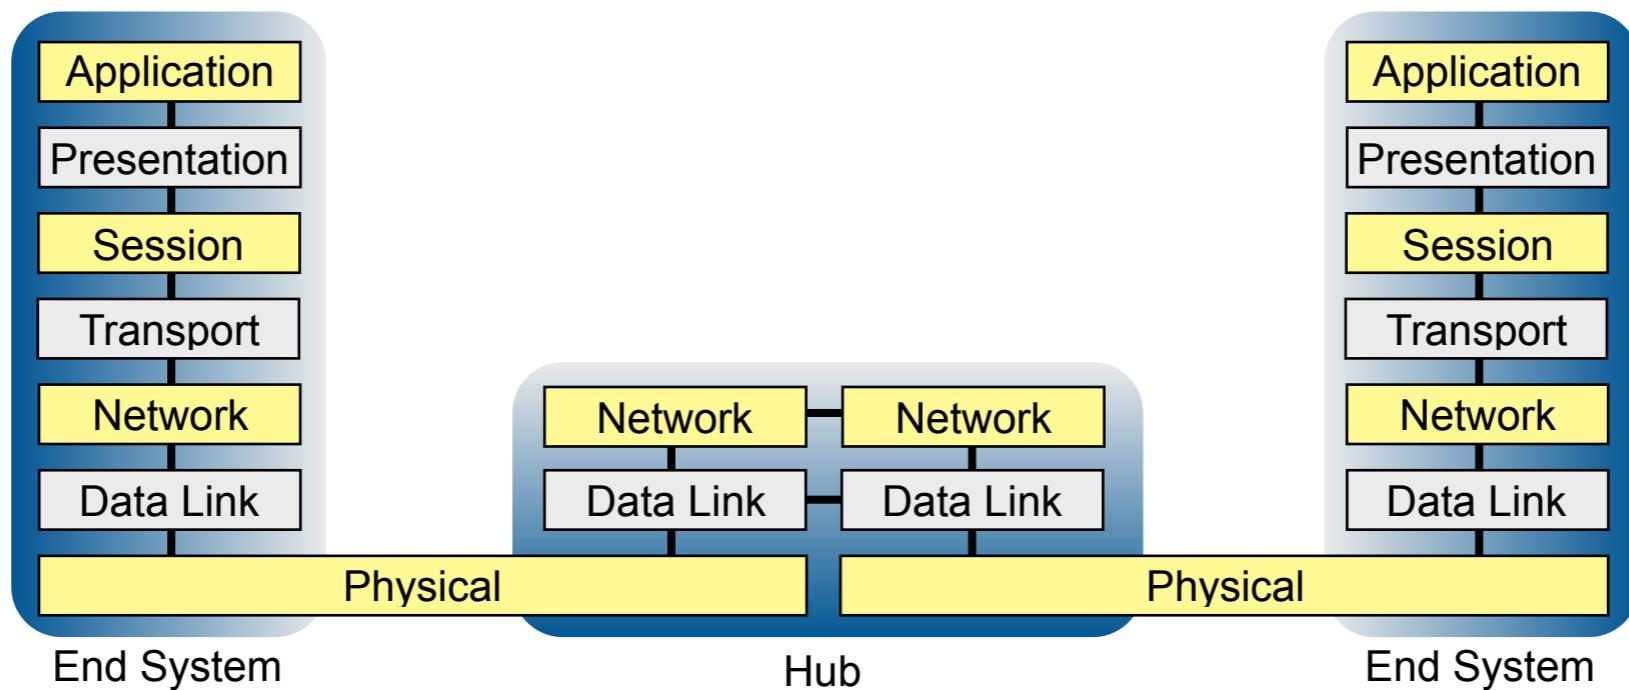


Internetworking

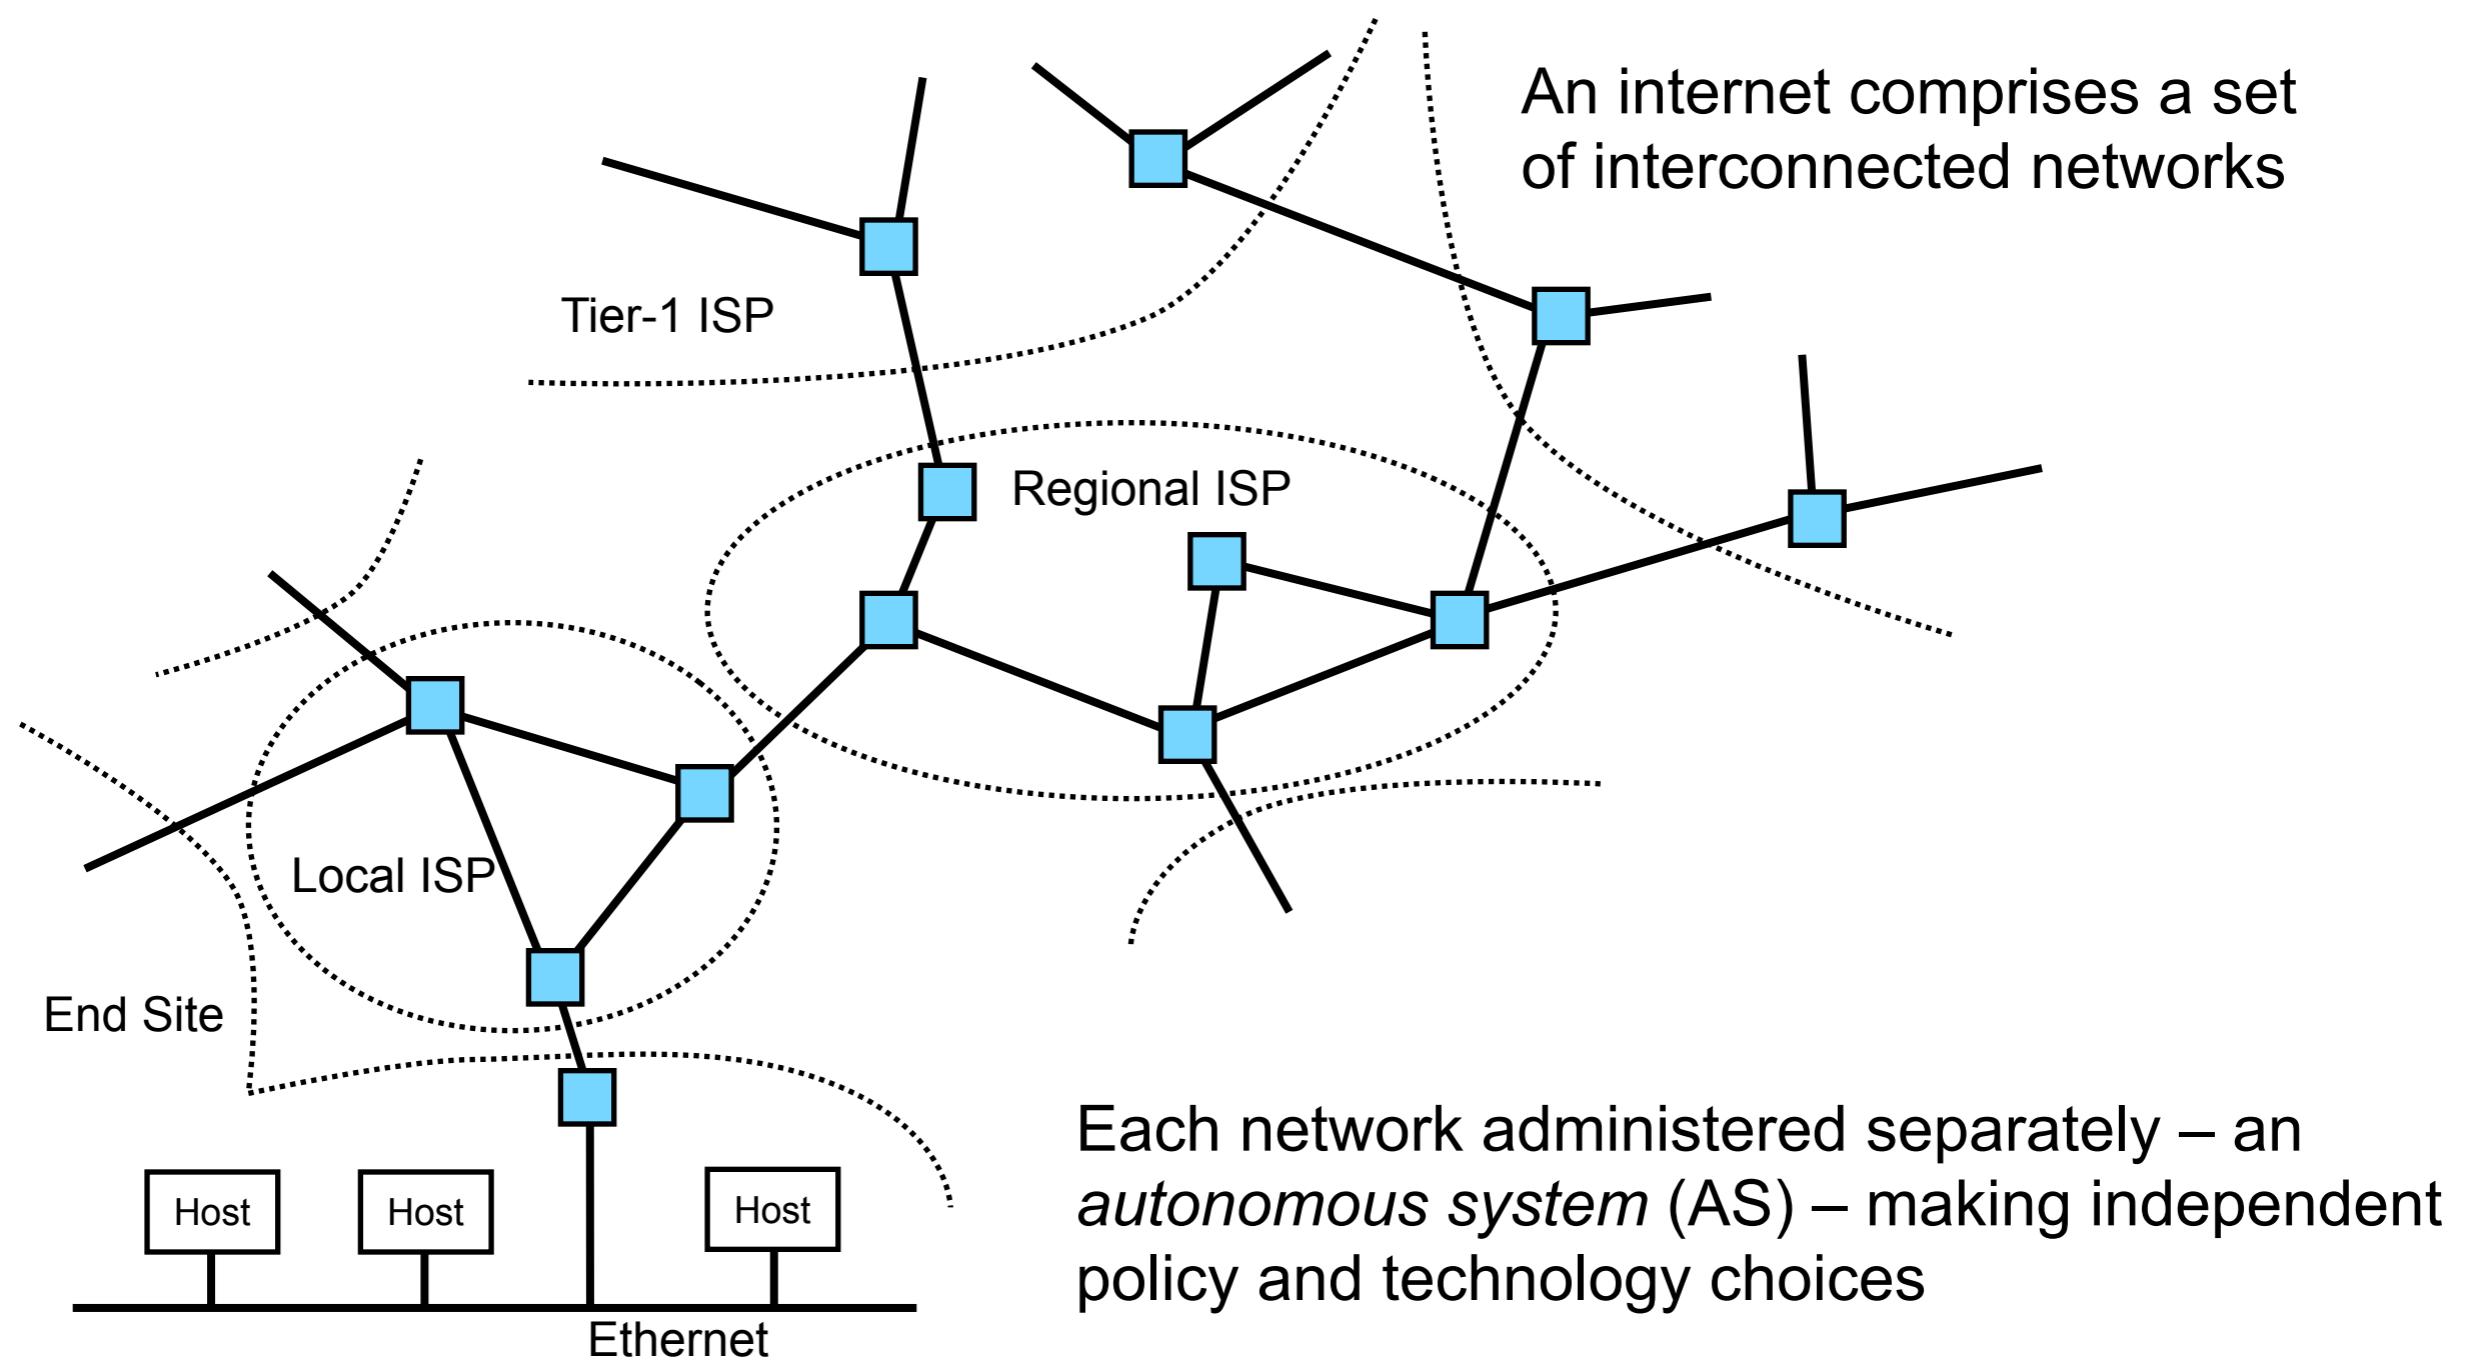
Networked Systems (H)
Lecture 7

Role of the Network Layer



- Network layer is first end-to-end layer in the OSI reference model
- Responsible for end-to-end delivery of data:
 - Across multiple link-layer hops and technologies
 - Across multiple *autonomous systems*
 - Building an *Internet*: a set of interconnected networks

Interconnecting Networks

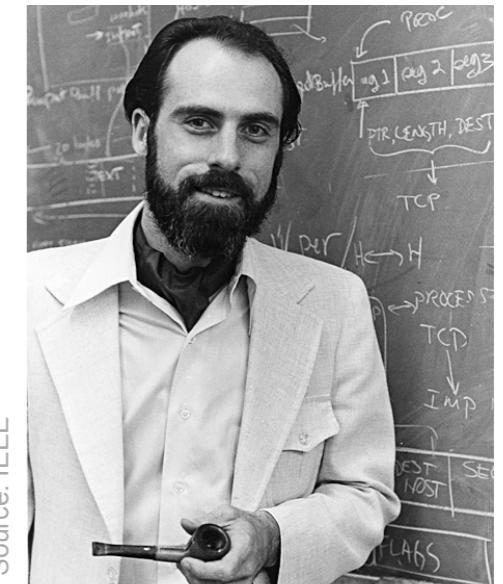


Components of an Internet

- A common end-to-end network protocol
 - Provide a single seamless service to transport layer
 - Delivery of data packets/provisioning of circuits
 - Addressing of end systems
- A set of gateway devices (a.k.a. *routers*)
 - Implement the common network protocol
 - Hide differences in link layer technologies
 - Framing, addressing, flow control, error detection and correction
 - Desire to perform the least amount of translation necessary

The Internet

- The globally interconnected networks running the *Internet Protocol* (IP)
 - Initial design by Vint Cerf and Robert Kahn, 1974
- IP provides an abstraction layer
 - Transport protocols and applications above
 - Assorted data link technologies and physical links below
 - A simple, best effort, connectionless, packet delivery service
 - Addressing, routing, fragmentation and reassembly



Source: IEEE

Vint Cerf

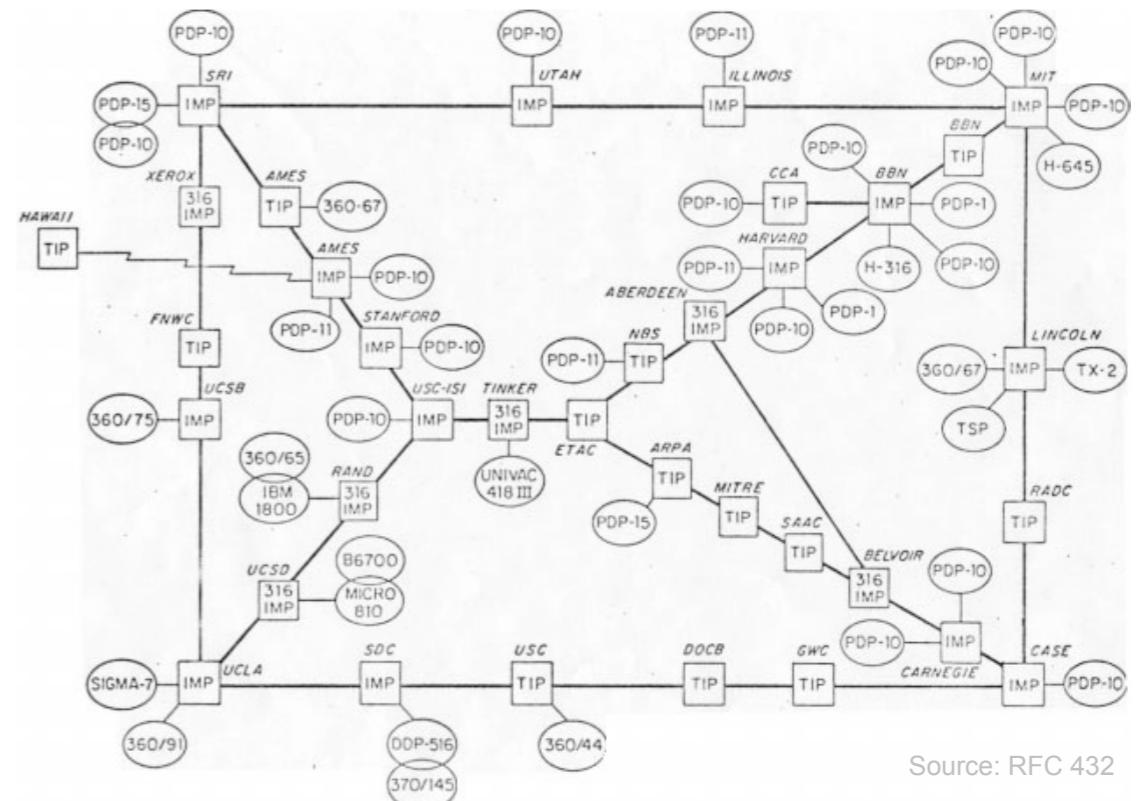


Source: IEEE

Robert Kahn

History and Development

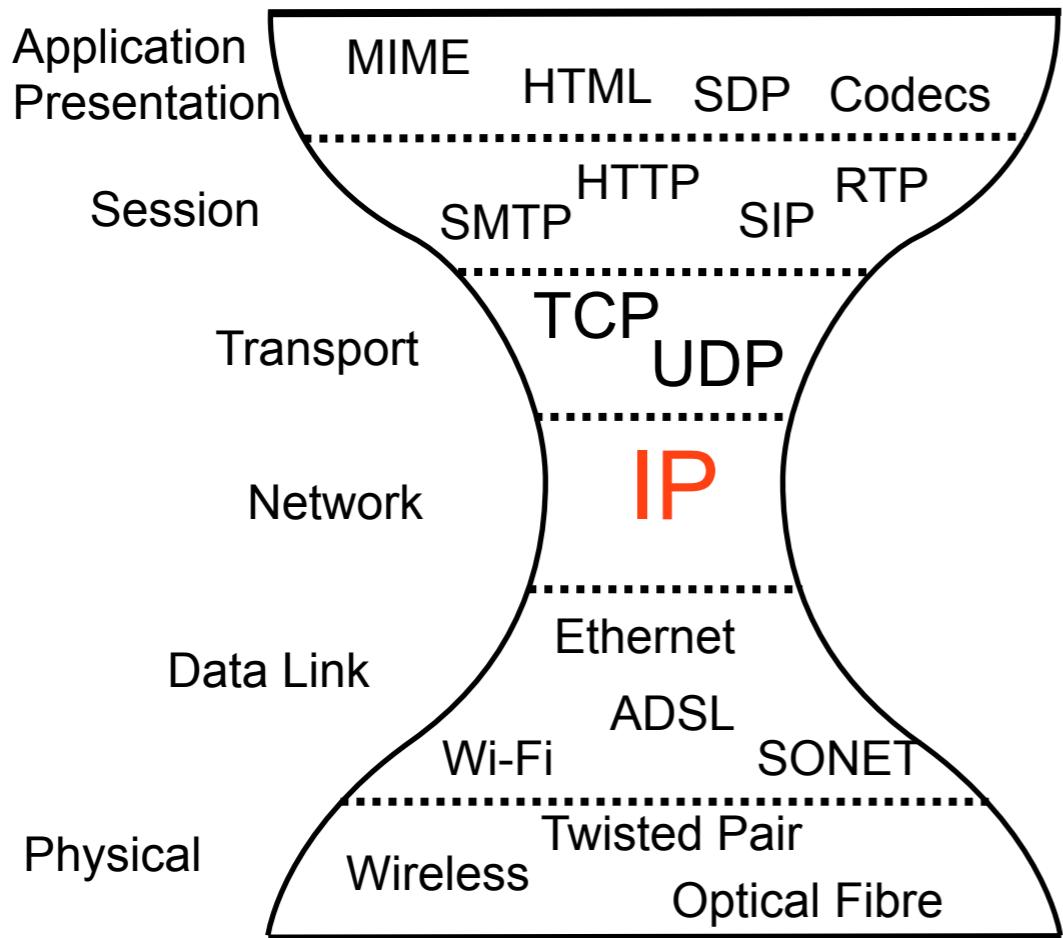
- 1965: Packet switching
 - Paul Baran (RAND), Donald Davies (NPL)
- 1969: ARPA funding
 - First link: UCLA – SRI
- 1973: First non-US sites
 - UCL, SICS
- 1983: Switch to IPv4
- 1990: World Wide Web
 - Tim Berners-Lee



Source: RFC 432

ARPA network map, December 1972

Basic Concepts

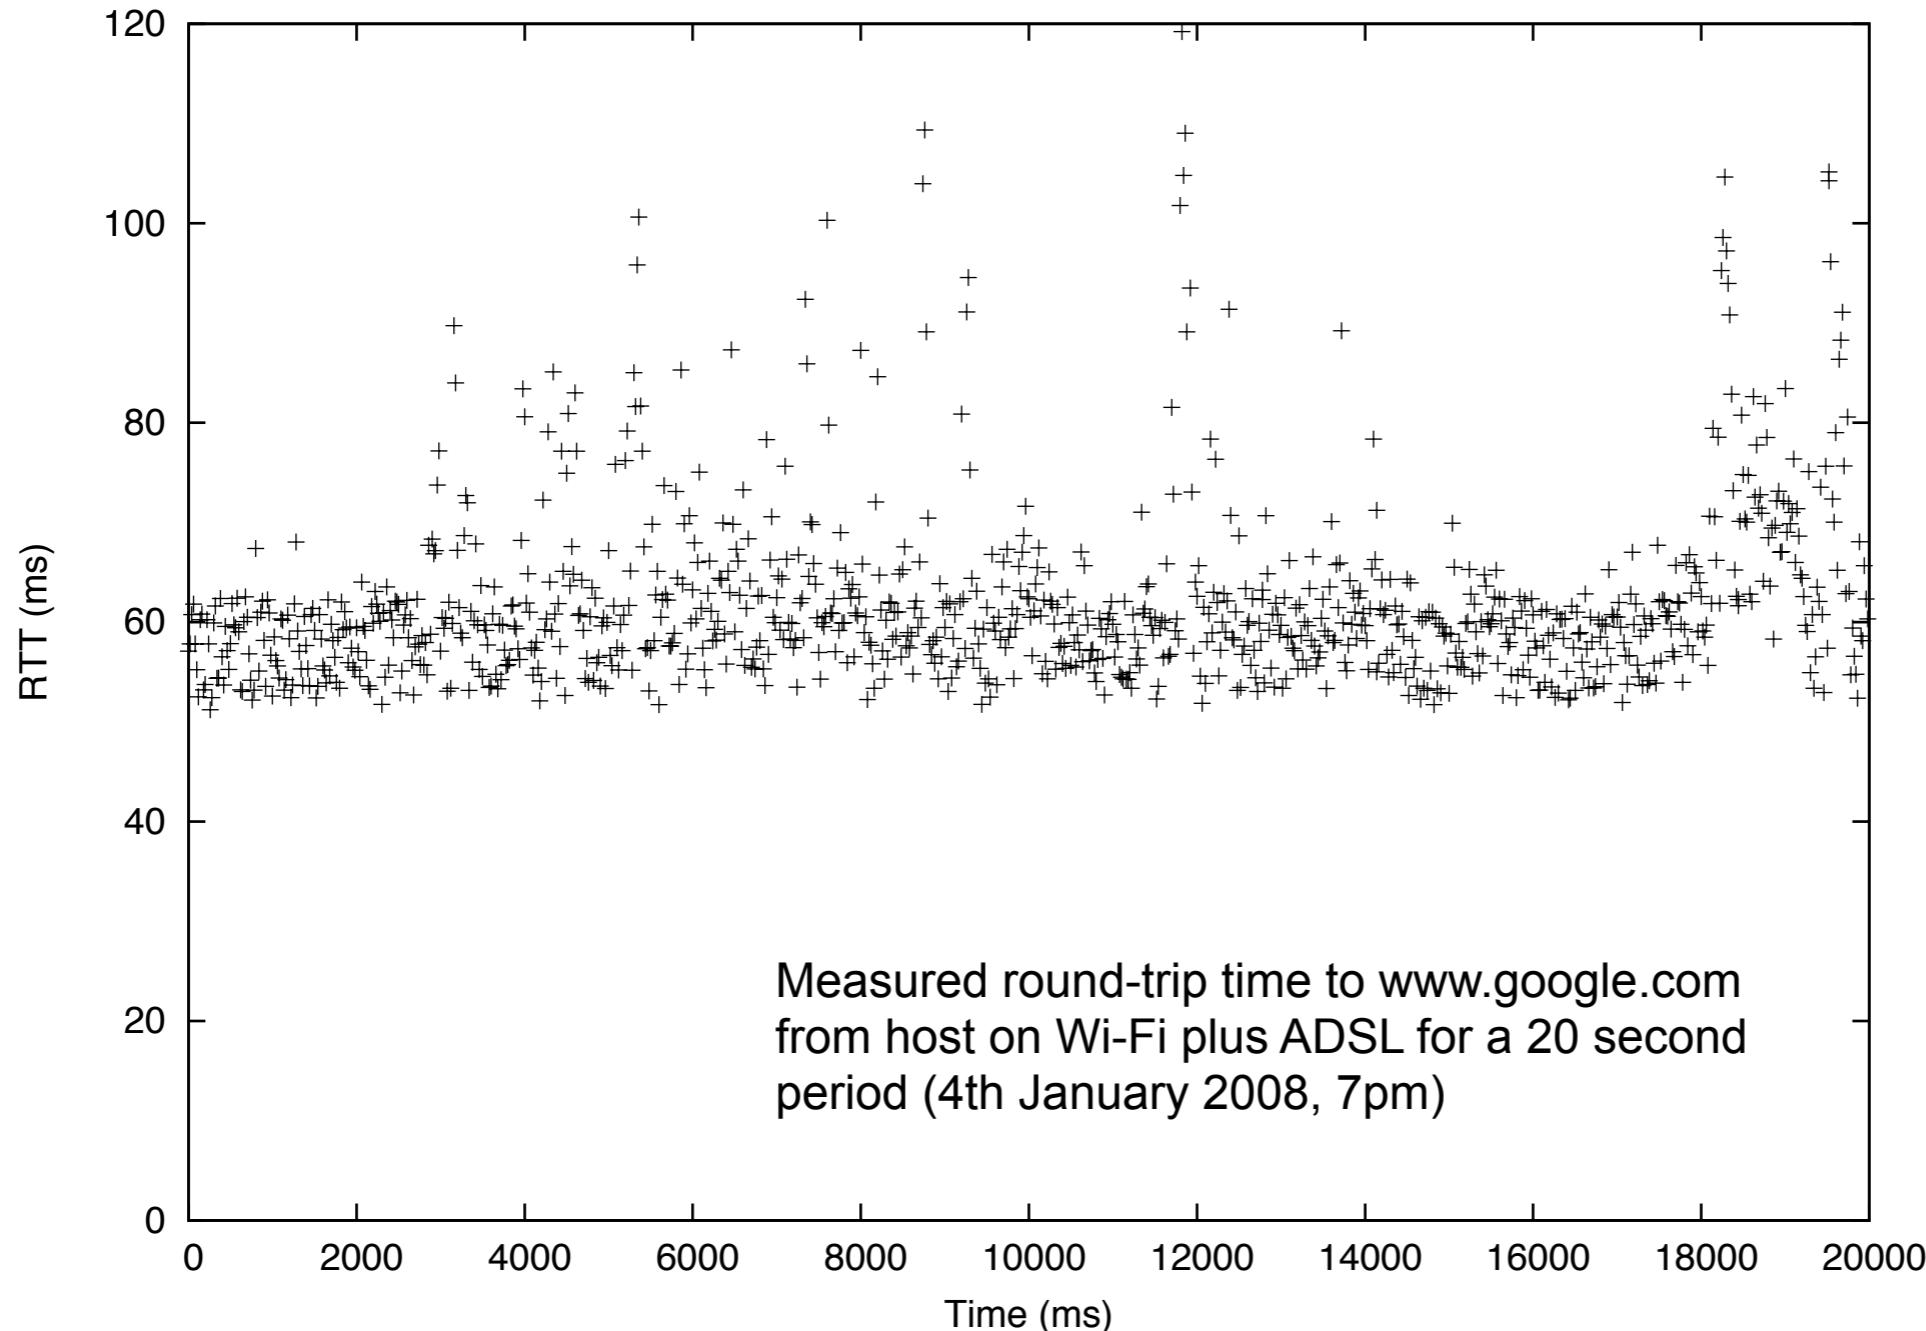


- Global inter-networking protocol
- Hour glass protocol stack
 - Single standard network layer protocol (IP)
 - Packet switched network, best effort service
 - Uniform network and host addressing
 - Uniform end-to-end connectivity (subject to firewall policy)
 - Many transport & application layer protocols
 - Range of link-layer technologies supported

IP Service Model

- Best effort, connectionless, packet delivery
 - Just send – no need to setup a connection first
 - Network makes its *best effort* to deliver packets, but provides no guarantees
 - Time taken to transit the network may vary
 - Packets may be lost, delayed, reordered, duplicated or corrupted
 - The network discards packets it can't deliver
 - Easy to run over any type of link layer
 - Fundamental service: can easily simulate a circuit over packets, but simulating packets over a circuit difficult

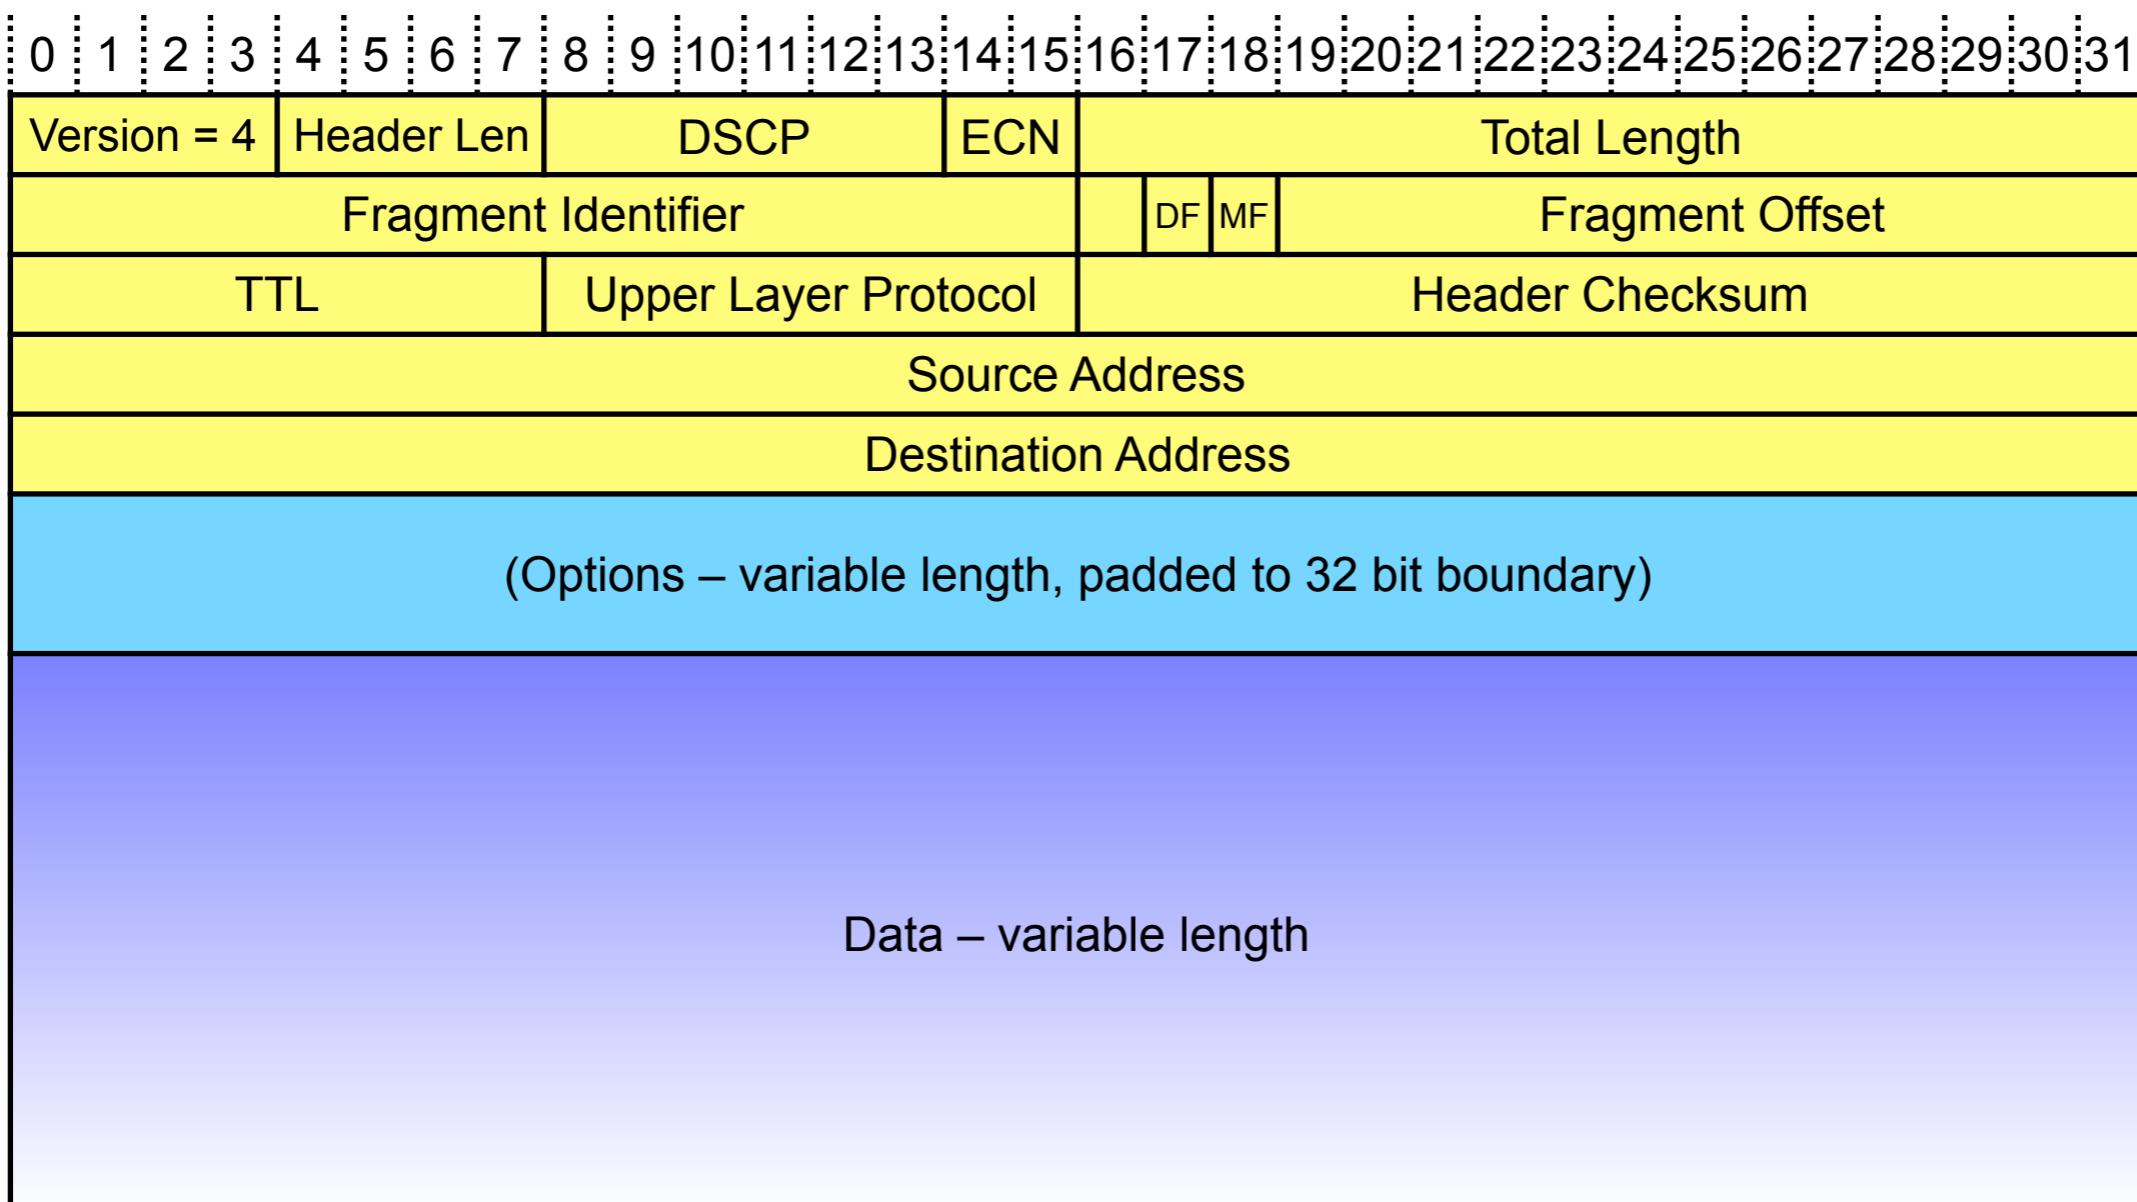
Best Effort Packet Delivery



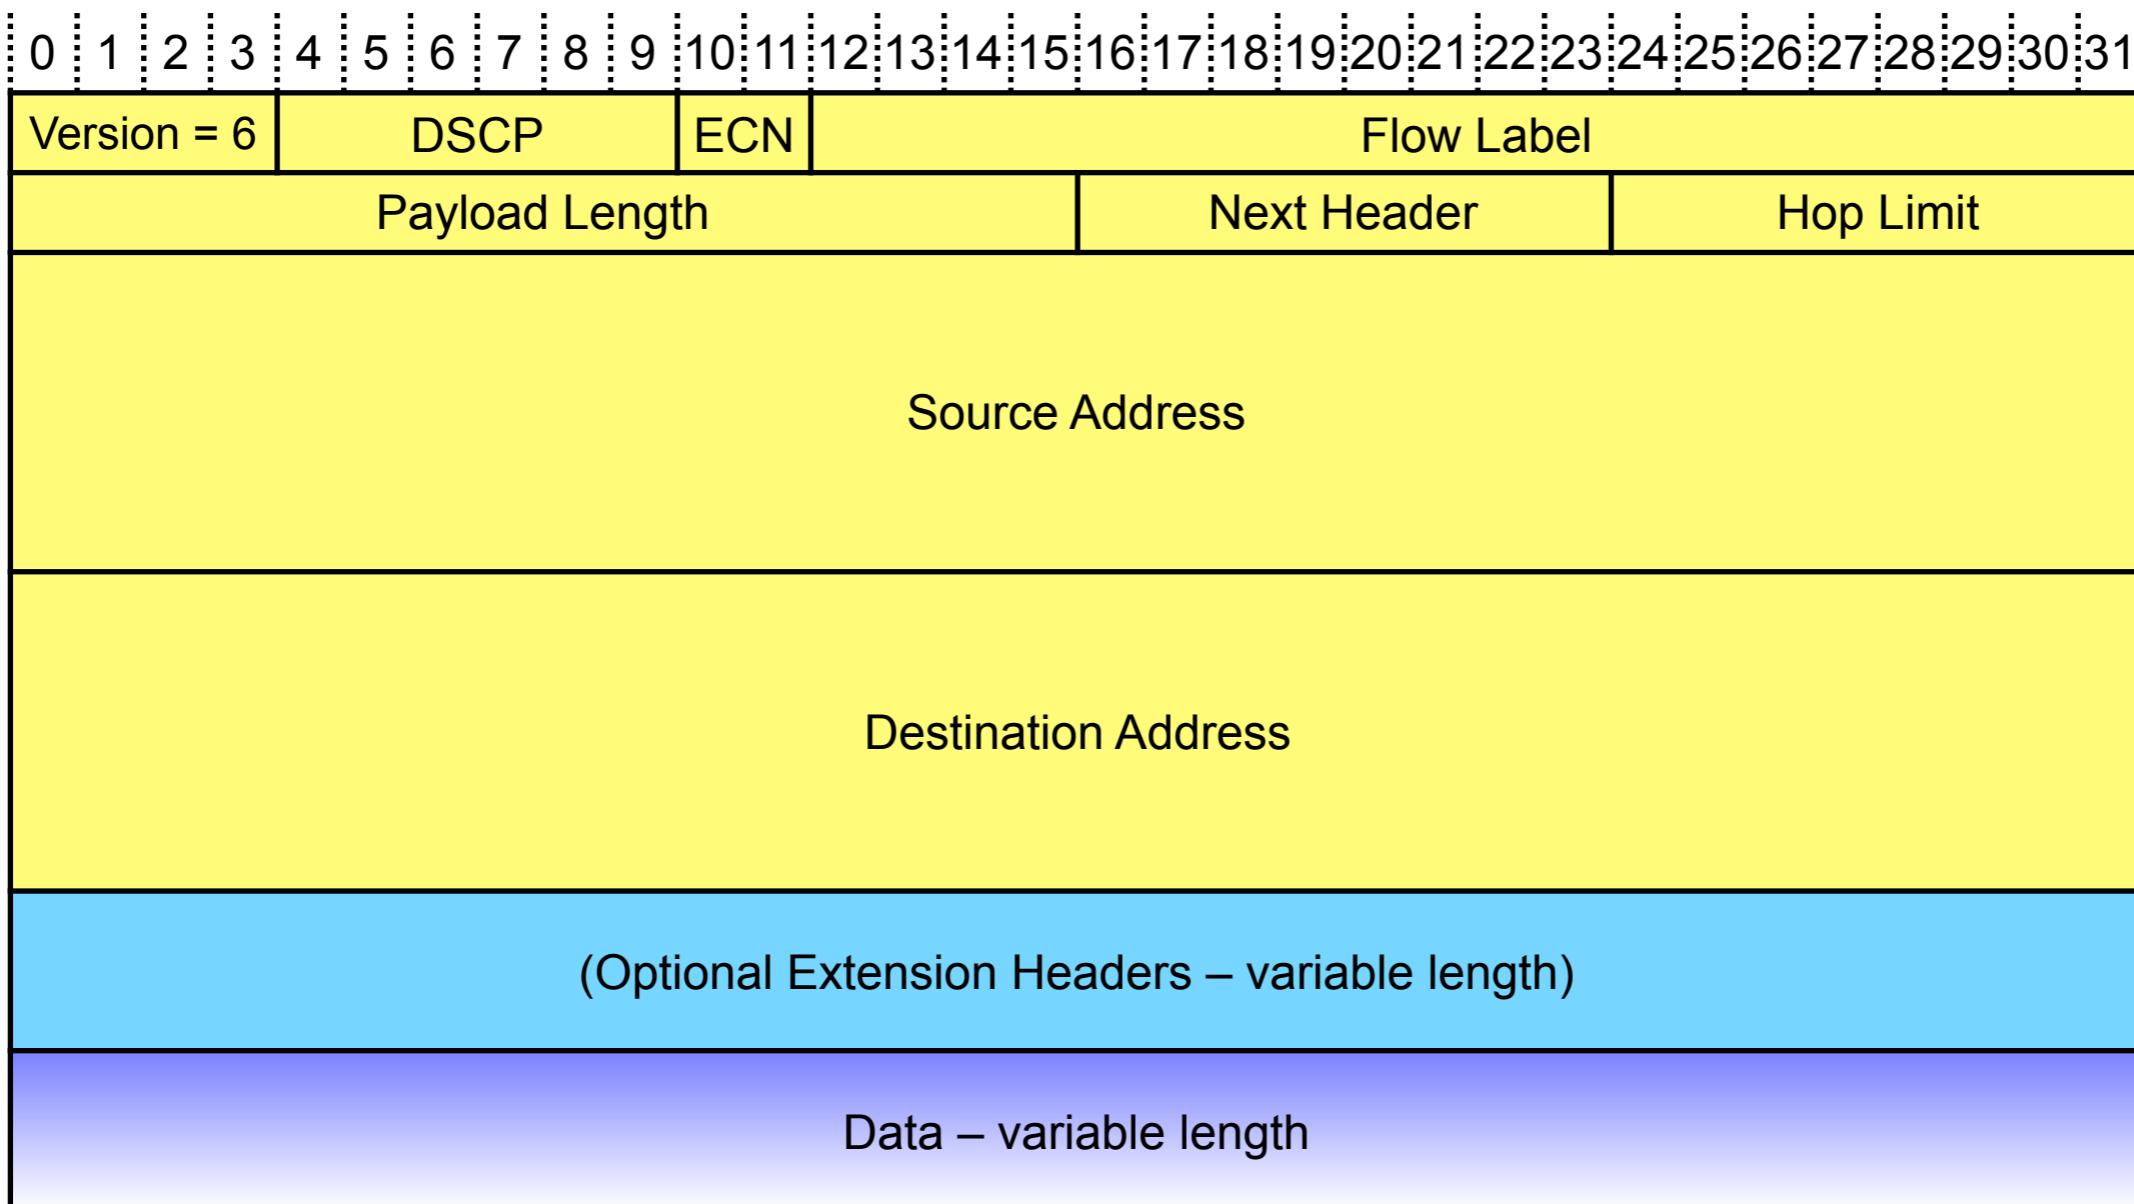
Internet Protocol

- Two versions of IP in use:
 - IPv4 – the current production Internet
 - IPv6 – the next generation Internet
- IPv5 was assigned to the Internet Stream Protocol
 - An experimental multimedia streaming protocol developed between 1979 and 1995 [<http://www.ietf.org/rfc/rfc1819.txt>], but no longer used

IPv4 Packet Format

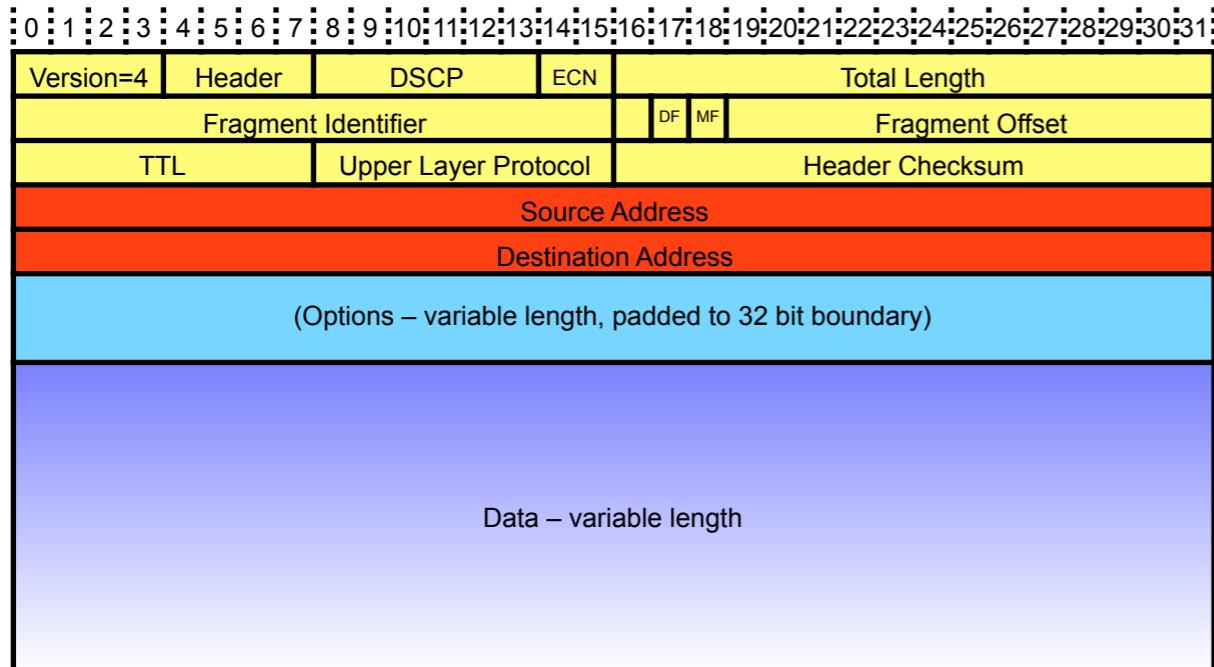


IPv6 Packet Format

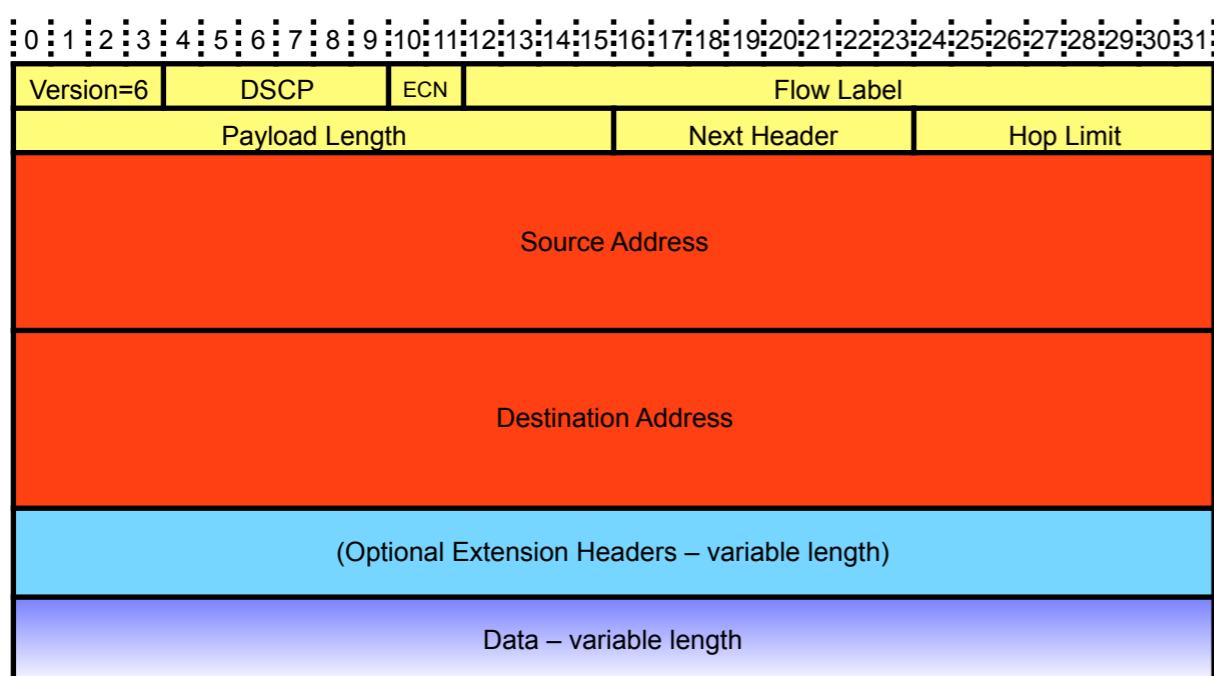


Compared to IPv4: simpler header format, larger addresses, removes support for fragmentation, adds flow label

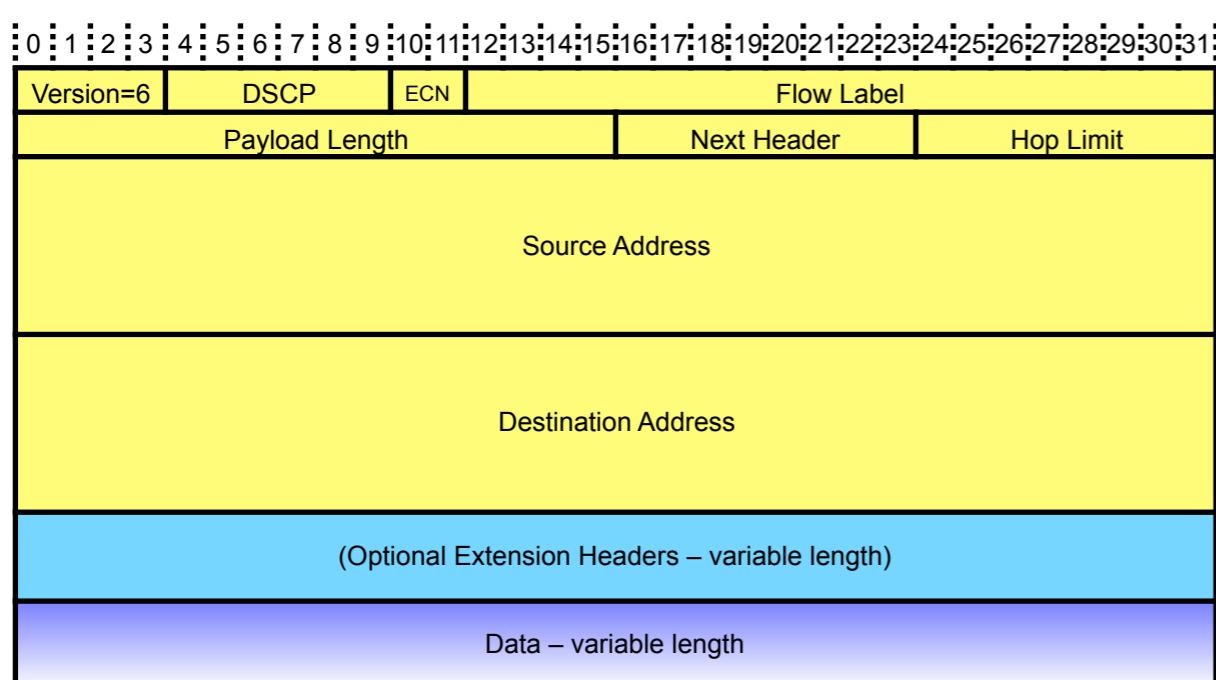
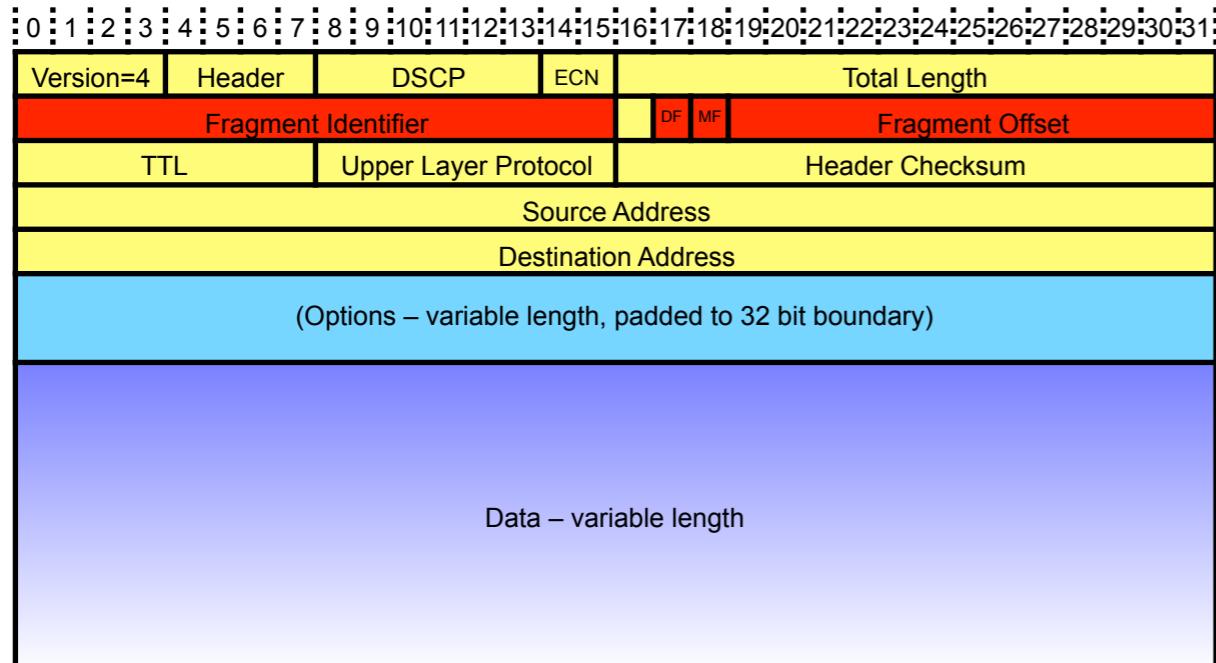
Addressing



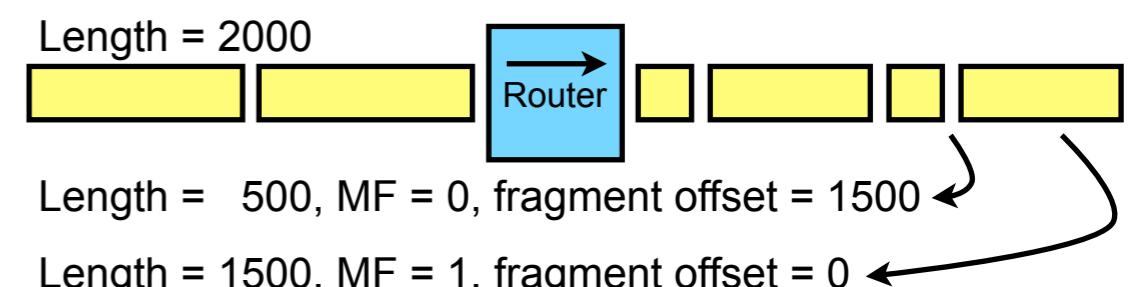
- Every network interface on every host is intended to have a unique address
 - Hosts may change address over time to give illusion of privacy
 - Addressable ≠ reachable: firewalls exist in both IPv4 and IPv6
- IPv4 addresses are 32 bits
 - Example: 130.209.247.112
 - Significant problems due to lack of IPv4 addresses → lecture 9
- IPv6 addresses are 128 bits
 - Example: 2001:4860:4860::8844



Fragmentation

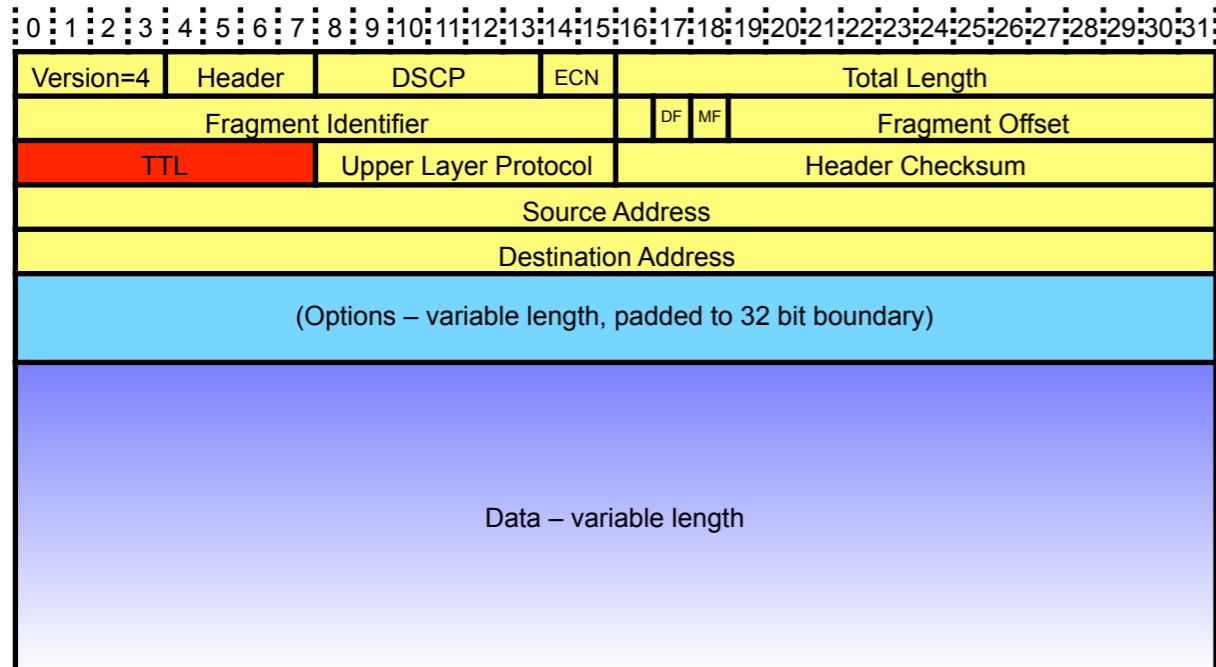


- Link layer has a maximum packet size (MTU)
- IPv4 will routers fragment packets that are larger than the MTU
 - MF bit is set if more fragments follow: reconstruct using fragment offset and fragment identifier

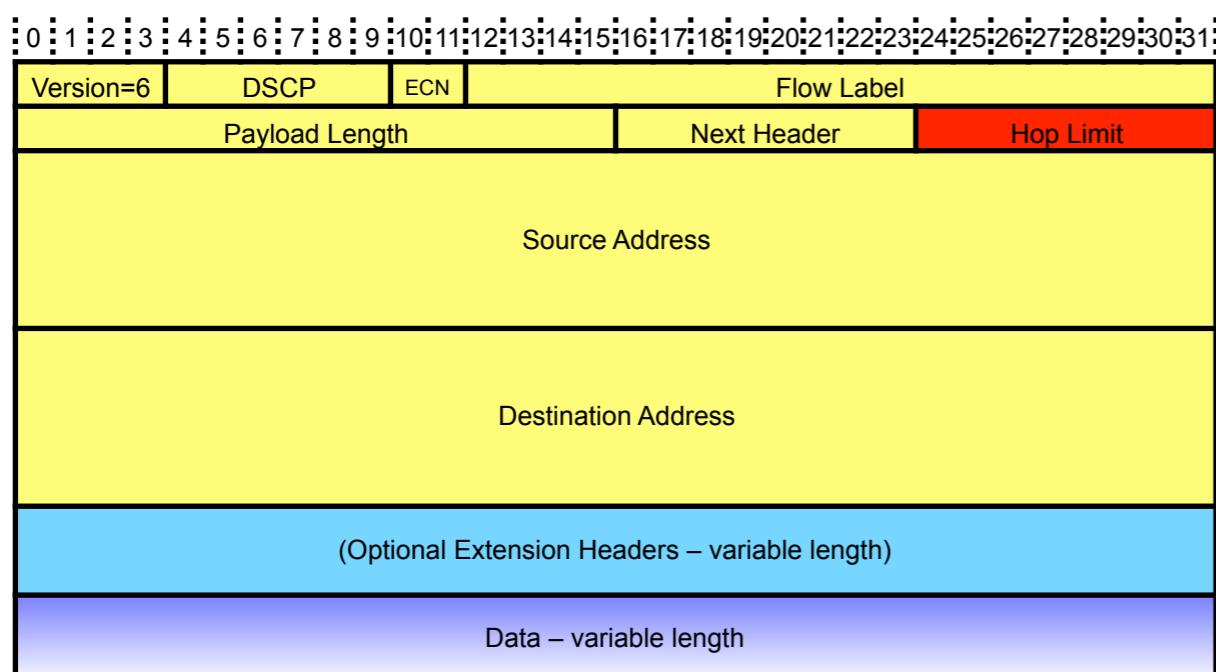


- DF bit is set to indicate routers shouldn't fragment, and must discard large packets
- IPv6 doesn't support fragmentation
 - Hard to implement for very high rate links
 - End-to-end principle

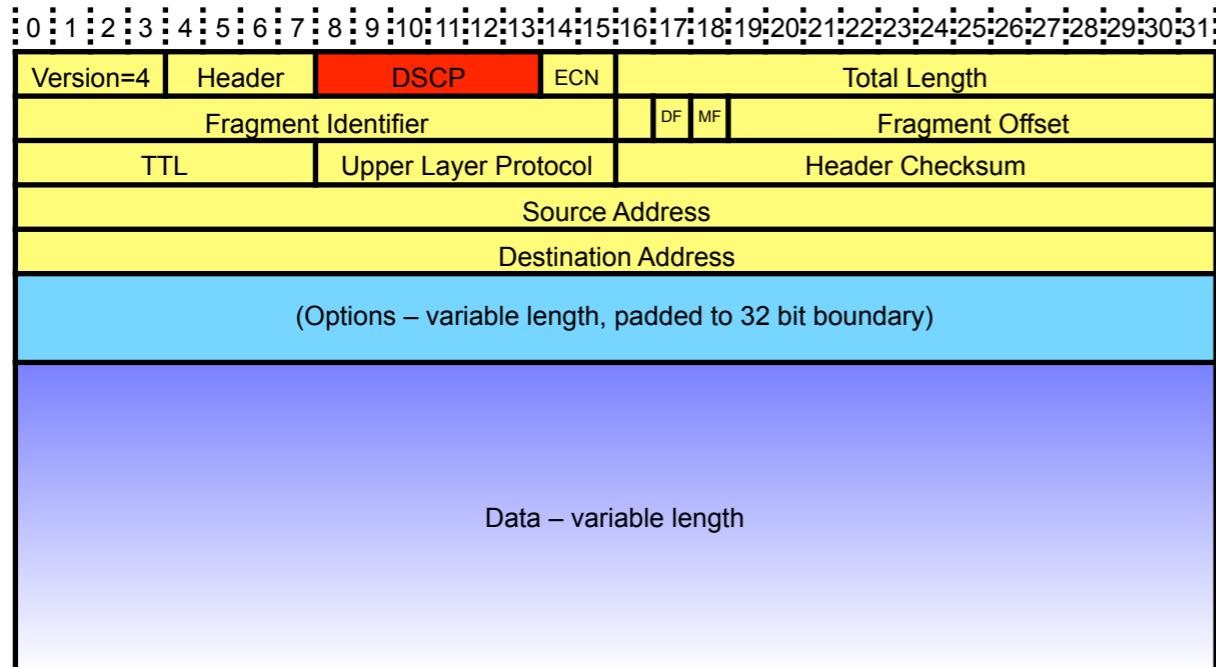
Loop Protection



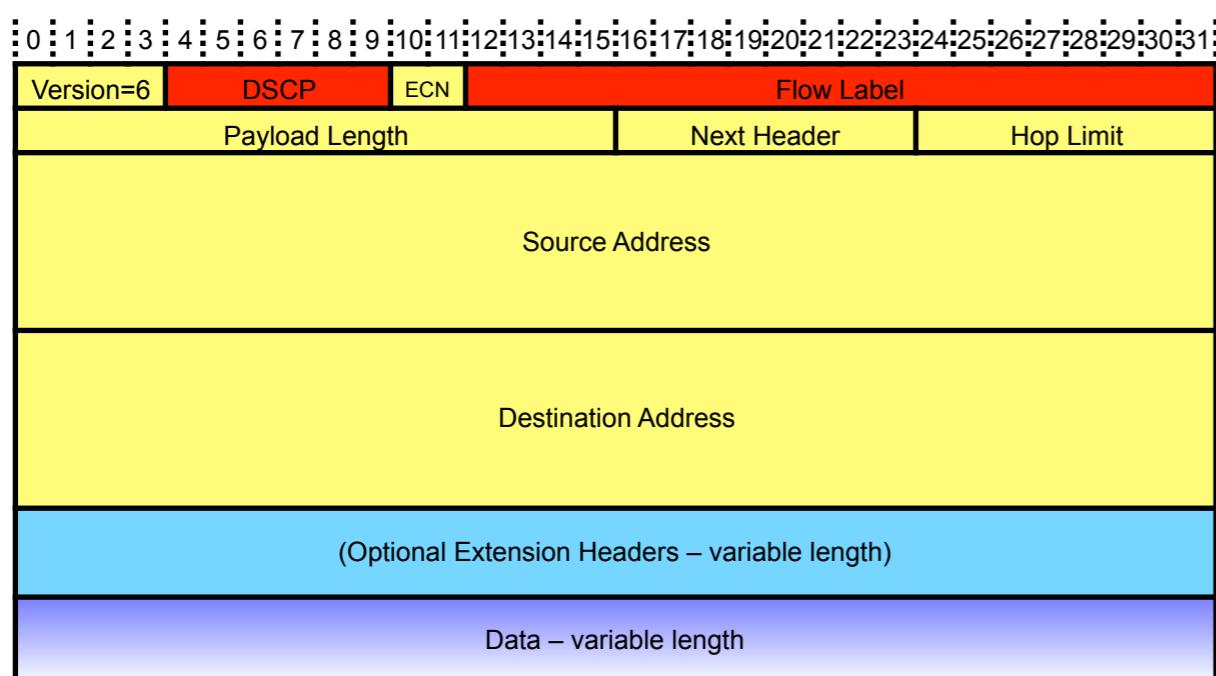
- Packets include a forwarding limit:
 - Set to a non-zero value when the packet is sent (typically 64 or 128)
 - Each router that forwards the packet reduces this value by 1
 - If zero is reached, packet is discarded
- Stops packets circling forever if a network problem causes a loop
- Assumption: network diameter is smaller than initial value of forwarding limit



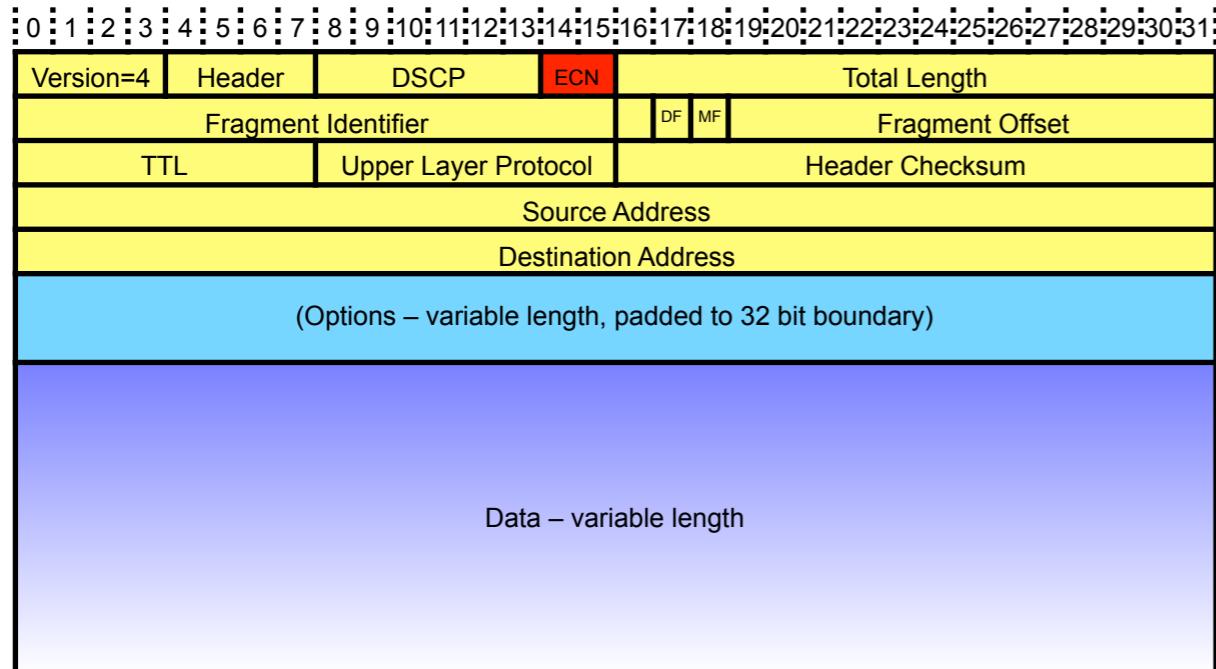
Differentiated Services



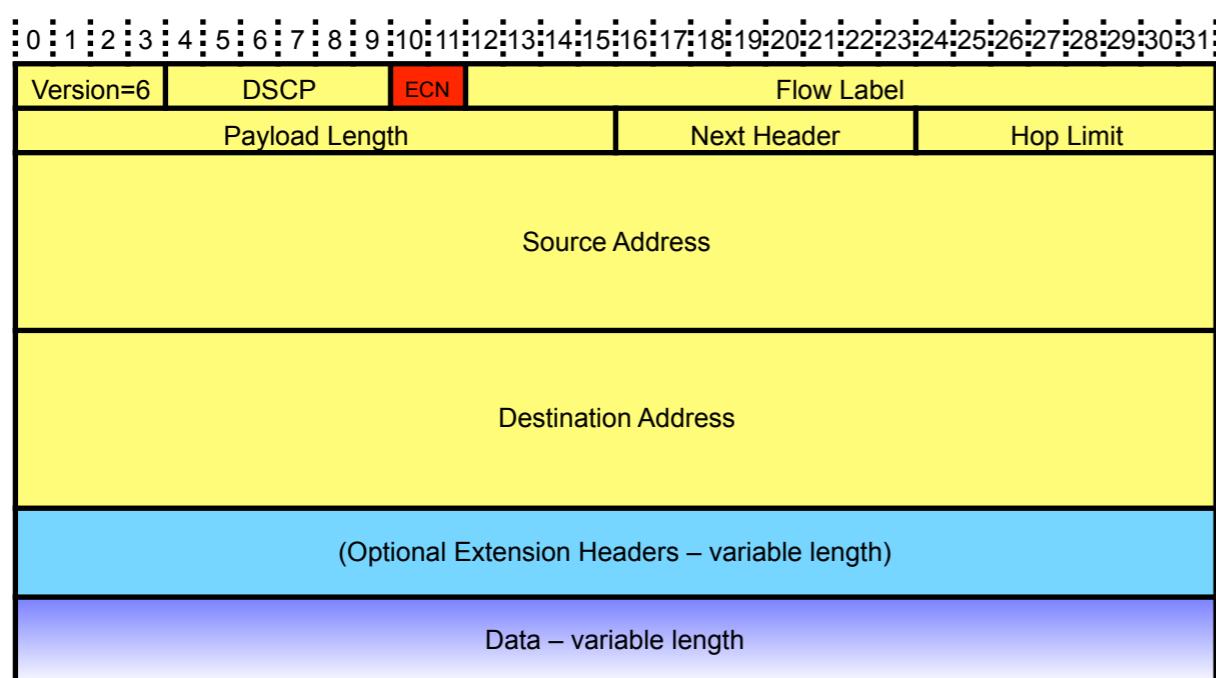
- End systems can request special service from the network
 - Telephony or gaming might prefer low latency over high bandwidth
 - Emergency traffic could be prioritised
 - Background software updates might ask for low priority
- Signalled by differentiated service code point (DSCP) field in header
- Provides a hint to the network, not a guarantee
 - Often stripped out at network boundaries
 - Difficult economic and network neutrality issues – who is allowed to set the DSCP and what are they charged for doing so?
- IPv6 provides a flow label to group related traffic flows together



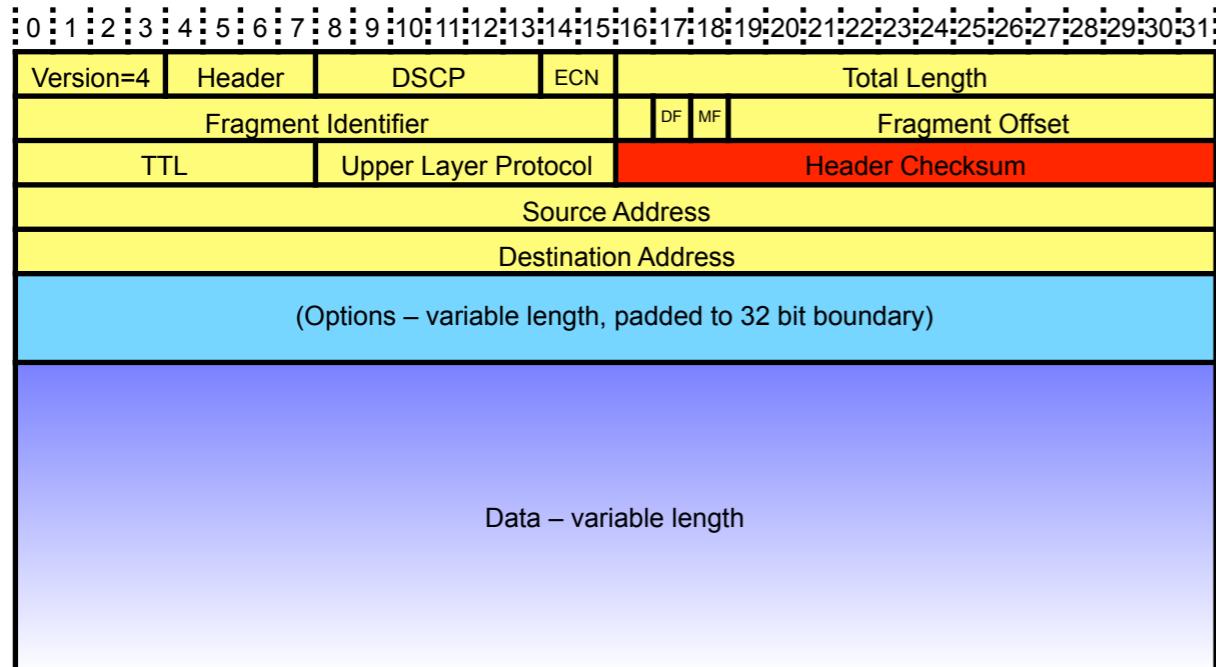
Explicit Congestion Notification



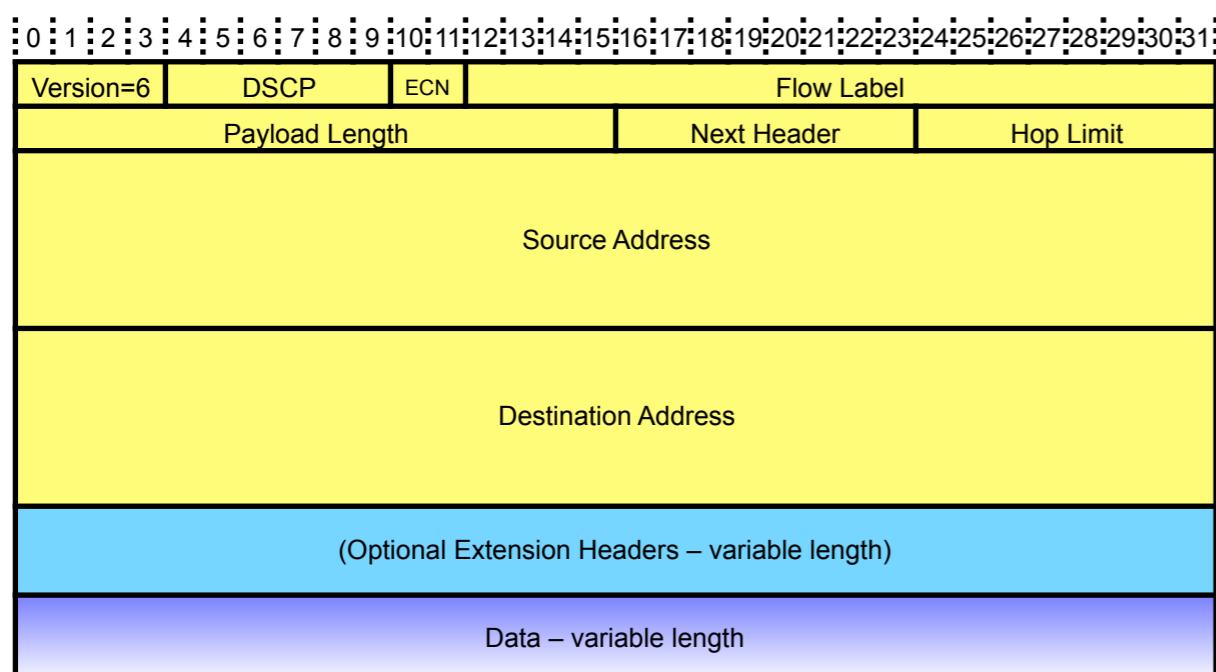
- Routers typically respond to network congestion by dropping packets
 - A “best effort” packet delivery service
 - Transport protocols detect the loss, and can request a retransmission if necessary
- Explicit congestion notification gives routers a way to signal congestion is approaching
 - If ECN=00 explicit congestion notification is disabled
 - If a sending host sets ECN=10 or ECN=01, routers monitor link usage, and can change the field to ECN=11 indicating congestion is imminent
 - A host receiving ECN=11 needs to reduce its sending rate – or the congested router will start dropping packets



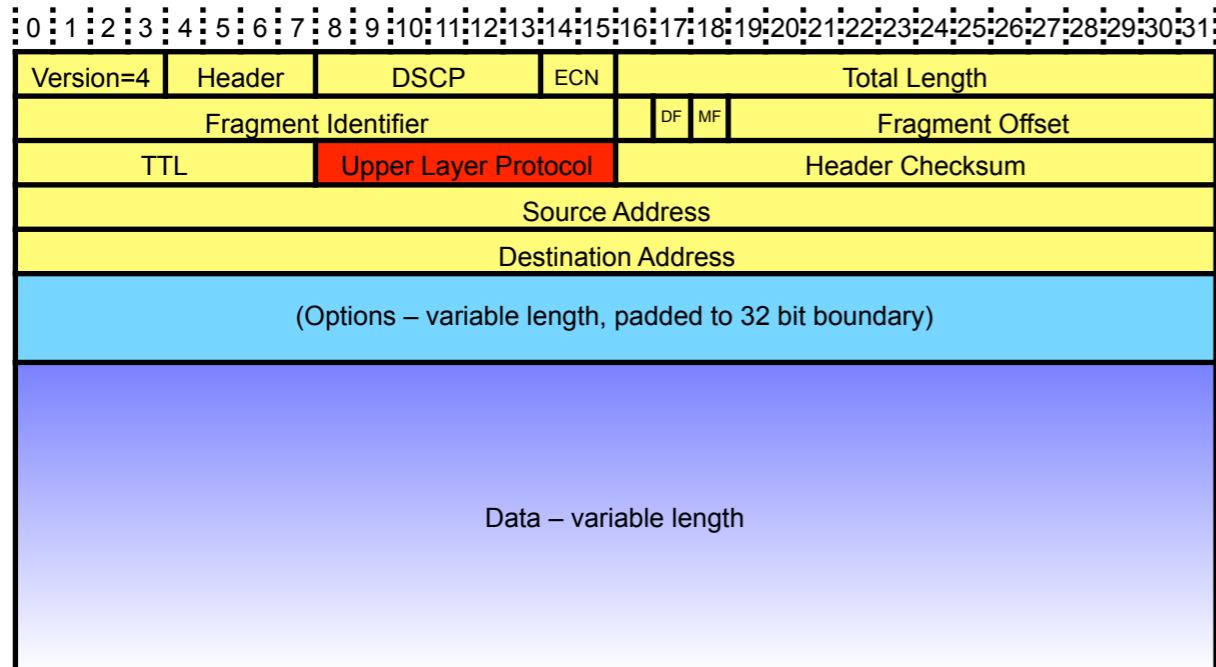
Header Checksum



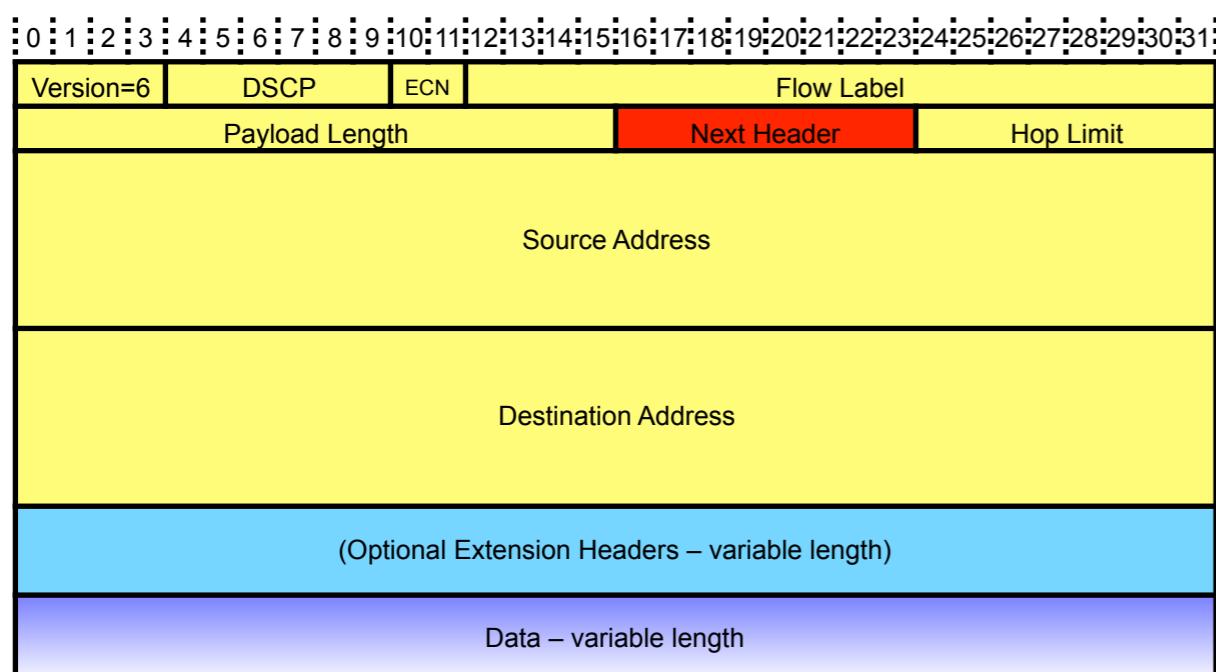
- IPv4 header contain a checksum to detect transmission errors
 - Conceptually similar to link-layer checksum, although uses a different algorithm
 - Protects the IP header only, not the payload data protected (must be protected by upper layer protocol, if needed)
- IPv6 does not contain checksum – assumes the data is protected by a link layer checksum



Transport Layer Protocol Identifier



- Network layer packet carry transport layer data as their payload
- Necessary to identify what transport protocol is used, to pass the data to the correct upper-layer protocol
 - TCP = 6
 - UDP = 17
 - DCCP = 33
 - ICMP = 1
- Legal values managed by the IANA
<http://www.iana.org/assignments/protocol-numbers/>



IPv4 or IPv6?

- IPv4 has reached end-of-life: insufficient addresses
- IPv6 intended as long term replacement for IPv4
 - Primary goal: increase the size of the address space, to allow more hosts on the network
 - Also simplifies the protocol, makes high-speed implementations easier
- Not yet clear if IPv6 will be widely deployed
 - But, straight-forward to build applications that work with both IPv4 and IPv6
 - DNS query using `getaddrinfo()` will return IPv6 address if it exists, else IPv4 address; all other socket calls use the returned value
 - Write new code to support both IPv6 and IPv4

Summary

- Role of the network layer
- From *an* internet to *the* Internet
- Internet service model
- IPv4 and IPv6