### Improving Running Components

Evan Weaver
Twitter, Inc.

QCon London, 2009

#### Many tools:

Rails Scala Java MySQL

#### Rails front-end:

rendering cache composition db querying

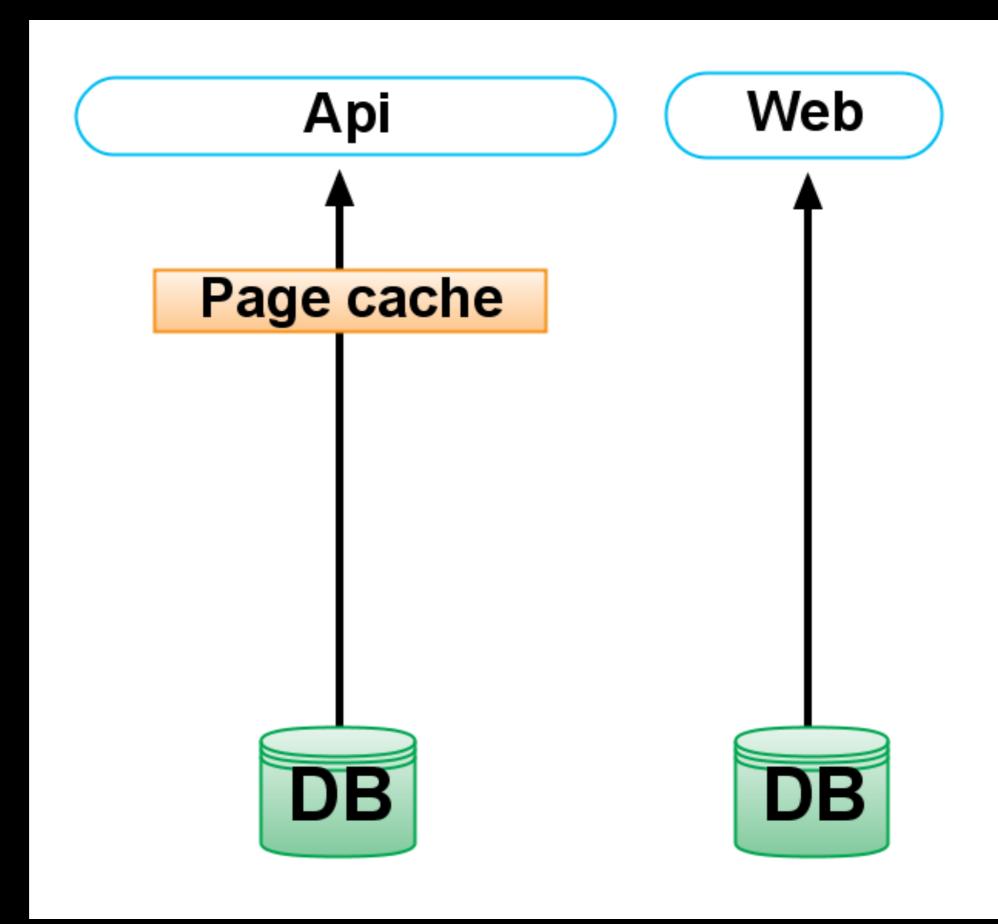
#### Middleware:

Memcached
Varnish (cache)
Kestrel (MQ)
comet server

#### Milestone 1: Cache policy

Optimization plan:

stop working
 share the work
 work faster



Old

## Everything runs from memory in Web 2.0.

### First policy change: vector cache

Stores arrays of tweet pkeys
Write-through
99% hit rate

### Second policy change: row cache

Store records from the db (Tweets and users)
Write-through
95% hit rate

## Third policy change: fragment cache

Stores rendered version of tweets for the API Read-through 95% hit rate

# Fourth policy change: giving the page cache its own cache pool

Generational keys Low hit rate (40%)

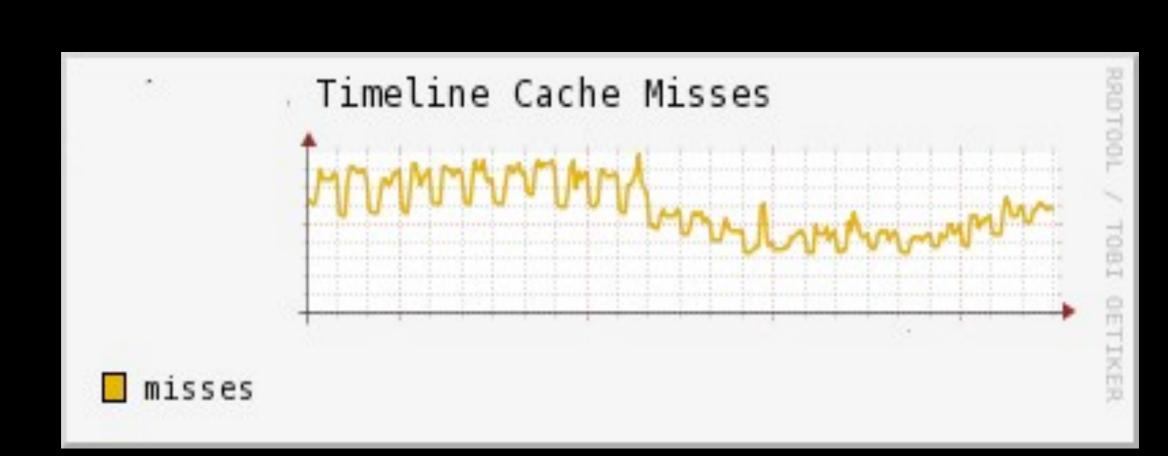
#### Visibility was lacking.

# Peep tool Dumps a live memcached heap

### Cache only was living five hours

mysql> select round(round(log10(3576669 - last\_read\_time) \* 5, 0) / 5, 1) as log, round(avg(3576669 - last\_read\_time), -2) as freshness, count(\*), rpad('', count(\*) / 2000, '\*') as bar from entries group by log order by log desc;

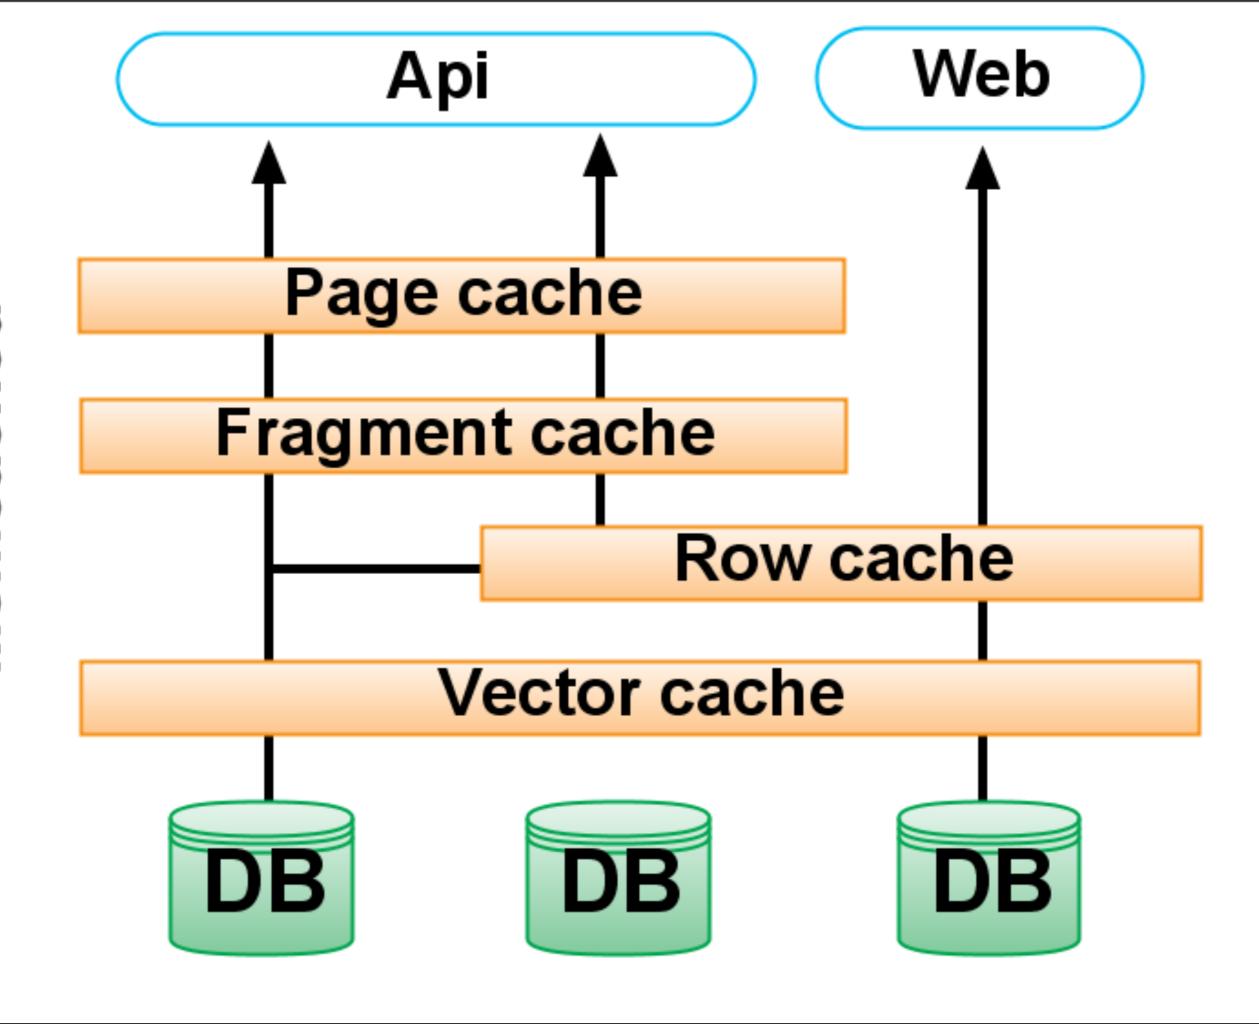
log   ++	freshness		bar
NULL	0	13400	****
6.6	3328300	940	
6.2	1623200	1	
5.2	126200	1	
5.0	81100	343	
4.8	64800	3200	**
4.6	34800	18064	*****
4.4	24200	96739	************************
4.2	15700	212865	**************************
4.0	10200	224703	***********************************
3.8	6500	158067	************
3.6	4100	108034	*************************
3.4	2600	82000	******************
3.2	1600	65637	**************
3.0	1000	49267	***********
2.8	600	34398	*******
2.6	400	24322	******
2.4	300	19865	*****
2.2	200	14810	*****
2.0	100	10108	****
1.8	100	8002	***
1.6	0	6479	***
1.4	0	4014	**
1.2	0	2297	*
1.0	0	1733	*
0.8	0	649	
0.6	0	710	
0.4	0	672	
0.0	0	319	
++	+	+	



### What does a timeline miss mean?

Container union

/home rebuild reads through your followings' profiles



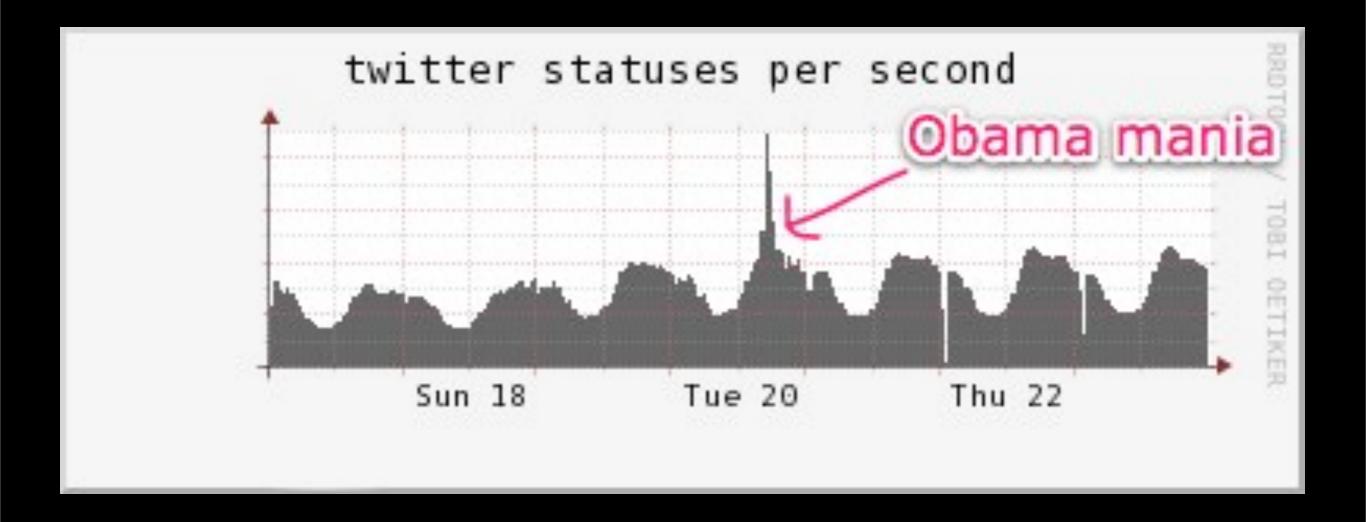
### Milestone 2: Message queue

A component with problems

#### Purpose in a web app:

Move operations out of the synchronous request cycle

Amortize load over time



### Inauguration, 2009

#### Simplest MQ ever:

Gives up constraints for scalability No strict ordering of jobs No shared state among servers Just like memcached Uses memcached protocol

## First version was written in Ruby

Ruby is "optimization-resistant"

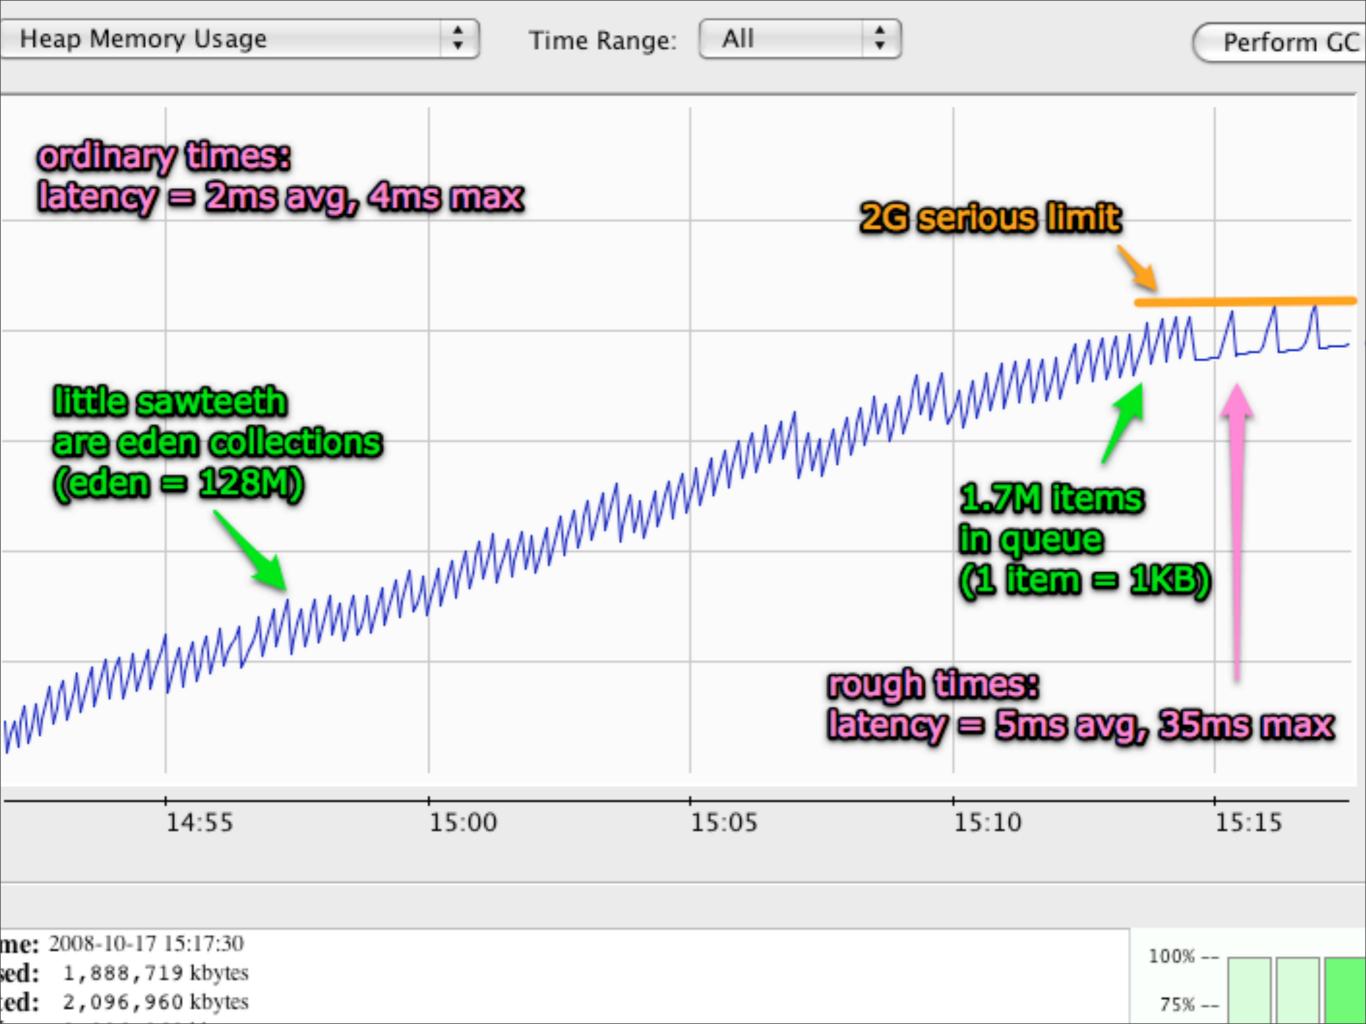
Mainly due to the GC

# If the consumers could not keep pace, the MQ would fill up and crash

Ported it to Scala for this reason

### Good tooling for the Java GC:

J Console Yourkit



### Poor tooling for the Ruby GC:

Railsbench w/patches
BleakHouse w/patches
Valgrind/Memcheck
MBARI 1.8.6 patches

#### Our Railsbench GC tunings

#### 35% speed increase

```
RUBY_HEAP_MIN_SLOTS=500000
RUBY_HEAP_SLOTS_INCREMENT=250000
RUBY_HEAP_SLOTS_GROWTH_FACTOR=I
RUBY_GC_MALLOC_LIMIT=50000000
RUBY_HEAP_FREE_MIN=4096
```

#### Situational decision:

Scala is a flexible language (But libraries a bit lacking)
We have experienced JVM engineers

### Big rewrites fail...?

Small rewrite:
No new features added
Well-defined interface
Already went over the wire

#### Deployed to I MQ host

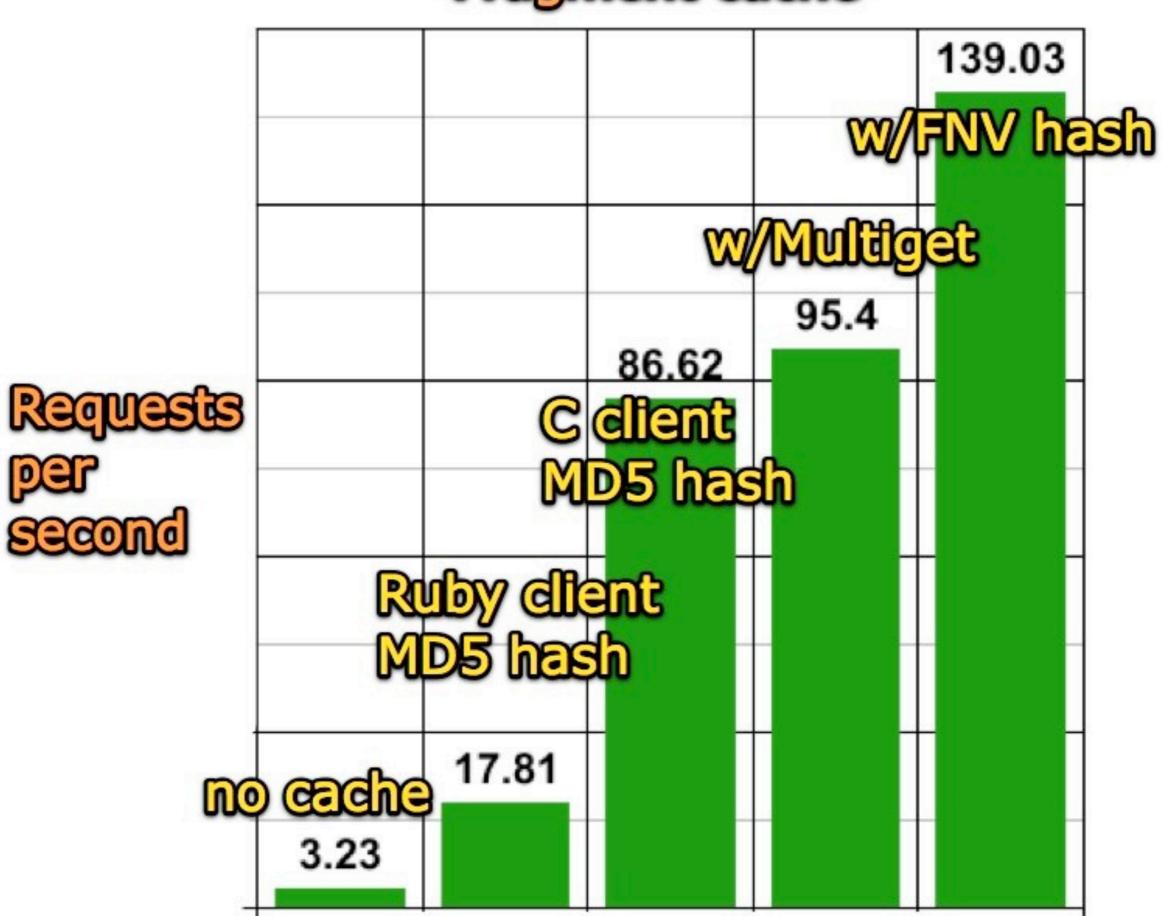
Fixed regressions

Eventually deployed to all hosts

### Milestone 3: the memcached client

Optimizing a critical path

#### Fragment cache



### Switched to libmemcached, a new C Memcached client

We are now the biggest user and biggest 3rd-party contributor

## Uses a SWIG Ruby binding I started a year or so ago

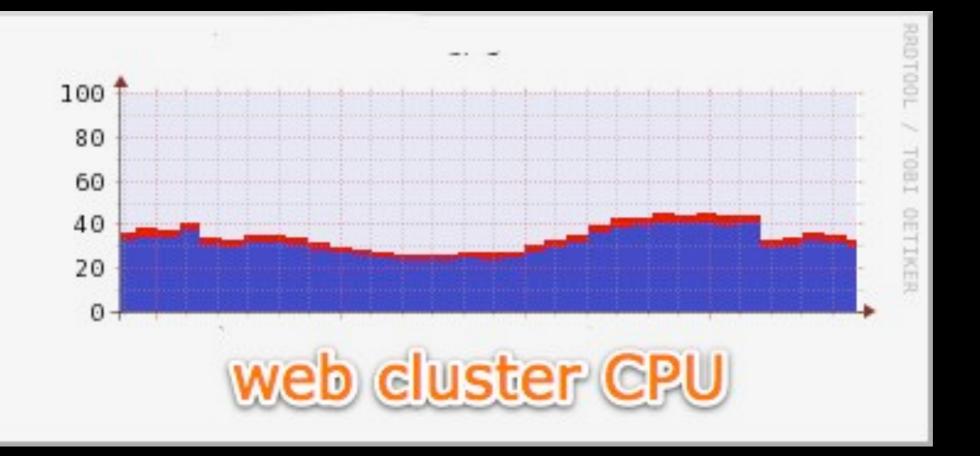
Compatibility among memcached clients is critical

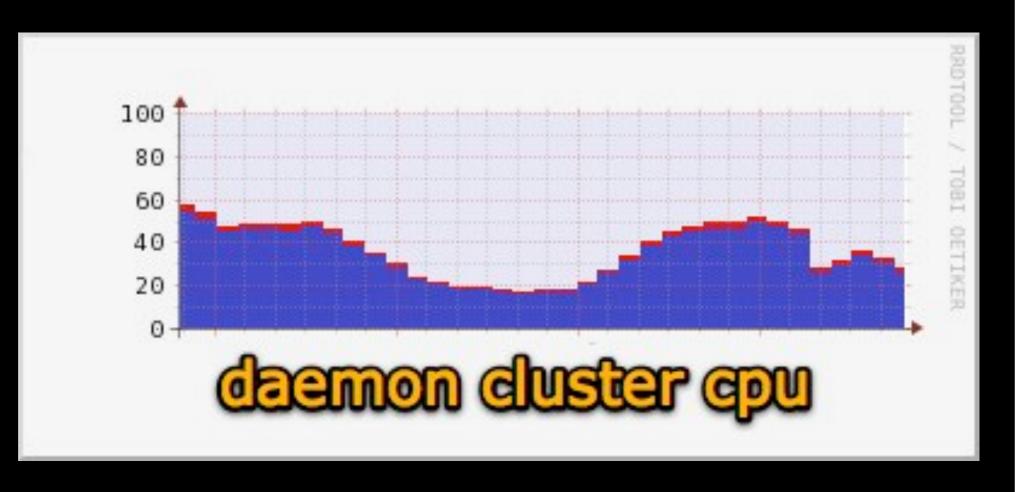
### Twitter is big, and runs hot

## Flushing the cache would be catastrophic

## Spent endless time on backwards compatibility

### A/B tested the new client over 3 months





#### MQ also benefitted

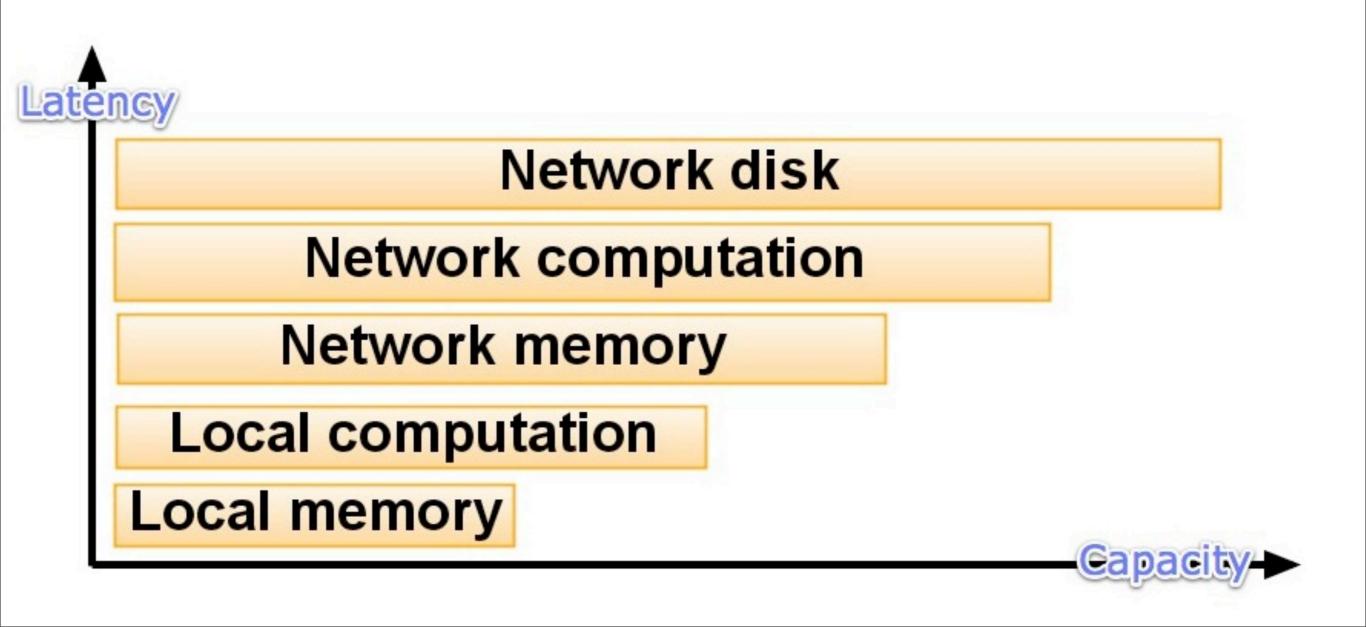
Memcached can be a generic lightweight service protocol

We also use Thrift and HTTP internally

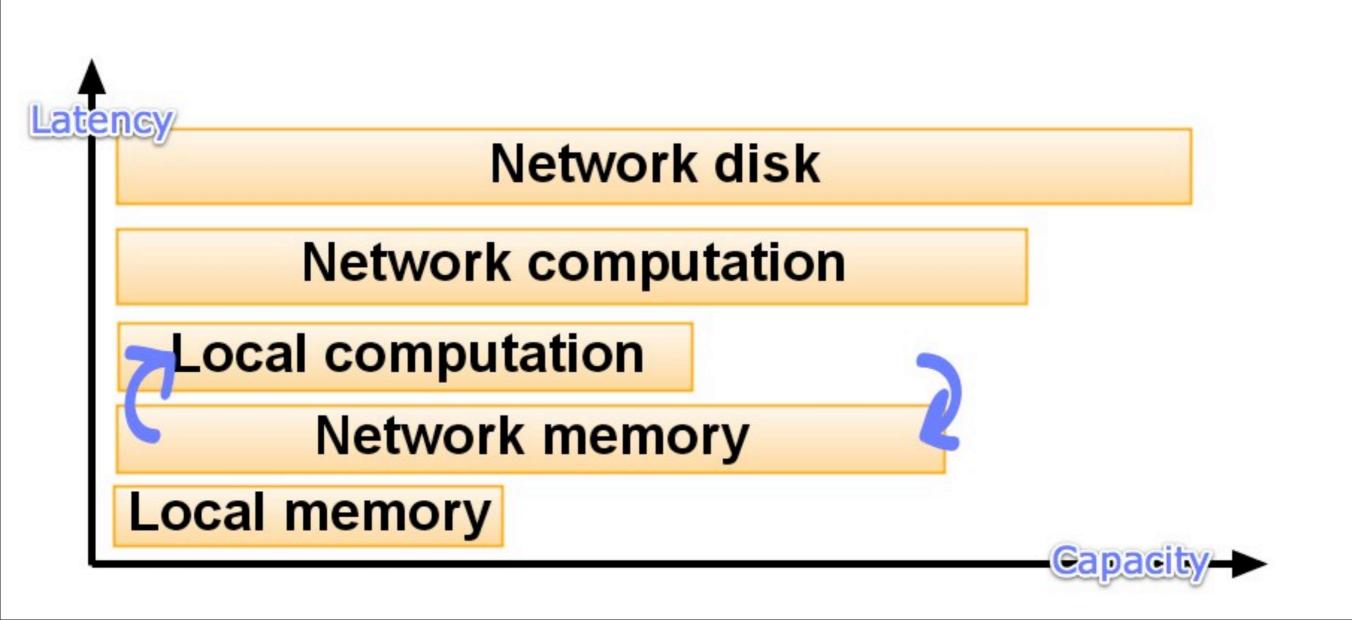
# So many RPCs! Sometimes 100s of Memcached round trips per request.

"As a memory device gets larger, it tends to get slower."

### Performance hierarchy is supposed to look like:



### At web scale, it looks more like:



### End

#### Links:

#### C tools:

- Peep <a href="http://github.com/fauna/peep/">http://github.com/fauna/peep/</a>
- Libmemcached <a href="http://tangent.org/552/libmemcached.html">http://tangent.org/552/libmemcached.html</a>
- Valgrind <a href="http://valgrind.org/">http://valgrind.org/</a>

#### JVM tools:

- Kestrel <a href="http://github.com/robey/kestrel/">http://github.com/robey/kestrel/</a>
- Smile <a href="http://github.com/robey/smile/">http://github.com/robey/smile/</a>
- Jconsole <a href="http://openjdk.java.net/tools/svc/jconsole/">http://openjdk.java.net/tools/svc/jconsole/</a>
- Yourkit <a href="http://www.yourkit.com/">http://www.yourkit.com/</a>

#### Ruby tools:

- BleakHouse <a href="http://github.com/fauna/bleak\_house/">http://github.com/fauna/bleak\_house/</a>
- Railsbench Ruby patches <a href="http://github.com/skaes/railsbench/">http://github.com/skaes/railsbench/</a>
- MBARI Ruby patches <a href="http://github.com/brentr/matzruby/tree/vI\_8\_6\_287-mbari">http://github.com/brentr/matzruby/tree/vI\_8\_6\_287-mbari</a>

#### **General:**

- Danga stack <a href="http://www.danga.com/words/2005\_oscon/oscon-2005.pdf">http://www.danga.com/words/2005\_oscon/oscon-2005.pdf</a>
- Seymour Cray quote <a href="http://books.google.com/books?client=safari&id=qM4Yzf8K9hwC&dq=rapid">http://books.google.com/books?client=safari&id=qM4Yzf8K9hwC&dq=rapid</a> +development&q=cray&pgis=I
  - Last.fm downtime <a href="http://blog.last.fm/2008/04/18/possible-lastfm-downtime">http://blog.last.fm/2008/04/18/possible-lastfm-downtime</a>

#### twitter.com/evan

blog.evanweaver.com

cloudbur.st