

SYNTHETIC AND RUM

BEST OF BOTH WORLDS



Mark Zeman
@SpeedCurve



Cliff Crocker
@cliffcrocker
SOASTA

ACTIVE MONITORING

- Simulated health-checks of the system and its parts
- Scripted journeys through an application
- “Creation” of a singular or set of events and/or actions
- Synthetic is a form of active monitoring



PASSIVE MONITORING

- Listening to actual traffic as it moves through a system
- Observing behavior of the system and/or its operator
- “Measurement” of actions or events
- RUM is a form of passive monitoring



Our focus is on the experience of the crowd



RUM 101


Real User Measurement (RUM) is a technology for collecting performance metrics directly from the browser of an end user.

- Involves instrumentation of your website via JavaScript
- Measurements are fired across the network to a collection point through a small request object (beacon)

<JS>



<beacon>




THE GOOD PARTS - RUM

- Always on
- Every user, every browser, every network, anywhere in the world
- Ability to capture human behavior/events with performance data
- Only getting better (waterfalls)

Synthetic 101

Synthetic monitoring (for purposes of this discussion) refers to the use of automated agents (bots) to measure your website from different physical locations.

- A set 'path' or URL is defined
- Tests are run either adhoc or scheduled and data is collected



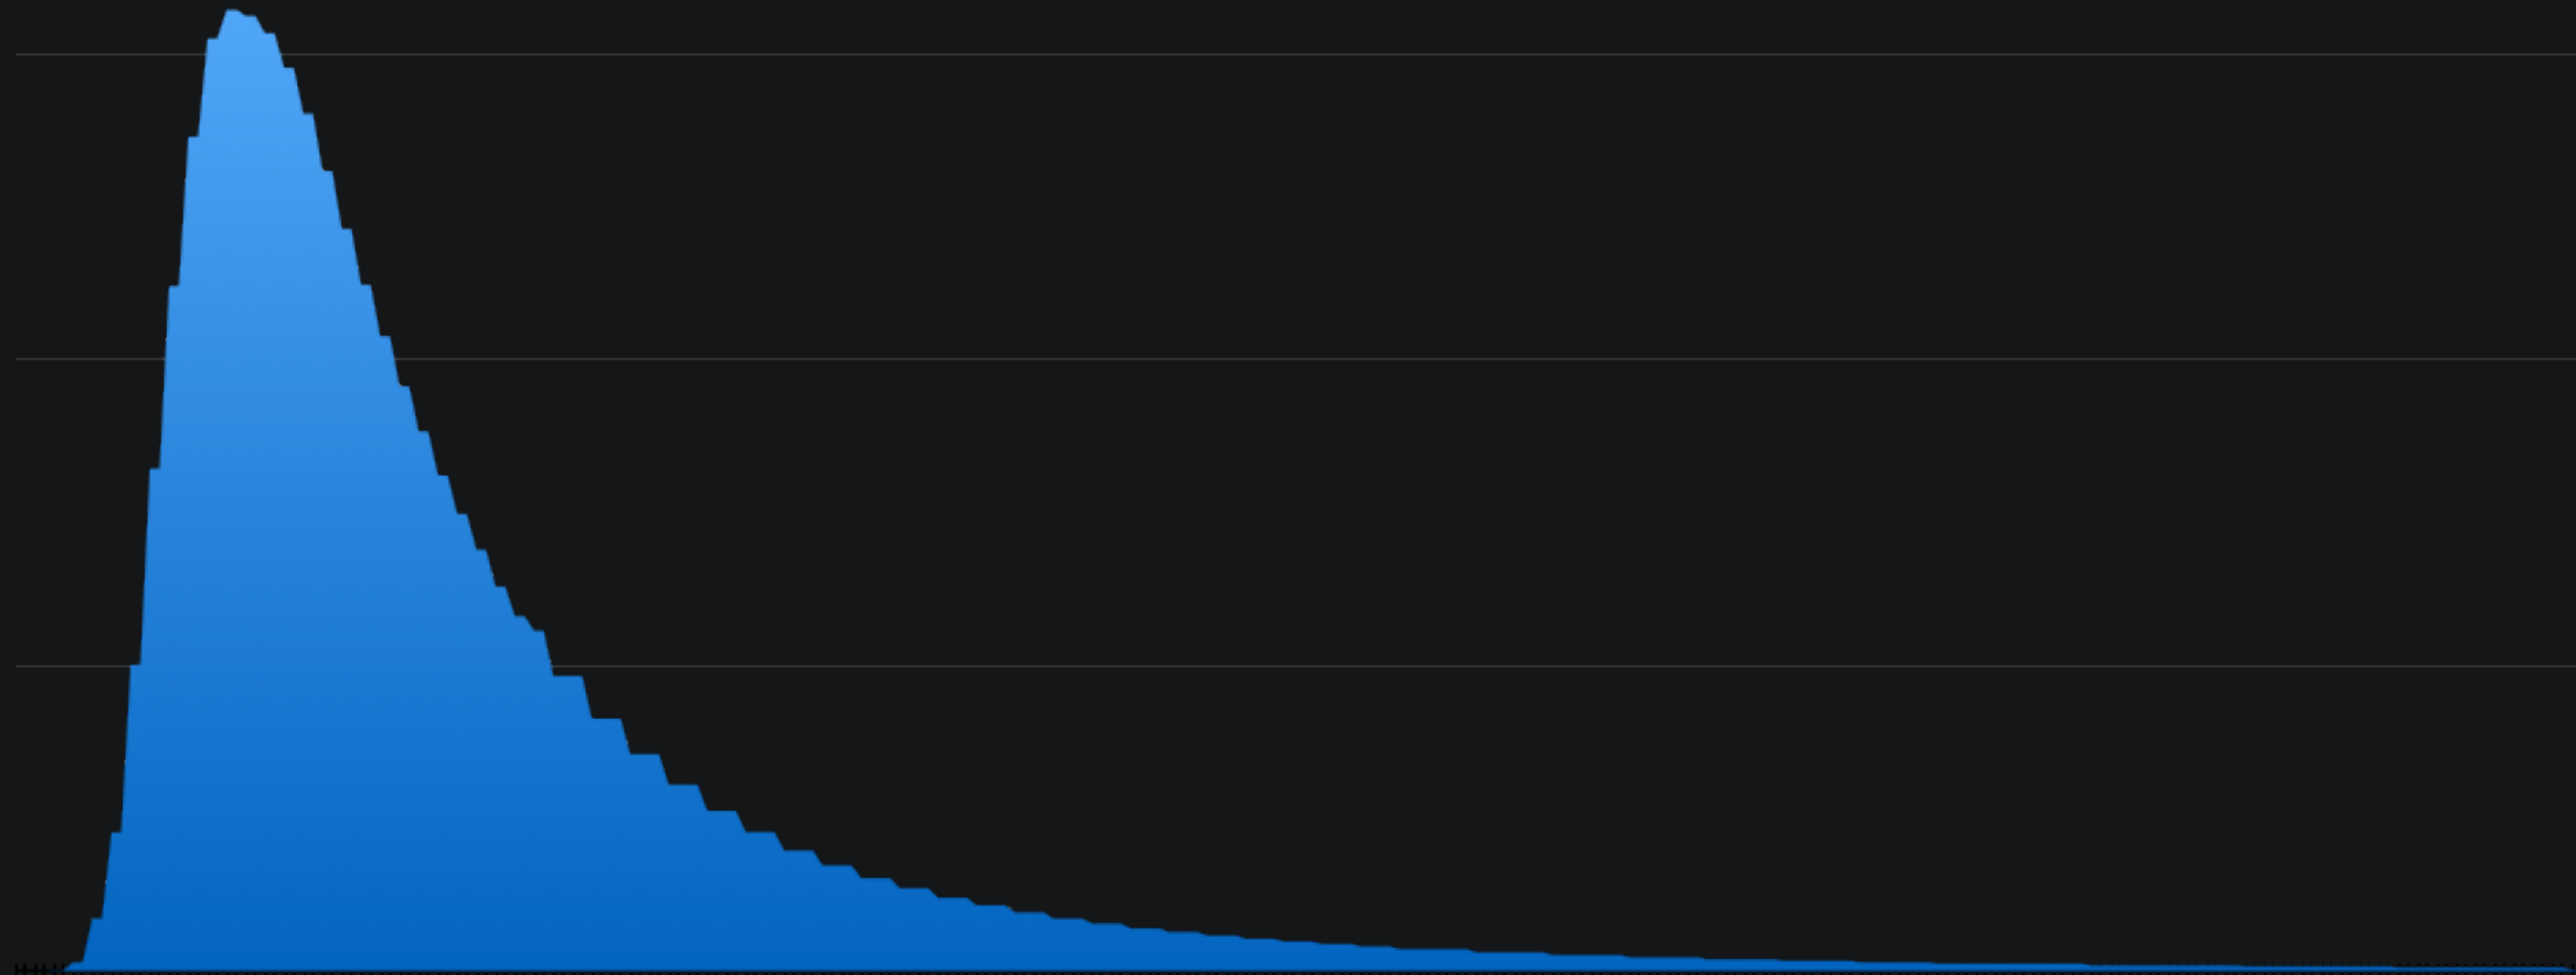
THE GOOD PARTS - SYNTHETIC

- Rich data collected (waterfall, videos/filmstrip, HTTP Headers)
- Consistent “clean room” baseline
- Nothing to install
- Doesn't require users/ability to measure pre-production, competition

COMMON THINGS WE HEAR ABOUT BOTH RUM AND SYNTHETIC

- “Why are these numbers so different?”
- “I don’t trust your data. Your numbers are wrong.”
- “How are you calculating Page Load time?”
- “I can’t share two sets of numbers with the business.”

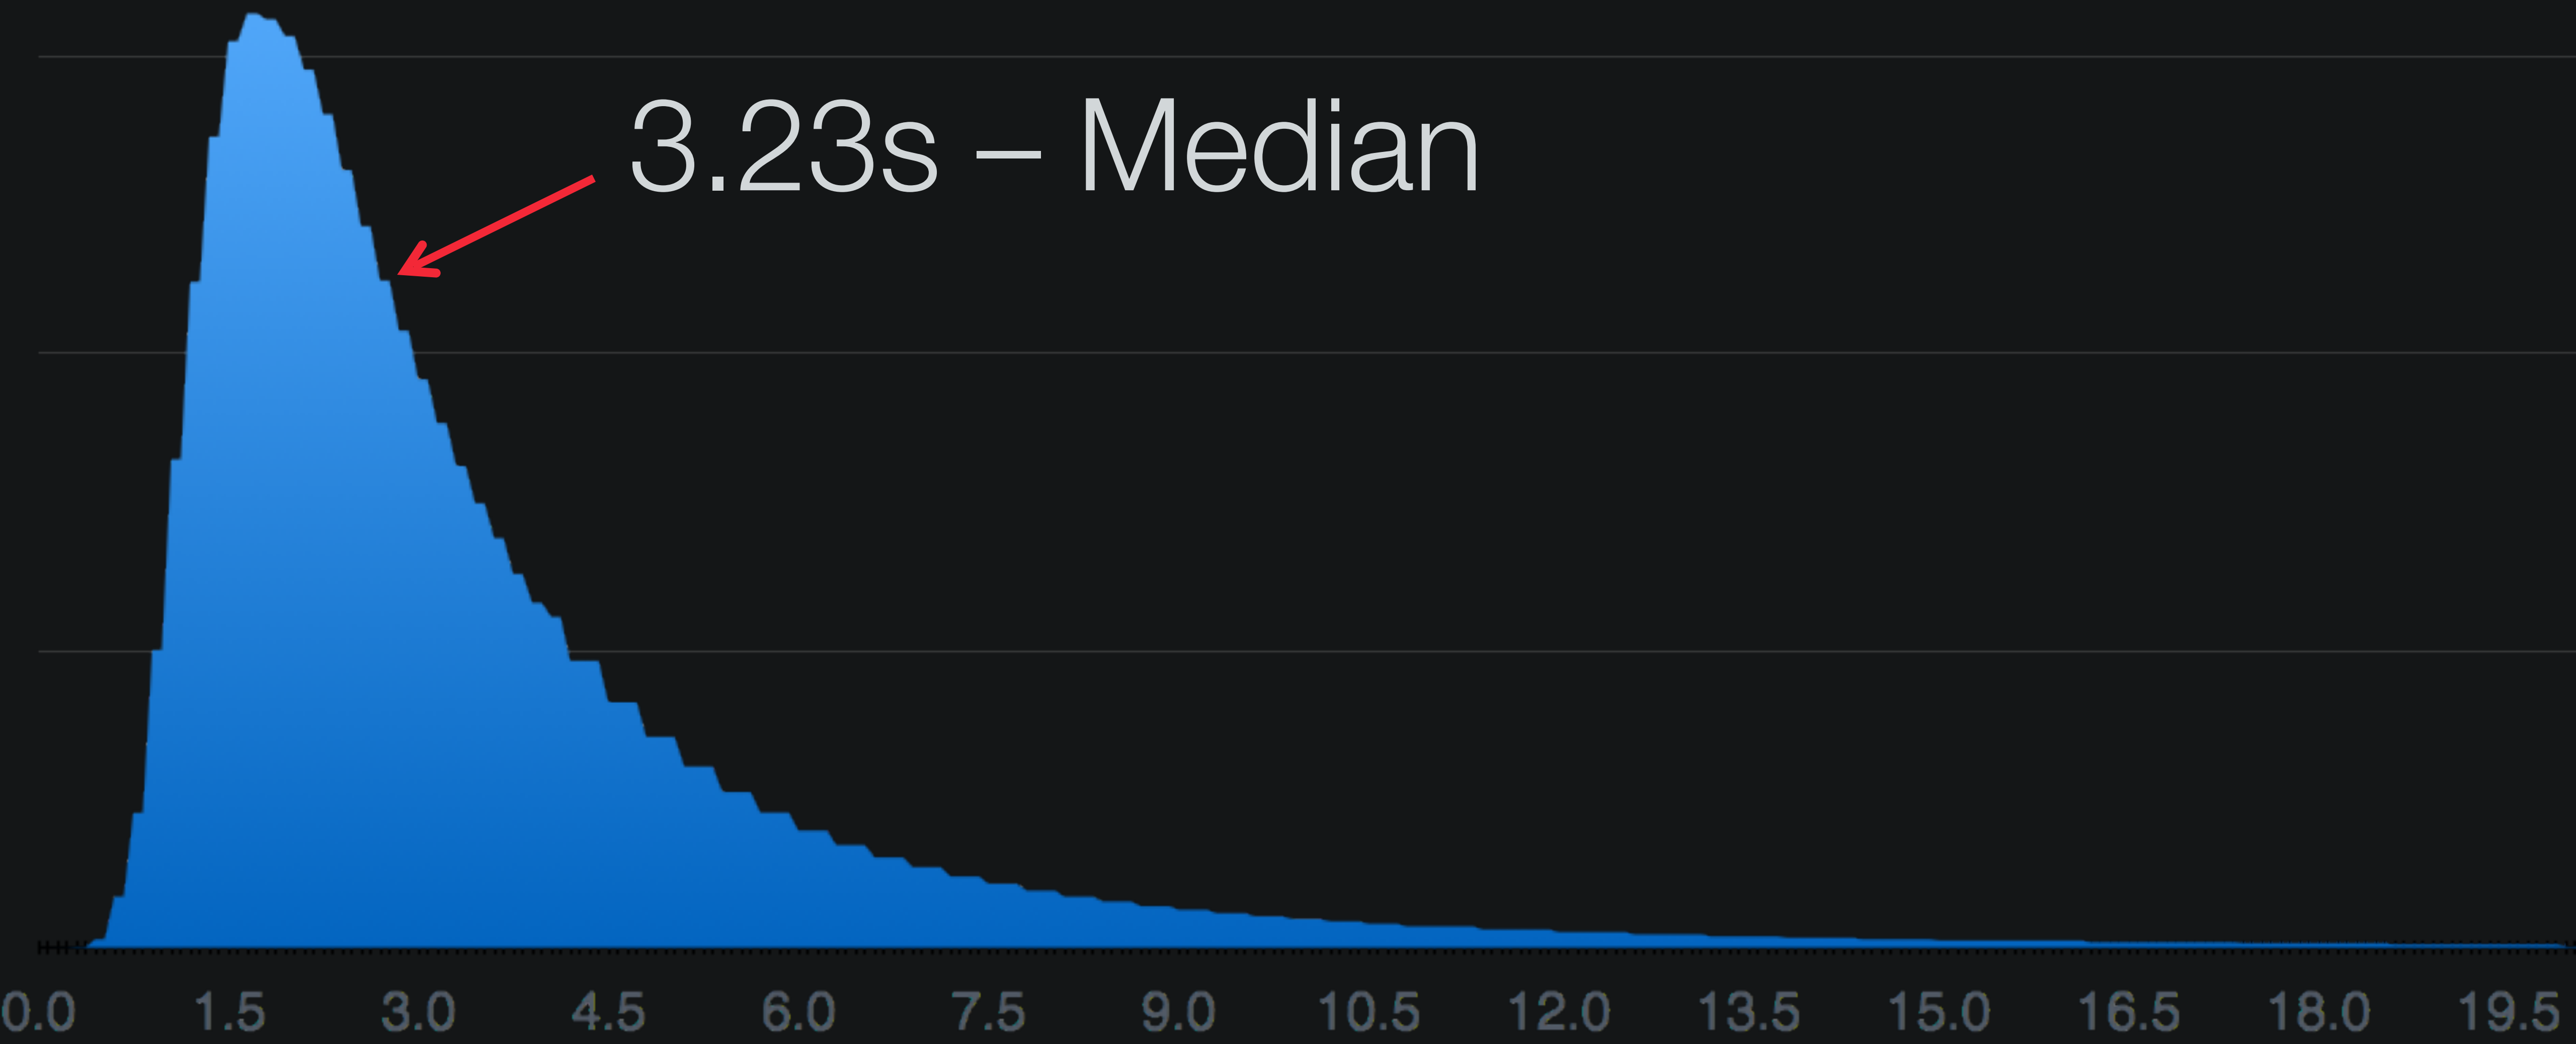
WHY ARE THE TWO NUMBERS SO FAR OFF?



Real users are not normal

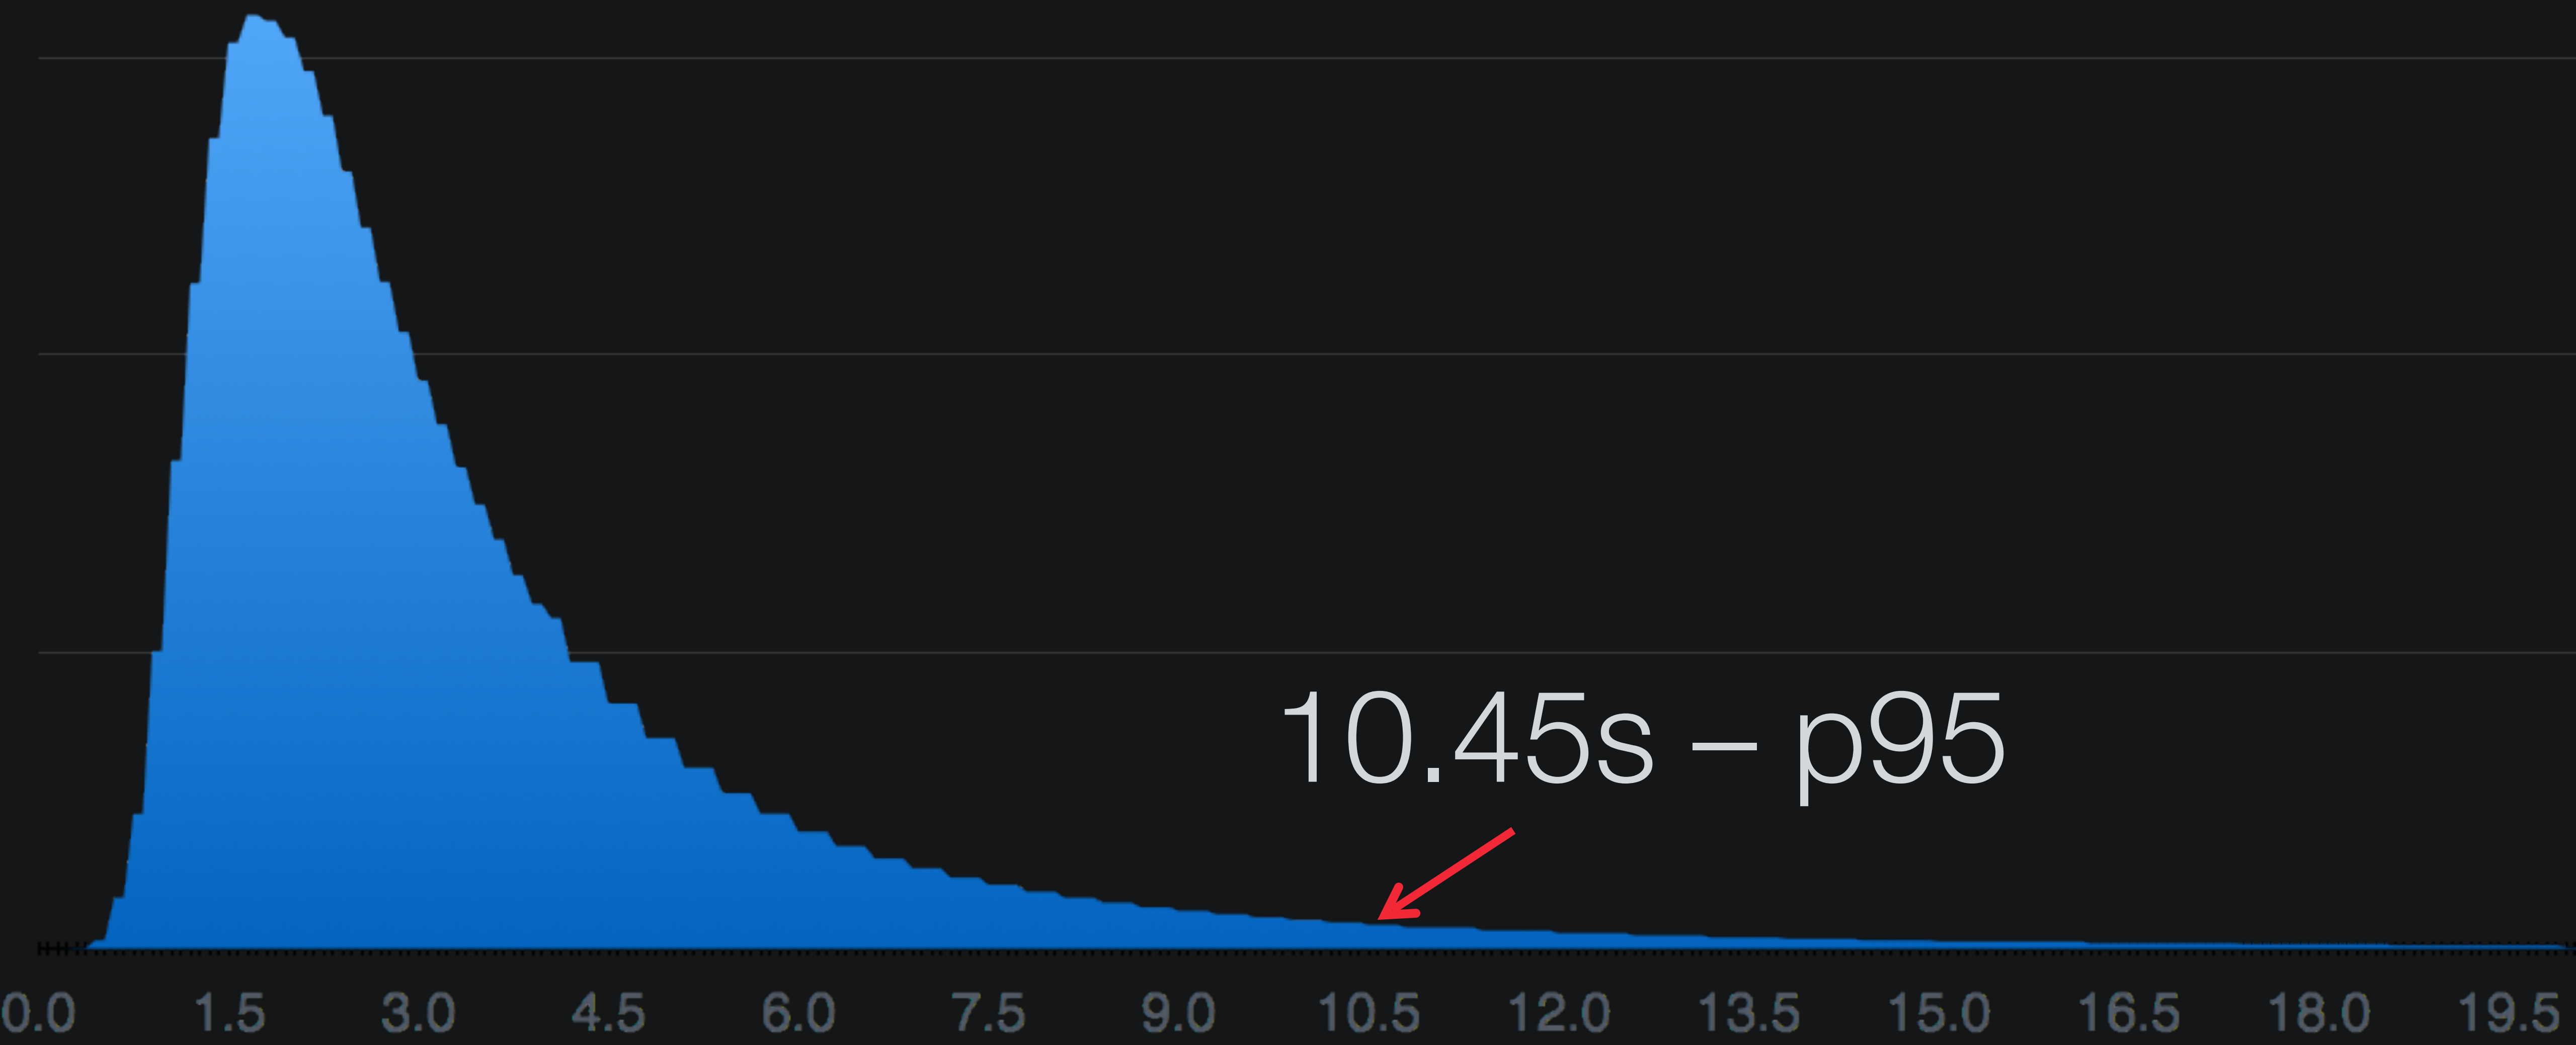
Page Load Times

All Countries, User Agents, Carriers, Operating Systems



Page Load Times

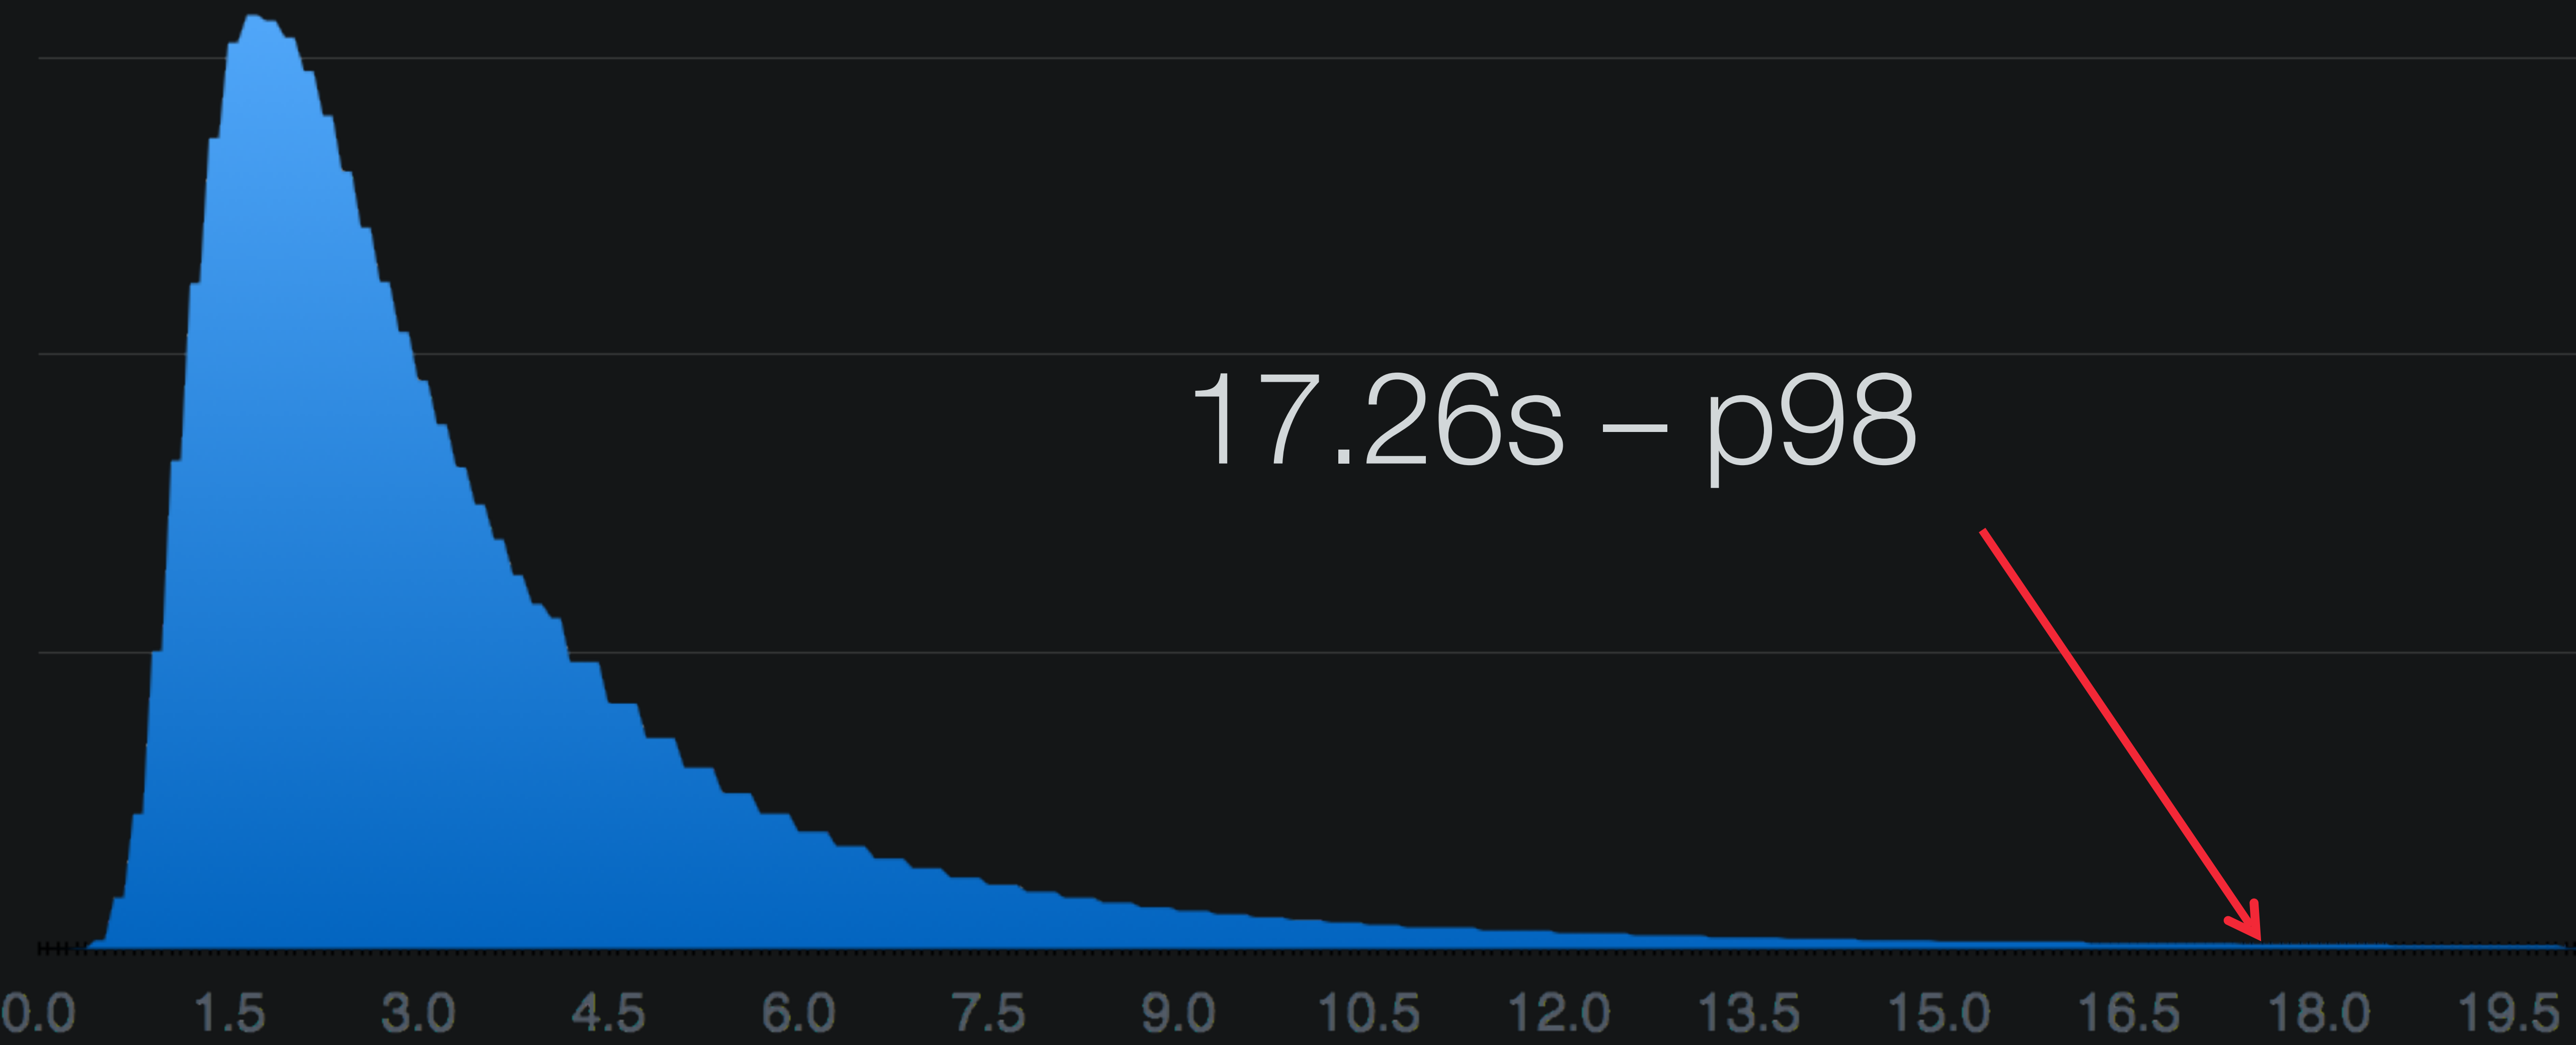
All Countries, User Agents, Carriers, Operating Systems



10.45s – p95


Page Load Times

All Countries, User Agents, Carriers, Operating Systems




Page Load Times

By Operating System



OS	Median	95 th Percentile	98 th Percentile
Windows 7	2.41s	9.29s	15.89s
Mac OS X/10	2.30s	8.11s	13.45s
iOS7	3.27s	10.64s	15.79s
Android 4	4.06s	14.30s	27.93s
iOS8	3.53s	11.54s	19.72s
Windows 8	2.67s	10.75s	18.74s

WHY ARE THE TWO NUMBERS SO FAR OFF?



“BUT, IT LOADS SO MUCH FASTER FOR ME??!!”

4 seconds



vs.

20 seconds



- 2015 Macbook Pro
- Warm browser cache
- FIOS

- x86 - Windows 7 VM
- Completely cold cache/dns
- Throttled bandwidth

RUM

SYNTHETIC

- Primary source of truth
- User behavior/ conversion
- Complete coverage

- Diagnostics
- Page Construction
- Pre-production & Competitive benchmarking

RUM

- Primary source of truth
- User behavior/
conversion
- Complete coverage

SYNTHETIC

- Diagnostics
- Page Construction
- Pre-production &
Competitive
benchmarking

You need both

THIS IS YOUR "ONE NUMBER"

START RENDER 1.58S	DNS 110MS	TCP 40MS	TTFB 650MS
DOM LOADING 680MS	DOM READY 2.75S	PAGE LOAD 3.05	FULLY LOADED 19.68S
USER TIMING 1.95S	RESOURCE TIMING 237MS	REQUESTS 116	BYTES IN 2,143KB
SPEED INDEX 3086	PAGESPEED SCORE 83	1S=\$27M	DOM ELEMENTS 3931
DOM SIZE 3,143KB	VISUALLY COMPLETE 15.2S	REDIRECT 0S	SSL NEGOTIATION 93MS

THE MANY FACES OF PERFORMANCE



C-Level



Operations



Developer



Designer







START RENDER	DNS	TCP	TTFB
DOM LOADING	DOM READY	PAGE LOAD	FULLY LOADED
USER TIMING	RESOURCE TIMING	REQUESTS	BYTES IN
SPEED INDEX	PAGESPEED SCORE	1S-\$\$	DOM ELEMENTS
DOM SIZE	VISUALLY COMPLETE	REDIRECT	SSL NEGOTIATION

START RENDER	DNS	TCP	TTFB
DOM LOADING	DOM READY	PAGE LOAD	FULLY LOADED
USER TIMING	RESOURCE TIMING	REQUESTS	BYTES IN
SPEED INDEX	PAGESPEED SCORE	1S-\$\$	DOM ELEMENTS
DOM SIZE	VISUALLY COMPLETE	REDIRECT	SSL NEGOTIATION

START RENDER	DNS	TCP	TTFB
DOM LOADING	DOM READY	PAGE LOAD	FULLY LOADED
USER TIMING	RESOURCE TIMING	REQUESTS	BYTES IN
SPEED INDEX	PAGESPEED SCORE	1S-\$\$	DOM ELEMENTS
DOM SIZE	VISUALLY COMPLETE	REDIRECT	SSL NEGOTIATION






START RENDER	DNS	TCP	TTFB
DOM LOADING	DOM READY	PAGE LOAD	FULLY LOADED
USER TIMING	RESOURCE TIMING	REQUESTS	BYTES IN
SPEED INDEX	PAGESPEED SCORE	1S-\$\$	DOM ELEMENTS
DOM SIZE	VISUALLY COMPLETE	REDIRECT	SSL NEGOTIATION



Analysis



START RENDER	DNS	TCP	TTFB
DOM LOADING	DOM READY	PAGE LOAD	FULLY LOADED
USER TIMING	RESOURCE TIMING	REQUESTS	BYTES IN
SPEED INDEX	PAGESPEED SCORE	1S=\$27M	DOM ELEMENTS
DOM SIZE	VIZ COMPLETE	REDIRECT	SSL NEGOTIATION



Analysis



START RENDER	DNS	TCP	TTFB
DOM LOADING	DOM READY	PAGE LOAD	FULLY LOADED
USER TIMING	RESOURCE TIMING	REQUESTS	BYTES IN
SPEED INDEX	PAGESPEED SCORE	1S=\$27M	DOM ELEMENTS
DOM SIZE	VIZ COMPLETE	REDIRECT	SSL NEGOTIATION

“What is the right number for me to communicate to the organization?”

– CTO



REALISTICALLY, THERE IS NOT ONE NUMBER



1s=\$27M

**Page Load
3.05s**

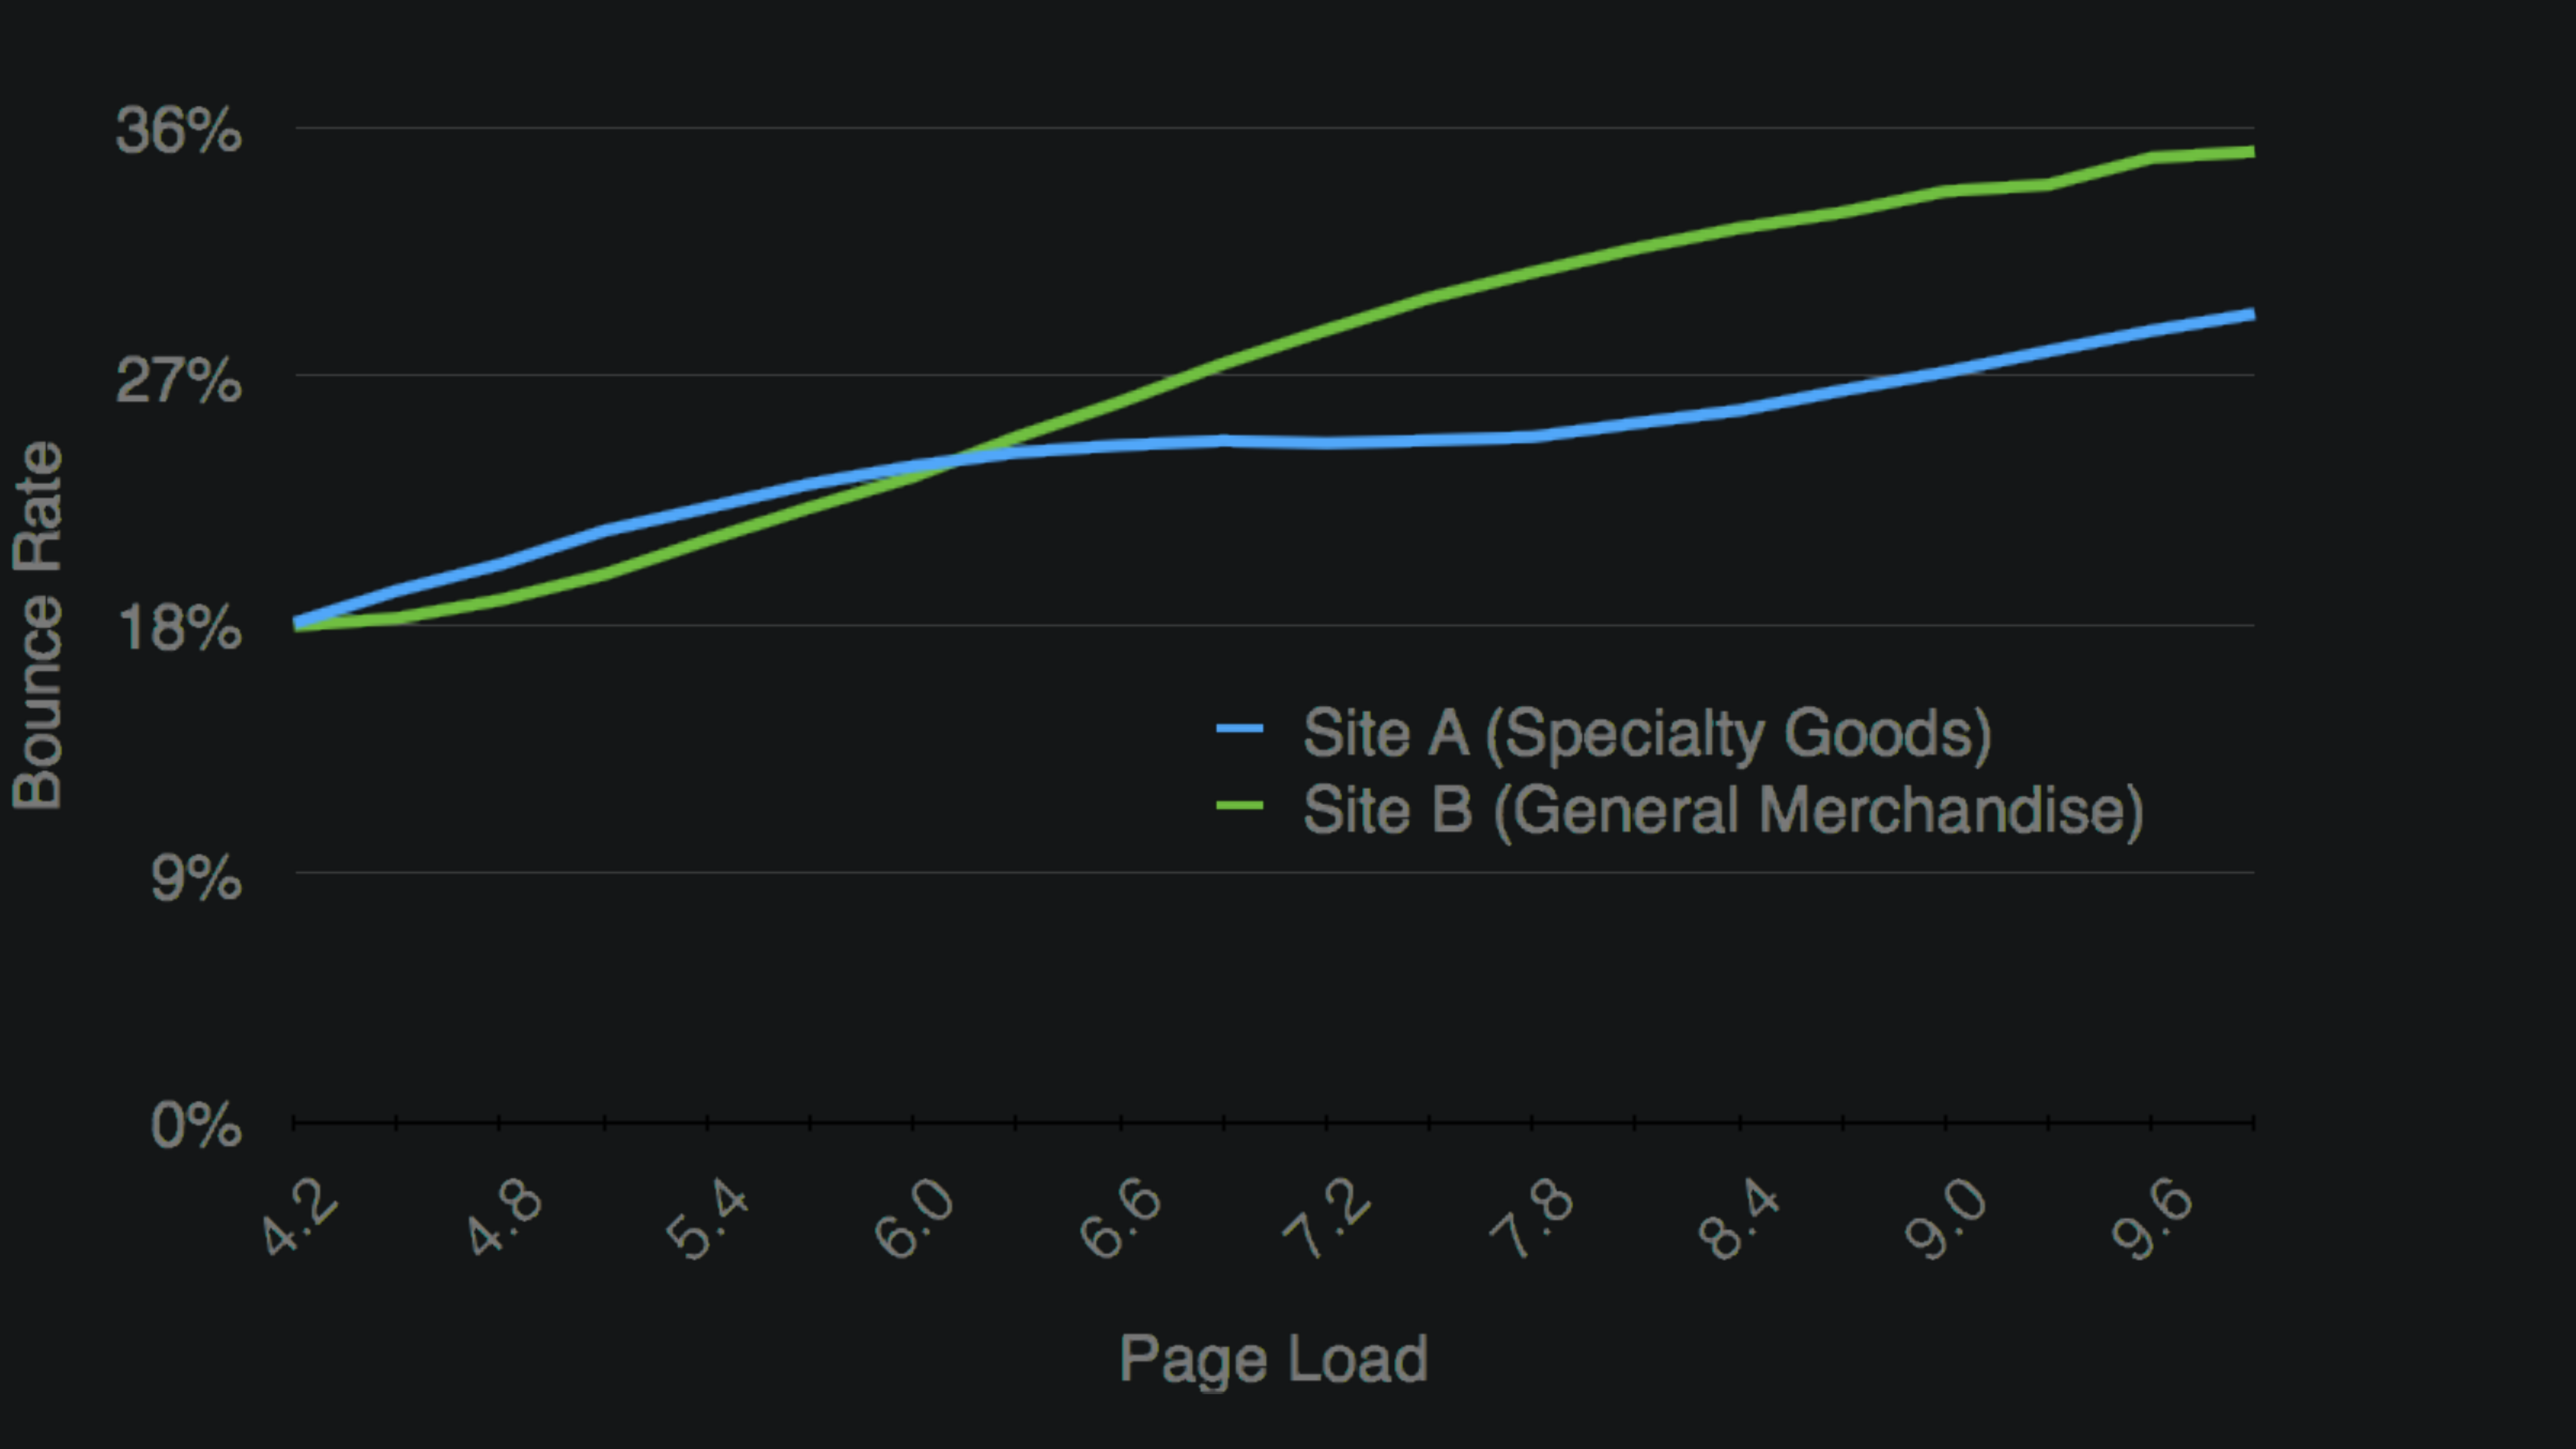
**Start Render
1.58s**

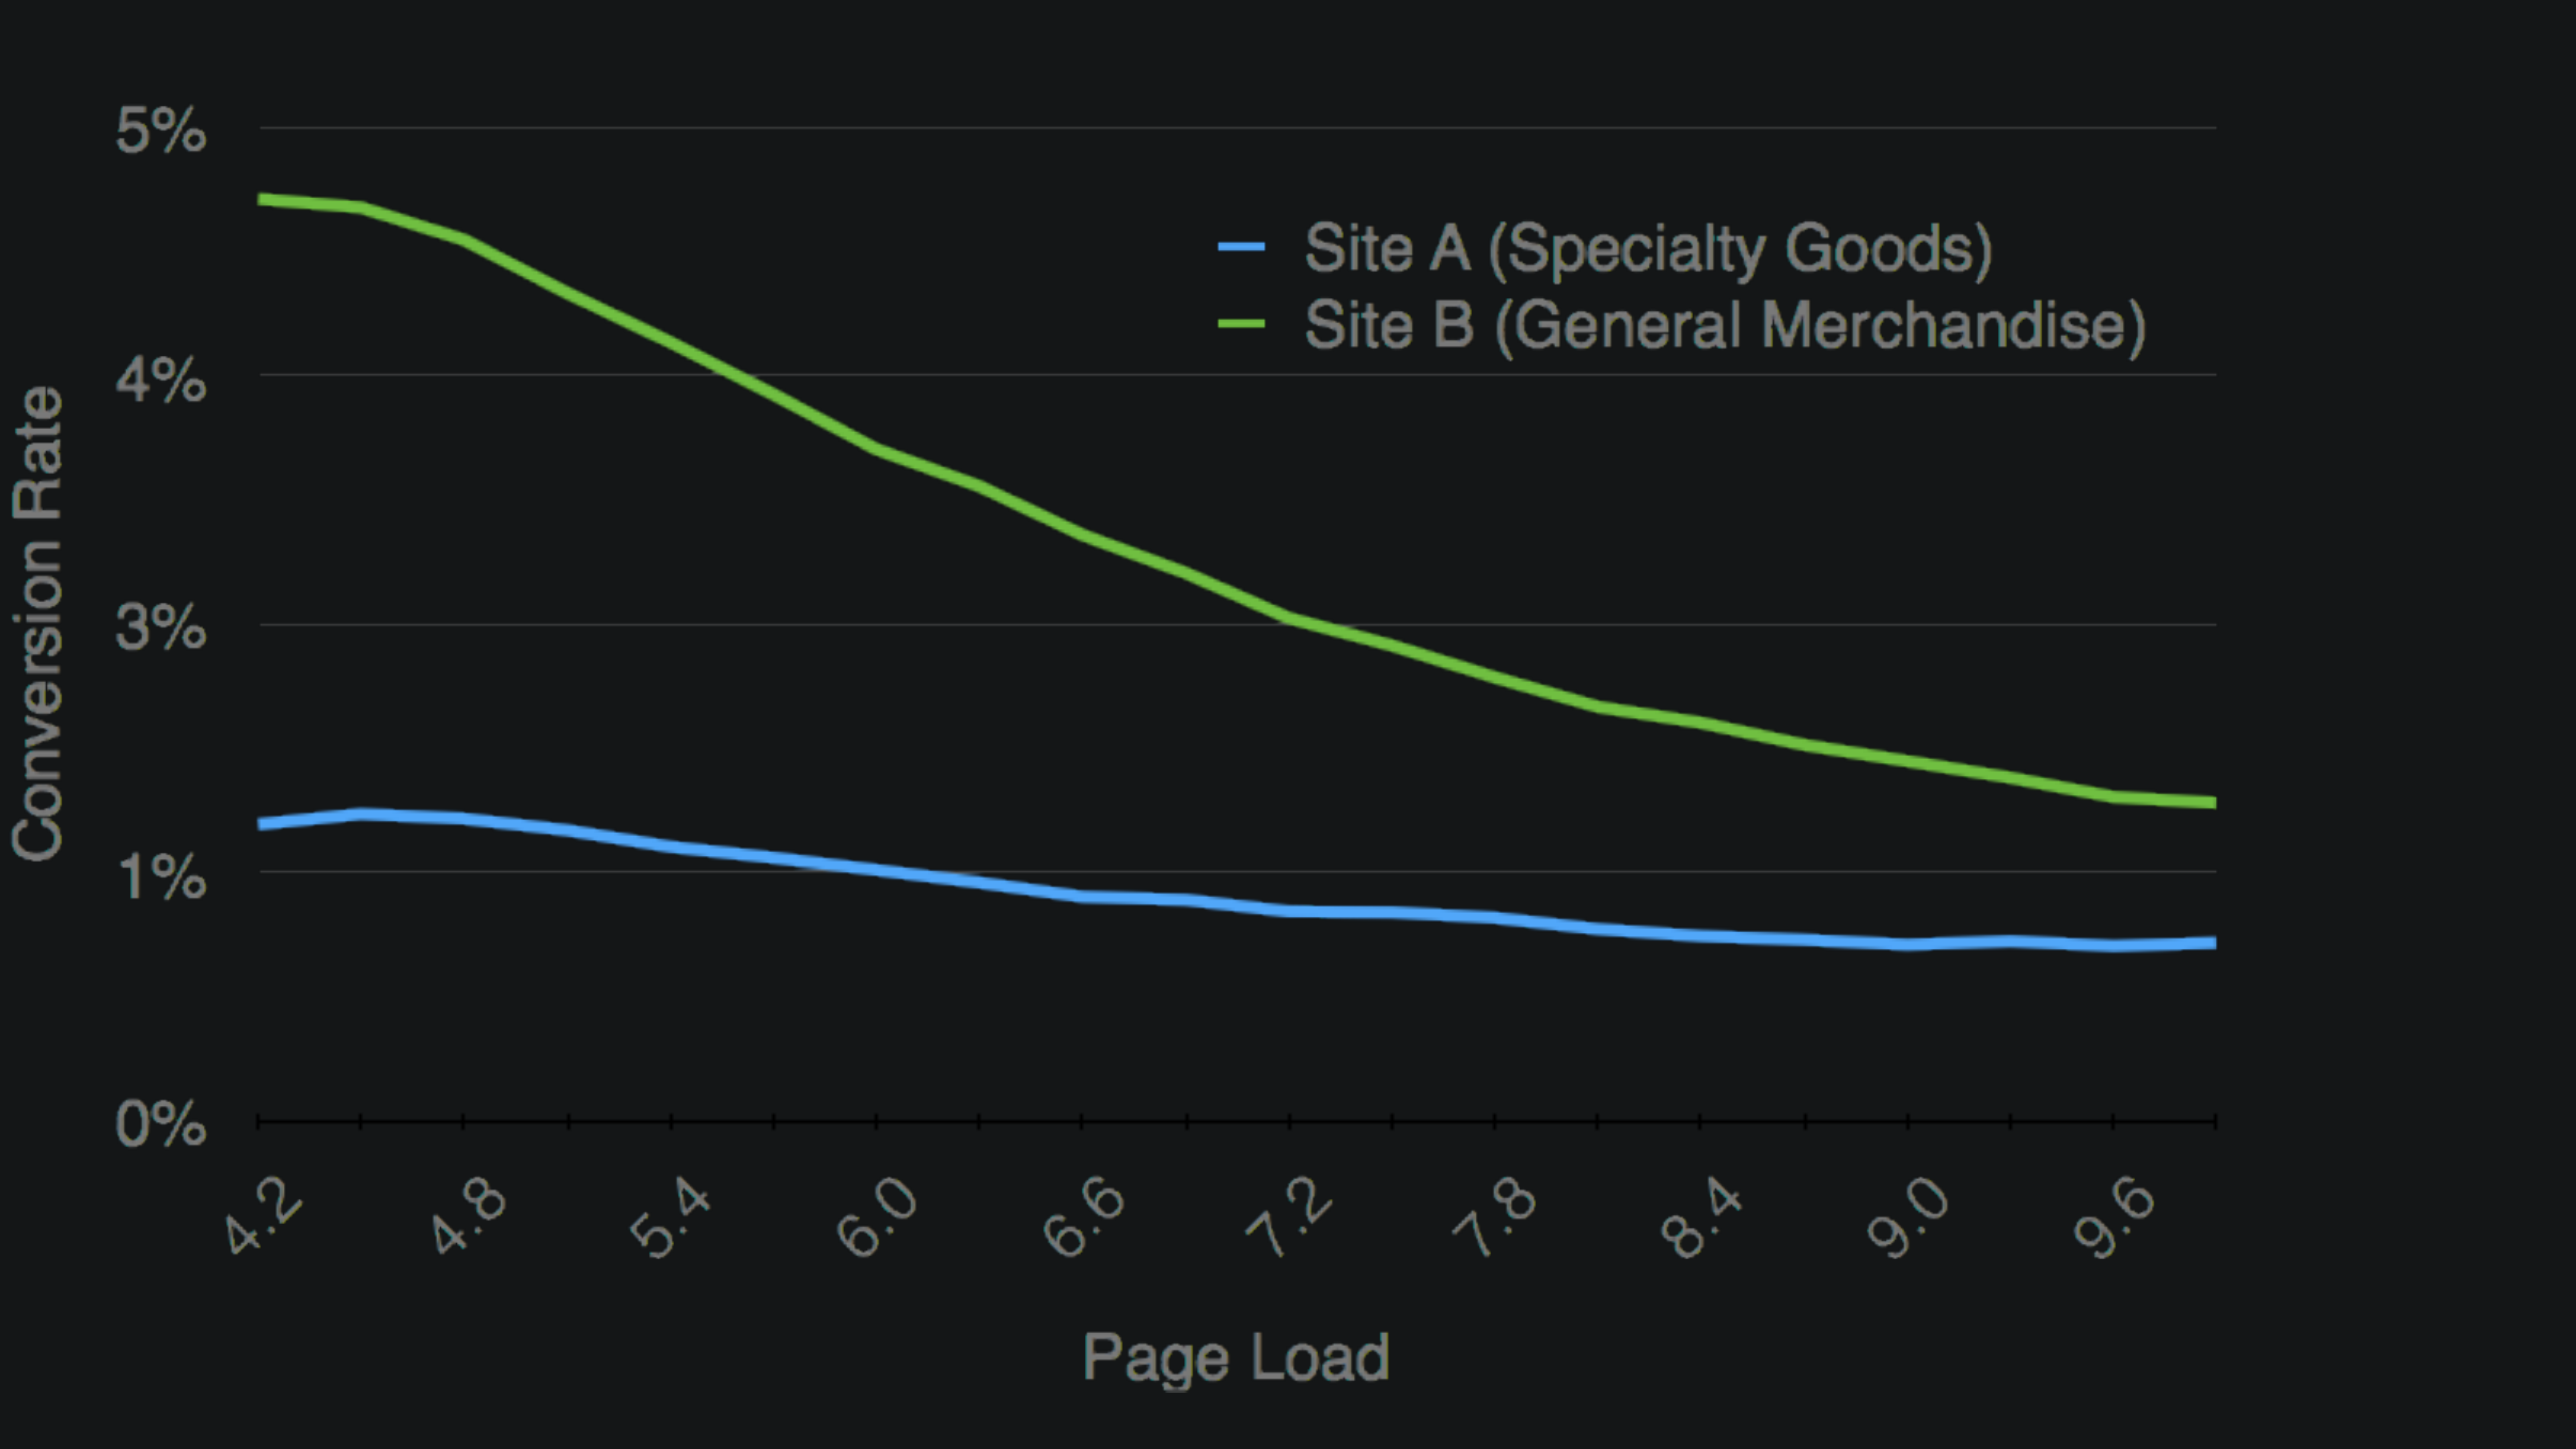
**Speed Index
3086**

“How is site performance impacting revenue?”

– CEO

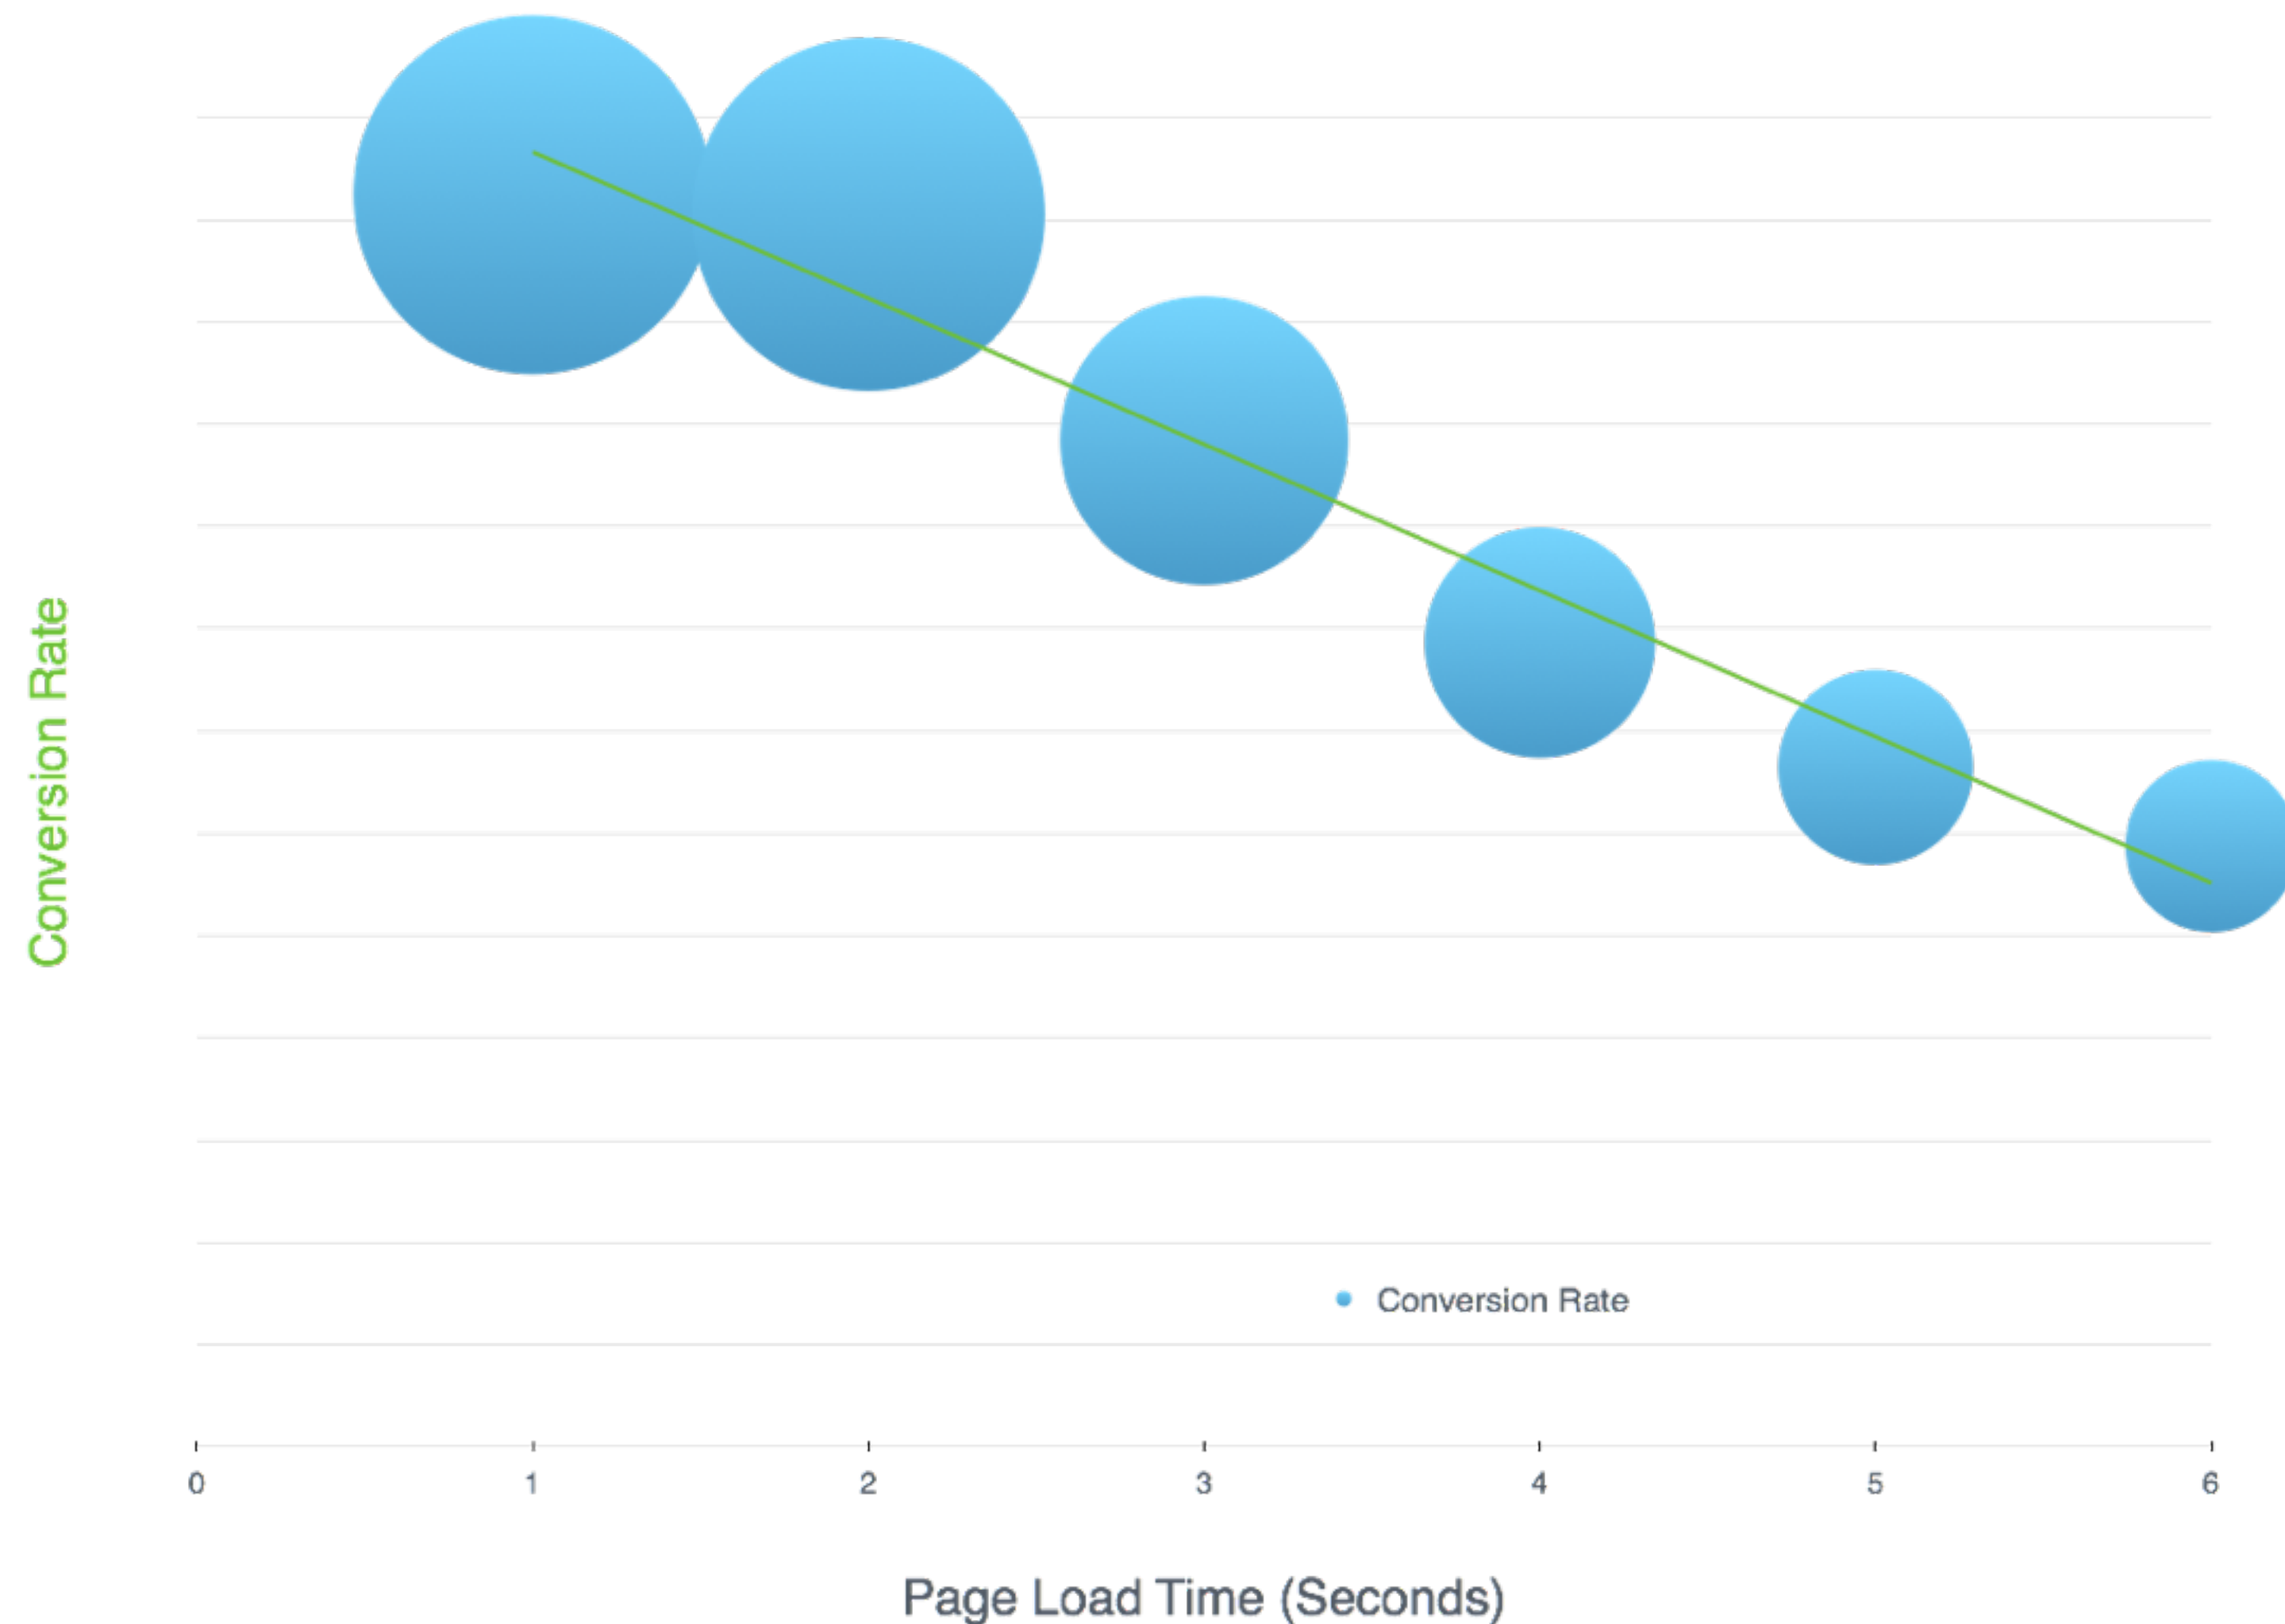






TOP OF FUNNEL IMPACT (BROWSE PAGES)


NOT ALL PAGES ARE CREATED EQUAL



For a typical eCommerce site, conversion rate drops by up to **50%** when “browse” pages increase from 1 to 6 seconds


BOTTOM OF FUNNEL IMPACT (CHECKOUT PAGES)

NOT ALL PAGES ARE CREATED EQUAL



However, there is much less impact to conversion when “checkout” pages degrade

HOW FAST SHOULD I BE?



“How do we compare to the competition?”

– CEO



MEDIAN

BACKEND | START RENDER | DOM | FULLY LOADED | SPEEDINDEX | PAGESPEED

Guardian Beta

Huffington

Mashable

NY Times

Readwrite

Smashing

TechCrunch

The Verge

0.8s

2.1s

1.2s

1.9s

1.4s

1s

3.2s

4.6s

#1

#6 - 163% slower

#3 - 50% slower


#5 - 138% slower

#4 - 75% slower

#2 - 25% slower

#7 - 300% slower

#8 - 475% slower



MEDIAN

BACKEND | START RENDER | DOM | FULLY LOADED | SPEEDINDEX | PAGESPEED

Guardian Beta

Huffington

Mashable

NY Times

Readwrite

Smashing

TechCrunch

The Verge

#1

#6

#3

#5

#4

#2

#7

#8

163% slower

50% slower


138% slower

75% slower

25% slower

300% slower

475% slower



“How is our site doing in Spain? Are we winning there?”

– CEO

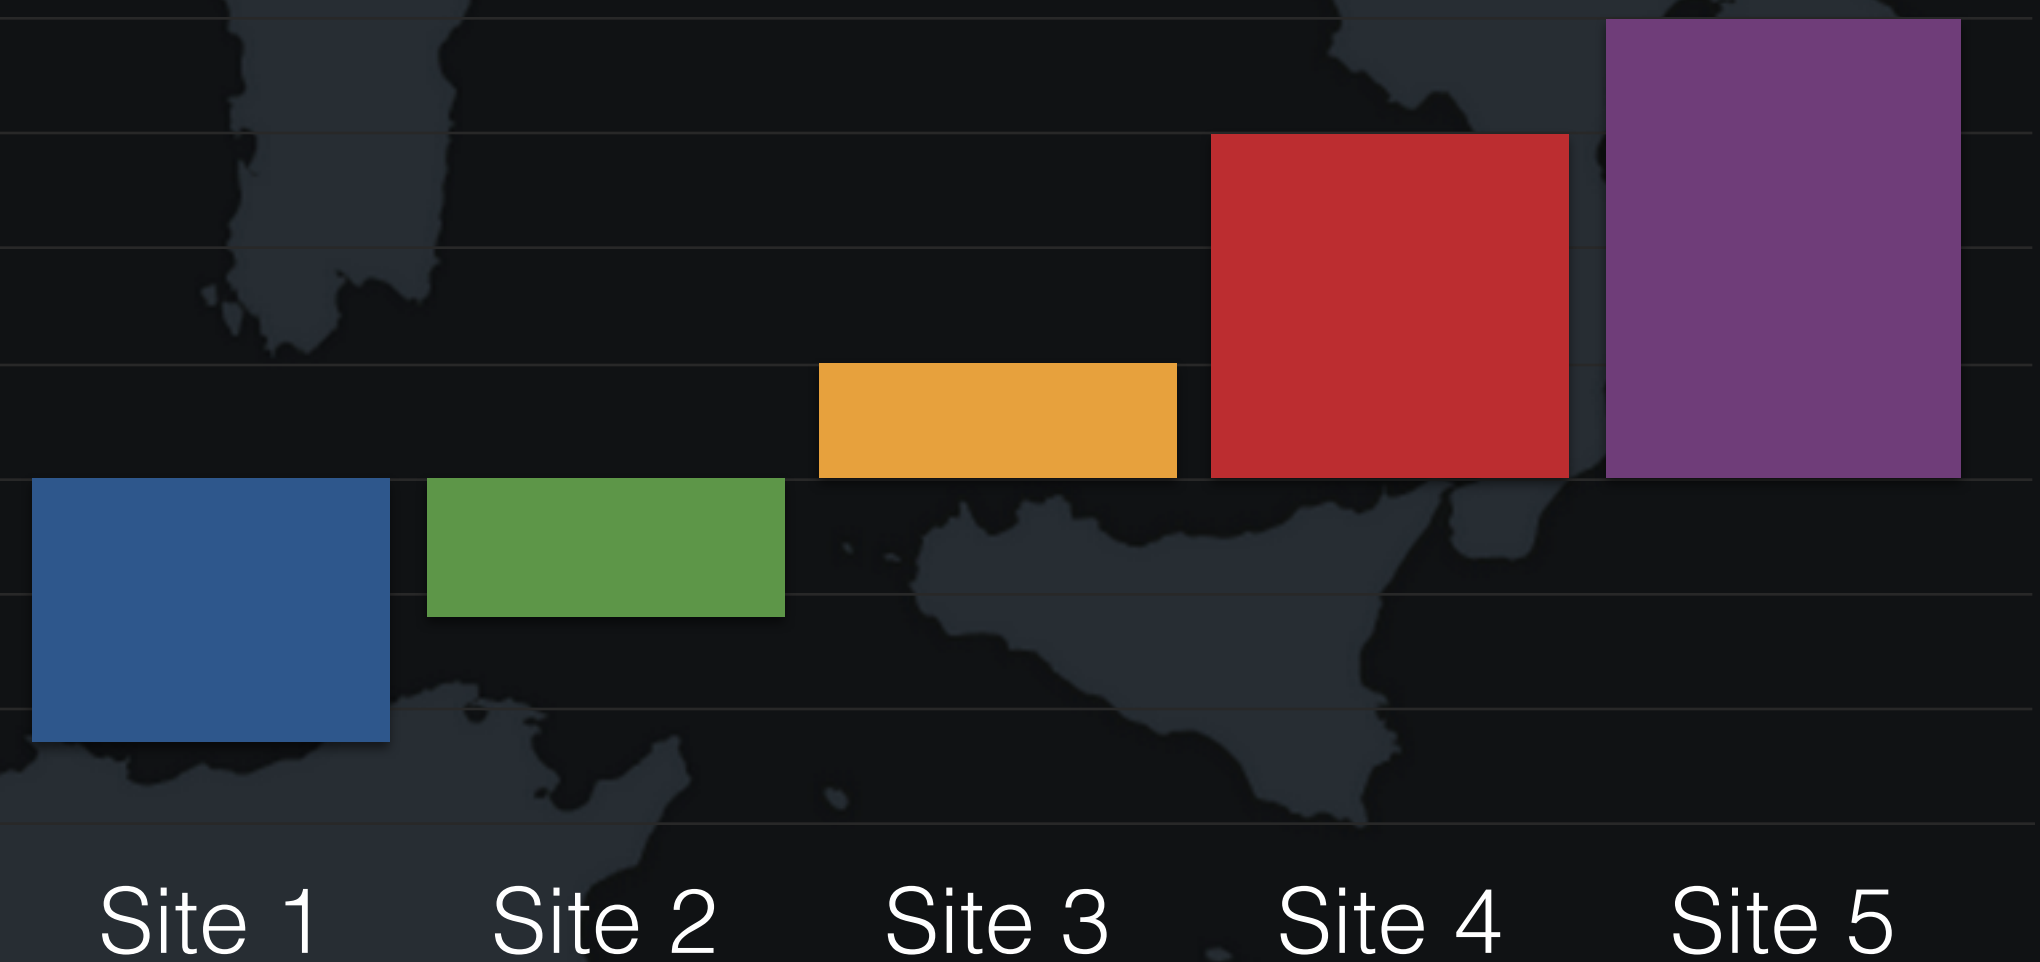


SITE PERFORMANCE - SPAIN

3.91s

- **Represents 1% of our overall traffic**
- **3rd overall in our competitive index**

Competitive Index



“How are third parties impacting my site?”

– VP ECOMMERCE




3RD PARTY STUDY - UNIQUE DOMAINS PER PAGE

■ Primary domains ■ 3rd Party domains


Domains:
Median - 32
Min - 25
Max - 87

Domains




Alexa Top 25 Retail

3RD PARTY STUDY



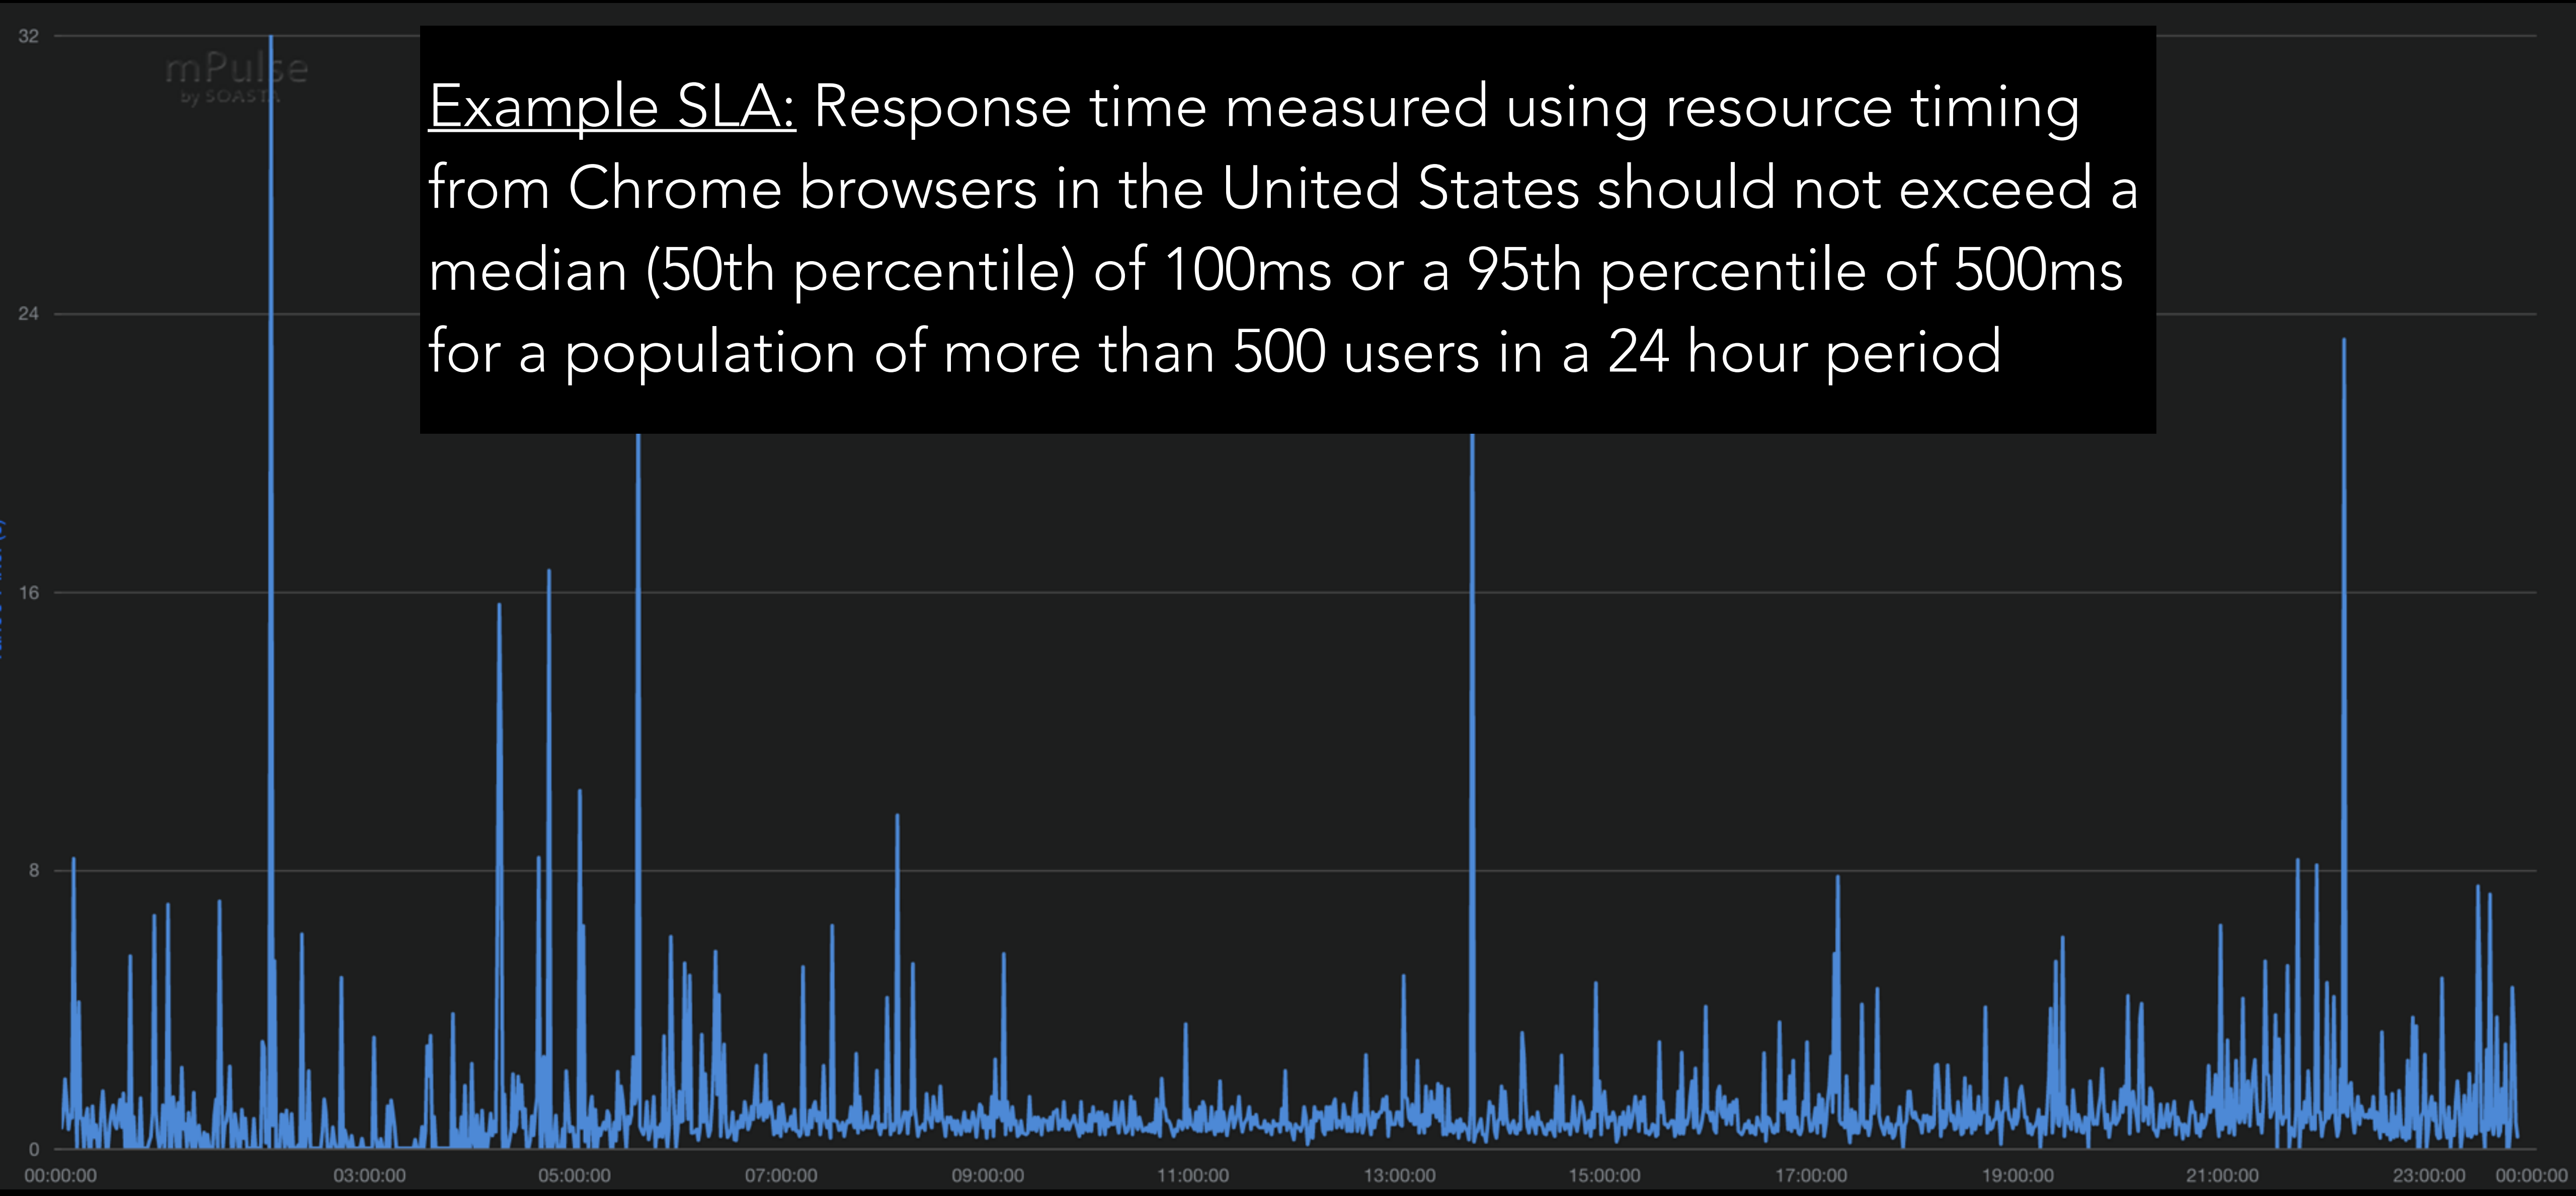
91% of domains
are 3rd parties

HUFFINGTON POST - THIRD PARTY IMPACT USING PETTY CASH



SETTING MEANINGFUL SLAS

Example SLA: Response time measured using resource timing from Chrome browsers in the United States should not exceed a median (50th percentile) of 100ms or a 95th percentile of 500ms for a population of more than 500 users in a 24 hour period



“How should I measure my site in pre-production?”

– FRONT END DEVELOPER

JAVASCRIPT SIZE BUDGET

Current JS Size

201KB

JS Size Budget


300KB

Decrease over 30 days (-62%)

-327KB

Remaining budget (33%)

99KB



CSS SIZE BUDGET

Current CSS Size

120KB

CSS Size Budget


100KB

Decrease over 30 days (-5%)

-6KB

Over budget (-20%)

-20KB




“Should I optimize my site for mobile?”

– VP ECOMMERCE

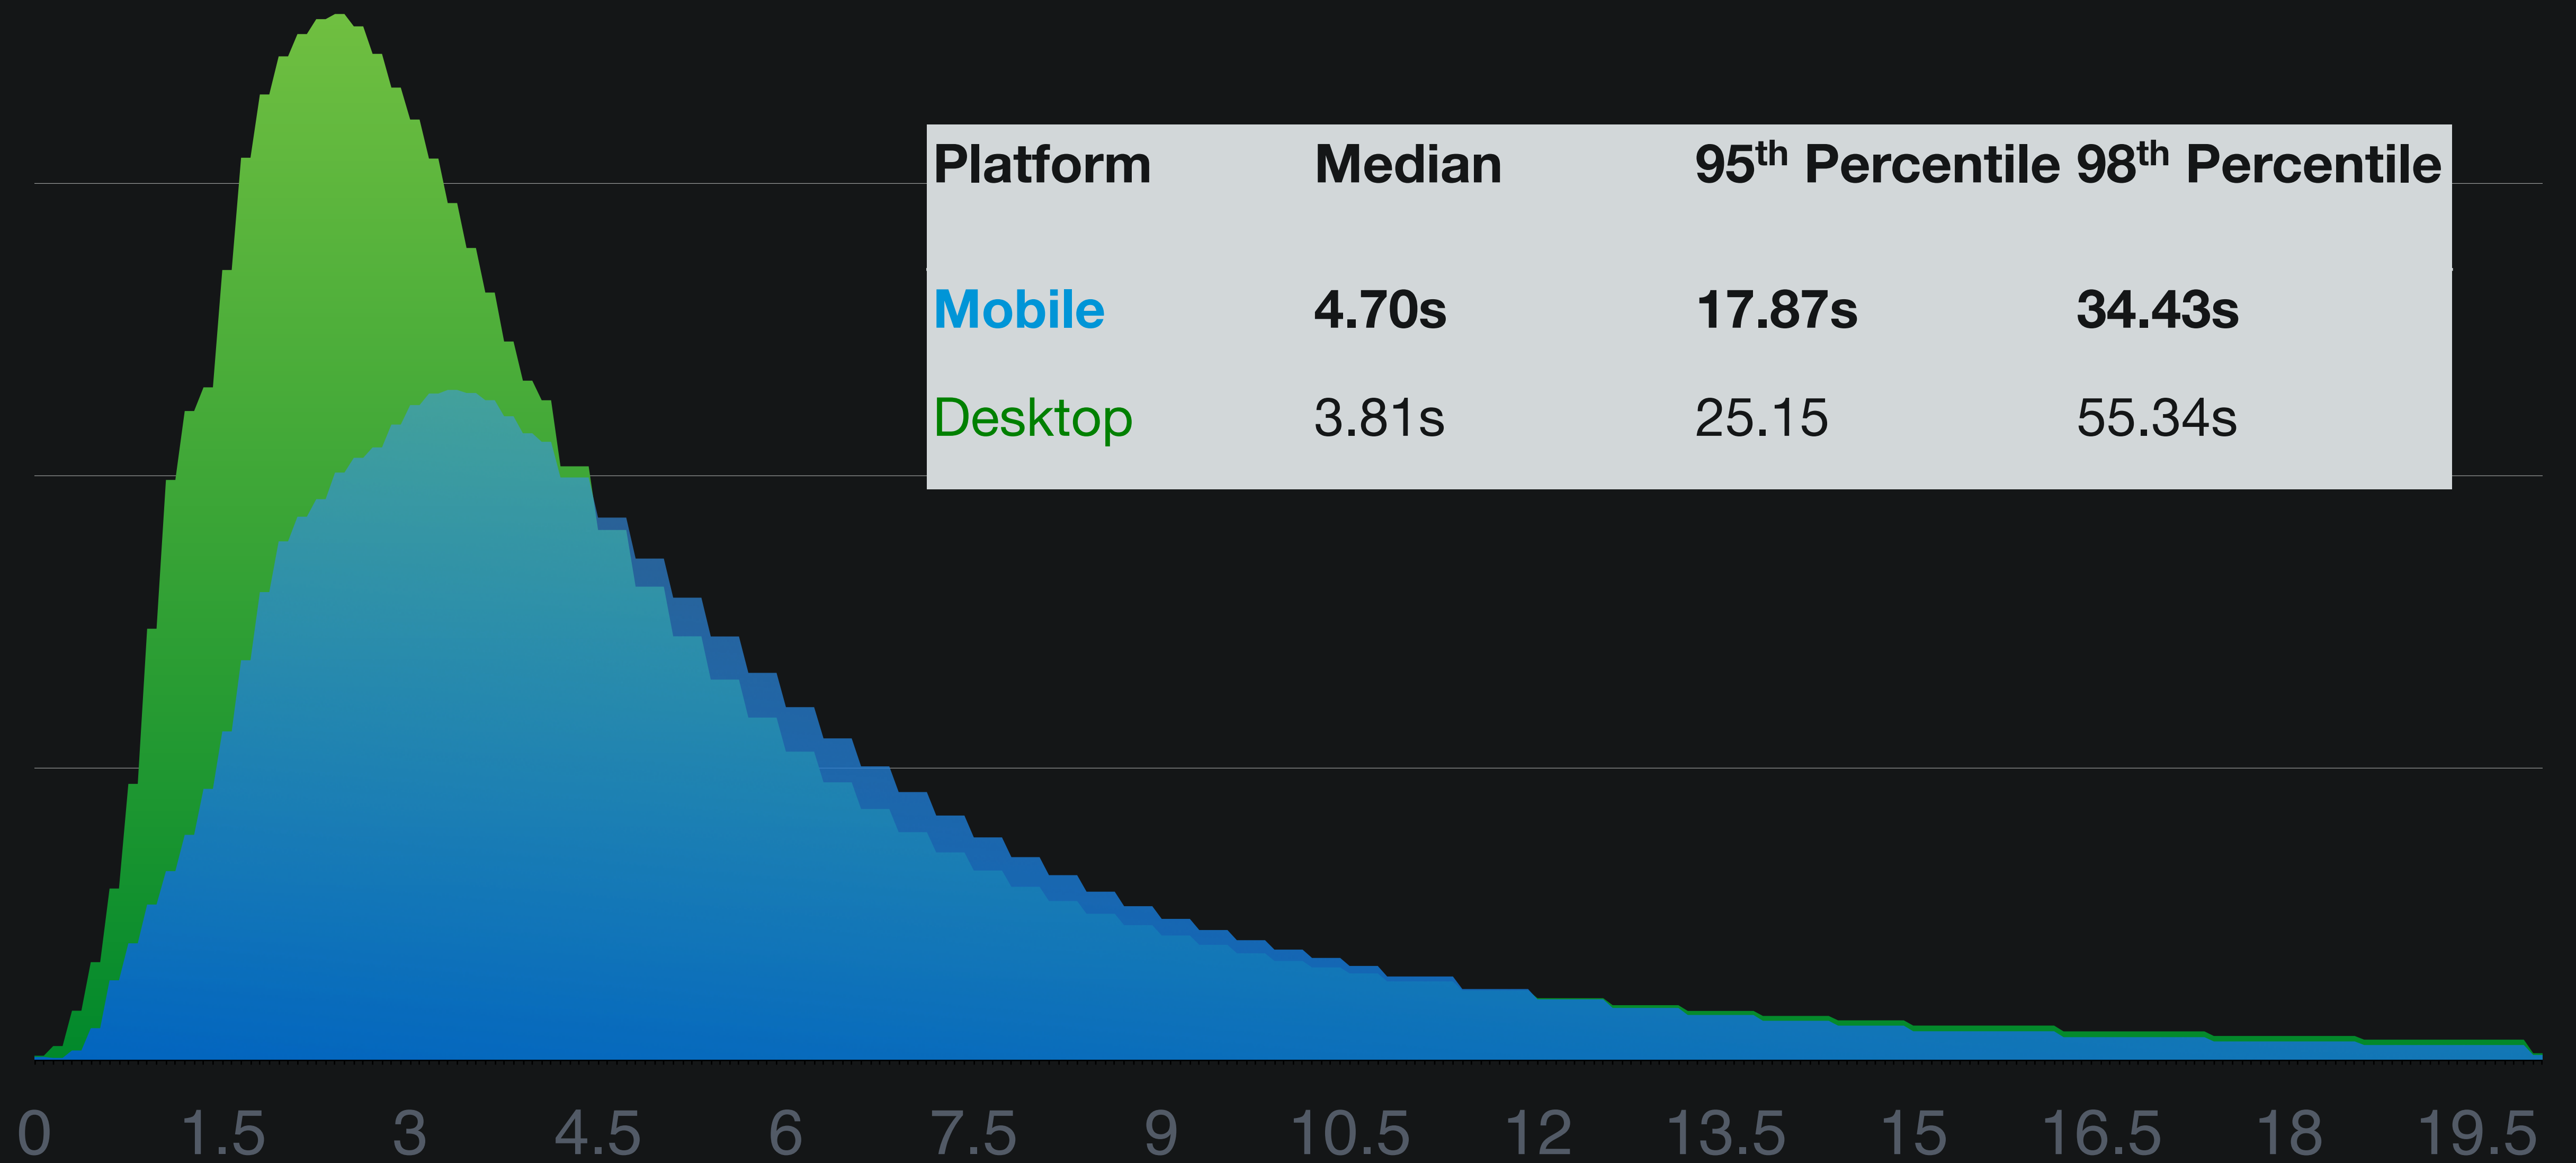
Page Load Times

Desktop vs. Mobile (Mobile and Tablet Optimized)



Page Load Times

Desktop vs. Mobile (Full site on Tablet)



Platform	Median	95 th Percentile	98 th Percentile
Mobile	4.70s	17.87s	34.43s
Desktop	3.81s	25.15	55.34s

3530

3530

IMG

JS


CSS


320

1024

1366

1920

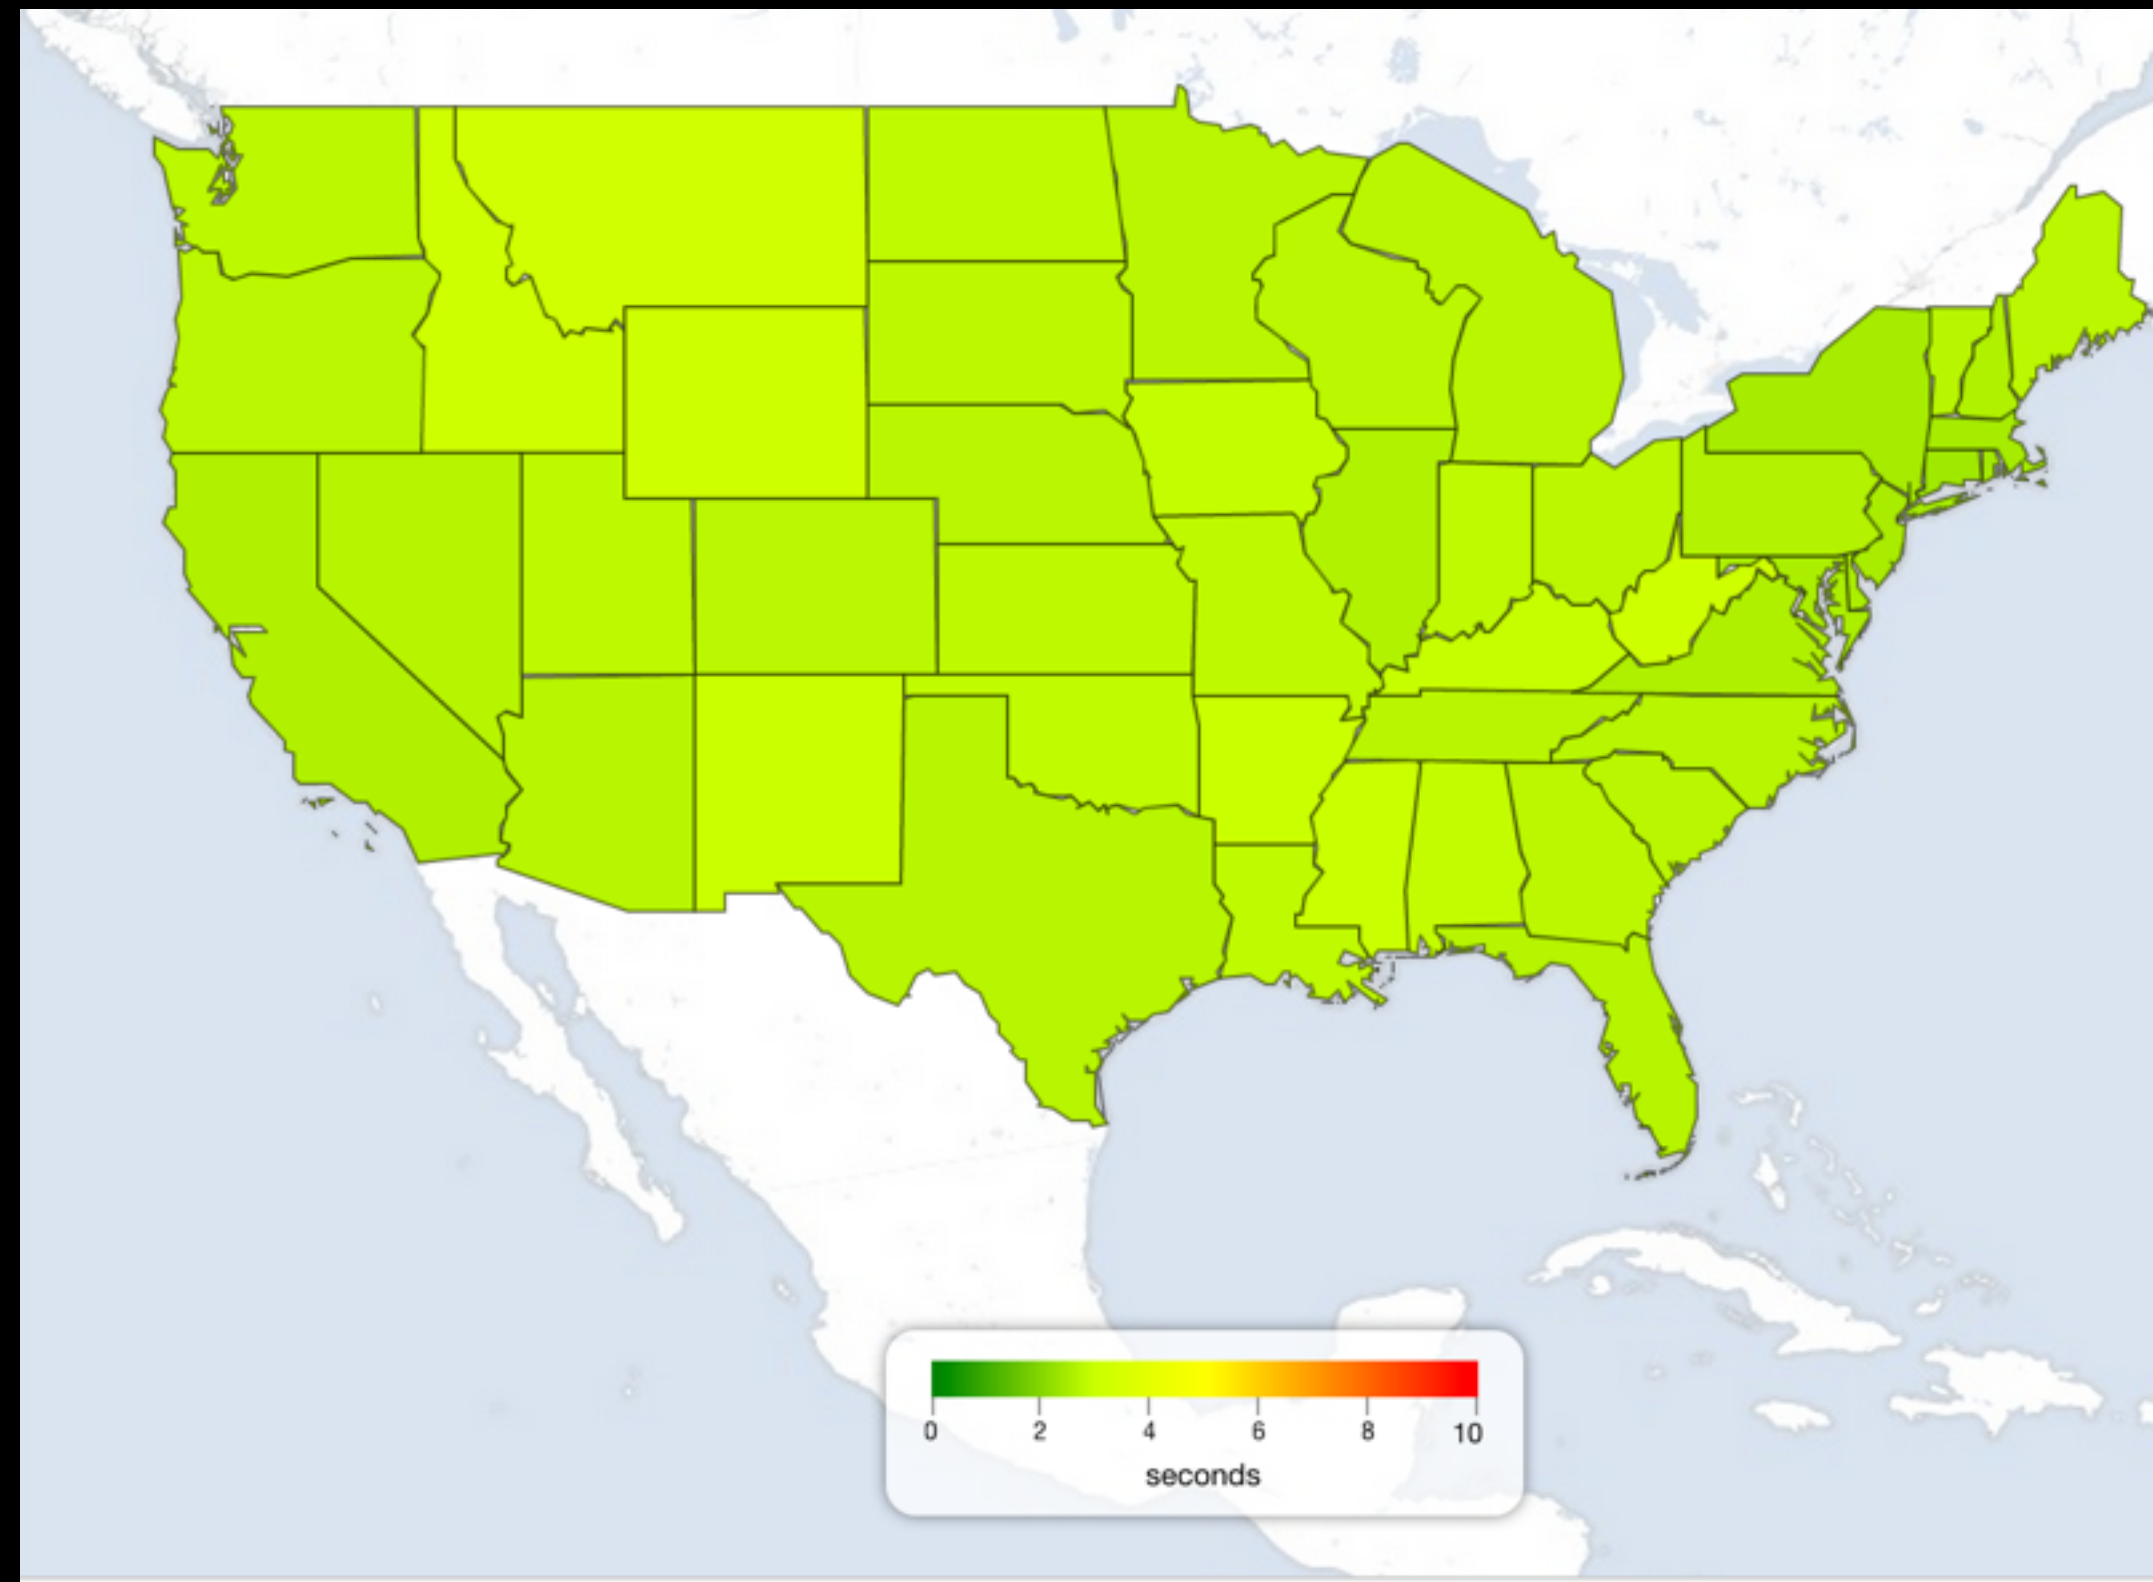




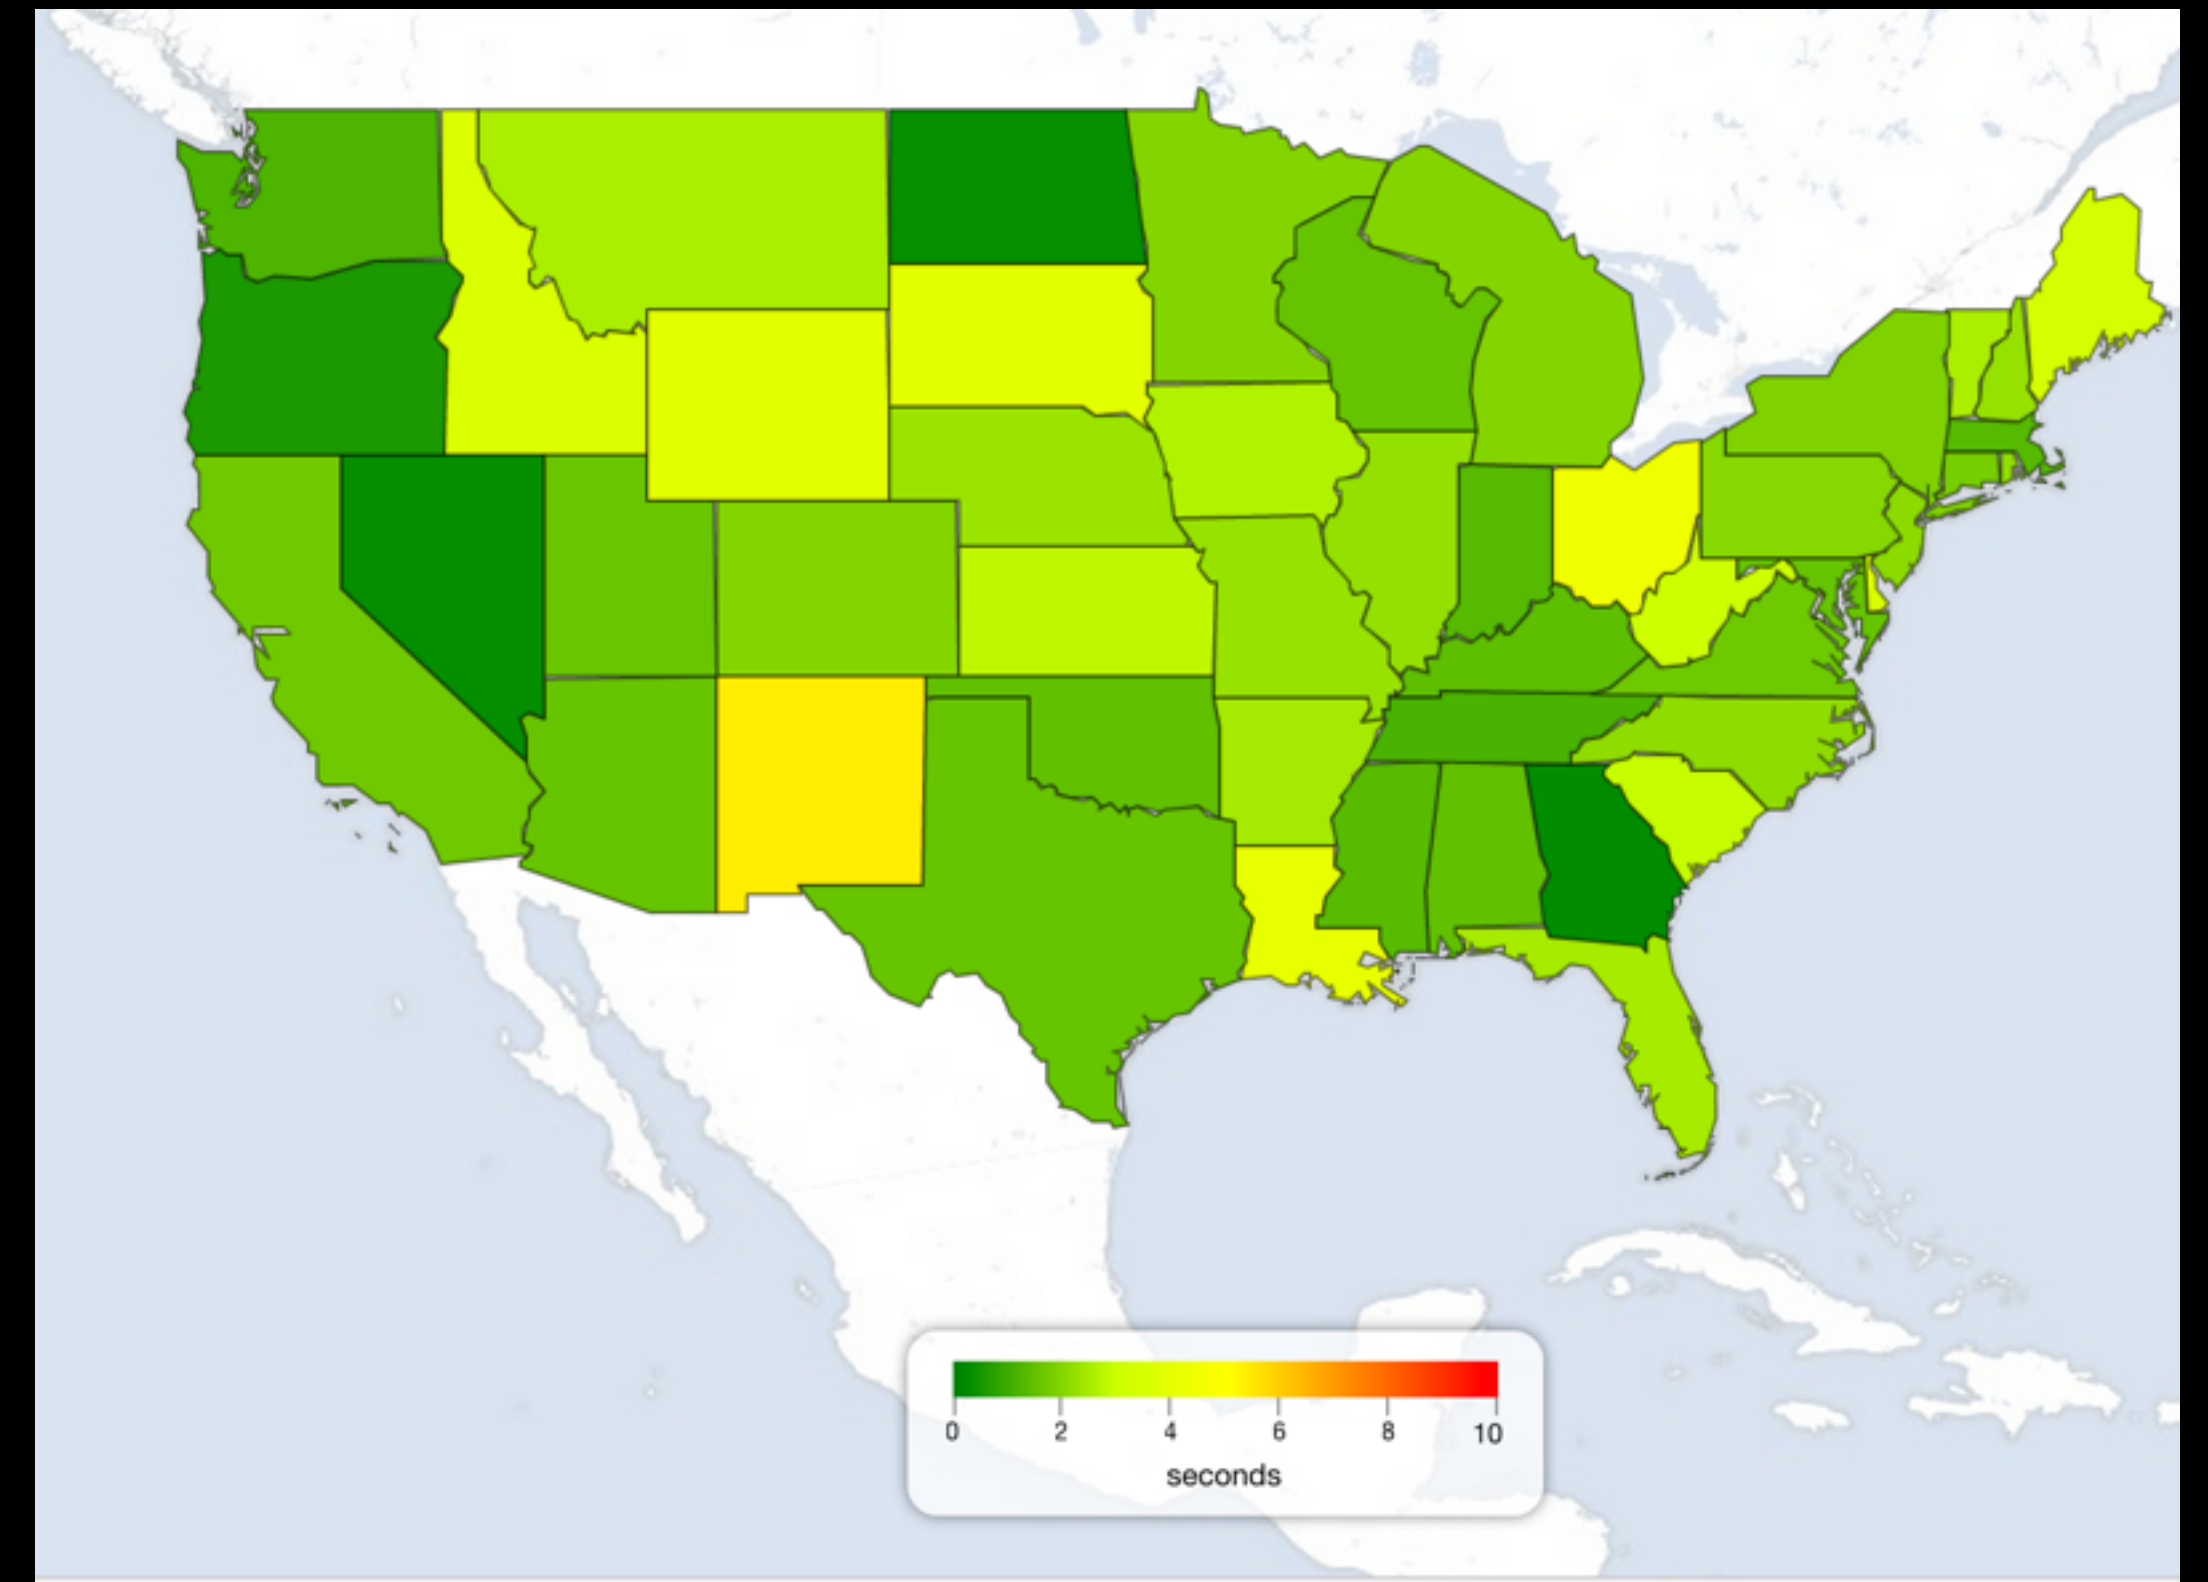
“How can I determine the effectiveness of my CDN provider?”

– OPERATIONS

WHICH SITE IS USING A CDN?



Site A





Site B

“How has page construction changed over the last month,
and what impact has that had on performance?”

– FRONT END DEVELOPER

CONTENT REQUESTS



“What is the performance of my single page web application?”

– FRONT END DEVELOPER

THE CHALLENGE OF SINGLE-PAGE APPLICATIONS

- Challenges with synthetic: Only measure first page load, not feasible to measure secondary loads via XHR
- Challenges with RUM: Determining whether call is a "page" or service call; Risk of diluting overall numbers


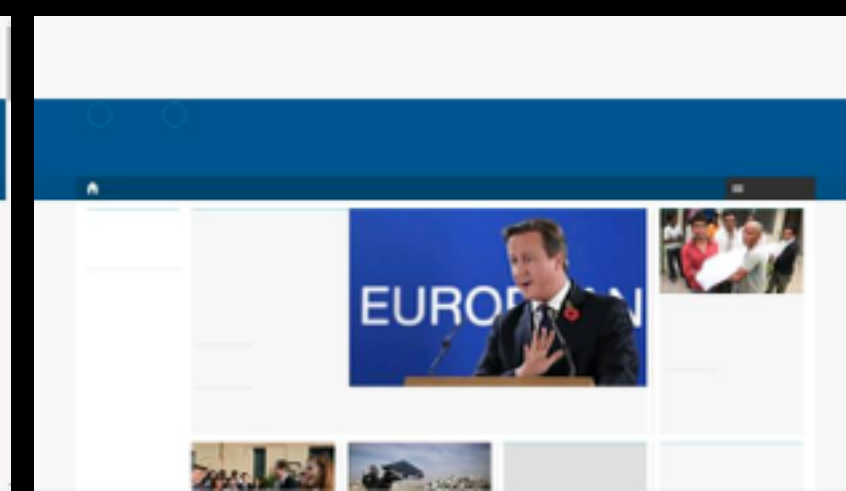
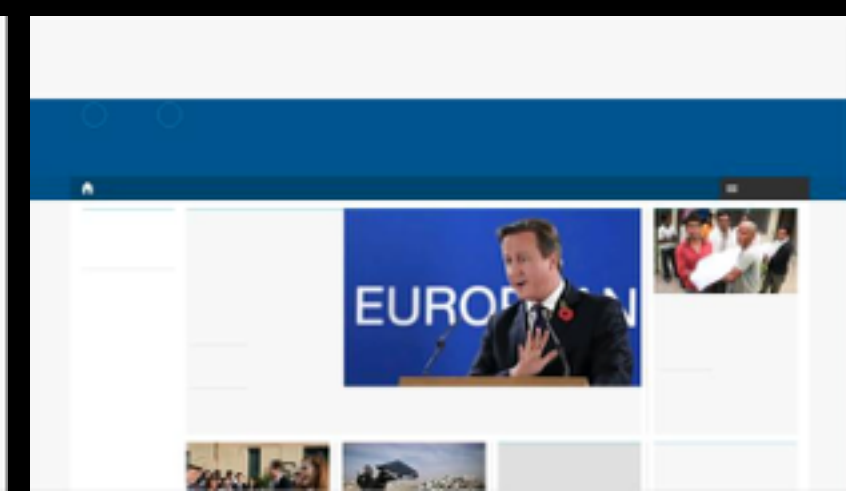

“How do I know how fast my page becomes usable?”

– DESIGNER

Start Render

Loaded

Usable



- Other questions?

SUMMARY

- You need rich sets of data
- There is not one number
- Choose the right number(s) for the right audience