

SOA-C02^{Q&As}

AWS Certified SysOps Administrator - Associate (SOA-C02)

Pass Amazon SOA-C02 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.leads4pass.com/soa-c02.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Amazon
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers



QUESTION 1

A global company operates out of five AWS Regions. A SysOps administrator wants to identify all the company's tagged and untagged Amazon EC2 instances.

The company requires the output to display the instance ID and tags.

What is the MOST operationally efficient way for the SysOps administrator to meet these requirements?

- A. Create a tag-based resource group in AWS Resource Groups.
- B. Use AWS Trusted Advisor. Export the EC2 On-Demand Instances check results from Trusted Advisor.
- C. Use Cost Explorer. Choose a service type of EC2-Instances, and group by Resource.
- D. Use Tag Editor in AWS Resource Groups. Select all Regions, and choose a resource type of AWS::EC2::Instance.

Correct Answer: D

With Tag Editor, you build a query to find resources in one or more AWS Regions that are available for tagging. You can choose up to 20 individual resource types, or build a query on All resource types. Your query can include resources that already have tags, or resources that have no tags. <https://docs.aws.amazon.com/ARG/latest/userguide/find-resources-to-tag.html>

QUESTION 2

A company needs to automatically monitor an AWS account for potential unauthorized AWS Management Console logins from multiple geographic locations. Which solution will meet this requirement?

- A. Configure Amazon Cognito to detect any compromised IAM credentials.
- B. Set up Amazon Inspector. Scan and monitor resources for unauthorized logins.
- C. Set up AWS Config. Add the iam-policy-blacklisted-check managed rule to the account.
- D. Configure Amazon GuardDuty to monitor the UnauthorizedAccess:IAMUser/ConsoleLoginSuccess finding.

Correct Answer: D

Guard duty IAM finding types: UnauthorizedAccess:IAMUser/ConsoleLoginSuccess.B
UnauthorizedAccess:IAMUser/InstanceCredentialExfiltration.OutsideAWS
UnauthorizedAccess:IAMUser/InstanceCredentialExfiltration.OutsideAWS
UnauthorizedAccess:IAMUser/MaliciousIPCaller UnauthorizedAccess:IAMUser/MaliciousIPCaller.Custom
UnauthorizedAccess:IAMUser/TorIPCaller

QUESTION 3

A company uses Amazon Elasticsearch Service (Amazon ES) to analyze sales and customer usage data. Members of the company's geographically dispersed sales team are traveling. They need to log in to Kibana by using their existing corporate credentials that are stored in Active Directory. The company has deployed Active Directory Federation Services (AD FS) to enable authentication to cloud services.

Which solution will meet these requirements?

- A. Configure Active Directory as an authentication provider in Amazon ES. Add the Active Directory server's domain name to Amazon ES. Configure Kibana to use Amazon ES authentication.
- B. Deploy an Amazon Cognito user pool. Configure Active Directory as an external identity provider for the user pool. Enable Amazon Cognito authentication for Kibana on Amazon ES.
- C. Enable Active Directory user authentication in Kibana. Create an IP-based custom domain access policy in Amazon ES that includes the Active Directory server's IP address.
- D. Establish a trust relationship with Kibana on the Active Directory server. Enable Active Directory user authentication in Kibana. Add the Active Directory server's IP address to Kibana.

Correct Answer: B

<https://aws.amazon.com/blogs/security/how-to-enable-secure-access-to-kibana-using-aws-single-sign-on/>
<https://docs.aws.amazon.com/elasticsearch-service/latest/developerguide/es-cognito-auth.html>

QUESTION 4

A SysOps administrator must ensure that a company's Amazon EC2 instances auto scale as expected. The SysOps administrator configures an Amazon EC2 Auto Scaling Lifecycle hook to send an event to Amazon EventBridge (Amazon CloudWatch Events), which then invokes an AWS Lambda function to configure the EC2 instances. When the configuration is complete, the Lambda function calls the complete Lifecycle-action event to put the EC2 instances into service. In testing, the SysOps administrator discovers that the Lambda function is not invoked when the EC2 instances auto scale.

What should the SysOps administrator do to resolve this issue?

- A. Add a permission to the Lambda function so that it can be invoked by the EventBridge (CloudWatch Events) rule.
- B. Change the lifecycle hook action to CONTINUE if the lifecycle hook experiences a failure or timeout.
- C. Configure a retry policy in the EventBridge (CloudWatch Events) rule to retry the Lambda function invocation upon failure.
- D. Update the Lambda function execution role so that it has permission to call the complete lifecycle-action event.

Correct Answer: A

To allow the EventBridge (CloudWatch Events) rule to invoke the Lambda function, the function's execution role needs to have the necessary permissions to be invoked by the rule. Specifically, the execution role needs to have an event pattern that matches the rule and an IAM policy that grants the necessary permissions to execute the Lambda function. By adding the necessary permissions to the Lambda function, the SysOps administrator can ensure that the function is invoked when the EC2 instances auto scale.

Option D is incorrect because updating the Lambda function execution role so that it has permission to call the complete-lifecycle-action event will not address the issue of the Lambda function not being invoked by the EventBridge (CloudWatch Events) rule.

QUESTION 5

A company runs its entire suite of applications on Amazon EC2 instances. The company plans to move the applications to containers and AWS Fargate. Within 6 months, the company plans to retire its EC2 instances and use only Fargate.

The company has been able to estimate its future Fargate costs.

A SysOps administrator needs to choose a purchasing option to help the company minimize costs. The SysOps administrator must maximize any discounts that are available and must ensure that there are no unused reservations.

Which purchasing option will meet these requirements?

- A. Compute Savings Plans for 1 year with the No Upfront payment option
- B. Compute Savings Plans for 1 year with the Partial Upfront payment option
- C. EC2 Instance Savings Plans for 1 year with the All Upfront payment option
- D. EC2 Reserved Instances for 1 year with the Partial Upfront payment option

Correct Answer: B

Given the company's plan to move to Fargate within 6 months and retire EC2 instances, it might be more cost-efficient to choose Option A (No Upfront payment). This way, the company avoids any upfront commitment and can easily transition to Fargate without being tied to EC2 instances. Savings Plans apply to both EC2 and Fargate, making it a suitable option for the planned migration.

QUESTION 6

CORRECT TEXT

If your AWS Management Console browser does not show that you are logged in to an AWS account, close the browser and relaunch the console by using the AWS Management Console shortcut from the VM desktop.

If the copy-paste functionality is not working in your environment, refer to the instructions file on the VM desktop and use Ctrl+C, Ctrl+V or Command-C, Command-V.

Use the following configuration requirements to create an Amazon DynamoDB Accelerator (DAX) cluster and modify an existing DynamoDB table.

1.

Use the us-east-2 Region for all resources.

2.

Use the default configuration settings unless different settings are specified in the following instructions.

3.

Configure a DAX cluster to expire cached data items after 240 seconds and to expire cached queries after 120 seconds. ***Note: Configure these values before you finalize creation of the cluster. Otherwise, you will have to wait until cluster creation is complete before you can do this step.

4.

Create a three-node DynamoDB DAX cluster that is named DaxLabCluster:

a.

Use dax.t3.small instances.

b.

Use the LabVPC VPC and the PrimaryPrivateSubnet and FailoverPrivateSubnet subnets.

c.

Use the LabDAXSG security group.

d.

Configure the DAX cluster to use the DynamoDBAccessRole IAM role.

5. Modify the LabDynamoDBTable DynamoDB table so that the table uses on-demand capacity.

Note: Do NOT wait until cluster creation is complete before you submit this exam lab.

Important: Click the Next button to complete this lab and continue to the next lab. Once you click the Next button, you will NOT be able to return to this lab.

A. Check the answer in explanation.

B. Place Holder

Correct Answer: A

Steps mentioned above.

Select us-east2 region Congure DAX eluser with 240 seconds and create 3 cluster.

Use LabDAX5G security group to be configured.

Recently visited [Info](#)

No recently visited services

Explore one of these commonly visited AWS services.

[IAM](#) [EC2](#) [S3](#) [RDS](#) [Lambda](#)

Recently visited

Favorites

All services

- Analytics
- Application Integration
- AR & VR
- AWS Cost Management
- Blockchain
- Business Applications
- Compute
- Containers
- Customer Enablement
- Database**
- Developer Tools
- End User Computing
- Front-end Web & Mobile
- Game Development

Database

- Amazon DocumentDB**
Fully-managed MongoDB-compatible database service
- DynamoDB**
Managed NoSQL Database
- ElastiCache**
In-Memory Cache
- Amazon Keyspaces**
Serverless Cassandra-compatible database
- Amazon MemoryDB for Redis**
Fully managed, Redis-compatible, in-memory database service
- Neptune**
Fast, reliable graph database built for the cloud
- Amazon QLDB**
Fully managed ledger database
- RDS**
Managed Relational Database Service

Share your feedback on Amazon DynamoDB [Share feedback](#)

Share your feedback on Amazon DynamoDB. Your feedback is an important part of helping us provide a better customer experience. Take this short survey to let us know how we're doing.

DynamoDB

- Dashboard
- Tables
- [Update settings](#)
- Explore Items
- PartiQL editor
- Backups
- Exports to S3
- Imports from S3
- Reserved capacity
- Settings

DAX

- Clusters
- Subnet groups
- Parameter groups
- Events

DynamoDB > Dashboard

Dashboard

Alarms (0) [Manage in CloudWatch](#)

Alarm name	Status
No custom alarms	

DAX clusters (0) [View details](#)

Cluster name	Status
No clusters	

The image displays three sequential screenshots of the Amazon DynamoDB console interface, illustrating the steps to access the 'Additional settings' for a specific table.

- Top Screenshot:** Shows the 'DynamoDB > Tables' view. A search filter 'Any table tag' is applied. The table 'LadyDynamoDBTable' is listed. The 'Actions' menu is open, and the 'Explore table items' button is highlighted.
- Middle Screenshot:** Shows the 'LabDynamoDBTable' overview page. The 'Overview' tab is selected. A notification box titled 'Protect your DynamoDB table from accidental writes and deletes' is visible, with an 'Edit PITR' button.
- Bottom Screenshot:** Shows the 'LabDynamoDBTable' 'Additional settings' page. The 'Read/write capacity' section is expanded, showing the 'Capacity mode' is set to 'Provisioned'. An 'Edit' button is visible.

DynamoDB > Tables > LabDynamoDBTable > Edit capacity

Edit read/write capacity

Capacity mode [Info](#)

On-demand
Simplify billing by paying for the actual reads and writes your application performs.

Provisioned
Manage and optimize the price by allocating read/write capacity in advance.

► **Capacity calculator**

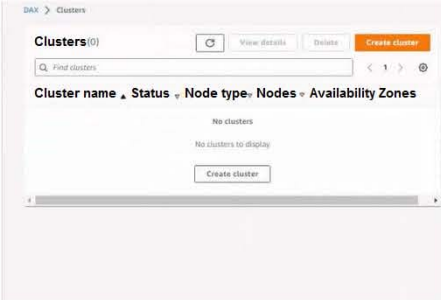
Table capacity

Read capacity

Auto scaling [Info](#)
Dynamically adjusts provisioned throughput capacity on your behalf in response to actual traffic patterns.

On

Off



DAX > Clusters > Create Cluster

Step 1 Choose cluster nodes

Step 2 Configure networks

Step 3 Configure security

Step 4 - optional Verify advanced settings

Step 5 Review and create

Choose cluster nodes

Cluster name

Cluster name
Provide a meaningful name that uniquely identifies your DAX cluster.

DaxLabCluster

Must be between 1 and 20 characters; begin with a letter; contain only ASCII letters, digits, and hyphens; and not end with a hyphen or contain two consecutive hyphens.

Cluster description - optional

Maximum 255 characters.

Node type family

All families
Compare all node families.

t-type family
Provides a baseline level of CPU performance with the ability to burst above the baseline when needed. Recommended for use cases requiring lower throughput.

r-type family
Each node is allocated with fixed resources, for always-ready capacity.

Node types (1/23)

Find node types

Node type	vCPU	Memory (GiB)	Network performance
<input checked="" type="radio"/> dax.r5.large	2	16.00	Moderate
<input type="radio"/> dax.r5.xlarge	4	32.00	Moderate
<input type="radio"/> dax.r5.2xlarge	8	64.00	High
<input type="radio"/> dax.r5.4xlarge	16	128.00	High
<input type="radio"/> dax.r5.8xlarge	32	256.00	10 Gigabit

Node types (1/4)

Find node types

Node type	vCPU	Memory (GiB)	Network performance
<input checked="" type="radio"/> dax.t2.small	1	2.00	Low to Moderate
<input type="radio"/> dax.t2.medium	2	4.00	Low to Moderate
<input type="radio"/> dax.t3.small	2	2.00	Low to Moderate
<input type="radio"/> dax.t3.medium	2	4.00	Low to Moderate

Cluster size

Number of nodes

For a cluster requiring high availability, we strongly recommend at least three nodes. You can scale the number of nodes up or down later.

3

Cancel **Next**

Step 4 - optional
Verify advanced settings

Step 5
Review and create

Choose existing
 Create new

New subnet group

Subnet group name

subnet group

Subnet group description - optional

Enter subnet group description

Maximum 255 characters.

VPC ID

The VPC environment where your DAX cluster will run.

vpc-0336e1c58b3b7b890(LabVPC)

View in VPC console

Subnets

Choose one or more subnets within the VPC. DAX will select IP addresses from these subnets, and assign the IP addresses to the nodes in your DAX cluster.

Choose subnets

View in VPC console

Subnets

Choose one or more subnets within the VPC. DAX will select IP addresses from these subnets, and assign the IP addresses to the nodes in your DAX cluster.

Choose subnets

subnet-0d90ece1d9391e88f x
Region-us-east-2a CIDR:
172.30.0.128/25

subnet-0f8c741641e6d1256 x
Region-us-east-2b CIDR:
172.30.0.128/25

Access control

Security Group

It is security group action or firewall that controls network access to your DAX cluster.

LabDAXSG(vpc-0336e1c58b3b7b890)

View in EC2 console

To access the DAX cluster from your application, you must turn on inbound access on port 8111 for this security group, or port 8111 if encrypted in transit. For detailed instructions, see [Configure Security Group Inbound Rules](#).

Availability Zones (AZ)

AZ allocation

An Availability Zone is one or more discrete data centers with redundant power, networking, and connectivity in an AWS Region. Your cluster nodes can be deployed to several zones to increase the availability of your cluster.

Automatic

We will spread your nodes evenly across AZs for best availability.

Specify manually

Distribute your nodes over the AZs you choose.

Cancel Previous **Next**

Step 2
Configure networks

Step 3
Configure security

Step 4 - optional
Verify advanced settings

Step 5
Review and create

IAM permissions

IAM Service role for DynamoDB access

Choose the IAM service role that your cluster nodes will assume when accessing your DynamoDB tables. The policy on this role controls which subnets they can access, and which operations they can run.

Choose existing
 Create new

DynamoDBAccessRole

View in IAM console

Encryption

Turn on encryption at rest

Protects your data while it is stored, at no additional cost. You cannot change this setting after the cluster is created. We recommend turning on encryption when possible.

Turn on encryption in transit

Protects your data in transit, at no additional cost. Only the latest versions of the DAX client are compatible with encryption in transit. You cannot change this setting after the cluster is created. We recommend turning on encryption when possible.

Cancel Previous **Next**

Step 2
Configure networks

Step 1
Configure security

Step 4 - optional
Verify advanced settings

Step 5
Review and create

Parameter group

Parameter group
Parameter groups are sets of configurations that apply to all the nodes in your cluster.

Choose existing
 Create new

default.dax1.0

Details of the chosen group

Group description	Item time-to-live (TTL)	Query time-to-live (TTL)
default.dax1.0	5 minutes	5 minutes

[Edit parameter group](#)

Maintenance window

DAX occasionally applies software upgrades and patches to all of your nodes during a weekly maintenance window. This operation might affect the cluster's performance.

No preference
Maintenance will be done during time windows DAX will determine.

Specify time window
Maintenance will only occur during the time window you specify.

Step 2
Configure networks

Step 1
Configure security

Step 4 - optional
Verify advanced settings

Step 5
Review and create

Parameter group

Parameter group
Parameter groups are sets of configurations that apply to all the nodes in your cluster.

Choose existing
 Create new

Create new parameter group

Group name

Must be between 1 and 255 characters; begin with a letter; contain only ASCII letters, digits, and hyphens; and has and with a hyphen or contain two consecutive hyphens.

Group description - optional

Enter description

Maximum: 255 characters.

Item time-to-live (TTL)
The maximum amount of time a data item can remain in the cache without being evicted.

Expire after... 240 seconds

An integer greater than zero.

Query time-to-live (TTL)
The maximum amount of time a query can remain in the cache without being evicted.

Expire after... 120 seconds

An integer greater than zero.

Tags

Tags are pairs of keys and optional values, that you can assign to AWS resources. You can use tags to control access to your resources or track your AWS spending.

No tags are associated with the resource.

[Add new tag](#)

You can add 50 more tags

Step 2
Configure networks

Step 3
Configure security

Step 4 - optional
Verify advanced settings

Step 5
Review and create

Step 1: Choose cluster node

Cluster name [Edit](#)

Cluster name	Cluster description
DaxLabCluster	--

Nodes [Edit](#)

Node type	Number of nodes
dax.t3.small	3

You cannot change the node type after cluster is created. Make sure this node type fits your needs.

Step 2: Configure networks

Subnets [Edit](#)

Subnet group	Virtual Private Cloud (VPC) ID	Subnets
subnetgroup	vpc-0336e1c58b3b7b890	2

Step 2
Configure networks

Step 3
Configure security

Step 4 – optional
Verify advanced settings

Step 5
Review and create

Step 4: Verify advanced settings – optional

[Edit](#)

Parameter group

[Edit](#)

Parameter group name	Parameter group description	Item time-to-live (TTL)	Query time-to-live (TTL)
parametergroup	-	240 seconds	120 seconds

Maintenance window

Maintenance
No preference

Tags (0)

Key	Value
No tags are associated with the resource.	

- Update settings
- Explore items
- PartiQL editor
- Backups**
- Exports to S3
- Imports from S3
- Reserved capacity
- Settings

Clusters (1)

Find clusters

Cluster name	Status	Node type	Nodes	Availability Zones
daxlabcluster	Creating	dax.t3.small	3	-

- ▼ DAX
 - Clusters
 - Subnet groups
 - Parameter groups
 - Events

QUESTION 7

A company is running production workloads that use a Multi-AZ deployment of an Amazon RDS for MySQL db.m6g.xlarge (general purpose) standard DB instance. Users report that they are frequently encountering a "too many connections" error. A SysOps administrator observes that the number of connections on the database is high.

The SysOps administrator needs to resolve this issue while keeping code changes to a minimum.

Which solution will meet these requirements MOST cost-effectively?

- A. Modify the RDS for MySQL DB instance to a larger instance size.
- B. Modify the RDS for MySQL DB instance to Amazon DynamoDB.
- C. Configure RDS Proxy. Modify the application configuration file to use the RDS Proxy endpoint.
- D. Modify the RDS for MySQL DB instance to a memory optimized DB instance.

Correct Answer: C

The "too many connections" error indicates that the Amazon RDS for MySQL DB instance is reaching its maximum allowed connections, causing users to encounter issues. RDS Proxy is a highly recommended solution to manage database

connections and improve scalability and availability.

By implementing RDS Proxy and updating the application's configuration to use the proxy endpoint, you can effectively manage connections and alleviate the "too many connections" issue without making significant code changes.

QUESTION 8

A company is implementing security and compliance by using AWS Trusted Advisor. The company's SysOps team is validating the list of Trusted Advisor checks that it can access. Which factor will affect the quantity of available Trusted Advisor checks?

- A. Whether at least one Amazon EC2 instance is in the running state
- B. The AWS Support plan
- C. An AWS Organizations service control policy (SCP)
- D. Whether the AWS account root user has multi-factor authentication (MFA) enabled

Correct Answer: B

AWS Basic Support and AWS Developer Support customers get access to 6 security checks (S3 Bucket Permissions, Security Groups - Specific Ports Unrestricted, IAM Use, MFA on Root Account, EBS Public Snapshots, RDS Public Snapshots) and 50 service limit checks. AWS Business Support, AWS Enterprise On-Ramp, and AWS Enterprise Support customers get access to all 115 Trusted Advisor checks (14 cost optimization, 17 security, 24 fault tolerance, 10 performance, and 50 service limits) and recommendations."

QUESTION 9

A SysOps administrator needs to delete an AWS CloudFormation stack that is no longer in use. The CloudFormation stack is in the DELETE_FAILED state. The SysOps administrator has validated the permissions that are required to delete the Cloud Formation stack.

- A. The configured timeout to delete the stack was too low for the delete operation to complete.
- B. The stack contains nested stacks that must be manually deleted first.
- C. The stack was deployed with the -disable rollback option.
- D. There are additional resources associated with a security group in the stack
- E. There are Amazon S3 buckets that still contain objects in the stack.

Correct Answer: BE

B. One possible cause of the DELETE_FAILED state is that the stack contains nested stacks, and the deletion of the parent stack cannot proceed until the nested stacks are manually deleted first. Nested stacks are separate CloudFormation stacks that are created and managed as part of the resources in the parent stack. When a parent stack is deleted, CloudFormation will attempt to delete the nested stacks, but if there are any issues, the parent stack deletion will fail.

E. Another possible cause of the DELETE_FAILED state is that there are Amazon S3 buckets that still contain objects in the stack. If there are objects (files) present in the S3 buckets that were created as part of the CloudFormation stack, the deletion of the stack will fail. CloudFormation cannot delete the S3 buckets that have objects in them, and those buckets must be emptied or manually deleted before the stack deletion can be completed successfully.

<https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/troubleshooting.html#troubleshooting-errors-delete-stack-fails-~:text=but%20not%20deleted-,Delete%20stack%20fails,-To%20resolve%20this>

QUESTION 10

A company manages its multi-account environment by using AWS Organizations. The company needs to automate the creation of daily incremental backups of any Amazon Elastic Block Store (Amazon EBS) volume that is marked with a Lifecycle: Production tag in one of its primary AWS accounts.

The company wants to prevent users from using Amazon EC2 * permissions to delete any of these production snapshots.

What should a SysOps administrator do to meet these requirements?

- A. Create a daily snapshot of all EBS volumes by using Amazon Data Lifecycle Manager. Specify Lifecycle as the tag key. Specify Production as the tag value.
- B. Associate a service control policy (SCP) with the account to deny users the ability to delete EBS snapshots. Create an Amazon EventBridge rule with a 24-hour cron schedule. Configure EBS Create Snapshot as the target. Target all EBS volumes with the specified tags.
- C. Create a daily snapshot of all EBS volumes by using AWS Backup. Specify Lifecycle as the tag key. Specify Production as the tag value.
- D. Create a daily Amazon Machine Image (AMI) of every production EC2 instance within the AWS account by using Amazon Data Lifecycle Manager.

Correct Answer: A

In this scenario, the objective is to automate the creation of daily incremental backups for EBS volumes marked with a specific tag and prevent users from deleting these snapshots using EC2 permissions. Amazon Data Lifecycle Manager (DLM) is a service that can automate the creation, retention, and deletion of EBS snapshots based on policies. By creating a DLM policy with a daily schedule and configuring it to target EBS volumes with the "Lifecycle: Production" tag, you can achieve the automated backup requirement.

QUESTION 11

A company runs a web application on three Amazon EC2 instances behind an Application Load Balancer (ALB). The company notices that random periods of increased traffic cause a degradation in the application's performance. A SysOps administrator must scale the application to meet the increased traffic.

Which solution meets these requirements?

- A. Create an Amazon CloudWatch alarm to monitor application latency and increase the size of each EC2 instance if the desired threshold is reached.
- B. Create an Amazon EventBridge (Amazon CloudWatch Events) rule to monitor application latency and add an EC2 instance to the ALB if the desired threshold is reached.
- C. Deploy the application to an Auto Scaling group of EC2 instances with a target tracking scaling policy. Attach the ALB to the Auto Scaling group.
- D. Deploy the application to an Auto Scaling group of EC2 instances with a scheduled scaling policy. Attach the ALB to the Auto Scaling group.

Correct Answer: C

QUESTION 12

A company requires that all IAM user accounts that have not been used for 90 days or more must have their access keys and passwords immediately disabled. A SysOps administrator must automate the process of disabling unused keys using the MOST operationally efficient method.

How should the SysOps administrator implement this solution?

- A. Create an AWS Step Functions workflow to identify IAM users that have not been active for 90 days. Run an AWS Lambda function when a scheduled Amazon EventBridge (Amazon CloudWatch Events) rule is invoked to automatically remove the AWS access keys and passwords for these IAM users.
- B. Configure an AWS Config rule to identify IAM users that have not been active for 90 days. Set up an automatic weekly batch process on an Amazon EC2 instance to disable the AWS access keys and passwords for these IAM users.
- C. Develop and run a Python script on an Amazon EC2 instance to programmatically identify IAM users that have not been active for 90 days. Automatically delete these IAM users.
- D. Set up an AWS Config managed rule to identify IAM users that have not been active for 90 days. Set up an AWS Systems Manager automation runbook to disable the AWS access keys for these IAM users.

Correct Answer: D

Checks if your AWS Identity and Access Management (IAM) users have passwords or active access keys that have not been used within the specified number of days you provided. The rule is NON_COMPLIANT if there are inactive accounts not recently used.

QUESTION 13

A company needs to archive all audit logs for 10 years. The company must protect the logs from any future edits.

Which solution will meet these requirements?

- A. Store the data in an Amazon Elastic Block Store (Amazon EBS) volume. Configure AWS Key Management Service (AWS KMS) encryption.
- B. Store the data in an Amazon S3 Glacier vault. Configure a vault lock policy for write- once, read-many (WORM) access.
- C. Store the data in Amazon S3 Standard-Infrequent Access (S3 Standard-IA). Configure server-side encryption.
- D. Store the data in Amazon S3 Standard-Infrequent Access (S3 Standard-IA). Configure multi-factor authentication (MFA).

Correct Answer: B

To meet the requirements of the workload, a company should store the data in an Amazon S3 Glacier vault and configure a vault lock policy for write-once, read-many (WORM) access. This will ensure that the data is stored securely and cannot be edited in the future. The other solutions (storing the data in an Amazon Elastic Block Store (Amazon EBS) volume and configuring AWS Key Management Service (AWS KMS) encryption, storing the data in Amazon S3 Standard-Infrequent Access (S3 Standard-IA) and configuring server-side encryption, or storing the data in Amazon S3 Standard-Infrequent Access (S3 Standard-IA) and configuring multi-factor authentication (MFA)) will not meet the requirements, as they do not provide a way to protect the audit logs from future edits.

https://docs.aws.amazon.com/zh_tw/AmazonS3/latest/userguide/object-lock.html

QUESTION 14

A company asks a SysOps administrator to ensure that AWS CloudTrail files are not tampered with after they are created. Currently, the company uses AWS Identity and Access Management (IAM) to restrict access to specific trails. The company's security team needs the ability to trace the integrity of each file.

What is the MOST operationally efficient solution that meets these requirements?

- A. Create an Amazon EventBridge (Amazon CloudWatch Events) rule that invokes an AWS Lambda function when a new file is delivered. Configure the Lambda function to compute an MD5 hash check on the file and store the result in an Amazon DynamoDB table. The security team can use the values that are stored in DynamoDB to verify the integrity of the delivered files.
- B. Create an AWS Lambda function that is invoked each time a new file is delivered to the CloudTrail bucket. Configure the Lambda function to compute an MD5 hash check on the file and store the result as a tag in an Amazon S3 object. The security team can use the information in the tag to verify the integrity of the delivered files.
- C. Enable the CloudTrail file integrity feature on an Amazon S3 bucket. Create an IAM policy that grants the security team access to the file integrity logs that are stored in the S3 bucket.

D. Enable the CloudTrail file integrity feature on the trail. The security team can use the digest file that is created by CloudTrail to verify the integrity of the delivered files.

Correct Answer: D

<https://docs.aws.amazon.com/awsccloudtrail/latest/userguide/cloudtrail-log-file-validation-intro.html> "When you enable log file integrity validation, CloudTrail creates a hash for every log file that it delivers. Every hour, CloudTrail also creates and delivers a file that references the log files for the last hour and contains a hash of each. This file is called a digest file. Validated log files are invaluable in security and forensic investigations"

QUESTION 15

A company stores critical data in Amazon S3 buckets. A SysOps administrator must build a solution to record all S3 API activity. Which action will meet this requirement?

- A. Configure S3 bucket metrics to record object access logs.
- B. Create an AWS CloudTrail trail to log data events for all S3 objects.
- C. Enable S3 server access logging for each S3 bucket.
- D. Use AWS IAM Access Analyzer for Amazon S3 to store object access logs.

Correct Answer: B

Amazon S3 is integrated with AWS CloudTrail, a service that provides a record of actions taken by a user, role, or an AWS service in Amazon S3. CloudTrail captures a subset of API calls for Amazon S3 as events, including calls from the Amazon S3 console and code calls to the Amazon S3 APIs.

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/cloudtrail-logging.html>

[SOA-C02 Practice Test](#)

[SOA-C02 Study Guide](#)

[SOA-C02 Exam Questions](#)