



# Amazon-Web-Services

## Exam Questions SCS-C02

AWS Certified Security - Specialty

#### NEW QUESTION 1

A company is using Amazon Macie, AWS Firewall Manager, Amazon Inspector, and AWS Shield Advanced in its AWS account. The company wants to receive alerts if a DDoS attack occurs against the account.

Which solution will meet this requirement?

- A. Use Macie to detect an active DDoS even
- B. Create Amazon CloudWatch alarms that respond to Macie findings.
- C. Use Amazon Inspector to review resources and to invoke Amazon CloudWatch alarms for any resources that are vulnerable to DDoS attacks.
- D. Create an Amazon CloudWatch alarm that monitors Firewall Manager metrics for an active DDoS event.
- E. Create an Amazon CloudWatch alarm that monitors Shield Advanced metrics for an active DDoS event.

**Answer: D**

#### Explanation:

This answer is correct because AWS Shield Advanced is a service that provides comprehensive protection against DDoS attacks of any size or duration. It also provides metrics and reports on the DDoS attack vectors, duration, and size. You can create an Amazon CloudWatch alarm that monitors Shield Advanced metrics such as DDoSAttackBitsPerSecond, DDoSAttackPacketsPerSecond, and DDoSAttackRequestsPerSecond to receive alerts if a DDoS attack occurs against your account.

For more information, see [Monitoring AWS Shield Advanced with Amazon CloudWatch and AWS Shield Advanced metrics and alarms](#).

#### NEW QUESTION 2

A company has a relational database workload that runs on Amazon Aurora MySQL. According to new compliance standards the company must rotate all database credentials every 30 days. The company needs a solution that maximizes security and minimizes development effort.

Which solution will meet these requirements?

- A. Store the database credentials in AWS Secrets Manager
- B. Configure automatic credential rotation for every 30 days.
- C. Store the database credentials in AWS Systems Manager Parameter Store
- D. Create an AWS Lambda function to rotate the credentials every 30 days.
- E. Store the database credentials in an environment file or in a configuration file
- F. Modify the credentials every 30 days.
- G. Store the database credentials in an environment file or in a configuration file
- H. Create an AWS Lambda function to rotate the credentials every 30 days.

**Answer: A**

#### Explanation:

To rotate database credentials every 30 days, the most secure and efficient solution is to store the database credentials in AWS Secrets Manager and configure automatic credential rotation for every 30 days. Secrets Manager can handle the rotation of the credentials in both the secret and the database, and it can use AWS KMS to encrypt the credentials. Option B is incorrect because it requires creating a custom Lambda function to rotate the credentials, which is more effort than using Secrets Manager. Option C is incorrect because it stores the database credentials in an environment file or a configuration file, which is less secure than using Secrets Manager. Option D is incorrect because it combines the drawbacks of option B and option C. Verified References:

➤ <https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotating-secrets.html>

➤ [https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotate-secrets\\_turn-on-for-other.html](https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotate-secrets_turn-on-for-other.html)

#### NEW QUESTION 3

A company's security team is building a solution for logging and visualization. The solution will assist the company with the large variety and velocity of data that it receives from IAM across multiple accounts. The security team has enabled IAM CloudTrail and VPC Flow Logs in all of its accounts. In addition, the company has an organization in IAM Organizations and has an IAM Security Hub master account.

The security team wants to use Amazon Detective. However, the security team cannot enable Detective and is unsure why. What must the security team do to enable Detective?

- A. Enable Amazon Macie so that Security Hub will allow Detective to process findings from Macie.
- B. Disable IAM Key Management Service (IAM KMS) encryption on CloudTrail logs in every member account of the organization
- C. Enable Amazon GuardDuty on all member accounts. Try to enable Detective in 48 hours
- D. Ensure that the principal that launches Detective has the organizations ListAccounts permission

**Answer: D**

#### NEW QUESTION 4

A company developed an application by using AWS Lambda, Amazon S3, Amazon Simple Notification Service (Amazon SNS), and Amazon DynamoDB. An external application puts objects into the company's S3 bucket and tags the objects with date and time. A Lambda function periodically pulls data from the company's S3 bucket based on date and time tags and inserts specific values into a DynamoDB table for further processing.

The data includes personally identifiable information (PII). The company must remove data that is older than 30 days from the S3 bucket and the DynamoDB table. Which solution will meet this requirement with the MOST operational efficiency?

- A. Update the Lambda function to add a TTL S3 flag to S3 object
- B. Create an S3 Lifecycle policy to expire objects that are older than 30 days by using the TTL S3 flag.
- C. Create an S3 Lifecycle policy to expire objects that are older than 30 days
- D. Update the Lambda function to add the TTL attribute in the DynamoDB table
- E. Enable TTL on the DynamoDB table to expire entries that are older than 30 days based on the TTL attribute.
- F. Create an S3 Lifecycle policy to expire objects that are older than 30 days and to add all prefixes to the S3 bucket
- G. Update the Lambda function to delete entries that are older than 30 days.
- H. Create an S3 Lifecycle policy to expire objects that are older than 30 days by using object tag
- I. Update the Lambda function to delete entries that are older than 30 days.

**Answer: B**

#### NEW QUESTION 5

A company uses a third-party identity provider and SAML-based SSO for its AWS accounts. After the third-party identity provider renewed an expired signing certificate, users saw the following message when trying to log in:

Error: Response Signature Invalid (Service: AWSSecurityTokenService; Status Code: 400; Error Code: InvalidIdentityToken)

A security engineer needs to provide a solution that corrects the error and minimizes operational overhead.

Which solution meets these requirements?

- A. Upload the third-party signing certificate's new private key to the AWS identity provider entity defined in AWS Identity and Access Management (IAM) by using the AWS Management Console.
- B. Sign the identity provider's metadata file with the new public key
- C. Upload the signature to the AWS identity provider entity defined in AWS Identity and Access Management (IAM) by using the AWS CLI.
- D. Download the updated SAML metadata file from the identity service provider
- E. Update the file in the AWS identity provider entity defined in AWS Identity and Access Management (IAM) by using the AWS CLI.
- F. Configure the AWS identity provider entity defined in AWS Identity and Access Management (IAM) to synchronously fetch the new public key by using the AWS Management Console.

**Answer:** C

#### Explanation:

This answer is correct because downloading the updated SAML metadata file from the identity service provider ensures that AWS has the latest information about the identity provider, including the new public key. Updating the file in the AWS identity provider entity defined in IAM by using the AWS CLI allows AWS to verify the signature of the SAML assertions sent by the identity provider. This solution also minimizes operational overhead because it can be automated with a script or a cron job.

#### NEW QUESTION 6

A Security Engineer creates an Amazon S3 bucket policy that denies access to all users. A few days later, the Security Engineer adds an additional statement to the bucket policy to allow read-only access to one other employee. Even after updating the policy, the employee still receives an access denied message.

What is the likely cause of this access denial?

- A. The ACL in the bucket needs to be updated
- B. The IAM policy does not allow the user to access the bucket
- C. It takes a few minutes for a bucket policy to take effect
- D. The allow permission is being overridden by the deny

**Answer:** D

#### NEW QUESTION 7

A company has several workloads running on AWS. Employees are required to authenticate using on-premises ADFS and SSO to access the AWS Management Console. Developers migrated an existing legacy web application to an Amazon EC2 instance. Employees need to access this application from anywhere on the internet, but currently, there is no authentication system built into the application.

How should the Security Engineer implement employee-only access to this system without changing the application?

- A. Place the application behind an Application Load Balancer (ALB). Use Amazon Cognito as authentication for the ALB.
- B. Define a SAML-based Amazon Cognito user pool and connect it to ADFS.
- C. Implement AWS SSO in the master account and link it to ADFS as an identity provider
- D. Define the EC2 instance as a managed resource, then apply an IAM policy on the resource.
- E. Define an Amazon Cognito identity pool, then install the connector on the Active Directory server
- F. Use the Amazon Cognito SDK on the application instance to authenticate the employees using their Active Directory user names and passwords.
- G. Create an AWS Lambda custom authorizer as the authenticator for a reverse proxy on Amazon EC2. Ensure the security group on Amazon EC2 only allows access from the Lambda function.

**Answer:** A

#### Explanation:

<https://docs.aws.amazon.com/elasticloadbalancing/latest/application/listener-authenticate-users.html>

#### NEW QUESTION 8

A development team is attempting to encrypt and decode a secure string parameter from the IAM Systems Manager Parameter Store using an IAM Key Management Service (IAM KMS) CMK. However, each attempt results in an error message being sent to the development team.

Which CMK-related problems possibly account for the error? (Select two.)

- A. The CMK used in the attempt does not exist.
- B. The CMK used in the attempt needs to be rotated.
- C. The CMK used in the attempt is using the CMK's key ID instead of the CMK ARN.
- D. The CMK used in the attempt is not enabled.
- E. The CMK used in the attempt is using an alias.

**Answer:** AD

#### Explanation:

<https://docs.aws.amazon.com/kms/latest/developerguide/services-parameter-store.html#parameter-store-cmk-fa>

#### NEW QUESTION 9

A company is attempting to conduct forensic analysis on an Amazon EC2 instance, but the company is unable to connect to the instance by using AWS Systems Manager Session Manager. The company has installed AWS Systems Manager Agent (SSM Agent) on the EC2 instance.

The EC2 instance is in a subnet in a VPC that does not have an internet gateway attached. The company has associated a security group with the EC2 instance. The security group does not have inbound or outbound rules. The subnet's network ACL allows all inbound and outbound traffic.

Which combination of actions will allow the company to conduct forensic analysis on the EC2 instance without compromising forensic data? (Select THREE.)

- A. Update the EC2 instance security group to add a rule that allows outbound traffic on port 443 for 0.0.0.0/0.
- B. Update the EC2 instance security group to add a rule that allows inbound traffic on port 443 to the VPC's CIDR range.
- C. Create an EC2 key pair
- D. Associate the key pair with the EC2 instance.
- E. Create a VPC interface endpoint for Systems Manager in the VPC where the EC2 instance is located.
- F. Attach a security group to the VPC interface endpoint
- G. Allow inbound traffic on port 443 to the VPC's CIDR range.
- H. Create a VPC interface endpoint for the EC2 instance in the VPC where the EC2 instance is located.

**Answer:** BCF

#### NEW QUESTION 10

A security engineer recently rotated all IAM access keys in an AWS account. The security engineer then configured AWS Config and enabled the following AWS Config managed rules; mfa-enabled-for-iam-console-access, iam-user-mfa-enabled, access-key-rotated, and iam-user-unused-credentials-check. The security engineer notices that all resources are displaying as noncompliant after the IAM GenerateCredentialReport API operation is invoked. What could be the reason for the noncompliant status?

- A. The IAM credential report was generated within the past 4 hours.
- B. The security engineer does not have the GenerateCredentialReport permission.
- C. The security engineer does not have the GetCredentialReport permission.
- D. The AWS Config rules have a MaximumExecutionFrequency value of 24 hours.

**Answer:** D

#### Explanation:

The correct answer is D. The AWS Config rules have a MaximumExecutionFrequency value of 24 hours. According to the AWS documentation<sup>1</sup>, the MaximumExecutionFrequency parameter specifies the maximum frequency with which AWS Config runs evaluations for a rule. For AWS Config managed rules, this value can be one of the following:

- One\_Hour
- Three\_Hours
- Six\_Hours
- Twelve\_Hours
- TwentyFour\_Hours

If the rule is triggered by configuration changes, it will still run evaluations when AWS Config delivers the configuration snapshot. However, if the rule is triggered periodically, it will not run evaluations more often than the specified frequency.

In this case, the security engineer enabled four AWS Config managed rules that are triggered periodically. Therefore, these rules will only run evaluations every 24 hours, regardless of when the IAM credential report is generated. This means that the resources will display as noncompliant until the next evaluation cycle, which could take up to 24 hours after the IAM access keys are rotated.

The other options are incorrect because:

- A. The IAM credential report can be generated at any time, but it will not affect the compliance status of the resources until the next evaluation cycle of the AWS Config rules.
- B. The security engineer was able to invoke the IAM GenerateCredentialReport API operation, which means they have the GenerateCredentialReport permission. This permission is required to generate a credential report that lists all IAM users in an AWS account and their credential status<sup>2</sup>.
- C. The security engineer does not need the GetCredentialReport permission to enable or evaluate AWS Config rules. This permission is required to retrieve a credential report that was previously generated by using the GenerateCredentialReport operation<sup>2</sup>.

References:

1: AWS::Config::ConfigRule - AWS CloudFormation 2: IAM: Generate and retrieve IAM credential reports

#### NEW QUESTION 10

A company uses AWS Organizations to manage a multi-account AWS environment in a single AWS Region. The organization's management account is named management-01. The company has turned on AWS Config in all accounts in the organization. The company has designated an account named security-01 as the delegated administrator for AWS Config.

All accounts report the compliance status of each account's rules to the AWS Config delegated administrator account by using an AWS Config aggregator. Each account administrator can configure and manage the account's own AWS Config rules to handle each account's unique compliance requirements.

A security engineer needs to implement a solution to automatically deploy a set of 10 AWS Config rules to all existing and future AWS accounts in the organization. The solution must turn on AWS Config automatically during account creation.

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

- A. Create an AWS CloudFormation template that contains the 10 required AWS Config rule
- B. Deploy the template by using CloudFormation StackSets in the security-01 account.
- C. Create a conformance pack that contains the 10 required AWS Config rule
- D. Deploy the conformance pack from the security-01 account.
- E. Create a conformance pack that contains the 10 required AWS Config rule
- F. Deploy the conformance pack from the management-01 account.
- G. Create an AWS CloudFormation template that will activate AWS Config
- H. Deploy the template by using CloudFormation StackSets in the security-01 account.
- I. Create an AWS CloudFormation template that will activate AWS Config
- J. Deploy the template by using CloudFormation StackSets in the management-01 account.

**Answer:** BE

#### NEW QUESTION 14

A company uses AWS Organizations to manage a small number of AWS accounts. However, the company plans to add 1,000 more accounts soon. The company allows only a centralized security team to create IAM roles for all AWS accounts and teams. Application teams submit requests for IAM roles to the security team. The security team has a backlog of IAM role requests and cannot review and provision the IAM roles quickly.

The security team must create a process that will allow application teams to provision their own IAM roles. The process must also limit the scope of IAM roles and prevent privilege escalation.



Which solution will meet these requirements with the LEAST operational overhead?

- A. Create an IAM group for each application team
- B. Associate policies with each IAM group
- C. Provision IAM users for each application team member
- D. Add the new IAM users to the appropriate IAM group by using role-based access control (RBAC).
- E. Delegate application team leads to provision IAM roles for each team
- F. Conduct a quarterly review of the IAM roles the team leads have provisioned
- G. Ensure that the application team leads have the appropriate training to review IAM roles.
- H. Put each AWS account in its own OU
- I. Add an SCP to each OU to grant access to only the AWS services that the teams plan to use
- J. Include conditions in the AWS account of each team.
- K. Create an SCP and a permissions boundary for IAM role
- L. Add the SCP to the root OU so that only roles that have the permissions boundary attached can create any new IAM roles.

**Answer: D**

**Explanation:**

To create a process that will allow application teams to provision their own IAM roles, while limiting the scope of IAM roles and preventing privilege escalation, the following steps are required:

➤ Create a service control policy (SCP) that defines the maximum permissions that can be granted to any IAM role in the organization. An SCP is a type of policy that you can use with AWS Organizations to manage permissions for all accounts in your organization. SCPs restrict permissions for entities in member accounts, including each AWS account root user, IAM users, and roles. For more information, see [Service control policies overview](#).

➤ Create a permissions boundary for IAM roles that matches the SCP. A permissions boundary is an advanced feature for using a managed policy to set the maximum permissions that an identity-based policy can grant to an IAM entity. A permissions boundary allows an entity to perform only the actions that are allowed by both its identity-based policies and its permissions boundaries. For more information, see [Permissions boundaries for IAM entities](#).

➤ Add the SCP to the root organizational unit (OU) so that it applies to all accounts in the organization.

This will ensure that no IAM role can exceed the permissions defined by the SCP, regardless of how it is created or modified.

➤ Instruct the application teams to attach the permissions boundary to any IAM role they create. This will prevent them from creating IAM roles that can escalate their own privileges or access resources they are not authorized to access.

This solution will meet the requirements with the least operational overhead, as it leverages AWS Organizations and IAM features to delegate and limit IAM role creation without requiring manual reviews or approvals.

The other options are incorrect because they either do not allow application teams to provision their own IAM roles (A), do not limit the scope of IAM roles or prevent privilege escalation (B), or do not take advantage of managed services whenever possible (C).

Verified References:

➤ [https://docs.aws.amazon.com/IAM/latest/UserGuide/access\\_policies\\_boundaries.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_boundaries.html)

**NEW QUESTION 19**

A company created an IAM account for its developers to use for testing and learning purposes. Because the account will be shared among multiple teams of developers, the company wants to restrict the ability to stop and terminate Amazon EC2 instances so that a team can perform these actions only on the instances it owns.

Developers were instructed to tag all their instances with a Team tag key and use the team name in the tag value. One of the first teams to use this account is Business Intelligence. A security engineer needs to develop a highly scalable solution for providing developers with access to the appropriate resources within the account. The security engineer has already created individual IAM roles for each team.

Which additional configuration steps should the security engineer take to complete the task?

A. For each team, create an IAM policy similar to the one that follows. Populate the `ec2:ResourceTag/Team` condition key with a proper team name. Attach resulting policies to the corresponding IAM roles.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "NotAction": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "ec2:ResourceTag/Team": "Businessintelligence"
        }
      }
    }
  ]
}
```

- B. For each team create an IAM policy similar to the one that follows. Populate the `IAM:TagKeys/Team` condition key with a proper team name.
- C. Attach the resulting policies to the corresponding IAM roles.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "NotAction": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*",
      "Condition": {
        "ForAnyValue:StringEquals": {
          "aws:TagKeys/Team": "BusinessIntelligence"
        }
      }
    }
  ]
}
```

- D. Tag each IAM role with a Team tag key  
 E. and use the team name in the tag value  
 F. Create an IAM policy similar to the one that follows, and attach it to all the IAM roles used by developers.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "NotAction": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "ec2:ResourceTag/Team": "${aws:PrincipalTag/Team}"
        }
      }
    }
  ]
}
```

- G. Tag each IAM role with the Team key, and use the team name in the tag value  
 H. Create an IAM policy similar to the one that follows, and attach it to all the IAM roles used by developers.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "NotAction": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*",
      "Condition": {
        "ForAnyValue:StringEquals": {
          "aws:TagKeys/Team": "${aws:PrincipalTag/Team}"
        }
      }
    }
  ]
}
```

**Answer: A**

## NEW QUESTION 20

A security engineer is configuring a mechanism to send an alert when three or more failed sign-in attempts to the AWS Management Console occur during a 5-minute period. The security engineer creates a trail in AWS CloudTrail to assist in this work. Which solution will meet these requirements?

- A. In CloudTrail, turn on Insights events on the trail  
 B. Configure an alarm on the insight with eventName matching ConsoleLogin and errorMessage matching "Failed authentication". Configure a threshold of 3 and a period of 5 minutes.  
 C. Configure CloudTrail to send events to Amazon CloudWatch Log  
 D. Create a metric filter for the relevant log group  
 E. Create a filter pattern with eventName matching ConsoleLogin and errorMessage matching "Failed authentication". Create a CloudWatch alarm with a threshold of 3 and a period of 5 minutes.  
 F. Create an Amazon Athena table from the CloudTrail event  
 G. Run a query for eventName matching ConsoleLogin and for errorMessage matching "Failed authentication". Create a notification action from the query to send an Amazon Simple Notification Service (Amazon SNS) notification when the count equals 3 within a period of 5 minutes.  
 H. In AWS Identity and Access Management Access Analyzer, create a new analyzer  
 I. Configure the analyzer to send an Amazon Simple Notification Service (Amazon SNS) notification when a failed sign-in event occurs 3 times for any IAM user within a period of 5 minutes.

**Answer: B**

## Explanation:

The correct answer is B. Configure CloudTrail to send events to Amazon CloudWatch Logs. Create a metric filter for the relevant log group. Create a filter pattern with eventName matching ConsoleLogin and errorMessage matching "Failed authentication". Create a CloudWatch alarm with a threshold of 3 and a period of 5 minutes.

This answer is correct because it meets the requirements of sending an alert when three or more failed sign-in attempts to the AWS Management Console occur during a 5-minute period. By configuring CloudTrail to send events to CloudWatch Logs, the security engineer can create a metric filter that matches the desired pattern of failed sign-in events. Then, by creating a CloudWatch alarm based on the metric filter, the security engineer can set a threshold of 3 and a period of 5

minutes, and choose an action such as sending an email or an Amazon Simple Notification Service (Amazon SNS) message when the alarm is triggered<sup>12</sup>. The other options are incorrect because:

- A. Turning on Insights events on the trail and configuring an alarm on the insight is not a solution, because Insights events are used to analyze unusual activity in management events, such as spikes in API call volume or error rates. Insights events do not capture failed sign-in attempts to the AWS Management Console<sup>3</sup>.
- C. Creating an Amazon Athena table from the CloudTrail events and running a query for failed sign-in events is not a solution, because it does not provide a mechanism to send an alert based on the query results. Amazon Athena is an interactive query service that allows analyzing data in Amazon S3 using standard SQL, but it does not support creating notifications or alarms from queries<sup>4</sup>.
- D. Creating an analyzer in AWS Identity and Access Management Access Analyzer and configuring it to send an Amazon SNS notification when a failed sign-in event occurs 3 times for any IAM user within a period of 5 minutes is not a solution, because IAM Access Analyzer is not a service that monitors sign-in events, but a service that helps identify resources that are shared with external entities. IAM Access Analyzer does not generate findings for failed sign-in attempts to the AWS Management Console<sup>5</sup>.

References:

1: Sending CloudTrail Events to CloudWatch Logs - AWS CloudTrail 2: Creating Alarms Based on Metric Filters - Amazon CloudWatch 3: Analyzing unusual activity in management events - AWS CloudTrail 4: What is Amazon Athena? - Amazon Athena 5: Using AWS Identity and Access Management Access Analyzer - AWS Identity and Access Management

## NEW QUESTION 22

A security engineer is designing an IAM policy to protect AWS API operations. The policy must enforce multi-factor authentication (MFA) for IAM users to access certain services in the AWS production account. Each session must remain valid for only 2 hours. The current version of the IAM policy is as follows:

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

Which combination of conditions must the security engineer add to the IAM policy to meet these requirements? (Select TWO.)

- A. "Bool" : { "aws : Multi FactorAuthPresent": "true" }
- B. "B001" : { "aws : MultiFactorAuthPresent": "false" }
- C. "NumericLessThan" : { "aws : Multi FactorAuthAge" : "7200" }
- D. "NumericGreaterThan" : { "aws : MultiFactorAuthAge" : "7200" }
- E. "NumericLessThan" : { "MaxSessionDuration" : "7200" }

**Answer:** AC

### Explanation:

The correct combination of conditions to add to the IAM policy is A and C. These conditions will ensure that IAM users must use MFA to access certain services in the AWS production account, and that each session will expire after 2 hours.

- Option A: "Bool" : { "aws:MultiFactorAuthPresent" : "true" } is a valid condition that checks if the principal (the IAM user) has authenticated with MFA before making the request. This condition will enforce MFA for the IAM users to access the specified services. This condition key is supported by all AWS services that support IAM policies<sup>1</sup>.
- Option B: "Bool" : { "aws:MultiFactorAuthPresent" : "false" } is the opposite of option A. This condition will allow access only if the principal has not authenticated with MFA, which is not the desired requirement. This condition key is supported by all AWS services that support IAM policies<sup>1</sup>.
- Option C: "NumericLessThan" : { "aws:MultiFactorAuthAge" : "7200" } is a valid condition that checks if the time since the principal authenticated with MFA is less than 7200 seconds (2 hours). This condition will enforce the session duration limit for the IAM users. This condition key is supported by all AWS services that support IAM policies<sup>1</sup>.
- Option D: "NumericGreaterThan" : { "aws:MultiFactorAuthAge" : "7200" } is the opposite of option C. This condition will allow access only if the time since the principal authenticated with MFA is more than 7200 seconds (2 hours), which is not the desired requirement. This condition key is supported by all AWS services that support IAM policies<sup>1</sup>.
- Option E: "NumericLessThan" : { "MaxSessionDuration" : "7200" } is not a valid condition key.

MaxSessionDuration is a property of an IAM role, not a condition key. It specifies the maximum session duration (in seconds) for the role, which can be between 3600 and 43200 seconds (1 to 12 hours). This property can be set when creating or modifying a role, but it cannot be used as a condition in a policy<sup>2</sup>.

## NEW QUESTION 23

Your company is planning on using bastion hosts for administering the servers in IAM. Which of the following is the best description of a bastion host from a security perspective?

Please select:

- A. A Bastion host should be on a private subnet and never a public subnet due to security concerns
- B. A Bastion host sits on the outside of an internal network and is used as a gateway into the private network and is considered the critical strong point of the network
- C. Bastion hosts allow users to log in using RDP or SSH and use that session to SSH into internal network to access private subnet resources.
- D. A Bastion host should maintain extremely tight security and monitoring as it is available to the public



**Answer:** C

**Explanation:**

A bastion host is a special purpose computer on a network specifically designed and configured to withstand attacks. The computer generally hosts a single application, for example a proxy server, and all other services are removed or limited to reduce the threat to the computer.

In IAM, A bastion host is kept on a public subnet. Users log on to the bastion host via SSH or RDP and then use that session to manage other hosts in the private subnets.

Options A and B are invalid because the bastion host needs to sit on the public network. Option D is invalid because bastion hosts are not used for monitoring. For more information on bastion hosts, just browse to the below URL:

<https://docs.IAM.amazon.com/quickstart/latest/linux-bastion/architecture.html>

The correct answer is: Bastion hosts allow users to log in using RDP or SSH and use that session to SSH into internal network to access private subnet resources. Submit your Feedback/Queries to our Experts

**NEW QUESTION 28**

A security team is working on a solution that will use Amazon EventBridge (Amazon CloudWatch Events) to monitor new Amazon S3 objects. The solution will monitor for public access and for changes to any S3 bucket policy or setting that result in public access. The security team configures EventBridge to watch for specific API calls that are logged from AWS CloudTrail. EventBridge has an action to send an email notification through Amazon Simple Notification Service (Amazon SNS) to the security team immediately with details of the API call.

Specifically, the security team wants EventBridge to watch for the s3:PutObjectAcl, s3:DeleteBucketPolicy, and s3:PutBucketPolicy API invocation logs from CloudTrail. While developing the solution in a single account, the security team discovers that the s3:PutObjectAcl API call does not invoke an EventBridge event. However, the s3:DeleteBucketPolicy API call and the s3:PutBucketPolicy API call do invoke an event.

The security team has enabled CloudTrail for AWS management events with a basic configuration in the AWS Region in which EventBridge is being tested.

Verification of the EventBridge event pattern indicates that the pattern is set up correctly. The security team must implement a solution so that the s3:PutObjectAcl API call will invoke an EventBridge event. The solution must not generate false notifications.

Which solution will meet these requirements?

- A. Modify the EventBridge event pattern by selecting Amazon S3. Select All Events as the event type.
- B. Modify the EventBridge event pattern by selecting Amazon S3. Select Bucket Level Operations as the event type.
- C. Enable CloudTrail Insights to identify unusual API activity.
- D. Enable CloudTrail to monitor data events for read and write operations to S3 buckets.

**Answer:** D

**Explanation:**

The correct answer is D. Enable CloudTrail to monitor data events for read and write operations to S3 buckets. According to the AWS documentation<sup>1</sup>, CloudTrail data events are the resource operations performed on or within a resource. These are also known as data plane operations. Data events are often high-volume activities. For example, Amazon S3 object-level API activity (such as GetObject, DeleteObject, and PutObject) is a data event.

By default, trails do not log data events. To record CloudTrail data events, you must explicitly add the

supported resources or resource types for which you want to collect activity. For more information, see Logging data events in the Amazon S3 User Guide<sup>2</sup>.

In this case, the security team wants EventBridge to watch for the s3:PutObjectAcl API invocation logs from CloudTrail. This API uses the acl subresource to set the access control list (ACL) permissions for a new or existing object in an S3 bucket<sup>3</sup>. This is a data event that affects the S3 object resource type. Therefore, the security team must enable CloudTrail to monitor data events for read and write operations to S3 buckets in order to invoke an EventBridge event for this API call.

The other options are incorrect because:

- A. Modifying the EventBridge event pattern by selecting Amazon S3 and All Events as the event type will not capture the s3:PutObjectAcl API call, because this is a data event and not a management event. Management events provide information about management operations that are performed on resources in your AWS account. These are also known as control plane operations<sup>4</sup>.
- B. Modifying the EventBridge event pattern by selecting Amazon S3 and Bucket Level Operations as the event type will not capture the s3:PutObjectAcl API call, because this is a data event that affects the S3 object resource type and not the S3 bucket resource type. Bucket level operations are management events that affect the configuration or metadata of an S3 bucket<sup>5</sup>.
- C. Enabling CloudTrail Insights to identify unusual API activity will not help the security team monitor new S3 objects or changes to any S3 bucket policy or setting that result in public access. CloudTrail Insights helps AWS users identify and respond to unusual activity associated with API calls and API error rates by continuously analyzing CloudTrail management events<sup>6</sup>. It does not analyze data events or generate EventBridge events.

References:

1: CloudTrail log event reference - AWS CloudTrail 2: Logging data events - AWS CloudTrail 3: PutObjectAcl - Amazon Simple Storage Service 4: [Logging management events - AWS CloudTrail] 5: [Amazon S3 Event Types - Amazon Simple Storage Service] 6: Logging Insights events for trails - AWS CloudTrail

**NEW QUESTION 33**

A security engineer needs to build a solution to turn IAM CloudTrail back on in multiple IAM Regions in case it is ever turned off.

What is the MOST efficient way to implement this solution?

- A. Use IAM Config with a managed rule to trigger the IAM-EnableCloudTrail remediation.
- B. Create an Amazon EventBridge (Amazon CloudWatch Events) event with a cloudtrail.amazonaws.com event source and a StartLogging event name to trigger an IAM Lambda function to call the StartLogging API.
- C. Create an Amazon CloudWatch alarm with a cloudtrail.amazonaws.com event source and a StopLogging event name to trigger an IAM Lambda function to call the StartLogging API.
- D. Monitor IAM Trusted Advisor to ensure CloudTrail logging is enabled.

**Answer:** B

**NEW QUESTION 36**

A company recently had a security audit in which the auditors identified multiple potential threats. These potential threats can cause usage pattern changes such as DNS access peak, abnormal instance traffic, abnormal network interface traffic, and unusual Amazon S3 API calls. The threats can come from different sources and can occur at any time. The company needs to implement a solution to continuously monitor its system and identify all these incoming threats in near-real time. Which solution will meet these requirements?

- A. Enable AWS CloudTrail logs, VPC flow logs, and DNS log
- B. Use Amazon CloudWatch Logs to manage these logs from a centralized account.
- C. Enable AWS CloudTrail logs, VPC flow logs, and DNS log
- D. Use Amazon Macie to monitor these logs from a centralized account.



- E. Enable Amazon GuardDuty from a centralized account
- F. Use GuardDuty to manage AWS CloudTrail logs, VPC flow logs, and DNS logs.
- G. Enable Amazon Inspector from a centralized account
- H. Use Amazon Inspector to manage AWS CloudTrail logs, VPC flow logs, and DNS logs.

**Answer:** C

**Explanation:**

Q: Which data sources does GuardDuty analyze? GuardDuty analyzes CloudTrail management event logs, CloudTrail S3 data event logs, VPC Flow Logs, DNS query logs, and Amazon EKS audit logs. GuardDuty can also scan EBS volume data for possible malware when GuardDuty Malware Protection is enabled and identifies suspicious behavior indicative of malicious software in EC2 instance or container workloads. The service is optimized to consume large data volumes for near real-time processing of security detections. GuardDuty gives you access to built-in detection techniques developed and optimized for the cloud, which are maintained and continuously improved upon by GuardDuty engineering.

**NEW QUESTION 41**

To meet regulatory requirements, a Security Engineer needs to implement an IAM policy that restricts the use of AWS services to the us-east-1 Region. What policy should the Engineer implement?

A.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "*",
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestedRegion": "us-east-1"
        }
      }
    }
  ]
}
```

B. A computer code with black text Description automatically generated

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "*",
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "ec2:Region": "us-east-1"
        }
      }
    }
  ]
}
```

C. A computer code with black text Description automatically generated

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": "*",
      "Resource": "*",
      "Condition": {
        "StringNotEquals": {
          "aws:RequestedRegion": "us-east-1"
        }
      }
    }
  ]
}
```

D. A computer code with text Description automatically generated

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "NotAction": "*",
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestedRegion": "us-east-1"
        }
      }
    }
  ]
}
```

**Answer:** C

**Explanation:**

[https://docs.aws.amazon.com/IAM/latest/UserGuide/reference\\_policies\\_examples\\_aws\\_deny-requested-region.h](https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_examples_aws_deny-requested-region.h)

#### NEW QUESTION 46

A security engineer is troubleshooting an AWS Lambda function that is named MyLambdaFunction. The function is encountering an error when the function attempts to read the objects in an Amazon S3 bucket that is named DOC-EXAMPLE-BUCKET. The S3 bucket has the following bucket policy:

```
{
  "Effect": "Allow",
  "Principal": {
    "Service": "lambda.amazonaws.com"
  },
  "Action": "s3:GetObject",
  "Resource": "arn:aws:s3:::DOC-EXAMPLE-BUCKET",
  "Condition": {
    "ArnLike": {
      "aws:SourceArn": "arn:aws:lambda:::function:MyLambdaFunction"
    }
  }
}
```

Which change should the security engineer make to the policy to ensure that the Lambda function can read the bucket objects?

- A. Remove the Condition element
- B. Change the Principal element to the following: {"AWS": "arn:aws:::lambda:::function:MyLambdaFunction"}
- C. Change the Action element to the following: " s3:GetObject\*" " s3:GetBucket\*"
- D. Change the Resource element to "arn:aws:s3:::DOC-EXAMPLE-BUCKET/\*".
- E. Change the Resource element to "arn:aws:lambda:::function:MyLambdaFunction". Change the Principal element to the following: {"Service": "s3.amazonaws.com"}

**Answer:** C

**Explanation:**

The correct answer is C. Change the Resource element to "arn:aws:s3:::DOC-EXAMPLE-BUCKET/\*".

The reason is that the Resource element in the bucket policy specifies which objects in the bucket are affected by the policy. In this case, the policy only applies to the bucket itself, not the objects inside it. Therefore, the Lambda function cannot access the objects with the s3:GetObject permission. To fix this, the Resource element should include a wildcard (\*) to match all objects in the bucket. This way, the policy grants the Lambda function permission to read any object in the bucket.

The other options are incorrect for the following reasons:

- > A. Removing the Condition element would not help, because it only restricts access based on the source IP address of the request. The Principal element should not be changed to the Lambda function ARN, because it specifies who is allowed or denied access by the policy. The policy should allow access to any principal ("\*") and rely on IAM roles or policies to control access to the Lambda function.
- > B. Changing the Action element to include s3:GetBucket\* would not help, because it would grant additional permissions that are not needed by the Lambda function, such as s3:GetBucketAcl or s3:GetBucketPolicy. The s3:GetObject\* permission is sufficient for reading objects in the bucket.
- > D. Changing the Resource element to the Lambda function ARN would not make sense, because it would mean that the policy applies to the Lambda function itself, not the bucket or its objects. The Principal element should not be changed to s3.amazonaws.com, because it would grant access to any AWS service that uses S3, not just Lambda.

#### NEW QUESTION 47

A company deployed Amazon GuardDuty In the us-east-1 Region. The company wants all DNS logs that relate to the company's Amazon EC2 instances to be

inspected. What should a security engineer do to ensure that the EC2 instances are logged?

- A. Use IPv6 addresses that are configured for hostnames.
- B. Configure external DNS resolvers as internal resolvers that are visible only to IAM.
- C. Use IAM DNS resolvers for all EC2 instances.
- D. Configure a third-party DNS resolver with logging for all EC2 instances.

**Answer: C**

**Explanation:**

To ensure that the EC2 instances are logged, the security engineer should do the following:

➤ Use AWS DNS resolvers for all EC2 instances. This allows the security engineer to use Amazon-provided DNS servers that resolve public DNS hostnames to private IP addresses within their VPC, and that log DNS queries in Amazon CloudWatch Logs.

**NEW QUESTION 48**

A company has multiple departments. Each department has its own IAM account. All these accounts belong to the same organization in IAM Organizations. A large .csv file is stored in an Amazon S3 bucket in the sales department's IAM account. The company wants to allow users from the other accounts to access the .csv file's content through the combination of IAM Glue and Amazon Athena. However, the company does not want to allow users from the other accounts to access other files in the same folder.

Which solution will meet these requirements?

- A. Apply a user policy in the other accounts to allow IAM Glue and Athena to access the .csv file.
- B. Use S3 Select to restrict access to the .csv file.
- C. In IAM Glue Data Catalog, use S3 Select as the source of the IAM Glue database.
- D. Define an IAM Glue Data Catalog resource policy in IAM Glue to grant cross-account S3 object access to the .csv file.
- E. Grant IAM Glue access to Amazon S3 in a resource-based policy that specifies the organization as the principal.

**Answer: A**

**NEW QUESTION 51**

A Security Engineer is building a Java application that is running on Amazon EC2. The application communicates with an Amazon RDS instance and authenticates with a user name and password.

Which combination of steps can the Engineer take to protect the credentials and minimize downtime when the credentials are rotated? (Choose two.)

- A. Have a Database Administrator encrypt the credentials and store the ciphertext in Amazon S3. Grant permission to the instance role associated with the EC2 instance to read the object and decrypt the ciphertext.
- B. Configure a scheduled job that updates the credential in AWS Systems Manager Parameter Store and notifies the Engineer that the application needs to be restarted.
- C. Configure automatic rotation of credentials in AWS Secrets Manager.
- D. Store the credential in an encrypted string parameter in AWS Systems Manager Parameter Store.
- E. Grant permission to the instance role associated with the EC2 instance to access the parameter and the AWS KMS key that is used to encrypt it.
- F. Configure the Java application to catch a connection failure and make a call to AWS Secrets Manager to retrieve updated credentials when the password is rotated.
- G. Grant permission to the instance role associated with the EC2 instance to access Secrets Manager.

**Answer: CE**

**Explanation:**

AWS Secrets Manager is a service that helps you manage, retrieve, and rotate secrets such as database credentials, API keys, and other sensitive information. By configuring automatic rotation of credentials in AWS Secrets Manager, you can ensure that your secrets are changed regularly and securely, without requiring manual intervention or application downtime. You can also specify the rotation frequency and the rotation function that performs the logic of changing the credentials on the database and updating the secret in Secrets Manager<sup>1</sup>.

\* E. Configure the Java application to catch a connection failure and make a call to AWS Secrets Manager to retrieve updated credentials when the password is rotated. Grant permission to the instance role associated with the EC2 instance to access Secrets Manager.

By configuring the Java application to catch a connection failure and make a call to AWS Secrets Manager to retrieve updated credentials, you can avoid hard-coding the credentials in your application code or configuration files. This way, your application can dynamically obtain the latest credentials from Secrets Manager whenever the password is rotated, without needing to restart or redeploy the application. To enable this, you need to grant permission to the instance role associated with the EC2 instance to access Secrets Manager using IAM policies<sup>2</sup>. You can also use the AWS SDK for Java to integrate your application with Secrets Manager<sup>3</sup>.

**NEW QUESTION 55**

A company purchased a subscription to a third-party cloud security scanning solution that integrates with AWS Security Hub. A security engineer needs to implement a solution that will remediate the findings from the third-party scanning solution automatically. Which solution will meet this requirement?

- A. Set up an Amazon EventBridge rule that reacts to new Security Hub findings.
- B. Configure an AWS Lambda function as the target for the rule to remediate the findings.
- C. Set up a custom action in Security Hub.
- D. Configure the custom action to call AWS Systems Manager Automation runbooks to remediate the findings.
- E. Set up a custom action in Security Hub.
- F. Configure an AWS Lambda function as the target for the custom action to remediate the findings.
- G. Set up AWS Config rules to use AWS Systems Manager Automation runbooks to remediate the findings.

**Answer: A**

**NEW QUESTION 58**

A business requires a forensic logging solution for hundreds of Docker-based apps running on Amazon EC2. The solution must analyze logs in real time, provide

message replay, and persist logs.

Which Amazon Web Offerings (IAM) services should be employed to satisfy these requirements? (Select two.)

- A. Amazon Athena
- B. Amazon Kinesis
- C. Amazon SQS
- D. Amazon Elasticsearch
- E. Amazon EMR

**Answer:** BD

#### NEW QUESTION 61

A team is using AWS Secrets Manager to store an application database password. Only a limited number of IAM principals within the account can have access to the secret. The principals who require access to the secret change frequently. A security engineer must create a solution that maximizes flexibility and scalability. Which solution will meet these requirements?

- A. Use a role-based approach by creating an IAM role with an inline permissions policy that allows access to the secret
- B. Update the IAM principals in the role trust policy as required.
- C. Deploy a VPC endpoint for Secrets Manager
- D. Create and attach an endpoint policy that specifies the IAM principals that are allowed to access the secret
- E. Update the list of IAM principals as required.
- F. Use a tag-based approach by attaching a resource policy to the secret
- G. Apply tags to the secret and the IAM principal
- H. Use the aws:PrincipalTag and aws:ResourceTag IAM condition keys to control access.
- I. Use a deny-by-default approach by using IAM policies to deny access to the secret explicitly
- J. Attach the policies to an IAM group
- K. Add all IAM principals to the IAM group
- L. Remove principals from the group when they need access
- M. Add the principals to the group again when access is no longer allowed.

**Answer:** C

#### NEW QUESTION 65

A company is using IAM Secrets Manager to store secrets for its production Amazon RDS database. The Security Officer has asked that secrets be rotated every 3 months. Which solution would allow the company to securely rotate the secrets? (Select TWO.)

- A. Place the RDS instance in a public subnet and an IAM Lambda function outside the VPC
- B. Schedule the Lambda function to run every 3 months to rotate the secrets.
- C. Place the RDS instance in a private subnet and an IAM Lambda function inside the VPC in the private subnet
- D. Configure the private subnet to use a NAT gateway
- E. Schedule the Lambda function to run every 3 months to rotate the secrets.
- F. Place the RDS instance in a private subnet and an IAM Lambda function outside the VPC
- G. Configure the private subnet to use an internet gateway
- H. Schedule the Lambda function to run every 3 months to rotate the secrets.
- I. Place the RDS instance in a private subnet and an IAM Lambda function inside the VPC in the private subnet
- J. Schedule the Lambda function to run quarterly to rotate the secrets.
- K. Place the RDS instance in a private subnet and an IAM Lambda function inside the VPC in the private subnet
- L. Configure a Secrets Manager interface endpoint
- M. Schedule the Lambda function to run every 3 months to rotate the secrets.

**Answer:** BE

#### Explanation:

these are the solutions that can securely rotate the secrets for the production RDS database using Secrets Manager. Secrets Manager is a service that helps you manage secrets such as database credentials, API keys, and passwords. You can use Secrets Manager to rotate secrets automatically by using a Lambda function that runs on a schedule. The Lambda function needs to have access to both the RDS instance and the Secrets Manager service. Option B places the RDS instance in a private subnet and the Lambda function in the same VPC in another private subnet. The private subnet with the Lambda function needs to use a NAT gateway to access Secrets Manager over the internet. Option E places the RDS instance and the Lambda function in the same private subnet and configures a Secrets Manager interface endpoint, which is a private connection between the VPC and Secrets Manager. The other options are either insecure or incorrect for rotating secrets using Secrets Manager.

#### NEW QUESTION 69

A company's public Application Load Balancer (ALB) recently experienced a DDoS attack. To mitigate this issue, the company deployed Amazon CloudFront in front of the ALB so that users would not directly access the Amazon EC2 instances behind the ALB.

The company discovers that some traffic is still coming directly into the ALB and is still being handled by the EC2 instances.

Which combination of steps should the company take to ensure that the EC2 instances will receive traffic only from CloudFront? (Choose two.)

- A. Configure CloudFront to add a cache key policy to allow a custom HTTP header that CloudFront sends to the ALB.
- B. Configure CloudFront to add a custom: HTTP header to requests that CloudFront sends to the ALB.
- C. Configure the ALB to forward only requests that contain the custom HTTP header.
- D. Configure the ALB and CloudFront to use the X-Forwarded-For header to check client IP addresses.
- E. Configure the ALB and CloudFront to use the same X.509 certificate that is generated by AWS Certificate Manager (ACM).

**Answer:** BC

#### Explanation:

To prevent users from directly accessing an Application Load Balancer and allow access only through CloudFront, complete these high-level steps: Configure CloudFront to add a custom HTTP header to requests that it sends to the Application Load Balancer. Configure the Application Load Balancer to only forward requests that contain the custom HTTP header. (Optional) Require HTTPS to improve the security of this solution.

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/restrict-access-to-load-balancer.html>



#### NEW QUESTION 74

During a manual review of system logs from an Amazon Linux EC2 instance, a Security Engineer noticed that there are sudo commands that were never properly alerted or reported on the Amazon CloudWatch Logs agent. Why were there no alerts on the sudo commands?

- A. There is a security group blocking outbound port 80 traffic that is preventing the agent from sending the logs
- B. The IAM instance profile on the EC2 instance was not properly configured to allow the CloudWatchLogs agent to push the logs to CloudWatch
- C. CloudWatch Logs status is set to ON versus SECURE, which prevents it from pulling in OS security event logs
- D. The VPC requires that all traffic go through a proxy, and the CloudWatch Logs agent does not support a proxy configuration.

**Answer: B**

#### Explanation:

the reason why there were no alerts on the sudo commands. Sudo commands are commands that allow a user to execute commands as another user, usually the superuser or root. CloudWatch Logs agent is a software agent that can send log data from an EC2 instance to CloudWatch Logs, a service that monitors and stores log data. The CloudWatch Logs agent needs an IAM instance profile, which is a container for an IAM role that allows applications running on an EC2 instance to make API requests to AWS services. If the IAM instance profile on the EC2 instance was not properly configured to allow the CloudWatch Logs agent to push the logs to CloudWatch, then there would be no alerts on the sudo commands. The other options are either irrelevant or invalid for explaining why there were no alerts on the sudo commands.

#### NEW QUESTION 79

A company uses AWS Organizations. The company wants to implement short-term credentials for third-party AWS accounts to use to access accounts within the company's organization. Access is for the AWS Management Console and third-party software-as-a-service (SaaS) applications. Trust must be enhanced to prevent two external accounts from using the same credentials. The solution must require the least possible operational effort. Which solution will meet these requirements?

- A. Use a bearer token authentication with OAuth or SAML to manage and share a central Amazon Cognito user pool across multiple Amazon API Gateway APIs.
- B. Implement AWS IAM Identity Center (AWS Single Sign-On), and use an identity source of choice. Grant access to users and groups from other accounts by using permission sets that are assigned by account.
- C. Create a unique IAM role for each external account
- D. Create a trust policy
- E. Use AWS Secrets Manager to create a random external key.
- F. Create a unique IAM role for each external account
- G. Create a trust policy that includes a condition that uses the sts:ExternalId condition key.

**Answer: D**

#### Explanation:

The correct answer is D.

To implement short-term credentials for third-party AWS accounts, you can use IAM roles and trust policies. A trust policy is a JSON policy document that defines who can assume the role. You can specify the AWS account ID of the third-party account as a principal in the trust policy, and use the sts:ExternalId condition key to enhance the security of the role. The sts:ExternalId condition key is a unique identifier that is agreed upon by both parties and included in the AssumeRole request. This way, you can prevent the "confused deputy" problem, where an unauthorized party can use the same role as a legitimate party.

Option A is incorrect because bearer token authentication with OAuth or SAML is not suitable for granting access to AWS accounts and resources. Amazon Cognito and API Gateway are used for building web and mobile applications that require user authentication and authorization.

Option B is incorrect because AWS IAM Identity Center (AWS Single Sign-On) is a service that simplifies the management of access to multiple AWS accounts and cloud applications for your workforce users. It does not support granting access to third-party AWS accounts.

Option C is incorrect because using AWS Secrets Manager to create a random external key is not necessary and adds operational complexity. You can use the sts:ExternalId condition key instead to provide a unique identifier for each external account.

#### NEW QUESTION 80

A company is using Amazon Route 53 Resolver for its hybrid DNS infrastructure. The company has set up Route 53 Resolver forwarding rules for authoritative domains that are hosted on on-premises DNS servers.

A new security mandate requires the company to implement a solution to log and query DNS traffic that goes to the on-premises DNS servers. The logs must show details of the source IP address of the instance from which the query originated. The logs also must show the DNS name that was requested in Route 53 Resolver. Which solution will meet these requirements?

- A. Use VPC Traffic Mirroring
- B. Configure all relevant elastic network interfaces as the traffic source, include amazon-dns in the mirror filter, and set Amazon CloudWatch Logs as the mirror target
- C. Use CloudWatch Insights on the mirror session logs to run queries on the source IP address and DNS name.
- D. Configure VPC flow logs on all relevant VPC
- E. Send the logs to an Amazon S3 bucket
- F. Use Amazon Athena to run SQL queries on the source IP address and DNS name.
- G. Configure Route 53 Resolver query logging on all relevant VPC
- H. Send the logs to Amazon CloudWatch Log
- I. Use CloudWatch Insights to run queries on the source IP address and DNS name.
- J. Modify the Route 53 Resolver rules on the authoritative domains that forward to the on-premises DNS server
- K. Send the logs to an Amazon S3 bucket
- L. Use Amazon Athena to run SQL queries on the source IP address and DNS name.

**Answer: C**

#### Explanation:

The correct answer is C. Configure Route 53 Resolver query logging on all relevant VPCs. Send the logs to Amazon CloudWatch Logs. Use CloudWatch Insights to run queries on the source IP address and DNS name.

According to the AWS documentation<sup>1</sup>, Route 53 Resolver query logging lets you log the DNS queries that Route 53 Resolver handles for your VPCs. You can send the logs to CloudWatch Logs, Amazon S3, or Kinesis Data Firehose. The logs include information such as the following:

- The AWS Region where the VPC was created

- The ID of the VPC that the query originated from
- The IP address of the instance that the query originated from
- The instance ID of the resource that the query originated from
- The date and time that the query was first made
- The DNS name requested (such as prod.example.com)
- The DNS record type (such as A or AAAA)
- The DNS response code, such as NoError or ServFail
- The DNS response data, such as the IP address that is returned in response to the DNS query

You can use CloudWatch Insights to run queries on your log data and analyze the results using graphs and statistics<sup>2</sup>. You can filter and aggregate the log data based on any field, and use operators and functions to perform calculations and transformations. For example, you can use CloudWatch Insights to find out how many queries were made for a specific domain name, or which instances made the most queries.

Therefore, this solution meets the requirements of logging and querying DNS traffic that goes to the on-premises DNS servers, showing details of the source IP address of the instance from which the query originated, and the DNS name that was requested in Route 53 Resolver.

The other options are incorrect because:

- A. Using VPC Traffic Mirroring would not capture the DNS queries that go to the on-premises DNS servers, because Traffic Mirroring only copies network traffic from an elastic network interface of an EC2 instance to a target for analysis<sup>3</sup>. Traffic Mirroring does not include traffic that goes through a Route 53 Resolver outbound endpoint, which is used to forward queries to on-premises DNS servers<sup>4</sup>. Therefore, this solution would not meet the requirements.
- B. Configuring VPC flow logs on all relevant VPCs would not capture the DNS name that was requested in Route 53 Resolver, because flow logs only record information about the IP traffic going to and from network interfaces in a VPC<sup>5</sup>. Flow logs do not include any information about the content or payload of a packet, such as a DNS query or response. Therefore, this solution would not meet the requirements.
- D. Modifying the Route 53 Resolver rules on the authoritative domains that forward to the on-premises DNS servers would not enable logging of DNS queries, because Resolver rules only specify how to forward queries for specified domain names to your network<sup>6</sup>. Resolver rules do not have any logging functionality by themselves. Therefore, this solution would not meet the requirements. References:

1: Resolver query logging - Amazon Route 53 2: Analyzing log data with CloudWatch Logs Insights - Amazon CloudWatch 3: What is Traffic Mirroring? - Amazon Virtual Private Cloud 4: Outbound Resolver endpoints - Amazon Route 53 5: Logging IP traffic using VPC Flow Logs - Amazon Virtual Private Cloud 6: Managing forwarding rules - Amazon Route 53

### NEW QUESTION 83

There is a requirement for a company to transfer large amounts of data between IAM and an on-premise location. There is an additional requirement for low latency and high consistency traffic to IAM. Given these requirements how would you design a hybrid architecture? Choose the correct answer from the options below

Please select:

- A. Provision a Direct Connect connection to an IAM region using a Direct Connect partner.
- B. Create a VPN tunnel for private connectivity, which increases network consistency and reduces latency.
- C. Create an iPSec tunnel for private connectivity, which increases network consistency and reduces latency.
- D. Create a VPC peering connection between IAM and the Customer gateway.

**Answer: A**

#### Explanation:

IAM Direct Connect makes it easy to establish a dedicated network connection from your premises to IAM. Using IAM Direct Connect you can establish private connectivity between IAM and your datacenter, office, or colocation environment which in many cases can reduce your network costs, increase bandwidth throughput and provide a more consistent network experience than Internet-based connections.

Options B and C are invalid because these options will not reduce network latency Options D is invalid because this is only used to connect 2 VPC's

For more information on IAM direct connect, just browse to the below URL: <https://iam.amazon.com/directconnect>

The correct answer is: Provision a Direct Connect connection to an IAM region using a Direct Connect partner. omit your Feedback/Queries to our Experts

### NEW QUESTION 84

A security engineer needs to configure an Amazon S3 bucket policy to restrict access to an S3 bucket that is named DOC-EXAMPLE-BUCKET. The policy must allow access to only DOC-EXAMPLE-BUCKET from only the following endpoint: vpce-1a2b3c4d. The policy must deny all access to DOC-EXAMPLE-BUCKET if the specified endpoint is not used.

Which bucket policy statement meets these requirements?

- A. A computer code with black text Description automatically generated

```
"Statement": [
  {
    "Sid": "Access-to-specific-VPCE-only",
    "Principal": "*",
    "Action": "s3:*",
    "Effect": "Allow",
    "Resource": ["arn:aws:s3::DOC-EXAMPLE-BUCKET",
                 "arn:aws:s3::DOC-EXAMPLE-BUCKET/*"],
    "Condition": {
      "StringNotEquals": {
        "aws:sourceVpce": "vpce-1a2b3c4d"
      }
    }
  }
]
```

- B. A computer code with black text Description automatically generated

```
"Statement": [
  {
    "Sid": "Access-to-specific-VPCE-only",
    "Principal": "*",
    "Action": "s3:*",
    "Effect": "Deny",
    "Resource": ["arn:aws:s3::DOC-EXAMPLE-BUCKET",
                 "arn:aws:s3::DOC-EXAMPLE-BUCKET/*"],
    "Condition": {
      "StringNotEquals": {
        "aws:sourceVpce": "vpce-1a2b3c4d"
      }
    }
  }
]
```

C. A computer code with black text Description automatically generated

```
"Statement": [
  {
    "Sid": "Access-to-specific-VPCE-only",
    "Principal": "*",
    "Action": "s3:*",
    "Effect": "Deny",
    "Resource": ["arn:aws:s3::DOC-EXAMPLE-BUCKET",
                 "arn:aws:s3::DOC-EXAMPLE-BUCKET/*"],
    "Condition": {
      "StringEquals": {
        "aws:sourceVpce": "vpce-1a2b3c4d"
      }
    }
  }
]
```

D. A computer code with black text Description automatically generated

```
"Statement": [
  {
    "Sid": "Access-to-specific-VPCE-only",
    "Principal": "*",
    "Action": "s3:*",
    "Effect": "Allow",
    "Resource": ["arn:aws:s3::DOC-EXAMPLE-BUCKET",
                 "arn:aws:s3::DOC-EXAMPLE-BUCKET/*"],
    "Condition": {
      "StringEquals": {
        "aws:sourceVpce": "vpce-1a2b3c4d"
      }
    }
  }
]
```

**Answer: B**

**Explanation:**

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/example-bucket-policies-vpc-endpoint.html>

**NEW QUESTION 87**

A company has enabled Amazon GuardDuty in all AWS Regions as part of its security monitoring strategy. In one of its VPCs, the company hosts an Amazon EC2 instance that works as an FTP server. A high number of clients from multiple locations contact the FTP server. GuardDuty identifies this activity as a brute force attack because of the high number of connections that happen every hour.

The company has flagged the finding as a false positive, but GuardDuty continues to raise the issue. A security engineer must improve the signal-to-noise ratio without compromising the company's visibility of potential anomalous behavior.

Which solution will meet these requirements?

- A. Disable the FTP rule in GuardDuty in the Region where the FTP server is deployed.
- B. Add the FTP server to a trusted IP list.
- C. Deploy the list to GuardDuty to stop receiving the notifications.
- D. Create a suppression rule in GuardDuty to filter findings by automatically archiving new findings that match the specified criteria.
- E. Create an AWS Lambda function that has the appropriate permissions to delete the finding whenever a new occurrence is reported.

**Answer: C**

**Explanation:**

"When you create an Amazon GuardDuty filter, you choose specific filter criteria, name the filter and can enable the auto-archiving of findings that the filter matches. This allows you to further tune GuardDuty to your unique environment, without degrading the ability to identify threats. With auto-archive set, all findings are still generated by GuardDuty, so you have a complete and immutable history of all suspicious activity."

**NEW QUESTION 89**

A company needs a forensic-logging solution for hundreds of applications running in Docker on Amazon EC2. The solution must perform real-time analytics on the logs and must support the replay of messages and must persist the logs.

Which IAM services should be used to meet these requirements? (Select TWO)



- A. Amazon Athena
- B. Amazon Kinesis
- C. Amazon SQS
- D. Amazon Elasticsearch
- E. Amazon EMR

**Answer:** BD

**Explanation:**

Amazon Kinesis and Amazon Elasticsearch are both suitable for forensic-logging solutions. Amazon Kinesis can collect, process, and analyze streaming data in real time. Amazon Elasticsearch can store, search, and analyze log data using the popular open-source tool Elasticsearch. The other options are not designed for forensic-logging purposes. Amazon Athena is a query service that can analyze data in S3, Amazon SQS is a message queue service that can decouple and scale microservices, and Amazon EMR is a big data platform that can run Apache Spark and Hadoop clusters.

**NEW QUESTION 93**

A company has an AWS account that includes an Amazon S3 bucket. The S3 bucket uses server-side encryption with AWS KMS keys (SSE-KMS) to encrypt all the objects at rest by using a customer managed key. The S3 bucket does not have a bucket policy.

An IAM role in the same account has an IAM policy that allows s3 List\* and s3 Get\* permissions for the S3 bucket. When the IAM role attempts to access an object in the S3 bucket the role receives an access denied message.

Why does the IAM role not have access to the objects that are in the S3 bucket?

- A. The IAM role does not have permission to use the KMS CreateKey operation.
- B. The S3 bucket lacks a policy that allows access to the customer managed key that encrypts the objects.
- C. The IAM role does not have permission to use the customer managed key that encrypts the objects that are in the S3 bucket.
- D. The ACL of the S3 objects does not allow read access for the objects when the objects are encrypted at rest.

**Answer:** C

**Explanation:**

When using server-side encryption with AWS KMS keys (SSE-KMS), the requester must have both Amazon S3 permissions and AWS KMS permissions to access the objects. The Amazon S3 permissions are for the bucket and object operations, such as s3:ListBucket and s3:GetObject. The AWS KMS permissions are for the key operations, such as kms:GenerateDataKey and kms:Decrypt. In this case, the IAM role has the necessary Amazon S3 permissions, but not the AWS KMS permissions to use the customer managed key that encrypts the objects. Therefore, the IAM role receives an access denied message when trying to access the objects. Verified References:

- > <https://docs.aws.amazon.com/AmazonS3/latest/userguide/troubleshoot-403-errors.html>
- > <https://repost.aws/knowledge-center/s3-access-denied-error-kms>
- > <https://repost.aws/knowledge-center/cross-account-access-denied-error-s3>

**NEW QUESTION 97**

A company is using AWS Organizations to manage multiple AWS accounts for its human resources, finance, software development, and production departments. All the company's developers are part of the software development AWS account.

The company discovers that developers have launched Amazon EC2 instances that were preconfigured with software that the company has not approved for use. The company wants to implement a solution to ensure that developers can launch EC2 instances with only approved software applications and only in the software development AWS account.

Which solution will meet these requirements?

- A. In the software development account, create AMIs of preconfigured instances that include only approved software
- B. Include the AMI IDs in the condition section of an AWS CloudFormation template to launch the appropriate AMI based on the AWS Region
- C. Provide the developers with the CloudFormation template to launch EC2 instances in the software development account.
- D. Create an Amazon EventBridge rule that runs when any EC2 RunInstances API event occurs in the software development account
- E. Specify AWS Systems Manager Run Command as a target of the rule
- F. Configure Run Command to run a script that will install all approved software onto the instances that the developers launch.
- G. Use an AWS Service Catalog portfolio that contains EC2 products with appropriate AMIs that include only approved software
- H. Grant the developers permission to portfolio access only the Service Catalog to launch a product in the software development account.
- I. In the management account, create AMIs of preconfigured instances that include only approved software
- J. Use AWS CloudFormation StackSets to launch the AMIs across any AWS account in the organization
- K. Grant the developers permission to launch the stack sets within the management account.

**Answer:** C

**NEW QUESTION 100**

A company has launched an Amazon EC2 instance with an Amazon Elastic Block Store (Amazon EBS) volume in the us-east-1 Region. The volume is encrypted with an AWS Key Management Service (AWS KMS) customer managed key that the company's security team created. The security team has created an IAM key policy and has assigned the policy to the key. The security team has also created an IAM instance profile and has assigned the profile to the instance. The EC2 instance will not start and transitions from the pending state to the shutting-down state to the terminated state.

Which combination of steps should a security engineer take to troubleshoot this issue? (Select TWO)

- A. Verify that the KMS key policy specifies a deny statement that prevents access to the key by using the aws:SourceIP condition key. Check that the range includes the EC2 instance IP address that is associated with the EBS volume.
- B. Verify that the KMS key that is associated with the EBS volume is set to the Symmetric key type.
- C. Verify that the KMS key that is associated with the EBS volume is in the Enabled state.
- D. Verify that the EC2 role that is associated with the instance profile has the correct IAM instance policy to launch an EC2 instance with the EBS volume.
- E. Verify that the key that is associated with the EBS volume has not expired and needs to be rotated.

**Answer:** CD

**Explanation:**

To troubleshoot the issue of an EC2 instance failing to start and transitioning to a terminated state when it has an EBS volume encrypted with an AWS KMS customer managed key, a security engineer should take the following steps:



\* C. Verify that the KMS key that is associated with the EBS volume is in the Enabled state. If the key is not enabled, it will not function properly and could cause the EC2 instance to fail.

\* D. Verify that the EC2 role that is associated with the instance profile has the correct IAM instance policy to launch an EC2 instance with the EBS volume. If the instance does not have the necessary permissions, it may not be able to mount the volume and could cause the instance to fail.

Therefore, options C and D are the correct answers.

#### NEW QUESTION 102

A company maintains an open-source application that is hosted on a public GitHub repository. While creating a new commit to the repository, an engineer uploaded their IAM access key and secret access key. The engineer reported the mistake to a manager, and the manager immediately disabled the access key. The company needs to assess the impact of the exposed access key. A security engineer must recommend a solution that requires the least possible managerial overhead.

Which solution meets these requirements?

- A. Analyze an IAM Identity and Access Management (IAM) use report from IAM Trusted Advisor to see when the access key was last used.
- B. Analyze Amazon CloudWatch Logs for activity by searching for the access key.
- C. Analyze VPC flow logs for activity by searching for the access key
- D. Analyze a credential report in IAM Identity and Access Management (IAM) to see when the access key was last used.

**Answer:** A

#### Explanation:

To assess the impact of the exposed access key, the security engineer should recommend the following solution:

➤ Analyze an IAM use report from AWS Trusted Advisor to see when the access key was last used. This allows the security engineer to use a tool that provides information about IAM entities and credentials in their account, and check if there was any unauthorized activity with the exposed access key.

#### NEW QUESTION 103

A security engineer needs to create an IAM Key Management Service (IAM KMS) key that will be used to encrypt all data stored in a company's Amazon S3 Buckets in the us-west-1 Region. The key will use server-side encryption. Usage of the key must be limited to requests coming from Amazon S3 within the company's account.

Which statement in the KMS key policy will meet these requirements?

A)

```
{
  "Effect": "Allow",
  "Principal": {
    "AWS": "*"
  },
  "Action": [
    "kms:Encrypt",
    "kms:Decrypt",
    "kms:ReEncrypt*",
    "kms:GenerateDataKey*",
    "kms:DescribeKey"
  ],
  "Resource": "*",
  "Condition": {
    "StringEquals": {
      "kms:ViaService": "s3.us-west-1.amazonaws.com",
      "kms:CallerAccount": "<CustomerAccountID>"
    }
  }
}
```

B)

```
{
  "Effect": "Allow",
  "Principal": {
    "AWS": "s3.us-west-1.amazonaws.com"
  },
  "Action": [
    "kms:Encrypt",
    "kms:Decrypt",
    "kms:ReEncrypt*",
    "kms:GenerateDataKey*",
    "kms:DescribeKey"
  ],
  "Resource": "*",
  "Condition": {
    "StringEquals": {
      "kms:CallerAccount": "<CustomerAccountID>"
    }
  }
}
```

C)

```
{
  "Effect": "Allow",
  "Principal": {
    "AWS": "*"
  },
  "Action": [
    "kms:Encrypt",
    "kms:Decrypt",
    "kms:ReEncrypt*",
    "kms:GenerateDataKey*",
    "kms:DescribeKey"
  ],
  "Resource": "*",
  "Condition": {
    "StringEquals": {
      "kms:EncryptionContext:aws:s3:arn": [
        "arn:aws:s3:::"
      ]
    }
  }
}
```

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

**Answer:** A

#### NEW QUESTION 107

A company is running its workloads in a single AWS Region and uses AWS Organizations. A security engineer must implement a solution to prevent users from launching resources in other Regions.

Which solution will meet these requirements with the LEAST operational overhead?

- A. Create an IAM policy that has an aws RequestedRegion condition that allows actions only in the designated Region Attach the policy to all users.
- B. Create an IAM policy that has an aws RequestedRegion condition that denies actions that are not in the designated Region Attach the policy to the AWS account in AWS Organizations.
- C. Create an IAM policy that has an aws RequestedRegion condition that allows the desired actions Attach the policy only to the users who are in the designated Region.
- D. Create an SCP that has an aws RequestedRegion condition that denies actions that are not in the designated Region
- E. Attach the SCP to the AWS account in AWS Organizations.

**Answer:** D

#### Explanation:

Although you can use a IAM policy to prevent users launching resources in other regions. The best practice is to use SCP when using AWS organizations.  
[https://docs.aws.amazon.com/organizations/latest/userguide/orgs\\_manage\\_policies\\_scps\\_examples\\_general.htm](https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_general.htm)

#### NEW QUESTION 110

A company hosts multiple externally facing applications, each isolated in its own IAM account The company's Security team has enabled IAM WAF, IAM Config, and Amazon GuardDuty on all accounts. The company's Operations team has also joined all of the accounts to IAM Organizations and established centralized logging for CloudTrail, IAM Config, and GuardDuty. The company wants the Security team to take a reactive remediation in one account, and automate implementing this remediation as proactive prevention in all the other accounts.

How should the Security team accomplish this?

- A. Update the IAM WAF rules in the affected account and use IAM Firewall Manager to push updated IAM WAF rules across all other accounts.
- B. Use GuardDuty centralized logging and Amazon SNS to set up alerts to notify all application teams of security incidents.
- C. Use GuardDuty alerts to write an IAM Lambda function that updates all accounts by adding additional NACLs on the Amazon EC2 instances to block known malicious IP addresses.
- D. Use IAM Shield Advanced to identify threats in each individual account and then apply the account-based protections to all other accounts through Organizations.

**Answer:** C

#### NEW QUESTION 111

A company has implemented IAM WAF and Amazon CloudFront for an application. The application runs on Amazon EC2 instances that are part of an Auto Scaling group. The Auto Scaling group is behind an Application Load Balancer (ALB).

The IAM WAF web ACL uses an IAM Managed Rules rule group and is associated with the CloudFront distribution. CloudFront receives the request from IAM WAF and then uses the ALB as the distribution's origin.

During a security review, a security engineer discovers that the infrastructure is susceptible to a large, layer 7 DDoS attack.

How can the security engineer improve the security at the edge of the solution to defend against this type of attack?

- A. Configure the CloudFront distribution to use the Lambda@Edge feature
- B. Create an IAM Lambda function that imposes a rate limit on CloudFront viewer request
- C. Block the request if the rate limit is exceeded.
- D. Configure the IAM WAF web ACL so that the web ACL has more capacity units to process all IAM WAF rules faster.
- E. Configure IAM WAF with a rate-based rule that imposes a rate limit that automatically blocks requests when the rate limit is exceeded.
- F. Configure the CloudFront distribution to use IAM WAF as its origin instead of the ALB.

**Answer:** C

**Explanation:**

To improve the security at the edge of the solution to defend against a large, layer 7 DDoS attack, the security engineer should do the following:

➤ Configure AWS WAF with a rate-based rule that imposes a rate limit that automatically blocks requests when the rate limit is exceeded. This allows the security engineer to use a rule that tracks the number of requests from a single IP address and blocks subsequent requests if they exceed a specified threshold within a specified time period.

**NEW QUESTION 112**

An international company has established a new business entity in South Korea. The company also has established a new AWS account to contain the workload for the South Korean region. The company has set up the workload in the new account in the ap-northeast-2 Region. The workload consists of three Auto Scaling groups of Amazon EC2 instances. All workloads that operate in this Region must keep system logs and application logs for 7 years.

A security engineer must implement a solution to ensure that no logging data is lost for each instance during scaling activities. The solution also must keep the logs for only the required period of 7 years.

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

- A. Ensure that the Amazon CloudWatch agent is installed on all the EC2 instances that the Auto Scaling groups launch
- B. Generate a CloudWatch agent configuration file to forward the required logs to Amazon CloudWatch Logs.
- C. Set the log retention for desired log groups to 7 years.
- D. Attach an IAM role to the launch configuration or launch template that the Auto Scaling groups use. Configure the role to provide the necessary permissions to forward logs to Amazon CloudWatch Logs.
- E. Attach an IAM role to the launch configuration or launch template that the Auto Scaling groups use. Configure the role to provide the necessary permissions to forward logs to Amazon S3.
- F. Ensure that a log forwarding application is installed on all the EC2 instances that the Auto Scaling groups launch
- G. Configure the log forwarding application to periodically bundle the logs and forward the logs to Amazon S3.
- H. Configure an Amazon S3 Lifecycle policy on the target S3 bucket to expire objects after 7 years.

**Answer:** ABC

**Explanation:**

The correct combination of steps that the security engineer should take to meet these requirements are A. Ensure that the Amazon CloudWatch agent is installed on all the EC2 instances that the Auto Scaling groups launch. Generate a CloudWatch agent configuration file to forward the required logs to Amazon CloudWatch Logs., B. Set the log retention for desired log groups to 7 years., and C. Attach an IAM role to the launch configuration or launch template that the Auto Scaling groups use. Configure the role to provide the necessary permissions to forward logs to Amazon CloudWatch Logs.

\* A. This answer is correct because it meets the requirement of ensuring that no logging data is lost for each instance during scaling activities. By installing the CloudWatch agent on all the EC2 instances, the security engineer can collect and send system logs and application logs to CloudWatch Logs, which is a service that stores and monitors log data. By generating a CloudWatch agent configuration file, the security engineer can specify which logs to forward and how often.

\* B. This answer is correct because it meets the requirement of keeping the logs for only the required period of 7 years. By setting the log retention for desired log groups, the security engineer can control how long CloudWatch Logs retains log events before deleting them. The security engineer can choose a predefined retention period of 7 years, or use a custom value.

\* C. This answer is correct because it meets the requirement of providing the necessary permissions to forward logs to CloudWatch Logs. By attaching an IAM role to the launch configuration or launch template that the Auto Scaling groups use, the security engineer can grant permissions to the EC2 instances that are launched by the Auto Scaling groups. By configuring the role to provide the necessary permissions, such as cloudwatch:PutLogEvents and cloudwatch:CreateLogStream, the security engineer can allow the EC2 instances to send log data to CloudWatch Logs.

**NEW QUESTION 113**

A company is deploying an Amazon EC2-based application. The application will include a custom health-checking component that produces health status data in JSON format. A Security Engineer must

implement a secure solution to monitor application availability in near-real time by analyzing the health status data.

Which approach should the Security Engineer use?

- A. Use Amazon CloudWatch monitoring to capture Amazon EC2 and networking metrics Visualizemetrics using Amazon CloudWatch dashboards.
- B. Run the Amazon Kinesis Agent to write the status data to Amazon Kinesis Data Firehose Store the streaming data from Kinesis Data Firehose in Amazon Redshif
- C. (hen run a script on the pool data and analyze the data in Amazon Redshift
- D. Write the status data directly to a public Amazon S3 bucket from the health-checking component Configure S3 events to invoke an IAM Lambda function that analyzes the data
- E. Generate events from the health-checking component and send them to Amazon CloudWatch Events. Include the status data as event payload
- F. Use CloudWatch Events rules to invoke an IAM Lambda function that analyzes the data.

**Answer:** A

**Explanation:**

Amazon CloudWatch monitoring is a service that collects and tracks metrics from AWS resources and applications, and provides visualization tools and alarms to monitor performance and availability<sup>1</sup>. The health status data in JSON format can be sent to CloudWatch as custom metrics<sup>2</sup>, and then displayed in CloudWatch dashboards<sup>3</sup>. The other options are either inefficient or insecure for monitoring application availability in near-real time.

**NEW QUESTION 116**

A security engineer logs in to the AWS Lambda console with administrator permissions. The security engineer is trying to view logs in Amazon CloudWatch for a Lambda function that is named my Function.

When the security engineer chooses the option in the Lambda console to view logs in CloudWatch, an "error loading Log Streams" message appears.

The IAM policy for the Lambda function's execution role contains the following:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "logs:CreateLogGroup",
      "Resource": "arn:aws:logs:us-east-1:111111111111:*"
    },
    {
      "Effect": "Allow",
      "Action": ["logs:PutLogEvents"],
      "Resource": ["arn:aws:logs:us-east-1:111111111111:log-
group:/aws/Lambda/myFunction:*"]
    }
  ]
}
```

How should the security engineer correct the error?

- A. Move the logs:CreateLogGroup action to the second Allow statement.
- B. Add the logs:PutDestination action to the second Allow statement.
- C. Add the logs:GetLogEvents action to the second Allow statement.
- D. Add the logs:CreateLogStream action to the second Allow statement.

**Answer:** D

**Explanation:**

CloudWatchLogsReadOnlyAccess doesn't include "logs:CreateLogStream" but it includes "logs:Get\*" <https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/iam-identity-based-access-control-cwl.html#:~:te>

**NEW QUESTION 117**

You have an S3 bucket defined in IAM. You want to ensure that you encrypt the data before sending it across the wire. What is the best way to achieve this. Please select:

- A. Enable server side encryption for the S3 bucket
- B. This request will ensure that the data is encrypted first.
- C. Use the IAM Encryption CLI to encrypt the data first
- D. Use a Lambda function to encrypt the data before sending it to the S3 bucket.
- E. Enable client encryption for the bucket

**Answer:** B

**Explanation:**

One can use the IAM Encryption CLI to encrypt the data before sending it across to the S3 bucket. Options A and C are invalid because this would still mean that data is transferred in plain text Option D is invalid because you cannot just enable client side encryption for the S3 bucket For more information on Encrypting and Decrypting data, please visit the below URL:

<https://IAM.amazonaws.com/blogs/security/how-to-encrypt-and-decrypt-your-data-with-the-IAM-encryption-cli> The correct answer is: Use the IAM Encryption CLI to encrypt the data first Submit your Feedback/Queries to our Experts

**NEW QUESTION 120**

A security administrator is setting up a new AWS account. The security administrator wants to secure the data that a company stores in an Amazon S3 bucket. The security administrator also wants to reduce the chance of unintended data exposure and the potential for misconfiguration of objects that are in the S3 bucket. Which solution will meet these requirements with the LEAST operational overhead?

- A. Configure the S3 Block Public Access feature for the AWS account.
- B. Configure the S3 Block Public Access feature for all objects that are in the bucket.
- C. Deactivate ACLs for objects that are in the bucket.
- D. Use AWS PrivateLink for Amazon S3 to access the bucket.

**Answer:** D

**NEW QUESTION 125**

A company has deployed Amazon GuardDuty and now wants to implement automation for potential threats. The company has decided to start with RDP brute force attacks that come from Amazon EC2 instances in the company's AWS environment. A security engineer needs to implement a solution that blocks the detected communication from a suspicious instance until investigation and potential remediation can occur.

Which solution will meet these requirements?

- A. Configure GuardDuty to send the event to an Amazon Kinesis data stream
- B. Process the event with an Amazon Kinesis Data Analytics for Apache Flink application that sends a notification to the company through Amazon Simple Notification Service (Amazon SNS). Add rules to the network ACL to block traffic to and from the suspicious instance.
- C. Configure GuardDuty to send the event to Amazon EventBridge (Amazon CloudWatch Events). Deploy an AWS WAF web ACL
- D. Process the event with an AWS Lambda function that sends a notification to the company through Amazon Simple Notification Service (Amazon SNS) and adds a web ACL rule to block traffic to and from the suspicious instance.
- E. Enable AWS Security Hub to ingest GuardDuty findings and send the event to Amazon EventBridge (Amazon CloudWatch Events). Deploy AWS Network Firewall
- F. Process the event with an AWS Lambda function that adds a rule to a Network Firewall firewall policy to block traffic to and from the suspicious instance.
- G. Enable AWS Security Hub to ingest GuardDuty findings
- H. Configure an Amazon Kinesis data stream as an event destination for Security Hub
- I. Process the event with an AWS Lambda function that replaces the security group of the suspicious instance with a security group that does not allow any connections.

**Answer:** C



**Explanation:**

<https://aws.amazon.com/blogs/security/automatically-block-suspicious-traffic-with-aws-network-firewall-and-a>

**NEW QUESTION 129**

A company's application team wants to replace an internal application with a new IAM architecture that consists of Amazon EC2 instances, an IAM Lambda function, and an Amazon S3 bucket in a single IAM Region. After an architecture review, the security team mandates that no application network traffic can traverse the public internet at any point. The security team already has an SCP in place for the company's organization in IAM Organizations to restrict the creation of internet gateways, NAT gateways, and egress-only gateways.

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

- A. Create an S3 endpoint that has a full-access policy for the application's VPC.
- B. Create an S3 access point for the S3 bucket
- C. Include a policy that restricts the network origin to VPCs.
- D. Launch the Lambda function
- E. Enable the block public access configuration.
- F. Create a security group that has an outbound rule over port 443 with a destination of the S3 endpoint. Associate the security group with the EC2 instances.
- G. Create a security group that has an outbound rule over port 443 with a destination of the S3 access point. Associate the security group with the EC2 instances.
- H. Launch the Lambda function in a VPC.

**Answer:** ADF

**NEW QUESTION 131**

A company uses a third-party application to store encrypted data in Amazon S3. The company uses another third-party application that decrypts the data from Amazon S3 to ensure separation of duties. Between the applications, a Security Engineer warns to separate the permissions using IAM roles attached to Amazon EC2 instances. The company prefers to use native IAM services.

Which encryption method will meet these requirements?

- A. Use encrypted Amazon EBS volumes with Amazon default keys (IAM EBS)
- B. Use server-side encryption with customer-provided keys (SSE-C)
- C. Use server-side encryption with IAM KMS managed keys (SSE-KMS)
- D. Use server-side encryption with Amazon S3 managed keys (SSE-S3)

**Answer:** C

**NEW QUESTION 134**

A developer signed in to a new account within an IAM Organization organizational unit (OU) containing multiple accounts. Access to the Amazon S3 service is restricted with the following SCP.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": "s3:*",
      "Resource": "*"
    }
  ]
}
```

How can the security engineer provide the developer with Amazon S3 access without affecting other accounts?

- A. Move the SCP to the root OU of organization to remove the restriction to access Amazon S3.
- B. Add an IAM policy for the developer, which grants S3 access.
- C. Create a new OU without applying the SCP restricting S3 access
- D. Move the developer account to this new OU.
- E. Add an allow list for the developer account for the S3 service.

**Answer:** C

**NEW QUESTION 138**

A company needs to use HTTPS when connecting to its web applications to meet compliance requirements. These web applications run in Amazon VPC on Amazon EC2 instances behind an Application Load Balancer (ALB). A security engineer wants to ensure that the load balancer will only accept connections over port 443, even if the ALB is mistakenly configured with an HTTP listener.

Which configuration steps should the security engineer take to accomplish this task?

- A. Create a security group with a rule that denies inbound connections from 0.0.0.0/0 on port 80. Attach this security group to the ALB to overwrite more permissive rules from the ALB's default security group.
- B. Create a network ACL that denies inbound connections from 0.0.0.0/0 on port 80. Associate the network ACL with the VPC's internet gateway.
- C. Create a network ACL that allows outbound connections to the VPC IP range on port 443 only. Associate the network ACL with the VPC's internet gateway.
- D. Create a security group with a single inbound rule that allows connections from 0.0.0.0/0 on port 443. Ensure this security group is the only one associated with the ALB.

**Answer:** D

**Explanation:**

To ensure that the load balancer only accepts connections over port 443, the security engineer should do the following:

- Create a security group with a single inbound rule that allows connections from 0.0.0.0/0 on port 443.

This means that the security group allows HTTPS traffic from any source IP address.

➤ Ensure this security group is the only one associated with the ALB. This means that the security group overrides any other rules that might allow HTTP traffic on port 80.

#### NEW QUESTION 141

A company uses identity federation to authenticate users into an identity account (987654321987) where the users assume an IAM role named IdentityRole. The users then assume an IAM role named JobFunctionRole in the target IAM account (123456789123) to perform their job functions. A user is unable to assume the IAM role in the target account. The policy attached to the role in the identity account is:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "sts:AssumeRole"
      ],
      "Resource": [
        "arn:aws:iam::*:role/JobFunctionRole"
      ],
      "Effect": "Allow"
    }
  ]
}
```

What should be done to enable the user to assume the appropriate role in the target account?

A Update the IAM policy attached to the role in the identity account to be:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "sts:AssumeRole"
      ],
      "Resource": [
        "arn:aws:iam::123456789123:role/JobFunctionRole"
      ],
      "Effect": "Allow"
    }
  ]
}
```

B Update the trust policy on the role in the target account to be:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::987654321987:role/IdentityRole"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

C Update the trust policy on the role in the identity account to be:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": { "AWS": "arn:aws:iam::987654321987:root" },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

D Update the IAM policy attached to the role in the target account to be:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Stmt1502946463000",
      "Effect": "Allow",
      "Action": "sts:AssumeRole",
      "Resource": "arn:aws:iam::123456789123:role/JobFunctionRole"
    }
  ]
}
```

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

**Answer: B**

**Explanation:**

<https://aws.amazon.com/blogs/security/how-to-use-trust-policies-with-iam-roles/>

**NEW QUESTION 146**

A Security Engineer is working with a Product team building a web application on AWS. The application uses Amazon S3 to host the static content, Amazon API Gateway to provide RESTful services; and Amazon DynamoDB as the backend data store. The users already exist in a directory that is exposed through a SAML identity provider.

Which combination of the following actions should the Engineer take to enable users to be authenticated into the web application and call APIs? (Choose three.)

- A. Create a custom authorization service using AWS Lambda.
- B. Configure a SAML identity provider in Amazon Cognito to map attributes to the Amazon Cognito user pool attributes.
- C. Configure the SAML identity provider to add the Amazon Cognito user pool as a relying party.
- D. Configure an Amazon Cognito identity pool to integrate with social login providers.
- E. Update DynamoDB to store the user email addresses and passwords.
- F. Update API Gateway to use a COGNITO\_USER\_POOLS authorizer.

**Answer: BCF**

**Explanation:**

The combination of the following actions should the Engineer take to enable users to be authenticated into the web application and call APIs are:

- B. Configure a SAML identity provider in Amazon Cognito to map attributes to the Amazon Cognito user pool attributes. This is a necessary step to federate the

existing users from the SAML identity provider to the Amazon Cognito user pool, which will be used for authentication and authorization<sup>1</sup>.

- C. Configure the SAML identity provider to add the Amazon Cognito user pool as a relying party. This is a necessary step to establish a trust relationship between the SAML identity provider and the Amazon Cognito user pool, which will allow the users to sign in using their existing credentials<sup>2</sup>.
- F. Update API Gateway to use a COGNITO\_USER\_POOLS authorizer. This is a necessary step to enable API Gateway to use the Amazon Cognito user pool as an authorizer for the RESTful services, which will validate the identity or access tokens that are issued by Amazon Cognito when a user signs in successfully<sup>3</sup>. The other options are incorrect because:
- A. Creating a custom authorization service using AWS Lambda is not a necessary step, because Amazon Cognito user pools can provide built-in authorization features, such as scopes and groups, that can be used to control access to API resources<sup>4</sup>.
- D. Configuring an Amazon Cognito identity pool to integrate with social login providers is not a necessary step, because the users already exist in a directory that is exposed through a SAML identity provider, and there is no requirement to support social login providers<sup>5</sup>.
- E. Updating DynamoDB to store the user email addresses and passwords is not a necessary step, because the user credentials are already stored in the SAML identity provider, and there is no need to duplicate them in DynamoDB<sup>6</sup>.

References:

1: Using Tokens with User Pools 2: Adding SAML Identity Providers to a User Pool 3: Control Access to a REST API Using Amazon Cognito User Pools as Authorizer 4: API Authorization with Resource Servers and OAuth 2.0 Scopes 5: Using Identity Pools (Federated Identities) 6: Amazon DynamoDB

#### NEW QUESTION 149

A company wants to ensure that its IAM resources can be launched only in the us-east-1 and us-west-2 Regions.

What is the MOST operationally efficient solution that will prevent developers from launching Amazon EC2 instances in other Regions?

- A. Enable Amazon GuardDuty in all Region
- B. Create alerts to detect unauthorized activity outside us-east-1 and us-west-2.
- C. Use an organization in IAM Organization
- D. Attach an SCP that allows all actions when the IAM: Requested Region condition key is either us-east-1 or us-west-2. Delete the FullIAMAccess policy.
- E. Provision EC2 resources by using IAM Cloud Formation templates through IAM CodePipeline
- F. Allow only the values of us-east-1 and us-west-2 in the IAM CloudFormation template's parameters.
- G. Create an IAM Config rule to prevent unauthorized activity outside us-east-1 and us-west-2.

**Answer: C**

#### NEW QUESTION 150

A company is evaluating its security posture. In the past, the company has observed issues with specific hosts and host header combinations that affected the company's business. The company has configured AWS WAF web ACLs as an initial step to mitigate these issues.

The company must create a log analysis solution for the AWS WAF web ACLs to monitor problematic activity. The company wants to process all the AWS WAF logs in a central location. The company must have the ability to filter out requests based on specific hosts.

A security engineer starts to enable access logging for the AWS WAF web ACLs.

What should the security engineer do next to meet these requirements with the MOST operational efficiency?

- A. Specify Amazon Redshift as the destination for the access log
- B. Deploy the Amazon Athena Redshift connecto
- C. Use Athena to query the data from Amazon Redshift and to filter the logs by host.
- D. Specify Amazon CloudWatch as the destination for the access log
- E. Use Amazon CloudWatch Logs Insights to design a query to filter the logs by host.
- F. Specify Amazon CloudWatch as the destination for the access log
- G. Export the CloudWatch logs to an Amazon S3 bucket
- H. Use Amazon Athena to query the logs and to filter the logs by host.
- I. Specify Amazon CloudWatch as the destination for the access log
- J. Use Amazon Redshift Spectrum to query the logs and to filter the logs by host.

**Answer: C**

#### Explanation:

The correct answer is C. Specify Amazon CloudWatch as the destination for the access logs. Export the CloudWatch logs to an Amazon S3 bucket. Use Amazon Athena to query the logs and to filter the logs by host.

According to the AWS documentation<sup>1</sup>, AWS WAF offers logging for the traffic that your web ACLs analyze. The logs include information such as the time that AWS WAF received the request from your protected AWS resource, detailed information about the request, and the action setting for the rule that the request matched. You can send your logs to an Amazon CloudWatch Logs log group, an Amazon Simple Storage Service (Amazon S3) bucket, or an Amazon Kinesis Data Firehose.

To create a log analysis solution for the AWS WAF web ACLs, you can use Amazon Athena, which is an interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL<sup>2</sup>. You can use Athena to query and filter the AWS WAF logs by host or any other criteria. Athena is serverless, so there is no infrastructure to manage, and you pay only for the queries that you run.

To use Athena with AWS WAF logs, you need to export the CloudWatch logs to an S3 bucket. You can do this by creating a subscription filter that sends your log events to a Kinesis Data Firehose delivery stream, which then delivers the data to an S3 bucket<sup>3</sup>. Alternatively, you can use AWS DMS to migrate your CloudWatch logs to S3<sup>4</sup>.

After you have exported your CloudWatch logs to S3, you can create a table in Athena that points to your S3 bucket and use the AWS service log format that matches your log schema<sup>5</sup>. For example, if you are using JSON format for your AWS WAF logs, you can use the AWSJSONSerDe serde. Then you can run SQL queries on your Athena table and filter the results by host or any other field in your log data.

Therefore, this solution meets the requirements of creating a log analysis solution for the AWS WAF web ACLs with the most operational efficiency. This solution does not require setting up any additional infrastructure or services, and it leverages the existing capabilities of CloudWatch, S3, and Athena.

The other options are incorrect because:

- A. Specifying Amazon Redshift as the destination for the access logs is not possible, because AWS WAF does not support sending logs directly to Redshift. You would need to use an intermediate service such as Kinesis Data Firehose or AWS DMS to load the data from CloudWatch or S3 to Redshift. Deploying the Amazon Athena Redshift connector is not necessary, because you can query Redshift data directly from Athena without using a connector<sup>6</sup>. This solution would also incur additional costs and operational overhead of managing a Redshift cluster.
- B. Specifying Amazon CloudWatch as the destination for the access logs is possible, but using Amazon CloudWatch Logs Insights to design a query to filter the logs by host is not efficient or scalable. CloudWatch Logs Insights is a feature that enables you to interactively search and analyze your log data in CloudWatch Logs<sup>7</sup>. However, CloudWatch Logs Insights has some limitations, such as a maximum query duration of 20 minutes, a maximum of 20 log groups per query, and a maximum retention period of 24 months<sup>8</sup>. These limitations may affect your ability to perform complex and long-running analysis on your AWS WAF logs.



➤ D. Specifying Amazon CloudWatch as the destination for the access logs is possible, but using Amazon Redshift Spectrum to query the logs and filter them by host is not efficient or cost-effective. Redshift Spectrum is a feature of Amazon Redshift that enables you to run queries against exabytes of data in S3 without loading or transforming any data<sup>9</sup>. However, Redshift Spectrum requires a Redshift cluster to process the queries, which adds additional costs and operational overhead. Redshift Spectrum also charges you based on the number of bytes scanned by each query, which can be expensive if you have large volumes of log data<sup>10</sup>.

References:

1: Logging AWS WAF web ACL traffic - Amazon Web Services 2: What Is Amazon Athena? - Amazon Athena 3: Streaming CloudWatch Logs Data to Amazon S3 - Amazon CloudWatch Logs 4: Migrate data from CloudWatch Logs using AWS Database Migration Service - AWS Database Migration Service 5: Querying AWS service logs - Amazon Athena 6: Querying data from Amazon Redshift - Amazon Athena 7: Analyzing log data with CloudWatch Logs Insights - Amazon CloudWatch Logs 8: CloudWatch Logs Insights quotas - Amazon CloudWatch 9: Querying external data using Amazon Redshift Spectrum - Amazon Redshift 10: Amazon Redshift Spectrum pricing - Amazon Redshift

#### NEW QUESTION 154

A company has two VPCs in the same AWS Region and in the same AWS account. Each VPC uses a CIDR block that does not overlap with the CIDR block of the other VPC. One VPC contains AWS Lambda functions that run inside a subnet that accesses the internet through a NAT gateway. The Lambda functions require access to a publicly accessible Amazon Aurora MySQL database that is running in the other VPC.

A security engineer determines that the Aurora database uses a security group rule that allows connections from the NAT gateway IP address that the Lambda functions use. The company's security policy states that no database should be publicly accessible.

What is the MOST secure way that the security engineer can provide the Lambda functions with access to the Aurora database?

- A. Move the Aurora database into a private subnet that has no internet access routes in the database's current VPC. Configure the Lambda functions to use the Aurora database's new private IP address to access the database. Configure the Aurora database's security group to allow access from the private IP addresses of the Lambda functions.
- B. Establish a VPC endpoint between the two VPCs in the Aurora database's VPC. Configure a service VPC endpoint for Amazon RDS. In the Lambda functions' VPC, configure an interface VPC endpoint that uses the service endpoint in the Aurora database's VPC. Configure the service endpoint to allow connections from the Lambda functions.
- C. Establish an AWS Direct Connect interface between the VPCs. Configure the Lambda functions to use a new route table that accesses the Aurora database through the Direct Connect interface. Configure the Aurora database's security group to allow access from the Direct Connect interface IP address.
- D. Move the Lambda functions into a public subnet in their VPC. Move the Aurora database into a private subnet in its VPC. Configure the Lambda functions to use the Aurora database's new private IP address to access the database. Configure the Aurora database to allow access from the public IP addresses of the Lambda functions.

**Answer: B**

#### Explanation:

This option involves creating a VPC Endpoint between the two VPCs that allows private communication between them without going through the internet or exposing any public IP addresses. In this option, a VPC endpoint for Amazon RDS will be established, and an interface VPC endpoint will be created that points to the service endpoint in the Aurora database's VPC. This way, the Lambda functions can use the private IP address of the Aurora database to access it through the VPC endpoint without exposing any public IP addresses or allowing public internet access to the database.

#### NEW QUESTION 159

A company has an application that uses dozens of Amazon DynamoDB tables to store data. Auditors find that the tables do not comply with the company's data protection policy.

The company's retention policy states that all data must be backed up twice each month: once at midnight on the 15th day of the month and again at midnight on the 25th day of the month. The company must retain the backups for 3 months.

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

- A. Use the DynamoDB on-demand backup capability to create a backup plan.
- B. Configure a lifecycle policy to expire backups after 3 months.
- C. Use AWS DataSync to create a backup plan.
- D. Add a backup rule that includes a retention period of 3 months.
- E. Use AWS Backup to create a backup plan.
- F. Add a backup rule that includes a retention period of 3 months.
- G. Set the backup frequency by using a cron schedule expression.
- H. Assign each DynamoDB table to the backup plan.
- I. Set the backup frequency by using a rate schedule expression.
- J. Assign each DynamoDB table to the backup plan.

**Answer: AD**

#### NEW QUESTION 162

A company has an application that uses an Amazon RDS PostgreSQL database. The company is developing an application feature that will store sensitive information for an individual in the database.

During a security review of the environment, the company discovers that the RDS DB instance is not encrypting data at rest. The company needs a solution that will provide encryption at rest for all the existing data and for any new data that is entered for an individual.

Which combination of options can the company use to meet these requirements? (Select TWO.)

- A. Create a snapshot of the DB instance.
- B. Copy the snapshot to a new snapshot, and enable encryption for the copy process.
- C. Use the new snapshot to restore the DB instance.
- D. Modify the configuration of the DB instance by enabling encryption.
- E. Create a snapshot of the DB instance.
- F. Use the snapshot to restore the DB instance.
- G. Use IAM Key Management Service (IAM KMS) to create a new default IAM managed AWS/RDS key. Select this key as the encryption key for operations with Amazon RDS.
- H. Use IAM Key Management Service (IAM KMS) to create a new CMK.
- I. Select this key as the encryption key for operations with Amazon RDS.
- J. Create a snapshot of the DB instance.
- K. Enable encryption on the snapshot. Use the snapshot to restore the DB instance.

**Answer:** CE

#### NEW QUESTION 163

A company uses infrastructure as code (IaC) to create AWS infrastructure. The company writes the code as AWS CloudFormation templates to deploy the infrastructure. The company has an existing CI/CD pipeline that the company can use to deploy these templates. After a recent security audit, the company decides to adopt a policy-as-code approach to improve the company's security posture on AWS. The company must prevent the deployment of any infrastructure that would violate a security policy, such as an unencrypted Amazon Elastic Block Store (Amazon EBS) volume. Which solution will meet these requirements?

- A. Turn on AWS Trusted Advisor
- B. Configure security notifications as webhooks in the preferences section of the CI/CD pipeline.
- C. Turn on AWS Config
- D. Use the prebuilt rules or customized rule
- E. Subscribe the CI/CD pipeline to an Amazon Simple Notification Service (Amazon SNS) topic that receives notifications from AWS Config.
- F. Create rule sets in AWS CloudFormation Guard
- G. Run validation checks for CloudFormation templates as a phase of the CI/CD process.
- H. Create rule sets as SCP
- I. Integrate the SCPs as a part of validation control in a phase of the CI/CD process.

**Answer:** C

#### Explanation:

The correct answer is C. Create rule sets in AWS CloudFormation Guard. Run validation checks for CloudFormation templates as a phase of the CI/CD process. This answer is correct because AWS CloudFormation Guard is a tool that helps you implement policy-as-code for your CloudFormation templates. You can use Guard to write rules that define your security policies, such as requiring encryption for EBS volumes, and then validate your templates against those rules before deploying them. You can integrate Guard into your CI/CD pipeline as a step that runs the validation checks and prevents the deployment of any non-compliant templates<sup>12</sup>.

The other options are incorrect because:

- A. Turning on AWS Trusted Advisor and configuring security notifications as webhooks in the preferences section of the CI/CD pipeline is not a solution, because AWS Trusted Advisor is not a policy-as-code tool, but a service that provides recommendations to help you follow AWS best practices. Trusted Advisor does not allow you to define your own security policies or validate your CloudFormation templates against them<sup>3</sup>.
- B. Turning on AWS Config and using the prebuilt or customized rules is not a solution, because AWS Config is not a policy-as-code tool, but a service that monitors and records the configuration changes of your AWS resources. AWS Config does not allow you to validate your CloudFormation templates before deploying them, but only evaluates the compliance of your resources after they are created<sup>4</sup>.
- D. Creating rule sets as SCPs and integrating them as a part of validation control in a phase of the CI/CD process is not a solution, because SCPs are not policy-as-code tools, but policies that you can use to manage permissions in your AWS Organizations. SCPs do not allow you to validate your CloudFormation templates, but only restrict the actions that users and roles can perform in your accounts<sup>5</sup>.

References:

1: What is AWS CloudFormation Guard? 2: Introducing AWS CloudFormation Guard 2.0 3: AWS Trusted Advisor 4: What Is AWS Config? 5: Service control policies - AWS Organizations

#### NEW QUESTION 167

A company is building a data processing application that uses AWS Lambda functions. The application's Lambda functions need to communicate with an Amazon RDS DB instance that is deployed within a VPC in the same AWS account. Which solution meets these requirements in the MOST secure way?

- A. Configure the DB instance to allow public access. Update the DB instance security group to allow access from the Lambda public address space for the AWS Region.
- B. Deploy the Lambda functions inside the VPC. Attach a network ACL to the Lambda subnet. Provide outbound rule access to the VPC CIDR range only. Update the DB instance security group to allow traffic from 0.0.0.0/0.
- C. Deploy the Lambda functions inside the VPC. Attach a security group to the Lambda functions. Provide outbound rule access to the VPC CIDR range only. Update the DB instance security group to allow traffic from the Lambda security group.
- D. Peer the Lambda default VPC with the VPC that hosts the DB instance to allow direct network access without the need for security groups.

**Answer:** C

#### Explanation:

This solution ensures that the Lambda functions are deployed inside the VPC and can communicate with the Amazon RDS DB instance securely. The security group attached to the Lambda functions only allows outbound traffic to the VPC CIDR range, and the DB instance security group only allows traffic from the Lambda security group. This solution ensures that the Lambda functions can communicate with the DB instance securely and that the DB instance is not exposed to the public internet.

#### NEW QUESTION 172

A company has two teams, and each team needs to access its respective Amazon S3 buckets. The company anticipates adding more teams that also will have their own S3 buckets. When the company adds these teams, team members will need the ability to be assigned to multiple teams. Team members also will need the ability to change teams. Additional S3 buckets can be created or deleted.

An IAM administrator must design a solution to accomplish these goals. The solution also must be scalable and must require the least possible operational overhead.

Which solution meets these requirements?

- A. Add users to groups that represent the team.
- B. Create a policy for each team that allows the team to access its respective S3 buckets only.
- C. Attach the policy to the corresponding group.
- D. Create an IAM role for each team.
- E. Create a policy for each team that allows the team to access its respective S3 buckets only.
- F. Attach the policy to the corresponding role.
- G. Create IAM roles that are labeled with an access tag value of a team.
- H. Create one policy that allows dynamic access to S3 buckets with the same tag.
- I. Attach the policy to the IAM role.

- J. Tag the S3 buckets accordingly.
- K. Implement a role-based access control (RBAC) authorization mode
- L. Create the corresponding policies, and attach them to the IAM users.

**Answer:** A

#### NEW QUESTION 173

There are currently multiple applications hosted in a VPC. During monitoring it has been noticed that multiple port scans are coming in from a specific IP Address block. The internal security team has requested that all offending IP Addresses be denied for the next 24 hours. Which of the following is the best method to quickly and temporarily deny access from the specified IP Address's.  
Please select:

- A. Create an AD policy to modify the Windows Firewall settings on all hosts in the VPC to deny access from the IP Address block.
- B. Modify the Network ACLs associated with all public subnets in the VPC to deny access from the IP Address block.
- C. Add a rule to all of the VPC Security Groups to deny access from the IP Address block.
- D. Modify the Windows Firewall settings on all AMI'S that your organization uses in that VPC to deny access from the IP address block.

**Answer:** B

#### Explanation:

NACL acts as a firewall at the subnet level of the VPC and we can deny the offending IP address block at the subnet level using NACL rules to block the incoming traffic to the VPC instances. Since NACL rules are applied as per the Rule numbers make sure that this rule number should take precedence over other rule numbers if there are any such rules that will allow traffic from these IP ranges. The lowest rule number has more precedence over a rule that has a higher number. The IAM Documentation mentions the following as a best practices for IAM users

For extra security, enable multi-factor authentication (MFA) for privileged IAM users (users who are allowed access to sensitive resources or APIs). With MFA, users have a device that generates a unique authentication code (a one-time password, or OTP). Users must provide both their normal credentials (like their user name and password) and the OTP. The MFA device can either be a special piece of hardware, or it can be a virtual device (for example, it can run in an app on a smartphone).

Options C is invalid because these options are not available Option D is invalid because there is not root access for users

For more information on IAM best practices, please visit the below URL: <https://docs.IAM.amazon.com/IAM/latest/UserGuide/best-practices.html>

The correct answer is: Modify the Network ACLs associated with all public subnets in the VPC to deny access from the IP Address block.

omit your Feedback/Queries to our Experts

#### NEW QUESTION 175

.....



## Thank You for Trying Our Product

### We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

### SCS-C02 Practice Exam Features:

- \* SCS-C02 Questions and Answers Updated Frequently
- \* SCS-C02 Practice Questions Verified by Expert Senior Certified Staff
- \* SCS-C02 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- \* SCS-C02 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

**100% Actual & Verified — Instant Download, Please Click**  
**[Order The SCS-C02 Practice Test Here](#)**