IPv6 Tutorial

MENOG 10 - April 2012

Matt Parker

RIPE NCC

About the RIPE NCC

Section 1

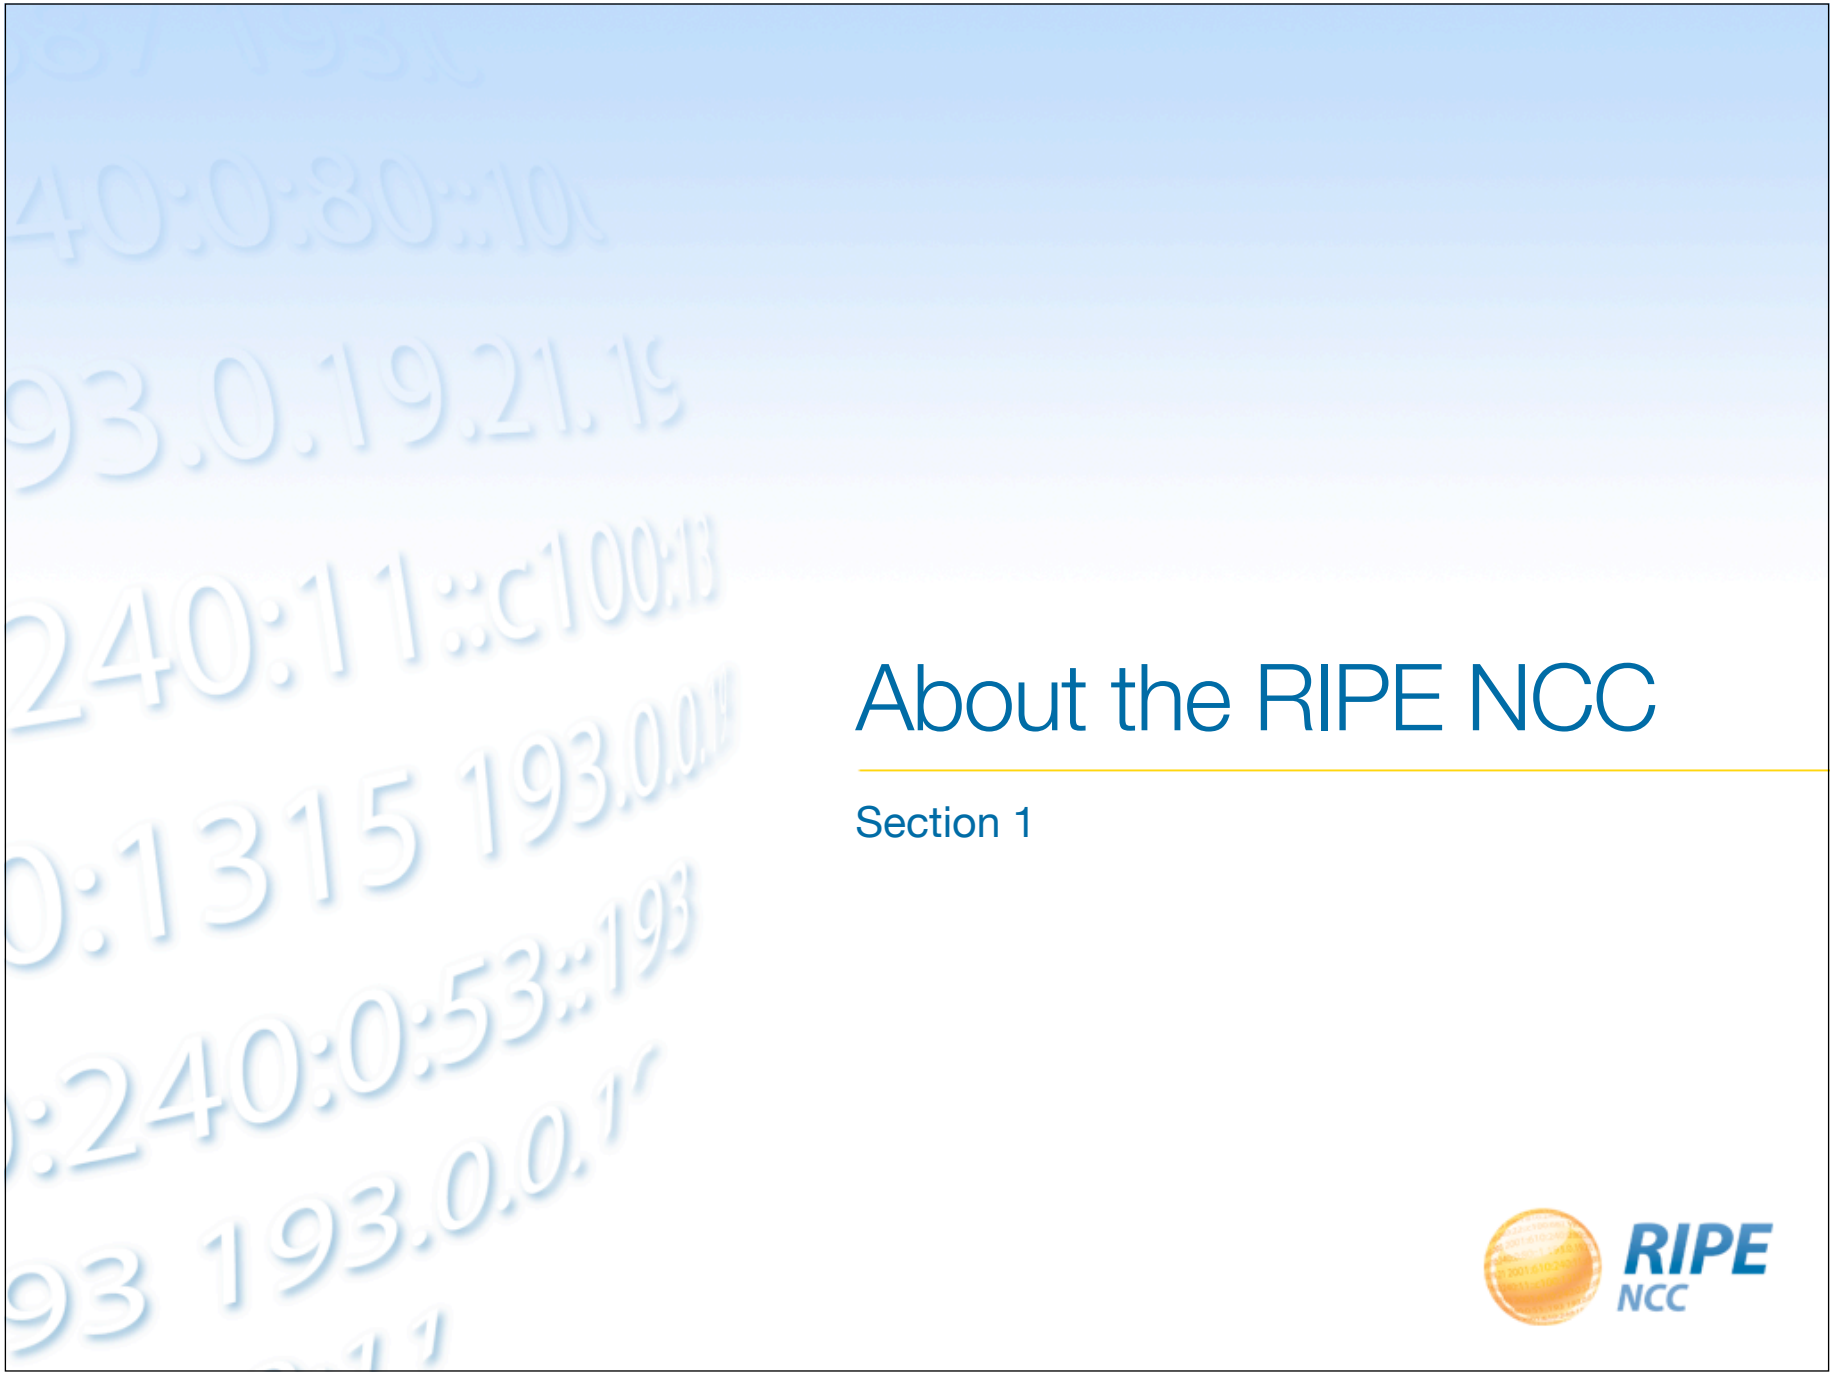

The IR system

- Five RIRs worldwide
  - Not for profit organisations
  - Funded by membership fees
  - Policies decided by regional communities

The five RIRs

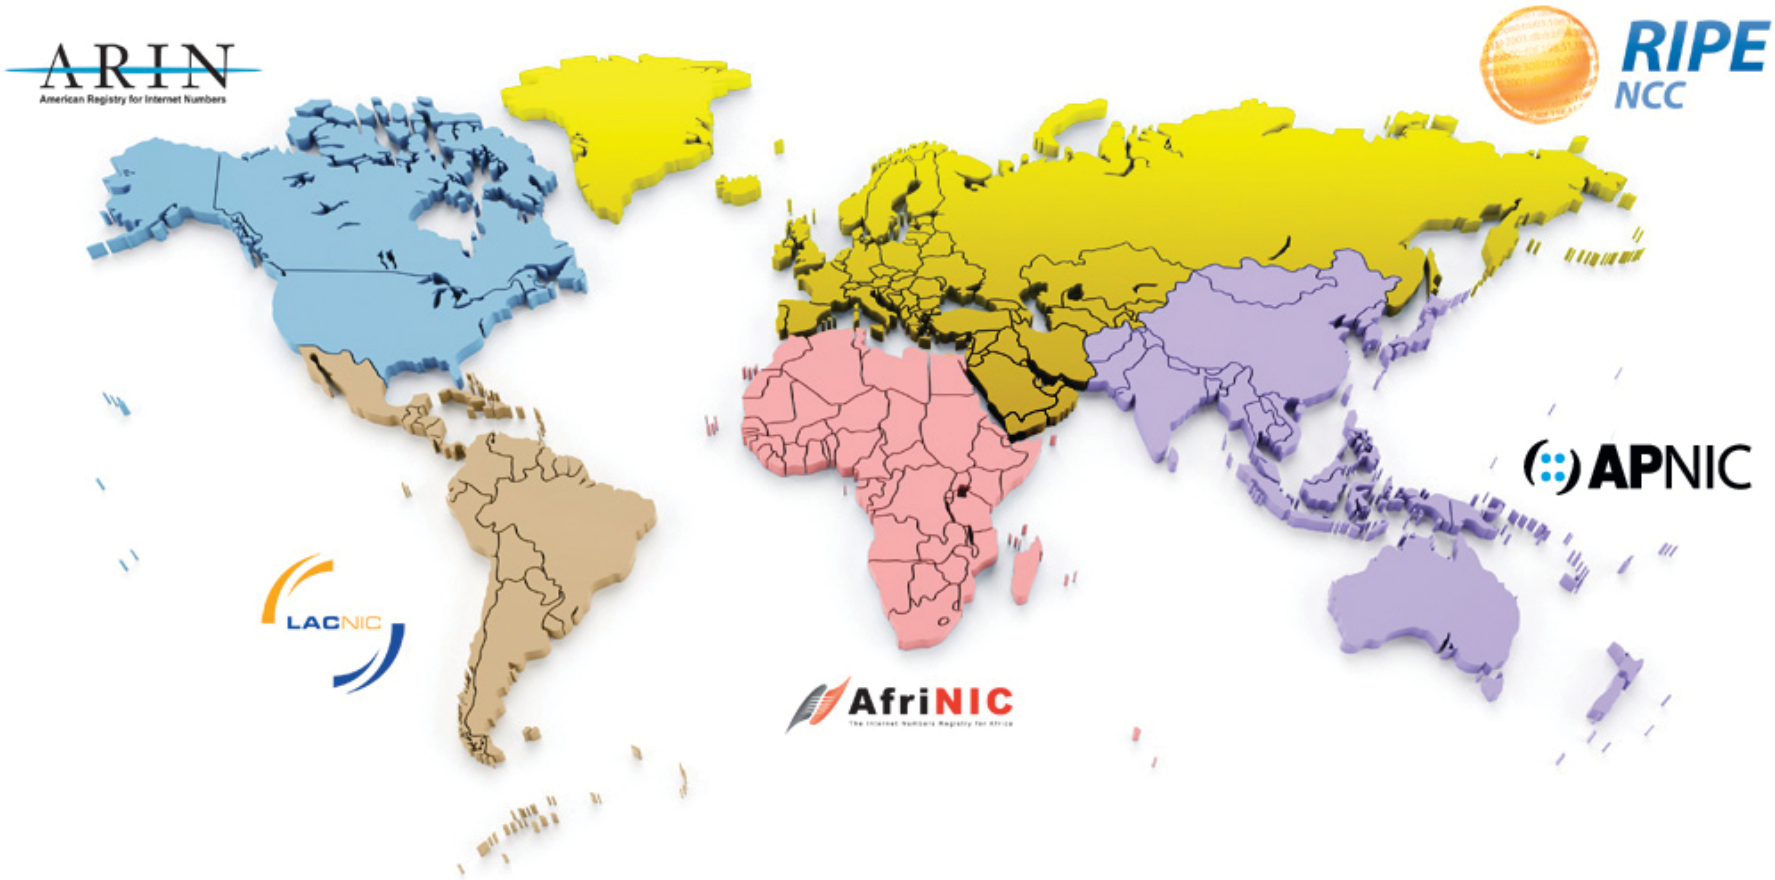

Goals: Registration

- Why?
  - Ensure uniqueness of Internet number resources
  - Provide contact information for users of Internet number resources
- How?
  - RIR whois databases
- Results:
  - IP address space used only by one organisation
  - Information available on users of Internet number resources

Goals: Aggregation

- Why?
  - Routing tables growing too fast
  - Provide scalable routing solution for Internet
- How?
  - Encourage announcement of whole allocations
  - Introduction of Classless Inter Domain Routing (CIDR)
- Result:
  - Growth of routing tables has slowed, but could still be better

Growth of the routing table

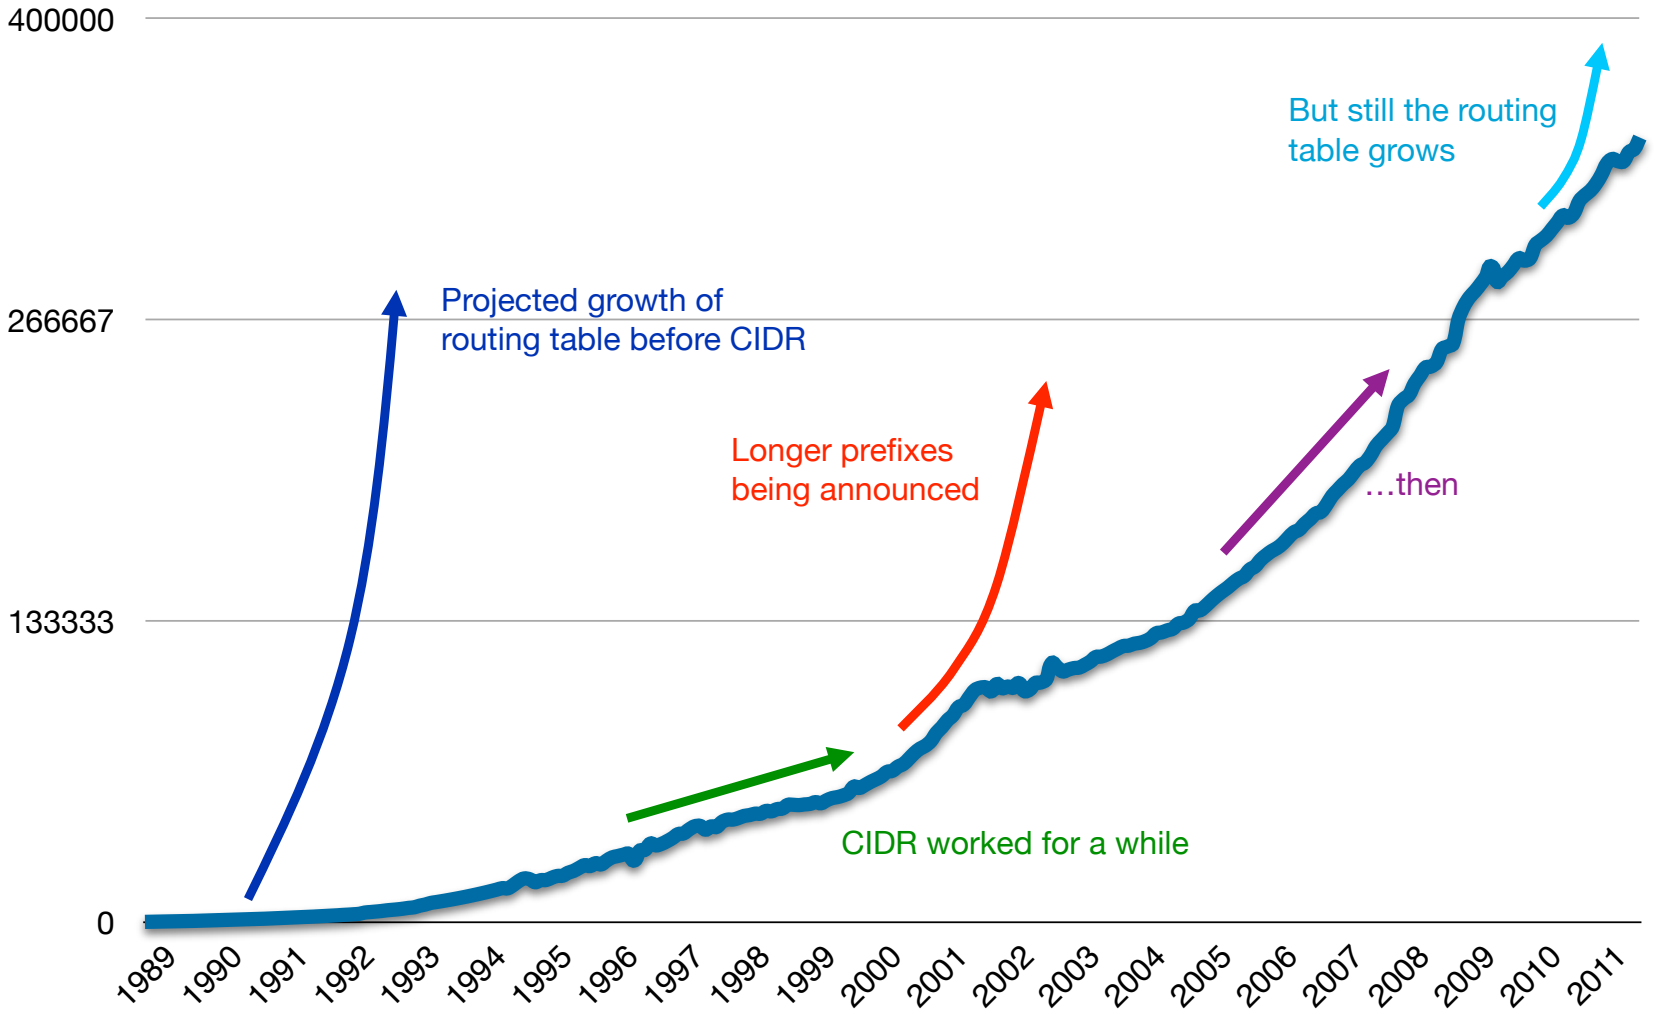

Goals: Conservation

- Why?
  - IP addresses and AS Numbers are limited resources
  - These resources were not used efficiently in the past
- How?
  - Introduction of CIDR
  - Policies to ensure fair usage
- Results:
  - Growth in IP address space usage slowed down
  - Internet number resources are distributed based on need

IP address distribution

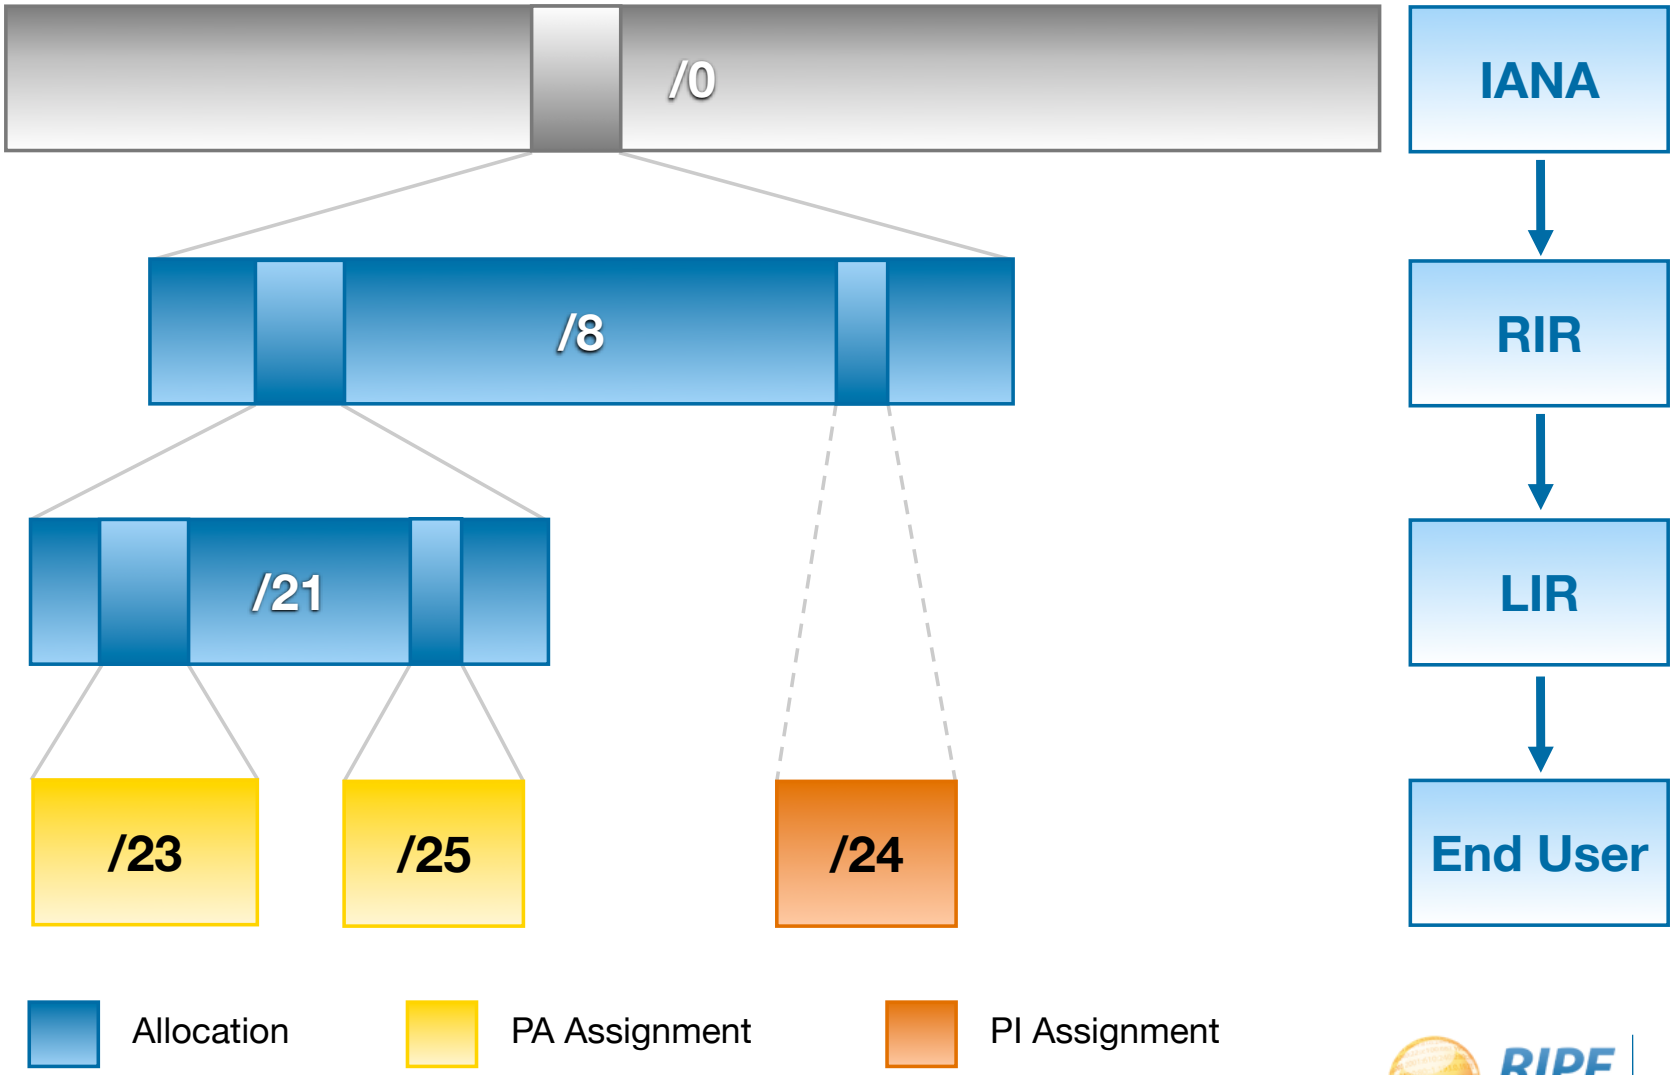

Who makes policies?

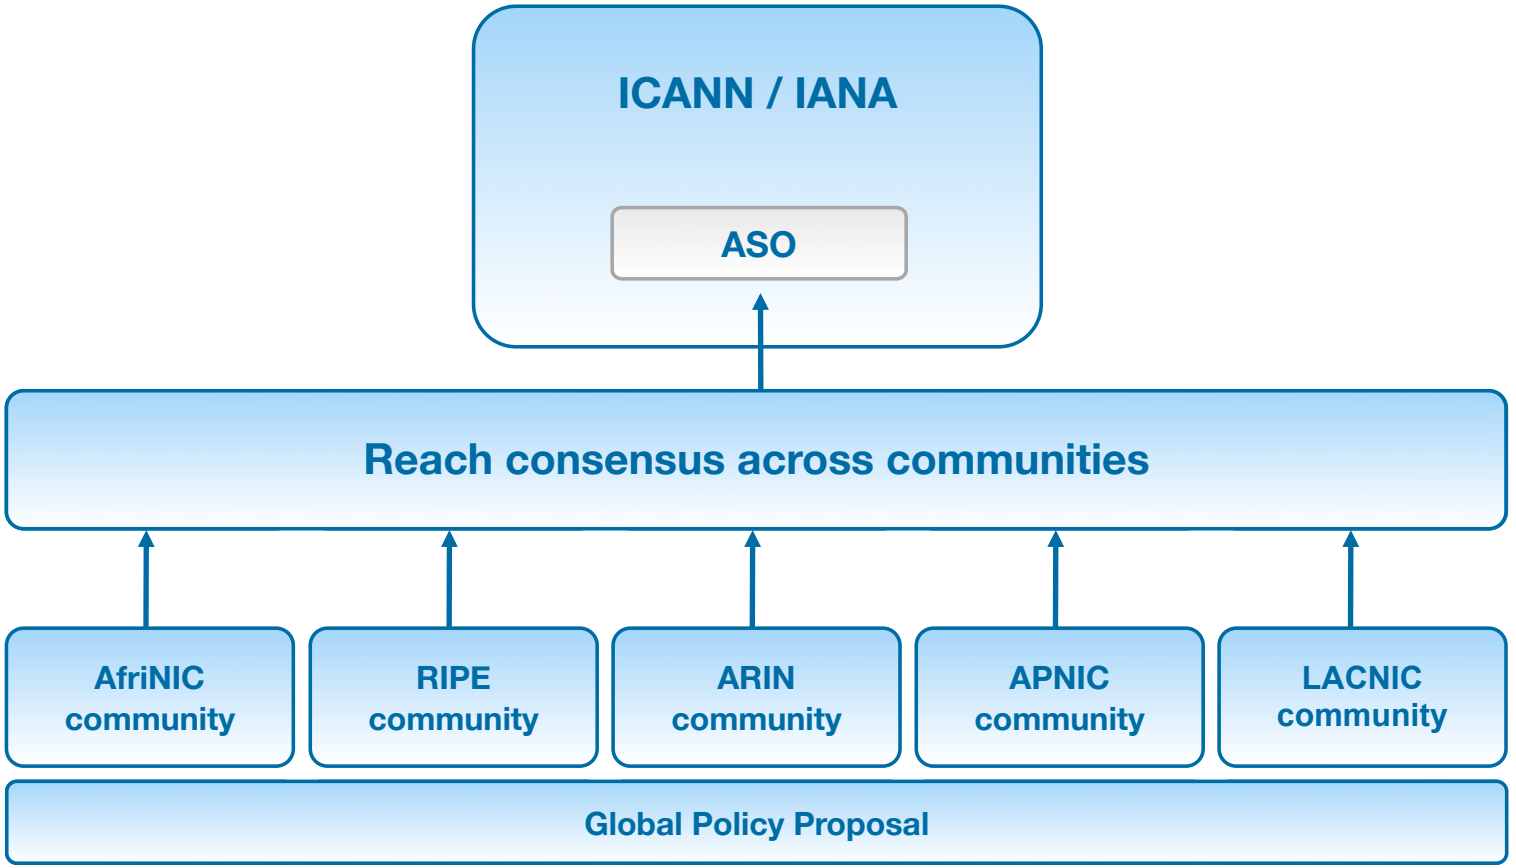

Why do policies change?

- Technological improvements
- Rapidly growing and evolving Internet industry
- Changing business requirements

Questions?

RIPE NCC

Exercise:

RIPE & RIPE NCC

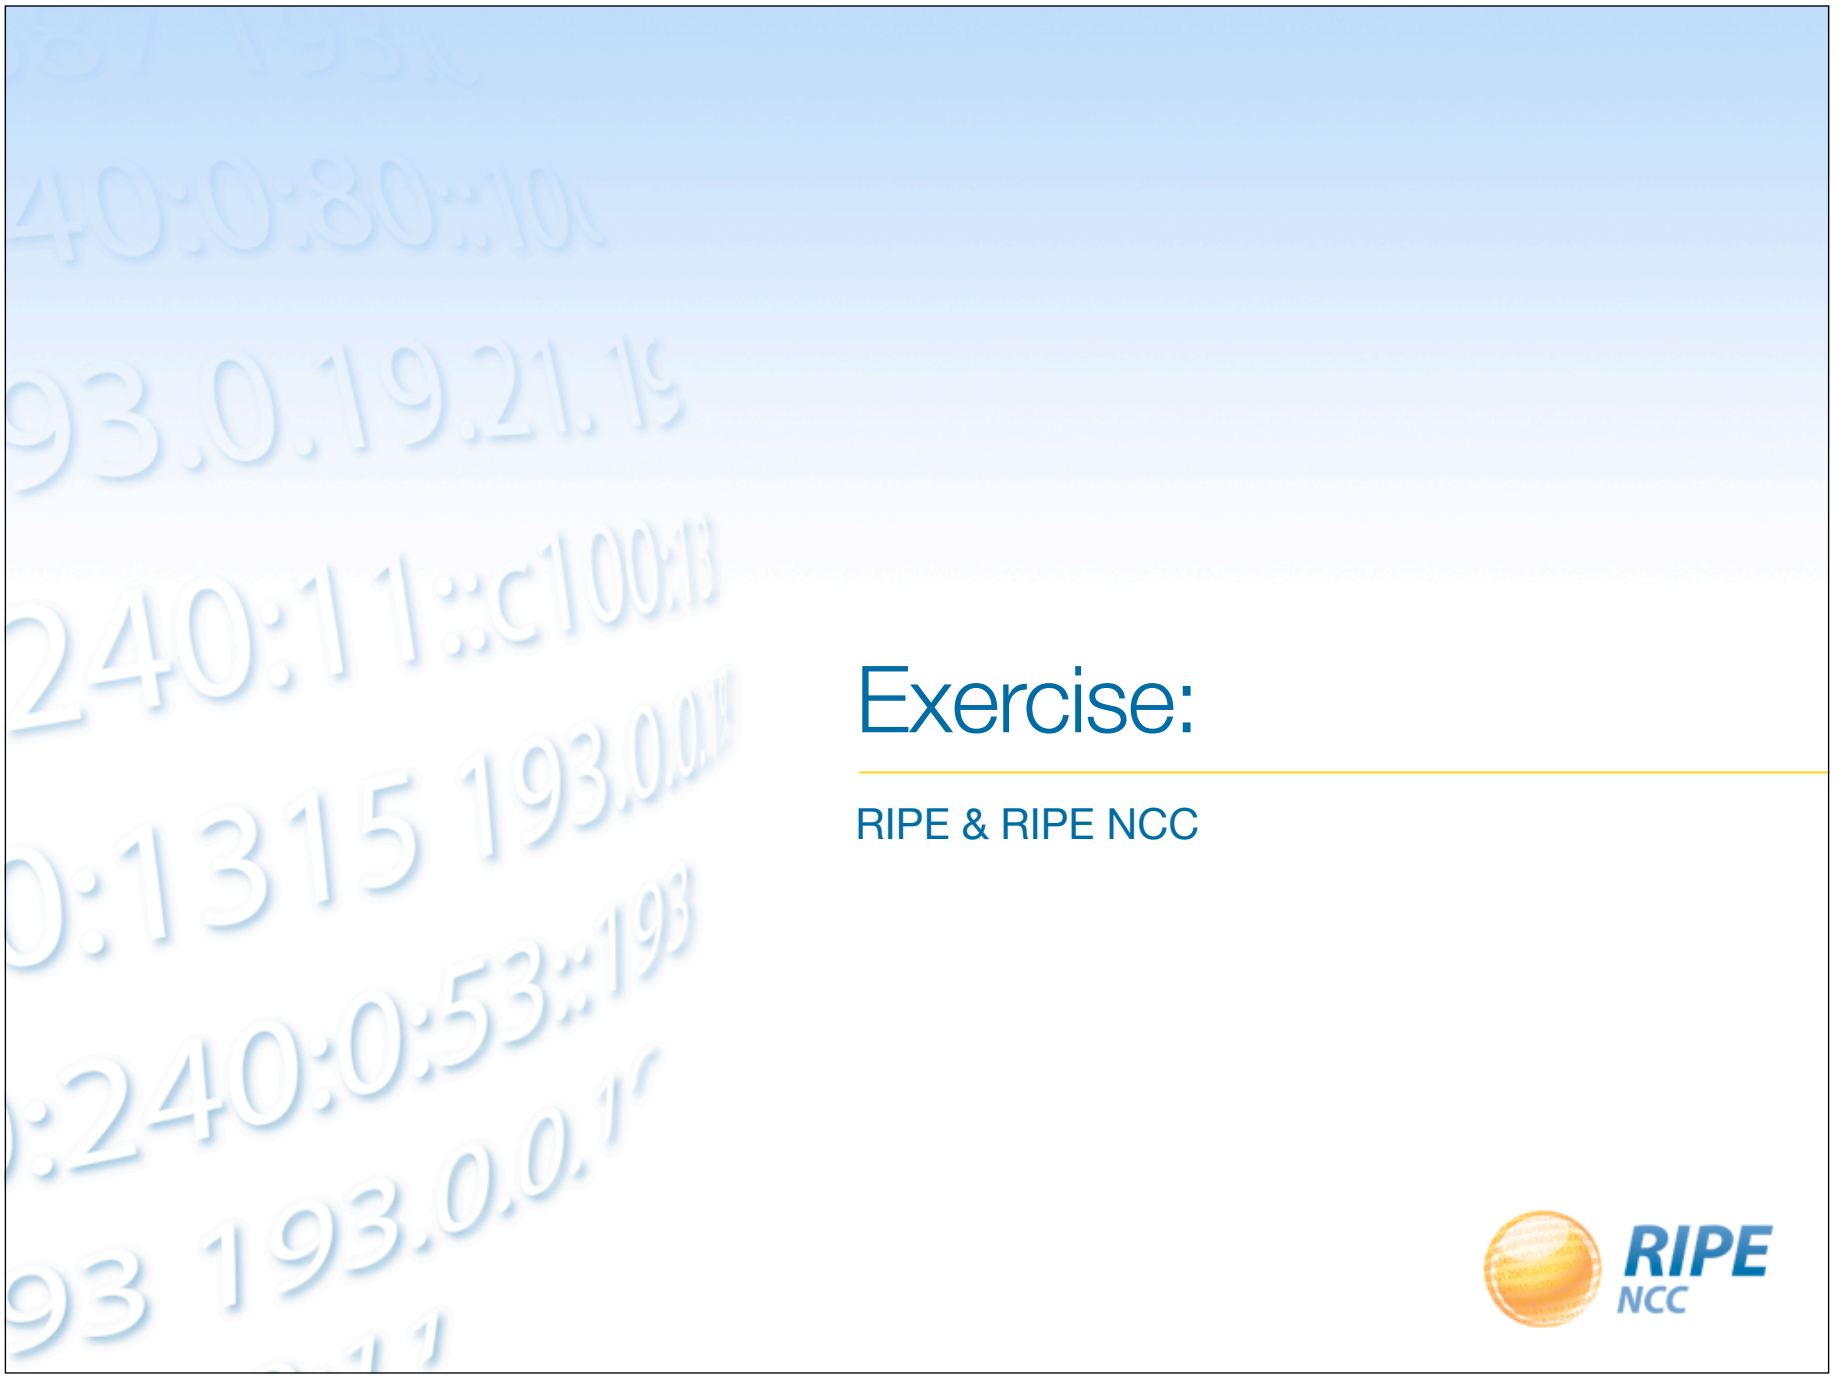

What about IPv4?

Section 2

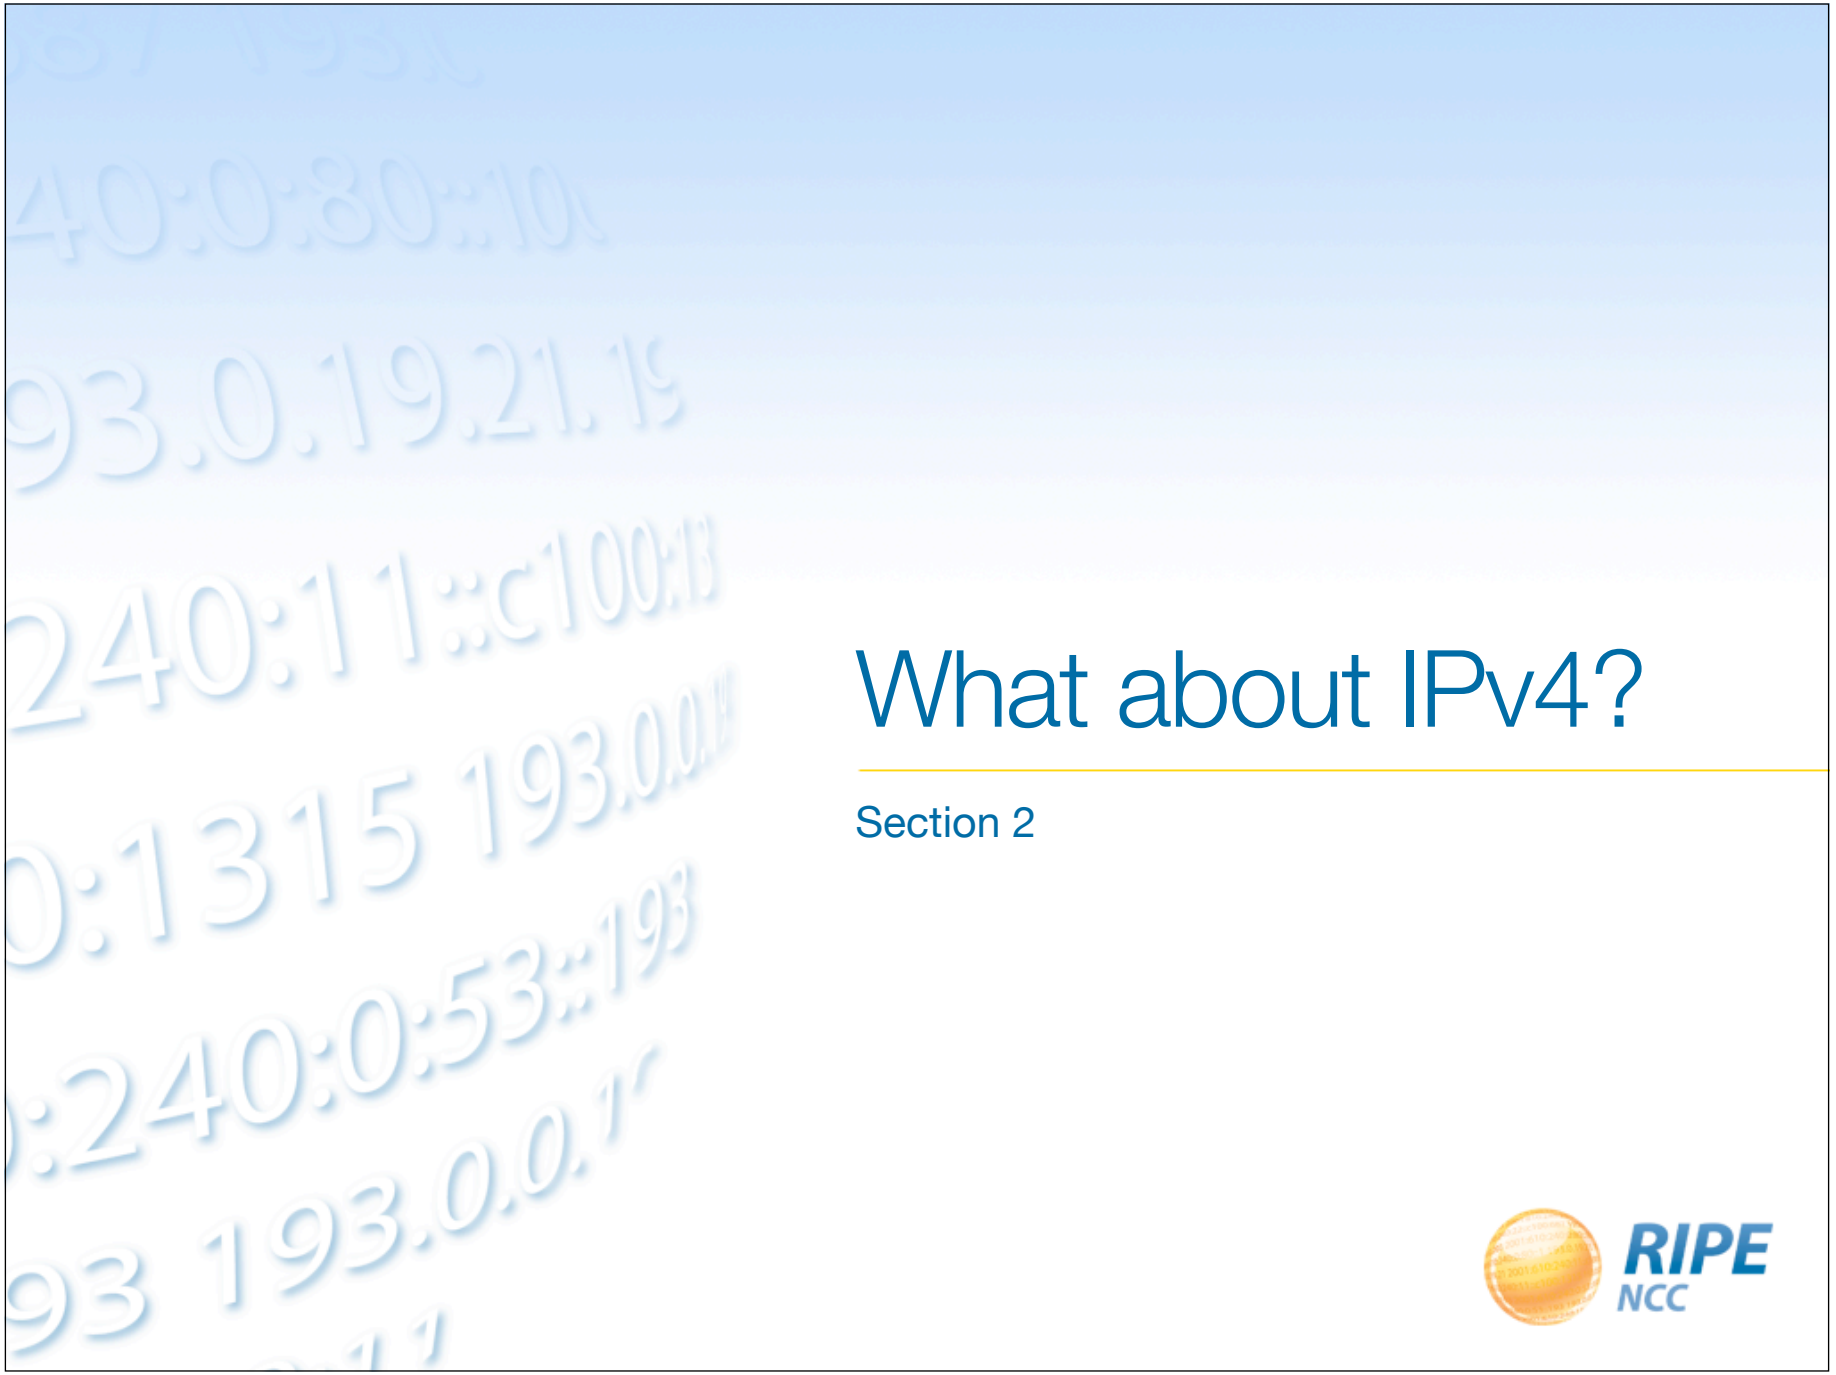

IANA IPv4 Pool

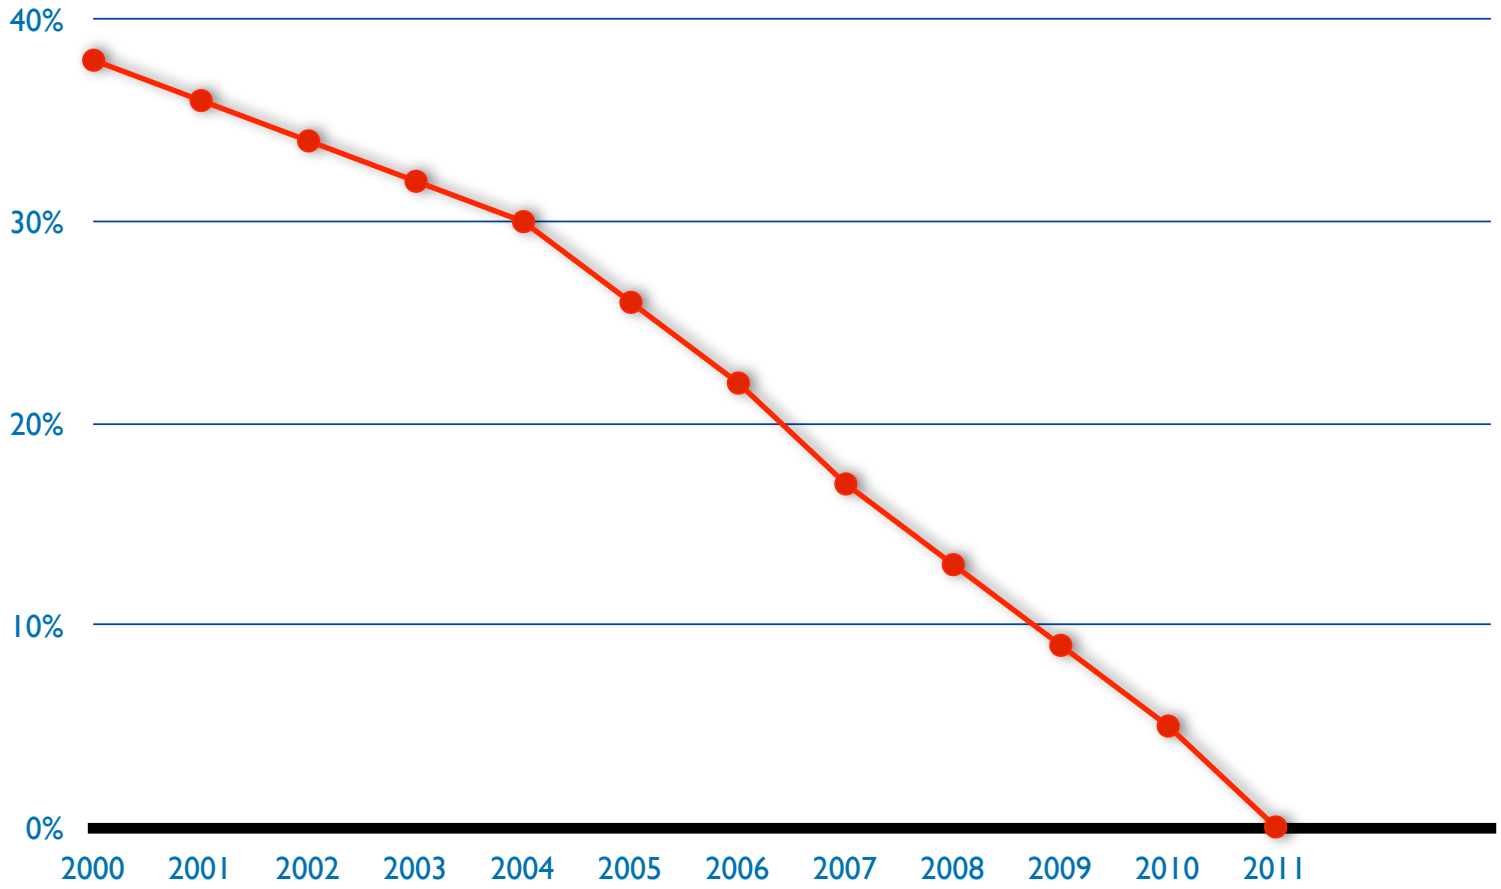

RIPE NCC IPv4 Pool - Now

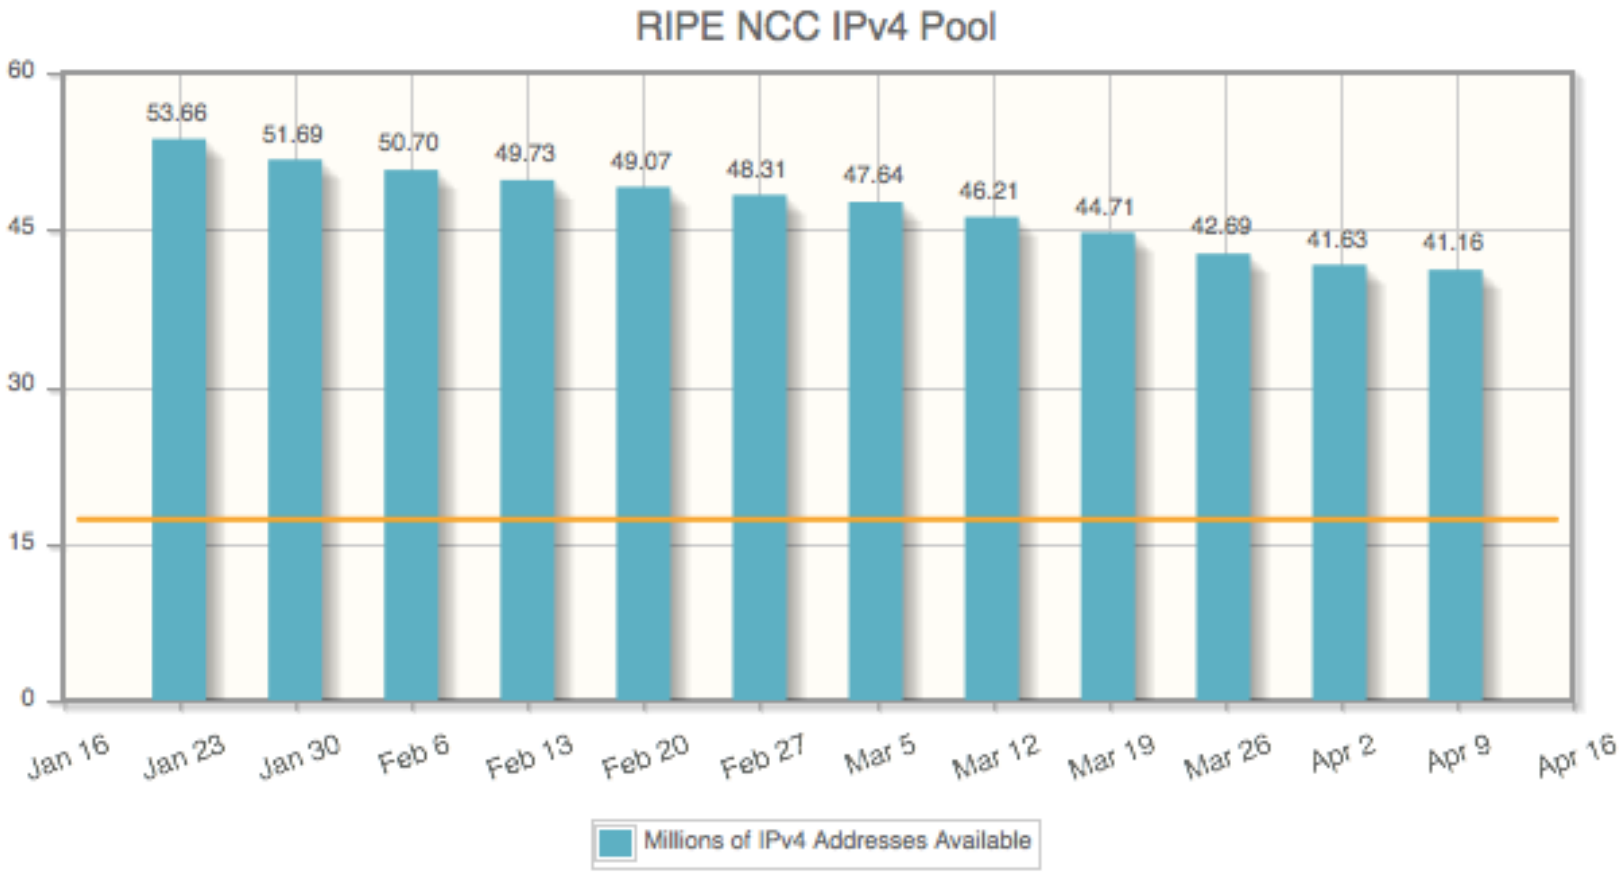

IPv4 exhaustion phases

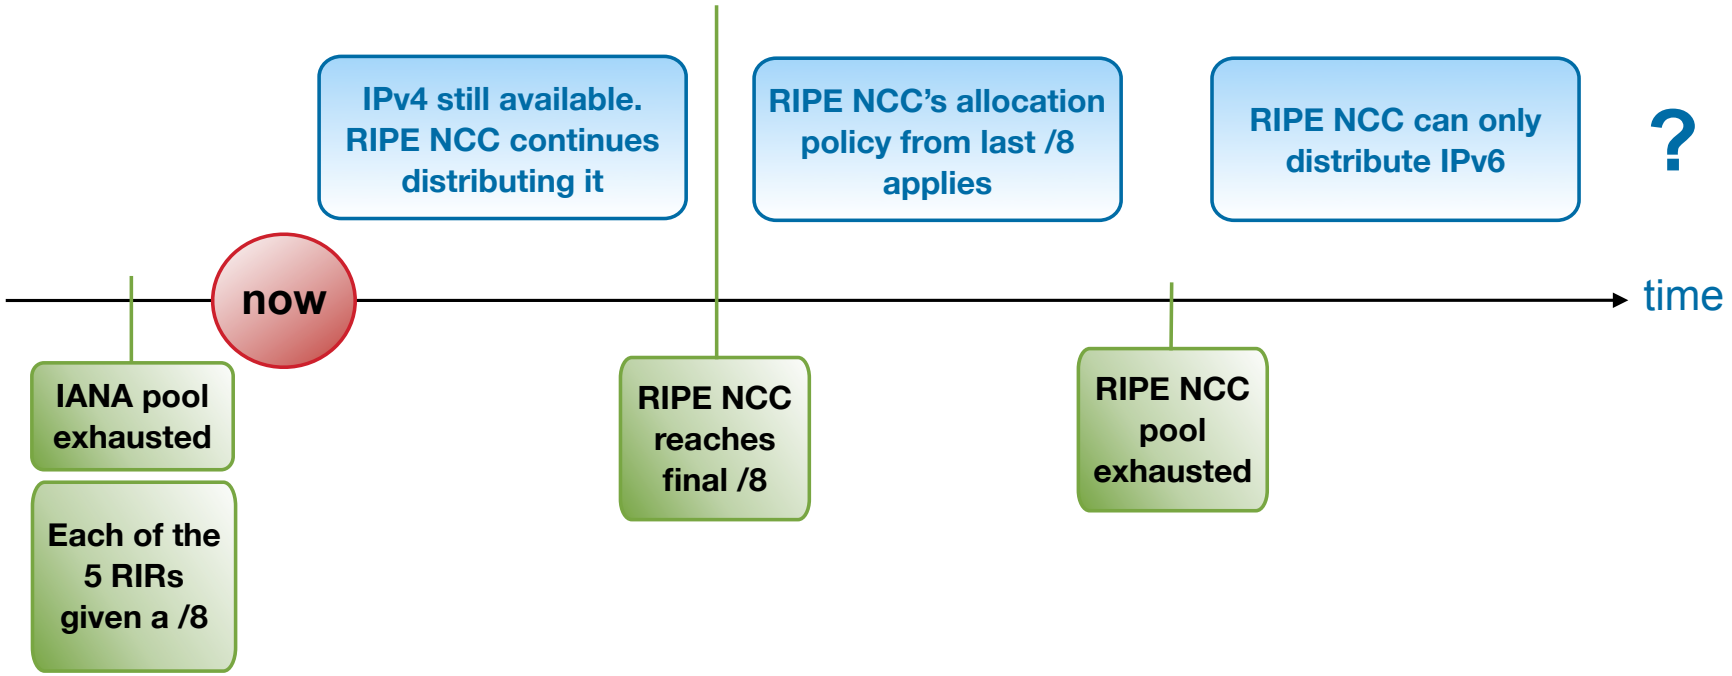

Run Out Fairly (of IPv4)

- Gradually reduced allocation / assignment periods
- Needs for “Entire Period” of up to...
  - 12 months (January 2010)
  - 9 months (July 2010)
  - 6 months (January 2011)
  - 3 months (July 2011)
- 50% has to be used up by half-period

Network Address Translation

- Extends the capacity of the IPv4 address space by sharing an IPv4 address between clients
- Fairly common technology, used everywhere
- Breaks the end to end connectivity model
- **It doesn't allow communication with IPv6!**
- You are probably going to need it in some form

Large Scale NAT

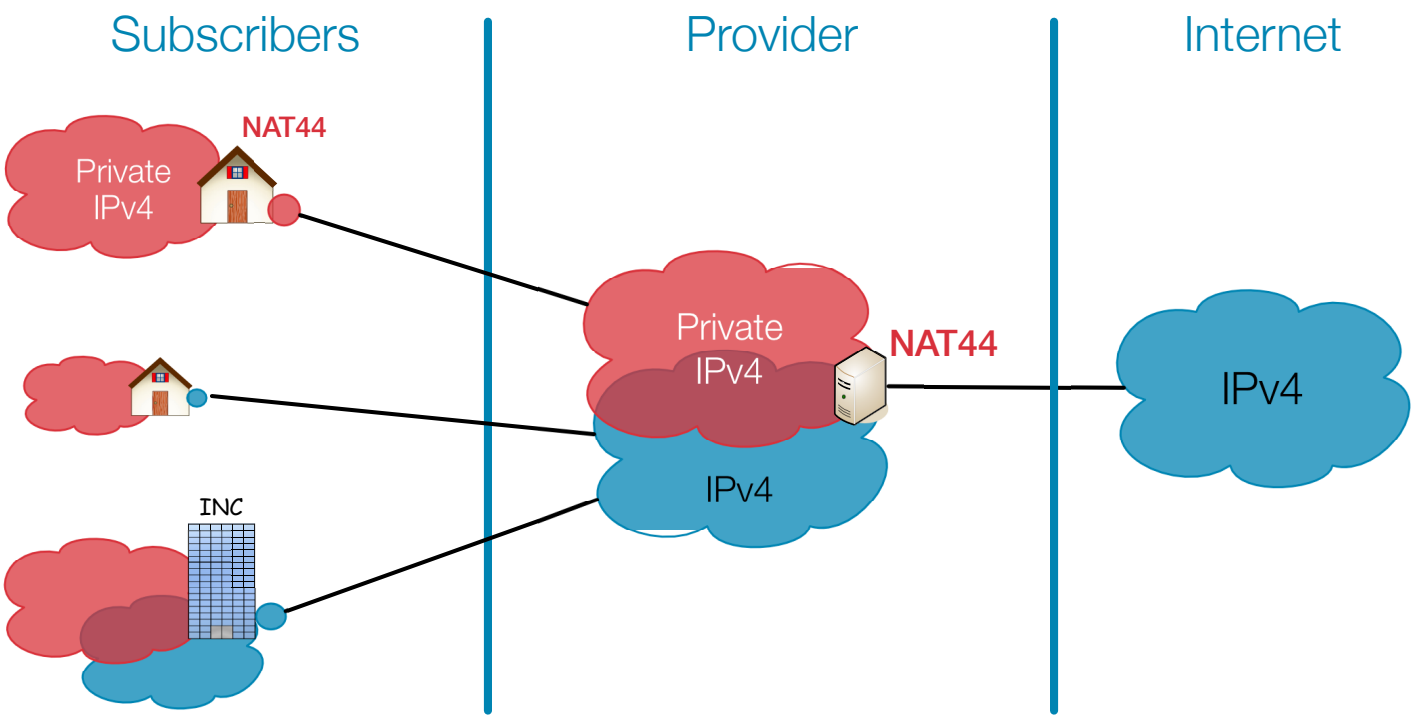

Questions?

RIPE NCC

The road to IPv6

Section 3

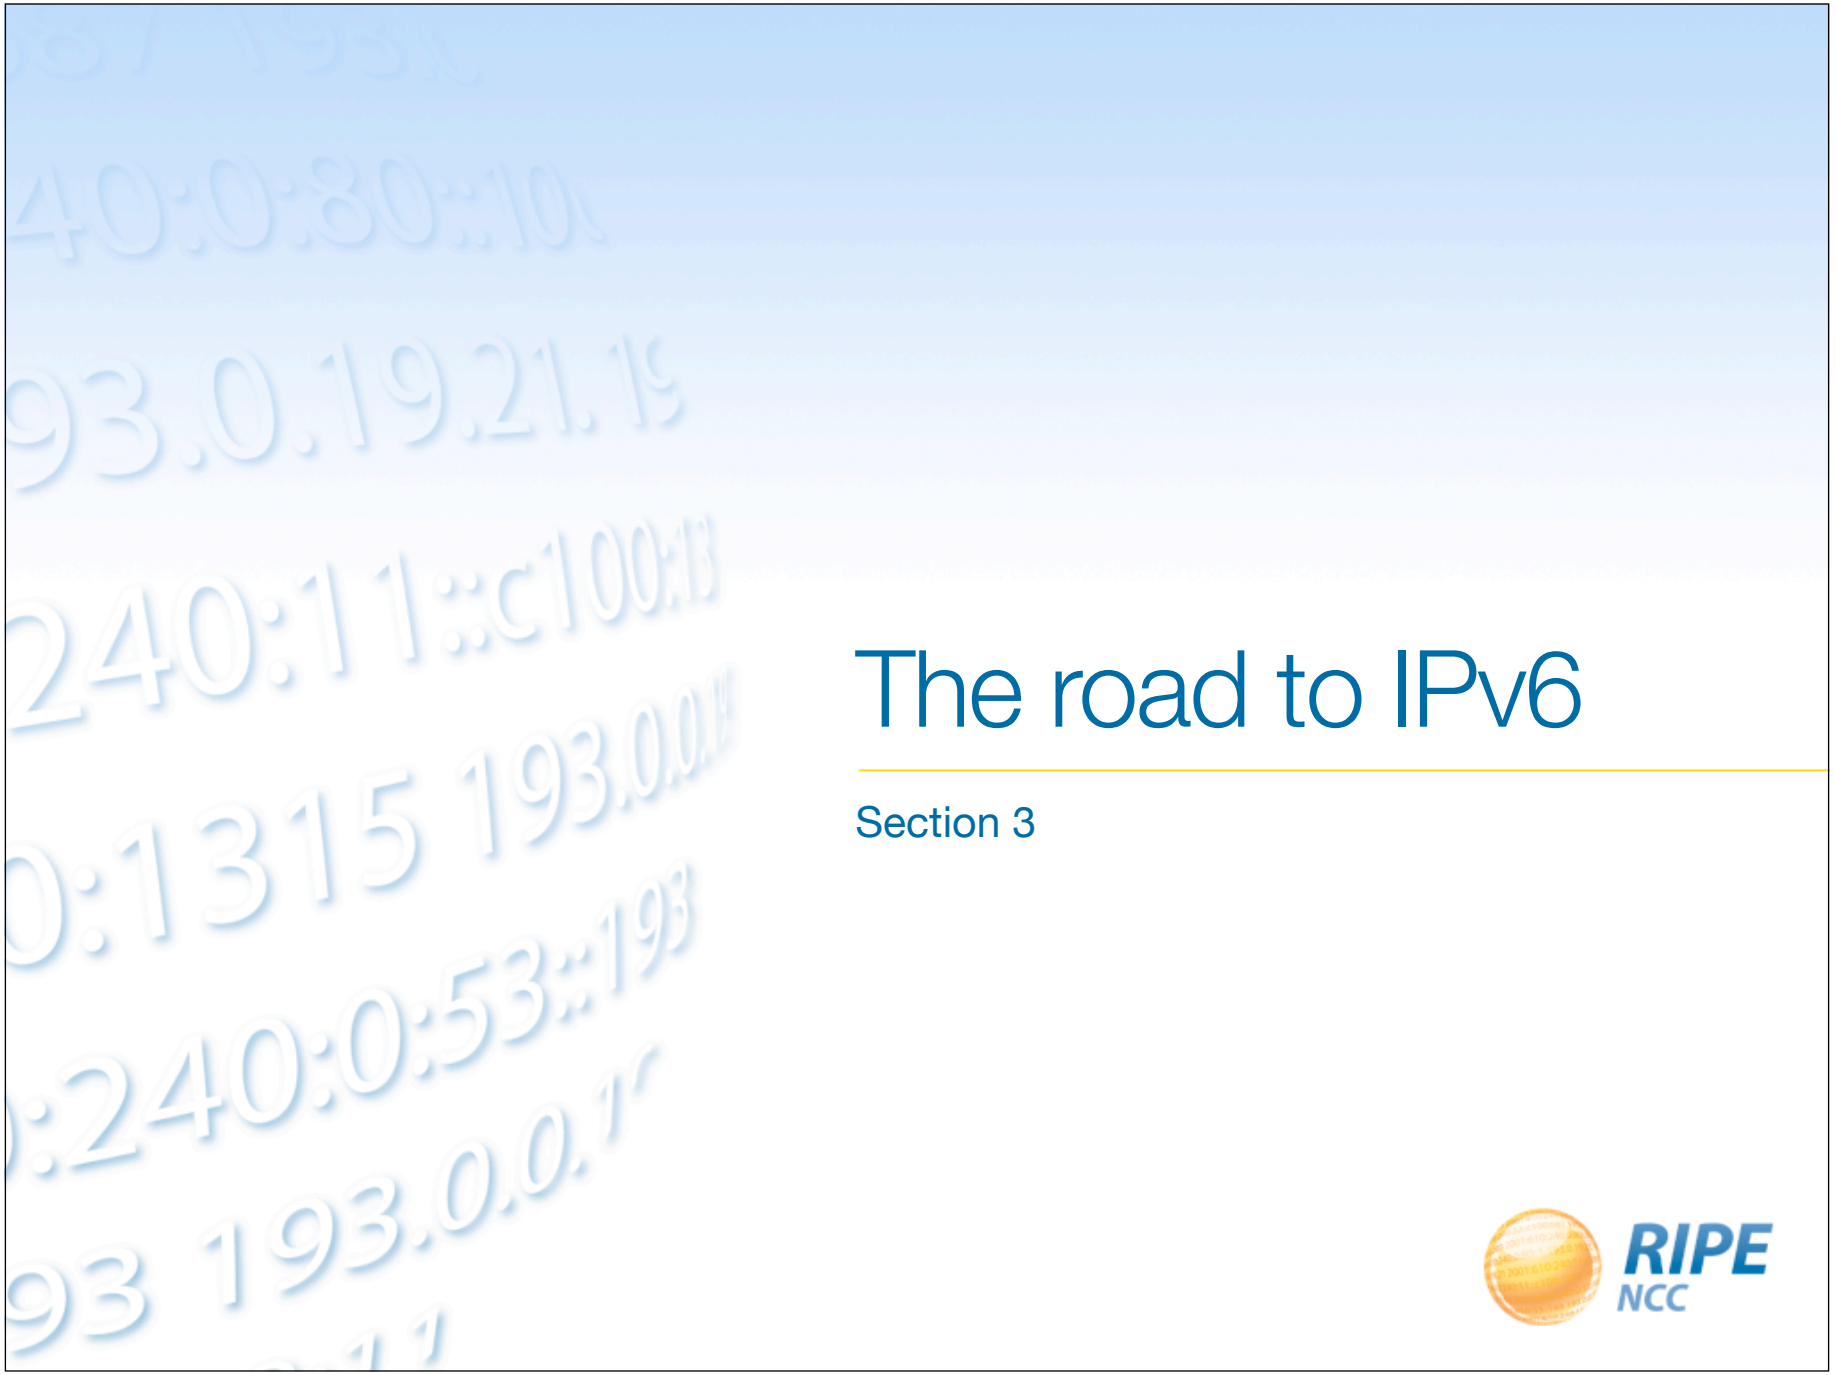

There Was a Plan

- The original idea was to have IPv6 deployed before we were out of IPv4 addresses
- By now the whole of the Internet should have been dual-stacked
- And we wouldn't be here today.

IPv6 is the End Goal

- Exhaustion of the IPv4 free pool is a permanent problem
- The only way to support the future growth of the Internet is by deploying IPv6
- This will take time, so an intermediate solution has to be found
- Eventually, be prepared to switch off IPv4

Dual Stack while you can

Transitioning Mechanisms

- The IETF has several RFCs and active drafts, and some that have been abandoned already:

6in4	NAT64
6to4	DS-lite
Teredo	A+P
6RD	4RD
ISATAP	SIIT
TSP	TRT
6over4	NAT-PT
IVI	....

Solving Two Problems

- Maintaining connectivity to IPv4 hosts by sharing IPv4 addresses between clients
  - Extending the address space with NAT/CGN/LSN
  - Translating between IPv6 and IPv4
- Provide a mechanism to connect to the emerging IPv6-only networks
  - Tunneling IPv6 packets over IPv4-only networks

Questions?

RIPE NCC

IPv6
Address Basics

Section 4

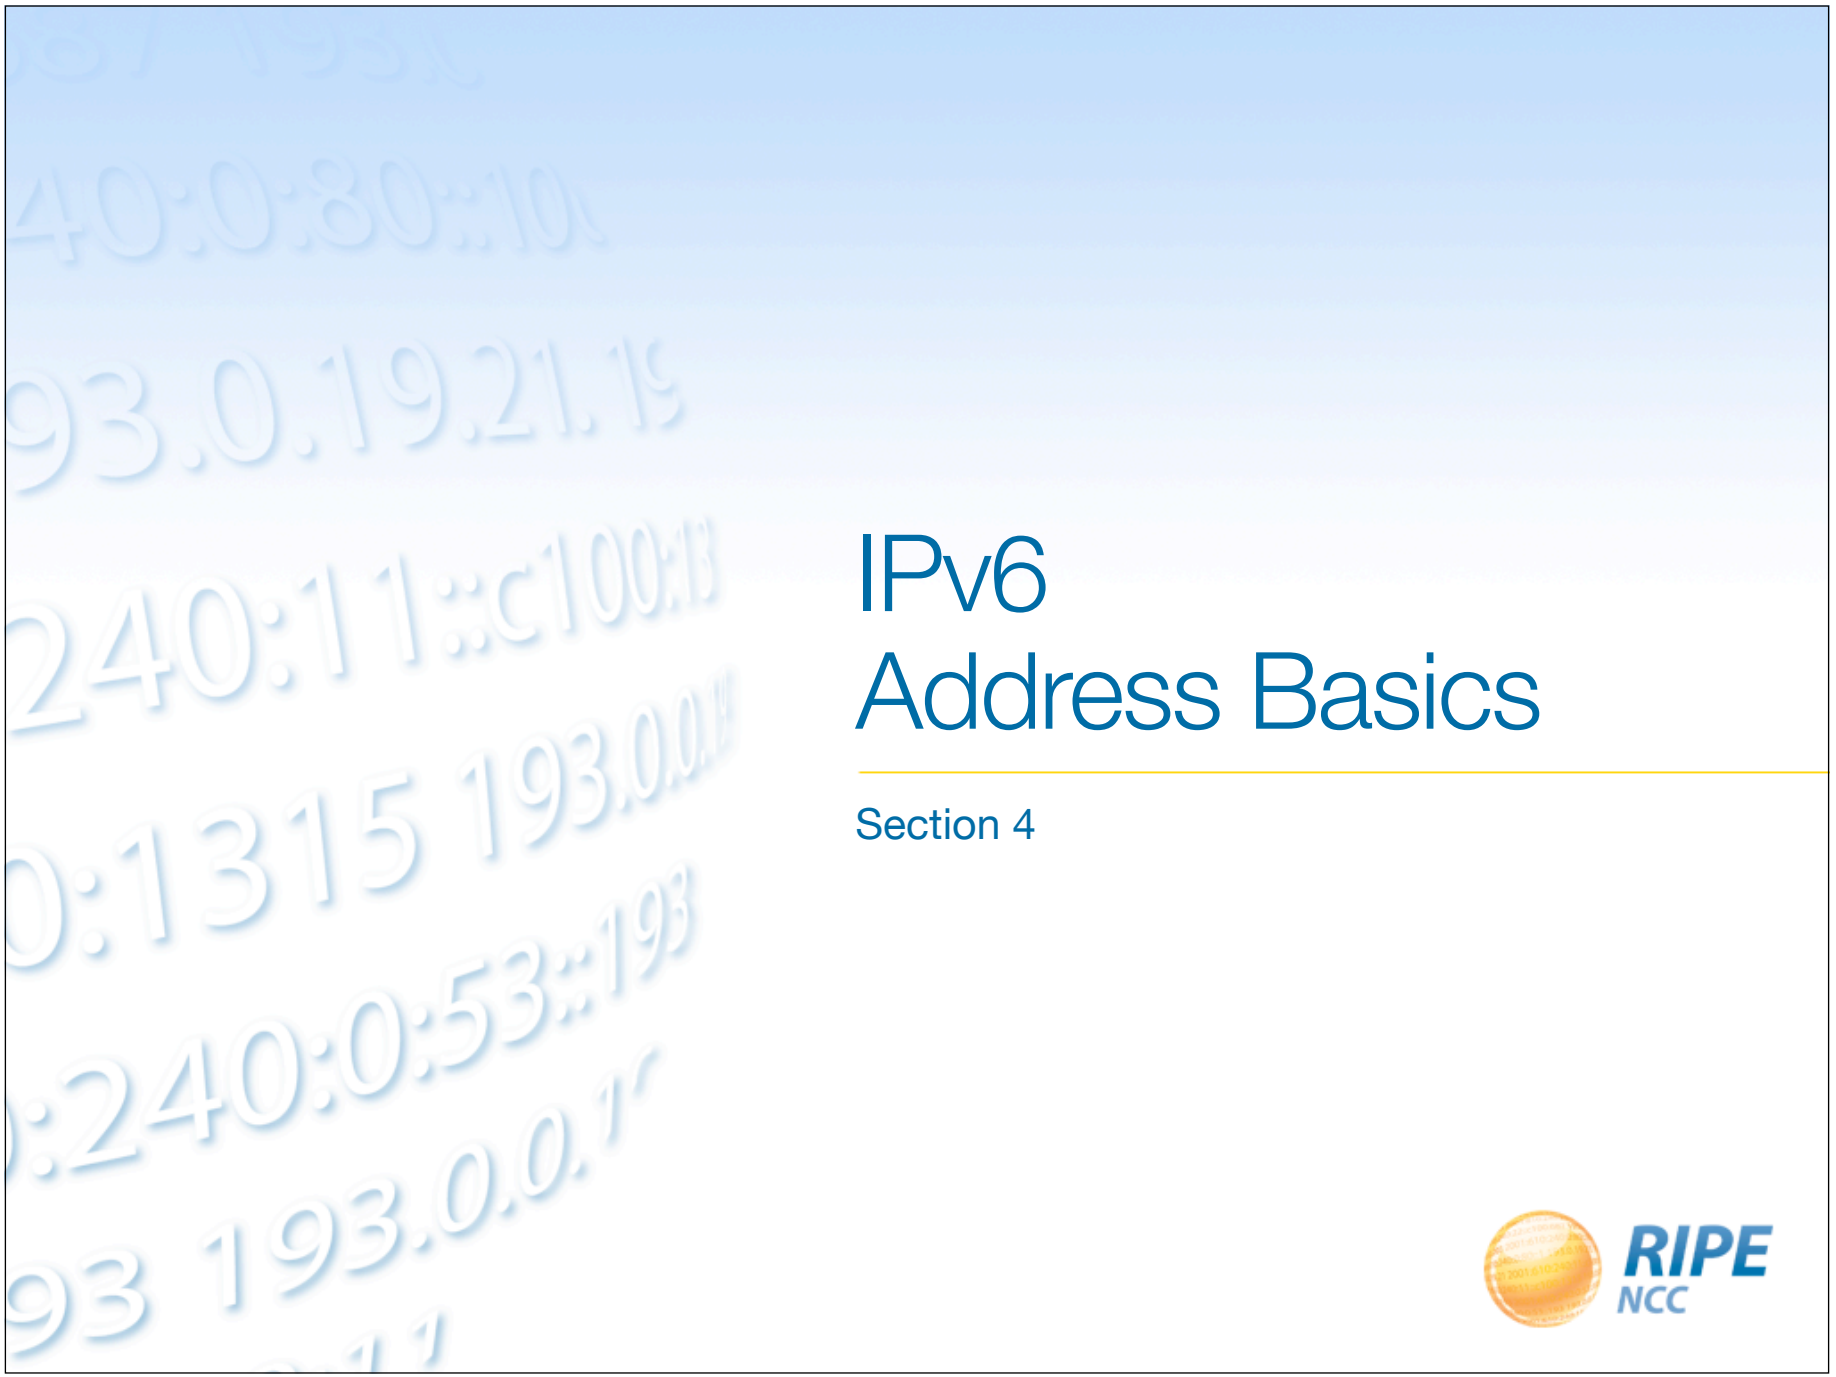

IP Address Distribution

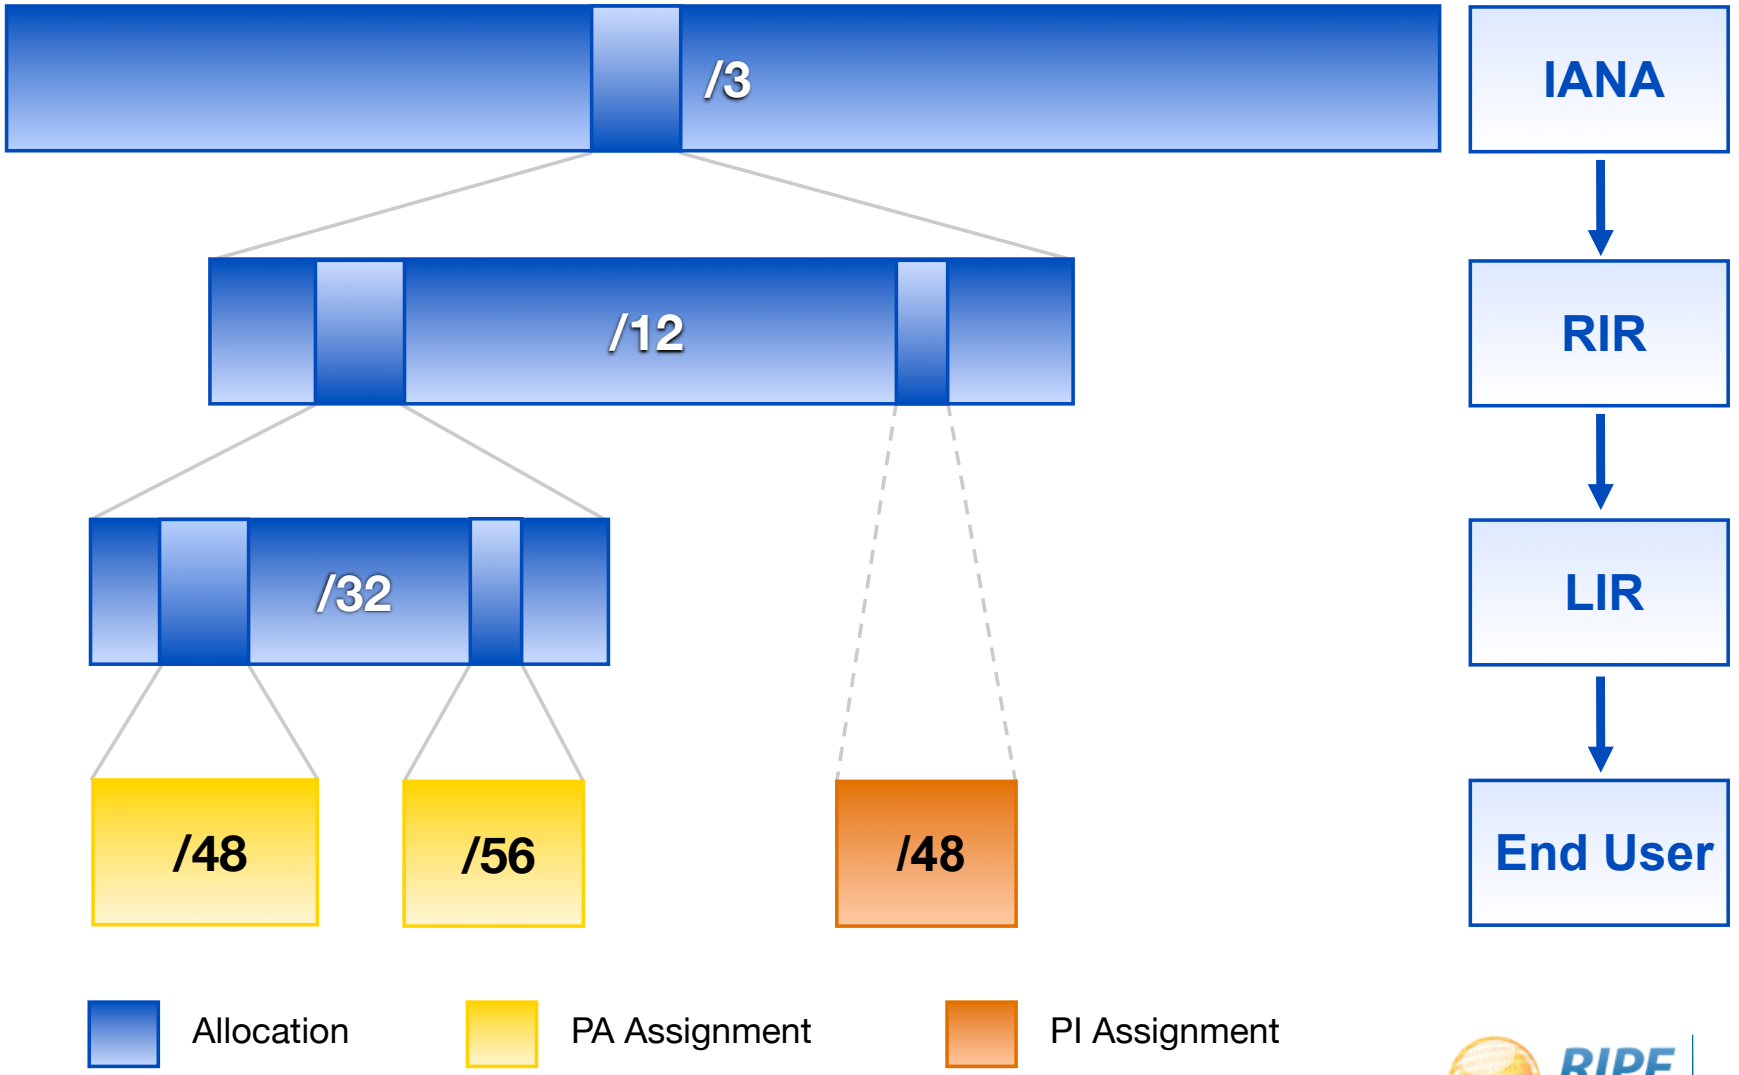

IPv6 Address Basics

- IPv6 address: 128 bits
  - 32 bits in IPv4
- Every subnet should be a /64
- Customer assignments (sites) between:
  - /64 (1 subnet)
  - /48 (65,536 subnets)
- Minimum allocation size /32
  - 65,536 /48s
  - 16,777,216 /56s

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

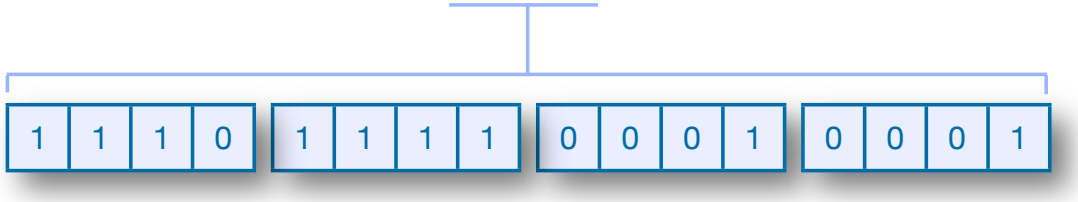

Multiple addresses

Addresses	Range	Scope
Loopback	::1	host
Link Local	fe80::/10	link
Unique Local	fc00::/7	global
Global Unicast	2000::/3	global
6to4	2002::/16	global
Teredo	2001::/32	global
Multicast	ff00::/8	variable

Exercise

IPv6 Address Notation

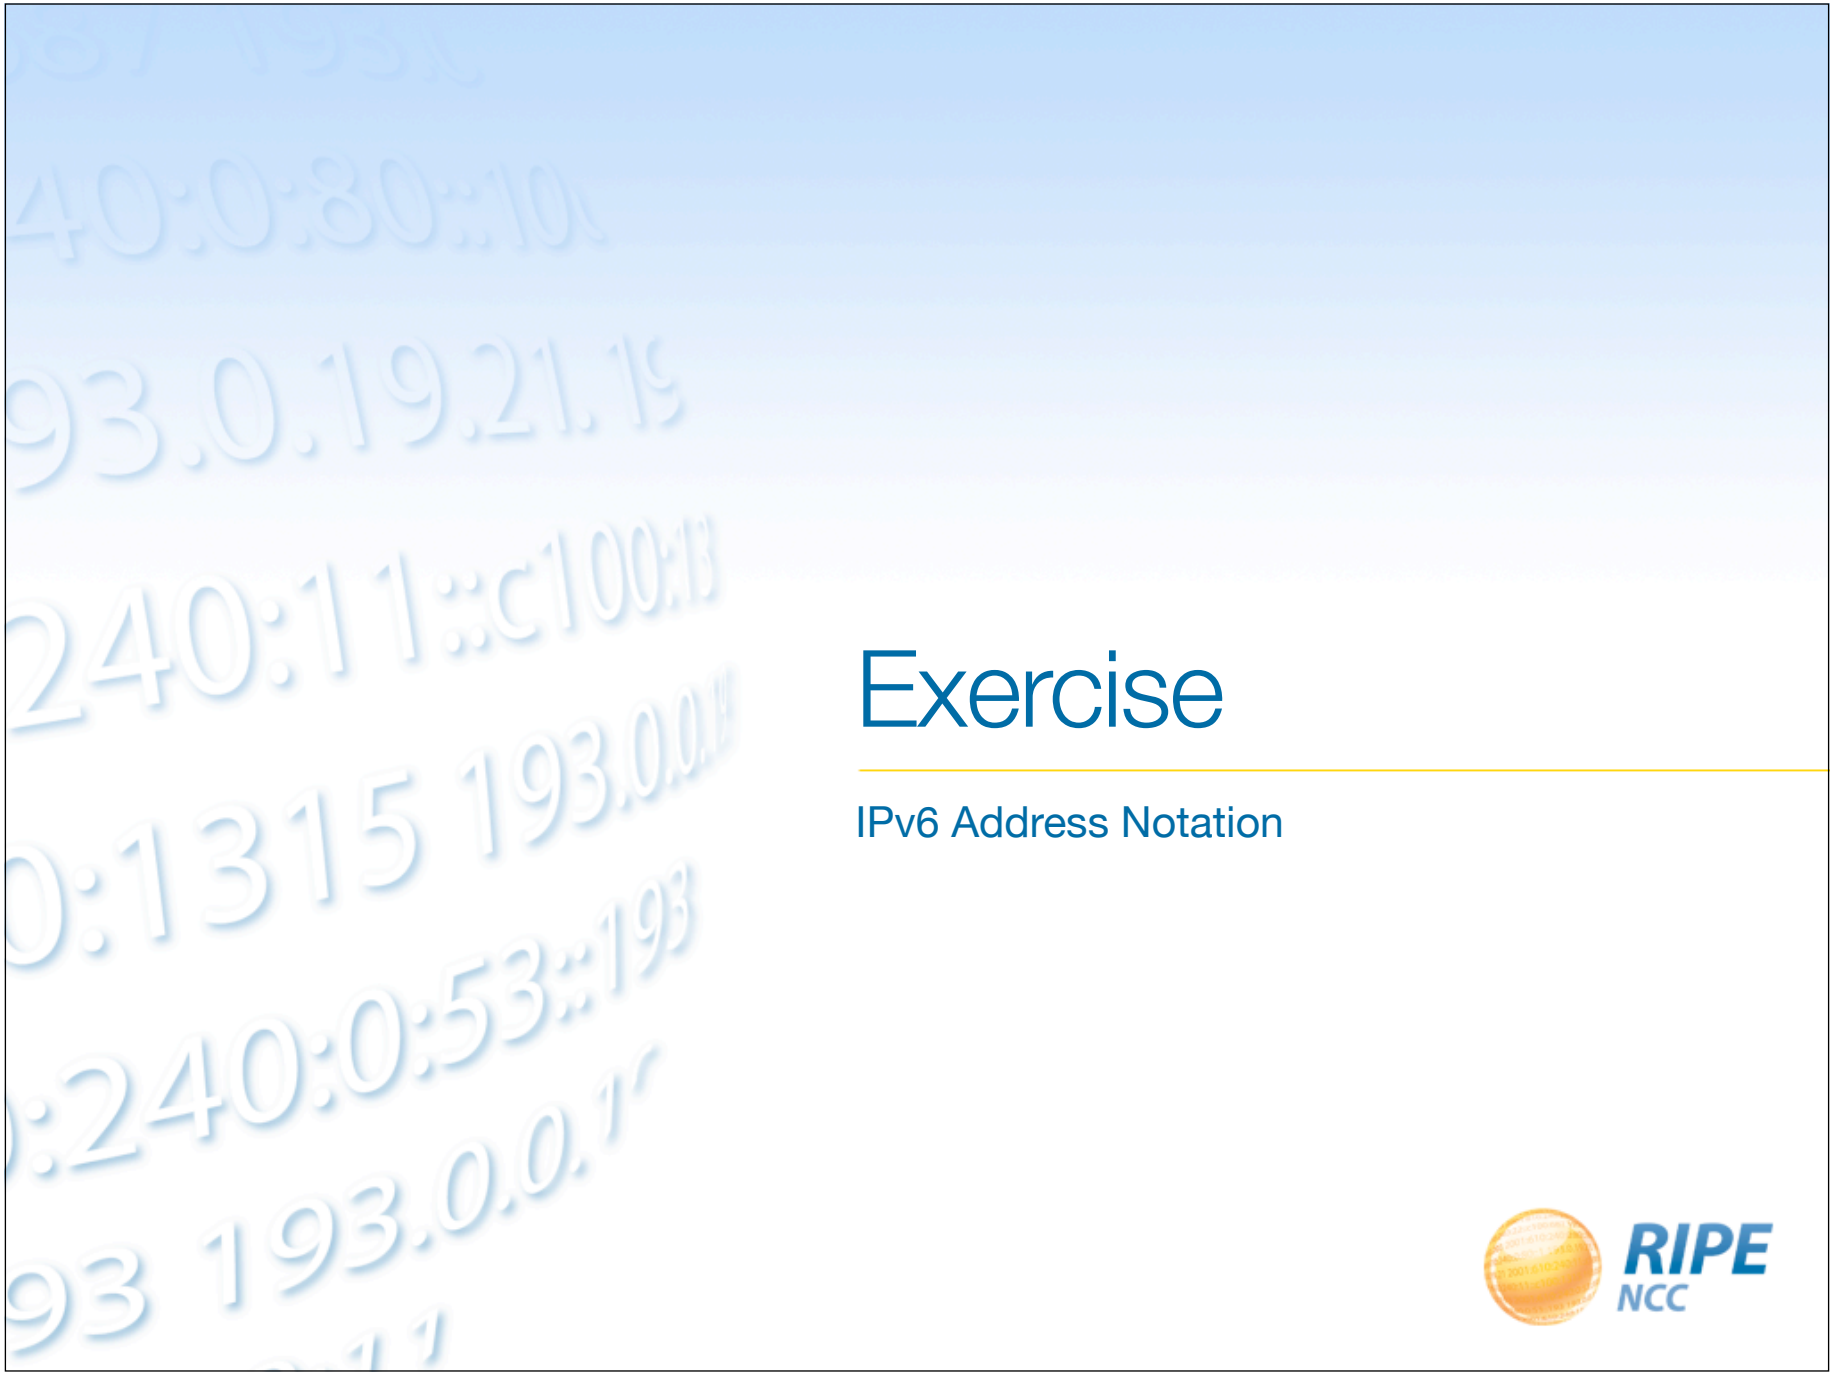

IPv6 Stateless Address Autoconfiguration

- Host will automatically start looking for a router
- Response will contain:
  - Router's address
  - One or more link prefixes
  - SLAAC allowed yes/no
  - MTU

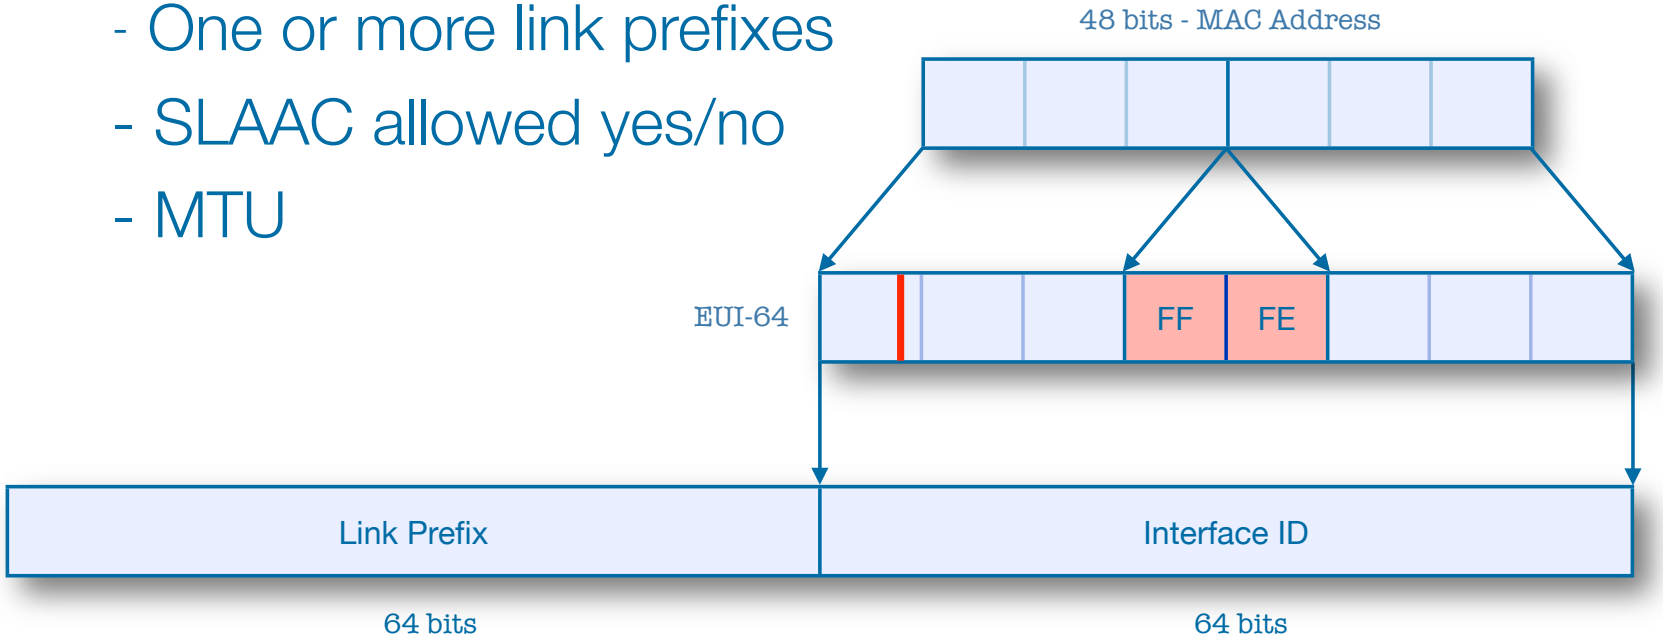

DHCPv6

- You can use DHCPv6 to get information like DNS servers
- Router message contains hints
  - If a DHCPv6 server is present
  - If the use of DHCPv6 is mandatory to get an address, the so called “managed” flag
  - Optionally the address of a DNS server (RFC 6106)
- With manual configuration subnet sizes other than /64 are possible, but please think twice

DNS in IPv6 is difficult?

- DNS is not IP layer dependent
- A record for IPv4
- AAAA record for IPv6

- Don't answer based on incoming protocol
- Only challenges are for translations
  - NAT64, proxies

Reverse DNS

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

Reverse DNS

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

. . . . . .ip6.arpa

d.4.0.0.0.0.1.c.0.0.0.0.0.0.0.0.0.1.1.f.e.e.
3.0.0.8.b.d.0.1.0.0.2.ip6.arpa PTR
yourname.domain.tld.

d.4.0.0.0.0.1.c.0.0.0.0.0.0.0.0.0.1.1.f.e.e.3.0.0.8.b.d.0.1.0.0.2.ip6.arpa PTR yourname.domain.tld.

Questions?

RIPE NCC

IPv6 Ripeness

Section 5

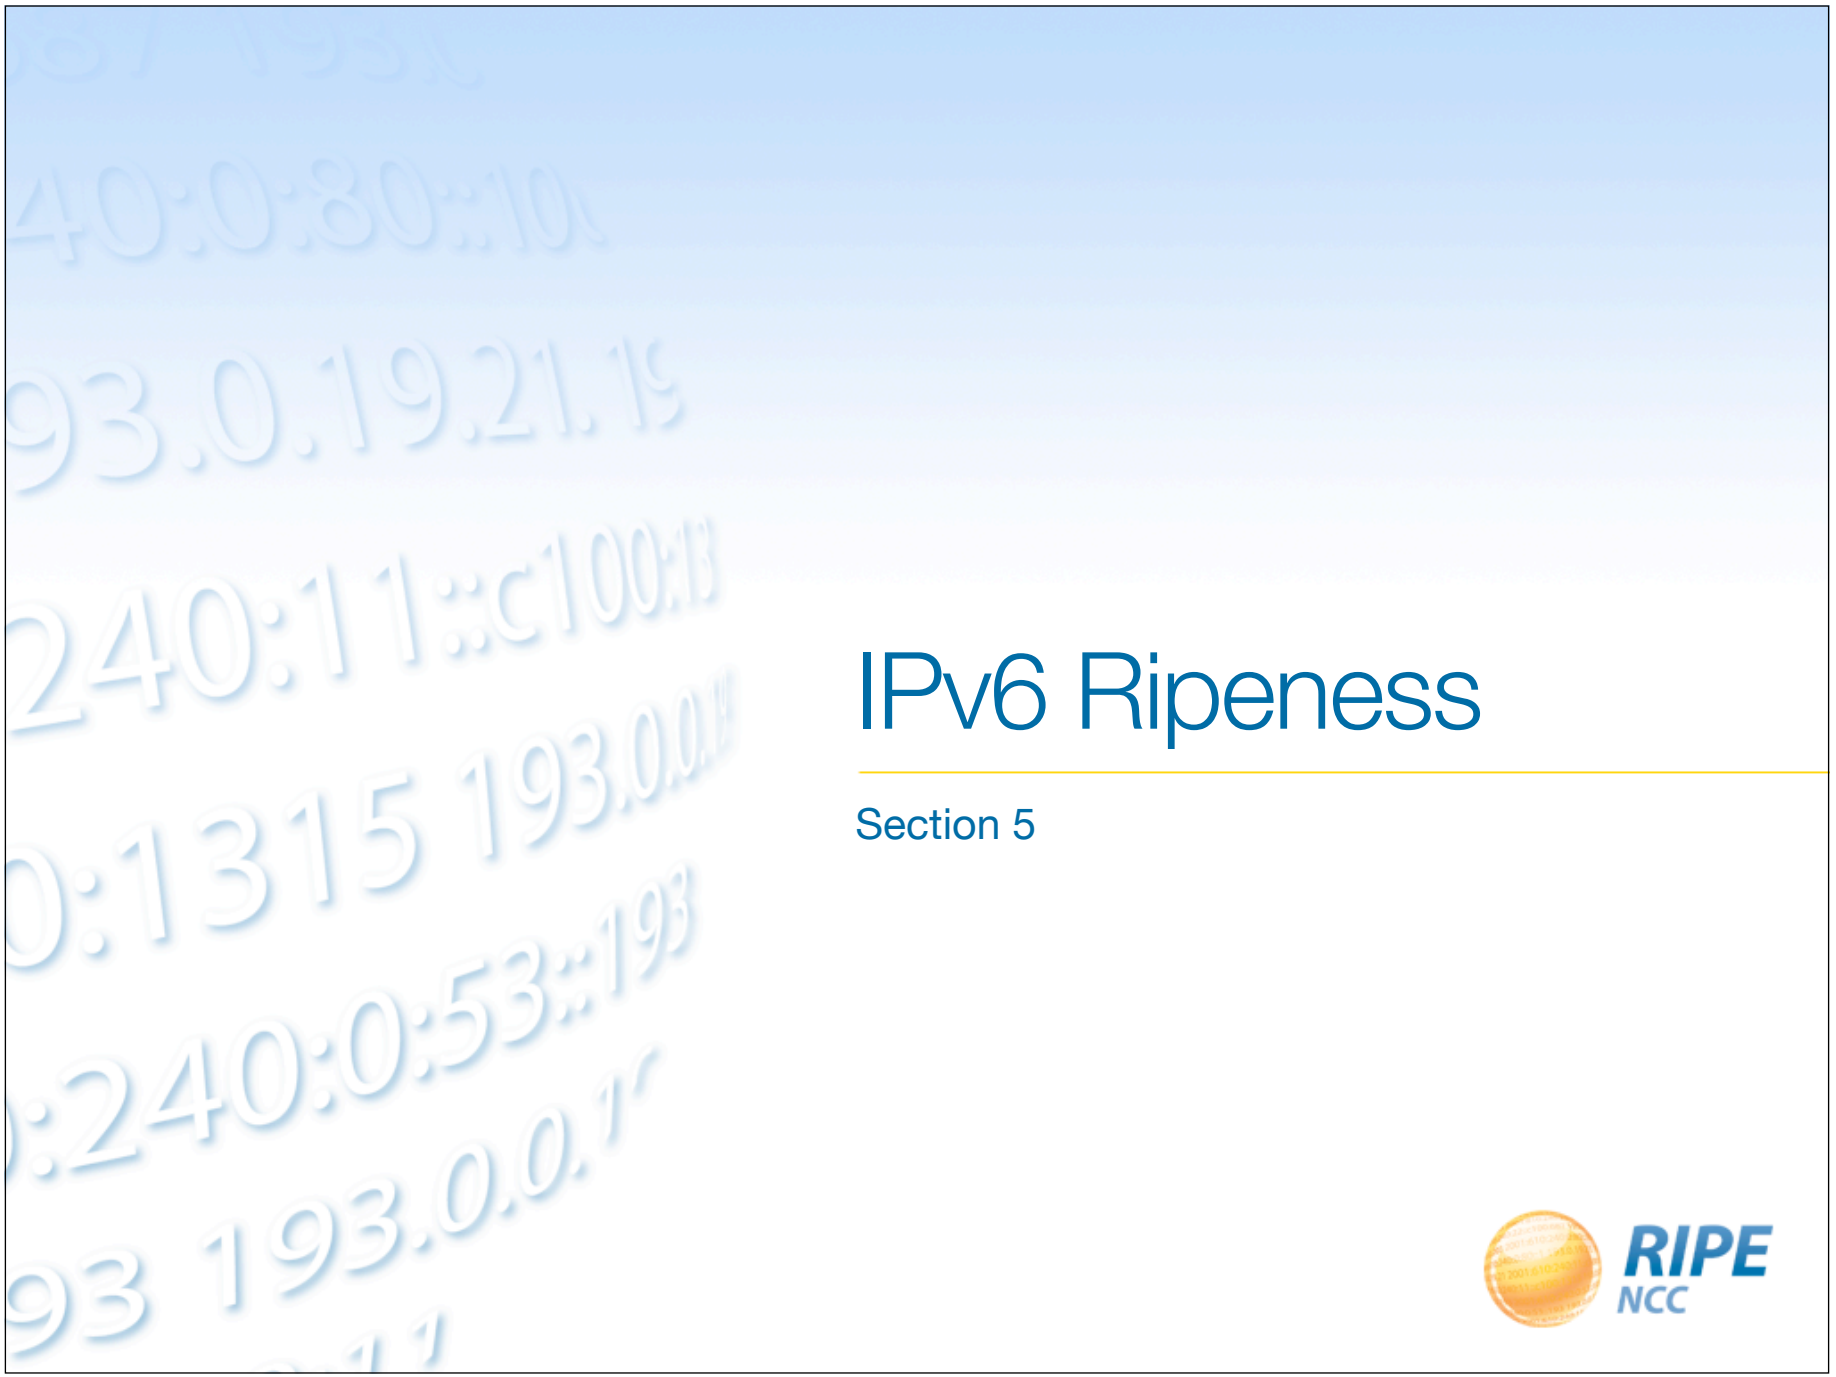

IPv6 Ripeness

- 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/

IPv6 RIPEness: 8154 LIRs

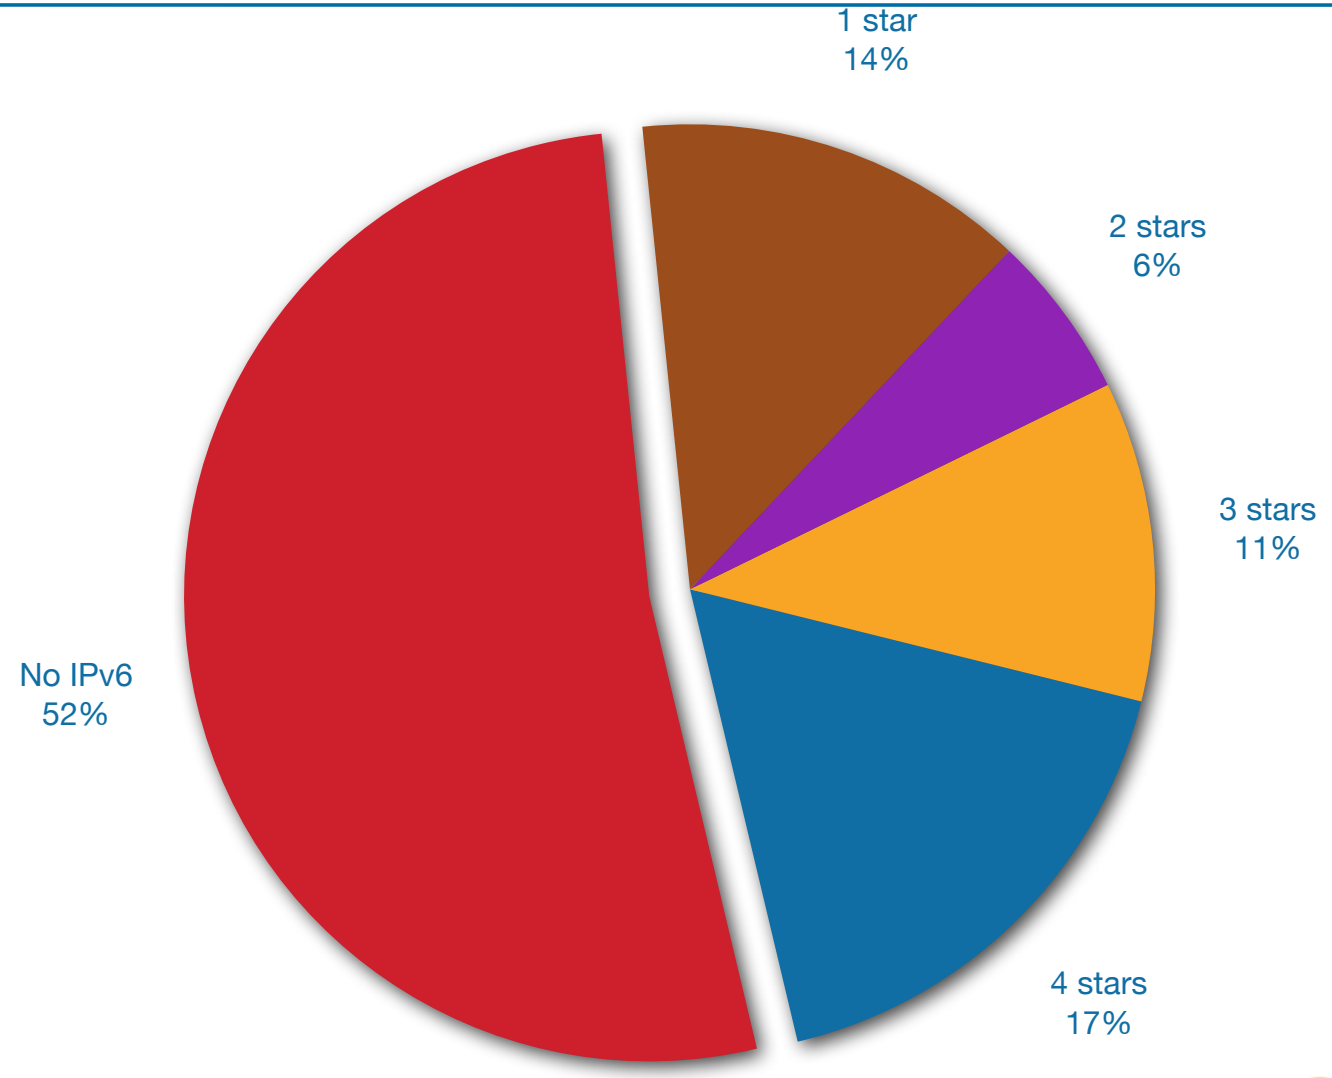

IPv6 RIPEness - MENOG Region: 583 LIRs

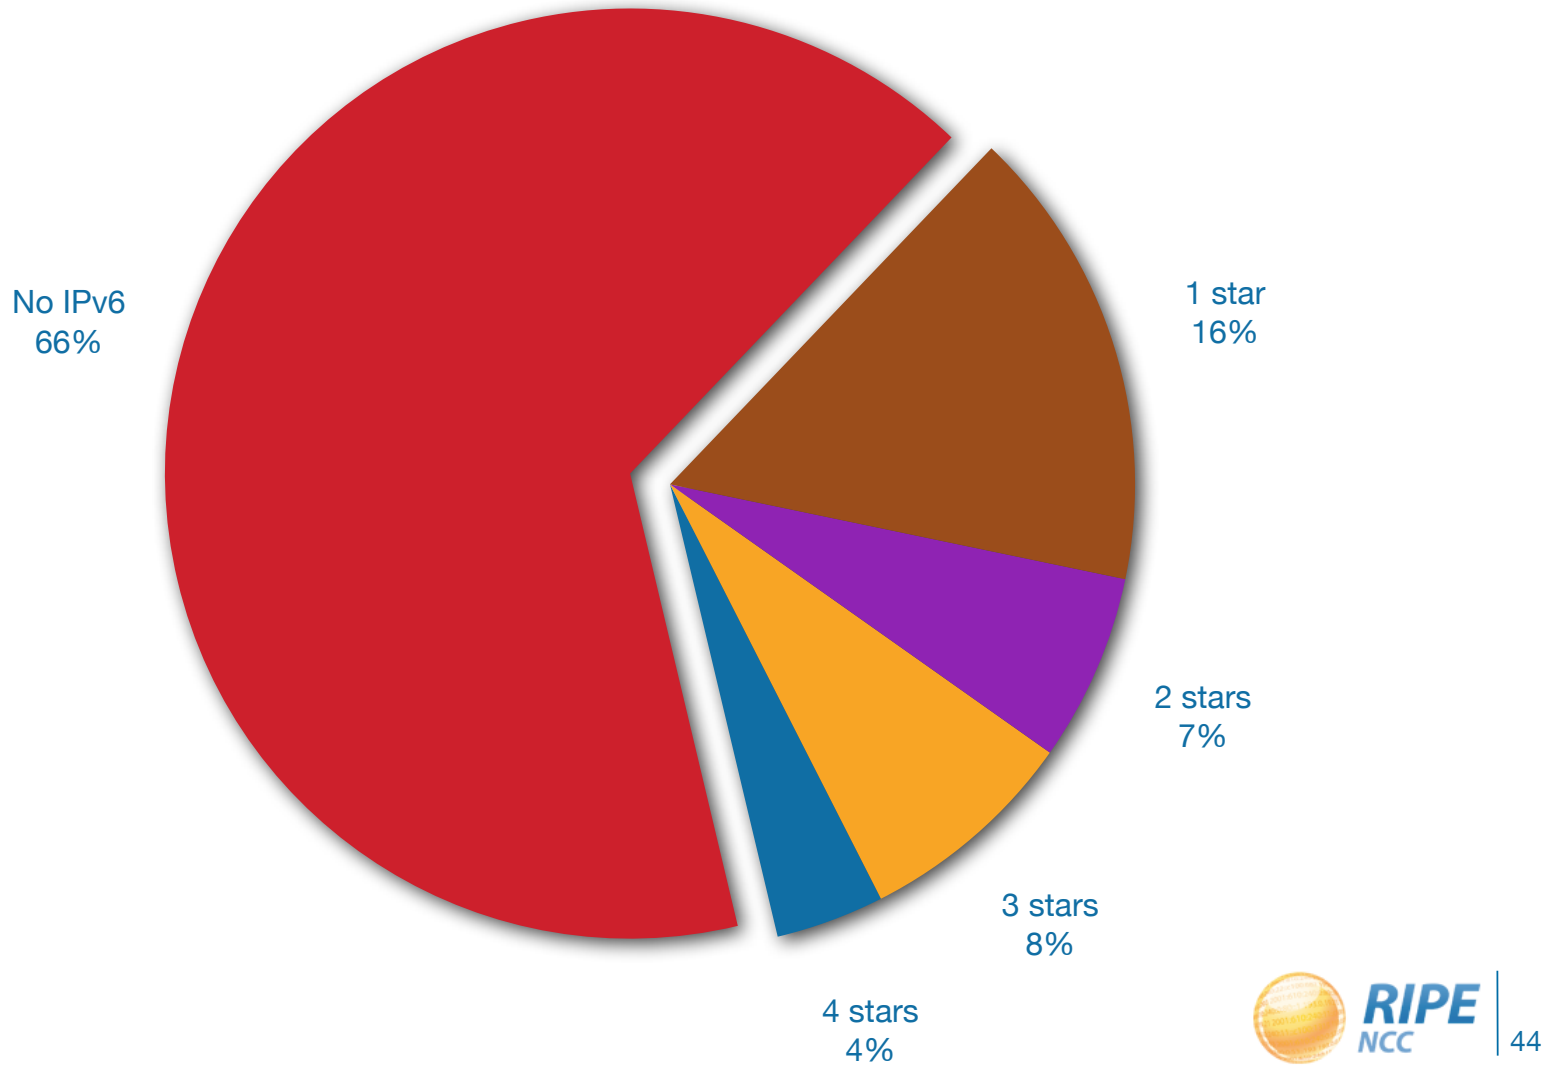

IPv6 Enabled Networks

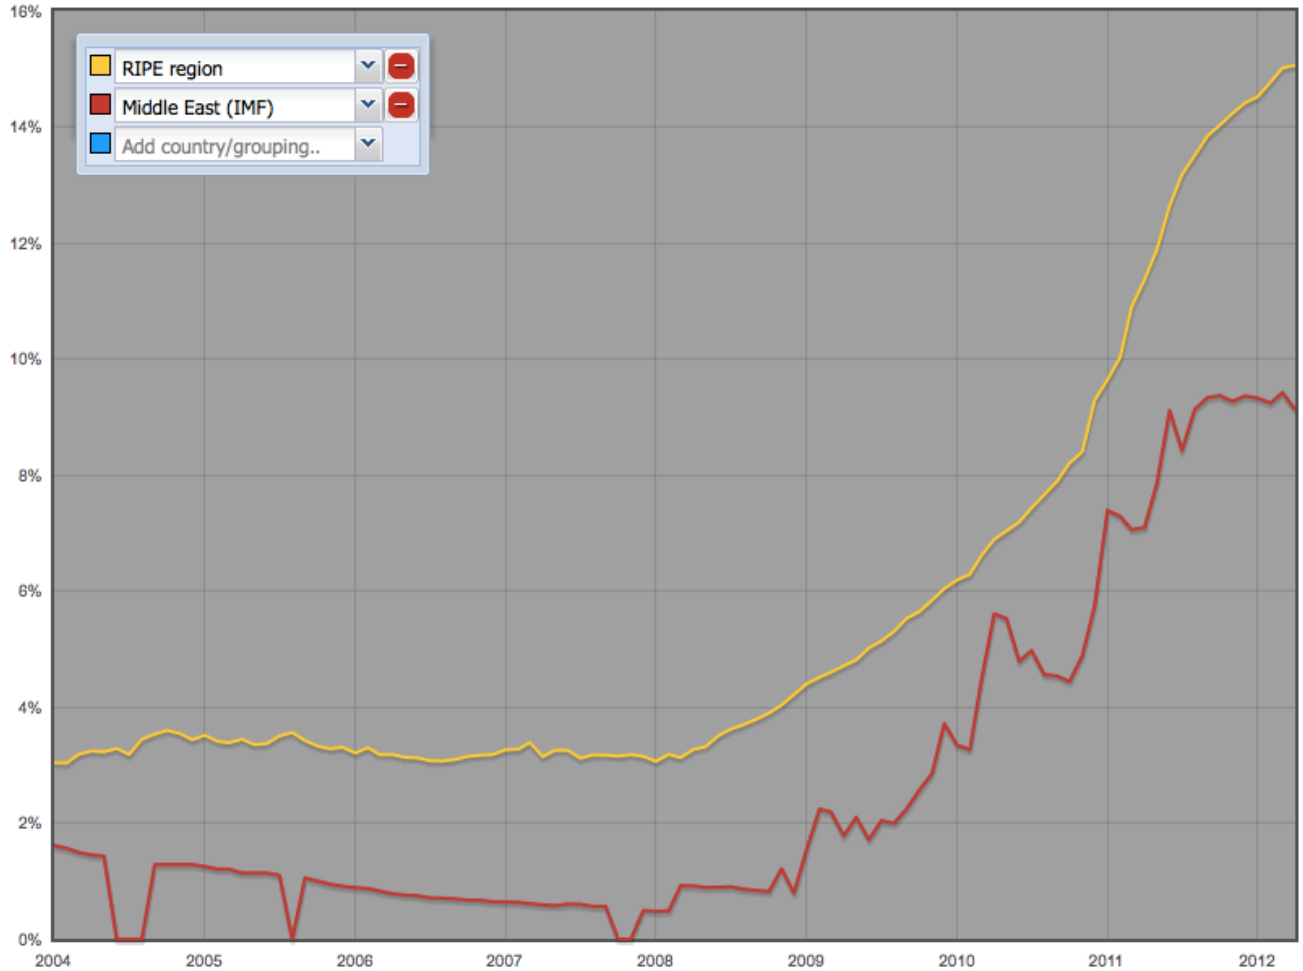

Transition Mechanisms

Section 6

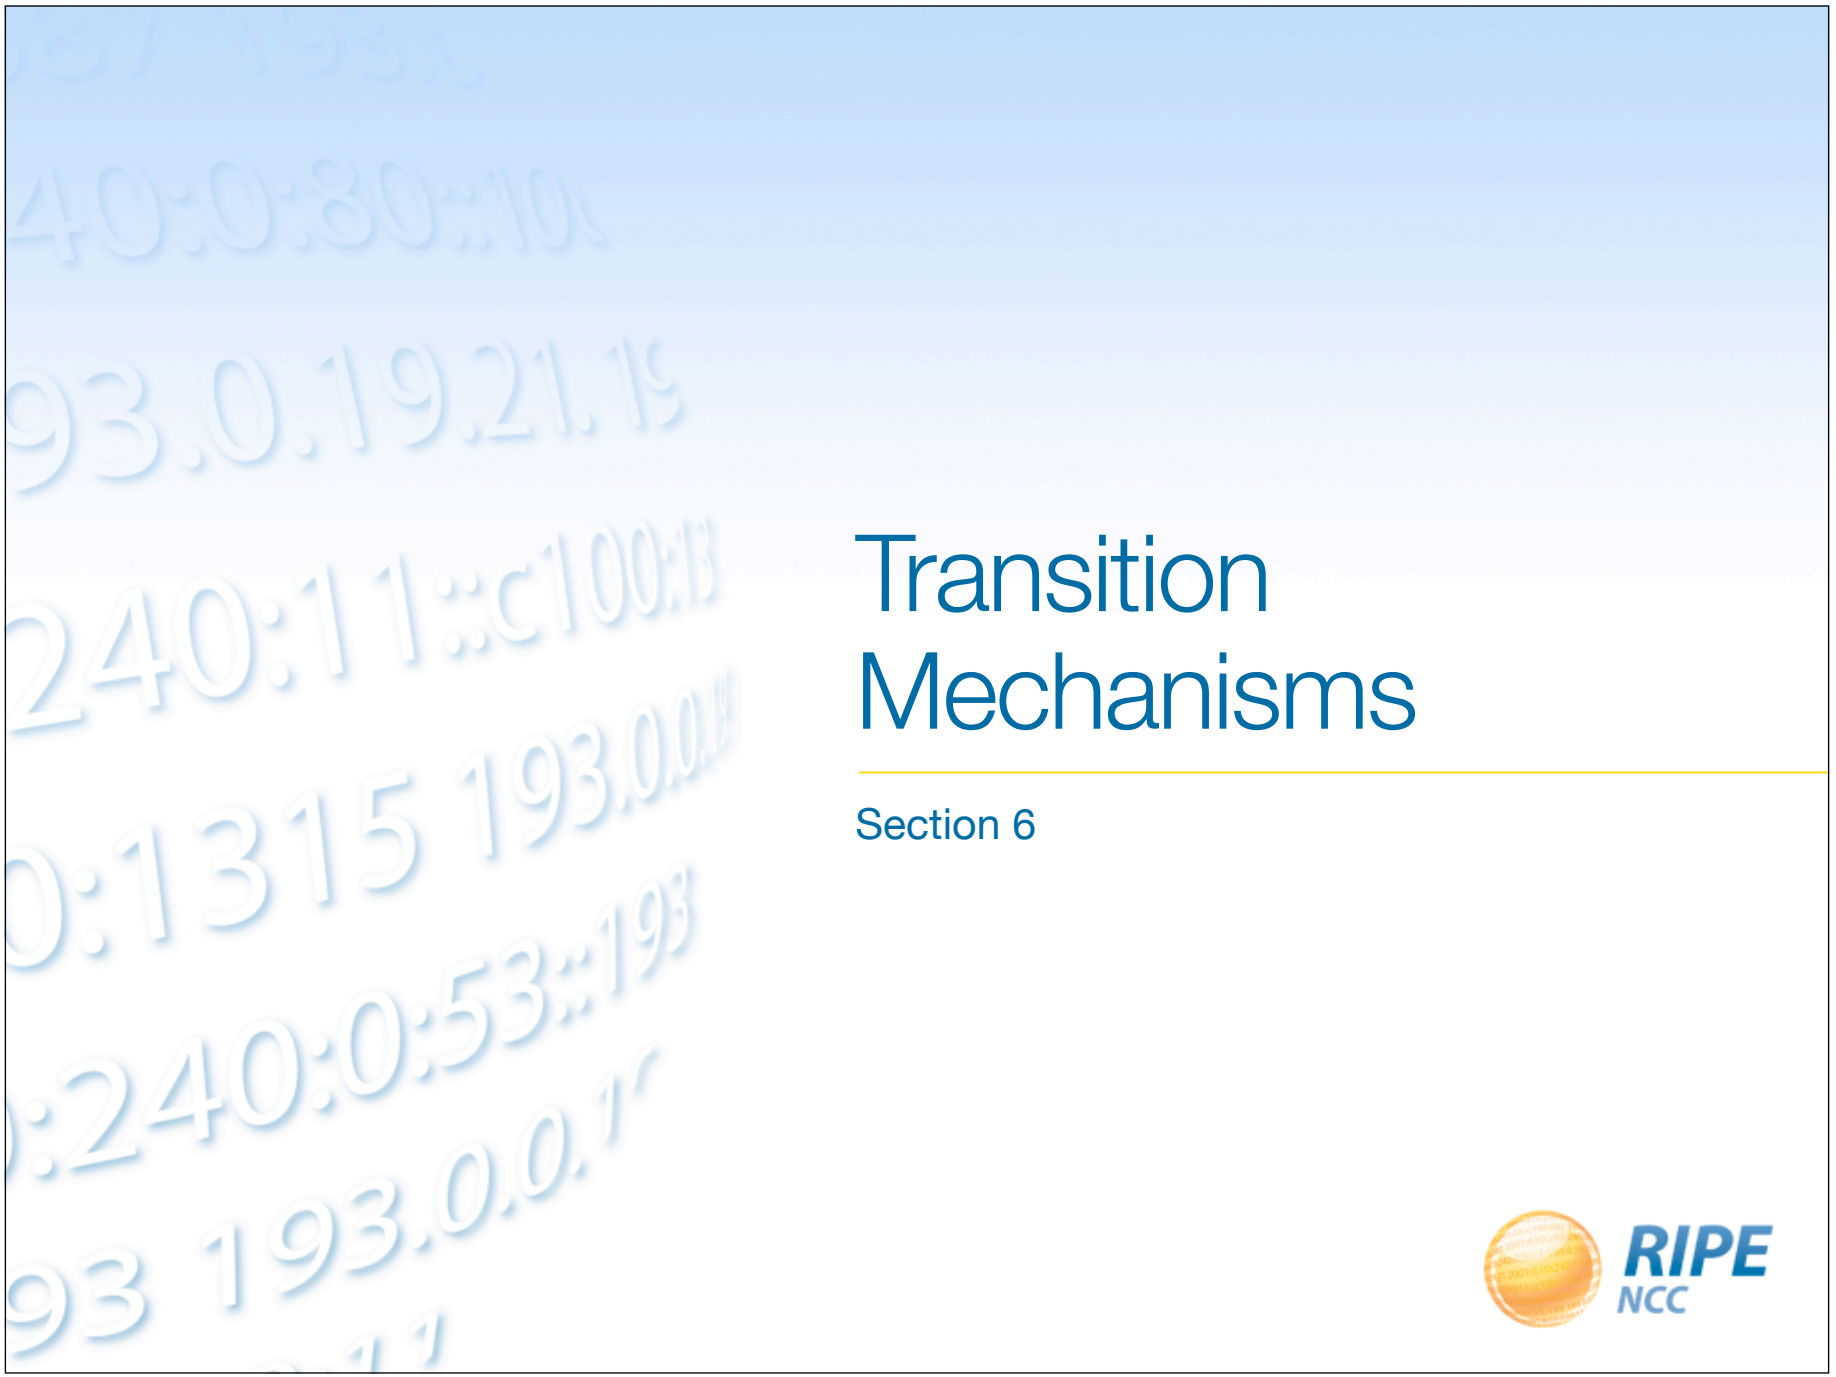

Transitioning: Two Main Methods

- Transporting IPv6 in IPv4
  - 6in4
  - 6to4
  - Teredo
  - 6RD
- Translating IPv6 into IPv4
  - NAT64/DNS64

6in4

- Manually configured tunnels towards a fixed tunnel broker like SixXS, Hurricane Electric or your own system
- Stable and predictable but not easily deployed to the huge residential markets
- MTU might cause issues

6in4

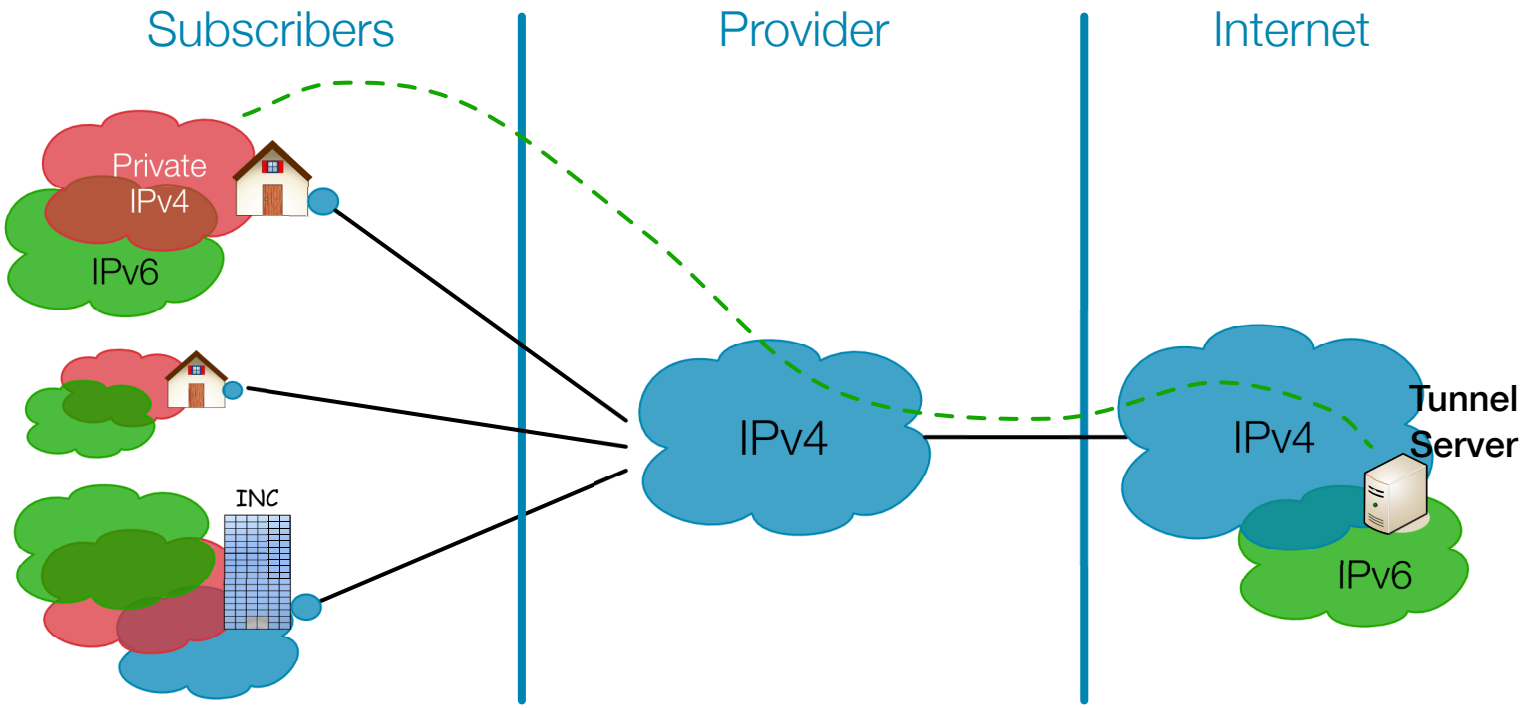
Subscribers
Provider
Internet
Private
IPv4
IPv6
INC
IPv4
IPv4
Tunnel
Server
IPv6

6to4 and Teredo

- 6to4
  - "Automatic" tunnel, system can configure itself
  - IPv4 address is part of the IPv6 address
  - Requires a public IPv4 address
  - Uses anycast to reach a nearby server
  - Return traffic might choose another server
- Teredo
  - Uses UDP to encapsulate packets
  - Works across (most) NAT implementations

6to4 and Teredo

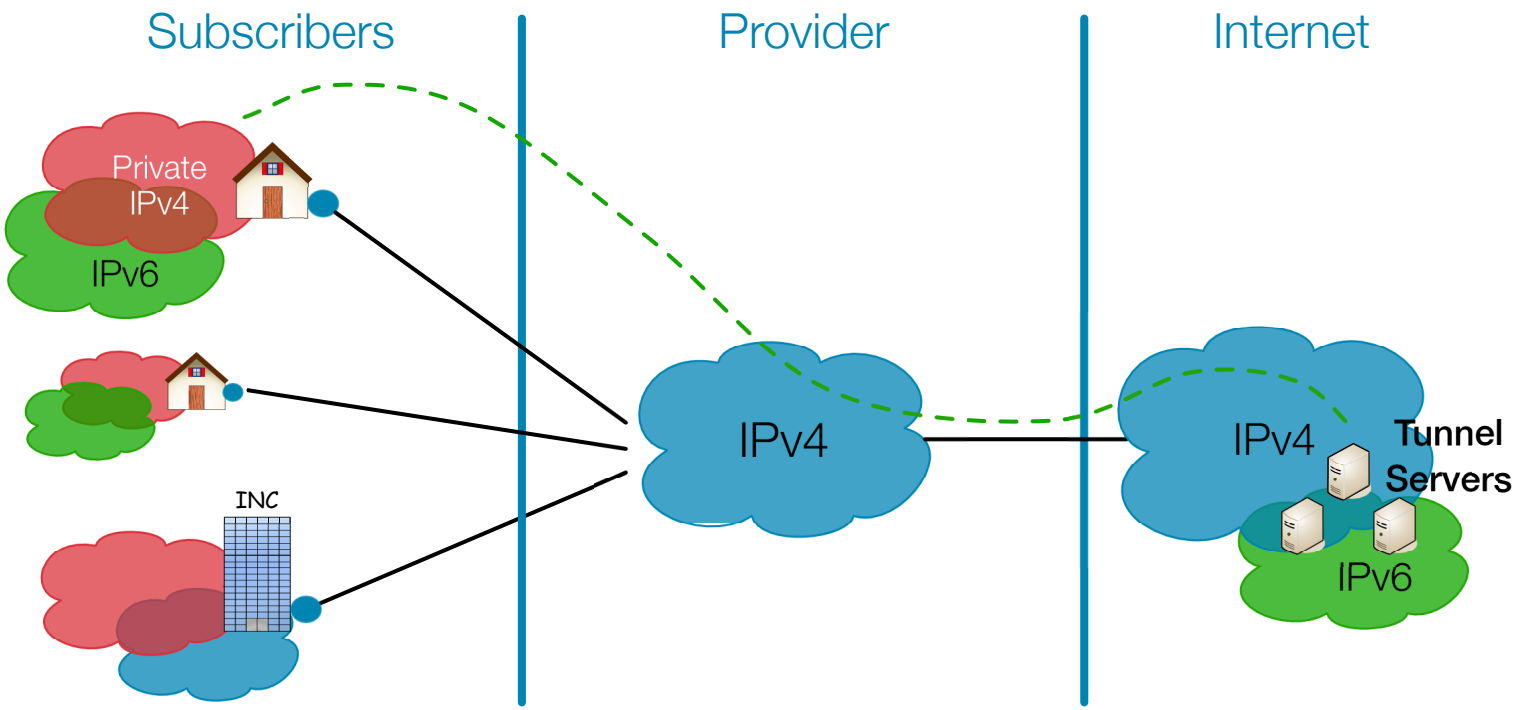

6RD

- Quite similar to 6to4
  - Encodes the IPv4 address in the IPv6 prefix
- Uses address space assigned to the operator
- The operator has full control over the relay
- Traffic is symmetric across a relay
  - Or at least stays in your domain
- Can work with both public and private space
- Needs additional software for signaling

6RD

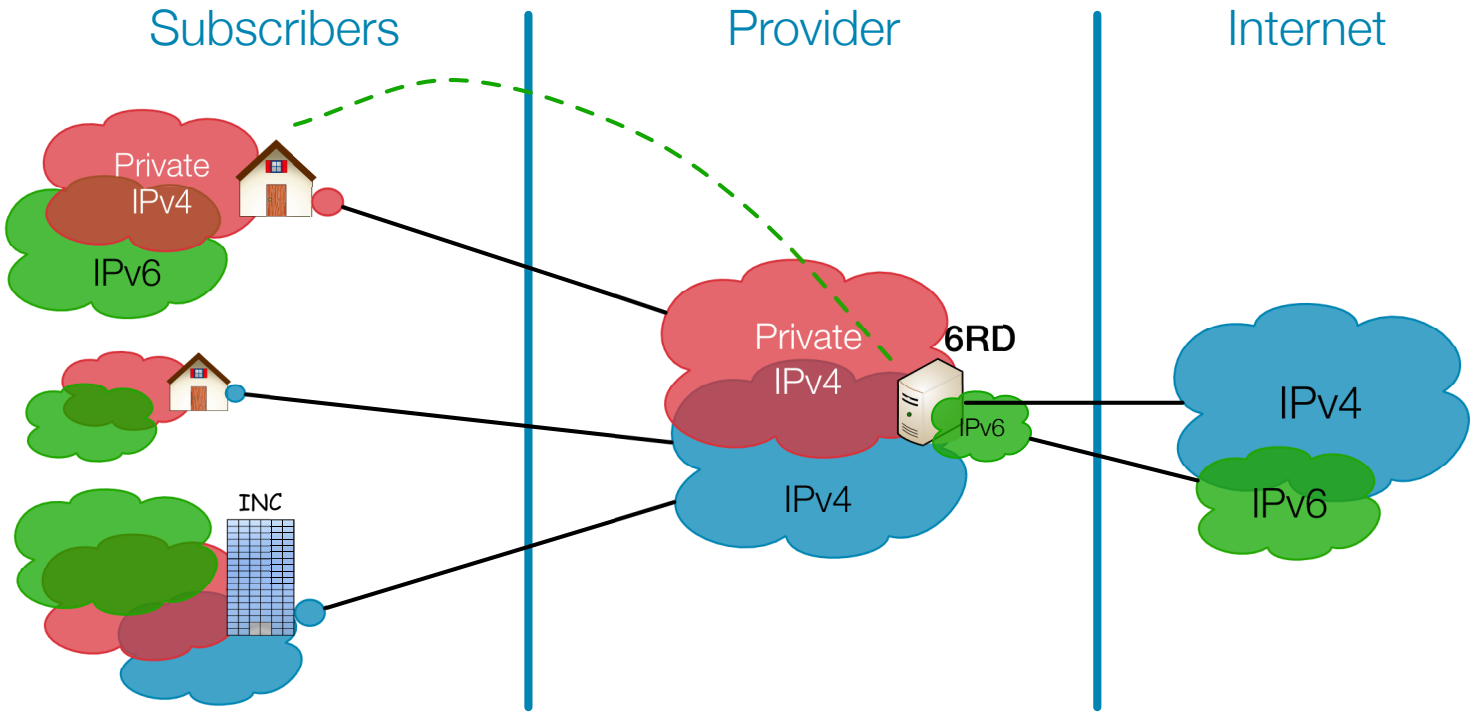
Subscribers
Provider
Internet
Private
IPv4
IPv6
Private
IPv4
6RD
IPv6
IPv4
IPv4
IPv6
INC

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

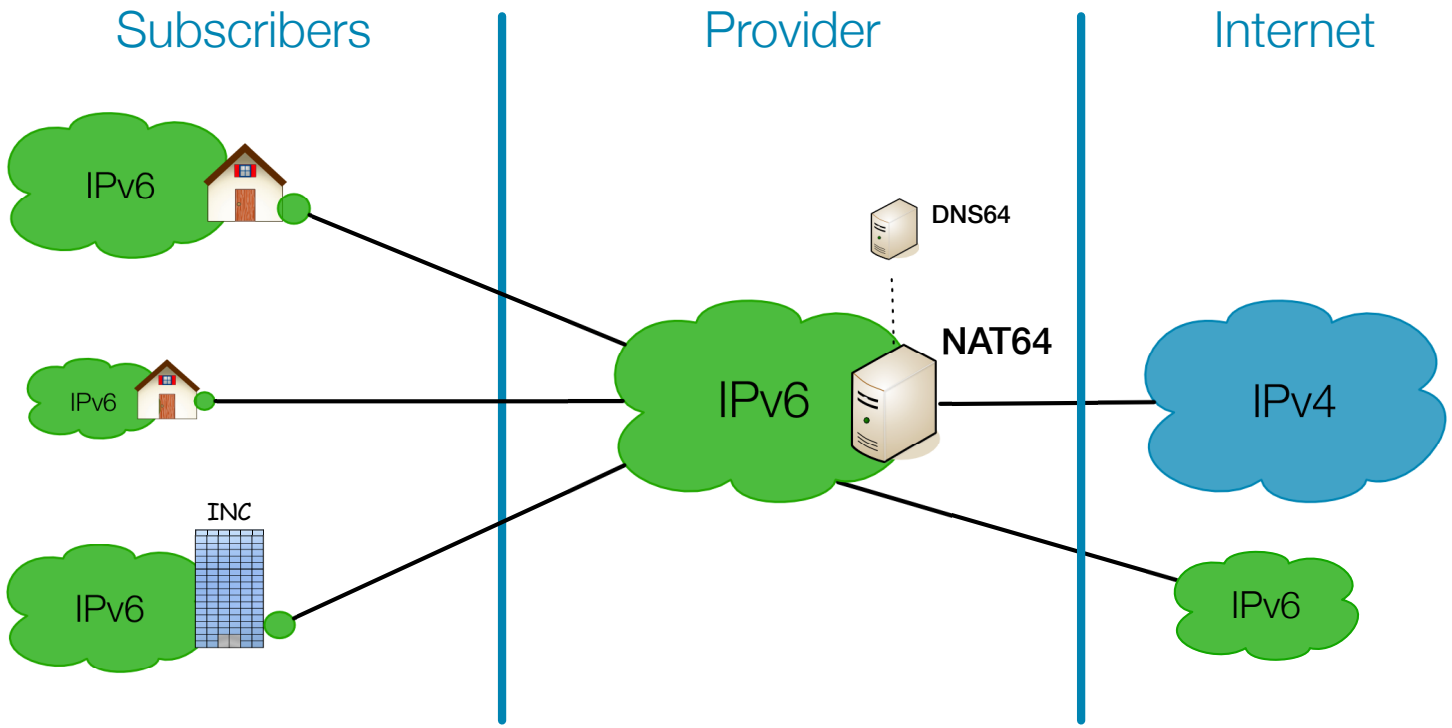
Subscribers
Provider
Internet
IPv6
IPv6
INC
IPv6
DNS64
NAT64
IPv6
IPv4
IPv6

DS-lite

- Tunneling IPv4 over IPv6
- Allows clients to use RFC1918 addresses without doing NAT themselves
- NAT is centrally located at the provider
- Client's IPv6 address is used to maintain state and to keep clients apart
  - Allows for duplicate IPv4 ranges

DS-lite

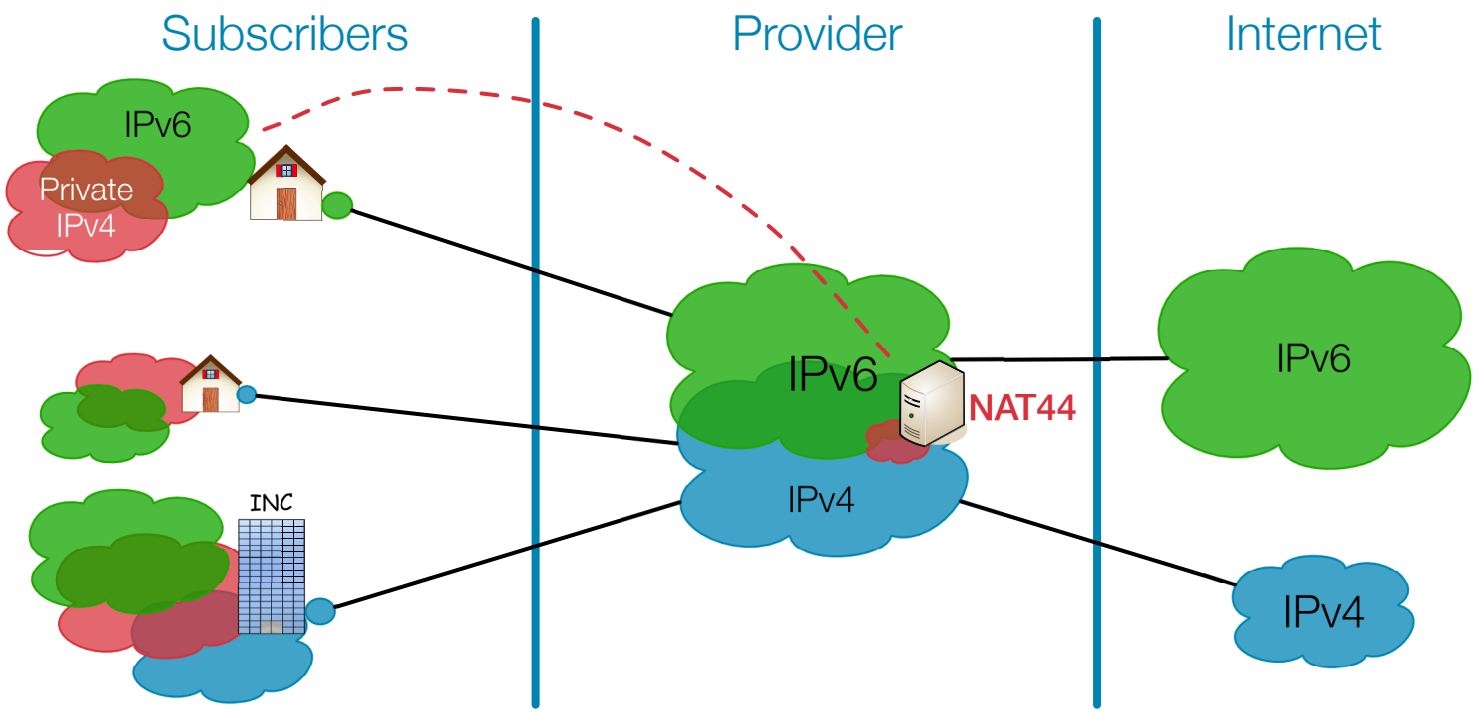
Subscribers
Provider
Internet
IPv6
Private
IPv4
IPv6
NAT44
IPv4
IPv6
IPv4
INC

Questions?

RIPE NCC

Tips

Section 7

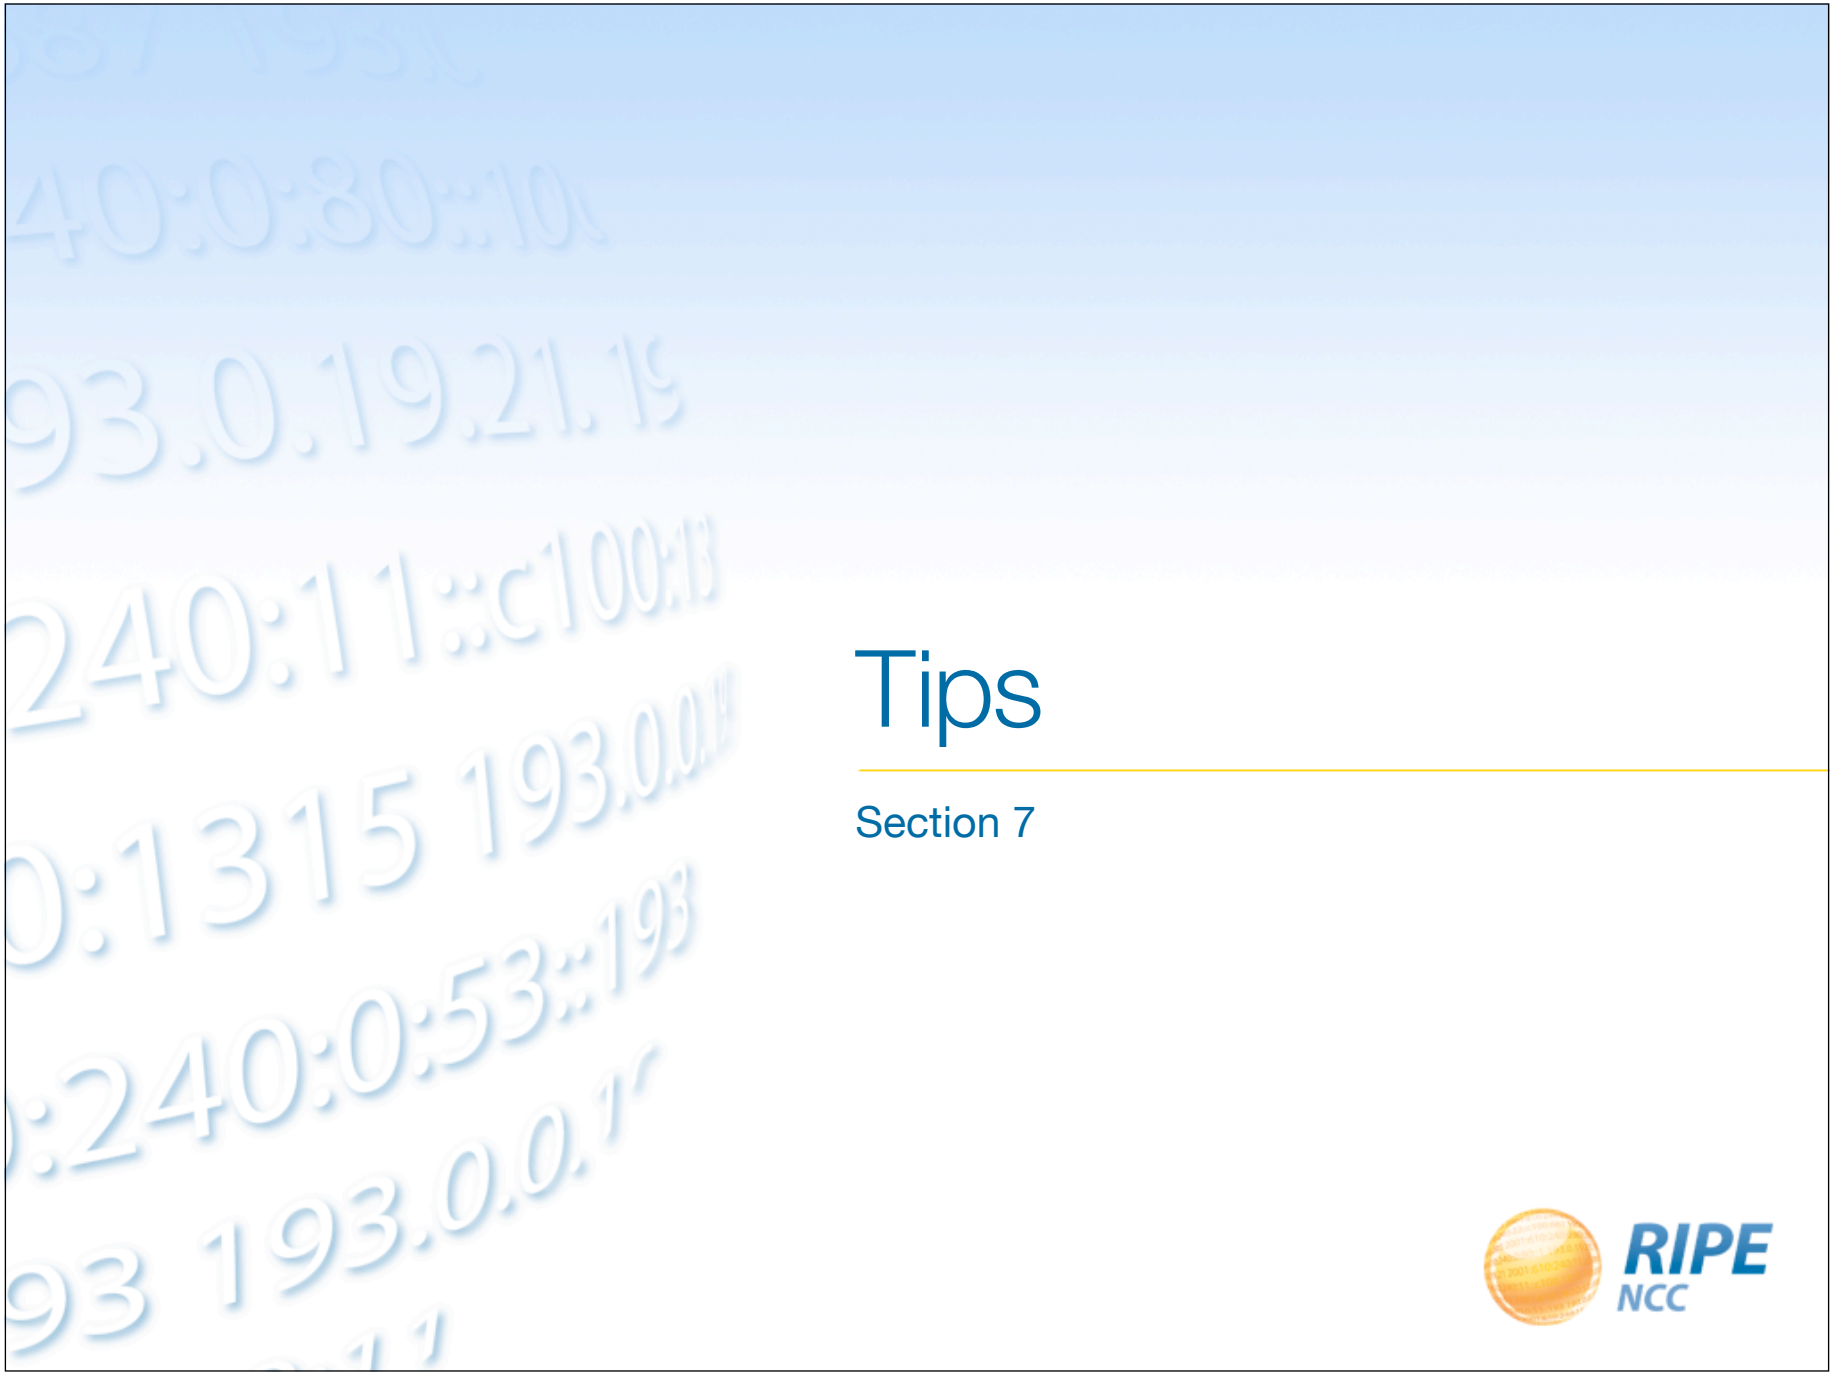

Best Scenario: Act Now, Phased Approach

- Change purchasing procedure (feature parity)
  - RIPE-501
- Check your current hardware and software
- Plan every step and test
- One service at a time
  - face first
  - core
  - customers
- Prepare to be able to switch off IPv4

Don'ts

- Don't separate IPv6 features from IPv4
- Don't do everything in one go
- Don't appoint an IPv6 specialist
  - do you have an IPv4 specialist?
- Don't see IPv6 as a product
  - the Internet is the product

Business Case

- IPv4 is no longer equal to “the Internet”
- Avoiding the issue does not make it go away
- How much are you willing to spend now to save money later?
- Only IPv6 allows continued IP networking growth
- What do you want the Internet to be like in 5 years?

“IPv6, act now!”

Customer Premises Equipment Survey

- CPE devices that support IPv6
- Based on feedback from users
- Use it as a guide
- labs.ripe.net: search for ‘IPv6 CPE’

Customers And Their /48

- Customers have no idea how to handle 65536 subnets!
- Provide them with information
  - https://www.ripe.net/lir-services/training/material/IPv6-for-LIRs-Training-Course/IPv6_addr_plan4.pdf

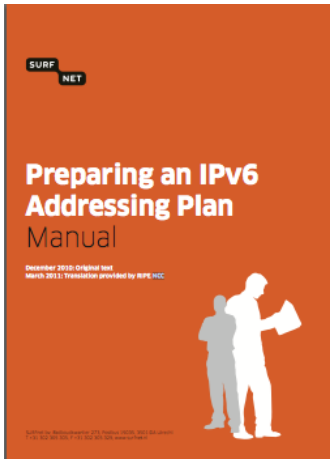


The End!
Край
Y Diwedd
Fí
Finis
النهاية
Соңы
վերջ
Liðugt
Ende
Finvezh
Кінець
Konec
Kraj
Ënn
Fund
پایان
Son
Kpaj
Lõpp
Beigas
Vége
An Críoch
הסוף
Endir
Fine
Sfârşit
Fin
Τέλος
Einde
Конец
Slut
Slutt
დასასრული
Pabaiga
Tmiem
Koniec
Amaia
Loppu
Fim