THIRD-PARTY CONTENT

THE WEAK LINK IN YOUR CHAIN?

Simon Hearne - Principle Engineer @ Akamai

#perfmatters conf 2018

@SimonHearne

AKA

BAD PEOPLE DOING HORRIBLE THINGS TO GOOD SITES

@SimonHearne

THE MODERN WEB WORKFLOW 101

make something 🤓

test it 👍

ship it 😎

...

put tags on it 😳

@SimonHearne

AMP GOT SOMETHING RIGHT 🙄

timkadlec.com/remembers/2018-03-19-how-fast-is-amp-really/

@SimonHearne

WHAT I'VE LEARNED IN 5 YEARS

we *seem* to have less control than ever

@SimonHearne

THE "BUSINESS" 

TAGS SERVE BUSINESS GOALS

- Measurement & Analytics
- Ads & Retargeting
- "Optimization" & Testing
- Comments & Live Chat
- Tag Management

THE *MONEY* 

“We know that Optimizely slows down the site, but it will get us $750k increased revenue this year”

Holiday website, UK

THE *IT'S NOT MY JOB* 🤷‍♀️

"We suspect it slows the site down, we haven't tested it. Marketing says it's critical to their latest TV campaign so there's no point arguing"

Budget airline, UK

THE *TAG MANAGER* 

"All the tags go through the tag manager, so they should be fine."

Clicks-and-mortar store, UK

BUT WHAT ABOUT THE

RISK?

PS: AVAILABILITY HEURISTIC 101 🤔

PS: AVAILABILITY HEURISTIC 101 🤔

Remember when Facebook went down?
Remember when Disqus went down?
Remember when Maxymiser went down?
Remember when **Dyn** went down?
left-pad 😞

RISK 1:

MALICIOUS CODE INJECTION 

HOW MUCH OF YOUR CODE HAS VULNERABILITIES?

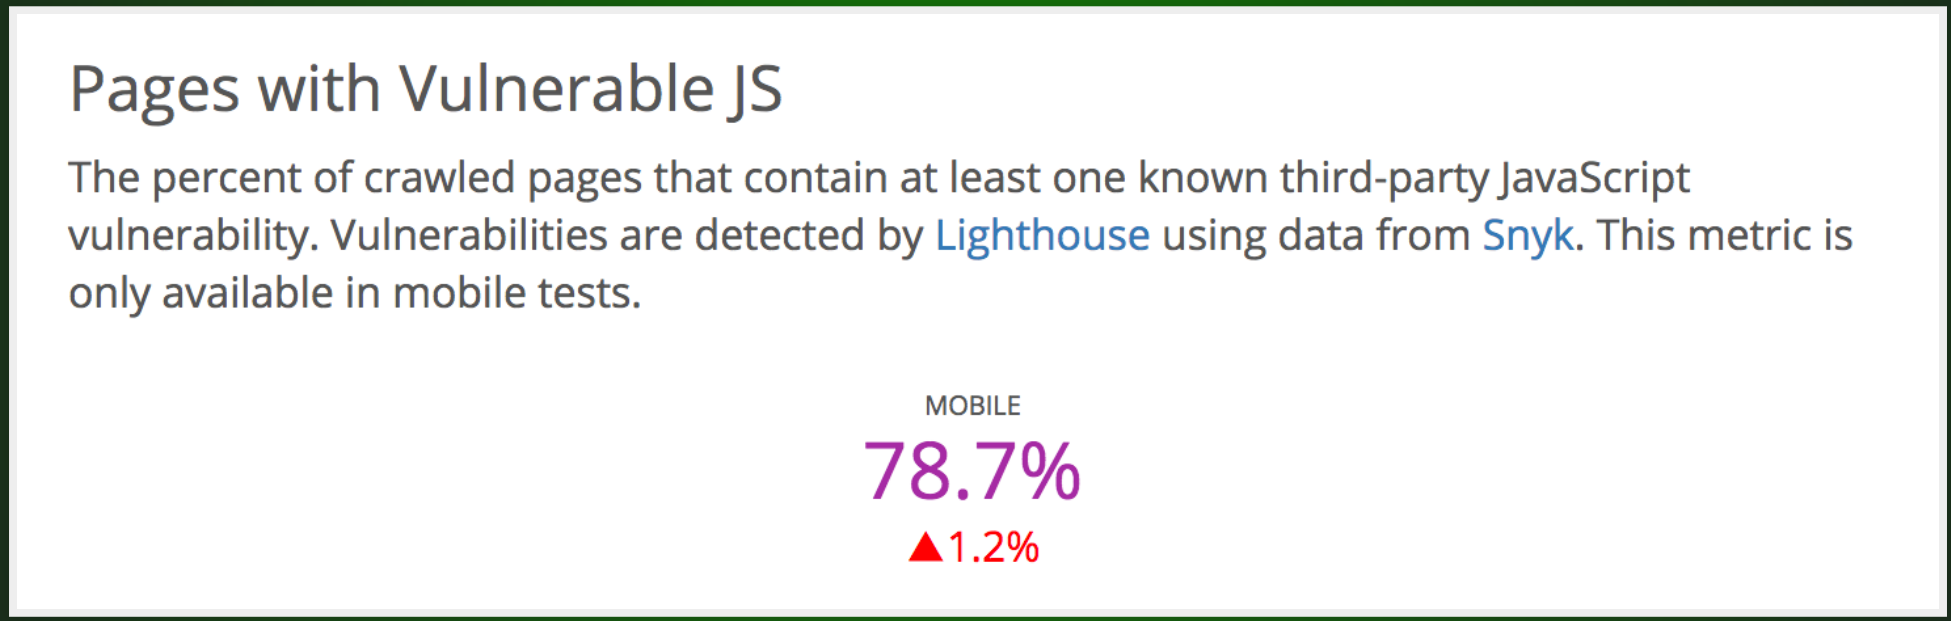

beta.httparchive.org/reports/state-of-the-web#pctVuln

CRYPTOJACKING

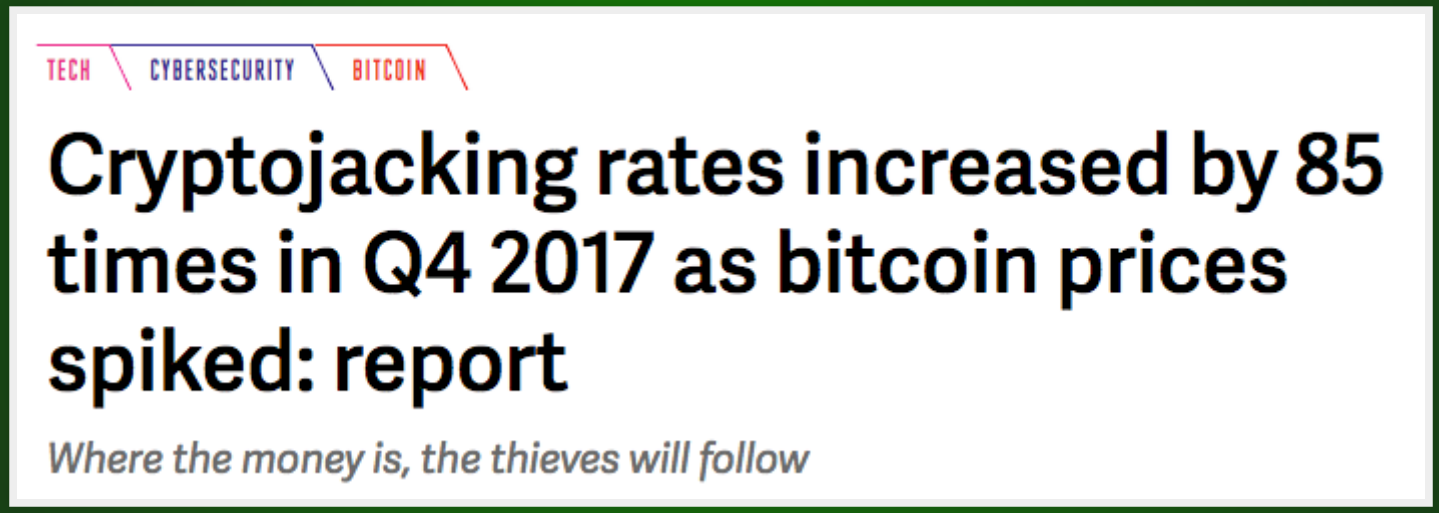

www.theverge.com/2018/3/22/17147320/cryptojacking-8500-percentage-points-bitcoin-monero-spike-symantec-security-mining

IT HAPPENS TO THE BIGGEST PLAYERS

twitter.com/nytimes/status/3958547840

INTERNET 'SERVICE' PROVIDERS

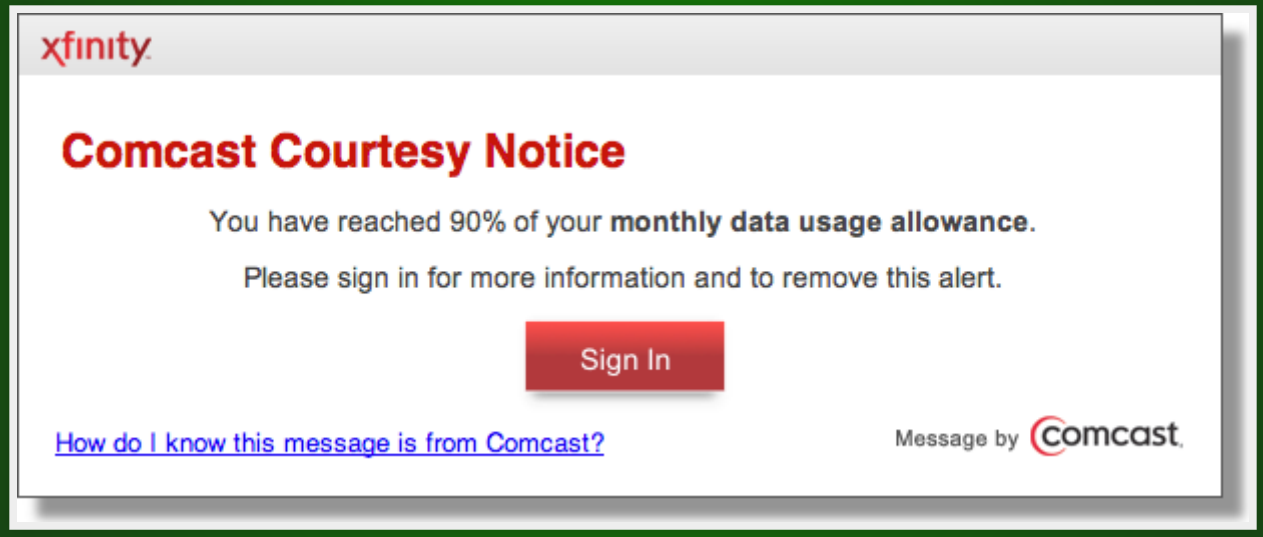

blog.ryankearney.com/2013/01/comcast-caught-intercepting-and-altering-your-web-traffic/

CONTENT DELIVERY NETWORKS

github.com/ampproject/amphtml/issues/2380

CONTENT DELIVERY NETWORKS

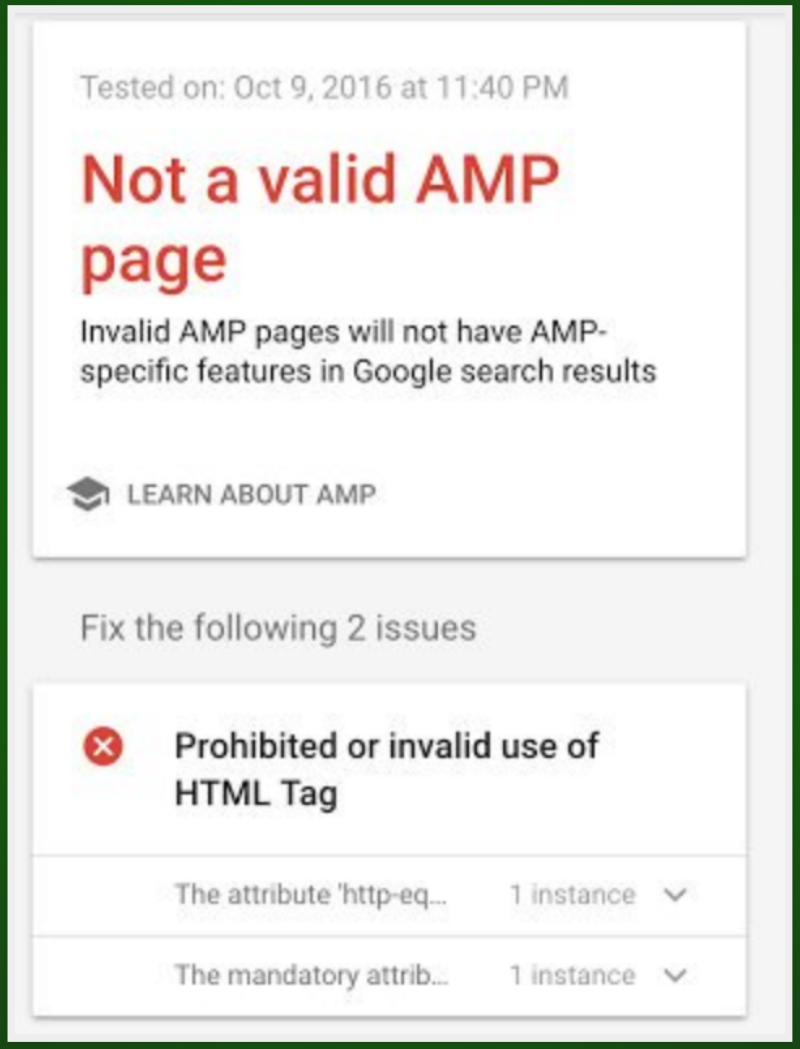

UNINTENTIONAL DATA COLLECTION

website tracking is a "security disaster waiting to happen"

freedom-to-tinker.com/2018/02/26/no-boundaries-for-credentials-password-leaks-to-mixpanel-and-session-replay-companies/

RISK 2:

AVAILABILITY 

DO THEY FAIL GRACEFULLY?

DO THEY FAIL GRACEFULLY?

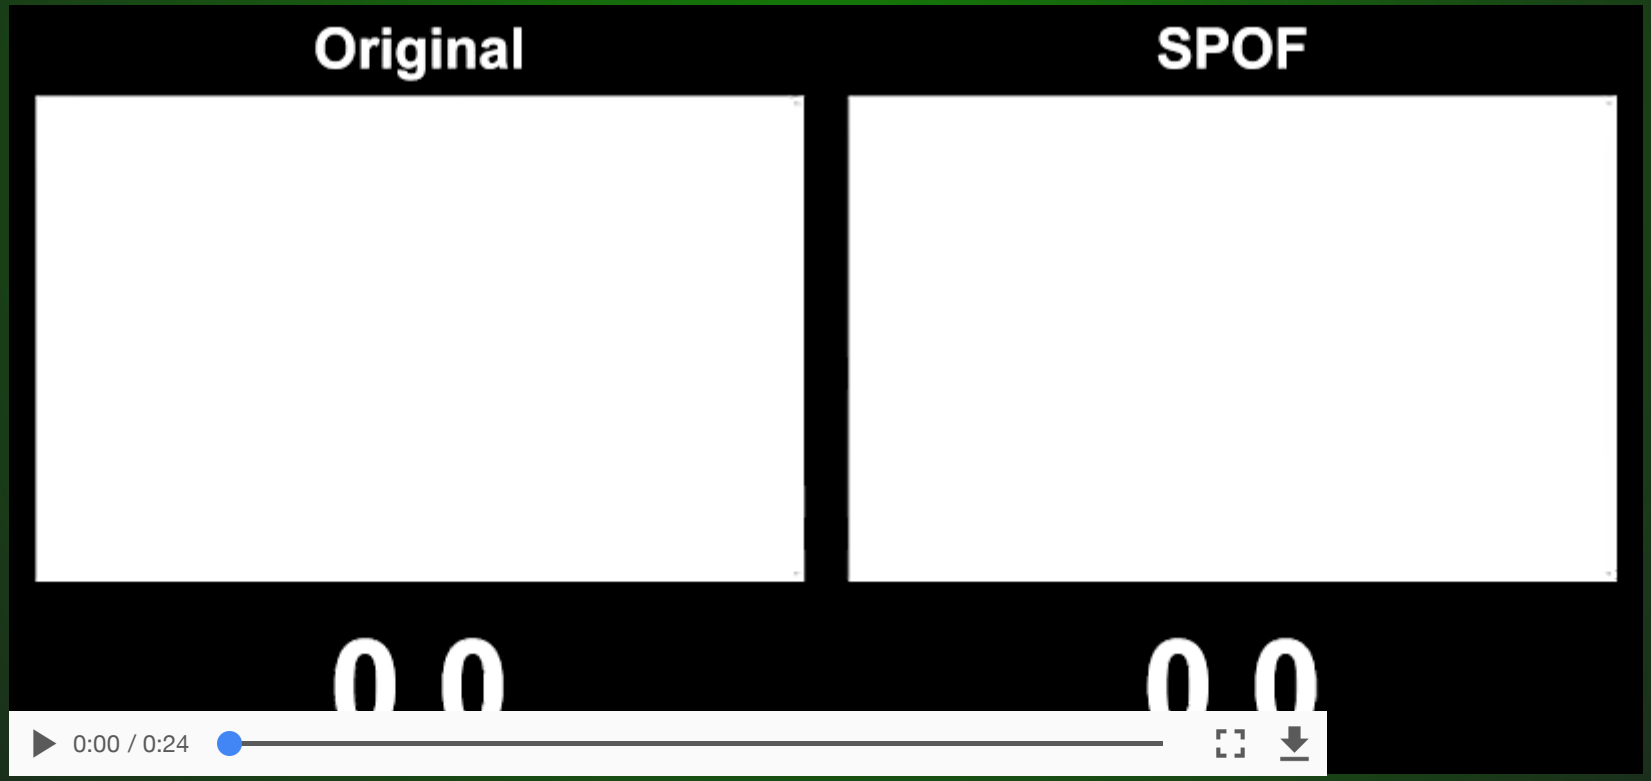

ARE THEY USING A CDN?

(& IS IT AS GOOD AS YOURS?)

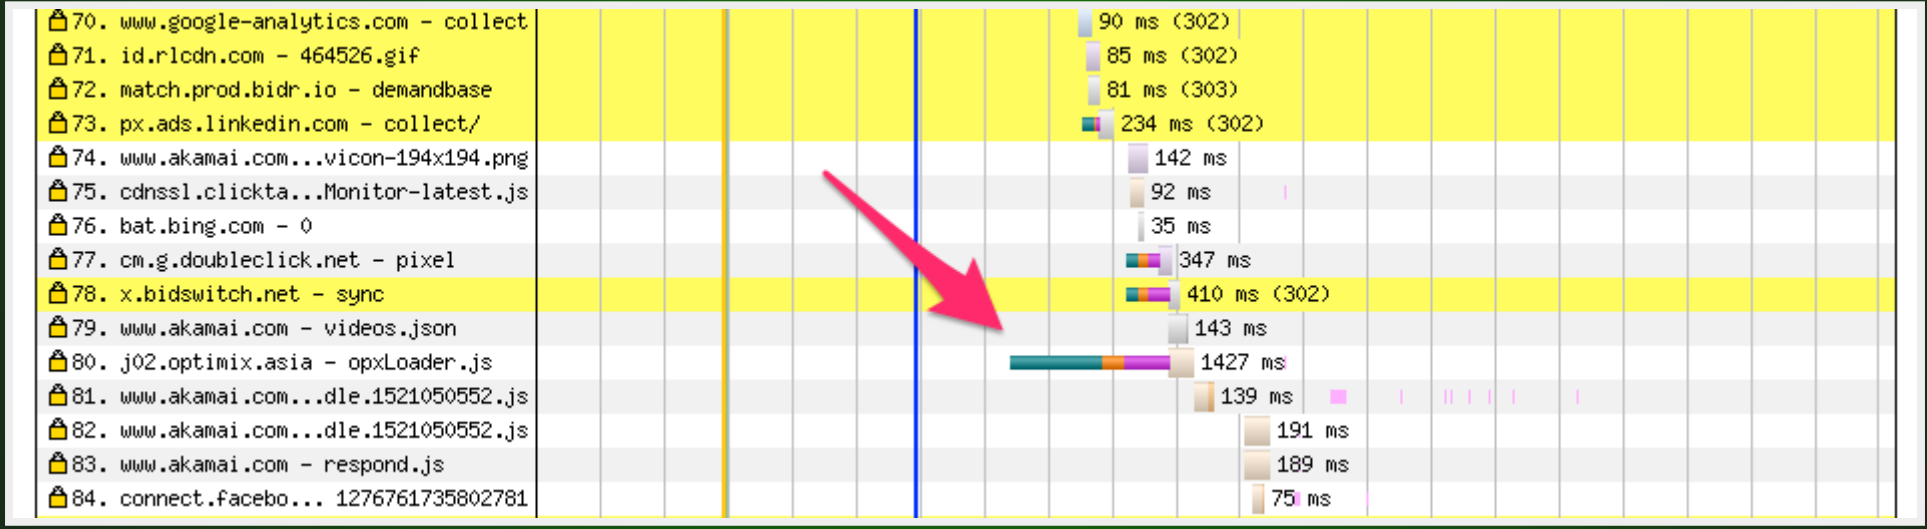

WHAT IS THEIR SLA FOR AVAILABILITY?

(& IS IT AS GOOD AS YOURS?)

Severity Level	Description	Standard Support First Response Time*	Priority Support First Response Time*
1 - Critical	**Customer website/application is unusable or unresponsive.** An Optimizely Service is causing a catastrophic problem to the Customer's production website or mobile application, such as a complete loss of availability. Customer is persistently unable to continue essential operations and no temporary workaround exists (e.g. pausing the experiment or campaign).	**1 business day** (during available working hours)	**1 hour** (24x7)

help.optimizely.com/Account_Settings/Optimizely_support_plans

AVAILABILITY TO THE USER

- Government / geo blocking
- Ad blocking
- Tracker blocking
- Random crap

RISK 3:

CODE QUALITY 

XSS VULNERABILITES

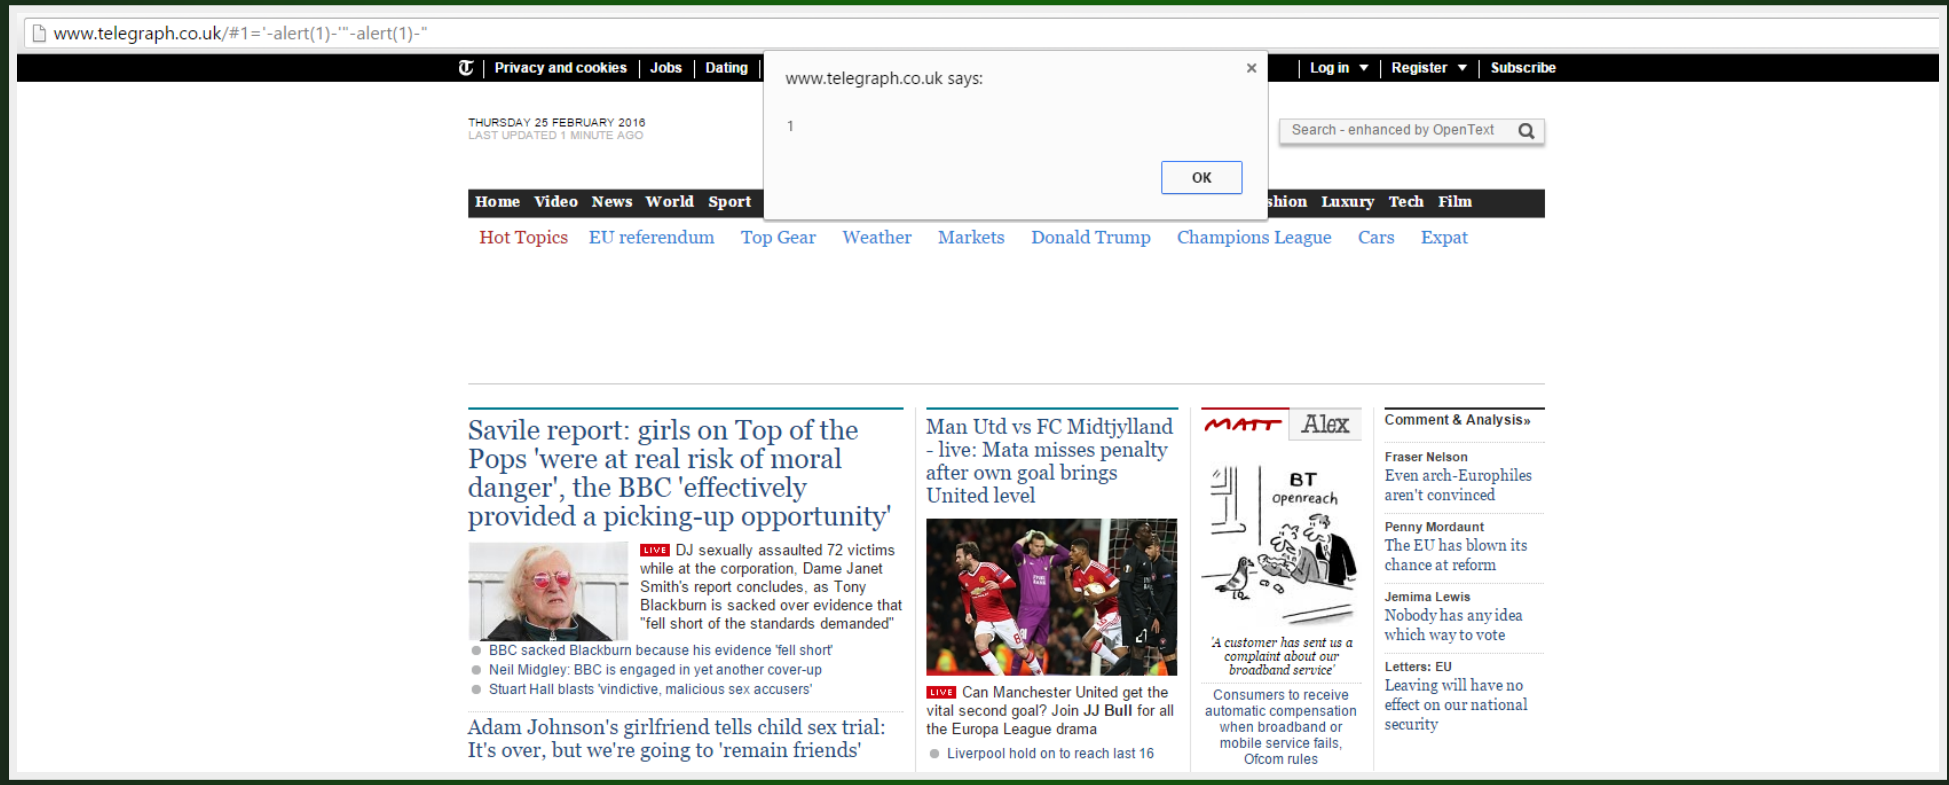

randywestergren.com/widespread-xss-vulnerabilities/

XSS VULNERABILITES

```
document.write("<div
        class='vdb_player vdb_565ec775e4b092ebc9685ce853180f5de4b066208a63279a'
        vdb_params='m.pub_id=606413&m.url=http://nypost.com/#1' - alert(1) - '-ale
        </div>");
```

randywestergren.com/widespread-xss-vulnerabilities/

DIFFERENT RELEASE SCHEDULES

```
<script src="//s7.addthis.com/addthis_widget.js" async></script>
```

HOW DO YOU KNOW WHEN IT CHANGES?

www.addthis.com

JUST PLAIN THOUGHTLESS

New Relic.

Do not clear the Resource Timing buffer

Support: Feature Ideas · Feature Ideas: Browser · javascript · bug

simonh — 1 · 3d

The Resource Timing API exposes resource-level performance metrics to the *window.performance* object. New Relic Browser consumes this data.

The specification sets a minimum recommended buffer size `https://www.w3.org/TR/resource-timing-2/#h-extensions-performance-interface` of 150. To avoid potential memory issues most browsers set the maximum buffer size to this value (150).

Many pages exceed 150 requests, especially on HTTP/2. This means that the buffer will reach its maximum value often, and stop recording resource timing data after the 150th resource.

The Browser script injected in to pages attempts to overcome this limitation by attaching an event listener to the *resourcetimingbufferfull* event, which processes the current buffer and then clears it in order to continue collection. While this is a good solution for New Relic (continued resource capture past the 150 limit), it means that any other tools that process Resource Timing data lose access to the prior 150 Resource Timing entries.

RISK 4:

PERFORMANCE 

SELF-POLICING ISN'T GOOD ENOUGH

*... the X Web Reference Snippet was available ... and the download time over **HTTP** did not exceed **500 ms**.*

*Snippet is sampled every minute from a variety of **U.S.** locations.*

status.optimizely.com

TOOLS AREN'T EQUAL

Top 9 Fastest AB testing tools

Ranked by time to complete test variation

Rank	Tool	Technology	CDN	Snippet type	Snippet placement	Variation complete (s)
1	**SiteSpect**	Server-side	n/a	n/a	n/a	1.33
2	**Maxymiser**	Client-side	Akamai	Synchronous	<head>	1.70
3	**Convert**	Client-side	Akamai	Synchronous	<head>	1.81
4	**Optimizely**	Client-side	Akamai/Edgecast	Synchronous	<head>	1.86
5	**Optimost**	Client-side	Akamai	Synchronous	<head>	1.93
6	**Marketizator**	Client-side	Cloudfront	Synchronous	<head>	2.13
7	**Qubit**	Client-side	Cloudfront	Synchronous	<head>	2.73
8	**AB Tasty**	Client-side	Amazon/MaxCDN /Cloudfront	Synchronous	<head>	2.88
9	**VWO**	Client-side	Private CDN	Asynchronous	<head>	4.29

orangevalley.nl/en/blog/9-ab-testing-tools-compared-on-site-speed-impact/

THE WEB IS VARIABLE

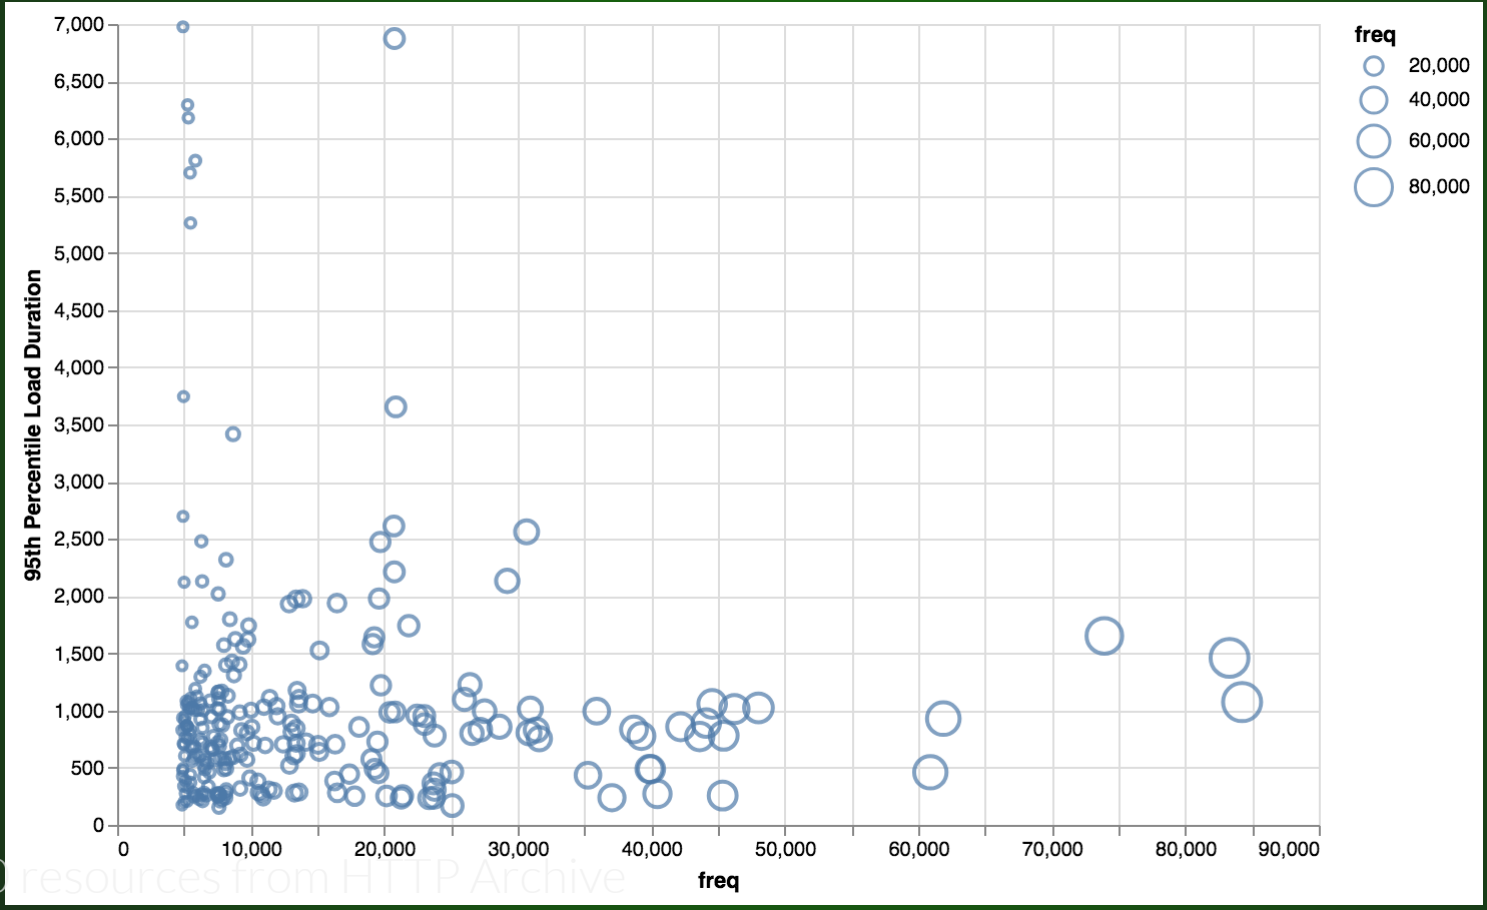

Top 250 resources from HTTP Archive

THE WEB IS VARIABLE

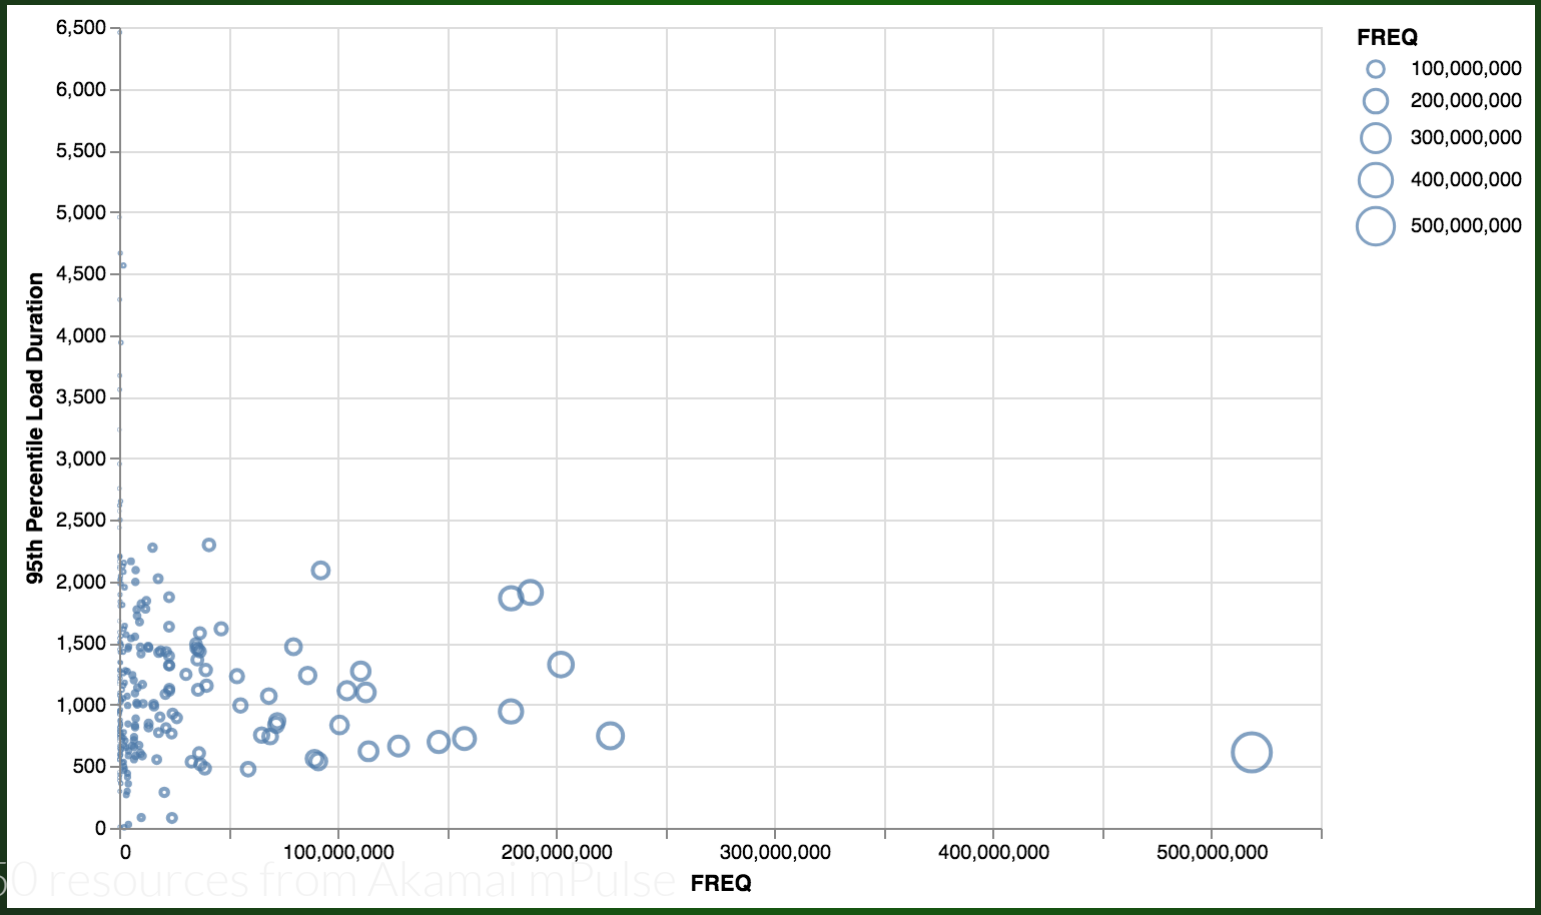

Top 250 resources from Akamai mPulse

RESOURCE TIMING IS THE HERO WE NEED

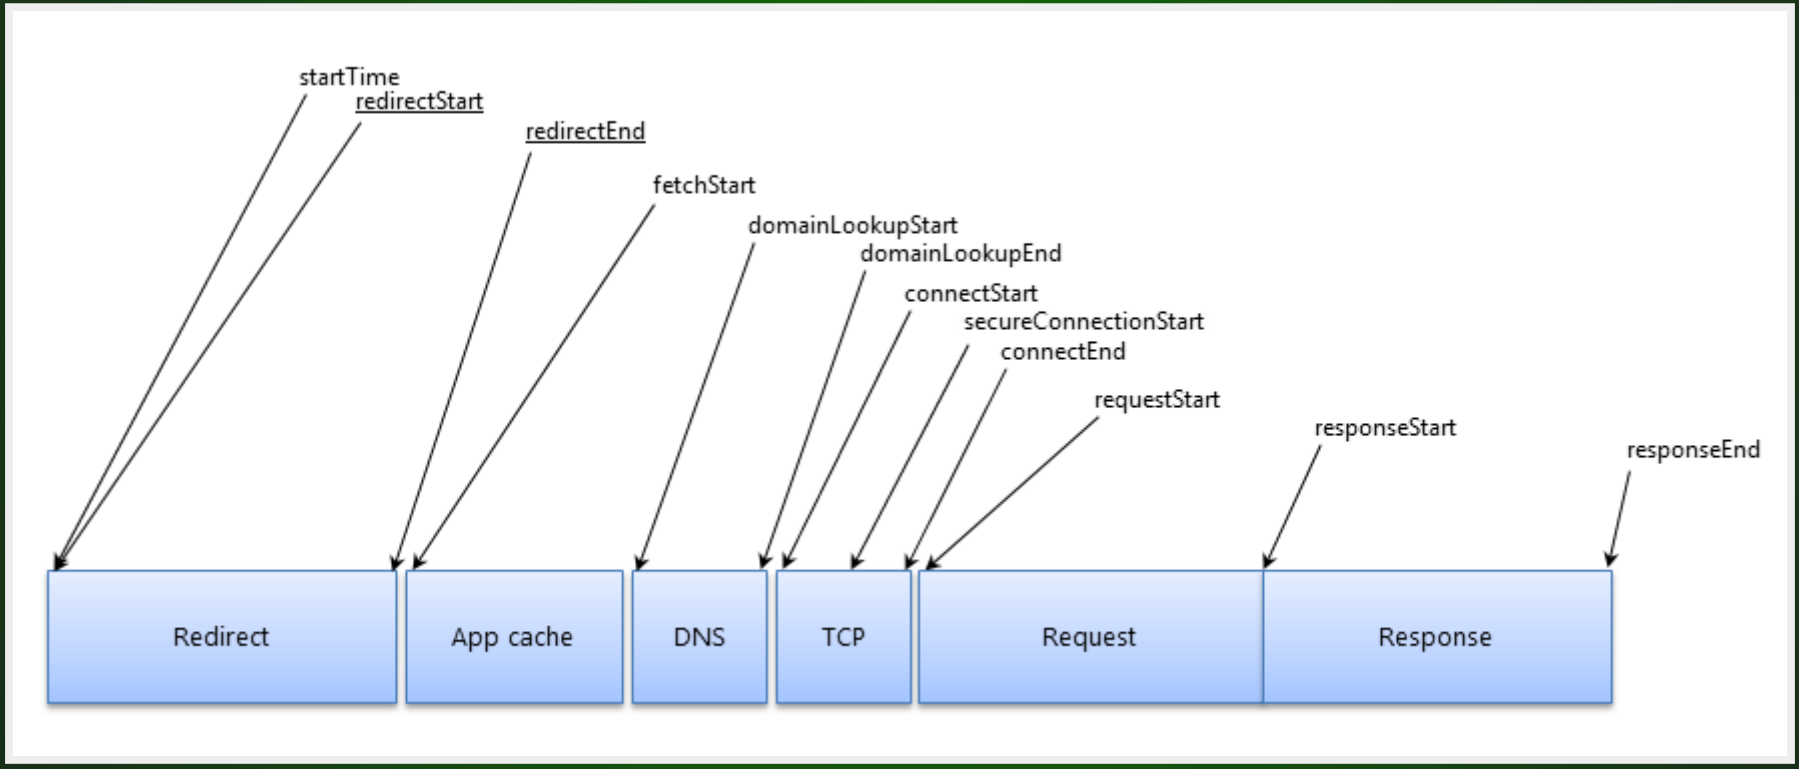

https://www.w3.org/TR/resource-timing-1/

RESOURCE TIMING IS THE HERO WE NEED

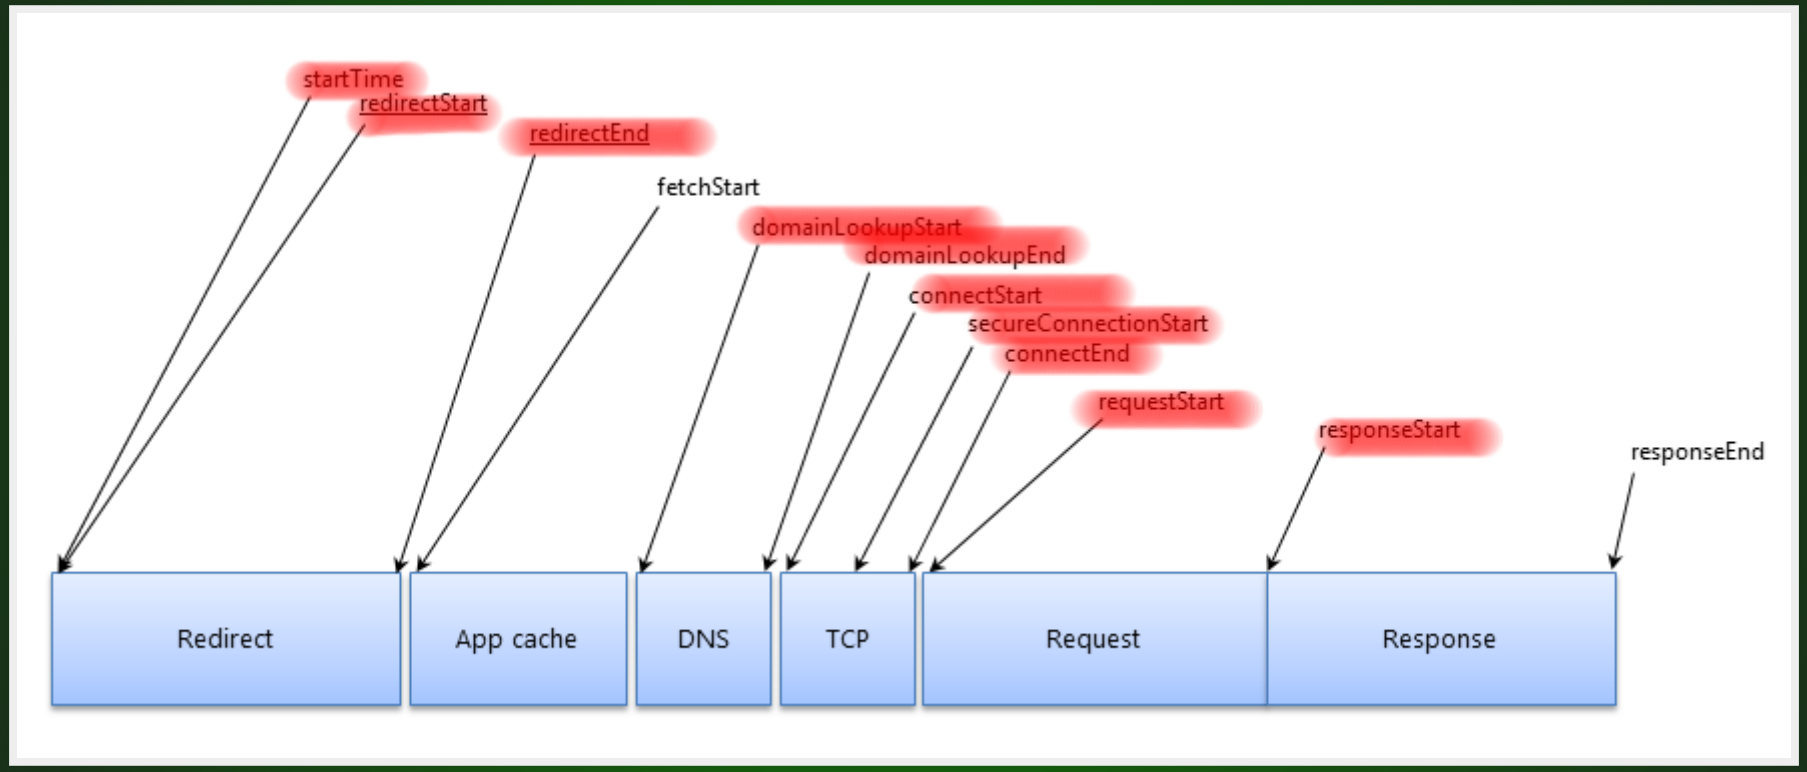

NOT WITHOUT TIMING-ALLOW-ORIGIN 🙈

nicj.net/resourcetiming-visibility-third-party-scripts-ads-and-page-weight/

RESOURCE TIMING WON'T SAVE US

- 🙈 no redirect information
- 🙈 limited data on **72%** of third-party content
- 🙈 only the first 150 requests *
- 🙈 no data on old iDevices
- 🙈 no data for cross-origin iframes

developer.akamai.com/blog/2017/07/26/measuring-performance-third-party-contributors/

* limit can be increased per pageview

RISK 4:

PERFORMANCE 

(FOR REAL THIS TIME)

CPU IS OUR BIGGEST BOTTLENECK ⏳

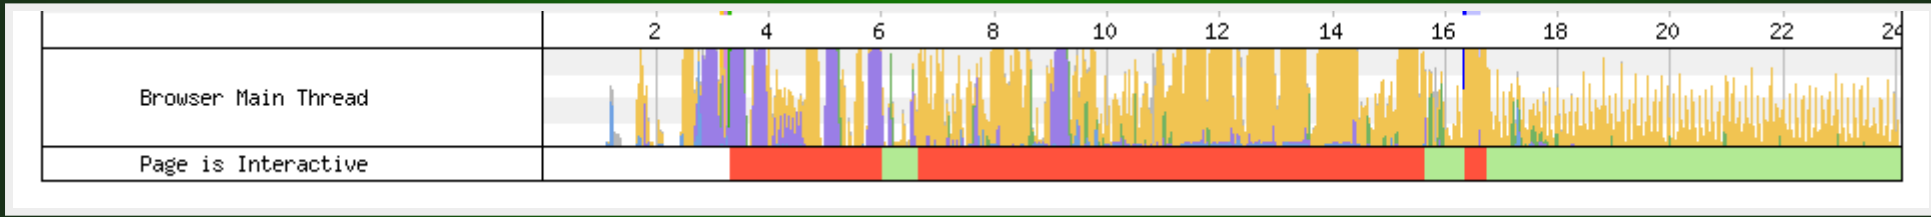

WHO'S POLICING THE THIRD-PARTIES?

Analyzing samples/VA-article.json

Total CPU busy time (ms)	10107.07
Total number of domains	57
Number of big offenders	40

CPU Time (ms)	domain name
1311.39	www.vogue.fr
865.91	z.moatads.com
835.78	pagead2.googlesyndication.com
788.86	securepubads.g.doubleclick.net
.....	.....
565.78	data05.adlooxtracking.com
559.52	tpc.googlesyndication.com

github.com/jpvincent/3rd-party-cpu-abuser

VARIABLE CPU TIME

ON SYNTHETIC AGENTS!

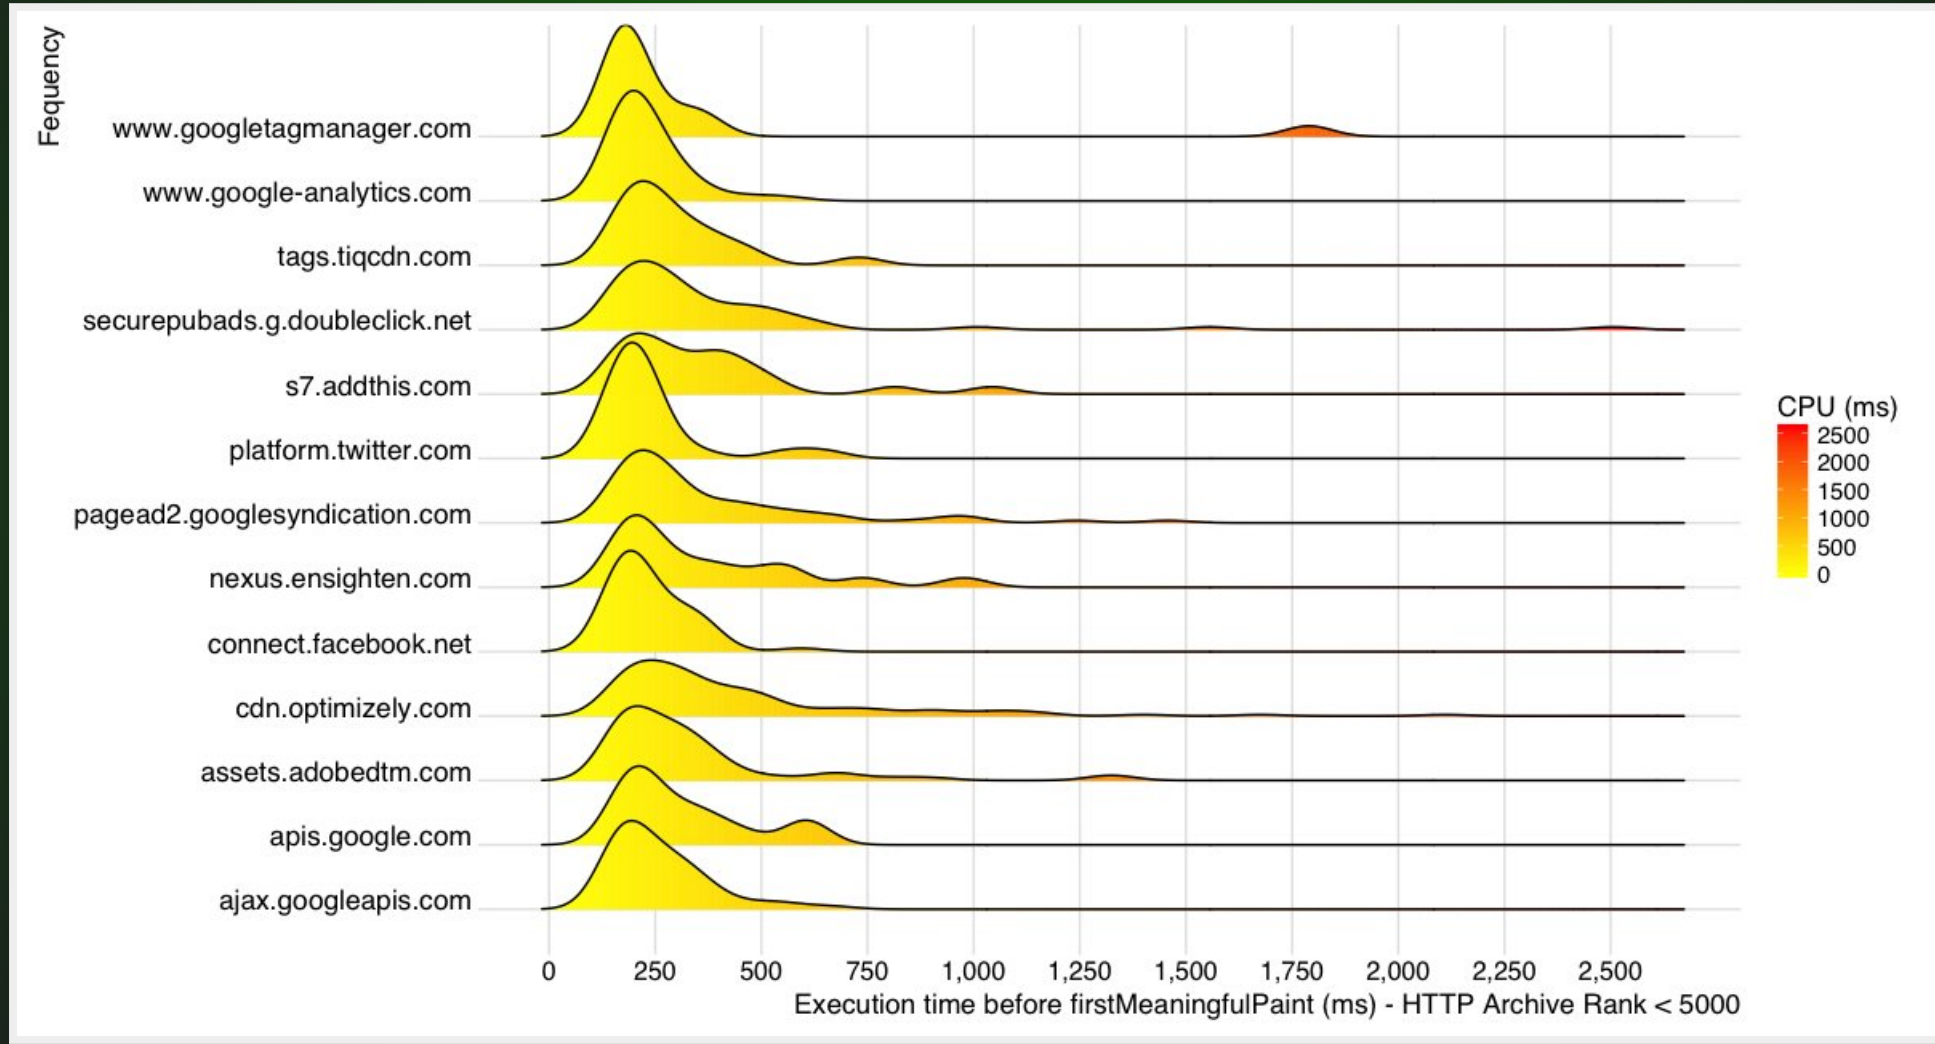

THE MOST FRUSTRATING PERF BUG, EVER

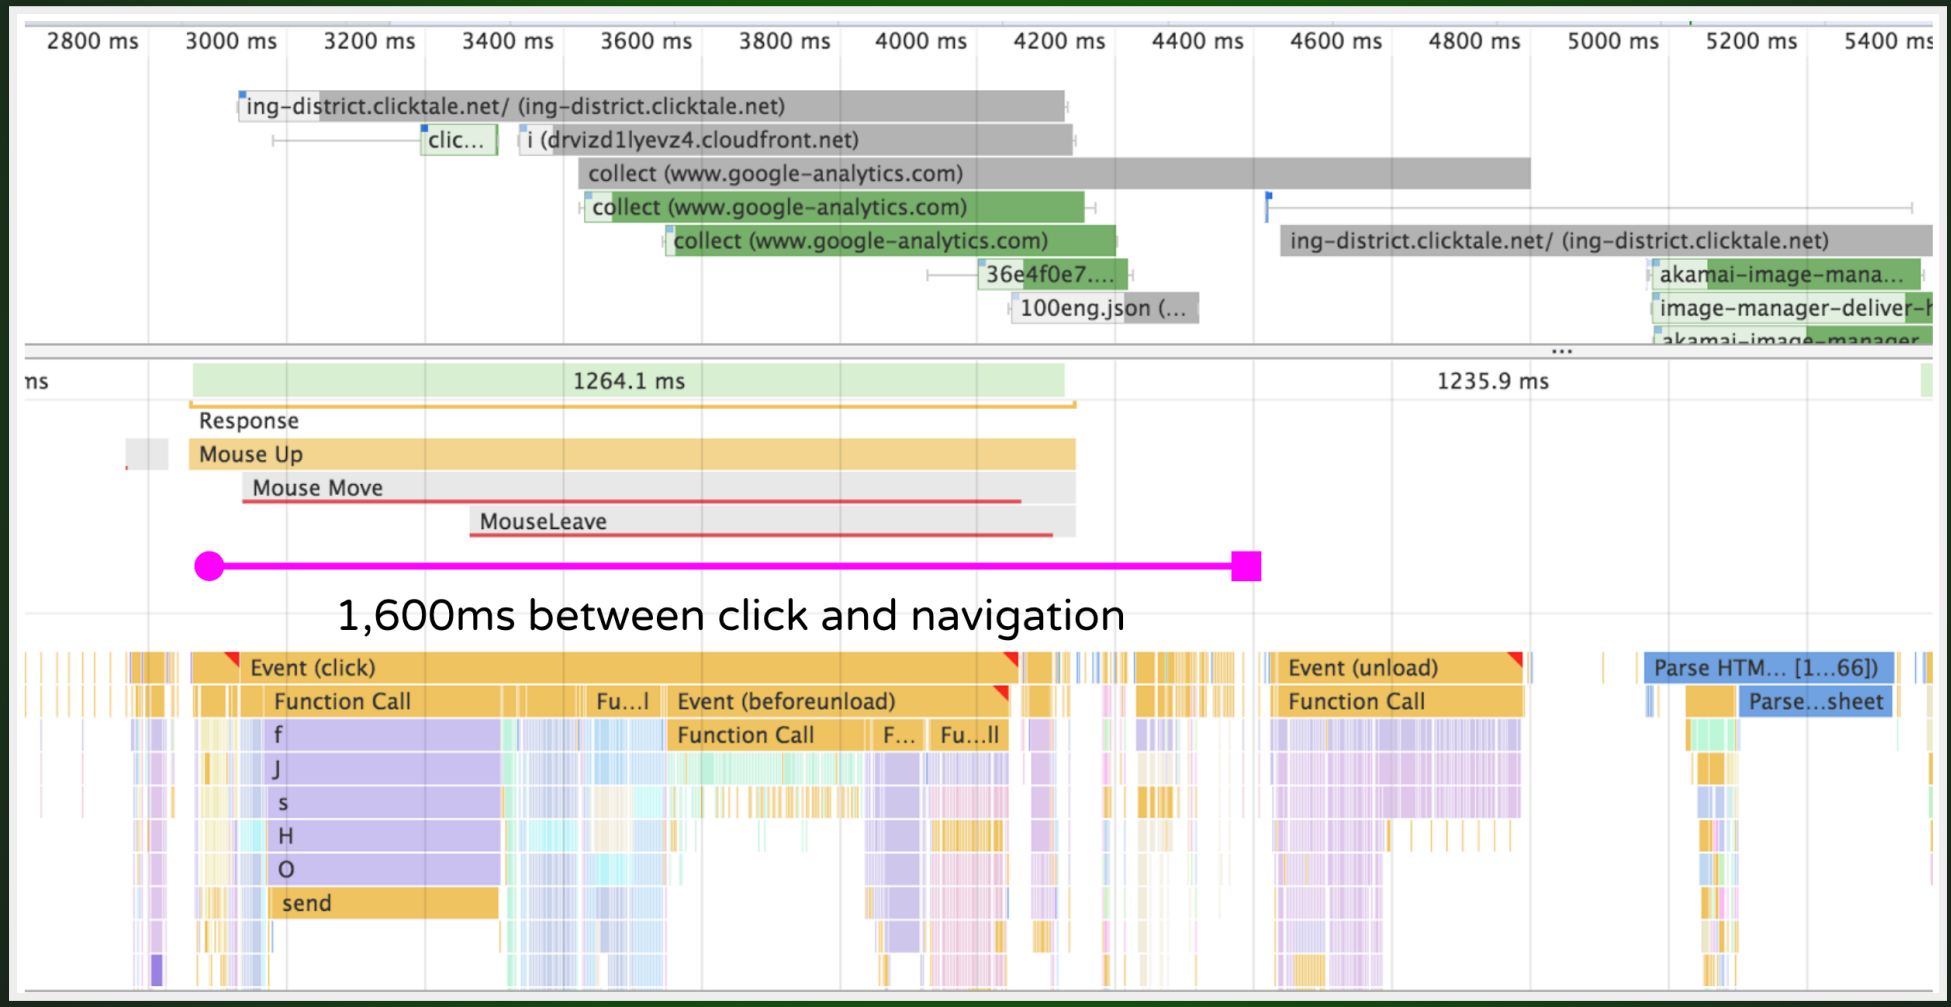

WHO WATCHES THE WATCHMEN?

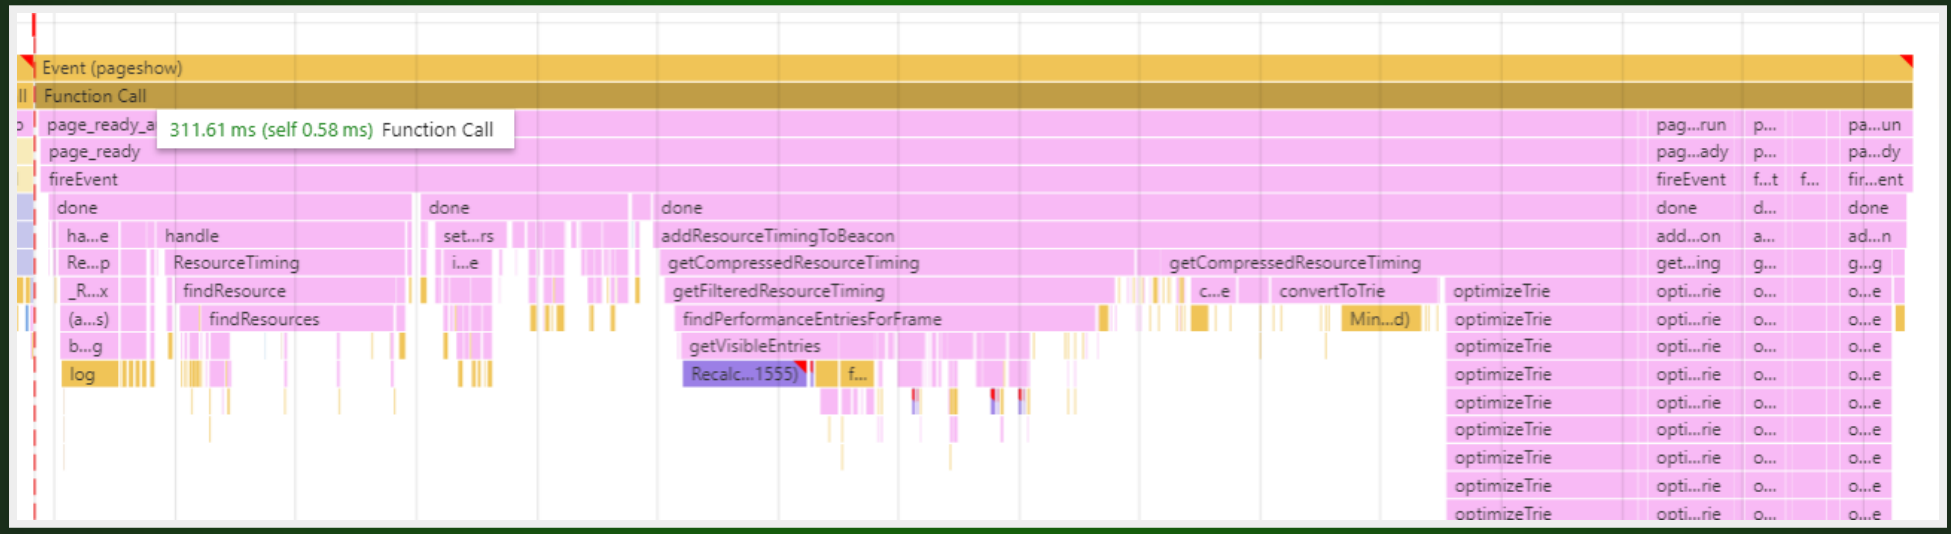

calendar.perfplanet.com/2017/an-audit-of-boomerangs-performance/

RISK 4:

PERFORMANCE 

(PART III)

DELAYING ONLOAD

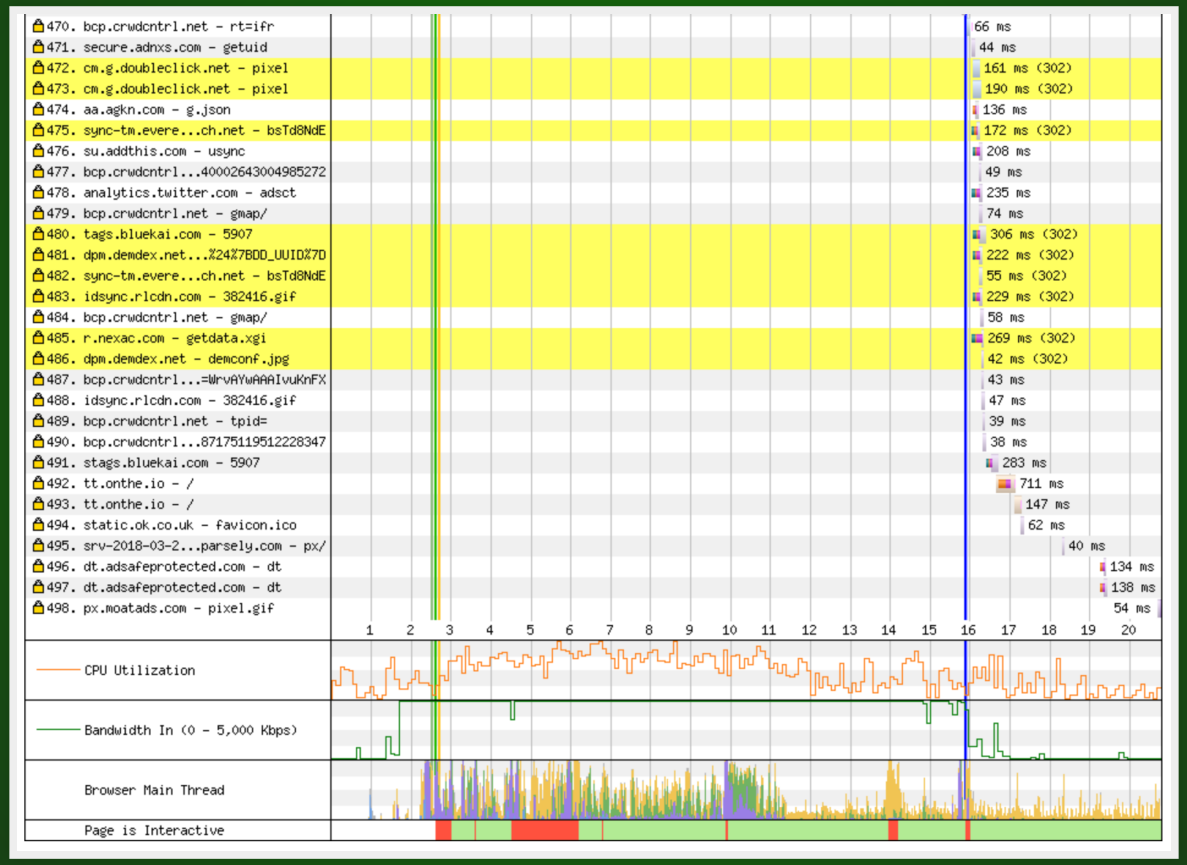

KEEPING THE RADIO AWAKE

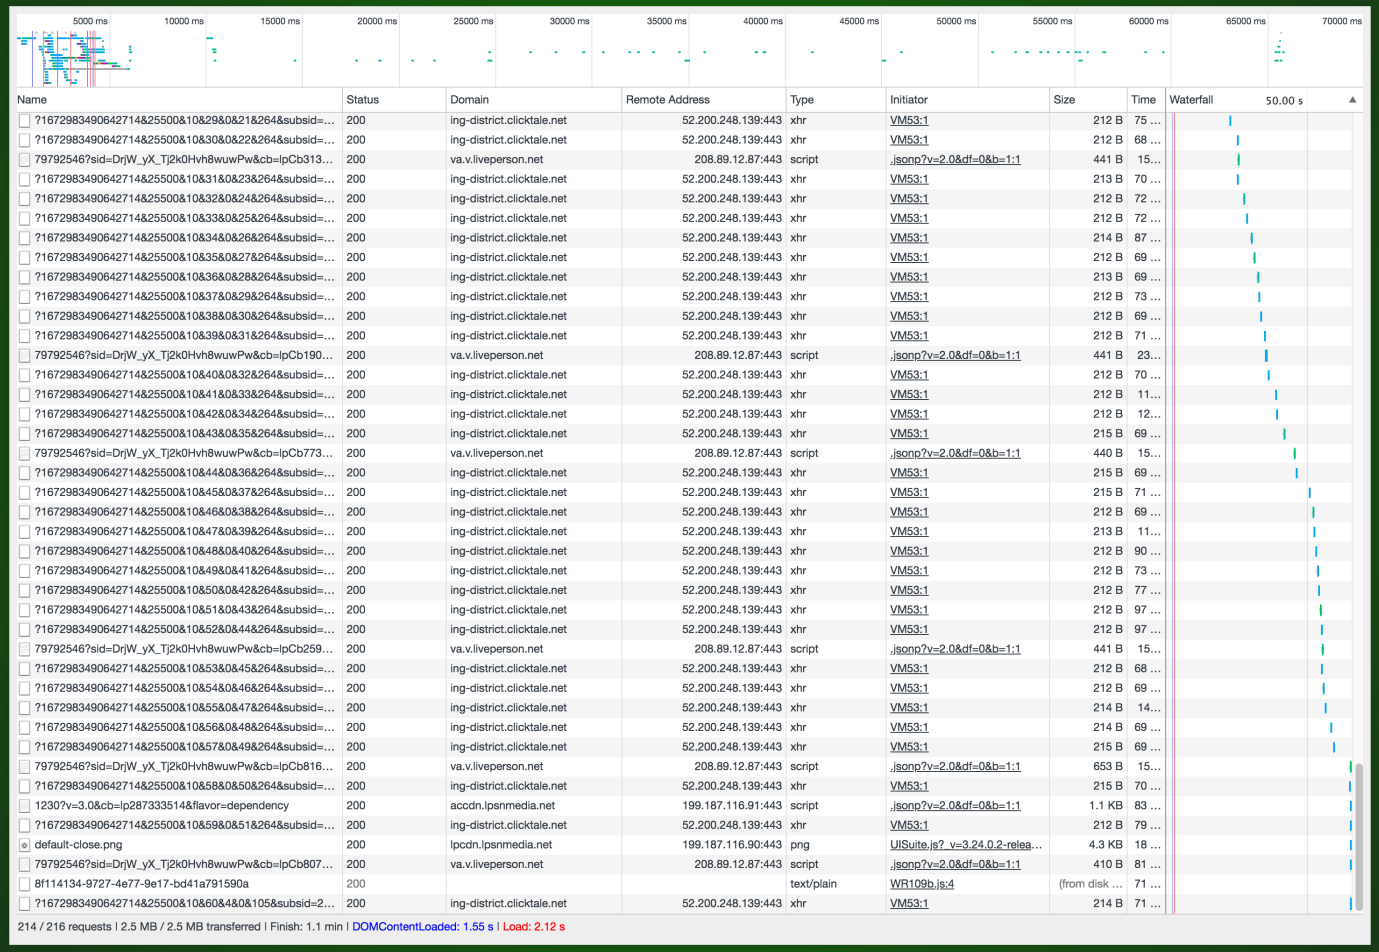

WE HAVE LITTLE CONTROL OVER WHICH ARE USED

But there are things we can do...

@SimonHearne

STAGE 1:
FIND OUT WHAT'S THERE

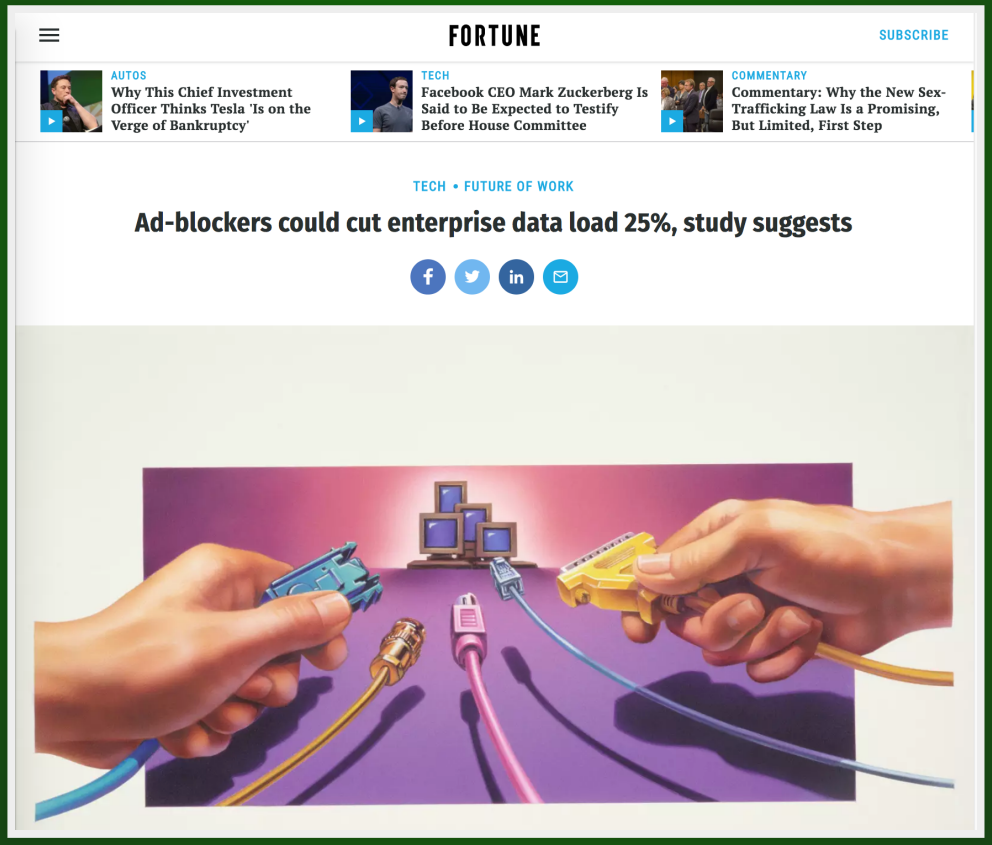

fortune.com/2015/07/08/adblock-data/

SYNTHETIC TESTING (WEBPAGETEST)

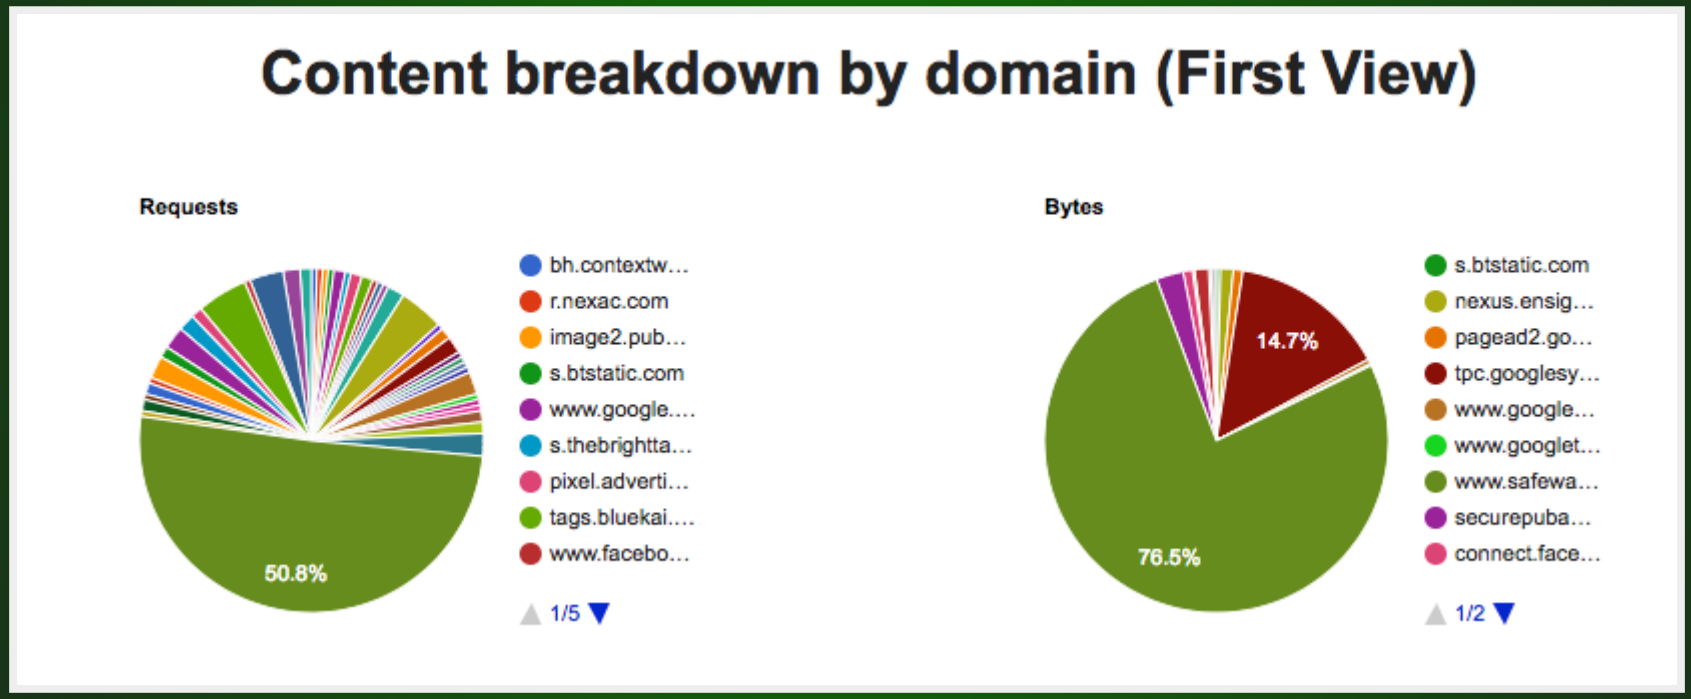

webpagetest.org/

SYNTHETIC TESTING (WEBPAGETEST)

Content breakdown by domain (First View)

Requests

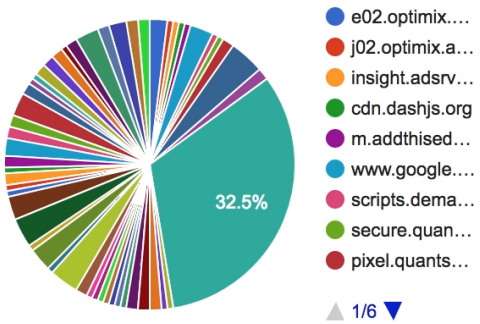

Bytes

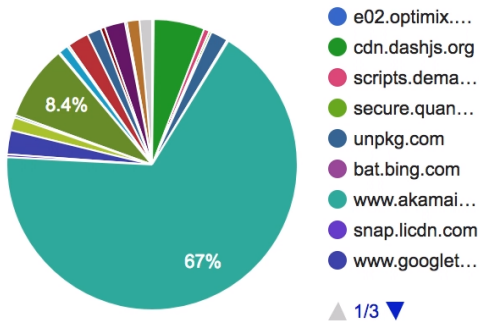

Domain	Requests
www.akamai.com	50
unpkg.com	6
www.google-analytics.com	5
secure.adnxs.com	5
www.google.com	4
s7.addthis.com	4
d.company-target.com	4
cdnssl.clicktale.net	4
va.v.liveperson.net	4
e02.optimix.asia	3
lpcdn.lpsnmedia.net	3
drvizd1lyevz4.cloudfront.net	3
pixel.quantserve.com	2
bat.bing.com	2
dc.ads.linkedin.com	2
us-east-1.dc.ads.linkedin.com	2

Domain	Bytes
www.akamai.com	1,733,898
s7.addthis.com	216,250
cdn.dashjs.org	146,209
www.googletagmanager.com	69,519
cdnssl.clicktale.net	61,197
lptag.liveperson.net	58,751
unpkg.com	49,522
c.go-mpulse.net	40,275
www.google-analytics.com	37,727
d26x5ounzdjojj.cloudfront.net	36,425
lpcdn.lpsnmedia.net	27,485
scripts.demandbase.com	16,077
79792546.va.cobrowse.liveperson.net	12,680
snap.licdn.com	8,124
www.googleadservices.com	6,671
secure.quantserve.com	6,552

BONUS: THIRD-PARTY CATEGORIZATION

github.com/simonhearne/thirdpartydb | thirdpartydb.appspot.com

REAL USER MONITORING 🥃

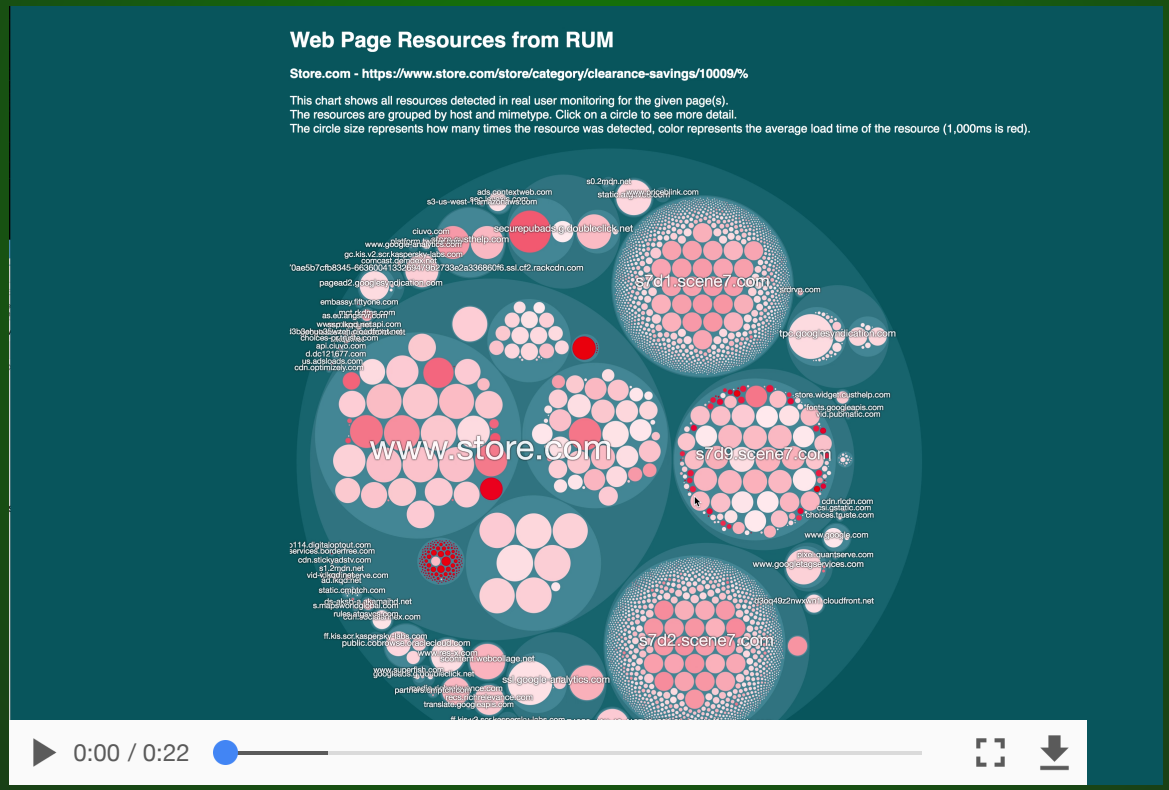

Akamai mPulse

STAGE 2:
DETERMINE THE IMPACT

SYNTHETIC TESTING (WEBPAGETEST)

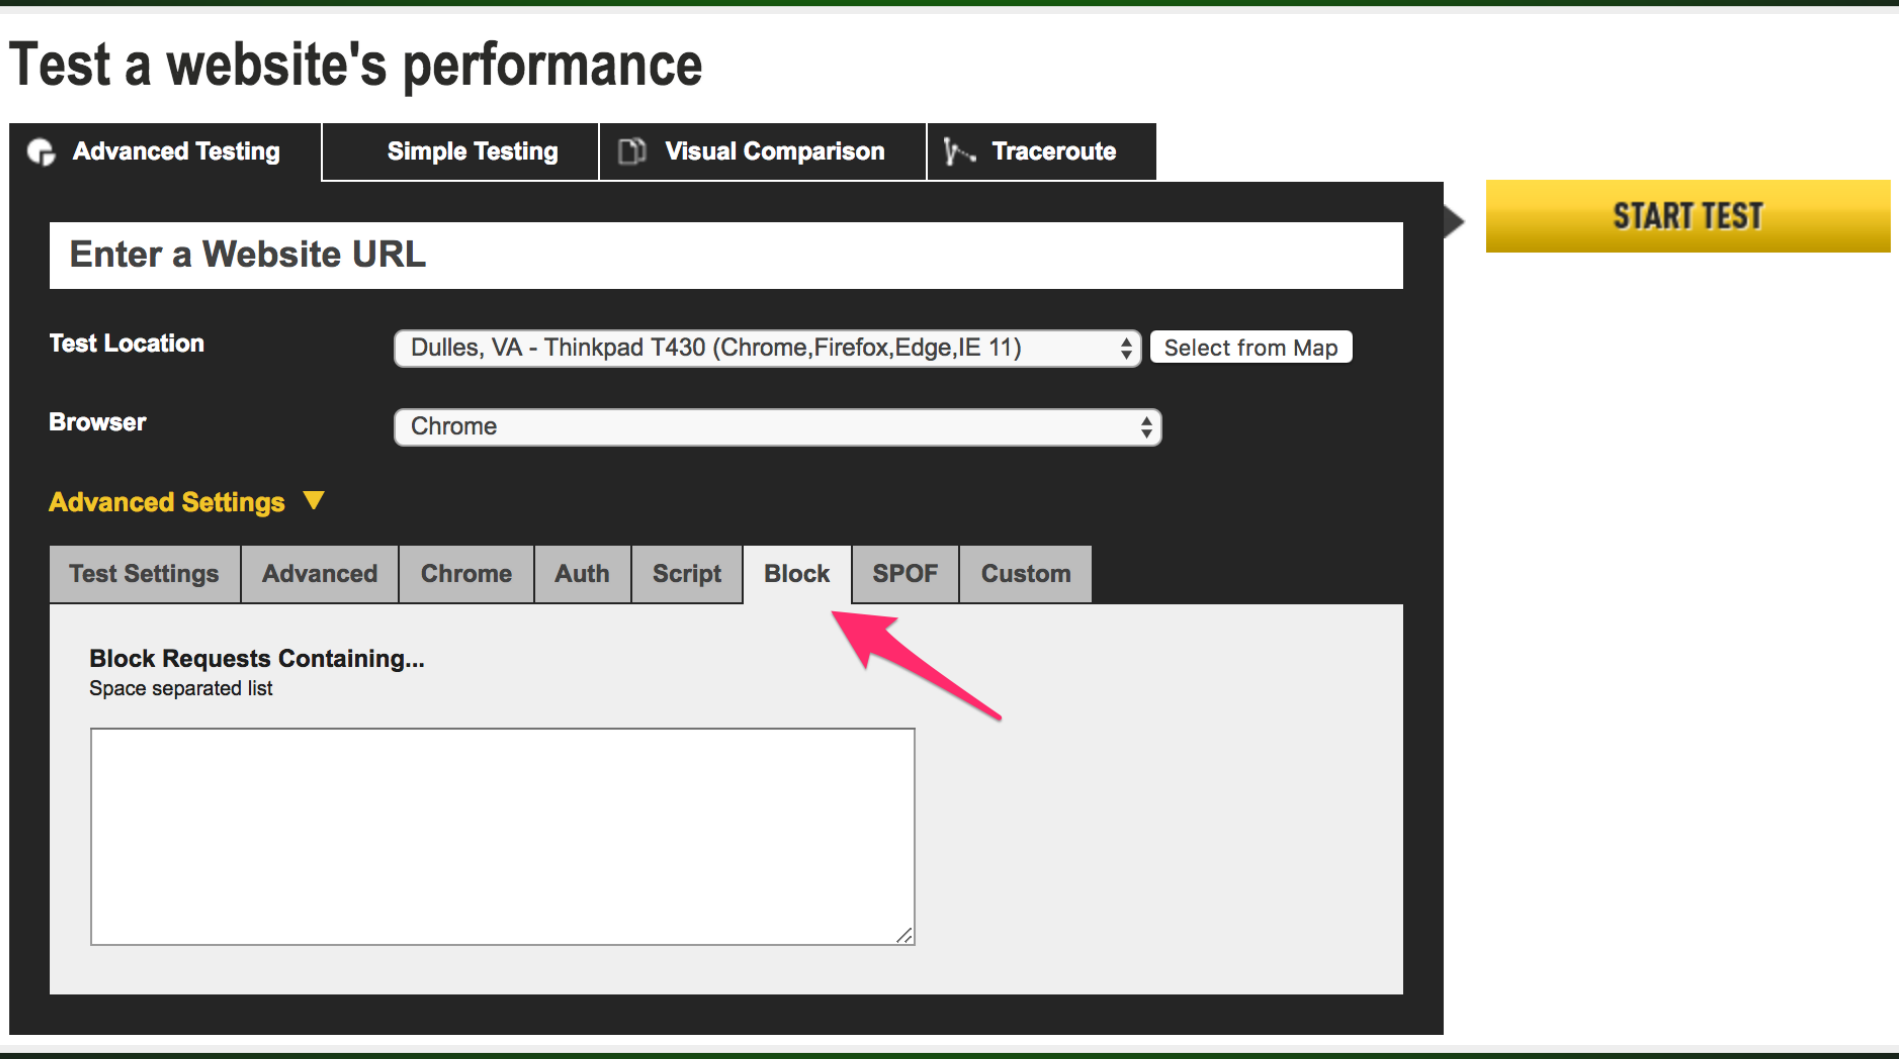

SYNTHETIC TESTING (MADE EASY)

Third-Party Impact Analysis

Third-party domain selection for https://www.akamai.com/

These are the domains detected on your page, along with some key stats on their performance impact.
Select which domains to include in the performance impact analysis then hit 'Run Analysis'.

Run Analysis

Domain	Bytes	Requests	Connections	CPU (total)	CPU (before interactive)	Include in Analysis
www.akamai.com	1,698,196	52	1	502	502	
s7.addthis.com	216,593	4	1	372	372	☑
c.go-mpulse.net	40,276	2	1	243	243	☑
www.googletagmanager.com	54,311	1	1	131	131	☐
www.google-analytics.com	49,181	6	1	42	42	☐
cdnssl.clicktale.net	62,766	4	3	29	29	☐
scripts.demandbase.com	16,856	1	1	25	25	☐
lpcdn.lpsnmedia.net	27,662	3	1	21	21	☐
79792546.va.cobrowse.liveperson.net	12,541	1	1	11	11	☐
79423.analytics.edgekey.net	126,849	1	1	8	8	☐

SYNTHETIC TESTING (MADE EASY)

Third-Party Impact Analysis

Results for https://www.akamai.com/

You can close this tab and come back later!

The results below show the impact of each third-party domain on the performance measurements of your page.
The results are ordered by the greatest impact on perceived performance (measured by Speed Index).

Blocked Domains	Speed Index (ms)	First Paint (ms)	First Interactive (ms)	Page Load (ms)	Fully Loaded (ms)	Page Size (B)	Requests
-- baseline --	4,360 *(-)*	722.7 *(-)*	*(-)*	2,444 *(-)*	8,658 *(-)*	2,615,205 *(-)*	140 *(-)*
c.go-mpulse.net	4,311 *(-49)*	931.7 *(+209)*		2,481 *(+37)*	8,732 *(+74)*	2,445,896 *(-169,309)*	143 *(+3)*
s7.addthis.com c.go-mpulse.net	4,332 *(-28)*	645.5 *(-77.2)*		2,420 *(-24)*	8,797 *(+139)*	2,218,087 *(-397,118)*	131 *(-9)*
s7.addthis.com	4,392 *(+32)*	701.1 *(-21.6)*		2,300 *(-144)*	10,563 *(+1,905)*	2,266,764 *(-348,441)*	135 *(-5)*

View comparison on webpagetest.org

Dulles, VA - Chrome - Cable on 3/26/2018, 2:35:27 PM

Made by Simon • Fork on GitHub

github.com/simonhearne/third-party-impact/

RESOURCE IMPACT FROM SYNTHETICS 🤖

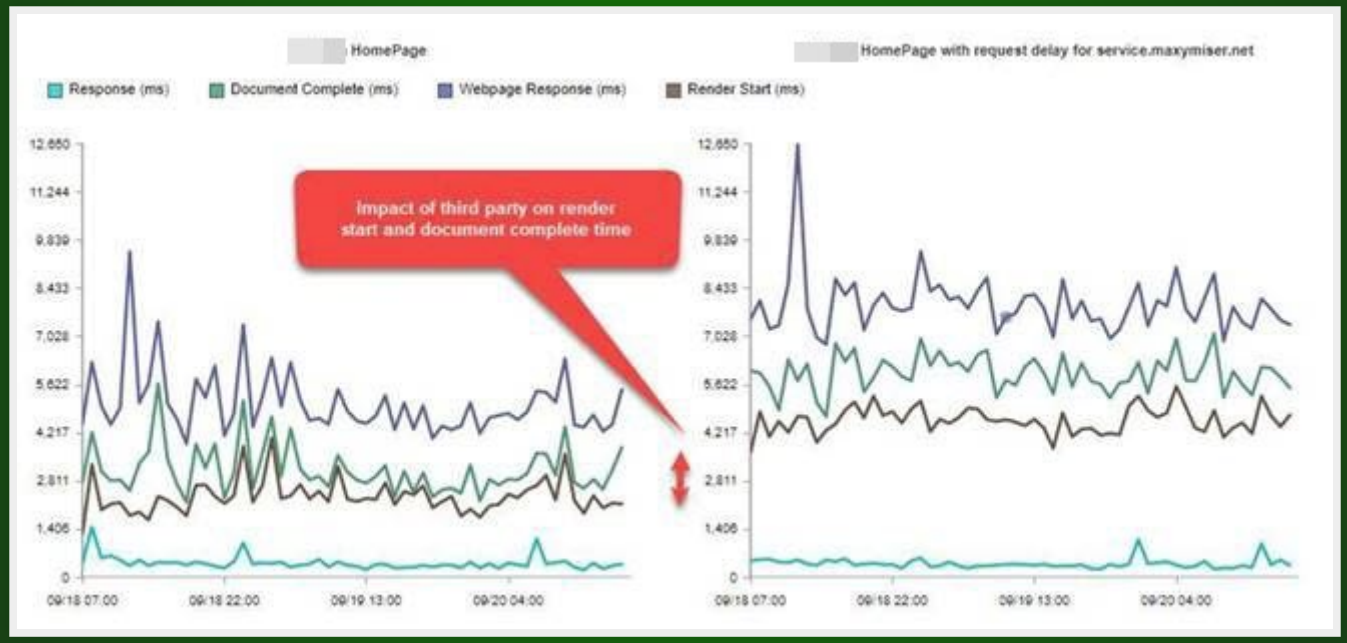

blog.catchpoint.com/2018/01/10/using-catchpoint-to-analyze-third-party-impact/

RESOURCE IMPACT FROM SYNTHETICS 🤖

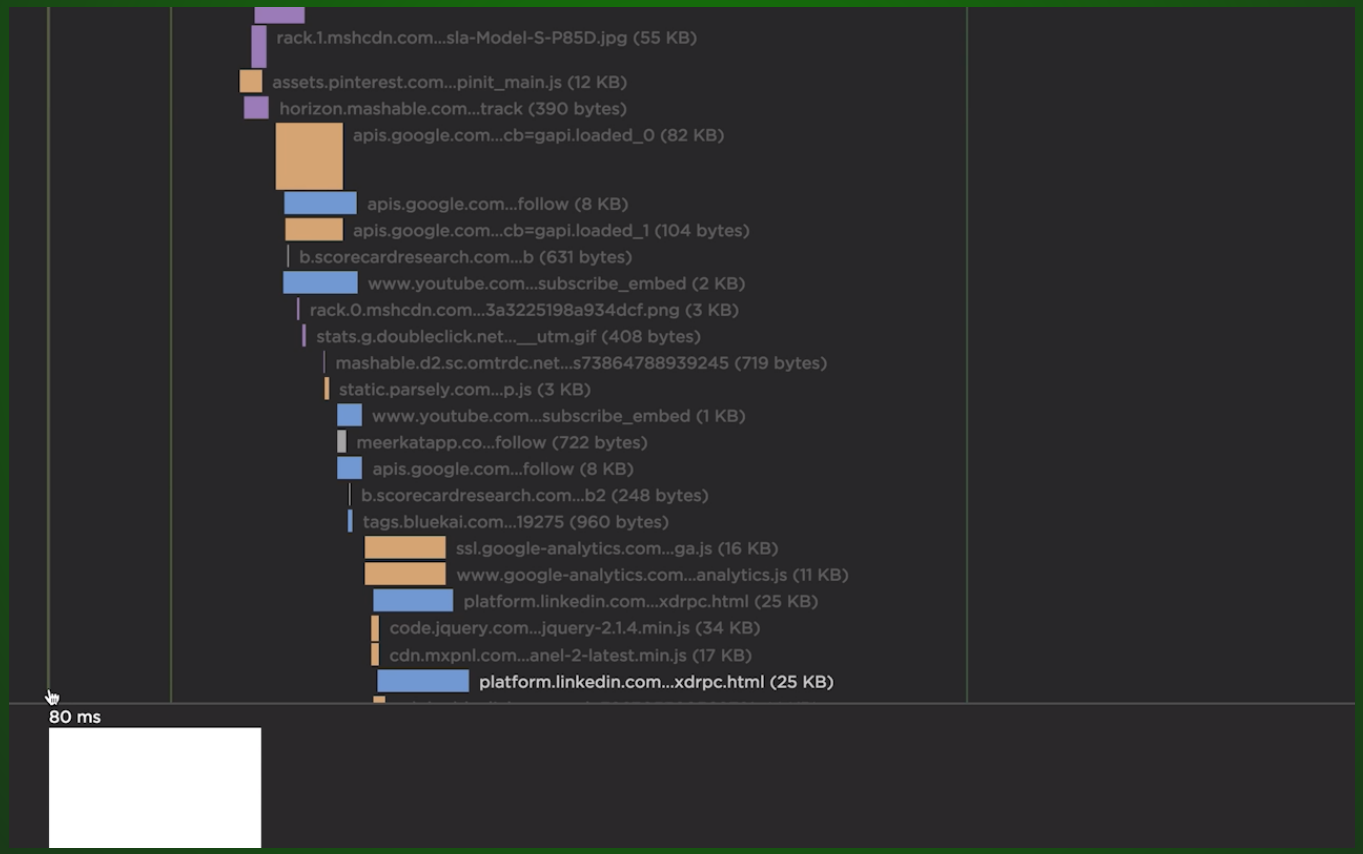

speedcurve.com/blog/ux-focus-for-waterfalls-and-third-parties/

RESOURCE IMPACT FROM RUM

ADVERTISING PARTNERS

Partner 1 = **~400ms slower** than partner 2

Migrating all ads = **220ms faster** page load

Additional revenue ~= **$12,000 per month**

Large US publishing company

LONGTASKS API

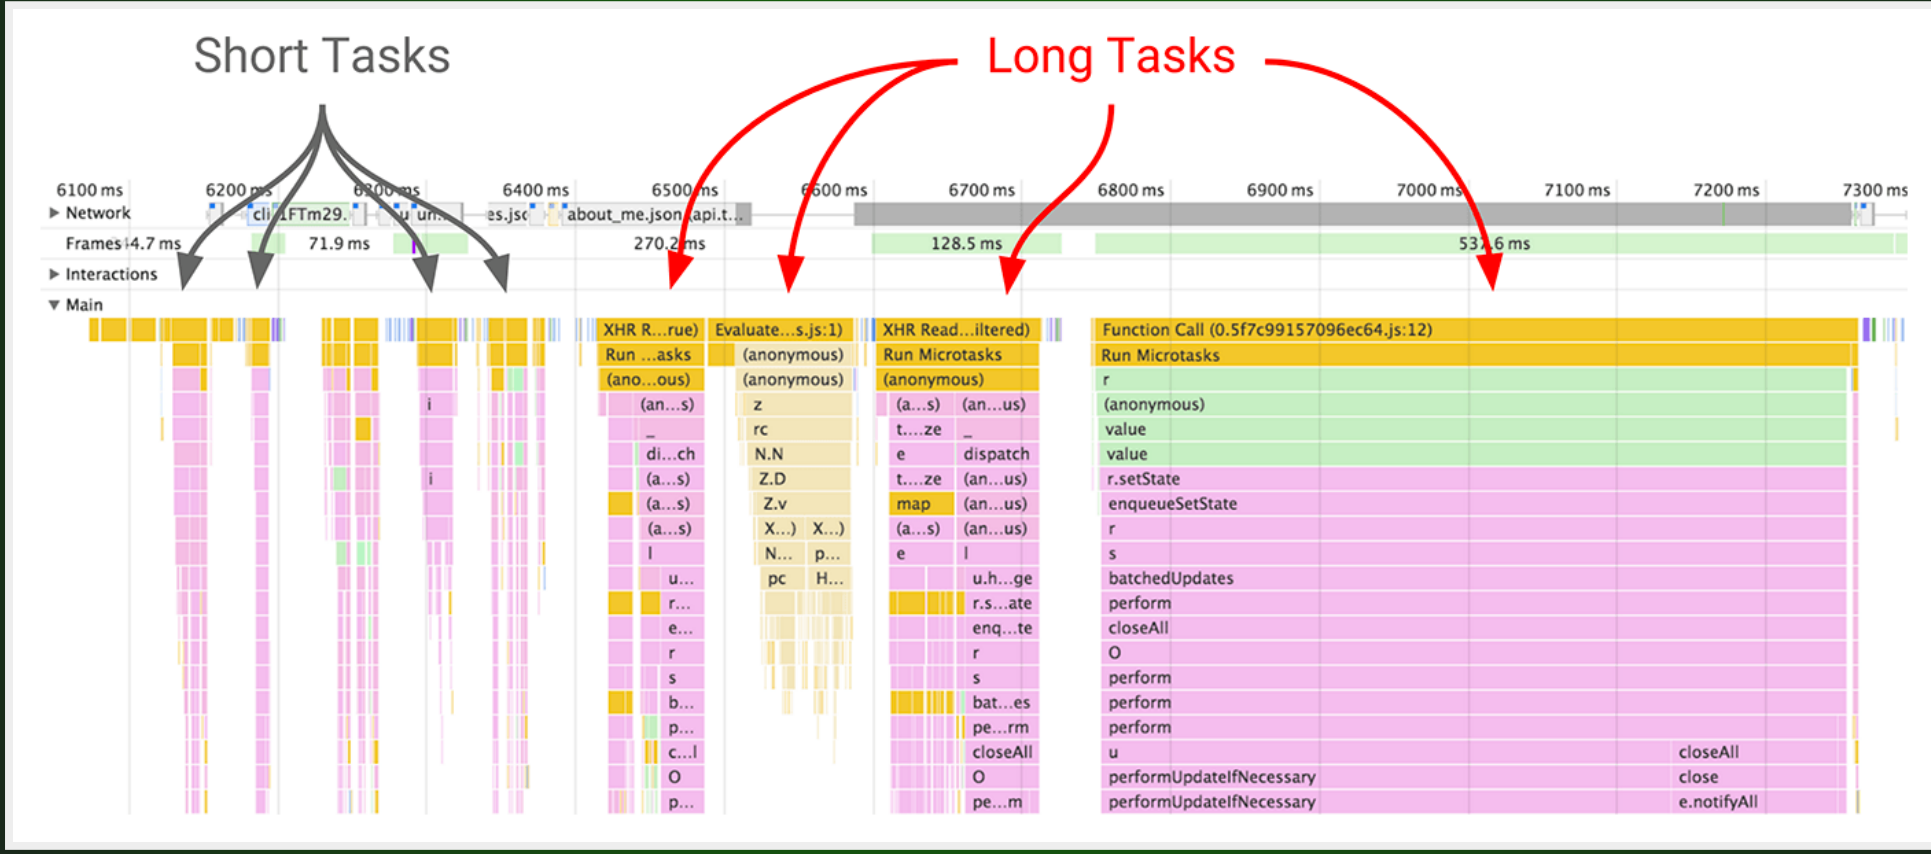

BONUS: DETERMINE THE VALUE!

Signup Click Homepage Edit Remove

The percentage of visitors who clicked on a tracked element.

Variation #1 is beating Variation #2 by +18.1%.

VARIATIONS	VISITORS	CONVERSIONS	CONVERSION RATE	IMPROVEMENT	CHANCE TO BEAT BASELINE
Variation #1	3,920	721	18.4% (±1.21%)	+18.1%	100.0%
Variation #2 BASELINE	3,999	623	15.6% (±1.12%)	---	---
Original	0	0	0.0% (±0.00%)	-100.0%	---

blog.sumall.com/journal/optimizely-got-me-fired.html

*"Everything should have a **value**, because everything has a **cost**"*

Tim Kadlec - freelance #webperf god

STAGE 3:

MEASURE THEM AND REPORT ON THEM

CONTENT SECURITY POLICY 👮

(REPORT-ONLY)


```
{
        "csp-report": {
                "document-uri": "https://yourwebsite.com/",
                "referrer": "",
                "violated-directive": "style-src",
                "effective-directive": "style-src",
                "original-policy": "",
                "disposition": "enforce",
                "blocked-uri": "inline",
                "line-number": 4,
                "column-number": 3,
                "source-file": "https://static.hotjar.com/c/hotjar-730716.js?sv=6
                "status-code": 0,
                "script-sample": ""
        }
```


report-uri.com

SYNTHETIC TESTING 🤖

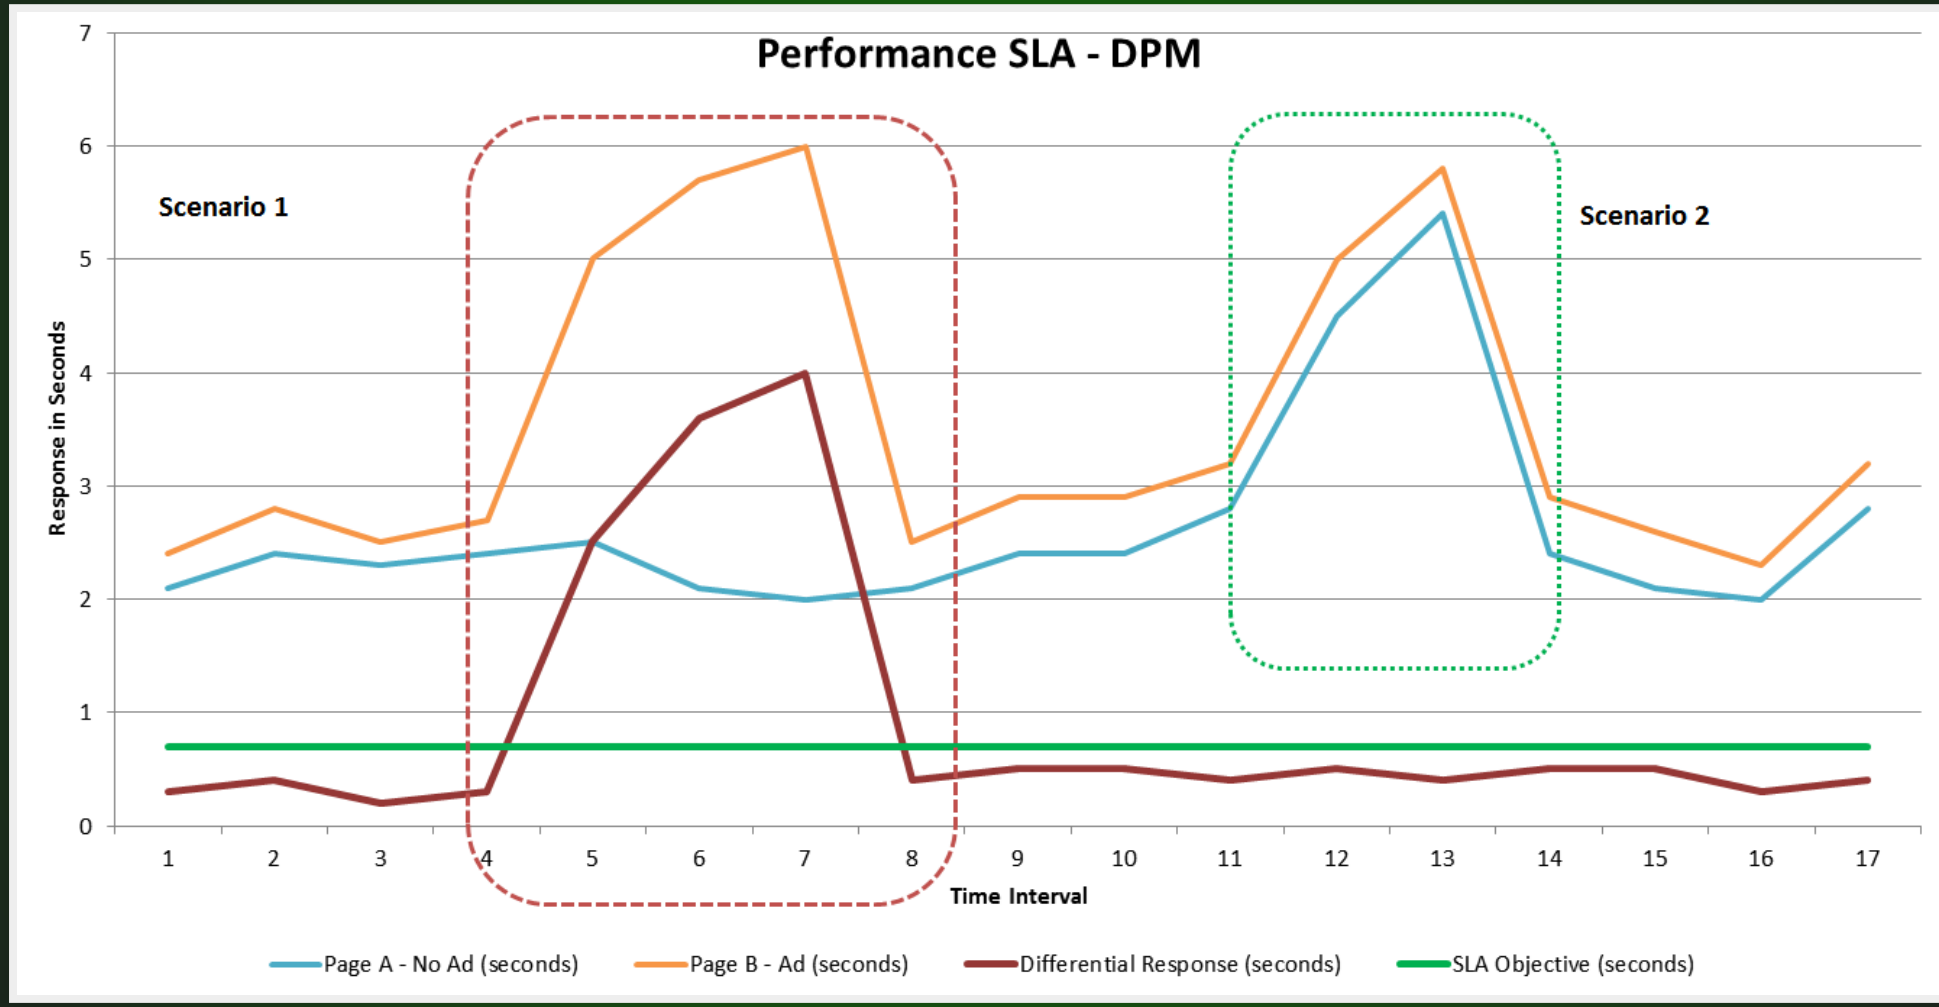

RUM 

The **best** way to monitor resources,
even with its limitations

WATERFALLS

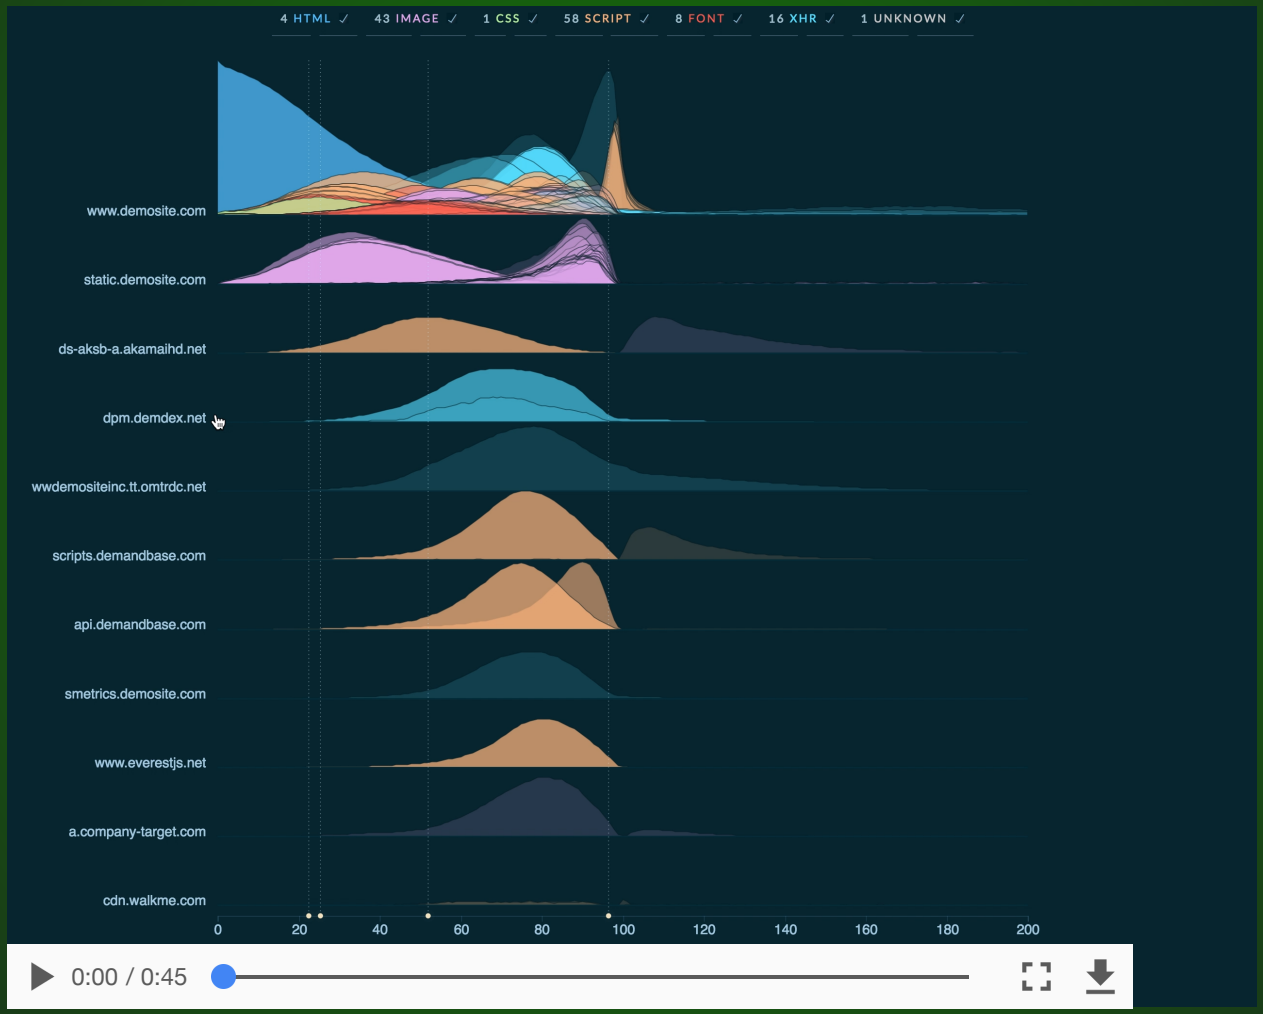
4 HTML ✓
43 IMAGE ✓
1 CSS ✓
58 SCRIPT ✓
8 FONT ✓
16 XHR ✓
1 UNKNOWN ✓
www.demosite.com
static.demosite.com
ds-aksb-a.akamaihd.net
dpm.demdex.net
wwdemositeinc.tt.omtrdc.net
scripts.demandbase.com
api.demandbase.com
smetrics.demosite.com
www.everestjs.net
a.company-target.com
cdn.walkme.com
0
20
40
60
80
100
120
140
160
180
200
0:00 / 0:45

STAGE 4:

DEFEND OURSELVES 

CONTENT SECURITY POLICY 👮

Content Security Policy Level 2 - CR

Mitigate cross-site scripting attacks by whitelisting allowed sources of script, style, and other resources. CSP 2 adds hash-source, nonce-source, and five new directives

Usage	% of all users	
Global	75.42% + 5.46% =	80.88%
U.K.	84.71% + 3.24% =	87.95%

Current aligned | Usage relative | Date relative | Show all

IE	Edge *	Firefox	Chrome	Safari	iOS Safari *	Opera Mini *	Chrome for Android	UC Browser for Android	Samsung Internet
			49						
			63		10.2				
		7 58	64		10.3				4
11	16	7 59	65	11	11.2	all	64	11.8	6.2
	17	7 60	66	11.1	11.3				
		7 61	67	TP					
			68						

caniuse.com/#feat=contentsecuritypolicy2

CONTENT SECURITY POLICY 👮

Introducing XSS Auditor reporting to Report URI

March 26, 2018

Whilst we already have support for CSP reports over at Report URI, there is another potential source of information about XSS attacks that may be attempted or happening on your site. The X-XSS-Protection header allows you to configure the XSS Auditor, deem what action it should take and request that the auditor send reports if action is required. We now support XSS Auditor reporting on Report URI!

The XSS Auditor

The XSS Auditor runs whilst HTML is being parsed and attempts to find reflected XSS attacks against the user. If it finds a possible attack the Auditor can take no action, it can filter what it thinks is the attack payload or it can refuse to render the page at all. You can find more details about the XSS

The Author

Scott Helme is a Security Researcher, international speaker and author of this blog. He is also the founder of securityheaders.io and report-uri.com, free tools to help organisations better deploy security.

CSP STRICT-DYNAMIC

'strict-dynamic'

The `strict-dynamic` source expression specifies that the trust explicitly given to a script present in the markup, by accompanying it with a nonce or a hash, shall be propagated to all the scripts loaded by that root script. At the same time, any whitelist or source expressions such as `'self'` or `'unsafe-inline'` will be ignored. See script-src for an example.

CONTENT SECURITY POLICY 

- 👍 CryptoJacking
- 👍 XSS
- 👎 Maintenance

SUB-RESOURCE INTEGRITY 🔍

```
<link
        rel="stylesheet"
        href="//maxcdn.bootstrapcdn.com/.../bootstrap.min.css"
        integrity="
                sha256-8EtRe6XWoFEEhWiaPkLaw...=
                sha512-/5KWJw2mvMO2ZM5fndVxU...=
                "

        crossorigin="anonymous">
```

```
<script
        src="//ajax.googleapis.com/.../jquery.min.js"
        integrity="
                sha256-ivk71nXhz9nsyFDoYoGf2...=
                sha512-7aMbXH03HUs6zO1R+pLye...=
                "

        crossorigin="anonymous"></script>
```

SUB-RESOURCE INTEGRITY 🔍

👍 Malicious Code

👍 Untested Changes

👎 Maintenance

⭐ Signature-based Restrictions...*

github.com/mikewest/signature-based-sri

SERVICE WORKER 💪

```
function timeout(delay) {
        return new Promise(function(resolve, reject) {
                setTimeout(function(){
                        resolve(new Response('', {
                                status: 408,
                                statusText: 'Request timed out.'
                        }));
                }, delay);
        });
}

self.addEventListener('fetch', function(event) {
        // Only fetch JavaScript files for now
        if (/\.js$/.test(event.request.url)) {
                event.respondWith(Promise.race([timeout(2000), fetch(event.request.url)]));
        } else {
                event.respondWith(fetch(event.request));
        }
});
```

calendar.perfplanet.com/2015/reducing-single-point-of-failure-using-service-workers/

SERVICE WORKER 💪

- 👍 CDN / Network outages
- 👎 Not on first pageview
- 👎 Maintenance
- 👎 You can break your site

SELF-HOSTING / PROXYING ☔

google-webfonts-helper.herokuapp.com/fonts

SELF-HOSTING / PROXYING ☔

Why Self-Host the Test Files

As a website owner, you may be skeptical about using external JavaScript from a third party. For your requirements, you may want to remove dependency on VWO servers and have all the test resources under your control. Using the self-hosting option, you can store the VWO JavaScript files locally to ensure that the test settings, files, and code for your tests reside on your server.

vwo.com/knowledge/host-vwo-javascript-files-on-your-server/

SELF-HOSTING / PROXYING ☔

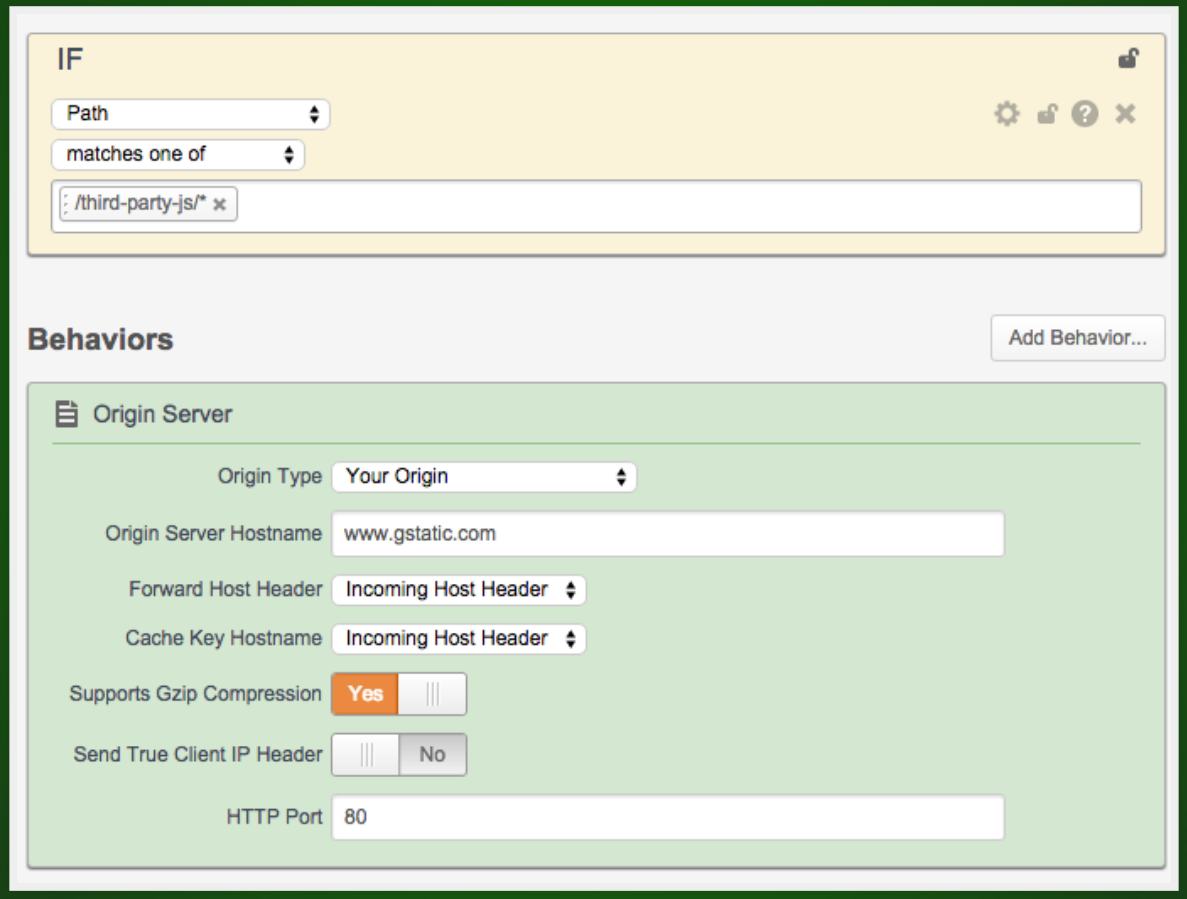

community.akamai.com/community/web-performance/blog/2016/01/13/5-ways-to-prevent-slow-3rd-party-front-end-services

SELF-HOSTING / PROXYING ☔

- 👍 CDN / Network outages
- 👍 Shared TCP connection
- 👎 Maintenance

STAGE 5:

HAVE A THIRD-PARTY POLICY 

- What does it do?
- Who uses it?
- What's the risk to the site?
- How do you remove it?

SHARE WITH OTHER TEAMS! 🗣️

speedcurve.com/demo/thirdparty/

THIRD-PARTY CONTENT MAY BE A WEAK LINK

BUT IT'S HERE TO STAY

FIVE THINGS YOU CAN DO TODAY:

- Know what's there
- Measure them
- Have a solid defense
- Share the data
- Have third-party policy

FURTHER READING

Loading Third-Party JavaScript

Contents

By Addy Osmani
Eng Manager, Web Developer Relations

By Arthur Evans
Arthur is a Tech Writer

You've optimized all of your code, but your site still loads too slowly. Who's the culprit?

THANK YOU, GOOD LUCK!

 @SimonHearne

 webperf.ninja/tools

 simonhearne.github.io/weak-links

@SimonHearne