PassQueen
a.k.a dumpsqueen.com

AWS Certified Security - Specialty

Amazon AWS SCS-C02

Version Demo

Total Demo Questions: 78

Total Premium Questions: 789

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Threat Detection and Incident Response	87
Topic 2, Security Logging and Monitoring	132
Topic 3, Infrastructure Security	172
Topic 4, Identity and Access Management	165
Topic 5, Data Protection	167
Topic 6, Management and Security Governance	58
Topic 7, Mix Questions	8
Total	789

QUESTION NO: 1

A corporation is preparing to acquire several companies. A Security Engineer must design a solution to ensure that newly acquired IAM accounts follow the corporation's security best practices. The solution should monitor each Amazon S3 bucket for unrestricted public write access and use IAM managed services.

What should the Security Engineer do to meet these requirements?

A. Configure Amazon Macie to continuously check the configuration of all S3 buckets.

B. Enable AWS Config to check the configuration of each S3 bucket.

C. Set up AWS Systems Manager to monitor S3 bucket policies for public write access.

D. Configure an Amazon EC2 instance to have an IAM role and a cron job that checks the status of all S3 buckets.

ANSWER: B

Explanation:

Enable AWS Config to check the configuration of each S3 bucket. AWS Config is an AWS managed service that continuously records resource configurations and evaluates them against compliance rules. For this requirement, AWS Config can use the AWS managed `s3-bucket-public-write-prohibited` rule to identify buckets that allow public write access through either bucket ACLs or bucket policies. This supplies ongoing rather than point-in-time monitoring, along with configuration history and compliance results that a central security team can review.

As companies are acquired, the corporation can enable AWS Config in the acquired AWS accounts and relevant Regions, then centrally consolidate compliance visibility by using AWS Config aggregators. AWS Organizations can help govern and deploy this baseline consistently across member accounts, including new accounts added to the organization. The approach requires no customer-managed polling infrastructure and directly uses AWS managed configuration monitoring and managed rules, satisfying both the security-control and operational-management requirements. AWS Config's managed-rule documentation describes `s3-bucket-public-write-prohibited` and its evaluation of public write permissions: AWS Config managed rule: s3-bucket-public-write-prohibited. General service capabilities are documented in What is AWS Config?.

QUESTION NO: 2

A security engineer creates an Amazon S3 bucket policy that denies access to all users. A few days later, the security engineer adds an additional statement to the bucket policy to allow read-only access to one other employee. Even after updating the policy, the employee still receives an access denied message.

What is the likely cause of this access denial?

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

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

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

A. Deploy a NAT gateway in a public subnet in each Availability Zone that is in use.

B. Place the DB instance in a public subnet.

C. Place the DB instance in a private subnet.

D. Configure the Auto Scaling group to place the EC2 instances in a public subnet.

E. Configure the Auto Scaling group to place the EC2 instances in a private subnet.

F. Deploy the ALB in a private subnet.

ANSWER: A C E

Explanation:

Deploying a NAT gateway in a public subnet in each Availability Zone that is in use gives workloads in private subnets controlled outbound internet access. Associating an Elastic IP address with each NAT gateway provides stable public source IP addresses that the external payment provider can add to its allow list. Using one NAT gateway per Availability Zone also supports resilient, zone-local egress as the Auto Scaling group grows or operates across multiple Availability Zones. The private-subnet route tables should send internet-bound traffic to the NAT gateway in the same Availability Zone.

Placing the DB instance in a private subnet keeps the database off the public internet. The application instances can connect to the database through private VPC addressing, with security groups limiting database access to the application tier. This supports network segmentation for the data tier while retaining the required application-to-database path.

Configuring the Auto Scaling group to place the EC2 instances in a private subnet prevents direct internet reachability to individual application instances. The internet-facing ALB can accept the required HTTP and HTTPS requests and forward them to the private targets. This design permits horizontal scaling without exposing changing instance addresses or requiring the payment provider to update its source-IP allow list. AWS documents this NAT gateway routing pattern in NAT gateway basics and describes private DB subnet deployment in Amazon RDS in a VPC.

QUESTION NO: 3

A company is building a data processing application mat uses AWS Lambda functions. The application's

Lambda functions need to communicate with an Amazon RDS OB instance that is deployed within a VPC in the same AWS account

Which solution meets these requirements in the MOST secure way?

A. Configure the DB instance to allow public access Update the DB instance security group to allow access from the Lambda public address space for the AWS Region

B. Deploy the Lambda functions inside the VPC Attach a network ACL to the Lambda subnet Provide outbound rule access to the VPC CIDR range only Update the DB instance security group to allow traffic from 0.0.0.0/0

C. Deploy the Lambda functions inside the VPC Attach a security group to the Lambda functions Provide outbound rule access to the VPC CIDR range only Update the DB instance security group to allow traffic from the Lambda security group

D. Peer the Lambda default VPC with the VPC that hosts the DB instance to allow direct network access without the need for security groups

ANSWER: C

Explanation:

Deploying the Lambda functions inside the VPC, associating a dedicated security group with those functions, restricting egress, and allowing the RDS DB instance to accept traffic from that Lambda security group provides private, least-privilege connectivity. When Lambda is configured with VPC subnets and security groups, Lambda creates and manages Hyperplane ENIs that allow the function to access resources reachable within that VPC, including a non-public RDS DB instance. The DB security group can use the Lambda security group as its inbound source, which authorizes only workloads associated with that group rather than relying on broad, change-prone IP address ranges. Security groups are stateful, so return traffic for an allowed database connection is automatically permitted. Limiting the Lambda security group's outbound access to the VPC CIDR range reduces unnecessary network reachability; in a production implementation, the rule can be narrowed further to the database listener port and destination as appropriate. This design keeps database traffic on private VPC networking and applies security controls at both the Lambda and RDS boundaries. AWS documents VPC configuration for Lambda at Configuring a Lambda function to access resources in a VPC and security-group source referencing at Security group rules.

QUESTION NO: 4

A company's engineering team is developing a new application that creates IAM Key Management Service (IAM KMS) CMK grants for users immediately after a grant IS created users must be able to use the CMK tu encrypt a 512-byte payload. During load testing, a bug appears |intermittently where AccessDeniedExceptions are occasionally triggered when a userï¬□rst attempts to encrypt using the CMK

Which solution should the c0mpanyâ€˜s security specialist recommendâ€˜?

A. Instruct users to implement a retry mechanism every 2 minutes until the call succeeds.

B. Instruct the engineering team to consume a random grant token from users, and to call the CreateGrant operation, passing it the grant token. Instruct use to use that grant token in their call to encrypt.

C. Instruct the engineering team to create a random name for the grant when calling the CreateGrant operation. Return the name to the users and instruct them to provide the name asthe grant token in the call to encrypt.

D. Instruct the engineering team to pass the grant token returned in the CreateGrant response to users. Instruct users to use that grant token in their call to encrypt.

ANSWER: D

Explanation:

Instruct the engineering team to pass the grant token returned in the CreateGrant response to users. Instruct users to use that grant token in their call to encrypt. AWS KMS grant creation is eventually consistent. Therefore, immediately after `CreateGrant` returns, the new grant might not yet be visible to all AWS KMS authorization systems. A principal that relies on the newly created grant can intermittently receive an `AccessDeniedException` if it calls `Encrypt` before grant propagation completes. The `CreateGrant` response includes a grant token specifically to address this situation. Supplying that token in the `GrantTokens` parameter of a subsequent cryptographic request tells AWS KMS to use the new grant immediately, without waiting for eventual consistency. The engineering service should securely provide the returned token to the authorized user or workload, and that caller should include it in its initial `Encrypt` request. This provides reliable immediate access while retaining the least-privilege permissions defined by the grant. AWS documents grant tokens as the mechanism for using new grants immediately after creation and lists `Encrypt` among the KMS operations that accept them. See Using grant tokens and the CreateGrant API reference.

QUESTION NO: 5

A company is developing a highly resilient application to be hosted on multiple Amazon EC2 instances . The application will store highly sensitive user data in Amazon RDS tables

The application must

• Include migration to a different IAM Region in the application disaster recovery plan.

• Provide a full audit trail of encryption key administration events

• Allow only company administrators to administer keys.

• Protect data at rest using application layer encryption

A Security Engineer is evaluating options for encryption key management

Why should the Security Engineer choose IAM CloudHSM over IAM KMS for encryption key management in this situation?

A. The key administration event logging generated by CloudHSM is significantly moreextensive than IAM KMS.

B. CloudHSM provides company administrators with exclusive control to administer encryption keys in dedicated HSMs.

C. The ciphertext produced by CloudHSM provides more robust protection against brute force decryption attacks than the ciphertext produced by IAM KMS

D. CloudHSM provides the ability to copy keys to a different Region, whereas IAM KMS does not

ANSWER: B

Explanation:

CloudHSM lets company administrators exclusively control cryptographic key administration within dedicated, single-tenant hardware security modules. The customer creates and manages the CloudHSM users, including Crypto Officers, who administer HSM users and keys. AWS manages the underlying service and hardware availability, but does not have access to the customer's HSM credentials or plaintext keys. This ownership model is appropriate when the organization requires direct, exclusive administrative control over encryption keys used by its application-layer encryption implementation.

CloudHSM also supports separation of duties through distinct HSM user roles and produces HSM audit logs that record HSM activity. Those logs can be retained and monitored as part of the organization's audit process, while the application can use standard cryptographic interfaces to perform encryption operations with keys held in the HSM. The organization remains responsible for designing resilient CloudHSM clusters and implementing its regional disaster-recovery process, including the required backup and recovery procedures.

For more detail, see the AWS CloudHSM User Guide introduction and AWS guidance on managing HSM users.

QUESTION NO: 6

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

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

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

A. Deploy a NAT gateway in a public subnet in each Availability Zone that is in use.

B. Place the DB instance in a public subnet.

C. Place the DB instance in a private subnet.

D. Configure the Auto Scaling group to place the EC2 instances in a public subnet.

E. Configure the Auto Scaling group to place the EC2 instances in a private subnet.

F. Deploy the ALB in a private subnet.

ANSWER: A C E

Explanation:

Deploying a NAT gateway in a public subnet in each Availability Zone used by the private application subnets provides controlled, highly available outbound internet access for the application tier. Each NAT gateway uses an Elastic IP address, so the payment provider can allowlist those stable egress addresses. The EC2 instances can therefore scale in or out without changing the source public IP addresses that the provider observes. Private-subnet route tables should direct external IPv4-bound traffic through the NAT gateway in the same Availability Zone, avoiding a single-AZ egress dependency.

Placing the EC2 instances in private subnets ensures that they do not accept direct inbound connections from the internet. The internet-facing Application Load Balancer remains the public entry point for HTTP and HTTPS and forwards approved traffic to the application instances through their security groups. Placing the RDS DB instance in private subnets keeps the database isolated from direct internet access; database security groups can permit connections only from the application tier. Together, these controls expose only the intended web endpoint while providing predictable payment-provider egress and private, layered access to data services. See the Amazon VPC NAT gateway documentation and Amazon RDS VPC documentation.

QUESTION NO: 7

A security engineer is investigating a malware infection that has spread across a set of Amazon EC2 instances. A key indicator of the compromise is outbound traffic on TCP port 2905 to a set of command and control hosts on the internet.

The security engineer creates a network ACL rule that denies the identified outbound traffic. The security engineer applies the network ACL rule to the subnet of the EC2 instances. The security engineer must identify any EC2 instances that are trying to communtcate on TCP port 2905.

Which solution will identify the affected EC2 instances with the LEAST operational effort?

A. Create a Network Access Scope in Amazon VPC Network Access Analyzer. Use the Network Access Scope to identify EC2 instances that try to send traffic to TCP port 2905.

B. Enable VPC flow logs for the VPC where the affected EC2 instances are located Configure the flow logs to capture rejected traffic. In the flow logs, search for REJECT records that have a destination TCP port of 2905.

C. Enable Amazon GuardDuty Create a custom GuardDuty IP list to create a finding when an EC2 instance tries to communicate with one of the command and control hosts. Use Amazon Detective to identify the EC2 instances that initiate the communication.

D. Create a firewall in AWS Network Firewall. Attach the firewall to the subnet of the EC2 instances. Create a custom rule to identify and log traffic from the firewall on TCP port 2905. Create an Amazon CloudWatch Logs metric filter to identify firewall logs that reference traffic on TCP port 2905.

ANSWER: B

Explanation:

Enable VPC flow logs for the VPC where the affected EC2 instances are located Configure the flow logs to capture rejected traffic. In the flow logs, search for REJECT records that have a destination TCP port of 2905. This directly uses the enforcement point already created by the network ACL. Because network ACLs are stateless and evaluate traffic at the subnet boundary, an outbound connection attempt that matches the deny rule is recorded in VPC Flow Logs with an action of `REJECT`. Flow log records include the source private IP address, destination IP address, destination port, protocol, action, and interface ID. The source address and network interface information can be used to map each rejected attempt back to the associated EC2 instance. Configuring the flow log to collect rejected traffic avoids collecting unnecessary accepted-flow records and requires no agents, packet capture appliances, routing changes, or additional inline inspection infrastructure. The engineer can deliver the logs to CloudWatch Logs, Amazon S3, or Amazon Data Firehose and query for TCP destination port 2905 and the `REJECT` action. AWS documents that Flow Logs can be created for a VPC, subnet, or network interface and can capture rejected traffic; see VPC Flow Logs and the Flow Log record examples.

QUESTION NO: 8

A security engineer configures Amazon S3 Cross-Region Replication (CRR) for all objects that are in an S3 bucket in the us-east-1. Region Some objects in this S3 bucket use server-side encryption with AWS KMS keys (SSE-KMS) for encryption at test. The security engineer creates a destination S3 bucket in the us-west-2 Region. The destination S3 bucket is in the same AWS account as the source S3 bucket.

The security engineer also creates a customer managed key in us-west-2 to encrypt objects at rest in the destination S3 bucket. The replication configuration is set to use the key in us-west-2 to encrypt objects in the destination S3 bucket. The security engineer has provided the S3 replication configuration with an IAM role to perform the replication in Amazon S3.

After a day, the security engineer notices that no encrypted objects from the source S3 bucket are replicated to the destination S3 bucket. However, all the unencrypted objects are replicated.

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

A. Change the replication configuration to use the key in us-east-1 to encrypt the objects that are in the destination S3 bucket.

B. Grant the IAM role the kms. Encrypt permission for the key in us-east-1 that encrypts source objects.

C. Grant the IAM role the s3 GetObjectVersionForReplication permission for objects that are in the source S3 bucket.

D. Grant the IAM role the kms. Decrypt permission for the key in us-east-1 that encrypts source objects.

E. Change the key policy of the key in us-east-1 to grant the kms. Decrypt permission to the security engineer's IAM account.

F. Grant the IAM role the kms Encrypt permission for the key in us-west-2 that encrypts objects that are in the destination S3 bucket.

ANSWER: C D F

Explanation:

Amazon S3 replication handles SSE-KMS encrypted objects differently from unencrypted objects because S3 must access the source object version through the replication role and use AWS KMS for cryptographic operations in both Regions. Granting `s3:GetObjectVersionForReplication` for objects in the source S3 bucket allows the replication role to retrieve the specific version of each source object that S3 must replicate. Granting `kms:Decrypt` for the key in us-east-1 that encrypts source objects allows S3, acting through the replication role, to decrypt the source data for replication. Granting `kms:Encrypt` for the key in us-west-2 that encrypts objects in the destination S3 bucket allows S3 to encrypt the replica by using the destination Region customer managed KMS key configured in the replication rule.

In a production policy, the replication role and applicable KMS key policies must also permit the required `kms:GenerateDataKey` operations for the relevant source and destination keys. These permissions are required by the documented SSE-KMS replication workflow and should be scoped to the replication role and the affected KMS keys. See the Amazon S3 guidance for replicating SSE-KMS objects and the replication permission overview.

QUESTION NO: 9

A security engineer is designing an IAM policy to protect AWS API operations. The policy must enforce multi-factor authentication (MFA) for IAM users to access certain services in the AWS production account. Each session must remain valid for only 2 hours. The current version of the IAM policy is as follows:

```
{
    "Version": "2012-10-17",
    "Statement": [{
        "Effect": "Allow",
        "Action": [
            "ec2:DescribeInstances",
            "ec2:StopInstances",
            "ec2:TerminateInstances"
        ],
        "Resource": ["*"]
    }]
}
```

Which combination of conditions must the security engineer add to the IAM policy to meet these

requirements? (Select TWO.)

A. "Bool" : { "aws:MultiFactorAuthPresent": "true" }

B. "Bool" : { "aws:MultiFactorAuthPresent": "false" }

C. "NumericLessThan" : { "aws:MultiFactorAuthAge" : "7200"}

D. "NumericGreaterThan" : { "aws:MultiFactorAuthAge" : "7200" }

E. "NumericLessThan" : { "MaxSessionDuration" : "7200"}

ANSWER: A C

Explanation:

The required conditions are `"Bool": {"aws:MultiFactorAuthPresent": "true"}` and `"NumericLessThan": {"aws:MultiFactorAuthAge": "7200"}`. The `aws:MultiFactorAuthPresent` global condition key evaluates whether the request was signed with temporary credentials that were issued after MFA authentication. Setting this Boolean condition to `true` makes MFA a prerequisite for the protected API calls.

The `aws:MultiFactorAuthAge` global condition key represents the number of seconds since MFA was performed for the temporary session credentials. Applying a numeric less-than comparison against `7200` allows requests only while the MFA age is below 7,200 seconds, which establishes the required two-hour session window. Together, these conditions ensure that access is both MFA-authenticated and recent enough to satisfy the production-account control. These keys are AWS global condition context keys and are evaluated during IAM authorization for requests made with MFA-backed temporary credentials. AWS documents MFA enforcement and MFA age-based controls in the IAM global condition keys reference and its guidance for requiring MFA for IAM API access.

QUESTION NO: 10

A company has an organization with SCPs in AWS Organizations. The root SCP for the organization is as follows:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "AllowsAllActions",
            "Effect": "Allow",
            "Action": "*",
            "Resource": "*"
        },
        {
            "Sid": "DenySES",
            "Effect": "Deny",
            "Action": "ses:*",
            "Resource": "*"
        }
    ]
}
```

The company's developers are members of a group that has an IAM policy that allows access to Amazon Simple Email Service (Amazon SES) by allowing ses:* actions. The account is a child to an OU that has an SCP that allows Amazon SES. The developers are receiving a not-authorized error when they try to access Amazon SES through the AWS Management Console.

Which change must a security engineer implement so that the developers can access Amazon SES?

A. Add a resource policy that allows each member of the group to access Amazon SES.

B. Add a resource policy that allows "Principal": {"AWS": "arn:aws:iam::account-number:group/Dev"}.

C. Remove the AWS Control Tower control (guardrail) that restricts access to Amazon SES.

D. Remove Amazon SES from the root SCP.

ANSWER: D

Explanation:

Remove Amazon SES from the root SCP.
is correct because SCPs define the maximum permissions available to identities in member accounts. An IAM policy granting `ses:*` does not grant usable access unless the action is also permitted by every applicable SCP in the account's inheritance path. SCP permissions are evaluated as an intersection: permissions must remain available at the organization root, each applicable organizational unit, and the account itself, in addition to being granted by IAM policies.

Although the account's OU SCP allows Amazon SES, that allowance cannot restore a service that the root-level SCP excludes or explicitly denies. A root SCP applies to all OUs and member accounts beneath the organization root, so its restriction continues to limit the effective permissions of the developers. Updating the root SCP to remove the Amazon SES restriction allows the OU-level SCP and the developers' IAM group policy to take effect, enabling console access to SES, subject to any remaining applicable IAM and SCP conditions.

AWS documents SCPs as permission guardrails rather than permission grants and explains that an action must be allowed through the relevant hierarchy. See Service control policies (SCPs) and SCP evaluation.

QUESTION NO: 11

A company plans to create individual child accounts within an existing organization in AWS Organizations for each of its DevOps teams. AWS CloudTrail has been enabled and configured on all accounts to write audit logs to an Amazon S3 bucket in a centralized AWS account. A security engineer needs to ensure that DevOps team members are unable to modify or disable this configuration.

How can the security engineer meet these requirements?

A. Create an IAM policy that prohibits changes to the specific CloudTrail trail and apply the policy to the AWS account root user.

B. Create an S3 bucket policy in the specified destination account for the CloudTrail trail that prohibits configuration changes from the AWS account root user in the source account.

C. Create an SCP that prohibits changes to the specific CloudTrail trail and apply the SCP to the appropriate organizational unit or account in Organizations.

D. Create an IAM policy that prohibits changes to the specific CloudTrail trail and apply the policy to a new IAM group. Have team members use individual IAM accounts that are members of the new IAM group.

ANSWER: C

Explanation:

Create an SCP that prohibits changes to the specific CloudTrail trail and apply the SCP to the appropriate organizational unit or account in Organizations. Service control policies (SCPs) are AWS Organizations permission guardrails that establish the maximum permissions available in member accounts. An explicit SCP deny applies even when a DevOps team member has broad IAM permissions, including administrator-level access in their account. This makes an SCP appropriate for enforcing a centrally mandated CloudTrail configuration across independently administered child accounts.

The policy can deny CloudTrail management actions that would stop, alter, or remove the organization's required logging, such as `cloudtrail:StopLogging`, `cloudtrail:DeleteTrail`, and relevant update actions. The SCP should be attached to the OU containing the DevOps accounts or to the individual accounts. Its scope can be constrained to the protected trail ARN and include an exception for an authorized centralized security administration role if that role must maintain the trail. This preserves the ability of the security team to administer logging while preventing DevOps principals from disabling or changing audit collection and delivery. SCPs therefore provide the required centralized, durable control for all affected accounts. See AWS Organizations service control policies and CloudTrail identity-based policy examples.

QUESTION NO: 12

A company in France uses Amazon Cognito with the Cognito Hosted UI as an identity broker for sign-in and sign-up processes. The company is marketing an application and expects that all the application's users will come from France.

When the company launches the application, the company's security team observes fraudulent sign-ups for the application. Most of the fraudulent registrations are from users outside of France.

The security team needs a solution to perform custom validation at sign-up. Based on the results of the validation, the solution must accept or deny the registration request.

Which combination of steps will meet these requirements? (Choose two.)

A. Create a pre sign-up AWS Lambda trigger. Associate the Lambda function with the Amazon Cognito user pool.

B. Use a geographic match rule statement to configure an AWS WAF web ACL. Associate the web ACL with the Amazon Cognito user pool.

C. Configure an app client for the application's Amazon Cognito user pool. Use the app client ID to validate the requests in the hosted UI.

D. Update the application's Amazon Cognito user pool to configure a geographic restriction setting.

E. Use Amazon Cognito to configure a social identity provider (IdP) to validate the requests on the hosted UI.

ANSWER: A B

Explanation:

Creating a pre sign-up AWS Lambda trigger and associating it with the Amazon Cognito user pool provides the required custom validation point during registration. Amazon Cognito invokes this trigger before it creates a local user or confirms a federated user. The Lambda function can inspect sign-up attributes and request context, apply company-specific validation logic, and query an external fraud or risk service if needed. When validation determines that a registration must not proceed, the function can raise an error, which causes Cognito to reject the sign-up request. This directly satisfies the requirement to make an accept-or-deny decision based on custom validation.

Using an AWS WAF web ACL with a geographic match rule and associating the web ACL with the Cognito user pool adds a complementary edge-control layer for the Hosted UI. A geo match rule can match the source country derived from the request IP address, allowing the web ACL to block sign-up traffic from locations outside France. This reduces the volume of fraudulent and automated requests that reach Cognito, while the pre sign-up trigger remains available for finer-grained business and fraud validation. Together, these controls provide geographic filtering and authoritative custom registration validation. See the Amazon Cognito pre sign-up Lambda trigger documentation and AWS WAF protection for Amazon Cognito user pools documentation.

QUESTION NO: 13

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 rote 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 rote to the new IAM user.

C. Use AWS 1AM Identity Center (AWS Single Sign-On) 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: A

Explanation:

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. This approach uses VPC sharing within AWS Organizations so that the dedicated account can deploy CloudHSM client instances into a subnet owned by the central networking account. Because the client instances and the CloudHSM cluster are then connected through the same shared VPC environment, the clients can communicate with the HSM over private network addresses. The CloudHSM security group must explicitly permit the required inbound client connectivity from those client instances; sharing a subnet does not itself grant network access. This design maintains centralized ownership and administration of the VPC and CloudHSM infrastructure while allowing the participating account to consume the centralized cryptographic service. AWS RAM is designed for this type of multi-account resource sharing, and AWS Organizations simplifies sharing with accounts in the organization. For implementation guidance, see the AWS Knowledge Center article on sharing AWS CloudHSM clusters across accounts and the Amazon VPC documentation for VPC sharing.

QUESTION NO: 14

An online media company has an application that customers use to watch events around the world. The application is hosted on a fleet of Amazon EC2 instances that run Amazon Linux 2. The company uses AWS Systems Manager to manage the EC2 instances. The company applies patches and application updates by using the AWS-AmazonLinux2DefaultPatchBaseline patching baseline in Systems Manager Patch Manager.

The company is concerned about potential attacks on the application during the week of an upcoming event. The company needs a solution that can immediately deploy patches to all the EC2 instances in response to a security incident or vulnerability. The solution also must provide centralized evidence that the patches were applied successfully.

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

A. Create a new patching baseline in Patch Manager. Specify Amazon Linux 2 as the product. Specify Security as the classification. Set the automatic approval for patches to 0 days. Ensure that the new patching baseline is the designated default for Amazon Linux 2.

B. Use the Patch Now option with the scan and install operation in the Patch Manager console to apply patches against the baseline to all nodes. Specify an Amazon S3 bucket as the patching log storage option.

C. Use the Clone function of Patch Manager to create a copy of the AWS-AmazonLinux2DefaultPatchBaseline built-in baseline. Set the automatic approval for patches to 1 day.

D. Create a patch policy that patches all managed nodes and sends a patch operation log output to an Amazon S3 bucket. Use a custom scan schedule to set Patch Manager to check every hour for new patches. Assign the baseline to the patch policy.

E. Use Systems Manager Application Manager to inspect the package versions that were installed on the EC2 instances. Additionally, use Application Manager to validate that the patches were correctly installed.

ANSWER: A B

Explanation:

Creating a new Patch Manager baseline for Amazon Linux 2 that includes the *Security* classification and uses a zero-day automatic approval delay makes newly released security patches eligible for installation immediately. Making that custom baseline the default Amazon Linux 2 baseline ensures that Patch Manager uses those approval rules when patch operations are run without a specifically selected alternative baseline. This provides the required readiness for an urgent vulnerability response.

Using Patch Now with the scan and install operation initiates an on-demand patch deployment to the selected managed nodes, rather than waiting for a maintenance-window or recurring schedule. Selecting all relevant EC2 managed nodes allows the organization to rapidly apply every patch currently approved by the baseline across the fleet. Configuring an Amazon S3 bucket for patching log output centralizes command output and patch-operation records, providing durable evidence that can be reviewed after the deployment. AWS documents Patch Now as the on-demand Patch Manager workflow and supports S3 output for patching operations. See AWS Systems Manager Patch Manager documentation and Patch Now documentation.

QUESTION NO: 15

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: B C E

Explanation:

https://d1.awsstatic.com/WWPS/pdf/aws_security_incident_response.pdf

---

A. Enable termination protection for the EC2 instance if termination protection has not been enabled.

B. Take snapshots of the Amazon Elastic Block Store (Amazon EBS) data volumes that are attached to the EC2 instance.

C. Remove all snapshots of the Amazon Elastic Block Store (Amazon EBS) data volumes that are attached to the EC2 instance.

D. Capture the EC2 instance metadata, and then tag the EC2 instance as under quarantine.

E. Immediately remove any entries in the EC2 instance metadata that contain sensitive information.

F. Disable termination protection for the EC2 instance if termination protection has not been disabled.

ANSWER: A B D

Explanation:

Before examining a potentially compromised instance, the response team should preserve the instance and its associated evidence while recording the contextual information needed to conduct a reliable investigation. **Enable termination protection for the EC2 instance if termination protection has not been enabled.** helps prevent an accidental API-based termination from destroying the running system and its evidence during triage. **Take snapshots of the Amazon Elastic Block Store (Amazon EBS) data volumes that are attached to the EC2 instance.** creates point-in-time copies of disk contents, allowing investigators to examine files, logs, and other artifacts from copies while retaining the original state. AWS notes that EBS snapshots are incremental backups and can be used to create new volumes for analysis.

Capture the EC2 instance metadata, and then tag the EC2 instance as under quarantine. records critical investigative context, such as instance identity, networking configuration, attached roles, and launch details, before containment or remediation changes the environment. A quarantine tag also provides an explicit operational marker that can coordinate

responders and trigger approved isolation workflows. Together, these actions support evidence preservation, chain of custody, and controlled incident handling before deeper host-level investigation begins. See the AWS Security Incident Response Guide and the Amazon EBS snapshot documentation.

QUESTION NO: 16

A company hosts an end user application on AWS Currently the company deploys the application on Amazon EC2 instances behind an Elastic Load Balancer The company wants to configure end-to-end encryption between the Elastic Load Balancer and the EC2 instances.

Which solution will meet this requirement with the LEAST operational effort?

A. Use Amazon issued AWS Certificate Manager (ACM) certificates on the EC2 instances and the Elastic Load Balancer to configure end-to-end encryption

B. Import a third-party SSL certificate to AWS Certificate Manager (ACM) Install the third-party certificate on the EC2 instances Associate the ACM imported third-party certificate with the Elastic Load Balancer

C. Deploy AWS CloudHSM Import a third-party certificate Configure the EC2 instances and the Elastic Load Balancer to use the CloudHSM imported certificate

D. Import a third-party certificate bundle to AWS Certificate Manager (ACM) Install the third-party certificate on the EC2 instances Associate the ACM imported third-party certificate with the Elastic Load Balancer.

ANSWER: D

Explanation:

Import a third-party certificate bundle to AWS Certificate Manager (ACM), install the third-party certificate on the EC2 instances, and associate the ACM imported third-party certificate with the Elastic Load Balancer. This approach enables TLS encryption for both connections: from clients to the load balancer and from the load balancer to the EC2 instances. The certificate installed on each instance allows the instances to authenticate and encrypt the HTTPS connection that the load balancer establishes with its targets. Importing the corresponding certificate bundle into ACM makes the certificate available for use by the load balancer listener without requiring the private key to be managed separately on the load balancer.

A certificate bundle includes the certificate, its private key, and, when applicable, the certificate chain. These are the required components when importing a certificate into ACM. This is operationally efficient because the company can use the same externally obtained certificate material for the backend instances and the load balancer, avoiding the need to build and operate dedicated cryptographic hardware or a separate private certificate authority solely for this traffic path. The company remains responsible for renewal and replacement of an imported certificate, but the solution directly supports the required deployment model. See the ACM certificate import documentation and the Elastic Load Balancing HTTPS and SSL configuration documentation.

QUESTION NO: 17

A company deployed Amazon GuardDuty In the us-east-1 Region. The company wants all DNS logs that relate to the company's Amazon EC2 instances to be inspected. What should a security engineer do to ensure that the EC2 instances are logged?

A. Use IPv6 addresses that are configured for hostnames.

B. Configure external DNS resolvers as internal resolvers that are visible only to IAM.

C. Use the Amazon VPC-provided DNS resolver (AmazonProvidedDNS) for all EC2 instances.

D. Configure a third-party DNS resolver with logging for all EC2 instances.

ANSWER: C

Explanation:

Use the Amazon VPC-provided DNS resolver (AmazonProvidedDNS) for all EC2 instances. GuardDuty automatically monitors DNS request activity made through the Amazon-provided DNS resolver as a foundational data source; no separate DNS log delivery or agent installation is required for this GuardDuty capability. In a VPC, the resolver is available at the VPC network range plus two—for example, `10.0.0.2` in a `10.0.0.0/16` VPC—and is also reachable through the IPv4 link-local address `169.254.169.253`. Instances that use this resolver generate DNS query telemetry that GuardDuty can analyze for suspicious behavior, such as queries to domains associated with command-and-control infrastructure. The VPC must have DNS resolution enabled so that instances can use the Amazon-provided resolver. This configuration preserves GuardDuty visibility into the DNS activity originating from the EC2 instances while retaining AWS-managed DNS resolution within the VPC. See the Amazon GuardDuty data sources documentation and the Amazon VPC DNS concepts documentation for details.

QUESTION NO: 18

A security engineer recently rotated the host keys for an Amazon EC2 instance. The security engineer is trying to access the EC2 instance by using the EC2 Instance. Connect feature. However, the security engineer receives an error (or failed host key validation. Before the rotation of the host keys EC2 Instance Connect worked correctly with this EC2 instance.

What should the security engineer do to resolve this error?

A. Import the key material into AWS Key Management Service (AWS KMS).

B. Manually upload the new host key to the AWS trusted host keys database.

C. Ensure that the AmazonSSMManagedInstanceCore policy is attached to the EC2 instance profile.

D. Create a new SSH key pair for the EC2 instance.

E. Reboot the EC2 instance to refresh the host key used by EC2 Instance Connect.

ANSWER: E

Explanation:

Reboot the EC2 instance to refresh the host key used by EC2 Instance Connect. EC2 Instance Connect validates the SSH server host key when establishing a browser-based connection, which helps ensure that the connection is made to the intended instance. Rotating the instance's SSH host keys changes the identity presented by the SSH service. Until EC2 Instance Connect refreshes its host-key information, validation can fail because the service has information associated with the prior host key. Rebooting the instance causes the updated host-key information to be made available for EC2 Instance Connect validation and resolves the mismatch. This preserves host-key verification rather than bypassing an important protection against connecting to an impersonated host. The reboot does not require replacing the instance's user SSH key pair; it addresses the server-side SSH host identity that changed during rotation. AWS documents EC2 Instance Connect connection methods and its host-key validation behavior in the EC2 Instance Connect methods guide. For prerequisites, supported configuration, and operational guidance, see the EC2 Instance Connect setup documentation.

QUESTION NO: 19

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

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

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

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

A. Configure Amazon Macie to identify any objects in DOC-EXAMPLE-BUCKET1, including .txt objects, that contain PII.

B. Use Amazon CloudWatch Logs Insights to identify any objects in DOC-EXAMPLE-BUCKET1 that contain PII and that were available to the access key.

C. Use Amazon OpenSearch Service (Amazon Elasticsearch Service) to query the CloudTrail logs in DOC-EXAMPLE-BUCKET2 for API calls that used the access key to access an object that contained PII.

D. Use Amazon Athena to query the CloudTrail logs in DOC-EXAMPLE-BUCKET2 for GetObject API calls that used the access key, and correlate the accessed object keys with Amazon Macie PII findings.

E. Use AWS Identity and Access Management Access Analyzer to identify any API calls that used the access key to access objects that contained PII in DOC-EXAMPLE-BUCKET1.

ANSWER: A D

Explanation:

Configure Amazon Macie to identify objects in DOC-EXAMPLE-BUCKET1, including .txt objects, that contain PII. Macie uses automated sensitive-data discovery to inspect eligible Amazon S3 objects and produce findings for detected sensitive data, such as personally identifiable information. Its results provide the required classification of the bucket's text objects independently of the identity that accessed them. See the Amazon Macie automated sensitive data discovery documentation.

Use Amazon Athena to query the CloudTrail logs in DOC-EXAMPLE-BUCKET2 for S3 object-level GetObject events performed by the access key during the 60-day period. CloudTrail log records contain the access key ID in the user identity information and include the requested S3 object key, enabling the engineer to identify exactly which objects the application accessed. The engineer can then correlate those object keys with Macie findings, filtering for keys with a .txt extension, to determine whether accessed text files contained PII. Athena can directly query CloudTrail log files stored in Amazon S3 after a table is created for the logs. See Querying AWS CloudTrail logs with Amazon Athena. This investigation assumes the trail was configured to capture S3 data events, which are the events that record object-level GetObject activity.

QUESTION NO: 20

A company wants to prevent SSH access through the use of SSH key pairs for any Amazon Linux 2 Amazon EC2 instances in its AWS account. However, a system administrator occasionally will need to access these EC2 instances through SSH in an emergency. For auditing purposes, the company needs to record any commands that a user runs in an EC2 instance.

What should a security engineer do to configure access to these EC2 instances to meet these requirements?

A. Use the EC2 serial console Configure the EC2 serial console to save all commands that are entered to an Amazon S3 bucket. Provide the EC2 instances with an IAM role that allows the EC2 serial console to access Amazon S3. Configure an IAM account for the system administrator. Provide an IAM policy that allows the IAM account to use the EC2 serial console.

B. Use EC2 Instance Connect Configure EC2 Instance Connect to save all commands that are entered to Amazon CloudWatch Logs. Provide the EC2 instances with an IAM role that allows the EC2 instances to access CloudWatch Logs Configure an IAM account for the system administrator. Provide an IAM policy that allows the IAM account to use EC2 Instance Connect.

C. Use an EC2 key pair with an EC2 instance that needs SSH access Access the EC2 instance with this key pair by using SSH. Configure the EC2 instance to save allcommands that are entered to Amazon CloudWatch Logs. Provide the EC2 instance with an IAM role that allows the EC2 instance to access Amazon S3 and CloudWatchLogs.

D. Use AWS Systems Manager Session Manager Configure Session Manager to save all commands that are entered in a session to an Amazon S3 bucket. Provide the EC2 instances with an IAM role that allows Systems Manager to manage the EC2 instances. Configure an IAM account for the system administrator Provide an IAM policy that allows the IAM account to use Session Manager.

ANSWER: D

Explanation:

Use AWS Systems Manager Session Manager Configure Session Manager to save all commands that are entered in a session to an Amazon S3 bucket. Provide the EC2 instances with an IAM role that allows Systems Manager to manage the EC2 instances. Configure an IAM account for the system administrator Provide an IAM policy that allows the IAM account to use Session Manager. This approach provides interactive shell access to managed Amazon Linux 2 instances without requiring an inbound SSH port, a bastion host, or an SSH key pair. The instance must be managed by AWS Systems Manager, which requires the SSM Agent and an instance profile that grants the required Systems Manager permissions, commonly through the AmazonSSMManagedInstanceCore managed policy. The administrator is granted least-privilege IAM permissions to start Session Manager sessions against approved instances. Session Manager preferences can be configured to send session logs to an Amazon S3 bucket. For Linux shell sessions, these logs provide an auditable record of commands and their output, subject to the configured logging settings. The S3 bucket can also enforce encryption and restrictive bucket policies to protect the audit trail. This creates centrally controlled, identity-based emergency access while preserving a durable record of activity. See AWS Systems Manager Session Manager and Logging Session Manager sessions.

QUESTION NO: 21

A systems engineer deployed containers from several custom-built images that an application team provided through a QA workflow The systems engineer used Amazon Elastic Container Service (Amazon ECS) with the Fargate launch type as the target platform The system engineer now needs to collect logs from all containers into an existing Amazon CloudWatch log group

Which solution will meet this requirement?

A. Turn on the awslogs log driver by specifying parameters for awslogs-group and awslogs-region in the LogConfiguration property

B. Download and configure the CloudWatch agent on the container instances

C. Set up Fluent Bit and FluentO as a DaemonSet to send logs to Amazon CloudWatch Logs

D. Configure an IAM policy that includes the logs:CreateLogGroup action. Assign the policy to the container instances.

ANSWER: A

Explanation:

Turn on the awslogs log driver by specifying parameters for awslogs-group and awslogs-region in the LogConfiguration property is correct because Amazon ECS on AWS Fargate has native integration with Amazon CloudWatch Logs through the `awslogs` log driver. The task definition can configure each container's `logConfiguration` to use this driver and identify the existing destination log group with `awslogs-group`, as well as the AWS Region with `awslogs-region`. Including `awslogs-stream-prefix` is also recommended so that ECS creates identifiable log streams for tasks and containers. When the task runs, the container's standard output and standard error are delivered to CloudWatch Logs without changes to the application image or management of any underlying host.

This is particularly appropriate for the Fargate launch type because AWS manages the compute infrastructure. Logging is therefore configured declaratively in the ECS task definition rather than by installing or operating software on container hosts. The task execution role must allow the required CloudWatch Logs operations, such as `logs:CreateLogStream` and `logs:PutLogEvents`, for the target log group. This configuration centrally collects logs from all containers that use the configured task definition in the specified existing CloudWatch log group. See the Amazon ECS documentation for the awslogs log driver and task execution IAM role permissions.

QUESTION NO: 22

A security engineer is configuring AWS. Config for an AWS account that uses a new 1AM entity When the security engineer tries to configure AWS. Config rules and automatic remediation options, errors occur in the AWS CloudTrail logs the security engineer sees the following error message "Insufficient delivery policy to s3 bucket DOC-EXAMPLE-BUCKET, unable to write to bucket provided s3 key prefix is 'null'."

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

A. Check the Amazon S3 bucket policy Verify that the policy allows the config amazon aws com service to write to the target bucket.

B. Verify that the 1AM entity has the permissions necessary to perform the s3 GetBucketAc1 and s3 PutObjecj operations to write to the target bucket.

C. Verify that the Amazon S3 bucket policy has the permissions necessary to perform the s3: GetBucketAcl and s3 PutObject" operations to write to the target bucket.

D. Check the policy that is associated with the 1AM entity Verify that the policy allows the config amazonaws com service to write to the target bucket.

E. Verify that the AWS Config service role has permissions to invoke the BatchGetResourceConfig action instead of the GetResourceConfigHistory action and s3 PutObject" operation.

ANSWER: A C

Explanation:

The required remediation is to ensure that the destination Amazon S3 bucket explicitly authorizes the AWS Config service principal and grants the exact S3 permissions AWS Config requires for delivery. AWS Config delivers configuration history files and snapshots directly to the configured delivery bucket using the `config.amazonaws.com` service principal. Therefore, the bucket policy must allow that service principal to write objects to the AWS Config delivery prefix in the bucket. AWS Config also validates bucket access before delivery, which requires permission to retrieve the bucket ACL. A valid policy statement consequently grants `s3:GetBucketAcl` on the bucket ARN and `s3:PutObject` on the appropriate object ARN, typically under the `AWSLogs/account-id/Config/` path. These permissions should be present in the bucket's resource-based policy, because the reported error identifies an insufficient S3 delivery policy rather than a lack of permissions on the identity used by the engineer to configure the service. AWS provides an example S3 bucket policy for AWS Config delivery that uses these actions and the AWS Config service principal. See the AWS Config S3 bucket policy documentation and the AWS Config IAM role documentation.

QUESTION NO: 23

A company wants to deny a specific federated user named Bob access to an Amazon S3 bucket named DOC-EXAMPLE-BUCKET. The company wants to meet this requirement by using a bucket policy. The company also needs to ensure that this bucket policy affects Bob's S3 permissions only. Any other permissions that Bob has must remain intact.

Which policy should the company use to meet these requirements?

A)

```
{
    "Version": "2012-10-17",
    "Statement": {
        "Principal": {"AWS": "arn:aws:sts::account-id:federated-user/Bob"},
        "Effect": "Deny",
        "Action": "s3:*",
        "Resource": "arn:aws:s3:::DOC-EXAMPLE-BUCKET"
    }
}
```

B)

```
{
    "Version": "2012-10-17",
    "Statement": {
        "Principal": {"AWS": "arn:aws:sts::account-id:federated-user/Bob"},
        "Effect": "Allow",
        "Action": "s3:*",
        "Resource": "arn:aws:s3:::DOC-EXAMPLE-BUCKET"
    }
}
```

C)

```
{
    "Version": "2012-10-17",
    "Statement": {
        "Principal": {"AWS": "arn:aws:iam::account-id:user/Bob"},
        "Effect": "Deny",
        "Action": "s3:*",
        "Resource": "arn:aws:s3:::DOC-EXAMPLE-BUCKET"
    }
}
```

D)

```
{
    "Version": "2012-10-17",
    "Statement": {
        "Principal": {"AWS": "arn:aws:sts::account-id:assumed-role/Bob/role-session-name"},
        "Effect": "Deny",
        "Action": "s3:*",
        "Resource": "arn:aws:s3:::DOC-EXAMPLE-BUCKET"
    }
}
```

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: B

Explanation:

The correct policy is the one that uses an explicit `Deny` in the bucket policy, identifies Bob by his federated-user principal ARN, and limits the policy resources to `arn:aws:s3:::DOC-EXAMPLE-BUCKET` and `arn:aws:s3:::DOC-EXAMPLE-BUCKET/*`. A federated user is represented by an AWS STS federated-user ARN in the form `arn:aws:sts::account-id:federated-user/Bob`. Specifying that principal ensures the deny applies to Bob's federated session identity rather than to an IAM user, an entire IAM role, or every principal that can access the bucket.

The policy must cover both the bucket ARN and the object ARN pattern because S3 bucket-level actions and object-level actions use different resource ARNs. An explicit deny is evaluated before any allow, so it prevents Bob from accessing this bucket even if Bob otherwise receives an applicable S3 allow through federation, an IAM role, or another policy. Because the deny statement is scoped solely to the named bucket and its objects, it does not modify Bob's identity policies or restrict Bob's access to other AWS services and resources. AWS documents federated-user principals in Principal policy-element documentation and explains explicit-deny precedence in its policy evaluation logic.

QUESTION NO: 24

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

wants to create a centralized custom dashboard to correlate these findings with operational data for deeper analysis and insights. The company needs an analytics tool to search and visualize Security Hub findings.

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

A. Designate an AWS account as a delegated administrator for Security Hub. Publish events to Amazon CloudWatch from the delegated administrator account,all member accounts, and required Regions that are enabled for Security Hub findings.

B. Designate an AWS account in an organization in AWS Organizations as a delegated administrator for Security Hub. Publish events to Amazon EventBridgefrom the delegated administrator account, all member accounts, and required Regions that are enabled for Security Hub findings.

C. In each Region, create an Amazon EventBridge rule to deliver findings to an Amazon Kinesis data stream. Configure the Kinesis data streams to output thelogs to a single Amazon S3 bucket.

D. In each Region, create an Amazon EventBridge rule to deliver findings to an Amazon Kinesis Data Firehose delivery stream. Configure the Kinesis DataFirehose delivery streams to deliver the logs to a single Amazon S3 bucket.

E. Use AWS Glue DataBrew to crawl the Amazon S3 bucket and build the schema. Use AWS Glue Data Catalog to query the data and create views to flattennested attributes. Build Amazon QuickSight dashboards by using Amazon Athena.

F. Partition the Amazon S3 data. Use AWS Glue to crawl the S3 bucket and build the schema. Use Amazon Athena to query the data and create views toflatten nested attributes. Build Amazon QuickSight dashboards that use the Athena views.

ANSWER: B D F

Explanation:

Designating an AWS account in an organization in AWS Organizations as a delegated administrator for Security Hub provides the centralized Security Hub administration model needed to manage member accounts and consolidate findings across the organization. Security Hub findings are automatically sent as events to Amazon EventBridge in the Region where Security Hub generates them, enabling event-driven collection from the accounts and Regions in scope. AWS documents Security Hub’s central configuration and cross-Region aggregation capabilities at Security Hub delegated administrator account.

Creating an Amazon EventBridge rule in each Region that sends matching findings to Amazon Kinesis Data Firehose provides managed, near-real-time delivery into a central Amazon S3 data lake. Firehose can deliver EventBridge event data directly to Amazon S3, making the findings available in durable centralized storage. Partitioning that S3 data, cataloging it with an AWS Glue crawler, and querying it through Amazon Athena supports efficient SQL searches and views that transform nested Security Hub finding JSON into dashboard-friendly fields. Amazon QuickSight can then use the Athena views as a dataset, allowing the company to create a custom dashboard and combine security findings with other operational datasets. See Delivering EventBridge events with Firehose and Using Athena data with QuickSight.

QUESTION NO: 25

A company allows application development teams to create IAM roles in their AWS accounts. The security team must ensure that developers cannot create roles with permissions that exceed the permissions approved for their business unit.

The company also needs to allow the security team to update the maximum permissions that developers can delegate without changing every developer policy.

Which solution will meet these requirements?

A. Create an IAM permissions boundary that contains the approved maximum permissions. Require developers to specify the permissions boundary when they create IAM roles.

B. Create an SCP that allows developers to create any IAM role. Attach the SCP to each development account.

C. Create an IAM group for each business unit. Attach the approved permissions directly to each IAM group.

D. Configure AWS Config managed rules to delete IAM roles that contain permissions not approved by the security team.

ANSWER: A

Explanation:

Create an IAM permissions boundary that defines the maximum permissions available to developer-created roles. Grant developers permission to create roles only when the request includes that specific permissions boundary. A permissions boundary sets the maximum permissions that an identity-based policy can grant to an IAM principal. Even if a developer attaches a broad policy to a new role, the role can perform only actions that are allowed by both its identity-based policies and its permissions boundary.

The developer policy can use the `iam:PermissionsBoundary` condition key to require the approved boundary ARN during `iam:CreateRole` and related IAM operations. The security team can centrally maintain the managed policy that is used as the boundary. Changes to that boundary affect the maximum permissions of all IAM roles that use it, without requiring modifications to individual developer policies. See the IAM permissions boundaries documentation and the IAM condition keys documentation.

QUESTION NO: 26

A company has two teams, and each team needs to access its respective Amazon S3 buckets. The company anticipates adding more teams that also will have their own S3 buckets. When the company adds these teams, team members will need the ability to be assigned to multiple teams. Team members also will need the ability to change teams. Additional S3 buckets can be created or deleted.

An IAM administrator must design a solution to accomplish these goals. The solution also must be scalable and must require the least possible operational overhead.

Which solution meets these requirements?

A. Add users to groups that represent the teams. Create a policy for each team that allows the team to access its respective S3 buckets only. Attach the policy to the corresponding group.

B. Create an IAM role for each team. Create a policy for each team that allows the team to access its respective S3 buckets only. Attach the policy to the corresponding role.

C. Create IAM roles that are labeled with an access tag value of a team. Create one policy that allows dynamic access to S3 buckets with the same tag. Attach the policy to the IAM roles. Tag the S3 buckets accordingly.

D. Implement a role-based access control (RBAC) authorization model. Create the corresponding policies, and attach them to the IAM users.

ANSWER: C

Explanation:

Create IAM roles that are labeled with an access tag value of a team. Create one policy that allows dynamic access to S3 buckets with the same tag. Attach the policy to the IAM roles. Tag the S3 buckets accordingly. This implements attribute-based access control (ABAC), which is designed to scale authorization according to attributes rather than requiring separate, resource-specific policies for every team. A reusable IAM policy can compare a principal tag, such as `team`, with the corresponding tag on the requested S3 bucket by using IAM policy variables and condition keys. Access is therefore automatically limited to resources whose team attribute matches the role being used.

This model minimizes administration as teams and buckets change. Administrators create and consistently tag new buckets, then use the same policy framework without building a new policy for each team. A person who belongs to more than one team can be granted permission to assume multiple appropriately tagged roles. When a person changes teams, administrators update the available role assignments instead of revising numerous bucket permissions. AWS identifies ABAC as particularly useful in rapidly changing environments because permissions can be managed through tags and a smaller set of scalable policies. See the AWS IAM ABAC documentation and AWS documentation for controlling access with tags.

QUESTION NO: 27

AWS CloudTrail is being used to monitor API calls in an organization. An audit revealed that CloudTrail is failing to deliver events to Amazon S3 as expected.

What initial actions should be taken to allow delivery of CloudTrail events to S3? (Select TWO.)

A. Verify that the S3 bucket policy allows CloudTrail to write objects.

B. Verify that the IAM role used by CloudTrail has access to write to Amazon CloudWatch Logs.

C. Remove any lifecycle policies on the S3 bucket that are archiving objects to S3 Glacier Flexible Retrieval.

D. Verify that the S3 bucket defined in CloudTrail exists.

E. Verify that the log file prefix defined in CloudTrail exists in the S3 bucket.

ANSWER: A D

Explanation:

CloudTrail log delivery requires an existing Amazon S3 bucket that is configured to authorize the CloudTrail service to write log files. Confirming that the S3 bucket defined in the trail exists is an essential first step, because CloudTrail cannot deliver logs if its configured destination bucket has been deleted, renamed, or was specified incorrectly. The bucket must also be in the same AWS Region as the trail only for certain legacy configurations; for a standard trail, CloudTrail supports delivering to an S3 bucket in any Region, provided the bucket configuration is valid.

The S3 bucket policy must explicitly permit the `cloudtrail.amazonaws.com` service principal to perform `s3:GetBucketAcl` on the bucket and `s3:PutObject` to the appropriate CloudTrail log prefix. AWS recommends conditioning object writes on the `bucket-owner-full-control` canned ACL. These permissions enable CloudTrail to validate bucket ownership and place log objects under the trail's configured prefix and AWS account log path. Checking the bucket's existence and policy directly targets the prerequisites for S3 log delivery. See the AWS guidance for the required CloudTrail S3 bucket policy and the trail configuration requirements.

QUESTION NO: 28

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

Please select:

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

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

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

D. {"Version":"2012-10-17","Statement":[{"Effect":"Deny","Principal":"*","Action":"s3:PutObject","Resource":"arn:aws:s3:::demo/*","Condition":{"Null":{"s3:x-amz-server-side-encryption":"true"}}}]}

ANSWER: D

Explanation:

A bucket policy that explicitly denies `s3:PutObject` requests when the `x-amz-server-side-encryption` request header is absent ensures that every successful upload supplies a server-side encryption setting. The policy applies to object ARNs under the `demo` bucket, rather than the bucket ARN itself, because `PutObject` is an object-level action. Using a `Null` condition with a value of `true` matches requests that do not include the encryption header. Since explicit IAM denies override applicable allows, a principal cannot upload an object without specifying server-side encryption, even if another identity or bucket policy otherwise grants upload permission. Clients can then request either Amazon S3 managed-key encryption or AWS KMS encryption through the supported header value and, where applicable, the KMS key header. This is an effective preventive control when the requirement is that upload requests explicitly use server-side encryption. Amazon

S3 also encrypts new uploads by default, but a deny policy remains useful when enforcing an explicit encryption-header requirement. See the Amazon S3 server-side encryption documentation and Amazon S3 policy condition keys documentation.

QUESTION NO: 29

A company wants to know when users make changes to IAM roles in the company's AWS account. The company uses Amazon CloudWatch and AWS CloudTrail in the account. The company has configured a CloudTrail trail to capture read and write API activity for management events. The company has an Amazon Simple Notification Service (Amazon SNS) topic for security notifications.

A security engineer must implement a solution that provides a notification when an IAM role is edited.

Which solution will meet this requirement?

A. Enable Amazon Detective. Run a Detective investigation for changes to IAM roles. Create an Amazon EventBridge rule that monitors the results of the Detective investigation. Set the SNS topic as the target of the EventBridge rule.

B. Create an Amazon EventBridge rule that monitors AWS API calls from CloudTrail. Scope the event pattern to monitor changes to IAM roles from the iam.amazonaws.com event source. Set the SNS topic as the target of the EventBridge rule.

C. Create a new CloudWatch log group. Configure the CloudTrail trail to send events to the new log group. Set up a CloudWatch metric to monitor changes to IAM roles from the lam.amazonaws.com event source. Create a subscription filter for the log group. Set the SNS topic as the target of the subscription filter.

D. Create a new CloudWatch log group. Configure the CloudTrail trail to send events to the new log group. Create a subscription filter that includes an event pattemn to monitor changes to IAM roles from the lam.amazonaws.com event source. Set the SNS topic as the target of the subscription filter.

ANSWER: B

Explanation:

Create an Amazon EventBridge rule that monitors AWS API calls from CloudTrail, filters for IAM role-change activity with the `iam.amazonaws.com` event source, and sends matching events to the Amazon SNS topic. This is an event-driven solution: CloudTrail records IAM management API calls, and EventBridge can receive those calls as events using the `AWS API Call via CloudTrail` detail type. The rule can further narrow matching events to IAM role edit operations, such as `UpdateRole` and `UpdateAssumeRolePolicy`, by filtering on `detail.eventSource` and relevant `detail.eventName` values. When an authorized user or service invokes one of those APIs, EventBridge evaluates the CloudTrail-derived event and publishes a notification to the configured SNS topic. Because the trail already captures management read and write activity, IAM write operations required for role modifications are available for this workflow. EventBridge is designed to route matching AWS service events to targets such as SNS without polling CloudTrail logs or running an investigation. See the Amazon EventBridge documentation for CloudTrail API events and the IAM documentation for AWS CloudTrail logging.

QUESTION NO: 30

A company is undergoing a layer 3 and layer 4 DDoS attack on its web servers running on IAM.

Which combination of IAM services and features will provide protection in this scenario? (Select THREE).

A. Amazon Route 53

B. IAM Certificate Manager (ACM)

C. Amazon S3

D. IAM Shield

E. Elastic Load Balancer

F. Amazon GuardDuty

G. AWS Shield

ANSWER: A E G

Explanation:

AWS Shield, Amazon Route 53, and Elastic Load Balancer provide complementary protection and resilience for network- and transport-layer DDoS events. AWS Shield Standard is automatically included for AWS customers and provides always-on detection and inline mitigation for common infrastructure-layer attacks, including attacks directed at protected AWS resources. AWS Shield Advanced adds enhanced detection, protections, visibility, and access to the AWS Shield Response Team for eligible protected resources.

Amazon Route 53 uses globally distributed authoritative DNS infrastructure and supports health checks and DNS routing policies. These capabilities help direct users to healthy endpoints and support highly available architectures during an attack. Elastic Load Balancer distributes incoming connections across healthy targets and multiple Availability Zones, enabling applications to scale and continue serving legitimate traffic while avoiding dependence on a single server or Availability Zone. Together, Shield supplies DDoS mitigation, Route 53 supplies resilient DNS and traffic steering, and Elastic Load Balancer supplies scalable, multi-AZ traffic distribution. The service is named AWS Shield; “IAM Shield” is not an AWS service name. See the AWS DDoS Resilience overview and the Elastic Load Balancing documentation.

QUESTION NO: 31

A company deploys its application as a service on an Amazon Elastic Container Service (Amazon ECS) cluster with theAWS Fargate launch type. A security engineer suspects that some incoming requests are malicious. The security engineer needs to inspect the running container by retrieving log files and memory dump flies.

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

A. Migrate the application to an ECS cluster with the Amazon EC2 launch type. Configure the EC2 instances with proper remote access. Log in and inspect the container.

B. Update the application to dump the required data to STDOUT. Use the awslogs log driver to pass the logs to Amazon CloudWatch Logs. Examine the log files in CloudWatch Logs.

C. Turn on Amazon CloudWatch Container Insights for the ECS cluster. Send the log data to Amazon CloudWatch Logs by using AWS Distro for OpenTelemetry. Examine the log data in CloudWatch Logs.

D. Update the ECS task role with AWS Systems Manager permissions. Enable the ECS Exec feature for the ECS service. Use ECS Exec to inspect the container.

ANSWER: D

Explanation:

"Update the ECS task role with AWS Systems Manager permissions. Enable the ECS Exec feature for the ECS service. Use ECS Exec to inspect the container.
" is the correct solution because Amazon ECS Exec provides controlled interactive shell access and command execution within a running container hosted on AWS Fargate. This enables the security engineer to investigate the live container filesystem, locate relevant application logs and memory-dump artifacts, and retrieve their contents through commands without changing the workload’s launch type or managing underlying hosts. ECS Exec uses AWS Systems Manager Session Manager capabilities, so access can be governed through IAM rather than by opening inbound network ports or deploying SSH infrastructure. The task role requires the Systems Manager messaging permissions used by ECS Exec, and the investigating identity must also be authorized to initiate the exec session. ECS Exec can additionally record session activity to Amazon CloudWatch Logs or Amazon S3, supporting auditability during an incident investigation. Enabling this capability is a targeted service and task-definition configuration change, making it the lowest-operational-effort approach for inspecting an already running Fargate task. For setup requirements and supported use cases, see the Amazon ECS Exec documentation and the required IAM permissions guidance.

QUESTION NO: 32

An Incident Response team is investigating an IAM access key leak that resulted in Amazon EC2 instances being launched. The company did not discover the incident until many months later The Director of Information Security wants to implement new controls that will alert when similar incidents happen in the future

Which controls should the company implement to achieve this? {Select TWO.)

A. Enable VPC Flow Logs in all VPCs Create a scheduled IAM Lambda function that downloads and parses the logs, and sends an Amazon SNS notification for violations.

B. Use AWS CloudTrail to create a trail, and apply it to all Regions. Specify an Amazon S3 bucket to receive all the CloudTrail log files.

C. Add the following bucket policy to the company's IAM CloudTrail bucket to prevent log tampering
{
"Version": "2012-10-17-,
"Statement": {
"Effect": "Deny",
"Action": "s3:PutObject",
"Principal": "-",
"Resource": "arn:IAM:s3:::cloudtrail/IAMLogs/111122223333/*"
}
}
Create an Amazon S3 data event for an PutObject attempts, which sends notifications to an Amazon SNS topic.

D. Create a Security Auditor role with permissions to access Amazon CloudWatch Logs m all Regions Ship the logs to an Amazon S3 bucket and make a lifecycle policy to ship the logs to Amazon S3 Glacier.

E. Verify that Amazon GuardDuty is enabled in all Regions, and create an Amazon EventBridge rule for Amazon GuardDuty findings. Add an Amazon SNS topic as the rule's target.

ANSWER: B E

Explanation:

Creating an AWS CloudTrail trail that applies to all AWS Regions and delivers log files to a centralized Amazon S3 bucket establishes the essential audit record for this scenario. CloudTrail records AWS API activity, including IAM credential use and Amazon EC2 management actions such as `RunInstances`. A multi-Region trail helps ensure that activity in any enabled Region is captured, while central S3 delivery provides durable evidence for investigations and enables downstream monitoring or analysis. CloudTrail log-file validation can additionally help establish whether delivered logs were modified or deleted after delivery.

Enabling Amazon GuardDuty in all Regions and routing its findings through an Amazon EventBridge rule to an Amazon SNS topic provides the managed detection and immediate notification path. GuardDuty continuously analyzes relevant AWS telemetry, including CloudTrail management events, to identify suspicious API activity and indicators that credentials may be compromised. EventBridge can match GuardDuty findings as they are generated, and SNS can promptly notify the incident response team through configured subscriptions. Together, centralized CloudTrail logging supplies authoritative visibility into resource-launch activity, while GuardDuty findings and SNS notifications reduce the time to discover and respond to suspicious use of leaked access keys. See the AWS CloudTrail trail documentation and the Amazon GuardDuty EventBridge notification documentation.

QUESTION NO: 33

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 that are in an Auto Scaling group. The EC2 instances are deployed in a private subnet Of a VPC that does not have internet access. The EC2 instances and the S3 buckets are in the same AWS account

The EC2 instances access the S3 buckets through an S3 gateway endpoint that has the default access policy. Each EC2 instance is associated With an instance profile role that has a policy that explicitly allows the s3:GetObject action and the s3:PutObject action for only the required S3 buckets.

The company learns that one or more of the EC2 instances are compromised and are exfiltrating data to an S3 bucket that is outside the companys organization in AWS Organizations. A security engtneer must implement a solution to stop this exfiltration of data and to keep the EC2 processing job functional.

Which solution will meet these requirements?

A. Update the policy on the S3 gateway endpoint to allow the S3 actions only if the values of the aws:ResourceOrgID and aws:PrincipalOrgID condition keys match the company's values.

B. Update the policy on the instance profile role to allow the S3 actions only if the value of the aws:ResourceOrgID condition key matches the company's value.

C. Add a network ACL rule to the subnet of the EC2 instances to block outgoing connections on port 443.

D. Apply an SCP on the AWS account to allow the S3 actions only if the values of the aws:ResourceOrgID and aws:PrincipalOrgID condition keys match the company's values.

ANSWER: D

Explanation:

Apply an SCP on the AWS account to allow the S3 actions only if the values of the aws:ResourceOrgID and aws:PrincipalOrgID condition keys match the company's values. An AWS Organizations service control policy provides a centralized permission boundary for every IAM principal in the account, including the instance profile credentials used by the Auto Scaling fleet. This is valuable for a compromised workload because the restriction is enforced outside the account's ordinary IAM administration and cannot be expanded by changing permissions on the instance role.

The `aws:ResourceOrgID` condition key identifies the AWS Organization that owns the target resource. Restricting S3 access to resources in the company's organization prevents writes to an external S3 bucket, which stops the identified exfiltration path. The `aws:PrincipalOrgID` condition key ensures that the calling principal is also associated with the company's organization. At the same time, the required S3 buckets remain accessible because they are owned by an account in that organization, so the existing processing workflow can continue through the S3 gateway endpoint.

A service control policy defines the maximum available permissions for accounts in an organization. AWS documents SCP evaluation and organization condition keys at AWS Organizations service control policies and AWS global condition context keys.

QUESTION NO: 34

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

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

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

Which combination of steps will meet these requirements? (Choose two.)

A. Create an AWS CloudFormation template that contains the 10 required AWS Config rules. Deploy the template by using CloudFormation StackSets in the security-01 account.

B. Create a conformance pack that contains the 10 required AWS Config rules. Deploy the conformance pack from the security-01 account.

C. Create a conformance pack that contains the 10 required AWS Config rules. Deploy the conformance pack from the management-01 account.

D. Create an AWS CloudFormation template that will activate AWS Config. Deploy the template by using CloudFormation StackSets in the security-01 account.

E. Create an AWS CloudFormation template that will activate AWS Config. Deploy the template by using CloudFormation StackSets in the management-01 account.

ANSWER: B E

Explanation:

Creating a conformance pack that contains the 10 required AWS Config rules and deploying the conformance pack from the security-01 account provides the organization-wide rule baseline. AWS Config conformance packs package multiple AWS Config rules into one YAML-based deployable unit. With AWS Config trusted access enabled and security-01 registered as the AWS Config delegated administrator, that account can deploy an organization conformance pack to the organization. AWS Config automatically deploys the pack to existing member accounts and to accounts that are subsequently added, while each account can still maintain rules needed for its individual compliance requirements.

Creating an AWS CloudFormation template that will activate AWS Config and deploying the template by using CloudFormation StackSets in the management-01 account establishes the account-creation baseline. The management account can create service-managed StackSets that target an organization or organizational units. Enabling automatic deployment causes StackSets to create stack instances for accounts added to the targeted scope, so the template can provision the AWS Config recorder, delivery channel, and required IAM permissions when a new account is created. This complements the organization conformance pack: StackSets enable AWS Config, and the pack deploys the required compliance rules. See AWS Config conformance packs and StackSets with AWS Organizations.

QUESTION NO: 35

A company uses AWS Organizations and AWS IAM Identity Center. The company wants to provide contractors with access to selected AWS accounts by using the companys existing external identity provider.

Contractor accounts must be created, updated, and removed automatically when changes occur in the external identity provider. The company must grant contractors only the permissions that are required in each AWS account.

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

A. Configure the external identity provider as a SAML 2.0 identity provider for IAM Identity Center.

B. Configure SCIM provisioning between the external identity provider and IAM Identity Center.

C. Create least-privilege permission sets. Assign contractor users or groups to the required AWS accounts.

D. Create an IAM user for each contractor in every AWS account. Require the contractors to rotate access keys every 90 days.

E. Create a shared IAM role in the management account. Give all contractors permission to assume the shared role.

F. Create an SCP that grants contractors access to the AWS accounts that they require.

ANSWER: A B C

Explanation:

Configure the external identity provider as a SAML 2.0 identity provider for IAM Identity Center. This allows users to authenticate through the existing corporate identity system instead of requiring separately managed IAM users in AWS.

Configure System for Cross-domain Identity Management (SCIM) provisioning between the identity provider and IAM Identity Center. SCIM synchronizes users and groups so that contractor onboarding, attribute updates, and deprovisioning in the identity provider are reflected in IAM Identity Center. Create least-privilege permission sets and assign the appropriate contractor users or groups to the required AWS accounts. IAM Identity Center provisions the corresponding account access roles and applies the permission set policy. See the IAM Identity Center external identity provider documentation and the IAM Identity Center permission sets documentation.

QUESTION NO: 36

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

Which solution will meet these requirements?

A. Create new S3 buckets with S3 Object Lock enabled in compliance mode. Place objects in the S3 buckets.

B. Use S3 Glacier Vault Lock to attach a Vault Lock policy to new S3 buckets. Wait 24 hours to complete the Vault Lock process. Place objects in the S3 buckets.

C. Create new S3 buckets with S3 Object Lock enabled in governance mode. Place objects in the S3 buckets.

D. Create new S3 buckets with S3 Object Lock enabled in governance mode. Add a legal hold to the S3 buckets. Place objects in the S3 buckets.

ANSWER: A

Explanation:

Create new S3 buckets with S3 Object Lock enabled in compliance mode. Place objects in the S3 buckets. This meets the WORM requirement because S3 Object Lock prevents a protected object version from being deleted or overwritten for the duration of its retention period. Compliance mode provides the strongest immutability guarantee: no user can shorten the retention period or delete the protected object version before retention expires, including the AWS account root user. This is specifically designed for regulatory and records-retention workloads that require non-erasable, non-rewritable storage. S3 Object Lock is enabled when a bucket is created and operates on versioned objects; after the bucket exists, a retention period can be applied to individual object versions or established through a default retention configuration for newly uploaded objects. S3 Object Lock works with S3 Standard, so the company can retain its required storage class while enforcing WORM controls. To ensure every uploaded object receives protection automatically, the security engineer should configure a default Object Lock retention rule in compliance mode with an appropriate retention duration. AWS documents that compliance-mode retention cannot be removed or reduced by any user, which includes the root user. See the Amazon S3 Object Lock User Guide and the S3 retention modes documentation.

QUESTION NO: 37

A company used AWS Organizations to set up an environment with multiple AWS accounts. The company's organization currently has two AWS accounts, and the companyexpects to add more than 50 AWS accounts during the next 12 months The company will require all existing and future AWS accounts to use Amazon GuardDuty. Eachexisting AWS account has GuardDuty active. The company reviews GuardDuty findings by logging into each AWS account individually.

The company wants a centralized view of the GuardDuty findings for the existing AWS accounts and any future AWS accounts. The company also must ensure that anynew AWS account has GuardDuty automatically turned on.

Which solution will meet these requirements?

A. Enable AWS Security Hub in the organization’s management account. Configure GuardDuty within the management account to send all GuardDuty findings toSecurity Hub.

B. Create a new AWS account in the organization. Enable GuardDuty in the new account. Designate the new account as the delegated administrator account forGuardDuty. Configure GuardDuty to add existing accounts as member accounts. Select the option to automatically add new AWS accounts to the organization

C. Create a new AWS account in the organization. Enable GuardDuty in the new account. Enable AWS Security Hub in each account. Select the option toautomatically add new AWS accounts to the organization.

D. Enable AWS Security Hub in the organization's management account. Designate the management account as the delegated administrator account for SecurityHub. Add existing accounts as member accounts. Select the option to automatically add new AWS accounts to the organization. Send all Security Hub findingsto the organization's GuardDuty account.

ANSWER: B

Explanation:

Create a dedicated AWS account, enable Amazon GuardDuty in it, and designate it as the GuardDuty delegated administrator. This provides a scalable, centralized GuardDuty administration model for an AWS Organization without

requiring administrators to sign in to every member account. The delegated administrator can associate the organization's existing accounts as GuardDuty member accounts and view and manage their findings centrally. It can also configure the organization's auto-enable settings so that GuardDuty is enabled automatically for accounts that join the organization in the future. These settings can include the core GuardDuty protection plan and, where applicable, additional protection plans. A dedicated security account is a strong operational practice because it separates security administration and investigation activities from the AWS Organizations management account and ordinary workload accounts. Using GuardDuty's native organization integration directly addresses both required outcomes: centralized visibility into findings across current and future accounts, and automatic GuardDuty activation for newly created or added accounts. AWS documents that an organization can have one GuardDuty delegated administrator per Region, so the configuration should be performed in every AWS Region where the company requires GuardDuty coverage. See Managing GuardDuty accounts with AWS Organizations and Automatically enabling GuardDuty for organization accounts.

QUESTION NO: 38

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

How should the Security team accomplish this?

A. Update the IAM WAF rules in the affected account and use IAM Firewall Manager to push updated IAM WAF rules across all other accounts.

B. Use GuardDuty centralized logging and Amazon SNS to set up alerts to notify all application teams of security incidents.

C. Use GuardDuty alerts to write an IAM Lambda function that updates all accounts by adding additional NACLs on the Amazon EC2 instances to block known malicious IP addresses.

D. Use IAM Shield Advanced to identify threats in each individual account and then apply the account-based protections to all other accounts through Organizations.

ANSWER: A

Explanation:

Update the IAM WAF rules in the affected account and use IAM Firewall Manager to push updated IAM WAF rules across all other accounts. This approach lets the security team convert an observed attack pattern into a centrally governed preventive control. For example, after identifying malicious request characteristics or source addresses against one internet-facing application, the team can add an AWS WAF rule to a Web ACL and use AWS Firewall Manager to distribute and enforce the corresponding WAF policy throughout the organization. Firewall Manager is designed for centralized security-policy administration across AWS Organizations accounts and organizational units, including AWS WAF policies for supported resources such as Application Load Balancers, API Gateway APIs, CloudFront distributions, and AWS AppSync APIs. The policy can automatically apply protections to in-scope resources and maintain compliance as accounts and resources are added. This provides consistent, scalable preventive enforcement rather than relying on each application team to manually reproduce the remediation. AWS WAF supplies the request-inspection and blocking capability, while Firewall Manager supplies the multi-account policy deployment and ongoing governance needed for the stated centralized operational model. See the AWS Firewall Manager Developer Guide and AWS WAF web ACL documentation.

QUESTION NO: 39

A company uses Amazon API Gateway to present REST APIs to users. An API developer wants to analyze API access patterns without the need to parse the log files.

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

A. Configure access logging for the required API stage.

B. Configure an AWS CloudTrail trail destination for API Gateway events. Configure filters on the userIdentity, userAgent, and sourceIPAddress fields.

C. Configure an Amazon S3 destination for API Gateway logs. Run Amazon Athena queries to analyze API access information.

D. Use Amazon CloudWatch Logs Insights to analyze API access information.

E. Select the Enable Detailed CloudWatch Metrics option on the required API stage.

ANSWER: A D

Explanation:

Configure access logging for the required API stage. is required because API Gateway stage access logs capture request-level information, such as request identifiers, source IP address, HTTP method, resource path, status, response size, and latency. The developer can define a structured access-log format, commonly JSON, which makes the request data available for query-based analysis rather than requiring custom code to read and parse individual log files.

Use Amazon CloudWatch Logs Insights to analyze API access information. provides the lowest-effort analysis path when the stage access logs are delivered to CloudWatch Logs. Logs Insights is a serverless, interactive query service that can search and aggregate log events directly, enabling analyses such as the most frequently requested paths, requests by source address, response-status distributions, and latency trends. It avoids building a separate log-export pipeline and avoids managing external parsing infrastructure. API Gateway access logging configuration and the CloudWatch Logs Insights query interface therefore provide the required structured visibility with minimal operational work. See the API Gateway logging documentation and the CloudWatch Logs Insights documentation.

QUESTION NO: 40

A company needs to delect unauthenticated access to its Amazon Elastic Kubernetes Service (Amazon EKS) clusters. The company needs a solution that requires no additional configuration ot the existing EKS deployment.

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

A. Install an Amazon EKS add-on from a security vendor.

B. Enable AWS Security Hub Monitor the Kubernetes findings

C. Monitor Amazon CloudWatch Container Insights metrics for Amazon EKS.

D. Enable Amazon GuardDuty Use EKS Audit Log Monitoring.

ANSWER: D

Explanation:

Enable Amazon GuardDuty Use EKS Audit Log Monitoring. Amazon GuardDuty EKS Protection continuously monitors Amazon EKS control-plane activity for suspicious and potentially unauthorized behavior, including Kubernetes API activity represented in EKS audit logs. It is designed to identify threats such as anonymous or unauthenticated access attempts and generates GuardDuty findings that can be investigated or routed to downstream incident-response workflows.

This is the lowest-effort solution because GuardDuty directly consumes and analyzes the EKS audit-log data for supported clusters; the company does not need to deploy an agent, install a Kubernetes add-on, alter application workloads, or separately enable and manage EKS control-plane audit logging for GuardDuty's use. The operational action is simply enabling EKS Protection in GuardDuty for the applicable AWS Region and account. GuardDuty then applies AWS-managed threat detection logic and produces findings when relevant activity is detected.

AWS documents that EKS Protection monitors EKS audit logs and that GuardDuty manages the audit-log collection independently of EKS control-plane logging configuration. See Amazon GuardDuty EKS Protection and GuardDuty Kubernetes finding types.

QUESTION NO: 41

A company uses SAML federation to grant users access to AWS accounts. A company workload that is in an isolated AWS account runs on immutable infrastructure with no human access to Amazon EC2.

The company requires a specialized user known as a break glass user to have access to the workload AWS account and instances in the case of SAML errors. A recent audit discovered that the company did not create the break glass user for the AWS account that contains the workload.

The company must create the break glass user. The company must log any activities of the break glass user and send the logs to a security team.

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

A. Create a local individual break glass IAM user for the security team. Create a trail in AWS CloudTrail that has Amazon CloudWatch Logs turned on. Use Amazon EventBridge to monitor local user activities.

B. Create a break glass EC2 key pair for the AWS account. Provide the key pair to the security team. Use AWS CloudTrail to monitor key pair activity. Send notifications to the security team by using Amazon Simple Notification Service (Amazon SNS).

C. Create a break glass IAM role for the account. Allow security team members to perform the AssumeRoleWithSAML operation. Create an AWS Cloud Trail trail that has Amazon CloudWatch Logs turned on. Use Amazon EventBridge to monitor security team activities.

D. Create a local individual break glass IAM user on the operating system level of each workload instance. Configure unrestricted security groups on the instances to grant access to the break glass IAM users.

E. Configure AWS Systems Manager Session Manager for Amazon EC2. Configure an AWS Cloud Trail filter based on Session Manager. Send the results to an Amazon Simple Notification Service (Amazon SNS) topic.

ANSWER: A E

Explanation:

Creating a local individual break glass IAM user for the security team provides an authentication path that is independent of the SAML identity provider. This satisfies the need to access the isolated AWS account when federation is unavailable. The IAM user should be tightly controlled, protected with strong credentials and MFA, and granted only the permissions needed for emergency recovery. An AWS CloudTrail trail records the user's AWS API activity, including account and EC2-management actions. Delivering CloudTrail events to Amazon CloudWatch Logs makes those events available for timely detection, and Amazon EventBridge can match CloudTrail activity associated with the emergency identity and route alerts or automation to the security team.

Configuring AWS Systems Manager Session Manager provides the appropriate controlled path from the AWS account to managed EC2 instances. Session Manager avoids the need to open inbound SSH or RDP access and does not require distributing or retaining EC2 key-pair private keys. The break glass identity can be authorized to start Systems Manager sessions when emergency instance access is required. CloudTrail records Systems Manager API events, including session-related activity; a CloudTrail-based filter and an Amazon SNS topic can therefore notify the security team when this access path is used. Together, these measures provide both emergency account access and auditable, alertable instance access. See the AWS CloudTrail User Guide and AWS Systems Manager Session Manager documentation.

QUESTION NO: 42

A company's Security Engineer is copying all application logs to centralized Amazon S3 buckets. Currently, each of the company's applications is in its own IAM account, and logs are pushed into S3 buckets associated with each account. The Engineer will deploy an IAMLambda function into each account that copies the relevant log files to the centralized S3 bucket.

The Security Engineer is unable to access the log files in the centralized S3 bucket. The Engineer's IAM user policy from the centralized account looks like this:

The centralized S3 bucket policy looks like this:

Why is the Security Engineer unable to access the log files?

A. The S3 bucket policy does not explicitly allow the Security Engineer access to the objects in the bucket.

B. The object ACLs are not being updated to allow the users within the centralized account to access the objects

C. The Security Engineers IAM policy does not grant permissions to read objects in the S3 bucket

D. The s3:PutObject and s3:PutObjectAcl permissions should be applied at the S3 bucket level

ANSWER: B

Explanation:

The object ACLs are not being updated to allow the users within the centralized account to access the objects. In this cross-account upload design, the IAM Lambda function in each application account is the principal that creates the log objects. Under the legacy S3 object-ownership model, the uploading account owns those newly created objects, even though it writes them to a bucket in the centralized account. Bucket-level permissions alone do not automatically give the centralized bucket owner or its IAM users access to objects owned by another AWS account when the object ACL does not grant that access.

The upload workflow must therefore ensure that ownership or access is granted when each object is written. A common legacy approach is for the uploader to include the `bucket-owner-full-control` canned ACL and to have permission to set that ACL. The preferred current approach is to configure S3 Object Ownership with Bucket owner enforced on the centralized bucket. That setting disables ACLs and makes the bucket owner own every object uploaded to the bucket, simplifying centralized log access. AWS documents cross-account object ownership and the bucket-owner-full-control ACL at Amazon S3 cross-account access walkthrough and explains Bucket owner enforced ownership at Controlling ownership of objects and disabling ACLs.

QUESTION NO: 43

A security engineer is defining the controls required to protect the IAM account root user credentials in an IAM Organizations hierarchy. The controls should also limit the impact in case these credentials have been compromised.

Which combination of controls should the security engineer propose?(Select THREE.)

A)

B)

C) Enable multi-factor authentication (MFA) for the root user.

D) Set a strong randomized password and store it in a secure location.

E) Create an access key ID and secret access key, and store them in a secure location.

F) Apply the following permissions boundary to the toot user:

A. Apply service control policies (SCPs) that explicitly deny sensitive actions for root users in member accounts.

B. Option B

C. Enable multi-factor authentication (MFA) for the root user.

D. Set a strong randomized password and store it in a secure location.

E. Create an access key ID and secret access key, and store them in a secure location.

F. Apply the following permissions boundary to the toot user:

ANSWER: A C D

Explanation:

Applying service control policies (SCPs) that explicitly restrict sensitive actions for member-account root users provides an Organizations-level guardrail. SCPs define the maximum permissions available in affected member accounts, so a compromised root credential cannot perform actions that the SCP explicitly denies. This is an effective defense-in-depth control for accounts governed by AWS Organizations. AWS documents that SCPs apply to IAM principals in member

accounts, including the account root user; the management account is an important exception and should therefore have especially strict root-user operational controls.

Enabling multi-factor authentication (MFA) for the root user adds a separate authentication factor, substantially reducing the risk that a stolen or guessed password alone can be used to access the account. A strong, unique, randomly generated root password stored only in an approved secure secrets location also reduces the chance of password reuse, guessing, or exposure through routine handling. These controls support AWS guidance to protect root credentials, avoid routine root-user use, and use MFA on the root user. See the AWS root user best practices and the AWS Organizations SCP documentation.

QUESTION NO: 44

A company wants to receive automated email notifications when AWS access keys from developer AWS accounts are detected on code repository sites.

Which solution will provide the required email notifications?

A. Create an Amazon EventBridge rule to send Amazon Simple Notification Service (Amazon SNS) email notifications for Amazon GuardDuty UnauthorizedAccessIAMUser/InstanceCredentialExfiltration OutsideAWS findings.

B. Change the AWS account contact information for the Operations type to a separate email address. Periodically poll this email address for notifications.

C. Create an Amazon EventBridge rule that reacts to AWS Health events that have a value of Risk for the service category Configure email notifications by using Amazon Simple Notification Service (Amazon SNS).

D. Implement new anomaly detection software. Ingest AWS CloudTrail logs. Configure monitoring for ConsoleLogin events in the AWS Management Console. Configure email notifications from the anomaly detection software.

E. Create an Amazon EventBridge rule that matches AWS Health AWS_IAM_EXPOSED_ACCESS_KEY events and sends them to an Amazon SNS topic with email subscriptions.

ANSWER: E

Explanation:

AWS generates an AWS Health event when it detects an IAM access key exposed in a public code repository or other public location. The event type code is `AWS_IAM_EXPOSED_ACCESS_KEY`, and it provides an event-driven signal specifically intended for responding to potentially compromised credentials. An Amazon EventBridge rule can match this AWS Health event, including the IAM service and exposed-access-key event type, then send the matching event to an Amazon SNS topic. Subscribing the required email addresses to that topic provides automated email notification whenever AWS detects a newly exposed access key. The recipients must confirm the SNS email subscription before messages can be delivered. This design avoids polling and provides a managed, near-real-time notification path directly from AWS's exposed-credential detection service. After notification, the security team should immediately deactivate, delete, or rotate the affected access key and investigate its activity. AWS documents the exposed access key AWS Health event and its available response actions at AWS Health event types reference. EventBridge can route AWS Health events to targets such as SNS as described in the AWS Health and Amazon EventBridge documentation.

QUESTION NO: 45

A company has a group of Amazon EC2 instances in a single private subnet of a VPC with no internet gateway attached. A security engineer has installed the Amazon CloudWatch agent on all instances in that subnet to capture logs from a specific application. To ensure that the logs flow securely, the company's networking team has created VPC endpoints for CloudWatch monitoring and CloudWatch logs. The networking team has attached the endpoints to the VPC.

The application is generating logs. However, when the security engineer queries CloudWatch, the logs do not appear.

Which combination of steps should the security engineer take to troubleshoot this issue? (Choose three.)

A. Ensure that the EC2 instance profile that is attached to the EC2 instances has permissions to create log streams and write logs.

B. Create a metric filter on the logs so that they can be viewed in the AWS Management Console.

C. Check the CloudWatch agent configuration file on each EC2 instance to make sure that the CloudWatch agent is collecting the proper log files.

D. Check the VPC endpoint policies of both VPC endpoints to ensure that the EC2 instances have permissions to use them.

E. Create a NAT gateway in the subnet so that the EC2 instances can communicate with CloudWatch.

F. Ensure that the security groups allow all the EC2 instances to communicate with each other to aggregate logs before sending.

ANSWER: A C D

Explanation:

“Ensure that the EC2 instance profile that is attached to the EC2 instances has permissions to create log streams and write logs,” “Check the CloudWatch agent configuration file on each EC2 instance to make sure that the CloudWatch agent is collecting the proper log files,” and “Check the VPC endpoint policies of both VPC endpoints to ensure that the EC2 instances have permissions to use them” validate the key authorization, collection, and private-connectivity controls for this design.

The CloudWatch agent obtains AWS credentials from the EC2 instance profile. Its role therefore needs CloudWatch Logs API permissions needed to publish the configured data, including creating log streams when required and calling `logs:PutLogEvents`. The agent configuration must also identify the real application-log path and specify the intended log group and stream settings. A running agent does not automatically collect arbitrary application logs.

Because the instances have no internet gateway, they depend on the CloudWatch Logs interface endpoint for private API access. Endpoint policies can restrict which principals and CloudWatch Logs actions are permitted through an endpoint, so those policies must allow the workload role’s required calls. The monitoring endpoint is similarly relevant if the agent is configured to publish metrics. AWS documents the required CloudWatch Logs permissions at CloudWatch Logs IAM identity-based policies and private endpoint access at CloudWatch Logs interface VPC endpoints.

QUESTION NO: 46

A company uses an external identity provider to allow federation into different AWS accounts. A security engineer for the company needs to identify the federated user that terminated a production Amazon EC2 instance a week ago.

What is the FASTEST way for the security engineer to identify the federated user?

A. Review the AWS CloudTrail event history logs in an Amazon S3 bucket and look for the TerminateInstances event to identify the federated user from the role session name.

B. Filter the AWS CloudTrail event history for the TerminateInstances event and identify the assumed IAM role. Review the AssumeRoleWithSAML event call in CloudTrail to identify the corresponding username.

C. Search the AWS CloudTrail logs for the TerminateInstances event and note the event time.
Review the IAM Access Advisor tab for all federated roles. The last accessed time should match the time when the instance was terminated.

D. Use Amazon Athena to run a SQL query on the AWS CloudTrail logs stored in an Amazon S3 bucket and filter on the TerminateInstances event. Identify the corresponding role and run another query to filter the AssumeRoleWithWebIdentity event for the user name.

ANSWER: B

Explanation:

Filter the AWS CloudTrail event history for the TerminateInstances event and identify the assumed IAM role. Review the AssumeRoleWithSAML event call in CloudTrail to identify the corresponding username. This is the fastest approach because CloudTrail Event history is intended for interactive investigation of recent management events and retains the most recent 90 days of events in each AWS Region. An EC2 instance termination is recorded as a `TerminateInstances` API event. Its

`userIdentity` details identify the role session that performed the action, including the assumed-role ARN and session context. For SAML federation from an external identity provider, the associated `AssumeRoleWithSAML` event establishes that role session. Reviewing the federation event and its session information provides the correlation needed to map the AWS assumed-role activity to the originating federated identity. Because the incident occurred one week ago, it is within Event history's available time range, so the engineer can perform both lookups directly in CloudTrail without first downloading, cataloging, or querying log files. CloudTrail records identity information for API activity specifically to support this kind of security investigation and audit trail correlation. See Viewing events with CloudTrail Event history and CloudTrail userIdentity element.

QUESTION NO: 47

A company recently adopted new compliance standards that require all user actions in AWS to be logged. The user actions must be logged for all accounts that belong to an organization in AWS Organizations. The company needs to set alarms that respond when specified actions occur. The alarms must forward alerts to an email distribution list. The alerts must occur in as close to real time as possible.

Which solution will meet these requirements?

A. Implement an AWS CloudTrail trail as an organizational trail. Configure the trail with Amazon CloudWatch Logs forwarding. In CloudWatch Logs, set a metric filter for any user action events that the company specifies. Create an Amazon CloudWatch alarm to provide alerts for occurrences within a reported period and to publish messages to an Amazon Simple Notification Service (Amazon SNS) topic.

B. Implement an AWS CloudTrail trail. Configure the trail with Amazon CloudWatch Logs forwarding. In CloudWatch Logs, set a metric filter for any user action events that the company specifies. Create an Amazon CloudWatch alarm to provide alerts for occurrences within a reported period and to send messages to an Amazon Simple Queue Service (Amazon SQS) queue.

C. Implement an AWS CloudTrail trail as an organizational trail. Configure the trail to store logs in an Amazon S3 bucket. Configure an Amazon EC2 instance to mount the S3 bucket as a file system to ingest new log files that are pushed to the S3 bucket. Configure the EC2 instance also to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic when one of the specified actions is found in the logs.

D. Implement an AWS CloudTrail trail. Configure the trail to store logs in an Amazon S3 bucket. Each hour, create an AWS Glue Data Catalog that references the S3 bucket. Configure Amazon Athena to initiate queries against the Data Catalog to identify the specified actions in the logs.

ANSWER: A

Explanation:

Implementing an AWS CloudTrail trail as an organizational trail, forwarding events to Amazon CloudWatch Logs, creating metric filters and CloudWatch alarms, and publishing notifications to an Amazon SNS topic meets every requirement with AWS-managed services and low operational overhead. An organization trail records activity for all accounts in the AWS Organizations organization, including accounts that are added later, providing centralized coverage for the compliance requirement. CloudTrail management events capture user-initiated AWS API activity, and delivery to CloudWatch Logs makes those events available for ongoing monitoring. CloudWatch Logs metric filters evaluate incoming log events for defined patterns, such as particular API calls or actions by specified principals, and publish matching occurrences as CloudWatch metrics. A CloudWatch alarm can evaluate that metric promptly and invoke an SNS notification when its threshold is met. An SNS topic supports email subscriptions, so the company can subscribe its email distribution list and receive alarm notifications. This architecture provides near-real-time detection because CloudTrail events are delivered to CloudWatch Logs continuously, rather than waiting for periodic batch analysis. See Creating a trail for an organization and Monitoring log data with metric filters and alarms.

QUESTION NO: 48

A company has an application that processes personally identifiable information (PII). The application runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The company's security policies require that data is encrypted in transit at all times to avoid the possibility of exposing any PII in plaintext.

Which solutions could a security engineer use to meet these requirements'? (Select TWO)

A. Terminate SSL from clients on the existing ALB. Use HTTPS to connect from the ALB to the EC2 instances.

B. Replace the existing ALB with a Network Load Balancer (NLB) On the NLB, configure an SSL listener and TCP passthrough to receive client connections Terminate HTTPS traffic from the NLB on the EC2 instances.

C. Replace the existing ALB with a Network Load Balancer (NLB) On the NLB, configure TCP passthrough to receive client connections Terminate SSL from the NLB on the EC2 instances

D. Configure a Network Load Balancer (NLB) with TCP passthrough to receive client connections Terminate SSL on the existing ALB.

E. Configure a Network Load Balancer (NLB) with a TLS listener to receive client connections Configure TCP passthrough on the existing ALB so that the NLB can reach the EC2 instances Terminate SSL from the ALB on the EC2 instances.

ANSWER: A C

Explanation:

“Terminate SSL from clients on the existing ALB. Use HTTPS to connect from the ALB to the EC2 instances.” meets the requirement because it establishes two encrypted network legs: TLS from each client to the Application Load Balancer, and HTTPS/TLS from the load balancer to each registered EC2 target. The ALB decrypts traffic only for request processing and then encrypts it again before transmitting it across the network to the application instances. AWS documents this as HTTPS listener traffic being forwarded to HTTPS targets.

“Replace the existing ALB with a Network Load Balancer (NLB) On the NLB, configure TCP passthrough to receive client connections Terminate SSL from the NLB on the EC2 instances” also preserves encryption in transit. A Network Load Balancer TCP listener can pass the TCP/TLS stream through without decrypting it. TLS is then terminated by the EC2 application, so the client-to-instance connection remains encrypted while traversing the load balancer and network. This approach is appropriate when the application instances must perform TLS termination themselves. See the AWS documentation for Application Load Balancer listeners and HTTPS target groups and Network Load Balancer TCP listeners.

QUESTION NO: 49

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 (1AM) 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 1AM 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 1AM and Cognito.

E. Create an 1AM policy that includes a condition for minimum password length Enforce the policy for 1AM and Cognito

ANSWER: B C

Explanation:

Update the password length policy in the Cognito configuration because Amazon Cognito user pools manage the credentials for users stored in the user pool. A user pool password policy includes a configurable minimum password length, along with optional requirements for uppercase characters, lowercase characters, numbers, symbols, temporary-password validity, and password-history settings. The security engineer should configure the required minimum directly in each applicable user pool’s password policy.

Update the password length policy in the on-premises Active Directory configuration because Active Directory remains the credential authority for users who authenticate through the company's federated directory. The minimum password length must therefore be enforced by the applicable Active Directory domain password policy, such as through the domain's Group Policy settings or a fine-grained password policy where appropriate. Federation to AWS does not transfer password-policy enforcement to AWS; authentication continues to rely on the directory's password controls. Configuring the policies in both credential stores ensures that the required password length applies to the distinct populations of Cognito-native users and Active Directory users. See the Amazon Cognito password policy documentation and Microsoft's minimum password length policy documentation.

QUESTION NO: 50

A company wants to implement host-based security for Amazon EC2 instances and containers in Amazon Elastic Container Registry (Amazon ECR). The company has deployed AWS Systems Manager Agent (SSM Agent) on the EC2 instances. All the company's AWS accounts are in one organization in AWS Organizations. The company will analyze the workloads for software vulnerabilities and unintended network exposure. The company will push any findings to AWS Security Hub, which the company has configured for the organization.

The company must deploy the solution to all member accounts, including new accounts, automatically. When new workloads come online, the solution must scan the workloads.

Which solution will meet these requirements?

A. Use SCPs to configure scanning of EC2 instances and ECR containers for all accounts in the organization.

B. Configure a delegated administrator for Amazon GuardDuty for the organization. Create an Amazon EventBridge rule to initiate analysis of ECR containers

C. Configure a delegated administrator for Amazon Inspector for the organization. Configure automatic scanning for new member accounts.

D. Configure a delegated administrator for Amazon Inspector for the organization. Create an AWS Config rule to initiate analysis of ECR containers.

ANSWER: C

Explanation:

Configuring a delegated administrator for Amazon Inspector for the organization and configuring automatic scanning for new member accounts meets the centralized deployment and continuous assessment requirements. Amazon Inspector is designed to continuously identify software vulnerabilities and unintended network exposure for supported Amazon EC2 workloads and Amazon ECR container images. For EC2 scanning, Inspector uses Systems Manager capabilities and the installed SSM Agent to gather the inventory and vulnerability data needed for assessment. This aligns with the company's existing SSM Agent deployment.

Amazon Inspector integration with AWS Organizations allows a delegated administrator to centrally manage member accounts. The delegated administrator can configure auto-enablement so new organization accounts are automatically enrolled for the relevant Inspector scan types, avoiding manual configuration as the organization grows. Inspector also automatically evaluates eligible new EC2 instances and newly pushed ECR images as they become available, providing the required ongoing coverage for new workloads. Amazon Inspector findings are natively integrated with AWS Security Hub, allowing findings to flow into the organization's Security Hub deployment for centralized visibility and correlation. See the Amazon Inspector AWS Organizations integration documentation and the Amazon Inspector ECR scanning documentation.

QUESTION NO: 51

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

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

A. Enable Amazon GuardDuty to automatically monitor for malicious activity and block unauthorized access.

B. Subscribe to AWS Shield Advanced and reach out to AWS Support in the event of an attack.

C. Use VPC Flow Logs to monitor network traffic and an AWS Lambda function to automatically block an attacker's IP using security groups.

D. Set up an Amazon EventBridge rule to monitor the AWS CloudTrail events in real time, use AWS Config rules to audit the configuration, and use AWS Systems Manager for remediation.

E. Use AWS WAF to create rules to respond to such attacks.

ANSWER: B E

Explanation:

Subscribe to AWS Shield Advanced and reach out to AWS Support in the event of an attack. provides enhanced DDoS protection and access to the AWS Shield Response Team (SRT). During an active event, the company can open an AWS Support case to engage the SRT, which can help investigate the attack and apply mitigations. Shield Advanced also supplies DDoS detection, near-real-time visibility, and protections for supported resources, helping the organization respond rapidly and reduce the duration and impact of an attack. AWS documents the incident-response process at Responding to a DDoS attack.

Use AWS WAF to create rules to respond to such attacks. enables application-layer protections in front of web applications. For example, rate-based rules can detect and rate-limit requests from IP addresses that exceed a configured request threshold within a five-minute evaluation window. WAF rules can also block known malicious patterns and unwanted request characteristics before they reach the application. Placing AWS WAF on a supported entry point such as an Amazon CloudFront distribution or Application Load Balancer provides automated, continuously enforced mitigation during a layer-7 DDoS event, avoiding reliance on manual intervention. See AWS WAF rate-based rules.

QUESTION NO: 52

A business stores website images in an Amazon S3 bucket. The firm serves the photos to end users through Amazon CloudFront. The firm learned lately that the photographs are being accessible from nations in which it does not have a distribution license.

Which steps should the business take to safeguard the photographs and restrict their distribution? (Select two.)

A. Update the S3 bucket policy to restrict access to a CloudFront origin access identity (OAI).

B. Update the website DNS record to use an Amazon Route 53 geolocation record deny list of countries where the company lacks a license.

C. Add a CloudFront geo restriction deny list of countries where the company lacks a license.

D. Update the S3 bucket policy with a deny list of countries where the company lacks a license.

E. Enable the Restrict Viewer Access option in CloudFront to create a deny list of countries where the company lacks a license.

ANSWER: A C

Explanation:

“Update the S3 bucket policy to restrict access to a CloudFront origin access identity (OAI).” protects the S3 origin by allowing CloudFront, rather than internet users, to retrieve the image objects. This is essential because viewers could otherwise bypass CloudFront’s edge controls by requesting public S3 object URLs directly. AWS now recommends Origin Access Control for new CloudFront distributions, but an OAI is still a valid mechanism for private S3 origins and is the correct choice as stated.

“Add a CloudFront geo restriction deny list of countries where the company lacks a license.” applies geographic restrictions at CloudFront edge locations. CloudFront determines the viewer’s country from the request IP address and blocks access for countries included in the distribution’s deny list. Combining an origin-access restriction with CloudFront geographic

restrictions ensures that the images are available only through the controlled CloudFront delivery path and are blocked in the unlicensed countries.

AWS documents geographic restriction as the CloudFront feature for allowing or blocking viewer access by country: CloudFront geographic restrictions. For protecting an S3 origin so users access content through CloudFront, see Restricting access to an Amazon S3 origin.

QUESTION NO: 53

A company is migrating one of its legacy systems from an on-premises data center to AWS. The application server will run on AWS, but the database must remain in the on-premises data center for compliance reasons. The database is sensitive to network latency. Additionally, the data that travels between the on-premises data center and AWS must have IPsec encryption.

Which combination of AWS solutions will meet these requirements? (Choose two.)

A. AWS Site-to-Site VPN
AWS Site-to-Site VPN is a service that allows you to securely connect your on-premises data center to yourAWS VPC over the internet using IPsec encryption. This solution meets the requirement of encrypting the data in transit between the on-premises data center and AWS.

B. AWS Direct Connect
AWS Direct Connect is a service that allows you to establish a dedicated network connection between youron-premises data center and your AWS VPC. This solution meets the requirement of reducing network latency between the on-premises data center and AWS.

C. AWS VPN CloudHub
AWS VPN CloudHub is a service that allows you to connect multiple VPN connections from differentlocations to the same virtual private gateway in your AWS VPThis solution is not relevant for this scenario, as there is only one on-premises data center involved.

D. VPC peering
VPC peering is a service that allows you to connect two or more VPCs in the same or different regionsusing private IP addresses. This solution does not meet the requirement of connecting an on-premises data center to AWS, as it only works for VPCs.

E. NAT gateway

NAT gateway is a service that allows you to enable internet access for instances in a private subnet in yourAWS VPC. This solution does not meet the requirement of connecting an on-premises data center to AWS, as it only works for outbound traffic from your VPC.

ANSWER: A B

Explanation:

AWS Direct Connect and AWS Site-to-Site VPN together meet the connectivity, performance, and encryption requirements. AWS Direct Connect provides a dedicated network connection between the on-premises data center and AWS. This avoids relying solely on public internet connectivity and delivers more consistent network performance, which is important for an application tier in AWS that frequently communicates with a latency-sensitive database that remains on premises. Direct Connect can be used with a private virtual interface to reach resources in a VPC through a Direct Connect gateway or virtual private gateway.

AWS Site-to-Site VPN provides the required IPsec protection for traffic in transit. It establishes IPsec tunnels between the on-premises customer gateway device and an AWS VPN endpoint. For this scenario, the VPN can be configured as an encrypted overlay on the Direct Connect connection, combining Direct Connect's predictable private connectivity with IPsec encryption. This architecture supports the compliance requirement without moving the database or exposing its traffic unencrypted. AWS documents VPN over Direct Connect as a supported design for encrypting traffic across a Direct Connect connection. See the AWS Site-to-Site VPN User Guide and AWS Direct Connect User Guide.

QUESTION NO: 54

A company wants to deploy a continuous security threat-detection service at scale to automatically analyze all the company's member accounts in AWS Organizations within the ap-east-1 Region. The company's organization includes a management account, a security account, and many member accounts. When the company creates a new member account, the threat-detection service should automatically analyze the new account so that the company can review any findings from the security account.

Which solution uses AWS security best practices and meets these requirements with the LEAST effort?

A. Activate Amazon GuardDuty in ap-east-1, enable trusted access with AWS Organizations, designate the security account as the GuardDuty delegated administrator, and configure automatic enrollment for organization accounts.

B. Activate Amazon GuardDuty in ap-east-1 with trusted access toAWS Organizations Designate the management account as the GuardDuty organization administrator.

C. Activate AWS Security Hub in ap-east-1 Designate the management account as the Security Hub delegated administrator by using the console.

D. Activate AWS Control Tower in ap-east-1 with trusted access to AWS Organizations Designate the security account as the organization administrator.

ANSWER: A

Explanation:

Activate Amazon GuardDuty in ap-east-1, enable trusted access with AWS Organizations, designate the security account as the GuardDuty delegated administrator, and configure organization auto-enable is the appropriate solution. Amazon GuardDuty is AWS's managed, continuous threat-detection service. It analyzes AWS CloudTrail management events, VPC Flow Logs, DNS query logs, and other supported sources to identify potentially malicious or unauthorized activity. GuardDuty is configured on a Regional basis, so it must be enabled in ap-east-1 to analyze activity there.

Using AWS Organizations with GuardDuty centralizes administration while following the security best practice of separating security operations from the organization management account. The management account enables trusted access and delegates GuardDuty administration to the dedicated security account. The delegated administrator can then manage GuardDuty member accounts and view organization findings centrally. Enabling automatic account enrollment ensures that both existing organization accounts and newly created member accounts receive GuardDuty coverage without requiring manual per-account activation. Findings from those accounts are available to the security account for centralized investigation and response. For implementation details, see the Amazon GuardDuty AWS Organizations documentation and the GuardDuty delegated administrator documentation.

QUESTION NO: 55

A company uses identity federation to authenticate users into an identity account (987654321987) where the users assume an IAM role named IdentityRole. The users then assume an IAM role named JobFunctionRole in the target AWS account (123456789123) to perform their job functions. A user is unable to assume the IAM role in the target account. The policy attached to the role in the identity account is:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Action": [
                "sts:AssumeRole"
            ],
            "Resource": [
                "arn:aws:iam::*:role/JobFunctionRole"
            ],
            "Effect": "Allow"
        }
    ]
}
```

What should be done to enable the user to assume the appropriate role in the target account?

A. Update the IAM policy attached to the role in the identity account to be:

B. Update the trust policy on the role in the target account to be:

C. Update the trust policy on the role in the identity account to be:

D. Update the IAM policy attached to the role in the target account to be:

ANSWER: B

Explanation:

Update the trust policy on the role in the target account to be the policy that trusts `arn:aws:iam::987654321987:role/IdentityRole`. Cross-account role assumption requires two complementary permissions: the calling principal must be permitted to invoke `sts:AssumeRole`, and the target role must trust that calling principal in its trust policy. Here, the user has already authenticated through federation and is operating as a session of `IdentityRole`. Therefore, when the user attempts to assume `JobFunctionRole`, AWS STS evaluates the target role's resource-based trust policy against the identity role principal. Adding the identity role ARN as the trusted AWS principal to `JobFunctionRole`'s trust relationship grants the required cross-account trust while retaining the permissions policy on `IdentityRole` as the authorization for making the STS call. This pattern supports role chaining from a central identity account to job-function roles in workload accounts. AWS documents that a role trust policy identifies the principals allowed to assume the role, and cross-account access requires the target account's role to trust the external account principal. See AWS IAM role trust policy documentation and AWS cross-account role access documentation.

QUESTION NO: 56

AWS CloudTrail is being used to monitor API calls in an organization. An audit revealed that CloudTrail is failing to deliver events to Amazon S3 as expected.

What initial actions should be taken to allow delivery of CloudTrail events to S3? (Select TWO.)

A. Verify thattheS3 bucket policy allows CloudTrail to write objects.

B. Verify thatthe1AM role used by CloudTrail has access to write to Amazon CloudWatch Logs.

C. Remove any lifecycle policies on the S3 bucket that are archiving objects to S3 Glacier Flexible Retrieval.

D. Verify thattheS3 bucket defined in CloudTrail exists.

E. Verify that the log file prefix defined in CloudTrail exists in the S3 bucket.

ANSWER: A D

Explanation:

“Verify thattheS3 bucket policy allows CloudTrail to write objects” is an essential initial action because CloudTrail delivers log files to the configured S3 bucket by using the CloudTrail service principal. The bucket policy must allow the required S3 operations, including permission for CloudTrail to write log objects to the appropriate AWSLogs path. AWS also recommends allowing CloudTrail to retrieve the bucket ACL, which lets the service validate its ability to deliver logs. A restrictive bucket policy is a common cause of delivery failures, particularly when policies use conditions for source account, source ARN, or encryption requirements.

“Verify thattheS3 bucket defined in CloudTrail exists” is also required because a trail can deliver logs only to an existing destination bucket. The bucket must be available in the expected AWS partition and be correctly configured as the trail’s S3 destination. Once the bucket exists and its policy grants CloudTrail the required access, CloudTrail can create its log-file key structure automatically and begin delivering event records. AWS documents the required bucket-policy permissions and delivery setup in the CloudTrail S3 bucket policy guide and the trail configuration documentation.

QUESTION NO: 57

A company runs a cuslom online gaming application. The company uses Amazon Cognito for user authentication and authorization.

A security engineer wants to use AWS to implement fine-grained authorization on resources in the custom application. The security engineer must implement a solution that uses the user attributes that exist in Cognito. The company has already set up a user pool and an identity pool in Cognito.

Which solution will meet these requirements?

A. Create a set of IAM roles and IAM policies. Configure the Cognito identity pool to assign users to the IAM roles.

B. Create a policy store in Amazon Verified Permissions. Configure the Cognito user pool as the identity source. Map Cognito access tokens to the Verified Permissions schema.

C. Create customer managed permissions by using AWS Resource Access Manager (AWS RAM). Configure the Cognito identity pool to assign users to the customer managed permissions.

D. Create a set of IAM users and IAM policies. Configure the Cognito user pool to assign users to the IAM users.

ANSWER: B

Explanation:

Create a policy store in Amazon Verified Permissions. Configure the Cognito user pool as the identity source. Map Cognito access tokens to the Verified Permissions schema.

Amazon Verified Permissions is purpose-built to make fine-grained authorization decisions for custom applications. A policy store provides a centralized location for Cedar authorization policies and the corresponding schema that defines principals, resources, actions, and attributes. The application can submit an authorization request to Verified Permissions when a player attempts an action, such as accessing a game resource or modifying an in-game object.

Amazon Cognito user pools can be configured as an identity source for the policy store. Verified Permissions validates Cognito user pool tokens and maps token claims, including relevant standard or custom user attributes, into the policy store schema. Policies can then use those attributes to evaluate permissions at request time. This separates application authentication, handled by Cognito, from granular and centrally managed authorization, handled by Verified Permissions. Cognito identity pools are primarily used to exchange identities for temporary AWS credentials, whereas the relevant integration for application authorization decisions uses the Cognito user pool token claims. See the Amazon Verified Permissions User Guide and Cognito identity sources for Verified Permissions.

QUESTION NO: 58

A company has AWS accounts in an organization in AWS Organizations. The company needs to install a corporate software package on all Amazon EC2 instances for all the accounts in the organization.

A central account provides base AMIs for the EC2 instances. The company uses AWS Systems Manager for software inventory and patching operations.

A security engineer must implement a solution that detects EC2 instances ttjat do not have the required software. The solution also must automatically install the software if the software is not present.

Which solution will meet these requirements?

A. Provide new AMIs that have the required software pre-installed. Apply a tag to the AMIs to indicate that the AMIs have the required software. Configure an SCP that allows new EC2 instances to be launched only if the instances have the tagged AMIs. Tag all existing EC2 instances.

B. Configure a custom patch baseline in Systems Manager Patch Manager. Add the package name for the required software to the approved packages list. Associate the new patch baseline with all EC2 instances. Set up a maintenance window for software deployment.

C. Centrally enable AWS Config. Set up the ec2-managedinstance-applications-required AWS Config rule for all accounts Create an Amazon EventBridge rule that reacts to AWS Config events. Configure the EventBridge rule to invoke an AWS Lambda function that uses Systems Manager Run Command to install the required software.

D. Create a new Systems Manager Distributor package for the required software. Specify the download location. Select all EC2 instances in the different accounts. Install the software by using Systems Manager Run Command.

ANSWER: C

Explanation:

“Centrally enable AWS Config. Set up the ec2-managedinstance-applications-required AWS Config rule for all accounts Create an Amazon EventBridge rule that reacts to AWS Config events. Configure the EventBridge rule to invoke an AWS Lambda function that uses Systems Manager Run Command to install the required software.” meets both the continuous detection and automatic remediation requirements. The `ec2-managedinstance-applications-required` managed rule evaluates Amazon EC2 instances that are managed by AWS Systems Manager and uses Systems Manager Inventory data to determine whether a specified application is installed. This directly identifies instances that lack the corporate package rather than relying solely on the configuration of a launch image or a scheduled deployment.

AWS Config can be deployed centrally across the organization, including through organization-wide AWS Config rules, to provide consistent compliance evaluation in every member account. When the rule records a noncompliant resource, an Amazon EventBridge rule can respond to the AWS Config compliance-change event. Invoking a Lambda function enables the remediation workflow to call Systems Manager Run Command on the affected managed instance and execute the package-installation command. Systems Manager provides the managed, auditable execution channel required to install the software without requiring direct administrative access to the instance. This creates a closed-loop control: inventory-based detection, compliance notification, and automated installation. See the AWS Config documentation for ec2-managedinstance-applications-required and the AWS Systems Manager Run Command documentation.

QUESTION NO: 59

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

What should the security engineer recommend?

A. Enable Amazon RDS encryption to encrypt the database and snapshots. Enable Amazon Elastic Block Store (Amazon EBS) encryption on Amazon EC2 instances. Include the database credential in the EC2 user data field. Use an AWS Lambda function to rotate database credentials. Set up TLS for the connection to the database.

B. Install a database on an Amazon EC2 instance. Enable third-party disk encryption to encrypt the Amazon Elastic Block Store (Amazon EBS) volume. Store the database credentials in AWS
CloudHSM with automatic rotation. Set up TLS for the connection to the database.

C. Enable Amazon RDS encryption to encrypt the database and snapshots. Enable Amazon Elastic Block Store (Amazon EBS) encryption on Amazon EC2 instances. Store the database credentials in AWS Secrets Manager with automatic rotation. Set up TLS for the connection to the RDS hosted database.

D. Set up an AWS CloudHSM cluster with AWS Key Management Service (AWS KMS) to store KMS keys. Set up Amazon RDS encryption using AWS KMS to encrypt the database. Store database credentials in the AWS Systems Manager Parameter Store with automatic rotation. Set up TLS for the connection to the RDS hosted database.

E. Use Amazon RDS encryption with AWS KMS, encrypt any Amazon EBS volumes used by EC2 instances, store database credentials in AWS Secrets Manager with scheduled automatic rotation, and enforce TLS for database connections.

ANSWER: C E

Explanation:

Enable Amazon RDS encryption to encrypt the database and snapshots, enable Amazon Elastic Block Store (Amazon EBS) encryption on Amazon EC2 instances, store database credentials in AWS Secrets Manager with automatic rotation, and use TLS for database connections. This is an AWS-managed, defense-in-depth approach that reduces operational burden while protecting data at rest, in transit, and through credential lifecycle management. Amazon RDS encryption uses AWS Key Management Service (AWS KMS) keys to protect the underlying DB storage, automated backups, read replicas, and snapshots. EBS encryption similarly provides managed encryption for EC2-attached storage. AWS Secrets Manager is designed to centrally store application secrets and supports scheduled automatic rotation for supported database engines through a rotation Lambda function. Applications can retrieve the current credential at runtime rather than exposing it in user data or static configuration. TLS encrypts connections between the application and database, protecting sensitive information and authentication material while traversing the network. These controls are complementary: storage encryption limits exposure from underlying media, Secrets Manager limits the lifetime and distribution of credentials, and TLS protects communications. AWS documentation describes Amazon RDS encryption and AWS Secrets Manager rotation.

QUESTION NO: 60

A company is running a new workload across accounts that are 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. Create an AWS Config Custom Policy rule by using AWS CloudFormation Guard. Include the tag key of CostCenter and the approved values. Create an SCP that denies the creation of resources when the value of the aws:RequestTag/CostCenter condition key is not one of the three approved values.

B. Create an AWS CloudTrail trail. Create an Amazon EventBridge rule that includes a rule statement that matches the creation of new resources. Configure the EventBridge rule to invoke an AWS Lambda function that checks for the CostCenter tag. Program the Lambda function to block creation in case of a noncompliant value.

C. Enable tag policies for the organization. Create a tag policy that specifies a tag key of CostCenter and the approved values. Configure the policy to enforce noncompliant operations. Create an SCP that denies the creation of resources when the aws:RequestTag/CostCenter condition key has a null value.

D. Enable tag policies for the organization. Create a tag policy that specifies a tag key of CostCenter and the approved values. Create an Amazon EventBridge rule that invokes an AWS Lambda function when a noncompliant tag is created. Program the Lambda function to block changes to the tag.

ANSWER: C

Explanation:

Enable tag policies for the organization. Create a tag policy that specifies a tag key of CostCenter and the approved values. Configure the policy to enforce noncompliant operations. Create an SCP that denies the creation of resources when the aws:RequestTag/CostCenter condition key has a null value.

A tag policy centrally defines the permitted capitalization and values for the `CostCenter` tag across organization accounts. When the policy is configured for enforcement, AWS prevents noncompliant tagging operations for supported resource types. This includes attempts to add or modify the `CostCenter` tag with a value outside the approved set, which satisfies the requirement to prevent later changes to an invalid value. Tag-policy enforcement is applied at the point of the supported API operation rather than relying on post-creation detection or remediation.

The SCP complements this control by requiring the tag at resource-creation time. A deny statement using the `aws:RequestTag/CostCenter` condition key with a null-value test prevents principals in affected organization accounts from creating resources without supplying `CostCenter`. Together, the SCP establishes tag presence at creation, while enforced tag policies constrain the tag to the three approved values during supported create and tag-update operations. AWS documents tag-policy enforcement and supported services at Tag policy enforcement, and documents request tag condition keys at AWS global condition context keys.

QUESTION NO: 61

While securing the connection between a company's VPC and its on-premises data center, a security engineer sent a ping command from an on-premises host (IP address 203.0.113.12) to an Amazon EC2 instance (IP address 172.31.16.139). The ping command did not return a response. The flow log in the VPC showed the following:

```
2 123456789010 eni-1235b8ca 203.0.113.12 172.31.16.139 0 0 1 4 336 1432917027 1432917142 ACCEPT OK

2 123456789010 eni-1235b8ca 172.31.16.139 203.0.113.12 0 0 1 4 336 1432917094 1432917142 REJECT OK
```

What action should be performed to allow the ping to work?

A. In the security group of the EC2 instance, allow inbound ICMP traffic.

B. In the security group of the EC2 instance, allow outbound ICMP traffic.

C. In the VPC's NACL, allow inbound ICMP traffic.

D. In the VPC's NACL, allow outbound ICMP traffic.

ANSWER: D

Explanation:

In the flow-log record, the rejected ICMP traffic is the return traffic from the EC2 instance to the on-premises host. A ping requires both the inbound ICMP echo request and the outbound ICMP echo reply to pass. The required action is therefore **"In the VPC's NACL, allow outbound ICMP traffic."** Network ACLs apply at the subnet boundary and are stateless. Consequently, an allow rule for traffic entering a subnet does not automatically permit the corresponding response to leave it. The network ACL associated with the EC2 instance's subnet must include an outbound allow rule for ICMP to the on-premises network, such as 203.0.113.12/32 or the relevant on-premises CIDR range. The rule number must be evaluated before any conflicting deny rule, because network ACL rules are processed in numerical order and the first matching rule is used. Once the outbound ICMP echo reply is allowed, the instance can respond across the established VPC-to-on-premises connectivity path. VPC Flow Logs record ACCEPT or REJECT based on the network interface traffic evaluation, making the REJECT entry useful evidence that the return packet is not being permitted. See the Amazon VPC network ACL documentation and the VPC Flow Logs documentation.

QUESTION NO: 62

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

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

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

A. Configure the key policy for the customer managed key in the production account to allow access to the Lambda service.

B. Configure the key policy for the customer managed key in the production account to allow access to the IAM role of the Lambda function in the development account.

C. Configure a new IAM policy in the production account with permissions to use the customer managed key. Apply the IAM policy to the IAM role that the Lambda function in the development account uses.

D. Configure a new key policy in the development account with permissions to use the customer managed key. Apply the key policy to the IAM role that the Lambda function in the development account uses.

E. Configure the IAM role for the Lambda function in the development account by attaching an IAM policy that allows access to the customer managed key in the production account.

ANSWER: B E

Explanation:

Cross-account use of an AWS KMS customer managed key requires authorization from both the account that owns the key and the identity that calls KMS. The production account, which owns the customer managed key, must add the Lambda execution role from the development account as a principal in the KMS key policy. This resource-based policy grants that external role permission to perform only the cryptographic operations needed by the workload, such as `kms:Encrypt` and, where applicable, `kms:GenerateDataKey`. The policy should reference the role ARN in the development account and the specific key remains owned and administered in production.

The Lambda execution role also needs an identity-based IAM policy in the development account that allows the necessary KMS actions on the ARN of the customer managed key in the production account. AWS KMS evaluates this cross-account request using both sides of authorization: the key policy enables the external account principal to use the key, and the role policy permits that principal to make the requested API calls. Applying least privilege by limiting actions and resources to the required key and encryption operations satisfies the stated compliance and access requirements. AWS documents this model in its guidance for allowing users in other accounts to use KMS keys and using IAM policies with AWS KMS.

QUESTION NO: 63

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 1AM access key and a secret access key. Program the Lambda function to store the credentials on the player's device. Create 1AM keys for existing players.B Migrate the player credentials from the Aurora database to AWS Secrets Manager. When a new player signs up. create a key-value pair in Secrets Manager for the player's user ID and password.

B. 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.

C. Instead of using usernames and passwords for authentication, issue API keys to new and existing players. Create an Amazon API Gateway API to give the game client access to the game's functionality.

ANSWER: B

Explanation:

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. This is the appropriate solution because Amazon

Cognito user pools provide a managed user directory and authentication service designed for application users at large scale. Federation lets players authenticate through an external identity provider, such as a social provider, rather than requiring the game operator to maintain a separate password for those users. Consequently, password recovery and login-assistance workflows are handled by the selected identity provider or by Cognito's managed account-recovery capabilities, substantially reducing the customer-service burden. Cognito also supports secure token-based authentication: after a successful sign-in, the game receives standard JSON Web Tokens that it can validate and use to authorize access to game resources. Migrating the application authentication mechanism to Cognito centralizes identity controls instead of retaining application credentials in the Aurora database. User pools support federation with social and OpenID Connect identity providers and integrate with application clients for web and mobile sign-in flows. See the Amazon Cognito user pools documentation and AWS guidance on social identity provider sign-in.

QUESTION NO: 64

A security engineer is configuring account-based access control (ABAC) to allow only specific principals to put objects into an Amazon S3 bucket. The principalsalready have access to Amazon S3.

The security engineer needs to configure a bucket policy that allows principals to put objects into the S3 bucket only if the value of the Team tag on the object matches the value of the Team tag that is associated with the principal. During testing, the security engineer notices that a principal can still put objects into the S3 bucket when the tag values do not match.

Which combination of factors are causing the PutObject operation to succeed when the tag values are different? (Select TWO.)

A. The principal's identity-based policy grants access to put objects into the S3 bucket with no conditions.

B. The principal's identity-based policy overrides the condition because the identity-based policy contains an explicit allow.

C. The S3 bucket's resource policy does not deny access to put objects.

D. The S3 bucket's resource policy cannot allow actions to the principal.

E. The bucket policy does not apply to principals in the same zone of trust.

ANSWER: A C

Explanation:

The principal's identity-based policy grants access to put objects into the S3 bucket with no conditions. This creates an applicable explicit allow for `s3:PutObject` regardless of the principal's `Team` tag or the object tag supplied with the request. In AWS policy evaluation, permissions granted by applicable identity-based and resource-based policies are combined. A bucket-policy allow statement whose tag-matching condition is not satisfied simply does not apply; this results in no allow from that statement, rather than an explicit denial. Consequently, the unconditional identity-policy allow can authorize the upload.

The S3 bucket's resource policy does not deny access to put objects. To enforce that mismatched tag values cannot be uploaded when principals may have other allows, the bucket policy must contain an applicable explicit `Deny` for requests where the object's `Team` tag does not match the caller's principal tag. Explicit denies override allows from identity-based policies and resource-based policies. An allow statement conditioned on matching tags alone is insufficient as a guardrail when another applicable policy already allows `s3:PutObject`. AWS documents this evaluation behavior in the IAM policy evaluation logic and documents using principal and object tags for S3 ABAC in the Amazon S3 tagging and policy guide.

QUESTION NO: 65

A company uses AWS Organizations and has Amazon Elastic Kubernetes Service (Amazon EKS) clusters in many AWS accounts. A security engineer integrates Amazon EKS with AWS CloudTrail.

The CloudTrail trails are stored in an Amazon S3 bucket in each account to monitor API calls. The security engineer observes that CloudTrail logs are not displaying Kubernetes pod creation events.

What should the security engineer do to view the Kubernetes events from Amazon CloudWatch?

A. Configure the EKS clusters to use private S3 VPC endpoints. Configure the S3 buckets for logging.

B. Enable Kubernetes API server component logs for each cluster.

C. Enable cross-origin resource sharing (CORS) in the S3 bucket that is used for logging.

D. Configure CloudWatch. View the events in the CloudWatch console.

ANSWER: B

Explanation:

Enable Kubernetes API server component logs for each cluster. AWS CloudTrail records activity through AWS APIs, such as Amazon EKS control-plane management operations, but Kubernetes pod creation occurs through the Kubernetes API server inside the cluster control plane. To deliver visibility for those Kubernetes API requests to Amazon CloudWatch Logs, Amazon EKS control plane logging must be enabled on every relevant cluster. The API server log type records API server activity, including requests made to create Kubernetes resources such as pods. For more complete security visibility, EKS audit logs can also be enabled because they record audit events for requests to the Kubernetes API server. EKS sends the selected control plane log types to a CloudWatch Logs log group for the cluster, where the security team can search and filter events, configure metric filters and alarms, or forward logs to a centralized monitoring platform. Since the organization has clusters in multiple AWS accounts, this setting must be configured per cluster, then CloudWatch Logs can be centrally monitored through an appropriate cross-account observability or log-aggregation design. AWS documents the available EKS control plane log types and their CloudWatch Logs delivery behavior in the Amazon EKS control plane logging guide. CloudTrail's scope for AWS API activity is described in the AWS CloudTrail User Guide.

QUESTION NO: 66

A company uses Amazon RDS for MySQL as a database engine for its applications. A recent security audit revealed an RDS instance that is not compliant with company policy for encrypting data at rest. A security engineer at the company needs to ensure that all existing RDS databases are encrypted using server-side encryption and that any future deviations from the policy are detected. Which combination of steps should the security engineer take to accomplish this? (Choose two.)

A. Create an AWS Config rule to detect the creation of unencrypted RDS databases. Create an Amazon EventBridge rule to trigger on the AWS Config rules compliance state change and use Amazon Simple Notification Service (Amazon SNS) to notify the security operations team.

B. Use AWS System Manager State Manager to detect RDS database encryption configuration drift. Create an Amazon EventBridge rule to track state changes and use Amazon Simple Notification Service (Amazon SNS) to notify the security operations team.

C. Create a read replica for the existing unencrypted RDS database and enable replica encryption in the process. Once the replica becomes active, promote it into a standalone database instance and terminate the unencrypted database instance.

D. Take a snapshot of the unencrypted RDS database. Copy the snapshot and enable snapshot encryption in the process. Restore the database instance from the newly created encrypted snapshot. Terminate the unencrypted database instance.

E. Enable encryption for the identified unencrypted RDS instance by changing the configurations of the existing database.

ANSWER: A D

Explanation:

Amazon RDS storage encryption cannot be enabled in place on an existing unencrypted DB instance. To remediate the current instance, take a DB snapshot, copy that snapshot while selecting encryption and an AWS KMS key, and restore a replacement DB instance from the encrypted snapshot. The restored instance has encrypted storage, automated backups, read replicas, and snapshots. After application connectivity has been migrated and validated, the original unencrypted instance can be terminated. This is the supported migration pattern for applying encryption at rest to existing RDS data.

For ongoing detection, AWS Config can continuously evaluate RDS DB instances against an encryption requirement, such as the AWS managed `rds-storage-encrypted` rule. AWS Config publishes compliance-change events to Amazon EventBridge. An EventBridge rule can match a transition to NON_COMPLIANT and invoke an Amazon SNS topic, providing

the security operations team with prompt notification whenever an unencrypted RDS instance is identified. Together, the encrypted snapshot restoration remediates the existing database, while AWS Config, EventBridge, and SNS provide continuous policy monitoring and alerting.

See Encrypting Amazon RDS resources and AWS Config rds-storage-encrypted managed rule.

QUESTION NO: 67

A development team is using an IAM Key Management Service (IAM KMS) CMK to try to encrypt and decrypt a secure string parameter from IAM Systems Manager Parameter Store. However, the development team receives an error message on each attempt.

Which issues that are related to the CMK could be reasons for the error? (Select TWO.)

A. The CMK that is used in the attempt does not exist.

B. The CMK that is used in the attempt needs to be rotated.

C. The CMK that is used in the attempt is using the CMK's key ID instead of the CMK ARN.

D. The CMK that is used in the attempt is not enabled.

E. The CMK that is used in the attempt is using an alias.

ANSWER: A D

Explanation:

The CMK that is used in the attempt does not exist is a valid cause of failure because Systems Manager Parameter Store sends the requested encryption or decryption operation to AWS KMS using the key reference supplied for the SecureString parameter. AWS KMS must be able to resolve that reference to an existing KMS key in the applicable AWS Region and account context. If the key was deleted, the identifier is incorrect, or the request cannot locate the referenced key, KMS cannot perform the cryptographic operation required to store or retrieve the parameter value.

The CMK that is used in the attempt is not enabled is also a valid cause. An enabled KMS key is required for normal Encrypt and Decrypt operations. KMS key states control whether cryptographic operations are permitted; for example, a disabled key or a key pending deletion cannot be used to encrypt a new SecureString value or decrypt an existing one. Consequently, Parameter Store operations that depend on that key will fail until an appropriate, enabled key is available. AWS documents the operational effects of each KMS key state in the AWS KMS key states guide. Parameter Store's use of KMS keys for SecureString parameters is described in the SecureString parameter documentation.

QUESTION NO: 68

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

The company must create the break glass user. The company must log any activities of the break glass user and send the logs to a security team.

Which combination of solutions will meet these requirements? (Choose two.)

A. Create a local individual break glass IAM user for the security team. Create a trail in AWS CloudTrail that has Amazon CloudWatch Logs turned on. Use Amazon EventBridge to monitor local user activities.

B. Create a break glass EC2 key pair for the AWS account. Provide the key pair to the security team. Use AWS CloudTrail to monitor key pair activity. Send notifications to the security team by using Amazon Simple Notification Service (Amazon SNS).

C. Create a break glass IAM role for the account. Allow security team members to perform the AssumeRoleWithSAML operation. Create an AWS CloudTrail trail that has Amazon CloudWatch Logs turned on. Use Amazon EventBridge to monitor security team activities.

D. Create a local individual break glass IAM user on the operating system level of each workload instance. Configure unrestricted security groups on the instances to grant access to the break glass IAM users.

E. Configure AWS Systems Manager Session Manager for Amazon EC2. Configure an AWS CloudTrail filter based on Session Manager. Send the results to an Amazon Simple Notification Service (Amazon SNS) topic.

ANSWER: A E

Explanation:

Creating a local individual break glass IAM user for the security team establishes an emergency authentication path that is independent of the SAML identity provider. The IAM user can be protected with strong credentials and tightly scoped IAM permissions, including permissions required to access the workload through Systems Manager. AWS CloudTrail records the IAM user's AWS API activity, and delivering trail events to CloudWatch Logs makes those events available for monitoring. Amazon EventBridge can then evaluate CloudTrail events and route matching break-glass activity to a notification or response workflow for the security team.

Configuring AWS Systems Manager Session Manager for the EC2 instances provides controlled, auditable instance access without requiring inbound SSH or RDP access, public exposure, or EC2 key-pair distribution. The break glass IAM user can start only authorized sessions when an emergency occurs. CloudTrail records Session Manager API activity, including session initiation actions, while a CloudTrail-based filter and Amazon SNS notification path can alert the security team about this use. Together, the local IAM user supplies the SAML-independent account access requirement, and Session Manager supplies secure access to the immutable workload instances with centralized monitoring and notification. See the AWS IAM user documentation and the AWS Systems Manager Session Manager documentation.

QUESTION NO: 69

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

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

A. Create an identity policy that allows the sts: AssumeRole action in the AWS account that contains the resources. Attach the identity policy to the IAM user.

B. Ensure that the sts: AssumeRole action is allowed by the SCPs of the organization that owns the resources that the IAM user needs to access.

C. Create a role in the AWS account that contains the resources. Create an entry in the role's trust policy that allows the IAM user to assume the role. Attach the trust policy to the role.

D. Establish a trust relationship between the IAM user and the AWS account that contains the resources.

E. Create a role in the IAM user's AWS account. Create an identity policy that allows the sts: AssumeRole action. Attach the identity policy to the role.

F. Create an identity policy in the IAM user's AWS account that allows sts:AssumeRole on the role in the AWS account that contains the resources. Attach the policy to the IAM user.

ANSWER: C F

Explanation:

Cross-account role access requires permissions on both sides of the role assumption request. The account that owns the target resources must create the IAM role and configure that role's trust policy to trust the IAM user, or the user's source account. This resource-side trust relationship authorizes the specified external principal to call AWS Security Token Service and obtain credentials for the target role. The role's own permissions policies then determine which resources and actions are available after assumption.

The IAM user also needs an identity-based policy in the user's own account that permits `sts:AssumeRole` on the target role ARN. IAM evaluates this caller-side authorization before AWS STS can issue the role session credentials. Together, the trust policy in the target account and the caller's identity policy provide the required two-way authorization model for cross-account access. Organization-level guardrails, such as service control policies, must not deny the relevant actions, but they do not replace either required IAM authorization.

See the AWS IAM cross-account role tutorial at AWS IAM User Guide and the cross-account access guidance at AWS re:Post Knowledge Center.

QUESTION NO: 70

A company is hosting a static website on Amazon S3 The company has configured an Amazon CloudFront distribution to serve the website contents The company has associated an IAM WAF web ACL with the CloudFront distribution. The web ACL ensures that requests originate from the United States to address compliance restrictions.

THE company is worried that the S3 URL might still be accessible directly and that requests can bypass the CloudFront distribution

Which combination of steps should the company take to remove direct access to the S3 URL? (Select TWO.)

A. Select "Restrict Bucket Access" in the origin settings of the CloudFront distribution

B. Create an origin access identity (OAI) for the S3 origin

C. Update the S3 bucket policy to allow s3 GetObject with a condition that the IAM Referer key matches the secret value Deny all other requests

D. Configure the S3 bucket policy so that only the origin access identity (OAI) has read permission for objects in the bucket

E. Add an origin custom header that has the name Referer to the CloudFront distribution Give the header a secret value.

ANSWER: B D

Explanation:

Creating an origin access identity (OAI) for the S3 origin and configuring the S3 bucket policy so that only the origin access identity (OAI) has read permission for objects in the bucket establish CloudFront as the only authorized reader of the private S3 content. An OAI is a special CloudFront principal that S3 can recognize in a bucket policy. After the distribution is configured to use that identity, the bucket policy grants `s3:GetObject` to the OAI while withholding public and direct-user access. Consequently, a viewer who requests an object through CloudFront receives it because CloudFront retrieves it using the OAI; a viewer attempting the S3 object URL directly is not using that principal and is denied. This ensures that all successful viewer requests traverse CloudFront, where the AWS WAF web ACL and its geographic restrictions are evaluated. For a new implementation, AWS recommends using Origin Access Control (OAC), which uses SigV4 signing and is the modern replacement for OAI, but OAI is the applicable correct mechanism among the listed choices. See the CloudFront documentation for restricting access to an S3 origin and the Amazon S3 bucket policy examples.

QUESTION NO: 71

A company's data scientists want to create artificial intelligence and machine learning (AI/ML) training models by using Amazon SageMaker. The training models will use large datasets in an Amazon S3 bucket. The datasets contain sensitive information.

On average. the data scientists need 30 days to train models. The S3 bucket has been secured appropriately The companfs data retention policy states that all data that is older than 45 days must be removed from the S3 bucket.

Which action should a security engineer take to enforce this data retention policy?

A. Configure an S3 Lifecycle rule on the S3 bucket to delete objects after 45 days.

B. Create an AWS Lambda function to check the last-modified date of the S3 objects and delete objects that are older than 45 days. Create an S3 event notification to invoke the Lambda function for each PutObject operation.

C. Create an AWS Lambda function to check the last-modified date of the S3 objects and delete objects that are older than 45 days. Create an Amazon EventBridge rule to invoke the Lambda function each month.

D. Configure S3 Intelligent-Tiering on the S3 bucket to automatically transition objects to another storage class.

ANSWER: A

Explanation:

Configure an S3 Lifecycle rule on the S3 bucket to delete objects after 45 days. An Amazon S3 Lifecycle expiration rule is the managed S3 feature designed to enforce object-retention periods automatically. The rule can apply to all objects in the bucket or to a defined subset selected by prefix, object tags, or both. For an expiration action, S3 calculates an object's age from its creation date and automatically expires the object after the configured number of days. This directly implements the requirement that data older than 45 days be removed without requiring custom code, scheduled jobs, or operational intervention.

The 45-day expiration period also leaves sufficient time for the stated average 30-day SageMaker training duration. Lifecycle expiration is applied asynchronously by Amazon S3; expiration eligibility occurs at the configured age, and physical removal is performed by S3 thereafter. For a data-retention requirement, this managed control provides consistent policy-based enforcement at the storage layer. If the bucket is versioned, the lifecycle configuration should also include expiration handling for noncurrent object versions as appropriate, so historical versions are governed by the retention policy as well.

For implementation details, see the Amazon S3 Lifecycle management documentation and the S3 lifecycle expiration considerations.

QUESTION NO: 72

A security engineer needs to develop a process to investigate and respond to po-tential security events on a company's Amazon EC2 instances. All the EC2 in-stances are backed by Amazon Elastic Block Store (Amazon EBS). The company uses AWS Systems Manager to manage all the EC2 instances and has installed Systems Manager Agent (SSM Agent) on all the EC2 instances.

The process that the security engineer is developing must comply with AWS secu-rity best practices and must meet the following requirements:

• A compromised EC2 instance's volatile memory and non-volatile memory must be preserved for forensic purposes.

• A compromised EC2 instance's metadata must be updated with corresponding inci-dent ticket information.

• A compromised EC2 instance must remain online during the investigation but must be isolated to prevent the spread of malware.

• Any investigative activity during the collection of volatile data must be cap-tured as part of the process.

Which combination of steps should the security engineer take to meet these re-quirements with the LEAST operational overhead? (Select THREE.)

A. Gather any relevant metadata for the compromised EC2 instance. Enable ter-mination protection. Isolate the instance by updating the instance's secu-rity groups to restrict access. Detach the instance from any Auto Scaling groups that the instance is a member of. Deregister the instance from any Elastic Load Balancing (ELB) resources.

B. Gather any relevant metadata for the compromised EC2 instance. Enable ter-mination protection. Move the instance to an isolation subnet that denies all source and destination traffic. Associate the instance with the subnet to restrict access. Detach the instance from any Auto Scaling groups that the instance is a member of. Deregister the instance from any Elastic Load Balancing (ELB) resources.

C. Use Systems Manager Run Command to invoke scripts that collect volatile data.

D. Establish a Linux SSH or Windows Remote Desktop Protocol (RDP) session to the compromised EC2 instance to invoke scripts that collect volatile data.

E. Create a snapshot of the compromised EC2 instance's EBS volume for follow-up investigations. Tag the instance with any relevant metadata and inci-dent ticket information.

F. Create a Systems Manager State Manager association to generate an EBS vol-ume snapshot of the compromised EC2 instance. Tag the instance with any relevant metadata and incident ticket information.

ANSWER: A C E

Explanation:

Gathering the instance metadata, enabling termination protection, restricting the instance's security groups, removing it from Auto Scaling, and deregistering it from load balancers provides an effective live-containment workflow. The instance stays running so volatile evidence remains available, while security-group rules can immediately limit inbound and outbound communication to prevent malware propagation. Preserving contextual metadata before containment also supports a defensible incident record.

AWS Systems Manager Run Command is appropriate for collecting volatile evidence because the SSM Agent is already installed and the service enables centrally controlled, repeatable command execution without opening SSH or RDP access. Run Command activity can be audited through AWS CloudTrail, and command output can be retained in Amazon S3 or Amazon CloudWatch Logs, creating an investigation record. See AWS Systems Manager Run Command documentation.

Creating an EBS snapshot preserves point-in-time non-volatile disk evidence for later examination while the source instance remains online. Applying incident-ticket and related forensic tags to the EC2 instance associates the resource with the response workflow and supports tracking and automation. EBS snapshots are incremental, point-in-time backups of EBS volumes; AWS documents their forensic utility in its Amazon EBS snapshot documentation.

QUESTION NO: 73

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

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

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

A)

B)

C)

D)

A. Option

B. Option

C. Option

D. Option D

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

ANSWER: E

Explanation:

The policy that explicitly denies `guardduty:UpdateDetector`, `guardduty:DeleteDetector`, and `securityhub:DisableSecurityHub` is correct. An SCP attached at the organization root is inherited by every member account and organizational unit beneath that root. Because the default `FullAWSAccess` SCP remains attached, the

organization continues to establish a broad maximum permission boundary; this additional SCP narrows that boundary only for the specific actions that could disable the required security services.

GuardDuty can be disabled by updating a detector to turn it off, and deleting a detector also removes the GuardDuty configuration. Denying both actions prevents either path. Denying `securityhub:DisableSecurityHub` prevents Security Hub from being disabled. An explicit SCP deny applies regardless of permissions granted by IAM identity policies, permission sets, or resource policies, while IAM permissions for all actions not listed in the deny statement remain unaffected. Using `Resource` set to `*` is appropriate because these service-control actions do not support a narrower resource scope for this purpose.

See the AWS Organizations SCP documentation and the GuardDuty UpdateDetector API reference.

QUESTION NO: 74

A company has an application that uses dozens of Amazon DynamoDB tables to store data. Auditors find that the tables do not comply with the company's data protection policy.

The company's retention policy states that all data must be backed up twice each month: once at midnight on the 15th day of the month and again at midnight on the 25th day of the month. The company must retain the backups for 3 months.

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

A. Use the DynamoDB on-demand backup capability to create a backup plan. Configure a lifecycle policy to expire backups after 3 months.

B. Use AWS DataSync to create a backup plan. Add a backup rule that includes a retention period of 3 months.

C. Use AWS Backup to create a backup plan. Add a backup rule that includes a retention period of 3 months.

D. Set the backup frequency by using a cron schedule expression. Assign each DynamoDB table to the backup plan.

E. Set the backup frequency by using a rate schedule expression. Assign each DynamoDB table to the backup plan.

ANSWER: C D

Explanation:

Use AWS Backup to centrally define and apply a policy-driven backup plan for the DynamoDB tables. AWS Backup supports DynamoDB and allows a backup rule to specify the lifecycle of recovery points, including a deletion setting that retains each backup for approximately 3 months (for example, 90 days). This gives the company a consistent and auditable way to enforce retention across dozens of tables rather than configuring backups individually.

The required dates and time need a cron schedule expression because the policy is based on specific calendar days of each month: midnight on the 15th and midnight on the 25th. AWS Backup backup rules support cron expressions for this type of schedule; a suitable expression is `cron(0 0 15,25 * ? *)`, with schedules evaluated in UTC unless a backup plan time zone is configured. The tables must then be assigned to the backup plan, either by explicitly selecting the DynamoDB resources or through a tag-based resource assignment. This ensures every in-scope table receives backups according to the same schedule and retention configuration. See Creating a backup plan in AWS Backup and Assigning resources to a backup plan.

QUESTION NO: 75

An application team wants to use IAM Certificate Manager (ACM) to request public certificates to ensure that data is secured in transit. The domains that are being used are not currently hosted on Amazon Route 53

The application team wants to use an IAM managed distribution and caching solution to optimize requests to its systems and provide better points of presence to customers Thedistribution solution will use a primary domain name that is customized The distribution solution also will use several alternative domain names The certificates must renew automatically over an indefinite period of time

Which combination of steps should the application team take to deploy this architecture? (Select THREE.)

A. Request a certificate from ACM in the us-west-2 Region. Add the domain names that the certificate will secure.

B. Send an email message to the domain administrators to request validation of the domains for ACM.

C. Request validation of the domains for ACM through DNS. Insert CNAME records into each domain's DNS zone.

D. Create an Application Load Balancer for the caching solution. Select the newly requested certificate from ACM to be used for secure connections.

E. Create an Amazon CloudFront distribution for the caching solution. Configure the application endpoint as the origin, add the primary and alternative domain names as alternate domain names, and select the newly requested ACM certificate for secure connections.

F. Request a certificate from ACM in the us-east-1 Region. Add the domain names that the certificate will secure.

ANSWER: C E F

Explanation:

Requesting the ACM public certificate in the `us-east-1` Region is required because Amazon CloudFront can use ACM certificates only when they are issued or imported in US East (N. Virginia). The certificate request must include the custom primary domain name and every alternate domain name that CloudFront will serve. CloudFront then associates that certificate with the distribution so that viewers receive HTTPS connections for each configured alias.

DNS validation is appropriate even when the domains are hosted outside Route 53. ACM provides a unique CNAME record for each domain or wildcard name; the team adds those records at its existing authoritative DNS provider. As long as the validation CNAME records remain in DNS and the certificate remains in use, ACM can automatically renew an eligible publicly issued certificate without recurring manual validation. The CloudFront distribution should use the application endpoint as its origin, configure the custom domain names as alternate domain names (aliases), and select the validated ACM certificate. This combines CloudFront edge caching and points of presence with HTTPS for the organization’s custom hostnames. See ACM certificates for CloudFront distributions and ACM DNS validation and renewal.

QUESTION NO: 76

A company's Security Team received an email notification from the Amazon EC2 Abuse team that one or more of the company's Amazon EC2 instances may have been compromised

Which combination of actions should the Security team take to respond to (be current modem? (Select TWO.)

A. Open a support case with the IAM Security team and ask them to remove the malicious code from the affected instance

B. Respond to the notification and list the actions that have been taken to address the incident

C. Delete all IAM users and resources in the account

D. Detach the internet gateway from the VPC, remove all rules that contain 0.0.0.0/0 from the security groups, and create a NACL rule to deny all inbound traffic from the internet

E. Delete the identified compromised instances and delete any associated resources that the Security team did not create.

ANSWER: B D

Explanation:

“Respond to the notification and list the actions that have been taken to address the incident” is required because an EC2 Abuse notification indicates AWS has observed potentially malicious activity associated with the account’s resources. The Security team should acknowledge the notification and communicate concrete containment and remediation actions through the associated abuse case or contact channel. This creates a clear record of the response and enables AWS to continue handling the report appropriately.

“Detach the internet gateway from the VPC, remove all rules that contain 0.0.0.0/0 from the security groups, and create a NACL rule to deny all inbound traffic from the internet” provides immediate network containment. A suspected compromised instance should be isolated quickly to stop internet-originated access and limit possible command-and-control, scanning,

malware propagation, or exfiltration while evidence is preserved and the incident is investigated. These controls reduce exposure at multiple network layers: the VPC's internet connectivity, security-group rules, and subnet network ACLs. Containment is a core initial phase of AWS incident response, followed by investigation, eradication, and recovery. AWS guidance on responding to security incidents is available in the AWS Security Incident Response Guide and AWS's AWS Support documentation for abuse reports.

QUESTION NO: 77

A company is hosting a static website on Amazon S3 The company has configured an Amazon CloudFront distribution to serve the website contents The company has associated an IAM WAF web ACL with the CloudFront distribution. The web ACL ensures that requests originate from the United States to address compliance restrictions.

THE company is worried that the S3 URL might still be accessible directly and that requests can bypass the CloudFront distribution

Which combination of steps should the company take to remove direct access to the S3 URL? (Select TWO.)

A. Select "Restrict Bucket Access" in the origin settings of the CloudFront distribution

B. Create an origin access identity (OAI) for the S3 origin

C. Update the S3 bucket policy to allow s3 GetObject with a condition that the IAM Referer key matches the secret value Deny all other requests

D. Configure the S3 bucket policy so that only the origin access identity (OAI) has read permission for objects in the bucket

E. Add an origin custom header that has the name Referer to the CloudFront distribution Give the header a secret value.

ANSWER: B D

Explanation:

Creating an origin access identity (OAI) for the S3 origin and configuring the S3 bucket policy so that only the origin access identity (OAI) has read permission for objects in the bucket together establish CloudFront as the authorized path for retrieving the website's objects. An OAI is a special CloudFront identity that CloudFront uses when it requests content from the S3 origin. The bucket policy must then grant that identity `s3:GetObject` permission for the required object paths while removing public read access. As a result, viewers can request content through the CloudFront distribution, where the associated AWS WAF web ACL evaluates the request, but unauthenticated requests sent directly to the S3 REST endpoint are not authorized to read the objects. This enforces the intended access path and prevents direct-origin access from circumventing controls applied at CloudFront. OAI is the legacy CloudFront mechanism; AWS recommends Origin Access Control for new distributions, but OAI remains a valid solution and is the mechanism explicitly represented in these choices. See Restricting access to an Amazon S3 origin and Amazon S3 bucket policy overview.

QUESTION NO: 78

A company's cloud operations team is responsible for building effective security for IAM cross-account access. The team asks a security engineer to help troubleshoot why some developers in the developer account

(123456789012) in the developers group are not able to assume a cross-account role (ReadS3) into a production account (999999999999) to read the contents of an Amazon S3 bucket (productionapp). The two account policies are as follows:

Developer account 123456789012:
Developer group permissions:

```
{
    "Version": "2012-10-17",
    "Statement": {
        "Effect": "Allow",
        "Action": "sts:AssumeRole",
        "Resource": "arn:aws:iam::999999999999:role/ReadS3"
    }
}
```

Production account 999999999999:
Production account ReadS3 role policy:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": "s3:ListAllMyBuckets",
            "Resource": "*"
        },
        {
            "Effect": "Allow",
            "Action": [
                "s3:ListBucket",
                "s3:GetBucketLocation"
            ],
```

Production account ReadS3 role policy - trust relationship:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Principal": {
                "AWS": "arn:aws:iam::888888888888:root"
            },
            "Action": "sts:AssumeRole",
            "Condition": {
                "BoolIfExists": {
                    "aws:MultiFactorAuthPresent": "true"
                }
            }
        }
    ]
}
```

Which recommendations should the security engineer make to resolve this issue? (Select TWO.)

A. Ask the developers to change their password and use a different web browser.

B. Ensure that developers are using multi-factor authentication (MFA) when they log in to their developer account as the developer role.

C. Modify the production account ReadS3 role policy to allow the PutBucketPolicy action on the productionapp S3 bucket.

D. Update the trust relationship policy on the production account S3 role to allow the account number of the developer account.

E. Update the developer group permissions in the developer account to allow access to the productionapp S3 bucket.

ANSWER: B D

Explanation:

“Ensure that developers are using multi-factor authentication (MFA) when they log in to their developer account as the developer role” is required when the role-assumption authorization includes an MFA condition. AWS evaluates the MFA context keys from the temporary session that makes the `sts:AssumeRole` request. Therefore, developers must authenticate using MFA and use a session that carries MFA context before they can meet the condition and obtain credentials for the production role.

“Update the trust relationship policy on the production account S3 role to allow the account number of the developer account” establishes the required resource-based side of cross-account role assumption. The ReadS3 role’s trust policy in account `999999999999` must name a principal from account `123456789012`, such as that account’s root principal, with any appropriate conditions to limit which identities may assume the role. The developers’ identity-based permissions can then authorize `sts:AssumeRole`, while the production role’s permissions grant the resulting session read access to the intended S3 resources. Together, MFA-backed authentication and a valid cross-account trust relationship allow AWS STS to issue the production-role session credentials securely. See AWS IAM cross-account role access documentation and AWS guidance for requiring MFA for API access.