

Amazon-Web-Services

Exam Questions SCS-C01

AWS Certified Security- Specialty



NEW QUESTION 1

A company deployed IAM Organizations to help manage its increasing number of IAM accounts. A security engineer wants to ensure only principals in the Organization structure can access a specific Amazon S3 bucket. The solution must also minimize operational overhead. Which solution will meet these requirements?

- A. Put all users into an IAM group with an access policy granting access to the bucket.
- B. Have the account creation trigger an IAM Lambda function that manages the bucket policy, allowing access to accounts listed in the policy only.
- C. Add an SCP to the Organizations master account, allowing all principals access to the bucket.
- D. Specify the organization ID in the global key condition element of a bucket policy, allowing all principals access.

Answer: D

NEW QUESTION 2

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 notices.
- 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 3

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 4

A large corporation is creating a multi-account strategy and needs to determine how its employees should access the IAM infrastructure. Which of the following solutions would provide the MOST scalable solution?

- A. Create dedicated IAM users within each IAM account that employees can assume through federation based upon group membership in their existing identity provider
- B. Use a centralized account with IAM roles that employees can assume through federation with their existing identity provider Use cross-account roles to allow the federated users to assume their target role in the resource accounts.
- C. Configure the IAM Security Token Service to use Kerberos tokens so that users can use their existing corporate user names and passwords to access IAM resources directly
- D. Configure the IAM trust policies within each account's role to set up a trust back to the corporation's existing identity provider allowing users to assume the role based off their SAML token

Answer: B

Explanation:

the most scalable solution for accessing the IAM infrastructure in a multi-account strategy. A multi-account strategy is a way of organizing your AWS resources into multiple IAM accounts for security, billing, and management purposes. Federation is a process that allows users to access AWS resources using credentials from an external identity provider such as Active Directory or SAML. IAM roles are sets of permissions that grant access to AWS resources. Cross-account roles are IAM roles that allow users in one account to access resources in another account. By using a centralized account with IAM roles that employees can assume through federation with their existing identity provider, you can simplify and streamline the access management process. By using cross-account roles to allow the federated users to assume their target role in the resource accounts, you can enable granular and flexible access control across multiple accounts. The other options are either less scalable or less secure for accessing the IAM infrastructure in a multi-account strategy.

NEW QUESTION 5

A company uses SAML federation to grant users access to AWS accounts. A company workload that is in an isolated AWS account runs on immutable infrastructure with no human access to Amazon EC2. The company requires a specialized user known as a break glass user to have access to the workload AWS account and instances in the case of SAML errors. A recent audit discovered that the company did not create the break glass user for the AWS account that contains the workload.

The company must create the break glass user. The company must log any activities of the break glass user and send the logs to a security team. Which combination of solutions will meet these requirements? (Select TWO.)

- A. Create a local individual break glass IAM user for the security tea
- B. Create a trail in AWS CloudTrail that has Amazon CloudWatch Logs turned o
- C. Use Amazon EventBridge to monitor local user activities.
- D. Create a break glass EC2 key pair for the AWS account
- E. Provide the key pair to the security tea
- F. Use AWS CloudTrail to monitor key pair activit
- G. Send notifications to the security team by using Amazon Simple Notification Service (Amazon SNS).
- H. Create a break glass IAM role for the account
- I. Allow security team members to perform the AssumeRoleWithSAML operatio
- J. Create an AWS Cloud Trail trail that has Amazon CloudWatch Logs turned o
- K. Use Amazon EventBridge to monitor security team activities.
- L. Create a local individual break glass IAM user on the operating system level of each workload instance. Configure unrestricted security groups on the instances to grant access to the break glass IAM users.
- M. Configure AWS Systems Manager Session Manager for Amazon EC2. Configure an AWS Cloud Trail filter based on Session Manage
- N. Send the results to an Amazon Simple Notification Service (Amazon SNS) topic.

Answer: AE

Explanation:

The combination of solutions that will meet the requirements are:

- > A. Create a local individual break glass IAM user for the security team. Create a trail in AWS CloudTrail that has Amazon CloudWatch Logs turned on. Use Amazon EventBridge to monitor local user activities. This is a valid solution because it allows the security team to access the workload AWS account and instances using a local IAM user that does not depend on SAML federation. It also enables logging and monitoring of the break glass user activities using AWS CloudTrail, Amazon CloudWatch Logs, and Amazon EventBridge123.
- > E. Configure AWS Systems Manager Session Manager for Amazon EC2. Configure an AWS CloudTrail filter based on Session Manager. Send the results to an Amazon Simple Notification Service (Amazon SNS) topic. This is a valid solution because it allows the security team to access the workload instances without opening any inbound ports or managing SSH keys or bastion hosts. It also enables logging and notification of the break glass user activities using AWS CloudTrail, Session Manager, and Amazon SNS456.

The other options are incorrect because:

- B. Creating a break glass EC2 key pair for the AWS account and providing it to the security team is not a valid solution, because it requires opening inbound ports on the instances and managing SSH keys, which increases the security risk and complexity7.
- C. Creating a break glass IAM role for the account and allowing security team members to perform the AssumeRoleWithSAML operation is not a valid solution, because it still depends on SAML federation, which might not work in case of SAML errors8.
- D. Creating a local individual break glass IAM user on the operating system level of each workload instance and configuring unrestricted security groups on the instances to grant access to the break glass IAM users is not a valid solution, because it requires opening inbound ports on the instances and managing multiple local users, which increases the security risk and complexity9.

References:

1: Creating an IAM User in Your AWS Account 2: Creating a Trail - AWS CloudTrail 3: Using Amazon EventBridge with AWS CloudTrail 4: Setting up Session Manager - AWS Systems Manager 5: Logging Session Manager sessions - AWS Systems Manager 6: Amazon Simple Notification Service 7: Connecting to your Linux instance using SSH - Amazon Elastic Compute Cloud 8: AssumeRoleWithSAML - AWS Security Token Service 9: IAM Users - AWS Identity and Access Management

NEW QUESTION 6

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

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

Answer: AD

Explanation:

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

NEW QUESTION 7

A company is designing a multi-account structure for its development teams. The company is using AWS Organizations and AWS Single Sign-On (AWS SSO). The company must implement a solution so that the development teams can use only specific AWS Regions and so that each AWS account allows access to only specific AWS services.

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

- A. Use AWS SSO to set up service-linked roles with IAM policy statements that include the Condition, Resource, and NotAction elements to allow access to only the Regions and services that are needed.
- B. Deactivate AWS Security Token Service (AWS STS) in Regions that the developers are not allowed to use.
- C. Create SCPs that include the Condition, Resource, and NotAction elements to allow access to only the Regions and services that are needed.
- D. For each AWS account, create tailored identity-based policies for AWS SS
- E. Use statements that include the Condition, Resource, and NotAction elements to allow access to only the Regions and services that are needed.

Answer: C

Explanation:

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_syntax.html#scp-eleme

NEW QUESTION 8

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 9

A company has a legacy application that runs on a single Amazon EC2 instance. A security audit shows that the application has been using an IAM access key within its code to access an Amazon S3 bucket that is named DOC-EXAMPLE-BUCKET1 in the same AWS account. This access key pair has the s3:GetObject permission to all objects in only this S3 bucket. The company takes the application offline because the application is not compliant with the company's security policies for accessing other AWS resources from Amazon EC2.

A security engineer validates that AWS CloudTrail is turned on in all AWS Regions. CloudTrail is sending logs to an S3 bucket that is named DOC-EXAMPLE-BUCKET2. This S3 bucket is in the same AWS account as DOC-EXAMPLE-BUCKET1. However, CloudTrail has not been configured to send logs to Amazon CloudWatch Logs.

The company wants to know if any objects in DOC-EXAMPLE-BUCKET1 were accessed with the IAM access key in the past 60 days. If any objects were accessed, the company wants to know if any of the objects that are text files (.txt extension) contained personally identifiable information (PII).

Which combination of steps should the security engineer take to gather this information? (Choose two.)

- A. Configure Amazon Macie to identify any objects in DOC-EXAMPLE-BUCKET1 that contain PII and that were available to the access key.
- B. Use Amazon CloudWatch Logs Insights to identify any objects in DOC-EXAMPLE-BUCKET1 that contain PII and that were available to the access key.
- C. Use Amazon OpenSearch Service (Amazon Elasticsearch Service) to query the CloudTrail logs in DOC-EXAMPLE-BUCKET2 for API calls that used the access key to access an object that contained PII.
- D. Use Amazon Athena to query the CloudTrail logs in DOC-EXAMPLE-BUCKET2 for any API calls that used the access key to access an object that contained PII.
- E. Use AWS Identity and Access Management Access Analyzer to identify any API calls that used the access key to access objects that contained PII in DOC-EXAMPLE-BUCKET1.

Answer: AD

NEW QUESTION 10

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 creatio
- B. Create a custom AWS Config rule to assess the key's scheduleddeletio
- C. Configure the rule to trigger upon a configuration chang
- 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 DeleteAlia
- F. Create an AWS Lambda function to send an Amazon Simple Notification Service (Amazon SNS) message to the compan
- 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 compan
- 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 compan
- 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 10

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

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

Answer: D

Explanation:

The correct answer is D. The AWS Config rules have a MaximumExecutionFrequency value of 24 hours. According to the AWS documentation¹, 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 status².
- 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 operation2.

References:

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

NEW QUESTION 14

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

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

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

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

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

Answer: BE

NEW QUESTION 18

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

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

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

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

Answer: D

Explanation:

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

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

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

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

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

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

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

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

Verified References:

➤ https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_boundaries.html

NEW QUESTION 20

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 in-stances in all accounts.

Answer: A

NEW QUESTION 21

A company used a lift-and-shift approach to migrate from its on-premises data centers to the AWS Cloud. The company migrated on-premises VMS to Amazon EC2 in-stances. Now the company wants to replace some of components that are running on the EC2 instances with managed AWS services that provide similar functionality.

Initially, the company will transition from load balancer software that runs on EC2 instances to AWS Elastic Load Balancers. A security engineer must ensure that after this transition, all the load balancer logs are centralized and searchable for auditing. The security engineer must also ensure that metrics are generated to show which ciphers are in use.

Which solution will meet these requirements?

- A. Create an Amazon CloudWatch Logs log group
- B. Configure the load balancers to send logs to the log group
- C. Use the CloudWatch Logs console to search the log
- D. Create CloudWatch Logs filters on the logs for the required metrics.
- E. Create an Amazon S3 bucket
- F. Configure the load balancers to send logs to the S3 bucket
- G. Use Amazon Athena to search the logs that are in the S3 bucket
- H. Create Amazon CloudWatch filters on the S3 log files for the required metrics.
- I. Create an Amazon S3 bucket
- J. Configure the load balancers to send logs to the S3 bucket
- K. Use Amazon Athena to search the logs that are in the S3 bucket
- L. Create Athena queries for the required metric
- M. Publish the metrics to Amazon CloudWatch.
- N. Create an Amazon CloudWatch Logs log group
- O. Configure the load balancers to send logs to the log group
- P. Use the AWS Management Console to search the log
- Q. Create Amazon Athena queries for the required metric
- R. Publish the metrics to Amazon CloudWatch.

Answer: C

Explanation:

- Amazon S3 is a service that provides scalable, durable, and secure object storage. You can use Amazon S3 to store and retrieve any amount of data from anywhere on the web¹
- AWS Elastic Load Balancing is a service that distributes incoming application or network traffic across multiple targets, such as EC2 instances, containers, or IP addresses. You can use Elastic Load Balancing to increase the availability and fault tolerance of your applications²
- Elastic Load Balancing supports access logging, which captures detailed information about requests sent to your load balancer. Each log contains information such as the time the request was received, the client's IP address, latencies, request paths, and server responses. You can use access logs to analyze traffic patterns and troubleshoot issues³
- You can configure your load balancer to store access logs in an Amazon S3 bucket that you specify. You can also specify the interval for publishing the logs, which can be 5 or 60 minutes. The logs are stored in a hierarchical folder structure by load balancer name, IP address, year, month, day, and time.
- Amazon Athena is a service that allows you to analyze data in Amazon S3 using standard SQL. You can use Athena to run ad-hoc queries and get results in seconds. Athena is serverless, so there is no infrastructure to manage and you pay only for the queries that you run.
- You can use Athena to search the access logs that are stored in your S3 bucket. You can create a table in Athena that maps to your S3 bucket and then run SQL queries on the table. You can also use the Athena console or API to view and download the query results.
- You can also use Athena to create queries for the required metrics, such as the number of requests per cipher or protocol. You can then publish the metrics to Amazon CloudWatch, which is a service that monitors and manages your AWS resources and applications. You can use CloudWatch to collect and track metrics, create alarms, and automate actions based on the state of your resources.
- By using this solution, you can meet the requirements of ensuring that all the load balancer logs are centralized and searchable for auditing and that metrics are generated to show which ciphers are in use.

NEW QUESTION 26

Your CTO is very worried about the security of your IAM account. How best can you prevent hackers from completely hijacking your account? Please select:

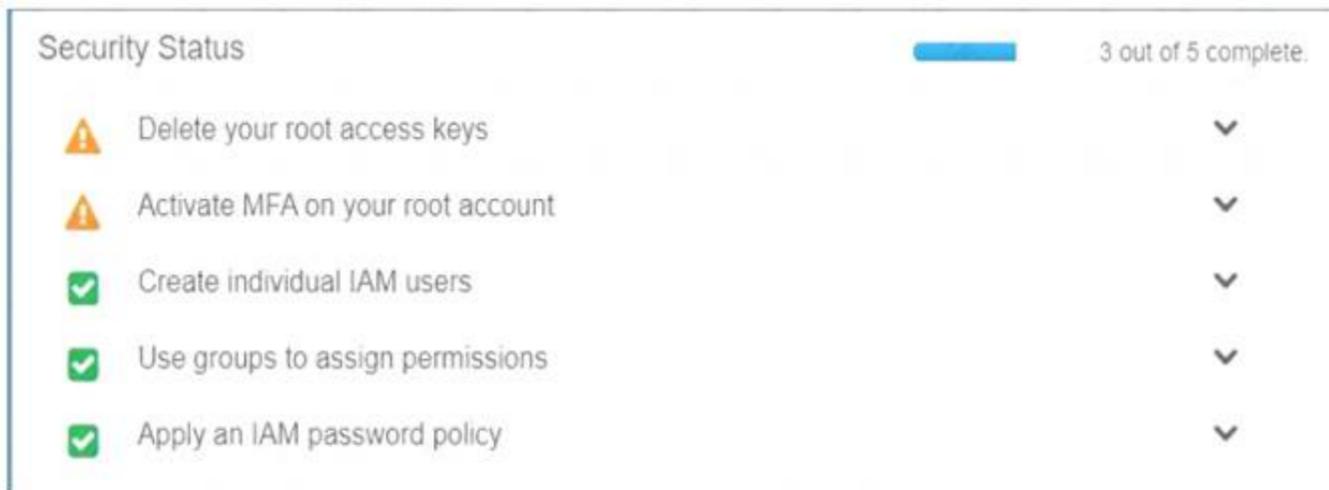
- A. Use short but complex password on the root account and any administrators.
- B. Use IAM IAM Geo-Lock and disallow anyone from logging in except for in your city.
- C. Use MFA on all users and accounts, especially on the root account.
- D. Don't write down or remember the root account password after creating the IAM account.

Answer: C

Explanation:

Multi-factor authentication can add one more layer of security to your IAM account Even when you go to your Security Credentials dashboard one of the items is to enable MFA on your root account

C:\Users\wk\Desktop\mudassar\Untitled.jpg



Option A is invalid because you need to have a good password policy Option B is invalid because there is no IAM Geo-Lock Option D is invalid because this is not a recommended practices For more information on MFA, please visit the below URL

http://docs.IAM.amazonaws.com/IAM/latest/UserGuide/id_credentials_mfa.html

The correct answer is: Use MFA on all users and accounts, especially on the root account. Submit your Feedback/Queries to our Experts

NEW QUESTION 27

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

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

The security team has enabled CloudTrail for AWS management events with a basic configuration in the AWS Region in which EventBridge is being tested. Verification of the EventBridge event pattern indicates that the pattern is set up correctly. The security team must implement a solution so that the s3:PutObjectAcl API call will invoke an EventBridge event. The solution must not generate false notifications.

Which solution will meet these requirements?

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

Answer: D

Explanation:

The correct answer is D. Enable CloudTrail to monitor data events for read and write operations to S3 buckets. According to the AWS documentation¹, CloudTrail data events are the resource operations performed on or within a resource. These are also known as data plane operations. Data events are often high-volume activities. For example, Amazon S3 object-level API activity (such as GetObject, DeleteObject, and PutObject) is a data event.

By default, trails do not log data events. To record CloudTrail data events, you must explicitly add the supported resources or resource types for which you want to collect activity. For more information, see Logging data events in the Amazon S3 User Guide².

In this case, the security team wants EventBridge to watch for the s3:PutObjectAcl API invocation logs from CloudTrail. This API uses the acl subresource to set the access control list (ACL) permissions for a new or existing object in an S3 bucket³. This is a data event that affects the S3 object resource type. Therefore, the security team must enable CloudTrail to monitor data events for read and write operations to S3 buckets in order to invoke an EventBridge event for this API call.

The other options are incorrect because:

- > A. Modifying the EventBridge event pattern by selecting Amazon S3 and All Events as the event type will not capture the s3:PutObjectAcl API call, because this is a data event and not a management event. Management events provide information about management operations that are performed on resources in your AWS account. These are also known as control plane operations⁴.
- > B. Modifying the EventBridge event pattern by selecting Amazon S3 and Bucket Level Operations as the event type will not capture the s3:PutObjectAcl API call, because this is a data event that affects the S3 object resource type and not the S3 bucket resource type. Bucket level operations are management events that affect the configuration or metadata of an S3 bucket⁵.
- > C. Enabling CloudTrail Insights to identify unusual API activity will not help the security team monitor new S3 objects or changes to any S3 bucket policy or setting that result in public access. CloudTrail Insights helps AWS users identify and respond to unusual activity associated with API calls and API error rates by continuously analyzing CloudTrail management events⁶. It does not analyze data events or generate EventBridge events.

References:

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

NEW QUESTION 31

A company uses Amazon GuardDuty. The company's security team wants all High severity findings to automatically generate a ticket in a third-party ticketing system through email integration.

Which solution will meet this requirement?

- A. Create a verified identity for the third-party ticketing email system in Amazon Simple Email Service (Amazon SES). Create an Amazon EventBridge rule that includes an event pattern that matches High severity GuardDuty finding
- B. Specify the SES identity as the target for the EventBridge rule.
- C. Create an Amazon Simple Notification Service (Amazon SNS) topic
- D. Subscribe the third-party ticketing email system to the SNS topic
- E. Create an Amazon EventBridge rule that includes an event pattern that matches High severity GuardDuty finding
- F. Specify the SNS topic as the target for the EventBridge rule.
- G. Use the GuardDuty CreateFilter API operation to build a filter in GuardDuty to monitor for High severity finding
- H. Export the results of the filter to an Amazon Simple Notification Service (Amazon SNS) topic
- I. Subscribe the third-party ticketing email system to the SNS topic.

- J. Use the GuardDuty CreateFilter API operation to build a filter in GuardDuty to monitor for High severity finding
- K. Create an Amazon Simple Notification Service (Amazon SNS) topic
- L. Subscribe the third-party ticketing email system to the SNS topic
- M. Create an Amazon EventBridge rule that includes an event pattern that matches GuardDuty findings that are selected by the filter
- N. Specify the SNS topic as the target for the EventBridge rule.

Answer: B

Explanation:

The correct answer is B. Create an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the third-party ticketing email system to the SNS topic. Create an Amazon EventBridge rule that includes an event pattern that matches High severity GuardDuty findings. Specify the SNS topic as the target for the EventBridge rule.

According to the AWS documentation¹, you can use Amazon EventBridge to create rules that match events from GuardDuty and route them to targets such as Amazon SNS topics. You can use event patterns to filter events based on criteria such as severity, type, or resource. For example, you can create a rule that matches only High severity findings and sends them to an SNS topic that is subscribed by a third-party ticketing email system. This way, you can automate the creation of tickets for High severity findings and notify the security team.

NEW QUESTION 33

A company wants to remove all SSH keys permanently from a specific subset of its Amazon Linux 2 Amazon EC2 instances that are using the same 1AM instance profile. However, three individuals who have IAM user accounts will need to access these instances by using an SSH session to perform critical duties. How can a security engineer provide the access to meet these requirements?

- A. Assign an 1AM policy to the instance profile to allow the EC2 instances to be managed by AWS Systems Manager. Provide the 1AM user accounts with permission to use Systems Manager. Remove the SSH keys from the EC2 instances. Use Systems Manager Inventory to select the EC2 instance and connect.
- B. Assign an 1AM policy to the 1AM user accounts to provide permission to use AWS Systems Manager Run Command. Remove the SSH keys from the EC2 instances. Use Run Command to open an SSH connection to the EC2 instance.
- C. Assign an 1AM policy to the instance profile to allow the EC2 instances to be managed by AWS Systems Manager. Provide the 1AM user accounts with permission to use Systems Manager. Remove the SSH keys from the EC2 instances. Use Systems Manager Session Manager to select the EC2 instance and connect.
- D. Assign an 1AM policy to the 1AM user accounts to provide permission to use the EC2 service in the AWS Management Console. Remove the SSH keys from the EC2 instances. Connect to the EC2 instance as the ec2-user through the AWS Management Console's EC2 SSH client method.

Answer: C

Explanation:

To provide access to the three individuals who have IAM user accounts to access the Amazon Linux 2 Amazon EC2 instances that are using the same IAM instance profile, the most appropriate solution would be to assign an IAM policy to the instance profile to allow the EC2 instances to be managed by AWS Systems Manager, provide the IAM user accounts with permission to use Systems Manager, remove the SSH keys from the EC2 instances, and use Systems Manager Session Manager to select the EC2 instance and connect.

References: : AWS Systems Manager Session Manager - AWS Systems Manager : AWS Systems Manager AWS Management Console : AWS Identity and Access Management - AWS Management Console : Amazon Elastic Compute Cloud - Amazon Web Services : Amazon Linux 2 - Amazon Web Services : AWS Systems Manager - AWS Management Console : AWS Systems Manager - AWS Management Console : AWS Systems Manager - AWS Management Console

NEW QUESTION 36

A company uses AWS Organizations and has production workloads across multiple AWS accounts. A security engineer needs to design a solution that will proactively monitor for suspicious behavior across all the accounts that contain production workloads.

The solution must automate remediation of incidents across the production accounts. The solution also must publish a notification to an Amazon Simple Notification Service (Amazon SNS) topic when a critical security finding is detected. In addition, the solution must send all security incident logs to a dedicated account.

Which solution will meet these requirements?

- A. Activate Amazon GuardDuty in each production account
- B. In a dedicated logging account
- C. aggregate all GuardDuty logs from each production account
- D. Remediate incidents by configuring GuardDuty to directly invoke an AWS Lambda function
- E. Configure the Lambda function to also publish notifications to the SNS topic.
- F. Activate AWS Security Hub in each production account
- G. In a dedicated logging account
- H. aggregate all Security Hub findings from each production account
- I. Remediate incidents by using AWS Config and AWS Systems Manager
- J. Configure Systems Manager to also publish notifications to the SNS topic.
- K. Activate Amazon GuardDuty in each production account
- L. In a dedicated logging account
- M. aggregate all GuardDuty logs from each production account. Remediate incidents by using Amazon EventBridge to invoke a custom AWS Lambda function from the GuardDuty finding
- N. Configure the Lambda function to also publish notifications to the SNS topic.
- O. Activate AWS Security Hub in each production account
- P. In a dedicated logging account
- Q. aggregate all Security Hub findings from each production account
- R. Remediate incidents by using Amazon EventBridge to invoke a custom AWS Lambda function from the Security Hub finding
- S. Configure the Lambda function to also publish notifications to the SNS topic.

Answer: D

Explanation:

The correct answer is D.

To design a solution that will proactively monitor for suspicious behavior across all the accounts that contain production workloads, the security engineer needs to use a service that can aggregate and analyze security findings from multiple sources. AWS Security Hub is a service that provides a comprehensive view of your security posture across your AWS accounts and enables you to check your environment against security standards and best practices. Security Hub also integrates with other AWS services, such as Amazon GuardDuty, AWS Config, and AWS Systems Manager, to collect and correlate security findings.

To automate remediation of incidents across the production accounts, the security engineer needs to use a service that can trigger actions based on events.

Amazon EventBridge is a serverless event bus service that allows you to connect your applications with data from a variety of sources. EventBridge can use rules to match events and route them to targets for processing. You can use EventBridge to invoke a custom AWS Lambda function from the Security Hub findings. Lambda is a serverless compute service that lets you run code without provisioning or managing servers.

To publish a notification to an Amazon SNS topic when a critical security finding is detected, the security engineer needs to use a service that can send messages to subscribers. 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 configure the Lambda function to also publish notifications to the SNS topic.

To send all security incident logs to a dedicated account, the security engineer needs to use a service that can aggregate and store log data from multiple sources. AWS Security Hub allows you to aggregate security findings from multiple accounts into a single account using the delegated administrator feature. This feature enables you to designate an AWS account as the administrator for Security Hub in an organization. The administrator account can then view and manage Security Hub findings from all member accounts.

Therefore, option D is correct because it meets all the requirements of the solution. Option A is incorrect because GuardDuty does not provide a comprehensive view of your security posture across your AWS accounts. GuardDuty is primarily a threat detection service that monitors for malicious or unauthorized behavior. Option B is incorrect because Config and Systems Manager are not designed to automate remediation of incidents based on Security Hub findings. Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources, while Systems Manager is a service that allows you to manage your infrastructure on AWS at scale. Option C is incorrect because GuardDuty does not provide a comprehensive view of your security posture across your AWS accounts.

References:

- > AWS Security Hub
- > Amazon EventBridge
- > AWS Lambda
- > Amazon SNS
- > Aggregating Security Hub findings across accounts

NEW QUESTION 40

A startup company is using a single AWS account that has resources in a single AWS Region. A security engineer configures an AWS Cloud Trail trail in the same Region to deliver log files to an Amazon S3 bucket by using the AWS CLI.

Because of expansion, the company adds resources in multiple Regions. The security engineer notices that the logs from the new Regions are not reaching the S3 bucket.

What should the security engineer do to fix this issue with the LEAST amount of operational overhead?

- A. Create a new CloudTrail trail
- B. Select the new Regions where the company added resources.
- C. Change the S3 bucket to receive notifications to track all actions from all Regions.
- D. Create a new CloudTrail trail that applies to all Regions.
- E. Change the existing CloudTrail trail so that it applies to all Regions.

Answer: D

Explanation:

The correct answer is D. Change the existing CloudTrail trail so that it applies to all Regions.

According to the AWS documentation¹, you can configure CloudTrail to deliver log files from multiple Regions to a single S3 bucket for a single account. To change an existing single-Region trail to log in all Regions, you must use the AWS CLI and add the `--is-multi-region-trail` option to the `update-trail` command². This will ensure that you log global service events and capture all management event activity in your account.

Option A is incorrect because creating a new CloudTrail trail for each Region will incur additional costs and increase operational overhead. Option B is incorrect because changing the S3 bucket to receive notifications will not affect the delivery of log files from other Regions. Option C is incorrect because creating a new CloudTrail trail that applies to all Regions will result in duplicate log files for the original Region and also incur additional costs.

NEW QUESTION 43

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 Config
- 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 48

A company is running an application in The eu-west-1 Region. The application uses an IAM Key Management Service (IAM KMS) CMK to encrypt sensitive data. The company plans to deploy the application in the eu-north-1 Region.

A security engineer needs to implement a key management solution for the application deployment in the new Region. The security engineer must minimize changes to the application code.

Which change should the security engineer make to the IAM KMS configuration to meet these requirements?

- A. Update the key policies in eu-west-1. Point the application in eu-north-1 to use the same CMK as the application in eu-west-1.
- B. Allocate a new CMK to eu-north-1 to be used by the application that is deployed in that Region.
- C. Allocate a new CMK to eu-north-1. Create the same alias name for both key
- D. Configure the application deployment to use the key alias.
- E. Allocate a new CMK to eu-north-1. Create an alias for eu-'-1. Change the application code to point to the alias for eu-'-1.

Answer: B

NEW QUESTION 52

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

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

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

Answer: CE

Explanation:

AWS Secrets Manager is a service that helps you manage, retrieve, and rotate secrets such as database credentials, API keys, and other sensitive information. By configuring automatic rotation of credentials in AWS Secrets Manager, you can ensure that your secrets are changed regularly and securely, without requiring manual intervention or application downtime. You can also specify the rotation frequency and the rotation function that performs the logic of changing the credentials on the database and updating the secret in Secrets Manager¹.

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

By configuring the Java application to catch a connection failure and make a call to AWS Secrets Manager to retrieve updated credentials, you can avoid hard-coding the credentials in your application code or configuration files. This way, your application can dynamically obtain the latest credentials from Secrets Manager whenever the password is rotated, without needing to restart or redeploy the application. To enable this, you need to grant permission to the instance role associated with the EC2 instance to access Secrets Manager using IAM policies². You can also use the AWS SDK for Java to integrate your application with Secrets Manager³.

NEW QUESTION 54

A company needs a security engineer to implement a scalable solution for multi-account authentication and authorization. The solution should not introduce additional user-managed architectural components. Native IAM features should be used as much as possible. The security engineer has set up IAM Organizations with all features activated and IAM SSO enabled.

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

- A. Use AD Connector to create users and groups for all employees that require access to IAM accounts. Assign AD Connector groups to IAM accounts and link to the IAM roles in accordance with the employees' job functions and access requirements. Instruct employees to access IAM accounts by using the IAM Directory Service user portal.
- B. Use an IAM SSO default directory to create users and groups for all employees that require access to IAM account
- C. Assign groups to IAM accounts and link to permission sets in accordance with the employees' job functions and access requirement
- D. Instruct employees to access IAM accounts by using the IAM SSO user portal.
- E. Use an IAM SSO default directory to create users and groups for all employees that require access to IAM account
- F. Link IAM SSO groups to the IAM users present in all accounts to inherit existing permission
- G. Instruct employees to access IAM accounts by using the IAM SSO user portal.
- H. Use IAM Directory Service for Microsoft Active Directory to create users and groups for all employees that require access to IAM accounts. Enable IAM Management Console access in the created directory and specify IAM SSO as a source of information for integrated accounts and permission set
- I. Instruct employees to access IAM accounts by using the IAM Directory Service user portal.

Answer: B

NEW QUESTION 57

A company's security engineer is developing an incident response plan to detect suspicious activity in an AWS account for VPC hosted resources. The security engineer needs to provide visibility for as many AWS Regions as possible.

Which combination of steps will meet these requirements MOST cost-effectively? (Select TWO.)

- A. Turn on VPC Flow Logs for all VPCs in the account.
- B. Activate Amazon GuardDuty across all AWS Regions.
- C. Activate Amazon Detective across all AWS Regions.
- D. Create an Amazon Simple Notification Service (Amazon SNS) topic
- E. Create an Amazon EventBridge rule that responds to findings and publishes the findings to the SNS topic.
- F. Create an AWS Lambda function
- G. Create an Amazon EventBridge rule that invokes the Lambda function to publish findings to Amazon Simple Email Service (Amazon SES).

Answer: BD

Explanation:

To detect suspicious activity in an AWS account for VPC hosted resources, the security engineer needs to use a service that can monitor network traffic and API calls across all AWS Regions. Amazon GuardDuty is a threat detection service that can do this by analyzing VPC Flow Logs, AWS CloudTrail event logs, and DNS logs. By activating GuardDuty across all AWS Regions, the security engineer can provide visibility for as many regions as possible. GuardDuty generates findings that contain details about the potential threats detected in the account. To respond to these findings, the security engineer needs to create a mechanism that can

notify the relevant stakeholders or take remedial actions. One way to do this is to use Amazon EventBridge, which is a serverless event bus service that can connect AWS services and third-party applications. By creating an EventBridge rule that responds to GuardDuty findings and publishes them to an Amazon Simple Notification Service (Amazon SNS) topic, the security engineer can enable subscribers of the topic to receive notifications via email, SMS, or other methods. This is a cost-effective solution that does not require any additional infrastructure or code.

NEW QUESTION 61

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 66

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 account
- 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/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

NEW QUESTION 69

A company accidentally deleted the private key for an Amazon Elastic Block Store (Amazon EBS)-backed Amazon EC2 instance. A security engineer needs to regain access to the instance.

Which combination of steps will meet this requirement? (Choose two.)

- A. Stop the instanc
- B. Detach the root volum
- C. Generate a new key pair.
- D. Keep the instance runnin
- E. Detach the root volum
- F. Generate a new key pair.
- G. When the volume is detached from the original instance, attach the volume to another instance as a data volum
- H. Modify the authorized_keys file with a new public ke
- I. Move the volume back to the original instanc
- J. Start the instance.
- K. When the volume is detached from the original instance, attach the volume to another instance as a data volum
- L. Modify the authorized_keys file with a new private ke
- M. Move the volume back to the original instanc
- N. Start the instance.
- O. When the volume is detached from the original instance, attach the volume to another instance as a data volum
- P. Modify the authorized_keys file with a new public ke
- Q. Move the volume back to the original instance that is running.

Answer: AC

Explanation:

If you lose the private key for an EBS-backed instance, you can regain access to your instance. You must stop the instance, detach its root volume and attach it to another instance as a data volume, modify the authorized_keys file with a new public key, move the volume back to the original instance, and restart the instance.

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/TroubleshootingInstancesConnecting.html#replacing>

NEW QUESTION 72

A company needs to follow security best practices to deploy resources from an AWS CloudFormation template. The CloudFormation template must be able to configure sensitive database credentials.

The company already uses AWS Key Management Service (AWS KMS) and AWS Secrets Manager. Which solution will meet the requirements?

- A. Use a dynamic reference in the CloudFormation template to reference the database credentials in Secrets Manager.
- B. Use a parameter in the CloudFormation template to reference the database credential
- C. Encrypt the CloudFormation template by using AWS KMS.
- D. Use a SecureString parameter in the CloudFormation template to reference the database credentials in Secrets Manager.
- E. Use a SecureString parameter in the CloudFormation template to reference an encrypted value in AWS KMS

Answer: A

Explanation:

➤ Option A: This option meets the requirements of following security best practices and configuring sensitive database credentials in the CloudFormation template. A dynamic reference is a way to specify external values that are stored and managed in other services, such as Secrets Manager, in the stack templates¹. When using a dynamic reference, CloudFormation retrieves the value of the specified reference when necessary during stack and change set operations¹. Dynamic references can be used for certain resources that support them, such as `AWS::RDS::DBInstance`¹. By using a dynamic reference to reference the database credentials in Secrets Manager, the company can leverage the existing integration between these services and avoid hardcoding the secret information in the template. Secrets Manager is a service that helps you protect secrets needed to access your applications, services, and IT resources². Secrets Manager enables you to rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle².

NEW QUESTION 73

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 VP
- B. Allow only the Lambda function's subnet to route traffic through the egress-only internet gateway.
- C. Add a NAT gateway to the VP
- 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 Manage
- 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 connection¹. 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 endpoint².

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 instances³. 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 addresses².
- 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 instances⁴. 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 address⁴.
- 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 addresses⁵. 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 devices².

NEW QUESTION 75

A company is using AWS Organizations to implement a multi-account strategy. The company does not have on-premises infrastructure. All workloads run on AWS. The company currently has eight member accounts. The company anticipates that it will have no more than 20 AWS accounts total at any time.

The company issues a new security policy that contains the following requirements:

- No AWS account should use a VPC within the AWS account for workloads.
- The company should use a centrally managed VPC that all AWS accounts can access to launch workloads in subnets.
- No AWS account should be able to modify another AWS account's application resources within the centrally managed VPC.
- The centrally managed VPC should reside in an existing AWS account that is named Account-A within an organization.

The company uses an AWS CloudFormation template to create a VPC that contains multiple subnets in Account-A. This template exports the subnet IDs through the CloudFormation Outputs section.

Which solution will complete the security setup to meet these requirements?

- A. Use a CloudFormation template in the member accounts to launch workload
- B. Configure the template to use the `Fn::ImportValue` function to obtain the subnet ID values.

- C. Use a transit gateway in the VPC within Account-
- D. Configure the member accounts to use the transit gateway to access the subnets in Account-A to launch workloads.
- E. Use AWS Resource Access Manager (AWS RAM) to share Account-A's VPC subnets with the remaining member account
- F. Configure the member accounts to use the shared subnets to launch workloads.
- G. Create a peering connection between Account-A and the remaining member account
- H. Configure the member accounts to use the subnets in Account-A through the VPC peering connection to launch workloads.

Answer: C

Explanation:

The correct answer is C. Use AWS Resource Access Manager (AWS RAM) to share Account-A's VPC subnets with the remaining member accounts. Configure the member accounts to use the shared subnets to launch workloads.

This answer is correct because AWS RAM is a service that helps you securely share your AWS resources across AWS accounts, within your organization or organizational units (OUs), and with IAM roles and users for supported resource types¹. One of the supported resource types is VPC subnets², which means you can share the subnets in Account-A's VPC with the other member accounts using AWS RAM. This way, you can meet the requirements of using a centrally managed VPC, avoiding duplicate VPCs in each account, and launching workloads in shared subnets. You can also control the access to the shared subnets by using IAM policies and resource-based policies³, which can prevent one account from modifying another account's resources.

The other options are incorrect because:

- > A. Using a CloudFormation template in the member accounts to launch workloads and using the Fn::ImportValue function to obtain the subnet ID values is not a solution, because Fn::ImportValue can only import values that have been exported by another stack within the same region⁴. This means that you cannot use Fn::ImportValue to reference the subnet IDs that are exported by Account-A's CloudFormation template, unless all the member accounts are in the same region as Account-A. This option also does not avoid creating duplicate VPCs in each account, which is one of the requirements.
- > B. Using a transit gateway in the VPC within Account-A and configuring the member accounts to use the transit gateway to access the subnets in Account-A to launch workloads is not a solution, because a transit gateway does not allow you to launch workloads in another account's subnets. A transit gateway is a network transit hub that enables you to route traffic between your VPCs and on-premises networks⁵, but it does not enable you to share subnets across accounts.
- > D. Creating a peering connection between Account-A and the remaining member accounts and configuring the member accounts to use the subnets in Account-A through the VPC peering connection to launch workloads is not a solution, because a VPC peering connection does not allow you to launch workloads in another account's subnets. A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them privately⁶, but it does not enable you to share subnets across accounts.

References:

1: What is AWS Resource Access Manager? 2: Shareable AWS resources 3: Managing permissions for shared resources 4: Fn::ImportValue 5: What is a transit gateway? 6: What is VPC peering?

NEW QUESTION 79

A company has multiple Amazon S3 buckets encrypted with customer-managed CMKs. Due to regulatory requirements the keys must be rotated every year. The company's Security Engineer has enabled automatic key rotation for the CMKs; however the company wants to verify that the rotation has occurred. What should the Security Engineer do to accomplish this?

- A. Filter IAM CloudTrail logs for KeyRotation events
- B. Monitor Amazon CloudWatch Events for any IAM KMS CMK rotation events
- C. Using the IAM CLI
- D. run the IAM kms get-key-rotation-status operation with the --key-id parameter to check the CMK rotation date
- E. Use Amazon Athena to query IAM CloudTrail logs saved in an S3 bucket to filter Generate New Key events

Answer: C

Explanation:

the aws kms get-key-rotation-status command returns a boolean value that indicates whether automatic rotation of the customer master key (CMK) is enabled¹. This command also shows the date and time when the CMK was last rotated². The other options are not valid ways to check the CMK rotation status.

NEW QUESTION 83

A company is using IAM Organizations to develop a multi-account secure networking strategy. The company plans to use separate centrally managed accounts for shared services, auditing, and security inspection. The company plans to provide dozens of additional accounts to application owners for production and development environments.

Company security policy requires that all internet traffic be routed through a centrally managed security inspection layer in the security inspection account. A security engineer must recommend a solution that minimizes administrative overhead and complexity.

Which solution meets these requirements?

- A. Use IAM Control Tower
- B. Modify the default Account Factory networking template to automatically associate new accounts with a centrally managed VPC through a VPC peering connection and to create a default route to the VPC peer in the default route table
- C. Create an SCP that denies the CreateInternetGateway action
- D. Attach the SCP to all accounts except the security inspection account.
- E. Create a centrally managed VPC in the security inspection account
- F. Establish VPC peering connections between the security inspection account and other account
- G. Instruct account owners to create default routes in their account route tables that point to the VPC peer
- H. Create an SCP that denies the AttachInternetGateway action
- I. Attach the SCP to all accounts except the security inspection account.
- J. Use IAM Control Tower
- K. Modify the default Account Factory networking template to automatically associate new accounts with a centrally managed transit gateway and to create a default route to the transit gateway in the default route table
- L. Create an SCP that denies the AttachInternetGateway action
- M. Attach the SCP to all accounts except the security inspection account.
- N. Enable IAM Resource Access Manager (IAM RAM) for IAM Organization
- O. Create a shared transit gateway, and make it available by using an IAM RAM resource share
- P. Create an SCP that denies the CreateInternetGateway action
- Q. Attach the SCP to all accounts except the security inspection account
- R. Create routes in the route tables of all accounts that point to the shared transit gateway.

Answer: C

NEW QUESTION 88

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

Which bucket policy statement meets these requirements?

A. A computer code with black text Description automatically generated

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

B. A computer code with black text Description automatically generated

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

C. A computer code with black text Description automatically generated

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

D. A computer code with black text Description automatically generated

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

Answer: B

Explanation:

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

NEW QUESTION 91

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 92

A company is building a data processing application that uses AWS Lambda functions. The application's Lambda functions need to communicate with an Amazon RDS DB instance that is deployed within a VPC in the same AWS account.

Which solution meets these requirements in the MOST secure way?

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

Answer: C

Explanation:

The AWS documentation states that you can deploy the Lambda functions inside the VPC and attach a security group to the Lambda functions. You can then provide outbound rule access to the VPC CIDR range only and update the DB instance security group to allow traffic from the Lambda security group. This method is the most secure way to meet the requirements.

References: : AWS Lambda Developer Guide

NEW QUESTION 94

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 95

A security engineer is working with a company to design an e-commerce application. The application will run on Amazon EC2 instances that run in an Auto Scaling group behind an Application Load Balancer (ALB). The application will use an Amazon RDS DB instance for its database.

The only required connectivity from the internet is for HTTP and HTTPS traffic to the application. The application must communicate with an external payment provider that allows traffic only from a preconfigured allow list of IP addresses. The company must ensure that communications with the external payment provider are not interrupted as the environment scales.

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

- A. Deploy a NAT gateway in each private subnet for every Availability Zone that is in use.
- B. Place the DB instance in a public subnet.
- C. Place the DB instance in a private subnet.
- D. Configure the Auto Scaling group to place the EC2 instances in a public subnet.
- E. Configure the Auto Scaling group to place the EC2 instances in a private subnet.
- F. Deploy the ALB in a private subnet.

Answer: ACE

NEW QUESTION 100

Which of the following bucket policies will ensure that objects being uploaded to a bucket called 'demo' are encrypted. Please select:

A. C:\Users\wk\Desktop\mudassar\Untitled.jpg

```

{
  "Version": "2012-10-17",
  "Id": "PutObj",
  "Statement": [
    {
      "Sid": "DenyUploads",
      "Effect": "Deny",
      "Principal": "*",
      "Action": "s3:PutObject",
      "Resource": "arn:aws:s3:::demo/*",
      "Condition": {
        "StringNotEquals": {
          "s3:x-amz-server-side-encryption": "aws:kms"
        }
      }
    }
  ]
}

```

B. C:\Users\wk\Desktop\mudassar\Untitled.jpg

```

{
  "Version": "2012-10-17",
  "Id": "PutObj",
  "Statement": [
    {
      "Sid": "DenyUploads",
      "Effect": "Deny",
      "Principal": "*",
      "Action": "s3:PutObject",
      "Resource": "arn:aws:s3:::demo/*",
      "Condition": {
        "StringEquals": {
          "s3:x-amz-server-side-encryption": "aws:kms"
        }
      }
    }
  ]
}

```

C. C:\Users\wk\Desktop\mudassar\Untitled.jpg

```

{
  "Version": "2012-10-17",
  "Id": "PutObj",
  "Statement": [
    {
      "Sid": "DenyUploads",
      "Effect": "Deny",
      "Principal": "*",
      "Action": "s3:PutObject",
      "Resource": "arn:aws:s3:::demo/*"
    }
  ]
}

```

D. C:\Users\wk\Desktop\mudassar\Untitled.jpg

```

"Version":"2012-10-17",
"Id":"PutObj",
"Statement":[{
  "Sid":"DenyUploads",
  "Effect":"Deny",
  "Principal":"*",
  "Action":["s3:PutObjectEncrypted"],
  "Resource":["arn:aws:s3:::demo/*"]
}]
}
]
}

```

Answer: A

Explanation:

The condition of "s3:x-amz-server-side-encryption":"IAM:kms" ensures that objects uploaded need to be encrypted. Options B,C and D are invalid because you have to ensure the condition of ns3:x-amz-server-side-encryption":"IAM:kms" is present. For more information on IAM KMS best practices, just browse to the below URL: <https://dl.IAMstatic.com/whitepapers/IAM-kms-best-practices.pdf>. Submit your Feedback/Queries to our Expert

NEW QUESTION 101

A security engineer needs to run an AWS CloudFormation script. The CloudFormation script builds AWS infrastructure to support a stack that includes web servers and a MySQL database. The stack has been deployed in pre-production environments and is ready for production. The production script must comply with the principle of least privilege. Additionally, separation of duties must exist between the security engineer's IAM account and CloudFormation. Which solution will meet these requirements?

- A. Use IAM Access Analyzer policy generation to generate a policy that allows the CloudFormation script to run and manage the stack.
- B. Attach the policy to a new IAM role.
- C. Modify the security engineer's IAM permissions to be able to pass the new role to CloudFormation.
- D. Create an IAM policy that allows ec2:* and rds:* permission.
- E. Attach the policy to a new IAM role. Modify the security engineer's IAM permissions to be able to assume the new role.
- F. Use IAM Access Analyzer policy generation to generate a policy that allows the CloudFormation script to run and manage the stack.
- G. Modify the security engineer's IAM permissions to be able to run the CloudFormation script.
- H. Create an IAM policy that allows ec2:* and rds:* permission.
- I. Attach the policy to a new IAM role.
- J. Use the IAM policy simulator to confirm that the policy allows the AWS API calls that are necessary to build the stack.
- K. Modify the security engineer's IAM permissions to be able to pass the new role to CloudFormation.

Answer: A

Explanation:

The correct answer is A. Use IAM Access Analyzer policy generation to generate a policy that allows the CloudFormation script to run and manage the stack. Attach the policy to a new IAM role. Modify the security engineer's IAM permissions to be able to pass the new role to CloudFormation. According to the AWS documentation, IAM Access Analyzer is a service that helps you identify the resources in your organization and accounts, such as Amazon S3 buckets or IAM roles, that are shared with an external entity. You can also use IAM Access Analyzer to generate fine-grained policies that grant least privilege access based on access activity and access attempts. To use IAM Access Analyzer policy generation, you need to enable IAM Access Analyzer in your account or organization. You can then use the IAM console or the AWS CLI to generate a policy for a resource based on its access activity or access attempts. You can review and edit the generated policy before applying it to the resource. To use IAM Access Analyzer policy generation with CloudFormation, you can follow these steps:

- > Run the CloudFormation script in a pre-production environment and monitor its access activity or access attempts using IAM Access Analyzer.
- > Use IAM Access Analyzer policy generation to generate a policy that allows the CloudFormation script to run and manage the stack. The policy will include only the permissions that are necessary for the script to function.
- > Attach the policy to a new IAM role that has a trust relationship with CloudFormation. This will allow CloudFormation to assume the role and execute the script.
- > Modify the security engineer's IAM permissions to be able to pass the new role to CloudFormation.

This will allow the security engineer to launch the stack using the role.

- > Run the CloudFormation script in the production environment using the new role.

This solution will meet the requirements of least privilege and separation of duties, as it will limit the permissions of both CloudFormation and the security engineer to only what is needed for running and managing the stack. Option B is incorrect because creating an IAM policy that allows ec2:* and rds:* permissions is not following the principle of least privilege, as it will grant more permissions than necessary for running and managing the stack. Moreover, modifying the security engineer's IAM permissions to be able to assume the new role is not ensuring separation of duties, as it will allow the security engineer to bypass CloudFormation and directly access the resources. Option C is incorrect because modifying the security engineer's IAM permissions to be able to run the CloudFormation script is not ensuring separation of duties, as it will allow the security engineer to execute the script without using CloudFormation. Option D is incorrect because creating an IAM policy that allows ec2:* and rds:* permissions is not following the principle of least privilege, as it will grant more permissions than necessary for running and managing the stack. Using the IAM policy simulator to confirm that the policy allows the AWS API calls that are necessary to build the stack is not sufficient, as it will not generate a fine-grained policy based on access activity or access attempts.

NEW QUESTION 105

A company is running its workloads in a single AWS Region and uses AWS Organizations. A security engineer must implement a solution to prevent users from launching resources in other Regions. Which solution will meet these requirements with the LEAST operational overhead?

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

Answer: D

Explanation:

Although you can use a IAM policy to prevent users launching resources in other regions. The best practice is to use SCP when using AWS organizations.
https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_general.htm

NEW QUESTION 109

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

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

Answer: C

NEW QUESTION 112

An international company has established a new business entity in South Korea. The company also has established a new AWS account to contain the workload for the South Korean region. The company has set up the workload in the new account in the ap-northeast-2 Region. The workload consists of three Auto Scaling groups of Amazon EC2 instances. All workloads that operate in this Region must keep system logs and application logs for 7 years. A security engineer must implement a solution to ensure that no logging data is lost for each instance during scaling activities. The solution also must keep the logs for only the required period of 7 years. Which combination of steps should the security engineer take to meet these requirements? (Choose three.)

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

Answer: ABC

Explanation:

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

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

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

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

NEW QUESTION 113

A company uses several AWS CloudFormation stacks to handle the deployment of a suite of applications. The leader of the company's application development team notices that the stack deployments fail with permission errors when some team members try to deploy the stacks. However, other team members can deploy the stacks successfully.

The team members access the account by assuming a role that has a specific set of permissions that are necessary for the job responsibilities of the team members. All team members have permissions to perform operations on the stacks.

Which combination of steps will ensure consistent deployment of the stacks MOST securely? (Select THREE.)

- A. Create a service role that has a composite principal that contains each service that needs the necessary permission
- B. Configure the role to allow the sts:AssumeRole action.
- C. Create a service role that has cloudformation.amazonaws.com as the service principal
- D. Configure the role to allow the sts:AssumeRole action.
- E. For each required set of permissions, add a separate policy to the role to allow those permissions

- F. Add the ARN of each CloudFormation stack in the resource field of each policy.
- G. For each required set of permissions, add a separate policy to the role to allow those permission
- H. Add the ARN of each service that needs the per-missions in the resource field of the corresponding policy.
- I. Update each stack to use the service role.
- J. Add a policy to each member role to allow the iam:PassRole actio
- K. Set the policy's resource field to the ARN of the service role.

Answer: BDF

NEW QUESTION 116

A company is evaluating the use of AWS Systems Manager Session Manager to gam access to the company's Amazon EC2 instances. However, until the company implements the change, the company must protect the key file for the EC2 instances from read and write operations by any other users. When a security administrator tries to connect to a critical EC2 Linux instance during an emergency, the security administrator receives the following error. "Error Unprotected private key file - Permissions for 'ssh/my_private_key.pem' are too open". Which command should the security administrator use to modify the private key file permissions to resolve this error?

- A. `chmod 0040 ssh/my_private_key.pem`
- B. `chmod 0400 ssh/my_private_key.pem`
- C. `chmod 0004 ssh/my_private_key.pem`
- D. `chmod 0777 ssh/my_private_key.pem`

Answer: B

Explanation:

The error message indicates that the private key file permissions are too open, meaning that other users can read or write to the file. This is a security risk, as the private key should be accessible only by the owner of the file. To fix this error, the security administrator should use the `chmod` command to change the permissions of the private key file to `0400`, which means that only the owner can read the file and no one else can read or write to it. The `chmod` command takes a numeric argument that represents the permissions for the owner, group, and others in octal notation. Each digit corresponds to a set of permissions: read (4), write (2), and execute (1). The digits are added together to get the final permissions for each category. For example, `0400` means that the owner has read permission (4) and no other permissions (0), and the group and others have no permissions at all (0). The other options are incorrect because they either do not change the permissions at all (D), or they give too much or too little permissions to the owner, group, or others (A, C).

Verified References:

- > <https://superuser.com/questions/215504/permissions-on-private-key-in-ssh-folder>
- > <https://www.baeldung.com/linux/ssh-key-permissions>

NEW QUESTION 120

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

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

Answer: B

Explanation:

One can use the IAM Encryption CLI to encrypt the data before sending it across to the S3 bucket. Options A and C are invalid because this would still mean that data is transferred in plain text. Option D is invalid because you cannot just enable client side encryption for the S3 bucket. For more information on Encrypting and Decrypting data, please visit the below URL: <https://IAM.amazonaws.com/blogs/security/how-to-encrypt-and-decrypt-your-data-with-the-IAM-encryption-cli>. The correct answer is: Use the IAM Encryption CLI to encrypt the data first. Submit your Feedback/Queries to our Experts

NEW QUESTION 122

A company's policy requires that all API keys be encrypted and stored separately from source code in a centralized security account. This security account is managed by the company's security team. However, an audit revealed that an API key is stored with the source code of an IAM Lambda function in an IAM CodeCommit repository in the DevOps account. How should the security team securely store the API key?

- A. Create a CodeCommit repository in the security account using IAM Key Management Service (IAMKMS) for encryption. Require the development team to migrate the Lambda source code to this repository.
- B. Store the API key in an Amazon S3 bucket in the security account using server-side encryption with Amazon S3 managed encryption keys (SSE-S3) to encrypt the key. Create a reserved URL for the S3 key.
- C. and specify the URL in a Lambda environmental variable in the IAM CloudFormation template. Update the Lambda function code to retrieve the key using the URL and call the API.
- D. Create a secret in IAM Secrets Manager in the security account to store the API key using IAM Key Management Service (IAM KMS) for encryption. Grant access to the IAM role used by the Lambda function so that the function can retrieve the key from Secrets Manager and call the API.
- E. Create an encrypted environment variable for the Lambda function to store the API key using IAM Key Management Service (IAM KMS) for encryption. Grant access to the IAM role used by the Lambda function so that the function can decrypt the key at runtime.

Answer: C

Explanation:

To securely store the API key, the security team should do the following:

- > Create a secret in AWS Secrets Manager in the security account to store the API key using AWS Key Management Service (AWS KMS) for encryption. This allows the security team to encrypt and manage the API key centrally, and to configure automatic rotation schedules for it.

➤ Grant access to the IAM role used by the Lambda function so that the function can retrieve the key from Secrets Manager and call the API. This allows the security team to avoid storing the API key with the source code, and to use IAM policies to control access to the secret.

NEW QUESTION 127

A company has several petabytes of data. The company must preserve this data for 7 years to comply with regulatory requirements. The company's compliance team asks a security officer to develop a strategy that will prevent anyone from changing or deleting the data. Which solution will meet this requirement MOST cost-effectively?

- A. Create an Amazon S3 bucket
- B. Configure the bucket to use S3 Object Lock in compliance mode
- C. Upload the data to the bucket
- D. Create a resource-based bucket policy that meets all the regulatory requirements.
- E. Create an Amazon S3 bucket
- F. Configure the bucket to use S3 Object Lock in governance mode
- G. Upload the data to the bucket
- H. Create a user-based IAM policy that meets all the regulatory requirements.
- I. Create a vault in Amazon S3 Glacier
- J. Create a Vault Lock policy in S3 Glacier that meets all the regulatory requirements
- K. Upload the data to the vault.
- L. Create an Amazon S3 bucket
- M. Upload the data to the bucket
- N. Use a lifecycle rule to transition the data to a vault in S3 Glacier
- O. Create a Vault Lock policy that meets all the regulatory requirements.

Answer: C

Explanation:

To preserve the data for 7 years and prevent anyone from changing or deleting it, the security officer needs to use a service that can store the data securely and enforce compliance controls. The most cost-effective way to do this is to use Amazon S3 Glacier, which is a low-cost storage service for data archiving and long-term backup. S3 Glacier allows you to create a vault, which is a container for storing archives. Archives are any data such as photos, videos, or documents that you want to store durably and reliably.

S3 Glacier also offers a feature called Vault Lock, which helps you to easily deploy and enforce compliance controls for individual vaults with a Vault Lock policy. You can specify controls such as "write once read many" (WORM) in a Vault Lock policy and lock the policy from future edits. Once a Vault Lock policy is locked, the policy can no longer be changed or deleted. S3 Glacier enforces the controls set in the Vault Lock policy to help achieve your compliance objectives. For example, you can use Vault Lock policies to enforce data retention by denying deletes for a specified period of time.

To use S3 Glacier and Vault Lock, the security officer needs to follow these steps:

- Create a vault in S3 Glacier using the AWS Management Console, AWS Command Line Interface (AWS CLI), or AWS SDKs.
- Create a Vault Lock policy in S3 Glacier that meets all the regulatory requirements using the IAM policy language. The policy can include conditions such as `aws:CurrentTime` or `aws:SecureTransport` to further restrict access to the vault.
- Initiate the lock by attaching the Vault Lock policy to the vault, which sets the lock to an in-progress state and returns a lock ID. While the policy is in the in-progress state, you have 24 hours to validate your Vault Lock policy before the lock ID expires. To prevent your vault from exiting the in-progress state, you must complete the Vault Lock process within these 24 hours. Otherwise, your Vault Lock policy will be deleted.
- Use the lock ID to complete the lock process. If the Vault Lock policy doesn't work as expected, you can stop the Vault Lock process and restart from the beginning.
- Upload the data to the vault using either direct upload or multipart upload methods. For more information about S3 Glacier and Vault Lock, see S3 Glacier Vault Lock.

The other options are incorrect because:

- Option A is incorrect because creating an Amazon S3 bucket and configuring it to use S3 Object Lock in compliance mode will not prevent anyone from changing or deleting the data. S3 Object Lock is a feature that allows you to store objects using a WORM model in S3. You can apply two types of object locks: retention periods and legal holds. A retention period specifies a fixed period of time during which an object remains locked. A legal hold is an indefinite lock on an object until it is removed. However, S3 Object Lock only prevents objects from being overwritten or deleted by any user, including the root user in your AWS account. It does not prevent objects from being modified by other means, such as changing their metadata or encryption settings. Moreover, S3 Object Lock requires that you enable versioning on your bucket, which will incur additional storage costs for storing multiple versions of an object.
- Option B is incorrect because creating an Amazon S3 bucket and configuring it to use S3 Object Lock in governance mode will not prevent anyone from changing or deleting the data. S3 Object Lock in governance mode works similarly to compliance mode, except that users with specific IAM permissions can change or delete objects that are locked. This means that users who have `s3:BypassGovernanceRetention` permission can remove retention periods or legal holds from objects and overwrite or delete them before they expire. This option does not provide strong enforcement for compliance controls as required by the regulatory requirements.
- Option D is incorrect because creating an Amazon S3 bucket and using a lifecycle rule to transition the data to a vault in S3 Glacier will not prevent anyone from changing or deleting the data. Lifecycle rules are actions that Amazon S3 automatically performs on objects during their lifetime. You can use lifecycle rules to transition objects between storage classes or expire them after a certain period of time. However, lifecycle rules do not apply any compliance controls on objects or prevent them from being modified or deleted by users. Moreover, transitioning objects from S3 to S3 Glacier using lifecycle rules will incur additional charges for retrieval requests and data transfers.

NEW QUESTION 129

A security engineer needs to implement a solution to create and control the keys that a company uses for cryptographic operations. The security engineer must create symmetric keys in which the key material is generated and used within a custom key store that is backed by an AWS CloudHSM cluster. The security engineer will use symmetric and asymmetric data key pairs for local use within applications. The security engineer also must audit the use of the keys. How can the security engineer meet these requirements?

- A. To create the keys use AWS Key Management Service (AWS KMS) and the custom key stores with the CloudHSM cluster
- B. For auditing, use Amazon Athena
- C. To create the keys use Amazon S3 and the custom key stores with the CloudHSM cluster
- D. For auditing use AWS CloudTrail.
- E. To create the keys use AWS Key Management Service (AWS KMS) and the custom key stores with the CloudHSM cluster
- F. For auditing, use Amazon GuardDuty.
- G. To create the keys use AWS Key Management Service (AWS KMS) and the custom key stores with the CloudHSM cluster

H. For auditing, use AWS CloudTrail.

Answer: D

Explanation:

AWS KMS supports asymmetric KMS keys that represent a mathematically related RSA, elliptic curve (ECC), or SM2 (China Regions only) public and private key pair. These key pairs are generated in AWS KMS hardware security modules certified under the FIPS 140-2 Cryptographic Module Validation Program, except in the China (Beijing) and China (Ningxia) Regions. The private key never leaves the AWS KMS HSMs unencrypted.
<https://docs.aws.amazon.com/kms/latest/developerguide/symmetric-asymmetric.html>

NEW QUESTION 130

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

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

Answer: D

NEW QUESTION 133

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

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

Answer: ADF

NEW QUESTION 134

.....

Thank You for Trying Our Product

We offer two products:

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

2nd - Questions and Answers in PDF Format

SCS-C01 Practice Exam Features:

- * SCS-C01 Questions and Answers Updated Frequently
- * SCS-C01 Practice Questions Verified by Expert Senior Certified Staff
- * SCS-C01 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * SCS-C01 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The SCS-C01 Practice Test Here](#)