

IPv6 Introduction and Why Prepare ?

Brussels, June 2015

Belnet
dedicated connectivity

.be



IPv6

Where do you want to be today ?

Purpose of this Session

■ Objectives

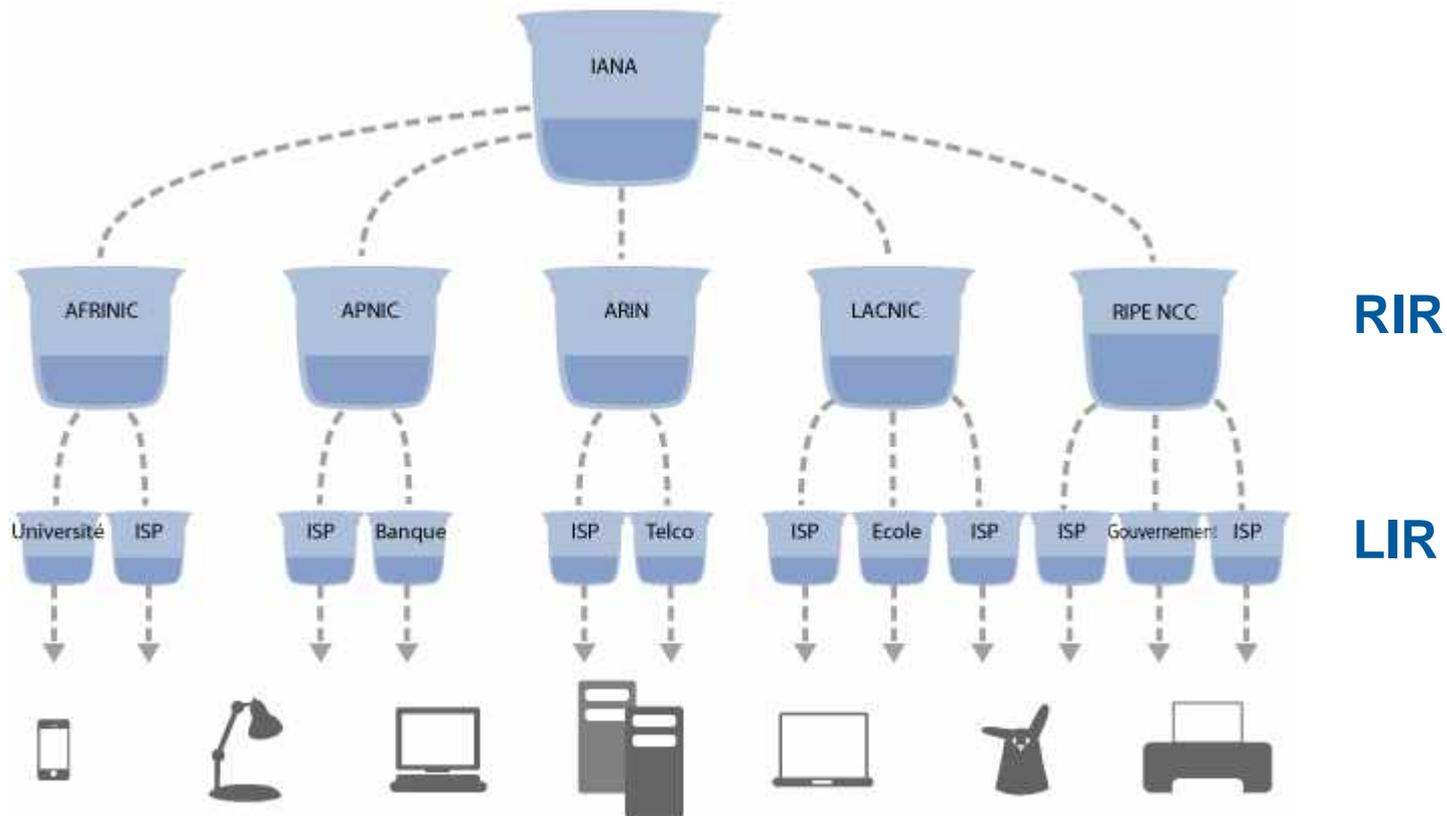
- IPv4 Problem statement
- Explain the solution & purpose of IPv6
- Acquire IPv6 awareness, & what all this space means
- Understand the syntax & semantics of IPv6 addressing
- Understand the different kinds of IPv6 addresses
- State where IPv6 is today
- Explain why you should think of adopting IPv6

The Problem with IPv4

- The Internet keeps growing
 - Mobile, IoT, Gaming, Home appliances,
 - China, India, ...
- Running out of IPv4 addresses
- Business opportunities are endangered
- Running out of time!

Who is the authority for IP ranges ?

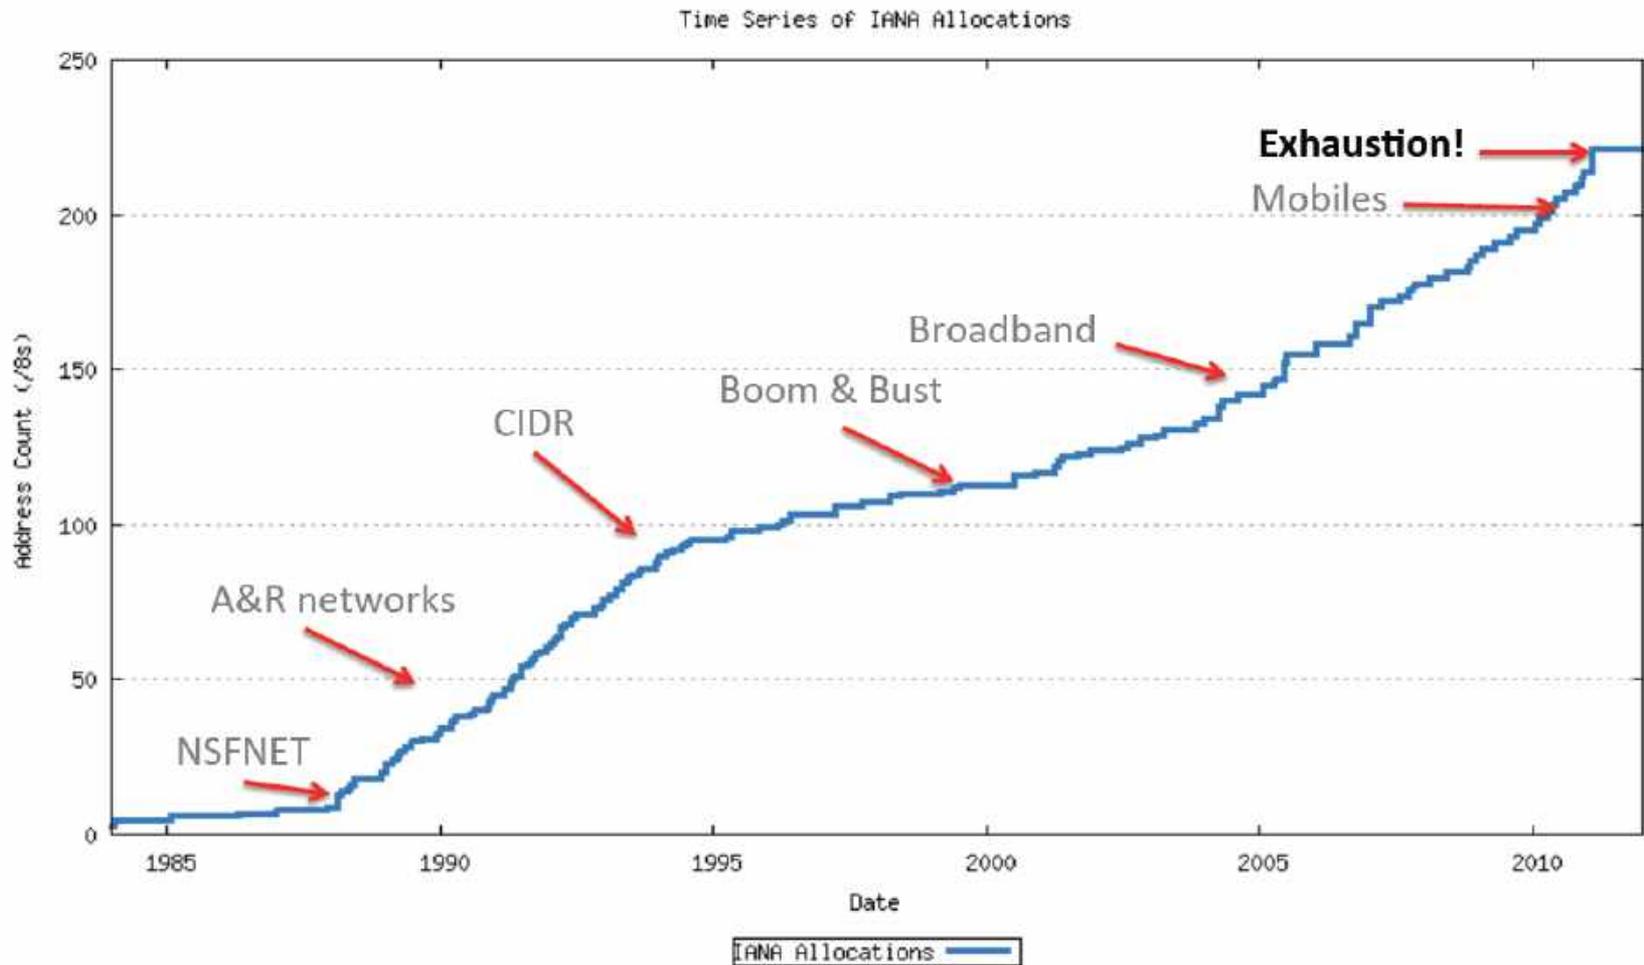
- *iana* = Internet Assigned Numbers Authority



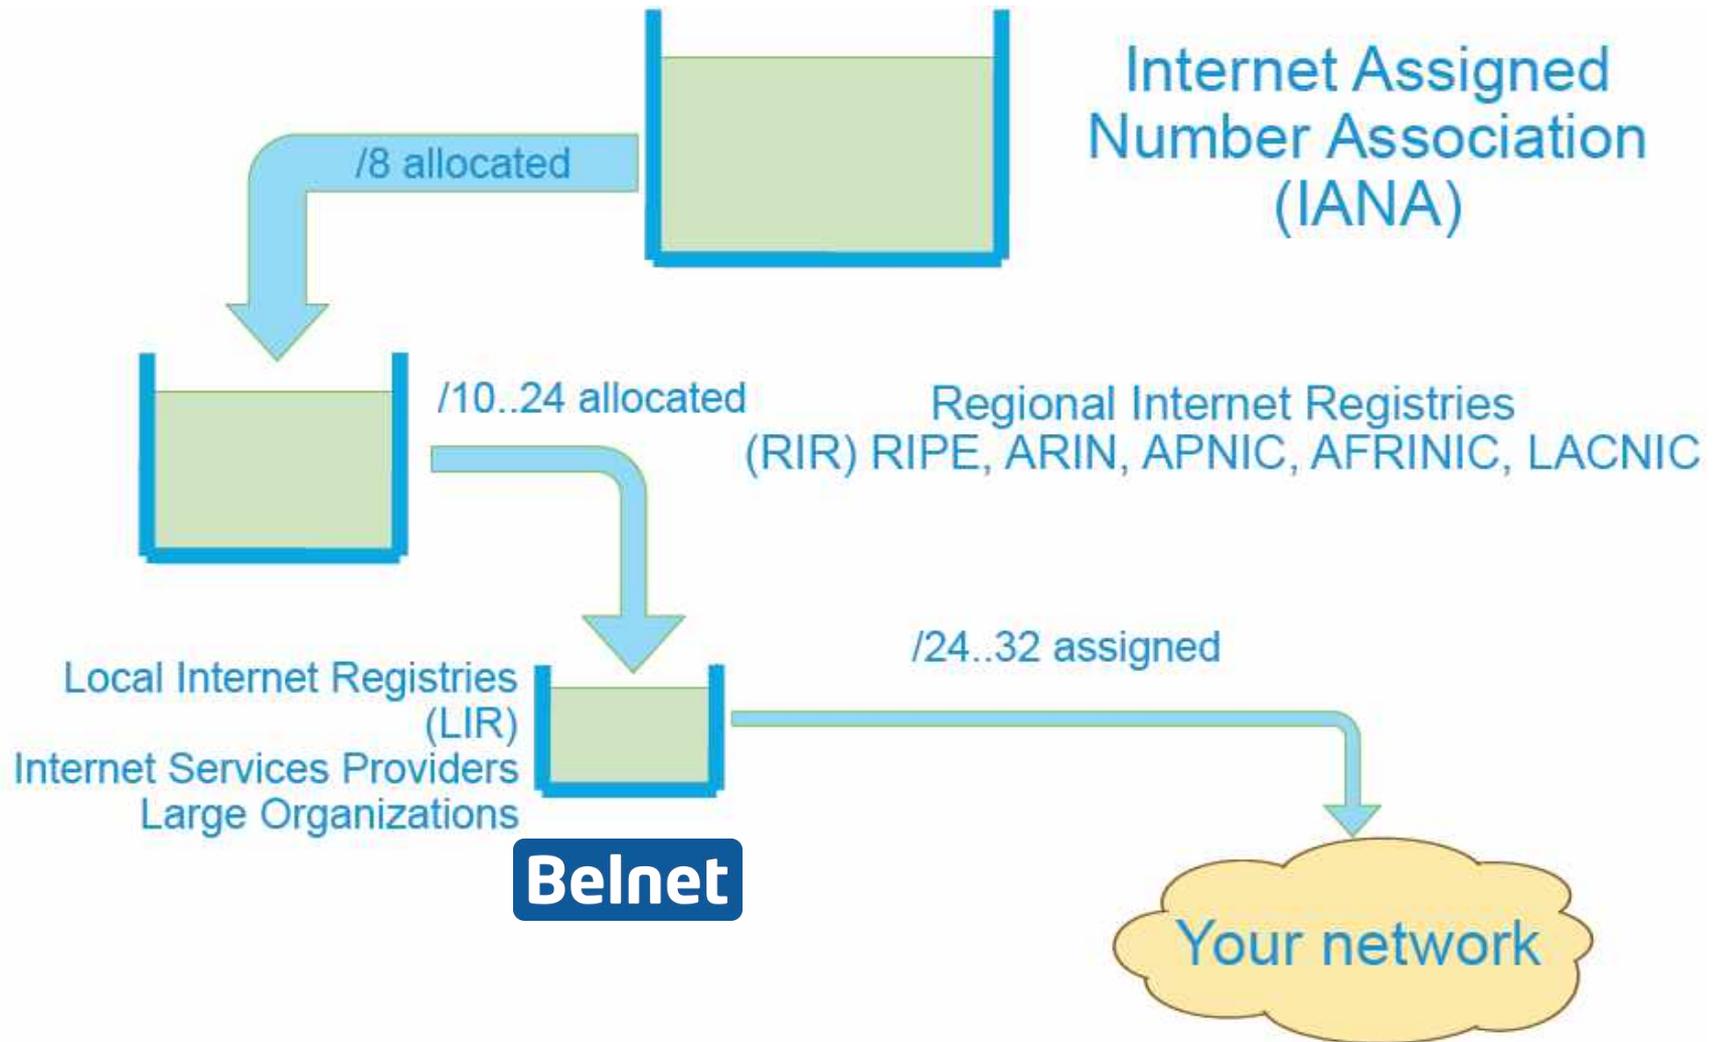
Regional Internet Registries (RIR)



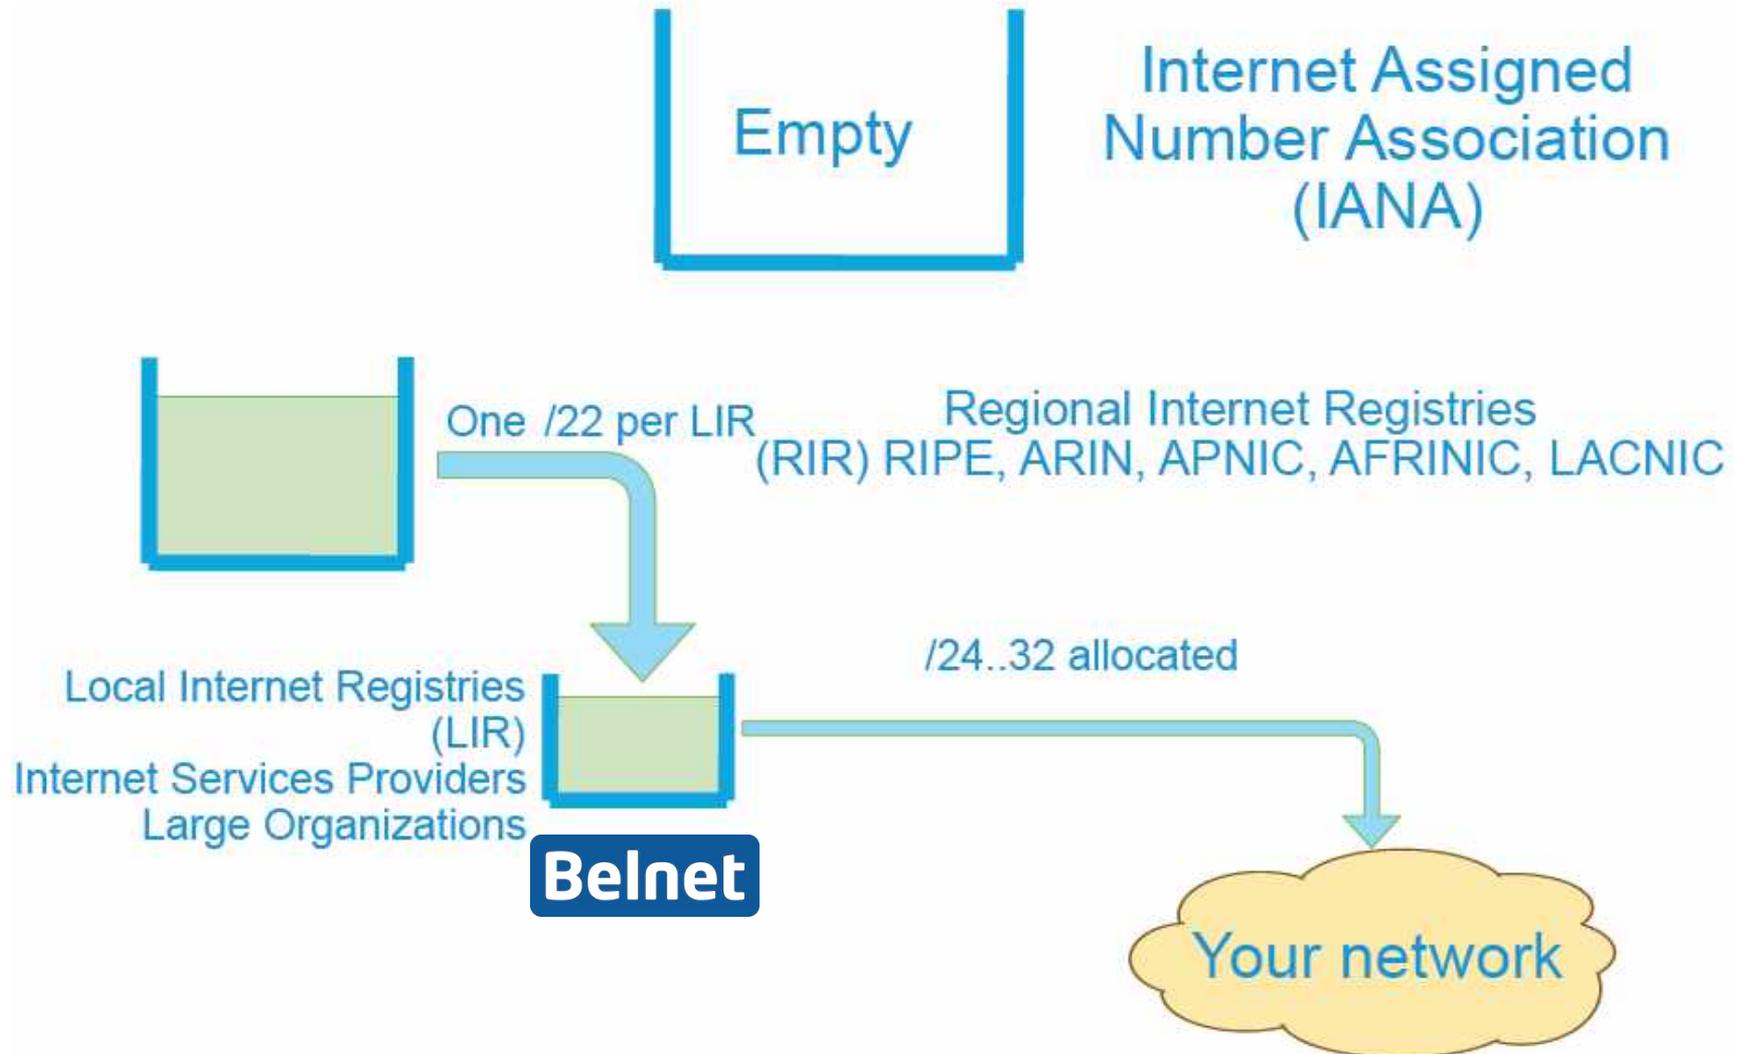
Problem: IANA has max 220 x /8s



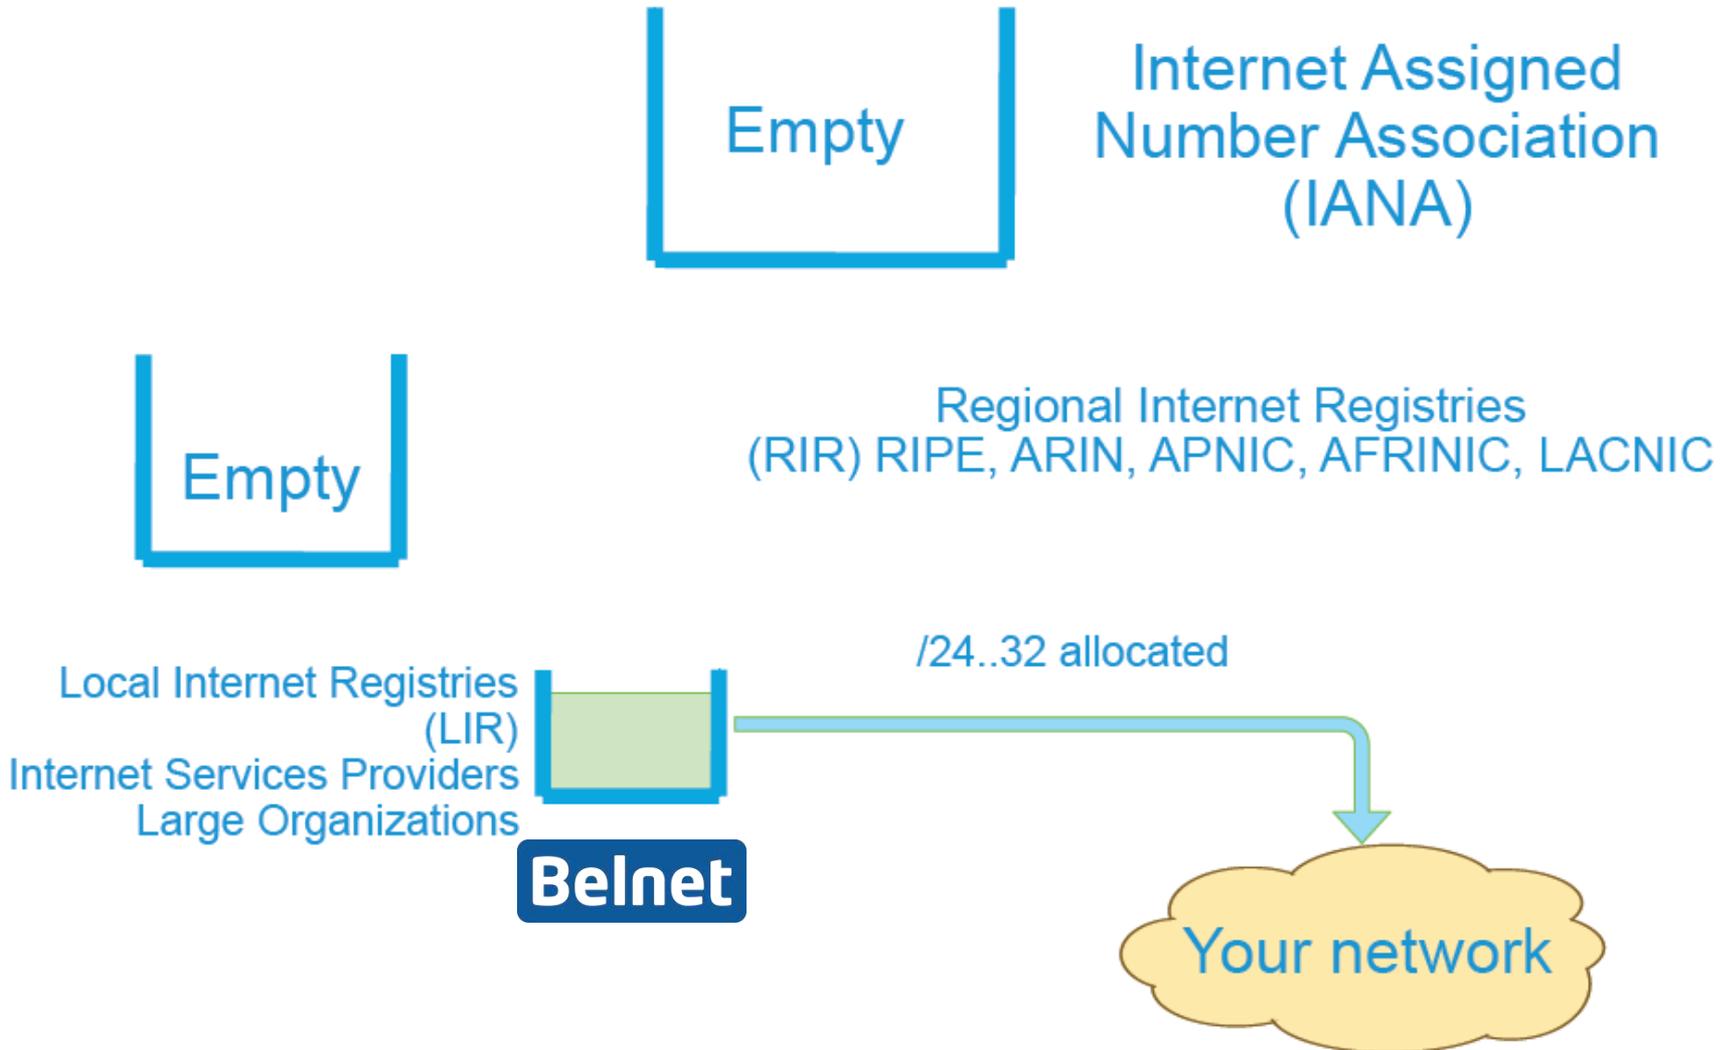
IPv4 Allocation before 2011



IPv4 Allocation : Q1 2011



IPv4 Allocation : 2016 ?



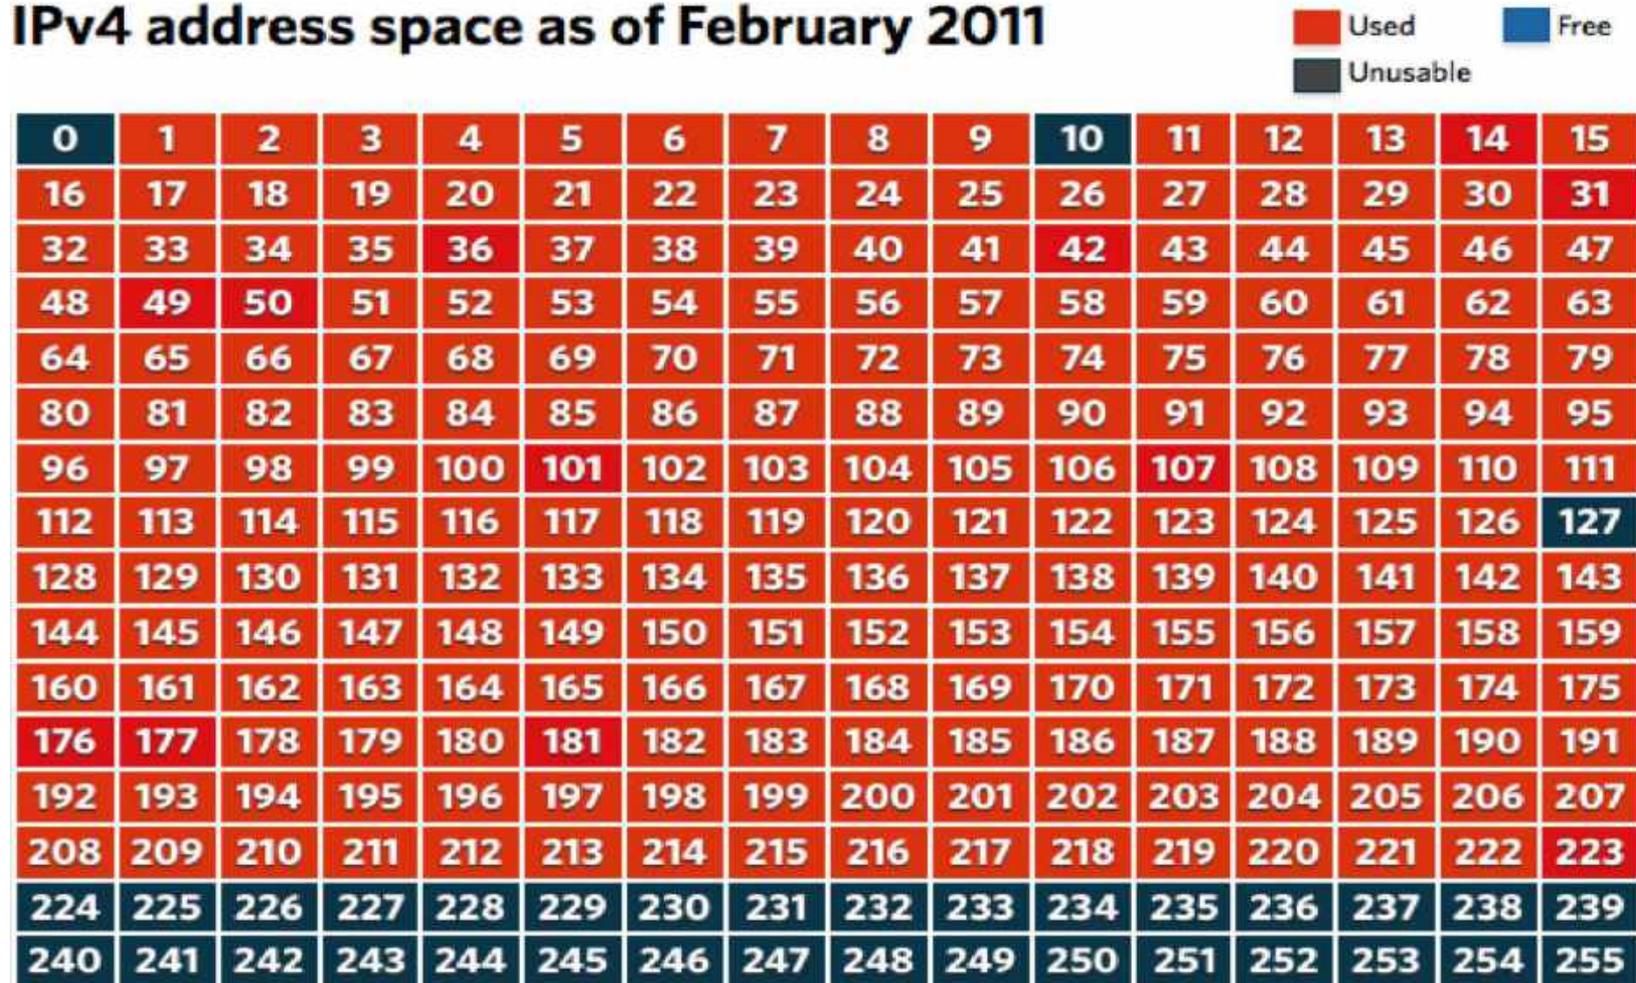
What is left ?

- The entire public IPv4 internet = 220 x /8
- In February 2011 IANA has allocated the last /8...
- And even RIR's are running out...
 - APNIC handed out last /8 in April 2012
 - RIPE is soon going to start allocating from its last /8
 - RIPE is asking legacy holders to become LIR
 - RIPE now allows transfers of allocations
 - IANA IPv4 Address Space Registry: list of all /8 attributions
 - Microsoft paid \$15M for 666K Nortel addresses

<http://www.iana.org/assignments/ipv4-address-space/ipv4-address-space.xml>

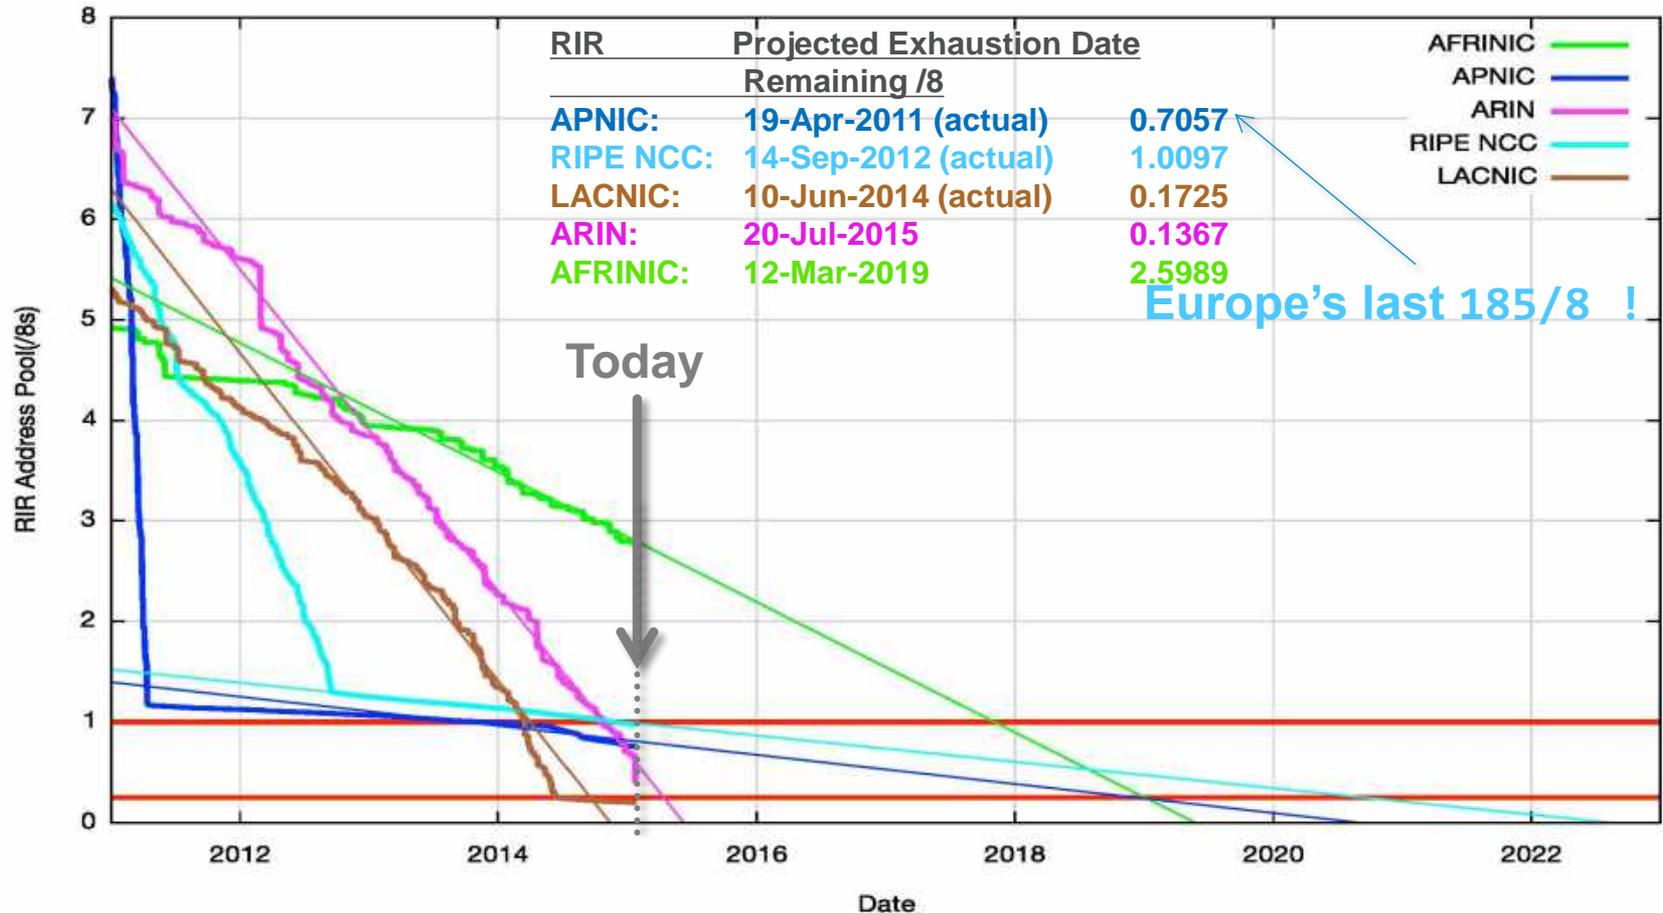
What is left ? NOTHING !

IPv4 address space as of February 2011



What is left?

RIR IPv4 Address Run-Down Model

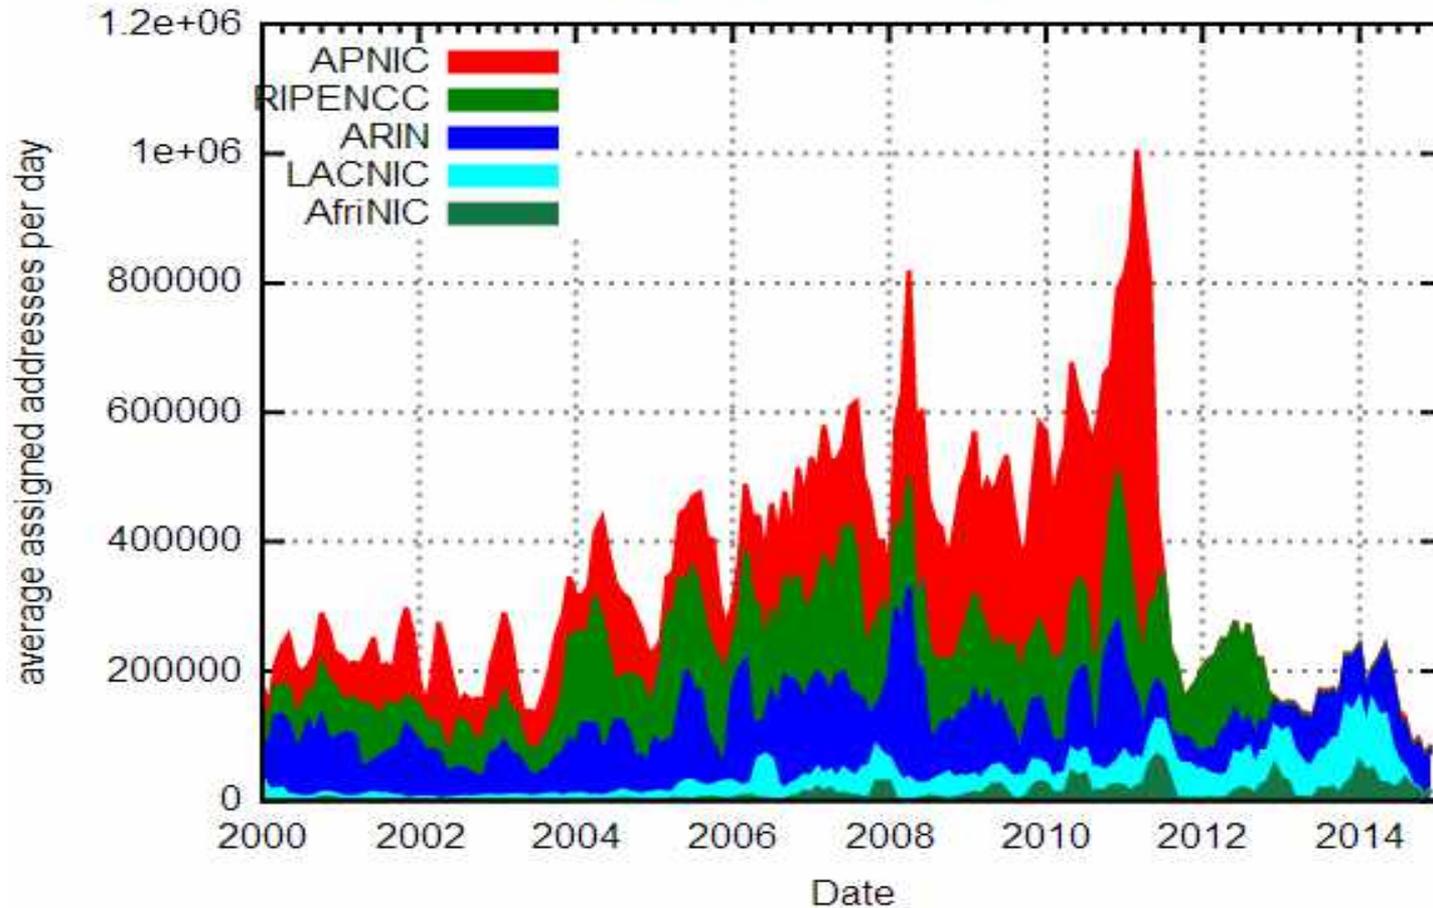


<http://www.potaroo.net/ispcol/2015-01/addressing2014.html>



Allocation timeline

Daily assignment rate per RIR





**You don't want to get disconnected,
do you ?**



IPv6

The Unavoidable Solution

But first, a bit of History

- 1995 : First introduction of IPv6
- Standard work finished in 1998
- Already 20 years ago ! In RFC 1883, RFC 1884
- Belnet full IPv6 since 2003 → Early adopter

- **What about IPv5 ?**

Yes, that did exist, as experimental ***Resource Reservation Protocol*** intended to provide QoS for multimedia, designed to co-exist with IPv4. Defined as *Internet Stream Protocol version 2* (ST2) in RFC 1819

IANA IP Version Numbers

- 4-bit IP version field in packet header

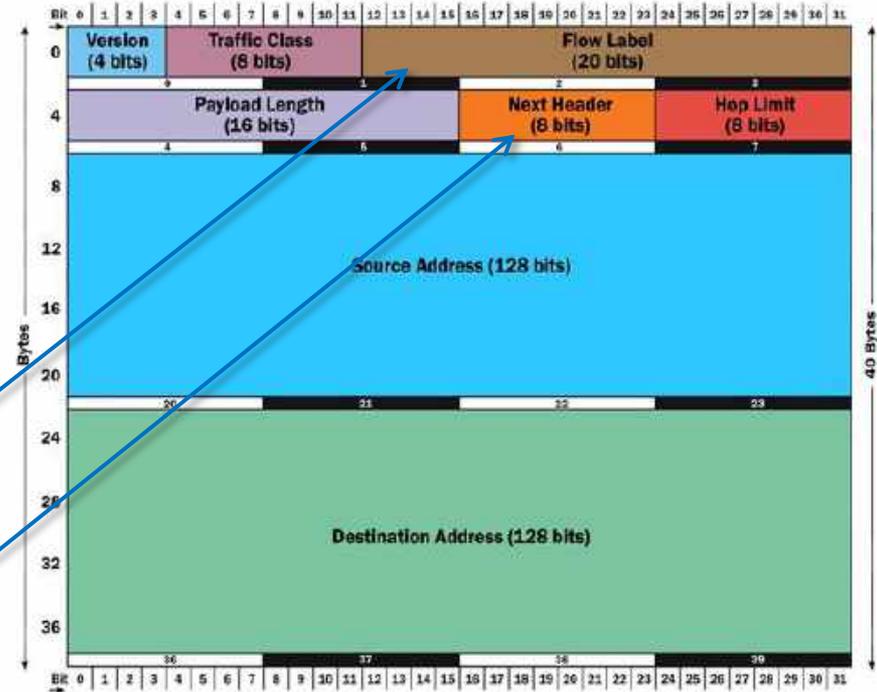
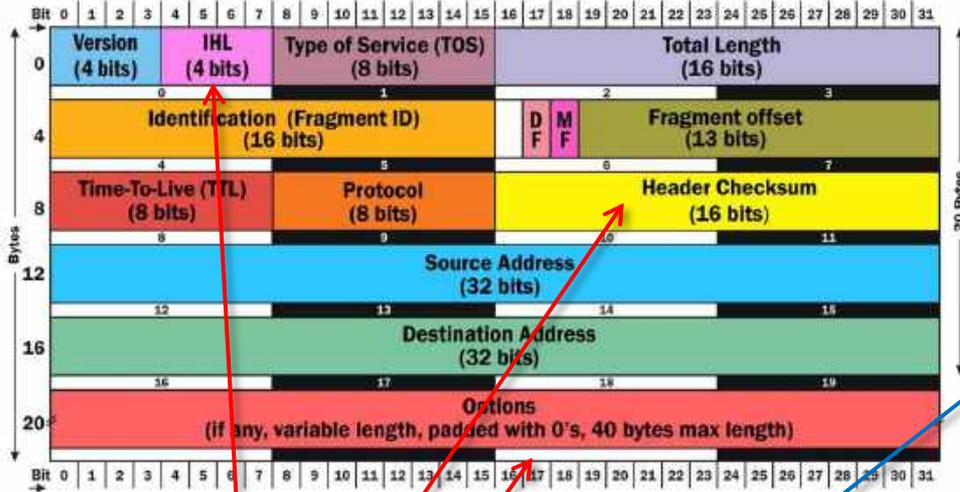
#	Keywor	Version	Reference
0-1		Reserved	[Jon Postel][RFC492]
2-3		Unassigned	[Jon Postel]
4	IP	Internet Protocol	[RFC791][Jon Poste]
5	ST	ST Datagram Mode	[RFC1190][Jim Forgi]
6	IPv6	Internet Protocol version	[RFC1752]
7	TP/IX	TP/IX: The Next Internet	[RFC1475]
8	PIP	The P Internet Protocol	[RFC1621]
9	TUBA	TUBA	[RFC1347]
10-14		Unassigned	[Jon Postel]
15		Reserved	[Jon Postel]

IPv6 Design Goals & Benefits

Objectives

- **128 bits** instead the 32 of IPv4
- **Global reachability and flexibility, end-to-end connectivity**
- **Auto-configuration:**
IPv6 host can auto-configure itself with an IPv6 globally unique address
- **Simpler packet header** and simpler processing in hardware.
 - Headers are aligned to 64 bits for easier storage & access in memory
 - Half of the previous IPv4 header fields were removed
- **Better support for QoS**, even when packet payload encrypted with IPsec
- **End-to-end security**, IPsec is mandatory **in theory, but never in practice**
RFC 4294 : “MUST”
RFC 6434 → became “SHOULD”
- **Mobility** is built in
↑
Nokia’s influence

Packet Header: IPv4 vs IPv6



Remove:

- IHL
- Header checksum
- Options

Add

- Flow Label
- Extension header

Jumbograms !

Packets with payload > 64K and < 4Gb
RFC 2675

Differences IPv4 – IPv6

- Different types and scope of addresses
- No broadcast, thus no ARP (replaced by NDP)
- Relies heavily on Multicasting
- Auto-configuration instead of DHCP?
 - Stateless vs Stateful
- Common to have multiple addresses on an interface. What IP will be used to source traffic?
- No Fragmentation along the way
 - Only at the source
 - Path MTU discovery is crucial (RFC 1981)

IPv6: Hexadecimal

- Hexadecimal system : 16 values

0, 1, 2, 3, 4, 5, 6, 7, 8, 9, a, b, c, d, e, f, ...
10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, ...

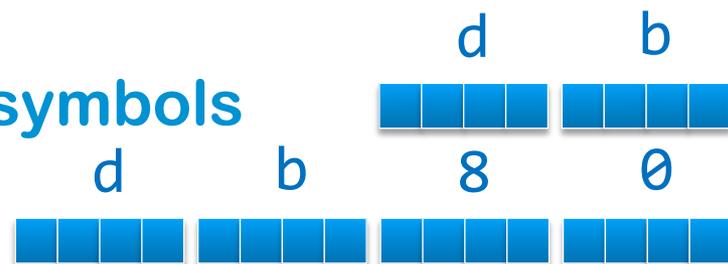


- We need 4 bits to represent 16 values

4 because $2^4 = 16 \rightarrow 1 \text{ hex symbols} = 4 \text{ bits}$

- 1 byte = 8 bits = two hex symbols

2 bytes = 16 bits, e.g.



IPv6 Address format

- IPv4 address = 4 groups of 1 byte, in decimal

– 193.190.0.1 ← Same size in # of IP addresses !

- IPv6 address = 8 groups of 2 bytes

– 20a0:db80:1234:ff23:a118:dead:beef:c911
– 2001:06a8:0000:0000:0000:0000:0000:0001

How many IPv6 addresses ?

- **IPv4 : each byte = 256 values**

$256 \times 256 \times 256 \times 256 = 4.294.967.296$ (4.3 billion) IP addresses

- **IPv6 : each group = 2 bytes = 65536 values**

2001 : 06a8 : 0000 : 0000 : 0000 : 0000 : 0000 : 0001

$65536 \times 65536 \times 65536 \times 65536 \times 65536 \times 65536 \times$

65536×65536

$= 340.282.366.920.938.463.463.374.607.431.768.211.456$

= 340 billion billion billion billion IP addresses

340.282.366.920.938.463.463.374.607.431.768.211.456

How much did you say ?

- Currently 130 million people are born each year.
- If this growth rate remains the same until the sun dies in 5 billion years,
- And all of these people live to be 72 years old...

→ They can all have 53x the address space of the IPv4 Internet...

...For every second of their lives...

And how many for Belnet Customers ?

- Every organisation connected to Belnet receives
 $a/48 = 65536 \times /64$
or 65K possible subnets / VLANs
- A $/64$ subnet
= 18.446.744.057.709.551.616 interface IDs (2^{64})
- Total # IPs : $65.536 \times 18.446.744.057.709.551.616$



IPv6 Address Syntax & Semantics

IPv6 Address Syntax

- IPv6 addresses can be very long

- Rules & Guidelines in [RFC5952](#)

“A Recommendation for IPv6 Address Text Representation”

- Simplifications

- These are rules to compress or shorten IPv6 addresses
- It is important to follow these rules for an uniform notation
- Wrong use of the rules may result in a different address when expanded!

IPv6 Address Syntax / RFC 5952

2001:0db8:0000:0000:0000:0000:0000:0001

Rule #1

Leading zeroes in a 16-bit group must be dropped

2001:0db8:0000:0000:0000:0000:0000:0001

→ 2001:db8:0:0:0:0:0:1

IPv6 Address Syntax / RFC 5952

2001:0db8:0000:0000:0000:0000:0000:0001

Rule #2

One series of **zeroes** and ‘:’ may be **compressed** to ‘::’

2001:db8:0:0:0:0:0:1

→ 2001:db8::1

‘::’ may only appear once

IPv6 Address Syntax / RFC 5952

Rule #2

(a) If several series possible, prefer the longest

2001:db8:0:5:0:0:0:1 → 2001:db8:0:5::1

(b) If equal size, prefer the leftmost

2001:db8:0:0:5:0:0:1 → 2001:db8::5:0:0:1

(c) No :: compression for only *one* block of zeroes

2001:db8:f:3:3:d:0000:1

→ 2001:db8:f:3:3:d:0:1 not 2001:db8:f:3:3:d::1

IPv6 Address Syntax / RFC 5952

Advise #1

Write hex values in **lowercase** for legibility

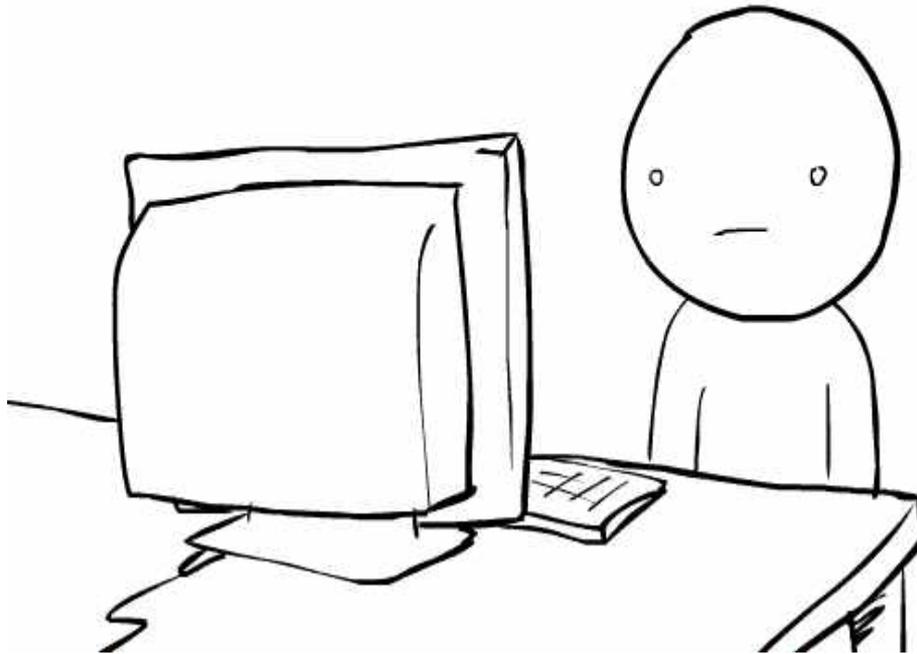
- Capital case D and 0 can quite often be misread
- Capital B and 8 can also be misread

2001:0DB8:0.0.0.1

2001:0db8:0.0.0.1 – **Most correct**

according to RFC5952

Let's do some exercises...



IPv6 Address Notation - Exercises

- Which are correct notations of a /32 prefix starting with 2001:0db8

- A. 2001:0db8 ✗
- B. 2001:0db8/32 ✗
- C. 2001:0db8::/32 ✓
- D. 2001:db8::/32 ✓

- How do you correctly compress this:

2001:0db8:0000:0000:0000:0000:0000:0c50

- A. 2001:0db8:0:0:0:0:0:0c50 ✓
- B. 2001:0db8::0c50 ✓
- C. 2001:db8::c50 ✓
- D. 2001:db8::c50 ✗

IPv6 Address Notation - Exercises

- How do you correctly compress this:

2001:0db8:0000:0000:b450:0000:0000:00b4

- A. 2001:db8::b450::b4 ❌
- B. 2001:db8::b450:0:0:b4 ✓ – Most correct according to RFC5952
- C. 2001:db8::b45:0000:0000:b4 ❌
- D. 2001:db8:0:0:b450::b4 ✓

- How do you correctly compress this:

2001:0db8:00f0:0000:0000:03d0:0000:00ff

- A. 2001:0db8:00f0::3d0:0:00ff ✓
- B. 2001:db8:f0:0:0:3d0:0:ff ✓
- C. 2001:db8:f0::3d0:0:ff ✓ – Most correct according to RFC5952
- D. 2001:0db8:0f0:0:0:3d0:0:0ff ✓

IPv6 Address Notation - Exercises

- How do you correctly compress this:

2001:0db8:0f3c:00d7:7dab:03d0:0000:00ff

- A. 2001:db8:f3c:d7:7dab:3d:0:ff ❌
- B. 2001:db8:f3c:d7:7dab:3d0:0:ff ✓
- C. 2001:db8:f3c:d7:7dab:3d0::ff ✓ – Most correct
- D. 2001:0db8:0f3c:00d7:7dab:03d::00ff ❌



IPv6 Addressing Architecture

IPv6 Addressing Architecture

- IPv6 : 128 bits

- 64 most significant bits are for the network part
- 64 least significant bits are for network interfaces



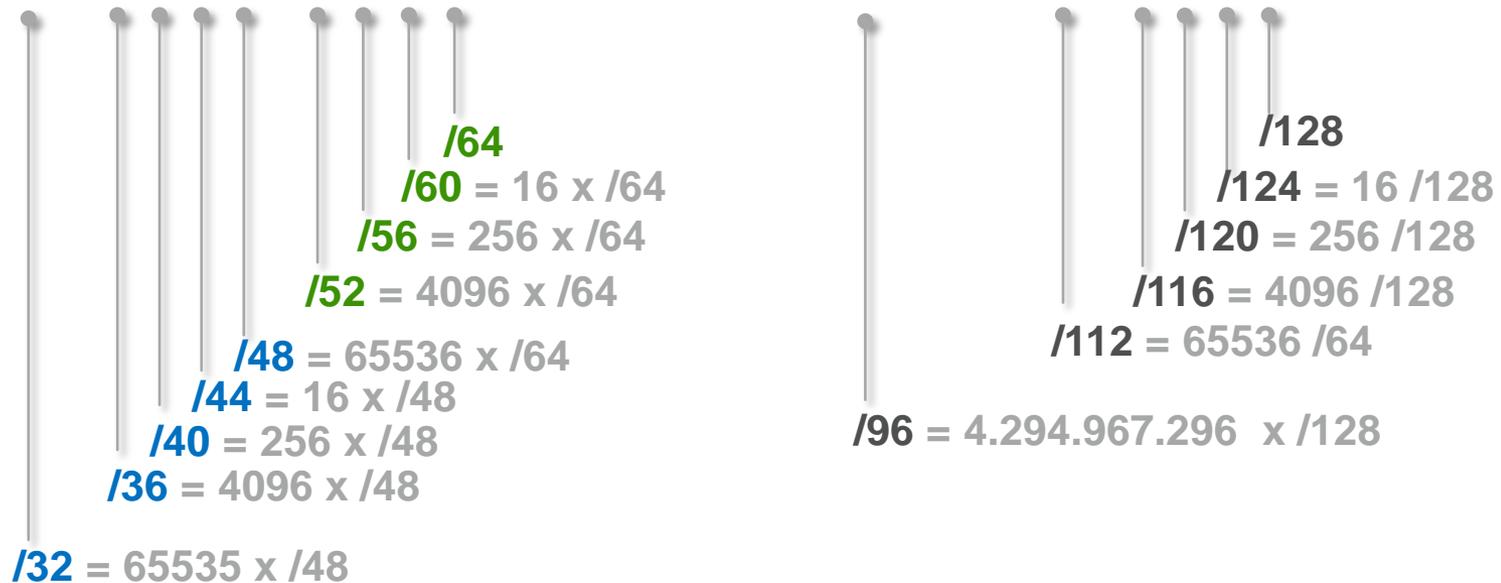
- Belnet is 2001:6a8:: /32



IPv6 Prefixes

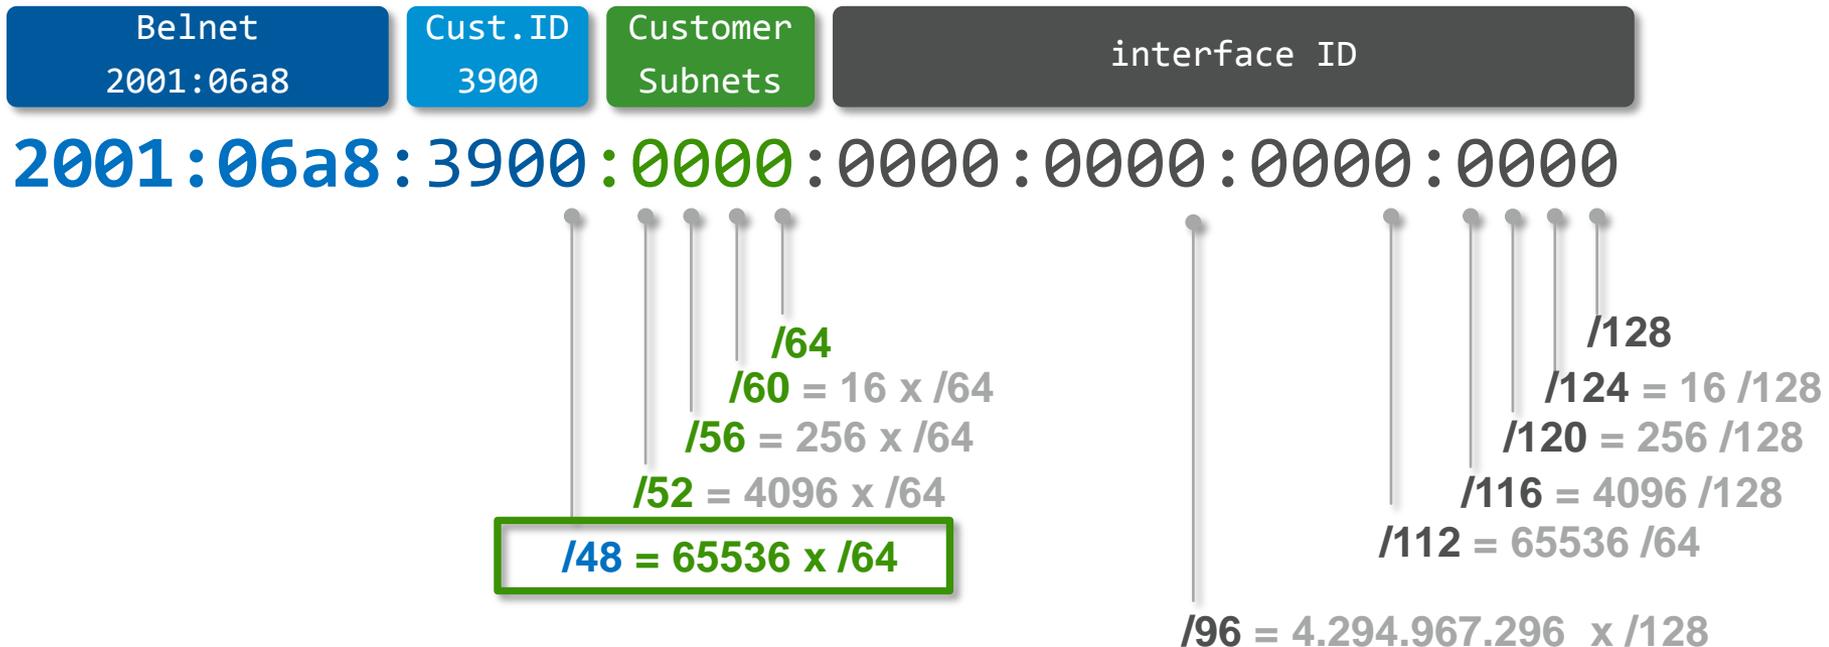
- For practical and readability reasons, it is advised to / on 4 bit boundaries (ie. on a hex digit)

2001:06a8:0000:0000:0000:0000:0000:0000



IPv6 Prefixes / Subnets

- Example: Université de Namur : 2001:06a8:3900::/48



$/48 = 16 \times /52$ or $256 \times /56$ or $4096 \times$

$/60 = 16 \times /56$ or $256 \times /60$

$/56 = 16 \times /60$

RFC 4291

IPv6 Addressing Architecture

- **Defines the IPv6 address format**
- **Three types:**
 - **Unicast**
 - Packets sent to unicast address are delivered to **only one** single interface identified by that address
 - **Anycast**
 - **One-to-Nearest:** Packets sent to anycast address are delivered to **one of** the interfaces identified by that address; usually the “nearest” one, according to the routing protocol’s measure of distance.
 - Multiple devices share same address (load balancing)
 - **Multicast**
 - **One-to-Many:** Packet sent to multicast address is delivered to **all** interfaces identified by that address
 - **No broadcast** addresses (superseded by multicast)

Main IPv6 Address Types (1/3)

Prefix	Explanation	IPv4
::/0	Default route	0.0.0.0/0
::/128	Unspecified This address may only be used as a source address by an initialising host before it has learned/received its own address. It indicates the absence of an address. It may never be assigned to a node. A router must not forward IPv6 packets with this as destination address.	0.0.0.0/32
::1/128	Loopback This address is used when a host talks to itself over IPv6. This often happens when one program sends data to another. A router must not forward IPv6 packets with this as destination address.	127.0.0.1
ff00::/8	Multicast These addresses are used to identify multicast groups. They should only be used as destination addresses, never as source addresses	224.0.0.0/4

Multicast in IPv6

- Nodes “subscribe” to multicast groups, like tuning in to a radio station
 - Some examples:
 - ff01::1 - All nodes in local interface
 - ff02::2 - All nodes in local link
 - ff02::5 - All nodes in local site
 - ff05::2 - All routers in the site
 - ff02::fb - All DNS servers on the local link
 - ff08::fb - All DNS servers in the organisation
 - ff02::1:3 - All DHCPv6 servers on the local link
 - ff05::1:3 - All DHCPv6 servers in the site

Notes on Multicast in IPv6

- **Uses UDP or SCTP**

- RFC 4960 - **Stream Control Transmission Protocol**
- Basically TCP over UDP, but handling multiple streams

- **You don't assign a multicast address to a node, you run an application that "joins" a multicast group**

Main IPv6 Address Types (2/3)

Prefix	Explanation	IPv4
fc00::/7	Unique Local Address (ULA) – RFC 4193 These addresses are reserved for local use in home and enterprise environments and are not public address space . They can be routed inside a site. fc00::/8 – planned to be globally managed (not used yet) fd00::/8 – can be assigned locally by administrator	Private or RFC 1918 space: 10.0.0.0 /8 172.16.0.0 /12 192.168.0.0 /16
fe80::/64	Link-Local Addresses (LLA) These addresses are automagically created and assigned to every IPv6 interface used on a single link or a non-routed common access network, such as an Ethernet LAN. They do not need to be unique outside of that link. (local scope) <u>All IPv6 routers prevent link-local addresses from being forwarded between interfaces.</u> Routers use this address type as next-hop address. (e.g.: fe80::215:60ff:fe00:f126)	IPv4 link-local 169.254.0.0/16 See RFC 3927

ULA $fc00::/7$ – Optional info



Prefix To identify Local IPv6 unicast addresses

- L** – **1** if prefix is locally assigned : $fd00 = 1111\ 1101$
- **0** may be defined in the future : $fc00 = 1111\ 1100$

Global ID 40-bit global identifier used to create a globally unique prefix (high probability) RFC 4193 defines an algorithm (using MAC address & time stamp) for the generation of pseudo-random global IDs.

<http://www.kame.net/~suz/gen-ula.html>

<http://www.sixxs.net/tools/grh/ula/>

Subnet ID 16-bit Identifier of a subnet within a site. Should be structured hierarchically, by site administrators

Interface ID 64-bit Interface ID

ULA `fc00::/7` – Optional info

- NOT OBLIGED to use the pseudo-random generator
- Default exterior routing protocol behaviour should be to **NOT to advertise `fc00::/7` prefixes**
- **BGP must filter out `fc00::/7` prefixes**, both in and out
- Site border routers and firewalls should be configured to filter out `fc00::/7`. Default behaviour should be to reject
- Site border routes should respond with ICMPv6 Destination Unreachable
- ULA do not aggregate due to the randomness of the Global ID
- To be used locally and independently of ISP-based IPv6 unicast allocations
- Can be used for inter-site VPNs

Main IPv6 Address Types (2 ¾ /3)

Prefix	Explanation	IPv4
fec0::/10 OBSOLETE	Site Local Address These addresses were defined in RFC 3513 and were destined for use only locally on sites. RFC 3879 deprecated the use of these addresses. This kind of addresses was superseded by the ULAs defined in RFC 4193	

Main IPv6 Address Types (3/3)

Prefix	Explanation	IPv4
2000::/3	<p>Global Unicast – RFC 3587</p> <p>The IANA public IPv6 address space distributed over all RIRs (Regional Internet Registry)</p> <p>This is only 12.5% of the total available IPv6 space !</p> <div style="display: flex; justify-content: space-around; align-items: center;"> <div style="text-align: center;"> <p><small>3</small></p> <p>001</p> </div> <div style="text-align: center;"> <p><small>45 bits</small></p> <p>global routing prefix</p> </div> <div style="text-align: center;"> <p><small>16 bits</small></p> <p>subnet ID</p> </div> <div style="text-align: center;"> <p><small>64 bits</small></p> <p>interface ID</p> </div> </div>	All IPv4 space
2001:db8::/32	<p>Documentation</p> <p>These addresses are used in examples and documentation. They should never be source or destination addresses.</p>	192.0.2.0/24 198.51.100.0/24 203.0.113.0/24 (RFC 5737)
2001:0002::/48	<p>Benchmarking – RFC 5180</p> <p>IPv6 complementary update of RFC 2544. Purpose is to evaluate IPv6 forwarding performance. Should be kept inside an ASN and not appear in global BGP.</p>	198.18.0.0/15 (RFC 2544)

2000::/3 – Optional info



001 First three bits should not be 000. That is reserved for global unicast addresses whose interface ID is not 64 bits. This means that the first usable prefix starts with binary 001, or 20 hex, and that no global unicast should start with 00 to 1F

global routing prefix Value assigned to a site (cluster of subnets/links). Designed to be hierarchically structured between RIRs and ISPs

subnet ID Identifier of a subnet within a site. Should be structured hierarchically, by site administrators

Interface ID 64-bit Interface ID

Only 12,5 % of IPv6 space

IPv6 Prefix	Allocation	Reference
-----	-----	-----
0000::/8	Reserved by IETF	[RFC4291]
0100::/8	Reserved by IETF	[RFC4291]
0200::/7	Reserved by IETF	[RFC4048]
0400::/6	Reserved by IETF	[RFC4291]
0800::/5	Reserved by IETF	[RFC4291]
1000::/4	Reserved by IETF	[RFC4291]
2000::/3	Global Unicast	[RFC4291]
4000::/3	Reserved by IETF	[RFC4291]
6000::/3	Reserved by IETF	[RFC4291]
8000::/3	Reserved by IETF	[RFC4291]
A000::/3	Reserved by IETF	[RFC4291]
C000::/3	Reserved by IETF	[RFC4291]
E000::/4	Reserved by IETF	[RFC4291]
F000::/5	Reserved by IETF	[RFC4291]
F800::/6	Reserved by IETF	[RFC4291]
FC00::/7	Unique Local Unicast	[RFC4193]
FE00::/9	Reserved by IETF	[RFC4291]
FE80::/10	Link Local Unicast	[RFC4291]
FEC0::/10	Reserved by IETF	[RFC3879]
FF00::/8	Multicast	[RFC4291]

Initial Allocations of the 2000::/3 block



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

Some other IPv6 Address Types (1/2)

Optional info

Prefix	Explanation	IPv4
2001:0000::/32	Teredo This is a mapped address allowing <u>IPv6 tunneling through IPv4 NATs</u> . (IPv6 over UDP). The address is formed using the Teredo prefix, the server's unique IPv4 address, flags describing the type of NAT, the obfuscated client port and the client IPv4 address, which is probably a private address. It is possible to reverse the process and identify the IPv4 address of the relay server.	No Equivalent
2002::/16	6to4 A 6to4 gateway adds its IPv4 address to this 2002::/16, creating a unique /48 prefix. As the IPv4 address of the gateway router is used to compose the IPv6 prefix, it is possible to reverse the process and identify the IPv4 address.	No equivalent but 192.88.99.0/24 has been reserved as the 6to4 relay anycast address
Note: Tunnelling with 6to4 Anycast has been officially abandoned by the IETF in May 2015		Prefix by the IETF.

RFC 7526: Deprecating Anycast Prefix for 6to4 Relay Routers

<http://www.potaroo.net/ispcol/2010-12/6to4fail.html>

Some other IPv6 Address Types (2/2)

Optional info

Prefix	Explanation	IPv4
2001:0010::/28 EXPERIMENTAL	ORCHIDv2 - RFC 7343 Overlay Routable Cryptographic Hash Identifiers These addresses are used for a fixed-term experiment. These identifiers are intended to be used as location-independent endpoint identifiers at applications and Application Programming Interfaces (APIs) and not as identifiers for network location at the IP layer. They should only be visible on an end-to-end basis and routers should not see packets using them as source or destination addresses.	No Equivalent

IPv6 Address Scope

Scopes define the limits within which packets with a given destination address can travel, and in which that address is meaningful.

- **Local or Global**
- **Determined by the first (leftmost) 10 bits**
- **One host can have several IPv6 addresses for the same interface !**
 - **Global** scope : how it is known to Internet:
`2001:418:1:0:230:5ff:fe0b:1f3c`
 - **Local** scope : how it is known locally on the link to the router:
`fe80::230:5ff:fe0b:1f3c`

Many IP addresses per node !

Not only because of scope but also different assignment methods
can the interface of a node can have many IP addresses

IPv6 Address	Description
fe80::2030:9139:9cd5:ab52	Automatic Link Local Unicast
2001:470:3d:3000::2:1	Manually assigned Global Unicast
2001:470:3d:3000::4:57db	DHCPv6 assigned Global Unicast
2001:470:3d:3000:2030:9139:9cd5:ab52	Global Unicast from SLAAC
fea4:73c2:e5b8:1000:2030:9139:9cd5:ab52	ULA Unicast from SLAAC

– e.g. you avoid SLAAC when using DHCP or Manual IP

● RFC 6724: Default Address Selection for IPv6

IPv6 Default Address Selection

– Optional info

● RFC 6724 In English

Address selection can be done for source and for destination

- **Address pairs** of same scope or type are preferred
- Use the **smallest scope** possible for the Destination address
- A non-deprecated address is preferred
- ISATAP or 6to4 addresses are not used if native IPv6 is available
- If all criteria are similar, address pairs with the longest common prefix are preferred
- For source address, **global addresses are preferred over temporary**
- In Mobile IP situations, home addresses are preferred over care-

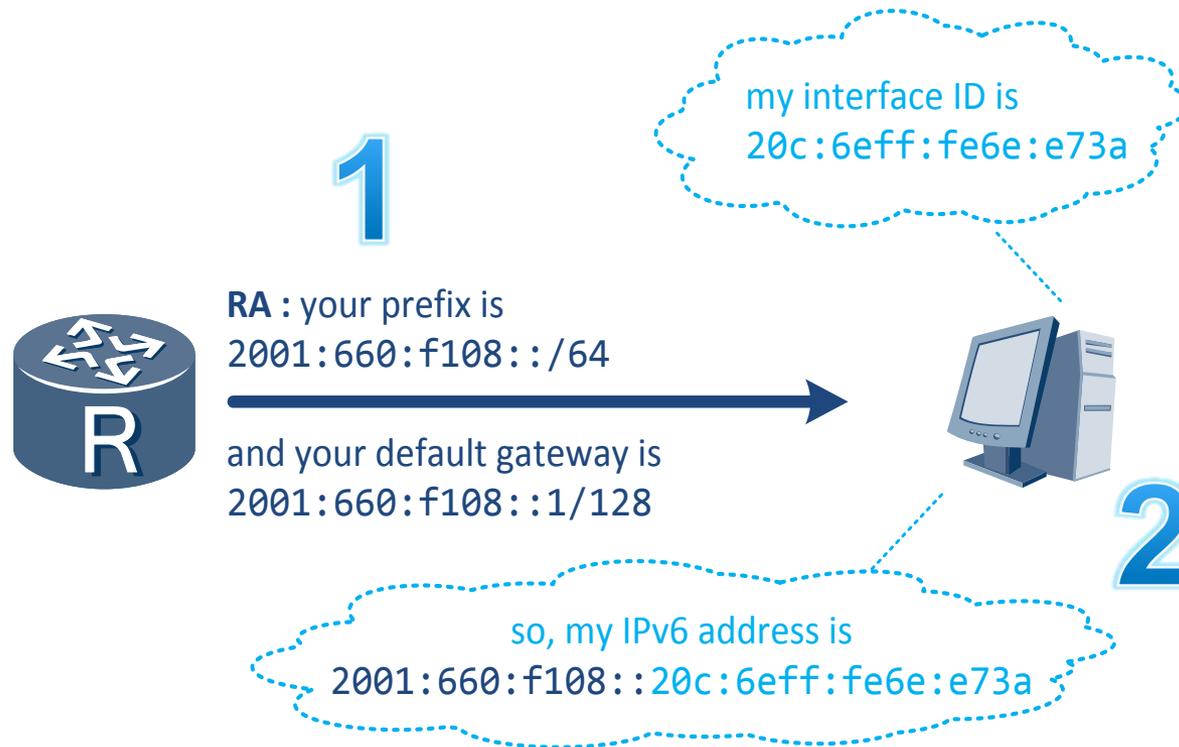


SLAAC & Privacy Issues

- **Stateless Address Auto-configuration : RFC 4862**
 - Stateless = no centralized structure/server storing everything
 - must have decentralized mechanism of assigning a globally unique IP address
 - This happens with **Router Advertisements** and **Router Solicitations**

How does SLAAC work ?

- RA : Router Advertisements



How does SLAAC work ?

- **RS : Router Solicitations (at node startup)**

- Not necessary to wait for RA:

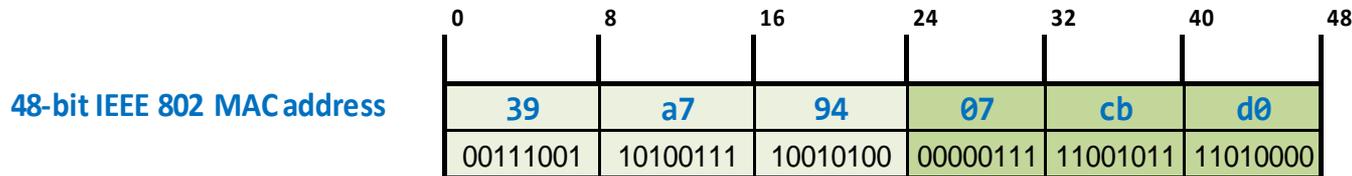


SLAAC automatic Interface ID

- Interface ID generated automatically
- Existing Mechanisms: (in evolutionary order)
 1. Modified EUI-64
 2. [RFC 4941](#) – *Privacy Extensions for SLAAC in IPv6*
 3. [RFC 7217](#) – *A Method for Generating Semantically Opaque Interface Identifiers with IPv6*

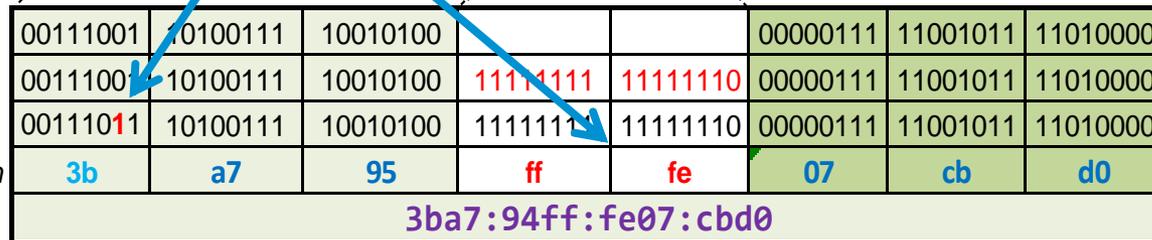
1. Modified EUI-64 Interface Identifier

= IPv6 address where the 64-bit interface ID is constructed from the 48-bit MAC address of an interface



3 2 1

1. Split MAC address
2. Insert "ffe" bit pattern
3. Invert bit 7 (u-bit)
4. Modified EUI-64 in hex notation



1. Modified EUI-64 Interface Identifier

■ Security Considerations for EUI-64

- **Your MAC can be followed**
- MAC = Vendor IDs : this makes it easier to scan IPv6 space
- See [RFC 5157](#) “IPv6 Implications for Network Scanning”,
which describes way the IPv6 search space can be reduced
 - Exploitation of address conventions, eg.
 - ::53 for DNS service, ::80 for web service
 - Use of words for easy reading :beef:, :dead:, :bad:, :dcaf:, ...
 - Linear search ::1, ::2, ...
 - See also [“Worm propagation Strategies in an IPv6 Internet”](#)

2. RFC 4941

Privacy Extensions for SLAAC in IPv6

- **Purpose: try to avoid MAC tracking**
- **How ?**
 - Generate (pseudo) random Interface IDs
 - Use Temporary Addresses
- **Side note : other easy ways to track you than MAC:**
 - Cookies
 - Remote physical device finger printing
 - <http://www.caida.org/publications/papers/2005/fingerprinting/>
 - You dynamically update your/a DNS sever with your address...

2. RFC 4941 – **Optional info** Privacy Extensions for SLAAC in IPv6

- **How is it generated ?**

1. Keep historic 64-bit value of previous generated address
2. Feed that value to the generator
3. Compute MD5 hash from the above result and
4. Take leftmost 64-bits and set bit 6 to 0
= Local significance (1=global))
5. Use this as generated interfaced ID

2. RFC 4941

Privacy Extensions for SLAAC in IPv6

- **Security Consideration for privacy extensions**

- Ingress filtering is a means of preventing the use of spoofed source addresses in DDoS attacks.
- This mechanism may fail with temp addresses: **hard to protect a machine that regularly changes IP** from another machine that also regularly changes (spoofed) IP..

- **Security Consideration for SLAAC**

- Anyone can send RAs
- Make sure to block rogue RAs

More on SLAAC & Privacy Issues

- **Implementations may vary with manufacturer**

- **Randomized** Interface ID generation.
If disabled, then modified EUI-64 is used.
 - **Temporary** address generation, unless disabled.
If enabled, you will get not just one unicast global address generated during SLAAC, **but two** - the second one will have a shorter lifetime, and automatically change over time
 - IP address recalculated every 24h
- All Major OSes, including Microsoft Windows (enabled by default since Vista) mainly based on Privacy Extensions RFC 4941
 - Linux: CONFIG_IPV6_PRIVACY option, but disabled by default (as requested by RFC)

3. RFC 7217 – A Method for Generating Semantically Opaque Interface IDs with SLAAC

- April 2014
- Improved, to guarantee anonymity (not MAC based)
- More stable privacy addressing
- Uses hashing of several parameters, including
 - a secret key,
 - IP prefix,
 - Interface name (eg. wlan0)
- Algorithm largely based on that in [RFC 6528](#)
“Defending against Sequence Number Attacks”

Privacy Considerations for Internet Protocols – **Optional info**

- **Since 2013 and PRISM Privacy came into the spotlight, and also at the IETF with RFC 6973 in July 2013**
 - *Every new TCP/IP protocol should, at its design stage, have a reflection on privacy concerns and how to limit exposure*
- **Before that, there was little RFC concern for privacy**
 - Except for IPv6 RFC 4941 (Privacy Extension for SLAAC) September 2007

Summarisation: IPv6 Address Configuration

- **3 Ways an interface can obtain an IPv6 address:**
 - Automatically generated (link-local, fe80) or SLAAC
 - Assigned by DHCPv6 ([RFC 3315](#))
 - Manually specified
- **Typical use:**
 - In Enterprise LANs it will be mostly DHCPv6 or SLAAC
 - For hosts that do not need well-known IP addresses
 - In ISP and infrastructure environments mostly manual / static configuration (routers, servers, ...)

IPv6 Address Configuration

■ Notes

- **If Manual, then SLAAC must be turned off**

Otherwise both may exist and machine will decide on its own what IP it will use for a given flow → unpredictable
(see [RFC 6724](#) “*Default Address Selection for IPv6*”)

- **SLAAC needs minimum /64**

Subnetting past /64 (ie. /96, /126 etc) **breaks** SLAAC



IPv6 Future

Who is already IPv6 ready ?

- **Equipment vendors** (routers, switches, firewalls, ...)
- **Software** (OS, applications, ...)
- **Smartphones** now ship with IPv6 capability
 - *“Facebook says it has seen users' News Feeds loading 20 percent to 40 percent faster on mobile devices using IPv6.”*
 - Computer World, April 2015
<http://www.computerworld.com/article/2909628/the-future-is-here-you-may-already-be-using-ipv6.html>
- **Networks**
 - Content: Google, Facebook, Bing, Cisco, Yahoo!, LinkedIn, Akamai, Belnet
 - ISPs: Comcast, T-Mobile, (US), XS4all (NL), Belnet, Proximus, Telenet
 - IXes

Why is it so slow to pick up ?

- **IPv6 is a topic that is often avoided, in practice**
 - many strategic and resources issues
 - lack of top management decision-making
 - lack of long-term vision
 - lack of IPv6 skilled engineers
 - lack IPv6 deployment best practices
 - very limited ISP IPv6 access deployment

→ **But not at Belnet !**

Why is it so slow to pick up ?

● Chicken and Eggs

- For applications to support IPv6, the network must first be able to deliver the packets
 - Yes but networks don't supply IPv6 connectivity because there are not many IPv6 applications that support it
- Content and Services need to adopt IPv6, but there are no users who can access it using IPv6
 - Yes but there are no users because there is no content

→ **But not at Belnet !**

Why is it so slow to pick up ?

- **IPv6 has momentarily lost some momentum because of hot topics absorbing all attention**
 - IoT
 - SDN / NFV
 - Cloud
 - 5G
- **But the need has not disappeared AT ALL**
 - It only gets worse

Where are we now today ?

■ June 2015: IPv6 enabled users

1. Belgium leads (worldwide) the users penetration with 32.36 %
2. USA 17.85 %
3. Germany 14.64 %
4. Peru 14.20 %
5. Luxembourg 12.56 %
6. Portugal 10.36 %
7. Switzerland 9.15 %
15. France 5.56 %
23. Netherlands 2.32%

Source: <https://www.vyncke.org/ipv6status/>

Where are we now today ?

- **If this trend continues**

- It is estimated we should achieve 50% by 2017
- This would be the inflection point when the full roll-out of IPv6 becomes a strategic plumbing decision of other network operators

- **Once they get there**

- Things will start going fast, and you better be there too.
- Others will...

Where is Belgium today ?

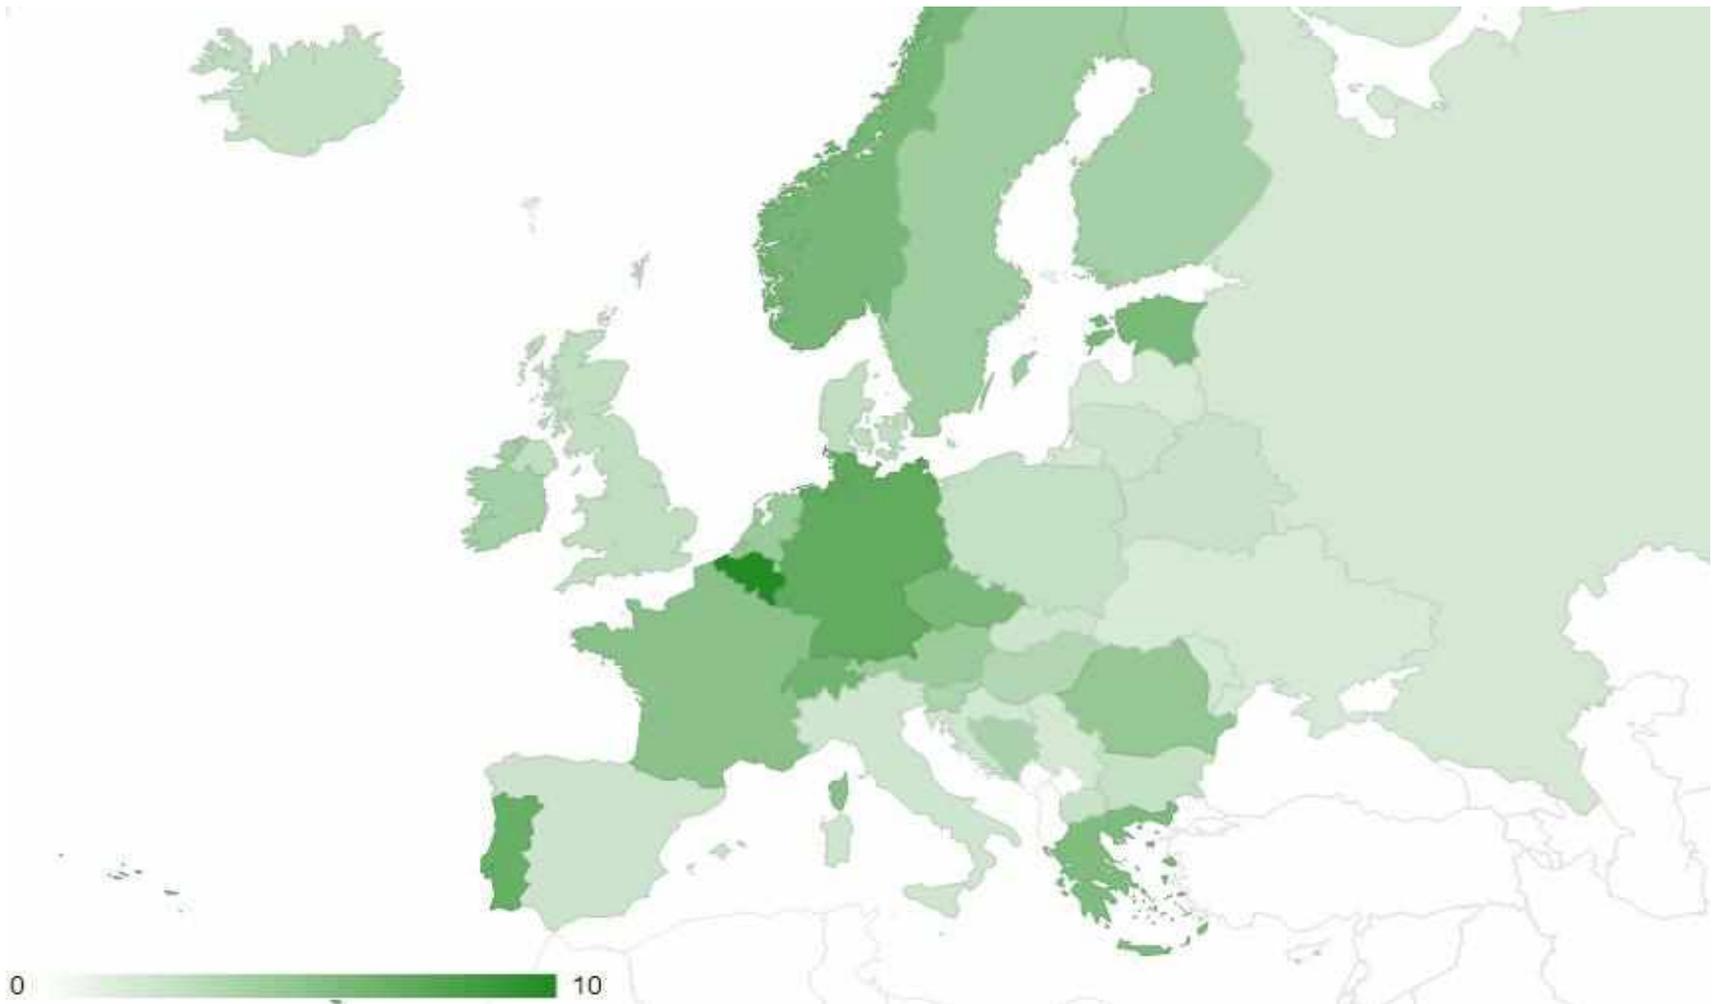
- <http://6lab.cisco.com/stats/>

Prefixes: 38.58 %
Transit AS: 78.24 %
Content: 50.14 %
Users 32.4 %
IPv6 Deployment: 49.79 %

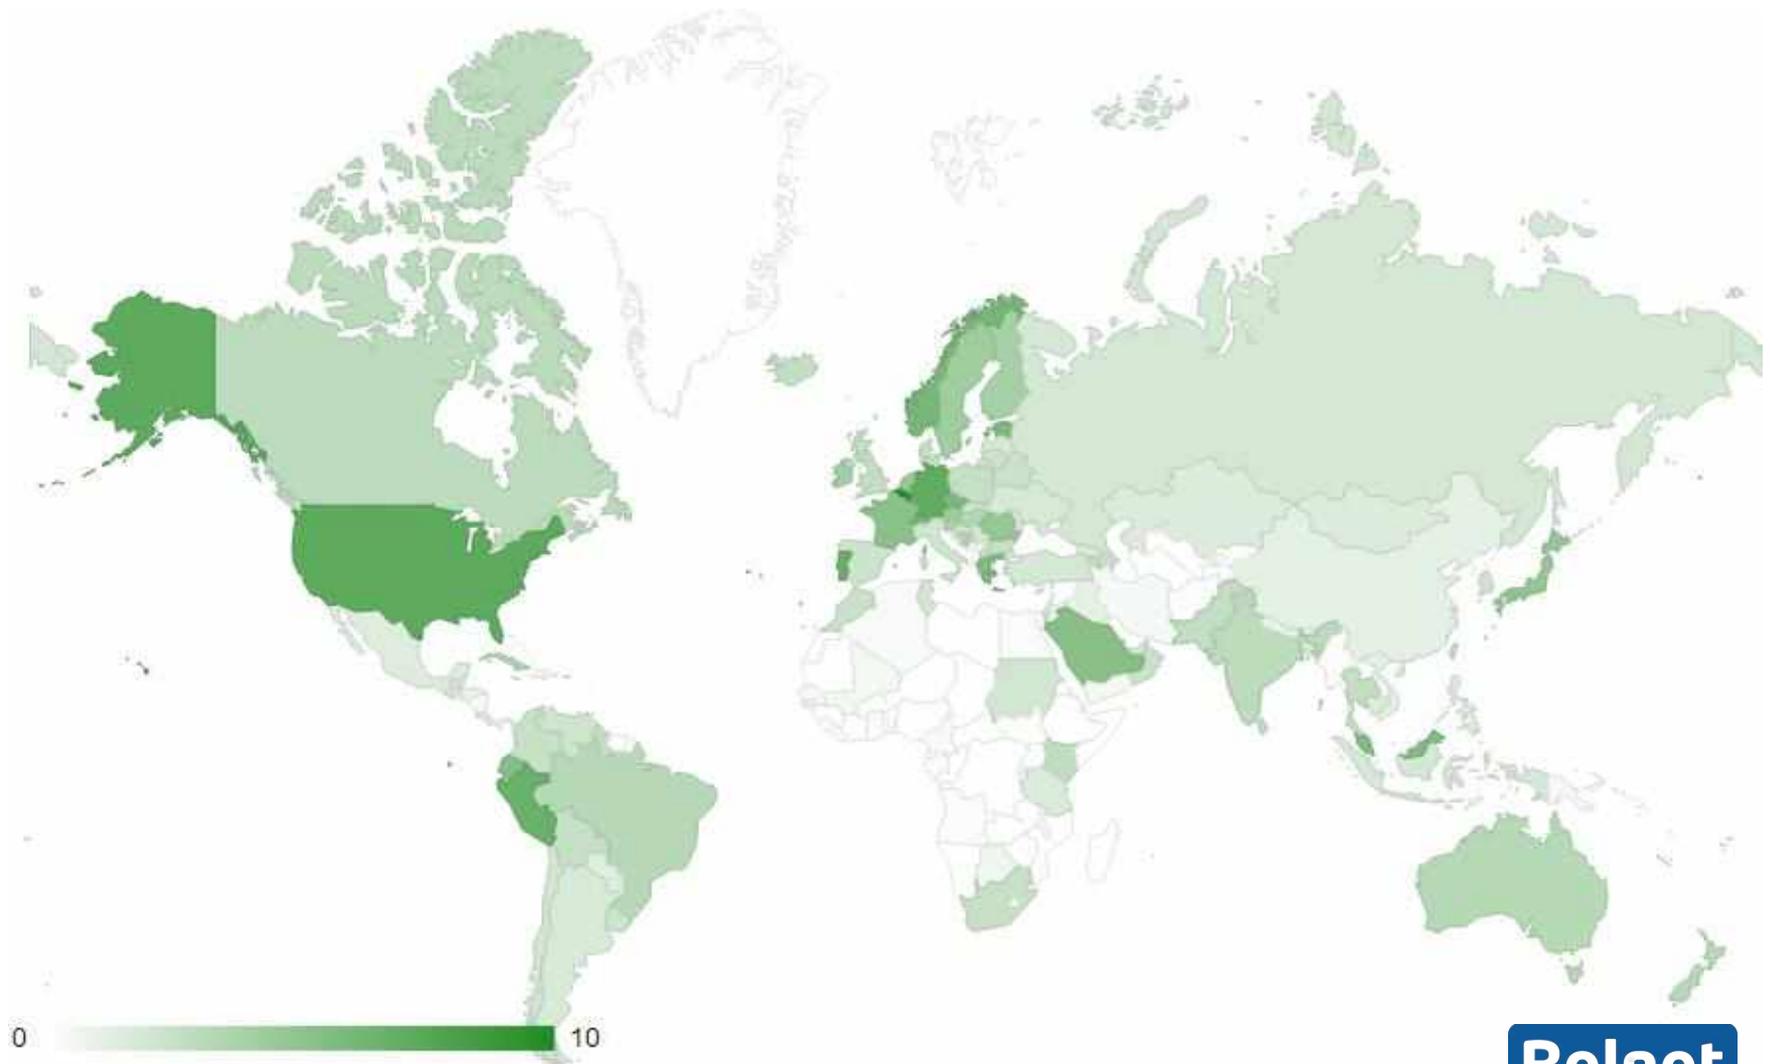
June 2015



Where is Europe today ?



Where is the World today ?



Belgium: Users penetration



Where is Belgium today ?

- **NANOG 64, June 5th 2015 in San Francisco**

- North Atlantic Network Operators Group
- Quoted Belgium !

*“Some other markets have also jumped, including Germany where IPv6 penetration is also at approximately one quarter of their national customer base, and **Belgium where IPv6 is used by almost one half of the users in that country.**”*

<https://blog.apnic.net/2015/06/05/thoughts-from-nanog-64/>



IPv6

Where do YOU want to be today ?

Why you should run IPv6

- **Experimental users**

- Take time to get to know it, be eager to learn now that you have time

- **Power users**

- Start exploiting it, experiment features, speed, applications

- **Global audience**

- Be reachable for the World who already moves on to IPv6
- Some are IPv6 only, and not by choice ! (APNIC area)

- **You have time to plan it properly,
don't wait until you have to go in rush mode !**

→ Get your content available over IPv6

Belnet can help you !

- We can advise and guide you on general approaches
- Help you find the best solution matching your needs
- Direct you to the right resources to do it yourself
- Assist you in gaining experience

IPv6 Dual Stack interoperability

Belnet

AS2611

193.190.0.0/15

2001:6a8::/32

All traffic to
193.190.162.0/24

All traffic to
2001:6a8:24c0::/48

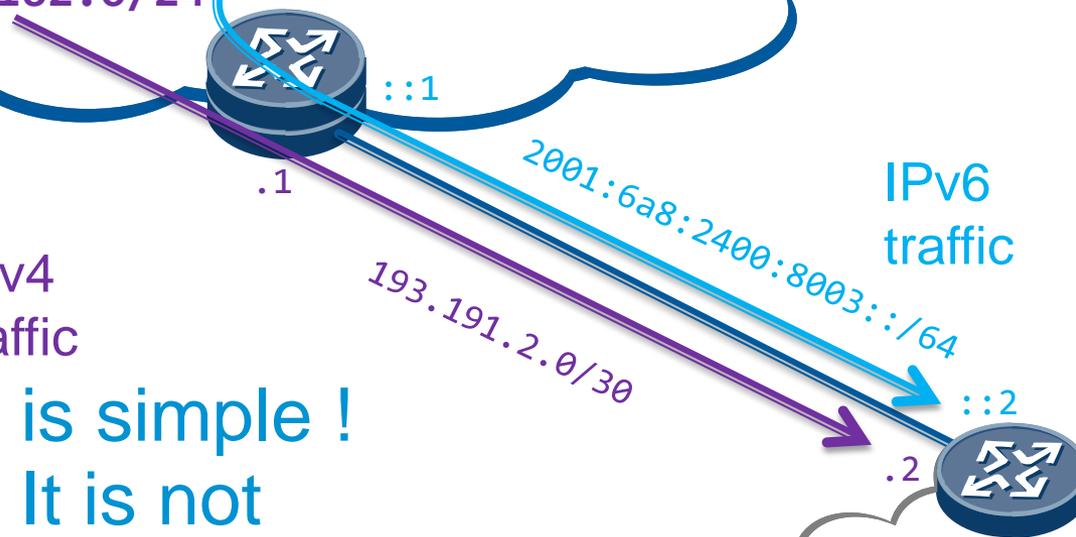
IPv4
traffic

It is simple !

It is not

more complicated

than that !



Customer

Wrap-up: IPv6 @ Belnet

- **Native, dual-stack since Jan 2003 : 2001:6a8::/32**
- **Customers all get /48 IPv6 range**
- **Network with Multiple IPv6 peerings**
 - GÉANT
 - IP Transit
 - BNIX
 - Other IXes, AMS-IX, LINX, SFINX
- **Belnet Services on already IPv6**
 - FTP, DNS, Jabber, NTP, WWW, SMTP, Antispam Pro...

Some Belnet Customers Actively using IPv6

Agence Wallone des Télécommunications	Diffusion Universitaire CIACO	CIRB-CIBG
Haute Ecole Libre Mosane	Hogeschool Westvlaanderen	iMinds
Instituut voor Tropische Geneeskunde	Katholieke Hogeschool Sint-Lieven	Katholieke Hogeschool Leuven
Bureau Fédéral du Plan	Katholieke Universiteit Leuven	NATO NCI
Technocité	Universiteit Antwerpen	RTBF
Universiteit Hasselt	Université de Liège	Université de Mons
Universitair Ziekenhuis Antwerpen	Université de Namur	Maybe you ?
Vlaamse Instelling voor Technologisch Onderzoek	Katholieke Hogeschool VIVES	

Some Belnet Customers Connected but not actively using it yet

Haute Ecole Louvain en Hainaut	Haute Ecole Hainaut
Institut Wallon de l'évaluation, de la prospective et de la statistique	Thomas More
Centrum voor Volwassenenonderwijs De Oranjerie	Provincie West-Vlaanderen
Groep T	KMS-ERM
IBPT-BIPT	Stad Leuven
PXL	Université de Namur

**Don't wait, get on this list too. It's already a start !
Request to activate IPv6 once you go back to the office:**

servicedesk@belnet.be

Your /48 is waiting to come alive !

Organisation	IPv6 prefix
ASTRID	pending
Grondwettelijk Hof	2001:06a8:be00::/48
BELSPO	2001:06a8:b400::/48
PXL	2001:06a8:2080::/48
VRT	2001:06a8:0a90::/48

Organisation	IPv6 prefix
Cenaero	2001:06a8:1480::/48
Province De Liege	2001:06a8:2d00::/48
Waterwegen en Zeekanaal	2001:06a8:0560::/48
Technifutur	2001:06a8:2d20::/48
Université de Namur	2001:06a8:3801::/48

servicedesk@belnet.be

Some IPv6 links with extra info

- **What IPv6 basics you need to know for CCNA**
 - <https://www.youtube.com/watch?v=RkhK-JliNUY>
- **Preparing an IPv6 Address Plan (SURFnet)**
 - <http://www.ipv6forum.com/dl/presentations/IPv6-addressing-plan-howto.pdf>
- **RIPE-554 “Requirements for IPv6 in ICT Equipment**
 - <https://www.ripe.net/publications/docs/ripe-554>
- **IETF RFC 7381 “Enterprise IPv6 Deployment Guidelines**
 - <https://tools.ietf.org/html/rfc7381>
- **IETF RFC 6434 “IPv6 Node Requirements”**
 - <https://tools.ietf.org/html/rfc6434>



Thank You!

Extra Information



IPv6 Address Plan

Why Create an Address Plan ?

- **Mental heath during implementation !**
- **Easier implementation of security policies**
- **Efficient address plans are scalable**
 - Group by Function **vs** Group by Location
- **More efficient route aggregation**

Address Plan

- **Number of hosts is irrelevant**
- **Philosophy change : think subnets, not IP addresses**
- **Work on 4-bit boundaries as much as possible**
- **/64 for all subnets**
 - Auto-configuration (SLAAC) works
 - Easier to renumber
 - Less typo errors because of simplicity
- **Use one /64 block (per site) for loopbacks**
 - One /128 per device

Address Plan

- For private networks, look at ULA
- Manual configuration for servers & routers
- Point-to-Point ?
 - RFC 4291: Interface IDs are required to be /64
 - [RFC 6164](#): Using /127 on Inter-Router links
 - Avoids ping-pong issue of forwarding loops where NDP not possible (eg SDHG/SONET)
 - But [RFC 4443](#) mitigates this (if implemented)
 - Avoids Neighbour Cache Exhaustion issue
 - But do reserve /64

Neighbour Cache Exhaustion (Optional)

- Use of a 64-bit prefix on inter-router link that used ND (eg Ethernet) potentially allows for DoS attacks on routers on the link
- Packet destined for a non-existing address on the point-to-point will bounce between the two routers
 - Consider an Ethernet link between two routers, A and B, to which a /64 subnet has been assigned. A packet sent to any address on the /64 (except the addresses of A and B) will cause the router attempting to forward it to create a new cache entry in INCOMPLETE state, send a Neighbor Solicitation message on the link, start a retransmit timer, and so on...
 - See RFC 6164 § 5.2
- NDP cache tables smaller than ARP tables (today)
- Sometimes NDP and ARP tables share same memory !



RFC 6724
Default Address Selection for IPv6
Optional Section

IPv6 Default Address Selection

- **Many IPv6 addresses per node: choose which ?**
 - Note: Microsoft uses the term **Source Address Selection**
 - Problem exist for source AND destination addresses
 - RFC 6724 defines algorithm for selecting best address for an interface
 - Source address selection for a given destination
 - Destination address selection to sort possible destinations

IPv6 Default Address Selection

- **From the RFC 6724:**

“If an implementation is not configurable or has not been configured, then it SHOULD operate according to the algorithms specified here in conjunction with the following default policy table:”

Prefix	Precedence	Label
::1/128	50	0
::/0	40	1
2002::/16	30	2
::/96	10	3
::ffff:0:0/96	0	4

IPv6 Default Address Selection

- **In English**

- Address pairs of same scope or type are preferred
- Use the smallest scope possible for the Destination address
- A non-deprecated address is preferred
- ISATAP or 6to4 addresses are not used if native IPv6 is available
- If all criteria are similar, address pairs with the longest common prefix are preferred
- For source address, global addresses are preferred over temporary
- In Mobile IP situations, home addresses are preferred over care-of

Lunch time

