IPv6 for LIRs

March 2012

IANA IPv4 Pool

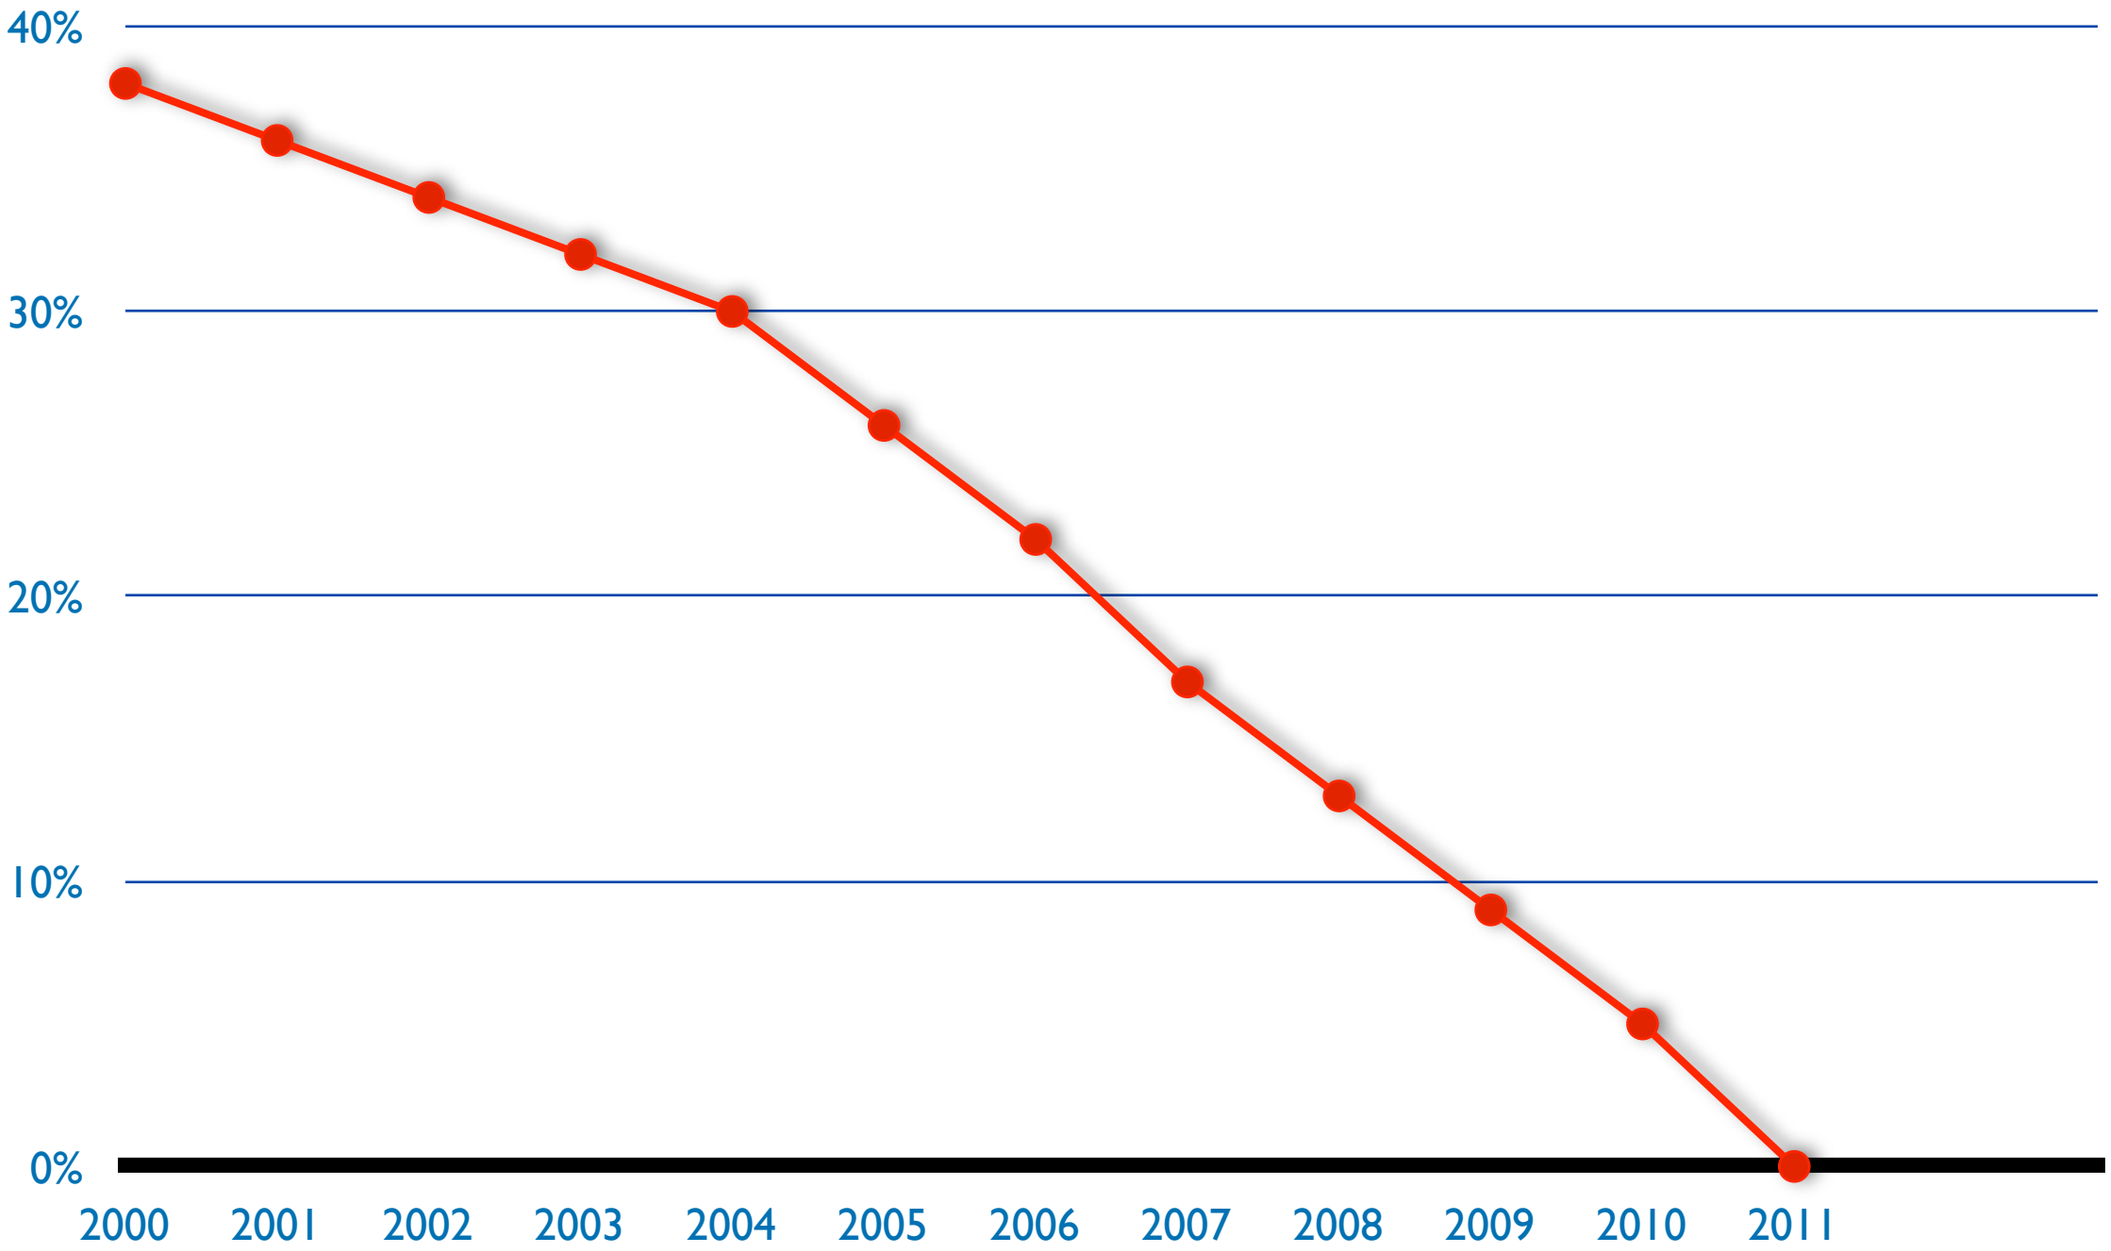

Reaching the next billion

- Around 2 billion Internet users now
  - around 30% of all people

- Mobile phones are becoming Internet devices

- The Internet of things

Schedule

- 09:00 - 09:30 Coffee, Tea
- 11:00 - 11:15 Break
- 13:00 - 14:00 Lunch
- 15:00 - 15:15 Break

- 17:30 End

Introductions

- Name
- Number in the list
- Experience with the RIPE NCC
- Goals

Overview

- The Registry System
- IPv4?
- The road to IPv6
- IPv6 Address Basics
- Transition Mechanisms
- Exercise: Addressing Plan
- Getting it
- Exercise: Making Assignments
- Deploying
- Exercise: Deployment Challenges
- Real Life IPv6 Deployment
- Tips

The Registry System

IP Address Distribution

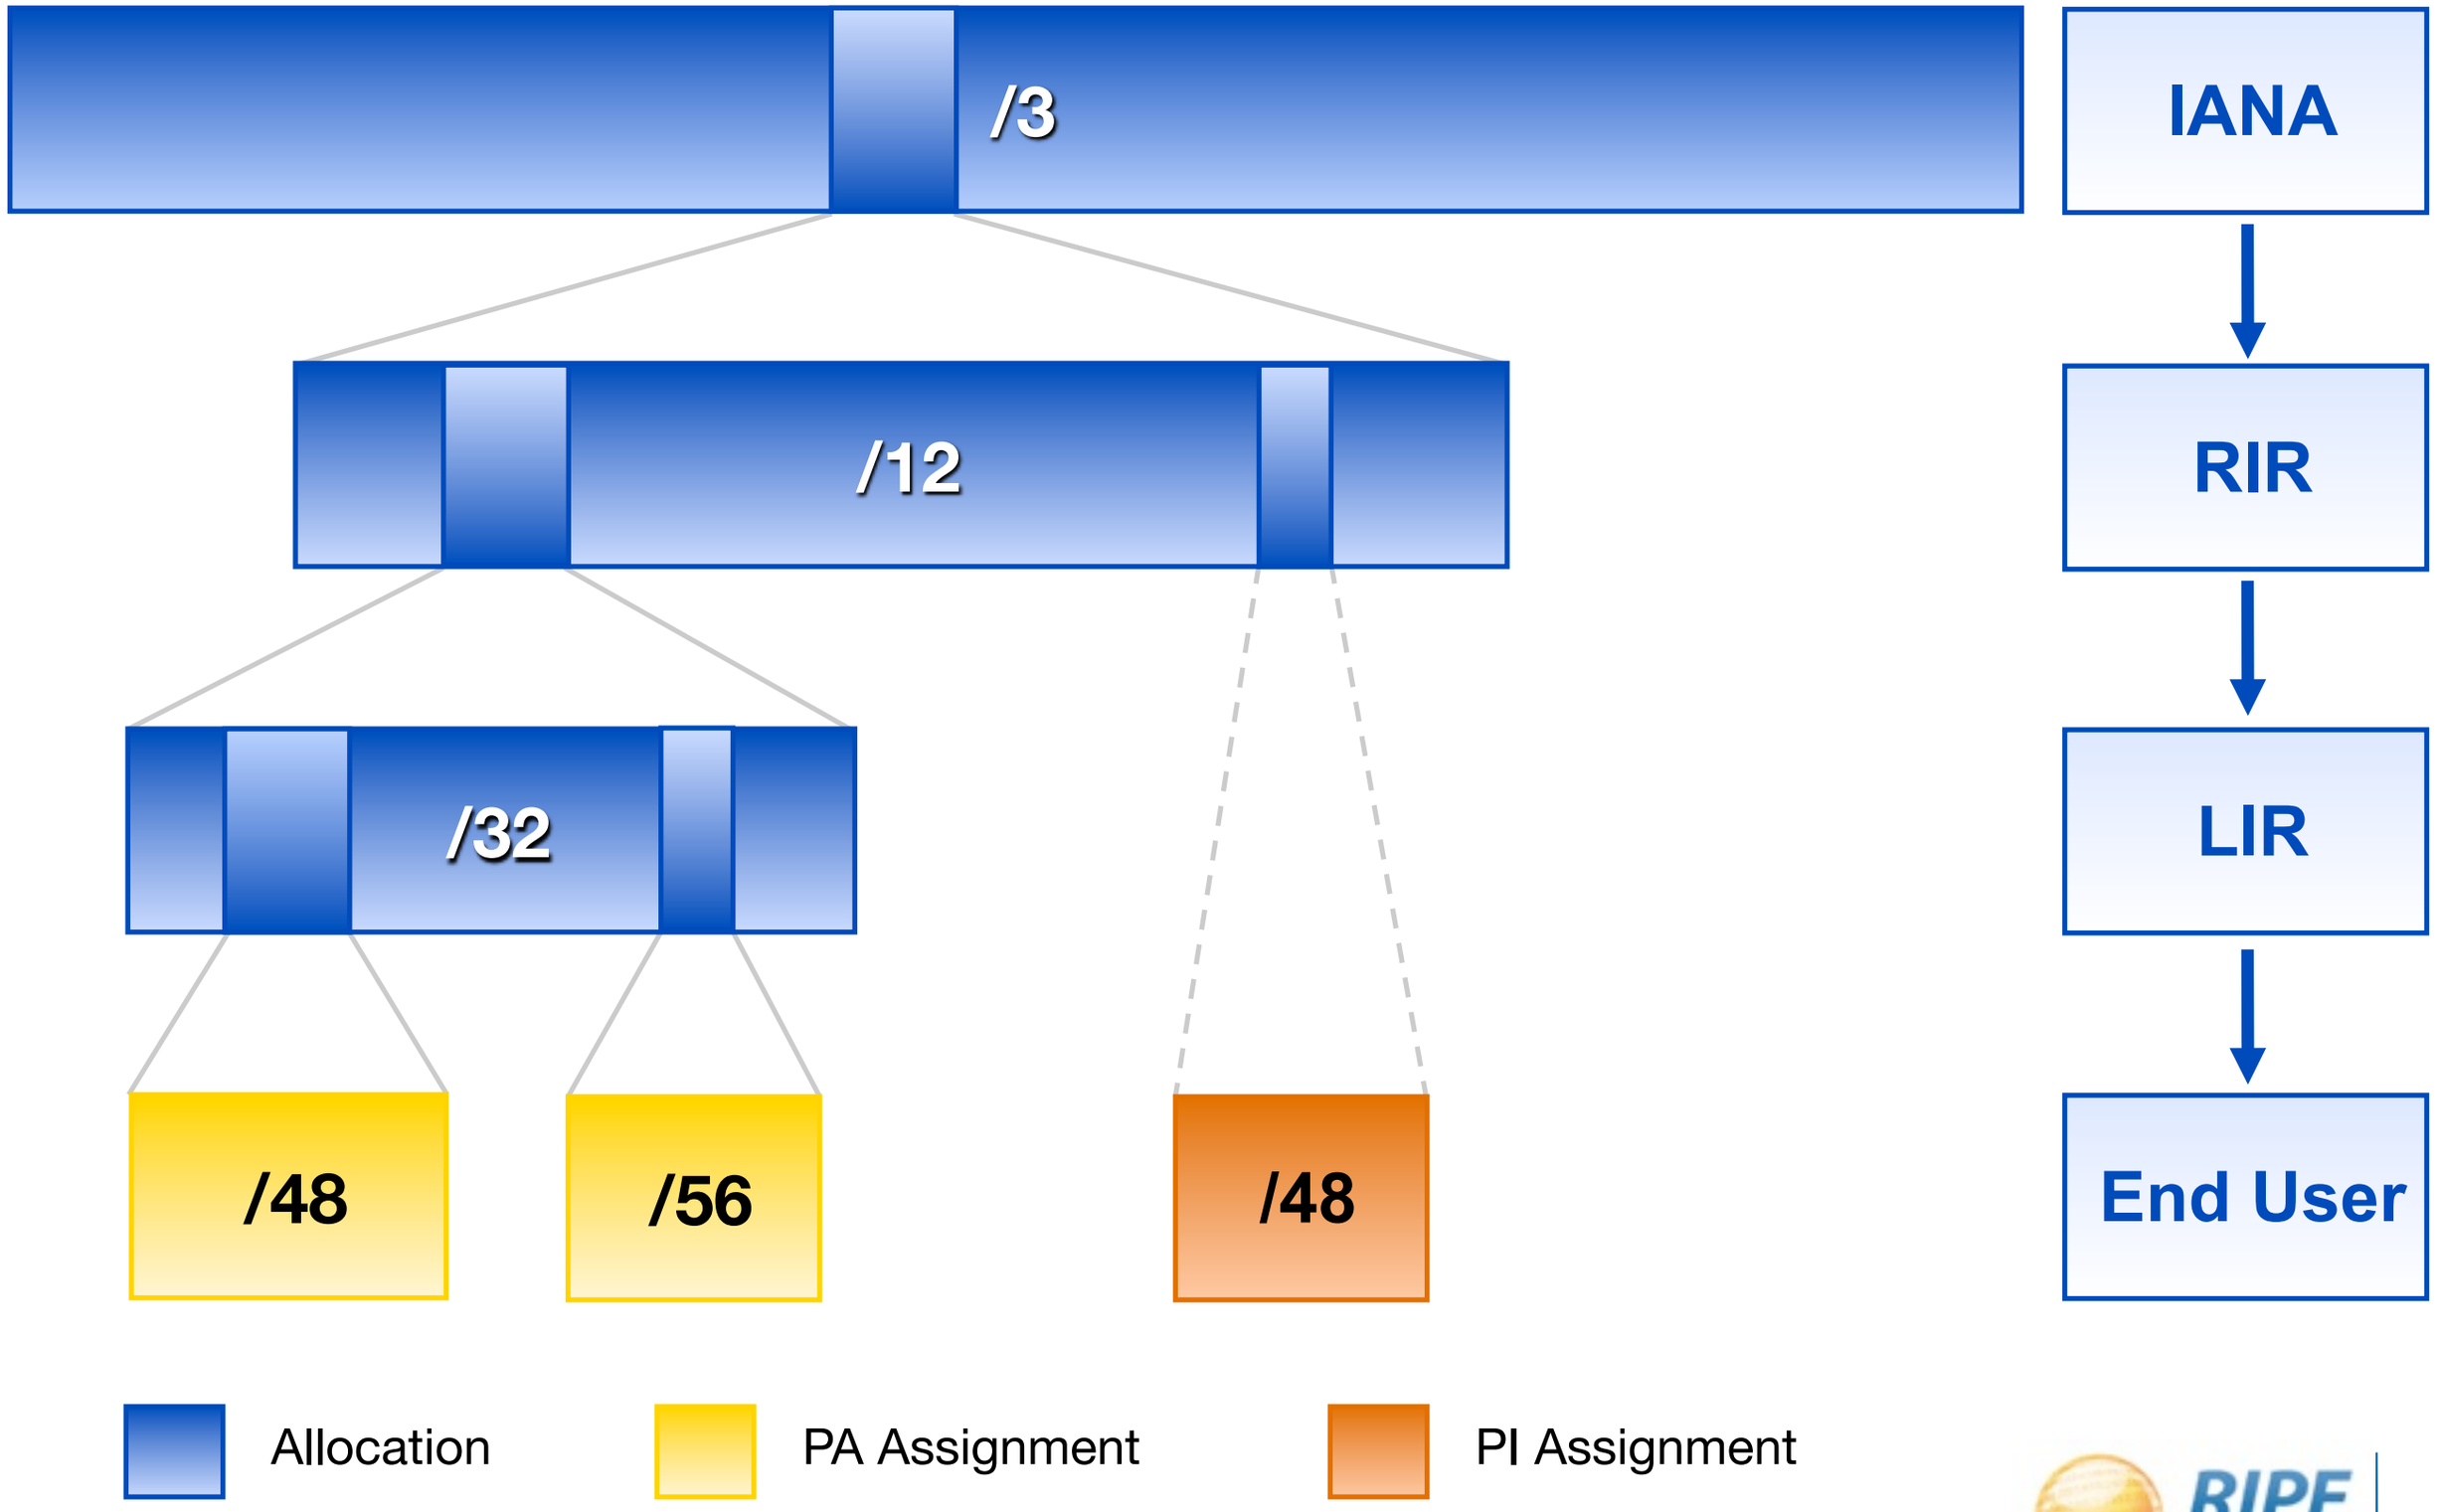

Registration

Conservation

Aggregation

IPv4?

Our Slice of the IPv4 Pie

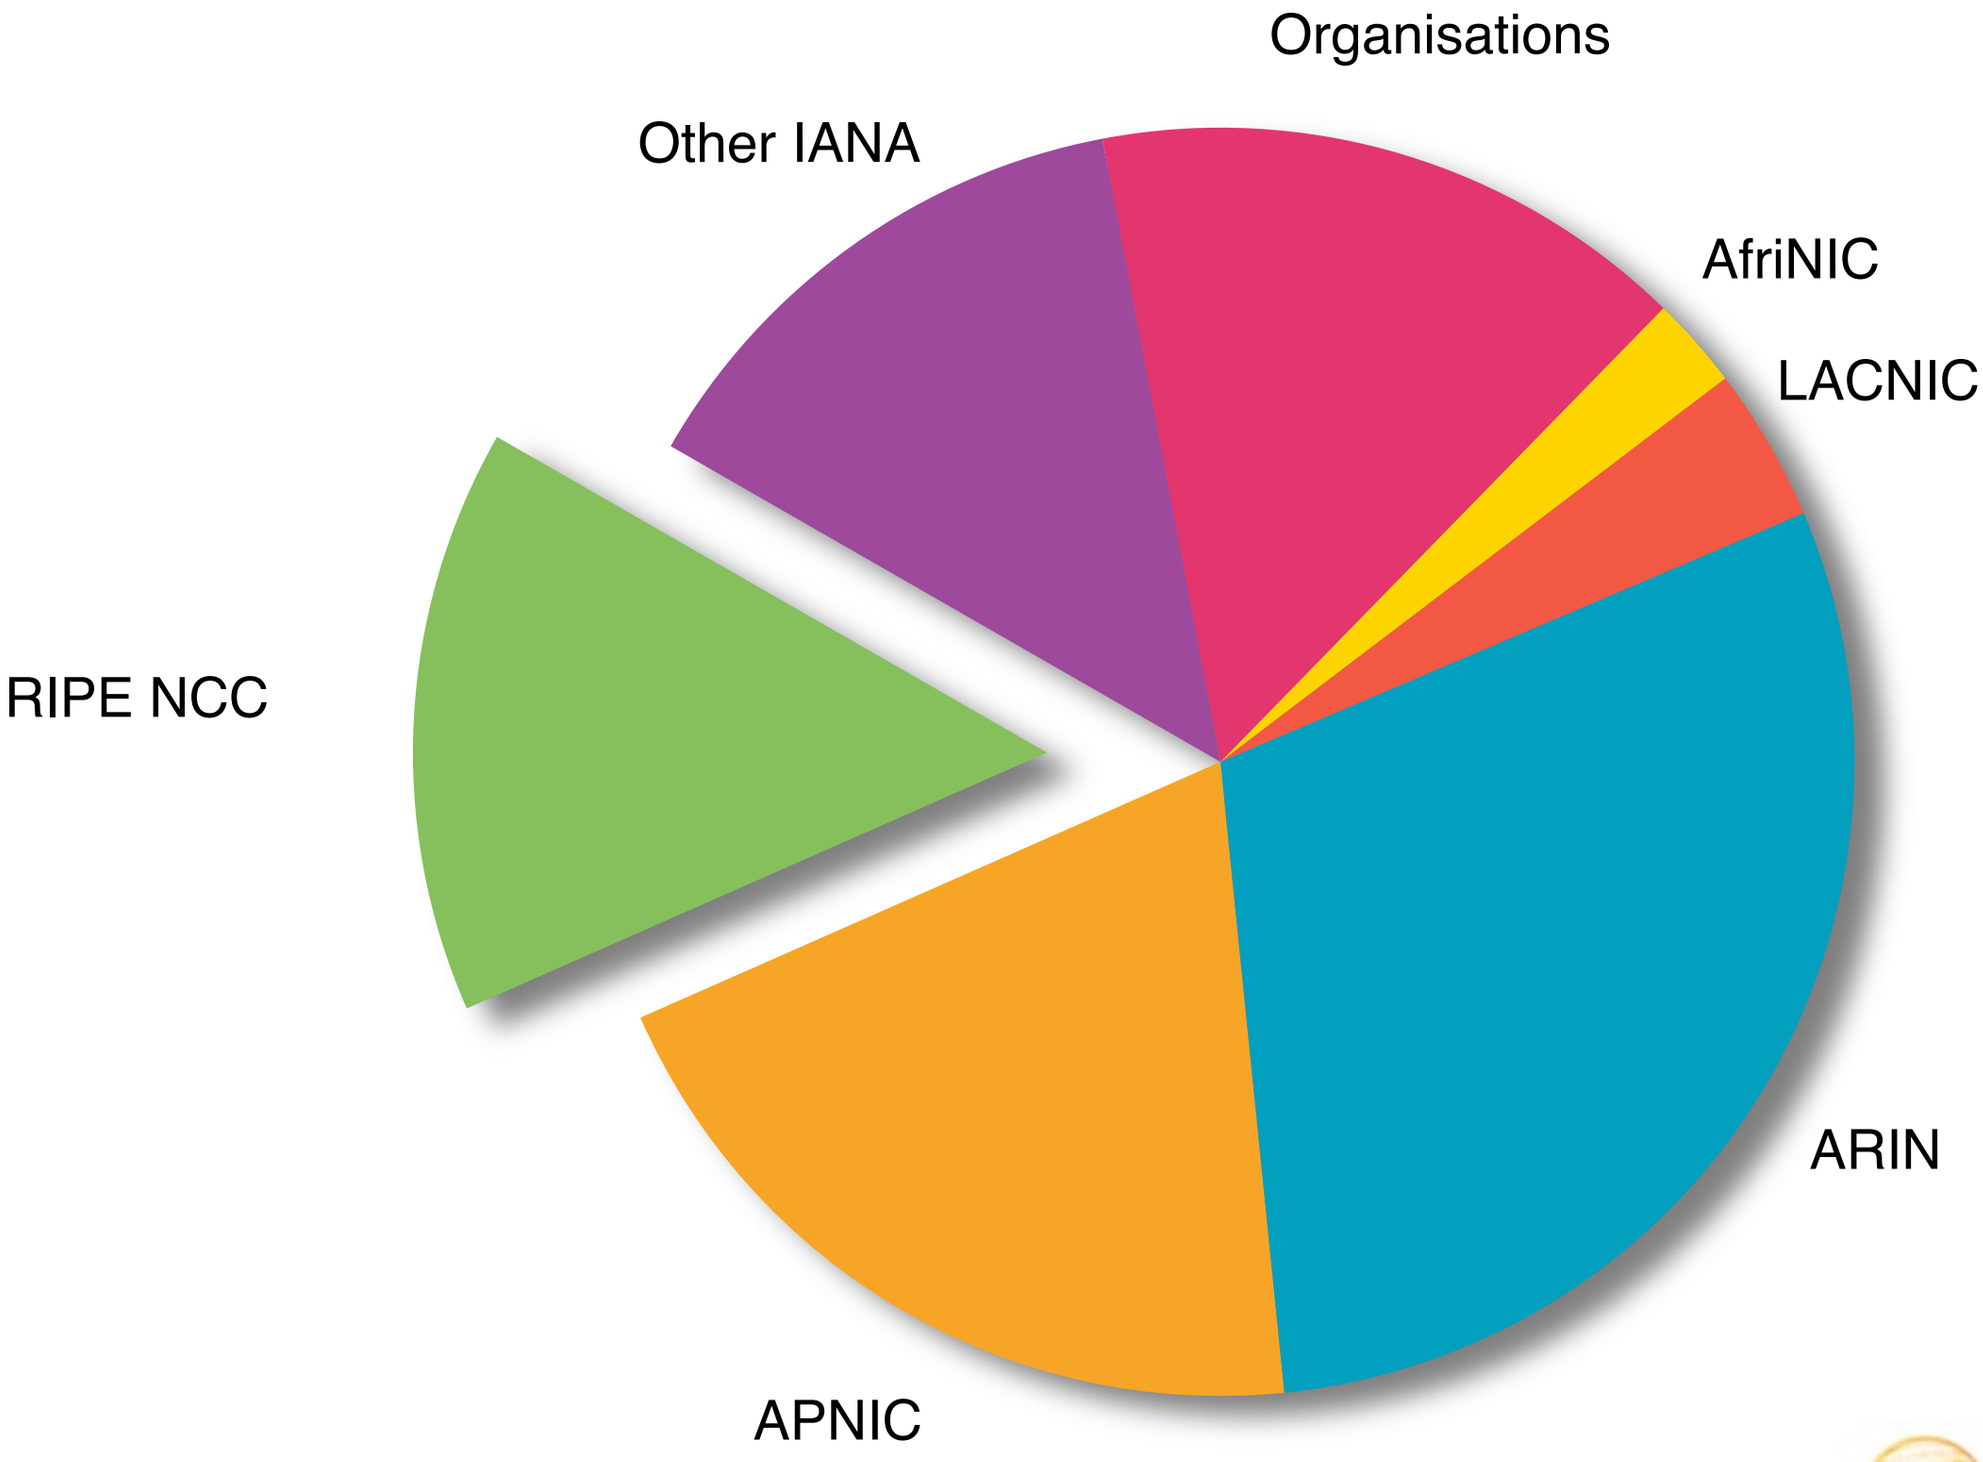

Wait and See?

Reduced Assignment Periods

- Used to be: 24 months
- January 2010: 12 months
- July 2010: 9 months
- January 2011: 6 months
- **July 2011: 3 months**

Hot IPv4 / IPv6 Policy Topics

- Allocations from the last /8 (2010-02)
  - new and existing LIRs can receive only one /22 allocation
  - only if they already have an IPv6 allocation!

- RIPE Document: RIPE-530

RIPE Policy Proposal 2011-04

- Extension of the Minimum Size for IPv6 Initial Allocation
  - Proposes initial allocation up to a /29
  - For example, for small LIRs to deploy IPv6 via 6RD (RFC 5969)

- Proposal currently in Review Phase
  - The RIPE NCC is working on impact analysis

Transfer of IPv4 Allocations

- Policy 2007-08: Allocation Transfer Policy
  - Don’t buy your IPv4 on eBay!
    - Transfer unused allocations to another LIR
    - Minimum allocation size /21
    - Evaluated by RIPE NCC
    - Update in RIPE Database

https://www.ripe.net/lir-services/resource-management/listing/

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

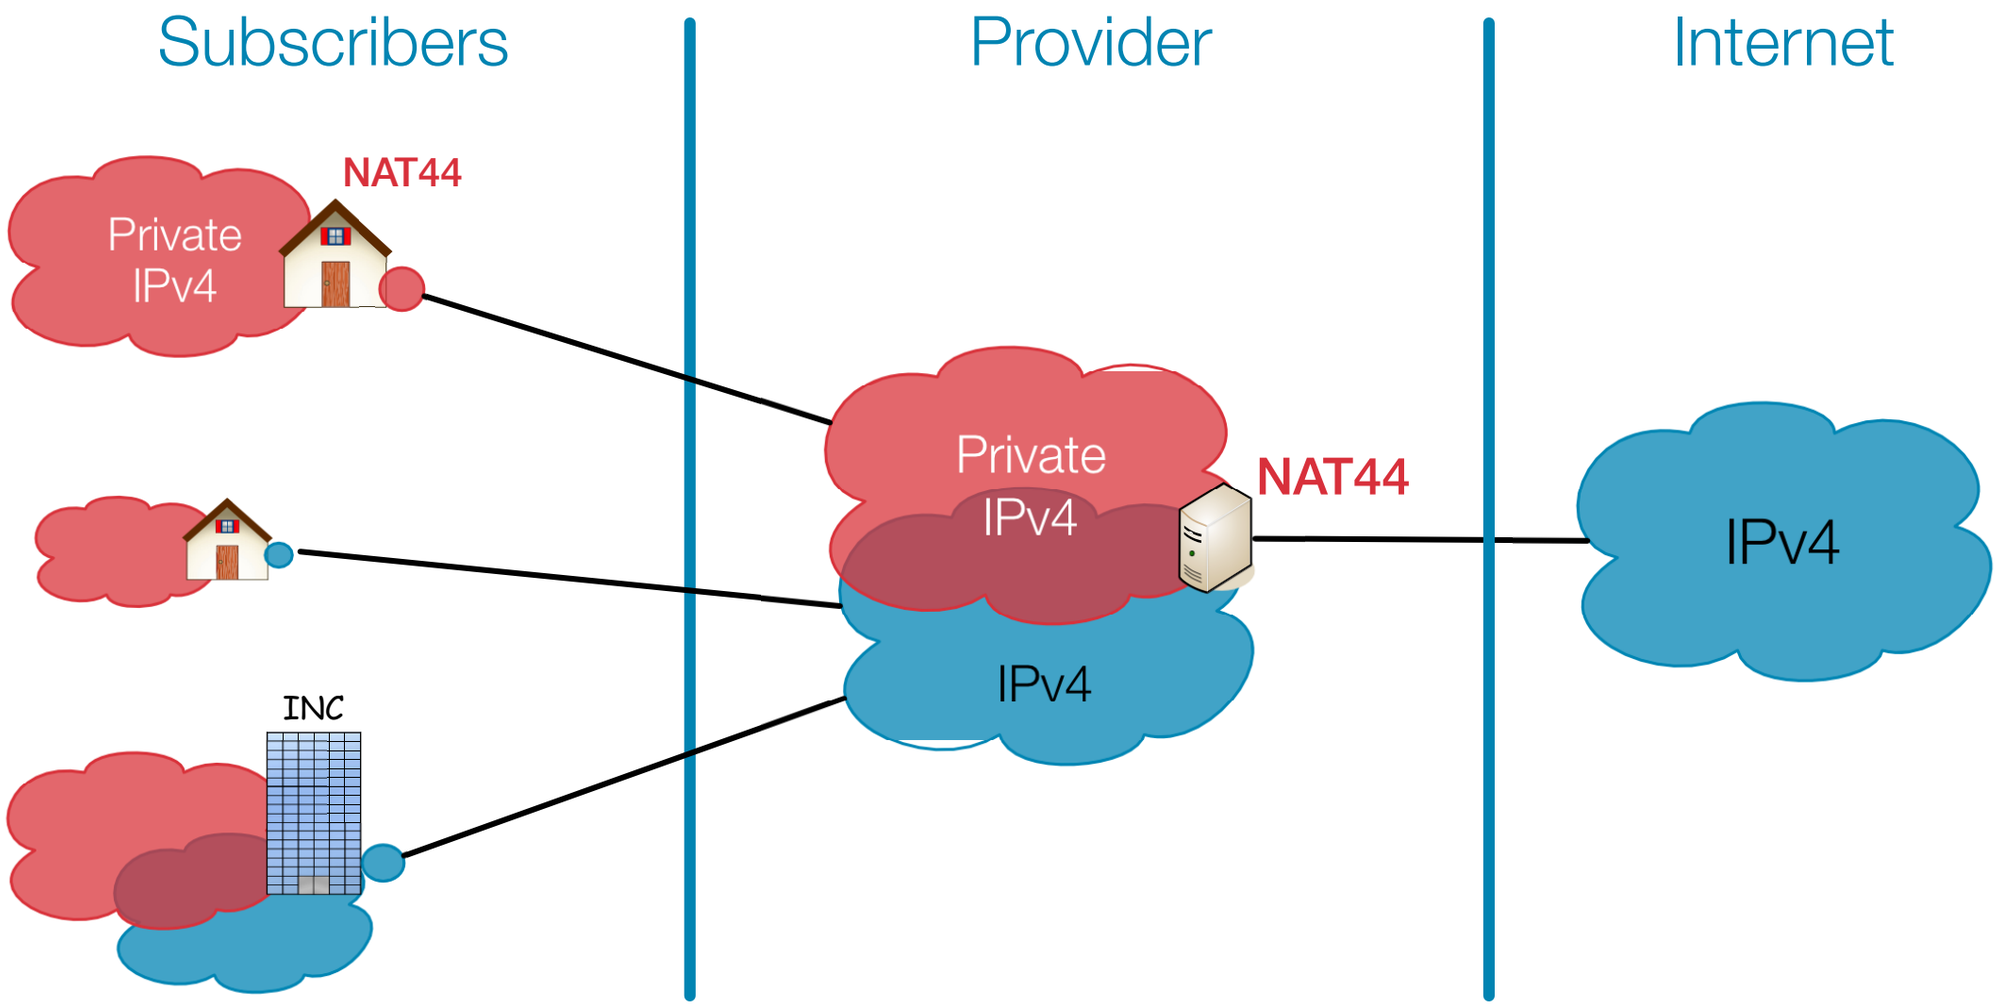

The road to IPv6

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 Techniques

- 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

IPv6 Address Basics

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

IPv6 Subnetting

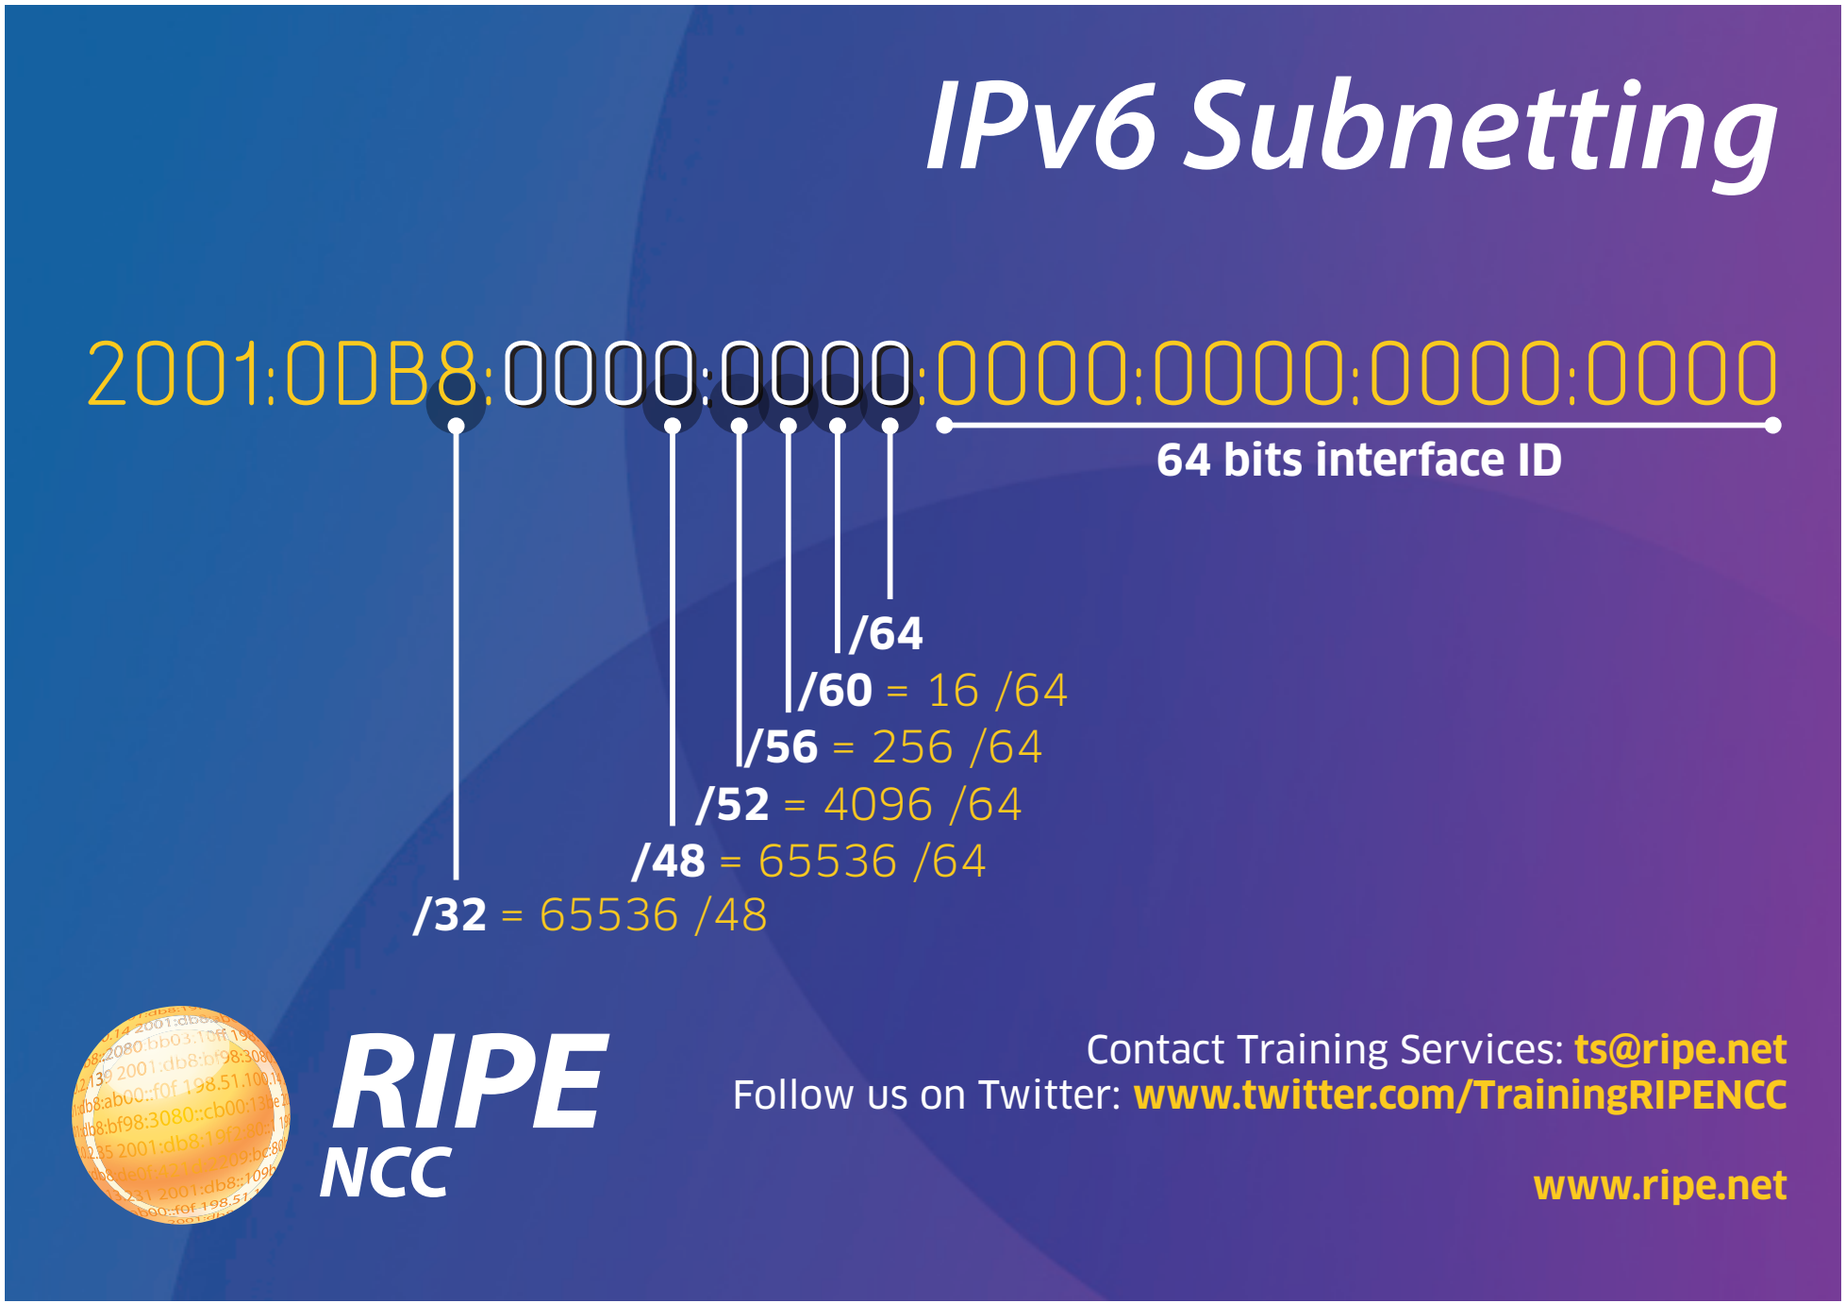

Address Notation

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

Address Notation

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

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

Address Notation

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

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

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

Address Notation

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

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

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

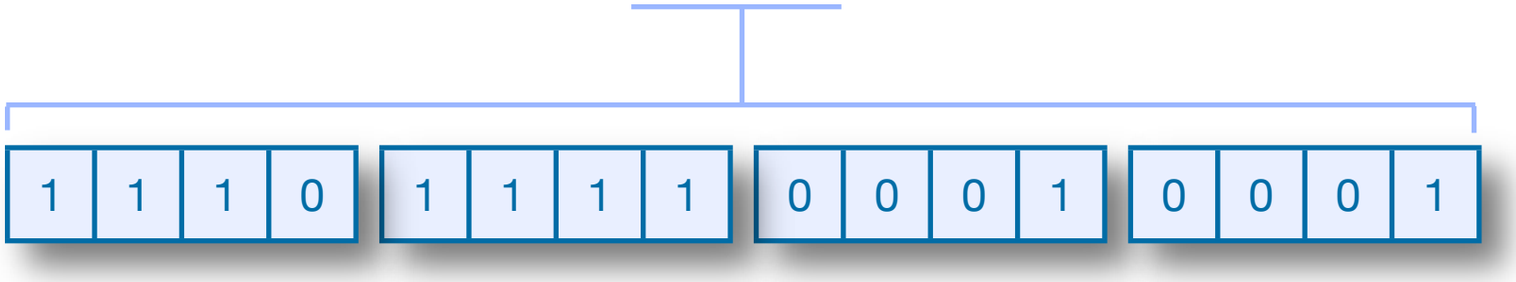

Multiple addresses

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

Transition Mechanisms

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

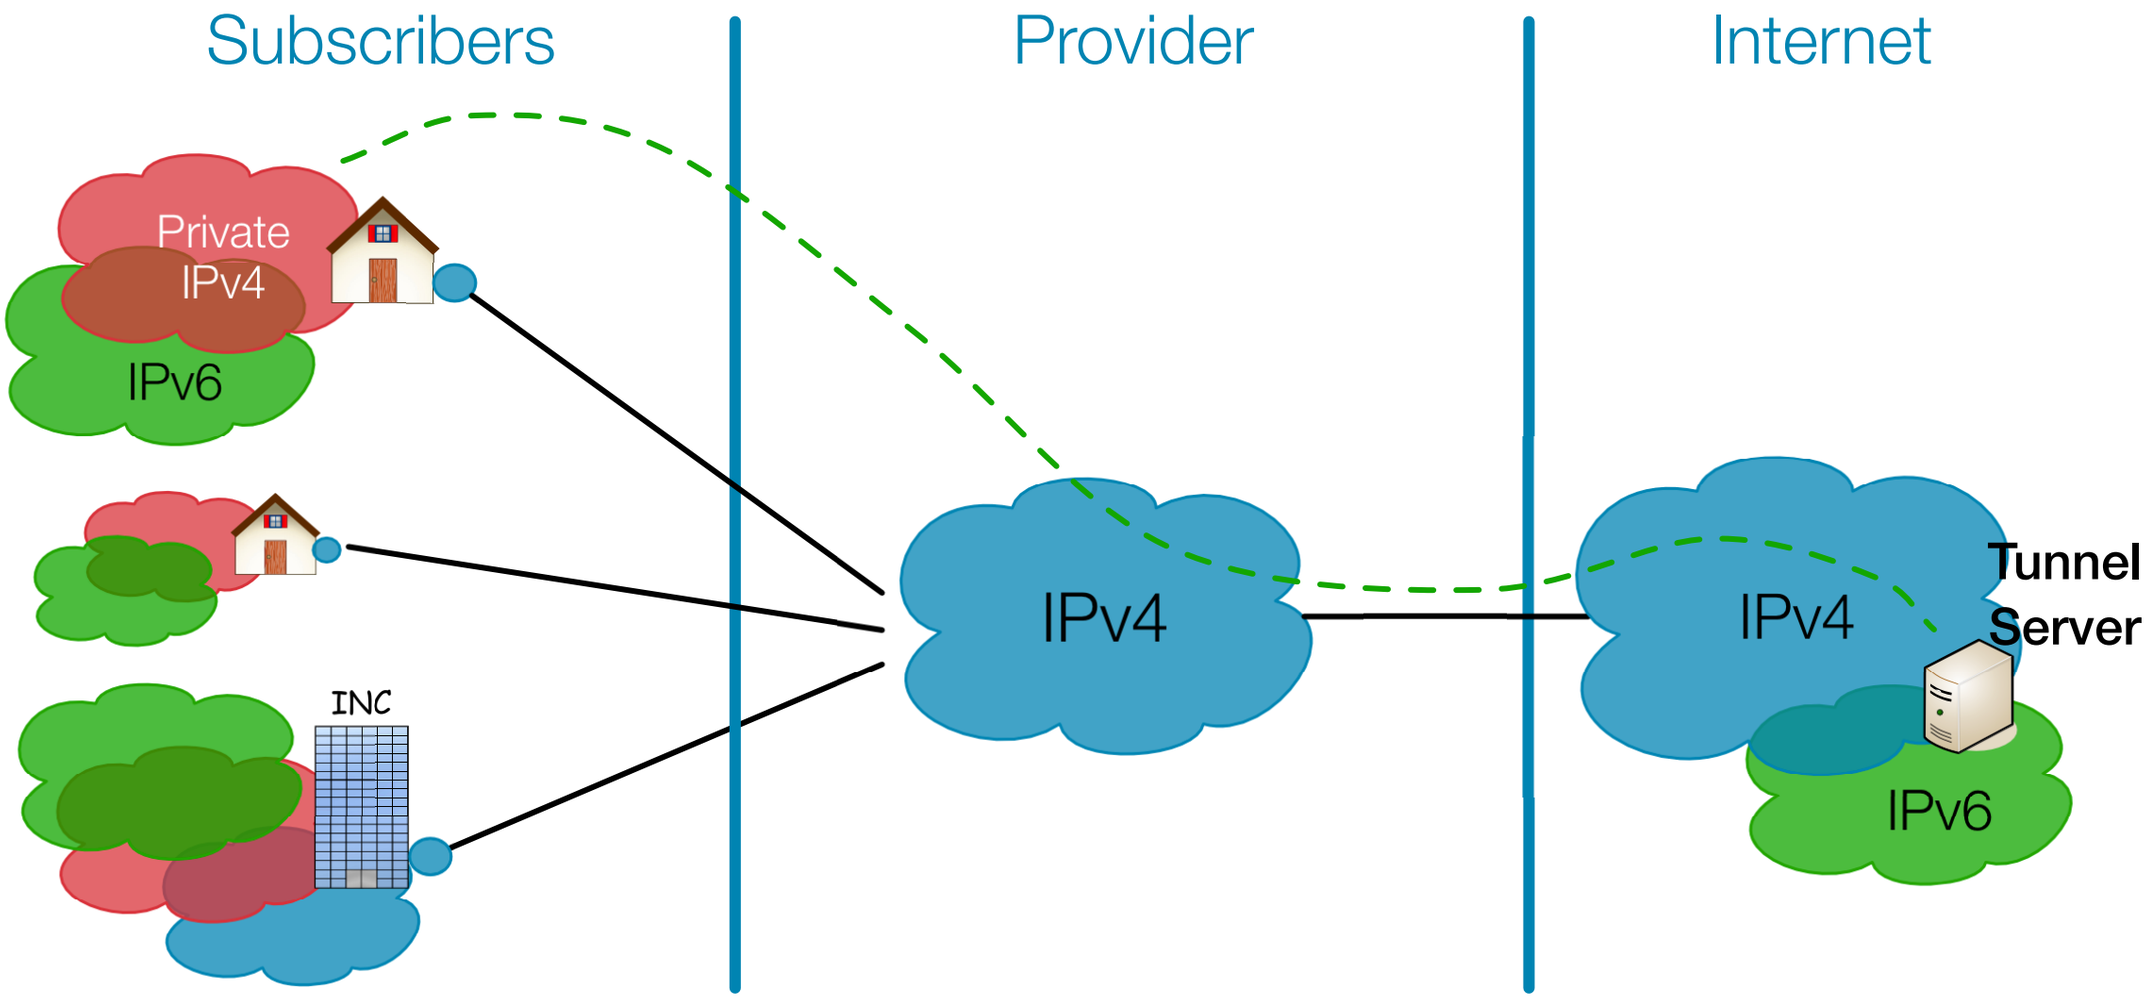
Subscribers
Provider
Internet
Private
IPv4
IPv6
INC
IPv4
IPv4
Tunnel
Server
IPv6

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

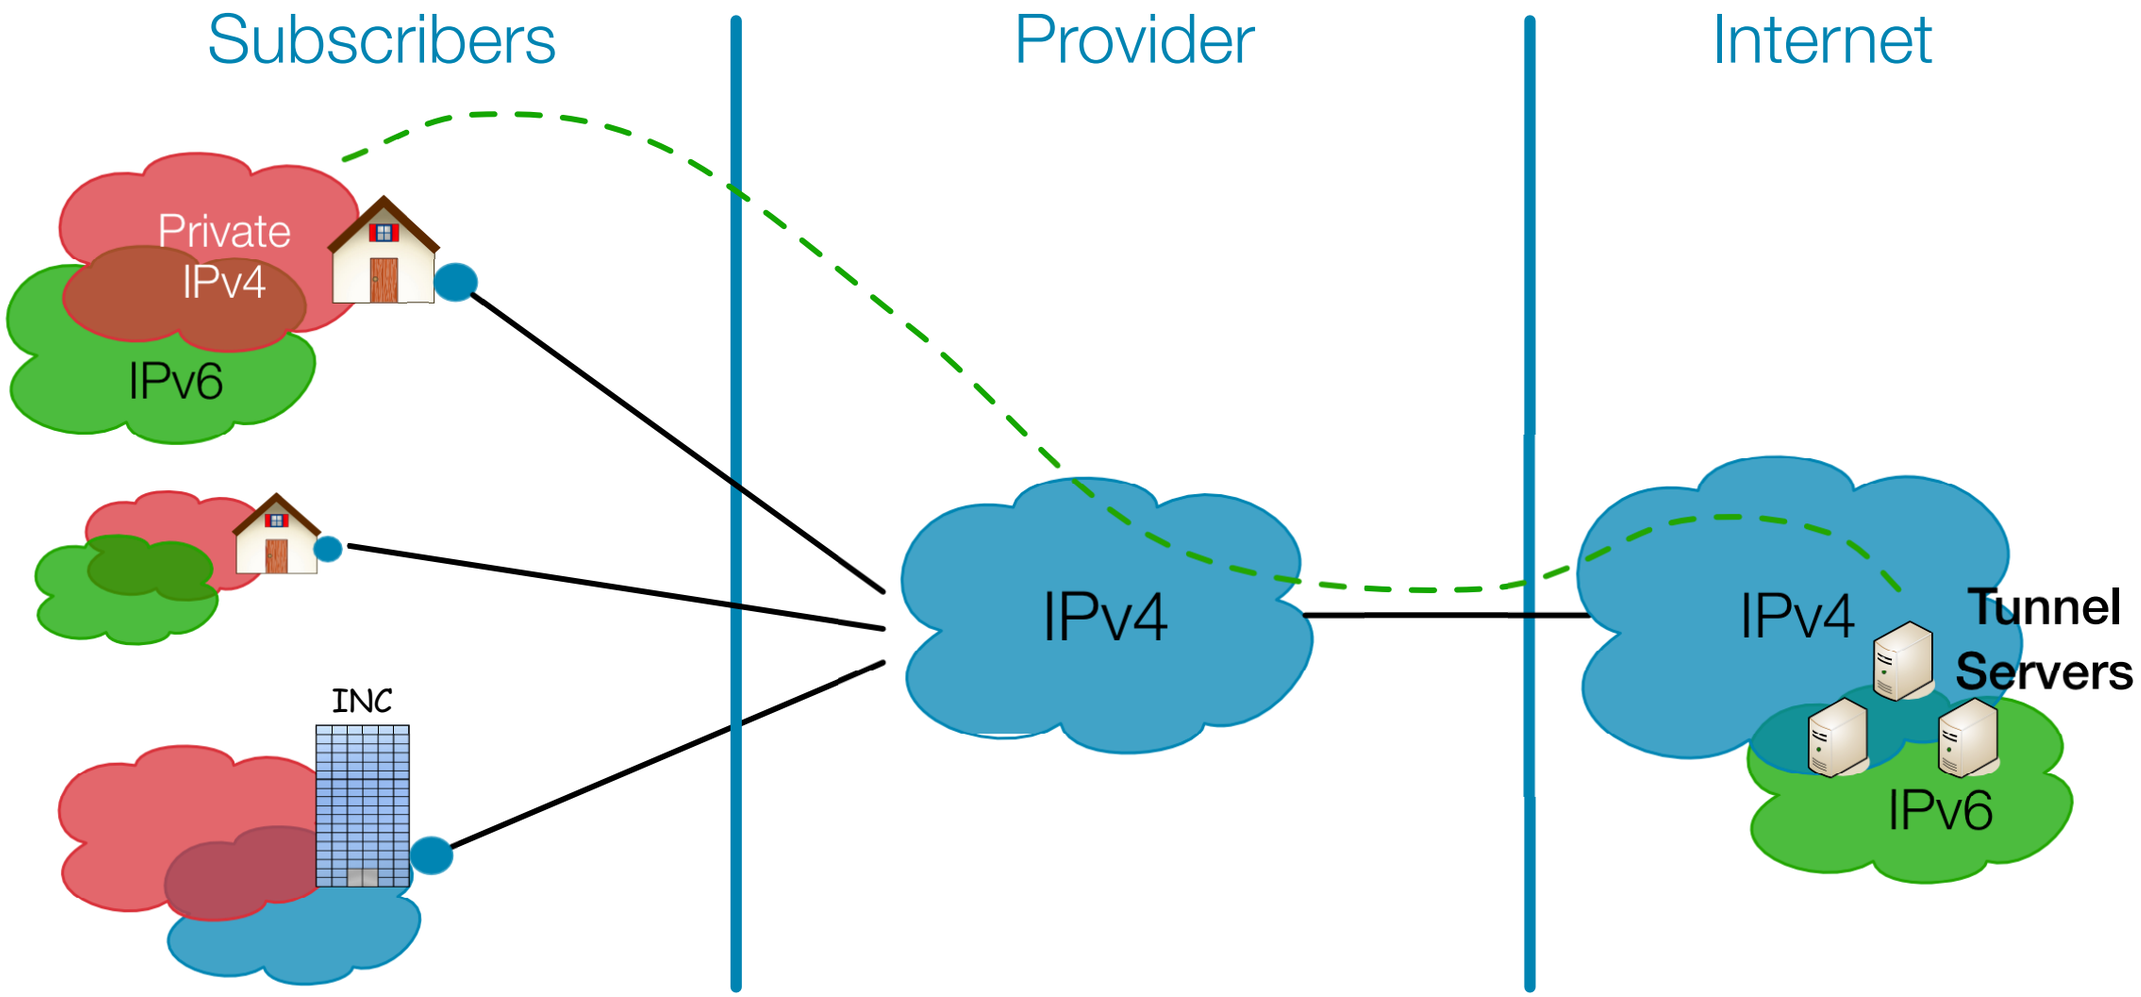

6RD

- 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

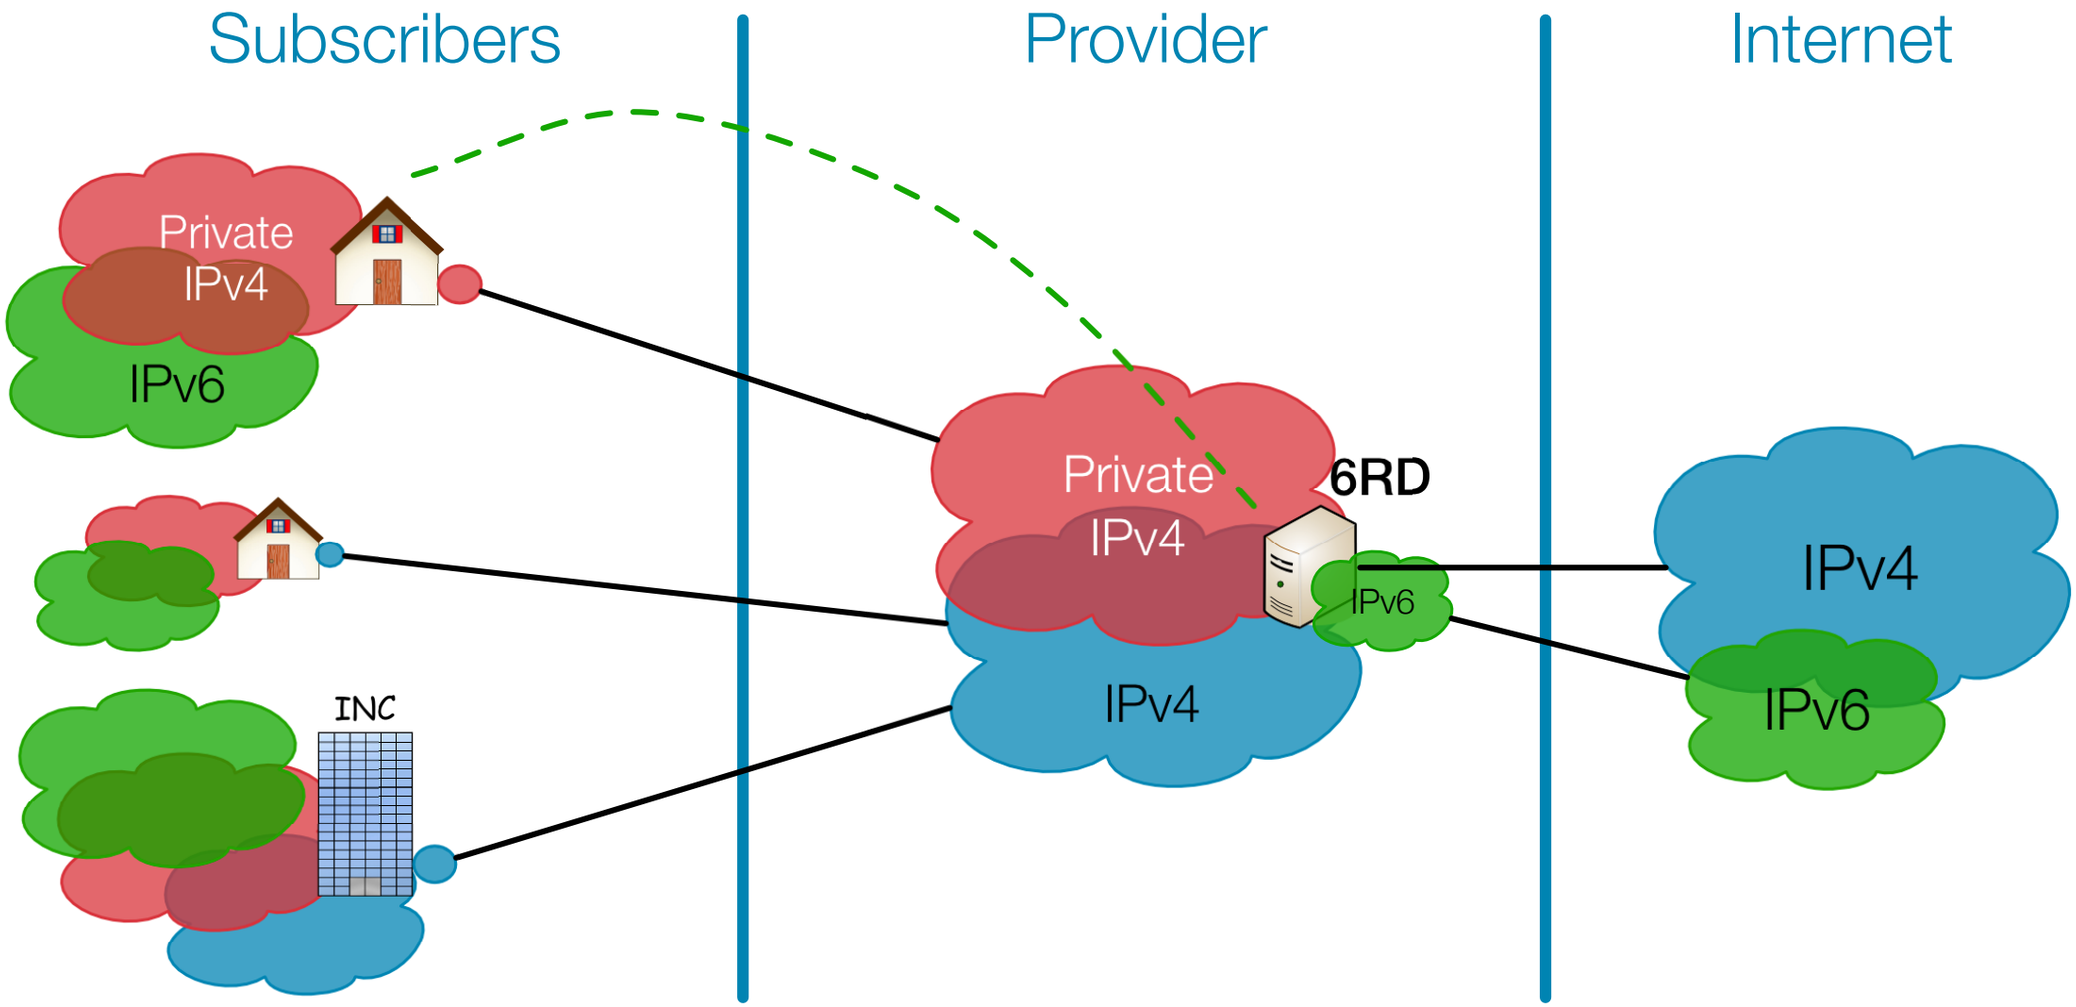
Subscribers
Provider
Internet
Private
IPv4
IPv6
INC
Private
IPv4
6RD
IPv6
IPv4
IPv4
IPv6

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

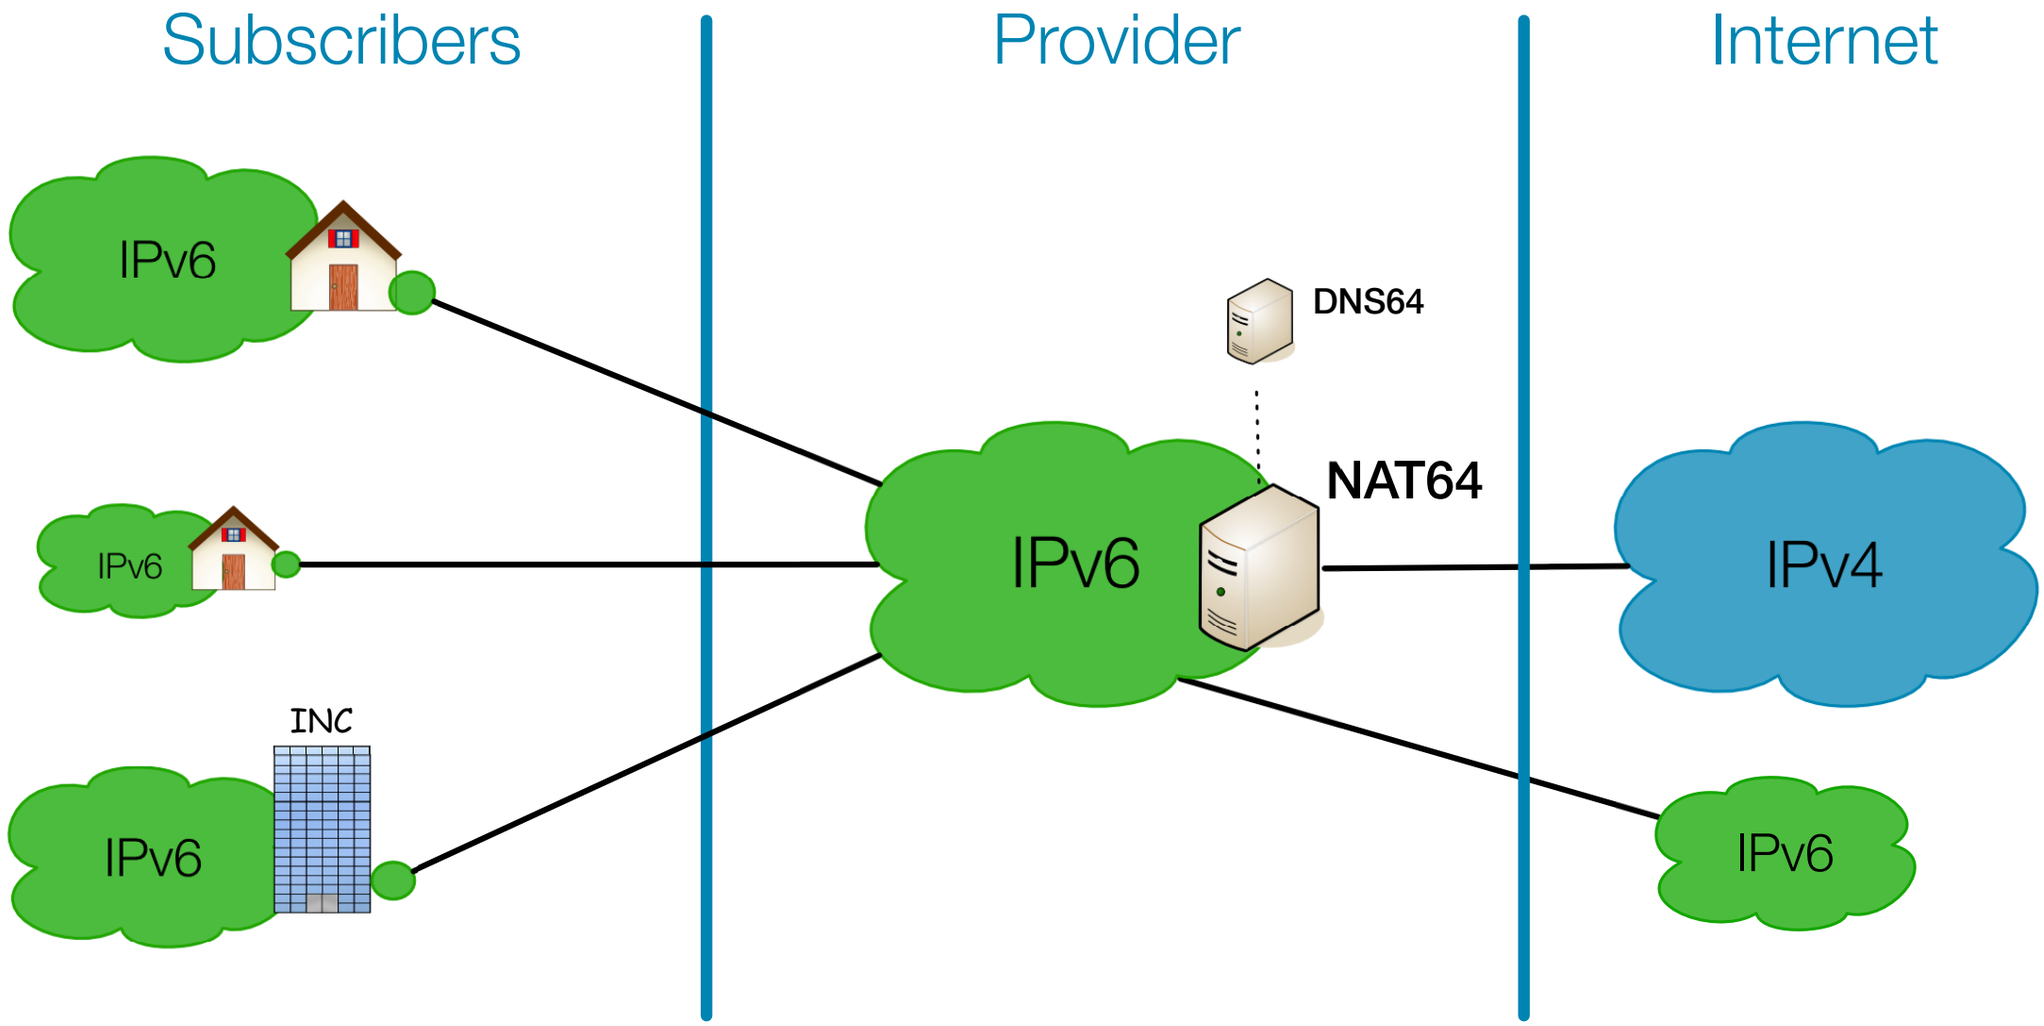
Subscribers
Provider
Internet
IPv6
IPv6
INC
IPv6
DNS64
NAT64
IPv6
IPv4
IPv6

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

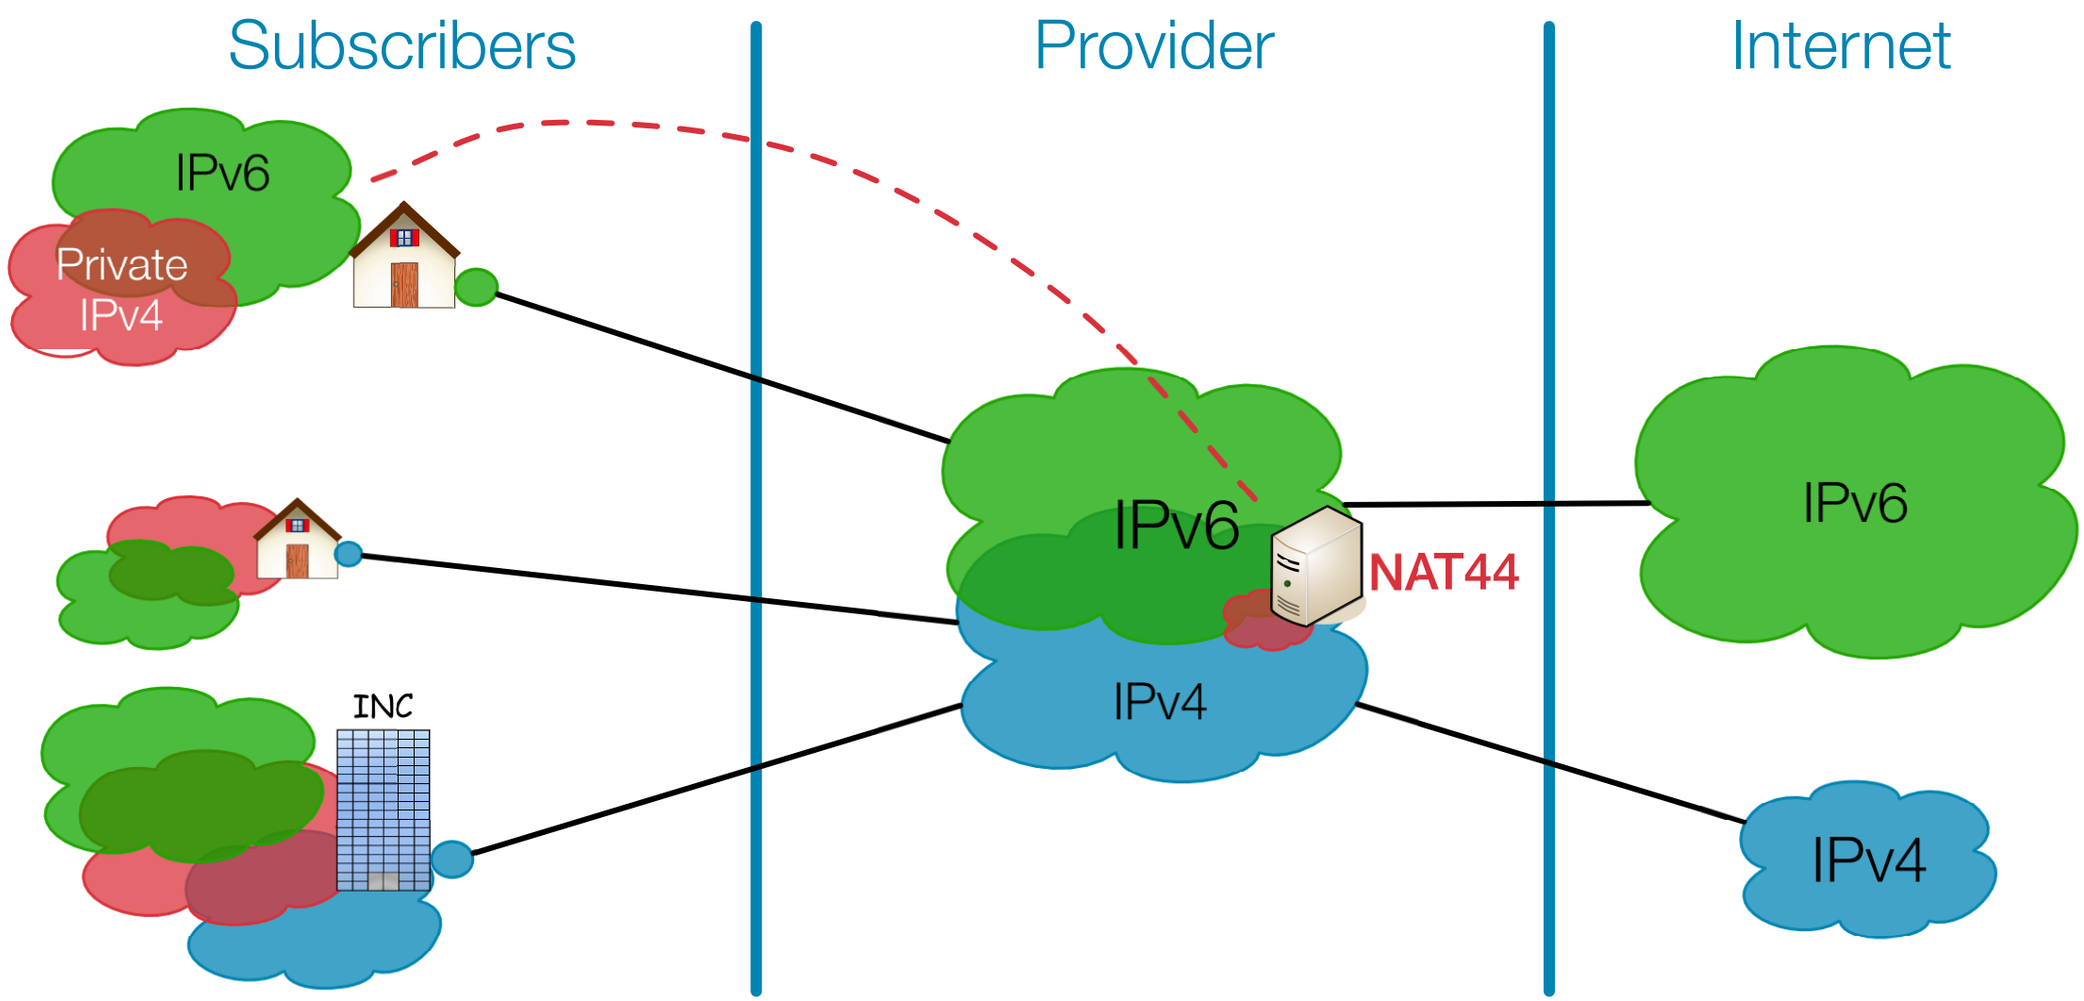
Subscribers
Provider
Internet
IPv6
Private
IPv4
IPv6
NAT44
IPv4
IPv6
INC
IPv4

Exercise

Addressing Plan

Addressing Plan Exercise

- Things to consider
  - administrative ease!
  - use assignments on 4 bit boundary
  - 2 possible scenarios for network
  - 5 possible scenarios for customer assignments

- 40 minutes preparation time
- 15 minutes discussion

Addressing Plans

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

More On Addressing Plans

- For private networks, get ULA
- For servers you might want manual config
- Use port numbers for addresses
  - pop server 2001:db8:1::110
  - dns server 2001:db8:1::53
  - etc...

Getting it

Getting an IPv6 allocation

- To qualify, an organisation must:
  - Be an LIR
  - Have a plan for making assignments within two years

- Minimum allocation size /32

- Allocation size is based on customer numbers and growth, not on transition technique!

IPv6 Allocations and Announcements

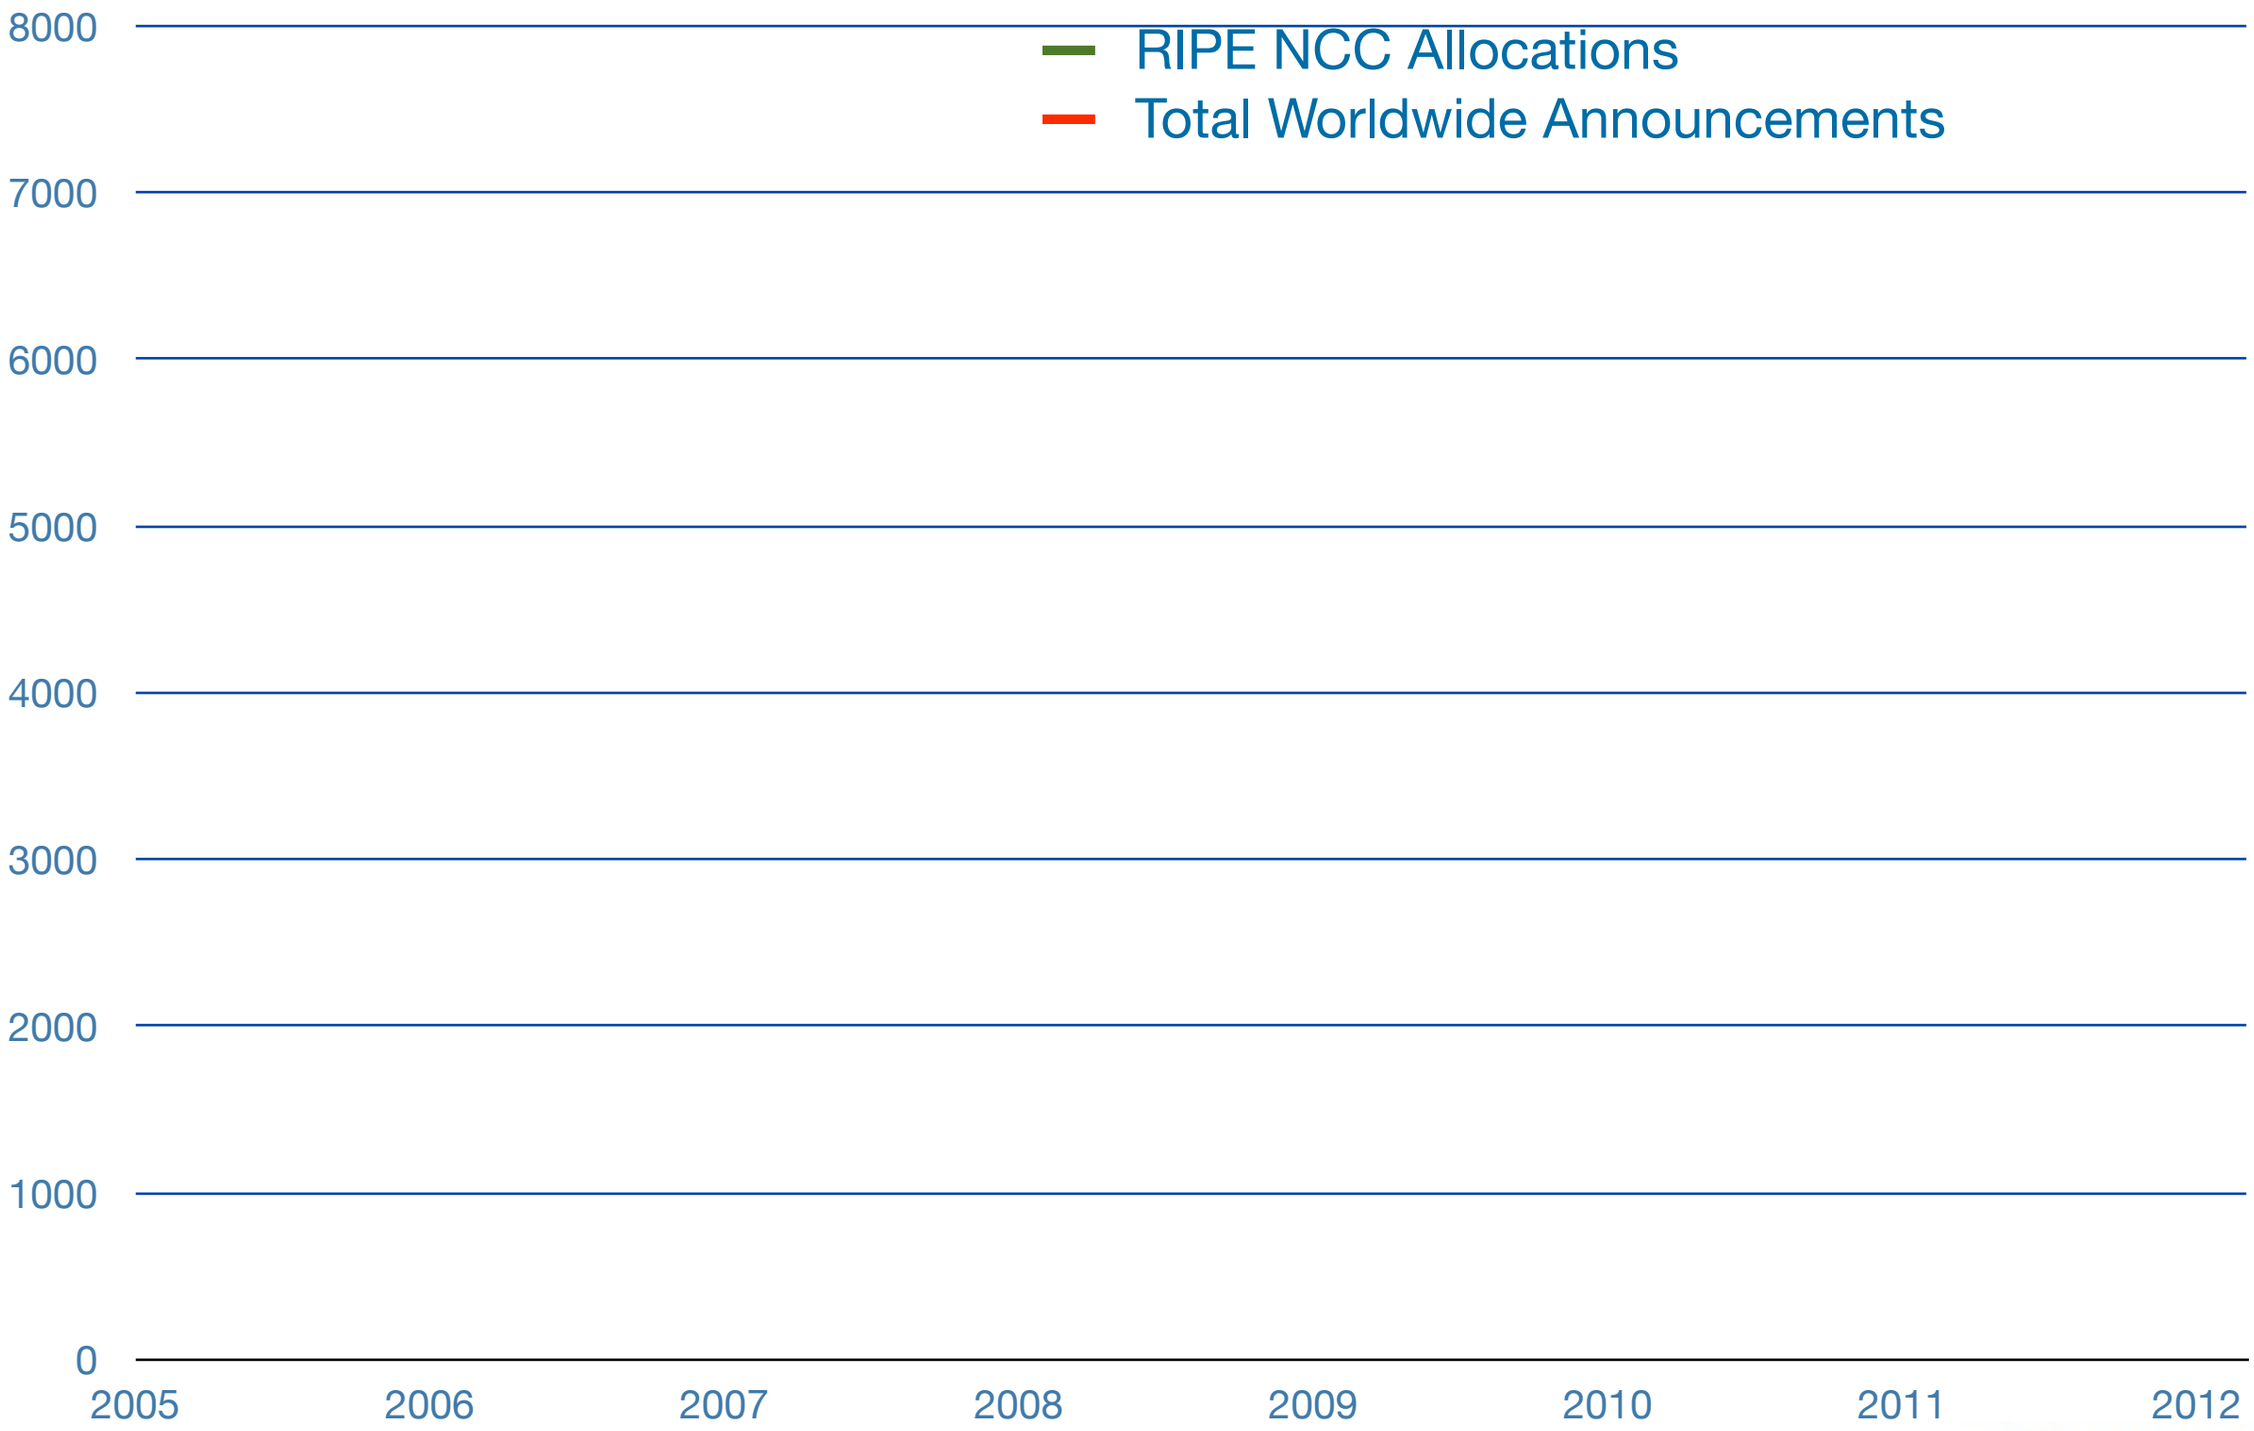

IPv6 Allocations and Announcements

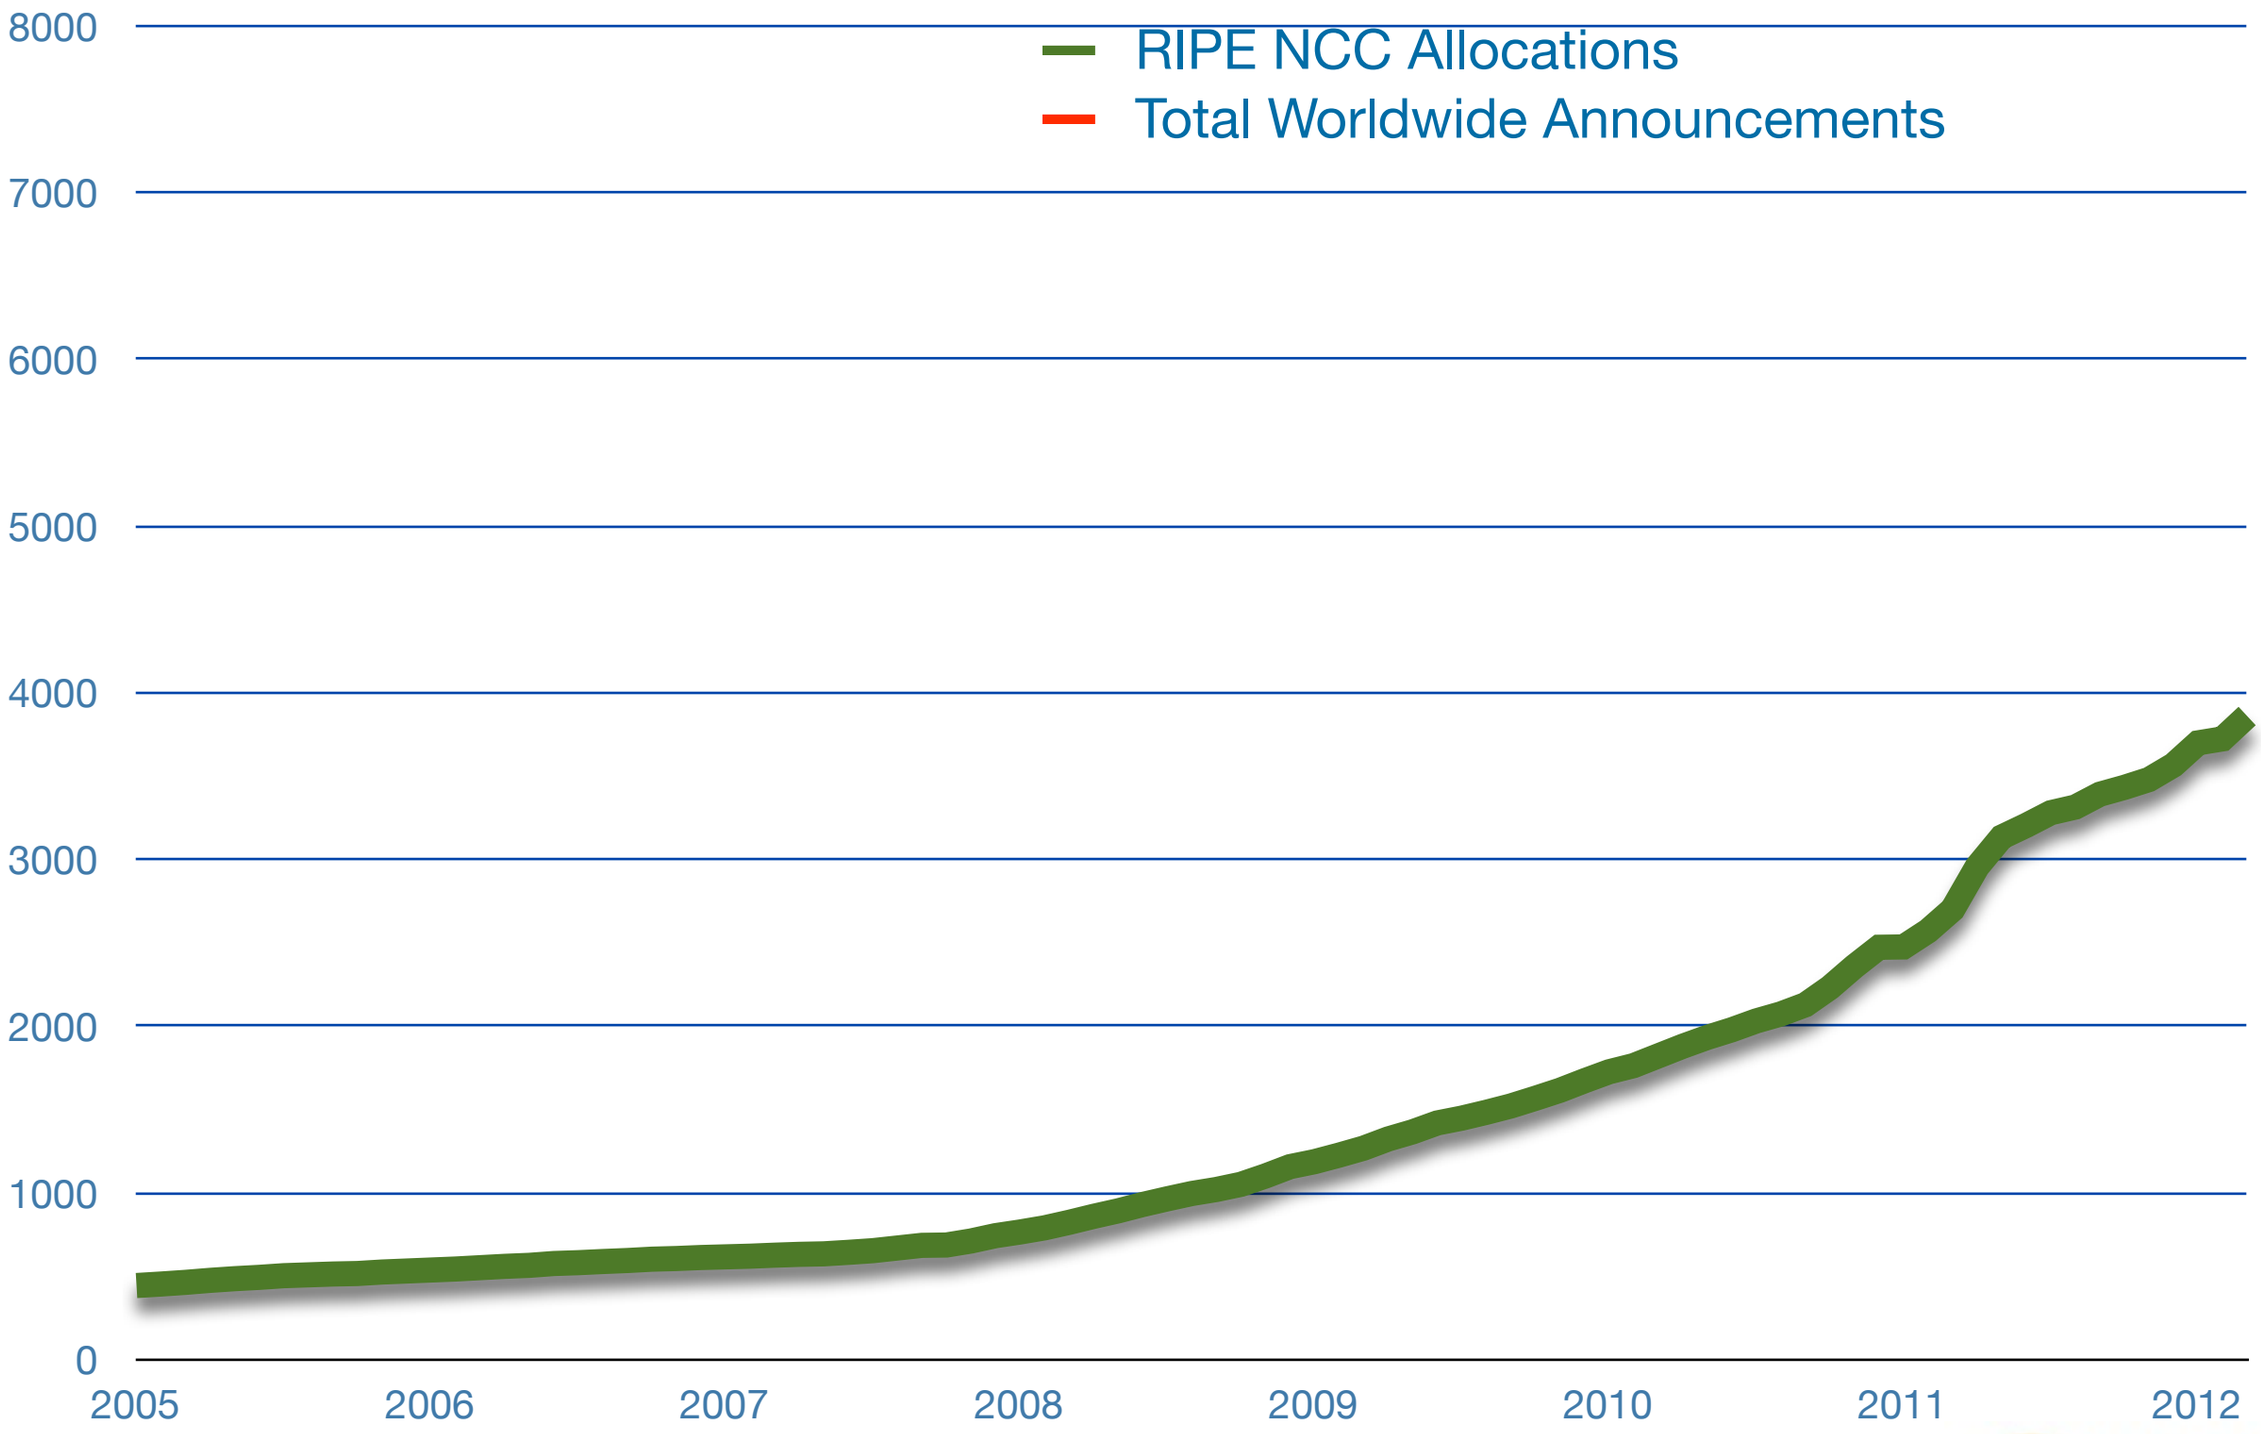

IPv6 Allocations and Announcements

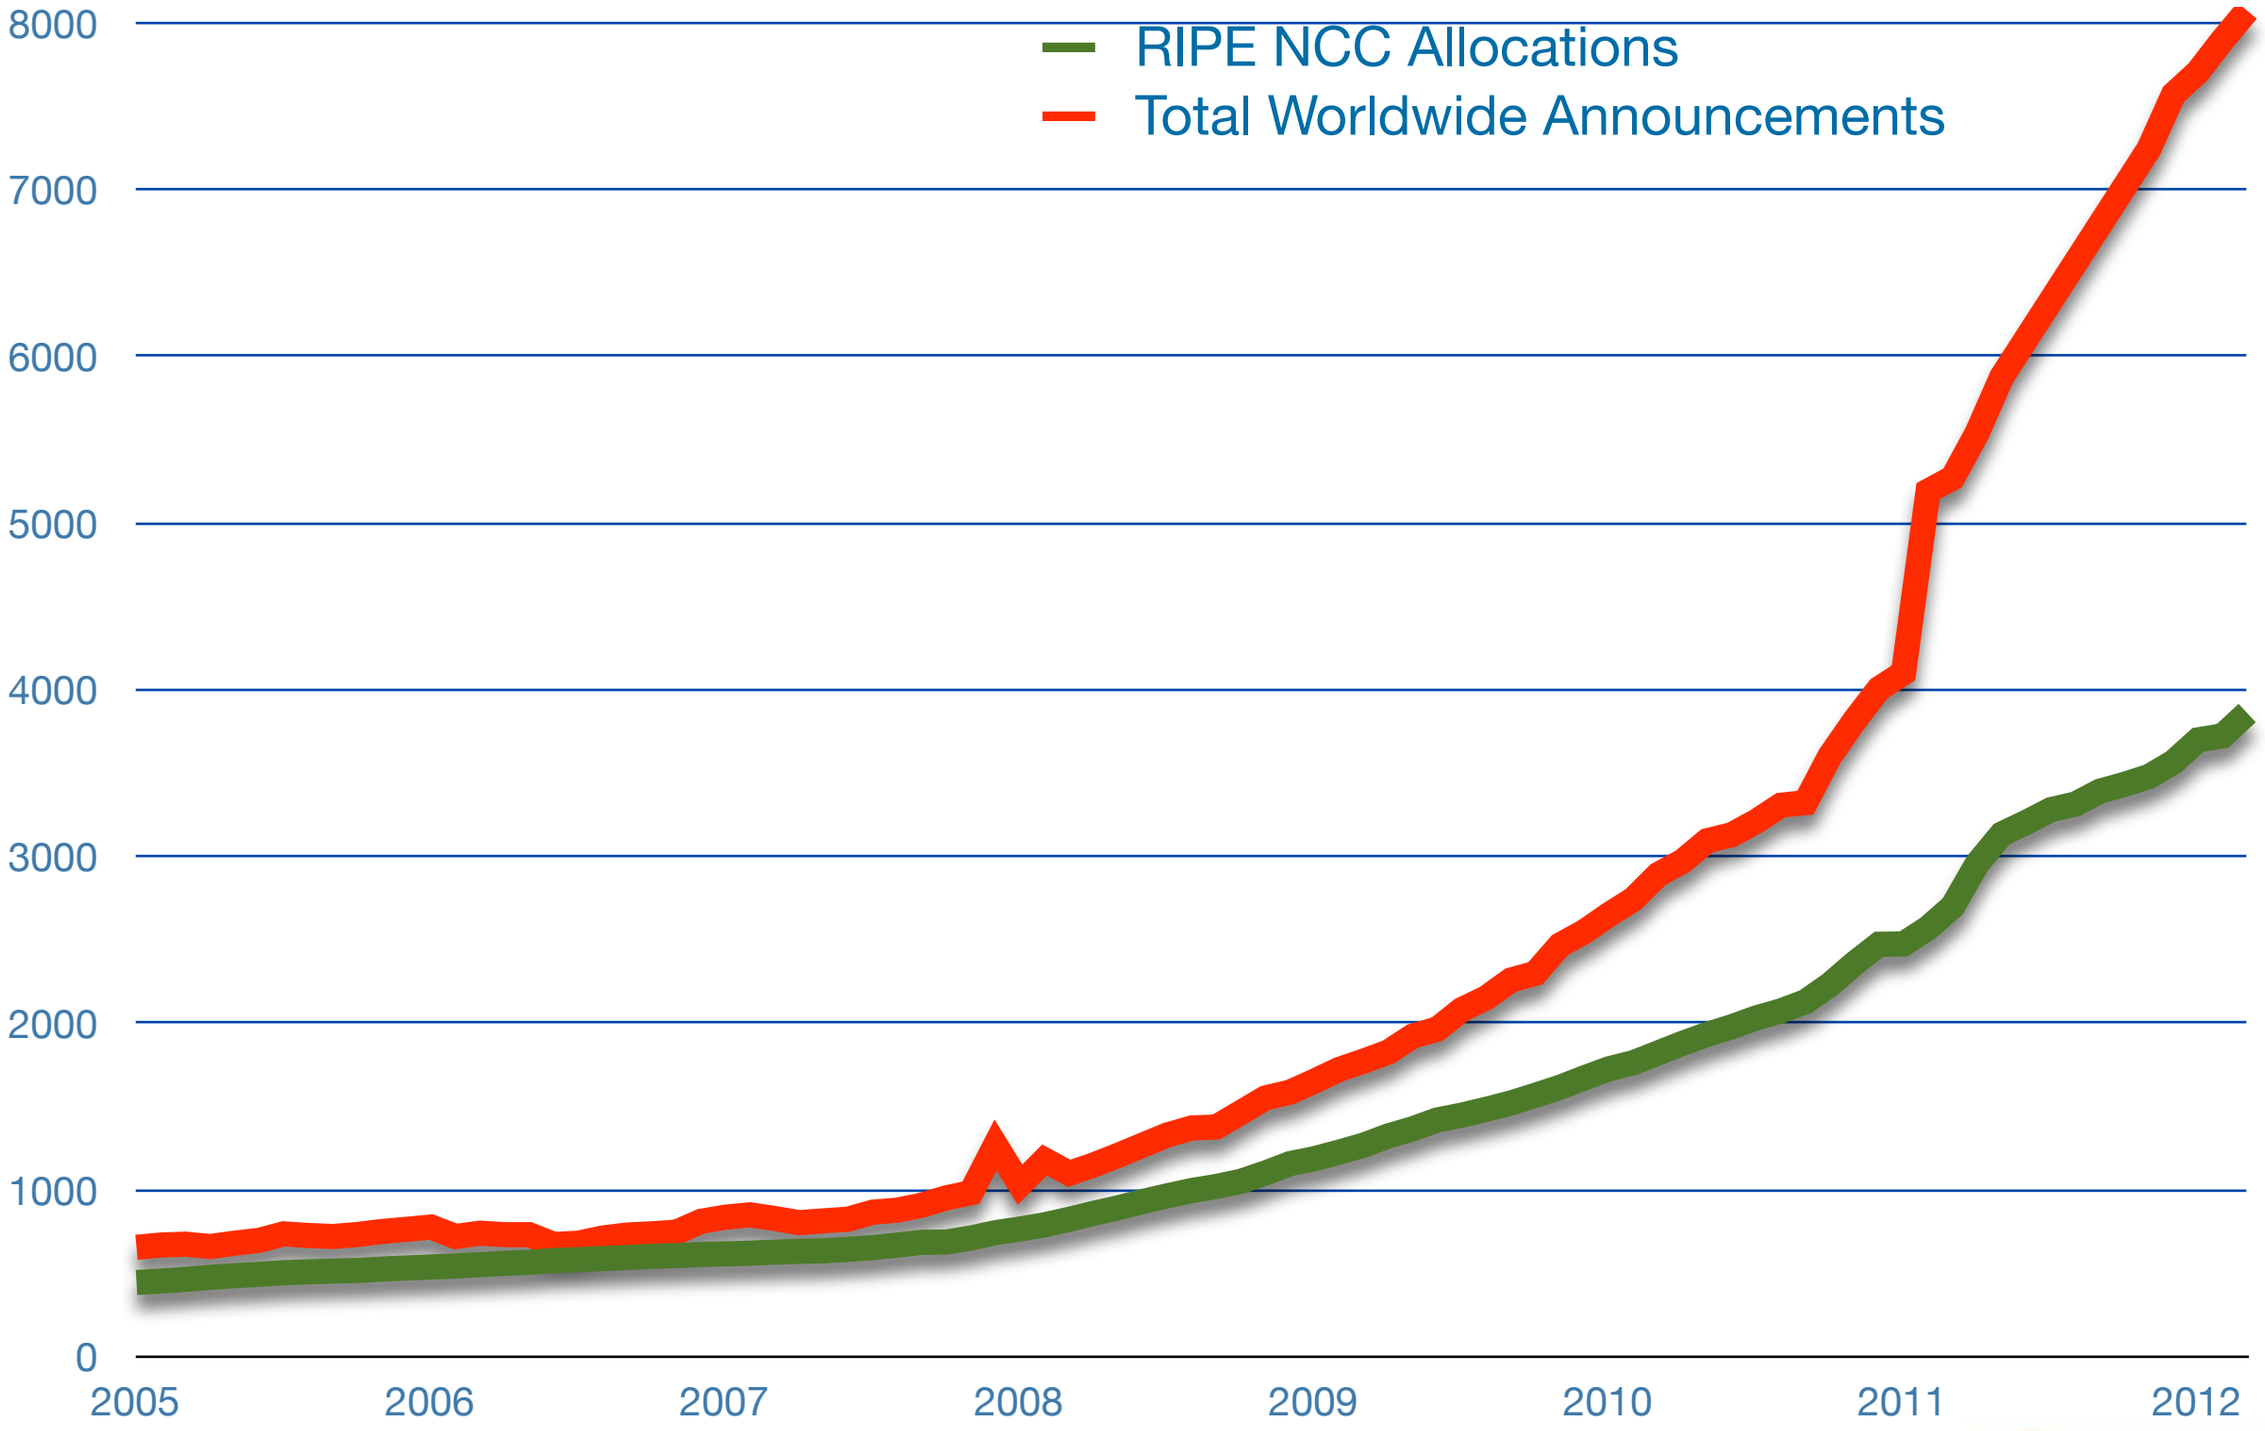

IPv6 Ripeness

- 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: 7964 LIRs

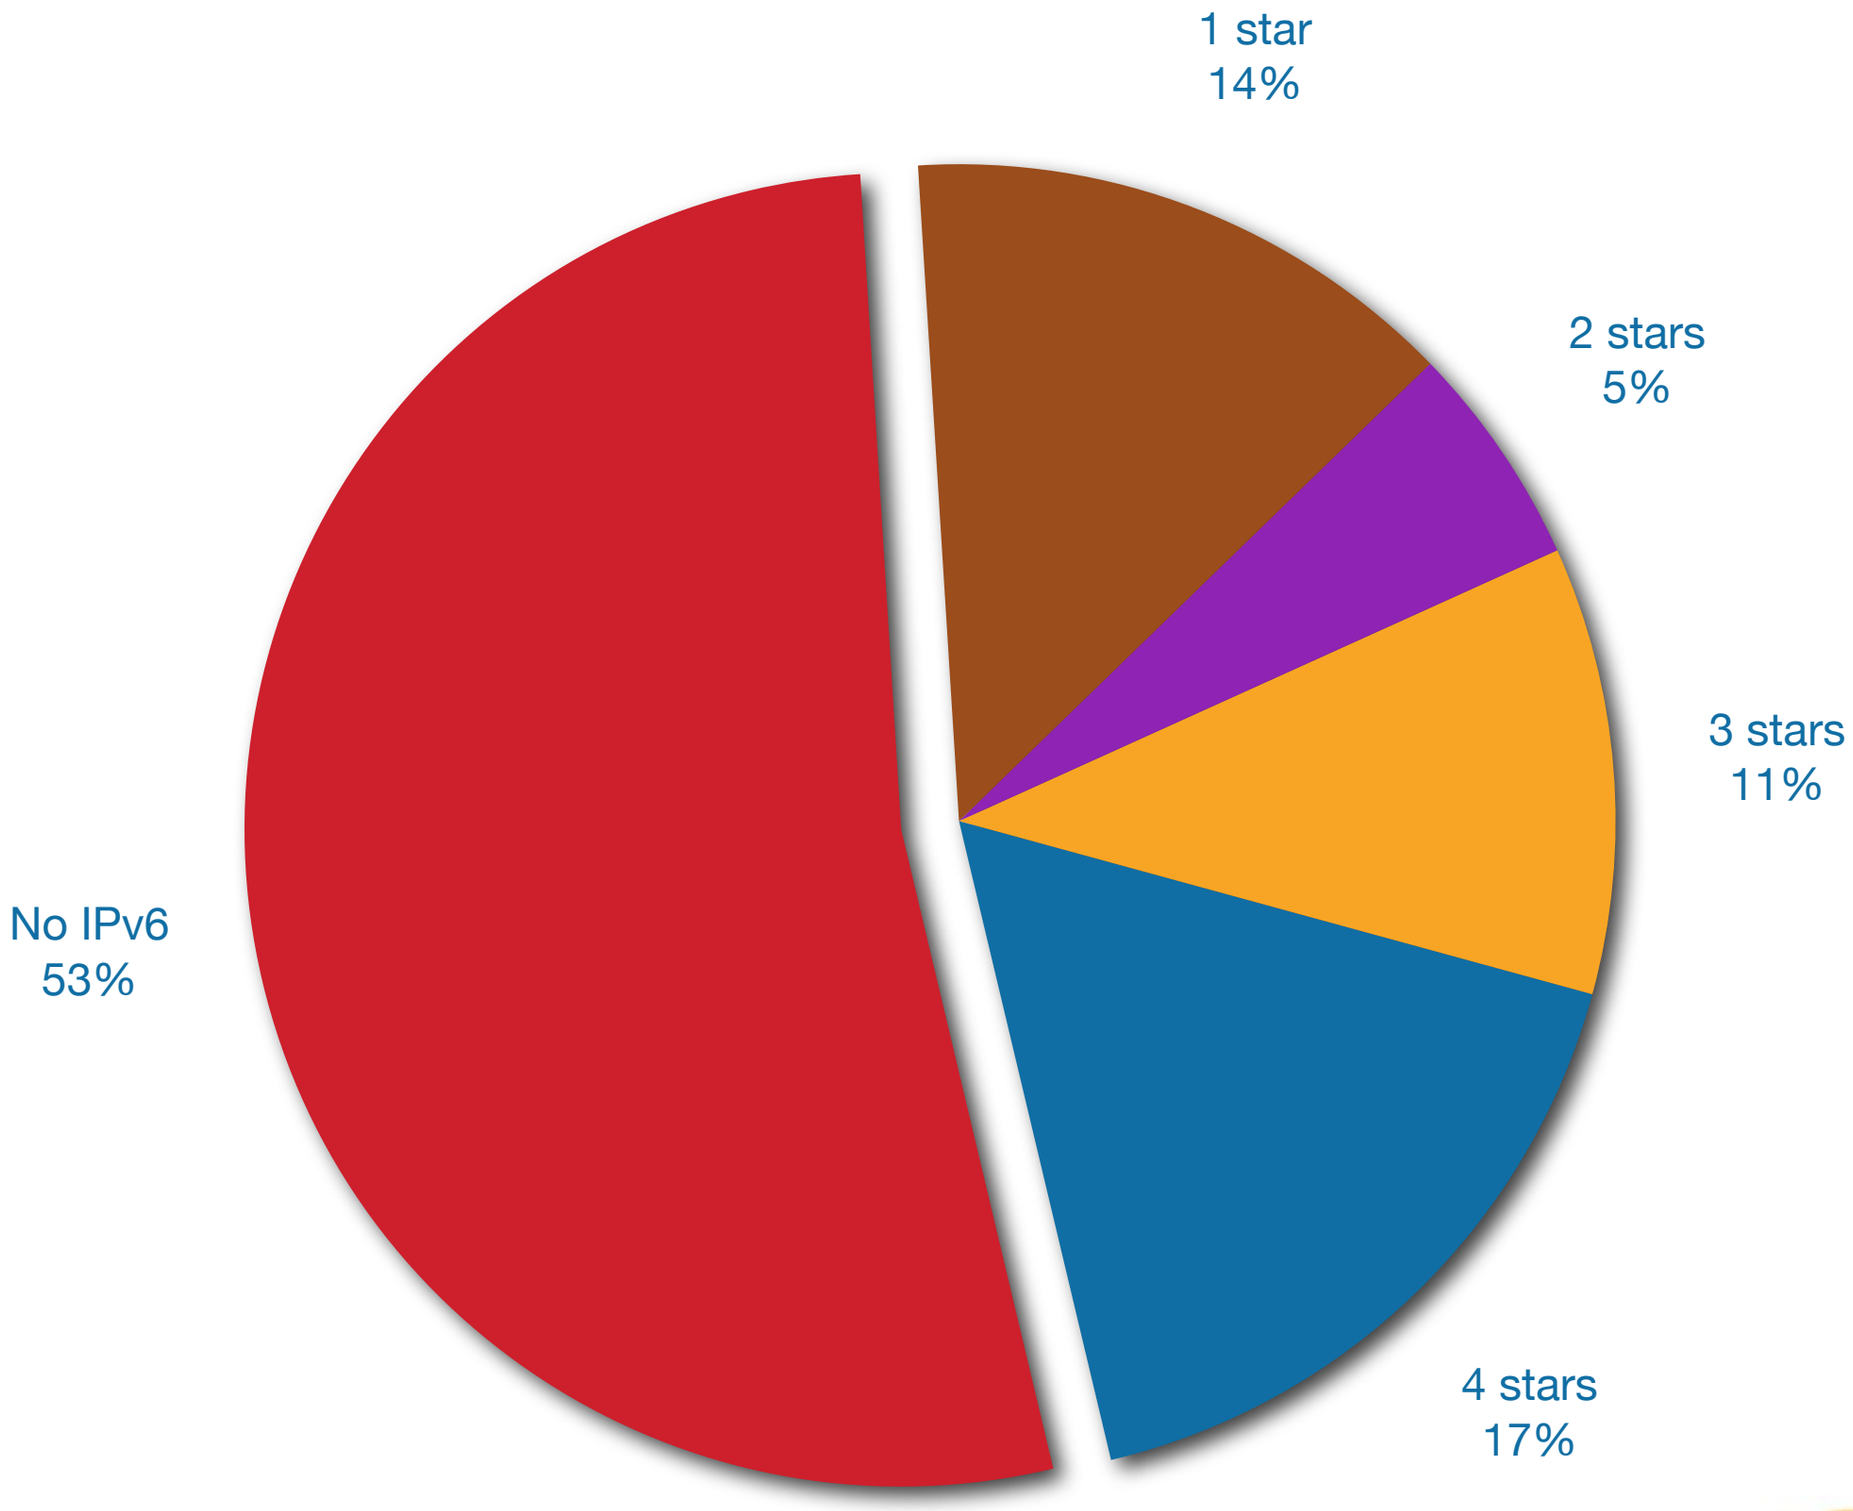

What does an IPv6 allocation cost?

- /32 = 1 scoring unit
- /31 = 2 scoring units
- points = ∑(2012-1992)x(scoring unit) =20x1+...

Category	Points	Fee 2012
Extra Small	0 - 16	€ 1300
Small	- 109	€ 1800
Medium	- 1040	€ 2550
Large	- 7728	€ 4100
Extra Large	> 7728	€ 5500

What does an IPv6 allocation cost?

- /32 = 1 scoring unit
- /31 = 2 scoring units
- points = ∑(2012-1992)x(scoring units) =20x1+....

Category	Points	Fee 2012
Small		€ 13
S	09	
Me n	0	€ 2550
La	- 7728	€ 4100
Extra Large	> 7	€ 5500

Getting IPv6 PI address space

- To qualify, an organisation must:
  - Meet the contractual requirements for provider independent resources
  - LIRs must demonstrate special routing requirements

- Minimum assignment size /48

- PI space can not be used for sub-assignments

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 now be registered in the RIPE database 

Using AGGREGATED-BY-LIR

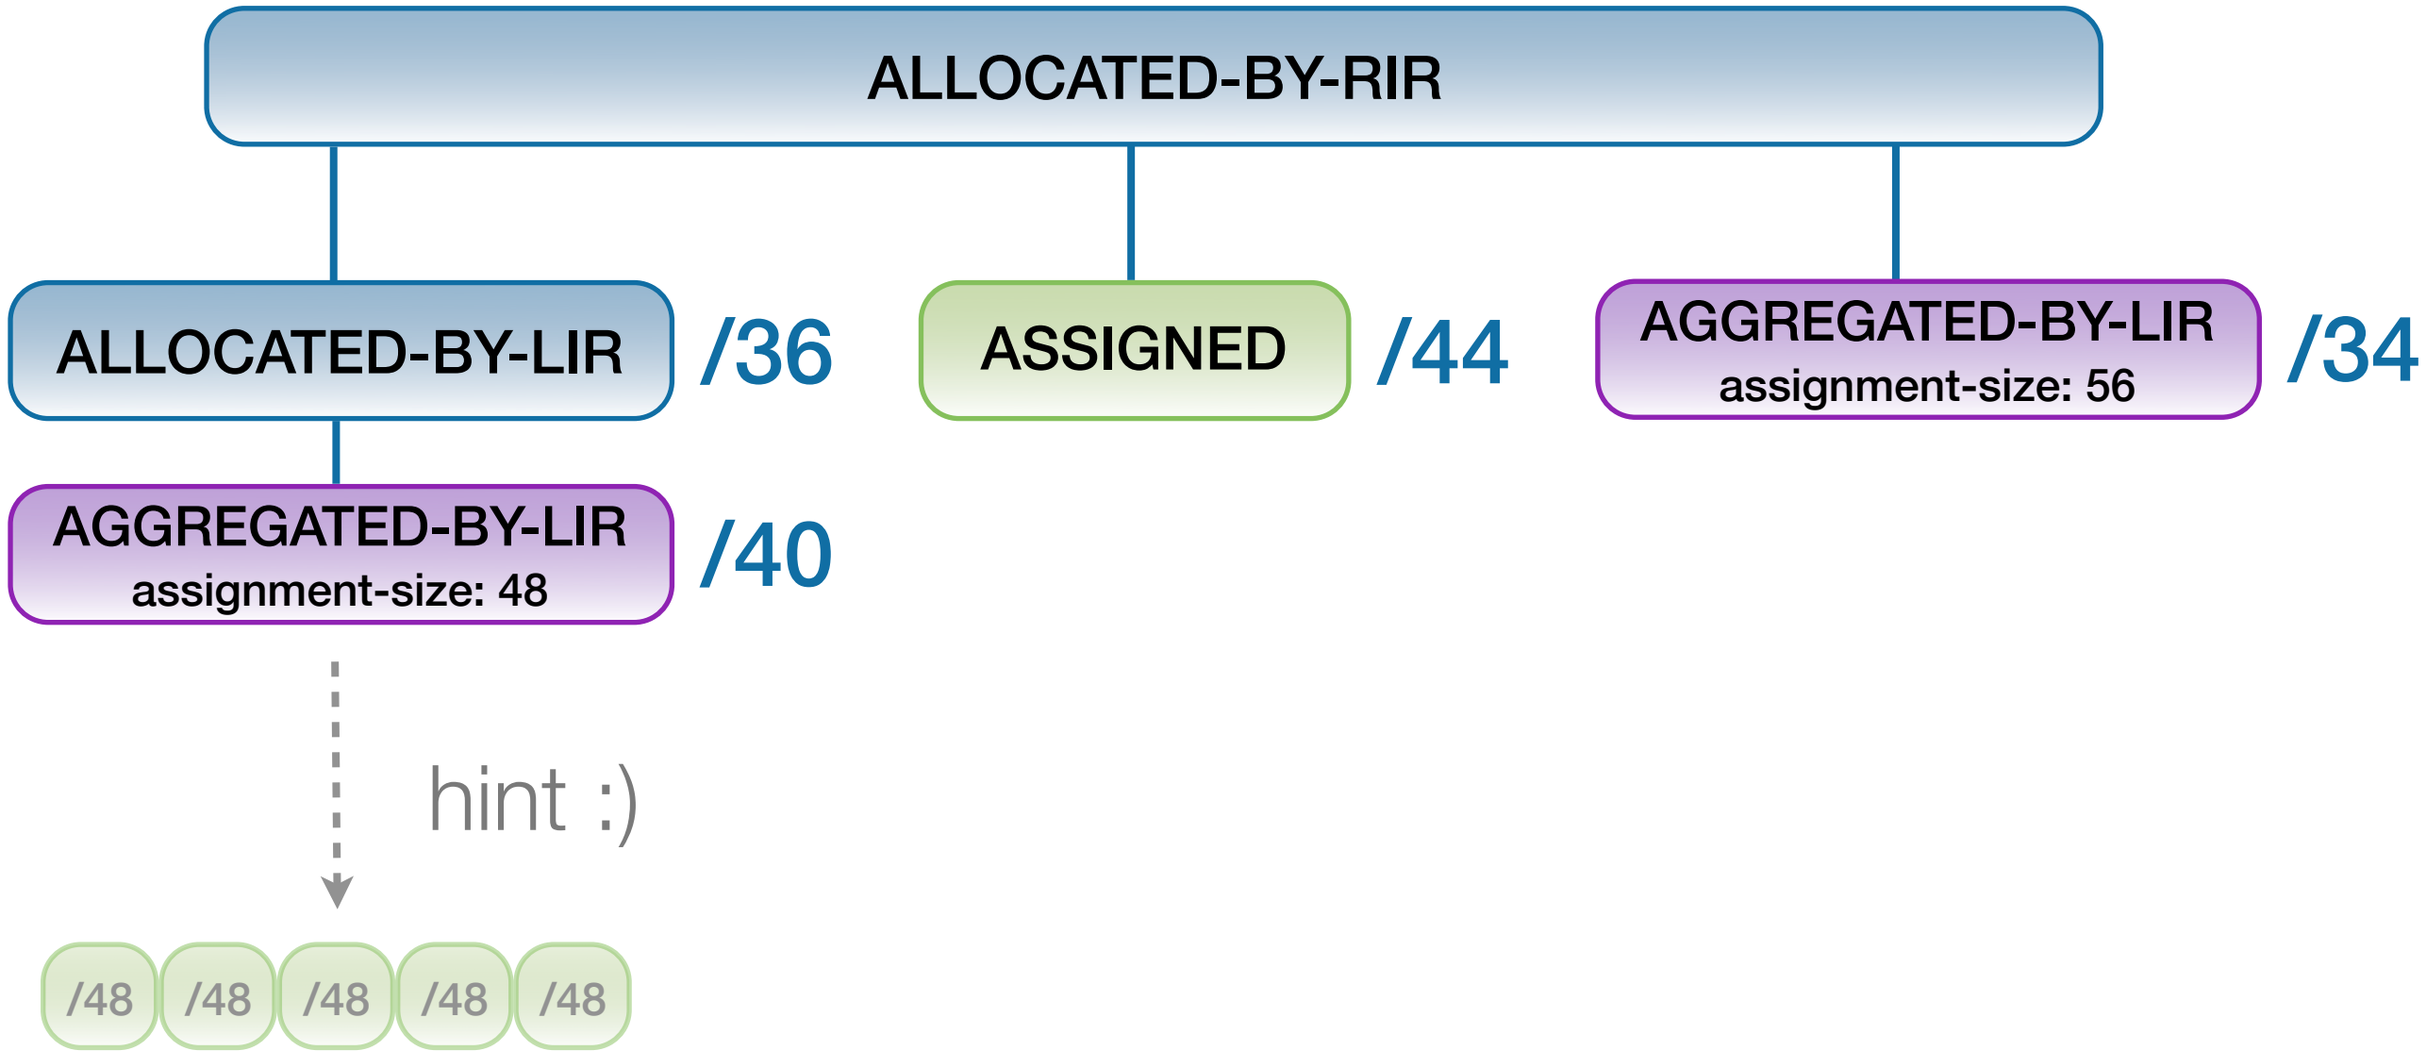

Group assignments in the RIPE DB

inet6num:	2001:db8:1000::/36
netname:	Bluelight
descr:	We want more Bluelight B.V.
descr:	Colocation services
country:	NL
admin-c:	BN649-RIPE
tech-c:	BN649-RIPE
status:	**AGGREGATED-BY-LIR**
assignment-size:	**48**
mnt-by:	BLUELIGHT-MNT
notify:	noc@example.net
changed:	noc@example.net 20110218
source:	RIPE

Exercise

Making Assignments

Making Assignments Exercise

Fridge6!

- 40 minutes preparation time
- 10 minutes discussion

Solution RIPE Database object

inet6num:	2001:db8:1000::/40
netname:	FREEZ
descr:	Freez Fridges
country:	NL
admin-c:	RM1204-RIPE
tech-c:	RM1204-RIPE
status:	**AGGREGATED-BY-LIR**
assignment-size:	**60**
mnt-by:	BLUELIGHT-MNT
notify:	noc@freez.net
changed:	noc@freez.net 20110801
source:	RIPE

Deploying

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

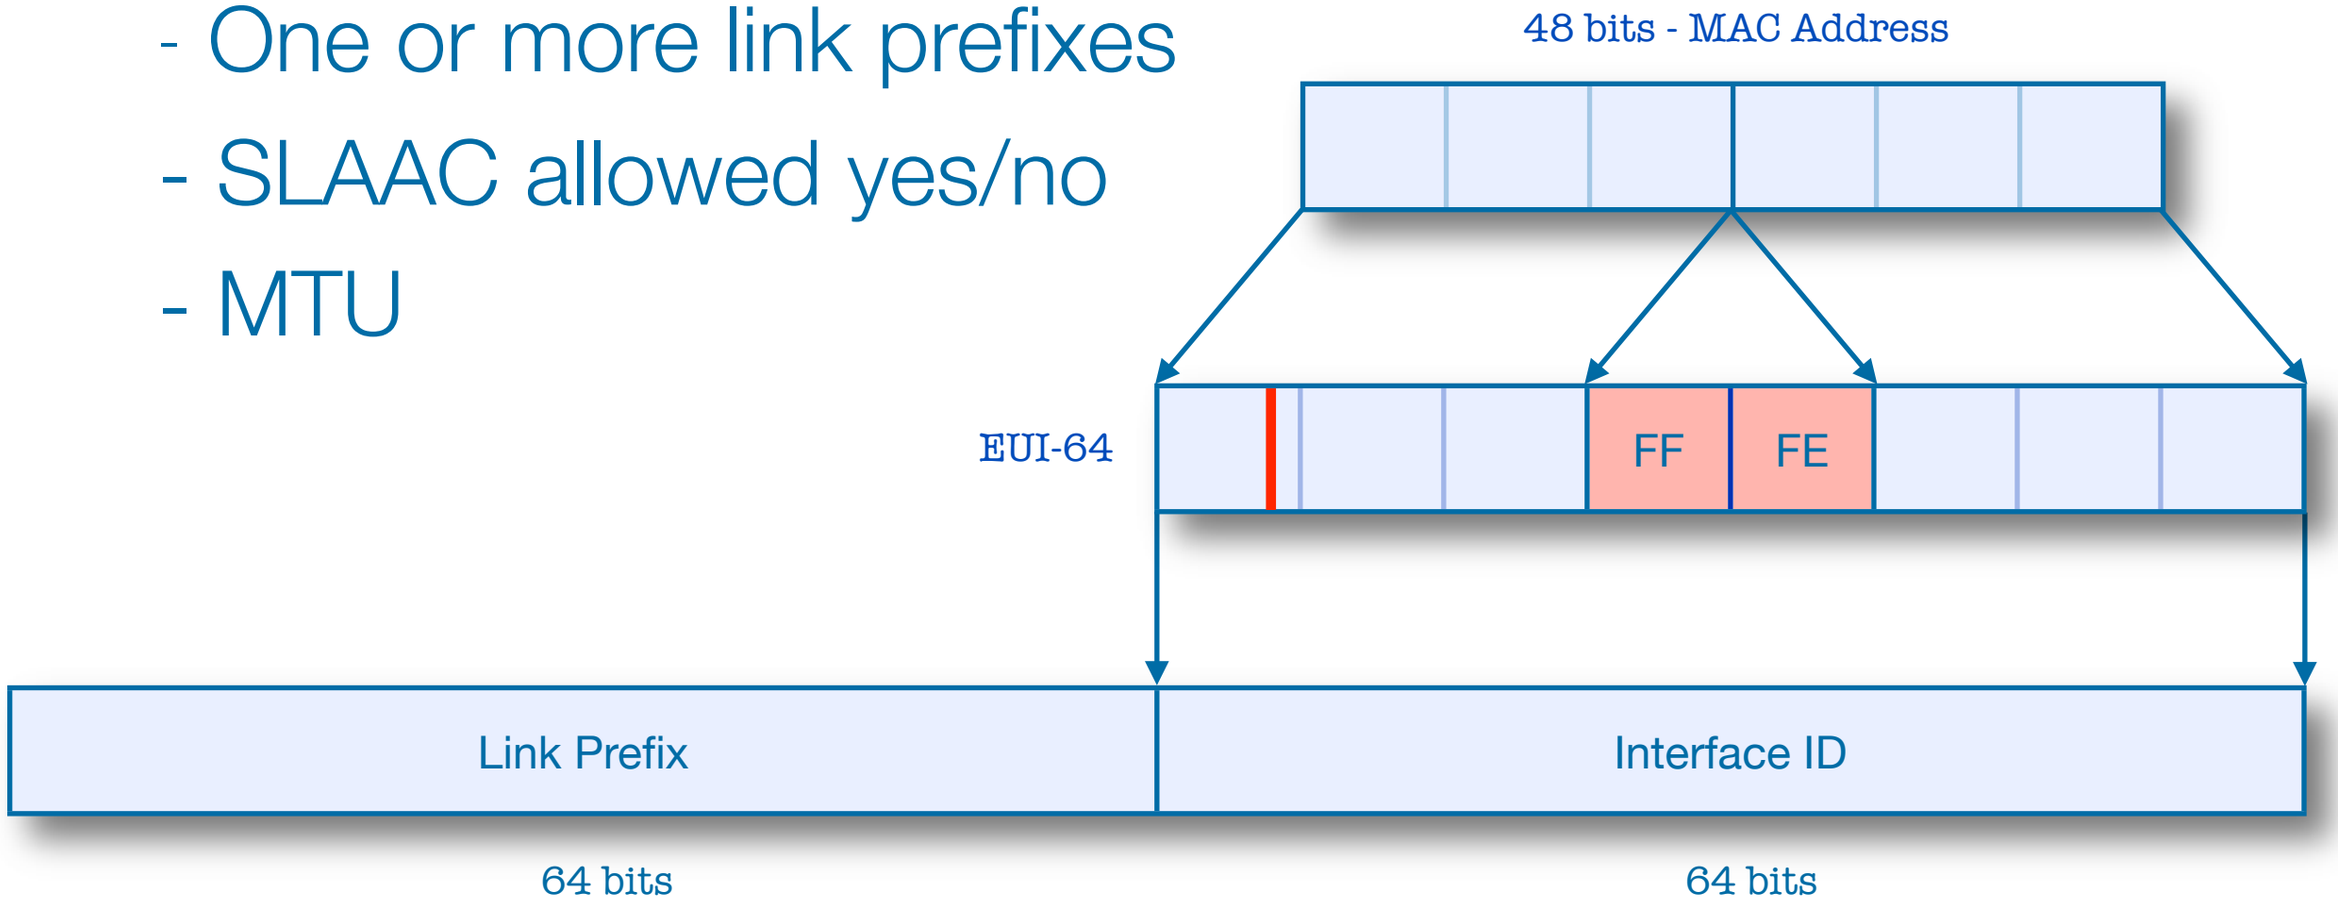

DHCPv6

- 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

Interfaces will have multiple addresses

- Unicast
  - Link Local fe80::5a55:caff:fef6:bdbf/64
  - Global Unicast 2001::5a55:caff:fef6:bdbf/64 (multiple)
- Multicast
  - All Nodes ff02::1 (scope: link)
  - Solicited Node ff02::1:fff6:bdbf (scope: link)
- Routers
  - All Routers ff02::2 (scope: link)

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

Reverse DNS

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

Reverse DNS

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

8.b.d.0.1.0.0.2.ip6.arpa

Reverse DNS

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

8.b.d.0.1.0.0.2.ip6.arpa

d.4.0.0.0.0.1.c.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.

Reverse DNS

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

8.b.d.0.1.0.0.2.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.

IPv6 in the Routing Registry

Route6 object:

```
route6:     2001:db8::/32
origin:     AS65550
```

Aut-num object:

```
aut-num:    AS65550
mp-import:  afi ipv6.unicast from AS64496 accept ANY
mp-export:  afi ipv6.unicast to AS64496 announce AS65550
```

Security considerations

- Privacy Extensions for Address Configuration in IPv6 - RFC 4941
- Secure Neighbor Discovery (SEND)
  - http://tools.ietf.org/html/rfc3971
- RA (Router Advertisement) Guard
  - http://tools.ietf.org/html/rfc6105

Exercise

Deployment Challenges

Deployment Challenges

- Think of 5 challenges/problems your organisation could have when you deploy IPv6 and write them on the exercise sheet
- Let's see if you can find solutions!

- 10 minutes for challenges
- 10 minutes for solutions

Real Life IPv6 Deployment

Colocation Provider

- 30 staff
- Routing
  - Dual stack or separate routers?
  - Possible IGP combinations are:
    - OSPFv2 for IPv4, IS-IS for IPv6 (only)
    - OSPFv2 for IPv4, OSPFv3 for IPv6
    - IS-IS for IPv4, OSPFv3 for IPv6
    - IS-IS for both IPv4 and IPv6 (their solution)
  - Check internal routing before going external!

Colocation Provider

- Checklist
  - set access lists on network equipment
  - set up monitoring (SNMP)
  - have working DNS
- Subnetting tools
  - sipcalc, IPv6calc
- Every customer gets a /48 assignment
  - and a /64 for the connection

Colocation Provider

- Points of attention:
  - stateless auto configuration can assign a subnet “unexpected”
  - Google IPv6 whitelist
  - not all firewalls support IPv6
  - be careful with statement “IPv6 ready”

ISP xDSL

- 200 staff
- 2 /32 prefixes (due to merger)
  - not enough
  - make a plan before request allocation

- /48 per POP
- /56 per router
- /64 per customer vlan

ISP xDSL

- Servers
  - no EUI-64
  - no autoconfig
  - port number for services (i.e. POP3 at ::110)
  - default gateway manually set to :1/64 (usually)

ISP xDSL

- Network links (point-to-point)
  - core
    - /64 per link
    - ::1 - ::2
    - no auto configuration
    - easy to remember

- You don’t want your router link at:

  2001:DB8:CF9D:7631:CD01:FE55:4532:AE60/127

  You want your router link at:

  2001:DB8:1:1::/127

Tips

Best Scenario: Act Now, Phased Approach

- Change purchasing procedure (feature parity)
- 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!”

Also useful

Websites

- http://www.getipv6.info/
- http://datatracker.ietf.org/wg/v6ops/
- http://www.ripe.net/ripe/docs/ripe-501.html

Mailing lists

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

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

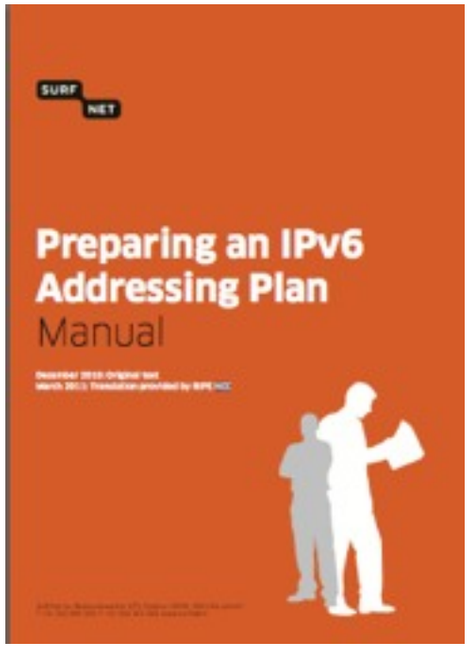

Survey!

http://www.ripe.net/training/lir/survey

Follow us!

@TrainingRIPENCC

The End!

Край

Y Diwedd

Fí

النهاية

Соңы

Finis

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