

40% saving by Migrating to Google Cloud

Customer: One of the top tech-based real estate platform

Summary

Case: Migration

Type: Re-host and Re-architect

Number of VM's: 140+

Number of applications migrated: 5

Approximate size of DB: 5 TB +

Tools used: GitHub, Jenkins, MongoDB, MySQL, MongoDB, Redis, ELK, Kafka, Elasticsearch.

Google Services used - Compute Engine, Container Build, Google Kubernetes Engine, Container Registry, Cloud SQL, Google Storage Bucket, Cloud Identity & Access Management, Cloud VPN, Cloud DNS, Cloud Load Balancing.

The customer is a technology-based real estate platform that has been built to buy, sell, rent, and find a paying guest or a flatmate without paying any brokerage. They enable owners and tenants to interact with each other directly by using their technologically advanced platform. The plan is to re-architect and migrate their entire infrastructure to Google Cloud Platform (GCP) to reduce costs and increase their portal's workload efficiency.

About Customer

Headquartered in Bengaluru, India, the customer is a technology-based real estate search portal that connects property owners and tenants and buyers and sellers directly, eliminating the concept of a broker. The organization does not charge any brokerage for the services it provides. It was founded by IIT alumnus in March 2014 and is a team of 350 people serving over 35 lakh customers today. They have worked endlessly to remove all asymmetric information caused by brokers.

Their goal is to lead India's real estate industry towards an era of doing convenient housing transactions, brokerage-free. They currently save their customers over 250 crores per year in brokerage. They are proving to be a new and disruptive force in the real estate industry.

The customer has been running their infrastructure on Digital Ocean and after much research and analysis, they have evaluated GCP and its solution to be more suited to their requirements. Powerup is proposing to work with the customer to help them migrate their infrastructure to Google Cloud from Digital Ocean, which in turn would help them in running their workload on cloud in a very efficient and cost-effective manner.

Problem Statement

The customer currently runs its infrastructure on Digital Ocean. This set up, however, restricts them from managing their network, containers and storage connectivity. They have also been unable to use multiple features like auto-scaling, managed Kubernetes, load balancer and Google Cloud Storage due to the current constricted set up.

Powerup has proposed to work with the customer team to help migrate their infrastructure from the Digital Ocean to Google Cloud Platform (GCP), which will help them in running their workload on the cloud more efficiently and cost-effectively. However, any application-level changes as part of the migration and on-premise or Digital Ocean VPN tunnel setup is out of scope for this project. Meanwhile, on Powerup's recommendation, the customer has evaluated Google Cloud Platform (GCP) and found the solution to be most suited to their requirements.

Proposed Solution

Pre-Migration

Understanding the current Digital Ocean setup, migration timelines and the business needs for it, application and network architectures were the key pre-migration tasks.

Project Set Up

To establish a hierarchy of ownership, GCP organization node was setup. Separate folders were created for production, UAT and management environments. Identity and access management (IAM) framework was adapted to access different GCP services. This was mainly to have control on user access to critical information within the organization, which meant, a GCP service could securely access data on another service.

The set up involved organizing shared Virtual Private Cloud (VPC). Multiple such VPCs were created to deploy production, UAT and management applications. Configuring VPC peering between the required VPCs, creating appropriate CIDR range, subnets, and route tables as per architecture requirements and establishing VPN tunnels between Digital Ocean/ On-premise and GCP were all part of network configuration.

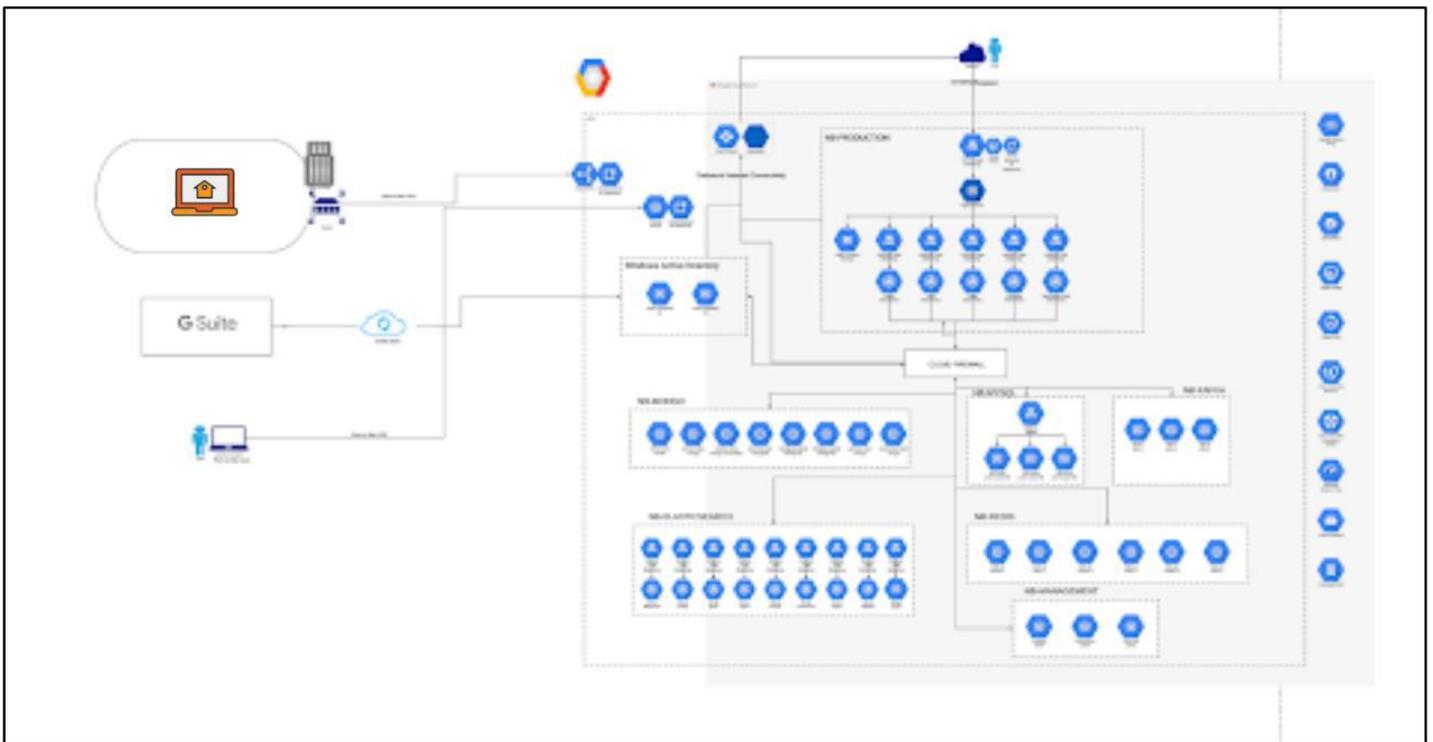
Cloud Network Address Translation (NAT) enabled Internet access for servers in the private subnet. Hypertext Transfer Protocol Secure (HTTPS) load balancers were created to distribute the traffic between the web and app servers.

Re-architecting and migrating to GKE

CloudEndure service was used to migrate the application and databases from Digital Ocean to GCP. The purpose was to install CloudEndure Agents on the required servers and ports to be opened to connect to GCP. A staging VPC in the respective folders was to host the Linux instance

(volumes) to which CloudEndure will replicate the data. CloudEndure will then replicate the data at the block level to the volumes on GCP. Once the replication is done CloudEndure will deploy the application and database servers to the required subnets mentioned in the replication settings.

All ElasticSearch clusters would be created and the data imported to the new cluster. Frontend application servers and Elastic search clusters were deployed on GKE. SQL cluster was deployed on GCP as a Master-Master setup with Percona. Redis Instance will be used to host the Redis Cache. MYSQL database will have multiple Read Replicas created. MongoDB will contain all the status data for emails and SMS sent to the customer. Kafka cluster will be recreated on GCP. Storm application will be migrated using CloudEndure. All static content will be migrated from S3 to Cloud Storage on GCP. Content Delivery Network (CDN) will be migrated from AWS Cloudfront to GCP Cloud CDN. Jenkins will be used to Setup Continuous integration/continuous deployment (CI/CD) for the application servers. Backups will be taken based on recommended best practices and customer requirements.



[Architecture diagram]

Monitoring and logging

Stackdriver tool to be used as the Monitoring tool for infrastructure/application monitoring, log analytics etc. Stackdriver also supports features like tracing, debugging, profiling etc. to monitor the overall performance of the application. Appropriate alarms and Stackdriver alerting policies

have to be configured based on customer requirements. Sensus will be used for additional monitoring of infrastructure. FluxDB will be used as the database to store all data and Graphana will be used to visualize and create the dashboards.

Cloud Audit logging will be enabled to capture all API activities in the account. VPC flow logs will be enabled to capture all network logs. Centralized Logging and monitoring will be done using the Stackdriver Logging module.

Security & Compliance

IAM users and groups will be created with the least permissible access. MFA will be enabled for account access which will provide an additional layer of security.

All application, database, Elastic Search (ES), Kafka, etc will be deployed in the private subnet. This ensures that there is no direct internet access to critical servers like database and app servers. Firewall rules will be configured to control traffic to and from the resources attached to the VPC. Firewall rules are applied at the virtual networking level, so they help provide effective protection and traffic control regardless of the operating system that the instances use. Only the required ports will be opened, and access will be provided only to the required IP addresses. Both data in transit and at rest are by default encrypted in GCP. All GCP API endpoints are SSL enabled. VPN Tunnels will be enabled between GCP, Customer Locations & Customer Data Centers.

Benefits

- Google Cloud Platform (GCP) helped the customer sail through a hassle-free digital transformation providing them with a highly efficient and cost-effective solution.
- Re-architecting cut down their costs by approximately 40% and helped them gain more scalability and high availability.
- Managing the current setup on GCP became more structured. This helped them gain more control from a security standpoint.
- With a combination of strategies like recreation/lift and shift, re-architecture of their current infrastructure running on Digital Ocean being adopted and planned with advanced technologies like Managed Kubernetes, App Engine, Big Query & Managed Load Balancer lead to a better experience for its customers. Additionally, the increase of web traffic on their site had no adverse impact on end-users and there were zero performance issues.
- The customer team could even scale up the infrastructure at any point in time as and when required.

Cloud platform

GCP.