

HTTP/2: Operable and Performant



Mark Nottingham @mnot (@akamai)

A close-up photograph of a vintage control panel, likely from a ship or aircraft. The panel is made of a light-colored, textured material, possibly aluminum or steel, and features three semi-circular ammeters in the top row. Each ammeter has a white face with black markings and a needle, and is labeled 'D.C. AMPERES'. Below the ammeters are two circular buttons, each with the word 'RAISE' embossed on its top edge. To the right of the second 'RAISE' button is a black toggle switch. At the bottom of the panel, there are several black rectangular components, possibly relays or switches, with various terminals and connectors. The overall appearance is that of a well-used, historical piece of machinery.

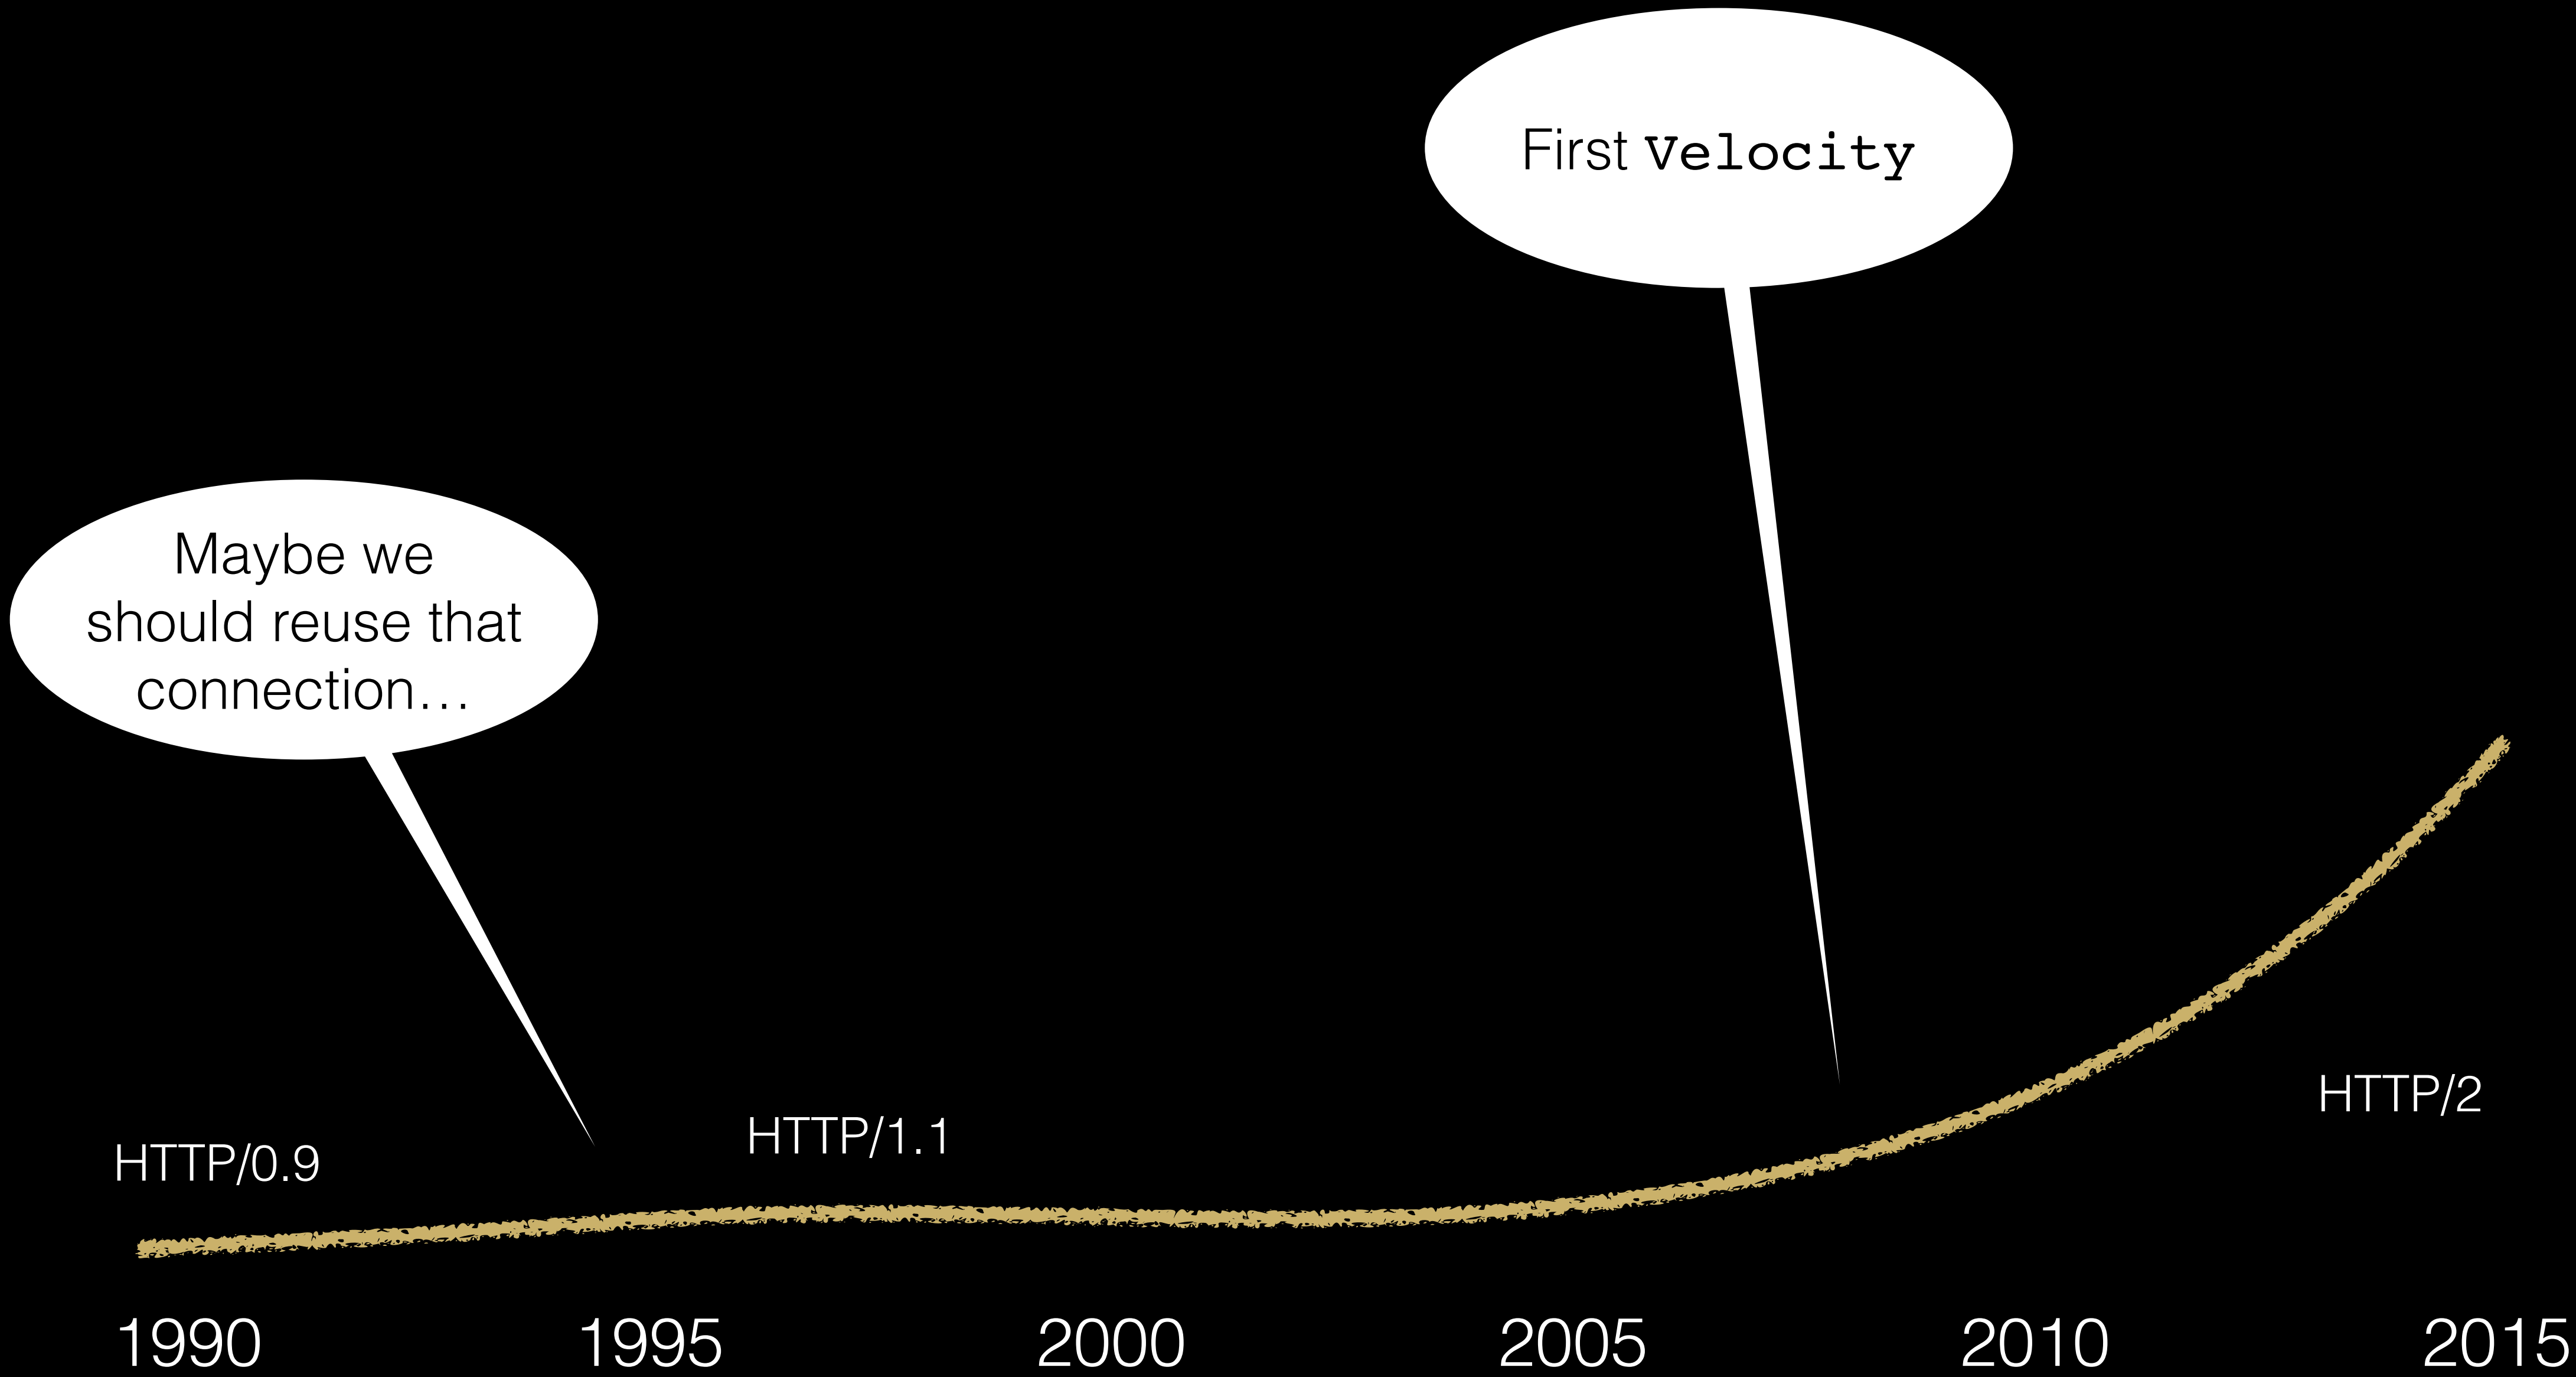
This talk may be disappointing.



Donald Rumsfeld

“As we know, there are **known knowns**; there are things we know we know. We also know there are **known unknowns**; that is to say we know there are some things we do not know. But there are also **unknown unknowns** — the ones we don’t know we don’t know.”

HTTP Perf Knowledge



HTTP/2 is almost here!

... and it's a lot like SPDY



HTTP/2 in One Slide

- Multiplexing + Header Compression + Server Push
- Goal: one connection from browser
- Post-WGLC
- Coming in Firefox, Chrome, IE, others very, **very** soon



Florian Bender 2014-09-15 14:46:18 PDT

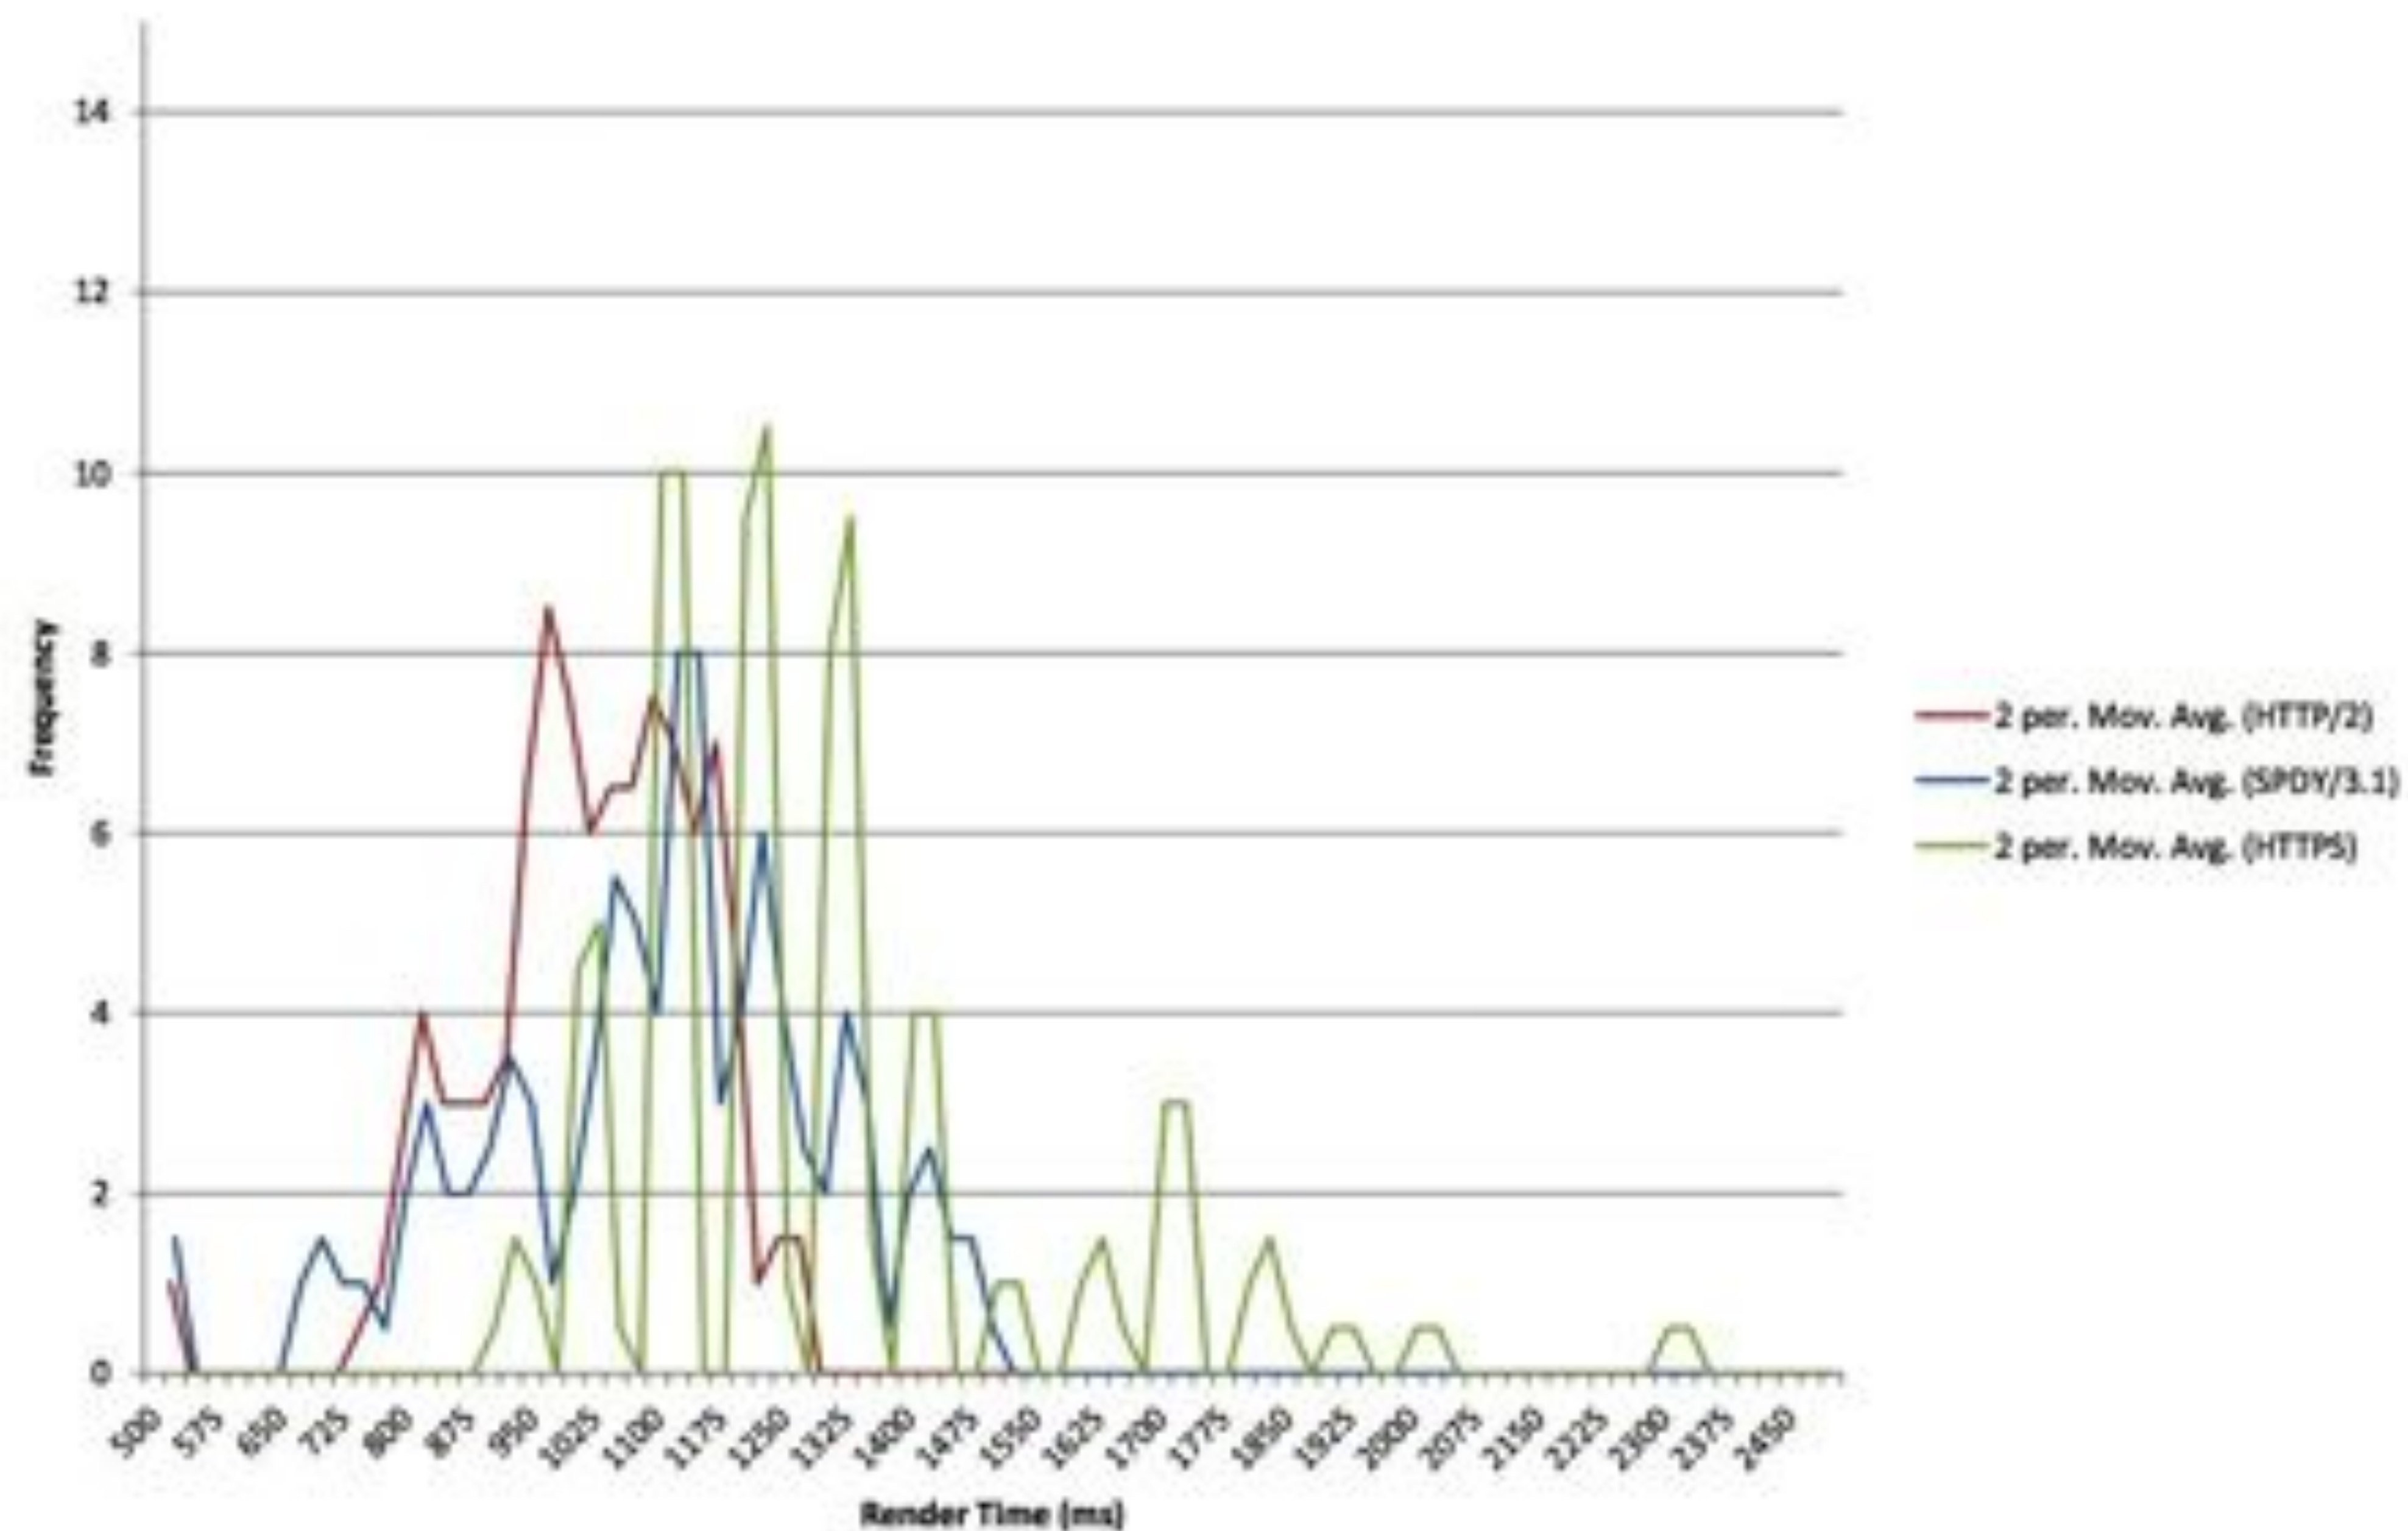
Release Note Request (optional, but appreciated)

[Why is this notable]: The next era in Web tech has begun. Seriously.

[Suggested wording]: Implemented HTTP/2 and APLN.

HTTP/2 vs SPDY/3.1 vs HTTPS

Render Time



{HTTP2, SPDY} and Performance

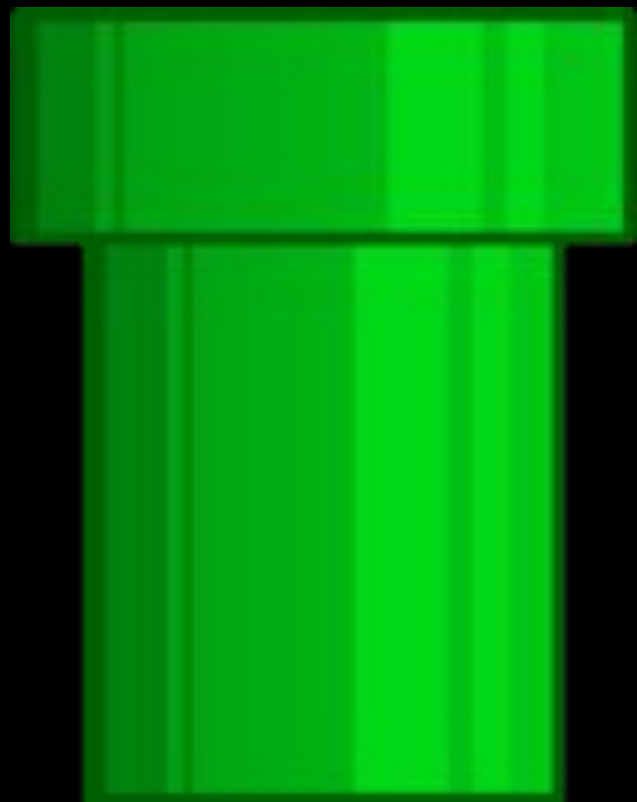
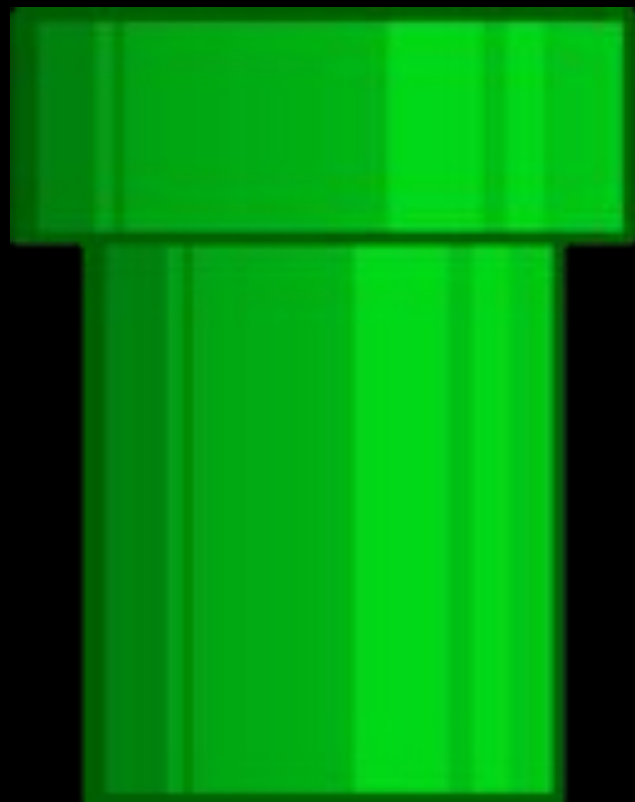
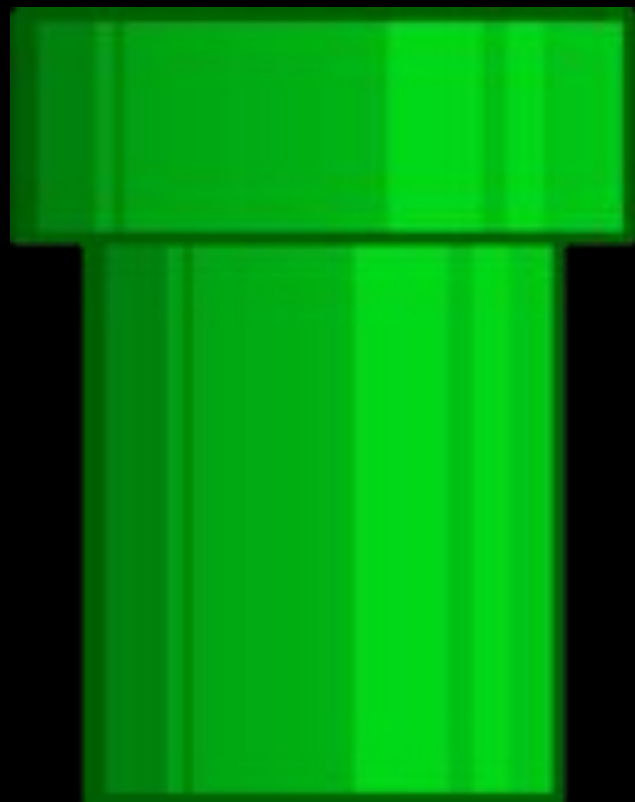
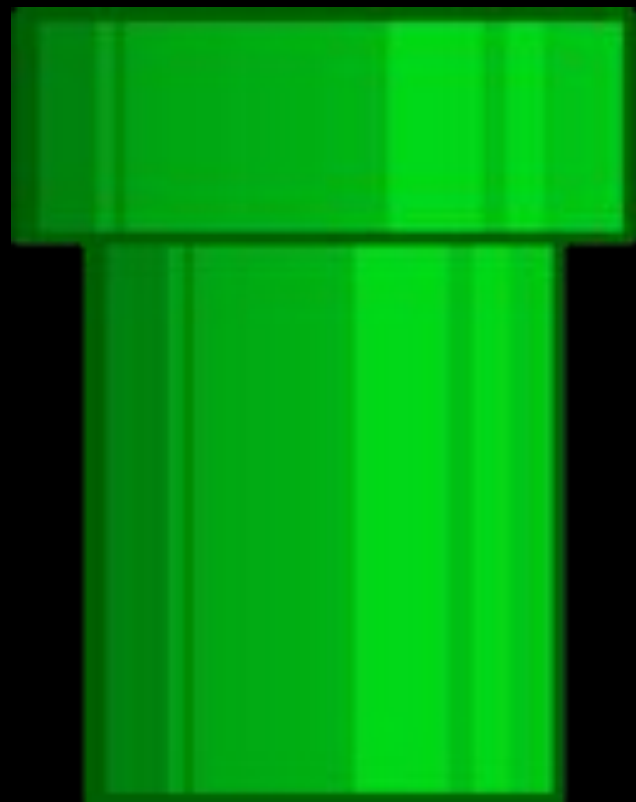
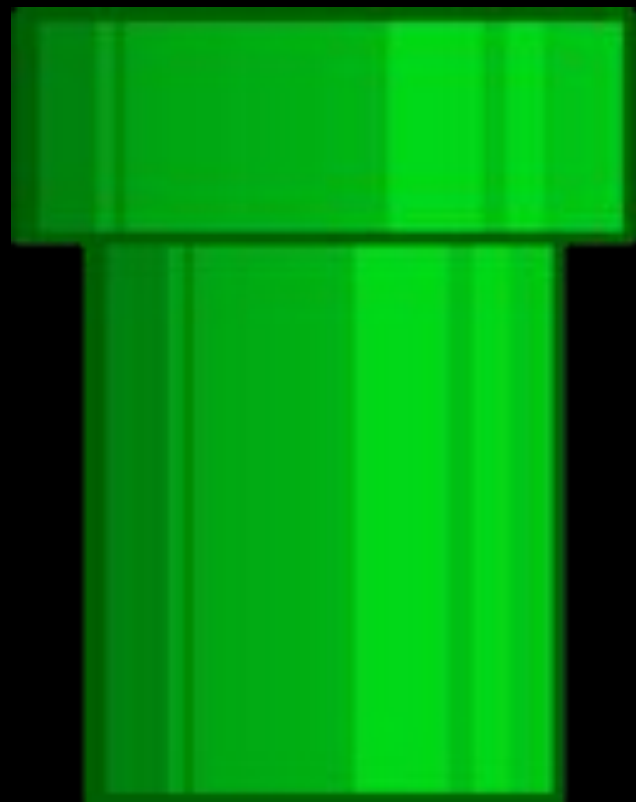
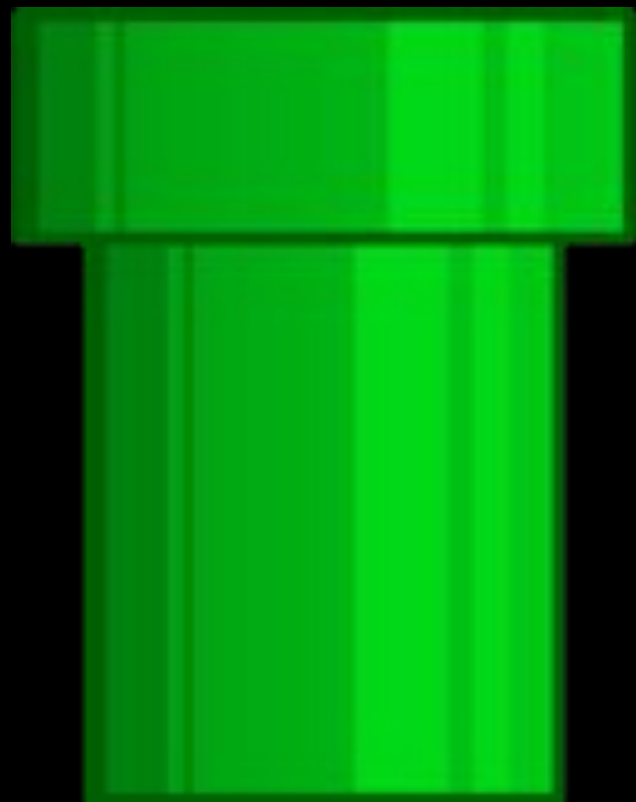
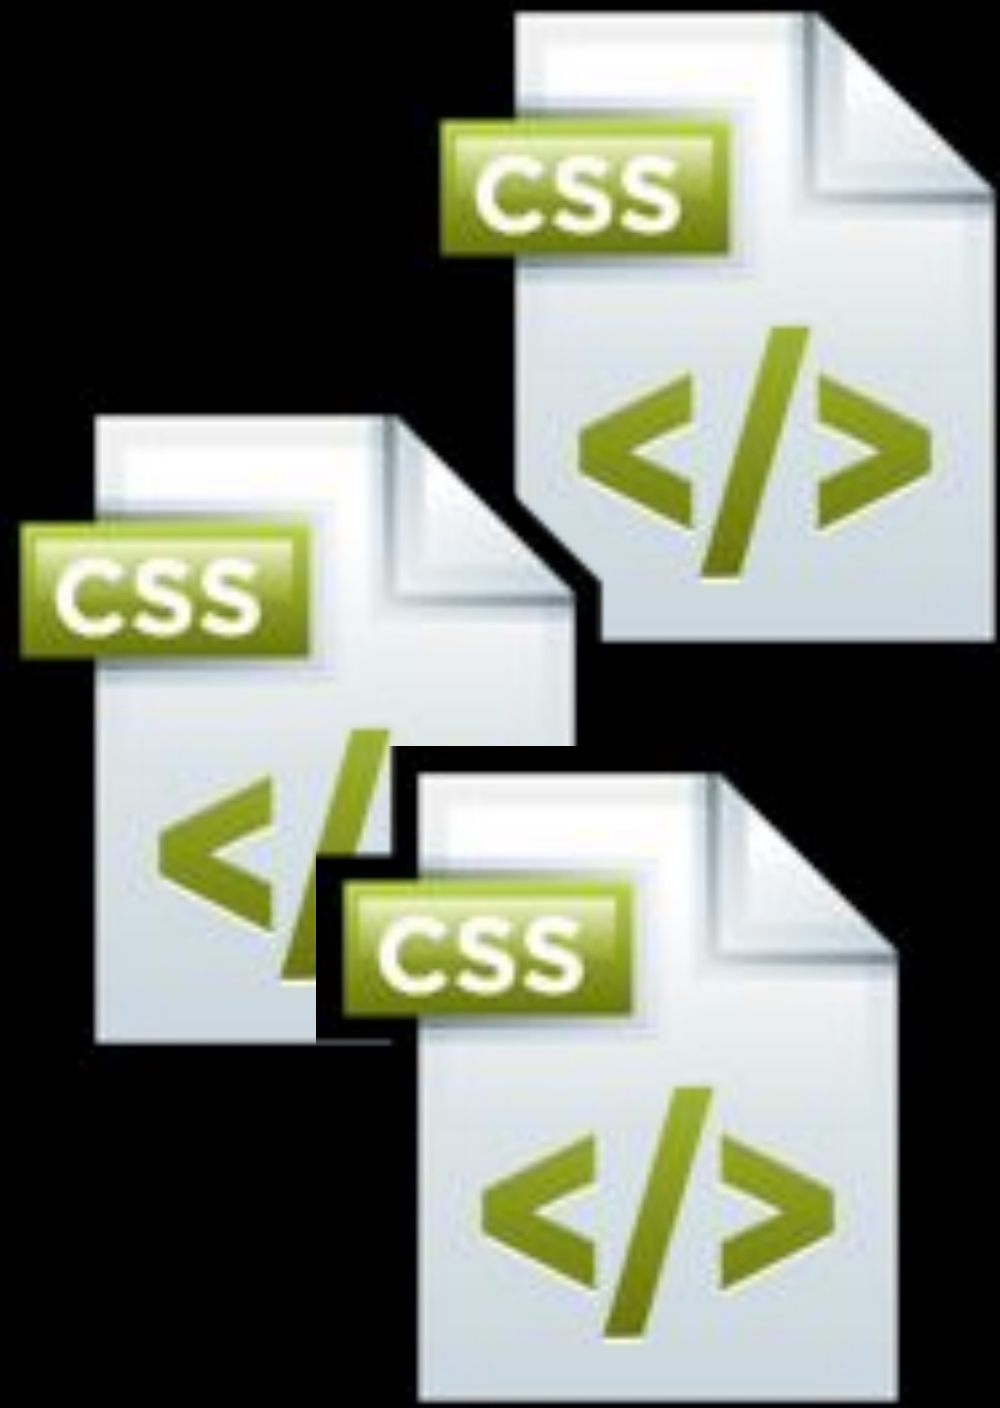


- **Prioritisation**
- **Server Push**
- **Header Compression**
- **TLS and TCP**

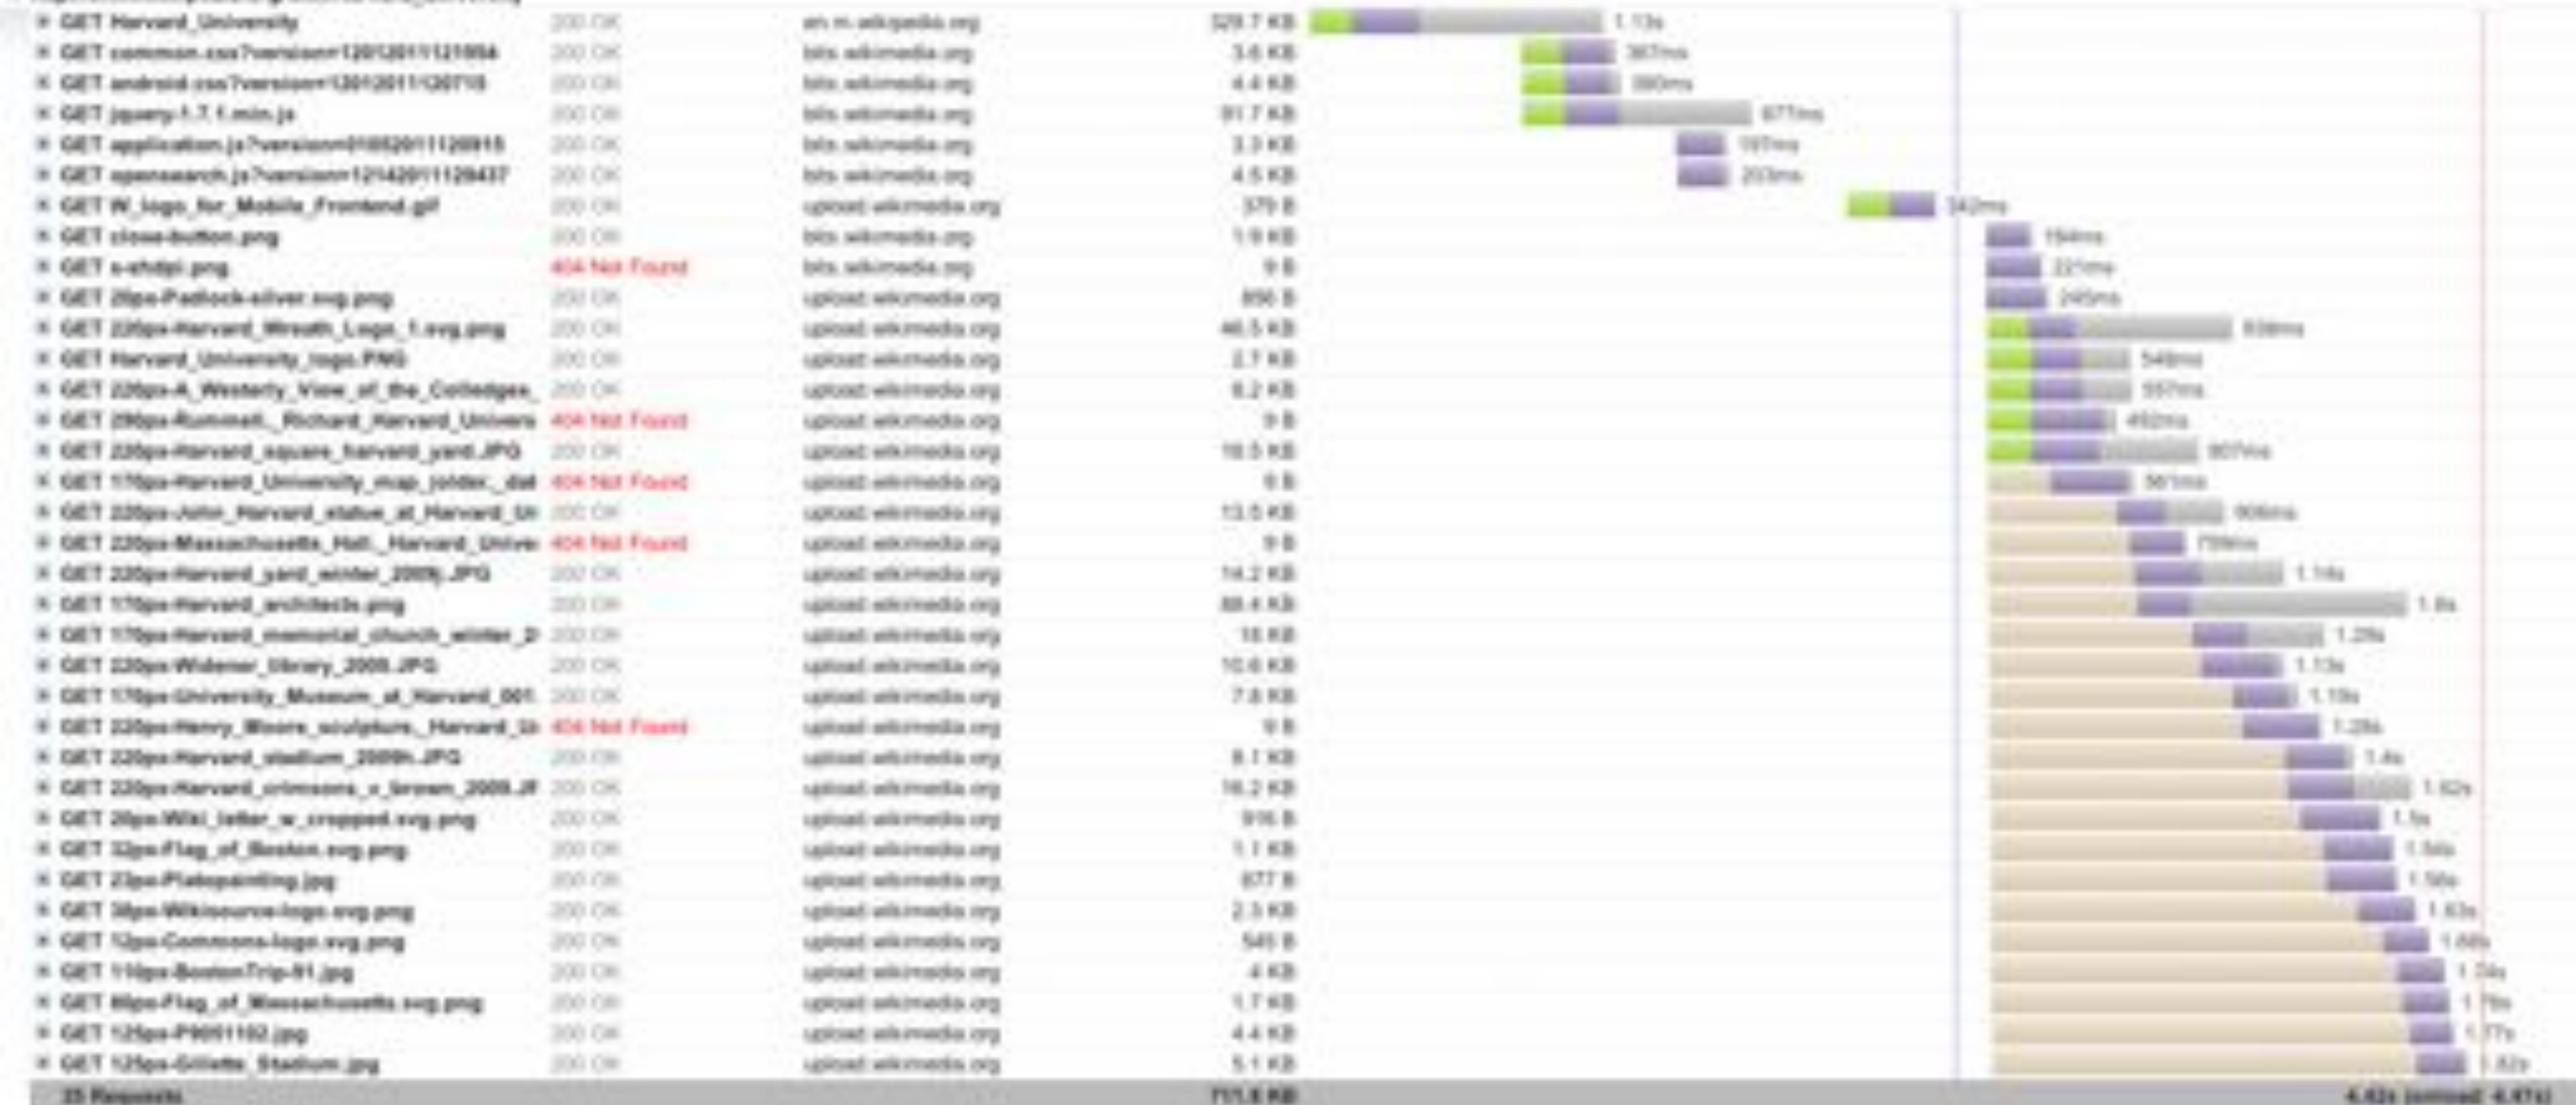


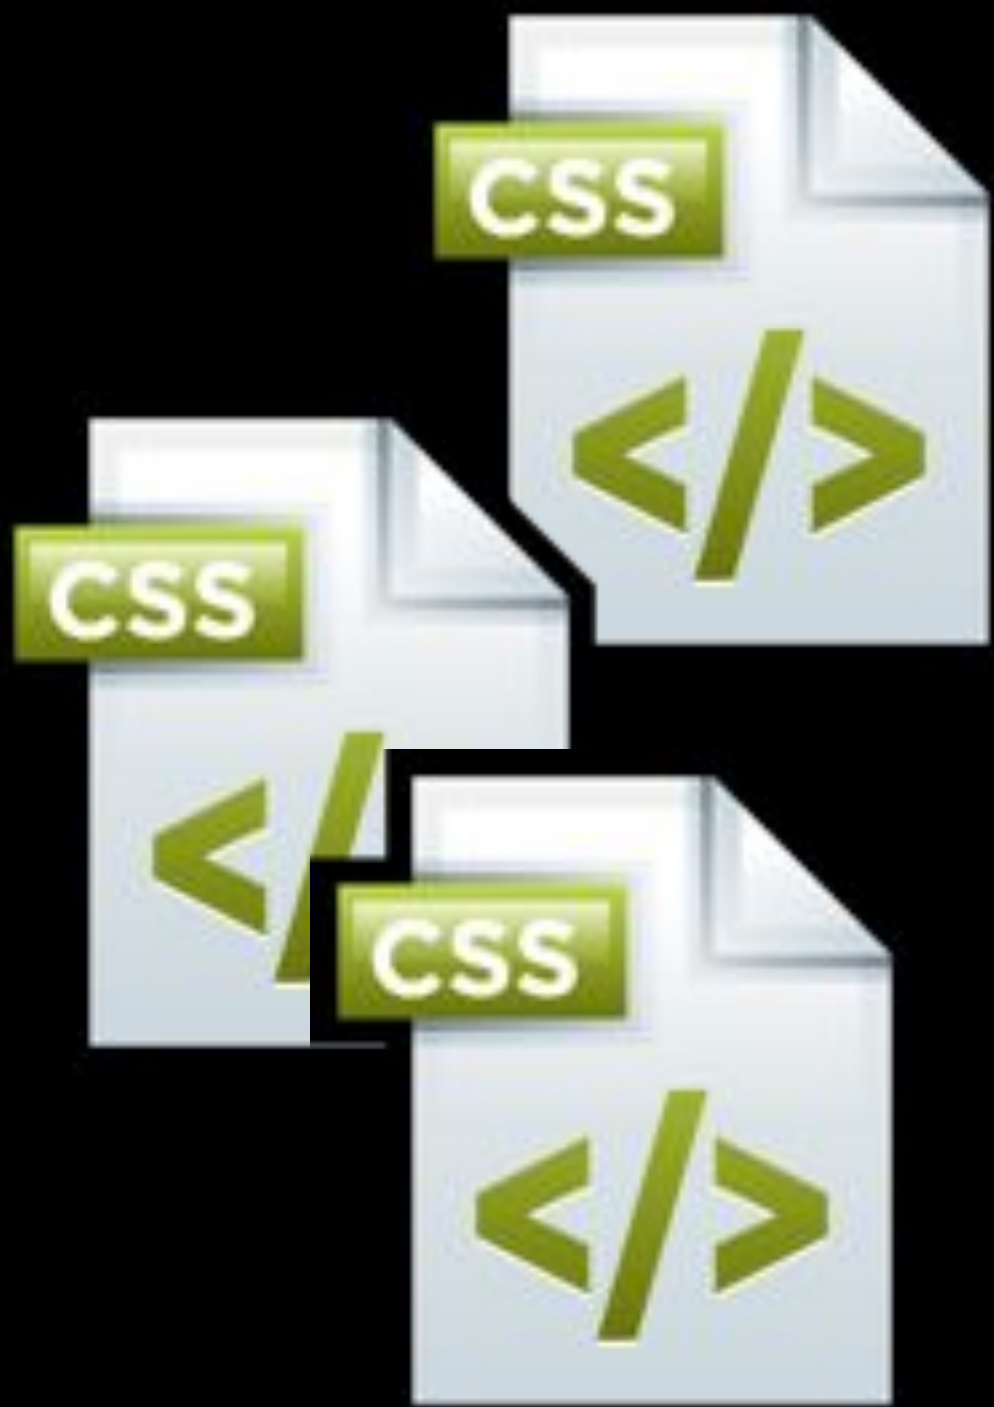
1. Prioritisation

HTTP/1

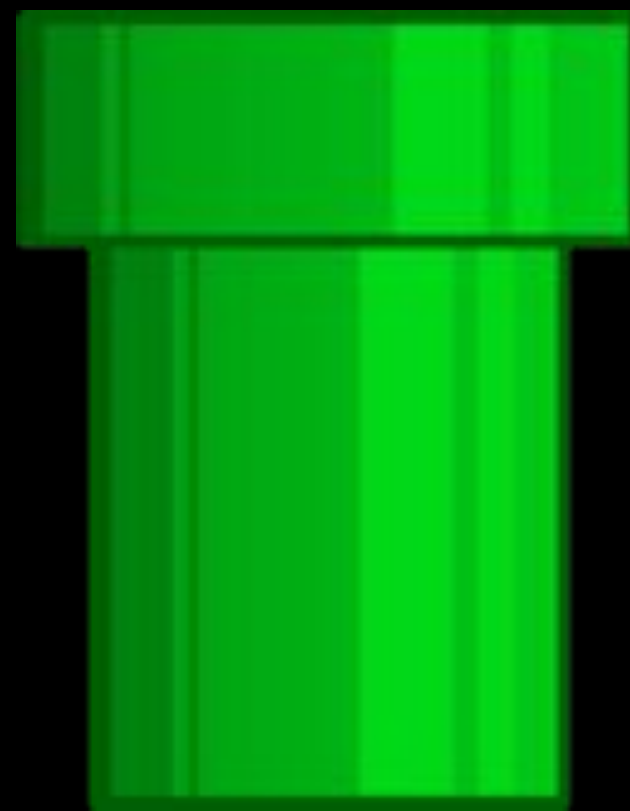


ii http://en.m.wikipedia.org/wiki/Harvard_University





HTTP/2



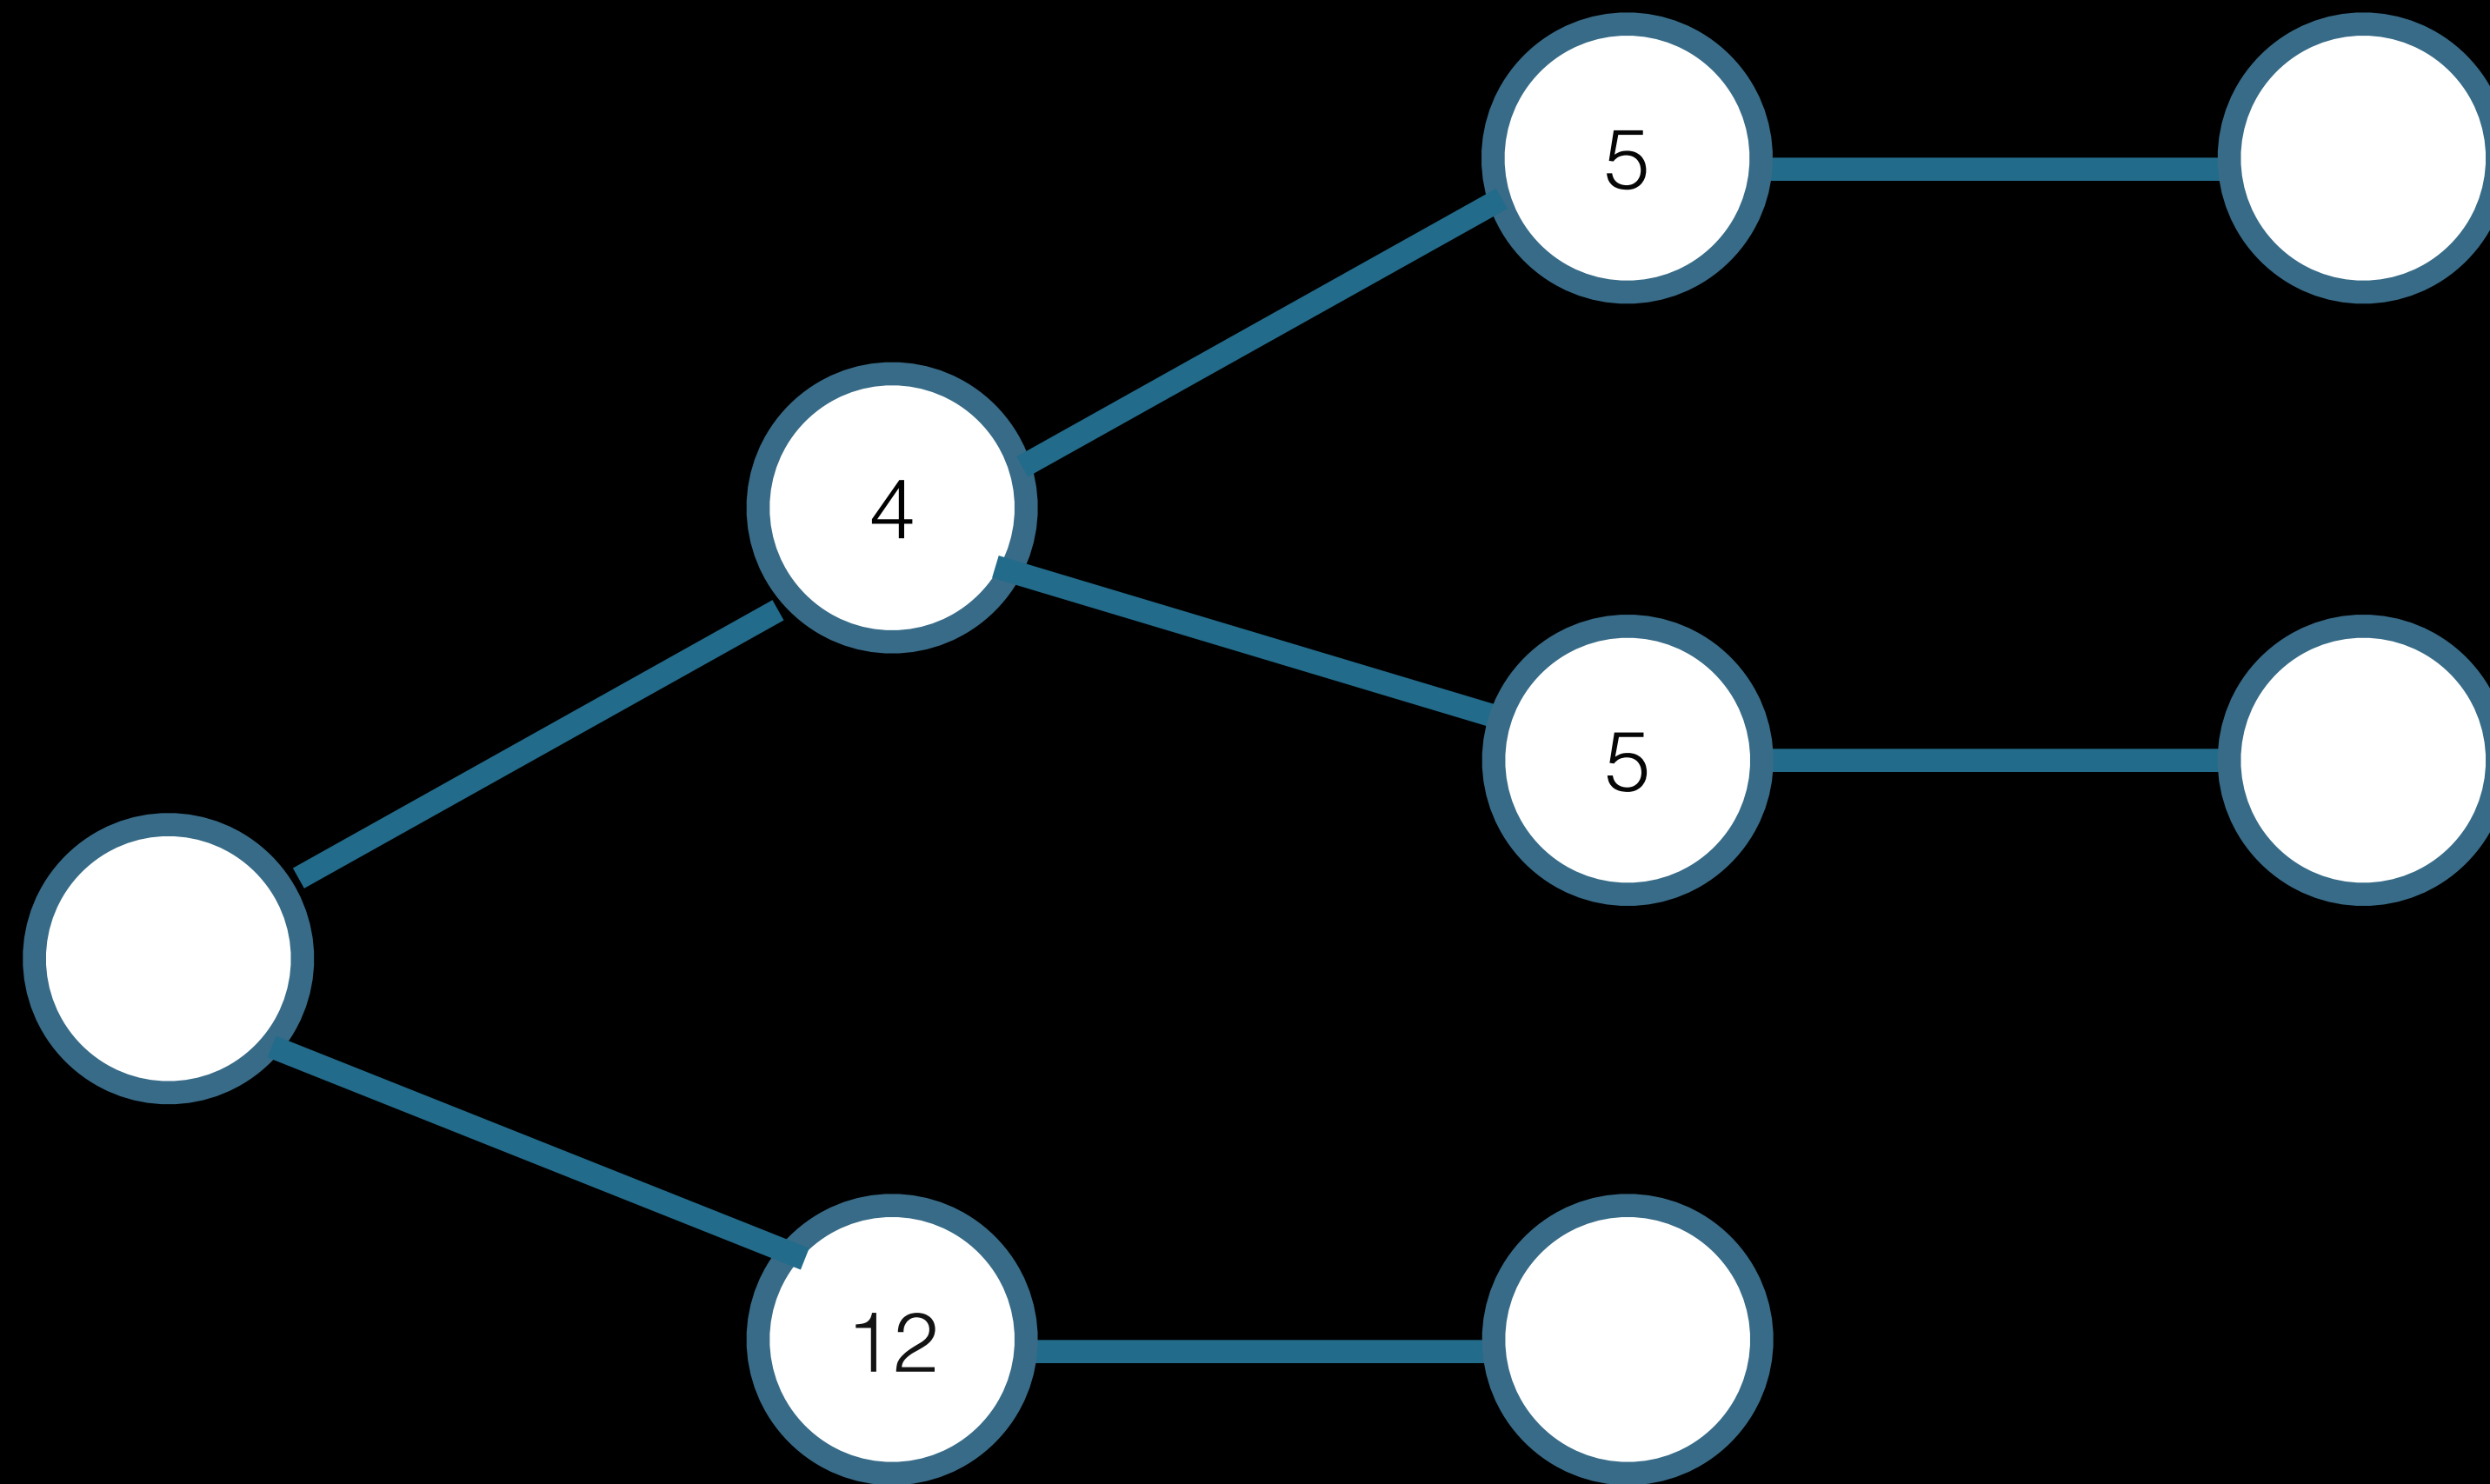
http://en.wikipedia.org/wiki/Harvard_University



- In HTTP/1, Prioritisation is a browser heuristic*
 - “CSS and JS first, then images...”
 - “This connection for that request...”
- In HTTP/2, it’s **hinted by the client, determined by the server.**

* a.k.a. “guessing”

HTTP/2 Priority Hinting





Most Hilarious Videos Of Cats 2014 - Best Funny Kitten Compilation

Advertisement
New Simons Cat
by Simon's Cat

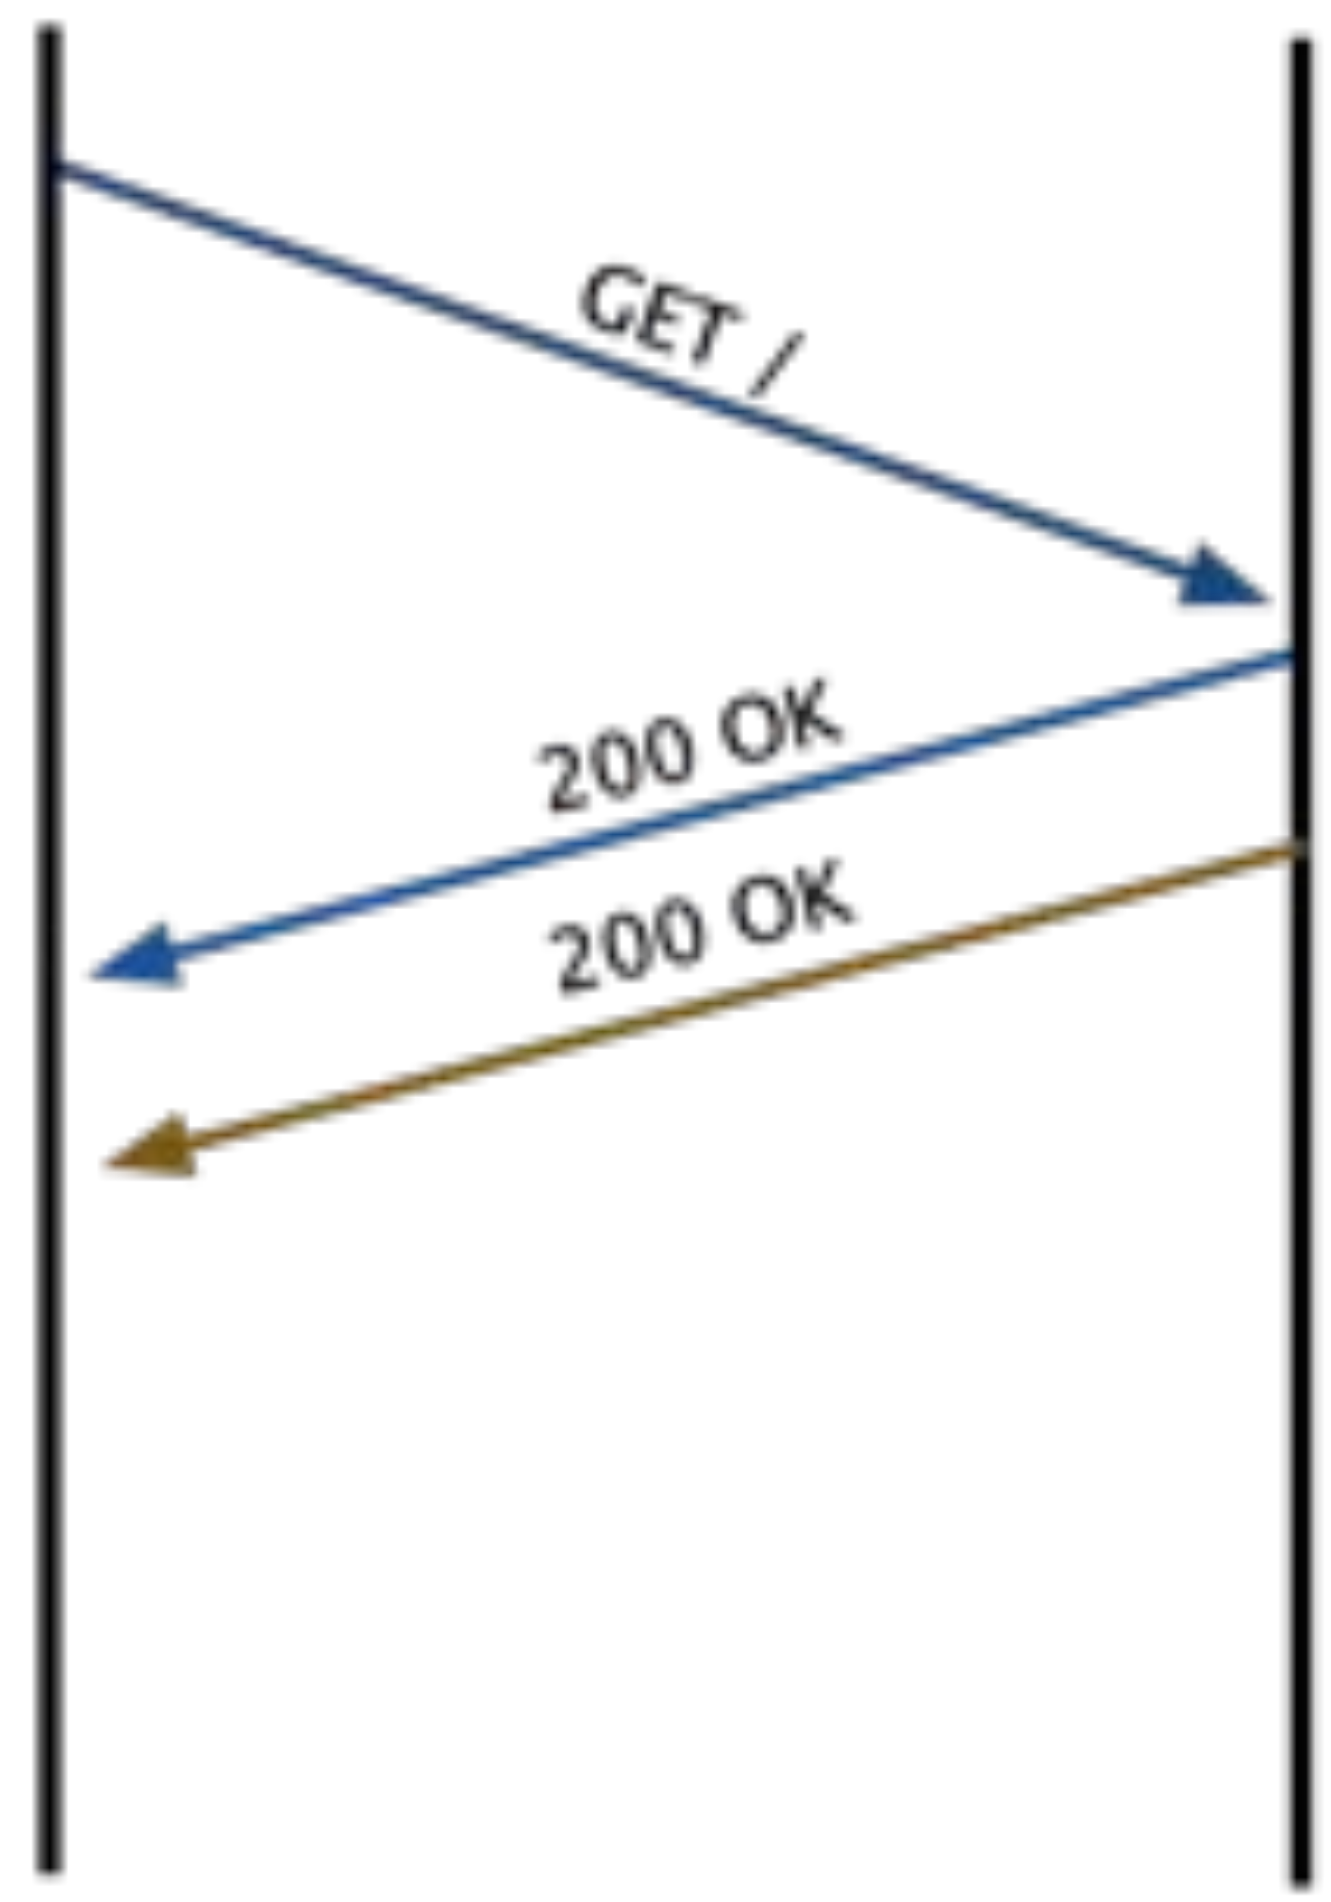
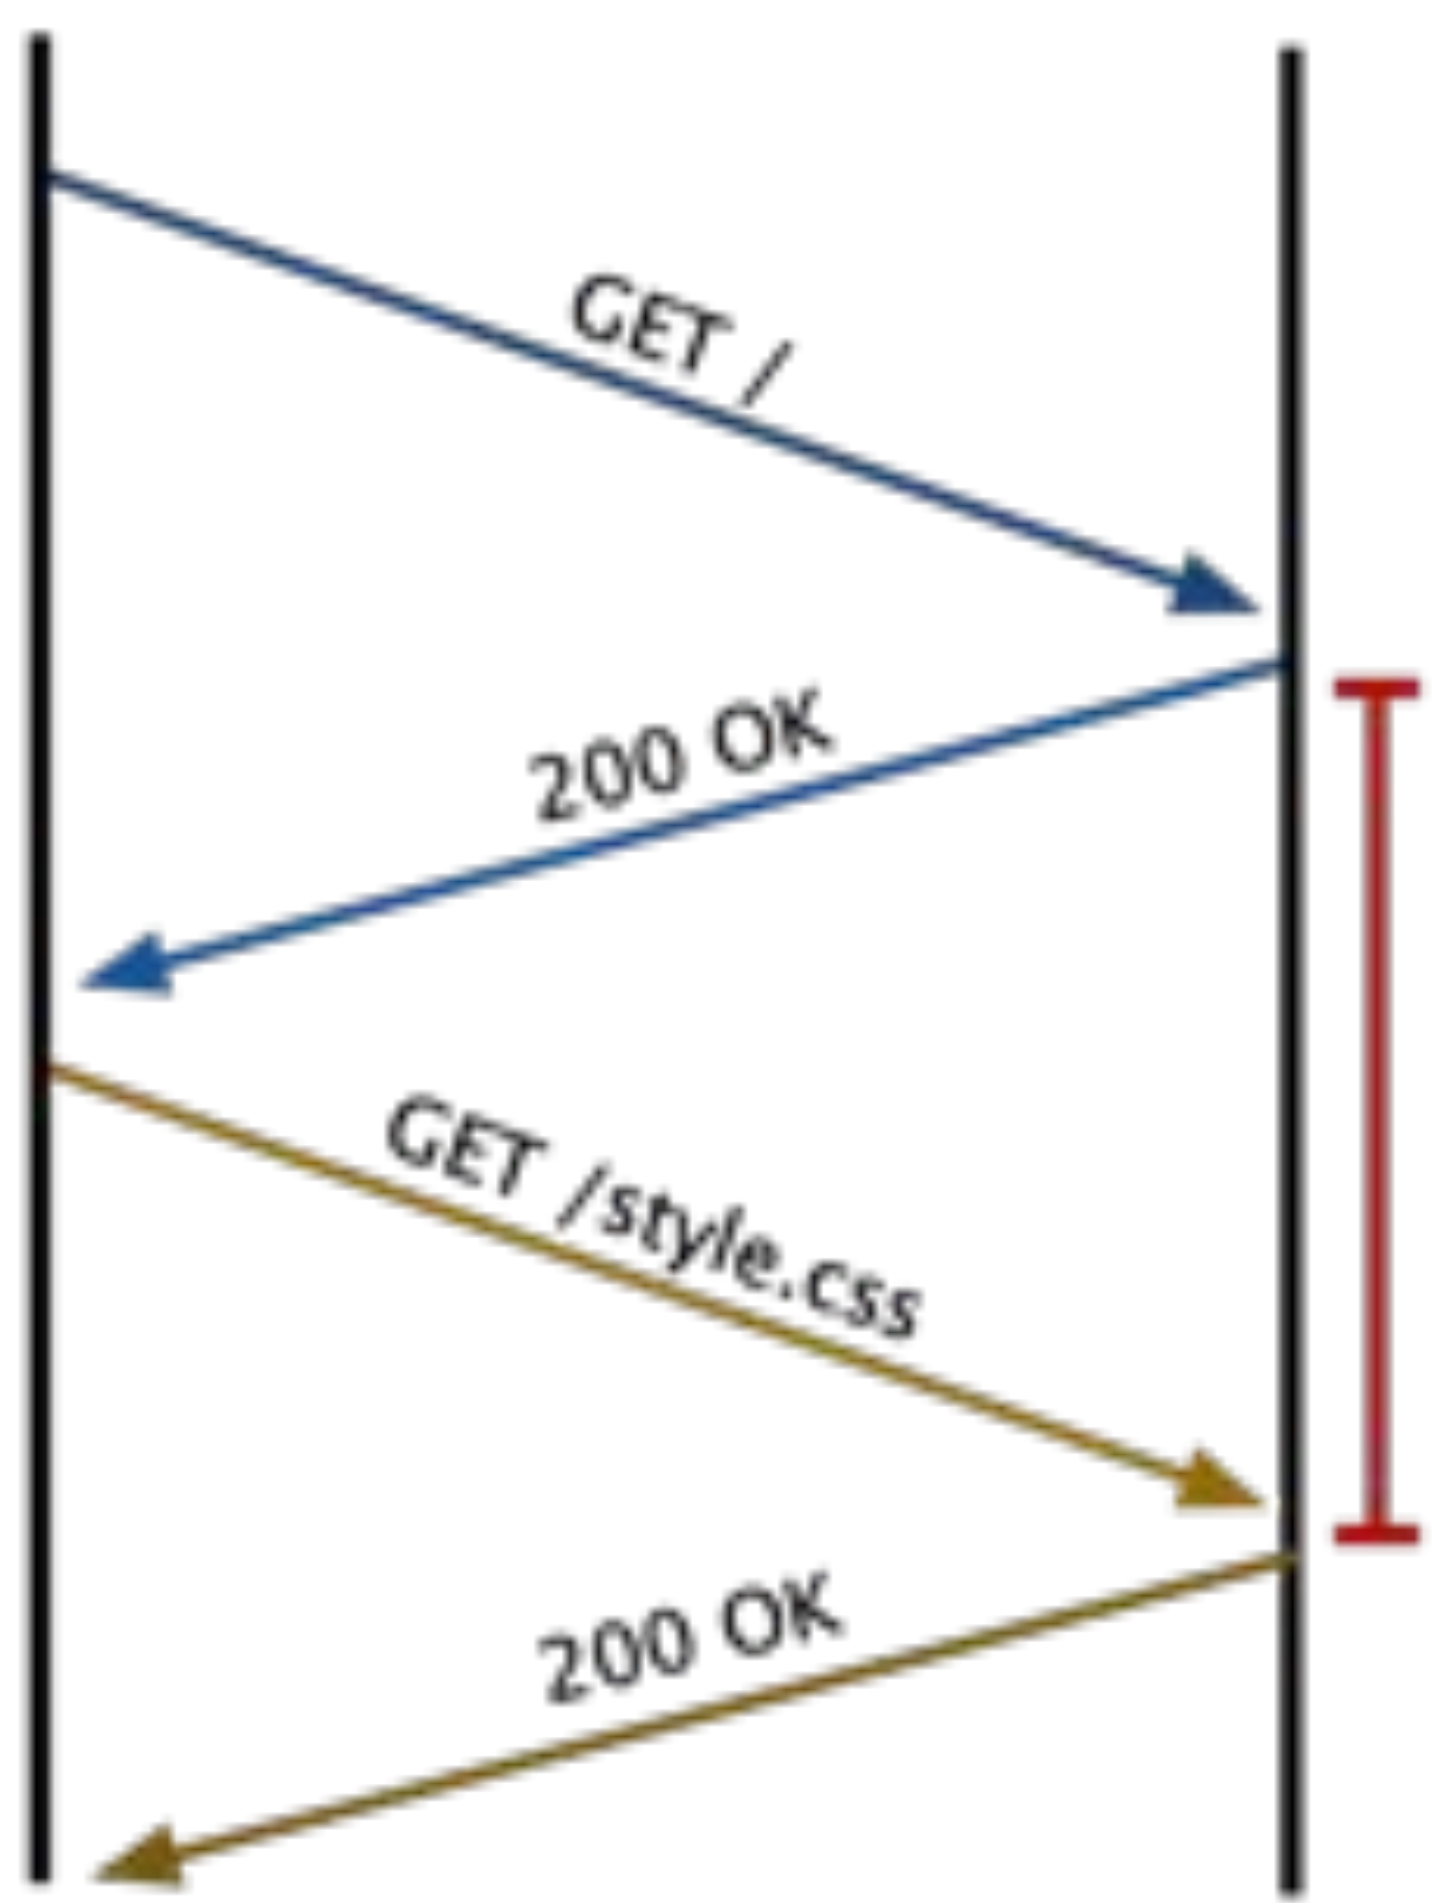
- Servers can do what they like with this information.
- Only one of several potential sources!
- Page analysis, RUM over time, etc.
- No prioritisation performs **worse** than HTTP/1.

Ask Your Implementation

- How do you handle priority by default?
- What APIs do you expose for effecting prioritisation?
 - How do you handle reprioritisation?
- Do you queue in user space and use TCP_NOTSENT_LOWAT*?

* <https://insouciant.org/tech/prioritization-only-works-when-theres-pending-data-to-prioritize/>

2. Server Push



Benefits of Push

- Avoid a RT without sacrificing **Resource Granularity!**
 - Better cache efficiency
 - Reduced parse / blocking
 - Load what you need
 - Modularity
- Client can refuse push with RST_STREAM

When do I Push?

- **Easy answer:** when you previously inlined / concatenated
- **Creative answer:** when you want to overload it for async data
- **Real answer:** we need research, metrics and tools!
 - *speculative push* is likely a very different beast
 - how will intermediaries handle server push?

3. Header Compression

HPACK - Header Compression for HTTP/2

draft-ietf-httpbis-header-compression-latest

Abstract

This specification defines HPACK, a compression format for efficiently representing HTTP header fields, to be used in HTTP/2.

Editorial Note (To be removed by RFC Editor)

Discussion of this draft takes place on the HTTPBIS working group mailing list (ietf-http-wg@w3.org), which is archived at <https://lists.w3.org/Archives/Public/ietf-http-wg/>.

Working Group information can be found at <http://tools.ietf.org/wg/httpbis/>; that specific to HTTP/2 are at <http://http2.github.io/>.

The changes in this draft are summarized in [Appendix D.1](#).

Status of This Memo

GET / HTTP/1.1

Host: www.etsy.com

User-Agent: Mozilla/5.0 (Macintosh; Intel Mac OS X 10_8_2) AppleWebKit/536.26.14 (KHTML, like Gecko) Version/6.0.1 Safari/536.26.14

Accept: text/html,application/xhtml+xml,application/xml;q=0.9,*/*;q=0.8

DNT: 1

Accept-Language: en-us

Accept-Encoding: gzip, deflate

Cookie: uaid=uaid%3DVdhk5W6sexG-_Y7ZBeQFa3cq7yMQ%26_now%3D1325204464%26_slit%3Ds_LCLVpU%26_kid%3D1%26_ver%3D1%26_mac%3D1VnlM3hMdb3Cs3hqMVuk_dQEixsqQzUlNYCs9H_Kj8c.; user_prefs=1&2596706699&q0tPzMlJLaoEAA==

Connection: keep-alive

GET /assets/dist/js/etsy.recent-searches.20121001205006.js HTTP/1.1

Host: www.etsy.com

User-Agent: Mozilla/5.0 (Macintosh; Intel Mac OS X 10_8_2) AppleWebKit/536.26.14 (KHTML, like Gecko) Version/6.0.1 Safari/536.26.14

Accept: */*

DNT: 1

Referer: http://www.etsy.com/

Accept-Language: en-us

Accept-Encoding: gzip, deflate

Cookie: autosuggest_split=1; etala=111461200.1476767743.1349274889.1349274889.1349274889.1.0;

etalb=111461200.1.10.1349274889; last_browse_page=%2F; uaid=uaid%3DVdhk5W6sexG-_Y7ZBeQFa3cq7yMQ%26_now%3D1325204464%26_slit%3Ds_LCLVpU%26_kid%3D1%26_ver%3D1%26_mac

%3D1VnlM3hMdb3Cs3hqMVuk_dQEixsqQzUlNYCs9H_Kj8c.; user_prefs=1&2596706699&q0tPzMlJLaoEAA==

Connection: keep-alive

- Simple replacement strategies work surprisingly well
 - e.g., LRU
- Tuning is more important when you mux clients
 - e.g., in a load balancer / reverse proxy
 - look at state commitment as well
- In the long run, **changing how we use HTTP headers will improve compression**

4. TLS and TCP

TLS has exactly one performance problem: it is not used widely enough.

Everything else can be optimized.

Data delivered over an unencrypted channel is insecure, untrustworthy, and trivially intercepted. We owe it to our users to protect the security, privacy, and integrity of their data — all data must be encrypted while in flight and at rest. Historically, concerns over performance have been the common excuse to avoid these obligations, but today that is a false dichotomy. Let's dispel some myths.

CPU & latency costs

The process of establishing and communicating over an encrypted channel introduces additional computational costs. First, there is the asymmetric (public key) encryption used during the TLS handshake. Then, once a shared secret is established,

```
# upgrade to latest
$> openssl version
OpenSSL 1.0.1j 6 Aug 2014
```

```
# run benchmarks
$> openssl speed sha
```


- initcwnd of 10
- Turn off tcp_slow_start_after_idle
- Experiment with congestion control algorithms

A photograph of a control room or operations center. In the foreground, there is a long console with numerous buttons, switches, and lights. The background features a large wall of windows with a grid pattern. The ceiling has several rows of cylindrical lights. The overall scene is brightly lit, suggesting a busy operational environment.

How does {SPDY, HTTP/2} affect Ops?



- **TLS**
- **Load Balancing/Failover**
- **DoS**
- **Tools**
- **Metrics**
- **Transition Strategies**

1. TLS

- HTTP/2 does **not** require https://
- However, Chrome and Firefox have said they won't do plaintext http://
 - ... Firefox is experimenting with "Opportunistic Security" for http://
- If you want to get the most users onto HTTP/2, **it has to be https://**
(for now)



- Minimum TLS 1.2
- SNI required
- Renegotiation during HTTP prohibited
- Ephemeral Key Exchange required (for forward secrecy)
- AEAD ciphers effectively required
- TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256 w/ P256



2. Failover and Load Balancing

- HTTP/2 flows are **longer-lived**
- So, you can't assume they'll go away soon
- This changes load balancing and failover strategies

- **GOAWAY** allows you to gracefully shut down a connection
- Existing requests drain

- **Alternative Services** acts like CNAME at the HTTP layer
- “Spool up a connection at host:port and talk to it like it’s this origin.”
- It’s graceful; should introduce no latency / pauses
- Load balancing, geo optimisation, taking server out of rotation
- Browser support coming (hopefully)

3. DoS

TCP Connections

- HTTP/2 uses **less connections** by a factor of 4x - 8x!
- Those connections are a **lot** more active, well-utilised
- If you think you're under attack, you've got options
 - Reduce SETTINGS_MAX_CONCURRENT_STREAMS
 - Flow control them
 - GOAWAY

Memory

- **Header compression** = More state per connection
 - default header compression state = **4k**
 - state commitment can be tuned using `SETTINGS_HEADER_TABLE_SIZE`
- **Buffering** e.g., by intermediaries
 - can be controlled by flow control
 - (stream-level and connection-level)

CPU

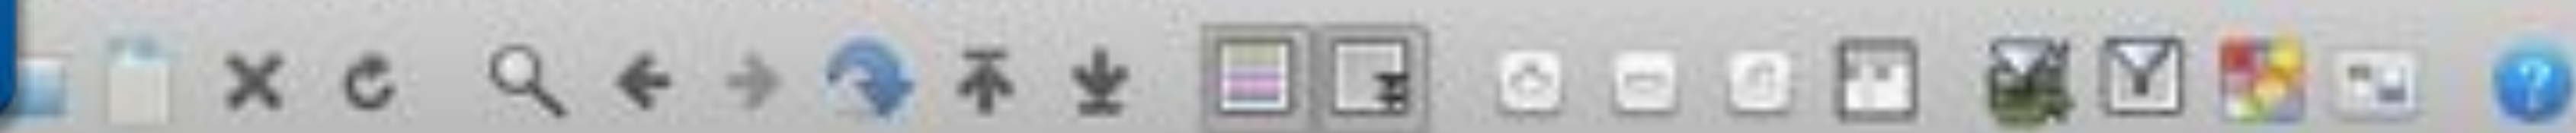
- Binary format is **easier to parse**
- **HPACK** is very low overhead (less than gzip)
 - can be dialled down if need be with `SETTINGS_HEADER_TABLE_SIZE`

Intermediaries

- There are some cases where an intermediary can get in trouble
 - Never forward a frame before you have it all!
 - Never forward a header-bearing block before you have it all!
 - ... unless you **really** trust the sender.



4. Tools



Filter: http2 Expression... Clear Apply Save

No.	Time	Source	Destination	Protocol	Length	Data Rate	Info
4	0.000375000	:::1	:::1	HTTP2	101		SETTINGS
6	0.000619000	:::1	:::1	HTTP2	197		Magic, SETTINGS, HEADERS
8	0.000738000	:::1	:::1	HTTP2	95		SETTINGS
10	0.001388000	:::1	:::1	HTTP2	21974		SETTINGS, HEADERS, DATA, DATA, I
11	0.001412000	:::1	:::1	HTTP2	984		DATA
13	0.002756000	:::1	:::1	HTTP2	155		SETTINGS, HEADERS
14	0.004259000	:::1	:::1	HTTP2	277		HEADERS, HEADERS, HEADERS, HEADERS
16	0.009281000	:::1	:::1	HTTP2	43365		SETTINGS, HEADERS, HEADERS, HEADERS
17	0.013286000	:::1	:::1	HTTP2	99		WINDOW UPDATE

[Header Length: 208]
▶ Header: :status: 200
▶ Header: server: nghttpd nghttp2/0.5.2-DEV
▶ Header: content-length: 22617
▶ Header: cache-control: max-age=3600
▶ Header: date: Sat, 02 Aug 2014 10:50:25 GMT
▶ Header: last-modified: Sat, 02 Aug 2014 07:58:59 GMT
Padding: <MISSING>
▶ Stream: DATA, Stream ID: 1, Length 4096
▶ Stream: DATA, Stream ID: 1, Length 4096
▶ Stream: DATA, Stream ID: 1, Length 4096
▶ Stream: DATA, Stream ID: 1, Length 4096
▶ Stream: DATA, Stream ID: 1, Length 4096

0010	00 00 55 a0 06 40 00 00 00 00 00 00 00 00 00 00
0020	00 00 00 00 00 01 00 00 00 00 00 00 00 00 00 00
0030	00 00 00 00 00 01 0b b8 dc bc d6 8f e
0040	23 fa 80 18 01 56 55 a8 00 00 01 01 e
0050	aa 63 00 89 aa 62 00 00 00 04 01 00 e
0060	00 58 01 04 00 00 00 01 88 76 93 aa e
0070	52 a9 a7 4a 6b 13 00 5d b5 c4 b5 fc 1
0080	84 10 9c 0b bf 58 89 a4 7e 56 1c c5 e
0090	61 96 dc 34 fd 28 01 29 0d b3 28 20 e
00a0	8d 82 e0 9b 53 16 8d ff 6c 96 dc 34 f
00b0	0d b3 28 20 05 a5 00 ed c6 de b8 db e
00c0	00 18 00 00 00 00 00 00 01 0a 0a 3c 2
00d0	5d 50 5a 45 7a 68 7d 6d 6c 3a 0a 3c 7

Frame (21974 bytes) Decompressed Header (208 bytes)

curls those URLs

?

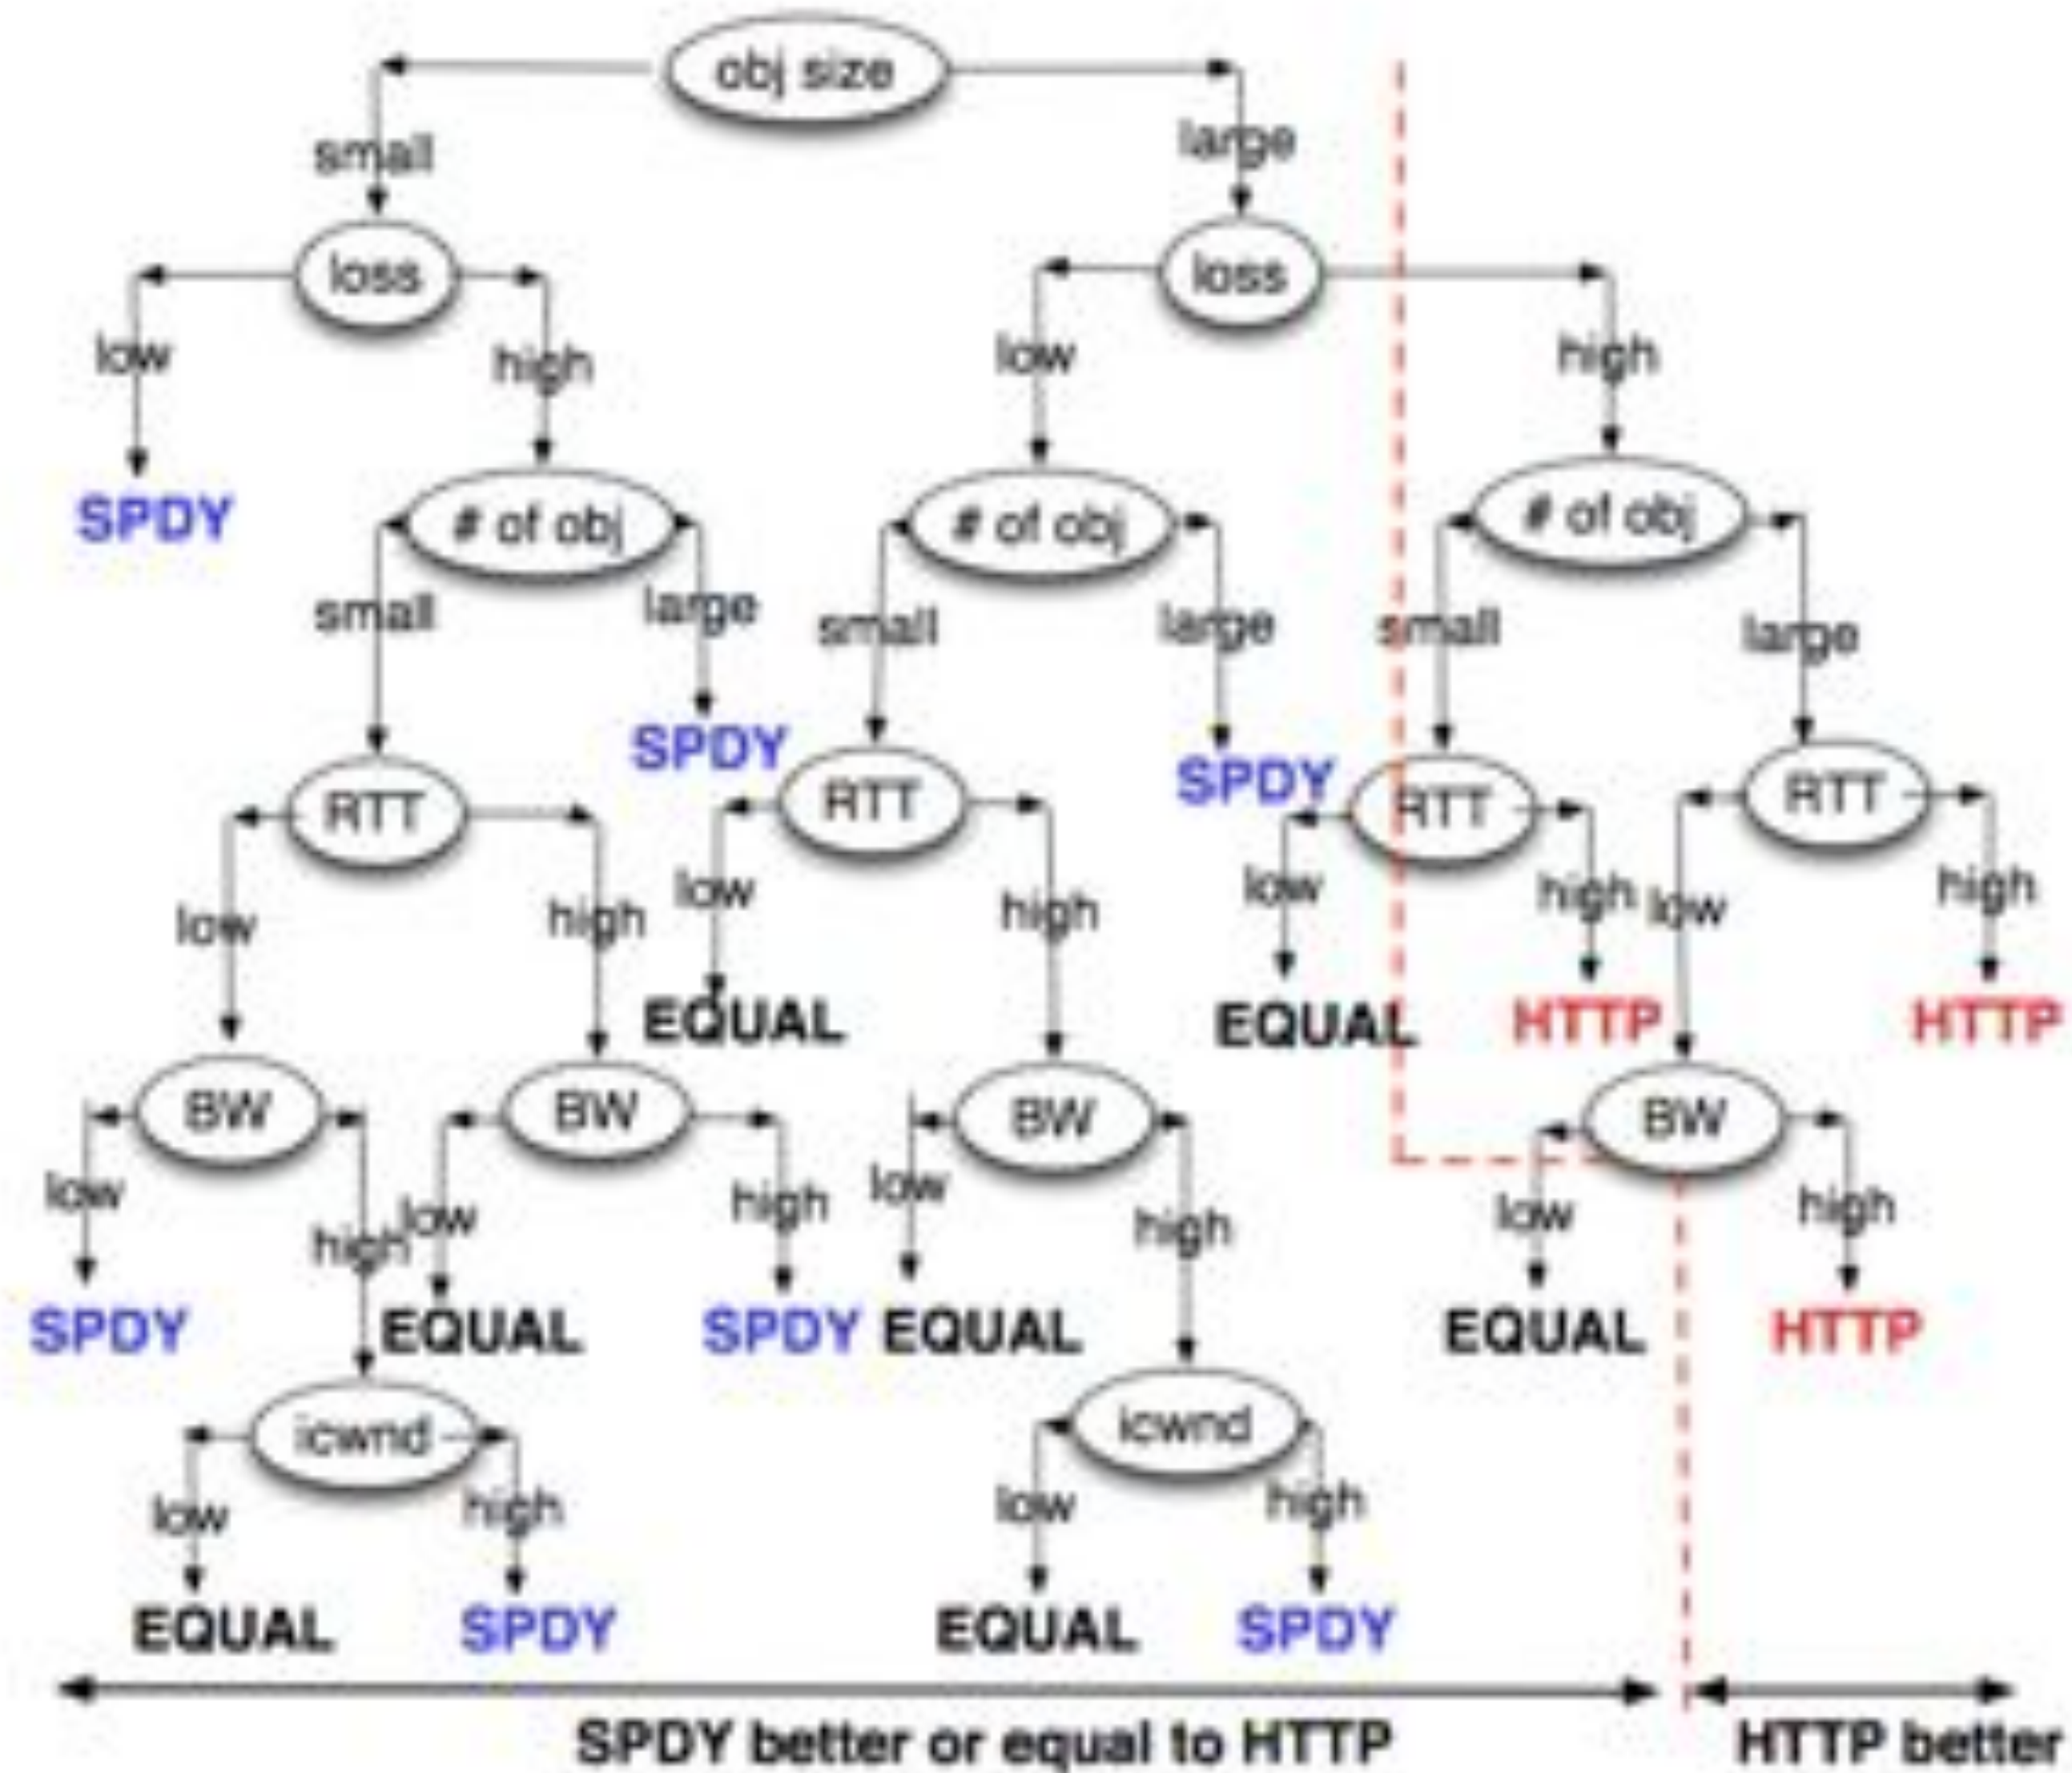
5. Metrics



- **Connection Utilisation**
 - idle periods, concurrency, reset streams, protocol errors
- **HPACK efficiency**
 - hit rates, memory consumption, lost opportunities
 - both directions!
- RUM RUM **RUM!**

6. Transition Strategies






“How Speedy is SPDY?” — Wang, Balasubramanian, Krishnamurthy and Wetherall

Un-Hacking Your Site

Spriting  Separate Resources

Inlining  Server Push

Sharding  Single Host

Concatenation  Separate Resources

Why Un-Hack?

- CSS Spriting delays image downloads (**indirection**)
- Concatenation / inlining / spriting reduce **cache efficiency**
- Concatenation / inlining / spriting encourages **wasted download**
- JS concatenation increases **parse/load overhead**
- Sharding **reduces header compression efficiency, tcp flow gains**

When to Un-Hack?

- Lots of choices:
 - When x% of your traffic supports HTTP/2
 - Dynamically hack per connection
 - Gradual rollout
 - Just Do It

- Switch to HTTPS first
- De-[sprite, concatenate, shard, inline]
- Browser rollout period is a unique opportunity
- Collect metrics!

THE TASMANIAN GOVERNMENT HYDRO-ELECTRIC POWER

The Master-key to Industry

The year Millennium celebrations for Centenary—Tasmania, has passed that celebrated for Centenary in 1901, and daily celebrates her government as the

**GREATEST
HYDRO-ELECTRIC
STATE
IN THE
COMMONWEALTH**

The Hydro-Electric Commission
1901—1902 TASMANIA
The Hydro-Electric Commission
1901—1902 TASMANIA
The Hydro-Electric Commission
1901—1902 TASMANIA

TASMANIA offers the greatest opportunities in Australia for industrial enterprises with unlimited power for all purposes.

TASMANIA—
Ideal Climate—
Deep Water Ports—
Cheap Factory Sites—
Abundant Fresh Water Supply





Questions?