Free Questions for DOP-C02

Shared by Mendez on 09-08-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

A company's organization in AWS Organizations has a single OU. The company runs Amazon EC2 instances in the OU accounts. The company needs to limit the use of each EC2 instance's credentials to the specific EC2 instance that the credential is assigned to. A DevOps engineer must configure security for the EC2 instances.

Which solution will meet these requirements?

Options:

A- Create an SCP that specifies the VPC CIDR block. Configure the SCP to check whether the value of the aws:VpcSourceIp condition key is in the specified block. In the same SCP check, check whether the values of the aws:EC2InstanceSourcePrivateIPv4 and aws:SourceVpc condition keys are the same. Deny access if either condition is false. Apply the SCP to the OU.
B- Create an SCP that checks whether the values of the aws:EC2InstanceSourceVPC and aws:SourceVpc condition keys are the same. Deny access if the values are not the same. In the same SCP check, check whether the values of the aws:EC2InstanceSourcePrivateIPv4 and awsVpcSourceIp condition keys are the same. Deny access if the values are not the same. Apply the SCP to the OU.
C- Create an SCP that includes a list of acceptable VPC values and checks whether the value of the aws:SourceVpc condition key is in the list. In the same SCP check, define a list of acceptable IP address values and check whether the value of the aws:VpcSourceIp condition key is in the list. Deny access if either condition is false. Apply the SCP to each account in the organization.
D- Create an SCP that checks whether the values of the aws:EC2InstanceSourceVPC and aws:VpcSourceIp condition keys are the same. Deny access if the values are not the same. In the same SCP check, check whether the values of the aws:EC2InstanceSourcePrivatoIPv4 and aws:SourceVpc condition keys are the same. Deny access if the values are not the same. Apply the SCP to each account in the organization.
Step 1: Using Service Control Policies (SCPs) for EC2 Security
To limit the use of EC2 instance credentials to the specific EC2 instance they are assigned to, you can create a Service Control Policy (SCP) that verifies specific conditions, such as whether the EC2 instance's source VPC and private IP match expected values.
Action: Create an SCP that checks whether the values of the aws:EC2InstanceSourceVPC and aws:SourceVpc condition keys are the same. Deny access if they are not.
Why: This ensures that credentials cannot be used outside the designated EC2 instance or VPC.
Step 2: Further Validation with Private IPs
The SCP should also verify that the EC2 instance's private IP matches the IP range specified for the VPC. If the instance's private IP does not match, access should be denied.
Action: In the same SCP, check whether the values of the aws:EC2InstanceSourcePrivateIP and aws:VpcSourceIP condition keys are the same. Deny access if they are not.
Why: This ensures that the credentials are only used within the specific EC2 instance and its

associated VPC.

Answer:

B

Explanation:

This corresponds to Option B: Create an SCP that checks whether the values of the aws:EC2InstanceSourceVPC and aws:SourceVpc condition keys are the same. Deny access if the values are not the same. In the same SCP check, check whether the values of the aws:EC2InstanceSourcePrivateIP and aws:VpcSourceIP condition keys are the same. Deny access if the values are not the same. Apply the SCP to the OU.

Question 2

Question Type: MultipleChoice

A company uses containers for its applications The company learns that some container Images are missing required security configurations

A DevOps engineer needs to implement a solution to create a standard base image The solution must publish the base image weekly to the us-west-2 Region, us-east-2 Region, and eu-central-1 Region.

Which solution will meet these requirements?

Options:

A- Create an EC2 Image Builder pipeline that uses a container recipe to build the image. Configure the pipeline to distribute the image to an Amazon Elastic Container Registry (Amazon ECR) repository in us-west-2. Configure ECR replication from us-west-2 to us-east-2 and from us-east-2 to eu-central-1 Configure the pipeline to run weekly

B- Create an AWS CodePipeline pipeline that uses an AWS CodeBuild project to build the image Use AWS CodeOeploy to publish the image to an Amazon Elastic Container Registry (Amazon ECR) repository in us-west-2 Configure ECR replication from us-west-2 to us-east-2 and from us-east-2 to eu-central-1 Configure the pipeline to run weekly

C- Create an EC2 Image Builder pipeline that uses a container recipe to build the Image Configure the pipeline to distribute the image to Amazon Elastic Container Registry (Amazon ECR) repositories in all three Regions. Configure the pipeline to run weekly.

D- Create an AWS CodePipeline pipeline that uses an AWS CodeBuild project to build the image

Use AWS CodeDeploy to publish the image to Amazon Elastic Container Registry (Amazon ECR) repositories in all three Regions. Configure the pipeline to run weekly.

Answer:

C

Explanation:

Create an EC2 Image Builder Pipeline that Uses a Container Recipe to Build the Image:

EC2 Image Builder simplifies the creation, maintenance, validation, and sharing of container images.

By using a container recipe, you can define the base image, components, and validation tests for your container image.

Configure the Pipeline to Distribute the Image to Amazon Elastic Container Registry (Amazon ECR) Repositories in All Three Regions:

Amazon ECR provides a secure, scalable, and reliable container registry.

Configuring the pipeline to distribute the image to ECR repositories in us-west-2, us-east-2, and eu-central-1 ensures that the image is available in all required regions.

Configure the Pipeline to Run Weekly:

Setting the pipeline to run on a weekly schedule ensures that the base image is regularly updated and published, incorporating any new security configurations or updates.

By using EC2 Image Builder to automate the creation and distribution of the container image, the solution ensures that the base image is consistently maintained and available across multiple regions with minimal management overhead.

EC2 Image Builder

Amazon ECR

Setting Up EC2 Image Builder Pipelines

Question 3

Question Type: MultipleChoice

A company uses an organization in AWS Organizations to manage its AWS accounts. The company recently acquired another company that has standalone AWS accounts. The acquiring

company's DevOps team needs to consolidate the administration of the AWS accounts for both companies and retain full administrative control of the accounts. The DevOps team also needs to collect and group findings across all the accounts to implement and maintain a security posture.

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

Options:

A- Invite the acquired company's AWS accounts to join the organization. Create an SCP that has full administrative privileges. Attach the SCP to the management account.
B- Invite the acquired company's AWS accounts to join the organization. Create the OrganizationAccountAccessRole 1AM role in the invited accounts. Grant permission to the management account to assume the role.
C- Use AWS Security Hub to collect and group findings across all accounts. Use Security Hub to automatically detect new accounts as the accounts are added to the organization.
D- Use AWS Firewall Manager to collect and group findings across all accounts. Enable all features for the organization. Designate an account in the organization as the delegated administrator account for Firewall Manager.
E- Use Amazon Inspector to collect and group findings across all accounts. Designate an account in the organization as the delegated administrator account for Amazon Inspector.

Answer:

B, C

Explanation:

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

Amazon Inspector is used to assess the security and compliance of applications running on Amazon EC2 instances, not to collect and group security findings across accounts.Reference:

Inviting an AWS account to join your organization

Enabling and disabling AWS Security Hub

Service control policies

AWS Firewall Manager

Amazon Inspector

Question 4

Question Type: MultipleChoice

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

the applications can access the data.

Which solution will meet these requirements?

Options:

A- Create an S3 bucket for each application. Configure S3 Same-Region Replication (SRR) from the raw data's S3 bucket to each application's S3 bucket. Configure each application to consume data from its own S3 bucket.

B- Create an Amazon Kinesis data stream. Create an AWS Lambda function that is invoked by object creation events in the raw data's S3 bucket. Program the Lambda function to redact data for each application. Publish the data on the Kinesis data stream. Configure each application to consume data from the Kinesis data stream.

C- For each application, create an S3 access point that uses the raw data's S3 bucket as the destination. Create an AWS Lambda function that is invoked by object creation events in the raw data's S3 bucket. Program the Lambda function to redact data for each application. Store the data in each application's S3 access point. Configure each application to consume data from its own S3 access point.

D- Create an S3 access point that uses the raw data's S3 bucket as the destination. For each application, create an S3 Object Lambda access point that uses the S3 access point. Configure the AWS Lambda function for each S3 Object Lambda access point to redact data when objects are retrieved. Configure each application to consume data from its own S3 Object Lambda access point.

Answer:

D

Explanation:

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

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

Question 5

Question Type: MultipleChoice

A company's DevOps team manages a set of AWS accounts that are in an organization in AWS Organizations

The company needs a solution that ensures that all Amazon EC2 instances use approved AMIs that the DevOps team manages. The solution also must remediate the usage of AMIs that are not approved The individual account administrators must not be able to remove the restriction to use approved AMIs.

Which solution will meet these requirements?

Options:

A- Use AWS CloudFormation StackSets to deploy an Amazon EventBridge rule to each account. Configure the rule to react to AWS CloudTrail events for Amazon EC2 and to send a notification to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the DevOps team to the SNS topic

B- Use AWS CloudFormation StackSets to deploy the approved-amis-by-id AWS Config managed rule to each account. Configure the rule with the list of approved AMIs. Configure the rule to run the the AWS-StopEC2Instance AWS Systems Manager Automation runbook for the noncompliant EC2 instances.

C- Create an AWS Lambda function that processes AWS CloudTrail events for Amazon EC2 Configure the Lambda function to send a notification to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the DevOps team to the SNS topic. Deploy the Lambda function in

each account in the organization Create an Amazon EventBridge rule in each account Configure the EventBridge rules to react to AWS CloudTrail events for Amazon EC2 and to invoke the Lambda function.
D- Enable AWS Config across the organization Create a conformance pack that uses the approved -amis-by-id AWS Config managed rule with the list of approved AMIs. Deploy the conformance pack across the organization. Configure the rule to run the AWS-StopEC2Instance AWS Systems Manager Automation runbook for the noncompliant EC2 instances.

Answer:

D

Explanation:

Enable AWS Config Across the Organization:

AWS Config provides a detailed view of the configuration of AWS resources in your AWS account. It can be used to assess, audit, and evaluate the configurations of your resources.

Enabling AWS Config across the organization ensures that all accounts are monitored for compliance.

Create a Conformance Pack Using the approved-amis-by-id AWS Config Managed Rule:

A conformance pack is a collection of AWS Config rules and remediation actions that can be easily deployed across an organization.

The approved-amis-by-id managed rule checks whether running instances are using approved AMIs.

Deploy the Conformance Pack Across the Organization:

Deploying the conformance pack across the organization ensures that all accounts adhere to the policy of using only approved AMIs.

The conformance pack can be deployed via the AWS Management Console, CLI, or SDKs.

Configure the Rule to Run the AWS-StopEC2Instance AWS Systems Manager Automation Runbook for Non-Compliant EC2 Instances:

The AWS-StopEC2Instance runbook can be configured to automatically stop any EC2 instances that are found to be non-compliant (i.e., not using approved AMIs).

This remediation action ensures that any unauthorized instances are promptly stopped, enforcing the policy without manual intervention.

By following these steps, the solution ensures that all EC2 instances across the organization use approved AMIs, and any non-compliant instances are remediated automatically.

AWS Config Conformance Packs

AWS Config Managed Rules

AWS Systems Manager Automation Runbooks

Question 6

Question Type: MultipleChoice

A company has a new AWS account that teams will use to deploy various applications. The teams will create many Amazon S3 buckets for application- specific purposes and to store AWS CloudTrail logs. The company has enabled Amazon Macie for the account.

A DevOps engineer needs to optimize the Macie costs for the account without compromising the account's functionality.

Which solutions will meet these requirements? (Select TWO.)

Options:

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

Answer:

A, D

Explanation:

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

Question 7

Question Type: MultipleChoice

A company has microservices running in AWS Lambda that read data from Amazon DynamoDB. The Lambda code is manually deployed by developers after successful testing The company now needs the tests and deployments be automated and run in the cloud Additionally, traffic to the new versions of each microservice should be incrementally shifted over time after deployment.

What solution meets all the requirements, ensuring the MOST developer velocity?

Options:

A- Create an AWS CodePipelme configuration and set up a post-commit hook to trigger the pipeline after tests have passed Use AWS CodeDeploy and create a Canary deployment configuration that specifies the percentage of traffic and interval
B- Create an AWS CodeBuild configuration that triggers when the test code is pushed Use AWS CloudFormation to trigger an AWS CodePipelme configuration that deploys the new Lambda versions and specifies the traffic shift percentage and interval
C- Create an AWS CodePipelme configuration and set up the source code step to trigger when code is pushed. Set up the build step to use AWS CodeBuild to run the tests Set up an AWS CodeDeploy configuration to deploy, then select the CodeDeployDefault.LambdaLinearlDPercentEvery3Minut.es Option.
D- Use the AWS CLI to set up a post-commit hook that uploads the code to an Amazon S3 bucket after tests have passed. Set up an S3 event trigger that runs a Lambda function that deploys the new version. Use an interval in the Lambda function to deploy the code over time at the required percentage

Answer:

C

Explanation:

https://docs.aws.amazon.com/codedeploy/latest/userguide/deployment-configurations.html

Question 8

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 9

Question Type: MultipleChoice

A company has an AWS Control Tower landing zone. The company's DevOps team creates a workload OU. A development OU and a production OU are nested under the workload OU. The company grants users full access to the company's AWS accounts to deploy applications.

The DevOps team needs to allow only a specific management 1AM role to manage the 1AM roles and policies of any AWS accounts In only the production OU.

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

Options:

A- Create an SCP that denies full access with a condition to exclude the management 1AM role for the organization root.
B- Ensure that the FullAWSAccess SCP is applied at the organization root
C- Create an SCP that allows IAM related actions Attach the SCP to the development OU
D- Create an SCP that denies IAM related actions with a condition to exclude the management I AM role Attach the SCP to the workload OU
E- Create an SCP that denies IAM related actions with a condition to exclude the management 1AM role Attach the SCP to the production OU

Answer:

B, E

Explanation:

You need to understand how SCP inheritance works in AWS. The way it works for Deny policies is different that allow policies.

Allow polices are passing down to children ONLY if they don't have an allow policy.

Deny policies always pass down to children.

That's why there is always an SCP set to the Root to allow everything by default. If you limit this policy, the whole organization will be limited, not matter what other policies are saying for the other OUs. So it's not A. It's not D because it restricts the wrong OU.

Question 10

Question Type: MultipleChoice

A company is reviewing its 1AM policies. One policy written by the DevOps engineer has been (lagged as too permissive. The policy is used by an AWS Lambda function that issues a stop command to Amazon EC2 instances tagged with Environment: NonProduccion over the weekend. The current policy is:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "ec2:*",
      "Resource": "*"
    }
  ]
}
```

What changes should the engineer make to achieve a policy ot least permission? (Select THREE.)

A.

A. Add the following conditional expression:

```
"Condition": {
  "StringEquals": {
    "aws:principaltype": "lambda.amazonaws.com"
  }
}
```

B.

Change `"Resource": "*"` to `"Resource": "arn:aws:ec2:*:*:instance/*"`

C.

Add the following conditional expression:

```
"Condition": {
  "StringNotEquals": {
    "ec2:ResourceTag/Environment": "Production"
  }
}
```

D.

Add the following conditional expression:

```
"Condition": {
  "StringEquals": {
    "ec2:ResourceTag/Environment": "NonProduction"
  }
}
```

E.

Change `"Action": "ec2:*"` to `"Action": "ec2:StopInstances"`

F.

Add the following conditional expression:

```
"Condition" :  {
  "DateGreaterThan" : {
    "aws:CurrentTime" : "$ ;{aws:DateTime:Friday}"
  },
  "DateLessThan": {
    "aws:CurrentTime" : "$ ;{aws:DateTime:Monday}"
  }
}
```

Options:

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

Answer:

A, B, D

Explanation:

The engineer should make the following changes to achieve a policy of least permission:

A:Add a condition to ensure that the principal making the request is an AWS Lambda function. This ensures that only Lambda functions can execute this policy.

B:Narrow down the resources by specifying the ARN of EC2 instances instead of allowing all resources. This ensures that the policy only affects EC2 instances.

D:Add a condition to ensure that this policy only applies to EC2 instances tagged with ''Environment: NonProduction''. This ensures that production environments are not affected by this policy.

AWS Identity and Access Management (IAM) - AWS Documentation

Certified DevOps Engineer - Professional (DOP-C02) Study Guide(page 179)

Question 11

Question Type: MultipleChoice

A company is examining its disaster recovery capability and wants the ability to switch over its daily operations to a secondary AWS Region. The company uses AWS CodeCommit as a source control tool in the primary Region.

A DevOps engineer must provide the capability for the company to develop code in the secondary Region. If the company needs to use the secondary Region, developers can add an additional remote URL to their local Git configuration.

Which solution will meet these requirements?

Options:

A- Create a CodeCommit repository in the secondary Region. Create an AWS CodeBuild project to perform a Git mirror operation of the primary Region's CodeCommit repository to the secondary Region's CodeCommit repository. Create an AWS Lambda function that invokes the CodeBuild project. Create an Amazon EventBridge rule that reacts to merge events in the primary Region's CodeCommit repository. Configure the EventBridge rule to invoke the Lambda function.
B- Create an Amazon S3 bucket in the secondary Region. Create an AWS Fargate task to perform a Git mirror operation of the primary Region's CodeCommit repository and copy the result to the S3 bucket. Create an AWS Lambda function that initiates the Fargate task. Create an Amazon EventBridge rule that reacts to merge events in the CodeCommit repository. Configure the EventBridge rule to invoke the Lambda function.
C- Create an AWS CodeArtifact repository in the secondary Region. Create an AWS CodePipeline pipeline that uses the primary Region's CodeCommit repository for the source action. Create a Cross-Region stage in the pipeline that packages the CodeCommit repository contents and stores the contents in the CodeArtifact repository when a pull request is merged into the CodeCommit repository.
D- Create an AWS Cloud9 environment and a CodeCommit repository in the secondary Region. Configure the primary Region's CodeCommit repository as a remote repository in the AWS Cloud9 environment. Connect the secondary Region's CodeCommit repository to the AWS Cloud9 environment.

Answer:

A

Explanation:

The best solution to meet the disaster recovery capability and allow developers to switch over to a secondary AWS Region for code development is option A. This involves creating aCodeCommit repositoryin the secondary Region and setting up anAWS CodeBuild projectto perform a Git mirror operation of the primary Region's CodeCommit repository to the secondary Region's repository. AnAWS Lambda functionis then created to invoke the CodeBuild project.Additionally,

anAmazon EventBridge ruleis configured to react to merge events in the primary Region's CodeCommit repository and invoke the Lambda function12.This setup ensures that the secondary Region's repository is always up-to-date with the primary repository, allowing for a seamless transition in case of a disaster recovery event1.

AWS CodeCommit User Guide on resilience and disaster recovery1.

AWS Documentation on monitoring CodeCommit events in Amazon EventBridge and Amazon CloudWatch Events2.