Internet Protocol version 6 and transition techniques

Piotr Szymański
053750

34352-Advanced Protocols

May 02 2006

Presentation overview

- Introduction to IPv6
- Addressing
- Packet structure
- Transition mechanisms

Motivation for creation of IPv6

- **Problems with IPv4**
  - Basically: shortage of IPv4 addresses

- **Solutions:**
  - Develop a completely new protocol
  - Develop a new version of IP, which will address the problems.

Overview of changes in IPv6

- Larger address space
- Hierarchical addressing
- Simpler autoconfiguration
- Better support for multicasting
- Added *anycast* – a new type of address
- Simplified packet header format
- Improved support for extensions and future options
- Quality of service support
- Built-in authentication and privacy capabilities
- Large („jumbo”) packets support

IPv6 timeline

- Early 1990s – IETF begins working on a successor to IPv4.
- Late 1994 – IPng adoption by IETF.
- December 1995 – RFC 1883: „Internet Protocol, Version 6 (IPv6) Specification”.
- December 1998 - RFC 2460, obsoletes 1883.
- 2006 – IPv6 still not widely adopted.

Problems with transition

- IPv6 is not compatible with IPv4
- Transition cannot happen „overnight”
- NAT and CIDR slowed down the adoption of IPv6
- Compatibility with TCP, UDP, ICMP and other Layer 4 and Layer 7 protocols.
- Software and hardware changes necessary.
- „US Dept. of Commerce: US shift to IPv6 could cost $75B” – *Internetnews.com*
- Chicken and egg problem

Addressing

ADDRESSING

Address size and address space

- An IPv6 address is a 128-bit number
- Supports 3.4×10^{38} addresses

IPv4 address space

Address notation and prefix notation

- Text representation of addresses using hexadecimal notation:

ABCD:EF01:0000:0000:6789:EF01:0345:6789

- No need to write leading zeros in a group:

ABCD:EF01:0:0:6789:EF01:345:6789

- Zero compression:

ABCD:EF01::6789:EF01:345:6789

Address notation and prefix notation

- Notation for mixed IPv4 and IPv6 environments:

0:0:0:0:0:FFFF:129.144.52.38

- Text representation of address prefixes:

ipv6_address/prefix_length

For example:

ABCD:EF01:0000:0000:6789:EF01:0000:0000/96

Types of IPv6 addresses

In IPv6 there are 3 types of addresses:

- *unicast* – An identifier for a single network interface.

  A packet is delivered to the interface identified by this address.

- *anycast* – An identifier for a set of interfaces.

  A packet is delivered to the „nearest" interface identified by this address.

- *multicast* – An identifier for a set of interfaces.

  A packet is delivered to all interfaces identified by this address.

Address space allocation

Address type	Binary prefix	IPv6 notation
Unspecified	000....000 (128 bits)	::/128
Loopback	000....001 (128 bits)	::1/128
Multicast	11111111 (8 bits)	FF00::/8
Link-Local unicast	1111111010 (10 bits)	FE80::/10
Global unicast	everything else	

Global unicast address format

Original TLA/NLA scheme (RFC 2374)

3	13	8	24	16	64
FP	TLA	RES	NLA	SLA	Interface ID

FP	Format Prefix (001)
TLA ID	Top-Level Aggregation Identifier
RES	Reserved for future use
NLA ID	Next-Level Aggregation Identifier
SLA ID	Site-Level Aggregation Identifier
Interface ID	Interface Identifier

Global unicast address format (contd.)

Current scheme (RFC 3587)

n bits	m bits	128-n-m bits
global routing prefix	subnet ID	Interface ID

- No complex inherent structure – more allocation freedom for Internet Registries
    - global routing prefix – value assigned to a site
    - subnet ID – identifier of a link within a site
    - interface ID – interface on a link
- All Global Unicast addresses with a global routing prefix starting with binary 000 must have a 64-bit Interface ID (i.e. n+m = 64).

Multicast addresses

8	4				4	112
11111111	0	R	P	T	scope	group ID

- A set of 4 binary flags:
  - The first bit is reserved, and set to 0.
  - The R and P bits are used for Rendezvous Point embedding and Unicast-Prefix-based IPv6 addresses, respectively.
  - The T bit indicates a permanently-assigned („well-known”) multicast address (0), or a non-permanently-assigned („dynamic”) one (1).
- Scope is a 4-bit multicast scope value used to limit the scope of the multicast group:

1	Interface-Local	5	Site-Local
2	Link-Local	8	Organization-Local
4	Admin-Local	E	Global

Multicast addresses (contd.)

Pre-Defined Multicast Addresses

Multicast Addresses	Description
FF0X:0:0:0:0:0:0:0	Reserved. X ranging from 0 to F.
FF01:0:0:0:0:0:0:1 FF02:0:0:0:0:0:0:1	All Nodes Addresses. Interface-Local and Link-Local, respectively.
FF01:0:0:0:0:0:0:2 FF02:0:0:0:0:0:0:2 FF05:0:0:0:0:0:0:2	All Routers Addresses. Interface-Local, Link-Local and Site-Local respectively.
FF02:0:0:0:0:1:FFXX:XXXX	Solicited-Node Address. Formed from node's Unicast or Anycast addresses.

Anycast addresses

- An address assigned to more than one interface.
- A packet sent to a *anycast* address is delivered to the „nearest” interface having this address.
- Allocated from the *unicast* address space.
- Expected use – identify a set of routers belonging to an organization providing Internet service.
- Subnet-Router address:

n bits	128-n bits
subnet prefix	0000000000000000000000000

Special addresses

- The Unspecified address:

0:0:0:0:0:0:0:0 (or ::)

- The Loopback address:

0:0:0:0:0:0:0:1 (or ::1)

- The Link-Local address:

10	54	64
1111111010	0	Interface ID

- IPv4-Mapped IPv6 address:

80	16	32
0000...000	FFFF	IPv4 address

Packet structure

PACKET STRUCTURE

Header format

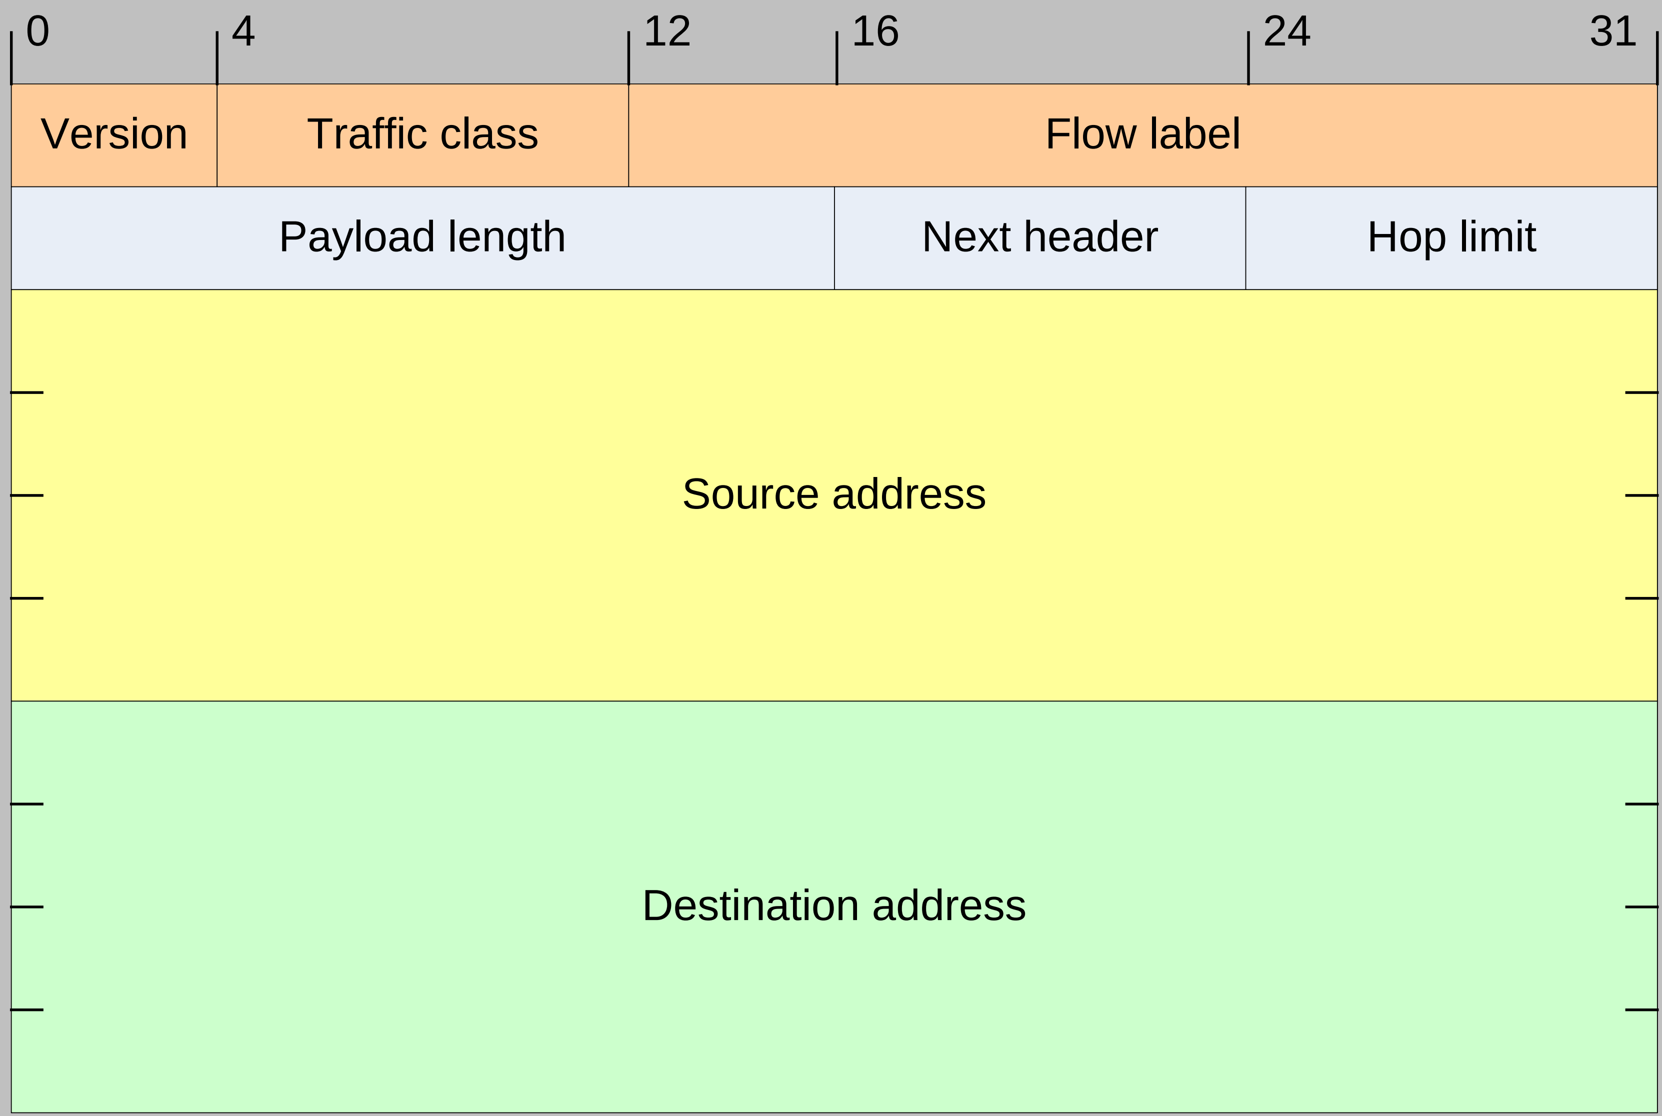

Header format (contd.)

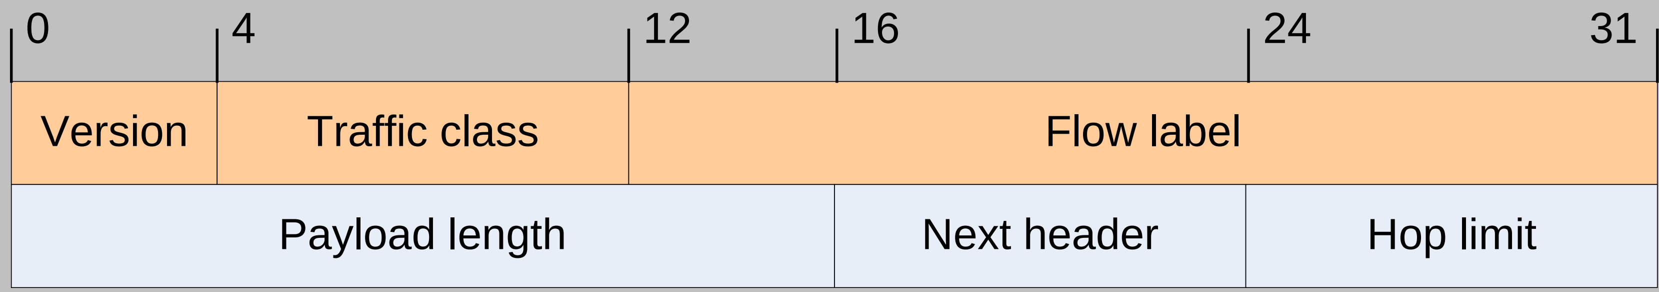

Version - 4 bit Internet Protocol version number = 6; the only common field with IPv4

Payload length

- 16 bit unsigned integer; length of the rest of the payload in octets (maximum size = 65,535 bytes)

Next header

- type of the **extension header** or type of upper layer protocol

Hop limit - IPv4 TTL field equivalent

Size comparison with IPv4

IPv4 header	IPv6 header
20 bytes	40 bytes

Extension headers

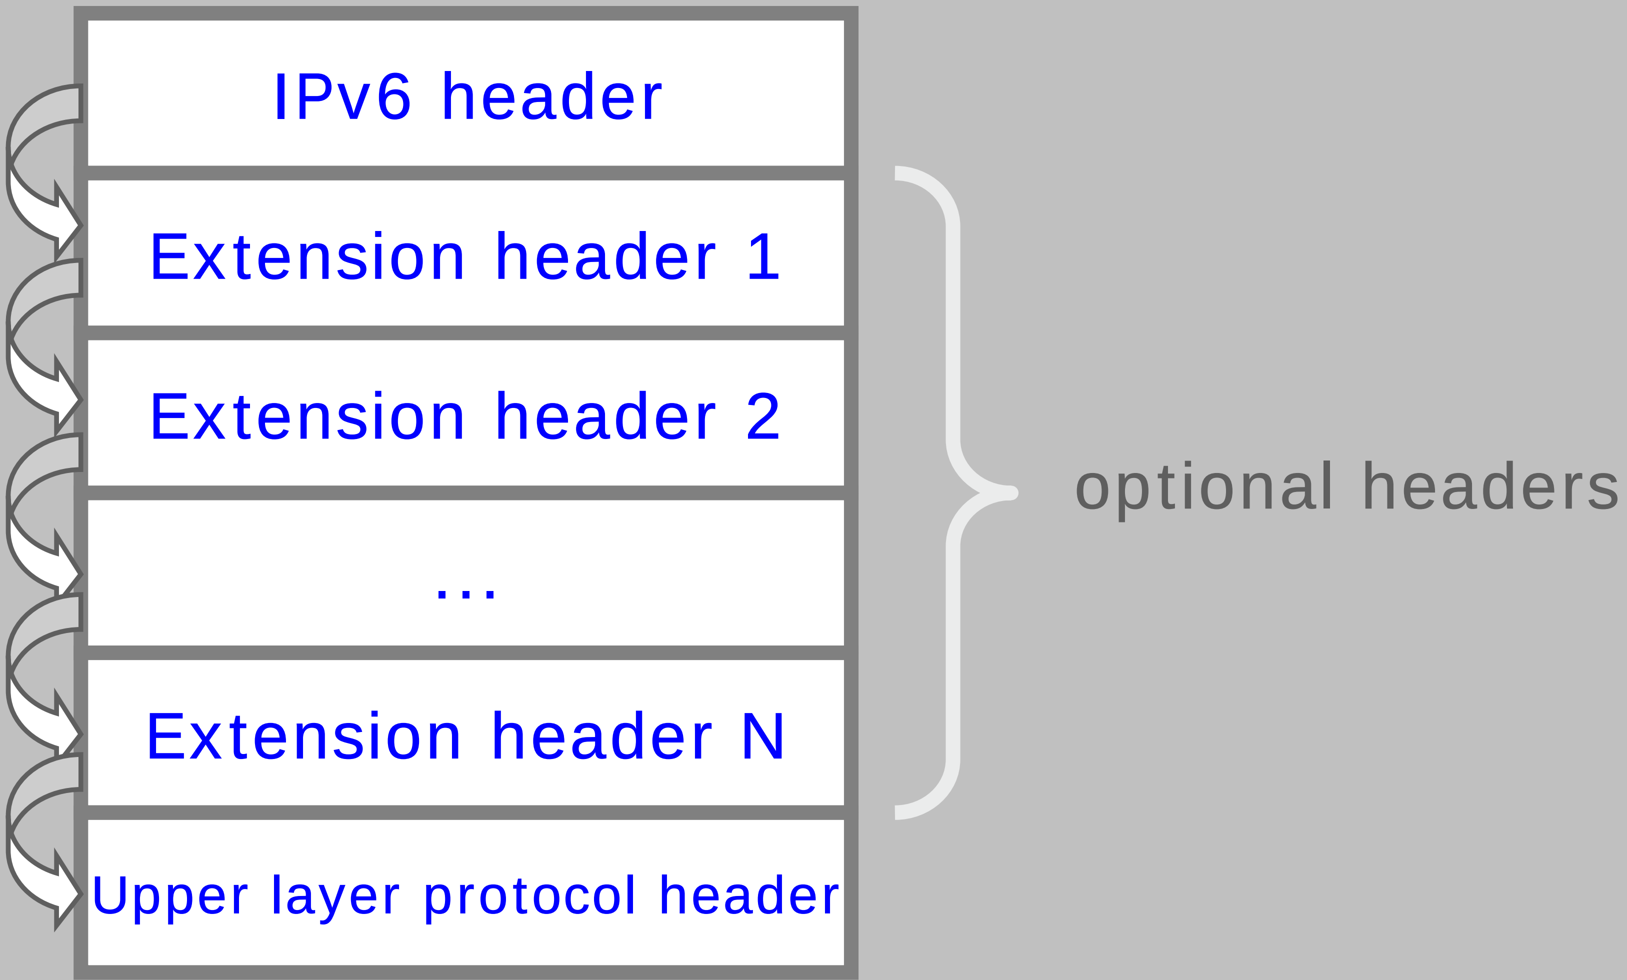

Extension headers (contd.)

Recommended order of headers

IPv6 header
Hop-by-Hop Options header
Destination Options header
Routing header
Fragment header
Authentication header
Encapsulating Security Payload
Destination Options header
upper-layer header
data

- Do not need to be examined by any intermediate node along the transmission path
- Must be processed in the order they appear in the packet
- Each header must appear only once
- A full IPv6 implementation must support all shown extension headers

Routing extension header

0	8	16	24 31
Next Header	Extension Hdr Len	Routing Type	Segments Left
Reserved			
Address 1			
...			
Address N			

Fragment extension header

0	8	16	29	31
Next Header	Reserved	Fragment Offset	R	M
Identification				

Fragment Offset – 13 bit; the offset in 8-octet units, of the data following this header, relative to the start of the Fragmentable part in the original packet

M – bit field; 1 = more fragments, 0 = last fragment

Identification – different than any fragmented packet with the same source and destination IP addresses sent recently

Fragmentation in IPv6

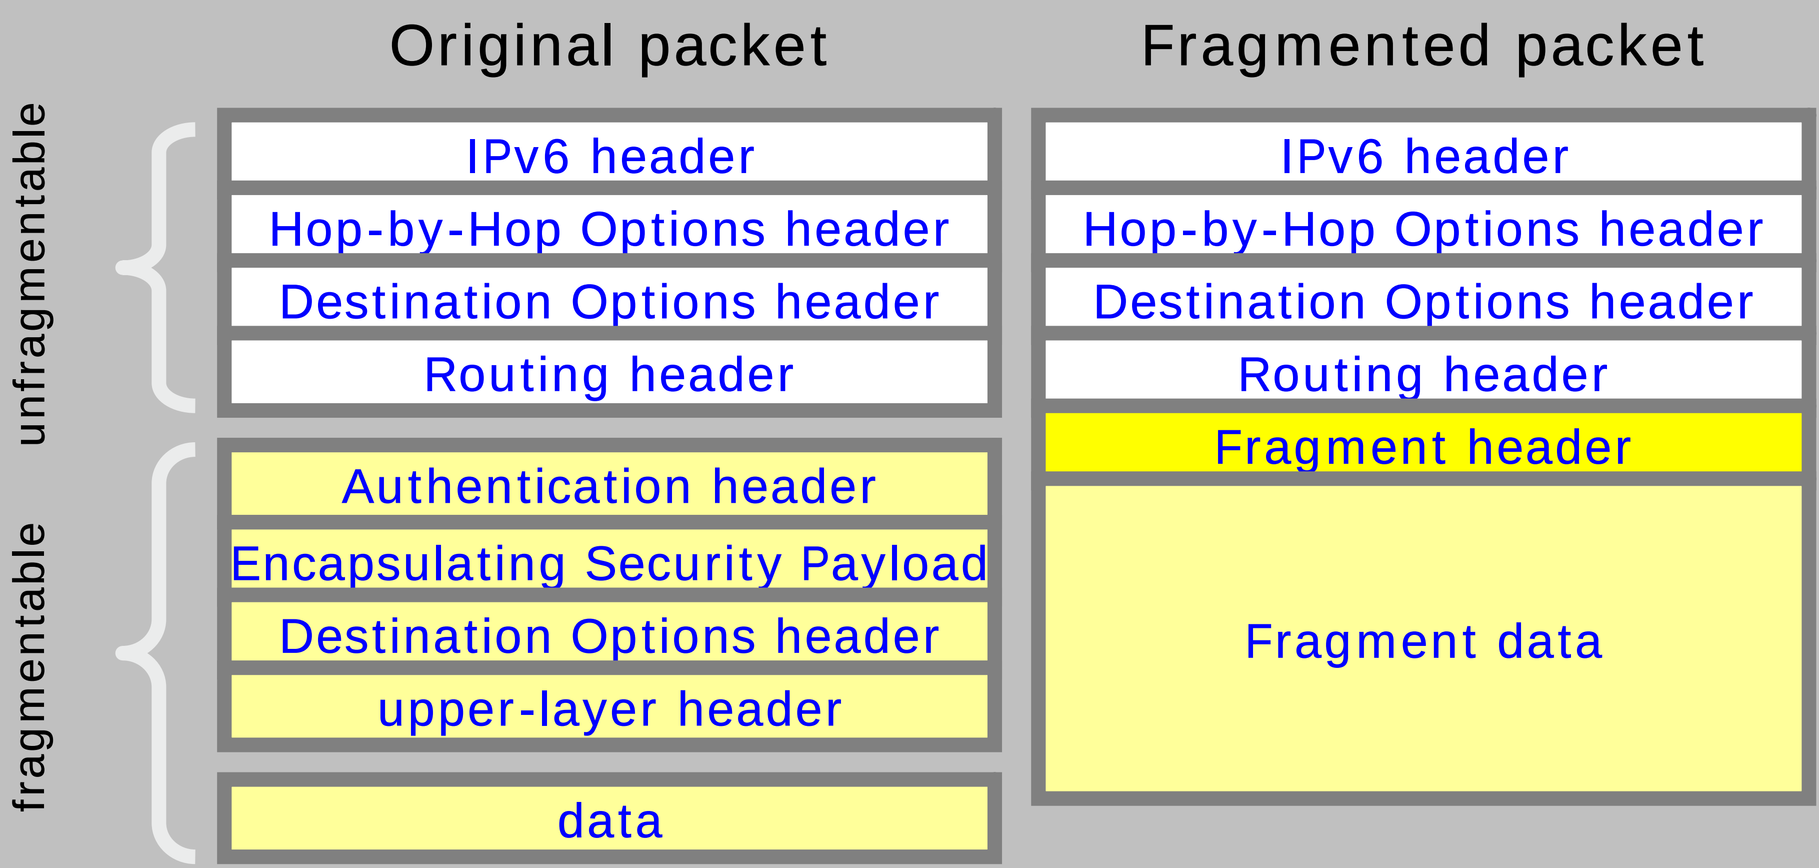

Fragmentation in IPv6 (contd.)

- Source-only fragmentation
  - Intermediate routers cannot perform fragmentation
  - ICMPv6 „Packet too big” message
- Path MTU Discovery
  - Separate protocol for discoverying minimum path MTU
  - Optional – doesn’t need to be implemented by simple embedded devices
- Implications
  - Routers can process packets faster
  - The route between source and destination must be reasonably static

Transition mechanisms

TRANSITION MECHANISMS

Dual IP Stack

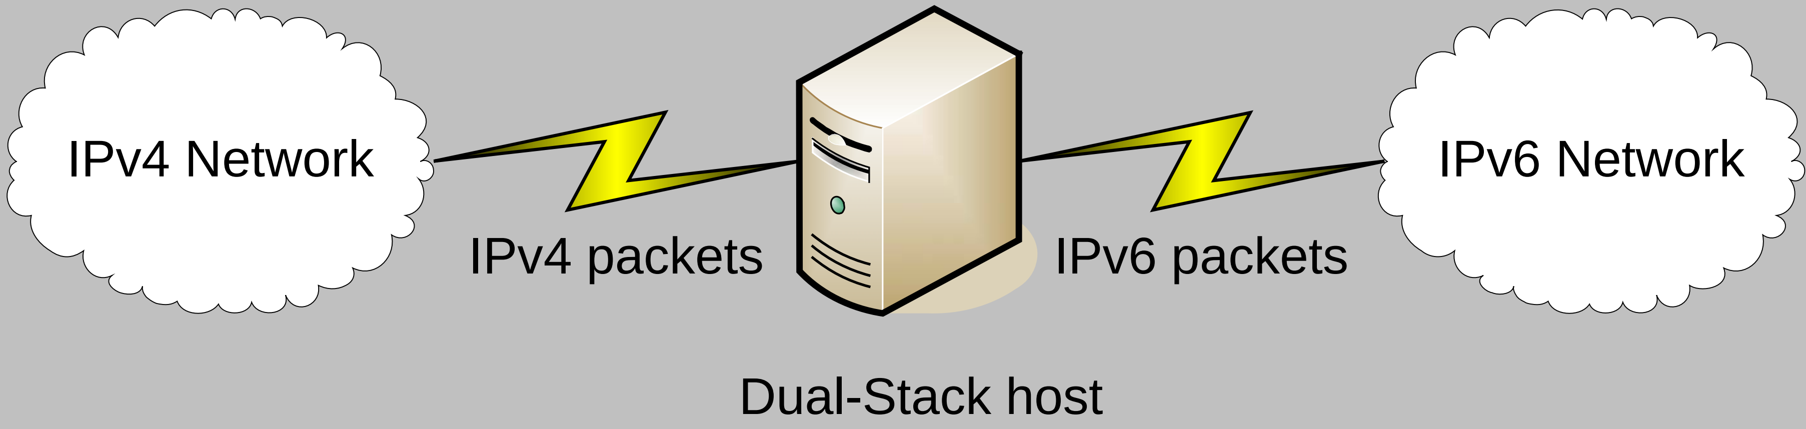

- Does not affect existing IPv4 applications.
- Host must have both IPv4 and IPv6 addresses and direct access to both networks.
- Use of IPv4-Mapped IPv6 addresses for IPv6-only applications communicating with IPv4 network.

Configured tunneling

Connect IPv6 regions by using existing IPv4 infrastructure:

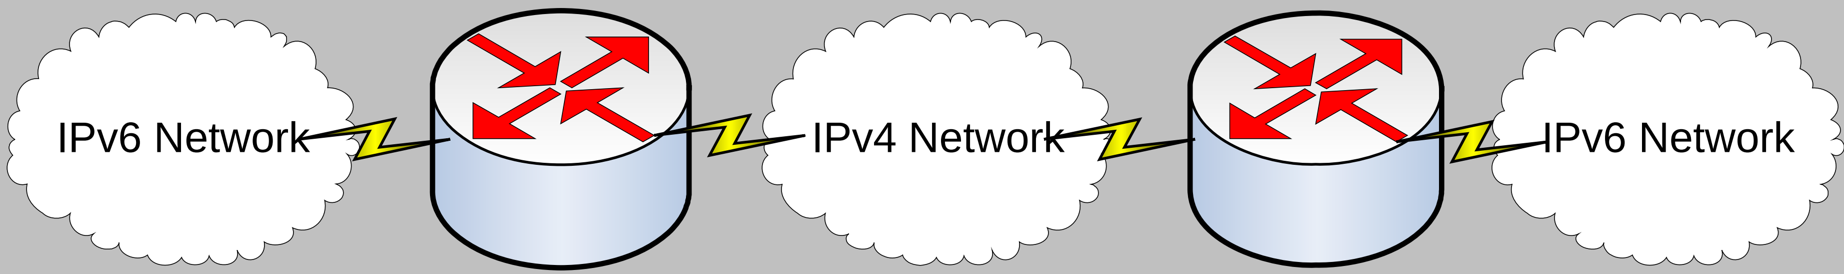

Packet encapsulation:

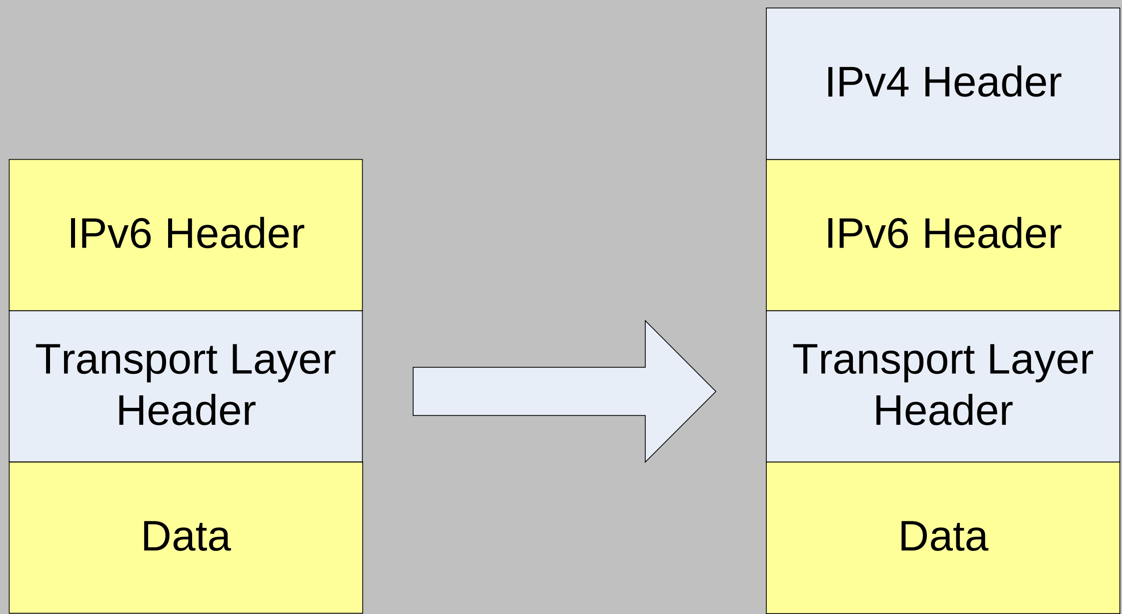

Automatic tunneling (6to4)

Assign a unique IPv6 prefix to any site that has at least one globally unique IPv4 address:

• A 2002::/16 prefix is assigned for this method. A site can construct its own IPv6 address:

16	32	16	64
2002	IPv4 address	SLA ID	Interface ID

Automatic tunneling (6to4)

6to4-only to 6to4+native-IPv6 connectivity scenario

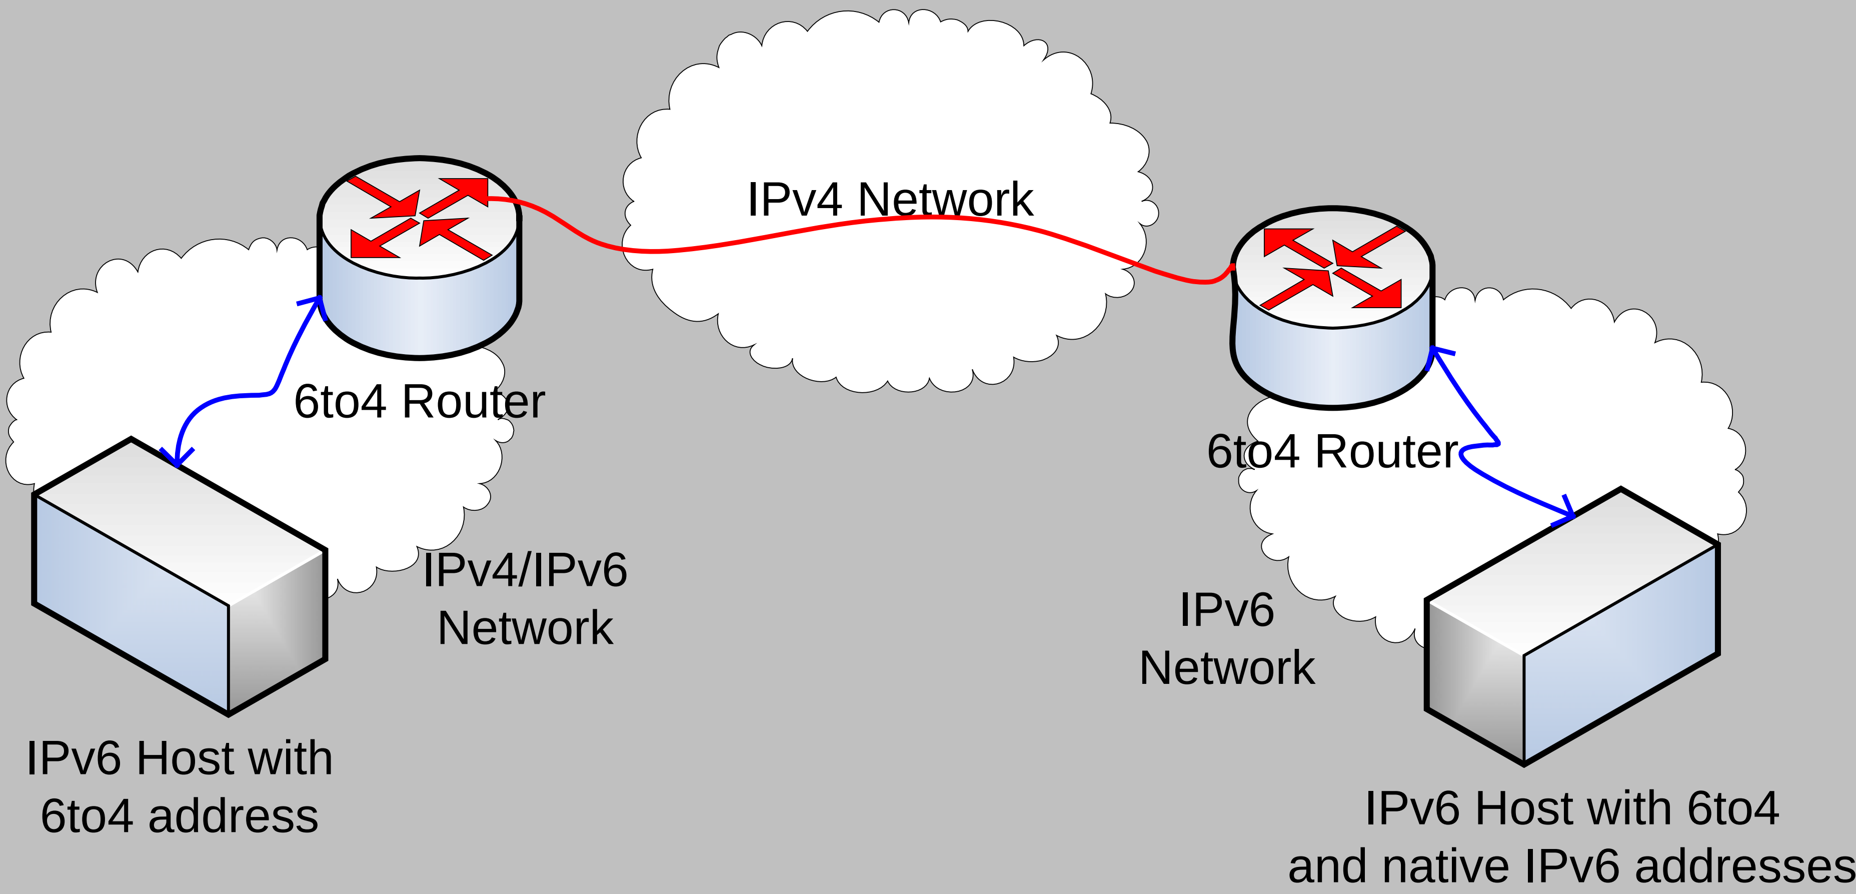

Automatic tunneling (6to4)

6to4-only to native-IPv6-only connectivity scenario

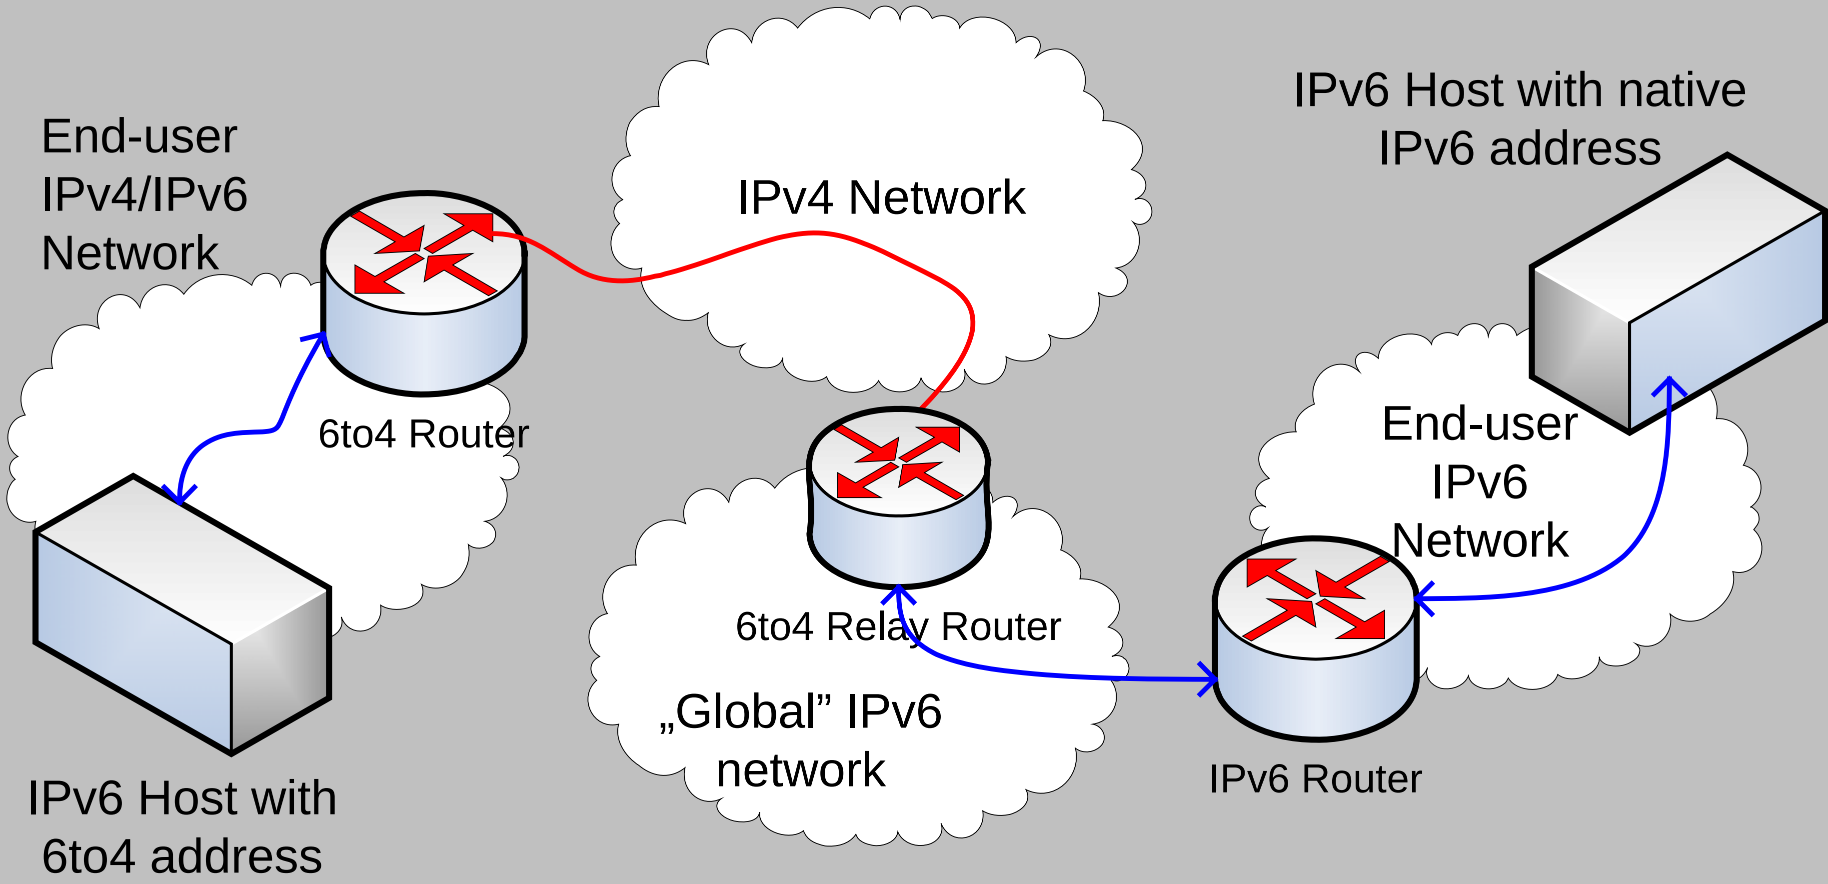

References

REFERENCES

References

1. RFC2460: Internet Protocol, Version 6 (IPv6) Specification. S. Deering, R. Hinden. December 1998.
2. RFC3587: IPv6 Global Unicast Address Format. R. Hinden, S. Deering, E. Nordmark. August 2003.
3. RFC2374 (*obsolete*): An IPv6 Aggregatable Global Unicast Address Format. R. Hinden, M. O'Dell, S. Deering. July 1998.
4. RFC4291: IP Version 6 Addressing Architecture. R. Hinden, S. Deering. February 2006.
5. „The TCP/IP Guide: Internet Protocol Version 6 (IPv6) / IP Next Generation (IPng)", http://www.tcpipguide.com/free/t_InternetProtocolVersion6IPv6IPNextGenerationIPng.htm
6. RFC1881: IPv6 Address Allocation Management. IAB and IESG. December 1995.
7. RFC1924: A Compact Representation of IPv6 Addresses. R. Elz. April 1996.
8. RFC4213: Basic Transition Mechanisms for IPv6 Hosts and Routers. E. Nordmark, R. Gilligan. October 2005.

References

9. „The IPv6 mess”. D.J. Bernstein.
   http://cr.yp.to/djbdns/ipv6mess.html
10. „Could a U.S. Shift to IPv6 Cost $75B?”. S.M. Kerner, Internetnews. December 2005.
   http://www.internetnews.com/infra/article.php/3570211
11. RFC4038: Application Aspects of IPv6 Transition. M-K. Shin, Ed., Y-G. Hong, J. Hagino, P. Savola, E. M. Castro. March 2005.
12. RFC3056: Connection of IPv6 Domains via IPv4 Clouds. B. Carpenter, K. Moore. February 2001.