Q&A

Up-to-date Questions and Answers from authentic resources to improve knowledge and pass the exam at very first attempt.
----- Guaranteed.

KILL EXAMS

SCS-C01 Dumps
SCS-C01 Braindumps
SCS-C01 Real Questions
SCS-C01 Practice Test
SCS-C01 Actual Questions

PASS

Amazon

SCS-C01

AWS Certified Security - Specialty (SCS-C01)

Question #229

A company's security officer is concerned about the risk of AWS account root user logins and has assigned a security engineer to implement a notification solution for near-real-time alerts upon account root user logins.
How should the security engineer meet these requirements?

- A. Create a cron job that runs a script to download the AWS IAM security credentials file, parse the file for account root user logins, and email the security team's distribution list.
- B. Run AWS CloudTrail logs through Amazon CloudWatch Events to detect account root user logins and trigger an AWS Lambda function to send an Amazon SNS notification to the security team's distribution list.
- C. Save AWS CloudTrail logs to an Amazon S3 bucket in the security team's account. Process the CloudTrail logs with the security engineer's logging solution for account root user logins. Send an Amazon SNS notification to the security team upon encountering the account root user login events.
- D. Save VPC Flow Logs to an Amazon S3 bucket in the security team's account, and process the VPC Flow Logs with their logging solutions for account root user logins. Send an Amazon SNS notification to the security team upon encountering the account root user login events.

Answer: B
Reference:
https://aws.amazon.com/blogs/mt/monitor-and-notify-on-aws-account-root-user-activity/

Question #230

A company wants to encrypt data locally while meeting regulatory requirements related to key exhaustion. The encryption key can be no more than 10 days old or encrypt more than 2^{16} objects. Any encryption key must be generated on a FIPS-validated hardware security module (HSM). The company is cost-conscious, as it plans to upload an average of 100 objects to Amazon S3 each second for sustained operations across 5 data producers.
Which approach MOST efficiently meets the company's needs?

- A. Use the AWS Encryption SDK and set the maximum age to 10 days and the maximum number of messages encrypted to 2^{16}. Use AWS Key Management Service (AWS KMS) to generate the master key and data key. Use data key caching with the Encryption SDK during the encryption process.
- B. Use AWS Key Management Service (AWS KMS) to generate an AWS managed CMK. Then use Amazon S3 client-side encryption configured to automatically rotate with every object.
- C. Use AWS CloudHSM to generate the master key and data keys. Then use Boto 3 and Python to locally encrypt data before uploading the object. Rotate the data key every 10 days or after 2^{16} objects have been uploaded to Amazon S3.
- D. Use server-side encryption with Amazon S3 managed encryption keys (SSE-S3) and set the master key to automatically rotate.

Answer: C

Question #231

A company is setting up products to deploy in AWS Service Catalog. Management is concerned that when users launch products, elevated IAM privileges will be required to create resources.
How should the company mitigate this concern?

- A. Add a template constraint to each product in the portfolio.
- B. Add a launch constraint to each product in the portfolio.
- C. Define resource update constraints for each product in the portfolio.
- D. Update the AWS CloudFormation template backing the product to include a service role configuration.

Answer: B
Reference:
https://aws.amazon.com/blogs/mt/how-to-launch-secure-and-governed-aws-resources-with-aws-cloudformation-and-aws-service-catalog/

Question #232

A company is implementing a new application in a new AWS 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 AWS 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 than 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 VPC. Add a new network ACL rule on the database subnets. Configure the rule to TCP port 1521 from the IP address range of the application VPC. Attach the new security group to the database instances that the application instances need to access.
- B. 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.
- C. Create a new security group in the application VPC with no inbound rules. 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 VPC. Attach the application security group to the application instances that need database access, and attach the database security group to the database instances.
- D. 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 subnets. Configure the rule to allow all traffic from the IP address range of the application VPC. Attach the new security group to the application instances that need database access.

Answer: A

Question #233

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

- A. Within the Auto Scaling lifecycle, add a hook to create an attach an Amazon Elastic Block Store (Amazon EBS) log volume each time an EC2 instance is created. When the instance is terminated, the EBS volume can be reattached to another instance for log review.
- B. Create an Amazon Elastic File System (Amazon EFS) file system and add a command in the user data section of the Auto Scaling launch template to mount the EFS file system during EC2 instance creation. Configure a process on the instance to copy the logs once a day from an instance Amazon Elastic Block Store (Amazon EBS) volume to a directory in the EFS file system.
- C. Build the Amazon CloudWatch agent into the AMI used in the Auto Scaling group. Configure the CloudWatch agent to send the logs to Amazon CloudWatch Logs for review.
- D. Within the Auto Scaling lifecycle, add a lifecycle hook at the terminating state transition and alert the engineering team by using a lifecycle notification to Amazon Simple Notification Service (Amazon SNS). Configure the hook to remain in the Terminating:Wait state for 1 hour to allow manual review of the security logs prior to instance termination.

Answer: A

Question #234

A company needs to retain log data archives for several years to be compliant with regulations. The log data is no longer used, but it must be retained.
What is the MOST secure and cost-effective solution to meet these requirements?

- A. Archive the data to Amazon S3 and apply a restrictive bucket policy to deny the s3:DeleteObject API.
- B. Archive the data to Amazon S3 Glacier and apply a Vault Lock policy.
- C. Archive the data to Amazon S3 and replicated it to a second bucket in a second AWS Region. Choose the S3 Standard-Infrequent Access (S3 Standard-IA) storage class and apply a restrictive bucket policy to deny the s3:DeleteObject API.
- D. Migrate the log data to a 16 TB Amazon Elastic Block Store (Amazon EBS) volume. Create a snapshot of the EBS volume.

Answer: C

Question #235

A company uses an Amazon S3 bucket to store reports. Management has mandated that all new objects stored in this bucket must be encrypted at rest using server-side encryption with a client specified AWS Key Management Service (AWS KMS) CMK owned by the same account as the S3 bucket. The AWS account number is 111122223333, and the bucket name is reportbucket. The company˅€™s security specialist must write the S3 bucket policy to ensure the mandate can be implemented.
Which statement should the security specialist include in the policy?

A.
```
{
        "Effect": "Deny",
        "Principal": "*",
        "Action": "s3:PutObject",
        "Resource": "arn:aws:s3:::reportbucket/*",
        "Condition": {
            "StringEquals": {
                "s3:x-amz-server-side-encryption": "AES256"
            }
        }
}
```

B.
```
{
        "Effect": "Deny",
        "Principal": "*",
        "Action": "s3:PutObject",
        "Resource": "arn:aws:s3:::reportbucket/*",
        "Condition": {
            "StringNotLike": {
                "s3:x-amz-server-side-encryption-aws-kms-key-id": "arn:aws:kms:
*:111122223333:key/*"
            }
        }
}
```

C.
```
{
        "Effect": "Deny",
        "Principal": "*",
        "Action": "s3:PutObject",
        "Resource": "arn:aws:s3:::reportbucket/*",
        "Condition": {
            "StringNotLike": {
                "s3:x-amz-server-side-encryption": "aws:kms"
            }
        }
}
```

D.

```
{
        "Effect": "Deny",
        "Principal": "*",
        "Action": "s3:PutObject",
        "Resource": "arn:aws:s3:::reportbucket/*",
        "Condition": {
            "StringNotLikeIfExists": {
                "s3:x-amz-server-side-encryption-aws-kms-key-id": "arn:aws:kms:
*:111122223333:key/*"
            }
        }
    }
}
```

Answer: A

Question #236

A company website runs on Amazon EC2 instances behind an Application Load Balancer (ALB). The instances run in an Auto Scaling group across multiple
Availability Zones. There is an Amazon CloudFront distribution in front of the ALB. Users are reporting performance problems. A security engineer discovers that the website is receiving a high rate of unwanted requests to the CloudFront distribution originating from a series of source IP addresses. How should the security engineer address this problem?

- A. Using AWS Shield, configure a deny rule with an IP match condition containing the source IPs of the unwanted requests.
- B. Using Auto Scaling, configure the maximum an instance value to an increased count that will absorb the unwanted requests.
- C. Using an Amazon VPC NACL, configure an inbound deny rule for each source IP CIDR address of the unwanted requests.
- D. Using AWS WAF, configure a web ACL rate-based rule on the CloudFront distribution with a rate limit below that of the unwanted requests.

Answer: D

Question #237

A developer is building a serverless application hosted on AWS that uses Amazon Redshift as a data store. The application has separate module for read/write 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? (Choose two.)

- A. Configure cluster security groups for each application module to control access to database users that are required for read-only and read-write.
- 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 IAM 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 IAM policy for each module. Specify the ARN of an IAM user that allows the GetClusterCredentials API call.

Answer: AD

Question #238

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

- A. Review the AWS CloudTrail event history logs in an Amazon S3 bucket and look for the TerminateInstances event to identify the federated user from the role session name.
- B. Filter the AWS CloudTrail event history for the TerminateInstances event and identify the assumed IAM role. Review the AssumeRoleWithSAML event call in CloudTrail to identify the corresponding username.
- C. Search the AWS CloudTrail logs for the TerminateInstances event and note the event time. Review the IAM Access Advisor tab for all federated roles. The last accessed time should match the time when the instance was terminated.
- D. Use Amazon Athena to run a SQL query on the AWS CloudTrail logs stored in an Amazon S3 bucket and filter on the TerminateInstances

event. Identify the corresponding role and run another query to filter the AssumeRoleWithWebIdentity event for the user name.

Answer: A
Reference:
https://aws.amazon.com/blogs/security/how-to-easily-identify-your-federated-users-by-using-aws-cloudtrail/

Question #239

A company has two software development teams that are creating applications that store sensitive data in Amazon S3. Each teamג€™s data must always be separate. The companyג€™s security team must design a data encryption strategy for both teams that provides the ability to audit key usage. The solution must also minimize operational overhead.
What should the security team recommend?

- A. Tell the application teams to use two different S3 buckets with separate AWS Key Management Service (AWS KMS) AWS managed CMKs. Limit the key policies to allow encryption and decryption of the CMKs to their respective teams only. Force the teams to use encryption context to encrypt and decrypt.
- B. Tell the application teams to use two different S3 buckets with a single AWS Key Management Service (AWS KMS) AWS managed CMK. Limit the key policy to allow encryption and decryption of the CMK only. Do not allow the teams to use encryption context to encrypt and decrypt.
- C. Tell the application teams to use two different S3 buckets with separate AWS Key Management Service (AWS KMS) customer managed CMKs. Limit the key policies to allow encryption and decryption of the CMKs to their respective teams only. Force the teams to use encryption context to encrypt and decrypt.
- D. Tell the application teams to use two different S3 buckets with a single AWS Key Management Service (AWS KMS) customer managed CMK. Limit the key policy to allow encryption and decryption of the CMK only. Do not allow the teams to use encryption context to encrypt and decrypt.

Answer: B

Question #240

A security engineer is designing a solution that will provide end-to-end encryption between clients and Docker containers running in Amazon Elastic Container
Service (Amazon ECS). This solution will also handle volatile traffic patterns.
Which solution would have the MOST scalability and LOWEST latency?

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

Answer: B

Question #241

A company uses an AWS Key Management Service (AWS KMS) CMK to encrypt application data before it is stored. The companyג€™s security policy was recently modified to require encryption key rotation annually. A security engineer must ensure that annual global key rotation is enabled for the key without making changes to the application.
What should the security engineer do to accomplish this requirement?

- A. Create new AWS managed keys. Configure the key schedule for the annual rotation. Create an alias to point to the new keys.
- B. Enable automatic annual key rotation for the existing customer managed CMKs. Update the application encryption library to use a new key ID for all encryption operations. Fall back to the old key ID to decrypt data that was encrypted with previous versions of the key.
- C. Create new AWS managed CMKs. Configure the key schedule for annual rotation. Create an alias to point to the new CMKs.
- D. Enable automatic annual key rotation for the existing customer managed CMKs. Update the application encryption library to use a new key ID for all encryption operations. Create a key grant for the old CMKs and update the code to point to the ARN of the grants.

Answer: D
Reference:
https://docs.aws.amazon.com/kms/latest/developerguide/rotate-keys.html

Question #242

A company is collecting AWS CloudTrail log data from multiple AWS accounts by managing individual trails in each account and forwarding log data to a centralized Amazon S3 bucket residing in a log archive account. After CloudTrail introduced support for AWS Organizations trails, the company decided to further centralize management and automate deployment of the CloudTrail logging capability across all of its AWS accounts.
The company's security engineer created an AWS Organizations trail in the master account, enabled server-side encryption with AWS KMS managed keys (SSE-
KMS) for the log files, and specified the same bucket as the storage location. However, the engineer noticed that logs recorded by the new trail were not delivered to the bucket.
Which factors could cause this issue? (Choose two.)

- A. The CMK key policy does not allow CloudTrail to make encrypt and decrypt API calls against the key.
- B. The CMK key policy does not allow CloudTrail to make GenerateDatakey API calls against the key.
- C. The IAM role used by the CloudTrail trail does not have permissions to make PutObject API calls against a folder created for the Organizations trail.
- D. The S3 bucket policy does not allow CloudTrail to make PutObject API calls against a folder created for the Organizations trail.
- E. The CMK key policy does not allow the IAM role used by the CloudTrail trail to use the key for cryptographical operations.

Answer: AD

Question #243

A company's AWS CloudTrail logs are all centrally stored in an Amazon S3 bucket. The security team controls the company's AWS account.
The security team must prevent unauthorized access and tampering of the CloudTrail logs.
Which combination of steps should the security team take? (Choose three.)

- A. Configure server-side encryption with AWS KMS managed encryption keys (SSE-KMS)
- B. Compress log file with secure gzip.
- C. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to notify the security team of any modifications on CloudTrail log files.
- D. Implement least privilege access to the S3 bucket by configuring a bucket policy.
- E. Configure CloudTrail log file integrity validation.
- F. Configure Access Analyzer for S3.

Answer: BCE

SAMPLE QUESTIONS

These questions are for demo purpose only. **Full version** *is up to date and contains actual questions and answers.*

Killexams.com is an online platform that offers a wide range of services related to certification exam preparation. The platform provides actual questions, exam dumps, and practice tests to help individuals prepare for various certification exams with confidence. Here are some key features and services offered by Killexams.com:

Actual Exam Questions: Killexams.com provides actual exam questions that are experienced in test centers. These questions are updated regularly to ensure they are up-to-date and relevant to the latest exam syllabus. By studying these actual questions, candidates can familiarize themselves with the content and format of the real exam.

Exam Dumps: Killexams.com offers exam dumps in PDF format. These dumps contain a comprehensive collection of questions and answers that cover the exam topics. By using these dumps, candidates can enhance their knowledge and improve their chances of success in the certification exam.

Practice Tests: Killexams.com provides practice tests through their desktop VCE exam simulator and online test engine. These practice tests simulate the real exam environment and help candidates assess their readiness for the actual exam. The practice tests cover a wide range of questions and enable candidates to identify their strengths and weaknesses.

Guaranteed Success: Killexams.com offers a success guarantee with their exam dumps. They claim that by using their materials, candidates will pass their exams on the first attempt or they will refund the purchase price. This guarantee provides assurance and confidence to individuals preparing for certification exams.

Updated Content: Killexams.com regularly updates its question bank and exam dumps to ensure that they are current and reflect the latest changes in the exam syllabus. This helps candidates stay up-to-date with the exam content and increases their chances of success.

Technical Support: Killexams.com provides free 24x7 technical support to assist candidates with any queries or issues they may encounter while using their services. Their certified experts are available to provide guidance and help candidates throughout their exam preparation journey.