IPv6 Protocol (headers & options)

6DEPLOY. IPv6 Deployment and Support

Copy …Rights

PDF files are available from www.6diss.org

Looking for a contact ?

- ***Mail to : martin.potts@martel-consulting.ch***
- ***Or bernard.tuy@renater.fr***

Droits d'auteur …

L'ensemble des présentations utilisées dans le cadre de cet atelier est la propriété de 6DISS, représenté par ses différents partenaires.

La version Powerpoint des présentations peut être réutilisée et modifiée après qu'une autorisation écrite ait été obtenue

L'usage de tout ou partie de ce matériel doit mentionné que sa source est le projet 6DISS

La version PDF des présentations est disponible sur www.6diss.org

Pour tout contact :

- ***Mail à Martin.Potts@martel-consulting.ch***
- ***Ou Bernard.Tuy@renater.fr***

Agenda

IPv6 Header

- Comparison with IPv4

IPv6 optional headers (extensions)

Processing IPv6 headers

- Comparison with IPv4

IPv6 Header

The IPv6 header is designed

- To minimize header overhead
- and reduce the header process for most of the packets
- Less important information and option fields are moved to ***extension headers***

⇒ **IPv6 & IPv4 headers are not interoperable**

IPv4 Header

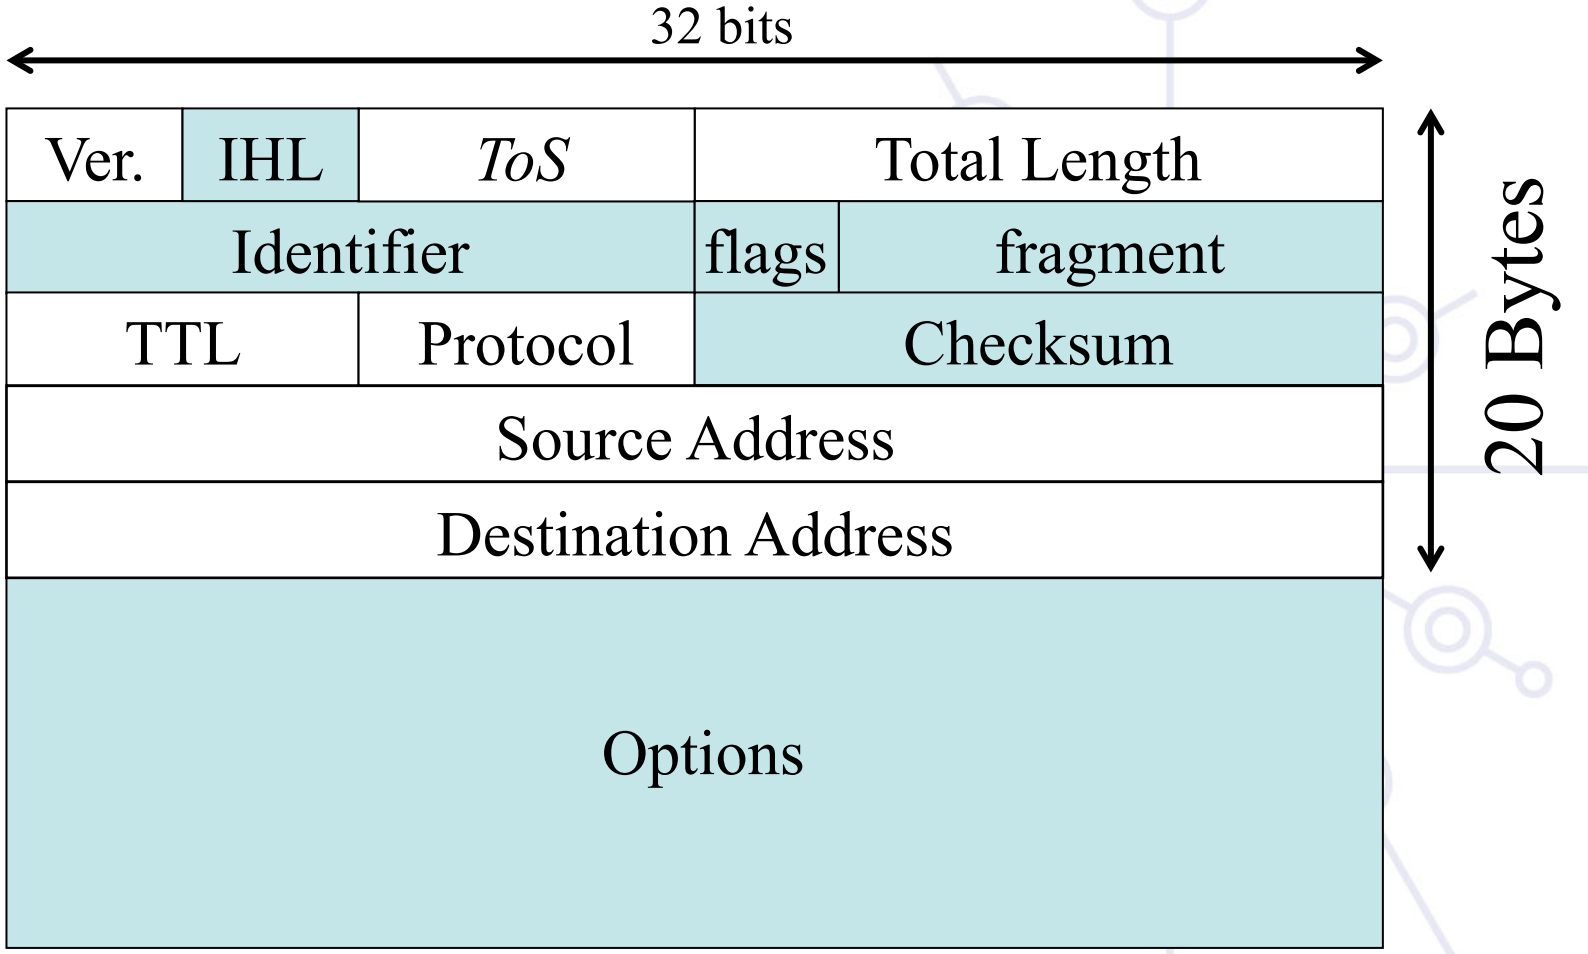
32 bits
Ver.
IHL
ToS
Total Length
Identifier
flags
fragment
TTL
Protocol
Checksum
Source Address
Destination Address
Options
20 Bytes

IPv6 Header simplification

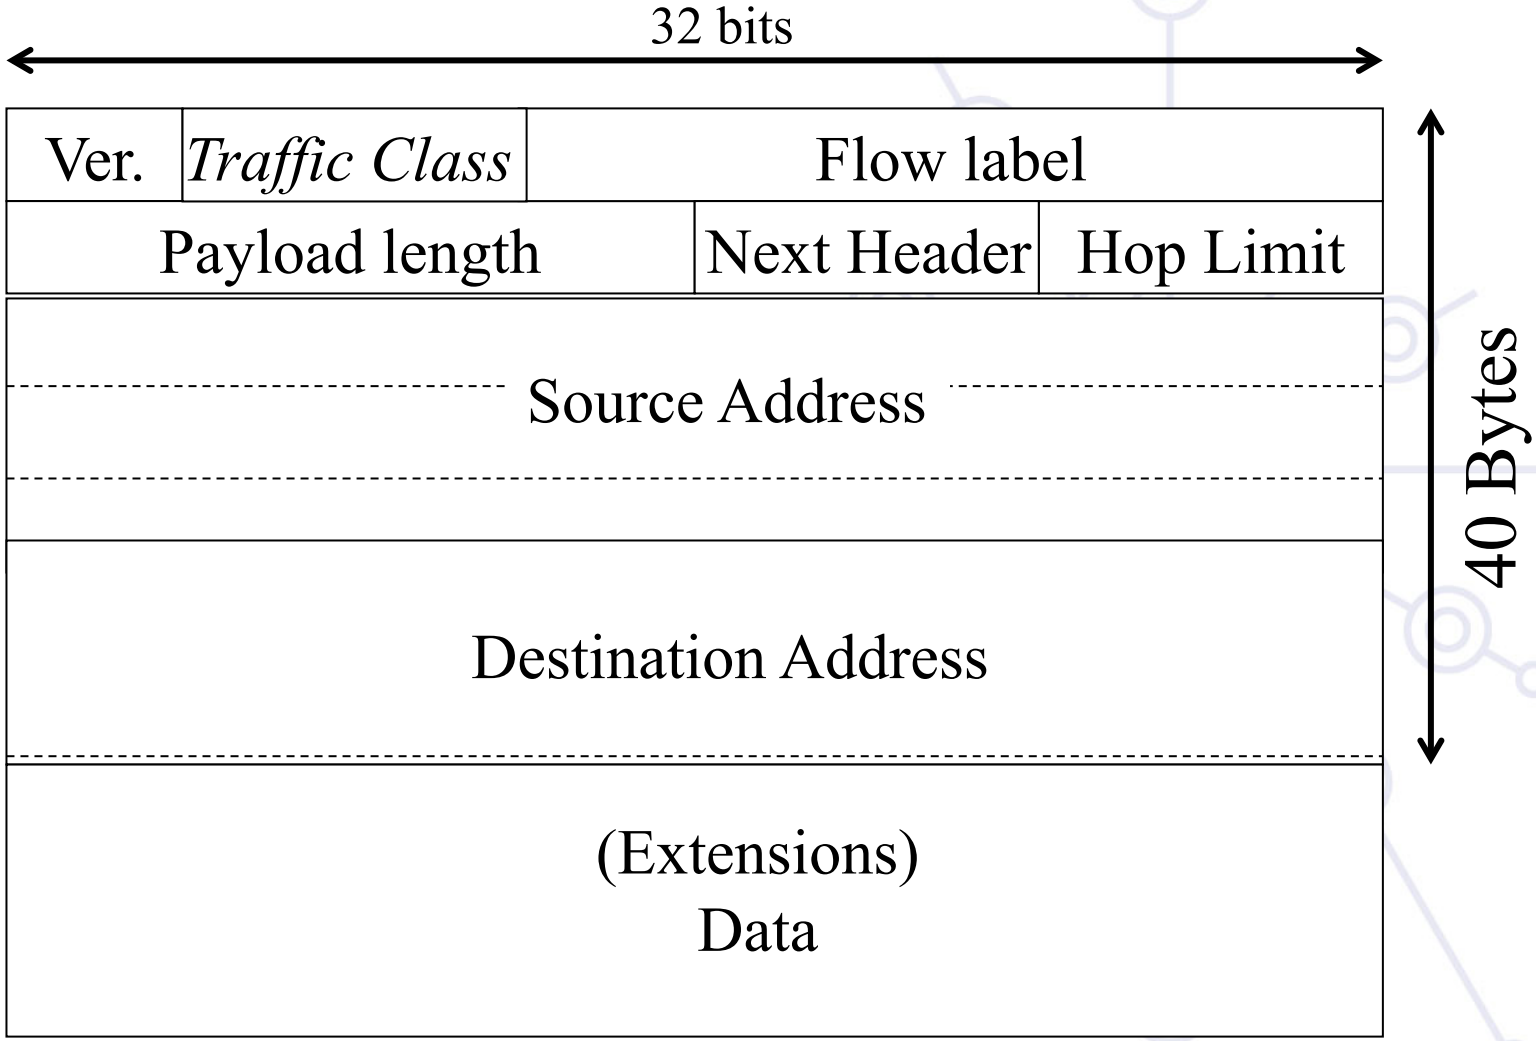

IPv6 header fields

Version

- 4 bits

Traffic class (*see next slide)*

- 8 bits

Flow label *(see next slide)*

- 20 bits

Payload length

- Use Jumbogram for specific cases (payload = 0)
- 16 bits

Hop limit

- 8 bits

Next header

- 8 bits

CoS support in IPv6

The Traffic Class field: *used as in IPv4*

- Work done in diffserv wg (closed): RFCs **2474**, 2475, 2597, 3260, …

CU is Currently Unused (reserved)

The Flow Label field: designed to enable classification of packets belonging to a specific flow

- **A flow** is a sequence of packets that should receive specific non-default handling from the network
- Intuitively: 5-tuple of the same source/destination address/port and transport protocol values
- Without the flow label the classifier must use transport next header value and port numbers
  - Less efficient (need to parse the option headers)
  - May be impossible (fragmentation or IPsec ESP)
- Further info:
  - RFC 3697 (PS)

IPv6: optional Extensions

New "mechanism" replacing IPv4 options

An IPv6 extension :

- Every extension has its own message format
- Is a n x 8-byte datagram
- Starts with a 1-byte 'Next Header' field
  - Pointing to either another extension or a L-4 protocol

Hop-by-hop (jumbogram, router alert)

- Always the first extension
- Analyzed by every router.

Destination

Routing (*loose source routing*)

Fragmentation

Security

- Authentication (AH)
- Encapsulating Security Payload (ESP) : confidentiality

Order is important (RFC 2460)

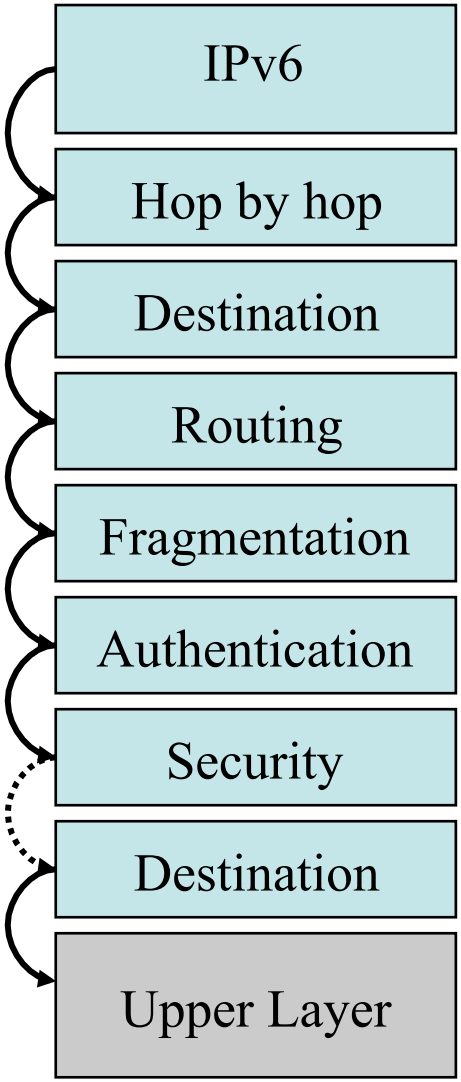

- Hop by hop: Processed by every router
- Destination: Processed by routers listed in *Routing extension*
- Routing: List of routers to cross
- Fragmentation: Processed by the destination
- Authentication: After reassembling the packet
- Security: Cipher the content of the remaining information
- Destination: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

IPv6: Optional headers

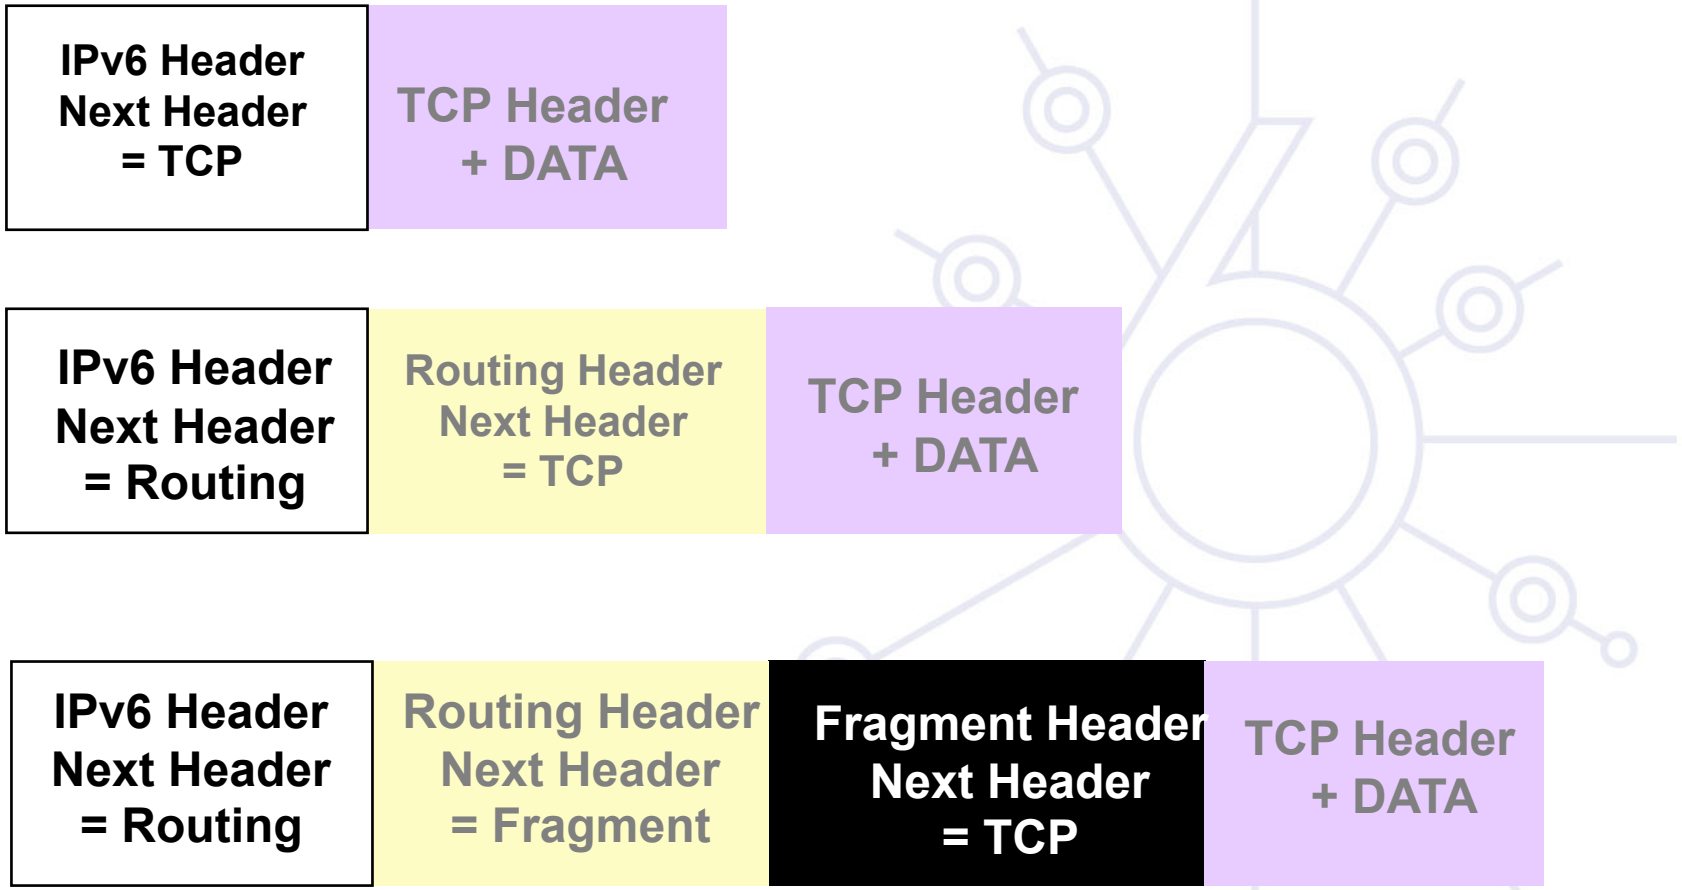

IPv4 header options processing

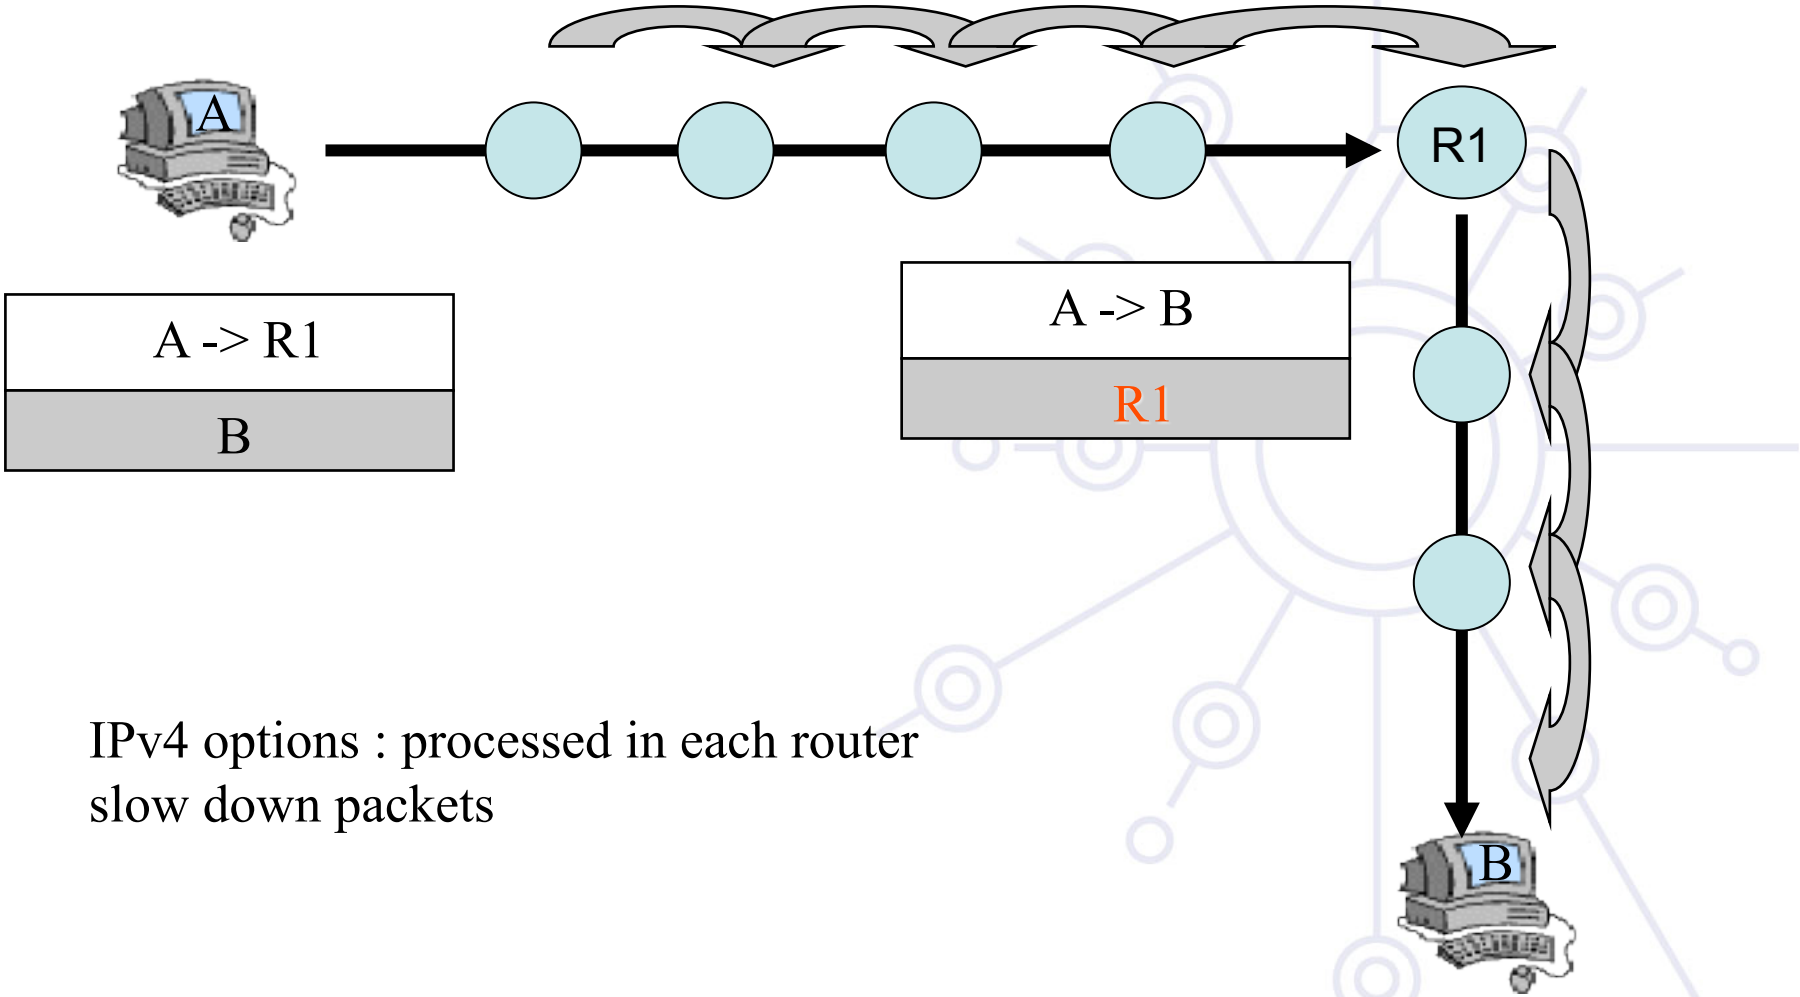

IPv4 options : processed in each router
slow down packets

IPv6 ext. header processing

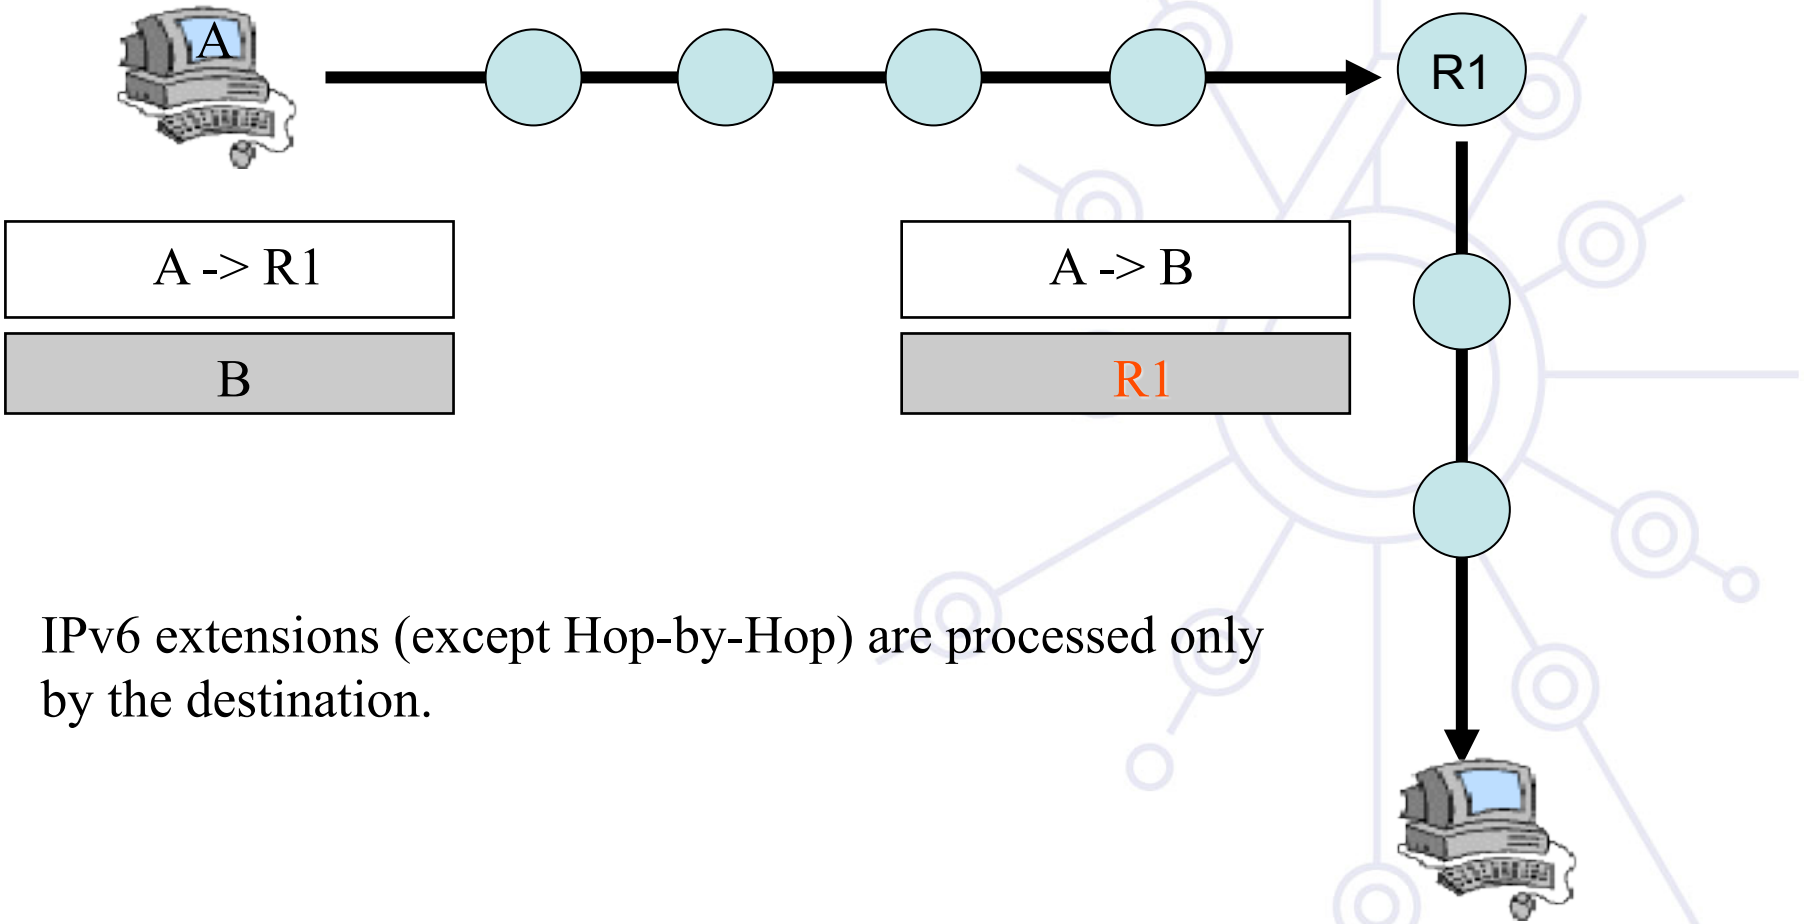

IPv6 extensions (except Hop-by-Hop) are processed only by the destination.

Conclusion

Main changes in IPv6 protocol are within address format and datagram headers

- **A lot of fields in the IPv6 header have disappeared**

  ⇒ More efficient processing in the (intermediate) routers

- **Optional extensions allow more functionalities (source routing, authentication, …)**
- **Optional header mechanism allows new options introduction without modifying the protocol**

IPv6 Addressing

6DEPLOY. IPv6 Deployment and Support

Copy …Rights

Looking for a contact ?

Mail to : martin.potts@martel-consulting.ch

Or bernard.tuy@renater.fr

IPv6 Addressing Scheme

RFC4291 defines IPv6 addressing scheme

RFC3587 defines IPv6 global unicast address format

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

IPv6 Address Types

Unicast (one-to-one)

- global
- link-local
- site-local (deprecated)
- Unique Local (ULA)
- IPv4-compatible (deprecated)
- IPv6-mapped

Multicast (one-to-many)

Anycast (one-to-nearest)

Reserved

Textual Address Format

Preferred Form (a 16-byte Global IPv6 Address):

2001:0DB8:3003:0001:0000:0000:6543:210F

Compact Format:

2001:DB8:3003:1::6543:210F

IPv4-mapped: ::FFFF:134.1.68.3

Literal representation

- [2001:DB8:3003:2:a00:20ff:fe18:964c]
- http://[2001:DB8::43]:80/index.html

IPv6 Address Type Prefixes

Address Type	Binary Prefix	IPv6 Notation
Unspecified	00…0 (128 bits)	::/128
Loopback	00…1 (128 bits)	::1/128
Multicast	1111 1111	FF00::/8
Link-Local Unicast	1111 1110 10	FE80::/10
ULA	1111 110	FC00::/7
Global Unicast	(everything else)	
IPv4-mapped	00…0:1111 1111:IPv4	::FFFF:IPv4/128
Site-Local Unicast (deprecated)	1111 1110 11	FEC0::/10
IPv4-compatible (deprecated)	00…0 (96 bits)	::IPv4/128

Global Unicast assigments actually use 2000::/3 (001 prefix)

Anycast addresses allocated from unicast prefixes

IPv6 Address Space

Aggregatable Global Unicast Addresses (001): 1/8
Unique Local Unicast addresses (1111 1110 00): 1/128
Link-Local Unicast Addresses (1111 1110 10): 1/1024
Multicast Addresses (1111 1111): 1/256

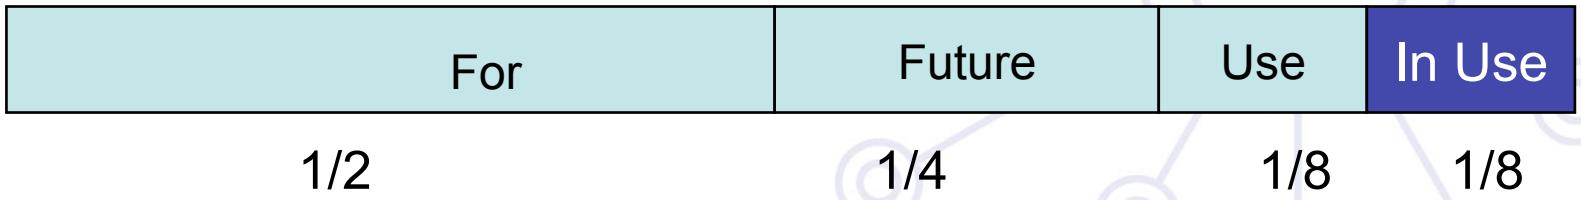

More info:
http://www.iana.org/assignments/ipv6-address-space

Some Special-Purpose Unicast Addresses

Listed in RFC5156

The **unspecified address,** used as a placeholder when no address is available:

0:0:0:0:0:0:0:0 (::/128)

The **loopback address,** for sending packets to itself:

0:0:0:0:0:0:0:1 (::1/128)

The **documentation prefix [RFC3849]:**

2001:db8::/32

Link-Local & Site-Local Unicast Addresses

Link-local addresses for use during auto-configuration and when no routers are present **(FE80::/10):**

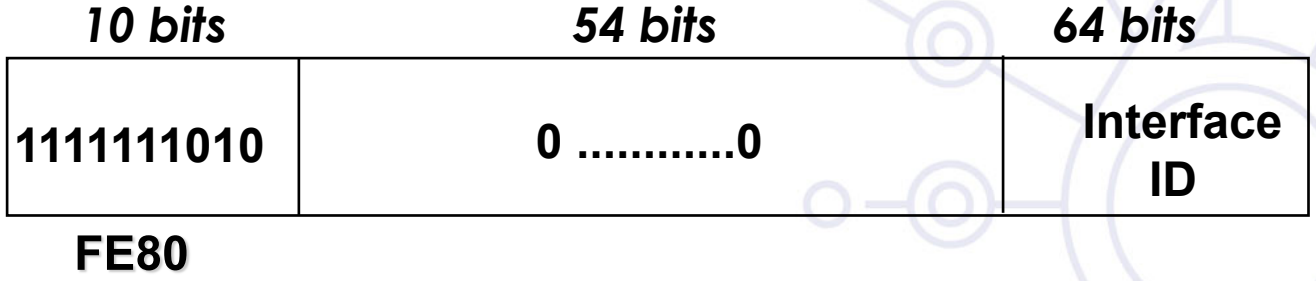

Site-local addresses for independence from changes of TLA / NLA* **(FEC0::/10):** (deprecated by RFC3879)

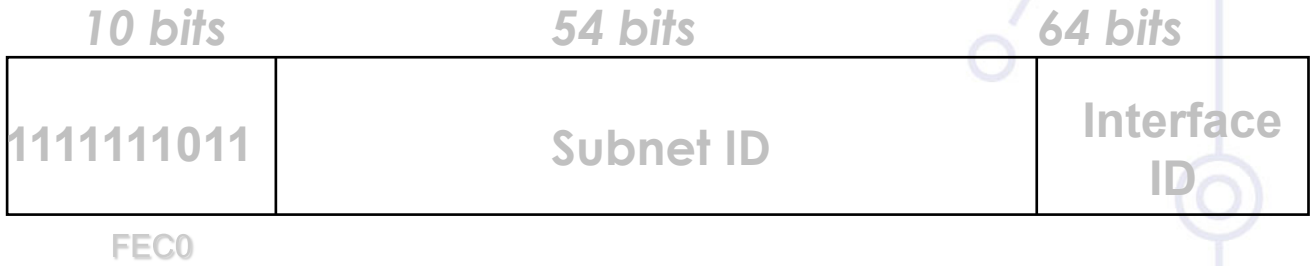

Interface IDs

The lowest-order 64-bit field of unicast addresses may be assigned in several different ways:

- auto-configured from a 64-bit MAC address
- auto-configured from a 48-bit MAC address (e.g., Ethernet) expanded into a 64-bit EUI-64 format
- assigned via DHCP
- manually configured
- auto-generated pseudo-random number (to counter some privacy concerns)
- CGA (Cryptographically Generated Address)
- possibly other methods in the future

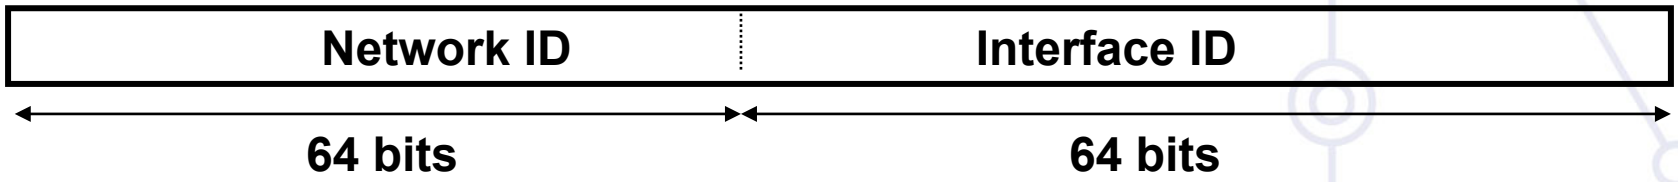

Autoconfigured Interface IDs (1)

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)

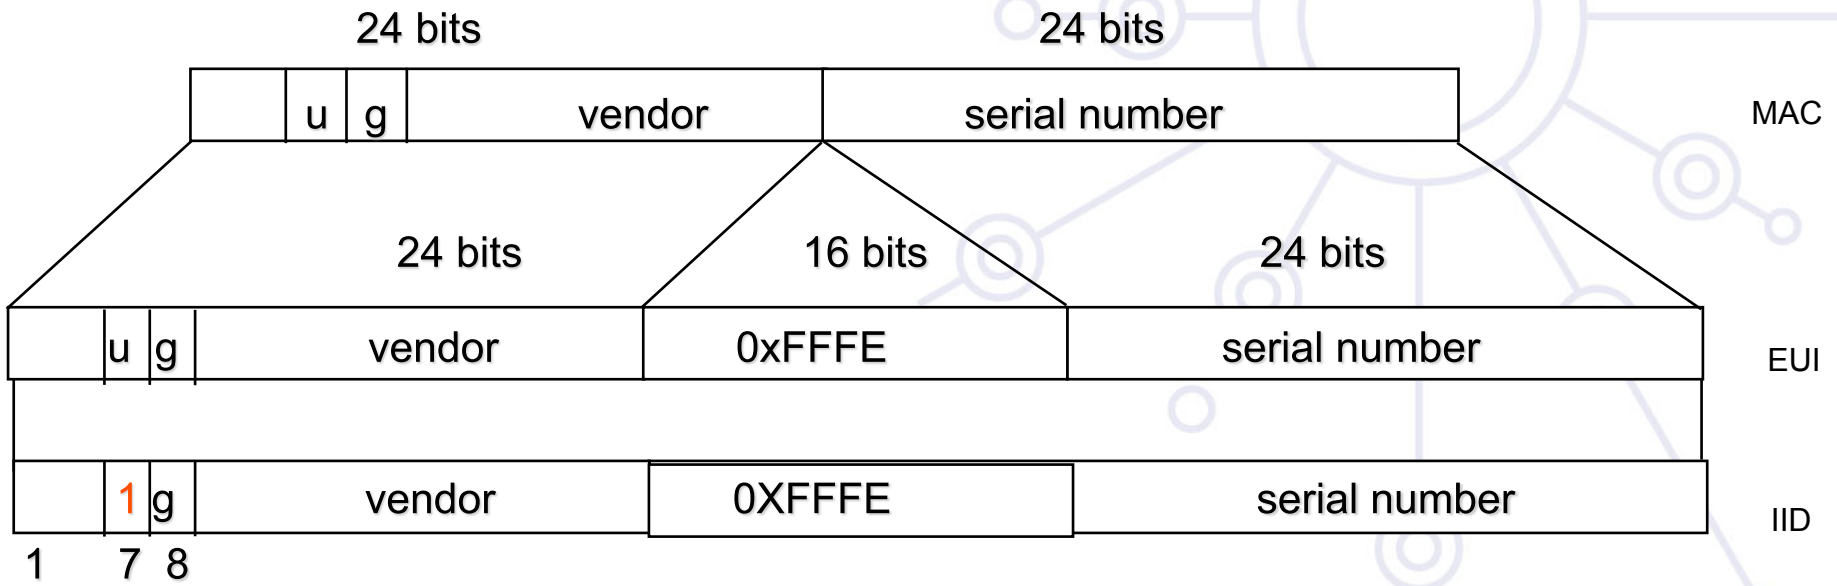

Autoconfigured Interface IDs (2)

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) to have a subnet-prefix-unique identifier:

- Choose the universal identifier of another interface
- Manual configuration
- Node Serial Number
- Other Node-Specific Token

Multicast Addresses

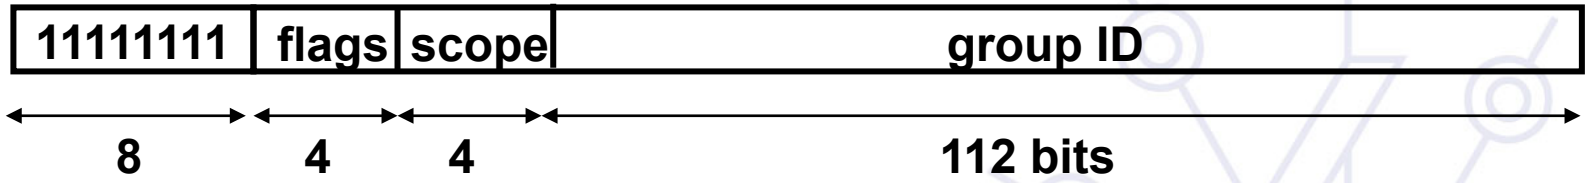

Flags: 0RPT: The high-order flag is reserved, and must be initialized to 0.

- **T**: Transient, or not, assignment
- **P**: Assigned, or not, based on network prefix
- **R**: Rendezvous Point Address embedded, or not

Scope field:
1 - Interface-Local
2 - link-local
4 - admin-local
5 - site-local
8 - organization-local
E - global

(3,F reserved)(6,7,9,A,B,C,D unassigned)

Unique Local IPv6 Unicast Addresses (1)

ULAs are defined in **RFC4193:**

- Globally unique prefix with high probability of uniqueness
- Intended for local communications, usually inside a site
- They are not expected to be routable on the Global Internet
- They are routable inside of a more limited area such as a site
- They may also be routed between a limited set of sites
- Locally-Assigned Local addresses vs. Centrally-Assigned Local addresses

Unique Local IPv6 Unicast Addresses (2)

ULA characteristics:

- Well-known prefix to allow for easy filtering at site boundaries
- ISP independent and can be used for communications inside of a site without having any permanent or intermittent Internet connectivity
- If accidentally leaked outside of a site via routing or DNS, there is no conflict with any other addresses
- In practice, applications may treat these addresses like global scoped addresses

Unique Local IPv6 Unicast Addresses (3)

Format:

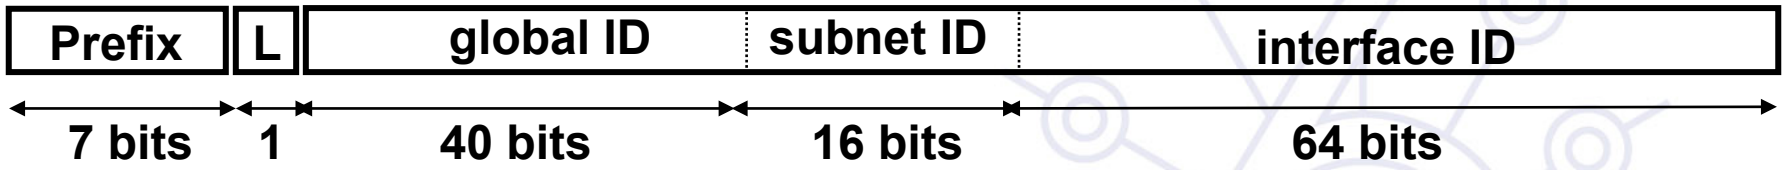

FC00::/7 Prefix identifies the Local IPv6 unicast addresses

L = 1 if the prefix is **locally assigned**

L = 0 may be defined in the future (in practice used for **centrally assigned** prefixes)

ULA are created using a pseudo-randomly allocated global ID

- This ensures that there is not any relationship between allocations and clarifies that these prefixes are not intended to be routed globally

Global Unicast Addresses

Defined in RFC3587

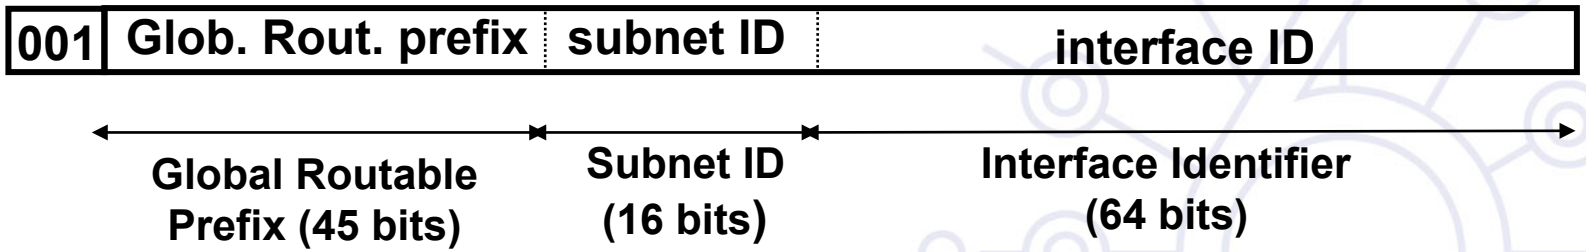

The global routing prefix is a value assigned to a zone (site, a set of subnetworks/links)

- It has been designed as an hierarchical structure from the Global Routing perspective

The subnetwork ID, identifies a subnetwork within a site

- Has been designed to be an hierarchical structure from the site administrator perspective

Anycast Addresses

- Identifier for a set of interfaces (typically in different nodes). A packet sent to an anycast address is delivered to the "nearest" interface (routing protocols' distance)
- Taken from the unicast address space (of any scope).**Not syntactically distinguishable from unicast addresses**
- A unicast address assigned to more than one interface, turning it into an anycast address, the nodes the address is assigned must be explicitly configured to know that it is an anycast address
- Reserved anycast addresses are defined in **RFC2526**
- The Subnet-Router anycast address is predefined (mandatory on all routers):

Production Addressing Scheme (1)

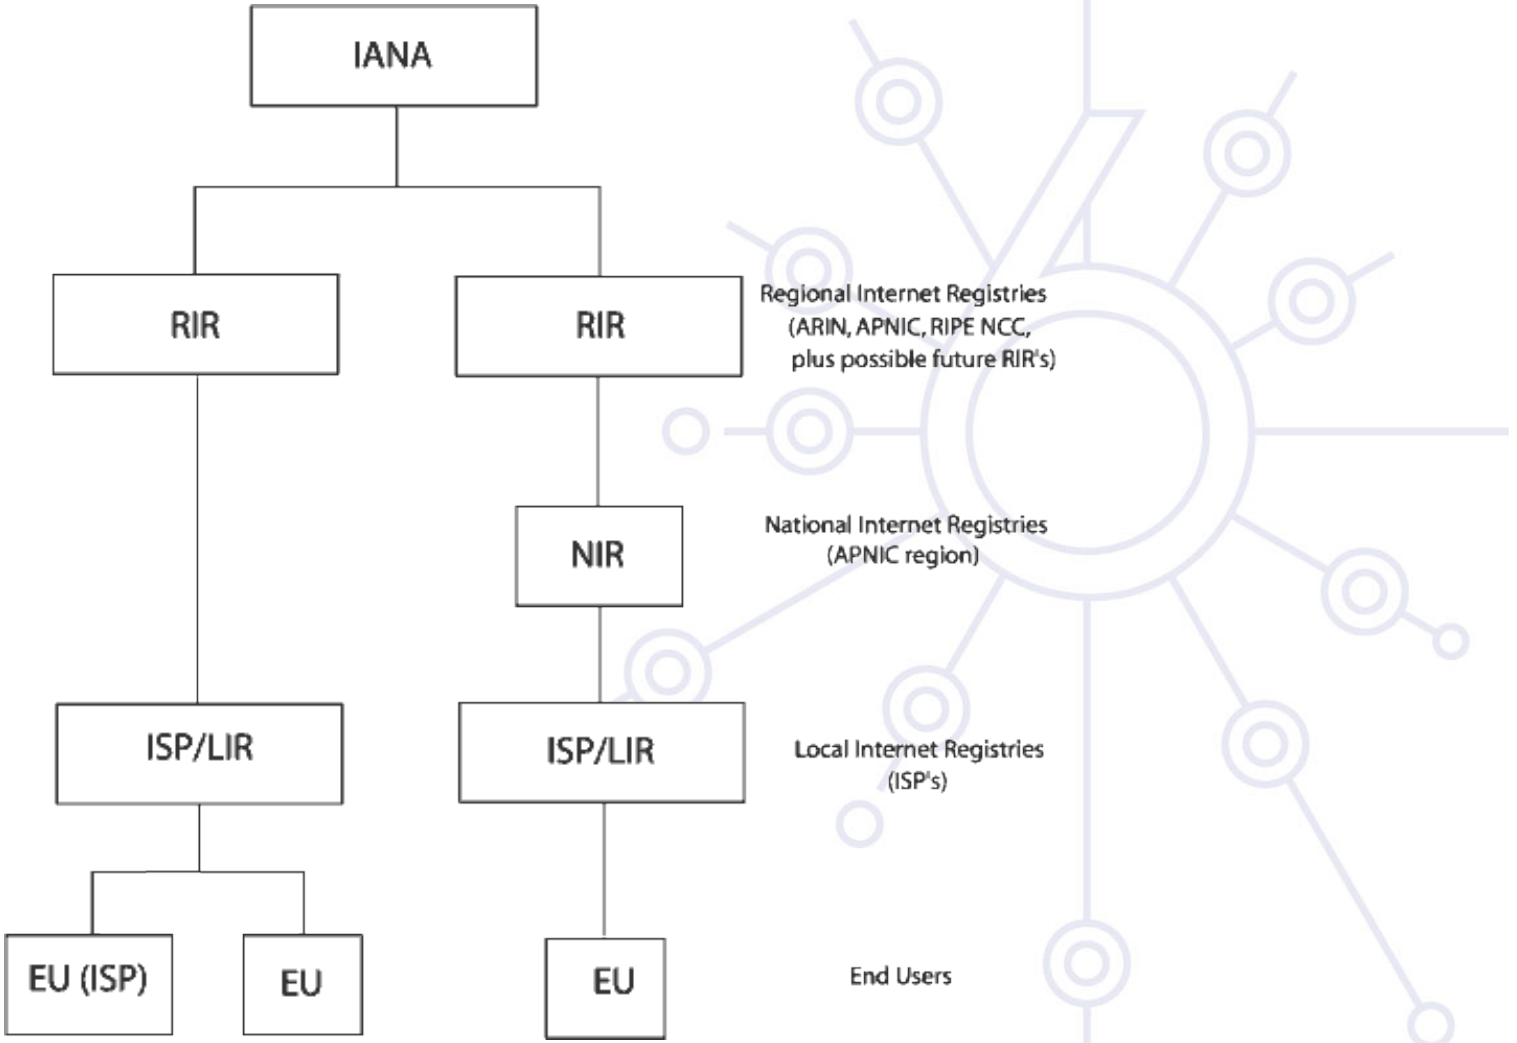

Production Addressing Scheme (2)

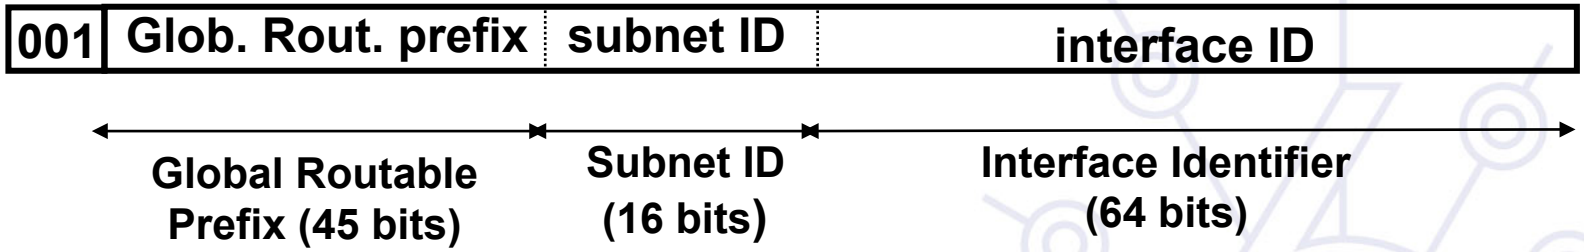

LIRs receive by default /32

- Production addresses today are from prefixes 2001, 2003, 2400, etc.
- Can request for more if justified

/48 used only within the LIR network, with some exceptions for critical infrastructures

/48 to /128 is delegated to end users

- Recommendations following RFC3177 and current policies
- /48 general case, /47 if justified for bigger networks
- /64 if one and only one network is required
- /128 if it is sure that one and only one device is going to be connected

Production Addressing Scheme (3)

Source:
http://www.iana.org/assignments/ipv6-unicast-address-assignments

IPv6 Global Unicast Address Assignments [0]
[last updated 2008-05-13]

Global Unicast Prefix	Assignment	Date	Note
2001:0000::/23	IANA	01 Jul 99	[1]
2001:0200::/23	APNIC	01 Jul 99	
2001:0400::/23	ARIN	01 Jul 99	
2001:0600::/23	RIPE NCC	01 Jul 99	
2001:0800::/23	RIPE NCC	01 May 02	
2001:0A00::/23	RIPE NCC	02 Nov 02	
2001:0C00::/23	APNIC	01 May 02	[2]
2001:0E00::/23	APNIC	01 Jan 03	
2001:1200::/23	LACNIC	01 Nov 02	

Production Addressing Scheme (4)

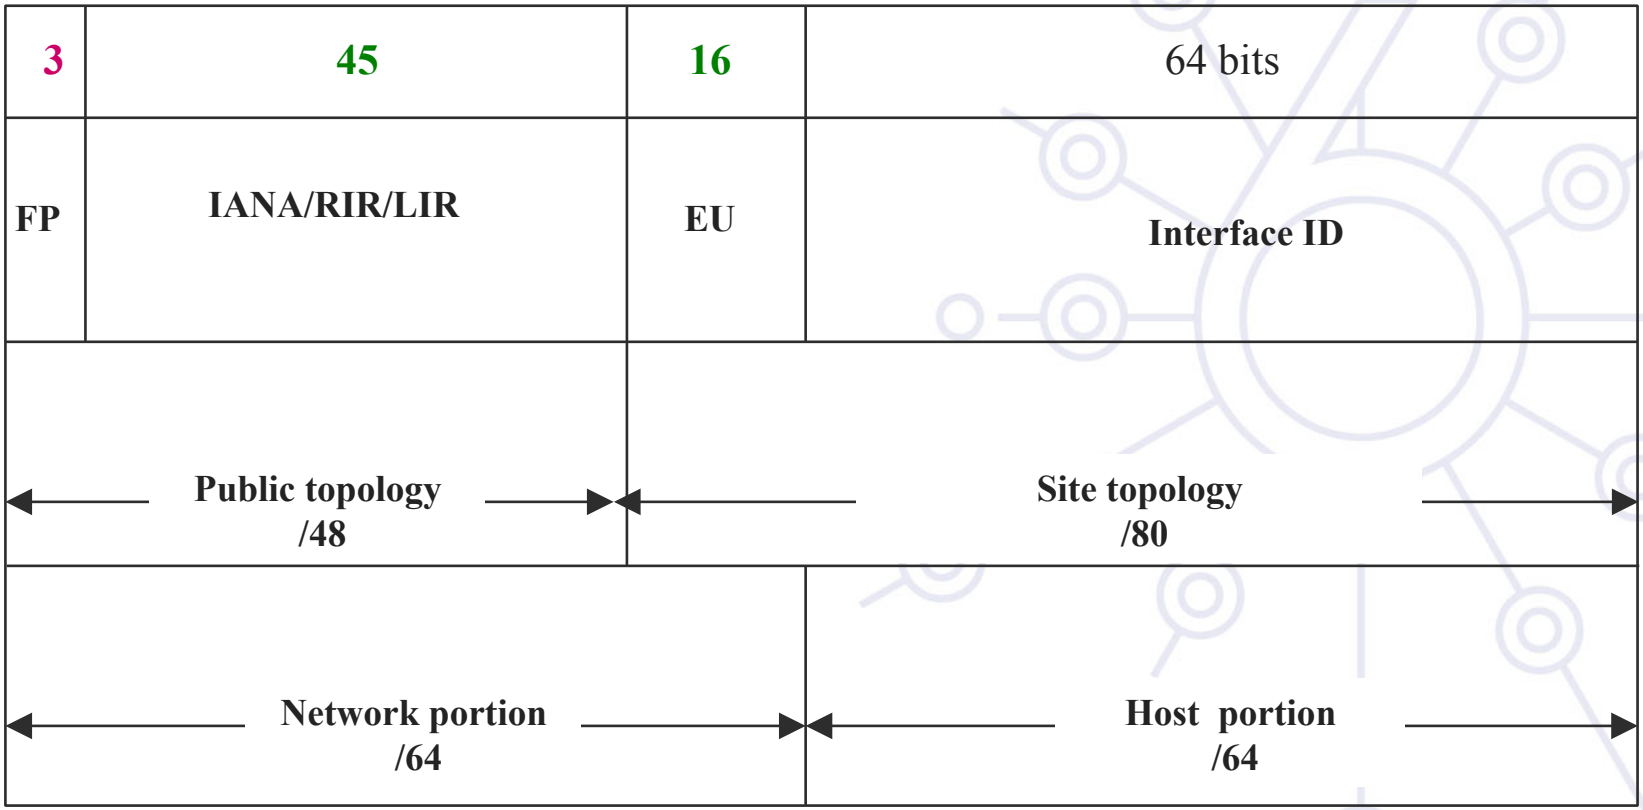

RIR Allocation Policies

AfriNIC:
http://www.afrinic.net/IPv6/index.htm
http://www.afrinic.net/docs/policies/afpol-v6200407-000.htm *
APNIC:
http://www.apnic.org/docs/index.html
http://www.apnic.org/policy/ipv6-address-policy.html *
ARIN:
http://www.arin.net/policy/index.html
http://www.arin.net/policy/nrpm.html#ipv6 *
LACNIC:
http://lacnic.net/sp/politicas/
http://lacnic.net/sp/politicas/ipv6.html *
RIPE-NCC:
http://www.ripe.net/ripe/docs/ipv6.html
http://www.ripe.net/ripe/docs/ipv6policy.html *

- *describes policies for the allocation and assignment of globally unique IPv6 address space

RIR Allocation Statistics

AfriNIC:

- http://www.afrinic.net/statistics/index.htm

APNIC:

- http://www.apnic.org/info/reports/index.html

ARIN:

- http://www.arin.net/statistics/index.html

LACNIC:

- http://lacnic.org/sp/est.html

RIPE-NCC:

- http://www.ripe.net/info/stats/index.html

See http://www.ripe.net/rs/ipv6/stats/

6DEPLOY Project Web Site:

http://www.6deploy.eu