

DataSynapse GridServer HPC Reference Architecture on GCP

Success story

Repeatable Reference Architecture
Deployment for DataSynapse
GridServer on GCP



Alter a full GridServer environment
at the touch of a button



Quicker time-to-market for new
HPC projects



Easily support exceptional
workloads

CLOUD



THE CHALLENGE

Migration of Equities DataSynapse GridServer processing workload to GCP

- Aside from Equities, the same bank also used DataSynapse GridServer for running HPC workloads for Fixed Income, FX Options, Market Risk and Credit Risk.
- GFT was engaged by Google to develop a deployable reference architecture for running DataSynapse GridServer on GCP. The purpose was for the reference architecture to be used in the Equities migration and later on for other use cases, including for other GCP clients wishing to migrate their on-premise HPC workloads to GCP.

THE ENGAGEMENT

GFT deployed a team consisting of an HPC Cloud Architect/Project Manager and two Cloud DevOps Engineers to the project, which was funded through Google PSF for 2 months

- GFT designed a target architecture for the deployable reference architecture and put together a backlog in order to deliver the project in an Agile fashion with weekly sprints and “show and tells”.
- The backlog included user stories to create a basic DataSynapse GCP deployment; automating the deployment using Terraform, Ansible and Powershell; incorporating networking and security best practices; implementing autoscaling functionality (using both CPU metrics and a custom metric reading the number of upcoming tasks from the GridServer Director); encryption of all data using Customer Managed Encryption Keys (CMEK); integrating all logging and monitoring activity with Stackdriver; implementing DevOps practices to deploy GridLibs using a Jenkins CI/CD pipeline and demonstrating visual analytics of GridServer workloads through TIBCO Spotfire integration.
- With some help from TIBCO, Google and the banking client, GFT managed to deliver all user stories in the backlog in the 2 month timeframe.

THE BENEFIT

The reference architecture is now available and has been used to migrate the Equities workloads to GCP

- Fixed Income is also planning to use the reference architecture to migrate. Other areas and GCP clients have also expressed an interest in using the reference architecture.
- Main benefits delivered by the reference architecture are the ability to “stand up” and “tear down” a full GridServer environment at the touch of a button, parameterizable engine types and initial engine pool sizes, enabling cost-optimization, quicker time-to-market for new HPC projects, ability to economically run a full-scale UAT and to easily support exceptional workloads, e.g. VaR re-runs. The reference architecture can be modified to support AWS & IBM Spectrum Symphony.