

Timelines @ Twitter

- Arya Asemanfar
 - twitter
 - track : architectures you're always wondered about
- figures
 - 100M active users wold wide
 - poll based : 200KQPS
 - search 30KQPS
 - 250M new tweets per day
 - 26 billion deliveries evey day
 - latency 3.5s

- timeline
 - poll-based timeline
 - search
 - streaming (aka push)

- architecture
 - http proxy
 - tweet api
 - fanout queue
 - timeline cache (redis RPUHX)
 - timeline api : thrift

- problems
 - flood on a timeline
 - in process cache
 - detects hight frequency requests
 - caching
 - "ingester"
 - with search index
 - RPC common library
 - finagle
 - distributed tracing...
 - configurable pools
 - connection management
 - chanel : lucene
 - users timeline : redis
 - delivery type
 - fan out
 - fan in
 - hybrid
 - users : OR/chanel AND
 - high frequency searches
 - precomputed
 - put into redis
 - Redis vs search index
 - both bounded memory
 - cache eviction