

Internet Protocol (IP)

Goals:


- Interconnect diverse LANs into one logical entity.
- Implement *best effort* (unreliable, connectionless) service model.

Specifies

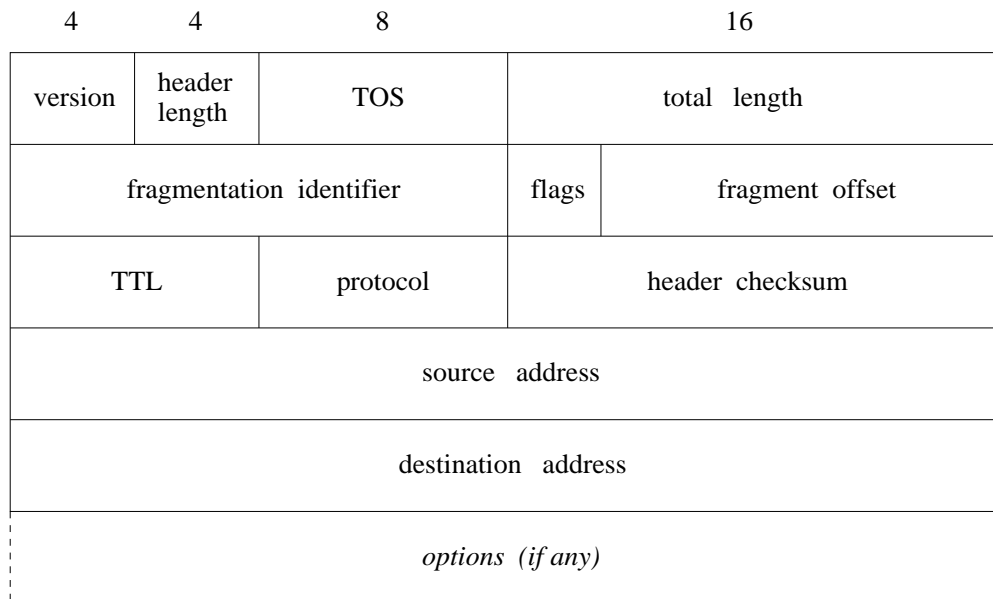
- Common language for carrying out non-LAN-specific conversations (protocol standards).
- Functionality and design philosophy.

Best effort vs. guaranteed service:

- Much easier to implement best effort service; no resource reservation.
- Simplifies router design but increases complexity of end stations \rightarrow trade-off
- Necessitates higher-up functional layer (transport layer) to achieve reliable transmission over unreliable medium.
- Duplication of work.
- Routers/switches already becoming more complex due to QoS; why not dispense with transport layer . . .

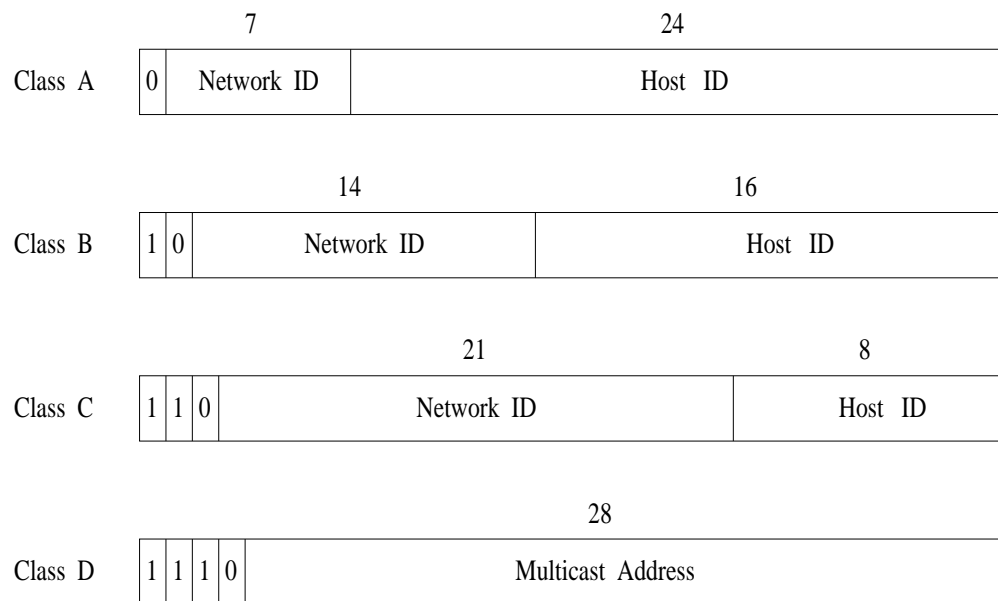


IP packet (datagram) format:



- Header length: in 4 byte (word) units.
- TOS (type-of-service): Most routers do not support.
- 4 bytes used for fragmentation.
- TTL (time-to-live): Prevent cycling (default 64).
- Protocol: demultiplexing key (TCP 6, UDP 17).

IP address format:



Dotted decimal notation: 10000000 00001011 00000011
00011111 \leftrightarrow 128.11.3.31

Symbolic name to IP address translation—domain name server (DNS).

Notice hierarchical organization (“2-level”).

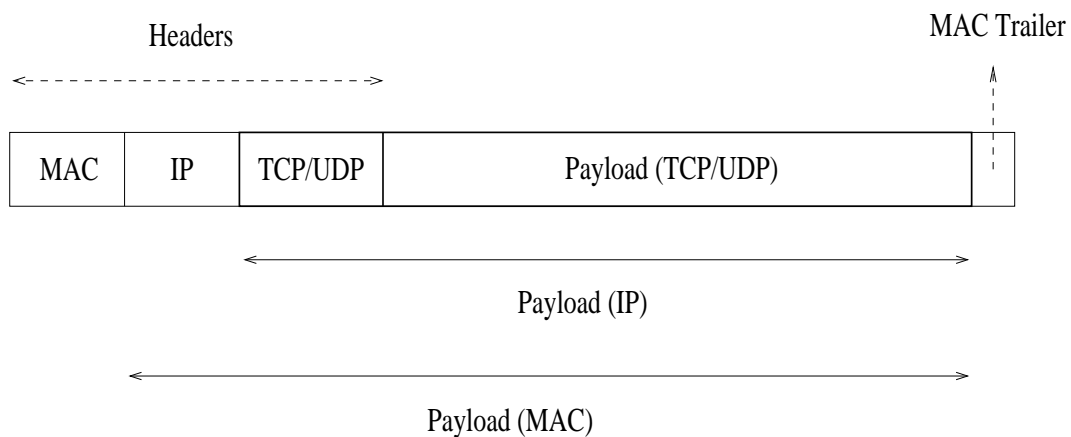
Each interface (NIU) has an IP address; single host can have multiple IP addresses.

Running out of unused addresses (IPv6).

Transport Protocols: TCP/UDP Structure

- end-to-end mechanism
- runs on top of link-based mechanism
- treat network layer as black box

Three-level encapsulation:



Network layer assumptions:


- unreliable
- out-of-order delivery (in general)
- absence of QoS guarantees (delay, throughput etc.)
- insecure (IPv4)

Additional (informal) performance properties:

- works “fine” under low load conditions
- can break down under high load conditions
- behavior range predictable (to certain extent)

Goal of UDP: Process identification (“multiplexing”).

→ port number as process demux key



- form of end host processing (O.S.)
- generally: end system support (e.g., scheduling)

UDP packet format:

2	2
Source Port	Destination Port
Length	Checksum
Payload	

Checksum calculation (pseudo header):

4		
Source Address		
Destination Address		
00 ... 0	Protocol	UDP Length

Goals of TCP:


- process identification
 - reliable communication (ARQ)
 - speedy communication (congestion/flow control)
 - segmentation
- connection-oriented (i.e., stateful)
- complex mixture of functionalities

Segmentation task: Provide “stream” interface to higher level protocols

—→ view: contiguous stream of bytes


- segment stream of bytes into blocks or *segments* of fixed size
- segment size determined by TCP MTU (Maximum Transmission Unit)
- use also for reliability mechanism

TCP packet format:

Source Port		Destination Port						
Sequence Number								
Acknowledgement Number								
Header Length		U R G	A C K	P S H	R S T	S S N	F I N	Window Size
Checksum				Urgent Pointer				
Options (if any)								
DATA (if any)								


- Sequence Number: position of first byte of payload
- Acknowledgement: next byte of data expected (receiver)
- Header Length (4 bits): 4 B units
- URG: urgent pointer flag
- ACK: ACK packet flag
- PSH: override TCP buffering
- RST: reset connection
- SYN: establish connection
- FIN: close connection
- Window Size: receiver's advertised window size
- Checksum: prepend pseudo-header
- Urgent Pointer: byte offset in current payload where urgent data begins
- Options: MTU; take min of sender & receiver (default 556 B)

TCP connection establishment (3-way handshake):




- X, Y are chosen randomly
- piggybacking
- sequence number prediction
- lingering packet problem

Call Collision:



- only single TCB gets allocated
- unique full association

TCP connection termination:



- full duplex
- half duplex