

EXAM PREPARATION MATERIAL

Practice files designed to help you understand exam patterns, question styles and the overall assessment format.

IMPORTANT NOTE



Please note that these files may not be up to date. However, the questions will help you understand the exam format and typical question patterns.



Stay connected with our support team for the latest file updates and exam-related information.

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 longterm 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

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 servicespecific 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 crossaccount 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 shortlived 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:

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 application. Amazon Detective allows investigators to pivot directly from a GuardDuty finding into a detailed activity graph, showing which IAM user made anomalous calls, what resources were accessed, and how behavior deviated from the baseline. This significantly accelerates incident investigation.

Options A and C involve applying DenyAll policies, which are containment actions and could affect application availability. Option D requires manual analysis and setup and is slower than using Amazon Detective, which is designed for immediate investigative workflows.

AWS incident response guidance recommends using Detective for rapid, non-intrusive analysis after GuardDuty findings.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

Amazon GuardDuty and Amazon Detective Integration

AWS Incident Response Investigation Best Practices

Question: 21

A security engineer needs to control access to data that is encrypted with an AWS Key Management Service (AWS KMS) customer managed key. The security engineer also needs to use additional authenticated data (AAD) to prevent tampering with ciphertext.

Which solution will meet these requirements?

- A. Pass the key alias to AWS KMS when calling the Encrypt and Decrypt API actions.
- B. Use IAM policies to restrict access to the Encrypt and Decrypt API actions.
- C. Use the kms:EncryptionContext condition key when defining IAM policies for the customer managed key.

D. Use key policies to restrict access to the appropriate IAM groups.

Answer: C

Explanation:

AWS KMS supports additional authenticated data (AAD) through the use of encryption context. According to the AWS Certified Security - Specialty documentation, encryption context is a set of key-value pairs that is cryptographically bound to the ciphertext. Any attempt to decrypt the data must include the same encryption context, or decryption will fail. This mechanism protects against ciphertext tampering and unauthorized reuse.

The kms:EncryptionContext condition key allows security engineers to enforce the use of specific encryption context values in IAM or key policies. By defining conditions that require particular encryption context attributes, access to encrypted data can be tightly controlled and bound to specific applications, environments, or workflows.

Option A does not provide integrity protection. Option B controls access but does not enforce the use of AAD. Option D restricts administrative access but does not address encryption context enforcement.

AWS documentation explicitly states that encryption context combined with policy conditions is the recommended method to implement authenticated encryption and fine-grained access control with KMS.

Referenced AWS Specialty Documents:

[AWS Certified Security - Specialty Official Study Guide](#)

[AWS KMS Encryption Context](#)

[AWS KMS Policy Condition Keys](#)

Question: 22

A company uses AWS IAM Identity Center to manage access to its AWS accounts. The accounts are in an organization in AWS Organizations. A security engineer needs to set up delegated administration of IAM Identity Center in the organization's management account.

Which combination of steps should the security engineer perform in IAM Identity Center before configuring delegated administration? (Select THREE.)

- A. Grant least privilege access to the organization's management account.
- B. Create a new IAM Identity Center directory in the organization's management account.
- C. Set up a second AWS Region in the organization's management account.
- D. Create permission sets for use only in the organization's management account.
- E. Create IAM users for use only in the organization's management account.
- F. Create user assignments only in the organization's management account.

Answer: B, D, F

Explanation:

AWS IAM Identity Center delegated administration requires foundational configuration to be completed in the organization's management account before delegation. According to the AWS Certified Security - Specialty documentation, IAM Identity Center must be enabled with a directory in the management account before any delegation can occur.

Permission sets must be created in the management account because they define the permissions that will later be delegated to member accounts. Additionally, user assignments must initially exist in the management account to establish baseline access control before delegation is configured.

Option A is too generic and not a required prerequisite step. Option C is unrelated to Identity Center delegation. Option E is incorrect because IAM Identity Center uses identities from its directory or external IdPs, not IAM users.

AWS guidance clearly outlines directory creation, permission set definition, and initial user assignments as mandatory preparatory steps for delegated administration.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

AWS IAM Identity Center Delegated Administration
AWS Organizations and Identity Center Integration

Question: 23

A security engineer needs to implement a solution to identify any sensitive data that is stored in an Amazon S3 bucket. The solution must report on sensitive data in the S3 bucket by using an existing Amazon Simple Notification Service (Amazon SNS) topic.

Which solution will meet these requirements with the LEAST implementation effort?

- A. Enable AWS Config. Configure AWS Config to monitor for sensitive data in the S3 bucket and to send notifications to the SNS topic.
- B. Create an AWS Lambda function to scan the S3 bucket for sensitive data that matches a pattern. Program the Lambda function to send notifications to the SNS topic.
- C. Configure Amazon Macie to use managed data identifiers to identify and categorize sensitive data. Create an Amazon EventBridge rule to send notifications to the SNS topic.
- D. Enable Amazon GuardDuty. Configure AWS CloudTrail S3 data events. Create an Amazon CloudWatch alarm that reacts to GuardDuty findings and sends notifications to the SNS topic.

Answer: C

Explanation:

Amazon Macie is the AWS service designed specifically to discover, classify, and report sensitive data stored in Amazon S3. According to the AWS Certified Security - Specialty Study Guide, Macie uses machine learning and managed data identifiers to automatically detect sensitive data types such as PII and financial information.

Macie integrates natively with Amazon EventBridge, allowing findings to be routed to other services such as Amazon SNS with minimal configuration. Creating an EventBridge rule to forward Macie findings to an existing SNS topic satisfies the notification requirement without custom code.

Option A is invalid because AWS Config does not inspect object contents. Option B requires custom development and ongoing maintenance. Option D is incorrect because Amazon GuardDuty focuses on threat detection, not sensitive data discovery.

AWS documentation emphasizes Macie as the lowest-effort and most accurate solution for sensitive data identification in S3.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

Amazon Macie Sensitive Data Discovery

Amazon EventBridge Integration with Security Services

Question: 24

An application is running on an Amazon EC2 instance that has an IAM role attached. The IAM role provides access to an AWS Key Management Service (AWS KMS) customer managed key and an Amazon S3 bucket. The key is used to access 2 TB of sensitive data that is stored in the S3 bucket. A security engineer discovers a potential vulnerability on the EC2 instance that could result in the compromise of the sensitive data.

a. Due to other critical operations, the security engineer cannot immediately shut down the EC2 instance for vulnerability patching.

What is the FASTEST way to prevent the sensitive data from being exposed?

- A. Download the data from the existing S3 bucket to a new EC2 instance. Then delete the data from the S3 bucket. Re-encrypt the data with a client-based key. Upload the data to a new S3 bucket.
- B. Block access to the public range of S3 endpoint IP addresses by using a host-based firewall. Ensure that internet-bound traffic from the affected EC2 instance is routed through the host-based firewall.
- C. Revoke the IAM role's active session permissions. Update the S3 bucket policy to deny access to the IAM role. Remove the IAM role from the EC2 instance profile.
- D. Disable the current key. Create a new KMS key that the IAM role does not have access to, and reencrypt all the data with the new key. Schedule the compromised key for deletion.

Answer: C

Explanation:

AWS incident response best practices emphasize rapid containment to prevent further data exposure. According to the AWS Certified Security - Specialty Study Guide, the fastest and least disruptive containment method for compromised compute resources is to immediately revoke credentials and permissions rather than modifying data or infrastructure.

Revoking the IAM role's active sessions prevents the EC2 instance from continuing to access AWS services. Updating the S3 bucket policy to explicitly deny access to the IAM role ensures immediate enforcement, even if temporary credentials remain cached. Removing the IAM role from the instance profile further prevents new credentials from being issued.

Option A and D involve large-scale data movement or re-encryption, which is time-consuming and operationally expensive. Option B relies on network-level controls that do not prevent access through private AWS endpoints.

AWS guidance explicitly recommends credential revocation and policy-based denial as the fastest containment step during active incidents.

Referenced AWS Specialty Documents:

[AWS Certified Security - Specialty Official Study Guide](#)

[AWS Incident Response Best Practices](#)

[AWS IAM Role Session Management](#)

Question: 25

A company needs a cloud-based, managed desktop solution for its workforce of remote employees. The company wants to ensure that the employees can access the desktops only by using company-provided devices. A security engineer must design a solution that will minimize cost and management overhead.

Which solution will meet these requirements?

- A. Deploy a custom virtual desktop infrastructure (VDI) solution with a restriction policy to allow access only from corporate devices.
- B. Deploy a fleet of Amazon EC2 instances. Assign an instance to each employee with certificate-based device authentication that uses Windows Active Directory.
- C. Deploy Amazon WorkSpaces. Set up a trusted device policy with IP blocking on the authentication gateway by using AWS Identity and Access Management (IAM).
- D. Deploy Amazon WorkSpaces. Create client certificates, and deploy them to trusted devices. Enable restricted access at the directory level.

Answer: D

Explanation:

Amazon WorkSpaces is a fully managed desktop-as-a-service solution designed to minimize infrastructure and operational overhead. According to AWS Certified Security - Specialty documentation, WorkSpaces supports device trust by using client certificates to restrict access to approved devices.

By deploying client certificates only to company-managed devices and enforcing restricted access at the directory level, the organization ensures that only trusted endpoints can authenticate. This approach avoids the cost and complexity of building and maintaining a custom VDI or managing individual EC2 instances.

Option A and B significantly increase management overhead. Option C is incorrect because IAM does not manage WorkSpaces authentication gateway policies or device trust.

AWS best practices highlight Amazon WorkSpaces with certificate-based device trust as the most efficient solution for secure, managed desktops.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

Amazon WorkSpaces Security Controls

Amazon WorkSpaces Device Trust

Question: 26

A company runs an application on an Amazon EC2 instance. The application generates invoices and stores them in an Amazon S3 bucket. The instance profile that is attached to the instance has appropriate access to the S3 bucket. The company needs to share each invoice with multiple clients that do not have AWS credentials. Each client must be able to download only the client's own invoices. Clients must download their invoices within 1 hour of invoice creation. Clients must use only temporary credentials to access the company's AWS resources.

Which additional step will meet these requirements?

A. Update the S3 bucket policy to ensure that clients that use pre-signed URLs have the S3:Get* permission and the S3:List* permission to access S3 objects in the bucket.

- B. Add a StringEquals condition to the IAM role policy for the EC2 instance profile. Configure the policy condition to restrict access based on the s3:ResourceTag/ClientId tag of each invoice. Tag each generated invoice with the ID of its corresponding client.
- C. Update the script to use AWS Security Token Service (AWS STS) to obtain new credentials each time the script runs by assuming a new role that has S3:GetObject permissions. Use the credentials to generate the pre-signed URLs.
- D. Generate an access key and a secret key for an IAM user that has S3:GetObject permissions on the S3 bucket. Embed the keys into the script. Use the keys to generate the pre-signed URLs.

Answer: B

Explanation:

Amazon S3 pre-signed URLs grant temporary access based on the permissions of the principal that generates them. AWS Certified Security - Specialty documentation explains that fine-grained authorization can be enforced by combining pre-signed URLs with IAM policy conditions.

By tagging each invoice object with a client identifier and adding a condition to the EC2 instance role policy using s3:ResourceTag/ClientId, the role can generate pre-signed URLs only for objects associated with a specific client. This ensures that each client can access only their own invoices, even though the URLs are temporary and unauthenticated.

Option A over-permissions clients. Option C is unnecessary because instance profiles already use temporary credentials. Option D violates AWS best practices by using long-term credentials.

AWS recommends resource tagging with IAM policy conditions for scalable, secure access control.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

Amazon S3 Pre-Signed URLs

IAM Policy Conditions and Resource Tags

Question: 27

A company is running an application in the eu-west-1 Region. The application uses an AWS Key Management Service (AWS KMS) customer managed key to encrypt sensitive data.

a. 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 AWS 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 customer managed key as the application in eu-west-1.
- B. Allocate a new customer managed key to eu-north-1 to be used by the application that is deployed in that Region.
- C. Allocate a new customer managed key to eu-north-1. Create the same alias name for both keys. Configure the application deployment to use the key alias.
- D. Allocate a new customer managed key to eu-north-1. Create an alias for eu--1. Change the application code to point to the alias for eu--1.

Answer: C

Explanation:

AWS KMS keys are regional resources and cannot be used across Regions. According to AWS Certified Security - Specialty documentation, applications that are deployed in multiple Regions should use region-specific customer managed keys while referencing keys by alias instead of key ID.

By creating a new customer managed key in eu-north-1 and assigning it the same alias as the key in eu-west-1, the application code can continue to reference the alias without modification. Each Region resolves the alias to the correct local key, ensuring encryption continues to function correctly.

Option A is invalid because KMS keys are regional. Option B requires application changes. Option D introduces unsupported alias patterns.

AWS best practices recommend alias-based key references for multi-Region deployments.

Referenced AWS Specialty Documents:

[AWS Certified Security - Specialty Official Study Guide](#)

[AWS KMS Regional Keys and Aliases](#)

[AWS KMS Best Practices](#)

Question: 28

A company that uses AWS Organizations is using AWS IAM Identity Center to administer access to AWS accounts. A security engineer is creating a custom permission set in IAM Identity Center. The company will use the permission set across multiple accounts. An AWS managed policy and a customer managed policy are attached to the permission set. The security engineer has full administrative permissions and is operating in the management account.

When the security engineer attempts to assign the permission set to an IAM Identity Center user who has access to multiple accounts, the assignment fails.

What should the security engineer do to resolve this failure?

- A. Create the customer managed policy in every account where the permission set is assigned. Give the customer managed policy the same name and same permissions in each account.
- B. Remove either the AWS managed policy or the customer managed policy from the permission set. Create a second permission set that includes the removed policy. Apply the permission sets separately to the user.
- C. Evaluate the logic of the AWS managed policy and the customer managed policy. Resolve any policy conflicts in the permission set before deployment.
- D. Do not add the new permission set to the user. Instead, edit the user's existing permission set to include the AWS managed policy and the customer managed policy.

Answer: A

Explanation:

AWS IAM Identity Center permission sets that include customer managed policies require those policies to exist in each target account. According to the AWS Certified Security - Specialty Study Guide, customer managed policies are account-scoped and are not automatically propagated across accounts by Identity Center.

When assigning a permission set across multiple accounts, Identity Center attempts to attach the referenced customer managed policy in each account. If the policy does not exist, the assignment fails. Creating the same customer managed policy with identical name and permissions in every target account resolves the issue.

Option B increases complexity. Option C does not address the root cause. Option D violates Identity Center management best practices.

AWS documentation clearly states that customer managed policies must be present in all accounts where permission sets are applied.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

AWS IAM Identity Center Permission Sets

AWS Organizations and Identity Center Policy Management

Question: 29

A company has a single AWS account and uses an Amazon EC2 instance to test application code. The company recently discovered that the instance was compromised and was serving malware. Analysis showed that the instance was compromised 35 days ago. A security engineer must implement a continuous monitoring solution that automatically notifies the security team by email for high severity findings as soon as possible.

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

- A. Enable AWS Security Hub in the AWS account.
- B. Enable Amazon GuardDuty in the AWS account.
- C. Create an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the security team's email distribution list to the topic.
- D. Create an Amazon Simple Queue Service (Amazon SQS) queue. Subscribe the security team's email distribution list to the queue.
- E. Create an Amazon EventBridge rule for GuardDuty findings of high severity. Configure the rule to publish a message to the topic.
- F. Create an Amazon EventBridge rule for Security Hub findings of high severity. Configure the rule to publish a message to the queue.

Answer: B, C, E

Explanation:

Amazon GuardDuty provides continuous threat detection for compromised instances by analyzing VPC Flow Logs, DNS logs, and CloudTrail events. According to AWS Certified Security - Specialty guidance, GuardDuty is the fastest service to enable for detecting malware and compromised EC2 instances.

To notify the security team, Amazon SNS provides a native email notification mechanism with minimal setup. Amazon EventBridge integrates directly with GuardDuty findings and can filter based on severity. Creating an EventBridge rule that matches high severity GuardDuty findings and publishes to SNS ensures immediate notification.

Security Hub is not required for this use case and adds additional setup time. Amazon SQS does not support email subscriptions.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

Amazon GuardDuty Findings and Severity

Amazon EventBridge Integration with GuardDuty

Question: 30

A company has a PHP-based web application that uses Amazon S3 as an object store for user files. The S3 bucket is configured for server-side encryption with Amazon S3 managed keys (SSE-S3). New requirements mandate full control of encryption keys.

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

- A. Create a new customer managed key in AWS Key Management Service (AWS KMS).
- B. Change the SSE-S3 configuration on the S3 bucket to server-side encryption with customer- provided keys (SSE-C).
- C. Configure the PHP SDK to use the SSE-S3 key before upload.
- D. Create an AWS managed key for Amazon S3 in AWS KMS.
- E. Change the SSE-S3 configuration on the S3 bucket to server-side encryption with AWS KMS managed keys (SSE-KMS).
- F. Change all the S3 objects in the bucket to use the new encryption key.

Answer: A, E, F

Explanation:

SSE-S3 uses AWS-managed keys and does not provide customer control. AWS Certified Security - Specialty documentation states that SSE-KMS with customer managed keys allows full control, auditing, and key rotation. The security engineer must first create a customer managed KMS key, then update the bucket to use SSE-KMS. Existing objects must be re-encrypted to ensure compliance.

SSE-C requires the application to manage keys, increasing complexity and risk. AWS managed keys do not meet the requirement for customer-controlled encryption.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

Amazon S3 Encryption Options

AWS KMS Customer Managed Keys

Question: 31

A company runs a public web application on Amazon EKS behind Amazon CloudFront and an Application Load Balancer (ALB). A security engineer must send a notification to an existing Amazon SNS topic when the

application receives 10,000 requests from the same end-user IP address within any 5-minute period.

Which solution will meet these requirements?

- A. Configure CloudFront standard logging and CloudWatch Logs metric filters.
- B. Configure VPC Flow Logs and CloudWatch Logs metric filters.
- C. Configure an AWS WAF web ACL with an ASN match rule and CloudWatch alarms.
- D. Configure an AWS WAF web ACL with a rate-based rule. Associate it with CloudFront. Create a CloudWatch alarm to notify SNS.

Answer: D

Explanation:

AWS WAF rate-based rules are designed specifically to track the number of requests from a single IP address over a configurable time window. According to AWS Certified Security - Specialty guidance, rate-based rules integrate natively with CloudFront and emit CloudWatch metrics that can trigger alarms.

CloudFront logs and VPC Flow Logs are not real-time detection tools. ASN match rules do not count request rates.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

AWS WAF Rate-Based Rules

CloudFront and AWS WAF Integration

Question: 32

A company needs to build a code-signing solution using an AWS KMS asymmetric key and must store immutable evidence of key creation and usage for compliance and audit purposes.

Which solution meets these requirements?

- A. Create an Amazon S3 bucket with S3 Object Lock enabled. Create an AWS CloudTrail trail with log file validation enabled for KMS events. Store logs in the bucket and grant auditors access.
- B. Log application events to Amazon CloudWatch Logs and export them.
- C. Capture KMS API calls using EventBridge and store them in DynamoDB.
- D. Track KMS usage with CloudWatch metrics and dashboards.

Answer: A

Explanation:

AWS CloudTrail provides authoritative records of KMS key creation, origin, and usage. Enabling log file validation ensures tamper detection. S3 Object Lock in compliance mode enforces immutability, which is a core audit requirement cited in AWS Certified Security - Specialty materials.

CloudWatch and DynamoDB do not provide immutable storage guarantees suitable for compliance evidence.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

AWS CloudTrail Log File Validation

Amazon S3 Object Lock

Question: 33

A company's application team needs a new AWS Key Management Service (AWS KMS) customer managed key to use with Amazon S3. The company's security policy requires separate keys for different AWS services to limit security exposure.

How can a security engineer limit the KMS customer managed key to work with only Amazon S3?

- A. Configure the key policy to allow only Amazon S3 to perform the kms:Encrypt action.
- B. Configure the key policy to allow KMS actions only when the value for the kms:ViaService condition key matches the Amazon S3 service name.
- C. Configure the application's IAM role policy to allow Amazon S3 to perform the iam:PassRole action.
- D. Configure the application's IAM role policy to allow only S3 operations when the operations are combined with the KMS customer managed key.

Answer: B

Explanation:

AWS KMS provides condition keys that can be used to tightly scope how and where a customer managed key can be used. According to the AWS Certified Security - Specialty Study Guide, the kms:ViaService condition key is specifically designed to restrict key usage to requests that originate from a particular AWS service in a specific Region.

By configuring the key policy to allow KMS cryptographic operations only when kms:ViaService equals s3.<region>.amazonaws.com, the security engineer ensures that the key can be used exclusively by Amazon S3. Even if other IAM principals have permissions to use the key, the key cannot be used by other services such as Amazon EC2, Amazon RDS, or AWS Lambda.

Option A is incorrect because AWS services do not assume identities in key policies. Options C and D modify IAM role policies, which do not control how a KMS key is used by AWS services. AWS documentation clearly states that service-level restrictions must be enforced at the KMS key policy level using condition keys.

This approach enforces strong separation of duties and limits blast radius, which aligns with AWS security best practices.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

AWS KMS Key Policy Condition Keys

AWS KMS Best Practices

Question: 34

A company uses AWS to run a web application that manages ticket sales in several countries. The company recently migrated the application to an architecture that includes Amazon API Gateway, AWS Lambda, and Amazon Aurora Serverless. The company needs the application to comply with Payment Card Industry Data Security Standard (PCI DSS) v4.0. A security engineer must generate a report that shows the effectiveness of the PCI DSS v4.0 controls that apply to the application. The company's compliance team must be able to add manual evidence to the report.

Which solution will meet these requirements?

- A. Enable AWS Trusted Advisor. Configure all the Trusted Advisor checks. Manually map the checks against the PCI DSS v4.0 standard to generate the report.
- B. Enable and configure AWS Config. Deploy the Operational Best Practices for PCI DSS conformance pack in AWS Config. Use AWS Config to generate the report.
- C. Enable AWS Security Hub. Enable the Security Hub PCI DSS security standard. Use the AWS Management Console to download the report from the security standard.
- D. Create an AWS Audit Manager assessment that uses the AWS managed PCI DSS v4.0 standard framework. Add all evidence to the assessment. Generate the report in Audit Manager for download.

Answer: D

Explanation:

AWS Audit Manager is specifically designed to help organizations continuously audit their AWS usage against

compliance frameworks and generate audit-ready reports. According to AWS Certified Security - Specialty documentation, Audit Manager includes AWS managed frameworks for compliance standards, including PCI DSS v4.0.

Audit Manager automatically collects evidence from AWS services such as API Gateway, Lambda, RDS, CloudTrail, and Config, and maps the evidence directly to PCI DSS controls. Importantly, Audit Manager allows compliance teams to upload and attach manual evidence, which is a key requirement in this scenario.

Option C provides visibility into control status but does not support adding manual evidence. Option B evaluates configuration compliance but does not generate formal compliance reports. Option A requires extensive manual effort and is not aligned with PCI reporting workflows.

AWS documentation positions Audit Manager as the authoritative service for compliance reporting and audit evidence management.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

AWS Audit Manager PCI DSS Framework

AWS Compliance Reporting Best Practices

Question: 35

A company is planning to deploy a new log analysis environment. The company needs to analyze logs from multiple AWS services in near real time. The solution must provide the ability to search the logs and must send alerts to an existing Amazon Simple Notification Service (Amazon SNS) topic when specific logs match detection rules.

Which solution will meet these requirements?

- A. Analyze the logs by using Amazon OpenSearch Service. Search the logs from the OpenSearch API. Use OpenSearch Service Security Analytics to match logs with detection rules and to send alerts to the SNS topic.
- B. Analyze the logs by using AWS Security Hub. Search the logs from the Findings page in Security Hub. Create custom actions to match logs with detection rules and to send alerts to the SNS topic.
- C. Analyze the logs by using Amazon CloudWatch Logs. Use a subscription filter to match logs with detection rules and to send alerts to the SNS topic. Search the logs manually by using CloudWatch Logs Insights.

D. Analyze the logs by using Amazon QuickSight. Search the logs by listing the query results in a dashboard. Run queries to match logs with detection rules and to send alerts to the SNS topic.

Answer: A

Explanation:

Amazon OpenSearch Service is designed for near real-time log ingestion, indexing, and search across large volumes of data. According to the AWS Certified Security - Specialty Study Guide, OpenSearch supports advanced log analytics use cases and integrates with OpenSearch Security Analytics, which provides prebuilt and custom detection rules.

Security Analytics can continuously evaluate incoming logs from multiple AWS services and generate alerts when detection rules are matched. These alerts can be forwarded to Amazon SNS with minimal configuration.

OpenSearch also provides powerful search and query capabilities through APIs and dashboards.

Option C supports detection but lacks advanced correlation and scalable search capabilities. Option B is not a log analytics service. Option D is a visualization service and does not support real-time detection.

AWS guidance recommends OpenSearch Service for centralized, near real-time log analysis and alerting.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

Amazon OpenSearch Service Security Analytics

AWS Logging and Monitoring Architecture

Question: 36

A company uses an organization in AWS Organizations to manage multiple AWS accounts. The company wants to centrally give users the ability to access Amazon Q Developer.

Which solution will meet this requirement?

A. Enable AWS IAM Identity Center and set up Amazon Q Developer as an AWS managed application.

- B. Enable Amazon Cognito and create a new identity pool for Amazon Q Developer.
- C. Enable Amazon Cognito and set up Amazon Q Developer as an AWS managed application.
- D. Enable AWS IAM Identity Center and create a new identity pool for Amazon Q Developer.

Answer: A

Explanation:

AWS IAM Identity Center is the recommended service for centrally managing workforce access across multiple AWS accounts within an organization. According to AWS Certified Security - Specialty documentation, Amazon Q Developer integrates natively with IAM Identity Center as an AWS managed application.

By enabling IAM Identity Center and assigning Amazon Q Developer to users or groups, the company can centrally control access using permission sets and organizational boundaries. This approach provides centralized authentication, authorization, and auditing with minimal overhead.

Amazon Cognito is intended for customer and application user authentication, not workforce access to AWS services. Identity pools are not applicable to IAM Identity Center integrations.

AWS best practices clearly recommend IAM Identity Center for workforce access to AWS-managed applications.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

AWS IAM Identity Center Integrations

Amazon Q Developer Access Management

Question: 37

A company stores sensitive data in an Amazon S3 bucket. The company encrypts the data at rest by using server-side encryption with Amazon S3 managed keys (SSE-S3). A security engineer must prevent any modifications to the data in the S3 bucket.

Which solution will meet this requirement?

- A. Configure S3 bucket policies to deny DELETE and PUT object permissions.
- B. Configure S3 Object Lock in compliance mode with S3 bucket versioning enabled.
- C. Change the encryption on the S3 bucket to use AWS Key Management Service (AWS KMS) customer managed keys.
- D. Configure the S3 bucket with multi-factor authentication (MFA) delete protection.

Answer: B

Explanation:

Amazon S3 Object Lock in compliance mode provides write-once-read-many (WORM) protection, which prevents objects from being modified or deleted for a specified retention period. According to the AWS Certified Security - Specialty Study Guide, compliance mode enforces immutability even for the root user and cannot be overridden.

Enabling S3 Object Lock requires S3 bucket versioning and ensures that once an object is written, it cannot be changed or removed until the retention period expires. This is the strongest protection against data modification and is commonly used for regulatory and legal retention requirements.

Option A can be bypassed by administrators. Option D only protects against deletions, not overwrites. Option C changes encryption but does not prevent modification.

AWS documentation explicitly identifies S3 Object Lock in compliance mode as the correct solution for immutable data storage.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

Amazon S3 Object Lock

Amazon S3 Data Protection and Compliance

Question: 38

A company is operating an open-source software platform that is internet facing. The legacy software platform no longer receives security updates. The software platform operates using Amazon Route 53 weighted load balancing to send traffic to two Amazon EC2 instances that connect to an Amazon RDS cluster. A recent report

suggests this software platform is vulnerable to SQL injection attacks, with samples of attacks provided. The company's security engineer must secure this system against SQL injection attacks within 24 hours. The solution must involve the least amount of effort and maintain normal operations during implementation.

What should the security engineer do to meet these requirements?

- A. Create an Application Load Balancer with the existing EC2 instances as a target group. Create an AWS WAF web ACL containing rules that protect the application from this attack, then apply it to the ALB. Test to ensure the vulnerability has been mitigated, then redirect the Route 53 records to point to the ALB. Update security groups on the EC2 instances to prevent direct access from the internet.
- B. Create an Amazon CloudFront distribution specifying one EC2 instance as an origin. Create an AWS WAF web ACL containing rules that protect the application from this attack, then apply it to the distribution. Test to ensure the vulnerability has been mitigated, then redirect the Route 53 records to point to CloudFront.
- C. Obtain the latest source code for the platform and make the necessary updates. Test the updated code to ensure that the vulnerability has been mitigated, then deploy the patched version of the platform to the EC2 instances.
- D. Update the security group that is attached to the EC2 instances, removing access from the internet to the TCP port used by the SQL database. Create an AWS WAF web ACL containing rules that protect the application from this attack, then apply it to the EC2 instances.

Answer: A

Explanation:

AWS WAF provides managed and custom rules that can immediately mitigate common web exploits such as SQL injection without modifying application code. According to AWS Certified Security - Specialty documentation, placing AWS WAF in front of an Application Load Balancer is a recommended rapid-response control for legacy applications with known vulnerabilities.

Creating an ALB in front of the existing EC2 instances allows seamless traffic migration. AWS WAF SQL injection rules can be deployed and tested without downtime. Updating Route 53 to point to the ALB preserves normal operations. Restricting EC2 security groups afterward prevents bypassing the WAF.

Option B introduces CloudFront changes and single-origin testing, increasing complexity. Option C cannot be completed within 24 hours and risks downtime. Option D is invalid because AWS WAF cannot be attached directly to EC2 instances.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

AWS WAF Web ACL Architecture

AWS Application Load Balancer Security

Question: 39

A company has a large fleet of Amazon Linux 2 Amazon EC2 instances that run an application processing sensitive data.

a. Compliance requirements include no exposed management ports, full session logging, and authentication through AWS IAM Identity Center. DevOps engineers occasionally need access for troubleshooting.

Which solution will provide remote access while meeting these requirements?

- A. Grant access to the EC2 serial console and allow IAM role access.
- B. Enable EC2 Instance Connect and configure security groups accordingly.
- C. Assign an EC2 instance role that allows access to AWS Systems Manager. Create an IAM policy that grants access to Systems Manager Session Manager and assign it to an IAM Identity Center role.
- D. Use Systems Manager Automation to temporarily open remote access ports.

Answer: C

Explanation:

AWS Systems Manager Session Manager provides secure, auditable shell access to EC2 instances without opening inbound ports. According to AWS Certified Security - Specialty guidance, Session Manager records all session activity to CloudWatch Logs or Amazon S3 and integrates with IAM Identity Center for centralized authentication.

This solution meets all requirements: no exposed ports, full audit logging, and identity-based access control. EC2 Instance Connect and serial console access do not integrate with Identity Center and may expose management paths.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

AWS Systems Manager Session Manager

AWS IAM Identity Center Integration

Question: 40

A company's data scientists use Amazon SageMaker with datasets stored in Amazon S3. Data older than 45 days must be removed according to policy.

Which action should enforce this policy?

- A. Configure an S3 Lifecycle rule to delete objects after 45 days.
- B. Create a Lambda function triggered on object upload to delete old data.
- C. Create a scheduled Lambda function to delete old objects monthly.
- D. Configure S3 Intelligent-Tiering.

Answer: A

Explanation:

Amazon S3 Lifecycle rules are the native and most efficient way to enforce data retention policies. AWS Certified Security - Specialty documentation recommends lifecycle rules over custom automation to reduce operational complexity and failure risk.

Lifecycle rules automatically and reliably delete objects after a specified age, ensuring compliance without additional compute services. Lambda-based solutions increase cost and management overhead. Intelligent-Tiering manages storage cost, not data deletion.

Referenced AWS Specialty Documents:

Question: 41

A company has a web application that reads from and writes to an Amazon S3 bucket. The company needs to authenticate all S3 API calls with AWS credentials.

Which solution will provide the application with AWS credentials?

- A. Use Amazon Cognito identity pools and the GetId API.
- B. Use Amazon Cognito identity pools and AssumeRoleWithWebIdentity.
- C. Use Amazon Cognito user pools with ID tokens.
- D. Use Amazon Cognito user pools with access tokens.

Answer: B

Explanation:

Amazon Cognito identity pools provide temporary AWS credentials by exchanging web identity tokens with AWS STS using AssumeRoleWithWebIdentity. According to AWS Certified Security - Specialty documentation, this is the correct mechanism for granting applications AWS credentials.

User pools authenticate users but do not issue AWS credentials. Identity pools integrate with IAM roles and STS, enabling secure, temporary access to AWS services.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

Amazon Cognito Identity Pools

AWS STS Web Identity Federation

Question: 42

A company runs ECS services behind an internet-facing ALB that is the origin for CloudFront. An AWS WAF web ACL is associated with CloudFront, but clients can bypass it by accessing the ALB directly.

Which solution will prevent direct access to the ALB?

- A. Use AWS PrivateLink with the ALB.
- B. Replace the ALB with an internal ALB.
- C. Restrict ALB listener rules to CloudFront IP ranges.
- D. Require a custom header from CloudFront and validate it at the ALB.

Answer: D

Explanation:

AWS best practices recommend using a shared secret header between CloudFront and ALB origins to prevent direct access. CloudFront injects a custom header, and the ALB listener rules validate its presence.

IP-based controls are brittle due to CloudFront IP changes. PrivateLink and internal ALBs are not supported as CloudFront origins. Header validation is the most reliable and widely recommended pattern.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

CloudFront Origin Protection

AWS WAF and ALB Integration

Question: 43

A company creates AWS Lambda functions from container images that are stored in Amazon Elastic Container Registry (Amazon ECR). The company needs to identify any software vulnerabilities in the container images and any code vulnerabilities in the Lambda functions.

Which solution will meet these requirements?

- A. Enable Amazon GuardDuty. Configure Amazon ECR scanning and Lambda code scanning in GuardDuty.
- B. Enable Amazon GuardDuty. Configure Runtime Monitoring and Lambda Protection in GuardDuty.
- C. Enable Amazon Inspector. Configure Amazon ECR enhanced scanning and Lambda code scanning in Amazon Inspector.
- D. Enable AWS Security Hub. Configure Runtime Monitoring and Lambda Protection in Security Hub.

Answer: C

Explanation:

Amazon Inspector is the AWS service designed specifically for vulnerability management across compute workloads, including Amazon ECR container images and AWS Lambda functions. According to the AWS Certified Security - Specialty documentation, Amazon Inspector provides automated vulnerability assessments for container images stored in ECR by performing enhanced image scanning that identifies common vulnerabilities and exposures (CVEs) in operating systems and application dependencies.

Inspector also supports Lambda code scanning to analyze function packages and container-based Lambda images for known software vulnerabilities. Findings include severity ratings and remediation guidance, allowing security teams to identify and prioritize risks efficiently.

Amazon GuardDuty focuses on threat detection using behavioral analysis and does not perform static vulnerability scanning of container images or Lambda code. AWS Security Hub aggregates findings from other services but does not perform scanning itself.

AWS best practices recommend Amazon Inspector for vulnerability detection in container images and serverless workloads.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

Amazon Inspector for ECR and Lambda

AWS Vulnerability Management Best Practices

Question: 44

A security engineer is designing a solution that will provide end-to-end encryption between clients and Docker containers running in Amazon Elastic Container Service (Amazon ECS). This solution must also handle volatile traffic patterns.

Which solution would have the MOST scalability and LOWEST latency?

- A. Configure a Network Load Balancer to terminate the TLS traffic and then re-encrypt the traffic to the containers.
- B. Configure an Application Load Balancer to terminate the TLS traffic and then re-encrypt the traffic to the containers.
- C. Configure a Network Load Balancer with a TCP listener to pass through TLS traffic to the containers.
- D. Configure Amazon Route 53 to use multivalue answer routing to send traffic to the containers.

Answer: C

Explanation:

Network Load Balancers operate at Layer 4 and are optimized for extreme performance, ultra-low latency, and handling sudden traffic spikes. According to AWS Certified Security - Specialty documentation, using a TCP listener on an NLB allows TLS traffic to pass through directly to backend containers without termination, preserving true end-to-end encryption.

This approach eliminates the overhead of decrypting and re-encrypting traffic at the load balancer, reducing latency and maximizing throughput. NLBs scale automatically to handle volatile traffic patterns and millions of requests per second.

Application Load Balancers operate at Layer 7 and introduce additional latency due to TLS termination and HTTP processing. Route 53 multivalue routing does not provide load balancing at the transport layer and does not ensure encryption handling.

AWS recommends NLB TCP pass-through for high-performance, end-to-end encrypted container workloads.

Referenced AWS Specialty Documents:

Elastic Load Balancing Architecture

Network Load Balancer Performance Characteristics

Question: 45

A security engineer has designed a VPC to segment private traffic from public traffic. The VPC includes two Availability Zones. Each Availability Zone contains one public subnet and one private subnet. Three route tables exist: one for the public subnets and one for each private subnet.

The security engineer discovers that all four subnets are routing traffic through the internet gateway that is attached to the VPC.

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

- A. Verify that a NAT gateway has been provisioned in the public subnet in each Availability Zone.
- B. Verify that a NAT gateway has been provisioned in the private subnet in each Availability Zone.
- C. Modify the route tables for the public subnets to add a local route to the VPC CIDR range.
- D. Modify the route tables for the private subnets to route 0.0.0.0/0 to the NAT gateway in the public subnet of the same Availability Zone.
- E. Modify the route tables for the private subnets to route 0.0.0.0/0 to the internet gateway.

Answer: A, D

Explanation:

AWS networking best practices require private subnets to access the internet only through NAT gateways located in public subnets. According to the AWS Certified Security - Specialty Study Guide, NAT gateways must be provisioned in public subnets and used as the default route for outbound traffic from private subnets.

Verifying NAT gateways in each Availability Zone ensures high availability and fault tolerance. Updating the private subnet route tables to send 0.0.0.0/0 traffic to the NAT gateway prevents direct internet access

while allowing outbound connectivity.

Routing private subnet traffic directly to an internet gateway violates subnet isolation principles. NAT gateways must never be placed in private subnets.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

Amazon VPC Routing and NAT Gateways

AWS Network Segmentation Best Practices

Question: 46

A company runs a web application on a fleet of Amazon EC2 instances in an Auto Scaling group. Amazon GuardDuty and AWS Security Hub are enabled. The security engineer needs an automated response to anomalous traffic that follows AWS best practices and minimizes application disruption.

Which solution will meet these requirements?

- A. Use EventBridge to disable the instance profile access keys.
- B. Use EventBridge to invoke a Lambda function that removes the affected instance from the Auto Scaling group and isolates it with a restricted security group.
- C. Use Security Hub to update the subnet network ACL to block traffic.
- D. Send GuardDuty findings to Amazon SNS for email notification.

Answer: B

Explanation:

AWS incident response best practices emphasize isolating compromised resources rather than immediately terminating them. According to AWS Certified Security - Specialty documentation, removing an instance from an Auto Scaling group prevents replacement loops, while applying a restrictive security group isolates the instance for forensic analysis.

Using Amazon EventBridge to trigger an AWS Lambda function enables automated, consistent responses to GuardDuty findings. This approach minimizes disruption to the application because healthy instances continue serving traffic while the affected instance is isolated.

Disabling credentials or modifying network ACLs can have broader impact on unrelated workloads. SNS notifications alone do not provide response automation.

AWS recommends isolate-and-investigate patterns for EC2 incident response.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

Amazon GuardDuty Automated Responses

AWS Incident Response Playbooks

Question: 47

A company's security engineer receives an alert that indicates that an unexpected principal is accessing a company-owned Amazon Simple Queue Service (Amazon SQS) queue. All the company's accounts are within an organization in AWS Organizations. The security engineer must implement a mitigation solution that minimizes compliance violations and investment in tools outside of AWS.

What should the security engineer do to meet these requirements?

- A. Create security groups and attach them to all SQS queues.
- B. Modify network ACLs in all VPCs to restrict inbound traffic.
- C. Create interface VPC endpoints for Amazon SQS. Restrict access using `aws:SourceVpce` and `aws:PrincipalOrgId` conditions.
- D. Use a third-party cloud access security broker (CASB).

Answer: C

Explanation:

Amazon SQS is a regional service that supports AWS PrivateLink through interface VPC endpoints. According to AWS Certified Security - Specialty documentation, the most secure and compliant way to restrict access to AWS services is by using VPC endpoints combined with resource-based policies.

By creating interface VPC endpoints for Amazon SQS in all VPCs, traffic to SQS remains on the AWS network and does not traverse the public internet. Using the `aws:SourceVpce` condition in the SQS queue policy ensures that only requests originating from approved VPC endpoints can access the queue. Adding the `aws:PrincipalOrgId` condition further restricts access to principals that belong to the same AWS Organization.

Security groups and network ACLs do not apply to SQS because SQS is not deployed inside a VPC. Third-party CASB tools add cost and operational overhead.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

Amazon SQS Security and VPC Endpoints

AWS Organizations Condition Keys

Question: 48

A company needs to deploy AWS CloudFormation templates that 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 encrypted parameters in the CloudFormation template.
- C. Use SecureString parameters to reference Secrets Manager.
- D. Use SecureString parameters encrypted by AWS KMS.

Answer: A

Explanation:

AWS CloudFormation supports dynamic references to AWS Secrets Manager, which allow sensitive values to be retrieved securely at stack runtime. According to AWS Certified Security - Specialty guidance, dynamic references prevent secrets from being stored in plaintext in templates, stack metadata, or logs.

Using dynamic references ensures that secrets remain encrypted at rest and are accessed only when required. CloudFormation does not support SecureString parameters for Secrets Manager references, and encrypting templates does not prevent exposure during execution.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

AWS CloudFormation Dynamic Reference

AWS Secrets Manager Best Practices

Question: 49

A company must immediately disable compromised IAM users across all AWS accounts and collect all actions performed by the user in the last 7 days.

Which solution will meet these requirements?

- A. Disable the IAM user and query CloudTrail logs in Amazon S3 using Athena.
- B. Remove IAM policies and query logs in Security Hub.
- C. Remove permission sets and query logs using CloudWatch Logs Insights.
- D. Disable the user in IAM Identity Center and query the organizational event data store.

Answer: D

Explanation:

AWS IAM Identity Center centrally manages user access across an AWS Organization. Disabling the user in Identity Center immediately revokes access to all AWS accounts. According to AWS Certified Security - Specialty documentation, organizational CloudTrail event data stores provide centralized, queryable access to all events across accounts.

Using CloudTrail Lake enables direct querying of activity without exporting logs. Disabling the user at the Identity Center level ensures full containment.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

AWS IAM Identity Center Incident Response

AWS CloudTrail Lake

Question: 50

A company detects bot activity targeting Amazon Cognito user pool endpoints. The solution must block malicious requests while maintaining access for legitimate users.

Which solution meets these requirements?

- A. Enable Amazon Cognito threat protection.
- B. Restrict access to authenticated users only.
- C. Associate AWS WAF with the Cognito user pool.
- D. Monitor requests with CloudWatch.

Answer: A

Explanation:

Amazon Cognito threat protection is purpose-built to detect and mitigate malicious authentication activity such as credential stuffing and bot traffic. It uses adaptive risk-based analysis without disrupting legitimate users.

AWS WAF cannot be directly associated with Cognito user pools.

Referenced AWS Specialty Documents:

Amazon Cognito Threat Protection

Question: 51

CloudFormation stack deployments fail for some users due to permission inconsistencies.

Which combination of steps will ensure consistent deployments MOST securely? (Select THREE.)

- A. Create a composite principal service role.
- B. Create a service role with cloudformation.amazonaws.com as the principal.
- C. Attach scoped policies to the service role.
- D. Attach service ARNs in policy resources.
- E. Update each stack to use the service role.
- F. Allow iam:PassRole to the service role.

Answer: B, E, F

Explanation:

AWS best practices require CloudFormation to assume a dedicated service role. This ensures consistent permissions regardless of the user. Users must have iam:PassRole permission to pass the role. Updating stacks to use the service role enforces uniform deployment behavior.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

AWS CloudFormation Service Roles

Question: 52

A company needs centralized log monitoring with automatic detection across hundreds of AWS accounts.

Which solution meets these requirements with the LEAST operational effort?

- A. Designate a GuardDuty administrator account and enable protections.
- B. Centralize CloudWatch logs and use Inspector.
- C. Centralize CloudTrail logs and query with Athena.
- D. Stream logs to Kinesis and process with Lambda.

Answer: A**Explanation:**

Amazon GuardDuty provides fully managed threat detection across accounts when configured with delegated administration. EKS and RDS protections enable workload-aware detection with minimal setup.

Other solutions require custom pipelines and higher operational overhead.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

Amazon GuardDuty Multi-Account Architecture

Question: 53

A company has decided to move its fleet of Linux-based web server instances to an Amazon EC2 Auto Scaling group. Currently, the instances are static and are launched manually. When an administrator needs to view log files, the administrator uses SSH to establish a connection to the instances and retrieves the logs manually.

The company often needs to query the logs to produce results about application sessions and user issues. The company does not want its new automatically scaling architecture to result in the loss of any log files when

instances are scaled in.

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

- A. Configure a cron job on the instances to forward the log files to Amazon S3 periodically.
- B. Configure AWS Glue and Amazon Athena to query the log files.
- C. Configure the Amazon CloudWatch agent on the instances to forward the logs to Amazon CloudWatch Logs.
- D. Configure Amazon CloudWatch Logs Insights to query the log files.
- E. Configure the instances to write the logs to an Amazon Elastic File System (Amazon EFS) volume.

Answer: C, D

Explanation:

Amazon CloudWatch Logs is designed to collect, store, and analyze log data from ephemeral compute resources such as EC2 instances in Auto Scaling groups. According to the AWS Certified Security - Specialty Study Guide, using the CloudWatch agent to stream logs off instances ensures log durability even when instances are terminated during scale-in events.

CloudWatch Logs Insights provides a fully managed, serverless query engine that enables ad hoc querying, filtering, and aggregation of log data without requiring additional infrastructure. This directly satisfies the requirement to query logs for application sessions and user troubleshooting.

Option A introduces operational risk because logs could be lost between cron executions. Option B requires additional services and data pipelines, increasing cost and complexity. Option E adds storage cost and management overhead and is not necessary for log analytics.

AWS best practices recommend CloudWatch Logs and Logs Insights as the most cost-effective and scalable

solution for centralized log retention and analysis in Auto Scaling environments.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

Amazon CloudWatch Logs and Logs Insights

AWS Logging Best Practices

Question: 54

A company's security engineer receives an abuse notification from AWS indicating that malware is being hosted from the company's AWS account. The security engineer discovers that an IAM user created a new Amazon S3 bucket without authorization.

Which combination of steps should the security engineer take to MINIMIZE the consequences of this compromise? (Select THREE.)

- A. Encrypt all AWS CloudTrail logs.
- B. Turn on Amazon GuardDuty.
- C. Change the password for all IAM users.
- D. Rotate or delete all AWS access keys.
- E. Take snapshots of all Amazon Elastic Block Store (Amazon EBS) volumes.
- F. Delete any resources that are unrecognized or unauthorized.

Answer: B, D, F

Explanation:

AWS incident response guidance emphasizes immediate containment, credential invalidation, and removal of malicious resources. According to the AWS Certified Security - Specialty documentation, compromised credentials must be rotated or deleted immediately to prevent further unauthorized actions. Rotating or deleting access keys directly mitigates ongoing abuse.

Deleting unrecognized or unauthorized resources, such as the malicious S3 bucket, removes the active threat and limits further damage. Enabling Amazon GuardDuty provides continuous monitoring and helps identify additional compromised resources or malicious behavior that may not yet be visible.

Changing passwords for all IAM users is disruptive and unnecessary if compromise scope is limited.

Encrypting CloudTrail logs does not reduce active impact. Taking EBS snapshots is primarily for forensic investigation, not immediate consequence minimization.

AWS best practices recommend GuardDuty activation, credential rotation, and removal of malicious resources as first-response actions.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

AWS Incident Response Best Practices

Amazon GuardDuty Threat Detection

Question: 55

A company hosts its public website on Amazon EC2 instances behind an Application Load Balancer (ALB). The website is experiencing a global DDoS attack from a specific IoT device brand that uses a unique user agent. A security engineer is creating an AWS WAF web ACL and will associate it with the ALB.

Which rule statement will mitigate the current attack and future attacks from these IoT devices **without** blocking legitimate customers?

A. Use an IP set match rule statement.

- B. Use a geographic match rule statement.
- C. Use a rate-based rule statement.
- D. Use a string match rule statement on the user agent.

Answer: D

Explanation:

AWS WAF string match rule statements allow inspection of HTTP headers, including the User-Agent header. According to AWS Certified Security - Specialty guidance, when malicious traffic can be uniquely identified by a consistent request attribute, such as a device-specific user agent, a string match rule provides precise mitigation with minimal false positives.

IP-based blocking is ineffective for globally distributed botnets. Geographic blocking risks denying access to legitimate users. Rate-based rules limit request volume but do not prevent low-and-slow attacks.

By matching the unique IoT device brand in the User-Agent header, the security engineer can block **only** malicious requests while preserving customer access.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

AWS WAF Rule Statements

AWS DDoS Mitigation Best Practices

Question: 56

A company's security team wants to receive near-real-time email notifications about AWS abuse reports related to DoS attacks. An Amazon SNS topic already exists and is subscribed to by the security team.

What should the security engineer do next?

- A. Poll Trusted Advisor for abuse notifications by using a Lambda function.
- B. Create an Amazon EventBridge rule that matches AWS Health events for `AWS_ABUSE_DOS_REPORT` and publishes to SNS.
- C. Poll the AWS Support API for abuse cases by using a Lambda function.
- D. Detect abuse reports by using CloudTrail logs and CloudWatch alarms.

Answer: B

Explanation:

AWS abuse notifications are delivered as AWS Health events. According to the AWS Certified Security - Specialty Study Guide, Amazon EventBridge integrates natively with AWS Health and can be used to detect specific event types such as `AWS_ABUSE_DOS_REPORT` in near real time.

By creating an EventBridge rule that filters for the abuse report event type and publishes directly to Amazon SNS, the solution remains fully managed, low latency, and cost effective.

Polling APIs introduces delay and complexity. CloudTrail does not log abuse notifications. EventBridge with AWS Health is the recommended mechanism for reacting to AWS service events.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

AWS Health and EventBridge Integration

AWS Abuse Notification Handling

Question: 57 A security engineer discovers that a company's user passwords have no required minimum length.

The company uses the following identity providers (IdPs):

- AWS Identity and Access Management (IAM) federated with on-premises Active Directory

- Amazon Cognito user pools that contain the user database for an AWS Cloud application

Which combination of actions should the security engineer take to implement a required minimum password length? (Select TWO.)

- A. Update the password length policy in the IAM configuration.
- B. Update the password length policy in the Amazon Cognito configuration.
- C. Update the password length policy in the on-premises Active Directory configuration.
- D. Create an SCP in AWS Organizations to enforce minimum password length.
- E. Create an IAM policy with a minimum password length condition.

Answer: B, C

Explanation:

Password policies are enforced at the identity provider where authentication occurs. According to the AWS Certified Security - Specialty Study Guide, when IAM is federated with an external identity provider such as on-premises Active Directory, IAM does not manage or enforce password policies. Instead, password requirements such as minimum length must be enforced directly in Active Directory Group Policy Objects.

Amazon Cognito user pools maintain their own user directory and authentication logic. Cognito provides configurable password policies, including minimum length, complexity, and expiration. To enforce a minimum password length for application users, the Cognito user pool password policy **must be updated**.

IAM password policies apply only to IAM users that authenticate directly with IAM and do not affect federated users or Cognito users. SCPs and IAM policies cannot enforce password length requirements.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

AWS IAM Federation and Password Policies

Amazon Cognito User Pool Security Settings

Question: 58

A company's web application runs on Amazon EC2 instances behind an Application Load Balancer (ALB) in an Auto Scaling group. An AWS WAF web ACL is associated with the ALB. Instance logs are lost after reboots. The operations team suspects malicious activity targeting a specific PHP file.

Which set of actions will identify the suspect attacker's IP address for future occurrences?

- A. Configure VPC Flow Logs and search for PHP file activity.
- B. Install the CloudWatch agent on the ALB and export application logs.
- C. Export ALB access logs to Amazon OpenSearch Service and search them.
- D. Configure the web ACL to send logs to Amazon Kinesis Data Firehose. Deliver logs to Amazon S3 and query them with Amazon Athena.

Answer: D

Explanation:

AWS WAF logs contain detailed request-level information, including source IP addresses, requested URIs, and rule matches. According to AWS Certified Security - Specialty guidance, enabling AWS WAF logging provides the most reliable and tamper-resistant method to investigate web-based attacks, especially when instance-level logs are unavailable.

By streaming WAF logs through Amazon Kinesis Data Firehose to Amazon S3, the company ensures durable, centralized log storage that is independent of EC2 lifecycle events. Amazon Athena can then query the logs efficiently to identify repeated requests to the new-user-creation.php endpoint and extract attacker IP addresses.

VPC Flow Logs do not capture HTTP-level details. ALB access logs alone may not capture blocked requests. WAF logs provide the best forensic visibility for future detection.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

AWS WAF Logging and Monitoring

Amazon Athena Log Analysis

Question: 59

A company has security requirements for Amazon Aurora MySQL databases regarding encryption, deletion protection, public access, and audit logging. The company needs continuous monitoring and real-time visibility into compliance status.

Which solution will meet these requirements?

- A. Use AWS Audit Manager with a custom framework.
- B. Enable AWS Config and use managed rules to monitor Aurora MySQL compliance.
- C. Use AWS Security Hub configuration policies.
- D. Use EventBridge and Lambda with custom metrics.

Answer: B

Explanation:

AWS Config is the AWS service designed to continuously evaluate resource configurations against defined rules. According to the AWS Certified Security - Specialty Study Guide, AWS Config managed rules exist specifically to check database encryption, public accessibility, deletion protection, and log exports for Amazon RDS and Aurora.

AWS Config provides a real-time compliance timeline and displays the compliance state of each resource against each rule at any point in time. This granular visibility is required to assess ongoing compliance with security policies.

Audit Manager generates reports but does not provide continuous compliance monitoring. Security Hub aggregates findings but does not track configuration drift. EventBridge and Lambda introduce unnecessary complexity.

Referenced AWS Specialty Documents:

[AWS Certified Security - Specialty Official Study Guide](#)

[AWS Config Managed Rules for RDS](#)

[AWS Continuous Compliance Monitoring](#)

Question: 60

A company runs a global ecommerce website using Amazon CloudFront. The company must block traffic from specific countries to comply with data regulations.

Which solution will meet these requirements MOST cost-effectively?

- A. Use AWS WAF IP match rules.
- B. Use AWS WAF geo match rules.
- C. Use CloudFront geo restriction to deny the countries.
- D. Use geolocation headers in CloudFront.

Answer: C

Explanation:

Amazon CloudFront includes a built-in geo restriction feature that allows content to be allowed or denied based on the viewer's country. According to AWS Certified Security - Specialty documentation, CloudFront geo restriction is the most cost-effective method for country-based blocking because it does not require AWS WAF or additional rule processing.

AWS WAF geo match rules incur additional cost and are more appropriate when advanced inspection or layered security controls are required. IP-based blocking is impractical due to frequent IP changes. Geolocation headers do not enforce access control.

CloudFront geo restriction is evaluated at the edge and efficiently blocks disallowed countries with minimal latency and cost.

Referenced AWS Specialty Documents:

[AWS Certified Security - Specialty Official Study Guide](#)

[Amazon CloudFront Geo Restriction](#)

[AWS Edge Security Best Practices](#)

Question: 61

An AWS Lambda function was misused to alter data, and a security engineer must identify who invoked the function and what output was produced. The engineer cannot find any logs created by the Lambda function in Amazon CloudWatch Logs.

Which of the following explains why the logs are not available?

- A. The execution role for the Lambda function did not grant permissions to write log data to CloudWatch Logs.
- B. The Lambda function was invoked by using Amazon API Gateway, so the logs are not stored in CloudWatch Logs.
- C. The execution role for the Lambda function did not grant permissions to write to the Amazon S3 bucket where CloudWatch Logs stores the logs.
- D. The version of the Lambda function that was invoked was not current.

Answer: A

Explanation:

AWS Lambda automatically sends function execution logs to Amazon CloudWatch Logs when logging is enabled in the function code. However, this logging capability depends on the Lambda execution role having the appropriate permissions. According to the AWS Certified Security - Specialty Study Guide, the execution role must include permissions such as `logs:CreateLogGroup`, `logs:CreateLogStream`, and `logs:PutLogEvents`.

If these permissions are missing, Lambda cannot create log groups or streams, and no execution logs will appear in CloudWatch Logs—even though the function was successfully invoked. This is the most common reason Lambda logs are unavailable during forensic investigations.

Option B is incorrect because Lambda logs are stored in CloudWatch Logs regardless of whether the invocation source is API Gateway, EventBridge, or another AWS service. Option C is incorrect because CloudWatch Logs does not require direct S3 permissions from the Lambda execution role. Option D is irrelevant because Lambda versions do not affect logging behavior.

AWS documentation emphasizes verifying execution role permissions as a first step when Lambda logs are missing.

Referenced AWS Specialty Documents:

AWS Lambda Execution Roles

Amazon CloudWatch Logs Integration with Lambda

Question: 62

A company has an encrypted Amazon Aurora DB cluster in the us-east-1 Region that uses an AWS KMS customer managed key. The company must copy a DB snapshot to the us-west-1 Region but cannot access the encryption key across Regions.

What should the company do to properly encrypt the snapshot in us-west-1?

- A. Store the customer managed key in AWS Secrets Manager in us-west-1.
- B. Create a new customer managed key in us-west-1 and use it to encrypt the snapshot.
- C. Create an IAM policy to allow access to the key in us-east-1 from us-west-1.
- D. Create an IAM policy that allows RDS in us-west-1 to access the key in us-east-1.

Answer: B

Explanation:

AWS KMS keys are strictly regional resources. According to AWS Certified Security - Specialty documentation, a KMS key created in one Region cannot be used to encrypt or decrypt data in another Region. This includes encrypted RDS and Aurora snapshots.

When copying an encrypted snapshot to a different Region, the destination Region must have its own KMS key. AWS automatically re-encrypts the snapshot using the specified KMS key in the destination Region during the copy operation.

Options C and D are invalid because IAM policies cannot extend a KMS key's scope across Regions.

Option A is incorrect because Secrets Manager does not store or manage KMS keys themselves.

AWS best practices require creating a new customer managed key in the target Region and using it during the

snapshot copy process.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

AWS KMS Regional Key Limitations

Amazon RDS Encrypted Snapshot Copy

Question: 63

A security engineer needs to prepare Amazon EC2 instances for quarantine during a security incident. AWS Systems Manager Agent (SSM Agent) is installed, and a script exists to install and update forensic tools.

Which solution will quarantine EC2 instances during a security incident?

- A. Track SSM Agent versions with AWS Config.
- B. Configure Session Manager to deny external connections.
- C. Store the script in Amazon S3 and grant read access.
- D. Configure IAM permissions for the SSM Agent to run the script as a Systems Manager Run Command document.

Answer: D

Explanation:

AWS Systems Manager Run Command enables secure, remote execution of commands on EC2 instances without requiring network access or inbound ports. According to the AWS Certified Security - Specialty Study Guide, Run Command is a recommended mechanism for incident response actions such as installing forensic tools, collecting evidence, or applying quarantine controls.

By granting the SSM Agent permission to execute a predefined Run Command document, the security engineer can immediately run the quarantine script across affected instances. This approach supports automation,

scalability, and auditability, all of which are critical during security incidents.

Options A, B, and C do not directly enforce quarantine or execute response actions. Tracking versions and storing scripts alone do not trigger incident response.

AWS documentation highlights Systems Manager Run Command as a core capability for automated containment and investigation.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

AWS Systems Manager Run Command

AWS Incident Response Automation

Question: 64

A company needs to identify the root cause of security findings and investigate IAM roles involved in those findings. The company has enabled VPC Flow Logs, Amazon GuardDuty, and AWS CloudTrail.

Which solution will meet these requirements?

- A. Use Amazon Detective to investigate IAM roles and visualize findings.
- B. Use Amazon Inspector and CloudWatch dashboards.
- C. Export GuardDuty findings to S3 and analyze with Athena.
- D. Use Security Hub custom actions to investigate IAM roles.

Answer: A

Explanation:

Amazon Detective is specifically designed to help security teams investigate and visualize the root cause of security findings. According to AWS Certified Security - Specialty documentation, Detective automatically aggregates and correlates data from GuardDuty, CloudTrail, and VPC Flow Logs to provide interactive visualizations and timelines.

Detective enables investigators to pivot from GuardDuty findings to IAM roles, API calls, network traffic, and resource behavior. This makes it the most efficient tool for understanding how IAM roles were used during suspicious activity.

Amazon Inspector focuses on vulnerability assessment, not behavioral investigation. Security Hub aggregates findings but does not provide deep investigation graphs. Manual analysis with Athena requires significantly more effort.

AWS guidance explicitly recommends Amazon Detective for root cause analysis and visualization of security incidents.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

Amazon Detective Investigation Capabilities

AWS Threat Detection and Analysis

Question: 65

A company uploads data files as objects into an Amazon S3 bucket. A vendor downloads the objects to perform data processing.

A security engineer must implement a solution that prevents objects from residing in the S3 bucket for longer than 72 hours.

- A. Configure S3 Versioning to expire object versions that have been in the bucket for 72 hours.
- B. Configure an S3 Lifecycle configuration rule on the bucket to expire objects after 72 hours.
- C. Use the S3 Intelligent-Tiering storage class and configure expiration after 72 hours.
- D. Generate presigned URLs that expire after 72 hours.

Answer: B

Explanation:

Amazon S3 Lifecycle configuration rules are the native, automated mechanism for managing object retention and deletion. According to AWS Certified Security - Specialty documentation, lifecycle rules can be configured to expire objects based on the number of days since object creation. Once the expiration time is reached, Amazon S3 permanently deletes the objects without manual intervention.

This solution directly enforces a maximum retention period of 72 hours and ensures compliance regardless of whether the vendor downloads the data or not. Lifecycle rules are evaluated continuously by Amazon S3 and do not require scripts, cron jobs, or additional services, making them the most operationally efficient and cost-effective solution.

S3 Versioning controls versions but does not enforce object deletion timelines. S3 Intelligent-Tiering optimizes storage cost but does not delete objects. Presigned URLs only control access duration and do not remove objects from storage.

AWS explicitly recommends lifecycle policies for automated data retention enforcement.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

Amazon S3 Lifecycle Management

Question: 66

A company experienced a security incident caused by a vulnerable container image that was pushed from an external CI/CD pipeline into Amazon ECR.

Which solution will prevent vulnerable images from being pushed?

- A. Enable ECR enhanced scanning with Lambda blocking.
- B. Use Amazon Inspector with EventBridge and Lambda.
- C. Integrate Amazon Inspector into the CI/CD pipeline using SBOM generation and fail the pipeline on critical findings.
- D. Enable basic continuous ECR scanning.

Answer: C

Explanation:

Amazon Inspector provides native CI/CD integration capabilities that allow security checks to occur before container images are pushed to Amazon ECR. According to AWS Certified Security - Specialty documentation, Inspector does not block image pushes automatically. Instead, prevention must occur inside the CI/CD pipeline itself.

By generating a Software Bill of Materials (SBOM) using the Amazon Inspector SBOM generator and submitting it to Inspector for scanning, the pipeline can detect critical vulnerabilities before the image is uploaded. If vulnerabilities exceed policy thresholds, the pipeline fails, preventing deployment.

Post-push scanning solutions only detect vulnerabilities after exposure. Event-driven blocking does not prevent the initial risk.

AWS best practices require “shift-left” security controls to prevent vulnerable artifacts from entering production.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

Amazon Inspector CI/CD Integration

Question: 67

AWS Config cannot deliver configuration snapshots to Amazon S3.

Which TWO actions will remediate this issue?

- A. Verify the S3 bucket policy allows config.amazonaws.com.
- B. Verify the IAM role has s3:GetBucketAcl and s3:PutObject permissions.
- C. Verify the S3 bucket can assume the IAM role.
- D. Verify IAM policy allows AWS Config to write logs.
- E. Modify AWS Config API permissions.

Answer: A, B

Explanation:

AWS Config requires permissions at two levels to deliver configuration data: the AWS Config service role and the S3 bucket policy. The AWS Certified Security - Specialty Study Guide states that the S3 bucket policy must explicitly allow the config.amazonaws.com service principal to write objects. Additionally, the IAM role used by AWS Config must allow s3:GetBucketAcl and s3:PutObject.

If either permission is missing, AWS Config cannot deliver snapshots and will log delivery errors in CloudTrail.

This dual-permission model ensures least privilege while maintaining secure delivery of compliance data.

Other options reference incorrect principals or irrelevant permissions.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

AWS Config Prerequisites

Question: 68

A company must inventory sensitive data across all Amazon S3 buckets in all accounts from a single security account.

- A. Delegate Amazon Macie and Security Hub administration.
- B. Use Amazon Inspector with Security Hub.
- C. Use Inspector with Trusted Advisor.
- D. Use Macie with Trusted Advisor.

Answer: A

Explanation:

Amazon Macie is the AWS service designed to discover and classify sensitive data in S3. Delegated administration enables centralized visibility across an organization. Security Hub aggregates Macie findings for a single-pane-of-glass view.

Inspector does not scan S3 data. Trusted Advisor is not a sensitive data discovery tool.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

Amazon Macie Multi-Account Architecture

Question: 69

A company needs to scan all AWS Lambda functions for code vulnerabilities.

- A. Use Amazon Macie.
- B. Enable Amazon Inspector Lambda scanning.
- C. Use GuardDuty and Security Hub.
- D. Use GuardDuty Lambda Protection.

Answer: B

Explanation:

Amazon Inspector provides native Lambda code vulnerability scanning. GuardDuty focuses on runtime threats, not static code analysis.

Question: 70

Notify when IAM roles are modified.

- A. Use Amazon Detective.
- B. Use EventBridge with CloudTrail events.
- C. Use CloudWatch metric filters.
- D. Use CloudWatch subscription filters.

Answer: B

Explanation:

EventBridge natively consumes CloudTrail management events and provides near-real-time notifications.

Question: 71

A company has several Amazon S3 buckets that do not enforce encryption in transit. A security engineer must implement a solution that enforces encryption in transit for all the company's existing and future S3 buckets.

Which solution will meet these requirements?

- A. Enable AWS Config. Create a proactive AWS Config Custom Policy rule. Create a Guard clause to evaluate the S3 bucket policies to check for a value of True for the aws:SecureTransport condition key. If the AWS Config rule evaluates to NON_COMPLIANT, block resource creation.
- B. Enable AWS Config. Configure the s3-bucket-ssl-requests-only AWS Config managed rule and set the rule trigger type to Hybrid. Create an AWS Systems Manager Automation runbook that applies a bucket policy to deny requests when the value of the aws:SecureTransport condition key is False. Configure automatic remediation. Set the runbook as the target of the rule.
- C. Enable Amazon Inspector. Create a custom AWS Lambda rule. Create a Lambda function that applies a bucket policy to deny requests when the value of the aws:SecureTransport condition key is False. Set the Lambda function as the target of the rule.

D. Create an AWS CloudTrail trail. Enable S3 data events on the trail. Create an AWS Lambda function that applies a bucket policy to deny requests when the value of the aws:SecureTransport condition key is False. Configure the CloudTrail trail to invoke the Lambda function.

Answer: B

Explanation:

To enforce encryption in transit for Amazon S3, AWS best practice is to require HTTPS (TLS) by using a bucket policy condition that denies any request where aws:SecureTransport is false. The requirement includes both existing buckets and future buckets, so the control must continuously evaluate configuration drift and automatically remediate. AWS Config is the service intended for continuous configuration compliance monitoring across resources, and AWS Config managed rules provide standardized checks with low operational overhead. The s3-bucket-ssl-requests-only managed rule evaluates whether S3 buckets enforce SSL-only requests, aligning directly with enforcing encryption in transit. Setting the trigger type to Hybrid ensures evaluation both on configuration changes and periodically. Automatic remediation with an AWS Systems Manager Automation runbook allows the organization to apply or correct the bucket policy consistently at scale without manual work. This approach also supports governance by maintaining a measurable compliance status while actively fixing noncompliance. Option A is not the best fit because a “proactive” custom policy rule does not by itself remediate existing buckets and “block resource creation” is not how AWS Config enforces controls. Option C is incorrect because Amazon Inspector is a vulnerability management service and does not govern S3 bucket transport policies. Option D is inefficient and indirect because CloudTrail data events are not a compliance engine and would require custom processing.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

AWS Config Managed Rules for S3 Compliance

Amazon S3 Security Best Practices for SSL-only Access

Question: 72

A company is using AWS Organizations with nested OUs to manage AWS accounts. The company has a custom compliance monitoring service for the accounts. The monitoring service runs as an AWS Lambda function

and is invoked by Amazon EventBridge Scheduler.

The company needs to deploy the monitoring service in all existing and future accounts in the organization. The company must avoid using the organization's management account when the management account is not required.

Which solution will meet these requirements?

- A. Create a CloudFormation stack set in the organization's management account and manually add new accounts.
- B. Configure a delegated administrator account for AWS CloudFormation. Create a CloudFormation StackSet in the delegated administrator account targeting the organization root with automatic deployment enabled.
- C. Use Systems Manager delegated administration and Automation to deploy the Lambda function and schedule.
- D. Create a Systems Manager Automation runbook in the management account and share it to accounts.

Answer: B

Explanation:

AWS Organizations and CloudFormation StackSets provide an organizational deployment mechanism for consistent infrastructure across accounts. AWS Certified Security - Specialty guidance emphasizes minimizing use of the management account and using delegated administrator capabilities where available for centralized governance while reducing blast radius. By configuring a delegated administrator account for AWS CloudFormation, the company can create and manage StackSets without performing day-to-day deployment operations from the management account. Targeting the organization root ensures the StackSet deploys to all existing accounts. Enabling automatic deployment ensures that any future accounts that join the organization (or move into targeted OUs, depending on configuration) automatically receive the monitoring service without manual intervention. This directly meets the requirement to deploy to all existing and future accounts with minimal effort. Option A requires ongoing manual updates when accounts are added, increasing operational overhead. Options C and D rely on Systems Manager Automation, which can work but introduces additional operational complexity and is not the standard AWS mechanism for organization-wide infrastructure rollout compared to StackSets with auto-deployment. StackSets also provide consistent change control, drift detection, and centralized update mechanisms, which align with governance expectations for compliance tooling.

Referenced AWS Specialty Documents:

AWS Organizations Delegated Administration

AWS CloudFormation StackSets for Multi-Account Governance

Question: 73

A company is building a secure solution that relies on an AWS Key Management Service (AWS KMS) customer managed key. The company wants to allow AWS Lambda to use the KMS key. However, the company wants to prevent Amazon EC2 from using the key.

Which solution will meet these requirements?

- A. Use IAM explicit deny for EC2 instance profiles and allow for Lambda roles.
- B. Use a KMS key policy with kms:ViaService conditions to allow Lambda usage and deny EC2 usage.
- C. Use aws:SourceIp and aws:AuthorizedService condition keys in the KMS key policy.
- D. Use an SCP to deny EC2 and allow Lambda.

Answer: B

Explanation:

AWS KMS access control is primarily enforced through key policies (and optionally grants), and AWS recommends using key policy condition keys to restrict how keys can be used. The kms:ViaService condition key is specifically designed to restrict KMS API usage to requests that come through a particular AWS service endpoint in a specific Region. This is the most robust way to ensure a key can be used only via AWS Lambda (for example, lambda.<region>.amazonaws.com) and not via Amazon EC2 (ec2.<region>.amazonaws.com), even if IAM permissions exist elsewhere. By writing a key policy that uses the Lambda execution role as the principal and conditions on kms:ViaService, the company can tightly bind key usage to Lambda-originated cryptographic operations while preventing use through EC2 service paths. Option A is weaker because EC2 is not the only way an IAM principal might use KMS, and relying on attaching explicit deny policies broadly is harder to manage and can miss principals. Option C is incorrect because aws:AuthorizedService is not the typical mechanism for KMS service restriction, and SourceIp is unreliable for service-to-service calls. Option D is not ideal because SCPs do not provide fine-grained service-path restrictions for KMS usage and cannot “allow” beyond IAM; key

policy controls still apply.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

AWS KMS Key Policies and Condition Keys

AWS KMS Best Practices for Service-Scoped Key Usage

Question: 74

A company has a web application that reads from and writes to an Amazon S3 bucket. The company needs to use AWS credentials to authenticate all S3 API calls to the S3 bucket.

Which solution will provide the application with AWS credentials to make S3 API calls?

- A. Integrate with Cognito identity pools and use GetId to obtain AWS credentials.
- B. Integrate with Cognito identity pools and use AssumeRoleWithWebIdentity to obtain AWS credentials.
- C. Integrate with Cognito user pools and use the ID token to obtain AWS credentials.
- D. Integrate with Cognito user pools and use the access token to obtain AWS credentials.

Answer: B

Explanation:

Amazon Cognito identity pools are designed to provide temporary AWS credentials for applications by exchanging an authenticated identity token for AWS Security Token Service (STS) credentials. AWS Certified Security - Specialty guidance distinguishes between Cognito user pools (authentication) and identity pools (authorization to AWS resources). A user pool can authenticate a user and issue tokens, but an identity pool is required to obtain AWS credentials that can be used to sign AWS API requests, such as S3 API calls. The correct mechanism is for the application to use AssumeRoleWithWebIdentity through STS (which is the underlying federation method used by identity pools) to receive temporary credentials for an IAM role that grants S3 permissions. GetId alone does not provide credentials; it returns an identity identifier that is used as part of the credential exchange flow. Options C and D are incorrect because user pool tokens are not AWS credentials and cannot directly sign S3 requests. The solution therefore must use identity pools to map users to IAM roles and retrieve temporary credentials, satisfying the requirement for authenticated API calls using

short-lived credentials.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

Amazon Cognito Identity Pools and STS Federation

AWS STS AssumeRoleWithWebIdentity

Question: 75

A company is running a new workload across accounts in an organization in AWS Organizations. All running resources must have a tag of CostCenter, and the tag must have one of three approved values. The company must enforce this policy and must prevent any changes of the CostCenter tag to a non-approved value.

Which solution will meet these requirements?

- A. Use AWS Config custom policy rule and an SCP to deny non-approved aws:RequestTag/CostCenter values.
- B. Use CloudTrail + EventBridge + Lambda to block creation.
- C. Enable tag policies, define allowed values, enforce noncompliant operations, and use an SCP to deny creation when aws:RequestTag/CostCenter is null.
- D. Enable tag policies and use EventBridge + Lambda to block changes.

Answer: C

Explanation:

AWS Organizations tag policies are designed to standardize and govern tag keys and allowed values across accounts. AWS Certified Security - Specialty documentation describes tag policies as a governance mechanism that helps enforce consistent tagging by specifying required tag keys and permitted values. To ensure every resource has the CostCenter tag at creation time, an SCP can deny create actions when aws:RequestTag/CostCenter is missing (null). This prevents resources from being created without the required

tag. Tag policies then define the three approved values and can be configured to enforce or report noncompliance depending on supported services, ensuring that tag values remain within the allowed set and preventing drift to unapproved values. Compared with custom Lambda-based enforcement, this approach minimizes operational overhead and keeps enforcement within AWS native governance services. Option A partially addresses allowed values at request time but does not address ongoing governance as cleanly across many services. Option B is not preventive because Lambda runs after events and cannot reliably block all creations. Option D still relies on custom logic and is not as operationally efficient as tag policies plus SCP guardrails.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

AWS Organizations Tag Policies

AWS Organizations SCP Condition Keys for Tag Enforcement

Question: 76

A company has AWS accounts in an organization in AWS Organizations. An Amazon S3 bucket in one account is publicly accessible. A security engineer must remove public access and ensure the bucket cannot be made public again.

Which solution will meet these requirements?

- A. Enforce KMS encryption and deny s3:GetObject by SCP.
- B. Enable PublicAccessBlock and deny s3:GetObject by SCP.
- C. Enable PublicAccessBlock and deny s3:PutPublicAccessBlock by SCP.
- D. Enable Object Lock governance and deny s3:PutPublicAccessBlock by SCP.

Answer: C

Explanation:

Amazon S3 Block Public Access provides centralized controls to prevent public access through bucket policies and ACLs. AWS Certified Security - Specialty guidance recommends enabling Block Public Access to reduce accidental exposure and to enforce guardrails that override public grants. Enabling Block Public Access on the bucket removes current public exposure when combined with correcting policies/ACLs and prevents future misconfiguration. To ensure the bucket cannot be made public again, the security engineer must prevent principals from disabling Block Public Access. An SCP that denies s3:PutPublicAccessBlock prevents changes that would remove or weaken the PublicAccessBlock configuration, enforcing the guardrail across the OU or account. Options A and D do not directly address public exposure control. Option B denies object reads but does not ensure public access cannot be re-enabled; it also does not address the root misconfiguration pathways and could disrupt legitimate access patterns. Option C specifically combines the correct preventive control (PublicAccessBlock) with organizational enforcement to stop future reversal.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

Amazon S3 Block Public Access

AWS Organizations SCP Guardrails for S3 Controls

Question: 77

A company uses AWS Organizations and has an SCP at the root that prevents sharing resources with external accounts. The company now needs to allow only the marketing account to share resources externally while preventing all other accounts from doing so. All accounts are in the same OU.

Which solution will meet these requirements?

- A. Create a new SCP in the marketing account to explicitly allow sharing.
- B. Edit the existing SCP to add a condition that excludes the marketing account.
- C. Edit the SCP to include an Allow statement for the marketing account.
- D. Use a permissions boundary in the marketing account.

Answer: B

Explanation:

Service control policies (SCPs) define the maximum available permissions for accounts and are evaluated as guardrails. AWS Certified Security - Specialty documentation states SCPs are typically used to apply organization-wide restrictions, and exceptions are commonly handled by using conditions (for example, excluding specific accounts) or by structuring OUs differently. Because all accounts are in the same OU and the company must continue blocking external sharing for everyone except one account, modifying the existing SCP to exclude the marketing account is the most direct solution. An SCP attached at the root affects all accounts unless conditions narrow its scope. Adding a condition that excludes the marketing account allows that account to retain the ability to share resources externally while the SCP continues to block sharing for other accounts. Option A is not feasible because account-level SCPs cannot override a deny applied by a parent SCP; explicit denies always win. Option C misunderstands SCP behavior because SCPs do not grant permissions; they only limit. Option D is an IAM control that cannot override an organization-level deny. Therefore, the only secure, scalable option is to modify the existing SCP with an exception condition for the marketing account.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

AWS Organizations SCP Evaluation Logic

SCP Deny Precedence and Exception Patterns

Question: 78

A security engineer configured VPC Flow Logs to publish to Amazon CloudWatch Logs. After 10 minutes, no logs appear. The issue is isolated to the IAM role associated with VPC Flow Logs.

What could be the reason?

- A. logs:GetLogEvents is missing.
- B. The engineer cannot assume the role.
- C. The vpc-flow-logs.amazonaws.com principal cannot assume the role.
- D. The role cannot tag the log stream.

Answer: C

Explanation:

VPC Flow Logs require an IAM role that CloudWatch Logs can use to publish flow log records. AWS documentation and AWS Certified Security - Specialty materials explain that the VPC Flow Logs service must be able to assume the IAM role through its trust policy. The trust relationship must include the service principal `vpc-flow-logs.amazonaws.com`. If the trust policy does not allow this principal to assume the role, flow logs cannot be delivered and no records will appear in the CloudWatch Logs log group even when traffic exists. `logs:GetLogEvents` is not required for delivery; it is used for reading logs. The security engineer's ability to assume the role is not relevant because the service, not the engineer, assumes it. Tagging permissions are not required for basic log delivery. Therefore, the most likely cause is an incorrect trust policy that prevents the VPC Flow Logs service principal from assuming the role.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

Amazon VPC Flow Logs IAM Role Requirements

IAM Trust Policies for AWS Services

Question: 79

A company requires a specific software application to be installed on all new and existing Amazon EC2 instances across an AWS Organization. SSM Agent is installed and active.

How can the company continuously monitor deployment status of the software application?

- A. Use AWS Config organization-wide with the `ec2-managedinstance-applications-required` managed rule and specify the application name.
- B. Use approved AMIs rule organization-wide.
- C. Use Distributor package and review output.
- D. Use Systems Manager Application Manager inventory filtering.

Answer: A

Explanation:

Continuous monitoring requires an always-on compliance service that evaluates resources over time. AWS Config provides managed rules that assess configuration state and compliance continuously. AWS Certified Security - Specialty guidance highlights AWS Config for continuous compliance across accounts and regions when used with AWS Organizations. The ec2-managedinstance-applications- required managed rule evaluates whether specified software is installed on managed instances, leveraging Systems Manager inventory/managed instance status. By enabling AWS Config organization-wide and deploying this managed rule across all accounts, the company can continuously evaluate both existing and newly launched instances for required application presence. This provides a consistent compliance dashboard and history of compliance changes. Option D can provide inventory lists, but it is not a compliance rule engine that flags noncompliance with the same governance reporting and remediation pathways. Options B and C are operational approaches but do not provide continuous compliance state across the organization.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

AWS Config Managed Rules for EC2 and SSM Managed Instances

AWS Organizations Integration with AWS Config

Question: 80

A company sends Apache logs from EC2 Auto Scaling instances to a CloudWatch Logs log group with 1-year retention. A suspicious IP address appears in logs. A security engineer needs to analyze the past week of logs to count requests from that IP and list requested URLs.

What should the engineer do with the LEAST effort?

- A. Export to S3 and use Macie.
- B. Stream to OpenSearch and analyze.
- C. Use CloudWatch Logs Insights with queries.

D. Export to S3 and use AWS Glue.

Answer: C

Explanation:

CloudWatch Logs Insights is a managed, on-demand query capability designed to search and analyze log data stored in CloudWatch Logs without moving the data elsewhere. AWS Certified Security - Specialty documentation highlights Logs Insights as the lowest-effort method for rapid investigations, because it supports filtering, parsing, aggregation, and time-range queries directly over existing log groups. In this scenario, the logs already exist in CloudWatch Logs with sufficient retention. The engineer can write a query that filters for the suspicious IP address, counts occurrences over the last 7 days, and extracts requested URLs using parsing functions. This satisfies both requirements (count and URLs) immediately, without building pipelines or exporting data. Option B adds operational overhead by provisioning and maintaining OpenSearch ingestion and indexing. Options A and D require exporting data and additional services that are not necessary for a one-week forensic query. Therefore, Logs Insights is the most efficient and cost-effective approach.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

Amazon CloudWatch Logs Insights Querying and Investigation Workflows

Question: 81

A company needs to detect unauthenticated access to its Amazon Elastic Kubernetes Service (Amazon EKS) clusters. The solution must require no additional configuration of the existing EKS deployment.

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

- A. Install a third-party security add-on.
- B. Enable AWS Security Hub and monitor Kubernetes findings.
- C. Monitor CloudWatch Container Insights metrics for EKS.
- D. Enable Amazon GuardDuty and use EKS Audit Log Monitoring.

Answer: D

Explanation:

Amazon GuardDuty provides managed threat detection and supports EKS protection features that analyze Kubernetes audit logs to detect suspicious activity, including unauthorized or unauthenticated access attempts. AWS Certified Security - Specialty documentation recommends GuardDuty for low-overhead detection because it is fully managed and does not require deploying agents or modifying application code. EKS Audit Log Monitoring is designed to consume and analyze relevant control plane audit events to identify anomalous or unauthorized actions against the cluster. Compared to third-party add-ons, GuardDuty reduces operational burden and remains fully within AWS managed services. Security Hub aggregates findings from services like GuardDuty but does not itself perform the detection. CloudWatch Container Insights focuses on performance and operational metrics, not authentication security detections. Therefore, enabling GuardDuty with EKS Audit Log Monitoring provides the required detection with the least operational effort and without requiring additional configuration to the existing EKS workload beyond enabling the feature.

Referenced AWS Specialty Documents:

AWS Certified Security - Specialty Official Study Guide

Amazon GuardDuty EKS Protection and Audit Log Monitoring

AWS Threat Detection Best Practices for Kubernetes on AWS