



# Free Questions for **SAA-C03** by **braindumpscollection**

Shared by **Estrada** on **22-07-2024**

For More Free Questions and Preparation Resources

**Check the Links on Last Page**

# Question 1

---

## Question Type: MultipleChoice

---

A company has an AWS Direct Connect connection from its on-premises location to an AWS account. The AWS account has 30 different VPCs in the same AWS Region. The VPCs use private virtual interfaces (VIFs). Each VPC has a CIDR block that does not overlap with other networks under the company's control.

The company wants to centrally manage the networking architecture while still allowing each VPC to communicate with all other VPCs and on-premises networks.

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

### Options:

---

- A-** Create a transit gateway and associate the Direct Connect connection with a new transit VIF. Turn on the transit gateway's route propagation feature.
- B-** Create a Direct Connect gateway. Recreate the private VIFs to use the new gateway. Associate each VPC by creating new virtual private gateways.
- C-** Create a transit VPC. Connect the Direct Connect connection to the transit VPC. Create a peering connection between all other VPCs in the Region. Update the route tables.
- D-** Create AWS Site-to-Site VPN connections from on-premises to each VPC. Ensure that both VPN tunnels are UP for each connection. Turn on the route propagation feature.

## Answer:

---

A

## Explanation:

---

This solution meets the following requirements:

It is operationally efficient, as it only requires one transit gateway and one transit VIF to connect the Direct Connect connection to all the VPCs in the same AWS Region. The transit gateway acts as a regional network hub that simplifies the network management and reduces the number of VIFs and gateways needed.

It is scalable, as it can support up to 5000 attachments per transit gateway, which can include VPCs, VPNs, Direct Connect gateways, and peering connections. The transit gateway can also be connected to other transit gateways in different Regions or accounts using peering connections, enabling cross-Region and cross-account connectivity.

It is flexible, as it allows each VPC to communicate with all other VPCs and on-premises networks using dynamic routing protocols such as Border Gateway Protocol (BGP). The transit gateway's route propagation feature automatically propagates the routes from the attached VPCs and VPNs to the transit gateway route table, eliminating the need to manually update the route tables.

[Transit Gateways - Amazon Virtual Private Cloud](#)

[Working with transit gateways - AWS Direct Connect](#)

[Amazon VPC-to-Amazon VPC connectivity options - Amazon Virtual Private Cloud Connectivity Options](#)

## Question 2

---

**Question Type:** MultipleChoice

---

A company has stored 10 TB of log files in Apache Parquet format in an Amazon S3 bucket. The company occasionally needs to use SQL to analyze the log files. Which solution will meet these requirements MOST cost-effectively?

### Options:

---

- A-** Create an Amazon Aurora MySQL database. Migrate the data from the S3 bucket into Aurora by using AWS Database Migration Service (AWS DMS). Issue SQL statements to the Aurora database.
- B-** Create an Amazon Redshift cluster. Use Redshift Spectrum to run SQL statements directly on the data in the S3 bucket.
- C-** Create an AWS Glue crawler to store and retrieve table metadata from the S3 bucket. Use Amazon Athena to run SQL statements directly on the data in the S3 bucket.
- D-** Create an Amazon EMR cluster. Use Apache Spark SQL to run SQL statements directly on the data in the S3 bucket.

### Answer:

---

C

## Explanation:

---

AWS Glue is a serverless data integration service that can crawl, catalog, and prepare data for analysis. AWS Glue can automatically discover the schema and partitioning of the data stored in Apache Parquet format in S3, and create a table in the AWS Glue Data Catalog. Amazon Athena is a serverless interactive query service that can run SQL queries directly on data in S3, without requiring any data loading or transformation. Athena can use the table metadata from the AWS Glue Data Catalog to query the data in S3. By using AWS Glue and Athena, you can analyze the log files in S3 most cost-effectively, as you only pay for the resources consumed by the crawler and the queries, and you do not need to provision or manage any servers or clusters.

[AWS Glue](#)

[Amazon Athena](#)

[Analyzing Data in S3 using Amazon Athena](#)

## Question 3

---

**Question Type:** MultipleChoice

---

A company has Amazon EC2 instances that run nightly batch jobs to process data.

a. The EC2 instances run in an Auto Scaling group that uses On-Demand billing. If a job fails on one instance, another instance will reprocess the job. The batch jobs run between 12:00 AM and 06:00 AM local time every day.

Which solution will provide EC2 instances to meet these requirements MOST cost-effectively'?

### Options:

---

- A-** Purchase a 1-year Savings Plan for Amazon EC2 that covers the instance family of the Auto Scaling group that the batch job uses.
- B-** Purchase a 1-year Reserved Instance for the specific instance type and operating system of the instances in the Auto Scaling group that the batch job uses.
- C-** Create a new launch template for the Auto Scaling group Set the instances to Spot Instances Set a policy to scale out based on CPU usage.
- D-** Create a new launch template for the Auto Scaling group Increase the instance size Set a policy to scale out based on CPU usage.

### Answer:

---

C

### Explanation:

---

This option is the most cost-effective solution because it leverages the Spot Instances, which are unused EC2 instances that are available at up to 90% discount compared to On-Demand prices. Spot Instances can be interrupted by AWS when the demand for On-Demand instances increases, but since the batch jobs are fault-tolerant and can be reprocessed by another instance, this is not a major issue. By using a launch template, the company can specify the configuration of the Spot Instances, such as the instance type, the operating system, and the user data. By using an Auto Scaling group, the company can automatically scale the number of Spot

Instances based on the CPU usage, which reflects the load of the batch jobs. This way, the company can optimize the performance and the cost of the EC2 instances for the nightly batch jobs.

A) Purchase a 1-year Savings Plan for Amazon EC2 that covers the instance family of the Auto Scaling group that the batch job uses. This option is not optimal because it requires a commitment to a consistent amount of compute usage per hour for a one-year term, regardless of the instance type, size, region, or operating system. This can limit the flexibility and scalability of the Auto Scaling group and result in overpaying for unused compute capacity. Moreover, Savings Plans do not provide a capacity reservation, which means the company still needs to reserve capacity with On-Demand Capacity Reservations and pay lower prices with Savings Plans.

B) Purchase a 1-year Reserved Instance for the specific instance type and operating system of the instances in the Auto Scaling group that the batch job uses. This option is not ideal because it requires a commitment to a specific instance configuration for a one-year term, which can reduce the flexibility and scalability of the Auto Scaling group and result in overpaying for unused compute capacity. Moreover, Reserved Instances do not provide a capacity reservation, which means the company still needs to reserve capacity with On-Demand Capacity Reservations and pay lower prices with Reserved Instances.

D) Create a new launch template for the Auto Scaling group Increase the instance size Set a policy to scale out based on CPU usage. This option is not cost-effective because it does not take advantage of the lower prices of Spot Instances. Increasing the instance size can improve the performance of the batch jobs, but it can also increase the cost of the On-Demand instances. Moreover, scaling out based on CPU usage can result in launching more instances than needed, which can also increase the cost of the system.

[1Spot Instances - Amazon Elastic Compute Cloud](#)

[2Launch templates - Amazon Elastic Compute Cloud](#)

[3Auto Scaling groups - Amazon EC2 Auto Scaling](#)

[\[4\] Savings Plans - Amazon EC2 Reserved Instances and Other AWS Reservation Models](#)

## Question 4

---

### Question Type: MultipleChoice

---

A company uses Amazon S3 as its data lake. The company has a new partner that must use SFTP to upload data files. A solutions architect needs to implement a highly available SFTP solution that minimizes operational overhead.

Which solution will meet these requirements?

#### Options:

---

- A-** Use AWS Transfer Family to configure an SFTP-enabled server with a publicly accessible endpoint. Choose the S3 data lake as the destination.
- B-** Use Amazon S3 File Gateway as an SFTP server. Expose the S3 File Gateway endpoint URL to the new partner. Share the S3 File Gateway endpoint with the new partner.
- C-** Launch an Amazon EC2 instance in a private subnet in a VPC. Instruct the new partner to upload files to the EC2 instance by using a VPN. Run a cron job script on the EC2 instance to upload files to the S3 data lake.
- D-** Launch Amazon EC2 instances in a private subnet in a VPC. Place a Network Load Balancer (NLB) in front of the EC2 instances. Create an SFTP listener port for the NLB. Share the NLB hostname with the new partner. Run a cron job script on the EC2 instances to upload files to the S3 data lake.



## Answer:

---

A

## Explanation:

---

This option is the most cost-effective and simple way to enable SFTP access to the S3 data lake. AWS Transfer Family is a fully managed service that supports secure file transfers over SFTP, FTPS, and FTP protocols. You can create an SFTP-enabled server with a public endpoint and associate it with your S3 bucket. You can also use AWS Identity and Access Management (IAM) roles and policies to control access to your S3 data lake. The service scales automatically to handle any volume of file transfers and provides high availability and durability. You do not need to provision, manage, or patch any servers or load balancers.

Option B is not correct because Amazon S3 File Gateway is not an SFTP server. It is a hybrid cloud storage service that provides a local file system interface to S3. You can use it to store and retrieve files as objects in S3 using standard file protocols such as NFS and SMB. However, it does not support SFTP protocol, and it requires deploying a file gateway appliance on-premises or on EC2.

Option C is not cost-effective or scalable because it requires launching and managing an EC2 instance in a private subnet and setting up a VPN connection for the new partner. This would incur additional costs for the EC2 instance, the VPN connection, and the data transfer. It would also introduce complexity and security risks to the solution. Moreover, it would require running a cron job script on the EC2 instance to upload files to the S3 data lake, which is not efficient or reliable.

Option D is not cost-effective or scalable because it requires launching and managing multiple EC2 instances in a private subnet and placing a NLB in front of them. This would incur additional costs for the EC2 instances, the NLB, and the data transfer. It would also introduce complexity and security risks to the solution. Moreover, it would require running a cron job script on the EC2 instances to

upload files to the S3 data lake, which is not efficient or reliable.Reference:

[What Is AWS Transfer Family?](#)

[What Is Amazon S3 File Gateway?](#)

[What Is Amazon EC2?](#)

[[What Is Amazon Virtual Private Cloud?](#)]

[[What Is a Network Load Balancer?](#)]

## Question 5

---

**Question Type: MultipleChoice**

---

A company uses AWS Organizations to run workloads within multiple AWS accounts A tagging policy adds department tags to AWS resources when the company creates tags.

An accounting team needs to determine spending on Amazon EC2 consumption The accounting team must determine which departments are responsible for the costs regardless of AWS account The accounting team has access to AWS Cost Explorer for all AWS accounts within the organization and needs to access all reports from Cost Explorer.

Which solution meets these requirements in the MOST operationally efficient way'?

## Options:

---

- A-** From the Organizations management account billing console, activate a user-defined cost allocation tag named department. Create one cost report in Cost Explorer grouping by tag name, and filter by EC2.
- B-** From the Organizations management account billing console, activate an AWS-defined cost allocation tag named department. Create one cost report in Cost Explorer grouping by tag name, and filter by EC2.
- C-** From the Organizations member account billing console, activate a user-defined cost allocation tag named department. Create one cost report in Cost Explorer grouping by the tag name, and filter by EC2.
- D-** From the Organizations member account billing console, activate an AWS-defined cost allocation tag named department. Create one cost report in Cost Explorer grouping by tag name and filter by EC2.

## Answer:

---

B

## Explanation:

---

This solution meets the following requirements:

It is operationally efficient, as it only requires one activation of the cost allocation tag and one creation of the cost report from the management account, which has access to all the member accounts' data and billing pReference.

It is consistent, as it uses the AWS-defined cost allocation tag named department, which is automatically applied to resources when the company creates tags using the tagging policy enforced by AWS Organizations. This ensures that the tag name and value are the same across all the resources and accounts, and avoids any discrepancies or errors that might arise from user-defined tags.

It is informative, as it creates one cost report in Cost Explorer grouping by the tag name, and filters by EC2. This allows the accounting team to see the breakdown of EC2 consumption and costs by department, regardless of the AWS account. The team can also use other features of Cost Explorer, such as charts, filters, and forecasts, to analyze and optimize the spending.

[Using AWS cost allocation tags - AWS Billing](#)

[User-defined cost allocation tags - AWS Billing](#)

[Cost Tagging and Reporting with AWS Organizations](#)

## Question 6

---

**Question Type:** MultipleChoice

---

A company is developing a mobile game that streams score updates to a backend processor and then posts results on a leaderboard. A solutions architect needs to design a solution that can handle large traffic spikes, process the mobile game updates in order of receipt, and store the processed updates in a highly available database. The company also wants to minimize the management overhead required to maintain the solution.

What should the solutions architect do to meet these requirements?

## Options:

---

- A-** Push score updates to Amazon Kinesis Data Streams Process the updates in Kinesis Data Streams with AWS Lambda Store the processed updates in Amazon DynamoDB.
- B-** Push score updates to Amazon Kinesis Data Streams. Process the updates with a fleet of Amazon EC2 instances set up for Auto Scaling Store the processed updates in Amazon Redshift.
- C-** Push score updates to an Amazon Simple Notification Service (Amazon SNS) topic Subscribe an AWS Lambda function to the SNS topic to process the updates. Store the processed updates in a SQL database running on Amazon EC2.
- D-** Push score updates to an Amazon Simple Queue Service (Amazon SQS) queue. Use a fleet of Amazon EC2 instances with Auto Scaling to process the updates in the SQS queue. Store the processed updates in an Amazon RDS Multi-AZ DB instance.

## Answer:

---

A

## Explanation:

---

Amazon Kinesis Data Streams is a scalable and reliable service that can ingest, buffer, and process streaming data in real-time. It can handle large traffic spikes and preserve the order of the incoming data records. AWS Lambda is a serverless compute service that can process the data streams from Kinesis Data Streams without requiring any infrastructure management. It can also scale automatically to match the throughput of the data stream. Amazon DynamoDB is a fully managed, highly available, and fast NoSQL database that can store the processed updates from Lambda. It can also handle high write throughput and provide consistent performance. By using these

services, the solutions architect can design a solution that meets the requirements of the company with the least operational overhead.

## Question 7

---

**Question Type:** MultipleChoice

---

A company runs analytics software on Amazon EC2 instances. The software accepts job requests from users to process data that has been uploaded to Amazon S3. Users report that some submitted data is not being processed. Amazon CloudWatch reveals that the EC2 instances have a consistent CPU utilization at or near 100%. The company wants to improve system performance and scale the system based on user load.

What should a solutions architect do to meet these requirements?

### Options:

---

- A-** Create a copy of the instance. Place all instances behind an Application Load Balancer.
- B-** Create an S3 VPC endpoint for Amazon S3. Update the software to reference the endpoint.
- C-** Stop the EC2 instances. Modify the instance type to one with a more powerful CPU and more memory. Restart the instances.
- D-** Route incoming requests to Amazon Simple Queue Service (Amazon SQS). Configure an EC2 Auto Scaling group based on queue size. Update the software to read from the queue.

## Answer:

---

D

## Explanation:

---

This option is the best solution because it allows the company to decouple the analytics software from the user requests and scale the EC2 instances dynamically based on the demand. By using Amazon SQS, the company can create a queue that stores the user requests and acts as a buffer between the users and the analytics software. This way, the software can process the requests at its own pace without losing any data or overloading the EC2 instances. By using EC2 Auto Scaling, the company can create an Auto Scaling group that launches or terminates EC2 instances automatically based on the size of the queue. This way, the company can ensure that there are enough instances to handle the load and optimize the cost and performance of the system. By updating the software to read from the queue, the company can enable the analytics software to consume the requests from the queue and process the data from Amazon S3.

A) Create a copy of the instance Place all instances behind an Application Load Balancer. This option is not optimal because it does not address the root cause of the problem, which is the high CPU utilization of the EC2 instances. An Application Load Balancer can distribute the incoming traffic across multiple instances, but it cannot scale the instances based on the load or reduce the processing time of the analytics software. Moreover, this option can incur additional costs for the load balancer and the extra instances.

B) Create an S3 VPC endpoint for Amazon S3 Update the software to reference the endpoint. This option is not effective because it does not solve the issue of the high CPU utilization of the EC2 instances. An S3 VPC endpoint can enable the EC2 instances to access Amazon S3 without going through the internet, which can improve the network performance and security. However, it cannot reduce the processing time of the analytics software or scale the instances based on the load.

C) Stop the EC2 instances. Modify the instance type to one with a more powerful CPU and more memory. Restart the instances. This option is not scalable because it does not account for the variability of the user load. Changing the instance type to a more powerful one can improve the performance of the analytics software, but it cannot adjust the number of instances based on the demand. Moreover, this option can increase the cost of the system and cause downtime during the instance modification.

[1Using Amazon SQS queues with Amazon EC2 Auto Scaling - Amazon EC2 Auto Scaling](#)

[2Tutorial: Set up a scaled and load-balanced application - Amazon EC2 Auto Scaling](#)

[3Amazon EC2 Auto Scaling FAQs](#)

## Question 8

---

**Question Type: MultipleChoice**

---

A company needs to create an AWS Lambda function that will run in a VPC in the company's primary AWS account. The Lambda function needs to access files that the company stores

in an Amazon Elastic File System (Amazon EFS) file system. The EFS file system is located in a secondary AWS account. As the company adds files to the file system the solution must scale to meet the demand.

Which solution will meet these requirements MOST cost-effectively?



## Options:

---

- A-** Create a new EFS file system in the primary account Use AWS DataSync to copy the contents of the original EFS file system to the new EFS file system
- B-** Create a VPC peering connection between the VPCs that are in the primary account and the secondary account
- C-** Create a second Lambda function In the secondary account that has a mount that is configured for the file system. Use the primary account's Lambda function to invoke the secondary account's Lambda function
- D-** Move the contents of the file system to a Lambda Layer's Configure the Lambda layer's permissions to allow the company's secondary account to use the Lambda layer.

## Answer:

---

B

## Explanation:

---

This option is the most cost-effective and scalable way to allow the Lambda function in the primary account to access the EFS file system in the secondary account. VPC peering enables private connectivity between two VPCs without requiring gateways, VPN connections, or dedicated network connections. The Lambda function can use the VPC peering connection to mount the EFS file system as a local file system and access the files as needed. The solution does not incur additional data transfer or storage costs, and it leverages the existing EFS file system without duplicating or moving the data.

Option A is not cost-effective because it requires creating a new EFS file system and using AWS DataSync to copy the data from the original EFS file system. This would incur additional storage and data transfer costs, and it would not provide real-time access to the files.

Option C is not scalable because it requires creating a second Lambda function in the secondary account and configuring cross-account permissions to invoke it from the primary account. This would add complexity and latency to the solution, and it would increase the Lambda invocation costs.

Option D is not feasible because Lambda layers are not designed to store large amounts of data or provide file system access. Lambda layers are used to share common code or libraries across multiple Lambda functions. Moving the contents of the EFS file system to a Lambda layer would exceed the size limit of 250 MB for a layer, and it would not allow the Lambda function to read or write files to the layer. Reference:

[What Is VPC Peering?](#)

[Using Amazon EFS file systems with AWS Lambda](#)

[What Are Lambda Layers?](#)

## Question 9

---

**Question Type: MultipleChoice**

---

A company needs to extract the names of ingredients from recipe records that are stored as text files in an Amazon S3 bucket. A web application will use the ingredient names to query an Amazon DynamoDB table and determine a nutrition score.

The application can handle non-food records and errors. The company does not have any employees who have machine learning knowledge to develop this solution.

Which solution will meet these requirements MOST cost-effectively?

### Options:

---

- A-** Use S3 Event Notifications to invoke an AWS Lambda function when PutObject requests occur Program the Lambda function to analyze the object and extract the ingredient names by using Amazon Comprehend Store the Amazon Comprehend output in the DynamoDB table.
- B-** Use an Amazon EventBridge rule to invoke an AWS Lambda function when PutObject requests occur. Program the Lambda function to analyze the object by using Amazon Forecast to extract the ingredient names Store the Forecast output in the DynamoDB table.
- C-** Use S3 Event Notifications to invoke an AWS Lambda function when PutObject requests occur Use Amazon Polly to create audio recordings of the recipe records. Save the audio files in the S3 bucket Use Amazon Simple Notification Service (Amazon SNS) to send a URL as a message to employees Instruct the employees to listen to the audio files and calculate the nutrition score Store the ingredient names in the DynamoDB table.
- D-** Use an Amazon EventBridge rule to invoke an AWS Lambda function when a PutObject request occurs Program the Lambda function to analyze the object and extract the ingredient names by using Amazon SageMaker Store the inference output from the SageMaker endpoint in the DynamoDB table.

### Answer:

---

A

### Explanation:

---

This solution meets the following requirements:

It is cost-effective, as it only uses serverless components that are charged based on usage and do not require any upfront provisioning or maintenance.

It is scalable, as it can handle any number of recipe records that are uploaded to the S3 bucket without any performance degradation or manual intervention.

It is easy to implement, as it does not require any machine learning knowledge or complex data processing logic. Amazon Comprehend is a natural language processing service that can automatically extract entities such as ingredients from text files. The Lambda function can simply invoke the Comprehend API and store the results in the DynamoDB table.

It is reliable, as it can handle non-food records and errors gracefully. Amazon Comprehend can detect the language and domain of the text files and return an appropriate response. The Lambda function can also implement error handling and logging mechanisms to ensure the data quality and integrity.

[Using AWS Lambda with Amazon S3 - AWS Lambda](#)

[What Is Amazon Comprehend? - Amazon Comprehend](#)

[Working with Tables - Amazon DynamoDB](#)

## Question 10

---

**Question Type:** MultipleChoice

---

A solutions architect needs to design the architecture for an application that a vendor provides as a Docker container image. The container needs 50 GB of storage available for temporary files. The infrastructure must be serverless.

Which solution meets these requirements with the LEAST operational overhead?

### Options:

---

- A-** Create an AWS Lambda function that uses the Docker container image with an Amazon S3 mounted volume that has more than 50 GB of space.
- B-** Create an AWS Lambda function that uses the Docker container image with an Amazon Elastic Block Store (Amazon EBS) volume that has more than 50 GB of space.
- C-** Create an Amazon Elastic Container Service (Amazon ECS) cluster that uses the AWS Fargate launch type. Create a task definition for the container image with an Amazon Elastic File System (Amazon EFS) volume. Create a service with that task definition.
- D-** Create an Amazon Elastic Container Service (Amazon ECS) cluster that uses the Amazon EC2 launch type with an Amazon Elastic Block Store (Amazon EBS) volume that has more than 50 GB of space. Create a task definition for the container image. Create a service with that task definition.

### Answer:

---

C

### Explanation:

---

The AWS Fargate launch type is a serverless way to run containers on Amazon ECS, without having to manage any underlying infrastructure. You only pay for the resources required to run your containers, and AWS handles the provisioning, scaling, and security of the cluster. Amazon EFS is a fully managed, elastic, and scalable file system that can be mounted to multiple containers, and provides high availability and durability. By using AWS Fargate and Amazon EFS, you can run your Docker container image with 50 GB of storage available for temporary files, with the least operational overhead. This solution meets the requirements of the question.

[AWS Fargate](#)

[Amazon Elastic File System](#)

[Using Amazon EFS file systems with Amazon ECS](#)

## Question 11

---

**Question Type:** MultipleChoice

---

A company uses Amazon S3 to store high-resolution pictures in an S3 bucket. To minimize application changes, the company stores the pictures as the latest version of an S3 object

The company needs to retain only the two most recent versions of the pictures.

The company wants to reduce costs. The company has identified the S3 bucket as a large expense.

Which solution will reduce the S3 costs with the LEAST operational overhead?

## Options:

---

- A- Use S3 Lifecycle to delete expired object versions and retain the two most recent versions.
- B- Use an AWS Lambda function to check for older versions and delete all but the two most recent versions
- C- Use S3 Batch Operations to delete noncurrent object versions and retain only the two most recent versions
- D- Deactivate versioning on the S3 bucket and retain the two most recent versions.

## Answer:

---

A

## Explanation:

---

S3 Lifecycle is a feature that allows you to automate the management of your S3 objects based on predefined rules. You can use S3 Lifecycle to delete expired object versions and retain the two most recent versions by creating a lifecycle configuration rule that applies to all objects in the bucket and specifies the expiration action for noncurrent versions. This way, you can reduce the storage costs of your S3 bucket without requiring any application changes or manual intervention. S3 Lifecycle runs once a day and marks the eligible object versions for deletion. You are no longer charged for the objects that are marked for deletion. S3 Lifecycle is the most cost-effective and simple solution among the options.

B) Use an AWS Lambda function to check for older versions and delete all but the two most recent versions. This option is not optimal because it requires you to write, test, and maintain a custom Lambda function that scans the S3 bucket for older versions and deletes

them. This can incur additional costs for Lambda invocations and increase the operational complexity and overhead. Moreover, you need to ensure that your Lambda function has the appropriate permissions and error handling mechanisms to perform the deletion operation.

C) Use S3 Batch Operations to delete noncurrent object versions and retain only the two most recent versions. This option is not ideal because S3 Batch Operations is designed for performing large-scale operations on S3 objects, such as copying, tagging, restoring, or invoking a Lambda function. To use S3 Batch Operations to delete noncurrent object versions, you need to provide a manifest file that lists the object versions that you want to delete. This can be challenging and time-consuming to generate and update. Moreover, S3 Batch Operations charges you for each operation that you perform, which can increase your costs.

D) Deactivate versioning on the S3 bucket and retain the two most recent versions. This option is not feasible because deactivating versioning on an S3 bucket does not delete the existing object versions. Instead, it prevents new versions from being created. Therefore, you still need to delete the older versions manually or use another method to do so. Additionally, deactivating versioning can compromise the data protection and recovery capabilities of your S3 bucket.

[1Considering four different replication options for data in Amazon S3 | AWS Storage Blog](#)

[2Using AWS Lambda with Amazon S3 batch operations - AWS Lambda](#)

[3Empty an Amazon S3 bucket with a lifecycle configuration rule](#)

[4Amazon S3 - Lifecycle Management - GeeksforGeeks](#)

## Question 12

---



**Question Type: MultipleChoice**

---

A company hosts an application used to upload files to an Amazon S3 bucket. Once uploaded, the files are processed to extract metadata which takes less than 5 seconds. The volume and frequency of the uploads varies from a few files each hour to hundreds of concurrent uploads. The company has asked a solutions architect to design a cost-effective architecture that will meet these requirements.

What should the solutions architect recommend?

**Options:**

---

- A-** Configure AWS CloudTrail trails to log S3 API calls. Use AWS AppSync to process the files.
- B-** Configure an object-created event notification within the S3 bucket to invoke an AWS Lambda function to process the files.
- C-** Configure Amazon Kinesis Data Streams to process and send data to Amazon S3. Invoke an AWS Lambda function to process the files.
- D-** Configure an Amazon Simple Notification Service (Amazon SNS) topic to process the files uploaded to Amazon S3. Invoke an AWS Lambda function to process the files.

**Answer:**

---

B

**Explanation:**

---

This option is the most cost-effective and scalable way to process the files uploaded to S3. AWS CloudTrail is used to log API calls, not to trigger actions based on them. AWS AppSync is a service for building GraphQL APIs, not for processing files. Amazon Kinesis Data Streams is used to ingest and process streaming data, not to send data to S3. Amazon SNS is a pub/sub service that can be used to notify subscribers of events, not to process files. Reference:

[Using AWS Lambda with Amazon S3](#)

[AWS CloudTrail FAQs](#)

[What Is AWS AppSync?](#)

[\[What Is Amazon Kinesis Data Streams?\]](#)

[\[What Is Amazon Simple Notification Service?\]](#)

**To Get Premium Files for SAA-C03 Visit**

<https://www.p2pexams.com/products/saa-c03>

**For More Free Questions Visit**

<https://www.p2pexams.com/amazon/pdf/saa-c03>

