

KillTest

高品質 より良いサービス



練習問題

<https://www.killtest.jp>

1年間無料アップデート

Exam : SAP-C02

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

Version : DEMO

1.A company runs an IoT application in the AWS Cloud. The company has millions of sensors that collect data from houses in the United States. The sensors use the MQTT protocol to connect and send data to a custom MQTT broker. The MQTT broker stores the data on a single Amazon EC2 instance. The sensors connect to the broker through the domain named `iot.example.com`. The company uses Amazon Route 53 as its DNS service. The company stores the data in Amazon DynamoDB.

On several occasions, the amount of data has overloaded the MQTT broker and has resulted in lost sensor data. The company must improve the reliability of the solution.

Which solution will meet these requirements?

- A. Create an Application Load Balancer (ALB) and an Auto Scaling group for the MQTT broker. Use the Auto Scaling group as the target for the ALB. Update the DNS record in Route 53 to an alias record. Point the alias record to the ALB. Use the MQTT broker to store the data.
- B. Set up AWS IoT Core to receive the sensor data. Create and configure a custom domain to connect to AWS IoT Core. Update the DNS record in Route 53 to point to the AWS IoT Core Data-ATS endpoint. Configure an AWS IoT rule to store the data.
- C. Create a Network Load Balancer (NLB). Set the MQTT broker as the target. Create an AWS Global Accelerator accelerator. Set the NLB as the endpoint for the accelerator. Update the DNS record in Route 53 to a multivalue answer record. Set the Global Accelerator IP addresses as values. Use the MQTT broker to store the data.
- D. Set up AWS IoT Greengrass to receive the sensor data. Update the DNS record in Route 53 to point to the AWS IoT Greengrass endpoint. Configure an AWS IoT rule to invoke an AWS Lambda function to store the data.

Answer: A

Explanation:

it describes a solution that uses an Application Load Balancer (ALB) and an Auto Scaling group for the MQTT broker. The ALB distributes incoming traffic across the instances in the Auto Scaling group and allows for automatic scaling based on incoming traffic. The use of an alias record in Route 53 allows for easy updates to the DNS record without changing the IP address. This solution improves the reliability of the MQTT broker by allowing it to automatically scale based on incoming traffic, reducing the likelihood of lost data due to broker overload.

Reference:

<https://aws.amazon.com/elasticloadbalancing/applicationloadbalancer/>

<https://aws.amazon.com/autoscaling/>

<https://aws.amazon.com/route53/>

2.A life sciences company is using a combination of open source tools to manage data analysis workflows and Docker containers running on servers in its on-premises data center to process genomics data Sequencing data is generated and stored on a local storage area network (SAN), and then the data is processed. The research and development teams are running into capacity issues and have decided to re-architect their genomics analysis platform on AWS to scale based on workload demands and reduce the turnaround time from weeks to days

The company has a high-speed AWS Direct Connect connection Sequencers will generate around 200 GB of data for each genome, and individual jobs can take several hours to process the data with ideal compute capacity. The end result will be stored in Amazon S3. The company is expecting 10-15 job requests each day

Which solution meets these requirements?

- A. Use regularly scheduled AWS Snowball Edge devices to transfer the sequencing data into AWS. When AWS receives the Snowball Edge device and the data is loaded into Amazon S3, use S3 events to trigger an AWS Lambda function to process the data.
- B. Use AWS Data Pipeline to transfer the sequencing data to Amazon S3. Use S3 events to trigger an Amazon EC2 Auto Scaling group to launch custom-AMI EC2 instances running the Docker containers to process the data.
- C. Use AWS DataSync to transfer the sequencing data to Amazon S3. Use S3 events to trigger an AWS Lambda function that starts an AWS Step Functions workflow. Store the Docker images in Amazon Elastic Container Registry (Amazon ECR) and trigger AWS Batch to run the container and process the sequencing data.
- D. Use an AWS Storage Gateway file gateway to transfer the sequencing data to Amazon S3. Use S3 events to trigger an AWS Batch job that runs on Amazon EC2 instances running the Docker containers to process the data.

Answer: C

Explanation:

AWS DataSync can be used to transfer the sequencing data to Amazon S3, which is a more efficient and faster method than using Snowball Edge devices. Once the data is in S3, S3 events can trigger an AWS Lambda function that starts an AWS Step Functions workflow. The Docker images can be stored in Amazon Elastic Container Registry (Amazon ECR) and AWS Batch can be used to run the container and process the sequencing data.

3. A company needs to optimize the cost of an AWS environment that contains multiple accounts in an organization in AWS Organizations. The company conducted cost optimization activities 3 years ago and purchased Amazon EC2 Standard Reserved Instances that recently expired.

The company needs EC2 instances for 3 more years. Additionally, the company has deployed a new serverless workload.

Which strategy will provide the company with the MOST cost savings?

- A. Purchase the same Reserved Instances for an additional 3-year term with All Upfront payment. Purchase a 3-year Compute Savings Plan with All Upfront payment in the management account to cover any additional compute costs.
- B. Purchase a 1-year Compute Savings Plan with No Upfront payment in each member account. Use the Savings Plans recommendations in the AWS Cost Management console to choose the Compute Savings Plan.
- C. Purchase a 3-year EC2 Instance Savings Plan with No Upfront payment in the management account to cover EC2 costs in each AWS Region. Purchase a 3-year Compute Savings Plan with No Upfront payment in the management account to cover any additional compute costs.
- D. Purchase a 3-year EC2 Instance Savings Plan with All Upfront payment in each member account. Use the Savings Plans recommendations in the AWS Cost Management console to choose the EC2 Instance Savings Plan.

Answer: A

Explanation:

The company should purchase the same Reserved Instances for an additional 3-year term with All Upfront payment. The company should purchase a 3-year Compute Savings Plan with All Upfront

payment in the management account to cover any additional compute costs. This solution will provide the company with the most cost savings because Reserved Instances and Savings Plans are both pricing models that offer significant discounts compared to On-Demand pricing. Reserved Instances are commitments to use a specific instance type and size in a single Region for a one- or three-year term. You can choose between three payment options: No Upfront, Partial Upfront, or All Upfront. The more you pay upfront, the greater the discount¹. Savings Plans are flexible pricing models that offer low prices on EC2 instances, Fargate, and Lambda usage, in exchange for a commitment to a consistent amount of usage (measured in \$/hour) for a one- or three-year term. You can choose between two types of Savings Plans: Compute Savings Plans and EC2 Instance Savings Plans. Compute Savings Plans apply to any EC2 instance regardless of Region, instance family, operating system, or tenancy, including those that are part of EMR, ECS, or EKS clusters, or launched by Fargate or Lambda. EC2 Instance Savings Plans apply to a specific instance family within a Region and provide the most savings². By purchasing the same Reserved Instances for an additional 3-year term with All Upfront payment, the company can lock in the lowest possible price for its EC2 instances that run continuously for 3 years. By purchasing a 3-year Compute Savings Plan with All Upfront payment in the management account, the company can benefit from additional discounts on any other compute usage across its member accounts.

The other options are not correct because:

Purchasing a 1-year Compute Savings Plan with No Upfront payment in each member account would not provide as much cost savings as purchasing a 3-year Compute Savings Plan with All Upfront payment in the management account. A 1-year term offers lower discounts than a 3-year term, and a No Upfront payment option offers lower discounts than an All Upfront payment option. Also, purchasing a Savings Plan in each member account would not allow the company to share the benefits of unused Savings Plan discounts across its organization.

Purchasing a 3-year EC2 Instance Savings Plan with No Upfront payment in the management account to cover EC2 costs in each AWS Region would not provide as much cost savings as purchasing Reserved Instances for an additional 3-year term with All Upfront payment. An EC2 Instance Savings Plan offers lower discounts than Reserved Instances for the same instance family and Region. Also, a No Upfront payment option offers lower discounts than an All Upfront payment option.

Purchasing a 3-year EC2 Instance Savings Plan with All Upfront payment in each member account would not provide as much flexibility or cost savings as purchasing a 3-year Compute Savings Plan with All Upfront payment in the management account. An EC2 Instance Savings Plan applies only to a specific instance family within a Region and does not cover Fargate or Lambda usage. Also, purchasing a Savings Plan in each member account would not allow the company to share the benefits of unused Savings Plan discounts across its organization.

Reference:

<https://aws.amazon.com/ec2/pricing/reserved-instances/>

<https://aws.amazon.com/savingsplans/>

4. A company has a latency-sensitive trading platform that uses Amazon DynamoDB as a storage backend. The company configured the DynamoDB table to use on-demand capacity mode. A solutions architect needs to design a solution to improve the performance of the trading platform. The new solution must ensure high availability for the trading platform.

Which solution will meet these requirements with the LEAST latency?

A. Create a two-node DynamoDB Accelerator (DAX) cluster. Configure an application to read and write

data by using DAX.

B. Create a three-node DynamoDB Accelerator (DAX) cluster. Configure an application to read data by using DAX and to write data directly to the DynamoDB table.

C. Create a three-node DynamoDB Accelerator (DAX) cluster. Configure an application to read data directly from the DynamoDB table and to write data by using DAX.

D. Create a single-node DynamoDB Accelerator (DAX) cluster. Configure an application to read data by using DAX and to write data directly to the DynamoDB table.

Answer: B

Explanation:

A DAX cluster can be deployed with one or two nodes for development or test workloads. One- and two-node clusters are not fault-tolerant, and we don't recommend using fewer than three nodes for production use. If a one- or two-node cluster encounters software or hardware errors, the cluster can become unavailable or lose cached data. A DAX cluster can be deployed with one or two nodes for development or test workloads. One- and two-node clusters are not fault-tolerant, and we don't recommend using fewer than three nodes for production use. If a one- or two-node cluster encounters software or hardware errors, the cluster can become unavailable or lose cached data.

<https://docs.aws.amazon.com/amazondynamodb/latest/developerguide/DAX.concepts.cluster.html>

5. A company is running a data-intensive application on AWS. The application runs on a cluster of hundreds of Amazon EC2 instances. A shared file system also runs on several EC2 instances that store 200 TB of data. The application reads and modifies the data on the shared file system and generates a report. The job runs once monthly, reads a subset of the files from the shared file system, and takes about 72 hours to complete. The compute instances scale in an Auto Scaling group, but the instances that host the shared file system run continuously. The compute and storage instances are all in the same AWS Region.

A solutions architect needs to reduce costs by replacing the shared file system instances. The file system must provide high performance access to the needed data for the duration of the 72-hour run.

Which solution will provide the LARGEST overall cost reduction while meeting these requirements?

A. Migrate the data from the existing shared file system to an Amazon S3 bucket that uses the S3 Intelligent-Tiering storage class. Before the job runs each month, use Amazon FSx for Lustre to create a new file system with the data from Amazon S3 by using lazy loading. Use the new file system as the shared storage for the duration of the job. Delete the file system when the job is complete.

B. Migrate the data from the existing shared file system to a large Amazon Elastic Block Store (Amazon EBS) volume with Multi-Attach enabled. Attach the EBS volume to each of the instances by using a user data script in the Auto Scaling group launch template. Use the EBS volume as the shared storage for the duration of the job. Detach the EBS volume when the job is complete.

C. Migrate the data from the existing shared file system to an Amazon S3 bucket that uses the S3 Standard storage class. Before the job runs each month, use Amazon FSx for Lustre to create a new file system with the data from Amazon S3 by using batch loading. Use the new file system as the shared storage for the duration of the job. Delete the file system when the job is complete.

D. Migrate the data from the existing shared file system to an Amazon S3 bucket. Before the job runs each month, use AWS Storage Gateway to create a file gateway with the data from Amazon S3. Use the file gateway as the shared storage for the job. Delete the file gateway when the job is complete.

Answer: A

Explanation:

<https://aws.amazon.com/blogs/storage/new-enhancements-for-moving-data-between-amazon-fsx-for-lustre-and-amazon-s3/>