Computer Architecture and Operating Systems
Lecture 12: Basics of Networking

Andrei Tatarnikov

atatarnikov@hse.ru
@andrewt0301

Client-Server Architecture

- Most network applications are based on the client-server model:
  - A ***server*** process and one or more ***client*** processes
  - Server manages some ***resource***
  - Server provides ***service*** by manipulating resource for clients
  - Server activated by request from client (vending machine analogy)

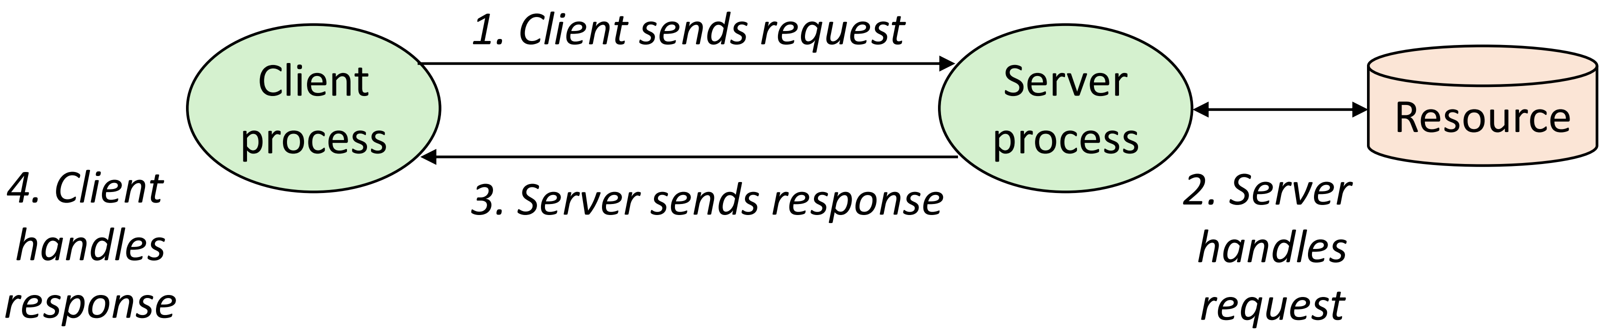

Note: clients and servers are processes running on hosts (can be the same or different hosts)

Hardware Organization of a Network Host

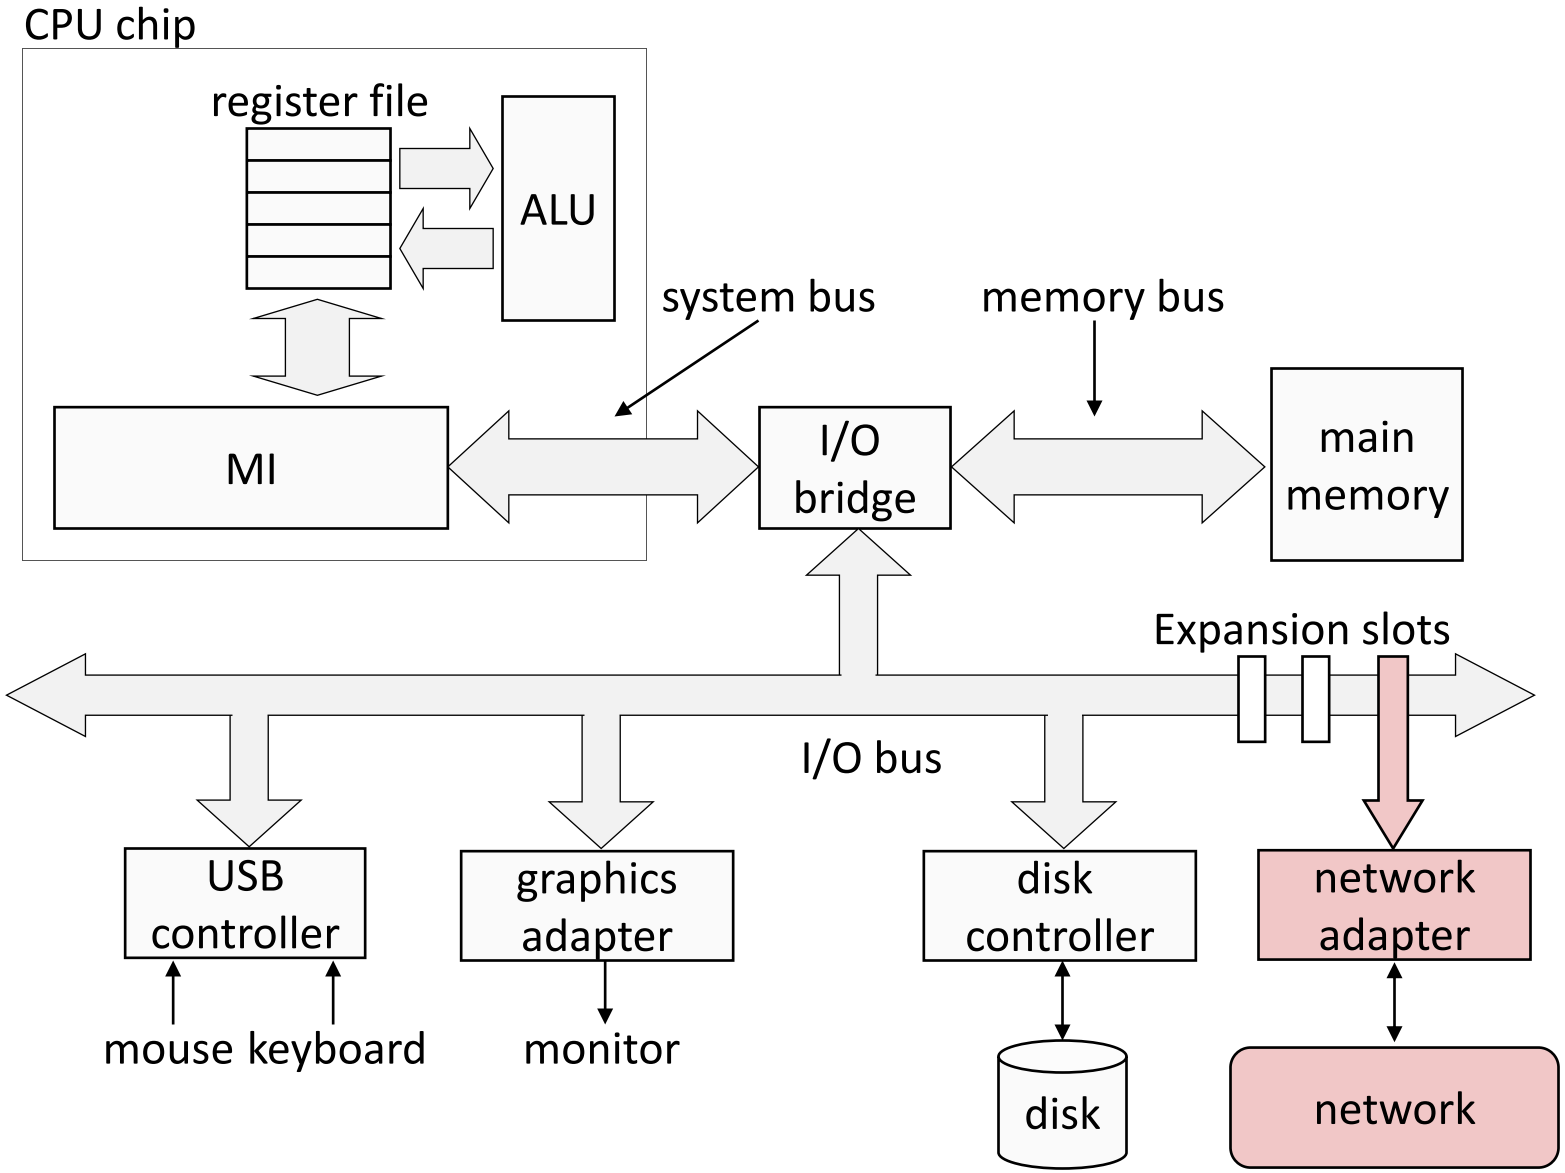

Computer Networks

- A ***network*** is a hierarchical system of boxes and wires organized by geographical proximity
  - SAN (System Area Network) spans cluster or machine room
    - Switched Ethernet, Quadrics QSW, …
  - LAN (Local Area Network) spans a building or campus
    - Ethernet is most prominent example
  - WAN (Wide Area Network) spans country or world
    - Typically high-speed point-to-point phone lines
- An ***internetwork (internet)*** is an interconnected set of networks
  - The Global IP Internet (**uppercase** "I") is the most famous example of an internet (**lowercase** "i")
- Let us see how an internet is built from the ground up

Lowest Level: Ethernet Segment

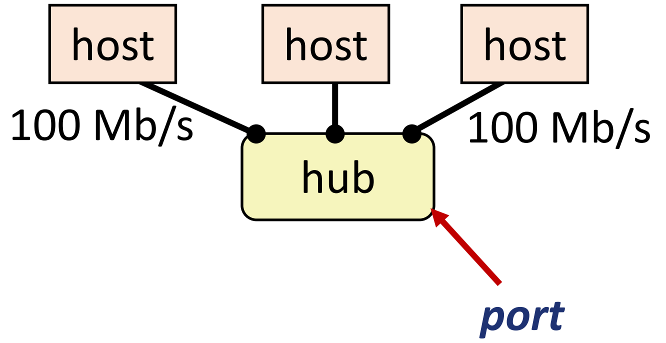

- Ethernet segment consists of a collection of ***hosts*** connected by wires (twisted pairs) to a ***hub***
- Spans room or floor in a building
- Operation
  - Each Ethernet adapter has a unique 48-bit address (MAC address)
    - E.g., 00:16:ea:e3:54:e6
  - Hosts send bits to any other host in chunks called ***frames***
  - Hub slavishly copies each bit from each port to every other port
    - Every host sees every bit
    - Note: Hubs are on their way out. Bridges (switches, routers) became cheap enough to replace them

Next Level: Bridged Ethernet Segment

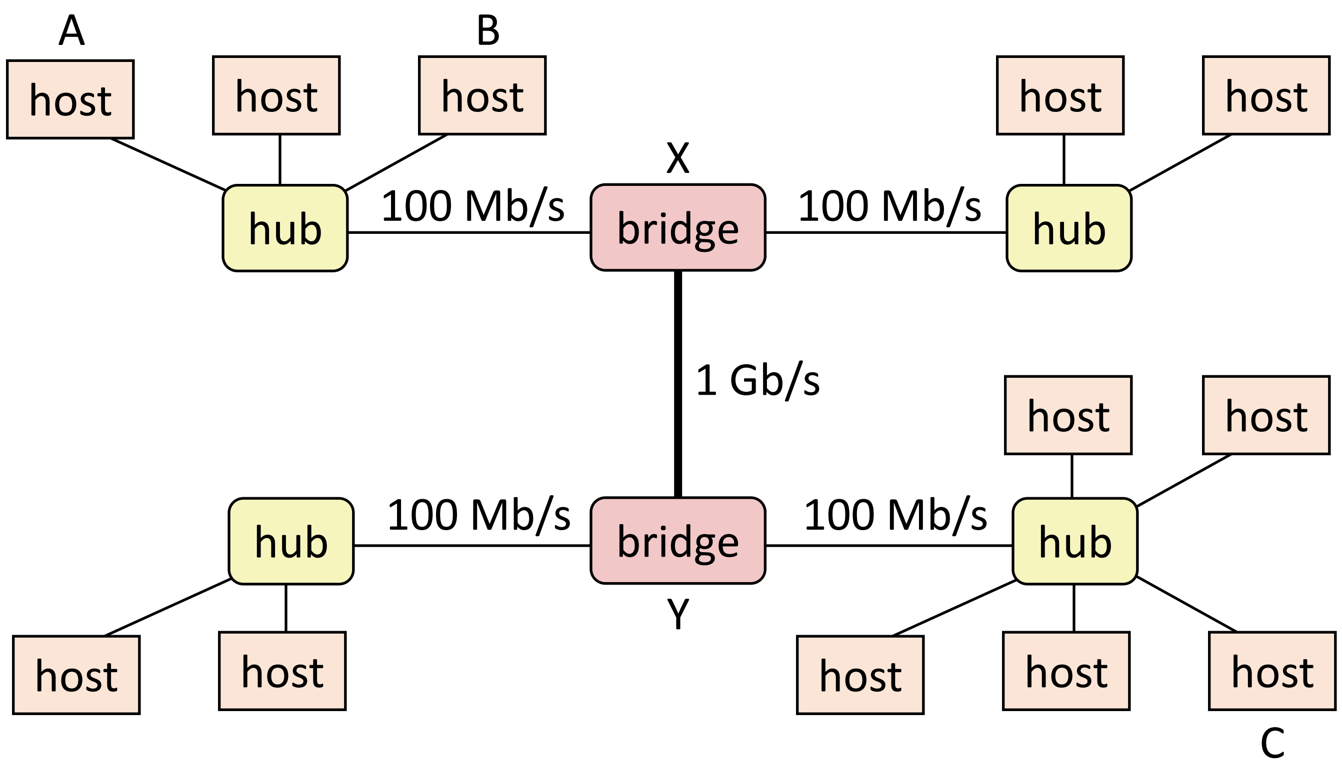

- Spans building or campus
- Bridges cleverly learn which hosts are reachable from which ports and then selectively copy frames from port to port

Conceptual View of LANs

- For simplicity, hubs, bridges, and wires are often shown as a collection of hosts attached to a single wire:

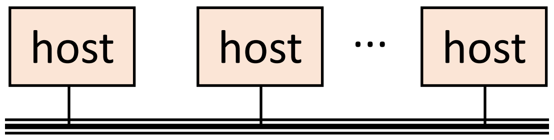

Next Level: Internets

- Multiple incompatible LANs can be physically connected by specialized computers called ***routers***
- The connected networks are called an ***internet*** (lower case)

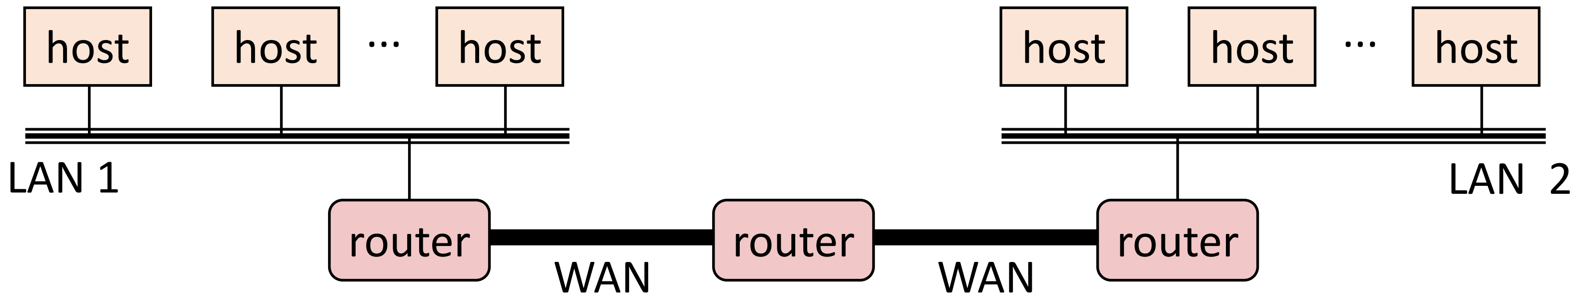

LAN 1 and LAN 2 might be completely different, totally incompatible

(e.g., Ethernet, Fibre Channel, 802.11, T1-links, DSL, …)*

Logical Structure of Internet

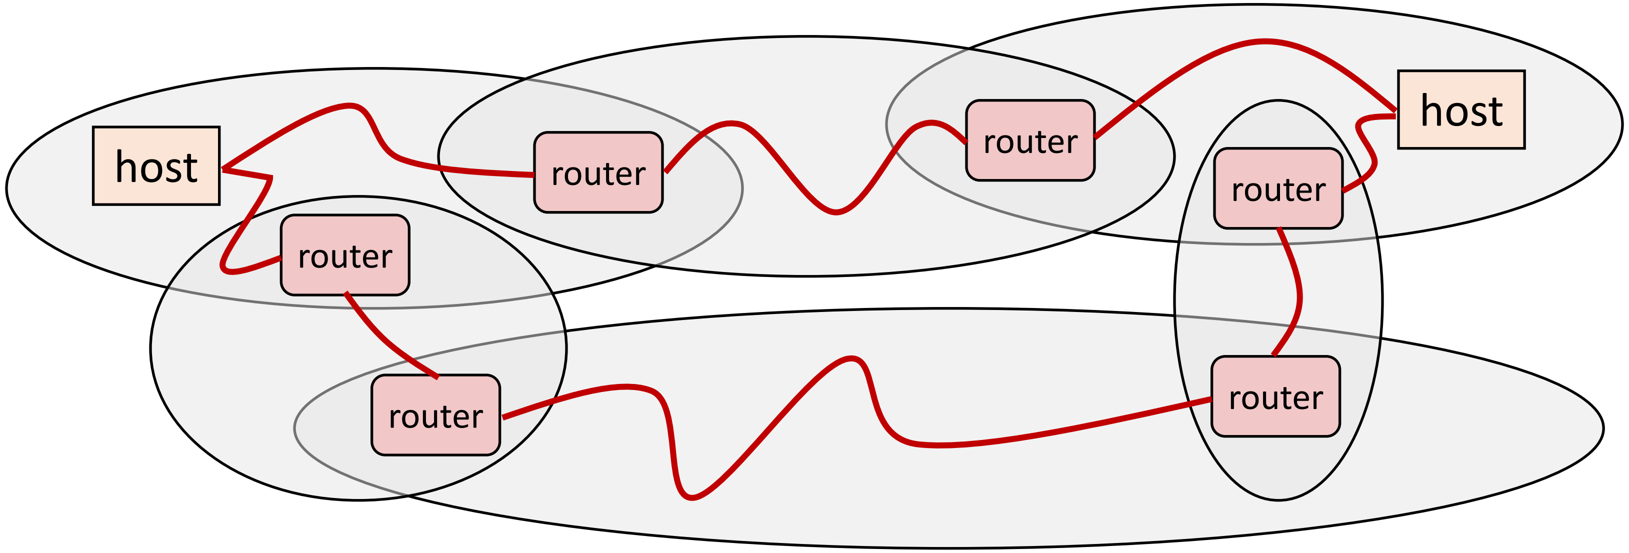

- Ad hoc interconnection of networks
  - No particular topology
  - Vastly different router & link capacities
- Send packets from source to destination by hopping through networks
  - Router forms bridge from one network to another
  - Different packets may take different routes

The Notion of an Internet Protocol

- How is it possible to send bits across incompatible LANs and WANs?

- Solution: ***protocol*** software running on each host and router
  - Protocol is a set of rules that governs how hosts and routers should cooperate when they transfer data from network to network.
  - Smooths out the differences between the different networks

What Does an internet Protocol Do?

- Provides a *naming scheme*
  - An internet protocol defines a uniform format for ***host addresses***
  - Each host (and router) is assigned at least one of these internet addresses that uniquely identifies it

- Provides a *delivery mechanism*
  - An internet protocol defines a standard transfer unit (***packet***)
  - Packet consists of ***header*** and ***payload***
    - Header: contains info such as packet size, source and destination addresses
    - Payload: contains data bits sent from source host

Transferring Internet Data Via Encapsulation

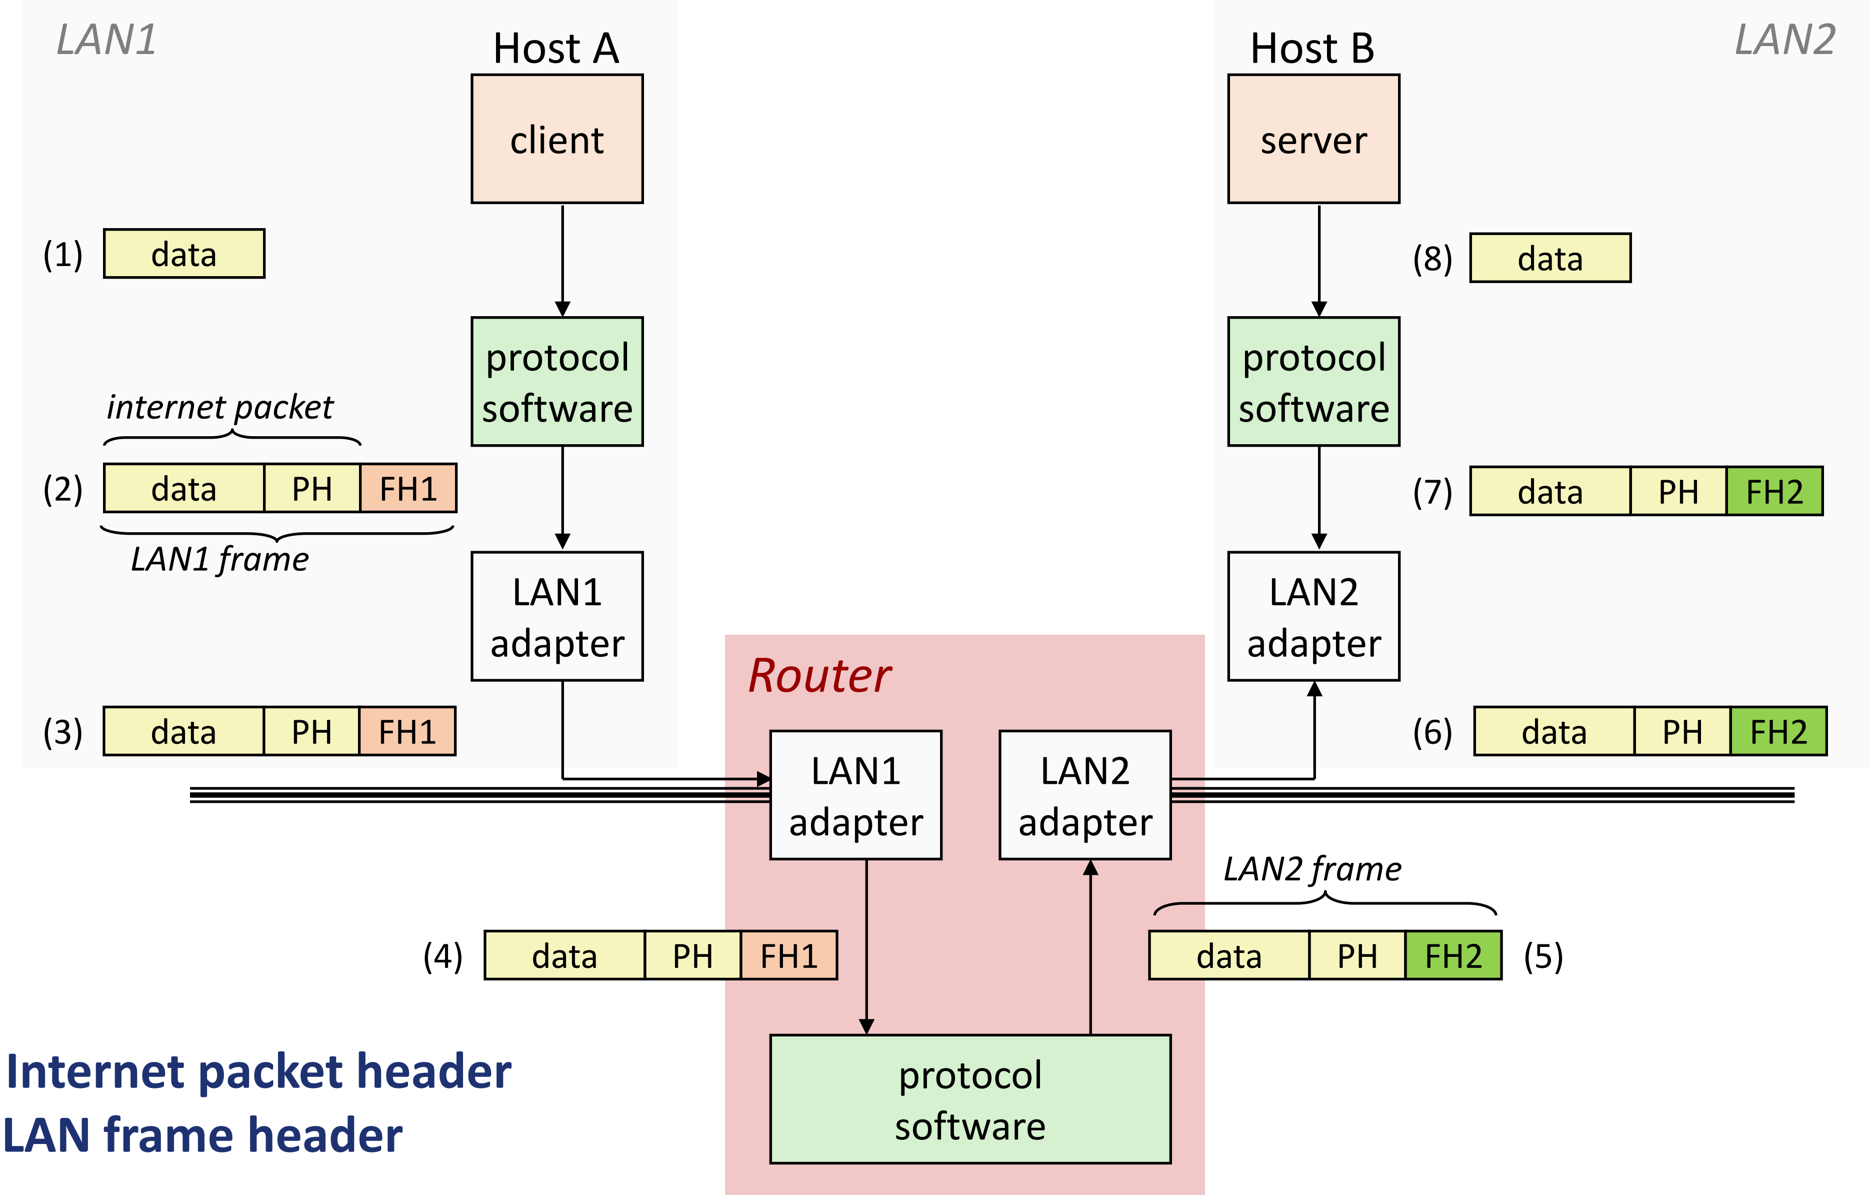

PH: Internet packet header
FH: LAN frame header

Other Issues

- We are glossing over a number of important questions:
  - What if different networks have different maximum frame sizes? (segmentation)
  - How do routers know where to forward frames?
  - How are routers informed when the network topology changes?
  - What if packets get lost?

- These (and other) questions are addressed by the area of systems known as ***computer networking***

Global IP Internet

- Most famous example of an Internet
- Based on the TCP/IP protocol family
  - IP (Internet Protocol) :
    - Provides ***basic naming scheme*** and unreliable ***delivery capability*** of packets (datagrams) from ***host-to-host***
  - UDP (Unreliable Datagram Protocol)
    - Uses IP to provide ***unreliable*** datagram delivery from ***process-to-process***
  - TCP (Transmission Control Protocol)
    - Uses IP to provide ***reliable*** byte streams from ***process-to-process*** over ***connections***
- Accessed via a mix of Unix file I/O and functions from the *sockets interface*

Organization of an Internet Application

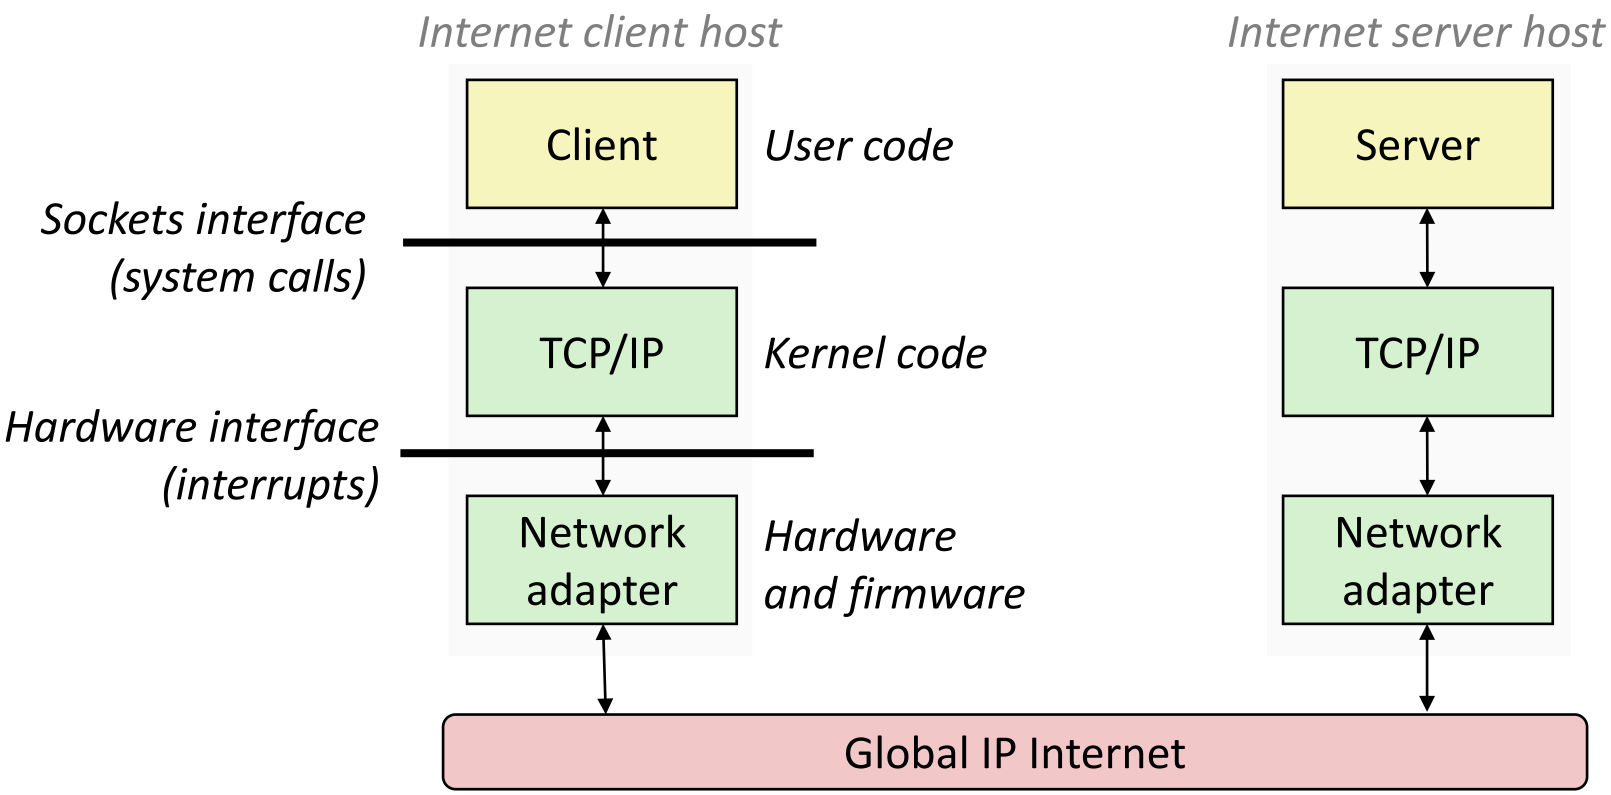

A Programmer's View of the Internet

1. Hosts are mapped to a set of 32-bit *IP addresses*
   - 128.2.203.179

2. The set of IP addresses is mapped to a set of identifiers called Internet *domain names*
   - 128.2.203.179 is mapped to www.cs.cmu.edu

3. A process on one Internet host can communicate with a process on another Internet host over a *connection*

Aside: IPv4 and IPv6

- The original Internet Protocol, with its 32-bit addresses, is known as *Internet Protocol Version 4* (**IPv4**)

- 1996: Internet Engineering Task Force (IETF) introduced *Internet Protocol Version 6* (**IPv6**) with 128-bit addresses
  - Intended as the successor to IPv4

- As of 2015, vast majority of Internet traffic still carried by IPv4
  - Only 4% of users access Google services using IPv6.

- We will focus on IPv4, but will show you how to write networking code that is protocol-independent.

IP Addresses

- 32-bit IP addresses are stored in an *IP address struct*
  - IP addresses are always stored in memory in *network byte order* (big-endian byte order)
  - True in general for any integer transferred in a packet header from one machine to another.
    - E.g., the port number used to identify an Internet connection.

```
/* Internet address structure */
struct in_addr {
    uint32_t  s_addr; /* network byte order (big-endian) */
};
```

Dotted Decimal Notation

- By convention, each byte in a 32-bit IP address is represented by its decimal value and separated by a period
  - IP address: `0x8002C2F2 = 128.2.194.242`

- Use `getaddrinfo` and `getnameinfo` functions (described later) to convert between IP addresses and dotted decimal format.

Internet Domain Names

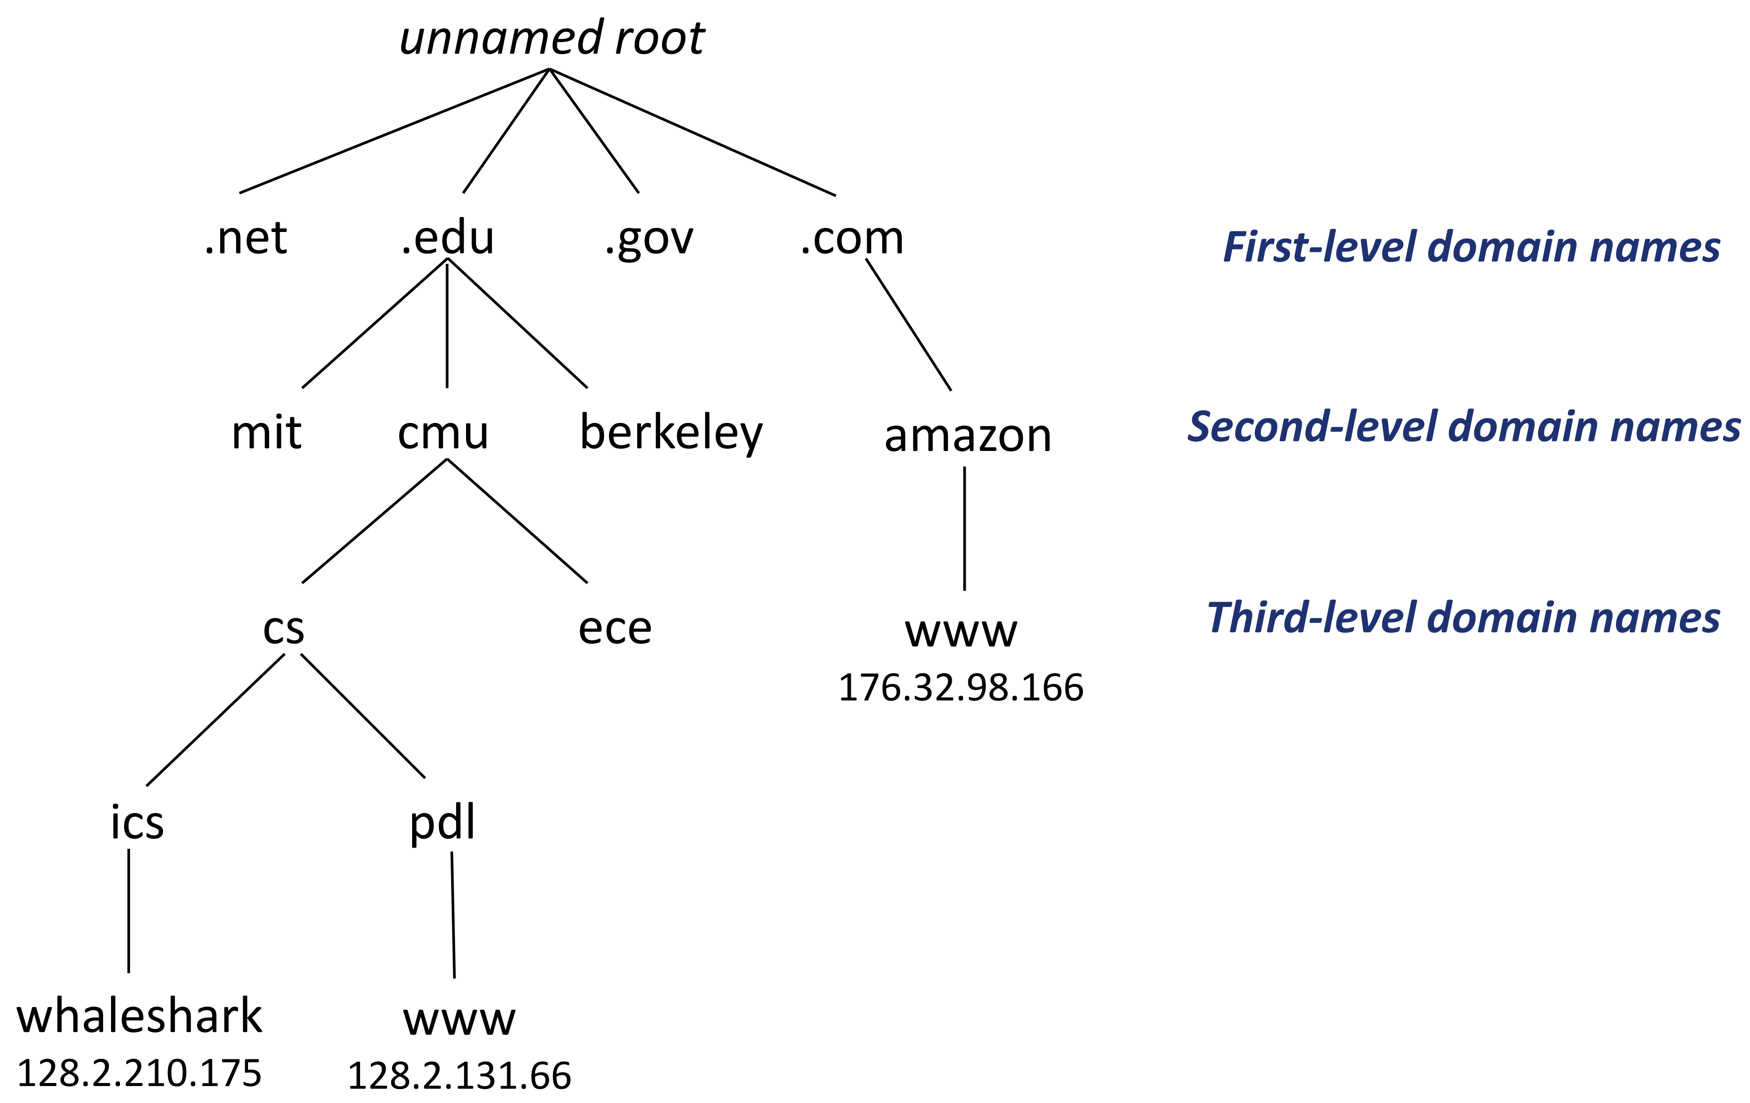

Domain Naming System (DNS)

- The Internet maintains a mapping between IP addresses and domain names in a huge worldwide distributed database called *DNS*

- Conceptually, programmers can view the DNS database as a collection of millions of *host entries.*
  - Each host entry defines the mapping between a set of domain names and IP addresses.
  - In a mathematical sense, a host entry is an equivalence class of domain names and IP addresses.

Properties of DNS Mappings

- Can explore properties of DNS mappings using `nslookup`
  - Output edited for brevity

```
linux> nslookup localhost
Address: 127.0.0.1
```

- Each host has a locally defined domain name `localhost` which always maps to the *loopback address* `127.0.0.1`

```
linux> hostname
whaleshark.ics.cs.cmu.edu
```

- Use `hostname` to determine real domain name of local host:

Properties of DNS Mappings

- Simple case: one-to-one mapping between domain name and IP address:

```
linux> nslookup whaleshark.ics.cs.cmu.edu
Address: 128.2.210.175
```

- Multiple domain names mapped to the same IP address:

```
linux> nslookup cs.mit.edu
Address: 18.62.1.6
linux> nslookup eecs.mit.edu
Address: 18.62.1.6
```

Properties of DNS Mappings

- Multiple domain names mapped to multiple IP addresses:

```
linux> nslookup www.twitter.com
Address: 199.16.156.6
Address: 199.16.156.70
Address: 199.16.156.102
Address: 199.16.156.230

linux> nslookup twitter.com
Address: 199.16.156.102
Address: 199.16.156.230
Address: 199.16.156.6
Address: 199.16.156.70
```

- Some valid domain names don’t map to any IP address:

```
linux> nslookup ics.cs.cmu.edu
*** Can't find ics.cs.cmu.edu: No answer
```

Internet Connections

- Clients and servers communicate by sending streams of bytes over *connections*. Each connection is:
  - *Point-to-point*: connects a pair of processes.
  - *Full-duplex*: data can flow in both directions at the same time,
  - *Reliable*: stream of bytes sent by the source is eventually received by the destination in the same order it was sent.

- *A socket* is an endpoint of a connection
  - *Socket address* is an **`IPaddress:port`** pair

- A *port* is a 16-bit integer that identifies a process:
  - ***Ephemeral port***: Assigned automatically by client kernel when client makes a connection request.
  - ***Well-known port:*** Associated with some *service* provided by a server (e.g., port 80 is associated with Web servers)

Well-known Ports and Service Names

- Popular services have permanently assigned *well-known ports and* corresponding *well-known service names*:
  - echo server: 7/echo
  - ssh servers: 22/ssh
  - email server: 25/smtp
  - Web servers: 80/http

- Mappings between well-known ports and service names is contained in the file `/etc/services` on each Linux machine.

Anatomy of a Connection

- A connection is uniquely identified by the socket addresses of its endpoints (*socket pair*)
  - **`(cliaddr:cliport, servaddr:servport)`**

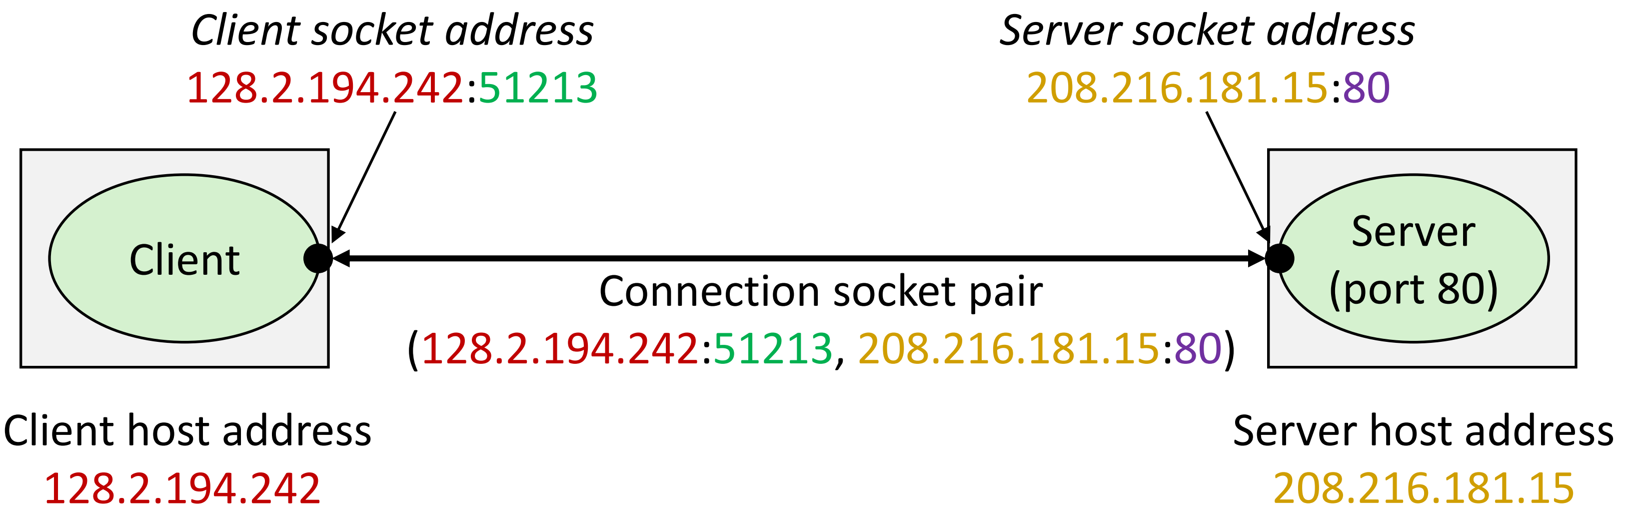

51213 is an ephemeral port allocated by the kernel

80 is a well-known port associated with Web servers

Using Ports to Identify Services

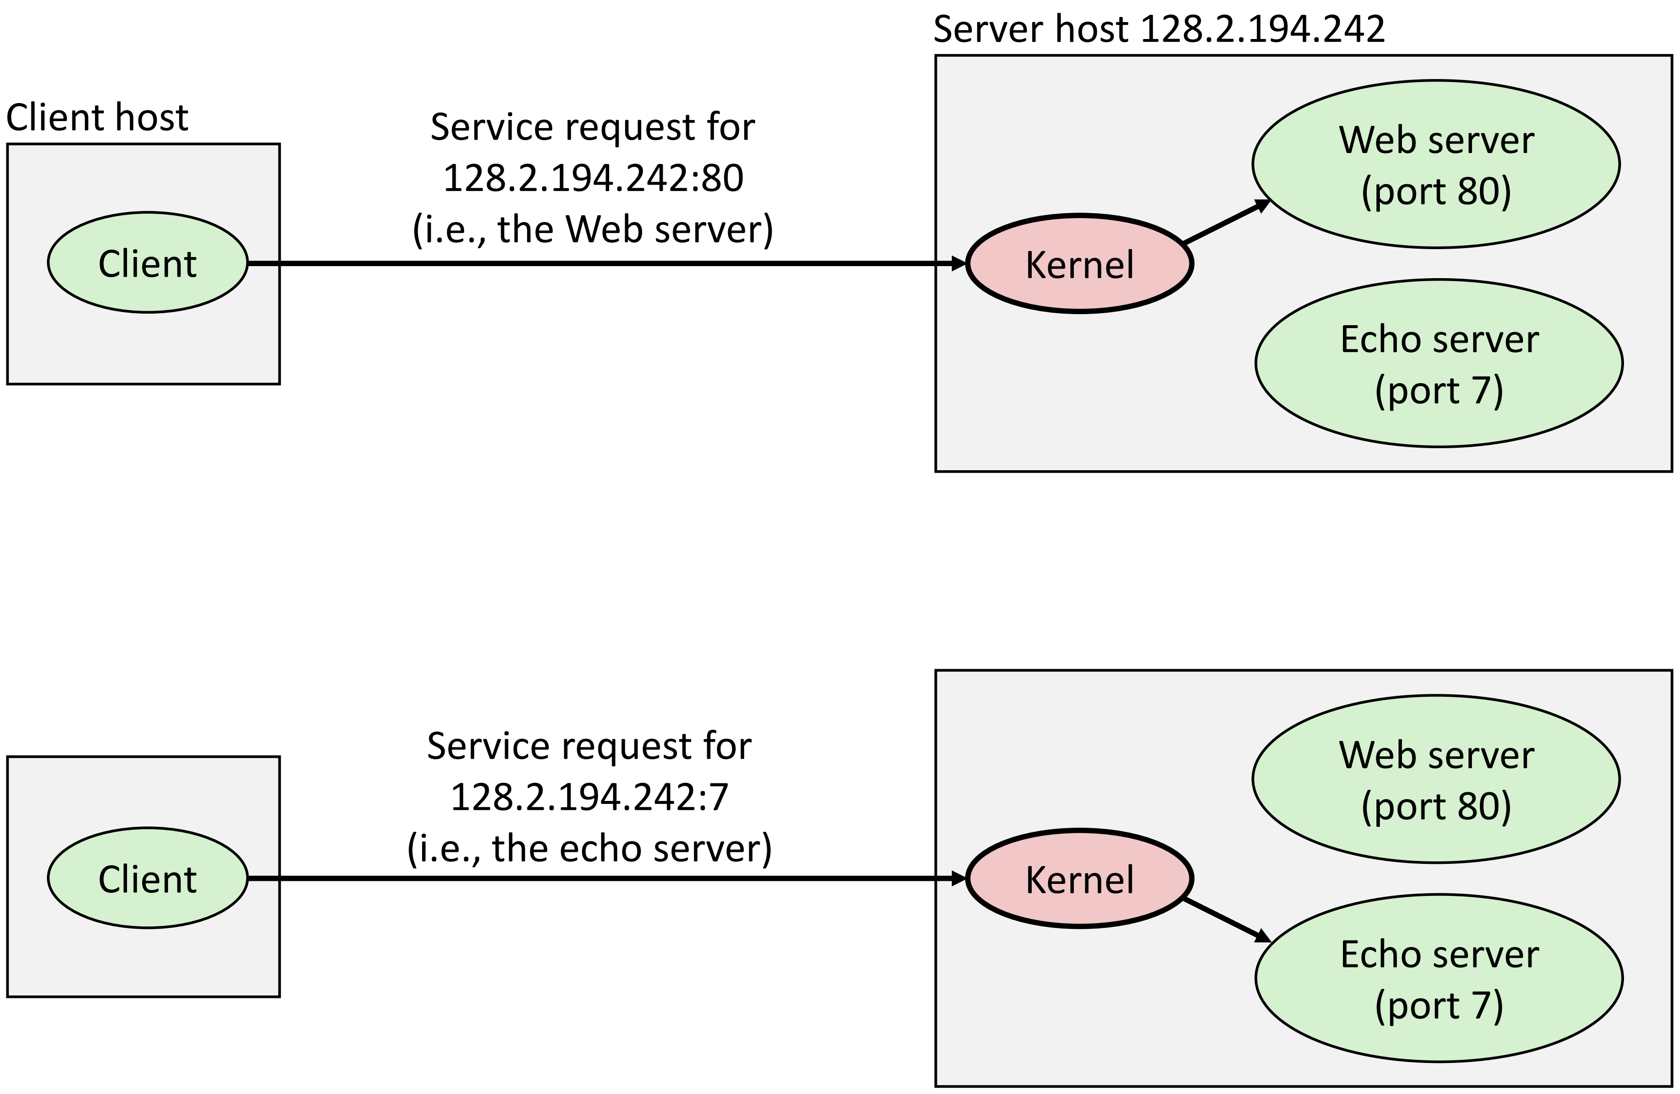

Basic Internet Components

- Internet backbone:
  - collection of routers (nationwide or worldwide) connected by high-speed point-to-point networks
- Internet Exchange Points (IXP):
  - router that connects multiple backbones (often referred to as peers)
  - Also called Network Access Points (NAP)
- Regional networks:
  - smaller backbones that cover smaller geographical areas (e.g., cities or states)
- Point of presence (POP):
  - machine that is connected to the Internet
- Internet Service Providers (ISPs):
  - provide dial-up or direct access to POPs

Internet Connection Hierarchy

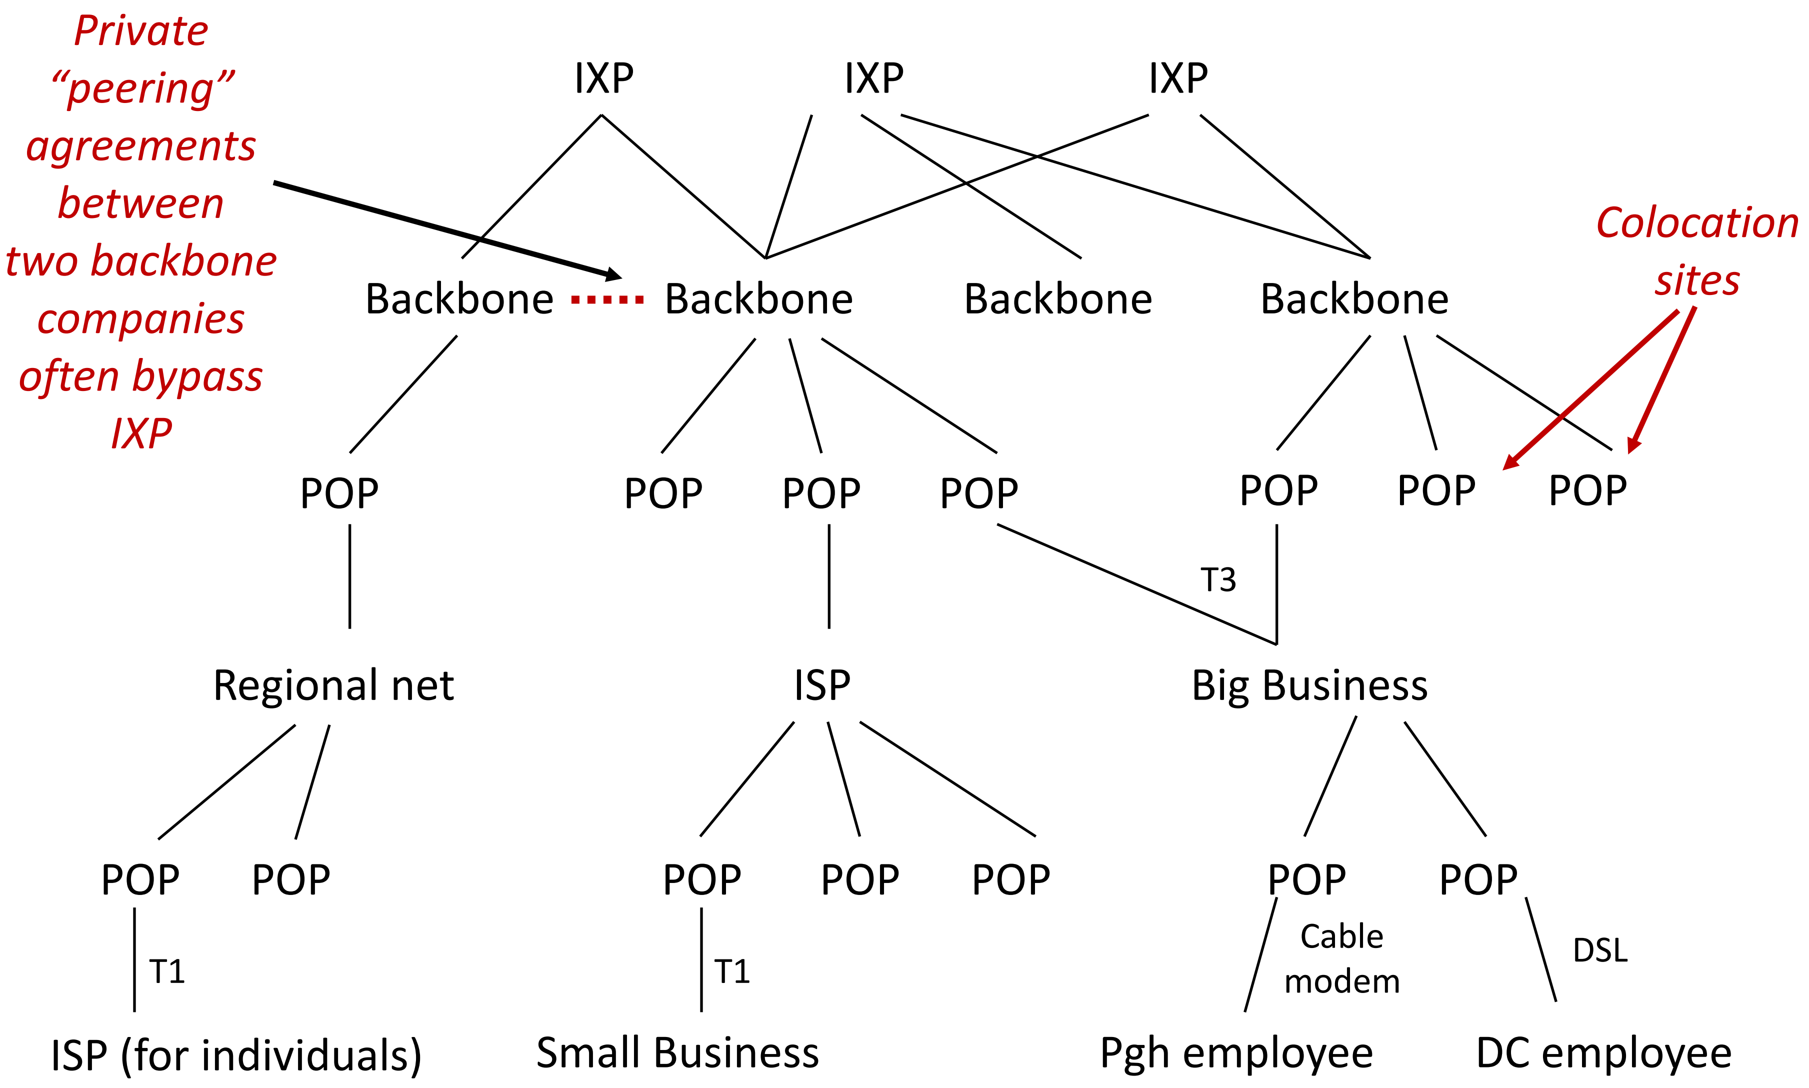

IP Address Structure

- IP (V4) Address space divided into classes:

	0	1	2	3	8	16	24	31
Class A	0	Net ID				Host ID		
Class B	1	0	Net ID				Host ID	
Class C	1	1	0	Net ID				Host ID
Class D	1	1	1	0	Multicast address			
Class E	1	1	1	1	Reserved for experiments			

- Network ID Written in form w.x.y.z/n
  - n = number of bits in host address
  - E.g., CMU written as 128.2.0.0/16
    - Class B address
- Unrouted (private) IP addresses:
  10.0.0.0/8   172.16.0.0/12   192.168.0.0/16

Any Questions?

```
            .text
__start:    addi t1, zero, 0x18
            addi t2, zero, 0x21
cycle:      beq t1, t2, done
            slt t0, t1, t2
            bne t0, zero, if_less
            nop
            sub t1, t1, t2
            j cycle
            nop
if_less:    sub t2, t2, t1
            j cycle
done:       add t3, t1, zero
```