RIPE NCC

RIPE NETWORK COORDINATION CENTRE

Basic IPv6

Training Course

January 2016

Schedule

09:00 - 09:30	**Coffee, Tea**
11:00 - 11:15	**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
- IPv6 Protocol Basics
- Exercise: Addressing Plan
- Deploying
- Transition Mechanisms
- Exercise: Configuring IPv6
- Real Life IPv6 Deployment
- Deployment Challenges
- Tips

IPv4?

Section 1

Reaching the next billion

- **Around 3,34 billion internet users now**
  - around 46,1 % of all people
- **Mobile phones are internet devices**
- **The Internet of Things**
  - How will the Internet look like in 5 - 10 years?

The Internet of Things

Libelium Smart World

http://www.libelium.com/top_50_iot_sensor_applications_ranking

© Libelium Comunicaciones Distribuidas S.L.

IANA IPv4 Pool

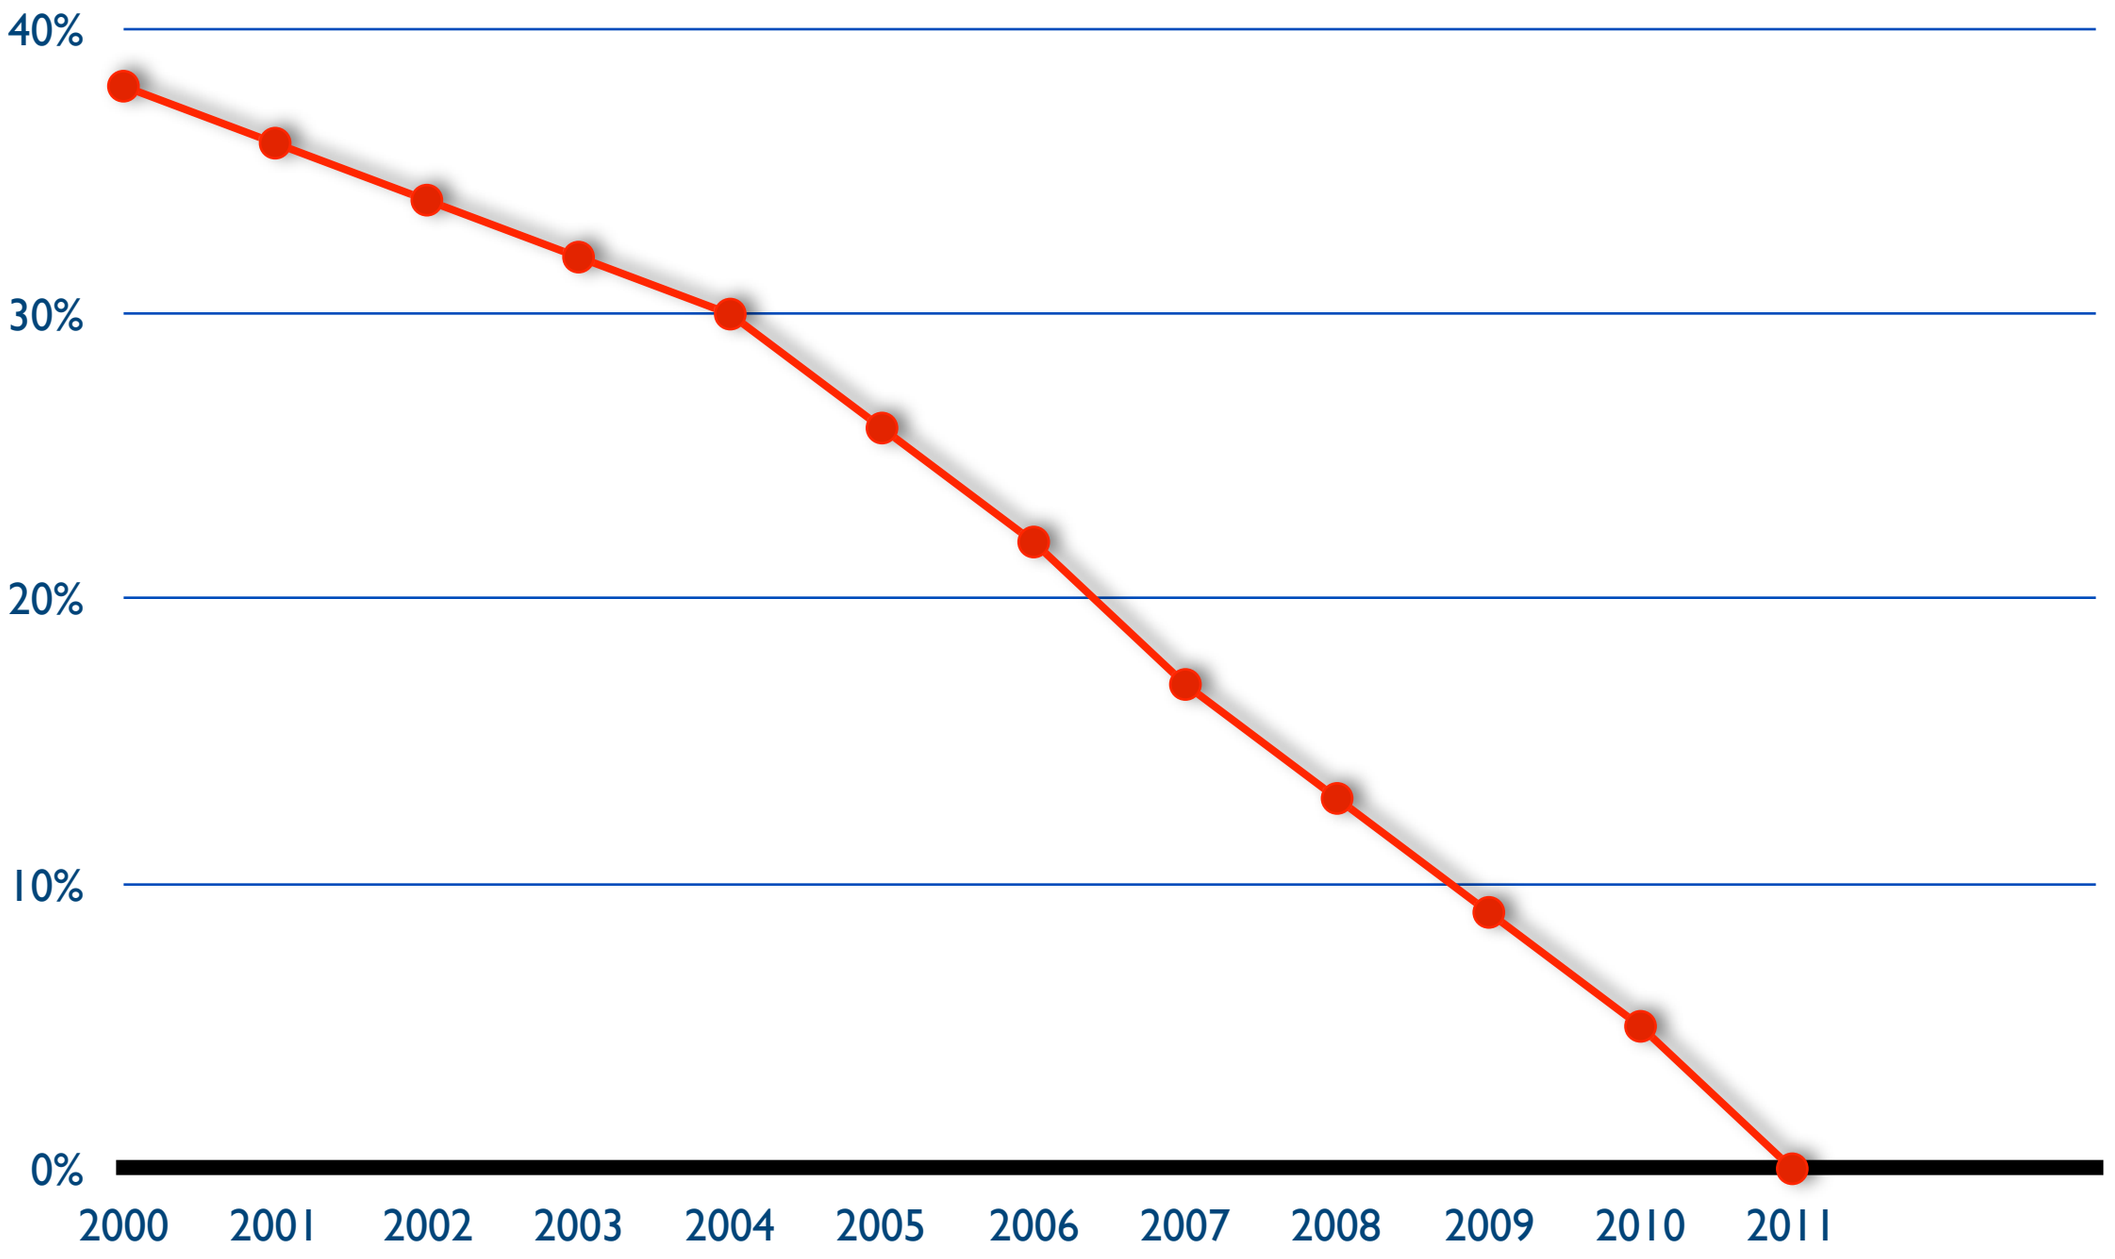

IPv4 Exhaustion

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

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

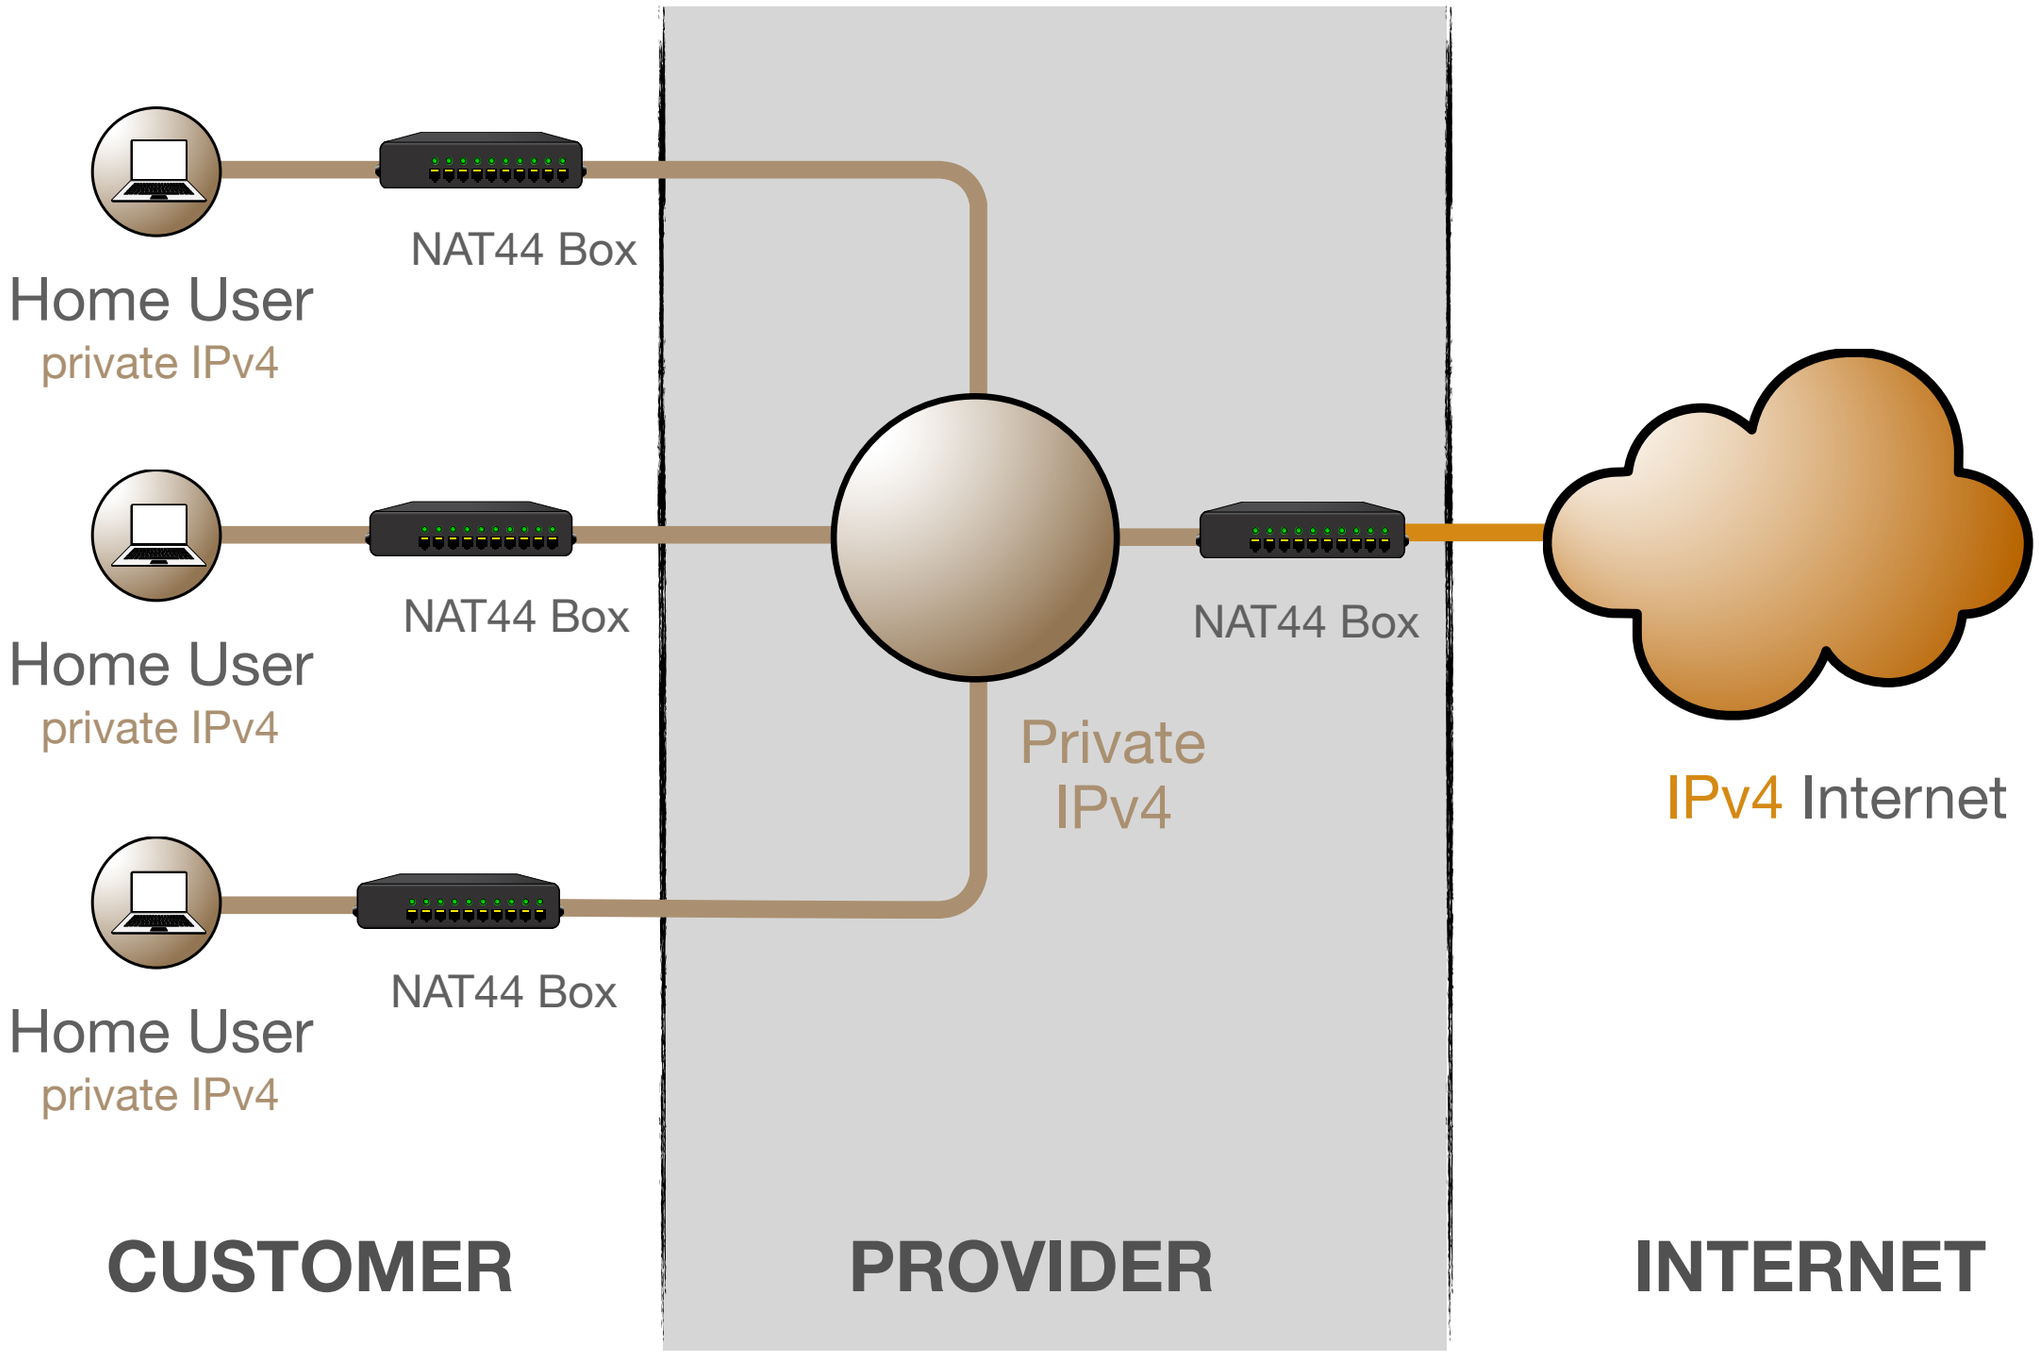

IPv6 Address Basics

Section 2

IP Address Distribution

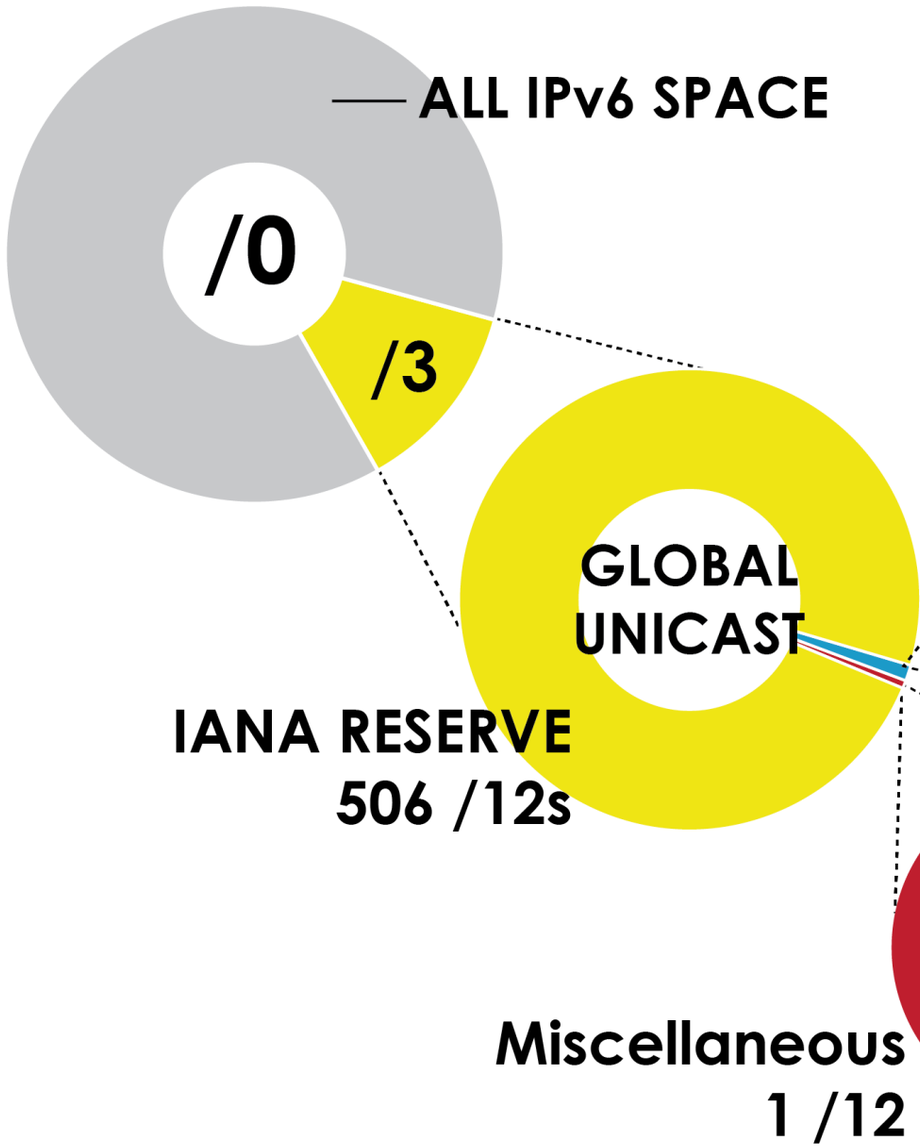

RIRs 5 /12s (October 2006)

RIR	IPv6 ADDRESS
AFRINIC	2C00:0000::/12
APNIC	2400:0000::/12
ARIN	2600:0000::/12
LACNIC	2800:0000::/12
RIPE NCC	2A00:0000::/12

Source: https://www.nro.net/statistics
Number Resource Organisation

IP Address Distribution

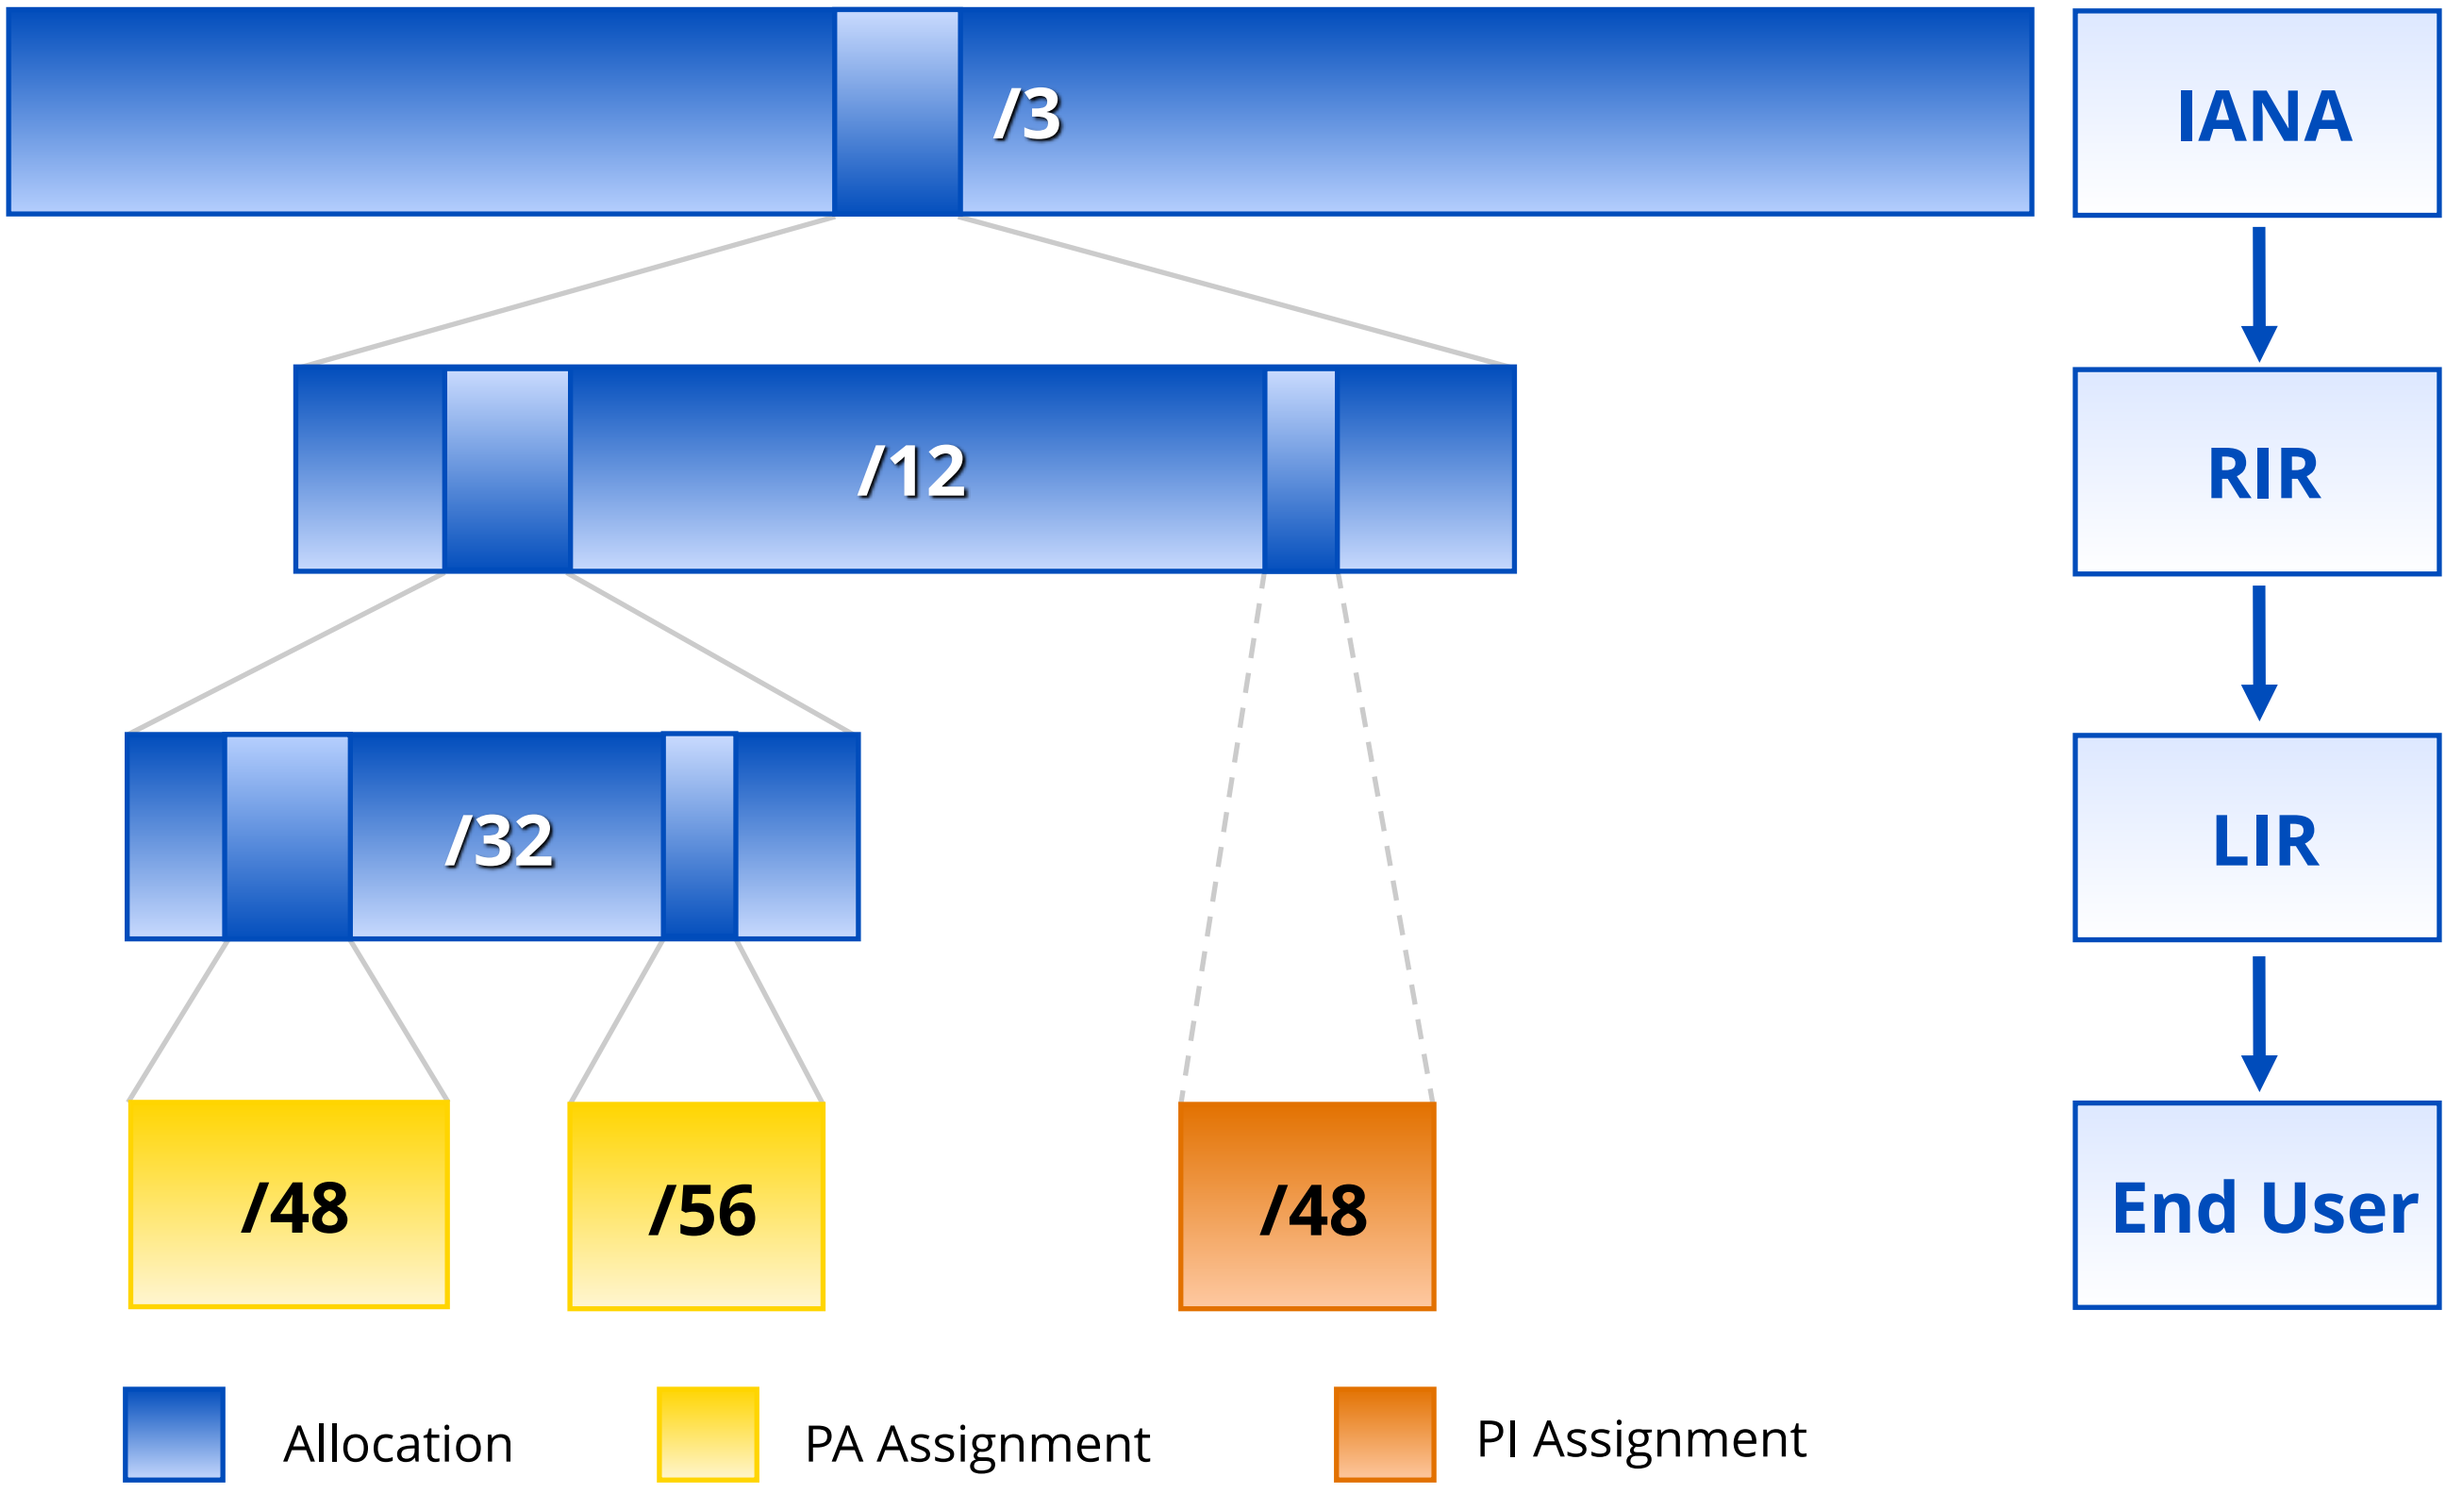

IPv6 Address Basics

- **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

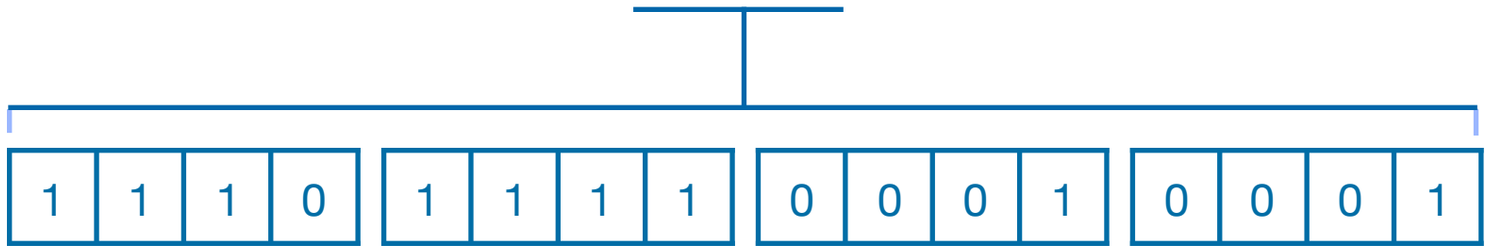

IPv6 Subnetting

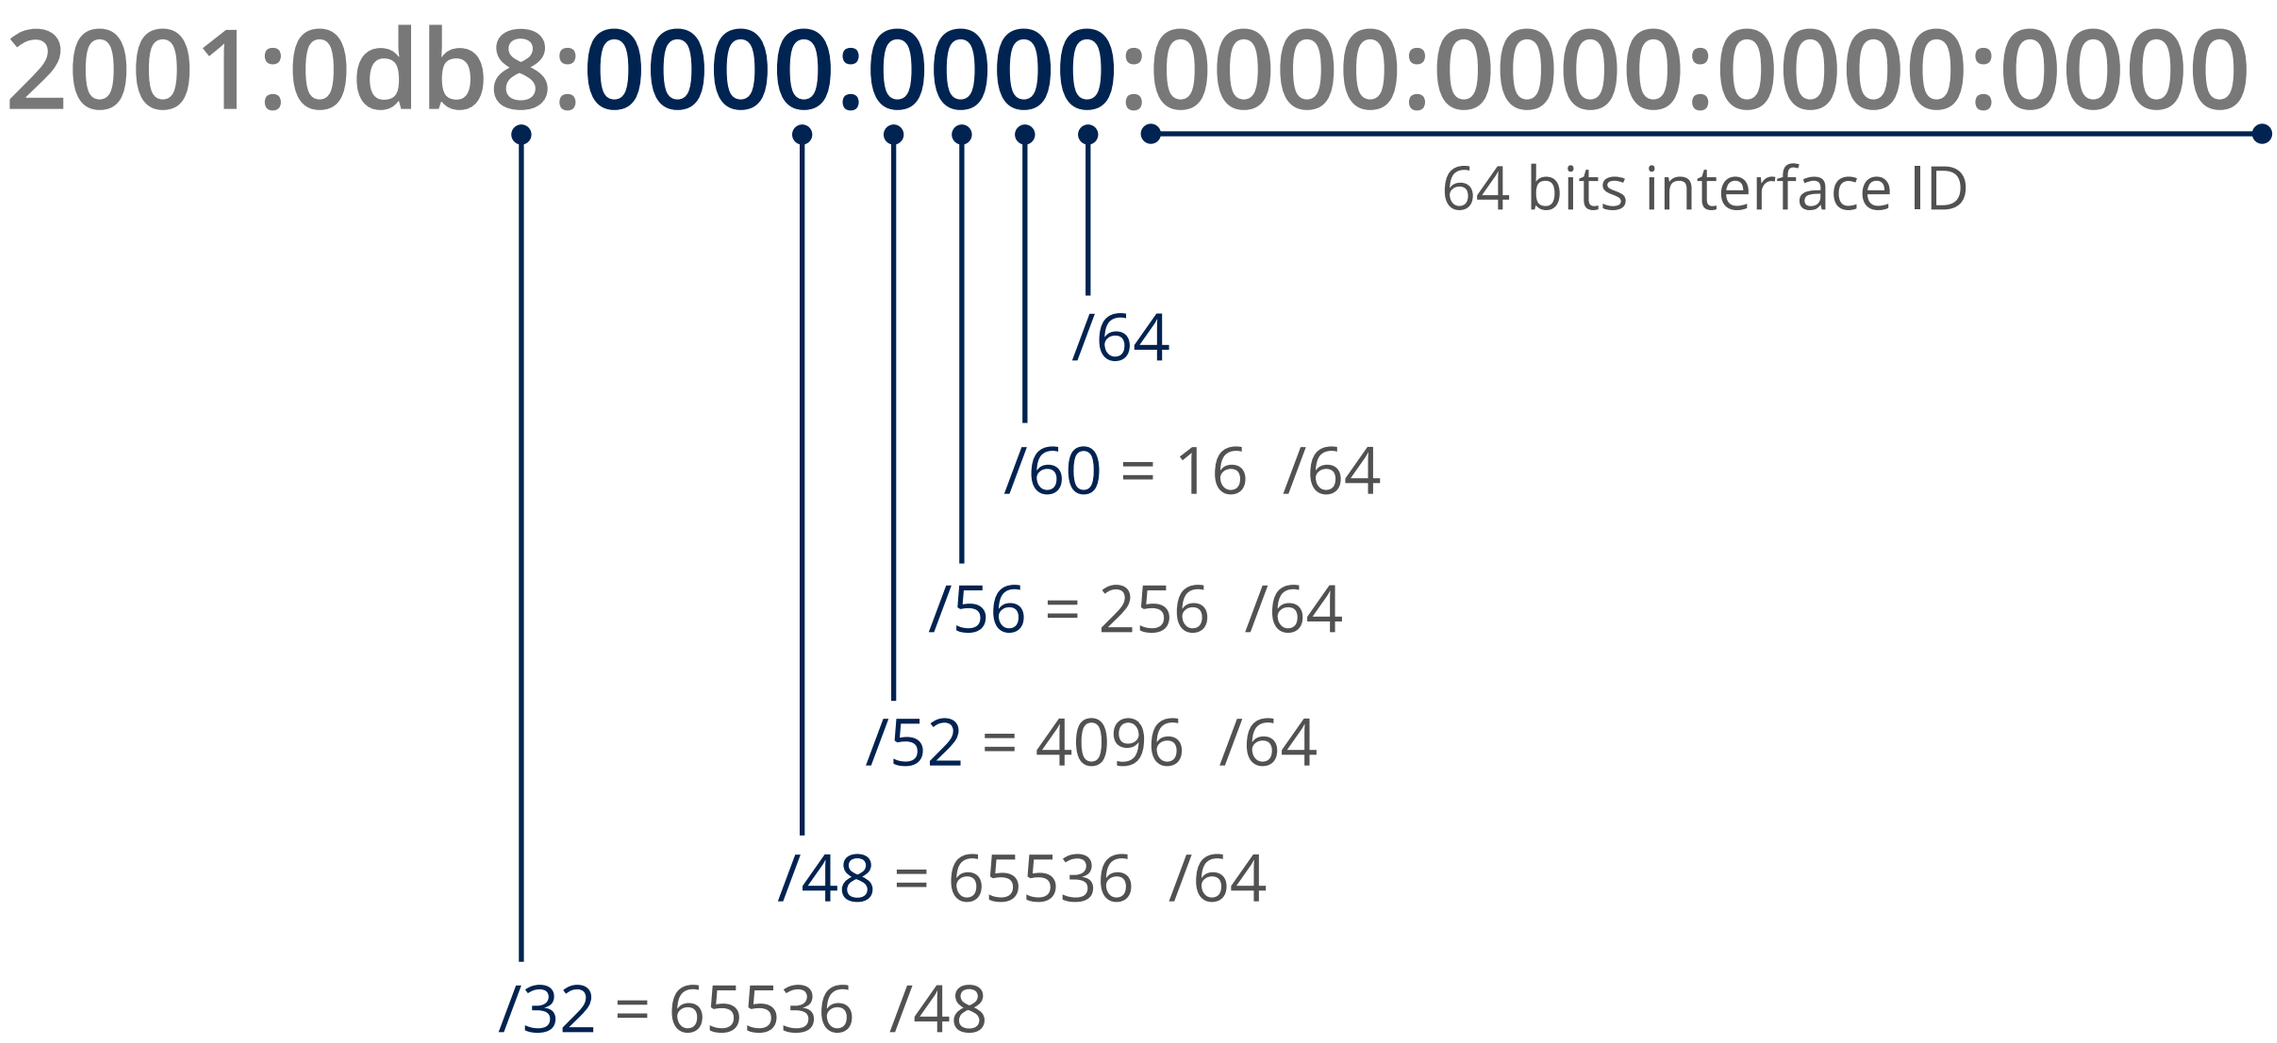

Multiple address types

Addresses	Range	Scope
Unspecified	::/128	n/a
Loopback	::1	host
Discard-Only	100::/64	n/a
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

Solicited Node Multicast Address

- A multicast address valid only on the local-link
- Used in Neighbor Discovery Protocol for obtaining the layer 2 link-layer (MAC) addresses
- Made from the ff02::1:ff00:0/104 prefix and the last 24 bits of an IPv6 address

IPv6 unicast address:	2001:0db8:b7f1:cf8d:23b1:8301:13**d9:aa6f**
Solicited Node address:	ff02::1:ff**d9:aa6f**

IPv6 Address Notation

Exercise

Questions

Getting It

Section 3

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 and network extension
  - Additional bits based on hierarchical and geographical structure, planned longevity and security levels

Customer Assignments

- **Give your customers enough addresses**
  - Minimum /64
  - Up to /48
- **More than /48, 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	Allocation	ALLOCATED-BY-RIR
ASSIGNED PA	Assignment	ASSIGNED
ASSIGNED PA	Group of Assignments	AGGREGATED-BY-LIR
SUB-ALLOCATED PA	Sub-Allocation	ALLOCATED-BY-LIR
ASSIGNED PI	PI Assignment	ASSIGNED PI

Using ASSIGNED

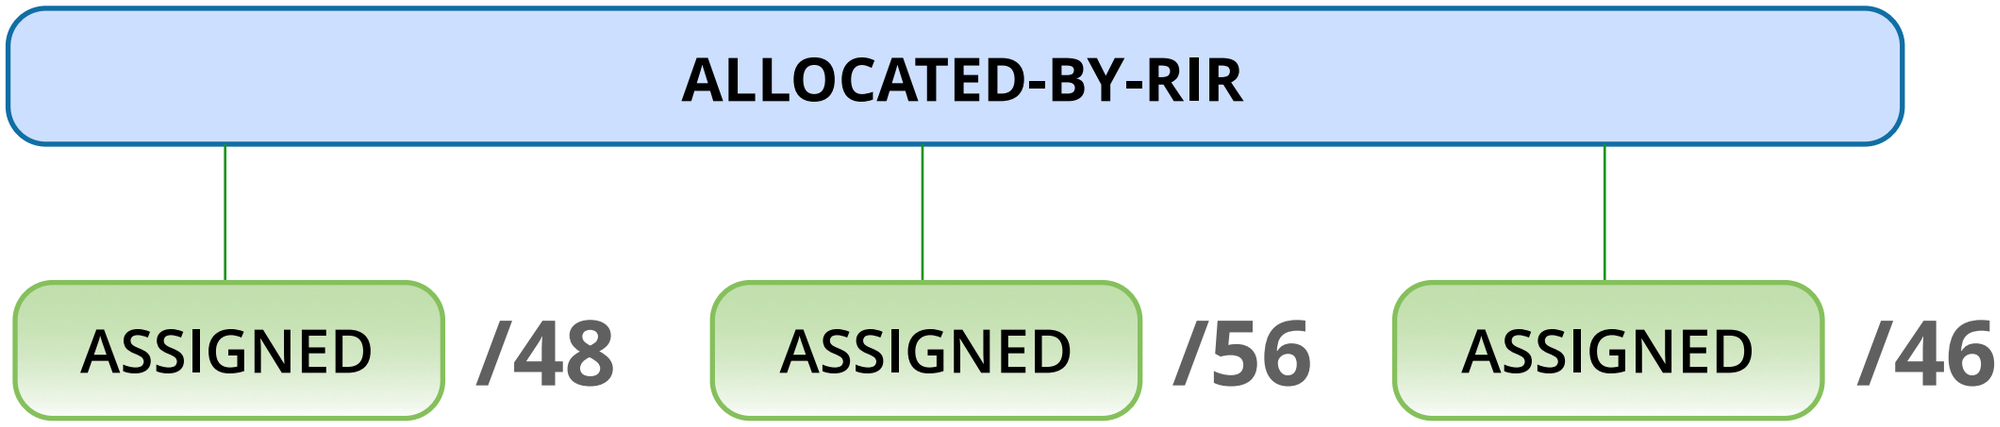

- Represents one assignment
- Minimum assignment size is a /64
- For more than a /48, send a request form

Using AGGREGATED-BY-LIR

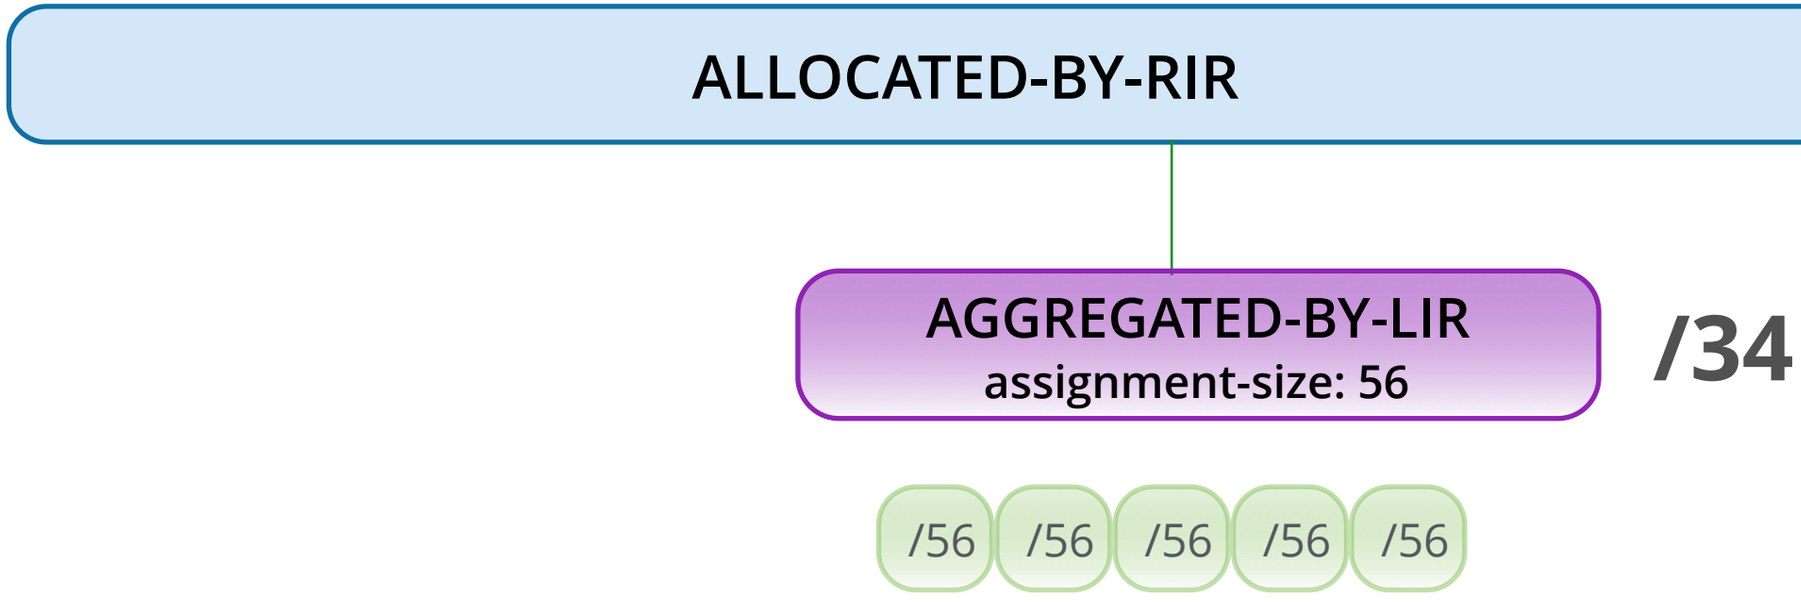

- **Can be used to group customers**
  - example: residential broadband customers
- **“assignment size:” = assignment of each customer**

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
created:	2015-05-31T12:34:01Z
last-modified:	2015-05-31T12:34:01Z
source:	RIPE

Using ALLOCATED-BY-LIR

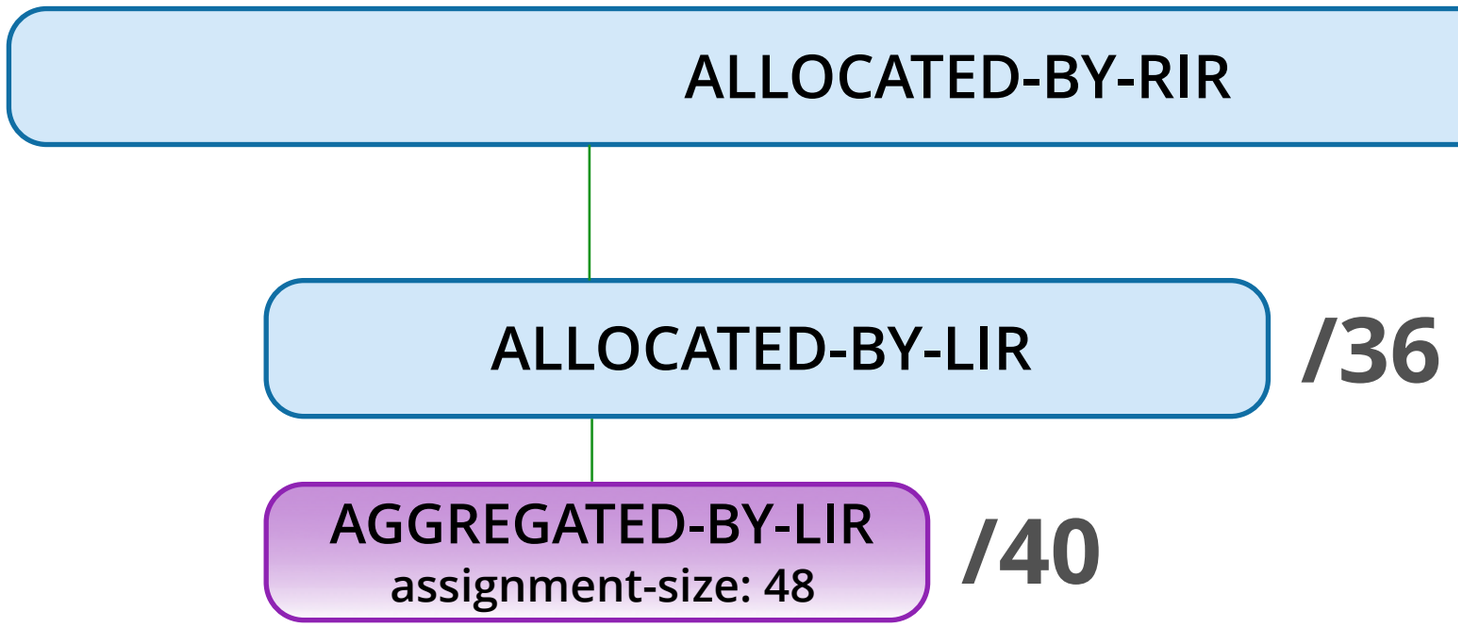

- Can be used for customers with potential for growth
  - or for your own infrastructure
  - or to delegate address space to a downstream ISP

Overview

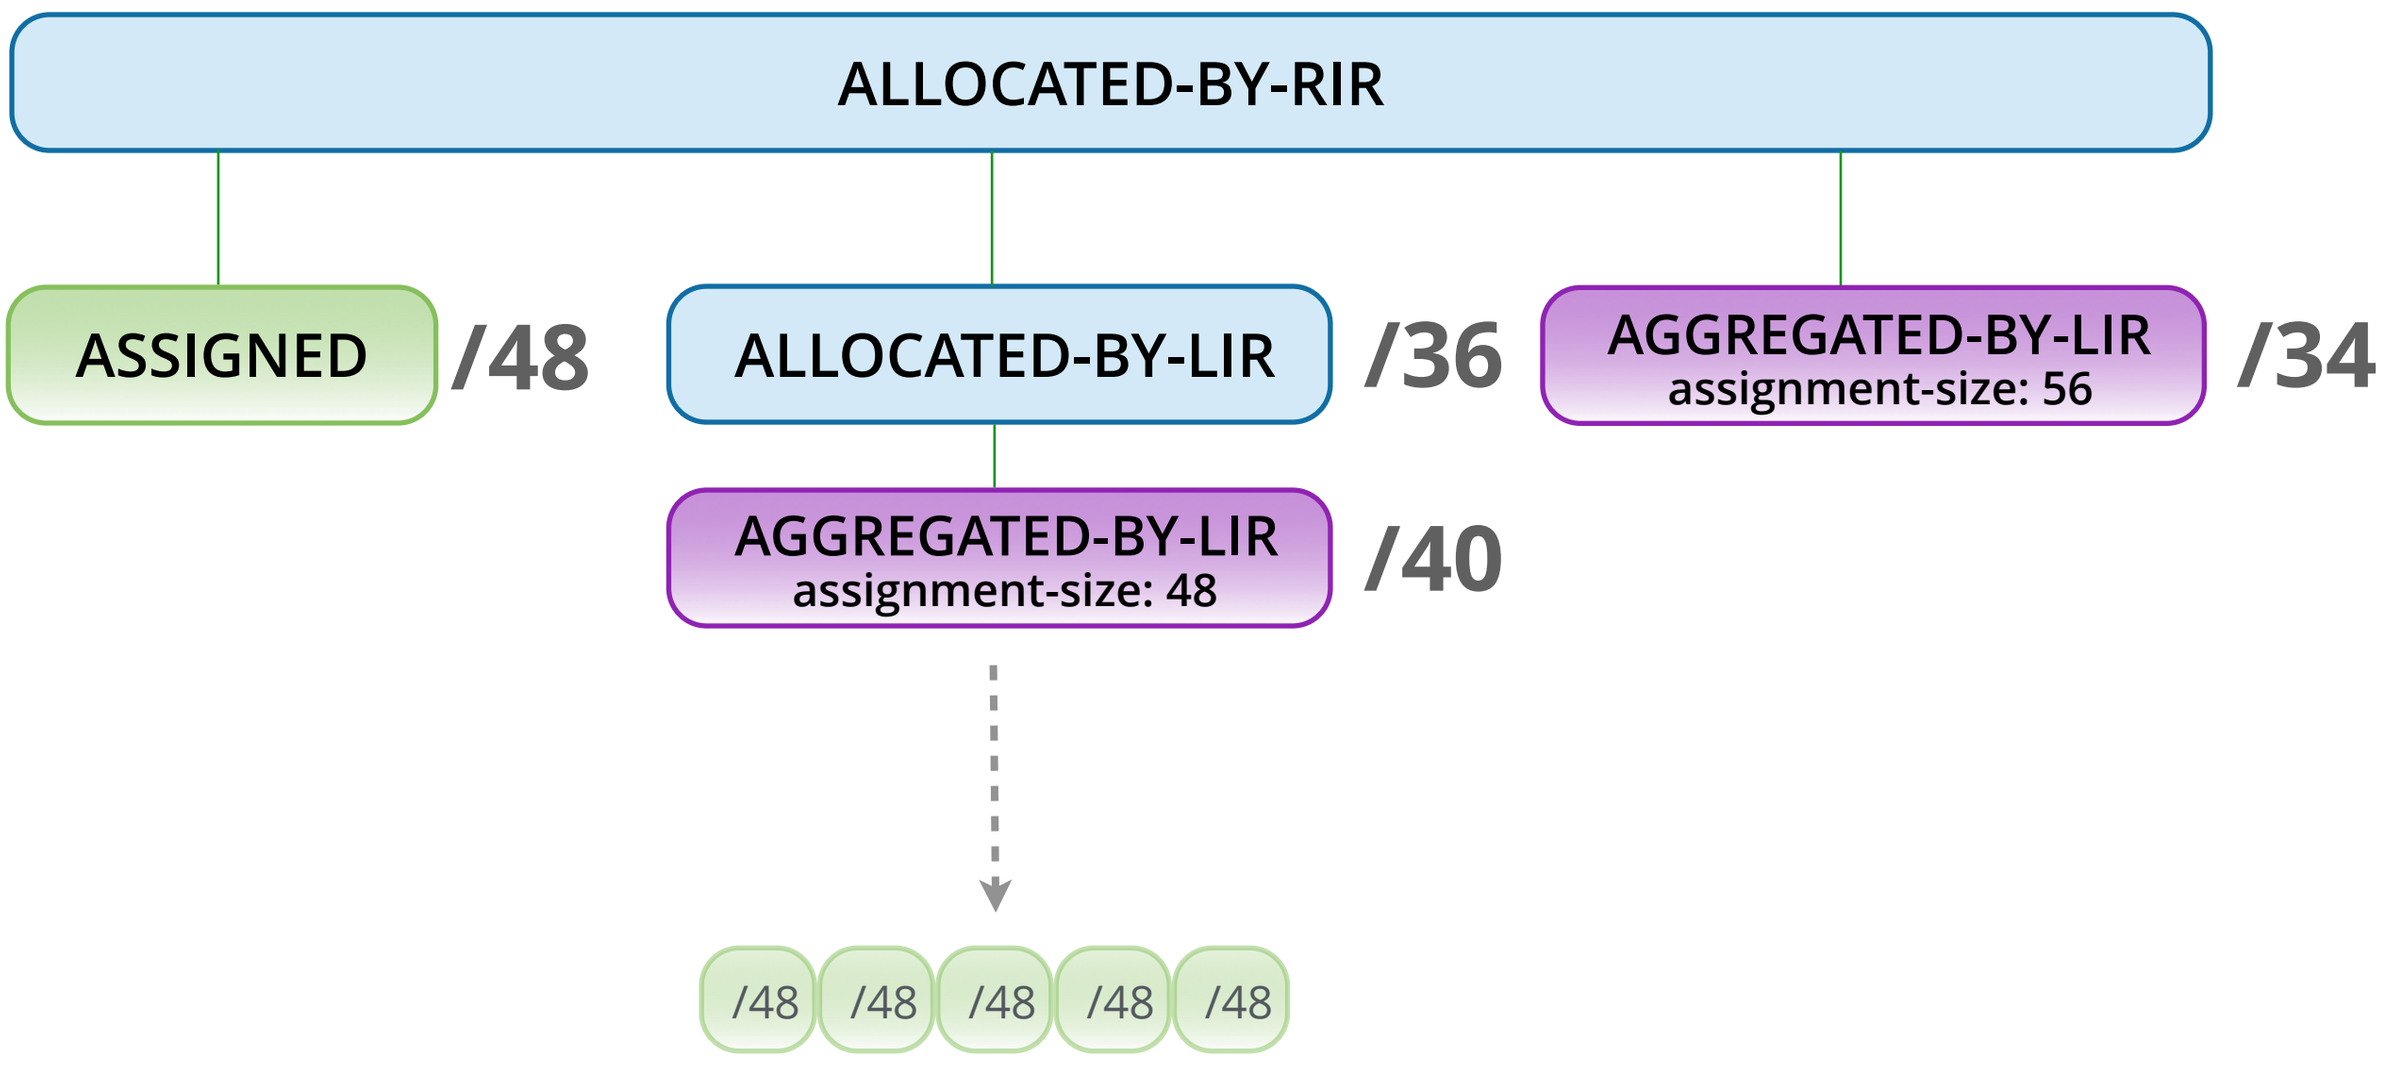
ALLOCATED-BY-RIR
ASSIGNED
/48
ALLOCATED-BY-LIR
/36
AGGREGATED-BY-LIR
assignment-size: 56
/34
AGGREGATED-BY-LIR
assignment-size: 48
/40
/48
/48
/48
/48
/48

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!

Unique Local Addresses

- **Prefixes from fc00::/7**
  - Only from the fd00::/8 block
- **Should not be routed on the Internet**
- **Generate a random 40-bit Global ID and insert it into fdxx:xxxx:xxxx**

Global ID:	da24154e1d
Prefix:	fdda:2415:4e1d::/48

- **You may want to register your prefix!**

Making Assignments

Exercise

Making Assignments Exercise

Smart Home 6!

- 20 minutes preparation time
- 10 minutes discussion

Solution RIPE Database object

inet6num:	2001:db8:1000::/36
netname:	SMART-HOME-6
descr:	Smart Home 6 network
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
created:	2015-05-31T12:34:01Z
last-modified:	2015-05-31T12:34:01Z
source:	RIPE

Solution RIPE Database object

inet6num:	2001:db8:1000::/36
netname:	SMART-HOME-6
descr:	Smart Home 6 network
country:	NL
admin-c:	RM1204-RIPE
tech-c:	RM1204-RIPE
status:	ALLOCATED-BY-LIR
mnt-by:	LIR-MNT
mnt-lower:	SMART-CASA-MNT
notify:	noc@lir-example.com
created:	2015-05-31T12:34:01Z
last-modified:	2015-05-31T12:34:01Z
source:	RIPE

IPv6 Protocol Basics

Section 4

IPv6 Protocol Functions

- **Address Autoconfiguration**
  - Supported by Neighbor Discovery
  - Stateless - with SLAAC
  - Stateful - with DHCPv6

- **Neighbor Discovery Protocol**
  - Replaces ARP from IPv4
  - Uses ICMPv6 and Multicast
  - Finds the other IPv6 devices on the link
  - Keeps track of reachability

The Autoconfiguration Process

1. Make a Link-Local address
2. Check for duplicates on the link
3. Search for a router
4. Make a Global Unicast address

Making a Link-Local Address

- Interface ID is made from the MAC address

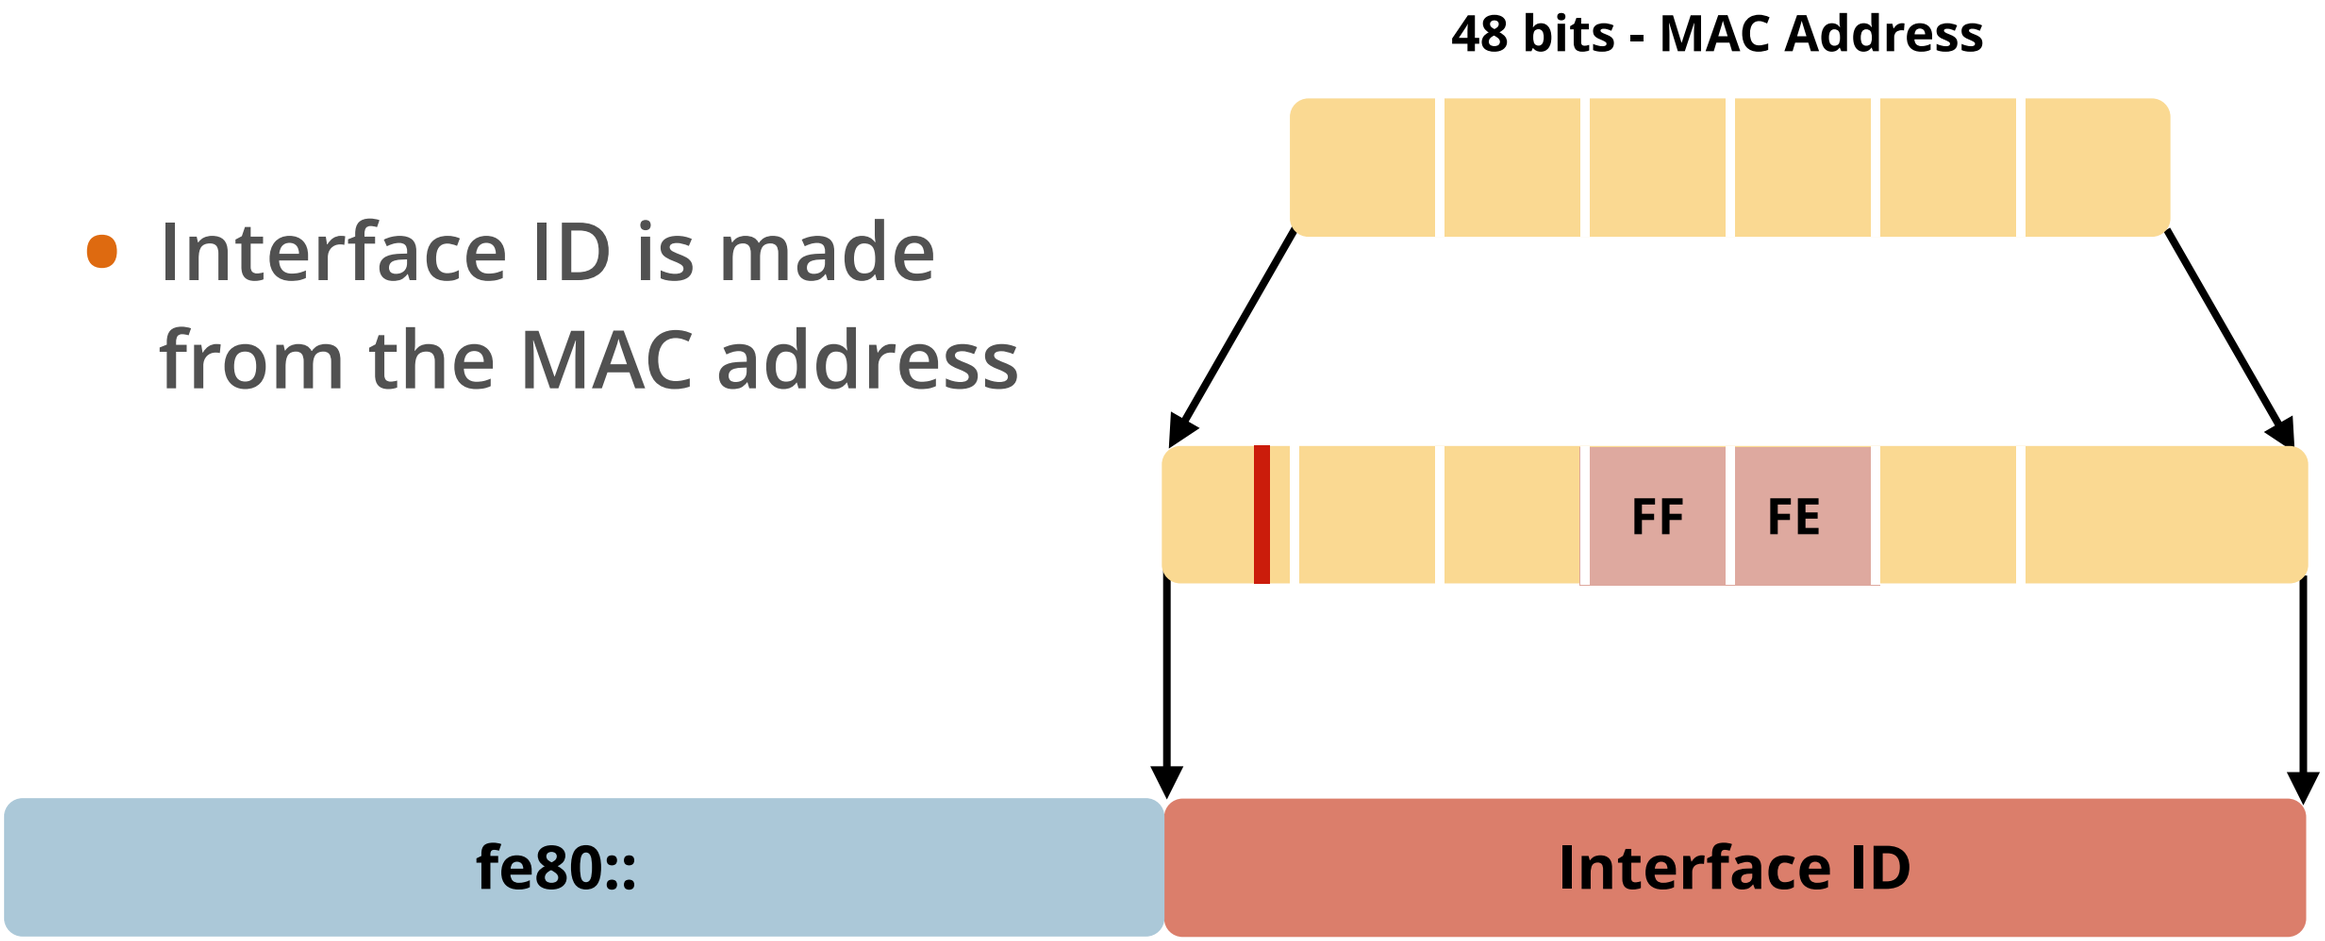

- fe80:: + Interface ID = Link-Local address for the host

Checking for Duplicates

Neighbor Solicitation

Hello! Is this IPv6 address in use?
Can you tell me your MAC address?

Neighbor Advertisement

Hello! Yes, I'm using that IPv6 address.
My MAC address is 72:D6:0C:2F:FC:01

If nobody replies to the Neighbor Solicitation, the host uses the generated link-local address

Searching for Routers

Router Solicitation

Hello! Is there a router out there?

Router Advertisement

Hello! I'm a router and I have some information for you…

The Router Advertisement gives the host more information to get an IPv6 address and set up a connection

Stateless Address Auto-Configuration

- **The Router Advertisement message tells the host:**
  - Router's address
  - Zero or more link prefixes
  - SLAAC allowed (yes/no)
  - DHCPv6 options
  - MTU size (optional)

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)

Verifying Reachability

Neighbor Solicitation

Hello! Are you still out there?
Is your MAC address still valid?

Neighbor Advertisement

Hello! Yes, I'm still online.
My MAC address is 72:D6:0C:2F:FC:01

If the target does not reply to the Neighbor Solicitation, the sender removes the MAC address from the cache

Redirects

IPv6 Packet

This packet is for an IPv6 host.

Redirect

Hello! That destination you wanted?
I know a better way to reach it.

- Hosts can be redirected to a better first-hop router
- They can also be informed that the destination is a neighbor on the link

Questions

Addressing Plans

Section 5

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 512K /48s in a /29
  - 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

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

- **/64 for all subnets**
- **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**

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…

Questions

Deploying IPv6

Section 6

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 collision, a new random address is generated

DHCPv6

- **Used to get information like DNS servers or to manage the address pool**

- **Router Advertisement message contains hints**

  - If “managed config” flag is set to ‘1’, the host can use DHCPv6 to get an address

  - Optionally the address of a DNS server (RFC 6106)

- **Using additional flags, the network admin can disable SLAAC and force DHCPv6**

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.0.0.0.0.0.0.0.0.1.1.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 Advertisement Attacks

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

Transition Mechanisms

Section 7

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**
  - Tunnelling IPv6 packets over IPv4-only networks

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

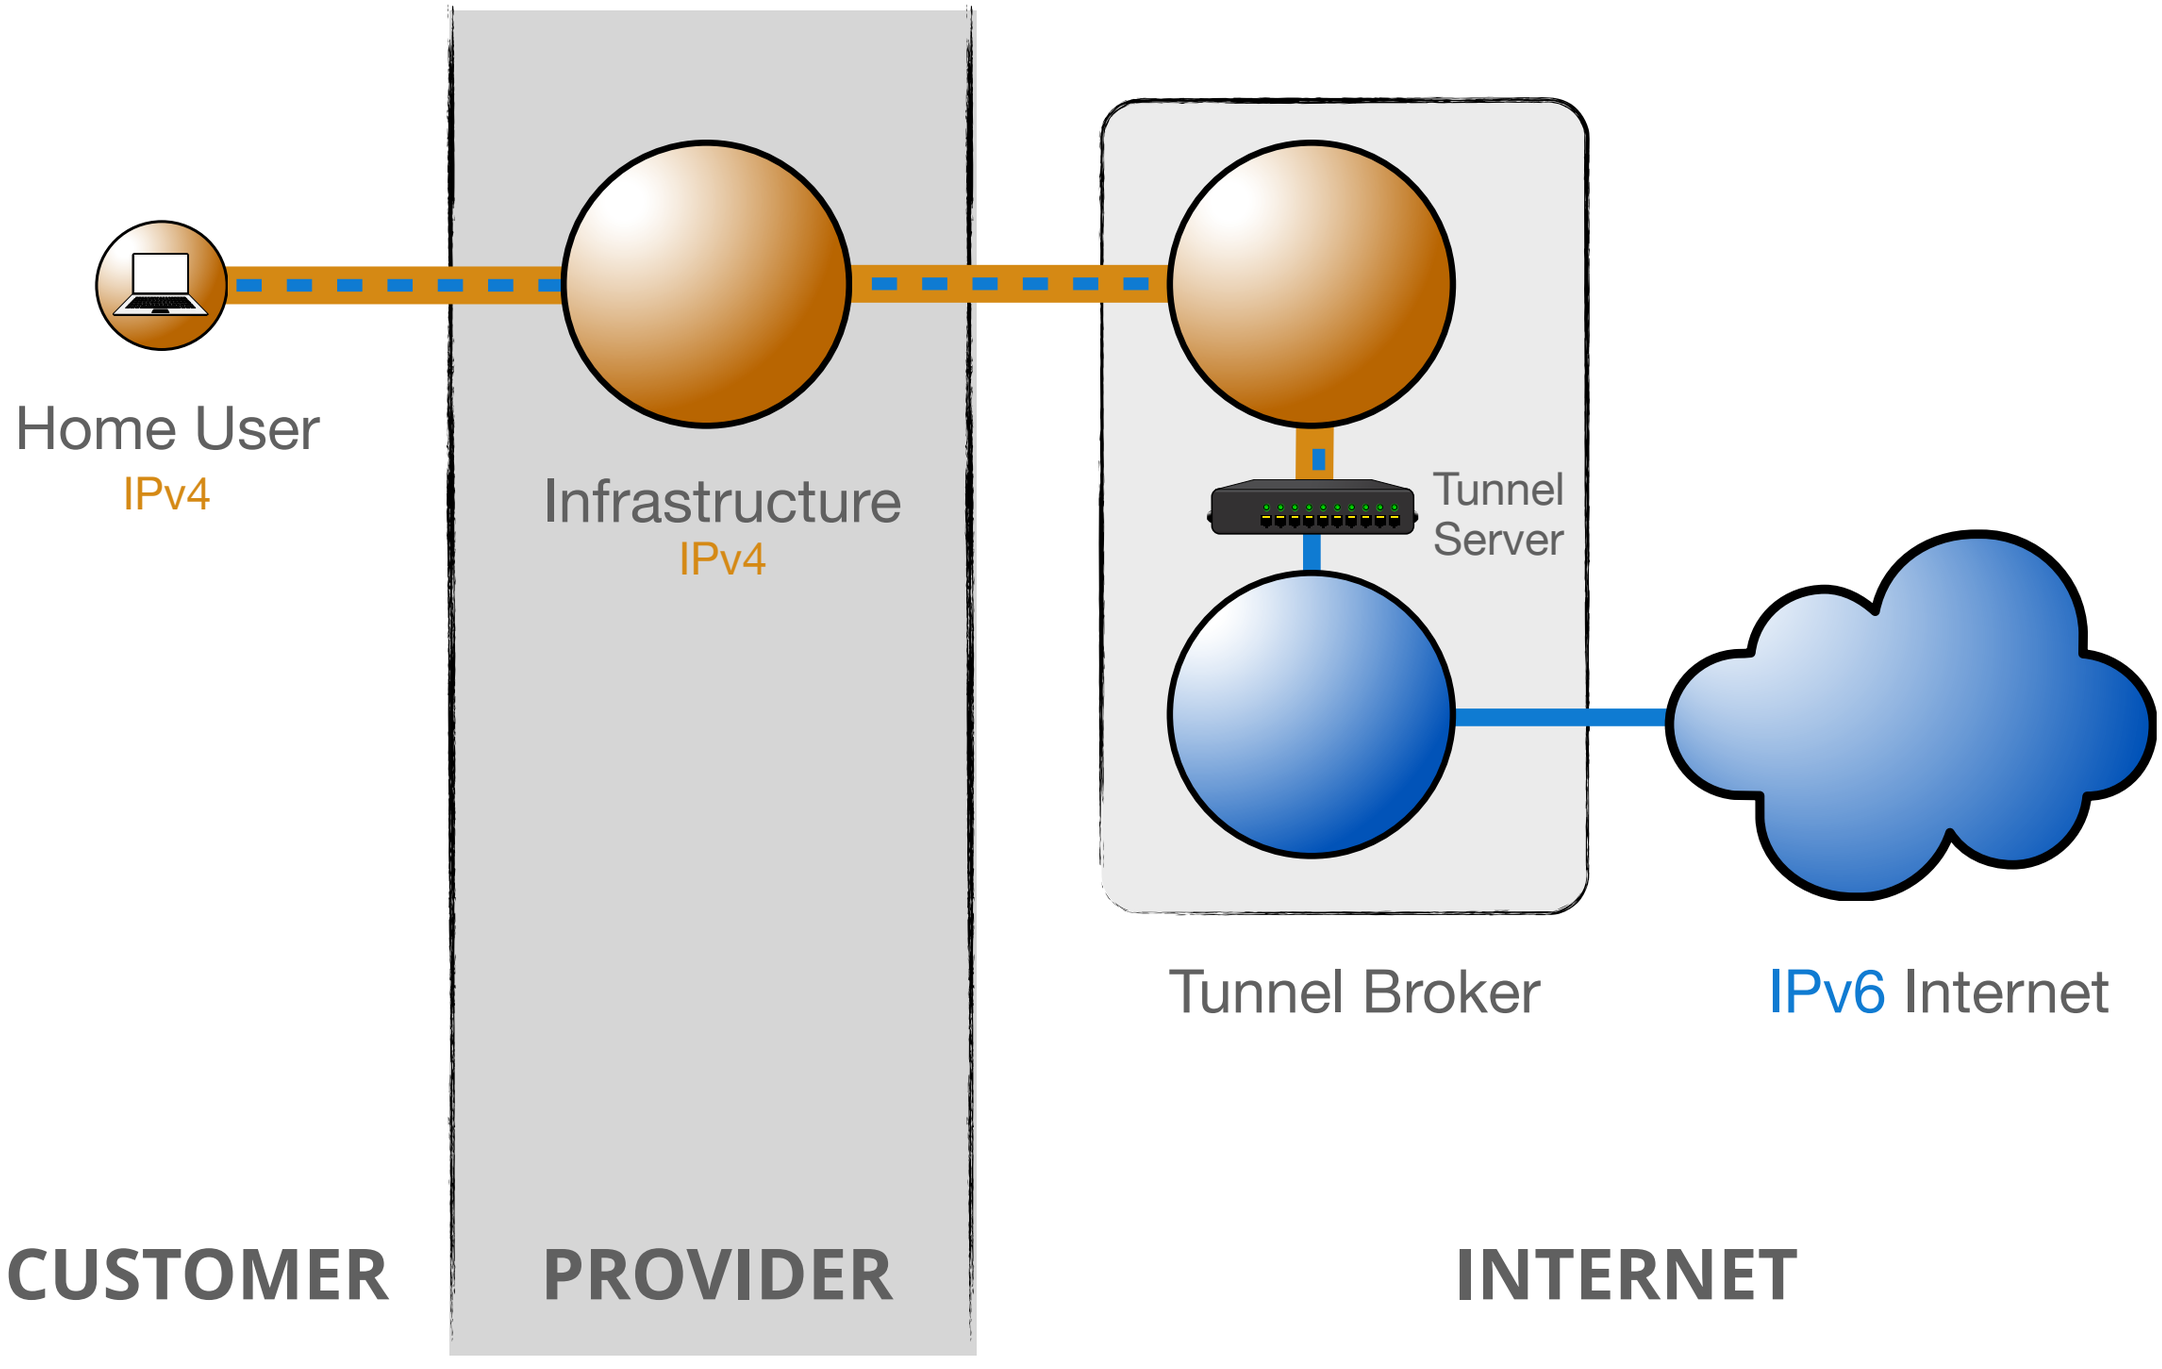
Home User
IPv4
Infrastructure
IPv4
Tunnel
Server
Tunnel Broker
IPv6 Internet
CUSTOMER
PROVIDER
INTERNET

6to4 and Teredo

- **6to4 (deprecated)**
  - "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

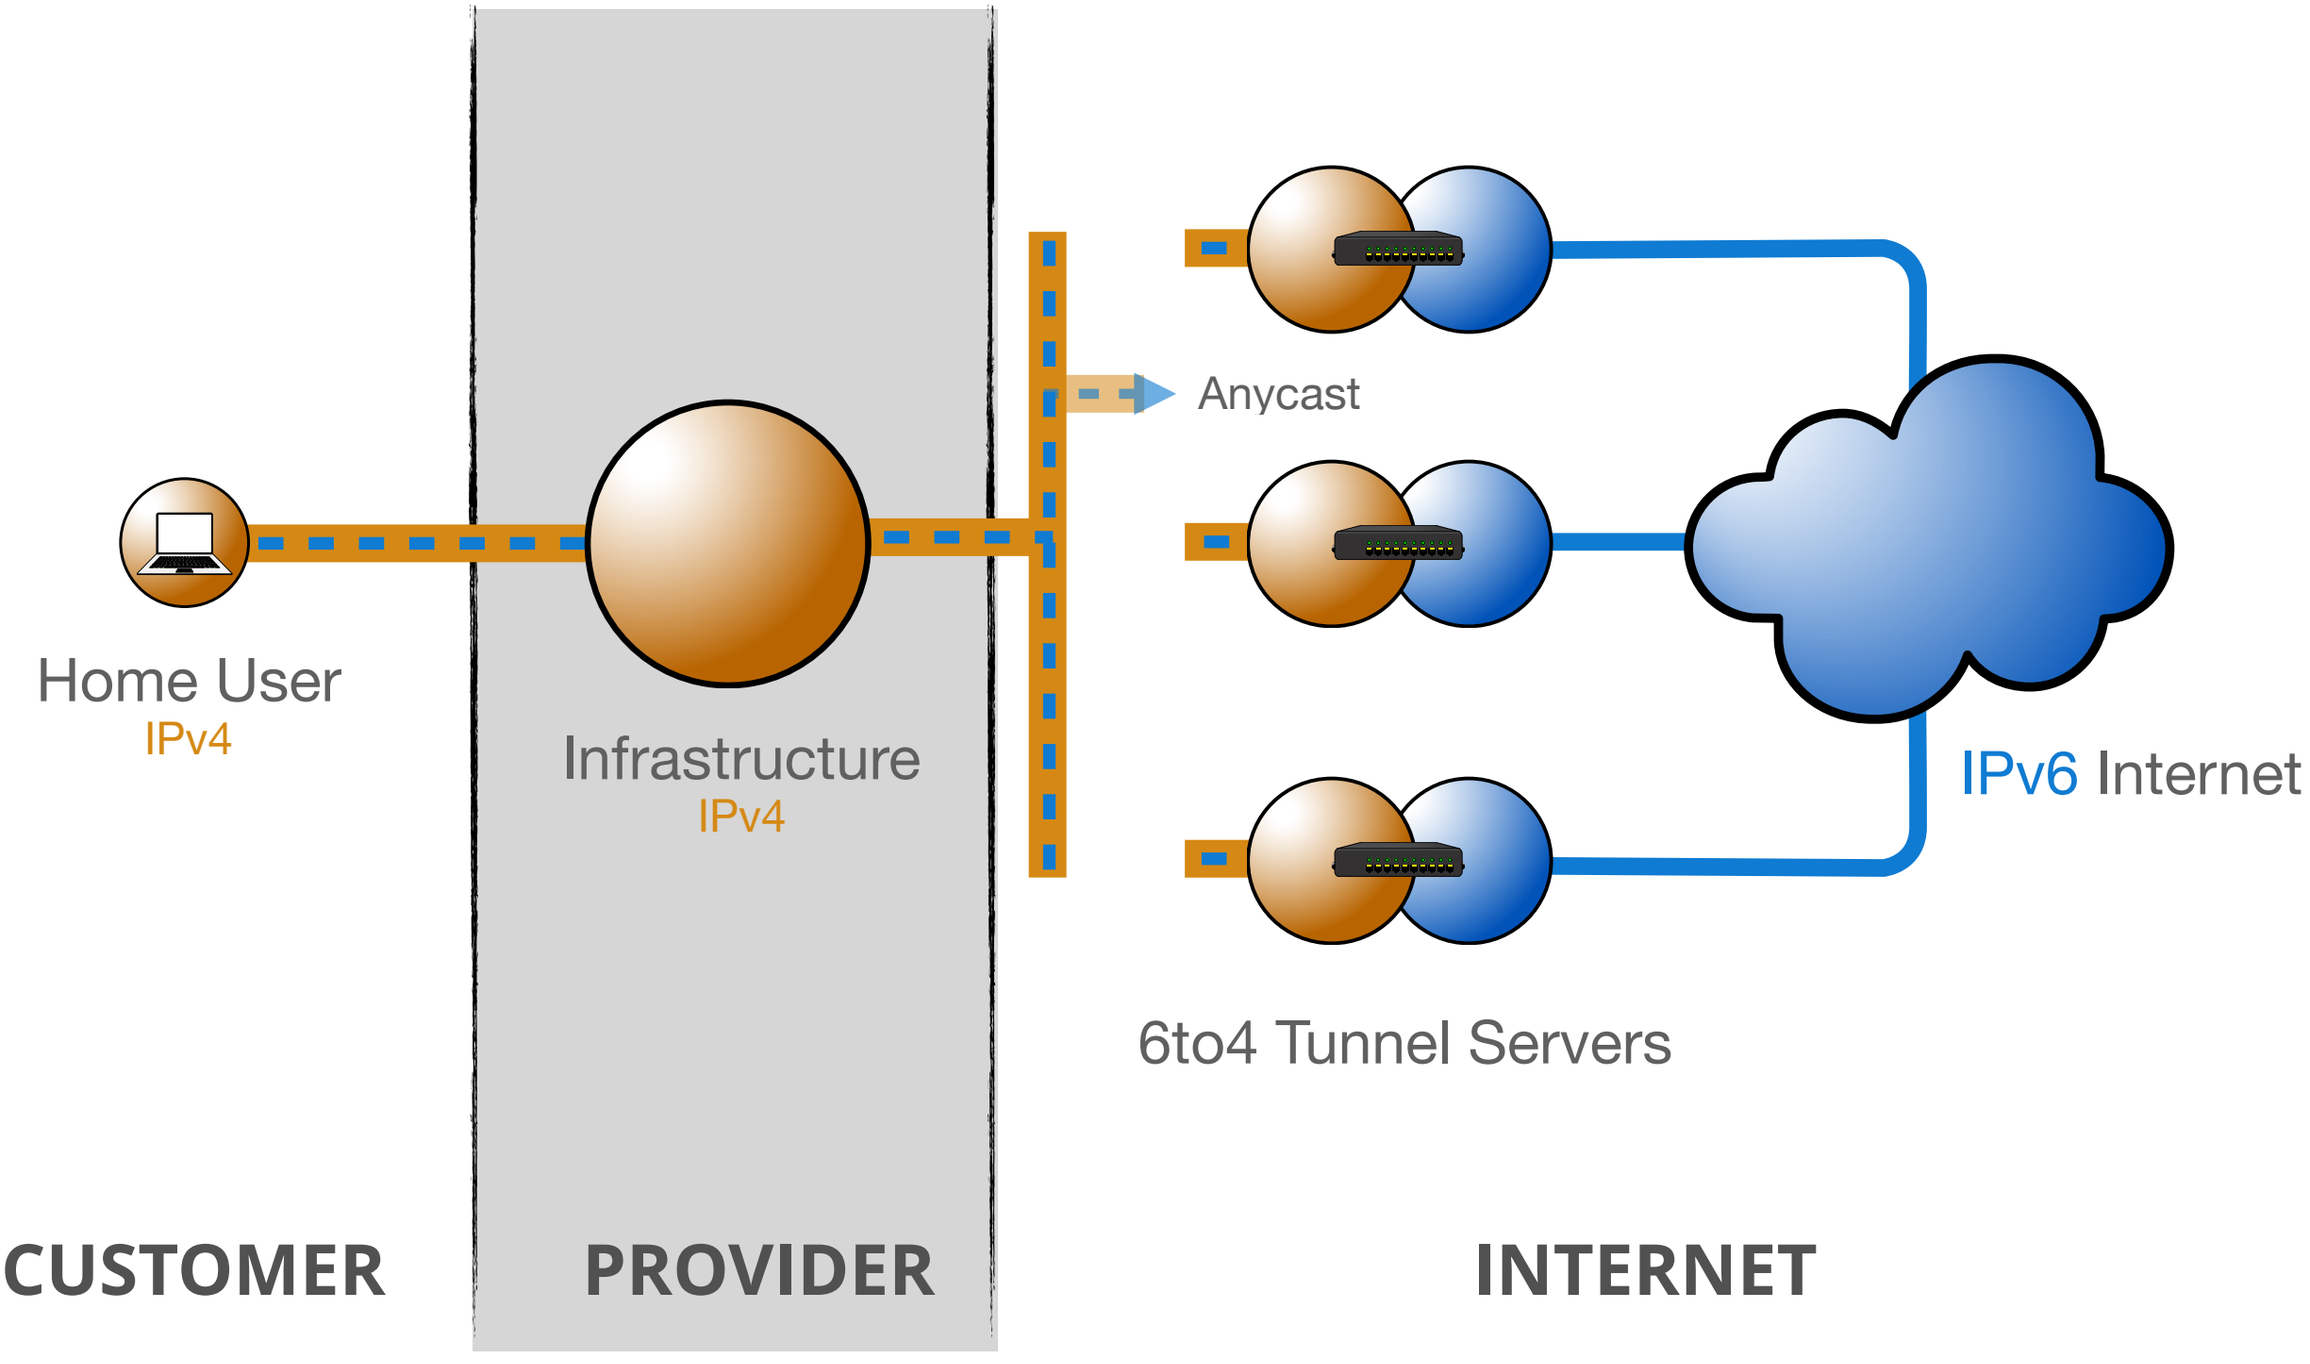

6RD

- **Encodes the IPv4 address in the IPv6 prefix**
  - Quite similar to 6to4
- **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 signalling**

6RD

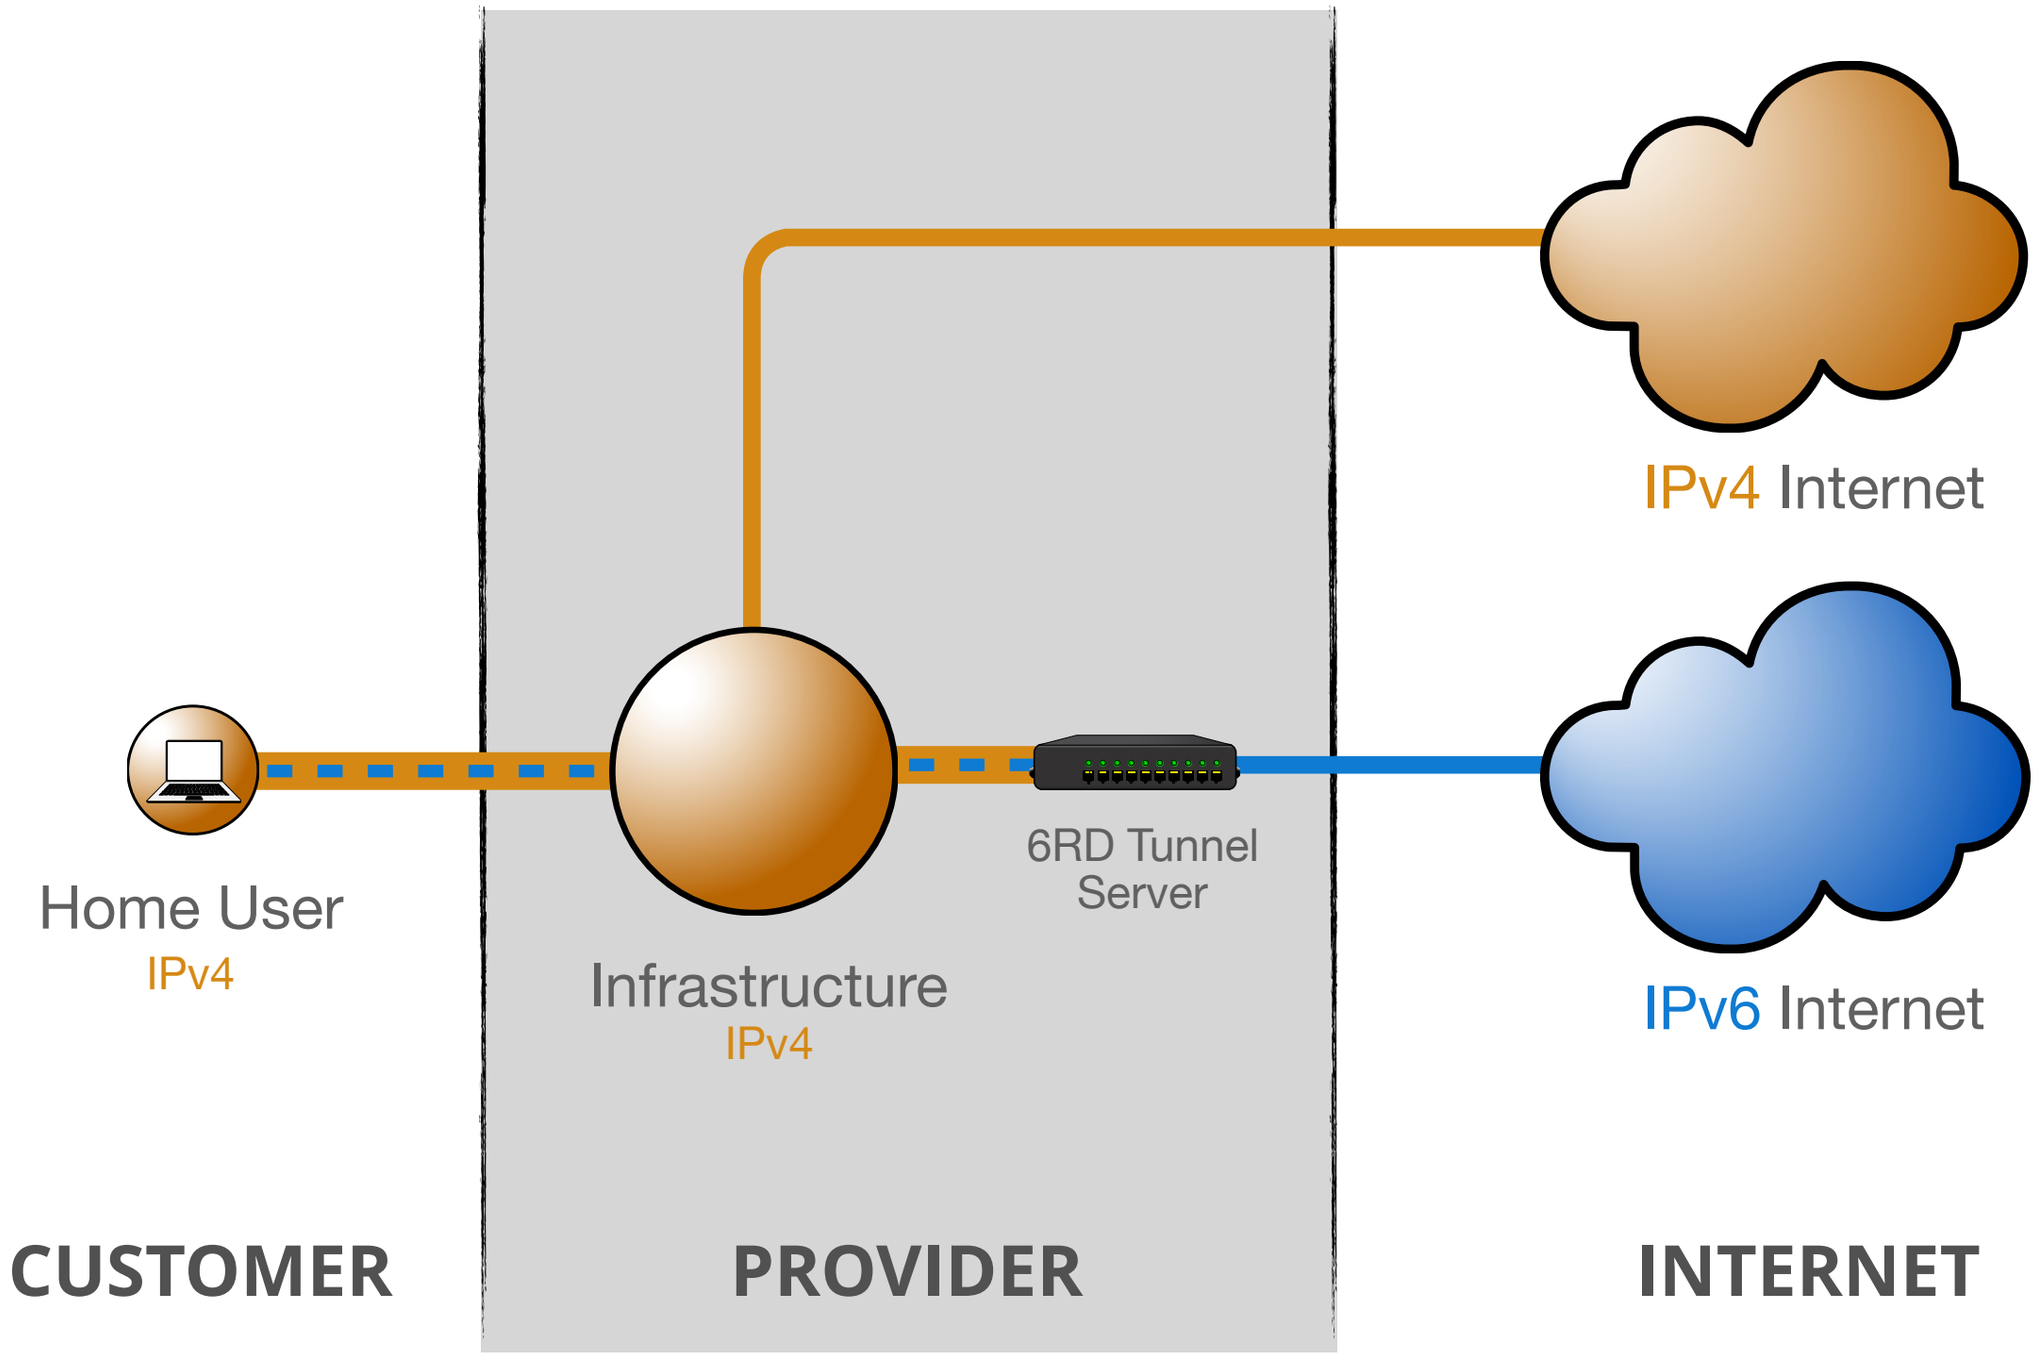
IPv4 Internet
6RD Tunnel Server
IPv6 Internet
Home User
IPv4
Infrastructure
IPv4
CUSTOMER
PROVIDER
INTERNET

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

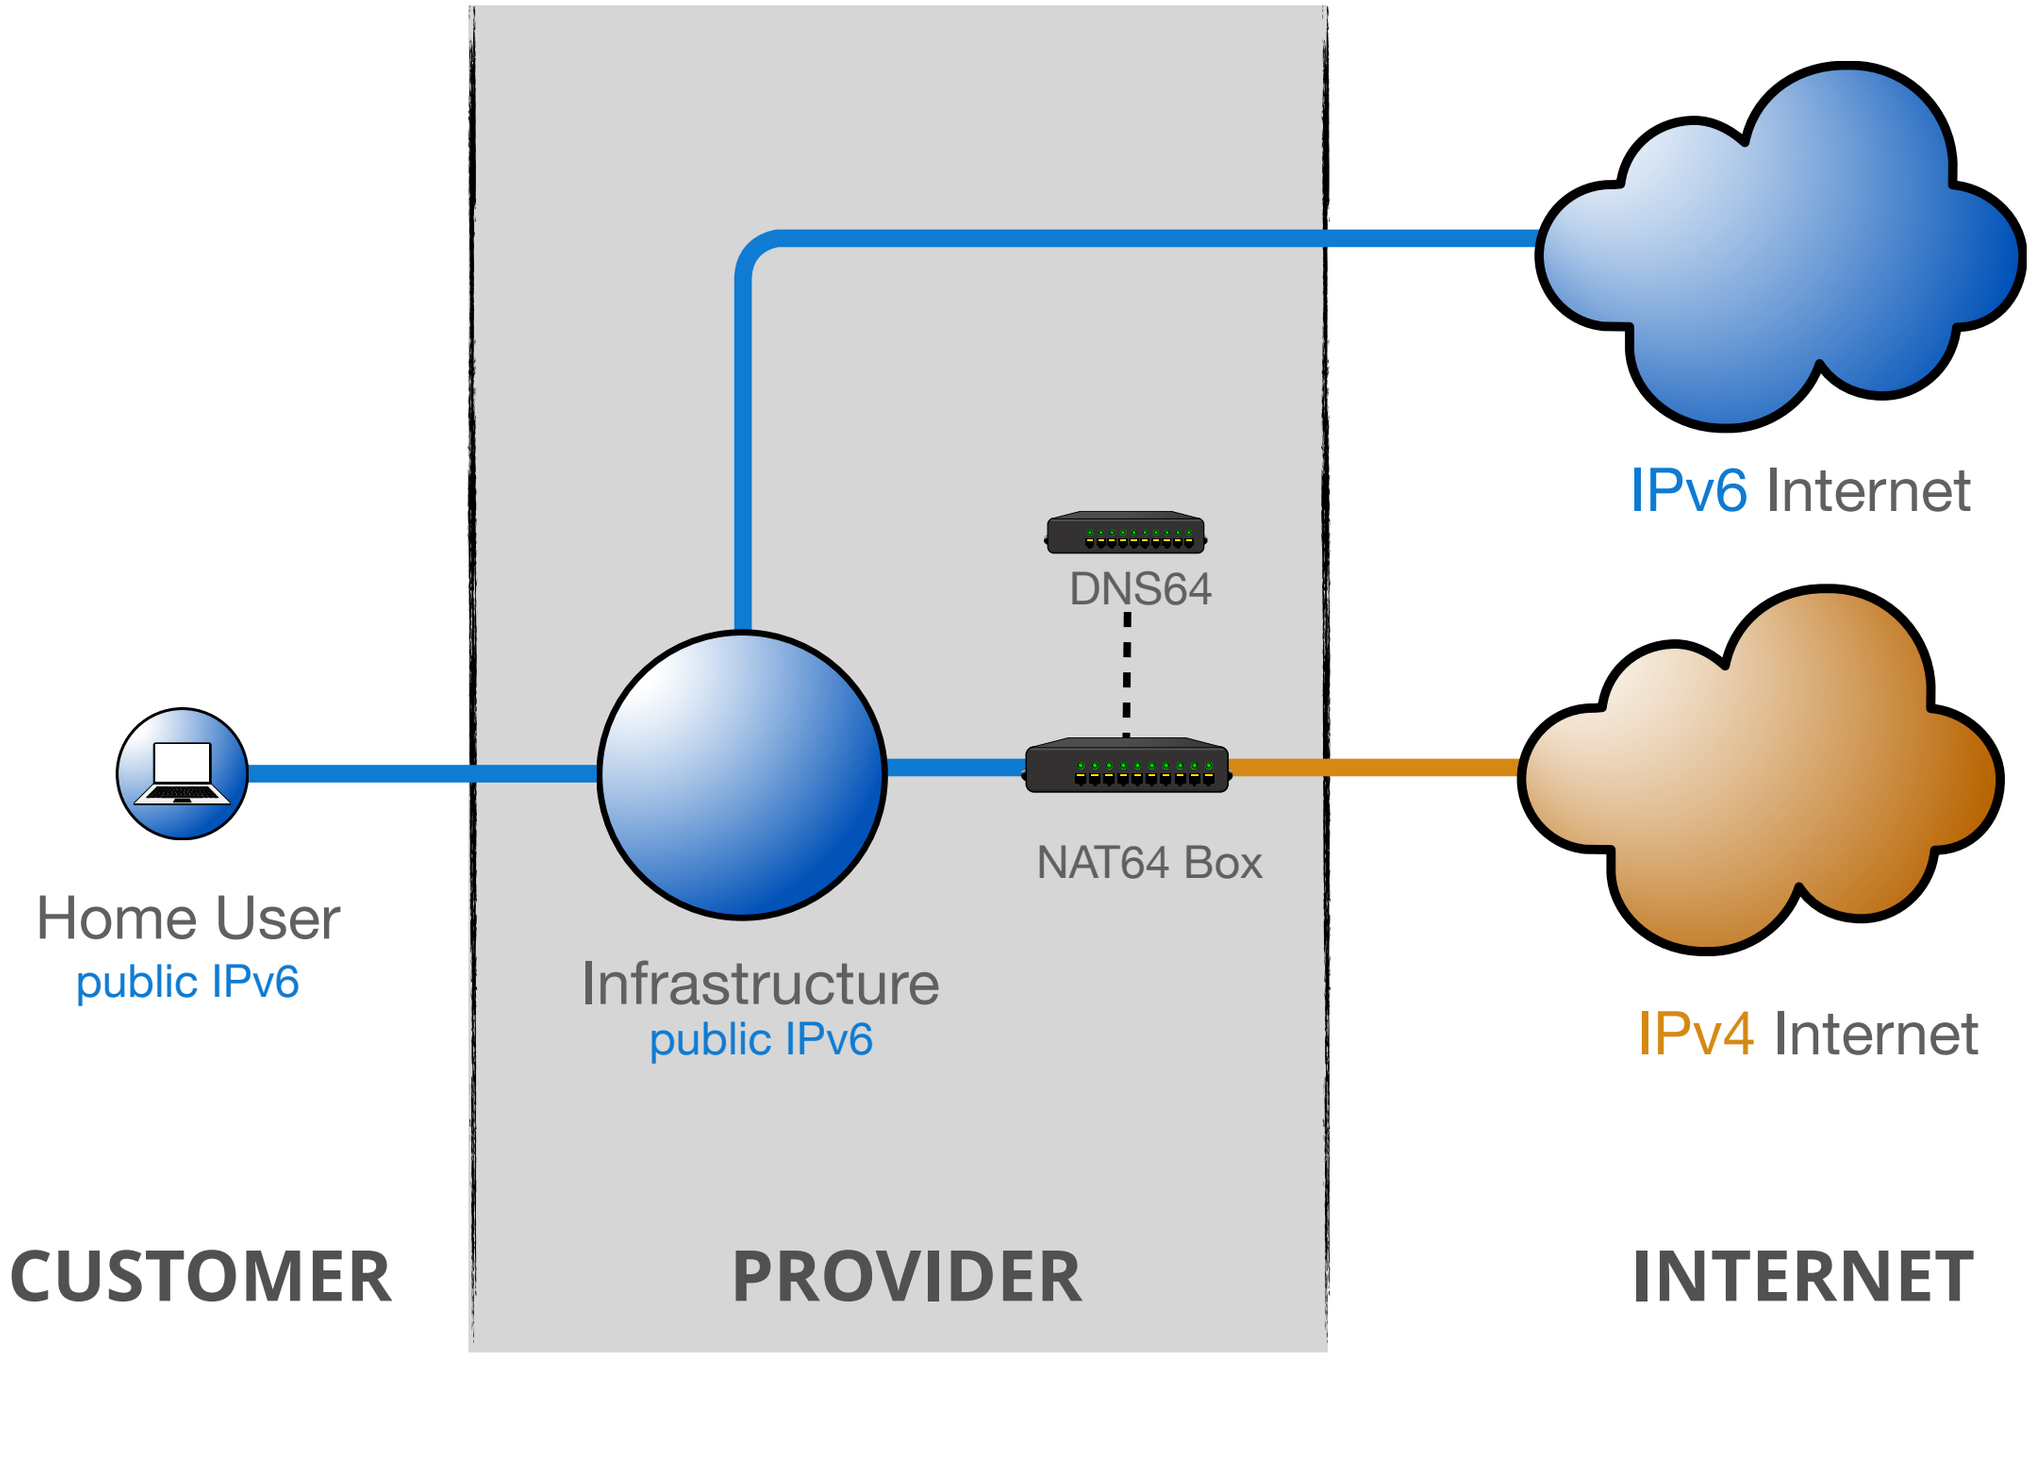
IPv6 Internet
DNS64
NAT64 Box
Home User
public IPv6
Infrastructure
public IPv6
IPv4 Internet
CUSTOMER
PROVIDER
INTERNET

464XLAT

- Extension to NAT64 to access IPv4-only applications (like Skype or Whatsapp)

- Handset pretends there is an IPv4 address (CLAT) and sends IPv4 packets in UDP over IPv6

464XLAT

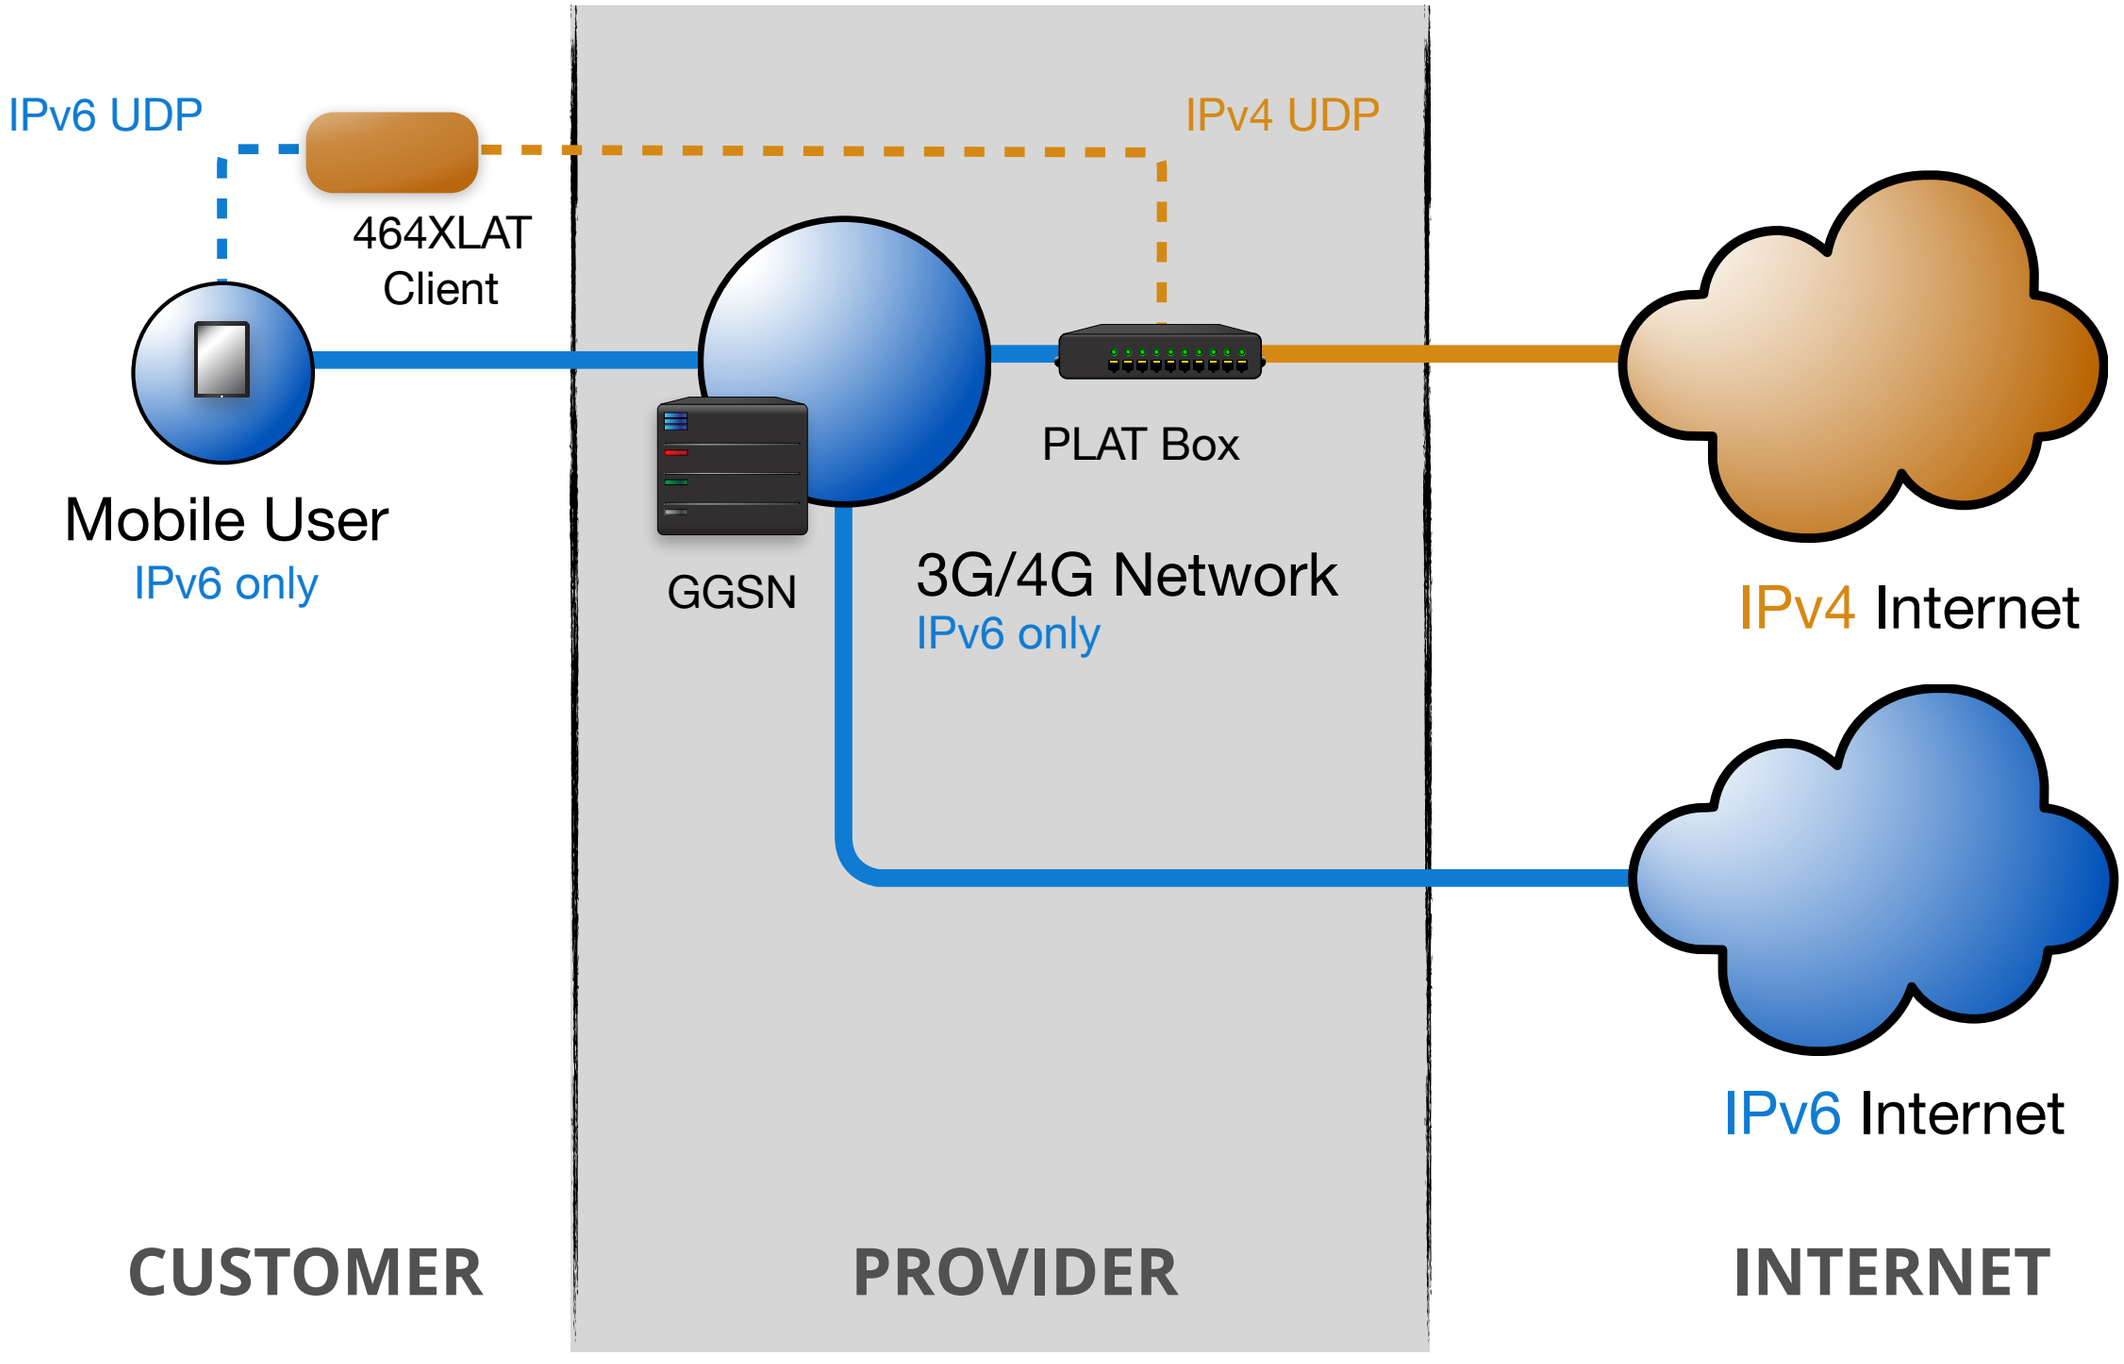
IPv6 UDP
IPv4 UDP
464XLAT
Client
PLAT Box
Mobile User
IPv6 only
GGSN
3G/4G Network
IPv6 only
IPv4 Internet
IPv6 Internet
CUSTOMER
PROVIDER
INTERNET

DS-lite

- Tunnelling 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

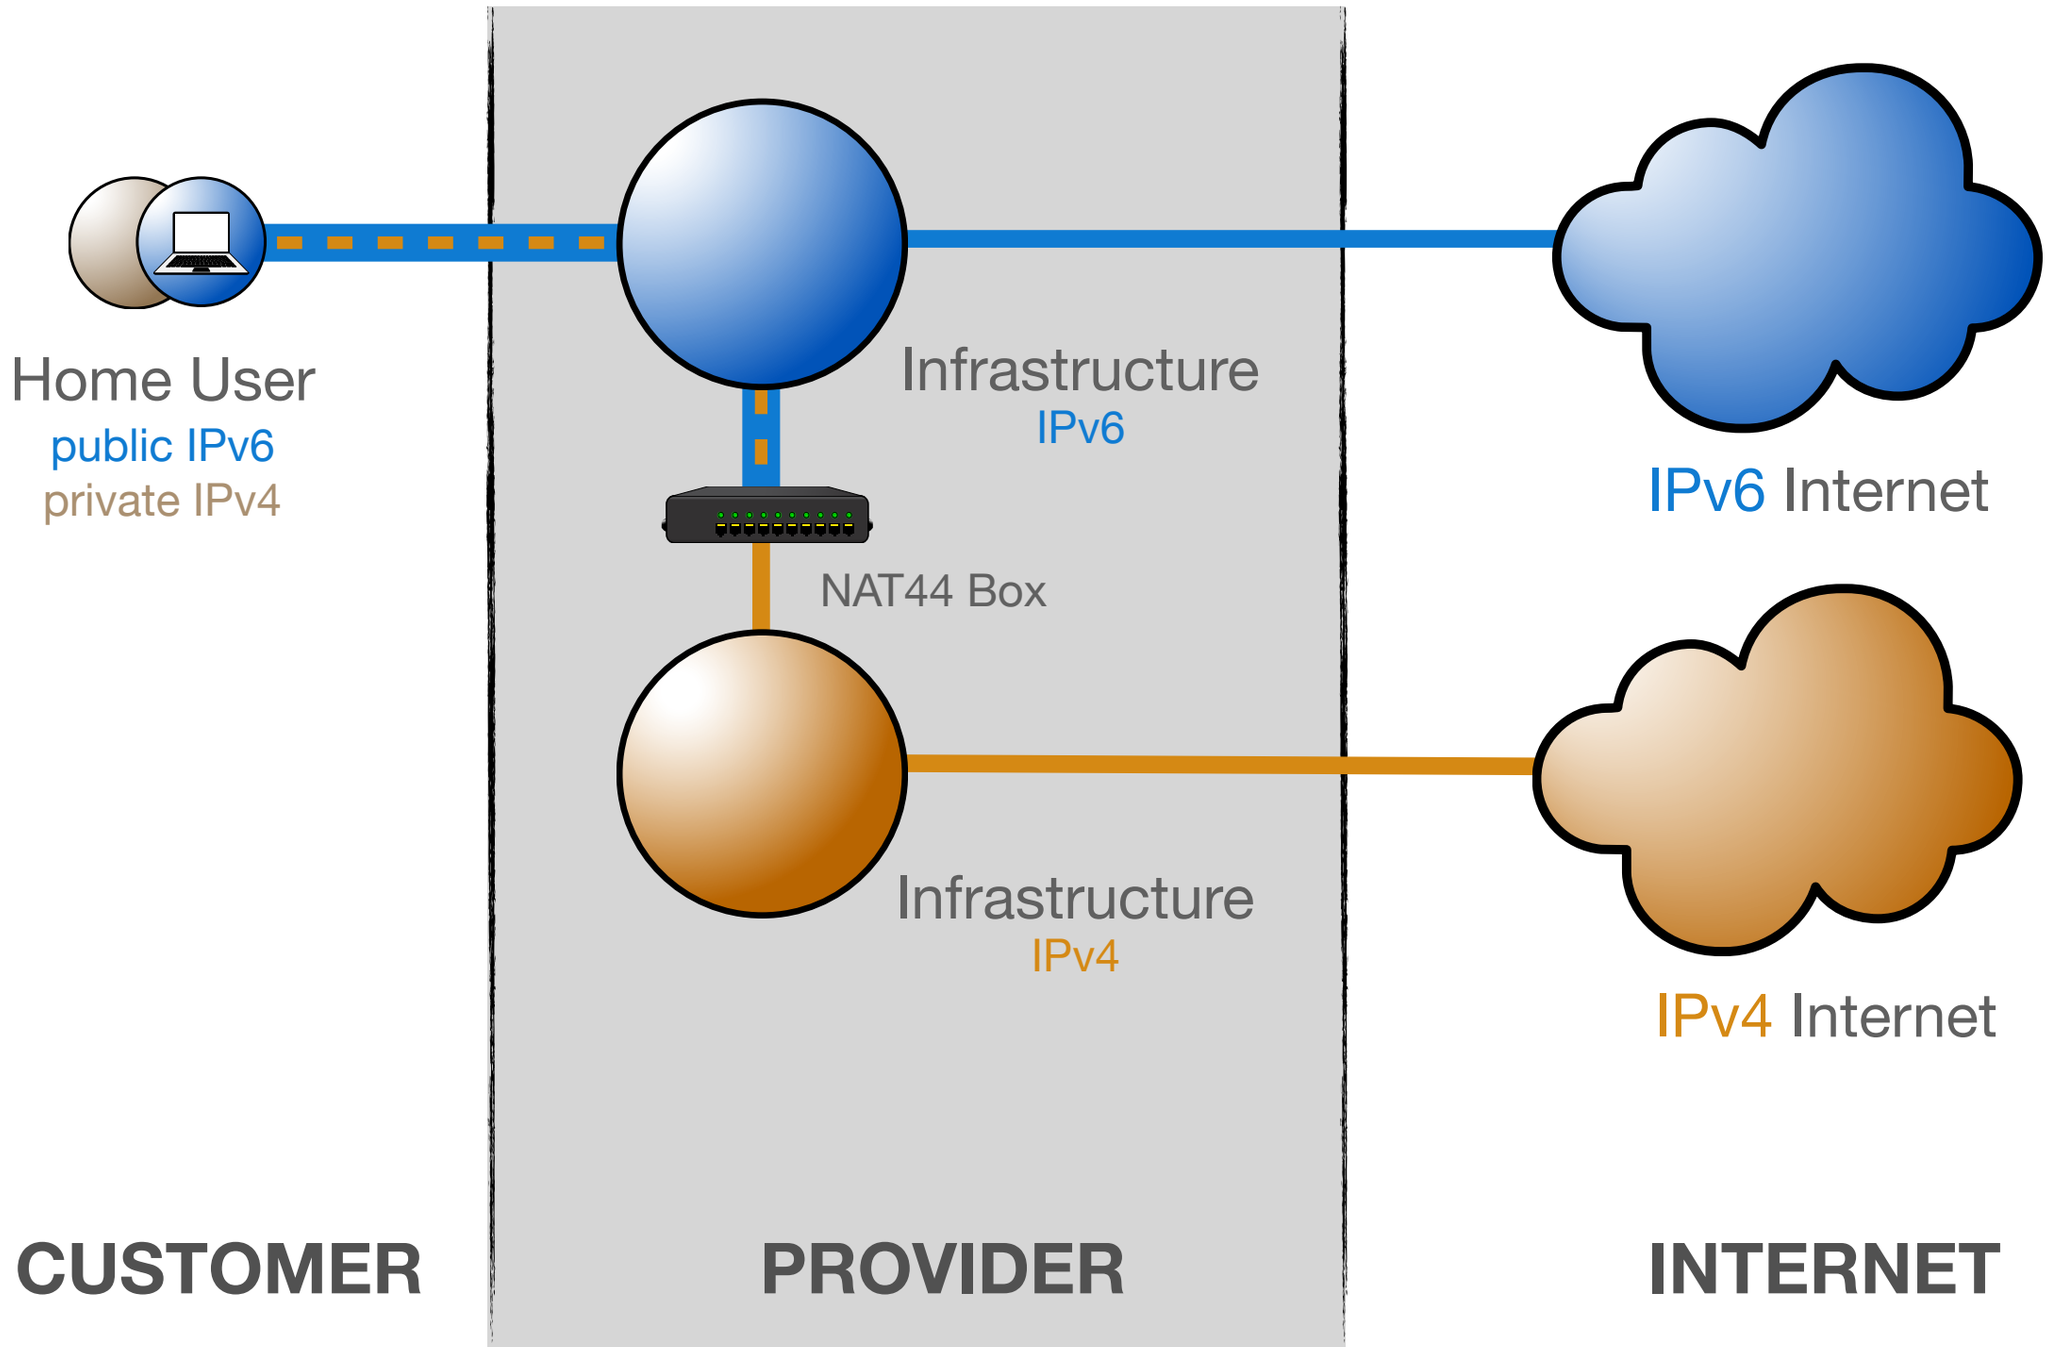
Home User
public IPv6
private IPv4
Infrastructure
IPv6
NAT44 Box
Infrastructure
IPv4
IPv6 Internet
IPv4 Internet
CUSTOMER
PROVIDER
INTERNET

Best transition mechanism?

Dual Stack

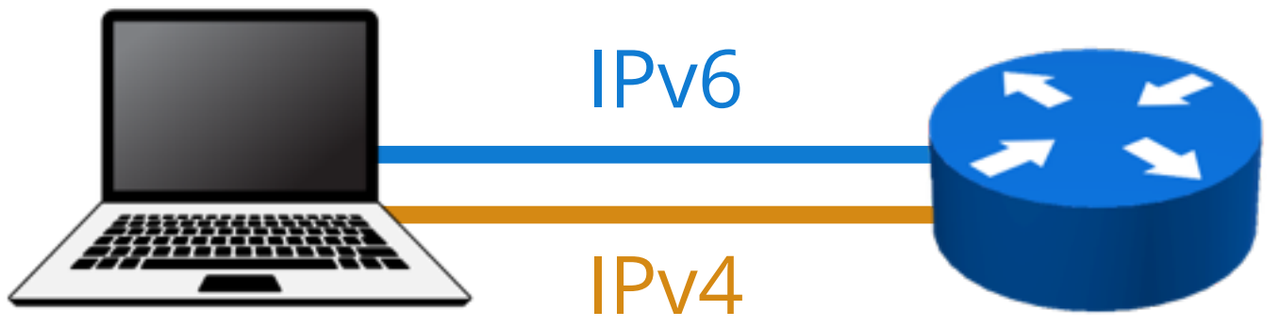

Configuring a 6in4 Tunnel

Exercise

Exercise: Configuring a 6in4 tunnel

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

Login: X

Password: ipv6

Choose “Tunnelling IPv6: 6in4” from the menu

Exercise: Configuring a 6in4 tunnel

- Login to the routers

Login: basicipv6_uX_r1

Password: ipv6

Login: basicipv6_uX_r2

Password: ipv6

Exercise: Configuring a 6in4 tunnel

- On both routers, execute:

```
configure terminal
ipv6 unicast-routing
ipv6 cef
```

Exercise: Configuring a 6in4 tunnel

- On router 1:

```
Interface tunnel0
no ip address
ipv6 address 2001:db8:X::1/64
tunnel source FastEthernet0/0
tunnel mode ipv6ip
tunnel destination 192.168.X.1
```

Exercise: Configuring a 6in4 tunnel

- On router 2:

```
Interface tunnel0
no ip address
ipv6 address 2001:db8:X::2/64
tunnel source FastEthernet0/0
tunnel mode ipv6ip
tunnel destination 10.0.X.1
```

Exercise: Configuring a 6in4 tunnel

- Testing the configuration

ping ipv6 2001:db8:X::1

and

ping ipv6 2001:db8:X::2

Real Life IPv6 Deployment

Section 8

Colocation Provider

- **30 staff**
- **Routing**
  - Dual Stack!
  - Possible IGP combinations were:
    - 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*”

ISP xDSL

- 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

ISP xDSL

- **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)*

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:ae60/64
- **You want your router link at:**
  - 2001:db8:1:1::/64

Large Enterprise

- Approx. 550 IT staff
- Several locations worldwide
- Most of their business processes rely heavily on the Internet
- Driven to IPv6 by need to continue doing business as usual

Large Enterprise

- **Make an inventory of IT needs**
  - Hardware / Software / Services
  - Talk to your ISPs early during preparation
- **Evaluate the current IPv6 offerings**
  - Don’t trust your vendor on “full IPv6 support”
  - Basic network functions are not the issue
  - Check cloud solutions
- **Train your IT staff**
  - Make them understand the WHY of IPv6
  - Focus on the people responsible for applications

Large Enterprise

- **Build a testlab (and start testing!)**
- **Make an IPv6 Roadmap**
  - Dedicated IT group approves roadmap and tracks status
  - “IPv6 Readiness” required for all new purchases
  - Plan replacement of solutions that don’t do IPv6
  - Point out the risks of apps not doing IPv6
- **Phased Approach to Deployment**
  - Phase 1: dual stack all external facing services
  - Phase 2: datacenter and internal network

Deployment Challenges

Discussion

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

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**
  - the Internet is the product!

Troubleshooting for ISP Helpdesks

- **Most ISP connectivity problems are not IPv6 related**
- **Helpdesks can get confused!**
  - IPv6 is new for them
  - They don't have experience with IPv6 issues

- **A generic troubleshooting guide can help!**
- **Based on the open source testipv6.com tool**
- **Customisable**

https://www.ripe.net/ripe/docs/ripe-631

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: 12583 LIRs

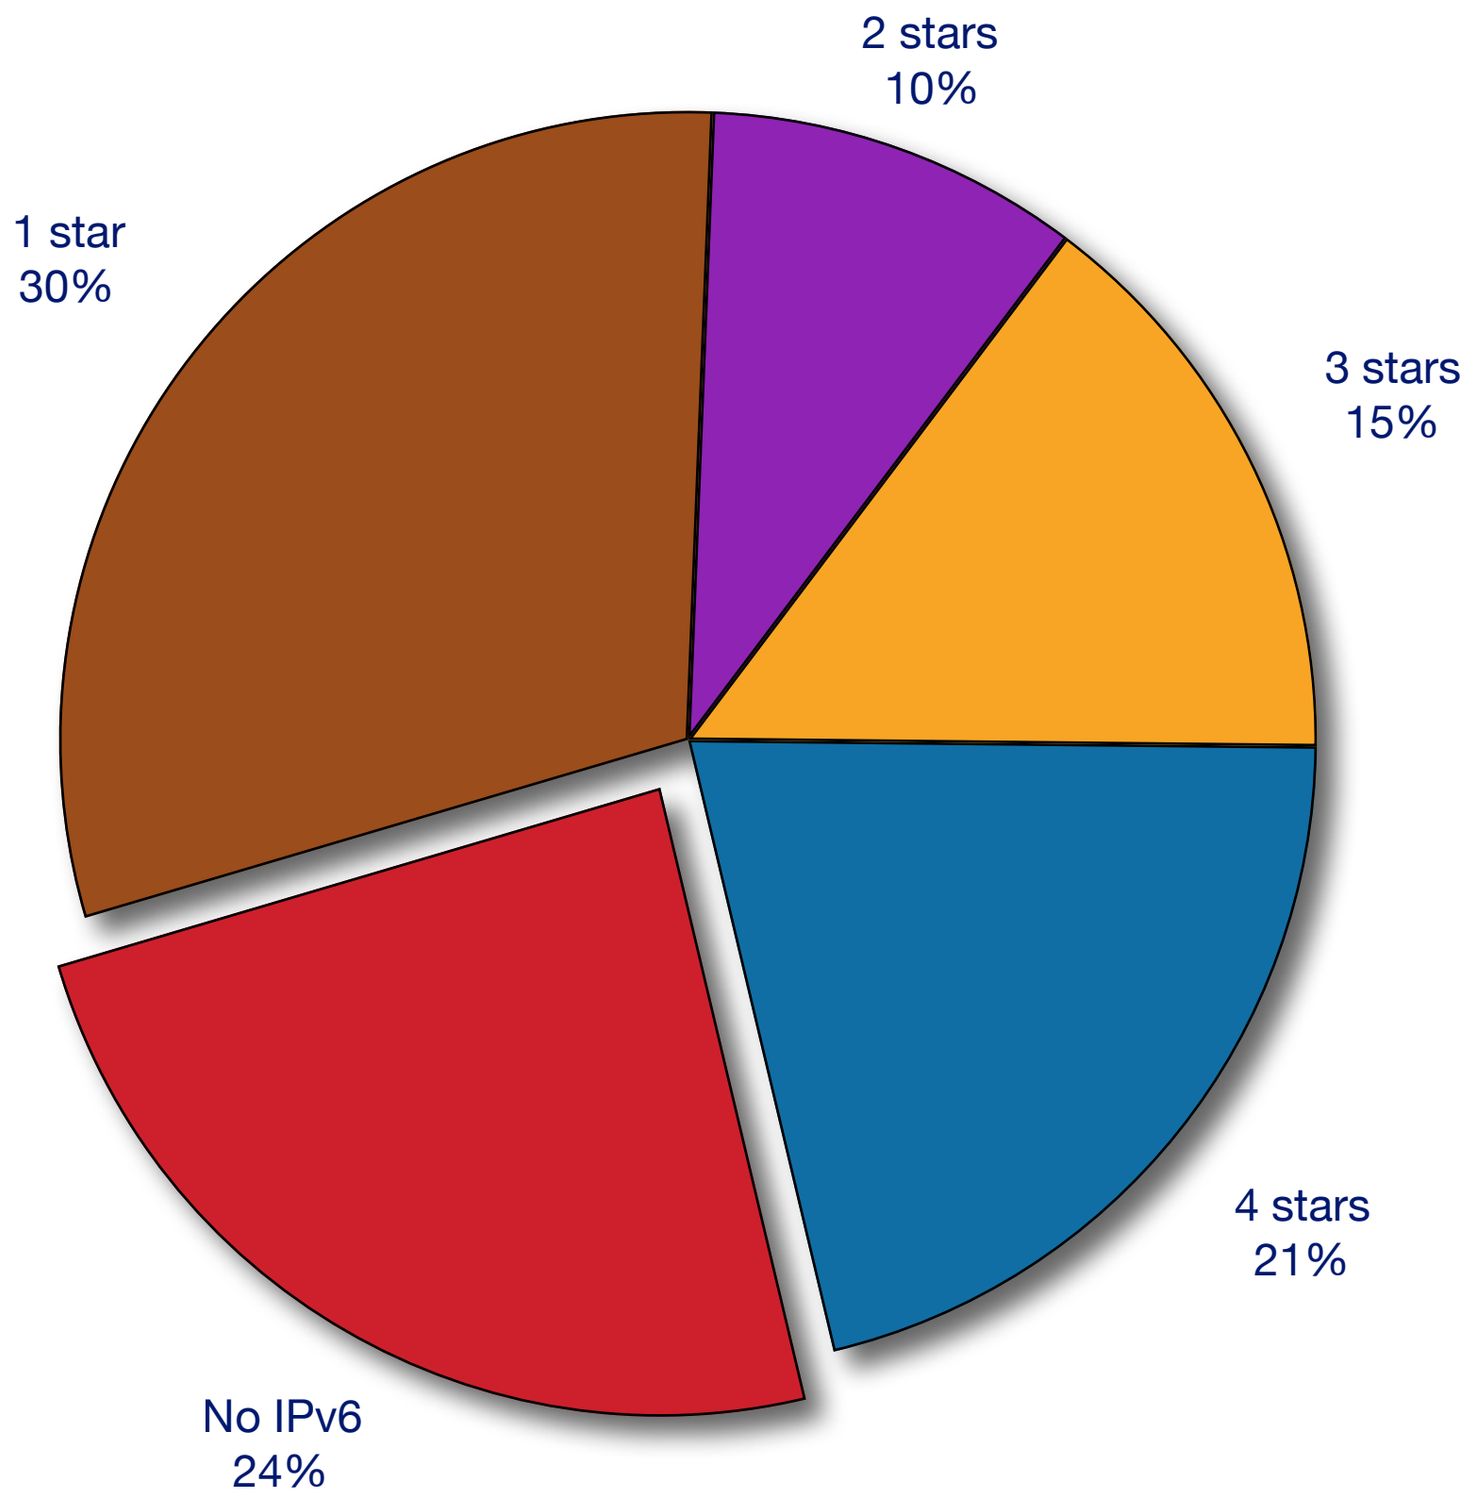

RIPE-554 Document

- **“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)

https://www.ripe.net/ripe/docs/ripe-554

Customers And Their /48

- Customers have no idea how to handle 65536 subnets!
- Provide them with information
  - https://www.ripe.net/support/training/material/basic-ipv6-training-course/Basic-IPv6-Addressing-Plan-Howto.pdf

Also useful

- **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

Questions

RIPE NCC
Academy

Graduate to the next level!

http://academy.ripe.net

Feedback!

https://www.ripe.net/training/basic-ipv6/feedback

Follow us!

@TrainingRIPENCC

The End!

Край

Y Diwedd

Fí

Finis

النهاية

Соңы

Վերջ

Liðugt

Ende

Finvezh

Кінець

Konec

Kraj

Ënn

Fund

پایان

Kpaj

Lõpp

Beigas

Vége

Son

An Críoch

הסוף

Endir

Fine

Sfârşit

Fin

Τέλος

Einde

Конец

Slut

Slutt

დასასრული

Pabaiga

Tmiem

Koniec

Fim

Amaia

Loppu