Address scarcity, NAT, and IPv6

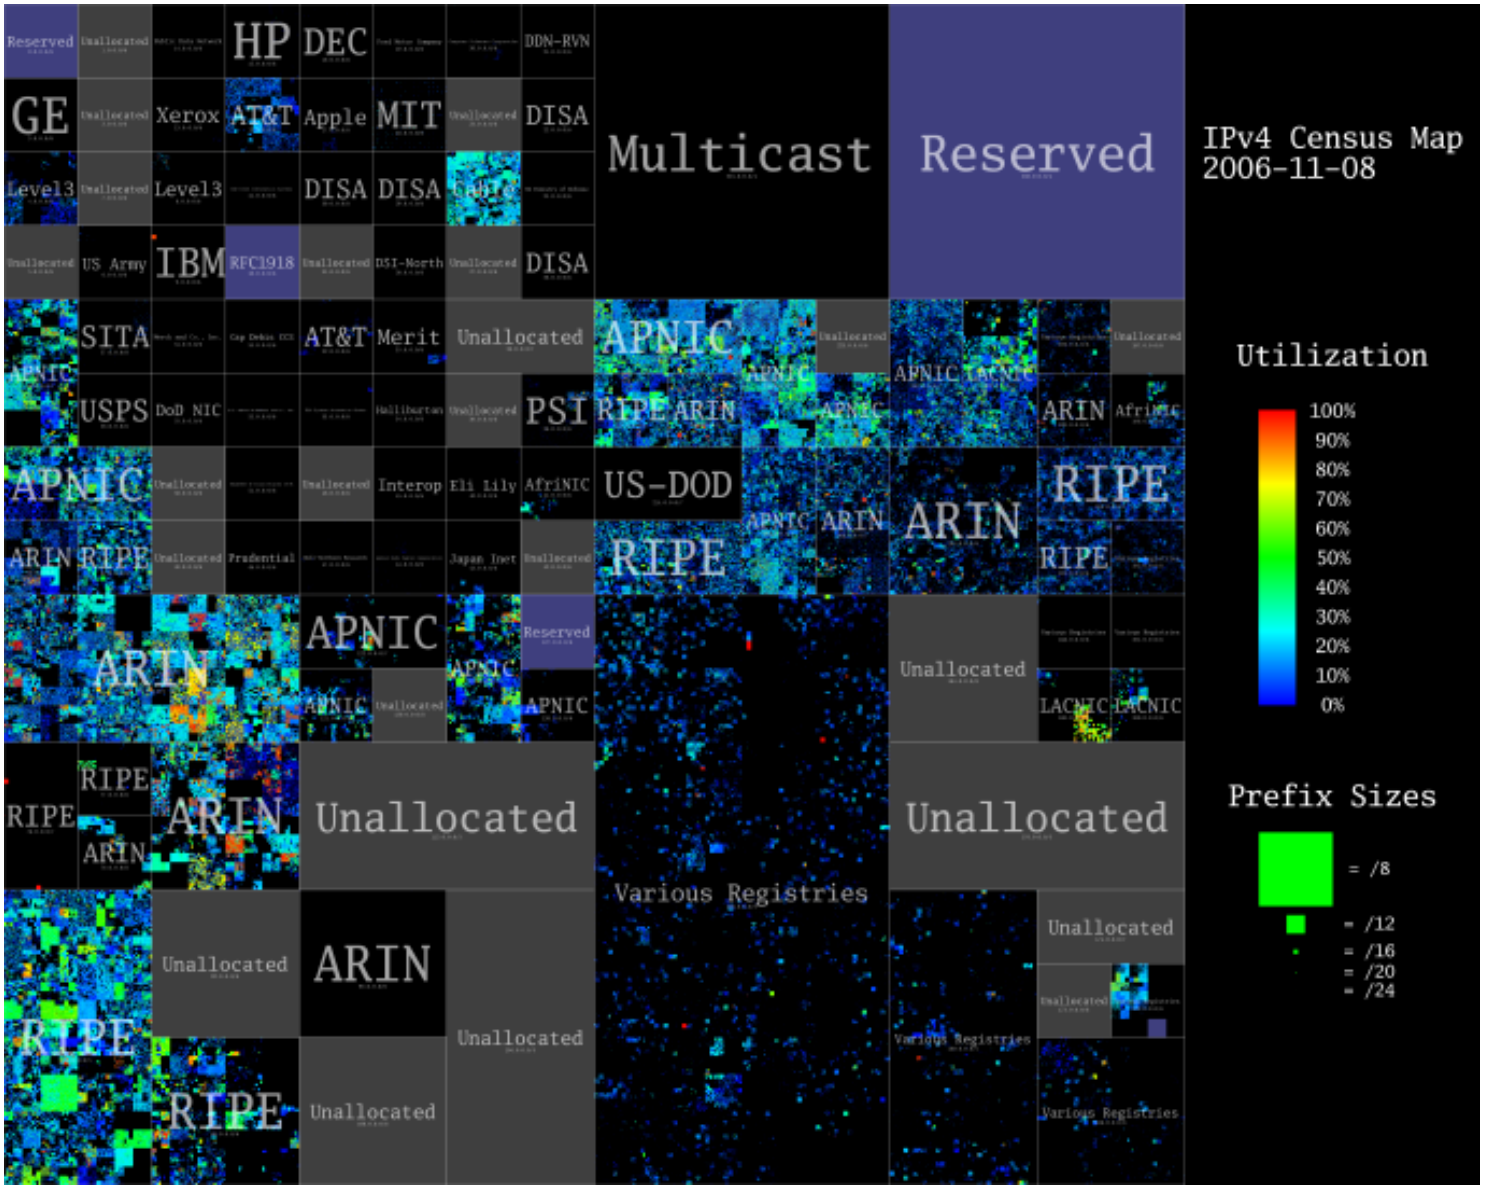

Overview

- Handling IPv4 address scarcity
  - Address space exhaustion
  - Classless addressing (CIDR)
  - Network Address Translation (NAT)
  - Bigger and better IP protocol (IPv6)

Internet Protocol

- IP service model
  - A global addressing scheme
  - Best effort delivery of datagrams
- Internet Protocol (IP) v4
  - 4-byte addresses allowing for hierarchical routing
  - Best case, 2^{32} = 4 billion unique hosts
  - Utilization is far from best case, ~250 million hosts

IPv4 address format

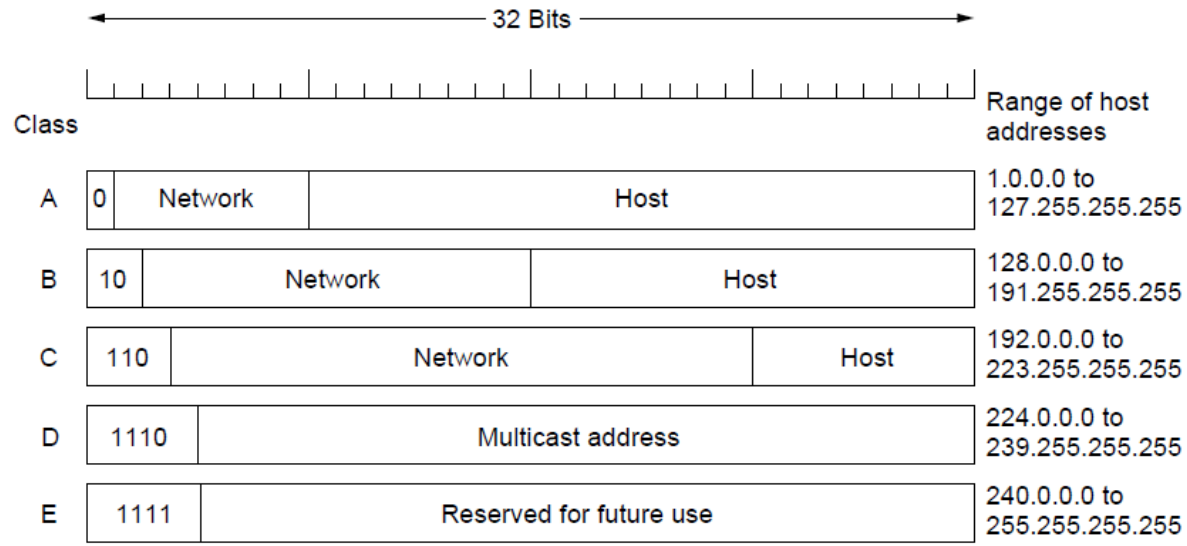

- Classful addressing (before 1993):
  - Class A: 128 networks with 16 million hosts
  - Class B: 16,384 networks with 65,536 hosts
  - Class C: 2 million networks, 256 hosts

Classless addressing

- Classless Interdomain Routing (CIDER)
  - We want:
    - Efficient address allocation
    - Small and fast forwarding tables
- Compromise:
  - Aggregate contiguous blocks of IP addresses
  - New /X notation
    - Specify how many prefix bits are network number
    - Like subnet mask, with X 1's and front then 0's
    - In CIDR 1's must be contiguous

MAP OF THE INTERNET

THE IPV4 SPACE, 2006

THIS CHART SHOWS THE IP ADDRESS SPACE ON A PLANE USING A FRACTAL MAPPING WHICH PRESERVES GROUPING -- ANY CONSECUTIVE STRING OF IPs WILL TRANSLATE TO A SINGLE COMPACT, CONTIGUOUS REGION ON THE MAP. EACH OF THE 256 NUMBERED BLOCKS REPRESENTS ONE /8 SUBNET (CONTAINING ALL IPs THAT START WITH THAT NUMBER). THE UPPER LEFT SECTION SHOWS THE BLOCKS SOLD DIRECTLY TO CORPORATIONS AND GOVERNMENTS IN THE 1990's BEFORE THE RIRs TOOK OVER ALLOCATION.

0 1 14 15 16 19 →
3 2 13 12 17 18
4 7 8 11
5 6 9 10

= UNALLOCATED BLOCK

http://xkcd.com/195/

IPv4 address exhaustion

- Jan 31, 2011
  - Last unreserved IANA /8 blocks allocated
  - 5 remaining blocks allocated to each of 5 Regional Internet registries (RIR)

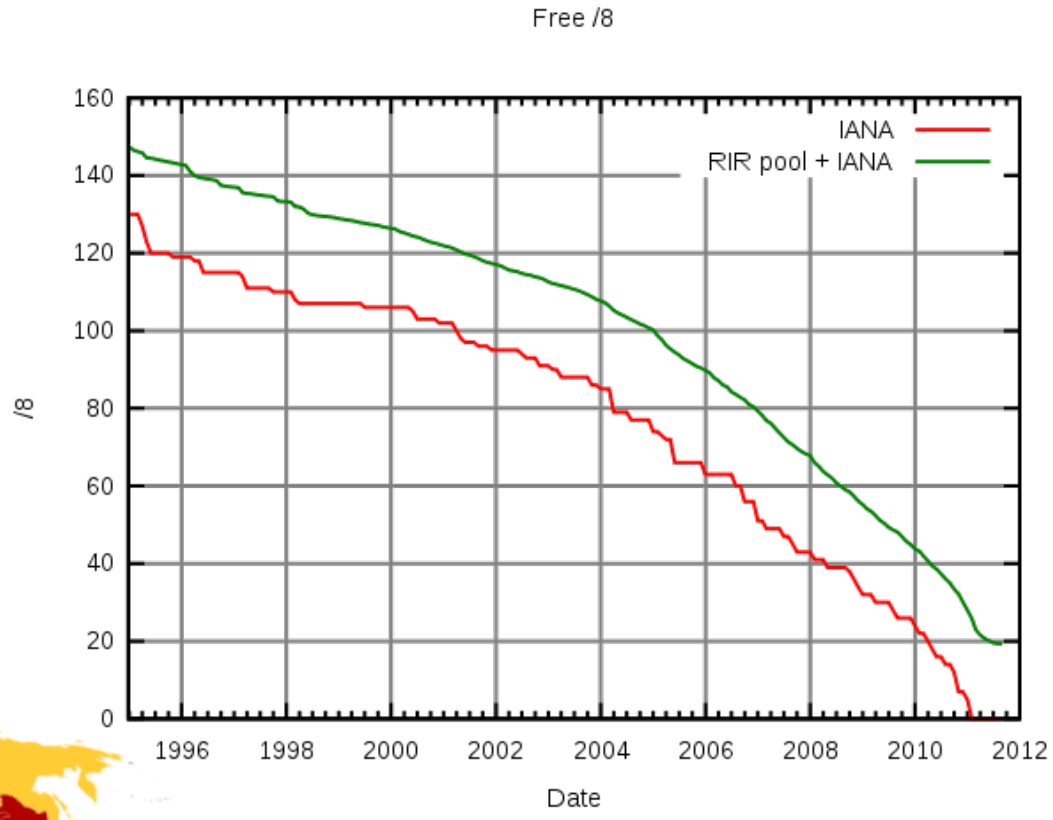

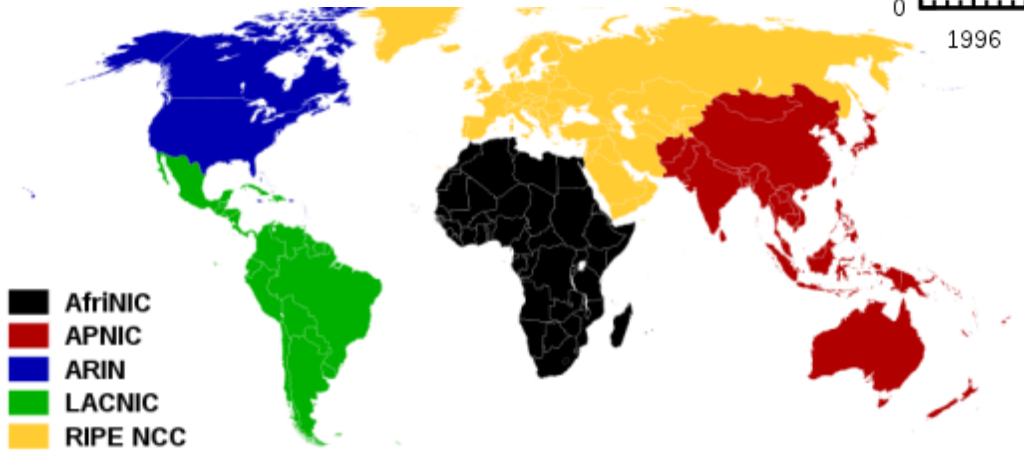

http://www.youtube.com/watch?v=y8WqJum_Gfg

BGP advertisements

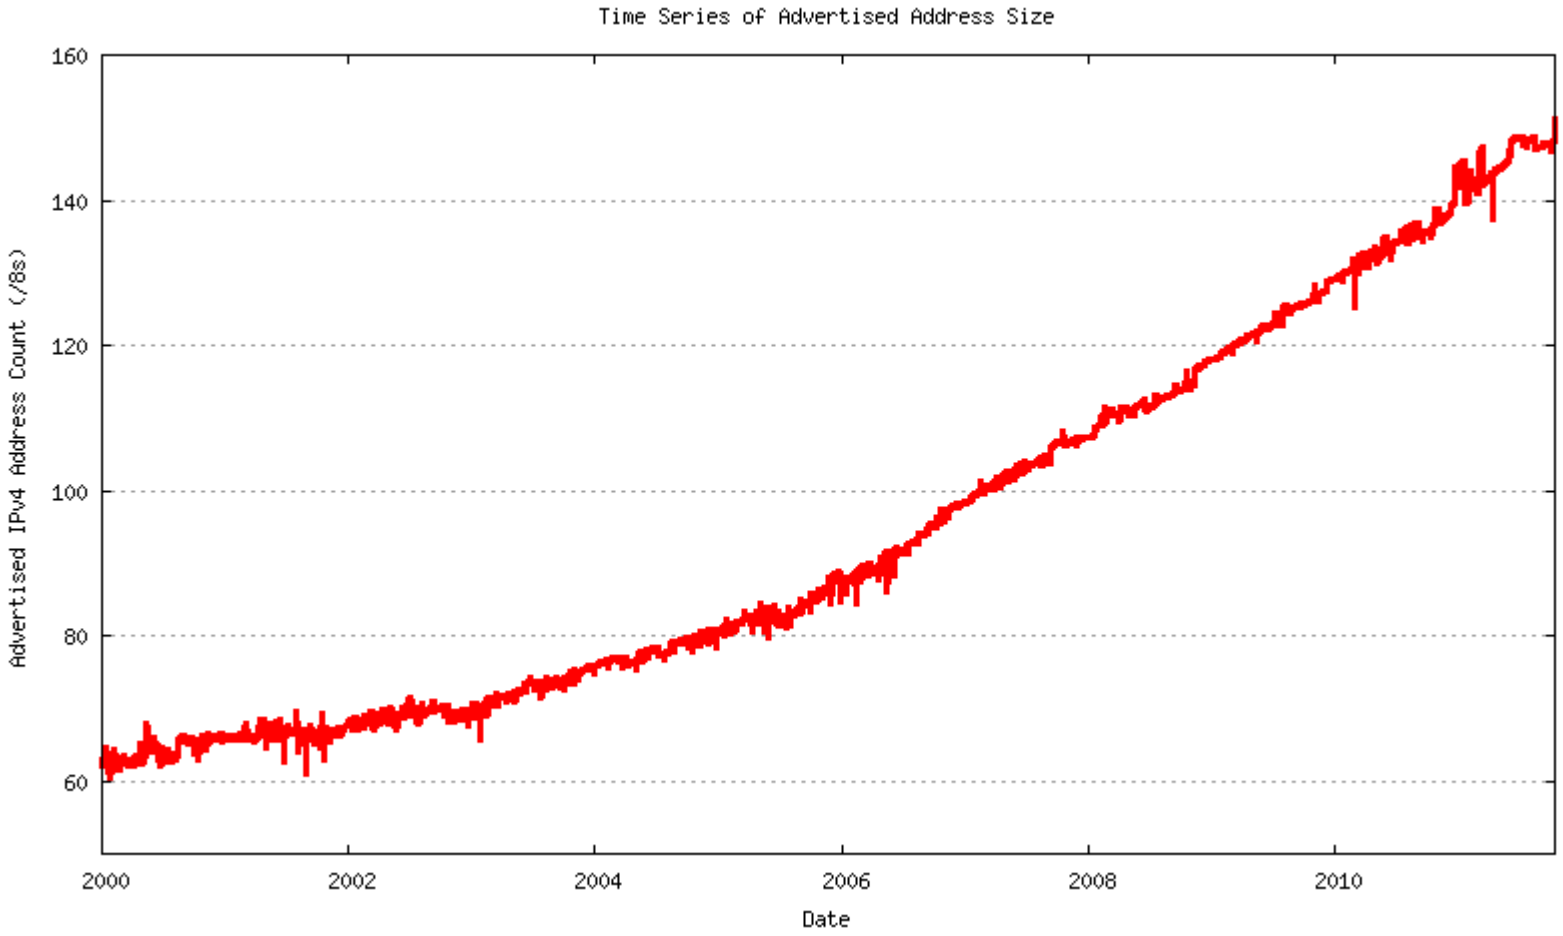

http://www.potaroo.net/tools/ipv4/index.html

NAT

- Network address translation (NAT)
  - Quick fix to address scarcity
  - Home/business gets one public IP
    - Private IP addresses for all hosts inside network
  - NAT box translates at boundary to public IP

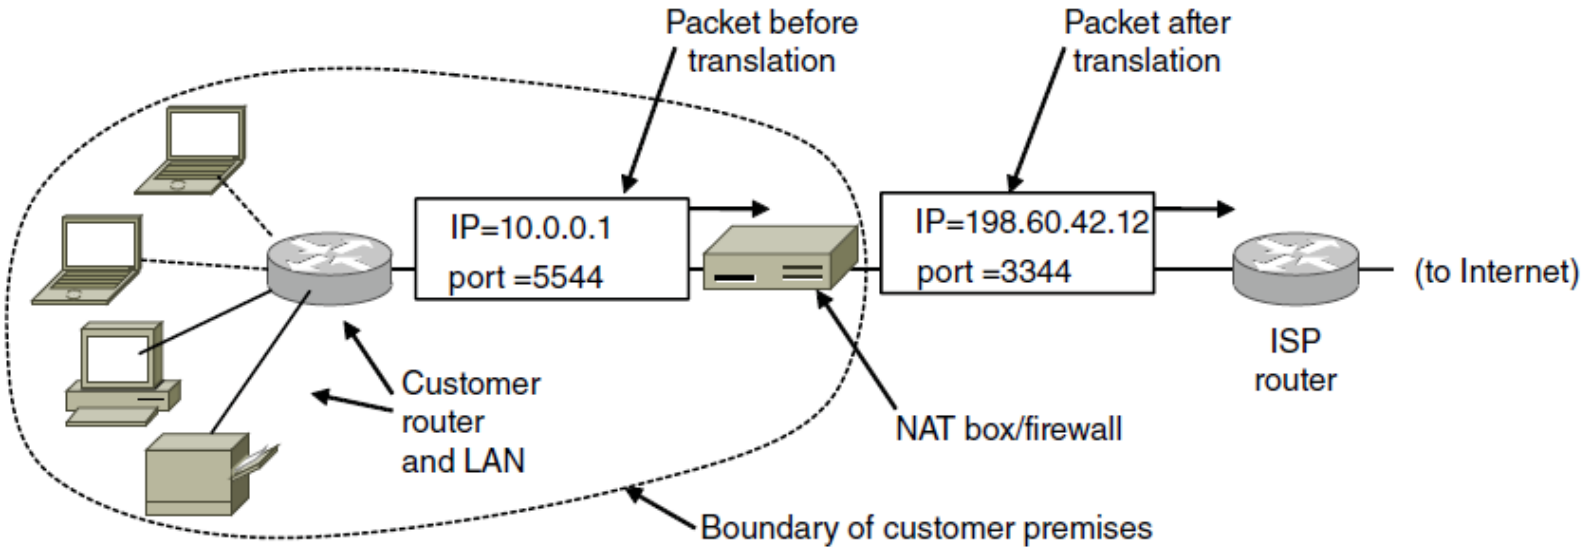

NAT design

- Problem: Where to route reply from remote server?
  - NAT designers observed:
    - Most IP traffic over TCP/UDP
    - TCP/UDP have a 16-bit integer port #
      - Source port and destination port (e.g. 80 for web)
- Solution: Use source port as an index into a translation table

NAT translation

- Map outgoing packets
  - Replace src addr → NAT box addr (public IP)
  - Replace src port # → new port #
- Maintain a translation table
  - (src address, port #) → (NAT addr, new port #)
  - Free up entry after timeout (frees up port #)
- Incoming packets
  - Consult translation table
  - Rewrite packet and send to local host

NAT example

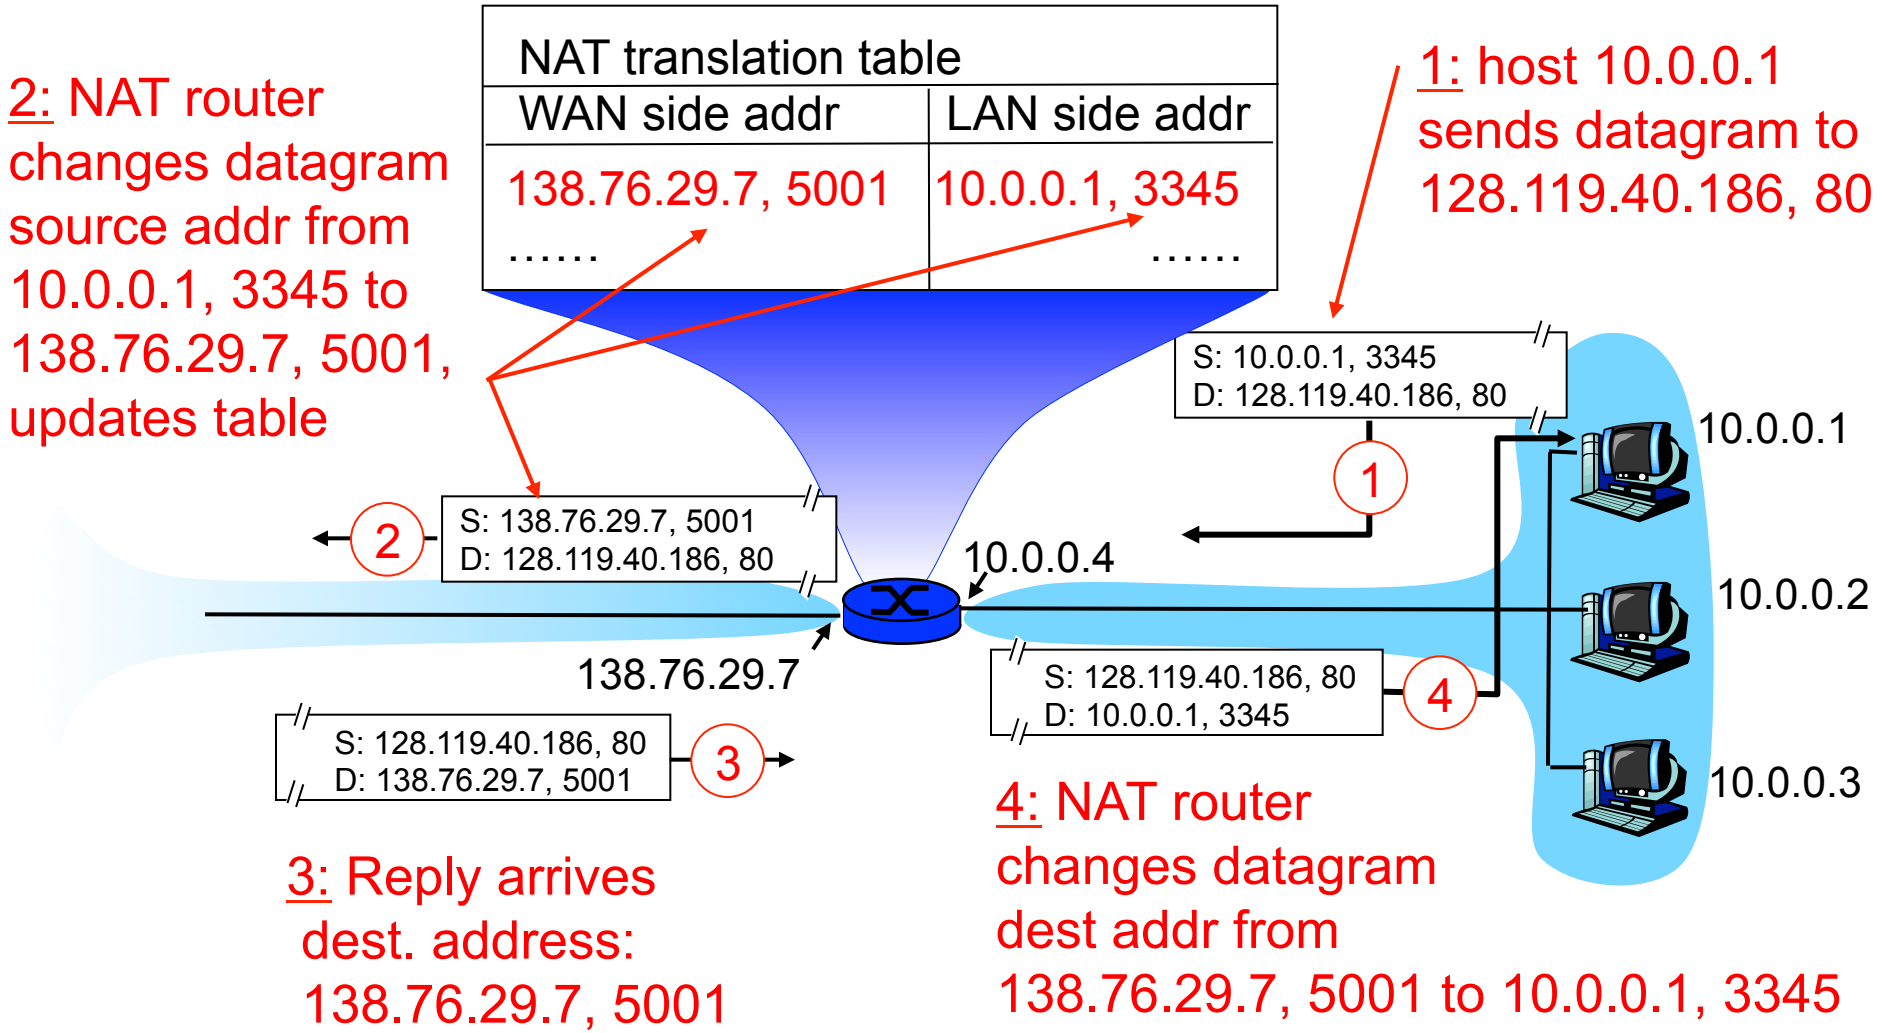
NAT translation table
WAN side addr
LAN side addr
138.76.29.7, 5001
10.0.0.1, 3345
……
……
2: NAT router changes datagram source addr from 10.0.0.1, 3345 to 138.76.29.7, 5001, updates table
1: host 10.0.0.1 sends datagram to 128.119.40.186, 80
S: 10.0.0.1, 3345
D: 128.119.40.186, 80
1
2
S: 138.76.29.7, 5001
D: 128.119.40.186, 80
10.0.0.4
138.76.29.7
10.0.0.1
10.0.0.2
10.0.0.3
S: 128.119.40.186, 80
D: 10.0.0.1, 3345
4
S: 128.119.40.186, 80
D: 138.76.29.7, 5001
3
3: Reply arrives dest. address: 138.76.29.7, 5001
4: NAT router changes datagram dest addr from 138.76.29.7, 5001 to 10.0.0.1, 3345

Where is NAT implemented?

- Home router
  - Integrates router, DHCP server, NAT, firewall, etc.
  - Single IP address on WAN side from service provider
- Campus or corporate network
  - NAT box at Internet connection point
  - Share a collection of public IP addresses
    - Allows many hosts inside network

NAT advantages

- Helps converse IPv4 addresses
- Easy to switch Internet providers
  - All your devices are using private IPs via DHCP
- Provides a measure of security
  - Outside computers cannot initiate connections
  - However, doesn't protect against:
    - Connections initiated from behind the NAT box to bad places
    - Attacks from hosts inside network

NAT an abomination?

1) Violates the IP model
   - Every host should have unique identifier
2) Breaks end-to-end connectivity model
   - Any host can send a packet to any other host at any time
3) Not connectionless
   - NAT box has state, effectively circuit switching
   - Single point of failure
4) Network layers are not independent
   - NAT looks into the payload

NAT an abomination?

5) Forces use of TCP/UDP protocols
   – Anything else, NAT fails to find TCP Source port
6) Breaks if multiple TCP/IP or UDP ports
   – e.g. FTP and H.323 Internet telephony
7) Limited number of hosts on NAT box
   – Only 16-bits in TCP Source port
   – Can't have > 64K machines on a single IP

NAT traversal

- Make connections through NAT boxes
  - Client-to-client apps:
    - Voice over IP, video conference, file sharing, gaming
  - One type: UDP hole punching
    - Goal: establish UDP connection between clients
    - Approach: Use central server with public IP to coordinate. Establish direct UDP connections between clients.

UDP hole punching

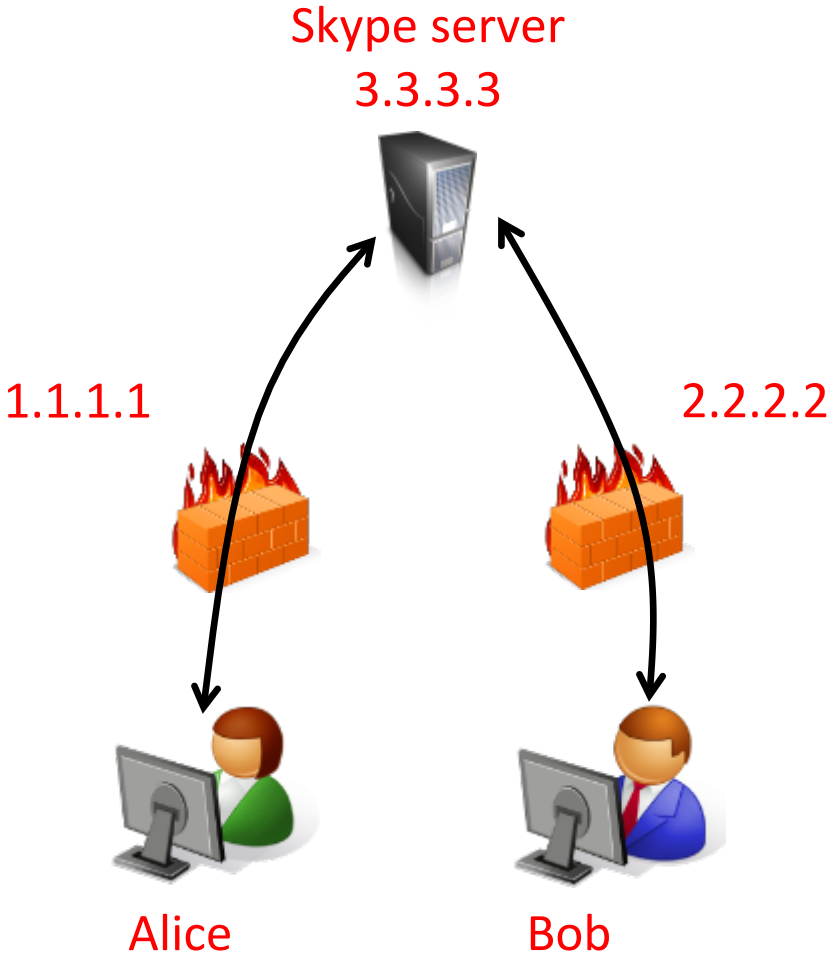

1. Permanent TCP connections to public central server.

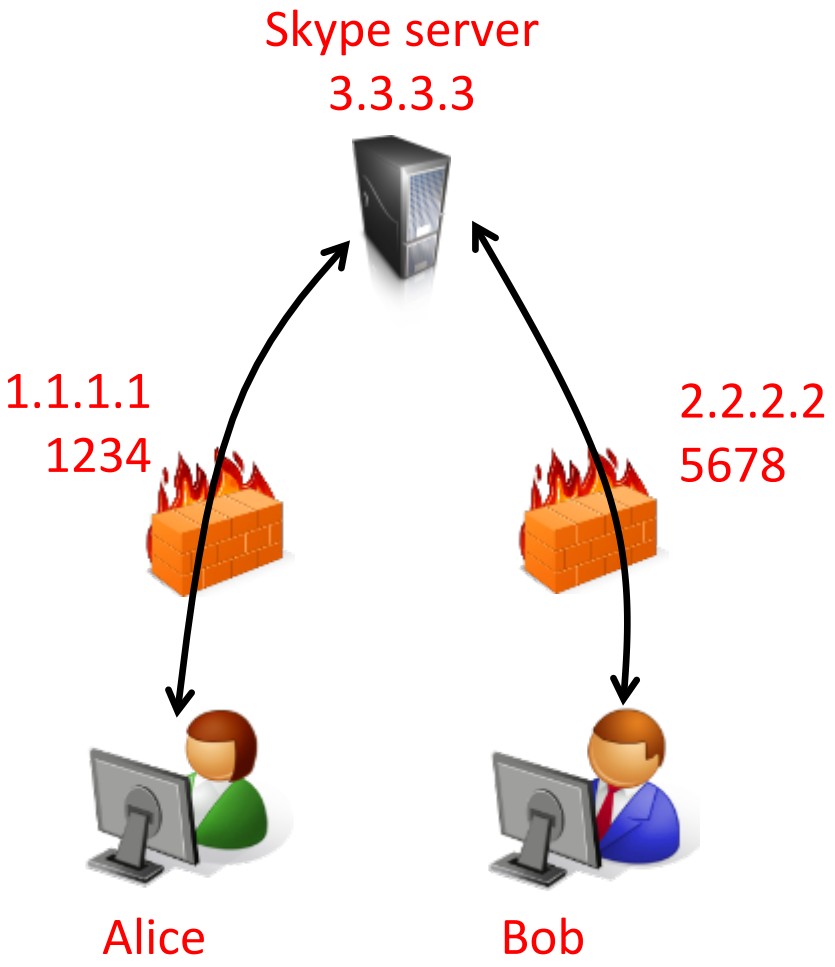

2. Tests reveals UDP port Alice and Bob use to send voice data.

UDP hole punching

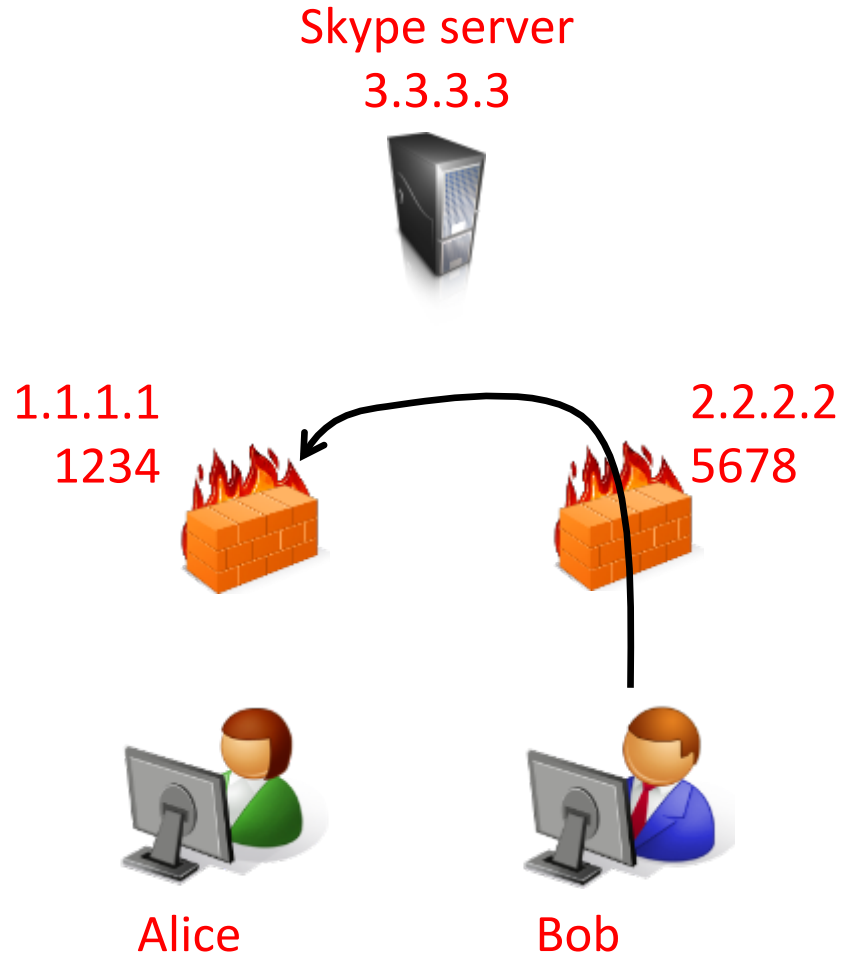

3. Bob sends Alice UDP packet on port 1234. Alice's firewall drops.

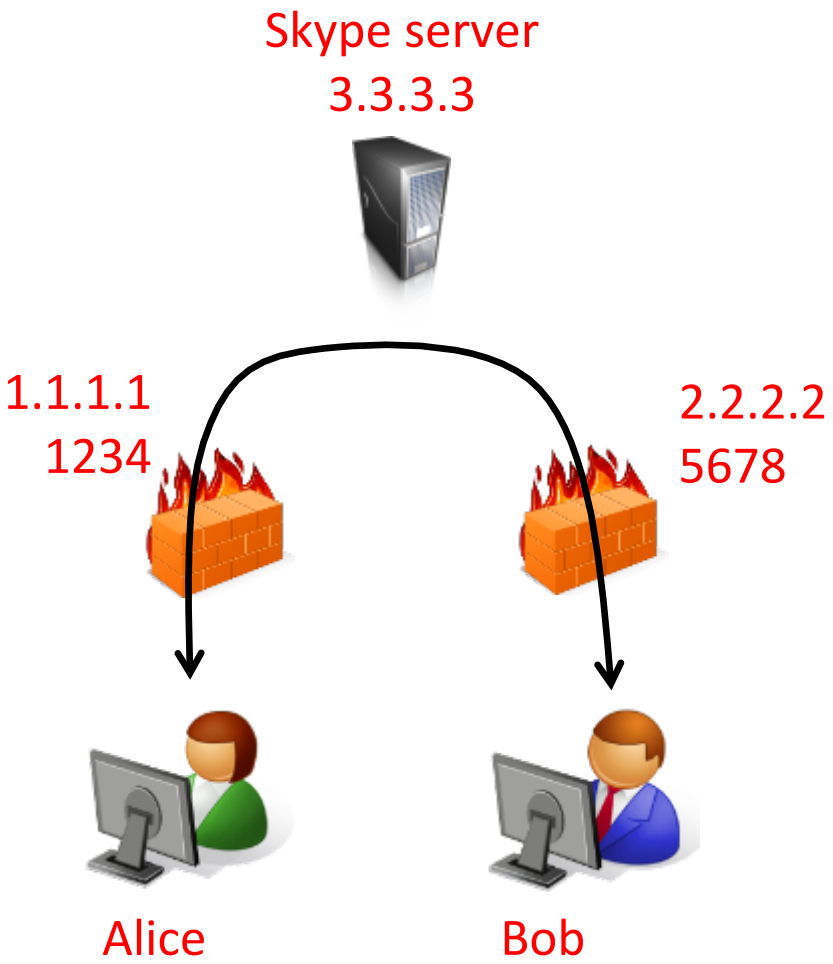

4. Alice sends Bob UDP packet on port 5678. Bob's firewall thinks it is a response to his blocked initial packet.

Internet Protocol: TNG

- Birth of IP version 6
  - Started looking at IPv4 exhaustion in 1991
  - Increase address size → new IP packet header
    - Thus new software for every Internet host/router
    - Might as well overhaul the whole thing
    - Draft standard in 1998

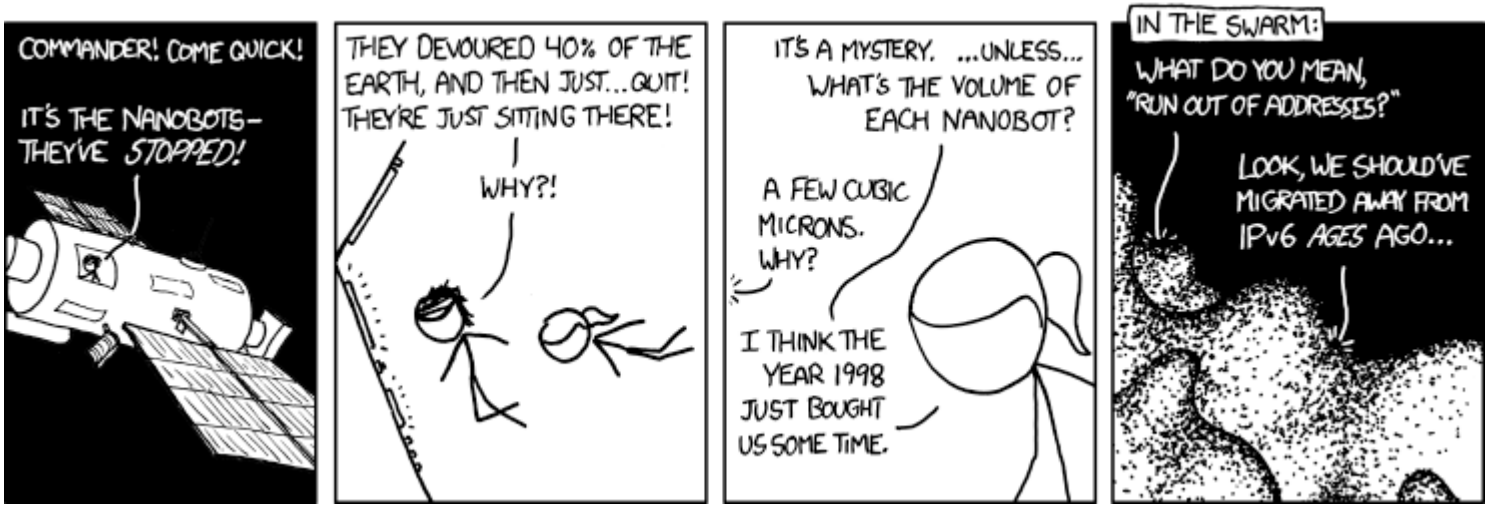

http://xkcd.com/865/

IPv6 goals & features

1. Support billions of hosts
   - 2^{128} addresses ≈ 3 x 10^{38}
   - If entire planet covered with computers:
     - 7 x 10^{23} IPs/ m^2, pessimistic utilization scenario: 1000 IPs / m^2
   - Address format: 8 groups of 4 hex digits

Full address	8000:0000:0000:0000:0123:4567:89AB:CDEF
Abbreviated	8000::0123:4567:89AB:CDEF
IPv4 mapped to IPv6	::FFFF:192.31.20.46

00…0 (128 bits)	Unspecified
00…1 (128 bits)	Loopback
1111 1111…	Multicast address
1111 1110 10…	Link-local unicast
Everything else	Global unicast addresses, 99% of the space

IPv6 goals & features

2. Simplify the protocol
   - Allow routers to process packets faster
   - Support gigabit/terabit routing
     - Predictable header size (40 bytes)
     - Removed little used fields
     - No checksum
   - Allow future protocol evolution
   - Extension headers

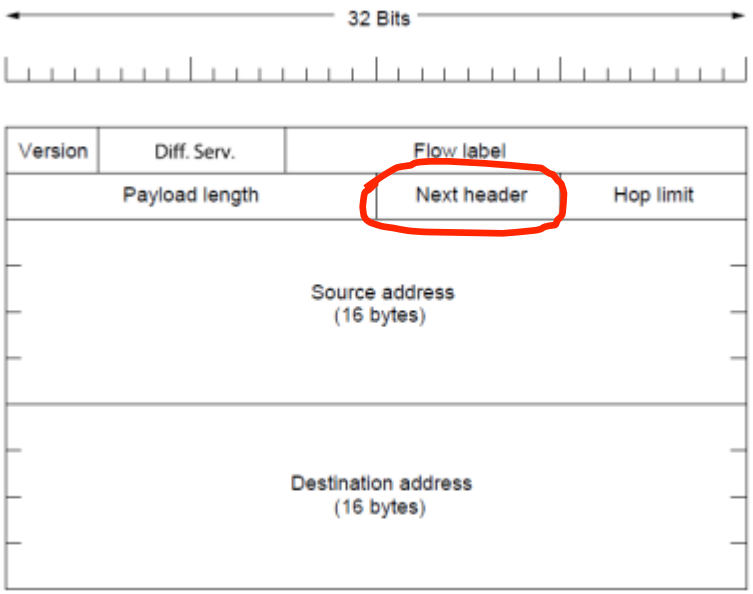

IPv6 fixed 40-byte header.

Extension headers

- Next header field
  - Allows chain of extension headers
  - Last one indicates payload protocol
    - e.g. 6 TCP, 17 UDP

Extension header	Description
Hop-by-hop options	Only extension that must be processed by all nodes. Support for datagrams exceeding 64 KB.
Destination options	Fields needed at destination host.
Routing	Lists one or more routers than must be visited on the way to destination. Similar to IPv4 loose source routing.
Fragmentation	Datagram identifier, fragment number, more fragments to follow. Must be done by source host, no fragmentation allowed in route. IPv6 requires MTU path discovery.
Authentication	Receiver can verify who sent it.
Encrypted security payload	Allows payload to be encrypted so only receiver can read it.

IPv6 goals & features

3. Autoconfiguration of hosts

- Guaranteed unique IPv6 address: prefix + 48-bit MAC
- Avoid users having to read/write 16 bytes addresses

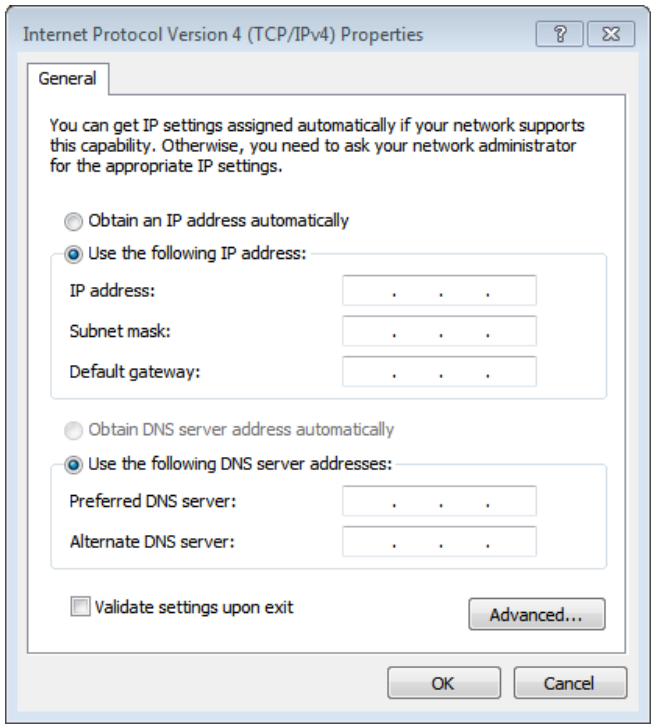

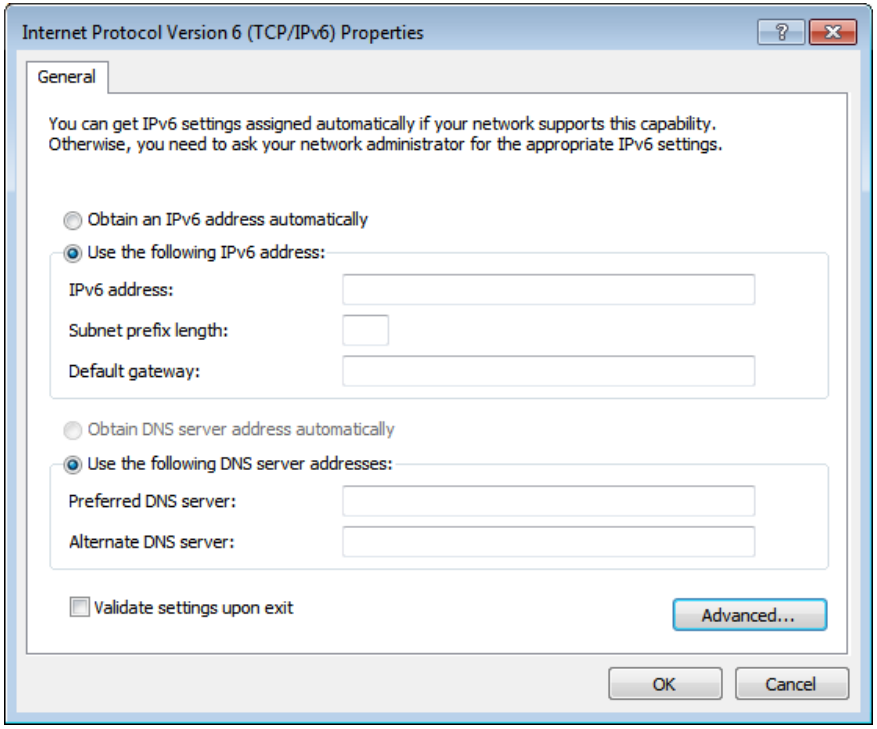

192.168.1.3

8000:0000:0000:0000:0123:4567:89AB:CDEF

IPv6 goals & features

4. Multicast/multimedia
   - Multicast a requirement, no longer optional
   - IPv4 DiffServ field + new 20-bit traffic flow field
   - Anycast, one address for a group of nodes
     - Delivery to only one node
     - Fault-tolerance, load balancing
     - Routing to closest node

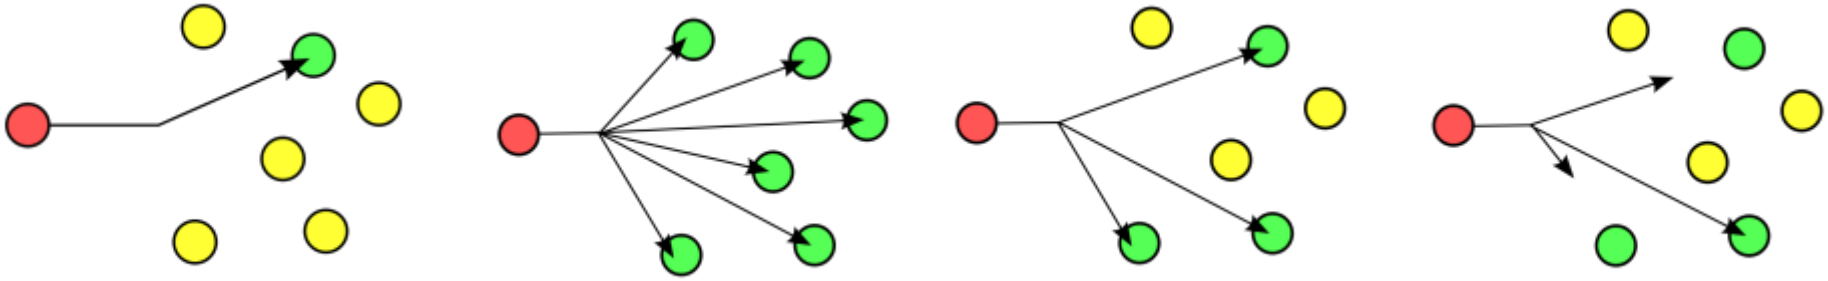

Unicast Broadcast Multicast Anycast

IPv6 goals & features

5. Improved security
   - IP security architecture (IPSec)
     - End-to-end security at the network layer
     - Must be in a IPv6 complaint node
     - An optional feature of an IPv4 node
   - Authentication header (AH)
     - Supports many different authentication techniques
     - Protects against attacks based on masquerading
   - Encapsulating security payload (ESP)
     - Integrity and confidentiality of datagram

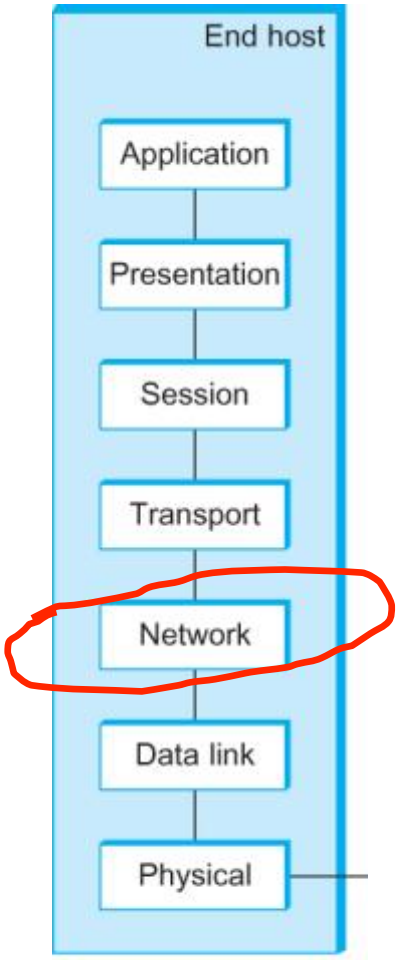

IPv6 goals & features

6. Support for mobile hosts
   - Mobile clients likely to be majority of IPv6 hosts
   - Mobile IPv6 (RFC 3775)
   - Use IPv6 features:
     - Stateless autoconfiguration
     - Neighbor discovery
     - Extension headers such as routing header

Mobile IP

- Routing to mobile hosts
  - Home address
    - Permanent IP of mobile host
  - Home agent
    - Router on your home network
    - Acts as your agent when you aren't attached to the home network
  - Foreign agent
    - Located on network mobile host connected to
    - Not always required

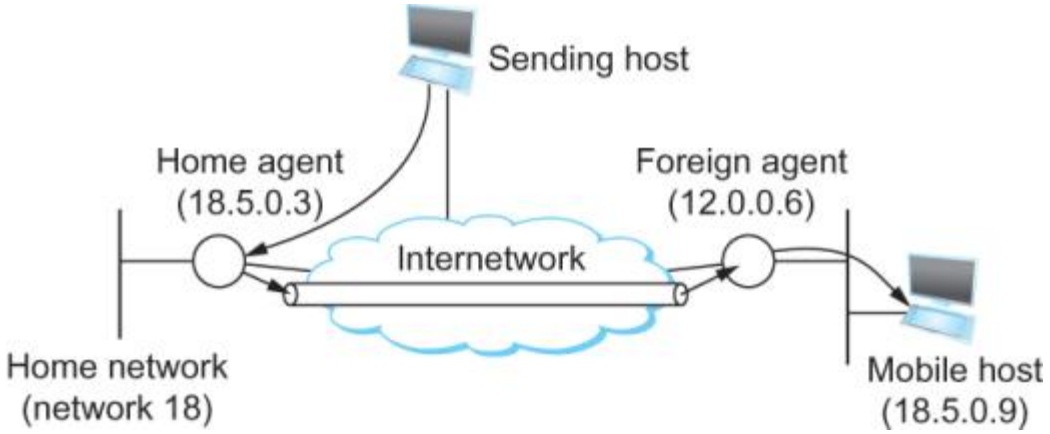

Mobile IP

- Mobile host contacts foreign agent
  - Foreign agent provides care-of address to home agent
- Home agent impersonates mobile host
  - Proxy ARP, intercepts traffic to mobile's permanent address
- Home agent tunnels traffic to foreign agent
  - Forwarded on to mobile host

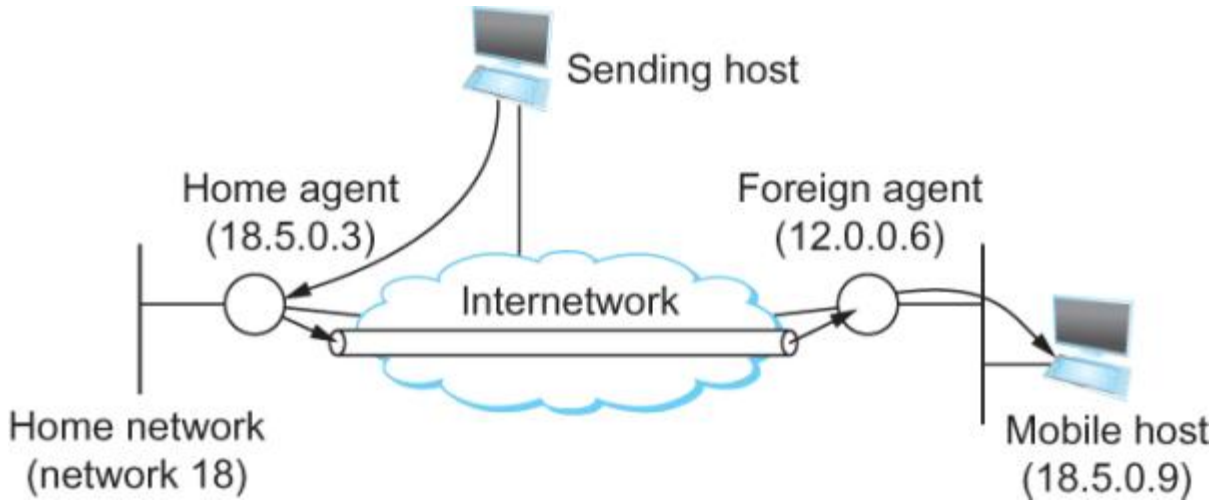

IPv6 goals & features

7. Ease of deployment
   - Achilles heel of IPv6
     - Google 2008 estimate, < 1% of traffic
   - We can't have a "flag" day to switch over

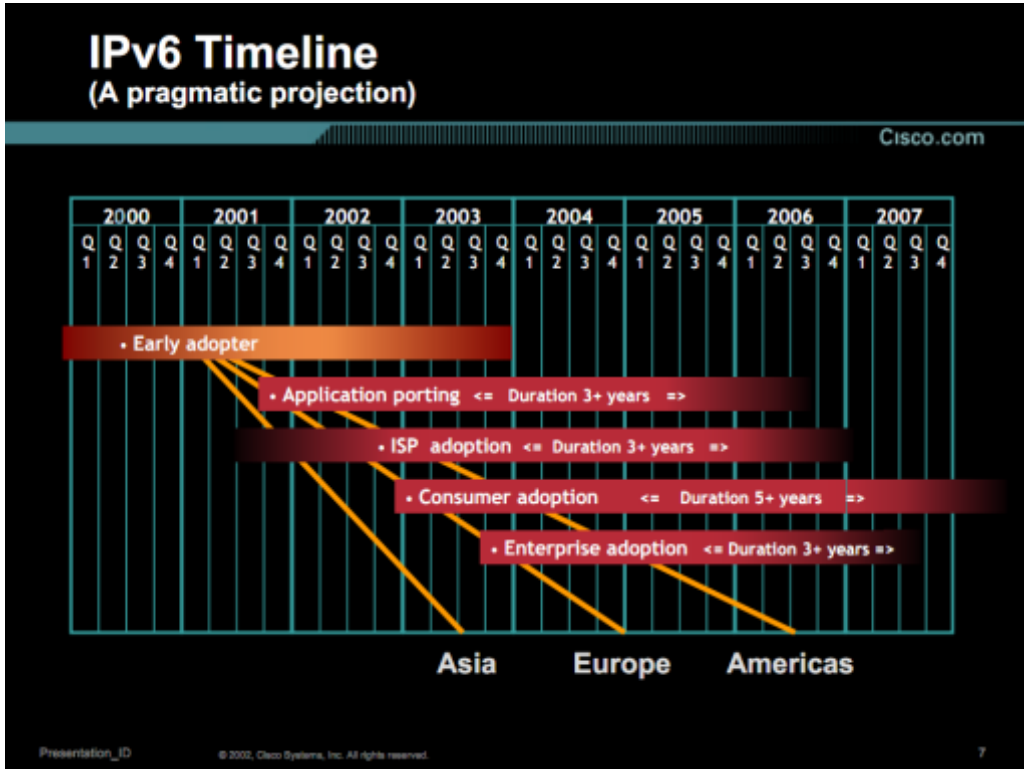

Deploying IPv6

- Dual-stack operation
  - IPv6 nodes also run IPv4
  - Consult version field in header to decide
  - Supported by major OS's for a long time

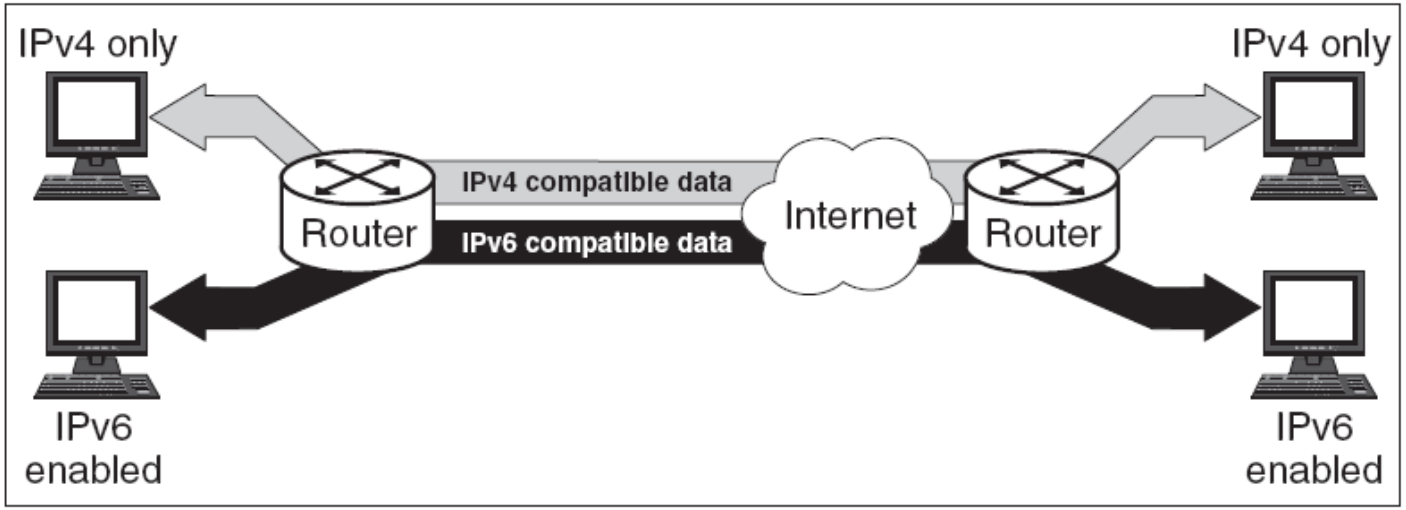

Source: GAO.

Deploying IPv6

- Tunneling IPv6 over IPv4 networks
  - Route IPv6 traffic over network segment that only understands IPv4

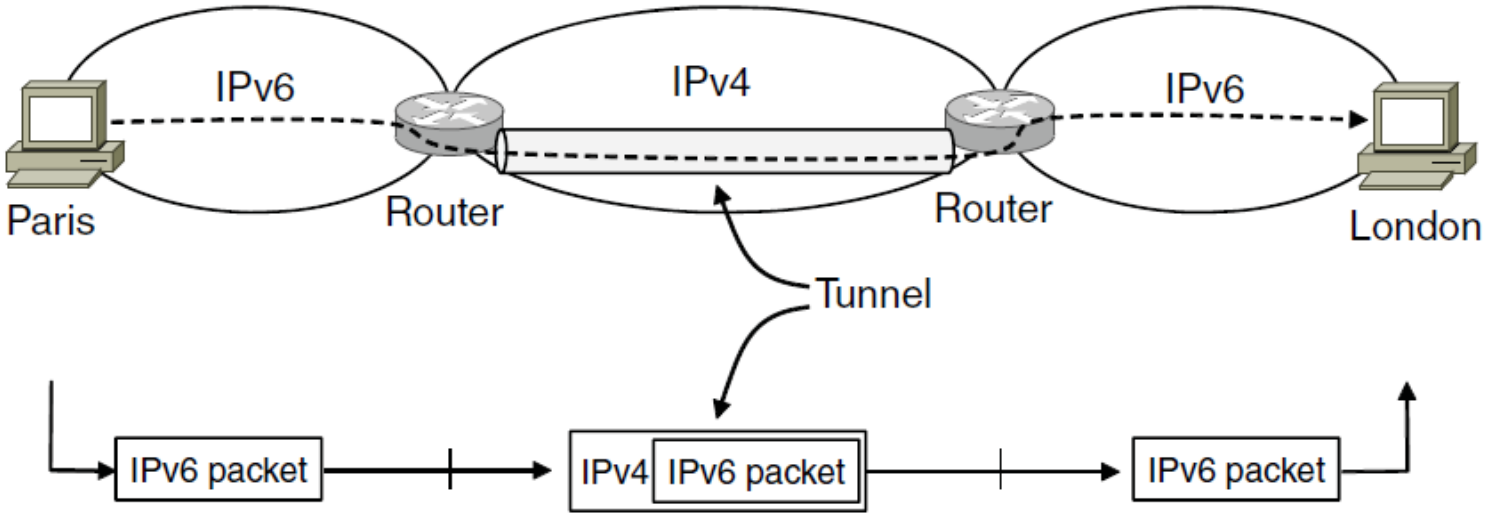

Summary

- Exhaustion of IPv4 address space
  - Temporary fixes:
    - CIDR, NAT, returning blocks
  - Permanent fix:
    - Migrate to IPv6
- IPv6
  - New version of protocol that runs the Internet
  - Deployment challenging
    - Dual stack support, tunneling over IPv4-only nodes