SCS-C02 Dumps

AWS Certified Security - Specialty

https://www.certleader.com/SCS-C02-dumps.html

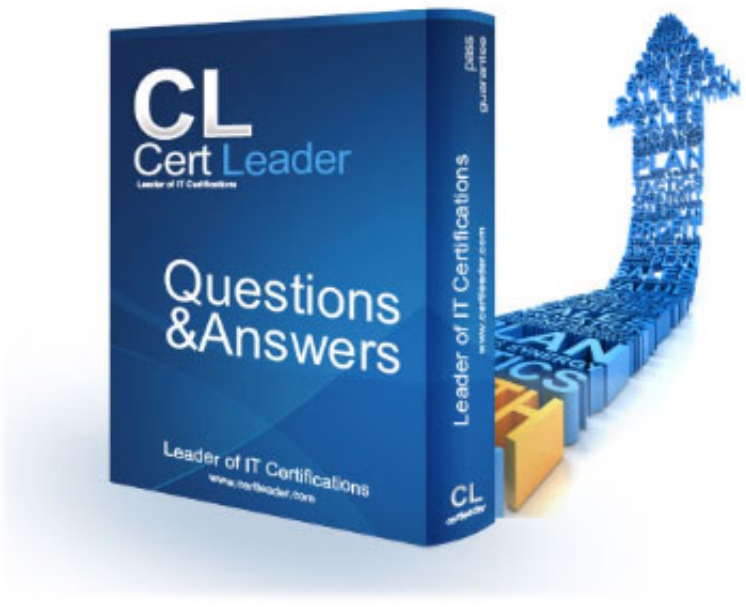

NEW QUESTION 1
A Security Architect has been asked to review an existing security architecture and identify why the application servers cannot successfully initiate a connection to the database servers. The following summary describes the architecture:
* 1 An Application Load Balancer, an internet gateway, and a NAT gateway are configured in the public subnet
* 2. Database, application, and web servers are configured on three different private subnets.
* 3 The VPC has two route tables: one for the public subnet and one for all other subnets The route table for the public subnet has a 0 0 0 0/0 route to the internet gateway The route table for all other subnets has a 0 0.0.0/0 route to the NAT gateway. All private subnets can route to each other
* 4 Each subnet has a network ACL implemented that limits all inbound and outbound connectivity to only the required ports and protocols
* 5 There are 3 Security Groups (SGs) database application and web Each group limits all inbound and outbound connectivity to the minimum required
Which of the following accurately reflects the access control mechanisms the Architect should verify1?

A. Outbound SG configuration on database servers Inbound SG configuration on application servers inbound and outbound network ACL configuration on the database subnet Inbound and outbound network ACL configuration on the application server subnet
B. Inbound SG configuration on database servers Outbound SG configuration on application serversInbound and outbound network ACL configuration on the database subnetInbound and outbound network ACL configuration on the application server subnet
C. Inbound and outbound SG configuration on database servers Inbound and outbound SG configuration on application servers Inbound network ACL configuration on the database subnet Outbound network ACL configuration on the application server subnet
D. Inbound SG configuration on database servers Outbound SG configuration on application servers Inbound network ACL configuration on the database subnet Outbound network ACL configuration on the application server subnet.

Answer: A

Explanation:
this is the accurate reflection of the access control mechanisms that the Architect should verify. Access control mechanisms are methods that regulate who can access what resources and how. Security groups and network ACLs are two types of access control mechanisms that can be applied to EC2 instances and subnets. Security groups are stateful, meaning they remember and return traffic that was previously allowed. Network ACLs are stateless, meaning they do not remember or return traffic that was previously allowed. Security groups and network ACLs can have inbound and outbound rules that specify the source, destination, protocol, and port of the traffic. By verifying the outbound security group configuration on database servers, the inbound security group configuration on application servers, and the inbound and outbound network ACL configuration on both the database and application server subnets, the Architect can check if there are any misconfigurations or conflicts that prevent the application servers from initiating a connection to the database servers. The other options are either inaccurate or incomplete for verifying the access control mechanisms.

NEW QUESTION 2
A company is running workloads in a single IAM account on Amazon EC2 instances and Amazon EMR clusters a recent security audit revealed that multiple Amazon Elastic Block Store (Amazon EBS) volumes and snapshots are not encrypted
The company's security engineer is working on a solution that will allow users to deploy EC2 Instances and EMR clusters while ensuring that all new EBS volumes and EBS snapshots are encrypted at rest. The solution must also minimize operational overhead
Which steps should the security engineer take to meet these requirements?

A. Create an Amazon Event Bridge (Amazon Cloud watch Events) event with an EC2 instance as the source and create volume as the event trigge
B. When the event is triggered invoke an IAM Lambda function to evaluate and notify the security engineer if the EBS volume that was created is not encrypted.
C. Use a customer managed IAM policy that will verify that the encryption ag of the Createvolume context is set to tru
D. Apply this rule to all users.
E. Create an IAM Config rule to evaluate the conguration of each EC2 instance on creation or modication.Have the IAM Cong rule trigger an IAM Lambdafunction to alert the security team and terminate the instance it the EBS volume is not encrypte
F. 5
G. Use the IAM Management Console or IAM CLi to enable encryption by default for EBS volumes in each IAM Region where the company operates.

Answer: D

Explanation:
To ensure that all new EBS volumes and EBS snapshots are encrypted at rest and minimize operational overhead, the security engineer should do the following:
> Use the AWS Management Console or AWS CLI to enable encryption by default for EBS volumes in each AWS Region where the company operates. This allows the security engineer to automatically encrypt any new EBS volumes and snapshots created from those volumes, without requiring any additional actions from users.

NEW QUESTION 3
A company hosts business-critical applications on Amazon EC2 instances in a VPC. The VPC uses default DHCP options sets. A security engineer needs to log all DNS queries that internal resources make in the VPC. The security engineer also must create a list of the most common DNS queries over time.
Which solution will meet these requirements?

A. Install the Amazon CloudWatch agent on each EC2 instance in the VP
B. Use the CloudWatch agent to stream the DNS query logs to an Amazon CloudWatch Logs log grou
C. Use CloudWatch metric filters to automatically generate metrics that list the most common ONS queries.
D. Install a BIND DNS server in the VP
E. Create a bash script to list the DNS request number of common DNS queries from the BIND logs.
F. Create VPC flow logs for all subnets in the VP
G. Stream the flow logs to an Amazon CloudWatch Logs log grou
H. Use CloudWatch Logs Insights to list the most common DNS queries for the log group in a custom dashboard.
I. Configure Amazon Route 53 Resolver query loggin
J. Add an Amazon CloudWatch Logs log group as the destinatio
K. Use Amazon CloudWatch Contributor Insights to analyze the data and create time series that display the most common DNS queries.

Answer: D

Explanation:
https://aws.amazon.com/blogs/aws/log-your-vpc-dns-queries-with-route-53-resolver-query-logs/

NEW QUESTION 4
A company is attempting to conduct forensic analysis on an Amazon EC2 instance, but the company is unable to connect to the instance by using AWS Systems Manager Session Manager. The company has installed AWS Systems Manager Agent (SSM Agent) on the EC2 instance.
The EC2 instance is in a subnet in a VPC that does not have an internet gateway attached. The company has associated a security group with the EC2 instance. The security group does not have inbound or outbound rules. The subnet's network ACL allows all inbound and outbound traffic.
Which combination of actions will allow the company to conduct forensic analysis on the EC2 instance without compromising forensic data? (Select THREE.)

A. Update the EC2 instance security group to add a rule that allows outbound traffic on port 443 for 0.0.0.0/0.
B. Update the EC2 instance security group to add a rule that allows inbound traffic on port 443 to the VPC's CIDR range.
C. Create an EC2 key pai
D. Associate the key pair with the EC2 instance.
E. Create a VPC interface endpoint for Systems Manager in the VPC where the EC2 instance is located.
F. Attach a security group to the VPC interface endpoin
G. Allow inbound traffic on port 443 to the VPC's CIDR range.
H. Create a VPC interface endpoint for the EC2 instance in the VPC where the EC2 instance is located.

Answer: BCF

NEW QUESTION 5
A Security Engineer receives alerts that an Amazon EC2 instance on a public subnet is under an SFTP brute force attack from a specific IP address, which is a known malicious bot. What should the Security Engineer do to block the malicious bot?

A. Add a deny rule to the public VPC security group to block the malicious IP
B. Add the malicious IP to IAM WAF backhsted IPs
C. Configure Linux iptables or Windows Firewall to block any traffic from the malicious IP
D. Modify the hosted zone in Amazon Route 53 and create a DNS sinkhole for the malicious IP

Answer: D

Explanation:
what the Security Engineer should do to block the malicious bot. SFTP is a protocol that allows secure file transfer over SSH. EC2 is a service that provides virtual servers in the cloud. A public subnet is a subnet that has a route to an internet gateway, which allows it to communicate with the internet. A brute force attack is a type of attack that tries to guess passwords or keys by trying many possible combinations. A malicious bot is a software program that performs automated tasks for malicious purposes. Route 53 is a service that provides DNS resolution and domain name registration. A DNS sinkhole is a technique that redirects malicious or unwanted traffic to a different destination, such as a black hole server or a honeypot. By modifying the hosted zone in Route 53 and creating a DNS sinkhole for the malicious IP, the Security Engineer can block the malicious bot from reaching the EC2 instance on the public subnet. The other options are either ineffective or inappropriate for blocking the malicious bot.

NEW QUESTION 6
A security engineer must troubleshoot an administrator's inability to make an existing Amazon S3 bucket public in an account that is part of an organization n IAM Organizations. The administrator switched the role from the master account to a member account and then attempted to make one S3 bucket public. This action was immediately denied
Which actions should the security engineer take to troubleshoot the permissions issue? (Select TWO.)

A. Review the cross-account role permissions and the S3 bucket policy Verify that the Amazon S3 block public access option in the member account is deactivated.
B. Review the role permissions m the master account and ensure it has sufficient privileges to perform S3 operations
C. Filter IAM CloudTrail logs for the master account to find the original deny event and update the cross-account role m the member account accordingly Verify that the Amazon S3 block public access option in the master account is deactivated.
D. Evaluate the SCPs covering the member account and the permissions boundary of the role in the member account for missing permissions and explicit denies.
E. Ensure the S3 bucket policy explicitly allows the s3 PutBucketPublicAccess action for the role m the member account

Answer: DE

Explanation:
- A is incorrect because reviewing the cross-account role permissions and the S3 bucket policy is not enough to troubleshoot the permissions issue. You also need to verify that the Amazon S3 block public access option in the member account is deactivated, as well as the permissions boundary and the SCPs of the role in the member account.
- D is correct because evaluating the SCPs and the permissions boundary of the role in the member account can help you identify any missing permissions or explicit denies that could prevent the administrator from making the S3 bucket public.
- E is correct because ensuring that the S3 bucket policy explicitly allows the s3 PutBucketPublicAccess action for the role in the member account can help you override any block public access settings that could prevent the administrator from making the S3 bucket public.

NEW QUESTION 7
A company used a lift-and-shift approach to migrate from its on-premises data centers to the AWS Cloud. The company migrated on-premises VMS to Amazon EC2 in-stances. Now the company wants to replace some of components that are running on the EC2 instances with managed AWS services that provide similar functionality.
Initially, the company will transition from load balancer software that runs on EC2 instances to AWS Elastic Load Balancers. A security engineer must ensure that after this transition, all the load balancer logs are centralized and searchable for auditing. The security engineer must also ensure that metrics are generated to show which ciphers are in use.
Which solution will meet these requirements?

A. Create an Amazon CloudWatch Logs log grou
B. Configure the load balancers to send logs to the log grou
C. Use the CloudWatch Logs console to search the log
D. Create CloudWatch Logs filters on the logs for the required met-rics.
E. Create an Amazon S3 bucke
F. Configure the load balancers to send logs to the S3 bucke
G. Use Amazon Athena to search the logs that are in the S3 bucke

H. Create Amazon CloudWatch filters on the S3 log files for the re-quired metrics.
I. Create an Amazon S3 bucke
J. Configure the load balancers to send logs to the S3 bucke
K. Use Amazon Athena to search the logs that are in the S3 bucke
L. Create Athena queries for the required metric
M. Publish the metrics to Amazon CloudWatch.
N. Create an Amazon CloudWatch Logs log grou
O. Configure the load balancers to send logs to the log grou
P. Use the AWS Management Console to search the log
Q. Create Amazon Athena queries for the required metric
R. Publish the metrics to Amazon CloudWatch.

Answer: C

Explanation:

- Amazon S3 is a service that provides scalable, durable, and secure object storage. You can use Amazon S3 to store and retrieve any amount of data from anywhere on the web1
- AWS Elastic Load Balancing is a service that distributes incoming application or network traffic across multiple targets, such as EC2 instances, containers, or IP addresses. You can use Elastic Load Balancing to increase the availability and fault tolerance of your applications2
- Elastic Load Balancing supports access logging, which captures detailed information about requests sent to your load balancer. Each log contains information such as the time the request was received, the client's IP address, latencies, request paths, and server responses. You can use access logs to analyze traffic patterns and troubleshoot issues3
- You can configure your load balancer to store access logs in an Amazon S3 bucket that you specify.

You can also specify the interval for publishing the logs, which can be 5 or 60 minutes. The logs are stored in a hierarchical folder structure by load balancer name, IP address, year, month, day, and time.

- Amazon Athena is a service that allows you to analyze data in Amazon S3 using standard SQL. You can use Athena to run ad-hoc queries and get results in seconds. Athena is serverless, so there is no infrastructure to manage and you pay only for the queries that you run.
- You can use Athena to search the access logs that are stored in your S3 bucket. You can create a table in Athena that maps to your S3 bucket and then run SQL queries on the table. You can also use the Athena console or API to view and download the query results.
- You can also use Athena to create queries for the required metrics, such as the number of requests per cipher or protocol. You can then publish the metrics to Amazon CloudWatch, which is a service that monitors and manages your AWS resources and applications. You can use CloudWatch to collect and track metrics, create alarms, and automate actions based on the state of your resources.
- By using this solution, you can meet the requirements of ensuring that all the load balancer logs are centralized and searchable for auditing and that metrics are generated to show which ciphers are in use.

NEW QUESTION 8
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 Hu
B. Publish events to Amazon CloudWatch from the delegated administrator account, all member accounts, and required Regions that are enabled for Security Hub findings.
C. Designate an AWS account in an organization in AWS Organizations as a delegated administrator for Security Hu
D. Publish events to Amazon EventBridge from the delegated administrator account, all member accounts, and required Regions that are enabled for Security Hub findings.
E. In each Region, create an Amazon EventBridge rule to deliver findings to an Amazon Kinesis data strea
F. Configure the Kinesis data streams to output the logs to a single Amazon S3 bucket.
G. In each Region, create an Amazon EventBridge rule to deliver findings to an Amazon Kinesis Data Firehose delivery strea
H. Configure the Kinesis Data Firehose delivery streams to deliver the logs to a single Amazon S3 bucket.
I. Use AWS Glue DataBrew to crawl the Amazon S3 bucket and build the schem
J. Use AWS Glue Data Catalog to query the data and create views to flatten nested attribute
K. Build Amazon QuickSight dashboards by using Amazon Athena.
L. Partition the Amazon S3 dat
M. Use AWS Glue to crawl the S3 bucket and build the schem
N. Use Amazon Athena to query the data and create views to flatten nested attribute
O. Build Amazon QuickSight dashboards that use the Athena views.

Answer: BDF

Explanation:
The correct answer is B, D, and F. Designate an AWS account in an organization in AWS Organizations as a delegated administrator for Security Hub. Publish events to Amazon EventBridge from the delegated administrator account, all member accounts, and required Regions that are enabled for Security Hub findings. In each Region, create an Amazon EventBridge rule to deliver findings to an Amazon Kinesis Data Firehose delivery stream. Configure the Kinesis Data Firehose delivery streams to deliver the logs to a single Amazon S3 bucket. 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 to flatten nested attributes. Build Amazon QuickSight dashboards that use the Athena views.
According to the AWS documentation, AWS Security Hub is a service that provides you with a comprehensive view of your security state across your AWS accounts, and helps you check your environment against security standards and best practices. You can use Security Hub to aggregate security findings from various sources, such as AWS services, partner products, or your own applications.
To use Security Hub with multiple AWS accounts and Regions, you need to enable AWS Organizations with all features enabled. This allows you to centrally manage your accounts and apply policies across your organization. You can also use Security Hub as a service principal for AWS Organizations, which lets you designate a delegated administrator account for Security Hub. The delegated administrator account can enable Security Hub automatically in all existing and future accounts in your organization, and can view and manage findings from all accounts.
According to the AWS documentation, Amazon EventBridge is a serverless event bus that makes it easy to connect applications using data from your own applications, integrated software as a service (SaaS) applications, and AWS services. You can use EventBridge to create rules that match events from various sources and route them to targets for processing.
To use EventBridge with Security Hub findings, you need to enable Security Hub as an event source in EventBridge. This will allow you to publish events from

Security Hub to EventBridge in the same Region. You can then create EventBridge rules that match Security Hub findings based on criteria such as severity, type, or resource. You can also specify targets for your rules, such as Lambda functions, SNS topics, or Kinesis Data Firehose delivery streams.
According to the AWS documentation, Amazon Kinesis Data Firehose is a fully managed service that delivers real-time streaming data to destinations such as Amazon S3, Amazon Redshift, Amazon Elasticsearch Service (Amazon ES), and Splunk. You can use Kinesis Data Firehose to transform and enrich your data before delivering it to your destination.
To use Kinesis Data Firehose with Security Hub findings, you need to create a Kinesis Data Firehose delivery stream in each Region where you have enabled Security Hub. You can then configure the delivery stream to receive events from EventBridge as a source, and deliver the logs to a single S3 bucket as a destination. You can also enable data transformation or compression on the delivery stream if needed.
According to the AWS documentation, Amazon S3 is an object storage service that offers scalability, data availability, security, and performance. You can use S3 to store and retrieve any amount of data from anywhere on the web. You can also use S3 features such as lifecycle management, encryption, versioning, and replication to optimize your storage.
To use S3 with Security Hub findings, you need to create an S3 bucket that will store the logs from Kinesis Data Firehose delivery streams. You can then partition the data in the bucket by using prefixes such as account ID or Region. This will improve the performance and cost-effectiveness of querying the data.
According to the AWS documentation, AWS Glue is a fully managed extract, transform, and load (ETL) service that makes it easy to prepare and load your data for analytics. You can use Glue to crawl your data sources, identify data formats, and suggest schemas and transformations. You can also use Glue Data Catalog as a central metadata repository for your data assets.
To use Glue with Security Hub findings, you need to create a Glue crawler that will crawl the S3 bucket and build the schema for the data. The crawler will create tables in the Glue Data Catalog that you can query using standard SQL.
According to the AWS documentation, Amazon Athena is an interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL. Athena is serverless, so there is no infrastructure to manage, and you pay only for the queries that you run. You can use Athena with Glue Data Catalog as a metadata store for your tables.
To use Athena with Security Hub findings, you need to create views in Athena that will flatten nested attributes in the data. For example, you can create views that extract fields such as account ID, Region, resource type, resource ID, finding type, finding title, and finding description from the JSON data. You can then query the views using SQL and join them with other tables if needed.
According to the AWS documentation, Amazon QuickSight is a fast, cloud-powered business intelligence
service that makes it easy to deliver insights to everyone in your organization. You can use QuickSight to create and publish interactive dashboards that include machine learning insights. You can also use QuickSight to connect to various data sources, such as Athena, S3, or RDS.
To use QuickSight with Security Hub findings, you need to create QuickSight dashboards that use the Athena views as data sources. You can then visualize and analyze the findings using charts, graphs, maps, or tables. You can also apply filters, calculations, or aggregations to the data. You can then share the dashboards with your users or embed them in your applications.

NEW QUESTION 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 : Multi FactorAuthPresent": "true" }
B. "B001 " : " aws : MultiFactorAuthPresent": "false" }
C. "NumericLessThan" : { " aws : Multi FactorAuthAge" : "7200"}
D. "NumericGreaterThan" : { " aws : MultiFactorAuthAge " : "7200"
E. "NumericLessThan" : { "MaxSessionDuration " : "7200"}

Answer: AC

Explanation:
The correct combination of conditions to add to the IAM policy is A and C. These conditions will ensure that IAM users must use MFA to access certain services in the AWS production account, and that each session will expire after 2 hours.

- Option A: “Bool” : { “aws:MultiFactorAuthPresent” : “true” } is a valid condition that checks if the principal (the IAM user) has authenticated with MFA before making the request. This condition will enforce MFA for the IAM users to access the specified services. This condition key is supported by all AWS services that support IAM policies1.
- Option B: “Bool” : { “aws:MultiFactorAuthPresent” : “false” } is the opposite of option A. This condition will allow access only if the principal has not authenticated with MFA, which is not the desired requirement. This condition key is supported by all AWS services that support IAM policies1.
- Option C: “NumericLessThan” : { “aws:MultiFactorAuthAge” : “7200” } is a valid condition that checks if the time since the principal authenticated with MFA is less than 7200 seconds (2 hours). This condition will enforce the session duration limit for the IAM users. This condition key is supported by all AWS services that support IAM policies1.
- Option D: “NumericGreaterThan” : { “aws:MultiFactorAuthAge” : “7200” } is the opposite of option C. This condition will allow access only if the time since the principal authenticated with MFA is more than 7200 seconds (2 hours), which is not the desired requirement. This condition key is supported by all AWS services that support IAM policies1.
- Option E: “NumericLessThan” : { “MaxSessionDuration” : “7200” } is not a valid condition key.

MaxSessionDuration is a property of an IAM role, not a condition key. It specifies the maximum session duration (in seconds) for the role, which can be between 3600 and 43200 seconds (1 to 12 hours). This property can be set when creating or modifying a role, but it cannot be used as a condition in a policy2.

NEW QUESTION 10
Your company is planning on using bastion hosts for administering the servers in IAM. Which of the following is the best description of a bastion host from a security perspective?
Please select:

A. A Bastion host should be on a private subnet and never a public subnet due to security concerns
B. A Bastion host sits on the outside of an internal network and is used as a gateway into the private network and is considered the critical strong point of the network
C. Bastion hosts allow users to log in using RDP or SSH and use that session to S5H into internal network to access private subnet resources.
D. A Bastion host should maintain extremely tight security and monitoring as it is available to the public

Answer: C

Explanation:
A bastion host is a special purpose computer on a network specifically designed and configured to withstand attacks. The computer generally hosts a single application, for example a proxy server, and all other services are removed or limited to reduce the threat to the computer.
In IAM, A bastion host is kept on a public subnet. Users log on to the bastion host via SSH or RDP and then use that session to manage other hosts in the private subnets.
Options A and B are invalid because the bastion host needs to sit on the public network. Option D is invalid because bastion hosts are not used for monitoring For more information on bastion hosts, just browse to the below URL:
https://docsIAM.amazon.com/quickstart/latest/linux-bastion/architecture.htl
The correct answer is: Bastion hosts allow users to log in using RDP or SSH and use that session to SSH into internal network to access private subnet resources.
Submit your Feedback/Queries to our Experts

NEW QUESTION 10
A Network Load Balancer (NLB) target instance is not entering the InService state. A security engineer determines that health checks are failing.
Which factors could cause the health check failures? (Select THREE.)

A. The target instance's security group does not allow traffic from the NLB.
B. The target instance's security group is not attached to the NLB.
C. The NLB's security group is not attached to the target instance.
D. The target instance's subnet network ACL does not allow traffic from the NLB.
E. The target instance's security group is not using IP addresses to allow traffic from the NLB.
F. The target network ACL is not attached to the NLB.

Answer: ACD

NEW QUESTION 11
A company is implementing a new application in a new IAM account. A VPC and subnets have been created for the application. The application has been peered to an existing VPC in another account in the same IAM Region for database access. Amazon EC2 instances will regularly be created and terminated in the application VPC, but only some of them will need access to the databases in the peered VPC over TCP port 1521. A security engineer must ensure that only the EC2 instances that need access to the databases can access them through the network.
How can the security engineer implement this solution?

A. Create a new security group in the database VPC and create an inbound rule that allows all traffic from the IP address range of the application VP
B. Add a new network ACL rule on the database subnet
C. Configure the rule to TCP port 1521 from the IP address range of the application VP
D. Attach the new security group to the database instances that the application instances need to access.
E. Create a new security group in the application VPC with an inbound rule that allows the IP address range of the database VPC over TCP port 1521. Create a new security group in the database VPC with an inbound rule that allows the IP address range of the application VPC over port 1521. Attach the new security group to the database instances and the application instances that need database access.
F. Create a new security group in the application VPC with no inbound rule
G. Create a new security group in the database VPC with an inbound rule that allows TCP port 1521 from the new application security group in the application VP
H. Attach the application security group to the application instances that need database access, and attach the database security group to the database instances.
I. Create a new security group in the application VPC with an inbound rule that allows the IP address range of the database VPC over TCP port 1521. Add a new network ACL rule on the database subnet
J. Configure the rule to allow all traffic from the IP address range of the application VP
K. Attach the new security group to the application instances that need database access.

Answer: C

NEW QUESTION 14
A developer is building a serverless application hosted on AWS that uses Amazon Redshift as a data store The application has separate modules for readwrite and read-only functionality The modules need their own database users for compliance reasons
Which combination of steps should a security engineer implement to grant appropriate access? (Select TWO.)

A. Configure cluster security groups for each application module to control access to database users that are required for read-only and readwrite
B. Configure a VPC endpoint for Amazon Redshift Configure an endpoint policy that maps database users to each application module, and allow access to the tables that are required for read-only and read/write
C. Configure an 1AM policy for each module Specify the ARN of an Amazon Redshift database user that allows the GetClusterCredentials API call
D. Create local database users for each module
E. Configure an 1AM policy for each module Specify the ARN of an 1AM user that allows the GetClusterCredentials API call

Answer: A

Explanation:

To grant appropriate access to separate modules for read-write and read-only functionality in a serverless application hosted on AWS that uses Amazon Redshift as a data store, a security engineer should configure cluster security groups for each application module to control access to database users that are required for read-only and readwrite, and configure an IAM policy for each module specifying the ARN of an IAM user that allows the GetClusterCredentials API call.
References: : Amazon Redshift - Amazon Web Services : Amazon Redshift - Amazon Web Services : Identity and Access Management - AWS Management Console : AWS Identity and Access Management - AWS Management Console

NEW QUESTION 18
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 IAM DNS resolvers for all EC2 instances.
D. Configure a third-party DNS resolver with logging for all EC2 instances.

Answer: C

Explanation:
To ensure that the EC2 instances are logged, the security engineer should do the following:
- Use AWS DNS resolvers for all EC2 instances. This allows the security engineer to use Amazon-provided DNS servers that resolve public DNS hostnames to private IP addresses within their VPC, and that log DNS queries in Amazon CloudWatch Logs.

NEW QUESTION 20
A company Is planning to use Amazon Elastic File System (Amazon EFS) with its on-premises servers. The company has an existing IAM Direct Connect connection established between its on-premises data center and an IAM Region Security policy states that the company's on-premises firewall should only have specific IP addresses added to the allow list and not a CIDR range. The company also wants to restrict access so that only certain data center-based servers have access to Amazon EFS
How should a security engineer implement this solution"

A. Add the file-system-id efs IAM-region amazonIAM com URL to the allow list for the data center firewall Install the IAM CLI on the data center-based servers to mount the EFS file system in the EFS security group add the data center IP range to the allow list Mount the EFS using the EFS file system name
B. Assign an Elastic IP address to Amazon EFS and add the Elastic IP address to the allow list for the data center firewall Install the IAM CLI on the data center-based servers to mount the EFS file system In the EFS security group, add the IP addresses of the data center servers to the allow list Mount the EFS using the Elastic IP address
C. Add the EFS file system mount target IP addresses to the allow list for the data center firewall In the EFS security group, add the data center server IP addresses to the allow list Use the Linux terminal to mount the EFS file system using the IP address of one of the mount targets
D. Assign a static range of IP addresses for the EFS file system by contacting IAM Support In the EFS security group add the data center server IP addresses to the allow list Use the Linux terminal to mount the EFS file system using one of the static IP addresses

Answer: B

Explanation:
To implement the solution, the security engineer should do the following:
- Assign an Elastic IP address to Amazon EFS and add the Elastic IP address to the allow list for the data center firewall. This allows the security engineer to use a specific IP address for the EFS file system that can be added to the firewall rules, instead of a CIDR range or a URL.
- Install the AWS CLI on the data center-based servers to mount the EFS file system. This allows the security engineer to use the mount helper provided by AWS CLI to mount the EFS file system with encryption in transit.
- In the EFS security group, add the IP addresses of the data center servers to the allow list. This allows the security engineer to restrict access to the EFS file system to only certain data center-based servers.
- Mount the EFS using the Elastic IP address. This allows the security engineer to use the Elastic IP address as the DNS name for mounting the EFS file system.

NEW QUESTION 22
Your CTO thinks your IAM account was hacked. What is the only way to know for certain if there was unauthorized access and what they did, assuming your hackers are very sophisticated IAM engineers and doing everything they can to cover their tracks?
Please select:

A. Use CloudTrail Log File Integrity Validation.
B. Use IAM Config SNS Subscriptions and process events in real time.
C. Use CloudTrail backed up to IAM S3 and Glacier.
D. Use IAM Config Timeline forensics.

Answer: A

Explanation:
The IAM Documentation mentions the following
To determine whether a log file was modified, deleted, or unchanged after CloudTrail delivered it you can use CloudTrail log file integrity validation. This feature is built using industry standard algorithms: SHA-256 for hashing and SHA-256 with RSA for digital signing. This makes it computationally infeasible to modify, delete or forge CloudTrail log files without detection. You can use the IAM CLI to validate the files in the location where CloudTrail delivered them
Validated log files are invaluable in security and forensic investigations. For example, a validated log file enables you to assert positively that the log file itself has not changed, or that particular user credentials performed specific API activity. The CloudTrail log file integrity validation process also lets you know if a log file has been deleted or changed, or assert positively that no log files were delivered to your account during a given period of time.
Options B.C and D is invalid because you need to check for log File Integrity Validation for cloudtrail logs For more information on Cloudtrail log file validation, please visit the below URL: http://docs.IAM.amazon.com/IAMcloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html
The correct answer is: Use CloudTrail Log File Integrity Validation. omit your Feedback/Queries to our Expert

NEW QUESTION 25
A security engineer must use AWS Key Management Service (AWS KMS) to design a key management solution for a set of Amazon Elastic Block Store (Amazon EBS) volumes that contain sensitive data. The solution needs to ensure that the key material automatically expires in 90 days.
Which solution meets these criteria?

A. A customer managed CMK that uses customer provided key material
B. A customer managed CMK that uses AWS provided key material
C. An AWS managed CMK
D. Operation system-native encryption that uses GnuPG

Answer: A

Explanation:
https://awscli.amazonaws.com/v2/documentation/api/latest/reference/kms/import-key-material.html aws kms import-key-material \
--key-id 1234abcd-12ab-34cd-56ef-1234567890ab \
--encrypted-key-material fileb://EncryptedKeyMaterial.bin \
--import-token fileb://ImportToken.bin \
--expiration-model KEY_MATERIAL_EXPIRES \
--valid-to 2021-09-21T19:00:00Z
The correct answer is A. A customer managed CMK that uses customer provided key material.
A customer managed CMK is a KMS key that you create, own, and manage in your AWS account. You have full control over the key configuration, permissions, rotation, and deletion. You can use a customer managed CMK to encrypt and decrypt data in AWS services that are integrated with AWS KMS, such as Amazon EBS1.
A customer managed CMK can use either AWS provided key material or customer provided key material. AWS provided key material is generated by AWS KMS and never leaves the service unencrypted. Customer provided key material is generated outside of AWS KMS and imported into a customer managed CMK. You can specify an expiration date for the imported key material, after which the CMK becomes unusable until you reimport new key material2.
To meet the criteria of automatically expiring the key material in 90 days, you need to use customer provided key material and set the expiration date accordingly. This way, you can ensure that the data encrypted with the CMK will not be accessible after 90 days unless you reimport new key material and re-encrypt the data.
The other options are incorrect for the following reasons:
* B. A customer managed CMK that uses AWS provided key material does not expire automatically. You can enable automatic rotation of the key material every year, but this does not prevent access to the data encrypted with the previous key material. You would need to manually delete the CMK and its backing key material to make the data inaccessible3.
* C. An AWS managed CMK is a KMS key that is created, owned, and managed by an AWS service on your behalf. You have limited control over the key configuration, permissions, rotation, and deletion. You cannot use an AWS managed CMK to encrypt data in other AWS services or applications. You also cannot set an expiration date for the key material of an AWS managed CMK4.
* D. Operation system-native encryption that uses GnuPG is not a solution that uses AWS KMS. GnuPG is a command line tool that implements the OpenPGP standard for encrypting and signing data. It does not integrate with Amazon EBS or other AWS services. It also does not provide a way to automatically expire the key material used for encryption5.
References:
1: Customer Managed Keys - AWS Key Management Service 2: [Importing Key Material in AWS Key Management Service (AWS KMS) - AWS Key Management Service] 3: [Rotating Customer Master Keys - AWS Key Management Service] 4: [AWS Managed Keys - AWS Key Management Service] 5: The GNU Privacy Guard

NEW QUESTION 30
A Security Engineer is building a Java application that is running on Amazon EC2. The application communicates with an Amazon RDS instance and authenticates with a user name and password.
Which combination of steps can the Engineer take to protect the credentials and minimize downtime when the credentials are rotated? (Choose two.)

A. Have a Database Administrator encrypt the credentials and store the ciphertext in Amazon S3. Grant permission to the instance role associated with the EC2 instance to read the object and decrypt the ciphertext.
B. Configure a scheduled job that updates the credential in AWS Systems Manager Parameter Store and notifies the Engineer that the application needs to be restarted.
C. Configure automatic rotation of credentials in AWS Secrets Manager.
D. Store the credential in an encrypted string parameter in AWS Systems Manager Parameter Stor
E. Grant permission to the instance role associated with the EC2 instance to access the parameter and the AWS KMS key that is used to encrypt it.
F. Configure the Java application to catch a connection failure and make a call to AWS Secrets Manager to retrieve updated credentials when the password is rotate
G. Grant permission to the instance role associated with the EC2 instance to access Secrets Manager.

Answer: CE

Explanation:
AWS Secrets Manager is a service that helps you manage, retrieve, and rotate secrets such as database credentials, API keys, and other sensitive information. By configuring automatic rotation of credentials in AWS Secrets Manager, you can ensure that your secrets are changed regularly and securely, without requiring manual intervention or application downtime. You can also specify the rotation frequency and the rotation function that performs the logic of changing the credentials on the database and updating the secret in Secrets Manager1.
* E. Configure the Java application to catch a connection failure and make a call to AWS Secrets Manager to retrieve updated credentials when the password is rotated. Grant permission to the instance role associated with the EC2 instance to access Secrets Manager.
By configuring the Java application to catch a connection failure and make a call to AWS Secrets Manager to retrieve updated credentials, you can avoid hard-coding the credentials in your application code or configuration files. This way, your application can dynamically obtain the latest credentials from Secrets Manager whenever the password is rotated, without needing to restart or redeploy the application. To enable this, you need to grant permission to the instance role associated with the EC2 instance to access Secrets Manager using IAM policies2. You can also use the AWS SDK for Java to integrate your application with Secrets Manager3.

NEW QUESTION 33
A business requires a forensic logging solution for hundreds of Docker-based apps running on Amazon EC2. The solution must analyze logs in real time, provide message replay, and persist logs.
Which Amazon Web Offerings (IAM) services should be employed to satisfy these requirements? (Select two.)

A. Amazon Athena
B. Amazon Kinesis

C. Amazon SQS
D. Amazon Elasticsearch
E. Amazon EMR

Answer: BD

NEW QUESTION 35
A team is using AWS Secrets Manager to store an application database password. Only a limited number of IAM principals within the account can have access to the secret. The principals who require access to the secret change frequently. A security engineer must create a solution that maximizes flexibility and scalability.
Which solution will meet these requirements?

A. Use a role-based approach by creating an IAM role with an inline permissions policy that allows access to the secre
B. Update the IAM principals in the role trust policy as required.
C. Deploy a VPC endpoint for Secrets Manage
D. Create and attach an endpoint policy that specifies the IAM principals that are allowed to access the secre
E. Update the list of IAM principals as required.
F. Use a tag-based approach by attaching a resource policy to the secre
G. Apply tags to the secret and the IAM principal
H. Use the aws:PrincipalTag and aws:ResourceTag IAM condition keys to control access.
I. Use a deny-by-default approach by using IAM policies to deny access to the secret explicitl
J. Attach the policies to an IAM grou
K. Add all IAM principals to the IAM grou
L. Remove principals from the group when they need acces
M. Add the principals to the group again when access is no longer allowed.

Answer: C

NEW QUESTION 36
A company is building an application on IAM 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 snapshot
B. Enable Amazon Elastic Block Store (Amazon EBS) encryption on Amazon EC2 instance
C. Include the database credential in the EC2 user data fiel
D. Use an IAM Lambda function to rotate database credential
E. Set up TLS for the connection to the database.
F. Install a database on an Amazon EC2 Instanc
G. Enable third-party disk encryption to encrypt the Amazon Elastic Block Store (Amazon EBS) volum
H. Store the database credentials in IAM CloudHSM with automatic rotatio
I. Set up TLS for the connection to the database.
J. Enable Amazon RDS encryption to encrypt the database and snapshot
K. Enable Amazon Elastic Block Store (Amazon EBS) encryption on Amazon EC2 instance
L. Store the database credentials in IAM Secrets Manager with automatic rotatio
M. Set up TLS for the connection to the RDS hosted database.
N. Set up an IAM CloudHSM cluster with IAM Key Management Service (IAM KMS) to store KMS keys.Set up Amazon RDS encryption using IAM KMS to encrypt the databas
O. Store database credentials in the IAM Systems Manager Parameter Store with automatic rotatio
P. Set up TLS for the connection to the RDS hosted database.

Answer: C

Explanation:
To protect the sensitive data against any data breach and minimize management overhead, the security engineer should recommend the following solution:

- Enable Amazon RDS encryption to encrypt the database and snapshots. This allows the security engineer to use AWS Key Management Service (AWS KMS) to encrypt data at rest for the database and any backups or replicas.
- Enable Amazon Elastic Block Store (Amazon EBS) encryption on Amazon EC2 instances. This allows the security engineer to use AWS KMS to encrypt data at rest for the EC2 instances and any snapshots or volumes.
- Store the database credentials in AWS Secrets Manager with automatic rotation. This allows the security engineer to encrypt and manage secrets centrally, and to configure automatic rotation schedules for them.
- Set up TLS for the connection to the RDS hosted database. This allows the security engineer to encrypt data in transit between the EC2 instances and the database.

NEW QUESTION 39
A company is using IAM Secrets Manager to store secrets for its production Amazon RDS database. The Security Officer has asked that secrets be rotated every 3 months. Which solution would allow the company to securely rotate the secrets? (Select TWO.)

A. Place the RDS instance in a public subnet and an IAM Lambda function outside the VP
B. Schedule the Lambda function to run every 3 months to rotate the secrets.
C. Place the RDS instance in a private subnet and an IAM Lambda function inside the VPC in the private subne
D. Configure the private subnet to use a NAT gatewa
E. Schedule the Lambda function to run every 3 months to rotate the secrets.
F. Place the RDS instance in a private subnet and an IAM Lambda function outside the VP
G. Configure the private subnet to use an internet gatewa
H. Schedule the Lambda function to run every 3 months lo rotate the secrets.
I. Place the RDS instance in a private subnet and an IAM Lambda function inside the VPC in the private subne
J. Schedule the Lambda function to run quarterly to rotate the secrets.

K. Place the RDS instance in a private subnet and an IAM Lambda function inside the VPC in the private subne
L. Configure a Secrets Manager interface endpoin
M. Schedule the Lambda function to run every 3 months to rotate the secrets.

Answer: BE

Explanation:
these are the solutions that can securely rotate the secrets for the production RDS database using Secrets Manager. Secrets Manager is a service that helps you manage secrets such as database credentials, API keys, and passwords. You can use Secrets Manager to rotate secrets automatically by using a Lambda function that runs on a schedule. The Lambda function needs to have access to both the RDS instance and the Secrets Manager service. Option B places the RDS instance in a private subnet and the Lambda function in the same VPC in another private subnet. The private subnet with the Lambda function needs to use a NAT gateway to access Secrets Manager over the internet. Option E places the RDS instance and the Lambda function in the same private subnet and configures a Secrets Manager interface endpoint, which is a private connection between the VPC and Secrets Manager. The other options are either insecure or incorrect for rotating secrets using Secrets Manager.

NEW QUESTION 44
A company's public Application Load Balancer (ALB) recently experienced a DDoS attack. To mitigate this issue. the company deployed Amazon CloudFront in front of the ALB so that users would not directly access the Amazon EC2 instances behind the ALB.
The company discovers that some traffic is still coming directly into the ALB and is still being handled by the EC2 instances.
Which combination of steps should the company take to ensure that the EC2 instances will receive traffic only from CloudFront? (Choose two.)

A. Configure CloudFront to add a cache key policy to allow a custom HTTP header that CloudFront sends to the ALB.
B. Configure CloudFront to add a custom: HTTP header to requests that CloudFront sends to the ALB.
C. Configure the ALB to forward only requests that contain the custom HTTP header.
D. Configure the ALB and CloudFront to use the X-Forwarded-For header to check client IP addresses.
E. Configure the ALB and CloudFront to use the same X.509 certificate that is generated by AWS Certificate Manager (ACM).

Answer: BC

Explanation:
To prevent users from directly accessing an Application Load Balancer and allow access only through CloudFront, complete these high-level steps: Configure CloudFront to add a custom HTTP header to requests that it sends to the Application Load Balancer. Configure the Application Load Balancer to only forward requests that contain the custom HTTP header. (Optional) Require HTTPS to improve the security of this solution.
https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/restrict-access-to-load-balancer.html

NEW QUESTION 46
The Security Engineer is managing a traditional three-tier web application that is running on Amazon EC2 instances. The application has become the target of increasing numbers of malicious attacks from the Internet.
What steps should the Security Engineer take to check for known vulnerabilities and limit the attack surface? (Choose two.)

A. Use AWS Certificate Manager to encrypt all traffic between the client and application servers.
B. Review the application security groups to ensure that only the necessary ports are open.
C. Use Elastic Load Balancing to offload Secure Sockets Layer encryption.
D. Use Amazon Inspector to periodically scan the backend instances.
E. Use AWS Key Management Services to encrypt all the traffic between the client and application servers.

Answer: BD

Explanation:
The steps that the Security Engineer should take to check for known vulnerabilities and limit the attack surface are:
B. Review the application security groups to ensure that only the necessary ports are open. This is a good practice to reduce the exposure of the EC2 instances to potential attacks from the Internet. Application security groups are a feature of Azure that allow you to group virtual machines and define network security policies based on those groups1.
D. Use Amazon Inspector to periodically scan the backend instances. This is a service that helps you to identify vulnerabilities and exposures in your EC2 instances and applications. Amazon Inspector can perform automated security assessments based on predefined or custom rules packages2.

NEW QUESTION 51
A company is running an Amazon RDS for MySQL DB instance in a VPC. The VPC must not send or receive network traffic through the internet.
A security engineer wants to use AWS Secrets Manager to rotate the DB instance credentials automatically. Because of a security policy, the security engineer cannot use the standard AWS Lambda function that Secrets Manager provides to rotate the credentials.
The security engineer deploys a custom Lambda function in the VPC. The custom Lambda function will be responsible for rotating the secret in Secrets Manager.
The security engineer edits the DB instance's security group to allow connections from this function. When the function is invoked, the function cannot communicate with Secrets Manager to rotate the secret properly.
What should the security engineer do so that the function can rotate the secret?

A. Add an egress-only internet gateway to the VP
B. Allow only the Lambda function's subnet to route traffic through the egress-only internet gateway.
C. Add a NAT gateway to the VP
D. Configure only the Lambda function's subnet with a default route through the NAT gateway.
E. Configure a VPC peering connection to the default VPC for Secrets Manage
F. Configure the Lambda function's subnet to use the peering connection for routes.
G. Configure a Secrets Manager interface VPC endpoin
H. Include the Lambda function's private subnet during the configuration process.

Answer: D

Explanation:
You can establish a private connection between your VPC and Secrets Manager by creating an interface VPC endpoint. Interface endpoints are powered by AWS PrivateLink, a technology that enables you to privately access Secrets Manager APIs without an internet gateway, NAT device, VPN connection, or AWS Direct

Connect connection. Reference:
https://docs.aws.amazon.com/secretsmanager/latest/userguide/vpc-endpoint-overview.html
The correct answer is D. Configure a Secrets Manager interface VPC endpoint. Include the Lambda function's private subnet during the configuration process.
A Secrets Manager interface VPC endpoint is a private connection between the VPC and Secrets Manager that does not require an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection1. By configuring a Secrets Manager interface VPC endpoint, the security engineer can enable the custom Lambda function to communicate with Secrets Manager without sending or receiving network traffic through the internet. The security engineer must include the Lambda function's private subnet during the configuration process to allow the function to use the endpoint2.
The other options are incorrect for the following reasons:

- A. An egress-only internet gateway is a VPC component that allows outbound communication over IPv6 from instances in the VPC to the internet, and prevents the internet from initiating an IPv6 connection with the instances3. However, this option does not meet the requirement that the VPC must not send or receive network traffic through the internet. Moreover, an egress-only internet gateway is for use with IPv6 traffic only, and Secrets Manager does not support IPv6 addresses2.
- B. A NAT gateway is a VPC component that enables instances in a private subnet to connect to the internet or other AWS services, but prevents the internet from initiating connections with those instances4. However, this option does not meet the requirement that the VPC must not send or receive network traffic through the internet. Additionally, a NAT gateway requires an elastic IP address, which is a public IPv4 address4.
- C. A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them using private IPv4 addresses or IPv6 addresses5. However, this option does not work because Secrets Manager does not have a default VPC that can be peered with. Furthermore, a VPC peering connection does not provide a private connection to Secrets Manager APIs without an internet gateway or other devices2.

NEW QUESTION 56
A company is using AWS Organizations to implement a multi-account strategy. The company does not have on-premises infrastructure. All workloads run on AWS. The company currently has eight member accounts. The company anticipates that it will have no more than 20 AWS accounts total at any time.
The company issues a new security policy that contains the following requirements:
• No AWS account should use a VPC within the AWS account for workloads.
• The company should use a centrally managed VPC that all AWS accounts can access to launch workloads in subnets.
• No AWS account should be able to modify another AWS account's application resources within the centrally managed VPC.
• The centrally managed VPC should reside in an existing AWS account that is named Account-A within an organization.
The company uses an AWS CloudFormation template to create a VPC that contains multiple subnets in Account-A. This template exports the subnet IDs through the CloudFormation Outputs section.
Which solution will complete the security setup to meet these requirements?

A. Use a CloudFormation template in the member accounts to launch workload
B. Configure the template to use the Fn::ImportValue function to obtain the subnet ID values.
C. Use a transit gateway in the VPC within Account-
D. Configure the member accounts to use the transit gateway to access the subnets in Account-A to launch workloads.
E. Use AWS Resource Access Manager (AWS RAM) to share Account-A's VPC subnets with the remaining member account
F. Configure the member accounts to use the shared subnets to launch workloads.
G. Create a peering connection between Account-A and the remaining member account
H. Configure the member accounts to use the subnets in Account-A through the VPC peering connection to launch workloads.

Answer: C

Explanation:
The correct answer is C. Use AWS Resource Access Manager (AWS RAM) to share Account-A's VPC
subnets with the remaining member accounts. Configure the member accounts to use the shared subnets to launch workloads.
This answer is correct because AWS RAM is a service that helps you securely share your AWS resources across AWS accounts, within your organization or organizational units (OUs), and with IAM roles and users for supported resource types1. One of the supported resource types is VPC subnets2, which means you can share the subnets in Account-A's VPC with the other member accounts using AWS RAM. This way, you can meet the requirements of using a centrally managed VPC, avoiding duplicate VPCs in each account, and launching workloads in shared subnets. You can also control the access to the shared subnets by using IAM policies and resource-based policies3, which can prevent one account from modifying another account's resources.
The other options are incorrect because:

- A. Using a CloudFormation template in the member accounts to launch workloads and using the Fn::ImportValue function to obtain the subnet ID values is not a solution, because Fn::ImportValue can only import values that have been exported by another stack within the same region4. This means that you cannot use Fn::ImportValue to reference the subnet IDs that are exported by Account-A's CloudFormation template, unless all the member accounts are in the same region as Account-A. This option also does not avoid creating duplicate VPCs in each account, which is one of the requirements.
- B. Using a transit gateway in the VPC within Account-A and configuring the member accounts to use the transit gateway to access the subnets in Account-A to launch workloads is not a solution, because a transit gateway does not allow you to launch workloads in another account's subnets. A transit gateway is a network transit hub that enables you to route traffic between your VPCs and on-premises networks5, but it does not enable you to share subnets across accounts.
- D. Creating a peering connection between Account-A and the remaining member accounts and configuring the member accounts to use the subnets in Account-A through the VPC peering connection to launch workloads is not a solution, because a VPC peering connection does not allow you to launch workloads in another account's subnets. A VPC peering connection is a networking connection between two VPCs that enables you to route traffic between them privately6, but it does not enable you to share subnets across accounts.

References:
1: What is AWS Resource Access Manager? 2: Shareable AWS resources 3: Managing permissions for shared resources 4: Fn::ImportValue 5: What is a transit gateway? 6: What is VPC peering?

NEW QUESTION 59
There is a requirement for a company to transfer large amounts of data between IAM and an on-premise location. There is an additional requirement for low latency and high consistency traffic to IAM. Given these requirements how would you design a hybrid architecture? Choose the correct answer from the options below
Please select:

A. Provision a Direct Connect connection to an IAM region using a Direct Connect partner.
B. Create a VPN tunnel for private connectivity, which increases network consistency and reduces latency.
C. Create an iPSec tunnel for private connectivity, which increases network consistency and reduces latency.
D. Create a VPC peering connection between IAM and the Customer gateway.

Answer: A

Explanation:
IAM Direct Connect makes it easy to establish a dedicated network connection from your premises to IAM. Using IAM Direct Connect you can establish private connectivity between IAM and your datacenter, office, or colocation environment which in many cases can reduce your network costs, increase bandwidth throughput and provide a more consistent network experience than Internet-based connections.
Options B and C are invalid because these options will not reduce network latency Options D is invalid because this is only used to connect 2 VPC's
For more information on IAM direct connect, just browse to the below URL: https://IAM.amazon.com/directconnect
The correct answer is: Provision a Direct Connect connection to an IAM region using a Direct Connect partner. omit your Feedback/Queries to our Experts

NEW QUESTION 61
A company is using IAM Organizations to develop a multi-account secure networking strategy. The company plans to use separate centrally managed accounts for shared services, auditing, and security inspection. The company plans to provide dozens of additional accounts to application owners for production and development environments.
Company security policy requires that all internet traffic be routed through a centrally managed security inspection layer in the security inspection account. A security engineer must recommend a solution that minimizes administrative overhead and complexity.
Which solution meets these requirements?

A. Use IAM Control Towe
B. Modify the default Account Factory networking template to automatically associate new accounts with a centrally managed VPC through a VPC peering connection and to create a default route to the VPC peer in the default route tabl
C. Create an SCP that denies the CreateInternetGateway actio
D. Attach the SCP to all accounts except the security inspection account.
E. Create a centrally managed VPC in the security inspection accoun
F. Establish VPC peering connections between the security inspection account and other account
G. Instruct account owners to create default routes in their account route tables that point to the VPC pee
H. Create an SCP that denies theAttach InternetGateway actio
I. Attach the SCP to all accounts except the security inspection account.
J. Use IAM Control Towe
K. Modify the default Account Factory networking template to automatically associate new accounts with a centrally managed transitgateway and to create a default route to the transit gateway in the default route tabl
L. Create an SCP that denies the AttachInternetGateway actio
M. Attach the SCP to all accounts except the security inspection account.
N. Enable IAM Resource Access Manager (IAM RAM) for IAM Organization
O. Create a shared transit gateway, and make it available by using an IAM RAM resource shar
P. Create an SCP that denies the CreateInternetGateway actio
Q. Attach the SCP to all accounts except the security inspection accoun
R. Create routes in the route tables of all accounts that point to the shared transit gateway.

Answer: C

NEW QUESTION 62
A security engineer needs to configure an Amazon S3 bucket policy to restrict access to an S3 bucket that is named DOC-EXAMPLE-BUCKET. The policy must allow access to only DOC-EXAMPLE-BUCKET from only the following endpoint: vpce-1a2b3c4d. The policy must deny all access to DOC-EXAMPLE-BUCKET if the specified endpoint is not used.
Which bucket policy statement meets these requirements?

A. A computer code with black text Description automatically generated

```
"Statement": [
    {
        "Sid": "Access-to-specific-VPCE-only",
        "Principal": "*",
        "Action": "s3:*",
        "Effect": "Allow",
        "Resource": ["arn:aws:s3:::DOC-EXAMPLE-BUCKET",
                     "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"],
        "Condition": {
            "StringNotEquals": {
                "aws:sourceVpce": "vpce-1a2b3c4d"
            }
        }
    }
]
```

B. A computer code with black text Description automatically generated

```
"Statement": [
    {
        "Sid": "Access-to-specific-VPCE-only",
        "Principal": "*",
        "Action": "s3:*",
        "Effect": "Deny",
        "Resource": ["arn:aws:s3:::DOC-EXAMPLE-BUCKET",
                     "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"],
        "Condition": {
            "StringNotEquals": {
                "aws:sourceVpce": "vpce-1a2b3c4d"
            }
        }
    }
]
```

C. A computer code with black text Description automatically generated

D. A computer code with black text Description automatically generated

```
"Statement": [
    {
        "Sid": "Access-to-specific-VPCE-only",
        "Principal": "*",
        "Action": "s3:*",
        "Effect": "Deny",
        "Resource": ["arn:aws:s3:::DOC-EXAMPLE-BUCKET",
                     "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"],
        "Condition": {
            "StringEquals": {
                "aws:sourceVpce": "vpce-1a2b3c4d"
            }
        }
    }
]
```

```
"Statement": [
    {
        "Sid": "Access-to-specific-VPCE-only",
        "Principal": "*",
        "Action": "s3:*",
        "Effect": "Allow",
        "Resource": ["arn:aws:s3:::DOC-EXAMPLE-BUCKET",
                     "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*"],
        "Condition": {
            "StringEquals": {
                "aws:sourceVpce": "vpce-1a2b3c4d"
            }
        }
    }
]
```

Answer: B

Explanation:
https://docs.aws.amazon.com/AmazonS3/latest/userguide/example-bucket-policies-vpc-endpoint.html

NEW QUESTION 65
A company has thousands of AWS Lambda functions. While reviewing the Lambda functions, a security engineer discovers that sensitive information is being stored in environment variables and is viewable as plaintext in the Lambda console. The values of the sensitive information are only a few characters long. What is the MOST cost-effective way to address this security issue?

A. Set up IAM policies from the Lambda console to hide access to the environment variables.
B. Use AWS Step Functions to store the environment variable
C. Access the environment variables at runtim
D. Use IAM permissions to restrict access to the environment variables to only the Lambda functions that require access.
E. Store the environment variables in AWS Secrets Manager, and access them at runtim
F. Use IAM permissions to restrict access to the secrets to only the Lambda functions that require access.
G. Store the environment variables in AWS Systems Manager Parameter Store as secure string parameters, and access them at runtim
H. Use IAM permissions to restrict access to the parameters to only the Lambda functions that require access.

Answer: D

Explanation:
Storing sensitive information in environment variables is not a secure practice, as anyone who has access to the Lambda console or the Lambda function code can view them as plaintext. To address this security issue, the security engineer needs to use a service that can store and encrypt the environment variables, and access them at runtime using IAM permissions. The most cost-effective way to do this is to use AWS Systems Manager Parameter Store, which is a service that provides secure, hierarchical storage for configuration data management and secrets management. Parameter Store allows you to store values as standard parameters (plaintext) or secure string parameters (encrypted). Secure string parameters use a AWS Key Management Service (AWS KMS) customer master key (CMK) to encrypt the parameter value. To access the parameter value at runtime, the Lambda function needs to have IAM permissions to decrypt the parameter using the KMS CMK.
The other options are incorrect because:

- Option A is incorrect because setting up IAM policies from the Lambda console to hide access to the environment variables will not prevent someone who has access to the Lambda function code from viewing them as plaintext. IAM policies can only control who can perform actions on AWS resources, not what they can see in the code or the console.
- Option B is incorrect because using AWS Step Functions to store the environment variables is not a secure or cost-effective solution. AWS Step Functions is a service that lets you coordinate multiple AWS services into serverless workflows. Step Functions does not provide any encryption or secrets management capabilities, and it will incur additional charges for each state transition in the workflow. Moreover, storing environment variables in Step Functions will make them visible in the execution history of the workflow, which can be accessed by anyone who has permission to view the Step Functions console or API.
- Option C is incorrect because storing the environment variables in AWS Secrets Manager and accessing them at runtime is not a cost-effective solution. AWS Secrets Manager is a service that helps you protect secrets needed to access your applications, services, and IT resources. Secrets Manager enables you to rotate, manage, and retrieve secrets throughout their lifecycle. While Secrets Manager can securely store and encrypt environment variables using KMS CMKs, it will incur higher charges than Parameter Store for storing and retrieving secrets. Unless the security engineer needs the advanced features of Secrets Manager, such as automatic rotation of secrets or integration with other AWS services, Parameter Store is a cheaper and simpler option.

NEW QUESTION 68
A company needs to store multiple years of financial records. The company wants to use Amazon S3 to store copies of these documents. The company must implement a solution to prevent the documents from being edited, replaced, or deleted for 7 years after the documents are stored in Amazon S3. The solution must

also encrypt the documents at rest.
A security engineer creates a new S3 bucket to store the documents. What should the security engineer do next to meet these requirements?

A. Configure S3 server-side encryptio
B. Create an S3 bucket policy that has an explicit deny rule for all users for s3:DeleteObject and s3:PutObject API call
C. Configure S3 Object Lock to use governance mode with a retention period of 7 years.
D. Configure S3 server-side encryptio
E. Configure S3 Versioning on the S3 bucke
F. Configure S3 ObjectLock to use compliance mode with a retention period of 7 years.
G. Configure S3 Versionin
H. Configure S3 Intelligent-Tiering on the S3 bucket to move the documents to S3 Glacier Deep Archive storag
I. Use S3 server-side encryption immediatel
J. Expire the objects after 7 years.
K. Set up S3 Event Notifications and use S3 server-side encryptio
L. Configure S3 Event Notifications to target an AWS Lambda function that will review any S3 API call to the S3 bucket and deny the s3:DeleteObject and s3:PutObject API call
M. Remove the S3 event notification after 7 years.

Answer: B

NEW QUESTION 70
A company has a web server in the AWS Cloud. The company will store the content for the web server in an Amazon S3 bucket. A security engineer must use an Amazon CloudFront distribution to speed up delivery of the content. None of the files can be publicly accessible from the S3 bucket direct.
Which solution will meet these requirements?

A. Configure the permissions on the individual files in the S3 bucket so that only the CloudFront distribution has access to them.
B. Create an origin access identity (OAI). Associate the OAI with the CloudFront distributio
C. Configure the S3 bucket permissions so that only the OAI can access the files in the S3 bucket.
D. Create an S3 role in AWS Identity and Access Management (IAM). Allow only the CloudFront distribution to assume the role to access the files in the S3 bucket.
E. Create an S3 bucket policy that uses only the CloudFront distribution ID as the principal and the Amazon Resource Name (ARN) as the target.

Answer: B

NEW QUESTION 72
A developer at a company uses an SSH key to access multiple Amazon EC2 instances. The company discovers that the SSH key has been posted on a public GitHub repository. A security engineer verifies that the key has not been used recently.
How should the security engineer prevent unauthorized access to the EC2 instances?

A. Delete the key pair from the EC2 consol
B. Create a new key pair.
C. Use the ModifyInstanceAttribute API operation to change the key on any EC2 instance that is using the key.
D. Restrict SSH access in the security group to only known corporate IP addresses.
E. Update the key pair in any AMI that is used to launch the EC2 instance
F. Restart the EC2 instances.

Answer: C

Explanation:
To prevent unauthorized access to the EC2 instances, the security engineer should do the following:
> Restrict SSH access in the security group to only known corporate IP addresses. This allows the security engineer to use a virtual firewall that controls inbound and outbound traffic for their EC2 instances, and limit SSH access to only trusted sources.

NEW QUESTION 74
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
```
"Version":"2012-10-17",
"Id":"PutObj",
"Statement":[{
"Sid":"DenyUploads",
"Effect":"Deny",
"Principal":"*",
"Action":"s3:PutObject",
"Resource":"arn:aws:s3:::demo/*",
"Condition":{
"StringNotEquals":{
"s3:x-amz-server-side-encryption":"aws:kms"
}
}
}
]
}
```
B. C:\Users\wk\Desktop\mudassar\Untitled.jpg

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

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

Answer: A

Explanation:
The condition of "s3:x-amz-server-side-encryption":"IAM:kms" ensures that objects uploaded need to be encrypted.
Options B,C and D are invalid because you have to ensure the condition of ns3:x-amz-server-side-encryption":"IAM:kms" is present
For more information on IAM KMS best practices, just browse to the below URL: https://dl.IAMstatic.com/whitepapers/IAM-kms-best-praaices.pdf
Submit your Feedback/Queries to our Expert

NEW QUESTION 77
A security engineer needs to run an AWS CloudFormation script. The CloudFormation script builds AWS infrastructure to support a stack that includes web servers and a MySQL database. The stack has been deployed in pre-production environments and is ready for production.
The production script must comply with the principle of least privilege. Additionally, separation of duties must exist between the security engineer's IAM account and CloudFormation.
Which solution will meet these requirements?

A. Use IAM Access Analyzer policy generation to generate a policy that allows the CloudFormation script to run and manage the stac
B. Attach the policy to a new IAM rol
C. Modify the security engineer's IAM permissions to be able to pass the new role to CloudFormation.
D. Create an IAM policy that allows ec2:* and rds:* permission
E. Attach the policy to a new IAM role.Modify the security engineer's IAM permissions to be able to assume the new role.
F. Use IAM Access Analyzer policy generation to generate a policy that allows the CloudFormation script to run and manage the stac
G. Modify the security engineer's IAM permissions to be able to run the CloudFormation script.
H. Create an IAM policy that allows ec2:* and rds:* permission
I. Attach the policy to a new IAM rol
J. Use the IAM policy simulator to confirm that the policy allows the AWS API calls that are necessary to build the stac

K. Modify the security engineer's IAM permissions to be able to pass the new role to CloudFormation.

Answer: A

Explanation:
The correct answer is A. Use IAM Access Analyzer policy generation to generate a policy that allows the CloudFormation script to run and manage the stack. Attach the policy to a new IAM role. Modify the security engineer's IAM permissions to be able to pass the new role to CloudFormation.
According to the AWS documentation, IAM Access Analyzer is a service that helps you identify the resources in your organization and accounts, such as Amazon S3 buckets or IAM roles, that are shared with an external entity. You can also use IAM Access Analyzer to generate fine-grained policies that grant least privilege access based on access activity and access attempts.
To use IAM Access Analyzer policy generation, you need to enable IAM Access Analyzer in your account or organization. You can then use the IAM console or the AWS CLI to generate a policy for a resource based on its access activity or access attempts. You can review and edit the generated policy before applying it to the resource.
To use IAM Access Analyzer policy generation with CloudFormation, you can follow these steps:

- Run the CloudFormation script in a pre-production environment and monitor its access activity or access attempts using IAM Access Analyzer.
- Use IAM Access Analyzer policy generation to generate a policy that allows the CloudFormation script to run and manage the stack. The policy will include only the permissions that are necessary for the script to function.
- Attach the policy to a new IAM role that has a trust relationship with CloudFormation. This will allow CloudFormation to assume the role and execute the script.
- Modify the security engineer's IAM permissions to be able to pass the new role to CloudFormation.

This will allow the security engineer to launch the stack using the role.

- Run the CloudFormation script in the production environment using the new role.

This solution will meet the requirements of least privilege and separation of duties, as it will limit the permissions of both CloudFormation and the security engineer to only what is needed for running and managing the stack.
Option B is incorrect because creating an IAM policy that allows ec2:* and rds:* permissions is not following the principle of least privilege, as it will grant more permissions than necessary for running and managing the stack. Moreover, modifying the security engineer's IAM permissions to be able to assume the new role is not ensuring separation of duties, as it will allow the security engineer to bypass CloudFormation and directly access the resources.
Option C is incorrect because modifying the security engineer's IAM permissions to be able to run the CloudFormation script is not ensuring separation of duties, as it will allow the security engineer to execute the script without using CloudFormation.
Option D is incorrect because creating an IAM policy that allows ec2:* and rds:* permissions is not following the principle of least privilege, as it will grant more permissions than necessary for running and managing the stack. Using the IAM policy simulator to confirm that the policy allows the AWS API calls that are necessary to build the stack is not sufficient, as it will not generate a fine-grained policy based on access activity or access attempts.

NEW QUESTION 82
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 accoun
B. Configure the CloudHSM security group to accept inbound traffic from the private IP addresses of client instances in the new dedicated account.
C. 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.
D. 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 accoun
E. Use thecross-account permissions that are assigned to the STS token to invoke an operation on the HSM in thecentral account.
F. 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 accoun
G. 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:
https://aws.amazon.com/premiumsupport/knowledge-center/cloudhsm-share-clusters/#:~:text=In%20the%20nav

NEW QUESTION 86
A company's security engineer is designing an isolation procedure for Amazon EC2 instances as part of an incident response plan. The security engineer needs to isolate a target instance to block any traffic to and from the target instance, except for traffic from the company's forensics team. Each of the company's EC2 instances has its own dedicated security group. The EC2 instances are deployed in subnets of a VPC. A subnet can contain multiple instances.
The security engineer is testing the procedure for EC2 isolation and opens an SSH session to the target instance. The procedure starts to simulate access to the target instance by an attacker. The security engineer removes the existing security group rules and adds security group rules to give the forensics team access to the target instance on port 22.
After these changes, the security engineer notices that the SSH connection is still active and usable. When the security engineer runs a ping command to the public IP address of the target instance, the ping command is blocked.
What should the security engineer do to isolate the target instance?

A. Add an inbound rule to the security group to allow traffic from 0.0.0.0/0 for all port
B. Add an outbound rule to the security group to allow traffic to 0.0.0.0/0 for all port
C. Then immediately delete these rules.
D. Remove the port 22 security group rul
E. Attach an instance role policy that allows AWS Systems Manager Session Manager connections so that the forensics team can access the target instance.
F. Create a network ACL that is associated with the target instance's subne
G. Add a rule at the top of the inbound rule set to deny all traffic from 0.0.0.0/0. Add a rule at the top of the outbound rule set to deny all traffic to 0.0.0.0/0.
H. Create an AWS Systems Manager document that adds a host-level firewall rule to block all inbound traffic and outbound traffi
I. Run the document on the target instance.

Answer: C

NEW QUESTION 91
A security engineer logs in to the AWS Lambda console with administrator permissions. The security engineer is trying to view logs in Amazon CloudWatch for a Lambda function that is named my Function.

When the security engineer chooses the option in the Lambda console to view logs in CloudWatch, an “error loading Log Streams" message appears.
The IAM policy for the Lambda function's execution role contains the following:

```
{
                "Version": "2012-10-17",
                "Statement": [
                  {
                       "Effect": "Allow",
                       "Action": "logs:CreateLogGroup",
                       "Resource": "arn:aws:logs:us-east-1:111111111111:*"
                  },
                  {
                       "Effect": "Allow",
                       "Action": ["logs:PutLogEvents"],
                       "Resource": ["arn:aws:logs:us-east-1:111111111111:log-
group:/aws/Lambda/myFunction:*"]
                  }
          ]
}
```

How should the security engineer correct the error?

A. Move the logs:CreateLogGroup action to the second Allow statement.
B. Add the logs:PutDestination action to the second Allow statement.
C. Add the logs:GetLogEvents action to the second Allow statement.
D. Add the logs:CreateLogStream action to the second Allow statement.

Answer: D

Explanation:
CloudWatchLogsReadOnlyAccess doesn't include "logs:CreateLogStream" but it includes "logs:Get*"
https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/iam-identity-based-access-control-cwl.html#:~:te

NEW QUESTION 94
A company's policy requires that all API keys be encrypted and stored separately from source code in a centralized security account. This security account is managed by the company's security team However, an audit revealed that an API key is steed with the source code of an IAM Lambda function m an IAM CodeCommit repository in the DevOps account
How should the security learn securely store the API key?

A. Create a CodeCommit repository in the security account using IAM Key Management Service (IAMKMS) tor encryption Require the development team to migrate the Lambda source code to this repository
B. Store the API key in an Amazon S3 bucket in the security account using server-side encryption with Amazon S3 managed encryption keys (SSE-S3) to encrypt the key Create a resigned URL tor the S3 ke
C. and specify the URL m a Lambda environmental variable in the IAM CloudFormation template Update the Lambda function code to retrieve the key using the URL and call the API
D. Create a secret in IAM Secrets Manager in the security account to store the API key using IAM Key Management Service (IAM KMS) tor encryption Grant access to the IAM role used by the Lambda function so that the function can retrieve the key from Secrets Manager and call the API
E. Create an encrypted environment variable for the Lambda function to store the API key using IAM Key Management Service (IAM KMS) tor encryption Grant access to the IAM role used by the Lambda function so that the function can decrypt the key at runtime

Answer: C

Explanation:
To securely store the API key, the security team should do the following:

- Create a secret in AWS Secrets Manager in the security account to store the API key using AWS Key Management Service (AWS KMS) for encryption. This allows the security team to encrypt and manage the API key centrally, and to configure automatic rotation schedules for it.
- Grant access to the IAM role used by the Lambda function so that the function can retrieve the key from Secrets Manager and call the API. This allows the security team to avoid storing the API key with the source code, and to use IAM policies to control access to the secret.

NEW QUESTION 95
A company needs to encrypt all of its data stored in Amazon S3. The company wants to use IAM Key Management Service (IAM KMS) to create and manage its encryption keys. The company's security policies require the ability to Import the company's own key material for the keys, set an expiration date on the keys, and delete keys immediately, if needed.
How should a security engineer set up IAM KMS to meet these requirements?

A. Configure IAM KMS and use a custom key stor
B. Create a customer managed CMK with no key material Import the company's keys and key material into the CMK
C. Configure IAM KMS and use the default Key store Create an IAM managed CMK with no key material Import the company's key material into the CMK
D. Configure IAM KMS and use the default key store Create a customer managed CMK with no key material import the company's key material into the CMK
E. Configure IAM KMS and use a custom key stor
F. Create an IAM managed CMK with no key material.Import the company's key material into the CMK.

Answer: A

Explanation:
To meet the requirements of importing their own key material, setting an expiration date on the keys, and deleting keys immediately, the security engineer should do the following:

- Configure AWS KMS and use a custom key store. This allows the security engineer to use a key manager outside of AWS KMS that they own and manage, such as an AWS CloudHSM cluster or an external key manager.
- Create a customer managed CMK with no key material. Import the company’s keys and key material into the CMK. This allows the security engineer to use their own key material for encryption and decryption operations, and to specify an expiration date for it.

NEW QUESTION 96
A security engineer is designing a cloud architecture to support an application. The application runs on Amazon EC2 instances and processes sensitive information, including credit card numbers.
The application will send the credit card numbers to a component that is running in an isolated environment. The component will encrypt, store, and decrypt the numbers.
The component then will issue tokens to replace the numbers in other parts of the application.
The component of the application that manages the tokenization process will be deployed on a separate set of EC2 instances. Other components of the application must not be able to store or access the credit card numbers.
Which solution will meet these requirements?

A. Use EC2 Dedicated Instances for the tokenization component of the application.
B. Place the EC2 instances that manage the tokenization process into a partition placement group.
C. Create a separate VP
D. Deploy new EC2 instances into the separate VPC to support the data tokenization.
E. Deploy the tokenization code onto AWS Nitro Enclaves that are hosted on EC2 instances.

Answer: D

Explanation:
AWS Nitro Enclaves are isolated and hardened virtual machines that run on EC2 instances and provide a secure environment for processing sensitive data. Nitro Enclaves have no persistent storage, interactive access, or external networking, and they can only communicate with the parent instance through a secure local channel. Nitro Enclaves also support cryptographic attestation, which allows verifying the identity and integrity of the enclave and its code. Nitro Enclaves are ideal for implementing data protection solutions such as tokenization, encryption, and key management.
Using Nitro Enclaves for the tokenization component of the application meets the requirements of isolating the sensitive data from other parts of the application, encrypting and storing the credit card numbers securely, and issuing tokens to replace the numbers. Other components of the application will not be able to access or store the credit card numbers, as they are only available within the enclave.

NEW QUESTION 97
A company has deployed Amazon GuardDuty and now wants to implement automation for potential threats. The company has decided to start with RDP brute force attacks that come from Amazon EC2 instances in the company's AWS environment. A security engineer needs to implement a solution that blocks the detected communication from a suspicious instance until investigation and potential remediation can occur.
Which solution will meet these requirements?

A. Configure GuardDuty to send the event to an Amazon Kinesis data strea
B. Process the event with an Amazon Kinesis Data Analytics for Apache Flink application that sends a notification to the company through Amazon Simple Notification Service (Amazon SNS). Add rules to the network ACL to block traffic to and from the suspicious instance.
C. Configure GuardDuty to send the event to Amazon EventBridge (Amazon CloudWatch Events). Deploy an AWS WAF web AC
D. Process the event with an AWS Lambda function that sends a notification to the company through Amazon Simple Notification Service (Amazon SNS) and adds a web ACL rule to block traffic to and from the suspicious instance.
E. Enable AWS Security Hub to ingest GuardDuty findings and send the event to Amazon EventBridge (Amazon CloudWatch Events). Deploy AWS Network Firewal
F. Process the event with an AWS Lambda function that adds a rule to a Network Firewall firewall policy to block traffic to and from the suspicious instance.
G. Enable AWS Security Hub to ingest GuardDuty finding
H. Configure an Amazon Kinesis data stream as an event destination for Security Hu
I. Process the event with an AWS Lambda function that replaces the security group of the suspicious instance with a security group that does not allow any connections.

Answer: C

Explanation:
https://aws.amazon.com/blogs/security/automatically-block-suspicious-traffic-with-aws-network-firewall-and-a

NEW QUESTION 102
A company is operating a website using Amazon CloudFornt. CloudFront servers some content from Amazon S3 and other from web servers running EC2 instances behind an Application. Load Balancer (ALB). Amazon DynamoDB is used as the data store. The company already uses IAM Certificate Manager (ACM) to store a public TLS certificate that can optionally secure connections between the website users and CloudFront. The company has a new requirement to enforce end-to-end encryption in transit.
Which combination of steps should the company take to meet this requirement? (Select THREE.)

A. Update the CloudFront distributio
B. configuring it to optionally use HTTPS when connecting to origins on Amazon S3
C. Update the web application configuration on the web servers to use HTTPS instead of HTTP when connecting to DynamoDB
D. Update the CloudFront distribution to redirect HTTP corrections to HTTPS
E. Configure the web servers on the EC2 instances to listen using HTTPS using the public ACM TLS certificate Update the ALB to connect to the target group using HTTPS
F. Update the ALB listen to listen using HTTPS using the public ACM TLS certificat
G. Update the CloudFront distribution to connect to the HTTPS listener.
H. Create a TLS certificate Configure the web servers on the EC2 instances to use HTTPS only with that certificat
I. Update the ALB to connect to the target group using HTTPS.

Answer: BCE

Explanation:
To enforce end-to-end encryption in transit, the company should do the following:

- Update the web application configuration on the web servers to use HTTPS instead of HTTP when connecting to DynamoDB. This ensures that the data is encrypted when it travels from the web servers to the data store.
- Update the CloudFront distribution to redirect HTTP requests to HTTPS. This ensures that the viewers always use HTTPS when they access the website through CloudFront.
- Update the ALB to listen using HTTPS using the public ACM TLS certificate. Update the CloudFront distribution to connect to the HTTPS listener. This ensures that the data is encrypted when it travels from CloudFront to the ALB and from the ALB to the web servers.

NEW QUESTION 105
A company wants to configure DNS Security Extensions (DNSSEC) for the company's primary domain. The company registers the domain with Amazon Route 53. The company hosts the domain on Amazon EC2 instances by using BIND.
What is the MOST operationally efficient solution that meets this requirement?

A. Set the dnssec-enable option to yes in the BIND configuratio
B. Create a zone-signing key (ZSK) and a key-signing key (KSK) Restart the BIND service.
C. Migrate the zone to Route 53 with DNSSEC signing enable
D. Create a zone-signing key (ZSK) and a key-signing key (KSK) that are based on an AW
E. Key Management Service (AWS KMS) customer managed key.
F. Set the dnssec-enable option to yes in the BIND configuratio
G. Create a zone-signing key (ZSK) and a key-signing key (KSK). Run the dnssec-signzone command to generate a delegation signer (DS) record Use AW
H. Key Management Service (AWS KMS) to secure the keys.
I. Migrate the zone to Route 53 with DNSSEC signing enable
J. Create a key-signing key (KSK) that is based on an AWS Key Management Service (AWS KMS) customer managed ke
K. Add a delegation signer (DS) record to the parent zone.

Answer: D

Explanation:
To configure DNSSEC for a domain registered with Route 53, the most operationally efficient solution is to migrate the zone to Route 53 with DNSSEC signing enabled, create a key-signing key (KSK) that is based on an AWS Key Management Service (AWS KMS) customer managed key, and add a delegation signer (DS) record to the parent zone. This way, Route 53 handles the zone-signing key (ZSK) and the signing of the records in the hosted zone, and the customer only needs to manage the KSK in AWS KMS and provide the DS record to the domain registrar. Option A is incorrect because it does not involve migrating the zone to Route 53, which would simplify the DNSSEC configuration. Option B is incorrect because it creates both a ZSK and a KSK based on AWS KMS customer managed keys, which is unnecessary and less efficient than letting Route 53 manage the ZSK. Option C is incorrect because it does not involve migrating the zone to Route 53, and it requires running the dnssec-signzone command manually, which is less efficient than letting Route 53 sign the zone automatically. Verified References:

- https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/domain-configure-dnssec.html
- https://aws.amazon.com/about-aws/whats-new/2020/12/announcing-amazon-route-53-support-dnssec/

NEW QUESTION 110
A company uses a third-party application to store encrypted data in Amazon S3. The company uses another third-party application trial decrypts the data from Amazon S3 to ensure separation of duties Between the applications A Security Engineer warns to separate the permissions using IAM roles attached to Amazon EC2 instances. The company prefers to use native IAM services.
Which encryption method will meet these requirements?

A. Use encrypted Amazon EBS volumes with Amazon default keys (IAM EBS)
B. Use server-side encryption with customer-provided keys (SSE-C)
C. Use server-side encryption with IAM KMS managed keys (SSE-KMS)
D. Use server-side encryption with Amazon S3 managed keys (SSE-S3)

Answer: C

NEW QUESTION 113
An ecommerce company is developing new architecture for an application release. The company needs to implement TLS for incoming traffic to the application. Traffic for the application will originate from the internet TLS does not have to be implemented in an end-to-end configuration because the company is concerned about impacts on performance. The incoming traffic types will be HTTP and HTTPS The application uses ports 80 and 443.
What should a security engineer do to meet these requirements?

A. Create a public Application Load Balance
B. Create two listeners one listener on port 80 and one listener on port 443. Create one target grou
C. Create a rule to forward traffic from port 80 to the listener on port 443 Provision a public TLS certificate in AWS Certificate Manager (ACM). Attach the certificate to the listener on port 443.
D. Create a public Application Load Balance
E. Create two listeners one listener on port 80 and one listener on port 443. Create one target grou
F. Create a rule to forward traffic from port 80 to the listener on port 443 Provision a public TLS certificate in AWS Certificate Manager (ACM). Attach the certificate to the listener on port 80.
G. Create a public Network Load Balance
H. Create two listeners one listener on port 80 and one listener on port 443. Create one target grou
I. Create a rule to forward traffic from port 80 to the listener on port 443. Set the protocol for the listener on port 443 to TLS.
J. Create a public Network Load Balance
K. Create a listener on port 443. Create one target grou
L. Create a rule to forward traffic from port 443 to the target grou
M. Set the protocol for the listener on port 443 to TLS.

Answer: A

Explanation:
An Application Load Balancer (ALB) is a type of load balancer that operates at the application layer (layer 7) of the OSI model. It can distribute incoming traffic based on the content of the request, such as the host
header, path, or query parameters. An ALB can also terminate TLS connections and decrypt requests from clients before sending them to the targets.
To implement TLS for incoming traffic to the application, the following steps are required:

- Create a public ALB in a public subnet and register the EC2 instances as targets in a target group.
- Create two listeners for the ALB, one on port 80 for HTTP traffic and one on port 443 for HTTPS traffic.
- Create a rule for the listener on port 80 to redirect HTTP requests to HTTPS using the same host, path, and query parameters.
- Provision a public TLS certificate in AWS Certificate Manager (ACM) for the domain name of the application. ACM is a service that lets you easily provision, manage, and deploy public and private SSL/TLS certificates for use with AWS services and your internal connected resources.

- Attach the certificate to the listener on port 443 and configure the security policy to negotiate secure connections between clients and the ALB.
- Configure the security groups for the ALB and the EC2 instances to allow inbound traffic on ports 80 and 443 from the internet and outbound traffic on any port to the EC2 instances.

This solution will meet the requirements of implementing TLS for incoming traffic without impacting performance or requiring end-to-end encryption. The ALB will handle the TLS termination and decryption, while forwarding unencrypted requests to the EC2 instances.
Verified References:

- https://docs.aws.amazon.com/elasticloadbalancing/latest/application/introduction.html
- https://docs.aws.amazon.com/elasticloadbalancing/latest/application/create-https-listener.html
- https://docs.aws.amazon.com/acm/latest/userguide/acm-overview.html

NEW QUESTION 117
A security team is using Amazon EC2 Image Builder to build a hardened AMI with forensic capabilities. An AWS Key Management Service (AWS KMS) key will encrypt the forensic AMI EC2 Image Builder successfully installs the required patches and packages in the security team's AWS account. The security team uses a federated IAM role m the same AWS account to sign in to the AWS Management Console and attempts to launch the forensic AMI. The EC2 instance launches and immediately terminates.
What should the security learn do lo launch the EC2 instance successfully

A. Update the policy that is associated with the federated IAM role to allow the ec2. Describelmages action for the forensic AMI.
B. Update the policy that is associated with the federated IAM role to allow the ec2 Start Instances action m the security team's AWS account.
C. Update the policy that is associated with the KMS key that is used to encrypt the forensic AM
D. Configure the policy to allow the km
E. Encrypt and kms Decrypt actions for the federated IAM role.
F. Update the policy that is associated with the federated IAM role to allow the km
G. DescribeKey action for the KMS key that is used to encrypt the forensic AMI.

Answer: C

Explanation:
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/troubleshooting-launch.html#troubleshooting-launch-i

NEW QUESTION 120
A Development team has built an experimental environment to test a simple stale web application It has built an isolated VPC with a private and a public subnet. The public subnet holds only an Application Load Balancer a NAT gateway, and an internet gateway. The private subnet holds ail of the Amazon EC2 instances There are 3 different types of servers Each server type has its own Security Group that limits access lo only required connectivity. The Security Groups nave both inbound and outbound rules applied Each subnet has both inbound and outbound network ACls applied to limit access to only required connectivity
Which of the following should the team check if a server cannot establish an outbound connection to the
internet? (Select THREE.)

A. The route tables and the outbound rules on the appropriate private subnet security group
B. The outbound network ACL rules on the private subnet and the Inbound network ACL rules on the public subnet
C. The outbound network ACL rules on the private subnet and both the inbound and outbound rules on the public subnet
D. The rules on any host-based firewall that may be applied on the Amazon EC2 instances
E. The Security Group applied to the Application Load Balancer and NAT gateway
F. That the 0.0.0./0 route in the private subnet route table points to the internet gateway in the public subnet

Answer: CEF

Explanation:
because these are the factors that could affect the outbound connection to the internet from a server in a private subnet. The outbound network ACL rules on the private subnet and both the inbound and outbound rules on the public subnet must allow the traffic to pass through8. The security group applied to the application load balancer and NAT gateway must also allow the traffic from the private subnet9. The 0.0.0.0/0 route in the private subnet route table must point to the NAT gateway in the public subnet, not the internet gateway10. The other options are either irrelevant or incorrect for troubleshooting the outbound connection issue.

NEW QUESTION 121
A company plans to create individual child accounts within an existing organization in IAM Organizations for each of its DevOps teams. IAM CloudTrail has been enabled and configured on all accounts to write audit logs to an Amazon S3 bucket in a centralized IAM 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 IAM account root user.
B. Create an S3 bucket policy in the specified destination account for the CloudTrail trail that prohibits configuration changes from the IAM 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 grou
E. Have team members use individual IAM accounts that are members of the new IAM group.

Answer: D

NEW QUESTION 123
A company's security engineer has been tasked with restricting a contractor's IAM account access to the company's Amazon EC2 console without providing access to any other IAM services The contractors IAM account must not be able to gain access to any other IAM service, even it the IAM account rs assigned additional permissions based on IAM group membership
What should the security engineer do to meet these requirements"

A. Create an mime IAM user policy that allows for Amazon EC2 access for the contractor's IAM user
B. Create an IAM permissions boundary policy that allows Amazon EC2 access Associate the contractor's IAM account with the IAM permissions boundary policy
C. Create an IAM group with an attached policy that allows for Amazon EC2 access Associate the contractor's IAM account with the IAM group

D. Create a IAM role that allows for EC2 and explicitly denies all other services Instruct the contractor to always assume this role

Answer: B

Explanation:
To restrict the contractor's IAM account access to the EC2 console without providing access to any other AWS services, the security engineer should do the following:

- Create an IAM permissions boundary policy that allows EC2 access. This is a policy that defines the maximum permissions that an IAM entity can have.
- Associate the contractor's IAM account with the IAM permissions boundary policy. This means that even if the contractor's IAM account is assigned additional permissions based on IAM group membership, those permissions are limited by the permissions boundary policy.

NEW QUESTION 126
An application is running on an Amazon EC2 instance that has an IAM role attached. The IAM role provides access to an AWS Key Management Service (AWS KMS) customer managed key and an Amazon S3 bucket. The key is used to access 2 TB of sensitive data that is stored in the S3 bucket.
A security engineer discovers a potential vulnerability on the EC2 instance that could result in the compromise of the sensitive data. Due to other critical operations, the security engineer cannot immediately shut down the EC2 instance for vulnerability patching.
What is the FASTEST way to prevent the sensitive data from being exposed?

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

Answer: D

NEW QUESTION 129
A company wants to ensure that its IAM resources can be launched only in the us-east-1 and us-west-2 Regions.
What is the MOST operationally efficient solution that will prevent developers from launching Amazon EC2 instances in other Regions?

A. Enable Amazon GuardDuty in all Region
B. Create alerts to detect unauthorized activity outside us-east-1 and us-west-2.
C. Use an organization in IAM Organization
D. Attach an SCP that allows all actions when the IAM: Requested Region condition key is either us-east-1 or us-west-2. Delete the FullIAMAccess policy.
E. Provision EC2 resources by using IAM Cloud Formation templates through IAM CodePipelin
F. Allow only the values of us-east-1 and us-west-2 in the IAM CloudFormation template's parameters.
G. Create an IAM Config rule to prevent unauthorized activity outside us-east-1 and us-west-2.

Answer: C

NEW QUESTION 132
A company hosts a web application on an Apache web server. The application runs on Amazon EC2 instances that are in an Auto Scaling group. The company configured the EC2 instances to send the Apache web server logs to an Amazon CloudWatch Logs group that the company has configured to expire after 1 year.
Recently, the company discovered in the Apache web server logs that a specific IP address is sending suspicious requests to the web application. A security engineer wants to analyze the past week of Apache web server logs to determine how many requests that the IP address sent and the corresponding URLs that the IP address requested.
What should the security engineer do to meet these requirements with the LEAST effort?

A. Export the CloudWatch Logs group data to Amazon S3. Use Amazon Macie to query the logs for the specific IP address and the requested URLs.
B. Configure a CloudWatch Logs subscription to stream the log group to an Am-azon OpenSearch Service cluste
C. Use OpenSearch Service to analyze the logs for the specific IP address and the requested URLs.
D. Use CloudWatch Logs Insights and a custom query syntax to analyze the CloudWatch logs for the specific IP address and the requested URLs.
E. Export the CloudWatch Logs group data to Amazon S3. Use AWS Glue to crawl the S3 bucket for only the log entries that contain the specific IP ad-dres
F. Use AWS Glue to view the results.

Answer: C

NEW QUESTION 136
A company needs complete encryption of the traffic between external users and an application. The company hosts the application on a fleet of Amazon EC2 instances that run in an Auto Scaling group behind an Application Load Balancer (ALB).
How can a security engineer meet these requirements?

A. Create a new Amazon-issued certificate in AWS Secrets Manage
B. Export the certificate from Secrets Manage
C. Import the certificate into the ALB and the EC2 instances.
D. Create a new Amazon-issued certificate in AWS Certificate Manager (ACM). Associate the certificate with the AL
E. Export the certificate from AC
F. Install the certificate on the EC2 instances.
G. Import a new third-party certificate into AWS Identity and Access Management (IAM). Export thecertificate from IA
H. Associate the certificate with the ALB and the EC2 instances.
I. Import a new third-party certificate into AWS Certificate Manager (ACM). Associate the certificate with the AL

J. Install the certificate on the EC2 instances.

Answer: D

Explanation:
The correct answer is D. Import a new third-party certificate into AWS Certificate Manager (ACM). Associate the certificate with the ALB. Install the certificate on the EC2 instances.
This answer is correct because it meets the requirements of complete encryption of the traffic between external users and the application. By importing a third-party certificate into ACM, the security engineer can use it to secure the communication between the ALB and the clients. By installing the same certificate on the EC2 instances, the security engineer can also secure the communication between the ALB and the instances. This way, both the front-end and back-end connections are encrypted with SSL/TLS1.
The other options are incorrect because:

- A. Creating a new Amazon-issued certificate in AWS Secrets Manager is not a solution, because AWS Secrets Manager is not a service for issuing certificates, but for storing and managing secrets such as database credentials and API keys2. AWS Secrets Manager does not integrate with ALB or EC2 for certificate deployment.
- B. Creating a new Amazon-issued certificate in AWS Certificate Manager (ACM) and exporting it from ACM is not a solution, because ACM does not allow exporting Amazon-issued certificates3. ACM only allows exporting private certificates that are issued by an AWS Private Certificate Authority (CA)4.
- C. Importing a new third-party certificate into AWS Identity and Access Management (IAM) is not a solution, because IAM is not a service for managing certificates, but for controlling access to AWS resources5. IAM does not integrate with ALB or EC2 for certificate deployment.

References:
1: How SSL/TLS works 2: What is AWS Secrets Manager? 3: Exporting an ACM Certificate 4: Exporting Private Certificates from ACM 5: What is IAM?

NEW QUESTION 137
A security engineer is trying to use Amazon EC2 Image Builder to create an image of an EC2 instance. The security engineer has configured the pipeline to send logs to an Amazon S3 bucket. When the security engineer runs the pipeline, the build fails with the following error: “AccessDenied: Access Denied status code: 403”.
The security engineer must resolve the error by implementing a solution that complies with best practices for least privilege access.
Which combination of steps will meet these requirements? (Choose two.)

A. Ensure that the following policies are attached to the IAM role that the security engineer is using: EC2InstanceProfileForImageBuilder, EC2InstanceProfileForImageBuilderECRContainerBuilds, and AmazonSSMManagedInstanceCore.
B. Ensure that the following policies are attached to the instance profile for the EC2 instance: EC2InstanceProfileForImageBuilder, EC2InstanceProfileForImageBuilderECRContainerBuilds, and AmazonSSMManagedInstanceCore.
C. Ensure that the AWSImageBuilderFullAccess policy is attached to the instance profile for the EC2 instance.
D. Ensure that the security engineer’s IAM role has the s3:PutObject permission for the S3 bucket.
E. Ensure that the instance profile for the EC2 instance has the s3:PutObject permission for the S3 bucket.

Answer: BE

Explanation:
The most likely cause of the error is that the instance profile for the EC2 instance does not have the s3:PutObject permission for the S3 bucket. This permission is needed to upload logs to the bucket. Therefore, the security engineer should ensure that the instance profile has this permission.
One possible solution is to attach the AWSImageBuilderFullAccess policy to the instance profile for the EC2 instance. This policy grants full access to Image Builder resources and related AWS services, including the s3:PutObject permission for any bucket with “imagebuilder” in its name. However, this policy may grant more permissions than necessary, which violates the principle of least privilege.
Another possible solution is to create a custom policy that only grants the s3:PutObject permission for the specific S3 bucket that is used for logging. This policy can be attached to the instance profile along with the other policies that are required for Image Builder functionality: EC2InstanceProfileForImageBuilder, EC2InstanceProfileForImageBuilderECRContainerBuilds, and AmazonSSMManagedInstanceCore. This solution follows the principle of least privilege more closely than the previous one.

- Ensure that the following policies are attached to the instance profile for the EC2 instance: EC2InstanceProfileForImageBuilder, EC2InstanceProfileForImageBuilderECRContainerBuilds, and AmazonSSMManagedInstanceCore.
- Ensure that the instance profile for the EC2 instance has the s3:PutObject permission for the S3 bucket.

This can be done by either attaching the AWSImageBuilderFullAccess policy or creating a custom policy with this permission.
1: Using managed policies for EC2 Image Builder - EC2 Image Builder 2: PutObject - Amazon Simple Storage Service 3: AWSImageBuilderFullAccess - AWS Managed Policy

NEW QUESTION 138
.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your SCS-C02 Exam with Our Prep Materials Via below:

https://www.certleader.com/SCS-C02-dumps.html