Pear cloud proposal

Team Name: Smarchitects

After carefully examining the requirements of the company. Considering the future plans and projections, the following proposal is created to migrate Pear's infrastructure to AWS Cloud.

The existing on-premises IT solution has the following problems.

1. The existing infrastructure hardware is more than 5 years old and projected Capex is $100 million for the next 2 years.
2. The current dataset is 100TB and is projected to increase by 100% in the next 5 years.
3. The current retail website is a tightly coupled classic monolithic solution that endured over 90 hours of downtime due to demand and hardware issues.
4. The application sustained a number of DDoS attacks, unable to diagnose due to the lack of logging and storing the logs for analysis.
5. The CFO also wanted a Realtime website analysis, which is not possible due to lack of an efficient solution.

After carefully considered all the current problems. A number of microservice architectures are recommended inline with the AWS Well-Architected Framework.

The solution is explained in the following steps.

1. Pear is suggested to move its workload and storage to AWS cloud. This allows the company to eliminate the Capex and the overhead in maintaining the huge infrastructure on premises.
2. The recommended service to migrate the on-premises datasets is AWS Database Migration Service (DMS). Oracle, Hadoop, and Spark clusters are migrated to Amazon EMR, SQL database is migrated to Amazon Aurora. 2x Storage optimized Snowball Edge devices with 80TB of storage are used to initially transfer the datasets to Amazon S3 and then to EMR and Aurora databases. The expected cost of migration for is $725 for Snowball Edge. Apache Sqoop service is used to regularly transfer and store aggregates and summaries in Amazon Aurora.
3. The retail website is rearchitected to use containers on Amazon Elastic Container Service (ECS). The containers are deployed on AWS Fargate to reduce the overhead of managing the underlying infrastructure. An application load balancer is placed in front of the Fargate compute layer. AWS Fargate is capable of scaling in and out automatically according to the demand and failed containers are replaced immediately without much downtime. Amazon Aurora is used for the database layer to take advantage of the flexibility and read replicas features to make the application highly available and also fault tolerant. Regular backups are stored in S3 Glacier Instant Retrieval, also replicated to a second region.
4. To mitigate the DDoS attacks, AWS Shield Advanced with WAF is used to filter the traffic and block known DDoS attacks. AWS Shield Advanced provides enhanced resource specific detection and employs advanced mitigation and routing techniques for sophisticated or larger attacks. AWS Shield Advanced also provides visibility and insights into all your DDoS incidents through AWS CloudWatch metrics and attack diagnostics. AWS Shield Advanced also protects the company from bill spikes due to the DDoS attacks. According to the best practices, the application is made highly scalable using services like Route 53, CloudFront and

Elastic Load Balancer. The application is secured with AWS Shield Advanced and WAF. The DDoS attacks and the traffic are monitored using CloudWatch.

5. The logs of the retail website are collected through Kinesis Data Firehose by installing the agent on AWS Fargate. The log data is streamed to AWS OpenSearch for real-time analysis as per the CFO's request. The data is also streamed to a S3 bucket for storage. AWS Glue Crawler and Glue Data Catalogue are used to create a schema and Amazon Athena is used to analyze for trends and meaningful information.

The architecture is presented below:

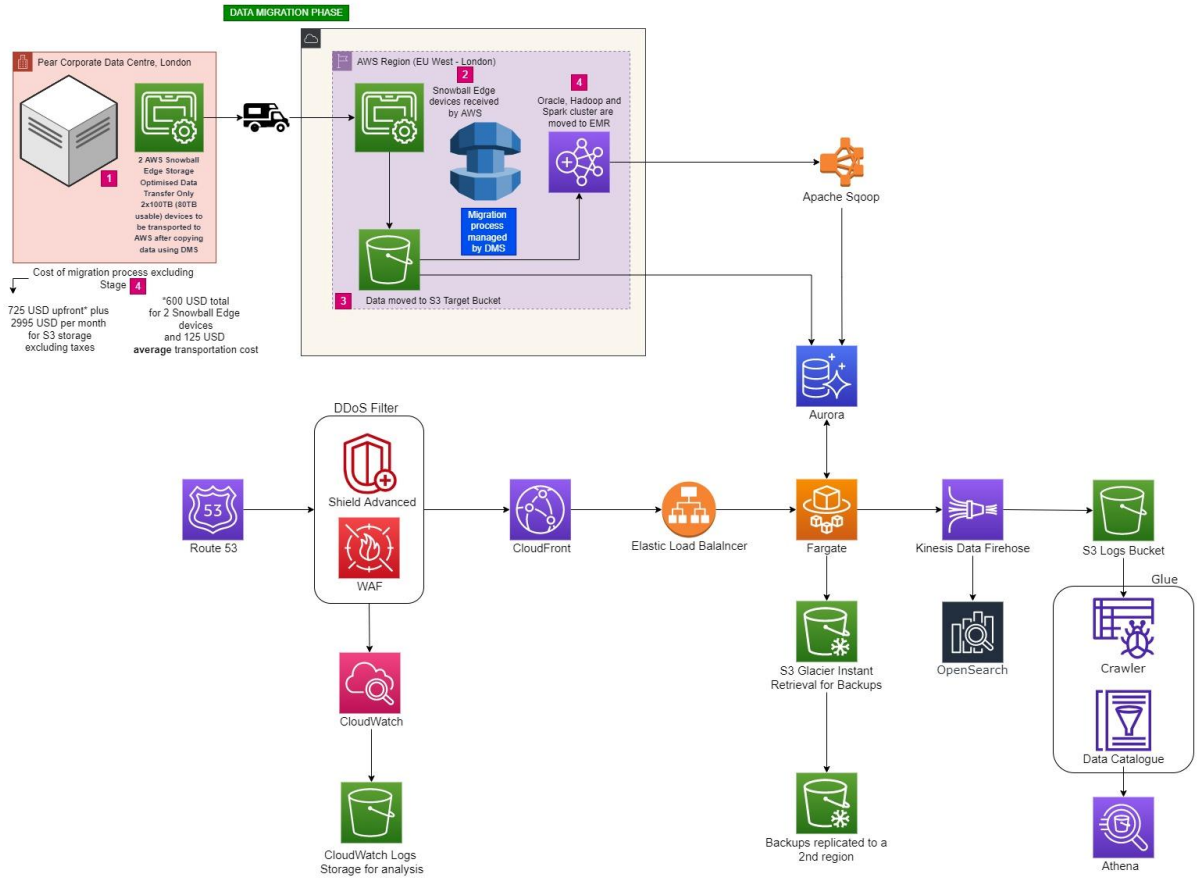

Figure 1. Architecture Diagram

Team Members:

- Aditya Ganji (Team Leader)
- Cristal Gallano
- Maciej Cwalina
- Mahesh Nair
- Phuntsok Dalu
- Sharayi Nyikavaranda
- Tony Sahajpal