IPv6 for LIRs

Training Course | April 2014

Training Services
RIPE NCC

Schedule

09:00 - 09:30	**Coffee, Tea**
11:00 - 11:45	**Break**
13:00 - 14:00	**Lunch**
15:30 - 15:45	**Break**
17:30	**End**

Introductions

- Name
- Number in the list
- Experience with IPv6
- Goals

Overview

- IPv4?
- IPv6 Address Basics
- Getting it
- Exercise: Making Assignments
- Transition Mechanisms
- Exercise: Addressing Plan
- Deploying
- Exercise: Configuring IPv6
- Real Life IPv6 Deployment
- Deployment Challenges
- Tips

IPv4?

Section 1

RIPE NCC

IANA IPv4 Pool

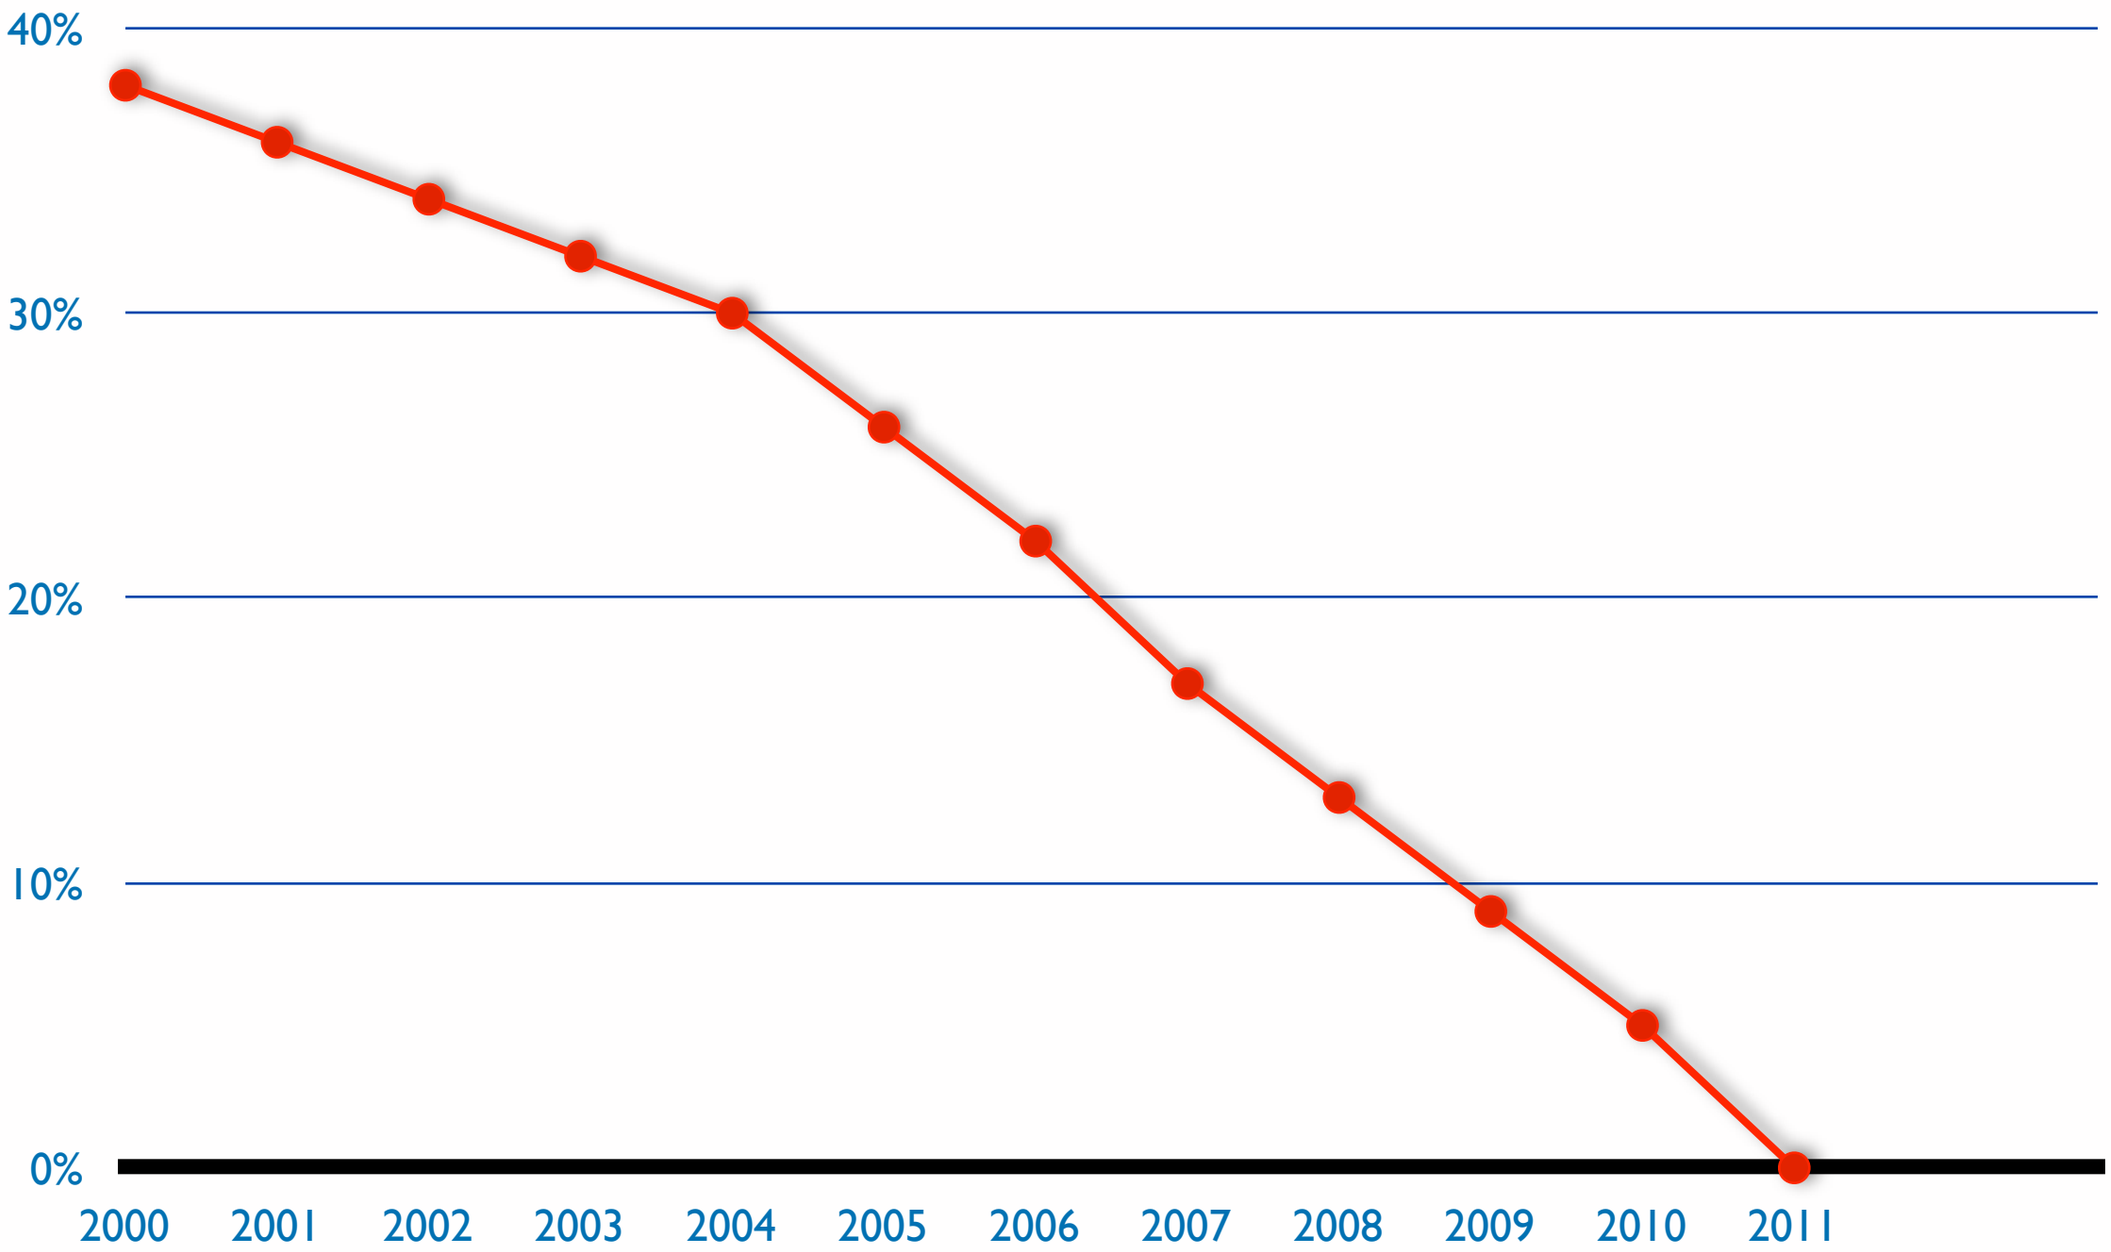

RIPE NCC

“On 14 September 2012, the RIPE NCC ran out of their regular pool of IPv4”

RIPE NCC

Reaching the next billion

- Around 2.4 billion internet users now
  - around 35% of all people
- Mobile phones are becoming internet devices
- The Internet of things
  - How will the Internet look like in 5 years?

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

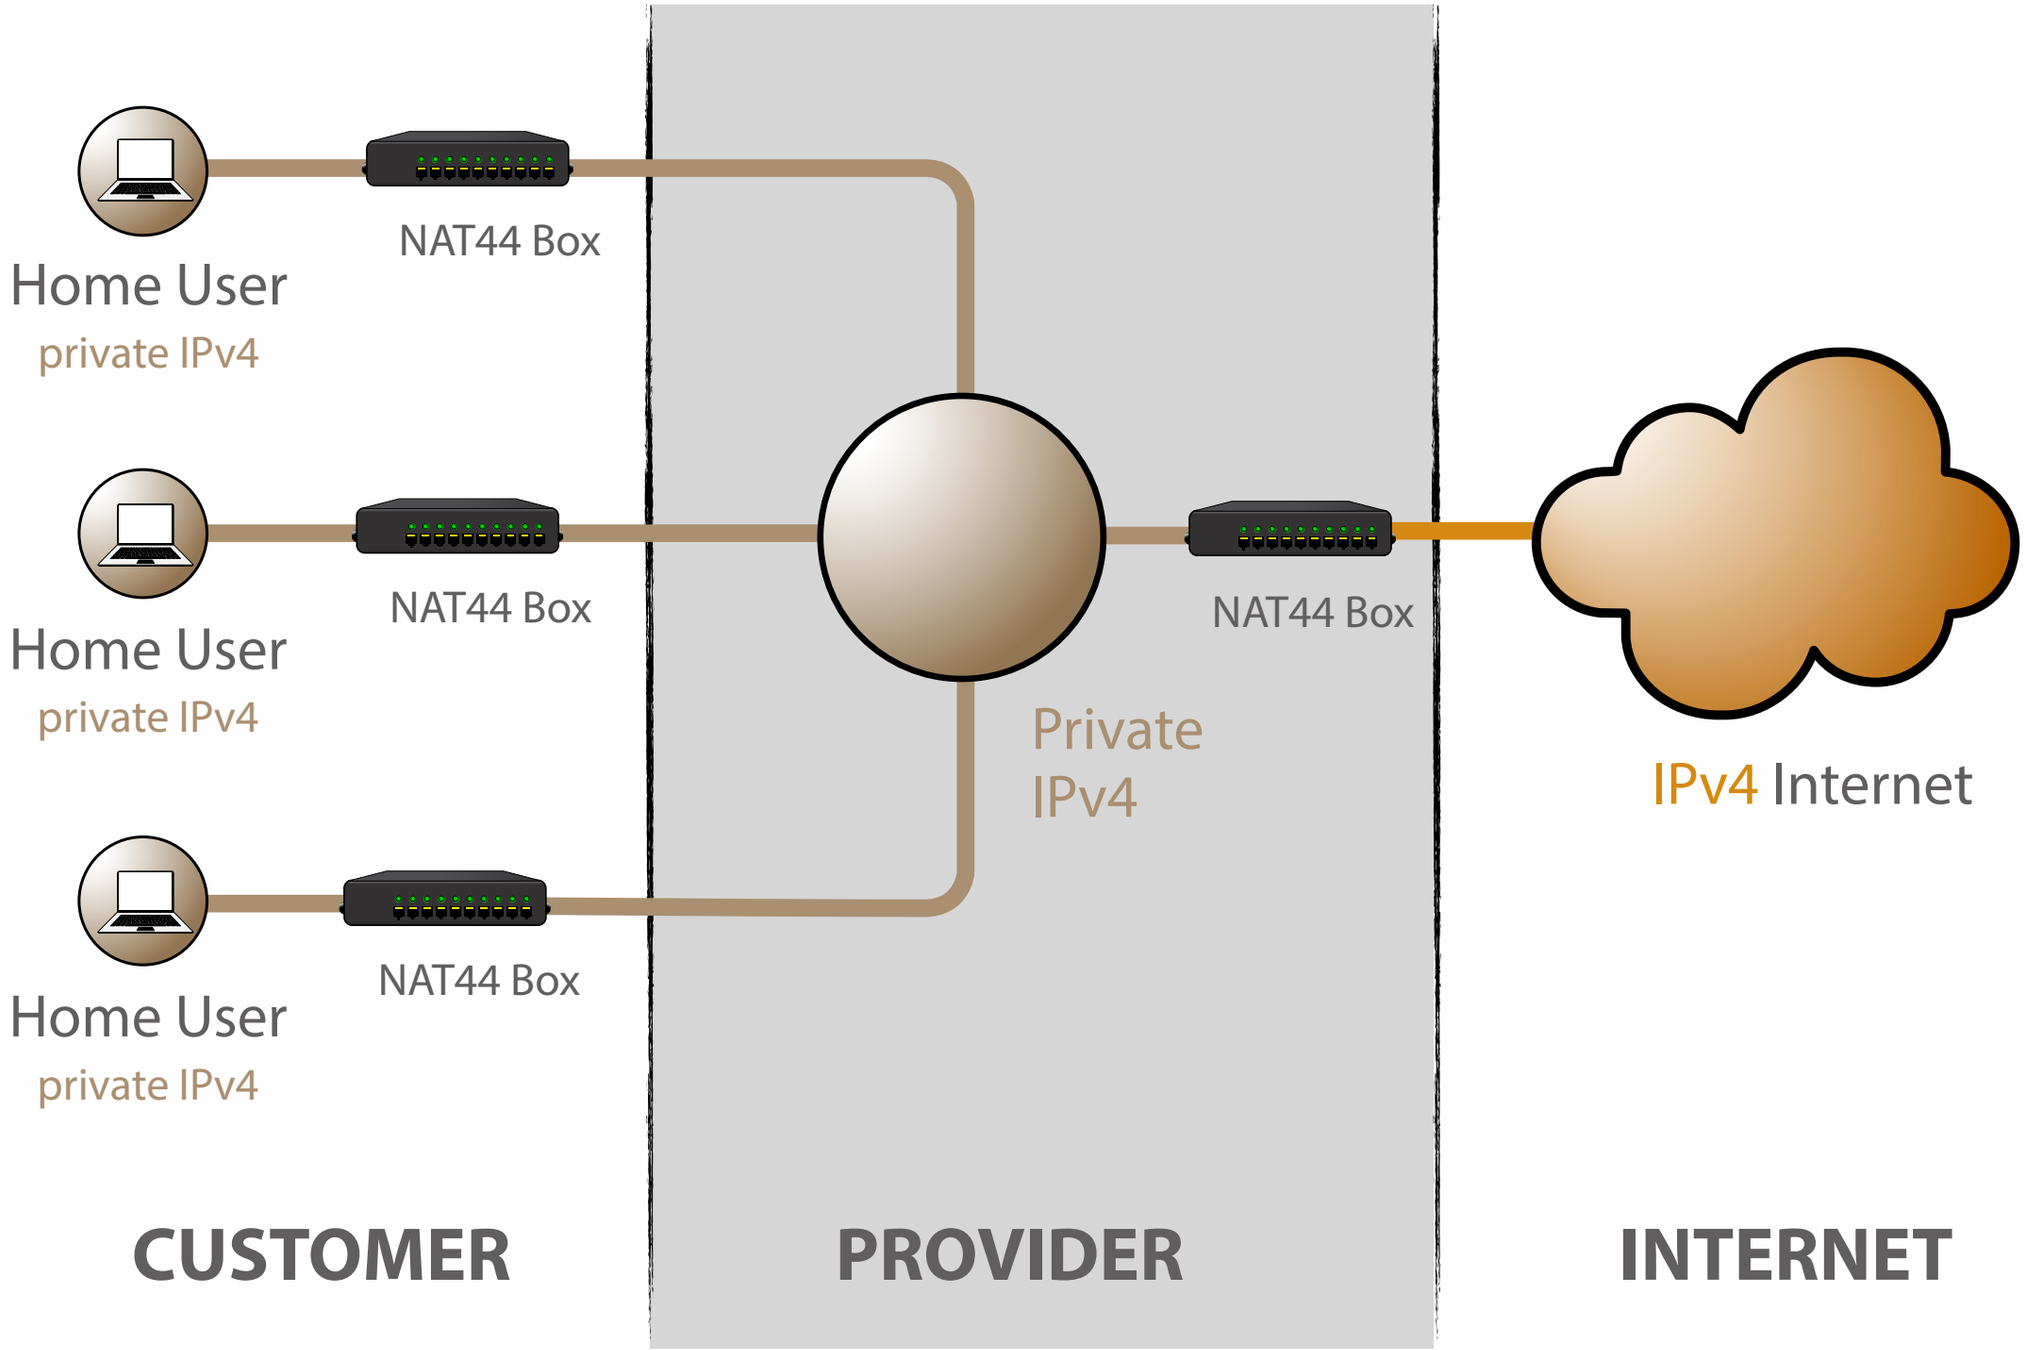

?

!

IPv6 Address Basics

Section 2

RIPE NCC

IP Address Distribution

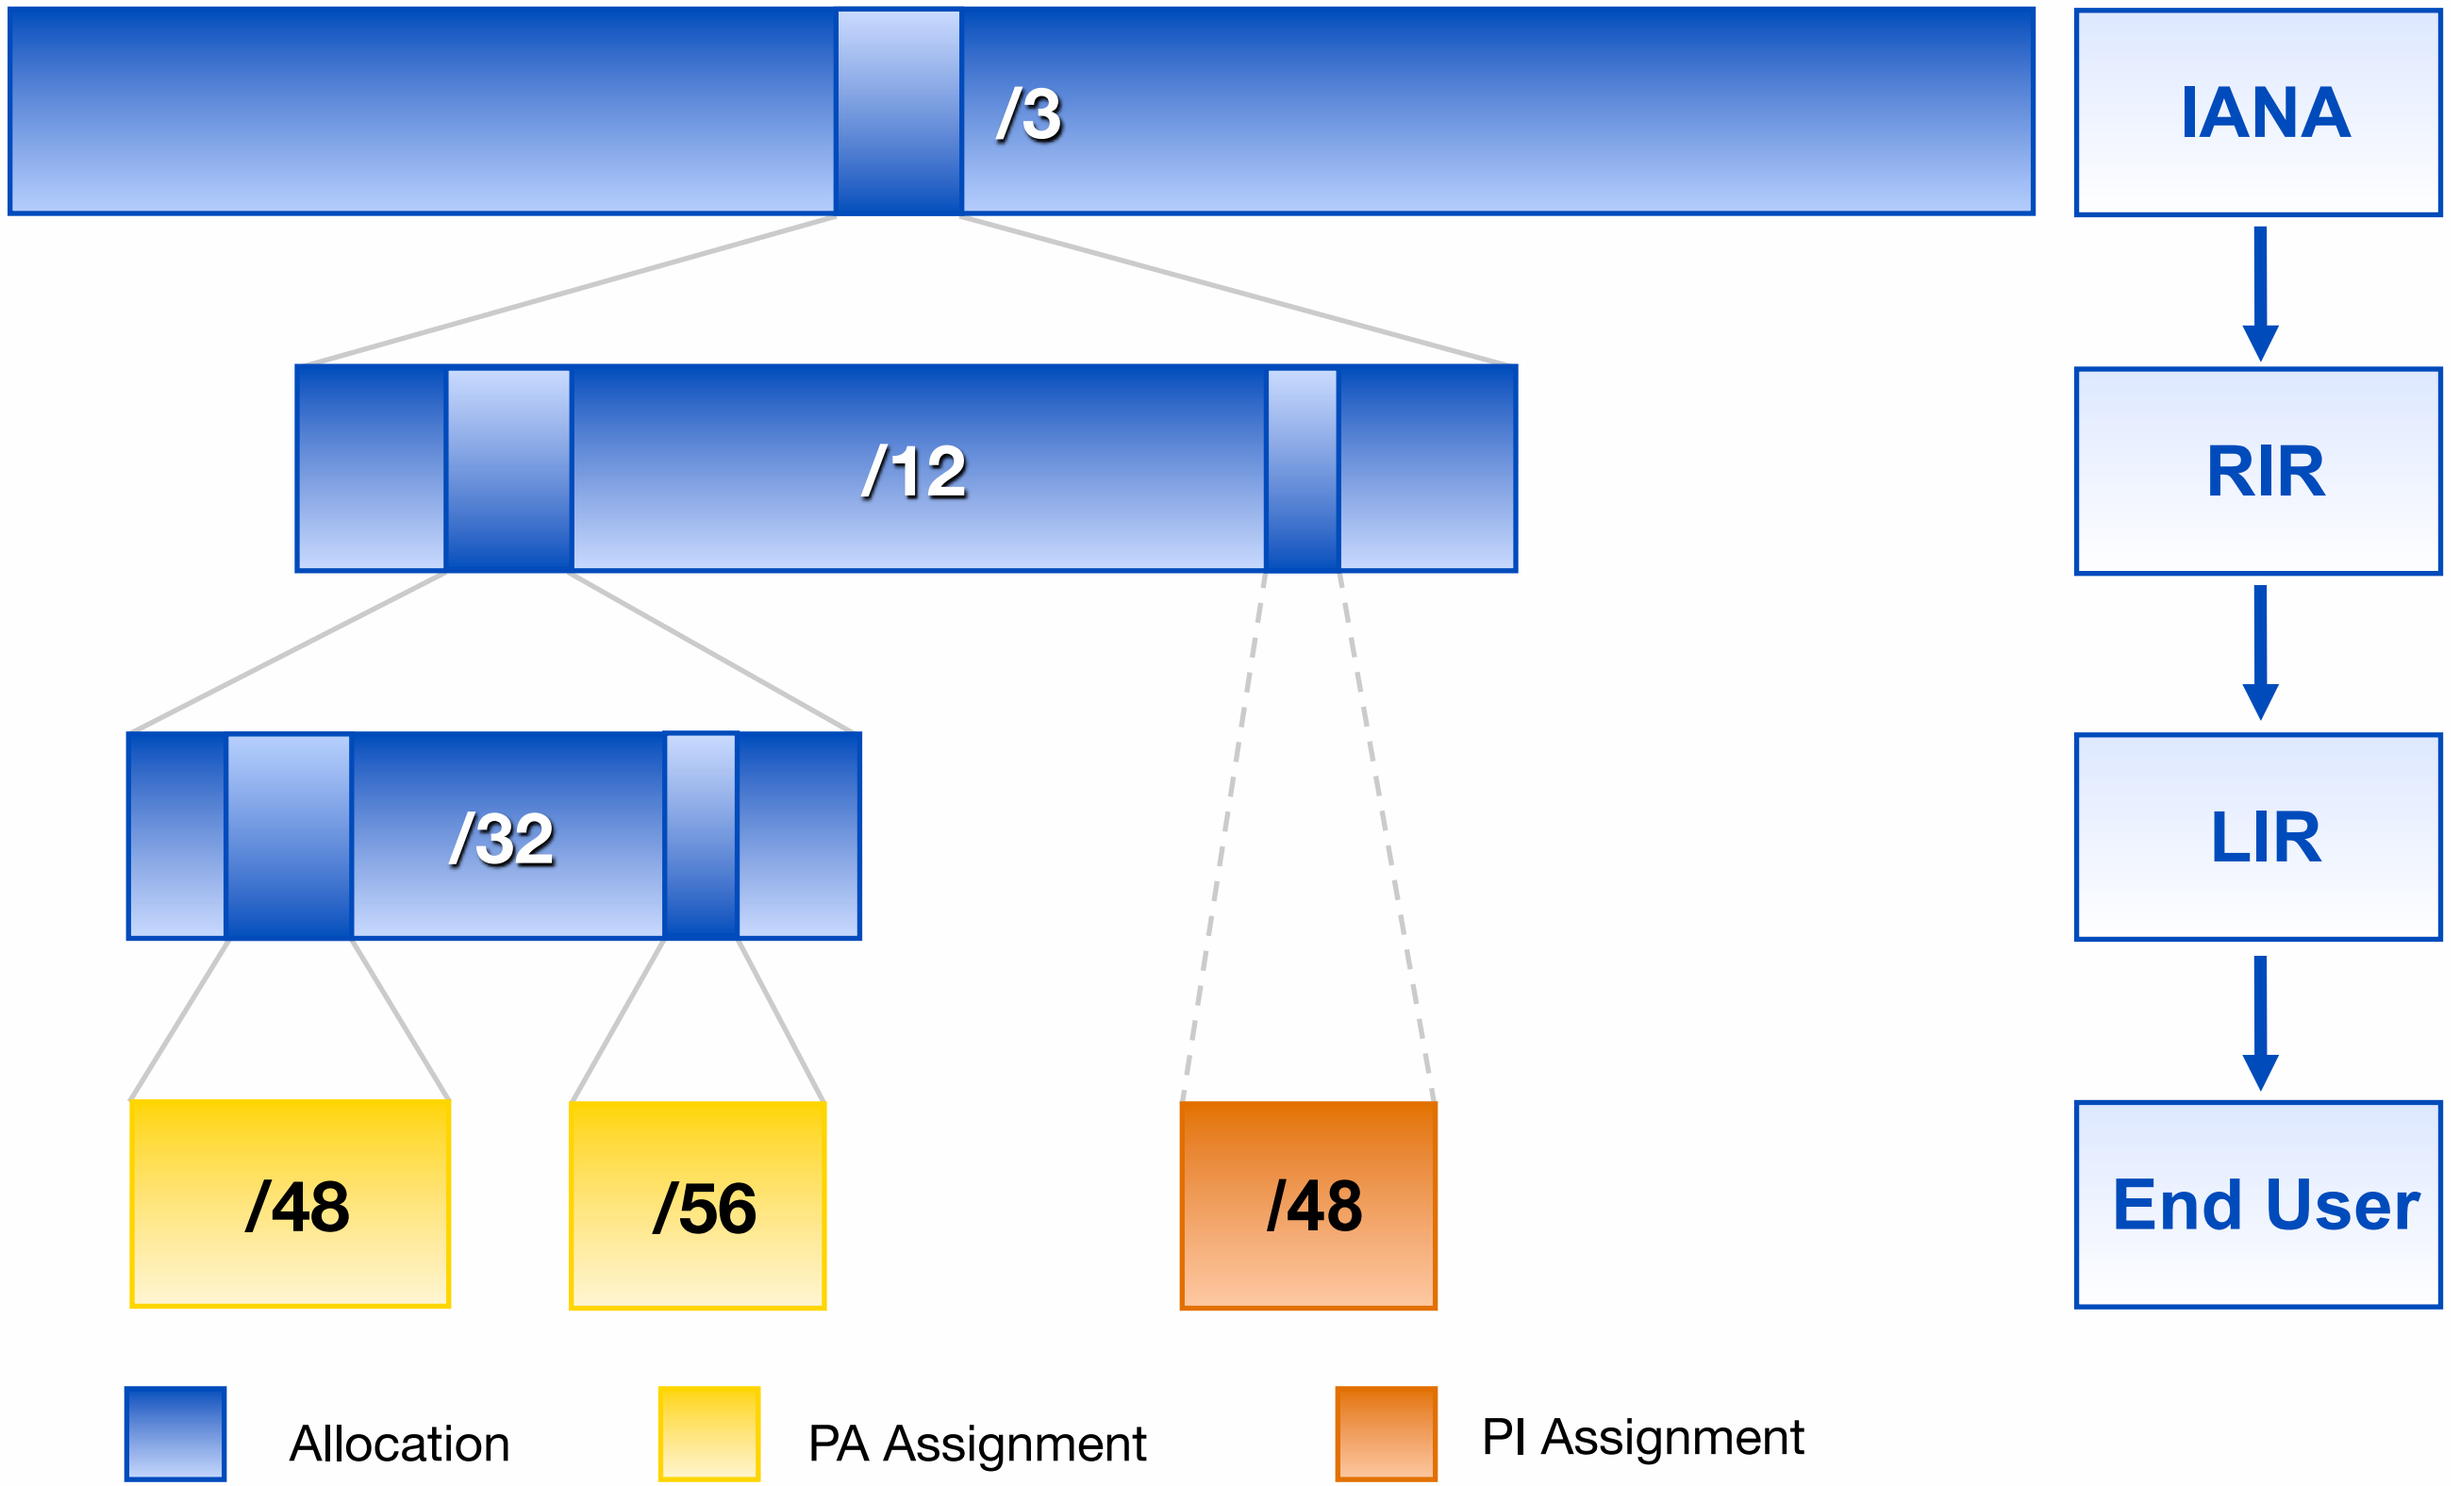

RIPE NCC

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

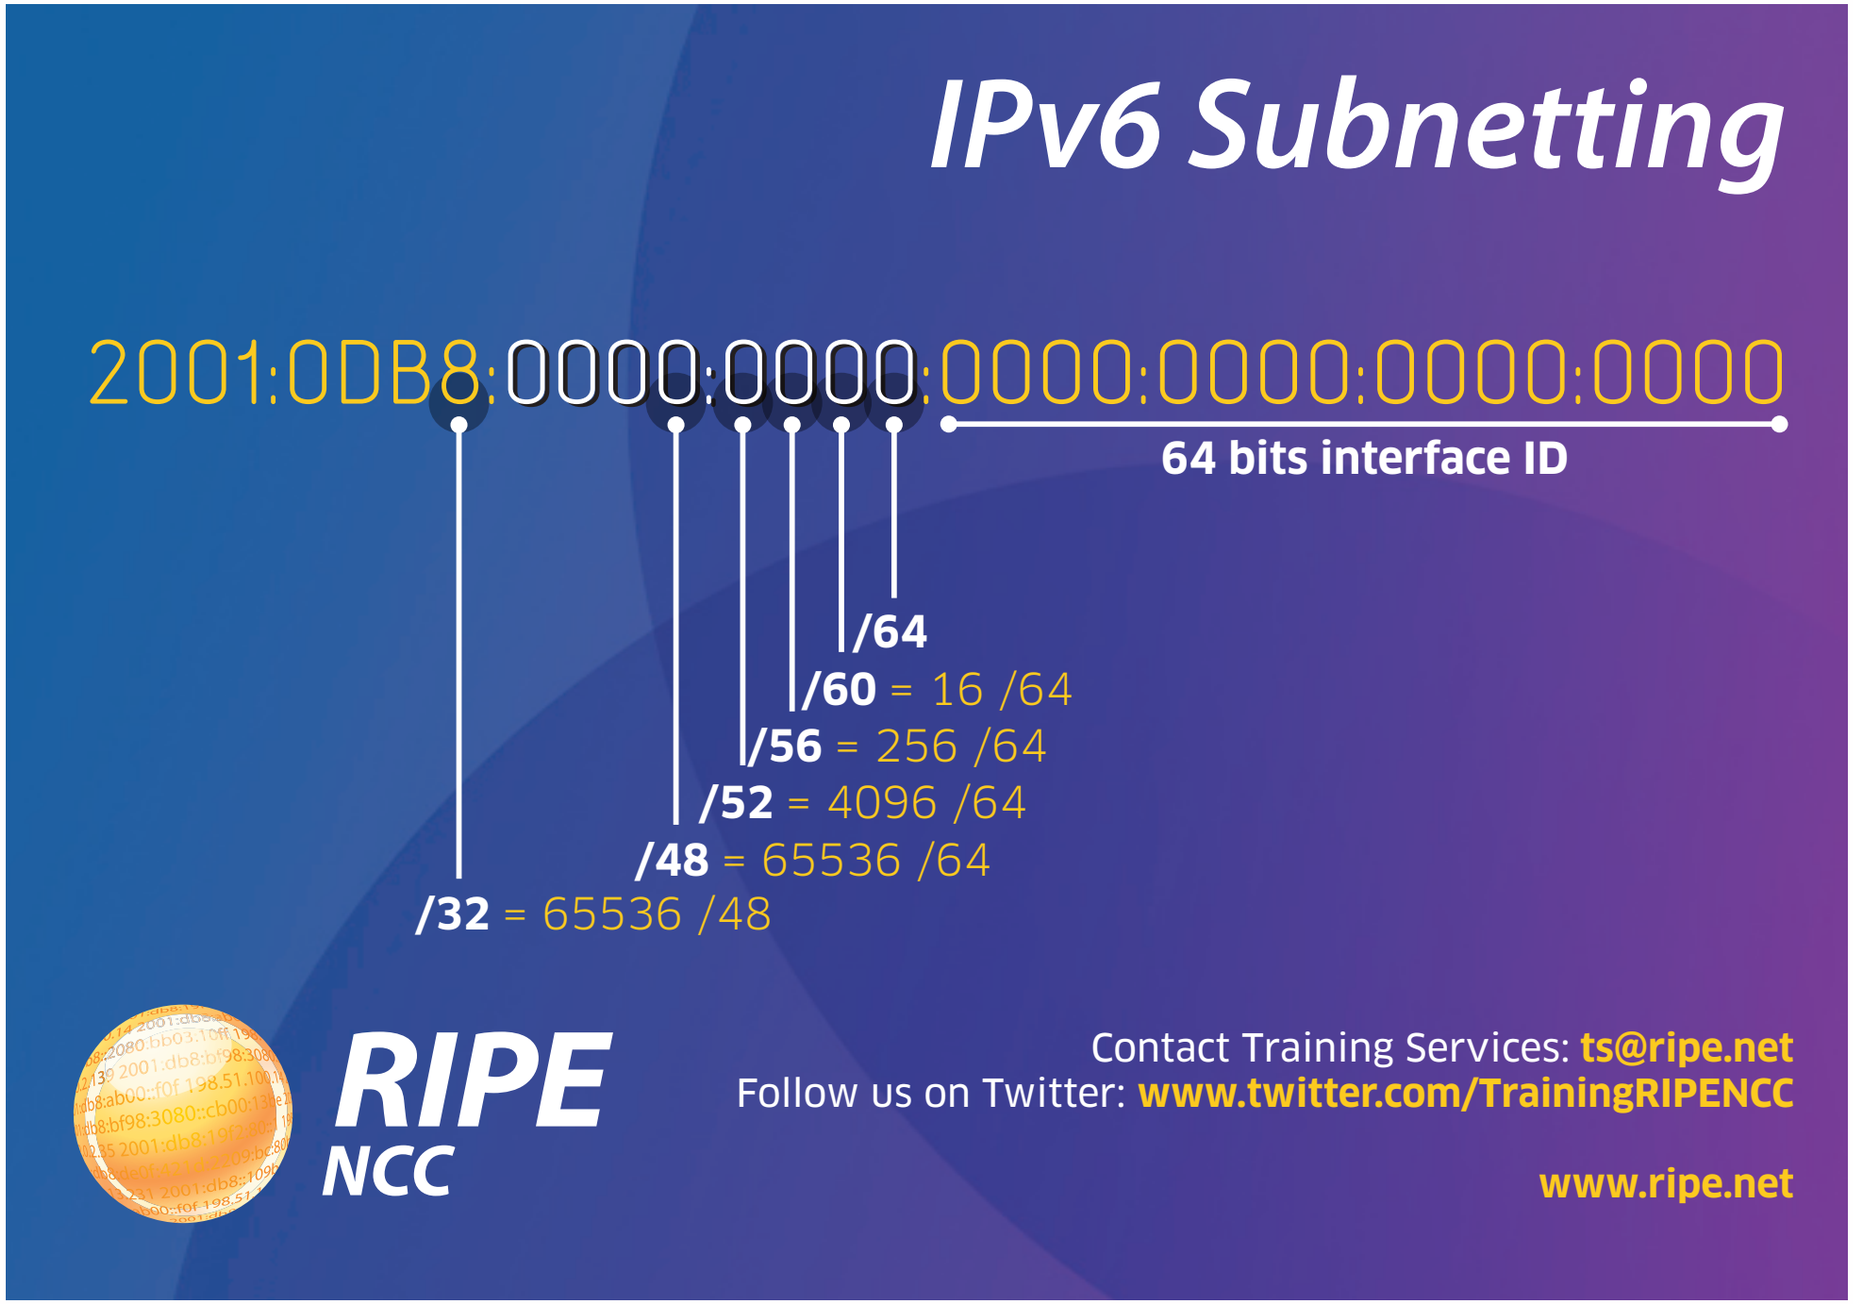
IPv6 Subnetting
2001:0DB8:0000:0000:0000:0000:0000:0000
64 bits interface ID
/64
/60 = 16 /64
/56 = 256 /64
/52 = 4096 /64
/48 = 65536 /64
/32 = 65536 /48
RIPE NCC
Contact Training Services: ts@ripe.net
Follow us on Twitter: www.twitter.com/TrainingRIPENCC
www.ripe.net

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

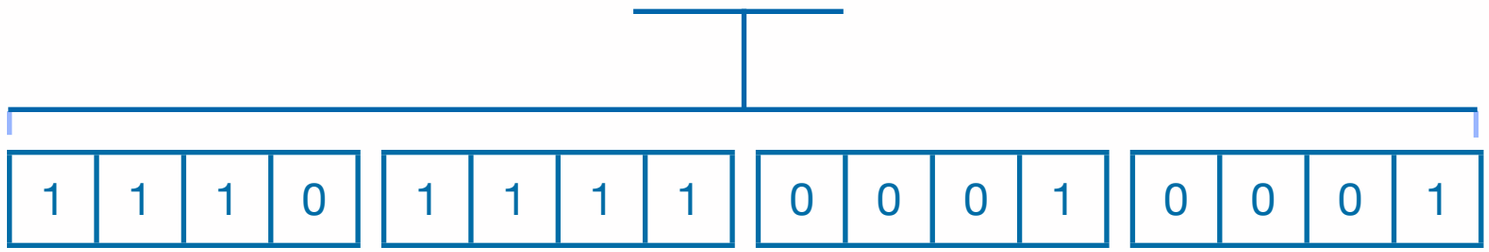

Multiple addresses

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

IPv6 Address Notation

Exercise

RIPE NCC

Getting it

Section 3

RIPE NCC

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
  - Up to a /29 without additional justification
  - More if justified by customer numbers

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

Comparison IPv4 and IPv6 status

IPv4	IPv6
ALLOCATED PA	ALLOCATED-BY-RIR
ASSIGNED PA	ASSIGNED
ASSIGNED PA	AGGREGATED-BY-LIR
SUB-ALLOCATED PA	ALLOCATED-BY-LIR
ASSIGNED PI	ASSIGNED PI

Using ASSIGNED

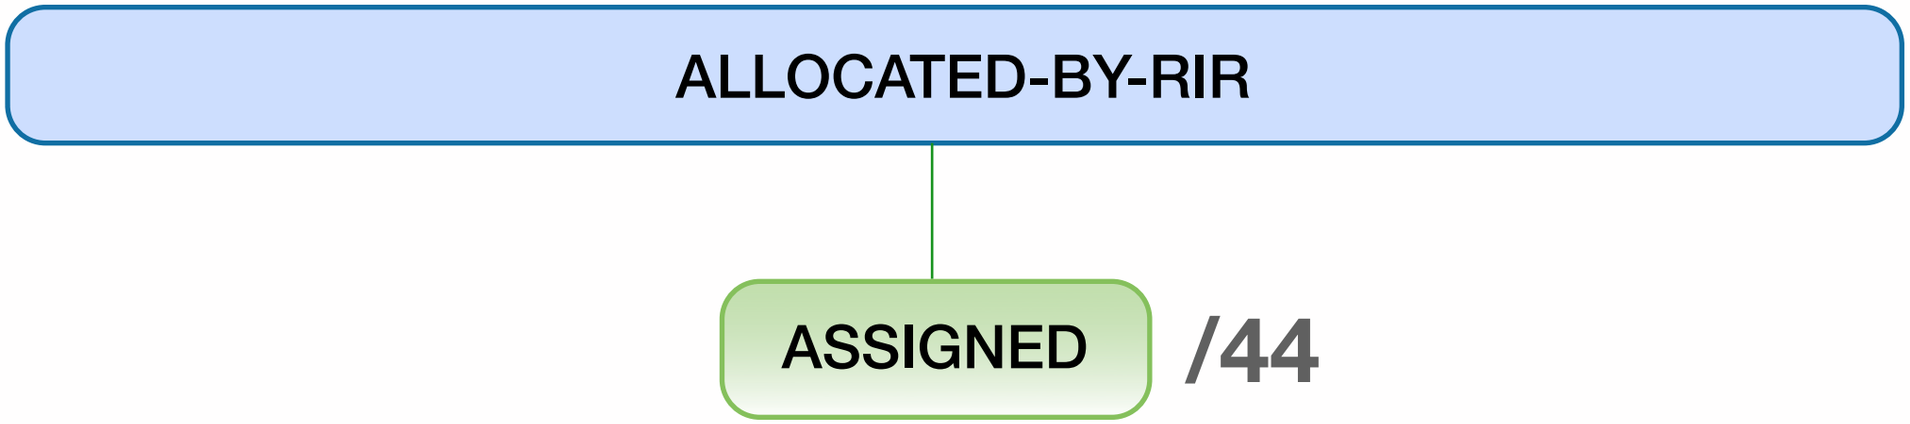

- Status is ASSIGNED
- Minimum assignment size is a/64
- For more than a /48, send a request form

Using AGGREGATED-BY-LIR

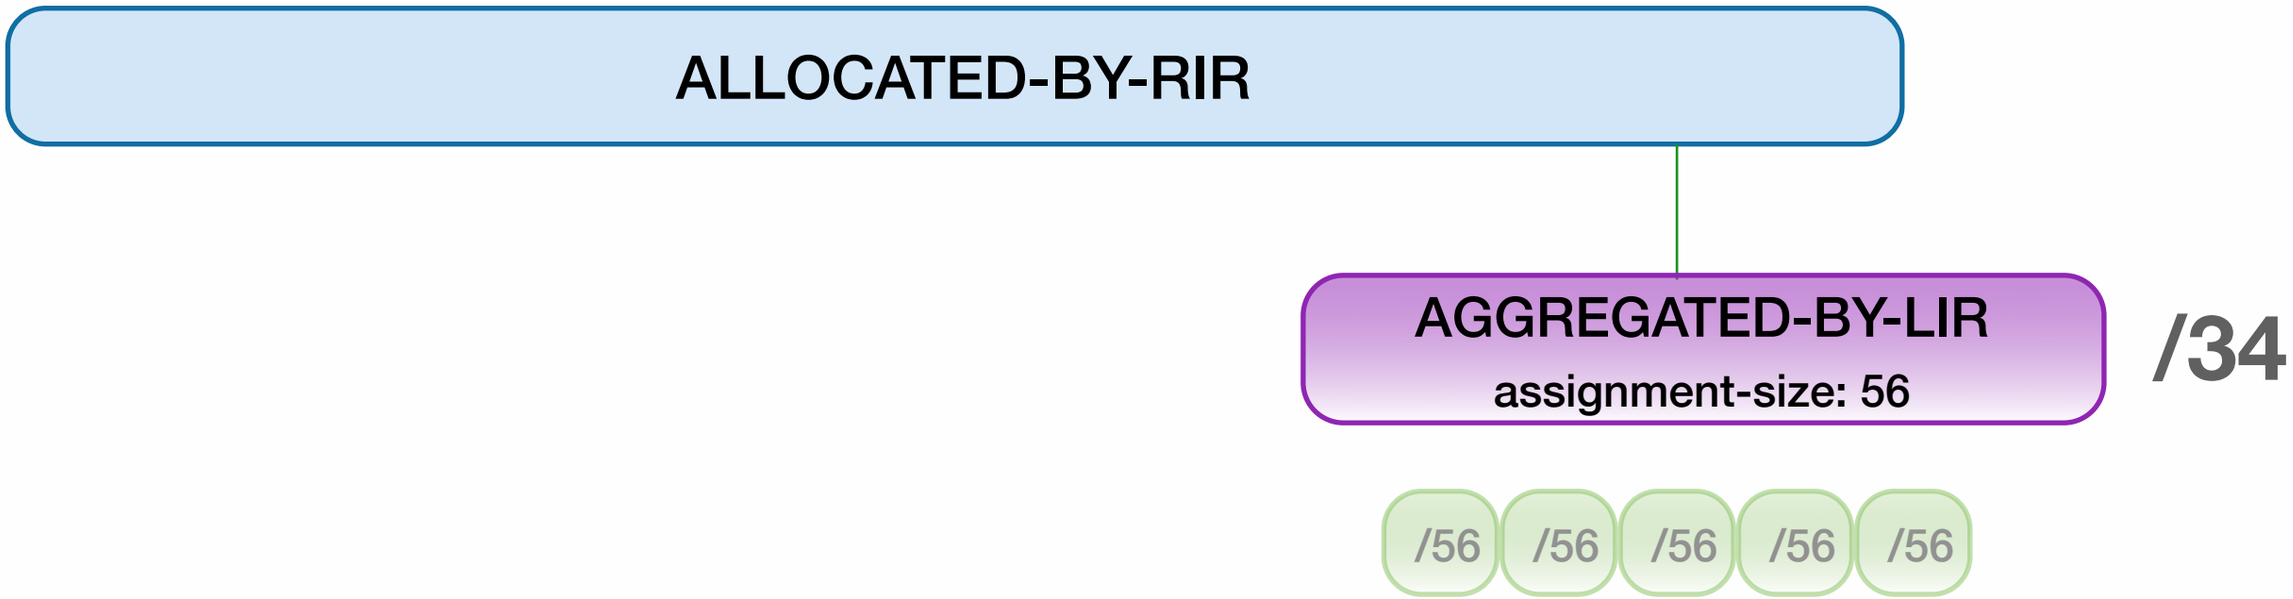

- Can be used to group customers
  - broadband, for example
- “assignment size” = assignment of each customer

AGGREGATED-BY-LIR in the RIPE DB


```
inet6num:        2001:db8:1000::/36
netname:         Brightlife
descr:           Broadband services
country:         NL
admin-c:         BN649-RIPE
tech-c:          BN649-RIPE
status:          AGGREGATED-BY-LIR
assignment-size: 48
mnt-by:          BRIGHTLIFE-MNT
notify:          noc@example.net
changed:         noc@example.net 20130218
source:          RIPE
```

Using ALLOCATED-BY-LIR

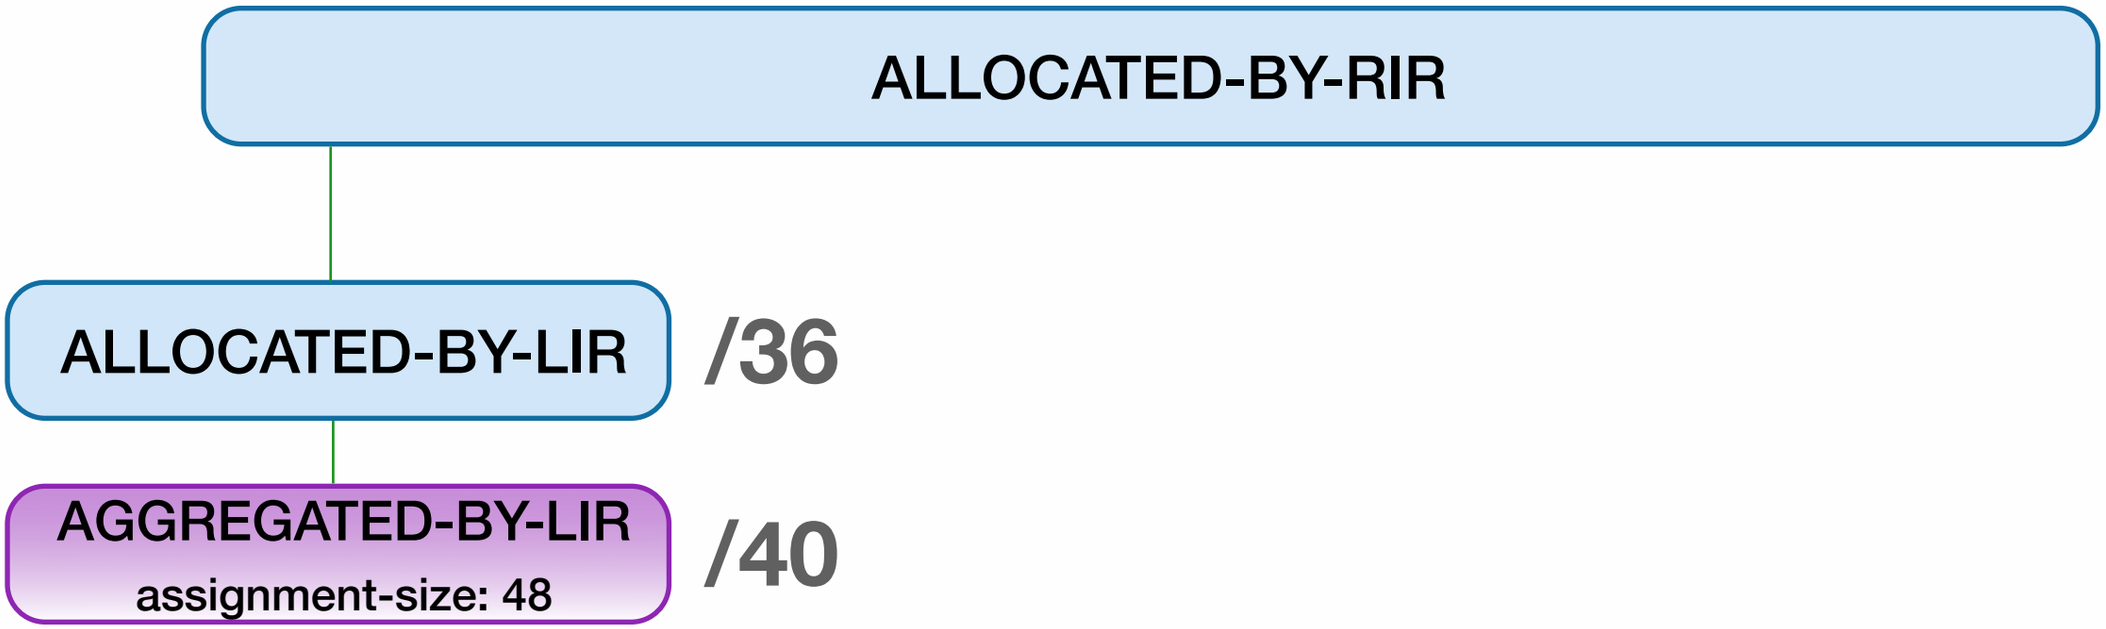

- Can be used for customers who expect large growth
  - or for your own infrastructure

Overview

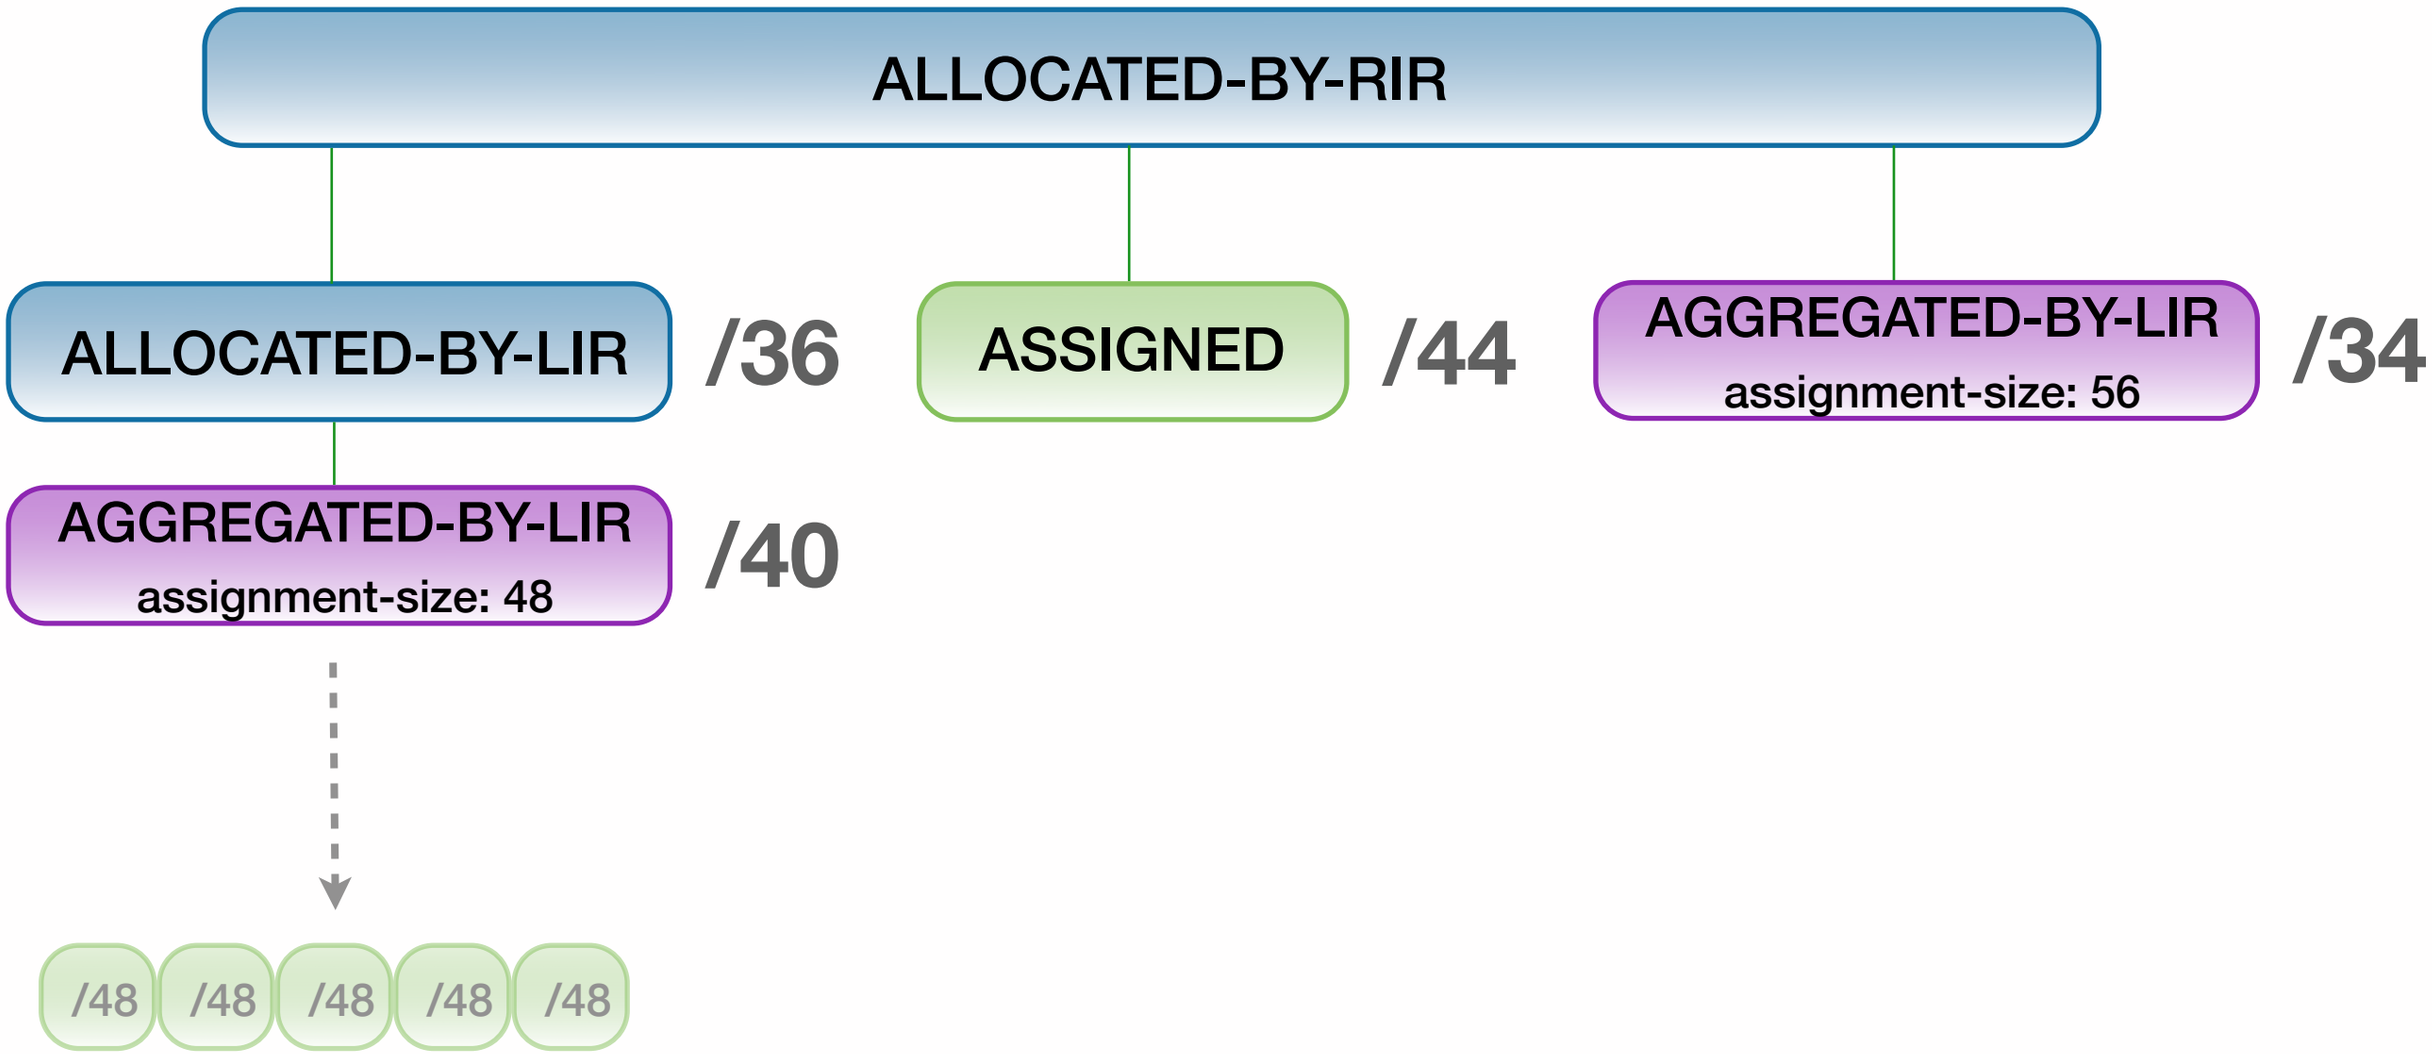
ALLOCATED-BY-RIR
ALLOCATED-BY-LIR
/36
AGGREGATED-BY-LIR
assignment-size: 48
/40
ASSIGNED
/44
AGGREGATED-BY-LIR
assignment-size: 56
/34
/48
/48
/48
/48
/48

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
  - not even 1 IP address

?

!

Making Assignments

Exercise

RIPE NCC

Fridge6!

- 20 minutes preparation time
- 10 minutes discussion

Solution RIPE Database object

inet6num:	2001:db8:1000::/36
netname:	FREEZ
descr:	Freez Fridges
country:	NL
admin-c:	RM1204-RIPE
tech-c:	RM1204-RIPE
status:	**AGGREGATED-BY-LIR**
assignment-size:	**56**
mnt-by:	LIR-MNT
notify:	noc@lir-example.com
changed:	noc@lir-example.com 20110801
source:	RIPE

Transition Mechanisms

Section 4

RIPE NCC

Transitioning: 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

- 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

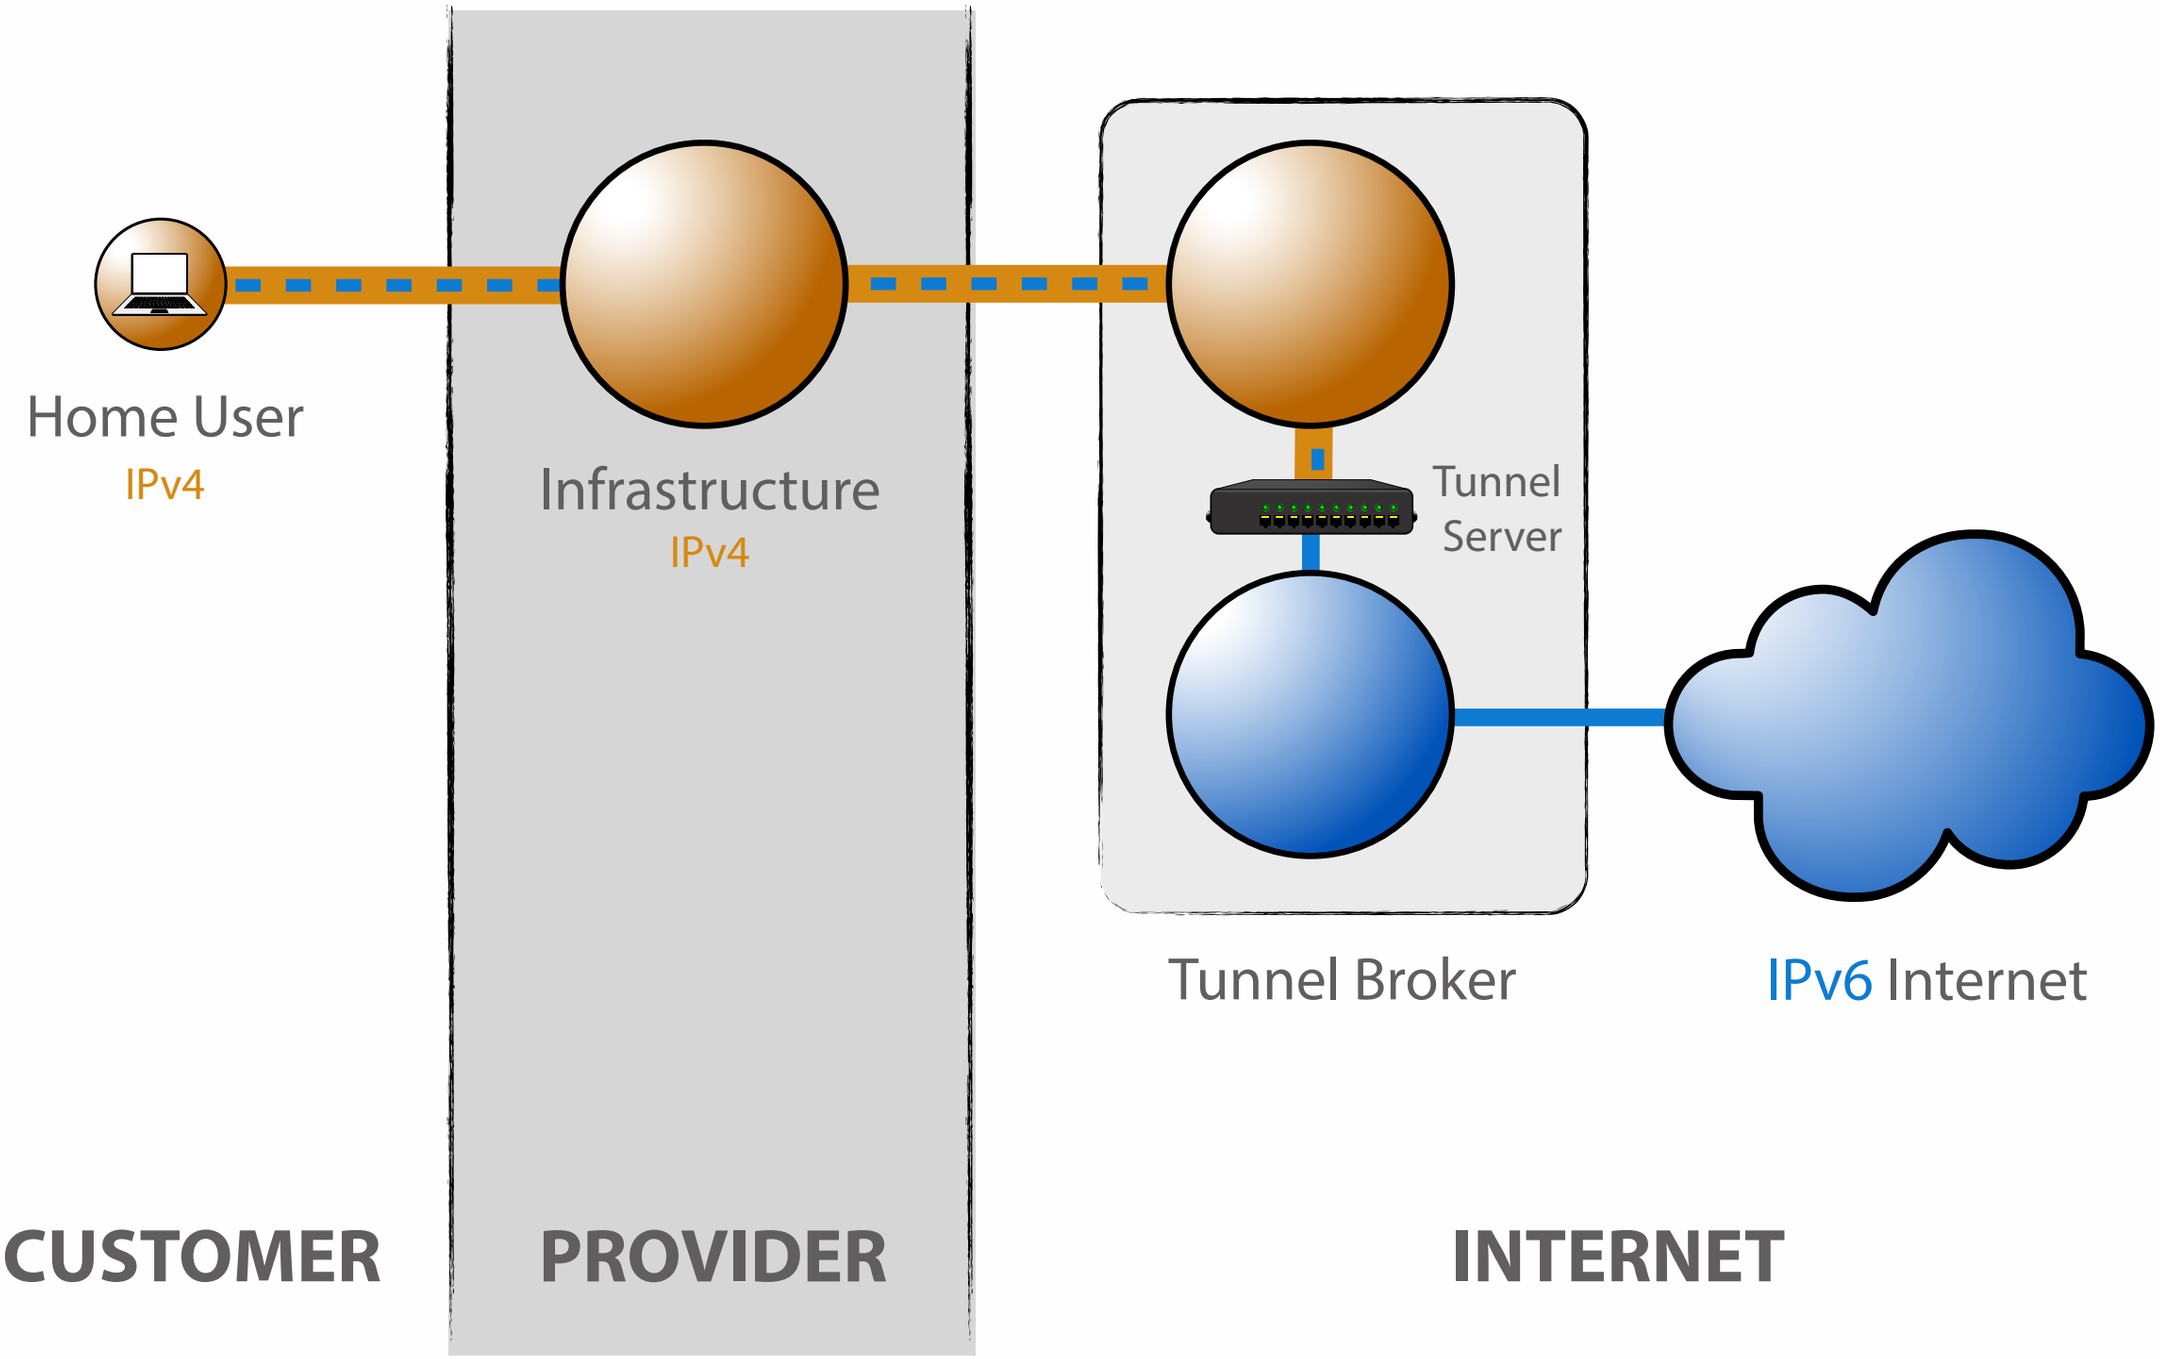
Home User
IPv4
Infrastructure
IPv4
Tunnel
Server
Tunnel Broker
IPv6 Internet
CUSTOMER
PROVIDER
INTERNET

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

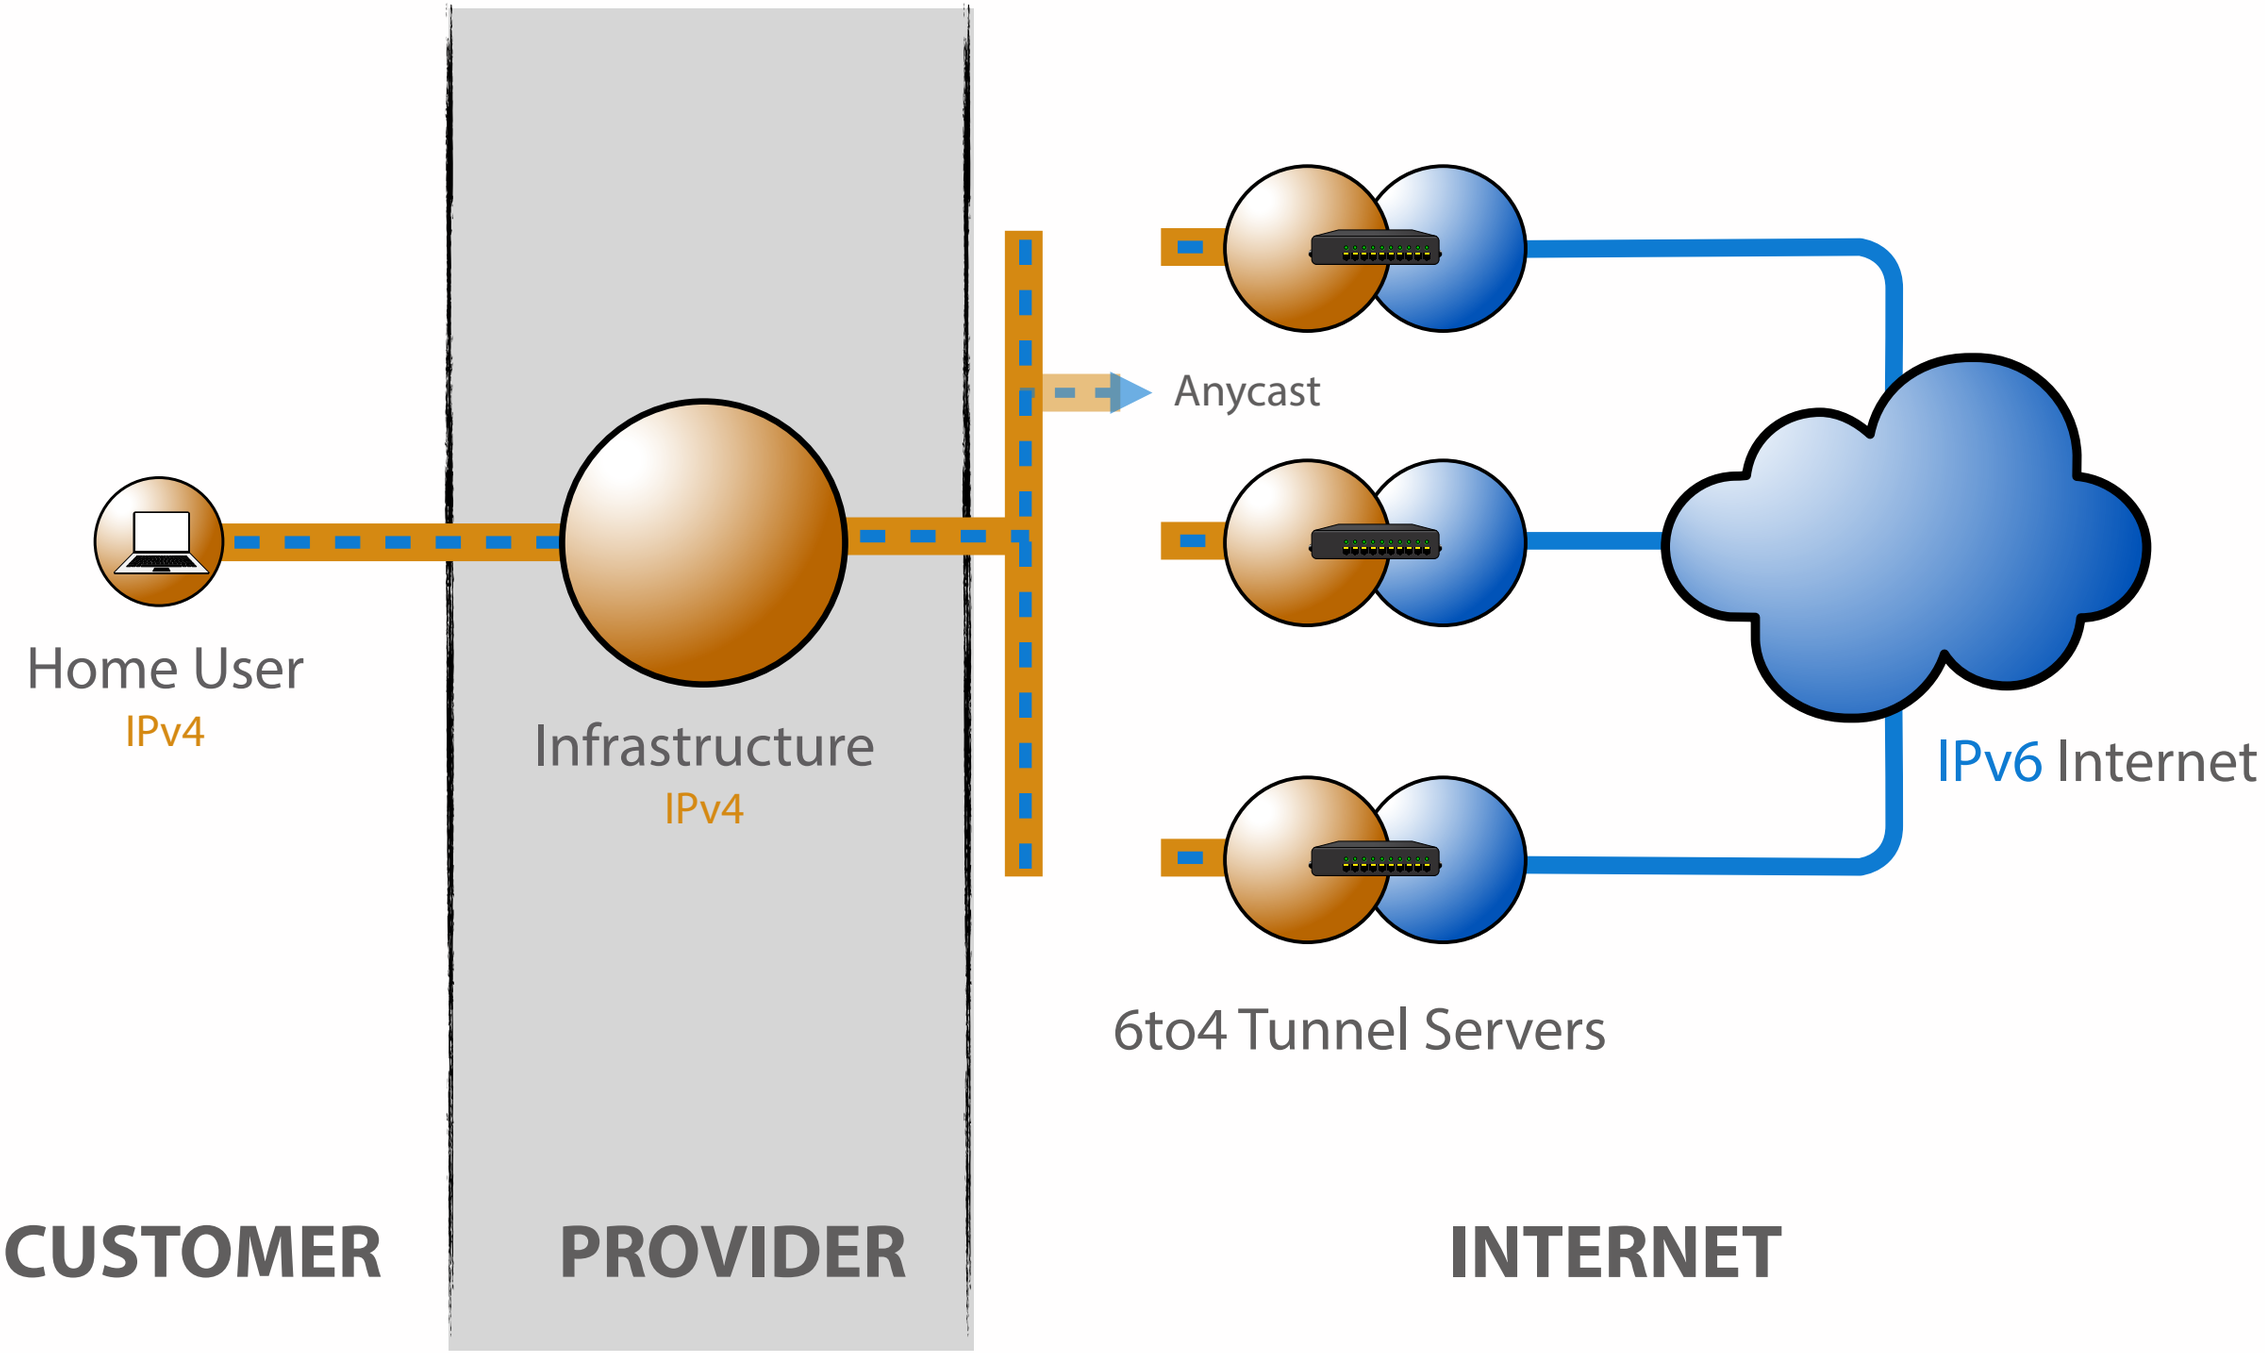

6RD

- **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**

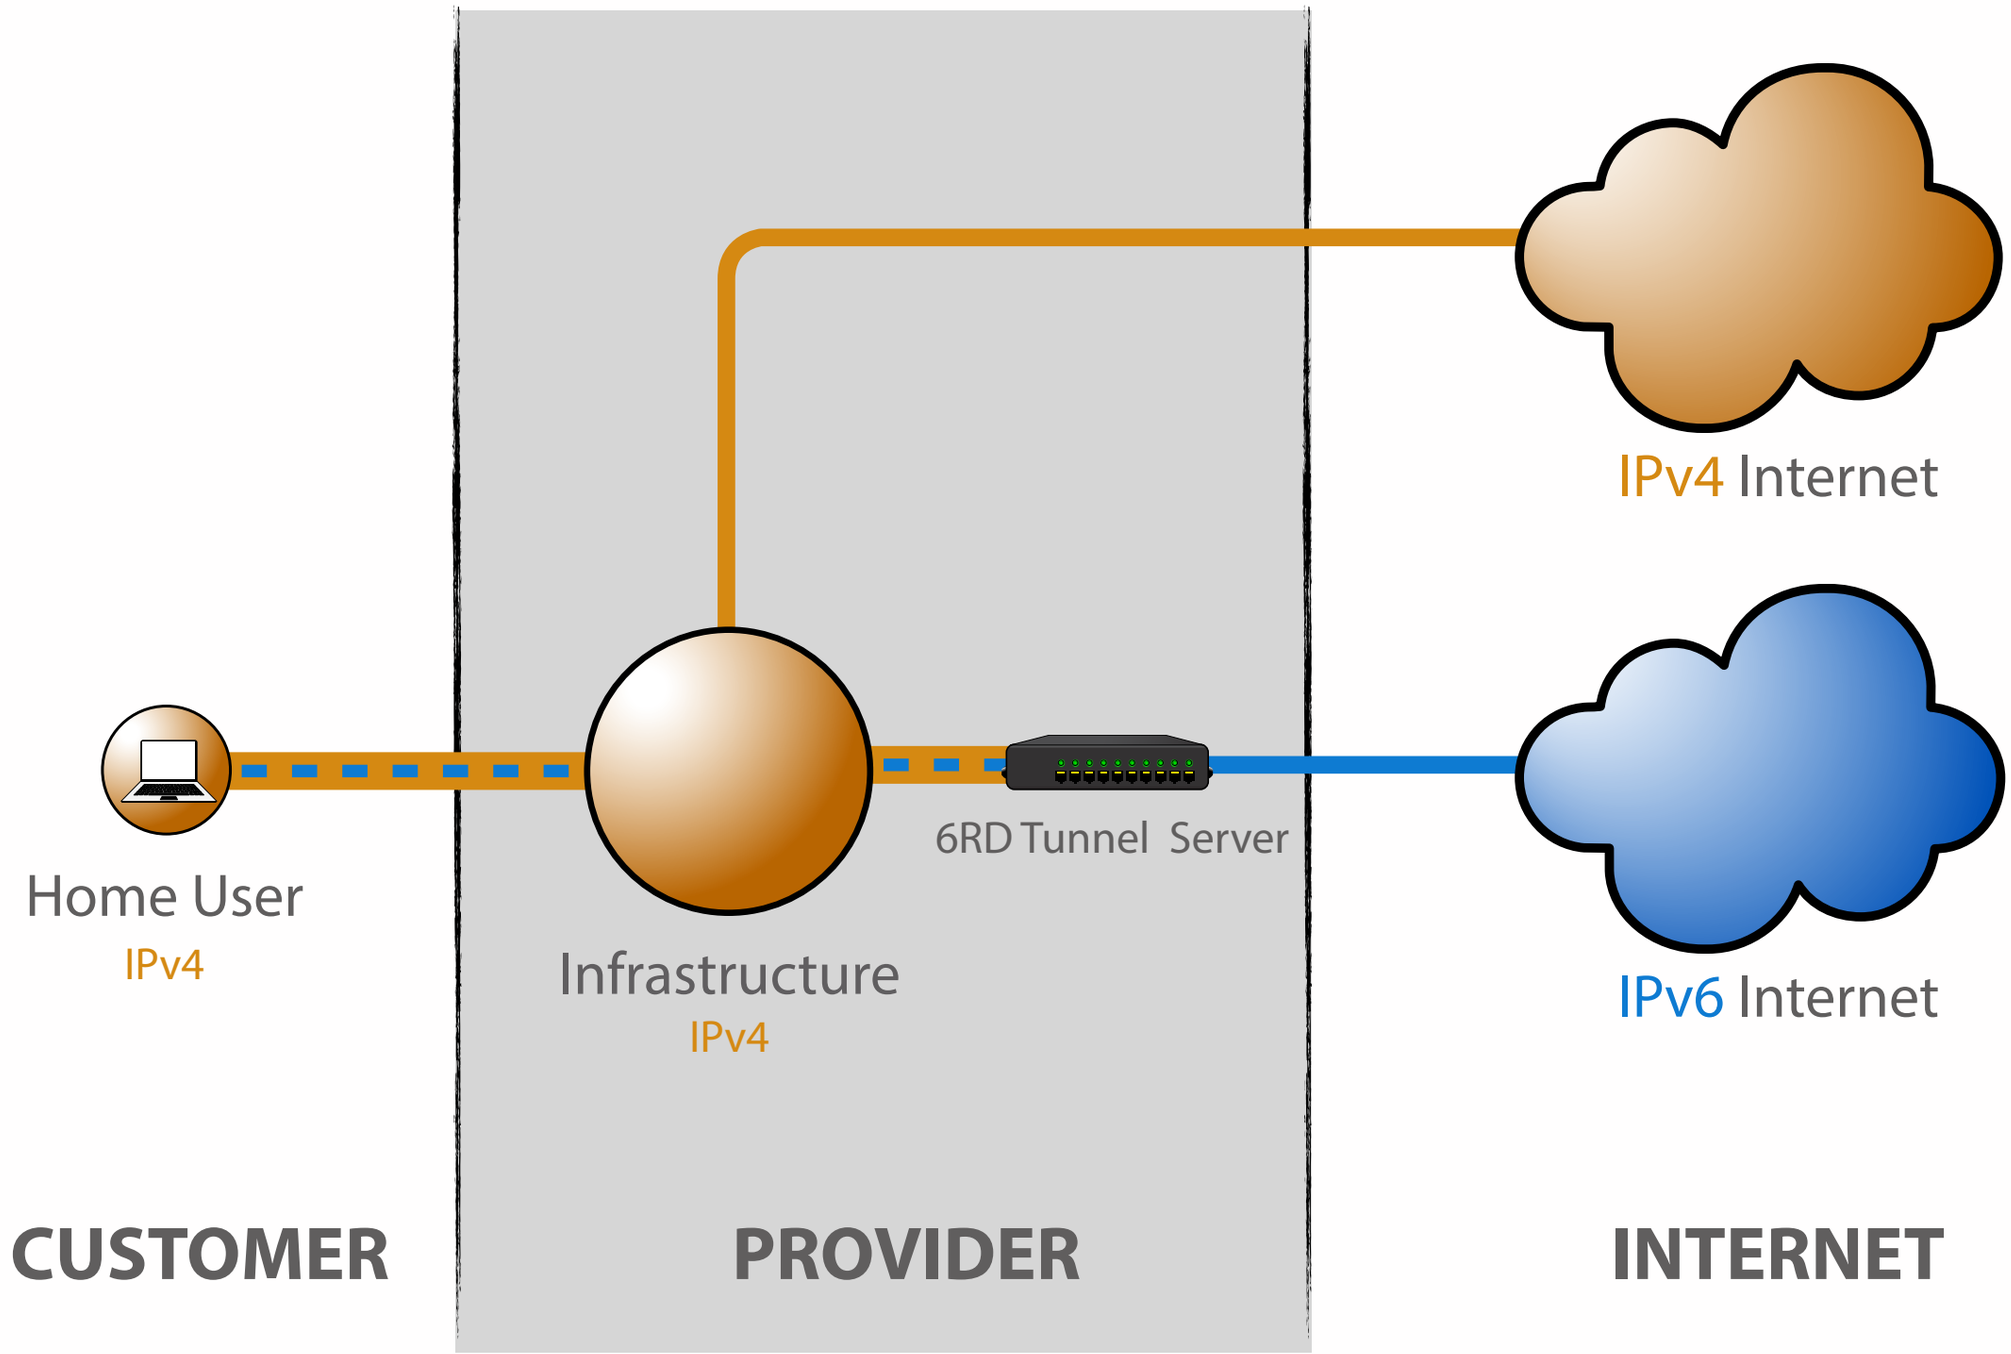
IPv4 Internet
Home User
IPv4
Infrastructure
IPv4
6RD Tunnel Server
IPv6 Internet
CUSTOMER
PROVIDER
INTERNET

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

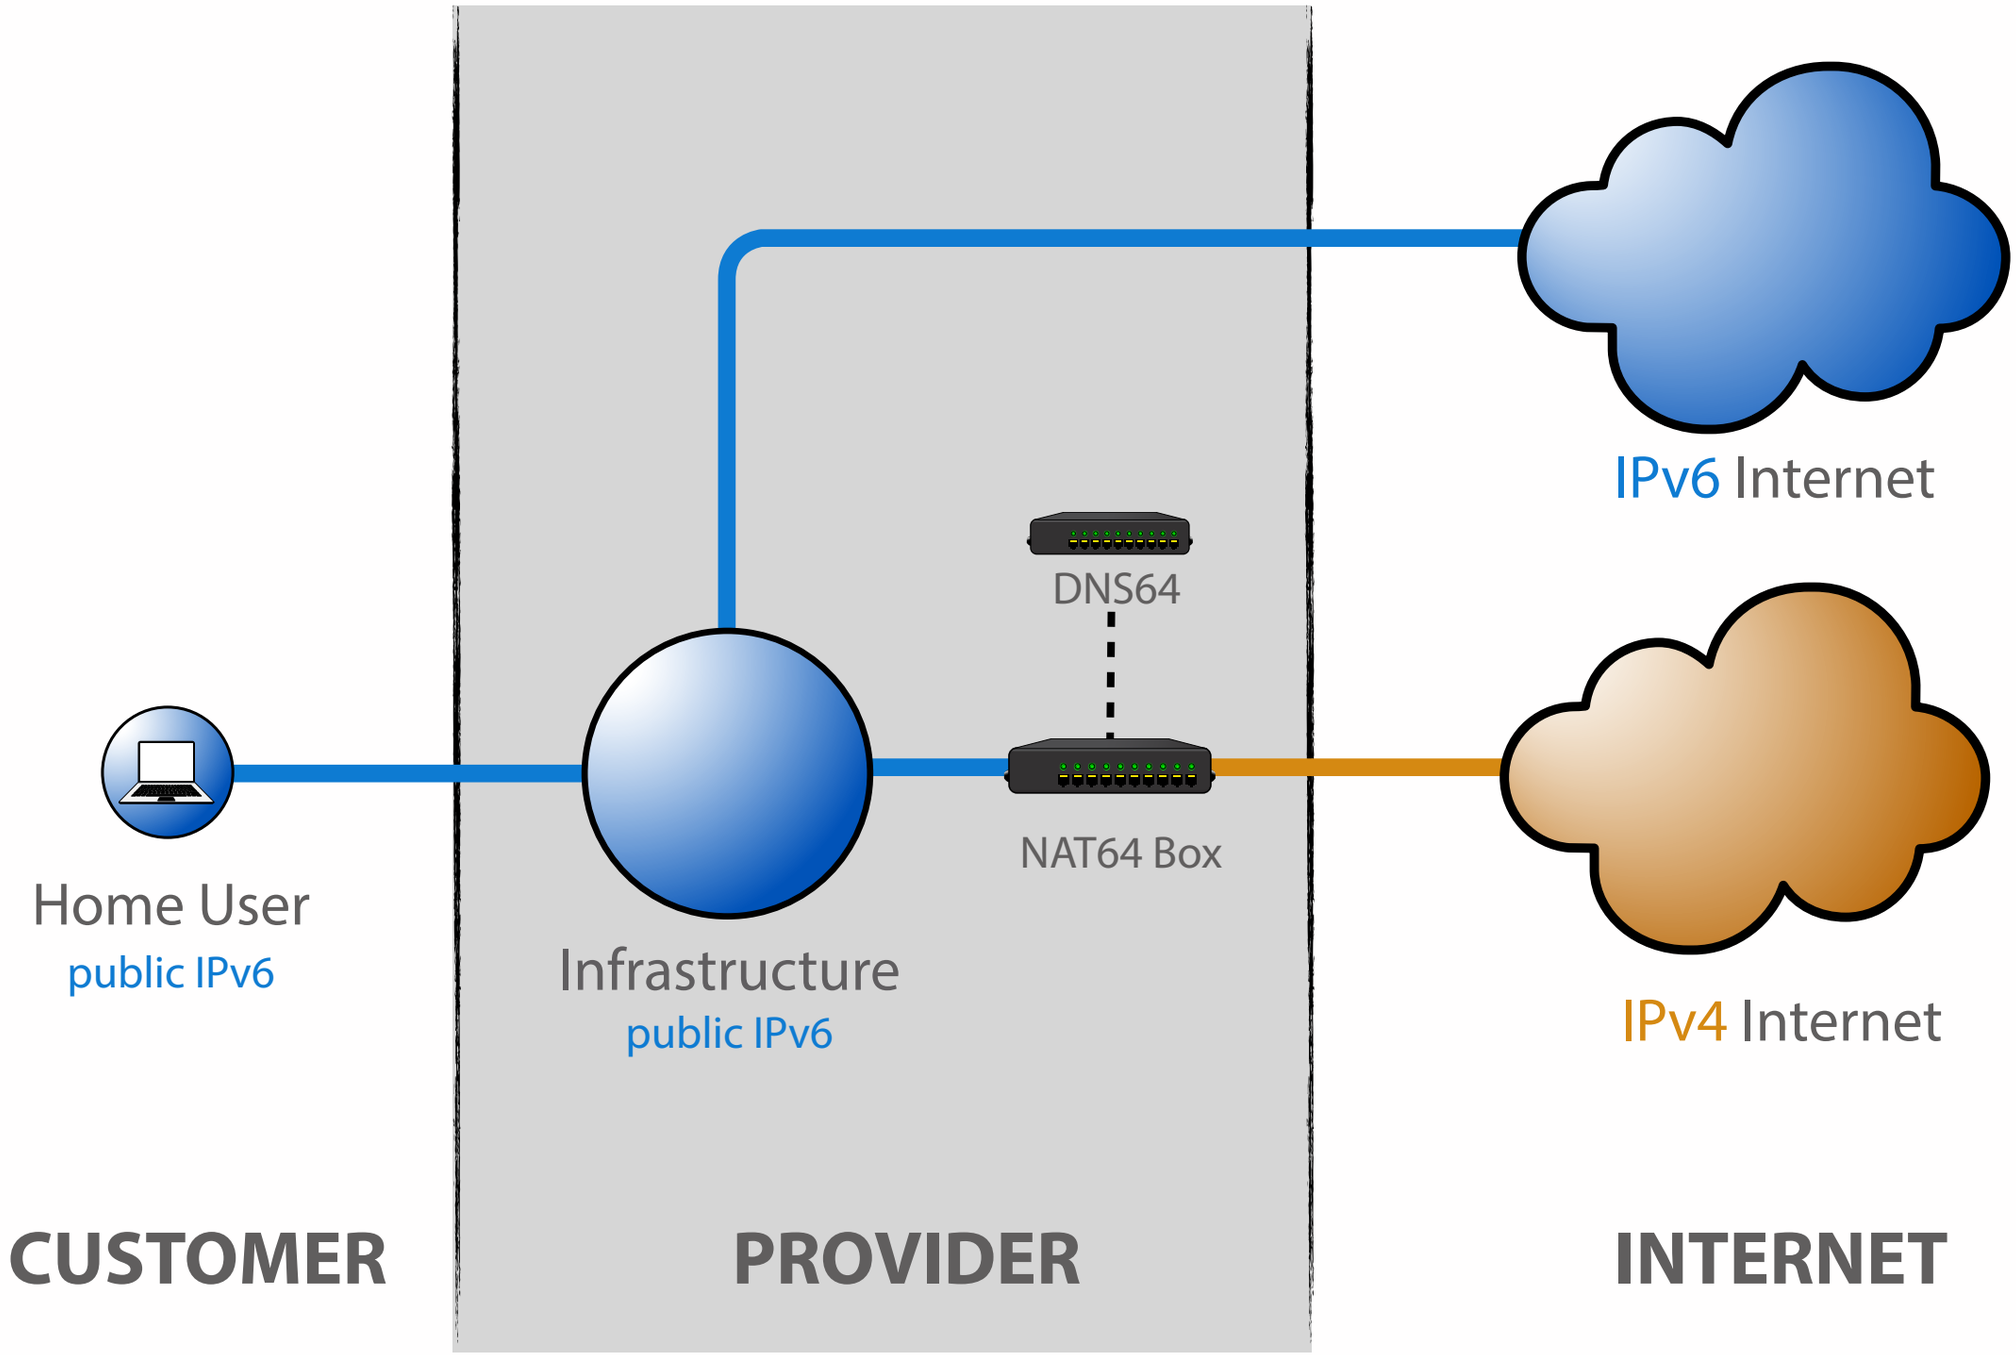
IPv6 Internet
DNS64
NAT64 Box
IPv4 Internet
Home User
public IPv6
Infrastructure
public IPv6
CUSTOMER
PROVIDER
INTERNET

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

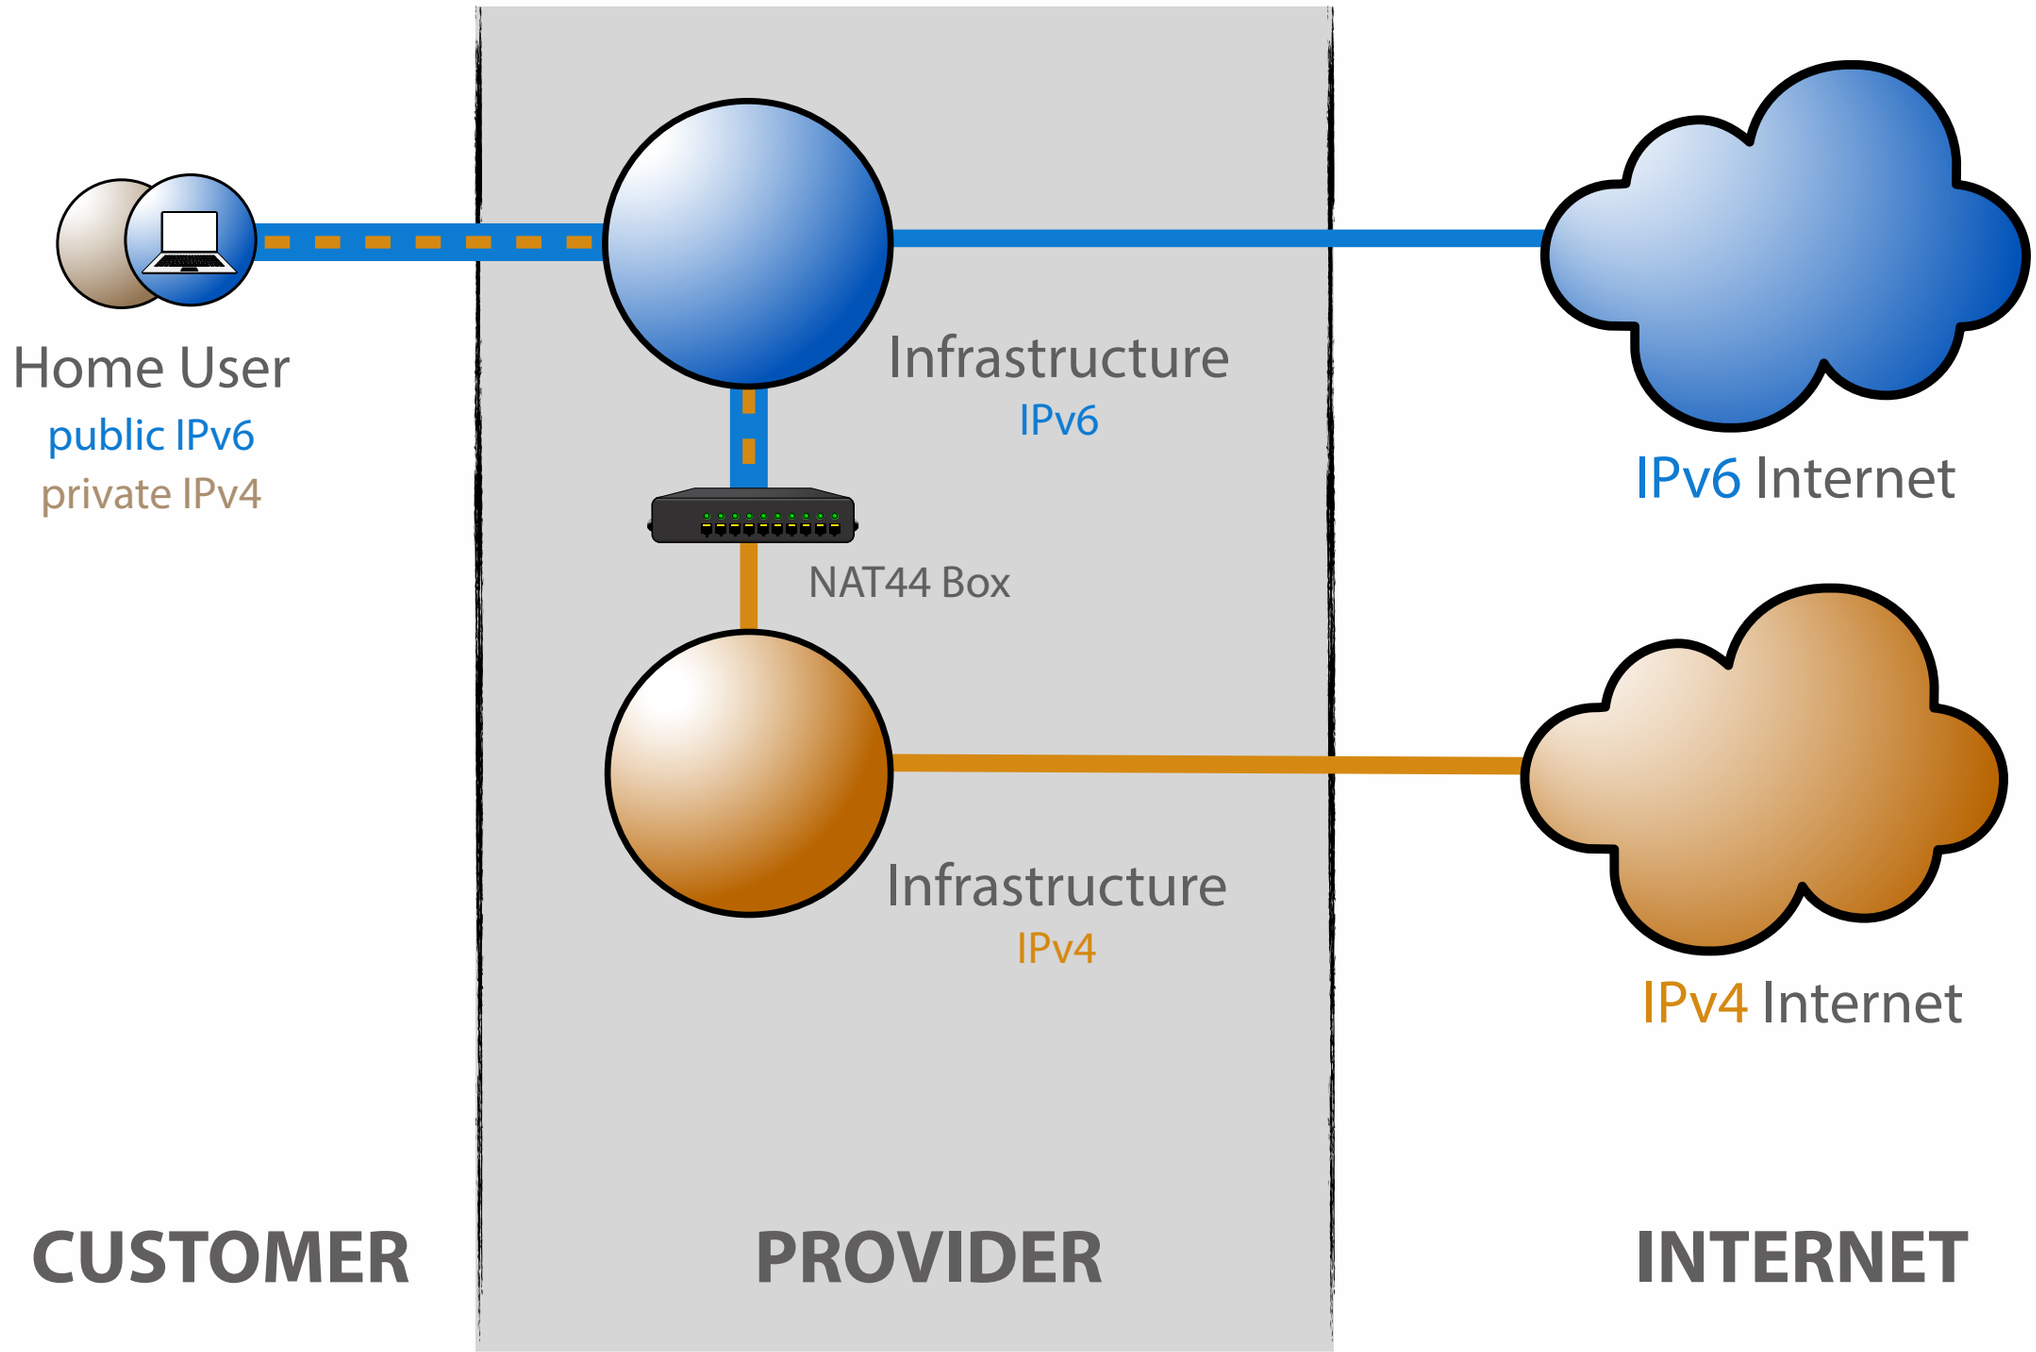
Home User
public IPv6
private IPv4
Infrastructure
IPv6
NAT44 Box
Infrastructure
IPv4
IPv6 Internet
IPv4 Internet
CUSTOMER
PROVIDER
INTERNET

Dual Stack

where you can

Addressing Plans

Section 5

RIPE NCC

Why Create an IPv6 Addressing Plan?

- Mental health during implementation(!)
- Easier implementation of security policies
- Efficient addressing plans are scalable
- More efficient route aggregation

IPv6 Address Management

- **Your spreadsheet 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**

?

Addressing Plan

Exercise

RIPE NCC

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
- **20 minutes preparation time**
- **10 minutes discussion**

Addressing plans

- **Number of hosts in a /64 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**

More on Addressing Plans

- For private networks, consider ULA
- For servers you want a manual configuration
- Use port numbers for addresses
  - pop server 2001:db8:1::110
  - dns server 2001:db8:1::53
  - etc…

Deploying

Section 6

RIPE NCC

IPv6 Stateless Address Autoconfiguration

- **Host will automatically start looking for a router**
  - **Response will contain:**
    - Router's address
    - Zero or more link prefixes
    - SLAAC allowed yes/no
    - MTU size (optional)

SLAAC - Generating IPv6 address

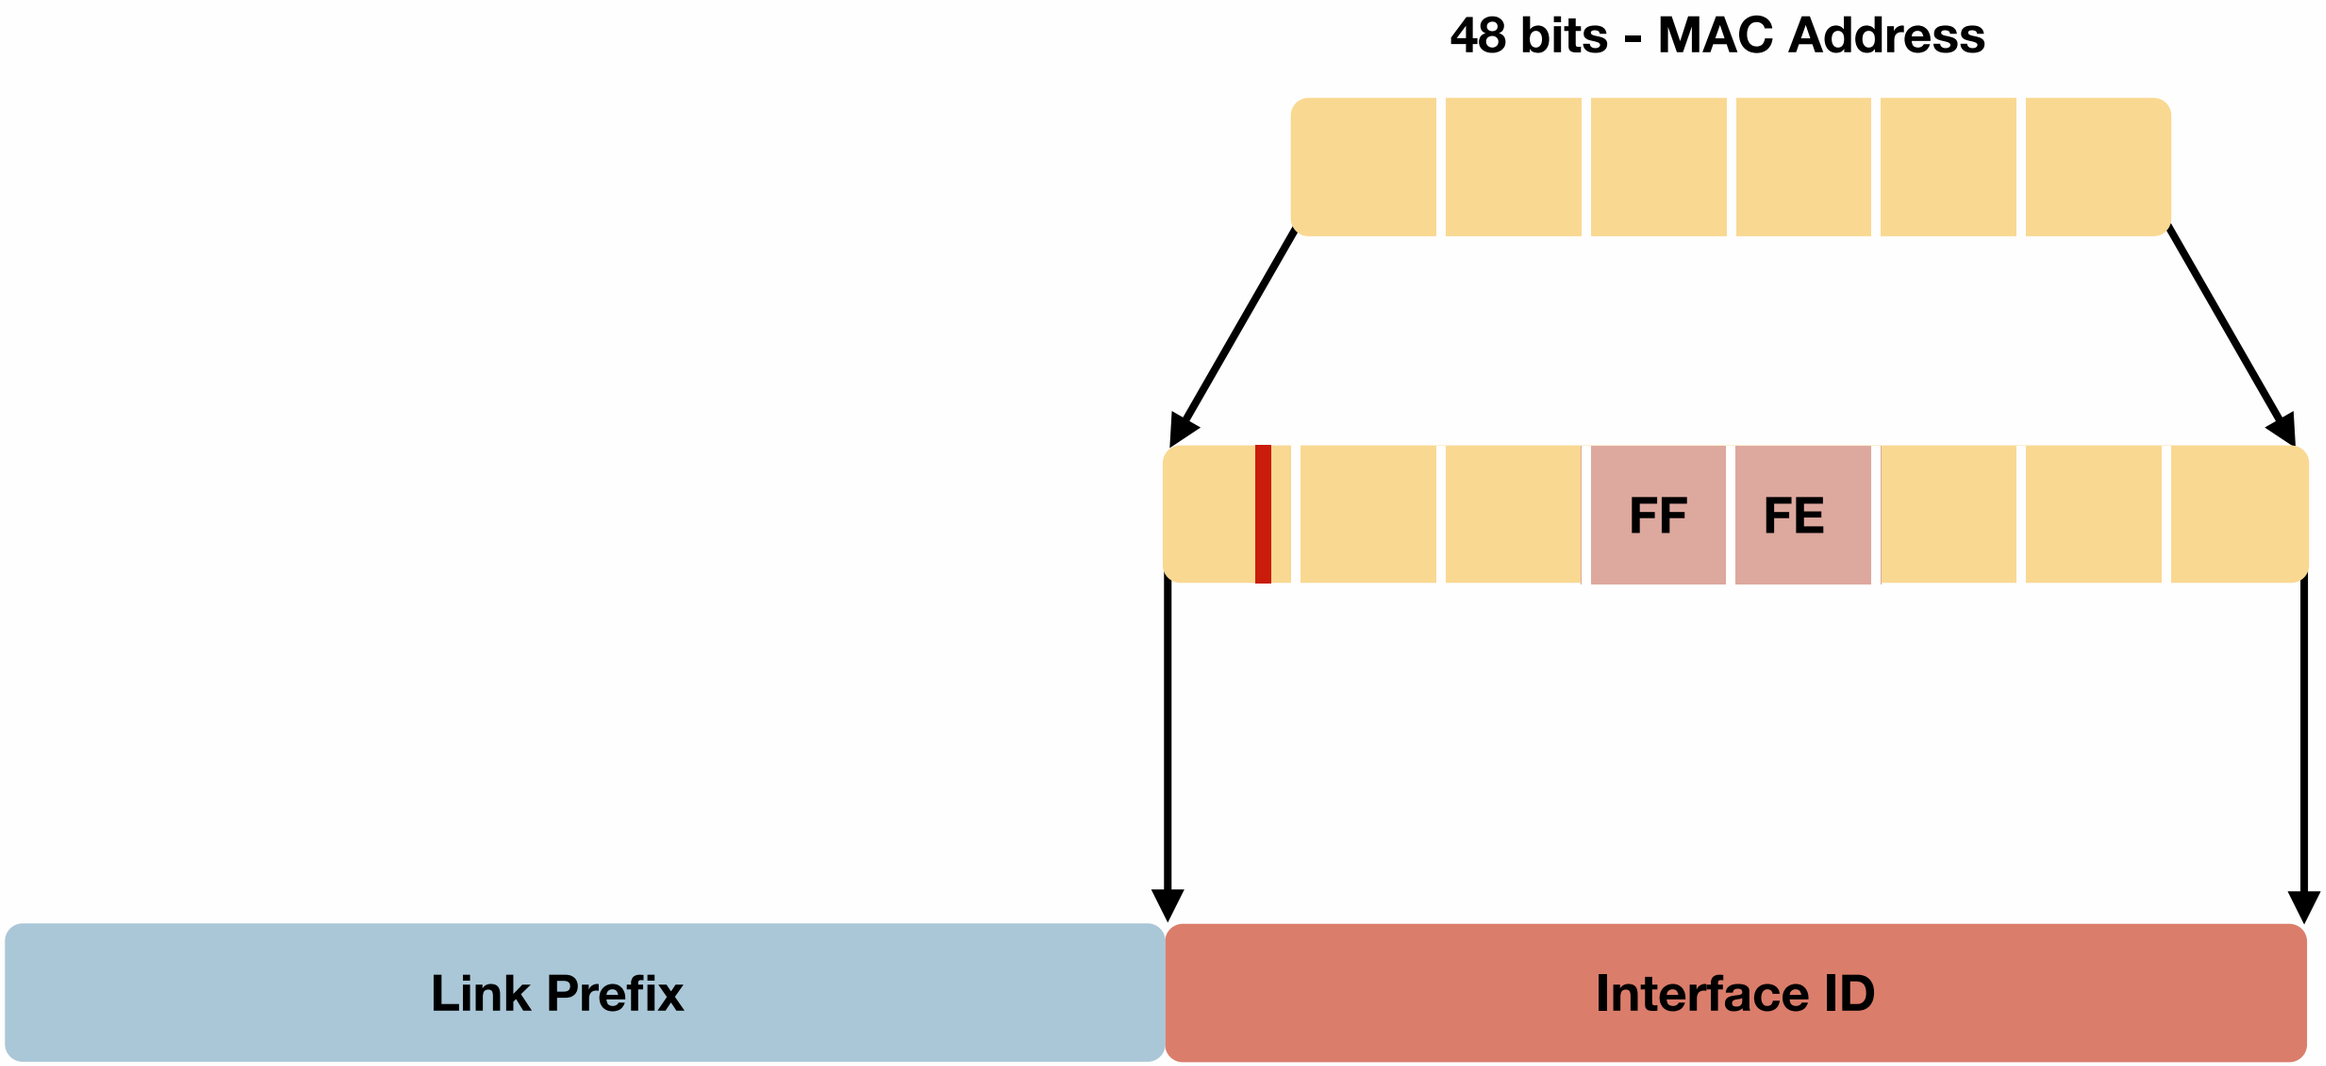

Privacy Extensions for SLAAC

- Provides privacy for users
- Changes the interface ID over time
- Interface ID must be locally unique
- Interface ID can be random
- Duplicate Address Detection ensures uniqueness
- In case of a collision a new random address is generated

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 config” 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:ca**ff:fe**f6:bdbf**/64
  - **Global Unicast** 2001::**5a55:ca**ff:fe**f6:bdbf**/64 (multiple)
- **Multicast**
  - **All Nodes** ff02::1 (scope: link)
  - **Solicited Node** ff02::1:ff**f6: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.0.0.0.0.0.0.0.0.1.1.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.0.0.0.0.0.0.0.0.1.1.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

- **Everybody can claim to be a router**
  - **Use RA Guard to filter unauthorised RAs**
    - RFC 6105
  - **Secure Neighbour Discovery (SEND)**
    - RFC 3971
    - Neighbour Solicitation/Advertisement spoofing
    - DoS Attack
    - Router Solicitation and Advertisment Attacks
    - No implementations (yet)

Security Considerations

- Leaking route advertisements
  - Cisco enables RA by default
  - Windows, OS X and others will default accept
  - A machine can easily get IPv6 unnoticed

- Big threat today in IPv6 is human error
  - lack of knowledge / training
  - typos
  - Maintaining of two protocols

Configuring IPv6

Exercise

RIPE NCC

Exercise: Configuring a 6in4 tunnel

- Make sure you have connectivity
- Go to: workbench.ripe.net
- Your login is your number on attendee list

Login: userx

Password: userx_secret

- Read instructions carefully
- First discover, then configure

Real Life IPv6 Deployment

Section 7

RIPE NCC

Colocation Provider

- 30 staff
- Routing
  - Dual Stack!
  - 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, apps
- Every customer gets a /48 assignment
  - and a /64 for the connection

Colocation Provider

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

- 200 staff
- 2 /32 prefixes (due to merger)
  - not enough
  - make a plan before requesting allocation
- /48 per POP
- /56 per router
- /64 per customer vlan

- Servers
  - no EUI-64
  - no autoconfig
  - port number for services (i.e. POP3 at ::110)
  - default gateway manually set to, for example:
    - 2001:db8::1/64 *(usually)*

- **Network links (point-to-point)**
  - **core**
    - /64 per link
    - ::1 - ::2
    - no auto configuration
    - easy to remember
- **You don’t want your router link”**
    - 2001:db8:cf9d:7631:cd01:fe55:4532:ae60/64
- **You want your router link at:**
    - 2001:db8:1:1::/64

Deployment Challenges

Section 8

RIPE NCC

Deployment Challenges

- Think of a challenge/problem your organisation could have when you deploy IPv6
- Let’s see if you can find solutions!

Tips

Section 9

RIPE NCC

How to get started

- 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

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

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 4 star LIRs:**
  - http://ripeness.ripe.net

IPv6 RIPEness: 10185 LIRs

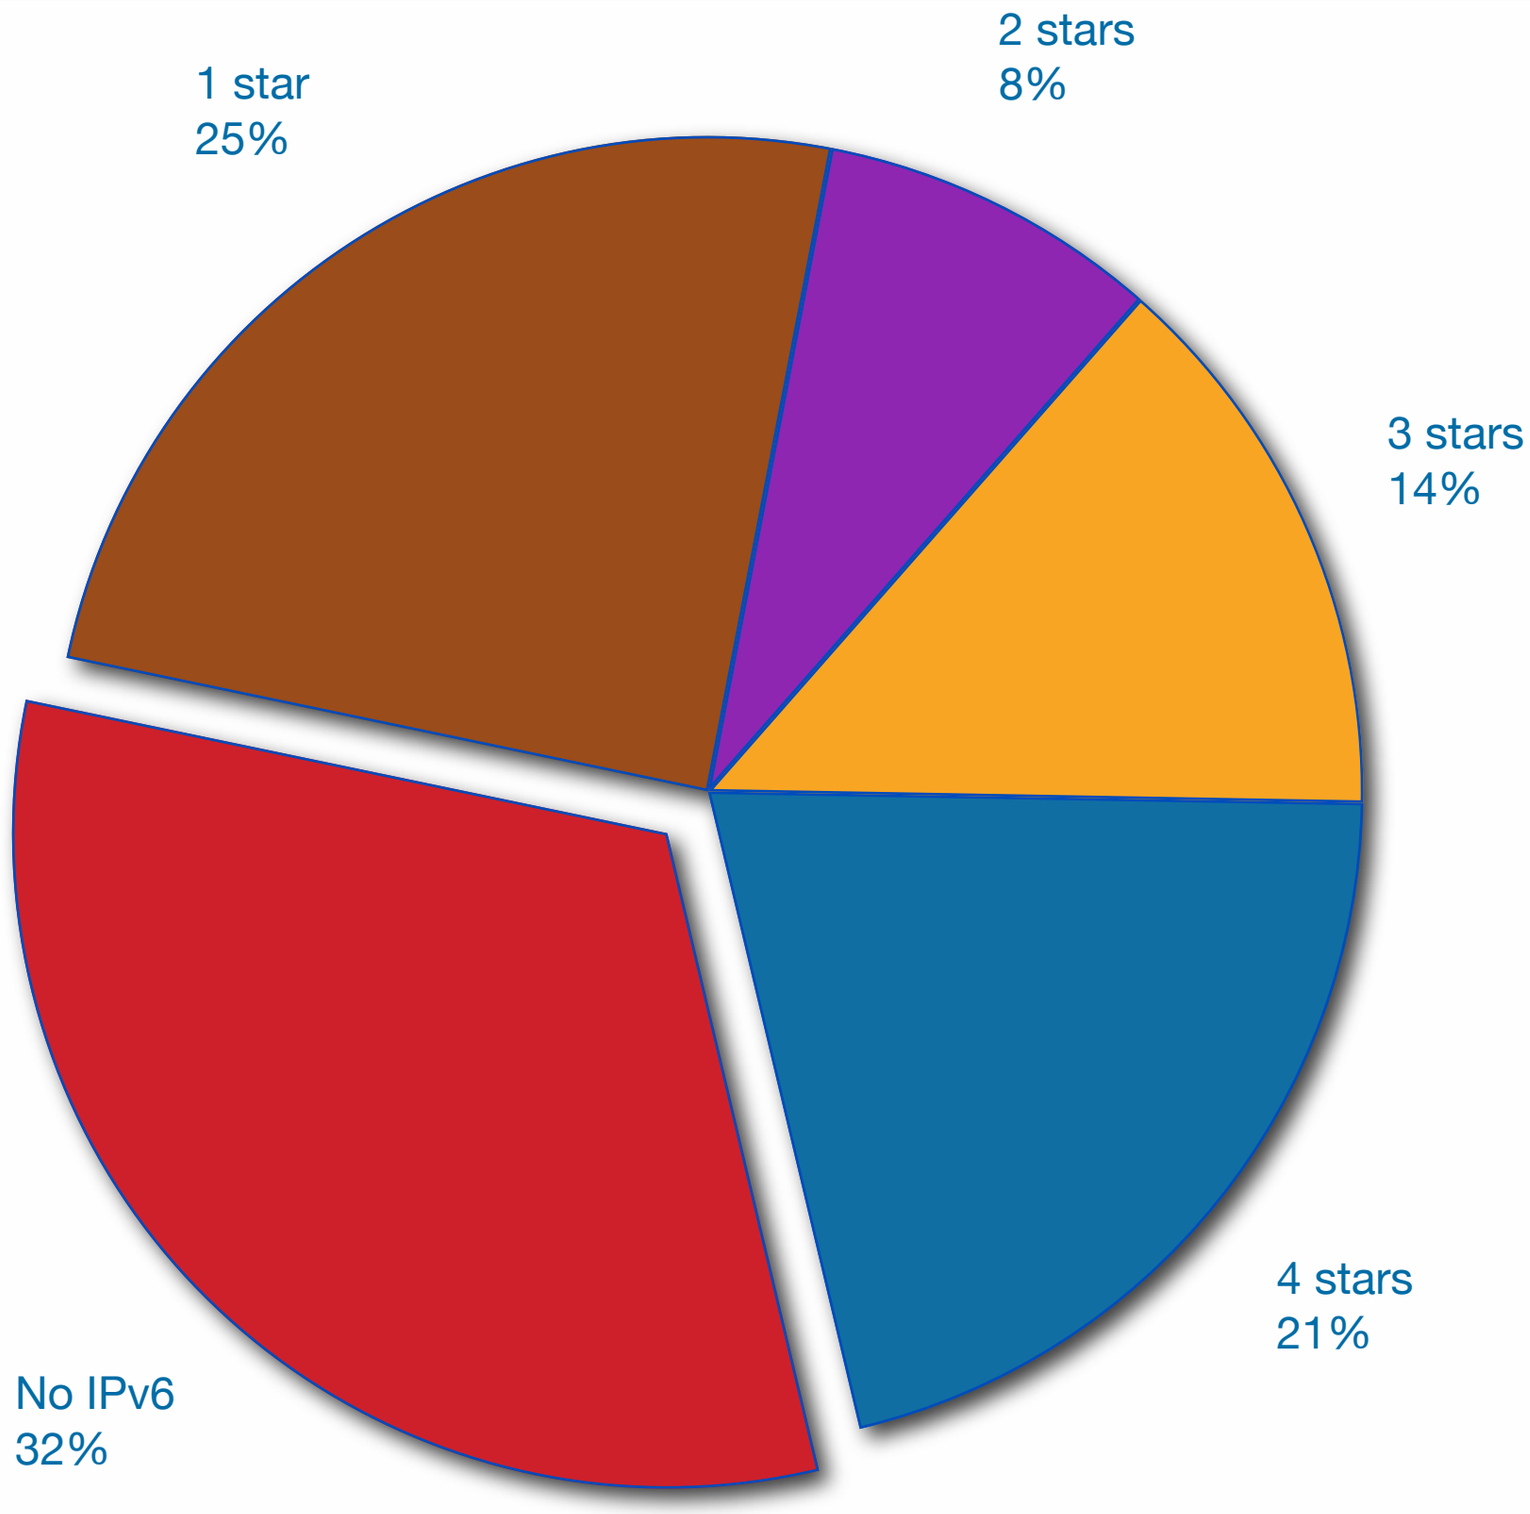

Tips

RIPE-554 Document

- "Requirements for IPv6 in ICT Equipment"
  - Best Current Practice describing what to ask for when requesting IPv6 Support
  - Useful for tenders and RFPs
  - Originated by the Slovenian Government
  - Adopted by various others (Germany, Sweden)

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

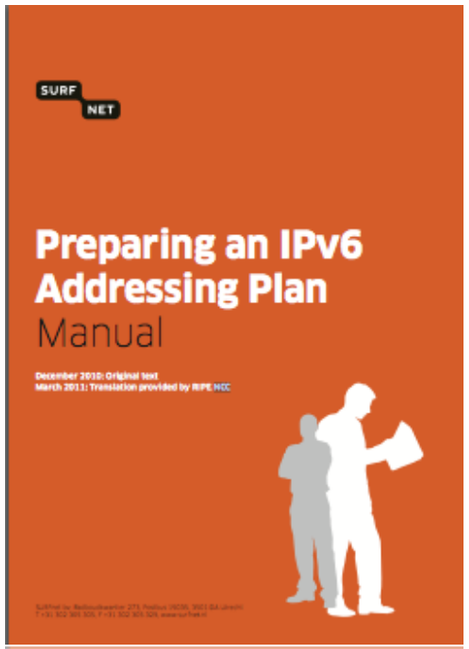

Also useful

- Websites
  - http://www.getipv6.info
  - http://www.getipv6.info
  - 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

?

Survey!

http://www.ripe.net/training/ipv6/survey

Follow us!

twitter

@TrainingRIPENCC

RIPE NCC

The End!

Край

Y Diwedd

Fí

Finis

Соңы

վերջ

النهاية

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