



Exam Code	SCS-C03
Total Questions	20

Dumpsgate: Real Questions. Proven Results.

Developed By Dumpsgate

Question: 1

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: A

Explanation:

Amazon S3 Block Public Access configured at the AWS account level is the recommended and most effective approach to protect data stored in Amazon S3 while minimizing operational overhead. AWS Security Specialty documentation explains that S3 Block Public Access provides centralized, preventative controls designed to block public access to S3 buckets and objects regardless of individual bucket policies or object-level ACL configurations. When enabled at the account level,

these controls automatically apply to all existing and newly created buckets, significantly reducing the risk of accidental exposure caused by misconfigured permissions.

The AWS Certified Security – Specialty Study Guide emphasizes that public access misconfiguration is a leading cause of data leaks in cloud environments. Account-level S3 Block Public Access acts as a guardrail by overriding any attempt to grant public permissions through bucket policies or ACLs. This eliminates the need to manage security settings on a per-bucket or per-object basis, thereby reducing administrative complexity and human error.

Configuring Block Public Access at the object level, as in option B, requires continuous monitoring and manual configuration, which increases operational overhead. Disabling ACLs alone, as described in option C, does not fully prevent public access because bucket policies can still allow public permissions. Using AWS PrivateLink, as in option D, controls network access but does not protect against public exposure through misconfigured S3 policies.

AWS security best practices explicitly recommend enabling S3 Block Public Access at the account level as the primary mechanism for preventing unintended public data exposure with minimal management effort.

Referenced AWS Specialty Documents:

AWS Certified Security – Specialty Official Study Guide

Amazon S3 Security Best Practices Documentation

Amazon S3 Block Public Access Overview

AWS Well-Architected Framework – Security Pillar

Question: 2

A company's developers are using AWS Lambda function URLs to invoke functions directly. The company must ensure that developers cannot configure or deploy unauthenticated functions in production accounts. The company wants to meet this requirement by using AWS Organizations. The solution must not require additional work for the developers.

Which solution will meet these requirements?

- A. Require the developers to configure all function URLs to support cross-origin resource sharing (CORS) when the functions are called from a different domain.
- B. Use an AWS WAF delegated administrator account to view and block unauthenticated access to

function URLs in production accounts, based on the OU of accounts that are using the functions.

C. Use SCPs to allow all `lambda:CreateFunctionUrlConfig` and `lambda:UpdateFunctionUrlConfig` actions that have a `lambda:FunctionUrlAuthType` condition key value of `AWS_IAM`.

D. Use SCPs to deny all `lambda:CreateFunctionUrlConfig` and `lambda:UpdateFunctionUrlConfig` actions that have a `lambda:FunctionUrlAuthType` condition key value of `NONE`.

Answer: D

Explanation:

AWS Organizations service control policies (SCPs) are designed to enforce preventive guardrails across accounts without requiring application-level changes. According to the AWS Certified Security – Specialty documentation, SCPs can restrict specific API actions or require certain condition keys to enforce security standards centrally. AWS Lambda function URLs support two authentication modes: `AWS_IAM` and `NONE`. When the authentication type is set to `NONE`, the function URL becomes publicly accessible, which introduces a significant security risk in production environments.

By using an SCP that explicitly denies the `lambda:CreateFunctionUrlConfig` and `lambda:UpdateFunctionUrlConfig` actions when the `lambda:FunctionUrlAuthType` condition key equals `NONE`, the organization ensures that unauthenticated function URLs cannot be created or modified in production accounts. This enforcement occurs at the AWS Organizations level and applies automatically to all accounts within the specified organizational units (OUs). Developers are not required to change their workflows or add additional controls, satisfying the requirement of no additional developer effort.

Option A relates to browser-based access controls and does not provide authentication or authorization enforcement. Option B is not valid because AWS WAF cannot be attached directly to AWS Lambda function URLs. Option C is incorrect because SCPs do not grant permissions; they only limit permissions. AWS documentation clearly states that SCPs define maximum available permissions and are evaluated before IAM policies.

This approach aligns with AWS best practices for centralized governance, least privilege, and preventive security controls.

Referenced AWS Specialty Documents:

AWS Certified Security – Specialty Official Study Guide

AWS Organizations Service Control Policies Documentation

AWS Lambda Security and Function URL Authentication Overview

Question: 3

A security engineer receives a notice about suspicious activity from a Linux-based Amazon EC2 instance that uses Amazon Elastic Block Store (Amazon EBS)-based storage. The instance is making connections to known malicious addresses.

The instance is in a development account within a VPC that is in the us-east-1 Region. The VPC contains an internet gateway and has a subnet in us-east-1a and us-east-1b. Each subnet is associated with a route table that uses the internet gateway as a default route. Each subnet also uses the default network ACL. The suspicious EC2 instance runs within the us-east-1b subnet. During an initial investigation, a security engineer discovers that the suspicious instance is the only instance that runs in the subnet.

Which response will immediately mitigate the attack and help investigate the root cause?

A. Log in to the suspicious instance and use the netstat command to identify remote connections. Use the IP addresses from these remote connections to create deny rules in the security group of the instance. Install diagnostic tools on the instance for investigation. Update the outbound network ACL for the subnet in us-east-1b to explicitly deny all connections as the first rule during the investigation of the instance.

B. Update the outbound network ACL for the subnet in us-east-1b to explicitly deny all connections as the first rule. Replace the security group with a new security group that allows connections only from a diagnostics security group. Update the outbound network ACL for the us-east-1b subnet to remove the deny all rule. Launch a new EC2 instance that has diagnostic tools. Assign the new security group to the new EC2 instance. Use the new EC2 instance to investigate the suspicious instance.

C. Ensure that the Amazon Elastic Block Store (Amazon EBS) volumes that are attached to the suspicious EC2 instance will not delete upon termination. Terminate the instance. Launch a new EC2 instance in us-east-1a that has diagnostic tools. Mount the EBS volumes from the terminated instance for investigation.

D. Create an AWS WAF web ACL that denies traffic to and from the suspicious instance. Attach the AWS WAF web ACL to the instance to mitigate the attack. Log in to the instance and install diagnostic tools to investigate the instance.

Answer: C

Explanation:

AWS incident response best practices emphasize immediate containment, preservation of evidence, and safe forensic investigation. According to the AWS Certified Security – Specialty Study Guide, when an EC2 instance is suspected of compromise, security teams should avoid logging in to the instance or installing additional tools, as these actions can alter evidence and increase risk.

Terminating the compromised instance after ensuring that its Amazon EBS volumes are preserved prevents further malicious activity immediately. By setting the EBS volumes to not delete on termination, all disk data is retained for forensic analysis. Launching a new, clean EC2 instance in a different subnet or Availability Zone with preinstalled diagnostic tools allows investigators to safely attach and analyze the compromised volumes without executing potentially malicious code.

Option A introduces significant risk by logging in to the compromised instance and modifying security controls during active compromise. Option B delays containment and allows continued outbound traffic during investigation steps. Option D is invalid because AWS WAF cannot be attached directly to Amazon EC2 instances and does not control outbound traffic.

AWS documentation strongly recommends isolating or terminating compromised resources and performing offline analysis using detached storage volumes. This approach ensures immediate mitigation, preserves forensic integrity, and aligns with AWS incident response frameworks.

Referenced AWS Specialty Documents:

AWS Certified Security – Specialty Official Study Guide

AWS Incident Response Best Practices

Amazon EC2 and EBS Forensics Guidance

AWS Well-Architected Framework – Security Pillar

Question: 4

A company has a VPC that has no internet access and has the private DNS hostnames option enabled. An Amazon Aurora database is running inside the VPC. A security engineer wants to use AWS Secrets Manager to automatically rotate the credentials for the Aurora database. The security

engineer configures the Secrets Manager default AWS Lambda rotation function to run inside the same VPC that the Aurora database uses. However, the security engineer determines that the password cannot be rotated properly because the Lambda function cannot communicate with the Secrets Manager endpoint.

What is the MOST secure way that the security engineer can give the Lambda function the ability to communicate with the Secrets Manager endpoint?

- A. Add a NAT gateway to the VPC to allow access to the Secrets Manager endpoint.
- B. Add a gateway VPC endpoint to the VPC to allow access to the Secrets Manager endpoint.
- C. Add an interface VPC endpoint to the VPC to allow access to the Secrets Manager endpoint.
- D. Add an internet gateway for the VPC to allow access to the Secrets Manager endpoint.

Answer: C

Explanation:

AWS Secrets Manager is a regional service that is accessed through private AWS endpoints. In a VPC without internet access, AWS recommends using AWS PrivateLink through interface VPC endpoints to enable secure, private connectivity to supported AWS services. According to AWS Certified Security – Specialty documentation, interface VPC endpoints allow resources within a VPC to communicate with AWS services without traversing the public internet, NAT devices, or internet gateways.

An interface VPC endpoint for Secrets Manager creates elastic network interfaces (ENIs) within the VPC subnets and assigns private IP addresses that route traffic directly to the Secrets Manager service. Because the VPC has private DNS enabled, the standard Secrets Manager DNS hostname resolves to the private IP addresses of the interface endpoint, allowing the Lambda rotation function to communicate securely and transparently.

Option A introduces unnecessary complexity and expands the attack surface by allowing outbound internet access. Option B is incorrect because gateway VPC endpoints are supported only for Amazon S3 and Amazon DynamoDB. Option D violates the security requirement by exposing the VPC to the internet.

AWS security best practices explicitly recommend interface VPC endpoints as the most secure connectivity method for private VPC workloads accessing AWS managed services.

Referenced AWS Specialty Documents:

AWS Certified Security – Specialty Official Study Guide

AWS Secrets Manager Security Architecture

AWS PrivateLink and Interface VPC Endpoints Documentation

Question: 5

A security engineer wants to forward custom application-security logs from an Amazon EC2 instance to Amazon CloudWatch. The security engineer installs the CloudWatch agent on the EC2 instance and adds the path of the logs to the CloudWatch configuration file.

However, CloudWatch does not receive the logs. The security engineer verifies that the awslogs service is running on the EC2 instance.

What should the security engineer do next to resolve the issue?

- A. Add AWS CloudTrail to the trust policy of the EC2 instance. Send the custom logs to CloudTrail instead of CloudWatch.
- B. Add Amazon S3 to the trust policy of the EC2 instance. Configure the application to write the custom logs to an S3 bucket that CloudWatch can use to ingest the logs.
- C. Add Amazon Inspector to the trust policy of the EC2 instance. Use Amazon Inspector instead of the CloudWatch agent to collect the custom logs.
- D. Attach the CloudWatchAgentServerPolicy AWS managed policy to the EC2 instance role.

Answer: D

Explanation:

The Amazon CloudWatch agent requires explicit IAM permissions to create log groups, create log streams, and put log events into Amazon CloudWatch Logs. According to the AWS Certified Security – Specialty Study Guide, the most common cause of CloudWatch agent log delivery failures is missing

or insufficient IAM permissions on the EC2 instance role.

The CloudWatchAgentServerPolicy AWS managed policy provides the required permissions, including logs:CreateLogGroup, logs:CreateLogStream, and logs:PutLogEvents. Attaching this policy to the EC2 instance role enables the CloudWatch agent to successfully deliver custom application logs without requiring changes to the application or logging configuration.

Options A, B, and C are incorrect because CloudTrail, Amazon S3, and Amazon Inspector are not designed to ingest custom application logs from EC2 instances in this manner. AWS documentation clearly states that IAM permissions must be granted to the EC2 role for CloudWatch Logs ingestion.

This approach aligns with AWS best practices for least privilege while ensuring reliable detection and monitoring capabilities.

Referenced AWS Specialty Documents:

AWS Certified Security – Specialty Official Study Guide

Amazon CloudWatch Logs Agent Configuration

AWS IAM Best Practices for Monitoring

Question: 6

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

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

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

A. Update the EC2 instance security group to add a rule that allows outbound traffic on port 443 for 0.0.0.0/0.

B. Update the EC2 instance security group to add a rule that allows inbound traffic on port 443 to the VPC's CIDR range.

- C. Create an EC2 key pair. Associate the key pair with the EC2 instance.
- D. Create a VPC interface endpoint for Systems Manager in the VPC where the EC2 instance is located.
- E. Attach a security group to the VPC interface endpoint. Allow inbound traffic on port 443 to the VPC's CIDR range.
- F. Create a VPC interface endpoint for the EC2 instance in the VPC where the EC2 instance is located.

Answer: A, D, E

Explanation:

AWS Systems Manager Session Manager requires secure outbound HTTPS connectivity from the EC2 instance to Systems Manager endpoints. In a VPC without internet access, AWS Certified Security – Specialty documentation recommends using interface VPC endpoints to enable private connectivity without exposing the instance to the internet.

Creating a VPC interface endpoint for Systems Manager allows the SSM Agent to communicate securely with the Systems Manager service. The endpoint must have an attached security group that allows inbound traffic on port 443 from the VPC CIDR range. Additionally, the EC2 instance security group must allow outbound HTTPS traffic on port 443 so the agent can initiate connections.

Option C is incorrect because creating or associating key pairs enables SSH access, which can alter forensic evidence and violates forensic best practices. Option B is unnecessary because Session Manager does not require inbound rules on the EC2 instance. Option F is invalid because EC2 does not use interface endpoints for management connectivity.

This combination ensures secure, private access for forensic investigation while preserving evidence integrity and adhering to AWS incident response best practices.

Referenced AWS Specialty Documents:

AWS Certified Security – Specialty Official Study Guide

AWS Systems Manager Session Manager Architecture

AWS Incident Response and Forensics Best Practices

Question: 7

A security team manages a company's AWS Key Management Service (AWS KMS) customer managed keys. Only members of the security team can administer the KMS keys. The company's application team has a software process that needs temporary access to the keys occasionally. The security team needs to provide the application team's software process with access to the keys.

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

- A. Export the KMS key material to an on-premises hardware security module (HSM). Give the application team access to the key material.
- B. Edit the key policy that grants the security team access to the KMS keys by adding the application team as principals. Revert this change when the application team no longer needs access.
- C. Create a key grant to allow the application team to use the KMS keys. Revoke the grant when the application team no longer needs access.
- D. Create a new KMS key by generating key material on premises. Import the key material to AWS KMS whenever the application team needs access. Grant the application team permissions to use the key.

Answer: C

Explanation:

AWS KMS key grants are specifically designed to provide temporary, granular permissions to use customer managed keys without modifying key policies. According to the AWS Certified Security – Specialty Study Guide, grants are the preferred mechanism for delegating key usage permissions to AWS principals for short-term or programmatic access scenarios. Grants allow permissions such as Encrypt, Decrypt, or GenerateDataKey and can be created and revoked dynamically.

Using a key grant avoids the operational risk and overhead of editing key policies, which are long-term control mechanisms and should remain stable. AWS documentation emphasizes that frequent key policy changes increase the risk of misconfiguration and accidental privilege escalation. Grants can be revoked immediately when access is no longer required, ensuring strong adherence to the principle of least privilege.

Options A and D violate AWS security best practices because AWS KMS does not allow direct export of key material unless the key was explicitly created as an importable key, and exporting key material

increases exposure risk. Option B requires manual policy changes and rollback, which introduces operational overhead and audit complexity.

AWS recommends key grants as the most efficient and secure way to provide temporary access to KMS keys for applications.

Referenced AWS Specialty Documents:

AWS Certified Security – Specialty Official Study Guide

AWS KMS Key Policies and Grants Documentation

AWS KMS Best Practices

Question: 8

A company is using AWS CloudTrail and Amazon CloudWatch to monitor resources in an AWS account. The company's developers have been using an IAM role in the account for the last 3 months.

A security engineer needs to refine the customer managed IAM policy attached to the role to ensure that the role provides least privilege access.

Which solution will meet this requirement with the LEAST effort?

- A. Implement AWS IAM Access Analyzer policy generation on the role.
- B. Implement AWS IAM Access Analyzer policy validation on the role.
- C. Search CloudWatch logs to determine the actions the role invoked and to evaluate the permissions.
- D. Use AWS Trusted Advisor to compare the policies assigned to the role against AWS best practices.

Answer: A

Explanation:

AWS IAM Access Analyzer policy generation is specifically designed to help security engineers generate least-privilege IAM policies based on actual usage recorded in AWS CloudTrail. According to the AWS Certified Security – Specialty documentation, policy generation analyzes historical CloudTrail data to identify the exact API actions and resources that a role has accessed over a specified time period.

Because the role has been actively used for three months, there is sufficient CloudTrail data for IAM Access Analyzer to generate a refined customer managed policy automatically. This significantly reduces manual effort and eliminates the need to analyze logs or infer permissions. The generated policy can be reviewed and attached directly to the role, ensuring least privilege access with minimal engineering effort.

Option B only validates existing policies for security warnings and does not reduce permissions. Option C requires manual analysis of CloudWatch logs, which is time-consuming and error-prone. Option D does not analyze real usage and cannot generate role-specific least privilege policies.

AWS documentation explicitly recommends IAM Access Analyzer policy generation as the fastest and most accurate method to refine IAM permissions based on observed behavior.

Referenced AWS Specialty Documents:

AWS Certified Security – Specialty Official Study Guide

AWS IAM Access Analyzer Policy Generation

AWS IAM Least Privilege Best Practices

Question: 9

A company uses AWS IAM Identity Center with SAML 2.0 federation. The company decides to change its federation source from one identity provider (IdP) to another. The underlying directory for both IdPs is Active Directory.

Which solution will meet this requirement?

A. Disable all existing users and groups within IAM Identity Center that were part of the federation with the original IdP.

B. Modify the attribute mappings within the IAM Identity Center trust relationship to match information that the new IdP sends.

- C. Reconfigure all existing IAM roles in the company's AWS accounts to explicitly trust the new IdP as the principal.
- D. Confirm that the Network Time Protocol (NTP) clock skew is correctly set between IAM Identity Center and the new IdP endpoints.

Answer: B

Explanation:

AWS IAM Identity Center relies on SAML assertions and attribute mappings to associate federated users with identities, groups, and permission sets. According to the AWS Certified Security – Specialty documentation, when changing identity providers while maintaining the same underlying directory, existing users and group identities can be preserved by updating attribute mappings to align with the new IdP's SAML assertions.

By modifying the attribute mappings, IAM Identity Center can correctly interpret usernames, group memberships, and unique identifiers sent by the new IdP without requiring changes to AWS account roles or permission sets. This approach minimizes operational effort and avoids disruption to access management.

Option A unnecessarily disables identities and causes access outages. Option C is incorrect because IAM Identity Center abstracts role trust relationships, and roles do not directly trust the IdP. Option D is unrelated to federation source configuration and only affects authentication timing issues.

AWS best practices recommend updating attribute mappings when switching IdPs that share the same directory source.

Referenced AWS Specialty Documents:

AWS Certified Security – Specialty Official Study Guide

AWS IAM Identity Center SAML Federation

AWS Identity Federation Best Practices

Question: 10

A company is running its application on AWS. The company has a multi-environment setup, and each environment is isolated in a separate AWS account. The company has an organization in AWS Organizations to manage the accounts. There is a single dedicated security account for the organization. The company must create an inventory of all sensitive data that is stored in Amazon S3 buckets across the organization's accounts. The findings must be visible from a single location.

Which solution will meet these requirements?

- A. Set the security account as the delegated administrator for Amazon Macie and AWS Security Hub. Enable and configure Macie to publish sensitive data findings to Security Hub.
- B. Set the security account as the delegated administrator for AWS Security Hub. In each account, configure Amazon Inspector to scan the S3 buckets for sensitive data. Publish sensitive data findings to Security Hub.
- C. In each account, configure Amazon Inspector to scan the S3 buckets for sensitive data. Enable Amazon Inspector integration with AWS Trusted Advisor. Publish sensitive data findings to Trusted Advisor.
- D. In each account, enable and configure Amazon Macie to detect sensitive data. Enable Macie integration with AWS Trusted Advisor. Publish sensitive data findings to Trusted Advisor.

Answer: A

Explanation:

Amazon Macie is the AWS service designed specifically to discover, classify, and inventory sensitive data stored in Amazon S3. According to the AWS Certified Security – Specialty Study Guide, Macie can be enabled organization-wide using AWS Organizations, with a delegated administrator account that centrally manages findings across all member accounts.

By designating the security account as the delegated administrator for both Amazon Macie and AWS Security Hub, the company can centralize sensitive data findings in a single location. Macie automatically scans S3 buckets for sensitive data such as personally identifiable information (PII) and publishes findings to Security Hub for centralized visibility and reporting.

Option B and C are incorrect because Amazon Inspector does not scan S3 objects for sensitive data. Option D is invalid because AWS Trusted Advisor does not ingest Macie sensitive data findings.

AWS best practices recommend Amazon Macie with delegated administration and Security Hub

integration for centralized sensitive data inventory across multi-account environments.

Referenced AWS Specialty Documents:

AWS Certified Security – Specialty Official Study Guide

Amazon Macie Sensitive Data Discovery

AWS Organizations Delegated Administrator Model

AWS Security Hub Integration Overview

Question: 11

A company must capture AWS CloudTrail data events and must retain the logs for 7 years. The logs must be immutable and must be available to be searched by complex queries. The company also needs to visualize the data from the logs.

Which solution will meet these requirements MOST cost-effectively?

- A. Create a CloudTrail Lake data store. Implement CloudTrail Lake dashboards to visualize and query the results.
- B. Use the CloudTrail Event History feature in the AWS Management Console. Visualize and query the results in the console.
- C. Send the CloudTrail logs to an Amazon S3 bucket. Provision a persistent Amazon EMR cluster that has access to the S3 bucket. Enable S3 Object Lock on the S3 bucket. Use Apache Spark to perform queries. Use Amazon QuickSight for visualizations.
- D. Send the CloudTrail logs to a log group in Amazon CloudWatch Logs. Set the CloudWatch Logs stream to send the data to an Amazon OpenSearch Service domain. Enable cold storage for the OpenSearch Service domain. Use OpenSearch Dashboards for visualizations and queries.

Answer: A

Explanation:

AWS CloudTrail Lake is purpose-built to store, query, and analyze CloudTrail events, including data

events, without requiring additional infrastructure. The AWS Certified Security – Specialty documentation explains that CloudTrail Lake provides immutable event storage with configurable retention periods, including multi-year retention, which satisfies long-term compliance requirements such as 7-year retention. Events are stored in an append-only, immutable format managed by AWS, reducing operational complexity.

CloudTrail Lake supports SQL-based queries for complex analysis directly against the event data, eliminating the need to export logs to other services for querying. Additionally, CloudTrail Lake includes built-in dashboards and integrations that enable visualization of event trends and patterns without standing up separate analytics or visualization platforms.

Option B is invalid because CloudTrail Event History only retains events for up to 90 days and does not support long-term retention or advanced querying. Option C introduces high operational overhead and cost by requiring persistent Amazon EMR clusters and additional services. Option D incurs ongoing ingestion, indexing, and storage costs for OpenSearch Service over a 7-year period, making it less cost-effective than CloudTrail Lake.

AWS documentation positions CloudTrail Lake as the most cost-effective and operationally efficient solution for long-term, queryable CloudTrail event storage and visualization.

Referenced AWS Specialty Documents:

AWS Certified Security – Specialty Official Study Guide

AWS CloudTrail Lake Architecture and Retention

AWS CloudTrail Data Events Overview

Question: 12

A company is planning to migrate its applications to AWS in a single AWS Region. The company's applications will use a combination of Amazon EC2 instances, Elastic Load Balancing (ELB) load balancers, and Amazon S3 buckets. The company wants to complete the migration as quickly as possible. All the applications must meet the following requirements:

- Data must be encrypted at rest.
- Data must be encrypted in transit.
- Endpoints must be monitored for anomalous network traffic.

Which combination of steps should a security engineer take to meet these requirements with the

LEAST effort? (Select THREE.)

- A. Install the Amazon Inspector agent on EC2 instances by using AWS Systems Manager Automation.
- B. Enable Amazon GuardDuty in all AWS accounts.
- C. Create VPC endpoints for Amazon EC2 and Amazon S3. Update VPC route tables to use only the secure VPC endpoints.
- D. Configure AWS Certificate Manager (ACM). Configure the load balancers to use certificates from ACM.
- E. Use AWS Key Management Service (AWS KMS) for key management. Create an S3 bucket policy to deny any PutObject command with a condition for x-amz-meta-side-encryption.
- F. Use AWS Key Management Service (AWS KMS) for key management. Create an S3 bucket policy to deny any PutObject command with a condition for x-amz-server-side-encryption.

Answer: B, D, F

Explanation:

Amazon GuardDuty provides continuous monitoring for anomalous and malicious network activity by analyzing VPC Flow Logs, DNS logs, and CloudTrail events. Enabling GuardDuty across accounts requires minimal configuration and immediately satisfies the requirement to monitor endpoints for anomalous network traffic, as described in the AWS Certified Security – Specialty Study Guide.

Encrypting data in transit for applications behind Elastic Load Balancing is most efficiently achieved by using AWS Certificate Manager (ACM). ACM provisions and manages TLS certificates automatically, and integrating ACM with ELB enables encrypted communication without manual certificate management.

For encryption at rest in Amazon S3, AWS best practices recommend enforcing server-side encryption using AWS KMS. An S3 bucket policy that denies PutObject requests unless the x-amz-server-side-encryption condition is present ensures that all uploaded objects are encrypted at rest using KMS-managed keys. This provides strong encryption guarantees with minimal operational effort.

Option A is unnecessary because Amazon Inspector focuses on vulnerability assessment, not encryption or network anomaly detection. Option C adds network complexity and is not required to

meet the stated requirements. Option E is incorrect because x-amz-meta-side-encryption is not a valid enforcement mechanism.

Referenced AWS Specialty Documents:

AWS Certified Security – Specialty Official Study Guide

Amazon GuardDuty Threat Detection

AWS Certificate Manager and ELB Integration

Amazon S3 Encryption Best Practices

Question: 13

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

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

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

Configure the Lambda function to publish messages to an Amazon Simple Notification Service (Amazon SNS) topic and to delete the unencrypted resource.

Answer: A

Explanation:

AWS Config provides managed rules that continuously evaluate resource configurations against compliance requirements. The AWS Certified Security – Specialty documentation highlights AWS Config managed rules as the preferred mechanism for enforcing configuration compliance at scale. The managed rule for encrypted RDS storage automatically detects DB instances and clusters that are created without encryption enabled.

By configuring automatic remediation, AWS Config can immediately invoke corrective actions without manual intervention. Integrating remediation with an Amazon SNS topic enables automated email notifications, while an AWS Lambda function can terminate the noncompliant resource. This creates a fully automated detect-alert-remediate workflow.

Option B requires manual remediation, which increases operational effort and delays enforcement. Options C and D rely on Amazon EventBridge, which evaluates events rather than configuration state and does not provide continuous compliance monitoring. AWS Config is explicitly designed for configuration compliance and governance use cases.

This solution aligns with AWS governance best practices by combining continuous monitoring, automated remediation, and centralized alerting with minimal operational overhead.

Referenced AWS Specialty Documents:

AWS Certified Security – Specialty Official Study Guide

AWS Config Managed Rules

AWS Config Automatic Remediation

Question: 14

A company uses AWS Organizations to manage an organization that consists of three workload OUs: Production, Development, and Testing. The company uses AWS CloudFormation templates to define

and deploy workload infrastructure in AWS accounts that are associated with the OUs. Different SCPs are attached to each workload OU.

The company successfully deployed a CloudFormation stack update to workloads in the Development OU and the Testing OU. When the company uses the same CloudFormation template to deploy the stack update in an account in the Production OU, the update fails. The error message reports insufficient IAM permissions.

What is the FIRST step that a security engineer should take to troubleshoot this issue?

- A. Review the AWS CloudTrail logs in the account in the Production OU. Search for any failed API calls from CloudFormation during the deployment attempt.
- B. Remove all the SCPs that are attached to the Production OU. Rerun the CloudFormation stack update to determine if the SCPs were preventing the CloudFormation API calls.
- C. Confirm that the role used by CloudFormation has sufficient permissions to create, update, and delete the resources that are referenced in the CloudFormation template.
- D. Make all the SCPs that are attached to the Production OU the same as the SCPs that are attached to the Testing OU.

Answer: A

Explanation:

AWS CloudTrail provides a record of all API calls made in an AWS account, including calls initiated by AWS CloudFormation. According to the AWS Certified Security – Specialty Study Guide, CloudTrail is the primary source for troubleshooting authorization failures because it records denied actions and the policy type that caused the denial, including service control policies.

Reviewing CloudTrail logs allows a security engineer to identify which specific API calls failed during the CloudFormation deployment and whether the denial was caused by an SCP, an IAM policy, or a permission boundary. This evidence-based approach is the recommended first step before making any configuration changes.

Option B is unsafe and violates governance best practices by removing SCPs in production. Option C may be necessary later, but it does not identify whether SCPs are the root cause. Option D introduces unnecessary risk and bypasses the purpose of differentiated controls across OUs.

AWS documentation emphasizes observing and validating before modifying security controls, making CloudTrail log analysis the correct initial troubleshooting step.

Referenced AWS Specialty Documents:

AWS Certified Security – Specialty Official Study Guide

AWS Organizations Service Control Policies

AWS CloudTrail Authorization Failure Analysis

Question: 15

A company stores infrastructure and application code in web-based, third-party, Git-compatible code repositories outside of AWS. The company wants to give the code repositories the ability to securely authenticate and assume an existing IAM role within the company's AWS account by using OpenID Connect (OIDC).

Which solution will meet these requirements?

- A. Create an OIDC identity provider (IdP) by using AWS Identity and Access Management (IAM) federation. Modify the trust policy of the IAM role to allow the code repositories to assume the IAM role.
- B. Use AWS Identity and Access Management (IAM) Roles Anywhere to create a trust anchor that uses OIDC. Modify the trust policy of the IAM role to allow the code repositories to assume the IAM role.
- C. Set up an account instance of AWS IAM Identity Center. Configure access to the code repositories as a customer managed OIDC application. Grant the application access to the IAM role.
- D. Use AWS Resource Access Manager (AWS RAM) to create a new resource share that uses OIDC. Limit the resource share to the specified code repositories. Grant the IAM role access to the resource share.

Answer: A

Explanation:

AWS IAM supports identity federation by allowing external identity providers that use OpenID Connect (OIDC) to authenticate and assume IAM roles. According to the AWS Certified Security – Specialty documentation, IAM OIDC identity providers are the recommended approach for enabling third-party systems, such as external CI/CD pipelines or Git-based repositories, to securely obtain temporary AWS credentials without using long-term access keys.

By creating an OIDC identity provider in IAM and configuring the IAM role trust policy to trust the external IdP, the company enables secure, token-based authentication. The trust policy can include conditions that restrict which repositories, branches, or workflows are allowed to assume the role, enforcing least privilege. AWS Security Specialty guidance emphasizes that this method eliminates static credentials and relies on short-lived tokens issued by the OIDC provider.

Option B is incorrect because IAM Roles Anywhere is designed for workloads running outside AWS that use X.509 certificates, not OIDC. Option C is intended for workforce identity federation, not machine-to-machine authentication. Option D is invalid because AWS RAM does not provide identity federation or authentication capabilities.

This solution aligns with AWS best practices for secure, scalable, and low-overhead authentication for external workloads.

Referenced AWS Specialty Documents:

AWS Certified Security – Specialty Official Study Guide

AWS IAM OIDC Identity Providers

AWS IAM Role Trust Policies

Question: 16

A company wants to establish separate AWS Key Management Service (AWS KMS) keys to use for different AWS services. The company's security engineer created a key policy to allow the infrastructure deployment team to create encrypted Amazon Elastic Block Store (Amazon EBS) volumes by assuming the InfrastructureDeployment IAM role. The security engineer recently discovered that IAM roles other than the InfrastructureDeployment role used this key for other services.

Which change to the policy should the security engineer make to resolve these issues?

A. In the statement block that contains the Sid "Allow use of the key", under the "Condition" block,

change StringEquals to StringLike.

B. In the policy document, remove the statement block that contains the Sid "Enable IAM User Permissions". Add key management policies to the KMS policy.

C. In the statement block that contains the Sid "Allow use of the key", under the "Condition" block, change the kms:ViaService value to ec2.us-east-1.amazonaws.com.

D. In the policy document, add a new statement block that grants the kms:Disable* permission to the security engineer's IAM role.

Answer: C

Explanation:

AWS KMS key policies can restrict how and where a key is used by leveraging condition keys such as kms:ViaService. According to the AWS Certified Security – Specialty documentation, kms:ViaService limits key usage to requests that originate from a specific AWS service in a specific Region. If this condition is overly broad or incorrect, other IAM roles and services may unintentionally use the key.

By explicitly setting the kms:ViaService condition value to ec2.us-east-1.amazonaws.com, the key policy ensures that the KMS key can only be used when requests are made through the Amazon EC2 service in that Region, such as for EBS volume encryption. This prevents other services or unintended IAM roles from using the key.

Option A weakens the condition logic and can broaden access. Option B removes essential permissions that allow IAM policies to function with KMS keys and is not recommended. Option D relates to administrative control of the key, not service-level usage restrictions.

AWS best practices recommend using kms:ViaService and precise condition values to enforce service-specific key usage and strong separation of duties.

Referenced AWS Specialty Documents:

AWS Certified Security – Specialty Official Study Guide

AWS KMS Key Policy Condition Keys

AWS KMS Best Practices

Question: 17

A consultant agency needs to perform a security audit for a company's production AWS account. Several consultants need access to the account. The consultant agency already has its own AWS account. The company requires multi-factor authentication (MFA) for all access to its production account. The company also forbids the use of long-term credentials.

Which solution will provide the consultant agency with access that meets these requirements?

- A. Create an IAM group. Create an IAM user for each consultant. Add each user to the group. Turn on MFA for each consultant.
- B. Configure Amazon Cognito on the company's production account to authenticate against the consultant agency's identity provider (IdP). Add MFA to a Cognito user pool.
- C. Create an IAM role in the consultant agency's AWS account. Define a trust policy that requires MFA. In the trust policy, specify the company's production account as the principal. Attach the trust policy to the role.
- D. Create an IAM role in the company's production account. Define a trust policy that requires MFA. In the trust policy, specify the consultant agency's AWS account as the principal. Attach the trust policy to the role.

Answer: D

Explanation:

AWS best practices strongly discourage the use of long-term credentials and recommend cross-account IAM roles with temporary credentials for third-party access. According to the AWS Certified Security – Specialty Study Guide, creating an IAM role in the resource-owning account and allowing a trusted external AWS account to assume that role is the recommended pattern for external access.

By creating the IAM role in the company's production account and specifying the consultant agency's AWS account as the trusted principal, the company retains full control over permissions. The trust policy can enforce MFA by using the `aws:MultiFactorAuthPresent` condition key, ensuring that all access requires MFA. Access is granted through AWS Security Token Service (STS), which issues short-lived credentials.

Option A violates the requirement to avoid long-term credentials. Option B is designed for

application user authentication, not AWS account access. Option C incorrectly places the role in the consultant's account, reducing the company's control over access.

This solution satisfies MFA enforcement, eliminates long-term credentials, and aligns with AWS third-party access best practices.

Referenced AWS Specialty Documents:

AWS Certified Security – Specialty Official Study Guide

AWS IAM Cross-Account Access

AWS STS and MFA Enforcement

Question: 18

A company runs an internet-accessible application on several Amazon EC2 instances that run Windows Server. The company used an instance profile to configure the EC2 instances. A security team currently accesses the VPC that hosts the EC2 instances by using an AWS Site-to-Site VPN tunnel from an on-premises office.

The security team issues a policy that requires all external access to the VPC to be blocked in the event of a security incident. However, during an incident, the security team must be able to access the EC2 instances to obtain forensic information on the instances.

Which solution will meet these requirements?

- A. Install EC2 Instance Connect on the EC2 instances. Update the IAM policy for the IAM role to grant the required permissions. Use the AWS CLI to open a tunnel to connect to the instances.
- B. Install EC2 Instance Connect on the EC2 instances. Configure the instances to permit access to the `ec2-instance-connect` command user. Use the AWS Management Console to connect to the EC2 instances.
- C. Create an EC2 Instance Connect endpoint in the VPC. Configure an appropriate security group to allow access between the EC2 instances and the endpoint. Use the AWS CLI to open a tunnel to connect to the instances.
- D. Create an EC2 Instance Connect endpoint in the VPC. Configure an appropriate security group to allow access between the EC2 instances and the endpoint. Use the AWS Management Console to connect to the EC2 instances.

Answer: D

Explanation:

EC2 Instance Connect endpoints provide secure, private connectivity to EC2 instances without requiring public IP addresses, inbound internet access, or VPN connectivity. According to AWS Certified Security – Specialty documentation, Instance Connect endpoints are designed specifically for incident response and secure administrative access scenarios.

By deploying an EC2 Instance Connect endpoint in the VPC, the security team can block all external network access while still maintaining controlled access to EC2 instances through the AWS Management Console. The endpoint uses AWS-managed infrastructure and private connectivity, and access is authorized using IAM policies and instance profiles.

Options A and B rely on direct EC2 Instance Connect installation and network paths that may still depend on external access. Option C is incorrect because tunneling is not required when using the console-based Instance Connect endpoint.

This solution enables forensic access during incidents without reopening external network paths, aligning with AWS incident response best practices.

Referenced AWS Specialty Documents:

AWS Certified Security – Specialty Official Study Guide

EC2 Instance Connect Endpoint Architecture

AWS Incident Response Best Practices

Question: 19

A company recently experienced a malicious attack on its cloud-based environment. The company successfully contained and eradicated the attack. A security engineer is performing incident response work. The security engineer needs to recover an Amazon RDS database cluster to the last known good version. The database cluster is configured to generate automated backups with a retention period of 14 days. The initial attack occurred 5 days ago at exactly 3:15 PM.

Which solution will meet this requirement?

- A. Identify the Regional cluster ARN for the database. Use the ARN to restore the Regional cluster by using the restore to point in time feature. Set a target time 5 days ago at 3:14 PM.
- B. Identify the Regional cluster ARN for the database. List snapshots that have been taken of the cluster. Restore the database by using the snapshot that has a creation time that is closest to 5 days ago at 3:14 PM.
- C. List all snapshots that have been taken of all the company's RDS databases. Identify the snapshot that was taken closest to 5 days ago at 3:14 PM and restore it.
- D. Identify the Regional cluster ARN for the database. Use the ARN to restore the Regional cluster by using the restore to point in time feature. Set a target time 14 days ago.

Answer: A

Explanation:

Amazon RDS supports point-in-time recovery (PITR) using automated backups within the configured retention window. According to the AWS Certified Security – Specialty Study Guide, PITR allows recovery to any second within the retention period, making it the most precise recovery method following a security incident.

By restoring the database cluster to a point just before the attack occurred, such as 3:14 PM, the security engineer ensures that the restored database reflects the last known good state without including malicious changes. This method is more accurate than restoring from snapshots, which are created at fixed intervals and may not align with the exact recovery time.

Options B and C rely on snapshot timing and may reintroduce compromised data. Option D restores to an arbitrary time and does not meet the requirement to recover to the last known good version.

AWS documentation explicitly recommends point-in-time recovery for incident response scenarios that require precise restoration.

Referenced AWS Specialty Documents:

AWS Certified Security – Specialty Official Study Guide

Amazon RDS Automated Backups and PITR

AWS Incident Response and Recovery Guidance

Question: 20

A company has an AWS account that hosts a production application. The company receives an email notification that Amazon GuardDuty has detected an `Impact:IAMUser/AnomalousBehavior` finding in the account. A security engineer needs to run the investigation playbook for this security incident and must collect and analyze the information without affecting the application.

Which solution will meet these requirements MOST quickly?

- A. Log in to the AWS account by using read-only credentials. Review the GuardDuty finding for details about the IAM credentials that were used. Use the IAM console to add a `DenyAll` policy to the IAM principal.
- B. Log in to the AWS account by using read-only credentials. Review the GuardDuty finding to determine which API calls initiated the finding. Use Amazon Detective to review the API calls in context.
- C. Log in to the AWS account by using administrator credentials. Review the GuardDuty finding for details about the IAM credentials that were used. Use the IAM console to add a `DenyAll` policy to the IAM principal.
- D. Log in to the AWS account by using read-only credentials. Review the GuardDuty finding to determine which API calls initiated the finding. Use AWS CloudTrail Insights and AWS CloudTrail Lake to review the API calls in context.

Answer: B

Explanation:

Amazon GuardDuty findings provide high-level detection of suspicious activity but are not designed for deep investigation on their own. The AWS Certified Security – Specialty documentation explains that Amazon Detective is purpose-built to support rapid investigations by automatically collecting, correlating, and visualizing data from GuardDuty, AWS CloudTrail, and VPC Flow Logs. Detective enables security engineers to analyze API calls, user behavior, and resource interactions in context without making any changes to the environment.

Using read-only credentials ensures that the investigation does not impact the production