

Migrating Legacy Applications to the Cloud

How SenTai supported Canaccord Genuity's
Move to Google Cloud Platform

David Senouf, PhD \ \ CEO SenTai



Google Cloud

SenTai

Who we are and what we do

SenTai \ \

SenTai: Consulting group specialized in FinTech, based in Paris, founded in 2013

Core Team: Group of (Data) Scientists, Engineers, and DevOps (developers & operations). Team can easily grow thanks to extensive network

Joint Work: Working together for 10 years in various Silicon Valley startups, Paris startups, as well as Paris based consulting projects

FinTech Experience: Each team member has extensive 10Y+ experience in the Financial Services industry - Platforms (Quest), Banking Systems (Banque de France), Infra & Big Data (AXA, SG, BNP)

Backgrounds: Academic Research (Mathematics & Physics), Computer Engineering (Infra / Big Data), Finance (Trading, Quant, Modelling & Data Science), Data Science (Algo Trading Algo, Aeronautics, GeoLoc)

Today's Agenda

01 Why move to the cloud?

02 The Quest Equity Valuation Platform

- What is Quest[®]?
- Quest limitations prior to migration?
- Reasons for limitations?

03 Migration

- Challenges
- Chronology
- Infrastructure comparison

04 Benefits

- Gains (Scalability, Performance, Agility)
- Costs
- Risk Mitigation (Security)
- Feedback

01 Why move to the cloud?

Why move key projects to the cloud?

Today \\ Legacy financial platform issues

- High Costs - Inefficiency - Difficulty to maintain - Lack of scalability
- Slow development cycles
- Complexity: Heterogeneous parts pieced together during lifespan of legacy applications

Tomorrow \\ Cloud computing platform advantages

- Managed Solutions (IaaS, PaaS): Scalability, Flexibility, Security, Economy and Performance
- Fast Implementation - Little maintenance
- Complete reset - Software architecture re-think thanks to Cloud tools (GCP)

02 The Quest Equity Valuation Platform

What is Quest®?

Quest® limitations prior to migration?

Reasons for limitations?

What is Quest?



An automated equity valuation computational platform

Quest is owned by **Canaccord Genuity**, a global investment bank

Quest is one of two (with CS's Holt) **major fully automated equity valuation tool available on the web**: major differentiator for a mid-tier investment bank

- Proprietary analytical equity research system
- Analyzes 8,800 listed companies (90% of world market cap.)
- End users are inst. investors, equity research analysts, portfolio managers, hedge funds
- Data intensive, automated, high end computational platform
- Computes thousands of performance metrics per company, sector, region
- Uses 30 years of annual financial reports (and daily stock prices)
- Website with numerous tables, charts and data also accessible via APIs



Apple



Quest® value per share
US\$ 110.1

Share price
US\$ 96.1

Quest® valuation (%)
15

98 QRR percentile
/100




6 triAngle decile
/10



Cash Flow Returns



Capital Growth



Long Run Averages (%)

	Current*	3yrs	5yrs	10yrs
CFROC	20.9	24.2	27.2	23.6
CFROA	21.4	24.7	27.6	23.9

*Rolling 12-month forward

Long Run Growth Averages (%)

	Current*	3yrs	5yrs	10yrs
Invested capital (real)	4.7	18.8	32.2	36.8
Equilibrium growth	5.7	5.6	15.2	16.6

*Rolling 12-month forward

Quest limitations prior to migration

Reaching an End of Life: Lack of Scalability, Maintainability, and Evolution

Evolution of financial model limited

Increase in **frequency of full computations** from weekly to daily impossible

Increase in **#companies, #industries, #regions** impossible

Antiquated **web interface**

Infrastructure maintenance difficult and costly

New features impossible to implement

Why such limitations?

The Obsolescence of Legacy Infrastructure

Hosted In-House: in 3 dedicated bare metal servers and multiple VMWare instances

Microsoft Technology: 15 year history, 4 Windows 2003 Servers, 11 SQL Server DBs 2008, ASP.net, Access Clients,...)

Overly complex valuation methodology and software architecture, with very slow in-DB computations (T-SQL)

Single weekly full computing cycle daily differentials, little contingency for errors or outages

03 Migration

Challenges

Chronology

Infrastructure Comparison

Migration Challenges (specific to Quest)

Remote Team: Communication between teams (SenTai/CG IT and Business, London and Paris)

Data Provider Changes:

- From Thomson Reuters to S&P (Company Data)
- From FTSE ICB to MSCI GICS (Industry Data)
- Required the creation of a complex ETL (Extract, Transform, Load)

Data Management System:

- A proprietary language was developed to assess data quality and flag outliers
- An extensive admin web interface was designed to allow for possible overrides
 - Input: S&P data
 - Output: Override computations

Migration Challenges (Continued)

Business Logic Documentation:

- Limited
- Needed to write documentation to adequately transform Business logic
- 200+ computational phases and 7,000+ comp. rules

Valuation Model Documentation: Limited. Needed to reverse engineer and rewrite

Valuation Model and Computation Algorithm: entirely rewritten

Valuation Methodology Change

- From SQL server transactional based computations (T-SQL)
- to Google Sheet defined proprietary syntax handling time series
- translated into Java (Antlr), run on GCE

Validation Phase: Lasted during entire development Phase (one year)

Migration Chronology

Documentation, Model Rewrite, Software Architecture Reset, and Validation

Reverse Engineering GSuite (GA4W)

Transformation of legacy valuation model

From regular conference calls to
Google Hangout

From Emails and Office docs to
Google Apps and Google Docs
shared documents (cf. BBVA)

5 months



Prototype Development GCP/GCE

New valuation methodology

Validation Phase 1

Once prototype was convincing
enough, **CG was ready for Google
Cloud!**

7 months



Production GCP/GCE

Full model rewrite and web development

FMR based on new syntax/new
software architecture

New web 2.0 site

Extensive Validation Phase II
(full scale)

1 year



Infrastructure Comparison


GCP vs MS bare metal

Before (Not Scalable) Microsoft Technology

3 Bare Metal Servers, fixed storage, RAM, CPU
3 VMWare Windows Servers per bare metal server
SSD & SCSI
No app (O.S.) virtualization
11 SQL Server DBs
SQL Server Access Clients
No APIs

After (Scalable) Google Technology & Open Source

3 Google Compute Engine instances, scalable specs
Single Virtual Machines: 32 vcpus/core, 208 GB RAM
SSD only persistent disks
Docker containers
NoDB, File System based (200+ computational stages)
Big Query. Tableau, Google Sheets
API driven
1 Windows / SQL Server for S&P Data Sync (reqd. By S&P)



Google Cloud Platform

04 Benefits

Scalability

Performance

Agility

Risk Mitigation

Costs and Rationale

Benefits

Scalability

Before (10 hr)



Weekly computations (frequency)
20 year history
2,200 companies
10 regions
100 industrial sectors in 2 level structure
Inability to scale up (F,Y,C,R,I.S.)
Static valuation model
Antiquated website
No new features

After (30mn)



Daily computations (even intra-day)
30 year history (**1.5x**)
8,800 companies (**4x**)
25 regions (**2.5x**)
275 industrial sectors in 4 level structure (**2.75x**)
Scalable (F,Y,C,R,I.S.)
Dynamic valuation model
Web 2.0 user friendly
New features

Benefits

Performance



50 X



Google CloudPlatform

Data (Quantity/Speed)

2.5x the amount of computed data (input)

1/20x of the time (30mn vs 10 hrs)

10x more aggregate computations (output)

120 billion data points per full run

7,000+ complex business rules using non trivial mathematical and numerical procedures

200+ computational stages

Benefits

Agility - Flexibility of Valuation Model Changes

Before

Model Changes: minor modifications to the model required **1-2 months**. Changes could only be performed by chosen external vendor

Deployment by developers only:

- **in-DB model and computations:** entire compute engine carried out in DB calculations - SQL Stored Procedures (T-SQL) Very slow, difficult to fix model

After

Model Changes: **Live editable model online by analysts simultaneously** without programming knowledge

Deployment of new model done by business analysts themselves!

- **Google Sheet:** Rewrite model within GS, adding a proprietary syntax for handling of Time-Series
- **Model Translation:** Into Java code from Google Sheet use ANTLR (language builder)
- **Java based computations on GCE**

A man with dark hair and glasses, wearing a checkered shirt, is looking down at a white tablet he is holding. He is standing in a server room, with rows of server racks visible in the background. The racks contain various electronic components and bundles of green and blue cables. The lighting is somewhat dim, typical of a server room.

Risk Mitigation



Security

Several layers of security provided by Google infrastructure



Backups

Use of high speed Disk Cloning and Imaging. Several layers of backup processes available via managed GCP



Disaster Recovery

Automatically redeploy from scratch entire production environment **in 5mn**, via single command line. **24h previously**



Confidentiality

Customer information is a part of CG's customer base, and minimal data is pushed to cloud platform for credentials, site preferences, site permissions



Intellectual Property

Permissioning (sharing) & restrictions to access to Valuation Model handled by Google Apps



Cloud Restrictions

Cloud restrictions imposed by S&P (financial data provider) required separate in-house SQL server

Rationale for Canaccord Genuity

Cost Reduction (4X) - Benefits

Total running costs vs. legacy costs

Infrastructure costs @ GCP: Low monthly costs (less than \$6500/month)

Maintenance: No maintenance of bare metal servers hosted in Canaccord datacenter

Autonomy: Canceling of costly technical outsourcing contract to maintain legacy in-house infrastructure

Resources: Reduction of data cleansing team by automation rules (ETL). Allows to reallocate resources

Benefits

Performance: Stability, scalability, maintainability, evolution, agility, security

Integration: API technology between GSuite & GCP allowed to develop a fully integrated and efficient environment for both business and technology needs

Autonomy: Recover technical knowledge: appropriation of platform by in-house developers trained by us

Resources: Focus its resources on core business (i.e. model improvement)

Successful Relaunch

Institutional Investor Feedback and Usage



- Investment in **Quest** has created a **state of the art platform** in the cloud
- **Relaunch has been a huge success.** Already we are witnessing increased broker votes in, and incremental revenues.

Portfolio Manager:

"It's an excellent tool and we love it. It fits all of our requirements and we appreciate all of the hard work done by the Quest® team."

"New Quest® is now better than Holt."

Senior Portfolio Manager:

"In my personal opinion, the system really is outstandingly powerful. The system is very quick and looks very good. This is a website capable of winning an award for speed, user-friendliness and - if I may say – intelligence of icon design / subtle function prompt."

Director of European Equities:

"CITN (Quest Feature) is the best product on the street by a country mile and new Quest® is a fantastic tool."

Research Analyst:

"Quest® product truly excellent. Product upgrade looks great."

Fund Manager:

"Quest® refresh excellent."



“With Google Cloud Platform, our product has deeper data, runs better analysis, is more customisable, handles more than 2.5x the load, and runs 20x faster than the old version.”

Rob Garwood

Head of IT Development, Canaccord Genuity

CANACCORD Genuity

Performance gains overview



100M+

Google BigQuery grinds through over 100 million data rows in only a few minutes, at speeds at least 25x faster than the previous system.



9,000

Analyzes 30 years of data for almost 9,000 companies, 275 aggregate industrial sectors, 25 regions.



Scale

Migrated from a 15-year-old version that couldn't scale.



120Bn

Over 120 billion data points calculated overnight, and now several times per day (30 minutes).

Google Cloud Platform

Few Financial Services Companies have taken the leap into the cloud for key projects!

Canaccord Genuity is one of them through the Quest platform:

- SenTai has adapted Quest to the cloud : leaving a Microsoft only, bare metal infrastructure, to Open Source and Google Cloud Platform, all at once!
- **SenTai** is uniquely positioned to accompany you with your Cloud migration or development projects on GCP, regardless of the industry

Read the full case study: <https://sentai.eu/quest/>
<https://cloud.google.com/customers/canaccord-genuity/>

Thank You!

david@sentai.eu

ANNEX

(Technical)

Data Availability

Like most financial applications, Quest is driven by the ways in which the processed data is made accessible to end users, and available for analysis:

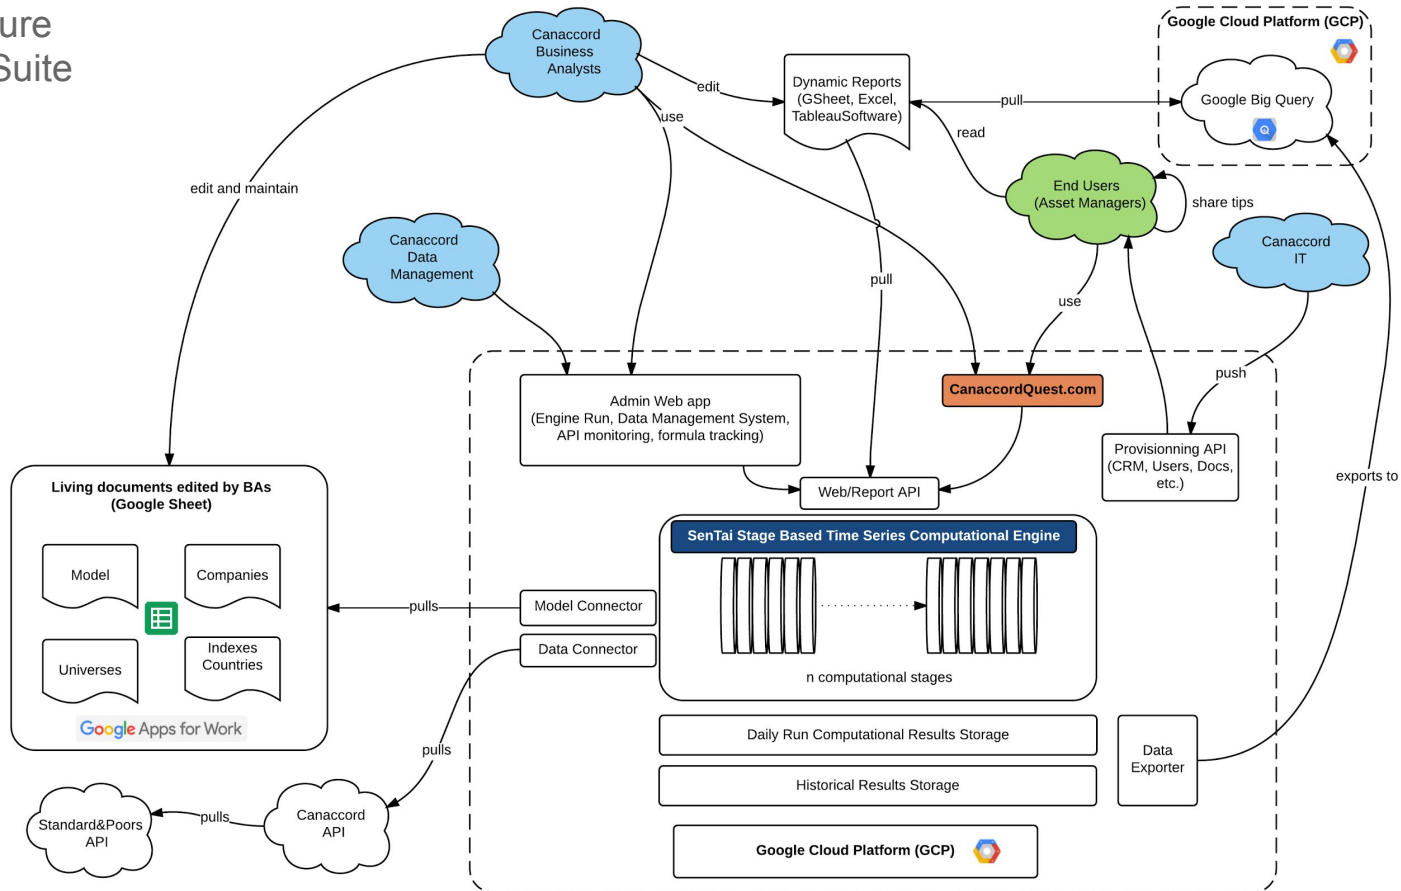
All users (Internal & External)

- Website (primary use)
- Download (xls, csv)
- [Big Query](#), TableauSoftware...
- API (access limitations for external clients)

Internal users only (Canaccord Genuity)

- [Google Sheets](#) and centralized [Google scripts](#) (using REST APIs and [Google Managed Library](#))
- Web Admin Interface

Architecture GCP & GSuite



Technology Stack (software)

- Docker Containers (OS virtualisation, portability and ease of deployment)
- Debian 8 (Linux operating system)
- Ansible, Jenkins, Hubot, Nginx (provisioning, automation, deployment, proxy)
- Java 8 (programming language), Antlr (language builder)
- **NoDB**, entirely File System based to optimize 200+ stage computational speed.
- StackDriver (monitoring), Swagger (API documentation)
- Web 2.0 (REST, Bootstrap, AngularJS, Coffeescript, Javascript)
- Post computational analysis: Google Big Query, Google Sheets, Tableau, APIs
- (GSuite/Non GCP): Google Sheets, Docs, Hangouts, Keep, Calendar, Drive
- (SaaS): Atlassian Jira, BitBucket, Slack, Trello (bug tracking, Git DVCS, collaboration)