistencies. By running the build and test stages in parallel, the pipeline can also speed up the delivery time and provide faster feedback.By using CodeDeploy as the deployment action, the pipeline can leverage the features of CodeDeploy, such as traffic shifting, health checks, rollback, and deployment configuration123

Create an AWS CodeDeploy application and a deployment group to deploy the packaged code to the EC2 instances. Configure the ALB for the deployment group. This action will improve the deployment reliability by using CodeDeploy to orchestrate the deployment across multiple EC2 instances behind an ALB. CodeDeploy can perform blue/green deployments or in-place deployments with traffic shifting, which can minimize downtime and reduce risks. CodeDeploy can also monitor the health of the instances during and after the

deployment, and automatically roll back if any issues are detected.By configuring the ALB for the deployment group, CodeDeploy can register and deregister instances from the load balancer as needed, ensuring that only healthy instances receive traffic45

The other options are not correct because they do not improve the deployment reliability or follow best practices. Creating separate pipeline stages that run a CodeBuild project to build and then test the application is not a good option because it will increase the pipeline execution time and delay the feedback loop. Creating individual Lambda functions that use CodeDeploy instead of Systems Manager to run build, test, and deploy actions is not a valid option because it will add unnecessary complexity and cost to the solution. Lambda functions are not designed for long-running tasks such as building or deploying applications. Creating an Amazon S3 bucket and modifying the CodeBuild project to store the packages in the S3 bucket instead of in CodeArtifact is not a necessary option because it will not affect the deployment reliability.CodeArtifact is a secure, scalable, and cost-effective package management service that can store and share software packages for application development67

References:

1: What is AWS CodePipeline? - AWS CodePipeline

2: Create a pipeline in AWS CodePipeline - AWS CodePipeline

3: Deploy an application with AWS CodeDeploy - AWS CodePipeline

4: What is AWS CodeDeploy? - AWS CodeDeploy

5: Configure an Application Load Balancer for your blue/green deployments - AWS CodeDeploy

6: What is AWS Lambda? - AWS Lambda

7: What is AWS CodeArtifact? - AWS CodeArtifact

Question 7

Question Type: MultipleChoice

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

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

Which storage solution will meet these requirements?

Options:

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

B- Use Amazon Elastic Block Store (Amazon EBS) for the application storage. Create a backup plan in AWS Backup that creates snapshots of the EBS volumes that are in the primary Region and replicates the snapshots to the DR Region.

C- Use a Volume Gateway in AWS Storage Gateway for the application storage. Configure Cross-Region Replication (CRR) of the Volume Gateway from the primary Region to the DR Region.

D- Use Amazon FSx for NetApp ONTAP for the application storage. Create an FSx for ONTAP instance in the DR Region. Configure NetApp SnapMirror replication from the primary Region to the DR Region.

Answer:

D

Explanation:

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

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

application storage is not a valid option because Volume Gateway does not support SMB protocol. Volume Gateway only supports iSCSI protocol, which means that only Linux applications can access the shared storage.

References:

1: What is Amazon FSx for NetApp ONTAP? - FSx for ONTAP

2: Amazon FSx for NetApp ONTAP

3: Amazon FSx for NetApp ONTAP | NetApp

4: AWS Announces General Availability of Amazon FSx for NetApp ONTAP

: Replicating Data with NetApp SnapMirror - FSx for ONTAP

: What Is Amazon S3? - Amazon Simple Storage Service

: What Is Amazon Elastic Block Store (Amazon EBS)? - Amazon Elastic Compute Cloud

: What Is AWS Storage Gateway? - AWS Storage Gateway

Question 8

Question Type: **MultipleChoice**

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?

Options:

A- Add the resource to the CloudFormation stack that creates the VPCs.

B- Create an organization in AWS Organizations. Add the company's AWS account to the organization. Create an SCP to prevent users from modifying VPC flow logs.

C- Turn on AWS Config. Create an AWS Config rule to check whether VPC flow logs are turned on. Configure automatic remediation to turn on VPC flow logs.

D- Create an IAM policy to deny the use of API calls for VPC flow logs. 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 isvpc-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

Question 9

Question Type: **MultipleChoice**

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

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

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

Options:

A- Create an IAM role that has the appropriate permissions for S3 buckets. Add the IAM role to an instance profile.

B- Update the launch template to include the IAM instance profile.

C- Create an IAM user that has the appropriate permissions for Amazon S3. Generate a secret key and token.

D- Create a trust anchor and profile. Attach the IAM role to the profile.

E- Update the launch template. Modify the user data to use the new secret key and token.

Answer:

A, B

Explanation:

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

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

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

The other options are not correct because they do not meet the requirements or follow best practices. Creating an IAM user and generating a secret key and token is not a good option because it involves managing long-term credentials that need to be rotated regularly. Moreover, embedding credentials in user data is not secure because user data is visible to anyone who can describe the EC2 instance. Creating a trust anchor and profile is not a valid option because trust anchors are used for certificate-based authentication, not for IAM roles or instance profiles. Modifying user data to use a new secret key and token is also not a good option because it requires updating user data every time the credentials change, which is not scalable or efficient.

References:

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

2: DevOps Resources - Amazon Web Services (AWS)

3: Exam Readiness: AWS Certified DevOps Engineer - Professional

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

: Working with Instance Profiles - AWS Identity and Access Management

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

: Temporary Security Credentials - AWS Identity and Access Management

Question 10

Question Type: **MultipleChoice**

A company's development team uses AVMS Cloud Formation to deploy its application resources The team must use for an changes to the environment The team cannot use AWS Management Console or the AWS CLI to make manual changes directly.

The team uses a developer IAM role to access the environment The role is configured with the Admnistratoraccess managed policy. The company has created a new Cloudformationdeployment IAM role that has the following policy.

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "elasticloadbalancing:*",
                "lambda:*",
                "dynamodb:*"
            ],
            "Resource": "*"
        }
    ]
}
```

The company wants ensure that only CloudFormation can use the new role. The development team cannot make any manual changes to the deployed resources.

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

Options:

A- Remove the AdministratorAccess policy. Assign the ReadOnlyAccess managed IAM policy to the developer role. Instruct the developers to use the CloudFormationDeployment role as a CloudFormation service role when the developers deploy new stacks.

B- Update the trust of CloudFormationDeployment role to allow the developer IAM role to assume the CloudFormationDepoyment role.

C- Configure the IAM to be to get and pass the CloudFormationDeployment role if
cloudformation actions for resources,

D- Update the trust Of the CloudFormationDepoyment role to anow the cloudformation.amazonaws.com AWS principal to perform the iam:AssumeR01e action

E- Remove me Administratoraccess policy. Assign the ReadOnly/Access managed IAM policy to the developer role Instruct the developers to assume the CloudFormatondeployment role when the developers new stacks

F- Add an IAM policy to CloudFormationDeplyment to allow cloudformation * on an Add a policy that allows the iam.PassR01e action for ARN of if iam PassedT0Service equal cloudformation.amazonaws.com

Answer:

A, D, F

Explanation:

A comprehensive and detailed explanation is:

Option A is correct because removing the AdministratorAccess policy and assigning the ReadOnlyAccess managed IAM policy to the developer role is a valid way to prevent the developers from making any manual changes to the deployed resources. The AdministratorAccess policy grants full access to all AWS resources and actions, which is not necessary for the developers. The ReadOnlyAccess policy grants read-only access to most AWS resources and actions, which is sufficient for the developers to view the status of their stacks. Instructing the developers to use the CloudFormationDeployment role as a CloudFormation service role when they

deploy new stacks is also a valid way to ensure that only CloudFormation can use the new role. A CloudFormation service role is an IAM role that allows CloudFormation to make calls to resources in a stack on behalf of the user1. The user can specify a service role when they create or update a stack, and CloudFormation will use that role's credentials for all operations that are performed on that stack1.

Option B is incorrect because updating the trust of CloudFormationDeployment role to allow the developer IAM role to assume the CloudFormationDeployment role is not a valid solution. This would allow the developers to manually assume the CloudFormationDeployment role and perform actions on the deployed resources, which is not what the company wants. The trust of CloudFormationDeployment role should only allow the cloudformation.amazonaws.com AWS principal to assume the role, as in option D.

Option C is incorrect because configuring the IAM user to be able to get and pass the CloudFormationDeployment role if cloudformation actions for resources is not a valid solution. This would allow the developers to manually pass the CloudFormationDeployment role to other services or resources, which is not what the company wants. The IAM user should only be able to pass the CloudFormationDeployment role as a service role when they create or update a stack with CloudFormation, as in option A.

Option D is correct because updating the trust of CloudFormationDeployment role to allow the cloudformation.amazonaws.com AWS principal to perform the iam:AssumeRole action is a valid solution. This allows CloudFormation to assume the CloudFormationDeployment role and access resources in other services on behalf of the user2. The trust policy of an IAM role defines which entities can assume the role2. By specifying cloudformation.amazonaws.com as the principal, you grant permission only to CloudFormation to assume this role.

Option E is incorrect because instructing the developers to assume the CloudFormationDeployment role when they deploy new stacks is not a valid solution. This would allow the developers to manually assume the CloudFormationDeployment role and perform actions on the deployed resources, which is not what the company wants. The developers should only use the CloudFormationDeployment role as a service role when they deploy new stacks with CloudFormation, as in option A.

Option F is correct because adding an IAM policy to CloudFormationDeployment that allows cloudformation:* on all resources and adding a policy that allows the iam:PassRole action for ARN of CloudFormationDeployment if iam:PassedToService equals cloudformation.amazonaws.com are valid solutions. The first policy grants permission for CloudFormationDeployment to perform any action with any resource using cloudformation.amazonaws.com as a service principal3. The second policy grants permission for passing this role only if it is passed by cloudformation.amazonaws.com as a service principal4. This ensures that only CloudFormation can use this role.

References:

1: AWS CloudFormation service roles

2: How to use trust policies with IAM roles

3: AWS::IAM::Policy

4: IAM: Pass an IAM role to a specific AWS service