

- **Vendor: Amazon**
- **Exam Code: SAA-C03**
- **Exam Name: AWS Certified Solutions Architect - Associate (SAA-C03)**
- **New Updated Questions from [Braindump2go](https://www.braindump2go.com)**
- **(Updated in [September/2022](https://www.braindump2go.com))**

Visit Braindump2go and Download Full Version SAA-C03 Exam Dumps

QUESTION 91

A company runs a global web application on Amazon EC2 instances behind an Application Load Balancer.

The application stores data in Amazon Aurora.

The company needs to create a disaster recovery solution and can tolerate up to 30 minutes of downtime and potential data loss.

The solution does not need to handle the load when the primary infrastructure is healthy.

What should a solutions architect do to meet these requirements?

- A. Deploy the application with the required infrastructure elements in place.
Use Amazon Route 53 to configure active-passive failover.
Create an Aurora Replica in a second AWS Region.
- B. Host a scaled-down deployment of the application in a second AWS Region.
Use Amazon Route 53 to configure active-active failover.
Create an Aurora Replica in the second Region.
- C. Replicate the primary infrastructure in a second AWS Region.
Use Amazon Route 53 to configure active-active failover.
Create an Aurora database that is restored from the latest snapshot.
- D. Back up data with AWS Backup.
Use the backup to create the required infrastructure in a second AWS Region.
Use Amazon Route 53 to configure active-passive failover.
Create an Aurora second primary instance in the second Region.

Answer: C

QUESTION 92

A company wants to measure the effectiveness of its recent marketing campaigns.

The company performs batch processing on csv files of sales data and stores the results in an Amazon S3 bucket once every hour.

The S3 bucket contains billions of objects. The company runs one-time queries in Amazon Athena to determine which products are most popular on a particular date for a particular region. Queries sometimes fail or take longer than expected to finish.

Which actions should a solutions architect take to improve the query performance and reliability? (Select TWO.)

- A. Reduce the S3 object sizes to less than 126 MB
- B. Partition the data by date and region in Amazon S3
- C. Store the files as large, single objects in Amazon S3.
- D. Use Amazon Kinesis Data Analytics to run the Queries as part of the batch processing operation
- E. Use an AWS Glue extract, transform, and load (ETL) process to convert the csv files into Apache

[SAA-C03 Exam Dumps](#) [SAA-C03 Exam Questions](#) [SAA-C03 PDF Dumps](#) [SAA-C03 VCE Dumps](#)

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

Parquet format.

Answer: CE

QUESTION 93

A company is running several business applications in three separate VPCs within the us-east-1 Region.

The applications must be able to communicate between VPCs.

The applications also must be able to consistently send hundreds to gigabytes of data each day to a latency-sensitive application that runs in a single on-premises data center.

A solutions architect needs to design a network connectivity solution that maximizes cost-effectiveness.

Which solution meets those requirements?

- A. Configure three AWS Site-to-Site VPN connections from the data center to AWS.
Establish connectivity by configuring one VPN connection for each VPC.
- B. Launch a third-party virtual network appliance in each VPC.
Establish an IPsec VPN tunnel between the Data center and each virtual appliance.
- C. Set up three AWS Direct Connect connections from the data center to a Direct Connect gateway in us-east-1.
Establish connectivity by configuring each VPC to use one of the Direct Connect connections.
- D. Set up one AWS Direct Connect connection from the data center to AWS.
Create a transit gateway, and attach each VPC to the transit gateway.
Establish connectivity between the Direct Connect connection and the transit gateway.

Answer: D

Explanation:

<https://docs.aws.amazon.com/whitepapers/latest/aws-vpc-connectivity-options/aws-direct-connect-aws-transit-gateway.html>

QUESTION 94

An online photo application lets users upload photos and perform image editing operations. The application offers two classes of service: free and paid. Photos submitted by paid users are processed before those submitted by free users. Photos are uploaded to Amazon S3 and the job information is sent to Amazon SQS.

Which configuration should a solutions architect recommend?

- A. Use one SQS FIFO queue.
Assign a higher priority to the paid photos so they are processed first.
- B. Use two SQS FIFO queues: one for paid and one for free.
Set the free queue to use short polling and the paid queue to use long polling.
- C. Use two SQS standard queues: one for paid and one for free.
Configure Amazon EC2 instances to prioritize polling for the paid queue over the free queue.
- D. Use one SQS standard queue.
Set the visibility timeout of the paid photos to zero.
Configure Amazon EC2 instances to prioritize visibility settings so paid photos are processed first.

Answer: C

Explanation:

Priority: Use separate queues to provide prioritization of work.

QUESTION 95

A company hosts its product information webpages on AWS. The existing solution uses multiple Amazon EC2 instances behind an Application Load Balancer in an Auto Scaling group. The website also uses a custom DNS name and communicates with HTTPS only using a dedicated SSL certificate. The company is planning a new product launch and wants to be sure that users from around the world have the best possible experience on the new website.

What should a solutions architect do to meet these requirements?

- A. Redesign the application to use Amazon CloudFront

[SAA-C03 Exam Dumps](#) [SAA-C03 Exam Questions](#) [SAA-C03 PDF Dumps](#) [SAA-C03 VCE Dumps](#)

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

- B. Redesign the application to use AWS Elastic Beanstalk
- C. Redesign the application to use a Network Load Balancer.
- D. Redesign the application to use Amazon S3 static website hosting

Answer: A

Explanation:

as CloudFront can help provide the best experience for global users. CloudFront integrates seamlessly with ALB and provides an option to use custom DNS and SSL certs.

QUESTION 96

A company has 150 TB of archived image data stored on-premises that needs to be moved to the AWS Cloud within the next month.

The company's current network connection allows up to 100 Mbps uploads for this purpose during the night only.

What is the MOST cost-effective mechanism to move this data and meet the migration deadline?

- A. Use AWS Snowmobile to ship the data to AWS.
- B. Order multiple AWS Snowball devices to ship the data to AWS.
- C. Enable Amazon S3 Transfer Acceleration and securely upload the data.
- D. Create an Amazon S3 VPC endpoint and establish a VPN to upload the data

Answer: B

Explanation:

eg. 6 hrs night

$6 \text{ hrs} \times 60 \text{ min/hr} = 360 \text{ min}$

$360 \text{ min} \times 60 \text{ sec/min} = 21600 \text{ sec}$

$100 \text{ Mbps} \times 21600 \text{ s} = 2160000 \text{ Mb}$

or 2160 Gb or 2.1 TB can only be done

So, for 150 TB, we can use 2 X Snowball Edge Storage Optimised devices.

Size of Snowball Edge Storage Optimised device = 80 TB
Size of Snowball Edge Compute Optimised device = 40 TB

Size of Snowcone = 8 TB

Size of Snowmobile = 100 PB (1 PB = 1000 TB)

Q: How should I choose between Snowmobile and Snowball?

To migrate large datasets of 10PB or more in a single location, you should use Snowmobile. For datasets less than 10PB or distributed in multiple locations, you should use Snowball. In addition, you should evaluate the amount of available bandwidth in your network backbone. If you have a high speed backbone with hundreds of Gb/s of spare throughput, then you can use Snowmobile to migrate the large datasets all at once. If you have limited bandwidth on your backbone, you should consider using multiple Snowballs to migrate the data incrementally.

QUESTION 97

A company hosts its web application on AWS using seven Amazon EC2 instances. The company requires that the IP addresses of all healthy EC2 instances be returned in response to DNS queries. Which policy should be used to meet this requirement?

- A. Simple routing policy
- B. Latency routing policy
- C. Multivalue routing policy
- D. Geolocation routing policy

Answer: C

Explanation:

<https://aws.amazon.com/premiumsupport/knowledge-center/multivalue-versus-simple-policies/>

"Use a multivalue answer routing policy to help distribute DNS responses across multiple resources.

For example, use multivalue answer routing when you want to associate your routing records with a Route 53 health check."

<https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/routing-policy.html#routing-policy-multivalue>

QUESTION 98

[SAA-C03 Exam Dumps](#) [SAA-C03 Exam Questions](#) [SAA-C03 PDF Dumps](#) [SAA-C03 VCE Dumps](#)

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

A company wants to use AWS Systems Manager to manage a fleet of Amazon EC2 instances. According to the company's security requirements, no EC2 instances can have internet access. A solutions architect needs to design network connectivity from the EC2 instances to Systems Manager while fulfilling this security obligation.

Which solution will meet these requirements?

- A. Deploy the EC2 instances into a private subnet with no route to the internet.
- B. Configure an interface VPC endpoint for Systems Manager.
Update routes to use the endpoint.
- C. Deploy a NAT gateway into a public subnet.
Configure private subnets with a default route to the NAT gateway.
- D. Deploy an internet gateway.
Configure a network ACL to deny traffic to all destinations except Systems Manager.

Answer: B

Explanation:

VPC Peering connections

VPC interface endpoints can be accessed through both intra-Region and inter-Region VPC peering connections.

VPC Gateway Endpoint connections can't be extended out of a VPC. Resources on the other side of a VPC peering connection in your VPC can't use the gateway endpoint to communicate with resources in the gateway endpoint service.

Reference: <https://docs.aws.amazon.com/systems-manager/latest/userguide/setup-create-vpc.html>

QUESTION 99

A company needs to build a reporting solution on AWS. The solution must support SQL queries that data analysts run on the data.

The data analysts will run lower than 10 total queries each day. The company generates 3 GB of new data daily in an on-premises relational database. This data needs to be transferred to AWS to perform reporting tasks.

What should a solutions architect recommend to meet these requirements at the LOWEST cost?

- A. Use AWS Database Migration Service (AWS DMS) to replicate the data from the on-premises database into Amazon S3.
Use Amazon Athena to query the data.
- B. Use an Amazon Kinesis Data Firehose delivery stream to deliver the data into an Amazon Elasticsearch Service (Amazon ES) cluster. Run the queries in Amazon ES.
- C. Export a daily copy of the data from the on-premises database.
Use an AWS Storage Gateway file gateway to store and copy the export into Amazon S3.
Use an Amazon EMR cluster to query the data.
- D. Use AWS Database Migration Service (AWS DMS) to replicate the data from the on-premises database and load it into an Amazon Redshift cluster.
Use the Amazon Redshift cluster to query the data.

Answer: D

Explanation:

https://docs.aws.amazon.com/dms/latest/userguide/CHAP_Target.Redshift.html

AWS DMS cannot migrate or replicate changes to a schema with a name that begins with underscore (_). If you have schemas that have a name that begins with an underscore, use mapping transformations to rename the schema on the target.

Amazon Redshift doesn't support VARCHARs larger than 64 KB. LOBs from traditional databases can't be stored in Amazon Redshift.

Applying a DELETE statement to a table with a multi-column primary key is not supported when any of the primary key column names use a reserved word. Go here to see a list of Amazon Redshift reserved words.

You may experience performance issues if your source system performs UPDATE operations on the primary key of a source table. These performance issues occur when applying changes to the target. This is because UPDATE (and DELETE) operations depend on the primary key value to identify the target row. If you update the primary key of a source table, your task log will contain messages like the following:

Update on table 1 changes PK to a PK that was previously updated in the same bulk update.

DMS doesn't support custom DNS names when configuring an endpoint for a Redshift cluster, and you need to use the Amazon provided DNS name. Since the Amazon Redshift cluster must be in the same AWS account and Region as the replication instance, validation fails if you use a custom DNS endpoint.

QUESTION 100

A company wants to monitor its AWS costs for financial review. The cloud operations team is designing an architecture in the AWS Organizations management account to query AWS Cost and Usage Reports for all member accounts.

The team must run this query once a month and provide a detailed analysis of the bill.

Which solution is the MOST scalable and cost-effective way to meet these requirements?

- A. Enable Cost and Usage Reports in the management account.
Deliver reports to Amazon Kinesis.
Use Amazon EMR for analysis.
- B. Enable Cost and Usage Reports in the management account.
Deliver the reports to Amazon S3.
Use Amazon Athena for analysis.
- C. Enable Cost and Usage Reports for member accounts.
Deliver the reports to Amazon S3.
Use Amazon Redshift for analysis.
- D. Enable Cost and Usage Reports for member accounts.
Deliver the reports to Amazon Kinesis.
Use Amazon QuickSight for analysis.

Answer: C

Explanation:

<https://docs.aws.amazon.com/cur/latest/userguide/what-is-cur.html>

If you are an administrator of an AWS Organizations management account and do not want any of the member accounts in your Organization to set-up a CUR you can do one of the following:

(Recommended) If you've opted into Organizations with all features enabled, you can apply a Service Control Policy (SCP). Note that SCPs only apply to member accounts and if you want to restrict any IAM users associated with the management account from setting up a CUR, you'll need to adjust their specific IAM permissions. SCPs also are not retroactive, so they will not de-activate any CURs a member account may have set-up prior to the SCP being applied. Submit a customer support case to block access to billing data in the Billing console for member accounts. This is a list of organizations where the payer account prevents member accounts in its organization from viewing billing data on the Bills and Invoices pages. This also prevents those accounts from setting up Cost and Usage Reports. This option is only available for organizations without all features enabled. Please note that if you have already opted into this to prevent member accounts from viewing bills and invoices in the Billing Console, you do not need to request this access again. Those same member accounts will also be prevented from setting up a Cost and Usage Report.

QUESTION 101

A company collects temperature, humidity, and atmospheric pressure data in cities across multiple continents. The average volume of data collected per site each day is 500 GB. Each site has a high-speed internet connection. The company's weather forecasting applications are based in a single Region and analyze the data daily.

What is the FASTEST way to aggregate data from all of these global sites?

- A. Enable Amazon S3 Transfer Acceleration on the destination bucket.
Use multipart uploads to directly upload site data to the destination bucket.
- B. Upload site data to an Amazon S3 bucket in the closest AWS Region.
Use S3 cross-Region replication to copy objects to the destination bucket.
- C. Schedule AWS Snowball jobs daily to transfer data to the closest AWS Region.
Use S3 cross-Region replication to copy objects to the destination bucket.
- D. Upload the data to an Amazon EC2 instance in the closest Region.
Store the data in an Amazon Elastic Block Store (Amazon EBS) volume.
Once a day take an EBS snapshot and copy it to the centralized Region.
Restore the EBS volume in the centralized Region and run an analysis on the data daily.

Answer: A

[SAA-C03 Exam Dumps](#) [SAA-C03 Exam Questions](#) [SAA-C03 PDF Dumps](#) [SAA-C03 VCE Dumps](#)

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

Explanation:

You might want to use Transfer Acceleration on a bucket for various reasons, including the following:

- You have customers that upload to a centralized bucket from all over the world.
- You transfer gigabytes to terabytes of data on a regular basis across continents.
- You are unable to utilize all of your available bandwidth over the Internet when uploading to Amazon S3.

<https://docs.aws.amazon.com/AmazonS3/latest/dev/transfer-acceleration.html>

[https://aws.amazon.com/s3/transfer-](https://aws.amazon.com/s3/transfer-acceleration/#:~:text=S3%20Transfer%20Acceleration%20(S3TA)%20reduces,to%20S3%20for%20remote%20applications)

[acceleration/#:~:text=S3%20Transfer%20Acceleration%20\(S3TA\)%20reduces,to%20S3%20for%20remote%20applications](https://aws.amazon.com/s3/transfer-acceleration/#:~:text=S3%20Transfer%20Acceleration%20(S3TA)%20reduces,to%20S3%20for%20remote%20applications)

"Amazon S3 Transfer Acceleration can speed up content transfers to and from Amazon S3 by as much as 50-500% for long-distance transfer of larger objects. Customers who have either web or mobile applications with widespread users or applications hosted far away from their S3 bucket can experience long and variable upload and download speeds over the Internet"

<https://docs.aws.amazon.com/AmazonS3/latest/userguide/mpuoverview.html>

"Improved throughput -You can upload parts in parallel to improve throughput."

QUESTION 102

A company needs the ability to analyze the log files of its proprietary application. The logs are stored in JSON format in an Amazon S3 bucket. Queries will be simple and will run on-demand.

A solutions architect needs to perform the analysis with minimal changes to the existing architecture.

What should the solutions architect do to meet these requirements with the LEAST amount of operational overhead?

- A. Use Amazon Redshift to load all the content into one place and run the SQL queries as needed
- B. Use Amazon CloudWatch Logs to store the logs
Run SQL queries as needed from the Amazon CloudWatch console
- C. Use Amazon Athena directly with Amazon S3 to run the queries as needed
- D. Use AWS Glue to catalog the logs
Use a transient Apache Spark cluster on Amazon EMR to run the SQL queries as needed

Answer: C

Explanation:

Amazon Athena can be used to query JSON in S3.

QUESTION 103

A company uses AWS Organizations to manage multiple AWS accounts for different departments. The management account has an Amazon S3 bucket that contains project reports.

The company wants to limit access to this S3 bucket to only users of accounts within the organization in AWS Organizations.

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

- A. Add the `aws:PrincipalOrgID` global condition key with a reference to the organization ID to the S3 bucket policy.
- B. Create an organizational unit (OU) for each department.
Add the `aws:PrincipalOrgPaths` global condition key to the S3 bucket policy.
- C. Use AWS CloudTrail to monitor the `CreateAccount`, `InviteAccountToOrganization`, `LeaveOrganization`, and `RemoveAccountFromOrganization` events.
Update the S3 bucket policy accordingly.
- D. Tag each user that needs access to the S3 bucket.
Add the `aws:PrincipalTag` global condition key to the S3 bucket policy.

Answer: A

Explanation:

<https://aws.amazon.com/blogs/security/control-access-to-aws-resources-by-using-the-aws-organization-of-iam-principals/>

The `aws:PrincipalOrgID` global key provides an alternative to listing all the account IDs for all AWS accounts in an organization.

For example, the following Amazon S3 bucket policy allows members of any account in the XXX organization to add an

object into the examtopics bucket.

```
{"Version": "2020-09-10",  
"Statement": {  
  "Sid": "AllowPutObject",  
  "Effect": "Allow",  
  "Principal": "*",  
  "Action": "s3:PutObject",  
  "Resource": "arn:aws:s3:::examtopics/*",  
  "Condition": {"StringEquals":  
    {"aws:PrincipalOrgID": ["XXX"]}}}}
```

https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_policies_condition-keys.html

QUESTION 104

An application runs on an Amazon EC2 instance in a VPC. The application processes logs that are stored in an Amazon S3 bucket. The EC2 instance needs to access the S3 bucket without connectivity to the internet.

Which solution will provide private network connectivity to Amazon S3?

- A. Create a gateway VPC endpoint to the S3 bucket.
- B. Stream the logs to Amazon CloudWatch Logs. Export the logs to the S3 bucket.
- C. Create an instance profile on Amazon EC2 to allow S3 access.
- D. Create an Amazon API Gateway API with a private link to access the S3 endpoint.

Answer: A

QUESTION 105

A company is hosting a web application on AWS using a single Amazon EC2 instance that stores user-uploaded documents in an Amazon EBS volume. For better scalability and availability, the company duplicated the architecture and created a second EC2 instance and EBS volume in another Availability Zone placing both behind an Application Load Balancer. After completing this change, users reported that, each time they refreshed the website, they could see one subset of their documents or the other, but never all of the documents at the same time.

What should a solutions architect propose to ensure users see all of their documents at once?

- A. Copy the data so both EBS volumes contain all the documents.
- B. Configure the Application Load Balancer to direct a user to the server with the documents
- C. Copy the data from both EBS volumes to Amazon EFS.
Modify the application to save new documents to Amazon EFS
- D. Configure the Application Load Balancer to send the request to both servers.
Return each document from the correct server.

Answer: C

Explanation:

Amazon EFS provides file storage in the AWS Cloud. With Amazon EFS, you can create a file system, mount the file system on an Amazon EC2 instance, and then read and write data to and from your file system. You can mount an Amazon EFS file system in your VPC, through the Network File System versions 4.0 and 4.1 (NFSv4) protocol. We recommend using a current generation Linux NFSv4.1 client, such as those found in the latest Amazon Linux, Redhat, and Ubuntu AMLs, in conjunction with the Amazon EFS Mount Helper. For instructions, see Using the amazon-efs-utils Tools.

For a list of Amazon EC2 Linux Amazon Machine Images (AMIs) that support this protocol, see NFS Support. For some AMIs, you'll need to install an NFS client to mount your file system on your Amazon EC2 instance. For instructions, see Installing the NFS Client. You can access your Amazon EFS file system concurrently from multiple NFS clients, so applications that scale beyond a single connection can access a file system. Amazon EC2 instances running in multiple Availability Zones within the same AWS Region can access the file system, so that many users can access and share a common data source.

QUESTION 106

A company uses NFS to store large video files in on-premises network attached storage. Each video file ranges in size from 1MB to 500 GB. The total storage is 70 TB and is no longer growing. The company decides to migrate the video files to Amazon S3. The company must migrate the video files as soon as possible while using the least possible

[SAA-C03 Exam Dumps](#) [SAA-C03 Exam Questions](#) [SAA-C03 PDF Dumps](#) [SAA-C03 VCE Dumps](#)

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

network bandwidth.

Which solution will meet these requirements?

- A. Create an S3 bucket.
Create an IAM role that has permissions to write to the S3 bucket.
Use the AWS CLI to copy all files locally to the S3 bucket.
- B. Create an AWS Snowball Edge job.
Receive a Snowball Edge device on premises.
Use the Snowball Edge client to transfer data to the device.
Return the device so that AWS can import the data into Amazon S3.
- C. Deploy an S3 File Gateway on premises.
Create a public service endpoint to connect to the S3 File Gateway.
Create an S3 bucket.
Create a new NFS file share on the S3 File Gateway.
Point the new file share to the S3 bucket.
Transfer the data from the existing NFS file share to the S3 File Gateway.
- D. Set up an AWS Direct Connect connection between the on-premises network and AWS.
Deploy an S3 File Gateway on premises.
Create a public virtual interface (VIF) to connect to the S3 File Gateway.
Create an S3 bucket.
Create a new NFS file share on the S3 File Gateway.
Point the new file share to the S3 bucket.
Transfer the data from the existing NFS file share to the S3 File Gateway.

Answer: C

QUESTION 107

A company has an application that ingests incoming messages. These messages are then quickly consumed by dozens of other applications and microservices. The number of messages varies drastically and sometimes spikes as high as 100,000 each second. The company wants to decouple the solution and increase scalability. Which solution meets these requirements?

- A. Persist the messages to Amazon Kinesis Data Analytics.
All the applications will read and process the messages.
- B. Deploy the application on Amazon EC2 instances in an Auto Scaling group, which scales the number of EC2 instances based on CPU metrics.
- C. Write the messages to Amazon Kinesis Data Streams with a single shard.
All applications will read from the stream and process the messages.
- D. Publish the messages to an Amazon Simple Notification Service (Amazon SNS) topic with one or more Amazon Simple Queue Service (Amazon SQS) subscriptions.
All applications then process the messages from the queues.

Answer: D

Explanation:

"SNS Standard Topic"

Maximum throughput: Standard topics support a nearly unlimited number of messages per second.

<https://aws.amazon.com/sns/features/>

"SQS Standard Queue"

Unlimited Throughput: Standard queues support a nearly unlimited number of transactions per second (TPS) per API action.

<https://aws.amazon.com/sqs/features/>

QUESTION 108

A company is migrating a distributed application to AWS. The application serves variable workloads. The legacy platform consists of a primary server that coordinates jobs across multiple compute nodes. The company wants to modernize the application with a solution that maximizes resiliency and scalability.

How should a solutions architect design the architecture to meet these requirements?

[SAA-C03 Exam Dumps](#) **[SAA-C03 Exam Questions](#)** **[SAA-C03 PDF Dumps](#)** **[SAA-C03 VCE Dumps](#)**

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

- A. Configure an Amazon Simple Queue Service (Amazon SQS) queue as a destination for the jobs.
Implement the compute nodes with Amazon EC2 instances that are managed in an Auto Scaling group.
Configure EC2 Auto Scaling to use scheduled scaling.
- B. Configure an Amazon Simple Queue Service (Amazon SQS) queue as a destination for the jobs.
Implement the compute nodes with Amazon EC2 Instances that are managed in an Auto Scaling group.
Configure EC2 Auto Scaling based on the size of the queue.
- C. Implement the primary server and the compute nodes with Amazon EC2 instances that are managed in an Auto Scaling group.
Configure AWS CloudTrail as a destination for the fobs .
Configure EC2 Auto Scaling based on the load on the primary server.
- D. implement the primary server and the compute nodes with Amazon EC2 instances that are managed in an Auto Scaling group.
Configure Amazon EventBridge (Amazon CloudWatch Events) as a destination for the jobs.
Configure EC2 Auto Scaling based on the load on the compute nodes.

Answer: B

QUESTION 109

A company is running an SMB file server in its data center. The file server stores large files that are accessed frequently for the first few days after the files are created. After 7 days the files are rarely accessed. The total data size is increasing and is close to the company's total storage capacity. A solutions architect must increase the company's available storage space without losing low-latency access to the most recently accessed files. The solutions architect must also provide file lifecycle management to avoid future storage issues. Which solution will meet these requirements?

- A. Use AWS DataSync to copy data that is older than 7 days from the SMB file server to AWS.
- B. Create an Amazon S3 File Gateway to extend the company's storage space.
Create an S3 Lifecycle policy to transition the data to S3 Glacier Deep Archive after 7 days.
- C. Create an Amazon FSx for Windows File Server file system to extend the company's storage space.
- D. Install a utility on each user's computer to access Amazon S3.
Create an S3 Lifecycle policy to transition the data to S3 Glacier Flexible Retrieval after 7 days.

Answer: D

QUESTION 110

A company is building an ecommerce web application on AWS. The application sends information about new orders to an Amazon API Gateway REST API to process. The company wants to ensure that orders are processed in the order that they are received. Which solution will meet these requirements?

- A. Use an API Gateway integration to publish a message to an Amazon Simple Notification Service (Amazon SNS) topic when the application receives an order.
Subscribe an AWS Lambda function to the topic to perform processing.
- B. Use an API Gateway integration to send a message to an Amazon Simple Queue Service (Amazon SQS) FIFO queue when the application receives an order.
Configure the SQS FIFO queue to invoke an AWS Lambda function for processing.
- C. Use an API Gateway authorizer to block any requests while the application processes an order.
- D. Use an API Gateway integration to send a message to an Amazon Simple Queue Service (Amazon SQS) standard queue when the application receives an order.
Configure the SQS standard queue to invoke an AWS Lambda function for processing.

Answer: A