

Amazon-Web-Services

Exam Questions SCS-C02

AWS Certified Security - Specialty



NEW QUESTION 1

A company is using Amazon Macie, AWS Firewall Manager, Amazon Inspector, and AWS Shield Advanced in its AWS account. The company wants to receive alerts if a DDoS attack occurs against the account.

Which solution will meet this requirement?

- A. Use Macie to detect an active DDoS even
- B. Create Amazon CloudWatch alarms that respond to Macie findings.
- C. Use Amazon Inspector to review resources and to invoke Amazon CloudWatch alarms for any resources that are vulnerable to DDoS attacks.
- D. Create an Amazon CloudWatch alarm that monitors Firewall Manager metrics for an active DDoS event.
- E. Create an Amazon CloudWatch alarm that monitors Shield Advanced metrics for an active DDoS event.

Answer: D

Explanation:

This answer is correct because AWS Shield Advanced is a service that provides comprehensive protection against DDoS attacks of any size or duration. It also provides metrics and reports on the DDoS attack vectors, duration, and size. You can create an Amazon CloudWatch alarm that monitors Shield Advanced metrics such as DDoSAttackBitsPerSecond, DDoSAttackPacketsPerSecond, and DDoSAttackRequestsPerSecond to receive alerts if a DDoS attack occurs against your account. For more information, see [Monitoring AWS Shield Advanced with Amazon CloudWatch and AWS Shield Advanced metrics and alarms](#).

NEW QUESTION 2

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 IAM Key Management Service (IAM KMS) CMK owned by the same account as the S3 bucket. The IAM account number is 111122223333, and the bucket name is report bucket. 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/*"
    }
  }
}
```

- E. Option A
- F. Option B
- G. Option C
- H. Option D

Answer: D

NEW QUESTION 3

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

All AWS account activity across all member accounts must be logged and reported to the dedicated security account. The company must retain all the activity logs

in a secure storage location within the dedicated security account for 2 years. No changes or deletions of the logs are allowed. Which combination of steps will meet these requirements with the LEAST operational overhead? (Select TWO.)

- A. In the dedicated security account, create an Amazon S3 bucket
- B. Configure S3 Object Lock in compliance mode and a retention period of 2 years on the S3 bucket
- C. Set the bucket policy to allow the organization's management account to write to the S3 bucket.
- D. In the dedicated security account, create an Amazon S3 bucket
- E. Configure S3 Object Lock in compliance mode and a retention period of 2 years on the S3 bucket
- F. Set the bucket policy to allow the organization's member accounts to write to the S3 bucket.
- G. In the dedicated security account, create an Amazon S3 bucket that has an S3 Lifecycle configuration that expires objects after 2 year
- H. Set the bucket policy to allow the organization's member accounts to write to the S3 bucket.
- I. Create an AWS Cloud Trail trail for the organization
- J. Configure logs to be delivered to the logging Amazon S3 bucket in the dedicated security account.
- K. Turn on AWS CloudTrail in each account
- L. Configure logs to be delivered to an Amazon S3 bucket that is created in the organization's management account
- M. Forward the logs to the S3 bucket in the dedicated security account by using AWS Lambda and Amazon Kinesis Data Firehose.

Answer: BD

Explanation:

The correct answer is B and D. In the dedicated security account, create an Amazon S3 bucket. Configure S3 Object Lock in compliance mode and a retention period of 2 years on the S3 bucket. Set the bucket policy to allow the organization's member accounts to write to the S3 bucket. Create an AWS CloudTrail trail for the organization. Configure logs to be delivered to the logging Amazon S3 bucket in the dedicated security account.

According to the AWS documentation, AWS CloudTrail is a service that enables governance, compliance, operational auditing, and risk auditing of your AWS account. With CloudTrail, you can log, continuously monitor, and retain account activity related to actions across your AWS infrastructure. CloudTrail provides event history of your AWS account activity, including actions taken through the AWS Management Console, AWS SDKs, command line tools, and other AWS services.

To use CloudTrail 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 CloudTrail as a service principal for AWS Organizations, which lets you create an organization trail that applies to all accounts in your organization. An organization trail logs events for all AWS Regions and delivers the log files to an S3 bucket that you specify.

To create an organization trail, you need to use an administrator account, such as the organization's management account or a delegated administrator account. You can then configure the trail to deliver logs to an S3 bucket in the dedicated security account. This will ensure that all account activity across all member accounts and regions is logged and reported to the security account.

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 CloudTrail logs, you need to create an S3 bucket in the dedicated security account that will store the logs from the organization trail. You can then configure S3 Object Lock on the bucket to prevent objects from being deleted or overwritten for a fixed amount of time or indefinitely. You can also enable compliance mode on the bucket, which prevents any user, including the root user in your account, from deleting or modifying a locked object until it reaches its retention date.

To set a retention period of 2 years on the S3 bucket, you need to create a default retention configuration for the bucket that specifies a retention mode (either governance or compliance) and a retention period (either a number of days or a date). You can then set the bucket policy to allow the organization's member accounts to write to the S3 bucket. This will ensure that all logs are retained in a secure storage location within the security account for 2 years and no changes or deletions are allowed.

Option A is incorrect because setting the bucket policy to allow the organization's management account to write to the S3 bucket is not sufficient, as it will not grant access to the other member accounts in the organization.

Option C is incorrect because using an S3 Lifecycle configuration that expires objects after 2 years is not secure, as it will allow users to delete or modify objects before they expire.

Option E is incorrect because using Lambda and Kinesis Data Firehose to forward logs from one S3 bucket to another is not necessary, as CloudTrail can directly deliver logs to an S3 bucket in another account. It also introduces additional operational overhead and complexity.

NEW QUESTION 4

A company has several workloads running on AWS. Employees are required to authenticate using on-premises ADFS and SSO to access the AWS Management Console. Developers migrated an existing legacy web application to an Amazon EC2 instance. Employees need to access this application from anywhere on the internet, but currently, there is no authentication system built into the application.

How should the Security Engineer implement employee-only access to this system without changing the application?

- A. Place the application behind an Application Load Balancer (ALB). Use Amazon Cognito as authentication for the ALB.
- B. Define a SAML-based Amazon Cognito user pool and connect it to ADFS.
- C. Implement AWS SSO in the master account and link it to ADFS as an identity provider
- D. Define the EC2 instance as a managed resource, then apply an IAM policy on the resource.
- E. Define an Amazon Cognito identity pool, then install the connector on the Active Directory server
- F. Use the Amazon Cognito SDK on the application instance to authenticate the employees using their Active Directory user names and passwords.
- G. Create an AWS Lambda custom authorizer as the authenticator for a reverse proxy on Amazon EC2. Ensure the security group on Amazon EC2 only allows access from the Lambda function.

Answer: A

Explanation:

<https://docs.aws.amazon.com/elasticloadbalancing/latest/application/listener-authenticate-users.html>

NEW QUESTION 5

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 CloudWatch Events) event with an EC2 instance as the source and create volume as the event trigger
- 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 6

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

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

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

- A. Create a local individual break glass IAM user for the security tea
- B. Create a trail in AWS CloudTrail that has Amazon CloudWatch Logs turned o
- C. Use Amazon EventBridge to monitor local user activities.
- D. Create a break glass EC2 key pair for the AWS accoun
- E. Provide the key pair to the security tea
- F. Use AWS CloudTrail to monitor key pair activit
- G. Send notifications to the security team by using Amazon Simple Notification Service (Amazon SNS).
- H. Create a break glass IAM role for the accoun
- I. Allow security team members to perform the AssumeRoleWithSAML operatio
- J. Create an AWS Cloud Trail trail that has Amazon CloudWatch Logs turned o
- K. Use Amazon EventBridge to monitor security team activities.
- L. Create a local individual break glass IAM user on the operating system level of each workload instance.Configure unrestricted security groups on the instances to grant access to the break glass IAM users.
- M. Configure AWS Systems Manager Session Manager for Amazon EC2. Configure an AWS Cloud Trail filter based on Session Manage
- N. Send the results to an Amazon Simple Notification Service (Amazon SNS) topic.

Answer: AE

Explanation:

The combination of solutions that will meet the requirements are:

- A. Create a local individual break glass IAM user for the security team. Create a trail in AWS CloudTrail that has Amazon CloudWatch Logs turned on. Use Amazon EventBridge to monitor local user activities. This is a valid solution because it allows the security team to access the workload AWS account and instances using a local IAM user that does not depend on SAML federation. It also enables logging and monitoring of the break glass user activities using AWS CloudTrail, Amazon CloudWatch Logs, and Amazon EventBridge123.
- E. Configure AWS Systems Manager Session Manager for Amazon EC2. Configure an AWS CloudTrail filter based on Session Manager. Send the results to an Amazon Simple Notification Service (Amazon SNS) topic. This is a valid solution because it allows the security team to access the workload instances without opening any inbound ports or managing SSH keys or bastion hosts. It also enables logging and notification of the break glass user activities using AWS CloudTrail, Session Manager, and Amazon SNS456.

The other options are incorrect because:

- B. Creating a break glass EC2 key pair for the AWS account and providing it to the security team is not a valid solution, because it requires opening inbound ports on the instances and managing SSH keys, which increases the security risk and complexity7.
- C. Creating a break glass IAM role for the account and allowing security team members to perform the AssumeRoleWithSAML operation is not a valid solution, because it still depends on SAML federation, which might not work in case of SAML errors8.
- D. Creating a local individual break glass IAM user on the operating system level of each workload instance and configuring unrestricted security groups on the instances to grant access to the break glass IAM users is not a valid solution, because it requires opening inbound ports on the instances and managing multiple local users, which increases the security risk and complexity9.

References:

1: Creating an IAM User in Your AWS Account 2: Creating a Trail - AWS CloudTrail 3: Using Amazon EventBridge with AWS CloudTrail 4: Setting up Session Manager - AWS Systems Manager 5: Logging Session Manager sessions - AWS Systems Manager 6: Amazon Simple Notification Service 7: Connecting to your Linux instance using SSH - Amazon Elastic Compute Cloud 8: AssumeRoleWithSAML - AWS Security Token Service 9: IAM Users - AWS Identity and Access Management

NEW QUESTION 7

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

What should the security engineer recommend?

- A. Enable Amazon RDS encryption to encrypt the database and 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 AWS 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 Amazon Elastic Block Store (Amazon EBS) volum
- H. Store the database credentials in AWS 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 AWS Secrets Manager with automatic rotation
- M. Set up TLS for the connection to the RDS hosted database.
- N. Set up an AWS CloudHSM cluster with AWS Key Management Service (AWS KMS) to store KMS key
- O. Set up Amazon RDS encryption using AWS KMS to encrypt the database
- P. Store the database credentials in AWS Systems Manager Parameter Store with automatic rotation
- Q. Set up TLS for the connection to the RDS hosted database.

Answer: C

NEW QUESTION 8

A security engineer is using AWS Organizations and wants to optimize SCPs. The security engineer needs to ensure that the SCPs conform to best practices. Which approach should the security engineer take to meet this requirement?

- A. Use AWS IAM Access Analyzer to analyze the policies
- B. View the findings from policy validation checks.
- C. Review AWS Trusted Advisor checks for all accounts in the organization.
- D. Set up AWS Audit Manager
- E. Run an assessment for all AWS Regions for all accounts.
- F. Ensure that Amazon Inspector agents are installed on all Amazon EC2 instances in all accounts.

Answer: A

NEW QUESTION 9

Your CTO is very worried about the security of your IAM account. How best can you prevent hackers from completely hijacking your account? Please select:


- A. Use short but complex password on the root account and any administrators.
- B. Use IAM Geo-Lock and disallow anyone from logging in except for in your city.
- C. Use MFA on all users and accounts, especially on the root account.
- D. Don't write down or remember the root account password after creating the IAM account.

Answer: C

Explanation:

Multi-factor authentication can add one more layer of security to your IAM account. Even when you go to your Security Credentials dashboard, one of the items is to enable MFA on your root account.

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



Option A is invalid because you need to have a good password policy. Option B is invalid because there is no IAM Geo-Lock. Option D is invalid because this is not a recommended practice. For more information on MFA, please visit the below URL:

http://docs.IAM.amazonaws.com/IAM/latest/UserGuide/id_credentials_mfa.html

The correct answer is: Use MFA on all users and accounts, especially on the root account. Submit your Feedback/Queries to our Experts

NEW QUESTION 10

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

Which solution will meet these requirements?

- A. Create new S3 buckets with S3 Object Lock enabled in compliance mode
- B. Place objects in the S3 buckets.
- C. Use S3 Glacier Vault Lock to attach a Vault Lock policy to new S3 bucket
- D. Wait 24 hours to complete the Vault Lock process
- E. Place objects in the S3 buckets.
- F. Create new S3 buckets with S3 Object Lock enabled in governance mode
- G. Place objects in the S3 buckets.
- H. Create new S3 buckets with S3 Object Lock enabled in governance mode
- I. Add a legal hold to the S3 bucket
- J. Place objects in the S3 buckets.

Answer: A

NEW QUESTION 10

A security engineer is troubleshooting an AWS Lambda function that is named MyLambdaFunction. The function is encountering an error when the function attempts to read the objects in an Amazon S3 bucket that is named DOC-EXAMPLE-BUCKET. The S3 bucket has the following bucket policy:

```
{
  "Effect": "Allow",
  "Principal": {
    "Service": "lambda.amazonaws.com"
  },
  "Action": "s3:GetObject",
  "Resource": "arn:aws:s3:::DOC-EXAMPLE-BUCKET",
  "Condition": {
    "ArnLike": {
      "aws:SourceArn": "arn:aws:lambda:::function:MyLambdaFunction"
    }
  }
}
```

Which change should the security engineer make to the policy to ensure that the Lambda function can read the bucket objects?

- A. Remove the Condition element
- B. Change the Principal element to the following: {"AWS": "arn:aws:::lambda:::function:MyLambdaFunction"}
- C. Change the Action element to the following: "s3:GetObject*" "s3:GetBucket*"
- D. Change the Resource element to "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*".
- E. Change the Resource element to "arn:aws:lambda:::function:MyLambdaFunction". Change the Principal element to the following: {"Service": "s3.amazonaws.com"}

Answer: C

Explanation:

The correct answer is C. Change the Resource element to "arn:aws:s3:::DOC-EXAMPLE-BUCKET/*".

The reason is that the Resource element in the bucket policy specifies which objects in the bucket are affected by the policy. In this case, the policy only applies to the bucket itself, not the objects inside it. Therefore, the Lambda function cannot access the objects with the s3:GetObject permission. To fix this, the Resource element should include a wildcard (*) to match all objects in the bucket. This way, the policy grants the Lambda function permission to read any object in the bucket.

The other options are incorrect for the following reasons:

- > A. Removing the Condition element would not help, because it only restricts access based on the source IP address of the request. The Principal element should not be changed to the Lambda function ARN, because it specifies who is allowed or denied access by the policy. The policy should allow access to any principal ("*") and rely on IAM roles or policies to control access to the Lambda function.
- > B. Changing the Action element to include s3:GetBucket* would not help, because it would grant additional permissions that are not needed by the Lambda function, such as s3:GetBucketAcl or s3:GetBucketPolicy. The s3:GetObject* permission is sufficient for reading objects in the bucket.
- > D. Changing the Resource element to the Lambda function ARN would not make sense, because it would mean that the policy applies to the Lambda function itself, not the bucket or its objects. The Principal element should not be changed to s3.amazonaws.com, because it would grant access to any AWS service that uses S3, not just Lambda.

NEW QUESTION 13

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 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: 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 16

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 20

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 24

A Security Engineer is troubleshooting an issue with a company's custom logging application. The application logs are written to an Amazon S3 bucket with event notifications enabled to send events to an Amazon SNS topic. All logs are encrypted at rest using an IAM KMS CMK. The SNS topic is subscribed to an encrypted Amazon SQS queue. The logging application polls the queue for new messages that contain metadata about the S3 object. The application then reads the content of the object from the S3 bucket for indexing.

The Logging team reported that Amazon CloudWatch metrics for the number of messages sent or received is showing zero. No togs are being received.

What should the Security Engineer do to troubleshoot this issue?

- A) Add the following statement to the IAM managed CMKs:

```
{
  "Sid": "Allow Amazon SNS to use this key",
  "Effect": "Allow",
  "Principal": {
    "Service": ["sns.amazonaws.com", "sqs.amazonaws.com", "s3.amazonaws.com"]
  },
  "Action": [
    "kms:Decrypt",
    "kms:GenerateDataKey*"
  ],
  "Resource": "*"
}
```

B)

Add the following statement to the CMK key policy:

```
{
  "Sid": "Allow Amazon SNS to use this key",
  "Effect": "Allow",
  "Principal": {
    "Service": "sns.amazonaws.com"
  },
  "Action": [
    "kms:Decrypt",
    "kms:GenerateDataKey*"
  ],
  "Resource": "*"
}
```

C)

Add the following statement to the CMK key policy:

```
{
  "Sid": "Allow Amazon SNS to use this key",
  "Effect": "Allow",
  "Principal": {
    "Service": "sqs.amazonaws.com"
  },
  "Action": [
    "kms:Decrypt",
    "kms:GenerateDataKey*"
  ],
  "Resource": "*"
}
```

D)

Add the following statement to the CMK key policy:

```
{
  "Sid": "Allow Amazon SNS to use this key",
  "Effect": "Allow",
  "Principal": {
    "Service": "s3.amazonaws.com"
  },
  "Action": [
    "kms:Decrypt",
    "kms:GenerateDataKey*"
  ],
  "Resource": "*"
}
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: D

NEW QUESTION 28

A company needs a security engineer to implement a scalable solution for multi-account authentication and authorization. The solution should not introduce additional user-managed architectural components. Native IAM features should be used as much as possible. The security engineer has set up IAM Organizations with all features activated and IAM SSO enabled.

Which additional steps should the security engineer take to complete the task?

- A. Use AD Connector to create users and groups for all employees that require access to IAM accounts. Assign AD Connector groups to IAM accounts and link to the IAM roles in accordance with the employees' job functions and access requirements. Instruct employees to access IAM accounts by using the IAM Directory Service user portal.
- B. Use an IAM SSO default directory to create users and groups for all employees that require access to IAM account
- C. Assign groups to IAM accounts and link to permission sets in accordance with the employees' job functions and access requirement
- D. Instruct employees to access IAM accounts by using the IAM SSO user portal.
- E. Use an IAM SSO default directory to create users and groups for all employees that require access to IAM account

- F. Link IAM SSO groups to the IAM users present in all accounts to inherit existing permission
- G. Instruct employees to access IAM accounts by using the IAM SSO user portal.
- H. Use IAM Directory Service for Microsoft Active Directory to create users and groups for all employees that require access to IAM accounts Enable IAM Management Console access in the created directory and specify IAM SSO as a source of information for integrated accounts and permission set
- I. Instruct employees to access IAM accounts by using the IAM Directory Service user portal.

Answer: B

NEW QUESTION 33

A security engineer receives a notice from the AWS Abuse team about suspicious activity from a Linux-based Amazon EC2 instance that uses Amazon Elastic Block Store (Amazon EBS)-based storage. The instance is making connections to known malicious addresses. The instance is in a development account within a VPC that is in the us-east-1 Region. The VPC contains an internet gateway and has a subnet in us-east-1a and us-east-1b. Each subnet is associated with a route table that uses the internet gateway as a default route. Each subnet also uses the default network ACL. The suspicious EC2 instance runs within the us-east-1b subnet. During an initial investigation, a security engineer discovers that the suspicious instance is the only instance that runs in the subnet. Which response will immediately mitigate the attack and help investigate the root cause?

- A. Log in to the suspicious instance and use the netstat command to identify remote connections. Use the IP addresses from these remote connections to create deny rules in the security group of the instance. Install diagnostic tools on the instance for investigation. Update the outbound network ACL for the subnet in us-east-1b to explicitly deny all connections as the first rule during the investigation of the instance.
- B. Update the outbound network ACL for the subnet in us-east-1b to explicitly deny all connections as the first rule. Replace the security group with a new security group that allows connections only from a diagnostics security group. Update the outbound network ACL for the us-east-1b subnet to remove the deny all rule. Launch a new EC2 instance that has diagnostic tools. Assign the new security group to the new EC2 instance. Use the new EC2 instance to investigate the suspicious instance.
- C. Ensure that the Amazon Elastic Block Store (Amazon EBS) volumes that are attached to the suspicious EC2 instance will not delete upon termination. Terminate the instance. Launch a new EC2 instance in us-east-1a that has diagnostic tools. Mount the EBS volumes from the terminated instance for investigation.
- D. Create an AWS WAF web ACL that denies traffic to and from the suspicious instance. Attach the AWS WAF web ACL to the instance to mitigate the attack. Log in to the instance and install diagnostic tools to investigate the instance.

Answer: B

Explanation:

This option suggests updating the outbound network ACL for the subnet in us-east-1b to explicitly deny all connections as the first rule, replacing the security group with a new one that only allows connections from a diagnostics security group, and launching a new EC2 instance with diagnostic tools to investigate the suspicious instance. This option will immediately mitigate the attack and provide the necessary tools for investigation.

NEW QUESTION 35

A company is undergoing a layer 3 and layer 4 DDoS attack on its web servers running on IAM. Which combination of IAM services and features will provide protection in this scenario? (Select THREE).

- A. Amazon Route 53
- B. IAM Certificate Manager (ACM)
- C. Amazon S3
- D. IAM Shield
- E. Elastic Load Balancer
- F. Amazon GuardDuty

Answer: DEF

NEW QUESTION 39

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 field
- 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 instance
- G. Enable third-party disk encryption to encrypt the Amazon Elastic Block Store (Amazon EBS) volume
- H. Store the database credentials in IAM CloudHSM with automatic rotation
- 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 rotation
- 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 database
- O. Store database credentials in the IAM Systems Manager Parameter Store with automatic rotation
- 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 44

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

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

- A. Create an identity policy that allows the sts: AssumeRole action in the AWS account that contains the resource
- B. Attach the identity policy to the IAM user.
- C. Ensure that the sts: AssumeRole action is allowed by the SCPs of the organization that owns the resources that the IAM user needs to access.
- D. Create a role in the AWS account that contains the resource
- E. Create an entry in the role's trust policy that allows the IAM user to assume the rol
- F. Attach the trust policy to the role.
- G. Establish a trust relationship between the IAM user and the AWS account that contains the resources.
- H. Create a role in the IAM user's AWS account
- I. Create an identity policy that allows the sts: AssumeRole actio
- J. Attach the identity policy to the role.

Answer: BC

Explanation:

To allow cross-account access to resources using IAM roles, the following steps are required:

- Create a role in the AWS account that contains the resources (the trusting account) and specify the AWS account that contains the IAM user (the trusted account) as a trusted entity in the role's trust policy. This allows users from the trusted account to assume the role and access resources in the trusting account.
- Ensure that the IAM user has permission to assume the role in their own AWS account. This can be done by creating an identity policy that allows the sts:AssumeRole action and attaching it to the IAM user or their group.
- Ensure that there are no service control policies (SCPs) in the organization that owns the resources that deny or restrict access to the sts:AssumeRole action or the role itself. SCPs are applied to all accounts in an organization and can override any permissions granted by IAM policies.

Verified References:

- <https://repost.aws/knowledge-center/cross-account-access-iam>
- https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_accounts_access.html
- https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_cross-account-with-roles.html

NEW QUESTION 49

A company accidentally deleted the private key for an Amazon Elastic Block Store (Amazon EBS)-backed Amazon EC2 instance. A security engineer needs to regain access to the instance.

Which combination of steps will meet this requirement? (Choose two.)

- A. Stop the instanc
- B. Detach the root volum
- C. Generate a new key pair.
- D. Keep the instance runnin
- E. Detach the root volum
- F. Generate a new key pair.
- G. When the volume is detached from the original instance, attach the volume to another instance as a data volum
- H. Modify the authorized_keys file with a new public ke
- I. Move the volume back to the original instanc
- J. Start the instance.
- K. When the volume is detached from the original instance, attach the volume to another instance as a data volum
- L. Modify the authorized_keys file with a new private ke
- M. Move the volume back to the original instanc
- N. Start the instance.
- O. When the volume is detached from the original instance, attach the volume to another instance as a data volum
- P. Modify the authorized_keys file with a new public ke
- Q. Move the volume back to the original instance that is running.

Answer: AC

Explanation:

If you lose the private key for an EBS-backed instance, you can regain access to your instance. You must stop the instance, detach its root volume and attach it to another instance as a data volume, modify the authorized_keys file with a new public key, move the volume back to the original instance, and restart the instance.

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/TroubleshootingInstancesConnecting.html#replacing>

NEW QUESTION 53

A company needs to follow security best practices to deploy resources from an AWS CloudFormation template. The CloudFormation template must be able to configure sensitive database credentials.

The company already uses AWS Key Management Service (AWS KMS) and AWS Secrets Manager. Which solution will meet the requirements?

- A. Use a dynamic reference in the CloudFormation template to reference the database credentials in Secrets Manager.
- B. Use a parameter in the CloudFormation template to reference the database credential
- C. Encrypt the CloudFormation template by using AWS KMS.
- D. Use a SecureString parameter in the CloudFormation template to reference the database credentials in Secrets Manager.

E. Use a SecureString parameter in the CloudFormation template to reference an encrypted value in AWS KMS

Answer: A

Explanation:

➤ Option A: This option meets the requirements of following security best practices and configuring sensitive database credentials in the CloudFormation template. A dynamic reference is a way to specify external values that are stored and managed in other services, such as Secrets Manager, in the stack templates¹. When using a dynamic reference, CloudFormation retrieves the value of the specified reference when necessary during stack and change set operations¹. Dynamic references can be used for certain resources that support them, such as AWS::RDS::DBInstance¹. By using a dynamic reference to reference the database credentials in Secrets Manager, the company can leverage the existing integration between these services and avoid hardcoding the secret information in the template. Secrets Manager is a service that helps you protect secrets needed to access your applications, services, and IT resources². Secrets Manager enables you to rotate, manage, and retrieve database credentials, API keys, and other secrets throughout their lifecycle².

NEW QUESTION 54

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 57

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

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

- A. Configure access logging for the required API stage.
- B. Configure an AWS CloudTrail trail destination for API Gateway event
- C. Configure filters on the userIdentity, userAgent, and sourceIPAddress fields.
- D. Configure an Amazon S3 destination for API Gateway log
- E. Run Amazon Athena queries to analyze API access information.
- F. Use Amazon CloudWatch Logs Insights to analyze API access information.
- G. Select the Enable Detailed CloudWatch Metrics option on the required API stage.

Answer: CD

NEW QUESTION 58

A company has an AWS account that includes an Amazon S3 bucket. The S3 bucket uses server-side encryption with AWS KMS keys (SSE-KMS) to encrypt all the objects at rest by using a customer managed key. The S3 bucket does not have a bucket policy.

An IAM role in the same account has an IAM policy that allows s3 List* and s3 Get* permissions for the S3 bucket. When the IAM role attempts to access an object in the S3 bucket the role receives an access denied message.

Why does the IAM role not have access to the objects that are in the S3 bucket?

- A. The IAM role does not have permission to use the KMS CreateKey operation.
- B. The S3 bucket lacks a policy that allows access to the customer managed key that encrypts the objects.
- C. The IAM role does not have permission to use the customer managed key that encrypts the objects that are in the S3 bucket.
- D. The ACL of the S3 objects does not allow read access for the objects when the objects are encrypted at rest.

Answer: C

Explanation:

When using server-side encryption with AWS KMS keys (SSE-KMS), the requester must have both Amazon S3 permissions and AWS KMS permissions to access the objects. The Amazon S3 permissions are for the bucket and object operations, such as s3:ListBucket and s3:GetObject. The AWS KMS permissions are for the key operations, such as kms:GenerateDataKey and kms:Decrypt. In this case, the IAM role has the necessary Amazon S3 permissions, but not the AWS KMS permissions to use the customer managed key that encrypts the objects. Therefore, the IAM role receives an access denied message when trying to access the objects. Verified References:

- <https://docs.aws.amazon.com/AmazonS3/latest/userguide/troubleshoot-403-errors.html>
- <https://repost.aws/knowledge-center/s3-access-denied-error-kms>
- <https://repost.aws/knowledge-center/cross-account-access-denied-error-s3>

NEW QUESTION 62

An IAM user receives an Access Denied message when the user attempts to access objects in an Amazon S3 bucket. The user and the S3 bucket are in the same AWS account. The S3 bucket is configured to use server-side encryption with AWS KMS keys (SSE-KMS) to encrypt all of its objects at rest by using a customer managed key from the same AWS account. The S3 bucket has no bucket policy defined. The IAM user has been granted permissions through an IAM policy that

allows the kms:Decrypt permission to the customer managed key. The IAM policy also allows the s3:List* and s3:Get* permissions for the S3 bucket and its objects.

Which of the following is a possible reason that the IAM user cannot access the objects in the S3 bucket?

- A. The IAM policy needs to allow the kms:DescribeKey permission.
- B. The S3 bucket has been changed to use the AWS managed key to encrypt objects at rest.
- C. An S3 bucket policy needs to be added to allow the IAM user to access the objects.
- D. The KMS key policy has been edited to remove the ability for the AWS account to have full access to the key.

Answer: D

Explanation:

The possible reason that the IAM user cannot access the objects in the S3 bucket is D. The KMS key policy has been edited to remove the ability for the AWS account to have full access to the key.

This answer is correct because the KMS key policy is the primary way to control access to the KMS key, and it must explicitly allow the AWS account to have full access to the key. If the KMS key policy has been edited to remove this permission, then the IAM policy that grants kms:Decrypt permission to the IAM user has no effect, and the IAM user cannot decrypt the objects in the S3 bucket¹².

The other options are incorrect because:

- > A. The IAM policy does not need to allow the kms:DescribeKey permission, because this permission is not required for decrypting objects in S3 using SSE-KMS. The kms:DescribeKey permission allows getting information about a KMS key, such as its creation date, description, and key state³.
- > B. The S3 bucket has not been changed to use the AWS managed key to encrypt objects at rest, because this would not cause an Access Denied message for the IAM user. The AWS managed key is a default KMS key that is created and managed by AWS for each AWS account and Region. The IAM user does not need any permissions on this key to use it for SSE-KMS⁴.
- > C. An S3 bucket policy does not need to be added to allow the IAM user to access the objects, because the IAM user already has s3:List* and s3:Get* permissions for the S3 bucket and its objects through an IAM policy. An S3 bucket policy is an optional way to grant cross-account access or public access to an S3 bucket⁵.

References:

- 1: Key policies in AWS KMS 2: Using server-side encryption with AWS KMS keys (SSE-KMS) 3: AWS KMS API Permissions Reference 4: Using server-side encryption with Amazon S3 managed keys (SSE-S3) 5: Bucket policy examples

NEW QUESTION 65

A developer is building a serverless application hosted on IAM that uses Amazon Redshift in a data store. The application has separate modules 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? (Select 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: CD

Explanation:

To grant appropriate access to the application modules, the security engineer should do the following:

- > Configure an IAM policy for each module. Specify the ARN of an Amazon Redshift database user that allows the GetClusterCredentials API call. This allows the application modules to use temporary credentials to access the database with the permissions of the specified user.
- > Create local database users for each module. This allows the security engineer to create separate users for read/write and read-only functionality, and to assign them different privileges on the database tables.

NEW QUESTION 69

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

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

- A. Update the S3 bucket policy to restrict access to a CloudFront origin access identity (OAI).
- B. Update the website DNS record to use an Amazon Route 53 geolocation record deny list of countries where the company lacks a license.
- C. Add a CloudFront geo restriction deny list of countries where the company lacks a license.
- D. Update the S3 bucket policy with a deny list of countries where the company lacks a license.
- E. Enable the Restrict Viewer Access option in CloudFront to create a deny list of countries where the company lacks a license.

Answer: AC

Explanation:

For Enable Geo-Restriction, choose Yes. For Restriction Type, choose Whitelist to allow access to certain countries, or choose Blacklist to block access from certain countries. <https://IAM.amazon.com/premiumsupport/knowledge-center/cloudfront-geo-restriction/>

NEW QUESTION 70

A company is evaluating the use of AWS Systems Manager Session Manager to gain access to the company's Amazon EC2 instances. However, until the company implements the change, the company must protect the key file for the EC2 instances from read and write operations by any other users.

When a security administrator tries to connect to a critical EC2 Linux instance during an emergency, the security administrator receives the following error. "Error Unprotected private key file - Permissions for 'ssh/my_private_key.pem' are too open".

Which command should the security administrator use to modify the private key file permissions to resolve this error?

- A. `chmod 0040 ssh/my_private_key.pem`
- B. `chmod 0400 ssh/my_private_key.pem`
- C. `chmod 0004 ssh/my_private_key.pem`

D. chmod 0777 ssh/my_private_key perm

Answer: B

Explanation:

The error message indicates that the private key file permissions are too open, meaning that other users can read or write to the file. This is a security risk, as the private key should be accessible only by the owner of the file. To fix this error, the security administrator should use the chmod command to change the permissions of the private key file to 0400, which means that only the owner can read the file and no one else can read or write to it.

The chmod command takes a numeric argument that represents the permissions for the owner, group, and others in octal notation. Each digit corresponds to a set of permissions: read (4), write (2), and execute (1). The digits are added together to get the final permissions for each category. For example, 0400 means that the owner has read permission (4) and no other permissions (0), and the group and others have no permissions at all (0).

The other options are incorrect because they either do not change the permissions at all (D), or they give too much or too little permissions to the owner, group, or others (A, C).

Verified References:

- > <https://superuser.com/questions/215504/permissions-on-private-key-in-ssh-folder>
- > <https://www.baeldung.com/linux/ssh-key-permissions>

NEW QUESTION 74

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 77

An organization wants to log all IAM API calls made within all of its IAM accounts, and must have a central place to analyze these logs. What steps should be taken to meet these requirements in the MOST secure manner? (Select TWO)

- A. Turn on IAM CloudTrail in each IAM account
- B. Turn on CloudTrail in only the account that will be storing the logs
- C. Update the bucket ACL of the bucket in the account that will be storing the logs so that other accounts can log to it
- D. Create a service-based role for CloudTrail and associate it with CloudTrail in each account
- E. Update the bucket policy of the bucket in the account that will be storing the logs so that other accounts can log to it

Answer: AE

Explanation:

these are the steps that can meet the requirements in the most secure manner. CloudTrail is a service that records AWS API calls and delivers log files to an S3 bucket. Turning on CloudTrail in each IAM account can help capture all IAM API calls made within those accounts. Updating the bucket policy of the bucket in the account that will be storing the logs can help grant other accounts permission to write log files to that bucket. The other options are either unnecessary or insecure for logging and analyzing IAM API calls.

NEW QUESTION 78

A company has several petabytes of data. The company must preserve this data for 7 years to comply with regulatory requirements. The company's compliance team asks a security officer to develop a strategy that will prevent anyone from changing or deleting the data.

Which solution will meet this requirement MOST cost-effectively?

- A. Create an Amazon S3 bucket
- B. Configure the bucket to use S3 Object Lock in compliance mod
- C. Upload the data to the bucket
- D. Create a resource-based bucket policy that meets all the regulatory requirements.
- E. Create an Amazon S3 bucket
- F. Configure the bucket to use S3 Object Lock in governance mod

- G. Upload the data to the bucket
- H. Create a user-based IAM policy that meets all the regulatory requirements.
- I. Create a vault in Amazon S3 Glacier
- J. Create a Vault Lock policy in S3 Glacier that meets all the regulatory requirement
- K. Upload the data to the vault.
- L. Create an Amazon S3 bucket
- M. Upload the data to the bucket
- N. Use a lifecycle rule to transition the data to a vault in S3 Glacier
- O. Create a Vault Lock policy that meets all the regulatory requirements.

Answer: C

Explanation:

To preserve the data for 7 years and prevent anyone from changing or deleting it, the security officer needs to use a service that can store the data securely and enforce compliance controls. The most cost-effective way to do this is to use Amazon S3 Glacier, which is a low-cost storage service for data archiving and long-term backup. S3 Glacier allows you to create a vault, which is a container for storing archives. Archives are any data such as photos, videos, or documents that you want to store durably and reliably.

S3 Glacier also offers a feature called Vault Lock, which helps you to easily deploy and enforce compliance controls for individual vaults with a Vault Lock policy. You can specify controls such as “write once read many” (WORM) in a Vault Lock policy and lock the policy from future edits. Once a Vault Lock policy is locked, the policy can no longer be changed or deleted. S3 Glacier enforces the controls set in the Vault Lock policy to help achieve your compliance objectives. For example, you can use Vault Lock policies to enforce data retention by denying deletes for a specified period of time.

To use S3 Glacier and Vault Lock, the security officer needs to follow these steps:

- Create a vault in S3 Glacier using the AWS Management Console, AWS Command Line Interface (AWS CLI), or AWS SDKs.
- Create a Vault Lock policy in S3 Glacier that meets all the regulatory requirements using the IAM policy language. The policy can include conditions such as `aws:CurrentTime` or `aws:SecureTransport` to further restrict access to the vault.
- Initiate the lock by attaching the Vault Lock policy to the vault, which sets the lock to an in-progress state and returns a lock ID. While the policy is in the in-progress state, you have 24 hours to validate your Vault Lock policy before the lock ID expires. To prevent your vault from exiting the in-progress state, you must complete the Vault Lock process within these 24 hours. Otherwise, your Vault Lock policy will be deleted.
- Use the lock ID to complete the lock process. If the Vault Lock policy doesn't work as expected, you can stop the Vault Lock process and restart from the beginning.
- Upload the data to the vault using either direct upload or multipart upload methods. For more information about S3 Glacier and Vault Lock, see S3 Glacier Vault Lock.

The other options are incorrect because:

- Option A is incorrect because creating an Amazon S3 bucket and configuring it to use S3 Object Lock in compliance mode will not prevent anyone from changing or deleting the data. S3 Object Lock is a feature that allows you to store objects using a WORM model in S3. You can apply two types of object locks: retention periods and legal holds. A retention period specifies a fixed period of time during which an object remains locked. A legal hold is an indefinite lock on an object until it is removed. However, S3 Object Lock only prevents objects from being overwritten or deleted by any user, including the root user in your AWS account. It does not prevent objects from being modified by other means, such as changing their metadata or encryption settings. Moreover, S3 Object Lock requires that you enable versioning on your bucket, which will incur additional storage costs for storing multiple versions of an object.
- Option B is incorrect because creating an Amazon S3 bucket and configuring it to use S3 Object Lock in governance mode will not prevent anyone from changing or deleting the data. S3 Object Lock in governance mode works similarly to compliance mode, except that users with specific IAM permissions can change or delete objects that are locked. This means that users who have `s3:BypassGovernanceRetention` permission can remove retention periods or legal holds from objects and overwrite or delete them before they expire. This option does not provide strong enforcement for compliance controls as required by the regulatory requirements.
- Option D is incorrect because creating an Amazon S3 bucket and using a lifecycle rule to transition the data to a vault in S3 Glacier will not prevent anyone from changing or deleting the data. Lifecycle rules are actions that Amazon S3 automatically performs on objects during their lifetime. You can use lifecycle rules to transition objects between storage classes or expire them after a certain period of time. However, lifecycle rules do not apply any compliance controls on objects or prevent them from being modified or deleted by users. Moreover, transitioning objects from S3 to S3 Glacier using lifecycle rules will incur additional charges for retrieval requests and data transfers.

NEW QUESTION 79

A security engineer wants to evaluate configuration changes to a specific AWS resource to ensure that the resource meets compliance standards. However, the security engineer is concerned about a situation in which several configuration changes are made to the resource in quick succession. The security engineer wants to record only the latest configuration of that resource to indicate the cumulative impact of the set of changes.

Which solution will meet this requirement in the MOST operationally efficient way?

- A. Use AWS CloudTrail to detect the configuration changes by filtering API calls to monitor the changes. Use the most recent API call to indicate the cumulative impact of multiple calls
- B. Use AWS Config to detect the configuration changes and to record the latest configuration in case of multiple configuration changes.
- C. Use Amazon CloudWatch to detect the configuration changes by filtering API calls to monitor the change
- D. Use the most recent API call to indicate the cumulative impact of multiple calls.
- E. Use AWS Cloud Map to detect the configuration change
- F. Generate a report of configuration changes from AWS Cloud Map to track the latest state by using a sliding time window.

Answer: B

Explanation:

AWS Config is a service that enables you to assess, audit, and evaluate the configurations of your AWS resources. AWS Config continuously monitors and records your AWS resource configurations and allows you to automate the evaluation of recorded configurations against desired configurations.

To evaluate configuration changes to a specific AWS resource and ensure that it meets compliance standards, the security engineer should use AWS Config to detect the configuration changes and to record the latest configuration in case of multiple configuration changes. This will allow the security engineer to view the current state of the resource and its compliance status, as well as its configuration history and timeline.

AWS Config records configuration changes as `ConfigurationItems`, which are point-in-time snapshots of the resource's attributes, relationships, and metadata. If multiple configuration changes occur within a short period of time, AWS Config records only the latest `ConfigurationItem` for that resource. This indicates the cumulative impact of the set of changes on the resource's configuration.

This solution will meet the requirement in the most operationally efficient way, as it leverages AWS Config's features to monitor, record, and evaluate resource configurations without requiring additional tools or services.

The other options are incorrect because they either do not record the latest configuration in case of multiple configuration changes (A, C), or do not use a valid

service for evaluating resource configurations (D).

Verified References:

- > <https://docs.aws.amazon.com/config/latest/developerguide/WhatIsConfig.html>
- > <https://docs.aws.amazon.com/config/latest/developerguide/config-item-table.html>

NEW QUESTION 84

A company's application team wants to replace an internal application with a new IAM architecture that consists of Amazon EC2 instances, an IAM Lambda function, and an Amazon S3 bucket in a single IAM Region. After an architecture review, the security team mandates that no application network traffic can traverse the public internet at any point. The security team already has an SCP in place for the company's organization in IAM Organizations to restrict the creation of internet gateways, NAT gateways, and egress-only gateways.

Which combination of steps should the application team take to meet these requirements? (Select THREE.)

- A. Create an S3 endpoint that has a full-access policy for the application's VPC.
- B. Create an S3 access point for the S3 bucket.
- C. Include a policy that restricts the network origin to VPCs.
- D. Launch the Lambda function.
- E. Enable the block public access configuration.
- F. Create a security group that has an outbound rule over port 443 with a destination of the S3 endpoint. Associate the security group with the EC2 instances.
- G. Create a security group that has an outbound rule over port 443 with a destination of the S3 access point. Associate the security group with the EC2 instances.
- H. Launch the Lambda function in a VPC.

Answer: ADF

NEW QUESTION 89

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 in 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 team do to launch the EC2 instance successfully?

- A. Update the policy that is associated with the federated IAM role to allow the ec2:DescribeImages action for the forensic AMI.
- B. Update the policy that is associated with the federated IAM role to allow the ec2:StartInstances action in the security team's AWS account.
- C. Update the policy that is associated with the KMS key that is used to encrypt the forensic AMI.
- D. Configure the policy to allow the kms:DescribeKey action.
- 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 kms:DescribeKey action.
- 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 93

A company that uses AWS Organizations is migrating workloads to AWS. The company's application team determines that the workloads will use Amazon EC2 instances, Amazon S3 buckets, Amazon DynamoDB tables, and Application Load Balancers. For each resource type, the company mandates that deployments must comply with the following requirements:

- All EC2 instances must be launched from approved AWS accounts.
- All DynamoDB tables must be provisioned with a standardized naming convention.
- All infrastructure that is provisioned in any accounts in the organization must be deployed by AWS CloudFormation templates.

Which combination of steps should the application team take to meet these requirements? (Select TWO.)

- A. Create CloudFormation templates in an administrator AWS account.
- B. Share the stack sets with an application AWS account.
- C. Restrict the template to be used specifically by the application AWS account.
- D. Create CloudFormation templates in an application AWS account.
- E. Share the output with an administrator AWS account to review compliant resources.
- F. Restrict output to only the administrator AWS account.
- G. Use permissions boundaries to prevent the application AWS account from provisioning specific resources unless conditions for the internal compliance requirements are met.
- H. Use SCPs to prevent the application AWS account from provisioning specific resources unless conditions for the internal compliance requirements are met.
- I. Activate AWS Config managed rules for each service in the application AWS account.

Answer: AD

NEW QUESTION 98

For compliance reasons a Security Engineer must produce a weekly report that lists any instance that does not have the latest approved patches applied. The Engineer must also ensure that no system goes more than 30 days without the latest approved updates being applied.

What would the MOST efficient way to achieve these goals?

- A. Use Amazon Inspector to determine which systems do not have the latest patches applied, and after 30 days, redeploy those instances with the latest AMI version.
- B. Configure Amazon EC2 Systems Manager to report on instance patch compliance and enforce updates during the defined maintenance windows.
- C. Examine IAM CloudTrail logs to determine whether any instances have not restarted in the last 30 days, and redeploy those instances.
- D. Update the AMIs with the latest approved patches and redeploy each instance during the defined maintenance window.

Answer: B

Explanation:

Amazon EC2 Systems Manager is a service that helps you automatically collect software inventory, apply OS patches, create system images, and configure Windows and Linux operating systems³. You can use Systems Manager to report on instance patch compliance and enforce updates during the defined maintenance windows⁴. The other options are either inefficient or not feasible for achieving the goals.

NEW QUESTION 102

Developers in an organization have moved from a standard application deployment to containers. The Security Engineer is tasked with ensuring that the containers are secure. Which strategies will reduce the attack surface and enhance the security of the containers? (Select TWO.)

- A. Use the containers to automate security deployments.
- B. Limit resource consumption (CPU, memory), networking connections, ports, and unnecessary container libraries.
- C. Segregate containers by host, function, and data classification.
- D. Use Docker Notary framework to sign task definitions.
- E. Enable container breakout at the host kernel.

Answer: AC

Explanation:

these are the strategies that can reduce the attack surface and enhance the security of the containers. Containers are a method of packaging and running applications in isolated environments. Using containers to automate security deployments can help ensure that security patches and updates are applied consistently and quickly across the container fleet. Segregating containers by host, function, and data classification can help limit the impact of a compromise and enforce the principle of least privilege. The other options are either irrelevant or risky for securing containers.

NEW QUESTION 103

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 contractor's IAM account must not be able to gain access to any other IAM service, even if the IAM account is assigned additional permissions based on IAM group membership. What should the security engineer do to meet these requirements?

- A. Create an 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 an 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 108

An audit determined that a company's Amazon EC2 instance security group violated company policy by allowing unrestricted incoming SSH traffic. A security engineer must implement a near-real-time monitoring and alerting solution that will notify administrators of such violations.

Which solution meets these requirements with the MOST operational efficiency?

- A. Create a recurring Amazon Inspector assessment run that runs every day and uses the Network Reachability package.
- B. Create an Amazon CloudWatch rule that invokes an IAM Lambda function when an assessment run starts.
- C. Configure the Lambda function to retrieve and evaluate the assessment run report when it completes.
- D. Configure the Lambda function also to publish an Amazon Simple Notification Service (Amazon SNS) notification if there are any violations for unrestricted incoming SSH traffic.
- E. Use the restricted-ssh IAM Config managed rule that is invoked by security group configuration changes that are not compliant.
- F. Use the IAM Config remediation feature to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic.
- G. Configure VPC Flow Logs for the VPC.
- H. and specify an Amazon CloudWatch Logs group.
- I. Subscribe the CloudWatch Logs group to an IAM Lambda function that parses new log entries, detects successful connections on port 22, and publishes a notification through Amazon Simple Notification Service (Amazon SNS).
- J. Create a recurring Amazon Inspector assessment run that runs every day and uses the Security Best Practices package.
- K. Create an Amazon CloudWatch rule that invokes an IAM Lambda function when an assessment run starts.
- L. Configure the Lambda function to retrieve and evaluate the assessment run report when it completes.
- M. Configure the Lambda function also to publish an Amazon Simple Notification Service (Amazon SNS) notification if there are any violations for unrestricted incoming SSH traffic.

Answer: B

Explanation:

The most operationally efficient solution to implement a near-real-time monitoring and alerting solution that will notify administrators of security group violations is to use the restricted-ssh AWS Config managed rule that is invoked by security group configuration changes that are not compliant. This rule checks whether security groups that are in use have inbound rules that allow unrestricted SSH traffic. If a violation is detected, AWS Config can use the remediation feature to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic.

Option A is incorrect because creating a recurring Amazon Inspector assessment run that uses the Network Reachability package is not operationally efficient, as it requires setting up an assessment target and template, running the assessment every day, and invoking a Lambda function to retrieve and evaluate the assessment report. It also does not provide near-real-time monitoring and alerting, as it depends on the frequency and duration of the assessment run.

Option C is incorrect because configuring VPC Flow Logs for the VPC and specifying an Amazon CloudWatch Logs group is not operationally efficient, as it requires creating a log group and stream, enabling VPC Flow Logs for each subnet or network interface, and subscribing a Lambda function to parse and analyze the log entries. It also does not provide proactive monitoring and alerting, as it only detects successful connections on port 22 after they have occurred.

Option D is incorrect because creating a recurring Amazon Inspector assessment run that uses the Security Best Practices package is not operationally efficient, for the same reasons as option A. It also does not provide specific monitoring and alerting for security group violations, as it covers a broader range of security issues. References:

- > [AWS Config Rules]
- > [AWS Config Remediation]
- > [Amazon Inspector]
- > [VPC Flow Logs]

NEW QUESTION 111

A security engineer is configuring a new website that is named example.com. The security engineer wants to secure communications with the website by requiring users to connect to example.com through HTTPS.

Which of the following is a valid option for storing SSL/TLS certificates?

- A. Custom SSL certificate that is stored in AWS Key Management Service (AWS KMS)
- B. Default SSL certificate that is stored in Amazon CloudFront.
- C. Custom SSL certificate that is stored in AWS Certificate Manager (ACM)
- D. Default SSL certificate that is stored in Amazon S3

Answer: C

NEW QUESTION 114

A company's Security Auditor discovers that users are able to assume roles without using multi-factor authentication (MFA). An example of a current policy being applied to these users is as follows:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Principal": {
        "AWS": "arn:aws:iam::555555555555:root"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "Bool": { "aws:MultiFactorAuthPresent": false}
      }
    }
  ]
}
```

The Security Auditor finds that the users who are able to assume roles without MFA are all coming from the IAM CLI. These users are using long-term IAM credentials. Which changes should a Security Engineer implement to resolve this security issue? (Select TWO.)

- A)


```
"Effect": "Deny",
"Condition": { "Bool": { "aws:MultiFactorAuthPresent": false} }
```
- B)


```
"Effect": "Allow",
"Condition": { "Bool": { "aws:MultiFactorAuthPresent": true} }
```
- C)


```
"Effect": "Allow","Condition": { "BoolIfExists": { "aws:MultiFactorAuthPresent": true} }
```
- D)


```
"Effect": "Deny","Condition": { "BoolIfExists": { "aws:MultiFactorAuthPresent": false} }
```
- E)


```
"Effect": "Deny","Condition": { "BoolIfNotExist": { "aws:MultiFactorAuthPresent": true} }
```

- A. Option A
- B. Option B
- C. Option C
- D. Option D
- E. Option E

Answer: AD

NEW QUESTION 119

A company is evaluating its security posture. In the past, the company has observed issues with specific hosts and host header combinations that affected the company's business. The company has configured AWS WAF web ACLs as an initial step to mitigate these issues. The company must create a log analysis solution for the AWS WAF web ACLs to monitor problematic activity. The company wants to process all the AWS WAF logs in a central location. The company must have the ability to filter out requests based on specific hosts. A security engineer starts to enable access logging for the AWS WAF web ACLs. What should the security engineer do next to meet these requirements with the MOST operational efficiency?

- A. Specify Amazon Redshift as the destination for the access log

- B. Deploy the Amazon Athena Redshift connecto
- C. Use Athena to query the data from Amazon Redshift and to filter the logs by host.
- D. Specify Amazon CloudWatch as the destination for the access log
- E. Use Amazon CloudWatch Logs Insights to design a query to filter the logs by host.
- F. Specify Amazon CloudWatch as the destination for the access log
- G. Export the CloudWatch logs to an Amazon S3 bucke
- H. Use Amazon Athena to query the logs and to filter the logs by host.
- I. Specify Amazon CloudWatch as the destination for the access log
- J. Use Amazon Redshift Spectrum to query the logs and to filter the logs by host.

Answer: C

Explanation:

The correct answer is C. Specify Amazon CloudWatch as the destination for the access logs. Export the CloudWatch logs to an Amazon S3 bucket. Use Amazon Athena to query the logs and to filter the logs by host.

According to the AWS documentation¹, AWS WAF offers logging for the traffic that your web ACLs analyze. The logs include information such as the time that AWS WAF received the request from your protected AWS resource, detailed information about the request, and the action setting for the rule that the request matched. You can send your logs to an Amazon CloudWatch Logs log group, an Amazon Simple Storage Service (Amazon S3) bucket, or an Amazon Kinesis Data Firehose.

To create a log analysis solution for the AWS WAF web ACLs, you can use Amazon Athena, which is an interactive query service that makes it easy to analyze data in Amazon S3 using standard SQL². You can use Athena to query and filter the AWS WAF logs by host or any other criteria. Athena is serverless, so there is no infrastructure to manage, and you pay only for the queries that you run.

To use Athena with AWS WAF logs, you need to export the CloudWatch logs to an S3 bucket. You can do this by creating a subscription filter that sends your log events to a Kinesis Data Firehose delivery stream, which then delivers the data to an S3 bucket³. Alternatively, you can use AWS DMS to migrate your CloudWatch logs to S3⁴.

After you have exported your CloudWatch logs to S3, you can create a table in Athena that points to your S3 bucket and use the AWS service log format that matches your log schema⁵. For example, if you are using JSON format for your AWS WAF logs, you can use the AWSJSONSerDe serde. Then you can run SQL queries on your Athena table and filter the results by host or any other field in your log data.

Therefore, this solution meets the requirements of creating a log analysis solution for the AWS WAF web ACLs with the most operational efficiency. This solution does not require setting up any additional infrastructure or services, and it leverages the existing capabilities of CloudWatch, S3, and Athena.

The other options are incorrect because:

- A. Specifying Amazon Redshift as the destination for the access logs is not possible, because AWS WAF does not support sending logs directly to Redshift. You would need to use an intermediate service such as Kinesis Data Firehose or AWS DMS to load the data from CloudWatch or S3 to Redshift. Deploying the Amazon Athena Redshift connector is not necessary, because you can query Redshift data directly from Athena without using a connector⁶. This solution would also incur additional costs and operational overhead of managing a Redshift cluster.
- B. Specifying Amazon CloudWatch as the destination for the access logs is possible, but using Amazon CloudWatch Logs Insights to design a query to filter the logs by host is not efficient or scalable. CloudWatch Logs Insights is a feature that enables you to interactively search and analyze your log data in CloudWatch Logs⁷. However, CloudWatch Logs Insights has some limitations, such as a maximum query duration of 20 minutes, a maximum of 20 log groups per query, and a maximum retention period of 24 months⁸. These limitations may affect your ability to perform complex and long-running analysis on your AWS WAF logs.
- D. Specifying Amazon CloudWatch as the destination for the access logs is possible, but using Amazon Redshift Spectrum to query the logs and filter them by host is not efficient or cost-effective. Redshift Spectrum is a feature of Amazon Redshift that enables you to run queries against exabytes of data in S3 without loading or transforming any data⁹. However, Redshift Spectrum requires a Redshift cluster to process the queries, which adds additional costs and operational overhead. Redshift Spectrum also charges you based on the number of bytes scanned by each query, which can be expensive if you have large volumes of log data¹⁰.

References:

1: Logging AWS WAF web ACL traffic - Amazon Web Services 2: What Is Amazon Athena? - Amazon Athena 3: Streaming CloudWatch Logs Data to Amazon S3 - Amazon CloudWatch Logs 4: Migrate data from CloudWatch Logs using AWS Database Migration Service - AWS Database Migration Service 5: Querying AWS service logs - Amazon Athena 6: Querying data from Amazon Redshift - Amazon Athena 7: Analyzing log data with CloudWatch Logs Insights - Amazon CloudWatch Logs 8: CloudWatch Logs Insights quotas - Amazon CloudWatch 9: Querying external data using Amazon Redshift Spectrum - Amazon Redshift 10: Amazon Redshift Spectrum pricing - Amazon Redshift

NEW QUESTION 122

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 123

A company's on-premises networks are connected to VPCs using an IAM Direct Connect gateway. The company's on-premises application needs to stream data using an existing Amazon Kinesis Data Firehose delivery stream. The company's security policy requires that data be encrypted in transit using a private network. How should the company meet these requirements?

- A. Create a VPC endpoint for Kinesis Data Firehose
- B. Configure the application to connect to the VPC endpoint.
- C. Configure an IAM policy to restrict access to Kinesis Data Firehose using a source IP condition. Configure the application to connect to the existing Firehose delivery stream.
- D. Create a new TLS certificate in IAM Certificate Manager (ACM). Create a public-facing Network Load Balancer (NLB) and select the newly created TLS certificate
- E. Configure the NLB to forward all traffic to Kinesis Data Firehose
- F. Configure the application to connect to the NLB.
- G. Peer the on-premises network with the Kinesis Data Firehose VPC using Direct Connect
- H. Configure the application to connect to the existing Firehose delivery stream.

Answer: A

Explanation:

To stream data using an existing Amazon Kinesis Data Firehose delivery stream and encrypt it in transit using a private network, the company should do the following:

- Create a VPC endpoint for Kinesis Data Firehose. This allows the company to use a private connection between their VPC and Kinesis Data Firehose without requiring an internet gateway or NAT device.
- Configure the application to connect to the VPC endpoint. This allows the application to stream data using Kinesis Data Firehose over AWS PrivateLink, which encrypts all traffic with TLS.

NEW QUESTION 128

A company has recently recovered from a security incident that required the restoration of Amazon EC2 instances from snapshots. The company uses an AWS Key

Management Service (AWS KMS) customer managed key to encrypt all Amazon Elastic Block Store (Amazon EBS) snapshots.

The company performs a gap analysis of its disaster recovery procedures and backup strategies. A security engineer needs to implement a solution so that the company can recover the EC2 instances if the AWS account is compromised and the EBS snapshots are deleted.

Which solution will meet this requirement?

- A. Create a new Amazon S3 bucket
- B. Use EBS lifecycle policies to move EBS snapshots to the new S3 bucket
- C. Use lifecycle policies to move snapshots to the S3 Glacier Instant Retrieval storage class
- D. Use S3 Object Lock to prevent deletion of the snapshots.
- E. Use AWS Systems Manager to distribute a configuration that backs up all attached disks to Amazon S3.
- F. Create a new AWS account that has limited privileges
- G. Allow the new account to access the KMS key that encrypts the EBS snapshots
- H. Copy the encrypted snapshots to the new account on a recurring basis.
- I. Use AWS Backup to copy EBS snapshots to Amazon S3. Use S3 Object Lock to prevent deletion of the snapshots.

Answer: C

Explanation:

This solution meets the requirement of recovering the EC2 instances if the AWS account is compromised and the EBS snapshots are deleted. By creating a new AWS account with limited privileges, the company can isolate the backup snapshots from the main account and reduce the risk of accidental or malicious deletion. By allowing the new account to access the KMS key that encrypts the EBS snapshots, the company can ensure that the snapshots are copied in an encrypted form and can be decrypted when needed. By copying the encrypted snapshots to the new account on a recurring basis, the company can maintain a consistent backup schedule and minimize data loss.

NEW QUESTION 129

A company has a web-based application using Amazon CloudFront and running on Amazon Elastic Container Service (Amazon ECS) behind an Application Load Balancer (ALB). The ALB is terminating TLS and balancing load across ECS service tasks. A security engineer needs to design a solution to ensure that application content is accessible only through CloudFront and that it is never accessible directly.

How should the security engineer build the MOST secure solution?

- A. Add an origin custom header Set the viewer protocol policy to HTTP and HTTPS Set the origin protocol policy to HTTPS only Update the application to validate the CloudFront custom header
- B. Add an origin custom header Set the viewer protocol policy to HTTPS only Set the origin protocol policy to match viewer Update the application to validate the CloudFront custom header.
- C. Add an origin custom header Set the viewer protocol policy to redirect HTTP to HTTPS Set the origin protocol policy to HTTP only Update the application to validate the CloudFront custom header.
- D. Add an origin custom header Set the viewer protocol policy to redirect HTTP to HTTP
- E. Set the origin protocol policy to HTTPS only Update the application to validate the CloudFront custom header

Answer: D

Explanation:

To ensure that application content is accessible only through CloudFront and not directly, the security engineer should do the following:

- Add an origin custom header. This is a header that CloudFront adds to the requests that it sends to the origin, but viewers cannot see or modify.
- Set the viewer protocol policy to redirect HTTP to HTTPS. This ensures that the viewers always use HTTPS when they access the website through CloudFront.

- Set the origin protocol policy to HTTPS only. This ensures that CloudFront always uses HTTPS when it connects to the origin.
- Update the application to validate the CloudFront custom header. This means that the application checks if the request has the custom header and only responds if it does. Otherwise, it denies or ignores the request. This prevents users from bypassing CloudFront and accessing the content directly on the origin.

NEW QUESTION 130

A company has deployed servers on Amazon EC2 instances in a VPC. External vendors access these servers over the internet. Recently, the company deployed a new application on EC2 instances in a new CIDR range. The company needs to make the application available to the vendors.

A security engineer verified that the associated security groups and network ACLs are allowing the required ports in the inbound direction. However, the vendors cannot connect to the application.

Which solution will provide the vendors access to the application?

- A. Modify the security group that is associated with the EC2 instances to have the same outbound rules as inbound rules.
- B. Modify the network ACL that is associated with the CIDR range to allow outbound traffic to ephemeral ports.
- C. Modify the inbound rules on the internet gateway to allow the required ports.
- D. Modify the network ACL that is associated with the CIDR range to have the same outbound rules as inbound rules.

Answer: B

Explanation:

The correct answer is B. Modify the network ACL that is associated with the CIDR range to allow outbound traffic to ephemeral ports.

This answer is correct because network ACLs are stateless, which means that they do not automatically allow return traffic for inbound connections. Therefore, the network ACL that is associated with the CIDR range of the new application must have outbound rules that allow traffic to ephemeral ports, which are the temporary ports used by the vendors' machines to communicate with the application servers. Ephemeral ports are typically in the range of 1024-65535. If the network ACL does not have such rules, the vendors will not be able to connect to the application.

The other options are incorrect because:

- A. Modifying the security group that is associated with the EC2 instances to have the same outbound rules as inbound rules is not a solution, because security groups are stateful, which means that they automatically allow return traffic for inbound connections. Therefore, there is no need to add outbound rules to the security group for the vendors to access the application.
- C. Modifying the inbound rules on the internet gateway to allow the required ports is not a solution, because internet gateways do not have inbound or outbound rules. Internet gateways are VPC components that enable communication between instances in a VPC and the internet. They do not filter traffic based on ports or protocols.
- D. Modifying the network ACL that is associated with the CIDR range to have the same outbound rules as inbound rules is not a solution, because it does not address the issue of ephemeral ports. The outbound rules of the network ACL must match the ephemeral port range of the vendors' machines, not necessarily the inbound rules of the network ACL.

References:

1: Ephemeral port - Wikipedia 2: Security groups for your VPC - Amazon Virtual Private Cloud 3: Internet gateways - Amazon Virtual Private Cloud 4: Network ACLs - Amazon Virtual Private Cloud

NEW QUESTION 134

There are currently multiple applications hosted in a VPC. During monitoring it has been noticed that multiple port scans are coming in from a specific IP Address block. The internal security team has requested that all offending IP Addresses be denied for the next 24 hours. Which of the following is the best method to quickly and temporarily deny access from the specified IP Address's.

Please select:

- A. Create an AD policy to modify the Windows Firewall settings on all hosts in the VPC to deny access from the IP Address block.
- B. Modify the Network ACLs associated with all public subnets in the VPC to deny access from the IP Address block.
- C. Add a rule to all of the VPC Security Groups to deny access from the IP Address block.
- D. Modify the Windows Firewall settings on all AMI'S that your organization uses in that VPC to deny access from the IP address block.

Answer: B

Explanation:

NACL acts as a firewall at the subnet level of the VPC and we can deny the offending IP address block at the subnet level using NACL rules to block the incoming traffic to the VPC instances. Since NACL rules are applied as per the Rule numbers make sure that this rule number should take precedence over other rule numbers if there are any such rules that will allow traffic from these IP ranges. The lowest rule number has more precedence over a rule that has a higher number. The IAM Documentation mentions the following as a best practices for IAM users

For extra security, enable multi-factor authentication (MFA) for privileged IAM users (users who are allowed access to sensitive resources or APIs). With MFA, users have a device that generates a unique authentication code (a one-time password, or OTP). Users must provide both their normal credentials (like their user name and password) and the OTP. The MFA device can either be a special piece of hardware, or it can be a virtual device (for example, it can run in an app on a smartphone).

Options C is invalid because these options are not available Option D is invalid because there is not root access for users

For more information on IAM best practices, please visit the below URL: <https://docs.IAM.amazon.com/IAM/latest/UserGuide/best-practices.html>

The correct answer is: Modify the Network ACLs associated with all public subnets in the VPC to deny access from the IP Address block.

omit your Feedback/Queries to our Experts

NEW QUESTION 138

Your company has a set of EC2 Instances defined in IAM. These EC2 Instances have strict security groups attached to them. You need to ensure that changes to the Security groups are noted and acted on accordingly. How can you achieve this?

Please select:

- A. Use Cloudwatch logs to monitor the activity on the Security Group
- B. Use filters to search for the changes and use SNS for the notification.
- C. Use Cloudwatch metrics to monitor the activity on the Security Group
- D. Use filters to search for the changes and use SNS for the notification.
- E. Use IAM inspector to monitor the activity on the Security Group
- F. Use filters to search for the changes and use SNS for the notification.
- G. Use Cloudwatch events to be triggered for any changes to the Security Group
- H. Configure the Lambda function for email notification as well.

Answer: D

Explanation:

The below diagram from an IAM blog shows how security groups can be monitored
 C:\Users\wk\Desktop\mudassar\Untitled.jpg



Option A is invalid because you need to use Cloudwatch Events to check for changes, Option B is invalid because you need to use Cloudwatch Events to check for changes

Option C is invalid because IAM inspector is not used to monitor the activity on Security Groups. For more information on monitoring security groups, please visit the below URL:

<https://iam.amazonaws.com/blogs/security/how-to-automatically-revert-and-receive-notifications-about-changes-to-vpc-security-groups/>

The correct answer is: Use Cloudwatch events to be triggered for any changes to the Security Groups. Configure the Lambda function for email notification as well. Submit your Feedback/Queries to our Experts

NEW QUESTION 139

A developer has created an AWS Lambda function in a company's development account. The Lambda function requires the use of an AWS Key Management Service (AWS KMS) customer managed key that exists in a security account that the company's security team controls. The developer obtains the ARN of the KMS key from a previous Lambda function in the development account. The previous Lambda function had been working properly with the KMS key. When the developer uses the ARN and tests the new Lambda function an error message states that access is denied to the KMS key in the security account. The developer tests the previous Lambda function that uses the same KMS key and discovers that the previous Lambda function still can encrypt data as expected. A security engineer must resolve the problem so that the new Lambda function in the development account can use the KMS key from the security account. Which combination of steps should the security engineer take to meet these requirements? (Select TWO.)

- A. In the security account configure an IAM role for the new Lambda function
- B. Attach an IAM policy that allows access to the KMS key in the security account.
- C. In the development account configure an IAM role for the new Lambda function
- D. Attach a key policy that allows access to the KMS key in the security account.
- E. In the development account configure an IAM role for the new Lambda function
- F. Attach an IAM policy that allows access to the KMS key in the security account.
- G. Configure a key policy for the KMS key in the security account to allow access to the IAM role of the new Lambda function in the security account.
- H. Configure a key policy for the KMS key in the security account to allow access to the IAM role of the new Lambda function in the development account.

Answer: CE

Explanation:

To allow cross-account access to a KMS key, the key policy of the KMS key must grant permission to the external account or principal, and the IAM policy of the external account or principal must delegate the key policy permission. In this case, the new Lambda function in the development account needs to use the KMS key in the security account, so the key policy of the KMS key must allow access to the IAM role of the new Lambda function in the development account (option E), and the IAM role of the new Lambda function in the development account must have an IAM policy that allows access to the KMS key in the security account (option C). Option A is incorrect because it creates an IAM role for the new Lambda function in the security account, not in the development account. Option B is incorrect because it attaches a key policy to an IAM role, which is not valid. Option D is incorrect because it allows access to the IAM role of the new Lambda function in the security account, not in the development account. Verified References:

> <https://docs.aws.amazon.com/autoscaling/ec2/userguide/key-policy-requirements-EBS-encryption.html>

NEW QUESTION 143

A company has an AWS Lambda function that creates image thumbnails from larger images. The Lambda function needs read and write access to an Amazon S3 bucket in the same AWS account.

Which solutions will provide the Lambda function this access? (Select TWO.)

- A. Create an IAM user that has only programmatic access
- B. Create a new access key pair
- C. Add environmental variables to the Lambda function with the access key ID and secret access key
- D. Modify the Lambda function to use the environmental variables at run time during communication with Amazon S3.
- E. Generate an Amazon EC2 key pair
- F. Store the private key in AWS Secrets Manager
- G. Modify the Lambda function to retrieve the private key from Secrets Manager and to use the private key during communication with Amazon S3.
- H. Create an IAM role for the Lambda function
- I. Attach an IAM policy that allows access to the S3 bucket.

- J. Create an IAM role for the Lambda function
- K. Attach a bucket policy to the S3 bucket to allow access. Specify the function's IAM role as the principal.
- L. Create a security group
- M. Attach the security group to the Lambda function
- N. Attach a bucket policy that allows access to the S3 bucket through the security group ID.

Answer: CD

NEW QUESTION 147

A System Administrator is unable to start an Amazon EC2 instance in the eu-west-1 Region using an IAM role. The same System Administrator is able to start an EC2 instance in the eu-west-2 and eu-west-3 Regions. The IAM System Administrator access policy attached to the System Administrator IAM role allows unconditional access to all IAM services and resources within the account.

Which configuration caused this issue?

A) An SCP is attached to the account with the following permission statement:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "All",
      "Action": "*",
      "Resource": "*"
    },
    {
      "Effect": "Deny",
      "NotAction": [
        "iam:*",
        "organizations:*",
        "route53:*",
        "budgets:*",
        "waf:*",
        "cloudfront:*",
        "globalaccelerator:*",
        "importexport:*",
        "support:*"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestedRegion": [
            "eu-west-*"
          ]
        }
      }
    }
  ]
}
```

B) A permission boundary policy is attached to the System Administrator role with the following permission statement:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:*"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Deny",
      "NotAction": [
        "iam:*",
        "organizations:*",
        "route53:*",
        "budgets:*",
        "waf:*",
        "cloudfront:*",
        "globalaccelerator:*",
        "importexport:*",
        "support:*",
        "ec2:*"
      ],
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestedRegion": [
            "eu-west-1"
          ]
        }
      }
    }
  ]
}
```

C) A permission boundary is attached to the System Administrator role with the following permission statement:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:*"
      ],
      "Resource": "*"
    },
    {
      "Version": "2012-10-17",
      "Statement": [
        {
          "Effect": "Allow",
          "Action": "ec2:*",
          "Resource": "*",
          "Condition": {
            "StringEquals": {
              "aws:RequestedRegion": [
                "eu-west-1"
              ]
            }
          }
        }
      ]
    }
  ]
}
```

D) An SCP is attached to the account with the following statement:

```

"Version": "2012-10-17",
"Statement": [
  {
    "Effect": "Allow",
    "Action": "*",
    "Resource": "*"
  },
  {
    "Effect": "Deny",
    "NotAction": [
      "iam:*",
      "organizations:*",
      "route53:*",
      "budgets:*",
      "waf:*",
      "cloudfront:*",
      "globalaccelerators:*",
      "lambda:*",
      "support:*",
      "ec2:*"
    ],
    "Resource": "*",
    "Condition": {
      "StringEquals": {
        "aws:RequestedRegion": "us-east-1"
      }
    }
  }
]

```

- A. Option A
- B. Option B
- C. Option C
- D. Option D

Answer: B

NEW QUESTION 150

A company has a batch-processing system that uses Amazon S3, Amazon EC2, and AWS Key Management Service (AWS KMS). The system uses two AWS accounts: Account A and Account B. Account A hosts an S3 bucket that stores the objects that will be processed. The S3 bucket also stores the results of the processing. All the S3 bucket objects are encrypted by a KMS key that is managed in Account A. Account B hosts a VPC that has a fleet of EC2 instances that access the S3 bucket in Account A by using statements in the bucket policy. The VPC was created with DNS hostnames enabled and DNS resolution enabled. A security engineer needs to update the design of the system without changing any of the system's code. No AWS API calls from the batch-processing EC2 instances can travel over the internet. Which combination of steps will meet these requirements? (Select TWO.)

- A. In the Account B VPC, create a gateway VPC endpoint for Amazon S3. For the gateway VPC endpoint, create a resource policy that allows the s3:GetObject, s3:ListBucket, s3:PutObject, and s3:PutObjectAcl actions for the S3 bucket.
- B. In the Account B VPC, create an interface VPC endpoint for Amazon S3. For the interface VPC endpoint, create a resource policy that allows the s3:GetObject, s3:ListBucket, s3:PutObject, and s3:PutObjectAcl actions for the S3 bucket.
- C. In the Account B VPC, create an interface VPC endpoint for AWS KM
- D. For the interface VPC endpoint, create a resource policy that allows the kms:Encrypt, kms:Decrypt, and kms:GenerateDataKey actions for the KMS ke
- E. Ensure that private DNS is turned on for the endpoint.
- F. In the Account B VPC, create an interface VPC endpoint for AWS KM
- G. For the interface VPC endpoint, create a resource policy that allows the kms:Encrypt, kms:Decrypt, and kms:GenerateDataKey actions for the KMS ke
- H. Ensure that private DNS is turned off for the endpoint.
- I. In the Account B VPC, verify that the S3 bucket policy allows the s3:PutObjectAcl action for cross-account us
- J. In the Account B VPC, create a gateway VPC endpoint for Amazon S3. For the gateway VPC endpoint, create a resource policy that allows the s3:GetObject, s3:ListBucket, and s3:PutObject actions for the S3 bucket.

Answer: BC

NEW QUESTION 151

A company is expanding its group of stores. On the day that each new store opens, the company wants to launch a customized web application for that store. Each store's application will have a non-production environment and a production environment. Each environment will be deployed in a separate AWS account. The company uses AWS Organizations and has an OU that is used only for these accounts. The company distributes most of the development work to third-party development teams. A security engineer needs to ensure that each team follows the company's deployment plan for AWS resources. The security engineer also must limit access to the deployment plan to only the developers who need access. The security engineer already has created an AWS CloudFormation template that implements the deployment plan. What should the security engineer do next to meet the requirements in the MOST secure way?

- A. Create an AWS Service Catalog portfolio in the organization's management account
- B. Upload the CloudFormation template
- C. Add the template to the portfolio's product list
- D. Share the portfolio with the OIJ.
- E. Use the CloudFormation CLI to create a module from the CloudFormation template
- F. Register the module as a private extension in the CloudFormation registry
- G. Publish the extension
- H. In the OU, create an SCP that allows access to the extension.
- I. Create an AWS Service Catalog portfolio in the organization's management account
- J. Upload the CloudFormation template
- K. Add the template to the portfolio's product list
- L. Create an IAM role that has a trust policy that allows cross-account access to the portfolio for users in the OU account

- M. Attach the AWSServiceCatalogEndUserFullAccess managed policy to the role.
- N. Use the CloudFormation CLI to create a module from the CloudFormation template.
- O. Register the module as a private extension in the CloudFormation registry.
- P. Publish the extension.
- Q. Share the extension with the OU.

Answer: A

Explanation:

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

According to the AWS documentation, AWS Service Catalog is a service that allows you to create and manage catalogs of IT services that are approved for use on AWS. You can use Service Catalog to centrally manage commonly deployed IT services and help achieve consistent governance and compliance requirements, while enabling users to quickly deploy only the approved IT services they need.

To use Service Catalog with multiple AWS accounts, 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 Service Catalog as a service principal for AWS Organizations, which lets you share your portfolios with organizational units (OUs) or accounts in your organization.

To create a Service Catalog portfolio, you need to use an administrator account, such as the organization's management account. You can upload your CloudFormation template as a product in your portfolio, and define constraints and tags for it. You can then share your portfolio with the OU that contains the accounts for the web applications. This will allow the developers in those accounts to launch products from the shared portfolio using the Service Catalog end user console.

Option B is incorrect because CloudFormation modules are reusable components that encapsulate one or more resources and their configurations. They are not meant to be used as templates for deploying entire stacks of resources. Moreover, sharing a module with an OU does not grant access to launch stacks from it.

Option C is incorrect because creating an IAM role that has a trust policy that allows cross-account access to the portfolio is not secure. It would allow any user in the OU accounts to assume the role and access the portfolio, regardless of their job function or access requirements.

Option D is incorrect because sharing a module with an OU does not grant access to launch stacks from it. It also does not limit access to the deployment plan to only the developers who need access.

NEW QUESTION 153

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

The company is worried that the S3 URL might still be accessible directly and that requests can bypass the CloudFront distribution. Which combination of steps should the company take to remove direct access to the S3 URL? (Select TWO.)

- A. Select "Restrict Bucket Access" in the origin settings of the CloudFront distribution.
- B. Create an origin access identity (OAI) for the S3 origin.
- C. Update the S3 bucket policy to allow s3 GetObject with a condition that the IAM Referer key matches the secret value. Deny all other requests.
- D. Configure the S3 bucket policy so that only the origin access identity (OAI) has read permission for objects in the bucket.
- E. Add an origin custom header that has the name Referer to the CloudFront distribution. Give the header a secret value.

Answer: AD

NEW QUESTION 157

A company has a set of EC2 instances hosted in IAM. The EC2 instances have EBS volumes which is used to store critical information. There is a business continuity requirement to ensure high availability for the EBS volumes. How can you achieve this?

- A. Use lifecycle policies for the EBS volumes.
- B. Use EBS snapshots.
- C. Use EBS volume replication.
- D. Use EBS volume encryption.

Answer: B

Explanation:

Data stored in Amazon EBS volumes is redundantly stored in multiple physical locations as part of normal operation of those services and at no additional charge. However, Amazon EBS replication is stored within the same availability zone, not across multiple zones; therefore, it is highly recommended that you conduct regular snapshots to Amazon S3 for long-term data durability. Option A is invalid because there is no lifecycle policy for EBS volumes. Option C is invalid because there is no EBS volume replication. Option D is invalid because EBS volume encryption will not ensure business continuity. For information on security for Compute Resources, please visit the below URL: https://d1.awsstatic.com/whitepapers/Security/Security_Compute_Services_Whitepaper.pdf

NEW QUESTION 161

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

SCS-C02 Practice Exam Features:

- * SCS-C02 Questions and Answers Updated Frequently
- * SCS-C02 Practice Questions Verified by Expert Senior Certified Staff
- * SCS-C02 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * SCS-C02 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The SCS-C02 Practice Test Here](#)