➢ **Vendor: Amazon**

➢ **Exam Code: SAA-C03**

➢ **Exam Name: AWS Certified Solutions Architect - Associate (SAA-C03)**

➢ **New Updated Questions from Braindump2go**

➢ **(Updated in September/2022)**

Visit Braindump2go and Download Full Version SAA-C03 Exam Dumps

QUESTION 42
A company uses 50 TB of data for reporting. The company wants to move this data from on premises to AWS A custom application in the company's data center runs a weekly data transformation job. The company plans to pause the application until the data transfer is complete and needs to begin the transfer process as soon as possible.
The data center does not have any available network bandwidth for additional workloads.
A solutions architect must transfer the data and must configure the transformation job to continue to run in the AWS Cloud.
Which solution will meet these requirements with the LEAST operational overhead?

A. Use AWS DataSync to move the data.
    Create a custom transformation job by using AWS Glue.
B. Order an AWS Snowcone device to move the data.
    Deploy the transformation application to the device.
C. Order an AWS Snowball Edge Storage Optimized device.
    Copy the data to the device.
    Create a custom transformation job by using AWS Glue.
D. Order an AWS D. Snowball Edge Storage Optimized device that includes Amazon EC2
    compute.
    Copy the data to the device.
    Create a new EC2 instance on AWS to run the transformation application.

Answer: C

QUESTION 43
A company has created an image analysis application in which users can upload photos and add photo frames to their images. The users upload images and metadata to indicate which photo frames they want to add to their images. The application uses a single Amazon EC2 instance and Amazon DynamoDB to store the metadata.
The application is becoming more popular, and the number of users is increasing. The company expects the number of concurrent users to vary significantly depending on the time of day and day of week. The company must ensure that the application can scale to meet the needs of the growing user base.
Which solution meats these requirements?

A. Use AWS Lambda to process the photos.
    Store the photos and metadata in DynamoDB.
B. Use Amazon Kinesis Data Firehose to process the photos and to store the photos and
    metadata.
C. Use AWS Lambda to process the photos.

SAA-C03 Exam Dumps SAA-C03 Exam Questions SAA-C03 PDF Dumps SAA-C03 VCE Dumps

https://www.braindump2go.com/saa-c03.html

Store the photos in Amazon S3.
Retain DynamoDB to store the metadata.

D. Increase the number of EC2 instances to three.
Use Provisioned IOPS SSD (io2) Amazon Elastic Block Store (Amazon EBS) volumes to store
the photos and metadata.

Answer: A

QUESTION 44
A medical records company is hosting an application on Amazon EC2 instances. The application processes customer
data files that are stored on Amazon S3. The EC2 instances are hosted in public subnets. The EC2 instances access
Amazon S3 over the internet, but they do not require any other network access.
A new requirement mandates that the network traffic for file transfers take a private route and not be sent over the
internet.
Which change to the network architecture should a solutions architect recommend to meet this requirement?

A. Create a NAT gateway.
Configure the route table for the public subnets to send traffic to Amazon S3 through the NAT
gateway.
B. Configure the security group for the EC2 instances to restrict outbound traffic so that only traffic
to the S3 prefix list is permitted.
C. Move the EC2 instances to private subnets.
Create a VPC endpoint for Amazon S3, and link the endpoint to the route table for the private
subnets
D. Remove the internet gateway from the VPC.
Set up an AWS Direct Connect connection, and route traffic to Amazon S3 over the Direct
Connect connection.

Answer: C

QUESTION 45
A company uses a popular content management system (CMS) for its corporate website. However, the required
patching and maintenance are burdensome. The company is redesigning its website and wants anew solution. The
website will be updated four times a year and does not need to have any dynamic content available. The solution must
provide high scalability and enhanced security.
Which combination of changes will meet these requirements with the LEAST operational overhead? (Choose two.)

A. Deploy an AWS WAF web ACL in front of the website to provide HTTPS functionality
B. Create and deploy an AWS Lambda function to manage and serve the website content
C. Create the new website and an Amazon S3 bucket Deploy the website on the S3 bucket with
static website hosting enabled
D. Create the new website.
Deploy the website by using an Auto Scaling group of Amazon EC2 instances behind an
Application Load Balancer.

Answer: AD

QUESTION 46
A company stores its application logs in an Amazon CloudWatch Logs log group.
A new policy requires the company to store all application logs in Amazon OpenSearch Service (Amazon Elasticsearch
Service) in near-real time.
Which solution will meet this requirement with the LEAST operational overhead?

A. Configure a CloudWatch Logs subscription to stream the logs to Amazon OpenSearch Service
(Amazon Elasticsearch Service).
B. Create an AWS Lambda function.
Use the log group to invoke the function to write the logs to Amazon OpenSearch Service

(Amazon Elasticsearch Service).

C. Create an Amazon Kinesis Data Firehose delivery stream.
   Configure the log group as the delivery stream's source.
   Configure Amazon OpenSearch Service (Amazon Elasticsearch Service) as the delivery
   stream's destination.
D. Install and configure Amazon Kinesis Agent on each application server to deliver the logs to
   Amazon Kinesis Data Streams.
   Configure Kinesis Data Streams to deliver the logs to Amazon OpenSearch Service (Amazon
   Elasticsearch Service)

Answer: C

QUESTION 47
A company is building a web-based application running on Amazon EC2 instances in multiple Availability Zones. The web application will provide access to a repository of text documents totaling about 900 TB in size. The company anticipates that the web application will experience periods of high demand. A solutions architect must ensure that the storage component for the text documents can scale to meet the demand of the application at all times. The company is concerned about the overall cost of the solution.
Which storage solution meets these requirements MOST cost-effectively?

A. Amazon Elastic Block Store (Amazon EBS)
B. Amazon Elastic File System (Amazon EFS)
C. Amazon Elasticsearch Service (Amazon ES)
D. Amazon S3

Answer: D
Explanation:
Amazon S3 is cheapest and can be accessed from anywhere.

QUESTION 48
A global company is using Amazon API Gateway to design REST APIs for its loyalty club users in the us-east-1 Region and the ap-southeast-2 Region. A solutions architect must design a solution to protect these API Gateway managed REST APIs across multiple accounts from SQL injection and cross-site scripting attacks.
Which solution will meet these requirements with the LEAST amount of administrative effort?

A. Set up AWS WAF in both Regions.
   Associate Regional web ACLs with an API stage.
B. Set up AWS Firewall Manager in both Regions.
   Centrally configure AWS WAF rules.
C. Set up AWS Shield in bath Regions.
   Associate Regional web ACLs with an API stage.
D. Set up AWS Shield in one of the Regions.
   Associate Regional web ACLs with an API stage.

Answer: A

QUESTION 49
A company has implemented a self-managed DNS solution on three Amazon EC2 instances behind a Network Load Balancer (NLB) in the us-west-2 Region. Most of the company's users are located in the United States and Europe. The company wants to improve the performance and availability of the solution. The company launches and configures three EC2 instances in the eu-west-1 Region and adds the EC2 instances as targets for a new NLB.
Which solution can the company use to route traffic to all the EC2 instances?

A. Create an Amazon Route 53 geolocation routing policy to route requests to one of the two
   NLBs.
   Create an Amazon CloudFront distribution.
   Use the Route 53 record as the distribution's origin.

B. Create a standard accelerator in AWS Global Accelerator.
   Create endpoint groups in us-west-2 and eu-west-1.
   Add the two NLBs as endpoints for the endpoint groups.
C. Attach Elastic IP addresses to the six EC2 instances.
   Create an Amazon Route 53 geolocation routing policy to route requests to one of the six EC2 instances.
   Create an Amazon CloudFront distribution.
   Use the Route 53 record as the distribution's origin.
D. Replace the two NLBs with two Application Load Balancers (ALBs).
   Create an Amazon Route 53 latency routing policy to route requests to one of the two ALBs.
   Create an Amazon CloudFront distribution.
   Use the Route 53 record as the distribution's origin.

Answer: A

QUESTION 50
A company is running an online transaction processing (OLTP) workload on AWS. This workload uses an unencrypted Amazon RDS DB instance in a Multi-AZ deployment. Daily database snapshots are taken from this instance.
What should a solutions architect do to ensure the database and snapshots are always encrypted moving forward?

A. Encrypt a copy of the latest DB snapshot.
   Replace existing DB instance by restoring the encrypted snapshot
B. Create a new encrypted Amazon Elastic Block Store (Amazon EBS) volume and copy the snapshots to it.
   Enable encryption on the DB instance.
C. Copy the snapshots and enable encryption using AWS Key Management Service (AWS KMS).
   Restore encrypted snapshot to an existing DB instance.
D. Copy the snapshots to an Amazon S3 bucket that is encrypted using server-side encryption with AWS Key Management Service (AWS KMS) managed keys (SSE-KMS)

Answer: A
Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/USER_RestoreFromSnapshot.html#US
ER_RestoreFromSnapshot.CON
Under "Encrypt unencrypted resources"
https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSEncryption.html

QUESTION 51
A company wants to build a scalable key management Infrastructure to support developers who need to encrypt data in their applications.
What should a solutions architect do to reduce the operational burden?

A. Use multifactor authentication (MFA) to protect the encryption keys.
B. Use AWS Key Management Service (AWS KMS) to protect the encryption keys
C. Use AWS Certificate Manager (ACM) to create, store, and assign the encryption keys
D. Use an IAM policy to limit the scope of users who have access permissions to protect the encryption keys

Answer: B

QUESTION 52
A company has a dynamic web application hosted on two Amazon EC2 instances. The company has its own SSL certificate, which is on each instance to perform SSL termination.
There has been an increase in traffic recently, and the operations team determined that SSL encryption and decryption is causing the compute capacity of the web servers to reach their maximum limit.
What should a solutions architect do to increase the application's performance?

A. Create a new SSL certificate using AWS Certificate Manager (ACM) install the ACM certificate on each instance.
B. Create an Amazon S3 bucket Migrate the SSL certificate to the S3 bucket.
Configure the EC2 instances to reference the bucket for SSL termination.
C. Create another EC2 instance as a proxy server Migrate the SSL certificate to the new instance and configure it to direct connections to the existing EC2 instances
D. Import the SSL certificate into AWS Certificate Manager (ACM).
Create an Application Load Balancer with an HTTPS listener that uses the SSL certificate from ACM.

Answer: D
Explanation:
https://aws.amazon.com/certificate-manager/:
"With AWS Certificate Manager, you can quickly request a certificate, deploy it on ACM-integrated AWS resources, such as Elastic Load Balancers, Amazon CloudFront distributions, and APIs on API Gateway, and let AWS Certificate Manager handle certificate renewals. It also enables you to create private certificates for your internal resources and manage the certificate lifecycle centrally."

QUESTION 53
A company has a highly dynamic batch processing job that uses many Amazon EC2 instances to complete it. The job is stateless in nature, can be started and stopped at any given time with no negative impact, and typically takes upwards of 60 minutes total to complete. The company has asked a solutions architect to design a scalable and cost-effective solution that meets the requirements of the job.
What should the solutions architect recommend?

A. Implement EC2 Spot Instances
B. Purchase EC2 Reserved Instances
C. Implement EC2 On-Demand Instances
D. Implement the processing on AWS Lambda

Answer: A

QUESTION 54
A company runs its two-tier ecommerce website on AWS. The web tier consists of a load balancer that sends traffic to Amazon EC2 instances. The database tier uses an Amazon RDS DB instance. The EC2 instances and the RDS DB instance should not be exposed to the public internet. The EC2 instances require internet access to complete payment processing of orders through a third-party web service. The application must be highly available.
Which combination of configuration options will meet these requirements? (Choose two.)

A. Use an Auto Scaling group to launch the EC2 instances in private subnets.
Deploy an RDS Multi-AZ DB instance in private subnets.
B. Configure a VPC with two private subnets and two NAT gateways across two Availability Zones.
Deploy an Application Load Balancer in the private subnets.
C. Use an Auto Scaling group to launch the EC2 instances in public subnets across two Availability Zones.
Deploy an RDS Multi-AZ DB instance in private subnets.
D. Configure a VPC with one public subnet, one private subnet, and two NAT gateways across two Availability Zones.
Deploy an Application Load Balancer in the public subnet.
E. Configure a VPC with two public subnets, two private subnets, and two NAT gateways across two Availability Zones.
Deploy an Application Load Balancer in the public subnets.

Answer: AE
Explanation:
Before you begin: Decide which two Availability Zones you will use for your EC2 instances. Configure your virtual private cloud (VPC) with at least one public subnet in each of these Availability Zones. These public subnets are used

to configure the load balancer. You can launch your EC2 instances in other subnets of these Availability Zones instead.

QUESTION 55
A solutions architect needs to implement a solution to reduce a company's storage costs. All the company's data is in the Amazon S3 Standard storage class. The company must keep all data for at least 25 years. Data from the most recent 2 years must be highly available and immediately retrievable.
Which solution will meet these requirements?

A. Set up an S3 Lifecycle policy to transition objects to S3 Glacier Deep Archive immediately.
B. Set up an S3 Lifecycle policy to transition objects to S3 Glacier Deep Archive after 2 years.
C. Use S3 Intelligent-Tiering. Activate the archiving option to ensure that data is archived in S3 Glacier Deep Archive.
D. Set up an S3 Lifecycle policy to transition objects to S3 One Zone-Infrequent Access (S3 One Zone-IA) immediately and to S3 Glacier Deep Archive after 2 years.

Answer: B

QUESTION 56
A company runs its ecommerce application on AWS. Every new order is published as a message in a RabbitMQ queue that runs on an Amazon EC2 instance in a single Availability Zone. These messages are processed by a different application that runs on a separate EC2 instance. This application stores the details in a PostgreSQL database on another EC2 instance. All the EC2 instances are in the same Availability Zone.
The company needs to redesign its architecture to provide the highest availability with the least operational overhead.
What should a solutions architect do to meet these requirements?

A. Migrate the queue to a redundant pair (active/standby) of RabbitMQ instances on Amazon MQ.
   Create a Multi-AZ Auto Scaling group (or EC2 instances that host the application.
   Create another Multi-AZ Auto Scaling group for EC2 instances that host the PostgreSQL database.
B. Migrate the queue to a redundant pair (active/standby) of RabbitMQ instances on Amazon MQ.
   Create a Multi-AZ Auto Scaling group for EC2 instances that host the application.
   Migrate the database to run on a Multi-AZ deployment of Amazon RDS for PostgreSQL.
C. Create a Multi-AZ Auto Scaling group for EC2 instances that host the RabbitMQ queue.
   Create another Multi-AZ Auto Scaling group for EC2 instances that host the application.
   Migrate the database to run on a Multi-AZ deployment of Amazon RDS fqjPostgreSQL.
D. Create a Multi-AZ Auto Scaling group for EC2 instances that host the RabbitMQ queue.
   Create another Multi-AZ Auto Scaling group for EC2 instances that host the application.
   Create a third Multi-AZ Auto Scaling group for EC2 instances that host the PostgreSQL database.

Answer: B

QUESTION 57
A reporting team receives files each day in an Amazon S3 bucket. The reporting team manually reviews and copies the files from this initial S3 bucket to an analysis S3 bucket each day at the same time to use with Amazon QuickSight.
Additional teams are starting to send more files in larger sizes to the initial S3 bucket.
The reporting team wants to move the files automatically analysis S3 bucket as the files enter the initial S3 bucket. The reporting team also wants to use AWS Lambda functions to run pattern-matching code on the copied data. In addition, the reporting team wants to send the data files to a pipeline in Amazon SageMaker Pipelines.
What should a solutions architect do to meet these requirements with the LEAST operational overhead?

A. Create a Lambda function to copy the files to the analysis S3 bucket.
   Create an S3 event notification for the analysis S3 bucket.
   Configure Lambda and SageMaker Pipelines as destinations of the event notification.
   Configure s3ObjectCreated:Put as the event type.
B. Create a Lambda function to copy the files to the analysis S3 bucket.
   Configure the analysis S3 bucket to send event notifications to Amazon EventBridge (Amazon

CloudWatch Events).
Configure an ObjectCreated rule in EventBridge (CloudWatch Events).
Configure Lambda and SageMaker Pipelines as targets for the rule.
C. Configure S3 replication between the S3 buckets.
Create an S3 event notification for the analysis S3 bucket.
Configure Lambda and SageMaker Pipelines as destinations of the event notification.
Configure s3ObjectCreated:Put as the event type.
D. Configure S3 replication between the S3 buckets.
Configure the analysis S3 bucket to send event notifications to Amazon EventBridge (Amazon CloudWatch Events).
Configure an ObjectCreated rule in EventBridge (CloudWatch Events).
Configure Lambda and SageMaker Pipelines as targets for the rule.

Answer: A

QUESTION 58
A solutions architect needs to help a company optimize the cost of running an application on AWS. The application will use Amazon EC2 instances, AWS Fargate, and AWS Lambda for compute within the architecture.
The EC2 instances will run the data ingestion layer of the application. EC2 usage will be sporadic and unpredictable. Workloads that run on EC2 instances can be interrupted at any time. The application front end will run on Fargate, and Lambda will serve the API layer. The front-end utilization and API layer utilization will be predictable over the course of the next year.
Which combination of purchasing options will provide the MOST cost-effective solution for hosting this application? (Choose two.)

A. Use Spot Instances for the data ingestion layer
B. Use On-Demand Instances for the data ingestion layer
C. Purchase a 1-year Compute Savings Plan for the front end and API layer.
D. Purchase 1-year All Upfront Reserved instances for the data ingestion layer.
E. Purchase a 1-year EC2 instance Savings Plan for the front end and API layer.

Answer: AC

QUESTION 59
A company runs a web-based portal that provides users with global breaking news, local alerts, and weather updates. The portal delivers each user a personalized view by using mixture of static and dynamic content. Content is served over HTTPS through an API server running on an Amazon EC2 instance behind an Application Load Balancer (ALB). The company wants the portal to provide this content to its users across the world as quickly as possible.
How should a solutions architect design the application to ensure the LEAST amount of latency for all users?

A. Deploy the application stack in a single AWS Region.
   Use Amazon CloudFront to serve all static and dynamic content by specifying the ALB as an origin.
B. Deploy the application stack in two AWS Regions.
   Use an Amazon Route 53 latency routing policy to serve all content from the ALB in the closest Region.
C. Deploy the application stack in a single AWS Region.
   Use Amazon CloudFront to serve the static content.
   Serve the dynamic content directly from the ALB.
D. Deploy the application stack in two AWS Regions.
   Use an Amazon Route 53 geolocation routing policy to serve all content from the ALB in the closest Region.

Answer: A
Explanation:
https://aws.amazon.com/blogs/networking-and-content-delivery/deliver-your-apps-dynamic-content-using-amazon-cloudfront-getting-started-template/

SAA-C03 Exam Dumps **SAA-C03 Exam Questions** **SAA-C03 PDF Dumps** **SAA-C03 VCE Dumps**

https://www.braindump2go.com/saa-c03.html

QUESTION 60
A gaming company is designing a highly available architecture. The application runs on a modified Linux kernel and supports only UDP-based traffic. The company needs the front-end tier to provide the best possible user experience. That tier must have low latency, route traffic to the nearest edge location, and provide static IP addresses for entry into the application endpoints.
What should a solutions architect do to meet these requirements?

A. Configure Amazon Route 53 to forward requests to an Application Load Balancer.
   Use AWS Lambda for the application in AWS Application Auto Scaling.
B. Configure Amazon CloudFront to forward requests to a Network Load Balancer.
   Use AWS Lambda for the application in an AWS Application Auto Scaling group.
C. Configure AWS Global Accelerator to forward requests to a Network Load Balancer.
   Use Amazon EC2 instances for the application in an EC2 Auto Scaling group.
D. Configure Amazon API Gateway to forward requests to an Application Load Balancer.
   Use Amazon EC2 instances for the application in an EC2 Auto Scaling group.

Answer: C

QUESTION 61
A company wants to migrate its existing on-premises monolithic application to AWS. The company wants to keep as much of the front-end code and the backend code as possible. However, the company wants to break the application into smaller applications. A different team will manage each application. The company needs a highly scalable solution that minimizes operational overhead.
Which solution will meet these requirements?

A. Host the application on AWS Lambda Integrate the application with Amazon API Gateway.
B. Host the application with AWS Amplify. Connect the application to an Amazon API Gateway API
   that is integrated with AWS Lambda.
C. Host the application on Amazon EC2 instances. Set up an Application Load Balancer with EC2
   instances in an Auto Scaling group as targets.
D. Host the application on Amazon Elastic Container Service (Amazon ECS) Set up an Application
   Load Balancer with Amazon ECS as the target.

Answer: D
Explanation:
https://aws.amazon.com/blogs/compute/microservice-delivery-with-amazon-ecs-and-application-load-balancers/

QUESTION 62
A company recently started using Amazon Aurora as the data store for its global ecommerce application.
When large reports are run developers report that the ecommerce application is performing poorly After reviewing metrics in Amazon CloudWatch, a solutions architect finds that the ReadlOPS and CPUUtilization metrics are spiking when monthly reports run.
What is the MOST cost-effective solution?

A. Migrate the monthly reporting to Amazon Redshift.
B. Migrate the monthly reporting to an Aurora Replica
C. Migrate the Aurora database to a larger instance class
D. Increase the Provisioned IOPS on the Aurora instance

Answer: B
Explanation:
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Replication.html
#Aurora.Replication.Replicas Aurora Replicas have two main purposes. You can issue queries to them to scale the read operations for your application. You typically do so by connecting to the reader endpoint of the cluster. That way, Aurora can spread the load for read-only connections across as many Aurora Replicas as you have in the cluster. Aurora Replicas also help to increase availability. If the writer instance in a cluster becomes unavailable, Aurora

automatically promotes one of the reader instances to take its place as the new writer.
https://docs.aws.amazon.com/AmazonRDS/latest/AuroraUserGuide/Aurora.Overview.html

QUESTION 63
A company hosts a website analytics application on a single Amazon EC2 On-Demand Instance. The analytics software is written in PHP and uses a MySQL database. The analytics software, the web server that provides PHP, and the database server are all hosted on the EC2 instance. The application is showing signs of performance degradation during busy times and is presenting 5xx errors.
The company needs to make the application scale seamlessly.
Which solution will meet these requirements MOST cost-effectively?

A. Migrate the database to an Amazon RDS for MySQL DB instance.
   Create an AMI of the web application.
   Use the AMI to launch a second EC2 On-Demand Instance.
   Use an Application Load Balancer to distribute the load to each EC2 instance.
B. Migrate the database to an Amazon RDS for MySQL DB instance.
   Create an AMI of the web application.
   Use the AMI to launch a second EC2 On-Demand Instance.
   Use Amazon Route 53 weighted routing to distribute the load across the two EC2 instances.
C. Migrate the database to an Amazon Aurora MySQL DB instance.
   Create an AWS Lambda function to stop the EC2 instance and change the instance type.
   Create an Amazon CloudWatch alarm to invoke the Lambda function when CPU utilization surpasses 75%.
D. Migrate the database to an Amazon Aurora MySQL DB instance.
   Create an AMI of the web application.
   Apply the AMI to a launch template.
   Create an Auto Scaling group with the launch template.
   Configure the launch template to use a Spot Fleet.
   Attach an Application Load Balancer to the Auto Scaling group.

Answer: D

QUESTION 64
A company runs a stateless web application in production on a group of Amazon EC2 On-Demand Instances behind an Application Load Balancer. The application experiences heavy usage during an 8-hour period each business day. Application usage is moderate and steady overnight Application usage is low during weekends.
The company wants to minimize its EC2 costs without affecting the availability of the application.
Which solution will meet these requirements?

A. Use Spot Instances for the entire workload.
B. Use Reserved instances for the baseline level of usage.
   Use Spot Instances for any additional capacity that the application needs.
C. Use On-Demand Instances for the baseline level of usage.
   Use Spot Instances for any additional capacity that the application needs.
D. Use Dedicated Instances for the baseline level of usage.
   Use On-Demand Instances for any additional capacity that the application needs.

Answer: B

QUESTION 65
A company needs to retain application logs files for a critical application for 10 years. The application team regularly accesses logs from the past month for troubleshooting, but logs older than 1 month are rarely accessed. The application generates more than 10 TB of logs per month. Which storage option meets these requirements MOST cost-effectively?

A. Store the logs in Amazon S3.
   Use AWS Backup to move logs more than 1 month old to S3 Glacier Deep Archive.

B. Store the logs in Amazon S3.
   Use S3 Lifecycle policies to move logs more than 1 month old to S3 Glacier Deep Archive.
C. Store the logs in Amazon CloudWatch Logs.
   Use AWS Backup to move logs more then 1 month old to S3 Glacier Deep Archive.
D. Store the logs in Amazon CloudWatch Logs.
   Use Amazon S3 Lifecycle policies to move logs more than 1 month old to S3 Glacier Deep Archive.

Answer: B
Explanation:
You need S3 to be able to archive the logs after one month. Cannot do that with CloudWatch Logs.