# Ustream’s experiences in NoSQL world

23/03/2016

USTREAM

An IBM company

# About Ustream

- ★ Founded in 2007, acquired by IBM in 2016
- ★ Live streaming and VOD
- ★ Web, Android, iOS
- ★ Freemium / Pro / Demand / Align

# Scale in numbers

- ★ 3 billion hours viewed
- ★ 80 million users per month
- ★ 1.5 Tbit/s peak throughput
- ★ More than 2 million peak concurrent viewers

# Main Components

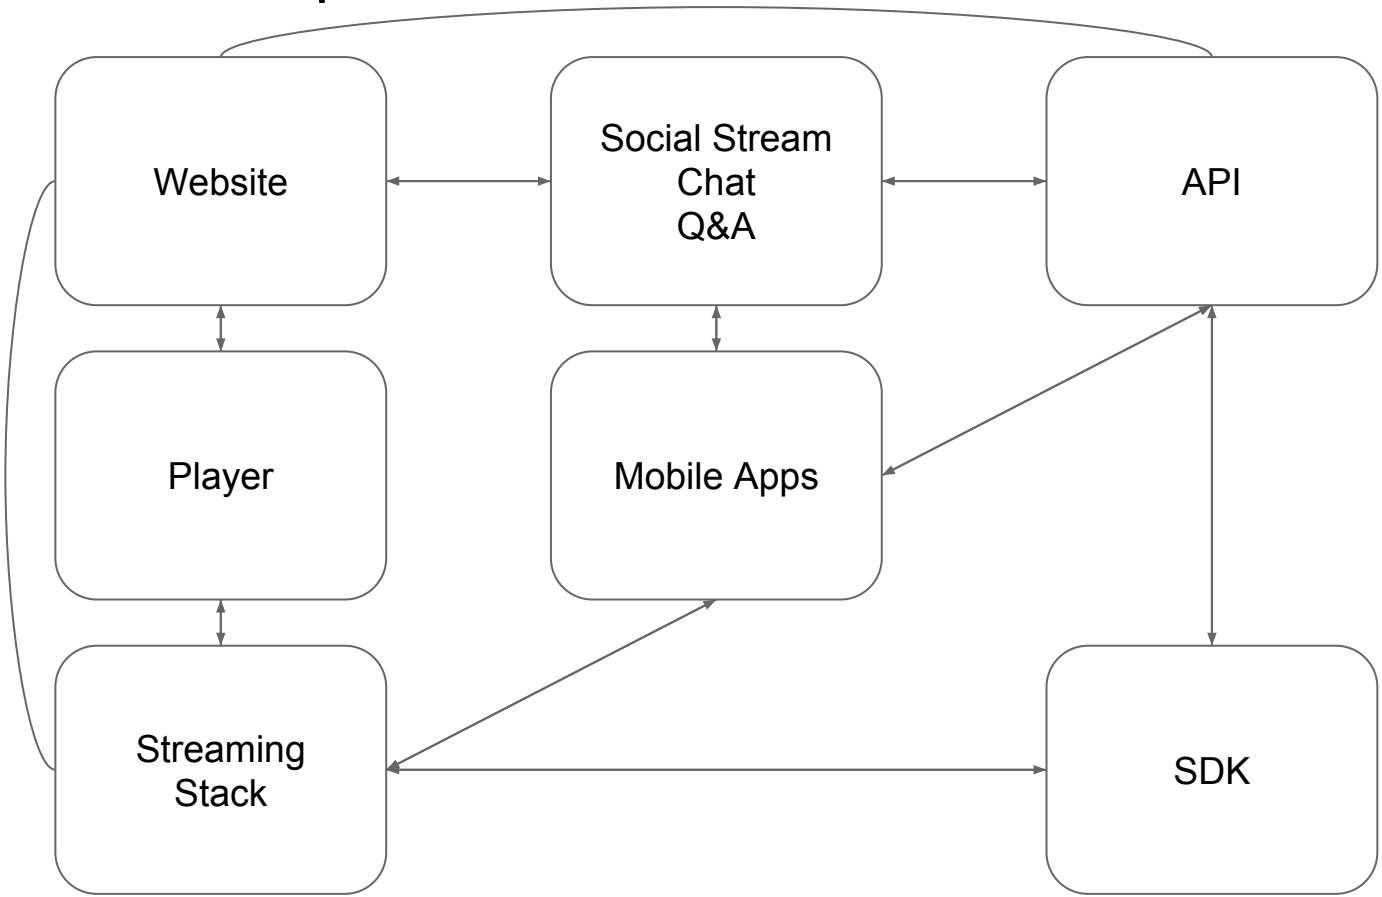

# Streaming Stack

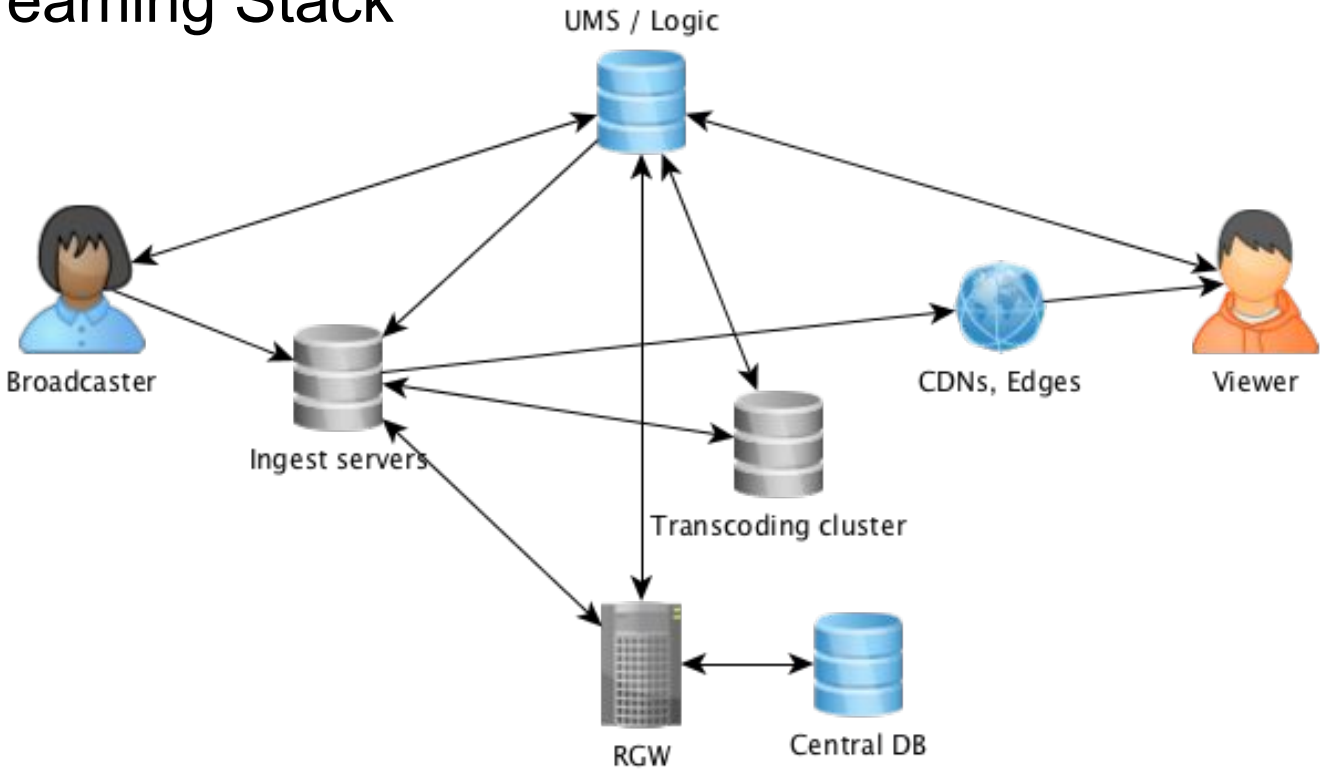

# NoSQL in other parts of the stack

- ★ Memcached
  - ○ Web data + sessions
  - ○ Counters
  - ○ Thumbnails
- ★ Redis
  - ○ Social stream
  - ○ Transcontroller

# Our Use Case

- ★ Scalable socket server cluster
- ★ Real-time bidirectional connection with Ustream users
- ★ Built on JVM stack
- ★ Real time aggregations of numeric values cluster wide
- ★ Real time cluster queries with eventual consistency

# Let’s put data into nosql cluster…

# Our candidates

- ★ Redis
- ★ FoundationDB
- ★ Cassandra
- ★ Riak
- ★ Aerospike

# Redis

- ★ easy to use
- ★ lot of experience with redis
- ★ Redis cluster was in alpha stage
- ★ Cluster wide numerical aggregations weren’t supported.

# FoundationDB

- ★ Cool new db
- ★ Impressive fault tolerance demo
- ★ Acquired by facebook and closed opensource! :((((

# Riak

- ★ Lot of conversation with support
- ★ Poc was too expensive
- ★ They suggested to use too much riak nodes.

# Cassandra

- ★ Positives: fast writes, fast reads when filtering with primary keys
- ★ Negatives: aggregations are very slow, very slow node synchronisation, driver issues, lot of scripts for maintainance, nodetool wasn’t stable when cluster was under pressure

# Aerospike

- ★ A good cache, not a good database
- ★ Lot of inconsistency issues when nodes are flapping
- ★ https://aphyr.com/posts/324-jepsen-aerospike
- ★ Talked with the CTO without success :)

# Hmm, do we have to do this way?

# Aggregator service 1.

- ★ Cqengine (https://github.com/npgall/cqengine)
- ★ NoSQL queries on Java Collections
- ★ Easy to plug into existing architecture
- ★ Written in pure java
- ★ Data remains on the node, no need to move into external db.

# Aggregator service 2.

- ★ Built on cqengine
- ★ stateless
- ★ scalable
- ★ api
- ★ Service discovery with consul

# Architecture

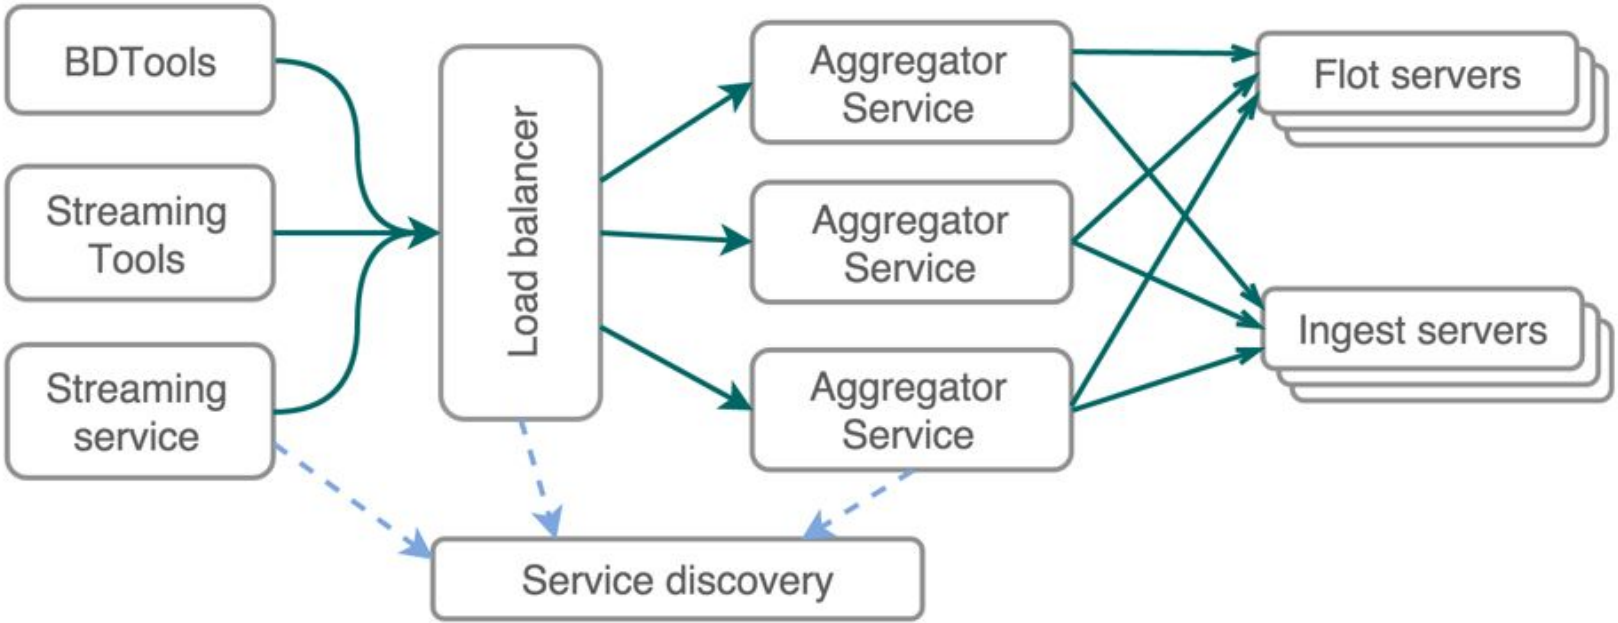

# Data flow

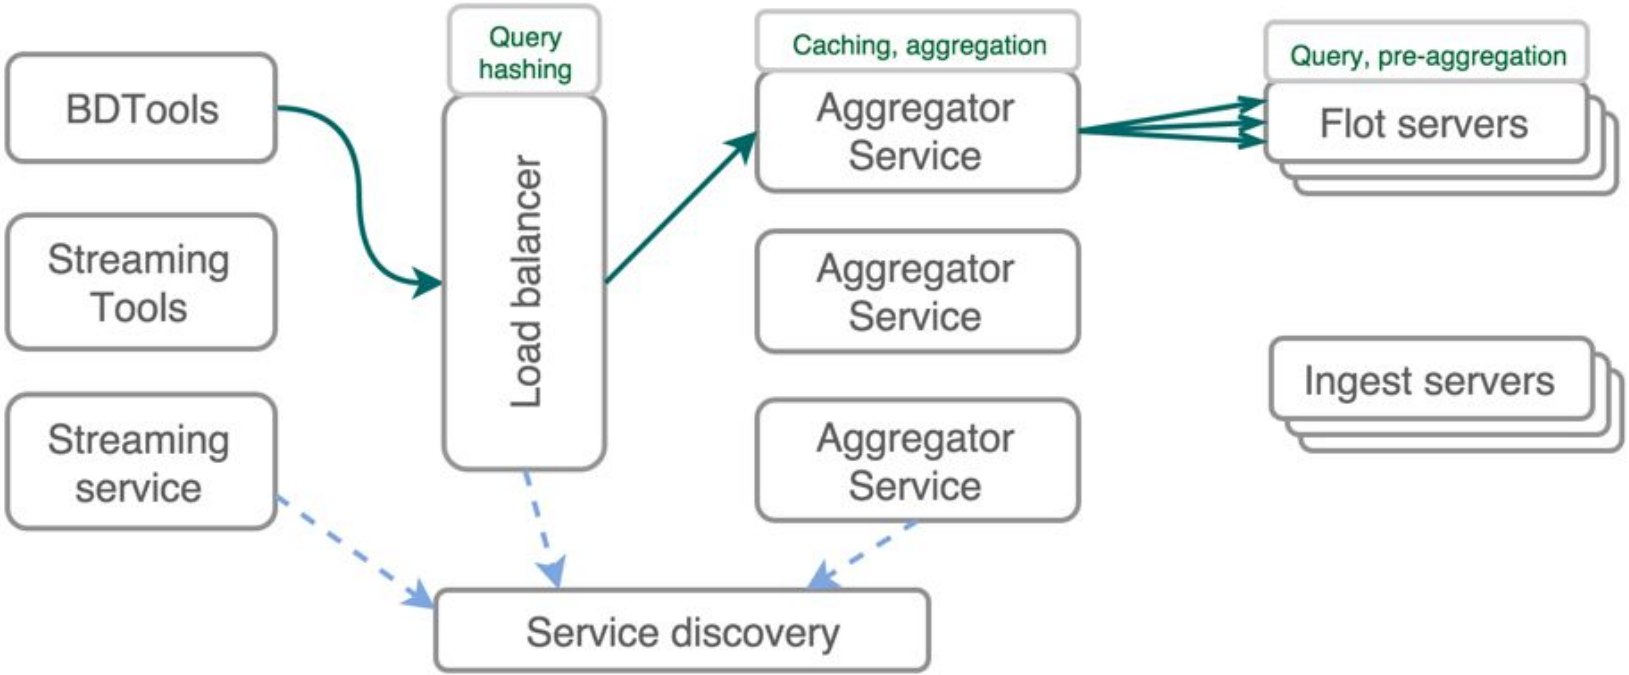
BDTools
Streaming Tools
Streaming service
Query hashing
Load balancer
Caching, aggregation
Aggregator Service
Aggregator Service
Aggregator Service
Query, pre-aggregation
Flot servers
Ingest servers
Service discovery

# Q&A

# THANK YOU!

ustream.tv