



PLATFORM ENGINEERING

Lessons Learned in Scaling Twitter

@bmdhacks
qcon london 2014



Scalability, actually all Software Engineering can be described in how you divide a problem into components. It is the scalpel you use to divide large problems into tractable ones. The abstractions are the boundaries of these cuts, and a well-divided problem is both scalable and manageable.

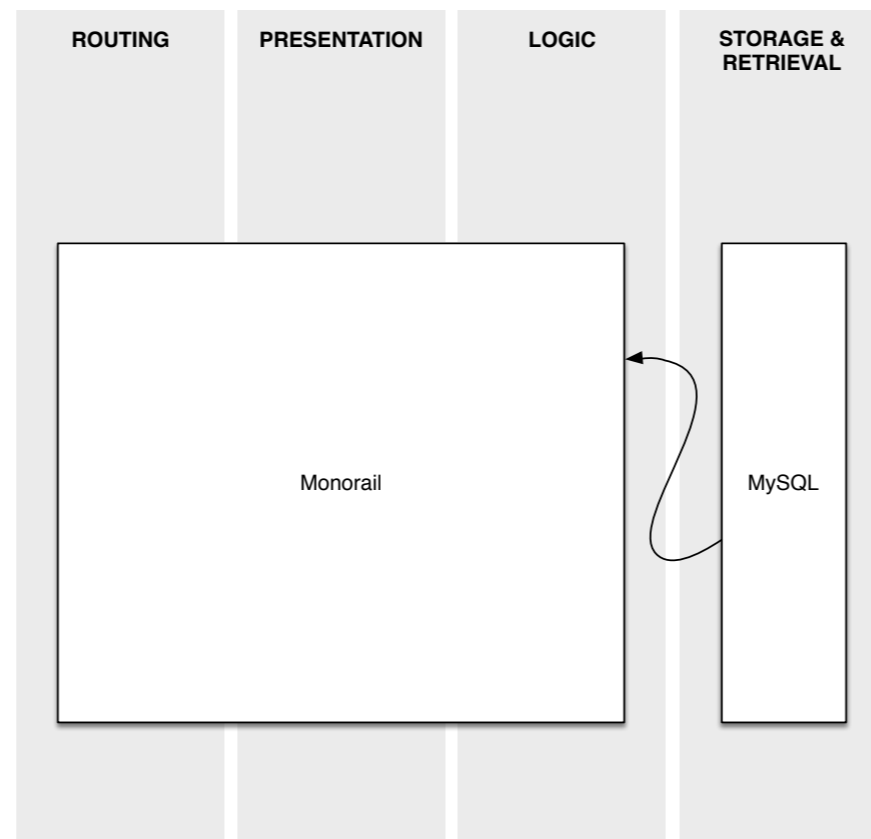
Lessons Learned

- Incrementally Implement SOA
- Separate semantics from execution
- Use statistics to monitor behavior



Service Oriented Architecture

@bmdhacks
qcon london 2014



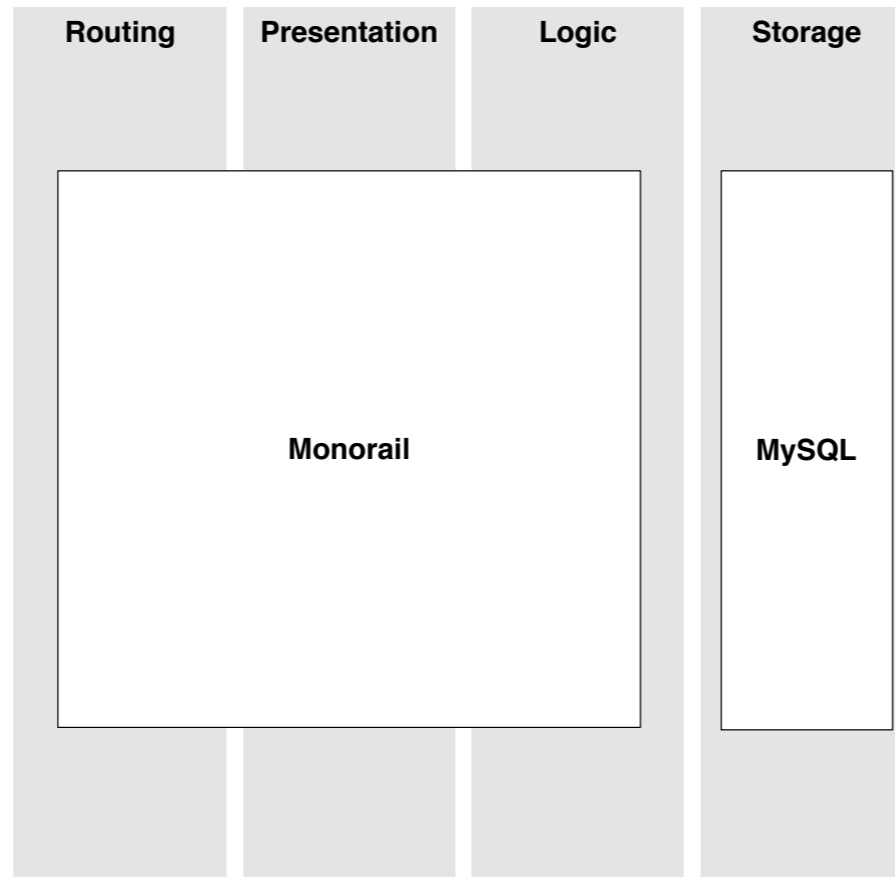
@raffi
qcon shanghai 2013

The monorail was our term for a single Rails application running on ruby 1.8.3

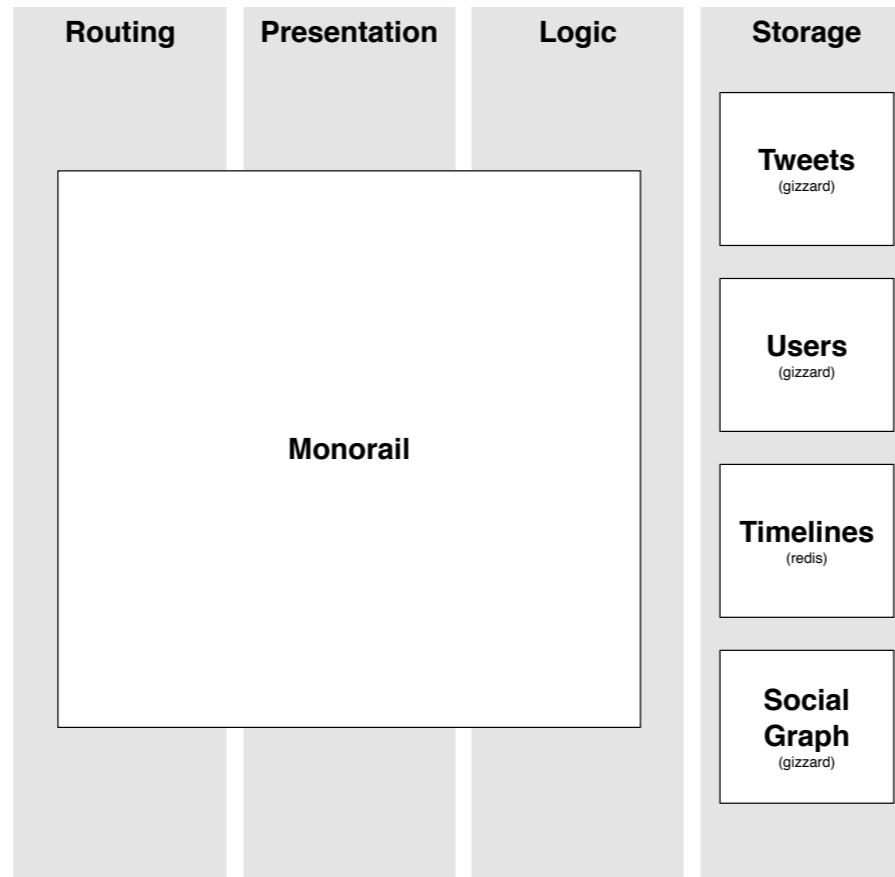
challenges

- difficult to scale storage model
- any change deploys to all servers
- poor concurrency, runtime performance
- leaky abstractions / tight-coupling

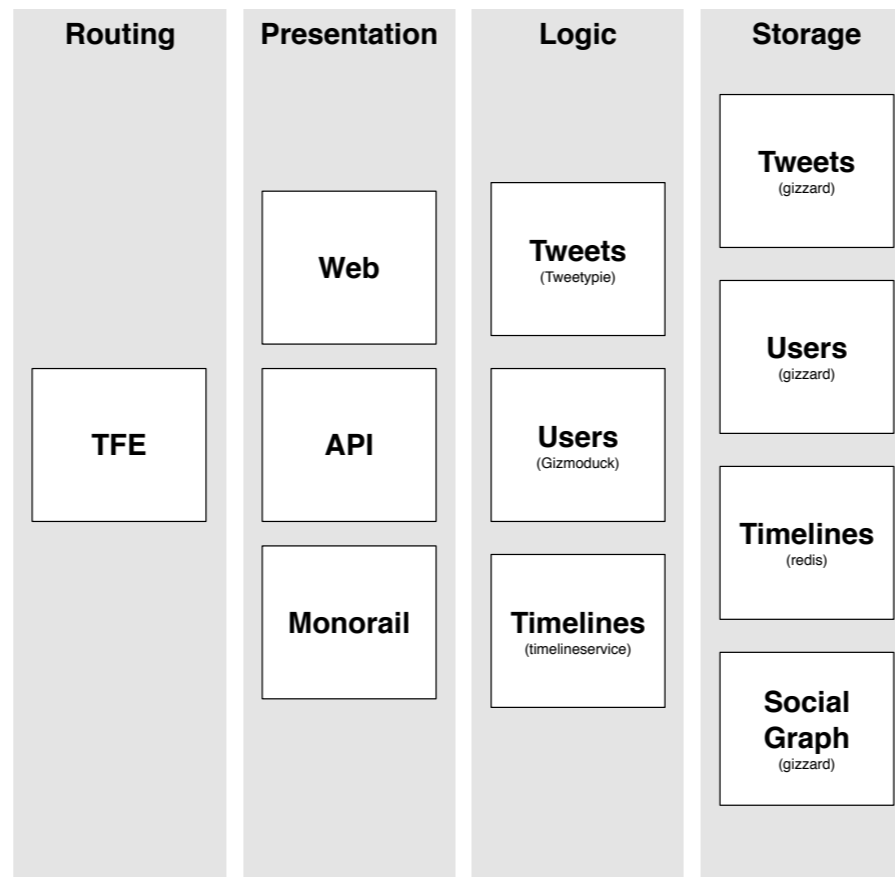




@bmdhacks
qcon london 2014




@bmdhacks
qcon london 2014



@bmdhacks
qcon london 2014

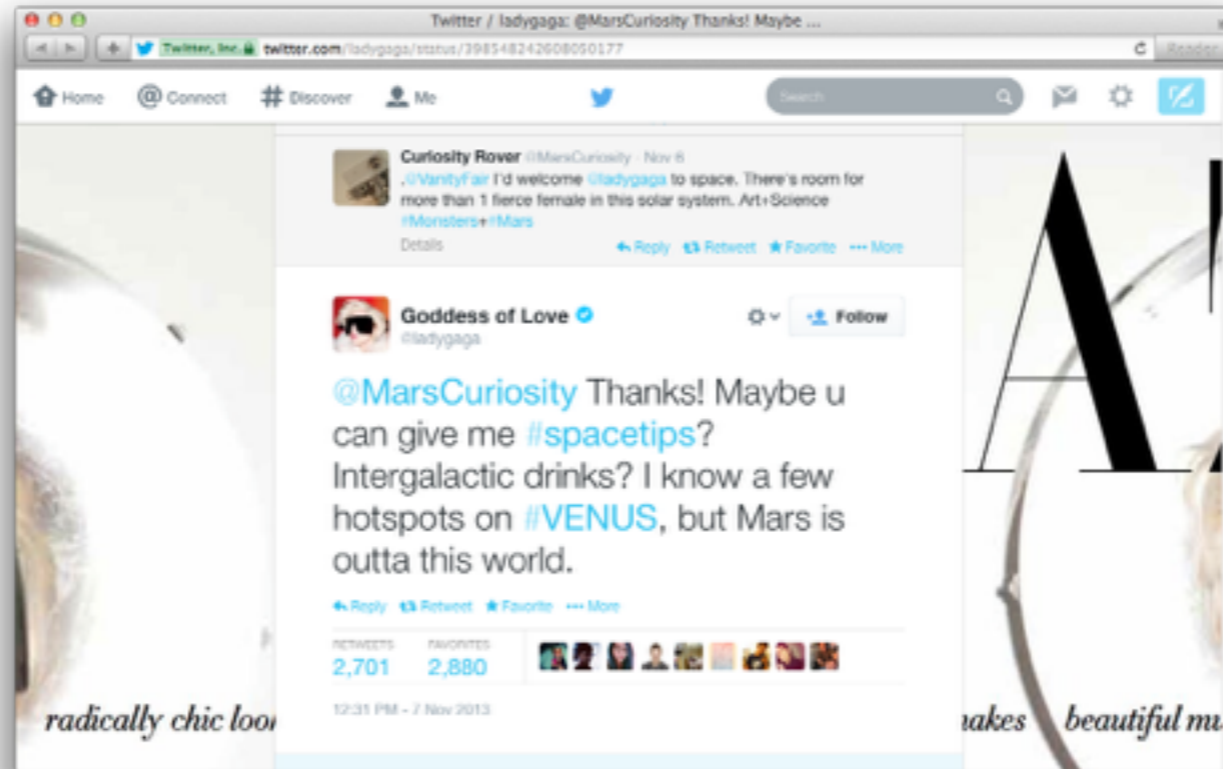
The final result of all services at twitter




ndhacks
ion 2014


Actually this is a map of all services at twitter

Sending a Tweet Today





@bmdhacks
qcon london 2014






Write API






Write API

Fanout



Write API


Fanout



Write API

Fanout

Timeline Cache





Write API


Fanout


Redis

Timeline Cache











Write API


Fanout






Write API

Fanout




Timeline Service






Write API

Fanout




Timeline Service






Ingester



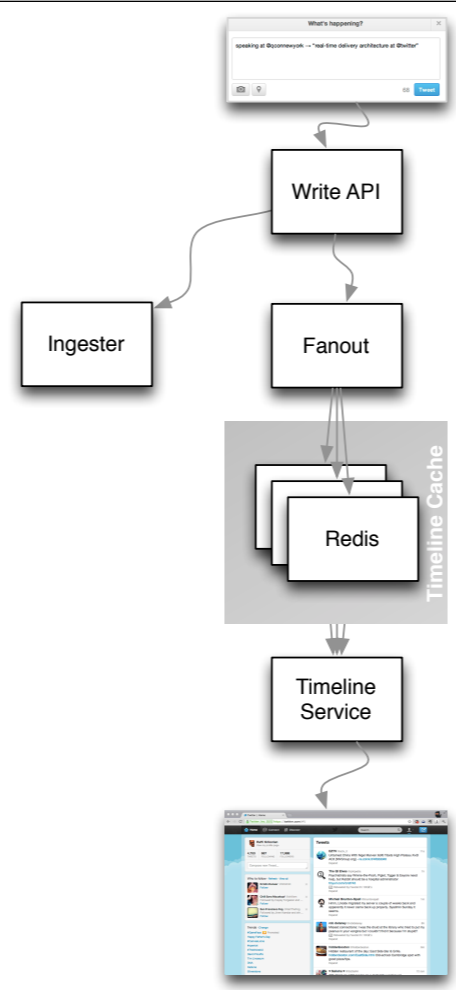
Write API

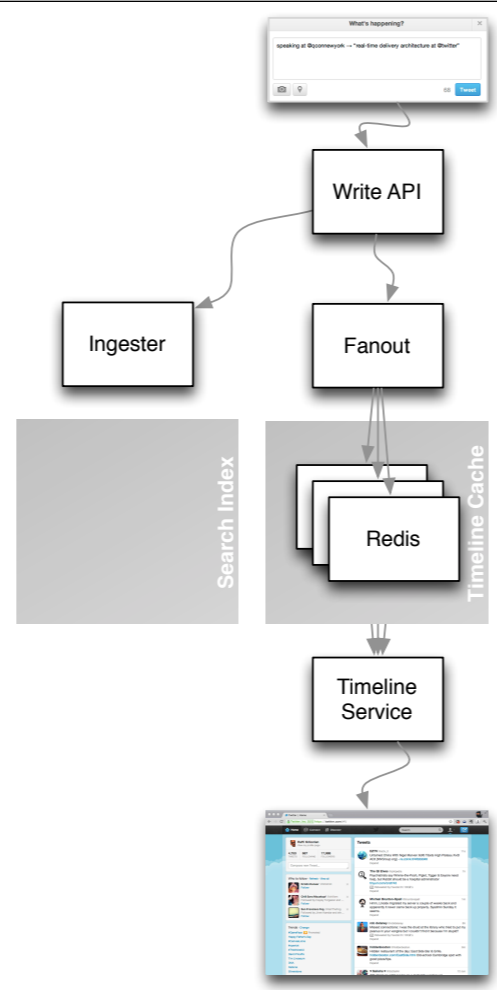
Fanout

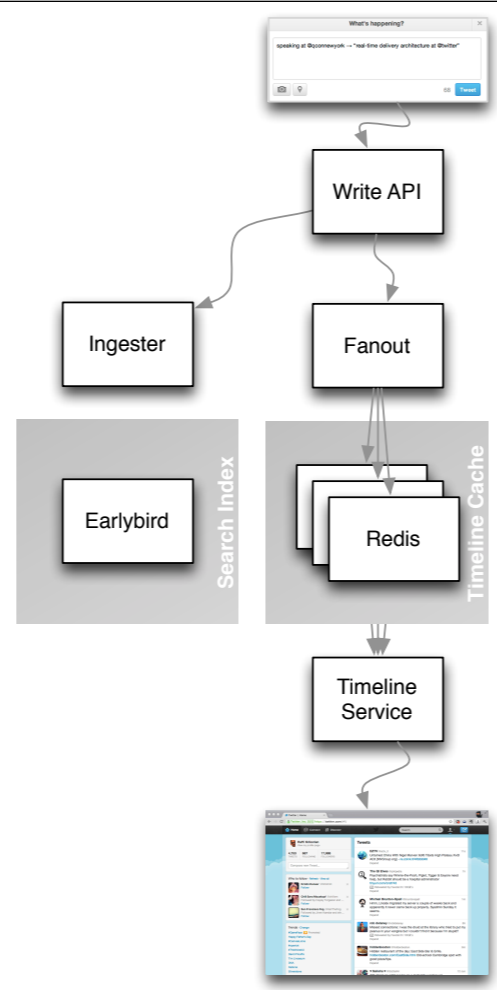


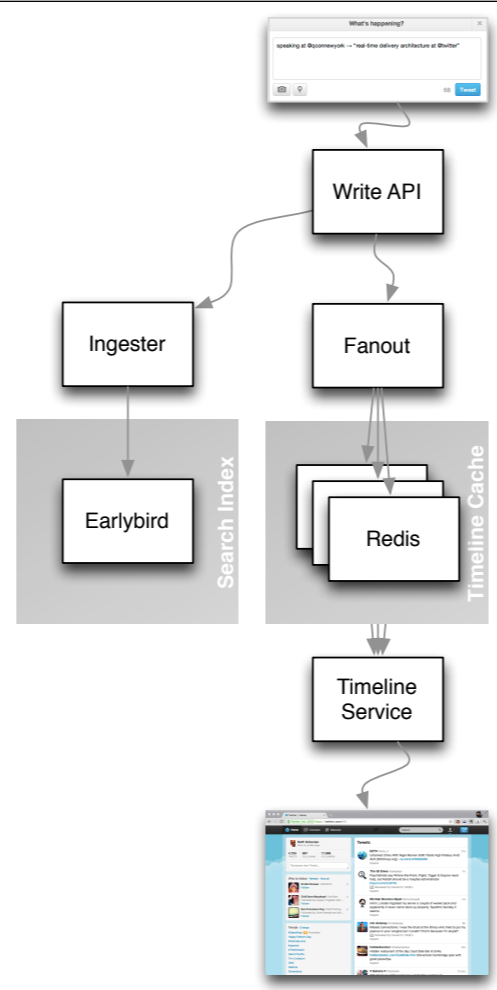
Timeline Service

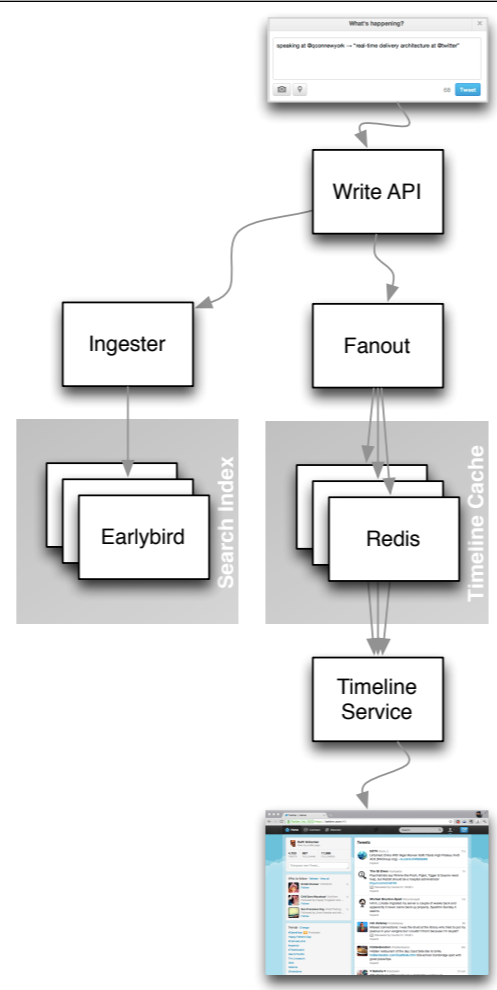


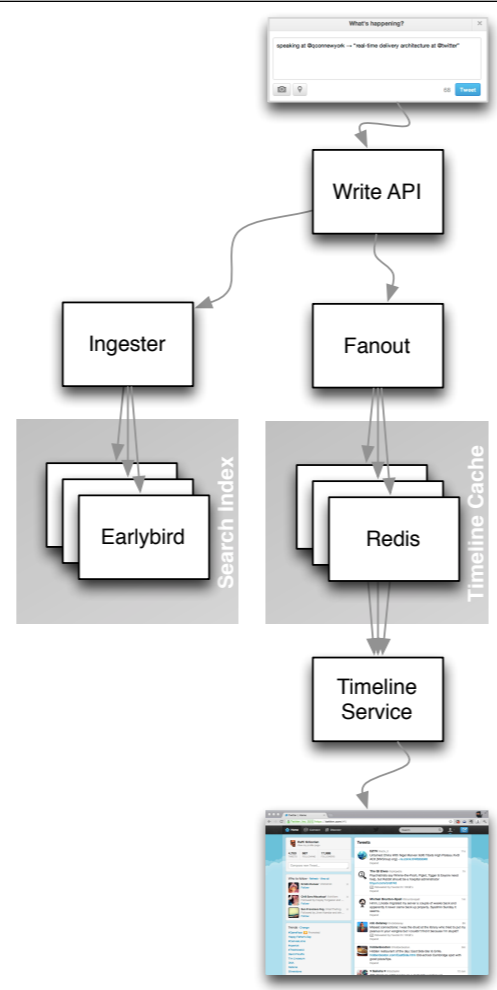


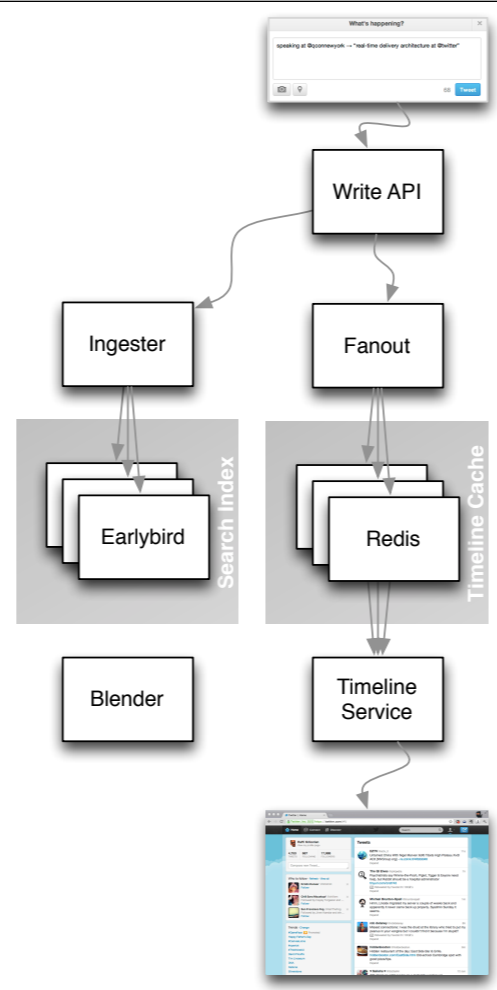


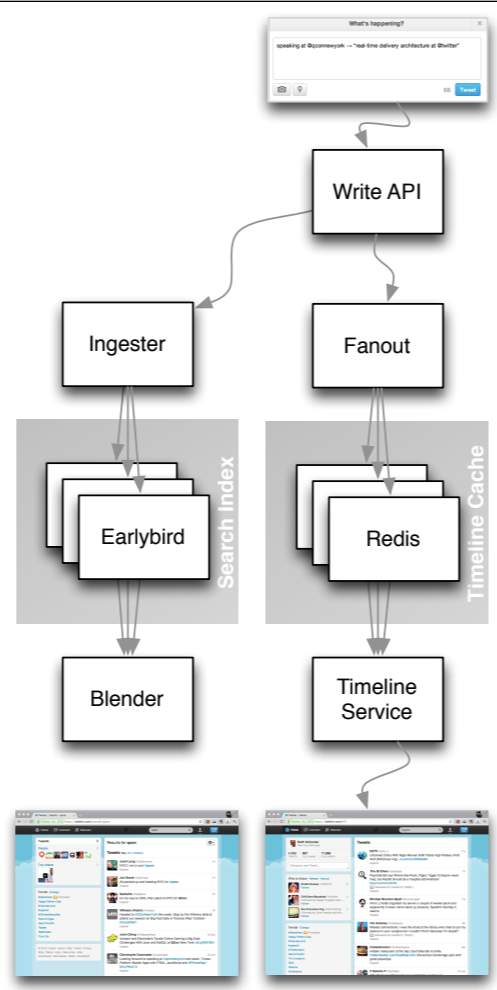


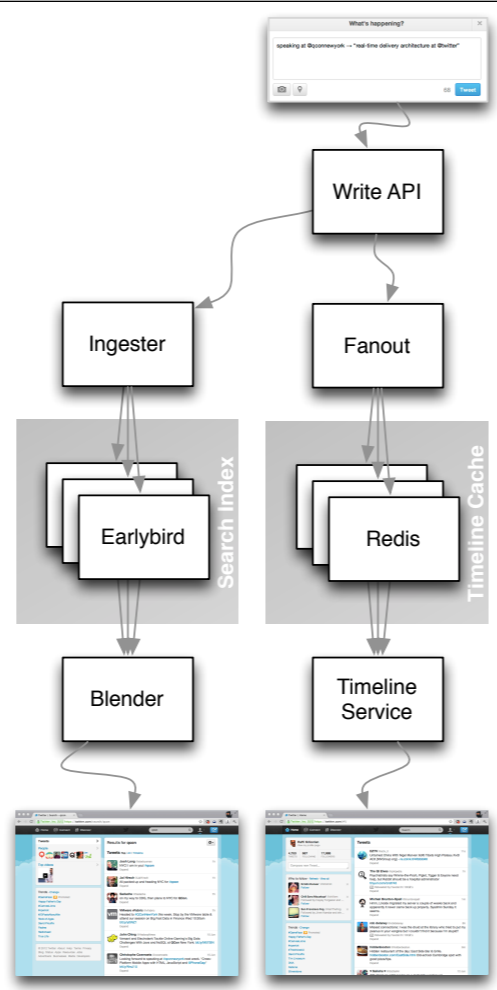


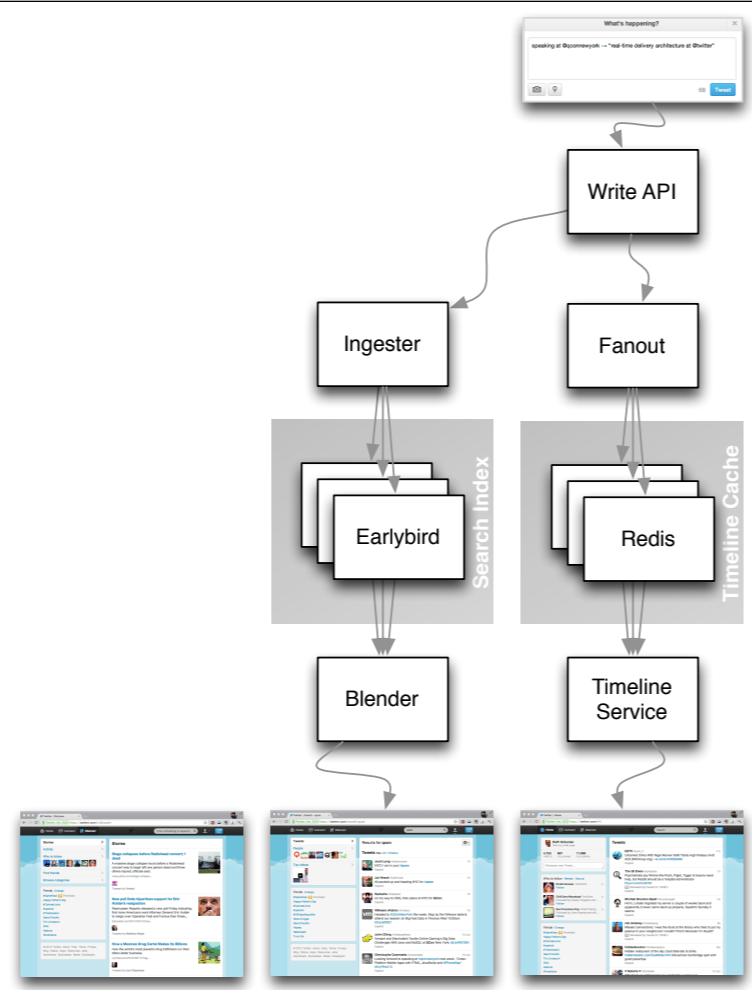


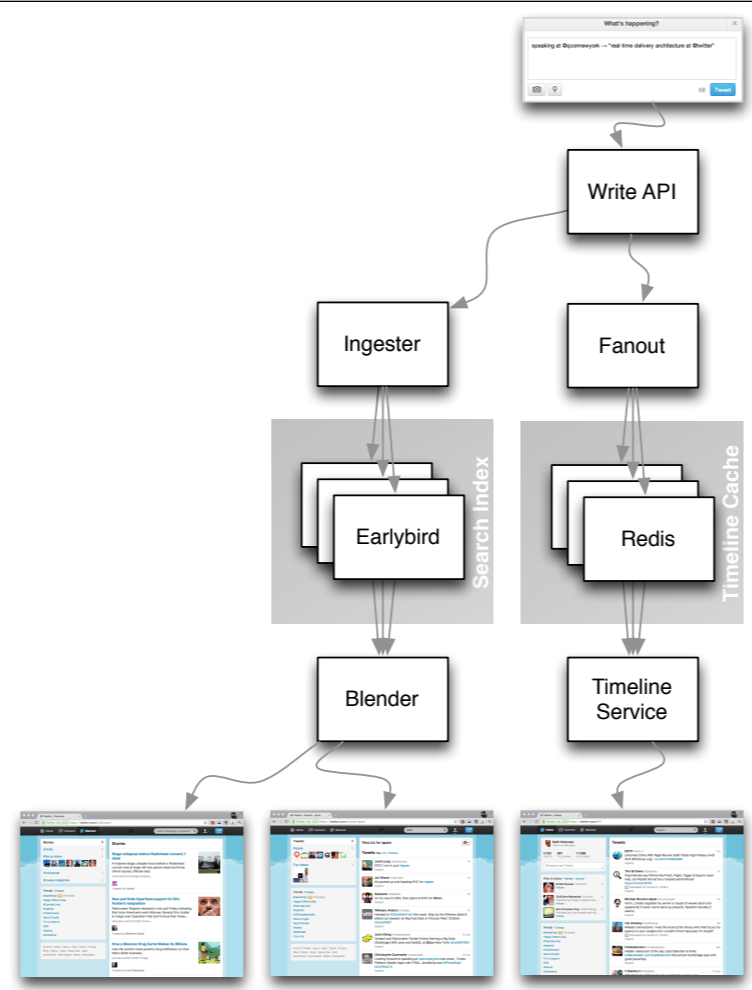


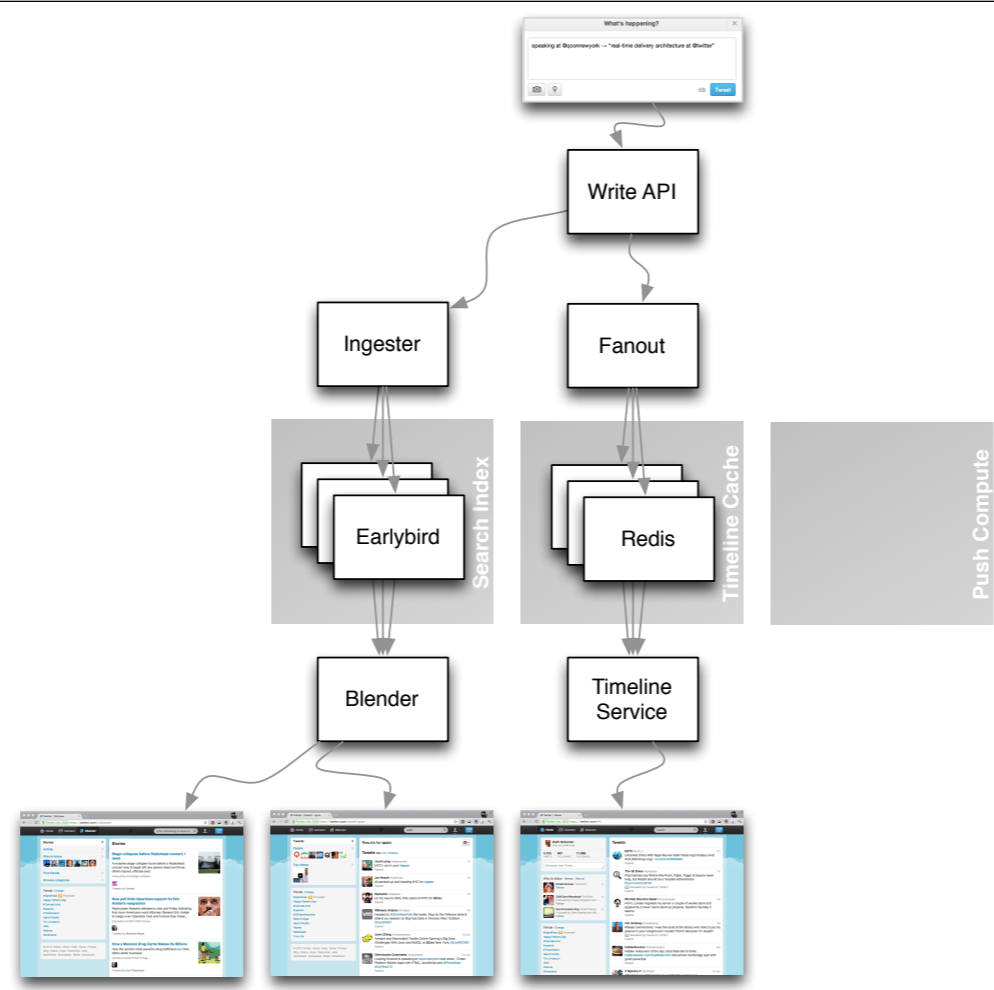


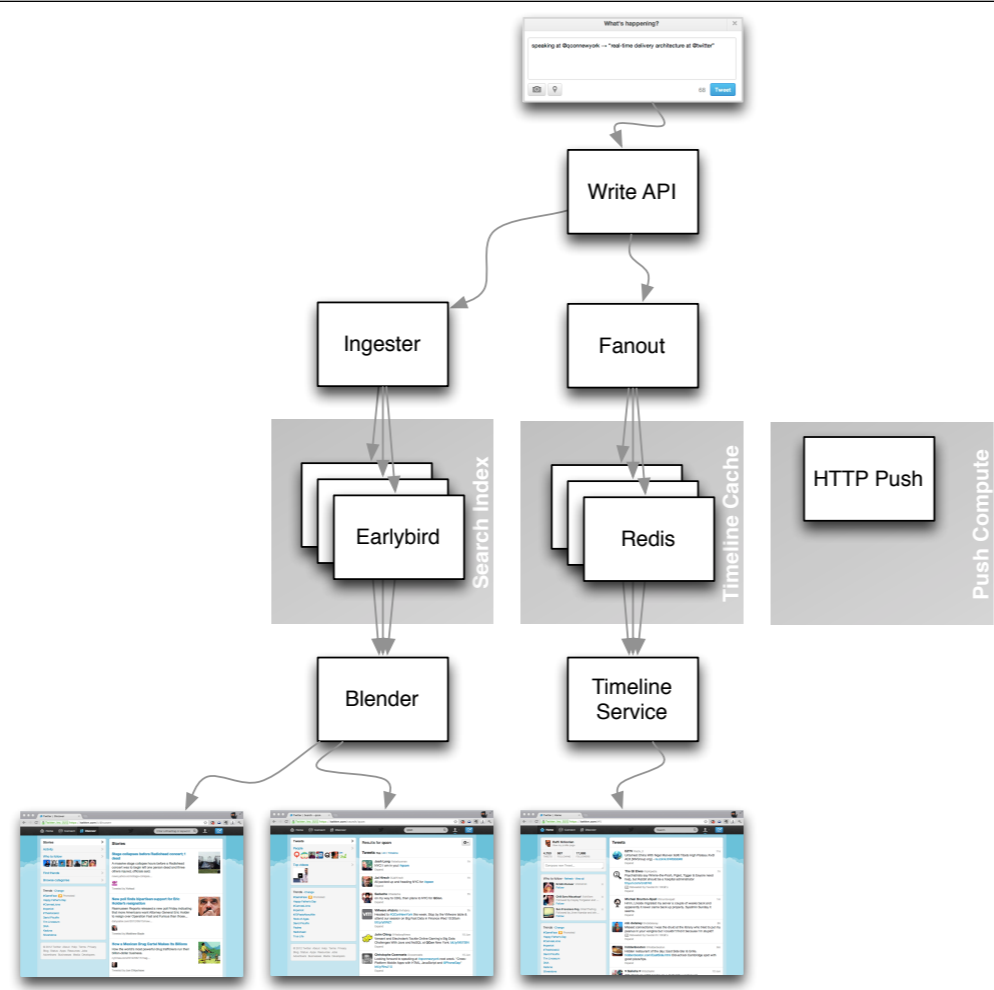


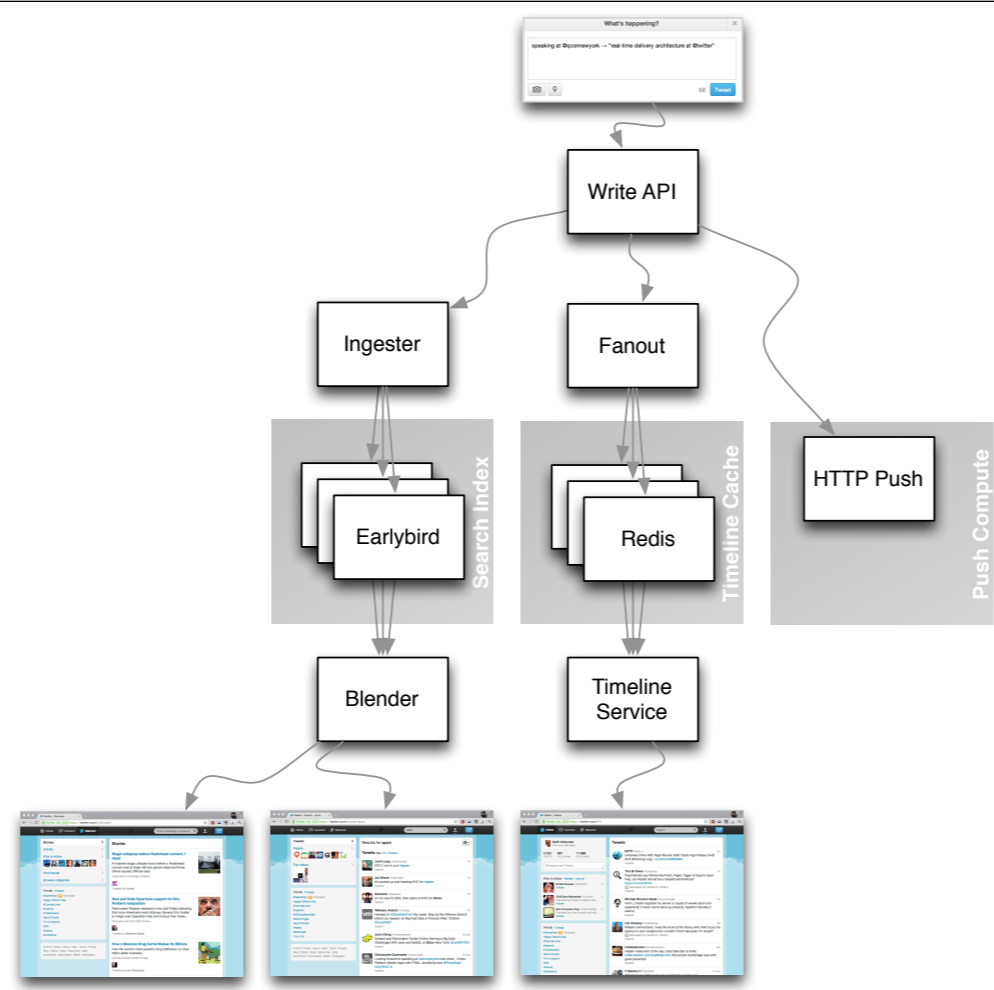


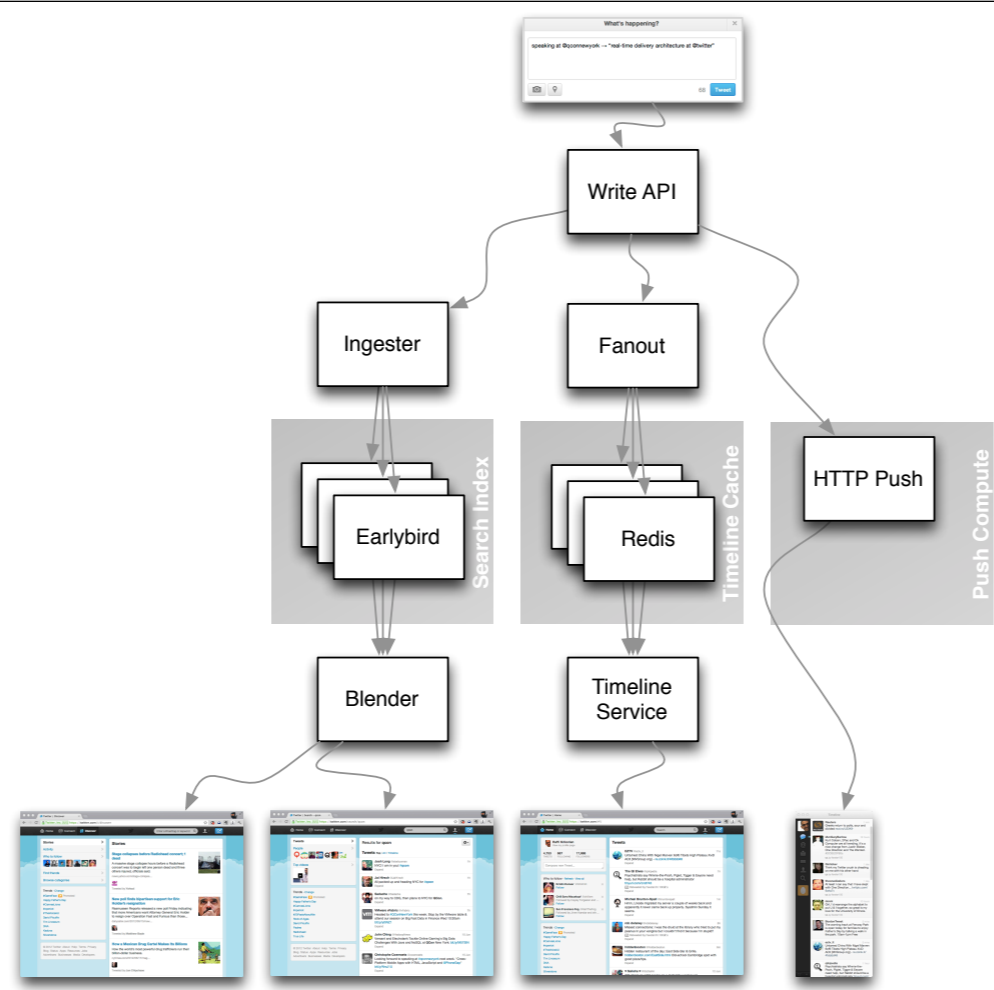


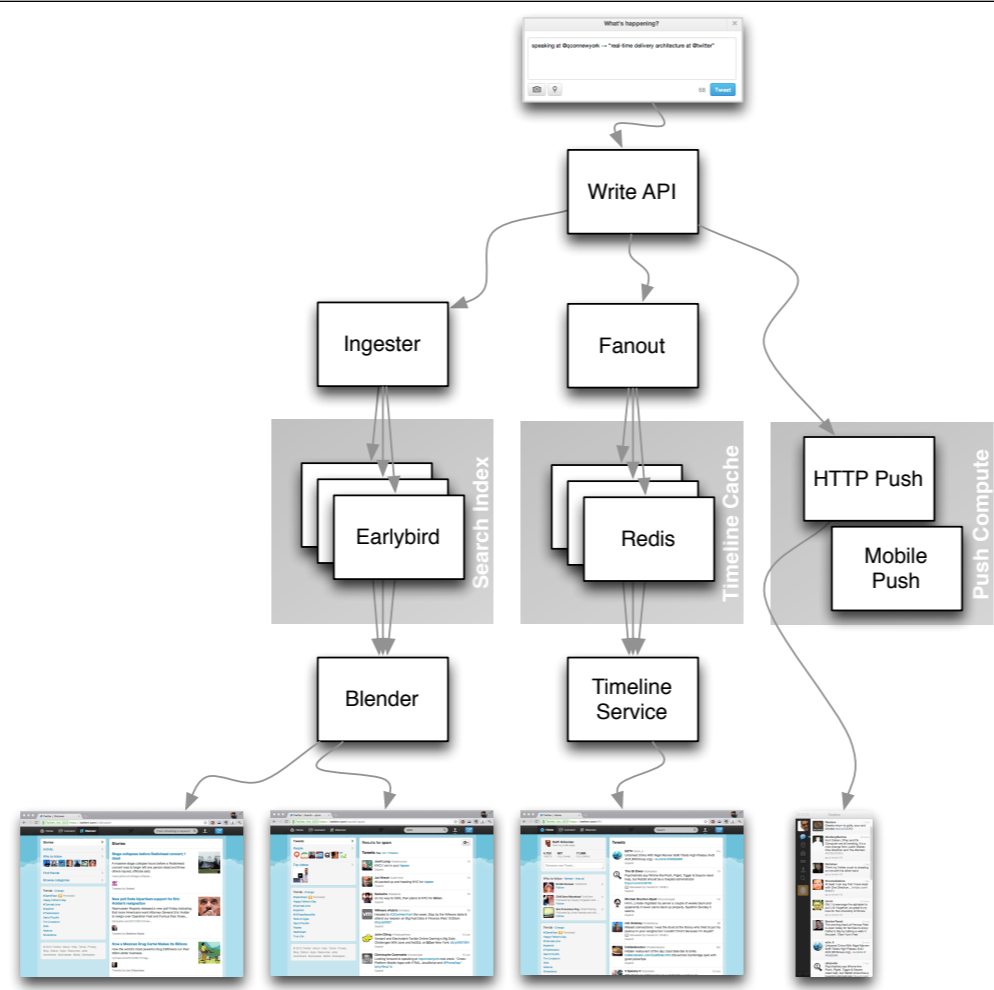


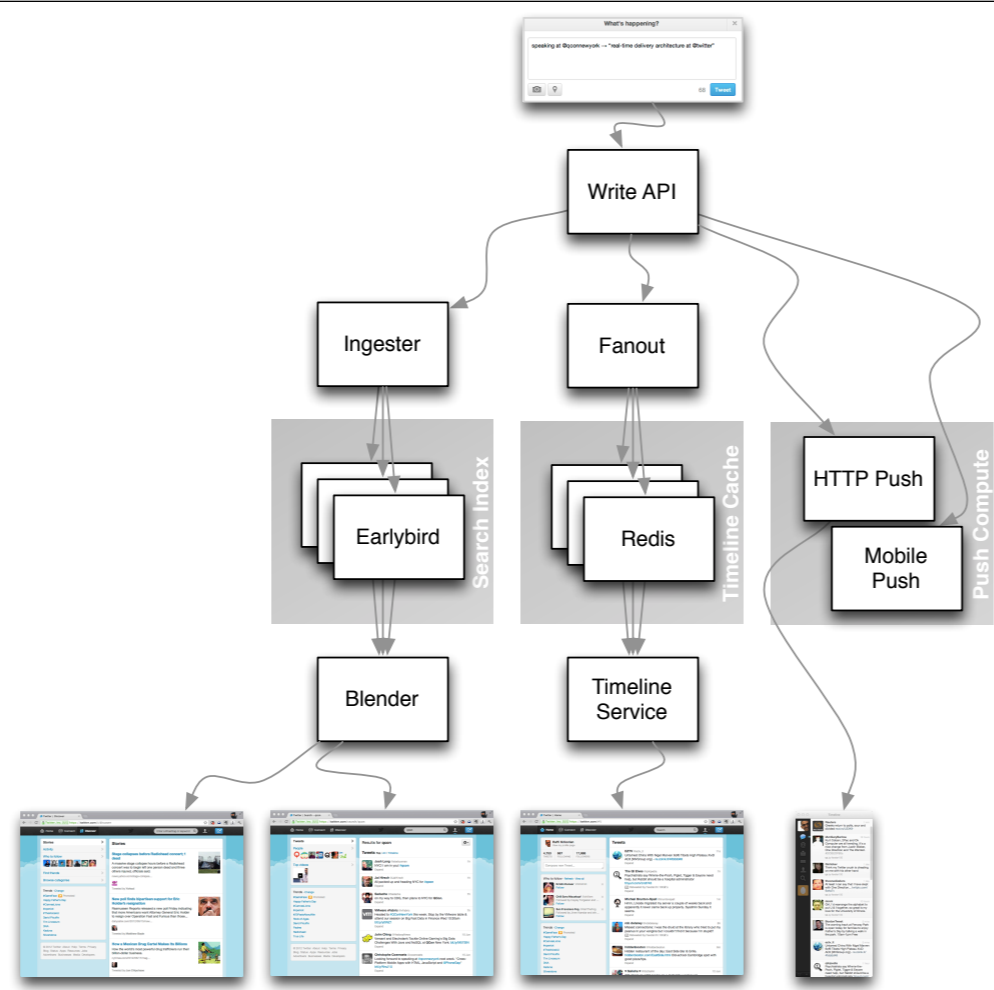


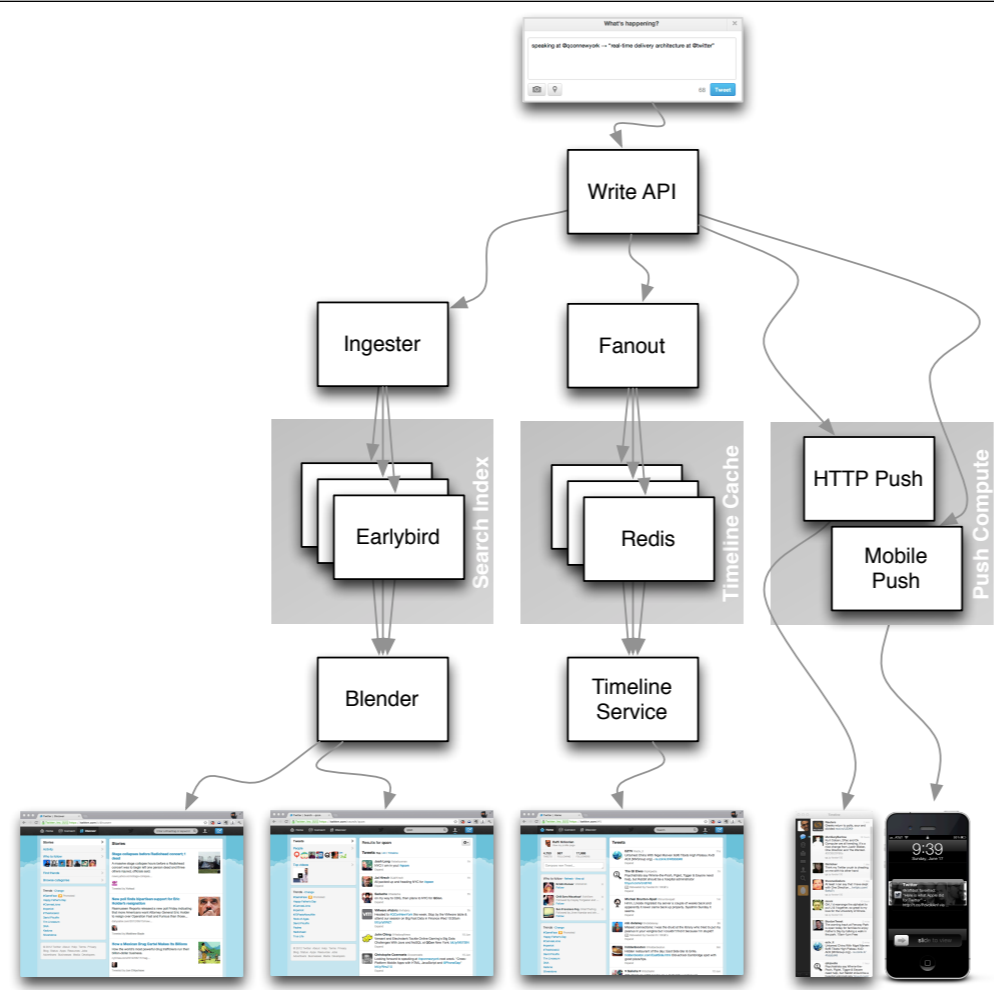












Timeline Query Statistics

→ >240m active users

→ 300k qps poll-based timelines
@1ms p50 / 4ms p99

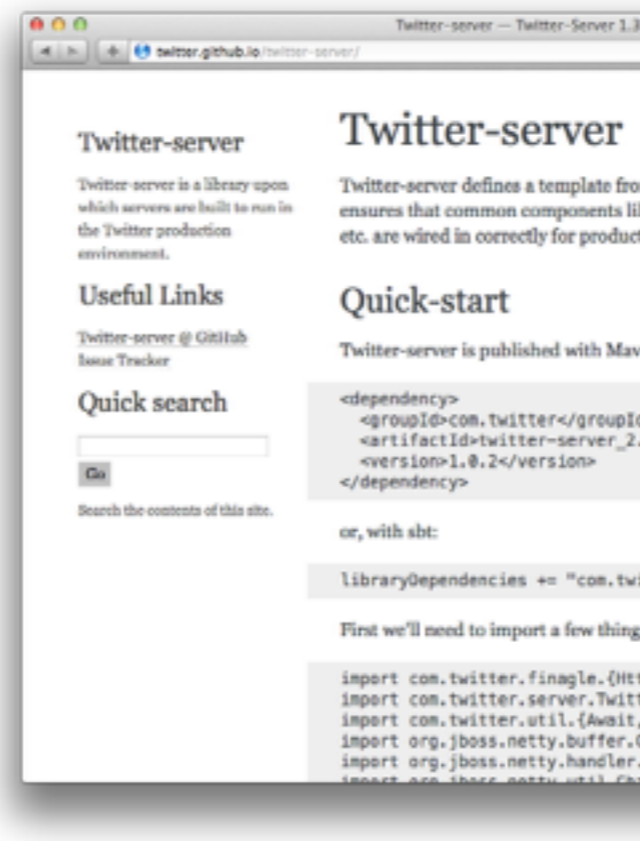


How to build a Service

@bmdhacks
qcon london 2014

Twitter-server

- configuration
- administration
- logging
- lifecycle
- metrics

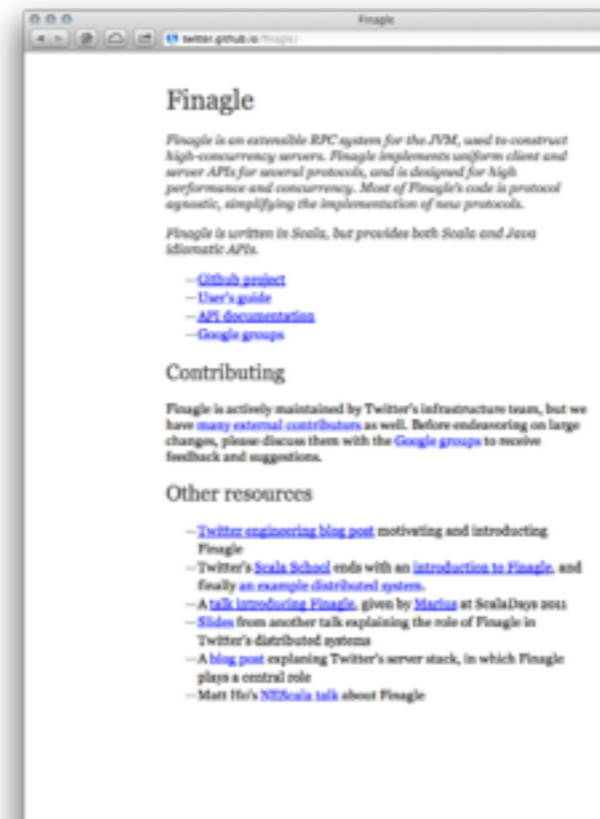


@bmdhacks
qcon london 2014

Finagle is an extensible RPC system for the JVM, used to construct high-concurrency servers

finagle

- service discovery
- load balancing
- retrying
- thread/connection pooling
- stats collection
- distributed tracing



@bmdhacks
qcon london 2014

Finagle is an extensible RPC system for the JVM, used to construct high-concurrency servers

Your Service as a Function

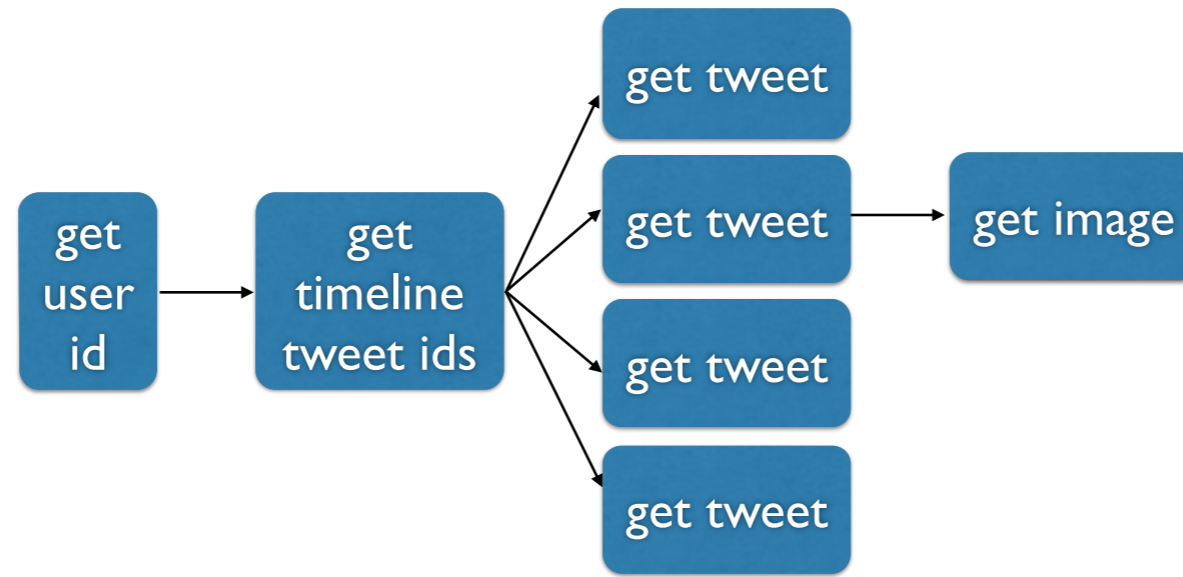
```
trait Service[Req,Rep] extends (Req => Future[Rep])
```

@bmdhacks
qcon london 2014

Future[T]

- Composable
- Concurrent
- Easy to reason about

Composable



@bmdhacks
qcon london 2014

Composable

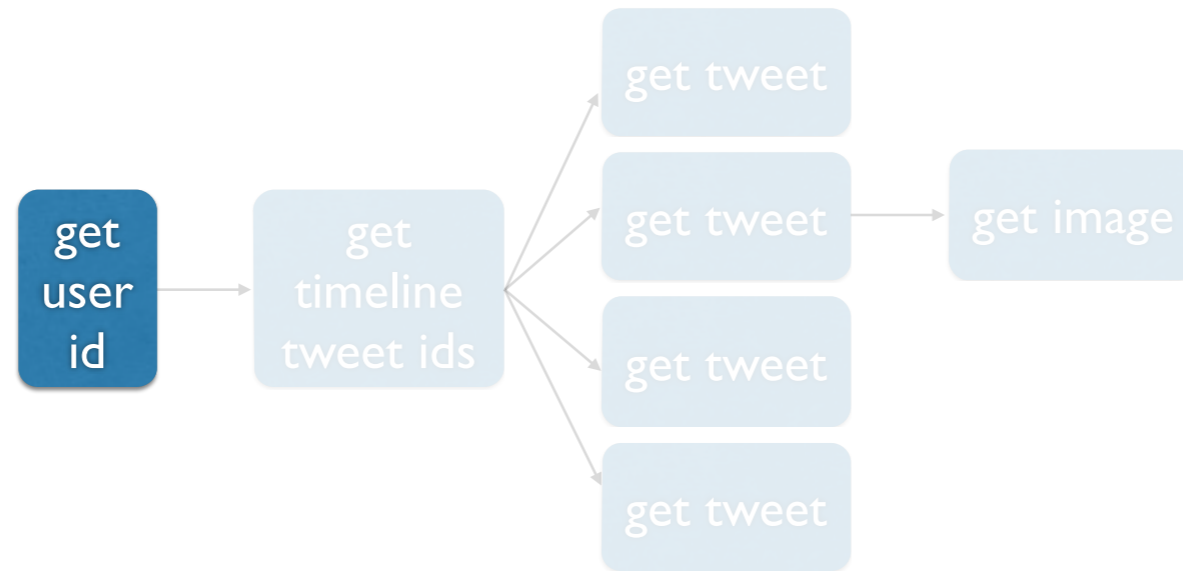
```
val user_id = Future(23)

val ftl = user_id.flatMap { id => getTimeline(id) }

val tweets: Future[Seq[Tweet]] =
  ftl.flatMap { tl =>
    val future_tweets = tl.tweet_ids.map { id =>
      val tweet = getTweet(id)
      tweet.map { t =>
        if (t.hasImage) getImage(t) else t
      }
    }
    Future.collect(future_tweets)
  }
}
```

Composable

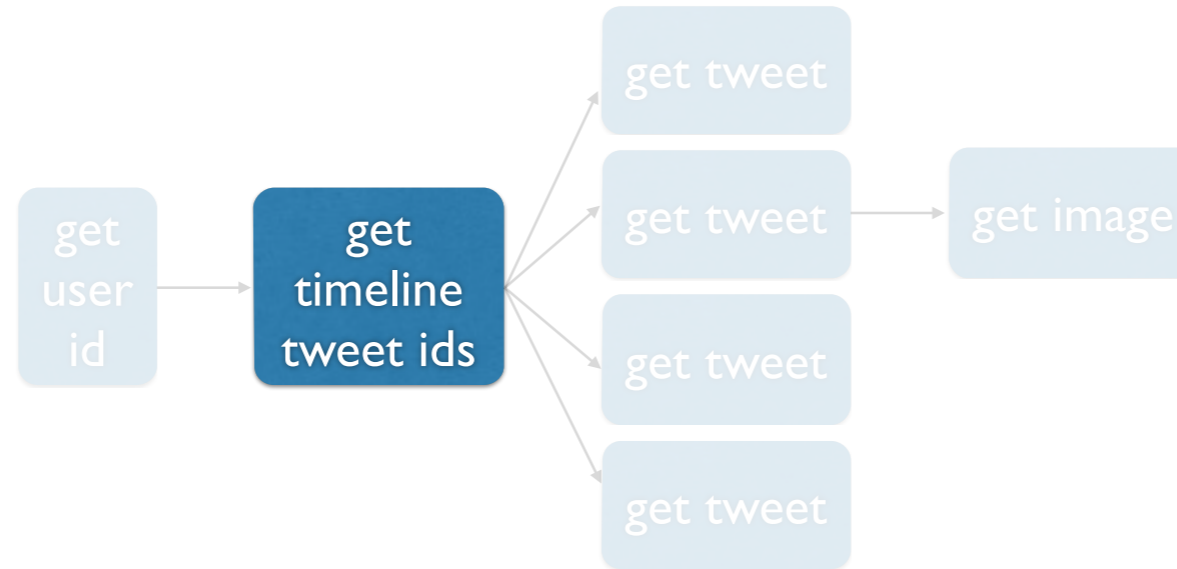
```
val user_id = Future(23)
```



@bmdhacks
qcon london 2014

Composable

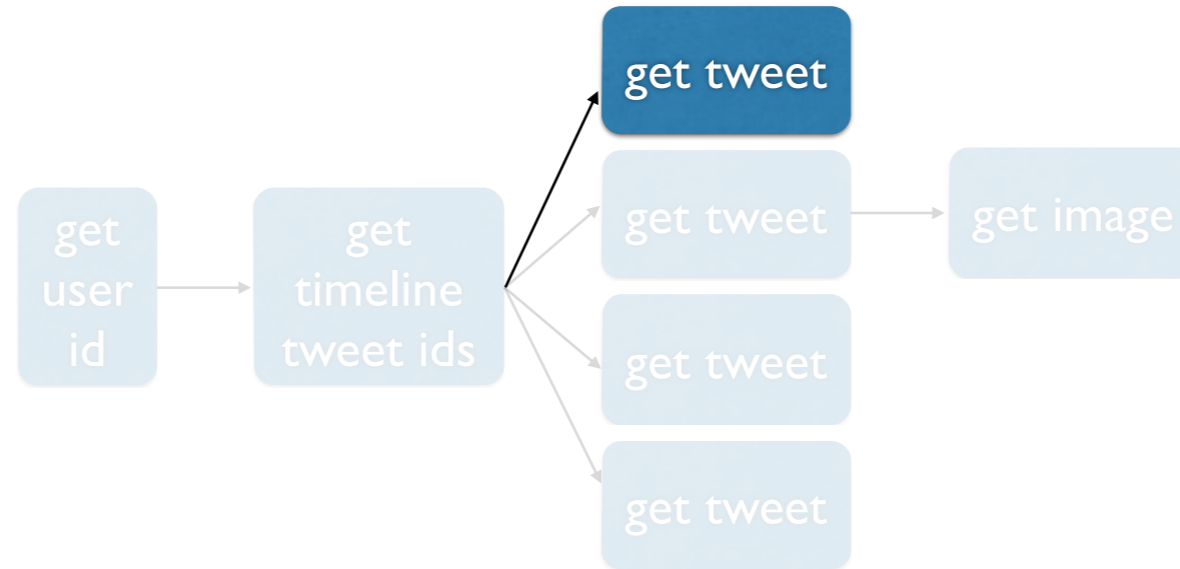
```
val ftl = user_id.flatMap { id => getTimeline(id) }
```



@bmdhacks
qcon london 2014

Composable

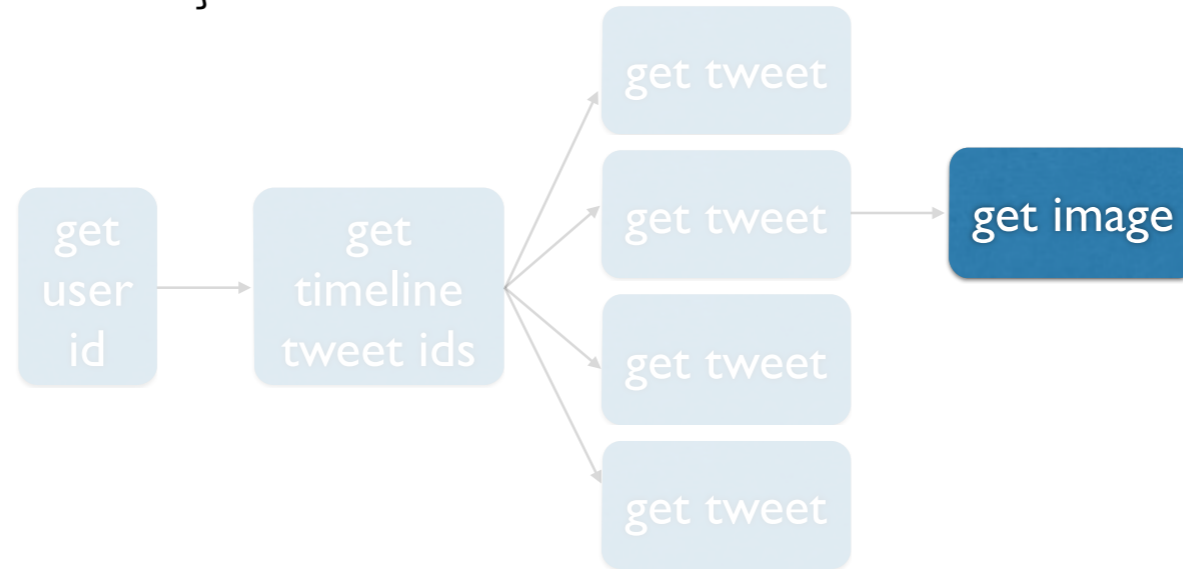
```
val future_tweets = tl.tweet_ids.map { id =>  
  val tweet = getTweet(id)
```



@bmdhacks
qcon london 2014

Composable


```
tweet.map { t =>
  if (t.hasImage) getImage(t) else t
}
```



@bmdhacks
qcon london 2014

Composable

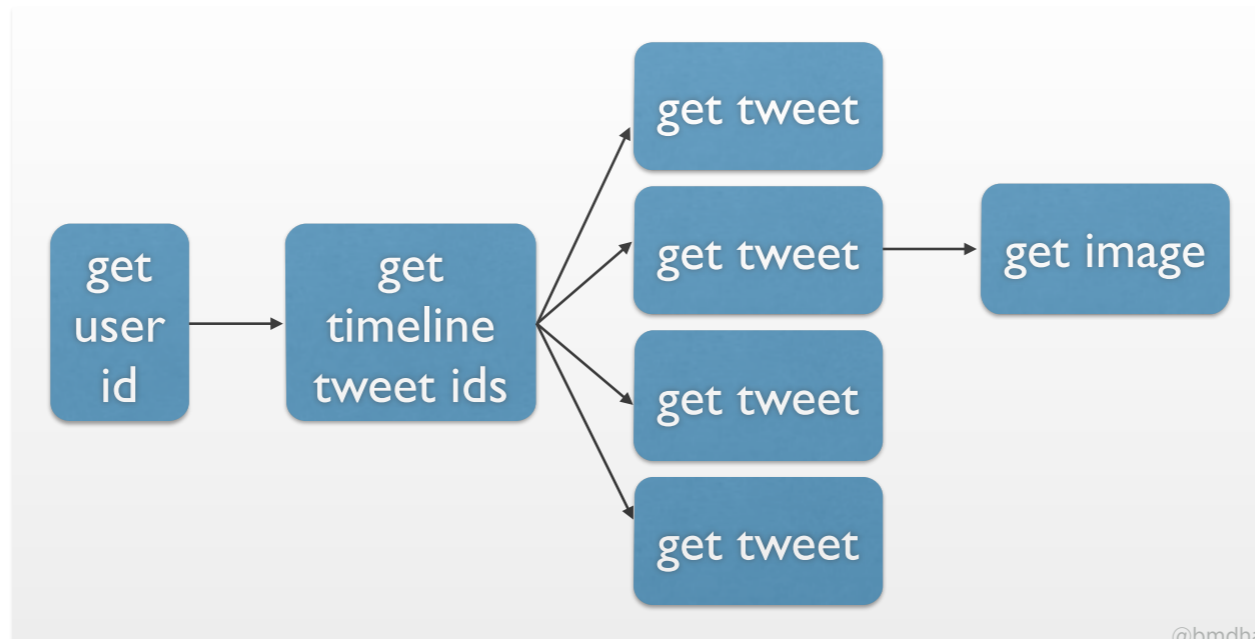
`Future.collect(future_tweets)`



@bmdhacks
qcon london 2014

Composable

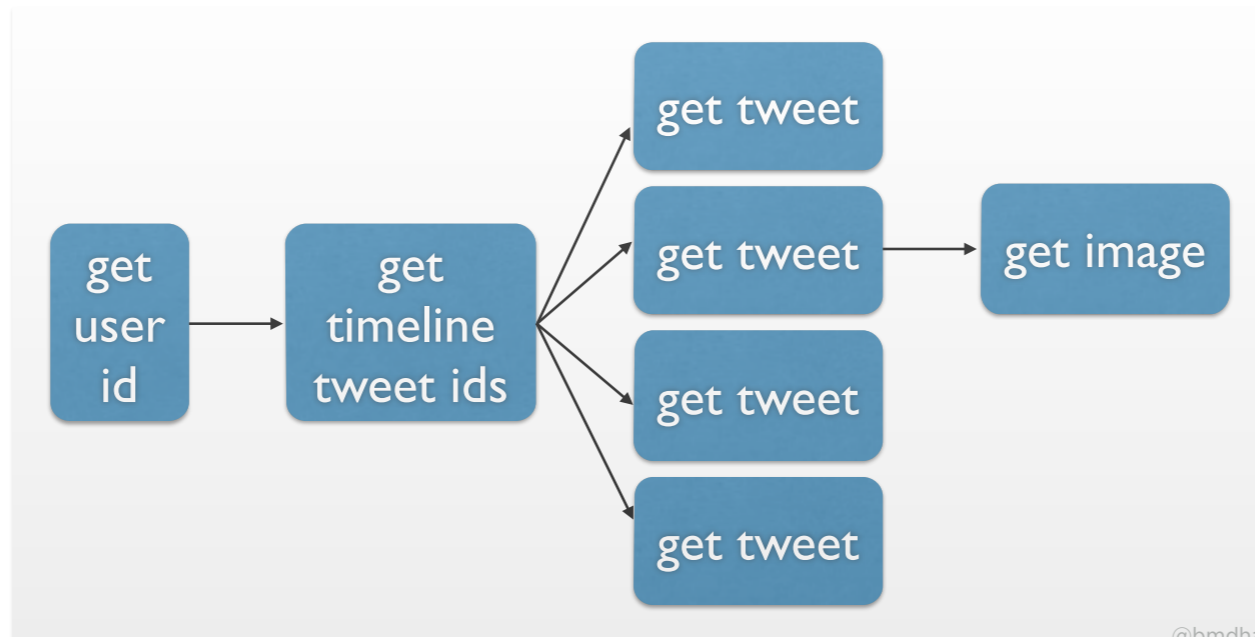
```
val timeline: Future[Timeline]
```



@bmdhacks
qcon london 2014

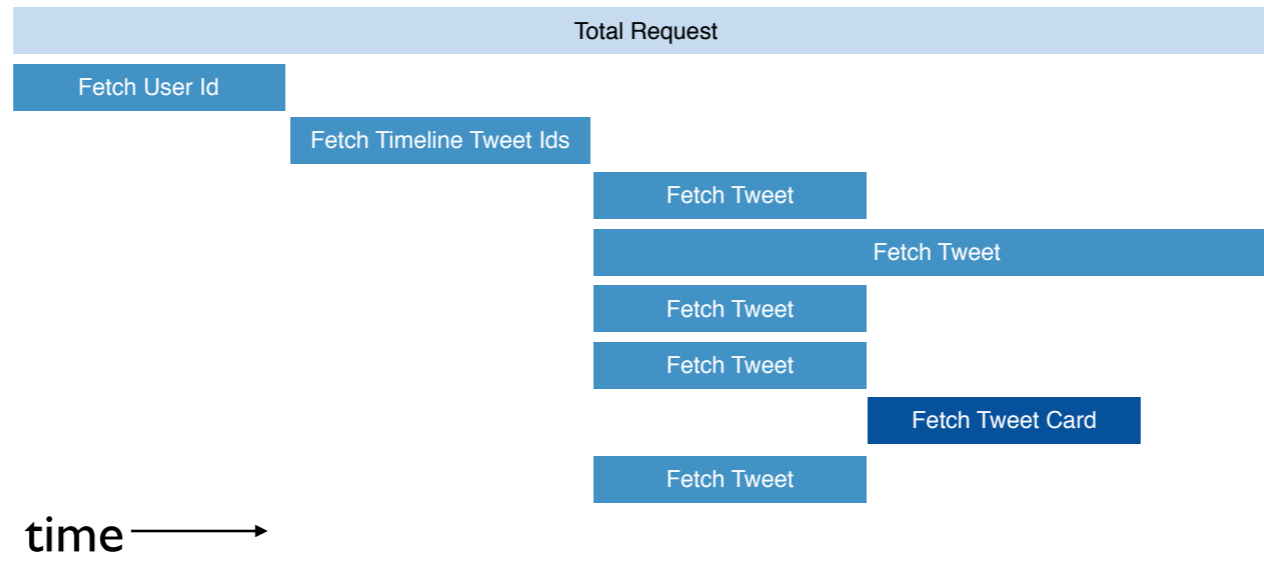
Concurrent

```
val timeline: Future[Timeline]
```



@bmdhacks
qcon london 2014

Concurrent



What we
want to do



How the threads
execute it

Separate semantics from execution

Your Service as a Function

```
trait Service[Req,Rep] extends (Req => Future[Rep])
```

@bmdhacks
qcon london 2014

(Also your Client)

```
trait Service[Req,Rep] extends (Req => Future[Rep])
```

@bmdhacks
qcon london 2014

Your Service

Twitter-Server

Finagle

JVM

Operating System

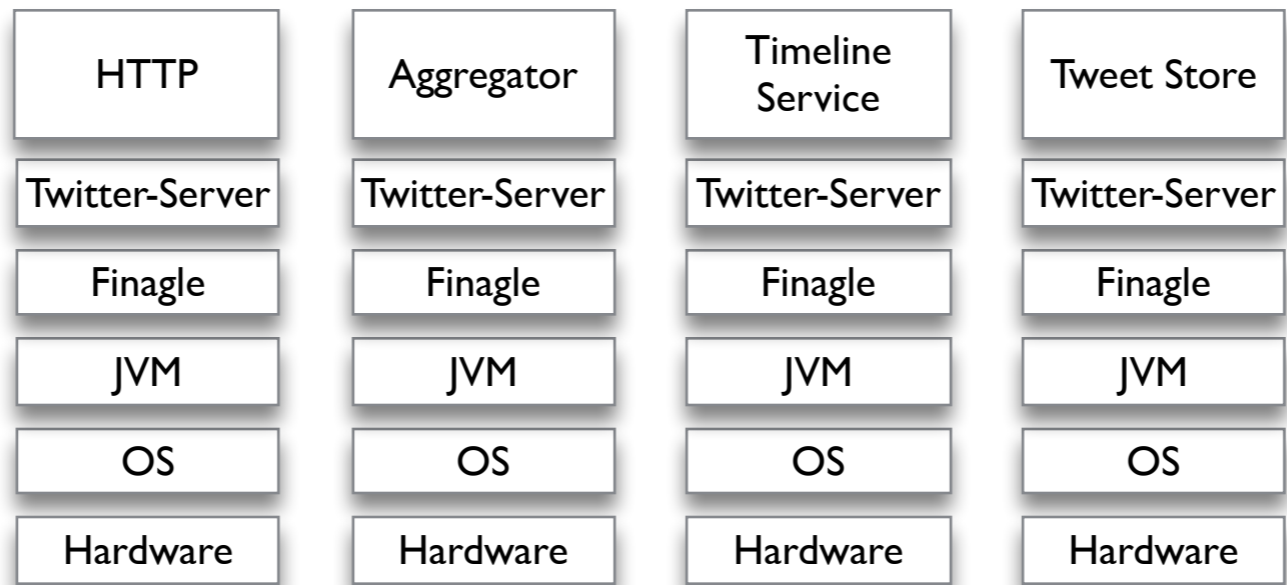
Hardware

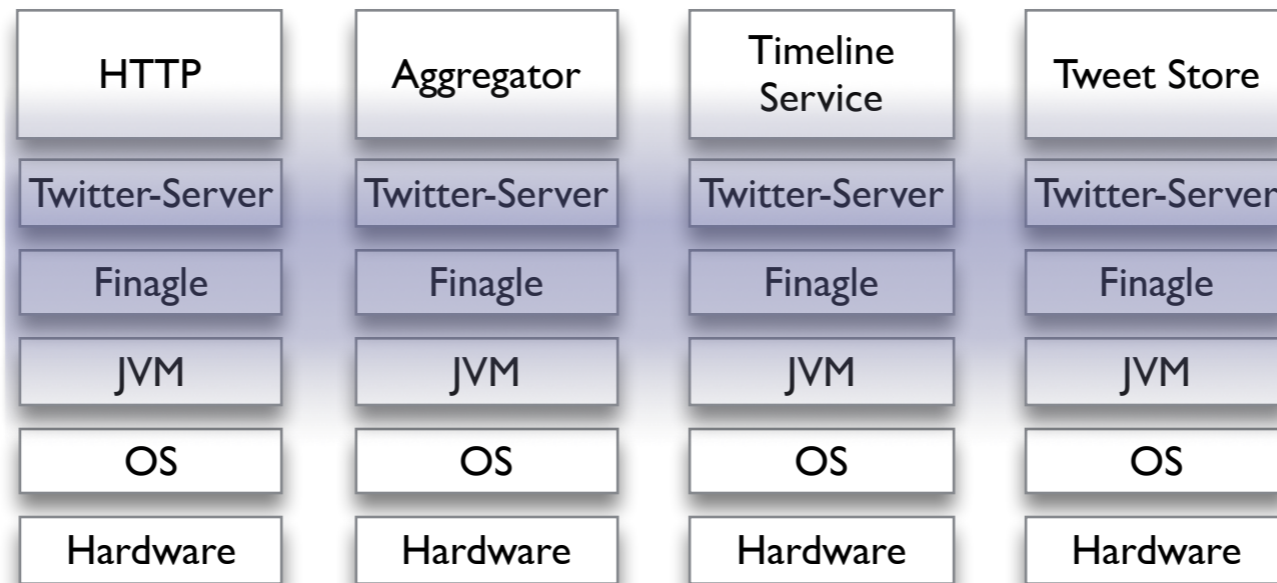
@bmdhacks
qcon london 2014

Monorail



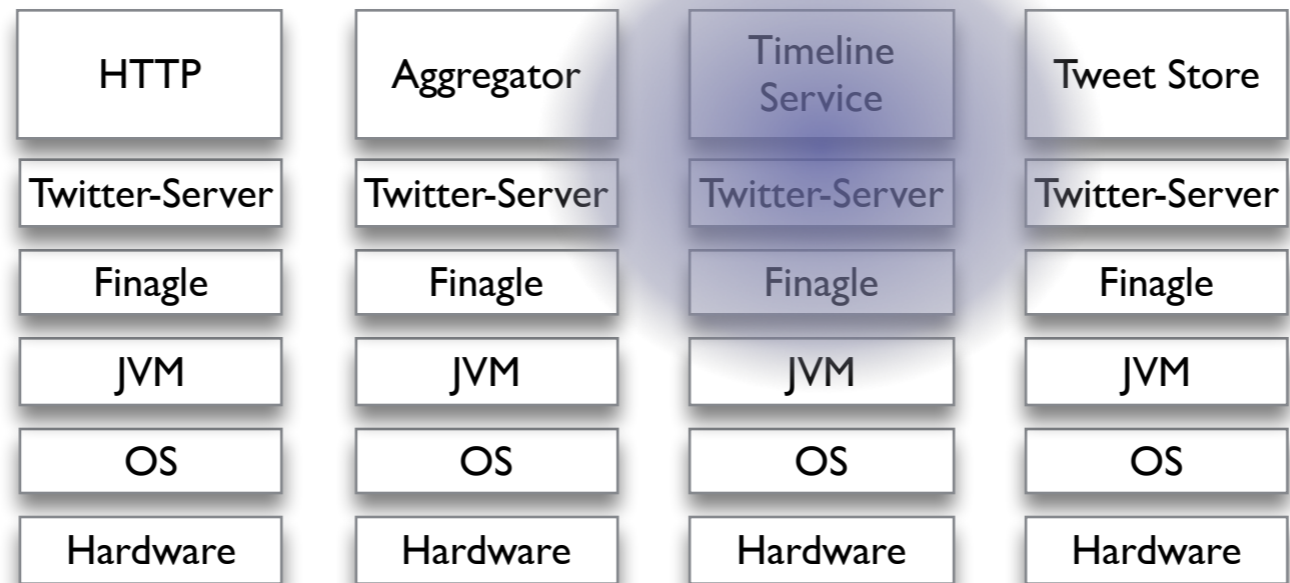
@bmdhacks
qcon london 2014





@bmdhacks
qcon london 2014

Expertise can be applied in both directions of our stack. My team, CSL, works on Twitter-Server and Finagle. These are used by all services, however we don't have a ton of insight into what those services do.



@bmdhacks
qcon london 2014

Other engineers focus on specific services, and extend their knowledge down the stack for their specific use case.

@bmdhacks
qcon london 2014



Use Statistics to Monitor Behavior

@bmdhacks
qcon london 2014


Why Stats Matter

- Increased traffic velocity makes individual events less important
- More vertical components requires more measurement
- Horizontal SOA needs more measurement

Failures are OK


- 300,000rps
- 99.99% success rate
- 30 failures per second
- Failure Rate is more informative than individual failures

Why Stats Matter




@bmdhacks
qcon london 2014

Why Stats Matter



@bmdhacks
qcon london 2014

Why Stats Matter



@bmdhacks
qcon london 2014

Why Stats Matter

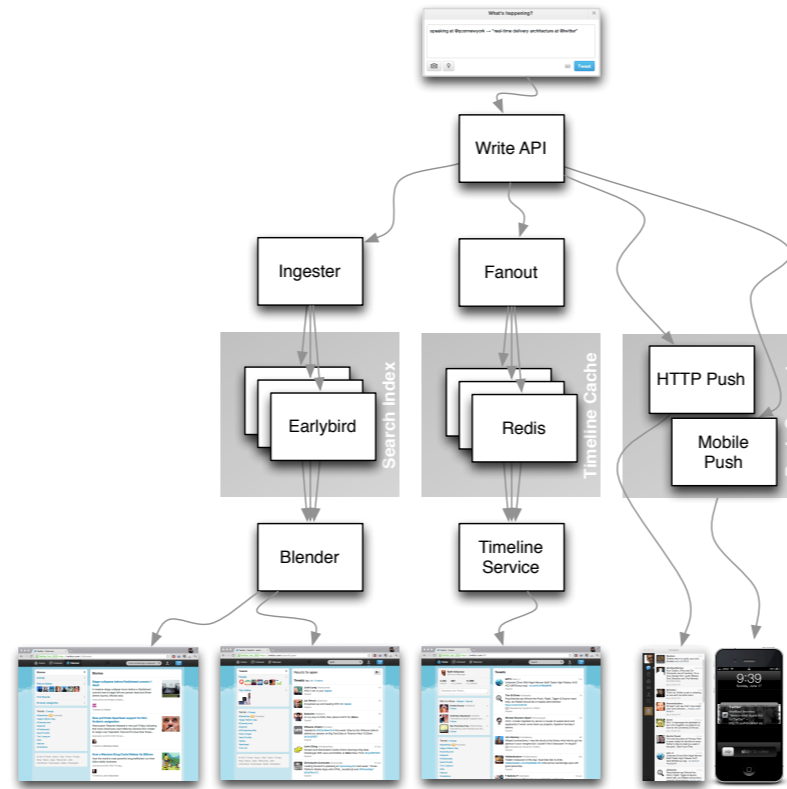
	Before	After
Median	10ms	11ms
p90	50ms	55ms
p99	100ms	400ms
p999	500ms	500ms

@bmdhacks
qcon london 2014

Measure the Long Tail

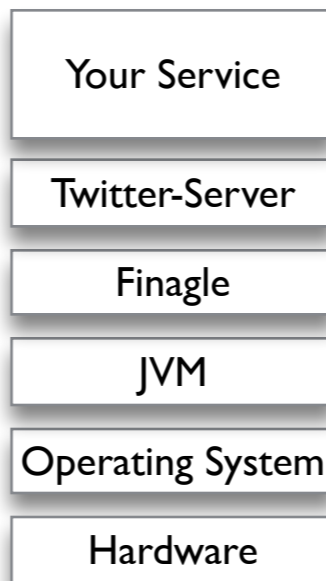
- Tail effects are cumulative
- This is why request-level concurrency is important

Tail effects are cumulative



This is also why concurrency is important.

Measuring the Vertical Stack

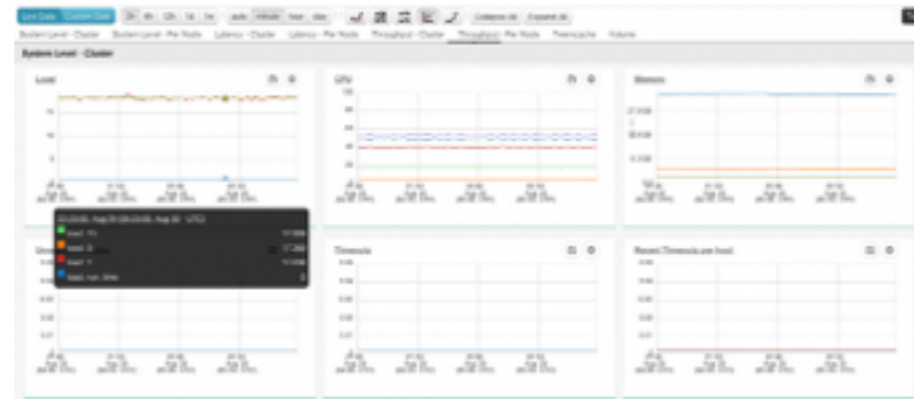


Viz: Measuring the Vertical Stack



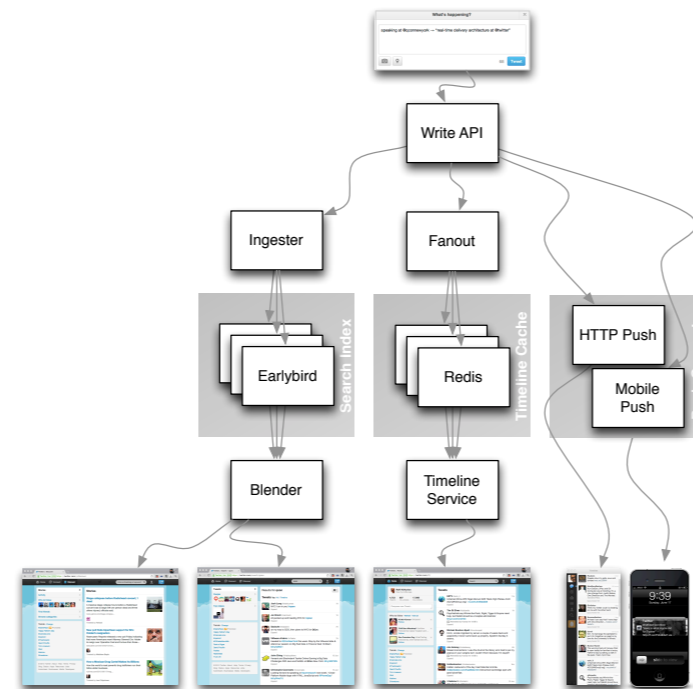
@bmdhacks
qcon london 2014

Viz: Measuring the Vertical Stack



@bmdhacks
qcon london 2014

Measuring the Horizontal SOA

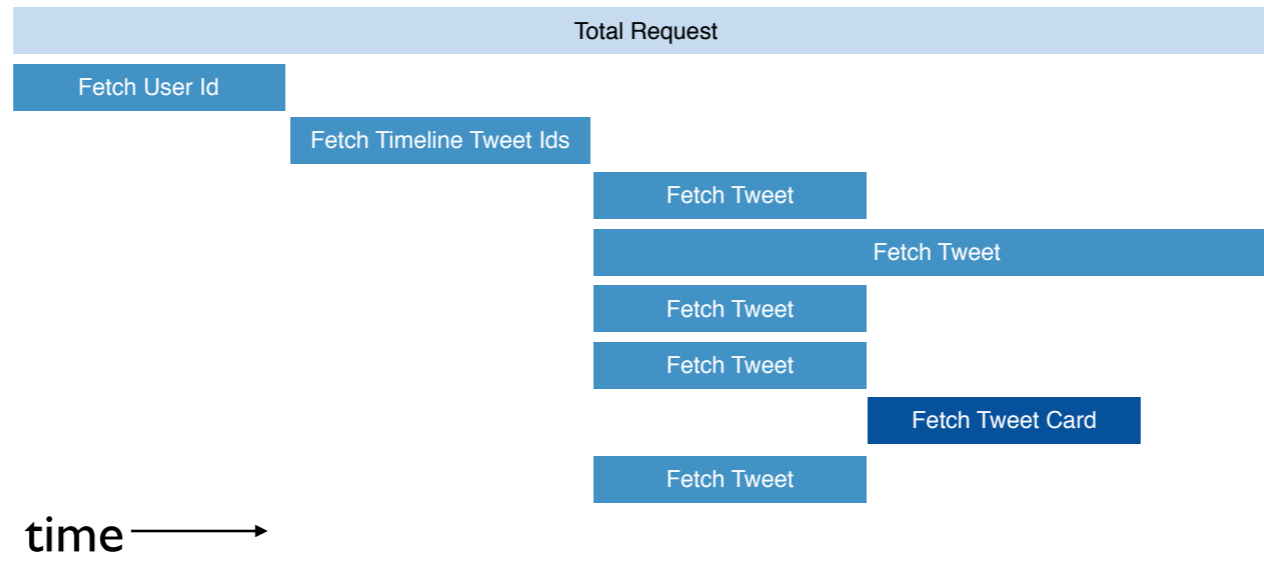


Zipkin

→ It's awesome

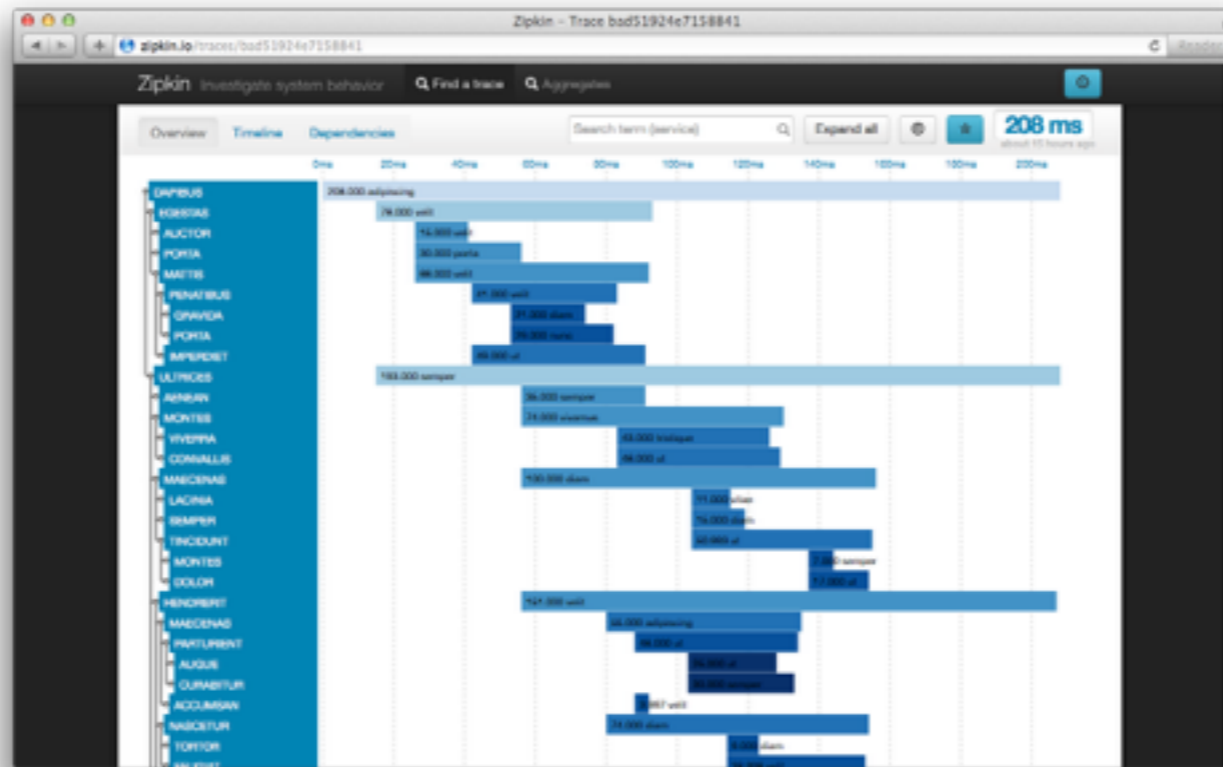
@bmdhacks
qcon london 2014

Remember This?



@bmdhacks
qcon london 2014


Zipkin



@bmdhacks
qcon london 2014

Zipkin

- Distributed Tracing System
- Open Source
- Modeled after Google's 'Dapper' paper



ndhacks
ion 2014



thanks!

References

- Timelines at Scale
<http://www.infoq.com/presentations/Twitter-Timeline-Scalability>
- Finagle
<http://twitter.github.io/finagle/>
- Twitter-Server
<http://twitter.github.io/twitter-server/>
- Zipkin
<http://twitter.github.io/zipkin/>
- Twitter's Observability Stack
<https://blog.twitter.com/2013/observability-at-twitter>
- Your Server as a Function
<http://monkey.org/~marius/funsrv.pdf>

Please evaluate
my talk via the
mobile app!

