Network layer

Chapter 4 in the textbook

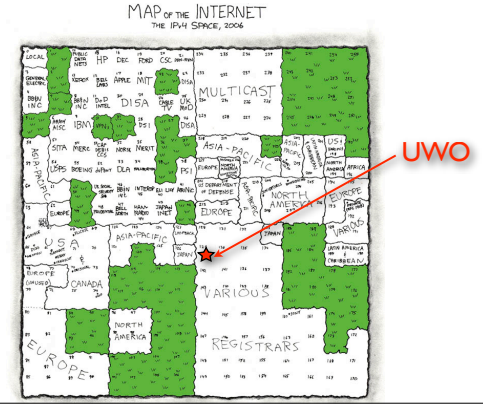

Assignment 3

- Abstractly, your server is essentially a simple router
  - Maybe more of a switch than a router
- Your server gets messages ("packets")
- These messages have target IDs ("addresses")
- You have to send ("forward") the message ("packet") to the right ID ("address")

Nickname/ID	File descriptor
<listening>	8

Nickname/ID	File descriptor
<listening>	8

Nickname/ID	File descriptor
<listening>	8
Joe	13

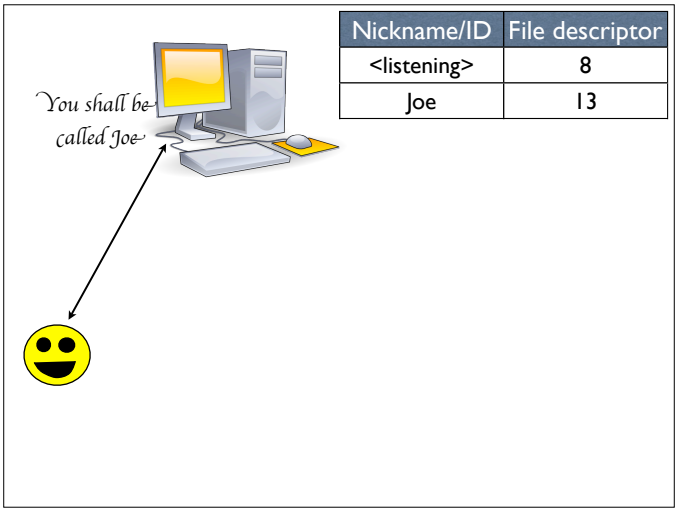
You shall be called Joe
Nickname/ID
File descriptor
<listening>
8
Joe
13

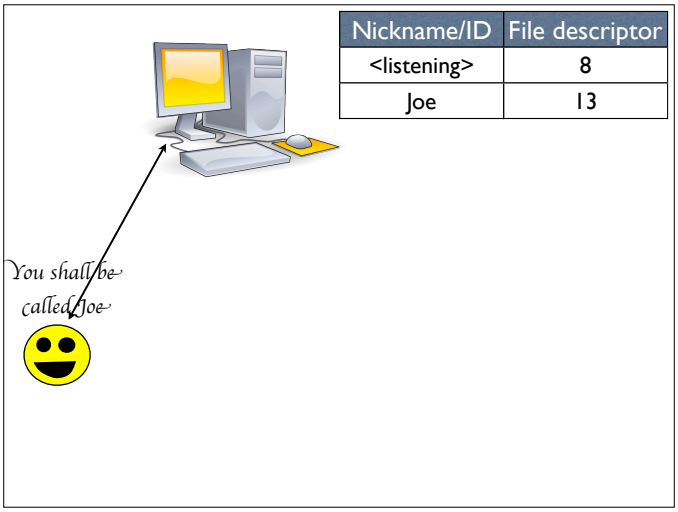
You shall be called Joe
Nickname/ID
File descriptor
<listening>
8
Joe
13

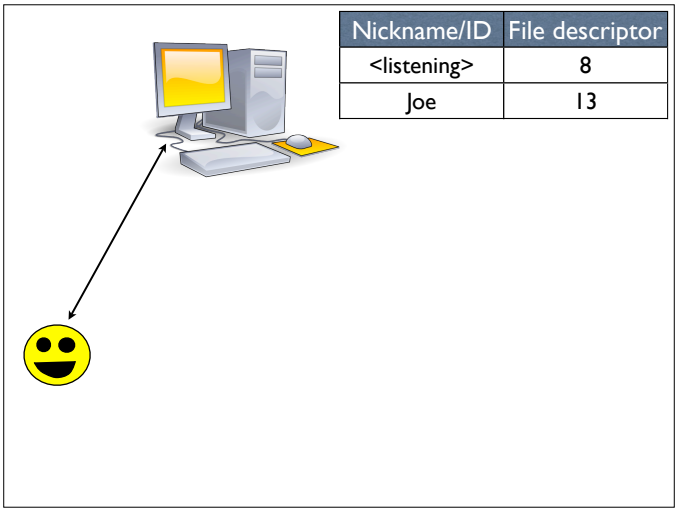
Nickname/ID
File descriptor
<listening>
8
Joe
13

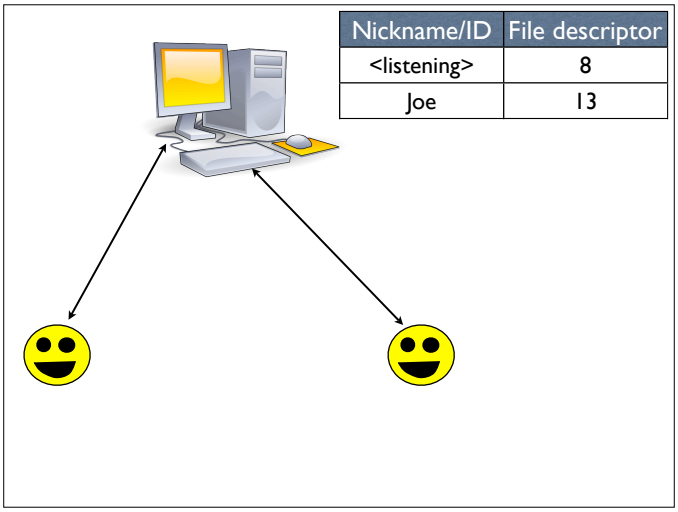
Nickname/ID
File descriptor
<listening>
8
Joe
13

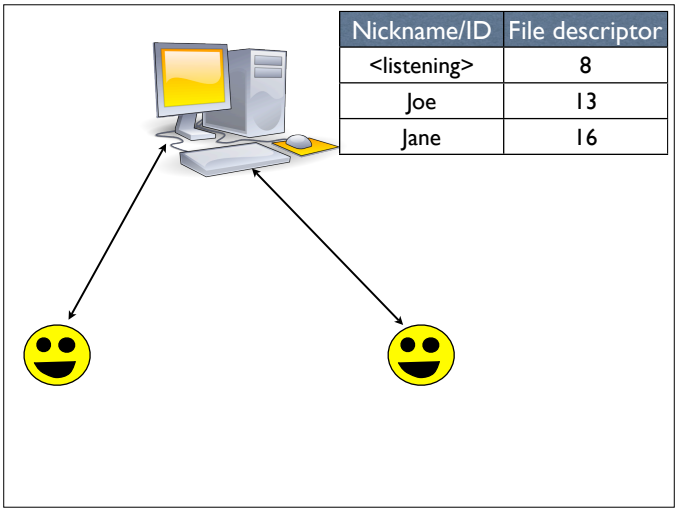
Nickname/ID
File descriptor
<listening>
8
Joe
13
Jane
16

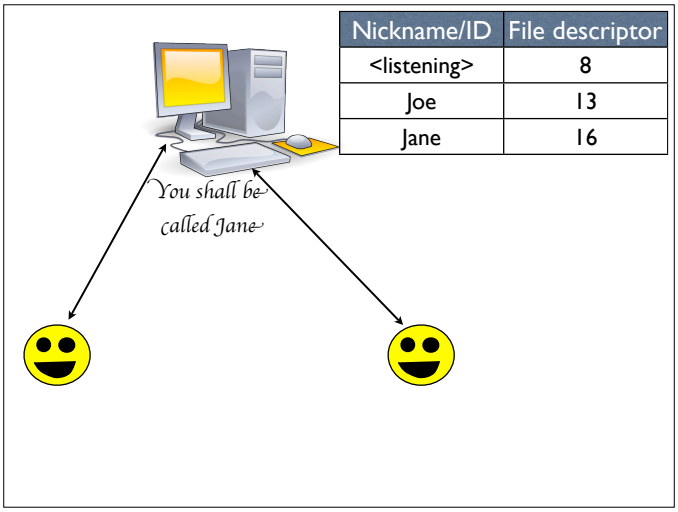
You shall be called Jane
Nickname/ID
File descriptor
<listening>
8
Joe
13
Jane
16

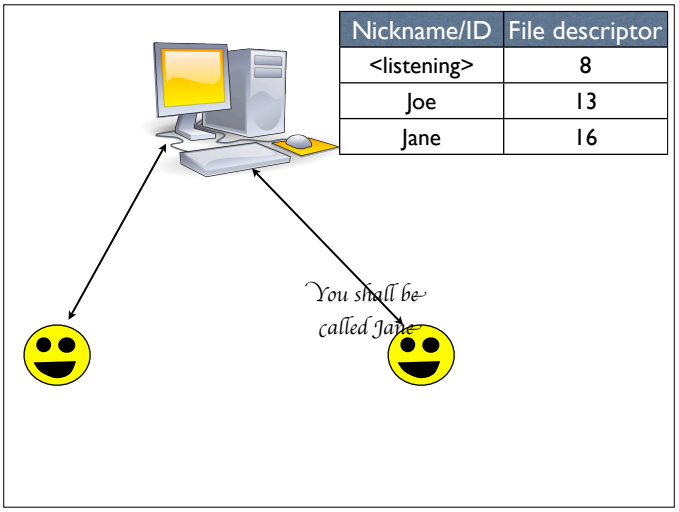
Nickname/ID
File descriptor
<listening>
8
Joe
13
Jane
16
You shall be called Jane

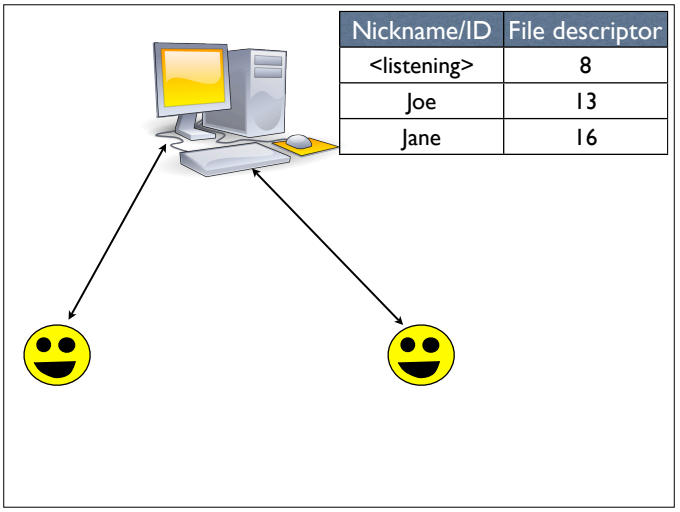
Nickname/ID
File descriptor
<listening>
8
Joe
13
Jane
16

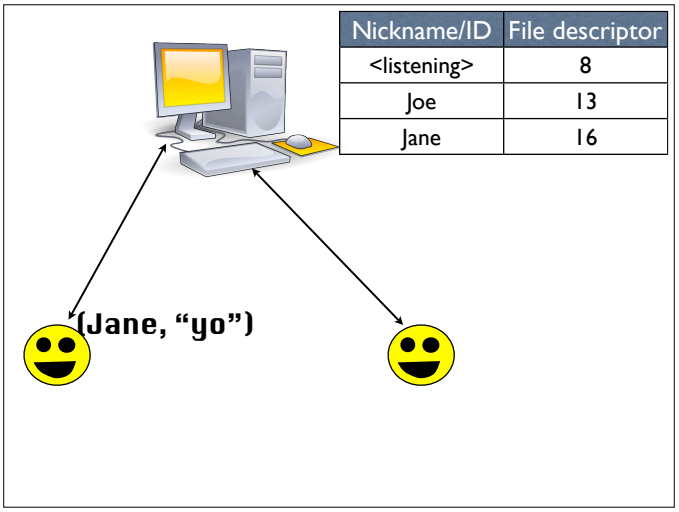
Nickname/ID
File descriptor
<listening>
8
Joe
13
Jane
16
(Jane, "yo")

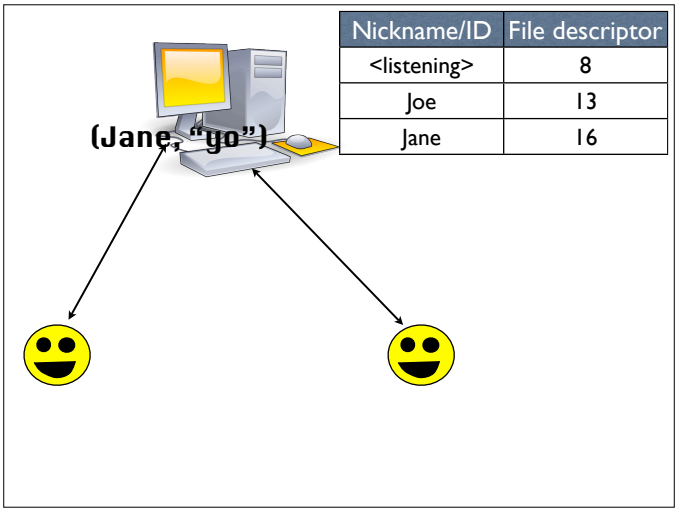
Nickname/ID
File descriptor
<listening>
8
Joe
13
Jane
16
(Jane, "yo")

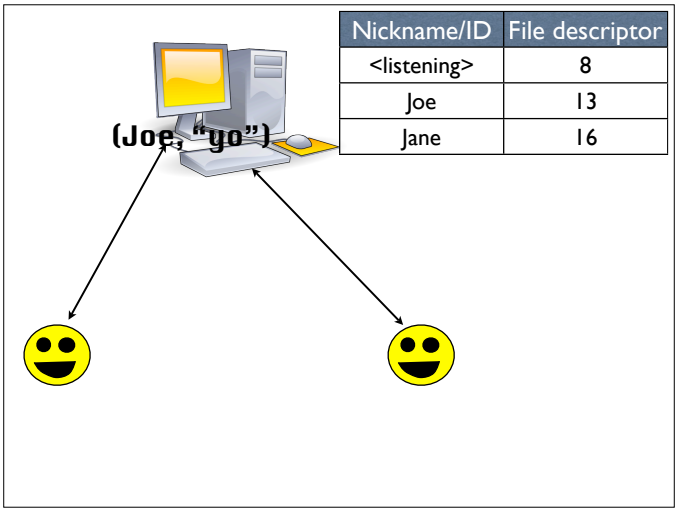
Nickname/ID
File descriptor
<listening>
8
Joe
13
Jane
16
(Joe, "yo")

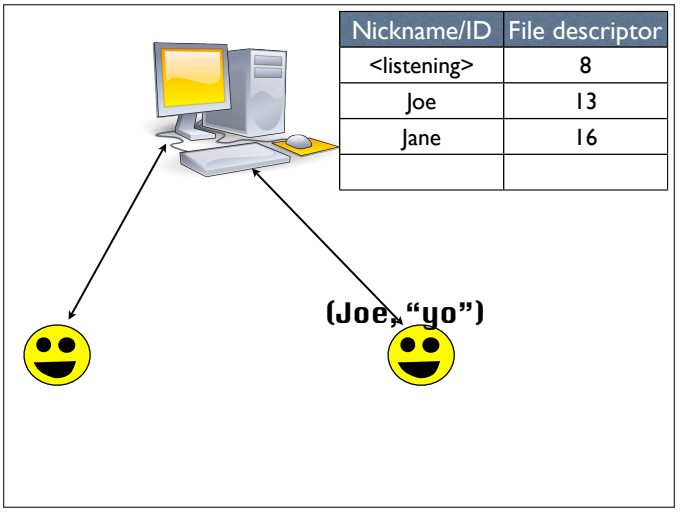
Nickname/ID
File descriptor
<listening>
8
Joe
13
Jane
16
(Joe, "yo")

Application
Transport
Network
Link
Physical

Application
Transport
Network
Link
Physical

Network vs Link

- Link layer forwards packets between hosts <u>on the same network</u>
  - Network layer routes packets <u>between networks</u>
- The link layer had no hope of scaling!
  - The forwarding table was linear in the number of hosts D:

IP (IPv4) addressing

- Addresses are 32 bits long
  - Typically given in "quad dotted" notation
    - E.g., 129.100.16.245
- Address space of ~4 billion addresses
- How does this help with routing?

IP (IPv4) addressing

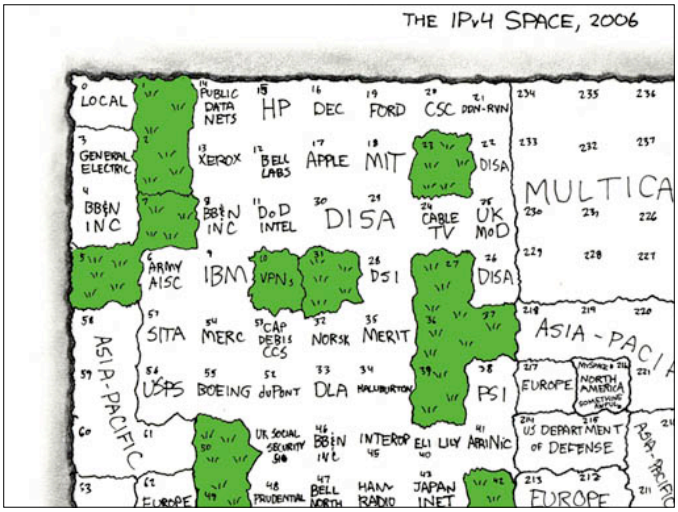

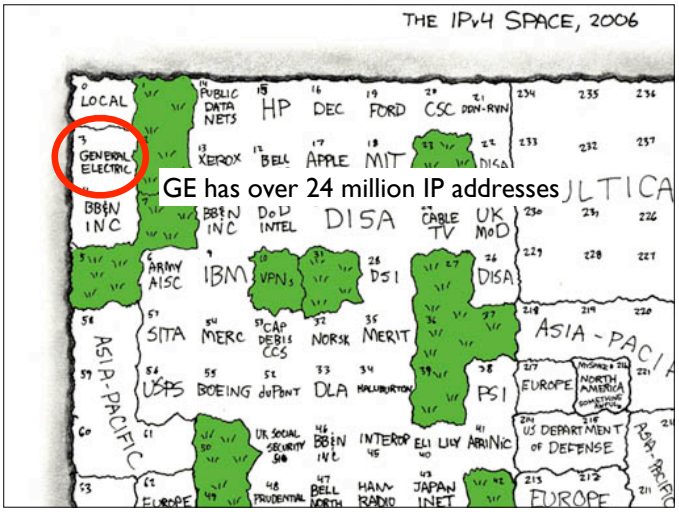

GE has over 24 million IP addresses

IP (IPv4) routing

- Prior to 1981, the most significant 8 bits were the "network number" and the least significant 24 bits were the "rest"
- This would be idiotic for a global network
  - Maximum of 255 networks in the entire world!

IP (IPv4) routing

- In 1981, "classful" routing was introduced
- Many "network numbers" had already been assigned (e.g., to GE)
  - These couldn't be reassigned!
- "Network numbers" 0 through 127 would remain as "class A" networks (8-bit network ID)
- 128 through 255 would be subdivided

IP Addresses

given notion of "network", let's re-examine IP addresses:

"class-full" addressing:

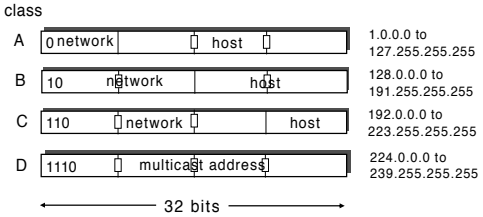

Side note: class E (240 to 255) was reserved, but never defined

"Classful" address space

- The 1981 change added a few other things
  - Three "private network" ranges were added (we'll talk about this later)
  - Multicast ranges were added (we'll talk about this later)
  - Loopback was added

Loopback

```
<Luigi30> someone ping flood
127.0.0.1
<Luigi30> he keeps trying to hack me
<Luigi30> nm
<Floach> This, folks, is why Mario
always gets the Princess.
```

- 127.0.0.0 through 127.255.255.255 are "loopback" addressing
- They refer to your own computer (no matter where you are)

IPv4 classful addressing

- 128 class A networks (each ~16 million hosts)
- ~16k class B networks (each ~65000 hosts)
- ~2 million class C networks (each ~250 hosts)
- Why didn't it work?

IPv4 classless addressing

- There was a large demand for networks of size, e.g., 300, 500, 1000, etc.
  - Everyone snapped up class B networks (~65000 hosts) which went 99% unused
- In 1993, we moved from classful addressing to classless addressing (CIDR)

IP addressing: CIDR

- Classful addressing:
  - inefficient use of address space, address space exhaustion
  - e.g., class B net allocated enough addresses for 65K hosts, even if only 2K hosts in that network
- CIDR: Classless InterDomain Routing
  - network portion of address of arbitrary length
  - address format: a.b.c.d/x, where x is # bits in network portion of address

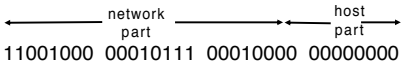

200.23.16.0/23

Transitioning classful to classless

- Class A had 8-bit network identifiers
  - Becomes /8. E.g., GE would be 3.0.0.0/8
- Class B had 16-bit network identifiers
  - E.g., 130.3.0.0/16
- Class C had 24-bit network identifiers
  - E.g., 200.278.0.0/24

Prefixes and netmasks

- /8, /16, /12, etc., are called "prefixes"
- They are often alternatively expressed as "netmasks"

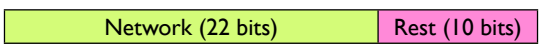

Prefixes and netmasks

- /8, /16, /12, etc., are called "prefixes"
- They are often alternatively expressed as "netmasks"

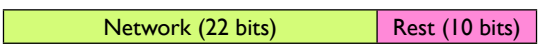

11111111 11111111 11111100 00000000

Prefixes and netmasks

- /8, /16, /12, etc., are called "prefixes"
- They are often alternatively expressed as "netmasks"

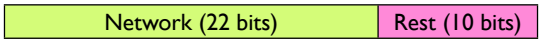

11111111 11111111 11111100 00000000

255.255.252.0

Netmasks

- Netmasks are strictly more flexible than prefixes
- 254.8.255.0
  - 11111110 00001000 11111111 00000000
  - Impossible to express as a prefix!
- In practice, it doesn't matter

Hooray for classless!

- Much more flexible
- An ISP can get a /16 (~65000 addresses)
- It can give a /31 (2 addresses) or /32 (1 address) to most of its customers
- If a company has 1500 computers, the ISP can hand out a /22 (2048 addresses)

IP addresses: how to get one?

Q: How does *host* get IP address?

- hard-coded by system administrator in a file
  - Wintel: control-panel->network->configuration->tcp/ip->properties
  - UNIX: /etc/rc.config
- DHCP: Dynamic Host Configuration Protocol: dynamically get address from as server
  - "plug-and-play"

(more shortly)

IP addresses: how to get one?

Q: How does a *network* get the network part of IP address?

A: gets allocated portion of its provider ISP's address space

ISP's block	11001000 00010111 00010000	00000000	200.23.16.0/20
Organization 0	11001000 00010111 00010000	00000000	200.23.16.0/23
Organization 1	11001000 00010111 00010010	00000000	200.23.18.0/23
Organization 2	11001000 00010111 00010100	00000000	200.23.20.0/23
...	.....	....	....
Organization 7	11001000 00010111 00011110	00000000	200.23.30.0/23

Routing problems

- Note that the smaller we make our networks, the better utilization we get
- An organization with 1500 computers needed a 20-bit subnet (2048 addresses)
- Optimally shouldn't we all just use /32s? (1-bit "networks" with only 1 address)

Hierarchical addressing: route aggregation

Hierarchical addressing allows efficient advertisement of routing information:

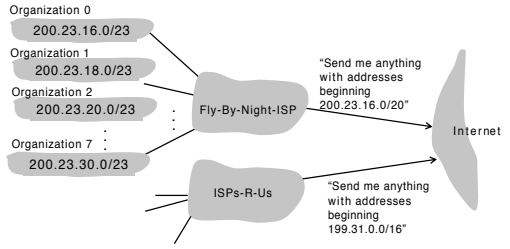

Hierarchical addressing: more specific routes

ISPs-R-Us has a more specific route to Organization 1

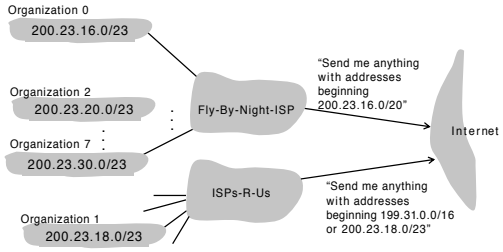

General concept behind routing

- The Internet wouldn't work if it weren't hierarchical
  - Massive routing tables
  - Loops in the routing?
- We route incrementally, only caring about how to get to the next hop
- Each router looks at only a few bits of each IP address
  - Routers at major exchange points may have 100k entries
  - Your home computer may have 2 entries

Routing example

- A router gets a packet destined for 195.43.86.9
- Router #1 looks at the last 7 bits and sends it to router #2
- Router #2 is a backbone router and looks at the first 15 bits before sending it off to a different backbone
- Router #3 ignores the first 15 bits and looks at the next 7 bits
- Router #4 checks the next 5 bits
- Router #5 can now ignore the first 27 bits and just look at the last 5 bits. It then decides to route it to your house

IP addressing: the last word...

Q: How does an ISP get block of addresses?

A: ICANN: Internet Corporation for Assigned Names and Numbers
- allocates addresses
- manages DNS
- assigns domain names, resolves disputes

More specifically, it's IANA (Internet Assigned Numbers Authority) working on behalf of ICANN

Note: ICANN is a US-owned company, created by act of the US government

ICANN

- ICANN was created in 1998 to handle IP address, domain names, etc., on behalf of the US government
- Up until then, there wasn't a clear authority
- In the early days of the Internet, gentleman's agreements reigned

Dramatization only


UC Berkley


MIT

Dramatization only

Hey I'm setting up my routing table. I'll put you guys in as network 1.


UC Berkley


MIT

Dramatization only

Sounds good. I'll put you guys in my routing table as network 2.


UC Berkley


MIT

Dramatization only


UC Berkley


MIT

Dramatization only

UWO


UC Berkley


MIT

Dramatization only

Hey guys, can I join in?

UWO


UC Berkley


MIT

Dramatization only

Guys? What network number are you going to give me?

UC Berkley

UWO

MIT

The need for ICANN

- It's a bigger deal with domain names than IP addresses, but a central authority started becoming more important to solve disputes
- There are proposals for replacing ICANN
  - UN?
  - No signs of it disappearing any time soon