

Fix the Website

A DevOps Success* Story

RANT



Help! The Website is Slow!

whoami

- Developer who ops
- former engineer (EE)
- ac·qui·hire
- LA.PM.org #perl
- Now at ZipRecruiter

@spazm

lowlevelmanager.com



#monitoring 

Dev who ops. Following #monitoringlove

acquired by DM. Worked on a great project for a year. After cancelled, I roamed the company fixing problems. Tasked with improving "the website"

Thanks ZR for sponsoring my talk.


About the Website

- ComScore top 30
- Originated the long tail how-to site
- dependent on google traffic
- monetized by google AdWords
- 3 Million pages, ~20 types

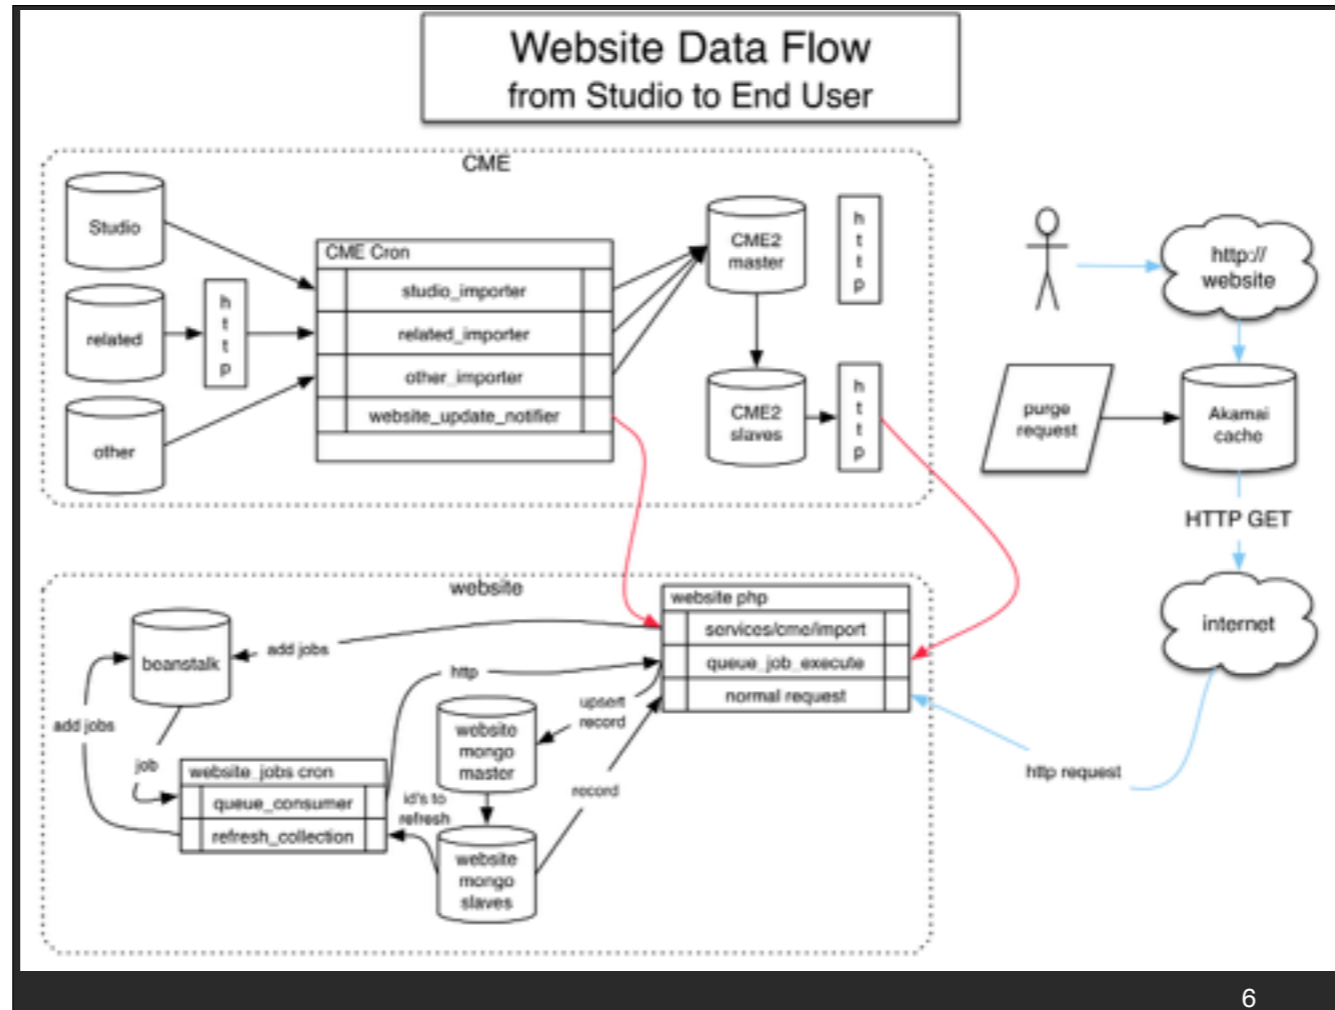


- originated its category. (content farm)
- May have been affected by google panda?
- In decline, but still a cash cow.
- long tail, some pages get 1 or fewer imp per day.
- about me: Dev who ops. Recently tasked with improving this site. Just finished migrating some transient data from the DB to a separate datastore.

A simple LAMP stack?



A nice simple LAMP stack, right? Nope.



Two levels of upstream

So many caches (foreshadowing)

Literally no one remembered how this worked.

complicated. Several generations of workers.

This metric is bad!

"FIX IT!" -- the PM

Avg(US) and Avg(DT)



ZipRecruiter

US and DT graphs.

Sources are summary graphs. Averages of averages!

Not a graph you want to go "up and to the right"

Notes indicating the PM's opinions on causality.

This is almost data-driven. More of a drive-by.

Hey PM, what does this graph mean?



HALP! I CAN HAZ METRIX! FIX IT!

Are all PM's hysterical?

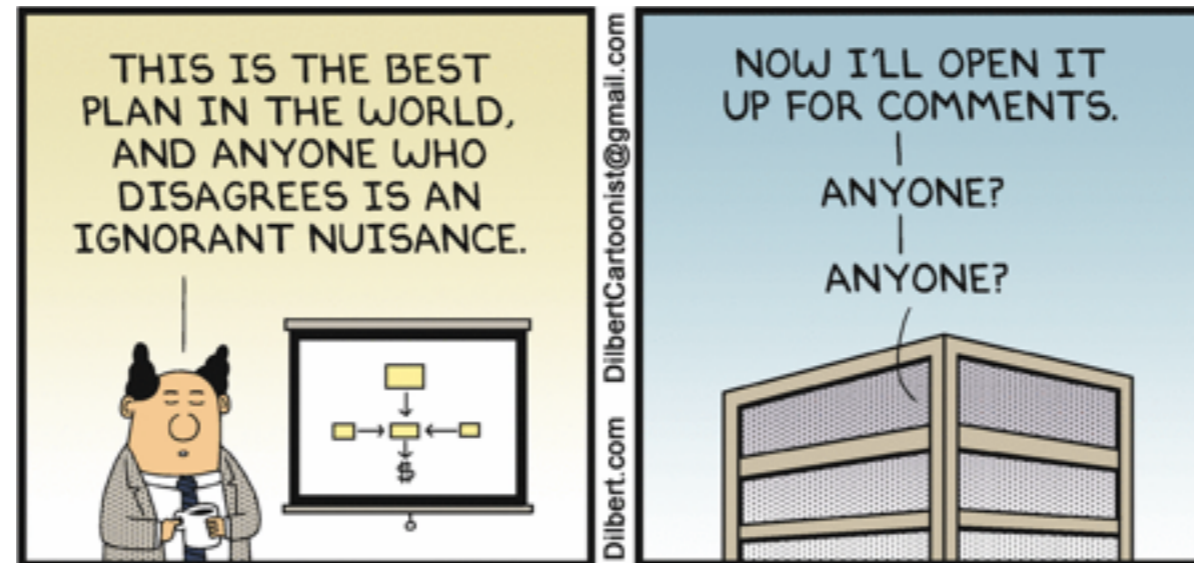
What does the metric mean?

10

Ask the PM his boss, VP-Eng

DT is Deficit Time

Deficit time must be the database!
Replace the DB!



The Database is to blame.

DT must be the database

We hate Mongo.

We know it's Mongo


I hate mongo. (Know your biases)

MongoDB is Web Scale



- I so want to kill Mongo, but it's not the culprit here.
- best use case: single key-value lookup for json doc, all in ram. It's quick.

What creates the metric?



13

splunk summary table => found that DT = US - PP

apache log config => found PP is php running time, and US is the standard apache timing metric. US vs US-FB.



Correlations?

14

Our logs are in splunk for analysis. Let's look! First, had to learn splunk query language. Turns out no one on the team had done deep analysis!
Not correlated to time, date, pagetype.

Non-Correlations

- time of day
- day of week
- week over week
- page subtype
- large data files don't have PP



don't have graphs to show, lost them.


241 results from 12:42:00 PM to 4:42:03 PM on Wednesday, April 9, 2014

Small Size (<15kB)

Options

50 per page

Formatting options



TITLE: all traffic, <15KB

SUMMARY:


* max(Delta) is -12ms steady-state.

* US and PP are steady-state at -60ms with worst case around a second (sometimes 5)

- Log axis means you're in trouble
- max(US) and max(PP) are indistinguishable in log axis.

241 results from 12:41:00 PM to 4:41:57 PM on Wednesday, April 9, 2014

Medium size (15-25kB)



TITLE: all 200 traffic, 15KB < size <= 25KB

SUMMARY:

- * max(delta) jumps to 1-4seconds steady-state.
- * FP stays steady-state at ~60-80ms with normal max of 0.3-1 second, spikes to 10s.
- * US avg jumps to by 150ms to 210-220ms, US max is dominates delta.

- Why does avg(US) jump by 150ms?
- Max(delta) jumps to 1-4s!
- Big max is throwing off the avg. Need %tile!


240 results from 12:41 PM to 4:41 PM on Wednesday, April 9, 2014

Large size (>25kB)

Options

50 per page


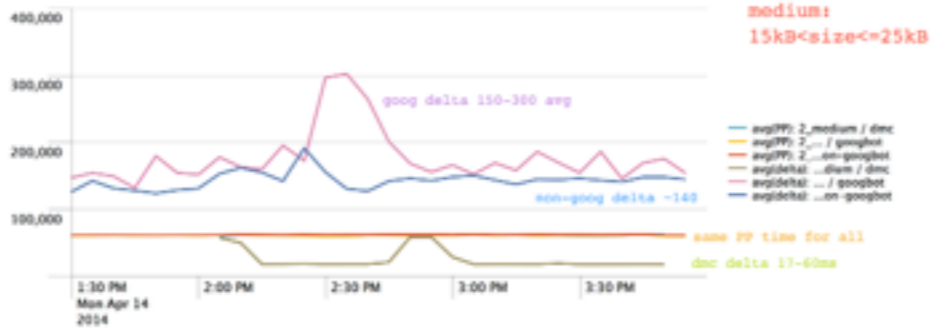
Formatting options



TITLE: all 200 traffic, 25KB < size
SUMMARY:
* max(delta) is about 1/2 a second, spikes to 3.
* PP stays steady-state at ~60-80ms with max < 0.5 second.
* US avg jumps again to 200-300ms, US max is dominates delta.
Not shown: avg size 36kB, max - kB

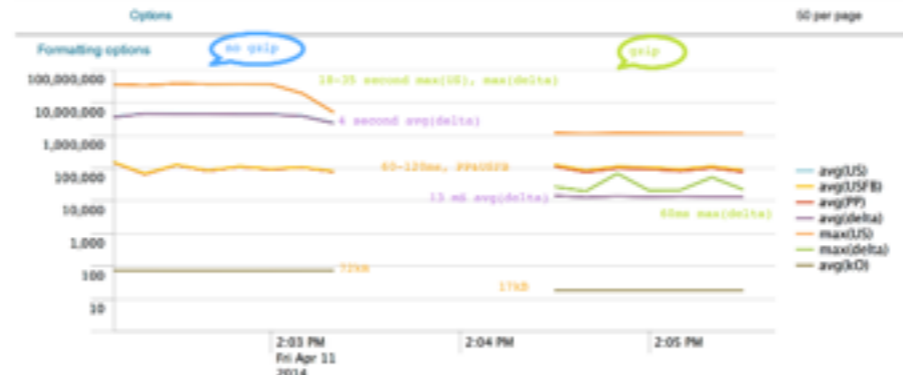
- Large is bad, but not as bad as medium?
- Sampling issue -- slow clients don't get the big pages

Compare Spiders



Can we Replicate?

21 results from 2:00 PM to 2:10 PM on Friday, April 11, 2014



Summary:
 * PP and USFB same for both runs.
 * US time goes way up when overloaded
 * gzip reduces overall load on server
 * max delta (deficit time) is lower in gzip than avg delta in non-gzip case!!
 * max delta = max us for no-gzip, lower than max us for gzip.

```

bin/httpperf-0.9.1-ens-patch.bin --hog --server 10.7.88.45 --port 80 \
--server-name www.show.com \
--loggy,requests_httpperf.log \
--num-conn=20000 \
--rate=300 \
--timeout=20 --add-header "Cookie: dp=desktop"
httpperf-0.9.1-ens-patch.bin --hog --timeout=20 --client=0/1 --server=10.7.88.45 --server_name=www.show.com --port=80 --ssl= \
--rate=300 --send-buffer=4096 --recv-buffer=4096 --add-header "Cookie: dp=desktop" --num-conn=20000 --num-calls=1

bin/httpperf-0.9.1-ens-patch.bin --hog --server 10.7.88.45 --port 80 \
--server-name www.show.com \
--loggy,requests_httpperf.log \
--num-conn=20000 \
--rate=300 \
--timeout=20 --add-header "Cookie: dp=desktop" --add-header "Accept-Encoding: gzip"
httpperf-0.9.1-ens-patch.bin --hog --timeout=20 --client=0/1 --server=10.7.88.45 --server_name=www.show.com --port=80 --ssl= \
--rate=300 --send-buffer=4096 --recv-buffer=4096 --add-header "Accept-Encoding: gzip" --num-conn=20000 --num-calls=1
    
```

20,000 70k requests, 300 rps, no gzip

20,000 70k requests, 300 rps, with gzip



- Duplicated high US, high Delta in httpperf test
- Low PP and low USFB -- time to first byte
- Modified and compiled httpperf
- Not same issue, but similar.
- Network buffering and latency drove up US time!!!

Plan of Action

THE PAGES

ARE TOO DAMN BIG!

- 16.5K compressed / 70k uncompressed
- Shrink pages - remove old cruft.
- Move CSS to common files that will cache better

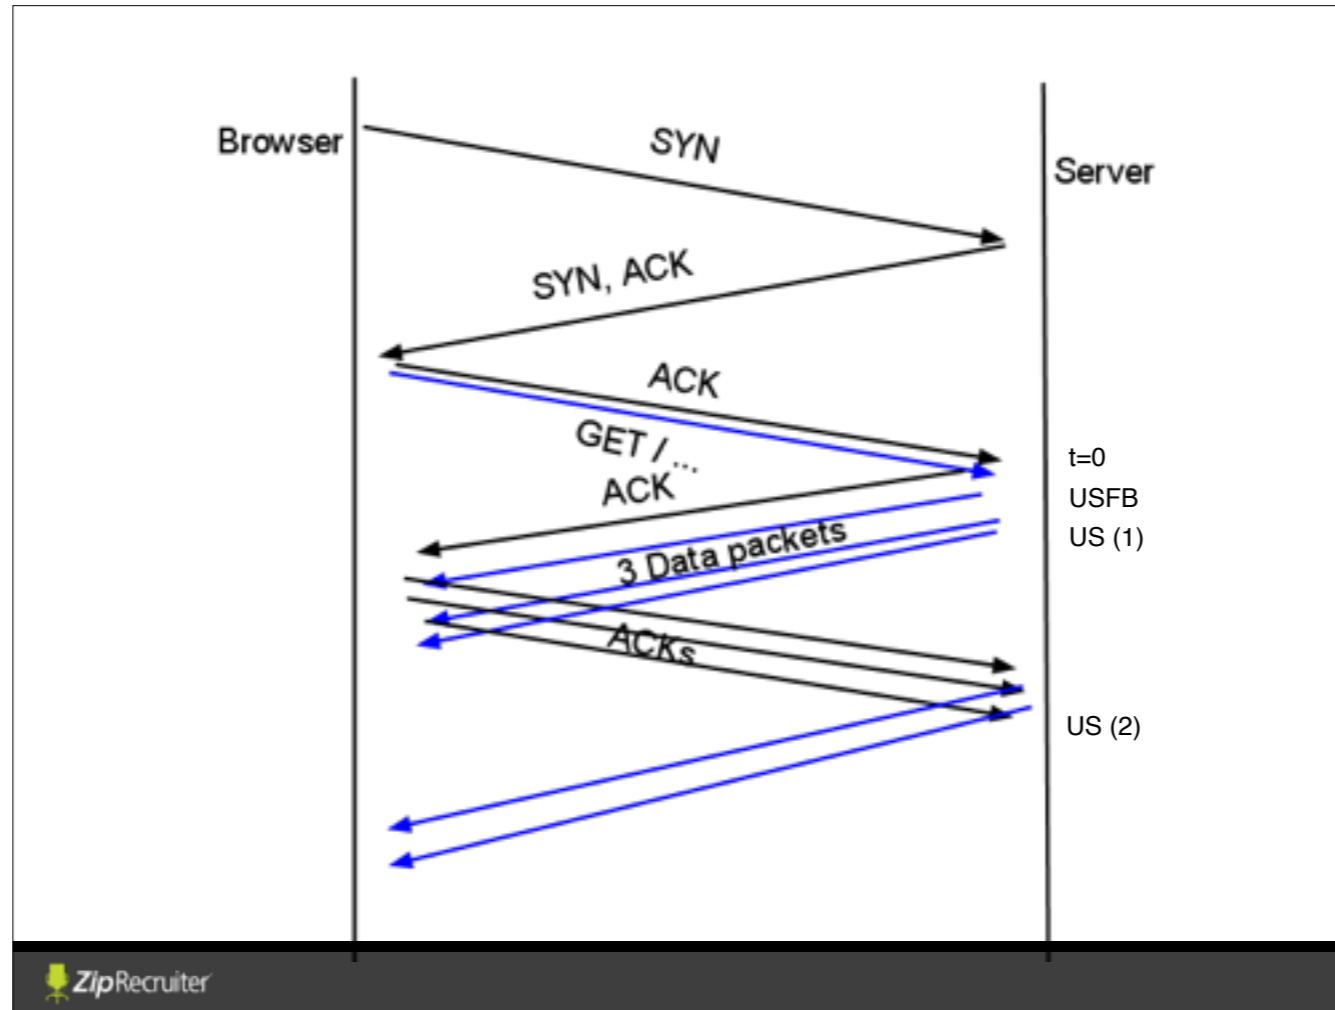
Why?

TCP receive window

TCP uses a **sliding window** flow control protocol. In each TCP segment, the receiver specifies in the *receive window* field the amount of additionally received data (in bytes) that it is willing to buffer for the connection. The sending host can send only up to that amount of data before it must wait for an acknowledgment and window update from the receiving host.

https://en.wikipedia.org/wiki/Transmission_Control_Protocol

- Watched HTTP TCP packets in Wireshark
- Akamai has a high starting window size, 12 packets.
- Can send 12 packets * 1430 bytes = 16.5kb without ACK.



eg: 5 packets with default of 3.

default is 3. Should be increased, is increased by CDN proxies.

akamai uses 10 or 12.

Apache records US time as soon as last packet is sent to local TCP buffer.

<http://www.cdnplanet.com/blog/tune-tcp-initwnd-for-optimum-performance/>

case 1: 3 packets

case 2: 5 packets

Success?

We fixed the glitch



fixed the metric.

marginally improved user experience

didn't change much for google, so not much bump.

LEARNED, DOCUMENTED, EXPLAINED, SHARED.

Learning

29

- misconfig meant akamai was claiming 99% cache hit, when it was more like a 90% cache miss
- Should set up local cache (varnish) in DC and use that for akamai origin.
- guesses become assumptions become institutional knowledge and lore. Check your assumptions.

**LEARN
DOCUMENT
IMPROVE**


**EXPLAIN
SHARE**

#devops

30

LEARNED, DOCUMENTED, IMPROVED, EXPLAINED, SHARED.

#DEVOPS IS CULTURE



Don't let truth be lost for myth.

Questions?



@spazm

about.me/spazm

www.lowlevelmanager.com

andrew.grangaard+scale14x@gmail.com



Follow me on the twitters, linked-in, facespace, etc.

Simba loves you!

Thank you ZipRecruiter for time for conferences. We're hiring!

Bonus Material

404 spike correlates perfectly with release push.



- PM complained about increase in daily 404 errors.
- I Isolated one specific php library error.
- Correlated against releases
- Cause Release script looped through servers updating the build and moving the symlink. then looped again restarting.
- php was in an unstable between move and restart, which grew as it took longer to install the build and more servers.



Thank You!

www.ZipRecruiter.com