











່ເmec

3

#### Web Performance

- HTTP/2 is the new kid in town!
- Replaces good old trusted HTTP/1.1
  - Currently 13.3% of sites is on H2\*

- How is HTTP/2 different from HTTP/1.1?
- Is h2 always faster than h1? And why?
- Do we need to change our sites to get the most out of h2?

lmec

## HTTP/1.1 : the bad and the ugly 1

- Single resource at a time per TCP connection
- "Head of Line" (HOL) blocking



NOWLEDGE IN ACTION



NOWLEDGE IN ACTION

#### HTTP/1.1: 6 is not enough!

- Best Practice 1 : Concatenation
- = merge files together to reduce # resources
- 1.js + 2.js + 3.js = scripts.js
- Downside: reduced cacheability



OWLEDGE IN ACTIO

#### HTTP/1.1 : 6 is not enough!

- Best Practice 2 : Domain Sharding
- Content Delivery Networks (CDNs)
- img1.site.com, img2.site.com, static.site.com
- Downside: serverside overhead



## HTTP/2 : The Challenger

- Started as SPDY, standardized in 2016
- Back to 1 TCP connection
  - Priority-based multiplexing of resource streams



WLEDGE IN ACTI

#### HTTP/2 : Solves HOL blocking!



### HTTP/1.1 : the bad and the ugly 2

- Round trip times cause most of the delay
- "dependency graph" / resource discovery



Index.html references Style.css references Font.woff2

-> min. 3 round-trips needed before rendering can start!

UHASSEL

OWLEDGE IN ACTION

umec



### HTTP/1.1 : Round trip times

- Best Practice 3 : Resource embedding
- Inline resource inside HTML file
- Ink rel="stylesheet" href="style.css" /> VS
- <style> #myDiv{ color: #FFBEEF; } </style>

Downside: reduced cacheability

#### HTTP/2 : The Challenger

#### Ability to Push resources along



KNOWLEDGE IN ACTION

#### HTTP/2 : The bad and the ugly

- In theory, H1 best practices are now obsolete
  - High-performance browser networking (Grigorik, 2013)
- But: Single TCP connection
  - More sensitive to loss than # parallel connections
- But: TCP in-order guarantee (kernel-level)
   ~Head-of-line blocking when retransmit
- Problems expected on bad connectionsDifficult to estimate actual impact...



#### Epic battles in contradictory previous work





Epic battles in contradictory previous work

- 1. Implementations change
  - SPDY vs H2
  - Server and Browser support
- 2. Test setups are limited
  - Single browser, server, metric, network, tool, platform, ...

? Protocol-inherent or implementation-based ?



#### The Speeder Framework

Table 1: Speeder software and metrics.

HTTP/1.1 (cleartext), HTTPS/1.1, HTTPS/2
Chrome (51 - 54), Firefox (45 - 49)
Sitespeed.io (3), Webpagetest (2.19)
Apache (2.4.20), NGINX (1.10), NodeJS (6.2.1)
<ul> <li>DUMMYNET (cable and cellular) (Webpagetest)</li> <li>fixed TC NETEM (cable and cellular)</li> <li>dynamic TC NETEM (cellular) (Goel et al., 2016)</li> </ul>
All Navigation Timing values (Wang, 2012), SpeedIn- dex, firstPaint, visualComplete, other Webpagetest metrics (Meenan, 2016)

#### TEST MOST OF THE THINGS!

UHASSEL

່ເmec

#### loadEventEnd vs SpeedIndex

#### IoadEventEnd

- document.onload
- Everything is downloaded and rendered



#### loadEventEnd vs SpeedIndex

#### SpeedIndex

- % of visual complete over time (video frames!)
- Quick to render = better for user experience
- Measured in milliseconds = LOWER IS BETTER



#### The Speeder Framework

Table 1: Speeder software and metrics.

Protocols	HTTP/1.1 (cleartext), HTTPS/1.1, HTTPS/2
Browsers	Chrome (51 - 54), Firefox (45 - 49)
Test drivers	Sitespeed.io (3), Webpagetest (2.19)
Servers	Apache (2.4.20), NGINX (1.10), NodeJS (6.2.1)
Network	<ul> <li>DUMMYNET (cable and cellular) (Webpagetest)</li> <li>fixed TC NETEM (cable and cellular)</li> <li>dynamic TC NETEM (cellular) (Goel et al., 2016)</li> </ul>
Metrics	All Navigation Timing values (Wang, 2012), SpeedIn dex, firstPaint, visualComplete, other Webpagetest metrics (Meenan, 2016)

#### TEST MOST OF THE THINGS!

່ເmec

#### The Speeder Framework

- Automate this setup using Docker
  - Light-weight stateless VMs based on install scripts
  - Easy to Update!
  - Easy to Share!



- Available on speeder.edm.uhasselt.be
- Go forth and reproduce!

umec

#### Recap

- 1. 3 Best Practices to combat H1 downsides
- 2. HTTP/2 performance is unclear
- 3. Test many permutations with Speeder

- ? Impact on developer workflow ?
- Synthetic test cases for best practices
- Realistic test cases



WLEDGE IN ACTIO

#### 1. Concatenation : 500 CSS/JS 7k HTTP/1 6k HTTPS/1 HTTPS/2 loadEventEnd (ms) 8 k k k 8 k HTTPS/2 (firefox) 1k 100 200 300 400 500 100 200 300 400 500 0 simple CSS files, total 16.4 KB complex JS files, total 395 KB (a) amount of files **CSS** Simple **JS Complex** UHASSEL ່ເກາec

OWLEDGE IN ACTIO

#### 1. Concatenation : 500 CSS/JS 7k HTTP/1 6k HTTPS/1 HTTPS/2 loadEventEnd (ms) 8 k k k 8 k HTTPS/2 (firefox) 1k 100 200 300 400 500 100 200 300 400 500 0 simple CSS files, total 16.4 KB complex JS files, total 395 KB (a) amount of files **CSS** Simple **JS** Complex UHASSEL ່ເກາec

WLEDGE IN ACTIO

#### 1. Concatenation : 30 CSS/JS









# 1. Concatenation : JS SpeedIndex500 individual files



#### 2. Sharding

- Images : 380 small, 42 medium, 30 large
- Sharded over 4 hosts (24 H1 conns, 4 H2)











WLEDGE IN ACTIC


WLEDGE IN ACTION



WLEDGE IN ACTIC



WLEDGE IN ACTIC











## 4. Realistic websites

How would a developer test H2 on his site?
No sharding for main assets (~H2 should win)







mec











# Conclusions : HTTP/2

- Better for large amounts of small files
- Worse for larger files
- Suffers from bad networks but so does HTTP/1
- For most "average" testcases, performance is similar between H1 and H2, with (much) less overhead (# connections) from H2



lmec

## Conclusions : HTTP/1 Best Practices

- Concatenation: implementation-limited
- Sharding: protocol-inherent
- Embedding: can be replaced by Push
- Most "average" H1 websites can be transferred to H2 without much work ☺



lmec

## Conclusions : Outlook

- More fine-grained caching (web components, Single Page Apps, ...)
- Bigger websites remain difficult (4k images, VR/gaming, ...)

QUIC is already being standardized!





່ເກາec

## Ask me about:

- Cellular network emulation
- QUIC protocol
- TCP slow start and embedding/push
- HPACK header compression
- Browser rendering behavior
- Website optimization anecdotes
- Docker / The Speeder Framework
- HTTPS vs HTTP cleartext
- Other uses of HTTP/2 Push

imer

## Cellular network emulation

- "Http/2 performance in cellular networks" (Goel et al., 2016) – MobiCOM
- Traces of real networks @ CDN edge
- Modulate TC NETEM parameters every 70 ms
- https://github.com/akamai/cell-emulation-util
- Fixed model based on Chrome Devtools
  - Add loss
  - Typically clearer but similar trends



## QUIC protocol

- QUIC = Quick UDP Internet Connections
  - "Re-implement TCP in application layer" ~TCP 2.0
  - Removes TCP HOL-blocking and adds smarter flow-control / less sensitive to loss
  - FEC, Connection UUID/handoffs, multipath, ...



## TCP slow start and embedding/push

- TCP slow start uses congestion window
  - Initial value is ~14KB in modern linux
  - Grows when data is ACK'ed (takes roundtrip)
- If data > 14KB limit, additional RTTs needed!



## HPACK header compression

- HTTP headers are often the same
  - Especially important for large cookies
- HPACK uses dynamic dictionary encoding
  - Learns during connection, builds dict at both sides

Table 4: Total bytes sent by Google Chrome ( $\sim$  HTTP headers) and ratio to total page size. For many small files, the HTTP header overhead is significant.

File count	Protocol	1 host BytesOut	4 hosts BytesOut	Total page size	4 hosts % of total page size			
42 files	h2s h1s	649 2786	1362 3346	1075000 1075000	0.1% 0.3%			
400 files	h2s h1s	29580 165300	38680 177600	610000 610000	6% 19%		_	
				6		mer	►►   UHA	SSELT

OWLEDGE IN AC

## Browser rendering behavior

- Render blocking"
  - CSS and sync JS : could change DOM/CSSOM!
- Mitigation: Place on bottom of <body>

Table 3: Observed browser render-blocking behavior. Chrome blocks on too much, while Firefox doesn't seem to block on anything.

	Expected <head></head>	Chrome <head></head>	Firefox <head></head>	
CSS blocking JavaScript blocking		blocking	progressive	
		progressive	progressive	
	Expected bottom	Chrome bottom	Firefox bottom	
CSS	progressive	blocking	progressive	
JavaScript progressive		progressive	progressive	
				<b>+</b> +
			imec	UHASSEL

OWLEDGE IN

## Website Optimization anecdotes

- Hero image halfway on the page
  - 6 HTTP/1 connections had been taken up, HTTP/2 gives all images same priority
  - Concatenation : only benefits HTTP/1
  - Moving image up in HTML : benefits both



## Website Optimization anecdotes

- Embedding is Best Practice!
  - Let's embed our fonts as well! Base64-style!
  - Embeds > 100KB in fonts that are not cached...
  - Every new page load, flash-of-unstyled-content despite embedding



## Docker / The Speeder Framework

- 3 Docker containers
  - Base (C&C), Backend and Agents
  - Dockerfile is easy to update, rebuild, deploy
- 5 private WebPageTest instances
- C&C starts instances and restarts all running software before each test
- Spread over 16 Linux VMs and 5 Windows VMs



# Docker / The Speeder FrameworkSimple GUI for starting tests

#### 0 tests sheduled

Testrun:	
debug_test	
Label:	

Run	count:	
2		

Option 1 : choose a locally served testcase (use Speeder backend setup)

#### Testcases: 0 / 286

```
□ 2v1 aggregation 🗹 0 / 57
□ 2v1 aggregationComplex 🗹 0 / 57
□ 2v1 aggregationLibrary  0 / 13
□ 2v1 domainSharding  0 / 10
□ 2v1 gzlib 🗹 0 / 16
□ 2v1 http2Priority 🗹 0 / 4
□ 2v1 images 🗹 0 / 11
□ 2v1 inlining  0 / 33
□ 2v1 inliningComplex  0 / 33
2v1 loadingJS 2 0 / 8
□ 2v1 redirect 🗹 0 / 9
2v1 sanitychecks Ø 0 / 1
□ 2v1 serverPush 🗹 0 / 24
□ 2v1 template  0 / 2
□ H2priorities H2oWithVsWithoutHttp2Prios  0 / 1
□ H2priorities test1debug 🗹 0 / 1
□ VoDLib NginxVsNode 🗹 0 / 4
□ drupal main 🗹 0 / 0
examples customconfigs @ 0 / 2
```

#### Backends:

Inode Inginx Inginx Inginx Inginx Ingina In

#### Frontends:

sitespeed chrome firefox slimer browsertime\_chrome SRUMdriver\_chrome SRUMextension\_chrome wpt chrome firefox ie

#### **Protocols:**

✓HTTPS 2
✓HTTPS 1.1
✓HTTP 1.1

#### Traffic:

none
wifi
wifi\_transatlantic
uvifi\_transatlantic\_loss
4g
4g\_transatlantic\_loss
3g\_loss
2g\_loss
cellular\_noloss
cellular\_fair
cellular\_fair
cellular\_passable
cellular\_poor
cellular\_verypoor

່<sub>ເກາຍc</sub> ບ



# Docker / The Speeder Framework Quick graphing of results



## HTTPS vs HTTP cleartext

- Push towards a safer web
- H2 cleartext isn't supported by any browser
- Many of the newer (performance) features only work over HTTPS



# Other uses of HTTP/2 Push

- Warm vs cold connections
- Single Page Apps / REST APIs
- Service Workers
- Video streaming
- CDN-side
- But: Caching problems!
- But: Server support!
- But: Re-prioritization!



lmec

## loadEventEnd vs SpeedIndex

## SpeedIndex

- % of visual complete over time (video frames!)
- Quick to render = better for user experience
- Measured in milliseconds = LOWER IS BETTER



WLEDGE IN ACTIC

67

## 1. Concatenation

- Images : 380 small, 42 medium, 30 large
- CSS/JavaScript : 1 500 files, simple and complex











WLEDGE IN ACTIC




Both are the same = rendering started at onLoad => Page was white until all CSS was loaded

> ▶►| UHASSE

່ເກາec

## 4. Realistic websites

- Many factors contribute to load times
  - Polaris (Netravali et al.), Shandian (Wang et al.)
- Other related work exists
  - <u>http://isthewebhttp2yet.com</u> (Varvello et al.)

## Can we simply transfer HTTP/1 sites to HTTP/2 ?



່ເຫາec

## Image Sources

- http://deliveryimages.acm.org/10.1145/2390000/2380673/figs/f5.jpg
- https://dab1nmslvvntp.cloudfront.net/wp-content/uploads/2017/01/1484692838webpack-dependencytree.png
- https://blog.cloudflare.com/content/images/2015/12/domain-sharding-1.png
- https://calendar.perfplanet.com/2016/http2-push-the-details/
- http://www.splicemarketing.co.uk/ blog/Splice Blog/post/the-arrival-of-http2/
- http://s1377.photobucket.com/user/barteks11/media/Marvel%20gifs/h20wkj2-iron-man-vs-captainamerica-who-sides-with-who-in-marvel-s-civil-war-jpeg-151871\_zpsarkko07f.jpg.html
- https://sites.google.com/a/webpagetest.org/docs/using-webpagetest/metrics/speed-index
- <u>https://twitter.com/officialrtsf</u>
- https://twitter.com/ThePracticalDev
- <u>http://www.netflix.com</u>



UHASSEL

່ເກາec