

Lecture 12:

Addressing and Aggregation

CSE 123: Computer Networks

Stefan Savage





Lecture 12 Overview

- Finish up addressing
 - ◆ Class-based addressing
 - ◆ Subnetting
 - ◆ Classless addressing
 - ◆ Address allocation



Class-based Addressing

- Most significant bits determines “class” of address

Class A

0	Network	Host
---	---------	------

127 nets, 16M hosts

Class B

		1	0	Network	Host
				14	16

16K nets, 64K hosts

Class C

			1	1	0	Network	Host
						21	8

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




IP Forwarding Tables

- Router needs to know where to forward a packet
- Forwarding table contains:
 - ◆ List of network names and **next hop** routers
 - ◆ Local networks have entries specifying which interface
 - » Link-local hosts can be delivered with Layer-2 forwarding
- E.g. `www.ucsd.edu` address is `132.239.180.101`
 - ◆ Class B address – class + network is `132.239`
 - ◆ Lookup `132.239` in forwarding table
 - ◆ Prefix – part of address that really matters for routing



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
 - ◆ Convention suggests it be a proper prefix
- Netmask written as an all-ones IP address
 - ◆ E.g., Class B netmask is 255.255.0.0
 - ◆ Sometimes expressed in terms of number of 1s, e.g., /16
- Need to size subnet appropriately for each LAN
 - ◆ Only have remaining bits to specify host addresses

IP Address Problem (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



IPv4 Census Map

June - October 2012

Utilization



Prefix Sizes



420 Million hosts that responded to ICMP Ping at least 2 times between June and October 2012
Source: Carna Botnet

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 10.0.0.0/8:
 - » 10.0.0.0 is network and remainder (95.1.2) is host
- Pro: Finer grained allocation; aggregation
- Con: More expensive lookup: **longest prefix match**



Longest Matching Prefix


- Forwarding table contains 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

PATRICIA Tree (Radix tree)



(Practical Algorithm to Retrieve Information Coded In Alphanumeric)


- Straightforward way to look up LMP
 - Arrange route entries into a series of bit tests
 - Worst case = 32 bit tests
 - Problem: memory speed is a bottleneck





Route Aggregation


- Combine adjacent networks in forwarding tables
 - ◆ Helps keep forwarding table size down





Most Specific Route

- But what if address range is not contiguous?



Aside: so where do IP addresses come from?



- It depends who's asking...
 - ◆ Your machine
 - » DHCP protocol (next time) gives you an individual IP address
 - ◆ Your company
 - » Ask ISP for chunk of address space
 - » Operate address space independently (bigger companies)
 - ◆ Your ISP
 - » Ask routing registries for big chunks of address space
 - » ARIN (North America), RIPE (Europe), AFRINIC (Africa), LANIC (South America), APNIC (Asia/Pacific)
- Big chunks of IPv4 address space are “legacy” blocks (not operated by routing registries). Ownership unclear. Private sales happen.



The space crunch...

- Still running out of IP addresses... what to do?
- Two solutions
 - ◆ Network Address Translation – multiple multiple hosts on a single IP address (future class)
 - ◆ Get bigger addresses -> IPv6
- IPv6: 128bit addresses... we won't run out
 - ◆ 64bit routing prefix, 64bit host id

An IPv6 address (in hexadecimal)

2001:0DB8:AC10:FE01:0000:0000:0000:0000

↓ ↓ ↓ ↓ |
2001:0DB8:AC10:FE01:: Zeroes can be omitted

0010000000000001:0000110110111000:1010110000010000:1111111000000001:
0000000000000000:0000000000000000:0000000000000000:0000000000000000



IPv6 vs IPv4 header

Ver	Traffic Class	Flow Label	
Payload Length		Next Header	Hop Limit
Source Address			
Destination Address			

Ver	Hdr Len	Type of Service	Total Length	
Identification			Flg	Fragment Offset
Time to Live		Protocol	Header Checksum	
Source Address				
Destination Address				
Options...				

Gray bits are unique to each header

Changes

- Eliminate fragmentation-related fields

- Eliminate header checksum

- Added flow label

- Quadruple size of addresses

- IPv6 header (40 bytes) vs IPv4 (20 bytes)



IPv6 Transition is slow

- Need to support both protocols at the same time
 - ◆ Complicated... if a destination has both a IPv4 and IPv6 address which to use?
- Less need in developed world -> slower adoption
- That said
 - ◆ All major operating systems now support IPv6
 - ◆ All major router vendors
 - ◆ US Mobile carriers (e.g., Tmobile, Verizon, etc)
 - ◆ Offered as option by many US ISPs
- In your lifetime it is likely that IPv6-based addressing will start to dominate

Summary



- You can't route efficiently on flat address spaces
 - ◆ You'd need a table the size of all hosts on the Internet
 - ◆ You'd need to send updates about that table to everyone
- Network-layer addressing is done hierarchically
 - ◆ Routing prefix + host suffix
 - ◆ Originally, this split was done statically (class-based addressing)
 - ◆ Now it is done dynamically (CIDR)
 - ◆ Requires more complex forwarding table lookup
 - ◆ Allows contiguous chunks of address space to be aggregated (for the purposes of routing) into fewer prefixes



For Next Time

- Read 3.2.6 and 9.3.1
- Project 1 & HW 2