



- Expert Verified, Online, **Free**.



CERTIFICATION TEST

- CertificationTest.net - Cheap & Quality Resources With Best Support

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? (Choose two.)

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

Suggested Answer: CD

Community vote distribution

CD (100%)

🗳️ 👤 **Raphaello** Highly Voted 1 year, 6 months ago

CD

Why the answers for many questions are categorically wrong?

upvoted 7 times

🗳️ 👤 **5pqo3pqn1l** 2 weeks, 5 days ago

thanks valid information Ref - <https://bitly.cx/o2VIQ>

upvoted 1 times

🗳️ 👤 **Robertwilliamm** Most Recent 4 days, 4 hours ago

Selected Answer: CD

C and D is Correct Options

Thanks to SkillCertExams I successfully cleared my SCS-C02 exam today.

upvoted 1 times

🗳️ 👤 **PrabhuKK** 3 months, 3 weeks ago

Selected Answer: CD

C. IAM role for Lambda with policy: This is the most common and recommended approach. You create an IAM role specifically for the Lambda function, then attach an IAM policy to the role that grants the necessary read and write permissions to the S3 bucket. Lambda functions assume this role when they are executed and can then interact with the S3 bucket.

D. IAM role for Lambda with a bucket policy: Another valid approach is to create an IAM role for the Lambda function and attach a bucket policy to the S3 bucket that allows access from the Lambda function's role. The bucket policy explicitly grants access to the function's role, so no IAM role-specific permissions are needed for the Lambda function itself, although this option is less common than the previous one.

upvoted 2 times

🗳️ 👤 **awssecuritynewbie** 9 months, 1 week ago

Selected Answer: CD

for sure this will be the use of AWS ROLES it is simple, you attach a role to the lambda to be able to access certain S3 bucket.

upvoted 2 times

🗳️ 👤 **FunkyFresco** 10 months, 1 week ago

Selected Answer: CD



C and D

upvoted 1 times

🗳️ 👤 **shailvardhan** 1 year, 1 month ago

Selected Answer: CD

100% it is C and D
upvoted 1 times

  **leaddev** 1 year, 3 months ago

C & D are the correct answers
upvoted 1 times

  **Raphaello** 1 year, 4 months ago

Selected Answer: CD

CD
Always create execution role for your lambda function, as a best practice.
upvoted 1 times



  **aescudero51** 1 year, 4 months ago

Respuesta seleccionada: CD
upvoted 1 times

  **Daniel76** 1 year, 7 months ago



Selected Answer: CD

Answer is CD
upvoted 3 times

  **lalee2** 1 year, 8 months ago



Selected Answer: CD

same account. for ABAC IAM role w/ policy will do. RBAC requires principal
upvoted 1 times



  **pupsik** 1 year, 8 months ago

Selected Answer: CD



Permission needs to be specified either from Lambda's role policy, or from S3 bucket's resource policy.
upvoted 2 times

  **KR693** 1 year, 8 months ago

c and d
upvoted 2 times

  **dhax12** 1 year, 8 months ago

Any type of key means a long term access. Always use IAM roles to keep access temporarily. Answer should be C&D
upvoted 2 times

  **RayHK** 1 year, 8 months ago

vote for C and D
upvoted 3 times

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

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **Malluchan** 3 months ago

Selected Answer: C

No Doubt, it is C

upvoted 1 times

🗳️ 👤 **leaddev** 7 months ago

Selected Answer: C

ACM is correct choice

upvoted 1 times

🗳️ 👤 **FunkyFresco** 10 months, 1 week ago

Selected Answer: C

C option

upvoted 1 times

🗳️ 👤 **Siva1971** 10 months, 3 weeks ago

It is correct -C

upvoted 1 times

🗳️ 👤 **shailvardhan** 1 year, 1 month ago

Selected Answer: C

yes. it is C

upvoted 1 times

🗳️ 👤 **[Removed]** 1 year, 1 month ago

Correct C

upvoted 1 times

🗳️ 👤 **Raphaello** 1 year, 4 months ago

Selected Answer: C

Correct answer is C

upvoted 1 times

🗳️ 👤 **aescudero51** 1 year, 4 months ago

Respuesta seleccionada: C

upvoted 1 times

🗳️ 👤 **Daniel76** 1 year, 7 months ago

Selected Answer: C

ACM is the natural choice.

upvoted 1 times

🗳️ 👤 **imymoco** 1 year, 8 months ago

Answer is C

upvoted 1 times

🗨️ 👤 **lalee2** 1 year, 8 months ago

Selected Answer: C

ACM do the job

upvoted 1 times

🗨️ 👤 **pupsik** 1 year, 8 months ago

Selected Answer: C

ACM is where you would normally store SSL certificates.

upvoted 1 times

🗨️ 👤 **KR693** 1 year, 8 months ago

Answer C

upvoted 1 times

🗨️ 👤 **100fold** 1 year, 8 months ago

Selected Answer: C

Answer C

upvoted 1 times

🗨️ 👤 **Certified101** 1 year, 8 months ago

Selected Answer: C

C is correct

upvoted 2 times

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

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

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

A compromised EC2 instance's metadata must be updated with corresponding incident ticket information.

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

Any investigative activity during the collection of volatile data must be captured as part of the process.

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

- A. Gather any relevant metadata for the compromised EC2 instance. Enable termination protection. Isolate the instance by updating the instance's security groups to restrict access. Detach the instance from any Auto Scaling groups that the instance is a member of. Deregister the instance from any Elastic Load Balancing (ELB) resources.
- B. Gather any relevant metadata for the compromised EC2 instance. Enable termination protection. Move the instance to an isolation subnet that denies all source and destination traffic. Associate the instance with the subnet to restrict access. Detach the instance from any Auto Scaling groups that the instance is a member of. Deregister the instance from any Elastic Load Balancing (ELB) resources.
- C. Use Systems Manager Run Command to invoke scripts that collect volatile data.
- D. Establish a Linux SSH or Windows Remote Desktop Protocol (RDP) session to the compromised EC2 instance to invoke scripts that collect volatile data.
- E. Create a snapshot of the compromised EC2 instance's EBS volume for follow-up investigations. Tag the instance with any relevant metadata and incident ticket information.
- F. Create a Systems Manager State Manager association to generate an EBS volume snapshot of the compromised EC2 instance. Tag the instance with any relevant metadata and incident ticket information.

Suggested Answer: ACE

Community vote distribution

ACE (100%)

🗳️ 👤 **pupsik** Highly Voted 1 year, 8 months ago

Selected Answer: ACE

The reason it is not "B" is because you cannot move a running instance into a different subnet.

upvoted 22 times

🗳️ 👤 **aguilar404** 1 year, 3 months ago

these are the comments that we, as people with less knowledge are looking for... thank you!

upvoted 7 times

🗳️ 👤 **Radhey_sen** Most Recent 1 month, 1 week ago

Selected Answer: ACE

ACE is the only correct answer

upvoted 1 times

🗳️ 👤 **bukkanni** 10 months ago

B is wrong because it says "Move the instance to an isolation subnet that denies all source and destination traffic. ". Even if you somehow snapshot it and move it to a different subnet, the fact that the isolation subnet denies all source and destination traffic would mean that the NACL will allow no inbound or outbound traffic. That includes any traffic to reach the instance for forensic evaluation.

upvoted 2 times

🗳️ 👤 **FunkyFresco** 10 months, 1 week ago

Selected Answer: ACE

ACE are correct

upvoted 1 times

🗳️ 👤 **RobWilliamsToronto** 11 months, 1 week ago

While A is still better than B (because you cannot update a running instances network placement/subnet). Updating an instances SG can be risky. Often times an instances SG is shared with multiple instances. A better wording for A would be to "replace the instances SG with a SG that restricts access" rather than 'updating'.

upvoted 2 times

🗨️ 👤 **Raphaello** 1 year, 4 months ago

Selected Answer: ACE

Correct answers: ACE

These describe the ideal steps to isolate an instance, and collect data required for forensics investigation, all limiting the spread of malware.

upvoted 1 times

🗨️ 👤 **csG13** 1 year, 6 months ago

A & C are correct. Since it's an SSM managed node already, why not F?

upvoted 1 times

🗨️ 👤 **Sab31** 1 year, 6 months ago

Can someone share why not F? As it automated the EBS backup process. Hence reducing the overhead.

upvoted 1 times

🗨️ 👤 **Daniel76** 1 year, 6 months ago

F is technically feasible but SSM state manager is used for routine backup of EC2 instances. In this case the snapshot is one-off and you cannot automate the second part that is tagging with metadata and incident ticket info. So it is not appropriate.

upvoted 6 times

🗨️ 👤 **Daniel76** 1 year, 6 months ago

Selected Answer: ACE

between C and D,

D is a traditional method which has more overhead: need to preconfigure instance connectivity to external storage medium for writing memory And it risk altering the memory and storage artifacts in the process. Using system manager is a comparatively better way.

https://d1.awsstatic.com/events/aws-reinforce-2022/TDR401_Instance-memory-acquisition-techniques-for-effective-incident-response.pdf

upvoted 2 times

🗨️ 👤 **Raphaello** 1 year, 6 months ago

A C E

It's not possible to move an existing instance to another subnet; rather, one can associate it with a restricted SG.

upvoted 1 times

🗨️ 👤 **awssecuritynewbie** 1 year, 7 months ago

Selected Answer: ACE

A C E, for sure.

upvoted 1 times

🗨️ 👤 **lalee2** 1 year, 8 months ago

Selected Answer: ACE

Gather info -> isolate -> detach -> snapshot

upvoted 1 times

🗨️ 👤 **KR693** 1 year, 8 months ago

A, C and E

upvoted 1 times

🗨️ 👤 **Odd** 1 year, 8 months ago

ACE. B is incorrect because once a EC2 instance created, it could not be moved to other subnets

upvoted 3 times

🗨️ 👤 **kk2000** 1 year, 8 months ago

ACE makes more sense. Updating Security group (Least operational overhead) rather than moving the EC2 to different subnet which needs more steps required.

upvoted 1 times

A company has an organization in AWS Organizations. The company wants to use AWS CloudFormation StackSets in the organization to deploy various AWS design patterns into environments. These patterns consist of Amazon EC2 instances, Elastic Load Balancing (ELB) load balancers, Amazon RDS databases, and Amazon Elastic Kubernetes Service (Amazon EKS) clusters or Amazon Elastic Container Service (Amazon ECS) clusters.

Currently, the company's developers can create their own CloudFormation stacks to increase the overall speed of delivery. A centralized CI/CD pipeline in a shared services AWS account deploys each CloudFormation stack.

The company's security team has already provided requirements for each service in accordance with internal standards. If there are any resources that do not comply with the internal standards, the security team must receive notification to take appropriate action. The security team must implement a notification solution that gives developers the ability to maintain the same overall delivery speed that they currently have.

Which solution will meet these requirements in the MOST operationally efficient way?

A. Create an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the security team's email addresses to the SNS topic. Create a custom AWS Lambda function that will run the `aws cloudformation validate-template` AWS CLI command on all CloudFormation templates before the build stage in the CI/CD pipeline. Configure the CI/CD pipeline to publish a notification to the SNS topic if any issues are found.


B. Create an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the security team's email addresses to the SNS topic. Create custom rules in CloudFormation Guard for each resource configuration. In the CI/CD pipeline, before the build stage, configure a Docker image to run the `cfn-guard` command on the CloudFormation template. Configure the CI/CD pipeline to publish a notification to the SNS topic if any issues are found.

C. Create an Amazon Simple Notification Service (Amazon SNS) topic and an Amazon Simple Queue Service (Amazon SQS) queue. Subscribe the security team's email addresses to the SNS topic. Create an Amazon S3 bucket in the shared services AWS account. Include an event notification to publish to the SQS queue when new objects are added to the S3 bucket. Require the developers to put their CloudFormation templates in the S3 bucket. Launch EC2 instances that automatically scale based on the SQS queue depth. Configure the EC2 instances to use CloudFormation Guard to scan the templates and deploy the templates if there are no issues. Configure the CI/CD pipeline to publish a notification to the SNS topic if any issues are found.

D. Create a centralized CloudFormation stack set that includes a standard set of resources that the developers can deploy in each AWS account. Configure each CloudFormation template to meet the security requirements. For any new resources or configurations, update the CloudFormation template and send the template to the security team for review. When the review is completed, add the new CloudFormation stack to the repository for the developers to use.

Suggested Answer: B

Community vote distribution



bluegardentower 1 month, 2 weeks ago

Selected Answer: B

CloudFormation Guard (cfn-guard) is specifically designed to validate CloudFormation templates against policy-as-code rules, making it perfect for enforcing internal security standards before deployment.

upvoted 1 times

FlyingHawk 4 months, 3 weeks ago

Selected Answer: B

not A: Aws cloudformation validate-template only checks syntax, not security compliance. It won't enforce internal security standards.

B: cfn-guard allows defining compliance rules for CloudFormation templates.

Running cfn-guard as a Docker image in the CI/CD pipeline ensures security standards are validated before deployment.

<https://docs.aws.amazon.com/cfn-guard/latest/ug/what-is-guard.html>

upvoted 1 times

nznzwell 6 months, 1 week ago

A is not correct: the `aws cloudformation validate-template` cli command only validates the JSON/YAML, rather than compliance. Should be B.

upvoted 1 times

StarBoy01 6 months, 3 weeks ago

Selected Answer: A

NOTE: The question is about "OPERATIONAL EFFICIENCY" while maintaining speed. AWS Lambda is a good service to efficiently automate and run commands. In this case, if the team uses Cloudformation Guard, they will need to run command on each stack (which is not efficient even though it is

possible).

However, using AWS Lambda will efficiently run the command on all stacks hence the BEST answer is A.

Another point to mention here is that the Security team wants to maintain same speed in delivery time which means they are not really concerned in stopping the process if compliance is not meant. They want to "know" so they can remediate it hence they've configured the SNS for notification and when the process is complete, they can fix the compliance issues by implementing necessary changes/updates/modification.

upvoted 1 times

🗳️ 👤 **ET1857** 9 months, 1 week ago

Selected Answer: A

because of line "configure a Docker image to run the cfn-guard command on the CloudFormation template."

Option B should not be considered because its not a optimal solution

upvoted 3 times

🗳️ 👤 **hro** 1 year, 3 months ago

The answer is A. What does docker have to do with the question? Furthermore, you can validate a template with validate-template - no need for cloudformation guard.

upvoted 1 times

🗳️ 👤 **DeaconStJohn** 8 months, 1 week ago

the question is asking about code compliance and governance, cfn validate-template will check the template for valid JSON, this isnt the end goal in this scenario. Guard is an open source tool that uses both detective and preventative measures to enforce compliance on the code.

upvoted 1 times

🗳️ 👤 **DeaconStJohn** 8 months, 1 week ago

I believe docker is being used as a distractor, its no less efficient. one line of code will pull the image and run the CFN Guard command, so operational efficiency is the same however the tool itself is more aligned with expected outcome in this scenario.

upvoted 2 times

🗳️ 👤 **Raphaello** 1 year, 4 months ago

Selected Answer: B

You may use a lambda function to validate the syntax and semantics of CloudFormation templates.

But when it comes to validate to a compliance policy, CloudFormation Guard makes more sense.

Option B is correct.

upvoted 3 times

🗳️ 👤 **Raphaello** 1 year, 6 months ago

Command "aws cloudformation validate-template" checks only the syntax of cfn template.

I'd go with answer B, CloudFormation Guard, as it evaluates and validates cfn templates.

upvoted 1 times

🗳️ 👤 **Aamee** 1 year, 7 months ago

Selected Answer: B

Def. going with B as it doesn't ask to improve the delivery speed. Had that been the case then probably would have gone with option A IMO.

upvoted 1 times

🗳️ 👤 **Christina666** 1 year, 7 months ago

Selected Answer: B

BBB The ask is to send notification and not affecting current delivery speed. So only need to add a validation step to send SNS

upvoted 1 times

🗳️ 👤 **Daniel76** 1 year, 7 months ago

Selected Answer: A

Option A likely is utilizing cfn-guard as well but by Lambda/cli command, instead of creating custom rules for each resource config. More operationally efficient than B.

upvoted 1 times

🗳️ 👤 **lalee2** 1 year, 8 months ago

Selected Answer: B

CloudFormation Guard and SNS

upvoted 2 times

🗳️ 👤 **pupsik** 1 year, 8 months ago

Selected Answer: B

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/best-practices.html#:~:text=Validate%20templates%20for,non%2Dcompliant%20resources.>

upvoted 4 times

🗲️ 👤 **leoaquinods** 1 year, 8 months ago

Selected Answer: A

most operationally efficient

upvoted 1 times

🗲️ 👤 **KR693** 1 year, 8 months ago

Option B.

The aws cloudformation validate-template command is designed to check only the syntax of your template.

Using Guard, you can write policy rules to validate any JSON- or YAML-formatted structured data against, including but not limited to AWS CloudFormation templates. Guard supports the entire spectrum of end-to-end evaluation of policy checks

upvoted 3 times

🗲️ 👤 **100fold** 1 year, 8 months ago

Answer A. Most operationally efficient

upvoted 1 times

🗲️ 👤 **100fold** 1 year, 7 months ago

Answer B. My original selection

upvoted 1 times

🗲️ 👤 **Lunga778** 1 year, 8 months ago

answerer is A

upvoted 2 times

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

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

- A. AWS Site-to-Site VPN
- B. AWS Direct Connect
- C. AWS VPN CloudHub
- D. VPC peering
- E. NAT gateway

Suggested Answer: AB

Community vote distribution

AB (100%)

🗳️ **FunkyFresco** 10 months, 1 week ago

Selected Answer: AB

A and B are the correct ones.

upvoted 2 times

🗳️ **Daniel76** 1 year, 7 months ago

Selected Answer: AB

C. AWS VPN CloudHub - this is hub-and-spoke solution to connect multiple sites.

D. VPC peering - this is to connect cloud VPC to VPC.

E. NAT gateway - public or private NAT gateway? the public type is to access the internet, and a private NAT gateway is often used for communication between VPCs or between VPCs and Transit Gateway.

upvoted 4 times

🗳️ **lalee2** 1 year, 8 months ago

Selected Answer: AB

A and B provide IPsec encryption and minimized latency

upvoted 1 times

🗳️ **pupsik** 1 year, 8 months ago

Selected Answer: AB

DirectConnect for reliable connection and VPN for IPSec tunnel.

upvoted 1 times

🗳️ **KR693** 1 year, 8 months ago

A and B

upvoted 1 times

🗳️ **100fold** 1 year, 8 months ago

Selected Answer: AB

Answer AB

<https://docs.aws.amazon.com/directconnect/latest/UserGuide/encryption-in-transit.html>

upvoted 2 times

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

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

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

- A. Use the DynamoDB on-demand backup capability to create a backup plan. Configure a lifecycle policy to expire backups after 3 months.
- B. Use AWS DataSync to create a backup plan. Add a backup rule that includes a retention period of 3 months.
- C. Use AWS Backup to create a backup plan. Add a backup rule that includes a retention period of 3 months.
- D. Set the backup frequency by using a cron schedule expression. Assign each DynamoDB table to the backup plan.
- E. Set the backup frequency by using a rate schedule expression. Assign each DynamoDB table to the backup plan.

Suggested Answer: CD

Community vote distribution

CD (93%)

7%

🗳️ 👤 **BryGraVel** 3 weeks ago

Selected Answer: CD

C por retención y D por Cron (Fecha y hora)

upvoted 1 times

🗳️ 👤 **SengarAshu** 6 months ago

Selected Answer: AC

Explanation:

AWS Backup (Option C) is the most suitable solution as it allows you to automate DynamoDB backups on specific schedules (like the 15th and 25th of each month) and ensures that backup retention policies (3 months) are met.

On-demand backups (Option A) can be used as an alternative, but AWS Backup provides a more streamlined, integrated solution for managing backups and retention policies at scale.

upvoted 1 times

🗳️ 👤 **ksg97031** 8 months, 3 weeks ago

While DynamoDB on-demand backup does offer a way to manually create backups, it does not support automated backup scheduling and lifecycle management in the same flexible way that AWS Backup does. Therefore, using on-demand backup wouldn't fully meet the requirement of automating the process based on the specified schedule (15th and 25th of each month).

upvoted 2 times

🗳️ 👤 **Davidng88** 9 months, 3 weeks ago

C, & D. Note that the cron expression is not a cron job but a term describing how the AWS backup plan configuration.

upvoted 1 times

🗳️ 👤 **hb0011** 10 months ago

Selected Answer: AC

D might work but the exam always wants you to choose native solutions if possible. Avoid things like manually configuring a cron job. that's why i think the answer is AC.

upvoted 1 times

🗳️ 👤 **hb0011** 10 months ago

Edit. I'm wrong. It says combination of steps. CD is correct.

upvoted 3 times

🗳️ 👤 **FunkyFresco** 10 months, 1 week ago

Selected Answer: CD

C and D

upvoted 2 times

🗳️ 👤 **tester6667** 1 year, 3 months ago

The question asked for a SCHEDULE

C and D

<https://aws.amazon.com/blogs/database/set-up-scheduled-backups-for-amazon-dynamodb-using-aws-backup/>

upvoted 1 times

🗨️ 👤 **Raphaello** 1 year, 6 months ago

A C

What happened to DynamoDB on-demand backup?

upvoted 2 times

🗨️ 👤 **Hisayuki** 1 year, 5 months ago

DynamoDB's Backup function is only 0 days ~ 35 days. So, Backup from Outside like Lambda or AWS Backup is fine for this requirement.

upvoted 1 times

🗨️ 👤 **Raphaello** 1 year, 4 months ago

Wrong!

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/BackupRestore.html>

It explicitly says "You can use the DynamoDB on-demand backup capability to create full backups of your tables for long-term retention,"

LONG TERM RETENTION.

upvoted 5 times

🗨️ 👤 **Aamee** 1 year, 6 months ago

Selected Answer: CD

AWS Backup and Cron Job schedule comb. makes sense...

upvoted 2 times

🗨️ 👤 **Christina666** 1 year, 7 months ago

Selected Answer: CD

DynamoDB on-demand backup capability to create full backups of your tables for long-term retention and archival for regulatory compliance needs.

upvoted 3 times

🗨️ 👤 **Daniel76** 1 year, 7 months ago

Selected Answer: CD

A. the DynamoDB on-demand backup capability creates full backup while the backup plan of 3 months is by AWS Backup.

B. AWS DataSync is for data discovery and migration, not backup DynamoDB.

E. Rate schedule expression is not suitable for fixed date scheduling.

upvoted 3 times

🗨️ 👤 **imymoco** 1 year, 8 months ago

C and D

upvoted 1 times

🗨️ 👤 **lalee2** 1 year, 8 months ago

Selected Answer: CD

C and D

upvoted 1 times

🗨️ 👤 **pupsik** 1 year, 8 months ago

Selected Answer: CD

C and D

upvoted 1 times

🗨️ 👤 **KR693** 1 year, 8 months ago

C and D

upvoted 1 times

🗨️ 👤 **AgboolaKun** 1 year, 8 months ago

Selected Answer: CD

<https://aws.amazon.com/blogs/database/set-up-scheduled-backups-for-amazon-dynamodb-using-aws-backup/>

upvoted 3 times

🗨️ 👤 **PrabhuGr** 1 year, 8 months ago

A and C are the answer.

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/BackupRestore.html>

upvoted 2 times

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 AWS features should be used as much as possible. The security engineer has set up AWS Organizations with all features activated and AWS IAM Identity Center (AWS Single Sign-On) 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 AWS accounts. Assign AD Connector groups to AWS accounts and link to the IAM roles in accordance with the employees' job functions and access requirements. Instruct employees to access AWS accounts by using the AWS Directory Service user portal.
- B. Use an IAM Identity Center default directory to create users and groups for all employees that require access to AWS accounts. Assign groups to AWS accounts and link to permission sets in accordance with the employees' job functions and access requirements. Instruct employees to access AWS accounts by using the IAM Identity Center user portal.
- C. Use an IAM Identity Center default directory to create users and groups for all employees that require access to AWS accounts. Link IAM Identity Center groups to the IAM users present in all accounts to inherit existing permissions. Instruct employees to access AWS accounts by using the IAM Identity Center user portal.
- D. Use AWS Directory Service for Microsoft Active Directory to create users and groups for all employees that require access to AWS accounts. Enable AWS Management Console access in the created directory and specify IAM Identity Center as a source of information for integrated accounts and permission sets. Instruct employees to access AWS accounts by using the AWS Directory Service user portal.

Suggested Answer: B

Community vote distribution

B (100%)

 **Daniel76** Highly Voted 1 year, 1 month ago

Selected Answer: B

- A. AD Connector only provides connectivity, not managing users.
 - C. IAM users should not need to be created in all accounts - results in admin overhead. assume role instead.
 - D. Letting end users DIY access in AWS Management Console, AWS Directory Service user portal is not a good idea.
- upvoted 6 times

 **Raphaello** Most Recent 10 months, 1 week ago

Selected Answer: B

Keywords: "Native AWS features should be used as much as possible"

Therefore choose to use Identity Center's own directory, plus there is no mention to on-prem AD and hence AD connector does not make sense. For the same reason, using AWS Directory Service - Managed MS AD does not fit with native AWS feature.

Option B is the right answer.


upvoted 2 times

 **Ialee2** 1 year, 2 months ago

Selected Answer: B

<https://aws.amazon.com/ko/iam/identity-center/faqs/>


upvoted 2 times

 **pupsik** 1 year, 2 months ago

Selected Answer: B

Normally we would use AD Connector to connect to on-premises AD. But option A doesn't come close to that. Hence option B.

upvoted 4 times

 **KR693** 1 year, 2 months ago

Option B

upvoted 3 times

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

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

Suggested Answer: C

Community vote distribution

C (70%)

D (30%)

 **Daniel76** Highly Voted 9 months, 1 week ago

Selected Answer: C

Let Guardduty detections be sent to Security Hub as findings is a simple and elegant way.

<https://docs.aws.amazon.com/guardduty/latest/ug/securityhub-integration.html>

Use eventbridge to respond by invoke Lambda. Amazon Kinesis data stream not needed.

<https://docs.aws.amazon.com/securityhub/latest/userguide/securityhub-cloudwatch-events.html>

Suggest to only block specific port 389 against thse suspicious EC2 instance instead of isolate it in a security group, to minimize the impact while it has not been verified as a confirmed attack.


upvoted 11 times

 **FLucas** Most Recent 2 months, 3 weeks ago

Selected Answer: C

C, because as security groups (in the option D) is stateful, changing it will not affect if the RDP was already successful.

upvoted 2 times

 **PrabhuKK** 3 months, 3 weeks ago

Selected Answer: D

Why is D the best option?

AWS Security Hub can aggregate security findings from GuardDuty and other AWS services.

EventBridge (or a similar event-driven system) can be used to trigger specific actions in response to GuardDuty findings.

Lambda can process the event and implement automatic remediation actions.

In this solution, the Lambda function would modify the security group of the suspicious EC2 instance, effectively blocking all incoming and outgoing traffic by applying a security group that does not allow connections. This is an effective way to isolate the instance until further investigation.

upvoted 1 times

 **cyber_m0nk** 4 months, 2 weeks ago

Selected Answer: C

Option C is the best answer. The Network Firewall has IPS/IDS and works at layers 3 and 4 providing the best protection against the RDP (3389) attack in this situation for the EC2. The WAF is optimal for protection against threats targeting HTTP/HTTPS (80, 443) web traffic.

D is incorrect Security Groups do not block outbound traffic.

upvoted 1 times

🗨️ **Almo89** 9 months, 1 week ago

Selected Answer: C

No A: Kinesis is for a stream of data. Guardduty will report (raise alarm) but no need of apache flink

Not B: WAF is for a web application (80, 443), the attack will be initiated from a EC2 (internal company machine) it doesn't specify the target location local or extern, but it will be RDP

Not D- SG will not block outbound traffic (stateful). SG will not be enough, the ec2 initiates the attack

C is the simplest and more correct as FW will block any traffic from/to ec2

upvoted 3 times

🗨️ **DLG_85** 1 year ago

Option B. We have to consider that Security team has to be noticed. There is nothing about that on the other answers.

upvoted 1 times

🗨️ **sabmel** 9 months, 4 weeks ago

its not required as per question, so not B.

upvoted 1 times

🗨️ **cumzle_com** 1 year ago

Selected Answer: D

While Option C is technically feasible and provides robust network-level protection using AWS Network Firewall, it is more complex and might be an overkill for the specific task of quickly isolating individual EC2 instances. Option D offers a more direct and simpler approach by replacing the security group of the suspicious instance, which is generally easier to manage and quicker to implement in the context of isolating instances based on GuardDuty findings.

Therefore, while Option C can meet the requirements, Option D is more appropriate and efficient for the specific task of blocking communication from suspicious instances quickly and effectively.

upvoted 1 times

🗨️ **cyber_m0nk** 4 months, 2 weeks ago

Security groups do not have IPS or IDS capabilities whereas the Network Firewall does. Option D is ineffective Security Groups do not have DENY statements for inbound or outbound traffic. Replacing the SG has no effect, it would require allowing each specific instance (IP addresses) you deem as legitimate.

upvoted 1 times

🗨️ **sabmel** 9 months, 4 weeks ago

In option D, all the connections will be blocked due to change in SG, not required so not D.

upvoted 1 times

🗨️ **Raphaello** 1 year, 4 months ago

Selected Answer: C

1. No point to use Kinesis Data Stream/analytics/Apache flink to stream and process event.
2. Neither WAF nor NACL is an effective solution to the mentioned case
3. GuardDuty findings can be sent directly to Amazon EventBridge to trigger action, but deploying SecurityHub is not entirely wrong.
4. AWS Network Firewall is better suited to block suspicious instances.

Option C is the correct answer.

upvoted 1 times

🗨️ **awssecuritynewbie** 1 year, 5 months ago

Selected Answer: C

I would go with C, as option D will block any connection to the Ec2 machine, which is not what you want, and security groups are easier and at the endpoint level.

upvoted 1 times

🗨️ **mynickc** 1 year, 5 months ago

Selected Answer: C

Here is some basics: WAF protects the port 443 / 80. RDP is different port and nothing to do with Layer 7 nor WAF

upvoted 4 times

🗨️ **happy34** 1 year, 5 months ago

D is the answer. we need identify the best method - tech and cost. implied. WAF is layer 7 prevention . FW is layer 3 - 7. WEB ACL can prevent layer 7.

RDP is mostly Layer 7. password guessing etc

<https://repost.aws/knowledge-center/waf-prevent-brute-force-attacks>

upvoted 1 times

🗨️ 👤 **shuaaaa** 1 year, 1 month ago

RDP is L3 (IP) + L4 (Port 3389).

it is NOT L7.

upvoted 2 times

🗨️ 👤 **brpjp** 1 year, 6 months ago

Hello, correct my understanding agree with answer C.

upvoted 1 times

🗨️ 👤 **brpjp** 1 year, 6 months ago

When GuardDuty is there, do not understand what is requirement to integrate Security Hub.

upvoted 2 times

🗨️ 👤 **brpjp** 1 year, 6 months ago

Answer B correct. Requested first scenario of RDP brute force attack. Neither NACL, Network Firewall, and Security Group support to block, only WAF help to block traffic based on pattern.

upvoted 2 times

🗨️ 👤 **cumzle_com** 1 year ago

Waf only for http traffic :)

upvoted 1 times

🗨️ 👤 **3633f8f** 1 year, 6 months ago

Selected Answer: C

C is the answer

upvoted 1 times

🗨️ 👤 **3633f8f** 1 year, 6 months ago

Selected Answer: D

To isolate there is nothing more powerful than an ACL at subnet level, which immediately denies traffic in any direction. Wishing to automate, there is no choice to use ACL, as you do not know the exact IP of the source is attacking, thus, you do apply security group restriction.

The need of Kinesis Data Streams is to process real-time events while happening.

A firewall you do not usually automate at this has complex features needs to be set via IaC or console.

upvoted 2 times

🗨️ 👤 **3633f8f** 1 year, 6 months ago

Correcting as RDP handles directly in layer 3. C

upvoted 1 times

🗨️ 👤 **Raphaello** 1 year, 6 months ago

Selected Answer: C

C

SecurityGroup is a simpler way of isolating a suspicious instance, unlike Network Firewall that is a paid service.

EventBridge is needed to relay events to Kinesis Data Stream. At that point, what is the need to Kinesis Data Stream? Lambda function could be invoked directly from EventBridge.

For that, I'd go with C.

upvoted 2 times

A company has an AWS account that hosts a production application. The company receives an email notification that Amazon GuardDuty has detected an Impact:IAMUser/AnomalousBehavior finding in the account. A security engineer needs to run the investigation playbook for this security incident and must collect and analyze the information without affecting the application.

Which solution will meet these requirements MOST quickly?

- A. Log in to the AWS account by using read-only credentials. Review the GuardDuty finding for details about the IAM credentials that were used. Use the IAM console to add a DenyAll policy to the IAM principal.
- B. Log in to the AWS account by using read-only credentials. Review the GuardDuty finding to determine which API calls initiated the finding. Use Amazon Detective to review the API calls in context.
- C. Log in to the AWS account by using administrator credentials. Review the GuardDuty finding for details about the IAM credentials that were used. Use the IAM console to add a DenyAll policy to the IAM principal.
- D. Log in to the AWS account by using read-only credentials. Review the GuardDuty finding to determine which API calls initiated the finding. Use AWS CloudTrail Insights and AWS CloudTrail Lake to review the API calls in context.

Suggested Answer: B

Community vote distribution

B (100%)

 **pupsik** Highly Voted 1 year, 8 months ago

Selected Answer: B

<https://aws.amazon.com/blogs/security/how-you-can-use-amazon-guardduty-to-detect-suspicious-activity-within-your-aws-account/#:~:text=Start%20an%20investigation%20with%20Amazon%20Detective>


upvoted 9 times

 **FLucas** Most Recent 2 months, 3 weeks ago

Selected Answer: B

Understand events from guardduty better always detective

upvoted 1 times

 **Davidng88** 9 months, 3 weeks ago

Selected Answer: B

Read-only credentials ensure that the investigation does not inadvertently affect the production application.

Reviewing the GuardDuty finding helps identify the specific IAM credentials and API calls involved.

Amazon Detective provides a comprehensive analysis of the API calls in context, allowing for a deeper investigation into the anomalous behavior123.

This approach balances thorough investigation with minimal impact on the production environment

upvoted 1 times

 **Raphaello** 1 year, 4 months ago

Selected Answer: B

Using CloudTrail (Insights & Lake) is not entirely wrong for the the aforementioned case, however, since the ask is to analyze the events "QUICKLY", I think Detective provides a good integration with GuardDuty to correlate data and analyze them.

I would go with B only for this.."quickly"!

upvoted 3 times

 **Aamee** 1 year, 6 months ago

Selected Answer: B

Option B since Detective is integrated with GuardDuty by native... contrast to option D where Insights and Lake are NA to GuardDuty..

upvoted 2 times

 **Daniel76** 1 year, 7 months ago

Selected Answer: B

A. Read-only login should not allow user to add DenyAllPolicy.

C. Add DenyAllPolicy to the principal is very intrusive intervention.

D. Use AWS CloudTrail Insights and AWS CloudTrail Lake are not integrated with GuardDuty (as opposed to AWS Detective) hence it might lack correlationship

upvoted 1 times

🗨️ 👤 **lalee2** 1 year, 8 months ago

Selected Answer: B

'Qs says collect and analyze the info' read-only credential is enough. Detective provides API activities.

upvoted 2 times

🗨️ 👤 **KR693** 1 year, 8 months ago

Option B

upvoted 1 times

🗨️ 👤 **Lunga778** 1 year, 8 months ago

correct answer is be

<https://aws.amazon.com/blogs/aws/new-aws-cloudtrail-lake-supports-ingesting-activity-events-from-non-aws-sources/>

<https://aws.amazon.com/about-aws/whats-new/2019/11/aws-cloudtrail-announces-cloudtrail-insights/>

upvoted 1 times

🗨️ 👤 **Lunga778** 1 year, 8 months ago

i mean is D

upvoted 1 times

🗨️ 👤 **100fold** 1 year, 8 months ago

Selected Answer: B

Answer B

upvoted 1 times

🗨️ 👤 **AgboolaKun** 1 year, 8 months ago

Selected Answer: B

<https://aws.amazon.com/blogs/security/how-you-can-use-amazon-guardduty-to-detect-suspicious-activity-within-your-aws-account/>

upvoted 2 times

Company A has an AWS account that is named Account A. Company A recently acquired Company B, which has an AWS account that is named Account B. Company B stores its files in an Amazon S3 bucket. The administrators need to give a user from Account A full access to the S3 bucket in Account B.

After the administrators adjust the IAM permissions for the user in Account A to access the S3 bucket in Account B, the user still cannot access any files in the S3 bucket.

Which solution will resolve this issue?

- A. In Account B, create a bucket ACL to allow the user from Account A to access the S3 bucket in Account B.
- B. In Account B, create an object ACL to allow the user from Account A to access all the objects in the S3 bucket in Account B.
- C. In Account B, create a bucket policy to allow the user from Account A to access the S3 bucket in Account B.
- D. In Account B, create a user policy to allow the user from Account A to access the S3 bucket in Account B.

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **FunkyFresco** 10 months, 1 week ago

Selected Answer: C

Answer C is the right option.

upvoted 1 times

🗳️ 👤 **liuyomz** 1 year, 2 months ago

Bucket ACL are deprecated, so probably a hint to ignore anything with it

upvoted 2 times

🗳️ 👤 **Raphaello** 1 year, 4 months ago

Selected Answer: C

Bucket POLICY.

upvoted 1 times

🗳️ 👤 **trashbox** 1 year, 6 months ago

Exam on 2023-12-18

upvoted 1 times

🗳️ 👤 **Raphaello** 1 year, 6 months ago

Selected Answer: C

S3 bucket policy should be edited to allow cross-account access.

upvoted 2 times

🗳️ 👤 **raj0011** 1 year, 7 months ago

why they correct answer A

upvoted 1 times

🗳️ 👤 **liuyomz** 1 year, 2 months ago

you can always ignore the site's "correct answer", go for the Most Voted and open the discussions

upvoted 1 times

🗳️ 👤 **Daniel76** 1 year, 7 months ago

Selected Answer: C

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/example-walkthroughs-managing-access-example2.html>

upvoted 2 times

🗳️ 👤 **[Removed]** 1 year, 8 months ago

Selected Answer: C

It's C

upvoted 1 times

🗳️ 👤 **lalee2** 1 year, 8 months ago

it is cross-account, requires S3 bucket policy update

upvoted 3 times

🗨️ 👤 **pupsik** 1 year, 8 months ago

Selected Answer: C

For cross-account access, got to update S3 bucket policy to allow principal from other account to access it.

upvoted 4 times

🗨️ 👤 **KR693** 1 year, 8 months ago

Option C

upvoted 2 times

🗨️ 👤 **100fold** 1 year, 8 months ago

Selected Answer: C

Answer C

upvoted 2 times

🗨️ 👤 **aragon_saa** 1 year, 9 months ago

<https://www.examttopics.com/discussions/amazon/view/88380-exam-aws-certified-security-specialty-topic-1-question-423/>

upvoted 4 times

A company wants to receive an email notification about critical findings in AWS Security Hub. The company does not have an existing architecture that supports this functionality.

Which solution will meet the requirement?

- A. Create an AWS Lambda function to identify critical Security Hub findings. Create an Amazon Simple Notification Service (Amazon SNS) topic as the target of the Lambda function. Subscribe an email endpoint to the SNS topic to receive published messages.
- B. Create an Amazon Kinesis Data Firehose delivery stream. Integrate the delivery stream with Amazon EventBridge. Create an EventBridge rule that has a filter to detect critical Security Hub findings. Configure the delivery stream to send the findings to an email address.
- C. Create an Amazon EventBridge rule to detect critical Security Hub findings. Create an Amazon Simple Notification Service (Amazon SNS) topic as the target of the EventBridge rule. Subscribe an email endpoint to the SNS topic to receive published messages.
- D. Create an Amazon EventBridge rule to detect critical Security Hub findings. Create an Amazon Simple Email Service (Amazon SES) topic as the target of the EventBridge rule. Use the Amazon SES API to format the message. Choose an email address to be the recipient of the message.

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ **c6ed25a** 3 months, 1 week ago

Selected Answer: C

SNS and eventbridge is the hint

upvoted 1 times

🗳️ **FunkyFresco** 10 months, 1 week ago

Selected Answer: C

C option.

upvoted 2 times

🗳️ **habros** 10 months, 3 weeks ago

Selected Answer: C

C is soooo much quicker... no need for complicated resources if all the administrator wants to find out who is responsible and to alert them

upvoted 1 times

🗳️ **hro** 1 year, 3 months ago

D - Eventbridge is essentially like SNS in this example - 'Amazon Simple Notification Service (SNS) and Amazon EventBridge are both AWS services that help with data processing and messaging [of events.]'

So then working directly with SNS is more efficient. And YES Eventbridge CAN work with SNS with just the same amount of effort but one less step of redundancy.

upvoted 2 times

🗳️ **Just_Ninja** 1 year, 1 month ago

It's not well architected. "C is better"

The Way SecurityHub >>EventBridge >>SNS is better, option D suggests using Amazon SES directly, which is not the most straightforward or typical approach for this use case. EventBridge and SNS together offer a more integrated and native solution for sending notifications from AWS Security Hub.

upvoted 2 times

🗳️ **hro** 1 year, 3 months ago

D. Indeed you can 'To subscribe an email address to an Amazon SNS topic using the AWS Management Console'

upvoted 1 times

🗳️ **helloworldabc** 9 months, 1 week ago

just C

upvoted 2 times

🗳️ **Raphaello** 1 year, 4 months ago

SecurityHub >>EventBridge >>SNS.

```
{
  "source": ["aws.securityhub"],
  "detail-type": ["Security Hub Findings - Imported"],
  "detail": {
    "findings": {"Severity": {"Label": ["Critical"]}}
  }
}
```

upvoted 3 times

 **trashbox** 1 year, 6 months ago

Exam on 2023-12-18

upvoted 1 times

 **3633f8f** 1 year, 6 months ago

Selected Answer: C

Marked D is definitely wrong

upvoted 1 times

 **Raphaello** 1 year, 6 months ago

C

SES? Seriously?

upvoted 1 times


 **Daniel76** 1 year, 7 months ago

Selected Answer: C

To filter for critical only finding:

<https://docs.aws.amazon.com/securityhub/latest/userguide/securityhub-cwe-all-findings.html>

upvoted 2 times

 **YR4591** 1 year, 7 months ago

Selected Answer: C

C is right

securityhub > eventbridge > sns

upvoted 1 times

 **Karamen** 1 year, 8 months ago

C is right. confirmed


upvoted 1 times

 **[Removed]** 1 year, 8 months ago

Selected Answer: C

C is right


upvoted 1 times

 **lalee2** 1 year, 8 months ago

Selected Answer: C

EventBridge -> SNS is right


upvoted 1 times

 **pupsik** 1 year, 8 months ago

Selected Answer: C


EventBridge Rule -> SNS -> Email delivery

upvoted 1 times

 **KR693** 1 year, 8 months ago

Option C

upvoted 1 times

 **jabilrn** 1 year, 8 months ago

I think C is right.

I dont believe SES can be the target for Eventbridge

upvoted 2 times

An international company has established a new business entity in South Korea. The company also has established a new AWS account to contain the workload for the South Korean region. The company has set up the workload in the new account in the ap-northeast-2 Region. The workload consists of three Auto Scaling groups of Amazon EC2 instances. All workloads that operate in this Region must keep system logs and application logs for 7 years.

A security engineer must implement a solution to ensure that no logging data is lost for each instance during scaling activities. The solution also must keep the logs for only the required period of 7 years.

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

- A. Ensure that the Amazon CloudWatch agent is installed on all the EC2 instances that the Auto Scaling groups launch. Generate a CloudWatch agent configuration file to forward the required logs to Amazon CloudWatch Logs.
- B. Set the log retention for desired log groups to 7 years.
- C. Attach an IAM role to the launch configuration or launch template that the Auto Scaling groups use. Configure the role to provide the necessary permissions to forward logs to Amazon CloudWatch Logs.
- D. Attach an IAM role to the launch configuration or launch template that the Auto Scaling groups use. Configure the role to provide the necessary permissions to forward logs to Amazon S3.
- E. Ensure that a log forwarding application is installed on all the EC2 instances that the Auto Scaling groups launch. Configure the log forwarding application to periodically bundle the logs and forward the logs to Amazon S3.
- F. Configure an Amazon S3 Lifecycle policy on the target S3 bucket to expire objects after 7 years.

Suggested Answer: ABC

Community vote distribution

ABC (82%)

ACF (18%)


  **FunkyFresco** 10 months, 1 week ago

Selected Answer: ABC

ABC adjust more to the question. No need of external applications to send logs.
upvoted 2 times

  **nischal7777** 10 months, 2 weeks ago

ABC IS CORRECT OPTION
upvoted 1 times

  **liuyomz** 1 year, 2 months ago

Selected Answer: ABC

ABC is the most straightfoward and elegant solution here
upvoted 1 times

  **Raphaello** 1 year, 4 months ago


Selected Answer: ABC

ABC.
EC2 (with a role allowing sending events) >> CloudWatch agent >> CloudWatch Logs >> CloudWatch Logs retention period
upvoted 3 times



  **Daniel76** 1 year, 7 months ago

Selected Answer: ACF

As the log data for 7 years will be expensive, use AWS S3 Lifecycle Management to transfer data to lower cost storage class will be more cost effective solution.
<https://medium.com/avmconsulting-blog/aws-s3-lifecycle-management-1ed2f67c3b73>
upvoted 3 times

  **mav3r1ck** 1 year, 6 months ago

I would go for ACF if the asked is "COST-EFFECTIVE" solution. But leaning to ABF, as Cloudwatch logs support up to 10yrs of retention as well. Feel free to disagree if you think I'm wrong. <https://docs.aws.amazon.com/managedservices/latest/userguide/log-customize-retention.html>
upvoted 4 times

  **Daniel76** 1 year, 6 months ago

Agree it should not be ACF but ABC. ACF does not explain how cloudwatch log ends up in s3. (seems that it requires a lambda function to automate)

In fact, if S3 is the chosen path then it can only be DEF for consistency. But this combination assumes that s3 bucket policy has been configured and the log forwarding application configured can reliably send all logs data without losing any.



upvoted 4 times

  **Karamen** 1 year, 8 months ago

ABC

there isn't good option to forwarding log from EC2 to S3 bucket.



upvoted 4 times

  **lalee2** 1 year, 8 months ago

Selected Answer: ABC

CloudWatch agent -> CloudWatch Logs, IAM role to launch template -> CloudWatch Logs



upvoted 1 times

  **pupsik** 1 year, 8 months ago

Selected Answer: ABC



ABC it is.

upvoted 2 times

  **KR693** 1 year, 8 months ago

A, B and C

upvoted 1 times

  **100fold** 1 year, 8 months ago

Selected Answer: ABC

Answer ABC

upvoted 1 times

  **aragon_saa** 1 year, 8 months ago

<https://www.examttopics.com/discussions/amazon/view/89514-exam-aws-certified-security-specialty-topic-1-question-451/>

upvoted 2 times

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

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

Which combination of conditions must the security engineer add to the IAM policy to meet these requirements? (Choose two.)

- A. "Bool": {"aws:MultiFactorAuthPresent": "true"}
- B. "Bool": {"aws:MultiFactorAuthPresent": "false"}
- C. "NumericLessThan": {"aws:MultiFactorAuthAge": "7200"}
- D. "NumericGreaterThan": {"aws:MultiFactorAuthAge": "7200"}
- E. "NumericLessThan": {"MaxSessionDuration": "7200"}

Suggested Answer: AC

Community vote distribution

AC (100%)

 **kk2000** Highly Voted 1 year, 8 months ago

Correct Answer is AC

upvoted 8 times

 **molerowan** Most Recent 3 months, 3 weeks ago

Selected Answer: AC

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:DescribeInstances",
        "ec2:StopInstances",
        "ec2:TerminateInstances"
      ],
      "Resource": "*",
      "Condition": {
        "Bool": {"aws:MultiFactorAuthPresent": "true"},
        "NumericLessThan": {"aws:MultiFactorAuthAge": 7200}
      }
    }
  ]
}
```

```
]
}
```

upvoted 1 times

🗨️ **youonebe** 5 months, 4 weeks ago

Selected Answer: AC

AC, not AE.

As mentioned in https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_mfa_configure-api-require.html

upvoted 1 times

🗨️ **102fe8d** 7 months, 1 week ago

its A & E

MFAAge checks how much time has passed since the user's last MFA usage

SessionAge makes sure the MFA session is NOT more than the selected timelimit.

upvoted 1 times

🗨️ **Gentlesoul** 10 months, 2 weeks ago

The correct answer is AC, the MFA age must be less than or equal to 7200

upvoted 1 times

🗨️ **nischal77777** 10 months, 2 weeks ago

Selected Answer: AC

Correct Answer is AC

upvoted 2 times

🗨️ **PareshBPatel** 1 year, 4 months ago

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

This condition ensures that the action is allowed only if the MFA session age is less than 7200 seconds (2 hours), meaning it enforces the requirement that each MFA session remains valid for only 2 hours. This is a correct choice as it directly addresses the session validity requirement.

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

This condition would allow the action only if the MFA session age is greater than 7200 seconds, which is contrary to the requirement. Therefore, this option is incorrect.

upvoted 2 times

🗨️ **Raphaello** 1 year, 4 months ago

Selected Answer: AC

AC ofc

upvoted 1 times

🗨️ **mtzanida** 1 year, 5 months ago

Selected Answer: AC

A and C

upvoted 1 times

🗨️ **Raphaello** 1 year, 6 months ago

AC

Action is ALLOW..as long as the auth. age is LESS 7200 seconds.

upvoted 1 times

🗨️ **Daniel76** 1 year, 7 months ago

Selected Answer: AC

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_mfa_configure-api-require.html

upvoted 2 times

🗨️ **kejam** 1 year, 7 months ago

A and C

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_mfa_configure-api-require.html#MFAProtectedAPI-overview



upvoted 3 times

🗨️ **lalee2** 1 year, 8 months ago

Selected Answer: AC

A and C



upvoted 1 times

  **denied** 1 year, 8 months ago

Selected Answer: AC

A and C

upvoted 2 times

  **KR693** 1 year, 8 months ago

A and C

upvoted 2 times

A company uses AWS Organizations and has production workloads across multiple AWS accounts. A security engineer needs to design a solution that will proactively monitor for suspicious behavior across all the accounts that contain production workloads.


The solution must automate remediation of incidents across the production accounts. The solution also must publish a notification to an Amazon Simple Notification Service (Amazon SNS) topic when a critical security finding is detected. In addition, the solution must send all security incident logs to a dedicated account.

Which solution will meet these requirements?

- A. Activate Amazon GuardDuty in each production account. In a dedicated logging account, aggregate all GuardDuty logs from each production account. Remediate incidents by configuring GuardDuty to directly invoke an AWS Lambda function. Configure the Lambda function to also publish notifications to the SNS topic.
- B. Activate AWS Security Hub in each production account. In a dedicated logging account, aggregate all Security Hub findings from each production account. Remediate incidents by using AWS Config and AWS Systems Manager. Configure Systems Manager to also publish notifications to the SNS topic.
- C. Activate Amazon GuardDuty in each production account. In a dedicated logging account, aggregate all GuardDuty logs from each production account. Remediate incidents by using Amazon EventBridge to invoke a custom AWS Lambda function from the GuardDuty findings. Configure the Lambda function to also publish notifications to the SNS topic.
- D. Activate AWS Security Hub in each production account. In a dedicated logging account, aggregate all Security Hub findings from each production account. Remediate incidents by using Amazon EventBridge to invoke a custom AWS Lambda function from the Security Hub findings. Configure the Lambda function to also publish notifications to the SNS topic.

Suggested Answer: C

Community vote distribution



Sat897 5 days, 2 hours ago

Selected Answer: C

Proactively Monitoring - Guard Duty

Automate remediation - Eventbridge (automation) and invoke Lambda policy for remediation. and SNS topic.

Security Hub - this will be used for centralized tool to gather all findings, but this is not required to current use case. So option C
upvoted 1 times

bluegardentower 1 month, 1 week ago

Selected Answer: D

- A GuardDuty findings can't directly invoke a Lambda function. Need EventBridge
 - B Systems Manager is not event-driven or as immediate as using EventBridge + Lambda
 - C GuardDuty alone does not aggregate findings like Security Hub
- upvoted 1 times

Kaps443 2 months, 2 weeks ago

Selected Answer: C

Proactively monitor for suspicious behavior ✓ (This is GuardDuty's job – not Security Hub's)
Option D is incorrect.
upvoted 2 times

Kaps443 2 months, 2 weeks ago

Selected Answer: D

Option: D
uses Security Hub (broad security visibility) + EventBridge (event-driven automation) + Lambda + SNS (custom remediation + notification) and aggregates findings in a central account, matching all requirements efficiently.

Why Not the Other Options?

A – GuardDuty can't invoke a Lambda directly; it needs EventBridge.

B – While Security Hub supports findings and aggregation, AWS Config + SSM is more about configuration compliance than real-time incident

remediation.

C – Similar to D but uses GuardDuty only, which covers only specific threat detection. Security Hub provides broader visibility by consolidating findings from multiple tools.

upvoted 2 times

  **DeaconStJohn** 8 months ago

Selected Answer: D

This question has had me going back and forth for days. I use Stephane Maareks course and AWS docs as my sources.

The bits I am getting tripped up by:

1. The question is not specifically about guard duty. For Security Hub to be enabled we can ascertain that AWS config has been enabled. Config rules could trigger findings. Security Hub can remediate this.

2. Option C says Guard Duty logs. Is this correct? Other options explicitly state "findings." Guard duty monitors various logs and pushes findings to Eventbridge.

In Stephane Maarek's course, he has this exact architecture for Security hub remediation.

upvoted 1 times

  **bukkanni** 10 months ago

Security Hub serves as a central repository to store security findings from various services including Guard Duty. Option C limits security findings ONLY to Guard Duty. The requirement is to alert for any security finding and not just those from Guard Duty. D is more appropriate in my opinion.

upvoted 3 times

  **hb0011** 10 months ago

I think the writer of the test would say you can't assume guardduty is enabled in D and security hub can't do it by itself. It would rely on guardduty. But I really don't know what they're looking for here.

upvoted 2 times

  **FunkyFresco** 10 months, 1 week ago

Selected Answer: C

Option C fits.

upvoted 1 times

  **Almo89** 12 months ago

Selected Answer: D

Difficult one, D could also be correct.



Securityhub can trigger remediation and sns from different sources (cloudwatch, guardduty)

Securityhub depends of guardduty. But guardduty depends of cloudwatch as well.

Securityhub has a predefined remediation based on best practice.

C and D are correct, question doesn't specify/differentiates

upvoted 1 times

  **hro** 1 year, 3 months ago

B

C never addresses Remediation and why would you Configure the Lambda function to also publish notifications to the SNS topic when Security Hub works with AWS Config and AWS Systems Manager and push notifications?

The answer is B

upvoted 1 times

  **Josh1217** 1 year, 3 months ago

GuardDuty can't directly invoke Lambda. Option C addresses remediation. Read the option properly.

upvoted 2 times

  **Raphaello** 1 year, 6 months ago

Best answer is C

One would not need SecurityHub to launch a response to GuardDuty finding.

SecurityHub is security posture management tool, but without it GuardDuty can still responds to findings.

https://docs.aws.amazon.com/guardduty/latest/ug/guardduty_findings_cloudwatch.html

upvoted 1 times

🗲️ 👤 **Daniel76** 1 year, 7 months ago

Selected Answer: C

Security Hub by itself does not detect suspicious activity, but GuardDuty.
Eventbridge rule is required to trigger remediation actions and SNS topic.

upvoted 4 times

🗲️ 👤 **[Removed]** 1 year, 8 months ago

Selected Answer: C

SecurityHub checks posture. GuardDuty monitors for malicious activity.

upvoted 2 times

🗲️ 👤 **[Removed]** 1 year, 8 months ago

It's C.

SecurityHub checks posture. GuardDuty monitors for malicious activity.

upvoted 1 times

🗲️ 👤 **lalee2** 1 year, 8 months ago

Selected Answer: C

Option C responds to all requirements; automate remediation, notification via SNS, send logs to a dedicated account

upvoted 1 times

🗲️ 👤 **bhui** 1 year, 8 months ago

I would say it is C as Guardduty must be turned on even for the security hub options. Also you can aggregate GuardDuty Findings and trigger Events.

<https://aws.amazon.com/blogs/security/how-to-manage-amazon-guardduty-security-findings-across-multiple-accounts/>

<https://repost.aws/knowledge-center/guardduty-eventbridge-sns-rule>

upvoted 3 times

🗲️ 👤 **pupsik** 1 year, 8 months ago

Selected Answer: C

Agree, it is C

upvoted 1 times

🗲️ 👤 **KR693** 1 year, 8 months ago

Option C

upvoted 1 times


A company is designing a multi-account structure for its development teams. The company is using AWS Organizations and AWS IAM Identity Center (AWS Single Sign-On). The company must implement a solution so that the development teams can use only specific AWS Regions and so that each AWS account allows access to only specific AWS services.

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

- A. Use IAM Identity Center to set up service-linked roles with IAM policy statements that include the Condition, Resource, and NotAction elements to allow access to only the Regions and services that are needed.
- B. Deactivate AWS Security Token Service (AWS STS) in Regions that the developers are not allowed to use.
- C. Create SCPs that include the Condition, Resource, and NotAction elements to allow access to only the Regions and services that are needed.
- D. For each AWS account, create tailored identity-based policies for IAM Identity Center. Use statements that include the Condition, Resource, and NotAction elements to allow access to only the Regions and services that are needed.

Suggested Answer: C

Community vote distribution



🗳️ 👤 **kkravets** 8 months, 1 week ago

Selected Answer: A

SCP is going to affect all, not only dev team
upvoted 1 times

🗳️ 👤 **nischal7777** 10 months, 2 weeks ago

Selected Answer: C

scp IS CORRECT
upvoted 1 times

🗳️ 👤 **komik_101** 11 months, 4 weeks ago

I would like to say A. :D Last week I did this, software team came to me, and they want to access another account S3 bucket permission.(Put, get, delete). and I went to the IAM Identity Center. created a group(put in the users group), and I did permission sets and they accessed other account s3 bucket.

SCP is huge topic. SCP very critical . if you doing something, will affect all accounts
upvoted 1 times

🗳️ 👤 **awssecuritynewbie** 1 year, 4 months ago

you cannot use "NOTACTION" with SCP though? Anyone can help?
upvoted 1 times

🗳️ 👤 **tester6667** 1 year, 3 months ago

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_syntax.html
You can see not action listed
upvoted 1 times

🗳️ 👤 **Raphaello** 1 year, 4 months ago

Selected Answer: C

SCP to allow certain services in certain regions for specific accounts.
upvoted 1 times

🗳️ 👤 **Raphaello** 1 year, 4 months ago

As explained here https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_syntax.html#scp-elements-table "Condition", "Resource", and "NotAction" elements can only be used with "Deny" effect, but answer C says "to allow access to only the Regions and services that are needed" as the ultimate outcome, not by the meaning with "Allow" effect.

It tries to trick you into thinking "those elements cannot be used with "Allow", then not C" !

Still believe C is the best answer here.

upvoted 1 times

🗉 👤 **NoCrapEva** 1 year, 4 months ago

Selected Answer: C

SCP is the GOTO solution for multiple accounts in AWS Organisations.

upvoted 1 times

🗉 👤 **habros** 1 year, 4 months ago

Selected Answer: C

C. If AWS organizations is enabled, why not take advantage of region deny feature? SCP is the actual mechanism to enforce this rule!

upvoted 2 times

🗉 👤 **mynickc** 1 year, 5 months ago

Selected Answer: A

C is wrong becoz notation, resource & condition can support deny only.

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_syntax.html

upvoted 1 times

🗉 👤 **Sab31** 1 year, 5 months ago

C seems a good option but can someone share if SCPs can have "NotAction" element?

upvoted 1 times

🗉 👤 **Raphaello** 1 year, 6 months ago

Correct answer is C.

SCP to control which organization node can operate on which region(s).

upvoted 1 times

🗉 👤 **Daniel76** 1 year, 7 months ago

Selected Answer: C

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_examples_general.html#example-scp-deny-region

upvoted 2 times

🗉 👤 **lalee2** 1 year, 8 months ago

Selected Answer: C

Under Organization SCP is the least operational overhead.

upvoted 1 times

🗉 👤 **KR693** 1 year, 8 months ago

Option C

upvoted 1 times

🗉 👤 **Sumi81** 1 year, 8 months ago

C is right

upvoted 1 times

🗉 👤 **100fold** 1 year, 8 months ago

Selected Answer: C

Agree answer C

upvoted 1 times

🗉 👤 **aragon_saa** 1 year, 8 months ago

<https://www.examttopics.com/discussions/amazon/view/88434-exam-aws-certified-security-specialty-topic-1-question-431/>

upvoted 2 times

A company is developing an ecommerce application. The application uses Amazon EC2 instances and an Amazon RDS MySQL database. For compliance reasons, data must be secured in transit and at rest. The company needs a solution that minimizes operational overhead and minimizes cost.

Which solution meets these requirements?

- A. Use TLS certificates from AWS Certificate Manager (ACM) with an Application Load Balancer. Deploy self-signed certificates on the EC2 instances. Ensure that the database client software uses a TLS connection to Amazon RDS. Enable encryption of the RDS DB instance. Enable encryption on the Amazon Elastic Block Store (Amazon EBS) volumes that support the EC2 instances.
- B. Use TLS certificates from a third-party vendor with an Application Load Balancer. Install the same certificates on the EC2 instances. Ensure that the database client software uses a TLS connection to Amazon RDS. Use AWS Secrets Manager for client-side encryption of application data.
- C. Use AWS CloudHSM to generate TLS certificates for the EC2 instances. Install the TLS certificates on the EC2 instances. Ensure that the database client software uses a TLS connection to Amazon RDS. Use the encryption keys from CloudHSM for client-side encryption of application data.
- D. Use Amazon CloudFront with AWS WAF. Send HTTP connections to the origin EC2 instances. Ensure that the database client software uses a TLS connection to Amazon RDS. Use AWS Key Management Service (AWS KMS) for client-side encryption of application data before the data is stored in the RDS database.

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ 👤 **FunkyFresco** 10 months, 1 week ago

Selected Answer: A

Option A.

upvoted 2 times

🗳️ 👤 **Raphaello** 1 year, 4 months ago

Selected Answer: A

A..ofc!

upvoted 1 times

🗳️ 👤 **awssecuritynewbie** 1 year, 5 months ago

Selected Answer: A

Non of them talk about the encryption at rest for the EBS apart from Option A

upvoted 1 times

🗳️ 👤 **3633f8f** 1 year, 6 months ago

Selected Answer: A

TLS In-Flight encryption is core functionality of ACM. Others are invalidated based on this.

upvoted 1 times

🗳️ 👤 **Daniel76** 1 year, 7 months ago

Selected Answer: A

Only A address data encryption at rest at RDS and EBS and is the most cost-effective and efficient method.

TLS certificates from a third-party vendor or generated by CloudHSM is unnecessarily increase cost and ops overhead.

CloudFront with WAF is irrelevant to the requirement.

upvoted 3 times

🗳️ 👤 **[Removed]** 1 year, 8 months ago

Selected Answer: A

Answer A



upvoted 1 times

🗳️ 👤 **lalee2** 1 year, 8 months ago

Selected Answer: A

Option A talks about ACM which is encryption in transit



upvoted 1 times

  **pupsik** 1 year, 8 months ago

Selected Answer: A



Option A

upvoted 1 times

  **KR693** 1 year, 8 months ago

Option A

upvoted 1 times

  **Sumi81** 1 year, 8 months ago

A. No other solution talks about encryption at rest

upvoted 1 times


  **100fold** 1 year, 8 months ago

Selected Answer: A

Agree answer A.

TLS certificates from (ACM) secures data in transit

upvoted 1 times

  **aragon_saa** 1 year, 8 months ago

<https://www.examttopics.com/discussions/amazon/view/60895-exam-aws-certified-security-specialty-topic-1-question-265/>

upvoted 2 times


A security engineer is working with a company to design an ecommerce application. The application will run on Amazon EC2 instances that run in an Auto Scaling group behind an Application Load Balancer (ALB). The application will use an Amazon RDS DB instance for its database. The only required connectivity from the internet is for HTTP and HTTPS traffic to the application. The application must communicate with an external payment provider that allows traffic only from a preconfigured allow list of IP addresses. The company must ensure that communications with the external payment provider are not interrupted as the environment scales.

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

- A. Deploy a NAT gateway in each private subnet for every Availability Zone that is in use.
- B. Place the DB instance in a public subnet.
- C. Place the DB instance in a private subnet.
- D. Configure the Auto Scaling group to place the EC2 instances in a public subnet.
- E. Configure the Auto Scaling group to place the EC2 instances in a private subnet.
- F. Deploy the ALB in a private subnet.

Suggested Answer: ACE

Community vote distribution



Ptol Highly Voted 1 year ago

Selected Answer: CE

This cannot possibly be an exam question. its poorly worded. None of the partial answers address the need to control the outgoing connections IP addresses. If A placed the NAT in a public subnet with an elastic IP address attached, then we would have a valid solution

upvoted 8 times

Kaps443 Most Recent 2 months, 2 weeks ago

Selected Answer: ACE

A. Deploy a NAT gateway in each private subnet for every Availability Zone that is in use.

✓ Ensures EC2s in private subnets can access the internet using static Elastic IPs via NAT — which are pre-approved by the payment provider.

C. Place the DB instance in a private subnet.

✓ Best practice: RDS should never be in a public subnet. Keeps DB isolated and secure.

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

✓ EC2s in private subnets use the NAT Gateway to talk to the internet, keeping them protected from direct public access. The ALB handles public traffic.

upvoted 1 times

phmeeeee 2 months, 2 weeks ago

Selected Answer: CE

NAT GW in private subnet is useless!

upvoted 2 times

Artemiy 2 months, 4 weeks ago

Selected Answer: CE

only two are valid

A can't be possible because nat gateway must be in public subnet

upvoted 1 times

molerowan 3 months, 3 weeks ago

Selected Answer: CE

C&E is the only legit answer A and F are a terrible idea.

upvoted 1 times

diegomarin 6 months, 2 weeks ago

Selected Answer: ACD

Nat GW em rede privada não tem saída para usar o EIP.

upvoted 2 times

  **bukkanni** 10 months ago

NAT Gateway in a private subnet. Seriously? How will EC2 instances in a private subnet get internet access if the NAT Gateway is in a private subnet and itself does not have internet access. Am I missing something?

upvoted 3 times

  **FunkyFresco** 10 months, 1 week ago

Selected Answer: ACE

ACE are correct.

upvoted 1 times

  **sema2232** 1 year ago

B,E,F are correct

upvoted 1 times

  **shailvardhan** 1 year, 1 month ago

Incorrect options as question is not clear enough.

upvoted 1 times

  **Raphaello** 1 year, 6 months ago

Keyword: " preconfigured allow list of IP addresses"

However, the question missed an important detail alongside these keywords..that the communication has to be routed to a VPN and through a virtual private gateway (VGW).

That's the only reason you can place NAT GW in a private subnet.

Without that piece of detail, placing NAT GW seems unreasonable.

Poorly written question, but if you considered that, A C E would make sense.

<https://docs.aws.amazon.com/vpc/latest/userguide/nat-gateway-scenarios.html#private-nat-allowed-range>

upvoted 4 times

  **Daniel76** 1 year, 7 months ago

Selected Answer: ACE

C and E are definitely correct.

A is confusing because NAT gateway should be created in public subnets. It should be referring to created *for each private subnets instead.

F is wrong because ALB should be in public subnet.

upvoted 2 times

  **confusedyeti69** 1 year, 6 months ago

There is private NAT gateway as well, but I'm not sure if private NAT gateway is deployed to public or private subnet.

<https://docs.aws.amazon.com/vpc/latest/userguide/vpc-nat-gateway.html>

Private – Instances in private subnets can connect to other VPCs or your on-premises network through a private NAT gateway. You can route traffic from the NAT gateway through a transit gateway or a virtual private gateway. You cannot associate an elastic IP address with a private NAT gateway. You can attach an internet gateway to a VPC with a private NAT gateway, but if you route traffic from the private NAT gateway to the internet gateway, the internet gateway drops the traffic.



upvoted 1 times

  **[Removed]** 1 year, 8 months ago

Selected Answer: ACE

A,C and E, except that NAT gateways are created in public subnets, so that private subnets can reach the Internet. Option A is worded wrong.



upvoted 3 times

  **lalee2** 1 year, 8 months ago

Selected Answer: ACE



A, C, E

upvoted 2 times

  **KR693** 1 year, 8 months ago



A, C and E

upvoted 2 times

  **Sumi81** 1 year, 8 months ago

I think its ACD

upvoted 3 times

  **100fold** 1 year, 8 months ago

Selected Answer: ACE

Answer ACE

upvoted 3 times

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

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

Which combination of steps will ensure consistent deployment of the stacks MOST securely? (Choose three.)

- A. Create a service role that has a composite principal that contains each service that needs the necessary permissions. Configure the role to allow the sts:AssumeRole action.
- B. Create a service role that has cloudformation.amazonaws.com as the service principal. Configure the role to allow the sts:AssumeRole action.
- C. For each required set of permissions, add a separate policy to the role to allow those permissions. Add the ARN of each CloudFormation stack in the resource field of each policy.
- D. For each required set of permissions, add a separate policy to the role to allow those permissions. Add the ARN of each service that needs the permissions in the resource field of the corresponding policy.
- E. Update each stack to use the service role.
- F. Add a policy to each member role to allow the iam:PassRole action. Set the policy's resource field to the ARN of the service role.

Suggested Answer: BDE

Community vote distribution



PareshBPatel Highly Voted 1 year, 4 months ago

BEF are the correct selection

Thought to consistent deployment of CloudFormation stacks would actually be

B. Create a service role that has cloudformation.amazonaws.com as the service principal. Configure the role to allow the sts:AssumeRole action.

E. Update each stack to use the service role.

F. Add a policy to each member role to allow the iam:PassRole action. Set the policy's resource field to the ARN of the service role.

These steps ensure that CloudFormation has the necessary permissions through a service role designed specifically for it (B), that each stack is configured to use this service role for deployments (E), and that users have the permission to pass this role to CloudFormation (F), aligning with best practices for security and consistency.

upvoted 17 times

cumzle.com Highly Voted 1 year ago

Selected Answer: BDE

B. Create a service role that has cloudformation.amazonaws.com as the service principal. Configure the role to allow the sts action.

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_create_for-service.html

D. For each required set of permissions, add a separate policy to the role to allow those permissions. Add the ARN of each service that needs the permissions in the resource field of the corresponding policy.

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-iam-servicerole.html#using-iam-servicerole-add>

E. Update each stack to use the service role.

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/using-iam-servicerole.html>

upvoted 8 times

Radhey_sen Most Recent 1 month, 1 week ago

Selected Answer: BE

BEF is the correct answer

upvoted 1 times

FlyingHawk 4 months, 3 weeks ago

Selected Answer: BDE

Why Not A, C, or F?



Option Reason for Rejection

A A composite principal is unnecessary here. CloudFormation is the only service assuming the role, so we only need cloudformation.amazonaws.com.

C Adding CloudFormation stack ARNs in the resource field is incorrect because policies should apply to the services being provisioned, not to CloudFormation stacks themselves.

F iam:PassRole is only needed when a user or service is delegating a role to another AWS service (e.g., EC2 assuming an IAM role). CloudFormation assumes the role directly, so this step is not required.

upvoted 2 times

  **youonebe** 5 months, 4 weeks ago

Selected Answer: BD

BDF

Option E: It's not necessary to explicitly update each stack to use the service role. CloudFormation automatically assumes the specified service role when performing stack operations, as long as the role is properly configured with the necessary permissions and trust relationships.

upvoted 1 times

  **ahirri** 6 months, 2 weeks ago

Selected Answer: BCE

Anyone that voted "D", how can you add "Service Arn" (not Resource ARN) to the "Resource" field in an IAM policy?

upvoted 2 times

  **IPLogic** 6 months, 3 weeks ago

Selected Answer: BE

BEF IS THE RIGHT ANSWER FOR THIS SCENARIO

upvoted 2 times

  **hb0011** 10 months ago

In a scenario where E and F are combined as one choice (E) as someone stated then the correct answer would be BCE.

upvoted 1 times

  **hb0011** 10 months ago

Selected Answer: BE

The voting buttons are messed up so it's showing the wrong answer. The answer is 100% definitely BEF but you can't vote for BEF.

upvoted 1 times

  **HunkyBunky** 10 months ago

Selected Answer: BDE

For me - BDE looks good.



upvoted 1 times

  **FunkyFresco** 10 months, 1 week ago

Selected Answer: BD

BDF make more sense to me.

upvoted 1 times



  **shyam87** 10 months, 1 week ago

B - the CloudFormation service needs to assume the role to create the resources

E - the stacks need to use the role to gain permissions

F - the IAM user needs the iam:PassRole permission to pass the role to the CloudFormation service

upvoted 3 times

  **5409b91** 11 months, 2 weeks ago

Selected Answer: BDE

B, D, E.

upvoted 5 times

  **shailvardhan** 1 year, 1 month ago

Selected Answer: BE

BEF are the correct answers.

upvoted 1 times

  **CloudHell** 1 year, 2 months ago



Selected Answer: BCE

B ensures that CloudFormation has the necessary permissions through a dedicated service role.

C restricts the permissions to the specific stacks, following the principle of least privilege.

E ensures that each stack uses the service role during deployment.

upvoted 4 times

  **Snape** 1 year, 2 months ago

Selected Answer: BE

BEF is correct

upvoted 4 times

  **Raphaello** 1 year, 4 months ago

Selected Answer: BDE

BDE

Create a service role to be used by CloudFormation.

For each service to be used by the CF stack, create the associated set of permissions.

Assign the service role to the stack.

The question does not feel right though, since it mentions all user assume an IAM role to access the account, therefore the stack they launch should use the permissions given to that IAM role, therefore the result should be the same for all users either (they don't launch the stack using their individual IAM users).

upvoted 1 times

  **Raphaello** 1 year, 4 months ago

Ok, looking again at the options of this question, option D is a bit tricky.

Yes you need to create permissions to CF service role, but there's nothing like "ARN of each service" to be added to the resource field. ARN's belong to resources not services, and in CF service role, resource element usually takes "*"; but even if you want to specify a resource it will be something like (arn:aws:s3:::my_bucket/*) NOT ARN OF EACH SERVICE!

ARN <--> Resource..not service.

For that, I would go with BEF.

"F" (users being able to iam:PassRole) is important and the option is worded correctly.

D is not worded correctly, as it starts with a correct part, but ended it with bogus!

BEF.

upvoted 3 times

A company used a lift-and-shift approach to migrate from its on-premises data centers to the AWS Cloud. The company migrated on-premises VMs to Amazon EC2 instances. Now the company wants to replace some of components that are running on the EC2 instances with managed AWS services that provide similar functionality.


Initially, the company will transition from load balancer software that runs on EC2 instances to AWS Elastic Load Balancers. A security engineer must ensure that after this transition, all the load balancer logs are centralized and searchable for auditing. The security engineer must also ensure that metrics are generated to show which ciphers are in use.

Which solution will meet these requirements?

- A. Create an Amazon CloudWatch Logs log group. Configure the load balancers to send logs to the log group. Use the CloudWatch Logs console to search the logs. Create CloudWatch Logs filters on the logs for the required metrics.
- B. Create an Amazon S3 bucket. Configure the load balancers to send logs to the S3 bucket. Use Amazon Athena to search the logs that are in the S3 bucket. Create Amazon CloudWatch filters on the S3 log files for the required metrics.
- C. Create an Amazon S3 bucket. Configure the load balancers to send logs to the S3 bucket. Use Amazon Athena to search the logs that are in the S3 bucket. Create Athena queries for the required metrics. Publish the metrics to Amazon CloudWatch.
- D. Create an Amazon CloudWatch Logs log group. Configure the load balancers to send logs to the log group. Use the AWS Management Console to search the logs. Create Amazon Athena queries for the required metrics. Publish the metrics to Amazon CloudWatch.

Suggested Answer: C

Community vote distribution



Ghe Highly Voted 1 year, 3 months ago

Selected Answer: C

You can't send ELB access logs into CloudWatch Logs, but to S3 only:

<https://docs.aws.amazon.com/elasticloadbalancing/latest/application/enable-access-logging.html>

Regarding the alarm, natively there is no way to use query result as a metric. We could use a Lambda Function for this.

C remains the most valid option.

upvoted 14 times

khuman_12 1 year, 1 month ago

There is nothing about "access logs" in the question.

upvoted 1 times

hro 1 year, 3 months ago

I think this is the key to the question -

save ELB access logs

To save ELB access logs, you can:

Create an S3 bucket

Attach a policy to the S3 bucket that allows Elastic Load Balancing to write the logs to the bucket

Configure access logs to capture and deliver log files to the S3 bucket

Verify bucket permissions

upvoted 1 times

Daniel76 Highly Voted 1 year, 7 months ago

Selected Answer: A

<https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-cloudwatch-metrics.html#view-metric-data>

s3 buckets and athena are not needed.

upvoted 12 times

Aamee 1 year, 7 months ago

No, if you look at the reqs. of this question, it specifically asks for this query:

"All the load balancer logs are centralized and searchable for auditing"..

So if you select A for CloudWatch Log groups, it has the default retention policy set. After which it will clear off all the saved logs!... so how would

you be able to do the audit on the logs after 14 days lets say??

That's why I'm going with Option C here..

upvoted 3 times

  **Daniel76** 1 year, 6 months ago

By default, cloudWatch log retention is indefinite unless you set it to limited duration due to audit requirement.

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/Working-with-log-groups-and-streams.html#SettingLogRetention>

upvoted 4 times

  **Daniel76** 1 year, 6 months ago

However i tend to agree it is C after reviewing this qn: enable access log can keep logs in s3 where you can search for ssl_cipher string using SQL like query:

<https://docs.aws.amazon.com/athena/latest/ug/application-load-balancer-logs.html>

you can indeed publish athena query metrics to cloudwatch by enabling the option.

<https://docs.aws.amazon.com/athena/latest/ug/query-metrics-viewing.html>

upvoted 3 times

  **03beafc** 1 year, 4 months ago

The article you posted is about the metrics behind the athena query, nothing there about generating metrics from an athena query. Answer is A

upvoted 3 times

  **hb0011** Most Recent 10 months ago

Selected Answer: C

The requirement for ciphers eliminates option A because you need ELB access logs which requires s3.

upvoted 2 times

  **FunkyFresco** 10 months, 1 week ago

Selected Answer: C

C fits with the requirements.

upvoted 1 times

  **shailvardhan** 1 year, 1 month ago

Selected Answer: C

C because elb access logs can not be sent directly to CloudWatch


upvoted 1 times

  **Just_Ninja** 1 year, 1 month ago

Selected Answer: C

Amazon Application Load Balancers (ALBs) cannot send logs directly to Amazon CloudWatch Logs. ALB access logs can be sent to Amazon S3 buckets for storage and analysis. <https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-access-logs.html>

upvoted 1 times

  **liuyomz** 1 year, 2 months ago

Selected Answer: A


Im between A and C, because probably A is less efficiente because searching through Athena is better? But i dont know, A still seems more correct.

upvoted 1 times

  **shailvardhan** 1 year, 1 month ago

but ELB access logs can not be sent directly to CloudWarch.



upvoted 1 times

  **CloudHell** 1 year, 2 months ago

Selected Answer: C

C for sure.

upvoted 2 times

  **Sodev** 1 year, 2 months ago

C. ALB only push directly to S3 + Amazon CloudWatch filters only for log on CW log stream not on S3

upvoted 1 times

  **PareshBPatel** 1 year, 4 months ago

B is not correct choice

Option B involves using S3 for storage and Athena for searching, but it suggests creating CloudWatch filters on S3 log files, which isn't directly possible as CloudWatch filters work on logs stored in CloudWatch Logs, not on S3.

upvoted 2 times

🗨️ 👤 **PareshBPatel** 1 year, 4 months ago

C

For ensuring centralized and searchable logging for auditing purposes after transitioning to AWS Elastic Load Balancers, and for generating metrics to show which ciphers are in use, the most effective solution among the provided options is:

C. Create an Amazon S3 bucket. Configure the load balancers to send logs to the S3 bucket. Use Amazon Athena to search the logs that are in the S3 bucket. Create Athena queries for the required metrics. Publish the metrics to Amazon CloudWatch.

upvoted 1 times

🗨️ 👤 **Raphaello** 1 year, 4 months ago

Selected Answer: B

Weird set of answers. Mixing between access logs and performance metrics.

Check out this

<https://docs.aws.amazon.com/elasticloadbalancing/latest/application/load-balancer-monitoring.html>

CloudWatch is responsible about collecting performance metrics.

Whereas, access logs are captured and sent to S3. You can use these logs to analyze traffic patterns, but NOT TO QUERY METRICS using Athena (it does not make sense even).

Therefore, the closest answer to correctness is B!

upvoted 2 times

🗨️ 👤 **Raphaello** 1 year, 4 months ago

Correct answer is C.

Athena is in fact capable of query metrics and publish them to CloudWatch.

<https://docs.aws.amazon.com/athena/latest/ug/athena-cloudwatch-metrics-enable.html>

upvoted 1 times

🗨️ 👤 **smanzana** 1 year, 5 months ago

A or C?? I choose C because I think that AWS ELB cant send logs to CloudWatch

upvoted 1 times

🗨️ 👤 **Gafa255** 1 year, 5 months ago

Selected Answer: C

ELB cant send log to CloudWatch.

upvoted 1 times

🗨️ 👤 **trashbox** 1 year, 6 months ago

Exam on 2023-12-18

upvoted 1 times

🗨️ 👤 **Raphaello** 1 year, 6 months ago

Correct Answer is B.

We're talking about ELB access logs, not metrics, which always get forwarded to S3 bucket.

From there one can use Athena for SQL querying.

upvoted 1 times

🗨️ 👤 **[Removed]** 1 year, 8 months ago

Selected Answer: C

Answer is C

upvoted 1 times

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

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


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

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

- A. Create an AWS CloudFormation template that contains the 10 required AWS Config rules. Deploy the template by using CloudFormation StackSets in the security-01 account.
- B. Create a conformance pack that contains the 10 required AWS Config rules. Deploy the conformance pack from the security-01 account.
- C. Create a conformance pack that contains the 10 required AWS Config rules. Deploy the conformance pack from the management-01 account.
- D. Create an AWS CloudFormation template that will activate AWS Config. Deploy the template by using CloudFormation StackSets in the security-01 account.
- E. Create an AWS CloudFormation template that will activate AWS Config. Deploy the template by using CloudFormation StackSets in the management-01 account.

Suggested Answer: BE

Community vote distribution



Christina666 Highly Voted 1 year, 7 months ago

Selected Answer: BE

Use management account to delegate accounts for auditing, security or compliance, then use delegated account to deploy conformance packs
upvoted 8 times

walter_white_008 Most Recent 1 year, 4 months ago

why is D not correct ? a delegated admin can deploy the stacks to enable the AWS config and its preferable to use the delegated account over admin account.

What am I missing ?

upvoted 1 times

Haykhay 7 months, 3 weeks ago

Delegated admin for AWS config. Not delegated admin for AWS org which manages stackset

upvoted 1 times

Daibin 1 year, 6 months ago

I'd go with B and E

<https://aws.amazon.com/blogs/mt/using-delegated-admin-for-aws-config-operations-and-aggregation/>

upvoted 3 times

trashbox 1 year, 6 months ago

Exam on 2023-12-18

upvoted 1 times

Raphaello 1 year, 6 months ago

"security-01 account" is a delegated administrator account, so let's agree that either Cfn stacksets or Config rules will be deployed from this account. Now, since there are multiple accounts, deploying AWS Config rules (conformance pack) would be through CloudFormation template/stackset. AD seems right choices for me, albeit B is not wrong but it misses the deployment part of Config rules.

upvoted 2 times

Raphaello 1 year, 4 months ago

Obviously I was wrong.
BE are the best answers.
upvoted 1 times

🗨️ 👤 **Daniel76** 1 year, 7 months ago

Selected Answer: BE

I go with B and E.

<https://aws.amazon.com/blogs/mt/deploying-conformance-packs-across-an-organization-with-automatic-remediation/>

Delegated administrator for AWS Organizations

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_delegate_policies.html

upvoted 4 times

🗨️ 👤 **Aamee** 1 year, 7 months ago

Selected Answer: CE

I'd probably go with C and E since the AWS documentation shows that it's only the management account from where the CFN stack can be deployed at along with the Conformance Packs which can also be deployed at the Management/Master account level.

But pls. correct me if I understood it incorrectly somewhere... Thnx!

upvoted 2 times

🗨️ 👤 **Josh1217** 1 year, 3 months ago

<https://aws.amazon.com/blogs/mt/deploy-aws-config-rules-and-conformance-packs-using-a-delegated-admin/>

upvoted 1 times

🗨️ 👤 **Aamee** 1 year, 7 months ago

<https://aws.amazon.com/blogs/mt/deploying-conformance-packs-across-an-organization-with-automatic-remediation/>

From the source above, it looks like the Conformance Packs can be setup only by the Master Account (Which probably in this usecase, it's the Management account I guess).

"These conformance packs and their underlying config rules and remediations actions are not modifiable by your organization's member accounts. Only master accounts can create, update, and delete organization conformance packs."

Still confused as to why we've the Security-01 account setup as the AWS Config Delegated administrator for all the member accounts?..

upvoted 1 times

🗨️ 👤 **[Removed]** 1 year, 8 months ago

Selected Answer: BE

B and E

upvoted 1 times

🗨️ 👤 **lalee2** 1 year, 8 months ago

Selected Answer: BE

B and E. Conformance should be set up in admin account but in the question it says 'security-01 as the delegated administrator for AWS Config'. I would pick B and E here.

upvoted 3 times

🗨️ 👤 **pupsik** 1 year, 8 months ago

Selected Answer: BE

Agree with @bhui.

upvoted 1 times

🗨️ 👤 **bhui** 1 year, 8 months ago

Selected Answer: BE

Should be BE

<https://aws.amazon.com/blogs/mt/deploying-conformance-packs-across-an-organization-with-automatic-remediation/>

B as security account is the AWS Config delegated admin

upvoted 4 times

🗨️ 👤 **bhui** 1 year, 8 months ago

Supplementing my thoughts with this blog.



<https://aws.amazon.com/blogs/mt/org-aggregator-delegated-admin/>

1. To enable AWS Config access to AWS Organizations

- Run the following command from your organization management account:

2. To set up an aggregator using delegated admin

upvoted 4 times

  **Sumi81** 1 year, 8 months ago

CE

<https://aws.amazon.com/blogs/mt/deploying-conformance-packs-across-an-organization-with-automatic-remediation/>

upvoted 3 times

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

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

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

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

- A. Use Amazon CloudWatch Logs Insights to identify any objects in DOC-EXAMPLE-BUCKET1 that contain PII and that were available to the access key.
- B. Use Amazon OpenSearch Service to query the CloudTrail logs in DOC-EXAMPLE-BUCKET2 for API calls that used the access key to access an object that contained PII.
- C. Use Amazon Athena to query the CloudTrail logs in DOC-EXAMPLE-BUCKET2 for any API calls that used the access key to access an object that contained PII.
- D. Use AWS Identity and Access Management Access Analyzer to identify any API calls that used the access key to access objects that contained PII in DOC-EXAMPLE-BUCKET1.
- E. Configure Amazon Macie to identify any objects in DOC-EXAMPLE-BUCKET1 that contain PII and that were available to the access key.

Suggested Answer: CE

Community vote distribution



🗳️ **hb0011** 10 months ago

Selected Answer: CE

User Macie to find the PII (E). Use Athena to search the CloudTrail bucket (bucket2) to find out who accessed the PII files (C).

upvoted 3 times

🗳️ **FunkyFresco** 10 months, 1 week ago

Selected Answer: CE

C and E

upvoted 2 times

🗳️ **shailvardhan** 1 year, 1 month ago

Selected Answer: CE

CE are correct options

upvoted 1 times

🗳️ **Jlee7** 1 year, 2 months ago

CE final answer. lol

upvoted 1 times

🗳️ **3633f8f** 1 year, 6 months ago

Selected Answer: CE

CE without further discussion.

upvoted 1 times

🗳️ **Raphaello** 1 year, 6 months ago

CE

A does not add anything.

upvoted 1 times

🗨️ 👤 **Raphaello** 1 year, 6 months ago

CE

CloudTrail is not configured to forward findings to CloudWatch in this scenario.

upvoted 1 times

🗨️ 👤 **Daniel76** 1 year, 7 months ago

Selected Answer: CE

D - using access analyzer seems to be a possible answer too:

<https://docs.aws.amazon.com/IAM/latest/UserGuide/logging-using-cloudtrail.html>

However, C should be better as Athena is used which able to return results to answer whether any object is accessed.

Macie is definitely needed to answer whether PII is present.

upvoted 1 times

🗨️ 👤 **Daniel76** 1 year, 6 months ago

Correction, C is only referring to the information captured in the cloudtrail through access analyzer api. You still need Athena for a quick and convenient search in the logs stored in s3.

<https://docs.aws.amazon.com/IAM/latest/UserGuide/logging-using-cloudtrail.html#service-name-info-in-cloudtrail>

upvoted 1 times

🗨️ 👤 **YR4591** 1 year, 7 months ago

Selected Answer: CE

Athena send the query results to s3 bucket > Macie can scan s3 bucket

upvoted 4 times

🗨️ 👤 **[Removed]** 1 year, 8 months ago

Selected Answer: C

C and E

upvoted 1 times

🗨️ 👤 **lalee2** 1 year, 8 months ago

my pick is C and E also

upvoted 1 times

🗨️ 👤 **pupsik** 1 year, 8 months ago

Selected Answer: CE

CE it is.

upvoted 2 times

🗨️ 👤 **Sumi81** 1 year, 8 months ago

CE is correct

upvoted 2 times

🗨️ 👤 **100fold** 1 year, 8 months ago

Selected Answer: CE

Agree answer CE

upvoted 2 times

🗨️ 👤 **tecte** 1 year, 8 months ago

CE is correct.

upvoted 2 times

🗨️ 👤 **aragon_saa** 1 year, 9 months ago

<https://www.examttopics.com/discussions/amazon/view/89893-exam-aws-certified-security-specialty-topic-1-question-450/>

upvoted 2 times

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

What is the likely cause of this access denial?

- A. The ACL in the bucket needs to be updated.
- B. The IAM policy does not allow the user to access the bucket.
- C. It takes a few minutes for a bucket policy to take effect.
- D. The allow permission is being overridden by the deny.

Suggested Answer: D

Community vote distribution

D (100%)

🗲️ 👤 **navid1365** 7 months, 3 weeks ago

Selected Answer: D

Explicit deny statements cannot be overridden by allow statements

upvoted 4 times

🗲️ 👤 **trashbox** 1 year ago

Exam on 2023-12-18

upvoted 1 times

🗲️ 👤 **[Removed]** 1 year, 2 months ago

Selected Answer: D

"An explicit deny in any policy overrides any allows."

https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_evaluation-logic.html#policy-eval-denyallow

upvoted 4 times

🗲️ 👤 **lalee2** 1 year, 2 months ago

Selected Answer: D

D is my answer

upvoted 1 times

🗲️ 👤 **100fold** 1 year, 2 months ago

Selected Answer: D

Agree answer D

upvoted 2 times

🗲️ 👤 **aragon_saa** 1 year, 2 months ago

<https://www.examttopics.com/discussions/amazon/view/30027-exam-aws-certified-security-specialty-topic-1-question-177/>

upvoted 2 times

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

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **FunkyFresco** 10 months, 1 week ago

D. Shield advance is for DDoS.
upvoted 1 times

🗳️ 👤 **navid1365** 1 year, 1 month ago

Selected Answer: D

Out of all the options, only Shield provides DDoS protection and detection.
upvoted 2 times

🗳️ 👤 **trashbox** 1 year, 6 months ago

Exam on 2023-12-18
upvoted 1 times

🗳️ 👤 **Raphaello** 1 year, 6 months ago

D is the correct answer.
upvoted 2 times

🗳️ 👤 **Daniel76** 1 year, 7 months ago

Selected Answer: D

<https://docs.aws.amazon.com/waf/latest/developerguide/ddos-cloudwatch-metrics.html>
upvoted 1 times

🗳️ 👤 **lalee2** 1 year, 8 months ago

Selected Answer: D

Shield Advanced for DDoS
upvoted 1 times

🗳️ 👤 **Sumi81** 1 year, 8 months ago

D is right
upvoted 2 times

🗳️ 👤 **100fold** 1 year, 8 months ago

Selected Answer: D

Agree answer D.
Shield Advanced for DDoS
upvoted 3 times

🗳️ 👤 **aragon_saa** 1 year, 8 months ago

<https://www.examttopics.com/discussions/amazon/view/89223-exam-aws-certified-security-specialty-topic-1-question-452/>
upvoted 1 times

A company hosts a web application on an Apache web server. The application runs on Amazon EC2 instances that are in an Auto Scaling group. The company configured the EC2 instances to send the Apache web server logs to an Amazon CloudWatch Logs group that the company has configured to expire after 1 year.

Recently, the company discovered in the Apache web server logs that a specific IP address is sending suspicious requests to the web application. A security engineer wants to analyze the past week of Apache web server logs to determine how many requests that the IP address sent and the corresponding URLs that the IP address requested.

What should the security engineer do to meet these requirements with the LEAST effort?

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

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **AgboolaKun** Highly Voted 1 year, 8 months ago

Selected Answer: C

The correct answer here is C. Please check the Queries for Apache server logs section of the following document -

https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/CWL_QuerySyntax-examples.html

upvoted 9 times

🗳️ 👤 **dhewa** Most Recent 9 months, 2 weeks ago

Selected Answer: C

CloudWatch Logs Insights allows you to run queries directly on your log data, making it a quick and efficient way to analyze logs without needing to export or set up additional services. You can use a custom query to filter logs by the specific IP address and extract the requested URLs

upvoted 2 times

🗳️ 👤 **shailvardhan** 1 year, 1 month ago

Selected Answer: C

no doubt it is C

upvoted 1 times

🗳️ 👤 **i7ovemyself** 1 year, 4 months ago

A is not the answer as macie is used to scan s3 buckets for PII.

upvoted 2 times

🗳️ 👤 **Raphaello** 1 year, 6 months ago

C

A classic usage of CloudWatch Logs Insights

https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/CWL_QuerySyntax-examples.html#CWL_QuerySyntax-examples-Apache

upvoted 2 times

🗳️ 👤 **Aamee** 1 year, 6 months ago

Selected Answer: C

Option C only among all others can only get this sol. done with the 'LEAST' effort as per the ask/reqs.



upvoted 1 times

🗳️ 👤 **[Removed]** 1 year, 8 months ago

Selected Answer: C

Answer is C


upvoted 1 times

  **lalee2** 1 year, 8 months ago

Selected Answer: C



C seems to be correct

upvoted 1 times

  **Sumi81** 1 year, 8 months ago

C is correct

upvoted 2 times

  **kk2000** 1 year, 8 months ago

Correct Answer is C

upvoted 2 times

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

```
2 123456789010 eni-1235b8ca 203.0.113.12 172.31.16.139 0 0 1 4 336 1432917027 1432917142 ACCEPT OK
2 123456789010 eni-1235b8ca 172.31.16.139 203.0.113.12 0 0 1 4 336 1432917094 1432917142 REJECT OK
```

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

- A. In the security group of the EC2 instance, allow inbound ICMP traffic.
- B. In the security group of the EC2 instance, allow outbound ICMP traffic.
- C. In the VPC's NACL, allow inbound ICMP traffic.
- D. In the VPC's NACL, allow outbound ICMP traffic.

Suggested Answer: D


Community vote distribution

D (100%)

 **i7ovemyself** Highly Voted 1 year, 4 months ago

For security group outbound rule is automatically allowed as security groups are stateful, NACL is stateless, so answer will be D as we need to allow the outbound rule in VPC's NACL.

upvoted 6 times

 **mnsait** 1 month, 2 weeks ago

Good answer

upvoted 1 times

 **braveheart22** Most Recent 3 weeks, 5 days ago

Selected Answer: D

The best option is D

This is why:

Looking closely at the Flow Log Analysis:

There are two entries:

ACCEPT – From 203.0.113.12 to 172.31.16.139 (on-premises to EC2)

REJECT – From 172.31.16.139 to 203.0.113.12 (EC2 to on-premises)

This tells us that:

The inbound request (ICMP Echo Request) from on-prem to EC2 is allowed.

The outbound reply (ICMP Echo Reply) from EC2 to on-prem is rejected.

Conclusion:

The ping fails because the outbound ICMP reply from the EC2 instance is being blocked.

Therefore, the correct answer will be:

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

This action will permit the Echo Reply back to the source, allowing the ping to succeed.

upvoted 2 times

 **piyushgupta9340** 2 months ago

Selected Answer: A

the most direct and likely solution is to configure the EC2 instance's security group to allow inbound ICMP traffic. The NACL is less likely to be the primary cause in this scenario.

upvoted 1 times

 **cyber_m0nk** 4 months, 1 week ago

Selected Answer: C

I failed to see how option D is correct and how this is an outbound issue. In the logs 203.0.113.2 is the source ip address sending a ping to 172.31.16.139, this is the outbound ICMP traffic which accepted. It is the inbound ICMP coming from 172.31.16.139 that is rejected, making this the

return inbound ICMP traffic. The question states the ping was initiated from the Host IP address 203.0.113.12 (source ip) making 172.31.16.139 the destination. Based on the logs Option C is the correct answer.

upvoted 1 times

🗲️ 👤 **FunkyFresco** 10 months, 1 week ago

Selected Answer: D

D is the right option.

upvoted 1 times

🗲️ 👤 **bkbaws** 1 year, 4 months ago

Selected Answer: D

see <https://docs.aws.amazon.com/vpc/latest/userguide/flow-logs-records-examples.html#flow-log-example-security-groups>

the NACL is blocking the outbound response

upvoted 2 times

🗲️ 👤 **Daniel76** 1 year, 7 months ago

Selected Answer: D

There are multiple possible cause.

<https://arcadian.cloud/aws/2022/07/01/4-reasons-you-cannot-ping-your-aws-ec2-instance-and-how-to-fix-them/>

Base on the logs, only one direction is not successful. Likely its #4 - NACL.

upvoted 1 times

🗲️ 👤 **Aamee** 1 year, 7 months ago

Selected Answer: D

It's the EC2 instance IP area from where the ping didn't get the response back to the on-prem location which is clearly a usecase of NACL area.

Therefore, def. going with 'D'.

upvoted 1 times

🗲️ 👤 **Christina666** 1 year, 7 months ago

Selected Answer: D

NACLs are stateless and do not track the state of a connection, while Security Groups are stateful and allow traffic based on the response to previous traffic.

Default rule: NACLs have a default rule that denies all traffic, while Security Groups have a default rule that allows all traffic.

upvoted 3 times

🗲️ 👤 **[Removed]** 1 year, 8 months ago

Selected Answer: D

Answer D

upvoted 1 times

🗲️ 👤 **lalee2** 1 year, 8 months ago

Selected Answer: D

Answer D

upvoted 1 times

🗲️ 👤 **pupsik** 1 year, 8 months ago

Selected Answer: D

Outbound communication on NACL is blocked.

upvoted 1 times

🗲️ 👤 **allcertcracker** 1 year, 8 months ago

it is D

upvoted 1 times

🗲️ 👤 **Sumi81** 1 year, 8 months ago

I think its B

upvoted 1 times

🗲️ 👤 **100fold** 1 year, 8 months ago

Selected Answer: D

Answer D

upvoted 1 times

🗲️ 👤 **aragon_saa** 1 year, 9 months ago

<https://www.examttopics.com/discussions/amazon/view/16473-exam-aws-certified-security-specialty-topic-2-question-8/>
upvoted 2 times

A company developed an application by using AWS Lambda, Amazon S3, Amazon Simple Notification Service (Amazon SNS), and Amazon DynamoDB. An external application puts objects into the company's S3 bucket and tags the objects with date and time. A Lambda function periodically pulls data from the company's S3 bucket based on date and time tags and inserts specific values into a DynamoDB table for further processing.

The data includes personally identifiable information (PII). The company must remove data that is older than 30 days from the S3 bucket and the DynamoDB table.

Which solution will meet this requirement with the MOST operational efficiency?

- A. Update the Lambda function to add a TTL S3 flag to S3 objects. Create an S3 Lifecycle policy to expire objects that are older than 30 days by using the TTL S3 flag.
- B. Create an S3 Lifecycle policy to expire objects that are older than 30 days. Update the Lambda function to add the TTL attribute in the DynamoDB table. Enable TTL on the DynamoDB table to expire entries that are older than 30 days based on the TTL attribute.
- C. Create an S3 Lifecycle policy to expire objects that are older than 30 days and to add all prefixes to the S3 bucket. Update the Lambda function to delete entries that are older than 30 days.
- D. Create an S3 Lifecycle policy to expire objects that are older than 30 days by using object tags. Update the Lambda function to delete entries that are older than 30 days.

Suggested Answer: B

Community vote distribution

B (100%)

  **dhewa** 9 months, 2 weeks ago

Selected Answer: B

B is the only answer that addresses Dynamo DB
upvoted 1 times

  **FunkyFresco** 10 months, 1 week ago

Selected Answer: B

B is correct.
upvoted 1 times

  **Raphaello** 1 year, 3 months ago

Selected Answer: B

B is a correct choice.
upvoted 1 times

  **i7ovemyself** 1 year, 4 months ago

Lifecycle Policy for s3 and TTL for Dynamo DB, best combination.
Answer will be B.
upvoted 1 times

  **Daniel76** 1 year, 7 months ago



Selected Answer: B

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/lifecycle-configuration-examples.html>
Just need to set expiration days in the LifecycleConfiguration- add prefix, object tags are not needed.
<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/TTL.html>
TTL is used for housekeeping data in DynamoDB by enabling TTL attribute by console or CLI, without the need for any lambda function/
upvoted 1 times

  **[Removed]** 1 year, 8 months ago

Selected Answer: B



Lifecycle policy for S3, TTL for dynamodb.
upvoted 2 times

  **lalee2** 1 year, 8 months ago

Selected Answer: B

B is right

upvoted 1 times

  **pupsik** 1 year, 8 months ago

Selected Answer: B



B is correct... although wording is not quite right.

upvoted 2 times

  **abhishek007** 1 year, 8 months ago


B is the correct answer

upvoted 1 times

  **Sumi81** 1 year, 8 months ago

B is right

upvoted 2 times

  **kk2000** 1 year, 8 months ago

Correct Answer is B

upvoted 2 times

What are the MOST secure ways to protect the AWS account root user of a recently opened AWS account? (Choose two.)

- A. Use the AWS account root user access keys instead of the AWS Management Console.
- B. Enable multi-factor authentication for the AWS IAM users with the AdministratorAccess managed policy attached to them.
- C. Use AWS KMS to encrypt all AWS account root user and AWS IAM access keys and set automatic rotation to 30 days.
- D. Do not create access keys for the AWS account root user; instead, create AWS IAM users.
- E. Enable multi-factor authentication for the AWS account root user.

Suggested Answer: DE

Community vote distribution

DE (100%)

  **i7ovemyself** 10 months, 1 week ago

A is not correct as it is not preferable to create access keys for root user.

B is not correct as enabling MFA for other user won't help the root user in any way.

C is not right as access keys shouldn't be created for root user.

Correct Answer is DE



upvoted 1 times

  **[Removed]** 1 year, 2 months ago

Selected Answer: DE

Answer D and E



upvoted 1 times

  **lalee2** 1 year, 2 months ago

Selected Answer: DE



Answer DE

upvoted 1 times

  **abhishek007** 1 year, 2 months ago



DE seems to be the correct answer

upvoted 2 times

  **Sumi81** 1 year, 2 months ago

DE is correct

upvoted 2 times

  **100fold** 1 year, 2 months ago

Selected Answer: DE

Answer DE

upvoted 3 times

  **aragon_saa** 1 year, 2 months ago

<https://www.examttopics.com/discussions/amazon/view/16530-exam-aws-certified-security-specialty-topic-1-question-135/>

upvoted 3 times

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. Upload the CloudFormation template. Add the template to the portfolio's product list. Share the portfolio with the OU.
- B. Use the CloudFormation CLI to create a module from the CloudFormation template. Register the module as a private extension in the CloudFormation registry. Publish the extension. In the OU, create an SCP that allows access to the extension.
- C. Create an AWS Service Catalog portfolio in the organization's management account. Upload the CloudFormation template. Add the template to the portfolio's product list. Create an IAM role that has a trust policy that allows cross-account access to the portfolio for users in the OU accounts. Attach the `AWSServiceCatalogEndUserFullAccess` managed policy to the role.
- D. Use the CloudFormation CLI to create a module from the CloudFormation template. Register the module as a private extension in the CloudFormation registry. Publish the extension. Share the extension with the OU.

Suggested Answer: A

Community vote distribution

A (76%)

D (24%)

 **Christina666** Highly Voted 1 year, 1 month ago

Selected Answer: A

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.

upvoted 9 times

 **Christina666** 1 year, 1 month ago

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.

upvoted 7 times

 **anandk180** Most Recent 8 months, 3 weeks ago

C - appears to be the most secure and effective way to meet the requirements. It leverages the AWS Service Catalog to manage and distribute the CloudFormation template, ensuring that only approved resources are deployed. Additionally, by creating an IAM role with a trust policy for cross-account access and attaching the appropriate managed policy, it ensures that only developers from within the OU who have the necessary permissions can access and use the portfolio, aligning with the need for strict access control to the deployment plan.

upvoted 1 times

 **Sodev** 8 months, 4 weeks ago

2 requirements:

- Deployment plan for AWS resources
- To only the developers who need access

=> That's a trap question, all people think about WS Service Catalog, However:

A. Share the portfolio with the OU. => not fit 2nd requirement

B. It match tho both req

C. `AWSServiceCatalogEndUserFullAccess` => user able to change resources in product list

D. Share the extension with the OU => not fit 2nd requirement

=> B is answer

upvoted 1 times

🗨️ **i7ovemyself** 10 months, 1 week ago

AWS Service Catalog is a service that enables organizations to create and manage catalogs of IT services that are approved for use on AWS. These catalogs can include everything from virtual machine images, servers, software, and databases to entire multi-tier application architectures.

Correct answer is A

C will not be correct as providing full access will not be secure as per AWS standards.

upvoted 2 times

🗨️ **Raphaello** 10 months, 2 weeks ago

Selected Answer: A

Sharing Service Catalog portfolio is more secure than allowing full access.

Correct answer is A.

upvoted 3 times

🗨️ **Raphaello** 10 months, 1 week ago

I spent some time looking into this question.

CF modules vs Service Catalog

It is resource centric vs service centric. Infra. provisioning vs service management. A low level building block vs service approval.

The scenario talks about ensuring "DEVELOPER" sticking to "deployment plan", which implies assurance at a lower level than service catalog.

Therefore, D could be the right answer.

Again, It all depends on the interpretation, but no clear "right" or "wrong" answer in this one.

[https://stackshare.io/stackups/aws-cloudformation-vs-aws-service-](https://stackshare.io/stackups/aws-cloudformation-vs-aws-service-catalog#:~:text=In%20summary%2C%20AWS%20CloudFormation%20is,catalogs%20of%20pre%2Dapproved%20services.)

[catalog#:~:text=In%20summary%2C%20AWS%20CloudFormation%20is,catalogs%20of%20pre%2Dapproved%20services.](https://stackshare.io/stackups/aws-cloudformation-vs-aws-service-catalog#:~:text=In%20summary%2C%20AWS%20CloudFormation%20is,catalogs%20of%20pre%2Dapproved%20services.)

upvoted 2 times

🗨️ **longns** 11 months ago

Selected Answer: A

C: incorrect because allows cross-account access to the portfolio for users in the OU accounts. Attach the AWSServiceCatalogEndUserFullAccess managed policy to the role --> this violate rule give least privilege

Privilege of <https://docs.aws.amazon.com/aws-managed-policy/latest/reference/AWSServiceCatalogEndUserFullAccess.html>

upvoted 1 times

🗨️ **Aamee** 1 year, 1 month ago

Still a bit ambiguous btw A and D... Not sure fully :/

upvoted 1 times

🗨️ **kejam** 1 year, 1 month ago

Selected Answer: A

You can share a Service Catalog portfolio to an Org OU

<https://docs.aws.amazon.com/organizations/latest/userguide/services-that-can-integrate-servicecatalog.html>

upvoted 2 times

🗨️ **[Removed]** 1 year, 1 month ago

Selected Answer: A

Service Catalog sounds like the right tool for the job

upvoted 1 times

🗨️ **M2ao** 1 year, 1 month ago

why not C

upvoted 1 times

🗨️ **bannium** 1 year, 2 months ago

Selected Answer: A

I think Option A is sufficient for our needs.

<https://aws.amazon.com/about-aws/whats-new/2022/11/aws-service-catalog-sharing-principal-names-portfolio-organization/>

upvoted 2 times

🗨️ 👤 **100fold** 1 year, 2 months ago

B instead of D

SCP limits the access

upvoted 2 times

🗨️ 👤 **Aamee** 1 year, 1 month ago

Are you sure that it should be B?... cuz will SCP going to work under the OU level?..

upvoted 1 times

🗨️ 👤 **100fold** 1 year, 2 months ago

Selected Answer: D

Answer D. You can use the CloudFormation (CLI).

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/registry.html>

<https://docs.aws.amazon.com/cloudformation-cli/latest/userguide/modules.html>

<https://aws.amazon.com/blogs/mt/introducing-aws-cloudformation-modules/>

upvoted 4 times

🗨️ 👤 **Raphaello** 10 months, 1 week ago

You probably have a point there. I learned something. Thank you.

upvoted 2 times

🗨️ 👤 **804a66d** 9 months, 2 weeks ago

Yes you are right D is a strong possible choice the module feature of AWS CloudFormation can accomplish the task but you are missing one aspect of the question. It asks for the MOST secure way and since the question states that AWS Organizations is being implemented and 3rd party users are accessing the environment then the MOST secure way is using AWS Service Catalog and establish that guardrail structure from the top down affecting the OU's.

Thanks for the links explaining modules that is a new concept I learned from your links!

upvoted 2 times

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

Which solution will meet these requirements?


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

Suggested Answer: C

Community vote distribution

C (85%)

A (15%)

 **100fold** Highly Voted 1 year, 8 months ago

Selected Answer: C

Answer C

https://docs.aws.amazon.com/IAM/latest/UserGuide/introduction_attribute-based-access-control.html

<https://aws.amazon.com/blogs/security/simplify-granting-access-to-your-aws-resources-by-using-tags-on-aws-iam-users-and-roles/>

upvoted 7 times

 **FLucas** Most Recent 1 month, 3 weeks ago

Selected Answer: C

Answer C - ABAC

upvoted 1 times

 **FunkyFresco** 10 months, 1 week ago

Selected Answer: C

C makes sense to me.

upvoted 1 times

 **rishikeshshukla7777** 1 year, 3 months ago

answer is C tag based is ABAC

upvoted 1 times

 **arvehisa** 1 year, 3 months ago

Selected Answer: C

C, as the principles changes, it's not good to always update the trusted policies but to tag the IAM principles correctly.

upvoted 1 times

 **Raphaello** 1 year, 4 months ago

Selected Answer: C

ABAC (aka tag-based policy control) provides flexible and scalable control.

upvoted 1 times

 **awssecuritynewbie** 1 year, 7 months ago

Selected Answer: A



A

upvoted 2 times

 **1c7c461** 1 year, 6 months ago

It is not A, it is not best-practice to use inline policies on a role.

upvoted 1 times

  **kejam** 1 year, 7 months ago

Selected Answer: C

https://docs.aws.amazon.com/secretsmanager/latest/userguide/managing-secrets_tagging.html

upvoted 2 times

A company is hosting a web application on Amazon EC2 instances behind an Application Load Balancer (ALB). The application has become the target of a DoS attack. Application logging shows that requests are coming from a small number of client IP addresses, but the addresses change regularly.

The company needs to block the malicious traffic with a solution that requires the least amount of ongoing effort.




Which solution meets these requirements?

- A. Create an AWS WAF rate-based rule, and attach it to the ALB.
- B. Update the security group that is attached to the ALB to block the attacking IP addresses.
- C. Update the ALB subnet's network ACL to block the attacking client IP addresses.
- D. Create an AWS WAF rate-based rule, and attach it to the security group of the EC2 instances.

Suggested Answer: A

Community vote distribution


A (100%)

  **Daniel76**  1 year, 7 months ago

Selected Answer: A




AWS WAF can be deployed on Amazon CloudFront, the Application Load Balancer (ALB), Amazon API Gateway, and AWS AppSync.

upvoted 5 times

  **smanzana** 1 year, 5 months ago

Indeed, AWS WAF cannot be directly attached to the Security Group of EC2 instances.

upvoted 1 times

  **FunkyFresco**  10 months, 1 week ago

Selected Answer: A

A without any doubt.



upvoted 1 times

  **navid1365** 1 year, 1 month ago

Selected Answer: A

Attach a WAF to the ALB

upvoted 1 times

  **Sodev** 1 year, 2 months ago

What's that mean "a small number of client IP addresses, but the addresses change regularly." : v if small number, why not use C =)) i thing this question not exactly, "small" is redundant words ?

upvoted 2 times

  **khuman_12** 1 year, 1 month ago

Exactly, the question is incorrect. And answer A: in other words "just create rate-limit and all DDoS will take up all the limit and the service will down"

upvoted 1 times


  **Raphaello** 1 year, 6 months ago

Selected Answer: A

A

WAF rate-based rule and attach it to ALB.

upvoted 1 times

  **Aamee** 1 year, 7 months ago

Selected Answer: A

WAF protects CloudFront, R53 and ALB as they're tightly integrated with WAF.

upvoted 1 times

  **awssecuritynewbie** 1 year, 7 months ago



They could of used the VPC flow logs to figure out the IPs that are attacking then use lambda to update the ACL NACL for the LB. it would be better than actually applying rate limiting.

upvoted 1 times

  **Josh1217** 1 year, 3 months ago

The IPs keep on changing.

upvoted 1 times

  **100fold** 1 year, 8 months ago

Selected Answer: A

Answer A

<https://www.examttopics.com/discussions/amazon/view/61173-exam-aws-certified-security-specialty-topic-1-question-259/>

upvoted 4 times

A company has hundreds of AWS accounts in an organization in AWS Organizations. The company operates out of a single AWS Region. The company has a dedicated security tooling AWS account in the organization. The security tooling account is configured as the organization's delegated administrator for Amazon GuardDuty and AWS Security Hub. The company has configured the environment to automatically enable GuardDuty and Security Hub for existing AWS accounts and new AWS accounts.

The company is performing control tests on specific GuardDuty findings to make sure that the company's security team can detect and respond to security events. The security team launched an Amazon EC2 instance and attempted to run DNS requests against a test domain, example.com, to generate a DNS finding. However, the GuardDuty finding was never created in the Security Hub delegated administrator account.

Why was the finding was not created in the Security Hub delegated administrator account?


- A. VPC flow logs were not turned on for the VPC where the EC2 instance was launched.
- B. The VPC where the EC2 instance was launched had the DHCP option configured for a custom OpenDNS resolver.
- C. The GuardDuty integration with Security Hub was never activated in the AWS account where the finding was generated.
- D. Cross-Region aggregation in Security Hub was not configured.

Suggested Answer: B

Community vote distribution

B (89%)

11%

 **[Removed]**  1 year, 7 months ago


Selected Answer: B

Guardduty cannot detect DNS requests if a custom resolver is setup

See below:

[https://repost.aws/knowledge-center/guardduty-finding-](https://repost.aws/knowledge-center/guardduty-finding-types#:~:text=Note%3A%20GuardDuty%20only%20processes%20DNS%20logs%20if%20you%20use%20the%20default%20VPC%20DNS%20resolver.%20All%20upvoted%209%20times)

types#:~:text=Note%3A%20GuardDuty%20only%20processes%20DNS%20logs%20if%20you%20use%20the%20default%20VPC%20DNS%20resolver.%20All%20upvoted 9 times

 **Daniel76** 1 year, 7 months ago

"GuardDuty only processes DNS logs if you use the default VPC DNS resolver. All other types of DNS resolvers won't generated DNS based findings."

upvoted 8 times


 **youonebe**  5 months, 3 weeks ago

Selected Answer: B

Answer B.

The company operates out of a single AWS Region, so D is out

upvoted 1 times


 **IPLogic** 6 months, 4 weeks ago

Selected Answer: B

The most likely reason why the GuardDuty finding was not created in the Security Hub delegated administrator account is option B: The VPC where the EC2 instance was launched had the DHCP option configured for a custom OpenDNS resolver.

When the VPC is configured with a custom DNS resolver, GuardDuty might not be able to detect DNS requests made by the EC2 instance, which would prevent the finding from being generated

upvoted 1 times


 **Davidng88** 9 months, 3 weeks ago

Selected Answer: B

AWS GuardDuty only detects DNS queries to Amazon provided DNS resolver only.

VPC flow not required for GuardDuty DNS, there will be no GuardDuty finding if GuardDuty not integrated with Security Hub, and only reside in one Region.

upvoted 1 times

 **Sodev** 1 year, 2 months ago

Why not A.

Backdoor:EC2/DenialOfService.Dns finding , evaluate VPC Flow logs, it it disable nothing to do lol

upvoted 1 times

🗨️ 👤 **hro** 1 year, 3 months ago

C... GuardDuty can access and process your request and response DNS logs through the internal ... Im going with C - it more likely that someone in the security team didnt activate Security Hub for the delegated admin account than any other of the answers.

upvoted 1 times

🗨️ 👤 **hro** 1 year, 3 months ago

C - Im going with this answer.

GuardDuty can access and process DNS logs through AWS DNS resolvers, such as Amazon Route 53, if the EC2 instances use them

upvoted 1 times

🗨️ 👤 **arvehisa** 1 year, 3 months ago

I don't understand why guarddduty should generate a finding against a test dns query.

upvoted 1 times

🗨️ 👤 **mynickc** 1 year, 5 months ago

Selected Answer: B

B is correct. <https://docs.aws.amazon.com/securityhub/latest/userguide/securityhub-internal-providers.html#integration-amazon-guarddduty>

upvoted 1 times

🗨️ 👤 **CloudRover** 1 year, 5 months ago

Selected Answer: B

"GuardDuty only processes DNS logs if you use the default VPC DNS resolver. All other types of DNS resolvers won't generated DNS based findings."

As Daniel76 pointed out, this is the correct answer.

upvoted 3 times

🗨️ 👤 **trashbox** 1 year, 6 months ago

Exam on 2023-12-18

upvoted 1 times

🗨️ 👤 **Raphaello** 1 year, 6 months ago

Selected Answer: D

Going with D.

<https://docs.aws.amazon.com/securityhub/latest/userguide/finding-aggregation-enable.html#finding-aggregation-enable-console>

upvoted 1 times

🗨️ 👤 **Raphaello** 1 year, 4 months ago

Correction: correct answer is B.

Missed that in the scenario the company operates from a single region, additionally, without using Route53 resolver and DNS query logs, GuardDuty would not be able to produce DNS findings.

upvoted 1 times

🗨️ 👤 **KaiW** 1 year, 6 months ago

but didn't the question said that the company operates out of a single region?

upvoted 1 times

🗨️ 👤 **kejam** 1 year, 7 months ago

Selected Answer: D

Choosing D through a process of elimination.

A. VPC flow logs are not required to be turned on.

<https://aws.amazon.com/guarddduty/faqs/>

B. Custom DNS resolver? GuardDuty should have picked that up from the VPC flow logs:

https://docs.aws.amazon.com/guarddduty/latest/ug/guarddduty_finding-types-ec2.html#defenseevasion-ec2-unusaldnsresolver

C. GuardDuty and Security Hub integration is enabled automatically:

<https://docs.aws.amazon.com/guarddduty/latest/ug/securityhub-integration.html#securityhub-integration-enable>

D. Cross-Region aggregation. Regions weren't mentioned in the question, but it is the only possible answer left.

<https://docs.aws.amazon.com/securityhub/latest/userguide/finding-aggregation.html>

upvoted 1 times

🗨️ 👤 **kejam** 1 year, 7 months ago

Just noticed "The company operates out of a single AWS Region." So changing my answer to none of the above ;-)

upvoted 1 times

🗨️ 👤 **bannium** 1 year, 8 months ago

Selected Answer: B

> Note: GuardDuty only processes DNS logs if you use the default VPC DNS resolver. All other types of DNS resolvers won't generated DNS based findings.

<https://repost.aws/knowledge-center/guardduty-finding-types>

upvoted 4 times

An ecommerce company has a web application architecture that runs primarily on containers. The application containers are deployed on Amazon Elastic Container Service (Amazon ECS). The container images for the application are stored in Amazon Elastic Container Registry (Amazon ECR). The company's security team is performing an audit of components of the application architecture. The security team identifies issues with some container images that are stored in the container repositories.


The security team wants to address these issues by implementing continual scanning and on-push scanning of the container images. The security team needs to implement a solution that makes any findings from these scans visible in a centralized dashboard. The security team plans to use the dashboard to view these findings along with other security-related findings that they intend to generate in the future. There are specific repositories that the security team needs to exclude from the scanning process.

Which solution will meet these requirements?

- A. Use Amazon Inspector. Create inclusion rules in Amazon ECR to match repositories that need to be scanned. Push Amazon Inspector findings to AWS Security Hub.
- B. Use ECR basic scanning of container images. Create inclusion rules in Amazon ECR to match repositories that need to be scanned. Push findings to AWS Security Hub.
- C. Use ECR basic scanning of container images. Create inclusion rules in Amazon ECR to match repositories that need to be scanned. Push findings to Amazon Inspector.
- D. Use Amazon Inspector. Create inclusion rules in Amazon Inspector to match repositories that need to be scanned. Push Amazon Inspector findings to AWS Config.

Suggested Answer: A

Community vote distribution



AgboolaKun Highly Voted 1 year, 8 months ago

Selected Answer: A

Amazon Inspector supports the configuration of inclusion rules to select which ECR repositories are scanned. Please see more information here - <https://aws.amazon.com/inspector/faqs/>
upvoted 7 times

IPLogic Most Recent 6 months, 4 weeks ago

Selected Answer: A

To meet the requirements of continual scanning and on-push scanning of container images, and to make the findings visible in a centralized dashboard, option A is the most suitable solution. Here's why:

Using Amazon Inspector for scanning container images stored in Amazon ECR provides automated, continuous scanning for vulnerabilities. By creating inclusion rules in Amazon ECR to match the repositories that need to be scanned, you can ensure that only the relevant repositories are included in the scanning process. Pushing Amazon Inspector findings to AWS Security Hub allows for centralized visibility of these findings along with other security-related findings, providing a comprehensive dashboard for the security team
upvoted 2 times

[Removed] 9 months, 4 weeks ago

Amazon Inspector provides a robust security assessment tool that can perform continuous and on-push scanning of container images stored in Amazon ECR. It integrates well with Amazon ECR, allowing for the detection of vulnerabilities and security issues within container images.

Inclusion Rules in Amazon ECR: By configuring inclusion rules in Amazon ECR, the security team can specify which repositories should be scanned, allowing them to exclude certain repositories from the scanning process as required.

Option : A

upvoted 1 times

Raphaello 1 year, 4 months ago

Selected Answer: A

For continual and on-push scanning, use Amazon Inspector. Push findings to Security Hub.
upvoted 1 times

Aamee 1 year, 7 months ago

Selected Answer: A

Def. it's A as per the features described for Amazon Inspector here:

<https://aws.amazon.com/inspector/faqs/>

upvoted 1 times

🗨️ 👤 **kejam** 1 year, 7 months ago

Selected Answer: A

Answer A

Inspector can continuously scan ECR and send findings to Security Hub

<https://docs.aws.amazon.com/inspector/latest/user/scanning-ecr.html>

upvoted 1 times

🗨️ 👤 **bannium** 1 year, 8 months ago

Selected Answer: A

using Amazon ECR integrates with Amazon Inspector with filters

<https://docs.aws.amazon.com/AmazonECR/latest/userguide/image-scanning.html#image-scanning-filters>

upvoted 2 times

🗨️ 👤 **bengalister** 1 year, 8 months ago

Answer A

Amazon inspector can definitely scan ECR repositories

<https://docs.aws.amazon.com/inspector/latest/user/scanning-ecr.html>

upvoted 2 times

🗨️ 👤 **pupsik** 1 year, 8 months ago

Selected Answer: B

Inspector scans EC2 instances.

upvoted 2 times

🗨️ 👤 **lightrod** 1 year, 4 months ago

it can scan ec2, ecr, and lambda

upvoted 2 times

🗨️ 👤 **angelsrp** 1 year, 8 months ago

B

ECR does provide basic image scanning functionality, which can detect software vulnerabilities in your container images. AWS Security Hub provides a centralized view for security alert and compliance posture across AWS accounts. This solution seems to meet the requirements. Amazon Inspector is used for analyzing EC2 instances and identifying potential security vulnerabilities, but not for container images.

upvoted 2 times

🗨️ 👤 **Daniel76** 1 year, 6 months ago

ECR basic scanning only can be configured to on push, or do manual. it does not support continuously scan as required.

ECR enhanced scanning integrates with AWS Inspector - so yes it covers not just EC2 instance but container.

upvoted 2 times

A company has a single AWS account and uses an Amazon EC2 instance to test application code. The company recently discovered that the instance was compromised. The instance was serving up malware. The analysis of the instance showed that the instance was compromised 35 days ago.

A security engineer must implement a continuous monitoring solution that automatically notifies the company's security team about compromised instances through an email distribution list for high severity findings. The security engineer must implement the solution as soon as possible.

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

- A. Enable AWS Security Hub in the AWS account.
- B. Enable Amazon GuardDuty in the AWS account.
- C. Create an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the security team's email distribution list to the topic.
- D. Create an Amazon Simple Queue Service (Amazon SQS) queue. Subscribe the security team's email distribution list to the queue.
- E. Create an Amazon EventBridge rule for GuardDuty findings of high severity. Configure the rule to publish a message to the topic.
- F. Create an Amazon EventBridge rule for Security Hub findings of high severity. Configure the rule to publish a message to the queue.

Suggested Answer: BCE

Community vote distribution

BCE (100%)

  **100fold** Highly Voted 1 year, 8 months ago

Selected Answer: BCE

Answer BCE

upvoted 6 times

  **WeepingMaple** Highly Voted 1 year, 6 months ago

AWS Security Hub does not have any scanning capabilities. It provides you with a comprehensive view of your security state only.

upvoted 5 times

  **c6ed25a** Most Recent 3 months, 1 week ago

Selected Answer: BCE

BCE is correct. Security hub is used for comprehensive view

upvoted 1 times

  **FunkyFresco** 10 months, 1 week ago

Selected Answer: BCE

BCE are the right options.

upvoted 1 times

  **Ved_1142** 1 year ago

BCE its

GuardDuty -- for monitoring

SNS- to Notify

GuardDuty in EventBridge as Rule as mentioned in Question.

upvoted 1 times



  **Raphaello** 1 year, 4 months ago

Selected Answer: BCE

BCE..obviously.

GuardDuty + EventBridge + SNS



upvoted 1 times

  **Daniel76** 1 year, 7 months ago

Selected Answer: BCE



GuardDuty, Eventbridge, SNS topics

upvoted 2 times

  **Daniel76** 1 year, 7 months ago



Selected Answer: BCE

<https://repost.aws/knowledge-center/guardduty-eventbridge-sns-rule>
upvoted 3 times

  **Aamee** 1 year, 7 months ago

Selected Answer: BCE

BCE options look most relevant.
upvoted 1 times

  **pupsik** 1 year, 8 months ago

Selected Answer: BCE

BCE it is.
upvoted 2 times

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

A user is unable to assume the IAM role in the target account. The policy attached to the role in the identity account is:

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

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

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

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "sts:AssumeRole"
      ],
      "Resource": [
        "arn:aws:iam::123456789123:role/JobFunctionRole"
      ],
      "Effect": "Allow"
    }
  ]
}
```

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

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "AWS": "arn:aws:iam::987654321987:role/IdentityRole"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

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

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": { "AWS": "arn:aws:iam::987654321987:root" },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

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

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Stmt1502946463000",
      "Effect": "Allow",
      "Action": "sts:AssumeRole",
      "Resource": "arn:aws:iam::123456789123:role/JobFunctionRole"
    }
  ]
}
```

Suggested Answer: B

Community vote distribution

B (100%)

 **Daniel76** Highly Voted 1 year, 7 months ago

Selected Answer: B

When an user is unable to assume a role in the target account, one should check the principal element in the trust policy in the JobFunctionRole in the target account.

Refer to this article to understand permission vs trust policies.

<https://www.linkedin.com/pulse/permission-policy-vs-trust-aws-rupesh-tiwari/>

upvoted 5 times

 **Lilik** Most Recent 6 months, 3 weeks ago

Selected Answer: B

In IAM roles, use the Principal element in the role trust policy to specify who can assume the role.


upvoted 1 times

 **arvehisa** 1 year, 3 months ago

Selected Answer: B

B is the correct one.

upvoted 3 times

 **kejam** 1 year, 7 months ago

Selected Answer: B

Answer B

In IAM roles, use the Principal element in the role trust policy to specify who can assume the role.



https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_elements_principal.html

upvoted 3 times

 **bengalister** 1 year, 8 months ago

B is the correct answer

upvoted 2 times

  **pupsik** 1 year, 8 months ago

Selected Answer: B

B it is.



upvoted 2 times

  **100fold** 1 year, 8 months ago

Selected Answer: B

Agree with answer B

upvoted 2 times

  **kk2000** 1 year, 8 months ago

Correct Answer is B

upvoted 3 times

A company is using AWS Organizations to manage multiple AWS accounts for its human resources, finance, software development, and production departments. All the company's developers are part of the software development AWS account. The company discovers that developers have launched Amazon EC2 instances that were preconfigured with software that the company has not approved for use. The company wants to implement a solution to ensure that developers can launch EC2 instances with only approved software applications and only in the software development AWS account. Which solution will meet these requirements?

- A. In the software development account, create AMIs of preconfigured instances that include only approved software. Include the AMI IDs in the condition section of an AWS CloudFormation template to launch the appropriate AMI based on the AWS Region. Provide the developers with the CloudFormation template to launch EC2 instances in the software development account.
- B. Create an Amazon EventBridge rule that runs when any EC2 RunInstances API event occurs in the software development account. Specify AWS Systems Manager Run Command as a target of the rule. Configure Run Command to run a script that will install all approved software onto the instances that the developers launch.
- C. Use an AWS Service Catalog portfolio that contains EC2 products with appropriate AMIs that include only approved software. Grant the developers permission to access only the Service Catalog portfolio to launch a product in the software development account.
- D. In the management account, create AMIs of preconfigured instances that include only approved software. Use AWS CloudFormation StackSets to launch the AMIs across any AWS account in the organization. Grant the developers permission to launch the stack sets within the management account.

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **AgboolaKun** Highly Voted 1 year, 8 months ago

Selected Answer: C

You can use AWS Service Catalog to build a customized AMI from which a team can only launch products from. Please see - <https://aws.amazon.com/blogs/mt/use-aws-service-catalog-to-build-a-custom-catalog-of-products-from-aws-marketplace/>
upvoted 9 times

🗳️ 👤 **c6ed25a** Most Recent 3 months, 1 week ago

Selected Answer: C

All others are wrong only C is correct
upvoted 1 times

🗳️ 👤 **Lilik** 6 months, 3 weeks ago

Selected Answer: C

Use the AWS Service Catalog to create and manage custom catalogs of products that you approve for use in your organization.
upvoted 1 times

🗳️ 👤 **navid1365** 1 year, 1 month ago

Selected Answer: C

C is correct
upvoted 2 times

🗳️ 👤 **arvehisa** 1 year, 3 months ago

Selected Answer: C

C is the correct one.
upvoted 1 times

🗳️ 👤 **Raphaello** 1 year, 6 months ago

Selected Answer: C



AWS Service Catalog is introduced for this specific purpose: govern and pre-configure provisioning of approved products
upvoted 1 times

🗳️ 👤 **WeepingMaple** 1 year, 6 months ago

Selected Answer: C

AWS Service Catalog is a service that enables organizations to govern, manage, and automate the provisioning of IT services across their AWS accounts.



upvoted 2 times

  **100fold** 1 year, 8 months ago

Selected Answer: C

Answer C

upvoted 4 times

  **kk2000** 1 year, 8 months ago

C should be the correct answer

upvoted 4 times

A company has enabled Amazon GuardDuty in all AWS Regions as part of its security monitoring strategy. In one of its VPCs, the company hosts an Amazon EC2 instance that works as an FTP server. A high number of clients from multiple locations contact the FTP server. GuardDuty identifies this activity as a brute force attack because of the high number of connections that happen every hour.

The company has flagged the finding as a false positive, but GuardDuty continues to raise the issue. A security engineer must improve the signal-to-noise ratio without compromising the company's visibility of potential anomalous behavior.

Which solution will meet these requirements?

- A. Disable the FTP rule in GuardDuty in the Region where the FTP server is deployed.
- B. Add the FTP server to a trusted IP list. Deploy the list to GuardDuty to stop receiving the notifications.
- C. Create a suppression rule in GuardDuty to filter findings by automatically archiving new findings that match the specified criteria.
- D. Create an AWS Lambda function that has the appropriate permissions to delete the finding whenever a new occurrence is reported.

Suggested Answer: C

Community vote distribution

C (100%)

100fold Highly Voted 1 year, 8 months ago

Selected Answer: C

Answer C

https://docs.aws.amazon.com/guardduty/latest/ug/findings_suppression-rule.html

upvoted 6 times

Sat897 Most Recent 5 days ago

Selected Answer: C

A suppression rule is a set of criteria, consisting of a filter attribute paired with a value, used to filter findings by automatically archiving new findings that match the specified criteria. Suppression rules can be used to filter low-value findings, false positive findings, or threats you do not intend to act on, to make it easier to recognize the security threats with the most impact to your environment.

upvoted 1 times

FunkyFresco 10 months, 1 week ago

Selected Answer: C

Option C

upvoted 1 times

Raphaello 1 year, 3 months ago

Selected Answer: C

Create a suppression rule for certain findings so it will be automatically forwarded to "Archive" and do not show up in Security Hub (if enabled), Detective, or EventBridge.

Option C is correct.

upvoted 1 times

WeepingMaple 1 year, 6 months ago

Selected Answer: C

A suppression rule is a set of criteria, consisting of a filter attribute paired with a value, used to filter findings by automatically archiving new findings that match the specified criteria.

upvoted 2 times

Daniel76 1 year, 7 months ago

Selected Answer: C

Suppression rule with specific criteria is the most appropriate solution.

Disable FTP entirely, trust everything from the FTP server reduces the GuardDuty effectiveness.

Creating lambda to delete the finding is counter productive and the finding might have already trigger SNS topic if there's one.

upvoted 3 times

Aamee 1 year, 7 months ago

Selected Answer: C

Self-explanatory from the link provided below.

upvoted 1 times

A company is running internal microservices on Amazon Elastic Container Service (Amazon ECS) with the Amazon EC2 launch type. The company is using Amazon Elastic Container Registry (Amazon ECR) private repositories.

A security engineer needs to encrypt the private repositories by using AWS Key Management Service (AWS KMS). The security engineer also needs to analyze the container images for any common vulnerabilities and exposures (CVEs).

Which solution will meet these requirements?

- A. Enable KMS encryption on the existing ECR repositories. Install Amazon Inspector Agent from the ECS container instances' user data. Run an assessment with the CVE rules.
- B. Recreate the ECR repositories with KMS encryption and ECR scanning enabled. Analyze the scan report after the next push of images.
- C. Recreate the ECR repositories with KMS encryption and ECR scanning enabled. Install AWS Systems Manager Agent on the ECS container instances. Run an inventory report.
- D. Enable KMS encryption on the existing ECR repositories. Use AWS Trusted Advisor to check the ECS container instances and to verify the findings against a list of current CVEs.

Suggested Answer: B

Community vote distribution

B (100%)

  **100fold** Highly Voted 1 year, 2 months ago

Selected Answer: B

Answer B

<https://docs.aws.amazon.com/AmazonECR/latest/userguide/repository-create.html>

<https://docs.aws.amazon.com/AmazonECR/latest/userguide/repository-edit.html>

upvoted 7 times

  **b5f86df** Most Recent 3 weeks, 5 days ago

Selected Answer: B

AWS requires you to recreate the ECR repository if you want to enable KMS encryption (you can't enable it on existing repositories).

- Vulnerability scanning: Amazon ECR supports image scanning (basic and enhanced) for CVEs. Scans occur when images are pushed, and you can analyze the results from the scan findings.

upvoted 1 times

  **Raphaello** 10 months, 1 week ago

Selected Answer: B

Need to recreate ECR to enable encryption using KMS key.

Option B is correct.

upvoted 3 times

  **Jonu** 11 months, 3 weeks ago

How about the CVE part of the question?

upvoted 4 times

  **Osirus** 1 year ago

Answer B

upvoted 1 times

  **Osirus** 1 year ago

This should be correct

upvoted 1 times

  **Daniel76** 1 year ago

Selected Answer: B

ECR repositories need to be recreated not enabled, with KMS encryption.

The inventory for AWS system manager does not contain anything about vulnerability.

upvoted 3 times

  **ahrentom** 1 year, 2 months ago

Selected Answer: B

should be the right answer

upvoted 2 times


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 AWS services. The contractor's IAM account must not be able to gain access to any other AWS 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 inline IAM user policy that allows for Amazon EC2 access for the contractor's IAM user.
- B. Create an IAM permissions boundary policy that allows Amazon EC2 access. Associate the contractor's IAM account with the IAM permissions boundary policy.
- C. Create an IAM group with an attached policy that allows for Amazon EC2 access. Associate the contractor's IAM account with the IAM group.
- D. Create a IAM role that allows for EC2 and explicitly denies all other services. Instruct the contractor to always assume this role.

Suggested Answer: B

Community vote distribution



WeepingMaple Highly Voted 1 year ago

Selected Answer: B

IAM permissions boundary policy is a managed policy that defines the maximum permissions that an identity-based policy can grant to an IAM entity (user or role). It essentially acts as a safety net to prevent users and roles from exceeding their intended permissions.

upvoted 9 times

c6ed25a Most Recent 3 months, 1 week ago

Selected Answer: B

B Boundary defines limitation

upvoted 1 times

navid1365 7 months, 2 weeks ago

Selected Answer: B

A permissions boundary defines the maximum level of access that an IAM identity can have.

upvoted 1 times

Raphaello 10 months, 3 weeks ago

Selected Answer: B

IAM permissions boundary definition use.

upvoted 1 times

trashbox 1 year ago

Exam on 2023-12-18

upvoted 2 times

Raphaello 10 months, 3 weeks ago

What do you mean?

upvoted 3 times

Daniel76 1 year ago

Selected Answer: B

Only B talks about restricting the access, by using permission boundary.

D - if you assign more than one role to the vendor, there's always risk that the instruction is not followed.

A, C- regardless of feasibility, by creating allow doesn't block the vendor from accessing services other than EC2 instance.

upvoted 2 times

Aamee 1 year, 1 month ago

Selected Answer: B

B makes more sense to me as it would explicitly define the specific service based IAM permissions policy which then can be associated with the contractor's IAM account which then help in restricting down his access to only at that service level in question.

upvoted 1 times

🗨️ 👤 **awssecuritynewbie** 1 year, 1 month ago

Selected Answer: C

he Answer should be C, creating a inline does not deny him access to everything else and it also makes it harder to manager and scale.

upvoted 1 times

🗨️ 👤 **awssecuritynewbie** 1 year, 1 month ago

the Answer should be C, creating a inline does not deny him access to everything else and it also makes it harder to manager and scale.

upvoted 1 times

🗨️ 👤 **YR4591** 1 year, 1 month ago

Selected Answer: B

B is right

upvoted 2 times

🗨️ 👤 **kejam** 1 year, 1 month ago

Answer B

https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies.html#policies_bound

upvoted 3 times

🗨️ 👤 **100fold** 1 year, 2 months ago

Selected Answer: B

Answer B

upvoted 2 times

🗨️ 👤 **kk2000** 1 year, 2 months ago

B is the correct answer

upvoted 2 times

A company manages multiple AWS accounts using AWS Organizations. The company's security team notices that some member accounts are not sending AWS CloudTrail logs to a centralized Amazon S3 logging bucket. The security team wants to ensure there is at least one trail configured for all existing accounts and for any account that is created in the future.

Which set of actions should the security team implement to accomplish this?

- A. Create a new trail and configure it to send CloudTrail logs to Amazon S3. Use Amazon EventBridge to send notification if a trail is deleted or stopped.
- B. Deploy an AWS Lambda function in every account to check if there is an existing trail and create a new trail, if needed.
- C. Edit the existing trail in the Organizations management account and apply it to the organization.
- D. Create an SCP to deny the cloudtrail:Delete* and cloudtrail:Stop* actions. Apply the SCP to all accounts.

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **100fold** Highly Voted 🍌 1 year, 2 months ago

Selected Answer: C

Answer C

<https://docs.aws.amazon.com/awsccloudtrail/latest/userguide/creating-trail-organization.html>

upvoted 8 times

🗳️ 👤 **Raphaello** Most Recent 🕒 10 months, 3 weeks ago

Selected Answer: C

OrganizationTrail

upvoted 1 times

🗳️ 👤 **Daniel76** 1 year ago

Selected Answer: C

Only C mention the use of management account.

By editing an existing trail in their account, and apply it to an organization, an organization trail is ready to log events for the management account and all member accounts in the organization. This do away with the need to manually create and monitor trail in every accounts.

upvoted 3 times

A company recently had a security audit in which the auditors identified multiple potential threats. These potential threats can cause usage pattern changes such as DNS access peak, abnormal instance traffic, abnormal network interface traffic, and unusual Amazon S3 API calls. The threats can come from different sources and can occur at any time. The company needs to implement a solution to continuously monitor its system and identify all these incoming threats in near-real time.

Which solution will meet these requirements?

- A. Enable AWS CloudTrail logs, VPC flow logs, and DNS logs. Use Amazon CloudWatch Logs to manage these logs from a centralized account.
- B. Enable AWS CloudTrail logs, VPC flow logs, and DNS logs. Use Amazon Macie to monitor these logs from a centralized account.
- C. Enable Amazon GuardDuty from a centralized account. Use GuardDuty to manage AWS CloudTrail logs, VPC flow logs, and DNS logs.
- D. Enable Amazon Inspector from a centralized account. Use Amazon Inspector to manage AWS CloudTrail logs, VPC flow logs, and DNS logs.

Suggested Answer: C

Community vote distribution

C (100%)

100fold Highly Voted 1 year, 8 months ago

Selected Answer: C

Answer C

upvoted 6 times

c6ed25a Most Recent 3 months, 1 week ago

Selected Answer: C

Guard Duty is used to analysis of VPC flow logs, cloud trail logs and dns logs

upvoted 1 times

723993f 7 months ago

Selected Answer: C

Absolutely hate the wording here, unnecessary confusion. Guarddduty doesn't "manage" logs, but yes its C.

upvoted 1 times

[Removed] 9 months, 4 weeks ago

Amazon GuardDuty is a fully managed threat detection service that continuously monitors your AWS accounts and workloads for malicious activity and unauthorized behavior. It uses machine learning, anomaly detection, and integrated threat intelligence to identify and prioritize potential threats.

Near-Real-Time Monitoring: GuardDuty is designed to analyze data sources like AWS CloudTrail logs, VPC flow logs, and DNS logs in near-real-time. It can detect unusual patterns such as DNS access peaks, abnormal instance traffic, abnormal network interface traffic, and unusual Amazon S3 API calls, which align with the threats identified during the security audit.

Answer : C

upvoted 3 times

Raphaello 1 year, 4 months ago

Selected Answer: C

GuardDuty (C)

upvoted 1 times

smanzana 1 year, 5 months ago

C- near-real time -> GuardDuty

upvoted 1 times

Daniel76 1 year, 7 months ago

Selected Answer: C

Guarddduty draw data sources from: AWS CloudTrail logs, VPC flow logs, and DNS logs

https://docs.aws.amazon.com/guardduty/latest/ug/guardduty_data-sources.html

Only GuardDuty detects abnormal and unusual activities among all choices.

upvoted 4 times

Aamee 1 year, 7 months ago

Selected Answer: C

Monitoring threats, abnormal traffic etc always leads towards GuardDuty.
upvoted 2 times

A company that uses AWS Organizations is using AWS IAM Identity Center (AWS Single Sign-On) to administer access to AWS accounts. A security engineer is creating a custom permission set in IAM Identity Center. The company will use the permission set across multiple accounts. An AWS managed policy and a customer managed policy are attached to the permission set. The security engineer has full administrative permissions and is operating in the management account.


When the security engineer attempts to assign the permission set to an IAM Identity Center user who has access to multiple accounts, the assignment fails.

What should the security engineer do to resolve this failure?

- A. Create the customer managed policy in every account where the permission set is assigned. Give the customer managed policy the same name and same permissions in each account.
- B. Remove either the AWS managed policy or the customer managed policy from the permission set. Create a second permission set that includes the removed policy. Apply the permission sets separately to the user.
- C. Evaluate the logic of the AWS managed policy and the customer managed policy. Resolve any policy conflicts in the permission set before deployment.
- D. Do not add the new permission set to the user. Instead, edit the user's existing permission set to include the AWS managed policy and the customer managed policy.

Suggested Answer: A

Community vote distribution



pupsik Highly Voted 1 year, 8 months ago

Selected Answer: A

<https://docs.aws.amazon.com/singlesignon/latest/userguide/permissionsetcustom.html#:~:text=When%20you%20create%20a%20permission%20set%20wi>
upvoted 8 times

[Removed] Highly Voted 1 year, 7 months ago

Selected Answer: A

Give this a read y'all. Answer indeed is A. You must create the CMP in each account unlike the AWS Managed Policies

<https://aws.amazon.com/blogs/security/how-to-use-customer-managed-policies-in-aws-single-sign-on-for-advanced-use-cases/#:~:text=Configure%20an%20IAM%20Identity%20Center%20permission%20set%20to%20use%20a%20CMP>
upvoted 7 times

IPLogic Most Recent 6 months, 4 weeks ago

Selected Answer: A

To resolve the issue of the permission set assignment failure in AWS IAM Identity Center, the security engineer should follow option A: Create the customer managed policy in every account where the permission set is assigned. Give the customer managed policy the same name and same permissions in each account.

This approach ensures that the customer managed policy is available in all the accounts where the permission set is being assigned, which is necessary for the assignment to succeed. By having the same policy name and permissions in each account, the security engineer can ensure consistency and avoid any conflicts that might arise from missing or mismatched policies
upvoted 1 times

TenaciousD 10 months, 2 weeks ago

Definitely A - Key word = 'custom permission set' - <https://docs.aws.amazon.com/singlesignon/latest/userguide/permissionsetcustom.html>
upvoted 1 times

arvehisa 1 year, 3 months ago

Selected Answer: A

The correct answer is A.

<https://docs.aws.amazon.com/singlesignon/latest/userguide/permissionsetcustom.html#permissionsetscmpconcept>

When you create a permission set with a customer managed policy, you must create an IAM policy with the same name and path in each AWS account where IAM Identity Center assigns your permission set.

upvoted 1 times

🗨️ 👤 **walter_white_008** 1 year, 4 months ago

Selected Answer: C

Answer is C.

Dont blindly accept the answer selected by most of the people, it may be wrong sometimes.

<https://docs.aws.amazon.com/singlesignon/latest/userguide/permissionsetsconcept.html>

IAM Identity Center assigns access to a user or group in one or more AWS accounts with permission sets. When you assign a permission set, IAM Identity Center creates corresponding IAM Identity Center-controlled IAM roles in each account, and attaches the policies specified in the permission set to those roles.

Solve the policy conflicts as per option C and you are good.

upvoted 1 times

🗨️ 👤 **2f5c7cd** 1 month, 4 weeks ago

Answer is still A using the link you provided. Go to the link Custom permissions for AWS managed and customer managed and in the section Customer managed policies

"...When you create a permission set with a customer managed policy, you must create an IAM policy with the same name and path in each AWS account where IAM Identity Center assigns your permission set. If you are specifying a custom path, make sure to specify the same path in each AWS account..."

upvoted 1 times

🗨️ 👤 **NoCrapEva** 1 year, 4 months ago

Selected Answer: A

Answer C - does not resolve the failure - It will only highlight where the issue is..

<https://docs.aws.amazon.com/singlesignon/latest/userguide/permissionsetcustom.html>

"You can attach customer managed policies to your permission set. Customer managed policies are IAM policies in your account that you create and maintain. In contrast, AWS managed policies are IAM policies in your account that AWS maintains. You can assign an customer managed policy as permissions for the role that IAM Identity Center creates, or as a permissions boundary.

When you create a permission set with a customer managed policy, you must create an IAM policy with the same name and path in each AWS account where IAM Identity Center assigns your permission set. If you are specifying a custom path, make sure to specify the same path in each AWS account. For more information, see Friendly names and paths in the IAM User Guide. IAM Identity Center attaches the IAM policy to the IAM role that it creates in your AWS account."

upvoted 1 times

🗨️ 👤 **Raphaello** 1 year, 4 months ago

Selected Answer: A

Identity Center's permission set actually creates IAM role in the target (member) AWS accounts. Therefore, when you include a customer managed policy into a permission set, you need to make sure that the member accounts recognize the customer managed policy, by creating the policy and giving it same name in every AWS member account.

Answer is A.

upvoted 1 times

🗨️ 👤 **mynickc** 1 year, 5 months ago

Selected Answer: C

Answer is C. Because, when you assign a permissionset via the identity center; it automatically creates IAM controlled role in all the org.

<https://docs.aws.amazon.com/singlesignon/latest/userguide/permissionsetsconcept.html>

upvoted 1 times

🗨️ 👤 **Daniel76** 1 year, 6 months ago

Selected Answer: A

<https://docs.aws.amazon.com/singlesignon/latest/userguide/howtocmp.html>

"Before you assign your permission set with IAM policies, you must prepare your member account. The name of an IAM policy in your member

account must be a case-sensitive match to name of the policy in your management account. IAM Identity Center fails to assign the permission set if the policy doesn't exist in your member account."

upvoted 3 times

  **AWSvad** 1 year, 7 months ago

The correct answer is:

C. Evaluate the logic of the AWS managed policy and the customer managed policy. Resolve any policy conflicts in the permission set before deployment.

In this scenario, the assignment of the permission set to an IAM Identity Center user is failing, indicating a potential conflict between the AWS managed policy and the customer managed policy. It is important to review and evaluate the logic of both policies and resolve any conflicts before deploying the permission set.



Options A and B suggest alternative actions but do not directly address the issue of policy conflicts. Option A involves creating the customer managed policy in every account, which may not resolve the underlying problem. Option B suggests removing either the AWS managed policy or the customer managed policy, which may not be the most appropriate solution.

Option D suggests editing the user's existing permission set, but it does not address the potential conflicts between the AWS managed policy and the customer managed policy.

Therefore, option C is the most appropriate choice to resolve the issue by thoroughly evaluating and resolving policy conflicts in the permission set before deployment.

- ChatGPT

upvoted 2 times

  **alexleely** 1 year, 5 months ago

The correct answer is:

A. Create the customer managed policy in every account where the permission set is assigned. Give the customer managed policy the same name and same permissions in each account.

Explanation:



When using IAM Identity Center (AWS Single Sign-On) to administer access to AWS accounts across multiple accounts, and attaching a customer managed policy to a permission set, it's essential to create the corresponding IAM policy with the same name and permissions in each AWS account where the permission set is assigned. This ensures consistency and avoids issues during the assignment process.

Option A aligns with this requirement by recommending the creation of the customer managed policy in every account where the permission set is assigned, with the same name and permissions. This approach helps in maintaining uniformity across accounts and resolving the assignment failure.

Options B, C, and D do not directly address the need to create the customer managed policy in each account or ensure consistency across accounts, making option A the appropriate solution in this scenario.

- ChatGPT

upvoted 1 times

  **kejam** 1 year, 7 months ago

Selected Answer: C

Answer C

Not A: AWS IAM Identity Center enables you to centrally manage permissions across multiple AWS accounts without configuring each account manually.

<https://docs.aws.amazon.com/singlesignon/latest/userguide/manage-your-accounts.html>

Not B: You can assign more than one permission set to a user.

<https://docs.aws.amazon.com/singlesignon/latest/userguide/permissionsets.html>

Not D: A custom permission set can use up to 10 AWS managed or customer managed policies.

<https://docs.aws.amazon.com/singlesignon/latest/userguide/howtocreatepermissionset.html>

upvoted 1 times

  **ahrentom** 1 year, 8 months ago

Selected Answer: C

I would go with C here, seems to be the most logical answer

upvoted 3 times

  **ahrentom** 1 year, 7 months ago

have to correct me, the right one here is A

upvoted 3 times

  **mynickc** 1 year, 5 months ago

when you're using permissionset you don't need to create the customer managed policy in every org manually. so its C

upvoted 1 times

A company has thousands of AWS Lambda functions. While reviewing the Lambda functions, a security engineer discovers that sensitive information is being stored in environment variables and is viewable as plaintext in the Lambda console. The values of the sensitive information are only a few characters long.

What is the MOST cost-effective way to address this security issue?

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

Suggested Answer: D

Community vote distribution

D (86%)

14%

🗳️ 👤 **kejam** Highly Voted 1 year, 1 month ago

Selected Answer: D

Answer D

There is no charge from Parameter Store to create a SecureString parameter.

<https://docs.aws.amazon.com/systems-manager/latest/userguide/systems-manager-parameter-store.html#what-is-a-parameter>

upvoted 5 times

🗳️ 👤 **100fold** Highly Voted 1 year, 2 months ago

Selected Answer: D

Store the environment variables as secure strings in Parameter Store. Most cost-effective way.

upvoted 5 times

🗳️ 👤 **Raphaello** Most Recent 10 months, 3 weeks ago

Selected Answer: D

SSM Parameter Store standard is the most cost-effective solution.

upvoted 1 times

🗳️ 👤 **trashbox** 1 year ago

Exam on 2023-12-18

upvoted 1 times

🗳️ 👤 **[Removed]** 1 year, 1 month ago

Selected Answer: D

cost efficient.

upvoted 1 times

🗳️ 👤 **stream3652** 1 year, 2 months ago

Selected Answer: C

Isn't C more secure?

upvoted 2 times

🗳️ 👤 **AgboolaKun** 1 year, 1 month ago

That is a good question. However, the emphasis here is MOST cost-effective. Using secure string in Parameter Store is free, hence D is cost effective.

upvoted 6 times

🗳️ 👤 **M2ao** 1 year, 2 months ago

the question is asking for cost-effective way

upvoted 3 times

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?


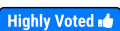
- A. Use AWS IAM Access Analyzer to analyze the policies. View the findings from policy validation checks.
- B. Review AWS Trusted Advisor checks for all accounts in the organization.
- C. Set up AWS Audit Manager. Run an assessment for all AWS Regions for all accounts.
- D. Ensure that Amazon Inspector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

Suggested Answer: A

Community vote distribution

A (90%)

10%

 **bannium**  1 year, 2 months ago

Selected Answer: A

It seems IAM Access Analyzer

> This powerful new feature will help you to construct IAM policies and SCPs that take advantage of time-tested AWS best practices.

<https://aws.amazon.com/jp/blogs/aws/iam-access-analyzer-update-policy-validation/>

upvoted 5 times

 **Raphaello**  10 months, 3 weeks ago

Selected Answer: A

IAM Access Analyzer.

upvoted 1 times

 **trashbox** 1 year ago

Exam on 2023-12-18

upvoted 1 times


 **AgboolaKun** 1 year, 1 month ago

Selected Answer: A

You can create AWS IAM Access Analyzer in AWS Organizations as the zone of trust.

<https://aws.amazon.com/blogs/aws/new-use-aws-iam-access-analyzer-in-aws-organizations/>

upvoted 2 times

 **pupsik** 1 year, 2 months ago

Selected Answer: A

I don't think Trusted Advisor would give insights regarding SCP configurations, hence A.

upvoted 1 times

 **ahrentom** 1 year, 2 months ago

Selected Answer: B

best practices --> Trusted Advisor

upvoted 1 times

 **Aamee** 1 year ago

No, cuz Trusted Advisor doesn't analyze anything with the SCP level so hence A is correct here.

upvoted 1 times


A company uses Amazon RDS for MySQL as a database engine for its applications. A recent security audit revealed an RDS instance that is not compliant with company policy for encrypting data at rest. A security engineer at the company needs to ensure that all existing RDS databases are encrypted using server-side encryption and that any future deviations from the policy are detected.

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

- A. Create an AWS Config rule to detect the creation of unencrypted RDS databases. Create an Amazon EventBridge rule to trigger on the AWS Config rules compliance state change and use Amazon Simple Notification Service (Amazon SNS) to notify the security operations team.
- B. Use AWS System Manager State Manager to detect RDS database encryption configuration drift. Create an Amazon EventBridge rule to track state changes and use Amazon Simple Notification Service (Amazon SNS) to notify the security operations team.
- C. Create a read replica for the existing unencrypted RDS database and enable replica encryption in the process. Once the replica becomes active, promote it into a standalone database instance and terminate the unencrypted database instance.
- D. Take a snapshot of the unencrypted RDS database. Copy the snapshot and enable snapshot encryption in the process. Restore the database instance from the newly created encrypted snapshot. Terminate the unencrypted database instance.
- E. Enable encryption for the identified unencrypted RDS instance by changing the configurations of the existing database.

Suggested Answer: AD

Community vote distribution



ahrentom Highly Voted 1 year, 2 months ago

Selected Answer: AD

A and D, here's another source

<https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/automatically-remediate-unencrypted-amazon-rds-db-instances-and-clusters.html>

upvoted 10 times

Zek Most Recent 7 months, 3 weeks ago

AD - Use AWS Config to track configuration changes and then take snapshot to enable encryption

upvoted 1 times

Boul 9 months, 2 weeks ago

AD

C not correct because you can't have an encrypted read replica of an unencrypted DB instance or an unencrypted read replica of an encrypted DB instance.

upvoted 3 times

Noexperience 10 months ago

Selected Answer: AC

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

Addresses Existing Issue: This directly resolves the compliance issue of the unencrypted RDS instance. It ensures data-at-rest encryption while minimizing application downtime.

Minimizes Disruption: Replicating with encryption avoids modifying the primary database directly, reducing the risk of application issues.

upvoted 1 times

Raphaello 10 months, 3 weeks ago

Selected Answer: AD

AD

Config to track drift, and taking snapshot to encrypt current RDS.

upvoted 1 times

trashbox 1 year ago

Exam on 2023-12-18

upvoted 1 times

  **100fold** 1 year, 2 months ago

Selected Answer: AD

Answer is AD.


<https://www.examttopics.com/discussions/amazon/view/60595-exam-aws-certified-security-specialty-topic-1-question-275/>
upvoted 2 times

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

Suggested Answer: C

Community vote distribution



Olaunfazed 1 year ago

Answer is D The best solution is to use AWS Backup. It allows you to copy EBS snapshots to Amazon S3 and provides built-in features like S3 Object Lock to prevent deletion of snapshots. By using AWS Backup, you can automate the process and ensure reliable disaster recovery even in case of account compromise
upvoted 3 times

urbanmonk 6 months, 1 week ago

As simplistic and plausible as Option D sounds, it is not the solution when you reread the question.

REQUIREMENTS: "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"

The question explicitly asks about the solution given a scenario that the "account is compromised". Hence, you NEED a backup/new account to remedy the situation. Correct answer is C
upvoted 3 times

lovekiller 8 months, 3 weeks ago

True, previously this was not an option but was a recent addition to AWS Backup.

Source: <https://aws.amazon.com/getting-started/hands-on/amazon-ebs-backup-and-restore-using-aws-backup/>

upvoted 1 times

Raphaello 1 year, 4 months ago

Selected Answer: A

Option A seems very good solution to me!

C is a fine solution, but why not A? What makes A less appealing than C?!!!

In fact, using Glacier Vault Lock is the ONLY way to protect against data deletion, and even after moving snapshots/backups to a different account, Glacier Vault Lock would be required to protect against data deletion from the new account.
upvoted 1 times

Raphaello 1 year, 4 months ago

Sorry, that's the answer for a different question.

For this one, C is the best option
upvoted 1 times

🗨️ 👤 **walter_white_008** 1 year, 4 months ago

Selected Answer: C

C makes sense.

upvoted 1 times

🗨️ 👤 **Raphaello** 1 year, 4 months ago

Selected Answer: C

This is a bit vague.

1. If the fear to lose account A, and subsequently the encrypted snapshots, that would apply to KMS keys used for snapshot encryption.
2. A solution to backup the encrypted snapshots to a different account, B, has to include creating new KMS key in account B, and not just access to KMS key in account A, cause it is subject to the fear of being compromised as well.
3. Answer C is the only one that taking KMS key into consideration, even if not in an ideal way. I would go with C only for that fact, and it mentioned a new account.

upvoted 2 times

🗨️ 👤 **Th3Dud3** 1 year, 5 months ago

c. You can add a vault lock to your AWS Backup Vault. So no need to use S3 object lock.

upvoted 1 times

🗨️ 👤 **confusedyeti69** 1 year, 6 months ago

How compromise is compromised? You wouldn't have access to KMS if you choose C and your snapshots are in the same account if you choose D.

upvoted 1 times

🗨️ 👤 **Aamee** 1 year, 6 months ago

This statement in option C "Allow the new account to access the KMS key that encrypts the EBS snapshots" clearly means that when you're creating a new account for a backup solution, you also have the appropriate 'Access' to encrypt and decrypt the keys as well. That's why it's further said to copy out the encrypted snapshots in the new account too for performing any future decrypt operations.

Hope it helps..

upvoted 1 times

🗨️ 👤 **confusedyeti69** 1 year, 6 months ago

If you store the snapshot in account B that is encrypted with account A's key, and then lose access to the key (compromised), would you still be able to use the snapshot?

upvoted 1 times

🗨️ 👤 **Daniel76** 1 year, 6 months ago

The key is in the KMS and account B has access to it.

if account A is gone, account B can still decrypt the snapshot, provided the account A did not have the right to delete this key in the KMS..

upvoted 1 times

🗨️ 👤 **Daniel76** 1 year, 6 months ago

Selected Answer: C

You cant use D because the snapshot can still be deleted even if under compliance mode, if the compromised AWS account is deleted.

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock.html>

upvoted 2 times

🗨️ 👤 **awssecuritynewbie** 1 year, 7 months ago

but if the AWS account is compromised and they are worried about the AWS account being deleted...so it will lose the KMS key as well. It is a tricky question the S3 lock is not enough because you will lose the KMS keys.. but it should have a solution to copy the keys into the new account as well.

upvoted 1 times

🗨️ 👤 **Boul** 1 year, 3 months ago

KMS is external to the accounts. It's either you have access to the service or not, regardless of the account.

upvoted 1 times

🗨️ 👤 **NoCrapEva** 1 year, 4 months ago

The question only says Account compromised (not deleted)... But the question specifically asks you "...to implement a solution so that the company can recover the EC2 instances" IF ..."the EBS snapshots are deleted"

Therefore only Answer C will allow this

upvoted 1 times

🗨️ 👤 **AgboolaKun** 1 year, 7 months ago

Selected Answer: C



C is the correct answer to me.

New AWS account with limited privileges - prevents the account from being compromised

Access to AWS KMS key - access to the key to decrypt data in the recovery account.

Copy snapshots to the recovery account (new account) on a recurring basis - This could be using AWS Backup as well or any other services.

upvoted 2 times

  **pupsik** 1 year, 8 months ago

Selected Answer: D

"Use S3 Object Lock to prevent deletion of the snapshots." makes this option very viable, even if account gets compromised.

upvoted 2 times

  **bannium** 1 year, 8 months ago

How I can export EBS Snapshot data to S3 bucket using AWS Backup?

upvoted 1 times

  **AgboolaKun** 1 year, 7 months ago

The only concern I have with D is that there is no mention of how to access the AWS KMS CMK key used for the encryption of EBS snapshots.

Therefore, I will go for C.

upvoted 2 times


  **confusedyeti69** 1 year, 6 months ago

It is creating a snapshot and storing it in S3 of the same account, there is no need for any KMS policy to be explicitly mentioned in the answer.

But like another previous comment mentioned, it would be better to backup the keys as well if storing the backup snapshot in another account.

I vote D as answer.

upvoted 1 times

  **100fold** 1 year, 8 months ago

Selected Answer: C

Answer C. The wording is rearranged, but same answer selections.

<https://www.examttopics.com/discussions/amazon/view/69464-exam-aws-certified-security-specialty-topic-1-question-315/>

upvoted 4 times

A company's security engineer is designing an isolation procedure for Amazon EC2 instances as part of an incident response plan. The security engineer needs to isolate a target instance to block any traffic to and from the target instance, except for traffic from the company's forensics team. Each of the company's EC2 instances has its own dedicated security group. The EC2 instances are deployed in subnets of a VPC. A subnet can contain multiple instances.

The security engineer is testing the procedure for EC2 isolation and opens an SSH session to the target instance. The procedure starts to simulate access to the target instance by an attacker. The security engineer removes the existing security group rules and adds security group rules to give the forensics team access to the target instance on port 22.


After these changes, the security engineer notices that the SSH connection is still active and usable. When the security engineer runs a ping command to the public IP address of the target instance, the ping command is blocked.

What should the security engineer do to isolate the target instance?

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

Suggested Answer: B

Community vote distribution



AgboolaKun Highly Voted 1 year, 7 months ago

Selected Answer: B

There is no need for SSH port 22 since Systems Manager Session Manager can give the necessary access that the security team needs to the EC2 instances.

upvoted 10 times

b5f86df Most Recent 3 weeks, 4 days ago

Selected Answer: B

Systems Manager Session Manager to connect without open 22 port

upvoted 1 times

maneeshsharma 2 months ago

Selected Answer: D

Correct Answer: D

Explanation: To effectively isolate the target EC2 instance, the security engineer should use an AWS Systems Manager (SSM) document to apply host-level firewall rules that block all inbound and outbound traffic, except for traffic from the forensics team. This approach ensures the existing SSH connection is terminated and provides granular control over the instance's traffic without affecting other instances in the subnet. The SSM document can be executed if the instance has the SSM agent installed and the appropriate IAM permissions, which is a common setup in AWS environments.

upvoted 1 times

Kaps443 2 months, 2 weeks ago

Selected Answer: B

B is correct.

upvoted 1 times

fcbflo 4 months ago

Selected Answer: B

The scenario specifically requires isolating a single target instance while maintaining forensic team access. Using a NACL would isolate all instances in the subnet, potentially disrupting other operations. (NACLs effect is on subnet borders - SGs on ENI)

Option B is more precise because:

It removes direct SSH access (port 22) which eventually terminates the existing connection

It provides an alternative access method (Systems Manager Session Manager) that doesn't rely on direct network access

It targets only the specific instance without affecting others

So while your point about stateful vs. stateless is technically accurate, in this specific scenario, the subnet-wide impact of NACLs makes option B more appropriate for the requirement to isolate only the target instance.

upvoted 1 times

🗳️ 👤 **AWSLoverLoverLoverLoverLover** 4 months, 1 week ago

Selected Answer: C

C. Create a network ACL that is associated with the target instance's subnet. Add a rule at the top of the inbound rule set to deny all traffic from 0.0.0.0/0. Add a rule at the top of the outbound rule set to deny all traffic to 0.0.0.0/0.

Explanation:

Why is the SSH session still active?

Security group changes do not affect existing connections (i.e., established TCP sessions remain open).

This is why the SSH session remains active, but new inbound connections (like ping/ICMP) are blocked.

upvoted 1 times

🗳️ 👤 **FlyingHawk** 4 months, 3 weeks ago

Selected Answer: C

Updating the security group of an EC2 instance does not affect existing connections because security groups are stateful. This means that if a connection is already established (e.g., an active SSH session), changes to the security group rules will not terminate that connection. The rules only apply to new connections.

This is why the security engineer's SSH session remained active even after modifying the security group rules. However, the ping command was blocked because the security group changes applied to new traffic (ICMP traffic for ping).

upvoted 2 times

🗳️ 👤 **FlyingHawk** 4 months, 3 weeks ago

A network ACL (Access Control List) can immediately affect current connections if you modify its rules while traffic is flowing, as it acts as a virtual firewall at the subnet level, meaning any changes to its allow/deny rules can instantly impact connectivity to resources within that subnet, potentially disrupting existing connections if not configured carefully.

upvoted 2 times

🗳️ 👤 **Wardove** 5 months, 1 week ago

Selected Answer: C

Only this option terminates active ssh sessions for 100% isolation

upvoted 1 times

🗳️ 👤 **TareDHakim** 5 months, 1 week ago

Selected Answer: B

C is wrong, NACL changes will impact entire subnet and all instances attached to it. to target the single instance, B makes most sense.

upvoted 2 times

🗳️ 👤 **IPLogic** 6 months, 3 weeks ago

Selected Answer: B

Bis correct as NACL will block traffic to all instances within the subnet.

At instance level right approach is option B.

upvoted 1 times

🗳️ 👤 **IPLogic** 6 months, 4 weeks ago

Selected Answer: C

Simply put Option B. does not block s other traffic and isolate the Instance completely.

upvoted 1 times

🗳️ 👤 **IPLogic** 7 months ago

Selected Answer: C

C. Create a network ACL that is associated with the target instance's subnet. Add a rule at the top of the inbound rule set to deny all traffic from 0.0.0.0/0. Add a rule at the top of the outbound rule set to deny all traffic to 0.0.0.0/0.

<https://docs.aws.amazon.com/whitepapers/latest/aws-security-incident-response-guide/destination-containment.html>

upvoted 1 times

🗳️ 👤 **mzeynalli** 7 months, 3 weeks ago

Selected Answer: C

Guys, keep it simple. The task is "What should the security engineer do to isolate the target instance?"

Adding deny rules for both inbound and outbound traffic ensures that all communications to and from the instance are blocked, effectively isolating the instance.

After the completion of the REQUIRED task, the engineer will allow the forensic team to add either AWS Systems Manager Session Manager or NACL to accomplish the SECOND part of the task.

upvoted 4 times

🗨️ 👤 **urbanmonk** 6 months, 1 week ago

Speaking of REQUIRED - "The security engineer needs to isolate a target instance to block any traffic to and from the target instance, except for traffic from the company's forensics team."

The question seeks for a solution that will isolate the target instance but also allow access to the Forensic Team. Option C offers the isolation but does not address the required access for the company's forensic team. We are not allowed to make assumptions in these exam questions. So B is the correct answer as it isolates the target and also provides access to the forensic team using SMSM.

upvoted 3 times

🗨️ 👤 **jakie22332** 7 months, 3 weeks ago

Selected Answer: C

C is correct

upvoted 2 times

🗨️ 👤 **Johnyw** 8 months, 3 weeks ago

Answer is C.

The question is definitely trying to test the difference between an SG and a NACL, the System Manager Session Manager is just a distraction. The engineer just need to implement a NACL to drop the sessions. the sessions are allowed because SG's are stateful.

upvoted 2 times

🗨️ 👤 **Just_Ninja** 9 months, 2 weeks ago

Selected Answer: B

B is correct, but if you still have an active session on port 22, the connection will remain open. You can easily test this. Security Groups (SGs) are stateful, which means an existing session remains active even if the rule is removed.

Launch an EC2 instance with SSH access restricted to your IP.

Connect to the instance via SSH.

Remove the SSH rule from the Security Group.

Your active session will continue without interruption! 😊

Also, from my point of view, AWS Systems Manager (SSM) is one of the coolest tools. You can easily log session data to an S3 bucket for reviews and analytics, making it a powerful tool for auditing and security operations.

upvoted 1 times

🗨️ 👤 **jamesf** 10 months ago

Selected Answer: B

keywords: "The security engineer removes the existing security group rules and adds security group rules to give the forensics team access to the target instance on port 22."

So just remove the Security Group SSH Port 22 and use AWS Systems Manager Session Manager

upvoted 1 times

A startup company is using a single AWS account that has resources in a single AWS Region. A security engineer configures an AWS CloudTrail trail in the same Region to deliver log files to an Amazon S3 bucket by using the AWS CLI. Because of expansion, the company adds resources in multiple Regions. The security engineer notices that the logs from the new Regions are not reaching the S3 bucket. What should the security engineer do to fix this issue with the LEAST amount of operational overhead?

- A. Create a new CloudTrail trail. Select the new Regions where the company added resources.
- B. Change the S3 bucket to receive notifications to track all actions from all Regions.
- C. Create a new CloudTrail trail that applies to all Regions.
- D. Change the existing CloudTrail trail so that it applies to all Regions.

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **kk2000** Highly Voted 1 year, 8 months ago

Correct Answer is D

<https://docs.aws.amazon.com/awsccloudtrail/latest/userguide/receive-cloudtrail-log-files-from-multiple-regions.html>

upvoted 10 times

🗳️ 👤 **Just_Ninja** Most Recent 9 months, 2 weeks ago

Selected Answer: D

An Update is the right solution. Here is the CLI for it :-)

`aws cloudtrail update-trail --name my-trail --is-multi-region-trail`

upvoted 3 times

🗳️ 👤 **unravikumar** 1 year, 3 months ago

Selected Answer: D

D is correct

upvoted 1 times

🗳️ 👤 **Raphaello** 1 year, 4 months ago

Selected Answer: D

Change CloudTrail trail to ne multi-region.

upvoted 1 times

🗳️ 👤 **brpjp** 1 year, 7 months ago

Question is new region trail not reaching to s3 bucket and not trail created from new region.

upvoted 1 times

🗳️ 👤 **smanzana** 1 year, 5 months ago

But each trail is independent and its region cannot be changed after the initial configuration...and the answer D says "change the trail" not "create a new trail"

upvoted 1 times

🗳️ 👤 **Aamee** 1 year, 6 months ago

Question specifically emphasis on the solution with 'LEAST amount of operational overhead' which is doable only through option D. All other options still involves some kind of operational overhead. Hope it helps..

upvoted 2 times

🗳️ 👤 **ahrentom** 1 year, 8 months ago

Selected Answer: D

go with D

upvoted 2 times

🗳️ 👤 **100fold** 1 year, 8 months ago

Selected Answer: D

Agree answer is D. Change the existing CloudTrail using AWS CLI, add the --is-multi-region-trail option to the update-trail command.

upvoted 4 times

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

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

Suggested Answer: BC

Community vote distribution

BC (100%)

100fold Highly Voted 1 year, 8 months ago

Selected Answer: BC

Answer is BC.

<https://www.examttopics.com/discussions/amazon/view/88447-exam-aws-certified-security-specialty-topic-1-question-437/>
upvoted 6 times

FunkyFresco Most Recent 10 months, 1 week ago

Selected Answer: BC

B and C.

upvoted 1 times

xusang 1 year, 4 months ago

Selected Answer: BC

Restricting access to Application Load Balancers <https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/restrict-access-to-load-balancer.html>
upvoted 3 times

Raphaello 1 year, 4 months ago

Selected Answer: BC

BC

Add custom origin-request header (CloudFront > ALB), set ALB to only accept request with such HTTP header.
upvoted 1 times

Aamee 1 year, 6 months ago

Selected Answer: BC

W/o any doubt..

upvoted 1 times

Daniel76 1 year, 6 months ago

Selected Answer: BC

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/example-function-add-true-client-ip-header.html>
<https://aws.amazon.com/blogs/security/three-most-important-aws-waf-rate-based-rules/>
upvoted 1 times

A company discovers a billing anomaly in its AWS account. A security consultant investigates the anomaly and discovers that an employee who left the company 30 days ago still has access to the account. The company has not monitored account activity in the past. The security consultant needs to determine which resources have been deployed or reconfigured by the employee as quickly as possible. Which solution will meet these requirements?

- A. In AWS Cost Explorer, filter chart data to display results from the past 30 days. Export the results to a data table. Group the data table by resource.
- B. Use AWS Cost Anomaly Detection to create a cost monitor. Access the detection history. Set the time frame to Last 30 days. In the search area, choose the service category.
- C. In AWS CloudTrail, filter the event history to display results from the past 30 days. Create an Amazon Athena table that contains the data. Partition the table by event source.
- D. Use AWS Audit Manager to create an assessment for the past 30 days. Apply a usage-based framework to the assessment. Configure the assessment to assess by resource.

Suggested Answer: C

Community vote distribution

C (100%)

100fold Highly Voted 1 year, 8 months ago

Selected Answer: C

Answer C. CloudTrail for as quickly as possible. Look up events related to the creation, modification, or deletion of resources in your AWS account.

AWS Cost Anomaly Detection: For a new service subscription, 10 days of historical service usage data is needed before anomalies can be detected for that service. If you create a new monitor, it can take up to 24 hours to begin detecting new anomalies.

upvoted 9 times

FunkyFresco Most Recent 10 months, 1 week ago

Selected Answer: C

Option C, cloud trail.

upvoted 1 times

Raphaello 1 year, 4 months ago

Selected Answer: C

To investigate a certain user and find what resources that IAM user created over past period, CloudTrail is the tool to use.

upvoted 1 times

awssecuritynewbie 1 year, 4 months ago

Selected Answer: C

I think C is the best answer but Athena is over kill

upvoted 1 times

WeepingMaple 1 year, 6 months ago

Think only Cloudtrail records down resources reconfigured by the employee.

upvoted 1 times

Daniel76 1 year, 6 months ago

Selected Answer: C

The investigation was triggered by cost anomaly but that is not the only concern. The security engineer needs to find out what has been deployed as well as reconfigured, so AWS Cost explorer or Anomaly detection will not do the job. Only CloudTrail and Anthena will be the most effective method. Cost should have nothing to do with compliance so audit manager will not help.

upvoted 3 times



mnsait 2 months, 4 weeks ago

Tricky question and a smart answer. Thank you @Daniel76. Nicely explained.

upvoted 1 times

Ernestokoro 1 year, 7 months ago

Ans is B! Please see link below: <https://docs.aws.amazon.com/cost-management/latest/userguide/getting-started-ad.html>
upvoted 1 times

  **Aamee** 1 year, 7 months ago

Don't think it can be B. See this comment above:

"AWS Cost Anomaly Detection: For a new service subscription, 10 days of historical service usage data is needed before anomalies can be detected for that service. If you create a new monitor, it can take up to 24 hours to begin detecting new anomalies."

upvoted 2 times

A security engineer is checking an AWS CloudFormation template for vulnerabilities. The security engineer finds a parameter that has a default value that exposes an application's API key in plaintext. The parameter is referenced several times throughout the template. The security engineer must replace the parameter while maintaining the ability to reference the value in the template.

Which solution will meet these requirements in the MOST secure way?

- A. Store the API key value as a SecureString parameter in AWS Systems Manager Parameter Store. In the template, replace all references to the value with `{{resolve:ssm:MySSMParameterName:1}}`.
- B. Store the API key value in AWS Secrets Manager. In the template, replace all references to the value with `{{resolve:secretsmanager:MySecretId:SecretString}}`.
- C. Store the API key value in Amazon DynamoDB. In the template, replace all references to the value with `{{resolve:dynamodb:MyTableName:MyPrimaryKey}}`.
- D. Store the API key value in a new Amazon S3 bucket. In the template, replace all references to the value with `{{resolve:s3:MyBucketName:MyObjectName}}`.

Suggested Answer: B

Community vote distribution

B (100%)

100fold Highly Voted 1 year, 8 months ago

Selected Answer: B

Agree answer B.

Not A. `{{resolve:ssm:MySSMParameterName:1}}`

ssm: Systems Manager Parameter Store plaintext parameter.

ssm-secure: Systems Manager Parameter Store secure string parameter.

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/dynamic-references.html>

upvoted 11 times

dexterryu Highly Voted 1 year, 6 months ago

This is a bit of a trick question. A is correct outside of the syntax in the ssm reference. Therefor B. Had it been `resolve:ssm-secure:MySSMParam` then A would be correct.

upvoted 6 times

Raphaello 1 year, 4 months ago

Spot on.

ssm: Systems Manager Parameter Store plaintext parameter.

ssm-secure: Systems Manager Parameter Store secure string parameter.

Option B is the right answer.

upvoted 2 times

d3717d5 Most Recent 2 weeks, 6 days ago

Selected Answer: B

Use secrets manager whenever you need the most secure way. Use Systems Manager Parameter store and store the value as a secure string if you need most cost effective solution.

upvoted 1 times

FunkyFresco 10 months, 1 week ago

Selected Answer: B

Option B is the right choice.

upvoted 1 times

Aamee 1 year, 6 months ago

Selected Answer: B

MOST secure way..

upvoted 3 times

🗨️ 👤 **confusedyeti69** 1 year, 6 months ago

Selected Answer: B

Secure. B

upvoted 1 times

🗨️ 👤 **kk2000** 1 year, 8 months ago

Answer: B using secrets manager.(More secure way)

https://docs.aws.amazon.com/secretsmanager/latest/userguide/cfn-example_reference-secret.html

upvoted 4 times

🗨️ 👤 **AgboolaKun** 1 year, 8 months ago

Agree. SSM can create secure API string as well but the emphasis here is on MOST secure. Therefore, Secret Manager is the answer.

upvoted 3 times

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

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

Suggested Answer: ADE

Community vote distribution

ADE (100%)

  **ahrentom**  1 year, 2 months ago

Selected Answer: ADE

ADE

<https://docs.aws.amazon.com/awsccloudtrail/latest/userguide/best-practices-security.html>



upvoted 5 times

  **d3717d5**  2 weeks, 5 days ago

Selected Answer: ADE

ADE, not F - because Access Analyzer is monitoring , will not prevent unauthorized access.

upvoted 1 times

  **youonebe** 5 months, 3 weeks ago

Selected Answer: ADE

Option B (compression) doesn't provide security benefits

Option C is reactive rather than preventive

Option F is a monitoring tool but doesn't directly prevent unauthorized access or tampering

upvoted 1 times

  **Zek** 7 months, 3 weeks ago

ADE. Agreed

upvoted 1 times

  **Raphaello** 10 months, 3 weeks ago

Selected Answer: ADE

ADE

We agree.

upvoted 1 times

  **Aamee** 1 year, 1 month ago

Selected Answer: ADE

Here's what it describes about the usage of log file integration and the SSE-KMS usecase scenario:

"If you use SSE-KMS and log file validation, and you have modified your Amazon S3 bucket policy to only allow SSE-KMS encrypted files, you will not be able to create trails that utilize that bucket unless you modify your bucket policy to specifically allow AES256 encryption, as shown in the following example policy line.

```
"StringNotEquals": { "s3:x-amz-server-side-encryption": ["aws:kms", "AES256"] }
```

"


upvoted 4 times

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. Configure the bucket to use S3 Object Lock in compliance mode. Upload the data to the bucket. Create a resource-based bucket policy that meets all the regulatory requirements.
- B. Create an Amazon S3 bucket. Configure the bucket to use S3 Object Lock in governance mode. Upload the data to the bucket. Create a user-based IAM policy that meets all the regulatory requirements.
- C. Create a vault in Amazon S3 Glacier. Create a Vault Lock policy in S3 Glacier that meets all the regulatory requirements. Upload the data to the vault.
- D. Create an Amazon S3 bucket. Upload the data to the bucket. Use a lifecycle rule to transition the data to a vault in S3 Glacier. Create a Vault Lock policy that meets all the regulatory requirements.

Suggested Answer: C

Community vote distribution



🗳️ 👤 **IPLogic** 6 months, 4 weeks ago

Selected Answer: C

To meet the requirement of preserving several petabytes of data for 7 years in a cost-effective manner while ensuring that the data cannot be changed or deleted, option C is the most suitable solution. Here's why:

Creating a vault in Amazon S3 Glacier and implementing a Vault Lock policy ensures that the data is stored using a write-once-read-many (WORM) model, which prevents any changes or deletions during the specified retention period. S3 Glacier is designed for long-term storage of infrequently accessed data, making it a cost-effective choice for this use case. The Vault Lock policy will enforce compliance with regulatory requirements by ensuring that the data remains immutable for the entire retention period.

This approach leverages the cost benefits of S3 Glacier for long-term storage while providing the necessary security and compliance features to meet the company's requirements.

upvoted 2 times

🗳️ 👤 **BBR01** 11 months, 1 week ago

Selected Answer: A

For the folks voting for C with the link in the comment, it clearly says "This page is only for existing customers of the S3 Glacier service using Vaults and the original REST API from 2012."

upvoted 1 times

🗳️ 👤 **Just_Ninja** 9 months, 2 weeks ago

Its C:

The claim that S3 Glacier Vaults are only for existing customers using the original 2012 API is incorrect. While newer S3 Glacier Storage Classes have been introduced, S3 Glacier Vaults with Vault Lock are still fully supported and available for both new and existing customers.

Vault Lock remains a cost-effective option for securing data with WORM (Write Once, Read Many) compliance, making it ideal for long-term data storage with regulatory requirements. AWS continues to support this solution for all customers, including those setting up new Glacier Vaults.

In short, Vault Lock is not restricted to legacy customers and is available for current use.

upvoted 2 times

🗳️ 👤 **Sodev** 1 year, 2 months ago

C. Right here !

<https://docs.aws.amazon.com/amazonglacier/latest/dev/getting-started-upload-archive.html>

upvoted 2 times

🗳️ 👤 **Boul** 1 year, 3 months ago

C.

Petabytes of data takes a lot of storage, therefore require the cheapest storage: Glacier, if looking for a cost efficiency. The vault policy will keep it safe.

D is not realistic. Uploading petabytes of data to a bucket would require a AWS Snowmobile
upvoted 2 times

🗨️ **walter_white_008** 1 year, 3 months ago

Selected Answer: C

<https://docs.aws.amazon.com/amazonglacier/latest/dev/vault-lock-policy.html>
upvoted 1 times

🗨️ **Raphaello** 1 year, 4 months ago

Selected Answer: C

The data is already there, we just want to keep it for COMPLIANCE for SEVEN years.
There's no need to place the data in S3 bucket then use lifecycle to move it to Glacier.
Option C is correct.
upvoted 1 times

🗨️ **trashbox** 1 year, 6 months ago

Exam on 2023-12-18
upvoted 1 times

🗨️ **kejam** 1 year, 7 months ago

Selected Answer: C

<https://docs.aws.amazon.com/amazonglacier/latest/dev/vault-lock-policy.html>
upvoted 1 times

🗨️ **cjkuga** 1 year, 8 months ago

Selected Answer: C

Both A and C work here but C is the MOST cost-effective.
upvoted 4 times

🗨️ **pupsik** 1 year, 8 months ago

Selected Answer: A

Question doesn't ask for a backup solution, so Glacier is not a good fir here.
upvoted 2 times

🗨️ **Aamee** 1 year, 7 months ago

No, it clearly states that "The company must preserve this data for 7 years"... so how would you keep such large data safe and specifically complied with all the regulatory reqs. That's why going with C here.
upvoted 1 times

🗨️ **100fold** 1 year, 8 months ago

Selected Answer: C

Correction, answer C
upvoted 2 times

🗨️ **AgboolaKun** 1 year, 8 months ago

Selected Answer: C

The correct answer here is C. This option ticks all the boxes.

Several petabytes of data + 7 years + Regulatory Compliance + MOST cost-effective solution.

D is close but we don't S3 at all.

upvoted 2 times

🗨️ **100fold** 1 year, 8 months ago



Thanks AgboolKun! What are your thoughts on #49?
Agree with answer C as well. Can set the policy on Vault Lock that cannot be altered.
<https://docs.aws.amazon.com/amazonglacier/latest/dev/vault-access-policy.html>
upvoted 2 times

🗨️ **AgboolaKun** 1 year, 7 months ago

@100fold, I agree with your answer (C) in #49. There is no better option to C!!



I upvoted your answer already!!

upvoted 1 times

  **100fold** 1 year, 7 months ago



@AgboolaKun, I sat the exam Friday and marked 926. 80% from this study were on my exam. 6-7 new questions, one related to AWS KMS keyrings. Good luck everyone!

upvoted 5 times

  **kejam** 1 year, 7 months ago



Were most questions from examtopics aws-certified-security-specialty-scs-c02 ? or also the previous one for aws-certified-security-specialty-scs-c01 ?

upvoted 1 times

  **100fold** 1 year, 7 months ago

@kejam I studied both SCS-C02 & SCS-01, but 80%+- were from this C02 study. You'll notice C01 questions are merged. The other % were totally new questions. My access expires today so I can't comment further, but good luck you'll be good with this C02 study ☺

upvoted 3 times

  **100fold** 1 year, 8 months ago



Selected Answer: A

Answer A.

Compliance mode will prevent anyone from changing or deleting the data including the root user. Requested by the company's compliance team.

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock-overview.html>

upvoted 2 times

  **100fold** 1 year, 8 months ago

Correction to Answer C.

<https://docs.aws.amazon.com/amazonglacier/latest/dev/vault-access-policy.html>

upvoted 1 times

A company uses a third-party identity provider and SAML-based SSO for its AWS accounts. After the third-party identity provider renewed an expired signing certificate, users saw the following message when trying to log in:

Error: Response Signature Invalid (Service: AWSSecurityTokenService; Status Code: 400; Error Code: InvalidIdentityToken)

A security engineer needs to provide a solution that corrects the error and minimizes operational overhead.



Which solution meets these requirements?

- A. Upload the third-party signing certificate's new private key to the AWS identity provider entity defined in AWS Identity and Access Management (IAM) by using the AWS Management Console.
- B. Sign the identity provider's metadata file with the new public key. Upload the signature to the AWS identity provider entity defined in AWS Identity and Access Management (IAM) by using the AWS CLI.
- C. Download the updated SAML metadata file from the identity service provider. Update the file in the AWS identity provider entity defined in AWS Identity and Access Management (IAM) by using the AWS CLI.
- D. Configure the AWS identity provider entity defined in AWS Identity and Access Management (IAM) to synchronously fetch the new public key by using the AWS Management Console.

Suggested Answer: C

Community vote distribution

C (100%)

  **100fold** Highly Voted 1 year, 2 months ago

Selected Answer: C

Answer C.

https://docs.aws.amazon.com/IAM/latest/UserGuide/troubleshoot_saml.html#troubleshoot_saml_invalid-metadata

<https://www.examttopics.com/discussions/amazon/view/69166-exam-aws-certified-security-specialty-topic-1-question-292/>

upvoted 11 times

  **Raphaello** Most Recent 10 months, 3 weeks ago

Selected Answer: C

Download the updated SAML metadata file from your identity service provider, then update it in AWS.

upvoted 3 times

  **yorkicurke** 1 year ago

Selected Answer: C

nice link by user: 100fold

thanks bud

upvoted 1 times

  **awssecuritynewbie** 1 year ago

Selected Answer: C

C is the correct answer

upvoted 1 times

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. Define a SAML-based Amazon Cognito user pool and connect it to ADFS.
- B. Implement AWS IAM Identity Center (AWS Single Sign-On) in the management account and link it to ADFS as an identity provider. Define the EC2 instance as a managed resource, then apply an IAM policy on the resource.
- C. Define an Amazon Cognito identity pool, then install the connector on the Active Directory server. Use the Amazon Cognito SDK on the application instance to authenticate the employees using their Active Directory user names and passwords.
- D. 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.

Suggested Answer: A

Community vote distribution

A (100%)

100fold **Highly Voted** 1 year, 2 months ago

Selected Answer: A

Answer A.

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

<https://www.examtotics.com/discussions/amazon/view/30063-exam-aws-certified-security-specialty-topic-1-question-143/>

upvoted 10 times

Raphaello **Most Recent** 10 months, 3 weeks ago

Selected Answer: A

Add SAML IdP to Cognito user pool.

<https://docs.aws.amazon.com/cognito/latest/developerguide/cognito-user-pools-configuring-federation-with-saml-2-0-idp.html>

upvoted 2 times

A company is using AWS to run a long-running analysis process on data that is stored in Amazon S3 buckets. The process runs on a fleet of Amazon EC2 instances that are in an Auto Scaling group. The EC2 instances are deployed in a private subnet of a VPC that does not have internet access. The EC2 instances and the S3 buckets are in the same AWS account.

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

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

Which solution will meet these requirements?

- A. Update the policy on the S3 gateway endpoint to allow the S3 actions only if the values of the `aws:ResourceOrgID` and `aws:PrincipalOrgID` condition keys match the company's values.
- B. Update the policy on the instance profile role to allow the S3 actions only if the value of the `aws:ResourceOrgID` condition key matches the company's value.
- C. Add a network ACL rule to the subnet of the EC2 instances to block outgoing connections on port 443.
- D. Apply an SCP on the AWS account to allow the S3 actions only if the values of the `aws:ResourceOrgID` and `aws:PrincipalOrgID` condition keys match the company's values.

Suggested Answer: D

Community vote distribution



kejam Highly Voted 1 year, 7 months ago

Selected Answer: D

Answer D based on the syntax of these answers.

A. This could work, but you don't need `aws:ResourceOrgID` and `aws:PrincipalOrgID`. You can add allowed buckets (internal or external) as needed which is much more flexible IMO. <https://docs.aws.amazon.com/vpc/latest/privatelink/vpc-endpoints-s3.html#edit-vpc-endpoint-policy-s3>

B. This doesn't prevent S3 actions on external accounts.

C. This does nothing as the S3 endpoint is inside the VPC.

D. This solution matches the answer exactly.

Example 3: <https://aws.amazon.com/blogs/security/how-to-control-access-to-aws-resources-based-on-aws-account-or-organization/>
upvoted 14 times

NoCrapEva 1 year, 4 months ago

Also the question states the company has AWS Organisations - therefore any policy restrictions SHOULD be done at the Organisation level - In this case with a SCP

upvoted 2 times

AgboolaKun 1 year, 7 months ago

I agree totally. I have always thought that D is the correct answer but I could not locate any supported documentation online. Thank you for providing the link. The example 3 in the link as you pointed out tallies with the scenario in this question.

upvoted 1 times

Raphaello 1 year, 4 months ago

In fact Example 3 Restrict access to AWS resources (in this case S3) within my organization, which means denying access from principals (e.g. EC2 instance roles) that do not belong to S3 Org. That example does not correspond to what we need to do here!

"Deny",

"Action": "s3:*",

"Resource": "arn:aws:s3:::/*/*",

"Condition": { "StringNotEquals": { "aws:ResourceOrgID": "\${aws:PrincipalOrgID}" }

Note the "PrincipalOrgID" is a variable.

Whereas, we basically want our own EC2 instances not to access S3 that belong to another account.

"Allow",

"Action": "s3:*",


"Resource": "arn:aws:s3:::/*",

"Condition": {"StringEquals": {"aws:PrincipalOrgID": ["o-yyyyyyyyyy"]}}

Or

maybe even add an explicit deny statement if the "aws:ResourceOrgID" does not equal my Org ID "o-yyyyyyyyyy".

upvoted 1 times

 **1c7c461** Highly Voted 1 year, 6 months ago

Selected Answer: B

The answer is B. You all missed the part that EC2 instance is compromised. The restriction has to be added to the instance profile of the ec2 instance to restrict which S3 buckets it can connect to. This question is about limiting access from EC2 to external S3 buckets.

upvoted 6 times

 **b5f86df** Most Recent 3 weeks, 3 days ago

Selected Answer: D

The company is using AWS organizations and the best method to apply a filter in the resources it's through an SCP

upvoted 1 times

 **2f5c7cd** 1 month ago

Selected Answer: A

[https://aws.amazon.com/blogs/security/establishing-a-data-perimeter-on-aws-allow-only-trusted-resources-from-my-](https://aws.amazon.com/blogs/security/establishing-a-data-perimeter-on-aws-allow-only-trusted-resources-from-my-organization/#:~:text=The%20condition%20key%20used%20to,to%20implement%20the%20resource%20perimeter.&text=In%20the%20next%20section%2C%2)

[organization/#:~:text=The%20condition%20key%20used%20to,to%20implement%20the%20resource%20perimeter.&text=In%20the%20next%20section%2C%2](https://aws.amazon.com/blogs/security/establishing-a-data-perimeter-on-aws-allow-only-trusted-resources-from-my-organization/#:~:text=The%20condition%20key%20used%20to,to%20implement%20the%20resource%20perimeter.&text=In%20the%20next%20section%2C%2)

upvoted 1 times

 **AWSLoverLoverLoverLoverLover** 4 months, 1 week ago

Selected Answer: A

The correct answer is A.

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

upvoted 1 times

 **AWSLoverLoverLoverLoverLover** 4 months, 1 week ago

Selected Answer: A

Why is updating the S3 Gateway Endpoint policy the best solution?

S3 Gateway Endpoints control access at the VPC level, ensuring that all S3 traffic passes through a controlled path.

By modifying the endpoint policy to allow access only to S3 buckets within the company's AWS Organization, we prevent data exfiltration to external S3 buckets.

D. Apply an SCP to the AWS account

SCPs apply only to IAM identities (users, roles, accounts) and do not affect VPC Gateway Endpoints. This means compromised instances can still exfiltrate data through the endpoint.

upvoted 1 times


 **IPLogic** 6 months, 4 weeks ago

Selected Answer: A

To stop the exfiltration of data from compromised EC2 instances to external S3 buckets while keeping the EC2 processing job functional, the best solution is option A: Update the policy on the S3 gateway endpoint to allow the S3 actions only if the values of the aws:ResourceOrgID and aws:PrincipalOrgID condition keys match the company's values.

This approach ensures that only resources and principals within the company's AWS Organization can perform S3 actions, effectively blocking any attempts to exfiltrate data to S3 buckets outside the organization. By updating the S3 gateway endpoint policy, you can enforce this restriction at the network level, providing a robust and centralized control mechanism

upvoted 3 times

 **mzeynali** 7 months, 3 weeks ago

Selected Answer: A

Using Service Control Policies (SCPs) can be a part of a broader security strategy, but there are specific reasons why it may not be the most effective or immediate solution for the scenario where EC2 instances are compromised and exfiltrating data to an S3 bucket outside the company's organization:

1. Scope of SCPs: SCPs are designed to manage permissions across AWS Organizations. They apply to AWS accounts rather than individual resources. If the compromised instances are operating under an IAM role that has certain permissions, an SCP may not have the granularity needed to effectively restrict access at the resource level.

upvoted 3 times

🗨️ **jakie22332** 7 months, 3 weeks ago

Selected Answer: B

<https://aws.amazon.com/blogs/security/how-to-control-access-to-aws-resources-based-on-aws-account-or-organization/>

upvoted 1 times

🗨️ **icecool36** 1 year, 1 month ago

Selected Answer: D

SCP is the right choice.

Not B: This is only effective if the nodes are replaced. The processing must not be interrupted.

Not A: This does not do anything against the exfiltration

Not C: This will not work

upvoted 1 times

🗨️ **9bb8cb3** 1 year, 2 months ago

Selected Answer: A

Blocking at the network layer would allow you still have other workloads that can talk to other buckets outside of the account whereas the option D as others have suggested would mean no workload in the account would ever be able to talk to a bucket outside the org which is IMO too restrictive and the solution is not asking for a general solution just to this specific problem. You can also create additional VPC endpoints and bound them to other route tables which don't have this policy as to allow for other architectural possibilities mentioned above,

upvoted 4 times

🗨️ **ion_gee** 1 year, 2 months ago

The correct answer should be B, as it directly addresses the issue.

Option D seems too broad, and might affect other roles in the Account.

upvoted 2 times

🗨️ **Noexperience** 1 year, 4 months ago

Selected Answer: B

"Version": "2012-10-17",

"Statement": [

{

"Effect": "Allow",

"Action": [

"s3:GetObject",

"s3:PutObject"

],

"Resource": [

"arn:aws:s3:::bucket-name/*", // Specific buckets to restrict

"arn:aws:s3:::another-bucket/*"

],

"Condition": {

"StringEquals": {

"aws:ResourceOrgID": "YOUR_AWS_ORGANIZATION_ID"

}

}

}

}

}

upvoted 2 times

🗨️ **bkbaws** 1 year, 4 months ago

Selected Answer: A

the EC2 role S3 GET/PUT are restricted to the appropriate buckets, so the exfiltration bucket access is being granted by the default S3 gateway resource policy.

Hence restricting the EC2 attached IAM role to the given organization (B) will do nothing and B is incorrect.

C would break everything.

For (D) - SCPs don't apply to resource policies, so exfiltration would continue through the S3 gateway.

Answer is A

upvoted 2 times

🗨️ 👤 **Raphaello** 1 year, 4 months ago

Selected Answer: B

The problem is that EC2 instance exfiltrating data to an S3 bucket that is outside the company's organization in AWS Organizations.

So we need to make sure those instance cannot put the data to an external account's bucket.

Therefore, we need to restrict access ONLY to resources within an organization using condition "aws:ResourceOrgID".

Remember, it is not about controlling access to our own S3 bucket. It is about stopping EC2 instances from exfiltrate our data to accounts outside our Org.

Option B is the correct answer.

upvoted 1 times

🗨️ 👤 **LazyAutonomy** 1 year, 4 months ago

Selected Answer: A

Answer is A.

D is wrong because attackers wont use EC2 instance credentials to exfil data - no attacker is that stupid.

upvoted 2 times

🗨️ 👤 **LazyAutonomy** 1 year, 4 months ago

<https://docs.aws.amazon.com/vpc/latest/privatelink/vpc-endpoints-s3.html#edit-vpc-endpoint-policy-s3>

<https://developer.squareup.com/blog/adopting-aws-vpc-endpoints-at-square/>

upvoted 1 times

🗨️ 👤 **mark16dc** 1 year, 4 months ago

Given the effectiveness and direct impact on preventing data exfiltration to external S3 buckets, Option D is the correct solution. It leverages the organizational control provided by AWS Organizations to enforce policy restrictions at the account level, ensuring that S3 actions are confined to the company's organizational boundaries, thus meeting the security requirements without disrupting the EC2 processing jobs.

upvoted 1 times

A company that operates in a hybrid cloud environment must meet strict compliance requirements. The company wants to create a report that includes evidence from on-premises workloads alongside evidence from AWS resources. A security engineer must implement a solution to collect, review, and manage the evidence to demonstrate compliance with company policy.

Which solution will meet these requirements?

- A. Create an assessment in AWS Audit Manager from a prebuilt framework or a custom framework. Upload manual evidence from the on-premises workloads. Add the evidence to the assessment. Generate an assessment report after Audit Manager collects the necessary evidence from the AWS resources.
- B. Install the Amazon CloudWatch agent on the on-premises workloads. Use AWS Config to deploy a conformance pack from a sample conformance pack template or a custom YAML template. Generate an assessment report after AWS Config identifies noncompliant workloads and resources.
- C. Set up the appropriate security standard in AWS Security Hub. Upload manual evidence from the on-premises workloads. Wait for Security Hub to collect the evidence from the AWS resources. Download the list of controls as a .csv file.
- D. Install the Amazon CloudWatch agent on the on-premises workloads. Create a CloudWatch dashboard to monitor the on-premises workloads and the AWS resources. Run a query on the workloads and resources. Download the results.

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ 👤 **Sodev** 8 months, 3 weeks ago

A

Right here : <https://docs.aws.amazon.com/audit-manager/latest/userguide/upload-evidence.html>

upvoted 3 times

🗳️ 👤 **as9347** 10 months ago

Selected Answer: A

voted A

upvoted 1 times

🗳️ 👤 **Raphaello** 10 months, 3 weeks ago

Selected Answer: A

AWS Audit Manager helps you to collect evidences as per a selected framework.

Option A is correct.

upvoted 1 times

🗳️ 👤 **WeepingMaple** 1 year ago

Selected Answer: A

The keyword in the question is "evidence". AWS Audit Manager manages evidences.

<https://docs.aws.amazon.com/audit-manager/latest/userguide/what-is.html>

upvoted 3 times

🗳️ 👤 **confusedyeti69** 1 year ago

Selected Answer: A

voted A

upvoted 1 times

🗳️ 👤 **brpjp** 1 year, 1 month ago

Requirement is company policy compliance, which security hub supports and not Audit Manager. C is correct answer.

upvoted 2 times

🗳️ 👤 **Aamee** 1 year ago

Source for this statement pls?... See the below Audit Manager link which justifies option A instead:

<https://aws.amazon.com/blogs/aws/aws-audit-manager-simplifies-audit-preparation/>



upvoted 3 times

🗳️ 👤 **100fold** 1 year, 2 months ago

Selected Answer: A

Answer A

upvoted 4 times

  **kk2000** 1 year, 2 months ago

A

<https://aws.amazon.com/blogs/aws/aws-audit-manager-simplifies-audit-preparation/>

upvoted 4 times

To meet regulatory requirements, a security engineer needs to implement an IAM policy that restricts the use of AWS services to the us-east-1 Region.

What policy should the engineer implement?

A.

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

B.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "*",
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "ec2:Region": "us-east-1"
        }
      }
    }
  ]
}
```

C.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": "*",
      "Resource": "*",
      "Condition": {
        "StringNotEquals": {
          "aws:RequestedRegion": "us-east-1"
        }
      }
    }
  ]
}
```

D.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Deny",
      "NotAction": "*",
      "Resource": "*",
      "Condition": {
        "StringEquals": {
          "aws:RequestedRegion": "us-east-1"
        }
      }
    }
  ]
}
```

Suggested Answer: C

Community vote distribution

C (100%)

 **100fold** Highly Voted 1 year, 8 months ago

Selected Answer: C

Answer C

https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_examples_aws_deny-requested-region.html

upvoted 8 times

 **zzyy** Most Recent 1 year, 2 months ago

Not sure even if C is correct... On the condition it says StringNotEquals to us-east-1. which means the Deny all the aws resources except us-east-1 which is not we want right..? I am confused..

upvoted 1 times

🗳️ 👤 **zzyy** 1 year, 2 months ago

Nevermind.. I got the actual picture now.. ChatGPT helped.

When you use "StringEquals": { "aws:RequestedRegion": "us-east-1" }, it means that the condition will only be satisfied if the requested region is exactly "us-east-1".

So, with this condition in place, any AWS service request made from a region other than us-east-1 will be denied. It's a strict policy that allows access only if the requested region matches "us-east-1". If the request comes from any other region, it will be denied, ensuring that all operations occur exclusively within the specified region.

This policy denies all actions ("Action": "*") on all resources ("Resource": "*"), but only if the requested region is not "us-east-1". This effectively restricts the use of AWS services to the us-east-1 Region.

upvoted 1 times

🗳️ 👤 **frankzeng** 1 year, 3 months ago

C is wrong. C is deny, no allow.

upvoted 1 times

🗳️ 👤 **helloworldabc** 9 months, 2 weeks ago

just C

upvoted 1 times

🗳️ 👤 **Raphaello** 1 year, 4 months ago

Selected Answer: C

The request is to restrict (deny) use of services outside a specific region, therefore an "allow" policy for that specific region is not enough.

Option C does just that, it denies all services if the "requested region" is not the specific one.

upvoted 3 times

🗳️ 👤 **jakie22332** 7 months, 3 weeks ago

this - exactly

upvoted 1 times

🗳️ 👤 **rahav** 1 year, 6 months ago

Selected Answer: C

Answer is C

upvoted 1 times

🗳️ 👤 **Imimi** 1 year, 7 months ago

Why not A? C is just not denied, but not explicit allow.

upvoted 1 times

🗳️ 👤 **Aamee** 1 year, 7 months ago

A can't be correct since the 'Deny' always takes the precedence over 'Allow' if any similar SID policy statement is defined. The option C looks correct since it denies the access of the aws resources explicitly through the condition that 'IF' the region is not equal to 'us-east-1'. Since the question states that the access restriction should be limited to just us-east-1 region only.

upvoted 2 times

🗳️ 👤 **kk2000** 1 year, 8 months ago

C is Correct

upvoted 2 times

A company has a web server in the AWS Cloud. The company will store the content for the web server in an Amazon S3 bucket. A security engineer must use an Amazon CloudFront distribution to speed up delivery of the content. None of the files can be publicly accessible from the S3 bucket directly.

Which solution will meet these requirements?

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

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **100fold** Highly Voted 🍌 1 year, 2 months ago

Selected Answer: B

Answer B.

Amazon CloudFront introduces Origin Access Control (OAC)

<https://aws.amazon.com/blogs/networking-and-content-delivery/amazon-cloudfront-introduces-origin-access-control-oac/>

upvoted 10 times

🗳️ 👤 **Raphaello** Most Recent ⌚ 10 months, 1 week ago

Selected Answer: B

Correct answer is B.

Set OAC to restrict access to S3 bucket to only CF distribution.

upvoted 1 times

🗳️ 👤 **nouns** 1 year ago

Selected Answer: B

Option B: When you associate an OAI with a CloudFront distribution, it acts as a pseudo-user for the distribution, and you can configure S3 bucket permissions to grant access only to that OAI. This allows CloudFront to fetch and serve objects from the S3 bucket on behalf of the end-users without making the objects directly accessible from the S3 bucket.

upvoted 2 times

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 myFunction. 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.

Suggested Answer: D

Community vote distribution

D (79%)

C (21%)

 **100fold** Highly Voted 1 year, 8 months ago

Selected Answer: D

Answer D

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/iam-identity-based-access-control-cwl.html>

upvoted 6 times

 **mnsait** Most Recent 2 months, 4 weeks ago

Selected Answer: D

I tend to agree with those saying D.

Rationale: The engineer gets "error loading Log Streams" message, which is when there is NO stream to load in the first place. Hence D is the answer.

Had it been that the stream is available and the engineer could not retrieve the logs, then option C (adding GetLogEvents) could help.

upvoted 1 times

 **Malluchan** 2 months, 4 weeks ago


Selected Answer: D

To enable Lambda functions to properly log to Cloud-Watch, their execution role needs the permissions to create log groups,

Create log streams, (Missing)

and put log events, which are typically granted by attaching the AWSLambdaBasicExecutionRole managed policy.

upvoted 1 times

 **sunsetalder** 3 months, 3 weeks ago

Selected Answer: D

For Lambda functions to properly log to CloudWatch, they need three key permissions:

logs:CreateLogGroup (already present)

logs:CreateLogStream (missing)
logs:PutLogEvents (already present)
upvoted 1 times

🗳️ 👤 **icecool36** 1 year, 1 month ago

Not C: It is about the function role policy, not the engineer role policy
upvoted 2 times

🗳️ 👤 **ajwksldldfgdsg** 3 weeks, 6 days ago
great point brother
upvoted 1 times

🗳️ 👤 **icecool36** 1 year, 1 month ago

Selected Answer: D
Needs to be able to create a logstream
upvoted 4 times

🗳️ 👤 **minTwin** 1 year, 2 months ago

Selected Answer: C
While logs:CreateLogStream might be needed by the Lambda function itself to write logs initially, the prompt emphasizes viewing logs, which requires logs:GetLogEvents.
upvoted 1 times

🗳️ 👤 **helloworldabc** 9 months, 1 week ago
just D
upvoted 1 times

🗳️ 👤 **Sodev** 1 year, 2 months ago

D.
View Lambda logs on Lambda console, not same mean with: Lambda need to fetch logs from CW. That's only a AWS UI feature, a CW log frame in Lambda UI.
upvoted 1 times

🗳️ 👤 **Raphaello** 1 year, 4 months ago

Selected Answer: D
Action "logs:GetLogEvents" gets log events from log stream
(https://docs.aws.amazon.com/AmazonCloudWatchLogs/latest/APIReference/API_GetLogEvents.html), but there no log stream on first place!

Correct answer is allowing the function to create log stream (through logs:CreateLogStream).

D.
upvoted 3 times

🗳️ 👤 **mynickc** 1 year, 5 months ago

Selected Answer: D
putlogevent require logstream. so it is createlogstream. <https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/permissions-reference-cwl.html>
upvoted 1 times

🗳️ 👤 **Gafa255** 1 year, 5 months ago

Selected Answer: C
C is correct options because Security Engineer wants to see LOG. logs:GetLogEvents
<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/iam-identity-based-access-control-cwl.html>
upvoted 2 times

🗳️ 👤 **Gafa255** 1 year, 5 months ago
Sorry the correct options is D because the issue is when the Lambda function want to create the streams.
upvoted 1 times

🗳️ 👤 **marlonchin** 1 year, 7 months ago

D I don't think adding logs is not the responsibility of the Lamda function here
upvoted 1 times

🗳️ 👤 **[Removed]** 1 year, 7 months ago

Seems poorly written to me.. The engineer is trying to view the logs. So they'll need "logs:GetLogEvents " permission. However the policy is also missing "logs:CreateLogStream" so they will also need that in order for the Lambda to "create a new log stream". Doubt something this bad will be on the exam

upvoted 1 times

🗨️ 👤 **YR4591** 1 year, 7 months ago

Selected Answer: D

When creating log group, there should be a permission to put log streams in the log group

upvoted 2 times

🗨️ 👤 **kejam** 1 year, 7 months ago

Selected Answer: C

Answer C. The security engineer wants to view logs in CloudWatch.

- A. logs:CreateLogGroup - Required to create a new log group
- B. logs:PutDestination - Required to create or update a destination log stream
- C. Add the logs:GetLogEvents - Required to retrieve log events from a log stream
- D. Add the logs:CreateLogStream - Required to create a new log stream in a log group

<https://docs.aws.amazon.com/AmazonCloudWatch/latest/logs/permissions-reference-cwl.html>

upvoted 1 times

🗨️ 👤 **Aamee** 1 year, 6 months ago

But w/o creating the LogStream, how can the PutLogStream going to work from option C?.. that looks missing to me though..

upvoted 1 times

🗨️ 👤 **kk2000** 1 year, 8 months ago

Correct Answer is D

upvoted 3 times

A company has a new partnership with a vendor. The vendor will process data from the company's customers. The company will upload data files as objects into an Amazon S3 bucket. The vendor will download the objects to perform data processing. The objects will contain sensitive data. A security engineer must implement a solution that prevents objects from residing in the S3 bucket for longer than 72 hours. Which solution will meet these requirements?

- A. Use Amazon Macie to scan the S3 bucket for sensitive data every 72 hours. Configure Macie to delete the objects that contain sensitive data when they are discovered.
- B. Configure an S3 Lifecycle rule on the S3 bucket to expire objects that have been in the S3 bucket for 72 hours.
- C. Create an Amazon EventBridge scheduled rule that invokes an AWS Lambda function every day. Program the Lambda function to remove any objects that have been in the S3 bucket for 72 hours.
- D. Use the S3 Intelligent-Tiering storage class for all objects that are uploaded to the S3 bucket. Use S3 Intelligent-Tiering to expire objects that have been in the S3 bucket for 72 hours.

Suggested Answer: B

Community vote distribution

B (100%)

100fold **Highly Voted** 1 year, 2 months ago

Selected Answer: B

B

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lifecycle-mgmt.html>

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/how-to-set-lifecycle-configuration-intro.html>

upvoted 6 times

navid1365 **Most Recent** 7 months, 2 weeks ago

Selected Answer: B

For sure it's B.

upvoted 1 times

Raphaello 10 months, 1 week ago

Selected Answer: B

Correct answer is B.

The question tries to trick to select Macie, since data contains sensitive data, but Macie discover and classify data, and send findings to SecurityHub or EventBridge for any action that might be needed. It does NOT delete objects.

Another thing is that S3 object lifecycle is used for 2 roles, both transition and EXPIRATION of objects.

B is the correct answer.

upvoted 2 times

Aamee 1 year ago

Selected Answer: B

B is self-explanatory and sufficient.

upvoted 1 times

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 instance. Detach the root volume. Generate a new key pair.
- B. Keep the instance running. Detach the root volume. Generate a new key pair.
- C. When the volume is detached from the original instance, attach the volume 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. Start the instance.
- D. When the volume is detached from the original instance, attach the volume to another instance as a data volume. Modify the `authorized_keys` file with a new private key. Move the volume back to the original instance. Start the instance.
- E. When the volume is detached from the original instance, attach the volume 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 that is running.

Suggested Answer: AC

Community vote distribution

AC (100%)

100fold Highly Voted 1 year, 2 months ago

Selected Answer: AC

Answer AC

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/TroubleshootingInstancesConnecting.html#replacing-lost-key-pair>
upvoted 10 times

Raphaello Most Recent 10 months, 1 week ago

Selected Answer: AC

AC are the correct answers.

There are other ways to add/replace pub key into "authorized_keys" file without stopping the instance, but within the context of this scenario, AC are good.

Remember, "authorized_keys" file resides on the root volume. You cannot keep the instance running without the root volume.
upvoted 3 times

jeff001 1 year ago

Selected Answer: AC

A & C. 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

upvoted 2 times

352ae9a 1 year, 1 month ago

Answer AC


upvoted 2 times

A company purchased a subscription to a third-party cloud security scanning solution that integrates with AWS Security Hub. A security engineer needs to implement a solution that will remediate the findings from the third-party scanning solution automatically. Which solution will meet this requirement?

- A. Set up an Amazon EventBridge rule that reacts to new Security Hub findings. Configure an AWS Lambda function as the target for the rule to remediate the findings.
- B. Set up a custom action in Security Hub. Configure the custom action to call AWS Systems Manager Automation runbooks to remediate the findings.
- C. Set up a custom action in Security Hub. Configure an AWS Lambda function as the target for the custom action to remediate the findings.
- D. Set up AWS Config rules to use AWS Systems Manager Automation runbooks to remediate the findings.

Suggested Answer: A

Community vote distribution



100fold Highly Voted 1 year, 8 months ago

Selected Answer: A

Answer A

<https://docs.aws.amazon.com/prescriptive-guidance/latest/patterns/automate-remediation-for-aws-security-hub-standard-findings.html>

upvoted 8 times

Raphaello Highly Voted 1 year, 4 months ago

Selected Answer: A

Another tricking question.

EventBridge integrates with SecurityHub in 3 different ways..

1. All findings (SH Imported)
2. Findings for custom actions (SH Custom Actions)
3. Insights for custom actions (SH Insights)

(<https://docs.aws.amazon.com/securityhub/latest/userguide/securityhub-cwe-integration-types.html>)

You do not always need custom actions for EB integration, and to automatically remediate findings as in this scenario, 1st type of integration is required.

Answer is A.

upvoted 5 times

IPLogic Most Recent 6 months, 4 weeks ago

Selected Answer: A

Option B involves setting up a custom action in Security Hub and configuring the custom action to call AWS Systems Manager Automation runbooks to remediate the findings. While this approach can be effective, it requires manual intervention to trigger the custom action. This means that someone would need to manually select the custom action for each finding, which does not fully automate the remediation process.

On the other hand, option A leverages Amazon EventBridge to automatically react to new Security Hub findings and triggers an AWS Lambda function to remediate the findings. This approach ensures that the remediation process is fully automated and immediate, without requiring any manual intervention. It provides a more seamless and efficient solution for automatically addressing security findings.

upvoted 1 times

Davidng88 9 months, 3 weeks ago

Selected Answer: A

The requirement is automate response. EventBridge rules automatically actions when triggered, while Security Hub actions are manually triggered.

upvoted 1 times

cumzle_com 1 year ago

Selected Answer: B

While Option A can work, it lacks the direct integration with custom actions in Security Hub that Option B provides. Custom actions allow you to define specific remediation steps based on findings, which is more efficient and streamlined. Therefore, I recommend considering Option B for automated remediation.

<https://amer.resources.awscloud.com/security-assets/aws-security-hub-automated-response-and-remediation-implementation-guide>

upvoted 1 times

🗨️ 👤 **helloworldabc** 9 months, 1 week ago

just A

upvoted 1 times

🗨️ 👤 **SamHan** 1 year, 2 months ago

Selected Answer: B

B is correct

upvoted 1 times

🗨️ 👤 **Derets** 1 year, 5 months ago

Selected Answer: B

Answer B

Custom action is a native feature for Security Hub when using a 3rd-party library. Then you need to use Systems Manager Automation runbooks.

Answer A (EventBridge+Lambda) can be used for standard findings.

upvoted 2 times

🗨️ 👤 **yorkicurke** 1 year, 6 months ago

Selected Answer: B

Very tricky one. A,B and C can all be implemented. and we havent asked for easy,quickly or something like that a solution. so reason for not picking others

A:- could also be used but would require additional steps to configure rules to route findings from this specific third-party source to the appropriate target. Custom actions provide a native option within Security Hub.

C:- identical to B. same reasoning that Custom actions provide a native option within Security Hub.

to be honest i could go for any out of these three. even though i chose B. Arghhhh

upvoted 4 times

🗨️ 👤 **Raphaello** 1 year, 4 months ago

I beg to disagree.

EventBridge integration type "SH Imported" with automatically send all findings to EB.

EB does not care how the findings ended up in SH. SH integration with the third party handle this, and as long as the third party tool actually integrates with SH, that means it can send findings to it.

Once done, findings in SH automatically sent over to EB

(<https://docs.aws.amazon.com/securityhub/latest/userguide/securityhub-cwe-integration-types.html#securityhub-cwe-integration-types-all-findings>)

There you can build the action that you want.

Answer is A.

upvoted 3 times

🗨️ 👤 **Aamee** 1 year, 6 months ago

Selected Answer: A

To remediate the findings automatically, option A describes about the best practices..

upvoted 1 times

An application is running on an Amazon EC2 instance that has an IAM role attached. The IAM role provides access to an AWS Key Management Service (AWS KMS) customer managed key and an Amazon S3 bucket. The key is used to access 2 TB of sensitive data that is stored in the S3 bucket.


A security engineer discovers a potential vulnerability on the EC2 instance that could result in the compromise of the sensitive data. Due to other critical operations, the security engineer cannot immediately shut down the EC2 instance for vulnerability patching.

What is the FASTEST way to prevent the sensitive data from being exposed?

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

Suggested Answer: C

Community vote distribution



AWSLoverLoverLoverLoverLover 4 months, 1 week ago

Selected Answer: C

The goal is to quickly prevent the compromised EC2 instance from accessing the sensitive data in S3. The fastest way to do this is:

Revoke the IAM role's active session permissions

This immediately removes temporary credentials from the EC2 instance.

Use AWS STS to revoke active session tokens:

Update the S3 bucket policy to deny access to the IAM role

Even if the instance has temporary credentials, a bucket policy denial rule will override all permissions.

Remove the IAM role from the EC2 instance profile

This prevents the EC2 instance from obtaining new credentials.

upvoted 1 times

sophire 6 months ago

Selected Answer: D

While Option C attempts to mitigate access by removing permissions, it does not directly address the potential misuse of the KMS key during an active session. Option D ensures an immediate and definitive block to sensitive data access, making it the more secure and fastest response to prevent exposure.

If the volume of objects is significant, re-encryption can be deferred without compromising data security because disabling the original key renders the data inaccessible.

upvoted 1 times

Sodev 8 months, 3 weeks ago

C, However, This is not a good question. what happen if "Critical operation" need to continue access to data on S3 ?

If i am a writer for this question. I will add more Answer "Deny outbound from EC2 and routing to S3 via s3 gateway endpoint only"

upvoted 1 times

xflare 9 months, 1 week ago

Selected Answer: C

D takes longer than it looks.

If you disable the KMS key, the data key is still usable until you try to encrypt it again. If the data key has been unencrypted within the instance, it will remain usable even after you disable the KMS and the rogue instance can keep reading info from S3.

<https://docs.aws.amazon.com/kms/latest/developerguide/concepts.html#unusable-kms-keys>.

upvoted 2 times

🗳️ 👤 **Raphaello** 10 months, 1 week ago

Selected Answer: C

D is a valid solution, but not the fastest as requested. Creating a batch operation to re-encrypt 2TB to data on S3 might take time.

Plus, old or new KMS key are both equally same for an attacker who has access to the EC2/role that's allowed to use the key.

The solution needs to be with the role itself to eliminate further access to sensitive data.

Revoke current active session permissions, set S3 bucket policy to deny the role, and remove the role altogether from EC2 instance profile.

C.

upvoted 1 times

🗳️ 👤 **Oralinux** 1 year ago

C. Revoke the IAM role's active session permissions.

upvoted 3 times

🗳️ 👤 **Daniel76** 1 year ago

Selected Answer: C

This contains more detail response. Refer to part 2 for containment step. The first step is always to deal with the role access first.

<https://www.bicarait.com/2021/04/27/aws-incident-response-unintended-access-to-s3-bucket/>

It only takes a few minutes for policy updates to effectively revoke the role's temporary security credentials to force all users assuming the role to reauthenticate and request new credentials. (as compare to re-encrypt entire s3 bucket data to a single new key)

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_temp_control-access_disable-perms.html#deny-access-to-all-sessions

Furthermore, though unrelated to the requirement:

The s3 bucket may be encrypted by multiple data keys which is intended. by re-encrypting the entire bucket, you may affect data that are encrypted by other legitimate keys unaffected by this vulnerable ec2.

upvoted 2 times

🗳️ 👤 **Aamee** 1 year, 1 month ago

Still v confusing btw C and D.... but I'd probably go with C.

upvoted 1 times

🗳️ 👤 **kejam** 1 year, 1 month ago

Selected Answer: D

Answer D. The fastest way to prevent sensitive data from being exposed is to disable the current key.

A. Not fast

B. Not fast

C. `AWSRevokeOlderSessions` is fast, however bad actors can immediately reconnect with new sessions before you remove the IAM role from the EC2 instance profile. If these steps were reversed to prevent that its no longer the fastest solution because its 2 steps.

D. Disable the current key... 1st step prevents sensitive data exposure and the rest of the steps to re-encrypt the data with a new key can follow.

upvoted 3 times

🗳️ 👤 **confusedyeti69** 1 year ago

If your bucket has millions of objects, re-encryption is slower. Ans is C

upvoted 1 times

🗳️ 👤 **Aamee** 1 year ago

No, that's not the point here. The req. is to implement it 'FASTER' to get it secured on the first attempt which I also feel Option D provides it.

Disabling the key right away can atleast help ensure that no sensitive data would get exposed further IMO... and then the rest of the steps to re-encrypting the data can be done as a 2nd step to follow...

upvoted 1 times

🗳️ 👤 **100fold** 1 year, 2 months ago

Selected Answer: C

Answer C

<https://www.examttopics.com/discussions/amazon/view/60659-exam-aws-certified-security-specialty-topic-1-question-287/>

upvoted 3 times

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

What should the security engineer recommend?

- A. Enable Amazon RDS encryption to encrypt the database and snapshots. Enable Amazon Elastic Block Store (Amazon EBS) encryption on Amazon EC2 instances. Include the database credential in the EC2 user data field. Use an AWS Lambda function to rotate database credentials. Set up TLS for the connection to the database.
- B. Install a database on an Amazon EC2 instance. Enable third-party disk encryption to encrypt the Amazon Elastic Block Store (Amazon EBS) volume. Store the database credentials in AWS CloudHSM with automatic rotation. Set up TLS for the connection to the database.
- C. Enable Amazon RDS encryption to encrypt the database and snapshots. Enable Amazon Elastic Block Store (Amazon EBS) encryption on Amazon EC2 instances. Store the database credentials in AWS Secrets Manager with automatic rotation. Set up TLS for the connection to the RDS hosted database.
- D. Set up an AWS CloudHSM cluster with AWS Key Management Service (AWS KMS) to store KMS keys. Set up Amazon RDS encryption using AWS KMS to encrypt the database. Store database credentials in the AWS Systems Manager Parameter Store with automatic rotation. Set up TLS for the connection to the RDS hosted database.

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **Raphaello** 10 months, 1 week ago

Selected Answer: C

Correct answer is C.

It covers the required areas to protect sensitive data, without least overhead.

upvoted 1 times

🗳️ 👤 **kejam** 1 year, 1 month ago

Selected Answer: C

Answer C

https://docs.aws.amazon.com/secretsmanager/latest/userguide/rotate-secrets_turn-on-for-db.html

upvoted 4 times

🗳️ 👤 **100fold** 1 year, 2 months ago

Selected Answer: C

Answer C

<https://www.examtactics.com/discussions/amazon/view/60739-exam-aws-certified-security-specialty-topic-1-question-268/>

upvoted 4 times

A company is using Amazon Route 53 Resolver for its hybrid DNS infrastructure. The company has set up Route 53 Resolver forwarding rules for authoritative domains that are hosted on on-premises DNS servers.

A new security mandate requires the company to implement a solution to log and query DNS traffic that goes to the on-premises DNS servers. The logs must show details of the source IP address of the instance from which the query originated. The logs also must show the DNS name that was requested in Route 53 Resolver.


Which solution will meet these requirements?

- A. Use VPC Traffic Mirroring. Configure all relevant elastic network interfaces as the traffic source, include amazon-dns in the mirror filter, and set Amazon CloudWatch Logs as the mirror target. Use CloudWatch Insights on the mirror session logs to run queries on the source IP address and DNS name.
- B. Configure VPC flow logs on all relevant VPCs. Send the logs to an Amazon S3 bucket. Use Amazon Athena to run SQL queries on the source IP address and DNS name.
- C. Configure Route 53 Resolver query logging on all relevant VPCs. Send the logs to Amazon CloudWatch Logs. Use CloudWatch Insights to run queries on the source IP address and DNS name.
- D. Modify the Route 53 Resolver rules on the authoritative domains that forward to the on-premises DNS servers. Send the logs to an Amazon S3 bucket. Use Amazon Athena to run SQL queries on the source IP address and DNS name.

Suggested Answer: C

Community vote distribution

C (100%)

 **Malluchan** 2 months, 4 weeks ago

Selected Answer: C

C is the answer, Resolver query logs include values such as the following:

The AWS Region where the VPC was created

The ID of the VPC that the query originated from

The IP address of the instance that the query originated from

The instance ID of the resource that the query originated from

The date and time that the query was first made

The DNS name requested (such as prod.example.com)

The DNS record type (such as A or AAAA)

The DNS response code, such as NoError or ServFail

The DNS response data, such as the IP address that is returned in response to the DNS query
upvoted 1 times

 **Daniel76** 1 year, 6 months ago

Selected Answer: C

C is the only answer as you need to config resolver query logging on all vpc, and cloudwatch log insight indeed allow you to query the source IP address.

<https://aws.amazon.com/blogs/aws/log-your-vpc-dns-queries-with-route-53-resolver-query-logs/>
upvoted 2 times

🗨️ 👤 **brpjp** 1 year, 6 months ago

Correct answer is D. <https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver.html>. Please read this statement : Resolver rules enable you to create one forwarding rule for each domain name and specify the name of the domain for which you want to forward DNS queries from your VPC to an on-premises DNS resolver and from your on-premises to your VPC. Rules are applied directly to your VPC and can be shared across multiple accounts. so correct answer, based on above statement is D and not C, as it does not specify the requirements to send outbound connections to on-premise.

upvoted 1 times

🗨️ 👤 **FlyingHawk** 4 months, 3 weeks ago

Route 53 Resolver rules do not generate logs by themselves.

There is no built-in feature to modify rules to send logs to S3.

The correct method is to enable Route 53 Resolver query logging, not modify forwarding rules.

upvoted 1 times

🗨️ 👤 **helloworldabc** 9 months, 1 week ago

just C

upvoted 1 times

🗨️ 👤 **ahrentom** 1 year, 6 months ago

Selected Answer: C

For me it's answer C

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver-query-logs.html>

upvoted 3 times

🗨️ 👤 **[Removed]** 1 year, 7 months ago

Selected Answer: C

Seems like it's C

<https://medium.com/@sisodiyapradeep/dns-query-logging-aggregation-control-tower-environment-well-architected-telemetry-workload-266dcdbf7195>

upvoted 2 times

🗨️ 👤 **oioi** 1 year, 7 months ago

Selected Answer: C

correct.

upvoted 1 times

🗨️ 👤 **marlonchin** 1 year, 7 months ago

<https://repost.aws/knowledge-center/route53-view-endpoint-traffic>

upvoted 2 times

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

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

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

- A. The principal's identity-based policy grants access to put objects into the S3 bucket with no conditions.
- B. The principal's identity-based policy overrides the condition because the identity-based policy contains an explicit allow.
- C. The S3 bucket's resource policy does not deny access to put objects.
- D. The S3 bucket's resource policy cannot allow actions to the principal.
- E. The bucket policy does not apply to principals in the same zone of trust.

Suggested Answer: AC

Community vote distribution

AC (100%)

🗳️ 👤 **rahav** 1 year ago

Selected Answer: AC

AC is the correct ones

upvoted 1 times

🗳️ 👤 **Daniel76** 1 year ago

Selected Answer: AC

https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_attribute-based-access-control.html

Look out for step 2- create ABAC policy example.

To ensure that resources is only granted when principal and resource tag match, there should be condition for the access, and disallow when tag not match.

upvoted 2 times

🗳️ 👤 **[Removed]** 1 year, 1 month ago

BDE Don't appear to make much sense. So A C it is.

The resource policy should be setup to DENY if a tag DOES NOT MATCH the desired tags. Not ALLOW the tag listed alone. Otherwise, a IAM policy without conditions may be enough to provide access. Not a fan of this question as it leaves a ton up in the air but hey..

upvoted 3 times

🗳️ 👤 **oioi** 1 year, 1 month ago

Selected Answer: AC

correct

upvoted 3 times

🗳️ 👤 **marlonchin** 1 year, 1 month ago

Sorry not D I need the explanation for E if it is right answer

upvoted 1 times

🗳️ 👤 **marlonchin** 1 year, 1 month ago

BE can some explain how it is possible for D to be answer

upvoted 1 times

A company is hosting multiple applications within a single VPC in its AWS account. The applications are running behind an Application Load Balancer that is associated with an AWS WAF web ACL. The company's security team has identified that multiple port scans are originating from a specific range of IP addresses on the internet.

A security engineer needs to deny access from the offending IP addresses.

Which solution will meet these requirements?

- A. Modify the AWS WAF web ACL with an IP set match rule statement to deny incoming requests from the IP address range.
- B. Add a rule to all security groups to deny the incoming requests from the IP address range.
- C. Modify the AWS WAF web ACL with a rate-based rule statement to deny the incoming requests from the IP address range.
- D. Configure the AWS WAF web ACL with regex match conditions. Specify a pattern set to deny the incoming requests based on the match condition.

Suggested Answer: A

Community vote distribution

A (86%)

14%

 **christianperezalmanza** 4 months, 3 weeks ago

Selected Answer: B

WAF operates at layer 7 of the OSI model. If we want to block port scans, then option B is the right answer although it's more work.
upvoted 1 times

 **IPLogic** 7 months ago

Selected Answer: A

The best solution to deny access from the offending IP addresses is Option A: Modify the AWS WAF web ACL with an IP set match rule statement to deny incoming requests from the IP address range.
upvoted 1 times

 **navid1365** 1 year, 1 month ago

Selected Answer: A

Different available categories for AWS WAF conditions:

- ip address
- geolocation
- string and RegEx matching
- SQLi
- XSS
- Size Constraints

In this example ip address range is correct

upvoted 1 times

 **Raphaello** 1 year, 4 months ago

Selected Answer: A

The offending IP's are specified, and the request is to "DENY" access from them. There's no reason to rate-limit the flow. Just block it.

A.

upvoted 1 times

 **MikeRach** 1 year, 4 months ago

AWS WAF--> Inspects the request against a set of IP addresses and address ranges.

<https://docs.aws.amazon.com/waf/latest/developerguide/waf-rule-statement-type-ipset-match.html>

upvoted 1 times

 **3633f8f** 1 year, 6 months ago

Selected Answer: A

As per below discussion.

upvoted 1 times

🗨️ 👤 **Daniel76** 1 year, 6 months ago

Selected Answer: C

<https://docs.aws.amazon.com/waf/latest/developerguide/listing-managed-ips.html>

upvoted 1 times

🗨️ 👤 **WeepingMaplte** 1 year, 6 months ago

You will use rate limit if it is normal HTTP/s traffic. For port scanning is a network reconnaissance technique used to identify which ports on a computer system are open and potentially vulnerable. You will want to block it 100%.

upvoted 2 times

🗨️ 👤 **3633f8f** 1 year, 6 months ago

The question is how to apply what the security engineer wants to do which is to block every single request coming from that IP set. Hence, I also think correct answer is A.

upvoted 2 times

🗨️ 👤 **Daniel76** 1 year, 6 months ago

agree, should be C.

upvoted 1 times

🗨️ 👤 **Daniel76** 1 year, 6 months ago

agree, should be A instead.

upvoted 2 times

🗨️ 👤 **[Removed]** 1 year, 7 months ago

Selected Answer: A

Agreed on A

upvoted 1 times

🗨️ 👤 **oioi** 1 year, 7 months ago

Selected Answer: A

correct

upvoted 2 times

A company has contracted with a third party to audit several AWS accounts. To enable the audit, cross-account IAM roles have been created in each account targeted for audit. The auditor is having trouble accessing some of the accounts.

Which of the following may be causing this problem? (Choose three.)

- A. The external ID used by the auditor is missing or incorrect.
- B. The auditor is using the incorrect password.
- C. The auditor has not been granted sts:AssumeRole for the role in the destination account.
- D. The Amazon EC2 role used by the auditor must be set to the destination account role.
- E. The secret key used by the auditor is missing or incorrect.
- F. The role ARN used by the auditor is missing or incorrect.

Suggested Answer: ACF

Community vote distribution

ACF (100%)

  **ahirri** 6 months, 3 weeks ago

Selected Answer: AC

I am not sure what is the 3rd option! if it was "F" then the auditor should be denied in ALL accounts!

"The auditor is having trouble accessing some of the accounts"

upvoted 1 times

  **TrilogueLand** 4 months, 4 weeks ago

I think the language used in the question "...IAM rols have been created in each account" indicates that there could have been a mistake in those accounts when configuring the role which resulted in a typo of the ARN. Hence why auditor may not be able to access some of the accounts. A human error when making the role policies multiple times.

upvoted 1 times

  **Raphaello** 1 year, 4 months ago

Selected Answer: ACF



ACF the correct answers to this case.

* ExternalID is incorrect or missing

* Auditor is not allowed to assume the role on first place

* The role ARN is incorrect

upvoted 1 times

  **rahav** 1 year, 6 months ago

Selected Answer: ACF

ACF for sure

upvoted 1 times


  **Daniel76** 1 year, 6 months ago

Selected Answer: ACF

B & E, if its due to external ID or secret key used by auditor, then access to all accounts shd be affected.

D, ec2 is irrelevant in account access.



upvoted 2 times

  **[Removed]** 1 year, 7 months ago

Selected Answer: ACF

ACF for sure

upvoted 3 times

  **oioi** 1 year, 7 months ago

Selected Answer: ACF

correct

upvoted 2 times

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

Which bucket policy statement meets these requirements?

- A.
- ```

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

```
- B.
- ```

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

```
- C.
- ```

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

```
- D.
- ```

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

```

Suggested Answer: B

Community vote distribution



B (100%)

B is NOT CORRECT, because it's mandatory to have an explicit allow!!

D seems not correct because of "*" in the resources, but this "bucket policy" is only applied to this specific bucket

So for me, the only correct answer is D because it allows requests from the vpce



upvoted 1 times

  **2f5c7cd** 1 month, 3 weeks ago

B is correct

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/example-bucket-policies-vpc-endpoint.html#example-bucket-policies-restrict-access-vpc>

upvoted 1 times

  **sz1234** 1 year, 4 months ago

B is correct

upvoted 1 times



  **awssecuritynewbie** 1 year, 4 months ago

Selected Answer: B

B

The option D does not even have the bucket in it lol



upvoted 1 times

  **rahav** 1 year, 6 months ago

Selected Answer: B

B is the correct one

upvoted 1 times

  **Aamee** 1 year, 7 months ago

Selected Answer: B

No doubt, it's B.



upvoted 1 times

  **[Removed]** 1 year, 7 months ago

Selected Answer: B

B is correcto

upvoted 2 times

  **oioi** 1 year, 7 months ago

Selected Answer: B

correct

upvoted 3 times

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

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

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

- A. Ensure that the EC2 instance profile that is attached to the EC2 instances has permissions to create log streams and write logs.
- B. Create a metric filter on the logs so that they can be viewed in the AWS Management Console.
- C. Check the CloudWatch agent configuration file on each EC2 instance to make sure that the CloudWatch agent is collecting the proper log files.
- D. Check the VPC endpoint policies of both VPC endpoints to ensure that the EC2 instances have permissions to use them.
- E. Create a NAT gateway in the subnet so that the EC2 instances can communicate with CloudWatch.
- F. Ensure that the security groups allow all the EC2 instances to communicate with each other to aggregate logs before sending.

Suggested Answer: ACD

Community vote distribution

ACD (100%)

🗳️ 👤 **rahav** 1 year ago

Selected Answer: ACD

ACD are the correct ones. for sure
upvoted 1 times

🗳️ 👤 **Aamee** 1 year ago

Selected Answer: ACD

ACD makes a good comb. logically.
upvoted 2 times

🗳️ 👤 **[Removed]** 1 year, 1 month ago

Selected Answer: ACD

ACD seem correct
upvoted 3 times

🗳️ 👤 **oioi** 1 year, 1 month ago

Selected Answer: ACD

correct
upvoted 1 times


A company uses AWS Signer with all of the company's AWS Lambda functions. A developer recently stopped working for the company. The company wants to ensure that all the code that the developer wrote can no longer be deployed to the Lambda functions.

Which solution will meet this requirement?

- A. Revoke all versions of the signing profile assigned to the developer.
- B. Examine the developer's IAM roles. Remove all permissions that grant access to Signer.
- C. Re-encrypt all source code with a new AWS Key Management Service (AWS KMS) key.
- D. Use Amazon CodeGuru to profile all the code that the Lambda functions use.

Suggested Answer: A

Community vote distribution



🗳️ 👤 **IPLogic** 6 months, 4 weeks ago

Selected Answer: A

Option B, which involves examining the developer's IAM roles and removing all permissions that grant access to AWS Signer, is not as effective as revoking the signing profile. While removing IAM permissions can prevent the developer from accessing AWS Signer in the future, it does not invalidate the signatures that have already been created using the developer's signing profile. This means that any code signed by the developer before their permissions were revoked could still be deployed to the Lambda functions.

Revoking all versions of the signing profile assigned to the developer ensures that any signatures generated by that profile are no longer valid, effectively preventing the deployment of the developer's code

upvoted 1 times

🗳️ 👤 **rahav** 1 year, 6 months ago

Selected Answer: A

<https://docs.aws.amazon.com/signer/latest/developerguide/revocation.html>

upvoted 1 times

🗳️ 👤 **vincentsr7** 1 year, 6 months ago

The privilege has to be removed from signer, for this option b is the right answer

upvoted 1 times

🗳️ 👤 **helloworldabc** 9 months, 1 week ago

just A

upvoted 1 times

🗳️ 👤 **vincentsr7** 1 year, 6 months ago

Option A suggests revoking all versions of the signing profile assigned to the developer, but this is not the most effective solution for preventing the developer from deploying code to Lambda functions. Signing profiles primarily deal with the integrity and authenticity of code, rather than controlling the ability to deploy code

upvoted 1 times

🗳️ 👤 **kejam** 1 year, 7 months ago

Selected Answer: A

Answer A

New URL: <https://docs.aws.amazon.com/signer/latest/developerguide/revocation.html>

upvoted 3 times

🗳️ 👤 **Imimi** 1 year, 7 months ago

A

Refer to <https://docs.aws.amazon.com/signer/latest/developerguide/revoking.html>

<https://docs.aws.amazon.com/signer/latest/developerguide/revoking.html>

upvoted 3 times

🗳️ 👤 **AgboolaKun** 1 year, 7 months ago



Selected Answer: A

A is the correct answer. Revoke the developer signing profile - <https://docs.aws.amazon.com/signer/latest/developerguide/revocation.html>
upvoted 2 times

  **[Removed]** 1 year, 7 months ago



Selected Answer: A

A will handle prevention
upvoted 2 times

  **oioi** 1 year, 7 months ago

Selected Answer: B

correct
upvoted 1 times

  **Aamee** 1 year, 7 months ago

Source pls?
upvoted 1 times

A company plans to use AWS Key Management Service (AWS KMS) to implement an encryption strategy to protect data at rest. The company requires client-side encryption for company projects. The company is currently conducting multiple projects to test the company's use of AWS KMS. These tests have led to a sudden increase in the company's AWS resource consumption. The test projects include applications that issue multiple requests each second to KMS endpoints for encryption activities.

The company needs to develop a solution that does not throttle the company's ability to use AWS KMS. The solution must improve key usage for client-side encryption and must be cost optimized.

Which solution will meet these requirements?

- A. Use keyrings with the AWS Encryption SDK. Use each keyring individually or combine keyrings into a multi-keyring. Decrypt the data by using a keyring that has the primary key in the multi-keyring.
- B. Use data key caching. Use the local cache that the AWS Encryption SDK provides with a caching cryptographic materials manager.
- C. Use KMS key rotation. Use a local cache in the AWS Encryption SDK with a caching cryptographic materials manager.
- D. Use keyrings with the AWS Encryption SDK. Use each keyring individually or combine keyrings into a multi-keyring. Use any of the wrapping keys in the multi-keyring to decrypt the data.

Suggested Answer: B

Community vote distribution

B (100%)

 **Daniel76** 1 year, 6 months ago

Selected Answer: B

"Caching can reduce your use of cryptographic services, such as AWS Key Management Service (AWS KMS). If you are hitting your AWS KMS requests-per-second limit, caching can help."

<https://docs.aws.amazon.com/encryption-sdk/latest/developer-guide/data-key-caching.html>

upvoted 2 times

 **WeepingMaple** 1 year, 6 months ago


Local key cache: Implement local key caching in your applications to store frequently used encryption keys, reducing the number of calls to KMS. Consider libraries like AWS Encryption SDK for secure key cache management.

upvoted 1 times

 **Aamee** 1 year, 7 months ago


Shouldn't it be 'D'? Couldn't verify the source for the option of 'C' being correct anywhere.

upvoted 1 times

 **helloworldabc** 9 months, 1 week ago

just B

upvoted 1 times

 **Imimi** 1 year, 7 months ago

I vote for B

upvoted 1 times

 **AgboolaKun** 1 year, 7 months ago

Selected Answer: B

Data key caching helps to improve performance, reduce cost, and help stay within limit as your key usage increases without throttling -

<https://docs.aws.amazon.com/encryption-sdk/latest/developer-guide/data-key-caching.html>

upvoted 4 times



 **[Removed]** 1 year, 7 months ago

Selected Answer: B

B seems right

<https://repost.aws/knowledge-center/kms-throttlingexception-error>



upvoted 1 times

  **oioi** 1 year, 7 months ago

Selected Answer: B

correct

upvoted 1 times

  **oioi** 1 year, 7 months ago

C is correct

upvoted 1 times

A security team is working on a solution that will use Amazon EventBridge to monitor new Amazon S3 objects. The solution will monitor for public access and for changes to any S3 bucket policy or setting that result in public access. The security team configures EventBridge to watch for specific API calls that are logged from AWS CloudTrail. EventBridge has an action to send an email notification through Amazon Simple Notification Service (Amazon SNS) to the security team immediately with details of the API call.

Specifically, the security team wants EventBridge to watch for the s3:PutObjectAcl, s3:DeleteBucketPolicy, and s3:PutBucketPolicy API invocation logs from CloudTrail. While developing the solution in a single account, the security team discovers that the s3:PutObjectAcl API call does not invoke an EventBridge event. However, the s3:DeleteBucketPolicy API call and the s3:PutBucketPolicy API call do invoke an event.

The security team has enabled CloudTrail for AWS management events with a basic configuration in the AWS Region in which EventBridge is being tested. Verification of the EventBridge event pattern indicates that the pattern is set up correctly. The security team must implement a solution so that the s3:PutObjectAcl API call will invoke an EventBridge event. The solution must not generate false notifications.


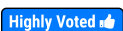
Which solution will meet these requirements?

- A. Modify the EventBridge event pattern by selecting Amazon S3. Select All Events as the event type.
- B. Modify the EventBridge event pattern by selecting Amazon S3. Select Bucket Level Operations as the event type.
- C. Enable CloudTrail Insights to identify unusual API activity.
- D. Enable CloudTrail to monitor data events for read and write operations to S3 buckets.

Suggested Answer: D

Community vote distribution

D (100%)

 **[Removed]**  1 year, 1 month ago

Selected Answer: D

You need to enable data events for that API event to trigger

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/cloudtrail-logging-s3-info.html#cloudtrail-object-level-tracking>
upvoted 8 times

 **WeepingMaple**  1 year ago

By default, CloudTrail only logs bucket-level API calls in S3, not object-level actions. This means it logs events like creating or deleting buckets, but not actions like uploading or downloading objects.

To enable object-level logging, you need to explicitly configure CloudTrail for your S3 buckets. You can do this in the S3 console, CLI, or SDK.
upvoted 6 times

 **AWSLoverLoverLoverLoverLover**  4 months, 1 week ago

Selected Answer: D

D. Enable CloudTrail to monitor data events for read and write operations to S3 buckets.

Explanation:

Why is s3:PutObjectAcl not triggering an EventBridge event?

CloudTrail logs two types of events:

Management events: Capture bucket-level operations (e.g., s3:PutBucketPolicy, s3:DeleteBucketPolicy).

Data events: Capture object-level operations (e.g., s3:PutObjectAcl, s3:GetObject).

The security team has only enabled management events, which is why s3:PutObjectAcl is not logged.



upvoted 2 times

 **rahav** 1 year ago

Selected Answer: D

you need to enable data events in cloudtrail

upvoted 1 times

  **oioi** 1 year, 1 month ago

Selected Answer: D

correct

upvoted 1 times

A company uses Amazon GuardDuty. The company's security team wants all High severity findings to automatically generate a ticket in a third-party ticketing system through email integration.

Which solution will meet this requirement?

- A. Create a verified identity for the third-party ticketing email system in Amazon Simple Email Service (Amazon SES). Create an Amazon EventBridge rule that includes an event pattern that matches High severity GuardDuty findings. Specify the SES identity as the target for the EventBridge rule.
- B. Create an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the third-party ticketing email system to the SNS topic. Create an Amazon EventBridge rule that includes an event pattern that matches High severity GuardDuty findings. Specify the SNS topic as the target for the EventBridge rule.
- C. Use the GuardDuty CreateFilter API operation to build a filter in GuardDuty to monitor for High severity findings. Export the results of the filter to an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the third-party ticketing email system to the SNS topic.
- D. Use the GuardDuty CreateFilter API operation to build a filter in GuardDuty to monitor for High severity findings. Create an Amazon Simple Notification Service (Amazon SNS) topic. Subscribe the third-party ticketing email system to the SNS topic. Create an Amazon EventBridge rule that includes an event pattern that matches GuardDuty findings that are selected by the filter. Specify the SNS topic as the target for the EventBridge rule.

Suggested Answer: B

Community vote distribution

B (90%)

10%

🗳️ 👤 **hro** 1 year, 3 months ago

D is essentially B & C together.

D is the most complete answer.

upvoted 1 times

🗳️ 👤 **Just_Ninja** 9 months, 1 week ago

Sorry hro, D is tu much work, if you can get a faster solution like B.

upvoted 1 times

🗳️ 👤 **helloworldabc** 9 months, 1 week ago

just B

upvoted 1 times

🗳️ 👤 **rahav** 1 year, 6 months ago

Selected Answer: B

B is the right one- you do it with EventBridge

upvoted 2 times

🗳️ 👤 **3633f8f** 1 year, 6 months ago

Selected Answer: B

Create a high severity Eventbridge based on GuardDuty High Severity Findings:

https://docs.aws.amazon.com/guardduty/latest/ug/guardduty_findings_cloudwatch.html

upvoted 2 times

🗳️ 👤 **WeepingMaplte** 1 year, 6 months ago

Key steps for integration:

1. Create a Eventbridge/CloudWatch event rule: Define the event pattern to trigger the rule based on specific GuardDuty findings.

2. Configure the rule target: Choose where to send the findings data, like an SNS topic

upvoted 1 times

🗳️ 👤 **Imimi** 1 year, 7 months ago

B.

Users can define filter in EventBridge. Not necessary to use GuardDuty CreateFilter API.

upvoted 2 times

🗨️ 👤 **Aamee** 1 year, 7 months ago

Selected Answer: D

D sounds pretty legit IMO.

upvoted 1 times

🗨️ 👤 **Aamee** 1 year, 7 months ago

Sorry, after reviewing it again, looks like Option 'B' looks to be more correct ans.

upvoted 1 times

🗨️ 👤 **[Removed]** 1 year, 7 months ago

Selected Answer: B

B.. Use Eventbridge (formally Cloudwatch Events)

https://docs.aws.amazon.com/guardduty/latest/ug/guardduty_findings_cloudwatch.html

upvoted 4 times

🗨️ 👤 **oioi** 1 year, 7 months ago

Selected Answer: B

correct

upvoted 1 times

🗨️ 👤 **oioi** 1 year, 7 months ago

D is correct

upvoted 1 times

A company is using AWS Organizations to implement a multi-account strategy. The company does not have on-premises infrastructure. All workloads run on AWS. The company currently has eight member accounts. The company anticipates that it will have no more than 20 AWS accounts total at any time.

The company issues a new security policy that contains the following requirements:

- No AWS account should use a VPC within the AWS account for workloads.
- The company should use a centrally managed VPC that all AWS accounts can access to launch workloads in subnets.
- No AWS account should be able to modify another AWS account's application resources within the centrally managed VPC.
- The centrally managed VPC should reside in an existing AWS account that is named Ac-count-A within an organization.

The company uses an AWS CloudFormation template to create a VPC that contains multiple subnets in Account-A. This template exports the subnet IDs through the CloudFormation Outputs section.

Which solution will complete the security setup to meet these requirements?

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

Suggested Answer: C

Community vote distribution

C (100%)

 **cumzle_com** 1 year ago

Selected Answer: C

I think this setup is called Core-Shared VPC
so C should be the answer ..
upvoted 1 times

 **vikasj1in** 1 year, 5 months ago


C. Explanation:

using AWS RAM, allows for centralized control over the VPC in Account-A, shared access to subnets with other member accounts, and isolation between the resources of different accounts within the VPC. This aligns well with the specified security requirements.

VPC Peering or Transit Gateway: While VPC peering (Option D) or Transit Gateway (Option B) could facilitate communication between VPCs, they might not provide the level of isolation required in this scenario. With VPC peering, all resources in the peered VPCs have the ability to communicate with each other directly.

CloudFormation Fn::ImportValue: While using CloudFormation's Fn::ImportValue (Option A) is a common approach to share values between stacks, in this scenario, AWS RAM provides a more structured and scalable way to share resources (subnets) across accounts.

upvoted 3 times

 **rahav** 1 year, 6 months ago

Selected Answer: C

RAM is used to share subnets
upvoted 2 times

 **AgboolaKun** 1 year, 7 months ago

Selected Answer: C

Yes, you can use the AWS Resource Access Manager (AWS RAM) to share your subnets and resources in VPC owner, Account-A with other accounts - <https://aws.amazon.com/blogs/networking-and-content-delivery/vpc-sharing-a-new-approach-to-multiple-accounts-and-vpc-management/>
upvoted 3 times

🗨️ 👤 **[Removed]** 1 year, 7 months ago

Selected Answer: C

C is correct. B and D will require multiple VPCs and A is not applicable here
upvoted 2 times

🗨️ 👤 **oioi** 1 year, 7 months ago

Selected Answer: C

correct
upvoted 2 times

A company's security team needs to receive a notification whenever an AWS access key has not been rotated in 90 or more days. A security engineer must develop a solution that provides these notifications automatically.


Which solution will meet these requirements with the LEAST amount of effort?

- A. Deploy an AWS Config managed rule to run on a periodic basis of 24 hours. Select the access-keys-rotated managed rule, and set the maxAccessKeyAge parameter to 90 days. Create an Amazon EventBridge rule with an event pattern that matches the compliance type of NON_COMPLIANT from AWS Config for the managed rule. Configure EventBridge to send an Amazon Simple Notification Service (Amazon SNS) notification to the security team.
- B. Create a script to export a .csv file from the AWS Trusted Advisor check for IAM access key rotation. Load the script into an AWS Lambda function that will upload the .csv file to an Amazon S3 bucket. Create an Amazon Athena table query that runs when the .csv file is uploaded to the S3 bucket. Publish the results for any keys older than 90 days by using an invocation of an Amazon Simple Notification Service (Amazon SNS) notification to the security team.
- C. Create a script to download the IAM credentials report on a periodic basis. Load the script into an AWS Lambda function that will run on a schedule through Amazon EventBridge. Configure the Lambda script to load the report into memory and to filter the report for records in which the key was last rotated at least 90 days ago. If any records are detected, send an Amazon Simple Notification Service (Amazon SNS) notification to the security team.
- D. Create an AWS Lambda function that queries the IAM API to list all the users. Iterate through the users by using the ListAccessKeys operation. Verify that the value in the CreateDate field is not at least 90 days old. Send an Amazon Simple Notification Service (Amazon SNS) notification to the security team if the value is at least 90 days old. Create an Amazon EventBridge rule to schedule the Lambda function to run each day.

Suggested Answer: A

Community vote distribution

A (100%)

  **aescudero51** 7 months, 4 weeks ago

Selected Answer: A

Correct, "A" AWS Config managed rule (access-keys-rotated):

<https://docs.aws.amazon.com/config/latest/developerguide/access-keys-rotated.html>

upvoted 2 times

  **vikasj1in** 11 months, 3 weeks ago

Selected Answer: A

AWS Config managed rule (access-keys-rotated): This managed rule checks whether IAM access keys have been rotated within a specified timeframe. By configuring it to run on a periodic basis of 24 hours and setting the maxAccessKeyAge parameter to 90 days, it will automatically detect access keys that haven't been rotated in 90 or more days.

Amazon EventBridge rule: Create an EventBridge rule with an event pattern that matches the compliance type of NON_COMPLIANT from AWS Config for the access-keys-rotated managed rule. This ensures that EventBridge triggers an action when the IAM access keys are found to be non-compliant.

Amazon SNS Notification: Configure EventBridge to send an SNS notification to the security team when the event pattern matches. This will automatically notify the security team when access keys have not been rotated within the specified timeframe.

upvoted 2 times

  **yorkicurke** 1 year ago

Selected Answer: A

The rest of the options are garbage



upvoted 3 times

  **[Removed]** 1 year, 1 month ago

Selected Answer: A

Yup.. A

upvoted 3 times

  **oioi** 1 year, 1 month ago

Selected Answer: A

correct

upvoted 2 times

A company maintains an open-source application that is hosted on a public GitHub repository. While creating a new commit to the repository, an engineer uploaded their AWS access key and secret access key. The engineer reported the mistake to a manager, and the manager immediately disabled the access key.


The company needs to assess the impact of the exposed access key. A security engineer must recommend a solution that requires the least possible managerial overhead.

Which solution meets these requirements?

- A. Analyze an AWS Identity and Access Management (IAM) use report from AWS Trusted Advisor to see when the access key was last used.
- B. Analyze Amazon CloudWatch Logs for activity by searching for the access key.
- C. Analyze VPC flow logs for activity by searching for the access key.
- D. Analyze a credential report in AWS Identity and Access Management (IAM) to see when the access key was last used.

Suggested Answer: D

Community vote distribution



🗳️ 👤 **46e75dd** 1 month, 1 week ago

Selected Answer: B

With Credentials Report you will get when last used. However, to gain insights about what were the keys used for you'll have to look at the API calls made using those keys. Hence, you should use cloudtrail.

upvoted 2 times

🗳️ 👤 **Just_Ninja** 9 months, 1 week ago

Selected Answer: D

If we will get an Hint in the Question like Business or Enterprise Support, then i will choose A.

But here is no Hint, so D ist the best Solution.

upvoted 1 times

🗳️ 👤 **Davidng88** 9 months, 2 weeks ago

Selected Answer: A

AWS Trusted Advisor provides a comprehensive IAM use report that includes details on when access keys were last used. This report can quickly help you determine if the exposed key was used after it was uploaded to the public repository.

upvoted 1 times

🗳️ 👤 **NSA_Poker** 8 months, 3 weeks ago

TrustAdvisor can just notify you if your key is exposed and also check for existence of IAM user to discourage root access

<https://docs.aws.amazon.com/awssupport/latest/user/trusted-advisor-check-reference.html#exposed-access-keys>

access_key_1_last_used_date is a column in the credential report CSV file.

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_getting-report.html#:~:text=access_key_1_last_used_date

upvoted 1 times

🗳️ 👤 **Sodev** 1 year, 2 months ago

"The company needs to assess the impact of the exposed access key" is not right sentence.

IAM credential report only show, last active, last active region. It is not show what resource or API is called after exposing keys. => To "Assess the impact" need analyzer Cloudtrail logs.

upvoted 1 times

🗳️ 👤 **ion_gee** 1 year, 2 months ago

Selected Answer: D

D is the best Answer. https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_getting-report.html

B is not a good answer in this scenario because cloudwatch should be used for monitoring and analyzing log data, and the access key may not be in plaintext due to security reasons. CloudWatch would be more useful for monitoring specific application or system activity rather than tracking the usage of an exposed access key.

upvoted 1 times

🗨️ 👤 **Raphaello** 1 year, 4 months ago

Selected Answer: B

Best answer is B.

The ask is to " assess the impact of the exposed access key", which mean how the exposed access key has been used.

Credential report does not include such information, it includes information about the credential itself..when it was created, last used, last changed.

Not useful to assess the impact.

upvoted 3 times

🗨️ 👤 **brpjp** 1 year, 5 months ago

Selected Answer: B is correct. Question is to analyze impact of exposed access key. From credential report you know only key last used, but not able to determine how many times key used and what activities performed.

upvoted 1 times

🗨️ 👤 **rahav** 1 year, 6 months ago

Selected Answer: D

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_getting-report.html

upvoted 2 times

🗨️ 👤 **3633f8f** 1 year, 6 months ago

Selected Answer: D

Least effort.

upvoted 2 times

🗨️ 👤 **kejam** 1 year, 7 months ago

Selected Answer: D

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_getting-report.html

upvoted 2 times

🗨️ 👤 **[Removed]** 1 year, 7 months ago

Selected Answer: D

D 99.999999% sure

upvoted 3 times

🗨️ 👤 **oioi** 1 year, 7 months ago

Selected Answer: D

correct

upvoted 1 times

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

How can the security engineer meet these requirements?

- A. Create an IAM policy that prohibits changes to the specific CloudTrail trail and apply the policy to the AWS account root user.
- B. Create an S3 bucket policy in the specified destination account for the CloudTrail trail that prohibits configuration changes from the AWS account root user in the source account.
- C. Create an SCP that prohibits changes to the specific CloudTrail trail and apply the SCP to the appropriate organizational unit or account in Organizations.
- D. Create an IAM policy that prohibits changes to the specific CloudTrail trail and apply the policy to a new IAM group. Have team members use individual IAM accounts that are members of the new IAM group.

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **7c84836** 11 months, 3 weeks ago

why not D?

upvoted 1 times

🗳️ 👤 **helloworldabc** 9 months, 1 week ago

just C

upvoted 1 times

🗳️ 👤 **Raphaello** 1 year, 4 months ago

Selected Answer: C

Correct answer is C.

Use SCP to deny changes to the specific trail, apply the policy to designated OU or accounts.

upvoted 3 times

🗳️ 👤 **vikasj1in** 1 year, 5 months ago

SCPs in AWS Organizations are used to set fine-grained permissions and restrictions on AWS accounts within an organization. They operate at the root level or organizational unit level.

the security engineer can enforce a policy at the organizational level, ensuring that no accounts under the specified organizational unit can make modifications or disable the CloudTrail configuration. While IAM policies and S3 bucket policies can control access to resources, they are typically more focused on granting permissions rather than restricting actions on CloudTrail trails globally across the organization.

Option C, using an SCP, provides centralized control and is well-suited for enforcing organization-wide policies. It ensures that even if DevOps team members have administrative permissions in their individual accounts, they won't be able to modify or disable the specified CloudTrail trail due to the SCP restrictions.

upvoted 1 times

🗳️ 👤 **3633f8f** 1 year, 6 months ago

Selected Answer: C

C is correct.

upvoted 1 times

🗳️ 👤 **Aamee** 1 year, 7 months ago

Selected Answer: C


For sure it should be 'D'.

upvoted 2 times

🗳️ 👤 **Aamee** 1 year, 7 months ago

typo: 'C'.



upvoted 2 times

  **[Removed]** 1 year, 7 months ago

Selected Answer: C

C sounds good

upvoted 2 times

  **oioi** 1 year, 7 months ago

Selected Answer: C

correct

upvoted 1 times

A company's policy requires that all API keys be encrypted and stored separately from source code in a centralized security account. This security account is managed by the company's security team. However, an audit revealed that an API key is stored with the source code of an AWS Lambda function in an AWS CodeCommit repository in the DevOps account.

How should the security team securely store the API key?

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

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **Raphaello** 10 months, 1 week ago

Selected Answer: C

Correct answer is C.

Secrets Manager to store API key securely using KMS key. Grant access to KSM key to Lambda exec role.
upvoted 2 times

🗳️ 👤 **3633f8f** 1 year ago

Selected Answer: C

C is correct among others. In case there is something related to AWS Systems Manager Parameter Store as Secure String would be even better more accurate choice.
upvoted 1 times

🗳️ 👤 **Raphaello** 1 year ago

C

Whenever there is an AWS can do the job, in this scenario Secrets Manager, then it is the right choice.
Environment variable is not the right answer, despite using KMS for encryption.
upvoted 1 times

🗳️ 👤 **Aamee** 1 year, 1 month ago

Selected Answer: C

Def. 'C' is the best practice way.
upvoted 3 times

🗳️ 👤 **[Removed]** 1 year, 1 month ago

Selected Answer: C

C for the win
upvoted 2 times

🗳️ 👤 **oioi** 1 year, 1 month ago

Selected Answer: C

correct
upvoted 1 times


A security engineer is asked to update an AWS CloudTrail log file prefix for an existing trail. When attempting to save the change in the CloudTrail console, the security engineer receives the following error message: "There is a problem with the bucket policy."

What will enable the security engineer to save the change?

- A. Create a new trail with the updated log file prefix, and then delete the original trail. Update the existing bucket policy in the Amazon S3 console with the new log file prefix, and then update the log file prefix in the CloudTrail console.
- B. Update the existing bucket policy in the Amazon S3 console to allow the security engineer's principal to perform PutBucketPolicy, and then update the log file prefix in the CloudTrail console.
- C. Update the existing bucket policy in the Amazon S3 console with the new log file prefix, and then update the log file prefix in the CloudTrail console.
- D. Update the existing bucket policy in the Amazon S3 console to allow the security engineer's principal to perform GetBucketPolicy, and then update the log file prefix in the CloudTrail console.

Suggested Answer: C

Community vote distribution



🗳️ **SHERLOCKAWS** 11 months ago

Selected Answer: C

Explained here >> <https://docs.aws.amazon.com/awscloudtrail/latest/userguide/create-s3-bucket-policy-for-cloudtrail.html#cloudtrail-add-change-or-remove-a-bucket-prefix>

upvoted 2 times

🗳️ **vikasj1in** 11 months, 3 weeks ago

C. Updating Bucket Policy: The error message indicates that there is a problem with the bucket policy. To resolve this, the security engineer needs to update the existing bucket policy in the Amazon S3 console with the new log file prefix. This ensures that the bucket policy is correctly configured to allow CloudTrail to write logs to the specified location.

Updating Log File Prefix in CloudTrail Console: Once the bucket policy is updated, the security engineer can then go back to the CloudTrail console and update the log file prefix there. This will ensure that CloudTrail knows the correct destination in the S3 bucket for storing the log files.

Option C is the correct choice as it addresses the issue by first updating the bucket policy and then updating the log file prefix in the CloudTrail console. The other options involve unnecessary steps or do not directly address the reported problem.

upvoted 1 times

🗳️ **WeepingMaple** 1 year ago

Verify the S3 bucket policy:

Open the S3 bucket policy in the S3 console, CLI, or SDK.

Ensure the policy grants CloudTrail permission to write log files with the new prefix to the specified S3 bucket. Look for permissions like s3:PutObject with the correct bucket prefix included.

upvoted 1 times

🗳️ **vincentsr7** 1 year ago

B. As C does not provide the required privileges to the security engineer

upvoted 2 times

🗳️ **Aamee** 1 year, 1 month ago

Selected Answer: C

C looks legit.


upvoted 2 times

🗳️ **AgboolaKun** 1 year, 1 month ago

Selected Answer: C

The correct answer is C - <https://docs.aws.amazon.com/awscloudtrail/latest/userguide/create-s3-bucket-policy-for-cloudtrail.html#cloudtrail-add-change-or-remove-a-bucket-prefix>



upvoted 3 times

  **[Removed]** 1 year, 1 month ago

Selected Answer: C

<https://docs.aws.amazon.com/awscloudtrail/latest/userguide/create-s3-bucket-policy-for-cloudtrail.html#cloudtrail-add-change-or-remove-a-bucket-prefix>

upvoted 3 times

  **oioi** 1 year, 1 month ago

Selected Answer: B

correct

upvoted 2 times

A company uses AWS Organizations. The company wants to implement short-term credentials for third-party AWS accounts to use to access accounts within the company's organization. Access is for the AWS Management Console and third-party software-as-a-service (SaaS) applications. Trust must be enhanced to prevent two external accounts from using the same credentials. The solution must require the least possible operational effort.

Which solution will meet these requirements?

- A. Use a bearer token authentication with OAuth or SAML to manage and share a central Amazon Cognito user pool across multiple Amazon API Gateway APIs.
- B. Implement AWS IAM Identity Center (AWS Single Sign-On), and use an identity source of choice. Grant access to users and groups from other accounts by using permission sets that are assigned by account.
- C. Create a unique IAM role for each external account. Create a trust policy Use AWS Secrets Manager to create a random external key.
- D. Create a unique IAM role for each external account. Create a trust policy that includes a condition that uses the sts:ExternalId condition key.

Suggested Answer: D

Community vote distribution

D (100%)

 **IPLogic** 7 months ago


Selected Answer: D

The best solution to meet these requirements is D.

Create a unique IAM role for each external account.

Create a trust policy that includes a condition that uses the sts:ExternalId condition key.

upvoted 1 times

 **723993f** 7 months ago

Selected Answer: B

why is everyone saying d, its b

are we forgetting about the saas requirement ? iam roles can't help you access saas apps. identity center will be able auth each user individually using whatever idp, from there onward they can sso into saas as well as mgmt console

upvoted 2 times

 **jakie22332** 7 months, 3 weeks ago

Selected Answer: B

B supports all the requirements with the LEAST operational overhead

upvoted 1 times

 **SHERLOCKAWS** 1 year, 4 months ago

Selected Answer: D

Explained here > https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_create_for-user_externalid.html

upvoted 2 times

 **vikasj1in** 1 year, 5 months ago

Creating a unique IAM role for each external account allows you to grant specific permissions to each external account independently. Including a condition in the trust policy that uses the sts:ExternalId condition key allows you to enhance the trust between the accounts and prevent one external account from using the credentials intended for another external account. The sts:ExternalId condition ensures that the request is accompanied by the expected external ID, adding an extra layer of security.

Options A, B, and C do not specifically address the requirement to prevent two external accounts from using the same credentials and may introduce unnecessary complexity or dependencies.

upvoted 1 times

 **WeepingMaple** 1 year, 6 months ago

Selected Answer: D

What is an external ID:

An external ID is a unique identifier that is managed by a third-party identity provider (IdP). It's used to verify the identity of a user without requiring them to have an AWS IAM account.



Creating a role with an external ID:

You can create a role in your AWS account and specify an external ID source (e.g., SAML provider, OIDC provider).

You can define trust relationships between the role and the external IdP. This ensures that only authorized users with the correct external ID can assume the role.

You can attach IAM policies to the role to grant specific permissions to access AWS resources.



upvoted 2 times

  **kejam** 1 year, 7 months ago

Selected Answer: D

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_create_for-user_externalid.html

upvoted 3 times

  **Aamee** 1 year, 7 months ago

Selected Answer: D

C looks a bit reasonable but with a condition on the role makes it more secured so going with 'D'.

upvoted 1 times

  **[Removed]** 1 year, 7 months ago

Selected Answer: D

D will do it. The rest are distractors / incorrect

upvoted 2 times

  **oioi** 1 year, 7 months ago

Selected Answer: D

correct

upvoted 1 times

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 logs. Deploy the Amazon Athena Redshift connector. Use Athena to query the data from Amazon Redshift and to filter the logs by host.
- B. Specify Amazon CloudWatch as the destination for the access logs. Use Amazon CloudWatch Logs Insights to design a query to filter the logs by host.
- 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.
- D. Specify Amazon CloudWatch as the destination for the access logs. Use Amazon Redshift Spectrum to query the logs and to filter the logs by host.

Suggested Answer: B

Community vote distribution



vikasj1in 11 months, 3 weeks ago

B.

CloudWatch Logs Insights is a fully managed service that enables you to search and analyze your log data efficiently. It allows you to interactively explore and analyze your logs directly in the CloudWatch console.

Operational Efficiency: CloudWatch is a native AWS service that can directly receive and store AWS WAF access logs. Using CloudWatch Logs Insights, you can design and run queries to filter logs based on specific hosts. This provides a quick and efficient way to analyze and monitor AWS WAF logs centrally.

Options C and D involve additional steps such as exporting logs to S3 or using Amazon Redshift Spectrum, which may introduce additional complexity and operational overhead. Option A suggests using Amazon Redshift directly, which may not be the most efficient option for log analysis in this scenario.

Therefore, option B is the most operationally efficient solution for analyzing and filtering AWS WAF access logs in a central location.

upvoted 3 times

3633f8f 1 year ago

Selected Answer: B

B indeed.

upvoted 1 times

WeepingMaplte 1 year ago

Selected Answer: B

Log Insights: Provides a powerful query interface for searching and analyzing WAF logs based on various criteria like IP addresses, user agents, and rule IDs.

upvoted 1 times

AgboolaKun 1 year, 1 month ago

Selected Answer: B

Agree. B is the MOST operational efficiency - <https://aws.amazon.com/blogs/mt/analyzing-aws-waf-logs-in-amazon-cloudwatch-logs/>

upvoted 3 times

🗨️ 👤 **Aamee** 1 year, 1 month ago

Selected Answer: B

Agreed. It asks specifically about the Operational Efficiency on this. Option C seems to be good as well but it takes a bit more steps to setup/configure those steps. Where from Option 'B', you can get it from the CW Insights features.

upvoted 1 times

🗨️ 👤 **Imimi** 1 year, 1 month ago

I think B is the MOST operational efficiency

upvoted 2 times

🗨️ 👤 **[Removed]** 1 year, 1 month ago

Selected Answer: B

voting B

upvoted 2 times

🗨️ 👤 **[Removed]** 1 year, 1 month ago

I'd argue B is more efficient. Less moving parts than C.

<https://aws.amazon.com/blogs/mt/analyzing-aws-waf-logs-in-amazon-cloudwatch-logs/>

upvoted 2 times

🗨️ 👤 **oioi** 1 year, 1 month ago

Selected Answer: C

correct

upvoted 1 times

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.

Suggested Answer: BE

Community vote distribution



Sodev 8 months, 3 weeks ago

In AWS docs:

"Select an IAM role that you want to associate with the instance profile for component permissions on your build and test instances. Image Builder uses these permissions to download and run your components, upload logs to CloudWatch, and perform any additional actions that the components in your recipe specify."

upvoted 1 times

ratsoft 9 months, 2 weeks ago

Selected Answer: BE

<https://docs.aws.amazon.com/imagebuilder/latest/userguide/troubleshooting.html#ts-access-denied>

ause

Possible causes include:

The instance profile does not have the required permissions to access APIs or component resources.

The instance profile role is missing permissions that are required for logging to Amazon S3. Most commonly, this occurs when the instance profile role does not have PutObject permissions for your S3 buckets.

Solution

Depending on the cause, this issue can be resolved as follows:

Instance profile is missing managed policies – Add the missing policies to your instance profile role. Then run the pipeline again.

Instance profile is missing write permissions for S3 bucket – Add a policy to your instance profile role that grants PutObject permissions to write to your S3 bucket. Then run the pipeline again.

upvoted 1 times

Mandla97 10 months, 1 week ago

BE

Cause

Possible causes include:

The instance profile does not have the required permissions to access APIs or component resources.

The instance profile role is missing permissions that are required for logging to Amazon S3. Most commonly, this occurs when the instance profile role does not have PutObject permissions for your S3 buckets.

Solution

Depending on the cause, this issue can be resolved as follows:

Instance profile is missing managed policies – Add the missing policies to your instance profile role. Then run the pipeline again.

Instance profile is missing write permissions for S3 bucket – Add a policy to your instance profile role that grants PutObject permissions to write to your S3 bucket. Then run the pipeline again.

upvoted 1 times

🗳️ 👤 **vikasj1in** 11 months, 3 weeks ago

A. The IAM role used by Amazon EC2 Image Builder needs to have the necessary policies attached to perform the required actions. In this case, the role needs policies such as EC2InstanceProfileForImageBuilder, EC2InstanceProfileForImageBuilderECRContainerBuilds, and AmazonSSMManagedInstanceCore.

D: The IAM role must have the s3:PutObject permission for the specified S3 bucket. This permission is required for storing logs in the S3 bucket.

Options B and E involve attaching policies directly to the instance profile, which is not the recommended approach for Amazon EC2 Image Builder. The IAM role associated with EC2 Image Builder is used for the build process, and it is the role that needs the required permissions.

Option C is not specific to the IAM role or instance profile used by Amazon EC2 Image Builder, and it's generally not recommended to attach broad policies like AWSImageBuilderFullAccess without following the principle of least privilege.

upvoted 1 times

🗳️ 👤 **giancesarini2023** 1 year ago

The correct answer is B/E.

upvoted 1 times

🗳️ 👤 **3633f8f** 1 year ago

Selected Answer: BE

BE choice as Instance Profile >> Role for the Instance on start up - usually -

upvoted 2 times

🗳️ 👤 **WeepingMaple** 1 year ago

Selected Answer: BE

EC2InstanceProfileForImageBuilder, EC2InstanceProfileForImageBuilderECRContainerBuilds, and AmazonSSMManagedInstanceCore are policies.
<https://docs.aws.amazon.com/imagebuilder/latest/userguide/security-iam-awsmanpol.html>

Ensure the IAM roles used by Image Builder have the necessary permissions to access resources involved in the build process, like S3 buckets, EC2 instances, and SSM automation documents.

upvoted 1 times

🗳️ 👤 **Aamee** 1 year ago

Selected Answer: BE

Thanks folks!... got the EC2 Instance profile concept now so def. going with B and E now.

upvoted 1 times

🗳️ 👤 **snowmaggdon** 1 year, 1 month ago

A pipeline is running so that means the engineer's role is not relevant. Answer is BE

upvoted 1 times

🗳️ 👤 **AgboolaKun** 1 year, 1 month ago

Selected Answer: BE

Please note that an instance profile is an IAM role for the EC2 instance. Therefore, the option A which states that "IAM role attached to the engineer" is wrong. Please check this link for more information - <https://docs.aws.amazon.com/imagebuilder/latest/userguide/troubleshooting.html#ts-access-denied>

upvoted 4 times

🗳️ 👤 **Aamee** 1 year, 1 month ago

"Solution:

Depending on the cause, this issue can be resolved as follows:

Instance profile is missing managed policies – Add the missing policies to your instance profile role. Then run the pipeline again.

Instance profile is missing write permissions for S3 bucket – Add a policy to your instance profile role that grants PutObject permissions to write to your S3 bucket. Then run the pipeline again."

The sol. states that it's EC2 Instance Profile "Role" as per their documentation. Whereas, in option B and E, it states EC2 Instance profile only. Does it mean the same thing? Can someone pls. help clarify on this.

upvoted 1 times

  **ion_gee** 8 months, 4 weeks ago

AN EC2 Instance profile is an IAM role created to be used by an EC2 instance and attached to the Ec2 instance.

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_use_switch-role-ec2.html

upvoted 2 times



  **[Removed]** 1 year, 1 month ago

Selected Answer: BE

B and E as per the following

<https://docs.aws.amazon.com/imagebuilder/latest/userguide/troubleshooting.html#ts-access-denied>

upvoted 3 times

  **oioi** 1 year, 1 month ago

Selected Answer: DE

correct

upvoted 1 times


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

Suggested Answer: A

Community vote distribution



IPLogic 7 months ago

Selected Answer: A

The best solution to meet these criteria is A.

A customer managed key that uses customer provided key material.

upvoted 1 times

cloudescalate 1 year, 1 month ago

Selected Answer: A

You can only schedule the deletion of a customer managed key. You cannot delete AWS managed keys or AWS owned keys.

<https://docs.aws.amazon.com/kms/latest/developerguide/deleting-keys.html>

upvoted 2 times

Raphaello 1 year, 4 months ago

Selected Answer: A

Correct answer is A.

You can select your KMS key with imported key material expiration date.

<https://docs.aws.amazon.com/kms/latest/developerguide/importing-keys-managing.html>

upvoted 1 times

vikasj1in 1 year, 5 months ago

AWS Managed Key (AWS KMS): AWS managed keys (AWS KMS keys) are created, managed, and rotated by AWS KMS. AWS automatically handles key rotation for these keys. With an AWS managed key, you can set the key rotation interval, and AWS KMS will automatically rotate the key material.

Expiration: While AWS managed keys don't have an explicit "expiration" property, you can achieve similar functionality by configuring key rotation every 90 days. This ensures that the key material is automatically rotated, effectively providing a new key every 90 days.

Options A and B refer to customer managed keys, and the expiration of key material would need to be managed manually by the customer. Option D mentions GnuPG, which is not applicable for managing AWS EBS volume encryption keys.

Therefore, option C (AWS managed key) is the most suitable choice for this scenario.

upvoted 1 times

rahav 1 year, 6 months ago

Selected Answer: A

You may set an expiration period for an imported key. AWS KMS will automatically delete the key material after the expiration period. You can also delete imported key material on demand. In both cases the key material itself is deleted but the KMS key reference in AWS KMS and associated metadata are retained so that the key material can be re-imported in the future. Keys generated by AWS KMS do not have an expiration time and cannot be deleted immediately; there is a mandatory 7 to 30 day wait period. All customer managed KMS keys, regardless of whether the key material was imported, can be manually disabled or scheduled for deletion. In this case the KMS key itself is deleted, not just the underlying key material.

<https://aws.amazon.com/kms/faqs/>

upvoted 3 times

🗨️ 👤 **WeepingMaple** 1 year, 6 months ago

Selected Answer: A

A will be the answer. The key word in the question is automatically expires. For answer B and C, it does not have the expiration date option. It only has the rotate option.

upvoted 1 times

🗨️ 👤 **vincentsr7** 1 year, 6 months ago

Option C.

A customer managed key (option A) that uses customer provided key material would not have the automatic expiration capability by default.

upvoted 1 times

🗨️ 👤 **Daniel76** 1 year, 6 months ago

When you import key material, you can set an optional expiration time.

<https://docs.aws.amazon.com/kms/latest/developerguide/importing-keys-managing.html>

upvoted 2 times

🗨️ 👤 **Aamee** 1 year, 7 months ago

Selected Answer: A

A sounds legit.

upvoted 1 times

🗨️ 👤 **[Removed]** 1 year, 7 months ago

Selected Answer: A

Definitely A

upvoted 1 times

🗨️ 👤 **AgboolaKun** 1 year, 7 months ago

Selected Answer: A

A is correct - <https://awscli.amazonaws.com/v2/documentation/api/latest/reference/kms/import-key-material.html>

upvoted 1 times

🗨️ 👤 **oioi** 1 year, 7 months ago

Selected Answer: B

correct

upvoted 1 times

A security engineer is building a Java application that is running on Amazon EC2. The application communicates with an Amazon RDS instance and authenticates with a user name and password.


Which combination of steps can the engineer take to protect the credentials and minimize downtime when the credentials are rotated? (Choose two.)

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

Suggested Answer: CE

Community vote distribution

CE (100%)

  **aescudero51** 6 months, 3 weeks ago

Selected Answer: CE

My answer is C - E

Configure automatic rotation of credentials in AWS Secrets Manager (Option C): This ensures that the credentials are updated automatically without requiring manual intervention. AWS Secrets Manager provides secure, managed secrets for applications, and it can be configured to rotate credentials at regular intervals.

Configure the Java application to catch a connection failure and make a call to AWS Secrets Manager to retrieve updated credentials when the password is rotated (Option E): This approach ensures that the application can seamlessly switch to the new credentials when they are rotated, minimizing downtime. The Java application can be designed to catch connection failures and then retrieve the updated credentials from AWS Secrets Manager, ensuring continuous operation.

upvoted 1 times

  **vikasj1in** 11 months, 3 weeks ago

AWS Secrets Manager (Option C and E): AWS Secrets Manager provides a solution for managing and rotating sensitive information, such as database credentials. You can configure automatic rotation of credentials in AWS Secrets Manager, and the Java application can catch a connection failure and make a call to Secrets Manager to retrieve updated credentials when the password is rotated.

Systems Manager Parameter Store (Option D): While storing credentials in an encrypted string parameter in AWS Systems Manager Parameter Store is a valid approach, Secrets Manager provides a more specialized solution for credential rotation.

Option A involves storing ciphertext in Amazon S3, which adds complexity and may not be as secure as using dedicated services like AWS Secrets Manager.

Option B suggests manually updating the credential in Systems Manager Parameter Store and notifying the engineer, which is less automated and may introduce downtime.

upvoted 2 times

  **3633f8f** 1 year ago

Selected Answer: CE

Systems Manager provides integration with RDS and in combination with Java Try and Catch makes possible rotate credentials as frequent as required.

upvoted 1 times

🗨️ 👤 **rxhan** 10 months, 1 week ago

Secrets*

upvoted 1 times

🗨️ 👤 **Aamee** 1 year ago

Selected Answer: CE

Def. C and E makes perfect comb.

upvoted 1 times

🗨️ 👤 **Aamee** 1 year, 1 month ago

V confusing btw CE and DE. The question states about protecting the creds which gives the hint towards option D since it talks about leveraging KMS keys with SSM parameter store options too... No doubt on option E cuz that looks more reasonable but there's a confusion over option C or D... :(

upvoted 1 times

🗨️ 👤 **AgboolaKun** 1 year, 1 month ago

C is correct because Systems Manager Parameter Store cannot be used for key rotation. Key rotation is a feature of Secrets Manager. I hope that helps.

upvoted 4 times

🗨️ 👤 **Aamee** 1 year ago

Ah ok, got it now, thanks so much AgboolaKun! :)

upvoted 1 times

🗨️ 👤 **[Removed]** 1 year, 1 month ago

Selected Answer: CE

I'll vote C and E. Secrets Manager for rotation

upvoted 2 times

🗨️ 👤 **oioi** 1 year, 1 month ago

Selected Answer: CE

correct

upvoted 2 times

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

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

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

- A. Create a local individual break glass IAM user for the security team. Create a trail in AWS CloudTrail that has Amazon CloudWatch Logs turned on. Use Amazon EventBridge to monitor local user activities.
- B. Create a break glass EC2 key pair for the AWS account. Provide the key pair to the security team. Use AWS CloudTrail to monitor key pair activity. Send notifications to the security team by using Amazon Simple Notification Service (Amazon SNS).
- C. Create a break glass IAM role for the account. Allow security team members to perform the AssumeRoleWithSAML operation. Create an AWS CloudTrail trail that has Amazon CloudWatch Logs turned on. Use Amazon EventBridge to monitor security team activities.
- D. Create a local individual break glass IAM user on the operating system level of each workload instance. Configure unrestricted security groups on the instances to grant access to the break glass IAM users.
- E. Configure AWS Systems Manager Session Manager for Amazon EC2. Configure an AWS CloudTrail filter based on Session Manager. Send the results to an Amazon Simple Notification Service (Amazon SNS) topic.

Suggested Answer: AE

Community vote distribution

AE (82%)

CE (18%)

kejam Highly Voted 1 year, 7 months ago

Selected Answer: AE

A and E

"Although the use and creation of AWS IAM users is highly discouraged, break glass users are an exception.

To ensure human break-glass access to your environment, we recommend that you create the following in your AWS organization:

At least two IAM users..."

<https://docs.aws.amazon.com/whitepapers/latest/organizing-your-aws-environment/break-glass-access.html>

upvoted 9 times

Aamee 1 year, 6 months ago

I think ur option A might be right here cuz it asks about the IAM break-glass user and not any roles for the whole security team if I understood the question correctly...

upvoted 1 times

yorkicurke Highly Voted 1 year, 6 months ago

Selected Answer: AE

and why i did not go for C;

Because it relies on SAML for the AssumeRoleWithSAML operation. Question mentions that there might be SAML errors. If SAML is not functioning correctly, then the AssumeRoleWithSAML operation would also fail. This means that the security team members would not be able to assume the break glass IAM role when needed, defeating the purpose of having a break glass user for emergency access.

Peace Out:)

upvoted 6 times

youonebe 5 months, 3 weeks ago

Good argument!

upvoted 1 times

IPLogic Most Recent 7 months ago

Selected Answer: AE


The best combination of solutions to meet these 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 and 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.

upvoted 1 times

  **catblack** 10 months, 1 week ago

A & E why:

Option A:

Creating a local individual break glass IAM user ensures that there is a dedicated user with the necessary permissions to access the AWS account in case of SAML errors.

Using AWS CloudTrail with Amazon CloudWatch Logs allows for detailed logging of all activities performed by the break glass user.

Amazon EventBridge can be used to monitor these activities and trigger alerts or actions as needed.

Option E:

AWS Systems Manager Session Manager provides a secure and auditable way to access Amazon EC2 instances without the need for direct SSH or RDP access.

Configuring AWS CloudTrail to filter based on Session Manager activities ensures that all access sessions are logged.

Sending these logs to an Amazon SNS topic ensures that the security team is promptly notified of any break glass user activities.

upvoted 1 times

  **aescudero51** 1 year ago

Selected Answer: CE

My answer is C & E

C. IAM Role with AssumeRoleWithSAML: This creates a secure break glass user in the form of an IAM role. Security team members can assume this role using their existing SAML credentials, eliminating the need for separate credentials and enhancing security. CloudTrail with CloudWatch Logs ensures activity logging, and EventBridge allows for further monitoring of security team actions within the assumed role.


E. Session Manager with CloudTrail Filter: Session Manager provides secure access to EC2 instances without needing SSH keys. Filtering CloudTrail logs based on Session Manager actions specifically captures break glass user activity on the instances. Sending these logs to SNS allows the security team to receive notifications.

upvoted 1 times

  **helloworldabc** 9 months, 1 week ago

just AE

upvoted 1 times

  **minTwin** 1 year, 2 months ago

Selected Answer: AE

A. Create a local individual break glass IAM user for the security team. This step allows the security team to have a user account they can use in emergencies. It's a dedicated account that's different from the standard federated user accounts.

E. Configure AWS Systems Manager Session Manager for Amazon EC2. Session Manager will allow the security team to access EC2 instances securely without the need for SSH keys or open security groups. This service also integrates with AWS CloudTrail to log all session activity, which meets the requirement for logging and sending logs to the security team.

upvoted 2 times

  **mynickc** 1 year, 5 months ago

Selected Answer: AE

read the question carefully, it is asking for a user incase SAML error so assuming role with AssumeRoleWithSAML is not gonna work.

upvoted 2 times

  **vikasj1in** 1 year, 5 months ago

C, E.

Creating a break glass IAM role allows for temporary access when needed. Allowing security team members to perform the AssumeRoleWithSAML operation ensures that the break glass user can assume the role during incidents. Configuring AWS CloudTrail with CloudWatch Logs turned on allows for the logging of activities, and EventBridge can be used to monitor those logs for security team activities. Configuring CloudTrail filters based on Session Manager actions allows logging of activities, and sending the results to an SNS topic can notify the security team.

A & B involve local user or key pair management, which may not be as scalable or auditable compared to using IAM roles and Systems Manager Session Manager.

D suggests creating local individual IAM users on the operating system level, which is not the recommended approach, as it's more challenging to manage and audit compared to IAM roles and System Manager Session Manager.

upvoted 1 times

🗳️ 👤 **Daniel76** 1 year, 6 months ago

Selected Answer: AE

The rest of the options:

B- key pair is very vulnerable. very often, breakglass is sealed in physical envelope and kept in safe.

C- the question requires breakglass user. This option did not provide any user but role.

D- the question require breakglass user at account level but this option provide instance level. Besides, unrestricted security group for all such instances make them vulnerable to password guessing.

The best approach is still create breakglass at account level, seal the breakglass accounts with procedure and physical security, use cloudtrail to ensure its usage is accountable and notification to the entire security team is sent via SNS topic. The account level user allows breakglass user to access to all EC2 instances through session manager.

upvoted 3 times

🗳️ 👤 **brpjp** 1 year, 6 months ago

Correct Answer D & E:

Question is to log any activities of the break glass user and send the logs to a security team. Because of sending logs to security team, security can not be a break glass user to have adequate segregation of duties. Answer A, B and C refer to security team be a break glass user.

So correct answer is D and E.

upvoted 1 times

🗳️ 👤 **tayman** 1 year, 6 months ago

Selected Answer: AE

Vote for A and E.

upvoted 1 times

🗳️ 👤 **dexterryu** 1 year, 6 months ago

Selected Answer: AE

AE are correct. C would not work in the case of SAML issues which the question specifically states is the purpose.

upvoted 2 times

🗳️ 👤 **dexterryu** 1 year, 6 months ago

AE are correct. C would not work in the case of SAML issues which the question specifically states is the purpose.

upvoted 2 times

🗳️ 👤 **AgboolaKun** 1 year, 6 months ago

Selected Answer: CE

CE are the correct answers to this question. Folks choosing AE need to read these "A break glass role that is deployed to all the accounts in the organization, and that can only be 'assumed' by the break glass users from the management account. These roles are needed to allow access from the management account to apply and update guardrails, to troubleshoot and resolve issues with the automation tooling from the security tooling account, or to remediate security and operational issues in one of the member accounts in the AWS organization." sentences from the following link - <https://docs.aws.amazon.com/whitepapers/latest/organizing-your-aws-environment/break-glass-access.html>

upvoted 1 times

🗳️ 👤 **Aamee** 1 year, 6 months ago

But if you read the question again, this is what they're asking specifically:

"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."

When even the break glass user doesn't exist then how can the role be a choice here for this usecase??... You must need to create a local users first inorder to grant the role to it right?... I know it's still v confusing but that's how I interpreted this question..

upvoted 2 times

🗳️ 👤 **[Removed]** 1 year, 6 months ago

Selected Answer: AE

Its A rather than C due to SAML att while using IAM roles.



upvoted 2 times

🗳️ 👤 **marco25** 1 year, 6 months ago

Selected Answer: AE

A is correct need local user in case same is broken

upvoted 3 times

  **Aamee** 1 year, 7 months ago

Selected Answer: CE

C and E makes a good combo imo.

upvoted 1 times

A security engineer is working with a product team building a web application on AWS. The application uses Amazon S3 to host the static content, Amazon API Gateway to provide RESTful services, and Amazon DynamoDB as the backend data store. The users already exist in a directory that is exposed through a SAML identity provider.

Which combination of the following actions should the engineer take to allow users to be authenticated into the web application and call APIs? (Choose three.)

- A. Create a custom authorization service using AWS Lambda.
- B. Configure a SAML identity provider in Amazon Cognito to map attributes to the Amazon Cognito user pool attributes.
- C. Configure the SAML identity provider to add the Amazon Cognito user pool as a relying party.
- D. Configure an Amazon Cognito identity pool to integrate with social login providers.
- E. Update DynamoDB to store the user email addresses and passwords.
- F. Update API Gateway to use a COGNITO_USER_POOLS authorizer.

Suggested Answer: BCF

Community vote distribution

BCF (100%)

 **Daniel76** 1 year ago

Selected Answer: BCF

For API to refer to Cognito user pool, use "COGNITO_USER_POOLS" authorizer

<https://docs.aws.amazon.com/apigateway/latest/developerguide/apigateway-enable-cognito-user-pool.html>

For Cognito user pool to act as relying party to SAML IdP

<https://docs.aws.amazon.com/cognito/latest/developerguide/cognito-user-pools-saml-idp.html>

Other options:

A - "As an alternative to using IAM roles and policies or Lambda authorizers (formerly known as custom authorizers), you can use an Amazon Cognito user pool to control who can access your API in Amazon API Gateway."

<https://docs.aws.amazon.com/apigateway/latest/developerguide/apigateway-integrate-with-cognito.html>

D - social login users not required for this question.

E - Dynamo DB is irrelevant- never store password in db without all the additional overheads required to keep them secure.

upvoted 2 times

 **rahav** 1 year ago

Selected Answer: BCF

BCF is logical here

upvoted 1 times

 **[Removed]** 1 year, 1 month ago

Selected Answer: BCF

BCF. This was on the other exam topics practice set

upvoted 2 times

 **oioi** 1 year, 1 month ago

Selected Answer: BCF

correct

upvoted 1 times

A company needs to improve its ability to identify and prevent IAM policies that grant public access or cross-account access to resources. The company has implemented AWS Organizations and has started using AWS Identity and Access Management Access Analyzer to refine overly broad access to accounts in the organization.

A security engineer must automate a response in the company's organization for any newly created policies that are overly permissive. The automation must remediate external access and must notify the company's security team.

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

- A. Create an AWS Step Functions state machine that checks the resource type in the finding and adds an explicit Deny statement in the trust policy for the IAM role. Configure the state machine to publish a notification to an Amazon Simple Notification Service (Amazon SNS) topic.
- B. Create an AWS Batch job that forwards any resource type findings to an AWS Lambda function. Configure the Lambda function to add an explicit Deny statement in the trust policy for the IAM role. Configure the AWS Batch job to publish a notification to an Amazon Simple Notification Service (Amazon SNS) topic.
- C. In Amazon EventBridge, create an event rule that matches active IAM Access Analyzer findings and invokes AWS Step Functions for resolution.
- D. In Amazon CloudWatch, create a metric filter that matches active IAM Access Analyzer findings and invokes AWS Batch for resolution.
- E. Create an Amazon Simple Queue Service (Amazon SQS) queue. Configure the queue to forward a notification to the security team that an external principal has been granted access to the specific IAM role and has been blocked.
- F. Create an Amazon Simple Notification Service (Amazon SNS) topic for external or cross-account access notices. Subscribe the security team's email addresses to the topic.

Suggested Answer: ACF

Community vote distribution

ACF (88%)

13%

🗳️ 👤 [Removed] **Highly Voted** 👍 1 year, 1 month ago

Selected Answer: ACF

ACF make thr most sense Herr. See below:

<https://aws.amazon.com/blogs/compute/orchestrating-a-security-incident-response-with-aws-step-functions/>
upvoted 5 times

🗳️ 👤 [Removed] 1 year, 1 month ago

the* here*

upvoted 1 times

🗳️ 👤 **tayman** **Most Recent** 🕒 1 year ago

Selected Answer: ACF

Vote for A C F

upvoted 1 times

🗳️ 👤 **Aamee** 1 year, 1 month ago

Selected Answer: ACF

ACF most likely.

upvoted 1 times

🗳️ 👤 **oioi** 1 year, 1 month ago

Selected Answer: CEF

correct

upvoted 1 times

A security engineer is configuring a mechanism to send an alert when three or more failed sign-in attempts to the AWS Management Console occur during a 5-minute period. The security engineer creates a trail in AWS CloudTrail to assist in this work.

Which solution will meet these requirements?

- A. In CloudTrail, turn on Insights events on the trail. Configure an alarm on the insight with eventName matching ConsoleLogin and errorMessage matching "Failed authentication". Configure a threshold of 3 and a period of 5 minutes.
- B. Configure CloudTrail to send events to Amazon CloudWatch Logs. Create a metric filter for the relevant log group. Create a filter pattern with eventName matching ConsoleLogin and errorMessage matching "Failed authentication". Create a CloudWatch alarm with a threshold of 3 and a period of 5 minutes.
- C. Create an Amazon Athena table from the CloudTrail events. Run a query for eventName matching ConsoleLogin and for errorMessage matching "Failed authentication". Create a notification action from the query to send an Amazon Simple Notification Service (Amazon SNS) notification when the count equals 3 within a period of 5 minutes.
- D. In AWS Identity and Access Management Access Analyzer, create a new analyzer. Configure the analyzer to send an Amazon Simple Notification Service (Amazon SNS) notification when a failed sign-in event occurs 3 times for any IAM user within a period of 5 minutes.

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ 👤 **IPLogic** 7 months ago

Selected Answer: B

The keyword in answer B is threshold for CloudWatch Alarm.

upvoted 2 times

🗳️ 👤 **cloudbusting** 1 year, 5 months ago

Because it says alert the answer is B

upvoted 1 times

🗳️ 👤 **brpjp** 1 year, 6 months ago

C is correct. Because SNS generated using detail information on finding that help security, while only CloudWatch Alarm generated do not have information as with SNS notification.

upvoted 1 times

🗳️ 👤 **Daniel76** 1 year, 6 months ago

Selected Answer: B

This is how it is done:

<https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudwatch-alarms-for-cloudtrail.html#cloudwatch-alarms-for-cloudtrail-signin>

upvoted 2 times

🗳️ 👤 **Daniel76** 1 year, 6 months ago

For option c, you will need Eventbridge, Lambda on top of SNS topic.

<https://stackoverflow.com/questions/62823521/i-need-to-create-alerts-based-on-the-results-returned-by-queries-in-amazon-athen>

Option b doesnt include SNS topic, but that is fine because the question ask for "alert" (you can find it in the console) but not "notification".

upvoted 1 times

🗳️ 👤 **rahav** 1 year, 6 months ago

Selected Answer: B

B is the answer. need an Alarm here

upvoted 1 times

🗳️ 👤 **Aamee** 1 year, 7 months ago

Selected Answer: B


CW alarm is best suited here for this scenario.

upvoted 1 times

🗳️ 👤 **[Removed]** 1 year, 7 months ago

Selected Answer: B

B it is. Insights does not do alarming
upvoted 3 times

  **oioi** 1 year, 7 months ago

Selected Answer: B

correct
upvoted 1 times

A company's security engineer is developing an incident response plan to detect suspicious activity in an AWS account for VPC hosted resources. The security engineer needs to provide visibility for as many AWS Regions as possible.

Which combination of steps will meet these requirements MOST cost-effectively? (Choose two.)

- A. Turn on VPC Flow Logs for all VPCs in the account.
- B. Activate Amazon GuardDuty across all AWS Regions.
- C. Activate Amazon Detective across all AWS Regions.
- D. Create an Amazon Simple Notification Service (Amazon SNS) topic. Create an Amazon EventBridge rule that responds to findings and publishes the findings to the SNS topic.
- E. Create an AWS Lambda function. Create an Amazon EventBridge rule that invokes the Lambda function to publish findings to Amazon Simple Email Service (Amazon SES).

Suggested Answer: BD

Community vote distribution

BD (67%)

AB (33%)

 **AgboolaKun** Highly Voted 1 year, 7 months ago

Selected Answer: BD

BD are the correct options here. The keywords here are "developing an incident response plan to detect suspicious activity". There is no better way to develop incident response plan without providing a way for the relevant stakeholders to take actions or respond to suspicious activities.

B is an obvious option because GuardDuty can monitor and analyze API calls across all AWS Regions, and network activities found in Amazon CloudTrail Events, Amazon VPC Flow Logs, and DNS Logs. Therefore, option A is not needed since GuardDuty monitoring activities include the VPC Flow Logs.

There is no better way to respond to the findings generated by GuardDuty than the services described in option D.

upvoted 10 times

 **61cfe5f** 11 months ago

https://docs.aws.amazon.com/guardduty/latest/ug/guardduty_data-sources.html#guardduty_vpc

"When you enable GuardDuty, it immediately starts analyzing your VPC flow logs from Amazon EC2 instances within your account. It consumes VPC flow log events directly from the VPC Flow Logs feature through an independent and duplicative stream of flow logs. This process does not affect any of your existing flow logs configuration."

upvoted 2 times

 **IPLogic** 7 months ago


So you have turn on the VPC Flow Logs. its A.

upvoted 1 times

 **TareDHakim** 6 months ago


No you don't need to enable it for Guard Duty will stream those regardless if you enable it or not.

upvoted 1 times

 **Aamee** 1 year, 6 months ago

Ok, but why the Detective svc. wasn't a good choice here as it's for developing an incident response plan to 'detect' right?... Agree with option D on the other hand cuz it makes sense..

upvoted 1 times

 **navid1365** Highly Voted 1 year, 1 month ago

Selected Answer: AB

A and B are correct for sure. The question does not mention anything about "notifications" or "communications", so D is incorrect.

A is correct because the question mentions visibility and detection. You get visibility into network traffic with VPC flow logs.

B is correct because because we need to "detection" threats. GaurdDuty is a threat detection capability.

upvoted 5 times

🗨️ **BabaP** Most Recent 1 week, 1 day ago

Selected Answer: AB

VPC Flow logs are not enabled by default. therefore you need to enable them before guardDuty

upvoted 1 times

🗨️ **Scunningham99** 4 weeks, 1 day ago

Selected Answer: BD

B. Activate Amazon GuardDuty across all AWS Regions

D. Create an Amazon Simple Notification Service (Amazon SNS) topic. Create an Amazon EventBridge rule that responds to findings and publishes the findings to the SNS topic.

Here's why:

Amazon GuardDuty (Option B):

Provides continuous security monitoring with minimal operational overhead

Uses machine learning to detect threats and suspicious activities

Can be enabled across all regions from a single management account

Has a pay-as-you-go pricing model with no upfront costs

Automatically analyzes various AWS logs (VPC Flow Logs, CloudTrail, DNS logs) without having to explicitly enable them

EventBridge rule with SNS (Option D):

Provides a serverless way to route GuardDuty findings

SNS is very cost-effective for notification delivery

Allows flexible integration with multiple notification endpoints

Can easily scale across regions

upvoted 1 times

🗨️ **bd38daf** 1 month, 4 weeks ago

Selected Answer: AB

VPC Flow Logs need to be enabled first in order for them to be analyzed by GuardDuty.

upvoted 1 times

🗨️ **GDuque** 5 months, 3 weeks ago

Selected Answer: AB

Using SNS + Evenbridge does not seem the most analytical and proper way to do it. It does not contribute to detect threats by itself. I think Guardduty + VPC flows is much more cost effective and straightforward solution. And with VPC flows you must only pay for data storage.

upvoted 1 times

🗨️ **TareDHakim** 6 months ago

Selected Answer: BD

You don't need to enable VPC flow logs for Amazon GuardDuty because GuardDuty already pulls data from VPC Flow Logs. GuardDuty is a threat detection service that monitors for malicious activity in AWS accounts and workloads. It integrates with other AWS security services, including: Amazon CloudTrail, Amazon VPC flow Logs, and AWS WAF.

upvoted 1 times

🗨️ **TareDHakim** 6 months ago

<https://aws.amazon.com/guardduty/faqs/#:~:text=No%2C%20GuardDuty%20pulls%20independent%20data%20streams%20directly%20from%20CloudTrai>

upvoted 1 times

🗨️ **IPLogic** 7 months ago

Selected Answer: AB

The most cost-effective combination of steps to meet these requirements.

A. Turn on VPC Flow Logs for all VPCs in the account.

B. Activate Amazon GuardDuty across all AWS Regions.

VPC Flow Logs provides detailed visibility into network traffic with your VPs. This is a cost effective way to monitor and log network activity, which is important for detecting suspicious behaviour.

Amazon GuardDuty is a managed threat detection service that continuously monitors for malicious activity and unauthorized behaviour. Activating across all AWS regions ensures comprehensive coverage and visibility into potential security threats.

upvoted 1 times

🗨️ 👤 **minTwin** 1 year, 2 months ago

Selected Answer: BD

B and D seems like the most straightforward and cost effective solution
upvoted 1 times

🗨️ 👤 **lightrod** 1 year, 4 months ago

Selected Answer: BD

GuardDuty analyzes VPC flow logs regardless of if you have turned them on or not
upvoted 3 times

🗨️ 👤 **rahav** 1 year, 6 months ago

Selected Answer: BD

VPC Flow logs are very expensive.... Guardduty is the right tool to do that with eventbridge
upvoted 1 times

🗨️ 👤 **WeepingMaple** 1 year, 6 months ago

Selected Answer: BD

A. Turn on VPC Flow Logs for all VPCs in the account: While VPC Flow Logs offer detailed information about network traffic, analyzing and storing logs for all VPCs across Regions can incur significant storage and processing costs.
C. Activate Amazon Detective across all AWS Regions: Detective focuses on root cause analysis and investigation, which might be overkill for initial detection and notification. Additionally, its per-hour billing model can quickly become expensive for continuous monitoring across multiple Regions.
E. Create an AWS Lambda function for publishing findings to SES: While Lambda offers flexibility, creating and maintaining a custom Lambda function specifically for publishing findings can add development and operational overhead compared to the readily available options with EventBridge and SNS.

upvoted 1 times

🗨️ 👤 **Aamee** 1 year, 7 months ago

Selected Answer: AB

AB options best suited. Self-explanatory too.
upvoted 1 times

🗨️ 👤 **[Removed]** 1 year, 7 months ago

Selected Answer: AB

AB are correcto
upvoted 1 times

🗨️ 👤 **oioi** 1 year, 7 months ago

Selected Answer: AB

correct
upvoted 1 times

A company stores images for a website in an Amazon S3 bucket. The company is using Amazon CloudFront to serve the images to end users. The company recently discovered that the images are being accessed from countries where the company does not have a distribution license.

Which actions should the company take to secure the images to limit their distribution? (Choose two.)

- A. Update the S3 bucket policy to restrict access to a CloudFront origin access control (OAC).
- 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.

Suggested Answer: AC

Community vote distribution



AC (93%)

7%

  **[Removed]**  1 year, 7 months ago

Selected Answer: AC

A so the object requests do not bypass Cloudfront, and C for georestrictions. Careful with this oioi fella
upvoted 9 times

  **Aamee** 1 year, 7 months ago

Fully agreed on ur 'oioi' feedback :D..
upvoted 1 times

  **DSEexam**  8 months ago

Selected Answer: AC

A and C makes more sense
upvoted 1 times



  **cumzle_com** 1 year ago

Selected Answer: AC

why not E : The "Restrict Viewer Access" option in CloudFront is designed to require signed URLs or signed cookies, which is not directly related to geo-restrictions. Geo-restriction settings in CloudFront should be used instead.
upvoted 1 times

  **sema2232** 1 year ago

why E is incorrect ?
upvoted 2 times

  **phmeeeee** 2 months, 3 weeks ago

S3 bucket policies don't natively support geo-based restrictions.
upvoted 1 times


  **Raphaello** 1 year, 4 months ago

Selected Answer: AC

Restrict access to CF distro through OAC might not directly help with geo-restriction, until you pick option C alongside it. It enforces flow to be only through CF, where the geo-restriction in place.


AC are correct.

upvoted 1 times

  **Gafa255** 1 year, 5 months ago



Selected Answer: AC

AC correct options. <https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/georestrictions.html>
upvoted 1 times

  **Gafa255** 1 year, 5 months ago

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/private-content-restricting-access-to-s3.html>



upvoted 1 times

  **kejam** 1 year, 7 months ago

Selected Answer: AC

<https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/georestrictions.html>

upvoted 3 times



  **Aamee** 1 year, 7 months ago

Selected Answer: AC

A describes how to limit the access via the policy to bound the access within OAC.



C describes about using the geo restriction based R53 policy you can use to limit the access on the unwanted countries.

upvoted 1 times

  **Aamee** 1 year, 7 months ago

typo: It's Geo restriction list in CloudFront and not R53 policy, my bad!...

upvoted 1 times

  **oioi** 1 year, 7 months ago

Selected Answer: CE

correct

upvoted 1 times

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.

Suggested Answer: B

Community vote distribution

B (90%)

10%

 **ykhan321** Highly Voted 1 year, 6 months ago

Did someone take the test recently? How many questions appeared from here?
upvoted 6 times

 **Davidng88** Most Recent 9 months, 3 weeks ago

Selected Answer: B

Network ACLs are stateless, meaning you need to explicitly allow both inbound and outbound traffic. While the inbound rules might be correctly configured, the outbound rules might be blocking the return traffic. Allowing outbound traffic to ephemeral ports (typically ports 1024-65535) ensures that the response traffic to the vendors.
upvoted 1 times


 **NoCrapEva** 1 year, 4 months ago

Selected Answer: B


Ephemeral ports are necessary for certain network responses and are dependant on the client type (O/S)...
The client that initiates the request chooses the ephemeral port range. The range varies depending on the client's operating system.
Many Linux kernels (including the Amazon Linux kernel) use ports 32768-61000.
Requests originating from Elastic Load Balancing use ports 1024-65535.
Windows operating systems through Windows Server 2003 use ports 1025-5000.
Windows Server 2008 and later versions use ports 49152-65535.
A NAT gateway uses ports 1024-65535.
AWS Lambda functions use ports 1024-65535.
For example, if a request comes into a web server in your VPC from a Windows 10 client on the internet, your network ACL must have an outbound rule to enable traffic destined for ports 49152-65535.
REF:
<https://docs.aws.amazon.com/vpc/latest/userguide/vpc-network-acls.html#nacl-ephemeral-ports>
upvoted 1 times

 **Jamshif01** 1 year, 5 months ago

I don't understand why they are called 'ephemeral' ports
upvoted 1 times

 **rxhan** 1 year, 4 months ago

They change and are random.
upvoted 1 times

 **tayman** 1 year, 6 months ago

Selected Answer: B

Definitely B.
upvoted 3 times

🗨️ 👤 **Aamee** 1 year, 7 months ago

Selected Answer: B

Agreed with B.

upvoted 1 times

🗨️ 👤 **[Removed]** 1 year, 7 months ago

Selected Answer: B

B. You must allow the ephemeral ports in the outbound NACL for the CIDR range.

upvoted 4 times

🗨️ 👤 **oioi** 1 year, 7 months ago

Selected Answer: D

correct

upvoted 1 times

A company uses infrastructure as code (IaC) to create AWS infrastructure. The company writes the code as AWS CloudFormation templates to deploy the infrastructure. The company has an existing CI/CD pipeline that the company can use to deploy these templates.

After a recent security audit, the company decides to adopt a policy-as-code approach to improve the company's security posture on AWS. The company must prevent the deployment of any infrastructure that would violate a security policy, such as an unencrypted Amazon Elastic Block Store (Amazon EBS) volume.

Which solution will meet these requirements?

- A. Turn on AWS Trusted Advisor. Configure security notifications as webhooks in the preferences section of the CI/CD pipeline.
- B. Turn on AWS Config. Use the prebuilt rules or customized rules. Subscribe the CI/CD pipeline to an Amazon Simple Notification Service (Amazon SNS) topic that receives notifications from AWS Config.
- C. Create rule sets in AWS CloudFormation Guard. Run validation checks for CloudFormation templates as a phase of the CI/CD process.
- D. Create rule sets as SCPs. Integrate the SCPs as a part of validation control in a phase of the CI/CD process.

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **RookieCSA** 4 months, 3 weeks ago

Selected Answer: C

It doesn't mention organizations, so SCP isn't needed. If organizations were mentioned, then D would also be correct.

<https://aws.amazon.com/tw/blogs/security/governance-at-scale-enforce-permissions-and-compliance-by-using-policy-as-code/>
upvoted 3 times

🗳️ 👤 **sendjawemail** 6 months ago

Selected Answer: D

If you see the following link, it talks about implementing SCP for preventing controls.

<https://aws.amazon.com/blogs/security/governance-at-scale-enforce-permissions-and-compliance-by-using-policy-as-code/>
For proactive controls Cloudformation guard is used.
upvoted 1 times

🗳️ 👤 **phmeeee** 2 months, 3 weeks ago

SCPs are not designed to be run in a CI/CD pipeline for template validation and it working at organization-level.
upvoted 2 times

🗳️ 👤 **Gzuiz** 9 months, 1 week ago

I think that is D
upvoted 1 times

🗳️ 👤 **Daniel76** 1 year ago

Selected Answer: C

<https://aws.amazon.com/blogs/mt/policy-as-code-for-securing-aws-and-third-party-resource-types/>
upvoted 2 times

🗳️ 👤 **Aamee** 1 year, 1 month ago

Selected Answer: C

C definitely. CFN Guard defined rule sets help in preventing from the derivation of Infrastructure resource security policies..
upvoted 4 times

🗳️ 👤 **[Removed]** 1 year, 1 month ago

Selected Answer: C

C for sure
upvoted 2 times

🗳️ 👤 **oioi** 1 year, 1 month ago

Selected Answer: C

correct

upvoted 1 times

A company is running an Amazon RDS for MySQL DB instance in a VPC. The VPC must not send or receive network traffic through the internet.

A security engineer wants to use AWS Secrets Manager to rotate the DB instance credentials automatically. Because of a security policy, the security engineer cannot use the standard AWS Lambda function that Secrets Manager provides to rotate the credentials.

The security engineer deploys a custom Lambda function in the VPC. The custom Lambda function will be responsible for rotating the secret in Secrets Manager. The security engineer edits the DB instance's security group to allow connections from this function. When the function is invoked, the function cannot communicate with Secrets Manager to rotate the secret properly.

What should the security engineer do so that the function can rotate the secret?

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

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **c6ed25a** 3 months, 1 week ago

Selected Answer: D

Since no internet we can only think of VPC Endpoint
upvoted 1 times

🗳️ 👤 **navid1365** 7 months, 1 week ago

Selected Answer: D

Since the RDS instance is in a private subnet without internet access, you need to configure a VPC endpoint.
upvoted 2 times

🗳️ 👤 **Daniel76** 1 year ago

Selected Answer: D

<https://docs.aws.amazon.com/secretsmanager/latest/userguide/vpc-endpoint-overview.html>
upvoted 1 times

🗳️ 👤 **Aamee** 1 year, 1 month ago

Selected Answer: D

D looks legit.
upvoted 3 times

🗳️ 👤 **[Removed]** 1 year, 1 month ago

Selected Answer: D

D is the winner
upvoted 1 times

🗳️ 👤 **oioi** 1 year, 1 month ago

Selected Answer: D

correct
upvoted 1 times

The security engineer is managing a traditional three-tier web application that is running on Amazon EC2 instances. The application has become the target of increasing numbers of malicious attacks from the internet.

What steps should the security engineer take to check for known vulnerabilities and limit the attack surface? (Choose two.)

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

Suggested Answer: BD

Community vote distribution

BD (100%)

🗳️ 👤 **AgboolaKun** Highly Voted 7 months ago

Selected Answer: BD

Security groups for reducing the attack surface,
Amazon Inspector to scan for and mitigate known vulnerabilities
upvoted 9 times

🗳️ 👤 **c6ed25a** Most Recent 3 months, 1 week ago

Selected Answer: BD

Security group and inspector is good answer
upvoted 1 times

🗳️ 👤 **Aamee** 7 months ago

Selected Answer: BD

B and D , self-explanatory..
upvoted 1 times

🗳️ 👤 **[Removed]** 7 months, 1 week ago

Selected Answer: BD

B D. Moderator, please correct the default answers
upvoted 1 times

🗳️ 👤 **oioi** 7 months, 1 week ago

Selected Answer: BD

correct
upvoted 1 times

A company is using Amazon Elastic Container Service (Amazon ECS) to run its container-based application on AWS. The company needs to ensure that the container images contain no severe vulnerabilities. The company also must ensure that only specific IAM roles and specific AWS accounts can access the container images.

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

- A. Pull images from the public container registry. Publish the images to Amazon Elastic Container Registry (Amazon ECR) repositories with scan on push configured in a centralized AWS account. Use a CI/CD pipeline to deploy the images to different AWS accounts. Use identity-based policies to restrict access to which IAM principals can access the images.
- B. Pull images from the public container registry. Publish the images to a private container registry that is hosted on Amazon EC2 instances in a centralized AWS account. Deploy host-based container scanning tools to EC2 instances that run Amazon ECS. Restrict access to the container images by using basic authentication over HTTPS.
- C. Pull images from the public container registry. Publish the images to Amazon Elastic Container Registry (Amazon ECR) repositories with scan on push configured in a centralized AWS account. Use a CI/CD pipeline to deploy the images to different AWS accounts. Use repository policies and identity-based policies to restrict access to which IAM principals and accounts can access the images.
- D. Pull images from the public container registry. Publish the images to AWS CodeArtifact repositories in a centralized AWS account. Use a CI/CD pipeline to deploy the images to different AWS accounts. Use repository policies and identity-based policies to restrict access to which IAM principals and accounts can access the images.

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **phmeeee** 2 months, 3 weeks ago

Selected Answer: C

1. ECR can scan on push.
 2. ECR can have resource-based policy to allow the specific IAM/principal
 3. LEAST management overhead.
- upvoted 1 times

🗳️ 👤 **ion_gee** 8 months, 1 week ago

Selected Answer: C

Answer is C. A just falls short due to absence of a required repository policy in addition to an IAM policy

upvoted 1 times

🗳️ 👤 **Gafa255** 11 months, 1 week ago

Selected Answer: C

<https://docs.aws.amazon.com/AmazonECR/latest/userguide/repository-policies.html>

upvoted 2 times

🗳️ 👤 **Daniel76** 1 year ago

Selected Answer: C

<https://docs.aws.amazon.com/AmazonECR/latest/userguide/repository-policies.html>

upvoted 1 times

🗳️ 👤 **AgboolaKun** 1 year, 1 month ago

Selected Answer: C

C is the correct answer. Please refer to <https://aws.amazon.com/premiumsupport/knowledge-center/secondary-account-access-ecr/>

upvoted 2 times

🗳️ 👤 **Aamee** 1 year, 1 month ago

Selected Answer: C

Should be C as it logically answers for not only the question of providing a solution of vulnerable free container image process but also covers the method of its access restrictions via IAM roles/principals and accounts as well.

upvoted 1 times

🗨️ 👤 **[Removed]** 1 year, 1 month ago

Selected Answer: C

I like C. More hardened than A?

upvoted 1 times

🗨️ 👤 **[Removed]** 1 year, 1 month ago

A also doesn't mention account restrictions so C for sure

upvoted 2 times

🗨️ 👤 **oioi** 1 year, 1 month ago

Selected Answer: C

correct

upvoted 1 times

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

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

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

- A. Configure an S3 Lifecycle rule on the S3 bucket to delete objects after 45 days.
- B. Create an AWS Lambda function to check the last-modified date of the S3 objects and delete objects that are older than 45 days. Create an S3 event notification to invoke the Lambda function for each PutObject operation.
- C. Create an AWS Lambda function to check the last-modified date of the S3 objects and delete objects that are older than 45 days. Create an Amazon EventBridge rule to invoke the Lambda function each month.
- D. Configure S3 Intelligent-Tiering on the S3 bucket to automatically transition objects to another storage class.

Suggested Answer: A

Community vote distribution

A (100%)

🗳️ 👤 **phmeeee** 2 months, 3 weeks ago

Selected Answer: A

You can set the TTL on S3 Lifecycle rule to remove the object by this criteria.

upvoted 1 times

🗳️ 👤 **navid1365** 7 months, 1 week ago

Selected Answer: A

A is the answer without a doubt

upvoted 2 times

🗳️ 👤 **rahav** 1 year ago

Selected Answer: A

A is correct. need to be removed

upvoted 1 times

🗳️ 👤 **ykhan321** 1 year ago

Selected Answer: A

correct

upvoted 1 times

🗳️ 👤 **Aamee** 1 year, 1 month ago

Selected Answer: A

Option A is sufficient to complete the required asks.

upvoted 3 times

🗳️ 👤 **[Removed]** 1 year, 1 month ago

Selected Answer: A

A. The rest are just pure comedy

upvoted 2 times

🗳️ 👤 **AgboolaKun** 1 year, 1 month ago

Laughing!!! I agree, they are just pure comedy.

upvoted 1 times

🗳️ 👤 **oioi** 1 year, 1 month ago

Selected Answer: A

correct

upvoted 1 times

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. Change the Principal element to the following:

```
{
  "AWS": "arn:aws:lambda:::function:MyLambdaFunction"
}
```

B. Change the Action element to the following:

```
[
  "s3:GetObject*",
  "s3:GetBucket*"
]
```

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

D. Change the Resource element to "arn:aws:lambda:::function:MyLambdaFunction". Change the Principal element to the following:

```
{
  "Service": "s3.amazonaws.com"
}
```

Suggested Answer: C

Community vote distribution

C (100%)

Aamee Highly Voted 1 year, 1 month ago

Selected Answer: C

Yup, option C is absolutely correct since there's nothing wrong on any other areas except the missing ' * ' under the Resource element.
upvoted 8 times

phmeeeee Most Recent 2 months, 3 weeks ago

Selected Answer: C

The object is stay inside S3 bucket so "DOC-EXAMPLE- BUCKET/*" is valid.
upvoted 1 times

Gafa255 11 months, 1 week ago

Selected Answer: C

missing the *
upvoted 1 times

[Removed] 1 year, 1 month ago

Selected Answer: C

C. Missing the *
upvoted 2 times

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.

Suggested Answer: D

Community vote distribution

D (100%)

 **azure4life** Highly Voted 1 year ago

Selected Answer: D

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

Since the S3 bucket is using SSE-KMS encryption with a customer managed key, the key policy for that KMS key needs to grant the appropriate permissions to allow decryption of the objects. The IAM policy grants the kms:Decrypt permission, but if the key policy no longer gives the AWS account full access, the decrypt permission will still be denied.

Options A and B relate to the kms:DescribeKey permission and AWS managed keys, but a customer managed key is being used here. Option C is incorrect because an S3 bucket policy is not required when using IAM policies for permissions. Therefore, option D that mentions the KMS key policy having inappropriate access for the account is the likely reason for the access being denied.

upvoted 5 times

 **vikasj1in** Most Recent 11 months, 2 weeks ago

The IAM user has been granted the kms:Decrypt permission for the customer managed key used for server-side encryption in the S3 bucket. If the KMS key policy has been modified to restrict access, it might override the IAM user's permissions, resulting in an Access Denied error.

It's crucial to ensure that the KMS key policy grants the necessary permissions to the AWS account (and by extension, the IAM user) to perform the required decryption operations.

upvoted 2 times

 **Daniel76** 1 year ago

Selected Answer: D

If you allow by IAM policy to a key, it still can be denied by key policy (which is another policy) unless you explicitly allows.

"Unless the key policy explicitly allows it, you cannot use IAM policies to allow access to a KMS key. Without permission from the key policy, IAM policies that allow permissions have no effect."

upvoted 4 times

 **Daniel76** 1 year ago

<https://docs.aws.amazon.com/kms/latest/developerguide/key-policies.html>

upvoted 2 times

 **kejam** 1 year ago

Selected Answer: D

<https://docs.aws.amazon.com/kms/latest/developerguide/key-policy-default.html#key-policy-default-allow-root-enable-iam>

upvoted 1 times

 **Aamee** 1 year, 1 month ago

Selected Answer: D

The following statement leads me to believe that option D could be the best option:

'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'.

upvoted 1 times

🗲️ 👤 **[Removed]** 1 year, 1 month ago

Selected Answer: D

D is correct

upvoted 2 times

🗲️ 👤 **oioi** 1 year, 1 month ago

Selected Answer: D

correct

upvoted 1 times

A company has a guideline that mandates the encryption of all Amazon S3 bucket data in transit. A security engineer must implement an S3 bucket policy that denies any S3 operations if data is not encrypted.

Which S3 bucket policy will meet this requirement?

- A.
- ```
{
 "Version": "2012-10-17",
 "Statement": [{
 "Sid": "AllowSSLRequestsOnly",
 "Action": "s3:*",
 "Effect": "Deny",
 "Resource": [
 "arn:aws:s3::DOC-EXAMPLE-BUCKET",
 "arn:aws:s3::DOC-EXAMPLE-BUCKET/*"
],
 "Condition": {
 "Bool": {
 "aws:SecureTransport": "true"
 }
 },
 "Principal": "*"
 }]
}
```
- B.
- ```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowSSLRequestsOnly",
    "Action": "s3:*",
    "Effect": "Deny",
    "Resource": [
      "arn:aws:s3::DOC-EXAMPLE-BUCKET",
      "arn:aws:s3::DOC-EXAMPLE-BUCKET/*"
    ],
    "Condition": {
      "Bool": {
        "aws:SecureTransport": "false"
      }
    },
    "Principal": "*"
  }]
}
```
- C.
- ```
{
 "Version": "2012-10-17",
 "Statement": [{
 "Sid": "AllowSSLRequestsOnly",
 "Action": "s3:*",
 "Effect": "Deny",
 "Resource": [
 "arn:aws:s3::DOC-EXAMPLE-BUCKET",
 "arn:aws:s3::DOC-EXAMPLE-BUCKET/*"
],
 "Condition": {
 "StringNotEquals": {
 "s3:x-amz-server-side-encryption": "AES256"
 }
 },
 "Principal": "*"
 }]
}
```
- D.
- ```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Sid": "AllowSSLRequestsOnly",
    "Action": "s3:*",
    "Effect": "Deny",
    "Resource": [
      "arn:aws:s3::DOC-EXAMPLE-BUCKET",
      "arn:aws:s3::DOC-EXAMPLE-BUCKET/*"
    ],
    "Condition": {
      "StringNotEquals": {
        "s3:x-amz-server-side-encryption": true
      }
    },
    "Principal": "*"
  }]
}
```

awssecuritynewbie 10 months, 1 week ago

Selected Answer: B

yeah is not good question, it should of been allow if thestring not equal : encrypted"
upvoted 1 times

awssecuritynewbie 10 months, 2 weeks ago

Selected Answer: B

It is defo B, you are denying access if the condition is set not to encrypt in trasnit
upvoted 1 times

cannoe 11 months, 2 weeks ago

The question is a bit tricky for me.

The requirement is: "A security engineer must implement an S3 bucket policy that denies any S3 operations if data is not encrypted."

Using HTTPS as a connection does not encrypt the data, it encrypts the connection. When using HTTPS to access an Amazon S3 bucket, the HTTPS encryption is de-encapsulated at the S3 service endpoint. This means the data transmitted between your application and the S3 endpoint is encrypted in transit using HTTPS, but once it reaches the S3 endpoint, the encryption is removed before the data is stored in S3.

upvoted 1 times

Aamee 1 year ago

Selected Answer: B

Updated selection: Def. B
upvoted 2 times

kejam 1 year ago

Selected Answer: B

Enforce encryption of data in transit
<https://docs.aws.amazon.com/AmazonS3/latest/userguide/security-best-practices.html#security-best-practices-prevent>
upvoted 1 times

AgboolaKun 1 year, 1 month ago

Selected Answer: B

This question is requesting to make objects accessible only through HTTPS. Option B is correct because it specifies the bucket policy condition with correct syntax.

Please refer to "Defense-in-depth requirement 2: Data must be accessible only by a limited set of public IP addresses" section in this link -
<https://aws.amazon.com/blogs/security/how-to-use-bucket-policies-and-apply-defense-in-depth-to-help-secure-your-amazon-s3-data/>

upvoted 4 times

Aamee 1 year ago

Thnx so much... got the concept cleared now so def. going with B here now..
upvoted 1 times

Aamee 1 year, 1 month ago

Selected Answer: C

Not sure if I'm fully correct here in selecting this ans. I'd go with C here cuz I feel like it is asked about no S3 bucket operation IF the data is not encrypted. It doesn't say about if the data is not securely in transit. That's why in my opinion, the AES256 encryption method should be mentioned under the conditional logic area in the bucket policy.

But I'd appreciate if anyone else would like to discuss and clarify my understandings on this if I'm incorrect here... Thnx so much!

upvoted 1 times

[Removed] 1 year, 1 month ago

Selected Answer: B

B. You want to deny where secure transport is false
upvoted 2 times

Wije1 1 year, 1 month ago

<https://repost.aws/knowledge-center/s3-bucket-policy-for-config-rule>
upvoted 2 times


A security engineer wants to use Amazon Simple Notification Service (Amazon SNS) to send email alerts to a company's security team for Amazon GuardDuty findings that have a High severity level. The security engineer also wants to deliver these findings to a visualization tool for further examination.

Which solution will meet these requirements?

- A. Set up GuardDuty to send notifications to an Amazon CloudWatch alarm with two targets in CloudWatch. From CloudWatch, stream the findings through Amazon Kinesis Data Streams into an Amazon Open Search Service domain as the first target for delivery. Use Amazon QuickSight to visualize the findings. Use OpenSearch queries for further analysis. Deliver email alerts to the security team by configuring an SNS topic as a second target for the CloudWatch alarm. Use event pattern matching with an Amazon EventBridge event rule to send only High severity findings in the alerts.
- B. Set up GuardDuty to send notifications to AWS CloudTrail with two targets in CloudTrail. From CloudTrail, stream the findings through Amazon Kinesis Data Firehose into an Amazon OpenSearch Service domain as the first target for delivery. Use OpenSearch Dashboards to visualize the findings. Use OpenSearch queries for further analysis. Deliver email alerts to the security team by configuring an SNS topic as a second target for CloudTrail. Use event pattern matching with a CloudTrail event rule to send only High severity findings in the alerts.
- C. Set up GuardDuty to send notifications to Amazon EventBridge with two targets. From EventBridge, stream the findings through Amazon Kinesis Data Firehose into an Amazon OpenSearch Service domain as the first target for delivery. Use OpenSearch Dashboards to visualize the findings. Use OpenSearch queries for further analysis. Deliver email alerts to the security team by configuring an SNS topic as a second target for EventBridge. Use event pattern matching with an EventBridge event rule to send only High severity findings in the alerts.
- D. Set up GuardDuty to send notifications to Amazon EventBridge with two targets. From EventBridge, stream the findings through Amazon Kinesis Data Streams into an Amazon OpenSearch Service domain as the first target for delivery. Use Amazon QuickSight to visualize the findings. Use OpenSearch queries for further analysis. Deliver email alerts to the security team by configuring an SNS topic as a second target for EventBridge. Use event pattern matching with an EventBridge event rule to send only High severity findings in the alerts.

Suggested Answer: C

Community vote distribution



mynickc 1 year, 5 months ago

Selected Answer: C

D is not right because you need to perform a few activities to establish connection b/w Opensearch and QuickSight which is missing in choice D
upvoted 1 times

vikasj1in 1 year, 5 months ago

Selected Answer: C

This option involves using EventBridge to handle the GuardDuty findings and then routing them to two targets:

Streaming the findings through Kinesis Data Firehose into an Amazon OpenSearch Service domain for visualization.

Sending email alerts to the security team via SNS, with event pattern matching to filter only high severity findings.

This approach leverages the flexibility of EventBridge to manage the workflow and routing of events to different services based on specific criteria.

upvoted 1 times

happy34 1 year, 5 months ago

C - it filters only High alerts.

upvoted 1 times

WeepingMaple 1 year, 6 months ago

<https://github.com/aws-samples/siem-on-amazon-opensearch-service>

upvoted 1 times

dexterryu 1 year, 6 months ago

D is correct due to Quicksight which is AWS's preferred visualization tool. The blog linked below is dated (from 2018), and while still valid not the best way to do visualization in AWS.

You can still use the OS query in Quicksight.

<https://docs.aws.amazon.com/quicksight/latest/user/connecting-to-os.html>

upvoted 1 times

  **happy34** 1 year, 5 months ago

Quick sight works with Cloudwatch Open Seach has its own dashboard. Not sure if Quicksight works with OpenSearch

upvoted 1 times



  **Daniel76** 1 year, 6 months ago

Selected Answer: C

According to this AWS article, it is GuardDuty -> EventBrige -> Firehouse -> OpenSearch -> OpenSearch visualization.

<https://aws.amazon.com/blogs/security/visualizing-amazon-guardduty-findings/>

upvoted 4 times

  **tayman** 1 year, 6 months ago

Selected Answer: C

Vote for C



upvoted 1 times

  **ykhan321** 1 year, 6 months ago

Selected Answer: D

QuickSight is the hint for Visualization.

upvoted 2 times

  **hb0011** 9 months, 4 weeks ago

Opensearch Dashboards is better for this use case. The Answer is C

upvoted 1 times

  **azure4life** 1 year, 6 months ago

Selected Answer: D

Option D is the correct solution.


GuardDuty can send findings to Amazon EventBridge. EventBridge can then stream to targets like Kinesis Data Streams to process and store the findings, and SNS to send email alerts. Using EventBridge event pattern matching allows filtering findings based on properties like severity. Kinesis Data Streams can feed findings into OpenSearch Service. OpenSearch Dashboards or Amazon QuickSight can visualize the findings, while OpenSearch queries can provide analysis.

Option A is incorrect because GuardDuty integrates with EventBridge, not CloudWatch Alarms.

Option B is incorrect because GuardDuty integrates with EventBridge, not CloudTrail.

Option C is incorrect because Kinesis Data Firehose would not allow querying and analysis of findings - Kinesis Data Streams enables this with OpenSearch.

upvoted 1 times

  **Th3Dud3** 1 year, 5 months ago

(C) - Kinesis Data *Streams* can't send to OpenSearch...



upvoted 3 times

  **Egle** 1 year, 6 months ago

Selected Answer: C

correct

upvoted 1 times

  **kejam** 1 year, 7 months ago

Selected Answer: C

https://docs.aws.amazon.com/guardduty/latest/ug/guardduty_settingup.html#setup-sns

<https://aws.amazon.com/blogs/big-data/audit-aws-service-events-with-amazon-eventbridge-and-amazon-kinesis-data-firehose/>

<https://aws.amazon.com/blogs/big-data/ingest-streaming-data-into-amazon-elasticsearch-service-within-the-privacy-of-your-vpc-with-amazon-kinesis-data-firehose/>



upvoted 4 times

  **[Removed]** 1 year, 7 months ago

Selected Answer: C

C gets the job done

upvoted 1 times

  **oioi** 1 year, 7 months ago

Selected Answer: C

correct

upvoted 1 times

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. Place objects in the S3 buckets.
- B. Use S3 Glacier Vault Lock to attach a Vault Lock policy to new S3 buckets. Wait 24 hours to complete the Vault Lock process. Place objects in the S3 buckets.
- C. Create new S3 buckets with S3 Object Lock enabled in governance mode. Place objects in the S3 buckets.
- D. Create new S3 buckets with S3 Object Lock enabled in governance mode. Add a legal hold to the S3 buckets. Place objects in the S3 buckets.

Suggested Answer: A

Community vote distribution

A (100%)

 **confusedyeti69** Highly Voted 1 year ago

Selected Answer: A

In compliance mode, a protected object version can't be overwritten or deleted by any user, including the root user in your AWS account. When an object is locked in compliance mode, its retention mode can't be changed, and its retention period can't be shortened. Compliance mode helps ensure that an object version can't be overwritten or deleted for the duration of the retention period.

In governance mode, users can't overwrite or delete an object version or alter its lock settings unless they have special permissions. With governance mode, you protect objects against being deleted by most users, but you can still grant some users permission to alter the retention settings or delete the objects if necessary. You can also use governance mode to test retention-period settings before creating a compliance-mode retention period.

upvoted 7 times

 **Jamshif01** Most Recent 11 months, 1 week ago

A

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock-configure.html>

upvoted 1 times

 **kejam** 1 year ago

Selected Answer: A

<https://aws.amazon.com/s3/features/object-lock/>

upvoted 2 times

 **Aamee** 1 year, 1 month ago

Selected Answer: A

Option A would work in this usecase.

upvoted 1 times

 **[Removed]** 1 year, 1 month ago

Selected Answer: A

A rings the most bells

upvoted 1 times

 **oioi** 1 year, 1 month ago

Selected Answer: A

coorect

upvoted 1 times

A company needs complete encryption of the traffic between external users and an application. The company hosts the application on a fleet of Amazon EC2 instances that run in an Auto Scaling group behind an Application Load Balancer (ALB).

How can a security engineer meet these requirements?

- A. Create a new Amazon-issued certificate in AWS Secrets Manager. Export the certificate from Secrets Manager. Import the certificate into the ALB and the EC2 instances.
- B. Create a new Amazon-issued certificate in AWS Certificate Manager (ACM). Associate the certificate with the ALB. Export the certificate from ACM. Install the certificate on the EC2 instances.
- C. Import a new third-party certificate into AWS Identity and Access Management (IAM). Export the certificate from IAM. Associate the certificate with the ALB and the EC2 instances.
- D. Import a new third-party certificate into AWS Certificate Manager (ACM). Associate the certificate with the ALB. Install the certificate on the EC2 instances.

Suggested Answer: D

Community vote distribution

D (88%)

12%

 **AgboolaKun** Highly Voted 1 year, 1 month ago

Selected Answer: D

Because of the wording of this question, I did not first know which of the options B and D is correct.

However, my conviction that you can't directly install Amazon-issued certificates on EC2 instances (refer to - <https://repost.aws/knowledge-center/associate-acm-certificate-alb-nlb> for more information) made me to study a few documentations to be sure D is the correct answer.

Please check the Accepted answer in the following thread - <https://repost.aws/questions/QUIo7PWvZ3T6aFYCByhZ5f0A/load-certificate-on-alb-and-ec2>

upvoted 9 times

 **ginseng** Most Recent 9 months, 2 weeks ago

Selected Answer: D

Q: Can I use certificates on Amazon EC2 instances or on my own servers?

You can use private certificates issued with Private CA with EC2 instances, containers, and on your own servers. At this time, public ACM certificates can be used only with specific AWS services, including AWS Nitro Enclaves. See ACM service integrations.

upvoted 4 times

 **vikasj1in** 11 months, 2 weeks ago

B.

To achieve complete encryption of the traffic between external users and an application hosted on Amazon EC2 instances behind an Application Load Balancer (ALB), you would typically use SSL/TLS encryption. AWS Certificate Manager (ACM) provides a managed service for provisioning and renewing SSL/TLS certificates.

Here's how the process works:

Create a new Amazon-issued certificate in ACM.

Associate the certificate with the ALB. This ensures that the ALB can terminate SSL/TLS connections on behalf of the EC2 instances.

Export the certificate from ACM.

Install the exported certificate on the EC2 instances. This ensures that the communication between the ALB and EC2 instances is also encrypted.

By using ACM, you benefit from the managed certificate service, automated certificate renewal, and easy integration with other AWS services like ALB. This approach ensures secure communication from external users to the ALB and between the ALB and EC2 instances.

upvoted 1 times

 **Daniel76** 1 year ago

Selected Answer: D

ACM should be used, so A and C are out.

Between B and D, B is out because Amazon-issued public cert cannot be installed on EC2 instances.

<https://repost.aws/knowledge-center/associate-acm-certificate-alb-nlb>

upvoted 1 times

🗨️ **azure4life** 1 year ago

Selected Answer: D

Option D is the correct solution.

To encrypt traffic between external users and the application behind the Application Load Balancer (ALB), a certificate should be imported into AWS Certificate Manager (ACM) and associated with the ALB. The same certificate should also be installed on the EC2 instances.

Option A is incorrect because Secrets Manager is used for storing secrets, not SSL/TLS certificates.

Option B is incorrect because Amazon-issued ACM certificates can only be used with Elastic Load Balancers and Amazon CloudFront. They cannot be exported and installed on EC2 instances.

Option C is incorrect because IAM does not support importing or managing SSL/TLS certificates.

Option D uses a third-party certificate imported into ACM, associated with the ALB, and installed on the EC2 instances. This provides complete encryption between the users and application.

upvoted 3 times

🗨️ **Oralinux** 1 year ago

Bad question; I think it should be B since AWS always tries to promote and use internal services and not go to third parties. We deploy SSL in ALB terminate and send non-SSL to EC2. In my opinion, the provided answers are incorrect.

upvoted 2 times

🗨️ **snowmaggodon** 1 year, 1 month ago

CAN'T use an ACM cert on ec2 instance. D is the right answer.

upvoted 1 times

🗨️ **Aamee** 1 year, 1 month ago

Selected Answer: B

I think it's asking about the key difference btw creating Amazon based Cert versus creating/using 3rd party Certs.... specially on leveraging the feature of 'exporting the Cert' from ACM which looks valid in option B only whereas on other choices, it's not a good fit. I could be wrong but that's what makes me feel to go with option B here..

upvoted 1 times

🗨️ **[Removed]** 1 year, 1 month ago

Kind of a bad question here, so I can't really make a proper decision between B and D.. What certificate is being applied? ALB does not pass encrypted traffic to a target. NLB will do that.

you must deploy at least one SSL server certificate on your load balancer. The load balancer uses a server certificate to terminate the front-end connection and then decrypt requests from clients before sending them to the targets. You must also specify a security policy, which is used to negotiate secure connections between clients and the load balancer.

If you need to pass encrypted traffic to targets without the load balancer decrypting it, you can create a Network Load Balancer or Classic Load Balancer with a TCP listener on port 443. With a TCP listener, the load balancer passes encrypted traffic through to the targets without decrypting it.

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

upvoted 2 times

🗨️ **oioi** 1 year, 1 month ago

Selected Answer: B

correct

upvoted 1 times

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

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

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

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

- A. Add a resource policy that allows each member of the group to access Amazon SES.
- B. Add a resource policy that allows "Principal": {"AWS": "arn:aws:iam::account-number:group/Dev"}.
- C. Remove the AWS Control Tower control (guardrail) that restricts access to Amazon SES.
- D. Remove Amazon SES from the root SCP.

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **navid1365** 7 months, 1 week ago

Selected Answer: D

D is correct. A SCP identifies the maximum level of access that IAM entity within that OU can have. Since SES is denied in the SCP, it does not matter if you allow it in other policies. It will simply not be allowed because the SCP does not allow it.

upvoted 2 times

🗳️ 👤 **awssecuritynewbie** 10 months, 2 weeks ago

Selected Answer: D

pay attention to the question... it states it wants to allow the dev to use SES... it does not ask to "only" allow them. so it would make sure to remove the SCP because SCP is never overwritten.

upvoted 2 times

🗳️ 👤 **rahav** 1 year ago

D is the correct one

upvoted 1 times

🗳️ 👤 **ykhan321** 1 year ago

Selected Answer: D

Why most of the answers are incorrect here.

upvoted 2 times

🗳️ 👤 **azure4life** 1 year ago

Selected Answer: D

Option D is the correct solution. The root SCP is denying access to Amazon SES across the organization. Even though the OU SCP and IAM policy allow SES access, the root SCP takes precedence and blocks it. Removing Amazon SES from the root SCP whitelist will resolve the issue and allow

the developers to access SES based on the permissions granted in their IAM policy.

Option A is incorrect because resource policies apply at the service level, not for IAM users/groups.

Option B is also related to resource policies, not the issue with the SCP whitelist.

Option C mentions AWS Control Tower which is not referenced in the question. The SCP is set through AWS Organizations.

So the root cause is the root SCP denying access to SES, and it needs to be removed from that SCP to allow access that is permitted in the lower levels of permissions.

upvoted 2 times

  **Oralinux** 1 year ago

Answer D:

a resource policy attached directly to an AWS resource (such as Amazon SES) cannot override an SCP (Service Control Policy) set at the root level in AWS Organizations.

Service Control Policies (SCPs) at the root level act as "guardrails" and define the maximum permissions that accounts within the organization can have. They are evaluated before resource policies.

If an SCP denies access to a particular service, even a resource policy allowing access on the specific resource won't take effect. The SCP at the root level will override any resource policy attached to individual resources.

So, while a resource policy can be useful for granting permissions on a specific resource, it cannot be used to override the restrictions imposed by an SCP at a higher level in the organization's hierarchy. In this scenario, removing the restriction for Amazon SES from the root SCP would be the effective solution.



upvoted 1 times

  **Aamee** 1 year ago

Selected Answer: D

Leads me towards option D only cuz it seems like the denial of ses* actions explicitly defined under the SCP is probably blocking their authorization requests... not sure if Control Tower here makes any big difference..

upvoted 1 times

  **[Removed]** 1 year, 1 month ago

Selected Answer: D

The answer is D

upvoted 3 times

A company hosts a public website on an Amazon EC2 instance. HTTPS traffic must be able to access the website. The company uses SSH for management of the web server.

The website is on the subnet 10.0.1.0/24. The management subnet is 192.168.100.0/24. A security engineer must create a security group for the EC2 instance.

Which combination of steps should the security engineer take to meet these requirements in the MOST secure manner? (Choose two.)

- A. Allow port 22 from source 0.0.0.0/0.
- B. Allow port 443 from source 0.0.0.0/0.
- C. Allow port 22 from 192.168.100.0/24.
- D. Allow port 22 from 10.0.1.0/24.
- E. Allow port 443 from 10.0.1.0/24.

Suggested Answer: BC

Community vote distribution

BC (100%)

🗳️ 👤 **rahav** 1 year ago

Selected Answer: BC

BC is the correct
upvoted 1 times

🗳️ 👤 **tayman** 1 year ago

Selected Answer: BC

BC 100%
upvoted 1 times

🗳️ 👤 **ykhan321** 1 year ago

Selected Answer: BC

Another wrong answer here. My head is spinning.
upvoted 1 times

🗳️ 👤 **Oralinux** 1 year ago

Answer: BC
upvoted 1 times

🗳️ 👤 **Aamee** 1 year, 1 month ago

Selected Answer: BC

Yup, agreed with B and C.
upvoted 4 times

🗳️ 👤 **[Removed]** 1 year, 1 month ago

Selected Answer: BC

Quite sure the subnet "10.0.1.0/24" is a distractor, so B and C are my vote.
upvoted 4 times

🗳️ 👤 **awssecuritynewbie** 10 months, 2 weeks ago

yup it is a twister lol you are allowing the world to access your website so it has to be 0.0.0.0/0
upvoted 1 times

🗳️ 👤 **AgboolaKun** 1 year, 1 month ago

Agree. I almost fell for it until I read the question again and paid attention to "HTTPS traffic must be able to access the website". This means any (0.0.0.0/0) https traffic.
upvoted 2 times

🗳️ 👤 **oioi** 1 year, 1 month ago

Selected Answer: BC

correct

upvoted 2 times

A security engineer wants to forward custom application-security logs from an Amazon EC2 instance to Amazon CloudWatch. The security engineer installs the CloudWatch agent on the EC2 instance and adds the path of the logs to the CloudWatch configuration file.

However, CloudWatch does not receive the logs. The security engineer verifies that the awslogs service is running on the EC2 instance.

What should the security engineer do next to resolve the issue?

- A. Add AWS CloudTrail to the trust policy of the EC2 instance. Send the custom logs to CloudTrail instead of CloudWatch.
- B. Add Amazon S3 to the trust policy of the EC2 instance. Configure the application to write the custom logs to an S3 bucket that CloudWatch can use to ingest the logs.
- C. Add Amazon Inspector to the trust policy of the EC2 instance. Use Amazon Inspector instead of the CloudWatch agent to collect the custom logs.
- D. Attach the CloudWatchAgentServerPolicy AWS managed policy to the EC2 instance role.

Suggested Answer: D

Community vote distribution

D (100%)

🗳️ 👤 **Jamshif01** 11 months, 1 week ago

D

All other answers are irrelevant

upvoted 2 times

🗳️ 👤 **rahav** 1 year ago

Selected Answer: D

D for sure

upvoted 2 times

🗳️ 👤 **yorkicurke** 1 year ago

Selected Answer: D

Uses of CloudWatchAgentServerPolicy ;

It allows the CloudWatch agent to publish metrics and logs to CloudWatch on behalf of the IAM role or user the policy is attached to.

It provides permissions for the agent to access and manage its own configuration files stored in S3.

The policy grants permissions across multiple AWS services like CloudWatch, S3, KMS etc. to allow end-to-end functionality of the monitoring agent.

upvoted 3 times

🗳️ 👤 **ykhan321** 1 year ago

Selected Answer: D

Only EC2 & Cloudwatch are in questions here.

upvoted 1 times

🗳️ 👤 **Oralinux** 1 year ago

Answer: D

upvoted 1 times

🗳️ 👤 **Aamee** 1 year, 1 month ago

Selected Answer: D

No doubt about D.

upvoted 3 times

🗳️ 👤 **[Removed]** 1 year, 1 month ago

Selected Answer: D

D is correcto

upvoted 3 times

🗳️ 👤 **oioi** 1 year, 1 month ago

Selected Answer: D

correct

upvoted 2 times

A systems engineer is troubleshooting the connectivity of a test environment that includes a virtual security appliance deployed inline. In addition to using the virtual security appliance, the development team wants to use security groups and network ACLs to accomplish various security requirements in the environment.

What configuration is necessary to allow the virtual security appliance to route the traffic?

- A. Disable network ACLs.
- B. Configure the security appliance's elastic network interface for promiscuous mode.
- C. Disable the Network Source/Destination check on the security appliance's elastic network interface.
- D. Place the security appliance in the public subnet with the internet gateway.

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ 👤 **vikasj1in** 11 months, 2 weeks ago

C,

When you deploy a virtual security appliance inline in a subnet, you need to ensure that it can effectively route traffic between different subnets. The "Network Source/Destination check" is a feature in Amazon EC2 that controls whether source/destination checking is enabled or disabled on a network interface.

In this context, the virtual security appliance acts as a router, and the "Network Source/Destination check" should be disabled on its elastic network interface. When this check is disabled, the network interface can handle traffic that is not specifically destined for the instance it is attached to, allowing it to route traffic between different subnets.

upvoted 3 times

🗳️ 👤 **rahav** 1 year ago

Selected Answer: C

C for sure

upvoted 1 times

🗳️ 👤 **Daniel76** 1 year ago

Selected Answer: C

Source/destination checking

You can enable or disable source/destination checks, which ensure that the instance is either the source or the destination of any traffic that it receives. Source/destination checks are enabled by default. You must disable source/destination checks if the instance runs services such as network address translation, routing, or firewalls.

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html>

upvoted 3 times

🗳️ 👤 **azure4life** 1 year ago

Selected Answer: C

Option C is the correct solution.

To allow a virtual security appliance deployed inline to route traffic between subnets, the Network Source/Destination Check needs to be disabled on its elastic network interface. This enables the appliance to receive traffic that is not specifically addressed to itself.

Option A is incorrect because disabling network ACLs is not required for a virtual appliance deployment and would reduce security.

Option B mentions promiscuous mode which applies to physical network interfaces, not virtual ones in AWS.

Option D places the appliance in the public subnet which may help route internet traffic but does not address routing between private subnets. Disabling the Source/Destination Check is required to enable that routing functionality.

upvoted 4 times

🗨️ 👤 **Oralinux** 1 year ago

Answer: C

upvoted 1 times

🗨️ 👤 **kejam** 1 year ago

Selected Answer: C

<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-eni.html>

upvoted 3 times

🗨️ 👤 **[Removed]** 1 year, 1 month ago

Selected Answer: C

C is correct

upvoted 3 times

🗨️ 👤 **oioi** 1 year, 1 month ago

Selected Answer: C

correct

upvoted 2 times

A security engineer needs to create an Amazon S3 bucket policy to grant least privilege read access to IAM user accounts that are named User1, User2, and User3. These IAM user accounts are members of the AuthorizedPeople IAM group. The security engineer drafts the following S3 bucket policy:

```
{
  "Version": "2012-10-17",
  "Id": "AuthorizedPeoplePolicy",
  "Statement": [
    {
      "Sid": "Actions-Authorized-People",
      "Effect": "Allow",
      "Action": [
        "s3:GetObject"
      ],
      "Resource": "arn:aws:s3::authorized-people-bucket/*"
    }
  ]
}
```

When the security engineer tries to add the policy to the S3 bucket, the following error message appears: "Missing required field Principal."

The security engineer is adding a Principal element to the policy. The addition must provide read access to only User1, User2, and User3.

Which solution meets these requirements?

- A.

```
"Principal": {
  "AWS": [
    "arn:aws:iam::1234567890:user/User1",
    "arn:aws:iam::1234567890:user/User2",
    "arn:aws:iam::1234567890:user/User3"
  ]
}
```
- B.

```
"Principal": {
  "AWS": [
    "arn:aws:iam::1234567890:root"
  ]
}
```
- C.

```
"Principal": {
  "AWS": [
    "*"
  ]
}
```
- D.



```
"Principal": {
  "AWS": "arn:aws:iam::1234567890:group/AuthorizedPeople"
}
```

Suggested Answer: A

Community vote distribution

A (91%)

9%

 [Removed]  1 year, 1 month ago

Selected Answer: A

Agree with AgboolaKun. What a lovely question

You can specify any of the following principals in a policy:

AWS account and root user

IAM roles

Role sessions

IAM users

Federated user sessions

AWS services

All principals

You cannot identify a user group as a principal in a policy (such as a resource-based policy) because groups relate to permissions, not authentication, and principals are authenticated IAM entities.

https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_elements_principal.html#Principal_specifying

upvoted 12 times

  **fcbflo** Most Recent 3 months, 4 weeks ago

Selected Answer: A

Option B uses the root account, which would grant access to all users in the account (too broad)

Option C uses a wildcard (*), which would grant access to all AWS users (too broad)

Option D attempts to use an IAM group as a principal, but IAM groups cannot be used as principals in resource-based policies like S3 bucket policies

upvoted 1 times

  **navid1365** 7 months, 1 week ago

Selected Answer: A

A is correct. Please not that the question says "only to user 1, user 2, and user 3". It does not say all members of the group "authorizedpeople". There may be other users part of that group who should NOT have access.

upvoted 3 times

  **Raphaello** 10 months, 1 week ago

Selected Answer: D

In AWS IAM, principals are authenticated IAM entities.

IAM entities are only IAM users and roles.

You cannot used IAM group as a principal in an IAM (resource) policy.



upvoted 2 times

  **Raphaello** 10 months, 1 week ago

Obviously error in selected answer.

Answer A is the correct one that is matching the provided explanation.

upvoted 1 times

  **Ritarocks** 9 months, 4 weeks ago

So D still ?

upvoted 1 times

  **Oralinix** 1 year ago

Answer A: " IAM user accounts that are named User1, User2, and User3. These IAM user accounts are members of the AuthorizedPeople IAM group"

we do not want to give read access to other accounts that are part of the AuthorizedPeople IAM group. => then only A satisfy this criteria

upvoted 2 times

  **Aamee** 1 year ago

Selected Answer: A

All others are not the valid choices since the Principal needs to be selected only for User1, User2 and User3 'only' explicitly... plus, groups can't be identified as a Principal anyways..

upvoted 4 times

  **AgboolaKun** 1 year, 1 month ago

Selected Answer: A

The correct answer is A - <https://www.examttopics.com/discussions/amazon/view/60777-exam-aws-certified-security-specialty-topic-1-question-258/>

upvoted 4 times

A security engineer recently rotated all IAM access keys in an AWS account. The security engineer then configured AWS Config and enabled the following AWS Config managed rules: mfa-enabled-for-iam-console-access, iam-user-mfa-enabled, access-keys-rotated, and iam-user-unused-credentials-check.


The security engineer notices that all resources are displaying as noncompliant after the IAM GenerateCredentialReport API operation is invoked.

What could be the reason for the noncompliant status?

- A. The IAM credential report was generated within the past 4 hours.
- B. The security engineer does not have the GenerateCredentialReport permission.
- C. The security engineer does not have the GetCredentialReport permission.
- D. The AWS Config rules have a MaximumExecutionFrequency value of 24 hours.

Suggested Answer: A

Community vote distribution



AgboolaKun Highly Voted 1 year, 1 month ago

Selected Answer: A

The report was generated within the past 4 hours - <https://repost.aws/knowledge-center/config-credential-report>
upvoted 8 times

navid1365 Highly Voted 7 months, 1 week ago

Selected Answer: A

AWS Config rules such as mfa-enabled-for-iam-console-access, iam-user-mfa-enabled, access-keys-rotated, and iam-user-unused-credentials-check rely on data from the IAM credential report. The IAM credential report is updated automatically every four hours, and changes in IAM (such as rotating access keys) may not be reflected in the report immediately. If the IAM credential report was generated within the past 4 hours, AWS Config might not yet have the updated information, causing the resources to display as noncompliant.
upvoted 5 times

Raphaello Most Recent 10 months, 1 week ago

Selected Answer: A

A.

You can generate a credential report as often as once every four hours
https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_getting-report.html
upvoted 1 times

awssecuritynewbie 10 months, 2 weeks ago

what a shit question btw! He has enabled a bunch of AWS configs, and then they are showing non-compliant. this can be anything really but yeah.
upvoted 2 times

vikasj1in 11 months, 2 weeks ago

Selected Answer: A

When these AWS Config rules are triggered, they rely on the latest IAM credential report to evaluate compliance. If the IAM credential report has been generated within the past 4 hours, it might not reflect the most recent changes, such as the rotation of access keys.
To address this, it's a good practice to ensure that the IAM credential report is generated and updated at regular intervals, and AWS Config rules are then evaluated against the most recent report. You can schedule the generation of the IAM credential report and the evaluation of AWS Config rules accordingly.

A & C are incorrect because the noncompliance is related to the timeliness of the IAM credential report rather than permissions. Option D is incorrect because the MaximumExecutionFrequency value doesn't affect the initial evaluation of the rules; it determines how often the rule is re-evaluated after its first evaluation.
upvoted 3 times

brpjp 1 year ago

Answer D may be correct, on assumption that if maximum execution frequency is 24 hours, then report is one day old rather than 4 hours mentioned on option A. Anyone can clarify my understanding.

upvoted 1 times

  **yorkicurke** 1 year ago

Selected Answer: A

Explained in the following link;

<https://repost.aws/knowledge-center/config-credential-report>

upvoted 1 times

  **yorkicurke** 1 year ago

oh shoot AgboolaKun already mentioned it. ok thumbs up for you AgboolaKun

upvoted 1 times

  **Aamee** 1 year, 1 month ago

Selected Answer: A

Agreed on A.



upvoted 1 times

  **[Removed]** 1 year, 1 month ago

Selected Answer: A

I agree with AgboolaKun. Read the link for some good insight

upvoted 1 times

  **oioi** 1 year, 1 month ago

Selected Answer: D

correct

upvoted 1 times

  **ykhan321** 1 year ago

Anything else besides correct?

upvoted 2 times

A company is using AWS WAF to protect a customized public API service that is based on Amazon EC instances. The API uses an Application Load Balancer.

The AWS WAF web ACL is configured with an AWS Managed Rules rule group. After a software upgrade to the API and the client application, some types of requests are no longer working and are causing application stability issues. A security engineer discovers that AWS WAF logging is not turned on for the web ACL.


The security engineer needs to immediately return the application to service, resolve the issue, and ensure that logging is not turned off in the future. The security engineer turns on logging for the web ACL and specifies Amazon CloudWatch Logs as the destination.

Which additional set of steps should the security engineer take to meet the requirements?

- A. Edit the rules in the web ACL to include rules with Count actions. Review the logs to determine which rule is blocking the request. Modify the IAM policy of all AWS WAF administrators so that they cannot remove the logging configuration for any AWS WAF web ACLs.
- B. Edit the rules in the web ACL to include rules with Count actions. Review the logs to determine which rule is blocking the request. Modify the AWS WAF resource policy so that AWS WAF administrators cannot remove the logging configuration for any AWS WAF web ACLs.
- C. Edit the rules in the web ACL to include rules with Count and Challenge actions. Review the logs to determine which rule is blocking the request. Modify the AWS WAF resource policy so that AWS WAF administrators cannot remove the logging configuration for any AWS WAF web ACLs.
- D. Edit the rules in the web ACL to include rules with Count and Challenge actions. Review the logs to determine which rule is blocking the request. Modify the IAM policy of all AWS WAF administrators so that they cannot remove the logging configuration for any AWS WAF web ACLs.

Suggested Answer: A

Community vote distribution



yorkicurke Highly Voted 1 year, 6 months ago

Selected Answer: A

As many have suggested of why it's unnecessary to go for 'challenge' so C&D -> OUT

As of why not picking B(resource-based) is because resource policy would only control access to that single web ACL.

The question asks to ensure logging is not turned off for any web ACLs[well that's what's implied], which modifying IAM policies globally achieves but modifying a single resource policy does not.

AWS documentation recommends applying least privilege permissions through IAM policies when managing access to resources across multiple accounts. This helps ensure permissions are restricted at the identity level rather than at the individual resource level.

upvoted 5 times

jsopra Most Recent 3 months, 1 week ago

Selected Answer: A

According to the link below, AWS WAF has no resource-based policy, hence the correct answer is A.

https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_aws-services-that-work-with-iam.html

upvoted 1 times

AWSLoverLoverLoverLoverLover 4 months, 1 week ago

Selected Answer: A

The correct answer is A:

Explanation:

Use "Count" action to debug the issue

Instead of outright blocking traffic, changing rules to Count mode helps identify which rule is causing issues without disrupting the API.

This allows normal API operation while collecting logs for analysis.

Review CloudWatch Logs for debugging

Since logging is now enabled, the security engineer can analyze logs in CloudWatch to pinpoint which rule is incorrectly blocking API requests.

Prevent logging from being turned off

The best way to enforce logging is by modifying IAM policies of all AWS WAF administrators.

This ensures that no admin user can disable logging, which aligns with AWS security best practices.

IAM policies are more effective than WAF resource policies for managing administrative actions like disabling logs.

upvoted 1 times

  **TareDHakim** 6 months ago

Selected Answer: A

There is no resource-based policy for AWS WAF. However, AWS WAF can be used with IAM permission policies to secure the service and its resources.

AWS WAF is a managed web application firewall that allows users to create custom rules to block, allow, count, or monitor web requests. These rules can be based on conditions such as IP address ranges, CIDR blocks, countries or regions, or malicious code.

upvoted 1 times

  **IPLogic** 7 months ago

Selected Answer: B

To address the issue and ensure logging remains enabled, the security engineer should:

Edit the rules in the web ACL to include rules with Count actions. This will help identify which rule is blocking the requests by reviewing the logs.

Review the logs to determine the specific rule causing the issue.

Modify the AWS WAF resource policy so that AWS WAF administrators cannot remove the logging configuration for any AWS WAF web ACLs.

Therefore, the correct answer is B.

upvoted 1 times

  **DSEexam** 8 months ago

Selected Answer: A

There is no resource base policy for WAF!

upvoted 4 times

  **kkravets** 8 months ago

Selected Answer: B

B as it is more scallable



upvoted 1 times

  **Pmktechno** 8 months, 1 week ago

Selected Answer: B

Answer B

upvoted 1 times

  **div05jkl** 8 months, 1 week ago

A

upvoted 1 times

  **VerRi** 10 months, 1 week ago

Selected Answer: A

I cannot even see resource policy setting within WAF. There is one in Firewall Manager to handle access control to WAF, but I will go with A in this scenario.


upvoted 2 times

  **cumzle_com** 1 year ago

Selected Answer: B

Both Option A and Option B are strong candidates, but Option B has a slight edge due to the use of AWS WAF resource policies, which are more direct for managing configurations. This method ensures that logging configurations cannot be easily altered by administrators, providing a more robust solution.

upvoted 3 times

  **SamHan** 1 year, 2 months ago

Selected Answer: B

B is correct

upvoted 1 times

🗨️ 👤 **March2023** 1 year, 2 months ago

Selected Answer: B

option B. This involves editing the rules in the web ACL to include rules with Count actions, reviewing the logs to identify the blocking rule, and then modifying the AWS WAF resource policy to prevent AWS WAF administrators from removing the logging configuration for any web ACLs in the future.

upvoted 1 times

🗨️ 👤 **Snape** 1 year, 2 months ago

Selected Answer: B

Resource based policy

upvoted 1 times

🗨️ 👤 **didorins** 1 year, 2 months ago

B - Resource based policy is better than identity policy here.

upvoted 1 times

🗨️ 👤 **hro** 1 year, 3 months ago

Im going with B - the issue is the resource policy and not an IAM policy issue. The AWS WAF resource policy allowed for the logging not to be turned on.

upvoted 1 times

🗨️ 👤 **walter_white_008** 1 year, 3 months ago

Selected Answer: B

Modifying Resource based policy is appropriate here. Purpose is to avoid WAF logging modification.

upvoted 2 times

A security engineer is creating an AWS Lambda function. The Lambda function needs to use a role that is named LambdaAuditRole to assume a role that is named AcmeAuditFactoryRole in a different AWS account.

When the code is processed, the following error message appears: "An error occurred (AccessDenied) when calling the AssumeRole operation."

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

- A. Ensure that LambdaAuditRole has the sts:AssumeRole permission for AcmeAuditFactoryRole.
- B. Ensure that LambdaAuditRole has the AWSLambdaBasicExecutionRole managed policy attached.
- C. Ensure that the trust policy for AcmeAuditFactoryRole allows the sts:AssumeRole action from LambdaAuditRole.
- D. Ensure that the trust policy for LambdaAuditRole allows the sts:AssumeRole action from the lambda.amazonaws.com service.
- E. Ensure that the sts:AssumeRole API call is being issued to the us-east-1 Region endpoint.

Suggested Answer: AC

Community vote distribution

AC (100%)

🗳️ 👤 **ale_brd_111** 9 months ago

Selected Answer: AC

A.

C.

upvoted 1 times

🗳️ 👤 **kejam** 1 year ago

Selected Answer: AC

<https://repost.aws/knowledge-center/lambda-function-assume-iam-role>

upvoted 2 times

🗳️ 👤 **Aamee** 1 year, 1 month ago

Selected Answer: AC

makes the perfect logic.

upvoted 2 times

🗳️ 👤 **[Removed]** 1 year, 1 month ago

Selected Answer: AC

A and C seems right

upvoted 1 times

🗳️ 👤 **oioi** 1 year, 1 month ago

Selected Answer: AC

coorect

upvoted 1 times

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? (Choose two.)

- A. 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 management account to write to the S3 bucket.
- B. 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.
- C. In the dedicated security account, create an Amazon S3 bucket that has an S3 Lifecycle configuration that expires objects after 2 years. Set the bucket policy to allow the organization's member accounts to write to the S3 bucket.
- D. Create an AWS CloudTrail trail for the organization. Configure logs to be delivered to the logging Amazon S3 bucket in the dedicated security account.
- E. Turn on AWS CloudTrail in each account. Configure logs to be delivered to an Amazon S3 bucket that is created in the organization's management account. Forward the logs to the S3 bucket in the dedicated security account by using AWS Lambda and Amazon Kinesis Data Firehose.

Suggested Answer: BD

Community vote distribution



ahrentom Highly Voted 1 year, 7 months ago

Selected Answer: BD

I go with BD, because each Member Account has to write into the security Account S3 bucket, not only the Organization Management Account.

<https://docs.aws.amazon.com/awscloudtrail/latest/userguide/cloudtrail-set-bucket-policy-for-multiple-accounts.html>

upvoted 10 times

kejam 1 year, 7 months ago

Agreed. CloudTrail for Org requires the destination S3 bucket to allow writes from each member account. Object Lock is enabled to prevent the data from being overwritten/deleted.

upvoted 2 times

itsjunukim Most Recent 2 months, 2 weeks ago

Selected Answer: AD

<https://docs.aws.amazon.com/awscloudtrail/latest/userguide/creating-trail-organization.html>

According to the official documentation, when an organization trail is created, the management account collects logs from all member accounts and stores them in the specified S3 bucket on their behalf.

upvoted 4 times

jsopra 3 months, 1 week ago

Selected Answer: AD

The management account's CloudTrail service will automatically write logs from all member accounts to the centralized bucket. Member accounts should not have direct write access to the logging bucket to maintain security and integrity of the logs.

"By default, organization log files are accessible only to the management account. For information about how to allow read access to the Amazon S3 bucket for IAM users in member accounts, see Sharing CloudTrail log files between AWS accounts."

<https://docs.aws.amazon.com/awscloudtrail/latest/userguide/create-s3-bucket-policy-for-cloudtrail.html#org-trail-bucket-policy>

upvoted 4 times

Pmktechno 8 months, 1 week ago

Selected Answer: AD

A and D

upvoted 2 times

🗨️ **lovekiller** 8 months, 3 weeks ago

Selected Answer: AD

To all who are choosing B, the answer is A.

Here's the reasoning:

In an AWS Organizations setup, the management account can be configured to collect CloudTrail logs from all member accounts and forward them to a centralized Amazon S3 bucket in a dedicated security account. This centralized logging approach ensures that all account activity across the organization is captured and securely stored.

upvoted 3 times

🗨️ **navid1365** 11 months, 1 week ago

Selected Answer: BD

B and D

upvoted 1 times

🗨️ **cumzle_com** 1 year ago

Selected Answer: AD

B increases the surface area for potential security issues since multiple member accounts have write access to the bucket.

upvoted 3 times

🗨️ **Ritarocks** 1 year, 3 months ago

AD.

A and not B because, member account number tracking does not make sense, when it's easy to use single Org as reference for Bucket policy.

upvoted 1 times

🗨️ **Ernestokoro** 1 year, 5 months ago

The organization includes a dedicated security account= Member account while ALL OTHER =Management account. this means to me that granting the permission from the Management account reduces operational overhead than doing it at individual member accounts. Therefore I go with option AD.

upvoted 1 times

🗨️ **vikasj1in** 1 year, 5 months ago

A, D

Option B covers the storage aspect by configuring a dedicated S3 bucket in the security account, allowing member accounts to write logs. S3 Object Lock in compliance mode ensures the retention requirements.

Option D complements this by configuring CloudTrail to capture the logs and deliver them to the dedicated S3 bucket directly.

Together, these options cover the log storage, retention, and collection requirements with the least operational overhead.

upvoted 1 times

🗨️ **WeepingMaple** 1 year, 6 months ago

Selected Answer: AD

Enable Organization Trail: In the Management Console or CLI, activate an organization trail that logs all events from all member accounts.

<https://docs.aws.amazon.com/awscloudtrail/latest/userguide/creating-trail-organization.html>

upvoted 2 times

🗨️ **jeff001** 1 year, 6 months ago

Selected Answer: BD

Member account needs to write to S3.

upvoted 2 times

🗨️ **marco25** 1 year, 6 months ago

Selected Answer: BD

trails across member accounts, needs permissions to the sender bucket

upvoted 4 times

🗨️ **Aamee** 1 year, 7 months ago

Selected Answer: AD

If I understand correctly, the reason why the option B can't be a correct one cuz the use case has asked about the logs which must not be deleted or changed which can't be met in option B if we opt for each member's accounts to be given with the full S3 logs access under an organization.

upvoted 2 times

🗨️ 👤 **ykhan321** 1 year, 6 months ago

A has only one account and option B has all the aws accounts.

upvoted 1 times

🗨️ 👤 **confusedyeti69** 1 year, 6 months ago

If following your logic, the management account can delete and change the logs too.

And the options also says to only give write access to S3 only.

It is not A because members need to write S3, not only management.

In compliance mode, a protected object version can't be overwritten or deleted by any user, including the root user in your AWS account.

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/object-lock.html>

upvoted 2 times

🗨️ 👤 **[Removed]** 1 year, 7 months ago

Selected Answer: AD

A and D are correct

upvoted 2 times

🗨️ 👤 **oioi** 1 year, 7 months ago

Selected Answer: AD

correct

upvoted 2 times

A company is testing its incident response plan for compromised credentials. The company runs a database on an Amazon EC2 instance and stores the sensitive database credentials as a secret in AWS Secrets Manager. The secret has rotation configured with an AWS Lambda function that uses the generic rotation function template. The EC2 instance and the Lambda function are deployed in the same private subnet. The VPC has a Secrets Manager VPC endpoint.

A security engineer discovers that the secret cannot rotate. The security engineer determines that the VPC endpoint is working as intended. The Amazon CloudWatch logs contain the following error: "setSecret: Unable to log into database".

Which solution will resolve this error?

- A. Use the AWS Management Console to edit the JSON structure of the secret in Secrets Manager so that the secret automatically conforms with the structure that the database requires.
- B. Ensure that the security group that is attached to the Lambda function allows outbound connections to the EC2 instance. Ensure that the security group that is attached to the EC2 instance allows inbound connections from the security group that is attached to the Lambda function.
- C. Use the Secrets Manager list-secrets command in the AWS CLI to list the secret. Identify the database credentials. Use the Secrets Manager rotate-secret command in the AWS CLI to force the immediate rotation of the secret.
- D. Add an internet gateway to the VPC. Create a NAT gateway in a public subnet. Update the VPC route tables so that traffic from the Lambda function and traffic from the EC2 instance can reach the Secrets Manager public endpoint.

Suggested Answer: B

Community vote distribution

B (100%)

 **yorkicurke** Highly Voted 1 year ago

Selected Answer: B

Hate questions like these which rather than testing your knowledge of technologies trick you into these weird worded questions.

the statement 'Ensure that the security group that is attached to the Lambda function allows outbound' threw me off as Lambda does not have SGs.

But then through some internet digging came across the fact that when a Lambda function needs to access resources inside a Virtual Private Cloud (VPC), it does so using ENI which resides in a subnet of the VPC and can have a security group associated with it. The security group acts as a virtual firewall for the ENI.

upvoted 7 times

 **Daniel76** Most Recent 1 year ago

Selected Answer: B

- a) when you use the console to store a database secret, Secrets Manager automatically creates it in the correct JSON structure.
- c) secret manager already configured as auto-rotation. also, secret id should have been known instead of listing secrets .
- d) accessing secret manager via public is not recommended.

upvoted 1 times

 **confusedyeti69** 1 year ago

Why would the lambda need access to the EC2? The question is unclear about the exact job of the lambda. It is worded as if the lambda job is to change the creds in secrets manager only.

upvoted 2 times

 **JPSWS** 1 year ago

The DB runs on the EC2 that's why the Lambda needs access to it to set the new credentials

upvoted 2 times



 **Aamee** 1 year, 1 month ago

Selected Answer: B

B is the only one that logically seems right. All others are distracters except option C. But option C describes the solution of this problem as a one time thing whereas, it's been asked to provide a permanent solution for this use case. That's why B looks much more secured and valid option among

all others.



upvoted 4 times

  **[Removed]** 1 year, 1 month ago

Selected Answer: B

I'll vote B. The rest are distractors but feel free to correct me if I'm wrong.

upvoted 2 times

  **oioi** 1 year, 1 month ago

Selected Answer: B

correct

upvoted 2 times

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

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

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

A.

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

B.

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

C.

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

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "NotAction": [
        "guardduty:DeleteDetector",
        "guardduty:UpdateDetector",
        "securityhub:DisableSecurityHub"
      ],
      "Resource": [
        "*"
      ]
    }
  ]
}

```

Suggested Answer: A

Community vote distribution

A (80%)

D (20%)

 **ahrentom** Highly Voted 1 year, 1 month ago

Selected Answer: A

A is correct, key word in SCP is to Deny, because it overwrites the FullAccessSCP Allow statement.
upvoted 5 times

 **ale_brd_111** Most Recent 9 months ago

Selected Answer: A

gotta be A
upvoted 2 times

 **AgboolaKun** 1 year, 1 month ago

Selected Answer: A

A is correct. The NotAction element cannot be used in this case.

You only need an explicit DENY here since all accounts and OUs already have a default FullAWSAccess SCP but you don't want them to be able to disable Amazon GuardDuty and AWS Security Hub.

upvoted 3 times

 **Sab31** 11 months, 4 weeks ago

Kindly correct me if I am wrong. When we attach a new SCP the default FullAWSAccess SCP is detached from the OU. isn't that right?
upvoted 1 times

 **Aamee** 1 year, 1 month ago

Selected Answer: D

Probably going with D but still not 100% sure how is it going to work that way... would appreciate if someone could help in understanding this question..
upvoted 2 times

 **LeoD** 1 year ago

SCPs do not support NotAction with effect Allow.
upvoted 2 times

 **Aamee** 1 year ago

Ah ok, got it...thnx so much... in this way, probably looks like all other options are invalid except option A since on all others they've used 'NotAction' attribute with Allow directly and indirectly which won't work..
upvoted 2 times

 **Zek** 7 months, 2 weeks ago

https://docs.aws.amazon.com/organizations/latest/userguide/orgs_manage_policies_scps_syntax.html#scp-elements-table
upvoted 1 times

 **[Removed]** 1 year, 1 month ago

A. OU level will still have access to other services outside of Guardduty and Security Hub due to the OU level policy. D could work but is not necessary

upvoted 1 times

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 credentials. Encrypt the CloudFormation template by using AWS KMS.
- C. Use a SecureString parameter in the CloudFormation template to reference the database credentials in Secrets Manager.
- D. Use a SecureString parameter in the CloudFormation template to reference an encrypted value in AWS KMS.

Suggested Answer: A

Community vote distribution

A (100%)

  **Aamee** Highly Voted 1 year ago

Since this is the last question here so maybe I can post it here. I've passed this exam with a score of 926. Only few of the questions were not from this exam material but else, everything came from here. Would like to thanks to all of you who helped in answering my queries and got my concept clarified!... Man_Kind, Agboola and others, you guys simply rock, thanks once again so much! :)

upvoted 19 times

  **giancesarini2023** 1 year ago

@Aamee, do you think there is a question from 1 to 50? I'm only studying from 50 to 115.

upvoted 2 times

  **saptati** Highly Voted 9 months, 2 weeks ago

I took the exam on 13 Feb 2024. Around 25 questions came from here. The rest I answered by myself. If you are an experienced AWS Professional, you won't find it difficult to pass the exam. If you are a novice, then wait for the exam topics to update the question bank. The 121 questions aren't enough. All the best, and thanks everyone for contributing to the discussion.

upvoted 6 times

  **saptati** 9 months, 2 weeks ago

Correction, I took the exam on 13 March 2024.

upvoted 2 times

  **phmeeeeee** Most Recent 2 months, 3 weeks ago

Selected Answer: A

all the invalid options are to store sensitive information (the RDS database password in this case) in PLAINTEXT format.

upvoted 1 times

  **nn67** 10 months, 4 weeks ago

A

keyword dynamic reference

upvoted 2 times

  **Pmktechno** 11 months ago

Yesterday I took this exam (Feb 1st) single question also wasn't came from this set of questions. Please wait examtopics team should be update soon new set of questions.

upvoted 3 times

  **brpjp** 11 months, 2 weeks ago

Hello, I passed exam with 956 score. Thank you all for contributing and correcting the answers.

upvoted 1 times

  **alexleely** 11 months ago

when did you take the exam?

upvoted 1 times

🗨️ 👤 **Aamee** 1 year, 1 month ago

Selected Answer: A

Yup, for sure it should be A. Here's the summary:

"Updating a secret in Secrets Manager doesn't automatically update the secret in CloudFormation. In order for CloudFormation to update a secretsmanager dynamic reference, you must perform a stack update that updates the resource containing the dynamic reference, either by updating the resource property that contains the secretsmanager dynamic reference, or updating another of the resource's properties."

upvoted 1 times

🗨️ 👤 **[Removed]** 1 year, 1 month ago

Selected Answer: A

A. See below

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/dynamic-references.html#dynamic-references-secretsmanager>

upvoted 4 times

🗨️ 👤 **oioi** 1 year, 1 month ago

Selected Answer: A

correct

upvoted 1 times

An international company wants to combine AWS Security Hub findings across all the company's AWS Regions and from multiple accounts. In addition, the company wants to create a centralized custom dashboard to correlate these findings with operational data for deeper analysis and insights. The company needs an analytics tool to search and visualize Security Hub findings.

Which combination of steps will meet these requirements? (Chose three.)

- A. Designate an AWS account as a delegated administrator for Security Hub. Publish events to Amazon CloudWatch from the delegated administrator account, all member accounts, and required Regions that are enabled for Security Hub findings.
- B. Designate an AWS account in an organization in AWS Organizations as a delegated administrator for Security Hub. Publish events to Amazon EventBridge from the delegated administrator account, all member accounts, and required Regions that are enabled for Security Hub findings.
- C. In each Region, create an Amazon EventBridge rule to deliver findings to an Amazon Kinesis data stream. Configure the Kinesis data streams to output the logs to a single Amazon S3 bucket.
- D. In each Region, create an Amazon EventBridge rule to deliver findings to an Amazon Kinesis Data Firehose delivery stream. Configure the Kinesis Data Firehose delivery streams to deliver the logs to a single Amazon S3 bucket.
- E. Use AWS Glue DataBrew to crawl the Amazon S3 bucket and build the schema. Use AWS Glue Data Catalog to query the data and create views to flatten nested attributes. Build Amazon QuickSight dashboards by using Amazon Athena.
- F. Partition the Amazon S3 data. Use AWS Glue to crawl the S3 bucket and build the schema. Use Amazon Athena to query the data and create views to flatten nested attributes. Build Amazon QuickSight dashboards that use the Athena views.

Suggested Answer: BDF

Community vote distribution

BDF (100%)

🗳️ 👤 **ion_gee** 8 months, 3 weeks ago

Selected Answer: BDF

Not A, do not need cloud watch.

Not C, Kinesis Data firehose(Now Amazon Data Firehose) is what we need here, not Kinesis Data Streams.

Not E . Rather use Athena to query S3

See Ref <https://aws.amazon.com/blogs/architecture/visualize-aws-security-hub-findings-using-analytics-and-business-intelligence-tools/>
upvoted 2 times

🗳️ 👤 **ale_brd_111** 9 months ago

Selected Answer: BDF

BDF probably

upvoted 1 times

🗳️ 👤 **nubliti** 10 months, 1 week ago

Selected Answer: BDF

BDF are the best options

upvoted 1 times

🗳️ 👤 **awssecuritynewbie** 10 months, 2 weeks ago

Selected Answer: BDF

BDF for sure,

upvoted 1 times

🗳️ 👤 **sarcactus** 10 months, 2 weeks ago

Selected Answer: BDF

BDF

Also agree with previous comment.

upvoted 1 times

🗳️ 👤 **MikeRach** 10 months, 2 weeks ago

BDF

The steps are literally provided in this Doc <https://aws.amazon.com/blogs/architecture/visualize-aws-security-hub-findings-using-analytics-and-business-intelligence-tools/>

upvoted 2 times

An AWS account administrator created an IAM group and applied the following managed policy to require that each individual user authenticate using multi-factor authentication:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "ec2:*",
      "Resource": "*"
    },
    {
      "Sid": "BlockAnyAccessUnlessSignedInWithMFA",
      "Effect": "Deny",
      "Action": "ec2:*",
      "Resource": "*",
      "Condition": {
        "BoolIfExists": {
          "aws:MultiFactorAuthPresent": false
        }
      }
    }
  ]
}
```

After implementing the policy, the administrator receives reports that users are unable to perform Amazon EC2 commands using the AWS CLI.

What should the administrator do to resolve this problem while still enforcing multi-factor authentication?

- A. Change the value of `aws:MultiFactorAuthPresent` to `true`.
- B. Instruct users to run the `aws sts get-session-token` CLI command and pass the multi-factor authentication `--serial-number` and `-token-code` parameters. Use these resulting values to make API/CLI calls.
- C. Implement federated API/CLI access using SAML 2.0, then configure the identity provider to enforce multi-factor authentication.
- D. Create a role and enforce multi-factor authentication in the role trust policy. Instruct users to run the `sts assume-role` CLI command and pass `--serial-number` and `--token-code` parameters. Store the resulting values in environment variables. Add `sts:AssumeRole` to `NotAction` in the policy.

Suggested Answer: B

Community vote distribution

B (100%)

🗳️ **molerowan** 3 months, 2 weeks ago

Selected Answer: B

`aws sts get-session-token` with MFA parameters creates temporary credentials that include MFA verification. These temporary credentials (access key, secret key, and session token) can be used for subsequent CLI calls. The AWS CLI will then recognize these requests as being made with MFA authentication, satisfying the policy condition.

upvoted 1 times

🗳️ **nublitt** 10 months, 1 week ago

Selected Answer: B

B for sure

upvoted 1 times

🗳️ **awssecuritynewbie** 10 months, 2 weeks ago

Selected Answer: B

B for sure they required to enable CLI

upvoted 1 times

🗳️ **sarcactus** 10 months, 2 weeks ago

Selected Answer: B

I agree with MikeRach comment.

upvoted 3 times

  **MikeRach** 10 months, 3 weeks ago

B

<https://www.examttopics.com/discussions/amazon/view/47596-exam-aws-certified-security-specialty-topic-1-question-225/>

upvoted 1 times

A company is developing a mechanism that will help data scientists use Amazon SageMaker to read, process, and output data to an Amazon S3 bucket. Data scientists will have access to a dedicated S3 prefix for each of their projects. The company will implement bucket policies that use the dedicated S3 prefixes to restrict access to the S3 objects. The projects can last up to 60 days.

The company's security team mandates that data cannot remain in the S3 bucket after the end of the projects that use the data.

Which solution will meet these requirements MOST cost-effectively?

- A. Create an AWS Lambda function to identify and delete objects in the S3 bucket that have not been accessed for 60 days. Create an Amazon EventBridge scheduled rule that runs every day to invoke the Lambda function.
- B. Create a new S3 bucket. Configure the new S3 bucket to use S3 Intelligent-Tiering. Copy the objects to the new S3 bucket.
- C. Create an S3 Lifecycle configuration for each S3 bucket prefix for each project. Set the S3 Lifecycle configurations to expire objects after 60 days.
- D. Create an AWS Lambda function to delete objects that have not been accessed for 60 days. Create an S3 event notification for S3 Intelligent-Tiering automatic archival events to invoke the Lambda function.

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ **rishikeshshukla777** Highly Voted 9 months, 2 weeks ago

anyone given exam in this week or last week can anyone confirm is this question is good for practise or required more question ?

upvoted 5 times

🗳️ **ion_gee** Most Recent 8 months, 3 weeks ago

Selected Answer: C

C makes the most sense

upvoted 1 times

🗳️ **ale_brd_111** 9 months ago

Selected Answer: C

no brainer C.

upvoted 1 times

🗳️ **Gzuiz** 9 months, 3 weeks ago

my vote is for C

upvoted 1 times

🗳️ **nublitt** 10 months, 1 week ago

Selected Answer: C

C for sure

upvoted 1 times

🗳️ **sarcactus** 10 months, 2 weeks ago

Selected Answer: C

C is the correct one.

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/lifecycle-configuration-examples.html#lifecycle-config-conceptual-ex3>

upvoted 1 times

🗳️ **anasbakla** 10 months, 2 weeks ago

C is the answer

upvoted 2 times

🗳️ **anasbakla** 10 months, 2 weeks ago

Me too



upvoted 1 times

🗳️ **awssecuritynewbie** 10 months, 2 weeks ago

Selected Answer: C

C makes sense you need to remove the data after 60 day so lifecycle will do that.

upvoted 1 times

  **jabilrn** 10 months, 2 weeks ago

C for me

upvoted 1 times

A company has AWS accounts that are in an organization in AWS Organizations. An Amazon S3 bucket in one of the accounts is publicly accessible.

A security engineer must change the configuration so that the S3 bucket is no longer publicly accessible. The security engineer also must ensure that the S3 bucket cannot be made publicly accessible in the future.

Which solution will meet these requirements?

- A. Configure the S3 bucket to use an AWS Key Management Service (AWS KMS) key. Encrypt all objects in the S3 bucket by creating a bucket policy that enforces encryption. Configure an SCP to deny the s3:GetObject action for the OU that contains the AWS account.
- B. Enable the PublicAccessBlock configuration on the S3 bucket. Configure an SCP to deny the s3:GetObject action for the OU that contains the AWS account.
- C. Enable the PublicAccessBlock configuration on the S3 bucket. Configure an SCP to deny the s3:PutPublicAccessBlock action for the OU that contains the AWS account.
- D. Configure the S3 bucket to use S3 Object Lock in governance mode. Configure an SCP to deny the s3:PutPublicAccessBlock action for the OU that contains the AWS account.

Suggested Answer: C

Community vote distribution

C (100%)

🗳️ **nznzwell** 5 months, 1 week ago

Selected Answer: C

For those who chose B, how can it be correct? an SCP denying s3:GetObject will deny access to the objects regardless if it is from the public or within the account... B will render the bucket inaccessible.

upvoted 1 times

🗳️ **helloworldabc** 9 months, 2 weeks ago

just B

upvoted 1 times

🗳️ **Olaunfazed** 1 year ago

Answer is B

Enabling the PublicAccessBlock configuration on the S3 bucket prevents public access.

Additionally, configuring an SCP (Service Control Policy) to deny the s3:GetObject action for the organizational unit (OU) containing the AWS account ensures that the bucket remains private.

upvoted 1 times

🗳️ **Davidng88** 9 months, 3 weeks ago

By configuring an SCP to deny the s3:GetObject action prevents accessing S3 objects from public, but deny the s3:PutPublicAccessBlock action, prevent any changes to the disable PublicAccessBlock settings, ensuring that the bucket remains private in future.

upvoted 1 times

🗳️ **sema2232** 1 year ago

why not B?

upvoted 1 times

🗳️ **aescudero51** 1 year, 1 month ago

Selected Answer: C

Enable PublicAccessBlock Configuration:

https://aws.amazon.com/s3/features/block-public-access/?nc1=h_ls

Configure an SCP (Service Control Policy):

An SCP is a policy that you can attach to an AWS Organization, organizational unit (OU), or an account. It acts as a guardrail to control permissions across accounts. In your case, you want to deny the s3:PutPublicAccessBlock action for the OU containing your AWS account.

Go to the AWS Organizations console.

Navigate to the OU that contains your account.



Create a new SCP or edit an existing one.

Add a statement that denies the s3:PutPublicAccessBlock action for the relevant S3 buckets.

Attach the SCP to the OU.

Ensure that your AWS account is part of the OU.

upvoted 2 times

  **Zek** 1 year, 1 month ago

C

Enable the PublicAccessBlock & use SCP to deny the s3:PutPublicAccessBlock action

upvoted 1 times

A company is designing a new application stack. The design includes web servers and backend servers that are hosted on Amazon EC2 instances. The design also includes an Amazon Aurora MySQL DB cluster.

The EC2 instances are in an Auto Scaling group that uses launch templates. The EC2 instances for the web layer and the backend layer are backed by Amazon Elastic Block Store (Amazon EBS) volumes. No layers are encrypted at rest. A security engineer needs to implement encryption at rest.

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

- A. Modify EBS default encryption settings in the target AWS Region to enable encryption. Use an Auto Scaling group instance refresh.
- B. Modify the launch templates for the web layer and the backend layer to add AWS Certificate Manager (ACM) encryption for the attached EBS volumes. Use an Auto Scaling group instance refresh.
- C. Create a new AWS Key Management Service (AWS KMS) encrypted DB cluster from a snapshot of the existing DB cluster.
- D. Apply AWS Key Management Service (AWS KMS) encryption to the existing DB cluster.
- E. Apply AWS Certificate Manager (ACM) encryption to the existing DB cluster.

Suggested Answer: AC


Community vote distribution

AC (100%)

  **b5f86df** 1 week, 4 days ago

Selected Answer: AC

The Auto Scaling group instance refresh will replace running instances so that the new ones launch with encrypted volumes.
upvoted 1 times

  **sema2232** 1 year ago

why not B?

upvoted 1 times

  **helloworldabc** 9 months, 2 weeks ago

just AC

upvoted 1 times

  **aescudero51** 1 year, 1 month ago

Selected Answer: AC

A/C

B - You don't use ACM for encryption, it's KMS

D - You can't encrypt an existing cluster, you need to snapshot, then encrypt with KMS

E - Same as B

upvoted 4 times

  **Zek** 1 year, 1 month ago

A,C - Agree

upvoted 1 times

  **danish1234** 1 year, 1 month ago

Selected Answer: AC

AC . You can not encrypt ebs with ACM.

upvoted 2 times

  **danish1234** 1 year, 1 month ago

AC . All other options are joke.

upvoted 1 times