



Main sponsor



From RoR & the Ruby VM to the JVM

Raffi Krikorian / @raffi



@DanaDanger
Dana Contreras

Tweet!

1 min ago via [Twitter for iPhone](#) ☆ Favorite ↻ Retweet ↩ Reply



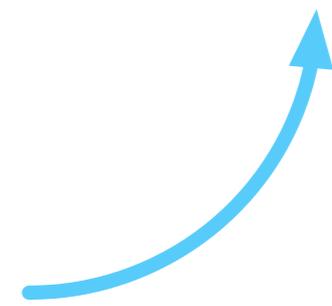
 **@DanaDanger**
Dana Contreras

Tweet!

1 min ago via [Twitter for iPhone](#) ☆ Favorite ↻ Retweet ↩ Reply



> 2.5E9 “deliveries” a day





Ruby on Rails

Sustainable productivity for web-application development



HAS...

- huge number of concurrent connections
- lots of I/O
- few persistent objects

Gosai Kiji

Phasianus versicolor also known as Japanese Pheasant is a bird of the lowlands. Closely related to the Common Pheasant, the cock is distinguished by dark green plumage on breast and mantle. The male has an iridescent violet neck, red bare facial skin and purplish green tail. The female is smaller than male, and has a dull brown plumage with dark spots.





Twitter Engineering: Building a Faster

1P + http://engineering.twitter.com/2011/03/building-faster-ruby

Share Report Abuse Next Blog»

Twitter Engineering



FRIDAY, MARCH 4, 2011

Building a Faster Ruby Garbage Collector

Since late 2009, much of `www.twitter.com` has run on Ruby Enterprise Edition, a modified version of the standard MRI 1.8.7 Ruby interpreter. At the time, we asked the REE team to integrate some third-party patches that allowed us to tune the garbage collector for long-lived workloads. We knew this was not a perfect choice, because a new runtime (even MRI 1.9x) would introduce compatibility problems, and we had indicated that alternative runtimes are not necessarily faster for our workload. Nevertheless, the CPU cost of REE remained too high.

To address this problem, we decided to explore options for optimizing the Ruby garbage collector. We called this effort Project Kiji, after the Japanese bird.

INEFFICIENT GARBAGE COLLECTION

Our performance measurements revealed that even after our patches, the Ruby garbage collector uses a significant fraction of the CPU for running the garbage collector on `www.twitter.com`. This is largely because MRI's garbage collector uses a single heap:

- The garbage collector's naive stop-the-world mark-and-sweep process accesses the entire memory set several times. It first marks all objects at the "root-set" level as "in-use" and then reexamines all the objects to release the memory of those not in use. Additionally, the collector suspends the program during every sweep, thereby periodically "freezing" some of the program's execution.
- The collection process is not generational. That is, the collector does not move objects between heaps; they all stay at the same address for their lifetime. The resulting fragmented memory extracts a penalty in bookkeeping cost because it can neither be consolidated nor discarded.

We needed to make the garbage collector more efficient but had limited options.





NEEDS...

- ability to handle server workloads
- an efficient language



NEEDS...

- ability to handle server workloads
- an efficient language

 **Scala**

The Scala logo consists of three horizontal red bars of varying lengths, stacked vertically, followed by the word 'Scala' in a bold, white, sans-serif font.



FIRST CHOSE Scala

- > fast
- > functional and expressive
- > statically typed
- > concurrent
- > beautiful



NEEDS...

- ability to handle server workloads
- flexibility in language

 **Scala**

The Scala logo consists of three horizontal red bars of varying lengths, stacked vertically, followed by the word 'Scala' in a bold, white, sans-serif font.



NEEDS...

- ability to handle server workloads
- flexibility in language

 **Scala**

 **Clojure**



NEEDS...

- ability to handle server workloads
- flexibility in language
- a real concurrency model



**IS AN EVENT DRIVEN &
REAL-TIME PROBLEM**



NEEDS...

ability to handle server

finagle |fə'nāgəl|

verb [trans.] informal

obtain (something) by devious or dishonest means : *Ted attended all the football games he could finagle tickets for.*

• [intrans.] act in a devious or dishonest manner : *they wrangled and finagled over the fine points.*

The screenshot shows a web browser window with the title "Finagle, from Twitter". The address bar contains "http://twitter.github.com/finagle/". The page has a dark navigation bar with links for "Finagle", "Quick Start", "Architecture", "Java patterns", and "Scala examples". The main content area has a blue background with the heading "Finagle, from Twitter" in large white text. Below the heading is a paragraph: "Finagle is a network stack for the JVM that you can use to build asynchronous Remote Procedure Call (RPC) clients and servers in Java, Scala, or any JVM-hosted language. Finagle provides a rich set of protocol-independent tools." At the bottom of this section, it says "Finagle is written in Scala on top of Netty." The footer of the page contains three buttons: "LEARN MORE", "JOIN IN", and "FORK ON GITHUB".

om Old

```
abstract class Service[-Req, +Rep]
  extends (Req => Future[Rep]) {
  ...

  /**
   * This is the method to override/implement
   * to create your own Service.
   */
  def apply(request: Req): Future[Rep]

  ...
}

val response = service(request)
```

```
service(request).onSuccess { response =>
  // compute result from response
  println("got response! " + response)
}
```



HAS A HOME TIMELINE

- figure out which tweets to show you
- get those tweets
- get the users that authored those tweets

```
timelineService(userId).andThen { tweetIds =>
  tweetService(tweetIds).andThen { tweets =>
    val userIds = getUserIds(tweets)
    userService(userIds).onSuccess { users =>
      // use users + tweets to produce JSON
    }
  }
}
```

'S FINAGLE = SUBSTRATE

- connection management
- protocol codecs
- transient error handling
- distributed tracing
- service discovery
- observability

```
ServerBuilder()  
  .name("ServiceName")  
  .reportTo(statsReceiver)  
  .tracer(traceReceiver)  
  .codec(Http())  
  .maxConcurrentRequests(1000)  
  .requestTimeout(500.milliseconds)  
  .build(Service[Request, Response])
```

```
ClientBuilder()  
  .cluster(TimelineServiceCluster)  
  .hostConnectionCoresize(5)  
  .hostConnectionLimit(10)  
  .hostConnectionIdleTime(5.seconds)  
  .retries(3)  
  .timeout(500.milliseconds)
```

Timeline
Storage

Tweet
Storage

Social
Graph

User
Storage

**Storage /
Retrieval**



Composition

**Models /
Biz Objects**

**Timeline
Storage**

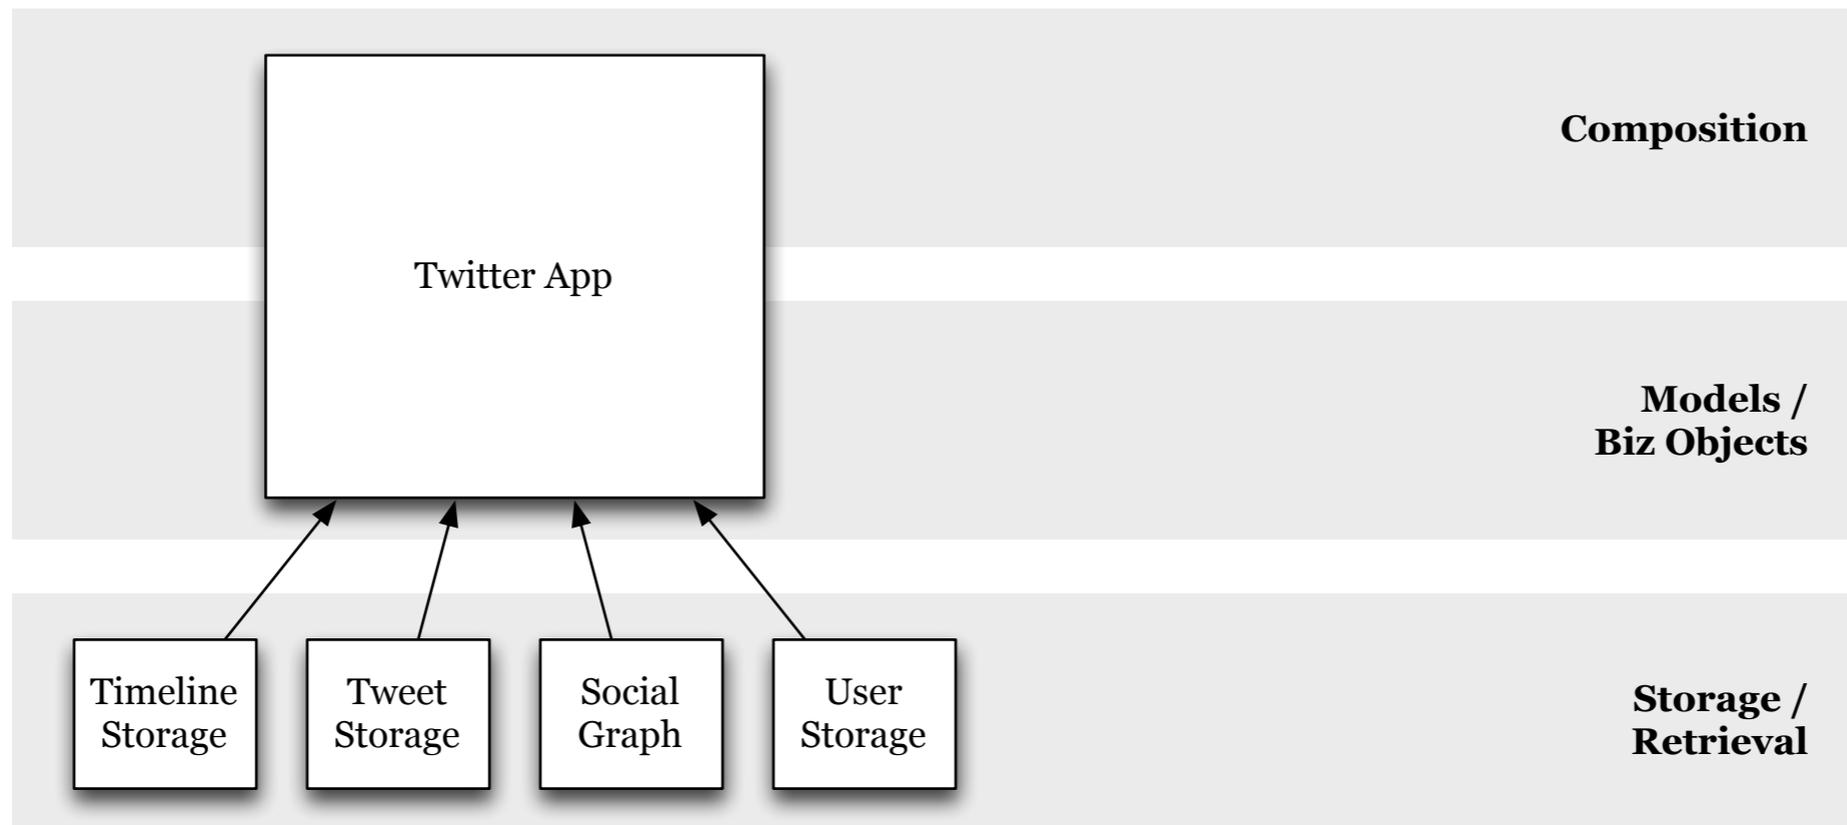
**Tweet
Storage**

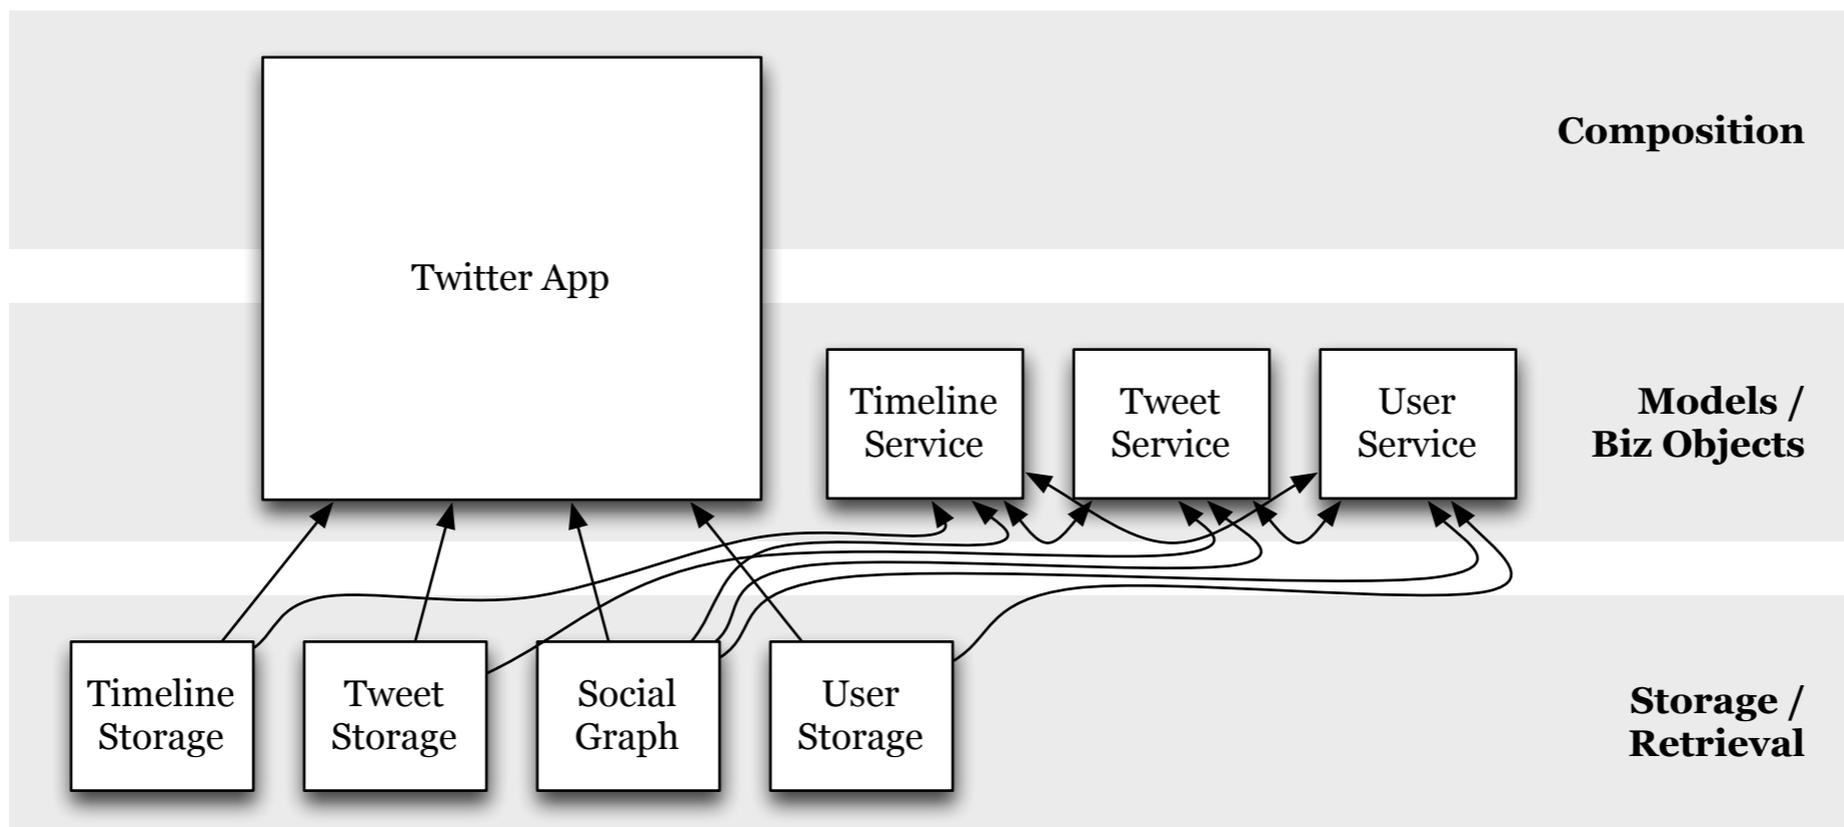
**Social
Graph**

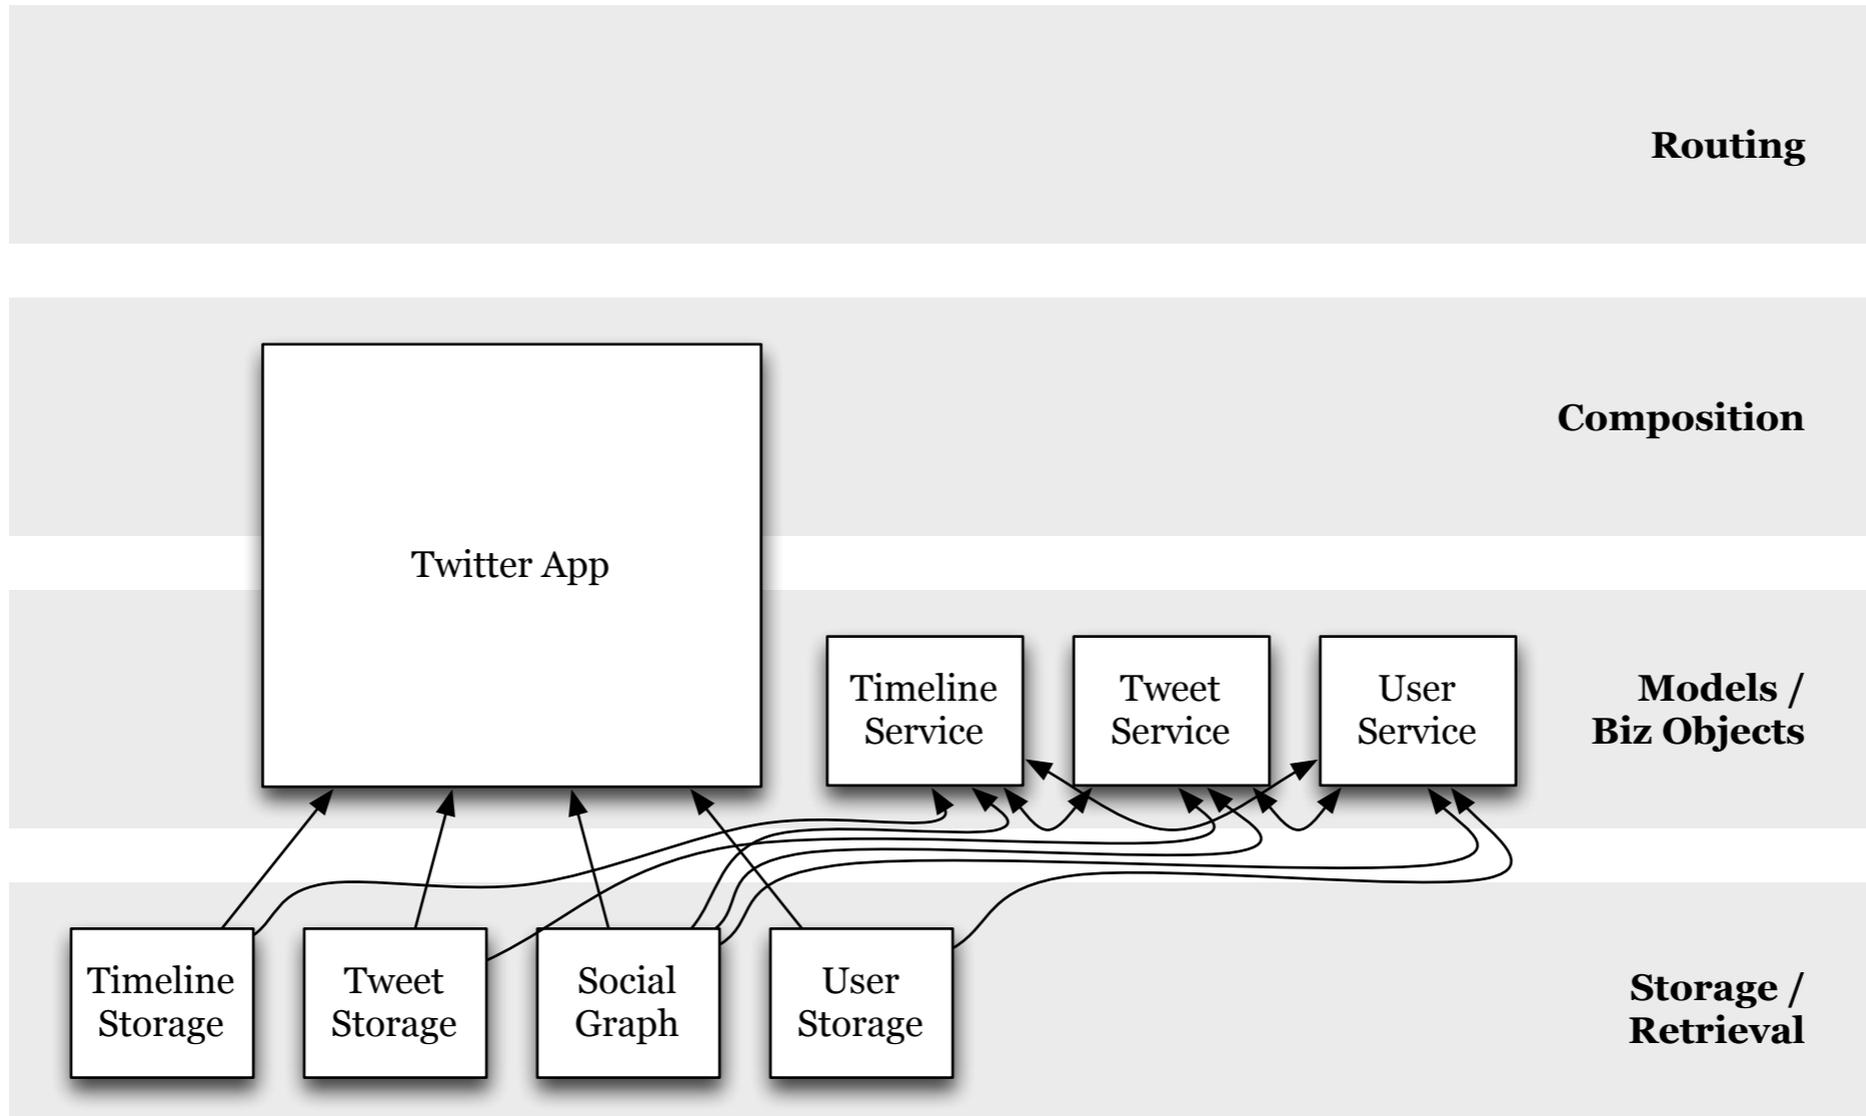
**User
Storage**

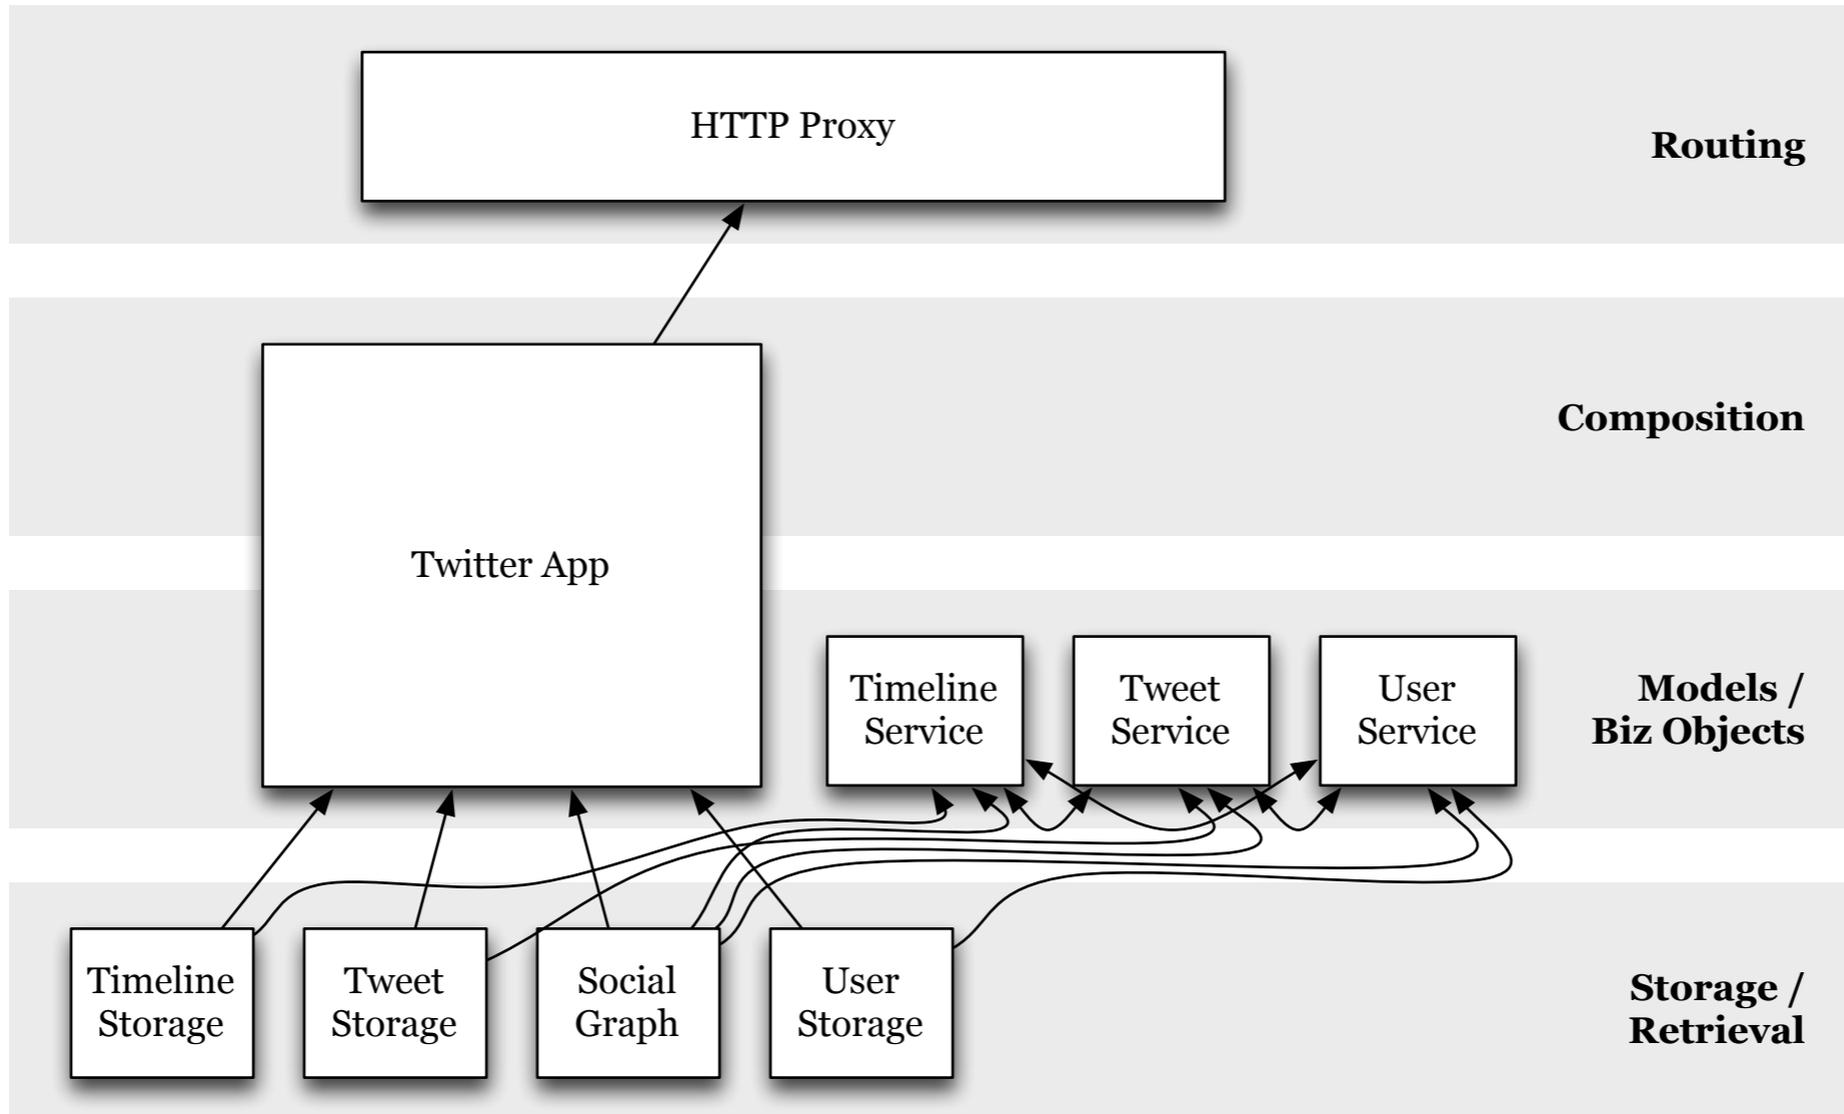
**Storage /
Retrieval**

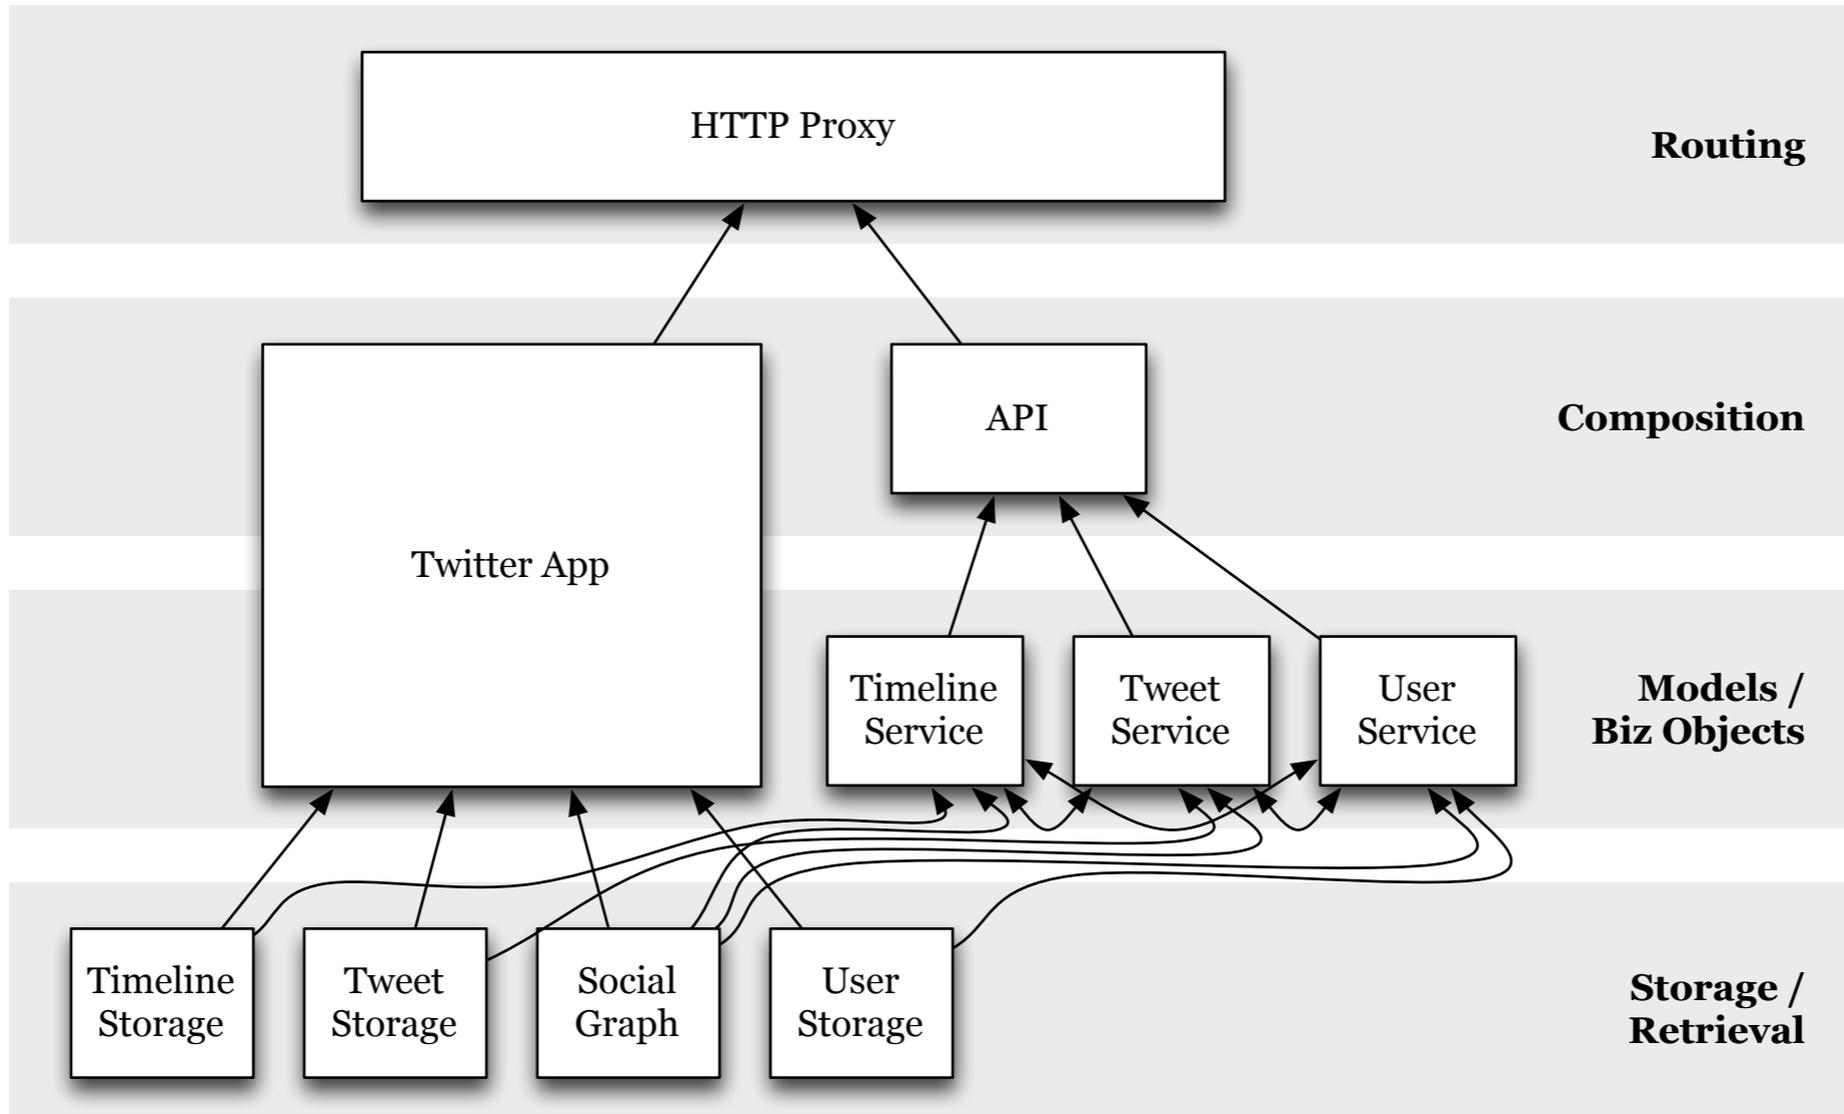


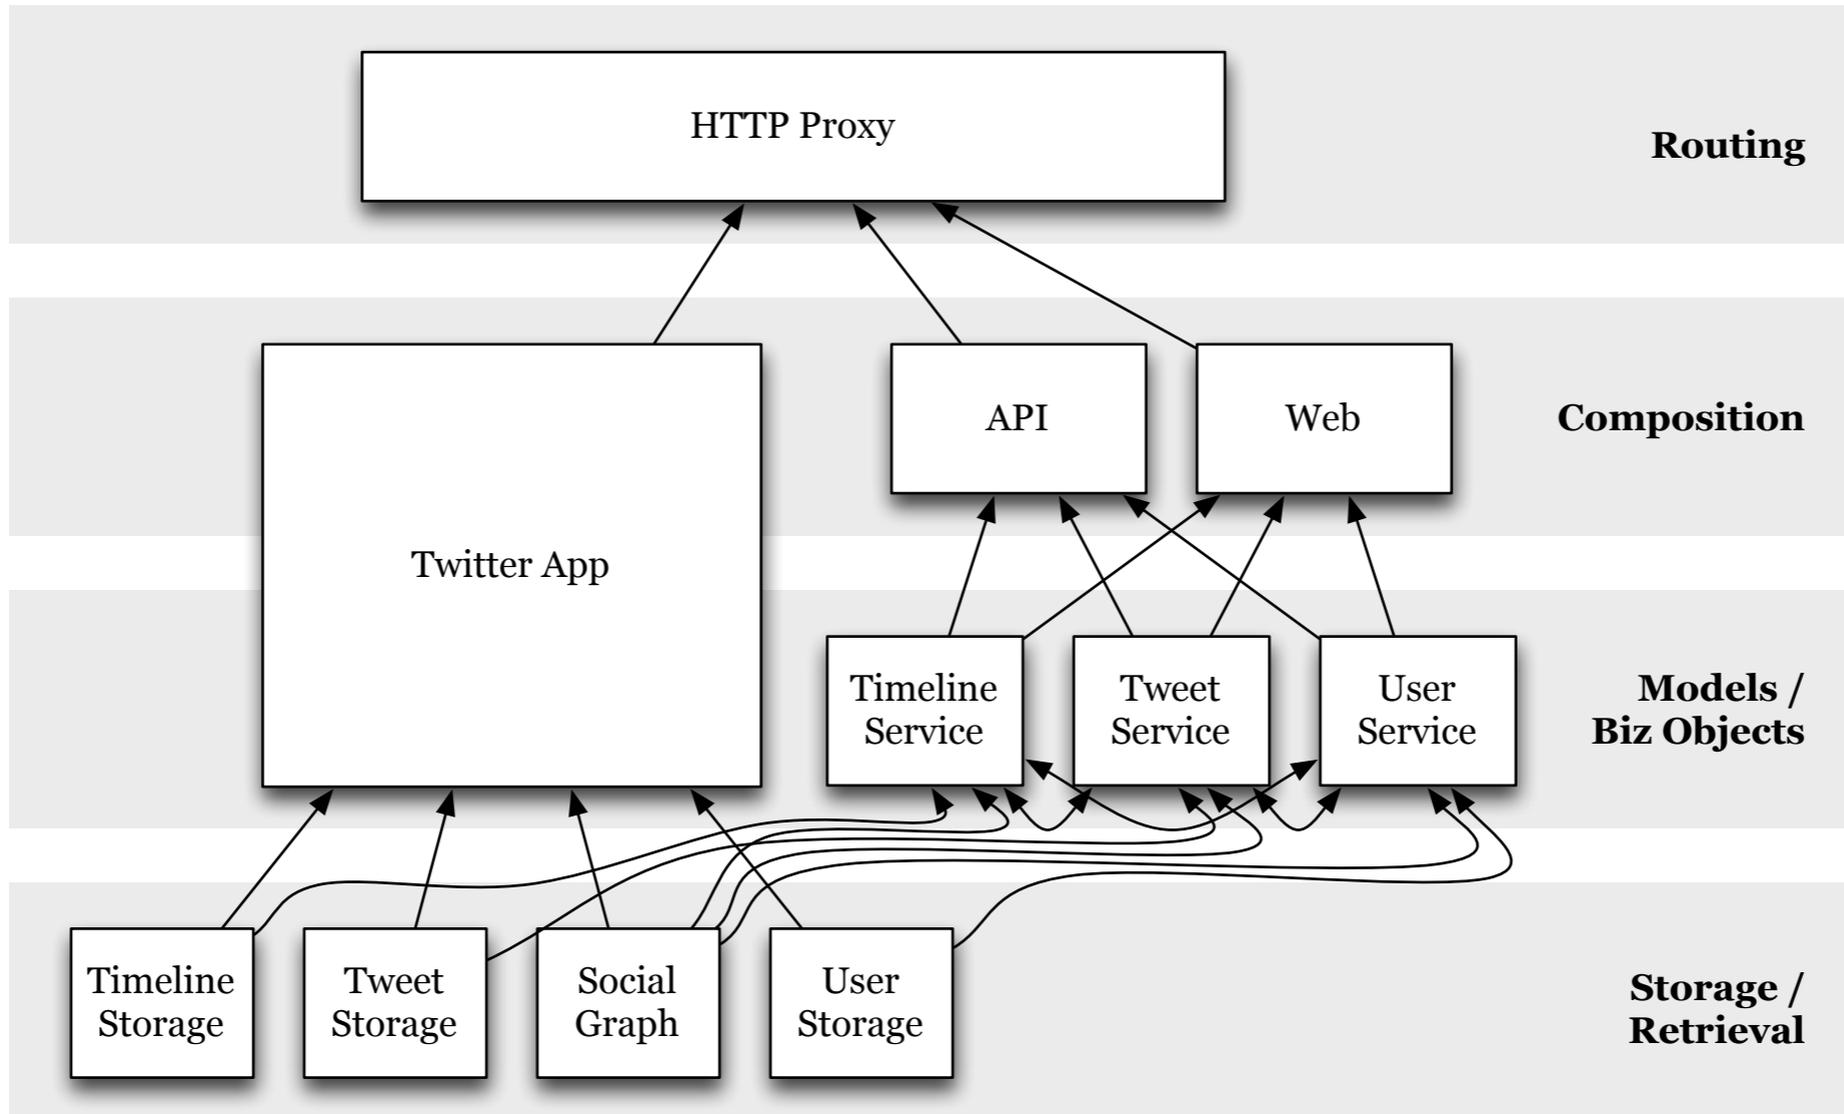


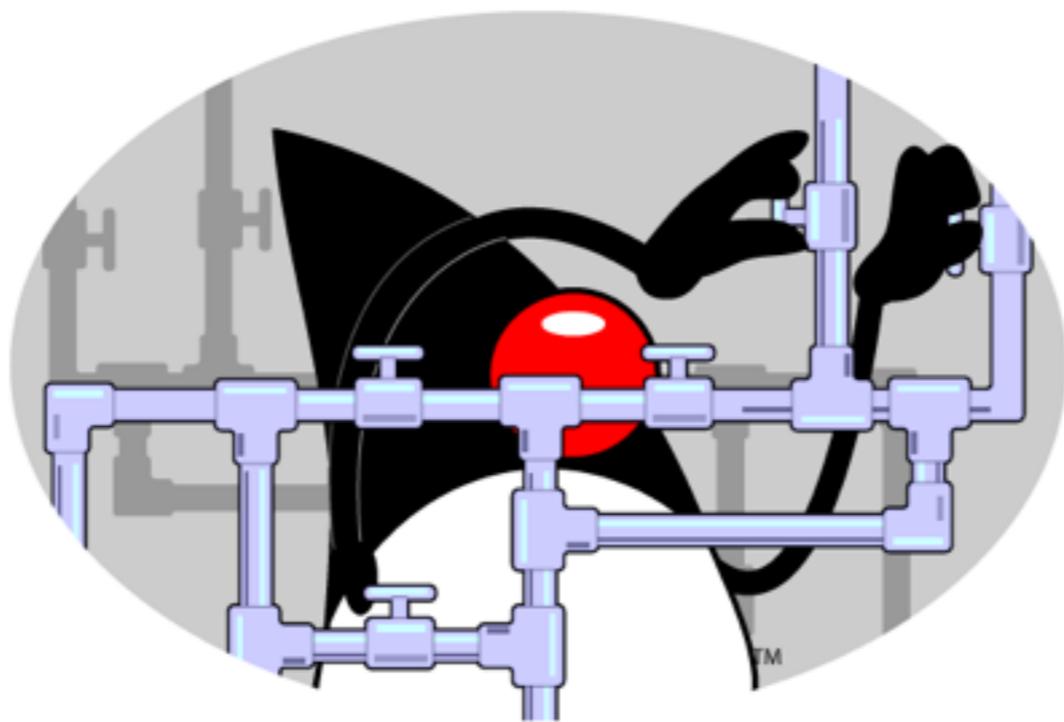












**SWITCHING TO
DOESN'T IMPLY THAT
IS A MISTAKE**





Following



In the final three minutes of the Super Bowl tonight, there were an average of 10,000 Tweets per second.

50+
RETWEETS

50+
FAVORITES



7:43 PM - 5 Feb 12 via Twitter for iPhone · Embed this Tweet

← Reply ↻ Retweet ★ Favorite

Reply to @twitter



Following

Madonna's performance during the Super Bowl's halftime show saw an average of 8,000 Tweets per second for five minutes.

50+
RETWEETS

50+
FAVORITES



7:43 PM - 5 Feb 12 via Twitter for iPhone · Embed this Tweet

Reply Retweet Favorite

Reply to @twitter



Following

The highest Tweets per second **#SuperBowl** peak came at the end of the game: 12,233. 2nd highest was during Madonna's performance: 10,245.

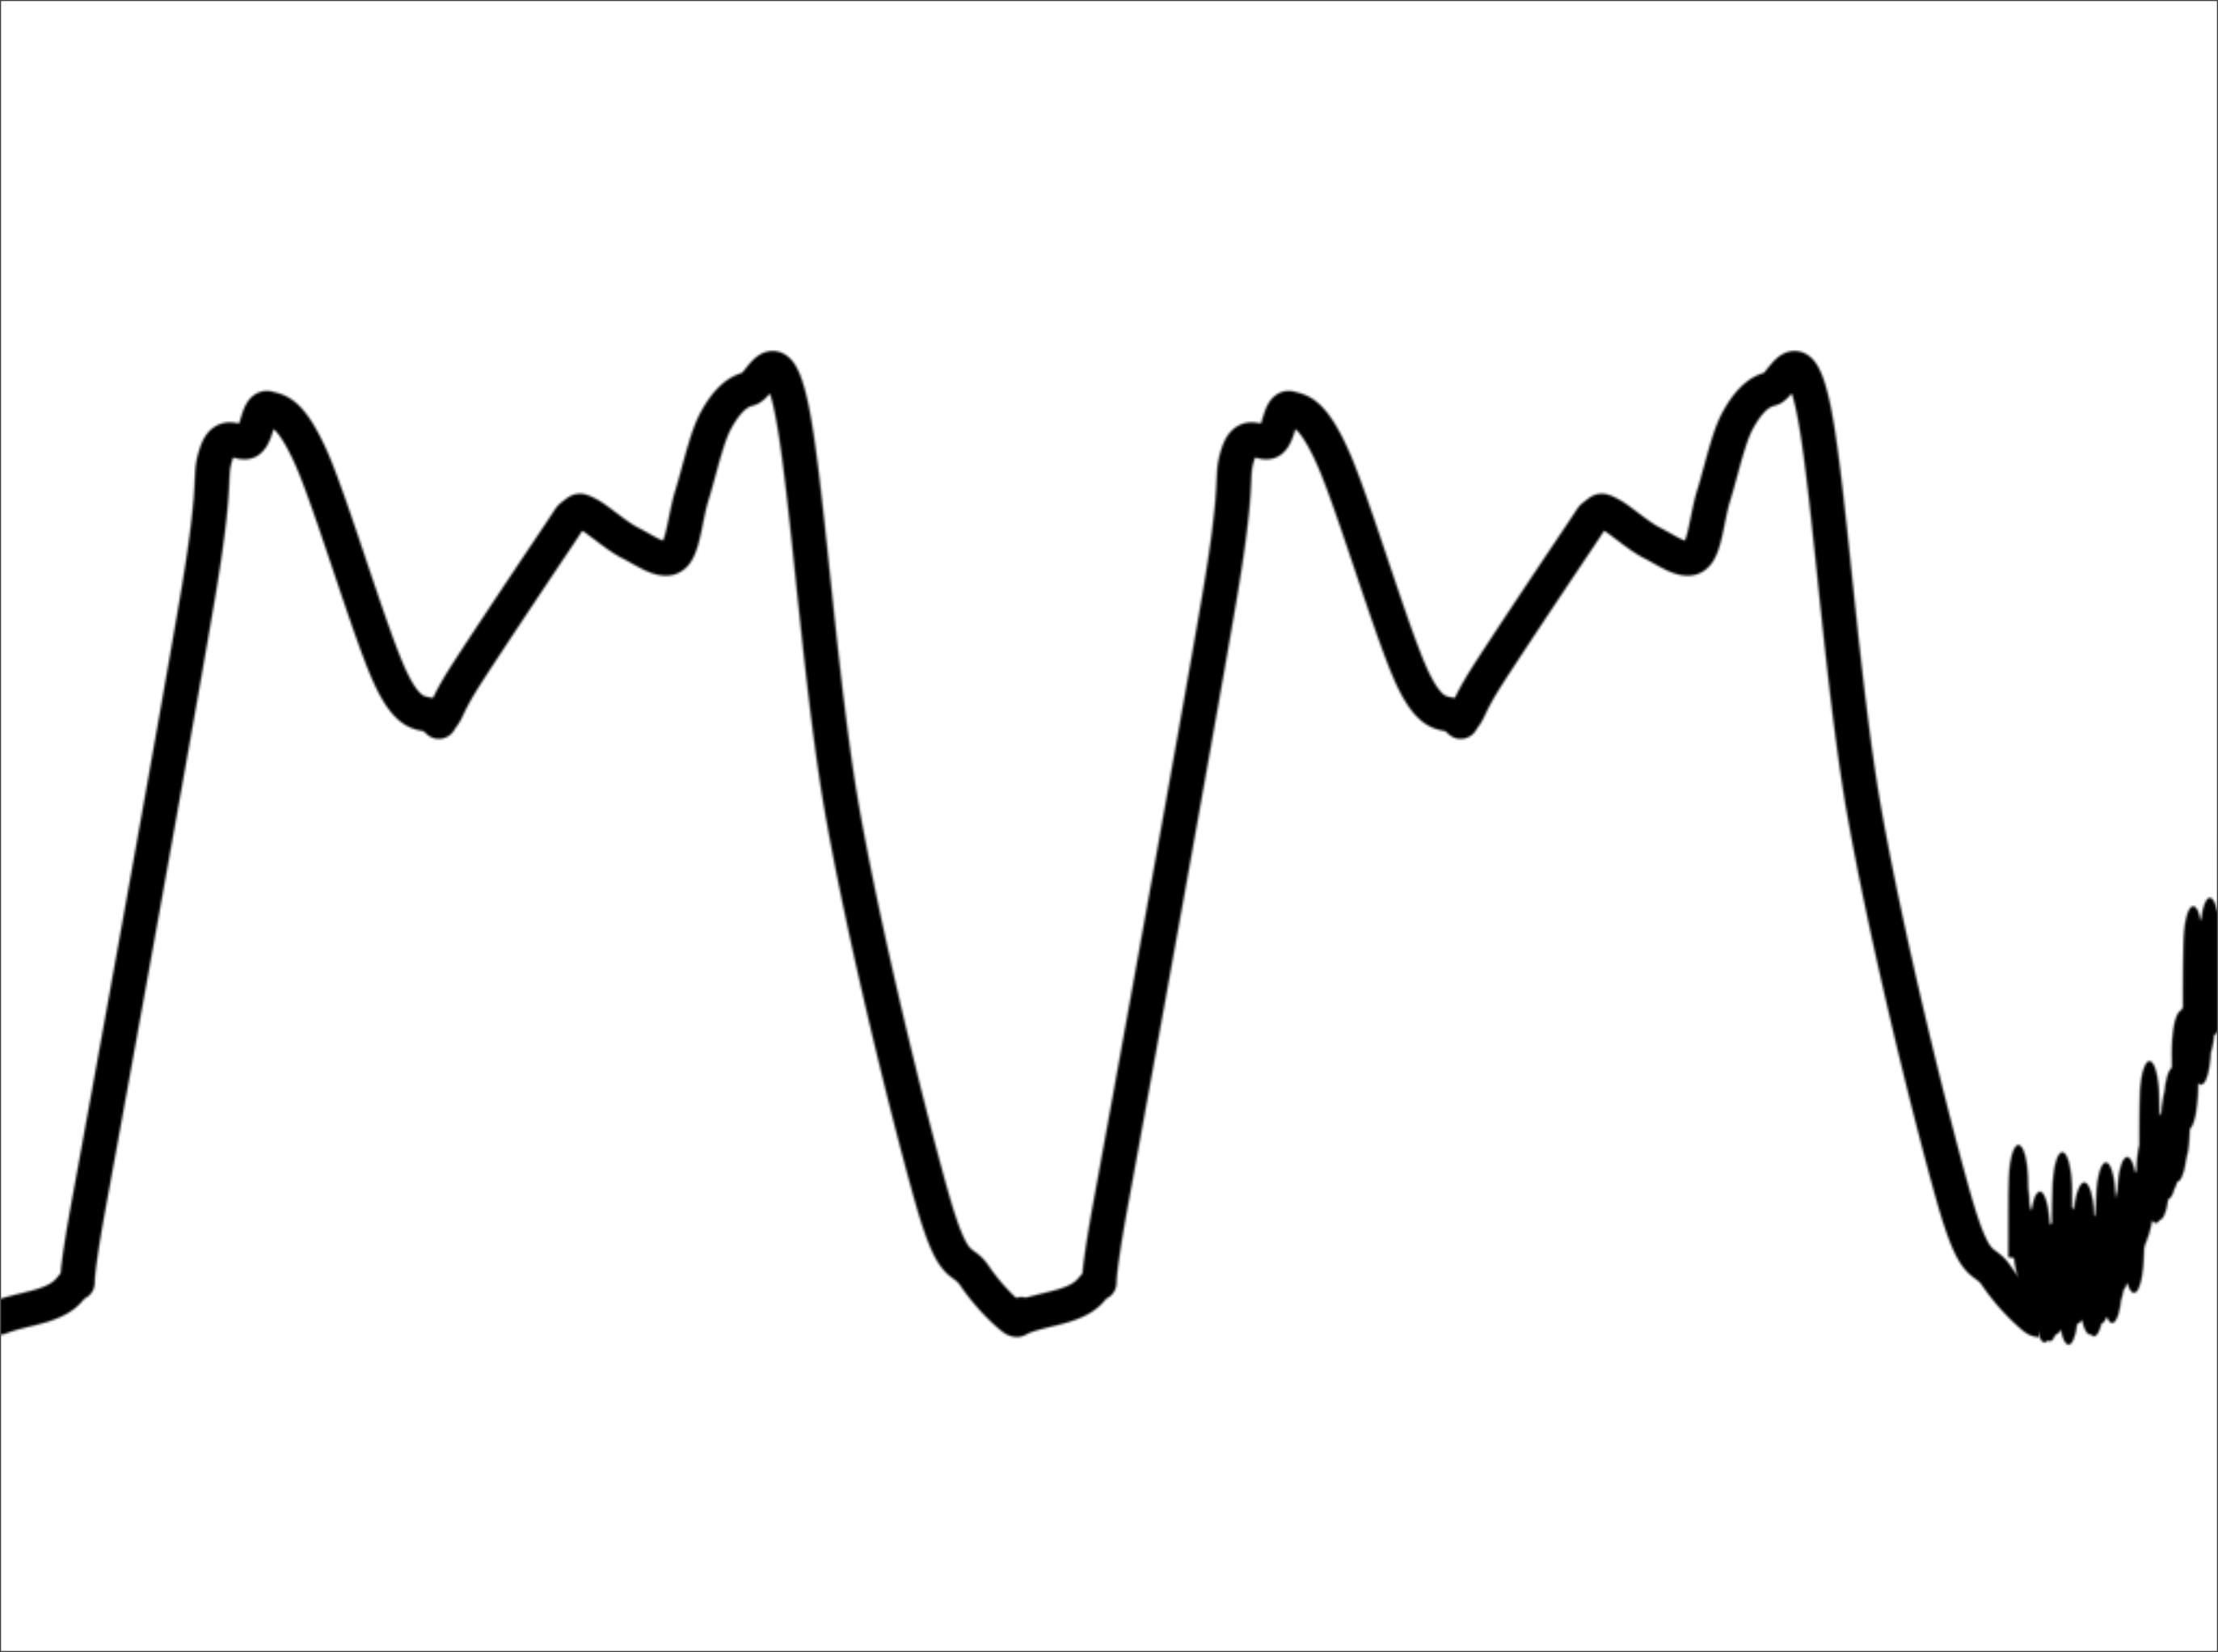
50+
RETWEETS

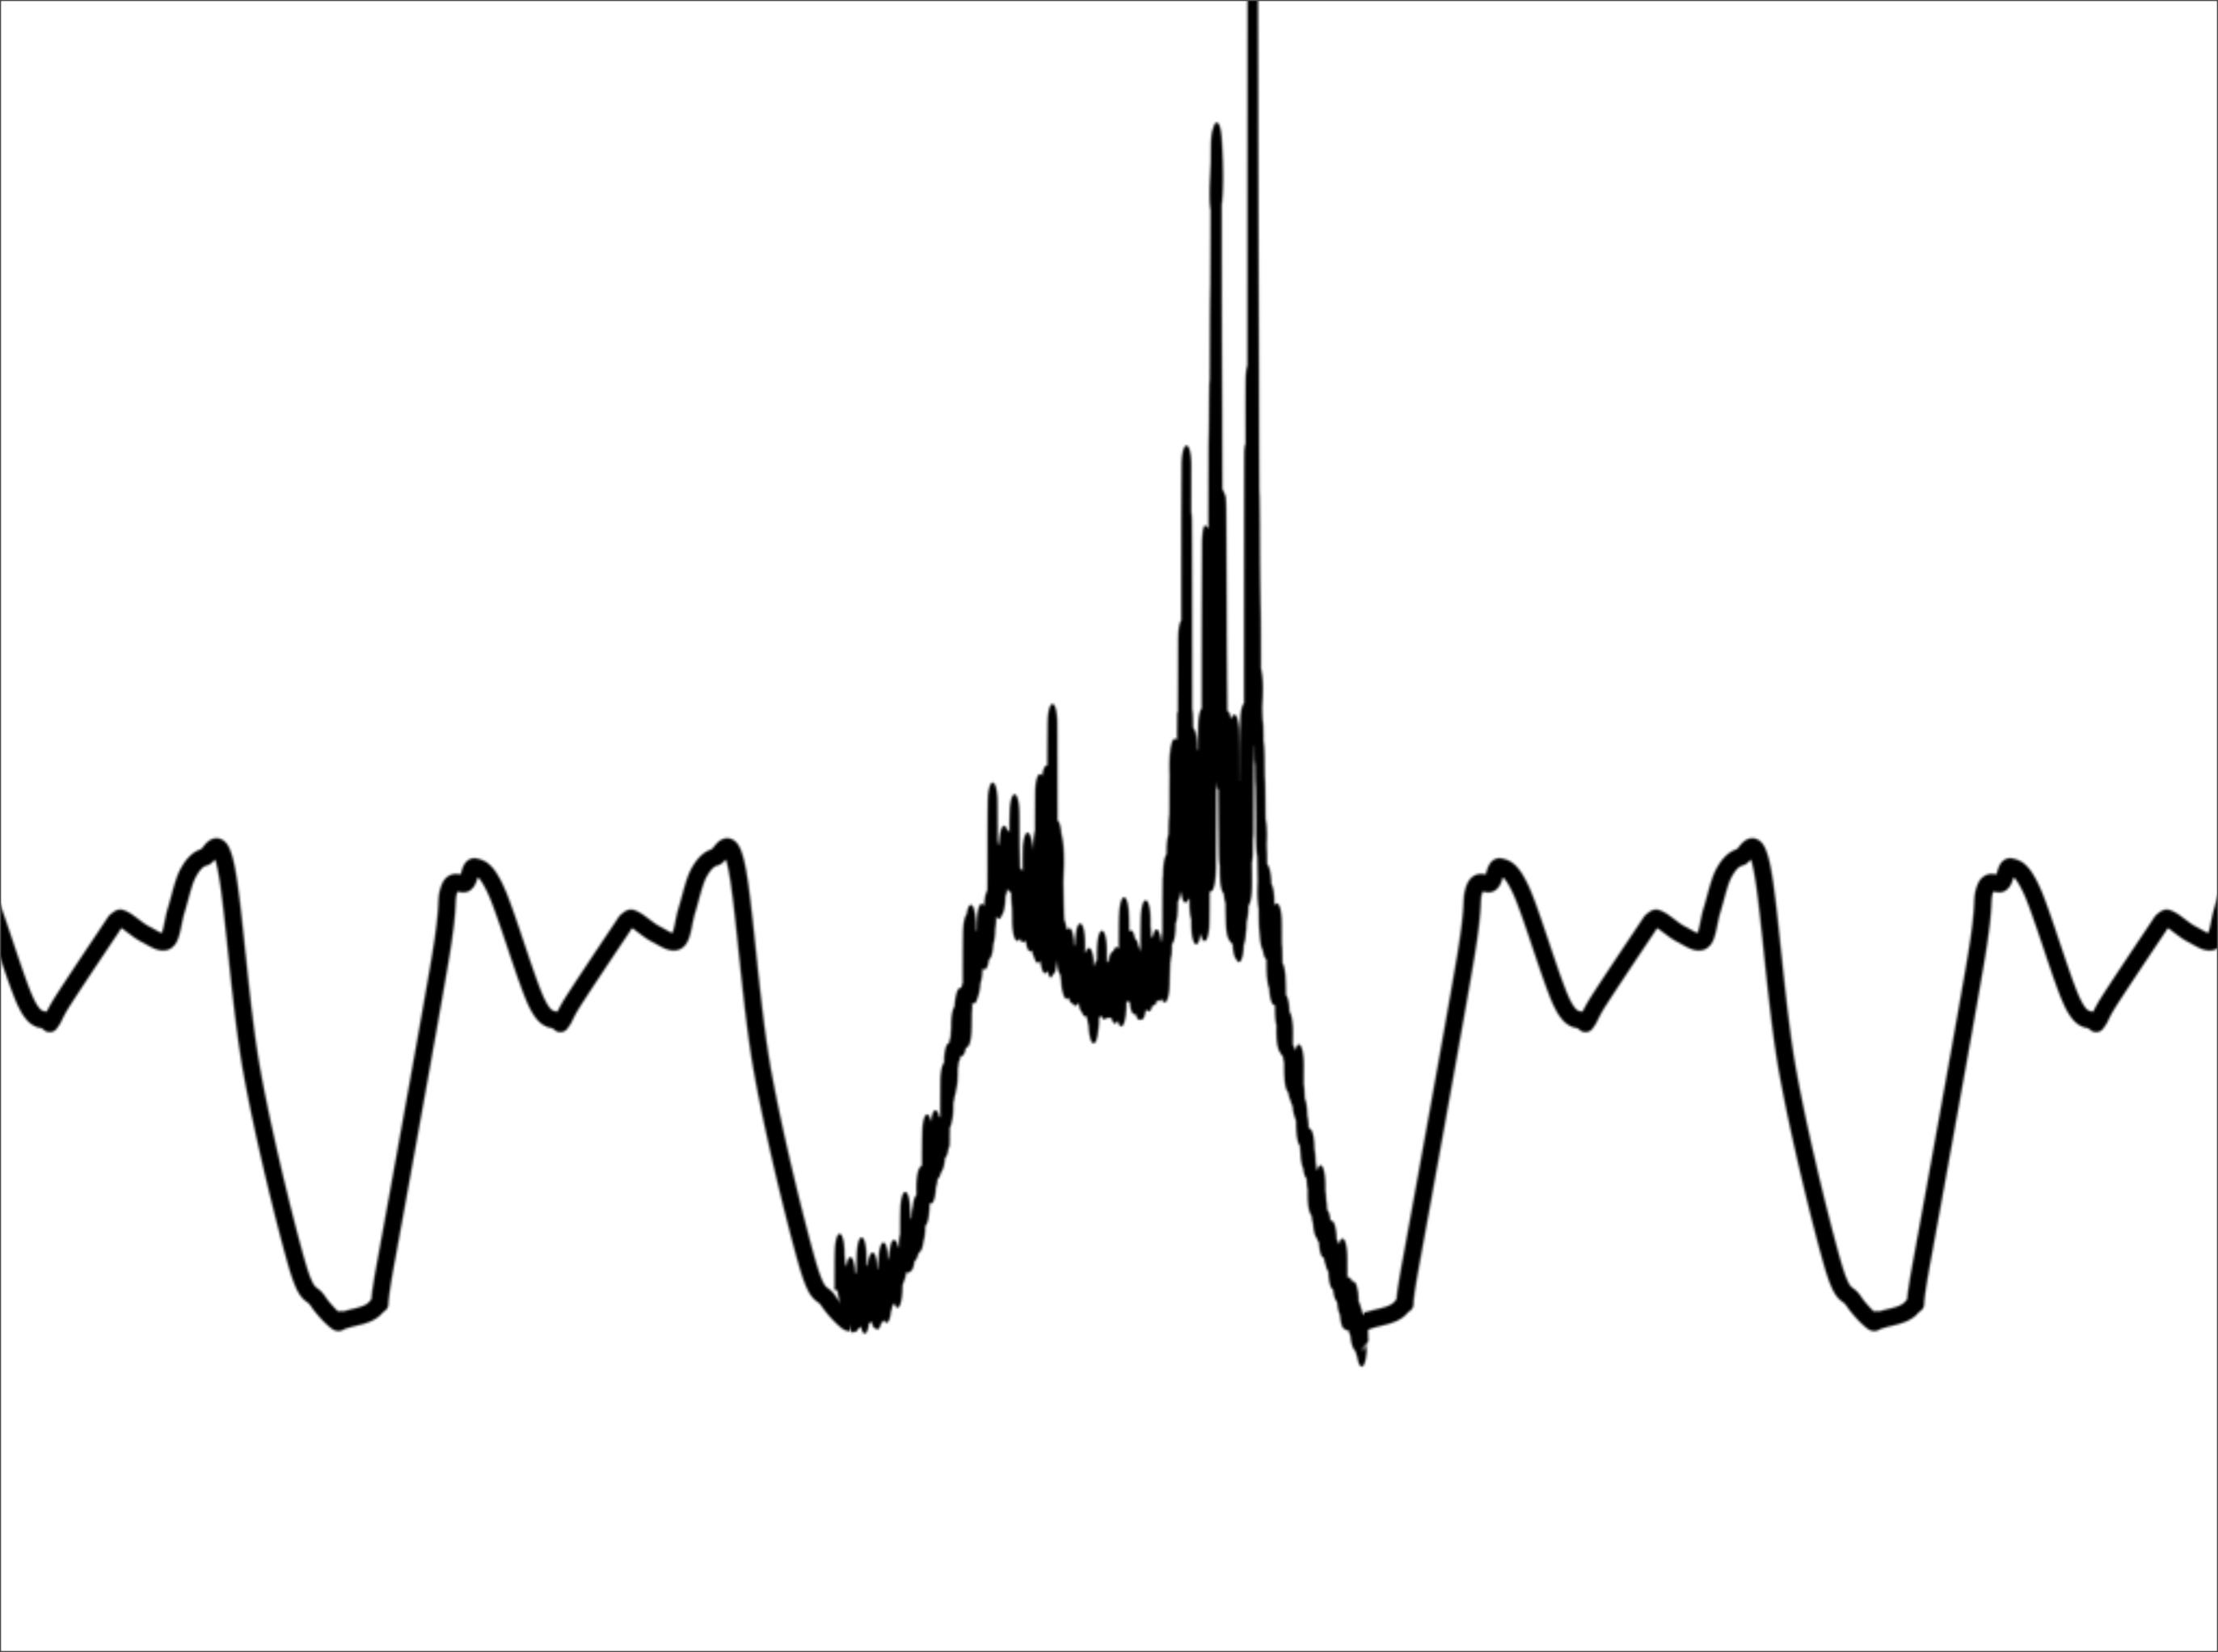
50+
FAVORITES



8:31 PM - 5 Feb 12 by RachaelRad via web · Embed this Tweet

Reply Retweet Favorite





 'S HIRING!
@JOINTHEFLOCK



FOLLOW ME
@RAFFI

