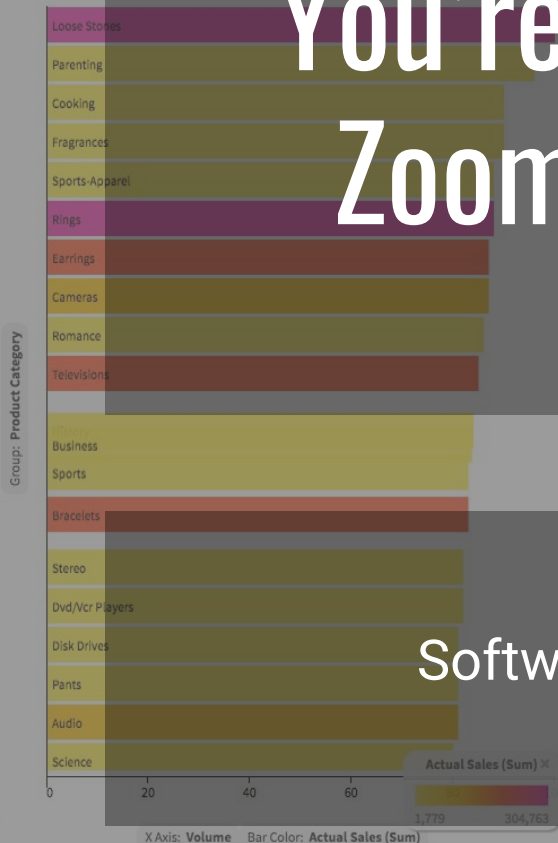




Product Category by Volume

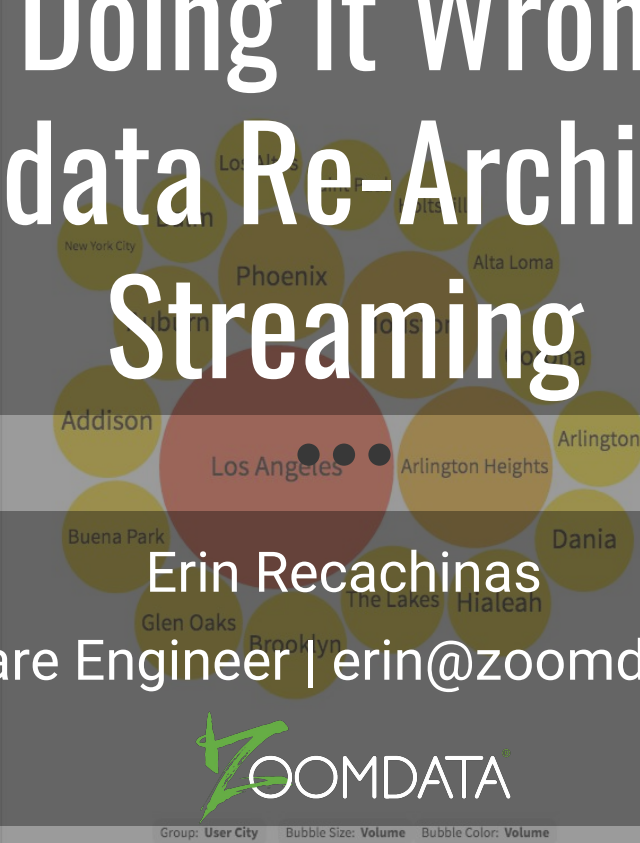


You're Doing it Wrong: How Zoomdata Re-Architected Streaming

Erin Recachinas

Software Engineer | erin@zoomdata.com

Real Time Sales



Actual vs. Planned Sales by Minute



Add New Chart

1/2: Select Data Source

- Real Time Sales
- Hotel Reviews
- Lending Data
- loans
- Micro Queries - Impala
- MongoDB Listing IDs Fusion
- My IMPALA Source
- MySQL Event Fusion
- S3 Sellers Fusion
- Ticket Marketplace Fusion
- Ticket Sales Table Fusion
- Twitter Data Source - Solr
- Vehicle Complaints





- Full stack engineer @ Zoomdata
- Built new streaming architecture
- #1 in Dresner Big Data Analytics Market Study 2016 & 2017

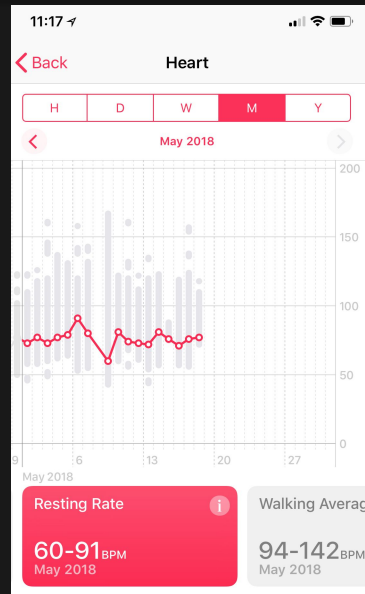
Use Case

Monitor



Single
Status
Alerts

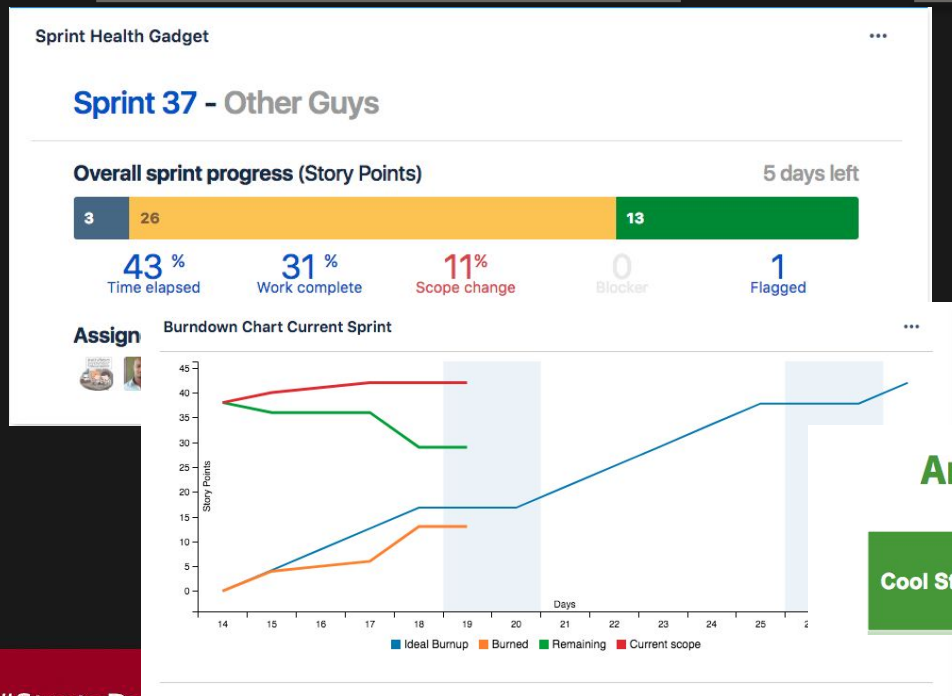
Aggregate



Multiple
Tracking
Analytics

Use Case

Monitor



Aggregate



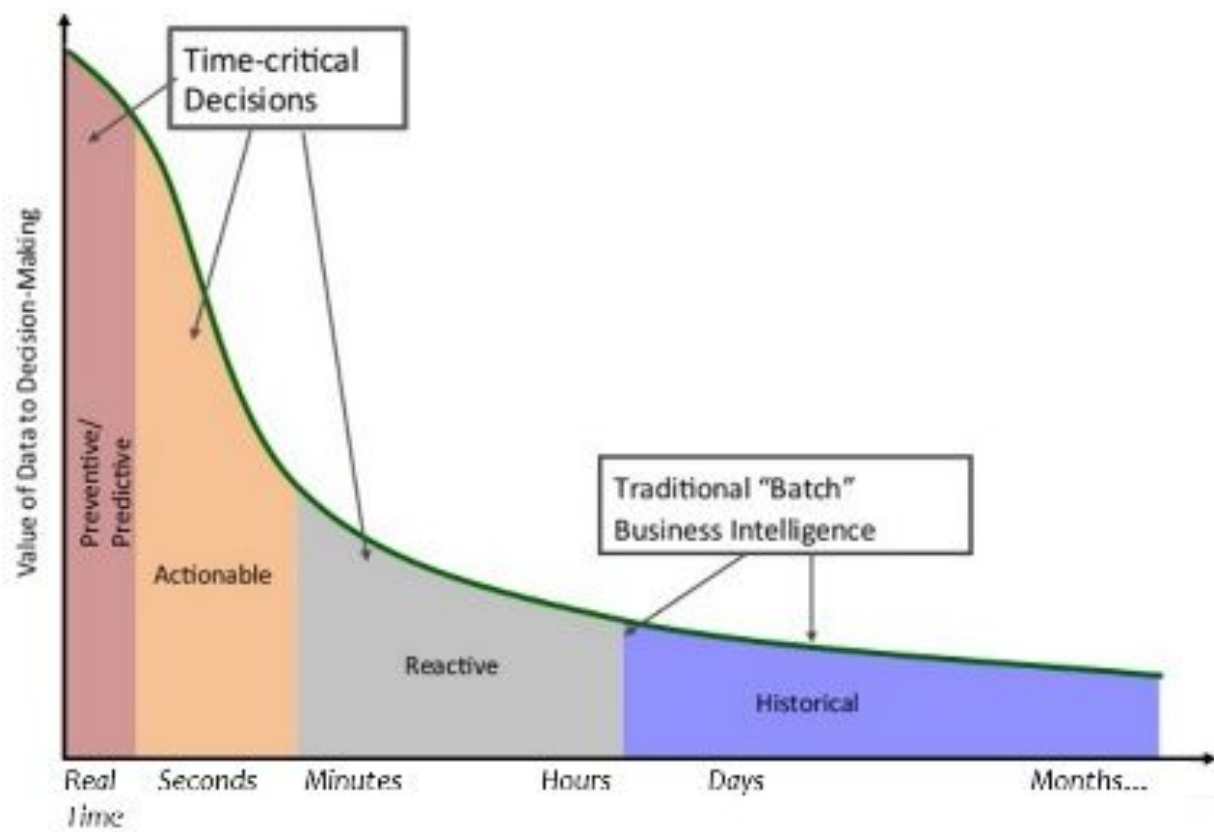
Architecture Roadmap (July 2017 Update)

Cool Stuff

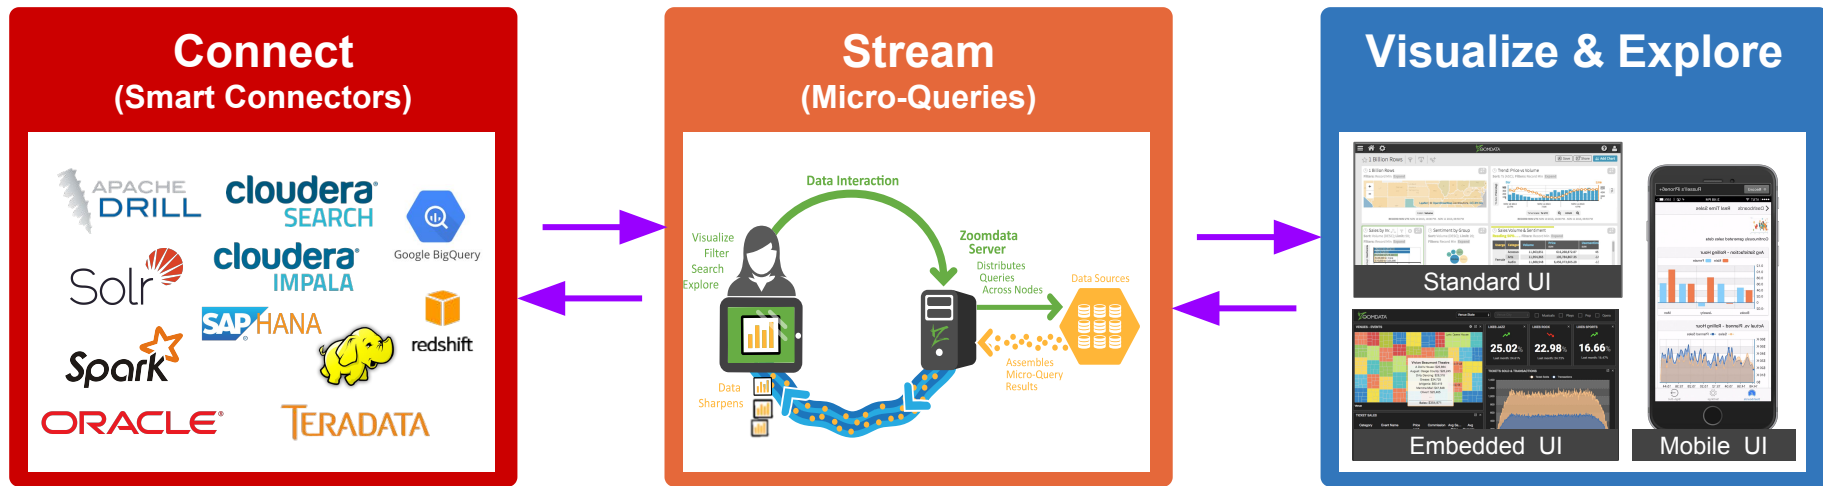
A better BI
Engine

Awesome
Stuff

Even More
Awesome
Stuff



Modern BI Components

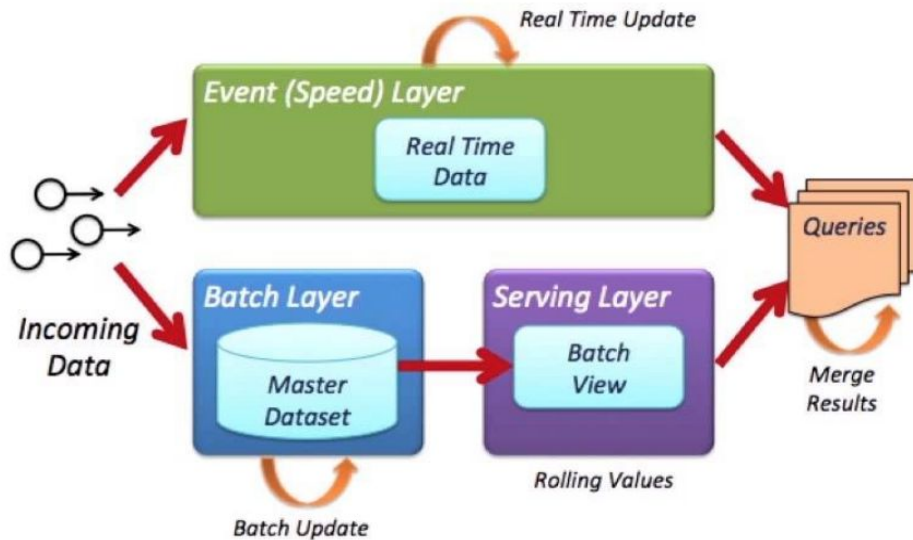


Traditional Streaming Architectures

Stream Processing



Lambda
(Stream + Batch)



How did we know we had a problem?

Customer use case

- Telecomm company
- 330M Subscribers
- 10 Billion transactions per day
- 4 Continents

Business Need


Single Platform

Real Time Analytics

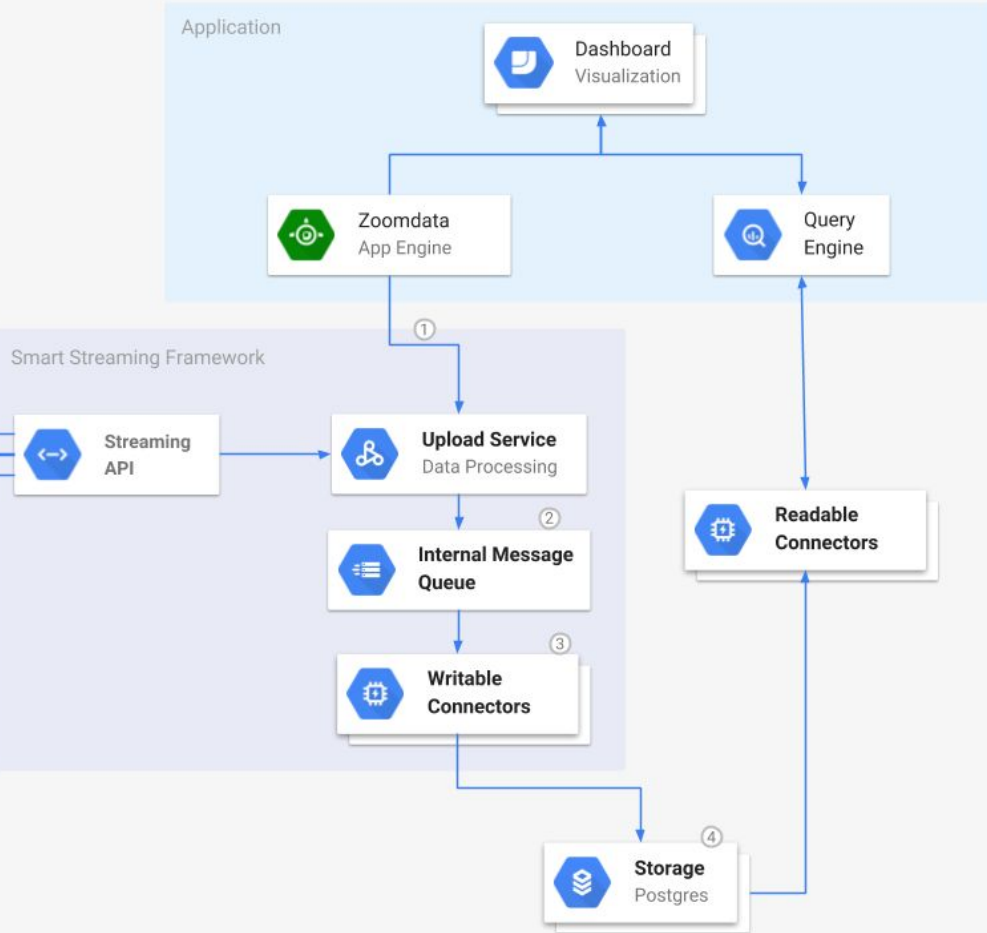

→ Velocity
→ Volume



- Data Storage
- Data Writing



Zoomdata Infrastructure



It was perfect and they lived happily ever after

~~It was perfect and they lived happily ever after~~

Lessons Learned

- Asynchronous vs Synchronous Workflows
- Stateless vs Stateful Microservices
- Data source security
- Achieving extensibility

Q & A



eatpraycodeblog.com/strata18

erin@zoomdata.com