

Lecture 10: Addressing

CSE 123: Computer Networks
Alex C. Snoeren



Lecture 10 overview



- Finish up ICMP
- Class-based addressing
- Subnetting
- Classless addressing



Common ICMP messages

- Destination unreachable
 - ♦ “Destination” can be host, network, port, or protocol – and *fragmentation needed*
- TTL Expired
 - ♦ Used by the “traceroute” program
 - » traceroute traces packet routes through Internet
- Echo request/reply
 - ♦ Used by the “ping” program
 - » ping just tests for host liveness



ICMP restrictions

- The generation of error messages is limited to avoid cascades
 - ◆ Error causes error that causes error...
- Don't generate ICMP error in response to:
 - ◆ An ICMP error
 - ◆ Broadcast/multicast messages (link or IP level)
 - ◆ IP header that is corrupt or has an invalid source address
 - ◆ Fragments, except the first
- ICMP messages are often rate limited, too
 - ◆ Don't waste valuable bandwidth sending tons of ICMP messages

IP addressing



- Need to be able to name hosts (sources/destinations)
 - ◆ Recall hosts are connected to a LAN of a certain technology
 - ◆ Have a link-local address that is useful on that LAN only
 - ◆ Need a globally routable (i.e., unique) name for every device on the Internet
- But also need to be able to name networks
 - ◆ Recall routers are just forwarding between networks
 - ◆ Moving within a given network is a L2 task
- Most importantly, need to know which network to which a host attaches
 - ◆ Ideally don't want to have to enumerate all hosts



Addressing considerations

- Fixed length or variable length addresses?
- Issues:
 - ◆ Flexibility
 - ◆ Processing costs
 - ◆ Header size
- Engineering choice: IP uses fixed length addresses

IP addresses



- 32-bits in an IPv4 address
 - ◆ Dotted decimal format a.b.c.d
 - ◆ Each represent 8 bits of address
- Hierarchical: Network part and host part
 - ◆ E.g. IP address 128.54.70.238
 - ◆ 128.54 refers to the UCSD campus network
 - ◆ 70.238 refers to the host ieng6.ucsd.edu
- Which part is network vs. host?



Class-based addressing

- Most significant bits determines “class” of address


Class A	0	Network	Host	127 nets, 16M hosts		
Class B	1	0	14 Network	16 Host	16K nets, 64K hosts	
Class C	1	1	0	21 Network	8 Host	2M nets, 254 hosts

- Special addresses
 - ◆ Class D (1110) for multicast, Class E (1111) experimental
 - ◆ 127.0.0.1: local host (a.k.a. the loopback address)
 - ◆ Host bits all set to 0: network address
 - ◆ Host bits all set to 1: broadcast address

L3 forwarding over LANs: Subnetting



- Individual networks may be composed of several LANs
 - ◆ Only want traffic destined to local hosts on physical network
 - ◆ Routers need a way to know which hosts on which LAN
- Networks can be arbitrarily decomposed into **subnets**
 - ◆ Each subnet is simply a prefix of the host address portion
 - ◆ Subnet prefix can be of any length, specified with **netmask**





Subnet addresses

- Every (sub)network has an address and a **netmask**
 - ◆ Netmask tells which bits of the network address is important
 - ◆ Called a “mask” because it is logically AND’d with IP address
 - ◆ Convention suggests it be a proper prefix
- Netmask written as an all-ones (bits) IP address
 - ◆ E.g., Class B netmask is 255.255.0.0 (i.e, 11111111111111110000000000000000)
- Need to size subnet appropriately for each LAN
 - ◆ Only have remaining bits to specify host addresses



IP address problem (circa 1991)

- Address space depletion
 - ◆ In danger of running out of classes A and B
- Why?
 - ◆ Class C too small for most organizations (only ~250 addresses)
 - ◆ Very few class A – very careful about giving them out (who has 16M hosts anyway?)
 - ◆ Class B – greatest problem

CIDR



- Classless Inter-Domain Routing (1993)
 - ♦ Networks described by variable-length prefix and length
 - ♦ Allows arbitrary allocation between network and host address



- ♦ e.g. 10.95.1.2 contained within 10.64.0.0/10:
 - » 10.64.0.0 is network and remainder (31.1.2) is host
 - » 0000101001 / 0111110000000100000010
- Pro: Finer grained allocation; aggregation
- Con: More expensive lookup: **longest prefix match**



Longest matching prefix

- FIBs contain many prefix/length tuples
 - ◆ They *need not* be disjoint!
 - ◆ E.g. 200.23.16.0/20 and 200.23.18.0/23
 - ◆ What to do if a packet arrives for destination 200.23.18.1?
 - ◆ Need to find the longest prefix in the table which matches it (200.23.18.0/23)
- Not a simple table, requires multiple memory lookups
 - ◆ Lots and lots of research done on this problem
 - ◆ Lots of this work was historically done by UCSD faculty

For next time



- Read 6.2.3-4
- Keep plugging on Project 1b