IP Addressing (cont'd)

* Some slides are adapted from *Computer Networking: A Top-Down Approach.* J.F Kurose and K.W. Ross. All Rights Reserved.

Learning Objectives

- Recap
- IPv6
  - think like a protocol designer
- NAT
- Middleboxes

© CS 438 Staff, University of Illinois

Where we left off — IP

Internet's "thin waist":

- one network layer protocol: IP
- must be implemented by every (billions) of Internet-connected devices

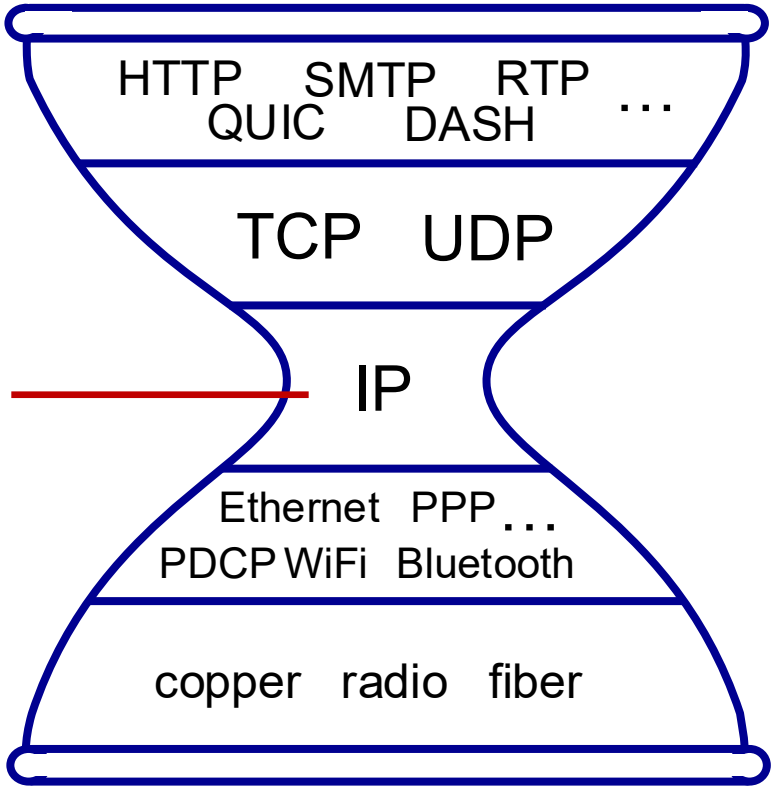

© CS 438 Staff, University of Illinois

Connecting the dots...

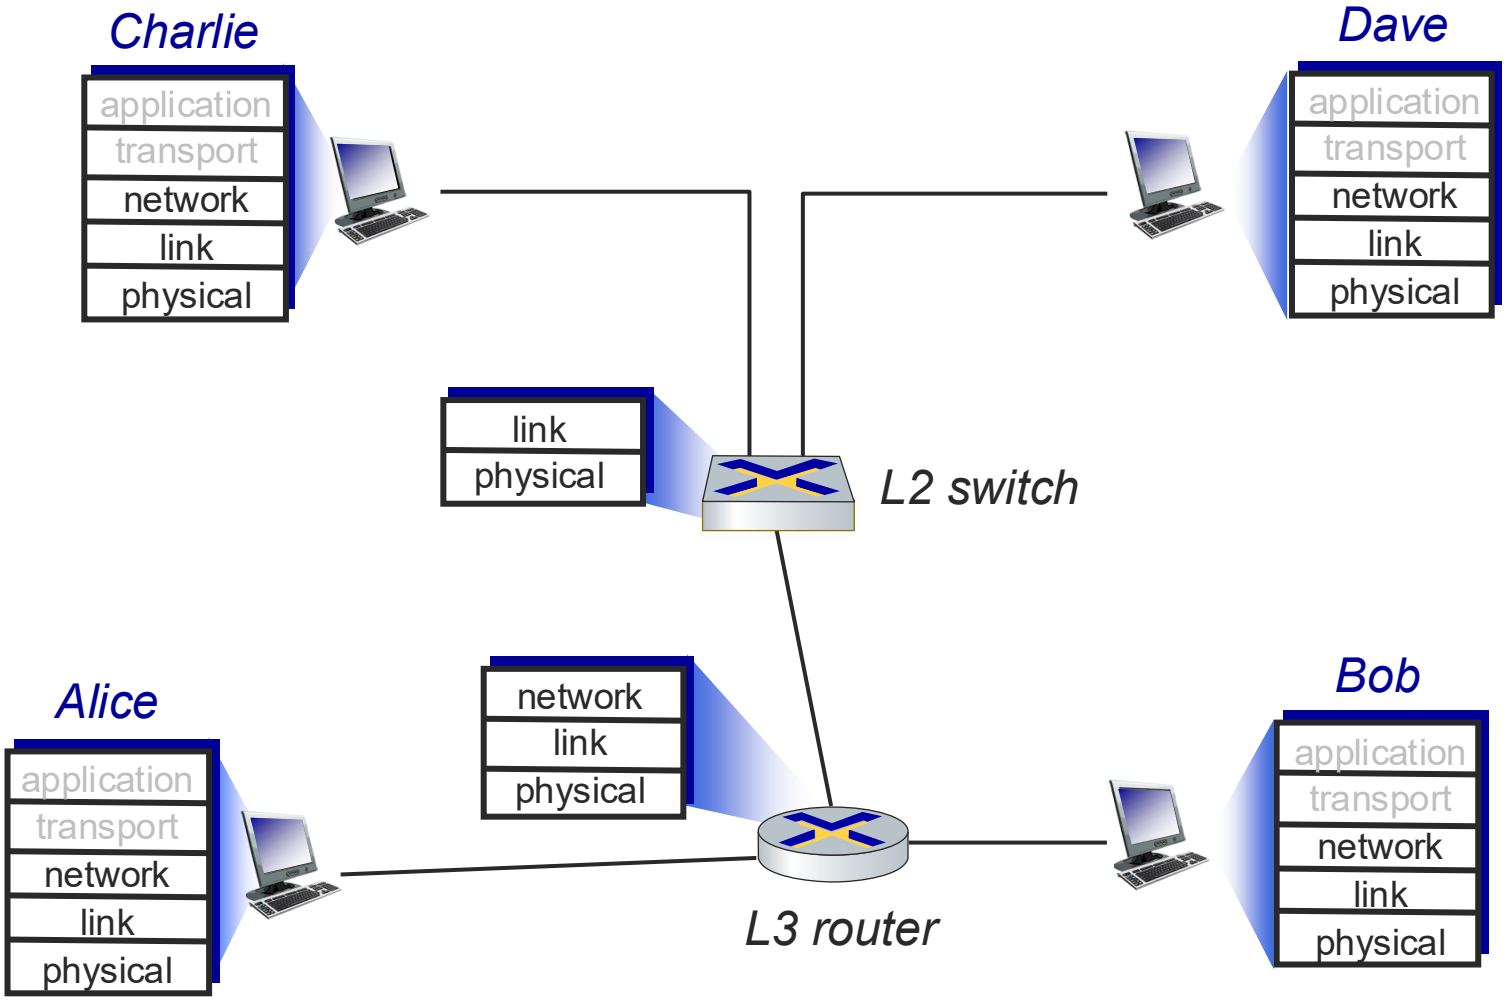

Connecting the dots...

- Network layer
  - *end-to-end delivery*
  - IP addressing, routing, forwarding
- Link layer
  - *a single link (or LAN)*
  - MAC addressing, framing, error detection, medium access control, switching
- Physical layer
  - *over the medium*
  - encoding, modulation

© CS 438 Staff, University of Illinois

IPv6 Motivation

- Initial motivation:
  - 32-bit IPv4 address space would be completely allocated
- Additional motivation:
  - speed processing/forwarding: fixed length header
  - enable different network-layer treatment of “flows”

© CS 438 Staff, University of Illinois

IPv6 Addresses

- 128-bit
  - 3.4 x 10^{38} addresses (as compared to 4 x 10^{9} in IPv4)
- Uses CIDR (prefix lengths, longest prefix match)
- Address notation
  - String of eight 16-bit hex values separated by colons
    - 5CFA:0002:0000:0000:CF07:1234:5678:FFCD
  - Set of contiguous 0's can be elided
    - 5CFA:0002::0000:CF07:1234:5678:FFCD
- Address allocation
  - Provider-based
  - Geographic

© CS 438 Staff, University of Illinois

IPv6 Addresses

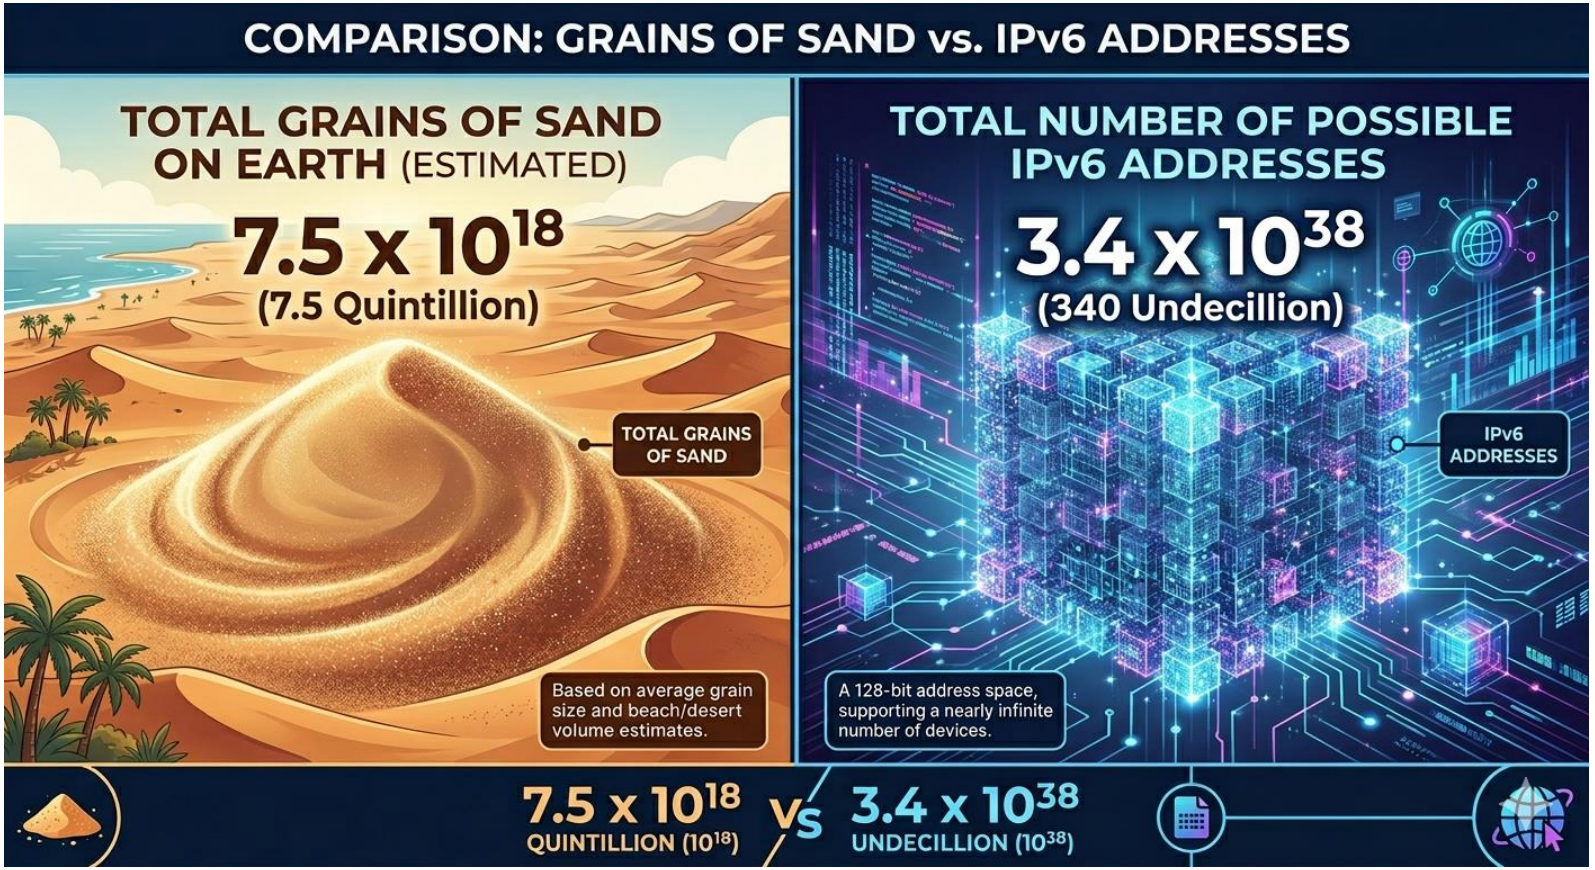

Illustration generated by Gemini

© CS 438 Staff, University of Illinois

IPv4 Packet Format

- 20 bytes minimum
- Mandatory fields are not always used
  - e.g., fragmentation
- Options are an unordered list of (name, value) pairs

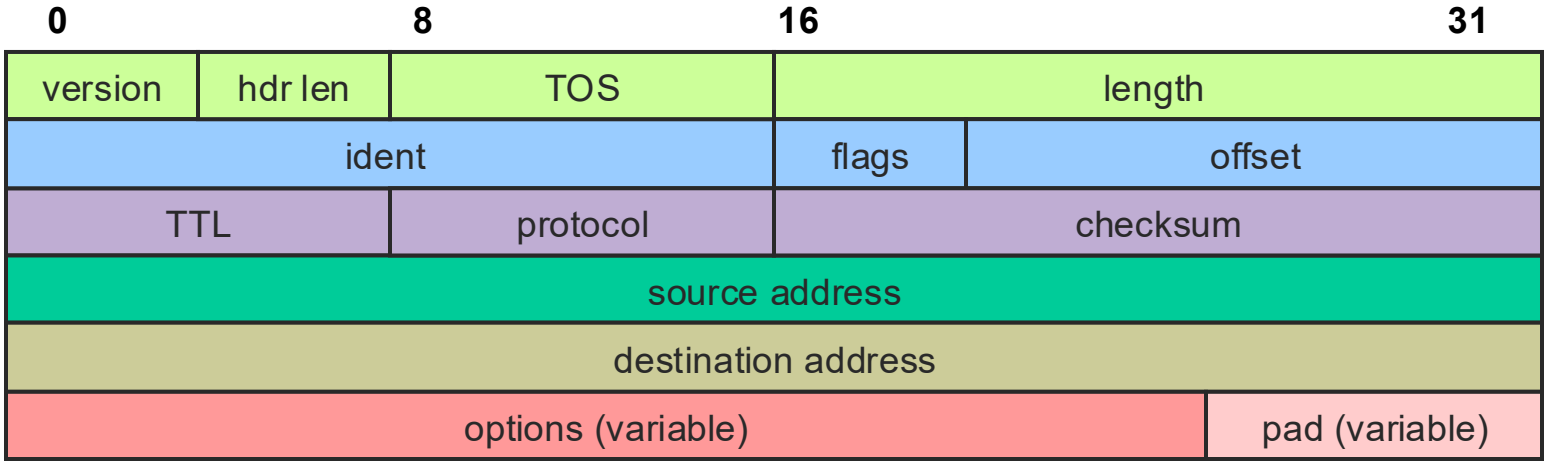

© CS 438 Staff, University of Illinois

IPv6 Packet Format

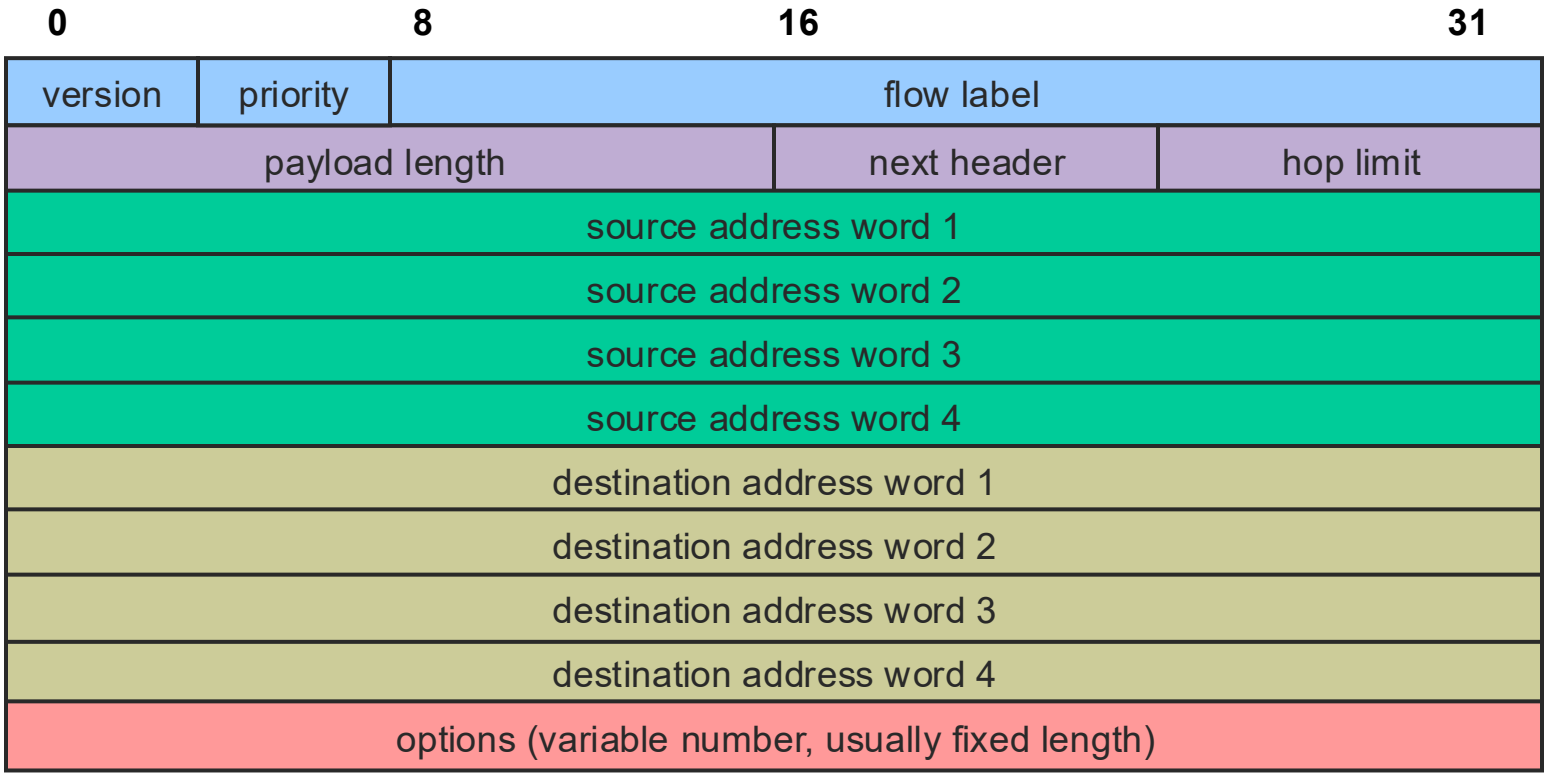

© CS 438 Staff, University of Illinois

IPv6 Packet Format

- 40 bytes minimum
- Mandatory fields (almost) always used
- Strict order on options reduces processing time
  - No need to parse irrelevant options

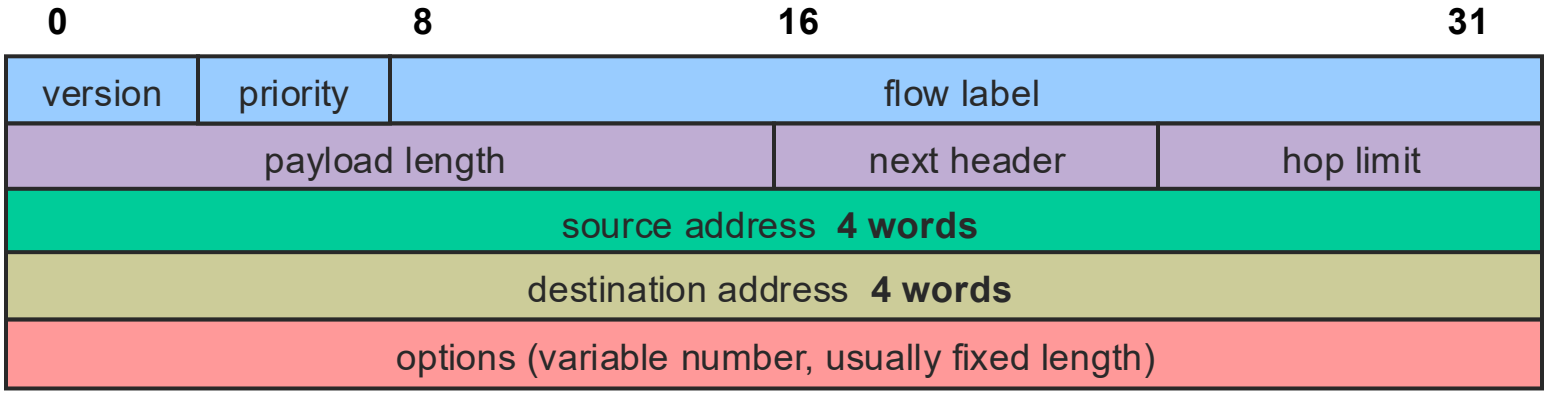

© CS 438 Staff, University of Illinois

IPv6 Packet Format

- Version
  - 6
- Priority and Flow Label
  - Support service guarantees
  - Allow “fair” bandwidth allocation
- Payload Length
  - Header not included
- Next Header
  - Combines options and protocol
  - Chain of extension headers
  - Ends with higher-level protocol header (e.g., TCP)
- Hop Limit
  - TTL renamed to match usage

© CS 438 Staff, University of Illinois

IPv6 Extension Headers

- Must appear in order
  - Hop-by-hop options
    - Miscellaneous information for routers
  - Routing
    - Full/partial route to follow
  - Fragmentation
    - IP fragmentation info
  - Authentication
    - Sender identification
  - Encrypted security payload
    - Information about contents
  - Destination options
    - Information for destination

© CS 438 Staff, University of Illinois

IPv6 Extension Headers

- Hop-by-Hop extension
  - Length is in bytes beyond mandatory 8

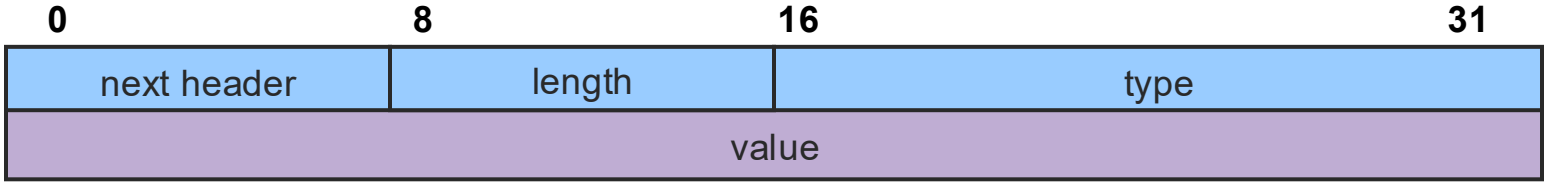

  - Example: Jumbogram option (packet longer than 65,535 bytes)
  - Payload length in main header set to 0

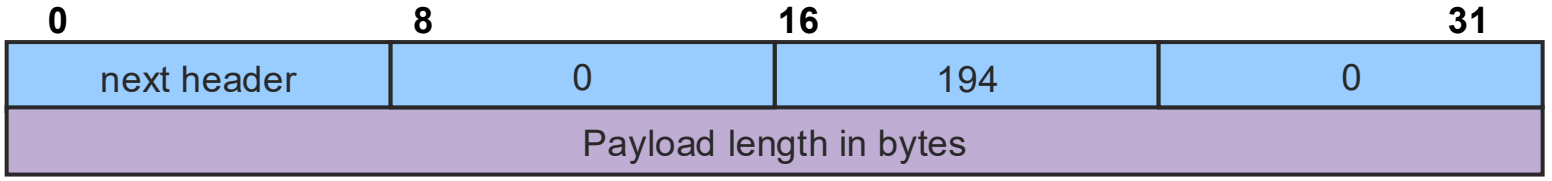

© CS 438 Staff, University of Illinois

IPv6 Extension Headers

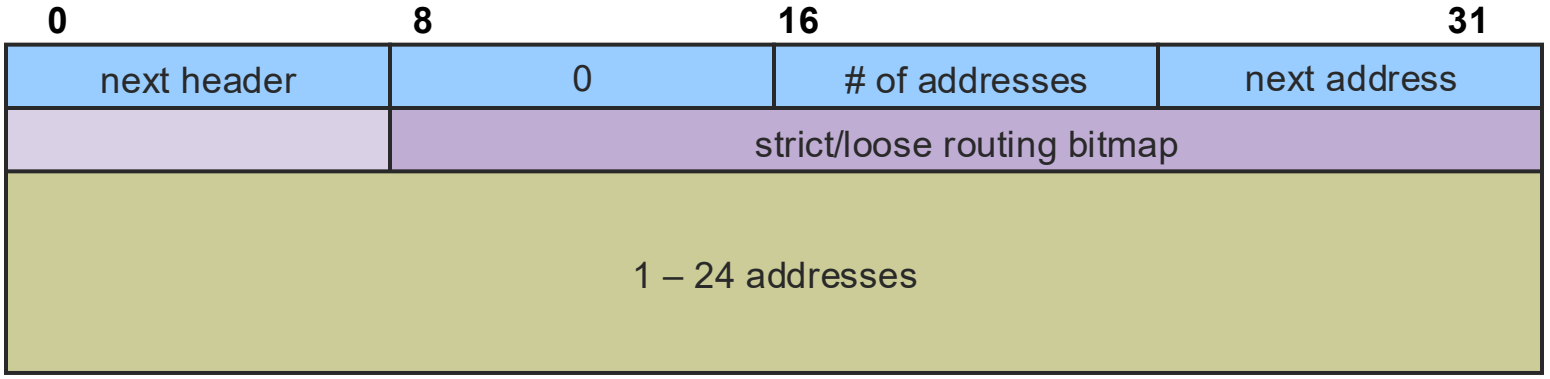

- Routing extension
  - Up to 24 “anycast” addresses target AS’s/providers
  - Next address tracks current target
  - Strict routing requires direct link
  - Loose routing allows intermediate nodes

© CS 438 Staff, University of Illinois

IPv6 Extension Headers

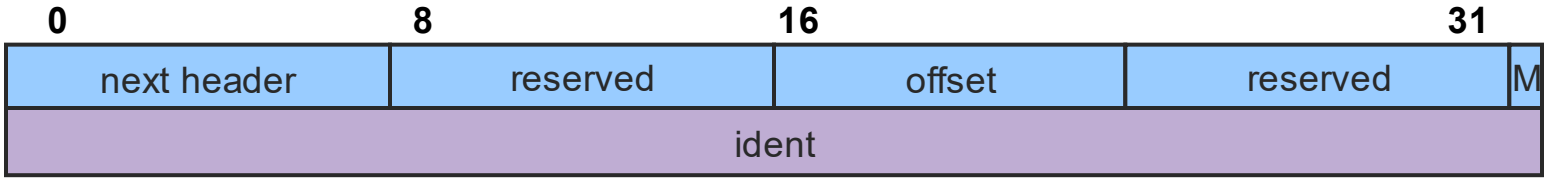

- Fragmentation extension
  - Similar to IPv4 fragmentation
    - 13-bit offset
    - Last fragment mark (M)
  - Larger fragment identification field

© CS 438 Staff, University of Illinois

IPv6 Extension Headers

- Authentication extension
  - Designed to be very flexible
  - Includes
    - Security parameters index (SPI)
    - Authentication data
- Encryption Extension
  - Called encapsulating security payload (ESP)
  - Includes an SPI
  - All headers and data after ESP are encrypted

© CS 438 Staff, University of Illinois

IPv6 Design Controversies

- Address length
  - 8 byte
    - Might run out in a few decades
    - Less header overhead
  - 16 byte
    - More overhead
    - Good for foreseeable future
  - 20 byte
    - Even more overhead
  - Variable length
    - Complexity in parsing, etc.

© CS 438 Staff, University of Illinois

IPv6 Design Controversies

- Hop limit
  - 65,535
    - 32 hop paths are common now
    - In a decade, we may see much longer paths
  - 255
    - Objective is to limit looping packet lifetime
    - Good network design makes long paths unlikely
      - Source to backbone
      - Across backbone
      - Backbone to destination

© CS 438 Staff, University of Illinois

IPv6 Design Controversies

- Greater than 64 KB data
  - Good for supercomputer/high bandwidth applications
  - 1 MB packet ties up a 1.5Mbps line for more than 5 seconds
  - Inconveniences interactive users
- 64 KB data
  - More compatible with low-bandwidth lines
  - More fragments => higher header overhead

© CS 438 Staff, University of Illinois

IPv6 Design Controversies

- Keep checksum
  - Light and faster
  - Unprepared for the unexpected
- Remove checksum
  - Typically duplicated in data link and transport layers
  - Very expensive in IPv4

© CS 438 Staff, University of Illinois

IPv6 Design Controversies

- Mobile hosts
  - Direct or indirect connectivity
    - Reconnect directly using canonical address
    - Use home and foreign agents to forward traffic
  - Mobility introduces asymmetry
    - Base station signal is strong, heard by mobile units
    - Mobile unit signal is weak and susceptible to interference, may not be heard by base station

© CS 438 Staff, University of Illinois

IPv6 Design Controversies

- Security
  - Where?
    - Network layer
      - A standard service
    - Application layer
      - No viable standard
      - Application susceptible to errors in network implementation
      - Expensive to turn on and off
  - How?
    - Political import/export issues
    - Cryptographic strength issues

© CS 438 Staff, University of Illinois

Transition from IPv4 to IPv6

- not all routers can be upgraded simultaneously
  - no “flag days”
  - how will network operate with mixed IPv4 & IPv6 routers?
- tunneling
  - IPv6 datagram carried as payload in IPv4 datagram among IPv4 routers (“packet within a packet”)
  - tunneling used extensively in other contexts (4G/5G)

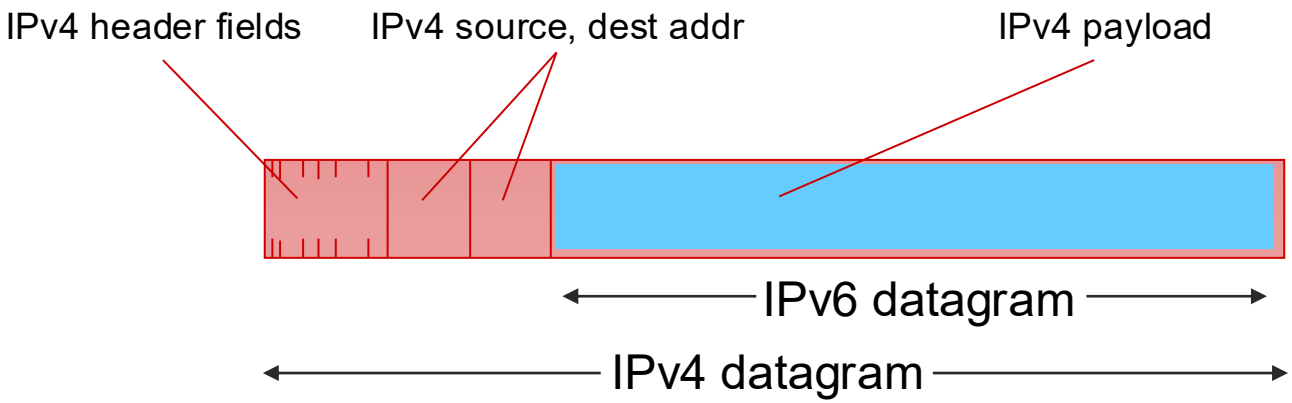

IPv6: adoption

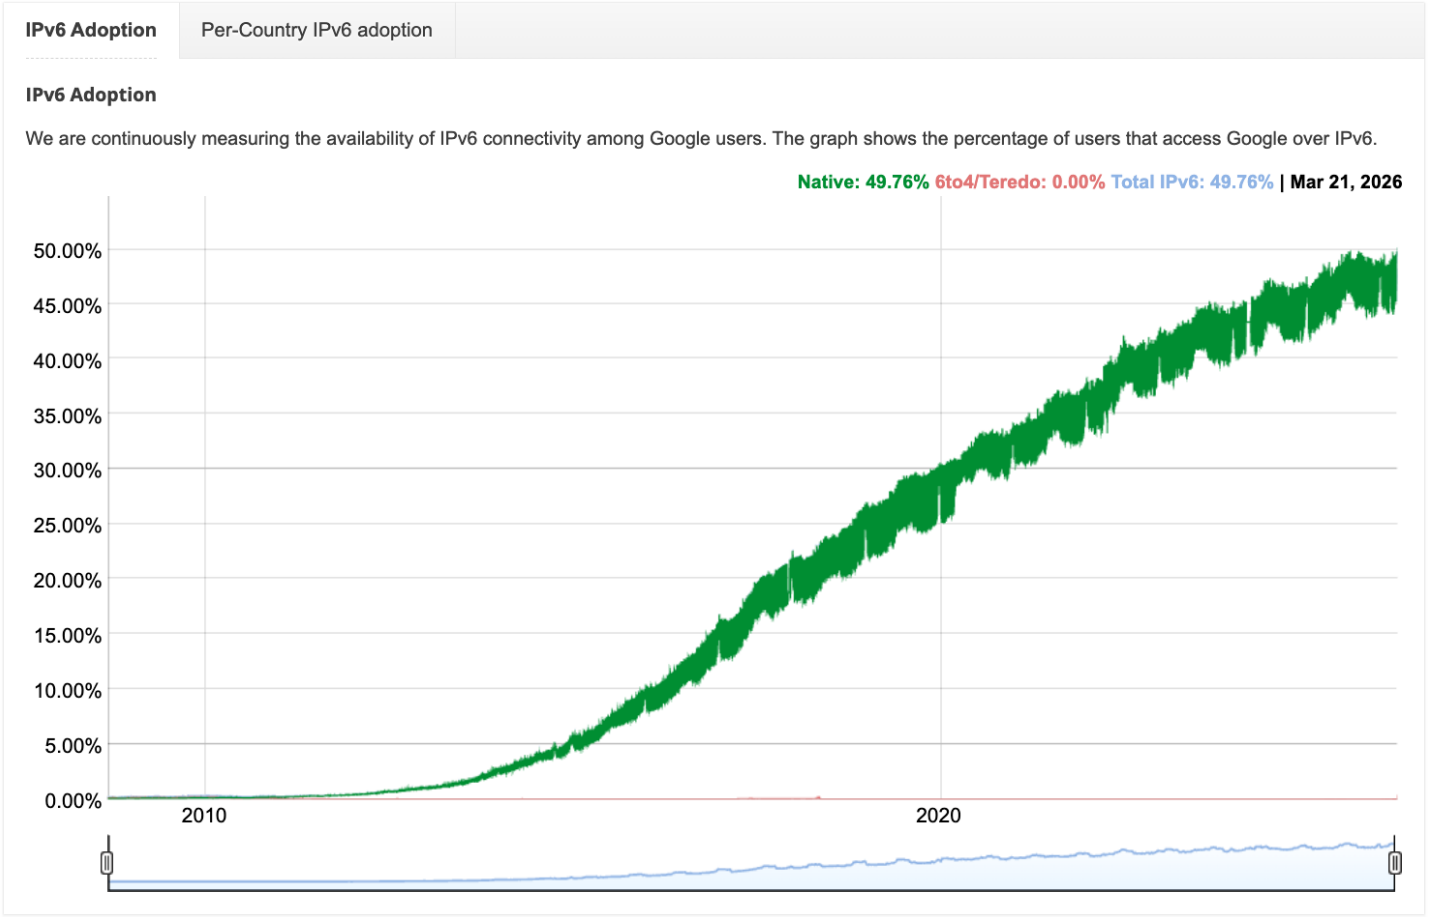

https://www.google.com/intl/en/ipv6/statistics.html

© CS 438 Staff, University of Illinois

Network Address Translation (NAT)

- Kludge but useful
- Sits between your network and the Internet
- Translates local (private) IP addresses to global (public) IP addresses
- Has a pool of global IP addresses (less than number of hosts on your network)

© CS 438 Staff, University of Illinois

NAT Illustration

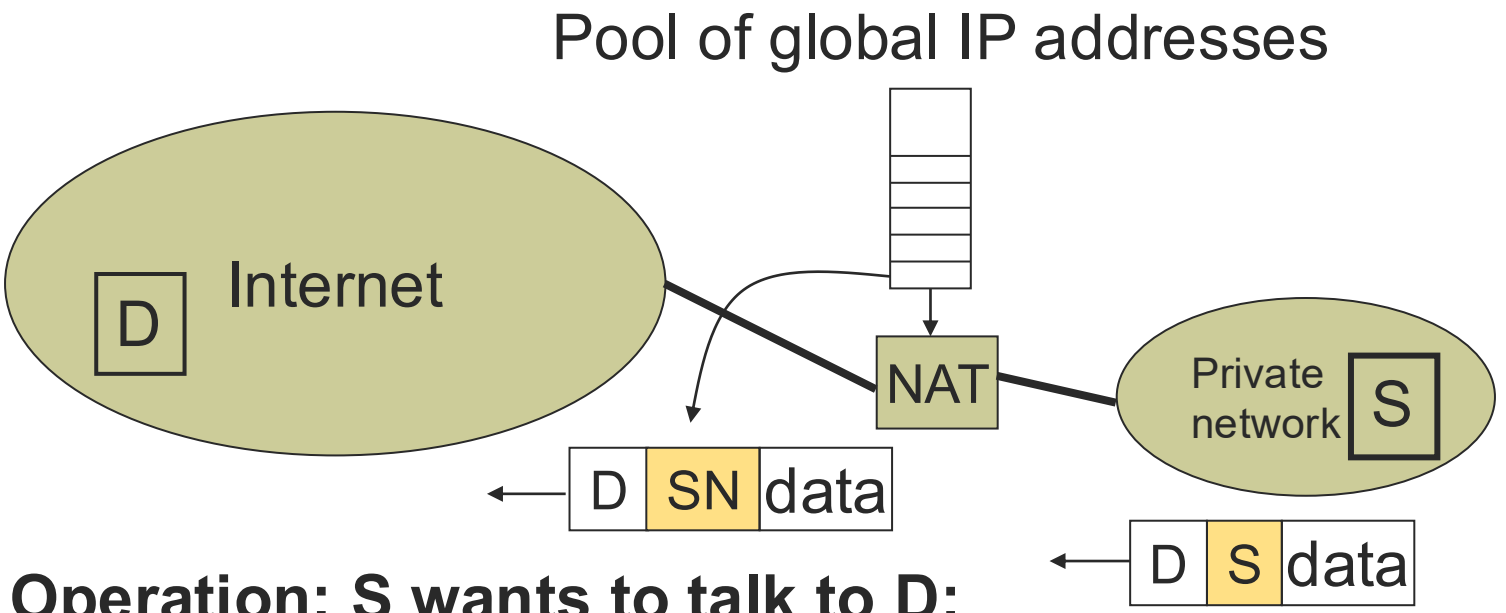

Operation: S wants to talk to D:

- Create S-SN mapping
- Replace S with SN for outgoing packets
- Replace SN with S for incoming packets

© CS 438 Staff, University of Illinois

What if we only have few (or just one) global IP addresses?

- Approach: use address + port
  - Assign one router a global IP address
  - Assign internal hosts local IP addresses
  - <PrivateAddr, PrivatePort> ⇔ <GlobalAddr, GlobalPort>
- Change IP Headers
  - IP addresses and possibly port numbers of IP datagrams are replaced at the boundary of a private network
  - Enables hosts on private networks to communicate with hosts on the Internet
  - Run on routers that connect private networks to the public Internet

© CS 438 Staff, University of Illinois

NAT

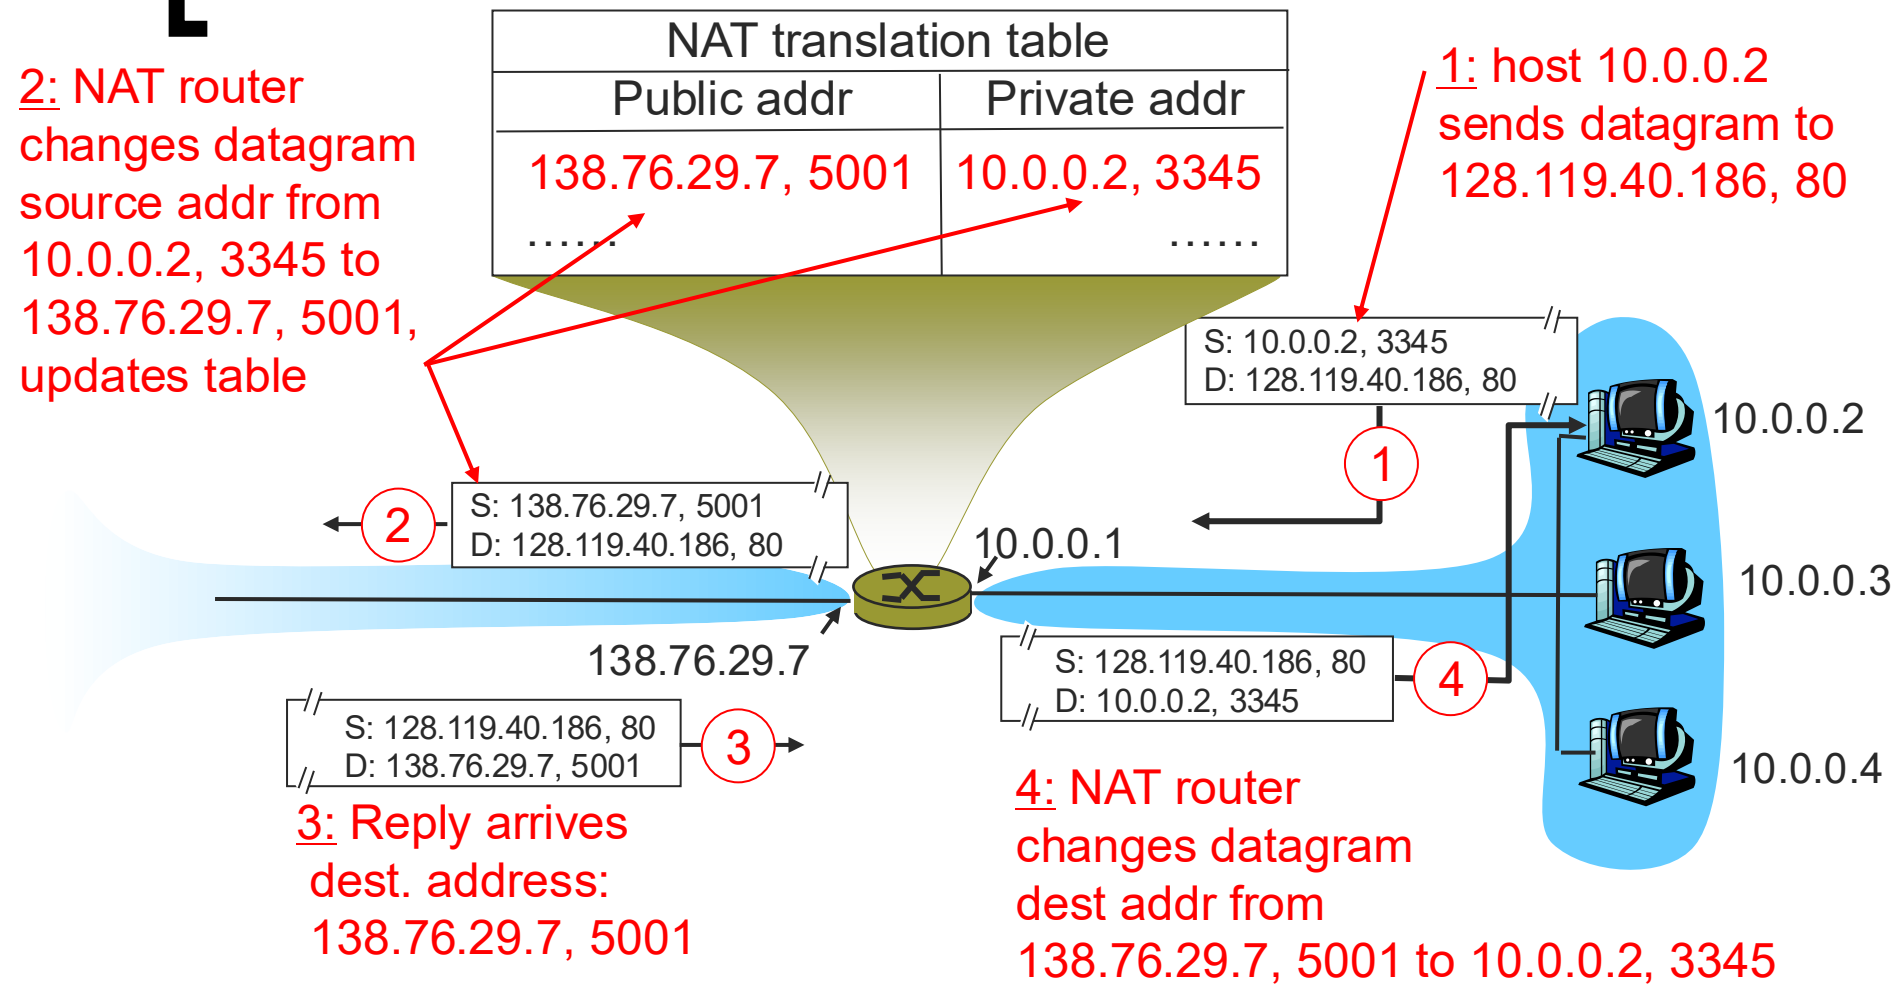

© CS 438 Staff, University of Illinois

Benefits of NAT

- Local network uses just one (or a few) IP address as far as outside world is concerned
  - No need to be allocated range of addresses from ISP
    - Just one IP address is used for all devices
  - Can change addresses of devices in local network without notifying outside world
  - Can change ISP without changing addresses of devices in local network
  - Devices inside local net not explicitly addressable, visible by outside world (a security plus)

© CS 438 Staff, University of Illinois

Use cases

- Address Pooling
  - Corporate network has many hosts
  - Only a small number of public IP addresses
- NAT solution
  - Manage corporate network with a private address space
  - NAT, at boundary between corporate network and public Internet, manages a pool of public IP addresses
  - When a host from corporate network sends an IP datagram to a host in public Internet, NAT picks a public IP address from the address pool, and binds this address to the private address of the host

© CS 438 Staff, University of Illinois

Use cases

- Load balancing
  - Balance the load on a set of identical servers, which are accessible from a single IP address
- NAT solution
  - Servers are assigned private addresses
  - NAT acts as a proxy for requests to the server from the public network
  - NAT changes the destination IP address of arriving packets to one of the private addresses for a server
  - Balances load on the servers by assigning addresses in a round-robin fashion

© CS 438 Staff, University of Illinois

Problems with NAT?

© CS 438 Staff, University of Illinois

NAT: Consequences

- End-to-end connectivity
  - NAT destroys universal end-to-end reachability of hosts on the Internet
  - A host in the public Internet often cannot initiate communication to a host in a private network
  - The problem is worse, when two hosts that are in different private networks need to communicate with each other
    - *How is it possible to have video calls then?*

© CS 438 Staff, University of Illinois

NAT: Consequences

- 16-bit port-number field
  - Only ~60,000 simultaneous connections with a single LAN-side address
- Performance
  - Modifying the IP header (address) requires that NAT boxes recalculate the IP header checksum
  - Modifying port number requires that NAT boxes recalculate TCP checksum
    - UDP checksum is optional in IPv4 (but mandatory in IPv6)
- Fragmentation
  - Datagrams fragmented before NAT device must not be assigned different IP addresses or different port numbers

© CS 438 Staff, University of Illinois

NAT: Consequences

- IP address in application data
  - Applications often carry IP addresses in the payload of the application data (e.g., FTP)
  - No longer work across a private-public network boundary
  - Hack: Some NAT devices inspect the payload of widely used application layer protocols and, if an IP address is detected in the application-layer header or the application payload, translate the address according to the address translation table

© CS 438 Staff, University of Illinois

NAT controversies

- NAT has been controversial:
  - port # manipulation by network-layer devices
  - routers “should” only process up to layer 3
  - address “shortage” should be solved by IPv6
  - violates end-to-end argument
- but NAT is here to stay:
  - extensively used in home and institutional networks, 4G/5G cellular networks

© CS 438 Staff, University of Illinois

Middleboxes

Middlebox (RFC 3234)

"any intermediary box performing functions apart from normal, standard functions of an IP router on the data path between a source host and destination host"

© CS 438 Staff, University of Illinois

Middleboxes everywhere!

NAT:
home, cellular, institutional

Proxies:
ISPs, institutional, CDNs

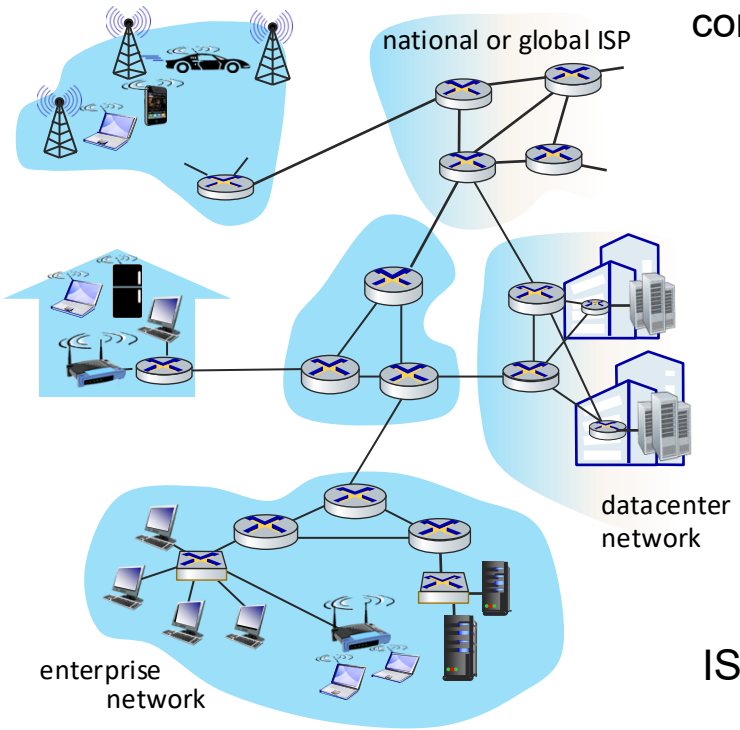

Firewalls, IDS:
corporate, institutional, ISPs

Load balancers:
corporate, ISPs, data center, mobile nets

Caches:
ISPs, mobile, CDNs

Middleboxes

- initially: proprietary (closed) hardware solutions
- move towards "whitebox" hardware implementing open API
  - move away from proprietary hardware solutions
  - programmable local actions via match+action
  - move towards innovation/differentiation in software
- network function virtualization (NFV)
  - programmable services over white box networking, computation, storage

© CS 438 Staff, University of Illinois