First Digital Computer

ENIAC (Electronic Numerical Integrator And Calculator)

•1946 - University of Pennsylvania

• Result of military research funding – the need to more quickly calculate trajectories for artillery shells. Not finished until end of World War 2.

60,000 pounds
18,000 vacuum tubes
Size of medium-sized house (1,600 square feet)

• Could only run one program at a time, which had to be fed into the computer as a sequence punch cards -- cards with holes punched into them to represent programming instructions.

• Some pictures are provided on the Web Site.

1960's – Minicomputers

- About the size of a refrigerator.

- Only cost about $20,000 so most universities could afford one.

- Multiple Terminals (keyboard+monitor) -- Each was a dumb terminal that didn't have it's own processor. But at least different people could run programs on the same computer at the same time.

- Time Sharing -- If multiple programs were running, the minicomputer would share the processor cycles among the different programs.

It's funny that these were called minicomputers, because we now literally have computers in our pockets (smartphones), each with it's own processor, and some with multi-core processors.

The First "Internet" -- ARPANET 1969

- Funded by the US department of Defense
  ARPA – Advanced Research Projects Agency
- Original 4 Hosts at Universities in the Western USA
- Original plans call for 128 Hosts

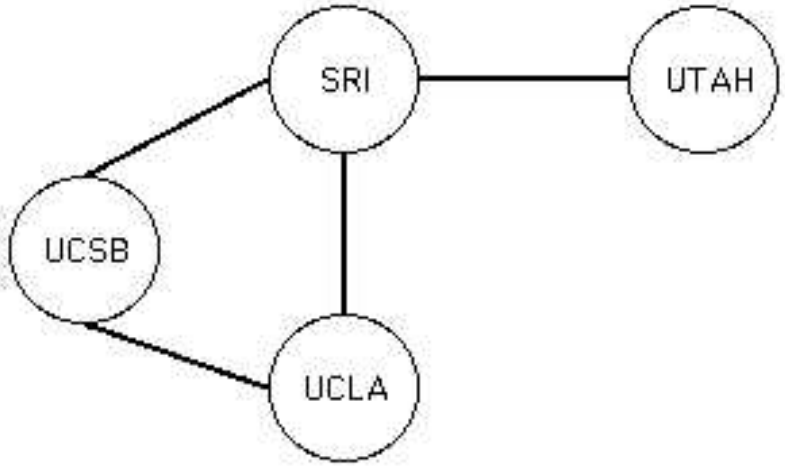

Exponential Internet Growth

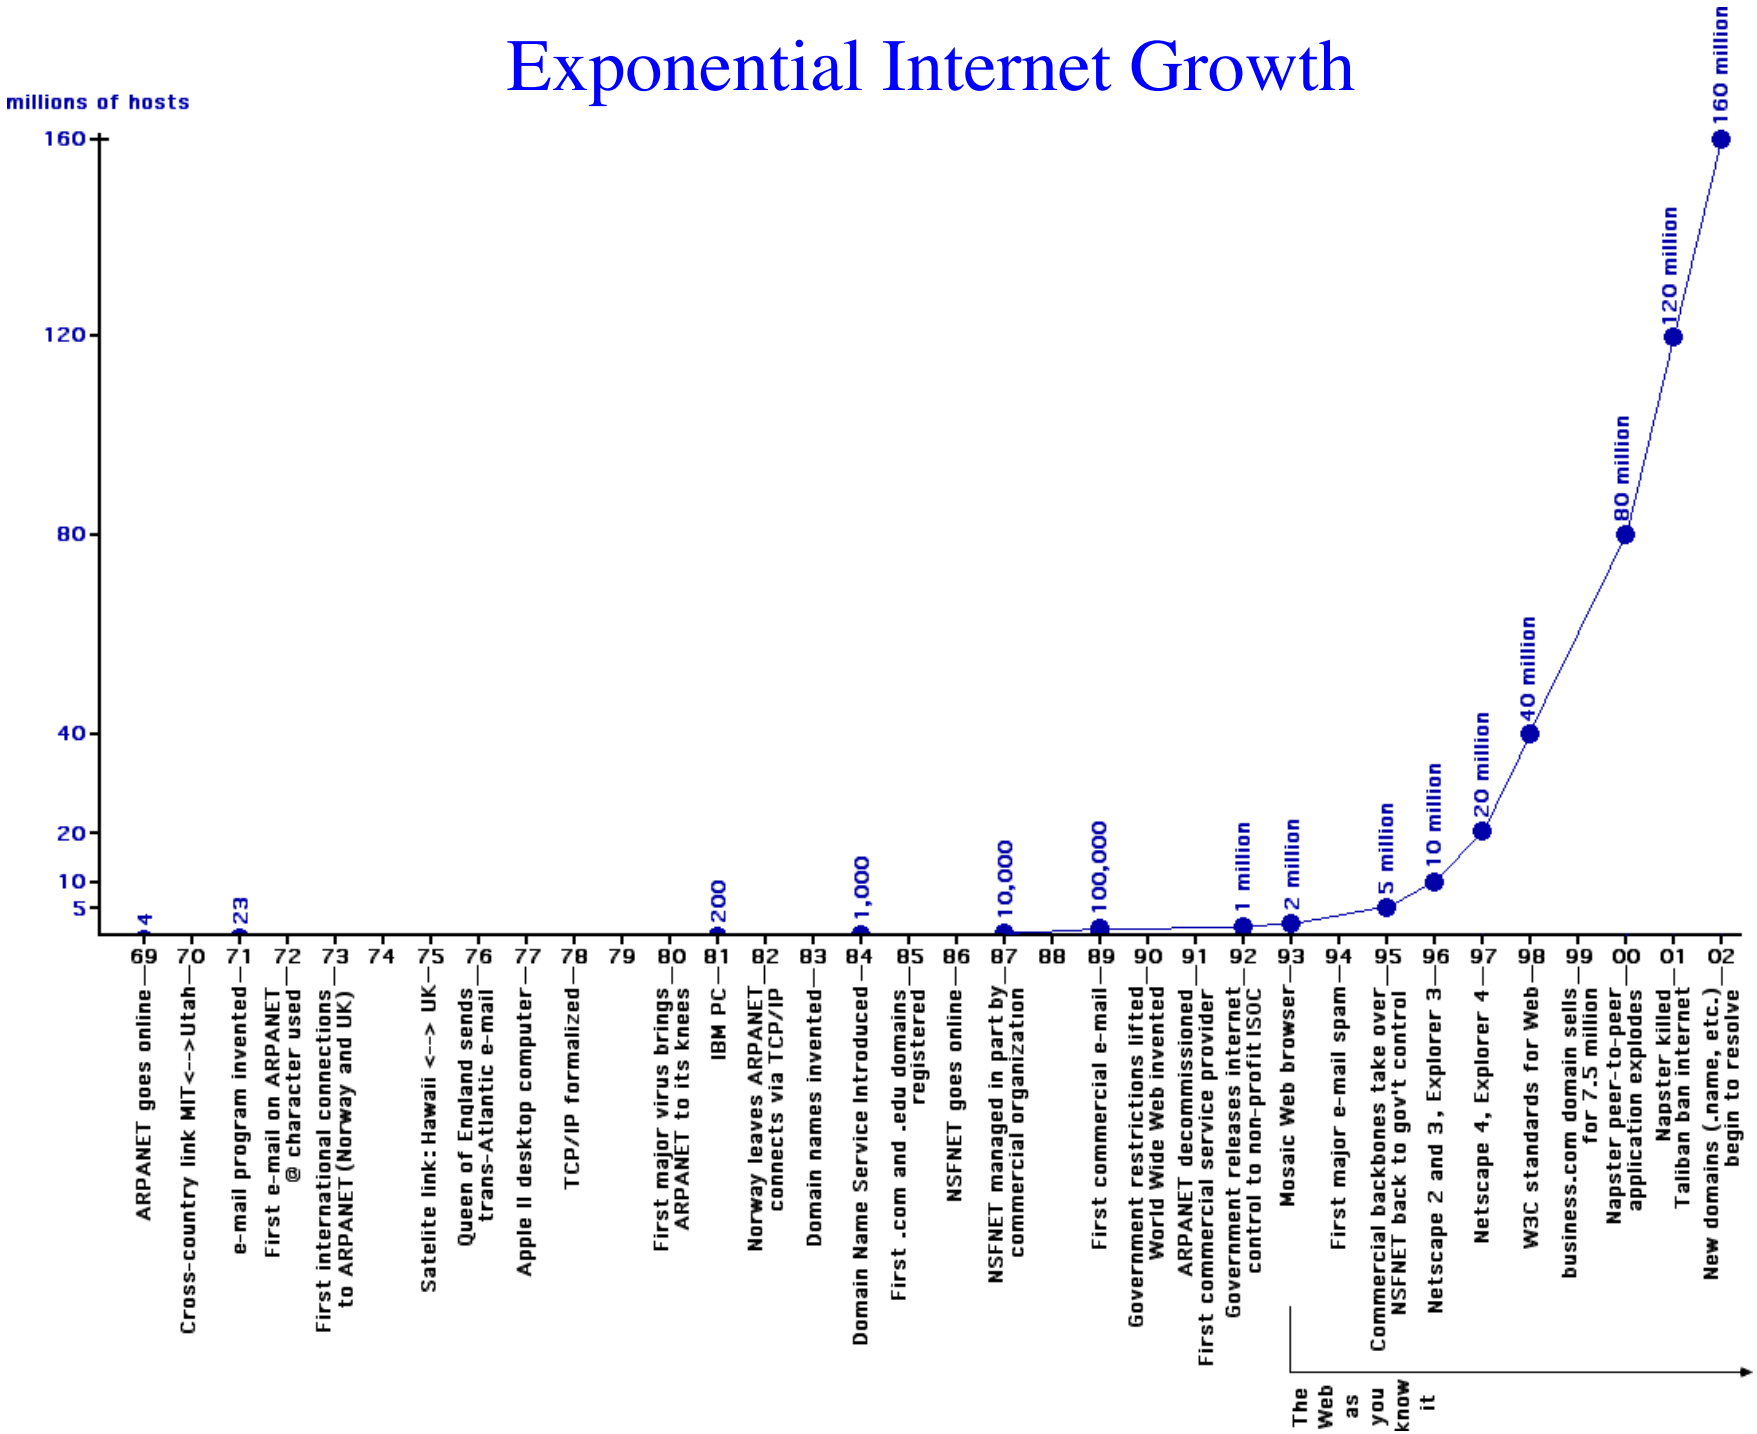

3 main things caused Internet use to achieve exponential growth beginning in the late 1980s

- Personal Computers (PCs) start becoming affordable in 1980s, so average people started owning computers, not just scientists.

- US Government releases control of Internet around 1990. This causes massive investment in Internet infrastructure by private companies.

- The World Wide Web (WWW) is invented in 1990 in Geneva Switzerland.

First desktop Computers available around 1980

- 1977 - Apple II

Apple computers (later called Macs) are the first desktop computers to feature

  – Floppy Disk Drive (1978)
  – Mouse (1983)
  – Windows Graphical User Interface (1983)

- 1981 IBM PC
  – Used Microsoft's DOS Operating System
  – Microsoft eventually patents the term *Windows*
  – Windows 1.1 (1985)

US Government releases control of the Internet

- 1986 -- Second generation Internet goes online. NSFNET, run by National Science Foundation.

- 1990 -- First year internet access could be obtained without filling out paperwork and requesting permission from the US Government.

- 1992 – Loose Internet oversight transferred to non-profit organization which eventually became the *ISOC (Internet SOCiety)*. It's membership includes researchers from major technology companies and research universities around the world.

The World Wide Web Invented around 1990.

• Many people, including prominent news commentators, and authors use the terms WWW and Internet as if they are the same thing.

• But they are radically different.

• The internet is over 20 years old when Tim Berners-Lee invents the WWW while working at the CERN nuclear physics research lab in Geneva Switzerland.

• The WWW is simply a software application that makes use of the Internet's vast infrastructure.

5-layer Internet Model

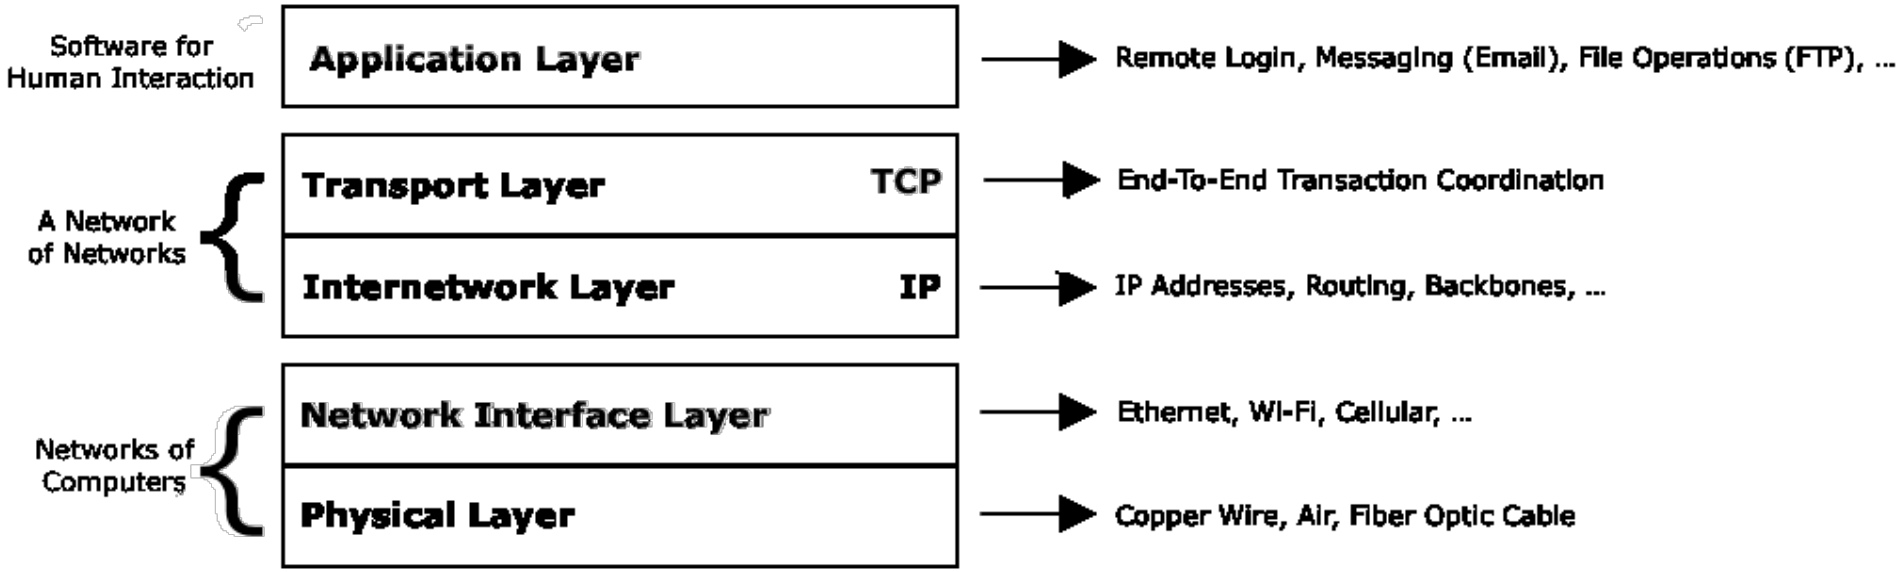

Physical Layer

- Electrons through copper cables.
- Visible light (or infrared) through fiber optic cables.
- Radio waves through the air.

This is the domain of electrical Engineers -- voltage, amperage, wavelength, etc.

Data can be lost in this layer. Electrons slam into the nuclei of copper atoms, magnetic disturbances block radio waves, etc.

The networking protocols in the next layer up ensure reliable data transfer between two or more computers.

Network Interface Layer

Ethernet – networks the size of rooms or small buildings. Computers are linked with copper ethernet cables, similar to traditional telephone cables but more thick.

Wi-Fi – wireless (though the air) networking. Most modern desktops/laptops can switch between Ethernet and Wi-Fi.

Cellular Networks (3G/4G) – Through air, but over much larger distances than Wi-Fi. Modern smart phones can switch between Wi-Fi and cellular network connections.

Internet Service Providers (ISPs) – variety of network types including cable TV modems and satellite dish connections.

Comparison of network-level data transfer rates.

Network Link	Approximate Transfer Time	Typical Bandwidth
Gigabit Ethernet, Wi-Fi (802.11ac), 4G LTE Cell	1/25th of a second	1000 Mbps (1 Gbps)
Wi-Fi (802.11n), 4G Cell, Cable Modem	.4 seconds	100 Mbps
Wi-Fi (802.11g), Satellite Dish Modem	1.25 seconds	50 Mbps
3G Cell	8 seconds	5 Mbps
Dialup Phone Modem	13 minutes	.05 Mbps (50 Kbps)

Approximate transfer times for a 5 megabyte MP3 file

For perspective, a 5 Meg is above average for an MP3 music file. (Let's say it's from Metallica or Mastodon as a matter of good taste.)

These transfer times are very approximate. In particular, several varieties of Wi-Fi have been sold widely, so when you walk into a Wi-Fi hotspot, you never know what network speed you might get.

Inter-Network Layer – *IP Protocol* facilitates packet routing between totally different networks.

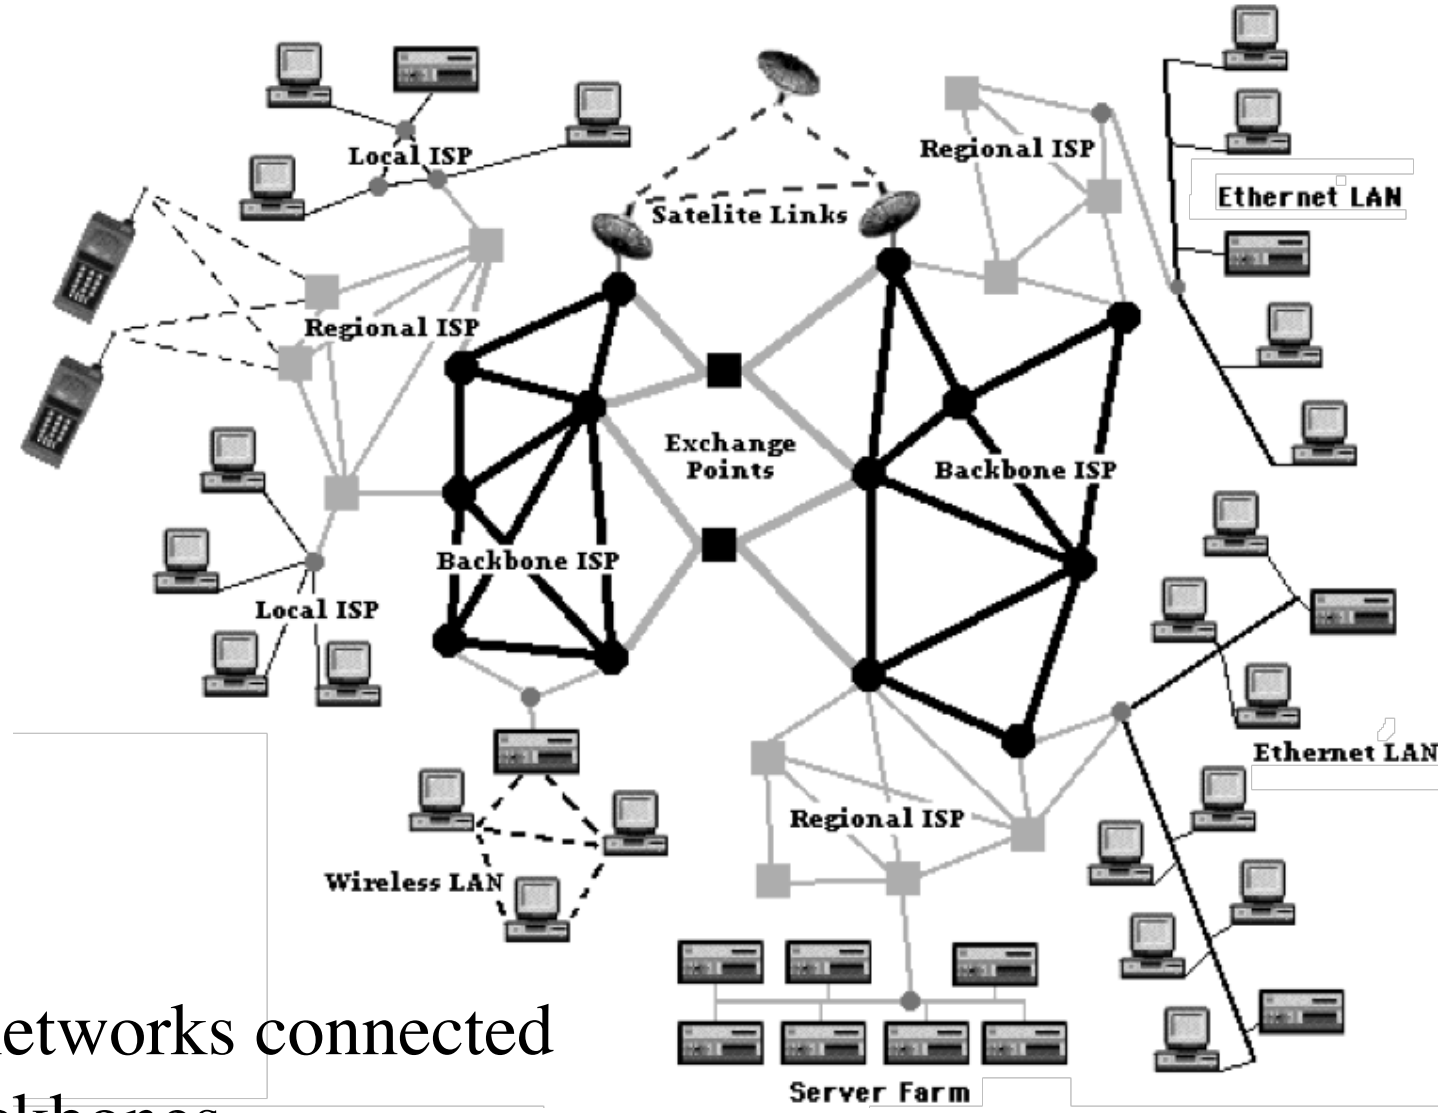

A network of networks connected by Internet Backbones

IP addresses

Example: 164.68.21.170

- Each number in the range 0-255
- Around 4 billion different IP addresses

Internet Assigned Numbers Authority (IANA) grants them for free in large blocks to Internet Service Providers (ISPs).

Class A block-- Example: 164.x.x.x (16 million addresses)
Class B block-- Example: 164.68.x.x (65 thousand addresses)
Class C block-- Example: 164.68.21.x (256 addresses)

Other IP Features

• Data is divided into small *packets*, averaging about 1.5 K in size. If data were sent in huge chunks, all of it would need to be resent if a small part if it is lost or damaged. Small packets solve that problem.

• *Flow Control* -- Packet routers try to pick the optimal path to the destination based upon how busy neighboring routers are, not necessarily the shortest path. Packets from same transaction might take different routes.

• *Time To Live* -- A packet is only allowed a certain number of router "hops." If a packet has exceeded its TTL, a router will simply delete it.

The robust design of IP is a major contributor to the success of the Internet!

Transport Layer

The *Transmission Control Protocol (TCP)* coordinates the end-to-end details of an Internet transaction. Since IP does not guarantee delivery, the computers on each end must maintain a "conversation" called a socket until all the packets for that transaction arrive undamaged and the transfer of data is complete.

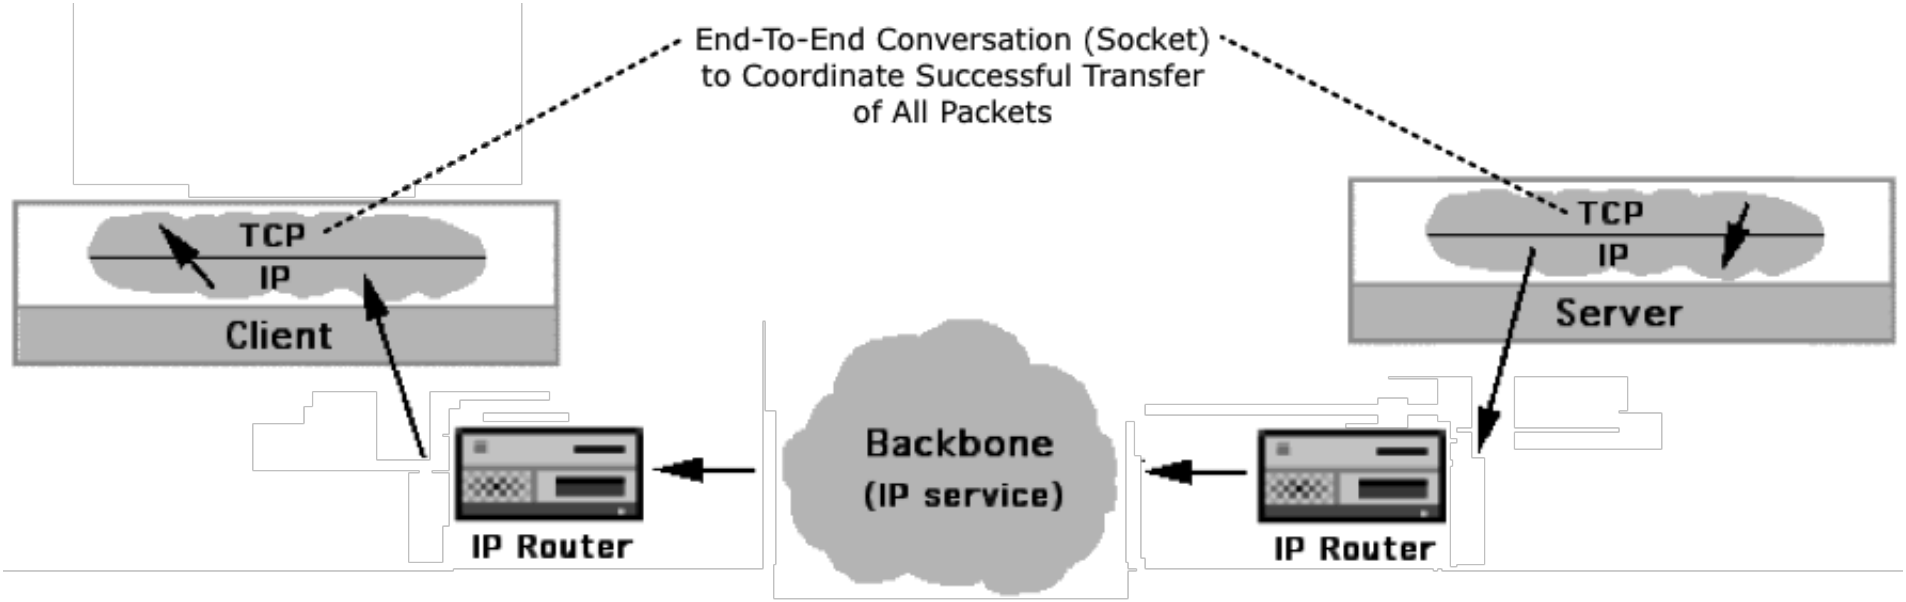

How TCP works:

On the Sending End, TCP does the following:

- Chops data into packets, each with a sequence number.
- Calculates a *checksum* for each packet and adds that to the packet. This is a count of the bits in the packet, used to test for data loss.
- Gives the packets the destination IP address.
- Gives the packets to the IP layer to start delivery process.

As the IP layer receives the packets on the other end (receiving end), it passes them to the TCP layer.

On the Receiving End, TCP does the following:

- Re-calculates the checksum for each packet and checks that against the actual data in the packet to see if it has been damaged.
- Makes requests back to the sending computer to resend a packet if it is damaged or is never received .
- Re-assembles the original data using the packet sequence numbers.

The combination of IP and TCP is referred to as the *TCP/IP Internet protocols*. The Internet is often characterized as a network of networks.

Inter-Networking (TCP/IP)
End-to-end Coordination (TCP) -- Two computers maintaining a conversation until the transaction is complete.
Inter-Networking (IP) -- Routing packets between networks. Backbones are mostly maintained by large phone companies.

The bottom two layers are usually lumped together as simply the network connection. Unless a computer can "talk" to another computer via networking, it is effectively isolated.

Networking
Networking protocols -- Two computers sharing data over Ethernet, Wi-Fi, Cellular, Cable Modem, Dish, Phone Modem, etc.
Physical – Different networking technologies use copper, air, fiber-optic..

Application Layer

• How useful is the Internet to humans? That depends upon how useful the internet-capable software applications (apps) are.

• Very early on, internet-capable software was developed for:

Remote Login (telnet/ssh)
Robust Remote Messaging (email)
File Transfer Operations (FTP)

The World Wide Web was a clever software invention (killer app) that helped revolutionize how humans could use the Internet infrastructure.

For many years, most packets on the Internet were generated by email software. Packets generated by WWW software quickly changed that!

World Wide Web Invented around 1990.

- The internet is over 20 years old.

- Tim Berners-Lee invents WWW while working at the CERN nuclear physics research lab (huge particle accelerator) in Geneva Switzerland.

- Initial goal was to enable physicists to share abstracts of physics research papers over the Internet as hypertext documents – documents with hyperlinks to other documents.

- Berners-Lee's larger goal was literally to create a Web of inter-connected information of World-Wide scope.

- He actually named his new software WWW, a term that collectively referred to the first Web browser and Web server software which worked together to deliver hypertext documents (web pages) over the Internet.

Netscape Navigator Browser

• 1993 -- Marc Andreessen, a Computer Science graduate student at the University of Illinois, creates *Mosaic* Web browser, the first Web browser that could render graphics.

• 1994 -- Andreessen and friends form *Netscape Communications Corporation*, originally a privately held company.

• 1995 -- Netscape corporation goes public, even though the company was not profitable and didn't have a solid business model. Stock price triples on first day, raising almost 3 billion dollars. This unprecedented IPO from a non-profitable company is a major news story.

• This is when the world starts be aware of the WWW. At that point, most people had not even heard of the Internet, so many people use the two terms interchangeably, still to this day.

Netscape Navigator 2 (NN2)

• Released in 1995, this was one of the most influential web browsers in history because many people first surfed the web with this browser, and because it introduced many important new features.

• Fill-in data forms to collect information in Web pages.

• Cookies so that the browser could store a small piece of data between separate page loads.

• JavaScript to increase interactive capabilities.

• Netscape company even created the secure sockets extension to the TCP protocol that allowed for packets to be encrypted during transit. Now that credit card information could be securely collected in web pages, companies formed like Amazon, originally only an online book store.

Browser Wars

• Microsoft releases Internet Explorer Browser (IE3) in 1996, the only real competitor to NN3, also just released.

• Microsoft uses it's dominance as a desktop operating system (90%+ market share) to try to kill Netscape. Made it very difficult to use Netscape on Windows OS.

• Microsoft hit with several anti-trust lawsuits and eventually starts playing nice with Netscape, and also releasing Office software suite for Mac OS (Apple). Browser wars mostly over by late 1990s.

• A terrible side effect of the browser wars was that the HTML language that creates web pages was diverging. By the time of IE4 and NN4, the HTML language features supported by each browser were becoming increasingly different. People were starting to build separate web pages to get the most out of each browser – "best viewed in Netscape or IE".

World Wide Web Consortium (W3C)

• Helped to end the browser wars by releasing uniform standards for HTML and other web-related technologies.

• If all browsers implement the W3C recommendations, then all browsers will work mostly the same way.

• W3C is non-profit organization formed by Tim-Berners Lee and other people that helped create the WWW. Major corporations like Microsoft and Netscape also get a seat at the table because their input is important.

• W3C still maintains the uniform standards today. That's why popular browsers such as Chrome, Safari, IE (now Edge), Firefox (the legacy of NN), and Opera all more or less work uniformly so web pages will (mostly) render the same way on all browsers.

Virtual Domain (Domain Name)

• A domain name like *craigknuckles.com* is just virtual property that you can buy. Then you basically own the name. Other than that, it's nothing.

• The Internet uses IP addresses exclusively, so a virtual domains are by themselves not viable addresses.

Fully Qualified Doman Name (FQDN)

• To make an FQDN, you assign a prefix such as *www* and associate the named address with an IP address.

www. craigknuckles.com <------- DNS Mapping -------> 164.68.21.170

• It's then called an FQDN or *named address* -- a human friendly address that points to an actual IP address (of a web server, for example).

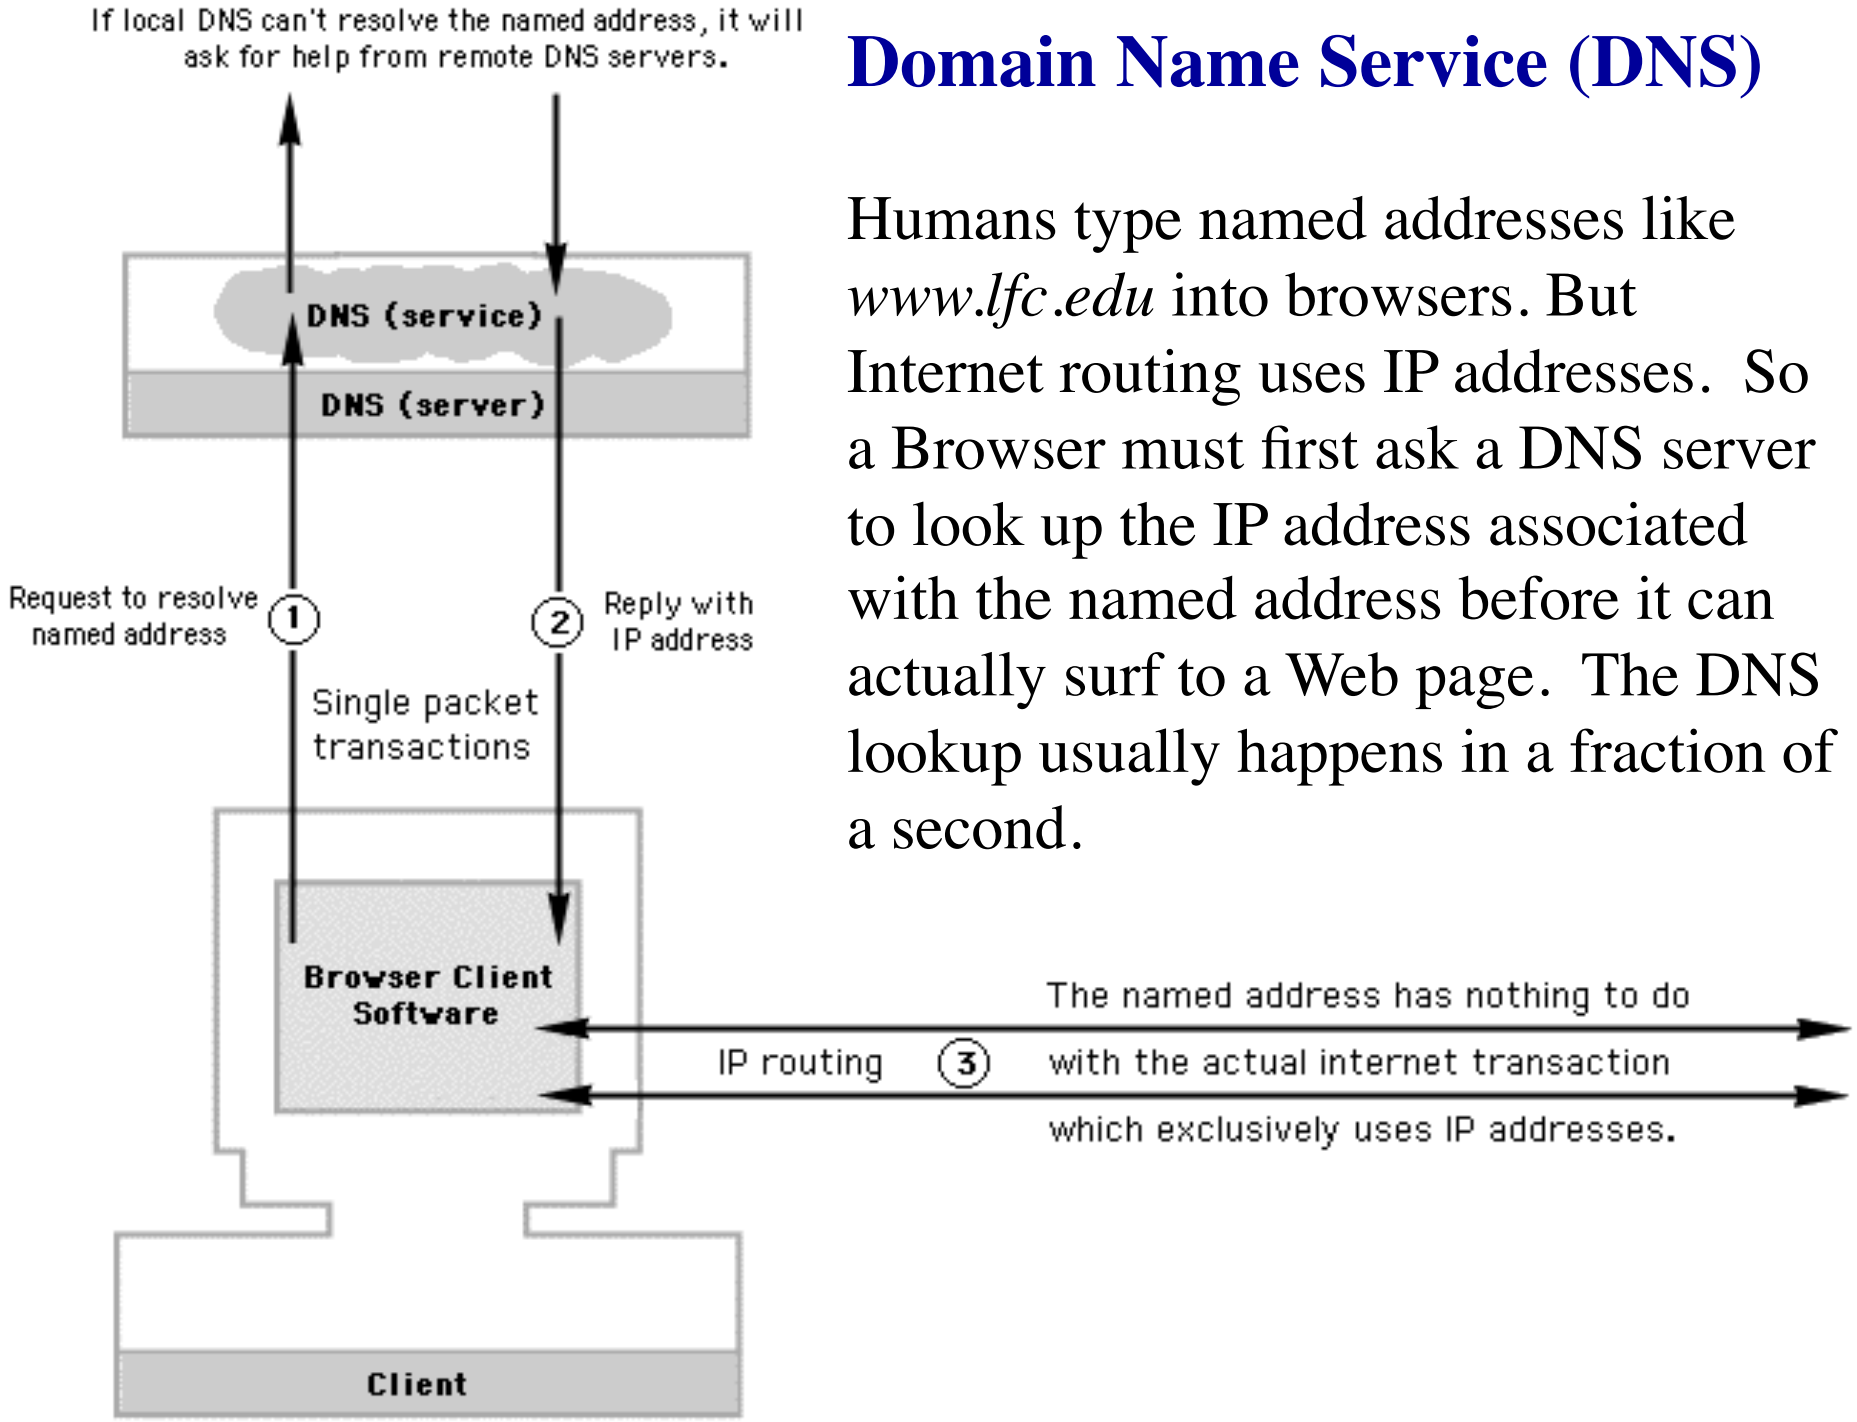

Domain Name Service (DNS)

Humans type named addresses like *www.lfc.edu* into browsers. But Internet routing uses IP addresses. So a Browser must first ask a DNS server to look up the IP address associated with the named address before it can actually surf to a Web page. The DNS lookup usually happens in a fraction of a second.

Top-level domains such as *.com* are at the top of *DNS hierarchy*.

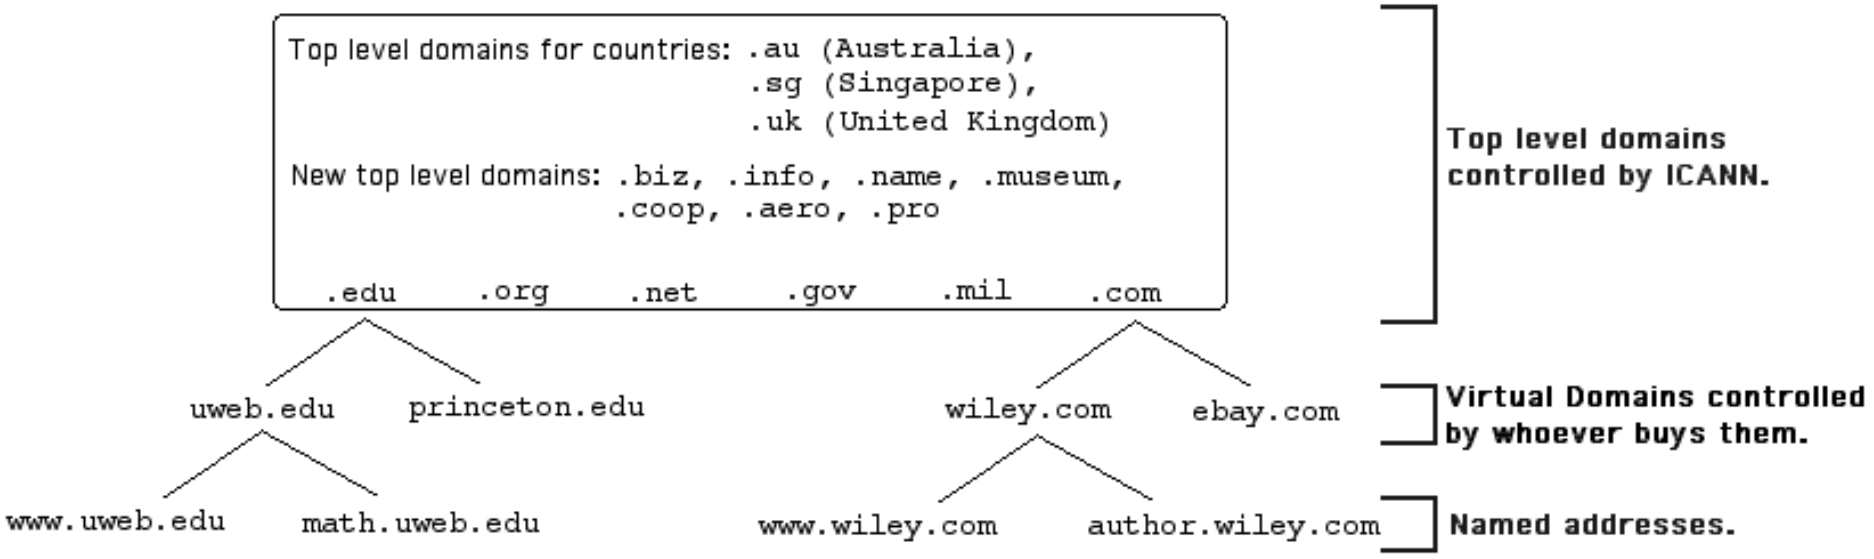

There's no way one DNS server can know all the IP addresses associated with named addresses. So for example, the local DNS server might need to ask a top-level DNS server (such as .com or .edu) where to locate another lower-level DNS server that might hold the sought after FQDN record. Certain DNS servers specialize in certain domains, so this is a division of labor where requests are sent up and down the hierarchy to find the pertinent DNS server.

Virtual Hosting refers to hosting multiple named addresses (Web sites) on the same server. The DNS records for each named address point to the same IP address.

Requests for different named addresses are sent to different folders by the *Web Server Software*, the software that answers the requests from the Web browser software on the client side.

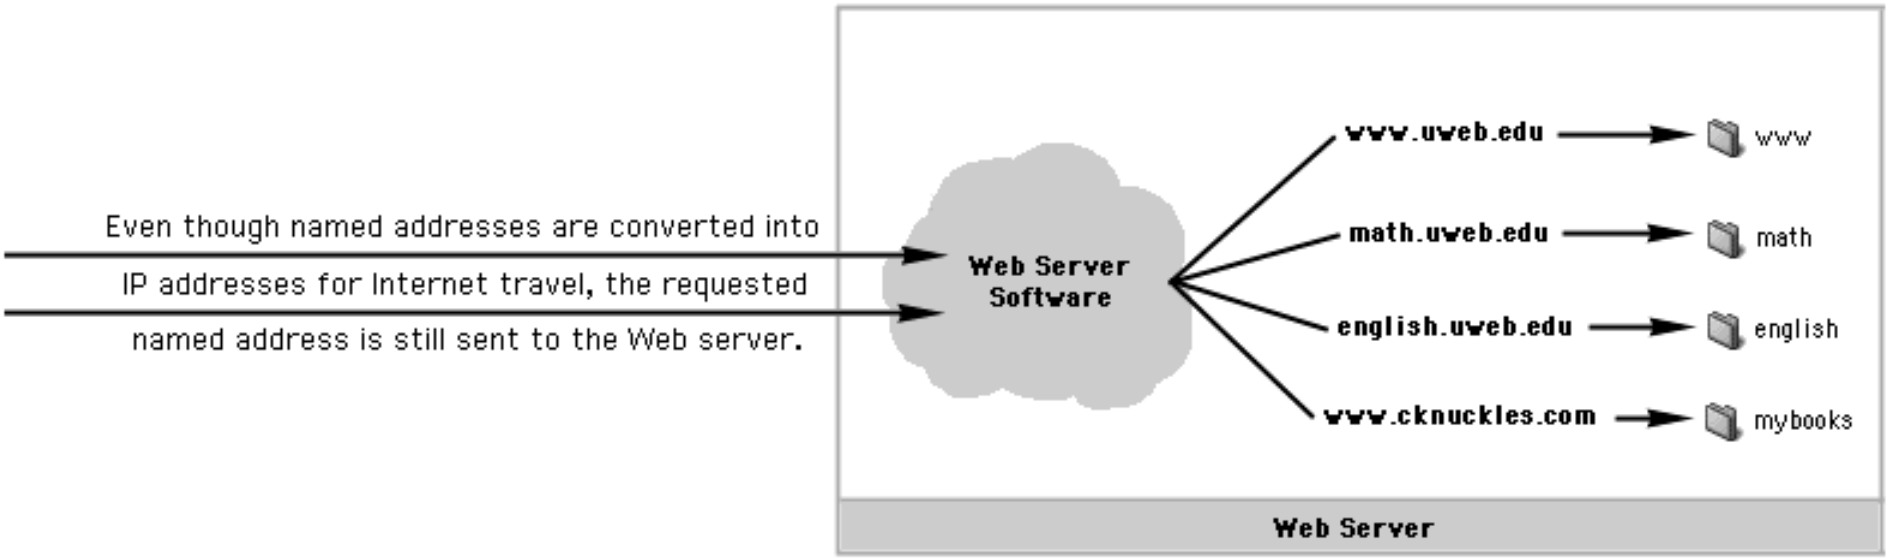

URL -- Uniform Resource Locator

You have seen URLs like *http://www.cknuckles.com*
A URL has three primary components

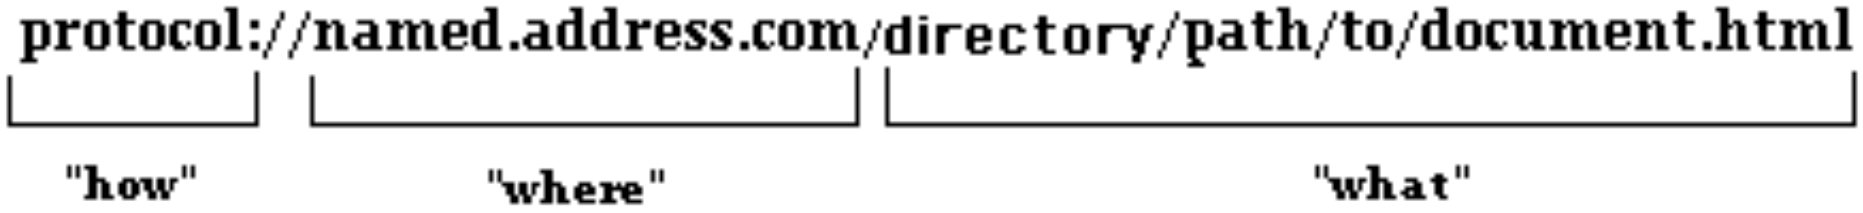

The *how* is a protocol like *http* (hypertext transfer protocol), or the secure version *https*, that specify how web pages are delivered.

Other protocols such as *ftp* in the how part of a URL initiate a complete file download, rather than temporary web page load.

The *where* part is an address, which is usually a named address, but can be a raw numeric IP address (which bypasses the DNS lookup).

The *what* part specifies a specific resource on the server.

1. http://www.uweb.edu/
2. http://www.uweb.edu/academics.html
3. http://www.uweb.edu/sports/
4. http://www.uweb.edu/sports/volleyball/
5. http://www.uweb.edu/sports/football/stats.html

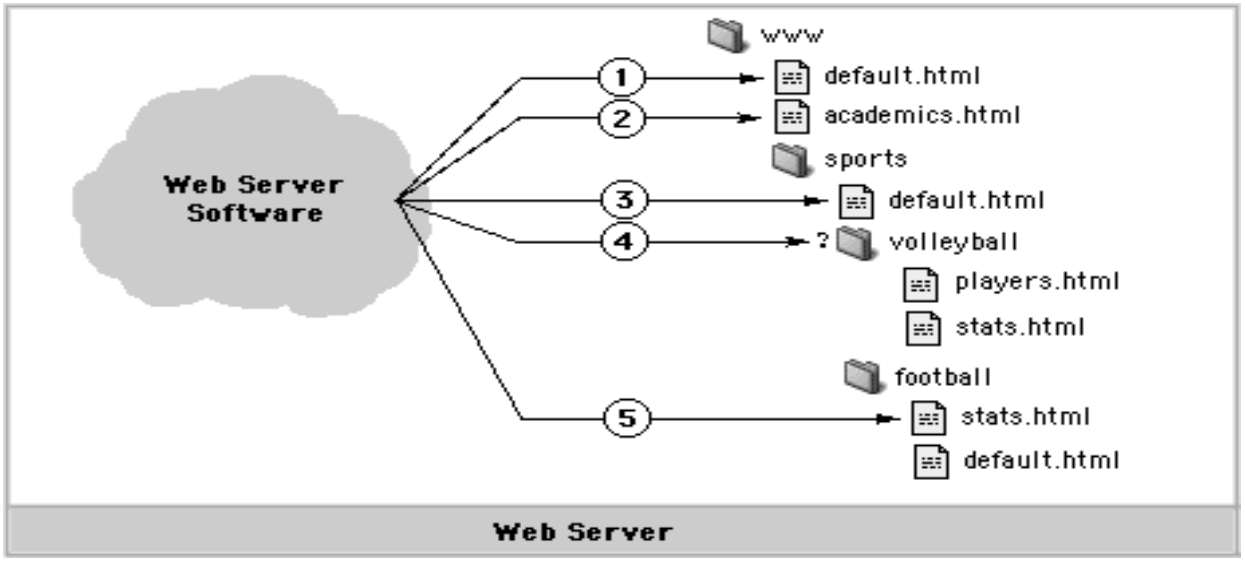

Directory Path

The *what* part of a URL is a path that descends into subfolders within the Web site.

Some URLs (2 and 5 above) request a specific page. But 1,3,4 actually request folders. If there is a *default file* present in the folder, that web page will be served to represent the folder. In URL 4, there is no default file in the requested folder. In that case, the server might give an "access denied" message, or serve a clickable listing of all files in the directory.