IPv6: Internet Protocol version 6

- Why?
- Addressing
- Packet format
- Fragmentation
- Coexistence
- Control messages (ICMPv6)
- Getting an address (DHCPv6, Autoconfig.)
- Finding neighbours (ND)
- Naming things (DNS)

Why we need IPv6

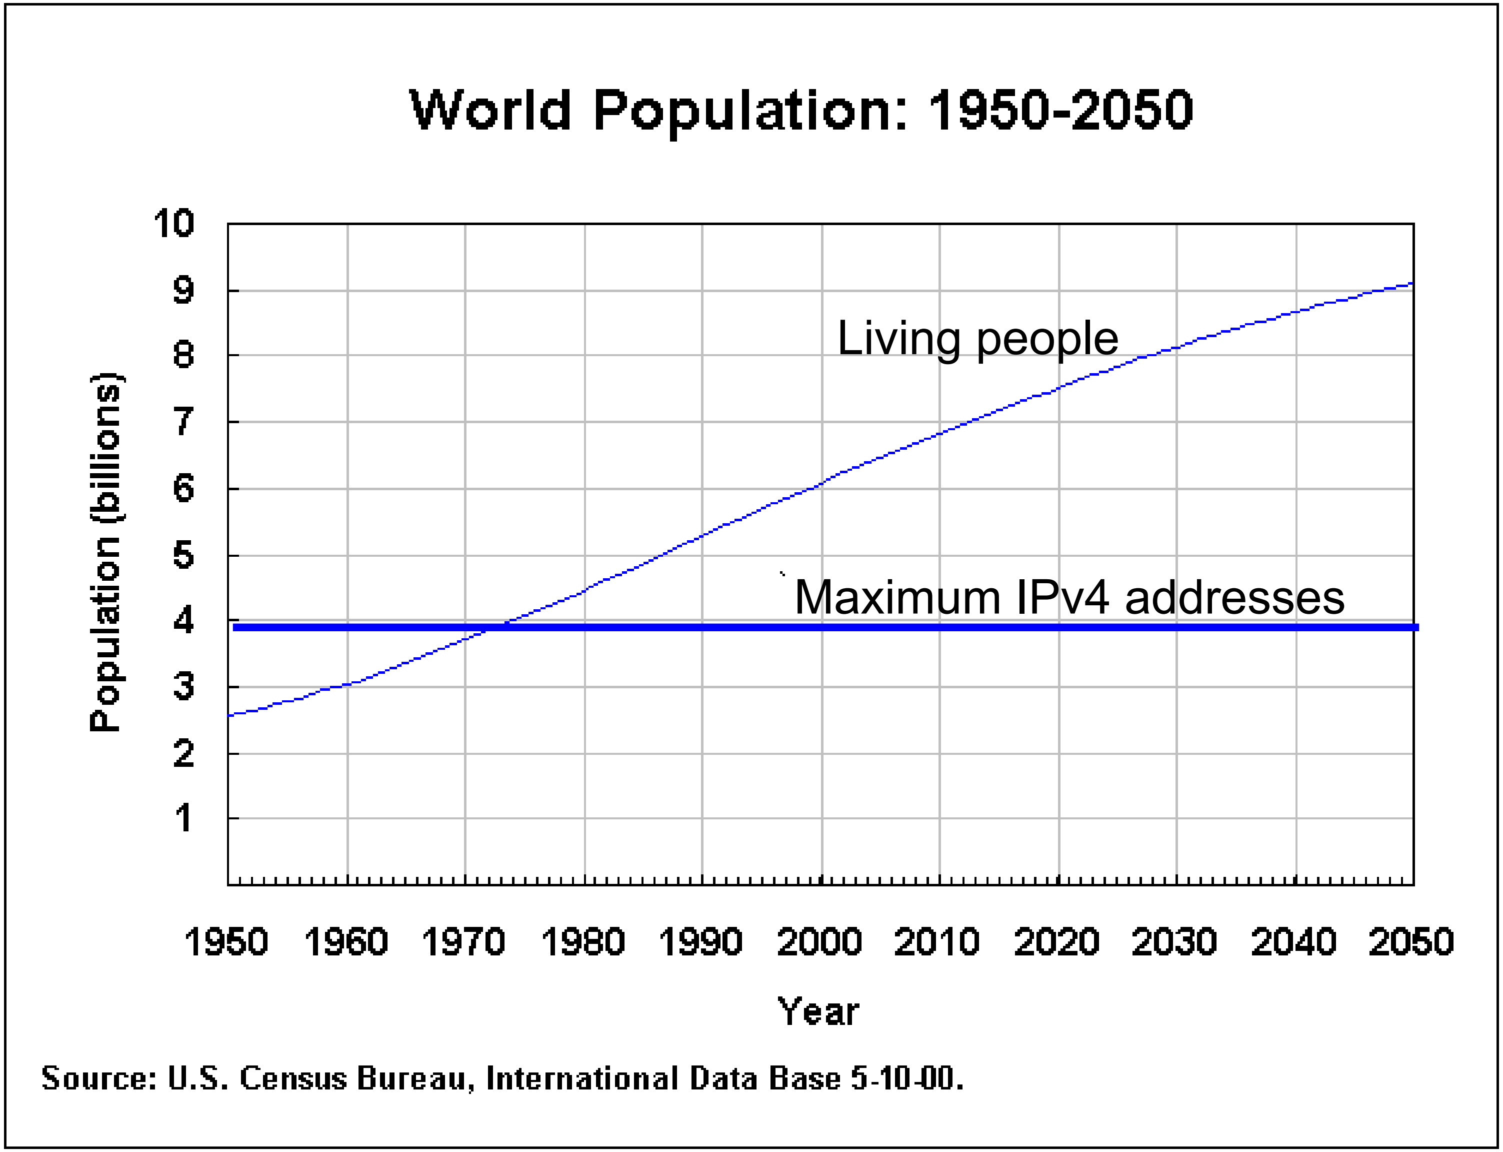

Obviously, having fewer addresses than people is silly

IPv6 in a nutshell

- New version of IP with bigger addresses
- Designed starting in 1994
  - operational experimentally in 1997
- Major deployments starting now
  - US Federal Government requirement in 2008
- Connectionless datagram approach doesn't change
- Will co-exist with IPv4 for many years

IPv6 Address Format

- In the abstract, it's just a 128 bit binary number
- Conventionally written in "colon-separated hexadecimal:"
  `2610:00a0:c779:000b:0000:0000:d1ad:35b4`
  abbreviated as
  `2610:a0:c779:b::d1ad:35b4`
- Obviously, the routing system has to treat it separately from IPv4

Location versus Identity

```
|          n bits                  |          128-n bits           |
+----------------------------------+-------------------------------+
|       routing prefix             |          interface ID         |
+----------------------------------+-------------------------------+

<----- high order bits indicate       low order bits indicate --->
       location for routing           identity on the LAN
```

- *In many cases the boundary is at /64*

```
| ISP part   ...     Site subnet|            64 bits             |
+-------------------------------+--------------------------------+
|       routing prefix          |          interface ID          |
+-------------------------------+--------------------------------+
```

- *An ISP might allocate a /48 prefix to a site*

```
|   48 bits              | 16 bits|            64 bits           |
+-------------------------------+--------------------------------+
|  ISP prefix            | Subnet |          interface ID        |
+-------------------------------+--------------------------------+
```

Special types of IPv6 address (1)

- IPv6 also supports *multicast* addressing and routing
  - Multicast IPv6 addresses are under prefix FF00::/8

```
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
|1 1 1 1 1 1 1 1 x x x x x x x x x x x x x x x x x x x x ...
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
```

- There is no *broadcast* address in IPv6
- *anycast* is a special use of unicast, as in IPv4

Special types of IPv6 address (2)

- ::/128 (all zeros) means “unspecified”
- ::1/128 is the loopback address (send a packet to yourself)
- FE80::/10 (1111111010xxx...) is “link local” space for isolated networks

Special types of IPv6 address (3)

- ULAs (Unique Local Addresses) are reserved for private use within a site, under prefix FD00::/7

```
| 8 bits  |  40 bits    |  16 bits  |          64 bits           |
+---------+-------------+-----------+----------------------------+
| 11111101| Global ID   | Subnet ID |        Interface ID        |
+---------+-------------+-----------+----------------------------+
```

  - Global ID is a unique pseudo-random value
  - ULAs are therefore unique, unlike IPv4 private addresses; can be safely routed locally

- IPv4 addresses mapped in IPv6 format:

```
|              80 bits                | 16 |      32 bits        |
+-------------------------------------+--------------------------+
|0000..............................0000|FFFF|    IPv4 address     |
+-------------------------------------+----+---------------------+
```

  - Not used on the network; used within IPv6+IPv4 hosts to exchange packets from IPv4 clients with applications

IPv6 Header Format

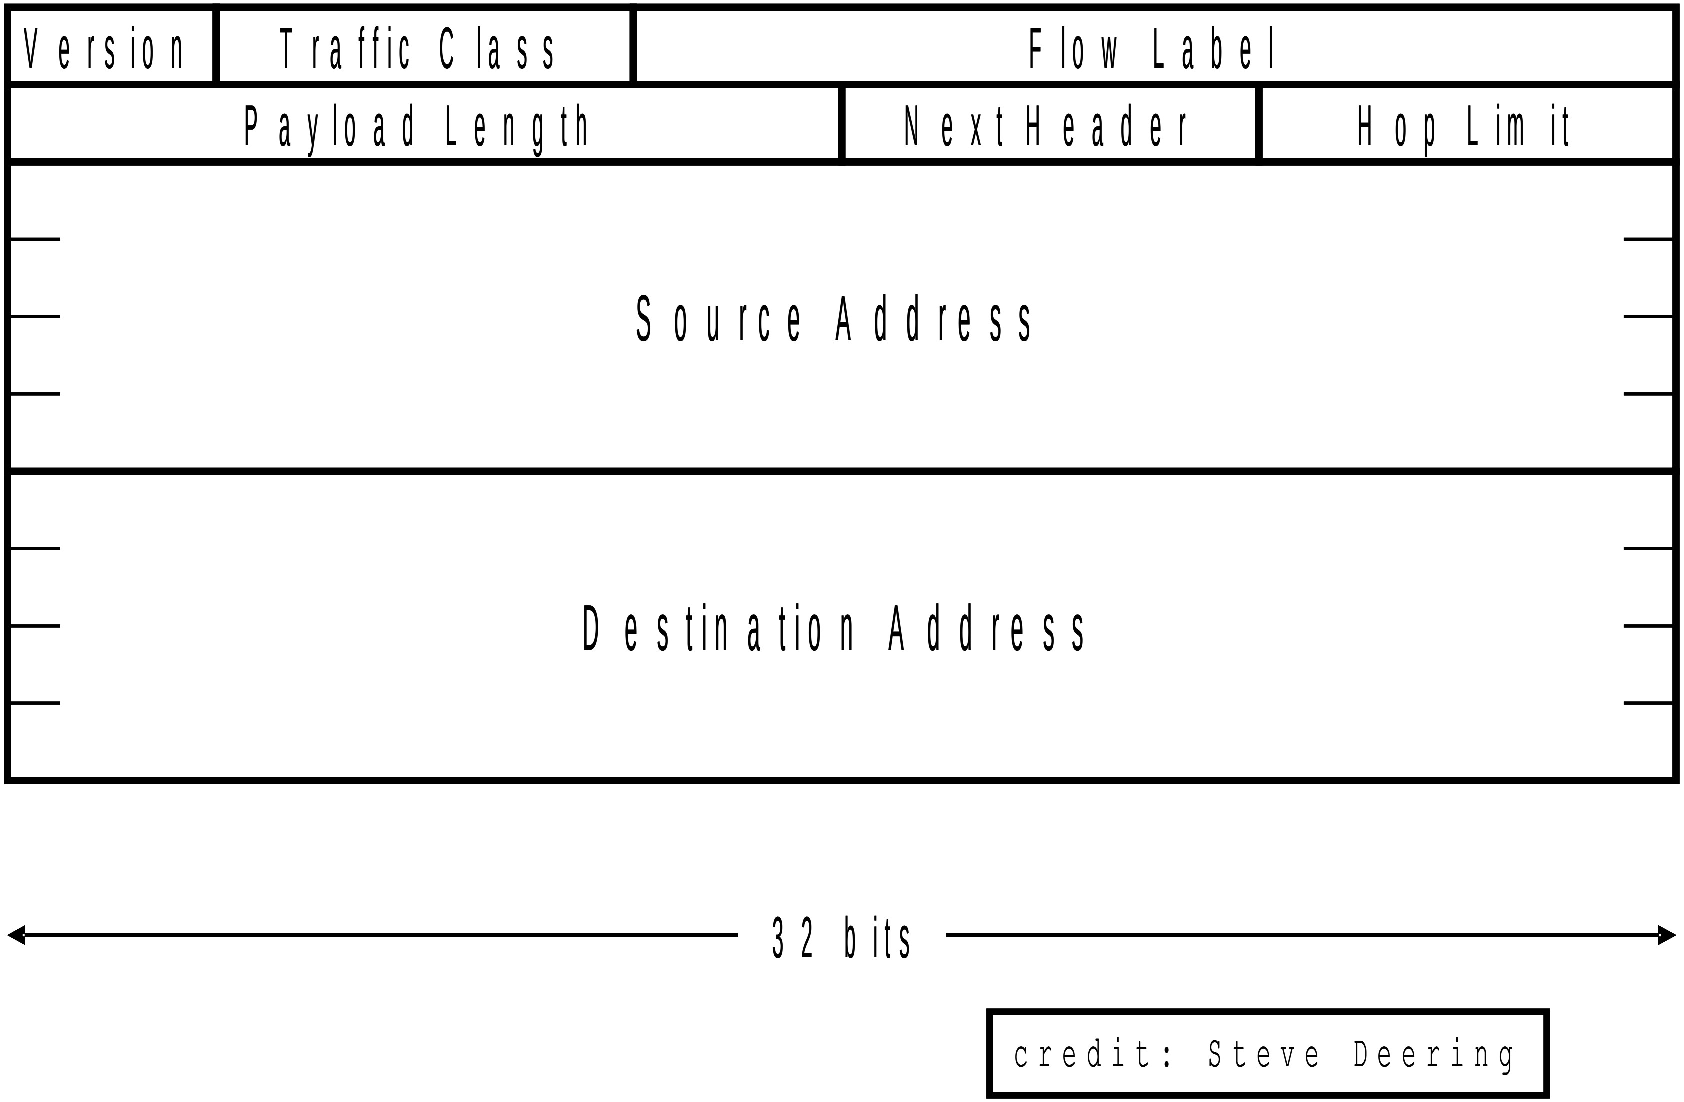

Back to the IPv4 header

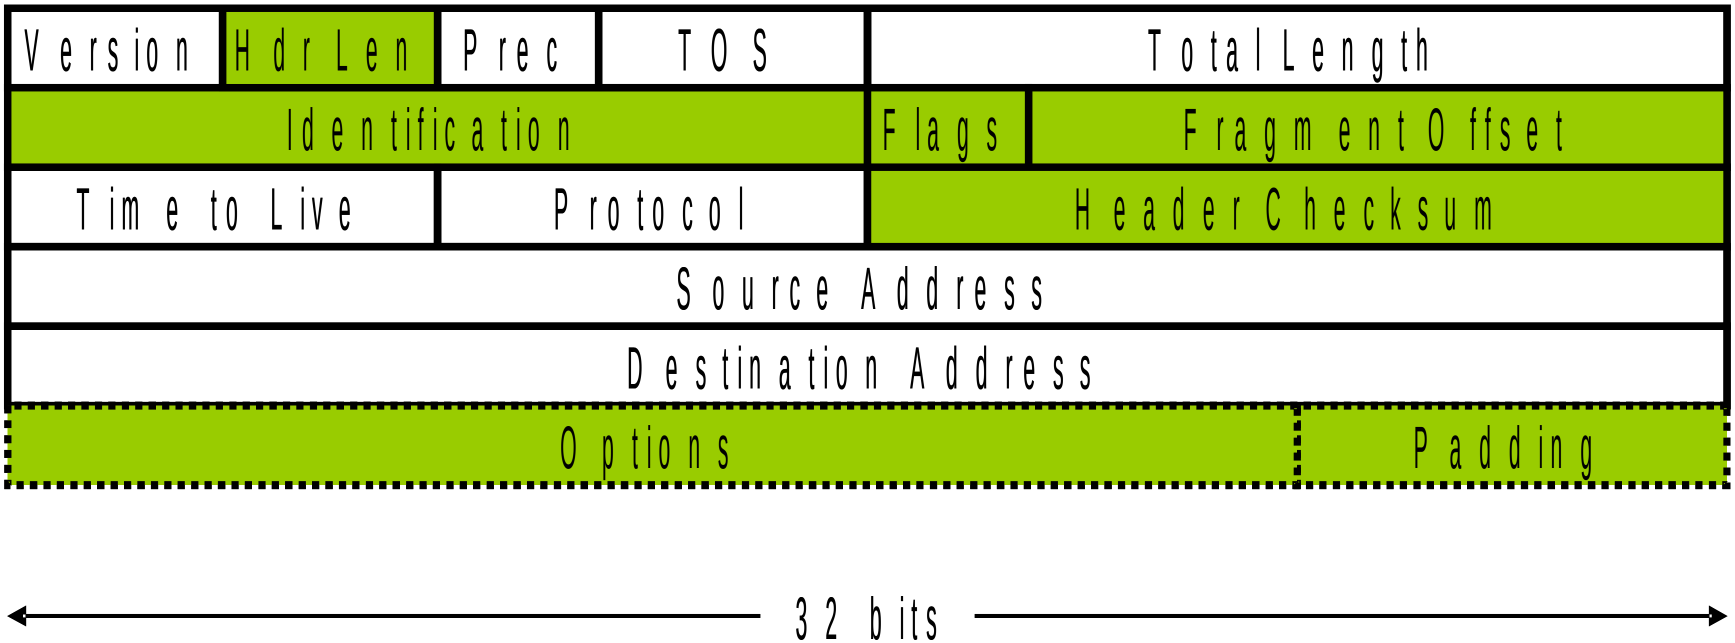

Green shaded parts have been dropped from IPv6.

credit: Steve Deering

Mapping to Layer 2

- The IP packet has to be sent inside a Layer 2 frame, such as an Ethernet frame.
- The exact way this is done depends on the type of Layer 2 link
  - e.g. using Ethertype 0x86DD on Ethernet

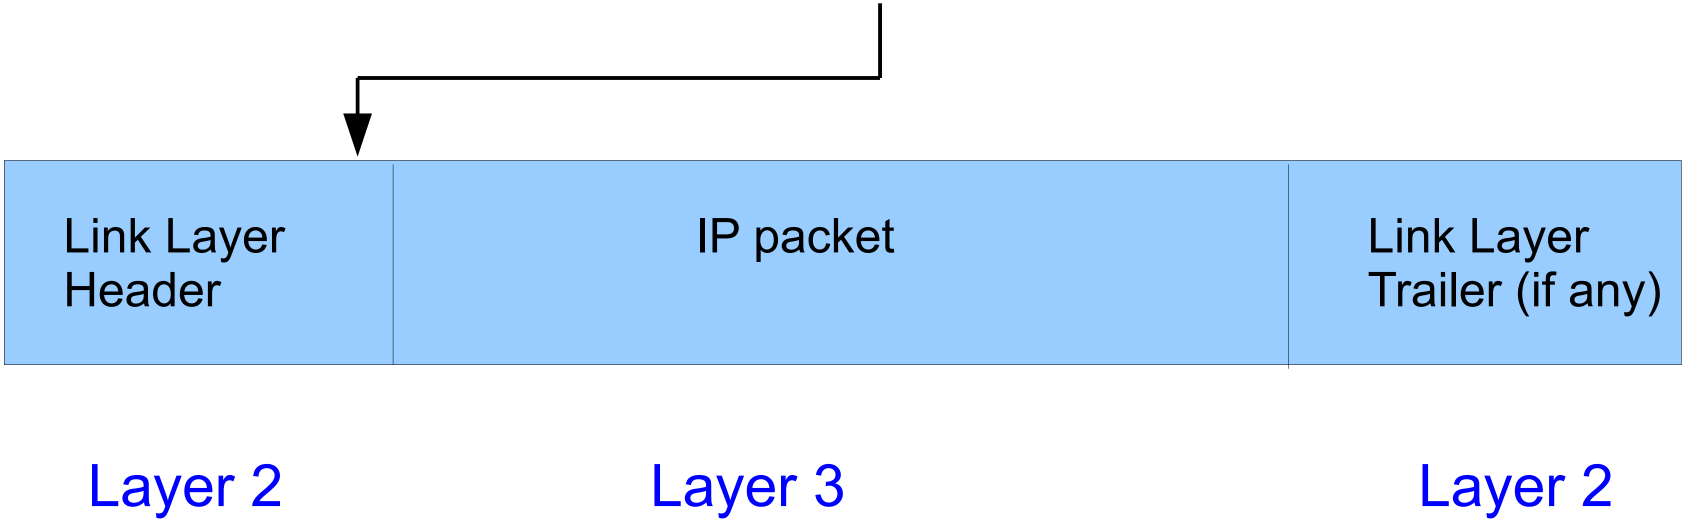

Explanation of IPv6 header

- Version: 6
- Traffic Class Field, identical to DS Field in IPv4
  - 8 bits used to manage quality of service

Wrong in Shay

- Flow Label
  - 20 bits intended for flow-based quality of Service
- Payload length
  - not including header
- Next Header
  - explained below
- Hop Limit
  - Same as IPv4 TTL

Next Header value

- An IPv6 packet can start with a string of headers
  - If there's only the basic header described so far, “Next Header” contains a protocol number just like IPv4, saying that the payload is TCP, UDP, etc.
- Various *optional* additional headers are defined
  - Hop-by-hop options header
  - Destination options header
  - Routing headers (several types)
  - Fragment header
  - and others
- Each one includes a new “Next Header” value
  - The last one is always the payload protocol

IPv6 Packets with Headers

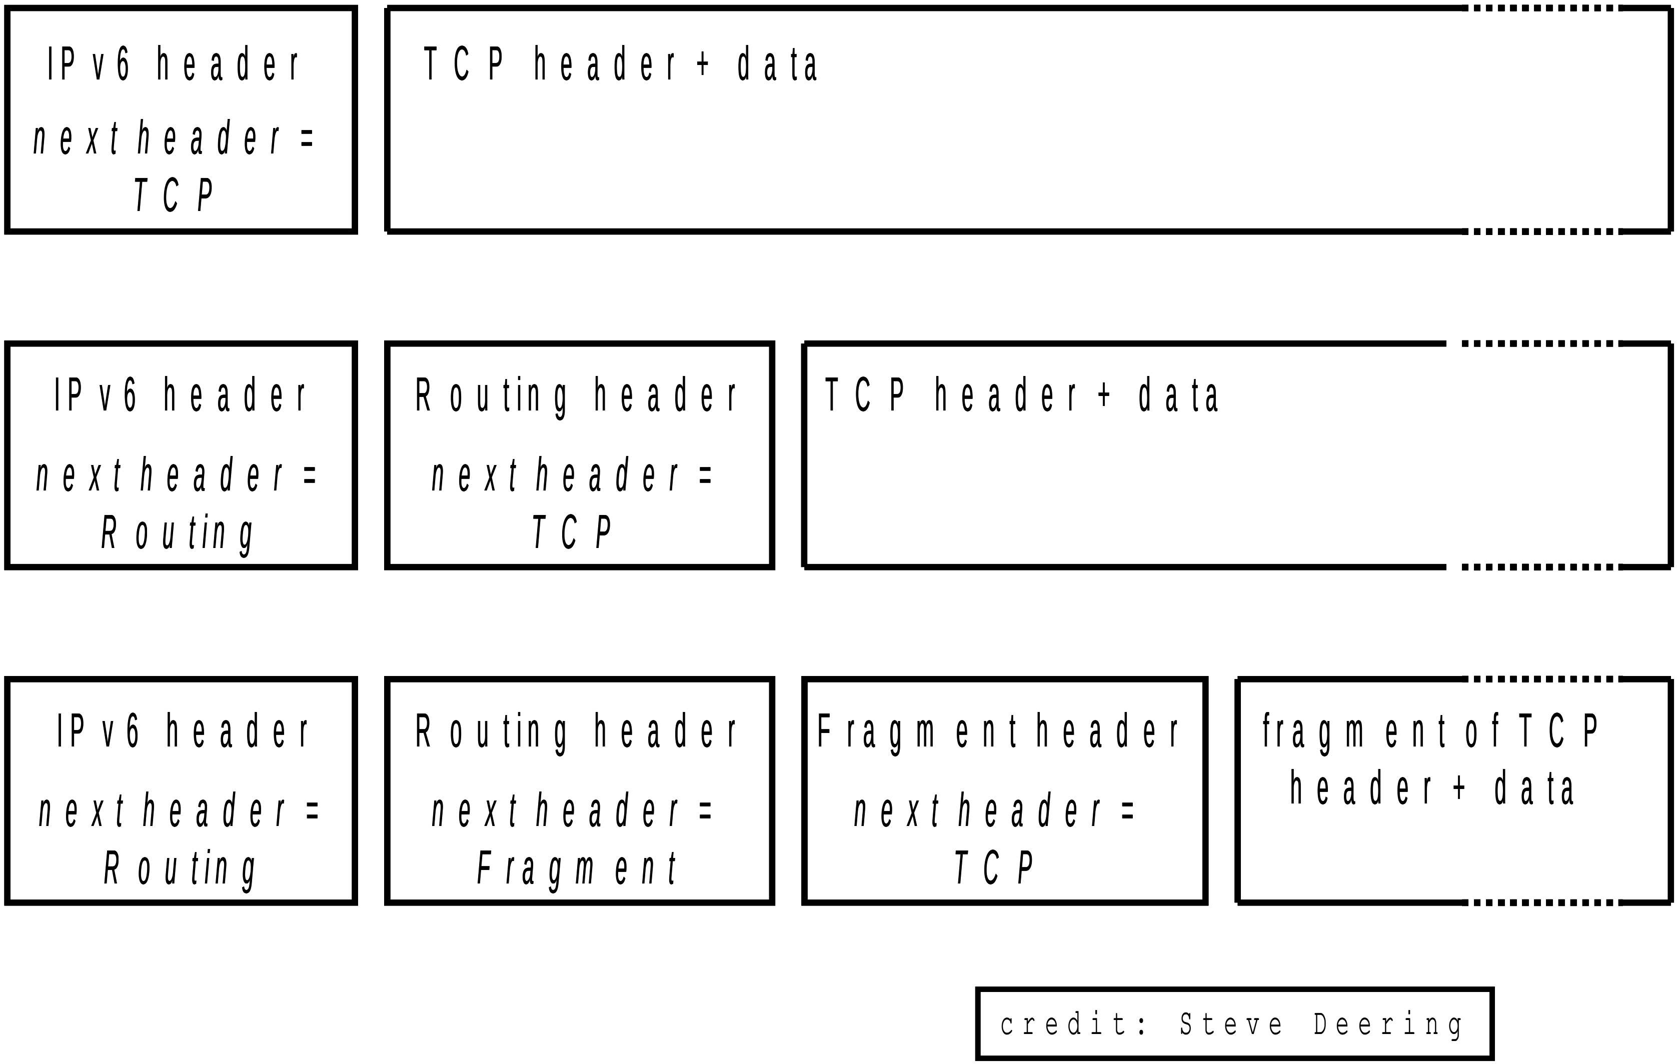

Fragmentation

- IPv6 requires that every link in the Internet has an MTU of 1280 bytes or greater
  - Any link incapable of this must fragment *at link level*
- IPv6 fragmentation is *only* done by the sending host, never by routers
  - Sender must determine path MTU size
- Fragmentation header details based on IPv4 experience

```
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Next Header  |   Reserved    |      Fragment Offset    |Res|M|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                         Identification                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

  - M=1 for more fragments, M=0 for last fragment
  - Res=Reserved

ICMPv6 and DHCPv6

- We'll skip the details
- They are both similar to but different in detail from the IPv4 versions

Routing for IPv6

- RIP, OSPF, BGP4 come in IPv6 versions
  - no change in principle
  - known as RIPng, OSPFv6 and BGP4+

Getting an address without DHCP: IPv6 Stateless Auto-configuration

- Intended for "dentist's office" scenario (i.e. no manual configuration needed)
- Nodes start by acquiring a Link Local address using the FE80::/10 prefix
- Router issues Router Advertisements to provide a routeable prefix for new nodes
  - unique global address formed from that prefix
- Nodes then use Neighbor Discovery and Duplicate Address Detection procedures to find neighbours
  - ARP experience showed that broadcast is not a good approach (risk of “broadcast storms”)
  - Therefore, IPv6 uses local multicast for ND

Auto-configuration functions

- Router Discovery
- Prefix Discovery
- Parameter Discovery
- Address Autoconfiguration
- Address Resolution
- Next-hop Determination
- Neighbour Unreachability Detection (NUD)
- Duplicate Address Detection (DAD)
- Redirect: router supplies better first-hop

Auto-configuration messages

- Router Solicitation*
- Router Advertisement*
- Neighbour Solicitation*
- Neighbour Advertisement*
- Redirect

All sent as ICMPv6 messages.

** May be sent to multicast addresses that don't "wake up" everybody, unlike ARP multicast*

Forming an address automatically

- Prefix (normally 64 bits)
  - Initially , FE80::/64 (link local)
  - Secondly, prefix received in Router Advertisement
- Interface Identifier (normally 64 bits)
  - Simplest: Ethernet address padded out to 64 bits
    `34 56 78 9A BC DE` becomes
    `3656:78FF:FE9A:BCDE`
    (16 bits inserted, and U/L bit inverted)
  - Privacy addresses: choose a pseudo-random value
  - Secure ND: a cryptographically generated value

DNS for IPv6

- A records carry 32-bit IPv4 addresses
- AAAA records carry 128-bit IPv6 addresses
- DNS queries for AAAA records can travel over IPv4 or IPv6
- A modern resolver returns both A and AAAA records

IPv4 and IPv6 coexistence

- The old and new versions will have to live together and work together for many years
- IPv6 can be carried over IPv4 in “tunnels”
  - IPv6 packets encapsulated in IPv4 packets
- Servers and ISPs will become “dual stack,” able to support IPv4 and IPv6 clients simultaneously
- Application proxies will be able to map IPv4 clients to IPv6 servers, or the opposite
- Direct translation of v4 to v6 at packet level doesn't work well

Coexistence mechanisms (simple version)

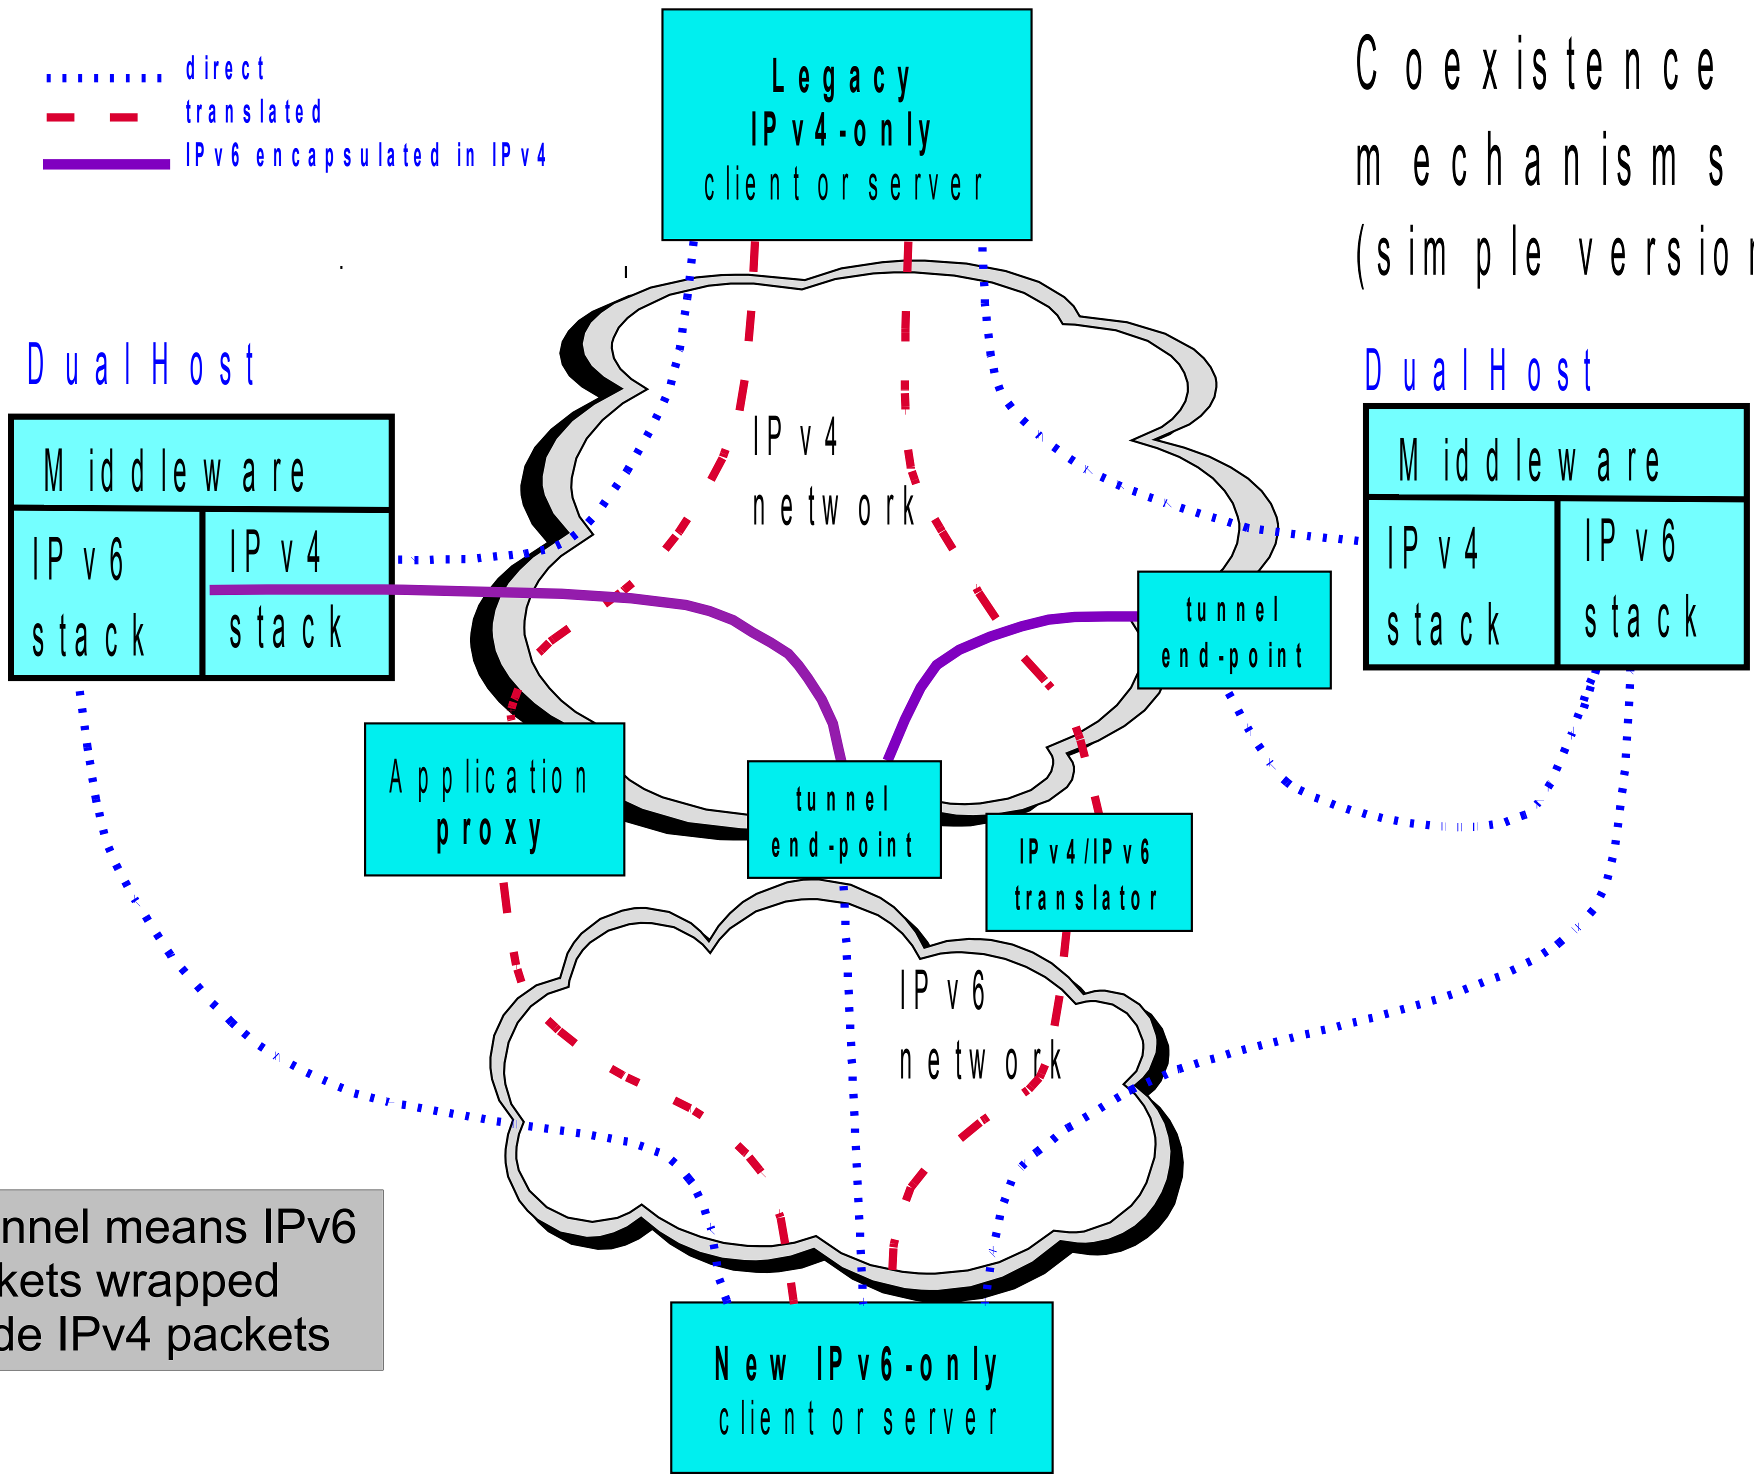

A tunnel means IPv6 packets wrapped inside IPv4 packets

References

- Shay 11.3
  - bugs:
    - "priority" and "flow label" out of date on page 562
    - ignore the "registry" bits in Fig.11.20 and page 568
    - IPv4-compatible format (Fig. 11.22(b)) is obsolete and the whole discussion of that figure is confused
- IPv6 Essentials by Silvia Hagen
- Lots of RFCs:
  2460 (protocol), 4861+4862 (autoconfig), 4291 (addressing), 4294 (node requirements - lists many important RFCs), etc., etc.