Big Data Infrastructure at Spotify

Wouter de Bie

Team Lead Data Infrastructure

June 12, 2013

Agenda

Let’s talk about Data Infrastructure, how we did it, what we learned and how we’ve failed

- Some **Context**
- Why **Data?**
- **Use** Cases
- Our **Infrastructure**
- Lesson **learned**

Spotify? Spotify!

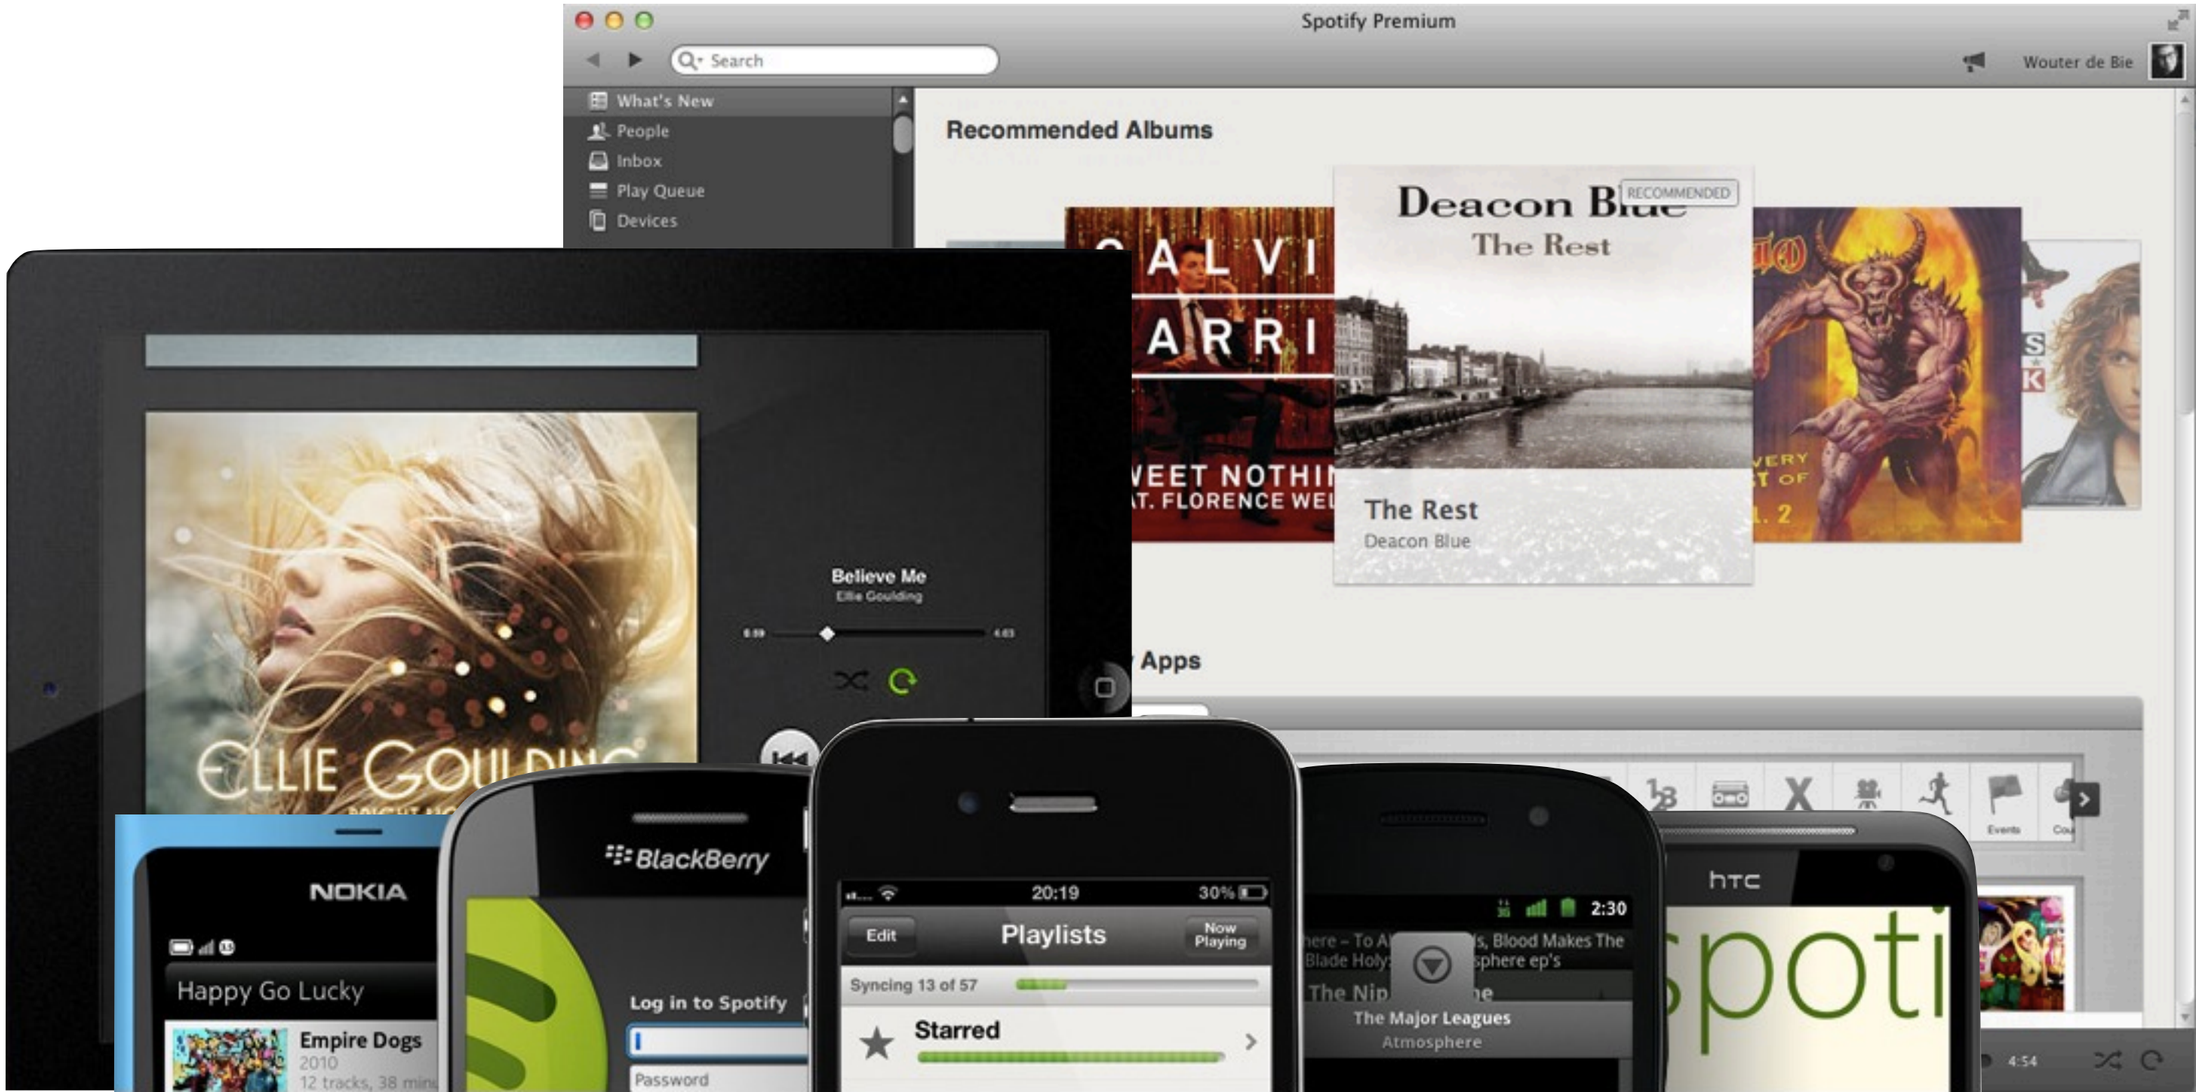

Some Context

- Spotify started in **2006**
- Now **850+** employees, **250+** engineers
- **26 million** monthly active users
- **20+ million** tracks available
- **4** data centers across the globe
- **12** data engineers building a **platform** for easy access to data

Why data?

We play music, right?

Why data?

We were the **first** company to do free music **streaming**. But now **everybody** can do it.

Reporting
Business Analytics
Operational Analytics
Product features

Use Cases

We’re a data-driven company, so data is used almost everywhere

Reporting

- Reporting to **labels**, **licensors**, **partners** and **advertisers**
- We support **our partners**

Business Analytics

- Analyzing **growth**, **user behavior**, **sign-up funnels**, etc
- Company **KPIs**
- **A/B** testing
- **NPS** analysis
- **Segmentation** analysis

Operational metrics

- **Root** cause analysis
- **Latency** analysis
- Better **capacity planning** (servers, people, bandwidth)

Product features

- **Radio**
- **Top** lists
- **Recommendations** (better then external parties, because of the amount of data)


Everybody should be
able to use data!
Spotify
June 12, 2013

So why Data Infrastructure?

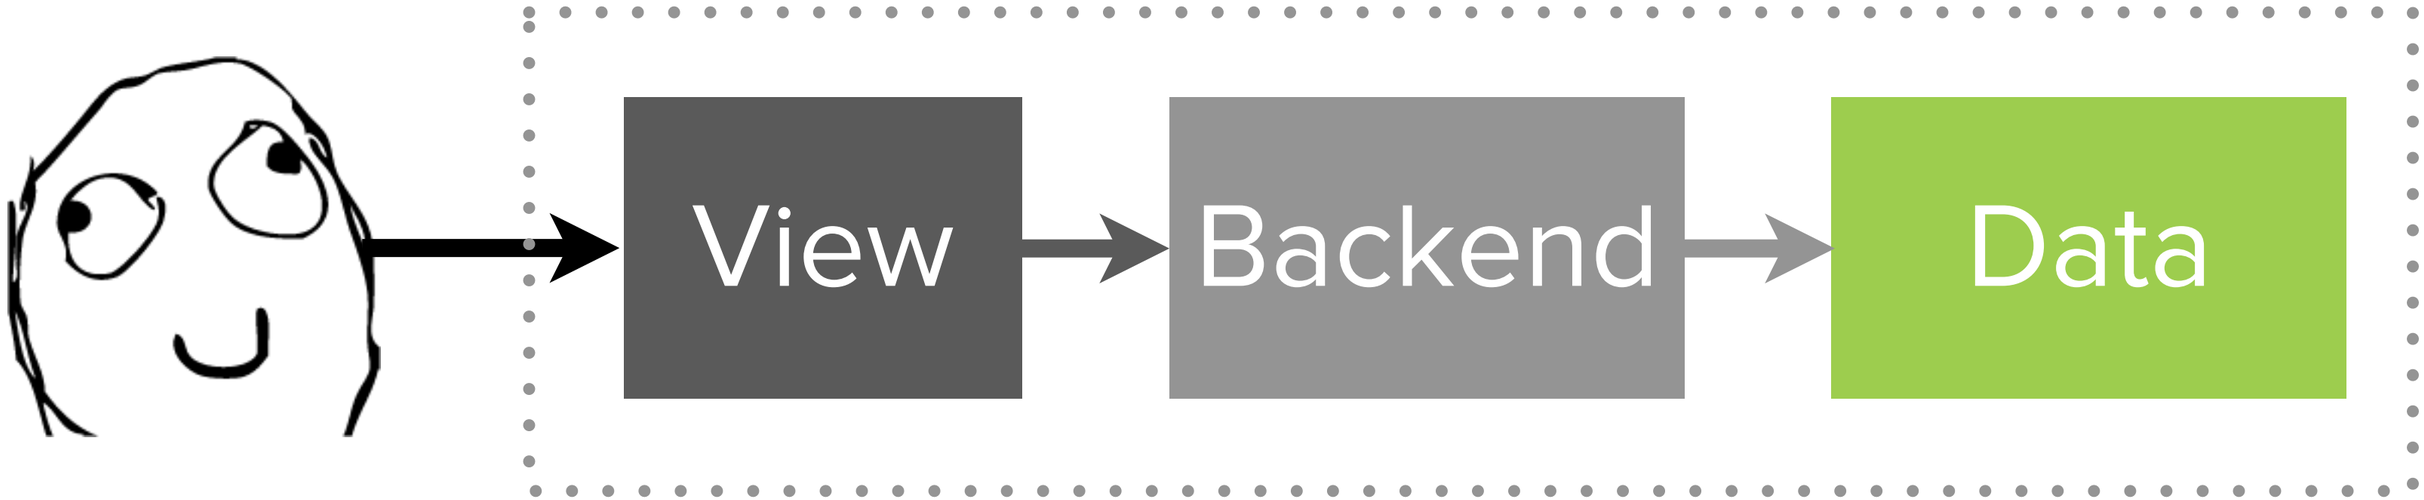

Our data infrastructure

Some geeky numbers

- **600 GB** of compressed data from users per day
- **150 GB** of data from services per day
- **4 TB** of data generated in Hadoop each day
- **190 node** Hadoop cluster
- Soon **690 nodes**
- **4 PB** of storage capacity (soon **28 PB**)

Spotify's data infrastructure

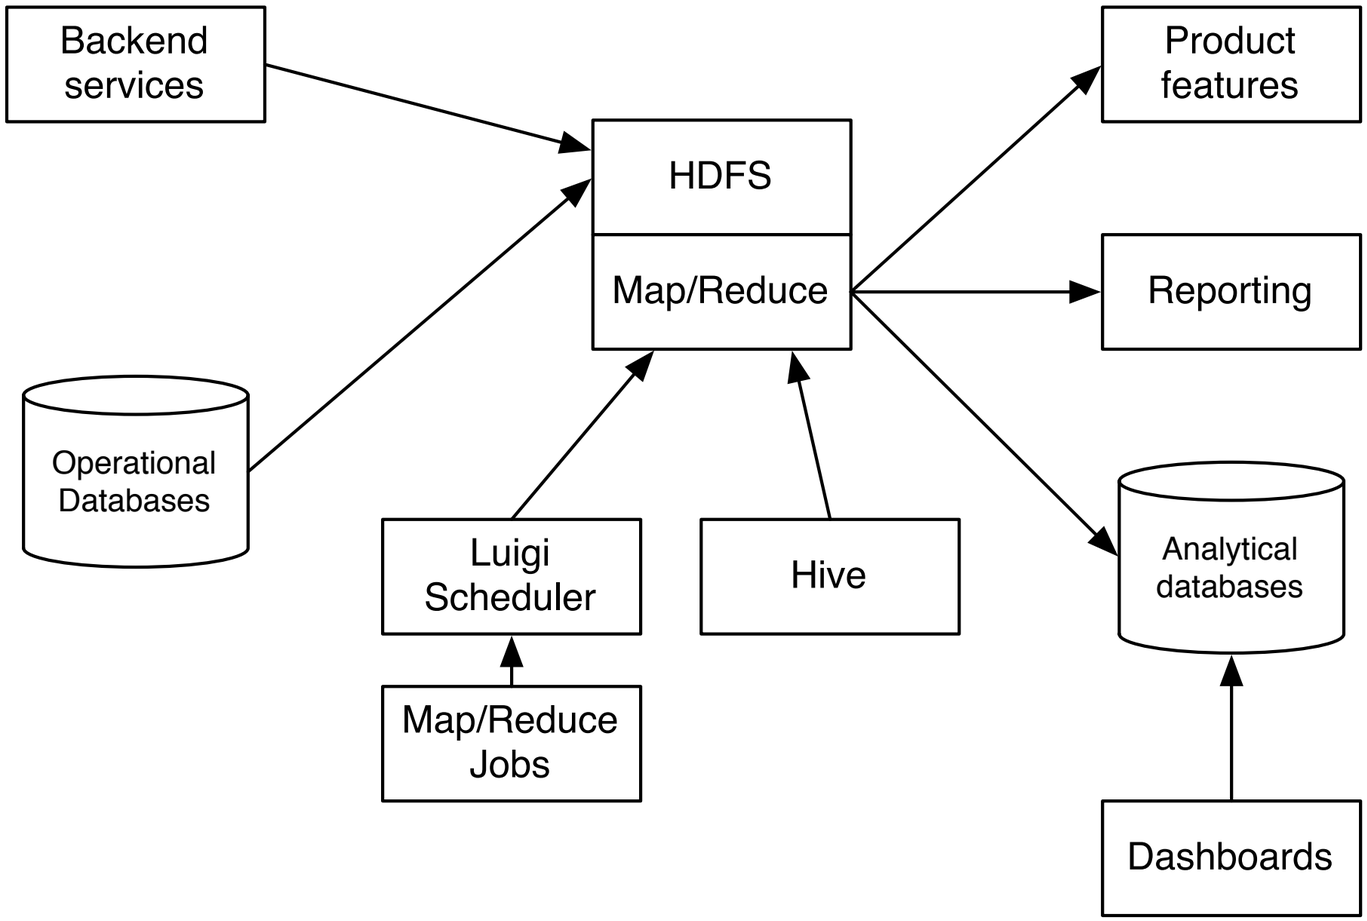

The thee pillars of our Data Infrastructure

Kafka

Collection

Hadoop

Processing

Databases

Analytics/Visualization

Data collection

Spotify

June 12, 2013

Data collection

Kafka: High volume pub-sub system

- Started with a store-and-forward system
- Evaluated Apache Flume
- Currently from Backend-to-HDFS, but in the future Backend-to-Backend
- It was almost a good fit, but...

Guaranteed delivery

Kafka doesn’t provide message acknowledgements..

- ... at least, not in 0.7 (stable)
- 0.8 has support, but no end-to-end acknowledgements
- A track streamed == monetary transaction

Hacking Kafka

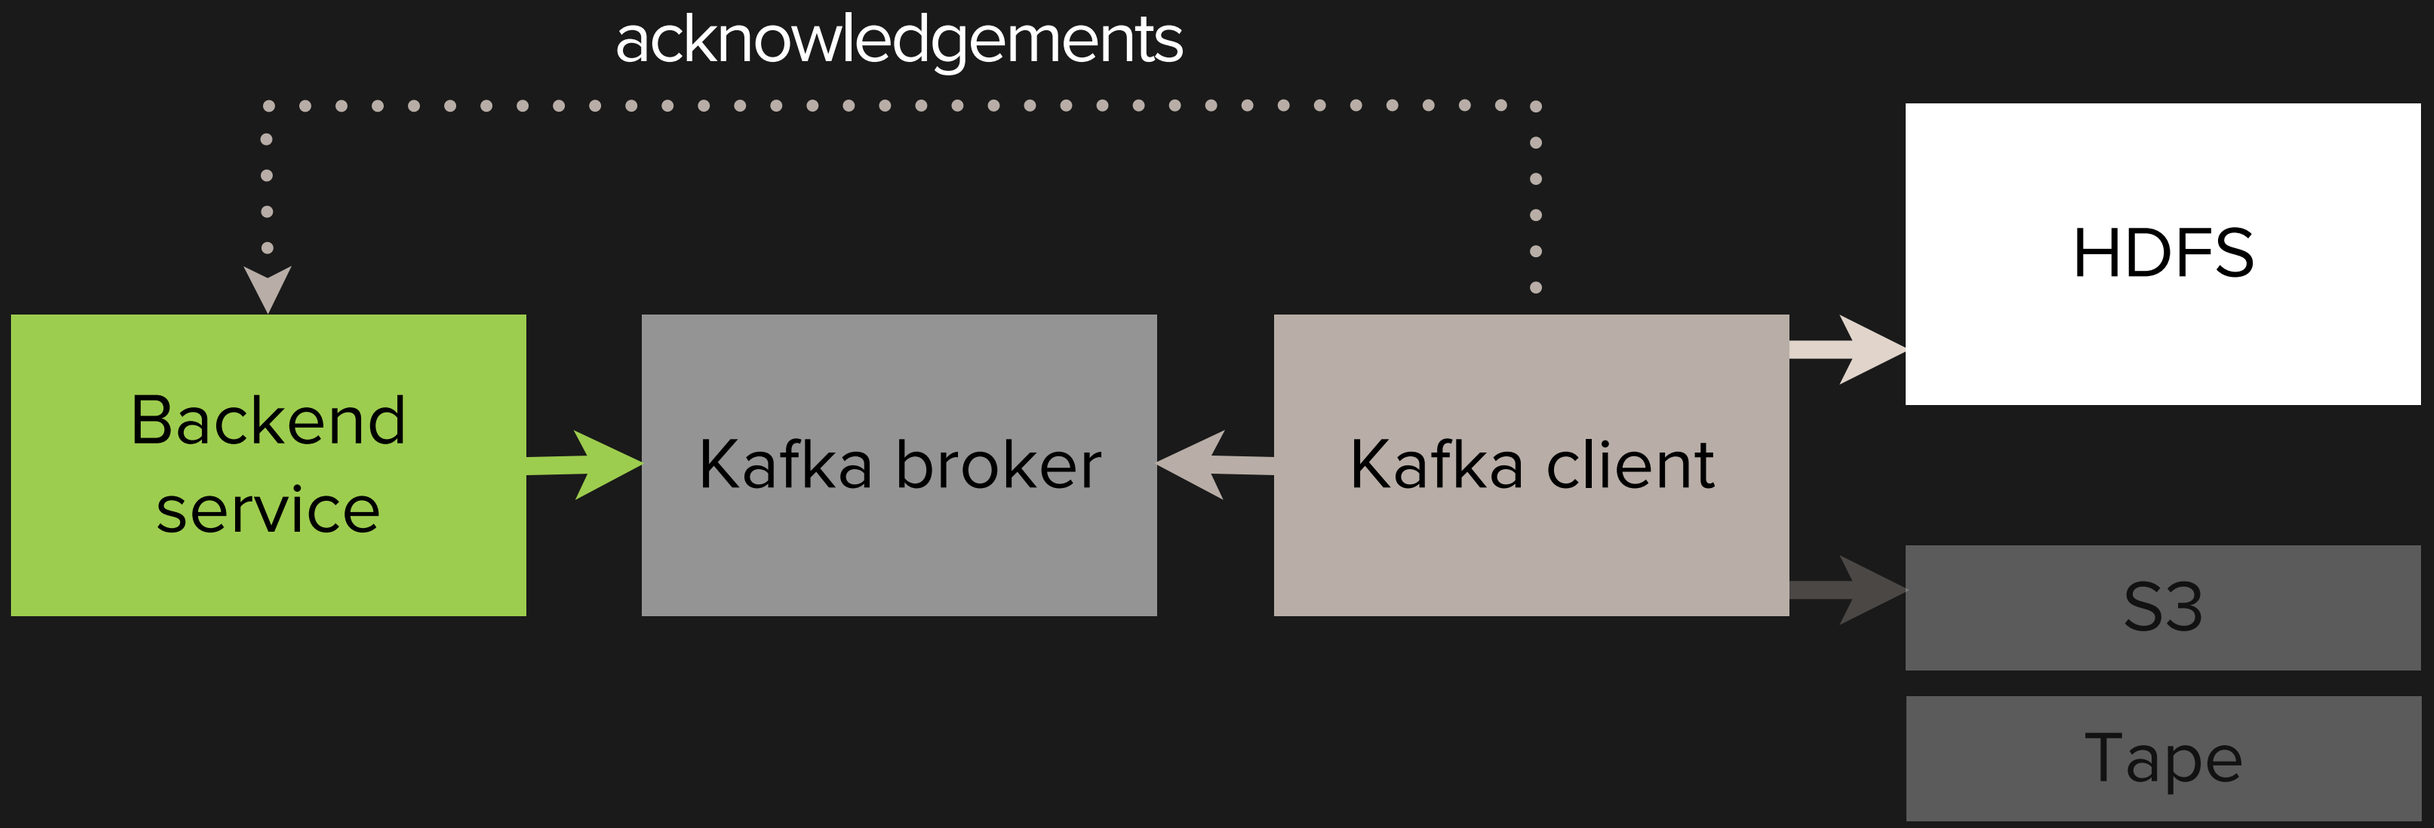

Dumping databases

Not only do we have log files, but also production databases

- Using Sqoop for dumping PostgreSQL
- Map/Reduce job for dumping Cassandra
- For large DB's we parse application logs and only dump deltas

Failures (or the hard stuff)

- We started with a store-and-forward system that didn’t scale
- Kafka has multiple components (client, broker, ZooKeeper)
- Internet weather
- As with many large Java systems: Garbage collection

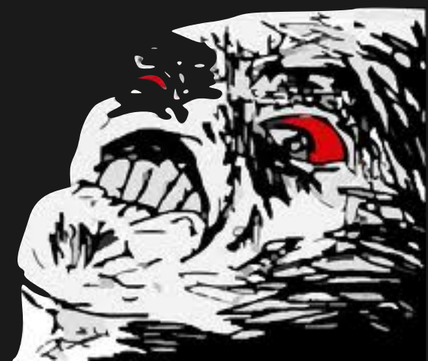

Hadoop: our trusted elephant

Spotify

June 12, 2013

Scheduling

We wrote and **open-sourced** our own scheduler: Luigi

- **Nothing** suitable out there.. (unless you really, really like the XML hell of **Oozie**)
- https://github.com/spotify/luigi
- Written in **Python**
- **Generic** scheduler and dependency system that supports Python M/R, Pig and Hive

Map/Reduce languages

Python with Hadoop Streaming

- Pros: **fast development**, many Spotify libraries available
- Cons: **slower then Java**, no access to Hadoop API

Java

- Pros: **fast**, access to Hadoop API
- Cons: **verbose** language, not many Spotify libraries available

PIG

- Pros: very **small scripts**, **faster** then streaming
- Cons: **yet another language** to learn, not many Spotify libs available

Hive

- Pros: **SQL like** syntax (easy for non-programmers) and relational data model
- Cons: more **moving parts** (not well suited for a whole pipe line)

Scaling Hadoop at Spotify

Our Journey

- Started with a small (scrap metal) cluster of **37 servers**
- Moved to Amazon **Elastic Map/Reduce** (EMR) and S3 to quickly scale
- Built an in-house cluster of **60 nodes** because of EMR costs
- Capacity planning every 6 months, grown to **190 nodes** today
- Just ordered **500 more** nodes
- Put in place **data-retention** policy and data archive

Hadoop failures

- “We just need good developers” - No, we need Hadoop experts
- We underestimated the complexity of Hadoop
- You can throw money at the problem of scaling, but at our scale, it pays off to optimize
- Give people easy tools early on

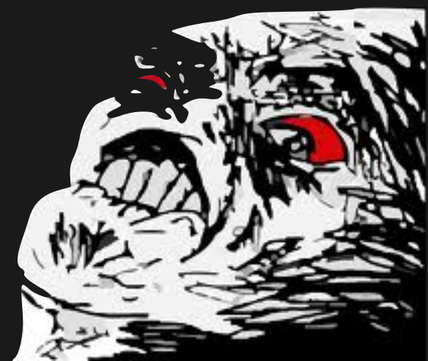

Lessons learned

4+ years of Hadoop taught us

- Hadoop has brought us **very far**. We would never be able to handle the current volume with a “cheap” RDBMS
- “**Commodity hardware**” doesn’t mean **cheap hardware**
- Hadoop isn’t a **silver bullet**
- Hadoop is a **complex** system that needs love and care
- You will have to **extend** Hadoop (and eco-system components) to tailor it to your needs

Databases and visualization

June 12, 2013

Databases: used for aggregates

- Aggregates from Hadoop are put into PostgreSQL or Cassandra
- PostgreSQL powering dashboards and empowering analysts
- Cassandra for columnar sets (Spotify Analytics for labels)
- Databases are used for low-latency access by systems or analysts

Spotify Data Warehouse

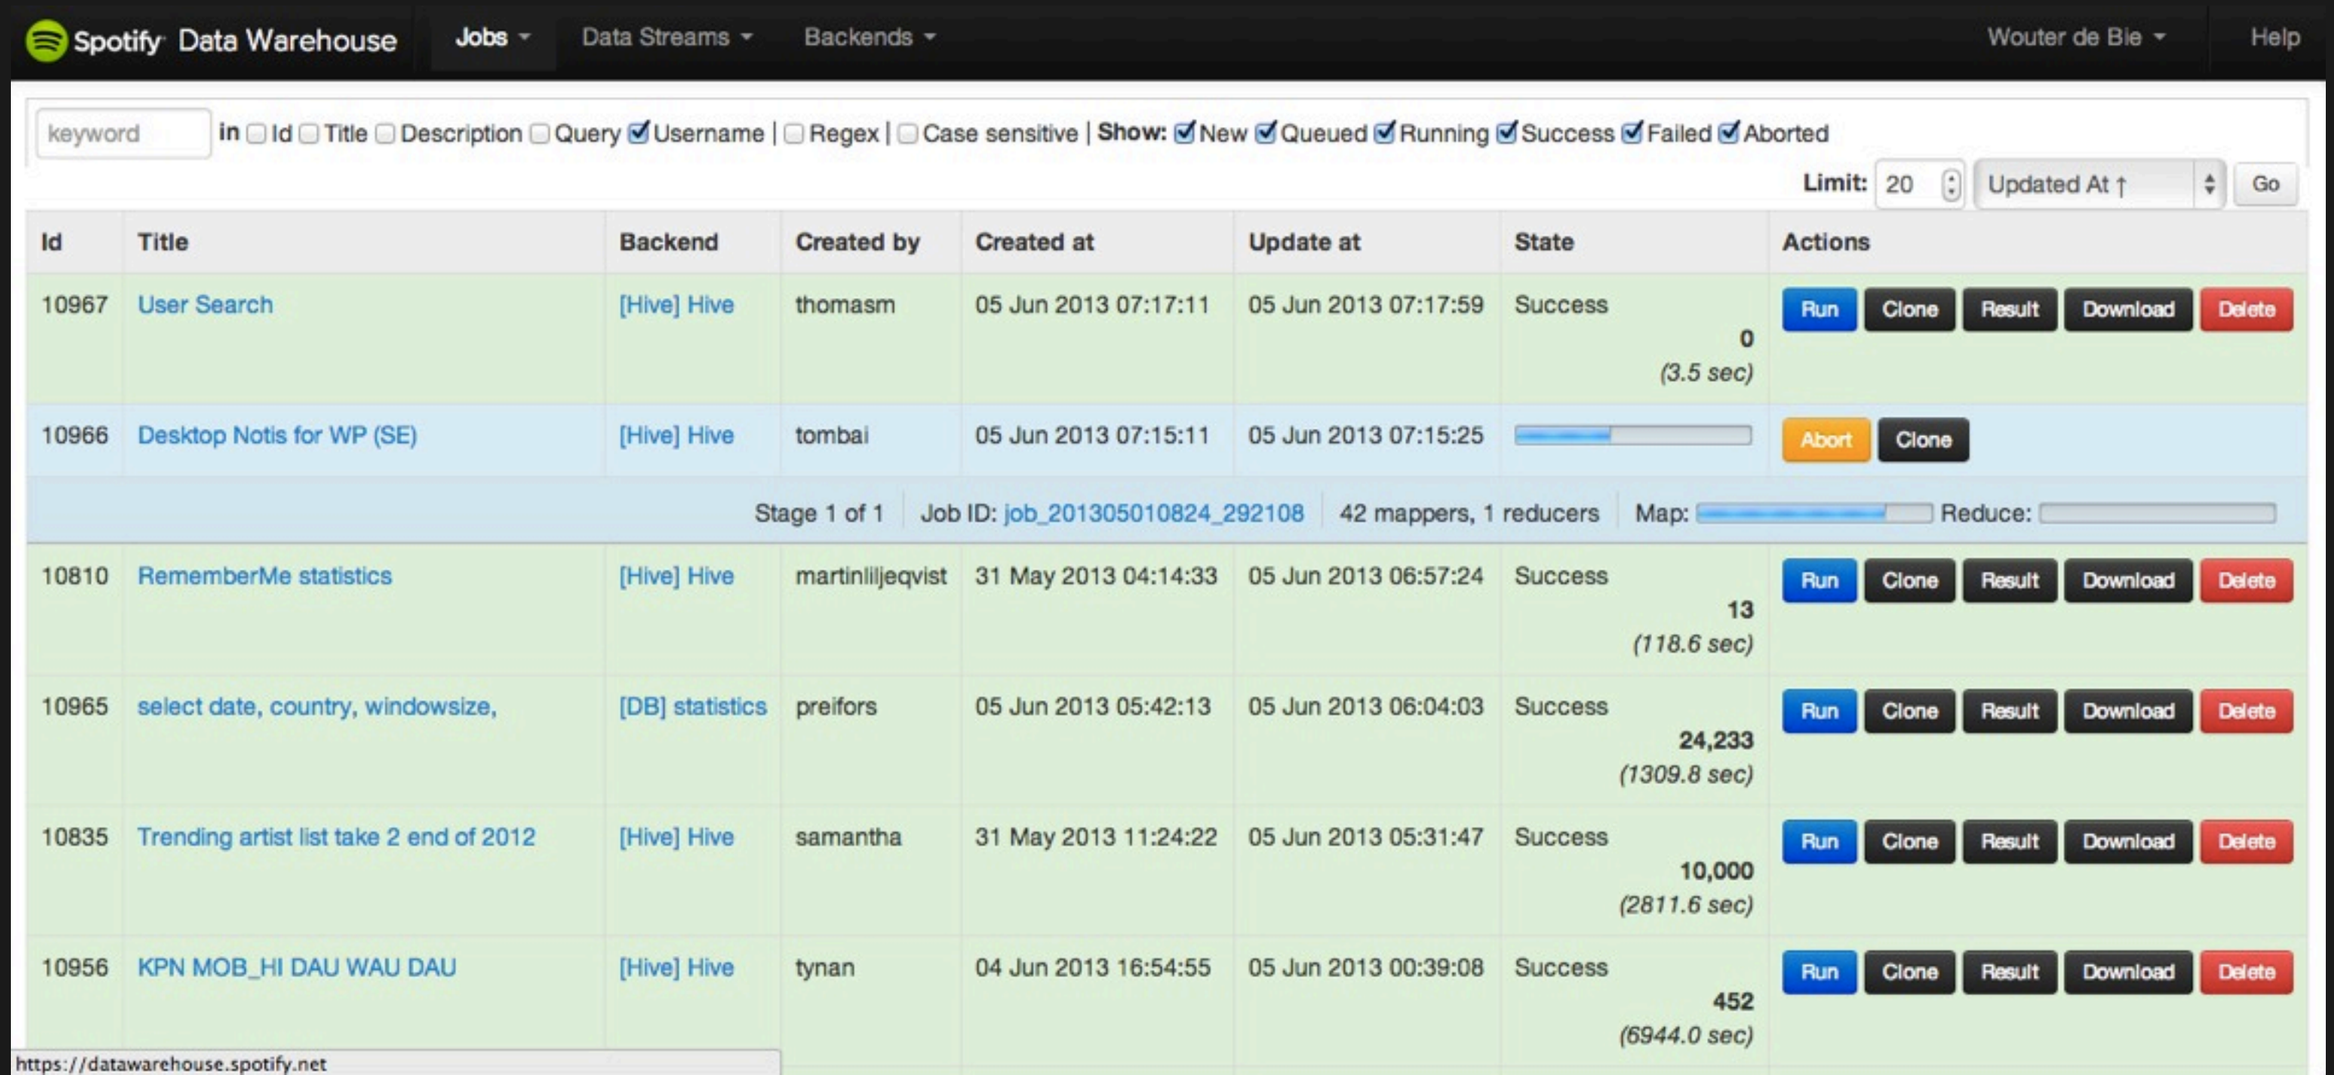

Spotify Data Warehouse

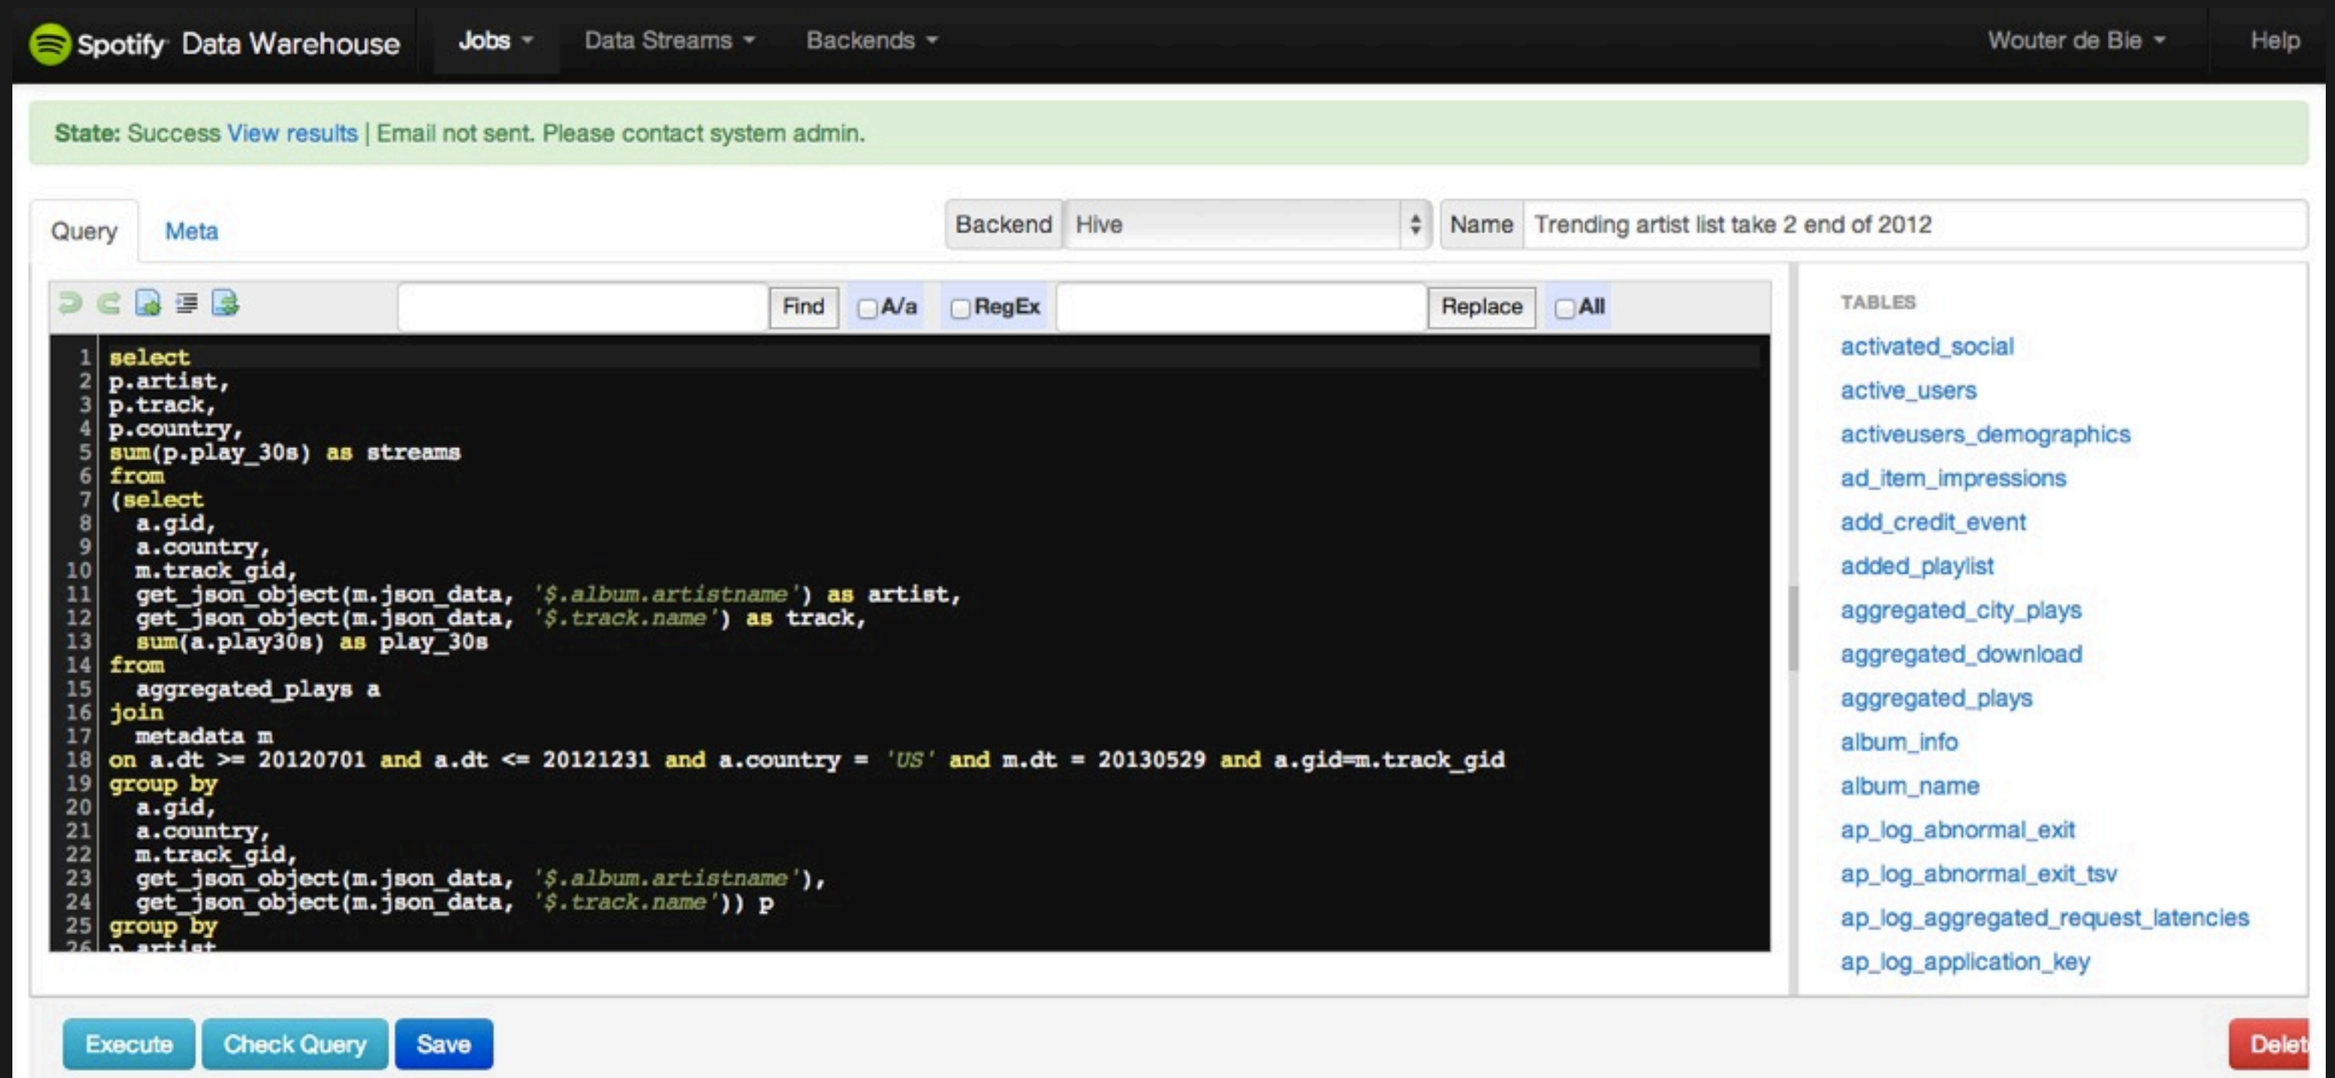

Spotify Analytics Dashboards

Spotify Analytics

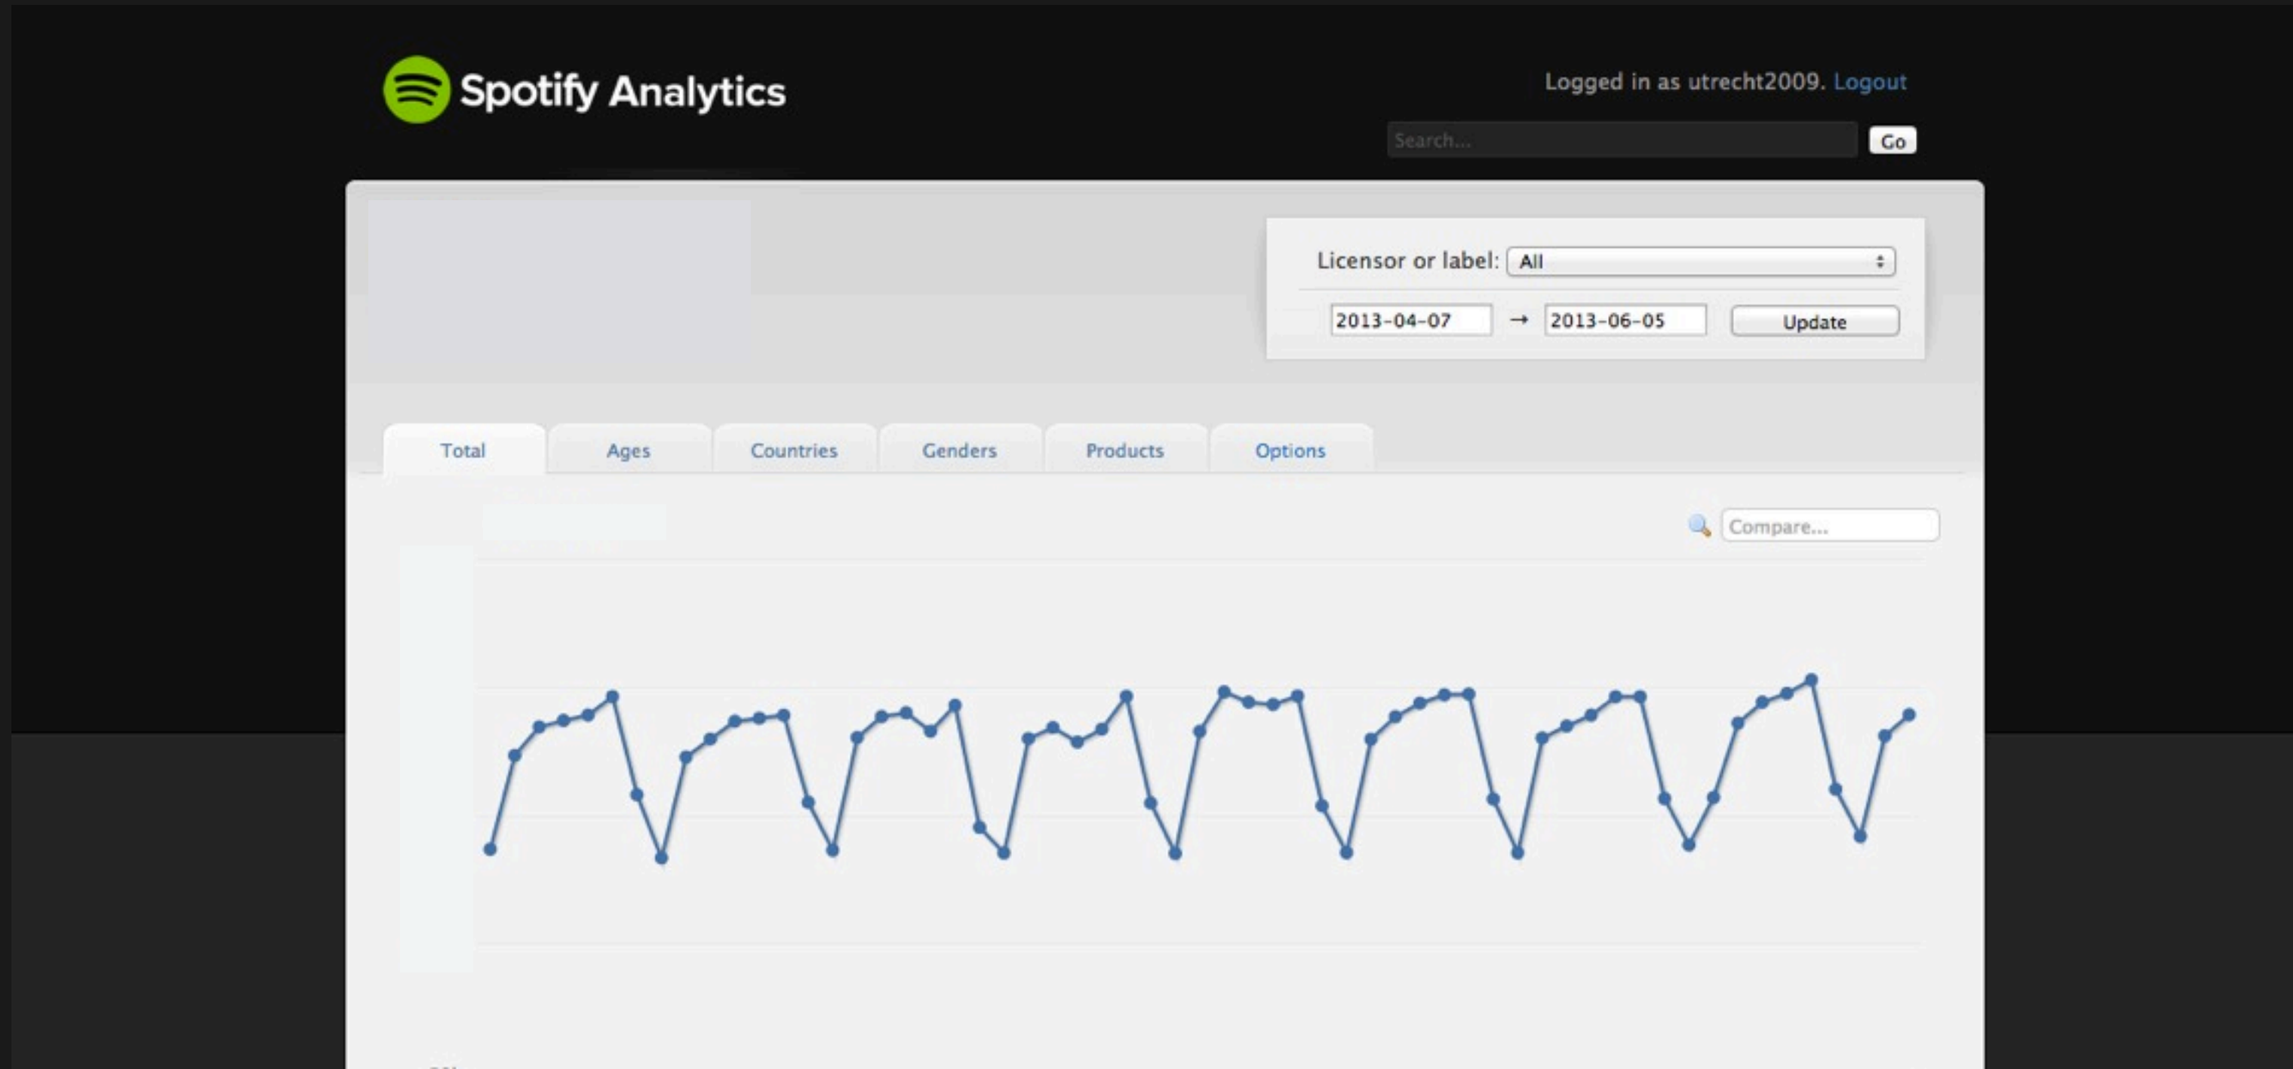

Spotify Analytics: Daft Punk

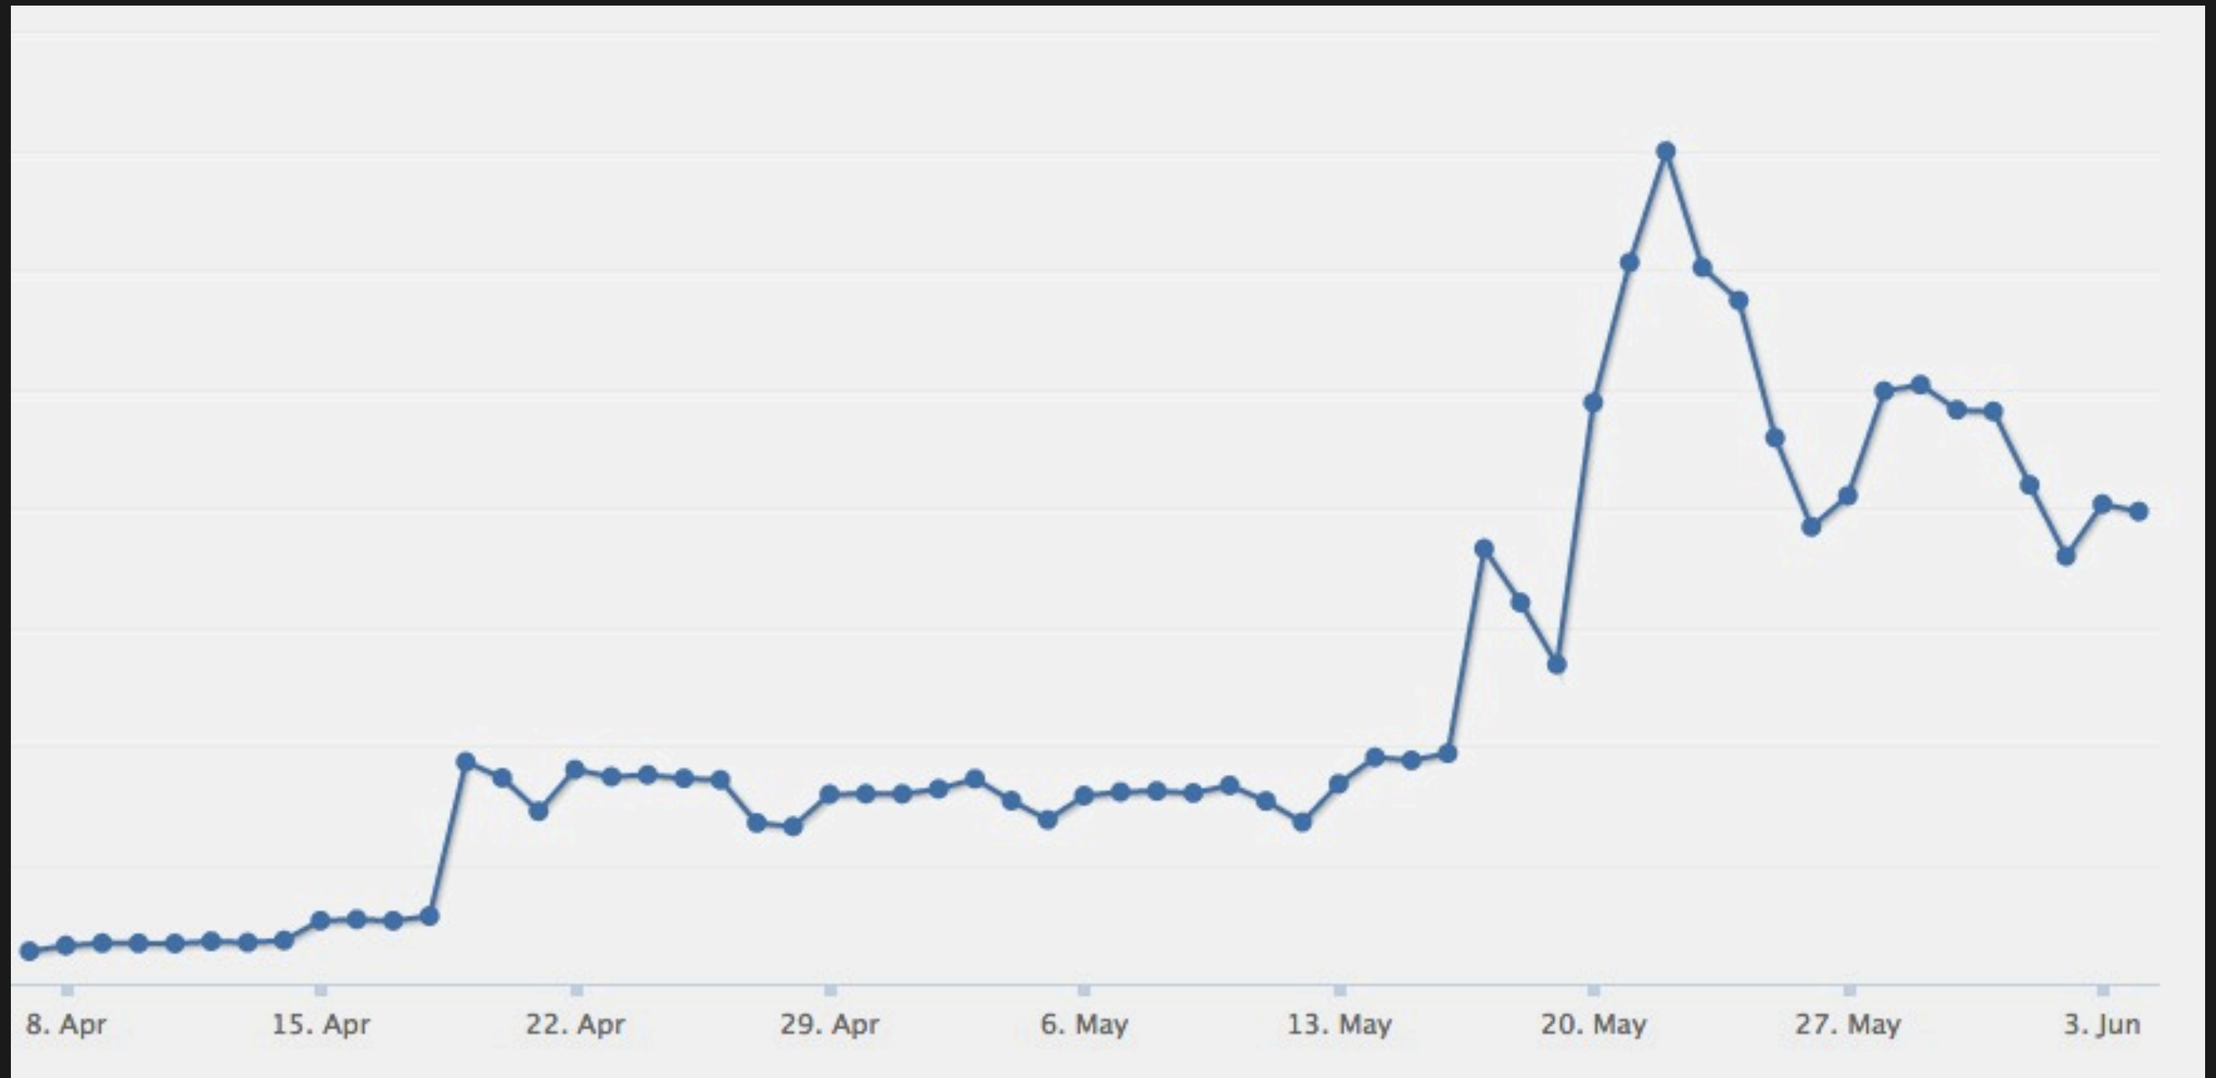

Spotify Analytics: Kendrick Lamar - Gábor was right :)

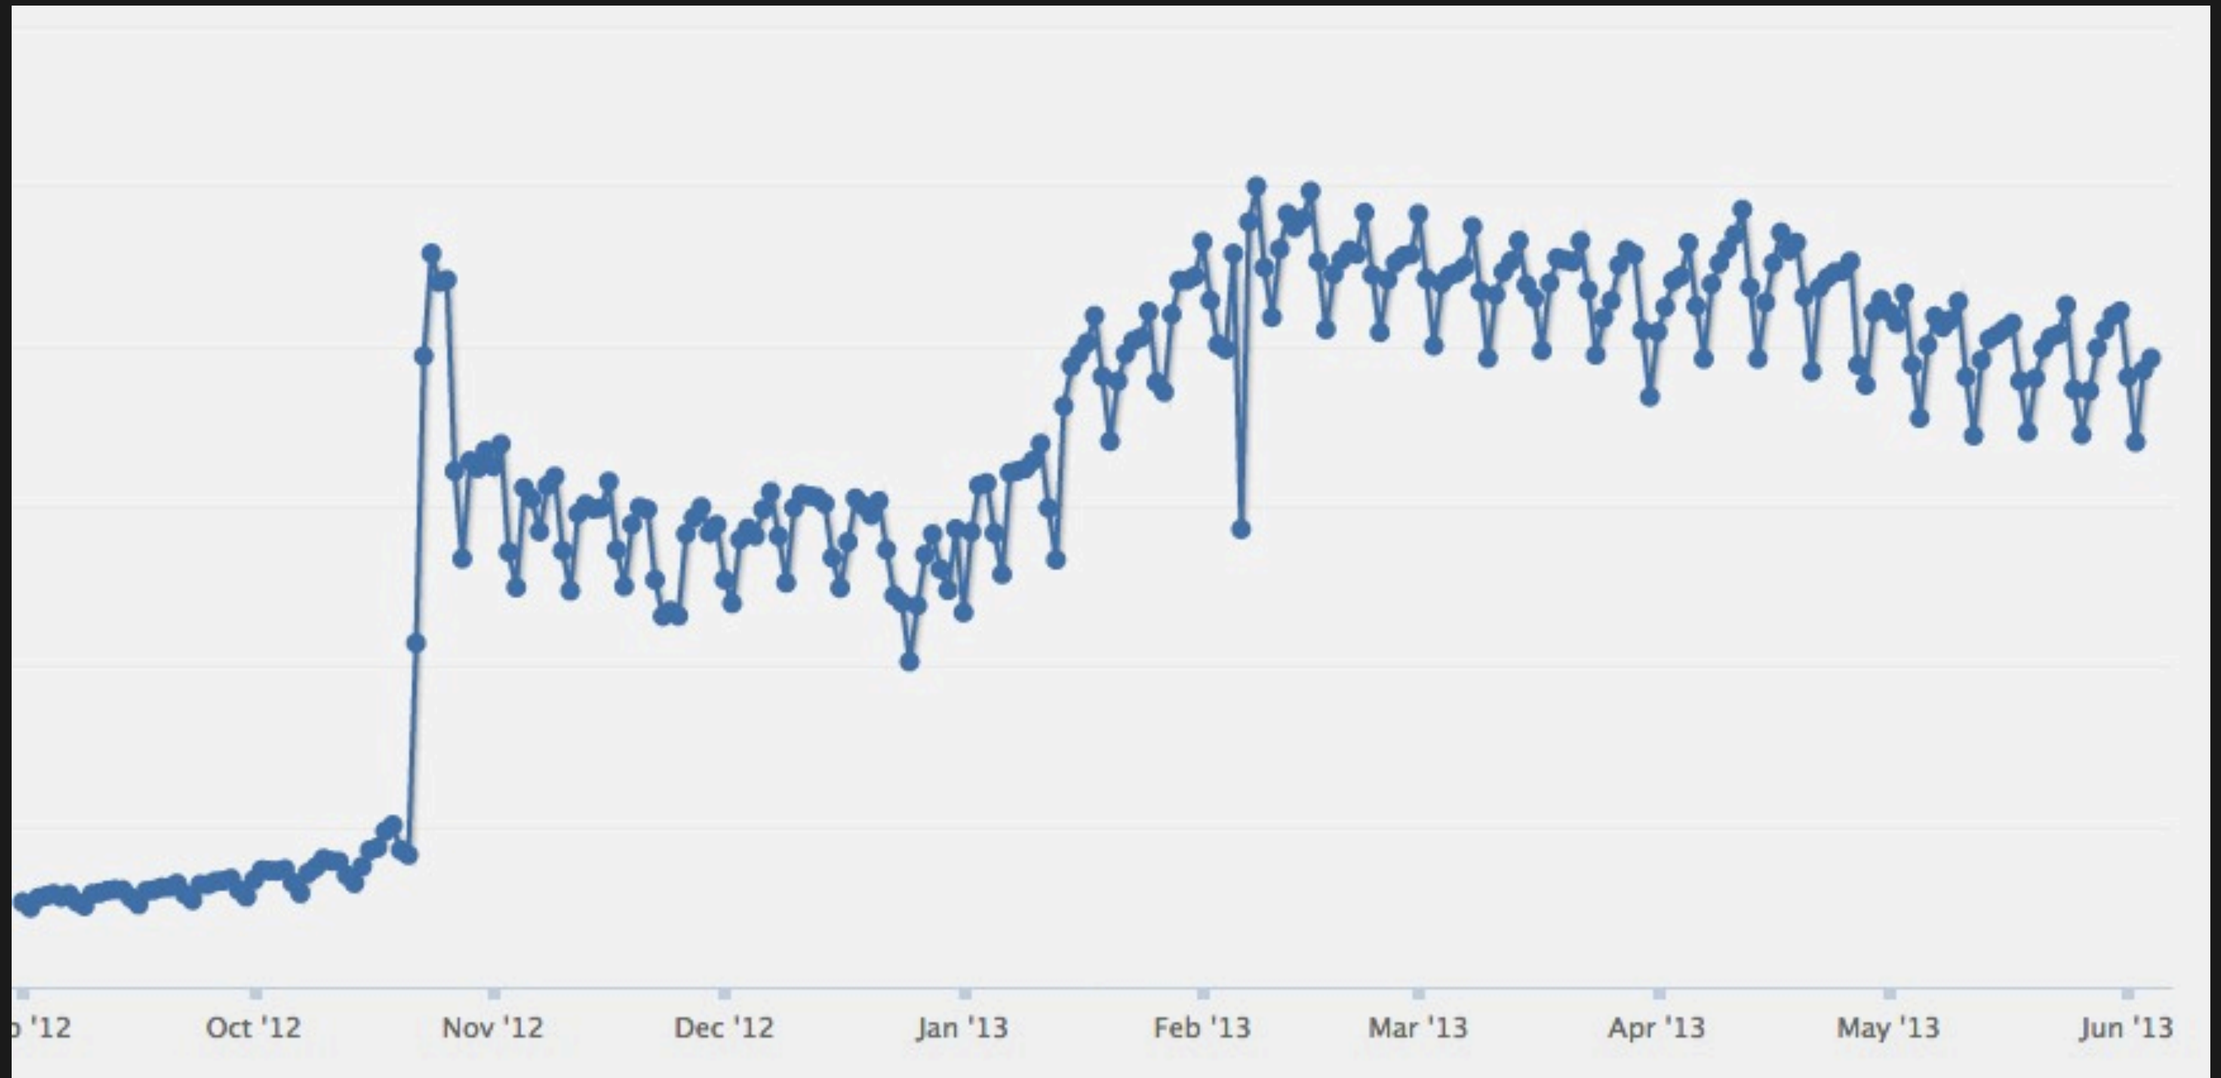

Spotify Analytics: Whitney Houston

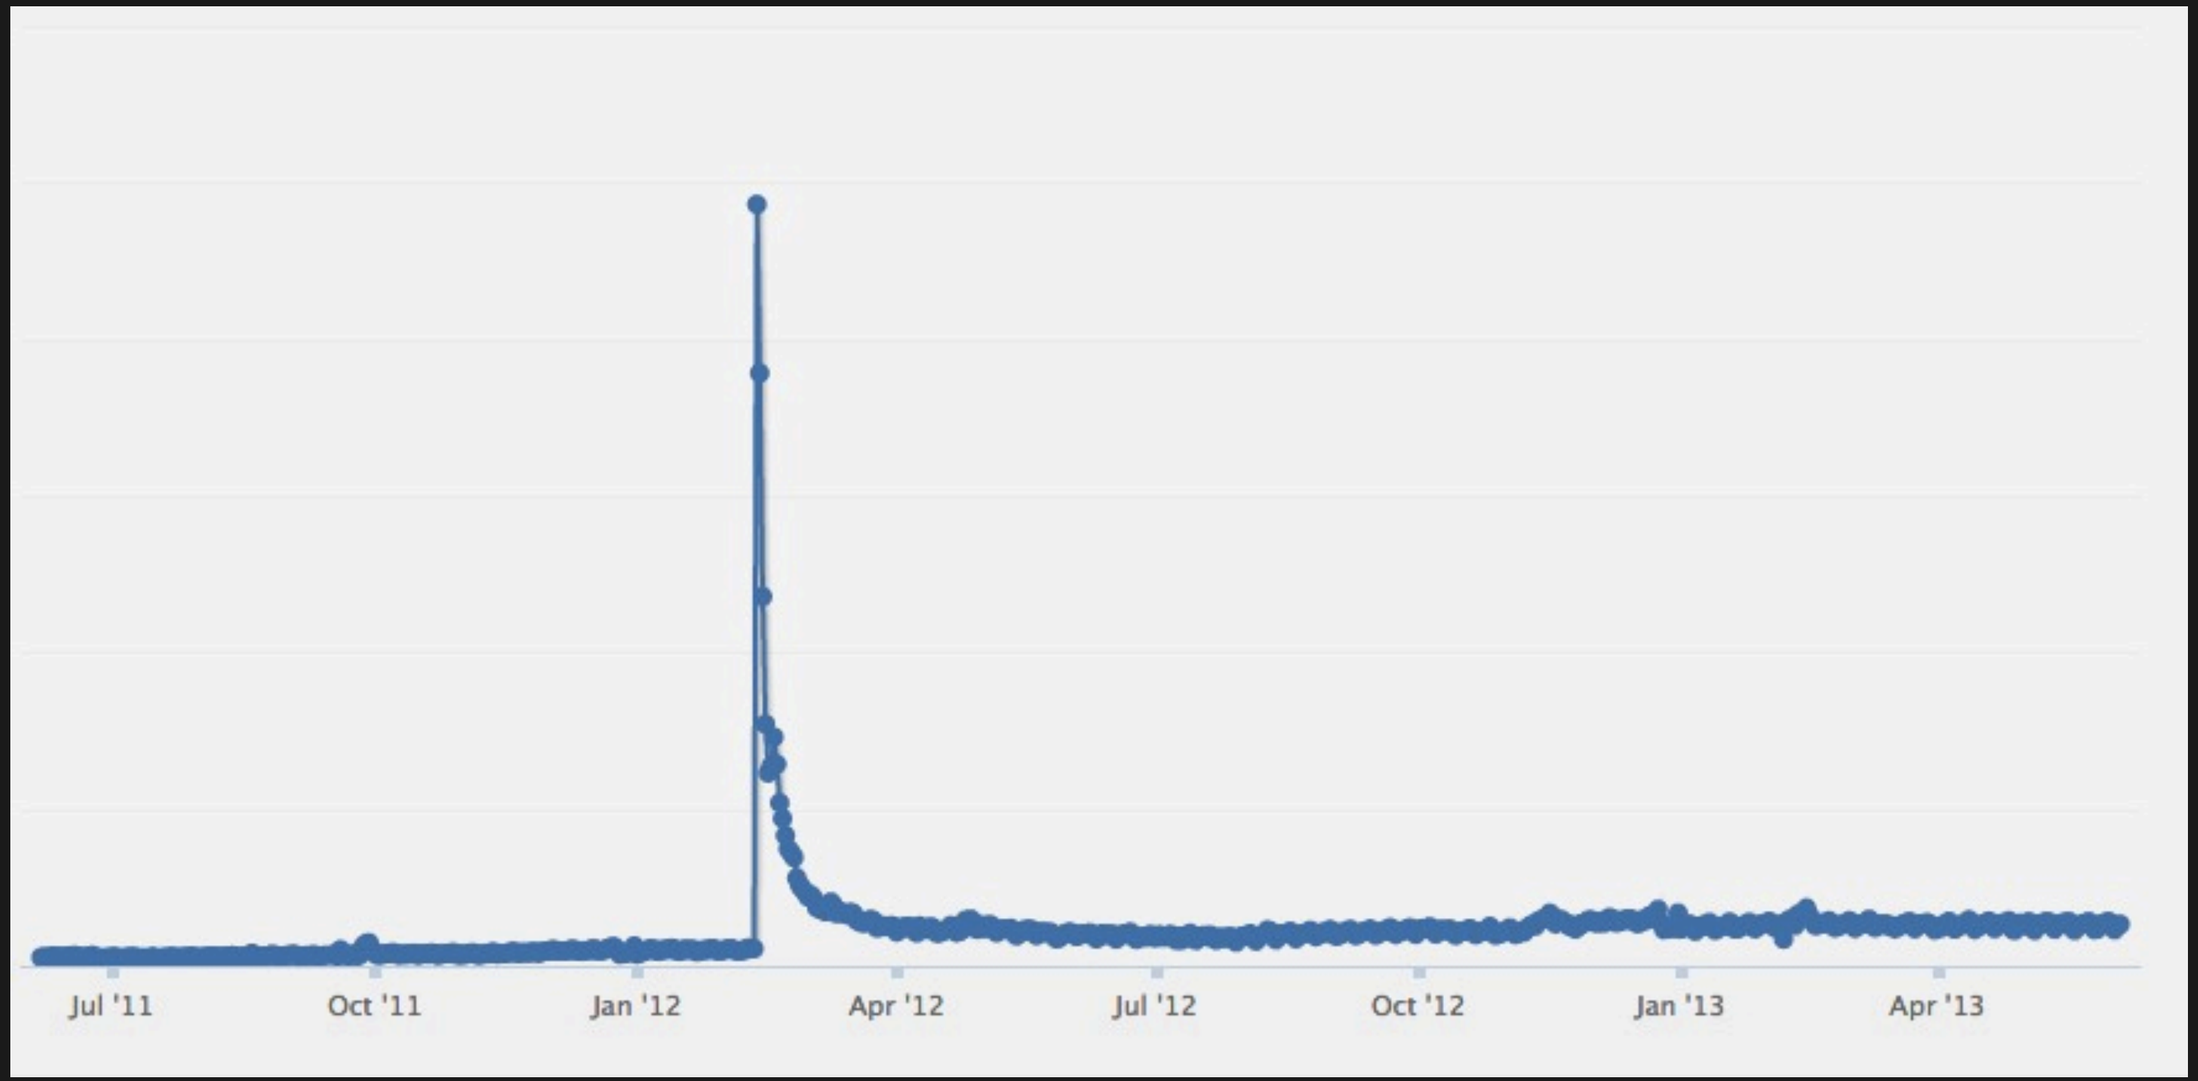

Lessons learned

- Data is crucial for our business
- Hadoop and other cutting-edge technology work pretty well for us
- Hiring technical analysts was a good idea!
- There is no “one-size-fits-all” data product
- Spotify has the “build it ourselves” mindset. Sometimes it’s better to buy then build

Want to join the band?

Check out **http://www.spotify.com/jobs** or **@Spotifyjobs** for more information.

Or mail: wouter@spotify.com
Or twitter: @xinit

Spotify

June 12, 2013