

Advanced Computer Networking

CYBR 230 – Jeff Shafer – University of the Pacific

HTTP/2

2



HTTP/0.9 (not an RFC: <u>https://www.w3.org/Protocols/HTTP/AsImplemented.html</u>)

- Initial version of HTTP a simple client-server, request-response, telnet-friendly protocol
- Request nature: single-line (method + path for requested document)
- Methods supported: GET only
- Response type: hypertext only
- Connection nature: terminated immediately after the response
- No HTTP headers (cannot transfer other content type files), No status/error codes, No URLs, No versioning



HTTP/1.0 – RFC 1945 (May 1996)

- Browser-friendly protocol
- Provided header fields including rich metadata about both request and response (HTTP version number, status code, content type)
- Response: not limited to hypertext (Content-Type header provided ability to transmit files other than plain HTML files e.g. scripts, stylesheets, media)
- Methods supported: GET, HEAD, POST
- Connection nature: terminated immediately after the response



HTTP/1.1 – RFC 2068 (January 1997)

- Performance optimizations and feature enhancements
 - Persistent and pipelined connections
 - Chunked transfers
 - Compression/decompression
 - Virtual hosting (a server with a single IP address hosting multiple domains)
- Methods supported: GET, HEAD, POST, PUT, DELETE, TRACE, OPTIONS
- Connection nature: long-lived

THE WORLD IS CHANGING

Motivations

What's wrong with HTTP/1.1?

- Complicated specification
- Many options / incomplete implementations
- Performance limitations
- Why does it take so long to load a web page?
 - https://bagder.gitbooks.io/http2-explained/en/

6

http archive



The **HTTP Archive** Tracks How the **Web is Built**.

We periodically crawl the top sites on the web and record detailed information about fetched resources, used web platform APIs and features, and execution traces of each page. We then crunch and analyze this data to identify trends — learn more about our methodology.

View Reports

Featured Report State of the Web

This report captures a long view of the web, including the adoption of techniques for efficient network utilization and usage of web standards like HTTPS.



View The Report

Total Kilobytes

The sum of transfer size kilobytes of all resources requested by the page.

See also: Page Weight



https://httparchive.org/reports/state-of-the-web

Total Requests

The number of resources requested by the page.

See also: Page Weight



Advanced Computer Networking

https://httparchive.org/reports/state-of-the-web

Fall 2018

TCP Connections Per Page

The number of TCP connections per page.



https://httparchive.org/reports/state-of-the-web

Fall 2018

Performance

- → Many files (Median of ~75 in 2018)
 - ↗ Some files are large, some files are small
- Many TCP sockets (Median of ~19 in 2018)
 - More than one file per socket
 - Resource penalty + time penalty to just open yet more sockets
- Problems
 - A Latency between requests
 - HTTP Pipelining typically disabled by default
 - **↗** Head-of-line blocking
 - Multiple requests over same socket, large file blocks small file
 - Not solved by HTTP Pipelining

HTTP Performance Hacks



Sprite Sheets

- Instead of sending many tiny images, send one big image and use JavaScript/CSS to pull out the desired pieces
- Inefficient if only a few pieces are needed

HTTP Performance Hacks

```
.icon1 {
```

```
background:
url(data:image/png;base64,<data>)
no-repeat;
}
.icon2 {
    background:
url(data:image/png;base64,<data>)
no-repeat;
```

Inlining

Embed data inside the CSS file instead of as separate files

Concatencation

 Combine multiple JS
 / CSS files into megafiles before sending to client

• 200	GE	Г 🗾 174.јрд		W	.cdn-expressen.se	jpeg	6.14 KB	→ 105 ms
• 200	GE	Г	174.ipa	V.	dn-expressen.se	jpeg	4.19 KB	→ 172 ms
• 200			1		dn-expressen.se	jpeg	4.48 KB	→ 223 ms
• 200		z.c	an-expressen.se		dn-expressen.se	jpeg	4.58 KB	→ 173 ms
• 200			1		dn-expressen.se	jpeg	35.18 KB	→ 56 ms
• 200		x.c	.con-expressen.se		dn-expressen.se	jpeg	12.97 KB	→ 165 ms
• 200	• 200		under ausernannen an		dn-expressen.se	jpeg	4.83 KB	→ 56 ms
• 200	• 200)		an-expressen.se		dn-expressen.se	jpeg	9.54 KB	→ 228 ms
• 200	200				dn-expressen.se	jpeg	182.50 KB	→ 285 ms
• 200	200 W		.con-expressen.se		dn-expressen.se	jpeg	5.66 KB	→ 104 ms
• 200	200 200 y				dn-expressen.se	jpeg	12.24 KB	→ 287 ms
• 200			.con-expressen.se		dn-expressen.se	jpeg	6.85 KB	→ 225 ms
• 200	200		z odn ovnroccon co		dn-expressen.se	jpeg	7.50 KB	→ 173 ms
• 200	2.0		un-expressen.se		dn-expressen.se	gif	2.85 KB	→ 227 ms
• 200			de evereses se		dn-expressen.se	jpeg	50.87 KB	→ 188 ms
• 200		w.c	un-expressen.se		dn-expressen.se	jpeg	6.65 KB	→ 55 ms
• 200	UE		zos.jpg	у.	dn-expressen.se	jpeg	6.09 KB	→ 196 ms
• 200	GE	Г	540.jpg	z.	.cdn-expressen.se	jpeg	16.14 KB	→ 67 ms
• 200	GE	r	540.jpg	W	.cdn-expressen.se	jpeg	19.89 KB	→ 112 ms
• 200	GE	г	274.jpg	z.	cdn-expressen.se	jpeg	5.03 KB	→ 55 ms
• 200	GE	Г	540.jpg	W	.cdn-expressen.se	jpeg	21.27 KB	→ 108 ms
• 200	GE	Г	540.jpg	x	cdn-expressen.se	jpeg	5.43 KB	→ 237 ms
• 200	GE	Г	174.jpg	у.	cdn-expressen.se	jpeg	6.08 KB	→ 169 ms
• 200	GE	Г	174.jpg	W	.cdn-expressen.se	jpeg	5.62 KB	→ 105 ms
• 200	GE	Г	540.jpg	×.	cdn-expressen.se	jpeg	20.32 KB	→ 241 ms
• 200	GE	Г	174.jpg	z.	cdn-expressen.se	jpeg	6.66 KB	→ 55 ms
• 200	GE	Г	540.jpg	×.	cdn-expressen.se	jpeg	11.13 KB	→ 237 ms
• 200	GE	Г	265.jpg	W	.cdn-expressen.se	jpeg	5.20 KB	→ 111 ms
• 200	GE	Г	265.jpg	X.	cdn-expressen.se	jpeg	6.93 KB	→ 288 ms
• 200	GE	Г	265.jpg	×.	cdn-expressen.se	Jpeg	12.09 KB	→ 249 ms
• 200	GE	Г	265.jpg	Ζ.	cdn-expressen.se	Jpeg	5.92 KB	→ 167 ms
• 200	GE	T	original.jpg	у.	cdn-expressen.se	Jpeg	64.28 KB	→ 192 ms
0 200	GE		original.jpg	W	.cdn-expressen.se	jpeg	21.88 KB	→ 106 ms
• 200	GE	-	540.jpg	W	.cdn-expressen.se	jpeg	18.// KB	→ 112 ms
200	GE	r I	128.jpg	Ζ.	can-expressen.se	jpeg	3.34 KB	→ 55 ms
200	GE		205.jpg	X.	cun-expressen.se	jpeg	13.00 KB	→ 245 ms
200	GE	r I	205.jpg	у.	can-expressen.se	jpeg	9.19 KB	→ 194 ms
200	GE		540.jpg	W	.con-expressen.se	jpeg	13.13 KB	→ 108 ms
200	GE	T	174.jpg	у.	can-expressen.se	jpeg	5.66 KB	→ 197 ms
200	GE	T	174.jpg	Z.	edn evpressen.se	jpeg	5.50 KB	> > > ms
200	GE	T	174.jpg	W	.cun-expressen.se	jpeg	5.07 KB	→ 111 ms
200	GE	r	174.jpg	Z.	edn expressen.se	jpeg	0.10 KB	-> 29 ms
200	GE	T	174.jpg	у.	odn.expressen.se	jpeg	0.57 KB	> 12 ms
200	GE		L/4.jpg	y.	con-expressen.se	Jpeg	4.58 KB	- 12 ms
200	GE		205.jpg	у.	con-expressen.se	jpeg	11.49 KB	⇒ 1/5 ms

HTTP Hacks

Domain Sharding

- Avoid per-host limits

 on # of connections by
 spreading website
 across many
 "separate" hostnames
- How does the browser (or OS) prioritize these connections?

Advanced Computer Networking

14

◄ IETF HTTPbis working group

- **7** Formed in 2007
- Finalized HTTP/2 RFC 7540 (May 2015)
- Finalize HPACK RFC 7541 (May 2015)
- Heavily derived from Google SPDY work
- https://http2.github.io/
- Google SPDY
 - **Released in Chrome in 2010**
 - Withdrawn in Chrome in 2016

HTTP/2

- Make the web faster *and* eliminate the design hacks
- Maintain high-level compatibility with HTTP/1.1
 - Methods, status codes, URIs, ...
 - ↗ No changes to web pages, web apps, ...
- Reduce Latency -> Reduce Page Load Times
 - Compress headers
 - Server Push
 - No head-of-line blocking
 - **7** Request pipelining
 - **7** Request multiplexing over single connection

HTTP/2 uses binary format 00111 0011100П 101110011000 00011010111

HTTP/2 Binary

- Advantages of Binary
 - Simplifies parsing (Case? Line endings? Whitespace?)
 - Compact representation
 - Simplifies multiplexing and prioritization
 - Impact: Reduces latency, improves throughput
- Disadvantages of Binary
 - Can't fire up Telnet to port 80 and demonstrate HTTP any more (couldn't do that with HTTPS anyway...)



https://developers.google.com/web/fundamentals/performance/http2/

- Stream = Bidirectional flow of data
 - **T** Streams can carry 1 or more messages
 - Binary format allows multiple streams to exist over single TCP connection
- Message = Collection of frames that, combined, form a *request* or *response* message
- Frame = Fundamental unit of communication
 - Note that these are frames inside the TCP connection, not Ethernet frames...

Connection



https://developers.google.com/web/fundamentals/performance/http2/

Frame Format

- Length
- 🛪 Туре
- **7** Flags
- Stream Identifier
- Payload

Frame Types

- **DATA**
- **HEADERS**
- **PRIORITY**
- → PUSH_PROMISE
 - Server push
- ↗ WINDOW_UPDATE
 - **7** Flow control
- (10 frame types in total)

HTTP/2 Multiplexing

HTTP 2.0 connection



https://developers.google.com/web/fundamentals/performance/http2/

- Interleave multiple requests / responses without head-of-line blocking
 - Ability for client or server to prioritize frames
- Single TCP connection to each "origin" (host)
 - Connection overhead and initial latency (especially TLS) is amortized over multiple file transfers
 - More efficient than opening multiple connections

HTTP/2 Compression



https://developers.google.com/web/fundamentals/performance/http2/

- HPACK Header Compression (RFC 7541)
- Fields can be encoded with static Huffman code
 - Reduces transfer size
- Client and server maintain list of previouslyseen fields
 - No need to re-transmit duplicate values (refer to them by index number)
- Operates with bounded memory requirements (i.e. embedded systems)
- Resistant to security flaws (e.g. CRIME – Compression Ratio Info-leak Made Easy)
 - Steal HTTP cookies from TLS connection by observing impact of random payloads on compressed ciphertext length

HTTP/2 Compression

↗ Is header compression worth the code?

- ↗ Typical page
 - ~75 assets
 - → ~1400 bytes per request (referrer tag, cookies, ...)
- **7**-8 round trips just to transmit requests
- Limited by TCP slow start congestion control
 - Can only have a few outstanding packets until ACKs begin returning
- Benefits increase the greater your network latency
 - ↗ Mobile LTE latency: 100+ ms (best case)

HTTP/2 Server Push

HTTP 2.0 connection



- Client requests Resource X
- Server anticipates that Resource Y will be requested next
 - Perhaps Resource Y is a dependency?
- Server pushes Resource Y to client despite never being requested
 - Client saves Resource Y in cache
- Client can elect not to accept pushed resources at all or limit total number/size

HTTP/2 Security

- HTTP/2 standard (RFC 7540) supports unencrypted connections
- Subject of much debate! (Should encryption be mandatory?)
 - **Pros:**
 - Security security security
 - Prevents tampering from *annoying* middleboxes assuming anything over port 80 is plain HTTP/1.1
 - **7** Cons:
 - Not all content has to be encrypted
 - → Performance / latency / etc...

HTTP/2 Security

- HTTP/2 standard (RFC 7540) supports unencrypted connections
 - Client starts with HTTP/1.1 and sends header: Upgrade: h2c
 - Server responds with HTTP 101 Switching Protocol status code
 - **7** Rare in wild! (*curl*?)
- Major browser implementations (Firefox, Chrome, Safari, etc...) <u>only</u> support HTTP/2 over TLS connections
 - Consistent with overall design philosophy of HTTPS everywhere: New features only enabled over HTTPS

- How to "enable" HTTP/2?
- Application Layer Protocol Negotiation (ALPN) – RFC 7301
 - Part of TLS handshake
 - Client provides server with list of protocols it supports
 - **7** Server picks one it prefers
 - Performance optimization that avoids additional roundtrip of starting with HTTP/1.1 and upgrading to HTTP/2

HTTP/2 Adoption

- **Widespread support!**
- Web browsers
 - Chrome, Safari, Firefox, Edge, ...
- Web servers
 - Apache, nginx, IIS, ...
- Content delivery networks
 - Akami, Azure, Cloudflare, AWS CloudFront, ...

HTTP/2 Requests

The percent of all requests in the crawl using HTTP/2.



Advanced Computer Networking

https://httparchive.org/reports/state-of-the-web

Fall 2018