Lecture 17

408/508 *Computational Techniques for Linguists*

Today's Topics

- Another term project idea
- Creating a webserver (`httpd`)
  - install the Apache2 webserver
    - installed by default on macOS,
    - install using apt install on Ubuntu
  - how to start/stop service
  - DocumentRoot
  - User web files

- Homework 8 (heads-up):
  - *set up your own webserver*
  - *given out on Thursday*

xkcd

https://xkcd.com/simplewriter/

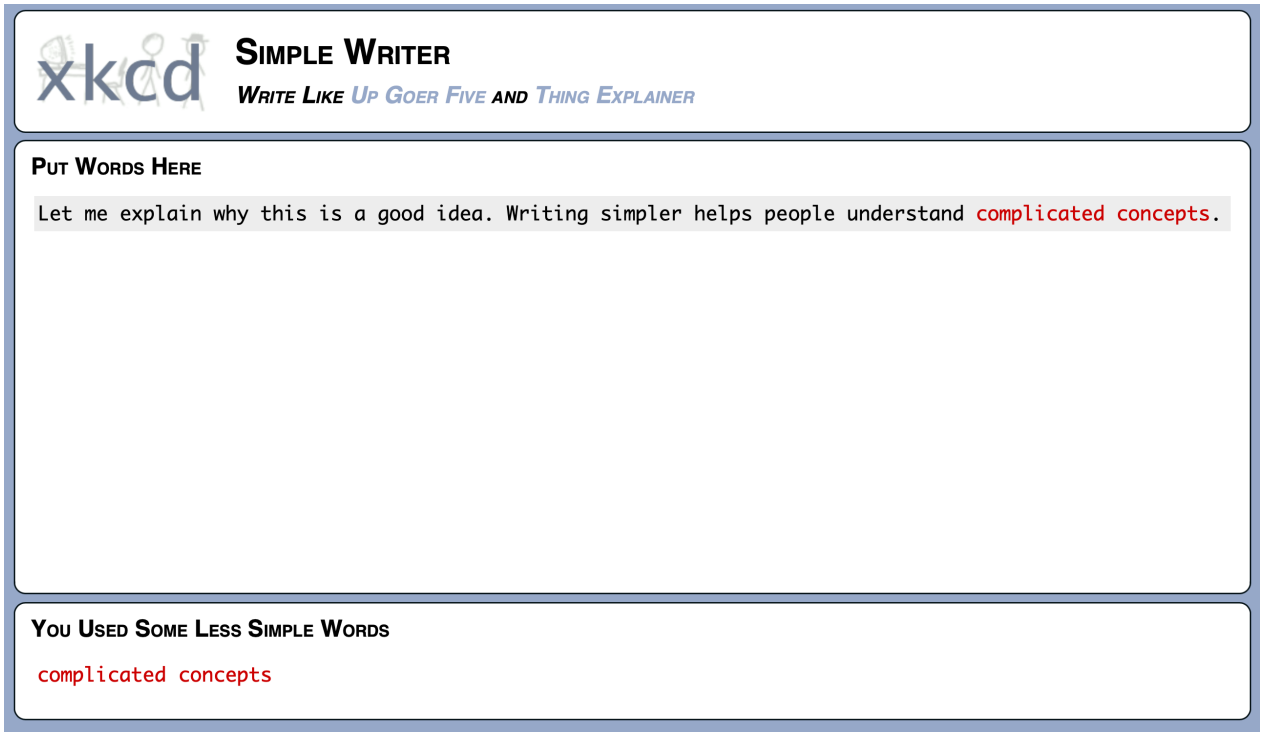

Demo

- as you type
- Javascript

The server side

- So far, all the web programming has been **client-side** only
  - i.e. the Javascript code is running *inside* the browser
- Let's build a webserver
  - the client-side will send form information to the **server-side** to be processed

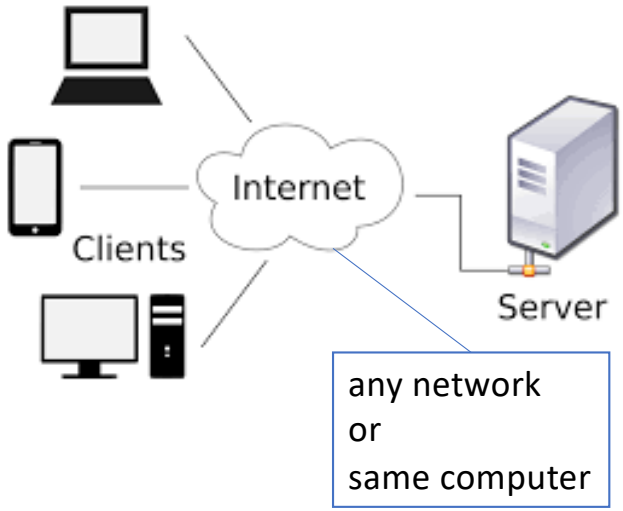

Building a Webserver

- We'll use `cgi-bin` and `bash` scripts initially …
- Apache2 is the most common webserver software
  - *https://httpd.apache.org*
  - has been around more than a quarter century
  - *unfortunately, configurations are similar but different on macOS and Ubuntu*

  (we will cover both today)

Common Gateway Interface (CGI)

- The glue between a webserver and programs that run on the computer (= server) hosting the webserver

1. Normally, a webserver sends out **static webpages** in response to (URL) requests from a client (your web browser).
2. Sometimes, we want the **request to run a program** (a script or binary) on the server that does some computation and generates some result to be displayed on the client (as a webpage).

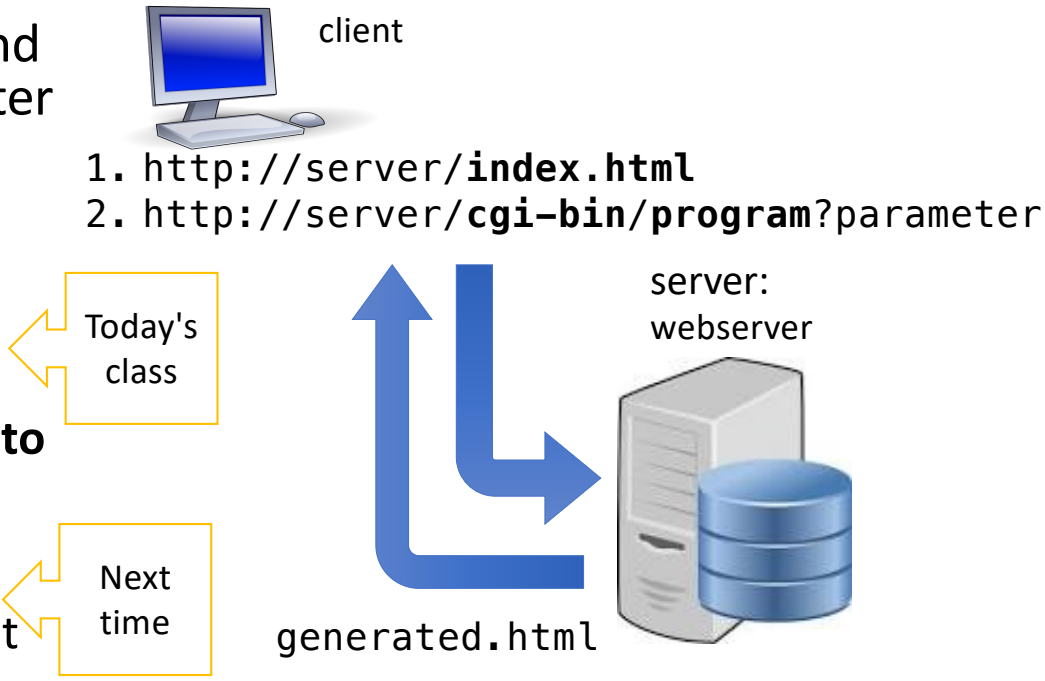

httpd

- `httpd` is the Apache HyperText Transfer Protocol (HTTP) server program.
  - It is designed to be run as a standalone **daemon** process. When used like this it will create a pool of child processes or threads to handle requests.
  - "We fancifully began to use the word **daemon** to describe background processes that worked tirelessly to perform system chores". (Prof. Corbató, MIT)
  - In general, `httpd` should not be invoked directly, but rather should be invoked via `apachctl` on Unix-based systems

Apache Webserver on macOS

Commands to be entered at a Terminal

```
~$ which httpd
/usr/sbin/httpd
```

- Apache version (macOS 14.0 *Sonoma*):
  - `~$ httpd -v`
  - `Server version: Apache/2.4.56 (Unix)`
  - `Server built:   Aug  5 2023 06:30:16`

- Apache version (OSX 10.13 *High Sierra*):
  - `~$ httpd -v`
  - `Server version: Apache/2.4.33 (Unix)`
  - `Server built:   Apr  3 2018 17:54:07`

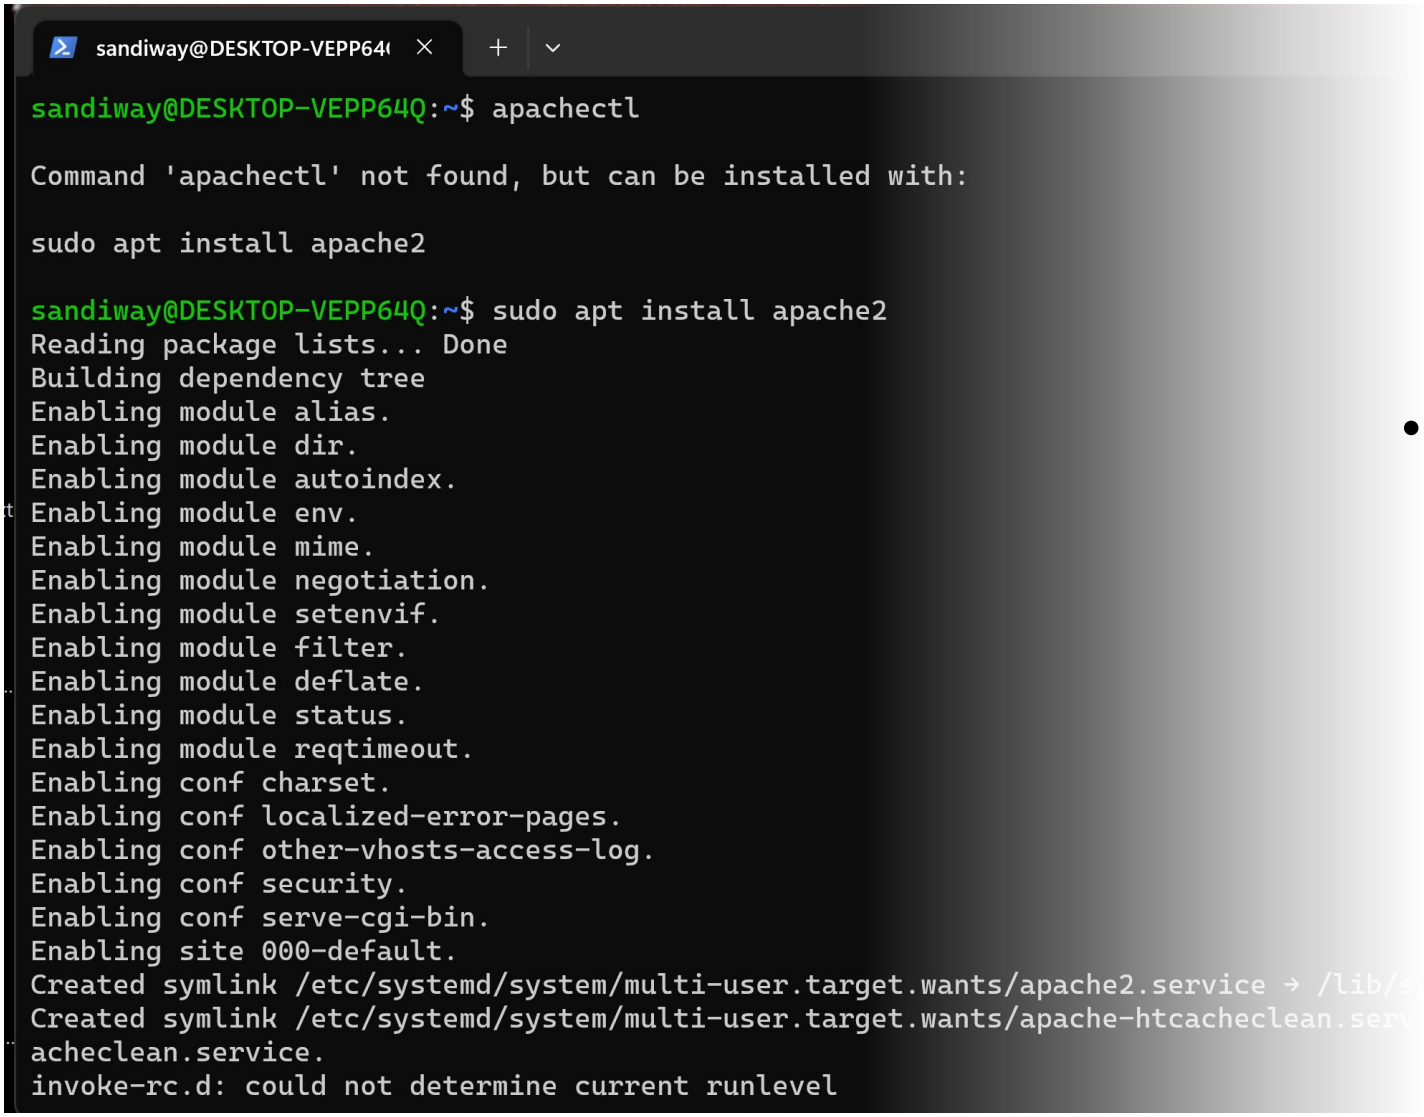

Apache Webserver on Ubuntu

- Ubuntu:
  - `sudo apt-get update`

Apache Webserver on Ubuntu

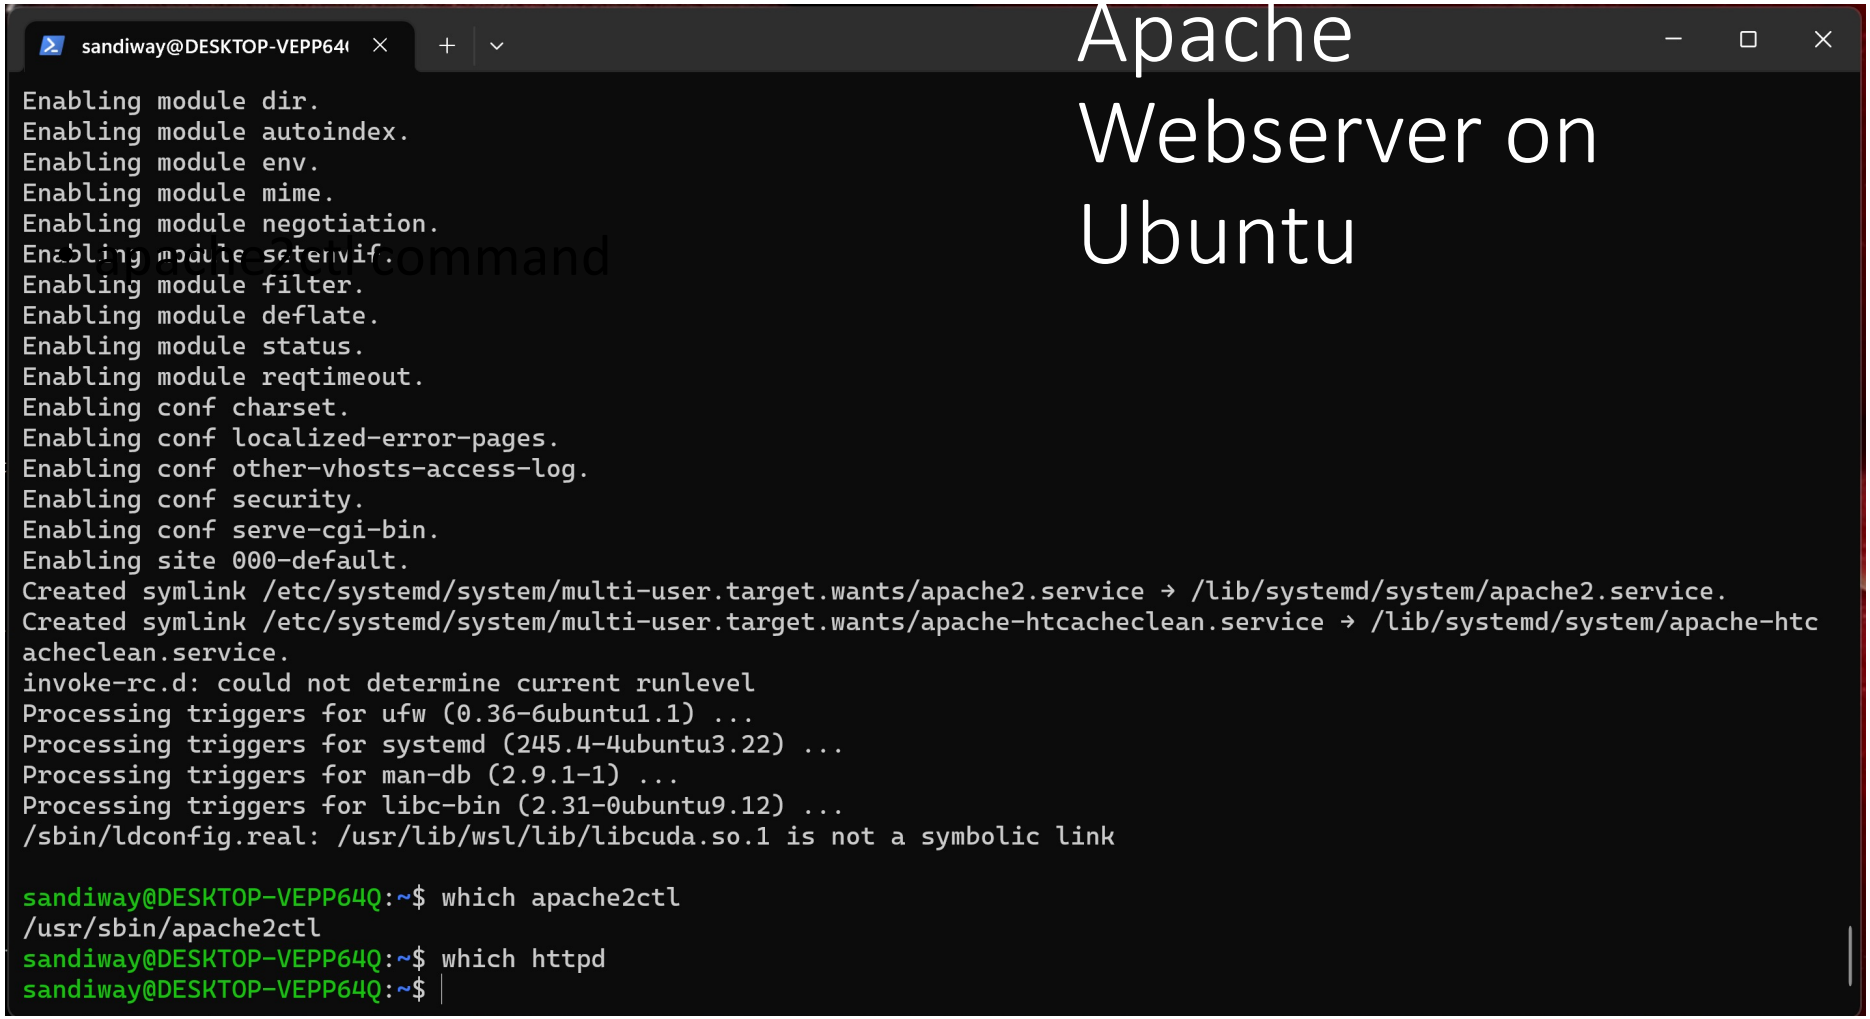

Apache Webserver on macOS

Commands to be entered at a Terminal

- Apache webserver control:
  - `~$ which apachectl`
  - `/usr/sbin/apachectl`
  - `sudo apachectl start`
  - `sudo apachectl stop`
  - `sudo apachectl –k restart` (after configuration change)
  - `apachectl configtest` (check configuration)
  - `Syntax OK`
  - `ps –ax | grep httpd` (`ps` = process status)
  - `sudo apachectl stop`
  - `ps –ax | grep httpd`

ps command

NAME

ps – process status

SYNOPSIS

ps [**–AaCcEefhjlMmrSTvwXx**] **…**

DESCRIPTION

The **ps** utility displays a header line, followed by lines containing information about all of your processes that have controlling terminals.

–a Display information about other users' processes as well as your own. This will skip any processes which do not have a controlling terminal, unless the **–x** option is also specified.

–x When displaying processes matched by other options, include processes which do not have a controlling terminal.

grep command

NAME

grep, egrep, fgrep, rgrep, bzgrep, bzegrep, bzfgrep, zgrep, zegrep, zfgrep – file pattern searcher

SYNOPSIS

grep … [pattern] [file ...]

DESCRIPTION

The **grep** utility searches any given input files, selecting lines that match one or more patterns. By default, a pattern matches an input line if the regular expression (RE) in the pattern matches the input line without its trailing newline. An empty expression matches every line. Each input line that matches at least one of the patterns is written to the standard output.

Apache Webserver on macOS

```
~$ ps -ax | grep httpd
  154 ??         0:19.63 /usr/sbin/httpd -D FOREGROUND
  517 ??         0:00.00 /usr/sbin/httpd -D FOREGROUND
17404 ttys001    0:00.00 grep httpd
~$ sudo apachectl stop
Password:
~$ ps -ax | grep httpd
17409 ttys001    0:00.00 grep httpd
(base) ~$ sudo apachectl start
(base) ~$ ps -ax | grep httpd
17414 ??         0:00.13 /usr/sbin/httpd -D FOREGROUND
17430 ??         0:00.00 /usr/sbin/httpd -D FOREGROUND
17432 ttys001    0:00.00 grep httpd
```

Apache Webserver on macOS

```
~$ sudo apachectl start
Password:
/System/Library/LaunchDaemons/org.apache.httpd.plist: service already loaded
Load failed: 37: Operation already in progress

~$ apachectl configtest
AH00558: httpd: Could not reliably determine the server's fully qualified domain name, using Sandiways-MacBook-4.local. Set the 'ServerName' directive globally to suppress this message
Syntax OK
```

Apache Webserver on macOS

- `sudo apachectl start`
- On a browser, enter: http://localhost/ this displays as default the file `/Library/WebServer/Documents/index.html`

not running…

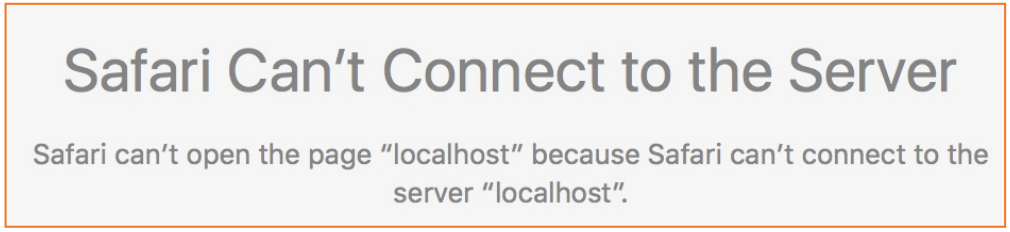

running

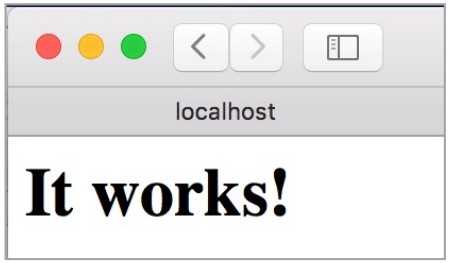

Apache Webserver on macOS

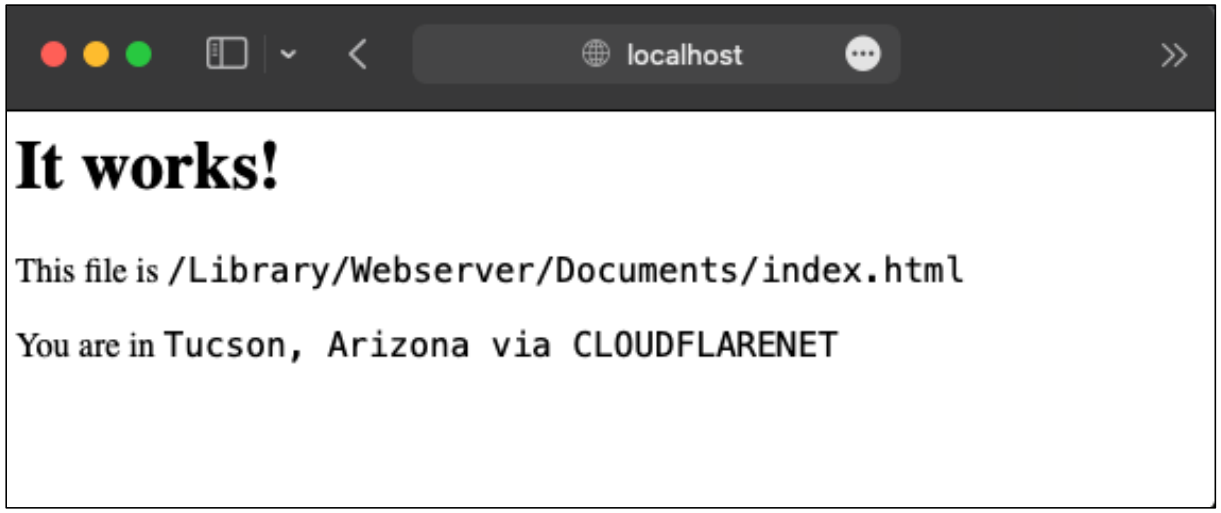

Apache Webserver on macOS

- On macOS:

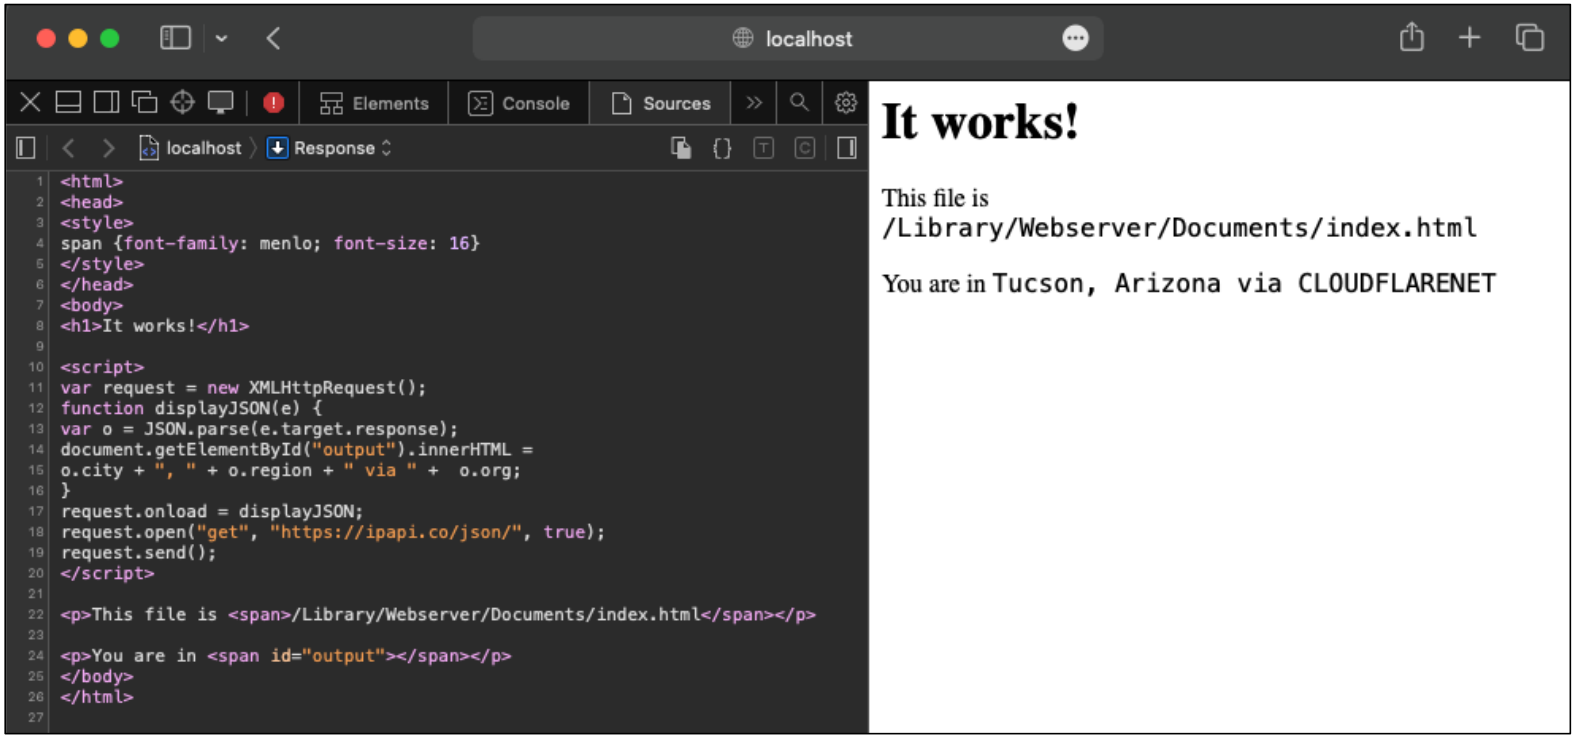

Sample Site webpage

- Normally, Javascript in the browser is sandboxed for security.
  - It has no access to your machine details, e.g. IP address or filesystem
- **So how did we do this?**
- File:
  - `sample-index.html`
  - I renamed it as `index.html` on my computer

```
<html>
<head>
<style>
span {font-family: menlo; font-size: 16}
</style>
</head>
<body>
<h1>It works!</h1>

<script>
var request = new XMLHttpRequest();
function displayJSON(e) {
var o = JSON.parse(e.target.response);
document.getElementById("output").innerHTML =
o.city + ", " + o.region + " via " +  o.org;
}
request.onload = displayJSON;
request.open("get", "https://ipapi.co/json/", true);
request.send();
</script>

<p>This file is <span>/Library/Webserver/Documents/index.html</span></p>

<p>You are in <span id="output"></span></p>
</body>
</html>
```

Lookup

`https://ipapi.co/json/`

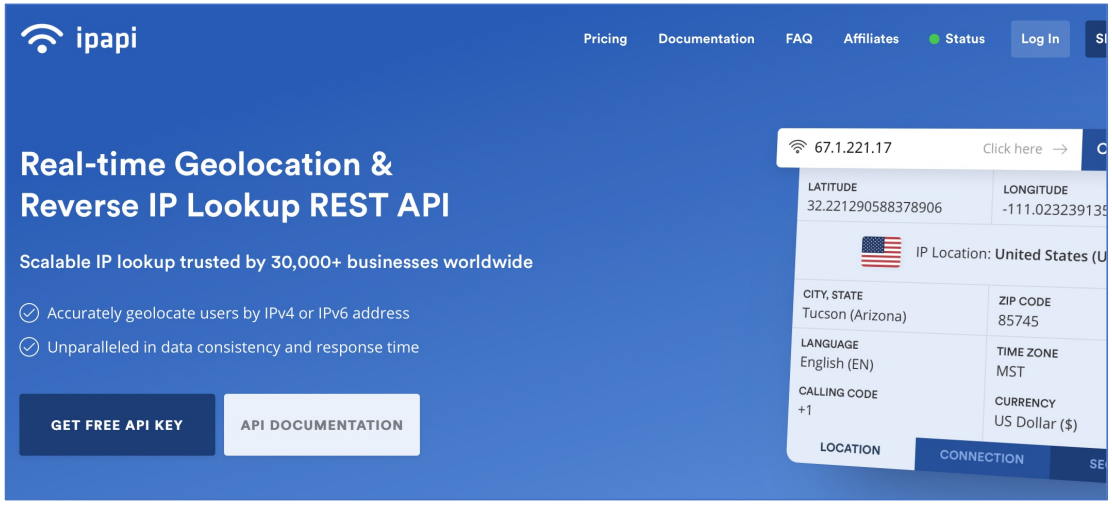

- Terminal:

```
~$ curl ifconfig.co/
2a09:bac3:6066:296::42:3d
```

Lookup

- JSON = Javascript Object Notation
- IPv6 format
  - x:x:x:x:x:x:x:x
  - each x is 4 hex digits
  - total: 128 bits
- IPv4 format (*since 1983*)
  - x.x.x.x each x is 2 hex digits
  - total: 32 bits

Apache Webserver on macOS

Default static webpage storage location:

- http://localhost/
- **`/Library/WebServer/Documents/index.html.en~orig`**
- Let's create **`index.html`** ourselves!
- **`sudo nano /Library/Webserver/Documents/index.html.en~orig`**

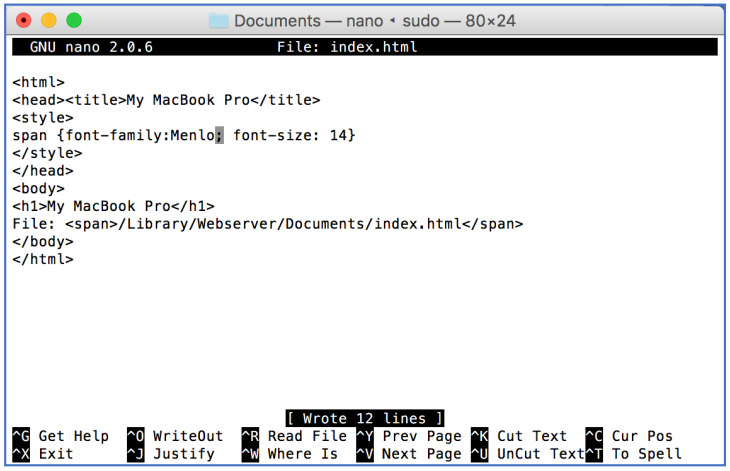

nano is a simple text editor
^ means use the Control key
save file as
`/Library/Webserver/Documents/index.html`

Apache Webserver on WSL2 (Ubuntu)

- From Windows Powershell:
  - `wsl`
- Default page (Document Root):
  - `/var/www/html/index.html`
  - owner of file is `root`; you need to use `sudo` prefix, e.g. `sudo nano`, to edit it
- Commands:
  - `sudo service apache2 start`
  - `sudo service apache2 stop`
  - `sudo service apache2 restart`

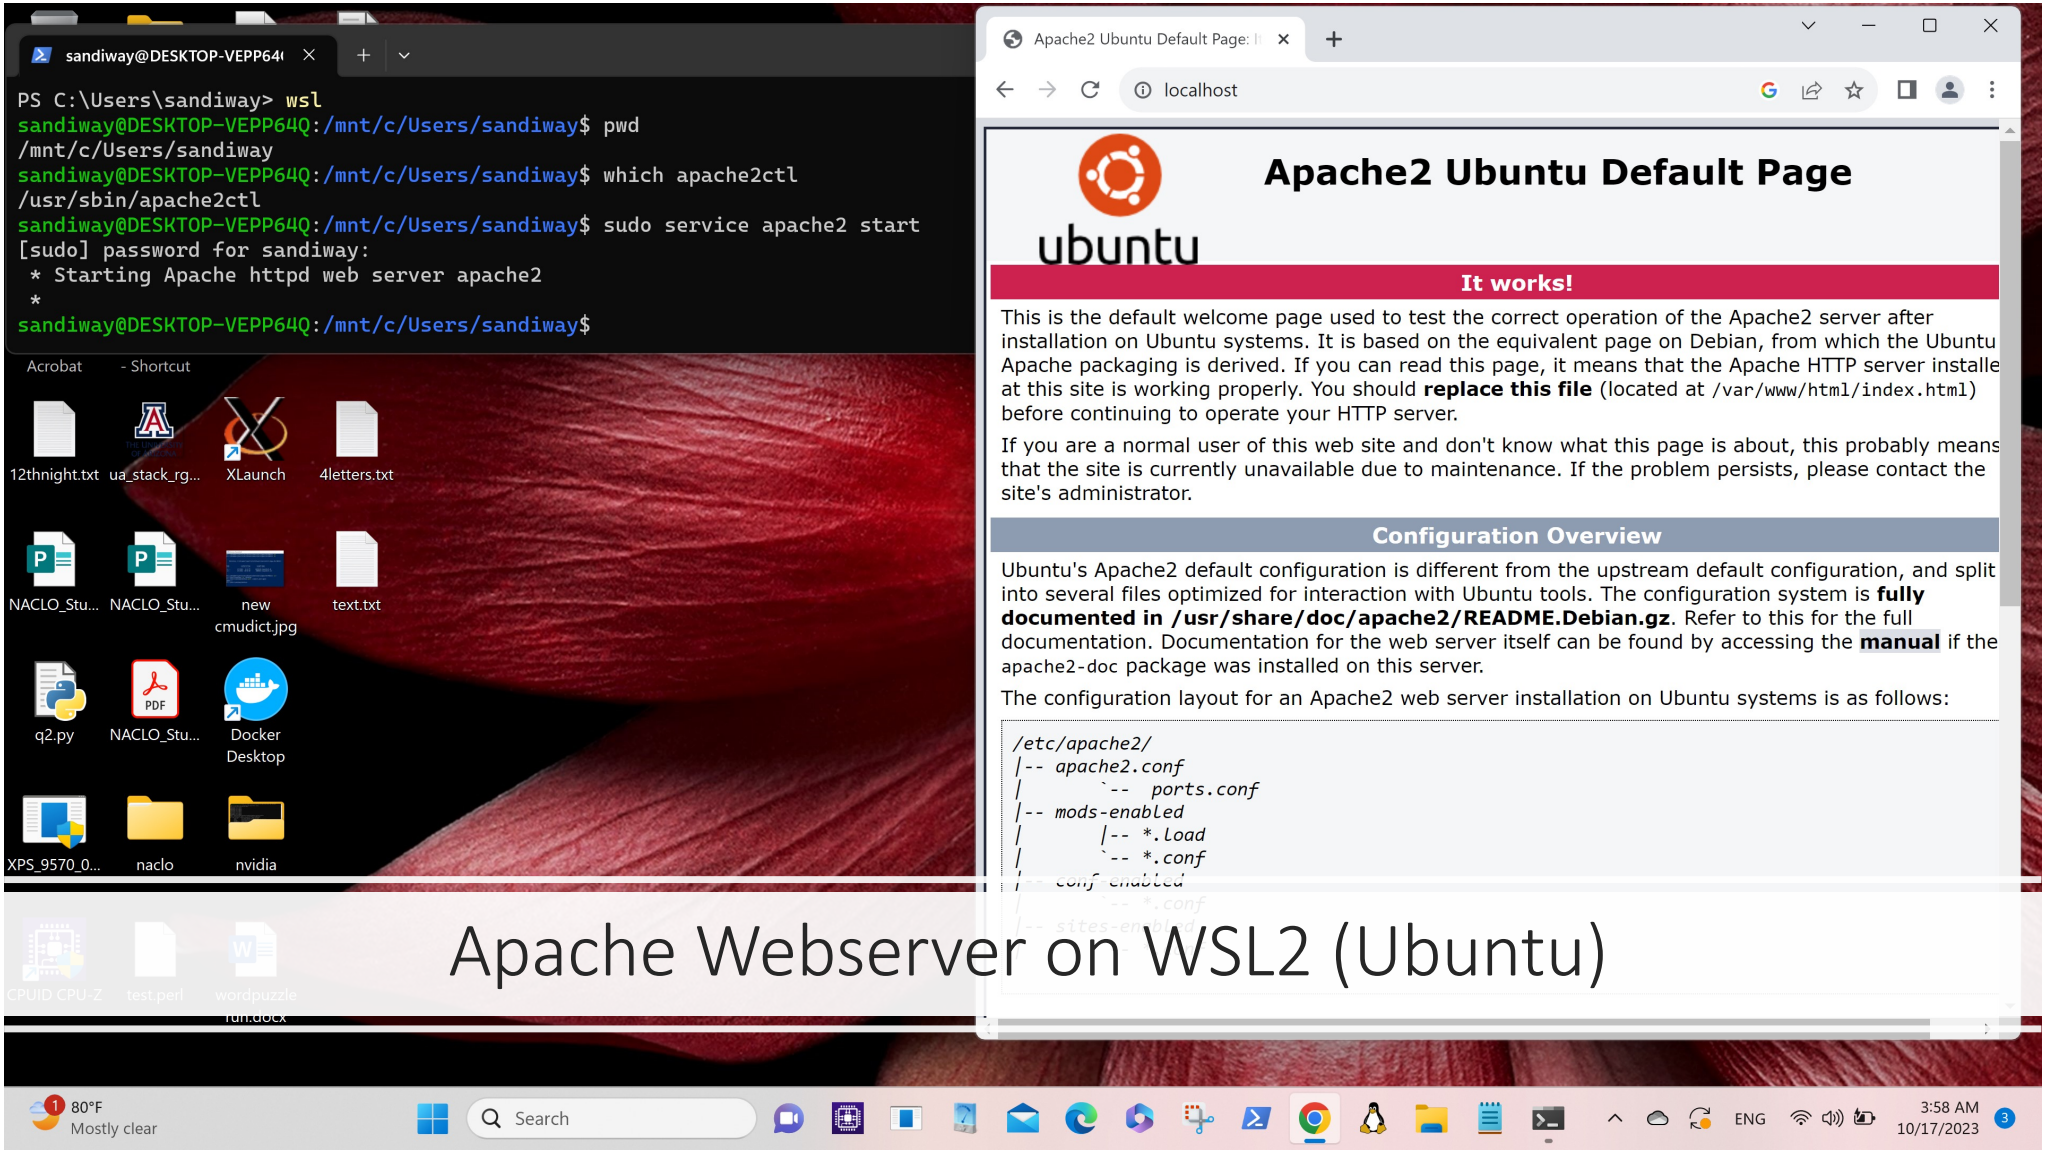

Apache Webserver on WSL2 (Ubuntu)

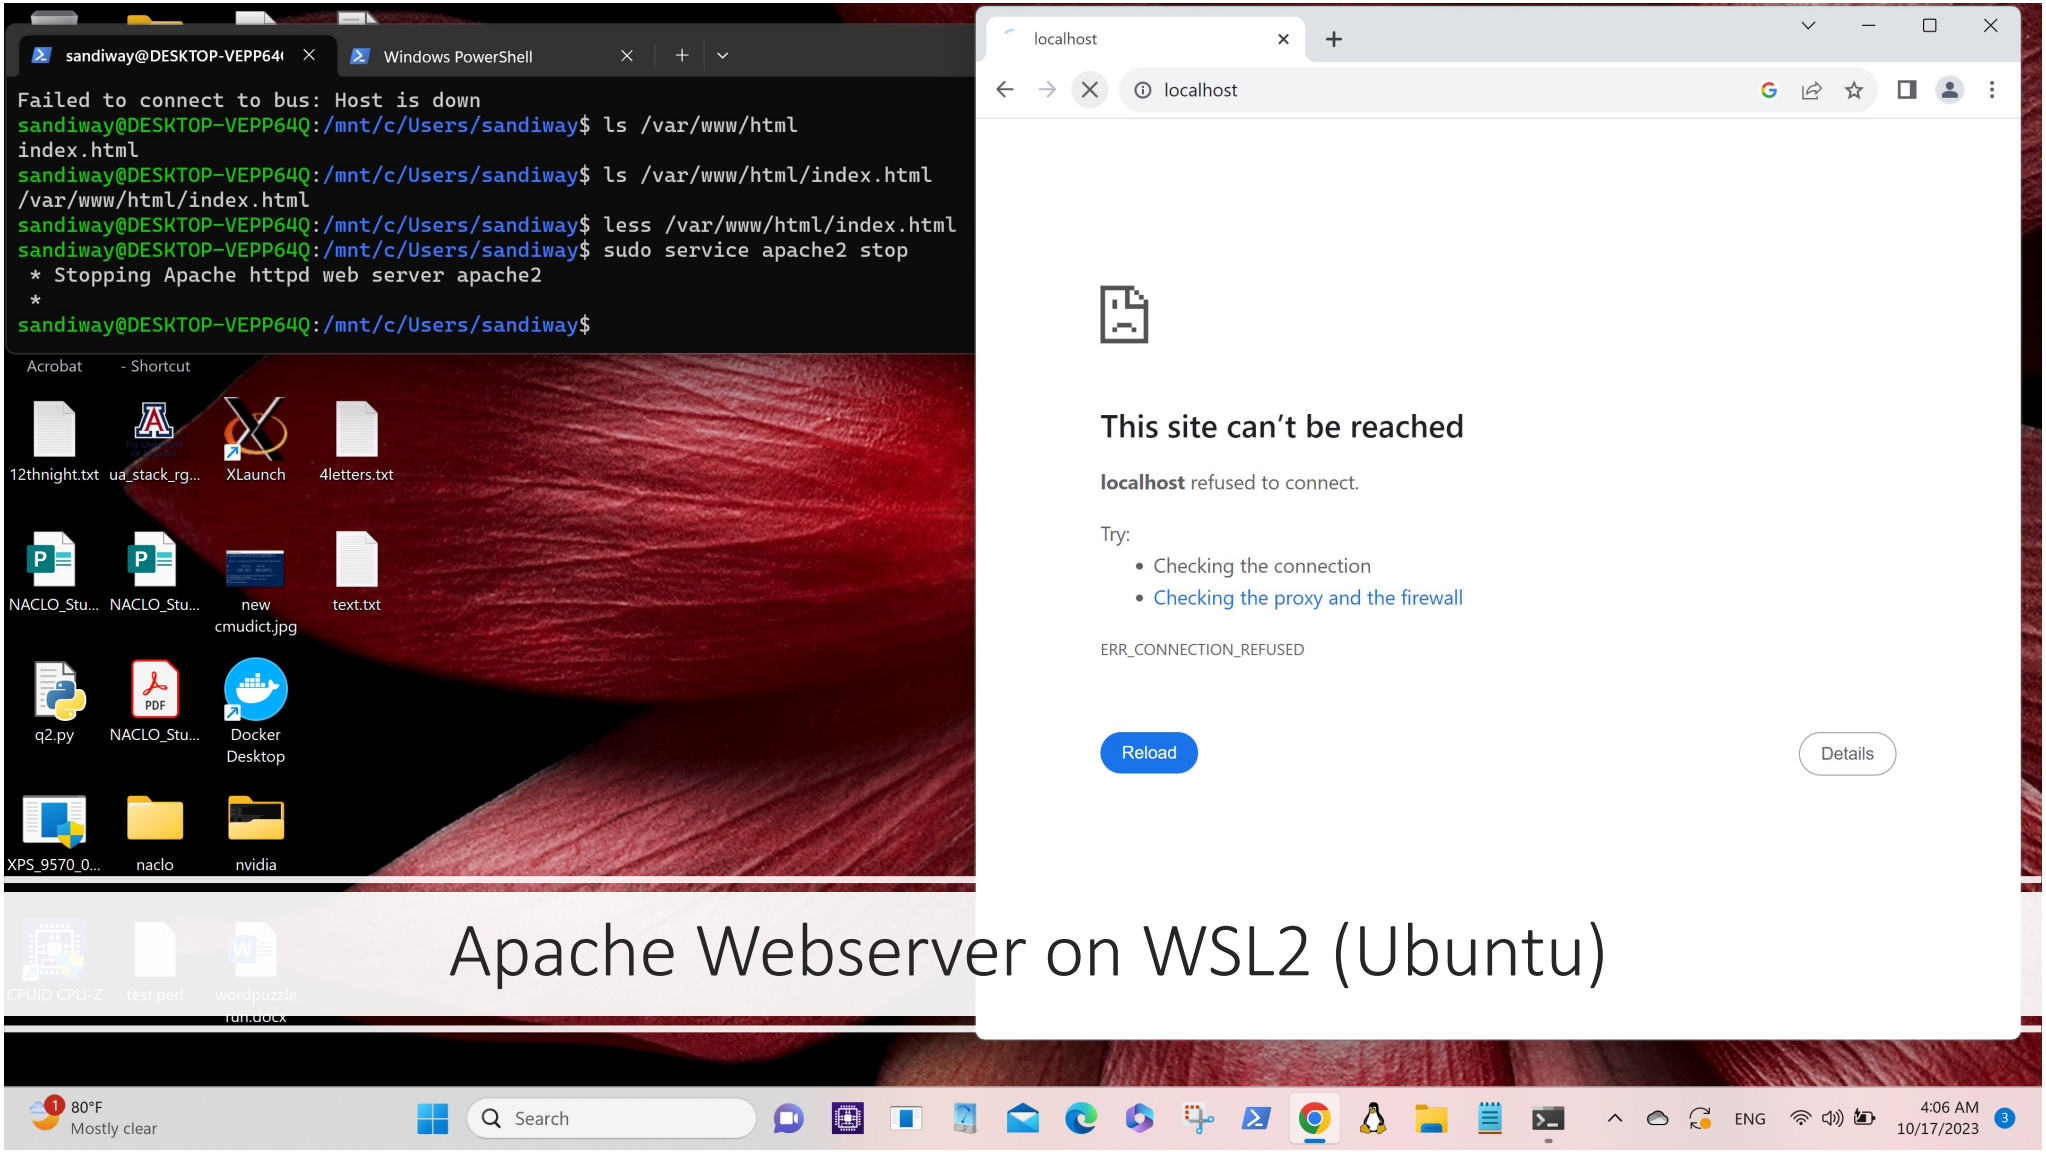

Apache Webserver on WSL2 (Ubuntu)

Apache2 on Ubuntu

- Apache2 webserver (on VirtualBox or WSL2):
  - `sudo apache2ctl start`
  - `sudo apache2ctl stop`
  - `sudo apache2ctl restart`

Note: macOS is different, uses `apachectl`

```
sandiway@sandiway-VirtualBox:~$ sudo apache2ctl start
Invoking 'systemctl start apache2'.
Use 'systemctl status apache2' for more info.
sandiway@sandiway-VirtualBox:~$ systemctl status apache2
● apache2.service - The Apache HTTP Server
   Loaded: loaded (/lib/systemd/system/apache2.service; enabled; vendor preset:
  Drop-In: /lib/systemd/system/apache2.service.d
           └─apache2-systemd.conf
   Active: active (running) since Wed 2018-10-24 20:43:13 MST; 4min 40s ago
 Main PID: 3488 (apache2)
    Tasks: 55 (limit: 4663)
   CGroup: /system.slice/apache2.service
           ├─3488 /usr/sbin/apache2 -k start
           ├─3490 /usr/sbin/apache2 -k start
           └─3491 /usr/sbin/apache2 -k start

Oct 24 20:43:13 sandiway-VirtualBox systemd[1]: Starting The Apache HTTP Server.
Oct 24 20:43:13 sandiway-VirtualBox apachectl[3477]: AH00558: apache2: Could not
Oct 24 20:43:13 sandiway-VirtualBox systemd[1]: Started The Apache HTTP Server.
sandiway@sandiway-VirtualBox:~$
```

Apache2 on Ubuntu

- Apache webserver (works on VirtualBox, doesn't work in WSL2):
  - `sudo systemctl start apache2`
  - `sudo systemctl stop apache2`
  - `sudo systemctl restart apache2`

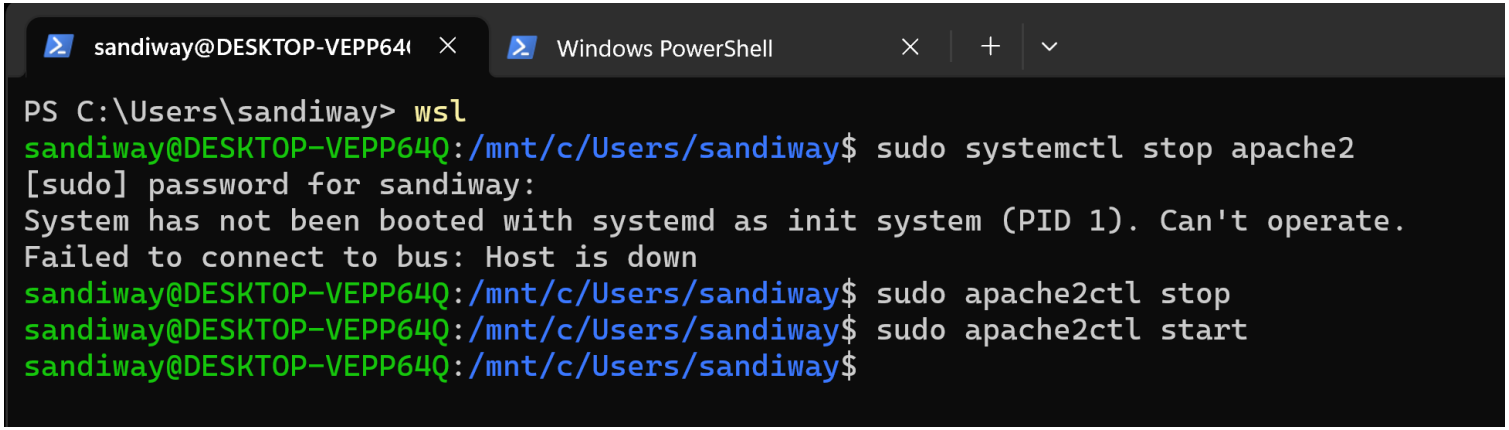

DocumentRoot

DocumentRoot Directive

Description:	Directory that forms the main document tree visible from the web
Syntax:	`DocumentRoot directory-path`
Default:	`DocumentRoot "/usr/local/apache/htdocs"`
Context:	server config, virtual host
Status:	Core
Module:	core

This directive sets the directory from which `httpd` will serve files. Unless matched by a directive like `Alias`, the server appends the path from the requested URL to the document root to make the path to the document. Example:

```
DocumentRoot "/usr/web"
```

Apache2 on Ubuntu

- Master configuration file:
  - **/etc/apache2/httpd.conf**

```
sandiway@sandiway-VirtualBox:~$ cd /etc/apache2/
sandiway@sandiway-VirtualBox:/etc/apache2$ ls
apache2.conf    conf-enabled  magic           mods-enabled  sites-available
conf-available  envvars       mods-available  ports.conf    sites-enabled
sandiway@sandiway-VirtualBox:/etc/apache2$ ls -l
total 80
-rw-r--r-- 1 root root  7224 Oct  3 07:41 apache2.conf
drwxr-xr-x 2 root root  4096 Oct 24 20:43 conf-available
drwxr-xr-x 2 root root  4096 Oct 24 20:43 conf-enabled
-rw-r--r-- 1 root root  1782 Jun 27 10:05 envvars
-rw-r--r-- 1 root root 31063 Jun 27 10:05 magic
drwxr-xr-x 2 root root 12288 Oct 24 20:43 mods-available
drwxr-xr-x 2 root root  4096 Oct 24 20:43 mods-enabled
-rw-r--r-- 1 root root   320 Jun 27 10:05 ports.conf
drwxr-xr-x 2 root root  4096 Oct 24 20:43 sites-available
drwxr-xr-x 2 root root  4096 Oct 24 20:43 sites-enabled
sandiway@sandiway-VirtualBox:/etc/apache2$
```

Apache2 on Ubuntu

- `cd /etc/apache2/`
- `grep –r DocumentRoot` (`grep –r` searches into subdirectories for string)

```
sandiway@sandiway-VirtualBox:/etc/apache2$ grep -r DocumentRoot
sites-available/000-default.conf:       DocumentRoot /var/www/html
sites-available/default-ssl.conf:               DocumentRoot /var/www/html
sandiway@sandiway-VirtualBox:/etc/apache2$
```

- For httpd, relevant file is:
`/etc/apache2/sites-enabled/000-default.conf`

Apache2 on Ubuntu

`nano /etc/apache2/sites-enabled/000-default.conf`

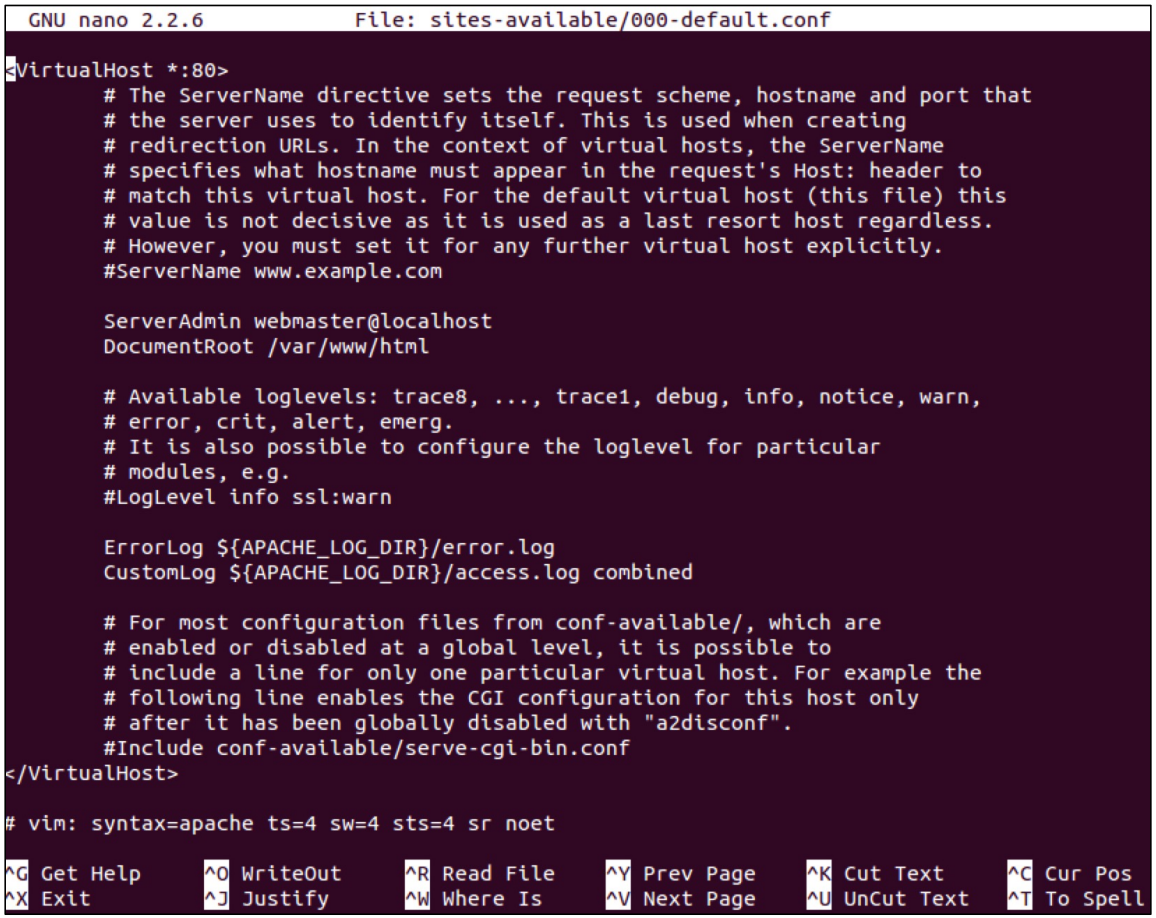

Apache Webserver on macOS

- Configuration file: `/etc/apache2/httpd.conf`
  - DocumentRoot

```
232 # DocumentRoot: The directory out of which you will serve your
233 # documents. By default, all requests are taken from this directory, but
234 # symbolic links and aliases may be used to point to other locations.
235 #
236 DocumentRoot "/Library/WebServer/Documents"
237 <Directory "/Library/WebServer/Documents">
238     #
239     # Possible values for the Options directive are "None", "All",
240     # or any combination of:
241     #   Indexes Includes FollowSymLinks SymLinksifOwnerMatch ExecCGI MultiViews
242     #
243     # Note that "MultiViews" must be named *explicitly* --- "Options All"
244     # doesn't give it to you.
245     #
246     # The Options directive is both complicated and important.  Please see
247     # http://httpd.apache.org/docs/2.4/mod/core.html#options
248     # for more information.
249     #
250     Options FollowSymLinks Multiviews
251     MultiviewsMatch Any
252
253     #
254     # AllowOverride controls what directives may be placed in .htaccess files.
255     # It can be "All", "None", or any combination of the keywords:
256     #   AllowOverride FileInfo AuthConfig Limit
257     #
258     AllowOverride None
259
260     #
261     # Controls who can get stuff from this server.
262     #
263     Require all granted
264 </Directory>
```

User web files

Per-user web directories

Available Languages: en | es | fr | ja | ko | tr

On systems with multiple users, each user can be permitted to have a web site in their home directory using the `UserDir` directive. Visitors to a URL `http://example.com/~username/` will get content out of the home directory of the user "`username`", out of the subdirectory specified by the `UserDir` directive.

Note that, by default, access to these directories is **not** enabled. You can enable access when using `UserDir` by uncommenting the line:

```
#Include conf/extra/httpd-userdir.conf
```

in the default config file `conf/httpd.conf`, and adapting the `httpd-userdir.conf` file as necessary, or by including the appropriate directives in a `<Directory>` block within the main config file.

Apache2 on Ubuntu vs macOS

- Logs are in directory: `/var/log/apache2/`
  - `access.log`
  - `error.log`
- User web files in:
  - `~/public_html`
- In your home directory do:
  - `mkdir public_html`
  - `nano public_html/index.html`

- Logs are in directory: `/var/log/apache2/`
  - `access_log`
  - `error_log`
- User web files in:
  - `~/Sites`
- In your home directory do:
  - `mkdir Sites`
  - `nano Sites/index.html`

Apache2 on Ubuntu

- To enable user web files in **~/public_html**
  - **sudo a2enmod userdir**
  - **sudo systemctl restart apache2**
  - **http://localhost/~sandiway/**

```
sandiway@sandiway-VirtualBox:~$ mkdir public_html
sandiway@sandiway-VirtualBox:~$ cd public_html
sandiway@sandiway-VirtualBox:~/public_html$ nano index.html
sandiway@sandiway-VirtualBox:~/public_html$ ls
index.html
sandiway@sandiway-VirtualBox:~/public_html$ ls -l
total 4
-rw-r--r-- 1 sandiway sandiway 92 Oct 24 21:12 index.html
sandiway@sandiway-VirtualBox:~/public_html$ sudo a2enmod userdir
[sudo] password for sandiway:
Enabling module userdir.
To activate the new configuration, you need to run:
  systemctl restart apache2
sandiway@sandiway-VirtualBox:~/public_html$ systemctl restart apache2
```

Apache2 on Ubuntu

- To enable user web files in **~/public_html**
  - `sudo a2enmod userdir` **(a2dismod)**
  - `sudo service apache2 restart`
  - http://localhost/~sandiway/

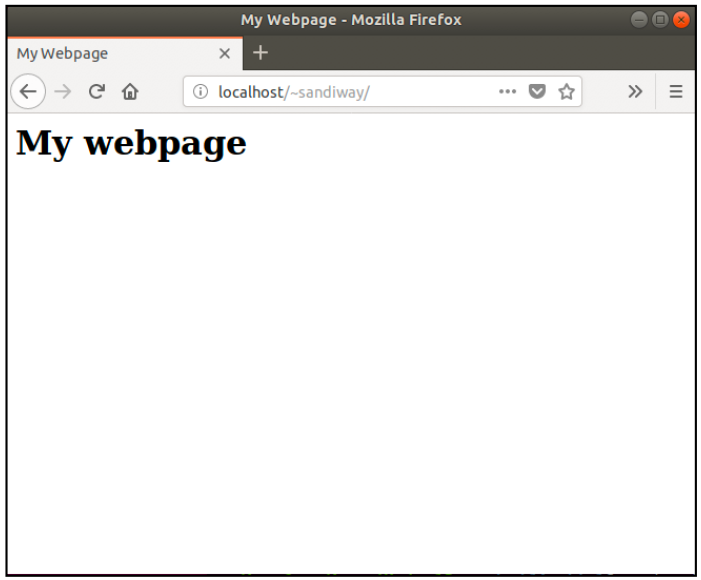

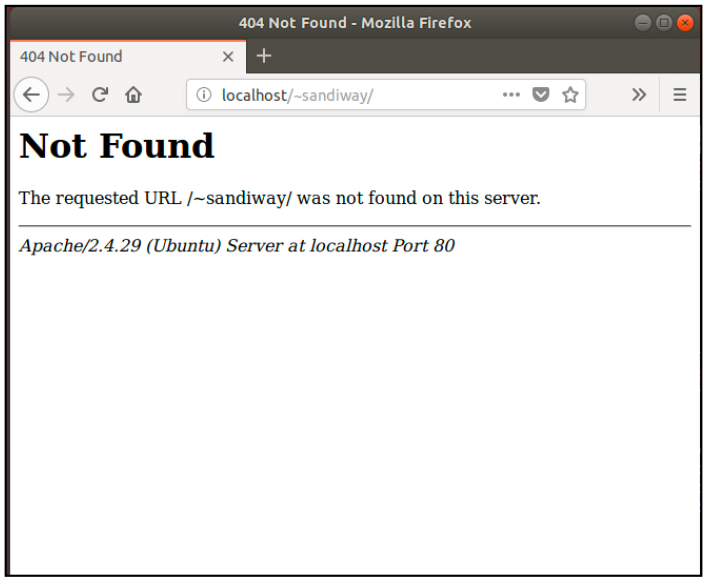

Apache Webserver on macOS

Static webpages

- storage locations:
  - http://localhost/~sandiway/ (**no need to be superuser**)
  - **`mkdir ~sandiway/Sites`** (**`/Users/`*`username`*`/Sites`**)
  - **`~/Sites/index.html`** (create this file!)
  - **`sudo nano /etc/apache2/users/sandiway.conf`**

create this file …

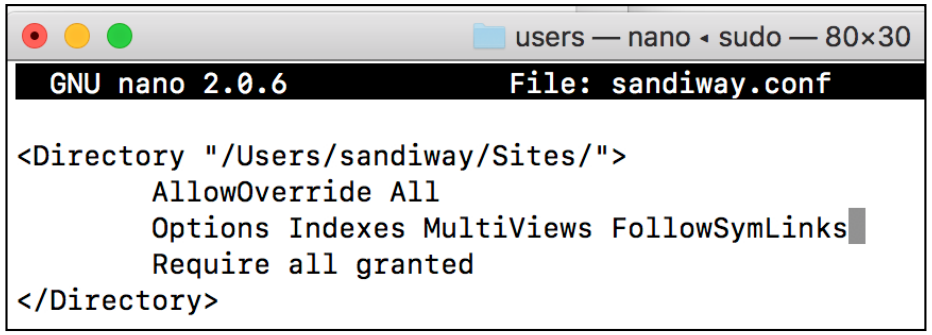

Apache Webserver on macOS

Static webpages

- http://localhost/~sandiway/

- storage locations:
  - `sudo nano /etc/apache2/httpd.conf`
- uncomment `mod_userdir.so` line
  - (*remove the comment char #*)

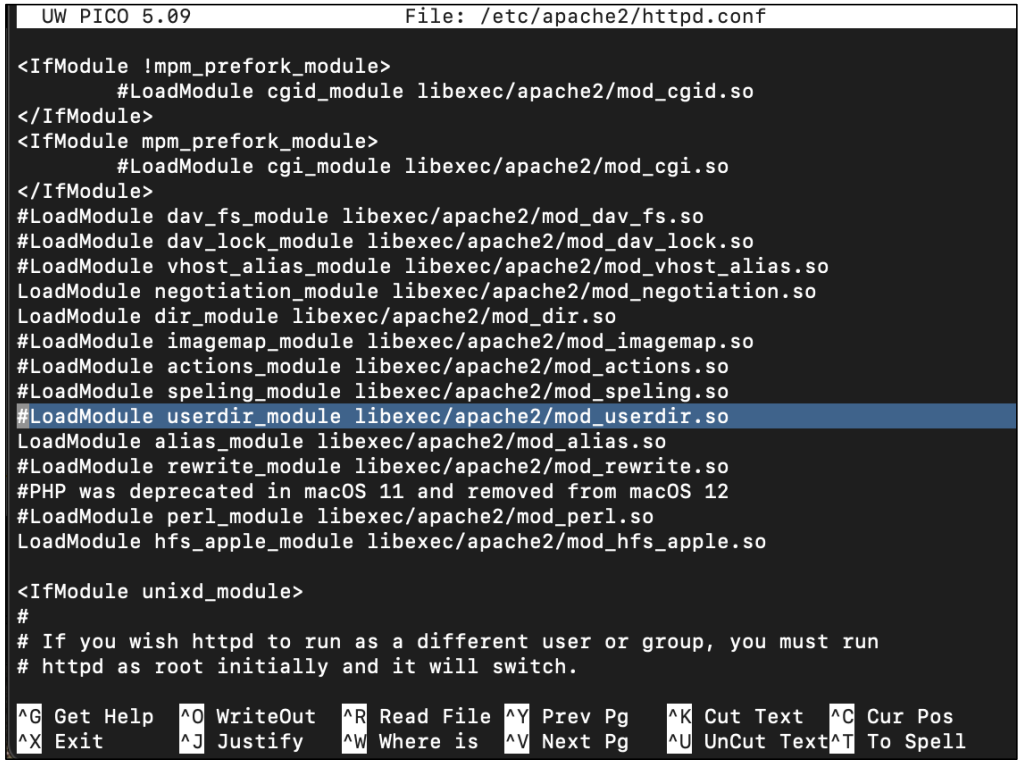

Apache Webserver on macOS

Static webpages

- also need to uncomment another line in `/etc/apache2/httpd.conf`
- uncomment the `httpd_userdir.conf` line
- (*remove the #*)

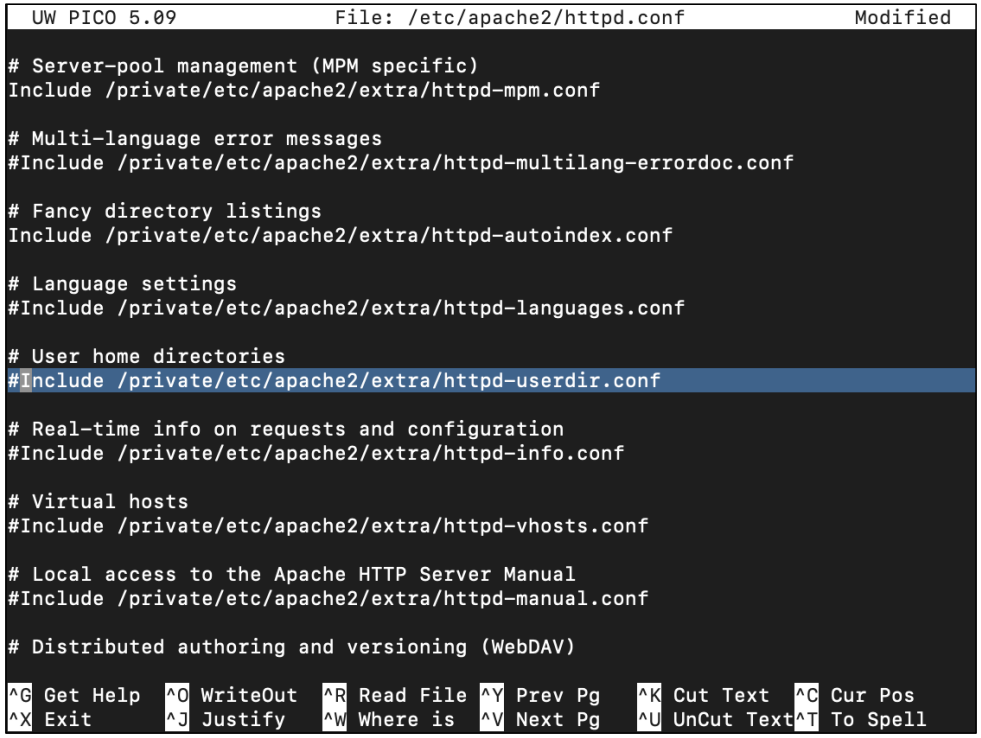

Apache Webserver on macOS

Static webpages

- `sudo nano /etc/apache2/extra/httpd-userdir.conf`
- uncomment this include
- (*remove the #*)

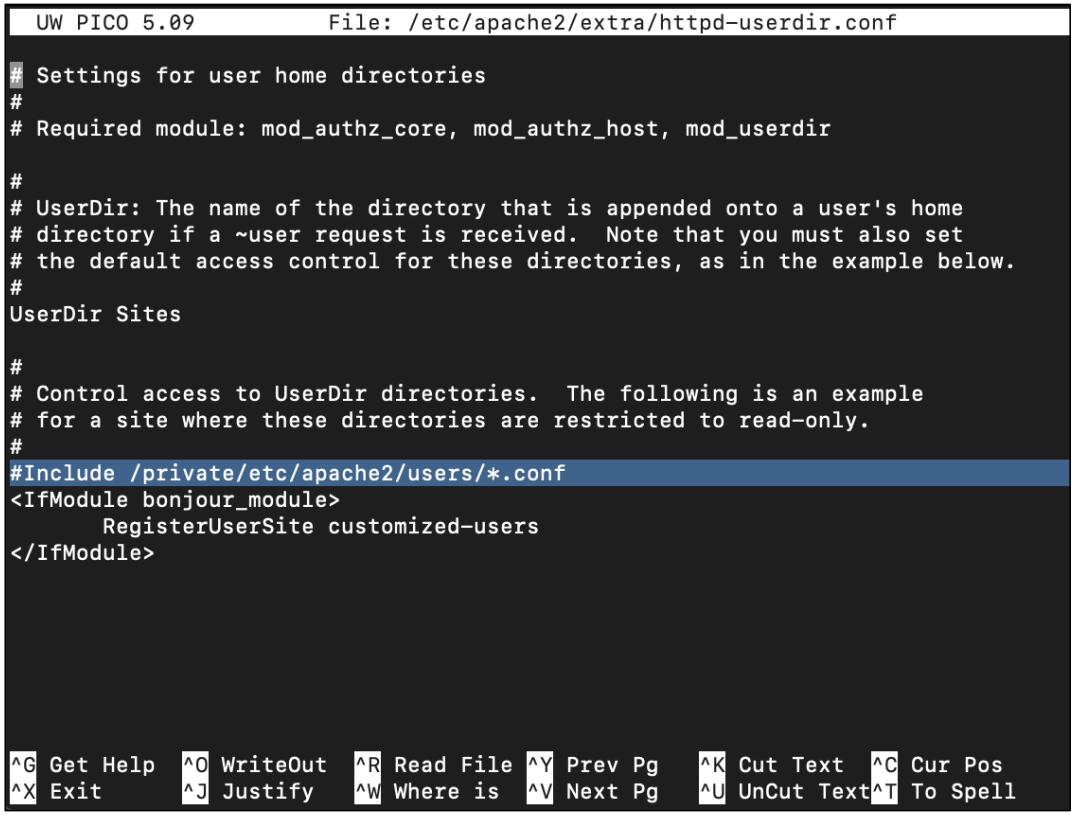

Apache Webserver on macOS

Static webpages

- storage locations:
  - `sudo apachectl –k restart`
  - create a file `~sandiway/Sites/index.html`

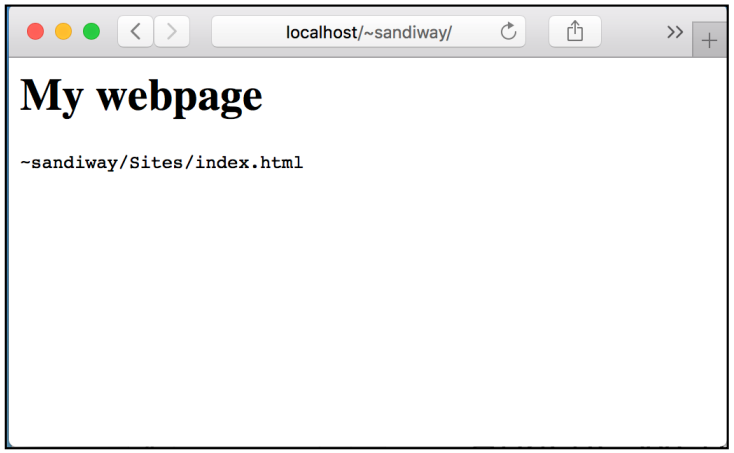

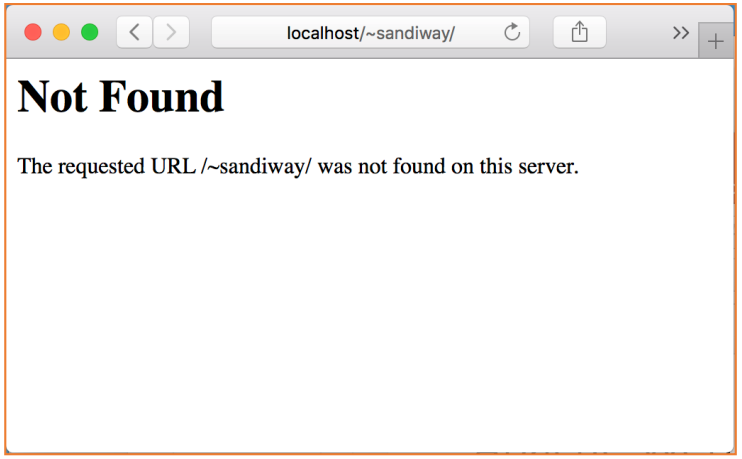

Apache Webserver on macOS

- `tail /var/log/apache2/access_log`

```
(base) apache2$ tail /var/log/apache2/access_log
127.0.0.1 - - [26/Sep/2022:05:56:09 -0700] "GET / HTTP/1.1" 200 555
127.0.0.1 - - [26/Sep/2022:05:56:29 -0700] "-" 408 -
::1 - - [30/Sep/2022:09:03:32 -0700] "GET / HTTP/1.1" 200 555
::1 - - [30/Oct/2022:21:38:44 -0700] "GET / HTTP/1.1" 200 555
::1 - - [30/Oct/2022:21:38:44 -0700] "GET /favicon.ico HTTP/1.1" 404 196
::1 - - [30/Oct/2022:21:59:51 -0700] "-" 408 -
::1 - - [30/Oct/2022:21:59:52 -0700] "GET /~sandiway HTTP/1.1" 404 196
::1 - - [30/Oct/2022:22:00:20 -0700] "GET /~sandiway HTTP/1.1" 404 196
::1 - - [30/Oct/2022:22:17:31 -0700] "GET /~sandiway HTTP/1.1" 301 235
::1 - - [30/Oct/2022:22:17:31 -0700] "GET /~sandiway/ HTTP/1.1" 200 101
(base) apache2$
```

tail command

```
TAIL(1)                    General Commands Manual                    TAIL(1)

NAME
     tail – display the last part of a file

SYNOPSIS
     tail [-F | -f | -r] [-qv] [-b number | -c number | -n number] [file ...]

DESCRIPTION
     The tail utility displays the contents of file or, by default, its
     standard input, to the standard output.
```

Apache Webserver on Ubuntu