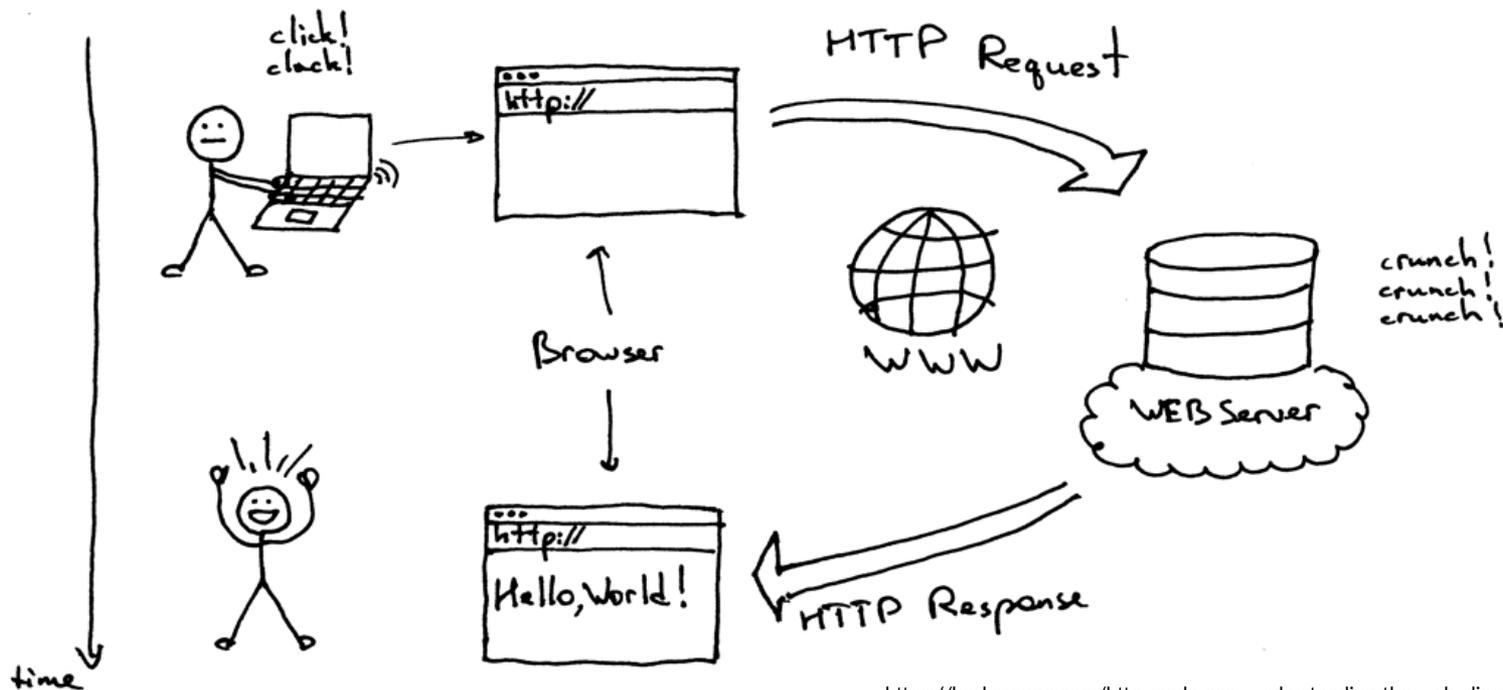


<http://www.>



# \*The\* Web Services Protocol

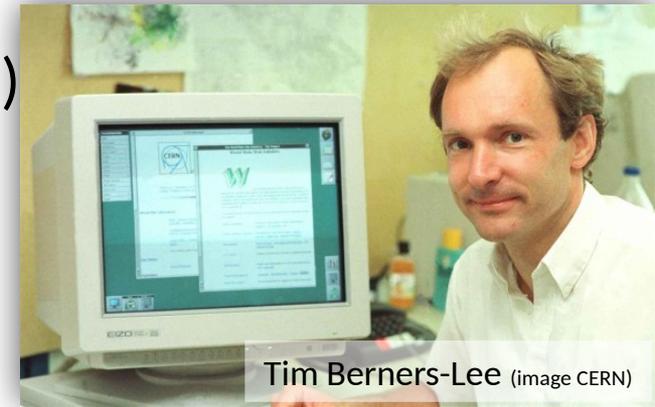
- Application-layer protocol
- Client sends service requests using HTTP messages
- Server replies using HTTP messages





# Hypertext Transport Protocol

- 1989-90, Tim Berners-Lee's problem at CERN:  
how to integrate and exchange information held on different computers in scattered places?
- Already exist:
  - TCP: reliable transport of information on the Internet
  - DNS: domain name ("www.centralesupelec.fr") ↔ IP @ ("138.195.9.117")
    - Human friendly
    - Computer friendly
  - object in a database that *references* others
- Put them all together: HTTP
  - Retrieve linked documents (resources)
  - Accessible via the Internet



Tim Berners-Lee (image CERN)

# Client-Server Protocol request / response

Client

(web browser or other application)

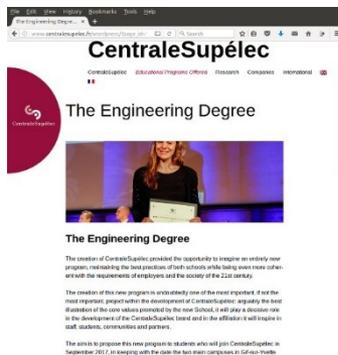


1. user clicks on hyperlink

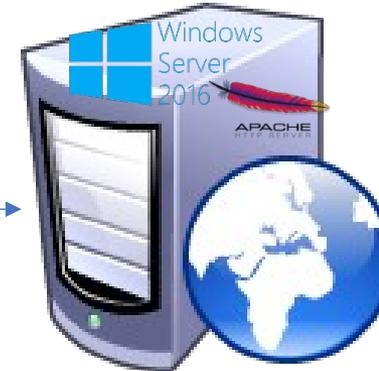
2. HTTP request message

4. HTTP response message

5. display file



web server



back-end  
(database...)

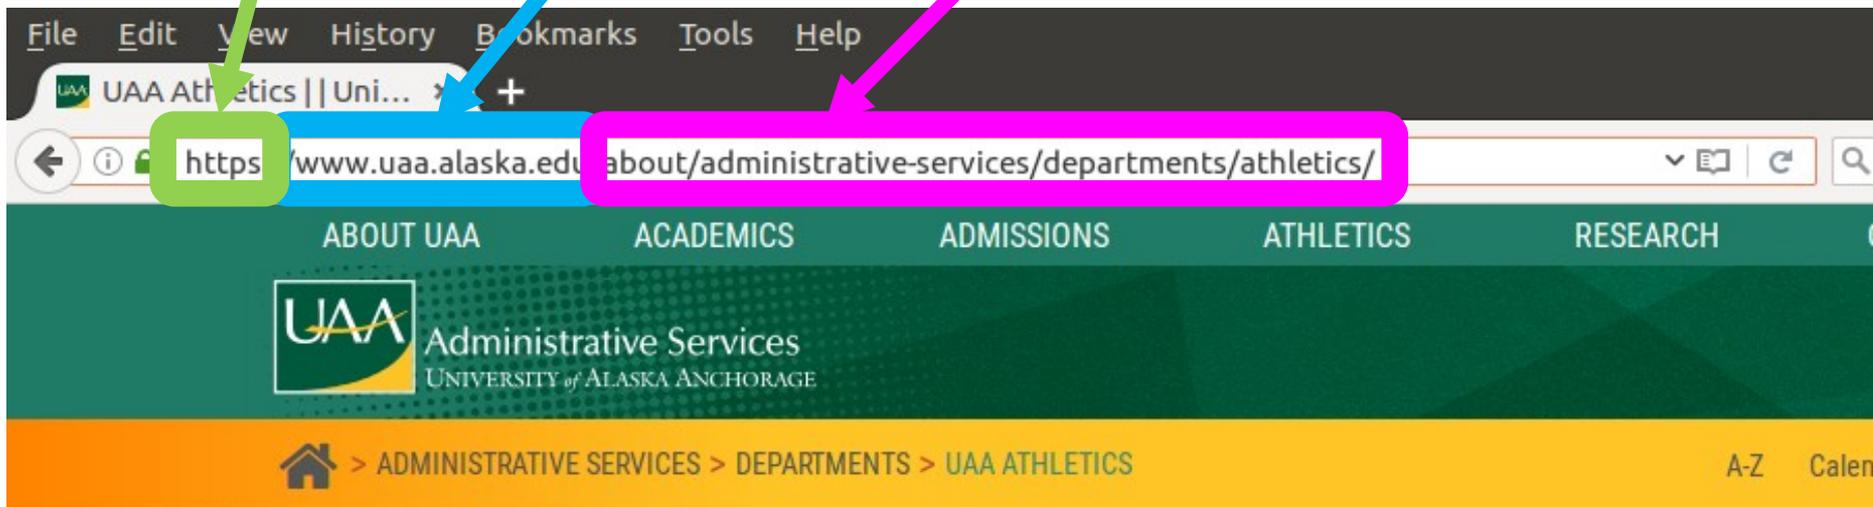


3. create or retrieve file



# Resources = addressable files

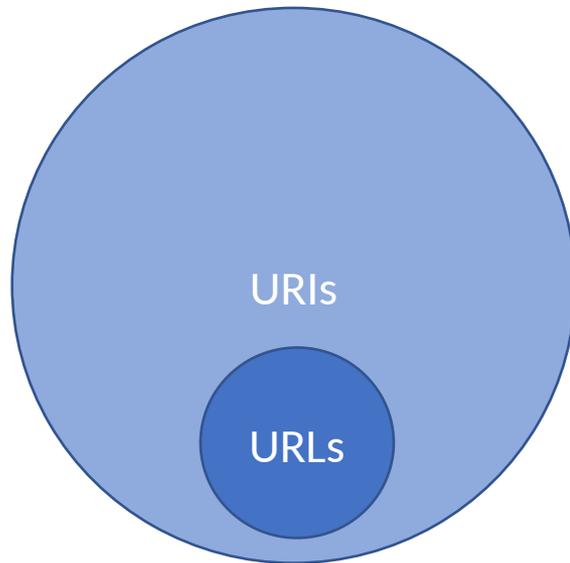
- Any kind of file: HTML file, JPG image file, binary file...
- URL (Uniform Resource Locator)  
= protocol + server host name + path on server



**UAA Athletics**

# Side note about URL and URI

- URI: identifier (name of a restaurant)
- URL: locator (GPS coordinates of the restaurant)



All URLs are URIs:

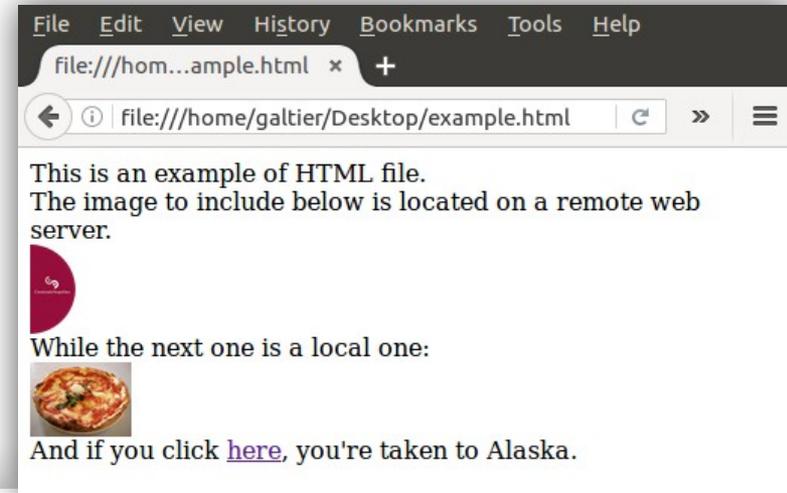
with the GPS coordinates I arrive to the right restaurant

Not all URIs are URLs:

the name of the restaurant gives no information on its location

# HTML file may include references to others resources

- 3 resources are required to display this web page:
  - HTML file
  - CentraleSupelec logo image
  - Pizza image



```
File Edit View Search Tools Documents Help
Open [icon] Save
<html>
<head><title="HTML Example"/></head>
<body>
This is an example of HTML file.</br>
The image to include below is located on a remote web server.</br>
</br>
While the next one is a local one:</br>
</br>
And if you click <a href="https://www.uaa.alaska.edu/about/administrative-services/departments/athletics/">here</a>, you're
taken to Alaska.
</body>
</html>
HTML Tab Width: 8 Ln 5. Col 20
```

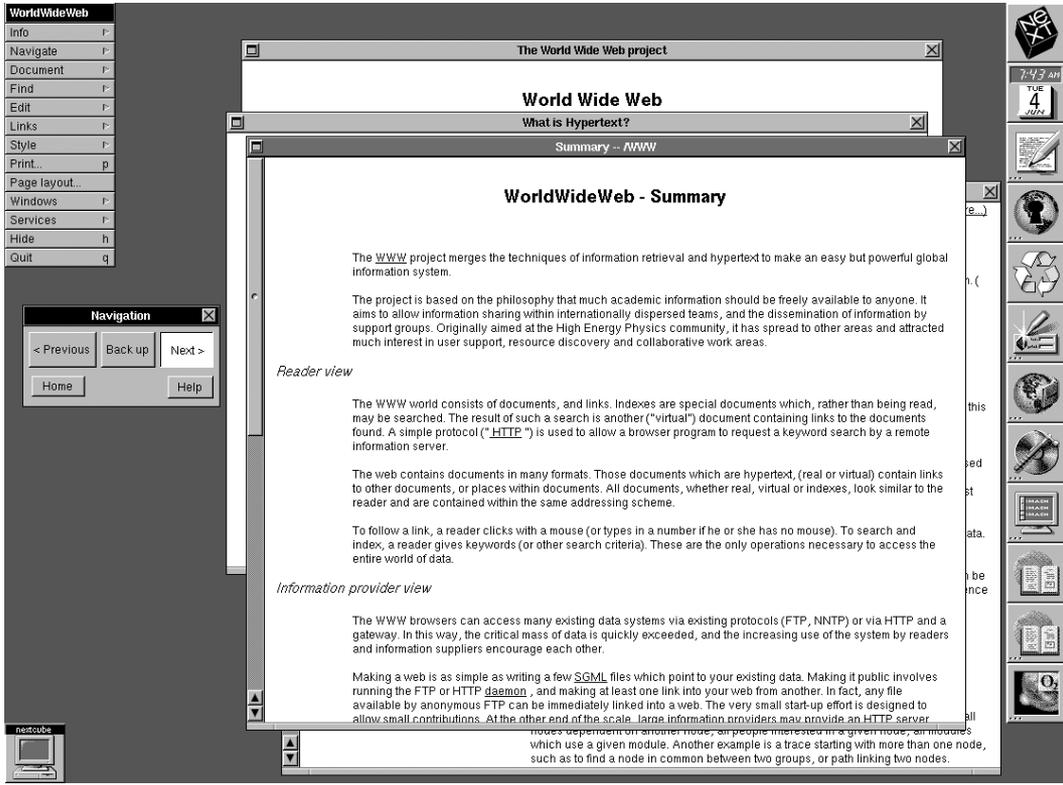
# HTTP Versions

HTTP/0.9

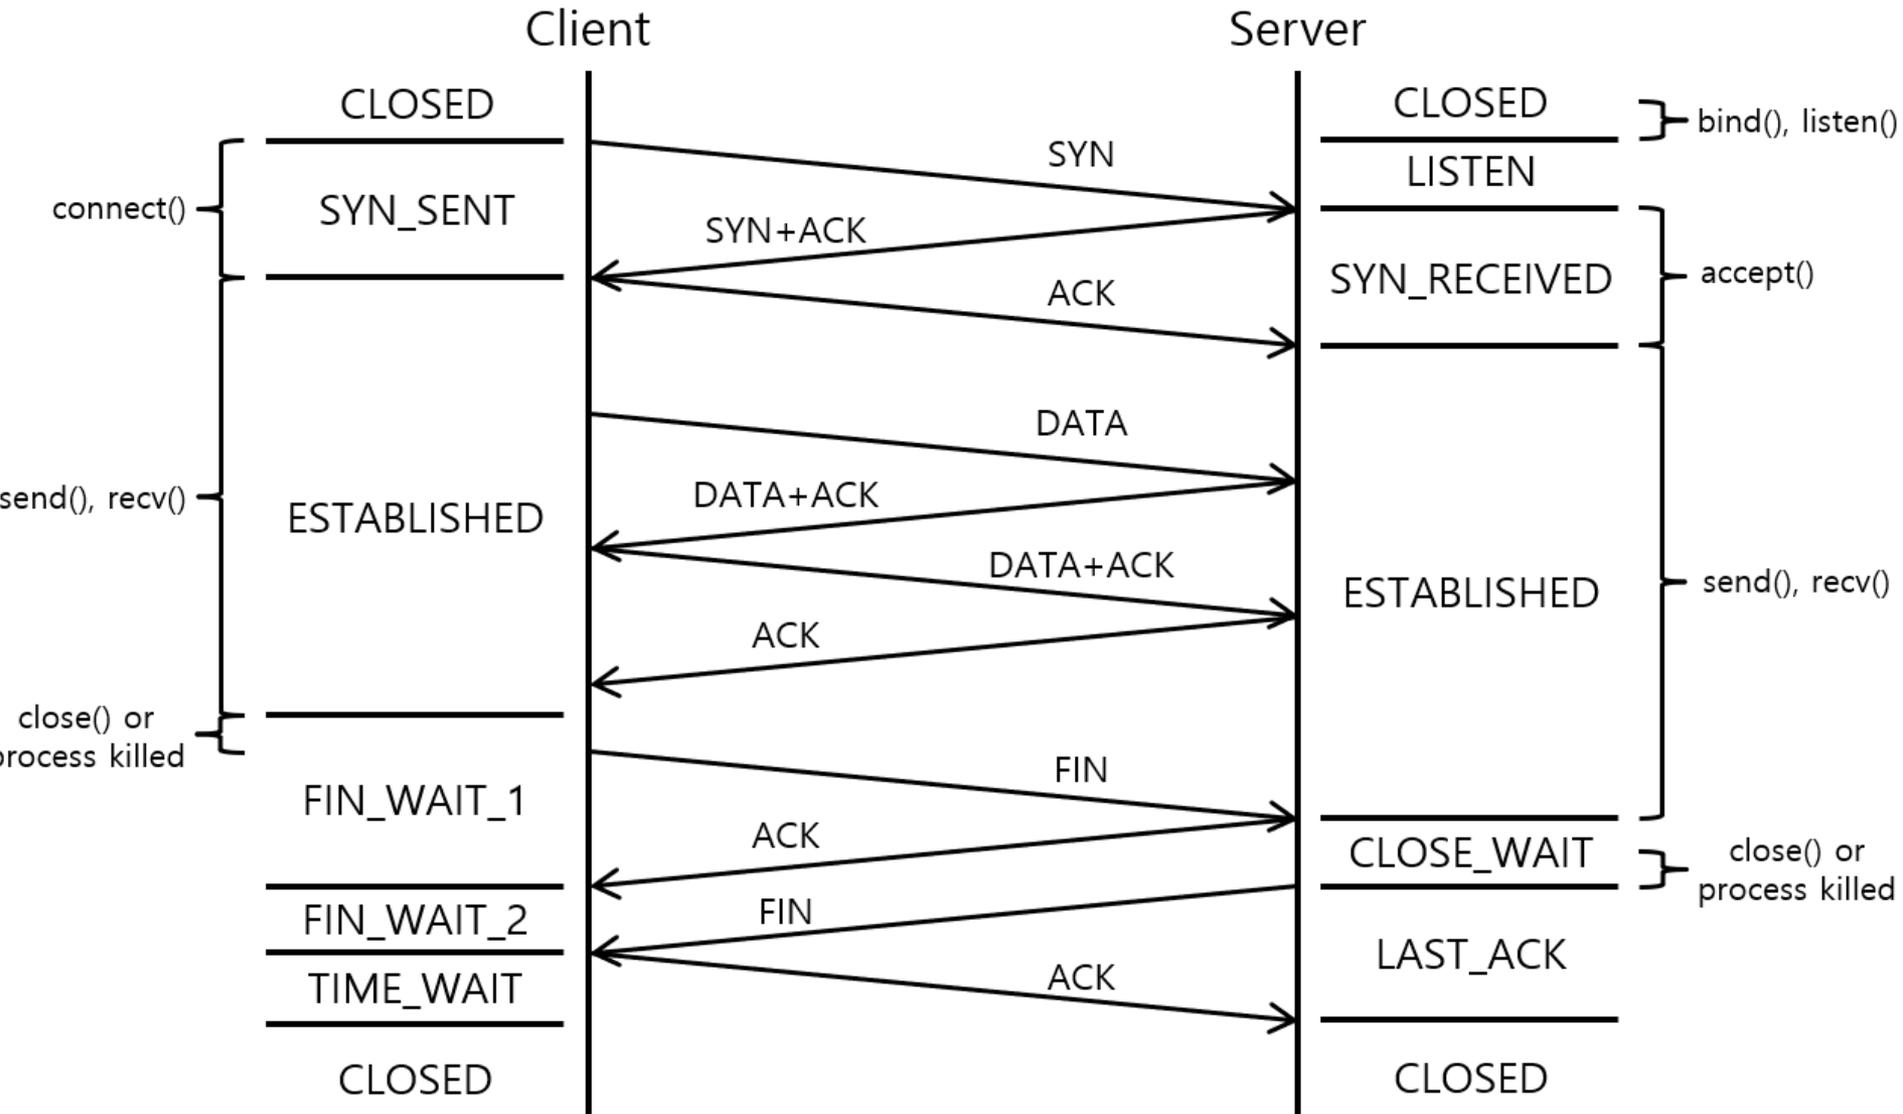


1991

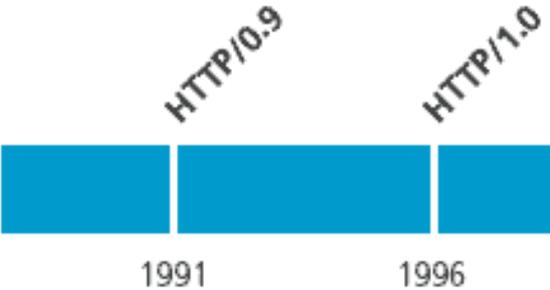
- 1991 – v0.9
  - First documented version
  - First web browser



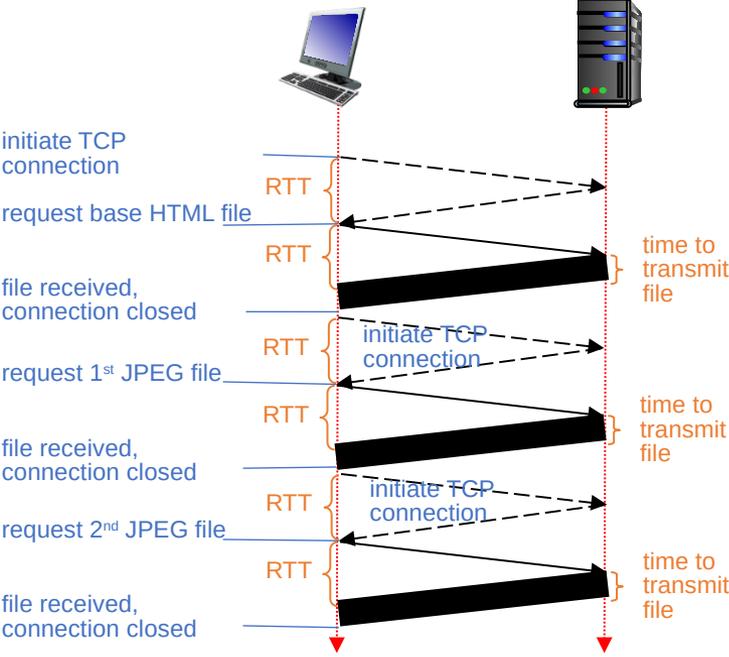
# Side note on TCP 3-way handshake



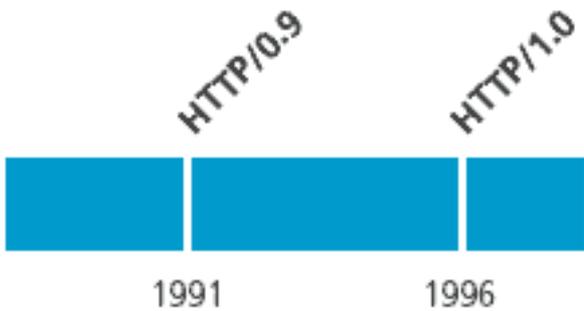
# HTTP Versions



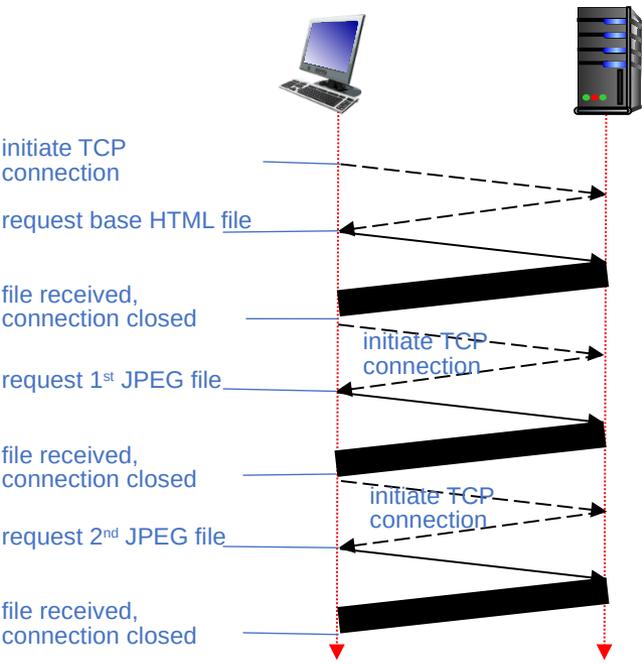
- 1996 – v1.0
  - One TCP connection per resource



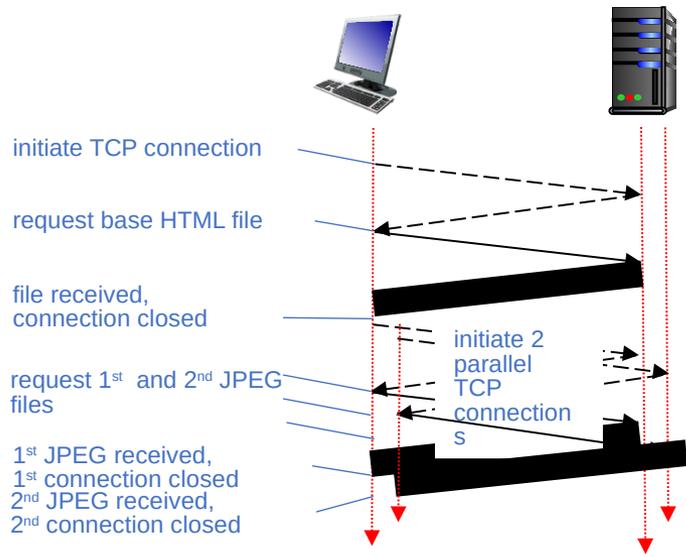
# HTTP Versions



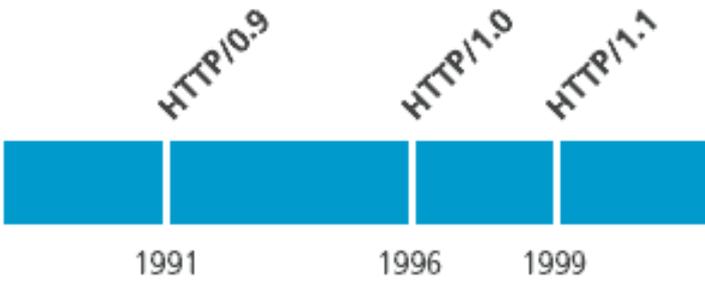
- 1996 – v1.0
  - One TCP connection per resource



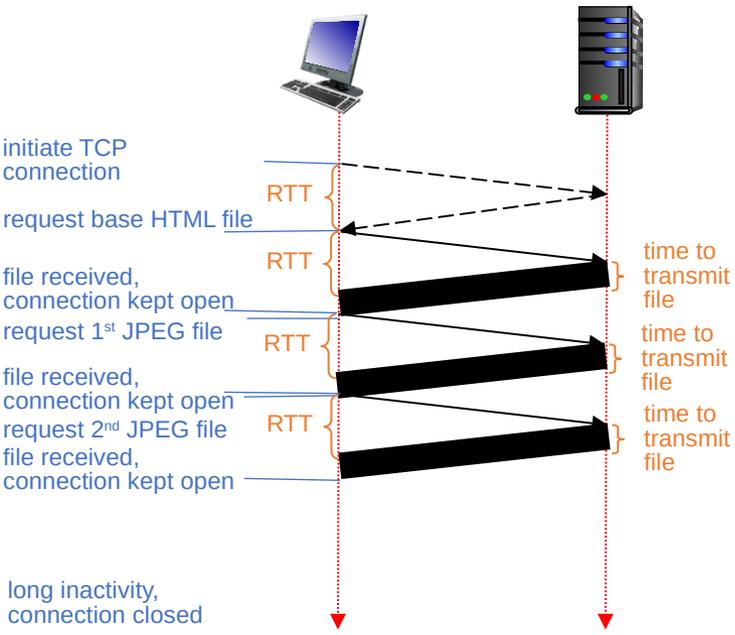
with parallel connections



# HTTP Versions



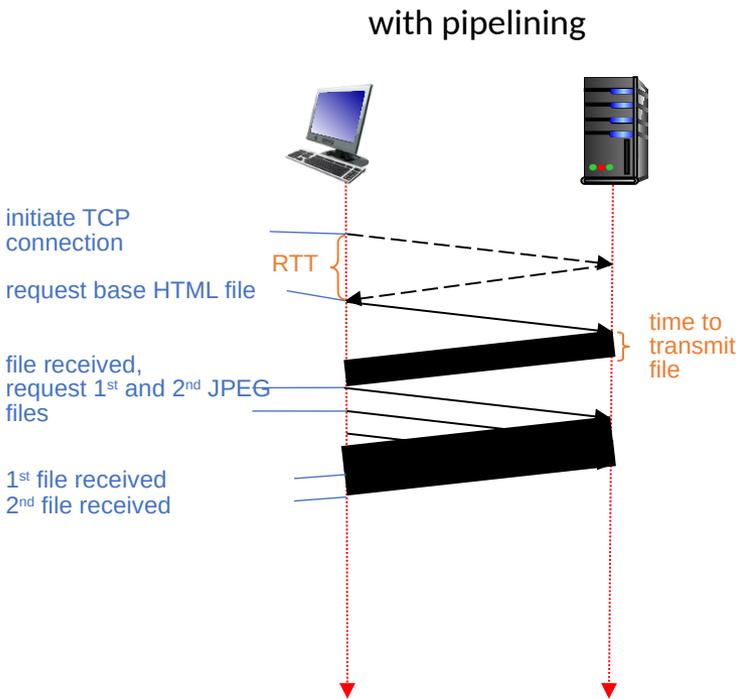
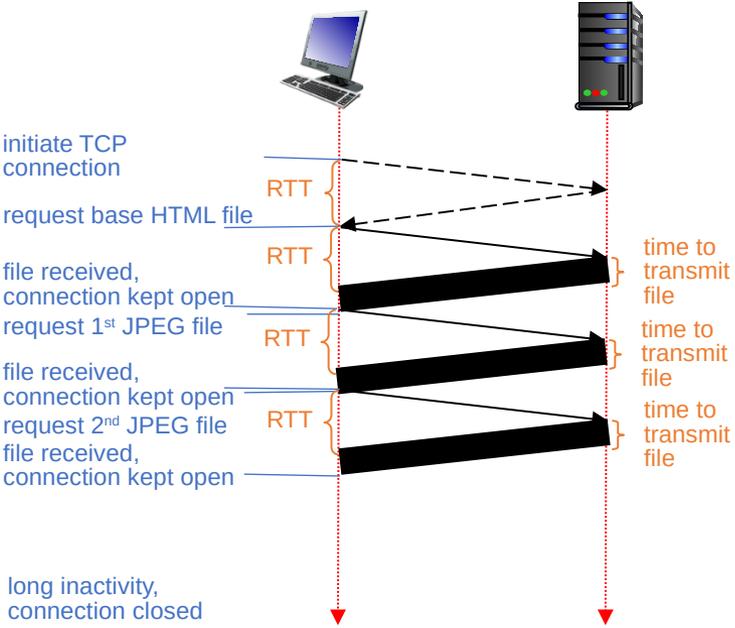
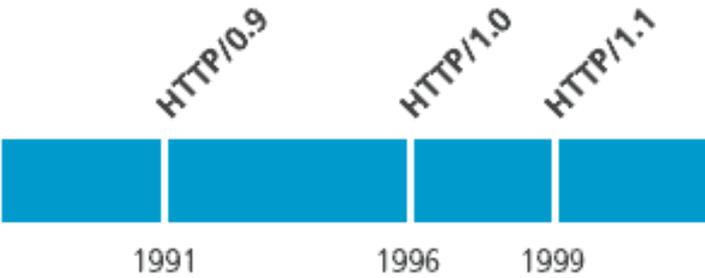
- 1999 - v1.1
  - Persistent connection



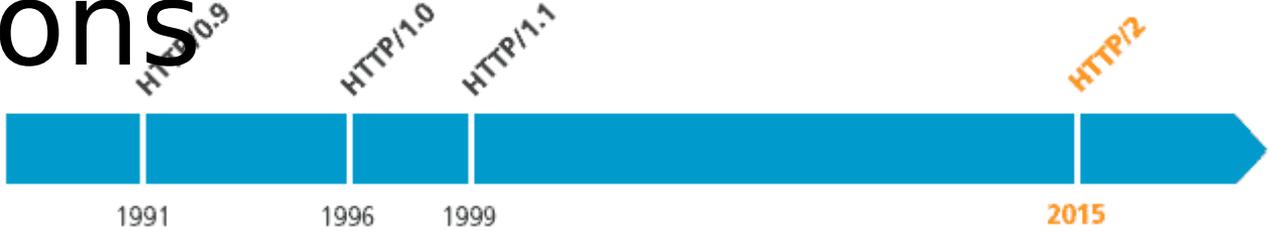
```
GET https://www.centralesupelec.fr/ HTTP/1.1
Host: www.centralesupelec.fr
User-Agent: Mozilla/5.0 (X11; Ubuntu; Linux x86
Accept: text/html,application/xhtml+xml,applica
Accept-Language: en-US,en;q=0.5
Accept-Encoding: gzip, deflate, br
Connection: keep-alive
Upgrade-Insecure-Requests: 1
```

# HTTP Versions

- 1999 - v1.1
  - Persistent connection

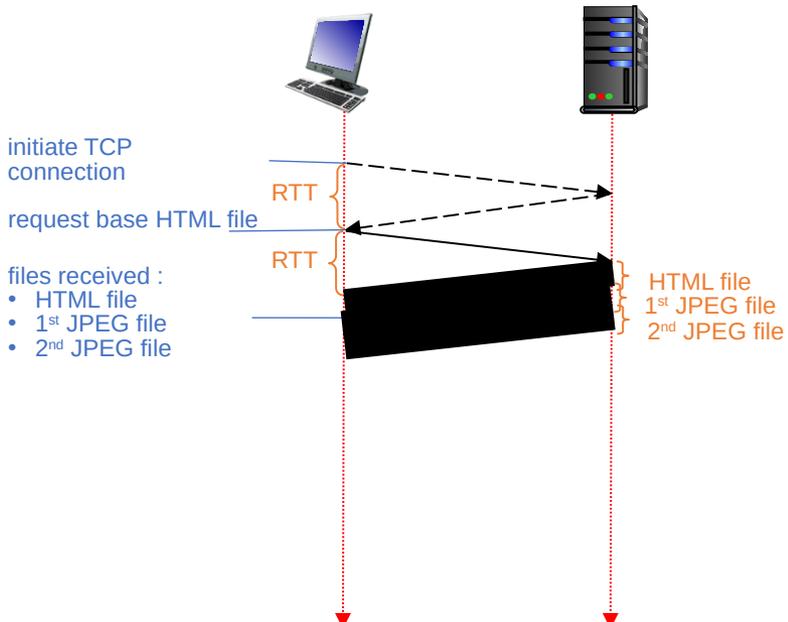


# HTTP Versions



- 2015 – v2
  - Server “pushes” content
  - [and other optimizations]

A screenshot of a web browser displaying a page from Akamai titled "HTTP/2: the Future of the Web, and it is here!". The page features the Akamai logo and a headline "HTTP/2 is the future of the Web, and it is here!". Below the headline, it says "Your browser supports HTTP/2!". A paragraph explains that this is a demo of HTTP/2's impact on downloading many small tiles for the Akamai Spinning Globe. Two side-by-side images of the globe are shown, each with a list of metrics: "Latency: 27ms" and "Load time: 3.23s" for HTTP/1.1, and "Latency: 27ms" and "Load time: 1.17s" for HTTP/2. At the bottom, it says "Demo content inspired by Galaxie".



# Optional Reading Exercise

- Find the document which describes HTTP/2.
- What is the “head-of-line blocking” (HOL blocking) problem observed in HTTP/1.1?
- Read the beginning of the FAQ at <https://http2.github.io/faq/>

# Reading: Results

- HTTP/2 is defined in RFC 7540.
- HOL blocking:
  - Imagine a HTTP client that sends to a server 2 requests over the same TCP connection, and that the first response is "large" in content length while the second response is "small" in content length.
  - Due to the nature of the HTTP 1.x protocol, the second response must wait for the first response to complete: the second response is *head-of-line blocked* by the first response.
  - HTTP/2 is fully multiplexed (instead of ordered and blocking), allowing multiple request and response messages to be in flight at the same time (it's even possible to intermingle parts of one message with another on the wire).

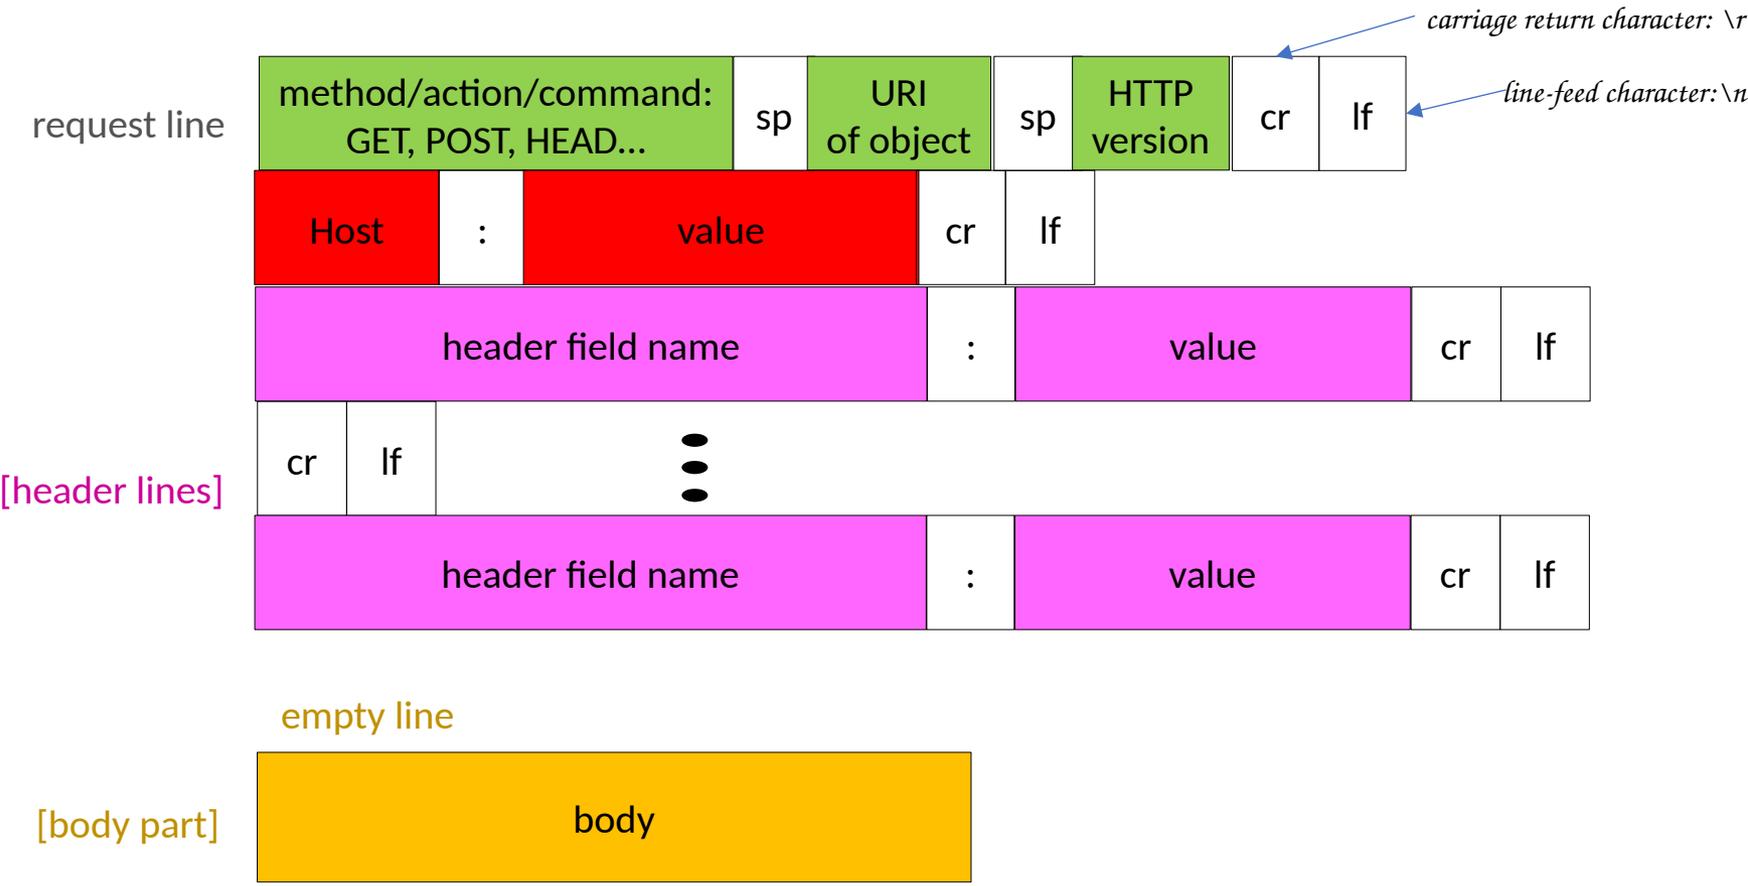
# HTTP Messages

- 2 kinds of messages
  - Request
  - Response
- In ASCII (HTTP 1.x)

# HTTP Requests Commands

- GET
  - retrieves an object
  - no request body
- HEAD
  - same response as GET but empty response body (used to test the access to or the "freshness" of the object without actually downloading it)
- POST
  - results in the creation of a new resource on the server
  - usual request: contains data
  - usual response: URL of the created resource
- PUT
  - updates an existing resource
  - request usually contains data
- DELETE
  - deletes a resource

# HTTP Request Format



# HTTP GET Request Example

GET /node/44 HTTP/1.1\r\n

Host: mapi.centralesupelec.fr\r\n

User-Agent: Mozilla/5.0 (X11; Ubuntu; Linux x86\_64; rv:50.0) Gecko/20100101 Firefox/50.0\r\n

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

Connection: keep-alive\r\n

\r\n

# HTTP POST Request Example

POST /post.php HTTP/1.1\r\n

Host: posttestserver.com\r\n

User-Agent: Mozilla/5.0 (X11; Ubuntu; Linux x86\_64; rv:50.0) Gecko/20100101 Firefox/50.0\r\n

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

Content-Type: text/xml\r\n

Content-Length: 27\r\n

Connection: keep-alive\r\n

\r\n

firstname=John\nlastname=Doe

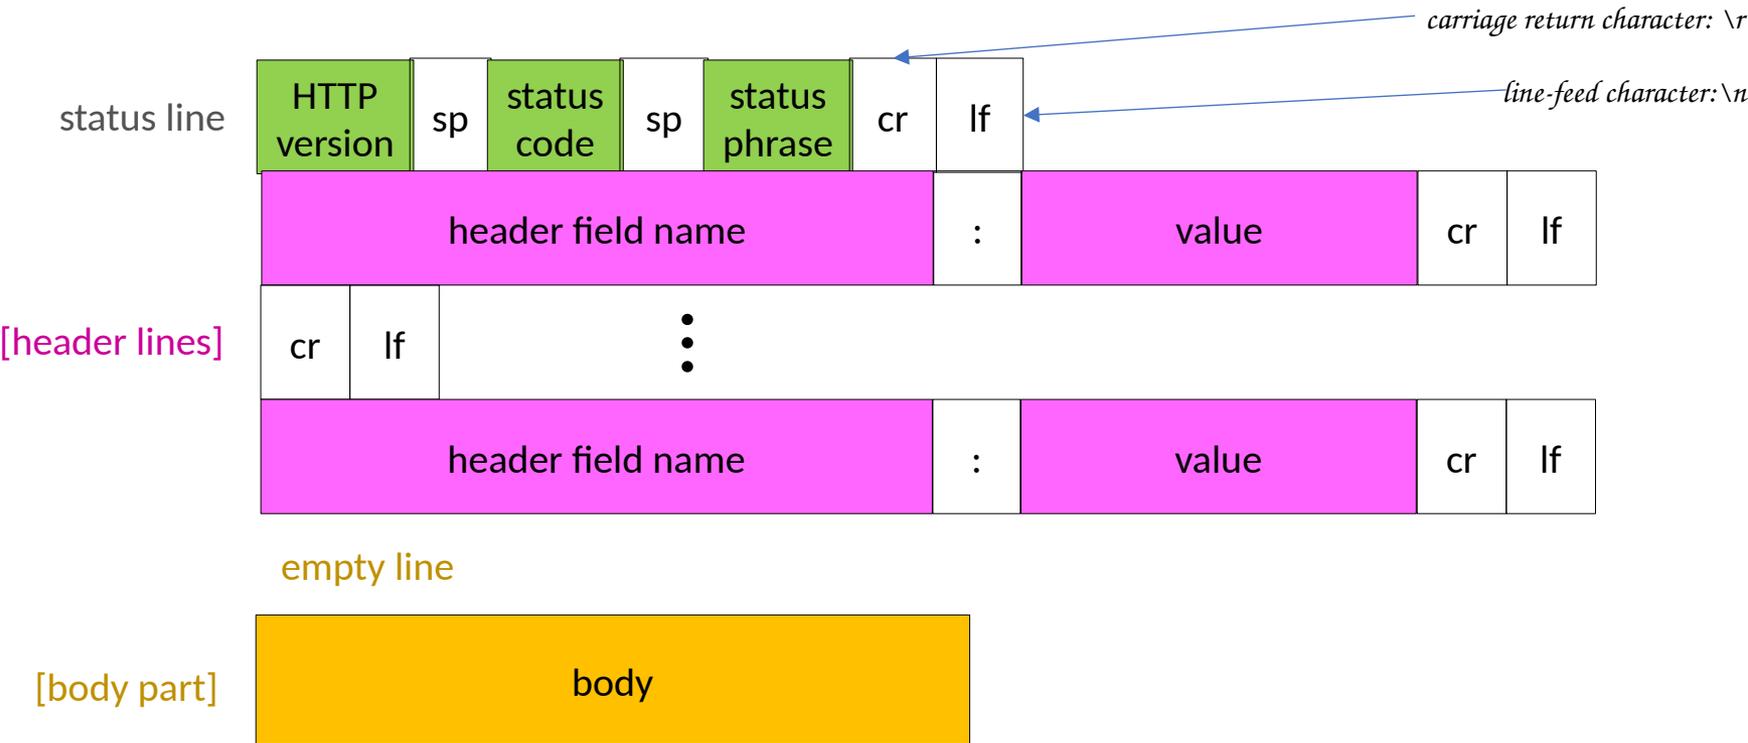
# Request Parameters

3 symbols to add parameters to an URL:

- **?** concatenates the URL and the string of parameters
- **&** separates multiple parameters
- **=** assigns a value to a parameter

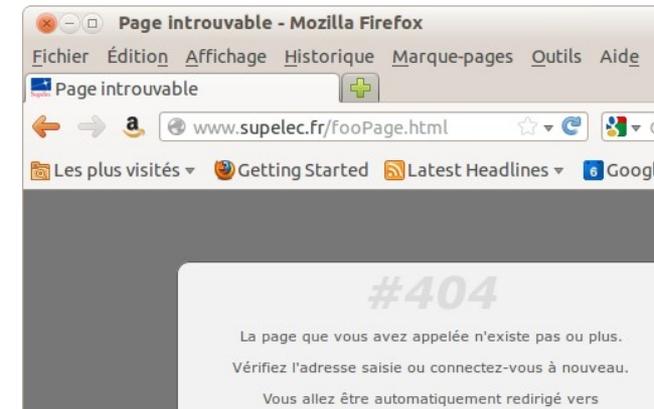
GET /products?priceMin=10&priceMax=40

# HTTP Response Format



# Status Codes

- 2xx: success
  - 200 OK
- 3xx: further action required
  - 301 Moved Permanently: the new URL is specified in a header field
- 4xx: client error
  - 400 Bad Request: badly formulated query
  - 404 Not Found: object does not exist on the server
- 5xx: server-side error
  - 505 HTTP Version Not Supported



# HTTP Response Example

HTTP/1.1 200 OK

Date: Wed, 01 Feb 2017 12:48:22 GMT

Server: Apache/2.4.10 (Debian)

[...]

Content-language: fr

Content-Encoding: gzip

Content-Length: 4740

Keep-Alive: timeout=5, max=100

Connection: Keep-Alive

Content-Type: text/html; charset=UTF-8

.....;r.8...W...-{(.KO...!K..-.\$..q;.(...D.4..T.v.\.....q.{...Z.2\_....b.....(l....Df" ..Hn... |....}.WQ..m....  
WD.sR)..J.....L:9.C..MC...X.I...  
J..(...'" .....J. D....d%bN,. \$..Y.....z.....y(.MS....#.qV.....>.9.j.0  
s&...v.M...'').....m8..<=.i..%B.....S.x}.J.:V..{." .HM..4b..!YJ.....X{i...l.;.TX}....N.r .<d...#.....S..  
..#Oa. ...V..EPj..G...A..D.K...Z1..c.h,b.4..b.3...l.6..La..>.L8#l.U.\.....2..y!...S.,.....%.....>..ID...  
N..^

# HTTP Response Example

HTTP/1.1 404 Not Found

Date: Wed, 01 Feb 2017 13:14:55 GMT

Server: Apache/2.4.10 (Debian)

[...]

Content-language: fr

Keep-Alive: timeout=5, max=100

Connection: Keep-Alive

Content-Type: text/html; charset=UTF-8

305d

```
<!DOCTYPE html>
```

```
<html lang="fr" dir="ltr" prefix="content: http://purl.org/rss/1.0/modules/content/ dc: ht
```

```
<div id="block-zircon-content" class="block block-system block-system-main-block">
```

```
  La page demand..e n'a pas pu ..tre trouv..e.
```

```
</div>
```

# Optional Reading

- What is HTTP error code 418?



# Hyper Text Coffee Pot Control Protocol

 14 languages 

[Article](#) [Talk](#)

[Read](#) [Edit](#) [View history](#)

From Wikipedia, the free encyclopedia

The **Hyper Text Coffee Pot Control Protocol** (**HTCPCP**) is a facetious [communication protocol](#) for controlling, monitoring, and diagnosing [coffee pots](#). It is specified in [RFC 2324](#)<sup>[ ]</sup>, published on 1 April 1998 as an [April Fools' Day RFC](#),<sup>[2]</sup> as part of an [April Fools prank](#).<sup>[3]</sup> An extension, HTCPCP-TEA, was published as RFC 7168 on 1 April 2014<sup>[4]</sup> to support brewing teas, which is also an April Fools' Day RFC.

## Protocol [[edit](#)]

RFC 2324 was written by [Larry Masinter](#), who describes it as a satire, saying "This has a serious purpose – it identifies many of the ways in which [HTTP](#) has been extended inappropriately."<sup>[5]</sup> The wording of the protocol made it clear that it was not entirely serious; for example, it notes that "there is a strong, dark, rich requirement for a protocol designed [espressoly](#) <sup>[sic]</sup> for the brewing of coffee".

Despite the joking nature of its origins, or perhaps because of it, the protocol has remained as a minor presence online. The editor [Emacs](#) includes a fully functional client side implementation of it,<sup>[6]</sup> and a number of bug reports exist complaining about [Mozilla](#)'s lack of support for the protocol.<sup>[7]</sup> Ten years after the publication of HTCPCP, the *Web-Controlled Coffee Consortium* (*WC3*) published a first draft of "HTCPCP Vocabulary in [RDF](#)"<sup>[8]</sup> in parody of the [World Wide Web Consortium](#)'s (W3C) "HTTP Vocabulary in RDF".<sup>[9]</sup>

On April 1, 2014, RFC 7168 extended HTCPCP to fully handle teapots.<sup>[4]</sup>

## Commands and replies [[edit](#)]

HTCPCP is an extension of [HTTP](#). HTCPCP requests are identified with the [Uniform Resource Identifier](#) (URI) scheme `coffee` (or the corresponding word in any other of the 29 listed languages) and contain several additions to the HTTP methods:

<b>BREW</b> or <b>POST</b>	Causes the HTCPCP server to brew <a href="#">coffee</a> . Using POST for this purpose is deprecated. A new <a href="#">HTTP request header field</a> "Accept-Additions" is proposed, supporting optional additions including Cream, Whole-milk, Vanilla, Raspberry, Whisky, Aquavit, etc.
<b>GET</b>	"Retrieves" coffee from the HTCPCP server.
<b>PROPFIND</b>	Returns <a href="#">metadata</a> about the coffee.
<b>WHEN</b>	<i>Says "when"</i> , causing the HTCPCP server to stop pouring <a href="#">milk</a> into the coffee (if applicable).

It also defines two [error responses](#):

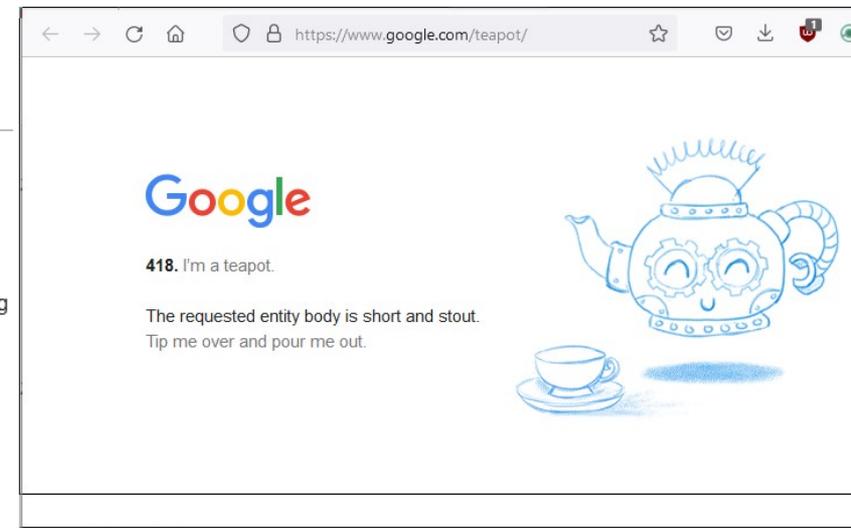
<b>406 Not Acceptable</b>	The HTCPCP server is unable to provide the requested addition for some reason; the response should indicate a list of available additions. The RFC observes that "In practice, most automated coffee pots cannot currently provide additions."
<b>418 I'm a teapot</b>	The HTCPCP server is a <a href="#">teapot</a> ; the resulting entity body "may be short and stout" (a reference to the song " <a href="#">I'm a Little Teapot</a> "). Demonstrations of this behaviour exist. <sup>[1][10]</sup>

## Hyper Text Coffee Pot Control Protocol



[Back-end infrastructure](#) of error418.net, which implements HTCPCP

<b>International standard</b>	<a href="#">Internet Engineering Task Force</a>
<b>Developed by</b>	<a href="#">Larry Masinter</a>
<b>Introduced</b>	April 1, 1998
<b>Website</b>	<a href="#">rfc2324</a> <sup>[<span> </span>]</sup>



# Optional Lab Exercise

- Use putty or telnet to connect to port 80 of a web server (<http://www.columbia.edu> for instance) and issue HTTP/1.x requests (get /~fdc/sample.html). Observe the responses.

```
telnet serverName 80
```

- For HTTPS, use:

```
openssl s_client -connect  
serverName:443
```

(note: this exercise is limited to HTTP/1.x because HTTP/2 is no longer textual but uses binary format commands)

# Lab: Results

```
telnet www.columbia.edu 80
Trying 128.59.105.24...
Connected to source.failover.cc.columbia.edu.
Escape character is '^]'.
```

opens a TCP connection on web server port 80 and sends everything that is typed

```
HEAD /~fdc/sample.html HTTP/1.1
Host: www.columbia.edu
```

typed out request

```
HTTP/1.1 200 OK
Date: Sun, 19 Feb 2023 09:15:26 GMT
Server: Apache
Last-Modified: Fri, 17 Sep 2021 19:26:14 GMT
Accept-Ranges: bytes
Content-Length: 34974
Vary: Accept-Encoding,User-Agent
Content-Type: text/html
Set-Cookie: BIGipServer~CUIT~www.columbia.edu-80-pool=1764244352
```

received response

# Lab: Results

```
openssl s_client -connect edition.cnn.com:443
```

```
CONNECTED(00000003)
```

```
...CERTIFICATE STUFF...
```

```
---
```

```
GET /travel HTTP/1.1
```

```
Host: edition.cnn.com
```

```
HTTP/1.1 200 OK
```

```
Connection: keep-alive
```

```
Content-Length: 220918
```

```
Content-Type: text/html; charset=utf-8
```

```
cache-control: max-age=60
```

```
Date: Sun, 19 Feb 2023 09:22:49 GMT
```

```
[...]
```

```
<!doctype html><html lang="en"><head><meta http-equiv="x-ua-compatible"
content="ie=edge"/><title data-rh="true">CNN Travel | Global Destinations, Tips
& Video</title><meta data-rh="true" name="theme-color"
content="#31315b"/><meta data-rh="true" charSet="utf-8"/><meta data-rh="true"
```

opens an SSL connection on web server port 443 and sends everything that is typed

typed out request

received response

# Lab: Results

```
openssl s_client -connect edition.cnn.com:443
CONNECTED(00000003)
...CERTIFICATE STUFF...
---
GET /travel HTTP/1.1
host edition.cnn.com

HTTP/1.1 400 Bad Request
Connection: close
Content-Length: 11
content-type: text/plain; charset=utf-8
x-served-by: cache-cdg20763

Bad Requestclosed
```

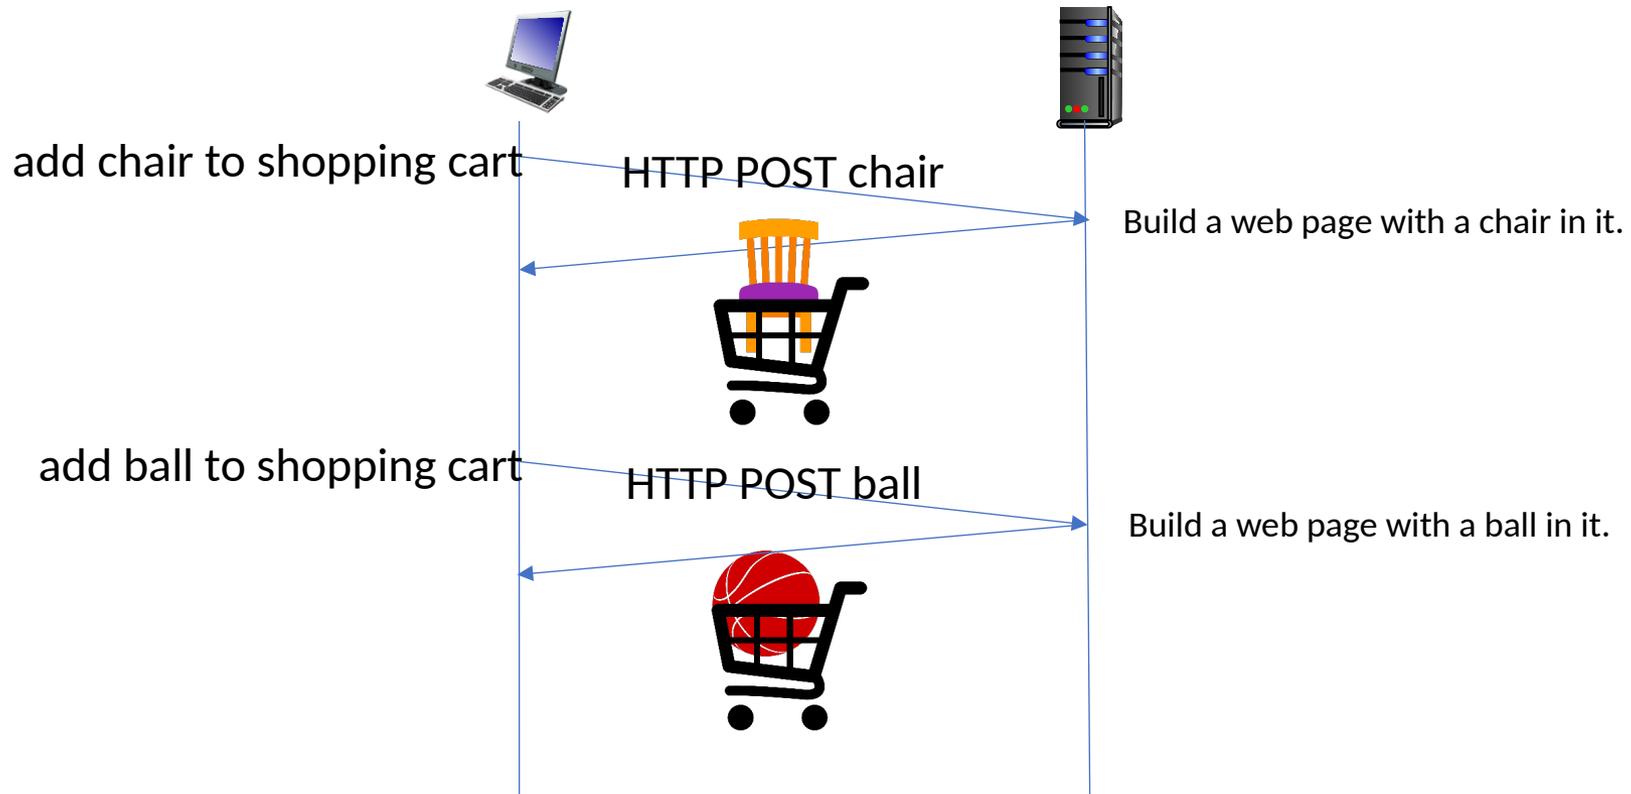
typed out request  
(correct syntax is  
Host: edition.cnn.com)

received response

# HTTP Server is Stateless

- A stateless protocol does not require the server to retain information or status about each user for the duration of multiple requests.
- Successive requests from a given client to a server are not treated as a chain but rather as separate requests, independent from the previous ones.

# What we get is not what we want.



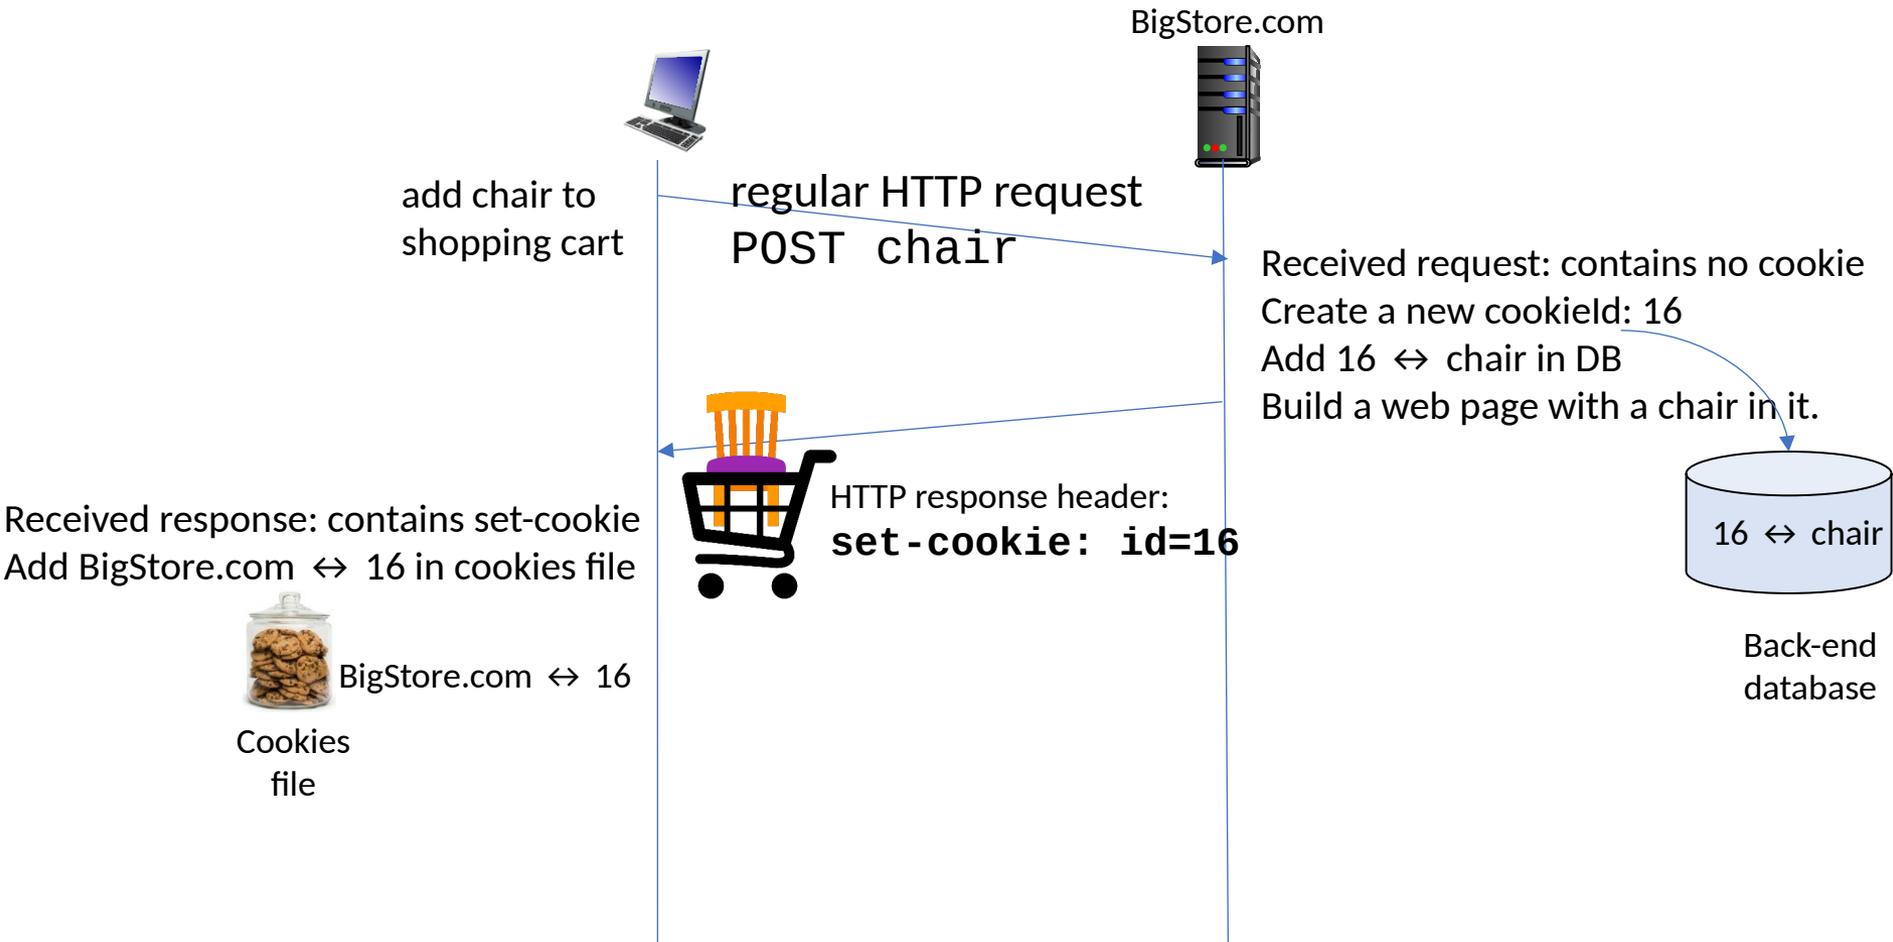
# Cookies

© Original Artist  
Reproduction rights obtainable from  
[www.CartoonStock.com](http://www.CartoonStock.com)

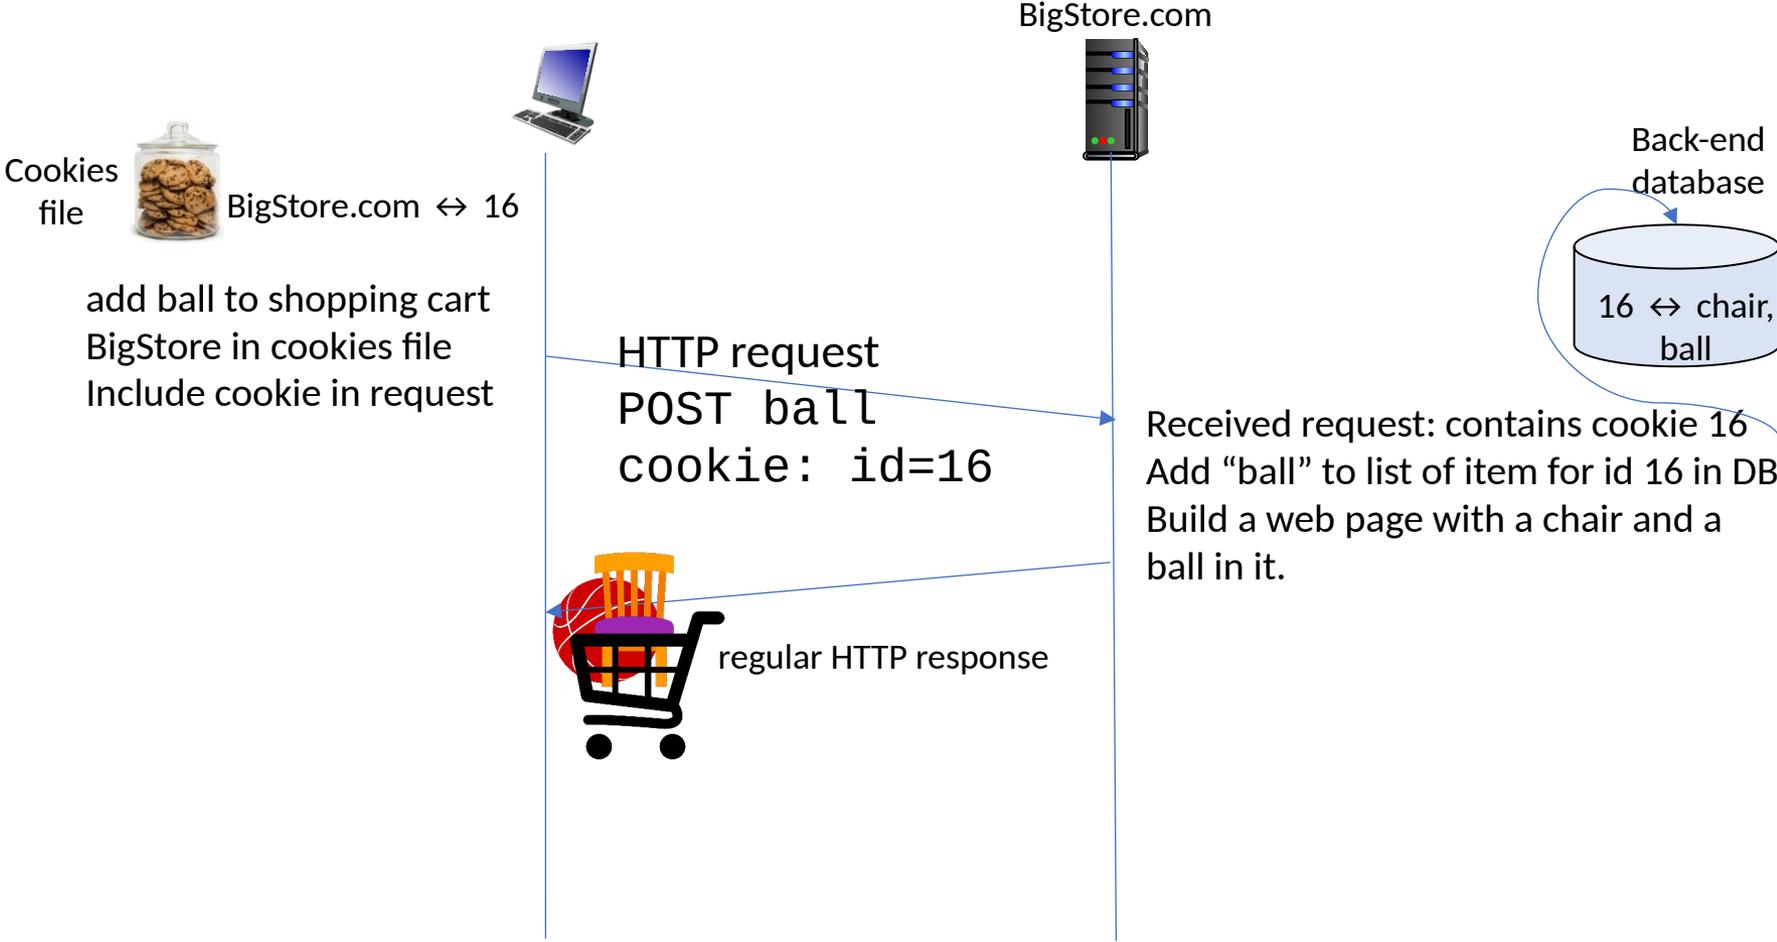
*Shhhhh... they're in here  
somewhere. Dad's always  
talking about the cookies  
in his computer.*



# Cookie Example



# Cookie Example



# Uses

## **create a user session layer on top of stateless HTTP**

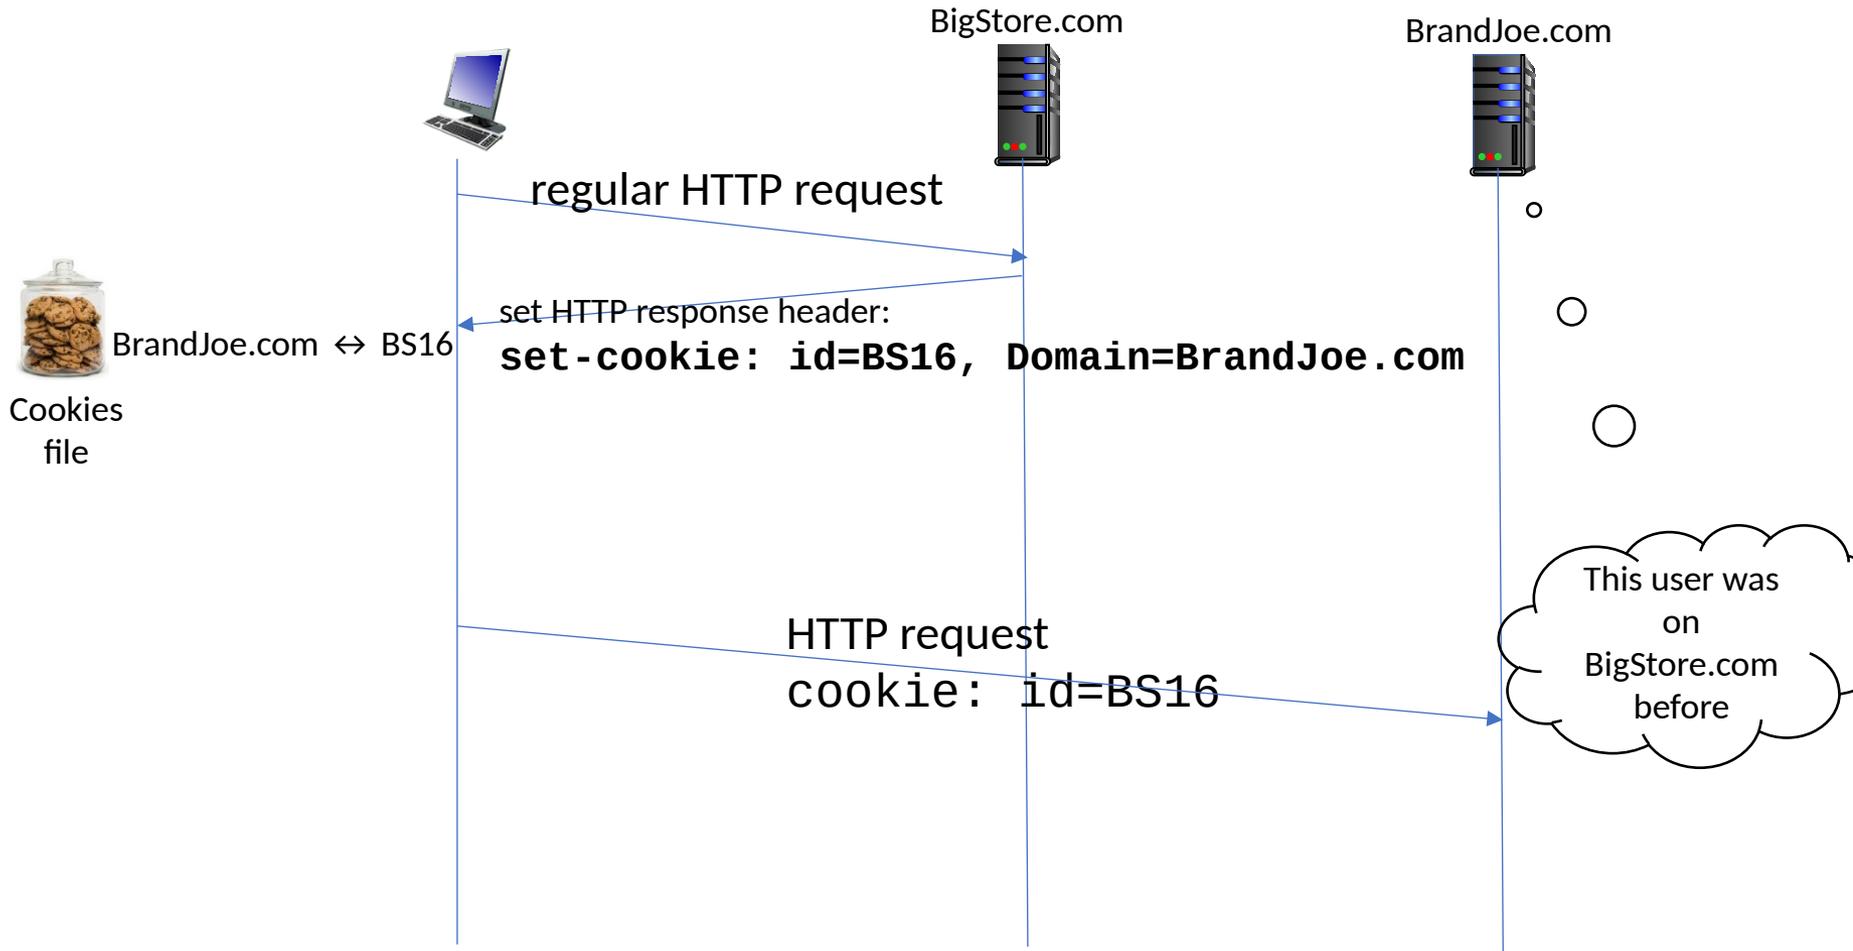
- content adaptation (recommendation based on previous visits etc.).
- shopping carts (e-business)
- session definition at application layer (Web mail)
- authorization
- ...

# Suspicious

- Invasion of privacy

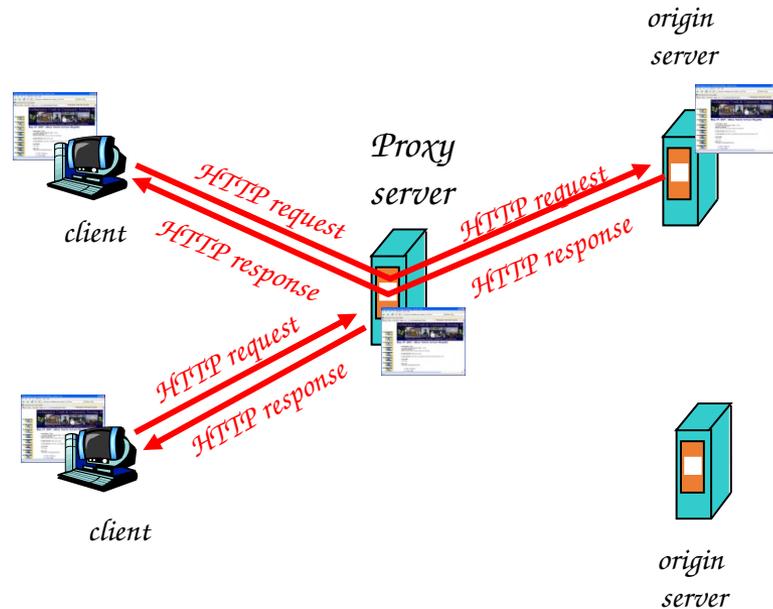


# Third-party advertising cookies



# Web Cache (proxy server)

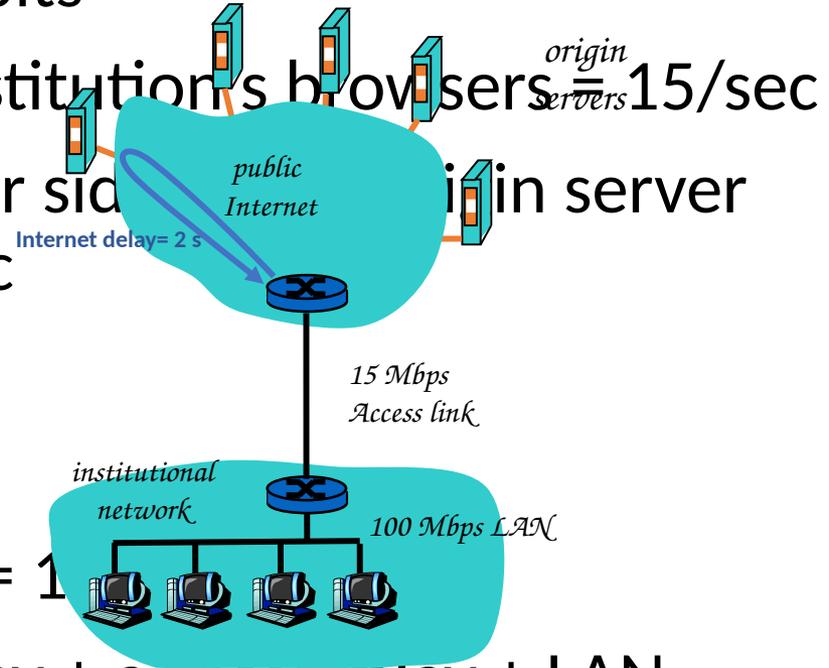
- to satisfy the requests without involving the real server
- browser must be configured to send all HTTP requests to cache
- reduced traffic on Internet, improved response time



# Caching Example

## assumptions

- average object size = 1Mbits
- avg. request rate from institution's browsers =  $15/\text{sec}$
- delay from Internet router side to origin server and back to router = 2 sec



## consequences

- utilization on LAN = 15%
  - utilization on access link = 15%
  - total delay = Internet delay + access delay + LAN delay
- = 2 sec + minutes + milliseconds

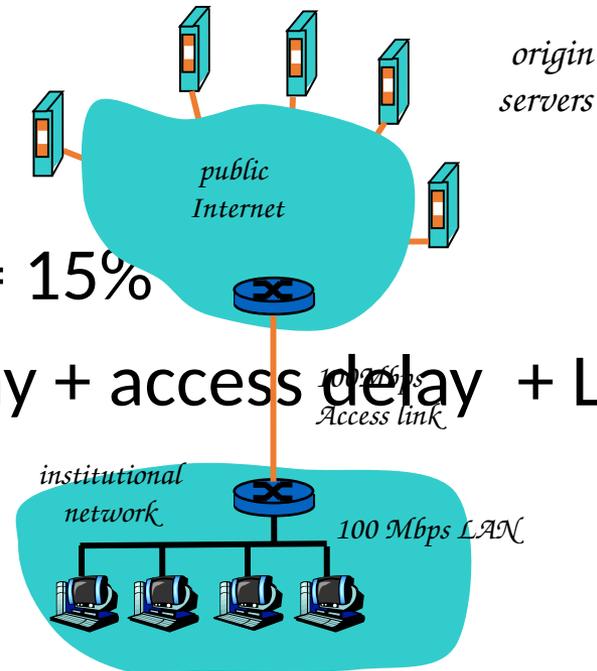
# Caching Example (cont)

possible solution

- increase bandwidth of access link to 100 Mbps

consequence

- utilization on LAN = 15%
- utilization on access link = 15%
- total delay = Internet delay + access delay + LAN delay  
= 2 sec + msec + msec  
often a costly upgrade



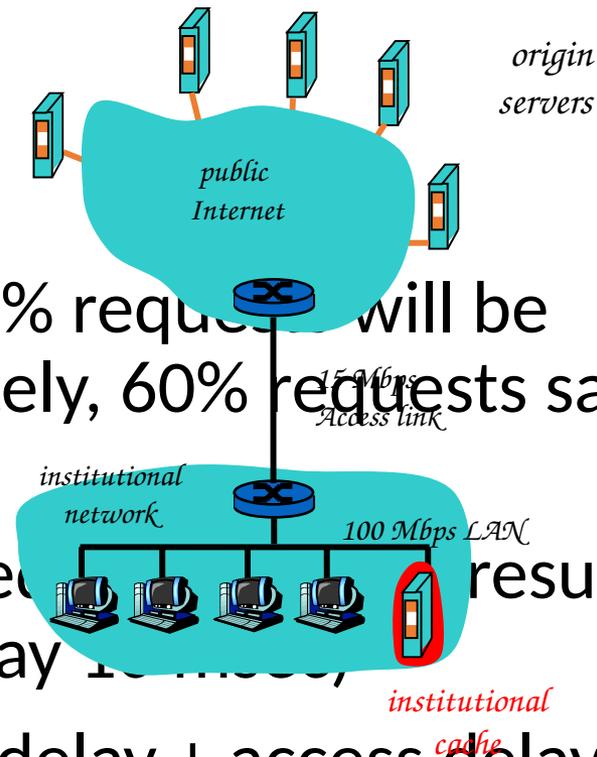
# Caching Example (cont)

possible solution:

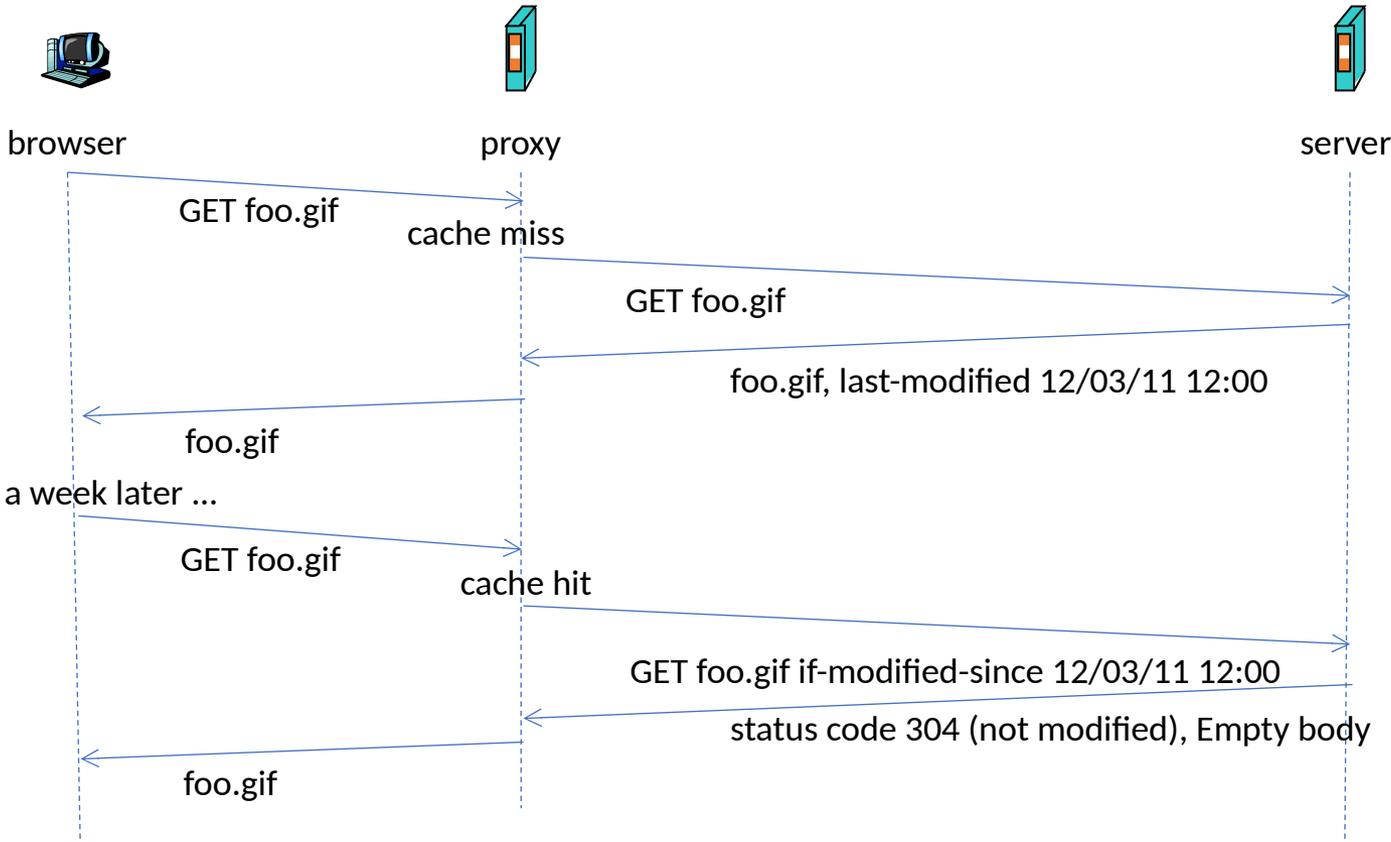
- install cache

consequence

- suppose hit rate is 0.4 (40% requests will be satisfied almost immediately, 60% requests satisfied by origin server)
- utilization of access link reduced, resulting in negligible delays (say 10 milliseconds)
- avg total delay = Internet delay + access delay + LAN delay =  $0.6 * 2.01 \text{ secs} * + 0.4 * 10 \text{ millisecs} < 1.3 \text{ secs}$



# Conditional GET



# Other Uses

- Allow multiple users to get a resource which access is limited to the proxy.
- Track and log web accesses.
- Deny access to a list of web sites.