Amazon.SCS-C03.v2026-05-07.q75

Exam Code:	SCS-C03
Exam Name:	AWS Certified Security - Specialty
Certification Provider:	Amazon
Free Question Number:	75
Version:	v2026-05-07
# of views:	133
# of Questions views:	750
https://www.freepdfdumps.com/Amazon.SCS-C03.v2026-05-07.q75.html	

NEW QUESTION: 1
A security engineer needs to implement a solution to create and control the keys that a company uses for cryptographic operations. The security engineer must create symmetric keys in which the key material is generated and used within a custom key store that is backed by an AWS CloudHSM cluster. The security engineer will use symmetric and asymmetric data key pairs for local use within applications. The security engineer also must audit the use of the keys.
How can the security engineer meet these requirements?
A. To create the keys, use AWS Key Management Service (AWS KMS) and the custom key stores with the CloudHSM cluster. For auditing, use Amazon Athena.
B. To create the keys, use Amazon S3 and the custom key stores with the CloudHSM cluster. For auditing, use AWS CloudTrail.
C. To create the keys, use AWS Key Management Service (AWS KMS) and the custom key stores with the CloudHSM cluster. For auditing, use Amazon GuardDuty.
D. To create the keys, use AWS Key Management Service (AWS KMS) and the custom key stores with the CloudHSM cluster. For auditing, use AWS CloudTrail.
Answer: D (LEAVE A REPLY)
The requirement is to havekey material generated and used inside a custom key store backed by an AWS CloudHSM cluster. This is exactly whatAWS KMS Custom Key Storesprovide: KMS manages the keys and policies, but the cryptographic operations for those KMS keys occur in the associatedCloudHSMcluster, keeping the key material within HSM boundaries. For applications that needlocal-use data keys(both symmetric data keys and asymmetric data key pairs), KMS supports generating data keys and data key pairs that applications can use for envelope encryption and local cryptographic operations, while the master key protections remain within KMS (and within CloudHSM when using a custom key store).
For auditing, AWS best practice isAWS CloudTrail, which records KMS API calls (such as CreateKey, GenerateDataKey, GenerateDataKeyPair, Encrypt/Decrypt, etc.) and provides an immutable event history for compliance and investigation. Athena can query logs, but it is not the

primary audit record source; GuardDuty is for threat detection, not authoritative key-usage auditing. Therefore, the correct combination isKMS with a CloudHSM-backed custom key storeplusCloudTrailfor auditability.

NEW QUESTION: 2
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 (LEAVE A REPLY)
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

NEW QUESTION: 3

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: (SHOW ANSWER)
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

NEW QUESTION: 4
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 (LEAVE A REPLY)
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

NEW QUESTION: 5
A company runs a web application on a fleet of Amazon EC2 instances that are in an Auto Scaling group. The EC2 instances are in the same VPC subnet as other workloads.
A security engineer deploys Amazon GuardDuty and integrates it with AWS Security Hub. The security engineer needs to implement anautomated solutionto detect and respond to anomalous traffic patterns. The solution must follow AWS best practices forinitial incident responseand mustminimize disruptionto the web application.
Which solution will meet these requirements?
A. Disable the instance profile access keys by using AWS Lambda.
B. Remove the affected instance from the Auto Scaling group and isolate it with a restricted security group by using AWS Lambda.
C. Update the network ACL to block the detected traffic source.
D. Send GuardDuty findings to Amazon SNS for email notification.
Answer: B (LEAVE A REPLY)
AWS incident response best practices emphasizecontainment with minimal blast radiuswhile preserving business continuity. According to the AWS Certified Security - Specialty Official Study Guide, isolating a compromised resource while allowing the application to continue operating is the recommended initial response.

By creating an Amazon EventBridge rule that reacts to GuardDuty anomalous traffic findings and invokes an AWS Lambda function, the security engineer can automaticallyremove the affected EC2 instance from the Auto Scaling groupand attach arestricted security group. This immediately stops malicious activity while allowing Auto Scaling to replace the instance and keep the application available.
Option A is inappropriate because EC2 instance profiles do not use long-term access keys.
Option C applies subnet-wide changes that could disrupt unrelated workloads. Option D provides notification only and does not meet the automated response requirement.
AWS documentation explicitly identifiesinstance isolation via security groupsas a preferred containment technique that preserves application availability and forensic integrity.
* AWS Certified Security - Specialty Official Study Guide
* Amazon GuardDuty User Guide
* AWS Incident Response Best Practices

NEW QUESTION: 6
A company uses AWS Organizations. The company has teams that use an AWS CloudHSM hardware security module (HSM) that is hosted in a central AWS account. One of the teams creates its own new dedicated AWS account and wants to use the HSM that is hosted in the central account.
How should a security engineer share the HSM that is hosted in the central account with the new dedicated account?
A. Use AWS Resource Access Manager (AWS RAM) to share the VPC subnet ID of the HSM that is hosted in the central account with the new dedicated account. Configure the CloudHSM security group to accept inbound traffic from the private IP addresses of client instances in the new dedicated account.
B. Use AWS Identity and Access Management (IAM) to create a cross-account role to access the CloudHSM cluster that is in the central account. Create a new IAM user in the new dedicated account.
Assign the cross-account role to the new IAM user.
C. Use AWS IAM Identity Center to create an AWS Security Token Service (AWS STS) token to authenticate from the new dedicated account to the central account. Use the cross-account permissions that are assigned to the STS token to invoke an operation on the HSM in the central account.
D. Use AWS Resource Access Manager (AWS RAM) to share the ID of the HSM that is hosted in the central account with the new dedicated account. Configure the CloudHSM security group to accept inbound traffic from the private IP addresses of client instances in the new dedicated account.
Answer: D (LEAVE A REPLY)
AWS CloudHSM is aVPC-scopedservice: the HSMs (and the CloudHSM cluster) live inside a VPC in the central account, and clients connect over the network to perform cryptographic operations. When another account needs to use a centrally managed CloudHSM cluster, the right

approach is toshare the CloudHSM clusterwith that account and allow network connectivity from the consuming account's clients. AWS provides cross-account resource sharing throughAWS Resource Access Manager (AWS RAM)for supported resources, including CloudHSM clusters. Sharing thecluster/HSM identifieris what grants the consuming account visibility/ability to create client configurations against that shared cluster.
After sharing, the consuming account's EC2 instances (CloudHSM clients) still must be able to reach the HSM ENIs over the network, so the CloudHSM security group in the central account must allow inbound connections from the client sources (typically by security group referencing via VPC connectivity, or by allowing the relevant IP range/ports as appropriate).
Option A is incorrect because sharing asubnet IDdoes not share the CloudHSM resource itself.
Options B and C misuse IAM/STS: CloudHSM cryptographic operations are not granted via assuming an IAM role into the central account; access is based oncluster sharing + network connectivity + CloudHSM user authenticationat the HSM level.

NEW QUESTION: 7
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 (LEAVE A REPLY)
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:
AWS Certified Security - Specialty Official Study Guide
Amazon S3 Lifecycle Management

NEW QUESTION: 8
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 (LEAVE A REPLY)
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

NEW QUESTION: 9
A company has a large fleet of Amazon Linux 2 Amazon EC2 instances that run an application processing sensitive data. 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 (LEAVE A REPLY)
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

NEW QUESTION: 10
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 (LEAVE A REPLY)
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

NEW QUESTION: 11
A company runs an application on a fleet of Amazon EC2 instances. The application is accessible to users around the world. The company associates an AWS WAF web ACL with an Application Load Balancer (ALB) that routes traffic to the EC2 instances.
A security engineer is investigating a sudden increase in traffic to the application. The security engineer discovers a significant amount of potentially malicious requests coming from hundreds of IP addresses in two countries. The security engineer wants to quickly limit the potentially malicious requests. The security engineer does not want to prevent legitimate users from accessing the application.
Which solution will meet these requirements?
A. Use AWS WAF to implement a rate-based rule for all incoming requests.
B. Use AWS WAF to implement a geographical match rule to block all incoming traffic from the two countries.
C. Edit the ALB security group to include a geographical match rule to block all incoming traffic from the two countries.
D. Add deny rules to the ALB security group that prohibit incoming requests from the IP addresses.
Answer: A (LEAVE A REPLY)
A rate-based rule in AWS WAF is designed to quickly mitigate spikes and potential layer 7 floods bytracking request rates per originating IPand temporarily blocking (or counting/challenging, depending on configuration) IPs that exceed a defined threshold within a 5-minute rolling window. In this scenario, the malicious traffic is distributed acrosshundreds of IPsin two countries, and the application still needs to remain available globally for legitimate users. A rate-based rule provides fast, targeted throttling that reduces abusive request patterns without permanently blocking entire geographies. This aligns with "quickly limit" while minimizing collateral impact.
Blocking both countries with a geo match rule (Option B) would likely block legitimate users located in those countries, which violates the requirement. Security groups (Options C and D) cannot natively enforcegeographicfiltering, and they are not well suited for large, rapidly changing sets of public source IPs at the application layer. Additionally, WAF operates at layer 7 with richer matching (rate limiting, URI/header patterns, bot controls), which is the appropriate control point when the ALB already has a web ACL associated. Therefore, implementing an AWS WAFrate-basedrule is the most effective and least disruptive immediate mitigation.

NEW QUESTION: 12
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: (SHOW ANSWER)
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

NEW QUESTION: 13
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 (LEAVE A REPLY)
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

NEW QUESTION: 14
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 (LEAVE A REPLY)
Amazon GuardDuty provides continuous monitoring for anomalous and malicious network activity by analyzing VPC Flow Logs, DNS logs, and CloudTrail events. Enabling GuardDuty across accounts requires minimal configuration and immediately satisfies the requirement to monitor endpoints for anomalous network traffic, as described in the AWS Certified Security - Specialty Study Guide.
Encrypting data in transit for applications behind Elastic Load Balancing is most efficiently achieved by using AWS Certificate Manager (ACM). ACM provisions and manages TLS certificates automatically, and integrating ACM with ELB enables encrypted communication without manual certificate management.
For encryption at rest in Amazon S3, AWS best practices recommend enforcing server-side encryption using AWS KMS. An S3 bucket policy that denies PutObject requests unless the x-amz-server-side-encryption condition is present ensures that all uploaded objects are encrypted at rest using KMS-managed keys. This provides strong encryption guarantees with minimal operational effort.
Option A is unnecessary because Amazon Inspector focuses on vulnerability assessment, not encryption or network anomaly detection. Option C adds network complexity and is not required to meet the stated requirements. Option E is incorrect because x-amz-meta-side-encryption is not a valid enforcement mechanism.
Referenced AWS Specialty Documents:
AWS Certified Security - Specialty Official Study Guide
Amazon GuardDuty Threat Detection
AWS Certificate Manager and ELB Integration

Amazon S3 Encryption Best Practices

NEW QUESTION: 15
A company is developing an application that runs across a combination of Amazon EC2 On-Demand Instances and Spot Instances. A security engineer needs to provide a logging solution that makes logs for all instances available from a single location. The solution must allow only a specific set of users to analyze the logs for events patterns. The users must be able to use SQL queries on the logs to perform root cause analysis.
Which solution will meet these requirements?
A. Configure the EC2 instances to send application logs to a single Amazon CloudWatch Logs log group.
Allow only specific users to access the log group. Use CloudWatch Logs Insights to query the log group.
B. Configure the EC2 instances to send application logs to a single Amazon S3 bucket. Allow only specific users to access the S3 bucket. Use Amazon CloudWatch Logs Insights to query the log files in the S3 bucket.
C. Configure each EC2 instance to send its application logs to its own specific Amazon CloudWatch Logs log group. Allow only specific users to access the log groups. Use Amazon Athena to query all the log groups.
D. Configure the EC2 instances to send application logs to a single Amazon CloudWatch Logs log group.Grant Amazon Detective access to the log group. Allow only specific users to use Detective to query the log group.
Answer: A (LEAVE A REPLY)
Option A satisfies all requirements with the most direct, purpose-built AWS logging workflow. By using the CloudWatch Agent (or fluent-bit / unified logging configuration) on each EC2 instance-regardless of whether it is On-Demand or Spot-the application logs can be centralized into asingle Amazon CloudWatch Logs log group. Centralization ensures the logs remain available even as Spot Instances are interrupted and replaced. Access control is handled withIAM policies(and optionally resource policies/KMS encryption) so that only a specific set of users can read/query the log group.
For analysis,CloudWatch Logs Insightsprovides an interactive query language that is SQL-like and commonly treated as "SQL queries" for troubleshooting. It enables fast filtering, aggregation, and pattern detection across large log volumes without building a separate data lake pipeline. This supports event-pattern analysis and root cause investigation directly from the centralized log group.
Option B is incorrect because Logs Insights queries CloudWatch Logs data, not arbitrary log files sitting in S3. Option C is inefficient (many log groups) and Athena cannot directly query CloudWatch log groups as a native data source. Option D is incorrect because Amazon Detective is for security investigations across supported data sources and is not the primary service for ad-hoc SQL-style querying of application logs.

NEW QUESTION: 16
A company is using AWS Organizations with the default SCP. The company needs to restrict AWS usage for all AWS accounts that are in a specific OU. Except for some desired global services, the AWS usage must occur only in theeu-west-1Region for all accounts in the OU. A security engineer must create an SCP that applies the restriction to existing accounts and any new accounts in the OU.
Which SCP will meet these requirements?
A. Deny with NotAction, but uses StringEquals for aws:RequestedRegion = eu-west-1
B. Allow with Action, scoped to desired global services in eu-west-1
C. Deny with NotAction for desired global services, and StringNotEquals aws:RequestedRegion = eu-west-
1
D. Allow with NotAction and StringNotEquals aws:RequestedRegion = eu-west-1
Answer: C (LEAVE A REPLY)
To restrict activity to a single Region in an OU using an SCP, the standard pattern is an explicitDenyfor requests madeoutsidethe allowed Region, while carving out exceptions forglobal servicesthat do not use aws:
RequestedRegion in the same way (or that must remain usable regardless of Region). This is done withEffect:
Deny, aConditionusing StringNotEquals on aws:RequestedRegion for the allowed Region (here, eu-west-1), andNotActionlisting the global services that should remain available.
This works because SCPs act asguardrails: an explicit Deny in an SCP overrides IAM Allow in member accounts, ensuring the restriction applies consistently to all existing and future accounts placed in the OU. The StringNotEquals condition ensures the deny triggers for any Region other than eu-west-1. The NotAction exception list ensures that the specified global services are not blocked by this deny statement.
Option A is wrong because StringEquals would deny actionsineu-west-1 rather than outside it.
Options B and D useAllowstatements, which do not enforce "only this Region" safely in SCPs unless combined with a comprehensive deny strategy; they would not reliably restrict all other services/regions. Therefore, option C is the correct SCP structure.

NEW QUESTION: 17

A public subnet contains two Amazon EC2 instances. The subnet has a custom network ACL. A security engineer is designing a solution to improve the subnet security. The solution must allow outbound traffic to an internet service that uses TLS through port 443. The solution also must deny inbound traffic that is destined for MySQL port 3306.
Which network ACL rule set meets these requirements?
A. Use inbound rule 100 to allow traffic on TCP port 443. Use inbound rule 200 to deny traffic on TCP port 3306. Use outbound rule 100 to allow traffic on TCP port 443.
B. Use inbound rule 100 to deny traffic on TCP port 3306. Use inbound rule 200 to allow traffic on TCP port range 1024-65535. Use outbound rule 100 to allow traffic on TCP port 443.
C. Use inbound rule 100 to allow traffic on TCP port range 1024-65535. Use inbound rule 200 to deny traffic on TCP port 3306. Use outbound rule 100 to allow traffic on TCP port 443.
D. Use inbound rule 100 to deny traffic on TCP port 3306. Use inbound rule 200 to allow traffic on TCP port 443. Use outbound rule 100 to allow traffic on TCP port 443.
Answer: B (LEAVE A REPLY)
Network ACLs arestateless, so you must allow both the outbound request and the inboundreturn traffic. For outbound TLS to an internet service on TCP443, you need an outbound allow rule permitting destination port
443. The return traffic from the internet service will come back to the instance'sephemeral port(typically in the range 1024-65535) on the inbound path. Therefore, you must allow inbound TCP traffic on the ephemeral port range to support established outbound connections.
At the same time, the requirement is todeny inbound MySQL (TCP 3306). Because NACLs process rules in order (lowest rule number first), placing an explicit deny for port 3306 as a low-numbered inbound rule ensures that traffic destined for MySQL is blocked even if there are broader allow rules later.
Option B does exactly this: it denies inbound TCP 3306 first, then allows inbound ephemeral ports for return traffic, and allows outbound TCP 443. Option A/D incorrectly allow inbound 443 (not needed for outbound- only TLS) and fail to explicitly allow ephemeral return traffic correctly.
Option C allows ephemeral inbound first, and then denies 3306 later; while 3306 is not in the ephemeral range, B is the clean, canonical ordering and matches the intended stateless-return-traffic pattern most directly.

NEW QUESTION: 18
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 (LEAVE A REPLY)

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

NEW QUESTION: 19
A company runs an online game on AWS. When players sign up for the game, their username and password credentials are stored in an Amazon Aurora database.
The number of users has grown to hundreds of thousands of players. The number of requests for password resets and login assistance has become a burden for the company ' s customer service team.
The company needs to implement a solution to give players another way to log in to the game. The solution must remove the burden of password resets and login assistance while securely protecting each player ' s credentials.
Which solution will meet these requirements?
A. When a new player signs up, use an AWS Lambda function to automatically create an IAM access key and a secret access key.
B. Migrate the player credentials from the Aurora database to AWS Secrets Manager.
C. Configure Amazon Cognito user pools to federate access to the game with third-party identity providers (IdPs), such as social IdPs. Migrate the game ' s authentication mechanism to Cognito.
D. Issue API keys to new and existing players and use Amazon API Gateway for authentication.
Answer: C (LEAVE A REPLY)
Amazon Cognito is a fully managed identity service that providesuser authentication, authorization, and user managementfor web and mobile applications. According to AWS Certified Security - Specialty documentation, Cognito user pools are specifically designed to offload authentication responsibilities from applications while maintaining strong security controls.

By federating authentication with third-party identity providers (such as social IdPs), Cognito eliminates the need for the company to manage user passwords directly. This dramatically reduces password reset requests and customer service overhead, while also improving security throughindustry-standard authentication mechanisms, including MFA and token-based access. Option A is insecure and incorrect because IAM access keys are not intended for end users. Option B simply relocates password storage and does not reduce operational burden. Option D uses API keys, which are not designed for user authentication and provide no identity verification. AWS guidance clearly states thatAmazon Cognito is the recommended service for scalable, secure user authentication, especially when reducing password management complexity is a requirement.
* AWS Certified Security - Specialty Official Study Guide
* Amazon Cognito User Pools Documentation
* AWS IAM Security Best Practices

NEW QUESTION: 20
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 (LEAVE A REPLY)
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

NEW QUESTION: 21

A company allows users to download its mobile app onto their phones. The app is MQTT based and connects to AWS IoT Core to subscribe to specific client-related topics. Recently, the company discovered that some malicious attackers have been trying to get a Trojan horse onto legitimate mobile phones. The Trojan horse poses as the authentic application and uses a client ID with injected special characters to gain access to topics outside the client ' s privilege scope.
Which combination of actions should the company take to prevent this threat? (Select TWO.)
A. In the application, use an IoT thing name as the client ID to connect the device to AWS IoT Core.
B. In the application, add a client ID check. Disconnect from the server if any special character is detected.
C. Apply an AWS IoT Core policy that allows " AWSIoTWirelessDataAccess " with the principal set to " client/${iot:Connection.Thing.ThingName} " .
D. Apply an AWS IoT Core policy to the device to allow " iot:Connect " with the resource set to " client
/${iot:ClientId} " .
E. Apply an AWS IoT Core policy to the device to allow " iot:Connect " with the resource set to " client
/${iot:Connection.Thing.ThingName} " .
Answer: A,E (LEAVE A REPLY)
The threat is client ID manipulation to break authorization boundaries. The strongest control is tobind the MQTT client identity to the authenticated device identity(the Thing) rather than trusting arbitrary client IDs provided by the client. Using theThing name as the client ID(Option A) removes ambiguity and makes the identifier predictable and tied to a registered identity.
On the authorization side, AWS IoT Core policies can use policy variables. Allowing iot:Connect only when the resource matches client/${iot:Connection.Thing.ThingName} (Option E) ensures the connection is permittedonlyif the client ID exactly equals the authenticated Thing name from the TLS certificate/Thing principal context. This prevents attackers from injecting special characters or choosing a different client ID to escalate access, because the policy evaluation ties the allowed client resource to the Thing identity, not the attacker-controlled string.
Option D is weaker because it effectively allows whatever client ID is presented (it matches the same value the client supplies), so it does not prevent crafted client IDs from being used. Option C is unrelated to the described MQTT connect authorization (and references an action not aligned with the scenario). Option B is an application-side check and can be bypassed by a malicious client; enforcement must be at AWS IoT Core policy level.

NEW QUESTION: 22
A company finds that one of its Amazon EC2 instances suddenly has a high CPU usage. The company does not know whether the EC2 instance is compromised or whether the operating system is performing background cleanup.
Which combination of steps should a security engineer take before investigating the issue? (Select THREE.)

A. Disable termination protection for the EC2 instance if termination protection has not been disabled.
B. Enable termination protection for the EC2 instance if termination protection has not been enabled.
C. Take snapshots of the Amazon Elastic Block Store (Amazon EBS) data volumes that are attached to the EC2 instance.
D. Remove all snapshots of the Amazon Elastic Block Store (Amazon EBS) data volumes that are attached to the EC2 instance.
E. Capture the EC2 instance metadata, and then tag the EC2 instance as under quarantine.
F. Immediately remove any entries in the EC2 instance metadata that contain sensitive information.
Answer: (SHOW ANSWER)
Before beginning an investigation, incident response best practice is topreserve evidence,prevent accidental loss of the asset, andclearly mark and control the potentially affected resource. Enablingtermination protection (Option B) helps ensure the instance is not accidentally terminated during triage, which would destroy volatile evidence and complicate forensics and recovery. TakingEBS snapshotsof all attached data volumes (Option C) preserves a point-in-time copy of disk evidence for later forensic analysis, malware scanning, or offline investigation. Snapshots allow responders to create forensic volumes or AMIs in an isolated environment without repeatedly touching the potentially compromised instance.
Capturinginstance metadataand tagging the instance asunder quarantine(Option E) supports both investigation and operational control. Metadata capture (instance ID, IAM role, network interfaces, security groups, user- data, tags, recent changes) provides context for responders. Quarantine tagging enables automated workflows (for example, incident runbooks that isolate the instance, restrict IAM, or move it to a quarantine security group) and signals to other teams/tools that the instance is under investigation.
Option A is the opposite of what you want. Option D destroys evidence. Option F is not an appropriate
"before investigation" step; altering metadata risks losing evidence and is not the primary containment approach.

NEW QUESTION: 23
A company ' s web application is hosted on Amazon EC2 instances running behind an Application Load Balancer (ALB) in an Auto Scaling group. An AWS WAF web ACL is associated with the ALB. AWS CloudTrail is enabled and stores logs in Amazon S3 and Amazon CloudWatch Logs.
The operations team has observed some EC2 instances reboot at random. After rebooting, all access logs on the instances have been deleted. During an investigation, the operations team found that each reboot happened just after a PHP error occurred on the new-user-creation.php file. The operations team needs to view log information to determine if the company is being attacked.
Which set of actions will identify the suspect attacker ' s IP address for future occurrences?

A. Configure VPC Flow Logs on the subnet where the ALB is located and stream the data to CloudWatch.
Search for the new-user-creation.php occurrences in CloudWatch.
B. Configure the CloudWatch agent on the ALB and send application logs to CloudWatch Logs.
C. Configure the ALB to export access logs to an Amazon OpenSearch Service cluster and search for the new-user-creation.php occurrences.
D. Configure the web ACL to send logs to Amazon Data Firehose, which delivers the logs to an S3 bucket.Use Amazon Athena to query the logs and find the new-user-creation.php occurrences.
Answer: D (LEAVE A REPLY)
AWS WAF logs capturedetailed request-level information, including source IP address, request URI, headers, and rule evaluation results. According to the AWS Certified Security - Specialty documentation, AWS WAF logging is acritical detection controlwhen application-level attacks are suspected, especially when host-based logs are unreliable or can be erased by attackers.
By configuring the AWS WAF web ACL to send logs toAmazon Data Firehose, the company ensures that all future requests are centrally captured and delivered to a durable storage service such as Amazon S3.
UsingAmazon Athena, the security team can query these logs to identify requests targeting specific application paths such as new-user-creation.php and extract the originating client IP addresses.
Option A is incorrect because VPC Flow Logs operate at the network layer and do not capture HTTP request paths. Option B is invalid because ALBs do not support CloudWatch agents.
Option C is viable but introduces additional operational complexity and cost, making it less appropriate than the native WAF logging solution.
AWS documentation highlightsAWS WAF logging combined with Athenaas a best practice for forensic analysis and attacker identification.
* AWS Certified Security - Specialty Official Study Guide
* AWS WAF Logging Documentation
* Amazon Athena User Guide
* AWS Detection and Monitoring Best Practices

NEW QUESTION: 24
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: (SHOW ANSWER)

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

NEW QUESTION: 25
A security engineer needs to implement AWS IAM Identity Center with an external identity provider (IdP).
Select and order the correct steps from the following list to meet this requirement. Select each step one time or not at all. (Select and order THREE.)
. Configure the external IdP as the identity source in IAM Identity Center.
. Create an IAM role that has a trust policy that specifies the IdP ' s API endpoint.
. Enable automatic provisioning in IAM Identity Center settings.
. Enable automatic provisioning in the external IdP.
. Obtain the SAML metadata from IAM Identity Center.
. Obtain the SAML metadata from the external IdP.

Step 1: Select...

- Select...
- Configure the external IdP as the identity source in IAM Identity Center.
- Create an IAM role that has a trust policy that specifies the IdP's API endpoint.
- Enable automatic provisioning in IAM Identity Center settings.
- Enable automatic provisioning in the external IdP.
- Obtain the SAML metadata from IAM Identity Center.
- Obtain the SAML metadata from the external IdP.

Step 2: Select...

- Select...
- Configure the external IdP as the identity source in IAM Identity Center.
- Create an IAM role that has a trust policy that specifies the IdP's API endpoint.
- Enable automatic provisioning in IAM Identity Center settings.
- Enable automatic provisioning in the external IdP.
- Obtain the SAML metadata from IAM Identity Center.
- Obtain the SAML metadata from the external IdP.

Step 3: Select...

- Select...
- Configure the external IdP as the identity source in IAM Identity Center.
- Create an IAM role that has a trust policy that specifies the IdP's API endpoint.
- Enable automatic provisioning in IAM Identity Center settings.
- Enable automatic provisioning in the external IdP.
- Obtain the SAML metadata from IAM Identity Center.
- Obtain the SAML metadata from the external IdP.

Answer:

Step 1: Select...

- Select...
- Configure the external IdP as the identity source in IAM Identity Center.
- Create an IAM role that has a trust policy that specifies the IdP's API endpoint.
- Enable automatic provisioning in IAM Identity Center settings.
- Enable automatic provisioning in the external IdP.
- Obtain the SAML metadata from IAM Identity Center.
- Obtain the SAML metadata from the external IdP.

Step 2: Select...

- Select...
- Configure the external IdP as the identity source in IAM Identity Center.
- Create an IAM role that has a trust policy that specifies the IdP's API endpoint.
- Enable automatic provisioning in IAM Identity Center settings.
- Enable automatic provisioning in the external IdP.
- Obtain the SAML metadata from IAM Identity Center.
- Obtain the SAML metadata from the external IdP.

Step 3: Select...

- Select...
- Configure the external IdP as the identity source in IAM Identity Center.
- Create an IAM role that has a trust policy that specifies the IdP's API endpoint.
- Enable automatic provisioning in IAM Identity Center settings.
- Enable automatic provisioning in the external IdP.
- Obtain the SAML metadata from IAM Identity Center.
- Obtain the SAML metadata from the external IdP.

Explanation:

Step 1:Obtain the SAML metadata fromIAM Identity Center.
Step 2:Obtain the SAML metadata from theexternal IdP.
Step 3:Configure theexternal IdP as the identity sourceinIAM Identity Center.
When integratingAWS IAM Identity Center (formerly AWS SSO)with anexternal identity provider (IdP) usingSAML 2.0, AWS requires a specific sequence of steps to establish trust and federation correctly.
Step 1: Obtain the SAML metadata from IAM Identity Center
IAM Identity Center acts as theservice provider (SP)in the SAML trust. The external IdP must trust IAM Identity Center, so the IdP needs IAM Identity Center's SAML metadata first. This metadata contains critical information such as the SP entity ID, ACS (Assertion Consumer Service) URL, and signing certificate.
Without this metadata, the external IdP cannot be configured to send assertions to AWS.
Step 2: Obtain the SAML metadata from the external IdP
After the external IdP is configured to trust IAM Identity Center, the IdP generates its own SAML metadata.
This metadata includes the IdP entity ID, SSO endpoint, and signing certificate. IAM Identity Center requires this information to validate authentication assertions coming from the external IdP.
Step 3: Configure the external IdP as the identity source in IAM Identity Center Once both metadata files are available, the security engineer configures the external IdP as theidentity sourcein IAM Identity Center. At this stage, IAM Identity Center imports the IdP metadata and establishes the SAML trust relationship. After this configuration, users authenticated by the external IdP can be federated into AWS accounts and applications via IAM Identity Center.
Why the other options are incorrect:
* Creating an IAM role with an IdP API endpoint is used forIAM federation, not IAM Identity Center.
* Automatic provisioning (SCIM) is optional and is configuredafterSAML federation is established.
* Automatic provisioning must be enabled onboth sides, but it is not required to complete the core IdP integration.
This sequence follows AWS best practices for SAML-based federation with IAM Identity Center.

NEW QUESTION: 26
A company's security team wants to receive email notification from AWS about any abuse reports regarding DoS attacks. A security engineer needs to implement a solution that will provide a near-real-time alert for any abuse reports that AWS sends for the account. The security engineer already has created an Amazon Simple Notification Service (Amazon SNS) topic and has subscribed the security team's email address to the topic.
What should the security engineer do next to meet these requirements?
A. Use the AWS Trusted Advisor API and a scheduled Lambda function to detect AWS_ABUSE_DOS_REPORT notifications.

B. Create an Amazon EventBridge rule that uses AWS Health and identifies a specific event for AWS_ABUSE_DOS_REPORT. Configure the rule action to publish a message to the SNS topic.
C. Use the AWS Support API and a scheduled Lambda function to detect abuse report cases.
D. Use AWS CloudTrail logs with metric filters to detect AWS_ABUSE_DOS_REPORT events.
Answer: B (LEAVE A REPLY)
AWS Health providesreal-time visibility into events that affect AWS accounts, including abuse notifications such asAWS_ABUSE_DOS_REPORT. According to the AWS Certified Security - Specialty Study Guide, AWS Health events are natively integrated withAmazon EventBridge, enabling automated, near-real-time responses without polling or custom code.
By creating an EventBridge rule that listens for AWS Health events related to abuse reports and configuring the rule to publish messages to an SNS topic, the security engineer ensures immediate notification to the security team whenever AWS issues a DoS-related abuse report for the account.
Option A and C rely on periodic polling using Lambda functions, which introduces latency and operational complexity. Option D is incorrect because CloudTrail does not log AWS abuse notifications.
AWS documentation explicitly identifiesAWS Health + EventBridge + SNSas the recommended architecture for near-real-time operational and security alerts originating from AWS.
* AWS Certified Security - Specialty Official Study Guide
* AWS Health User Guide
* Amazon EventBridge Documentation
* AWS Incident Response Best Practices

NEW QUESTION: 27
A security engineer is designing security controls for a fleet of Amazon EC2 instances that run sensitive workloads in a VPC. The security engineer needs to implement a solution to detect and mitigate software vulnerabilities on the EC2 instances.
Which solution will meet this requirement?
A. Scan the EC2 instances by using Amazon Inspector. Apply security patches and updates by using AWS Systems Manager Patch Manager.
B. Install host-based firewall and antivirus software on each EC2 instance. Use AWS Systems Manager Run Command to update the firewall and antivirus software.
C. Install the Amazon CloudWatch agent on the EC2 instances. Enable detailed logging. Use Amazon EventBridge to review the software logs for anomalies.
D. Scan the EC2 instances by using Amazon GuardDuty Malware Protection. Apply security patches and updates by using AWS Systems Manager Patch Manager.
Answer: A (LEAVE A REPLY)
To addresssoftware vulnerabilities, you need both (1) a vulnerability assessment capability and (2) a consistent patching mechanism.Amazon Inspectorcontinuously scans EC2 instances for known software vulnerabilities and exposures (CVEs), package-level issues, and security

misconfigurations relevant to the supported scan types. It provides prioritized findings and helps the security team understand which instances are exposed and why.
To mitigate those vulnerabilities,AWS Systems Manager Patch Managerprovides automated, policy-driven patching for fleets of EC2 instances. Patch Manager can schedule patch windows, control reboots, enforce baselines, and report compliance, allowing the company to remediate issues at scale with controlled operational impact.
Option B focuses on firewall/AV tooling, which can be helpful, but it is not a complete vulnerability detection- and-patching solution and is heavier to manage across large fleets. Option C is centered on log anomaly detection, not vulnerability management. Option D mixes GuardDuty Malware Protection (malware detection) with patching; GuardDuty is not a vulnerability scanner and does not replace Inspector for CVE detection. Therefore, Inspector + Patch Manager is the correct combined solution to detect and mitigate software vulnerabilities.

NEW QUESTION: 28
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 (LEAVE A REPLY)
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

NEW QUESTION: 29
A company uses AWS Config rules to identify Amazon S3 buckets that are not compliant with the company's data protection policy. The S3 buckets are hosted in several AWS Regions and several AWS accounts. The accounts are in an organization in AWS Organizations. The company needs a solution to remediate the organization ' s existing noncompliant S3 buckets and any noncompliant S3 buckets that are created in the future.
Which solution will meet these requirements?
A. Deploy an AWS Config aggregator with organization-wide resource data aggregation. Create an AWS Lambda function that responds to AWS Config findings of noncompliant S3 buckets by deleting or reconfiguring the S3 buckets.
B. Deploy an AWS Config aggregator with organization-wide resource data aggregation. Create an SCP that contains a Deny statement that prevents the creation of new noncompliant S3 buckets. Apply the SCP to all OUs in the organization.
C. Deploy an AWS Config aggregator that scopes only the accounts and Regions that the company currently uses. Create an AWS Lambda function that responds to AWS Config findings of noncompliant S3 buckets by deleting or reconfiguring the S3 buckets.
D. Deploy an AWS Config aggregator that scopes only the accounts and Regions that the company currently uses. Create an SCP that contains a Deny statement that prevents the creation of new noncompliant S3 buckets. Apply the SCP to all OUs in the organization.
Answer: A (LEAVE A REPLY)
The requirement includesremediating existing noncompliant bucketsand also handlingfuture noncompliant bucketsacross multiple accounts and Regions. Anorganization-wide AWS Config aggregatorprovides centralized visibility into compliance status across all member accounts/Regions. To remediate, AWS Config can trigger automation (commonly via EventBridge/SNS) that invokes anAWS Lambda function(or SSM Automation) to take corrective action-such as enabling default encryption, blocking public access, or applying required bucket policies-whenever a bucket is evaluated as noncompliant. This addresses both existing findings (by running remediation on current noncompliant resources) and new ones (by automatically reacting when new buckets appear or configurations drift).
Options C and D are insufficient because they scope only to "accounts and Regions the company currently uses," which fails the multi-Region/multi-account organization requirement and would miss future expansion.
Option B helps prevent creation of new noncompliant resources but does not remediate existing buckets, which is explicitly required. Therefore, the combination of an org-wide aggregator plus automated remediation on Config noncompliance is the best fit.

NEW QUESTION: 30

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 (LEAVE A REPLY)

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

NEW QUESTION: 31
A company has multiple accounts in the AWS Cloud. Users in the developer account need to have access to specific resources in the production account.
What is the MOST secure way to provide this access?
A. Create one IAM user in the production account. Grant the appropriate permissions to the resources that are needed. Share the password only with the users that need access.
B. Create cross-account access with an IAM role in the developer account. Grant the appropriate permissions to this role. Allow users in the developer account to assume this role to access the production resources.
C. Create cross-account access with an IAM user account in the production account. Grant the appropriate permissions to this user account. Allow users in the developer account to use this user account to access the production resources.
D. Create cross-account access with an IAM role in the production account. Grant the appropriate permissions to this role. Allow users in the developer account to assume this role to access the production resources.
Answer: D (LEAVE A REPLY)
The most secure and AWS-recommended pattern for cross-account access is to create anIAM role in the target account (production)and allow trusted principals from the source account (developer) toassume the roleby using AWS STS. This avoids long-term credentials in the production account, supports short-lived session credentials, and enables strong controls such as MFA requirements, session duration limits, and precise least-privilege permissions attached to the role. It also centralizes ownership of production permissions in the production account, which is important for separation of duties and governance.
Option A is insecure because it requires password sharing and uses a long-lived IAM user credential, which is against AWS best practices. Option C is also poor because it relies on a long-lived IAM user in the production account and encourages credential sharing/duplication. Option B places the role in the developer account; while you can attach permissions there, access to production resources is governed by the production account.

The standard approach is a production-account role with a trust policy that names the developer account principals (or a role) as allowed to assume it. Therefore, Option D is the most secure solution.

NEW QUESTION: 32
A security engineer needs to prepare for a security audit of an AWS account.
Select the correct AWS resource from the following list to meet each requirement. Select each resource one time or not at all. (Select THREE.)
* AWS Artifact reports
* AWS Audit Manager controls
* AWS Config conformance packs
* AWS Config rules
* Amazon Detective investigations
* AWS Identity and Access Management Access Analyzer internal access analyzers

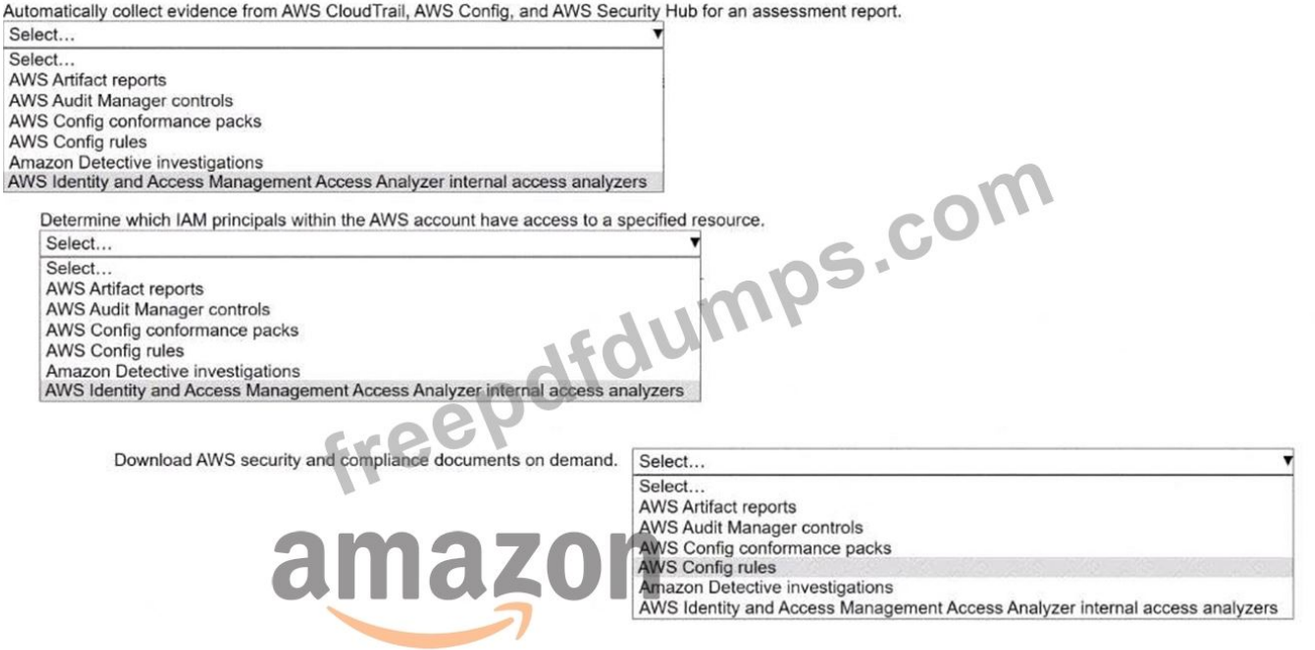

Answer:

Explanation:

Requirements and Correct Selections

Automatically collect evidence from AWS CloudTrail, AWS Config, and AWS Security Hub for an assessment report.

answer:

AWS Audit Manager controls

Why:

AWS Audit Manager is specifically designed toautomatically collect, map, and organize evidencefrom AWS services such as CloudTrail, AWS Config, and AWS Security Hub. Audit Manager controls are used within audit frameworks to continuously gather evidence and generate assessment reports for compliance audits.

Determine which IAM principals within the AWS account have access to a specified resource.

answer:

AWS Identity and Access Management Access Analyzer internal access analyzers Why:

IAM Access Analyzer internal access analyzers are used toidentify which IAM users, roles, or services within an account or organization have access to a specific resource. This is a core access visibility and audit requirement for IAM reviews.

Download AWS security and compliance documents on demand.

answer:

AWS Artifact reports

Why:

AWS Artifact provideson-demand access to AWS security, compliance, and audit reports, including SOC reports, ISO certifications, and compliance attestations. This service is explicitly intended for audit preparation and regulatory documentation.

NEW QUESTION: 33

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 (LEAVE A REPLY)
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

NEW QUESTION: 34
A company uses Amazon EC2 instances to host frontend services behind an Application Load Balancer.
Amazon Elastic Block Store (Amazon EBS) volumes are attached to the EC2 instances. The company uses Amazon S3 buckets to store large files for images and music. The company has implemented a security architecture on AWS to prevent, identify, and isolate potential ransomware attacks. The company now wants to further reduce risk. A security engineer must develop a disaster recovery solution that can recover to normal operations if an attacker bypasses preventive and detective controls. The solution must meet an RPO of1 hour.
Which solution will meet these requirements?

A. Use AWS Backup to create backups of the EC2 instances and S3 buckets every hour. Create AWS CloudFormation templates that replicate existing architecture components. Use a Git repository to store the CloudFormation templates alongside application configuration code.
B. Use AWS Backup to create backups of the EBS volumes and S3 objects every day. Use Amazon Security Lake to create a centralized data lake for AWS CloudTrail logs and VPC flow logs. Use the logs for automated response.
C. Use Amazon Security Lake to create a centralized data lake for AWS CloudTrail logs and VPC flow logs. Use the logs for automated response. Enable AWS Security Hub to establish a single location for recovery procedures. Create AWS CloudFormation templates that replicate existing architecture components. Use a Git repository to store the CloudFormation templates alongside application configuration code.
D. Create EBS snapshots every 4 hours. Enable Amazon GuardDuty Malware Protection. Create automation to immediately restore the most recent snapshot for any EC2 instances that produce an Execution:EC2/MaliciousFile finding in GuardDuty.
Answer: A (LEAVE A REPLY)
An RPO of1 hourmeans the company must be able to restore data with at most60 minutes of loss. Option A directly meets this by usingAWS Backupto takehourly backupsof both the compute layer (EC2) and the data layer (S3). AWS Backup provides centralized policy-based scheduling, retention, and (when configured) immutable protections such as Backup Vault Lock to help defend backups from tampering-important in ransomware recovery scenarios. Backing up the S3 buckets hourly also addresses recovery of critical objects such as images and music that users rely on.
In addition, recovery to "normal operations" is not only about data restoration; it also requires rapidly re- creating infrastructure reliably. UsingAWS CloudFormation templatesstored in aversion-controlled Git repositorysupports consistent, repeatable rebuilds of the ALB, EC2 fleet configuration, IAM roles, security groups, and related components. This infrastructure-as-code approach reduces human error under incident pressure and accelerates disaster recovery execution.
Option B fails the RPO because daily backups can lose up to 24 hours of data. Option C focuses on logging
/governance rather than backups and does not establish a 1-hour recovery point. Option D also fails the RPO (4-hour snapshots) and is reactive to a specific finding type rather than a comprehensive DR plan for EC2 and S3.

NEW QUESTION: 35
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 (LEAVE A REPLY)
KMS grantsare purpose-built for givingtemporary, scoped permissionsto use a customer managed key without continually editing key policies. A grant can allow only the specific cryptographic operations the application process needs (for example, Encrypt/Decrypt/GenerateDataKey) and can be constrained to a particular AWS principal and (optionally) conditions. When access is no longer required, the security team canrevokethe grant immediately, returning the key to its previous access posture. This meets the "temporary access occasionally" requirement with minimal operational work and lower risk than policy churn.
Option B requires repeated key policy edits and rollbacks, which is operationally noisy and increases the chance of misconfiguration or leaving access in place longer than intended. Option A is not supported in the way described-KMS key material for standard KMS customer managed keys is not something you "export" for sharing, and exposing key material breaks managed key controls. Option D introduces significant overhead and complexity by generating/importing key material repeatedly and is not the intended model for occasional access. Grants are the standard AWS pattern for delegating KMS key usage temporarily and safely.

NEW QUESTION: 36
A company has two AWS accounts: Account A and Account B. Each account has a VPC. An application that runs in the VPC in Account A needs to write to an Amazon S3 bucket in Account B. The application in Account A already has permission to write to the S3 bucket in Account B. The application and the S3 bucket are in the same AWS Region. The company cannot send network traffic over the public internet.
Which solution will meet these requirements?
A. In both accounts, create a transit gateway and VPC attachments in a subnet in each Availability Zone.Update the VPC route tables.
B. Deploy a software VPN appliance in Account A. Create a VPN connection between the software VPN appliance and a virtual private gateway in Account B.
C. Create a VPC peering connection between the VPC in Account A and the VPC in Account B. Update the VPC route tables, network ACLs, and security groups to allow network traffic between the peered IP ranges.

D. In Account A, create a gateway VPC endpoint for Amazon S3. Update the VPC route table in Account A.

Answer: (SHOW ANSWER)

To keep S3 accessoff the public internet, the standard AWS approach is to use anAmazon S3 gateway VPC endpoint(AWS PrivateLink for S3 is not used; S3 uses gateway endpoints). A gateway endpoint adds routes in the VPC route tables so traffic destined for S3 stays on the AWS backbone network rather than traversing an internet gateway, NAT gateway, or public IP paths. This satisfies the "cannot send traffic over the public internet" requirement while allowing the application in Account A to reach S3 in the same Region.

Cross-account bucket access is controlled byIAM and the S3 bucket policy, not by networking between the two accounts' VPCs. The bucket resides in S3 (a regional service), not inside Account B's VPC, so connecting VPC-to-VPC (peering, transit gateway, VPN) does not inherently provide private access to S3.

Those options would add complexity and still typically require internet/NAT unless S3 endpoints are used.

With the gateway endpoint in Account A, the application can privately reach S3, and because permissions are already granted to write to the bucket in Account B, the write operations will succeed without public internet routing.

NEW QUESTION: 37

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 (LEAVE A REPLY)

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 Certified Security - Specialty Official Study Guide
AWS Lambda Execution Roles
Amazon CloudWatch Logs Integration with Lambda

NEW QUESTION: 38
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: (SHOW ANSWER)
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

NEW QUESTION: 39

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 (LEAVE A REPLY)
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

NEW QUESTION: 40
A company recently set up Amazon GuardDuty and is receiving a high number of findings from IP addresses within the company. A security engineer has verified that these IP addresses are trusted and allowed.

Which combination of steps should the security engineer take to configure GuardDuty so that it does not produce findings for these IP addresses? (Select TWO.)

A. Create a plaintext configuration file that contains the trusted IP addresses.

B. Create a JSON configuration file that contains the trusted IP addresses.

C. Upload the configuration file directly to GuardDuty.

D. Upload the configuration file to Amazon S3. Add a new trusted IP list to GuardDuty that points to the file.

E. Manually copy and paste the configuration file data into the trusted IP list in GuardDuty.

Answer: A,D (LEAVE A REPLY)

GuardDuty supports "Trusted IP lists" to suppress findings that would otherwise be generated for activity originating from known safe IP addresses (for example, corporate NAT egress IPs, security scanners, or monitoring systems). To use a trusted IP list, you create aplain textfile that contains the IP addresses (typically one per line or in supported list form) and store it inAmazon S3. You then configure GuardDuty to reference that S3 object as a trusted IP list. GuardDuty periodically retrieves the file from S3 and uses it to adjust finding generation accordingly.

That maps directly to Option A (create a plaintext file) and Option D (upload to S3 and create a trusted IP list in GuardDuty pointing to the file).

Options B and E are incorrect because GuardDuty trusted IP lists are not configured by pasting JSON into the console; they are sourced from an S3-hosted text list. Option C is not supported because GuardDuty does not accept direct file uploads into the service as the configuration source; S3 is the expected integration point for IP lists and threat intel lists.

NEW QUESTION: 41

A company has a large fleet of Amazon Linux 2 Amazon EC2 instances that run an application. The application processes sensitive data and has the following compliance requirements:

* No remote access management ports to the EC2 instances can be exposed internally or externally.
* All remote session activity must be recorded in an audit log.
* All remote access to the EC2 instances must be authenticated and authorized by AWS IAM Identity Center.

The company's DevOps team occasionally needs to connect to one of the EC2 instances to troubleshoot issues.

Which solution will provide remote access to the EC2 instances while meeting the compliance requirements?

A. Grant access to the EC2 serial console at the account level.

B. Enable EC2 Instance Connect and configure security group rules.

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. Assign the policy to an IAM role of the DevOps team.

D. Use AWS Systems Manager Automation runbooks to open remote access ports.

Answer: C (LEAVE A REPLY)

AWS Systems Manager Session Manager providessecure, auditable, and portless accessto EC2 instances.
According to the AWS Certified Security - Specialty Study Guide, Session Manager allows administrators to connect to instanceswithout opening inbound SSH or RDP ports, fully satisfying strict compliance requirements.
Session Manager integrates directly withAWS IAM Identity Center, ensuring that all access is authenticated and authorized using centralized identity management. Additionally, Session Manager automatically records session activity and can send logs to Amazon CloudWatch Logs or Amazon S3, providing a complete audit trail of all commands executed during a session.
Option A (EC2 serial console) does not provide comprehensive auditing and is intended for recovery scenarios. Option B requires inbound network access and security group rules, violating the "no exposed management ports" requirement. Option D explicitly opens ports, which directly violates compliance constraints.
AWS documentation clearly identifiesSystems Manager Session Manager as the recommended solution for secure, auditable, and identity-integrated instance accessin regulated environments.
* AWS Certified Security - Specialty Official Study Guide
* AWS Systems Manager Session Manager Documentation
* AWS IAM Identity Center Best Practices

NEW QUESTION: 42
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 (LEAVE A REPLY)
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

NEW QUESTION: 43
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 (LEAVE A REPLY)
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

NEW QUESTION: 44
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 (LEAVE A REPLY)
EventBridge natively consumes CloudTrail management events and provides near-real-time notifications.

NEW QUESTION: 45
A company runs an application on a fleet of Amazon EC2 instances. The company can remove instances from the fleet without risk to the application. All EC2 instances use the same security group named ProdFleet.
Amazon GuardDuty and AWS Config are active in the company's AWS account.
A security engineer needs to provide a solution that will prevent an EC2 instance from sending outbound traffic if GuardDuty generates a cryptocurrency finding event. The security engineer creates a new security group named Isolate that contains no outbound rules. The security engineer configures an AWS Lambda function to remove an EC2 instance from the ProdFleet security group and add it to the Isolate security group.
Which additional step will meet this requirement?
A. Configure GuardDuty to directly invoke the Lambda function if GuardDuty generates a CryptoCurrency:EC2/* finding event.
B. Configure an AWS Config rule that invokes the Lambda function if a CryptoCurrency:EC2/* configuration change event occurs for an EC2 instance.
C. Configure an Amazon EventBridge rule that invokes the Lambda function if GuardDuty generates a CryptoCurrency:EC2/* finding event.
D. Configure an Amazon EventBridge rule that invokes the Lambda function if AWS Config detects a CryptoCurrency:EC2/* configuration change event for an EC2 instance.
Answer: C (LEAVE A REPLY)
Amazon GuardDuty generates security findings when it detects suspicious or malicious activity, including CryptoCurrency:EC2/* findings that indicate an EC2 instance may be involved in

unauthorized cryptocurrency mining. According to AWS Certified Security - Specialty documentation, GuardDuty findings are published as events toAmazon EventBridge(formerly Amazon CloudWatch Events).
Amazon EventBridge is the recommended service for buildingautomated incident response workflows. By creating an EventBridge rule that listens for GuardDuty findings of type CryptoCurrency:EC2/*, the security engineer can automatically invoke a Lambda function to isolate the affected EC2 instance by modifying its security group attachments.
Option A is incorrect because GuardDuty does not directly invoke Lambda functions. Option B and Option D are incorrect because AWS Config tracks configuration compliance and resource changes, not real-time threat detection events. Cryptocurrency findings are security detections, not configuration changes.
AWS documentation explicitly describes this pattern-GuardDuty # EventBridge # Lambda # remediation action-as a best practice for automated threat response and containment.
* AWS Certified Security - Specialty Official Study Guide
* Amazon GuardDuty User Guide - Findings
* Amazon EventBridge User Guide
* AWS Incident Response Best Practices

NEW QUESTION: 46
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 (LEAVE A REPLY)
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

NEW QUESTION: 47
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 security engineer's 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. Test to ensure the vulnerability has been mitigated, then restore the security group to the original setting.
Answer: A (LEAVE A REPLY)
The fastest, least-effort way to mitigate SQL injection at the edge-without modifying legacy application code-is to place the application behind a component that supportsAWS WAFand applymanaged SQL injection protections. AnApplication Load Balancerintegrates directly with AWS WAF, allowing the security engineer to deploy a web ACL (including AWS Managed Rules for SQL injection and custom patterns based on the provided samples) and immediately start blocking malicious payloads before they reach the EC2 instances and the database.
Option A also preserves normal operations during rollout: you can create the ALB, register the existing EC2 instances as targets, validate health checks and traffic behavior, apply WAF protections, and then shift Route
53 weighted records to the ALB with minimal downtime. Finally, tightening the EC2 security groups to prevent direct internet access ensures all inbound web traffic is forced through the ALB + WAF inspection point, reducing exposure quickly.
Option B is risky because it uses only one EC2 origin (reducing availability) and adds CloudFront origin configuration complexity under a 24-hour deadline. Option C requires code changes on unsupported software and is unlikely to be safely delivered in time. Option D is invalid because AWS WAF cannot be attached directly to EC2 instances, and changing DB-port exposure doesn't address SQL injection on the web layer.

NEW QUESTION: 48
A security engineer is responding to an incident that is affecting an AWS account. The ID of the account is
123456789012. The attack created workloads that are distributed across multiple AWS Regions. The security engineer contains the attack and removes all compute and storage resources from all affected Regions. However, the attacker also created an AWS KMS key. The key policy on the KMS key explicitly allows IAM principal kms:* permissions.
The key was scheduled to be deleted the previous day. However, the key is still enabled and usable. The key has an ARN of arn:aws:kms:us-east-2:123456789012:key/mrk-0bb0212cd9864fdea0dcamzo26efb5670.
The security engineer must delete the key as quickly as possible.
Which solution will meet this requirement?
A. Log in to the account by using the account root user credentials. Re-issue the deletion request for the KMS key with a waiting period of 7 days.

B. Identify the other Regions where the KMS key ID is present and schedule the key for deletion in 7 days.
C. Update the IAM principal to allow kms:* permissions on the KMS key ARN. Re-issue the deletion request for the KMS key with a waiting period of 7 days.
D. Disable the KMS key. Re-issue the deletion request for the KMS key in 30 days.
Answer: (SHOW ANSWER)
AWS KMS enforces amandatory minimum waiting period of 7 daysbefore a customer managed key can be deleted. According to AWS Certified Security - Specialty incident response guidance,no method exists to immediately delete a KMS key. The fastest possible deletion is achieved by scheduling deletion with the minimum 7-day waiting period.
In this scenario, although deletion was previously scheduled, the key remains enabled and usable. The most authoritative and reliable method to regain control and reissue deletion immediately is touse the AWS account root user, which has implicit permissions to manage KMS keys regardless of compromised IAM principals.
Option B is incorrect because KMS keys are regional resources; multi-Region keys require coordinated deletion but do not shorten the waiting period. Option C is unnecessary because the key policy already allows kms:*. Option D increases the deletion waiting period to 30 days, which violates the requirement to delete the key as quickly as possible.
AWS documentation clearly states thatroot user access is the ultimate authority for KMS key management and that7 days is the minimum deletion window, making this the fastest valid option.
* AWS Certified Security - Specialty Official Study Guide
* AWS Key Management Service Developer Guide
* AWS Incident Response Best Practices

NEW QUESTION: 49
A startup company is using a single AWS account that has resources in a single AWS Region. A security engineer configures an AWS CloudTrail trail in the same Region to deliver log files to an Amazon S3 bucket by using the AWS CLI. Because of expansion, the company adds resources in multiple Regions. The security engineer notices that the logs from the new Regions are not reaching the S3 bucket.
What should the security engineer do to fix this issue with the LEAST amount of operational overhead?
A. Create a new CloudTrail trail. Select the new Regions where the company added resources.
B. Change the S3 bucket to receive notifications to track all actions from all Regions.
C. Create a new CloudTrail trail that applies to all Regions.
D. Change the existing CloudTrail trail so that it applies to all Regions.
Answer: (SHOW ANSWER)
CloudTrail trails can be configured as eithersingle-Regionormulti-Region. A single-Region trail delivers events only for the Region in which it is created; when the company begins using additional Regions, those events will not be logged to the existing S3 destination unless the trail is updated. The least operational overhead is tomodify the existing trailand enablemulti-

Regionlogging (that is, configure the trail to "apply to all Regions"). This preserves the current S3 bucket destination, encryption settings, log file validation configuration, IAM roles, and any integrations that already depend on the existing trail name and S3 prefix structure.
Creating a new trail (Options A or C) increases overhead because the engineer would need to manage multiple trails or migrate downstream processes (alerting, Athena tables, SIEM ingestion, retention policies, and lifecycle rules). Option B is not how CloudTrail works-S3 notifications do not cause CloudTrail to collect events from other Regions; CloudTrail must be configured to generate and deliver those logs.
Therefore, updating the existing trail to be multi-Region is the simplest and most maintainable fix.

NEW QUESTION: 50
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 (LEAVE A REPLY)
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

NEW QUESTION: 51
A security engineer needs to protect a public web application that runs in a VPC. The VPC hosts the origin for an Amazon CloudFront distribution. The application has experienced multiple layer 7 DDoS attacks. An AWS WAF web ACL is associated with the CloudFront distribution. The web

ACL contains one AWS managed rule to protect against known IP addresses that have bad reputations.
The security engineer must configure an automated solution that detects and mitigates layer 7 DDoS attacks in real time with no manual effort.
Which solution will meet these requirements?
A. Enable AWS Shield Advanced on the CloudFront distribution. Configure alerts in Amazon CloudWatch for DDoS indicators.
B. Enable AWS Shield Advanced and configure proactive engagement with the AWS DDoS Response Team (DRT).
C. Deploy AWS Network Firewall in the VPC. Create security policies that detect DDoS indicators.
Create an AWS Lambda function to automatically update the web ACL rules during an attack.
D. Add a rate-based rule to the web ACL. Enable AWS Shield Advanced. Enable automatic application layer DDoS mitigation on the CloudFront distribution.
Answer: D (LEAVE A REPLY)
Option D is the correct solution because it provides fully automated, real-time detection and mitigation of application-layer (Layer 7) DDoS attacks with no manual intervention. AWS Shield Advanced includes automatic application layer DDoS mitigation when it is enabled for supported resources such as Amazon CloudFront distributions. This feature continuously monitors traffic patterns and, when an attack is detected, automatically deploys AWS WAF rules to mitigate malicious requests.
Adding a rate-based rule to the AWS WAF web ACL further strengthens protection by automatically blocking IP addresses that exceed a defined request threshold, which is a common characteristic of Layer 7 DDoS attacks. This combination aligns directly with AWS best practices for protecting web applications against volumetric and application-layer threats.
Option A only provides alerting and visibility but does not ensure automated mitigation. Option B includes proactive engagement with the AWS DDoS Response Team, which is valuable for complex or large-scale attacks but still involves human interaction and therefore does not meet the "no manual effort" requirement.
Option C introduces unnecessary complexity and is not recommended for protecting CloudFront-based applications against Layer 7 DDoS attacks.
AWS Security Specialty documentation explicitly recommends AWS Shield Advanced with automatic application layer DDoS mitigation and AWS WAF rate-based rules for fully automated, real-time protection of public web applications.

NEW QUESTION: 52
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 (LEAVE A REPLY)
During an incident, the company wants to block "external access to the VPC" (for example, shutting down VPN ingress or internet-exposed paths) yet still allow the security team to access instances for forensics.
AnEC2 Instance Connect Endpoint (EIC Endpoint)provides a managed, private connectivity path that lets authorized users connect to instances in a VPCwithout requiring inbound access from the internet or from on- premises. The endpoint lives inside the VPC, and access is controlled by IAM permissions plus security group rules between the endpoint and the instances. This supports incident containment (no external network entry) while preserving controlled administrative access for investigation.
Options A and B require installing Instance Connect on the instances and typically rely on network reachability patterns that may be blocked when external access is cut off; they also do not provide the same VPC-resident endpoint model. With an EIC endpoint, the security team can use theAWS Management Consoleto initiate connections (Option D) even while the VPC is isolated from on-prem and the public internet, because the connectivity is mediated through AWS control plane and the endpoint inside the VPC.
Option C mentions using the CLI to open a tunnel, but the most straightforward and commonly used operational method for responders is console-based access via the EIC endpoint.
Therefore, creating an Instance Connect Endpoint and using the console meets the requirement.

NEW QUESTION: 53
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: (SHOW ANSWER)
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 Certified Security - Specialty Official Study Guide
AWS Organizations Delegated Administration
AWS CloudFormation StackSets for Multi-Account Governance

NEW QUESTION: 54
A security engineer needs to prepare a company ' s Amazon EC2 instances for quarantine during a security incident. The AWS Systems Manager Agent (SSM Agent) has been deployed to all EC2 instances. The security engineer has developed a script to install and update forensics tools on the EC2 instances.
Which solution will quarantine EC2 instances during a security incident?

A. Create a rule in AWS Config to track SSM Agent versions.
B. Configure Systems Manager Session Manager to deny all connection requests from external IP addresses.
C. Store the script in Amazon S3 and grant read access to the instance profile.
D. Configure IAM permissions for the SSM Agent to run the script as a predefined Systems Manager Run Command document.
Answer: (SHOW ANSWER)
AWS Systems Manager Run Command enables security engineers toremotely and securely execute scripts on EC2 instanceswithout requiring SSH or inbound network access. According to AWS Certified Security - Specialty incident response guidance, Run Command is a foundational tool forinstance quarantine and forensic preparation.
By configuring IAM permissions that allow the SSM Agent to execute a predefined Run Command document, the security engineer can rapidly deploy forensic tools, disable services, or modify system configurations across affected EC2 instances during an incident. This approach aligns with AWS best practices forcontainment and evidence preservation, while maintaining auditability through Systems Manager logs.
Option A only provides visibility, not quarantine capability. Option B restricts access but does not allow forensic tooling. Option C enables access to the script but does not execute it.
AWS documentation emphasizes thatSystems Manager Run Command is the recommended mechanism for incident response automation and quarantine actionson EC2 instances.
* AWS Certified Security - Specialty Official Study Guide
* AWS Systems Manager Run Command Documentation
* AWS Incident Response Best Practices

NEW QUESTION: 55
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 (LEAVE A REPLY)
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

NEW QUESTION: 56
A company is using AWS to run a long-running analysis process on data that is stored in Amazon S3 buckets.
The process runs on a fleet of Amazon EC2 instances in an Auto Scaling group. The EC2 instances are deployed in a private subnet that does not have internet access.
The EC2 instances access Amazon S3 through an S3 gateway endpoint that has the default access policy.
Each EC2 instance uses an instance profile role that allows s3:GetObject and s3:PutObject only for required S3 buckets.
The company learns that one or more EC2 instances are compromised and are exfiltrating data to an S3 bucket that isoutside the company's AWS Organization. The processing job must continue to function.
Which solution will meet these requirements?
A. Update the policy on the S3 gateway endpoint to allow S3 actions only if aws:ResourceOrgId and aws:
PrincipalOrgId match the company's organization.
B. Update the instance profile role policy to require aws:ResourceOrgId.
C. Add a network ACL rule to block outbound traffic on port 443.
D. Apply an SCP that restricts S3 actions using organization condition keys.
Answer: A (LEAVE A REPLY)
Amazon S3 gateway endpoints supportendpoint policiesthat can restrict which S3 resources are accessible through the endpoint. According to AWS Certified Security - Specialty documentation, endpoint policies are evaluatedin addition to IAM policiesand are ideal for enforcingdata exfiltration controlswithout breaking legitimate workloads.
By updating the S3 gateway endpoint policy to require both aws:ResourceOrgId and aws:PrincipalOrgId to match the company's AWS Organization, the security engineer ensures that EC2 instances can accessonly S3 buckets that belong to the organization. This immediately blocks exfiltration to external S3 buckets while allowing legitimate internal data access to continue uninterrupted.

Option B is insufficient because IAM policies alone cannot prevent access when the endpoint allows it.
Option C would break all S3 access and stop the processing job. Option D applies too broadly and can impact unrelated services across the account.
AWS documentation highlightsS3 VPC endpoint policies with organization condition keysas a best practice for preventing S3 data exfiltration in private VPC environments.
* AWS Certified Security - Specialty Official Study Guide
* Amazon S3 VPC Endpoint Policy Documentation
* AWS Organizations Condition Keys Documentation

NEW QUESTION: 57
A security engineer discovers that a company ' s user passwords have no required minimum length. The company is using the following two 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 that the company developed Which combination of actions should the security engineer take to implement a required minimum length for the passwords? (Select TWO.)
A. Update the password length policy in the IAM configuration.
B. Update the password length policy in the Cognito configuration.
C. Update the password length policy in the on-premises Active Directory configuration.
D. Create an SCP in AWS Organizations. Configure the SCP to enforce a minimum password length for IAM and Cognito.
E. Create an IAM policy that includes a condition for minimum password length. Enforce the policy for IAM and Cognito.
Answer: (SHOW ANSWER)
The company uses two different identity systems, and password policy must be enforcedat the system that actually stores and manages the passwords. For users authenticating throughIAM federation with on-premises Active Directory, IAM is not storing the users' passwords; the password policy is enforced byActive Directory. Therefore, the minimum password length must be configured in theon-premises AD password policyso federated users are subject to the requirement during password creation/changes.
For the cloud application that usesAmazon Cognito user poolsas its user database, Cognitodoesstore and manage user passwords for those users. Cognito user pools include a configurable password policy (minimum length and complexity requirements). Updating the Cognito user pool password policy enforces the required minimum length for the application's users going forward.
Options D and E are not applicable. Service control policies (SCPs) restrict AWS API actions; they cannot enforce end-user password-length rules inside AD or Cognito. Similarly, IAM policies control authorization to AWS resources and APIs, not password complexity/length requirements across external IdPs or Cognito user databases. Updating IAM password policy (Option A) would

apply only toIAM users(local users in AWS), which is not the authentication model described for the federated workforce.

NEW QUESTION: 58
A security engineer for a company is investigating suspicious traffic on a web application in the AWS Cloud.
The web application is protected by an Application Load Balancer (ALB) behind an Amazon CloudFront distribution. There is an AWS WAF web ACL associated with the ALB. The company stores AWS WAF logs in an Amazon S3 bucket.
The engineer notices that all incoming requests in the AWS WAF logs originate from a small number of IP addresses that correspond to CloudFront edge locations. The security engineer must identify the source IP addresses of the clients that are initiating the suspicious requests.
Which solution will meet this requirement?
A. Enable VPC Flow Logs in the VPC where the ALB is deployed. Examine the source field to capture the client IP addresses.
B. Inspect the X-Forwarded-For header in the AWS WAF logs to determine the original client IP addresses.
C. Modify the CloudFront distribution to disable ALB connection reuse. Examine the clientIp field in the AWS WAF logs to identify the original client IP addresses.
D. Configure CloudFront to add a custom header named Client-IP to origin requests that are sent to the ALB.
Answer: B (LEAVE A REPLY)
When Amazon CloudFront is used in front of an Application Load Balancer, CloudFront becomes the immediate source of incoming requests to the ALB. As a result, AWS WAF logs record theCloudFront edge location IP addressesas the client IPs, not the original viewer IP addresses. This behavior is explicitly documented in the AWS Certified Security - Specialty Study Guide and the AWS WAF and CloudFront integration documentation.
To preserve the original client IP address, CloudFront automatically adds theX-Forwarded-For HTTP header, which contains the IP address of the originating client followed by any proxy addresses involved in forwarding the request. AWS WAF logs include this header, making it the authoritative source for identifying true client IP addresses when CloudFront is used.
Option A is incorrect because VPC Flow Logs capture network-level metadata and will only show CloudFront IP addresses, not the original client IPs. Option C is incorrect because disabling connection reuse does not change how client IPs are logged in AWS WAF. Option D is unnecessary and unsupported as a requirement because CloudFront already provides the required information through standard headers.
AWS documentation consistently states thatX-Forwarded-Foris the correct and supported mechanism for tracing client IPs in CloudFront-protected applications.
* AWS Certified Security - Specialty Official Study Guide
* AWS WAF Developer Guide - Logging
* Amazon CloudFront Developer Guide - Request Headers

NEW QUESTION: 59
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 (LEAVE A REPLY)
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

NEW QUESTION: 60
A company is developing an application that runs across a combination of Amazon EC2 On-Demand Instances and Spot Instances. A security engineer needs to provide a logging solution that makes logs for all instances available from a single location. The solution must allow only a specific set of users to analyze the logs for event patterns. The users must be able to use SQL queries on the logs to perform root cause analysis.
Which solution will meet these requirements?
A. Configure the EC2 instances to send application logs to a single Amazon CloudWatch Logs log group.
Allow only specific users to access the log group. Use CloudWatch Logs Insights to query the log group.

B. Configure the EC2 instances to send application logs to a single Amazon S3 bucket. Allow only specific users to access the S3 bucket. Use Amazon CloudWatch Logs Insights to query the log files in the S3 bucket.
C. Configure each EC2 instance to send its application logs to its own specific Amazon CloudWatch Logs log group. Allow only specific users to access the log groups. Use Amazon Athena to query all the log groups.
D. Configure the EC2 instances to send application logs to a single Amazon CloudWatch Logs log group.Grant Amazon Detective access to the log group. Allow only specific users to use Detective to analyze the logs.
Answer: A (LEAVE A REPLY)
Amazon CloudWatch Logs provides a centralized, scalable service for collecting and storing logs from Amazon EC2 instances, regardless of whether the instances are On-Demand or Spot Instances. According to the AWS Certified Security - Specialty Official Study Guide, CloudWatch Logs is therecommended service for centralized log aggregation and near-real-time analysisof application and system logs.
By configuring all EC2 instances to send logs to asingle CloudWatch Logs log group, the security engineer ensures that logs from all instances are available in one centralized location. Access to the log group can be restricted by using IAM policies, ensuring that only authorized users can view and analyze the logs.
CloudWatch Logs Insights provides apowerful query language with SQL-like syntax, enabling users to search, filter, aggregate, and analyze log data efficiently. This directly satisfies the requirement for SQL-style queries to identify event patterns and perform root cause analysis without requiring data movement or additional services.
Option B is incorrect because CloudWatch Logs Insights cannot query log files stored in Amazon S3. Option C is inefficient and operationally complex, as Athena cannot directly query CloudWatch Logs log groups.
Option D is invalid because Amazon Detective is designed for security investigations using GuardDuty findings, not for general application log analysis.
AWS documentation explicitly states thatCloudWatch Logs combined with CloudWatch Logs Insightsis the most efficient and secure approach for centralized log analysis in EC2-based architectures.
* AWS Certified Security - Specialty Official Study Guide
* Amazon CloudWatch Logs Documentation
* CloudWatch Logs Insights Query Guide

NEW QUESTION: 61
A company is running an application on Amazon EC2 instances in an Auto Scaling group. The application stores logs locally. A security engineer noticed that logs were lost after a scale-in event. The security engineer needs to recommend a solution to ensure the durability and availability of log data. All logs must be kept for a minimum of 1 year for auditing purposes. What should the security engineer recommend?

A. Within the Auto Scaling lifecycle, add a hook to create and attach an Amazon Elastic Block Store (Amazon EBS) log volume each time an EC2 instance is created. When the instance is terminated, the EBS volume can be reattached to another instance for log review.
B. Create an Amazon Elastic File System (Amazon EFS) file system and add a command in the user data section of the Auto Scaling launch template to mount the EFS file system during EC2 instance creation.
Configure a process on the instance to copy the logs once a day from an instance Amazon Elastic Block Store (Amazon EBS) volume to a directory in the EFS file system.
C. Add an Amazon CloudWatch agent into the AMI used in the Auto Scaling group. Configure the CloudWatch agent to send the logs to Amazon CloudWatch Logs for review.
D. Within the Auto Scaling lifecycle, add a lifecycle hook at the terminating state transition and alert the engineering team by using a lifecycle notification to Amazon Simple Notification Service (Amazon SNS). Configure the hook to remain in the Terminating:Wait state for 1 hour to allow manual review of the security logs prior to instance termination.
Answer: (SHOW ANSWER)
In an Auto Scaling group, instances are ephemeral-local disks and instance-level log files can disappear during scale-in or replacement. The most durable, operationally simple pattern is tostream logs off-host continuouslyto a managed log service. Installing and configuring theCloudWatch agent(or unified agent) to ship application logs toAmazon CloudWatch Logsensures logs are centralized and remain available regardless of instance lifecycle events. This directly solves the "logs lost after scale-in" problem and provides high availability for audit and investigation.
CloudWatch Logs also supports retention controls. The security engineer can set the log group retention toat least 1 year(or longer), meeting the audit requirement without building custom storage workflows. Access can be controlled with IAM to restrict who can view or export logs, and CloudWatch logs can be further integrated with Athena/OpenSearch/SIEM tools if needed.
Option A adds complexity and still ties durability to managing volumes across instance churn, with operational risk and scaling challenges. Option B requires daily copy jobs and can still lose logs between copy intervals; it also adds shared filesystem management overhead. Option D is manual and does not ensure durability, and it introduces operational friction during scale-in.
Therefore, centralized log shipping to CloudWatch Logs is the best recommendation.

NEW QUESTION: 62
A company begins to use AWS WAF after experiencing an increase in traffic to the company's public web applications. A security engineer needs to determine if the increase in traffic is because of application-layer attacks. The security engineer needs a solution to analyze AWS WAF traffic.
Which solution will meet this requirement?
A. Send AWS WAF logs to AWS CloudTrail and analyze them with OpenSearch.
B. Send AWS WAF logs to Amazon S3 and query them directly with OpenSearch.
C. Send AWS WAF logs to Amazon S3. Create an Amazon Athena table with partition projection. Use Athena to query the logs.
D. Send AWS WAF logs to AWS CloudTrail and analyze them with Amazon Athena.
Answer: C (LEAVE A REPLY)
AWS WAF supports logging of detailed HTTP request information, including source IP addresses, request URIs, headers, and rule evaluation results. According to the AWS Certified Security - Specialty documentation,Amazon S3 combined with Amazon Athenais the recommended and most cost-effective solution for ad hoc and forensic analysis of AWS WAF logs.
By configuring AWS WAF to deliver logs to Amazon S3 and usingAthena with partition projection, the security engineer can efficiently query large volumes of log data without maintaining partitions manually.
This enables rapid identification of application-layer attacks such as SQL injection, cross-site scripting, and bot activity.
Options A and D are incorrect because AWS WAF logs are not delivered to CloudTrail. Option B is invalid because OpenSearch cannot directly query data stored in S3 without ingestion or additional tooling.
AWS documentation highlightsS3 + Athenaas a best practice for scalable, serverless analysis of AWS WAF logs.
* AWS Certified Security - Specialty Official Study Guide
* AWS WAF Logging Documentation
* Amazon Athena Best Practices

NEW QUESTION: 63
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 the following key policy to allow the infrastructure deployment team to create encrypted Amazon Elastic Block Store (Amazon EBS) volumes by assuming the InfrastructureDeployment IAM role:
{
" Version " : " 2012-10-17 " ,
" Id " : " key-policy-ebs " ,
" Statement " : [
{
" Sid " : " Enable IAM User Permissions " ,

" Effect " : " Allow " ,
" Principal " : {
" AWS " : " arn:aws:iam::123456789012:root "
},
" Action " : " kms:* " ,
" Resource " : " * "
},
{
" Sid " : " Allow use of the key " ,
" Effect " : " Allow " ,
" Principal " : {
" AWS " : " arn:aws:iam::123456789012:role/aws-reserved/sso.amazonaws.com/InfrastructureDeployment "
},
" Action " : [
" kms:Encrypt " ,
" kms:Decrypt " ,
" kms:ReEncrypt* " ,
" kms:GenerateDataKey* " ,
" kms:DescribeKey " ,
" kms:CreateGrant " ,
" kms:ListGrants " ,
" kms:RevokeGrant "
],
" Resource " : " * " ,
" Condition " : {
" StringEquals " : {
" kms:ViaService " : " ec2.us-west-2.amazonaws.com "
}
}
}
]
}

The security engineer recently discovered that IAM rolesother thanthe 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 theConditionblock, 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 theConditionblock, 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: A (LEAVE A REPLY)
AWS KMS key policies can restrict how and when a key is used by applyingconditions such as kms:
ViaService, which limits usage to requests that originate from a specific AWS service. According to the AWS Certified Security - Specialty Official Study Guide and AWS KMS documentation, the kms:ViaService condition is evaluated against the service that calls KMS on behalf of the principal.
Using StringEquals with kms:ViaService restricts usage toexactly one service endpoint. However, AWS services can invoke KMS throughservice variants, internal endpoints, or additional service integrations. When StringEquals is used, these variations can unintentionally bypass the condition, allowing the key to be used by other services through different internal service paths.
Changing the condition operator from StringEquals to StringLike ensures thatonly EC2-related service callsthat match the intended service pattern are allowed, while still preventing use by unrelated AWS services.
This aligns with AWS guidance to use StringLike when service invocation patterns may vary.
Option B is incorrect because the root principal statement is required to retain administrative control over the key. Option C is invalid because changing Regions does not address unauthorized service usage. Option D does not restrict key usage and does not mitigate the issue.
AWS documentation explicitly recommendstightening condition operatorsin KMS key policies to prevent unintended service access while maintaining required functionality.
* AWS Certified Security - Specialty Official Study Guide
* AWS Key Management Service Developer Guide
* AWS KMS Key Policy Best Practices

NEW QUESTION: 64
A company runs several applications on Amazon Elastic Kubernetes Service (Amazon EKS). The company needs a solution to detect any Kubernetes security risks by monitoring Amazon EKS audit logs in addition to operating system, networking, and file events. The solution must send email alerts for any identified risks to a mailing list that is associated with a security team.
Which solution will meet these requirements?
A. Deploy AWS Security Hub and enable security standards that contain EKS controls. Create an Amazon Simple Notification Service (Amazon SNS) topic and set the security team ' s mailing list as a subscriber. Use an Amazon EventBridge rule to send relevant Security Hub events to the SNS topic.

B. Enable Amazon Inspector container image scanning. Configure Amazon Detective to analyze EKS security logs. Create Amazon CloudWatch log groups for EKS audit logs. Use an AWS Lambda function to process the logs and to send email alerts to the security team.
C. Enable Amazon GuardDuty. Enable EKS Protection and Runtime Monitoring for Amazon EKS in GuardDuty. Create an Amazon Simple Notification Service (Amazon SNS) topic and set the security team ' s mailing list as a subscriber. Use an Amazon EventBridge rule to send relevant GuardDuty events to the SNS topic.
D. Install the AWS Systems Manager Agent (SSM Agent) on all EKS nodes. Configure Amazon CloudWatch Logs to collect EKS audit logs. Create an Amazon Simple Notification Service (Amazon SNS) topic and set the security team ' s mailing list as a subscriber. Configure a CloudWatch alarm to publish a message to the SNS topic when new audit logs are generated.
Answer: C (LEAVE A REPLY)
Option C best meets the requirements because Amazon GuardDuty provides Kubernetes-focused threat detection for Amazon EKS by analyzingEKS control plane audit logs(EKS Protection) and combining that signal withruntime telemetryfrom the worker nodes (Runtime Monitoring). EKS audit logs capture Kubernetes API activity and authorization decisions, allowing GuardDuty to detect suspicious cluster actions such as unusual API calls, unexpected access patterns, or indicators of compromise within the cluster.
Runtime Monitoring extends coverage tooperating system/process activity, network connections, and file activityon the nodes, which directly aligns with the need to monitor OS, networking, and file events in addition to audit logs.
For notifications, GuardDuty generatesfindingsthat can be delivered throughAmazon EventBridgerules.
EventBridge can route relevant GuardDuty findings to anAmazon SNS topic, and SNS can sendemail alertsto the security team by subscribing the team's mailing list to the topic. This approach is fully managed, near real time, and avoids building custom log-parsing pipelines while still providing actionable alerts based on GuardDuty's curated EKS threat detections.

NEW QUESTION: 65
A company has contracted with a third party to audit several AWS accounts. To enable the audit, cross- account IAM roles have been created in each account targeted for audit. The auditor is having trouble accessing some of the accounts.
Which of the following may be causing this problem? (Select THREE.)
A. The external ID used by the auditor is missing or incorrect.
B. The auditor is using the incorrect password.
C. The auditor has not been grantedsts:AssumeRolefor the role in the destination account.
D. The Amazon EC2 role used by the auditor must be set to the destination account role.
E. The secret key used by the auditor is missing or incorrect.
F. The role ARN used by the auditor is missing or incorrect.
Answer: A,C,F (LEAVE A REPLY)

Cross-account access for an external auditor typically usesSTS AssumeRoleinto a role that exists in each target account. Three common failure points are: (1) therole ARNbeing wrong or missing, (2) the auditor lacking permission to callsts:AssumeRoleon that role, and (3) anexternal IDmismatch when the role trust policy requires it. Theexternal IDis a best practice for third-party access to mitigate the confused-deputy problem; if the trust policy includes an sts:ExternalId condition, the auditor must supply the exact expected value. If it's absent or incorrect, AssumeRole will be denied.
The auditor must also be authorized on their side (IAM user/role policy) to call sts:AssumeRole (permission), and the destination role'strust policymust trust the auditor principal. If either side is misconfigured, the assume operation fails. Finally, the auditor must reference the correctrole ARNfor each account; using an incorrect ARN (wrong account ID/role name/path) is a frequent reason only "some accounts" fail.
Incorrect password (B) is irrelevant if the auditor is using API/STS federation rather than console password auth, and EC2 instance roles (D) are not required for an external auditor.
Missing/incorrect secret key (E) is possible for API auth, but the question centers on cross-account role access issues; the canonical causes are external ID, AssumeRole permission, and role ARN.

NEW QUESTION: 66
A security engineer wants to evaluate configuration changes to a specific AWS resource to ensure that the resource meets compliance standards. However, the security engineer is concerned about a situation in which several configuration changes are made to the resource in quick succession. The security engineer wants to record only the latest configuration of that resource to indicate the cumulative impact of the set of changes.
Which solution will meet this requirement in the MOST operationally efficient way?
A. Use AWS CloudTrail to detect the configuration changes by filtering API calls to monitor the changes.
Use the most recent API call to indicate the cumulative impact of multiple calls.
B. Use AWS Config to detect the configuration changes and to record the latest configuration in case of multiple configuration changes.
C. Use Amazon CloudWatch to detect the configuration changes by filtering API calls to monitor the changes. Use the most recent API call to indicate the cumulative impact of multiple calls.
D. Use AWS Cloud Map to detect the configuration changes. Generate a report of configuration changes from AWS Cloud Map to track the latest state by using a sliding time window.
Answer: (SHOW ANSWER)
AWS Config is purpose-built torecord resource configuration stateand track how that state changes over time.
When multiple updates occur close together, AWS Config captures configuration items that represent the resource'scurrent recorded configuration, which is exactly what a compliance engineer needs to evaluate the final (cumulative) state after a burst of changes. AWS Config also integrates directly with compliance evaluation throughConfig rules, which can continuously

assess whether the latest configuration meets required standards. This is the most operationally efficient approach because it avoids building custom log filtering, state reconstruction, or timing logic.
CloudTrail and CloudWatch-based approaches (Options A and C) captureAPI events, not authoritative
"current configuration state." Reconstructing the final configuration from a series of API calls can be error- prone, especially when changes are made by different services, via consoles, or through chained automation.
CloudTrail is excellent for "who did what and when," but AWS Config is the service that maintains thelatest configuration snapshotsuitable for compliance posture evaluation.
AWS Cloud Map (Option D) is for service discovery and naming, not compliance configuration history.
Therefore, AWS Config is the correct and most efficient solution.

NEW QUESTION: 67
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 (LEAVE A REPLY)
AmazonMacieis the AWS service purpose-built todiscover and classify sensitive data in S3(PII, financial data, credentials, etc.) and produce findings that can be aggregated centrally. In a multi-account organization, the recommended centralized model is to designate adelegated administrator accountfor Macie so the security team can manage discovery across member accounts from one place.
To make the findings visible from a single location and integrate them with broader security visibility,AWS Security Hubprovides centralized aggregation of security findings across accounts

and services. By also configuring the security account as thedelegated administrator for Security Hub, the company can aggregate findings across the organization. Macie integrates with Security Hub so that sensitive data discovery findings flow into Security Hub's centralized view, giving the security team a single console and API surface to build an "inventory" of sensitive data locations and severity.
Inspector (options B and C) is focused on vulnerability management (EC2, ECR, and related scanning use cases), not sensitive data classification in S3. Trusted Advisor is not the primary destination for sensitive data discovery findings at organizational scale. Therefore, Macie + Security Hub with delegated administration in the security account is the correct solution.

NEW QUESTION: 68
A company runs its microservices architecture in Kubernetes containers on AWS by using Amazon Elastic Kubernetes Service (Amazon EKS) and Amazon Aurora. The company has an organization in AWS Organizations to manage hundreds of AWS accounts that host different microservices.
The company needs to implement a monitoring solution for logs from all AWS resources across all accounts.
The solution must include automatic detection of security-related issues.
Which solution will meet these requirements with theLEAST operational effort?
A. Designate an Amazon GuardDuty administrator account in the organization's management account.
Enable GuardDuty for all accounts. Enable EKS Protection and RDS Protection in the GuardDuty administrator account.
B. Designate a monitoring account. Share Amazon CloudWatch Logs from all accounts. Use Amazon Inspector to evaluate the logs.
C. Centralize CloudTrail logs in Amazon S3 and analyze them with Amazon Athena.
D. Stream CloudWatch Logs to Amazon Kinesis and analyze them with custom AWS Lambda functions.
Answer: (SHOW ANSWER)
Amazon GuardDuty is afully managed, organization-aware threat detection servicethat continuously analyzes AWS logs such as CloudTrail events, VPC Flow Logs, DNS logs, EKS audit logs, and RDS activity.
According to the AWS Certified Security - Specialty Official Study Guide, GuardDuty is designed to operate atscale across AWS Organizations with minimal operational overhead.
By designating a GuardDuty administrator account in the organization's management account and enabling GuardDuty organization-wide, the company can automatically enable threat detection across hundreds of AWS accounts. EnablingEKS Protectionallows GuardDuty to analyze Kubernetes audit logs for suspicious activity, whileRDS Protectionprovides anomaly detection for Amazon Aurora databases.
Options B, C, and D require custom log aggregation, processing, and analytics pipelines, which significantly increase operational effort and maintenance complexity. Amazon Inspector does not

analyze logs, Athena- based analysis is manual, and Kinesis plus Lambda requires custom detection logic.
AWS documentation explicitly identifiesGuardDuty with AWS Organizations integrationas the recommended solution for centralized, automated threat detection across multi-account environments with minimal operational effort.
* AWS Certified Security - Specialty Official Study Guide
* Amazon GuardDuty User Guide
* GuardDuty Organization Administration Documentation

NEW QUESTION: 69
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 (LEAVE A REPLY)
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

NEW QUESTION: 70
A company has installed a third-party application that is distributed on several Amazon EC2 instances and on- premises servers. Occasionally, the company ' s IT team needs to use SSH to connect to each machine to perform software maintenance tasks. Outside these time slots, the machines must be completely isolated from the rest of the network. The company does not want to maintain any SSH keys. Additionally, the company wants to pay only for machine hours when there is an SSH connection.

Which solution will meet these requirements?

A. Create a bastion host with port forwarding to connect to the machines.

B. Set up AWS Systems Manager Session Manager to allow temporary connections.

C. Use AWS CloudShell to create serverless connections.

D. Set up an interface VPC endpoint for each machine for private connection.

Answer: (SHOW ANSWER)

AWS Systems ManagerSession Managerprovides interactive shell access to managed instanceswithout inbound SSH,without bastion hosts, andwithout managing SSH keys. Access is controlled through IAM policies, and every session can be logged to CloudWatch Logs/S3 for auditability. This directly satisfies the

"no SSH keys" requirement and reduces the network exposure surface because you can keep port 22 closed and still obtain shell access when needed.

To meet the isolation requirement, the instances can be placed in private subnets with no inbound access, and you can use Systems Manager connectivity (via SSM endpoints/agents) for administrative sessions only when required. On-premises servers can also be managed by Systems Manager by registering them as managed instances (hybrid activations), allowing the same no-SSH-key operational model across EC2 and on-prem environments.

Options A and D still require network paths and do not eliminate key management; a bastion host is additional infrastructure that must be secured and maintained. CloudShell (Option C) is an AWS-managed shell environment but does not provide a direct, managed, keyless session channel into arbitrary EC2/on-prem hosts by itself. Therefore, Session Manager is the best solution.

NEW QUESTION: 71

A company has a compliance requirement to encrypt all data in transit. The company recently discovered an Amazon Aurora cluster that does not meet this requirement.

How can the company enforce encryption for all connections to the Aurora cluster?

A. In the Aurora cluster configuration, set therequire_secure_transportDB cluster parameter toON.

B. Use AWS Directory Service for Microsoft Active Directory to create a user directory and to enforce Kerberos authentication with Aurora.

C. Configure the Aurora cluster to use AWS Certificate Manager (ACM) to provide encryption certificates.

D. Create an Amazon RDS proxy. Connect the proxy to the Aurora cluster to enable encryption.

Answer: (SHOW ANSWER)

To enforce encryption in transit to Aurora, you must require clients to useTLS/SSLwhen connecting to the database. Aurora (depending on engine flavor) supports a parameter that enforces secure transport by rejecting non-TLS connections. Setting the DB (cluster/instance) parameter such asrequire_secure_transporttoONforces clients to negotiate SSL/TLS; otherwise the server refuses the connection, ensuring all data in transit is encrypted. This directly enforces the compliance requirement at the database endpoint itself, independent of client behavior.

Kerberos authentication (Option B) provides centralized authentication and can be useful for IAM/AD integration, but it does not by itself guarantee that the network session is encrypted. Option C is incorrect because Aurora/RDS usesRDS-provided certificatesfor TLS; you don't attach ACM certificates to Aurora the way you would for ALB/CloudFront. Option D is not the right enforcement mechanism: RDS Proxy can help with connection pooling and IAM auth patterns, but it does not inherently force all client-to-proxy or proxy-to- db connections to be encrypted in the way a DB parameter enforcement does (and you'd still need to ensure clients use TLS). Therefore, enabling the parameter that requires secure transport is the correct solution.

NEW QUESTION: 72
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 (LEAVE A REPLY)
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

NEW QUESTION: 73
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 (LEAVE A REPLY)
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

NEW QUESTION: 74
A company runs an application on a fleet of Amazon EC2 instances. The application is accessible to users around the world. The company associates an AWS WAF web ACL with an Application Load Balancer (ALB) that routes traffic to the EC2 instances.
A security engineer is investigating a sudden increase in traffic to the application. The security engineer discovers a significant amount of potentially malicious requests coming from hundreds of IP addresses in two countries. The security engineer wants to quickly limit the potentially malicious requests but does not want to prevent legitimate users from accessing the application.
Which solution will meet these requirements?
A. Use AWS WAF to implement a rate-based rule for all incoming requests.

B. Use AWS WAF to implement a geographical match rule to block all incoming traffic from the two countries.
C. Edit the ALB security group to include a geographical match rule to block all incoming traffic from the two countries.
D. Add deny rules to the ALB security group that prohibit incoming requests from the IP addresses.
Answer: A (LEAVE A REPLY)
AWS WAFrate-based rulesare specifically designed to protect applications from traffic floods and distributed attacks that originate from large numbers of IP addresses. According to the AWS Certified Security - Specialty Official Study Guide, rate-based rules automatically track the number of requests coming from individual IP addresses and temporarily block IPs that exceed a defined threshold.
In this scenario, the malicious traffic originates fromhundreds of IP addresses across two countries, mixed with legitimate user traffic. A rate-based rule allows the security engineer tolimit excessive request rates without fully blocking access from entire geographic regions, ensuring that legitimate users can still access the application.
Option B is incorrect because geographic match rules blockalltraffic from selected countries, which would deny access to legitimate users and violate the stated requirement. Option C is invalid because security groups do not support geographic filtering. Option D is not scalable, as manually blocking hundreds of IP addresses is operationally inefficient and ineffective against rapidly changing attacker IPs.
AWS documentation emphasizes thatrate-based rules are the recommended first-line mitigationfor sudden traffic spikes and potential application-layer DDoS attacks when business continuity must be preserved.
* AWS Certified Security - Specialty Official Study Guide
* AWS WAF Developer Guide - Rate-Based Rules
* AWS DDoS Resiliency Best Practices

NEW QUESTION: 75
A company operates an Amazon EC2 instance that is registered as a target of a Network Load Balancer (NLB). The NLB is associated with a security group. The security group allows inbound TCP traffic on port
22 from 10.0.0.0/23.
The company maps the NLB to two subnets that share the same network ACL and route table. The route table has a route for 0.0.0.0/0 to an internet gateway. The network ACL has one inbound rule that has a priority of
20 and that allows TCP traffic on port 22 from 10.0.0.0/16.
A security engineer receives an alert that there is an unauthorized SSH session on the EC2 instance. The unauthorized session originates from 10.0.1.5. The company's incident response procedure requires unauthorized SSH sessions to beimmediately interrupted. The instance must remain running, and its memory must remain intact.

Which solution will meet these requirements?

A. Restart the EC2 instance from either the AWS Management Console or the AWS CLI.

B. Add a new inbound rule that has a priority of 10 to the network ACL to deny TCP traffic on port 22 from 10.0.1.5.

C. Remove the security group rule that allows inbound TCP traffic on port 22 from 10.0.0.0/16.

D. Update the route table to remove the route to the internet gateway.

Answer: (SHOW ANSWER)

Network ACLs arestatelessand are evaluated in order based on rule number, with lower rule numbers taking precedence. According to AWS Certified Security - Specialty incident response guidance, network ACLs can be used toimmediately block traffic at the subnet levelwithout restarting instances or modifying their runtime state.

By adding a deny rule with alower priority number (10)that explicitly denies TCP traffic on port 22 from the offending IP address (10.0.1.5), the unauthorized SSH session is immediately interrupted. This approach satisfies the requirement to keep the instance running and to preserve memory for forensic analysis.

Option A violates the requirement because restarting the instance clears memory. Option C would disrupt all legitimate SSH access, not just the unauthorized session. Option D would block all internet access and could cause widespread service disruption.

AWS documentation emphasizes usingnetwork ACL deny rules for rapid, targeted containmentwhen immediate interruption is required without altering instance state.

* AWS Certified Security - Specialty Official Study Guide
* Amazon VPC Network ACL Documentation
* AWS Incident Response Best Practices