


## AWS-Certified-Security-Specialty Dumps

### Amazon AWS Certified Security - Specialty

<https://www.certleader.com/AWS-Certified-Security-Specialty-dumps.html>



**NEW QUESTION 1**

A company manages three separate IAM accounts for its production, development, and test environments. Each Developer is assigned a unique IAM user under the development account. A new application hosted on an Amazon EC2 instance in the developer account requires read access to the archived documents stored in an Amazon S3 bucket in the production account.

How should access be granted?

- A. Create an IAM role in the production account and allow EC2 instances in the development account to assume that role using the trust policy.
- B. Provide read access for the required S3 bucket to this role.
- C. Use a custom identity broker to allow Developer IAM users to temporarily access the S3 bucket.
- D. Create a temporary IAM user for the application to use in the production account.
- E. Create a temporary IAM user in the production account and provide read access to Amazon S3. Generate the temporary IAM user's access key and secret key and store these on the EC2 instance used by the application in the development account.

**Answer:** A

**Explanation:**

<https://IAM.amazon.com/premiumsupport/knowledge-center/cross-account-access-s3/>

**NEW QUESTION 2**

A company wants to protect its website from man-in-the-middle attacks by using Amazon CloudFront. Which solution will meet these requirements with the LEAST operational overhead?

- A. Use the SimpleCORS managed response headers policy.
- B. Use a Lambda@Edge function to add the Strict-Transport-Security response header.
- C. Use the SecurityHeadersPolicy managed response headers policy.
- D. Include the X-XSS-Protection header in a custom response headers policy.

**Answer:** C

**Explanation:**

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/using-managed-response-headers-policy.html> The SecurityHeadersPolicy is a managed policy provided by Amazon CloudFront that includes a set of recommended security headers to enhance the security of your website. These headers help protect against various types of attacks, including man-in-the-middle attacks. By applying the SecurityHeadersPolicy to your CloudFront distribution, the necessary security headers will be automatically added to the responses sent by CloudFront. This reduces operational overhead because you don't have to manually configure or manage the headers yourself.

**NEW QUESTION 3**

A company needs to improve its ability to identify and prevent IAM policies that grant public access or cross-account access to resources. The company has implemented AWS Organizations and has started using AWS Identity and Access Management Access Analyzer to refine overly broad access to accounts in the organization.

A security engineer must automate a response in the company's organization for any newly created policies that are overly permissive. The automation must remediate external access and must notify the company's security team.

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

- A. Create an AWS Step Functions state machine that checks the resource type in the finding and adds an explicit Deny statement in the trust policy for the IAM role.
- B. Configure the state machine to publish a notification to an Amazon Simple Notification Service (Amazon SNS) topic.
- C. Create an AWS Batch job that forwards any resource type findings to an AWS Lambda function. Configure the Lambda function to add an explicit Deny statement in the trust policy for the IAM role.
- D. Configure the AWS Batch job to publish a notification to an Amazon Simple Notification Service (Amazon SNS) topic.
- E. In Amazon EventBridge, create an event rule that matches active IAM Access Analyzer findings and invokes AWS Step Functions for resolution.
- F. In Amazon CloudWatch, create a metric filter that matches active IAM Access Analyzer findings and invokes AWS Batch for resolution.
- G. Create an Amazon Simple Queue Service (Amazon SQS) queue.
- H. Configure the queue to forward a notification to the security team that an external principal has been granted access to the specific IAM role and has been blocked.
- I. Create an Amazon Simple Notification Service (Amazon SNS) topic for external or cross-account access notice.
- J. Subscribe the security team's email addresses to the topic.

**Answer:** ACF

**Explanation:**

The correct answer is A, C, and F.

To automate a response for any newly created policies that are overly permissive, the security engineer needs to use a combination of services that can monitor, analyze, remediate, and notify the security incidents.

Option A is correct because creating an AWS Step Functions state machine that checks the resource type in the finding and adds an explicit Deny statement in the trust policy for the IAM role is a valid way to remediate external access. AWS Step Functions is a service that allows you to coordinate multiple AWS services into serverless workflows. You can use Step Functions to invoke AWS Lambda functions, which can modify the IAM policies programmatically. You can also use Step Functions to publish a notification to an Amazon SNS topic, which can send messages to subscribers such as email addresses.

Option B is incorrect because creating an AWS Batch job that forwards any resource type findings to an AWS Lambda function is not a suitable way to automate a response. AWS Batch is a service that enables you to run batch computing workloads on AWS. Batch is designed for large-scale and long-running jobs that can benefit from parallelization and dynamic provisioning of compute resources. Batch is not intended for event-driven and real-time workflows that require immediate response.

Option C is correct because creating an Amazon EventBridge event rule that matches active IAM Access Analyzer findings and invokes AWS Step Functions for resolution is a valid way to monitor and analyze the security incidents. Amazon EventBridge is a serverless event bus service that allows you to connect your applications with data from various sources. EventBridge can use rules to match events and route them to targets for processing. You can use EventBridge to invoke AWS Step Functions state machines from the IAM Access Analyzer findings.

Option D is incorrect because creating an Amazon CloudWatch metric filter that matches active IAM Access Analyzer findings and invokes AWS Batch for resolution is not a suitable way to monitor and analyze the security incidents. Amazon CloudWatch is a service that provides monitoring and observability for your AWS resources and applications. CloudWatch can collect metrics, logs, and events from various sources and perform actions based on alarms or filters. However, CloudWatch cannot directly invoke AWS Batch jobs from the IAM Access Analyzer findings. You would need to use another service such as EventBridge or SNS.

to trigger the Batch job.

Option E is incorrect because creating an Amazon SQS queue that forwards a notification to the security team that an external principal has been granted access to the specific IAM role and has been blocked is not a valid way to notify the security incidents. Amazon SQS is a fully managed message queue service that enables you to decouple and scale microservices, distributed systems, and serverless applications. SQS can deliver messages to consumers that poll the queue for messages. However, SQS cannot directly forward a notification to the security team's email addresses. You would need to use another service such as SNS or SES to send email notifications.

Option F is correct because creating an Amazon SNS topic for external or cross-account access notices and subscribing the security team's email addresses to the topic is a valid way to notify the security incidents. Amazon SNS is a fully managed messaging service that enables you to decouple and scale microservices, distributed systems, and serverless applications. SNS can deliver messages to a variety of endpoints, such as email, SMS, or HTTP. You can use SNS to send email notifications to the security team when a critical security finding is detected.

References:

- AWS Step Functions
- AWS Batch
- Amazon EventBridge
- Amazon CloudWatch
- Amazon SQS
- Amazon SNS

#### NEW QUESTION 4

A company has retail stores The company is designing a solution to store scanned copies of customer receipts on Amazon S3 Files will be between 100 KB and 5 MB in PDF format Each retail store must have a unique encryption key Each object must be encrypted with a unique key Which solution will meet these requirements?

- A. Create a dedicated AWS Key Management Service (AWS KMS) customer managed key for each retail store Use the S3 Put operation to upload the objects to Amazon S3 Specify server-side encryption with AWS KMS keys (SSE-KMS) and the key ID of the store's key
- B. Create a new AWS Key Management Service (AWS KMS) customer managed key every day for each retail store Use the KMS Encrypt operation to encrypt objects Then upload the objects to Amazon S3
- C. Run the AWS Key Management Service (AWS KMS) GenerateDataKey operation every day for each retail store Use the data key and client-side encryption to encrypt the objects Then upload the objects to Amazon S3
- D. Use the AWS Key Management Service (AWS KMS) ImportKeyMaterial operation to import new key material to AWS KMS every day for each retail store Use a customer managed key and the KMS Encrypt operation to encrypt the objects Then upload the objects to Amazon S3

**Answer:** A

#### Explanation:

To meet the requirements of storing scanned copies of customer receipts on Amazon S3, where files will be between 100 KB and 5 MB in PDF format, each retail store must have a unique encryption key, and each object must be encrypted with a unique key, the most appropriate solution would be to create a dedicated AWS Key Management Service (AWS KMS) customer managed key for each retail store. Then, use the S3 Put operation to upload the objects to Amazon S3, specifying server-side encryption with AWS KMS keys (SSE-KMS) and the key ID of the store's key.

References: : Amazon S3 - Amazon Web Services : AWS Key Management Service - Amazon Web Services : Amazon S3 - Amazon Web Services : AWS Key Management Service - Amazon Web Service

#### NEW QUESTION 5

A company deploys a set of standard IAM roles in AWS accounts. The IAM roles are based on job functions within the company. To balance operational efficiency and security, a security engineer implemented AWS Organizations SCPs to restrict access to critical security services in all company accounts.

All of the company's accounts and OUs within AWS Organizations have a default FullAWSAccess SCP that is attached. The security engineer needs to ensure that no one can disable Amazon GuardDuty and AWS Security Hub. The security engineer also must not override other permissions that are granted by IAM policies that are defined in the accounts.

Which SCP should the security engineer attach to the root of the organization to meet these requirements? A)

```
"Version": "2012-10-17",
"Statement": [
  {
    "Effect": "Deny",
    "Action": [
      "guardduty:DeleteDetector",
      "guardduty:UpdateDetector",
      "securityhub:DisableSecurityHub"
    ],
    "Resource": [
      "*"
    ]
  }
]
```

B)

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": "*",
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "NotAction": [
        "guardduty:DeleteDetector",
        "guardduty:UpdateDetector",
        "securityhub:DisableSecurityHub"
      ],
      "Resource": [
        "*"
      ]
    }
  ]
}
```

C)

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "*",
      "Resource": "*"
    },
    {
      "Effect": "Deny",
      "NotAction": [
        "guardduty:DeleteDetector",
        "guardduty:UpdateDetector",
        "securityhub:DisableSecurityHub"
      ],
      "Resource": [
        "*"
      ]
    }
  ]
}
```

D)

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "NotAction": [
        "guardduty:DeleteDetector",
        "guardduty:UpdateDetector",
        "securityhub:DisableSecurityHub"
      ],
      "Resource": [
        "*"
      ]
    }
  ]
}
```

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

**Answer: A****NEW QUESTION 6**

A Security Engineer is asked to update an AWS CloudTrail log file prefix for an existing trail. When attempting to save the change in the CloudTrail console, the Security Engineer receives the following error message: `There is a problem with the bucket policy.` What will enable the Security Engineer to save the change?

- A. Create a new trail with the updated log file prefix, and then delete the original trail
- B. Update the existing bucket policy in the Amazon S3 console with the new log file prefix, and then update the log file prefix in the CloudTrail console.



- C. Update the existing bucket policy in the Amazon S3 console to allow the Security Engineer's Principal to perform PutBucketPolicy, and then update the log file prefix in the CloudTrail console.
- D. Update the existing bucket policy in the Amazon S3 console with the new log file prefix, and then update the log file prefix in the CloudTrail console.
- E. Update the existing bucket policy in the Amazon S3 console to allow the Security Engineer's Principal to perform GetBucketPolicy, and then update the log file prefix in the CloudTrail console.

**Answer: C**

**Explanation:**

The correct answer is C. Update the existing bucket policy in the Amazon S3 console with the new log file prefix, and then update the log file prefix in the CloudTrail console.

According to the AWS documentation<sup>1</sup>, a bucket policy is a resource-based policy that you can use to grant access permissions to your Amazon S3 bucket and the objects in it. Only the bucket owner can associate a policy with a bucket. The permissions attached to the bucket apply to all of the objects in the bucket that are owned by the bucket owner.

When you create a trail in CloudTrail, you can specify an existing S3 bucket or create a new one to store your log files. CloudTrail automatically creates a bucket policy for your S3 bucket that grants CloudTrail write-only access to deliver log files to your bucket. The bucket policy also grants read-only access to AWS services that you can use to view and analyze your log data, such as Amazon Athena, Amazon CloudWatch Logs, and Amazon QuickSight.

If you want to update the log file prefix for an existing trail, you must also update the existing bucket policy in the S3 console with the new log file prefix. The log file prefix is part of the resource ARN that identifies the objects in your bucket that CloudTrail can access. If you don't update the bucket policy with the new log file prefix, CloudTrail will not be able to deliver log files to your bucket, and you will receive an error message when you try to save the change in the CloudTrail console.

The other options are incorrect because:

- A. Creating a new trail with the updated log file prefix, and then deleting the original trail is not necessary and may cause data loss or inconsistency. You can simply update the existing trail and its associated bucket policy with the new log file prefix.
- B. Updating the existing bucket policy in the S3 console to allow the Security Engineer's Principal to perform PutBucketPolicy is not relevant to this issue. The PutBucketPolicy action allows you to create or replace a policy on a bucket, but it does not affect CloudTrail's ability to deliver log files to your bucket. You still need to update the existing bucket policy with the new log file prefix.
- D. Updating the existing bucket policy in the S3 console to allow the Security Engineer's Principal to perform GetBucketPolicy is not relevant to this issue. The GetBucketPolicy action allows you to retrieve a policy on a bucket, but it does not affect CloudTrail's ability to deliver log files to your bucket. You still need to update the existing bucket policy with the new log file prefix.

References:

1: Using bucket policies - Amazon Simple Storage Service

**NEW QUESTION 7**

A security engineer is creating an AWS Lambda function. The Lambda function needs to use a role that is named LambdaAuditRole to assume a role that is named AcmeAuditFactoryRole in a different AWS account.

When the code is processed, the following error message appears: "An error occurred (AccessDenied) when calling the AssumeRole operation."

Which combination of steps should the security engineer take to resolve this error? (Select TWO.)

- A. Ensure that LambdaAuditRole has the sts:AssumeRole permission for AcmeAuditFactoryRole.
- B. Ensure that LambdaAuditRole has the AWSLambdaBasicExecutionRole managed policy attached.
- C. Ensure that the trust policy for AcmeAuditFactoryRole allows the sts:AssumeRole action from LambdaAuditRole.
- D. Ensure that the trust policy for LambdaAuditRole allows the sts:AssumeRole action from the lambda.amazonaws.com service.
- E. Ensure that the sts:AssumeRole API call is being issued to the us-east-1 Region endpoint.

**Answer: AC**

**NEW QUESTION 8**

A company has an organization in AWS Organizations. The company wants to use AWS CloudFormation StackSets in the organization to deploy various AWS design patterns into environments. These patterns consist of Amazon EC2 instances, Elastic Load Balancing (ELB) load balancers, Amazon RDS databases, and Amazon Elastic Kubernetes Service (Amazon EKS) clusters or Amazon Elastic Container Service (Amazon ECS) clusters.

Currently, the company's developers can create their own CloudFormation stacks to increase the overall speed of delivery. A centralized CI/CD pipeline in a shared services AWS account deploys each CloudFormation stack.

The company's security team has already provided requirements for each service in accordance with internal standards. If there are any resources that do not comply with the internal standards, the security team must receive notification to take appropriate action. The security team must implement a notification solution that gives developers the ability to maintain the same overall delivery speed that they currently have.

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

- A. Create an Amazon Simple Notification Service (Amazon SNS) topic
- B. Subscribe the security team's email addresses to the SNS topic
- C. Create a custom AWS Lambda function that will run the aws cloudformation validate-template AWS CLI command on all CloudFormation templates before the build stage in the CI/CD pipeline
- D. Configure the CI/CD pipeline to publish a notification to the SNS topic if any issues are found.
- E. Create an Amazon Simple Notification Service (Amazon SNS) topic
- F. Subscribe the security team's email addresses to the SNS topic
- G. Create custom rules in CloudFormation Guard for each resource configuration
- H. In the CI/CD pipeline, before the build stage, configure a Docker image to run the cfn-guard command on the CloudFormation template
- I. Configure the CI/CD pipeline to publish a notification to the SNS topic if any issues are found.
- J. Create an Amazon Simple Notification Service (Amazon SNS) topic and an Amazon Simple Queue Service (Amazon SQS) queue
- K. Subscribe the security team's email addresses to the SNS topic
- L. Create an Amazon S3 bucket in the shared services AWS account
- M. Include an event notification to publish to the SQS queue when new objects are added to the S3 bucket
- N. Require the developers to put their CloudFormation templates in the S3 bucket
- O. Launch EC2 instances that automatically scale based on the SQS queue depth
- P. Configure the EC2 instances to use CloudFormation Guard to scan the templates and deploy the templates if there are no issues
- Q. Configure the CI/CD pipeline to publish a notification to the SNS topic if any issues are found.
- R. Create a centralized CloudFormation stack set that includes a standard set of resources that the developers can deploy in each AWS account
- S. Configure each CloudFormation template to meet the security requirement
- T. For any new resources or configurations, update the CloudFormation template and send the template to the security team for review
- . When the review is completed, add the new CloudFormation stack to the repository for the developers to use.

Answer: B

#### NEW QUESTION 9

A company has a batch-processing system that uses Amazon S3, Amazon EC2, and AWS Key Management Service (AWS KMS). The system uses two AWS accounts: Account A and Account B.

Account A hosts an S3 bucket that stores the objects that will be processed. The S3 bucket also stores the results of the processing. All the S3 bucket objects are encrypted by a KMS key that is managed in Account A.

Account B hosts a VPC that has a fleet of EC2 instances that access the S3 bucket in Account A by using statements in the bucket policy. The VPC was created with DNS hostnames enabled and DNS resolution enabled.

A security engineer needs to update the design of the system without changing any of the system's code. No AWS API calls from the batch-processing EC2 instances can travel over the internet.

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

- A. In the Account B VPC, create a gateway VPC endpoint for Amazon S3. For the gateway VPC endpoint, create a resource policy that allows the s3:GetObject, s3:ListBucket, s3:PutObject, and s3:PutObjectAcl actions for the S3 bucket.
- B. In the Account B VPC, create an interface VPC endpoint for Amazon S3. For the interface VPC endpoint, create a resource policy that allows the s3:GetObject, s3:ListBucket, s3:PutObject, and s3:PutObjectAcl actions for the S3 bucket.
- C. In the Account B VPC, create an interface VPC endpoint for AWS KM
- D. For the interface VPC endpoint, create a resource policy that allows the kms:Encrypt, kms:Decrypt, and kms:GenerateDataKey actions for the KMS ke
- E. Ensure that private DNS is turned on for the endpoint.
- F. In the Account B VPC, create an interface VPC endpoint for AWS KM
- G. For the interface VPC endpoint, create a resource policy that allows the kms:Encrypt, kms:Decrypt, and kms:GenerateDataKey actions for the KMS ke
- H. Ensure that private DNS is turned off for the endpoint.
- I. In the Account B VPC, verify that the S3 bucket policy allows the s3:PutObjectAcl action for cross-account us
- J. In the Account B VPC, create a gateway VPC endpoint for Amazon S3. For the gateway VPC endpoint, create a resource policy that allows the s3:GetObject, s3:ListBucket, and s3:PutObject actions for the S3 bucket.

Answer: BC

#### NEW QUESTION 10

A System Administrator is unable to start an Amazon EC2 instance in the eu-west-1 Region using an IAM role. The same System Administrator is able to start an EC2 instance in the eu-west-2 and eu-west-3 Regions. The IAMSystemAdministrator access policy attached to the System Administrator IAM role allows unconditional access to all IAM services and resources within the account.

Which configuration caused this issue?

A) An SCP is attached to the account with the following permission statement:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "All",
      "Action": "*",
      "Resource": "*"
    },
    {
      "Effect": "Deny",
      "NotAction": [
        "iam:*",
        "organizations:*",
        "route53:*",
        "budgets:*",
        "waf:*",
        "cloudfront:*",
        "globalaccelerator:*",
        "importexport:*",
        "support:*"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestedRegion": [
            "eu-west-*"
          ]
        }
      }
    }
  ]
}
```

- B)
- A permission boundary policy is attached to the System Administrator role with the following permission statement:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:*"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Deny",
      "NotAction": [
        "iam:*",
        "organizations:*",
        "route53:*",
        "budgets:*",
        "waf:*",
        "cloudfront:*",
        "globalaccelerator:*",
        "importexport:*",
        "support:*",
        "ec2:*"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestedRegion": [
            "eu-west-1"
          ]
        }
      }
    }
  ]
}
```

- C)  
A permission boundary is attached to the System Administrator role with the following permission statement:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:*"
      ],
      "Resource": "*"
    },
    {
      "Version": "2012-10-17",
      "Statement": [
        {
          "Effect": "Allow",
          "Action": "ec2:*",
          "Resource": "*",
          "Condition": {
            "StringEquals": {
              "aws:RequestedRegion": [
                "eu-west-1"
              ]
            }
          }
        }
      ]
    }
  ]
}
```

- D)  
An SCP is attached to the account with the following statement:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "*",
      "Resource": "*"
    },
    {
      "Effect": "Deny",
      "NotAction": [
        "iam:*",
        "organizations:*",
        "route53:*",
        "budgets:*",
        "waf:*",
        "cloudfront:*",
        "globalaccelerator:*",
        "importexport:*",
        "support:*",
        "ec2:*"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestedRegion": "us-east-1"
        }
      }
    }
  ]
}
```

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

**Answer: B**

#### NEW QUESTION 10

An ecommerce website was down for 1 hour following a DDoS attack. Users were unable to connect to the website during the attack period. The ecommerce company's security team is worried about future potential attacks and wants to prepare for such events. The company needs to minimize downtime in its response to similar attacks in the future.

Which steps would help achieve this? (Select TWO.)

- A. Enable Amazon GuardDuty to automatically monitor for malicious activity and block unauthorized access.
- B. Subscribe to IAM Shield Advanced and reach out to IAM Support in the event of an attack.
- C. Use VPC Flow Logs to monitor network traffic and an IAM Lambda function to automatically block an attacker's IP using security groups.
- D. Set up an Amazon CloudWatch Events rule to monitor the IAM CloudTrail events in real time, use IAM Config rules to audit the configuration, and use IAM Systems Manager for remediation.
- E. Use IAM WAF to create rules to respond to such attacks.

**Answer: BE**

#### Explanation:

To minimize downtime in response to DDoS attacks, the company should do the following:

- Subscribe to AWS Shield Advanced and reach out to AWS Support in the event of an attack. This provides access to 24x7 support from the AWS DDoS Response Team (DRT), as well as advanced detection and mitigation capabilities for network and application layer attacks.
- Use AWS WAF to create rules to respond to such attacks. This allows the company to filter web requests based on IP addresses, headers, body, or URI strings, and block malicious requests before they reach the web applications.

#### NEW QUESTION 12

A company has two AWS accounts. One account is for development workloads. The other account is for production workloads. For compliance reasons, the production account contains all the AWS Key Management Service (AWS KMS) keys that the company uses for encryption.

The company applies an IAM role to an AWS Lambda function in the development account to allow secure access to AWS resources. The Lambda function must access a specific KMS customer managed key that exists in the production account to encrypt the Lambda function's data.

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

- A. Configure the key policy for the customer managed key in the production account to allow access to the Lambda service.
- B. Configure the key policy for the customer managed key in the production account to allow access to the IAM role of the Lambda function in the development account.
- C. Configure a new IAM policy in the production account with permissions to use the customer managed key.
- D. Apply the IAM policy to the IAM role that the Lambda function in the development account uses.
- E. Configure a new key policy in the development account with permissions to use the customer managed key.
- F. Apply the key policy to the IAM role that the Lambda function in the development account uses.
- G. Configure the IAM role for the Lambda function in the development account by attaching an IAM policy that allows access to the customer managed key in the production account.

**Answer: BE**

#### Explanation:

To allow a Lambda function in one AWS account to access a KMS customer managed key in another AWS account, the following steps are required:

- Configure the key policy for the customer managed key in the production account to allow access to the IAM role of the Lambda function in the development account. A key policy is a resource-based policy that defines who can use or manage a KMS key. To grant cross-account access to a KMS key, you must specify the AWS account ID and the IAM role ARN of the external principal in the key policy statement. For more information, see [Allowing users in other accounts to use a KMS key](#).
- Configure the IAM role for the Lambda function in the development account by attaching an IAM policy that allows access to the customer managed key in the production account. An IAM policy is an identity-based policy that defines what actions an IAM entity can perform on which resources. To allow an IAM role to use



a KMS key in another account, you must specify the KMS key ARN and the kms:Encrypt action (or any other action that requires access to the KMS key) in the IAM policy statement. For more information, see Using IAM policies with AWS KMS.

This solution will meet the requirements of allowing secure access to a KMS customer managed key across AWS accounts.

The other options are incorrect because they either do not grant cross-account access to the KMS key (A, C), or do not use a valid policy type for KMS keys (D).

Verified References:

➤ <https://docs.aws.amazon.com/kms/latest/developerguide/iam-policies.html>

### NEW QUESTION 13

A company stores images for a website in an Amazon S3 bucket. The company is using Amazon CloudFront to serve the images to end users. The company recently discovered that the images are being accessed from countries where the company does not have a distribution license.

Which actions should the company take to secure the images to limit their distribution? (Select TWO.)

- A. Update the S3 bucket policy to restrict access to a CloudFront origin access identity (OAI).
- B. Update the website DNS record to use an Amazon Route 53 geolocation record deny list of countries where the company lacks a license.
- C. Add a CloudFront geo restriction deny list of countries where the company lacks a license.
- D. Update the S3 bucket policy with a deny list of countries where the company lacks a license.
- E. Enable the Restrict Viewer Access option in CloudFront to create a deny list of countries where the company lacks a license.

**Answer:** AC

#### Explanation:

To secure the images to limit their distribution, the company should take the following actions:

- Update the S3 bucket policy to restrict access to a CloudFront origin access identity (OAI). This allows the company to use a special CloudFront user that can access objects in their S3 bucket, and prevent anyone else from accessing them directly.
- Add a CloudFront geo restriction deny list of countries where the company lacks a license. This allows the company to use a feature that controls access to their content based on the geographic location of their viewers, and block requests from countries where they do not have a distribution license.

### NEW QUESTION 15

Amazon GuardDuty has detected communications to a known command and control endpoint from a company's Amazon EC2 instance. The instance was found to be running a vulnerable version of a common web framework. The company's security operations team wants to quickly identify other compute resources with the specific version of that framework installed.

Which approach should the team take to accomplish this task?

- A. Scan all the EC2 instances for noncompliance with IAM Confi
- B. Use Amazon Athena to query IAM CloudTrail logs for the framework installation
- C. Scan all the EC2 instances with the Amazon Inspector Network Reachability rules package to identify instances running a web server with RecognizedPortWithListener findings
- D. Scan all the EC2 instances with IAM Systems Manager to identify the vulnerable version of the web framework
- E. Scan an the EC2 instances with IAM Resource Access Manager to identify the vulnerable version of the web framework

**Answer:** C

#### Explanation:

To quickly identify other compute resources with the specific version of the web framework installed, the team should do the following:

- Scan all the EC2 instances with AWS Systems Manager to identify the vulnerable version of the web framework. This allows the team to use AWS Systems Manager Inventory to collect and query information about the software installed on their EC2 instances, and to filter the results by software name and version.

### NEW QUESTION 17

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 re-quirements? (Select TWO.)

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

**Answer:** AD

### NEW QUESTION 22

A company has a web-based application using Amazon CloudFront and running on Amazon Elastic Container Service (Amazon ECS) behind an Application Load Balancer (ALB). The ALB is terminating TLS and balancing load across ECS service tasks A security engineer needs to design a solution to ensure that application content is accessible only through CloudFront and that I is never accessible directly.

How should the security engineer build the MOST secure solution?

- A. Add an origin custom header Set the viewer protocol policy to HTTP and HTTPS Set the origin protocol pokey to HTTPS only Update the application to validate the CloudFront custom header
- B. Add an origin custom header Set the viewer protocol policy to HTTPS only Set the origin protocol policy to match viewer Update the application to validate the

CloudFront custom header.

C. Add an origin custom header Set the viewer protocol policy to redirect HTTP to HTTPS Set the origin protocol policy to HTTP only Update the application to validate the CloudFront custom header.

D. Add an origin custom header Set the viewer protocol policy to redirect HTTP to HTTP

E. Set the origin protocol policy to HTTPS only Update the application to validate the CloudFront custom header

**Answer: D**

**Explanation:**

To ensure that application content is accessible only through CloudFront and not directly, the security engineer should do the following:

- Add an origin custom header. This is a header that CloudFront adds to the requests that it sends to the origin, but viewers cannot see or modify.
- Set the viewer protocol policy to redirect HTTP to HTTPS. This ensures that the viewers always use HTTPS when they access the website through CloudFront.
- Set the origin protocol policy to HTTPS only. This ensures that CloudFront always uses HTTPS when it connects to the origin.
- Update the application to validate the CloudFront custom header. This means that the application checks if the request has the custom header and only responds if it does. Otherwise, it denies or ignores the request. This prevents users from bypassing CloudFront and accessing the content directly on the origin.

**NEW QUESTION 23**

A security engineer needs to create an Amazon S3 bucket policy to grant least privilege read access to IAM user accounts that are named User=1, User2. and User3. These IAM user accounts are members of the AuthorizedPeople IAM group. The security engineer drafts the following S3 bucket policy:

```
{
  "Version": "2012-10-17",
  "Id": "AuthorizedPeoplePolicy",
  "Statement": [
    {
      "Sid": "Actions-Authorized-People",
      "Effect": "Allow",
      "Action": [
        "s3:GetObject"
      ],
      "Resource": "arn:aws:s3:::authorized-people-bucket/*"
    }
  ]
}
```

When the security engineer tries to add the policy to the S3 bucket, the following error message appears: "Missing required field Principal." The security engineer is adding a Principal element to the policy. The addition must provide read access to only User1. User2, and User3. Which solution meets these requirements?

A)

```
{
  "Principal": {
    "AWS": [
      "arn:aws:iam::1234567890:user/User1",
      "arn:aws:iam::1234567890:user/User2",
      "arn:aws:iam::1234567890:user/User3"
    ]
  }
}
```

B)

```
{
  "Principal": {
    "AWS": [
      "arn:aws:iam::1234567890:root"
    ]
  }
}
```

C)

```
{
  "Principal": {
    "AWS": [
      "*"
    ]
  }
}
```

D)

```
{
  "Principal": {
    "AWS": "arn:aws:iam::1234567890:group/AuthorizedPeople"
  }
}
```

A. Option A

B. Option B

C. Option C

D. Option D

**Answer: A**

**NEW QUESTION 28**

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 32**

Your company has just set up a new central server in a VPC. There is a requirement for other teams who have their servers located in different VPC's in the same region to connect to the central server. Which of the below options is best suited to achieve this requirement. Please select:

- A. Set up VPC peering between the central server VPC and each of the teams VPCs.
- B. Set up IAM DirectConnect between the central server VPC and each of the teams VPCs.
- C. Set up an IPSec Tunnel between the central server VPC and each of the teams VPCs.
- D. None of the above options will work.

**Answer:** A

**Explanation:**

A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them using private IPv4 addresses or IPv6 addresses. Instances in either VPC can communicate with each other as if they are within the same network. You can create a VPC peering connection between your own VPCs, or with a VPC in another IAM account within a single region.

Options B and C are invalid because you need to use VPC Peering Option D is invalid because VPC Peering is available

For more information on VPC Peering please see the below Link:

<http://docs.IAM.amazon.com/AmazonVPC/latest/UserGuide/vpc-peering.html>

The correct answer is: Set up VPC peering between the central server VPC and each of the teams VPCs. Submit your Feedback/Queries to our Experts

**NEW QUESTION 37**

A company hosts a public website on an Amazon EC2 instance. HTTPS traffic must be able to access the website. The company uses SSH for management of the web server.

The website is on the subnet 10.0.1.0/24. The management subnet is 192.168.100.0/24. A security engineer must create a security group for the EC2 instance. Which combination of steps should the security engineer take to meet these requirements in the MOST secure manner? (Select TWO.)

- A. Allow port 22 from source 0.0.0.0/0.
- B. Allow port 443 from source 0.0.0.0/0.
- C. Allow port 22 from 192.168.100.0/24.
- D. Allow port 22 from 10.0.1.0/24.
- E. Allow port 443 from 10.0.1.0/24.

**Answer:** BC

**Explanation:**

The correct answer is B and C.

\* B. Allow port 443 from source 0.0.0.0/0.

This is correct because port 443 is used for HTTPS traffic, which must be able to access the website from any source IP address.

\* C. Allow port 22 from 192.168.100.0/24.

This is correct because port 22 is used for SSH, which is the management protocol for the web server. The management subnet is 192.168.100.0/24, so only this subnet should be allowed to access port 22.

\* A. Allow port 22 from source 0.0.0.0/0.

This is incorrect because it would allow anyone to access port 22, which is a security risk. SSH should be restricted to the management subnet only.

\* D. Allow port 22 from 10.0.1.0/24.

This is incorrect because it would allow the website subnet to access port 22, which is unnecessary and a security risk. SSH should be restricted to the management subnet only.

\* E. Allow port 443 from 10.0.1.0/24.

This is incorrect because it would limit the HTTPS traffic to the website subnet only, which defeats the purpose of having a public website.

**NEW QUESTION 39**

A company is hosting a web application on Amazon EC2 instances behind an Application Load Balancer (ALB). The application has become the target of a DoS attack. Application logging shows that requests are coming from small number of client IP addresses, but the addresses change regularly.

The company needs to block the malicious traffic with a solution that requires the least amount of ongoing effort.

Which solution meets these requirements?

- A. Create an AWS WAF rate-based rule, and attach it to the ALB.
- B. Update the security group that is attached to the ALB to block the attacking IP addresses.
- C. Update the ALB subnet's network ACL to block the attacking client IP addresses.
- D. Create a AWS WAF rate-based rule, and attach it to the security group of the EC2 instances.



**Answer:** A

#### NEW QUESTION 42

A Security Architect has been asked to review an existing security architecture and identify why the application servers cannot successfully initiate a connection to the database servers. The following summary describes the architecture:

- \* 1 An Application Load Balancer, an internet gateway, and a NAT gateway are configured in the public subnet
  - \* 2. Database, application, and web servers are configured on three different private subnets.
  - \* 3 The VPC has two route tables: one for the public subnet and one for all other subnets The route table for the public subnet has a 0 0 0 0/0 route to the internet gateway The route table for all other subnets has a 0 0.0.0/0 route to the NAT gateway. All private subnets can route to each other
  - \* 4 Each subnet has a network ACL implemented that limits all inbound and outbound connectivity to only the required ports and protocols
  - \* 5 There are 3 Security Groups (SGs) database application and web Each group limits all inbound and outbound connectivity to the minimum required
- Which of the following accurately reflects the access control mechanisms the Architect should verify?

- A. Outbound SG configuration on database servers Inbound SG configuration on application servers inbound and outbound network ACL configuration on the database subnet Inbound and outbound network ACL configuration on the application server subnet
- B. Inbound SG configuration on database servers Outbound SG configuration on application servers Inbound and outbound network ACL configuration on the database subnet Inbound and outbound network ACL configuration on the application server subnet
- C. Inbound and outbound SG configuration on database servers Inbound and outbound SG configuration on application servers Inbound network ACL configuration on the database subnet Outbound network ACL configuration on the application server subnet
- D. Inbound SG configuration on database servers Outbound SG configuration on application servers Inbound network ACL configuration on the database subnet Outbound network ACL configuration on the application server subnet.

**Answer:** A

#### Explanation:

this is the accurate reflection of the access control mechanisms that the Architect should verify. Access control mechanisms are methods that regulate who can access what resources and how. Security groups and network ACLs are two types of access control mechanisms that can be applied to EC2 instances and subnets. Security groups are stateful, meaning they remember and return traffic that was previously allowed. Network ACLs are stateless, meaning they do not remember or return traffic that was previously allowed. Security groups and network ACLs can have inbound and outbound rules that specify the source, destination, protocol, and port of the traffic. By verifying the outbound security group configuration on database servers, the inbound security group configuration on application servers, and the inbound and outbound network ACL configuration on both the database and application server subnets, the Architect can check if there are any misconfigurations or conflicts that prevent the application servers from initiating a connection to the database servers. The other options are either inaccurate or incomplete for verifying the access control mechanisms.

#### NEW QUESTION 44

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 45

A company is implementing new compliance requirements to meet customer needs. According to the new requirements the company must not use any Amazon RDS DB instances or DB clusters that lack encryption of the underlying storage. The company needs a solution that will generate an email alert when an unencrypted DB instance or DB cluster is created. The solution also must terminate the unencrypted DB instance or DB cluster.

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

- A. Create an AWS Config managed rule to detect unencrypted RDS storage
- B. Configure an automatic remediation action to publish messages to an Amazon Simple Notification Service (Amazon SNS) topic that includes an AWS Lambda function and an email delivery target as subscriber
- C. Configure the Lambda function to delete the unencrypted resource.
- D. Create an AWS Config managed rule to detect unencrypted RDS storage
- E. Configure a manual remediation action to invoke an AWS Lambda function
- F. Configure the Lambda function to publish messages to an Amazon Simple Notification Service (Amazon SNS) topic and to delete the unencrypted resource.
- G. Create an Amazon EventBridge rule that evaluates RDS event patterns and is initiated by the creation of DB instances or DB clusters Configure the rule to publish messages to an Amazon Simple Notification Service (Amazon SNS) topic that includes an AWS Lambda function and an email delivery target as subscriber
- H. Configure the Lambda function to delete the unencrypted resource.
- I. Create an Amazon EventBridge rule that evaluates RDS event patterns and is initiated by the creation of DB instances or DB cluster
- J. Configure the rule to invoke an AWS Lambda function
- K. Configure the Lambda function to publish messages to an Amazon Simple Notification Service (Amazon SNS) topic and to delete the unencrypted resource.

**Answer:** A



**Explanation:**

<https://docs.aws.amazon.com/config/latest/developerguide/rds-storage-encrypted.html>

**NEW QUESTION 50**

A company uses an Amazon S3 bucket to store reports Management has mandated that all new objects stored in this bucket must be encrypted at rest using server-side encryption with a client-specified IAM Key Management Service (IAM KMS) CMK owned by the same account as the S3 bucket. The IAM account number is 111122223333, and the bucket name is report bucket. The company's security specialist must write the S3 bucket policy to ensure the mandate can be Implemented

Which statement should the security specialist include in the policy?

- A. 

```
{
  "Effect": "Deny",
  "Principal": "*",
  "Action": "s3:PutObject",
  "Resource": "arn:aws:s3:::reportbucket/*",
  "Condition": {
    "StringEquals": {
      "s3:x-amz-server-side-encryption": "AES256"
    }
  }
}
```
- B. 

```
{
  "Effect": "Deny",
  "Principal": "*",
  "Action": "s3:PutObject",
  "Resource": "arn:aws:s3:::reportbucket/*",
  "Condition": {
    "StringNotLike": {
      "s3:x-amz-server-side-encryption-aws-kms-key-id": "arn:aws:kms:*:111122223333:key/*"
    }
  }
}
```
- C. 

```
{
  "Effect": "Deny",
  "Principal": "*",
  "Action": "s3:PutObject",
  "Resource": "arn:aws:s3:::reportbucket/*",
  "Condition": {
    "StringNotLike": {
      "s3:x-amz-server-side-encryption": "aws:kms"
    }
  }
}
```
- D. 

```
{
  "Effect": "Deny",
  "Principal": "*",
  "Action": "s3:PutObject",
  "Resource": "arn:aws:s3:::reportbucket/*",
  "Condition": {
    "StringNotLikeIfExists": {
      "s3:x-amz-server-side-encryption-aws-kms-key-id": "arn:aws:kms:*:111122223333:key/*"
    }
  }
}
```
- E. Option A  
F. Option B  
G. Option C  
H. Option D

**Answer: D**

**NEW QUESTION 51**

A company is using an AWS Key Management Service (AWS KMS) AWS owned key in its application to encrypt files in an AWS account The company's security team wants the ability to change to new key material for new files whenever a potential key breach occurs A security engineer must implement a solution that gives the security team the ability to change the key whenever the team wants to do so

Which solution will meet these requirements?

- A. Create a new customer managed key Add a key rotation schedule to the key Invoke the key rotation schedule every time the security team requests a key change
- B. Create a new AWS managed key Add a key rotation schedule to the key Invoke the key rotation schedule every time the security team requests a key change
- C. Create a key alias Create a new customer managed key every time the security team requests a key change Associate the alias with the new key
- D. Create a key alias Create a new AWS managed key every time the security team requests a key change Associate the alias with the new key

**Answer: A**

**Explanation:**

To meet the requirement of changing the key material for new files whenever a potential key breach occurs, the most appropriate solution would be to create a new customer managed key, add a key rotation schedule to the key, and invoke the key rotation schedule every time the security team requests a key change.

References: : Rotating AWS KMS keys - AWS Key Management Service

**NEW QUESTION 55**

A company plans to use AWS Key Management Service (AWS KMS) to implement an encryption strategy to protect data at rest. The company requires client-side

encryption for company projects. The company is currently conducting multiple projects to test the company's use of AWS KMS. These tests have led to a sudden increase in the company's AWS resource consumption. The test projects include applications that issue multiple requests each second to KMS endpoints for encryption activities.

The company needs to develop a solution that does not throttle the company's ability to use AWS KMS. The solution must improve key usage for client-side encryption and must be cost optimized. Which solution will meet these requirements?

- A. Use keyrings with the AWS Encryption SD
- B. Use each keyring individually or combine keyrings into a multi-keyring
- C. Decrypt the data by using a keyring that has the primary key in the multi-keyring.
- D. Use data key caching
- E. Use the local cache that the AWS Encryption SDK provides with a caching cryptographic materials manager.
- F. Use KMS key rotation
- G. Use a local cache in the AWS Encryption SDK with a caching cryptographic materials manager.
- H. Use keyrings with the AWS Encryption SD
- I. Use each keyring individually or combine keyrings into a multi-keyring
- J. Use any of the wrapping keys in the multi-keyring to decrypt the data.

**Answer: B**

**Explanation:**

The correct answer is B. Use data key caching. Use the local cache that the AWS Encryption SDK provides with a caching cryptographic materials manager. This answer is correct because data key caching can improve performance, reduce cost, and help the company stay within the service limits of AWS KMS. Data key caching stores data keys and related cryptographic material in a cache, and reuses them for encryption and decryption operations. This reduces the number of requests to AWS KMS endpoints and avoids throttling. The AWS Encryption SDK provides a local cache and a caching cryptographic materials manager (caching CMM) that interacts with the cache and enforces security thresholds that the company can set<sup>1</sup>.

The other options are incorrect because:

- A. Using keyrings with the AWS Encryption SDK does not address the problem of throttling or cost optimization. Keyrings are used to generate, encrypt, and decrypt data keys, but they do not cache or reuse them. Using each keyring individually or combining them into a multi-keyring does not reduce the number of requests to AWS KMS endpoints<sup>2</sup>.
- C. Using KMS key rotation does not address the problem of throttling or cost optimization. Key rotation is a security practice that creates new cryptographic material for a KMS key every year, but it does not affect the data that the KMS key protects. Key rotation does not reduce the number of requests to AWS KMS endpoints, and it might incur additional costs for storing multiple versions of key material<sup>3</sup>.
- D. Using keyrings with the AWS Encryption SDK does not address the problem of throttling or cost optimization, as explained in option A. Moreover, using any of the wrapping keys in the multi-keyring to decrypt the data is not a valid option, because only one of the wrapping keys can decrypt a given data key. The wrapping key that encrypts a data key is stored in the encrypted data key structure, and only that wrapping key can decrypt it<sup>4</sup>.

References:

1: Data key caching - AWS Encryption SDK 2: Using keyrings - AWS Encryption SDK 3: Rotating AWS KMS keys - AWS Key Management Service 4: How keyrings work - AWS Encryption SDK

**NEW QUESTION 56**

A company has AWS accounts in an organization in AWS Organizations. The organization includes a dedicated security account.

All AWS account activity across all member accounts must be logged and reported to the dedicated security account. The company must retain all the activity logs in a secure storage location within the dedicated security account for 2 years. No changes or deletions of the logs are allowed.

Which combination of steps will meet these requirements with the LEAST operational overhead? (Select TWO.)

- A. In the dedicated security account, create an Amazon S3 bucket
- B. Configure S3 Object Lock in compliance mode and a retention period of 2 years on the S3 bucket
- C. Set the bucket policy to allow the organization's management account to write to the S3 bucket.
- D. In the dedicated security account, create an Amazon S3 bucket
- E. Configure S3 Object Lock in compliance mode and a retention period of 2 years on the S3 bucket
- F. Set the bucket policy to allow the organization's member accounts to write to the S3 bucket.
- G. In the dedicated security account, create an Amazon S3 bucket that has an S3 Lifecycle configuration that expires objects after 2 year
- H. Set the bucket policy to allow the organization's member accounts to write to the S3 bucket.
- I. Create an AWS CloudTrail trail for the organization
- J. Configure logs to be delivered to the logging Amazon S3 bucket in the dedicated security account.
- K. Turn on AWS CloudTrail in each account
- L. Configure logs to be delivered to an Amazon S3 bucket that is created in the organization's management account
- M. Forward the logs to the S3 bucket in the dedicated security account by using AWS Lambda and Amazon Kinesis Data Firehose.

**Answer: BD**

**Explanation:**

The correct answer is B and D. In the dedicated security account, create an Amazon S3 bucket. Configure S3 Object Lock in compliance mode and a retention period of 2 years on the S3 bucket. Set the bucket policy to allow the organization's member accounts to write to the S3 bucket. Create an AWS CloudTrail trail for the organization. Configure logs to be delivered to the logging Amazon S3 bucket in the dedicated security account.

According to the AWS documentation, AWS CloudTrail is a service that enables governance, compliance, operational auditing, and risk auditing of your AWS account. With CloudTrail, you can log, continuously monitor, and retain account activity related to actions across your AWS infrastructure. CloudTrail provides event history of your AWS account activity, including actions taken through the AWS Management Console, AWS SDKs, command line tools, and other AWS services.

To use CloudTrail with multiple AWS accounts and regions, you need to enable AWS Organizations with all features enabled. This allows you to centrally manage your accounts and apply policies across your organization. You can also use CloudTrail as a service principal for AWS Organizations, which lets you create an organization trail that applies to all accounts in your organization. An organization trail logs events for all AWS Regions and delivers the log files to an S3 bucket that you specify.

To create an organization trail, you need to use an administrator account, such as the organization's management account or a delegated administrator account. You can then configure the trail to deliver logs to an S3 bucket in the dedicated security account. This will ensure that all account activity across all member accounts and regions is logged and reported to the security account.

According to the AWS documentation, Amazon S3 is an object storage service that offers scalability, data availability, security, and performance. You can use S3 to store and retrieve any amount of data from anywhere on the web. You can also use S3 features such as lifecycle management, encryption, versioning, and replication to optimize your storage.

To use S3 with CloudTrail logs, you need to create an S3 bucket in the dedicated security account that will store the logs from the organization trail. You can then configure S3 Object Lock on the bucket to prevent objects from being deleted or overwritten for a fixed amount of time or indefinitely. You can also enable

compliance mode on the bucket, which prevents any user, including the root user in your account, from deleting or modifying a locked object until it reaches its retention date.

To set a retention period of 2 years on the S3 bucket, you need to create a default retention configuration for the bucket that specifies a retention mode (either governance or compliance) and a retention period (either a number of days or a date). You can then set the bucket policy to allow the organization's member accounts to write to the S3 bucket. This will ensure that all logs are retained in a secure storage location within the security account for 2 years and no changes or deletions are allowed.

Option A is incorrect because setting the bucket policy to allow the organization's management account to write to the S3 bucket is not sufficient, as it will not grant access to the other member accounts in the organization.

Option C is incorrect because using an S3 Lifecycle configuration that expires objects after 2 years is not secure, as it will allow users to delete or modify objects before they expire.

Option E is incorrect because using Lambda and Kinesis Data Firehose to forward logs from one S3 bucket to another is not necessary, as CloudTrail can directly deliver logs to an S3 bucket in another account. It also introduces additional operational overhead and complexity.

#### NEW QUESTION 59

A company is running workloads in a single IAM account on Amazon EC2 instances and Amazon EMR clusters. A recent security audit revealed that multiple Amazon Elastic Block Store (Amazon EBS) volumes and snapshots are not encrypted.

The company's security engineer is working on a solution that will allow users to deploy EC2 instances and EMR clusters while ensuring that all new EBS volumes and EBS snapshots are encrypted at rest. The solution must also minimize operational overhead.

Which steps should the security engineer take to meet these requirements?

- A. Create an Amazon Event Bridge (Amazon CloudWatch Events) event with an EC2 instance as the source and create volume as the event trigger.
- B. When the event is triggered, invoke an IAM Lambda function to evaluate and notify the security engineer if the EBS volume that was created is not encrypted.
- C. Use a customer managed IAM policy that will verify that the encryption flag of the CreateVolume context is set to true.
- D. Apply this rule to all users.
- E. Create an IAM Config rule to evaluate the configuration of each EC2 instance on creation or modification. Have the IAM Config rule trigger an IAM Lambda function to alert the security team and terminate the instance if the EBS volume is not encrypted.
- F. 5
- G. Use the IAM Management Console or IAM CLI to enable encryption by default for EBS volumes in each IAM Region where the company operates.

**Answer: D**

#### Explanation:

To ensure that all new EBS volumes and EBS snapshots are encrypted at rest and minimize operational overhead, the security engineer should do the following:

➤ Use the AWS Management Console or AWS CLI to enable encryption by default for EBS volumes in each AWS Region where the company operates. This allows the security engineer to automatically encrypt any new EBS volumes and snapshots created from those volumes, without requiring any additional actions from users.

#### NEW QUESTION 62

A company is building an application on AWS that will store sensitive information. The company has a support team with access to the IT infrastructure, including databases. The company's security engineer must introduce measures to protect the sensitive data against any data breach while minimizing management overhead. The credentials must be regularly rotated.

What should the security engineer recommend?

- A. Enable Amazon RDS encryption to encrypt the database and snapshot.
- B. Enable Amazon Elastic Block Store (Amazon EBS) encryption on Amazon EC2 instance.
- C. Include the database credential in the EC2 user data field.
- D. Use an AWS Lambda function to rotate database credential.
- E. Set up TLS for the connection to the database.
- F. Install a database on an Amazon EC2 instance.
- G. Enable third-party disk encryption to encrypt Amazon Elastic Block Store (Amazon EBS) volume.
- H. Store the database credentials in AWS CloudHSM with automatic rotation.
- I. Set up TLS for the connection to the database.
- J. Enable Amazon RDS encryption to encrypt the database and snapshot.
- K. Enable Amazon Elastic Block Store (Amazon EBS) encryption on Amazon EC2 instance.
- L. Store the database credentials in AWS Secrets Manager with automatic rotation.
- M. Set up TLS for the connection to the RDS hosted database.
- N. Set up an AWS CloudHSM cluster with AWS Key Management Service (AWS KMS) to store KMS key.
- O. Set up Amazon RDS encryption using AWS KMS to encrypt the database.
- P. Store the database credentials in AWS Systems Manager Parameter Store with automatic rotation.
- Q. Set up TLS for the connection to the RDS hosted database.

**Answer: C**

#### NEW QUESTION 66

A company wants to monitor the deletion of AWS Key Management Service (AWS KMS) customer managed keys. A security engineer needs to create an alarm that will notify the company before a KMS key is deleted. The security engineer has configured the integration of AWS CloudTrail with Amazon CloudWatch.

What should the security engineer do next to meet these requirements?

- A. Specify the deletion time of the key material during KMS key creation.
- B. Create a custom AWS Config rule to assess the key's scheduled deletion.
- C. Configure the rule to trigger upon a configuration change.
- D. Send a message to an Amazon Simple Notification Service (Amazon SNS) topic if the key is scheduled for deletion.
- E. Create an Amazon EventBridge rule to detect KMS API calls of DeleteAlias.
- F. Create an AWS Lambda function to send an Amazon Simple Notification Service (Amazon SNS) message to the company.
- G. Add the Lambda function as the target of the EventBridge rule.
- H. Create an Amazon EventBridge rule to detect KMS API calls of DisableKey and ScheduleKeyDeletion. Create an AWS Lambda function to send an Amazon Simple Notification Service (Amazon SNS) message to the company.
- I. Add the Lambda function as the target of the EventBridge rule.
- J. Create an Amazon Simple Notification Service (Amazon SNS) policy to detect KMS API calls of RevokeGrant and ScheduleKeyDeletion. Create an AWS



Lambda function to generate the alarm and send the notification to the company.  
K. Add the Lambda function as the target of the SNS policy.

**Answer:** C

**Explanation:**

The AWS documentation states that you can create an Amazon EventBridge rule to detect KMS API calls of DisableKey and ScheduleKeyDeletion. You can then create an AWS Lambda function to send an Amazon Simple Notification Service (Amazon SNS) message to the company. You can add the Lambda function as the target of the EventBridge rule. This method will meet the requirements.

References: : AWS KMS Developer Guide

**NEW QUESTION 68**

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 69**

A security engineer is using AWS Organizations and wants to optimize SCPs. The security engineer needs to ensure that the SCPs conform to best practices. Which approach should the security engineer take to meet this requirement?

- A. Use AWS IAM Access Analyzer to analyze the policies
- B. View the findings from policy validation checks.
- C. Review AWS Trusted Advisor checks for all accounts in the organization.
- D. Set up AWS Audit Manager
- E. Run an assessment for all AWS Regions for all accounts.
- F. Ensure that Amazon Inspector agents are installed on all Amazon EC2 instances in all accounts.

**Answer:** A

**NEW QUESTION 72**

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 76**

A security engineer needs to implement a write-once-read-many (WORM) model for data that a company will store in Amazon S3 buckets. The company uses the S3 Standard storage class for all of its S3 buckets. The security engineer must ensure that objects cannot be overwritten or deleted by any user, including the AWS account root user.

Which solution will meet these requirements?

- A. Create new S3 buckets with S3 Object Lock enabled in compliance mod
- B. Place objects in the S3 buckets.
- C. Use S3 Glacier Vault Lock to attach a Vault Lock policy to new S3 bucket
- D. Wait 24 hours to complete the Vault Lock proces
- E. Place objects in the S3 buckets.
- F. Create new S3 buckets with S3 Object Lock enabled in governance mod
- G. Place objects in the S3 buckets.
- H. Create new S3 buckets with S3 Object Lock enabled in governance mod
- I. Add a legal hold to the S3 bucket
- J. Place objects in the S3 buckets.

**Answer:** A

**NEW QUESTION 81**

An international company wants to combine AWS Security Hub findings across all the company's AWS Regions and from multiple accounts. In addition, the company

wants to create a centralized custom dashboard to correlate these findings with operational data for deeper analysis and insights. The company needs an analytics tool to search and visualize Security Hub findings. Which combination of steps will meet these requirements? (Select THREE.)

- A. Designate an AWS account as a delegated administrator for Security Hu
- B. Publish events to Amazon CloudWatch from the delegated administrator account, all member accounts, and required Regions that are enabled for Security Hub findings.
- C. Designate an AWS account in an organization in AWS Organizations as a delegated administrator for Security Hu
- D. Publish events to Amazon EventBridge from the delegated administrator account, all member accounts, and required Regions that are enabled for Security Hub findings.
- E. In each Region, create an Amazon EventBridge rule to deliver findings to an Amazon Kinesis data strea
- F. Configure the Kinesis data streams to output the logs to a single Amazon S3 bucket.
- G. In each Region, create an Amazon EventBridge rule to deliver findings to an Amazon Kinesis Data Firehose delivery strea
- H. Configure the Kinesis Data Firehose delivery streams to deliver the logs to a single Amazon S3 bucket.
- I. Use AWS Glue DataBrew to crawl the Amazon S3 bucket and build the schem
- J. Use AWS Glue Data Catalog to query the data and create views to flatten nested attribute
- K. Build Amazon QuickSight dashboards by using Amazon Athena.
- L. Partition the Amazon S3 dat
- M. Use AWS Glue to crawl the S3 bucket and build the schem
- N. Use Amazon Athena to query the data and create views to flatten nested attribute
- O. Build Amazon QuickSight dashboards that use the Athena views.

**Answer:** BDF

**Explanation:**

The correct answer is B, D, and F. Designate an AWS account in an organization in AWS Organizations as a delegated administrator for Security Hub. Publish events to Amazon EventBridge from the delegated administrator account, all member accounts, and required Regions that are enabled for Security Hub findings. In each Region, create an Amazon EventBridge rule to deliver findings to an Amazon Kinesis Data Firehose delivery stream. Configure the Kinesis Data Firehose delivery streams to deliver the logs to a single Amazon S3 bucket. Partition the Amazon S3 data. Use AWS Glue to crawl the S3 bucket and build the schema. Use Amazon Athena to query the data and create views to flatten nested attributes. Build Amazon QuickSight dashboards that use the Athena views.

According to the AWS documentation, AWS Security Hub is a service that provides you with a comprehensive view of your security state across your AWS accounts, and helps you check your environment against security standards and best practices. You can use Security Hub to aggregate security findings from various sources, such as AWS services, partner products, or your own applications.

To use Security Hub with multiple AWS accounts and Regions, you need to enable AWS Organizations with all features enabled. This allows you to centrally manage your accounts and apply policies across your organization. You can also use Security Hub as a service principal for AWS Organizations, which lets you designate a delegated administrator account for Security Hub. The delegated administrator account can enable Security Hub automatically in all existing and future accounts in your organization, and can view and manage findings from all accounts.

According to the AWS documentation, Amazon EventBridge is a serverless event bus that makes it easy to connect applications using data from your own applications, integrated software as a service (SaaS) applications, and AWS services. You can use EventBridge to create rules that match events from various sources and route them to targets for processing.

To use EventBridge with Security Hub findings, you need to enable Security Hub as an event source in EventBridge. This will allow you to publish events from Security Hub to EventBridge in the same Region. You can then create EventBridge rules that match Security Hub findings based on criteria such as severity, type, or resource. You can also specify targets for your rules, such as Lambda functions, SNS topics, or Kinesis Data Firehose delivery streams.

According to the AWS documentation, Amazon Kinesis Data Firehose is a fully managed service that delivers real-time streaming data to destinations such as Amazon S3, Amazon Redshift, Amazon Elasticsearch Service (Amazon ES), and Splunk. You can use Kinesis Data Firehose to transform and enrich your data before delivering it to your destination.

To use Kinesis Data Firehose with Security Hub findings, you need to create a Kinesis Data Firehose delivery stream in each Region where you have enabled Security Hub. You can then configure the delivery stream to receive events from EventBridge as a source, and deliver the logs to a single S3 bucket as a destination. You can also enable data transformation or compression on the delivery stream if needed.

According to the AWS documentation, Amazon S3 is an object storage service that offers scalability, data availability, security, and performance. You can use S3 to store and retrieve any amount of data from anywhere on the web. You can also use S3 features such as lifecycle management, encryption, versioning, and replication to optimize your storage.

To use S3 with Security Hub findings, you need to create an S3 bucket that will store the logs from Kinesis Data Firehose delivery streams. You can then partition the data in the bucket by using prefixes such as account ID or Region. This will improve the performance and cost-effectiveness of querying the data.

According to the AWS documentation, AWS Glue is a fully managed extract, transform, and load (ETL) service that makes it easy to prepare and load your data for analytics. You can use Glue to crawl your data sources, identify data formats, and suggest schemas and transformations. You can also use Glue Data Catalog as a central metadata repository for your data assets.

To use Glue with Security Hub findings, you need to create a Glue crawler that will crawl the S3 bucket and build the schema for the data. The crawler will create tables in the Glue Data Catalog that you can query using standard SQL.

According to the AWS documentation, Amazon Athena is an interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL.

Athena is serverless, so there is no infrastructure to manage, and you pay only for the queries that you run. You can use Athena with Glue Data Catalog as a metadata store for your tables.

To use Athena with Security Hub findings, you need to create views in Athena that will flatten nested attributes in the data. For example, you can create views that extract fields such as account ID, Region, resource type, resource ID, finding type, finding title, and finding description from the JSON data. You can then query the views using SQL and join them with other tables if needed.

According to the AWS documentation, Amazon QuickSight is a fast, cloud-powered business intelligence service that makes it easy to deliver insights to everyone in your organization. You can use QuickSight to create and publish interactive dashboards that include machine learning insights. You can also use QuickSight to connect to various data sources, such as Athena, S3, or RDS.

To use QuickSight with Security Hub findings, you need to create QuickSight dashboards that use the Athena views as data sources. You can then visualize and analyze the findings using charts, graphs, maps, or tables. You can also apply filters, calculations, or aggregations to the data. You can then share the dashboards with your users or embed them in your applications.

## NEW QUESTION 82

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 86

A security engineer must use AWS Key Management Service (AWS KMS) to design a key management solution for a set of Amazon Elastic Block Store (Amazon EBS) volumes that contain sensitive data. The solution needs to ensure that the key material automatically expires in 90 days.

Which solution meets these criteria?

- A. A customer managed CMK that uses customer provided key material
- B. A customer managed CMK that uses AWS provided key material
- C. An AWS managed CMK
- D. Operation system-native encryption that uses GnuPG

**Answer:** A

#### Explanation:

<https://awscli.amazonaws.com/v2/documentation/api/latest/reference/kms/import-key-material.html> aws kms import-key-material \

```
--key-id 1234abcd-12ab-34cd-56ef-1234567890ab \  
--encrypted-key-material fileb://EncryptedKeyMaterial.bin \  
--import-token fileb://ImportToken.bin \  
--expiration-model KEY_MATERIAL_EXPIRES \  
--valid-to 2021-09-21T19:00:00Z
```

The correct answer is A. A customer managed CMK that uses customer provided key material.

A customer managed CMK is a KMS key that you create, own, and manage in your AWS account. You have full control over the key configuration, permissions, rotation, and deletion. You can use a customer managed CMK to encrypt and decrypt data in AWS services that are integrated with AWS KMS, such as Amazon EBS<sup>1</sup>.

A customer managed CMK can use either AWS provided key material or customer provided key material. AWS provided key material is generated by AWS KMS and never leaves the service unencrypted. Customer provided key material is generated outside of AWS KMS and imported into a customer managed CMK. You can specify an expiration date for the imported key material, after which the CMK becomes unusable until you reimport new key material<sup>2</sup>.

To meet the criteria of automatically expiring the key material in 90 days, you need to use customer provided key material and set the expiration date accordingly. This way, you can ensure that the data encrypted with the CMK will not be accessible after 90 days unless you reimport new key material and re-encrypt the data.

The other options are incorrect for the following reasons:

- \* B. A customer managed CMK that uses AWS provided key material does not expire automatically. You can enable automatic rotation of the key material every year, but this does not prevent access to the data encrypted with the previous key material. You would need to manually delete the CMK and its backing key material to make the data inaccessible<sup>3</sup>.
- \* C. An AWS managed CMK is a KMS key that is created, owned, and managed by an AWS service on your behalf. You have limited control over the key configuration, permissions, rotation, and deletion. You cannot use an AWS managed CMK to encrypt data in other AWS services or applications. You also cannot set an expiration date for the key material of an AWS managed CMK<sup>4</sup>.
- \* D. Operation system-native encryption that uses GnuPG is not a solution that uses AWS KMS. GnuPG is a command line tool that implements the OpenPGP standard for encrypting and signing data. It does not integrate with Amazon EBS or other AWS services. It also does not provide a way to automatically expire the key material used for encryption<sup>5</sup>.

References:

1: Customer Managed Keys - AWS Key Management Service 2: [Importing Key Material in AWS Key Management Service (AWS KMS) - AWS Key Management Service] 3: [Rotating Customer Master Keys - AWS Key Management Service] 4: [AWS Managed Keys - AWS Key Management Service] 5: The GNU Privacy Guard

#### NEW QUESTION 87

A security engineer has enabled IAM Security Hub in their IAM account, and has enabled the Center for internet Security (CIS) IAM Foundations compliance standard. No evaluation results on compliance are returned in the Security Hub console after several hours. The engineer wants to ensure that Security Hub can evaluate their resources for CIS IAM Foundations compliance.

Which steps should the security engineer take to meet these requirements?

- A. Add full Amazon Inspector IAM permissions to the Security Hub service role to allow it to perform the CIS compliance evaluation
- B. Ensure that IAM Trusted Advisor Is enabled in the account and that the Security Hub service role has permissions to retrieve the Trusted Advisor security-related recommended actions
- C. Ensure that IAM Confi
- D. is enabled in the account, and that the required IAM Config rules have been created for the CIS compliance evaluation
- E. Ensure that the correct trail in IAM CloudTrail has been configured for monitoring by Security Hub and that the Security Hub service role has permissions to perform the GetObject operation on CloudTrails Amazon S3 bucket

**Answer:** C

#### Explanation:

To ensure that Security Hub can evaluate their resources for CIS AWS Foundations compliance, the security engineer should do the following:

- Ensure that AWS Config is enabled in the account. This is a service that enables continuous assessment and audit of your AWS resources for compliance.
- Ensure that the required AWS Config rules have been created for the CIS compliance evaluation. These are rules that represent your desired configuration settings for specific AWS resources or for an entire AWS account.

#### NEW QUESTION 90

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 95

What are the MOST secure ways to protect the AWS account root user of a recently opened AWS account? (Select TWO.)

- A. Use the AWS account root user access keys instead of the AWS Management Console.
- B. Enable multi-factor authentication for the AWS IAM users with the Adminis-tratorAccess managed policy attached to them.
- C. Enable multi-factor authentication for the AWS account root user.
- D. Use AWS KMS to encrypt all AWS account root user and AWS IAM access keys and set automatic rotation to 30 days.
- E. Do not create access keys for the AWS account root user; instead, create AWS IAM users.

**Answer:** CE

#### NEW QUESTION 99

A company is using AWS Organizations to manage multiple accounts. The company needs to allow an IAM user to use a role to access resources that are in another organization's AWS account.

Which combination of steps must the company perform to meet this requirement? (Select TWO.)

- A. Create an identity policy that allows the sts: AssumeRole action in the AWS account that contains the resource
- B. Attach the identity policy to the IAM user.
- C. Ensure that the sts: AssumeRole action is allowed by the SCPs of the organization that owns the resources that the IAM user needs to access.
- D. Create a role in the AWS account that contains the resource
- E. Create an entry in the role's trust policy that allows the IAM user to assume the rol
- F. Attach the trust policy to the role.
- G. Establish a trust relationship between the IAM user and the AWS account that contains the resources.
- H. Create a role in the IAM user's AWS accoun
- I. Create an identity policy that allows the sts: AssumeRole actio
- J. Attach the identity policy to the role.

**Answer:** BC

#### Explanation:

To allow cross-account access to resources using IAM roles, the following steps are required:

- Create a role in the AWS account that contains the resources (the trusting account) and specify the AWS account that contains the IAM user (the trusted account) as a trusted entity in the role's trust policy. This allows users from the trusted account to assume the role and access resources in the trusting account.
- Ensure that the IAM user has permission to assume the role in their own AWS account. This can be done by creating an identity policy that allows the sts:AssumeRole action and attaching it to the IAM user or their group.
- Ensure that there are no service control policies (SCPs) in the organization that owns the resources that deny or restrict access to the sts:AssumeRole action or the role itself. SCPs are applied to all accounts in an organization and can override any permissions granted by IAM policies.

Verified References:

- <https://repost.aws/knowledge-center/cross-account-access-iam>
- [https://docs.aws.amazon.com/organizations/latest/userguide/orgs\\_manage\\_accounts\\_access.html](https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_access.html)
- [https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial\\_cross-account-with-roles.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html)

#### NEW QUESTION 102

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 VP
- 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 subne
- D. Configure the private subnet to use a NAT gatewa
- 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 VP
- G. Configure the private subnet to use an internet gatewa
- H. Schedule the Lambda function to run every 3 months lo rotate the secrets.
- I. Place the RDS instance in a private subnet and an IAM Lambda function inside the VPC in the private subne
- 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 subne
- L. Configure a Secrets Manager interface endpoin
- 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 106**

A security engineer needs to set up an Amazon CloudFront distribution for an Amazon S3 bucket that hosts a static website. The security engineer must allow only specified IP addresses to access the website. The security engineer also must prevent users from accessing the website directly by using S3 URLs. Which solution will meet these requirements?

- A. Generate an S3 bucket policy
- B. Specify cloudfront.amazonaws.com as the principal
- C. Use the aws:SourceIp condition key to allow access only if the request comes from the specified IP addresses.
- D. Create a CloudFront origin access identity (OAI). Create the S3 bucket policy so that only the OAI has access
- E. Create an AWS WAF web ACL and add an IP set rule
- F. Associate the web ACL with the CloudFront distribution.
- G. Implement security groups to allow only the specified IP addresses access and to restrict S3 bucket access by using the CloudFront distribution.
- H. Create an S3 bucket access point to allow access from only the CloudFront distribution
- I. Create an AWS WAF web ACL and add an IP set rule
- J. Associate the web ACL with the CloudFront distribution.

**Answer: B**

**NEW QUESTION 110**

A company is running an Amazon RDS for MySQL DB instance in a VPC. The VPC must not send or receive network traffic through the internet.

A security engineer wants to use AWS Secrets Manager to rotate the DB instance credentials automatically. Because of a security policy, the security engineer cannot use the standard AWS Lambda function that Secrets Manager provides to rotate the credentials.

The security engineer deploys a custom Lambda function in the VPC. The custom Lambda function will be responsible for rotating the secret in Secrets Manager.

The security engineer edits the DB instance's security group to allow connections from this function. When the function is invoked, the function cannot communicate with Secrets Manager to rotate the secret properly.

What should the security engineer do so that the function can rotate the secret?

- A. Add an egress-only internet gateway to the VPC
- B. Allow only the Lambda function's subnet to route traffic through the egress-only internet gateway.
- C. Add a NAT gateway to the VPC
- D. Configure only the Lambda function's subnet with a default route through the NAT gateway.
- E. Configure a VPC peering connection to the default VPC for Secrets Manager
- F. Configure the Lambda function's subnet to use the peering connection for routes.
- G. Configure a Secrets Manager interface VPC endpoint
- H. Include the Lambda function's private subnet during the configuration process.

**Answer: D**

**Explanation:**

You can establish a private connection between your VPC and Secrets Manager by creating an interface VPC endpoint. Interface endpoints are powered by AWS PrivateLink, a technology that enables you to privately access Secrets Manager APIs without an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection. Reference:

<https://docs.aws.amazon.com/secretsmanager/latest/userguide/vpc-endpoint-overview.html>

The correct answer is D. Configure a Secrets Manager interface VPC endpoint. Include the Lambda function's private subnet during the configuration process.

A Secrets Manager interface VPC endpoint is a private connection between the VPC and Secrets Manager that does not require an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection<sup>1</sup>. By configuring a Secrets Manager interface VPC endpoint, the security engineer can enable the custom Lambda function to communicate with Secrets Manager without sending or receiving network traffic through the internet. The security engineer must include the Lambda function's private subnet during the configuration process to allow the function to use the endpoint<sup>2</sup>.

The other options are incorrect for the following reasons:

- A. An egress-only internet gateway is a VPC component that allows outbound communication over IPv6 from instances in the VPC to the internet, and prevents the internet from initiating an IPv6 connection with the instances<sup>3</sup>. However, this option does not meet the requirement that the VPC must not send or receive network traffic through the internet. Moreover, an egress-only internet gateway is for use with IPv6 traffic only, and Secrets Manager does not support IPv6 addresses<sup>2</sup>.
- B. A NAT gateway is a VPC component that enables instances in a private subnet to connect to the internet or other AWS services, but prevents the internet from initiating connections with those instances<sup>4</sup>. However, this option does not meet the requirement that the VPC must not send or receive network traffic through the internet. Additionally, a NAT gateway requires an elastic IP address, which is a public IPv4 address<sup>4</sup>.
- C. A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them using private IPv4 addresses or IPv6 addresses<sup>5</sup>. However, this option does not work because Secrets Manager does not have a default VPC that can be peered with. Furthermore, a VPC peering connection does not provide a private connection to Secrets Manager APIs without an internet gateway or other devices<sup>2</sup>.

**NEW QUESTION 111**

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 112**

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 114**

An IAM user receives an Access Denied message when the user attempts to access objects in an Amazon S3 bucket. The user and the S3 bucket are in the same AWS account. The S3 bucket is configured to use server-side encryption with AWS KMS keys (SSE-KMS) to encrypt all of its objects at rest by using a customer managed key from the same AWS account. The S3 bucket has no bucket policy defined. The IAM user has been granted permissions through an IAM policy that allows the kms:Decrypt permission to the customer managed key. The IAM policy also allows the s3:List\* and s3:Get\* permissions for the S3 bucket and its objects.

Which of the following is a possible reason that the IAM user cannot access the objects in the S3 bucket?

- A. The IAM policy needs to allow the kms:DescribeKey permission.
- B. The S3 bucket has been changed to use the AWS managed key to encrypt objects at rest.
- C. An S3 bucket policy needs to be added to allow the IAM user to access the objects.
- D. The KMS key policy has been edited to remove the ability for the AWS account to have full access to the key.

**Answer:** D

**Explanation:**

The possible reason that the IAM user cannot access the objects in the S3 bucket is D. The KMS key policy has been edited to remove the ability for the AWS account to have full access to the key.

This answer is correct because the KMS key policy is the primary way to control access to the KMS key, and it must explicitly allow the AWS account to have full access to the key. If the KMS key policy has been edited to remove this permission, then the IAM policy that grants kms:Decrypt permission to the IAM user has no effect, and the IAM user cannot decrypt the objects in the S3 bucket.

The other options are incorrect because:

- > A. The IAM policy does not need to allow the kms:DescribeKey permission, because this permission is not required for decrypting objects in S3 using SSE-KMS. The kms:DescribeKey permission allows getting information about a KMS key, such as its creation date, description, and key state.
- > B. The S3 bucket has not been changed to use the AWS managed key to encrypt objects at rest, because this would not cause an Access Denied message for the IAM user. The AWS managed key is a default KMS key that is created and managed by AWS for each AWS account and Region. The IAM user does not need any permissions on this key to use it for SSE-KMS.
- > C. An S3 bucket policy does not need to be added to allow the IAM user to access the objects, because the IAM user already has s3:List\* and s3:Get\* permissions for the S3 bucket and its objects through an IAM policy. An S3 bucket policy is an optional way to grant cross-account access or public access to an S3 bucket.

References:

1: Key policies in AWS KMS 2: Using server-side encryption with AWS KMS keys (SSE-KMS) 3: AWS KMS API Permissions Reference 4: Using server-side encryption with Amazon S3 managed keys (SSE-S3) 5: Bucket policy examples

**NEW QUESTION 118**

A developer at a company uses an SSH key to access multiple Amazon EC2 instances. The company discovers that the SSH key has been posted on a public GitHub repository. A security engineer verifies that the key has not been used recently.

How should the security engineer prevent unauthorized access to the EC2 instances?

- A. Delete the key pair from the EC2 console.
- B. Create a new key pair.
- C. Use the ModifyInstanceAttribute API operation to change the key on any EC2 instance that is using the key.
- D. Restrict SSH access in the security group to only known corporate IP addresses.
- E. Update the key pair in any AMI that is used to launch the EC2 instance.
- F. Restart the EC2 instances.

**Answer:** C

**Explanation:**

To prevent unauthorized access to the EC2 instances, the security engineer should do the following:

➤ Restrict SSH access in the security group to only known corporate IP addresses. This allows the security engineer to use a virtual firewall that controls inbound and outbound traffic for their EC2 instances, and limit SSH access to only trusted sources.

**NEW QUESTION 119**

A company's IAM account consists of approximately 300 IAM users. Now there is a mandate that an access change is required for 100 IAM users to have unlimited privileges to S3. As a system administrator, how can you implement this effectively so that there is no need to apply the policy at the individual user level? Please select:

- A. Create a new role and add each user to the IAM role
- B. Use the IAM groups and add users, based upon their role, to different groups and apply the policy to group
- C. Create a policy and apply it to multiple users using a JSON script
- D. Create an S3 bucket policy with unlimited access which includes each user's IAM account ID

**Answer: B**

**Explanation:**

Option A is incorrect since you don't add a user to the IAM Role Option C is incorrect since you don't assign multiple users to a policy Option D is incorrect since this is not an ideal approach

An IAM group is used to collectively manage users who need the same set of permissions. By having groups, it becomes easier to manage permissions. So if you change the permissions on the group scale, it will affect all the users in that group

For more information on IAM Groups, just browse to the below URL: [https://docs.IAM.amazon.com/IAM/latest/UserGuide/id\\_eroups.html](https://docs.IAM.amazon.com/IAM/latest/UserGuide/id_eroups.html)

The correct answer is: Use the IAM groups and add users, based upon their role, to different groups and apply the policy to group

Submit your Feedback/Queries to our Experts

**NEW QUESTION 122**

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.amazon.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 127**

An organization wants to log all IAM API calls made within all of its IAM accounts, and must have a central place to analyze these logs. What steps should be taken to meet these requirements in the MOST secure manner? (Select TWO)

- A. Turn on IAM CloudTrail in each IAM account
- B. Turn on CloudTrail in only the account that will be storing the logs
- C. Update the bucket ACL of the bucket in the account that will be storing the logs so that other accounts can log to it
- D. Create a service-based role for CloudTrail and associate it with CloudTrail in each account
- E. Update the bucket policy of the bucket in the account that will be storing the logs so that other accounts can log to it

**Answer: AE**

**Explanation:**

these are the steps that can meet the requirements in the most secure manner. CloudTrail is a service that records AWS API calls and delivers log files to an S3 bucket. Turning on CloudTrail in each IAM account can help capture all IAM API calls made within those accounts. Updating the bucket policy of the bucket in the account that will be storing the logs can help grant other accounts permission to write log files to that bucket. The other options are either unnecessary or insecure for logging and analyzing IAM API calls.

**NEW QUESTION 129**

A Security Engineer has been tasked with enabling IAM Security Hub to monitor Amazon EC2 instances for CVE in a single IAM account. The Engineer has already enabled IAM Security Hub and Amazon Inspector in the IAM Management Console and has installed the Amazon Inspector agent on an EC2 instances that need to be monitored.

Which additional steps should the Security Engineer take to meet this requirement?

- A. Configure the Amazon Inspector agent to use the CVE rule package
- B. Configure the Amazon Inspector agent to use the CVE rule package Configure Security Hub to ingest from IAM Inspector by writing a custom resource policy
- C. Configure the Security Hub agent to use the CVE rule package Configure IAM Inspector to ingest from Security Hub by writing a custom resource policy
- D. Configure the Amazon Inspector agent to use the CVE rule package Install an additional Integration library Allow the Amazon Inspector agent to communicate with Security Hub

**Answer: D**

**Explanation:**

you need to configure the Amazon Inspector agent to use the CVE rule package, which is a set of rules that check for vulnerabilities and exposures on your EC2



instances5. You also need to install an additional integration library that enables communication between the Amazon Inspector agent and Security Hub6. Security Hub is a service that provides you with a comprehensive view of your security state in AWS and helps you check your environment against security industry standards and best practices7. The other options are either incorrect or incomplete for meeting the requirement.

**NEW QUESTION 134**

A security engineer wants to evaluate configuration changes to a specific AWS resource to ensure that the resource meets compliance standards. However, the security engineer is concerned about a situation in which several configuration changes are made to the resource in quick succession. The security engineer wants to record only the latest configuration of that resource to indicate the cumulative impact of the set of changes.

Which solution will meet this requirement in the MOST operationally efficient way?

- A. Use AWS CloudTrail to detect the configuration changes by filtering API calls to monitor the changes. Use the most recent API call to indicate the cumulative impact of multiple calls
- B. Use AWS Config to detect the configuration changes and to record the latest configuration in case of multiple configuration changes.
- C. Use Amazon CloudWatch to detect the configuration changes by filtering API calls to monitor the change
- D. Use the most recent API call to indicate the cumulative impact of multiple calls.
- E. Use AWS Cloud Map to detect the configuration change
- F. Generate a report of configuration changes from AWS Cloud Map to track the latest state by using a sliding time window.

**Answer:** B

**Explanation:**

AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. AWS Config continuously monitors and records your AWS resource configurations and allows you to automate the evaluation of recorded configurations against desired configurations.

To evaluate configuration changes to a specific AWS resource and ensure that it meets compliance standards, the security engineer should use AWS Config to detect the configuration changes and to record the latest configuration in case of multiple configuration changes. This will allow the security engineer to view the current state of the resource and its compliance status, as well as its configuration history and timeline.

AWS Config records configuration changes as ConfigurationItems, which are point-in-time snapshots of the resource's attributes, relationships, and metadata. If multiple configuration changes occur within a short period of time, AWS Config records only the latest ConfigurationItem for that resource. This indicates the cumulative impact of the set of changes on the resource's configuration.

This solution will meet the requirement in the most operationally efficient way, as it leverages AWS Config's features to monitor, record, and evaluate resource configurations without requiring additional tools or services.

The other options are incorrect because they either do not record the latest configuration in case of multiple configuration changes (A, C), or do not use a valid service for evaluating resource configurations (D).

Verified References:

- > <https://docs.aws.amazon.com/config/latest/developerguide/WhatIsConfig.html>
- > <https://docs.aws.amazon.com/config/latest/developerguide/config-item-table.html>

**NEW QUESTION 138**

.....



## Thank You for Trying Our Product

\* **100% Pass or Money Back**

All our products come with a 90-day Money Back Guarantee.

\* **One year free update**

You can enjoy free update one year. 24x7 online support.

\* **Trusted by Millions**

We currently serve more than 30,000,000 customers.

\* **Shop Securely**

All transactions are protected by VeriSign!

**100% Pass Your AWS-Certified-Security-Specialty Exam with Our Prep Materials Via below:**

<https://www.certleader.com/AWS-Certified-Security-Specialty-dumps.html>