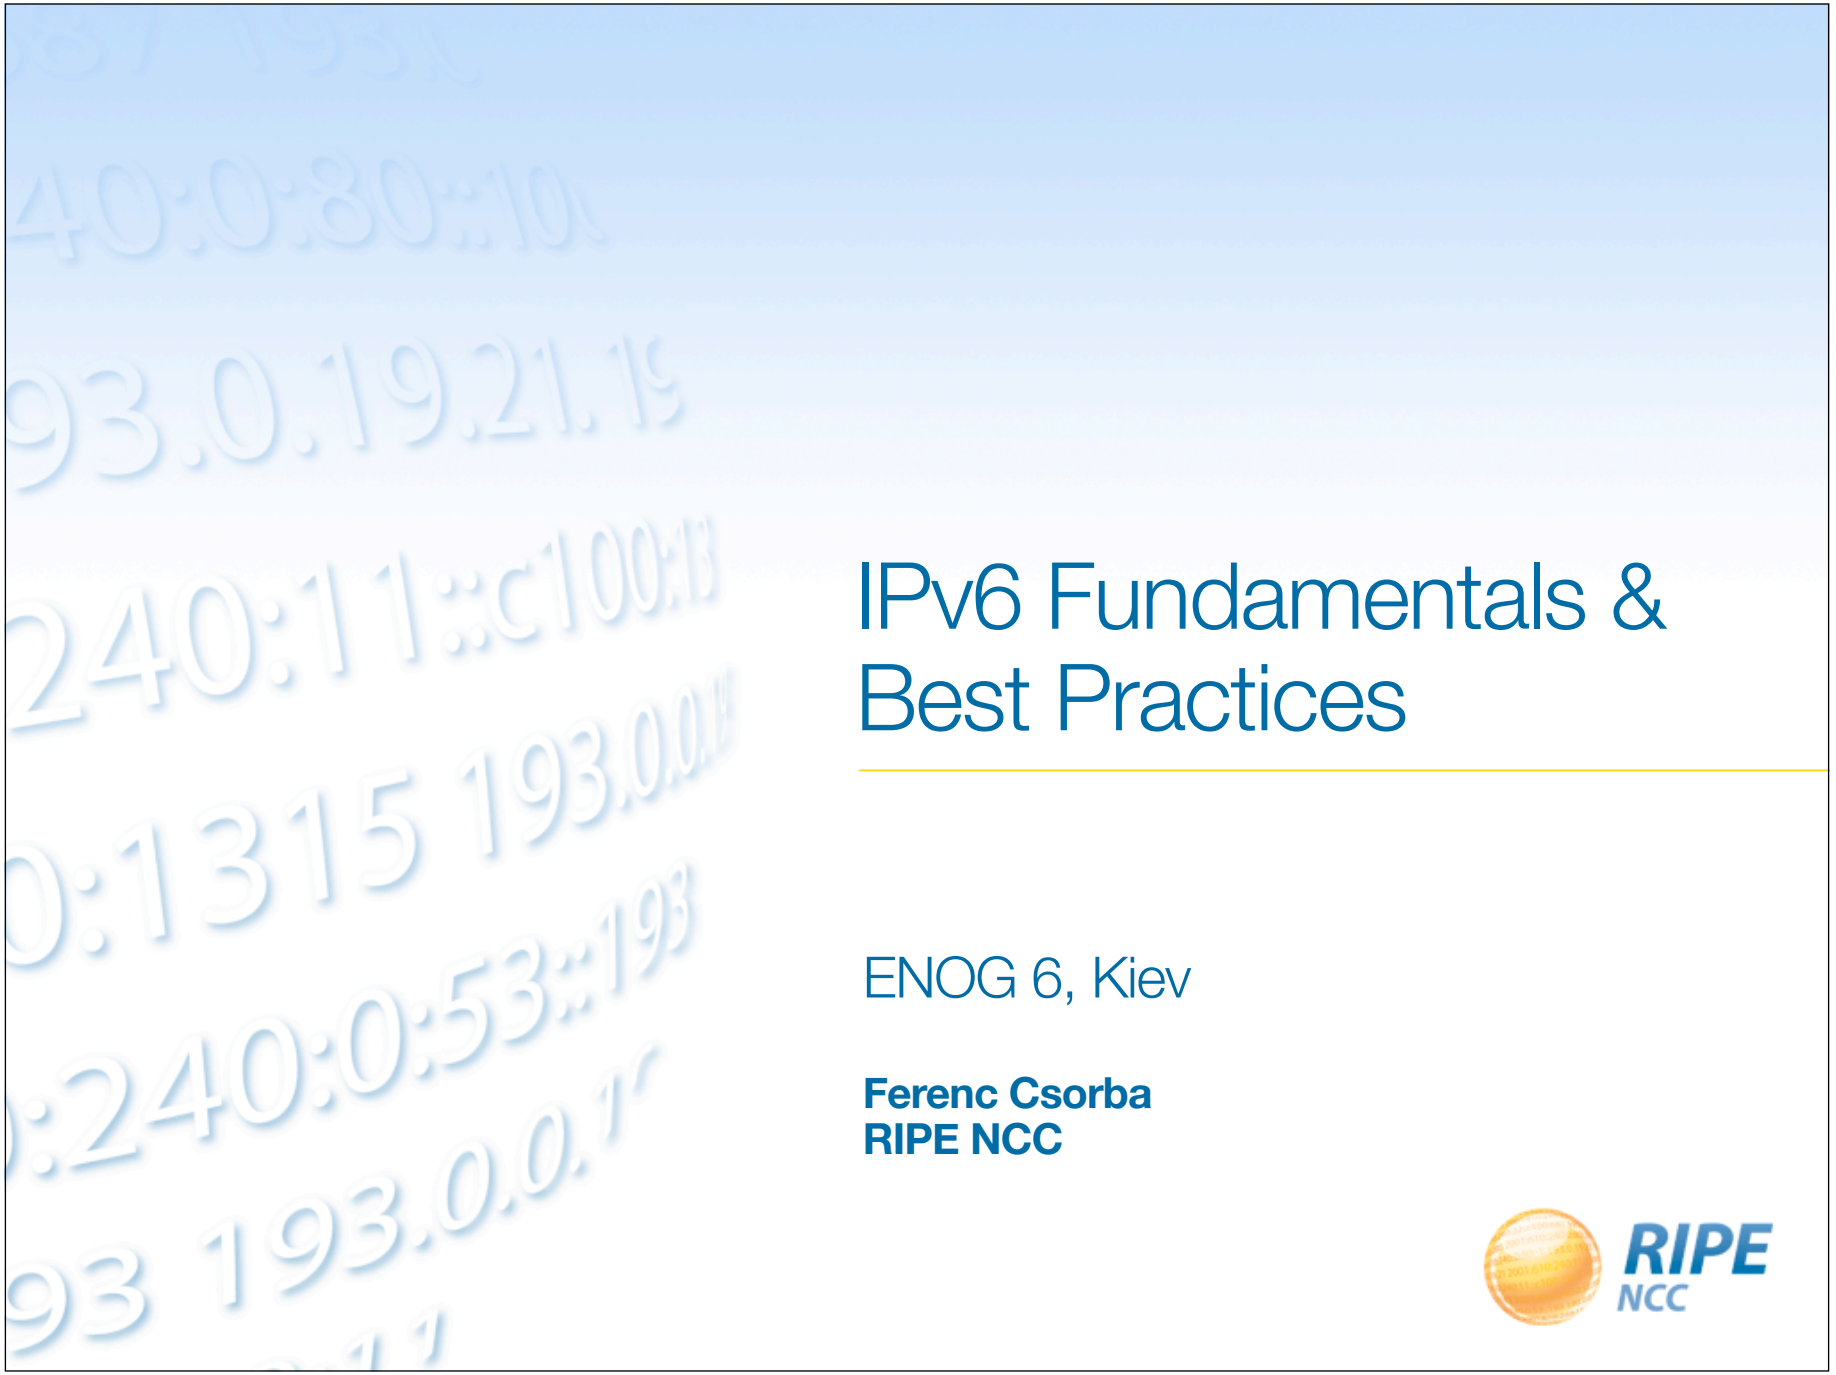

IPv6 Fundamentals & Best Practices

ENOG 6, Kiev

Ferenc Csorba
RIPE NCC

RIPE NCC

Schedule

- RIPE NCC: Who are we?
- IPv4 exhaustion
- IPv6 address space
- IPv6 for mobile telephony
- Tips and hints

Who are we?

RIPE NCC

- Located in Amsterdam
- Not for profit membership organisation
- One of five RIRs
- **Distribute IP addresses, ASNs** etc

RIPE

- Open community
- Develops addressing policies
- Working group mailing lists

The five RIRs

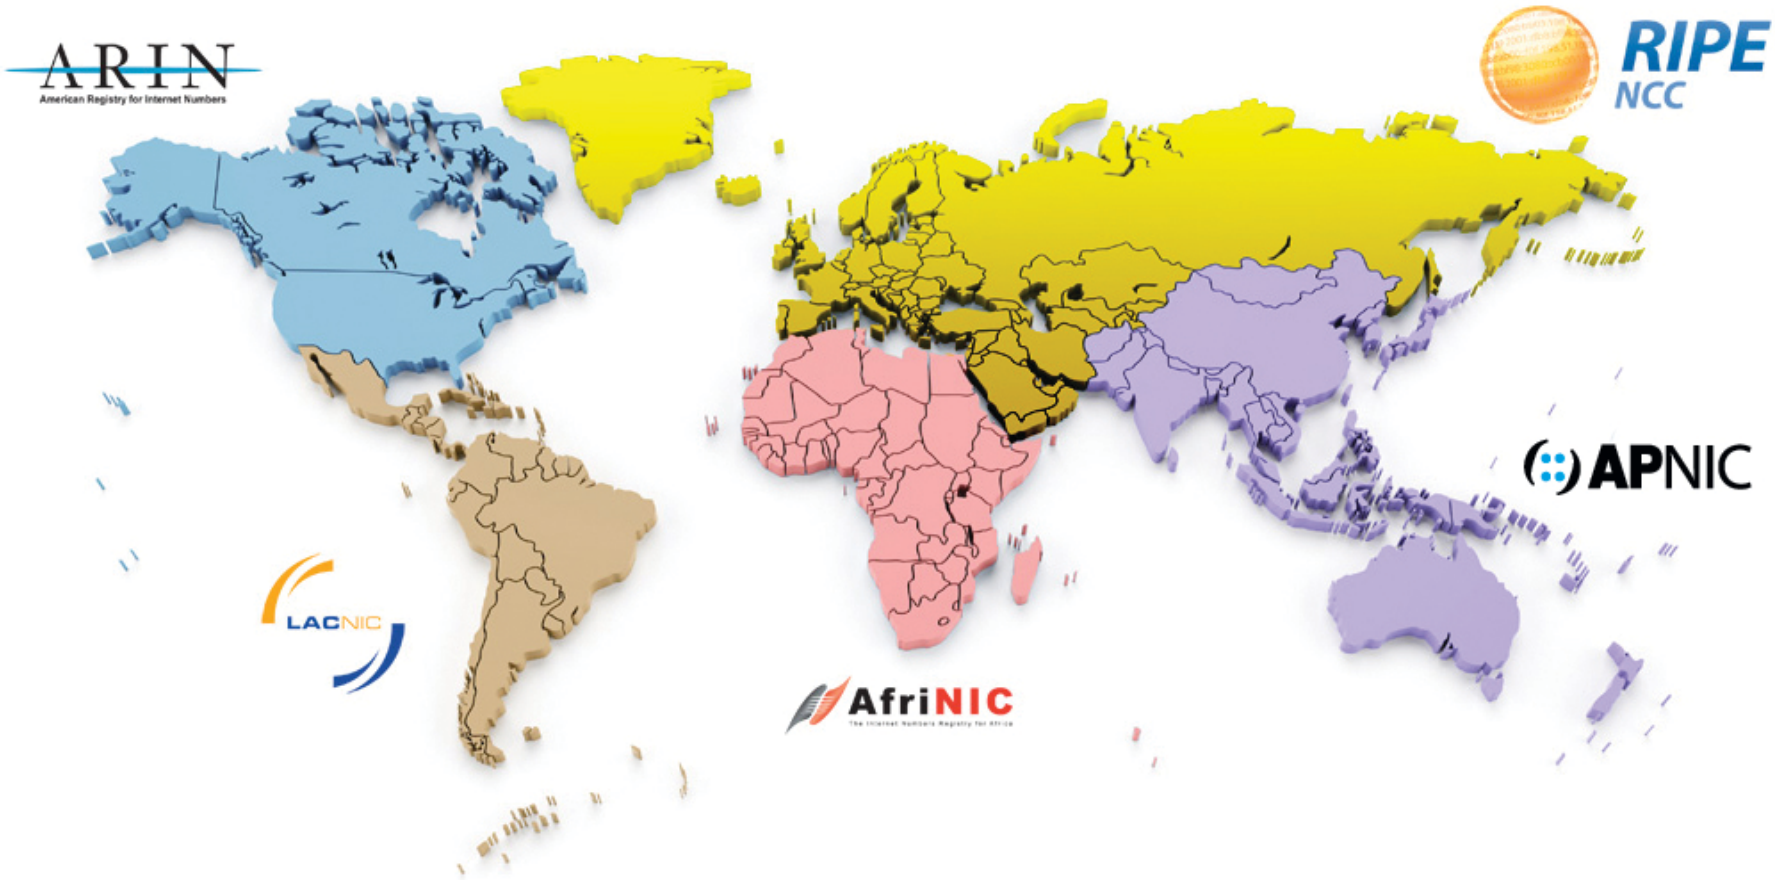

IPv4 Address Pool Exhaustion

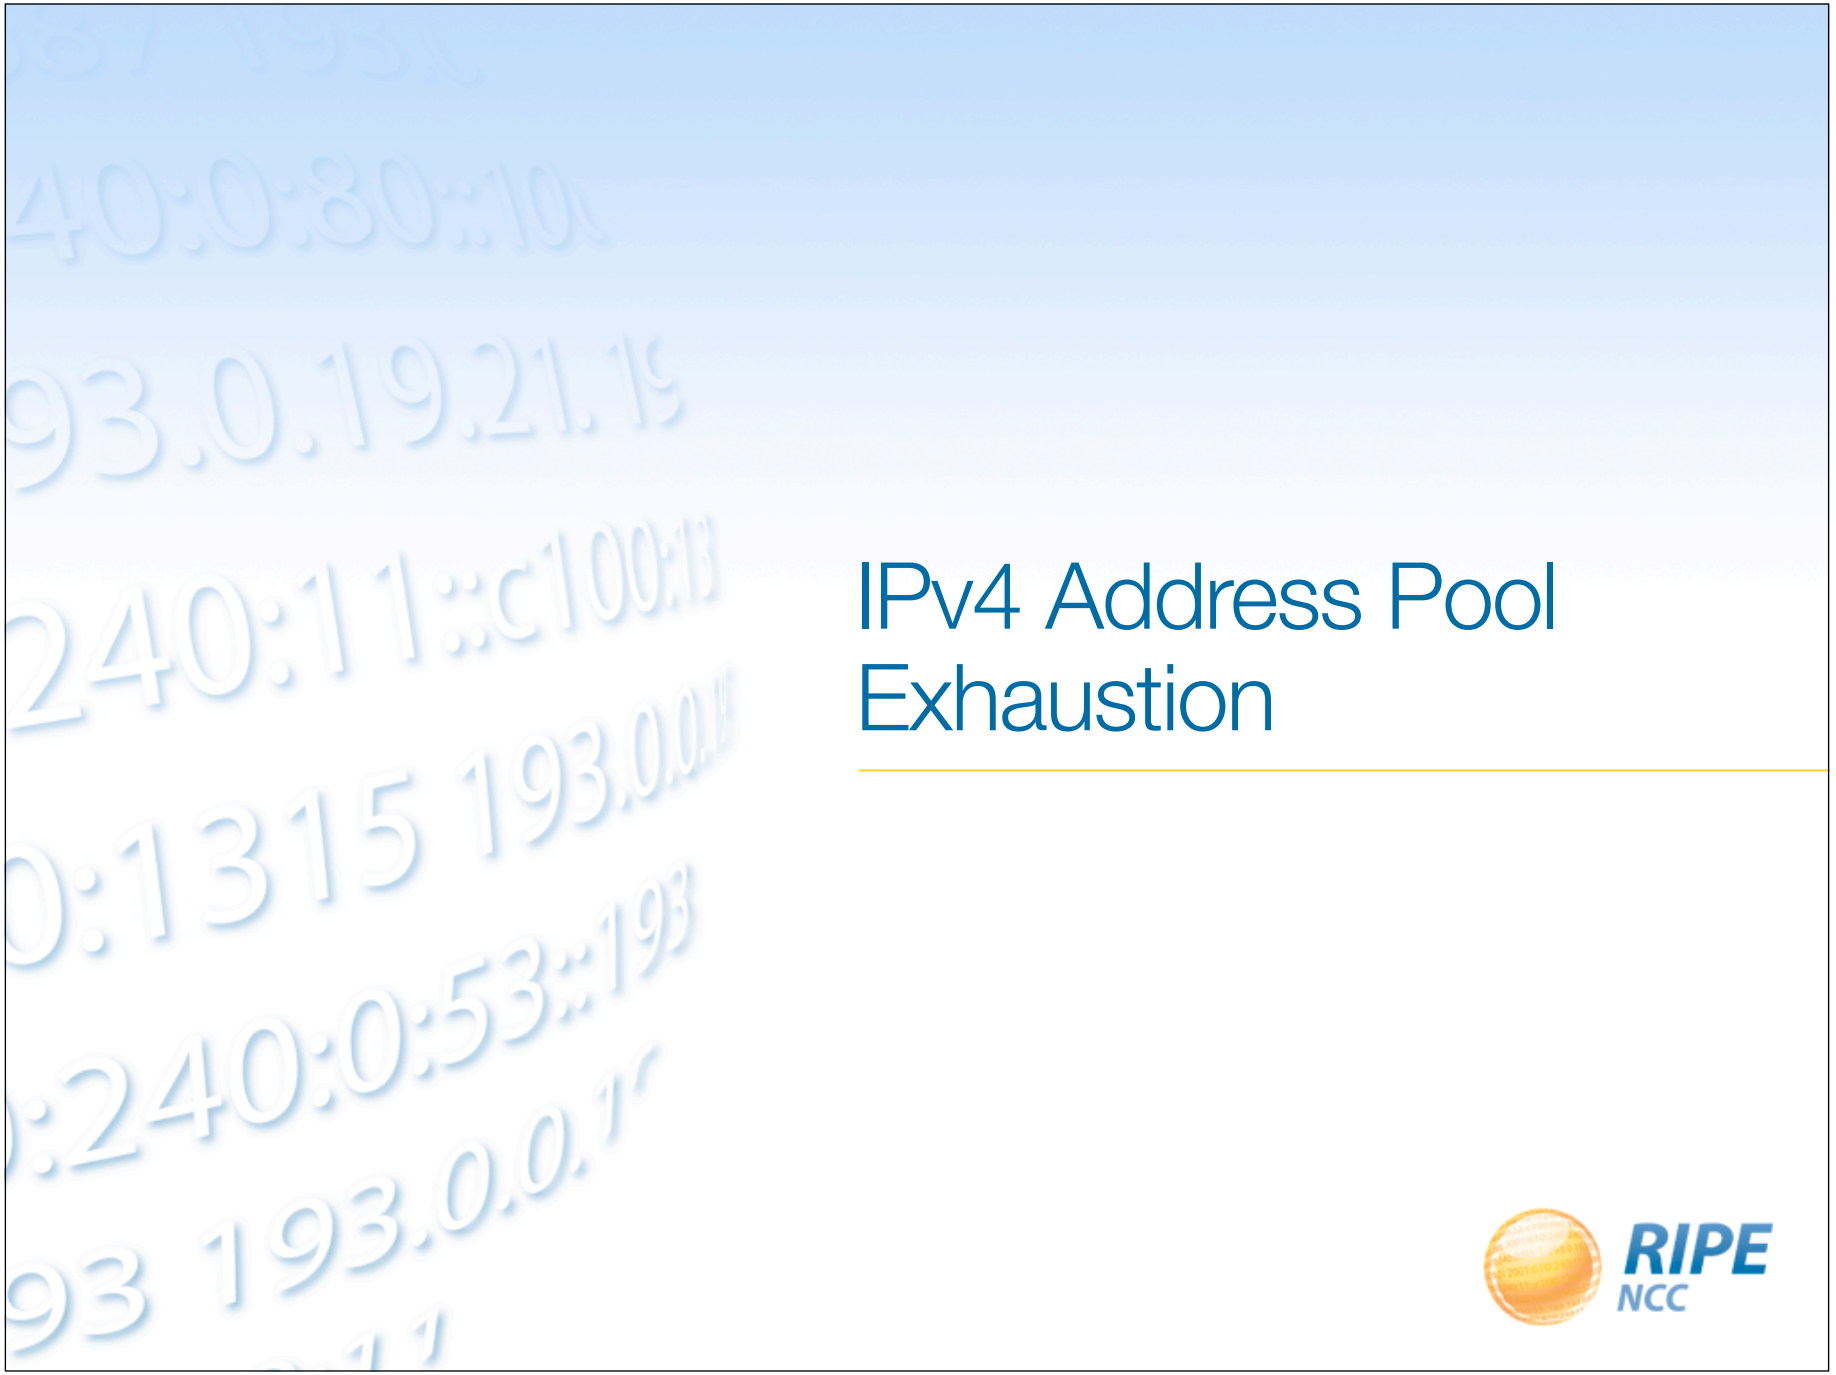

IANA IPv4 Pool

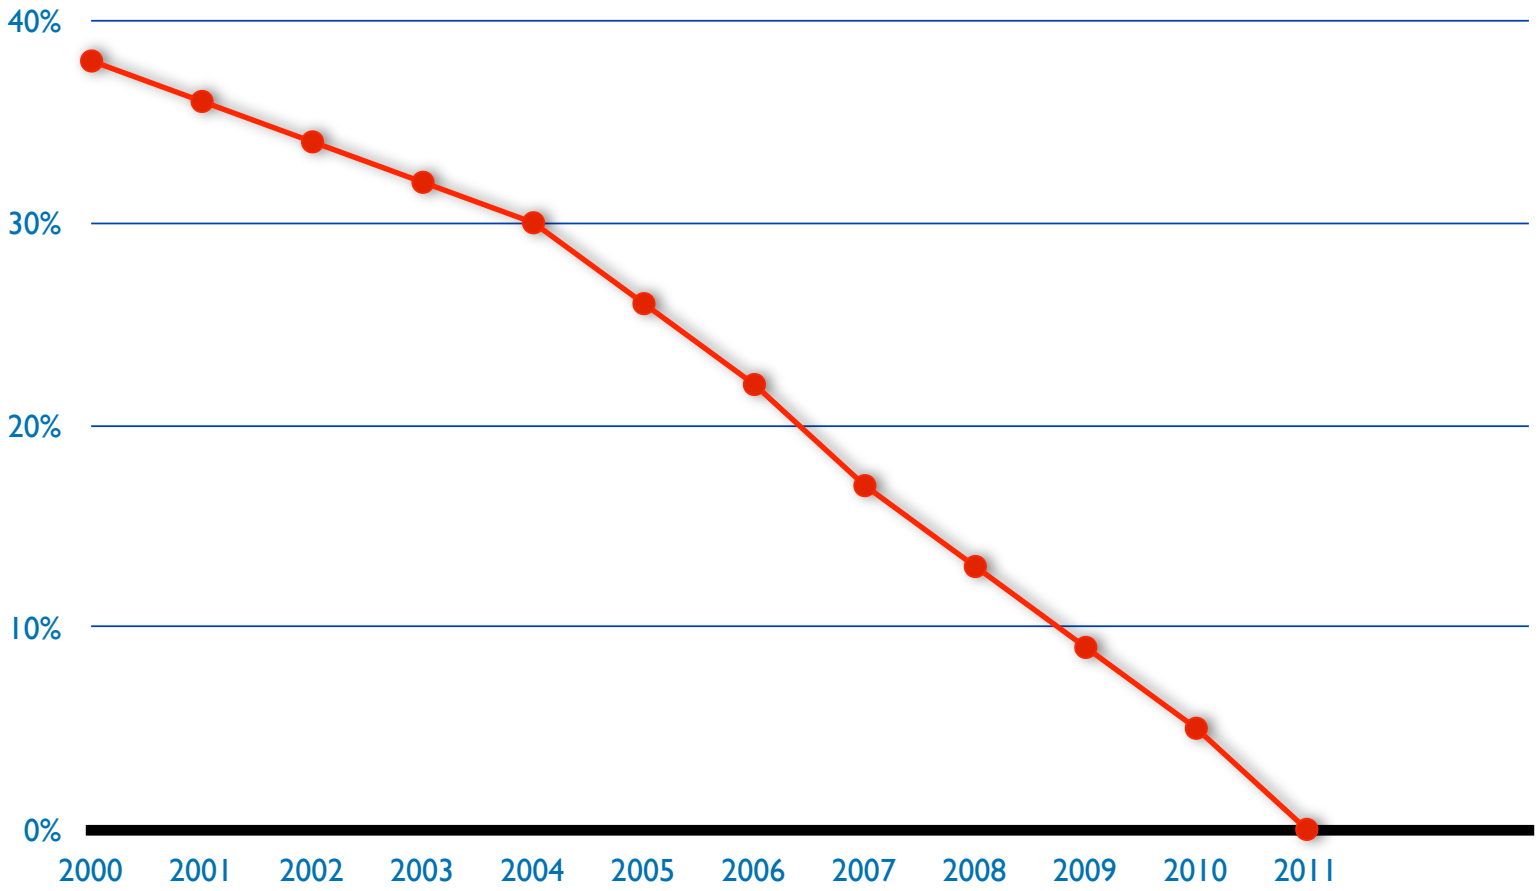

IPv4 Address Distribution

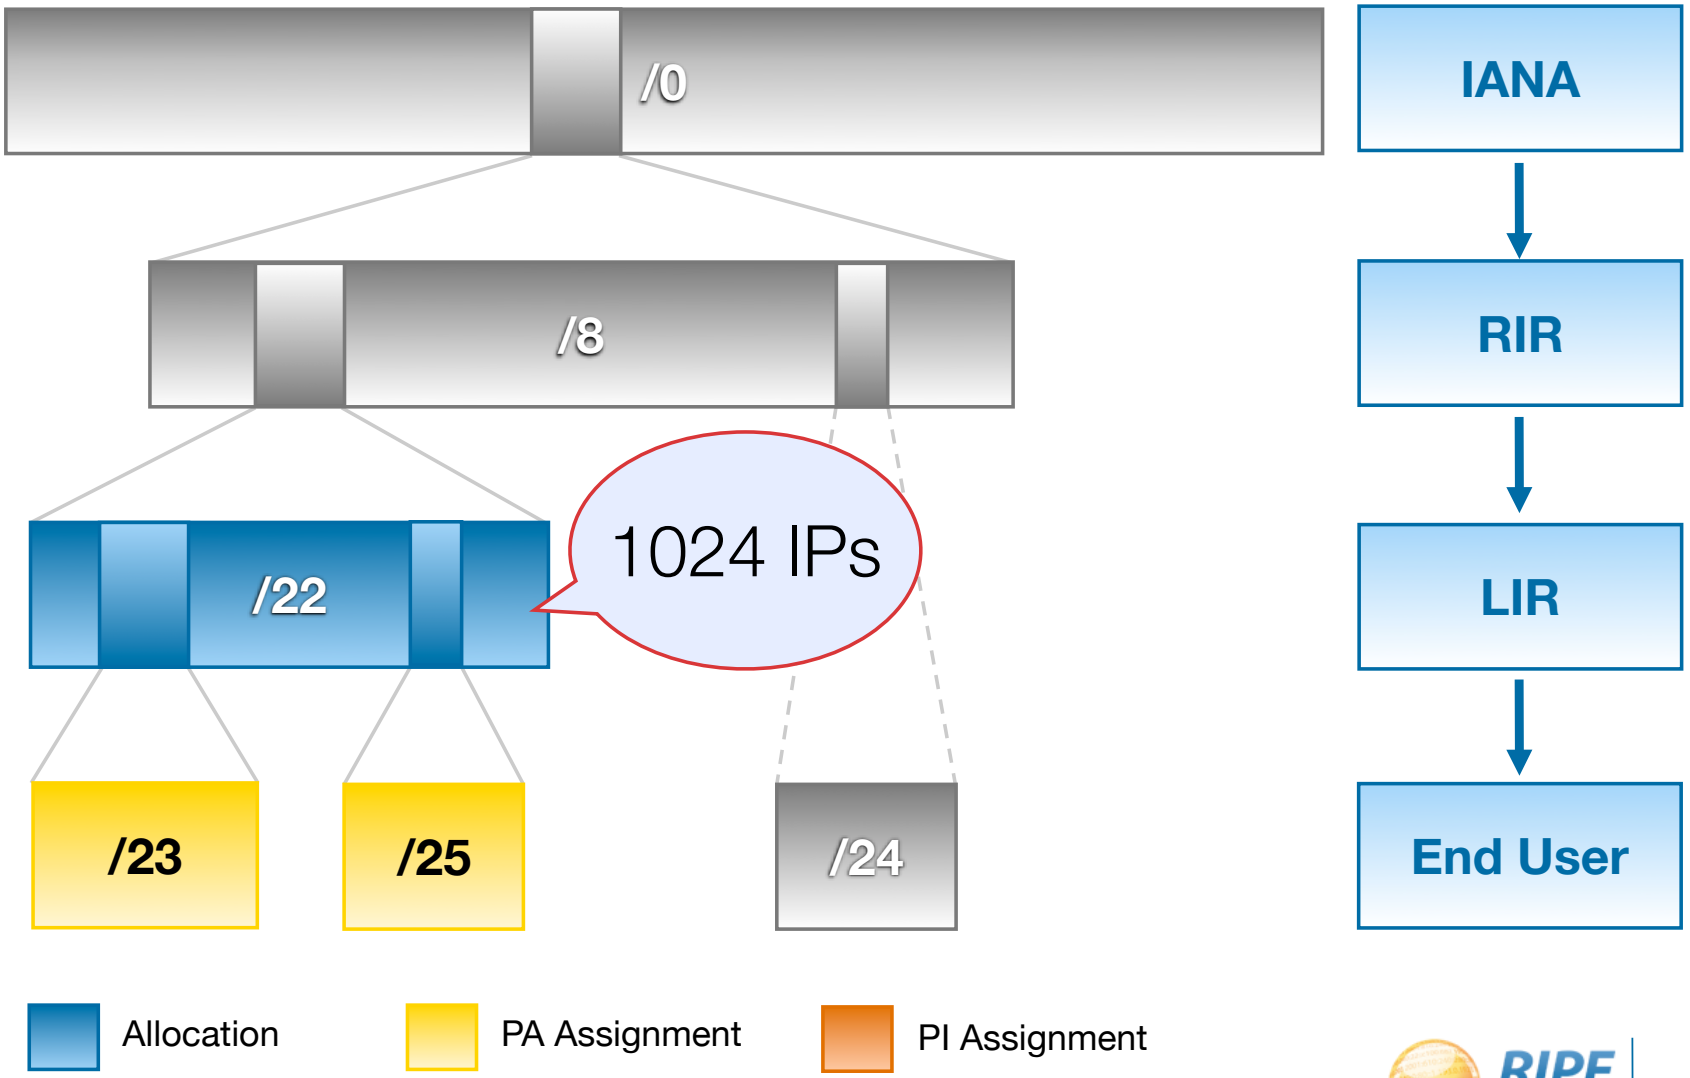

RIPE NCC's last /8

- Each LIR gets **one /22** (=1024 addresses)
- No PI

2187 /22 have been allocated since September 2012

IPv6 Address Space

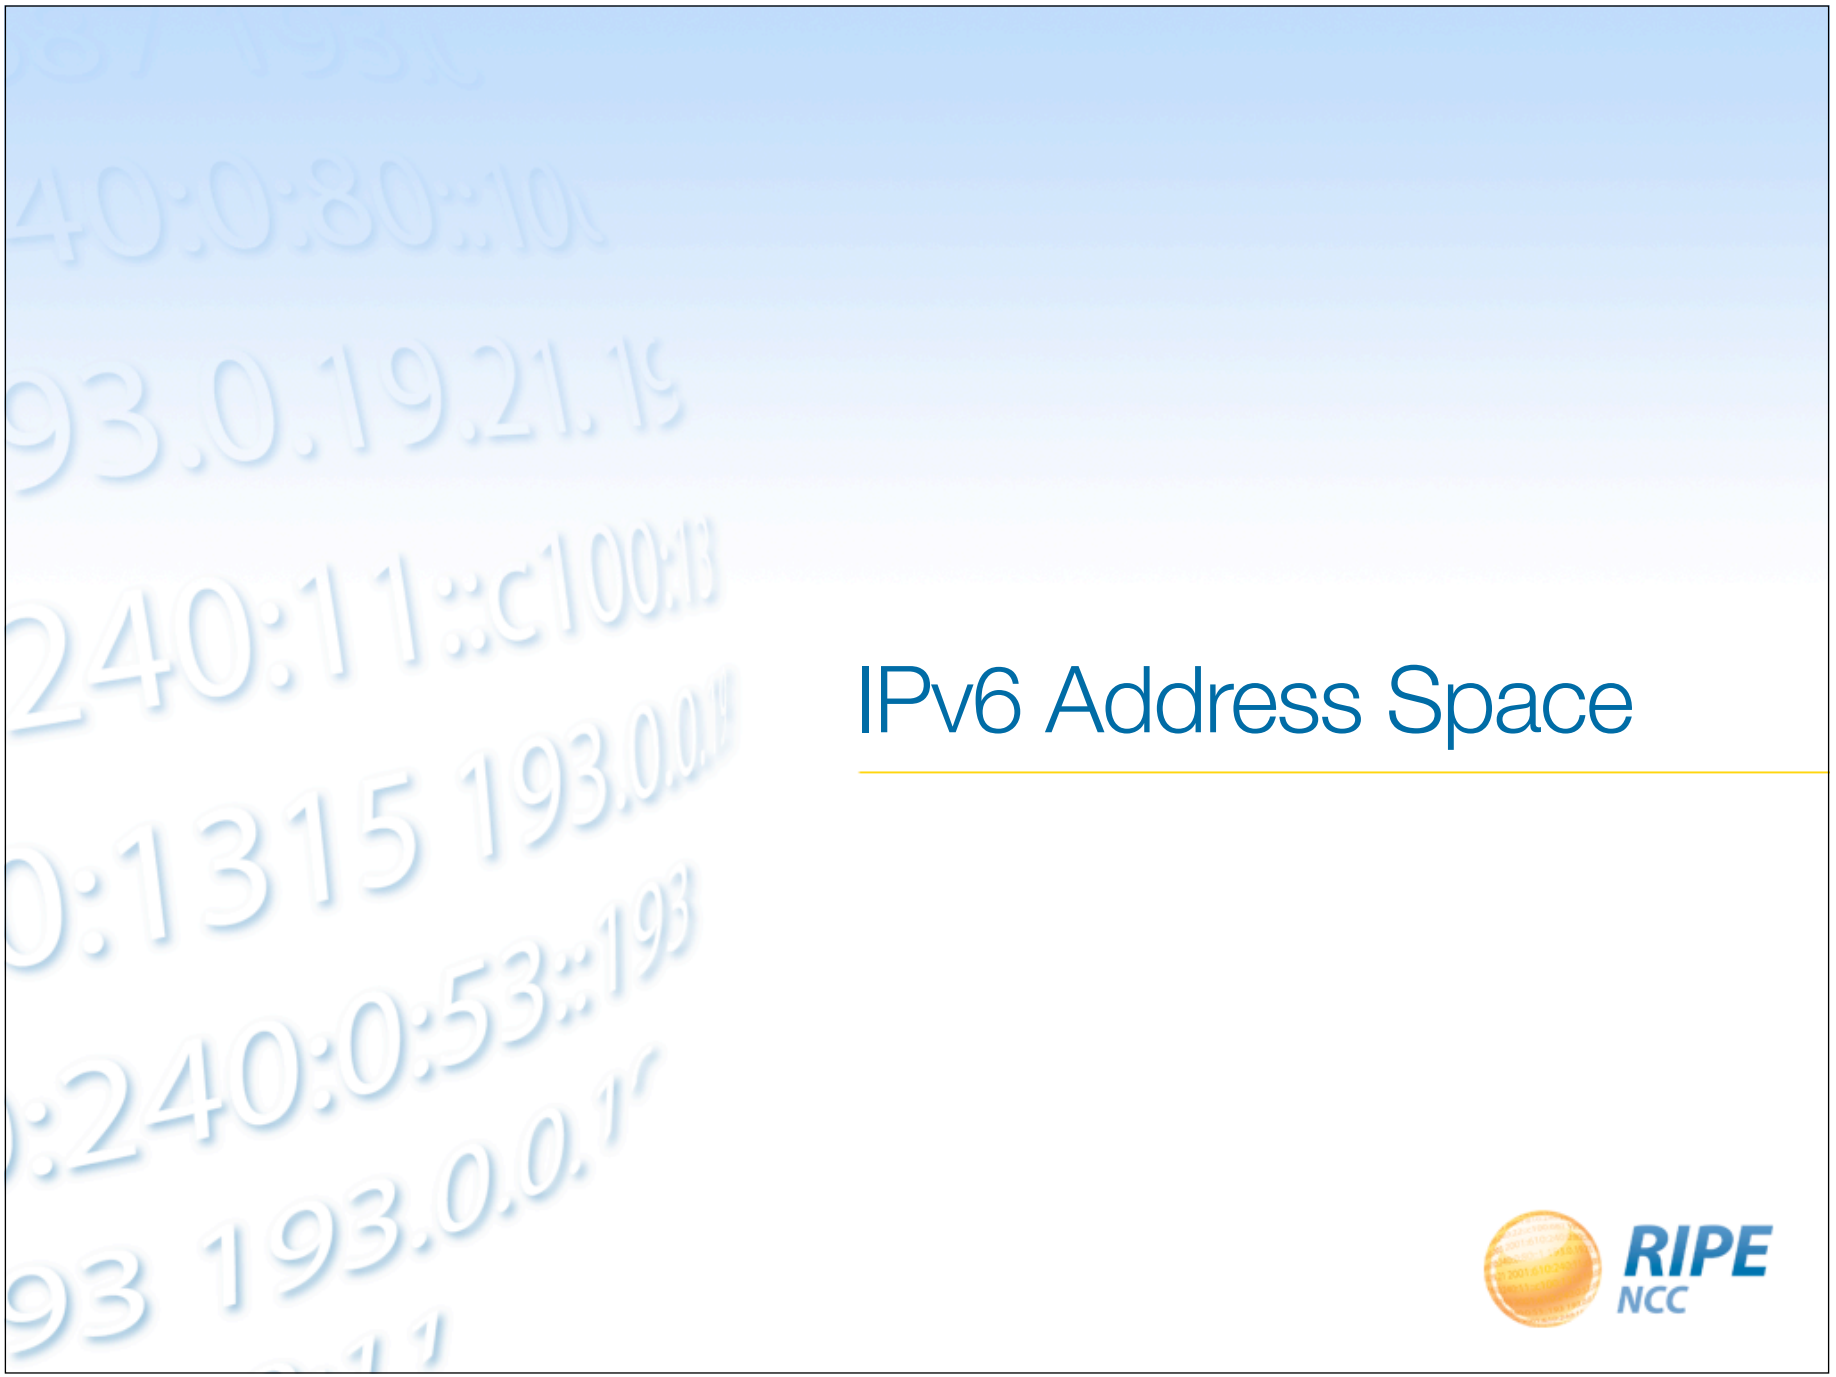

IPv6 Basics

- 128 bits in IPv6
- 32 bits in IPv4

Address Notation

2001:0db8:003e:ef11:0000:0000:c100:004d

2001:0db8:003e:ef11:0000:0000:c100:004d

2001:db8:3e:ef11:0:0:c100:4d

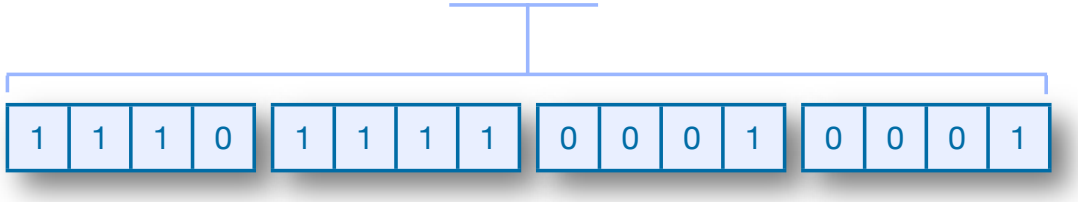

Number of addresses (rounded off)

- IPv4
  - 4000000000
- IPv6
  - 300000000000000000000000000000000

IPv6 Address Distribution

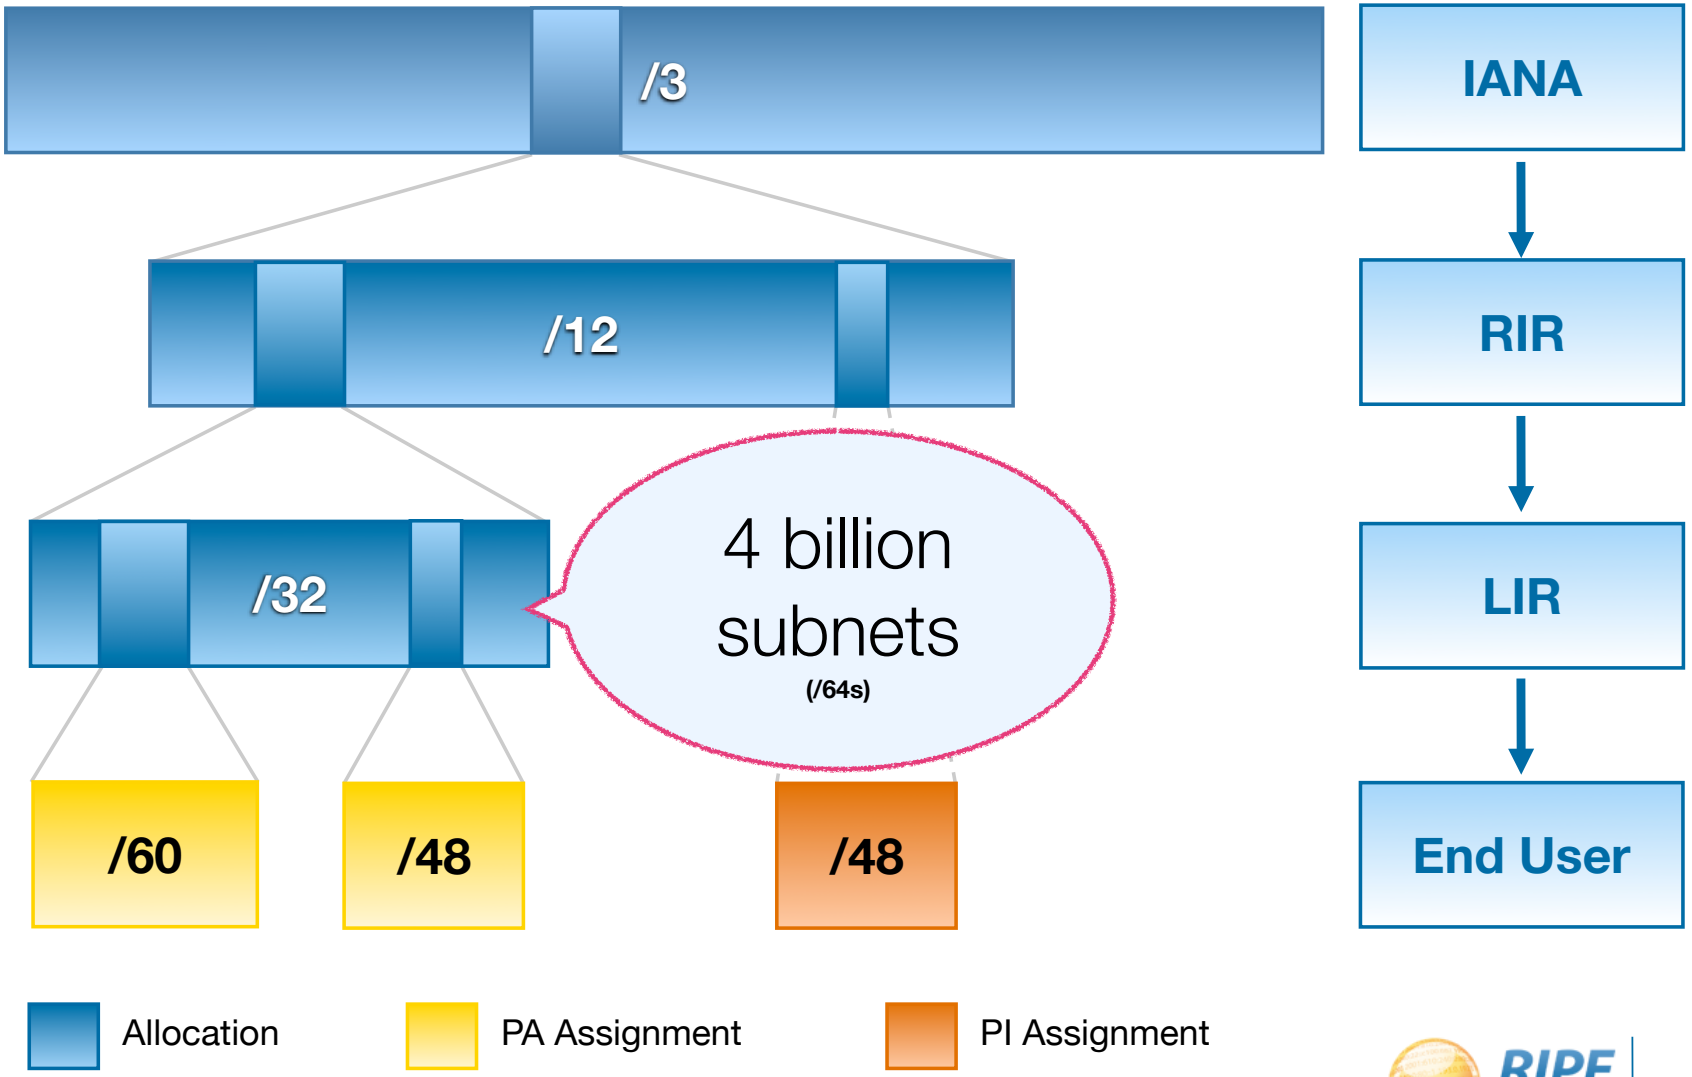

IPv6 Ripeness

- Rating system:
  - One star if the LIR has an IPv6 allocation
  - Additional stars if:
    - IPv6 Prefix is announced on router
    - A route6 object is in the RIPE Database
    - Reverse DNS is set up
  - A list of all 4 star LIRs: http://ripeness.ripe.net/

2013 IPv6 RIPEness: **ALL** 9579 LIRs

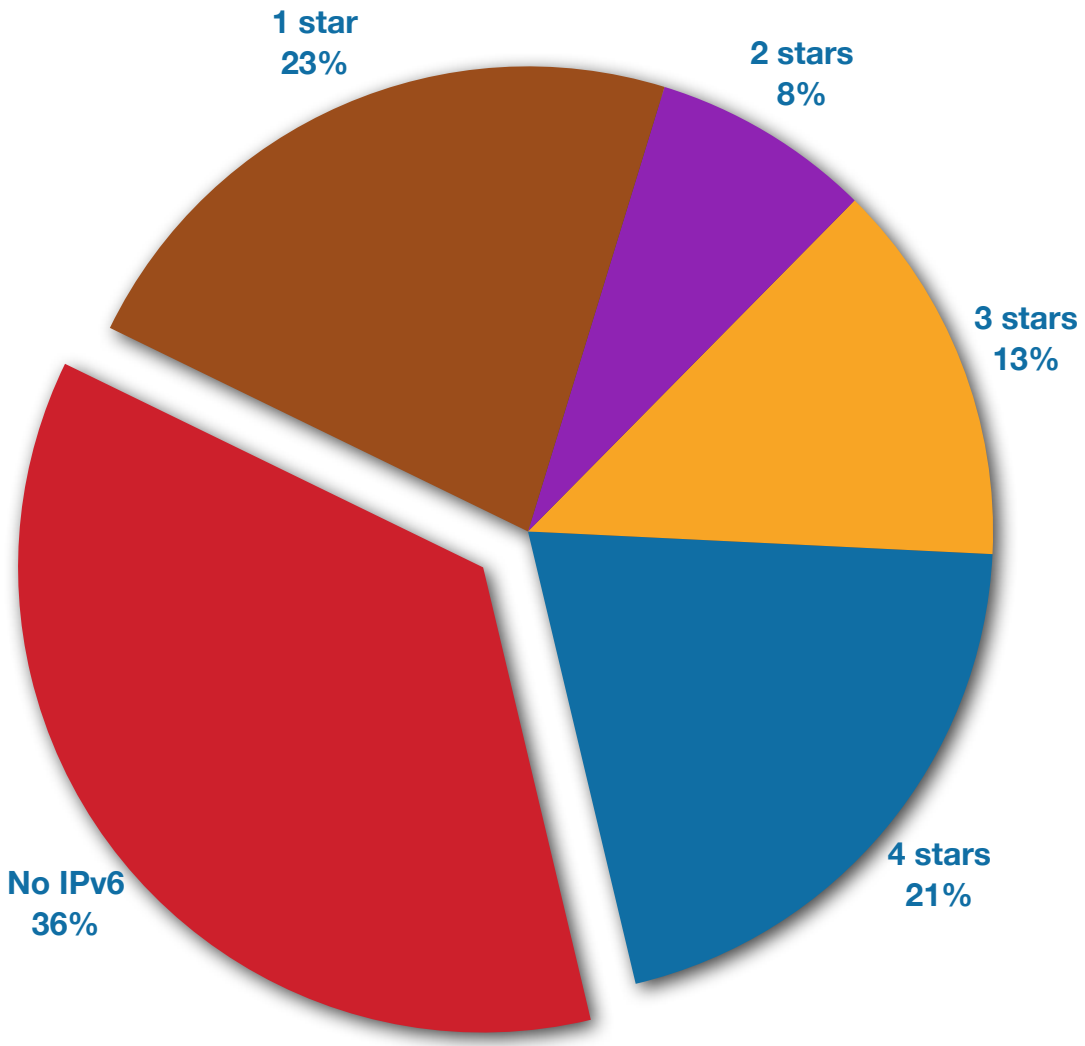

IPv6 RIPEness in the region

Russia(1101)

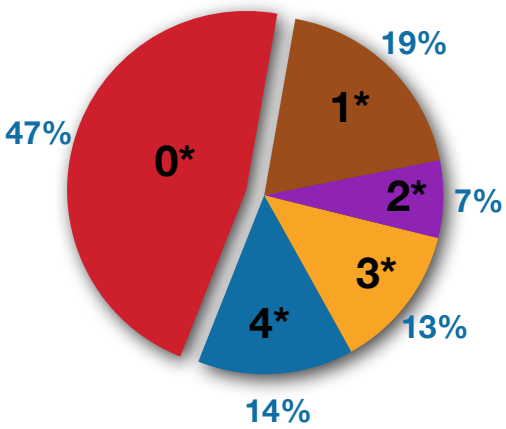

Ukraine(168)

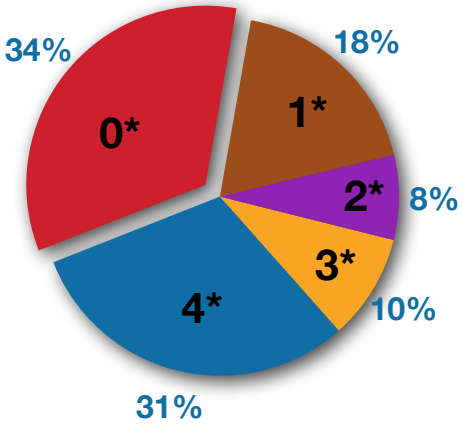

Belarus(17)

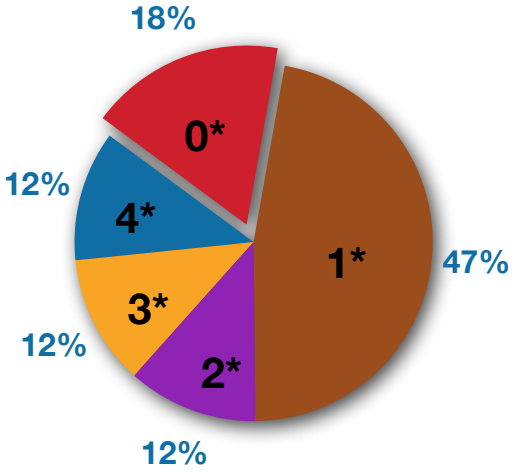

Kazakhstan(50)

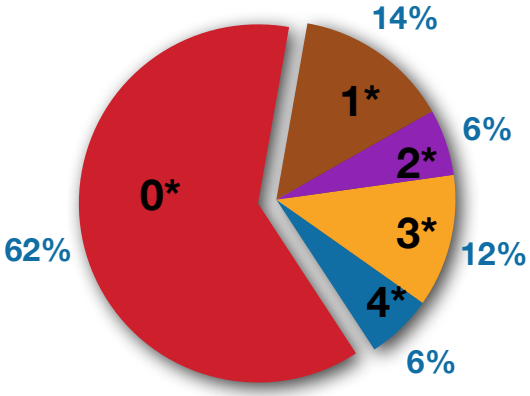

Uzbekistan(18)

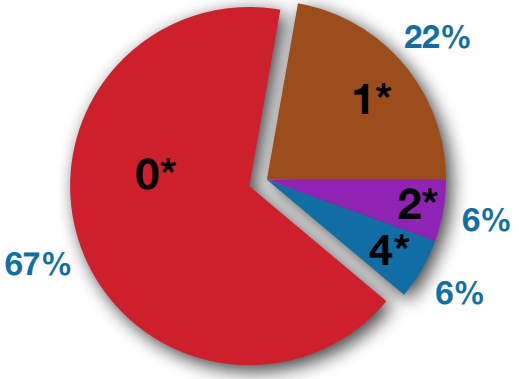

How to get an IPv6 allocation

- To qualify, an organisation mujst:
  - Be an LIR
  - Have a plan for making assignments within two years
- Minimum allocation size /32
- Up to a /29
- Announcement as a single prefix recommended

Now what?

Philosophy change

RIPE NCC

IPv4 -> IPv6 : What philosophy change?

~~How many IP addresses do I need?~~

How many subnets do I need?

Subnet always = /64

IPv4 -> IPv6 : Как изменилься подход ?

~~Сколько адресов мне нужно ?~~

Сколько подсетей мне нужно ?

Каждая подсеть = /64

Assignments to customers

- How many subnets do I give my customers?
  - /64 (1 subnet)
  - /60 (16 subnets)
  - **/56 (256 subnets)**
  - /52 (4096 subnets)
  - **/48 (65536 subnets)**

Default Allocation size = /32

- How many assignments can I make ?
  - 4 billion /64's
  - 268 million /60's
  - **17 million /56's**
  - 1million /52's
  - **65536 /48's**

Classless Inter-Domain Routing (CIDR)

IPv6 Chart

RIPE NCC

Prefix	/48s	/56s	/64s	Bits
/24	16M	4G	1T	104
/25	8M	2G	512G	103
/26	4M	1G	256G	102
/27	2M	512M	128G	101
/28	1M	256M	64G	100
/29	512K	128M	32G	99
/30	256K	64M	16G	98
/31	128K	32M	8G	97
/32	64K	16M	4G	96
/33	32K	8M	2G	95
/34	16K	4M	1G	94
/35	8K	2M	512M	93
/36	4K	1M	256M	92
/37	2K	512K	128M	91
/38	1K	256K	64M	90
/39	512	128K	32M	89
/40	256	64K	16M	88
/41	128	32K	8M	87
/42	64	16K	4M	86
/43	32	8K	2M	85
/44	16	4K	1M	84
/45	8	2K	512K	83
/46	4	1K	256K	82
/47	2	512	128K	81
/48	1	256	64K	80
/49		128	32K	79
/50		64	16K	78
/51		32	8K	77
/52		16	4K	76
/53		8	2K	75
/54		4	1K	74
/55		2	512	73
/56		1	256	72
/57			128	71
/58			64	70
/59			32	69
/60			16	68
/61			8	67
/62			4	66
/63			2	65
/64			1	64

K = 1,024 • M = 1,048,576 • G = 1,073,741,824 • T = 1,099,511,627,776

IPv4 CIDR Chart — **RIPE NCC**

IP Addresses	Bits	Prefix	Subnet Mask
1	0	/32	255.255.255.255
2	1	/31	255.255.255.254
4	2	/30	255.255.255.252
8	3	/29	255.255.255.248
16	4	/28	255.255.255.240
32	5	/27	255.255.255.224
64	6	/26	255.255.255.192
128	7	/25	255.255.255.128
256	8	/24	255.255.255.0
512	9	/23	255.255.254.0
1 K	10	/22	255.255.252.0
2 K	11	/21	255.255.248.0
4 K	12	/20	255.255.240.0
8 K	13	/19	255.255.224.0
16 K	14	/18	255.255.192.0
32 K	15	/17	255.255.128.0
64 K	16	/16	255.255.0.0
128 K	17	/15	255.254.0.0
256 K	18	/14	255.252.0.0
512 K	19	/13	255.248.0.0
1 M	20	/12	255.240.0.0
2 M	21	/11	255.224.0.0
4 M	22	/10	255.192.0.0
8 M	23	/9	255.128.0.0
16 M	24	/8	255.0.0.0
32 M	25	/7	254.0.0.0
64 M	26	/6	252.0.0.0
128 M	27	/5	248.0.0.0
256 M	28	/4	240.0.0.0
512 M	29	/3	224.0.0.0
1024 M	30	/2	192.0.0.0
2048 M	31	/1	128.0.0.0
4096 M	32	/0	0.0.0.0

K = 1,024 • M = 1,048,576

Contact Registration Services:
hostmaster@ripe.net • lir-help@ripe.net

www.ripe.net

Why use only multiples of 4?

If /x is a multiple of 4

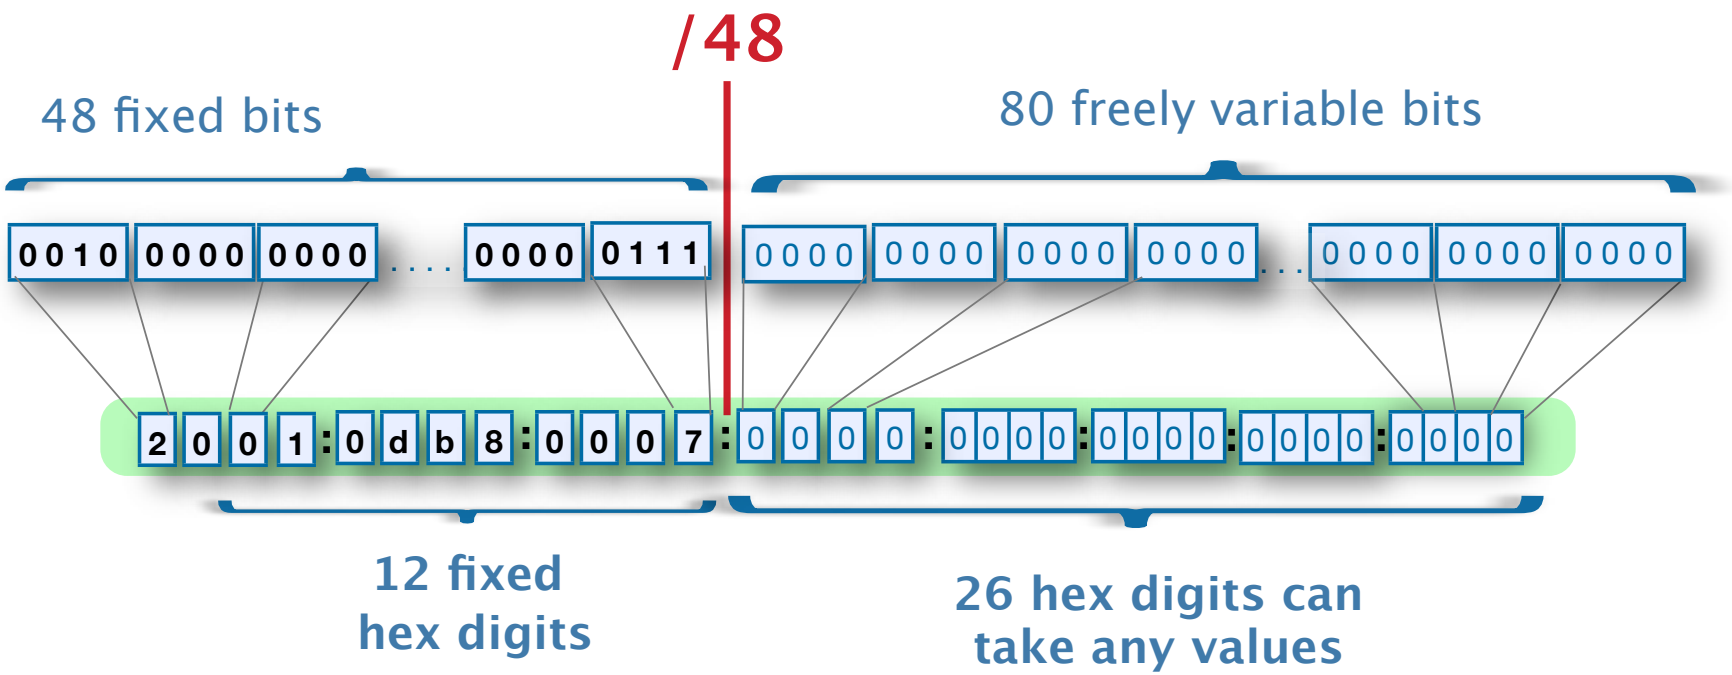

If /x is NOT a multiple of 4

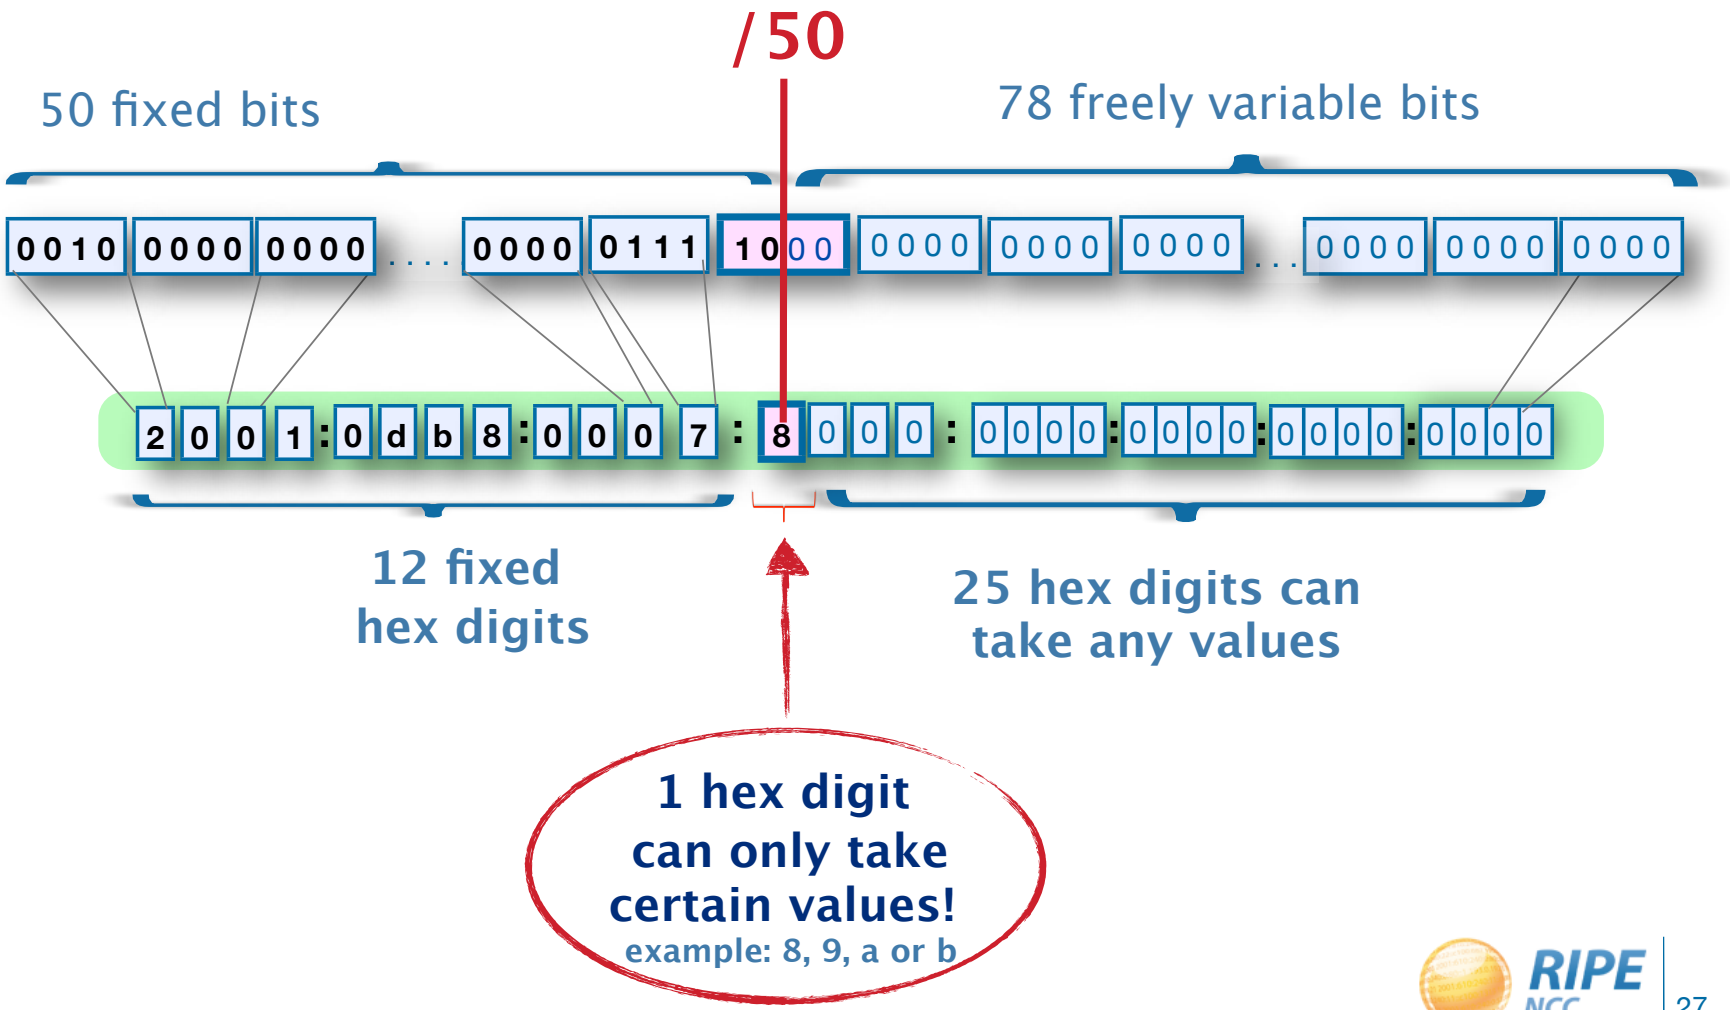

Only certain hex values possible

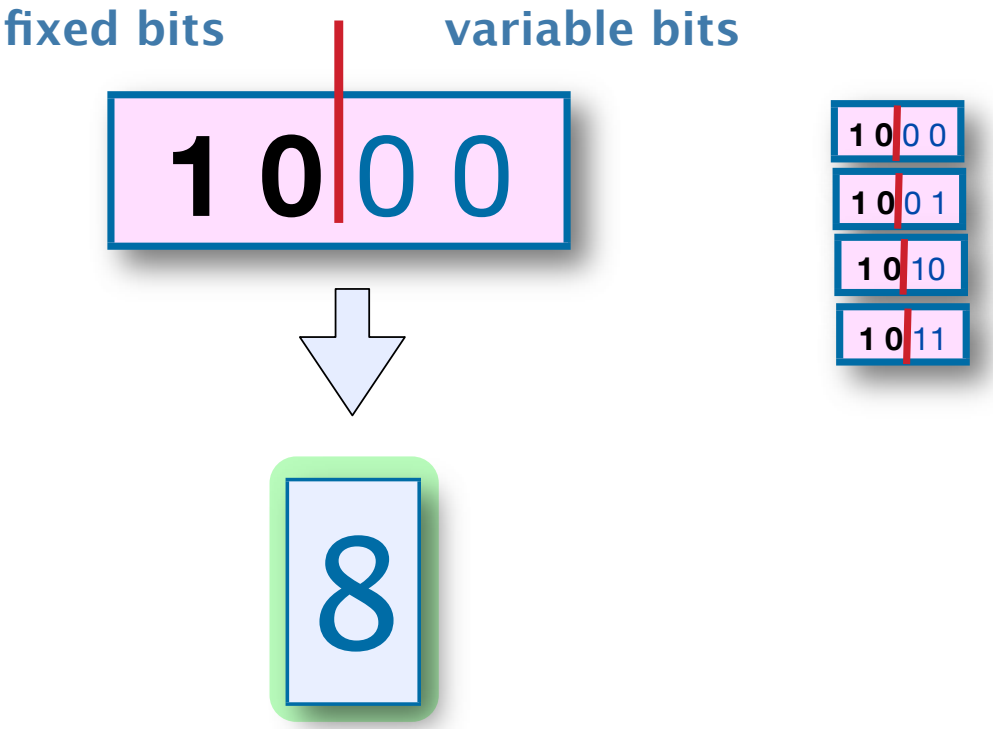

8, 9, a or b only!

“Easy” & “complicated” ranges

- 2001:db8:7::/48
  - 2001:db8:7:xxxx:xxxx:xxxx:xxxx:xxxx
- 2001:db8:7:8000::/50
  - 2001:db8:7:8xxx:xxxx:xxxx:xxxx:xxxx
  - 2001:db8:7:9xxx:xxxx:xxxx:xxxx:xxxx
  - 2001:db8:7:axxx:xxxx:xxxx:xxxx:xxxx
  - 2001:db8:7:bxxx:xxxx:xxxx:xxxx:xxxx

IPv6 Subnetting

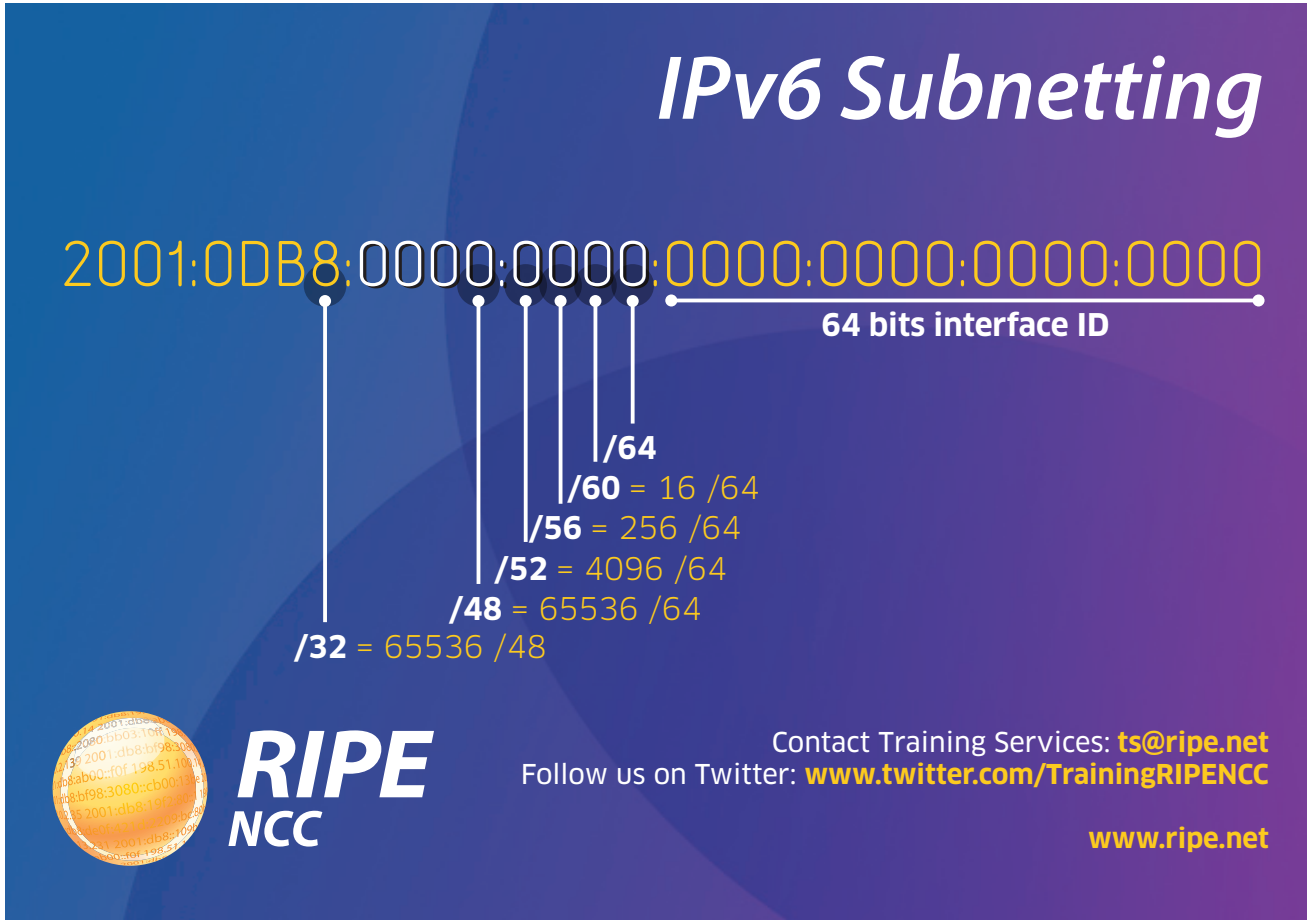

IPv4 vs IPv6 (rounded off)

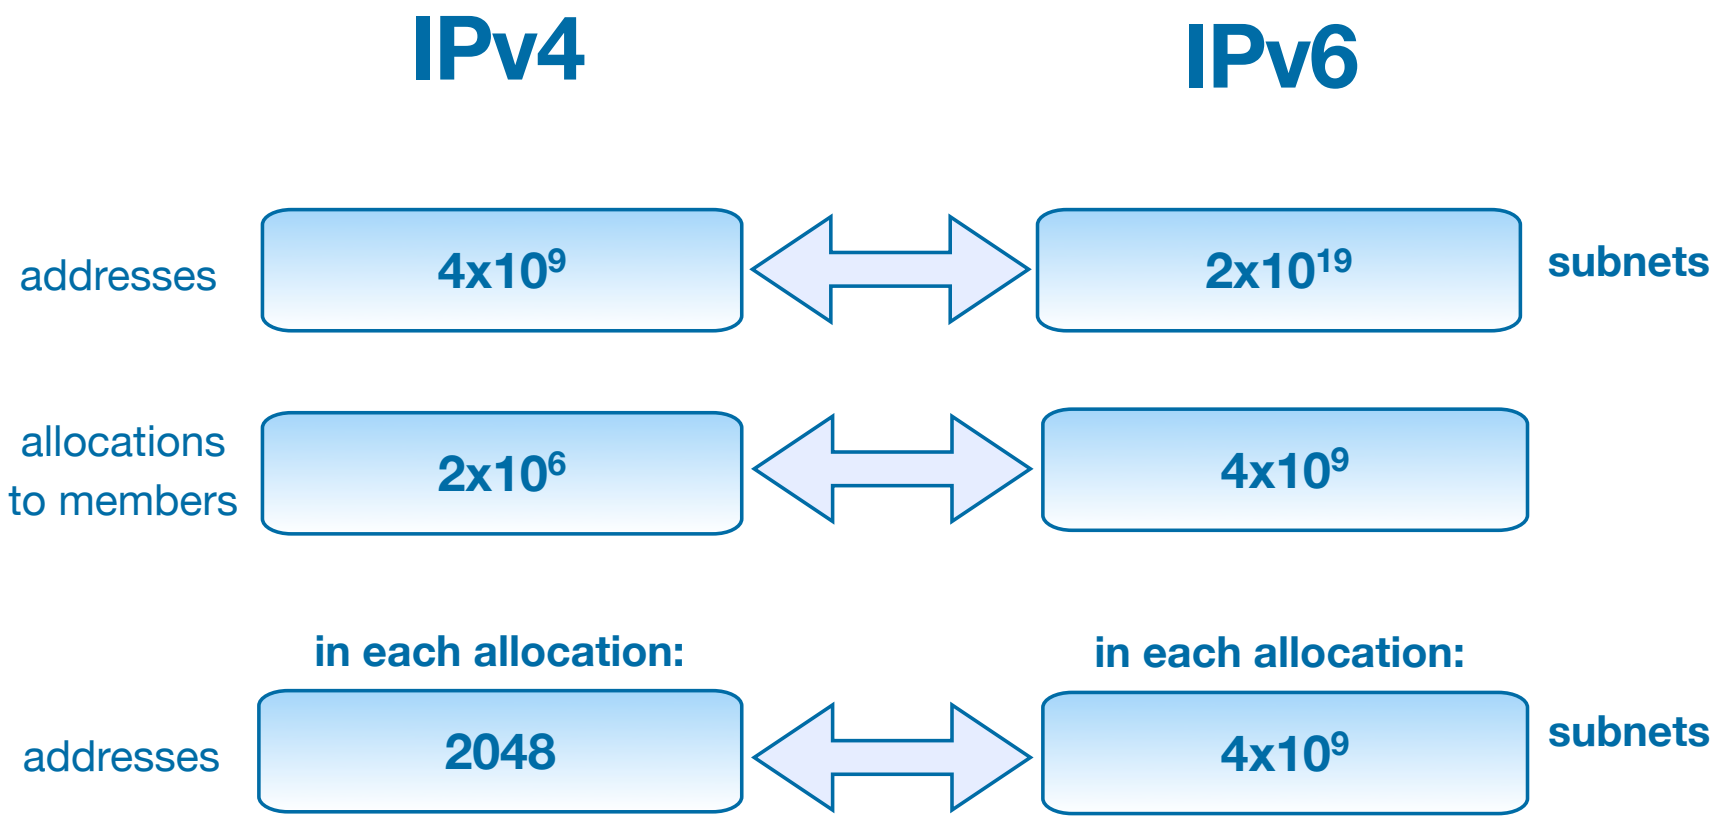

Addressing Plans

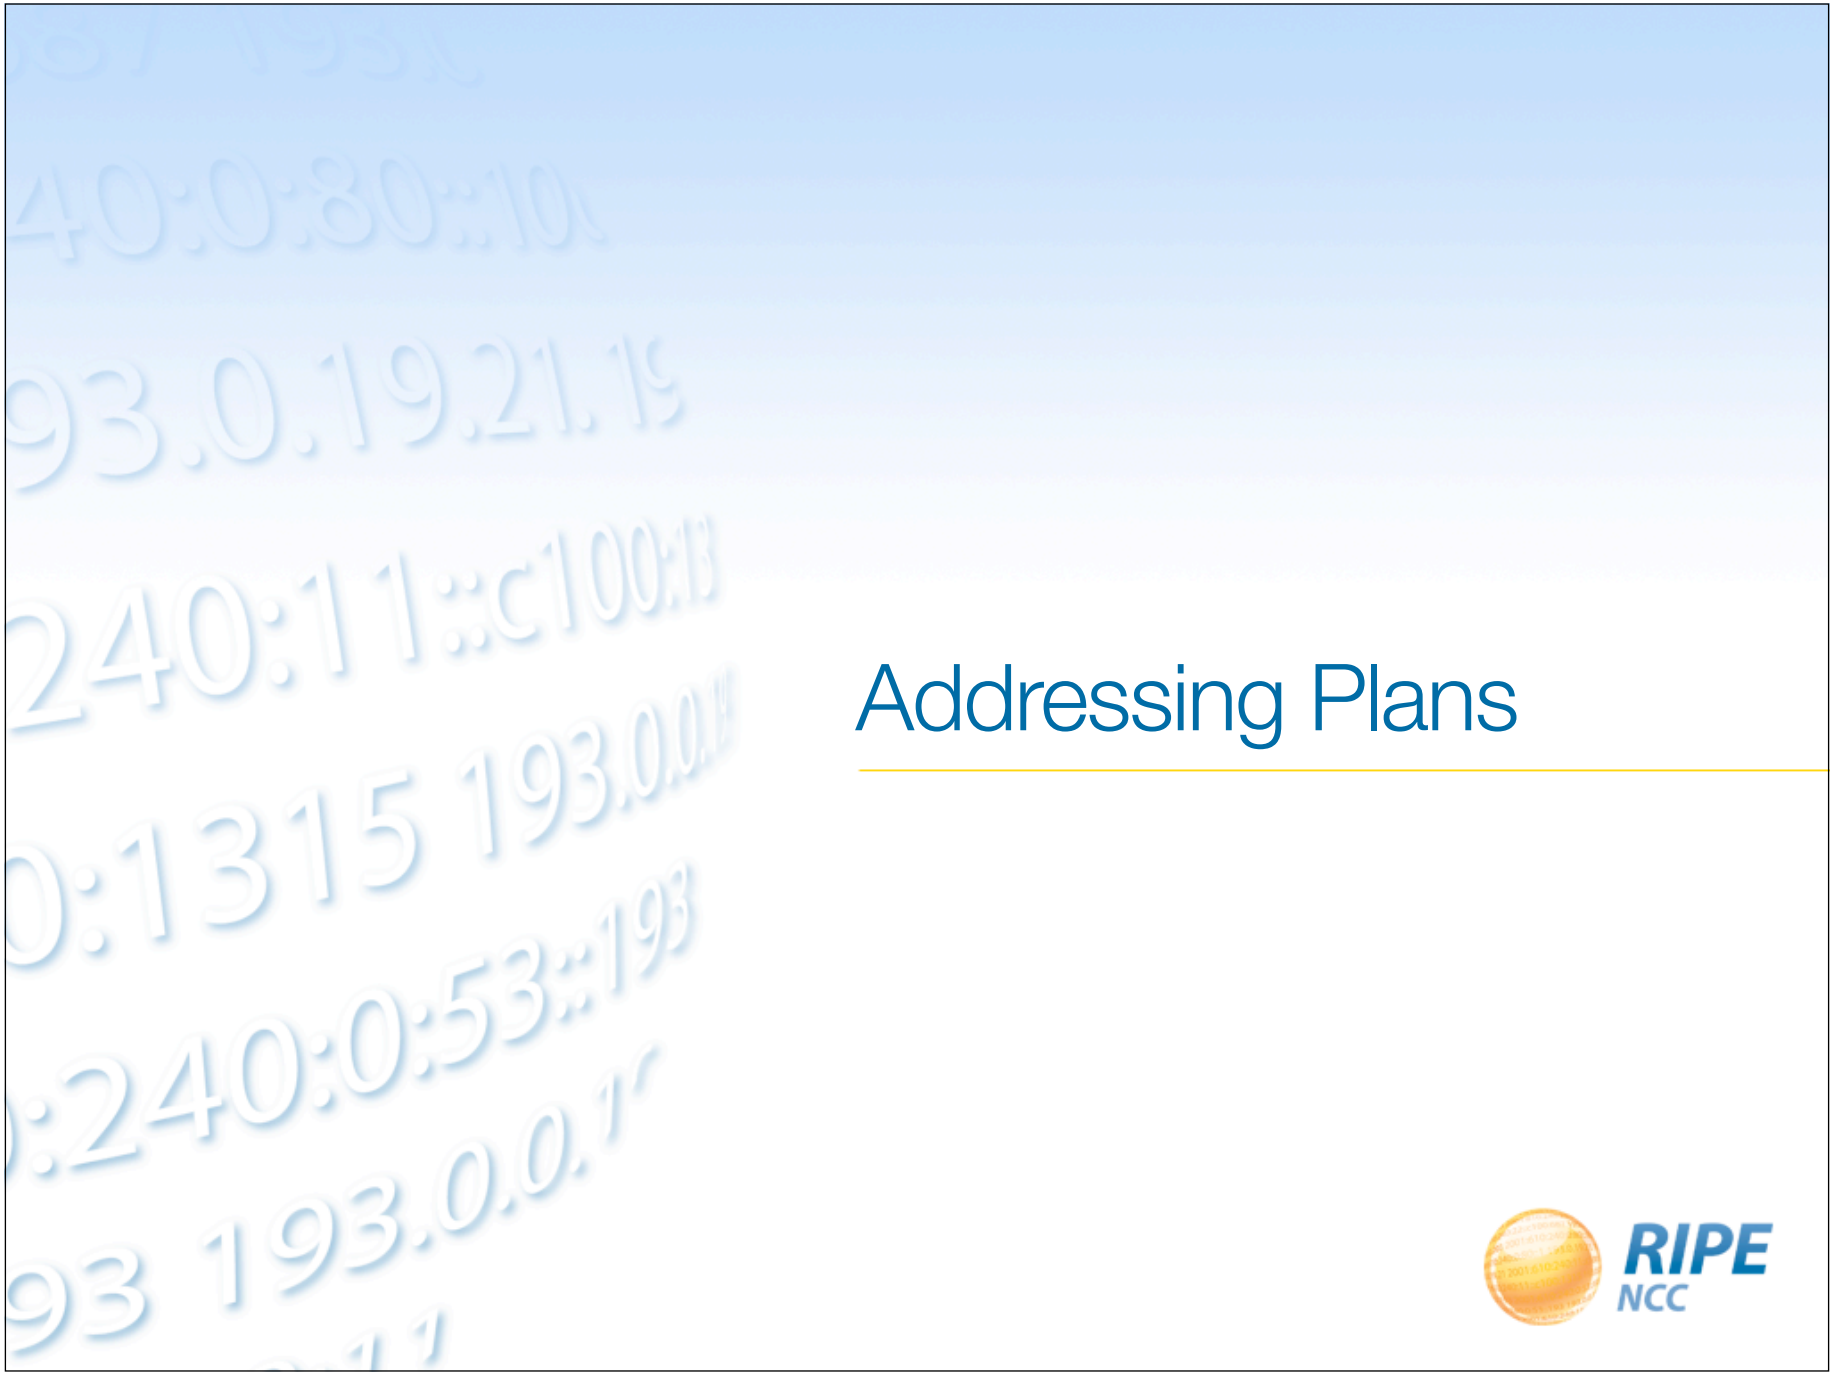

Why create an IPv6 addressing plan?

- Easier implementation of security policies
- Efficient addressing plans are scalable
- More efficient route aggregation

and most important...

Keep your mental health!

Image source: http://bit.ly/11YvFVC

RIPE NCC

Addressing plan example

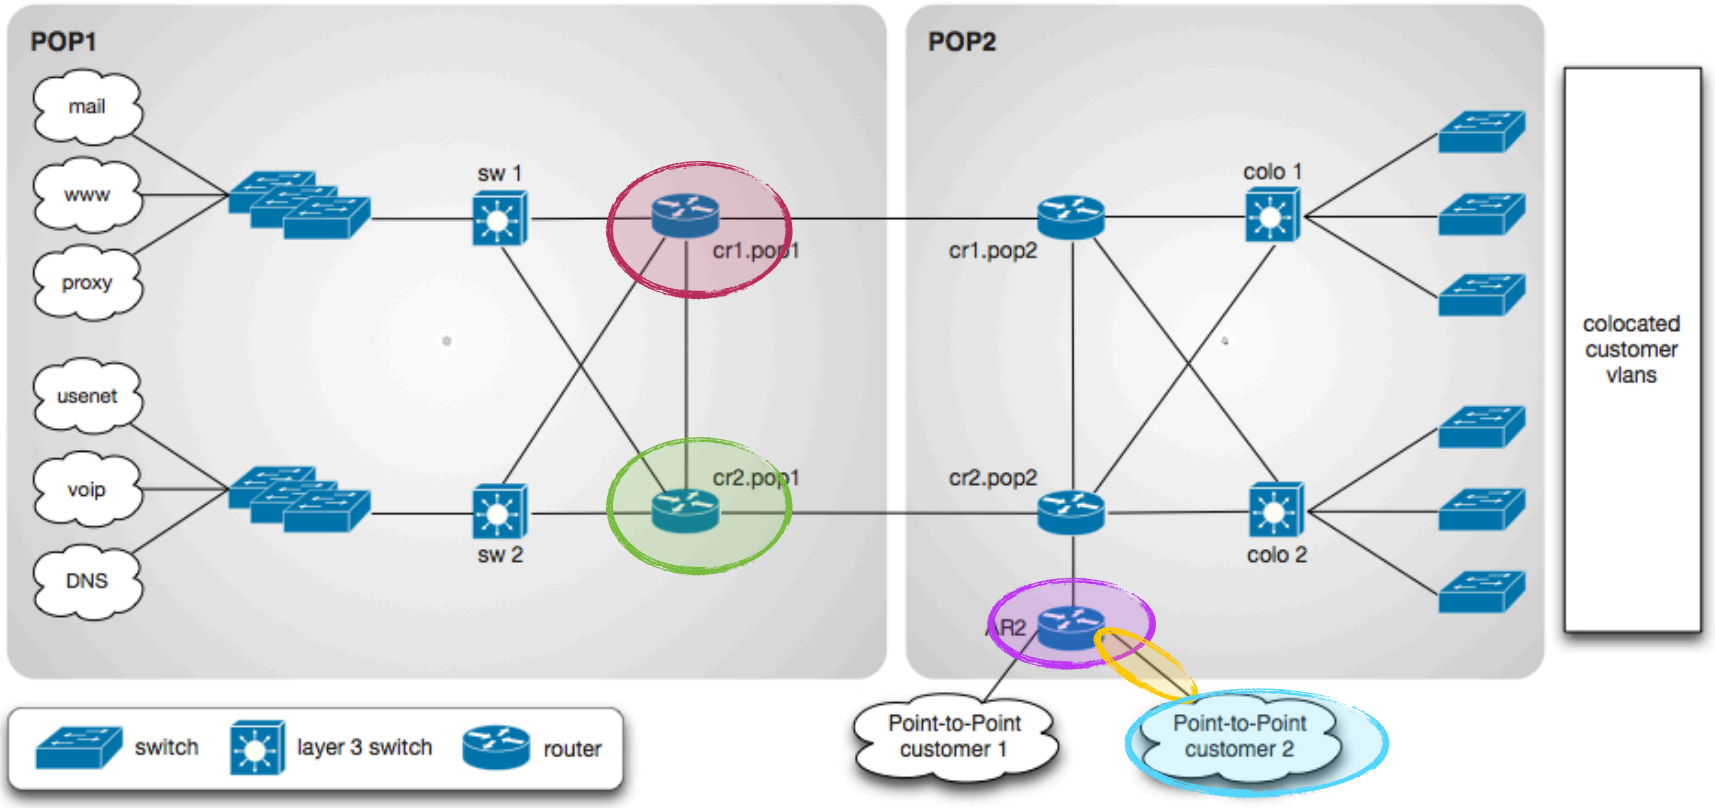

Solution POP1

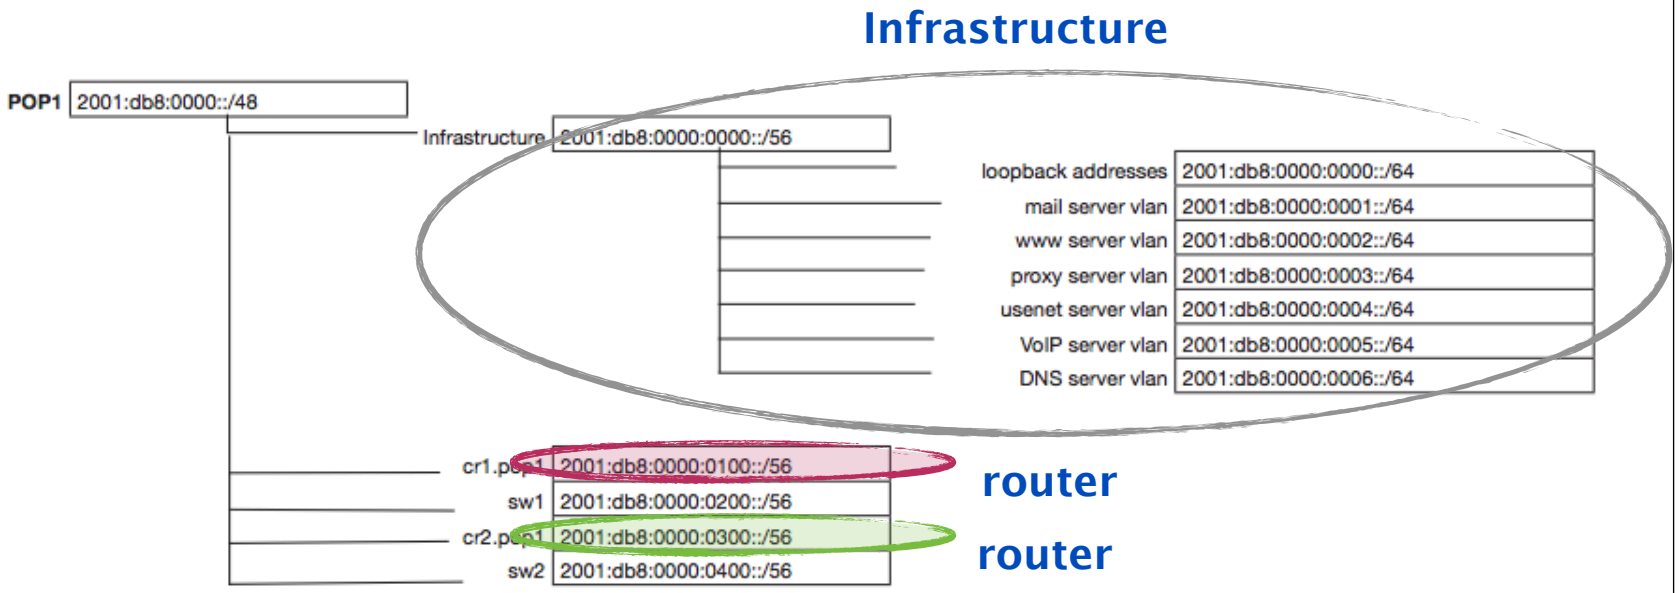

Solution POP2

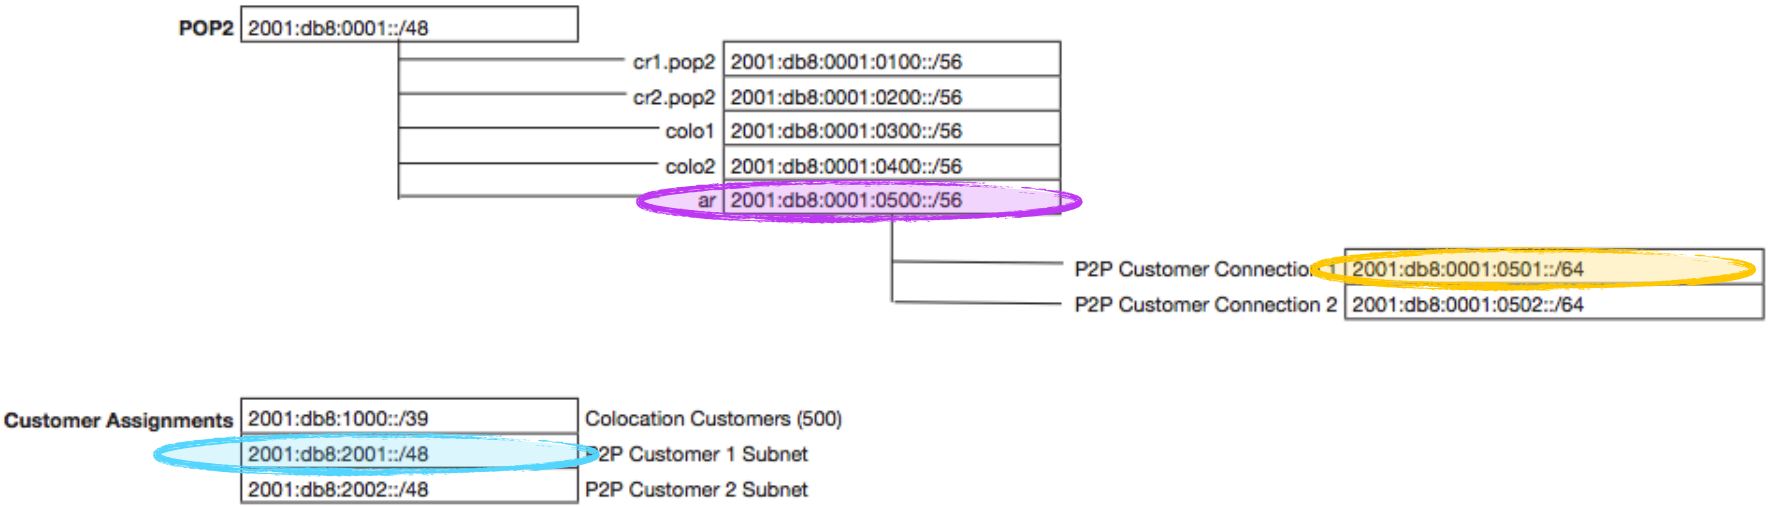

Make an addressing plan (I)

- Number of hosts is irrelevant
- Multiple /48s per pop can be used
  - separate blocks for infrastructure and customers
  - document address needs for allocation criteria
- /64 for all subnets
  - autoconfiguration works
  - renumbering easier
  - less typo errors because of simplicity

Make an addressing plan (II)

- Routers:
- Give all routers the same size block
- Minimum: One /64 per interface
- Allow for more interfaces in future
- /56 or /52 typical for a router

Make an addressing plan (II)

- Use one /64 block (per site) for loopbacks
  - One /128 per device

One /64 =

18.446.744.073.709.551.616

IPv6 addresses

More On Addressing Plans for ISPs

- For servers you want manual configuration
- Use port numbers for addresses

- pop server 2001:db8:1::110
- dns server 2001:db8:1::53
- etc...

Point-to-Point Connections

- Reserve a /64, assign a /127

Customer assignments

- Give your customers enough addresses
  - Up to a /48
- For more addresses, send in request form
  - Alternatively, make a sub-allocation
- Every assignment must be registered in the RIPE database

Customers And Their /48

- Customers have no idea how to handle 65536 subnets!
- Give them information

http://bit.ly/116HCTg

IPv6 Address Management

- Your Excel sheet might not scale
  - There are 65.536 /48s in a /32
  - There are 65.536 /64s in a /48
  - There are **16.777.216** /56s in a /32

- Find a suitable IPAM solution

IPv6 & Address Translation for Mobile Telephony

IPv6 and IPv4 compatibility?

- IPv6 is a different protocol from IPv4
- IPv6 hosts cannot talk to IPv4 hosts directly
- Tools like 6in4 and other transition mechanisms

  let IPv6 hosts talk to each other

  - tunneling
  - translation

NAT64/DNS64

- Single-stack clients will only have IPv6
- Translator box will strip all headers and replace them with IPv4
- Requires some DNS “magic”
  - Capture responses and replace A with AAAA
  - Response is crafted based on target IPv4 address
- Usually implies address sharing on IPv4

NAT64/DNS64

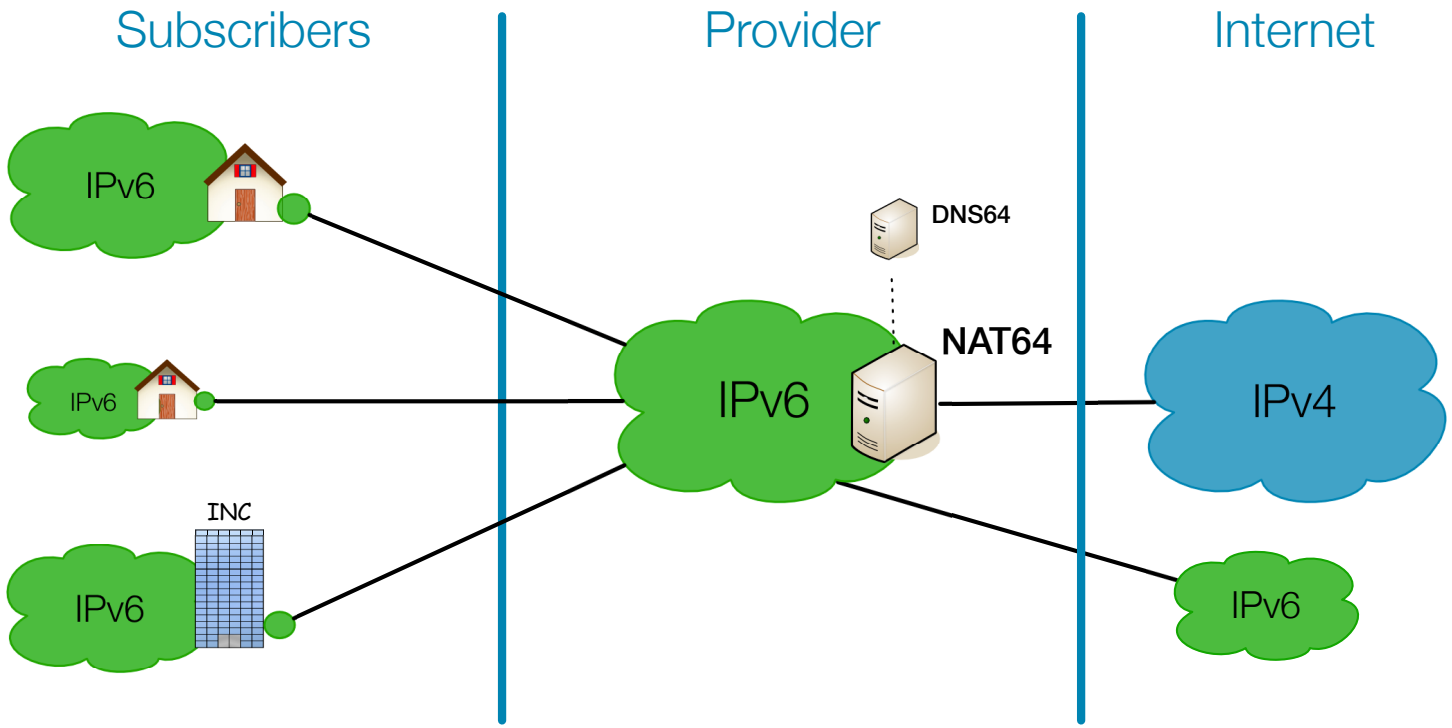
Subscribers
Provider
Internet
IPv6
IPv6
INC
IPv6
DNS64
NAT64
IPv6
IPv4
IPv6

Drawback

- Some applications don't work on IPv6 only devices
  - Spotify, Netflix, **Skype**

- Solution?
  - **464XLAT**
  - makes **IPv4-only** applications work on **IPv6-only** device

464XLAT

- NAT64+Stateless IP translation on device

- on IPv6 only mobile devices
  - Install **CLAT** demon locally
  - 464XLAT gives the mobile dummy IPv4 address
  - IPv4 only application can use IPv4 interface
    - and works!
  - CLAT translates IPv4 to IPv6 locally
  - **NAT64** for accessing IPv4 networks

Deployment?

- T-Mobile US, Verizon
- phones
  - Nexus S, Galaxy Nexus, Galaxy S, Galaxy Note, Verizon LTE
- Android
  - CLAT open source
- Android 4.3
  - CLAT built in

Useful links & hints

- 464xlat details:
  - https://sites.google.com/site/tmoipv6/464xlat
- RFC 6877 (464XLAT)
- RFC 6146 (NAT64-in the core)
- RFC 6145 (IP/ICMP translation on the edge)
- CLAT installation on Android platform
  - http://dan.drown.org/android/clat/index.html
- Video of live demo at 3rd World IPv6 Congress, Paris
  - http://www.internetsociety.org/deploy360/blog/2013/04/video-464xlat-live-demo-at-world-ipv6-congress-in-paris

Useful IPv6 info

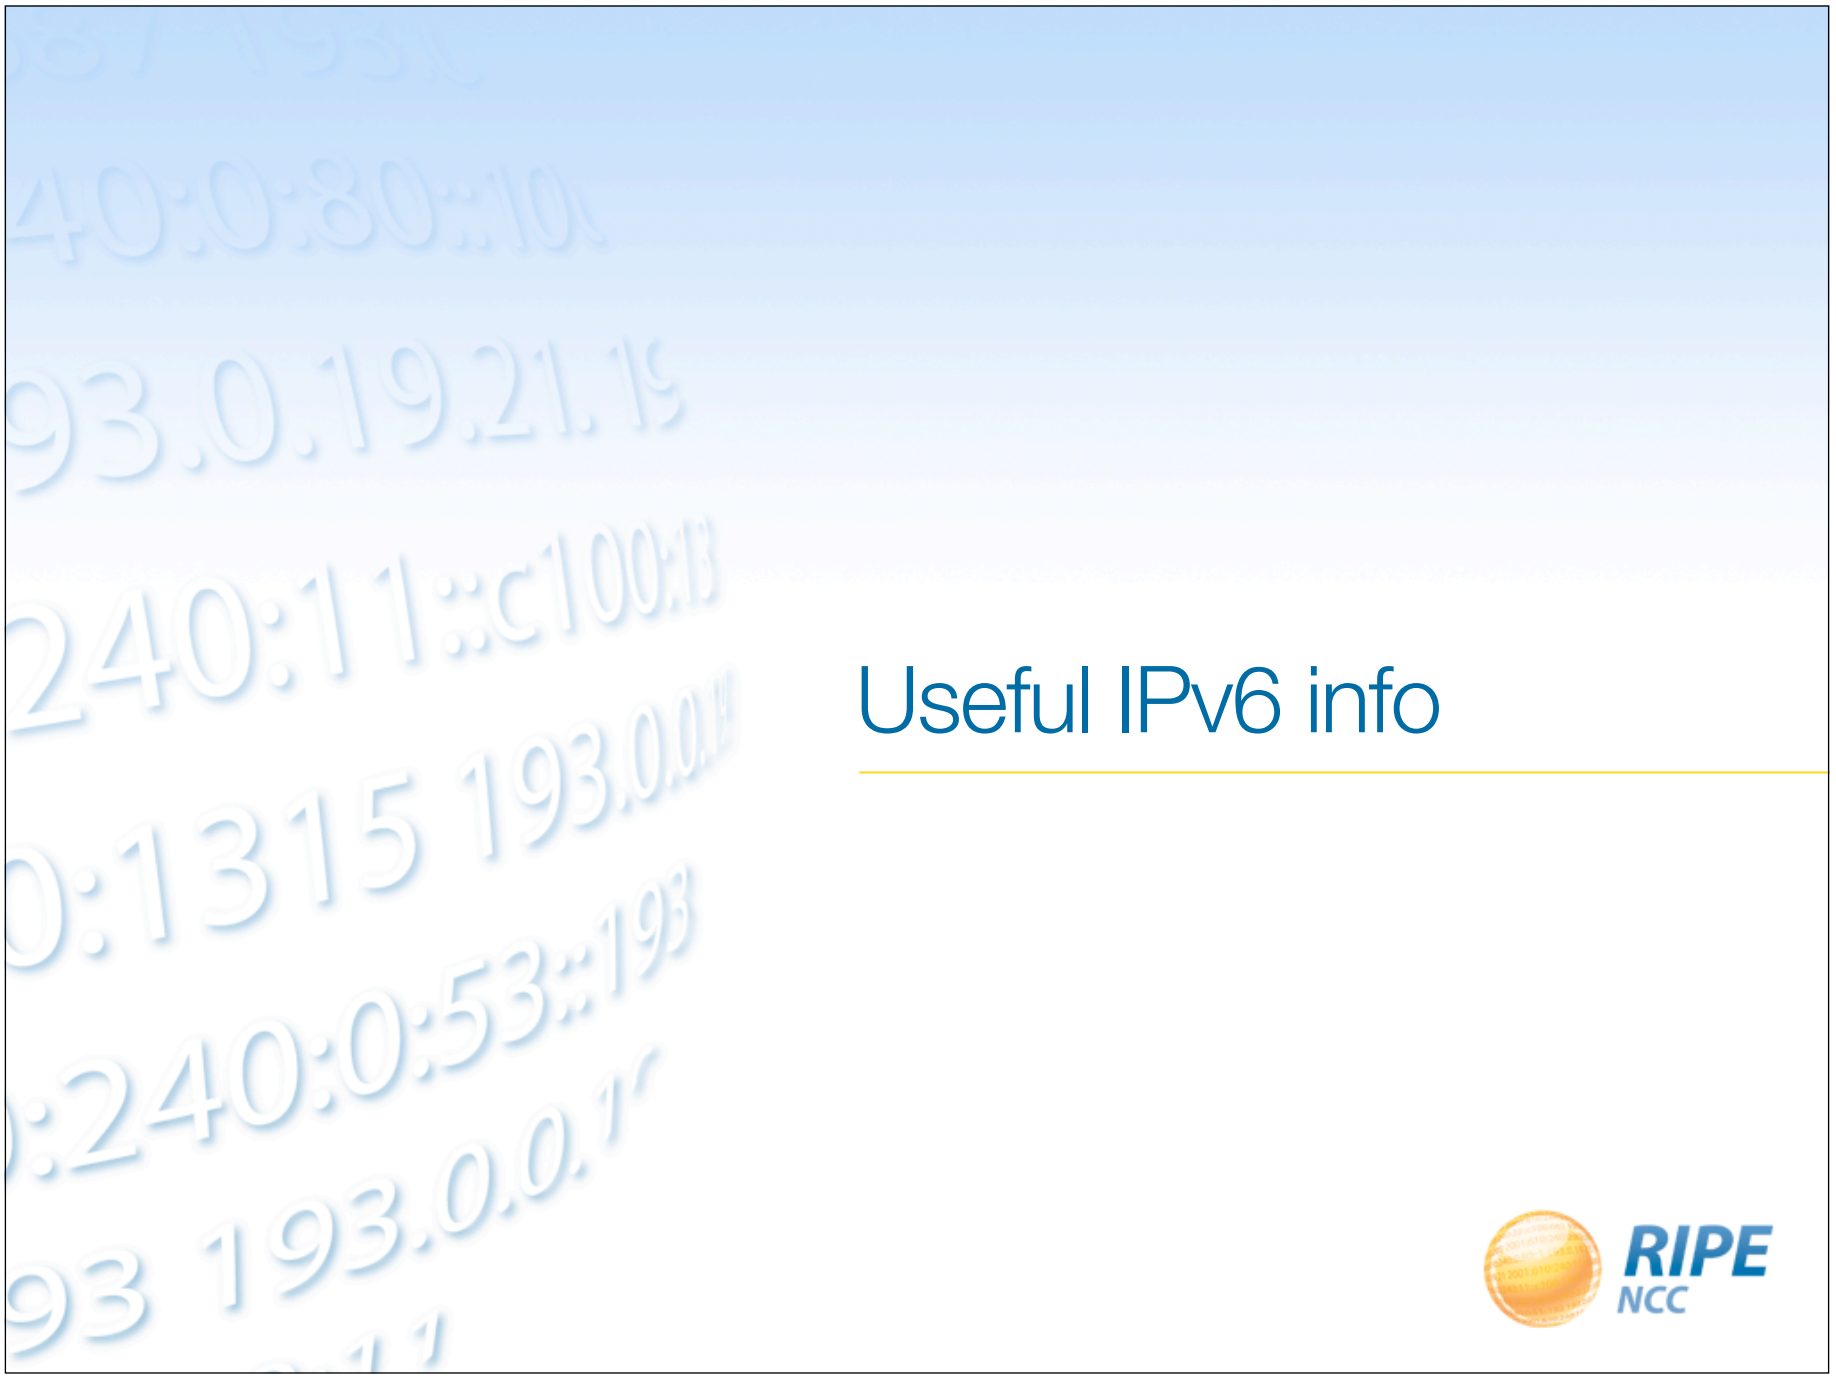

Useful information

Websites

- http://www.getipv6.info/
- http://www.ipv6actnow.org
- http://datatracker.ietf.org/wg/v6ops/
- http://www.ripe.net/ripe/docs/ripe-554.html

Mailing lists

- http://lists.cluenet.de/mailman/listinfo/ipv6-ops
- http://www.ripe.net/mailman/listinfo/ipv6-wg

More Questions?

Come to our 1 day free IPv6 training!

Only for RIPE NCC members:

www.ripe.net/training


The End!
Край
Y Diwedd
Fí
Finis
النهاية
Соңы
վերջ
Liðugt
Ende
Finvezh
Кінець
Konec
Kraj
Ënn
Fund
پایان
Lõpp
Beigas
Vége
Son
Kpaj
An Críoch
הסוף
Fine
Endir
Sfârşit
Fin
Τέλος
Einde
Конец
Канец
Slut
Slutt
დასასრული
Pabaiga
Tmiem
Amaia
Koniec
Fim
Loppu