Today's Host

David Ronald
Director of Product Marketing

- BSc in Applied Physics from Strathclyde University, MSc in Optoelectronic & Laser Devices from St Andrews

- Prior work in artificial intelligence, natural linguistic programming and telecommunications technology

- 18+ years in tech industry

TigerGraph

Today's Presenter

Richard Henderson
Lead Sales Engineer

- 18+ years in data architecture
- Supply-chain management expertise gained by working at Toyota and elsewhere
- Favourite expression, "Pushing data like it's on fire"

TigerGraph

Some Housekeeping Items

- Although your phone is muted we do want to answer your questions - **submit your questions at any time** using the Q&A tab in the menu
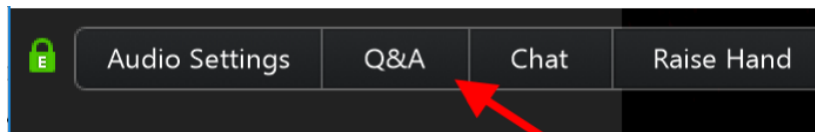
- The webinar is being recorded and will uploaded to our website shortly (https://www.tigergraph.com/webinars/) and the URL will be emailed you

- If you have issues with Zoom please contact the panelists via chat

TigerGraph

Value: Forecast versus Orders Supplier Impact

- A large manufacturer identified that they would benefit (potentially by **£10millions**) from a timely analysis of impact to their supply-chain of changes to their forecast orders.
  - Sales forecasts are typically years in advance so suppliers can tool-up
  - Minimum buy volumes are committed from forecast to support investment
  - Demand can vary widely and quickly from the forecast
  - Costs to the business can significantly impact margins
  - Having good information allows the executive to put mitigations in place

TigerGraph

Proving the Value of TigerGraph

Benefits of a pilot:

Rapid deployment of graph to understand the costs & benefits (typically 4-weeks)

Brings the key stakeholders together so they understand the positive impact this solution can have

We work alongside your team so they have a realistic basis for assessing TigerGraph

We map the success criteria to the results as an input to the technology selection process

Plan & Initiate	Get the Data	Process the data	Show Value	Build the business case
Ph 1: Setup		**Ph 2: Process**		**Ph 3: Case**

Key Steps:

Plan & Initiate	Get the Data	Process the data	Show Value	Build the business case
We plan the project ahead of arrival. A kick-off event agrees the requirements, success-criteria, stakeholder map, the ways of working, aligns diaries and agrees checkpoints	We analyse the the data sources you have, design the schema and ensure data can be loaded via any necessary data-engineering	We load the data and process it in the graph to check quality, enhance connectivity and add structure to serve the functional queries.	We design and code the queries, check for correctness, optimise for performance. Each PoC proof point is checked to validate for the next phase. Sizing and architecture is generated.	We generate results versus success criteria. We use outputs to agree sizing and proposed architecture. We assist team to present results and answer any questions.

**TigerGraph**

The Plan

Business stakeholders were identified and polled to identify overarching concerns

- Infrastructure: Enterprise architecture and IT Ops engaged to ensure solution fit into cloud-based strategy
- Strategic Alignment: Head of data engaged to ensure strategic alignment
- Business Engagement: Departmental teams engaged to validate business context and provide data
- Technical Engagement: Head of data-science provided two resources to work with TigerGraph over 2 weeks

Specific deliverables with success criteria were identified

- Value: Generate the supply-chain report with the highest value identified
- Integration: Show integration with existing BI toolset
- Data Input/Scalability: Data transformation, loading and reporting performance to be demonstrated
- Resourcing/skills: Proving that existing data-science resources could quickly add TigerGraph to their suite of tools

Infrastructure: Direct Integration with GCP
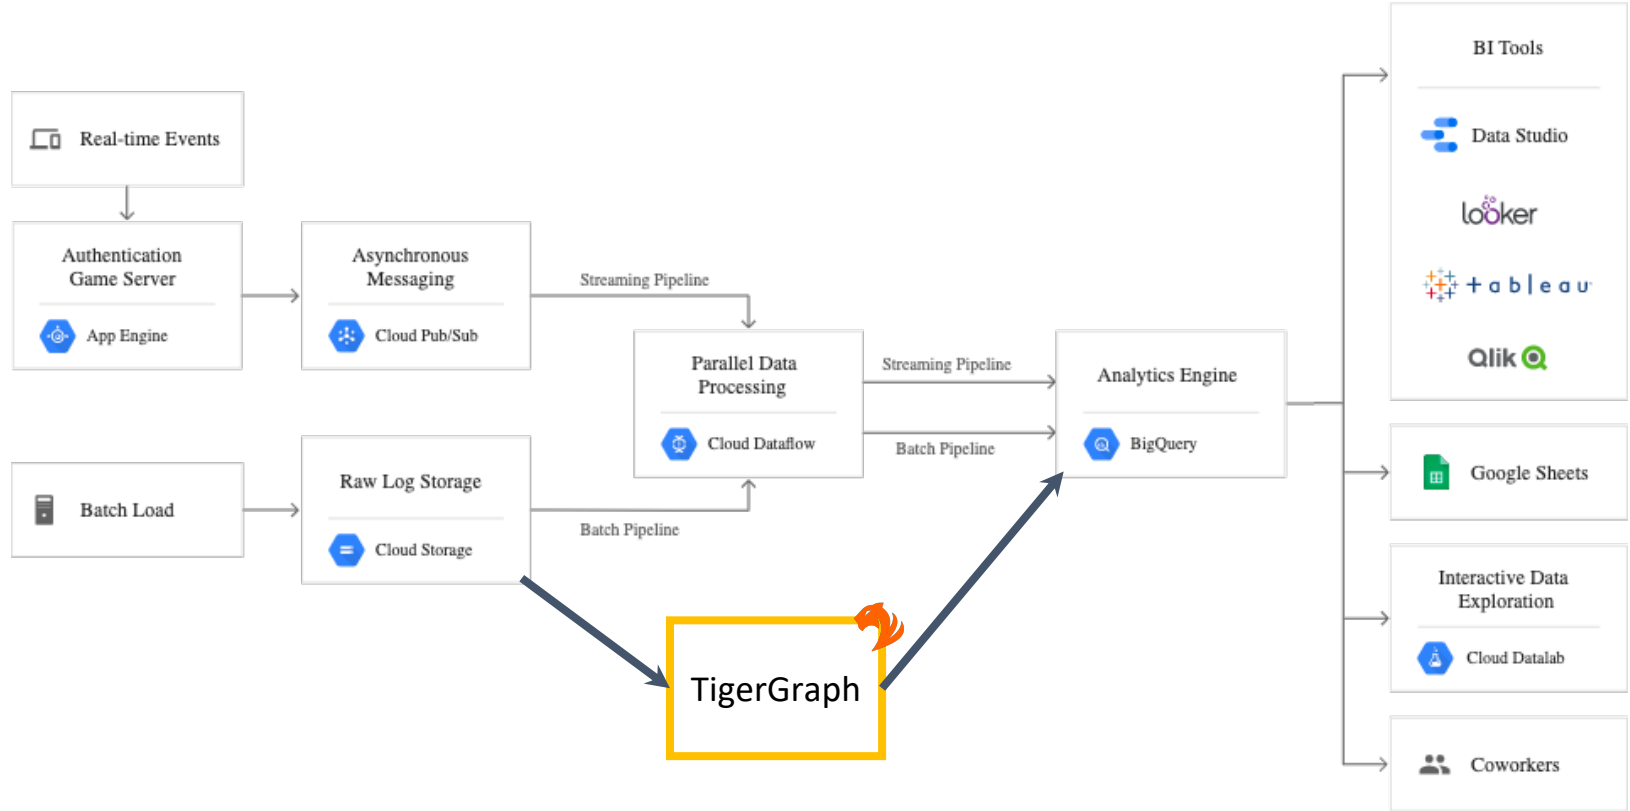
# Data Integration and Scalability

- 1 year's order forecast across two markets

- 8 complex data sources from multiple departments loaded into graph schema with 11 vertex types and 14 edge types relating them (see schema view)

- Used a combination of existing data engineering (in Python), shell scripts and TigerGraph native parallel loading

- Load took around 2 minutes for a few million vertices and a few tens of millions of edges including pre-processing (~2.5GB raw csv)

- Post-processing of graph (to about 1.5 billion derived edges) took 30 minutes

- Hardware was a single 8 core cloud server with 32GB of memory

TigerGraph

Use Existing Resources

- Worked with two members of data-science team

- Certified in GSQL1011 before we arrived

- Infra team installed software themselves with only limited need for help

- Schema designed and data transformed and loaded with off-site support from us

- Queries designed and run by prospect team while we mentored on-site

- Prospect team implemented BI tool integration and report generation completely by themselves

- Full-scale data optimization performed by us in the following week

TigerGraph

Q&A

Please submit your questions via the Q&A tab in Zoom
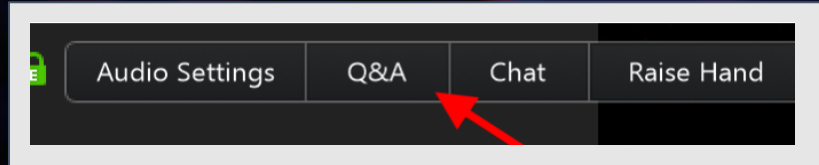
More Questions?

Join our Developer Forum

https://groups.google.com/a/opengsql.org/forum/#!forum/gsql-users

Sign up for our Developer Office Hours (every Thursday at 11 AM PST)

https://info.tigergraph.com/officehours

TigerGraph

Additional Resources

Start Free at TigerGraph Cloud Today!

https://www.tigergraph.com/cloud/

Test Drive Online Demo

https://www.tigergraph.com/demo

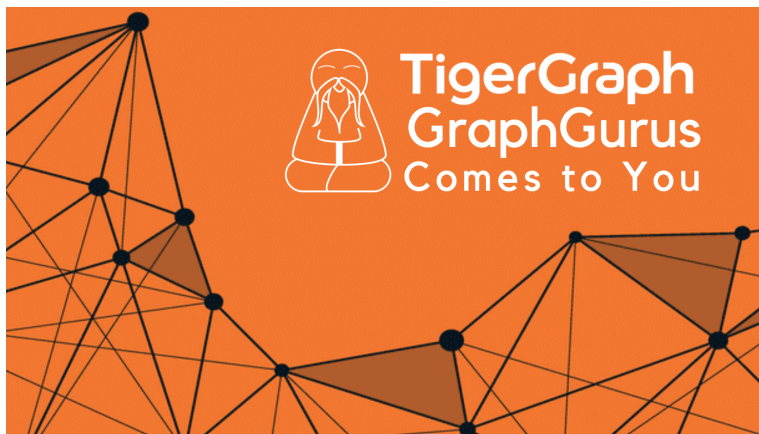**Download the Developer Edition**

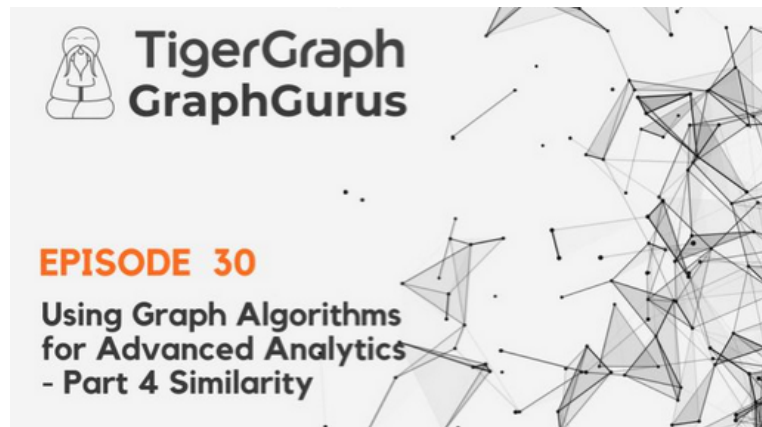https://www.tigergraph.com/download/

Guru Scripts

https://github.com/tigergraph/ecosys/tree/master/guru_scripts

TigerGraph

Upcoming Graph Guru Events

Coming to Atlanta, Charlotte, London, and more. View all events and request your own here:

https://www.tigergraph.com/graphguruscomestoyou/

Graph Gurus 30: Using Graph Algorithms for Advanced Analytics - Part 4 Similarity

Wednesday, March 4, at 11am PST

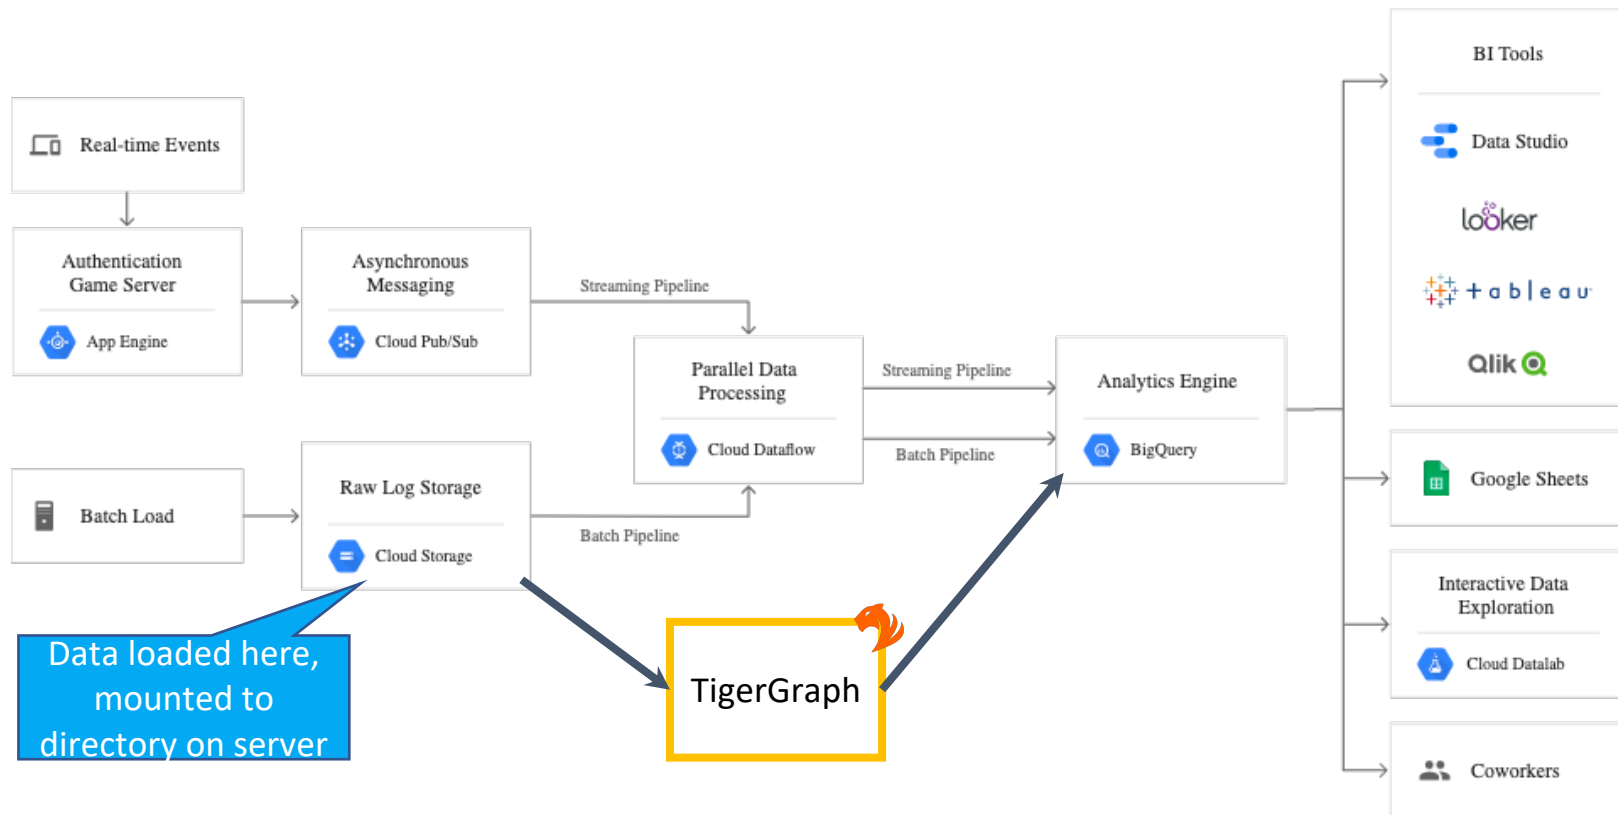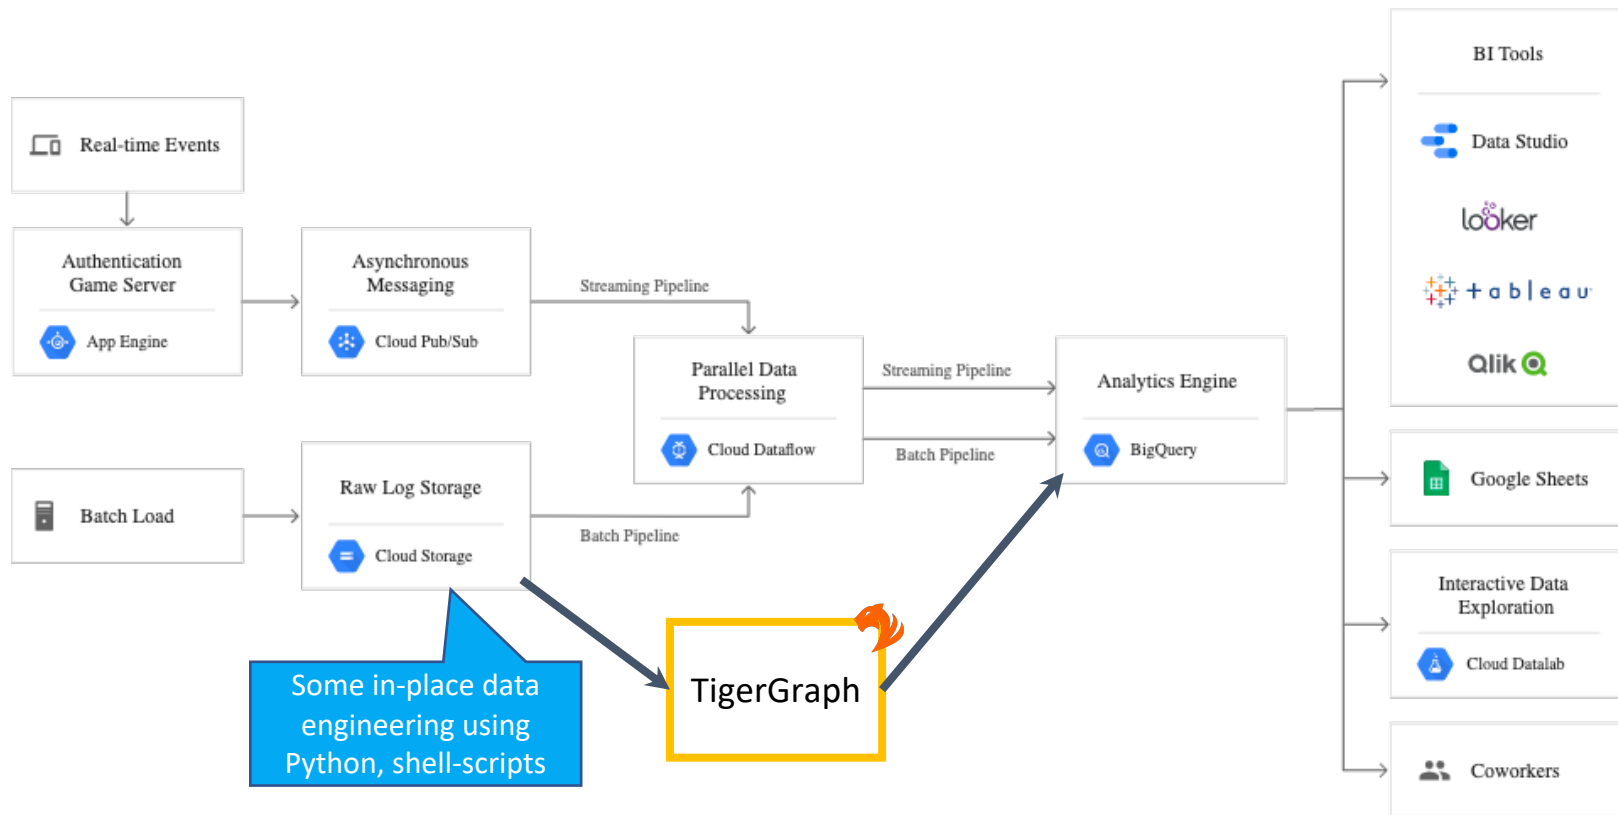https://info.tigergraph.com/graph-gurus-30

Thank You

Appendices

- Some additional technical/architectural commentary

Detailed Technical Work

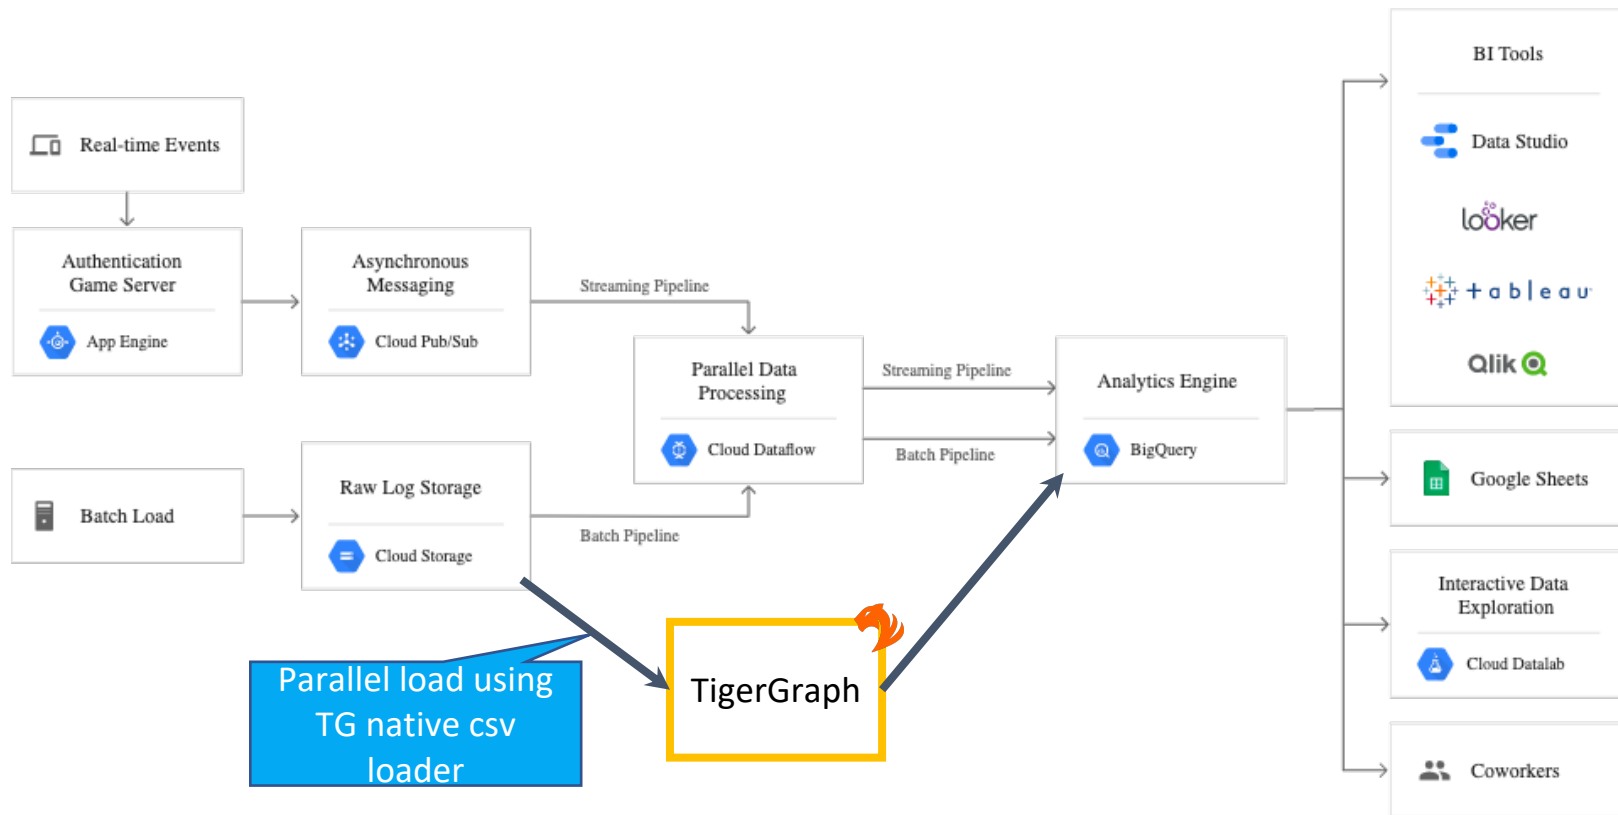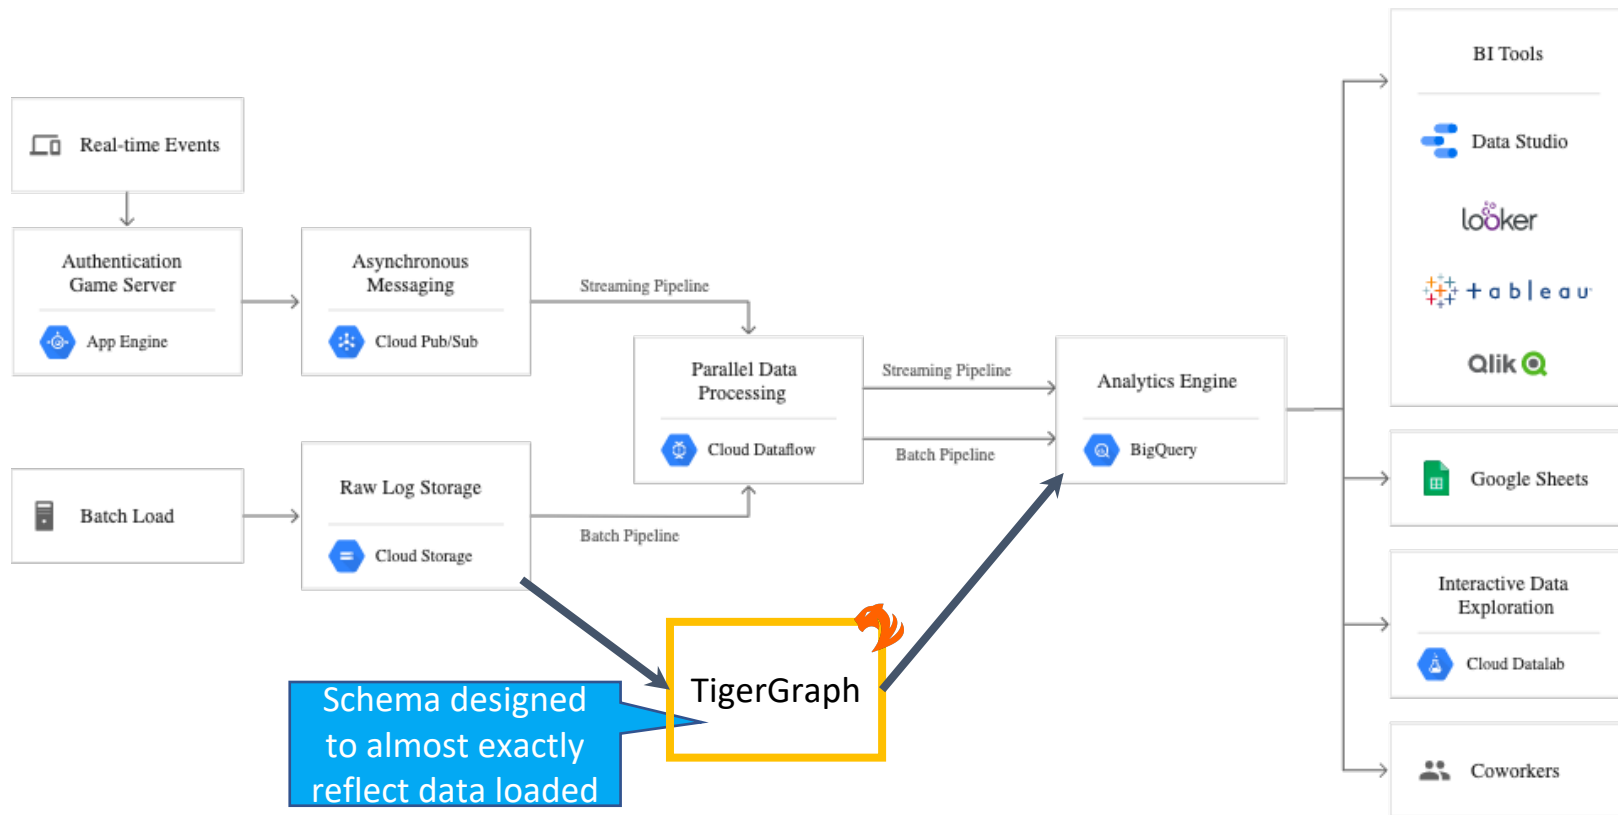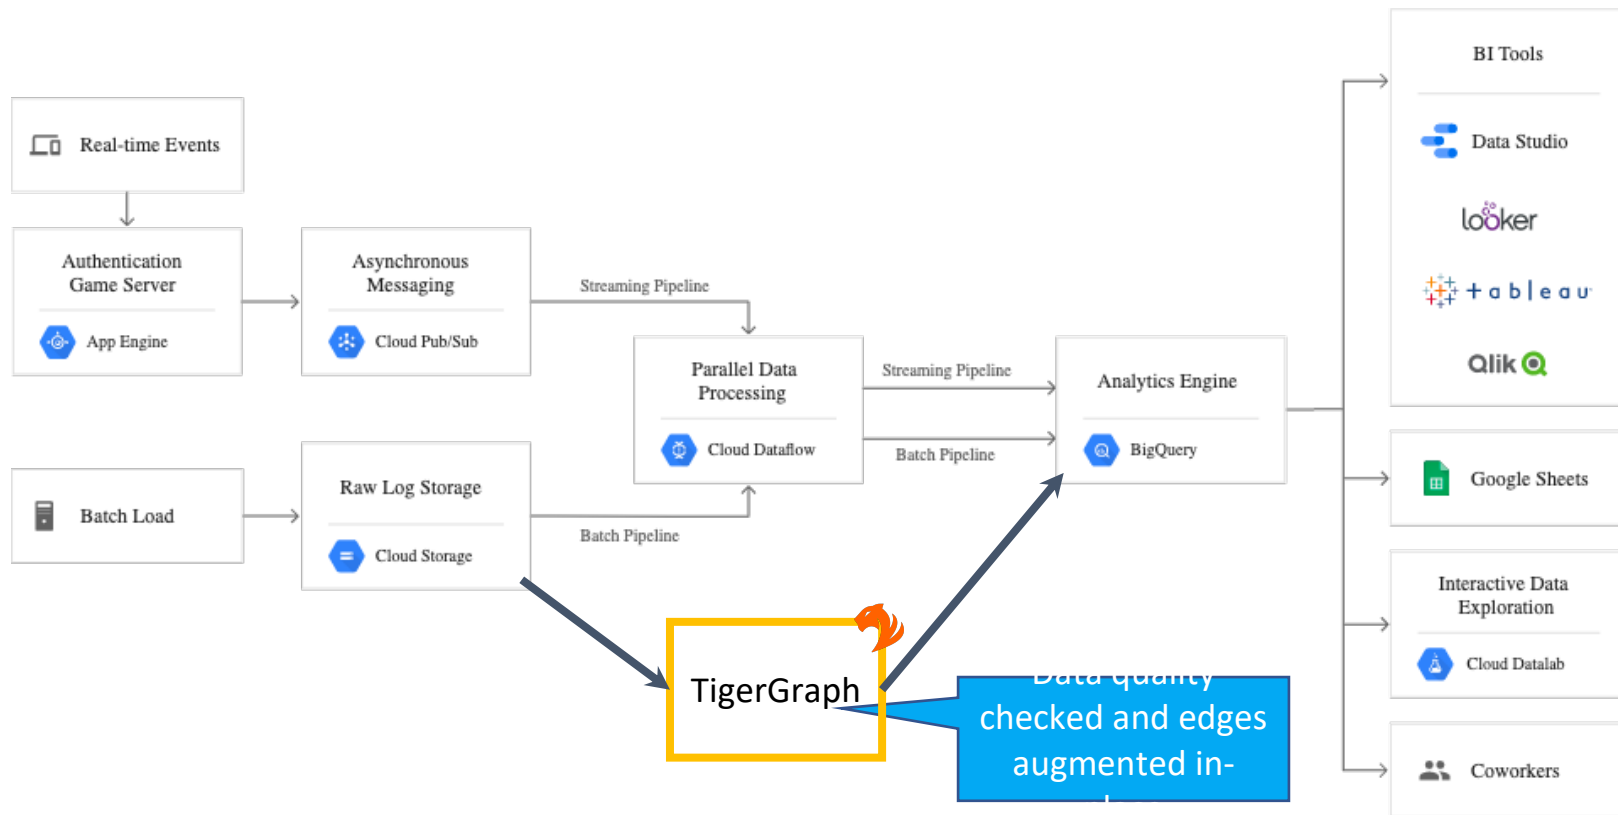Data loaded here, mounted to directory on server

TigerGraph

Detailed Technical Work
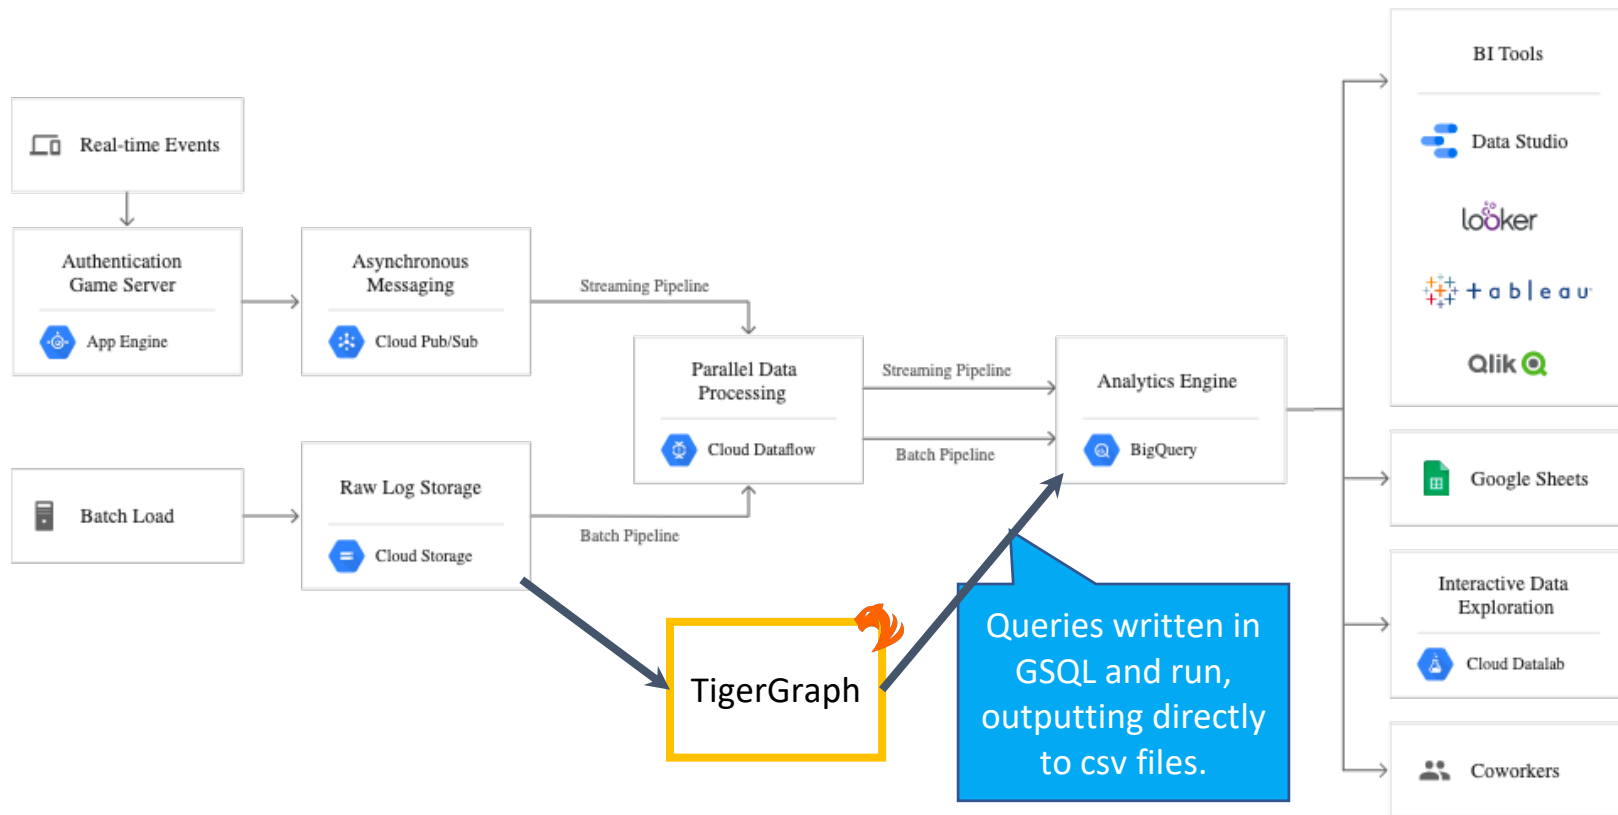
Detailed Technical Work

Detailed Technical Work

Detailed Technical Work

Detailed Technical Work

Detailed Technical Work

Real-time Events

Authentication
Game Server

App Engine

Asynchronous
Messaging

Cloud Pub/Sub

Streaming Pipeline

Batch Load

Raw Log Storage

Cloud Storage

Batch Pipeline

Parallel Data
Processing

Cloud Dataflow

Streaming Pipeline

Batch Pipeline

Analytics Engine

BigQuery

BigQuery directly imported reporting output into existing database, augmenting existing data.

TigerGraph

BI Tools

Data Studio

looker

tableau

Qlik

Google Sheets

Interactive Data
Exploration

Cloud Datalab

Coworkers

TigerGraph

Detailed Technical Work

BI tools allowed flexible reporting of augmented data-set

Real-time Events

Authentication Game Server
App Engine

Asynchronous Messaging
Cloud Pub/Sub

Streaming Pipeline

Parallel Data Processing
Cloud Dataflow

Streaming Pipeline

Analytics Engine
BigQuery

Batch Pipeline

Batch Load

Raw Log Storage
Cloud Storage

Batch Pipeline

TigerGraph

BI Tools
Data Studio
looker
+ableau
Qlik Q

Google Sheets

Interactive Data Exploration
Cloud Datalab

Coworkers

TigerGraph

Future Work

Extend data–quality checking to ensure configuration matched part availability and was complete

Add automated delta capability so can perform monthly upload (less than a minute including graph processing)

Extend queries to what-if simulation

Extend queries to analyse unit profitability down to order/part level

Extend queries to analyse sensitivity of supply-chain links in terms of resilience, and with mitigation costs

Extend queries to "next best investment" top-10 ROI in supply-chain

TigerGraph