Networking Potpourri: Plug-n-Play, Next Gen

14-740: Fundamentals of Computer Networks
Bill Nace

Material from Computer Networking: A Top Down Approach, 6th edition. J.F. Kurose and K.W. Ross

Administrivia

- Next time: 45-minute quiz in 75 minutes
- Paper Review cycle starts up next week
  - 2 reviews in the next 2 lectures
- Historical: 29 Oct 69, ARPAnet first linked SRI to UCLA (briefly)

traceroute

- DHCP
- NAT
- IPv6

DHCP

- Dynamic Host Configuration Protocol
  - Client-server mechanism to get config data
  - IP address, DNS server, etc
  - Client is usually a recently-booted host
- RFC 1531 (1993) replaced BOOTP
- Obsoleted by RFC 2131(1997)
- DHCPv6, an IPv6 version in RFC 3315

Principle of Operation

- Newly booted client needs configuration details
  - Most important: its own IP address
  - Others: next-hop gateway, subnet mask, DNS server, time servers, static routes, TCP TTL value, etc
- Broadcasts requests
- Server, somewhere, responds with details

DHCP Server

- Controls a pool of IP addresses
- And a repository of network details
  - Provides these details upon request or by default
- Multiple servers possible
- Also, a single server may serve multiple subnets

Server: IP Allocation

- Automatic allocation: A permanent IP address is assigned to the client
- Dynamic allocation: IP address is assigned for a limited period of time
  - Allows for automatic reuse
- Manual allocation: Sys admin decides which IP addresses will be assigned to each client

Leases

- Used for dynamic allocation
- Solution for control of when an address can be given to another client
  - Because most clients won't tell the server when they disappear
- Server allows use of addr for a set period
- Client will need to reacquire permission before lease period expires

DHCP Client

- How does a client who doesn't know anything about the network (like its own IP address) send messages to a DHCP server?

DHCP Messages

- Fields include
  - type: discover, offer, request, ack, release
  - xid: Random transaction value
  - chaddr: client hardware identifier
    - MAC or other opaque key
  - siaddr: server's IP address
  - yiaddr: "your" address
  - options: lots of optional parameters

Discovery Process

- All communication is broadcast
  - Thus the funky arrows in our sequence diagram
- Multiple servers may respond
  - Client chooses whichever offer it wishes
- DHCP Request / Ack is repeated to renew a lease

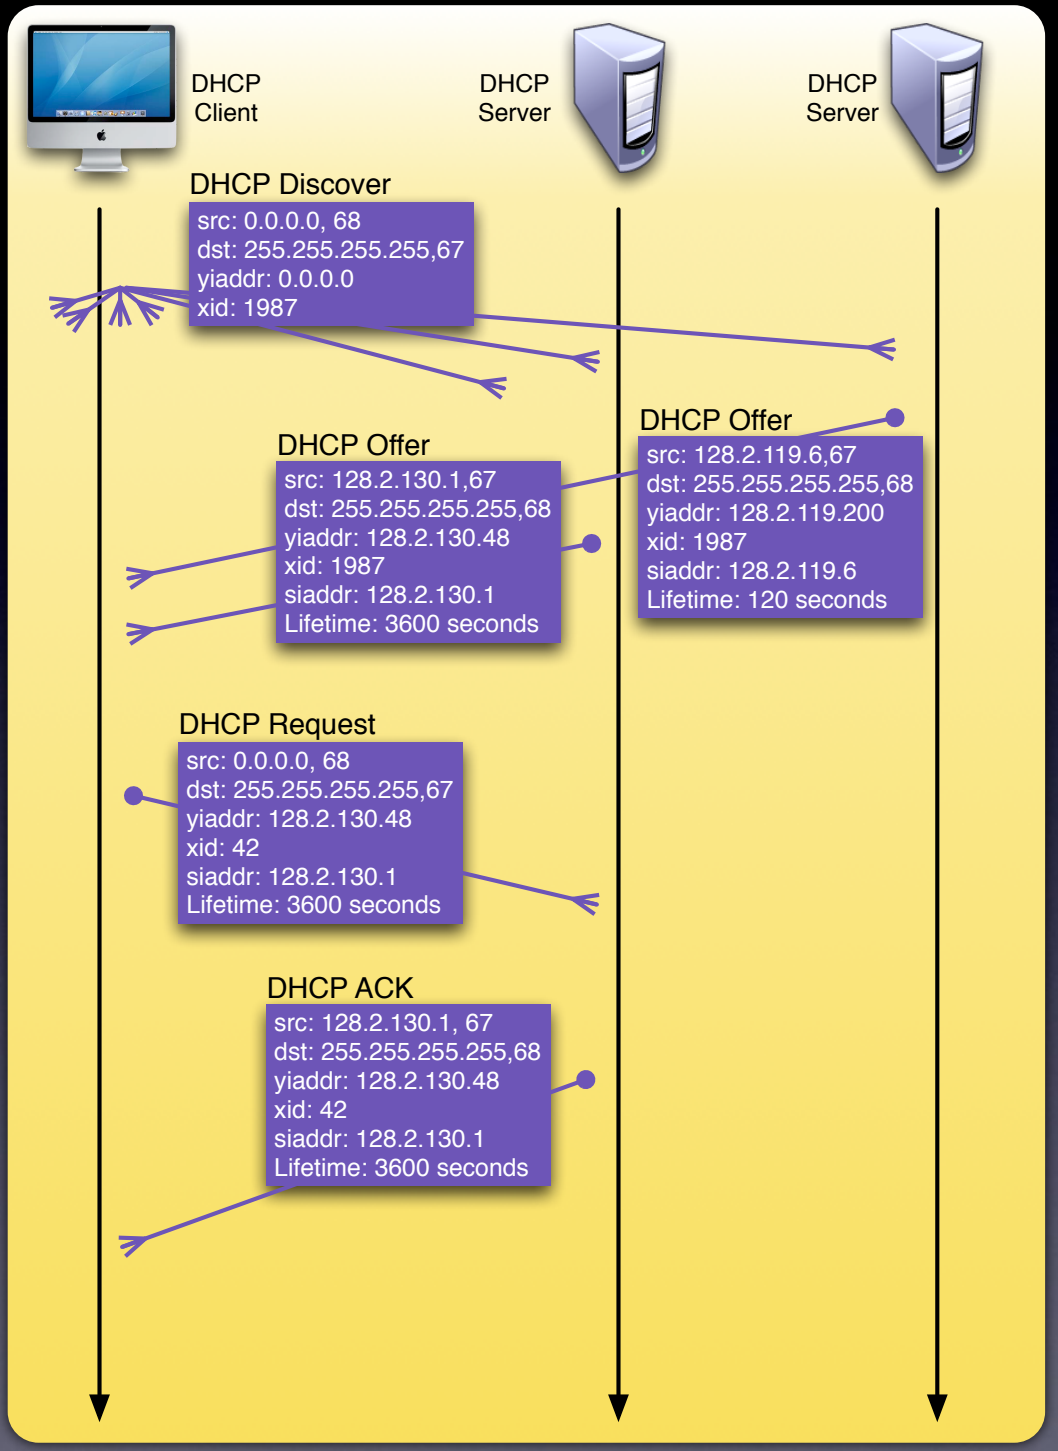

Security or lack thereof

- Unauthorized server
  - Can get client to use your configuration values (i.e. malicious DNS server)
- Unauthorized client
  - Can get access to the network
  - Can get server to exhaust IP address pool, and thus DOS the subnet

The Bottom Line

- DHCP is critical for “plug and play” network capabilities
- Saves administrative hassles
- ... and errors from mis-configured values
  - Finding which host is re-using an IP address can be a struggle, for instance

traceroute

- DHCP
- NAT
- IPv6

Network Address Translation

- Mechanism to simplify IP address allocation
- Basic idea: Router appears as a single IP address to the world, but manages a complete subnet with many hosts
  - Maps one address space into another
- Also called IP masquerading or IP spoofing

NAT Benefits

- Work-around to the impending exhaustion of IP addresses
  - Entire networks can operate with an allocation of just a single IP address
- Also allows for simple address allocation for the subnet
  - No need to contact the ISP to add additional end hosts
- “Security”: internal network structure obscured

Mechanics

- Hosts on private network use “non-routable” IP addresses
- Defined in RFC 1918
  - 10.0.0.0/8 prefix
  - 172.16.0.0/12 prefix
  - 192.168.0.0/16 prefix
- Packets restricted to the private subnet
  - Why?

xkcd.com/742

Mechanics (2)

- Router shows a single external IP address
- Translation table maps external IP / port combinations to internal IP / port
- Rewrites all packets in each direction, changing IP / port based on translation table
  - Other fixes also needed to the packet

Example

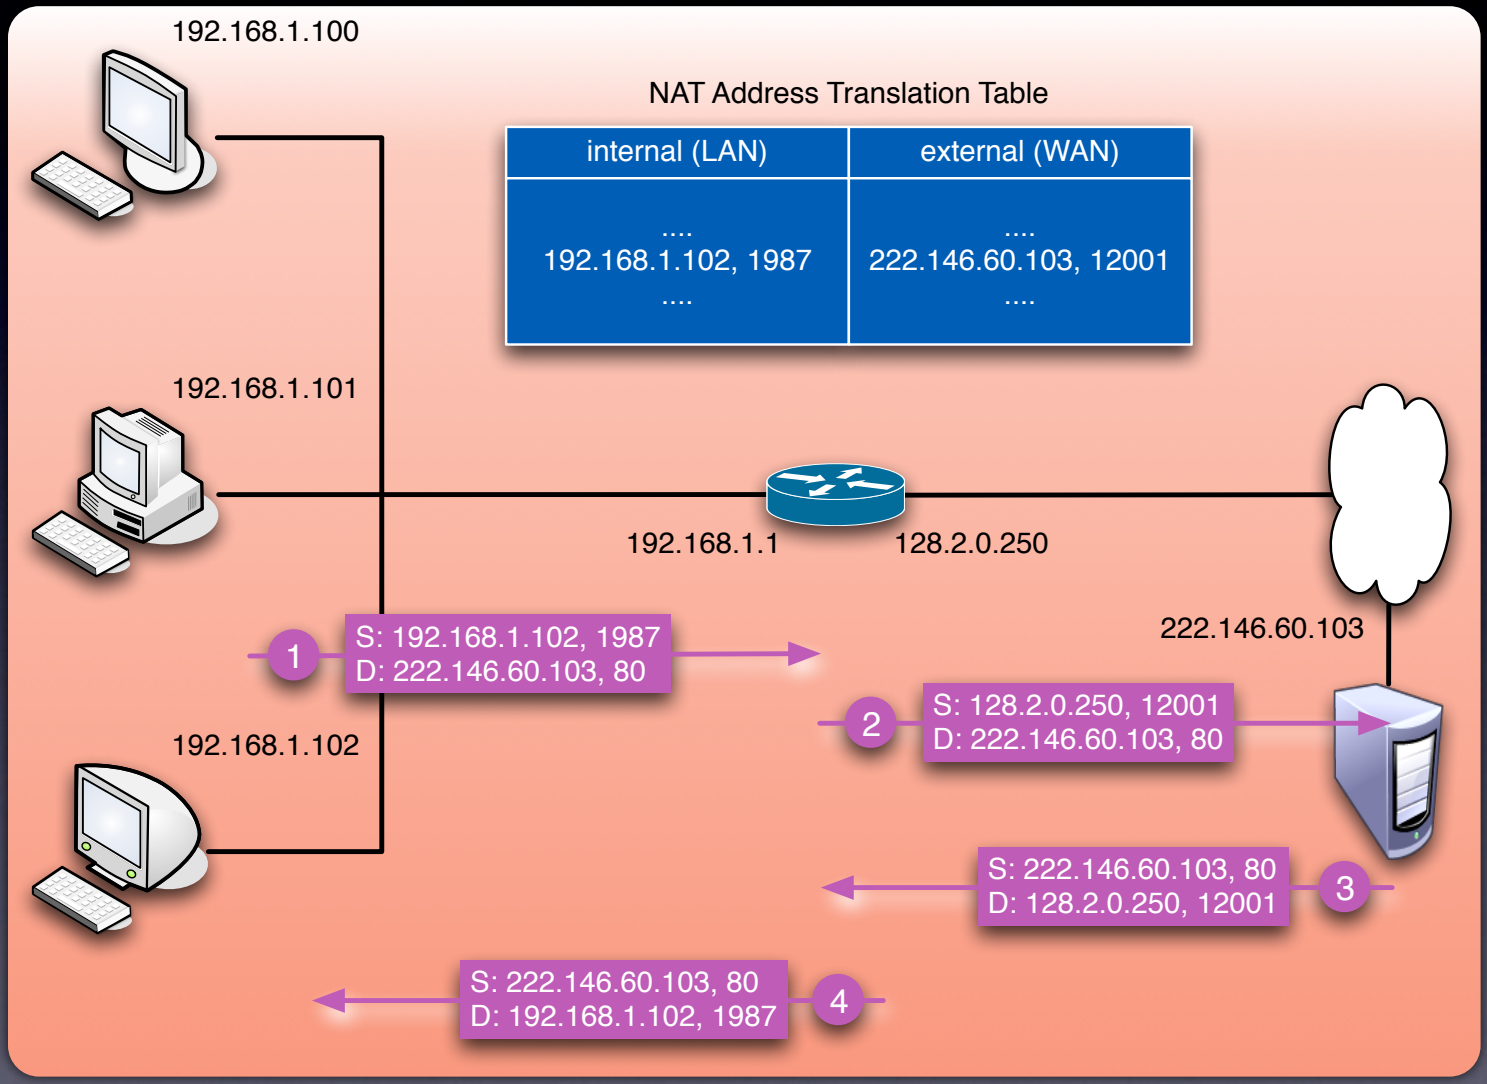

Operations vary based on contents of table
-- other possibilities exist!

NAT Versions

- Depending on contents of the table, may get differing effects
  - Basic NAT: IP address translation only
  - Port Address Trans: IP and port translations
  - Source NAT: rewrites sender's IP/Port
  - Dest NAT: rewrites destination's IP/Port
  - Symmetric NAT: mapping corresponds to {send, rcvr} pair
    - Requests sent from same sender but to different destinations get different mappings ➙ only an external receiver can reply

Packet Fix-up

- Router must do more than simply change address/port values
- Fix checksums
- Some application protocols need fixing
  - FTP and SIP send IP/port values in the data stream of the control channel
  - Must reassemble fragmented packets
  - Especially problematic if encryption has been applied

Port Forwarding

- Translation table is normally initialized by internal traffic
  - Which means no external host can initiate communication
  - One solution: connection reversal, involves ongoing communication with external server
- Port forwarding specifies values ahead of time
  - Example: BitTorrent traffic will be handled by a particular internal host with addr intIP, so forward all traffic to extIP:6881 to intIP:6881

Purists love this, right?

- Objection 1: IPv6 should be used to solve addressing problem
  - Believe NAT has staved off adoption of IPv6
- Objection 2: Violates end-to-end principle
- Objection 3: Routers shouldn’t process packets higher than network layer
- Objection 4: Using port numbers to address hosts

IPv6

- Purpose
- Addressing
- Autoconfiguration

Purpose: More Addresses

- Early 90s: "IP addr exhaustion coming"
  - How can that be? 2^{32} is 4 Billion hosts?
- 1996: 100% Class A, 62% Class B, 37% Class C addresses assigned
- 2011: IANA allocated last remaining blocks to regional registries

While we're at it...

- Bigger addresses mean breaking the IPv4 header, so may as well do some other stuff at the same time...
  - Streamline the header
  - Improve option processing
  - Self-configuration
- BTW, breaking the header format means everyone (all hosts and all routers) needs to change -- leading to slow adoption rate

IPv6 Addresses

- 128 bit addresses
  - 340 billion, billion, billion, billion addrs
    - $2^{128} = 3.4 \times 10^{38}$
  - That should be enough for a while

xkcd.com/865/

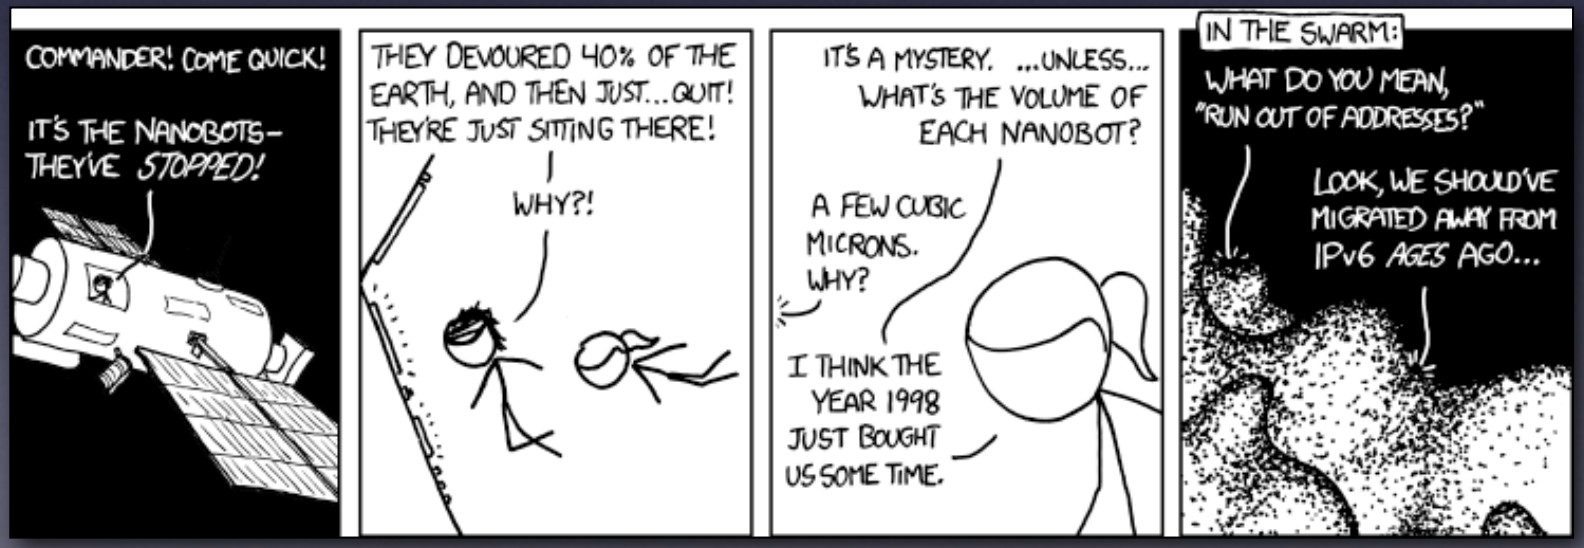

Address Notation

- Write IPv6 addresses using 4 hex digit groups, separated by colons

1987:A456:2B2B:1234:BEEF:5678:CAFE:D82F

- Leading zeros are dropped, largest string of contiguous zeros are not written

0124:A245:0000:0000:0000:0000:0001:0023

124:A245::1:23

Address Classes

- IPv6 Addresses are classless, like CIDR
- Some addresses are special
  - :: (all zeros) is special (unspecified)
  - ::1 is for loopback
  - Starts with 1111 1111 → multicast
  - Starting with 1111 1110 10 → link-local
    - more on this in a minute

Header Changes

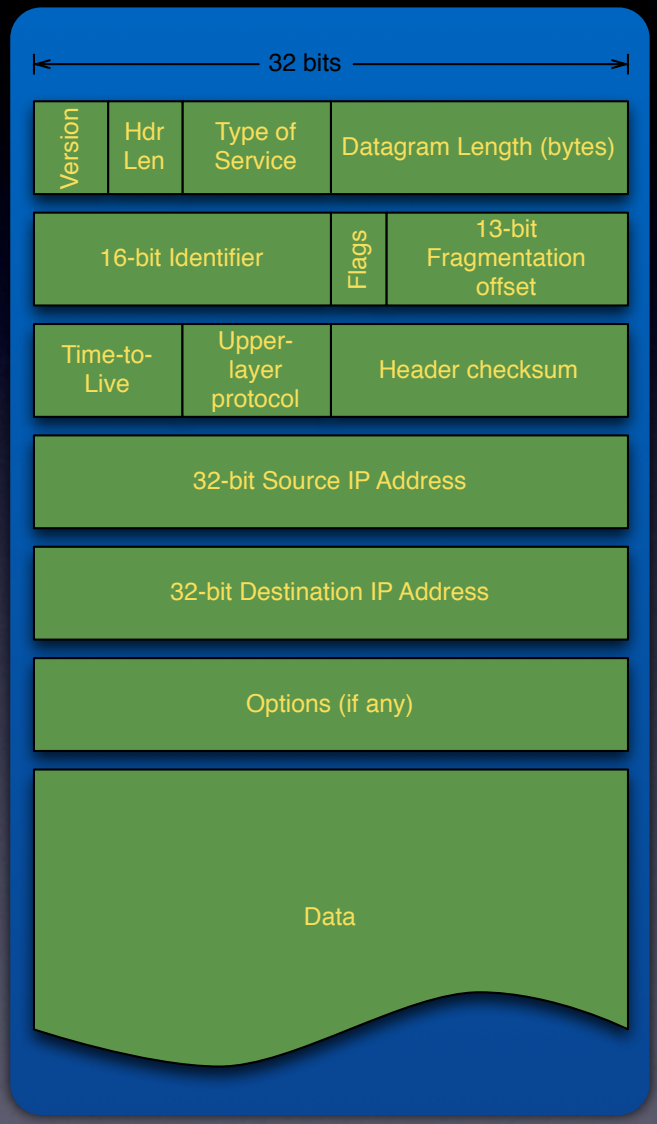

IPv4 Header

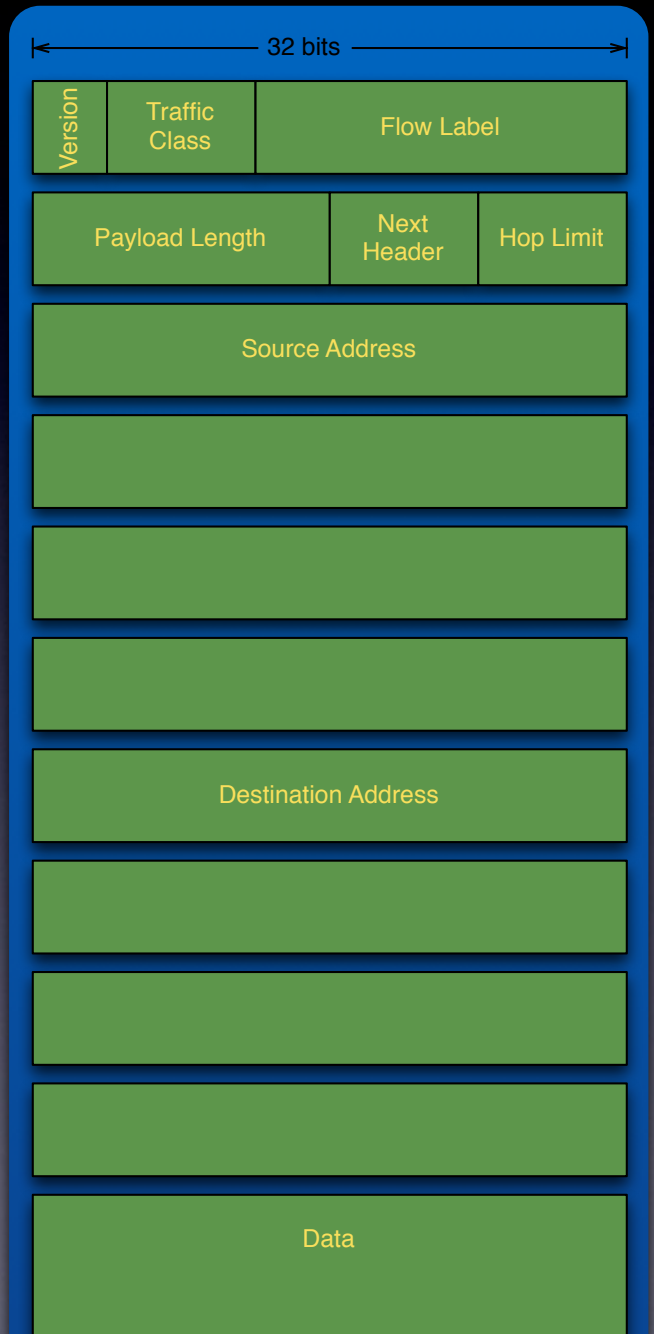

IPv6 Header

- Version → 0110 (i.e. 6)
- Traffic Class, Flow Label
  - Used to distinguish (give priority to) "flows" of data
  - No definition of how to use
- Payload Length (in bytes)

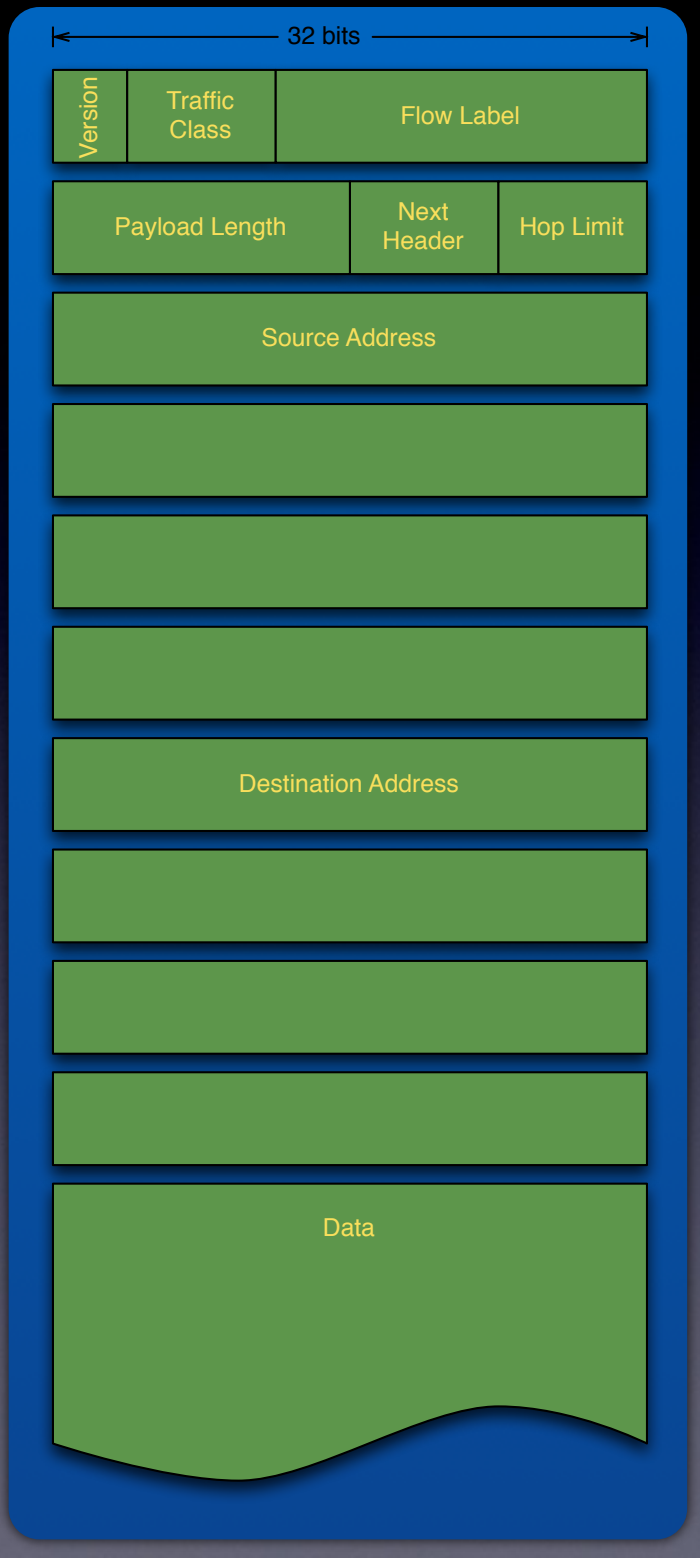

IPv6 Header

- Next Header: specifies the type of the data
  - Is it UDP? → 17 TCP? → 6
    - Same values and purpose as IPv4's Upper Layer Protocol field
  - Or, is it an options field?
    - Specifies the type of an *options header* placed at the beginning of payload
    - Options headers may be chained, last one specifies next header value for payload

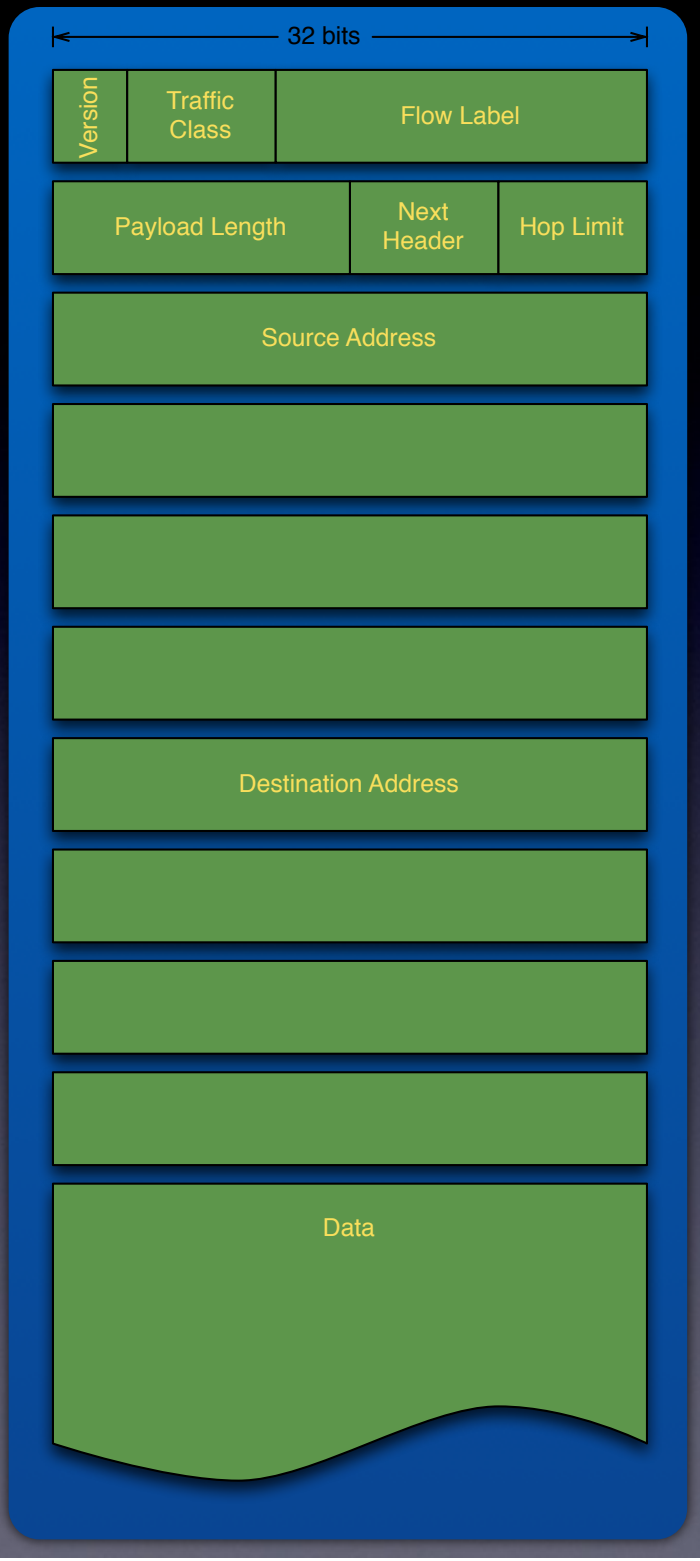

IPv6 Header

- Hop Limit: Renamed version of IPv4's TTL
- Src, Dest Addresses
  - 128 bit, so they take up a large portion of header
- Total header size: 40 bytes

IPv6

- Purpose
- Addressing
- Autoconfiguration

Configuration

- Process of getting an IP address (and other data) to a newly booted end host
- IP address needs to be unique
  - Often has no other requirement
    - A printer, for instance
- Autoconfiguration is then possible

Autoconfiguration

- Host wishing configuration can choose an IPv6 address for themselves...
  - ... as long as they can guarantee it is unique
- But, most hosts already have a "guaranteed to be unique number" built into their link-layer
  - Ethernet MAC address, for instance

IPv6 Autoconfiguration

- Stateless Autoconfiguration
  - Use Link-local address prefix (1111 1110 10), ...
  - ... followed by some zeros, ...
  - ... followed by Ethernet MAC (48 bits)
- Can be followed with Neighbor Discovery Protocol (NDP) broadcast messages (RFC 4861)
  - to find network prefix IP addresses, next-hop routers, DNS servers, etc
- Stateful Configuration: DHCPv6 exists (RFC 3315)

Lesson Objectives

- Now, you should be able to:
  - describe DHCP, including information carried, methods of communication, leases, message format and the discovery process
  - describe NAT, including benefits / objections, operations and port forwarding
  - describe situations where NAT needs to modify values other than IP address and TCP/UDP port fields

- You should be able to:
  - describe IPv6, including differences with IPv4, benefits, datagram format, and address notation
  - describe IPv6 address autoconfiguration