Amazon

SAA-C03 Exam

AWS Certified Solutions Architect - Associate
Questions & Answers
Demo

Version: 11.0

Topic 1, Exam Pool A

Question: 1

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

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-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: 2

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: 3

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: 4

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

Explanation:

VPC endpoint allows you to connect to AWS services using a private network instead of using the public Internet

Question: 5

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: A

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

AMIs, 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.

https://docs.aws.amazon.com/efs/latest/ug/how-it-works.html#how-it-works-ec2