

ESE 150 – Lab 11: Networking

LAB 11

Note: This lab will be in Ketterer Lab (not Detkin)

Today's Lab has the following objectives:

1. Learn about the OSI Network Model
2. Learn basics of TCP/IP
3. Learn the basics of networking on a Linux machine
4. Learn how applications communicate over a network

Background:

Recall from lecture that a network is a collection of independent computers that exchange information with each other over a shared communication medium. There are different types of networks, namely: LANs, WLANs, WANs,

Network Types

LAN:

The acronym LAN stands for Local Area network and it describes a configuration for a network that is typically confined to a limited geographical area (like a single building or a campus). LANs can be small (3 computers) but can link hundreds of computers as well.

WLANs:

Wireless LANs use radio frequency (RF) technology to transmit and receive data over the air. This minimizes the need for wired connections. The IEEE developed the 802.11 specification for wireless LAN technology, which specifies over-the-air interface between a wireless client and a base station, or between two wireless clients directly

WAN:

Wide Area Networks (WANs) combines multiple LANs (or WLANs) that are geographically separate. Typically a dedicated connection over a hard phone line, satellite link, fiber optic connection provides physical the interconnection, but the amount of data that travels over a WLAN vs. a LAN is far less, typically special protocols govern what a WAN will send to minimize the amount of data flowing over a WAN as opposed to a LAN.

ESE 150 – Lab 11: Networking

What is a network protocol?

A network protocol defines rules and conventions for communication between network devices. It sets up a standard way for devices to communicate on a network. In the same way that there is a “protocol” (aka a set of rules) for meeting say the Queen of England: perhaps your name being officially announced and then you bow or courtesy afterwards...it is a way to establish a standard repeatable and compatible form of communication!

After a physical connection between two computers has been established on a network (e.g. – say a cable has been connected between them), network protocols define the standards that allow computers to communicate over that physical connection. A protocol establishes the rules and encoding specifications for sending data. The protocol defines how computers identify one another on a network, the form that the data should take in transit, and how this information is processed once it reaches its final destination.

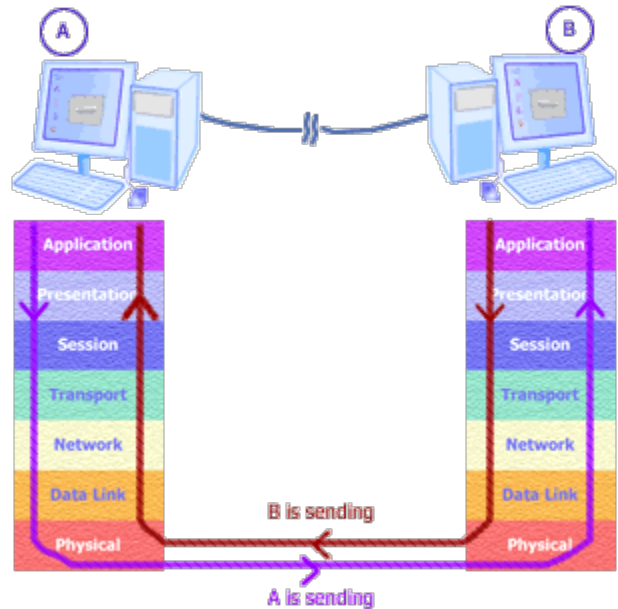
The most famous network protocol that we are probably all familiar with just by its name is: **TCP/IP (Transmission Control Protocol / Internet Protocol)**. But there are many others: AppleTalk, NetBIOS/NetBUI to name a few. Although each protocol is different, they can all share the same physical cabling (or wireless connection). This is called “protocol independence” which means devices which are compatible at the physical and data link layers in the OSI Model (talked about further below) allow the user to run many different protocols over the same medium.

ESE 150 – Lab 11: Networking

What is the TCP/IP Protocol?

The TCP/IP protocol defines a 4-layer “model” for how devices should communicate over a network. It is compared below to the OSI (Open Systems Interconnect) Model discussed in lecture.

TCP/IP Model	OSI Model
Application Layer	Application Layer
	Presentation Layer
	Session Layer
Transport Layer	Transport Layer
Internet Layer	Network Layer
Network Access Layer	Data Link Layer
	Physical Layer



TCP/IP lacks the “presentation and session” layer of the OSI model, and the datalink and physical layer of the OSI model are merged into one “network access layer” in the OSI Model.

What is Ethernet?

Ethernet falls into the “physical layer” category of LANs. It is the most popular because it strikes a good balance between speed, cost, and ease of installation (most of you are probably familiar with Ethernet ports and cables!). The Ethernet standard, specified in IEEE Standard 802.3, defines the number of conductors that are required for a connection, a framework for data transmission, and the expected performance thresholds. The standard Ethernet can transmit data up to 10 Megabits per second (other LAN types, like Fast Ethernet or Gigabit Ethernet, can transmit at 100 Mbps or 1000 Mbps). By following standard definitions of how elements of Ethernet networks are allowed to interact with each other, network equipment and protocols can communicate effectively.

IP Addresses:


Every device that wishes to use the “Internet Protocol” must have a unique number assigned to it known as an IP Address. The number is assigned to whomever administers the network one is a part of. Here at Penn, CETs is in charge of assigning IP Addresses to different devices on the network. The numbers are unique and make it so one computer can identify another computer by its number. It’s a lot like a phone number.

ESE 150 – Lab 11: Networking

IP Addresses come in two lengths: 32-bit and 128-bit. The 32-bit version is called “Ip Addressing scheme version 4 (IPv4)” and the 128-bit version is called “IP Addressing scheme version 6 (IPv6).”

IPv4 addresses have the following breakdown and you can see why they are 32-bits long, because there are 32 binary numbers. Notice the first example: 172 is the decimal representation of the binary number: 10101100. So each decimal number we see has “8-bits” underneath it. There are always 4 zero or positive decimal numbers in an IPv4 address, the highest number for any single 8-bit field is obviously 255!

An IPv4 address (dotted-decimal notation)



The first two numbers (172.16 in the example above) are typically referred to as the network identifier. And the last two (254.1 in this example) are known as the host identifier. Typically an organization (like Penn) will be identified by the network identifier and individual machines are assigned host identifiers.

Let's say my I purchase internet service from Verizon for my organization and they assign me the network identifier: 145.45. That would be the root for all my IP addresses in my organization. Let's say my first computer would be given the IP Address: 145.45.0.1, and the second: 145.45.0.2 and so on. But let's say I wish to further sub-divide my internal network. I could make the 3rd “octet” of my IP Address represent different departments, say:


145.45.0 represents the English Department

145.45.1 represents the Math Department.

Then I could have nearly 256 computers in each department and I'd know which was in which department by looking at that third “field” in the IP Address

ESE 150 – Lab 11: Networking

In IPv6 there are simply more IP addresses available to assign to different devices. Notice the format of the IPv6:



It is no longer in “decimal” it’s in what is known as hexadecimal (base 16).

IP Addresses and DNS Names

As you’ll learn later in the lab, IP Addresses can be associated with what are called DNS names or network names. It’s kind of like a phone # /name contact list. IP addresses go from very general to very specific. Let’s say my organization is “tom.com” and I get the network identifier: 145.45 from Verizon when I purchase internet service from them. I can associate my computer called “www” with an IP Address, let’s say I use 145.45.0.1 for my “www” machine. DNS will “map” these two things together:

Example DNS Record for tom.com:

www.tom.com 145.45.0.1

Notice, 145.45 is really associated with “tom.com” while the host identifier: 0.1 is associated with the computer named: “www”. So while DNS names go from specific (www) to general: (tom, then .com), IP Addresses go in the opposite direction: 145.45 (that’s tom.com), .0.1 (that’s www).

ESE 150 – Lab 11: Networking

PreLab:

PreLab – Section 1: IP Addresses

- Learn the basics of IP Addresses and their use in the TCP/IP Protocol

This section assumes you have read the background up above if you haven't please do now!

1. Login to the ENIAC server using your personal computer. Linux and MAC users can directly ssh into the ENIAC server.

```
ssh pennkey@eniac.seas.upenn.edu
```

If you are a windows user you can use PuTTY, you can download it from here

<https://www.putty.org/>.

Install PuTTY and open the application. Next click on SSH under Connection type and under Host name you can type the following

```
pennkey@eniac.seas.upenn.edu
```

Then click open. You might encounter with another window opening, press Ok and continue. In the PuTTY application enter your password and you will be logged in.

(or see instructions from Lab 9.)

2. Determine your "IP Address" by typing the following command at the shell:

```
/sbin/ifconfig
```

- a. "ifconfig" stands for "interface configuration", it shows you all the "network interfaces" you have on your computer. Think of your own laptop, you likely have a LAN connection you can plug a cable into, a Wi-Fi connection or a Bluetooth interface.
- b. You'll see the results of ifconfig are grouped by their interface. You may see "eth0", "em1", "em2", "lo", "wlan" as examples. "eth0" would represent the "0th" or first network card in the computer with an Ethernet connection.
- c. Look at the top of each grouping, after the words: "Link encap" This tells you how the network interface is connected to the network. For hardwired and wireless connections you'll commonly see the word "Ethernet" next to it
- d. Next you'll notice "HW Address" It is in HEX format, separated by colons:
Example: C8:1F:66:C4:B6:3A, that's actually 48-bits long, as each number/letter represents 4-bits each. This address is a unique hardware address that the network interface has assigned to it from its manufacturer. This is very different from an "IP address" and its often called a MAC address.
- e. Finally look for the address after the words "**inet addr**". The number that follows is your machine's 32-bit address. At Penn this number typically starts with: 158.
- f. Look lastly for your 128-bit address. It is located after the words "inet6 addr" and typically begins with "fe06" There may be many of them, but only one will begin with "fe06".

ESE 150 – Lab 11: Networking

- g. Lastly, notice the “Local Loopback” interface. This isn’t an actual physical interface, it is more like a logical “test” interface that always “loops back” to the network card itself. It always contains the 32-bit address: “127.0.0.1”
 3. Answer the following questions (in your writeup)
 - a. What was the name (eg eth0, vmnet1, em1, etc) of the network interface that contained the 32-bit IP Address starting with 158?
 - b. What was the 32-bit IP Address?
 - c. What is that 48-bit hardware address in hex form?
 - d. What was the 128-bit IP Address (in hex form, as it is displayed in the output)?

Prelab – Section 2: Names of Computers

1. IP Addresses are long and hard to remember, kind of like phone #'s. We instead like to associate IP addresses with names! Much like you put phone #'s in your phone and associate them with people's names.
2. Not every IP Address is associated with a name, but it can be if it's an IP Address people need often. Here at Penn, you are familiar with our school's web server: www.seas.upenn.edu.
3. There is something called a “domain name server (DNS)” that keeps the association between names and ip addresses. Each organization maintains their own “name server.” For example, PENN has a DNS that maps IP Addresses to names of computers.
4. You can ask the DNS for the IP address of a computer if you know its name as follows:

```
nslookup www.seas.upenn.edu
```

This simply means “name server” lookup eniac.seas.upenn.edu, it will return with:

```
Server:          128.91.18.1      (this is Penn's DNS computer's IP Address)
Address:         128.91.18.1#53
eniac.seas.upenn.edu canonical name = eniac-1.seas.upenn.edu.
Name:   eniac-1.seas.upenn.edu
Address: 158.130.64.112      (this is ENIAC's IP Address)
```

Some computers have many IP Addresses, this can simply mean that they have lots of “network interfaces” (e.g. connections to the network)

5. The “nslookup” utility also works in reverse. If you know the IP Address of a computer, if it has a name, it can tell you what it is! As an example:

```
nslookup 158.130.64.112
```

This returns:

```
Server:          128.91.18.1      (this is Penn's DNS computer's IP Address)
Address:         128.91.18.1#53
112.64.130.158.in-addr.arpa name = www.seas.upenn.edu
```

You may see something besides “eniac” in the name because eniac is actually composed of more than one server!

6. Find the IP Address of Apple's website server: www.apple.com
 - a. –submit this with your report too!
 7. Your computer can also have a private name that can be different from its official “network name.” Sometimes only the computer itself knows its private name. On Linux you can find out your computer's “host” name as its called by typing in:

ESE 150 – Lab 11: Networking

`hostname` (submit this with your report)

8. Also, you can find it by typing in:

`uname -a` (submit this with your report)

- a. *but this also return all sort of cool stuff, like what version of Linux you have, what type of computer architecture you have: x86_64 would be an Intel based 80x86 computer with a 64-bit word size for example!*
9. *Include a screenshot of all output above (make sure to label this in your writeup)*

ESE 150 – Lab 11: Networking

Prelab – Section 3: Testing Connections

1. PING is a utility that can send a “test” signal from one computer to another computer via the IP Address or its DNS registered name. Administrators like to use this utility to see if another computer is on, and properly on a network. But you may find it useful to if you’re trying to do the same!
2. Try “pinging” `www.seas.upenn.edu` by typing in the following command:

```
ping -c 4 www.seas.upenn.edu
```

```
PING eniac-1.seas.upenn.edu (158.130.64.112) 56(84) bytes of data.  
64 bytes from plus.seas.upenn.edu (158.130.64.112): icmp_seq=1 ttl=64 time=0.035 ms  
64 bytes from plus.seas.upenn.edu (158.130.64.112): icmp_seq=2 ttl=64 time=0.036 ms  
64 bytes from plus.seas.upenn.edu (158.130.64.112): icmp_seq=3 ttl=64 time=0.039 ms  
64 bytes from plus.seas.upenn.edu (158.130.64.112): icmp_seq=4 ttl=64 time=0.036 ms  
  
--- www.seas.upenn.edu ping statistics ---  
4 packets transmitted, 4 received, 0% packet loss, time 2999ms  
rtt min/avg/max/mdev = 0.035/0.036/0.039/0.006 ms
```

3. How to interpret the above output?
 - a. You asked ping to send “4” packets of data over to eniac
 - b. Each “packet” of data contained 64 bytes of data
 - c. The “round trip” time it took to for the 1st packet get back and forth was 0.035ms
 - d. *This can be different each time depending on what route through the network the packet took, or how busy the router was that routed the packet to eniac*
 - e. But on the average I transmitted 4 test packets of data to eniac, and eniac responded each time with an average speed of about 0.036ms!
 - f. It’s not a precise timing measurement, but it’s a good average number
4. Include a Screenshot of your output from above
 - a. What is your average roundtrip time?
5. Try pinging `www.apple.com` from eniac:
 - a. Save the results to hand in with your report (screenshot)
 - b. What was the average time?

ESE 150 – Lab 11: Networking

Lab Procedure:

Lab – Section 1: Routes between computers

1. Login to the Window Computer at your station
2. Bring up a command terminal (type command in the search bar under Start)
3. *Traceroute* is another handy shell utility that works like PING, except it records what machines (hubs/switches/routers) it sends the packets through. So let's say you "ping" apple.com, traceroute can show you the "route" the packets took to get there.
4. Run traceroute between your computer and www.ucla.edu:
 - a. Type the command:

```
tracert www.ucla.edu
```

Following is an example from tracert www.apple.com

The utility will output a list, showing each piece of equipment it went through to get to apple.com, and then to the "www" machine on apple.com's network:

```
tracert to www.apple.com (104.97.94.156), 30 hops max, 60 byte packets
 1  158.130.0.233 (158.130.0.233)  25.881 ms  34.633 ms  25.846 ms
 2  isc-ist.seas.upenn.edu (158.130.0.250)  0.278 ms  0.293 ms  0.284 ms
 3  vag-brdr.i2trcps-ashb.router.upenn.edu (128.91.238.222)  4.480 ms  4.495 ms  4.4ms
 4  18.ae15.pr0.dca10.tbone.rr.com (107.14.16.81)  4.478 ms  4.469 ms  4.460 ms
```

... (I've deleted some of the route)

```
 8  bu-ether12.nycmny837aw-bcr00.tbone.rr.com (66.109.6.27)  7.463 ms  12.549 ms
12.983 ms
 9  0.ae0.pr1.nyc20.tbone.rr.com (107.14.17.216)  12.039 ms  7.553 ms  66.109.1.59
(66.109.1.59)  7.525 ms
10 a104-97-94-156.deploy.static.akamaitechnologies.com (104.97.94.156)  7.724 ms
7.754 ms  7.544 ms
```

- b. This is just the first 4 "hops" on your packet's route to apple.com and the last 3 as well (8, 9, 10). Yours may be different than mine. The packet leaves your computer and goes most likely to a switch identified by IP address: 158.130.0.233. Then that device forwards your packet and it forwards it to a UPENN router named: va-brdr. This router then forwards your packet to our internet provider (rr = road-runner, part of time-warner cable) and it keeps on going all the way to "akamaitechnologies.com" that a content provider that is likely hosting apple.com's webserver!
- c. Each of these IP addresses is located somewhere different around the country. Trace the geographical locations for your tracert to www.ucla.edu using <https://www.iplocation.net/>, and create a map displaying each hop to turn in with your

ESE 150 – Lab 11: Networking

lap report. (You can use a tool like this to plot all points on a single map:
<http://www.darrinward.com/lat-long/>)

5. Run tracert between your computer and another lab team's computer
 - a. Ask a group next to you for their IP address or their hostname
 - b. First, ping their IP or hostname, to verify you have a connection between your computers, **record these results...**
 - c. Next use tracert to see how many devices you must go through to get over the network to your neighbor's computer. It shouldn't be many if the network is properly setup.
 - d. **Record your route and explain each hop in your writeup**

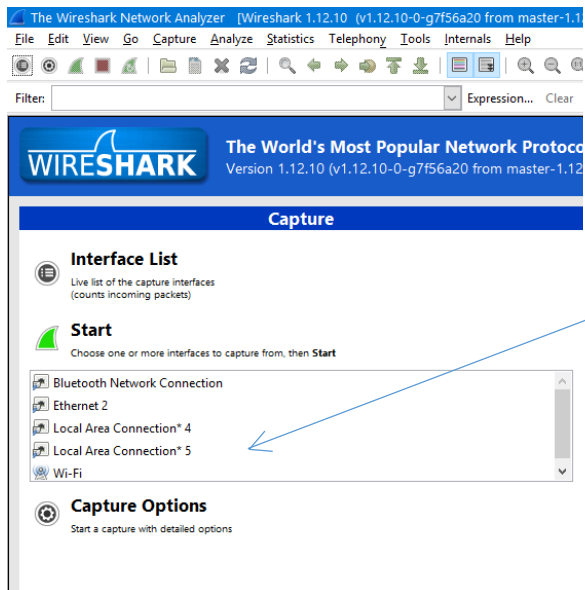
6. Run tracert between your computer and "www.princeton.edu"
 - a. Record the route. For your lab report, look up the locations of each hop using the IP address just like in step 3. You do not have to plot them on a map.
 - b. In your lab report, note the differences between your route and that in step 4

ESE 150 – Lab 11: Networking

Lab – Section 2: Investigating the packets

- Recall from lecture that a packet is the actual data that you send from one computer to another. Usually it is a fixed size. Let's say you're sending 128 kBytes to another computer, TCP/IP may break that up into 128 packets with about 1 kBytes each. This way if something goes wrong with one of the packets, it can just re-send one small packet, instead of re-sending the entire 128kBytes in one shot.
 - Also recall from lecture that a packet not only contains the data you want to send, but information about where the packet came from and where it went to!
 - We can look at the packets if we like using either an oscilloscope, or something a bit higher level, a software "packet sniffer" or a "network analyzer"
 - In this section, we'll use a utility called: wireshark (it is a software level packet sniffer)
 - You don't need to, but you can obtain this software for free from **wireshark.org**
1. Switch to the linux machine at your station.
 2. Use ifconfig to determine the IP address for your linux workstation and the name of the configured network interface (like you did on eniac in the prelab).
 3. On your linux machine, from the terminal type in:
`wireshark-gtk &`

If properly installed you'll see something like this:



Click on the network interface name that corresponds to the IP address from Step 2: Next, click on the "start" or green sharkfin button

If things are seeing is a growing list. It is a copy of every packet being handled by your network interface!

ESE 150 – Lab 11: Networking

- Back in your linux terminal type in **(or your computer's terminal if you are using your own version of wireshark)**

```
nslookup www.seas.upenn.edu
```

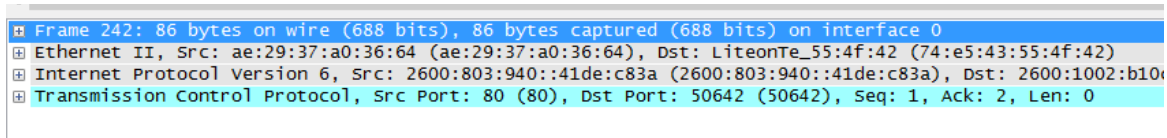
Record the IP address

```
ping -c 4 (the IPv4 address of eniac)
```

or

```
ping (the IPv4 address of eniac) (if you are using a PC)
```

- Press the “stop” button in wireshark
 - Wireshark has just copied all of the ping “packets” you sent to eniac
 - It’s kind of like a software/network oscscope!
- Next, find the PING packets...
 - There are many packets leaving/coming to your computer all the time. So finding them can be difficult.
 - Sort the list at the top by “destination” and scroll until you see eniac’s IP address
 - Make sure that the “source” is the IP address of your computer
 - Note, the protocol is “ICMP” that’s for Internet Control Message Protocol, it’s a protocol used with “IP” in TCP/IP for “control” signals send between computers, like “ping”
 - You should see 4 packets (so four entries)
- Click on the first packet: look at the “info” column on the right hand side, look for “echo (ping) request” but the one with the lowest “sequence #”
 - With the packet highlighted, look to the middle window, which shows you four components of the packet:



```
Frame 242: 86 bytes on wire (688 bits), 86 bytes captured (688 bits) on interface 0
Ethernet II, Src: ae:29:37:a0:36:64 (ae:29:37:a0:36:64), Dst: LiteonTe_55:4f:42 (74:e5:43:55:4f:42)
Internet Protocol Version 6, Src: 2600:803:940::41de:c83a (2600:803:940::41de:c83a), Dst: 2600:1002:b10
Transmission Control Protocol, Src Port: 80 (80), Dst Port: 50642 (50642), Seq: 1, Ack: 2, Len: 0
```

- Let’s examine these components
 - Click on the first line: **“frame”** and look to the window at the bottom:
Notice, all that “hexadecimal” data? That is the actual message that was sent between your computer and eniac
When you click on “frame” it highlights the entire message, as this packet is considered a “frame”
 - Now click on second component in the middle window: **“Ethernet II”**
 - Notice how it only highlights some of your packet in the window below?
 - That’s because only that part of the packet is for the “Ethernet II” layer of the TCP/IP model
 - Click on the little “plus” or “arrow” sign next to Ethernet II to expand it
 - Notice the “source” information: that’s your computer. Click on it and see which part of the packet is highlighted

ESE 150 – Lab 11: Networking

- 2) Look at the destination information as well! Which part of the packet makes up the destination?
- c. Click on the “Internet Protocol Version 4...” component to note which part is highlighted below. Expand this section and scroll through to look at the different components.
- d. Lastly, expand and look at the “Internet Control Message Protocol” component
 - i. This is the Data inside the ICMP (or ping) packet. Identify the sequence number and type encoding that says it is a ping packet. Include these in your report.
 - ii. Identify all the ping response packets. Inside a ping response packet, identify the type encoding for the response. Identify the round trip time and relate it to what ping printed in your console window. Include these in your report.
- e. Lastly, record the entire request packet (the hex data).
 - i. Indicate which parts of the hex belong to the following parts of the frame:
 - 1) Ethernet II
 - a. Source
 - b. Destination
 - 2) IP Protocol Version 4
 - a. Source
 - b. Destination
 - 3) ICMP
 - a. Data
 - b. Sequence number
 - c. Type encoding
- f. Draw a picture of the TCP/IP model. Show which parts of the frame belong to which part of the model (e.g. – application/transport/etc – see the background section). Submit this with your report.

ESE 150 – Lab 11: Networking

Lab – Section 3: Connecting to a webserver on port 80

- A webserver is just a “server” program running on remote computer (say apple.com)
 - Its purpose is to “serve” up webpages to clients requesting a webpage within a website.
 - A webserver program is set to “listen” for incoming clients on port #80 on a host (like www.apple.com)
 - When a client program (like a web browser - Chrome for example) opens up a socket to www.apple.com on port #80, the web server program on www.apple.com is identified by port #80 and allows it to connect. Then data is sent back (in the form of web pages) to the client.
1. Continuing to work on the linux machine, make a new directory called ese150_lab10b and cd into it
 2. Open up a web browser (any one will do) and type in the following address manually:
`http://ic.ese.upenn.edu:80/research`
 - a. What does this mean? By specifying the “:80” your browser is acting as a client, asking the server named: ic.ese.upenn.edu to talk directly to the process listening on PORT #:80
You can also simply type in: <http://ic.ese.upenn.edu/research> and the same thing will happen, web server’s typically use port 80 to listen for incoming requests
 - b. Next, type the webserver’s IP Address directly: (Try it for apple.com)
--this also works, your browser doesn’t need to work as hard, you specified the IP Address of the host, so now it doesn’t need to go to the DNS server first!
 - c. What does “http” mean? It stands for “**h**yper **t**ext **t**ransfer **p**rotocol, h t t p”
You have specified what protocol you’d like the browser to communicate with the remote server. This “http” protocol is used for exchanging “hyper text” files, or what is called: “HTML” files (Hyper Text Markup Language)
 - d. A “web browser” is a client that requests and can even display “.HTML” files. When it receives a .HTML file from a remote server, it opens it up an interprets the information inside as the text and graphics you see on your screen!
 - e. Include a screenshot of your output in your report

ESE 150 – Lab 11: Networking

3. Let's use a different "client" to look at the ic website.
 - a. Open a terminal window and type in:
`telnet 158.130.67.172 80`
 - b. "TELNET" is a simple client program that can open up a socket to the IP address you specify and to the PORT # you've specified. Above, you've opened up a socket to www.seas.upenn.edu, and telnet to be connected directly with the process running on port 80 (which is the webserver!)
 - c. Telnet shows its progress:

```
Trying 158.130.67.172...
Connected to alliance-vhosts.seas.upenn.edu.
Escape character is '^]'.
```

--This means it connected to ic.ese.upenn.edu:80
--But, ic.ese.upenn.edu:80 expects a "web browser" as a client, not telnet.
--So we must pretend to be a web browser and send the type of information a web browser typically sends to a web server

- d. Send the following information to ic.ese.upenn.edu, by typing it in (or copying and pasting):

```
GET /research.html HTTP/1.1
Host: ic.ese.upenn.edu
User-Agent: telnet
```

- e. Press <enter> twice and watch what happens!
 - f. If you typed things in correctly, you'll see the contents of the ".HTML" file research.html on the ic.ese.upenn.edu server.
 - g. TELNET doesn't really know how to display that data, other than just showing its contents to you. It doesn't understand that the .HTML file that has returned could be interpreted as graphics and formatted text, only a web browser knows how to do that
 - h. You can exit out of telnet by typing the word "exit"
 4. Let's "save" what telnet gets back from ic.ese.upenn.edu
 - a. Repeat the command from before, but this time we'll "save" the output to a file:

```
telnet 158.130.67.172 80 > research.txt
```

- b. This time, you won't see telnet's status, instead, that output is being stored in detkin.txt
 - c. **Include this file in your report writeup**
 - d. Manually type in the following (EXACTLY AS IT'S SHOWN)

```
GET /research.html HTTP/1.1
```


ESE 150 – Lab 11: Networking

Host: ic.ese.upenn.edu

User-Agent: telnet

- e. Press <enter> twice, wait for about 60 seconds and it will close the connection
- f. You won't see any data returned from ic.seas.penn.edu, but it was all stored in the file: research.txt
- g. Open up research.txt in a text editor as follows:

```
nano research.txt
```

- h. Notice, its that "HTML" stuff from before! This has all come back from ic.ese.upenn.edu and was stored in your file: "research.txt"
- i. Delete the top lines of the file (your version may be a little different):

```
Trying 158.130.67.172...
Connected to alliance-vhosts.seas.upenn.edu.
Escape character is '^]'.
HTTP/1.1 200 OK
Date: Sun, 15 Apr 2018 00:16:58 GMT
Server: Apache
Last-Modified: Sun, 15 Jan 2017 02:50:39 GMT
ETag: "1c24-5461921daa0cf"
Accept-Ranges: bytes
Content-Length: 7204
Content-Type: text/html
```

- j. Now, save the file by typing: <ctrl> X
 - i. Answer "Y" when it asks you to save
 - ii. Press <enter> when it shows you the file name, change it to "research.html" and press enter.
 - iii. Note, you'll now have 2 files: research.txt and research.html
- k. Now, let's open the file...we'll need a web browser! Type in:
firefox research.html

Look at the "address" bar and see the address looks like this:

```
file:///mnt/castor/seas\_home/t/tfarmer/research.html
```

You've asked your web browser to look at a file stored on your computer (not a website ic.ese.upenn.edu)

Notice, some of the images just show up. That's because telnet didn't download all the graphics associated with the webpage, but for the most part, we've manually done what a web browser does with .HTML files manually!

ESE 150 – Lab 11: Networking

I. Turn in research.txt and research.html with your writeup. Explain what they are and what they contain.

5. Let's re-examine how telnet & the webserver are communicating.

a. Open up two terminal windows so that they are side-by-side

b. In the first window, type in:

```
telnet 158.130.67.172 80
```

c. In the second window, type in:

```
ps -fu
```

d. Look for the "telnet" process you are running, write down this #

i. Also, take a screenshot of this window (a TRUE screenshot using <print screen>)

ii. Paste it into a document so you can submit it later

e. If you took a long time to take your screenshot copy/paste, "telnet" may have timed out!

i. You'll need to repeat steps b & c (as the process # will change)

ii. You don't need to take a new screenshot however!

f. Next, type in the following command in the second window to look at the PORT#'s in use:

```
netstat -an |grep 158.130.67.172
```

i. NETSTAT is a program that can investigate the status of the network on your machine

ii. GREP is a program that can "filter" the output of netstat and show only lines that contain the IP address: 158.130.68.91

iii. You'll see something like this:

```
tcp 0 0 158.130.68.44:57011      158.130.67.172:80      ESTABLISHED
```

iv. Take a screenshot (turn in with your report)

v. What does it mean? It means you have "established" a network connection between your computer: 158.130.68.44 (that's the computer I used when I wrote this tutorial, but yours should have the IP address you found in step 1)

vi. Notice, the # 57011 (you will have a different #, make a note of it for the next step). This is the PORT # telnet's socket is using on your computer. So when ic.ese.upenn.edu returns data back to you, your computer will know which program to route that data back to!

vii. Lastly, notice, the computer that your computer's PORT 57011 is connected to, it's connected to the computer: 158.130.67.172 on its PORT 80! That's ic.ese.upenn.edu and that makes sense.

ESE 150 – Lab 11: Networking

- g. Lastly, the computer can tell you which process PORT 57011 belongs. Type in the following command:

```
lsof -i:57011 (replace 57011)
```

- i. **Take a screenshot of what comes back**, if nothing came back, you need to “reopen” the connection to ic.ese.upenn.edu, repeat steps b, c, & f above to reestablish the connection
 - ii. LSOF is called “list of open files” It thinks that a PORT that’s open is just like a process having a file open, so it helps us see what PORTs our process has open, or the reverse as we’ve done above
 - iii. Does the returned PID match what you noted in a previous step?
6. **Exit checkoff** : Show a TA the following and answer questions:
- a. Packet recorded from WireShark in Section 2: step 8e; be able to point out the fields identified there.
 - b. `research.html` file you created in step 4 by editing the `research.txt` that you captured.

ESE 150 – Lab 11: Networking

Postlab (optional):

Most of the answers to these questions come from the lab procedure itself, or the lab background. You won't need google for much of anything here!

- 1) Can a computer communicate using different network protocols over the same physical medium (i.e. – cables?)
- 2) What is the advantage of IPv4 vs. IPv6?
 - a. How many devices total could be represented with IPv4? How many with IPv6?
- 3) Let's say I have a computer with the DNS record: www.seas.upenn.edu. Let's say nslookup returns the IP Address: 158.130.68.91 for that computer. What "parts" of the IP address map to the different parts of the name? www, seas, upenn.edu...what parts are what numerically?
- 4) From your wireshark message, what exactly did a "ping" message contain? It is actual data, readable by a human, what is it?
- 5) As a "network" administrator (the administrator of a network at a company/school), if the following information was returned from netstat:

```
tcp 0 0 158.130.5.2:61234      158.130.68.91:80      ESTABLISHED
```

and the lsof returned the following:

```
FIREFOX 2887 tfarmer      3u  IPv4 2774174      0t0  TCP
big05.seas.upenn.edu:61234->158.130.5.2:http (ESTABLISHED)
```

What would this mean? Draw a picture like the one shown in the "background" section, specifically in the NETWORK SOCKETS section. Fill in all the details you see from the output of the NETSTAT and LSOF commands

- 6) Why is a PORT # necessary? Why can't a client process simply indicate the IP Address of the computer it would like to communicate with?

ESE 150 – Lab 11: Networking

HOW TO TURN IN THE LAB

- Each individual student should submit a PDF document to canvas containing:
 - Recorded data and descriptions/explanations where asked
 - There is output we require in each section, it your responsibility to look through this writeup and ensure you turn in each component required
 - **Clearly label what each screenshot is in your report (or you lose credit)**
 - Answers to questions asked in the prelab and lab.
 - Note that the postlab is optional. We suggest you work through it, but given the Wednesday due date for the lab, we are not making this required.
 - Note this lab is due on Wednesday 4/25 (last day of classes).