


Wireshark Lab: HTTP

SOLUTION

Supplement to *Computer Networking: A Top-Down Approach, 7th ed.*, J.F. Kurose and K.W. Ross

© 2005-2012, J.F. Kurose and K.W. Ross, All Rights Reserved



The following screen shots showing the HTTP GET and HTTP reply answer these questions:

1. Is your browser running HTTP version 1.0 or 1.1? What version of HTTP is the server running?
2. What languages (if any) does your browser indicate that it can accept to the server?
3. What is the IP address of your computer? Of the gaia.cs.umass.edu server?
4. What is the status code returned from the server to your browser?
5. When was the HTML file that you are retrieving last modified at the server?
6. How many bytes of content are being returned to your browser?
7. By inspecting the raw data in the packet content window, do you see any headers within the data that are not displayed in the packet-listing window? If so, name one. *Answer: no, I don't see any in the HTTP Message below*

No.	Time	Source	Destination	Protocol	Info
133	4.098946	192.168.1.101	128.119.245.12	HTTP	GET /wireshark-labs/HTTP-wire

Frame 133 (488 bytes on wire, 488 bytes captured)
 Ethernet II, Src: IntelCor dc:36:d0 (00:22:fa:dc:36:d0), Dst: Cisco-Li 45:1f:1b (00:22:6b:45:1f:1b)
 Internet Protocol, Src: 192.168.1.101 (192.168.1.101), Dst: 128.119.245.12 (128.119.245.12)
 Version: 4
 Header length: 20 bytes
 Differentiated Services Field: 0x00 (DSCP 0x00: Default; ECN: 0x00)
 Total Length: 474
 Identification: 0x036e (878)
 Flags: 0x02 (Don't Fragment)
 Fragment offset: 0
 Time to live: 128
 Protocol: TCP (0x06)
 Header checksum: 0xbebe [correct]
 Source: 192.168.1.101 (192.168.1.101)
 Destination: 128.119.245.12 (128.119.245.12)

Transmission Control Protocol, Src Port: 55428 (55428), Dst Port: http (80), Seq: 1, Ack: 1, Len: 434
 Source port: 55428 (55428)
 Destination port: http (80)
 [Stream index: 27]
 Sequence number: 1 (relative sequence number)
 [Next sequence number: 435 (relative sequence number)]
 Acknowledgement number: 1 (relative ack number)
 Header length: 20 bytes
 Flags: 0x18 (PSH, ACK)
 Window size: 64240
 Checksum: 0xe737 [validation disabled]
 [SEQ/ACK analysis]

Hypertext Transfer Protocol
 GET /wireshark-labs/HTTP-wireshark-file1.html HTTP/1.1\r\n
 Host: gaia.cs.umass.edu\r\n
 User-Agent: Mozilla/5.0 (Windows; U; Windows NT 6.1; en-US; rv:1.9.2.11) Gecko/20101012 Firefox/3.
 Accept: text/html,application/xhtml+xml,application/xml;q=0.9,*/*;q=0.8\r\n
 Accept-Language: en-us,en;q=0.5\r\n
 Accept-Encoding: gzip,deflate\r\n
 Accept-Charset: ISO-8859-1,utf-8;q=0.7,*;q=0.7\r\n
 Keep-Alive: 115\r\n
 Connection: keep-alive\r\n
 \r\n

Client IP address

Gaia server IP address

Client running http 1.1

languages accepted

No.	Time	Source	Destination	Protocol	Info
135	4.126437	128.119.245.12	192.168.1.101	HTTP	HTTP/1.1 200 OK (text/html)

Frame 135 (488 bytes on wire, 488 bytes captured)
 Ethernet II, Src: Cisco-Li 45:1f:1b (00:22:6b:45:1f:1b), Dst: IntelCor dc:36:d0 (00:22:fa:dc:36:d0)
 Internet Protocol, Src: 128.119.245.12 (128.119.245.12), Dst: 192.168.1.101 (192.168.1.101)
 Transmission Control Protocol, Src Port: http (80), Dst Port: 55428 (55428), Seq: 1, Ack: 435, Len: 434

Hypertext Transfer Protocol
 HTTP/1.1 200 OK\r\n
 [Expert Info (Chat/Sequence): HTTP/1.1 200 OK\r\n]
 Request Version: HTTP/1.1
 Response Code: 200
 Date: Wed, 27 Oct 2010 11:26:58 GMT\r\n
 Server: Apache/2.0.52 (CentOS)\r\n
 Last-Modified: Wed, 27 Oct 2010 11:26:01 GMT\r\n
 ETag: "8734d-80-7d/4e440"\r\n
 Accept-Ranges: bytes\r\n
 Content-Length: 128\r\n
 [Content length: 128]
 Keep-Alive: timeout=10, max=100\r\n
 Connection: Keep-Alive\r\n
 Content-Type: text/html; charset=ISO-8859-1\r\n
 \r\n

Line-based text data: text/html
 <html>\n
 Congratulations. You've downloaded the file \n
 http://gaia.cs.umass.edu/wireshark-labs/HTTP-wireshark-file1.html!\n
 </html>\n

Return status:

server running http

document last modified

content: 128

2. The HTTP CONDITIONAL GET/response interaction

Here's a screenshot after doing the two identical HTTP GETs:

Microsoft - Wireshark

Filter: http

No.	Time	Source	Destination	Protocol	Info
110	7.087745	192.168.1.1	239.255.255.250	SSDP	NOTIFY * HTTP/1.1
111	7.191861	192.168.1.1	239.255.255.250	SSDP	NOTIFY * HTTP/1.1
112	7.292517	192.168.1.1	239.255.255.250	SSDP	NOTIFY * HTTP/1.1
142	7.396096	192.168.1.1	239.255.255.250	SSDP	NOTIFY * HTTP/1.1
167	7.740235	192.168.1.101	128.119.245.12	HTTP	GET /wireshark-labs/HTTP-wireshark-file2.html HTTP/1.1
170	7.773450	128.119.245.12	192.168.1.101	HTTP	HTTP/1.1 200 OK (text/html)
182	11.154651	192.168.1.101	128.119.245.12	HTTP	GET /wireshark-labs/HTTP-wireshark-file2.html HTTP/1.1
183	11.185432	128.119.245.12	192.168.1.101	HTTP	HTTP/1.1 304 Not Modified

Frame 167 (488 bytes on wire, 488 bytes captured)

Ethernet II, Src: IntelCor_dc:36:d0 (00:22:fa:dc:36:d0), Dst: Cisco-Li_45:1f:1b (00:22:60:45:1f:1b)

Internet Protocol, Src: 192.168.1.101 (192.168.1.101), Dst: 128.119.245.12 (128.119.245.12)

Transmission Control Protocol, Src Port: 55449 (55449), Dst Port: http (80), Seq: 1, Ack: 1, Len: 434

Hypertext Transfer Protocol

- GET /wireshark-labs/HTTP-wireshark-file2.html HTTP/1.1\r\n
- [Expert Info (Chat,Sequence): GET /wireshark-labs/HTTP-wireshark-file2.html HTTP/1.1
- Request Method: GET
- Request URI: /wireshark-labs/HTTP-wireshark-file2.html
- Request Version: HTTP/1.1
- Host: gaia.cs.umass.edu\r\n
- User-Agent: Mozilla/5.0 (Windows; U; Windows NT 6.1; en-US; rv:1.9.2.11) Gecko/2010101
- Accept: text/html,application/xhtml+xml,application/xml;q=0.9,*/*;q=0.8\r\n
- Accept-Language: en-us,en;q=0.5\r\n
- Accept-Encoding: gzip,deflate\r\n
- Accept-Charset: ISO-8859-1,utf-8;q=0.7,*;q=0.7\r\n
- Keep-Alive: 115\r\n
- Connection: keep-alive\r\n
- \r\n

First GET, then a reply, then another identical GET, then a reply (304 not modified)

Answer the following questions:

- Inspect the contents of the first HTTP GET request from your browser to the server. Do you see an "IF-MODIFIED-SINCE" line in the HTTP GET?
- Inspect the contents of the server response. Did the server explicitly return the contents of the file? How can you tell?
- Now inspect the contents of the second HTTP GET request from your browser to the server. Do you see an "IF-MODIFIED-SINCE:" line in the HTTP GET? If so, what information follows the "IF-MODIFIED-SINCE:" header?
- What is the HTTP status code and phrase returned from the server in response to this second HTTP GET? Did the server explicitly return the contents of the file? Explain.

Here are the four captures packets (two GETs and two REPLIES, in chronological order):

No.	Time	Source	Destination	Protocol	Info
167	7.740235	192.168.1.101	128.119.245.12	HTTP	GET /wireshark-labs/HTTP-wire

Frame 167 (488 bytes on wire, 488 bytes captured)
Ethernet II, Src: IntelCor dc:36:d0 (00:22:fa:dc:36:d0), Dst: Cisco-Li 45:1f:1b (00:22:6b:45:1f:1b)
Internet Protocol, Src: 192.168.1.101 (192.168.1.101), Dst: 128.119.245.12 (128.119.245.12)
Transmission Control Protocol, Src Port: 55449 (55449), Dst Port: http (80), Seq: 1, Ack: 1, Len: 434
Hypertext Transfer Protocol

```

GET /wireshark-labs/HTTP-wireshark-file2.html HTTP/1.1\r\n
[Expert Info (Chat/Sequence): GET /wireshark-labs/HTTP-wireshark-file2.html HTTP/1.1\r\n]
[Message: GET /wireshark-labs/HTTP-wireshark-file2.html HTTP/1.1\r\n]
[Severity level: Chat]
[Group: Sequence]
Request Method: GET
Request URI: /wireshark-labs/HTTP-wireshark-file2.html
Request Version: HTTP/1.1
Host: gaia.cs.umass.edu\r\n
User-Agent: Mozilla/5.0 (Windows; U; Windows NT 6.1; en-US; rv:1.9.2.11) Gecko/20101012 Firefox/3.
Accept: text/html,application/xhtml+xml,application/xml;q=0.9,*/*;q=0.8\r\n
Accept-Language: en-us,en;q=0.5\r\n
Accept-Encoding: gzip,deflate\r\n
Accept-Charset: ISO-8859-1,utf-8;q=0.7,*;q=0.7\r\n
Keep-Alive: 115\r\n
Connection: keep-alive\r\n
\r\n

```

There is no IF-MODIFIED-SINCE in the first GET

No.	Time	Source	Destination	Protocol	Info
170	7.773450	128.119.245.12	192.168.1.101	HTTP	HTTP/1.1 200 OK (text/html)

Frame 170 (425 bytes on wire, 425 bytes captured)
Ethernet II, Src: Cisco-Li 45:1f:1b (00:22:6b:45:1f:1b), Dst: IntelCor dc:36:d0 (00:22:fa:dc:36:d0)
Internet Protocol, Src: 128.119.245.12 (128.119.245.12), Dst: 192.168.1.101 (192.168.1.101)
Transmission Control Protocol, Src Port: http (80), Dst Port: 55449 (55449), Seq: 308, Ack: 435, Len: [Reassembled TCP Segments (678 bytes): #169(307), #170(371)]
Hypertext Transfer Protocol

```

HTTP/1.1 200 OK\r\n
[Expert Info (Chat/Sequence): HTTP/1.1 200 OK\r\n]
Request Version: HTTP/1.1
Response Code: 200
Date: Wed, 27 Oct 2010 11:54:25 GMT\r\n
Server: Apache/2.0.52 (CentOS)\r\n
Last-Modified: Wed, 27 Oct 2010 11:54:02 GMT\r\n
ETag: "d6c96-173-ela6ea80"\r\n
Accept-Ranges: bytes\r\n
Content-Length: 371\r\n
[Content length: 371]
Keep-Alive: timeout=10, max=100\r\n
Connection: Keep-Alive\r\n
Content-Type: text/html; charset=ISO-8859-1\r\n
\r\n

```

Line-based text data: text/html

```

\r\n
<html>\r\n
\r\n
\r\n
Congratulations again! Now you've downloaded the file lab2-2.html. <br>\r\n
This file's last modification date will not change. <p>\r\n
Thus if you download this multiple times on your browser, a complete copy <br>\r\n
will only be sent once by the server due to the inclusion of the IN-MODIFIED-SINCE<br>\r\n
field in your browser's HTTP GET request to the server.\r\n
\r\n
</html>\r\n

```

text returned in response to first GET

No.	Time	Source	Destination	Protocol	Info
182	11.154651	192.168.1.101	128.119.245.12	HTTP	GET /wireshark-labs/HTTP-wire

Frame 182 (575 bytes on wire, 575 bytes captured)
Ethernet II, Src: IntelCor dc:36:d0 (00:22:fa:dc:36:d0), Dst: Cisco-Li 45:1f:1b (00:22:6b:45:1f:1b)
Internet Protocol, Src: 192.168.1.101 (192.168.1.101), Dst: 128.119.245.12 (128.119.245.12)
Transmission Control Protocol, Src Port: 55449 (55449), Dst Port: http (80), Seq: 435, Ack: 679, Len:
Hypertext Transfer Protocol
GET /wireshark-labs/HTTP-wireshark-file2.html HTTP/1.1\r\n
[Expert Info (Chat/Sequence): GET /wireshark-labs/HTTP-wireshark-file2.html HTTP/1.1\r\n
Request Method: GET
Request URI: /wireshark-labs/HTTP-wireshark-file2.html
Request Version: HTTP/1.1
Host: gaia.cs.umass.edu\r\n
User-Agent: Mozilla/5.0 (Windows; U; Windows NT 6.1; en-US; rv:1.9.2.11) Gecko/20101012 Firefox/3.
Accept: text/html,application/xhtml+xml,application/xml;q=0.9,*/*;q=0.8\r\n
Accept-Language: en-us,en;q=0.5\r\n
Accept-Encoding: gzip,deflate\r\n
Accept-Charset: ISO-8859-1,utf-8;q=0.7,*;q=0.7\r\n
Keep-Alive: 115\r\n
Connection: keep-alive\r\n
If-Modified-Since: Wed, 27 Oct 2010 11:54:02 GMT\r\n
If-None-Match: "d6c96-173-e1a6ea80"\r\n
\r\n

2nd GET has IF-MODIFIED-SINCE

No.	Time	Source	Destination	Protocol	Info
183	11.185432	128.119.245.12	192.168.1.101	HTTP	HTTP/1.1 304 Not Modified

Frame 183 (236 bytes on wire, 236 bytes captured)
Ethernet II, Src: Cisco-Li 45:1f:1b (00:22:6b:45:1f:1b), Dst: IntelCor dc:36:d0 (00:22:fa:dc:36:d0)
Internet Protocol, Src: 128.119.245.12 (128.119.245.12), Dst: 192.168.1.101 (192.168.1.101)
Transmission Control Protocol, Src Port: http (80), Dst Port: 55449 (55449), Seq: 679, Ack: 956, Len:
Hypertext Transfer Protocol
HTTP/1.1 304 Not Modified\r\n
[Expert Info (Chat/Sequence): HTTP/1.1 304 Not Modified\r\n
Request Version: HTTP/1.1
Response Code: 304
Date: Wed, 27 Oct 2010 11:54:28 GMT\r\n
Server: Apache/2.0.52 (CentOS)\r\n
Connection: Keep-Alive\r\n
Keep-Alive: timeout=10, max=99\r\n
ETag: "d6c96-173-e1a6ea80"\r\n
\r\n

**The file has not been modified!
So the text of the file is NOT
returned in the HTTP message**

3. Retrieving Long Documents

In our answer below, we use the http-ethereal-trace-3 packet trace file. The HTTP GET for the long document is packet 8 in the trace (at t=4.623732); the HTTP OK reply is packet 14 (at t=6.680432).

No.	Time	Source	Destination	Protocol	Length	Info
5	4.602642	192.168.1.102	128.119.245.12	TCP	62	4272 > http [SYN] Seq=0 Win=64240 Len=0 MSS=1460 SACK_PERM=1
6	4.623285	128.119.245.12	192.168.1.102	TCP	62	http > 4272 [SYN, ACK] Seq=0 Ack=1 Win=5840 Len=0 MSS=1460 SACK_P
7	4.623313	192.168.1.102	128.119.245.12	TCP	54	4272 > http [ACK] Seq=1 Ack=1 Win=64240 Len=0
8	4.623732	192.168.1.102	128.119.245.12	HTTP	555	GET /ethereal-labs/lab2-5.html HTTP/1.1
9	4.652711	128.119.245.12	192.168.1.102	TCP	60	http > 4272 [ACK] Seq=1 Ack=502 Win=6432 Len=0
10	4.657459	128.119.245.12	192.168.1.102	TCP	1514	[TCP segment of a reassembled PDU]
11	4.658792	128.119.245.12	192.168.1.102	TCP	1514	[TCP segment of a reassembled PDU]
12	4.658828	192.168.1.102	128.119.245.12	TCP	54	4272 > http [ACK] Seq=502 Ack=2921 Win=64240 Len=0
13	4.680438	128.119.245.12	192.168.1.102	TCP	1514	[TCP segment of a reassembled PDU]
14	4.680920	128.119.245.12	192.168.1.102	HTTP	490	HTTP/1.1 200 OK (text/html)
15	4.680948	192.168.1.102	128.119.245.12	TCP	54	4272 > http [ACK] Seq=502 Ack=4817 Win=64240 Len=0

Frame 10: 1514 bytes on wire (12112 bits), 1514 bytes captured (12112 bits) on interface 0
 Ethernet II, Src: LinksysG_da:af:73 (00:06:25:da:af:73), Dst: DellComp_4f:36:23 (00:08:74:4f:36:23)
 Internet Protocol Version 4, Src: 128.119.245.12 (128.119.245.12), Dst: 192.168.1.102 (192.168.1.102)
 Transmission Control Protocol, Src Port: http (80), Dst Port: http (80), Seq: 1, Ack: 502, Len: 1460
 Source port: http (80)

3270 68 65 20 43 6f 6e 76 65 6e 74 69 6f 6e 73 20 6f 68 65 20 43 6f 6e 76 65 6e 74 69 6f 6e 73 20 6f
 3280 66 20 61 20 6e 75 6d 62 65 72 20 6f 66 20 74 68 66 20 6f 66 20 74 68 66 20 6f 66 20 74 68 66
 3290 65 20 53 74 61 61 74 65 73 20 68 61 76 69 6e 67 2c 6e 69 6e 67 2c 6e 69 6e 67 2c 6e 69 6e
 32a0 20 61 74 20 74 68 65 20 74 69 6d 65 20 6f 66 20 74 69 6d 65 20 6f 66 20 74 69 6d 65 20 6f
 32b0 61 64 6f 70 74 69 6e 67 0a 74 68 65 20 43 6f 6e 67 6e 67 6e 67 6e 67 6e 67 6e 67 6e 67 6e
 32c0 73 74 69 74 75 74 69 6f 6e 2c 20 65 78 70 72 65 6e 2c 20 65 78 70 72 65 6e 2c 20 65 78
 32d0 73 73 65 64 20 61 20 64 65 73 69 72 65 2c 20 69 65 73 69 72 65 2c 20 69 65 73 69 72 65
 32e0 6e 20 6f 72 64 65 72 20 74 6f 20 70 72 65 76 65 6e 20 70 72 65 76 65 6e 20 70 72 65 76
 32f0 6e 74 20 6d 69 73 63 6f 6e 73 74 72 75 63 74 69 6e 73 74 72 75 63 74 69 6e 73 74 72 75
 3300 6f 6e 0a 6f 72 20 61 62 75 73 65 20 6f 66 20 69 6e 73 65 20 6f 66 20 69 6e 73 65 20 6f
 3310 74 73 20 70 6f 77 65 72 73 2c 20 74 68 61 74 20 73 2c 20 74 68 61 74 20 73 2c 20 74 68 61
 3320 66 75 72 74 68 65 72 20 64 65 63 6c 61 72 61 74 66 75 72 74 68 65 72 20 64 65 63 6c 61
 3330 6f 72 79 20 61 6e 64 20 72 65 73 74 72 69 63 74 66 75 72 79 20 61 6e 64 20 72 65 73 74
 3340 69 76 65 20 63 6c 61 75 73 65 73 0a 73 68 6f 75 66 65 65 65 65 65 65 65 65 65 65 65 65
 3350 66 65

The HTTP repl7 carrying the text of the Bill of Rights are packets 10, 11, and 13. If you look into the ASCII content of packet 10, you can see the beginning of the text of the Bill or Rights. Note that packet 12 is a client-to-server TCP ACK.

Answer the following questions:

- How many HTTP GET request messages did your browser send? *Answer: 1.* Which packet number in the trace contains the GET message for the Bill or Rights? *Answer: 8.*
- Which packet number in the trace contains the status code and phrase associated with the response to the HTTP GET request? *Answer: packet 10.*
- What is the status code and Phrase in the response? *Answer: 200 (OK)*
- How many data-containing TCP segments were needed to carry the single HTTP response and the text of the Bill of Rights? *Answer: three packets (10, 11, 13 in the trace)*

4. HTML Documents with Embedded Objects

In our answers below, we use the http-ethereal-trace-4 packet trace file.

Answer the following questions:

- How many HTTP GET request messages did your browser send? *Answer: there were three HTTP GET messages sent: packet 10 in the trace (to get the base file), packet 17 (to get the Pearson logo) and packet 20 (to get the 5th edition textbook cover).* To which Internet addresses were these GET requests sent? Each of these three GET messages were sent to different IP addresses! Packet 10 was sent to 128.119.245.12, packet 17 to 165.193.123.218, and packet 20 to 134.241.6.82.

17. Can you tell whether your browser downloaded the two images serially, or whether they were downloaded from the two web sites in parallel? Explain. *Answer: The downloads occurred in parallel. Note that the two GET messages for the images are in packets 17 and 20. The 200OK reply containing the images show up as packets 25, and 54. Thus the request for the second image file (packet 20) was made BEFORE packet 25, the first image file was received.*

5 HTTP Authentication

Finally, let's try visiting a web site that is password-protected and examine the sequence of HTTP message exchanged for such a site. The URL http://gaia.cs.umass.edu/wireshark-labs/protected_pages/HTTP-wireshark-file5.html is password protected. The username is "wireshark-students" (without the quotes), and the password is "network" (again, without the quotes). So let's access this "secure" password-protected site. Do the following:

- Make sure your browser's cache is cleared, as discussed above, and close down your browser. Then, start up your browser
- Start up the Wireshark packet sniffer
- Enter the following URL into your browser
http://gaia.cs.umass.edu/wireshark-labs/protected_pages/HTTP-wireshark-file5.html
Type the requested user name and password into the pop up box.
- Stop Wireshark packet capture, and enter "http" in the display-filter-specification window, so that only captured HTTP messages will be displayed later in the packet-listing window.
- (*Note: If you are unable to run Wireshark on a live network connection, you can use the http-ethereal-trace-5 packet trace to answer the questions below; see footnote 2. This trace file was gathered while performing the steps above on one of the author's computers.*)

Now let's examine the Wireshark output. You might want to first read up on HTTP authentication by reviewing the easy-to-read material on "HTTP Access Authentication Framework" at [http://frontier.userland.com/stories/storyReader\\$2159](http://frontier.userland.com/stories/storyReader$2159)

Answer the following questions:

18. What is the server's response (status code and phrase) in response to the initial HTTP GET message from your browser? *Answer: Packet 6 in the trace contains the first GET and packet 9 contains the REPLY. The server's in packet 9 is: 401 Authorization Required*
19. When your browser's sends the HTTP GET message for the second time, what new field is included in the HTTP GET message? *Answer: The HTTP GET includes the Authorization: Basic: field*