

# IPng (IPv6)

## Motivation

Currently: mainly servers, workstations, SLIP/PPP  
Address space exhaustion + new uses:

- permanently attached home computers (CATV)
- game consoles
- mobile terminals (PDAs)
- embedded controllers (light switches)

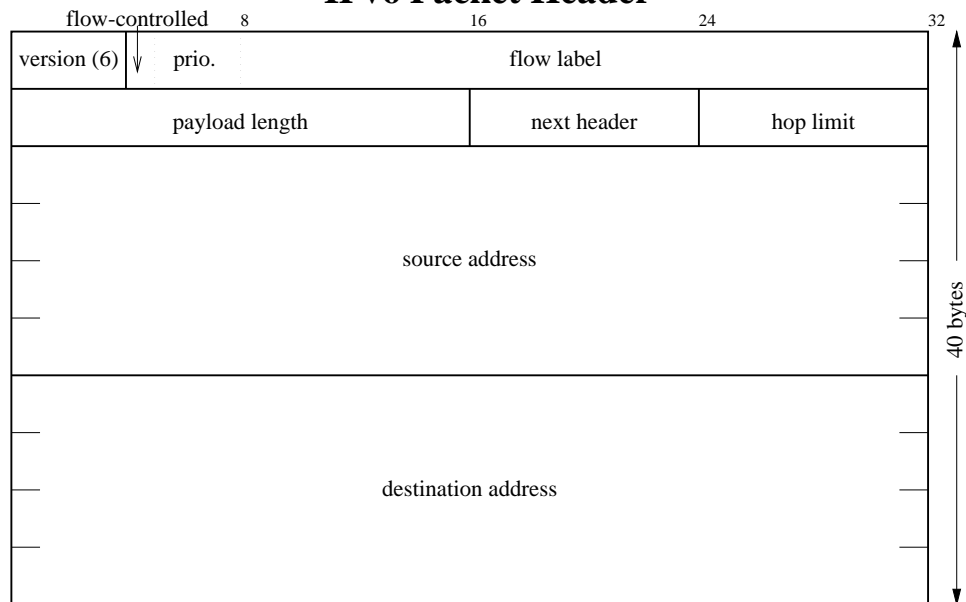
⇒ new IP ⇒ fix other problems

CLNP considered (variable-length addresses!), but not chosen

## Address Space

- 32 bits  $\Rightarrow 4 \cdot 10^9$  hosts, but can't number sequentially
- 128 bits instead of 32 bits  $\Rightarrow$   
340,282,366,920,938,463,463,374,607,431,768,211,456 ( $3 \cdot 10^{38}$ ) hosts  
or 665,570,793,348,866,943,898,599 hosts/m<sup>2</sup> of earth surface
- hierarchical assignment  $\Rightarrow$  inefficiency  $\Rightarrow$  only  $8 \cdot 10^{17}$  to  $2 \cdot 10^{33}$   
addresses if similar to telephones, etc.
- write addresses hex, with zero suppression:  
4123:::ACD9:4571:D5:F3:7


## IPv6 Packet Header



## IPv6 Packet Header

IPv4	IPv6
4-byte addresses	16-byte addresses
final destination	intermediate destination
precedence, TOS field	flow id, priority
header length, total length	payload length
≤ 44 bytes options	“unlimited”
options	options ∈ header extensions
router fragmentation	e-e fragmentation
header checksum	–
higher-layer protocol	next header type
TTL time-based	TTL hops only
20 bytes fixed header	40 bytes

## IPv6 Packet Structure



- next header can be IPv4 “tunneling”
- header extension: hop-by-hop options (HHO), routing, fragment, destination options (DO)
- DO, HHO: type-length-value TLV options
- HHO:
  - looked at by each node, immediately after header
  - jumbo payload option (32 bits)
- routing header:
  - fixed header may not contain final address if routing header!

- mixed loose/strict source route (bitmask)
- swap destination address and next address from routing header
- fragment header: like IPv4 (32 bit identification, offset, more fragments flag)
- explicit MTU message rather than try-until-fit


## IPv6 Flows

- explicit and implicit state
- router may cache based on flow identifier
- implicit state discarded after < 6 seconds
- priority ranges:
  - congestion controlled (TCP) drop priority
  - not controlled (UDP; fixed-bandwidth services)
    - 0 uncharacterized traffic
    - 1 filler traffic (netnews)
    - 2 unattended data transfer (email)
    - 4 attended bulk transfer (ftp, nfs)
    - 6 interactive traffic (telnet, X)
    - 7 internet control traffic (routing, network management)

## Address Assignment

Administrative and topological:

0000 0000	IPv4 addresses	1/256
0000 0001	reserved	1/256
0000 001	NSAP addresses	1/128
0000 010	IPX addresses	1/128
0000 011	reserved	1/128
0000 1	reserved	3/128
0001	reserved	1/16
001	reserved	1/16
010	provider-based	1/8
011	reserved	1/8
100	geographic	1/8
...		
1111 1111	scoped multicast	1/256



## Security

- authentication: insure that packet comes from right source → prevent “address spoofing”
- encryption: insure privacy (DES)

## Transition

- cannot have flag day
- transition period may last decade or more
- dual-stack → IPv6 can send IPv4
- IPv4 address range
- DNS servers need to understand IPv6 addresses
- avoid translation, use IPv6-in-IPv4 tunneling
- use IPv6 on same Ethernet
- end-to-end tunnel or concatenated tunnels




## UDP



## UDP Services

- datagram (like IP)
- protocol identifier 17
- unreliable - discarded by network or host (without notification)
- demultiplexing by process (not host)
- each datagram received by one process (“socket”) only
- no connection setup
- to answer: reverse source-destination ports
- also: audio, video, multicast

## UDP Header

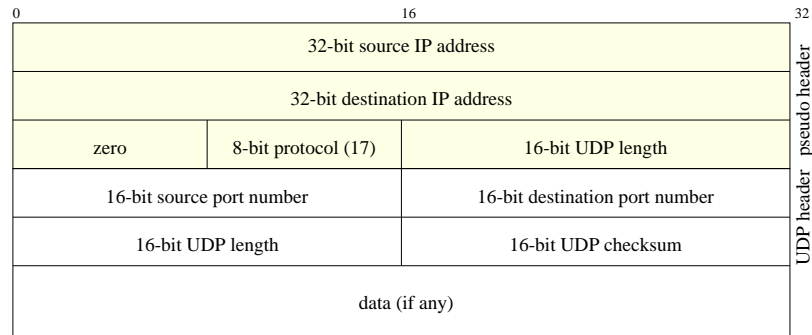


**port:** “TSAP” - identifies process sending/receiving datagram

**length:** including header (IP length → redundant)

**checksum:** across *pseudo-header*, UDP header, data for end-to-end reliability

## UDP Checksum Computation



- checksum may be turned off for better performance *at your peril*
- similar for IPng (but mandatory – why?)

## UDP Port Numbers

- < 512: reserved for particular services (one per host)
  - 512 . . . 1023: privileged port (Unix superuser only)
  - 1023 . . . 65536: available for applications
- |     |                             |
|-----|-----------------------------|
| 7   | echo                        |
| 9   | discard                     |
| 19  | character generator         |
| 37  | time                        |
| 53  | domain name service         |
| 123 | Network Time Protocol (NTP) |