

# Web Server Design

## Lecture 1 – Introduction to HTTP

Old Dominion University

Department of Computer Science

CS 431/531 Spring 2026

**David Calano** <[dcalo005@odu.edu](mailto:dcalo005@odu.edu)>

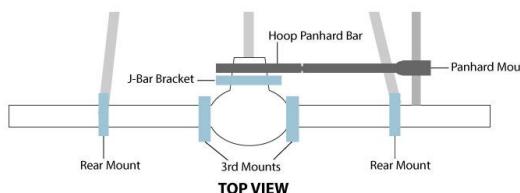
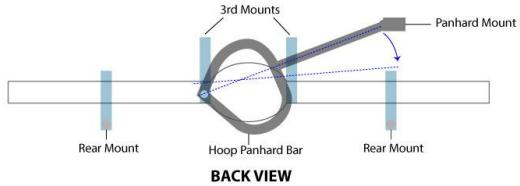
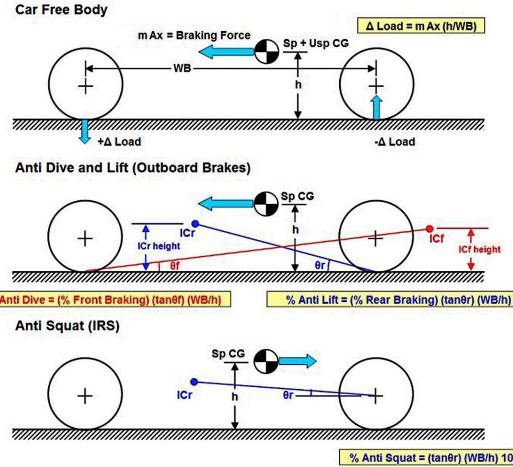
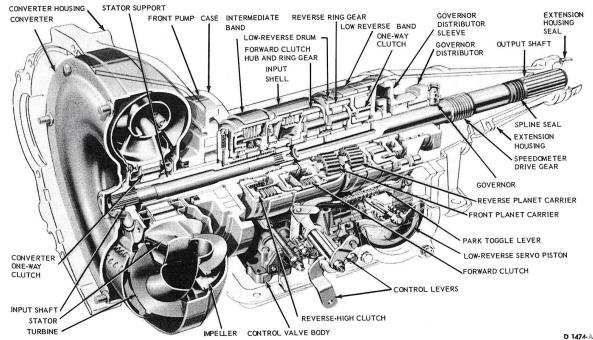
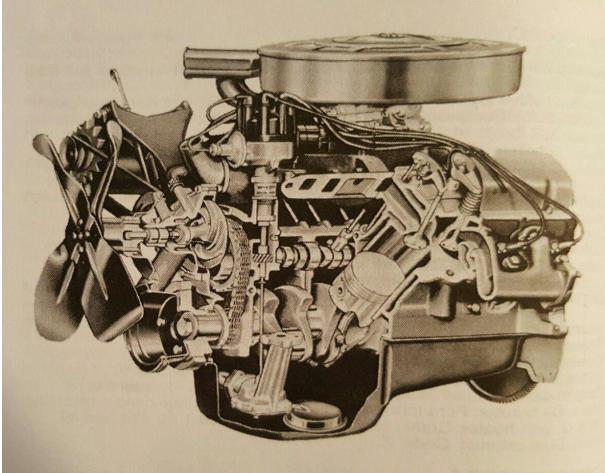
Original slides by Michael L. Nelson



Want to do this?

[https://www.youtube.com/watch?v=RJl\\_WfU5rE](https://www.youtube.com/watch?v=RJl_WfU5rE)

# It will be better/safer if you know this...



# Want to do this?

## Twitter Developer Documentation

Docs / REST APIs / Reference Documentation / GET search/tweets

### Products & Services

[Best practices](#)

[API overview](#)

[Twitter for Websites](#)

[Twitter Kit](#)

[Cards](#)

[OAuth](#)

[REST APIs](#)

[API Rate Limits](#)

[Rate Limits: Chart](#)

[The Search API](#)

## GET search/tweets

Returns a collection of relevant [Tweets](#) matching a specified query.

Please note that Twitter's search service and, by extension, the Search API is not meant to be an exhaustive source of Tweets. Not all Tweets will be indexed or made available via the search interface.

In API v1.1, the response format of the Search API has been improved to return [Tweet objects](#) more similar to the objects you'll find across the REST API and platform. However, perspectival attributes (fields that pertain to the perspective of the authenticating user) are not currently supported on this endpoint.

To learn how to use [Twitter Search](#) effectively, consult our guide to [Using the Twitter Search API](#). See [Working with Timelines](#) to learn best practices for navigating results by [since\\_id](#) and [max\\_id](#).

### Resource URL

<https://api.twitter.com/1.1/search/tweets.json>

# It will be better/safer if you know this...

```
$ telnet www.cs.odu.edu 80 | tee 6-1.out
Trying 128.82.4.2...
Connected to xenon.cs.odu.edu.
Escape character is '^].
HEAD /~mln/teaching/cs595-s06/a1-test/ HTTP/1.1
Host: www.cs.odu.edu

HTTP/1.1 200 OK
Date: Sun, 12 Feb 2006 20:58:49 GMT
Server: Apache/1.3.26 (Unix) ApacheJServ/1.1.2 PHP/4.3.4
Content-Type: text/html

HEAD /~mln/teaching/cs595-s06/a1-test/1/ HTTP/1.1
Host: www.cs.odu.edu

HTTP/1.1 200 OK
Date: Sun, 12 Feb 2006 20:58:55 GMT
Server: Apache/1.3.26 (Unix) ApacheJServ/1.1.2 PHP/4.3.4
Content-Type: text/html

HEAD /~mln/teaching/cs595-s06/a1-test/2/ HTTP/1.1
Host: www.cs.odu.edu

HTTP/1.1 200 OK
Date: Sun, 12 Feb 2006 20:59:01 GMT
Server: Apache/1.3.26 (Unix) ApacheJServ/1.1.2 PHP/4.3.4
Last-Modified: Sun, 29 Jan 2006 18:43:15 GMT
ETag: "1f4de2-790-43dd0cc3"
Accept-Ranges: bytes
Content-Length: 1936
Content-Type: text/html
X-Pad: avoid browser bug

Connection closed by foreign host.
```

# Goals

- We will write a web (HTTP) server from scratch
  - we will not use Apache, IIS, Nginx, or other existing web servers
  - the point is to *learn basic HTTP* and have a working server at the end of the class
    - your server won't be as “good” as Apache -- and that's ok...
- We will use industry standard tools/environments/systems/etc.
  - GitHub/Git
  - Docker

# I'm not teaching Web Application Development

- If you want to learn LAMP, you need to take Dr. Jian Wu's 418/518 (Web Programming) class



Instead of LAMP, you'll be learning the basis of:

REST: Representational State Transfer &

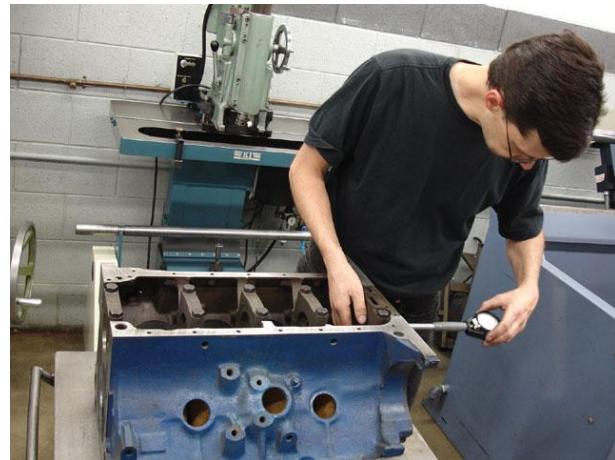
HATEOAS: Hypermedia as the Engine of Application State

# To Reiterate:



CS 418/518 – Make it Pretty

<https://www.hotrod.com/articles/fairlane-finale-finish-2016-road-tour-ford/>



CS 431/531 – Under the Hood

<https://www.hotrod.com/articles/ccrp-0808-ford-390-fe/>

# REST vs. RPC

RPC: foo.com/bigApp.jsp?verb=showThing&id=123

REST: foo.com/things/123 (w/ GET method)

RPC: foo.com/bigApp.jsp?verb=editThing&id=123

REST: foo.com/things/123 (w/ PUT method)

RPC: foo.com/bigApp.jsp?verb=newThing

REST: foo.com/things/ (w/ POST method)

Quick-n-dirty summary:

in REST, URIs are *nouns* and HTTP provides the *verbs*

this will make more sense as we go through the semester, and there's actually a lot more to REST:

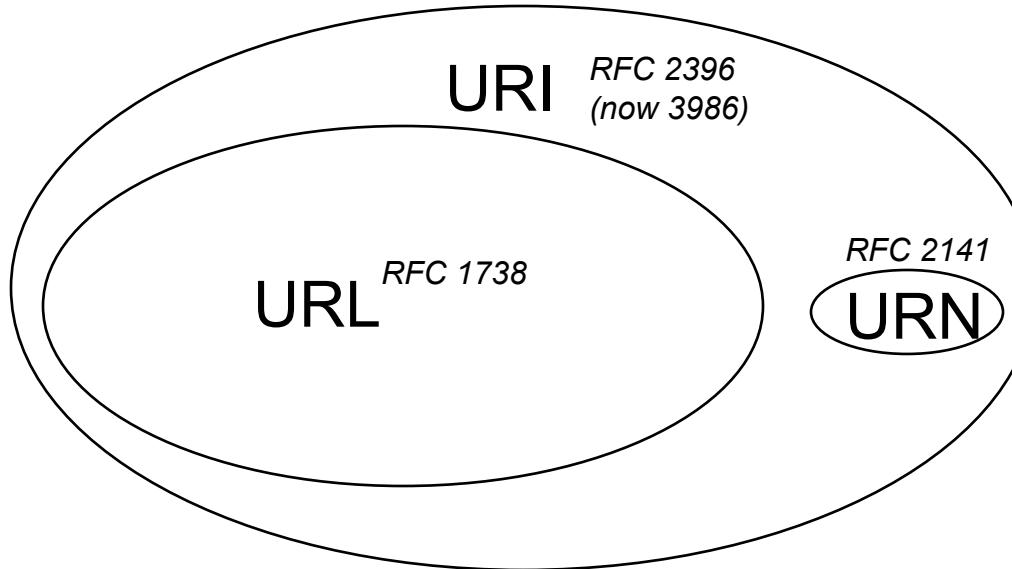
<https://www.ics.uci.edu/~fielding/pubs/dissertation/top.htm>

<https://research.google.com/pubs/archive/46310.pdf>

# Defining the Web / HTTP

- HTTP was originally defined by Request for Comments (RFCs) 1945, 2068, 2616
  - and several others for defining URLs, URIs, etc.
- Venerable RFC 2616 was replaced in 2014 with:
  - RFC7230 - [HTTP/1.1: Message Syntax and Routing](#) - low-level message parsing and connection management
  - RFC7231 - [HTTP/1.1: Semantics and Content](#) - methods, status codes and headers
  - RFC7232 - [HTTP/1.1: Conditional Requests](#) - e.g., If-Modified-Since
  - RFC7233 - [HTTP/1.1: Range Requests](#) - getting partial content
  - RFC7234 - [HTTP/1.1: Caching](#) - browser and intermediary caches
  - RFC7235 - [HTTP/1.1: Authentication](#) - a framework for HTTP authentication
  - see: [https://www.mnot.net/blog/2014/06/07/rfc2616\\_is\\_dead](https://www.mnot.net/blog/2014/06/07/rfc2616_is_dead)
- Further refactored and replaced in 2022
  - RFC9110 - HTTP Semantics - those core, versionless semantics
  - RFC9111 - HTTP Caching - split into a separate document for convenience, but also versionless
  - RFC9112 - HTTP/1.1 - everything that's specific to that textual wire protocol
  - see: <https://www.mnot.net/blog/2022/06/06/http-core>
- We also have a slightly revisionist but ultimately useful unifying document, ca. 2004:
  - The Architecture of the World Wide Web, Volume One: <http://www.w3.org/TR/webarch/>

# Uniform Resource Identifiers



URI & URL: <http://www.cs.odu.edu/>

URL: <ftp://ftp.isi.edu/pub/>

URI: info:pmid/12376099

URN: urn:uuid:6e8bc430-9c3a-11d9-9669-0800200c9a66

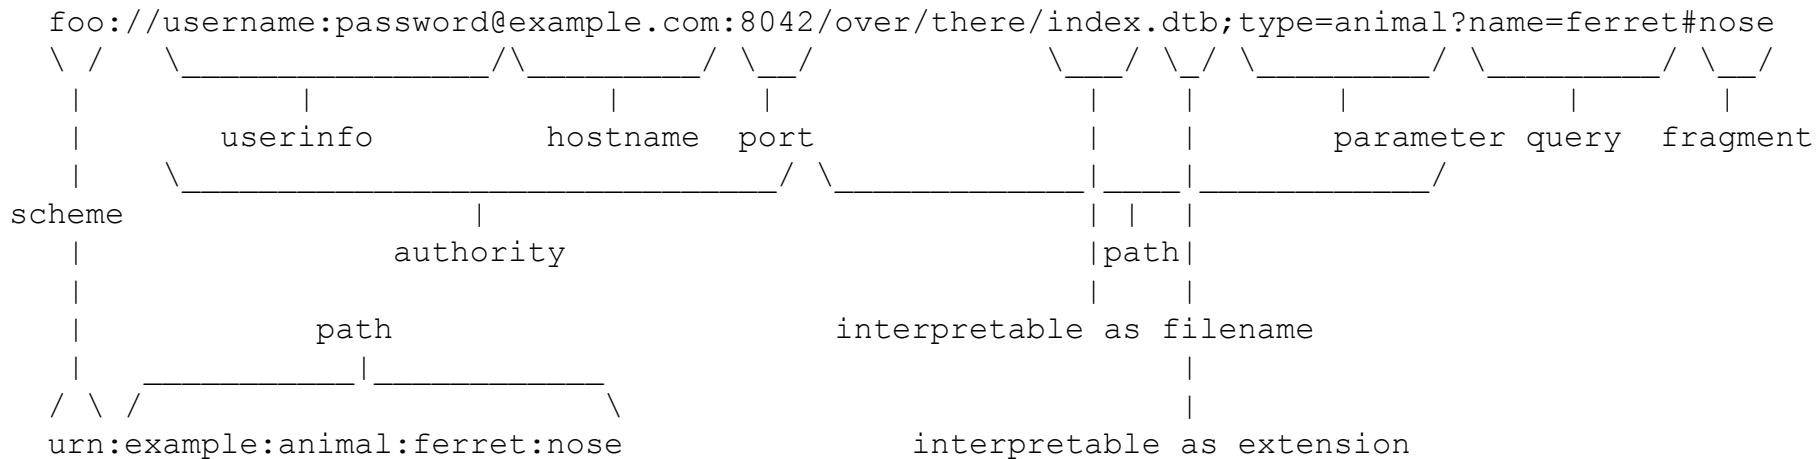
# From RFC 3986

“A URI can be further classified as a locator, a name, or both. The term "Uniform Resource Locator" (URL) refers to the subset of URIs that, in addition to identifying a resource, provide a means of locating the resource by describing its primary access mechanism (e.g., its network "location"). The term "Uniform Resource Name" (URN) has been used historically to refer to both URIs under the "urn" scheme [RFC2141], which are required to remain globally unique and persistent even when the resource ceases to exist or becomes unavailable, and to any other URI with the properties of a name.”

# URIs & URNs

- Registered URI schemes
  - <http://www.iana.org/assignments/uri-schemes>
- Registered URN namespaces
  - <http://www.iana.org/assignments/urn-namespaces>

# URI Schemes



taken from: [http://en.wikipedia.org/wiki/URI\\_scheme](http://en.wikipedia.org/wiki/URI_scheme)  
note: “scheme”, *not* “protocol”

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119.

1. **MUST** This word, or the terms "REQUIRED" or "SHALL", mean that the definition is an absolute requirement of the specification.
2. **MUST NOT** This phrase, or the phrase "SHALL NOT", mean that the definition is an absolute prohibition of the specification.
3. **SHOULD** This word, or the adjective "RECOMMENDED", mean that there may exist valid reasons in particular circumstances to ignore a particular item, but the full implications must be understood and carefully weighed before choosing a different course.
4. **SHOULD NOT** This phrase, or the phrase "NOT RECOMMENDED" mean that there may exist valid reasons in particular circumstances when the particular behavior is acceptable or even useful, but the full implications should be understood and the case carefully weighed before implementing any behavior described with this label.
5. **MAY** This word, or the adjective "OPTIONAL", mean that an item is truly optional. One vendor may choose to include the item because a particular marketplace requires it or because the vendor feels that it enhances the product while another vendor may omit the same item. An implementation which does not include a particular option MUST be prepared to interoperate with another implementation which does include the option, though perhaps with reduced functionality. In the same vein an implementation which does include a particular option MUST be prepared to interoperate with another implementation which does not include the option (except, of course, for the feature the option provides.)

# How To Read RFCs (quoting from RFC 2119)

# Important Web Architecture Concepts

## URIs

`http://www.cs.odu.edu/~mln/`

Identify

## Resources



## Representations\*

```
<html>
<head>
<title>
Home:: Michael L. Nelson, Old Dominion University
</title>
<link rel="stylesheet" type="text/css"
href="mln.css"/>
<script type="text/javascript"
src="mln.js"></script>
...
```

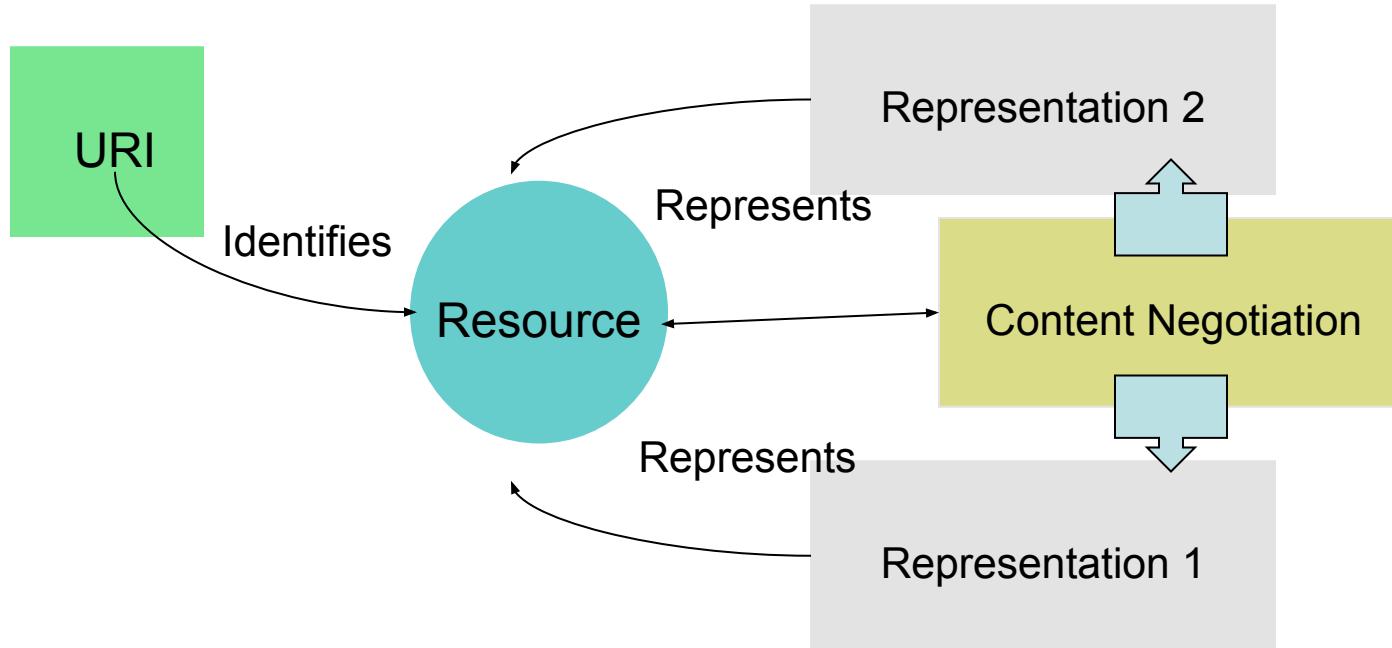
Represent

As defined by the Web Architecture

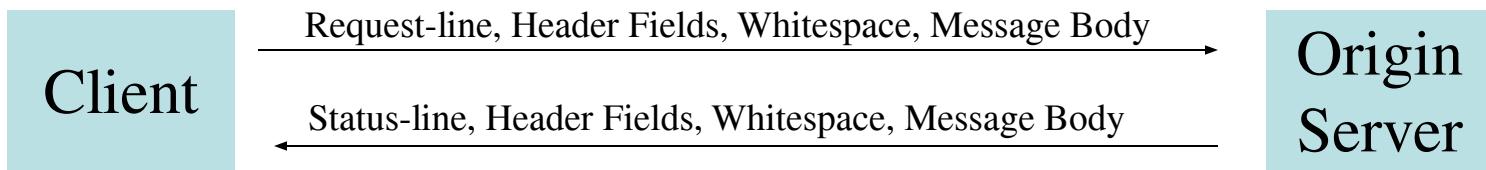
<http://www.w3.org/TR/webarch/>

\* = “message” or “message body” in RFC 7231,  
“entity”/“entity-body” in RFC-2616

# Resources can have multiple, simultaneous *representations*



# HTTP Operation



# General Template, CR/LF, and Magic Blank Lines

**Client:**

```
Method URI HTTP/1.1
Some-Request-Header-1: value1
Some-Request-Header-2: value2
...
(1st magic blank line)
```

**Server:**

```
HTTP/1.1 Code String
Some-Response-Header-1: value1
Some-Response-Header-2: value2
...
(2nd magic blank line)
message-body
```

Client's "request-line" and Server's "status-line" are the format exceptions; otherwise headers are in a flat, key-value syntax, followed by a blank line, followed by an optional message-body

# Modern Browsers (aka “user-agents”) are nice...



But they hide important details from us.  
As programmers, we care about those details.

# Talking to HTTP servers with “curl”

```
$ curl --head http://www.cs.odu.edu/~mln/
HTTP/1.1 200 OK
Date: Mon, 12 Jan 2009 15:44:19 GMT
Server: Apache/2.2.0
Last-Modified: Fri, 09 Jan 2009 17:18:37 GMT
ETag: "88849-1c71-f28dd540"
Accept-Ranges: bytes
Content-Length: 7281
Content-Type: text/html
```

```
$ curl -I http://www.google.com/
HTTP/1.1 200 OK
Cache-Control: private, max-age=0
Date: Mon, 12 Jan 2009 15:45:57 GMT
Expires: -1
Content-Type: text/html; charset=ISO-8859-1
Set-Cookie: PREF=ID=9a80d3f602b685f3:TM=1231775157:LM=1231775157:S=imGxRyNsTD0Zczm5;
expires=Wed, 12-Jan-2011 15:45:57 GMT; path=/; domain=.google.com
Server: gws
Content-Length: 0
```

# default curl returns message body, no headers...

```
$ curl https://www.cs.odu.edu/~mln/
<html>
<head>
<title>
Home::: Michael L. Nelson, Old Dominion University
</title>
<!-- CSS stuff largely stolen from Carl Lagoze's Page -->
<link rel="stylesheet" type="text/css" href="mln.css"/>

<meta property="fb:admins"
content="michael.lloyd.nelson"/>
<meta property="og:title" content="Michael L. Nelson"/>
[lots of html removed]
```

# curl –i shows response headers + message body:

```
$ curl -i https://www.cs.odu.edu/~mln/
HTTP/1.1 200 OK
Server: nginx
Date: Wed, 29 Aug 2018 02:34:15 GMT
Content-Type: text/html
Transfer-Encoding: chunked
Connection: keep-alive
Vary: Accept-Encoding
Front-End-Https: on

<html>
<head>
<title>
Home::: Michael L. Nelson, Old Dominion
University
</title>
[deletia]
```

```
$ curl -v http://ws-dl.blogspot.com/2018/08/2018-08-25-four-ws-dl-classes-offered.html
* Adding handle: conn: 0x7fa59b004000
* Adding handle: send: 0
* Adding handle: recv: 0
* Curl_addHandleToPipeline: length: 1
* - Conn 0 (0x7fa59b004000) send_pipe: 1, recv_pipe: 0
* About to connect() to ws-dl.blogspot.com port 80 (#0)
*   Trying 172.217.5.65...
* Connected to ws-dl.blogspot.com (172.217.5.65) port 80 (#0)
> GET /2018/08/2018-08-25-four-ws-dl-classes-offered.html HTTP/1.1
> User-Agent: curl/7.30.0
> Host: ws-dl.blogspot.com
> Accept: */*
>
< HTTP/1.1 200 OK
< Content-Type: text/html; charset=UTF-8
< Expires: Wed, 29 Aug 2018 01:28:50 GMT
< Date: Wed, 29 Aug 2018 01:28:50 GMT
< Cache-Control: private, max-age=0
< Last-Modified: Tue, 28 Aug 2018 23:33:07 GMT
< X-Content-Type-Options: nosniff
< X-XSS-Protection: 1; mode=block
* Server GSE is not blacklisted
< Server: GSE
< Accept-Ranges: none
< Vary: Accept-Encoding
< Transfer-Encoding: chunked
<
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Strict//EN" "http://www.w3.org/TR/xhtml1/DTD/xhtml1-strict.dtd">
<html dir='ltr' xmlns='http://www.w3.org/1999/xhtml' xmlns:b='http://www.google.com/2005/gml/b'
xmlns:data='http://www.google.com/2005/gml/data' xmlns:expr='http://www.google.com/2005/gml/expr'>
<head>
[much deletia]
```

# curl -v shows more

**\* = TCP activity**

**> = what the user-agent (i.e., curl) sent**

**< = what the server sent**

```
$ curl -IL https://t.co/LSUT8iaB4o
HTTP/1.1 301 Moved Permanently
cache-control: private,max-age=300
content-length: 0
date: Wed, 29 Aug 2018 16:42:22 GMT
expires: Wed, 29 Aug 2018 16:47:22 GMT
location: http://bit.ly/2wrdv27
server: tsa_b
set-cookie: muc=a5182687-df3f-415d-b0f7-c9bfc80ecf80; Expires=Fri, 28 Aug 2020 16:42:22 GMT; Domain=t.co
strict-transport-security: max-age=0
vary: Origin
x-connection-hash: 281d50e3381e8db894369efba3ef525d
x-response-time: 9

HTTP/1.1 301 Moved Permanently
Server: nginx
Date: Wed, 29 Aug 2018 16:42:22 GMT
Content-Type: text/html; charset=utf-8
Content-Length: 170
Connection: keep-alive
Cache-Control: private, max-age=90
Location: http://www.dailypress.com/sports/dp-spt-acc-football-injury-reports-0829-story.html

HTTP/1.1 200 OK
Content-Length: 427
Content-Type: text/html; charset=UTF-8
Httpd-Identifier: web-0a454bb1c0b26e0bc62060302d005aa2
Server: Apache-Coyote/1.1
x-Instance-Name: i15prod-c1fdbad-15-98.1
X-UA-Compatible: IE=Edge
Cache-Control: public, max-age=177
Date: Wed, 29 Aug 2018 16:42:23 GMT
Connection: keep-alive
```

curl has many,  
many flags...

# wget crawls and saves sites

```
$ wget https://t.co/LSUT8iaB4o
--2018-08-29 15:04:04--  https://t.co/LSUT8iaB4o
Resolving t.co (t.co)... 199.16.156.75, 199.16.156.11
Connecting to t.co (t.co)|199.16.156.75|:443... connected.
HTTP request sent, awaiting response... 301 Moved Permanently
Location: http://bit.ly/2wrdv27 [following]
--2018-08-29 15:04:04--  http://bit.ly/2wrdv27
Resolving bit.ly (bit.ly)... 67.199.248.11, 67.199.248.10
Connecting to bit.ly (bit.ly)|67.199.248.11|:80... connected.
HTTP request sent, awaiting response... 301 Moved Permanently
Location: http://www.dailypress.com/sports/dp-spt-acc-football-injury-reports-0829-story.html [following]
--2018-08-29 15:04:04--  http://www.dailypress.com/sports/dp-spt-acc-football-injury-reports-0829-story.html
Resolving www.dailypress.com (www.dailypress.com)... 184.84.171.209
Connecting to www.dailypress.com (www.dailypress.com)|184.84.171.209|:80... connected.
HTTP request sent, awaiting response... 200 OK
Length: unspecified [text/html]
Saving to: 'LSUT8iaB4o'

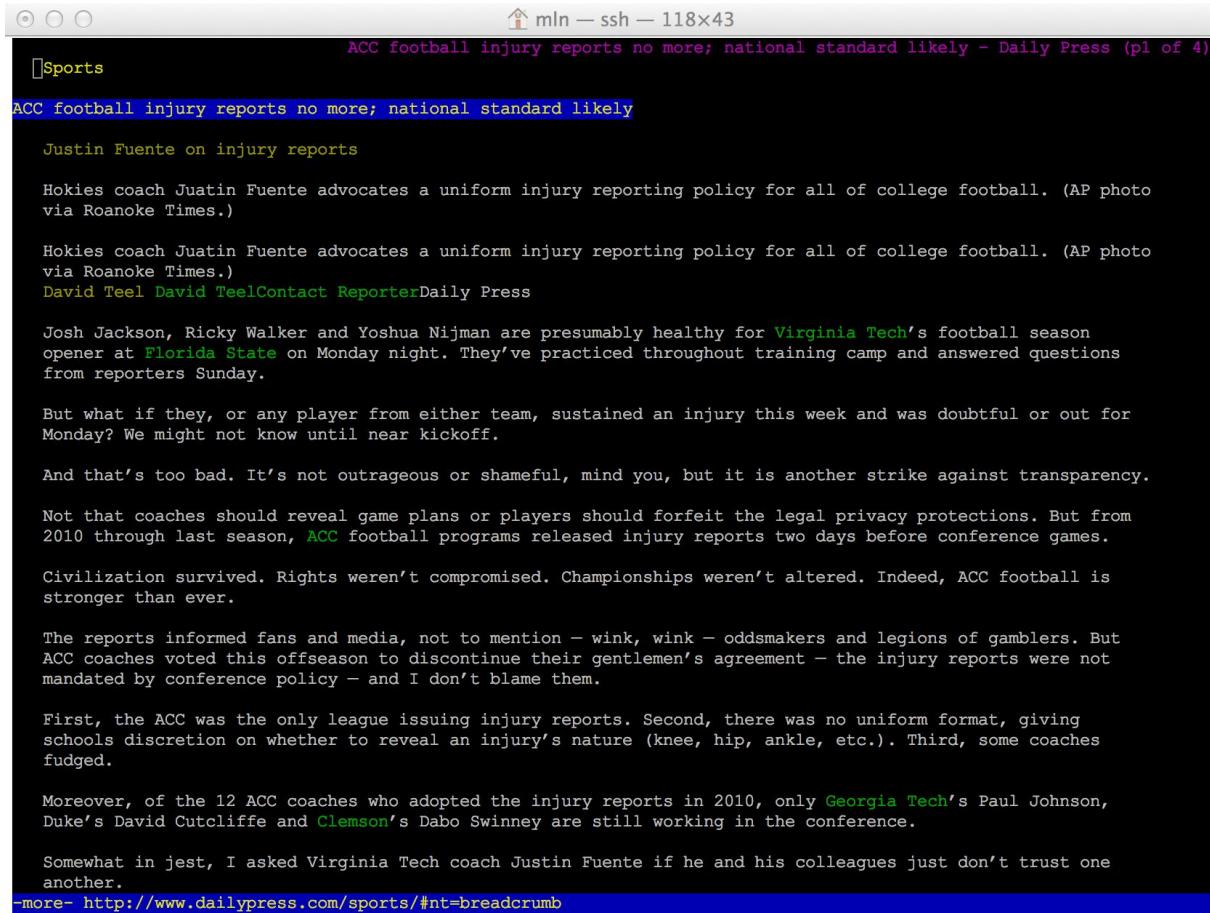
[=>] 149,323  --.-K/s  in 0.04s

2018-08-29 15:04:04 (3.77 MB/s) - 'LSUT8iaB4o' saved [149323]
```

```
$ head LSUT8iaB4o
<!DOCTYPE html>
<html lang="en-US">
<head itemid="" itemscope itemtype="http://schema.org/SpeakableSpecification">
<title>ACC football injury reports no more; national standard likely - Daily Press</title>
<meta charset="utf-8" />

<link rel="dns-prefetch" href="//www.trbimg.com"
      /><link rel="preconnect" href="//www.trbimg.com"
      /><link rel="dns-prefetch" href="//static.chartbeat.com"
      /><link rel="dns-prefetch" href="//m.trb.com"
```

```
$ lynx https://t.co/LSUT8iaB4o
```



# curl/wget/lynx are awesome

but they are still user-agents,  
*and the nature of user-agents  
is to hide details.*

we'll frequently use  
“telnet” or “openssl”  
*to expose details*

# GET

```
$ telnet www.cs.odu.edu 80
Trying 128.82.4.2...
Connected to xenon.cs.odu.edu.
Escape character is '^]'.
GET /~mln/index.html HTTP/1.1
Connection: close
Host: www.cs.odu.edu
```

Request (ends w/ CRLF)

```
HTTP/1.1 200 OK
Date: Mon, 09 Jan 2006 17:07:04 GMT
Server: Apache/1.3.26 (Unix) ApacheJServ/1.1.2 PHP/4.3.4
Last-Modified: Sun, 29 May 2005 02:46:53 GMT
ETag: "1c52-14ed-42992d1d"
Accept-Ranges: bytes
Content-Length: 5357
Connection: close
Content-Type: text/html
```

```
<html>
<head>
<title>Home Page for Michael L. Nelson</title>
<style type="text/css">
<!--
[lots of html deleted]
Connection closed by foreign host.
```

Response

# HEAD

```
$ telnet www.cs.odu.edu 80
Trying 128.82.4.2...
Connected to xenon.cs.odu.edu.
Escape character is '^]'.
HEAD /~mln/index.html HTTP/1.1
Connection: close
Host: www.cs.odu.edu

HTTP/1.1 200 OK
Date: Mon, 09 Jan 2006 17:14:39 GMT
Server: Apache/1.3.26 (Unix) ApacheJServ/1.1.2
PHP/4.3.4
Last-Modified: Sun, 29 May 2005 02:46:53 GMT
ETag: "1c52-14ed-42992d1d"
Accept-Ranges: bytes
Content-Length: 5357
Connection: close
Content-Type: text/html

Connection closed by foreign host.
```

# OPTIONS

## (many methods)

```
$ telnet www.cs.odu.edu 80
```

```
Trying 128.82.4.2...
```

```
Connected to xenon.cs.odu.edu.
```

```
Escape character is '^]'.
```

```
OPTIONS /~mln/index.html HTTP/1.1
```

```
Connection: close
```

```
Host: www.cs.odu.edu
```

```
HTTP/1.1 200 OK
```

```
Date: Mon, 09 Jan 2006 17:16:46 GMT
```

```
Server: Apache/1.3.26 (Unix) ApacheJServ/1.1.2 PHP/4.3.4
```

```
Content-Length: 0
```

```
Allow: GET, HEAD, POST, PUT, DELETE, CONNECT, OPTIONS, PATCH, PROPFIND,  
PROPPATCH, MKCOL, COPY, MOVE, LOCK, UNLOCK, TRACE
```

```
Connection: close
```

```
Connection closed by foreign host.
```

# OPTIONS

## (fewer methods)

```
$ telnet www.cs.odu.edu 80
Trying 128.82.4.2...
Connected to xenon.cs.odu.edu.
Escape character is '^]'.
OPTIONS /~mln/index.html HTTP/1.1
Connection: close
Host: www.cs.odu.edu

HTTP/1.1 200 OK
Date: Tue, 10 Jan 2012 17:26:44 GMT
Server: Apache/2.2.17 (Unix) PHP/5.3.5 mod_ssl/2.2.17 OpenSSL/0.9.8q
Allow: GET,HEAD,POST,OPTIONS
Content-Length: 0
Connection: close
Content-Type: text/html

Connection closed by foreign host.
```

# HTTPS is supplanting HTTP

this is mostly a good thing\*

but it does mean we can't use telnet for “https” sites

\* <https://www.theverge.com/2018/2/8/16991254/chrome-not-secure-marked-http-encryption-ssl>

# bye bye “telnet to port 80”

```
$ telnet www.cs.odu.edu 80
Trying 128.82.4.2...
Connected to xenon.cs.odu.edu.
Escape character is '^]'.
HEAD /~mln/ HTTP/1.1
Host: www.cs.odu.edu
Connection: close
```

```
HTTP/1.1 301 Moved Permanently
Server: nginx
Date: Wed, 29 Aug 2018 03:45:36 GMT
Content-Type: text/html
Connection: close
Location:
https://www.cs.odu.edu/~mln/
Connection closed by foreign host.
```

```
$ telnet www.cs.odu.edu 443
Trying 128.82.4.2...
Connected to xenon.cs.odu.edu.
Escape character is '^]'.
HEAD /~mln/ HTTP/1.1
Host: www.cs.odu.edu
Connection: close
```

```
HTTP/1.1 400 Bad Request
Server: nginx
Date: Wed, 29 Aug 2018 03:45:57
GMT
Content-Type: text/html
Connection: close
Connection closed by foreign host.
```

# hello “openssl to port 443”

```
$ openssl s_client -connect www.cs.odu.edu:443
CONNECTED(00000003)
[much, much SSL deletia]
SSL handshake has read 6270 bytes and written 328 bytes
---
New, TLSv1/SSLv3, Cipher is DHE-RSA-AES128-SHA
Server public key is 2048 bit
Secure Renegotiation IS supported
Compression: NONE
Expansion: NONE
SSL-Session:
Protocol : TLSv1
Cipher   : DHE-RSA-AES128-SHA
Session-ID: E19FD48AA69A296996B958877C48C28391ED217761F1E2023C7471ACB89B2694
Session-ID-ctx:
Master-Key: 0A9A3DC0C66F99FF85A480ADEC42A7EB74EEC1D391D9AF4A026CF27C16A19480C42A75B6CD283BFE68ADAB32D07D7242
Key-Arg   : None
Start Time: 1535514923
Timeout   : 300 (sec)
Verify return code: 0 (ok)
---
HEAD /~mln/ HTTP/1.1
Host: www.cs.odu.edu
Connection: close

HTTP/1.1 200 OK
Server: nginx
Date: Wed, 29 Aug 2018 03:55:35 GMT
Content-Type: text/html
Connection: close
Vary: Accept-Encoding
Front-End-Https: on

closed
```

# HTTP semantics don't change

```
$ openssl s_client -connect www.cs.odu.edu:443
[all SSL portions deleted]
OPTIONS /~mln/ HTTP/1.1
Host: www.cs.odu.edu
Connection: close

HTTP/1.1 200 OK
Server: nginx
Date: Wed, 29 Aug 2018 04:02:05 GMT
Content-Type: text/html
Content-Length: 0
Connection: close
Allow: POST,OPTIONS,GET,HEAD
Front-End-Https: on

closed
```

# Response Codes

from section 6 of RFC 7231

- 1xx: Informational - The request was received, continuing process
- 2xx: Success - The action was successfully received, understood, and accepted
- 3xx: Redirection - Further action must be taken in order to complete the request
- 4xx: Client Error - The request contains bad syntax or cannot be fulfilled
- 5xx: Server Error - The server failed to fulfill an apparently valid request

not “error” codes!!!

# 501 - Method Not Implemented

```
$ telnet www.cs.odu.edu 80
Trying 128.82.4.2...
Connected to xenon.cs.odu.edu.
Escape character is '^].
NOTAREALMETHOD /index.html HTTP/1.1
Connection: close
Host: www.cs.odu.edu

HTTP/1.1 501 Method Not Implemented
Date: Mon, 09 Jan 2006 17:22:40 GMT
Server: Apache/1.3.26 (Unix) ApacheJServ/1.1.2 PHP/4.3.4
Allow: GET, HEAD, POST, PUT, DELETE, CONNECT, OPTIONS, PATCH, PROPFIND, PROPPATCH, MKCOL, COPY,
MOVE, LOCK, UNLOCK, TRACE
Connection: close
Transfer-Encoding: chunked
Content-Type: text/html; charset=iso-8859-1

15f
<!DOCTYPE HTML PUBLIC "-//IETF//DTD HTML 2.0//EN">
<HTML><HEAD>
<TITLE>501 Method Not Implemented</TITLE>
</HEAD><BODY>
<H1>Method Not Implemented</H1>
NOTAREALMETHOD to /index.html not supported.<P>
Invalid method in request NOTAREALMETHOD /index.html HTTP/1.1<P>
<HR>
<ADDRESS>Apache/1.3.26 Server at www.cs.odu.edu Port 80</ADDRESS>
</BODY></HTML>

0

Connection closed by foreign host.
```

# 301 - Moved Permanently

```
$ telnet www.cs.odu.edu 80
```

```
Trying 128.82.4.2...
```

```
Connected to xenon.cs.odu.edu.
```

```
Escape character is '^]'.
```

```
GET /~mln HTTP/1.1
```

```
Connection: close
```

```
Host: www.cs.odu.edu
```

```
Connection closed by foreign host.
```

```
HTTP/1.1 301 Moved Permanently
```

```
Date: Mon, 09 Jan 2006 19:32:24 GMT
```

```
Server: Apache/1.3.26 (Unix) ApacheJServ/1.1.2 PHP/4.3.4
```

```
Location: http://www.cs.odu.edu/~mln/
```

```
Connection: close
```

```
Transfer-Encoding: chunked
```

```
Content-Type: text/html; charset=iso-8859-1
```

```
12e
```

```
<!DOCTYPE HTML PUBLIC "-//IETF//DTD HTML 2.0//EN">
```

```
<HTML><HEAD>
```

```
<TITLE>301 Moved Permanently</TITLE>
```

```
</HEAD><BODY>
```

```
<H1>Moved Permanently</H1>
```

```
The document has moved <A HREF="http://www.cs.odu.edu/~mln/">here</A>.<P>
```

```
<HR>
```

```
<ADDRESS>Apache/1.3.26 Server at www.cs.odu.edu Port 80</ADDRESS>
```

```
</BODY></HTML>
```

# 301- Moved Permanently

```
$ telnet bit.ly 80
Trying 69.58.188.39...
Connected to bit.ly.
Escape character is '^]'.
HEAD http://bit.ly/s2FPFa HTTP/1.1
Host: bit.ly
Connection: close
```

**HTTP/1.1 301 Moved**

```
Server: nginx
Date: Tue, 10 Jan 2012 17:34:29 GMT
Content-Type: text/html; charset=utf-8
Connection: close
Set-Cookie: __bit=4f0c76a5-002b9-048b1-331cf10a;domain=.bit.ly;
    expires=Sun Jul 8 17:34:29 2012;path=/; HttpOnly
Cache-control: private; max-age=90
Location: http://bit.ly/bundles/phonedude/e
MIME-Version: 1.0
Content-Length: 125
```

the response code is REQUIRED;  
phrase is RECOMMENDED

# 302 - Found

```
$ telnet doi.acm.org 80
Trying 64.238.147.57...
Connected to doi.acm.org.
Escape character is '^]'.
HEAD http://doi.acm.org/10.1145/1998076.1998100 HTTP/1.1
Host: doi.acm.org
Connection: close
```

## HTTP/1.1 302 Found

```
Date: Tue, 10 Jan 2012 17:53:36 GMT
Server: Apache/2.2.3 (Red Hat)
Location:
http://dl.acm.org/citation.cfm?doid=1998076.1998100
Connection: close
Content-Type: text/html; charset=iso-8859-1
```

# 303 - See Other

```
$ telnet dx.doi.org 80
Trying 38.100.138.149...
Connected to dx.doi.org.
Escape character is '^]'.
HEAD http://dx.doi.org/10.1007/978-3-642-24469-8_16 HTTP/1.1
Host: dx.doi.org
Connection: close
```

**HTTP/1.1 303 See Other**

```
Server: Apache-Coyote/1.1
Location:
http://www.springerlink.com/index/10.1007/978-3-642-24469-8_16
Expires: Wed, 11 Jan 2012 12:04:29 GMT
Content-Type: text/html; charset=utf-8
Content-Length: 210
Date: Tue, 10 Jan 2012 17:56:41 GMT
Connection: close
```

# 404 - Not Found

```
$ telnet www.cs.odu.edu 80
Trying 128.82.4.2...
Connected to xenon.cs.odu.edu.
Escape character is '^]'.
HEAD /lasdkfjalsdkfjldaskfj HTTP/1.1
Host: www.cs.odu.edu
Connection: close
```

**HTTP/1.1 404 Not Found**

```
Date: Tue, 10 Jan 2012 17:39:15 GMT
Server: Apache/2.2.17 (Unix) PHP/5.3.5 mod_ssl/2.2.17
 OpenSSL/0.9.8q
Connection: close
Content-Type: text/html; charset=iso-8859-1
```

Connection closed by foreign host.

# 401 - Unauthorized

```
$ telnet www4.cs.odu.edu 80
Trying 128.82.5.93...
Connected to www4.cs.odu.edu.
Escape character is '^]'.
HEAD http://www4.cs.odu.edu/Conference/index.aspx HTTP/1.1
Host: www4.cs.odu.edu
Connection: close
```

```
HTTP/1.1 401 Unauthorized
Content-Length: 1656
Content-Type: text/html
Server: Microsoft-IIS/6.0
WWW-Authenticate: Basic realm="www4.cs.odu.edu"
MicrosoftOfficeWebServer: 5.0_Pub
X-Powered-By: ASP.NET
Date: Tue, 10 Jan 2012 17:43:57 GMT
Connection: close
```

# 400 - Bad Request

```
$ telnet www.cs.odu.edu 80
Trying 128.82.4.2...
Connected to xenon.cs.odu.edu.
Escape character is '^]'.
HEAD http://www.cs.odu.edu/~mln/ HTTP/1.1
Connection: close
```

**HTTP/1.1 400 Bad Request**

```
Date: Tue, 10 Jan 2012 18:24:17 GMT
Server: Apache/2.2.17 (Unix) PHP/5.3.5 mod_ssl/2.2.17
 OpenSSL/0.9.8q
Connection: close
Content-Type: text/html; charset=iso-8859-1
```

# 505 - HTTP Version Not Supported

```
$ telnet www.cs.odu.edu 80
Trying 128.82.4.2...
Connected to xenon.cs.odu.edu.
Escape character is '^]'.
HEAD / HTTP/9.9
Host: www.cs.odu.edu
Connection: close
```

our servers will be more picky!

```
HTTP/1.1 200 OK
Date: Tue, 10 Jan 2012 17:40:05 GMT
Server: Apache/2.2.17 (Unix) PHP/5.3.5 mod_ssl/2.2.17
 OpenSSL/0.9.8q
Accept-Ranges: bytes
Connection: close
Content-Type: text/html

Connection closed by foreign host.
```

# 505 - HTTP Version Not Supported

```
% telnet www.w3c.org 80
Trying 128.30.52.45...
Connected to dolph.w3.org.
Escape character is '^]'.
```

```
HEAD / HTTP/9.9
```

```
Host: www.w3c.org
Connection: close
```

```
HTTP/1.0 403 Forbidden
Cache-Control: no-cache
Connection: close
Content-Type: text/html
```

```
<html><body><h1>403 Forbidden</h1>
Request forbidden by administrative rules.
</body></html>
```

a curious response...  
505 not defined in HTTP 1.0!

Code	Reason-Phrase	Defined in...
100	Continue	Section 6.2.1
101	Switching Protocols	Section 6.2.2
200	OK	Section 6.3.1
201	Created	Section 6.3.2
202	Accepted	Section 6.3.3
203	Non-Authoritative Information	Section 6.3.4
204	No Content	Section 6.3.5
205	Reset Content	Section 6.3.6
206	Partial Content	Section 4.1 of [RFC7233]
300	Multiple Choices	Section 6.4.1
301	Moved Permanently	Section 6.4.2
302	Found	Section 6.4.3
303	See Other	Section 6.4.4
304	Not Modified	Section 4.1 of [RFC7232]
305	Use Proxy	Section 6.4.5
307	Temporary Redirect	Section 6.4.7
400	Bad Request	Section 6.5.1
401	Unauthorized	Section 3.1 of [RFC7235]
402	Payment Required	Section 6.5.2
403	Forbidden	Section 6.5.3
404	Not Found	Section 6.5.4
405	Method Not Allowed	Section 6.5.5
406	Not Acceptable	Section 6.5.6
407	Proxy Authentication Required	Section 3.2 of [RFC7235]
408	Request Timeout	Section 6.5.7
409	Conflict	Section 6.5.8
410	Gone	Section 6.5.9
411	Length Required	Section 6.5.10
412	Precondition Failed	Section 4.2 of [RFC7232]
413	Payload Too Large	Section 6.5.11
414	URI Too Long	Section 6.5.12
415	Unsupported Media Type	Section 6.5.13
416	Range Not Satisfiable	Section 4.4 of [RFC7233]
417	Expectation Failed	Section 6.5.14
426	Upgrade Required	Section 6.5.15
500	Internal Server Error	Section 6.6.1
501	Not Implemented	Section 6.6.2
502	Bad Gateway	Section 6.6.3
503	Service Unavailable	Section 6.6.4
504	Gateway Timeout	Section 6.6.5
505	HTTP Version Not Supported	Section 6.6.6

Many,  
Many  
Response  
Codes...

## 7.1.1.1. Date/Time Formats

...

An example of the preferred format is

Sun, 06 Nov 1994 08:49:37 GMT ; IMF-fixdate

IMF-fixdate – formerly known as RFC 1123 date/times

Examples of the two obsolete formats are

Sunday, 06-Nov-94 08:49:37 GMT ; obsolete RFC 850 format  
Sun Nov 6 08:49:37 1994 ; ANSI C's asctime() format

A recipient that parses a timestamp value in an HTTP header field MUST accept all three HTTP-date formats. When a sender generates a header field that contains one or more timestamps defined as HTTP-date, the sender MUST generate those timestamps in the IMF-fixdate format.

An HTTP-date value represents time as an instance of Coordinated Universal Time (UTC). The first two formats indicate UTC by the three-letter abbreviation for Greenwich Mean Time, "GMT", a predecessor of the UTC name; values in the asctime format are assumed to be in UTC. A sender that generates HTTP-date values from a local clock ought to use NTP ([RFC5905]) or some similar protocol to synchronize its clock to UTC.

# Things to Think About for Your Server

- Claim HTTP/1.1
  - even though we'll not fully satisfy all requirements
- Configuration files
  - should not have to recompile or edit source code for trivial changes
- MIME types
  - most servers use a separate file (specified in your config file!) to map file extensions to MIME types
- Logging
  - real http servers log their events
    - we'll use “common log format”
  - you'll need logging for debugging
    - consider concurrent logs with varying verbosity

# More Things To Think About...

- A resource is more than just a file in the file system
  - content negotiation is in your future
  - sometimes we'll give respond with only a “slice” of a file
  - What does it mean to GET a directory?
  - eventually we'll execute scripts

# In the future, some methods will allow a client to send an entity body to the server...

Client:	Method	URI	HTTP/1.1	
	Some-Request-Header-1:	value1		
	Some-Request-Header-2:	value2		
	...			
			(1st magic blank line)	
	message-body			
Server:	HTTP	1.1	Code	String
	Some-Response-Header-1:	value1		
	Some-Response-Header-2:	value2		
	...			
			(2nd magic blank line)	
	message-body			

# Revisiting What You Will Learn

- Fundamental knowledge about how http works
  - your future career is likely to involve web programming
- Working with others, explaining your results to colleagues
  - in real life, tasks are rarely performed in isolation
- How to read & interpret technical specifications and translate them into code
  - in real life, interesting problems are ambiguous & messy
- Using GitHub/Git, Docker, AWS, and other modern tools
- The importance of good, extensible design early in a software project
  - in real life, writing code from scratch is an uncommon luxury