HTTP/2, HTTP/3 the State of the Art in Our Servers

Jean-Frederic Clere @jfclere

What I will cover

- HTTP/2
  - HTTP/2 and ALPN
- HTTP/3
- Servers
  - Apache HTTPD
  - Tomcat
  - Traffic server
- Demos
- Questions?

Who I am

Jean-Frederic Clere

Red Hat

Years writing JAVA code and server software

Tomcat committer since 2001

Doing OpenSource since 1999

Cyclist/Runner etc

Lived 15 years in Spain (Barcelona)

Now in Neuchâtel (CH)

Why HTTP/2

- HTTP/1.1: June 1999 (RFC 2616)
  - 1999:
    - 1 page ~ 1kB HTML
  - 2019:
    - 1 page ~ 3MB HTML + IMAGES + JS + CSS etc
- Protocol:
  - Not adapted / inefficient / etc

HTTP/2 general

- HTTP/2:
  - Binary
  - Frame
  - Multiplex
  - Based on SPDY
  - TLS everywhere:
    - Browers use https and strong ciphers
  - No forward proxy
  - h2c: Clear text only with reverse proxy (proxy to back-end server)

HTTP/2 general

- Two specifications:
  - Hypertext Transfer Protocol version 2 - RFC7540
  - HPACK - Header Compression for HTTP/2 - RFC7541
- By the Internet Engineering Task Force
- ALPN Application-Layer Protocol Negotiation - RFC 7301

HTTP/2 Multiplexed

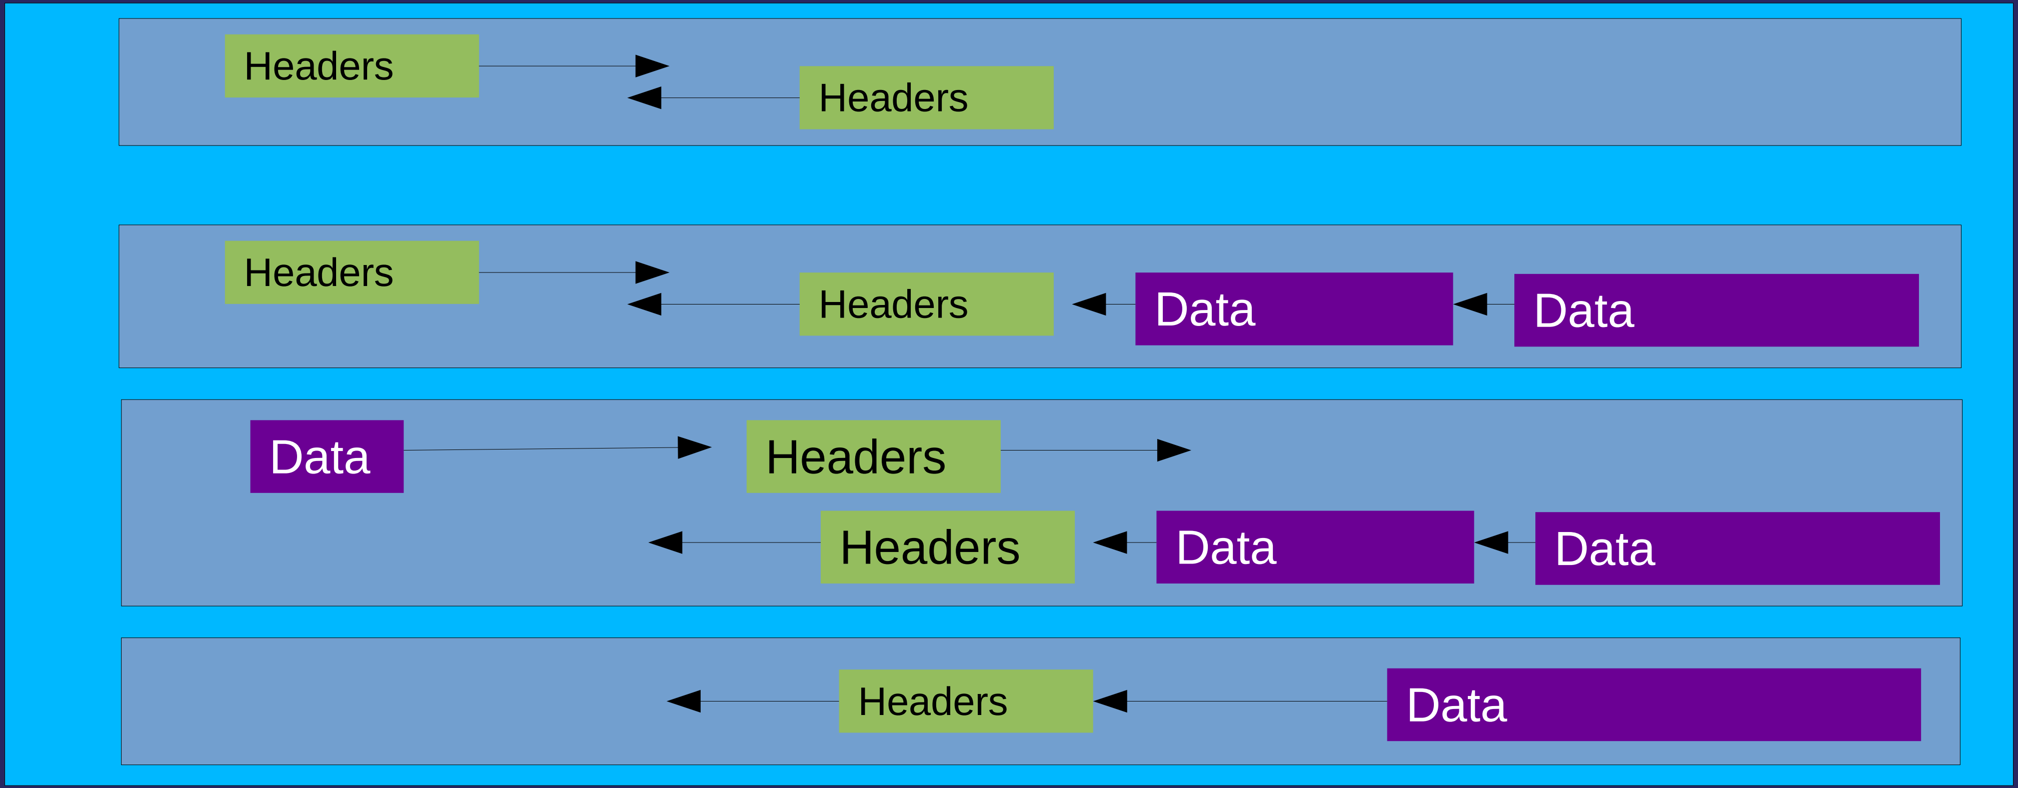

HTTP/2 : more

- HTTP headers compression
  - ~ 80 % save
- Request priority
  - Both sides
- Server Push
  - Prevent round trip to get element of a page
  - Faster / better rendering on browsers.

HTTP/2 With Browsers

- Browser with HTTP/2 and TLS
  - FireFox 34
  - Chrome 40 (with ALPN before was NPN)
  - IE 11
  - Opera and Safari 9
- Stats from docs.trafficserver and ci.trafficserver:
  - 80% is over HTTP/2 (data from last year)
- → go for it now!

ALPN Client Hello (Firefox)

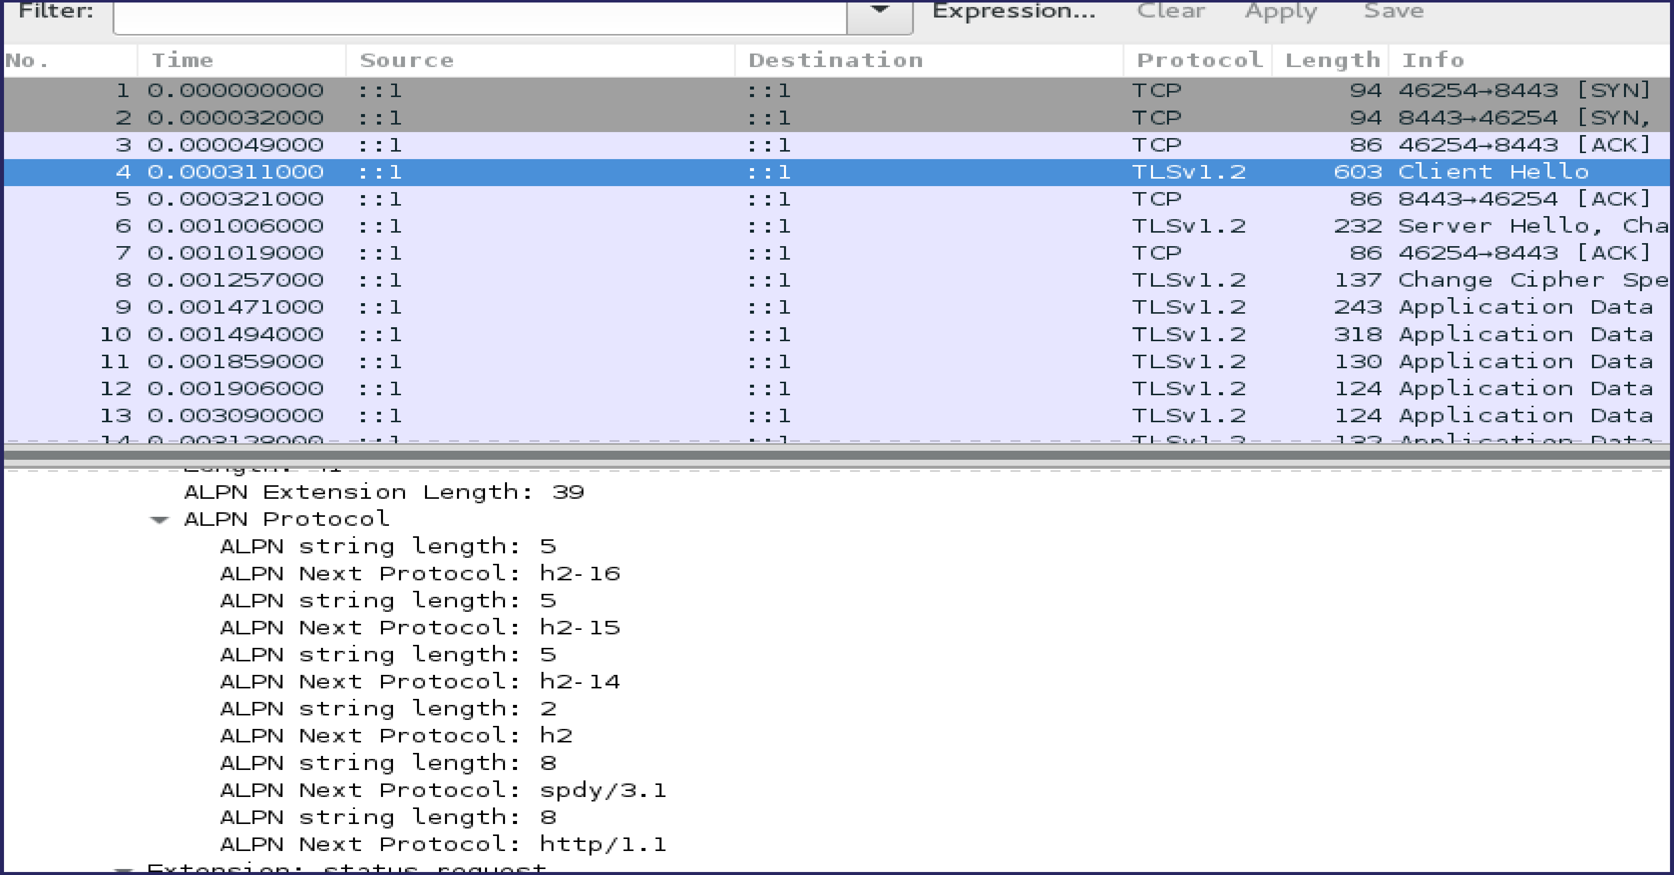

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000000	::1	::1	TCP	94	46254→8443 [SYN]
2	0.000032000	::1	::1	TCP	94	8443→46254 [SYN,
3	0.000049000	::1	::1	TCP	86	46254→8443 [ACK]
4	0.000311000	::1	::1	TLSv1.2	603	Client Hello
5	0.000321000	::1	::1	TCP	86	8443→46254 [ACK]
6	0.001006000	::1	::1	TLSv1.2	232	Server Hello, Cha
7	0.001019000	::1	::1	TCP	86	46254→8443 [ACK]
8	0.001257000	::1	::1	TLSv1.2	137	Change Cipher Spe
9	0.001471000	::1	::1	TLSv1.2	243	Application Data
10	0.001494000	::1	::1	TLSv1.2	318	Application Data
11	0.001859000	::1	::1	TLSv1.2	130	Application Data
12	0.001906000	::1	::1	TLSv1.2	124	Application Data
13	0.003090000	::1	::1	TLSv1.2	124	Application Data
14	0.003138000	::1	::1	TLSv1.2	133	Application Data

```
    Length: 41
    ALPN Extension Length: 39
  ▾ ALPN Protocol
      ALPN string length: 5
      ALPN Next Protocol: h2-16
      ALPN string length: 5
      ALPN Next Protocol: h2-15
      ALPN string length: 5
      ALPN Next Protocol: h2-14
      ALPN string length: 2
      ALPN Next Protocol: h2
      ALPN string length: 8
      ALPN Next Protocol: spdy/3.1
      ALPN string length: 8
      ALPN Next Protocol: http/1.1
▾ Extension: status_request
```

ALPN Server Hello (tomcat)

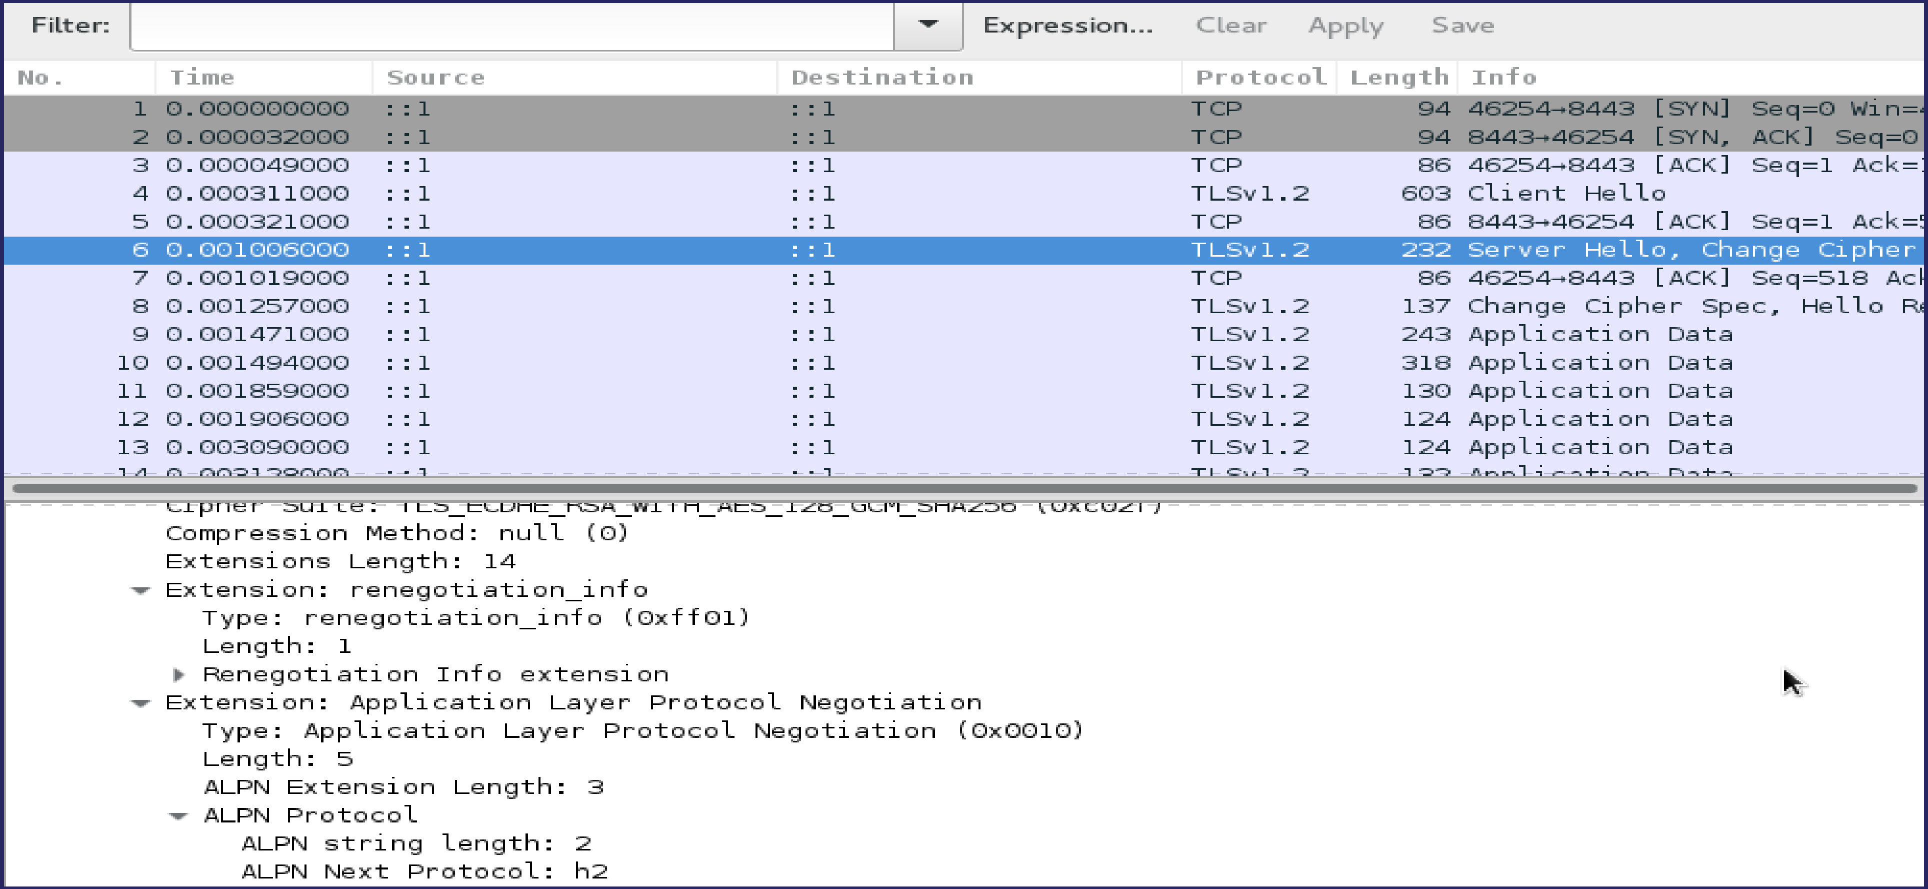

```
No.   Time          Source   Destination   Protocol  Length  Info
    1 0.000000000   ::1      ::1           TCP           94  46254→8443 [SYN] Seq=0 Win=
    2 0.000032000   ::1      ::1           TCP           94  8443→46254 [SYN, ACK] Seq=0
    3 0.000049000   ::1      ::1           TCP           86  46254→8443 [ACK] Seq=1 Ack=
    4 0.000311000   ::1      ::1           TLSv1.2      603  Client Hello
    5 0.000321000   ::1      ::1           TCP           86  8443→46254 [ACK] Seq=1 Ack=
    6 0.001006000   ::1      ::1           TLSv1.2      232  Server Hello, Change Cipher
    7 0.001019000   ::1      ::1           TCP           86  46254→8443 [ACK] Seq=518 Ac
    8 0.001257000   ::1      ::1           TLSv1.2      137  Change Cipher Spec, Hello R
    9 0.001471000   ::1      ::1           TLSv1.2      243  Application Data
   10 0.001494000   ::1      ::1           TLSv1.2      318  Application Data
   11 0.001859000   ::1      ::1           TLSv1.2      130  Application Data
   12 0.001906000   ::1      ::1           TLSv1.2      124  Application Data
   13 0.003090000   ::1      ::1           TLSv1.2      124  Application Data
   14 0.003138000   ::1      ::1           TLSv1.2      133  Application Data
```

```
    Cipher Suite: TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256 (0xc02f)
    Compression Method: null (0)
    Extensions Length: 14
  ▾ Extension: renegotiation_info
      Type: renegotiation_info (0xff01)
      Length: 1
    ▸ Renegotiation Info extension
  ▾ Extension: Application Layer Protocol Negotiation
      Type: Application Layer Protocol Negotiation (0x0010)
      Length: 5
      ALPN Extension Length: 3
    ▾ ALPN Protocol
        ALPN string length: 2
        ALPN Next Protocol: h2
```

Requirements

- OpenSSL for our 3 servers
  - At least 1.0.2c
- Tomcat (8.5 / trunk)
  - Tomcat-native (1.2.6 / trunk) or java9
- Httpd (2.4.17 / trunk)
  - HTTP/2 C Library (libnghttp2)
- TrafficServer (since ATS v5.3.2).
  - Nothing except openssl.

Status

- Tomcat (trunk/8.5)
  - Full support / released as stable.
  - Needs servlet 4.0 (JSR 369) for server PUSH API
  - Can't be full JAVA until JDK9 (ALPN support)
- Httpd (available since 2.4.17)
  - Full support (since 2.4.20)
- TrafficServer (since 5.3.0) (flow control 6.1)
  - Priorities (6.2.0) and Server PUSH (7.0.0)

TC connector server.xml

```
<Connector
     port="8002"
    protocol="org.apache.coyote.http11.Http11NioProtocol"
     MaxThreads="150"
     SSLEnabled="true">
     <SSLHostConfig>
         <Certificate
              certificateFile="/home/jfclere/H3/certs/pubcert.pem"
              certificateKeyFile="/home/jfclere/H3/certs/privkey.pem"/>
     </SSLHostConfig>
    <UpgradeProtocol className="org.apache.coyote.http2.Http2Protocol" />
 <Connector/>
```

Tomcat / configuration

In bin/setenv.sh:

```
LD_LIBRARY_PATH=/home/jfclere/tomcat-native/native/.libs

export LD_LIBRARY_PATH
```

And the libtcnative-1.so linked with openssl-1.0.2c, checking with ldd:

```
libssl.so.1.0.0 => /home/jfclere/OPENSSL-1.0.2c/lib/libssl.so.1.0.0 (0x00007f6ab147b000)

libcrypto.so.1.0.0 => /home/jfclere/OPENSSL-1.0.2c/lib/libcrypto.so.1.0.0 (0x00007f6ab1028000)

libapr-1.so.0 => /home/jfclere/APR-1.4.x/lib/libapr-1.so.0 (0x00007f6ab0dfa000)
```

Usually the openssl of recent distribution (fedora 23) will work.

Tomcat / Performances

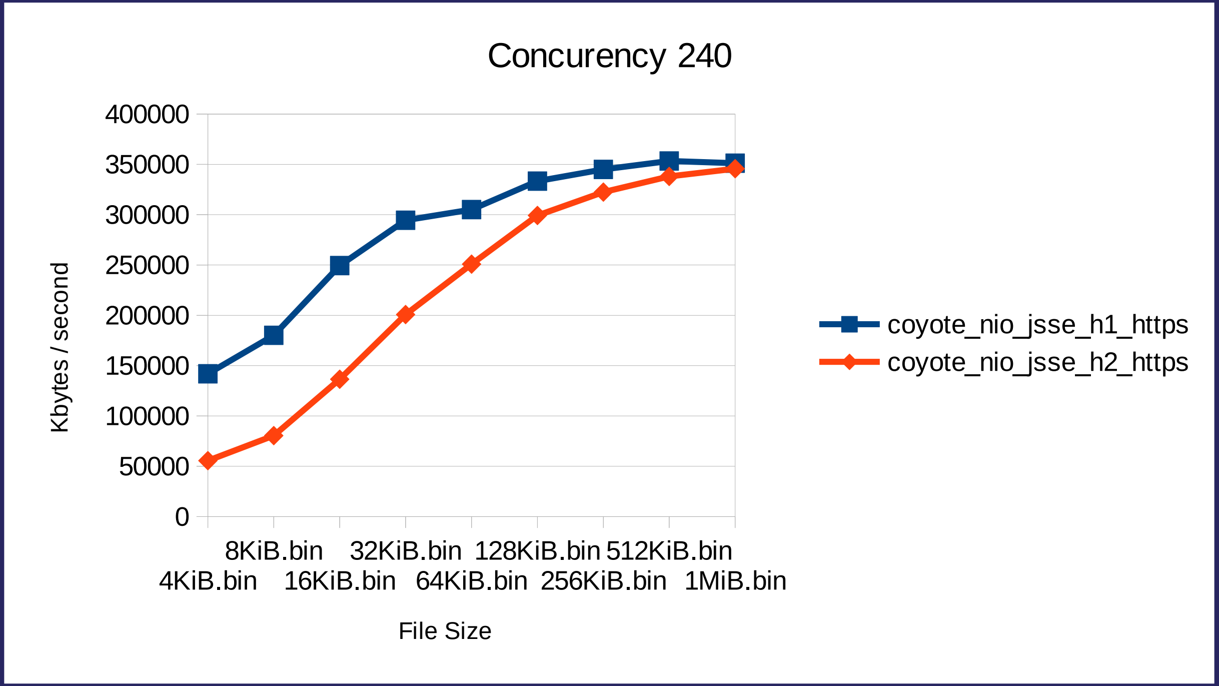

Tomcat / Performances

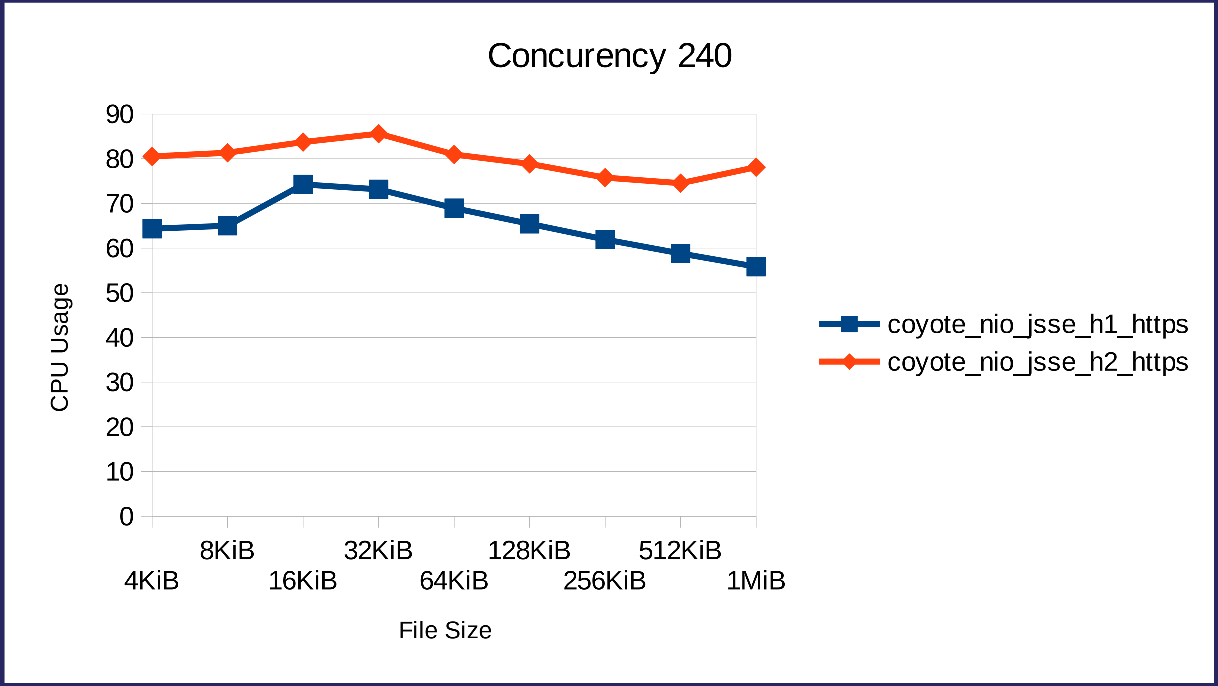

Tomcat / Demo

- No server push (anyway the browsers stop supporting it :-()
- Multiplexing
- headers compression
- HTML page:
  - That requires a lot (~100) of (~4Kbytes) images to render.

TrafficServer / Configuration

- records.config
  - CONFIG proxy.config.ssl.number.threads INT 0
  - **CONFIG proxy.config.http.server_ports STRING 8888:ssl**
  - CONFIG proxy.config.url_remap.pristine_host_hdr INT 1
  - CONFIG proxy.config.http2.enabled INT 1
  - CONFIG proxy.config.ssl.TLSv1_2 INT 1
- ssl_multicert.config:
  - **dest_ip=* ssl_cert_name=newcert.pem ssl_key_name=newkey.txt.pem**
- remap.config:
  - **map / http://127.0.0.1:8080**
- ip_allow.config:
  - src_ip=**192.168.1.38**   action=ip_allow method=ALL
  - src_ip=::-ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff action=ip_allow method=ALL

TrafficServer / Demo

- Like tomcat one
- Uses http/1.1 tomcat nio connector on 8080 as back-end.

HTTPd / Configuration

- httpd.conf:

```
LoadModule h2_module modules/mod_h2.so
Listen 8006
<VirtualHost *:8006>
  Protocols h2 http/1.1
  ProtocolsHonorOrder on
  SSLEngine on
  SSLCertificateFile "/home/jfclere/CERTS/newcert.pem"
  SSLCertificateKeyFile "/home/jfclere/CERTS/newkey.pem"
  SSLCACertificateFile "/etc/pki/CA/cacert.pem"
</VirtualHost>
```

HTTPd / Configuration proxy

- httpd.conf:

```
LoadModule http2_module modules/mod_http2.so
LoadModule proxy_http2_module modules/mod_proxy_http2.so
Listen 8006
<VirtualHost *:8006>
  Protocols h2 http/1.1
  ProtocolsHonorOrder on
  SSLEngine on
  …
  ProxyPass "/" "h2c://localhost:8003/"
</VirtualHost>
```

HTTPd / Demo

- Like the tomcat one:
  - htdocs/http2.html
  - htdocs/images/ the images.

HTTP/2 move to it?

- Conclusion:
  - Using HTTP/2 without PUSH is already good.
  - “safer” crypto is good but expensive.
  - No need to rewrite application to get the gains.

HTTP/2 : GO FOR IT

Then Why HTTP/3?

- TCP/IP:
  - Windows acks: 1 packet lost → all the channels blocked.
- UPD:
  - Channels are independent.
  - Need higher protocol level to insure integrity.
  - Packets might not be received in other.
- Security:
  - Need a patched version of OpenSSL (and use TLS-1.3)
  - UDP: cloud → no… but DNS → used every where

HTTP/3 (RFC 9114 published June 2022)

- Use QUIC / TLS-1.3 / UDP
- To transport HTTP like HTTP/2
- Initial connection TCP + Alt-Svc or HTTP/2
  - Response Alt-Svc: h3=":56666":
  - HTTP/2   ALTSVC frame
- problems:
  - UDP ports closed
  - UDP slower than TCP in Kernels
  - Needs extra CPU (?)
- Specifications:
  - RC 9114

Features: HTTP/2 vs HTTP/3

	HTTP/2	HTTP/3
Transport	TCP	UPD/QUIC
Streams	HTTP/2	QUIC
Clear text	yes (h2c: reverse proxy)	no
Independent streams	no	yes
Header compression	HPACK	QPACK
Server push	yes	yes
Early data	no	yes
0-RTT handshake	no (TLS-1.2)	Yes (TLS-1.3+)

HTTP/3 implementations

- quiche:
  - https://docs.quic.tech/quiche/
- Curl: https://curl.se/docs/http3.html
  - ngtcp2 (nghttp3/ngtcp2/patched openssl, GnuTLS etc)
  - quiche
  - msh3
  - In experimental at build time.
- Browser: chrome / firefox (active by default: Apr 2021).

HTTP/3 in our servers:

- Apache Tomcat: Problem UDP socket API incomplete (java 15)
- Apache HTTPD: need time probably like http/2
- Traffic Server: in the 9.1.x experimental (need patched openssl)
  - See ATS docs / curl docs
  - 10-dev: boringSSL or quiche

TrafficServer / Configuration

- records.config
  - CONFIG proxy.config.udp.threads INT 1
  - CONFIG proxy.config.http.server_ports STRING 4433:quic
  - CONFIG proxy.config.diags.debug.enabled INT 1
  - CONFIG proxy.config.diags.debug.tags STRING quic
- ssl_multicert.config:
  - **dest_ip=* ssl_cert_name=newcert.pem ssl_key_name=newkey.txt.pem**
- remap.config:
  - **map / http://127.0.0.1:8080**

TrafficServer / H3 Demo

- Uses tomcat as backend
- Uses http/1.1 tomcat nio connector on 8080 as back-end.
- Uses Apache HTTPD https + mod_header to create the alt-svc

TrafficServer / Demo

- https://jfclere.myddns.me:4433/
- Response HTTP/1.1 (HTTP/2) header alt-svc
- alt-svc: h3=":4433"; ma=60, h3-29=":4433"; ma=60
- H3-29 (HTTP/3 draft 29)
- ma=60 seconds = 1 minute.
- Next requests → HTTP/3

TrafficServer / Demo

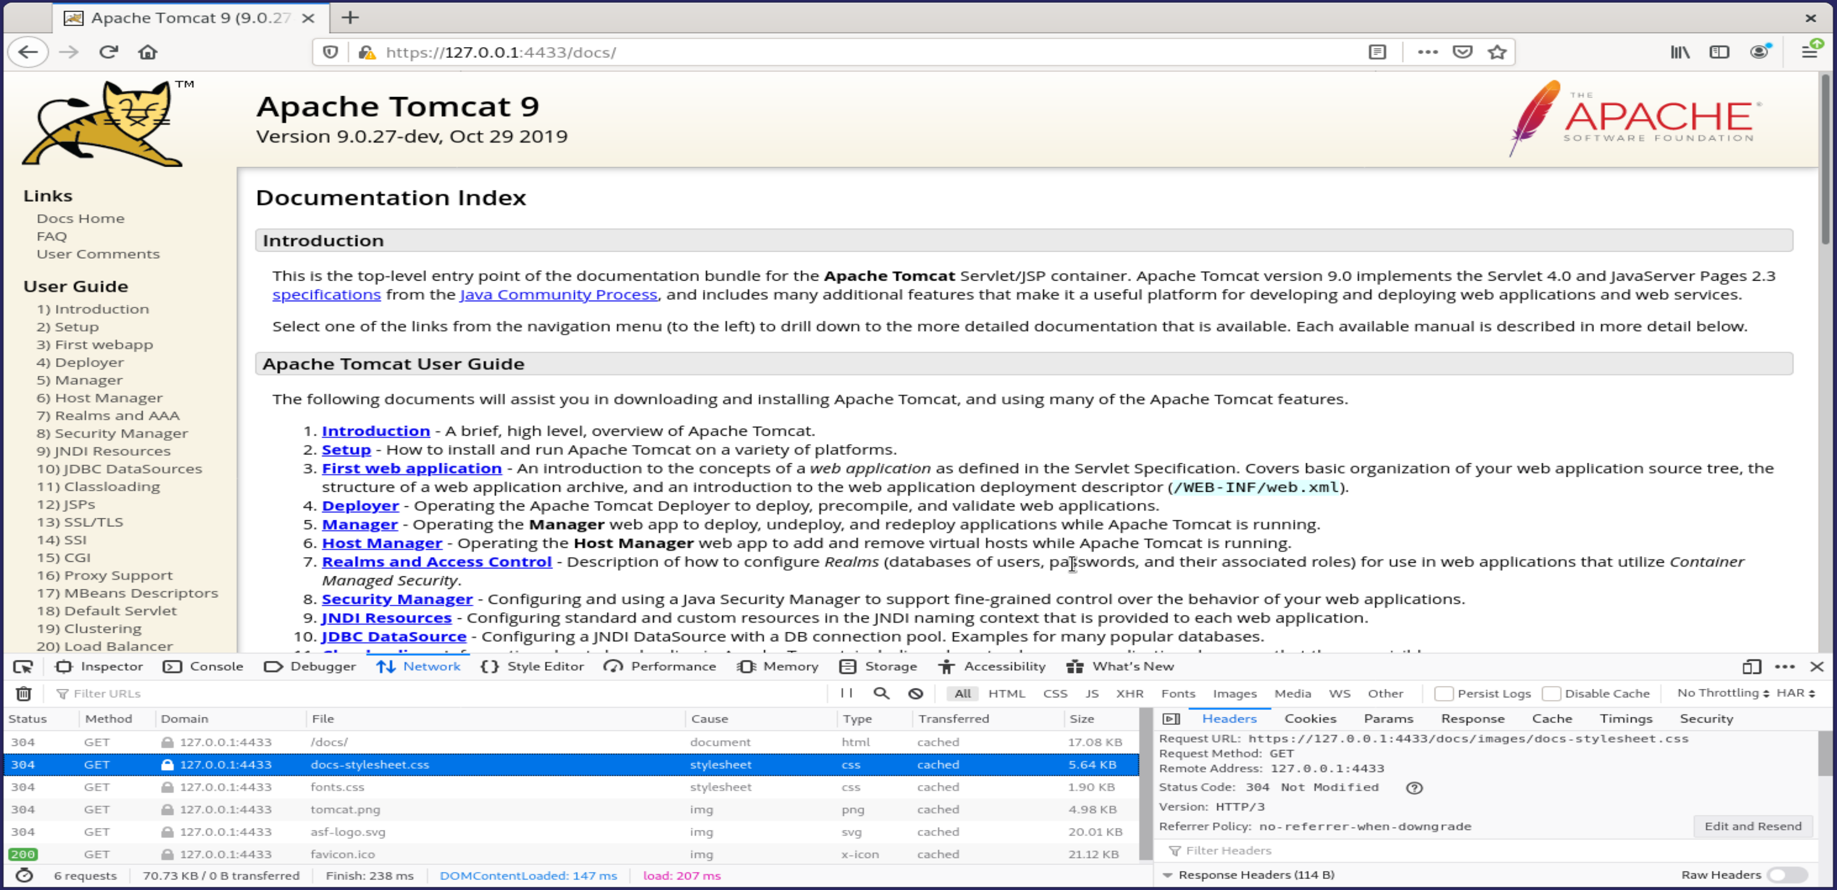

Apache Tomcat 9

Version 9.0.27-dev, Oct 29 2019

Links

- Docs Home
- FAQ
- User Comments

User Guide

- 1) Introduction
- 2) Setup
- 3) First webapp
- 4) Deployer
- 5) Manager
- 6) Host Manager
- 7) Realms and AAA
- 8) Security Manager
- 9) JNDI Resources
- 10) JDBC DataSources
- 11) Classloading
- 12) JSPs
- 13) SSL/TLS
- 14) SSI
- 15) CGI
- 16) Proxy Support
- 17) MBeans Descriptors
- 18) Default Servlet
- 19) Clustering
- 20) Load Balancer

Documentation Index

Introduction

This is the top-level entry point of the documentation bundle for the **Apache Tomcat** Servlet/JSP container. Apache Tomcat version 9.0 implements the Servlet 4.0 and JavaServer Pages 2.3 specifications from the Java Community Process, and includes many additional features that make it a useful platform for developing and deploying web applications and web services.

Select one of the links from the navigation menu (to the left) to drill down to the more detailed documentation that is available. Each available manual is described in more detail below.

Apache Tomcat User Guide

The following documents will assist you in downloading and installing Apache Tomcat, and using many of the Apache Tomcat features.

1. **Introduction** - A brief, high level, overview of Apache Tomcat.
2. **Setup** - How to install and run Apache Tomcat on a variety of platforms.
3. **First web application** - An introduction to the concepts of a *web application* as defined in the Servlet Specification. Covers basic organization of your web application source tree, the structure of a web application archive, and an introduction to the web application deployment descriptor (`/WEB-INF/web.xml`).
4. **Deployer** - Operating the Apache Tomcat Deployer to deploy, precompile, and validate web applications.
5. **Manager** - Operating the **Manager** web app to deploy, undeploy, and redeploy applications while Apache Tomcat is running.
6. **Host Manager** - Operating the **Host Manager** web app to add and remove virtual hosts while Apache Tomcat is running.
7. **Realms and Access Control** - Description of how to configure *Realms* (databases of users, passwords, and their associated roles) for use in web applications that utilize *Container Managed Security*.
8. **Security Manager** - Configuring and using a Java Security Manager to support fine-grained control over the behavior of your web applications.
9. **JNDI Resources** - Configuring standard and custom resources in the JNDI naming context that is provided to each web application.
10. **JDBC DataSource** - Configuring a JNDI DataSource with a DB connection pool. Examples for many popular databases.

Status	Method	Domain	File	Cause	Type	Transferred	Size
304	GET	127.0.0.1:4433	/docs/	document	html	cached	17.08 KB
304	GET	127.0.0.1:4433	docs-stylesheet.css	stylesheet	css	cached	5.64 KB
304	GET	127.0.0.1:4433	fonts.css	stylesheet	css	cached	1.90 KB
304	GET	127.0.0.1:4433	tomcat.png	img	png	cached	4.98 KB
304	GET	127.0.0.1:4433	asf-logo.svg	img	svg	cached	20.01 KB
200	GET	127.0.0.1:4433	favicon.ico	img	x-icon	cached	21.12 KB

6 requests | 70.73 KB / 0 B transferred | Finish: 238 ms | DOMContentLoaded: 147 ms | load: 207 ms

```
Request URL: https://127.0.0.1:4433/docs/images/docs-stylesheet.css
Request Method: GET
Remote Address: 127.0.0.1:4433
Status Code: 304 Not Modified
Version: HTTP/3
Referrer Policy: no-referrer-when-downgrade
```

Response Headers (114 B)

TrafficServer / Demo

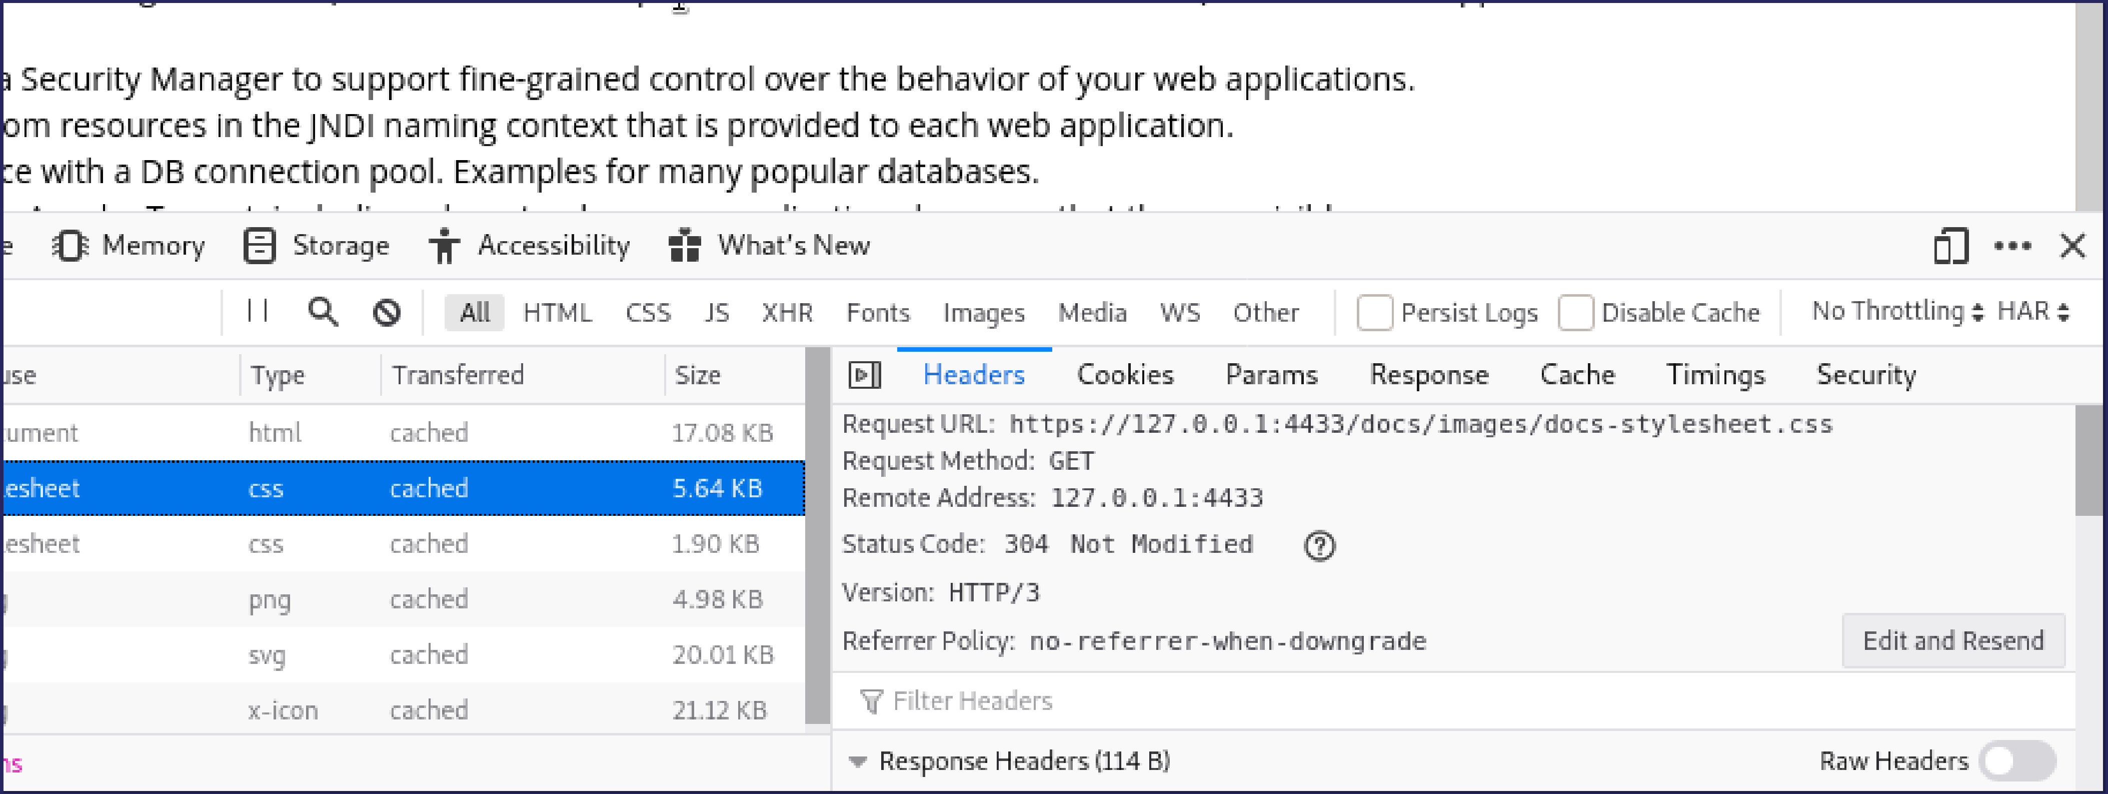

a Security Manager to support fine-grained control over the behavior of your web applications.
om resources in the JNDI naming context that is provided to each web application.
ce with a DB connection pool. Examples for many popular databases.

use	Type	Transferred	Size
ument	html	cached	17.08 KB
esheet	css	cached	5.64 KB
esheet	css	cached	1.90 KB
	png	cached	4.98 KB
	svg	cached	20.01 KB
	x-icon	cached	21.12 KB

ns

Request URL: https://127.0.0.1:4433/docs/images/docs-stylesheet.css
Request Method: GET
Remote Address: 127.0.0.1:4433
Status Code: 304 Not Modified
Version: HTTP/3
Referrer Policy: no-referrer-when-downgrade

Response Headers (114 B)

HTTP/3 more info:

- Playing with browsers:
  - Interop matrix
  - H3 activated by default in recent (2021) Firefox/Chrome
- OpenSSL 3.0.x (with patches)!!!

HTTP/3 ready?

- Conclusion:
  - Not more a draft, last draft was H3-34.
  - UDP versus TCP.
  - Needs forked version of openssl… (0-RTT).
  - Or BoringSSL.
  - No need to rewrite application to get the gains.

HTTP/3 : wait

Questions?

Thank you!

- jfclere@gmail.com
- users@tomcat.apache.org
- users@httpd.apache.org
- users@trafficserver.apache.org
- https://http2.github.io/
- Demo generator: https://github.com/jfclere/h2_demos
- HTTP/3 see curl docs: http3-explained by Daniel
- More on HTP/3: https://github.com/jfclere/AC2022/tree/main/h3