

KillTest

Higher Quality, Better Service!



Q&A

<http://www.killtest.com>

We offer free update service for one year.

Exam : SAA-C03

**Title : AWS Certified Solutions
Architect - Associate**

Version : DEMO

1.A company's software development team needs an Amazon RDS Multi-AZ cluster. The RDS cluster will serve as a backend for a desktop client that is deployed on premises. The desktop client requires direct connectivity to the RDS cluster.

The company must give the development team the ability to connect to the cluster by using the client when the team is in the office.

Which solution provides the required connectivity MOST securely?

- A. Create a VPC and two public subnets. Create the RDS cluster in the public subnets. Use AWS Site-to-Site VPN with a customer gateway in the company's office.
- B. Create a VPC and two private subnets. Create the RDS cluster in the private subnets. Use AWS Site-to-Site VPN with a customer gateway in the company's office.
- C. Create a VPC and two private subnets. Create the RDS cluster in the private subnets. Use RDS security groups to allow the company's office IP ranges to access the cluster.
- D. Create a VPC and two public subnets. Create the RDS cluster in the public subnets. Create a cluster user for each developer. Use RDS security groups to allow the users to access the cluster.

Answer: B

Explanation:

Requirement Analysis: Need secure, direct connectivity from an on-premises client to an RDS cluster, accessible only when in the office.

VPC with Private Subnets: Ensures the RDS cluster is not publicly accessible, enhancing security.

Site-to-Site VPN: Provides secure, encrypted connection between on-premises office and AWS VPC.

Implementation:

Create a VPC with two private subnets.

Launch the RDS cluster in the private subnets.

Set up a Site-to-Site VPN connection with a customer gateway in the office.

Conclusion: This setup ensures secure and direct connectivity with minimal exposure, meeting the requirement for secure access from the office.

Reference

AWS Site-to-Site VPN: [AWS Site-to-Site VPN Documentation](#)

Amazon RDS: [Amazon RDS Documentation](#)

2.A social media company wants to store its database of user profiles, relationships, and interactions in the AWS Cloud. The company needs an application to monitor any changes in the database. The application needs to analyze the relationships between the data entities and to provide recommendations to users.

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

- A. Use Amazon Neptune to store the information. Use Amazon Kinesis Data Streams to process changes in the database.
- B. Use Amazon Neptune to store the information. Use Neptune Streams to process changes in the database.
- C. Use Amazon Quantum Ledger Database (Amazon QLDB) to store the information. Use Amazon Kinesis Data Streams to process changes in the database.
- D. Use Amazon Quantum Ledger Database (Amazon QLDB) to store the information. Use Neptune Streams to process changes in the database.

Answer: B

Explanation:

Amazon Neptune: Neptune is a fully managed graph database service that is optimized for storing and querying highly connected data. It supports both property graph and RDF graph models, making it suitable for applications that need to analyze relationships between data entities.

Neptune Streams: Neptune Streams captures changes to the graph and streams these changes to other AWS services. This is useful for applications that need to monitor and respond to changes in real-time, such as providing recommendations based on user interactions and relationships.

Least Operational Overhead: Using Neptune Streams directly with Amazon Neptune ensures that the solution is tightly integrated, reducing the need for additional components and minimizing operational overhead. This integration simplifies the architecture by eliminating the need for a separate service like Kinesis for change processing.

Reference: Amazon Neptune Documentation

Neptune Streams Documentation

3.A company uses an Amazon S3 bucket as its data lake storage platform. The S3 bucket contains a massive amount of data that is accessed randomly by multiple teams and hundreds of applications. The company wants to reduce the S3 storage costs and provide immediate availability for frequently accessed objects

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

- A. Create an S3 Lifecycle rule to transition objects to the S3 Intelligent-Tiering storage class
- B. Store objects in Amazon S3 Glacier Use S3 Select to provide applications with access to the data.
- C. Use data from S3 storage class analysis to create S3 Lifecycle rules to automatically transition objects to the S3 Standard-Infrequent Access (S3 Standard-IA) storage class.
- D. Transition objects to the S3 Standard-Infrequent Access (S3 Standard-IA) storage class Create an AWS Lambda function to transition objects to the S3 Standard storage class when they are accessed by an application

Answer: A

Explanation:

Amazon S3 Intelligent-Tiering: This storage class is designed to optimize costs by automatically moving data between two access tiers (frequent and infrequent) when access patterns change. It provides cost savings without performance impact or operational overhead.

S3 Lifecycle Rules: By creating an S3 Lifecycle rule, the company can automatically transition objects to the Intelligent-Tiering storage class. This eliminates the need for manual intervention and ensures that objects are moved to the most cost-effective storage tier based on their access patterns.

Operational Efficiency: Intelligent-Tiering requires no additional management and delivers immediate availability for frequently accessed objects. This makes it the most operationally efficient solution for the given requirements.

Reference: Amazon S3 Intelligent-Tiering

S3 Lifecycle Policies

4.A company needs to optimize its Amazon S3 storage costs for an application that generates many files that cannot be recreated Each file is approximately 5 MB and is stored in Amazon S3 Standard storage. The company must store the files for 4 years before the files can be deleted. The files must be immediately accessible. The files are frequently accessed in the first 30 days of object creation, but they

are rarely accessed after the first 30 days.

Which solution will meet these requirements MOST cost-effectively?

- A. Create an S3 Lifecycle policy to move the files to S3 Glacier Instant Retrieval 30 days after object creation. Delete the files 4 years after object creation.
- B. Create an S3 Lifecycle policy to move the files to S3 One Zone-Infrequent Access (S3 One Zone-IA) 30 days after object creation. Delete the files 4 years after object creation.
- C. Create an S3 Lifecycle policy to move the files to S3 Standard-Infrequent Access (S3 Standard-IA) 30 days after object creation. Delete the files 4 years after object creation.
- D. Create an S3 Lifecycle policy to move the files to S3 Standard-Infrequent Access (S3 Standard-IA) 30 days after object creation. Move the files to S3 Glacier Flexible Retrieval 4 years after object creation.

Answer: C

Explanation:

Amazon S3 Standard-IA: This storage class is designed for data that is accessed less frequently but requires rapid access when needed. It offers lower storage costs compared to S3 Standard while still providing high availability and durability.

Access Patterns: Since the files are frequently accessed in the first 30 days and rarely accessed afterward, transitioning them to S3 Standard-IA after 30 days aligns with their access patterns and reduces storage costs significantly.

Lifecycle Policy: Implementing a lifecycle policy to transition the files to S3 Standard-IA ensures automatic management of the data lifecycle, moving files to a lower-cost storage class without manual intervention. Deleting the files after 4 years further optimizes costs by removing data that is no longer needed.

Reference: Amazon S3 Storage Classes

S3 Lifecycle Configuration

5. A company runs an AWS Lambda function in private subnets in a VPC. The subnets have a default route to the internet through an Amazon EC2 NAT instance. The Lambda function processes input data and saves its output as an object to Amazon S3.

Intermittently, the Lambda function times out while trying to upload the object because of saturated traffic on the NAT instance's network. The company wants to access Amazon S3 without traversing the internet.

Which solution will meet these requirements?

- A. Replace the EC2 NAT instance with an AWS managed NAT gateway.
- B. Increase the size of the EC2 NAT instance in the VPC to a network optimized instance type
- C. Provision a gateway endpoint for Amazon S3 in the VPC. Update the route tables of the subnets accordingly.
- D. Provision a transit gateway. Place transit gateway attachments in the private subnets where the Lambda function is running.

Answer: C

Explanation:

Gateway Endpoint for Amazon S3: A VPC endpoint for Amazon S3 allows you to privately connect your VPC to Amazon S3 without requiring an internet gateway, NAT device, VPN connection, or AWS Direct Connect connection.

Provisioning the Endpoint:

Navigate to the VPC Dashboard.

Select "Endpoints" and create a new endpoint.

Choose the service name for S3 (com.amazonaws.region.s3).

Select the appropriate VPC and subnets.

Adjust the route tables of the subnets to include the new endpoint.

Update Route Tables: Modify the route tables of the subnets to direct traffic destined for S3 to the newly created endpoint. This ensures that traffic to S3 does not go through the NAT instance, avoiding the saturated network and eliminating timeouts.

Operational Efficiency: This solution minimizes operational overhead by removing dependency on the NAT instance and avoiding internet traffic, leading to more stable and secure S3 interactions.

Reference: VPC Endpoints for Amazon S3

Creating a Gateway Endpoint