Introduction to IPv6

Jerome.Durand@renater.fr

Agenda

IPv6 header

IPv6 extensions

IPv6 addressing

IPv6 Header

IPv4 Header

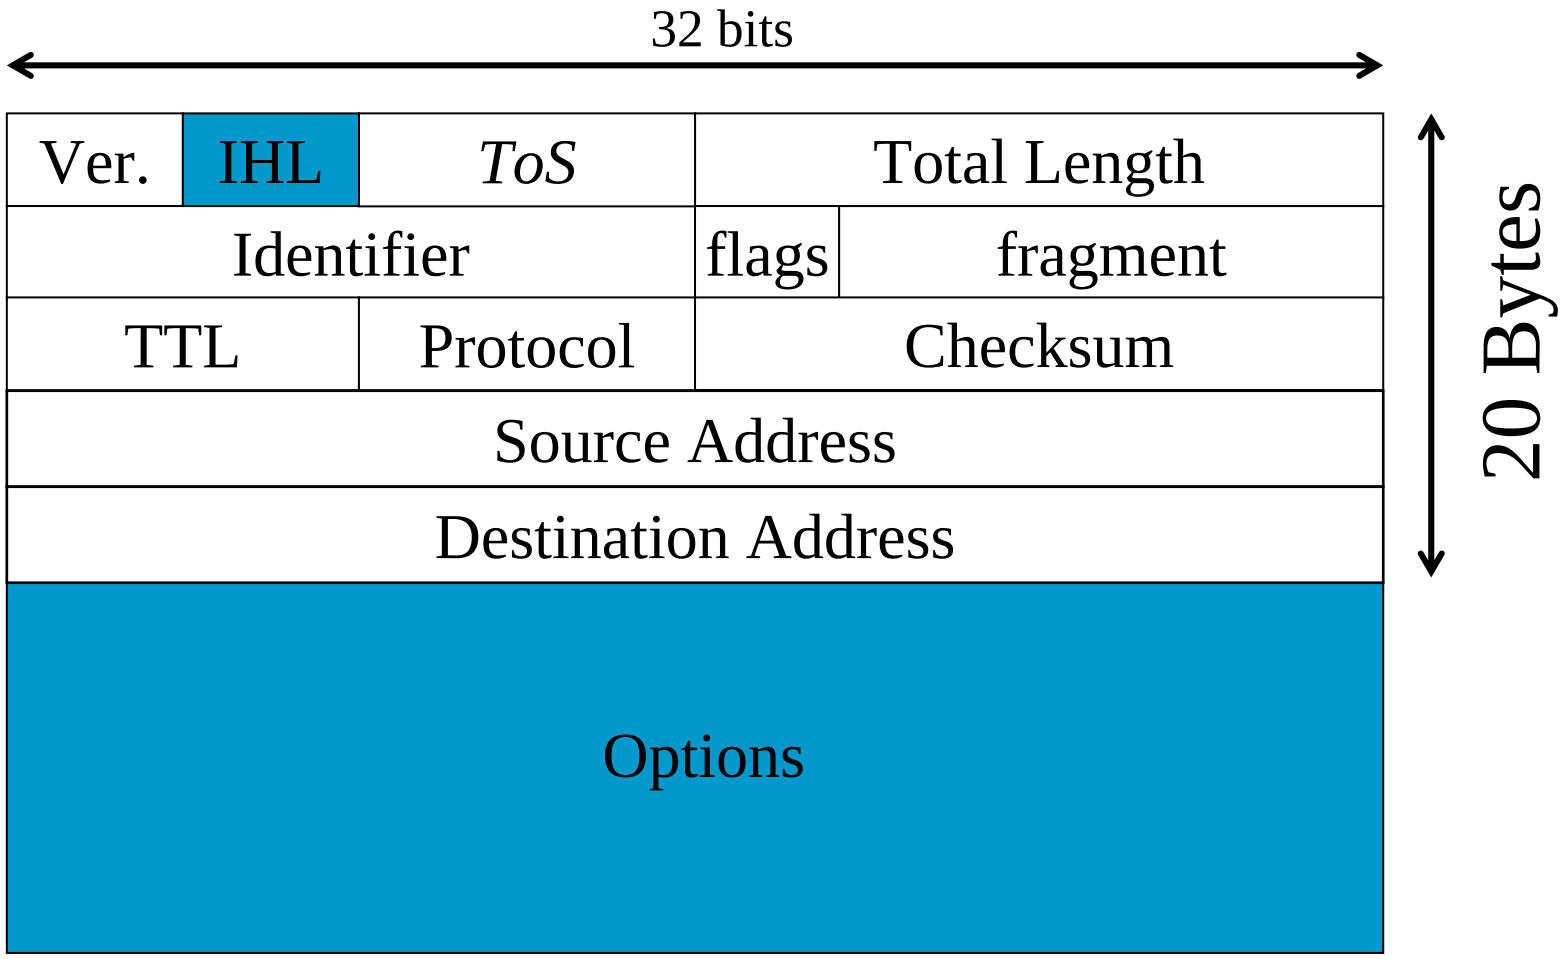
32 bits
Ver.
IHL
ToS
Total Length
Identifier
flags
fragment
TTL
Protocol
Checksum
Source Address
Destination Address
Options
20 Bytes

IPv4 Header

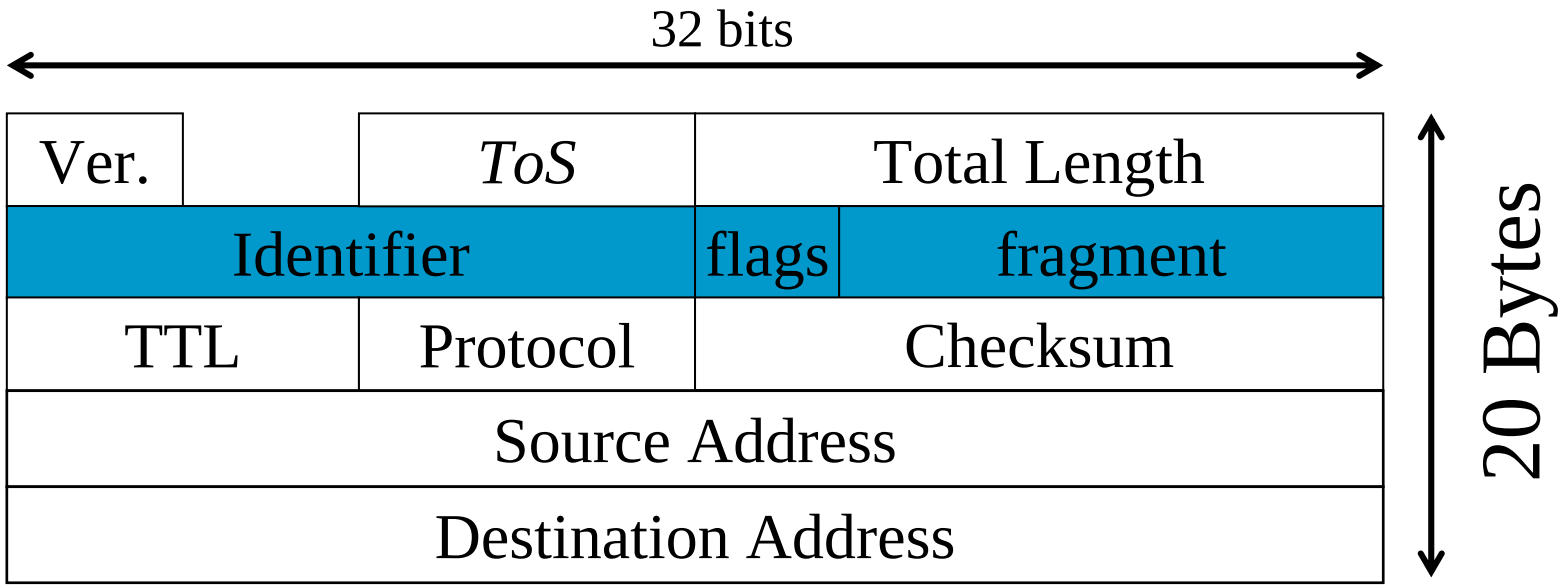
32 bits
Ver.
ToS
Total Length
Identifier
flags
fragment
TTL
Protocol
Checksum
Source Address
Destination Address
20 Bytes

IPv4 Header

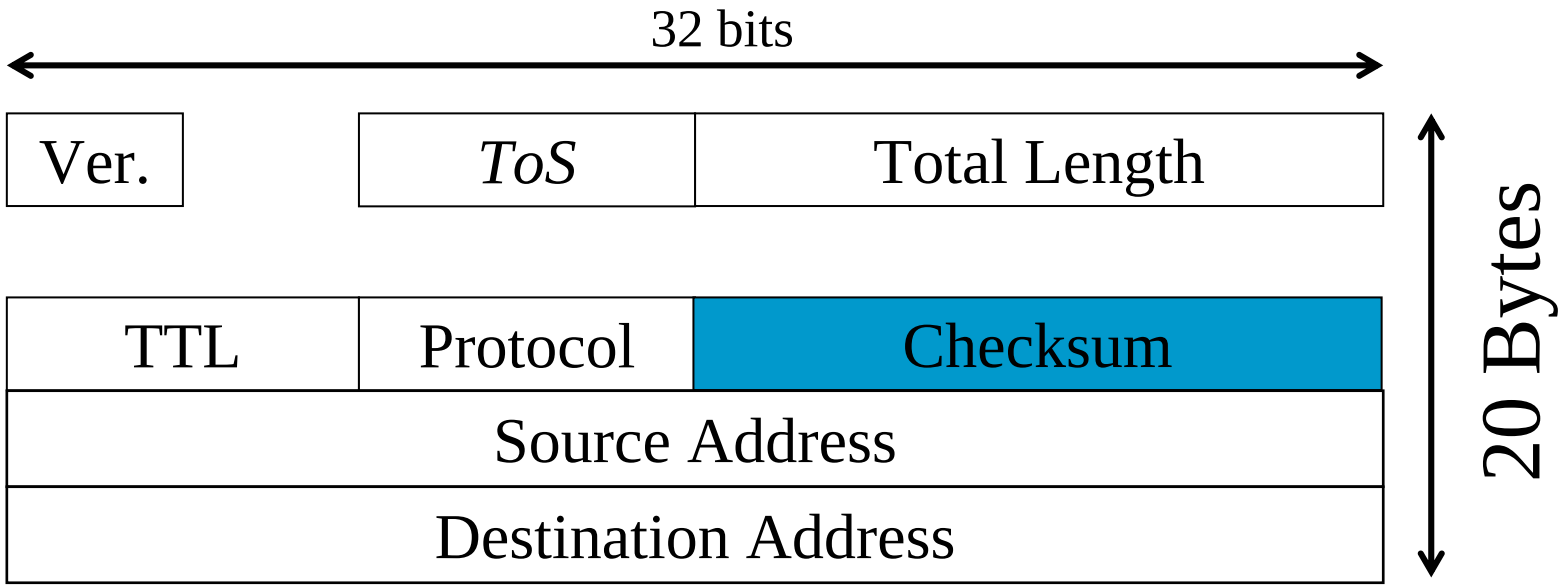
32 bits
Ver.
ToS
Total Length
TTL
Protocol
Checksum
Source Address
Destination Address
20 Bytes

IPv6: Header simplification

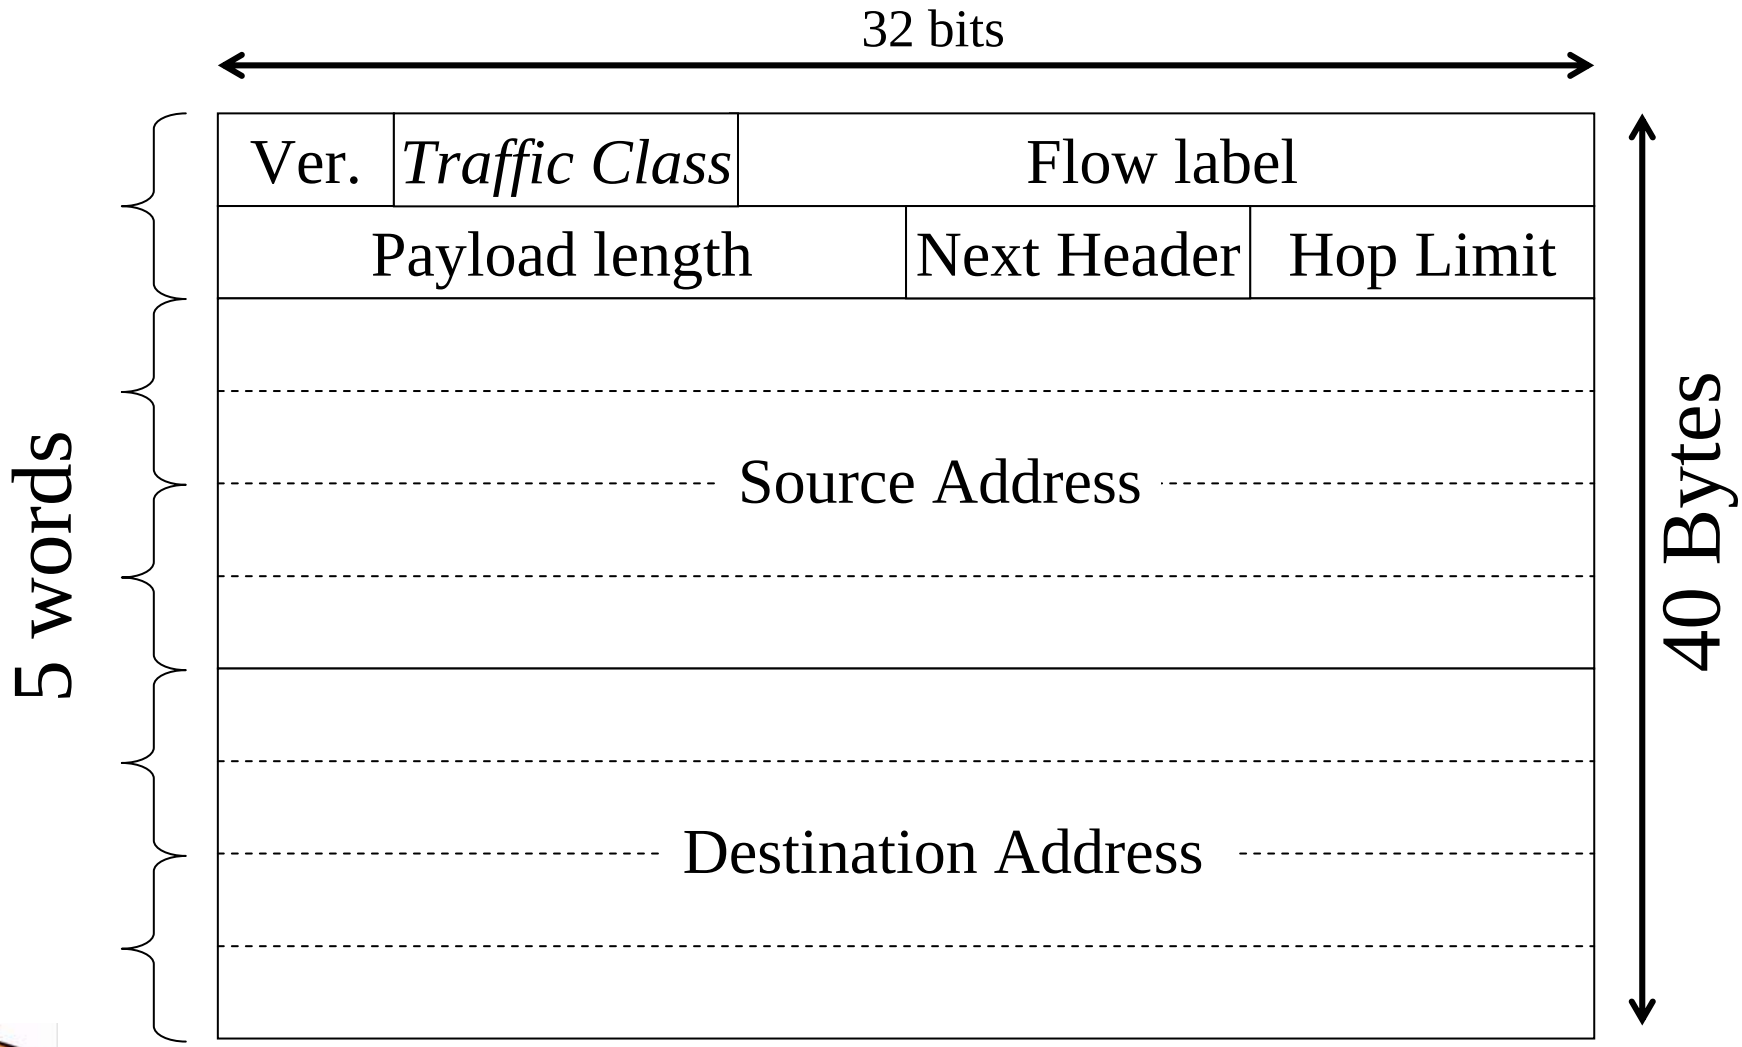

Is it enough for the future ?

- Address length
  - Between 1 564 and 3 911 873 538 269 506 102 addresses by m^2
  - Justification of a fix address length
- Hop Limit
  - Should not be a problem
- Payload Length
  - Use Jumbogram for specific cases

IPv6 extensions

IPv6: Optional headers

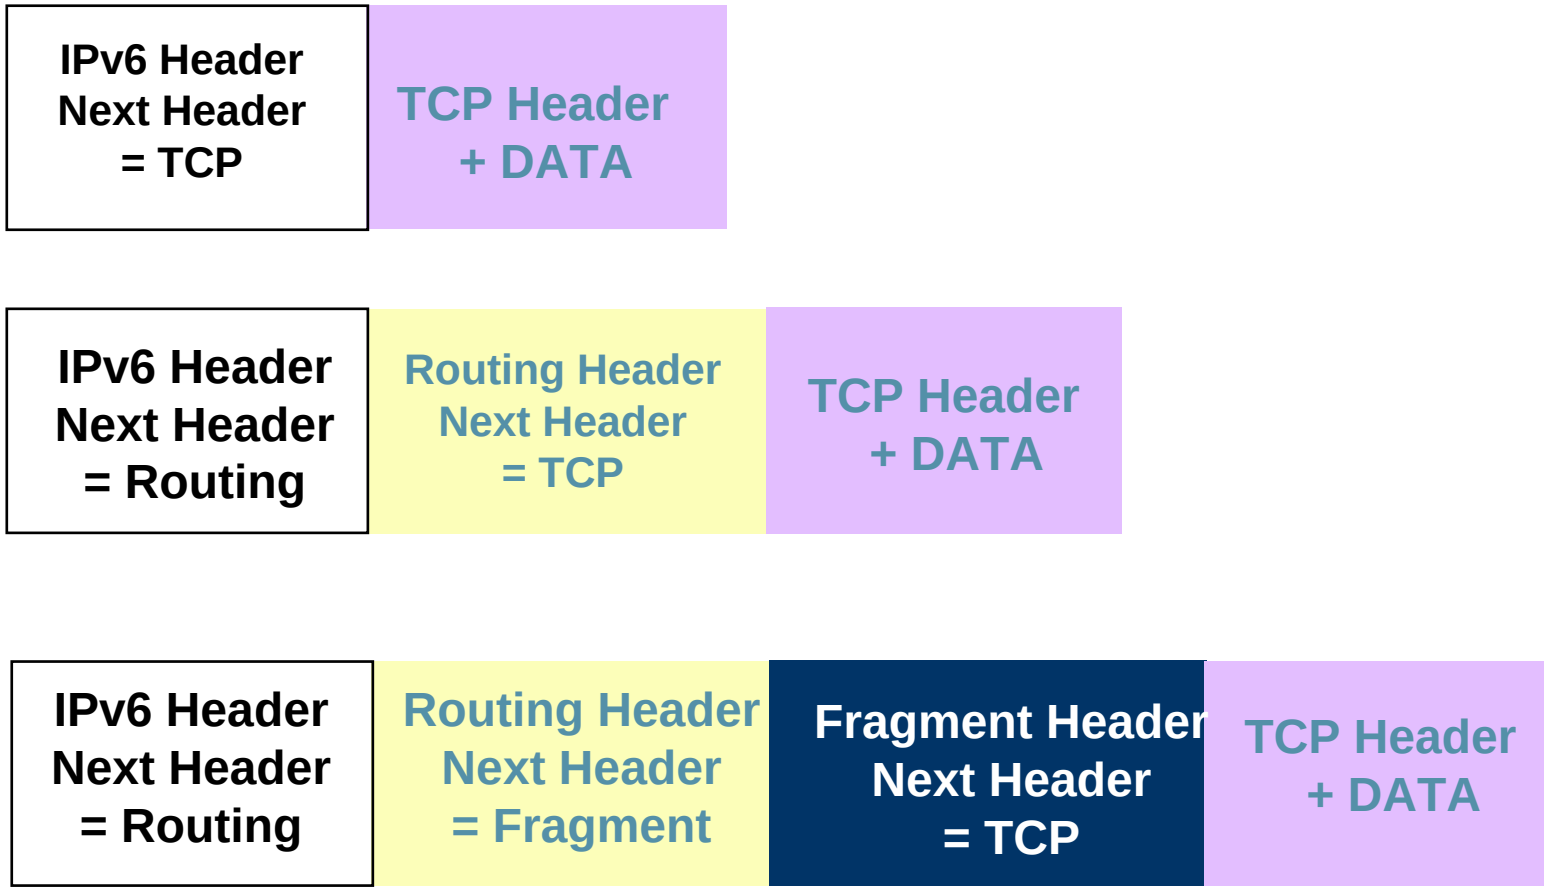

IPv6: Optional extensions

- Hop-by-hop (jumbogram, router alert)
  - Always the first extension
  - Replace IPv4 options,
  - Analyzed by every router.
- Destination
- Routing (loose source routing)
- Fragmentation
- Authentication
- Security

v4 options vs. v6 extensions

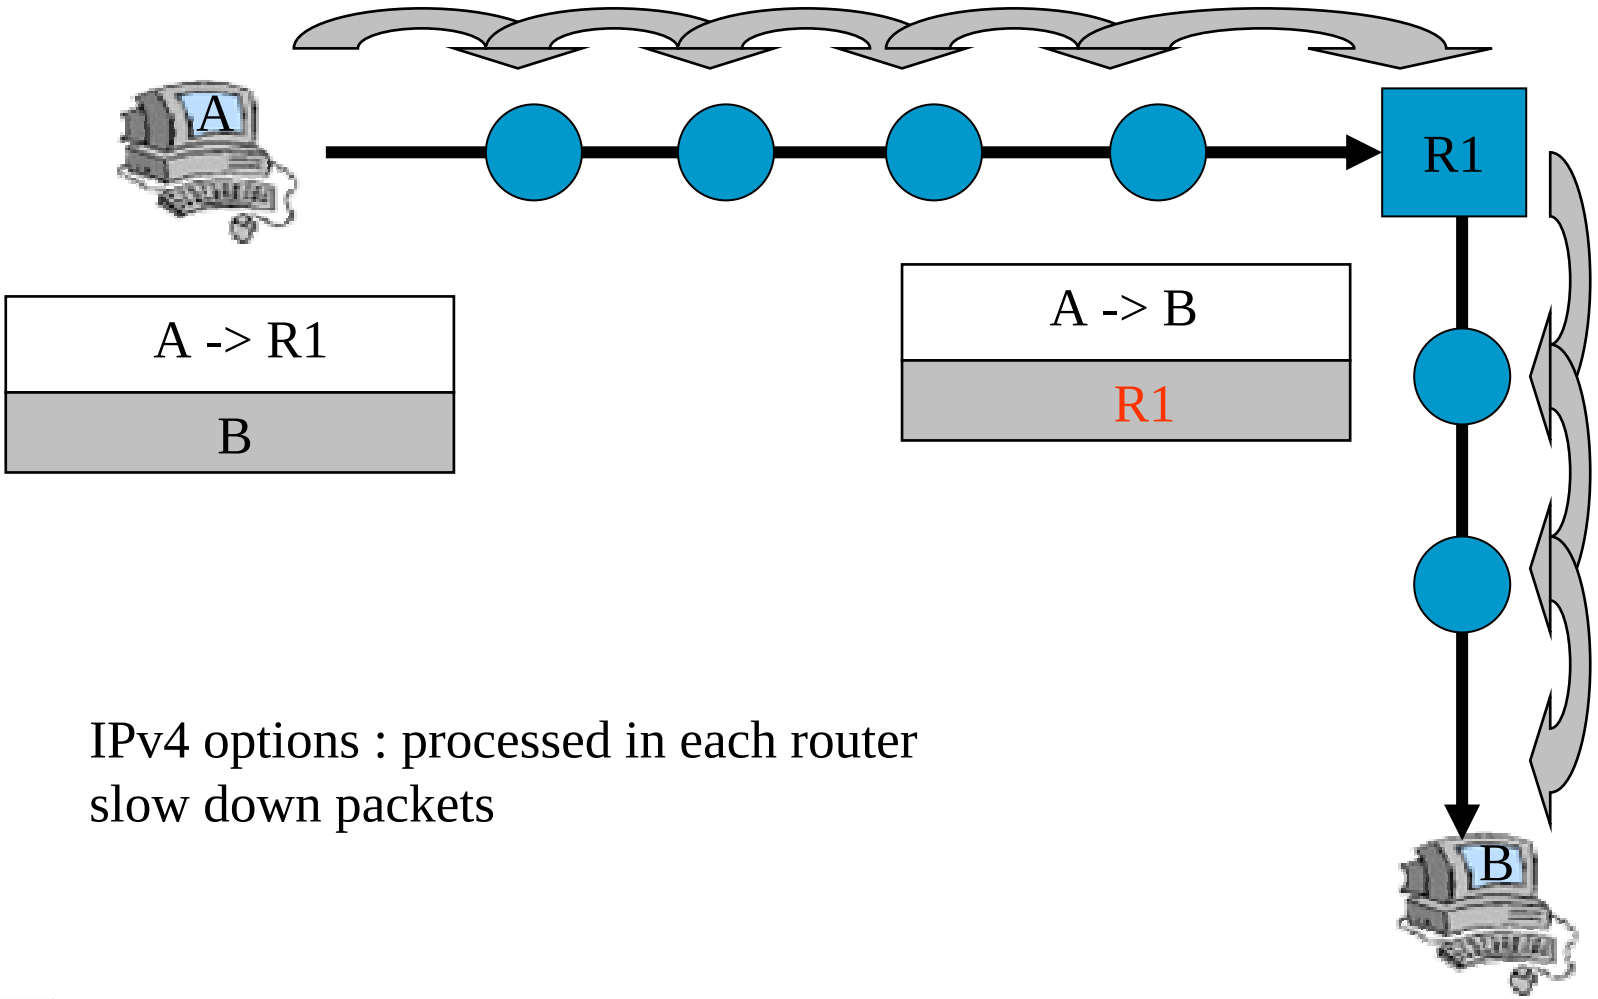

IPv4 options : processed in each router
slow down packets

v4 options vs. v6 extensions

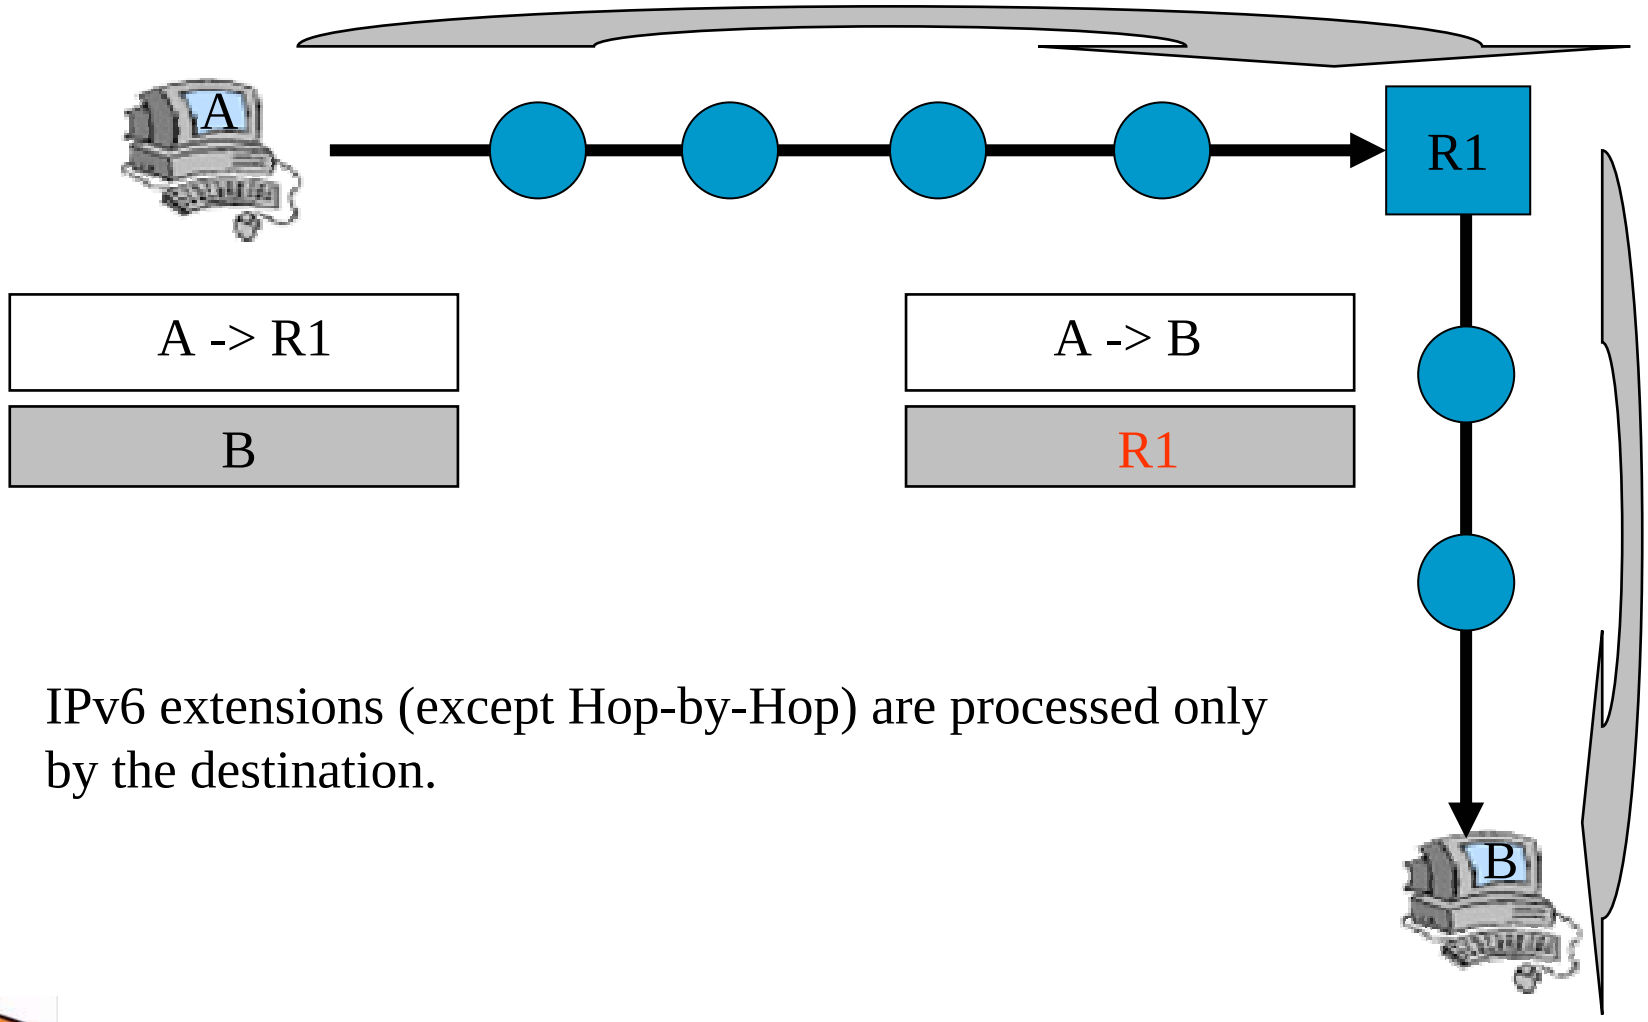

IPv6 extensions (except Hop-by-Hop) are processed only by the destination.

Order is important

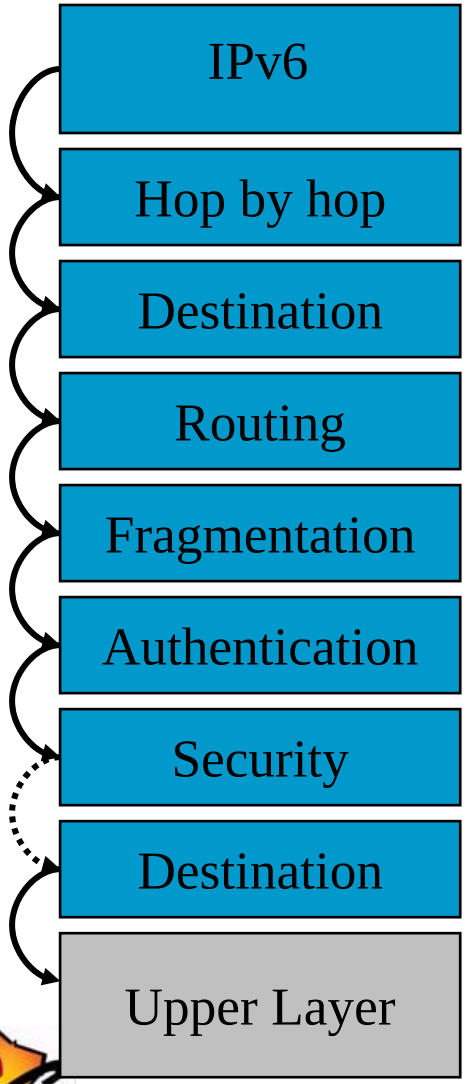

- Hop by hop: Processed by every router
- Destination: Processed by routers listed in Routing extension
- Routing: List of routers to cross
- Fragmentation: Processed by the destination
- Authentication: After reassembling the packet
- Security: Cipher the content of the remaining information
- Destination: Processed only by the destination
- Upper Layer

IPv6 Addressing

Addressing scheme

- RFC 3513 (obsoletes RFC 2373)
- RFC 3587 (obsoletes RFC 2374)
- 128 bit long addresses
  - Allow hierarchy
  - Flexibility for network evolutions
- Use CIDR principles:
  - Prefix / prefix length
    - 2001:660:3003::/48
    - 2001:660:3003:2:a00:20ff:fe18:964c/64
  - Aggregation reduces routing table size
- Hexadecimal representation
- Interfaces have several IPv6 addresses

Textual Address Format

- Base format (a 16-byte **Global IPv6 Address) :**
  - **2001:0660:3003:0001:0000:0000:6543:210F**
- Compact Format:

2001:660:3003:1::6543:210F

  - In order to avoid ambiguity, "::" can occur only once

IPv6 Addresses

- Loopback ::1
- Link local FE80::/10
- Site local FEC0::/10
- Unique local (ULA's) FC00::/7
- Global
  - 6bone: 3FFE::/16
  - Official: 2001::/16
  - IPv4 mapped
  - 6to4: 2002::/16

- Unicast
- Multicast FF00::/8
- Anycast

specific to IPv4/IPv6 integration

Local Addresses

Link-local

10 bits	*54 bits*	*64 bits*
1111111010	**0 .....0**	**Interface ID**

FE80

Site-local (in the process of being deprecated)

10 bits	*54 bits*	*64 bits*
1111111011	Subnet ID	Interface ID

FEC0

Local Addresses

Unique local Addresses (ULA's)

draft-ietf-ipv6-unique-local-addr-09

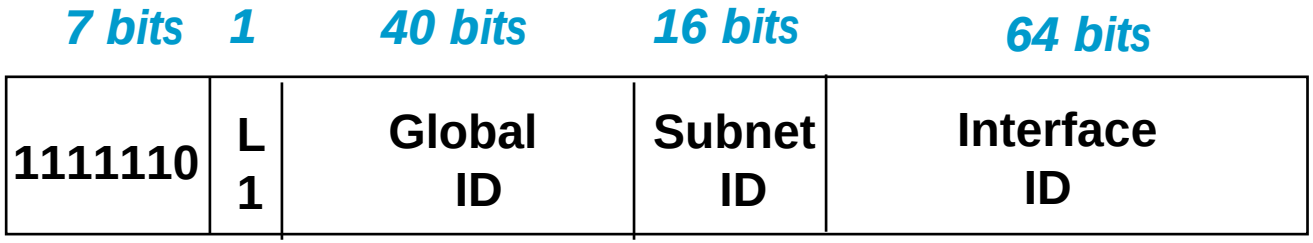

- Sites compute the global-ID (random, pseudo-random…)
- /48 for the site, can be used for internal communication
- Could be exchanged with some neighbors, but not in the entire IPv6 interdomain

Interface Identifier

- 64 bits to be compatible with IEEE 1394 (FireWire)
- Eases auto-configuration
- IEEE defines the mechanism to create an EUI-64 from IEEE 802 MAC addresses (Ethernet, FDDI)

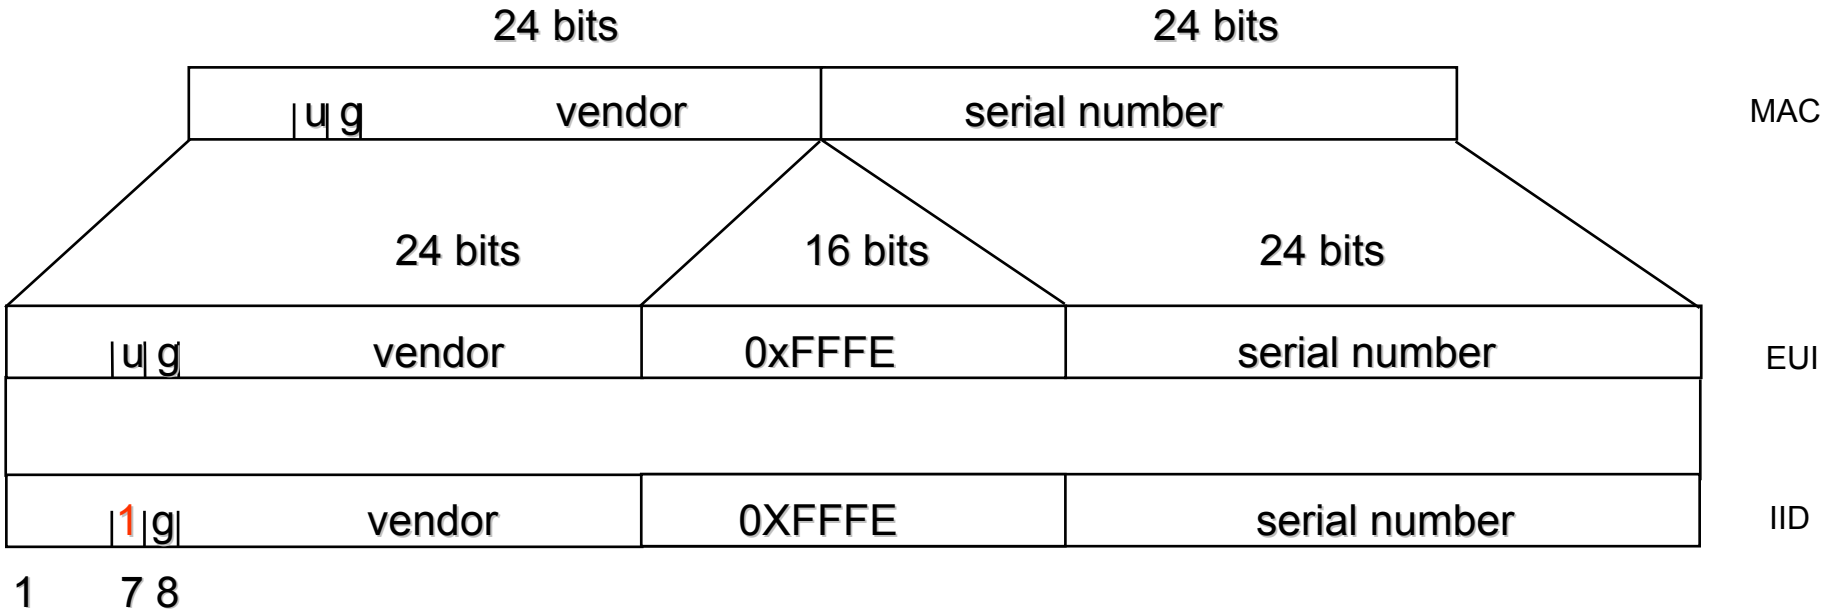

Interface Identifier (2)

- Links with non global identifier (e.g., the Localtalk 8 bit node identifier) → fill first left bits with 0
- For links without identifiers, there are different ways to proceed (e.g., tunnels, PPP):
  - Choose the identifier of another interface
  - Random number
  - Manual configuration
- **THEN :** Invert IEEE EUI-64 “u” bit to become an “interface identifier”
  - so that admins can more easily form local IID (with a zero)

Interface Identifier (3)
(Privacy issues)

- IEEE 24 bit OUI can be used to identify HW:
  - http://standards.ieee.org/regauth/oui/oui.txt
- Interface Identifier can be used to trace a user:
  - The prefix changes, but the interface ID remains the same,
  - Psychological issue.
- Possibility to change Interface ID (RFC 3041 PS):
  - If local storage, use MD5 algorithm
  - Otherwise draw a random number

Multicast Addresses

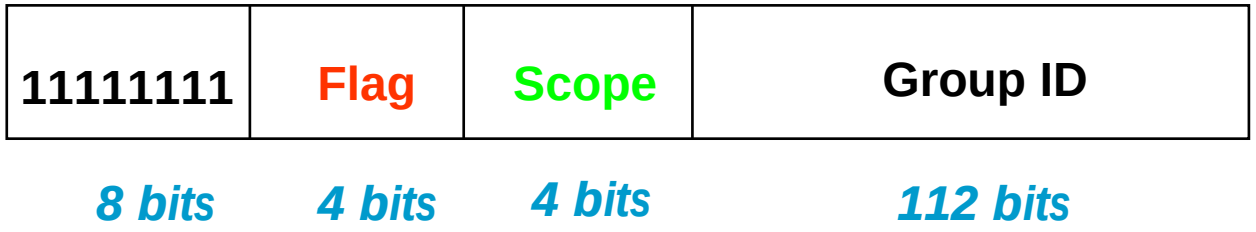

FF00::/8 addresses are multicast addresses

Flag bits: 0 R P T

T = 0 *permanent addresses (managed by IANA)*

T = 1 *transient multicast addresses*

- P = 1 *derived from unicast prefix (RFC3306)*
  - R = 1 *embedded RP addresses (I-D)*

Scope

0 : Reserved
1 : Interface-local
2 : Link-local
3 : Subnet-local
4 : Admin-local
5 : Site-local
8 : Organization-local
E : Global
F : Reserved

Anycast Addresses (RFC 3513)

- **It cannot be distinguished from a Unicast address**

- Subnet anycast router address is :
  - `subnet_ID::0/subnet_prefix_length`

- Can be used for distributed resources (DNS root servers, sites exit routers...)

RFC 3587: Aggregatable Global Unicast (obsoletes RFC 2374)

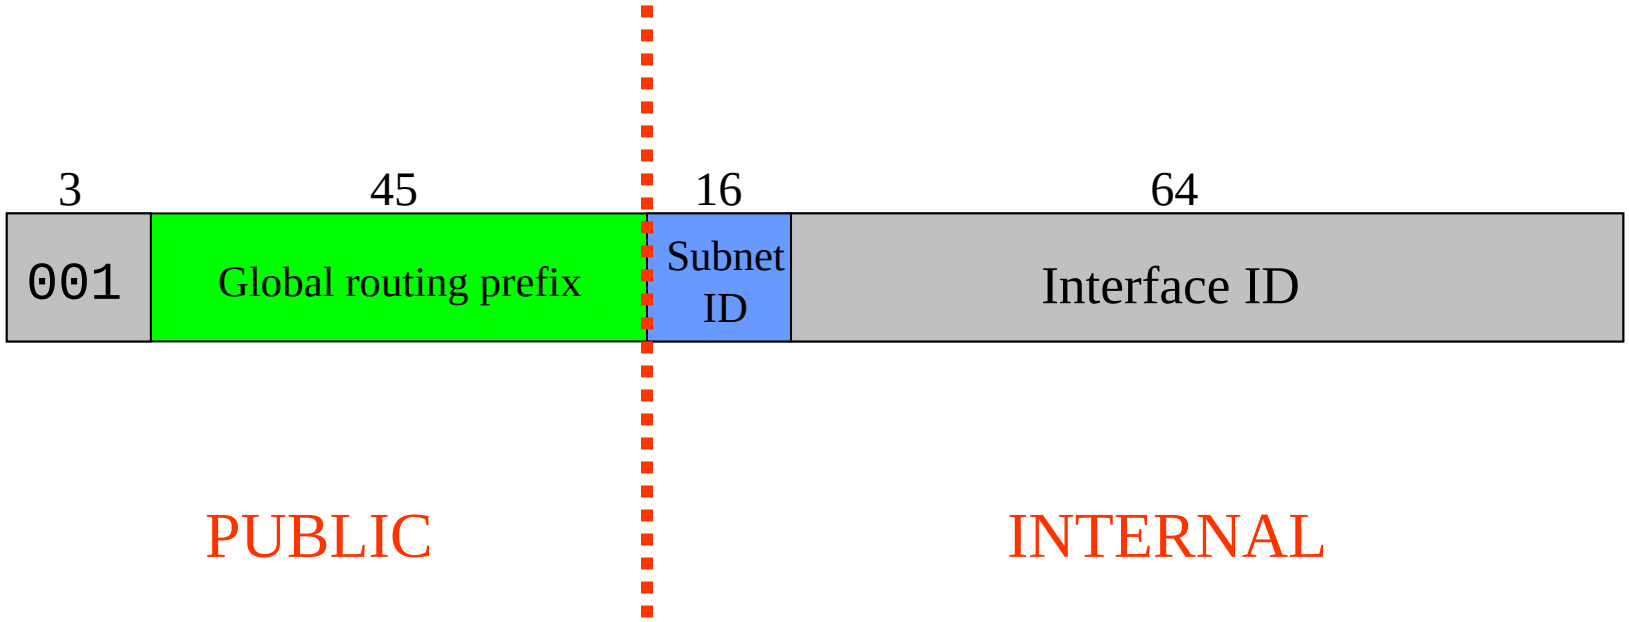

Production Addressing Scheme

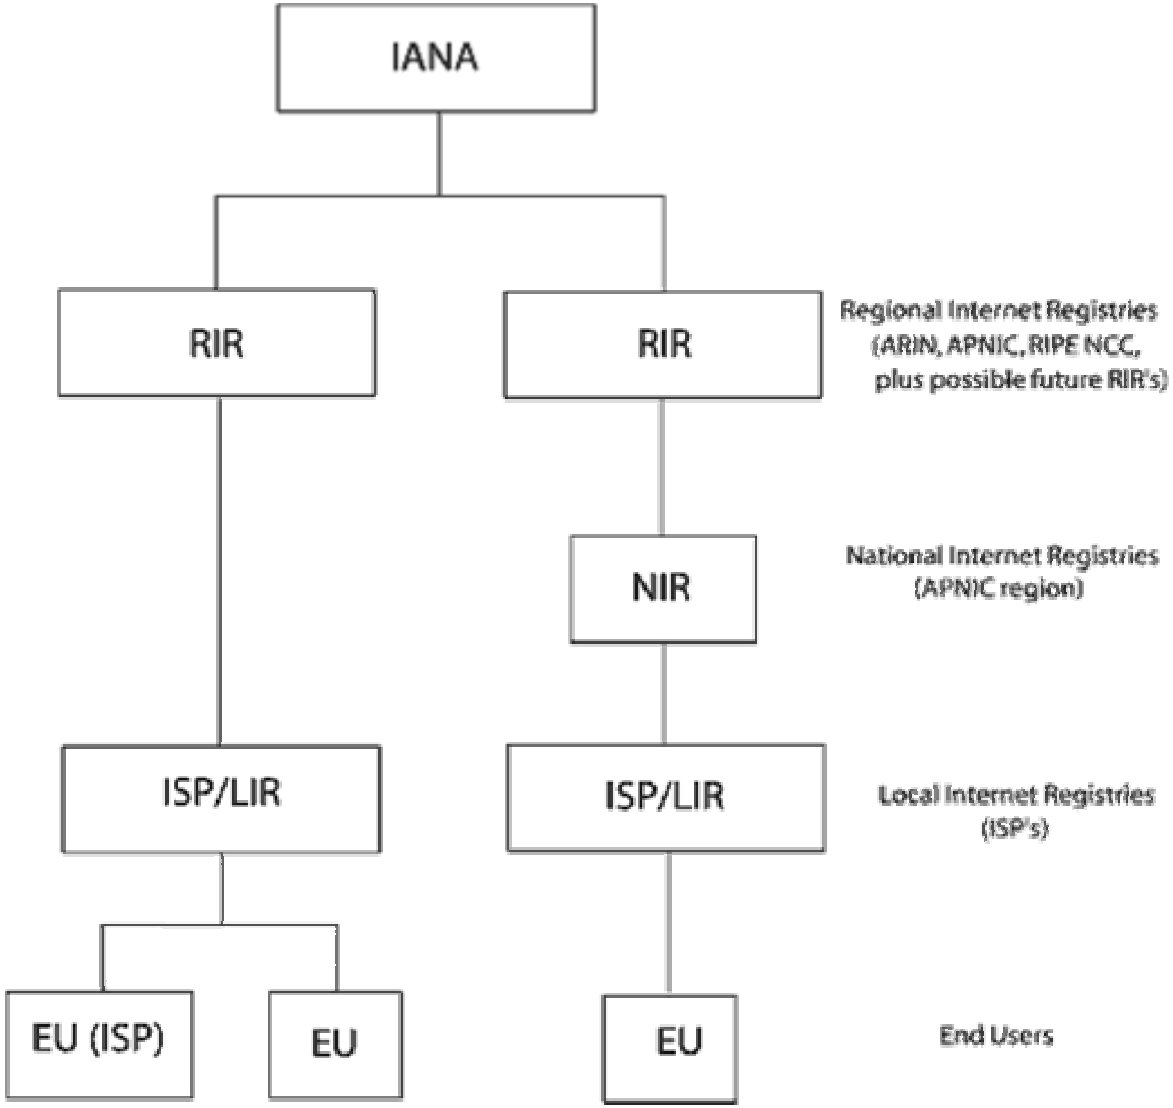

Initial RIR allocation Policy & Procedure

- Get the RIPE documents [246-250, 256, 261, 267, 274, 275, 280-282]
  - http://www.ripe.net/ripe/docs/ipv6.html
- Criteria: RIPE-267
  - http://www.ripe.net/ripe/docs/ipv6plicy.html
- To qualify for an initial allocation of IPv6 address space, an organization must:
  - be a LIR : *not be an end site*
  - plan to provide IPv6 connectivity to organizations to which it will assign /48s, by advertising that connectivity through its single aggregated address allocation (/32 prefix)

  and
  - have a plan for making at least 200 x /48 assignments to other organizations within two years.

IPv6

Size DOES matter…