PassQueen
a.k.a dumpsqueen.com

AWS Certified Security – Specialty

Amazon Web Services SCS-C03

Version Demo

Total Demo Questions: 26

Total Premium Questions: 261

Buy Premium PDF

https://passqueen.com

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Threat Detection and Incident Response	55
Topic 2, Security Logging and Monitoring	40
Topic 3, Infrastructure Security	61
Topic 4, Identity and Access Management	43
Topic 5, Data Protection	44
Topic 6, Security Management and Governance	17
Topic 7, Mix Questions	1
Total	261

QUESTION NO: 1

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 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 SNS topic.

ANSWER: A E

Explanation:

Creating a local individual break-glass IAM user provides a distinct emergency access path that does not depend on the company's SAML identity provider or federation workflow. The IAM user can be tightly controlled with strong authentication, least-privilege permissions, and credentials stored under an emergency-access process. AWS CloudTrail records the user's AWS API activity, while delivery to CloudWatch Logs provides a centralized logging destination. Amazon EventBridge can then detect relevant CloudTrail events and route notifications or automated handling to the security team. This produces attributable and reviewable evidence whenever emergency credentials are used.

Configuring AWS Systems Manager Session Manager provides the appropriate controlled path for emergency access to the immutable EC2 workload. Session Manager avoids distributing SSH keys or opening inbound administrative ports, and access is authorized through IAM. Session Manager API activity, including session creation and termination, is available through CloudTrail. Filtering for Session Manager activity and publishing the resulting alert to an Amazon SNS topic ensures the security team receives prompt notification when break-glass access to an instance occurs. Together, these controls cover emergency AWS account access and auditable, security-team-visible instance access. See AWS IAM emergency access guidance and AWS Systems Manager Session Manager documentation.

QUESTION NO: 2

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

Which solution will meet these requirements?

A. Use AWS Backup to create backups of the EC2 instances and S3 buckets every hour. Create AWS CloudFormation templates that replicate existing architecture components. Use a Git repository to store the CloudFormation templates alongside application configuration code.

B. Use AWS Backup to create backups of the EBS volumes and S3 objects every day. Use Amazon Security Lake to create a centralized data lake for AWS CloudTrail logs and VPC flow logs. Use the logs for automated response.

C. Use Amazon Security Lake to create a centralized data lake for AWS CloudTrail logs and VPC flow logs. Use the logs for automated response. Enable AWS Security Hub to establish a single location for recovery procedures. Create AWS CloudFormation templates that replicate existing architecture components. Use a Git repository to store the CloudFormation templates alongside application configuration code.

D. Create EBS snapshots every 4 hours. Enable Amazon GuardDuty Malware Protection. Create automation to immediately restore the most recent snapshot for any EC2 instances that produce an Execution:EC2/MaliciousFile finding in GuardDuty.

ANSWER: A

Explanation:

Use AWS Backup to create backups of the EC2 instances and S3 buckets every hour. Create AWS CloudFormation templates that replicate existing architecture components. Use a Git repository to store the CloudFormation templates alongside application configuration code. This solution establishes an hourly recovery point for the application's persistent data, limiting potential data loss to no more than 1 hour and therefore meeting the stated RPO. AWS Backup centrally manages backup plans, schedules, retention, and recovery for supported resources, including Amazon EC2 and Amazon S3. EC2 backups protect the attached EBS-backed data through the backup workflow, while S3 backups provide a recoverable copy of the image and music objects.

The recovery requirement also includes returning the service to normal operation after a ransomware incident. CloudFormation templates provide repeatable infrastructure as code for rebuilding the load-balanced EC2 environment and its dependent configuration. Keeping those templates and application configuration in version control preserves a known-good, reviewable deployment definition that can be used during recovery. Together, current hourly backups and reproducible infrastructure deployment enable restoration of both data and service architecture after an attacker has bypassed the primary security controls. See the AWS Backup Developer Guide and the AWS CloudFormation User Guide.

QUESTION NO: 3

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

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

E.

)

A. Encrypt all AWS CloudTrail logs.

B. Turn on Amazon GuardDuty.

C. Change the password for all IAM users.

D. Rotate or delete all AWS access keys.

E. Take snapshots of all Amazon Elastic Block Store (Amazon EBS) volumes.

F. Delete any resources that are unrecognized or unauthorized.

ANSWER: B D F

Explanation:

Turning on Amazon GuardDuty, rotating or deleting all AWS access keys, and deleting any resources that are unrecognized or unauthorized are appropriate containment actions for this active compromise. GuardDuty provides continuous threat detection by analyzing AWS logs and events, helping the security team identify related malicious activity, compromised

credentials, suspicious API calls, and affected resources that might not yet have been discovered. It also produces findings that can be routed to incident-response workflows.

Because an IAM user performed an unauthorized action, access keys must be treated as potentially compromised. Rotating keys replaces credentials with new secrets, while deleting keys immediately prevents their continued use. This limits an attacker's ability to create additional resources, host further malware, or access account data. The security team should ensure that legitimate workloads are updated to use replacement credentials as part of the containment process.

Deleting unrecognized or unauthorized resources directly stops the known abuse, including removing the unauthorized bucket and any related malicious infrastructure. AWS incident-response guidance prioritizes containment and eradication of attacker-controlled resources while investigating the scope of activity through logging and detection services. See the AWS Security Incident Response Guide and the Amazon GuardDuty User Guide.

QUESTION NO: 4

A company uses AWS Security Hub to aggregate findings from its AWS accounts. The company needs to automatically isolate an Amazon EC2 instance when a high severity finding identifies the instance as potentially compromised. The company must preserve the instance for investigation and minimize disruption to other application instances.

Which combination of steps should the company take? (Select TWO.)

A. Create an Amazon EventBridge rule that matches the required Security Hub findings and invokes an AWS Lambda function.

B. Configure a Security Hub automation rule to replace the security group on affected EC2 instances.

C. Configure the Lambda function to apply a restrictive security group and remove the affected instance from application traffic management.

D. Configure Amazon SNS to send an email notification for each high severity finding.

E. Configure AWS Config to terminate any EC2 instance that receives a high severity finding.

ANSWER: A C

Explanation:

AWS Security Hub findings are delivered to Amazon EventBridge. An EventBridge rule can filter for the required finding severity, product, finding type, and affected EC2 resource, and then invoke a Lambda function without requiring manual intervention.

The Lambda function can apply a restrictive isolation security group to the affected instance and remove the instance from the application load balancer target group or Auto Scaling group as appropriate. This contains the affected instance while leaving it running for evidence collection and allowing healthy application instances to continue serving traffic.

Security Hub automation rules can update finding fields such as workflow status or severity, but they do not directly perform EC2 network containment. Sending notifications to an Amazon SNS topic requires a responder to act manually and does not meet the automated-response requirement.

QUESTION NO: 5

A company has multiple accounts in the AWS Cloud. Users in the developer account need to have access to specific resources in the production account.

What is the MOST secure way to provide this access?

A. Create one IAM user in the production account. Grant the appropriate permissions to the resources that are needed. Share the password only with the users that need access.

B. Create cross-account access with an IAM role in the developer account. Grant the appropriate permissions to this role. Allow users in the developer account to assume this role to access the production resources.

C. Create cross-account access with an IAM user account in the production account. Grant the appropriate permissions to this user account. Allow users in the developer account to use this user account to access the production resources.

D. Create cross-account access with an IAM role in the production account. Grant the appropriate permissions to this role. Allow users in the developer account to assume this role to access the production resources.

ANSWER: D

Explanation:

Create cross-account access with an IAM role in the production account. Grant the appropriate permissions to this role. Allow users in the developer account to assume this role to access the production resources. This is the AWS-recommended cross-account access pattern because the production account retains control of permissions to its resources while explicitly defining which trusted principals from the developer account can obtain access. The production role's trust policy authorizes the developer-account user, group, or—preferably—an IAM role to call AWS Security Token Service `AssumeRole`. IAM then issues temporary, short-lived credentials for the role session rather than requiring shared or long-term production credentials. The production role's permissions policy can be scoped precisely to only the required resources and actions, enforcing least privilege and maintaining separation of duties between accounts. Additional controls, such as conditions on the trust policy, external identity attributes, MFA requirements, session tags, and limited session duration, can further constrain access and provide clear audit records in AWS CloudTrail. AWS documents cross-account role assumption and the respective trust-policy and permissions-policy responsibilities at IAM cross-account roles and AWS STS AssumeRole.

QUESTION NO: 6 - (SIMULATION)

A security engineer needs to prepare for a security audit of an AWS account.

Select the correct AWS resource from the following list to meet each requirement. Select each resource one time or not at all. (Select THREE.)

• AWS Artifact reports

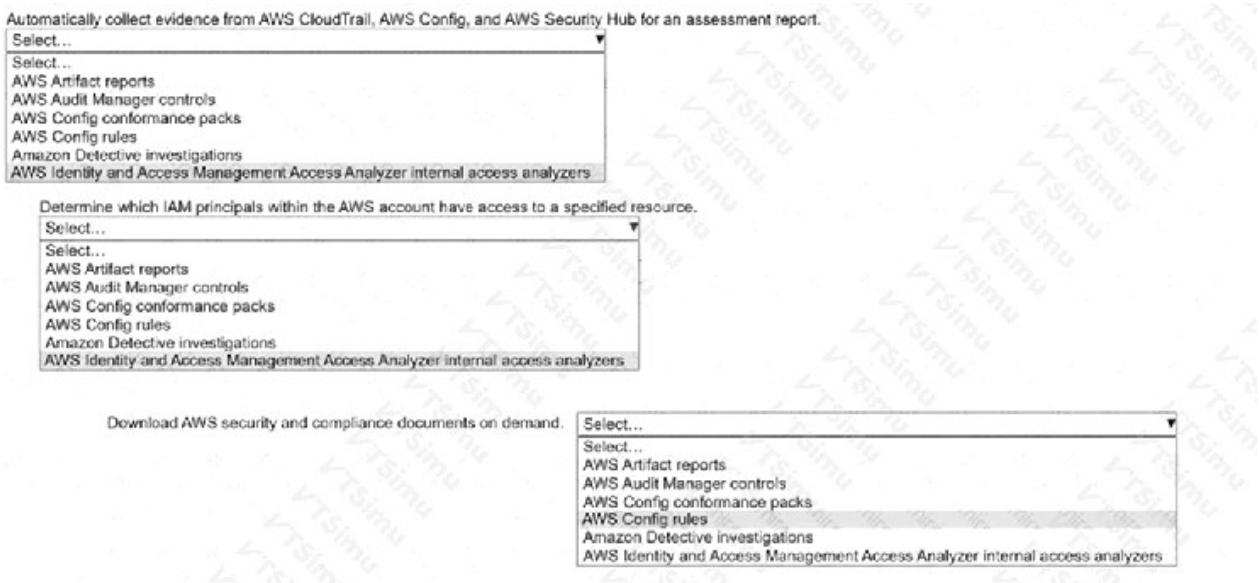
• AWS Audit Manager controls

• AWS Config conformance packs

• AWS Config rules

• Amazon Detective investigations

• AWS Identity and Access Management Access Analyzer internal access analyzers



ANSWER: See the explanation for the answer

Explanation:

Select the following three resources:

Automatically collect evidence from AWS CloudTrail, AWS Config, and AWS Security Hub for an assessment report:
AWS Audit Manager controls

AWS Audit Manager uses controls in an assessment framework to automatically collect and organize evidence from supported AWS services, including AWS CloudTrail, AWS Config, and AWS Security Hub.

Determine which IAM principals within the AWS account have access to a specified resource:
AWS Identity and Access Management Access Analyzer internal access analyzers

An IAM Access Analyzer internal access analyzer identifies internal principals, such as IAM users and roles, that have access to supported resources.

Download AWS security and compliance documents on demand:
AWS Artifact reports

AWS Artifact provides on-demand AWS compliance reports and security documents, such as SOC reports and ISO certifications.

QUESTION NO: 7

A security engineer needs to configure DDoS protection for a Network Load Balancer (NLB) with an Elastic IP address. The security engineer wants to set up an AWS WAF web ACL with a rate-based rule statement to protect the NLB.

The security engineer needs to determine a rate limit that will not block legitimate traffic. The security engineer has configured the rule statement to aggregate based on the source IP address.

How should the security engineer configure the rule to protect the NLB?

A. Configure the rule to use theCountaction.

B. Configure the rule to use theBlockaction.

C. Configure the rule to use theMonitoraction.

D. Configure the rule to use theAllowaction.

E. AWS WAF cannot be associated with an NLB. Protect the Elastic IP address with AWS Shield Advanced instead.

ANSWER: E

Explanation:

AWS WAF cannot be associated with an NLB. Protect the Elastic IP address with AWS Shield Advanced instead.
AWS WAF web ACLs can be associated only with supported application-layer resources, such as an Application Load Balancer, Amazon CloudFront distribution, Amazon API Gateway API, AWS AppSync API, Amazon Cognito user pool, AWS Verified Access instance, or AWS Amplify application. A Network Load Balancer is not a supported AWS WAF association target, so a rate-based rule cannot inspect, count, or block requests sent to the NLB, regardless of its source-IP aggregation configuration or rule action.

For DDoS protection of an NLB that uses Elastic IP addresses, AWS Shield Advanced is the applicable AWS service. Shield Advanced supports protecting Elastic IP addresses and provides enhanced DDoS detection and mitigation, access to the AWS Shield Response Team for eligible events, and DDoS-related cost protection. The engineer should create a Shield Advanced protection for the Elastic IP address or addresses used by the NLB. AWS Shield Standard also provides automatic baseline network and transport-layer DDoS protection, while Shield Advanced adds the enhanced features appropriate for workloads requiring additional DDoS safeguards. See the AWS WAF resource association documentation and the AWS Shield protected resources documentation.

QUESTION NO: 8

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? (Select 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 the Lambda function with the Amazon Cognito user pool provides the required synchronous custom validation point in the registration workflow. Cognito invokes this trigger before it completes the sign-up operation, allowing the function to apply company-specific controls such as validating an allowlist, consulting a fraud-detection service, assessing IP-address or device-risk signals, or enforcing location-related rules. The function can allow the registration to continue when validation succeeds or reject the sign-up by returning an error, directly meeting the requirement to accept or deny requests based on custom logic.

Using an AWS WAF web ACL with a geographic match rule and associating that web ACL with the Cognito user pool adds an effective perimeter control for the Hosted UI. A geo match rule can identify request origin countries and block traffic that does not originate in France. This reduces the volume of non-French fraudulent registration attempts before they reach the Cognito sign-up flow, while the pre sign-up trigger retains the flexibility to make the final custom validation decision for requests that reach Cognito. Together, these controls provide both broad geographic filtering and detailed application-specific registration enforcement. See the Amazon Cognito pre sign-up Lambda trigger documentation and the AWS WAF web ACL association documentation.

QUESTION NO: 9

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

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

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

A. Create a service role that has a composite principal that contains each service that needs the necessary permissions.

B. Create a service role that has cloudformation.amazonaws.com as the service principal.

C. Add policies that reference each CloudFormation stack ARN.

D. Add policies that reference the ARNs of each AWS service that requires permissions.

E. Update each stack to use the service role.

F. Add a policy to each member role to allow the iam:PassRole action for the service role.

ANSWER: B E F

Explanation:

A CloudFormation service role provides a consistent, centralized permission set for provisioning and updating all resources in a stack. Creating a service role that has cloudformation.amazonaws.com as the service principal establishes the required trust relationship: CloudFormation can assume the role and use its permissions while performing stack operations. This avoids making successful deployment dependent on the individual permissions available to whichever team member initiates the operation.

Updating each stack to use the service role is necessary because CloudFormation uses the configured service role for future stack actions. Once configured, CloudFormation continues to use that role for the stack rather than relying on the caller's permissions to create, modify, or delete the underlying resources. The service role should be granted only the resource actions required by the stack, following least privilege.

Adding permission for iam:PassRole on the specific service role to each member role allows authorized team members to supply or use that role during CloudFormation stack operations. Restricting this permission to the designated role ARN prevents users from passing arbitrary, more-privileged roles. Together, these controls provide reliable deployments without granting every developer direct permissions to all deployed AWS resources. See AWS CloudFormation service roles and IAM PassRole guidance.

QUESTION NO: 10

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

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

A. Grant least privilege access to the organization ' s management account.

B. Create a new IAM Identity Center directory in the organization ' s management account.

C. Set up a second AWS Region in the organization ' s management account.

D. Create permission sets for use only in the organization ' s management account.

E. Create IAM users for use only in the organization ' s management account.

F. Create user assignments only in the organization ' s management account.

ANSWER: A D F

Explanation:

Before delegated administration is configured, the management account needs an appropriately authorized IAM Identity Center administrator who can perform the initial configuration and register a member account for delegated administration. Granting least privilege access to the organization ' s management account ensures that this administrative access is limited to the required IAM Identity Center and AWS Organizations actions, rather than relying on unrestricted management-account access.

Creating permission sets for use only in the organization ' s management account provides the permission model that IAM Identity Center will apply when an administrator signs in to the management account through the AWS access portal. Permission sets define the IAM permissions that are provisioned into the target AWS account for federated users. Creating user assignments only in the organization ' s management account then associates the intended Identity Center users or groups with those permission sets and the management account. Together, the permission set and assignment establish who can enter the management account and carry out the initial delegated-administration configuration.

AWS documents that IAM Identity Center uses permission sets and account assignments to provide workforce access to AWS accounts, and that delegated administration is configured from an already enabled IAM Identity Center instance in an AWS Organization. See Delegated administration in IAM Identity Center and Create and manage permission sets.

QUESTION NO: 11

A company stores highly sensitive documents in an Amazon S3 bucket. The company needs AWS CloudTrail to log all read and write operations on objects only in the `restricted/` prefix of that bucket. The company does not want to incur data event charges for activity in other prefixes of the bucket.

Which solution will meet these requirements?

A. Configure a CloudTrail management event selector for read and write events on the S3 bucket.

B. Enable Amazon S3 server access logging on the bucket and configure the log destination to use the `restricted/` prefix.

C. Configure a CloudTrail advanced event selector for S3 object data events whose resource ARN starts with the ARN of the bucket's `restricted/` prefix. Include both read and write activity.

D. Enable CloudTrail data events for all objects in the S3 bucket. Use an Amazon EventBridge rule to discard events outside the `restricted/` prefix.

ANSWER: C

Explanation:

Configure a CloudTrail advanced event selector that includes Amazon S3 object data events whose resource ARN begins with the ARN for the required bucket prefix, and configure the selector for both read and write activity. Advanced event selectors provide granular filtering for CloudTrail data events, including filtering S3 object activity by resource ARN prefix.

A management event selector would not record object-level `GetObject` and `PutObject` activity. Enabling data events for the entire bucket would capture the required activity but would also capture events for every other prefix. Amazon S3 server access logging records access requests but does not provide the CloudTrail data-event audit record required by the scenario.

QUESTION NO: 12

A security engineer has designed a VPC to segment private traffic from public traffic. The VPC includes two Availability Zones. The security engineer has provisioned each Availability Zone with one private subnet and one public subnet. The security engineer has created three route tables for use with the environment. One route table is for the public subnets, and two route tables are for the private subnets (one route table for the private subnet in each Availability Zone).

The security engineer discovers that all four subnets are attempting to route traffic out through the internet gateway that is attached to the VP

C.

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

A. Verify that a NAT gateway has been provisioned in the public subnet in each Availability Zone.

B. Verify that a NAT gateway has been provisioned in the private subnet in each Availability Zone.

C. Modify the route tables that are associated with each of the public subnets. Create a new route for local destinations to the VPC CIDR range.

D. Modify the route tables that are associated with each of the private subnets. Create a new route for the destination 0.0.0.0/0. Specify the NAT gateway in the public subnet of the same Availability Zone as the target of the route.

E. Modify the route tables that are associated with each of the private subnets. Create a new route for the destination 0.0.0.0/0. Specify the internet gateway as the target of the route.

ANSWER: A D

Explanation:

Private subnets must not use an internet gateway as the target of their default route. An internet gateway provides direct internet connectivity for resources that have public IPv4 addresses and are placed in subnets whose route tables direct internet-bound traffic to that gateway. To preserve the intended separation, outbound internet access from private-subnet workloads should instead traverse a NAT gateway. The NAT gateway is deployed in a public subnet, where its route table has a default route to the internet gateway, and it uses an Elastic IP address for internet communication.

For resilient multi-AZ operation, a NAT gateway should be provisioned in the public subnet in each Availability Zone. Each private subnet's dedicated route table should then route `0.0.0.0/0` to the NAT gateway located in the same Availability Zone. This design allows private resources to initiate outbound connections for activities such as patching and package downloads, while preventing unsolicited inbound connections directly from the internet. Using a same-AZ NAT gateway also avoids unnecessary cross-AZ traffic and prevents an Availability Zone failure from disrupting outbound connectivity for private resources in the other zone. AWS documents this architecture as the standard private-subnet NAT gateway routing pattern. See the Amazon VPC NAT gateway scenarios documentation and Amazon VPC route table documentation.

QUESTION NO: 13

A company needs to deploy AWS CloudFormation templates that configure sensitive database credentials. The company already uses AWS Key Management Service (AWS KMS) and AWS Secrets Manager.

Which solution will meet the requirements?

A. Use a dynamic reference in the CloudFormation template to reference the database credentials in Secrets Manager.

B. Use encrypted parameters in the CloudFormation template.

C. Use SecureString parameters to reference Secrets Manager.

D. Use SecureString parameters encrypted by AWS KMS.

ANSWER: A

Explanation:

Use a dynamic reference in the CloudFormation template to reference the database credentials in Secrets Manager. CloudFormation dynamic references let a template identify a Secrets Manager secret without embedding the secret value in the template itself. CloudFormation resolves the reference only when it creates or updates the resource that needs the credential, such as a database configuration. This approach keeps the sensitive value out of the template body and avoids persisting it in CloudFormation stack metadata, change sets, or outputs.

The dynamic-reference syntax for Secrets Manager is `{{resolve:secretsmanager:secret-id:SecretString:json-key}}`, where the secret identifier can be a name or ARN and the JSON key selects a value such as a username or password. The CloudFormation service role, or the identity used for deployment, must have permission to retrieve the secret through `secretsmanager:GetSecretValue`; access to the KMS key is also required when the secret uses a customer managed key. This uses Secrets Manager as the source of truth for credentials while maintaining least-privilege access through IAM and KMS policies. AWS documents dynamic references specifically for securely supplying Secrets Manager values to CloudFormation-managed resources. See CloudFormation dynamic references for Secrets Manager and AWS Secrets Manager best practices.

QUESTION NO: 14

A company runs Amazon Elastic Kubernetes Service (Amazon EKS) clusters in multiple AWS accounts. The company needs centralized threat detection for suspicious Kubernetes API activity and for suspicious runtime behavior in containers. The company uses AWS Organizations and has a dedicated security account.

Which combination of actions should the company take? (Select TWO.)

A. Designate the security account as the Amazon GuardDuty delegated administrator and enable GuardDuty EKS Protection for the EKS clusters.

B. Enable Amazon GuardDuty Runtime Monitoring for the EKS workloads.

C. Enable AWS Config recording for Amazon EKS resources in each account.

D. Enable Amazon Inspector scanning for Amazon ECR repositories.

E. Create Amazon CloudWatch alarms for changes to EKS cluster security groups.

ANSWER: A B

Explanation:

Amazon GuardDuty EKS Protection analyzes EKS audit log activity to identify suspicious Kubernetes API operations, such as unusual credential use or potentially malicious changes to Kubernetes resources. It addresses the control-plane activity requirement.

Amazon GuardDuty Runtime Monitoring provides runtime threat detection for supported EKS workloads. It monitors runtime activity and produces findings for suspicious behavior in containers, addressing the requirement for workload-level detection.

A delegated GuardDuty administrator in the dedicated security account can centrally manage GuardDuty member accounts and findings. AWS Config records resource configuration changes but does not inspect Kubernetes audit activity or container runtime behavior. Amazon Inspector identifies software vulnerabilities rather than suspicious runtime activity.

QUESTION NO: 15

A security engineer needs to develop a process to investigate and respond to potential security events on a company's Amazon EC2 instances. All the EC2 instances are backed by Amazon EBS. The company uses AWS Systems Manager to manage all the EC2 instances and has installed Systems Manager Agent on all the EC2 instances.

The process that the security engineer is developing must comply with AWS security 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 incident 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 captured as part of the process.

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

A. Gather any relevant metadata for the compromised EC2 instance. Enable termination protection. Isolate the instance by updating the instance's security 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 resources.

B. Gather any relevant metadata for the compromised EC2 instance. Enable termination 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 resources.

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

D. Establish a Linux SSH or Windows Remote Desktop Protocol 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 incident ticket information.

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

ANSWER: A C E

Explanation:

The appropriate process combines network containment, auditable volatile-data acquisition, and durable evidence preservation. “Gather any relevant metadata for the compromised EC2 instance. Enable termination protection. Isolate the instance by updating the instance’s security 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 resources.” keeps the workload running while restricting its communications and preventing normal scaling or load-balancing behavior from interfering with containment. This supports live investigation without exposing other resources to potential malware propagation.

“Use Systems Manager Run Command to invoke scripts that collect volatile data.” provides a managed, noninteractive mechanism for memory and other live-data collection. Run Command records command execution details and output in Systems Manager, supporting an audit trail for investigative actions. “Create a snapshot of the compromised EC2 instance’s EBS volume for follow-up investigations. Tag the instance with any relevant metadata and incident ticket information.” preserves a point-in-time copy of non-volatile disk evidence and associates the live resource with the incident record. Together, these actions provide containment, evidence collection, traceability, and low operational overhead. See the AWS Run Command documentation and AWS guidance for Amazon EBS snapshots.

QUESTION NO: 16

A company runs a global ecommerce website that is hosted on AWS. The company uses Amazon CloudFront to serve content to its user base. The company wants to block inbound traffic from a specific set of countries to comply with recent data regulation policies.

Which solution will meet these requirements MOST cost-effectively?

A. Create an AWS WAF web ACL with an IP match condition to deny the countries ' IP ranges. Associate the web ACL with the CloudFront distribution.

B. Create an AWS WAF web ACL with a geo match condition to deny the specific countries. Associate the web ACL with the CloudFront distribution.

C. Use the geo restriction feature in CloudFront to deny the specific countries.

D. Use geolocation headers in CloudFront to deny the specific countries.

ANSWER: C

Explanation:

Use the geo restriction feature in CloudFront to deny the specific countries. CloudFront provides built-in geographic restrictions, also called geoblocking, that can be configured as a deny list for the countries from which viewers must not access a distribution. CloudFront determines the viewer’s country from the source IP address of the request and blocks a restricted viewer at the CloudFront edge before content is served. This directly satisfies a country-based regulatory access-control requirement without requiring custom code, country IP-range maintenance, or an additional security service.

This is the most cost-effective approach because geographic restriction is a native CloudFront distribution setting. The company can configure the relevant ISO country codes in the distribution’s restrictions settings and allow CloudFront to perform the location decision automatically. It is well suited when the sole policy requirement is to allow or deny access according to the viewer’s country. The capability applies to requests for the distribution, including cached and origin-fetched content, enabling consistent enforcement for the ecommerce site’s audience.

AWS documents geographic restrictions as the CloudFront mechanism for restricting content delivery by viewer location and supports both allow-list and deny-list modes. For implementation details, see Restricting the geographic distribution of your content and the CloudFront GeoRestriction API reference.

QUESTION NO: 17

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

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

A. Export the KMS key material to an on-premises hardware security module (HSM). Give the application team access to the key material.

B. Edit the key policy that grants the security team access to the KMS keys by adding the application team as principals. Revert this change when the application team no longer needs access.

C. Create a key grant to allow the application team to use the KMS keys. Revoke the grant when the application team no longer needs access.

D. Create a new KMS key by generating key material on premises. Import the key material to AWS KMS whenever the application team needs access. Grant the application team permissions to use the key.

ANSWER: C

Explanation:

Create a key grant to allow the application team to use the KMS keys. Revoke the grant when the application team no longer needs access. AWS KMS grants are purpose-built for delegating a defined set of KMS permissions, such as `Decrypt`, `Encrypt`, and `GenerateDataKey`, to a specific AWS principal without changing the KMS key policy. This lets the security team retain key-administration control while authorizing only the application process's required key-usage actions. A grant can also include constraints that limit cryptographic operations to a particular encryption context, providing more precise least-privilege access when appropriate.

For occasional temporary access, grants have lower operational overhead than repeatedly editing and rolling back a key policy. The security team can programmatically create the grant when the process needs access and retire it afterward with `RevokeGrant`. The grant mechanism is designed for this type of delegated authorization and supports auditing through AWS CloudTrail KMS events. The application process should use an IAM role as the grant grantee principal, rather than relying on broad human-user access, so credentials are temporary and permissions are attributable to the workload. See the AWS documentation on AWS KMS grants and managing grants.

QUESTION NO: 18

A company has AWS accounts in an organization in AWS Organizations. The organization includes a dedicated security account.

All AWS account activity across all member accounts must be logged and reported to the dedicated security account. The company must retain all the activity logs in a secure storage location within the dedicated security account for2 years.No changes or deletions of the logs are allowed.

Which combination of steps will meet these requirements with theLEAST operational overhead? (Select TWO.)

A. In the dedicated security account, create an Amazon S3 bucket with S3 Object Lock in compliance mode and a default retention period of 2 years. Set the bucket policy to allow AWS CloudTrail to write to the S3 bucket.

B. In the dedicated security account, create an Amazon S3 bucket. Configure S3 Object Lock incompliance modewith a retention period of 2 years. Set the bucket policy to allow the organization'smember accountsto write to the S3 bucket.

C. In the dedicated security account, create an Amazon S3 bucket with an S3 Lifecycle configuration that expires objects after 2 years. Allow member accounts to write to the bucket.

D. Create anAWS CloudTrail organization trail. Configure logs to be delivered to the Amazon S3 bucket in the dedicated security account.

E. Turn on AWS CloudTrail in each account and forward logs to the dedicated security account by using AWS Lambda and Amazon Data Firehose.

ANSWER: A D

Explanation:

Create an AWS CloudTrail organization trail and deliver its logs to a centralized Amazon S3 bucket in the dedicated security account. An organization trail automatically applies to all current and future member accounts, which avoids the ongoing administration of configuring and maintaining separate trails in every account. The trail can be created by the organization's management account or an appropriately registered delegated administrator, while the destination bucket remains isolated in the security account. AWS CloudTrail must be granted access through the bucket policy as the service principal that writes the log files.

The destination bucket should have S3 Object Lock enabled with a default retention period of 2 years in compliance mode. Compliance mode provides WORM retention: protected log objects cannot be overwritten or deleted by any principal, including the AWS account root user, until their retention period expires. Applying the retention by default ensures that each CloudTrail log object receives the required protection as it arrives, without per-object operational work. Together, an organization trail and an Object Lock-enabled centralized bucket provide organization-wide collection, secure cross-account storage, and immutable retention with minimal overhead. See AWS CloudTrail organization trails and Amazon S3 Object Lock documentation.

QUESTION NO: 19

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

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

E.

)

A. Create a new customer managed key in AWS Key Management Service (AWS KMS).

B. Change the SSE-S3 configuration on the S3 bucket to server-side encryption with customer-provided keys (SSE-C).

C. Configure the PHP SDK to use the SSE-S3 key before upload.

D. Create an AWS managed key for Amazon S3 in AWS KMS.

E. Change the SSE-S3 configuration on the S3 bucket to server-side encryption with AWS KMS managed keys (SSE-KMS).

F. Change all the S3 objects in the bucket to use the new encryption key.

ANSWER: A E F

Explanation:

Full control of the encryption keys requires using server-side encryption with AWS KMS keys (SSE-KMS) and selecting a customer managed KMS key. The security engineer must create that customer managed key in AWS KMS, where the company can administer the key policy, grants, rotation configuration, enablement state, and eventual deletion schedule. The bucket's default encryption configuration must then be changed from SSE-S3 to SSE-KMS and configured to use the newly created customer managed key. This ensures that newly uploaded objects are encrypted under the company-controlled KMS key by default. Changing a bucket encryption setting does not retroactively alter the encryption of objects already stored in the bucket. Therefore, all existing objects must be rewritten or copied so that S3 encrypts the replacement objects with SSE-KMS using the new customer managed key. For a large bucket, this migration can be performed at scale with Amazon S3 Batch Operations. After migration, KMS key usage can be monitored through AWS CloudTrail, providing auditable records of encryption-key activity. AWS documents SSE-KMS behavior and default bucket encryption at Using SSE-KMS with Amazon S3, and describes the administrative controls available for customer managed KMS keys at AWS KMS customer managed keys.

QUESTION NO: 20

A company is investigating an increase in its AWS monthly bill. The company discovers that bad actors compromised some Amazon EC2 instances and served webpages for a large email phishing campaign. A security engineer must implement a solution to monitor for cost increases in the future to help detect malicious activity.

Which solution will offer the company the EARLIEST detection of cost increases?

A. Create an Amazon EventBridge rule that invokes an AWS Lambda function hourly. Program the Lambda function to download an AWS usage report from AWS Data Exports about usage of all services. Program the Lambda function to analyze the report and to send a notification when anomalies are detected.

B. Create a cost monitor in AWS Cost Anomaly Detection. Configure an individual alert to notify an Amazon Simple Notification Service (Amazon SNS) topic when the percentage above the expected cost exceeds a threshold.

C. Review AWS Cost Explorer daily to detect anomalies in cost from prior months. Review the usage of any services that experience a significant cost increase from prior months.

D. Capture VPC flow logs from the VPC where the EC2 instances run. Use a third-party network analysis tool to analyze the flow logs and to detect anomalies in network traffic that might increase cost.

ANSWER: B

Explanation:

Create a cost monitor in AWS Cost Anomaly Detection. Configure an individual alert to notify an Amazon Simple Notification Service (Amazon SNS) topic when the percentage above the expected cost exceeds a threshold. AWS Cost Anomaly Detection is purpose-built to identify unusual spending rapidly by using machine learning to establish expected cost patterns from historical AWS usage and then continuously evaluating new cost data against that expected baseline. A cost monitor can cover the relevant AWS services, linked accounts, cost allocation tags, or cost categories, allowing the security team to focus alerts on resources associated with the affected environment. An individual alert subscription can use a percentage-based threshold and send notifications through Amazon SNS as soon as an anomaly exceeds the configured impact threshold. This gives responders an automated signal of unexpected spend without waiting for a scheduled manual review or building and operating a custom reporting and analysis workflow. In this scenario, compromised EC2 instances may create unusual compute, storage, or data-transfer charges. Detecting the resulting deviation in spend promptly helps the company investigate malicious use and contain it before charges grow further. AWS documents that Cost Anomaly Detection continuously monitors costs and usage and sends alerts when anomalous spend is detected. See AWS Cost Anomaly Detection and AWS Cost Anomaly Detection alerts.

QUESTION NO: 21

A corporate cloud security policy states that communications between the company ' s VPC and KMS must travel entirely within the AWS network and not use public service endpoints.

Which combination of the following actions MOST satisfies this requirement? (Select TWO.)

A. Add theaws:sourceVpcecondition to the AWS KMS key policy referencing the company ' s VPC endpoint ID.

B. Remove the VPC internet gateway from the VPC and add a virtual private gateway to the VPC to prevent direct, public internet connectivity.

C. Create a VPC endpoint for AWS KMS withprivate DNS enabled.

D. Use the KMS Import Key feature to securely transfer the AWS KMS key over a VPN.

E. Add the following condition to the AWS KMS key policy: " aws:SourceIp " : " 10.0.0.0/16 " .

ANSWER: A C

Explanation:

Creating a VPC endpoint for AWS KMS with private DNS enabled provides the required private network path. AWS KMS supports interface VPC endpoints powered by AWS PrivateLink. When private DNS is enabled, workloads in the VPC can continue to use the standard regional KMS hostname, such as `kms.us-east-1.amazonaws.com`, but that name resolves to the private IP addresses of the interface endpoint. KMS API requests therefore traverse the endpoint and AWS's private network rather than requiring access to the public KMS service endpoint. The endpoint must also be available in the subnets and Availability Zones from which the workloads make KMS calls.

Adding the `aws:sourceVpce` condition to the AWS KMS key policy referencing the company ' s VPC endpoint ID complements that network path with key-level authorization enforcement. AWS KMS supports the global `aws:SourceVpce` condition key in key policies and IAM policies. Restricting access to the specific endpoint means requests to use the key must originate through that approved interface endpoint, helping ensure that identities with otherwise valid permissions cannot use the key through an unintended endpoint path. Together, private DNS routing and a key-policy endpoint condition address both private connectivity and enforcement. See the AWS KMS interface VPC endpoint documentation and the AWS KMS aws:SourceVpce condition documentation.

QUESTION NO: 22

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

Which solution will meet these requirements?

A. Use EventBridge to disable the instance profile access keys.

B. Use EventBridge to invoke a Lambda function that removes the affected instance from the Auto Scaling group and isolates it with a restricted security group.

C. Use Security Hub to update the subnet network ACL to block traffic.

D. Send GuardDuty findings to Amazon SNS for email notification.

ANSWER: B

Explanation:

Use EventBridge to invoke a Lambda function that removes the affected instance from the Auto Scaling group and isolates it with a restricted security group. Amazon GuardDuty publishes findings to Amazon EventBridge, which supports event-driven automation for defined finding patterns. A Lambda function can then execute a targeted containment workflow without requiring manual intervention. For an instance that might be compromised, isolating it with a dedicated restrictive security group preserves the instance for investigation, evidence collection, and remediation while immediately limiting its network communications.

Removing the instance from the Auto Scaling group is important because it prevents the group from treating the containment action as an ordinary unhealthy-instance replacement scenario or continuing to manage the affected workload as part of the production fleet. Healthy instances in the group can continue serving the web application, so the response limits the operational effect to the specific instance implicated by the finding. This is aligned with AWS incident-response guidance to contain affected resources in a controlled manner while retaining forensic visibility, and with GuardDuty's documented EventBridge integration for automated remediation workflows. See Amazon GuardDuty findings in EventBridge and AWS Security Incident Response Guide.

QUESTION NO: 23

A company is expanding its group of stores. On the day that each new store opens, the company wants to launch a customized web application for that store. Each store ' s application will have a non-production environment and a production environment. Each environment will be deployed in a separate AWS account. The company uses AWS Organizations and has an OU that is used only for these accounts.

The company distributes most of the development work to third-party development teams. A security engineer needs to ensure that each team follows the company ' s deployment plan for AWS resources. The security engineer also must limit access to the deployment plan to only the developers who need access. The security engineer already has created an AWS CloudFormation template that implements the deployment plan.

What should the security engineer do next to meet the requirements in theMOST secureway?

A. Create an AWS Service Catalog portfolio in the organization ' s management account. Upload the CloudFormation template. Add the template to the portfolio ' s product list. Share the portfolio with the OU.

B. Use the CloudFormation CLI to create a module from the CloudFormation template. Register the module as a private extension in the CloudFormation registry. Publish the extension. In the OU, create an SCP that allows access to the extension.

C. Create an AWS Service Catalog portfolio in the organization ' s management account. Upload the CloudFormation template. Add the template to the portfolio ' s product list. Create an IAM role that has a trust policy that allows cross-account access to the portfolio for users in the OU accounts. Attach the AWSServiceCatalogEndUserFullAccess managed policy to the role.

D. Use the CloudFormation CLI to create a module from the CloudFormation template. Register the module as a private extension in the CloudFormation registry. Publish the extension. Share the extension with the OU.

ANSWER: A

Explanation:

Create an AWS Service Catalog portfolio in the organization ' s management account. Upload the CloudFormation template. Add the template to the portfolio ' s product list. Share the portfolio with the OU. This approach makes the approved CloudFormation-based deployment plan available as a governed Service Catalog product to only the dedicated organizational unit that contains the store accounts. A portfolio is the Service Catalog construct used to organize products and control which AWS accounts can receive them; sharing it with an OU automatically extends availability to the current accounts in that OU and can accommodate additional store accounts as they are created.

Service Catalog is designed for centrally governed self-service provisioning. The security team retains control of the product definition and versions, while developers launch the approved product rather than independently building an arbitrary resource configuration. Access can then be further limited to the required third-party developers through IAM permissions in the receiving accounts, following least privilege. This separates central template governance from developer provisioning access and provides a scalable operating model for repeated production and non-production account deployments.

AWS documents portfolio sharing through AWS Organizations, including sharing with an organizational unit, in the AWS Service Catalog Administrator Guide. The Service Catalog product model and CloudFormation provisioning behavior are described in the Service Catalog documentation.

QUESTION NO: 24

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

Which solution will meet these requirements?

A. Enable AWS Trusted Advisor. Configure all the Trusted Advisor checks. Manually map the checks against the PCI DSS v4.0 standard to generate the report.

B. Enable and configure AWS Config. Deploy the Operational Best Practices for PCI DSS conformance pack in AWS Config. Use AWS Config to generate the report.

C. Enable AWS Security Hub. Enable the Security Hub PCI DSS security standard. Use the AWS Management Console to download the report from the security standard.

D. Create an AWS Audit Manager assessment that uses the AWS managed PCI DSS v4.0 standard framework. Add all evidence to the assessment. Generate the report in Audit Manager for download.

ANSWER: D

Explanation:

Create an AWS Audit Manager assessment that uses the AWS managed PCI DSS v4.0 standard framework. Add all evidence to the assessment. Generate the report in Audit Manager for download. This approach directly meets both the control-assessment and evidence-management requirements. AWS Audit Manager provides managed frameworks that

organize compliance requirements into controls and can automatically collect evidence from relevant AWS data sources. For this application, the assessment can evaluate controls applicable to services such as API Gateway, Lambda, and Aurora, while retaining the evidence in a structured, audit-ready assessment.

Crucially, Audit Manager supports manual evidence. The compliance team can upload files or provide other evidence to a specific control set or control, allowing documentation that cannot be collected automatically—such as policies, attestations, or review records—to be included alongside automatically collected evidence. After evidence is collected and reviewed, Audit Manager can generate an assessment report containing the selected controls and their supporting evidence. This produces a downloadable package suitable for demonstrating the effectiveness of the applicable PCI DSS controls to internal stakeholders or auditors. See the AWS Audit Manager assessments documentation and AWS Audit Manager manual evidence documentation.

QUESTION NO: 25

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

Which additional step will meet these requirements?

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

B. Add a StringEquals condition to the IAM role policy for the EC2 instance profile. Configure the policy condition to restrict access based on the s3:ResourceTag/ClientId tag of each invoice. Tag each generated invoice with the ID of its corresponding client.

C. Update the script to use AWS Security Token Service (AWS STS) to obtain new credentials each time the script runs by assuming a new role that has S3:GetObject permissions. Use the credentials to generate the pre-signed URLs.

D. Generate an access key and a secret key for an IAM user that has S3:GetObject permissions on the S3 bucket. Embed the keys into the script. Use the keys to generate the pre-signed URLs.

E. Update the application script to generate an S3 presigned URL for each individual invoice with S3:GetObject access and an expiration of 1 hour, using the EC2 instance profile credentials.

ANSWER: E

Explanation:

Update the application script to generate an Amazon S3 presigned `GetObject` URL for each individual invoice, with an expiration of 1 hour, using the credentials supplied by the EC2 instance profile. A presigned URL grants access to the specific S3 object identified in the URL, so the application can issue each client URLs only for that client’s invoice objects. The client does not need an AWS account, IAM user, access key, or any other AWS credentials; the signed URL is the time-limited authorization mechanism for downloading that particular object.

Because the EC2 instance profile obtains short-term role credentials automatically from AWS STS, the application does not need to store or embed long-term access keys. The presigned URL’s expiration should be configured to no more than 3,600 seconds. S3 evaluates the request using the permissions of the role that created the URL, so the instance profile must retain permission to read the invoice objects. Keeping the bucket private and distributing each URL only to its intended client provides the required object-level sharing model.

See Sharing objects with presigned URLs and IAM roles for Amazon EC2.

QUESTION NO: 26

A company wants to implement a content delivery network (CDN) for an upcoming product launch. The origin for distribution is a web server outside the AWS Cloud. The origin requires an authorization header from each request.

Which solution will meet these requirements?

A. Use AWS Global Accelerator to create a custom routing accelerator. Configure the accelerator to use TCP. Specify the web server's IP address. Include the required authorization header in the request. Configure the web server to respond to requests from authorized users by using a signed cookie in the response header.

B. Create an Amazon CloudFront distribution. Configure CloudFront to use an origin access control (OAC). Specify the web server as the origin. Include the required authorization header in the OAC request. Configure the web server to respond to requests from authorized users by using a signed URL.

C. Create an Amazon CloudFront distribution. Configure CloudFront to forward a custom header to the origin. Specify trusted key groups for the distribution. Configure the web server to respond to requests from authorized users by using a signed URL.

D. Use AWS Global Accelerator to create an origin access control (OAC). Specify the web server as the origin. Include a custom header in the request. Configure the web server to respond to requests from authorized users by using a signed cookie in the response header.

ANSWER: C

Explanation:

Creating an Amazon CloudFront distribution, configuring CloudFront to forward a custom header to the origin, and specifying trusted key groups meets both the CDN and request-authorization requirements. CloudFront supports custom origins, including web servers that are hosted outside AWS. A custom origin header can be configured on the distribution so that CloudFront includes the required authorization header whenever it sends a request to the external origin. This enables the origin to verify that requests arriving through the distribution carry the expected authorization value.

Trusted key groups provide viewer-access protection through CloudFront signed URLs. The application signs URLs with a private key, and CloudFront uses the corresponding public key in the trusted key group to validate the signature before serving the requested content. This permits access to be limited to authorized users while retaining CloudFront caching and global edge delivery. CloudFront forwards the authorized request to the custom origin with the configured header, satisfying the origin's requirement on every origin request. For implementation details, see the CloudFront custom origin headers documentation and the CloudFront signed URLs and trusted key groups documentation.