



RIPE NCC
RIPE NETWORK COORDINATION CENTRE

Basic IPv6 Course

Training Course

March 2023

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



RIPE NCC Training Material

Please find your training material at the following link

<https://www.ripe.net/training-material>





Introductions

- Name
- Experience with IPv6
- Goals



Overview

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



IPv4?

Section 1



Reaching the next billion

- Around 5,385 billion Internet users now
 - around 67.9 % of all people in the world
- Phones, IP Cameras, “Smart” devices / Gateways 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_iiot_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”



IPv4 run-out



“Today, at 15:35 (UTC+1) on 25 November 2019, we made our final /22 IPv4 allocation from the last remaining addresses in our available pool. We have now run out of IPv4 addresses.”





Our Reality: The Waiting List

1. Submit the IPv4 allocation request form at the LIR Portal (/24)
2. Wait

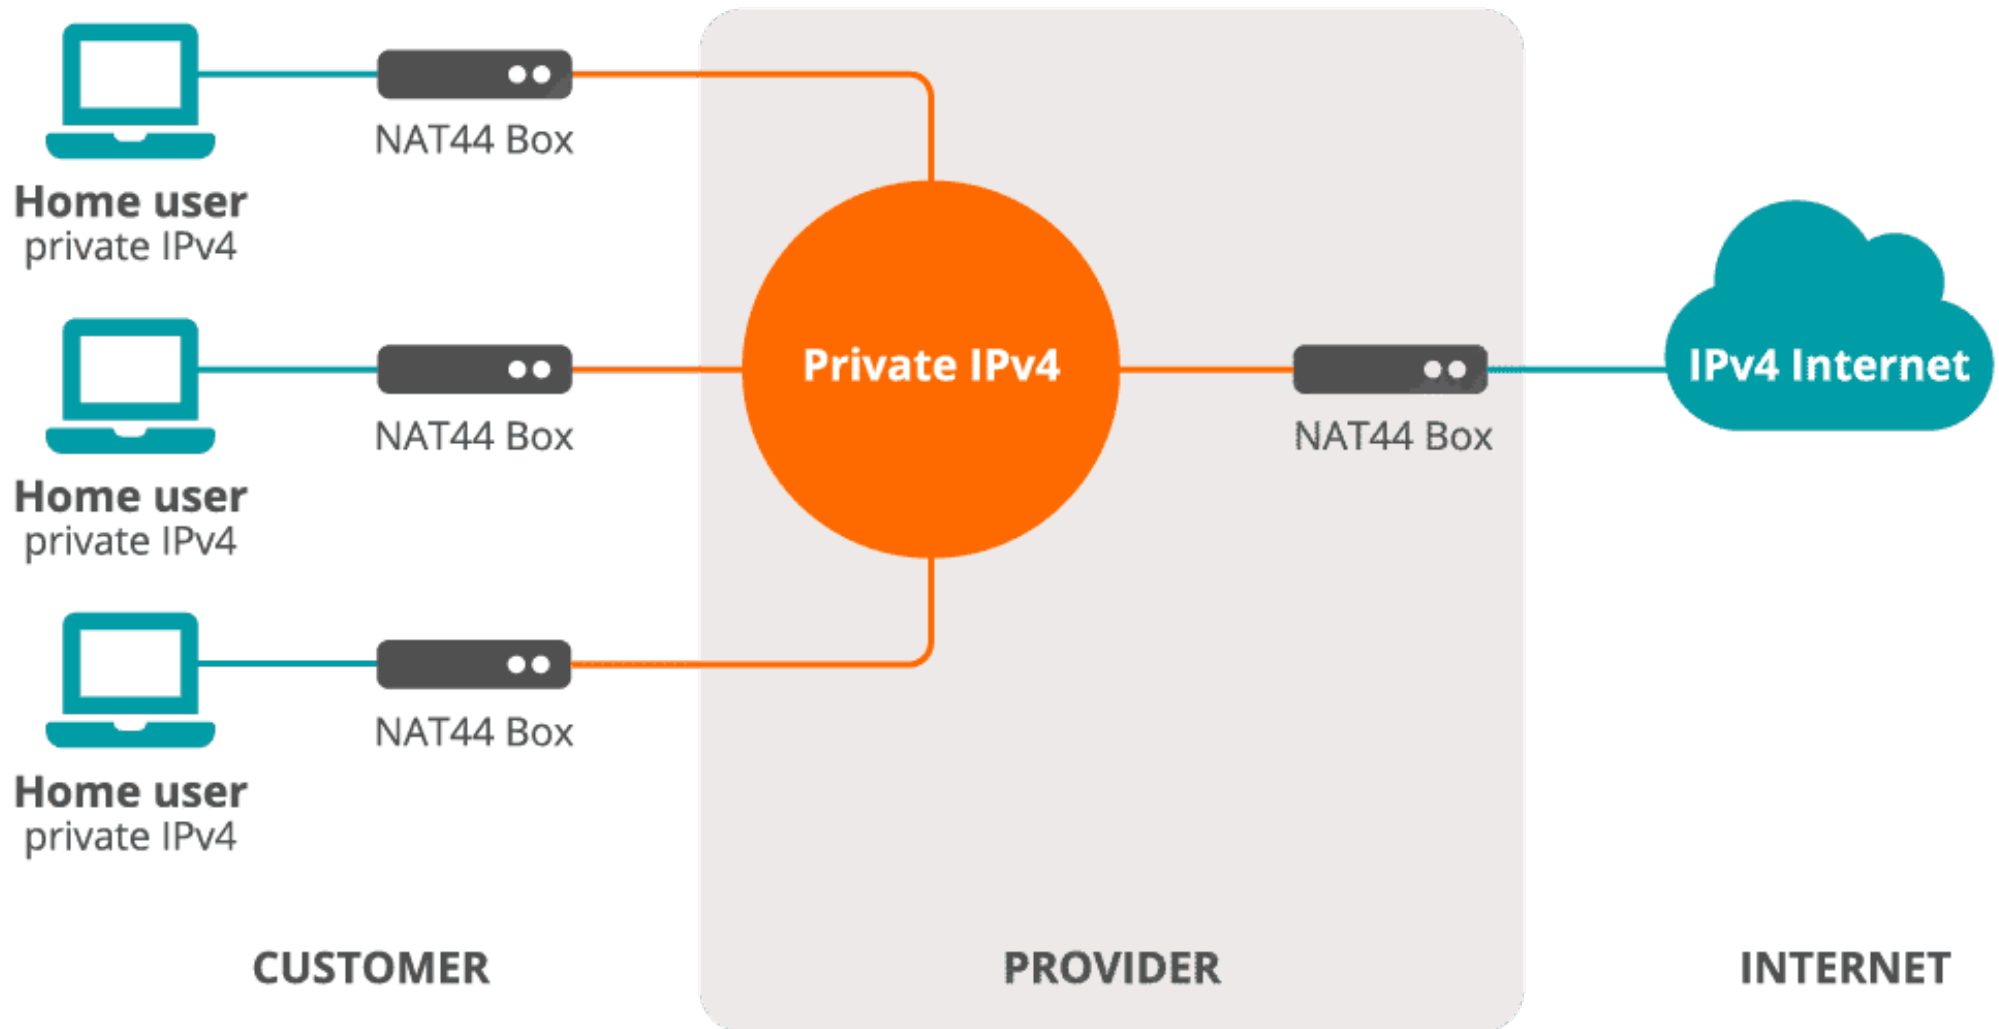




Network Address Translation

- 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



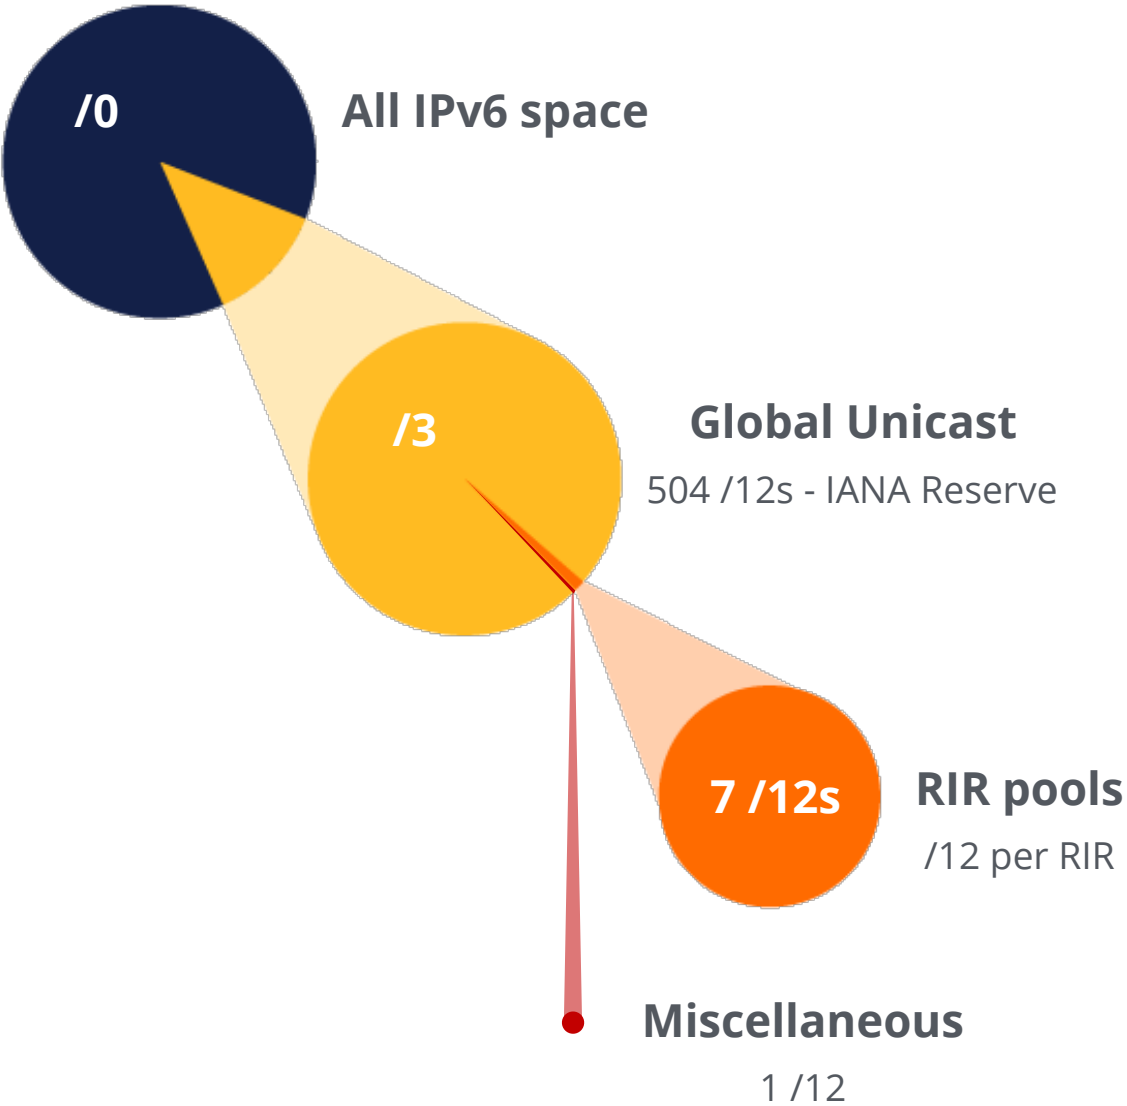


IPv6 Address Basics

Section 2



IP Address Distribution



RIR Pools



October 2006

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

June 2019

RIPE NCC	2A10:0000::/12
-----------------	----------------

November 2019

ARIN	2630:0000::/12
-------------	----------------



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

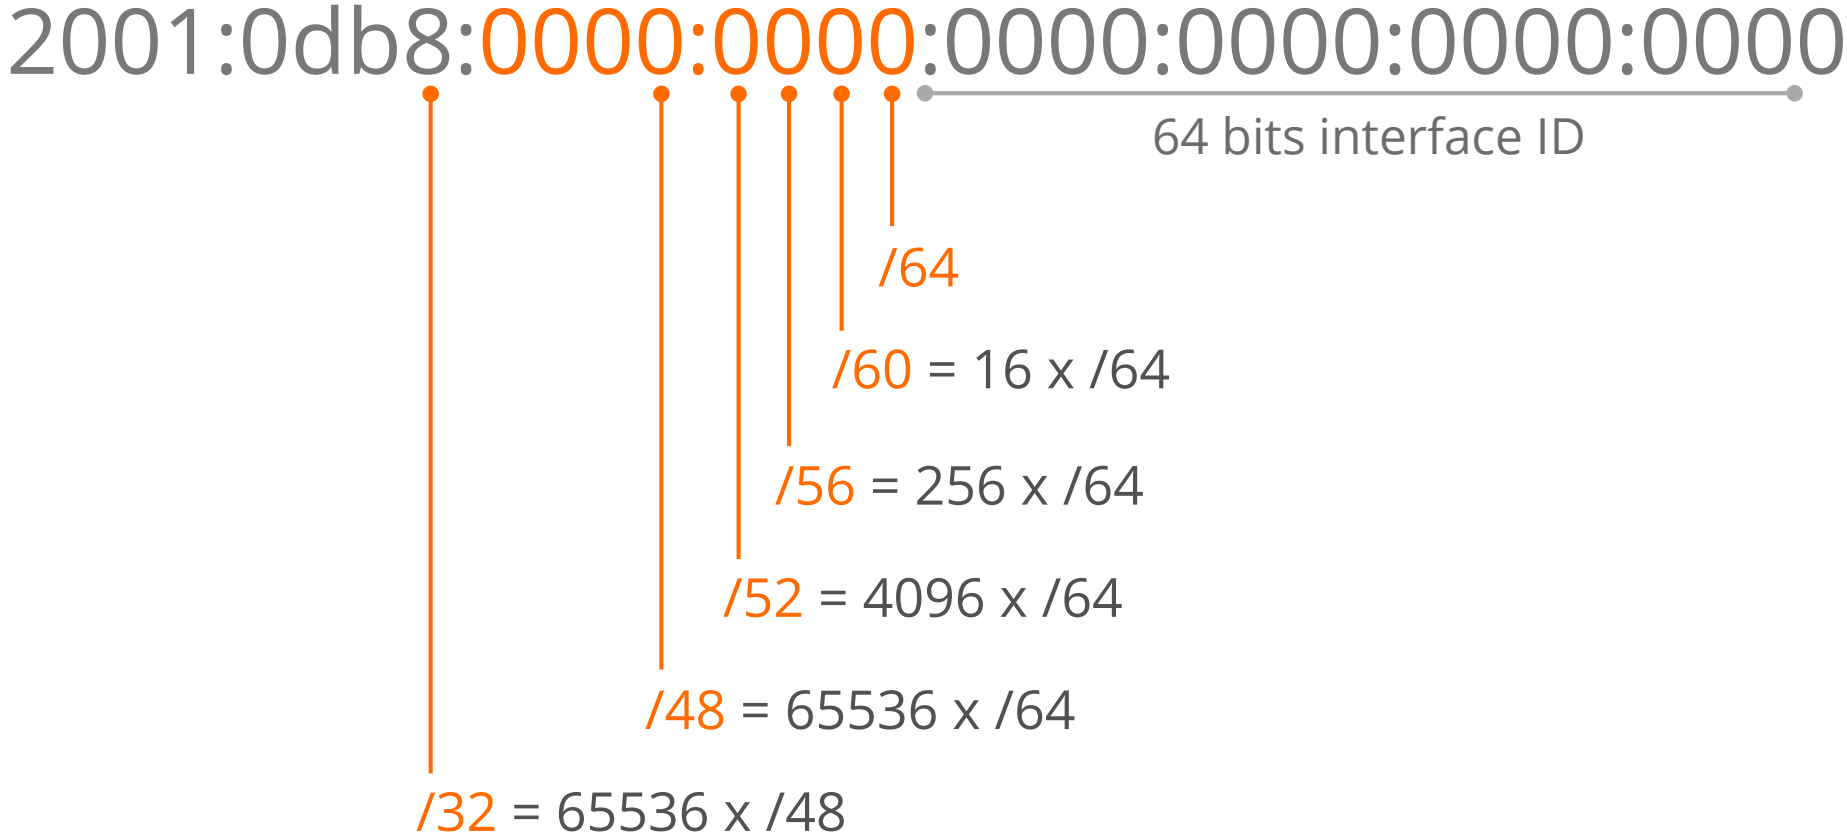
2001:0db8:003e:ef11:0000:0000:c100:004d

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





IPv6 Subnetting

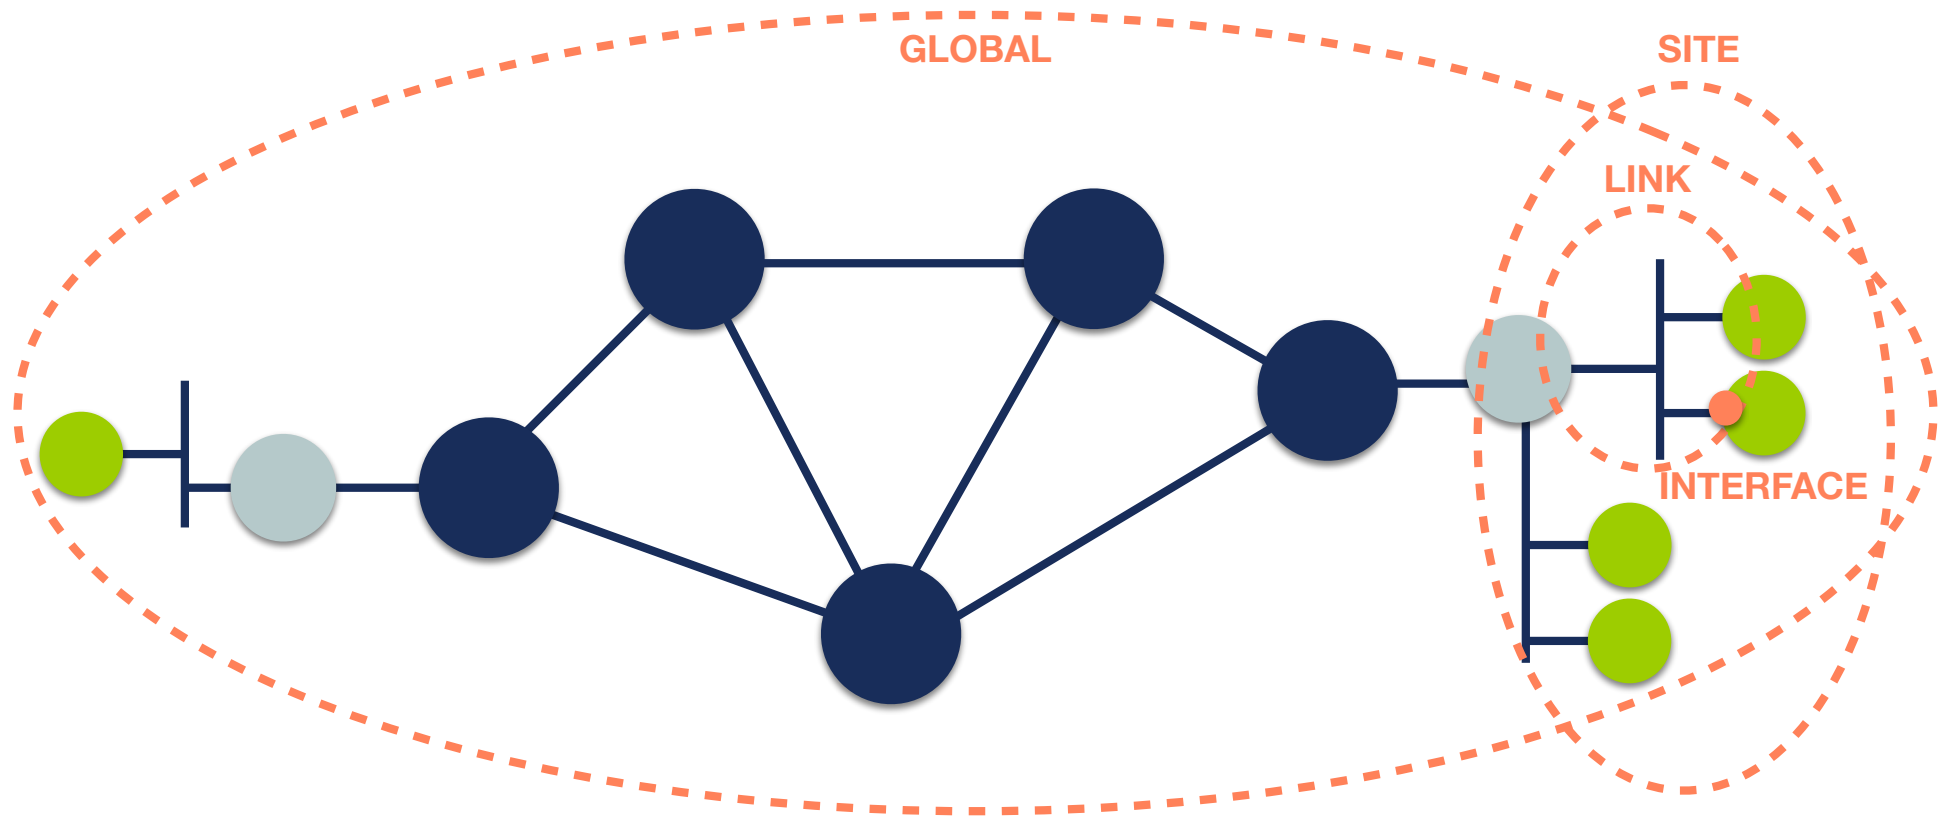




Multiple address types

Addresses	Range	Scope
Unspecified	::/128	n/a
Loopback	::1	host
IPv4-Embedded	64:ff9b::/96	n/a
Discard-Only	100::/64	n/a
Link Local	fe80::/10	link
Global Unicast	2000::/3	global
Unique Local	fc00::/7	global
Multicast	ff00::/8	variable

IPv6 Address Scope



FE80::A:B:100

FF01::2

2001:67C:2E:1::C1

FD00:A:B::100

FF05::1:3

FF02::1



IPv6 Address Notation

Exercise

Question #1

You have a /32 prefix starting with **2001:0db8**.

How do you search for it in the RIPE Database?

- a. 2001:0db8
- b. 2001:0db8/32
- c. 2001:0db8::/32
- d. 2001:db8::/32

Question #1 Answer

You have a /32 prefix starting with **2001:0db8**.

How do you search for it in the RIPE Database?

- a. 2001:0db8
- b. 2001:0db8/32
- c. 2001:0db8::- d. 2001:db8::

Question #2

How do you correctly compress the following IPv6 address:

2001:0db8:0000:0000:0000:0000:0000:0c50

- a. 2001:0db8:0:0:0:0:0:0c50
- b. 2001:0db8::0c50
- c. 2001:db8::c50
- d. 2001:db8::c5

Question #2 Answer

How do you correctly compress the following IPv6 address:

2001:0db8:0000:0000:0000:0000:0000:0c50

- a. 2001:0db8:0:0:0:0:0:0c50
- b. 2001:0db8::0c50
- c. 2001:db8::c50 *
- d. 2001:db8::c5

Question #3

How do you correctly compress the following IPv6 address:

2001:0db8:0000:0000:b450:0000:0000:00b4

- a. 2001:db8::b450::b4
- b. 2001:db8::b450:0:0:b4
- c. 2001:db8::b45:0000:0000:b4
- d. 2001:db8:0:0:b450::b4

Question #3 Answer

How do you correctly compress the following IPv6 address:

2001:0db8:0000:0000:b450:0000:0000:00b4

- a. 2001:db8::b450::b4
- b. 2001:db8::b450:0:0:b4 *
- c. 2001:db8::b45:0000:0000:b4
- d. 2001:db8:0:0:b450::b4

Question #4

How do you correctly compress the following IPv6 address:

2001:0db8:00f0:0000:0000:03d0:0000:00ff

- a. 2001:0db8:00f0::3d0:0:00ff
- b. 2001:db8:f0:0:0:3d0:0:ff
- c. 2001:db8:f0::3d0:0:ff
- d. 2001:0db8:0f0:0:0:3d0:0:0ff

Question #4 Answer

How do you correctly compress the following IPv6 address:

2001:0db8:00f0:0000:0000:03d0:0000:00ff

- a. 2001:0db8:00f0::3d0:0:00ff
- b. 2001:db8:f0:0:0:3d0:0:ff
- c. 2001:db8:f0::3d0:0:ff *
- d. 2001:0db8:0f0:0:0:3d0:0:0ff

Question #5

How do you correctly compress the following IPv6 address:

2001:0db8:0f3c:00d7:7dab:03d0:0000:00ff

- a. 2001:db8:f3c:d7:7dab:3d:0:ff
- b. 2001:db8:f3c:d7:7dab:3d0:0:ff
- c. 2001:db8:f3c:d7:7dab:3d0::ff
- d. 2001:0db8:0f3c:00d7:7dab:03d::00ff

Question #5 Answer

How do you correctly compress the following IPv6 address:

2001:0db8:0f3c:00d7:7dab:03d0:0000:00ff

- a. 2001:db8:f3c:d7:7dab:3d:0:ff
- b. 2001:db8:f3c:d7:7dab:3d0:0:ff *
- c. 2001:db8:f3c:d7:7dab:3d0::ff
- d. 2001:0db8:0f3c:00d7:7dab:03d::00ff

Question #6

How do you access your IPv6 web server at **2001:db8::8080** on port 8080 using a web browser?

- a. `http://2001:db8::8080:8080`
- b. `http://2001:0db8:0000:0000:0000:0000:8080:8080`
- c. `http://[2001:db8::8080]:8080`
- d. You cannot use the IPv6 address, you have to rely on DNS

Question #6 Answer

How do you access your IPv6 web server at

2001:db8::8080 on port 8080 using a web browser?


- a. `http://2001:db8::8080:8080`
- b. `http://2001:0db8:0000:0000:0000:0000:8080:8080`
- c. `http://[2001:db8::8080]:8080`
- d. You cannot use the IPv6 address, you have to rely on DNS



IPv6 Notation - RFC 5952

For more information, please read RFC 5952:

“A Recommendation for IPv6 Address Text Representation”



Link to the RFC:

<https://datatracker.ietf.org/doc/html/rfc5952>



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
- Originally, for more than /48, send in request form
- Every assignment **must be registered** in the RIPE Database



RIPE Policy Proposal 2019-06

- LIR can create assignments larger than /48 **without a request**
- **Will need to justify** it if there is an audit or if LIR requests subsequent allocation

Comparison IPv4 and IPv6 status

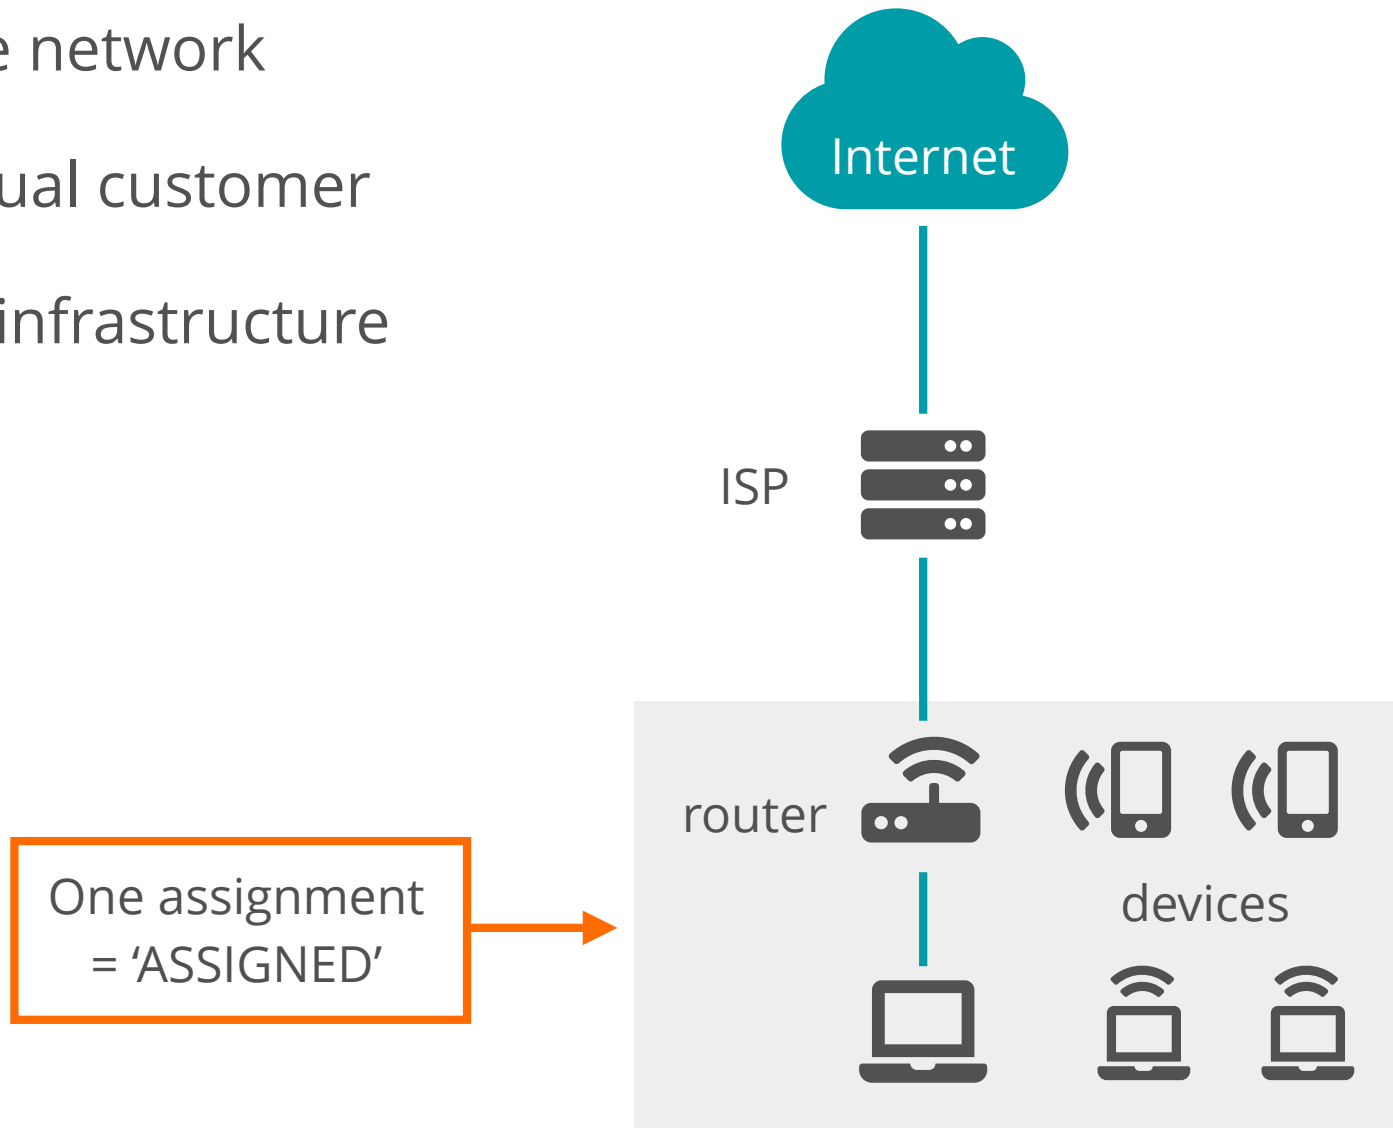


IPv4		IPv6
ALLOCATED PA	Allocation	ALLOCATED-BY-RIR
ASSIGNED PA	Assignment	ASSIGNED
	Group of Assignments	AGGREGATED-BY-LIR
SUB-ALLOCATED PA	Sub-Allocation	ALLOCATED-BY-LIR
ASSIGNED PI	PI Assignment	ASSIGNED PI



Examples ASSIGNED


- One single network
- An individual customer
- Your own infrastructure





Using ASSIGNED

- Represents one assignment
- Minimum assignment size is a /64



Using ASSIGNED - Example Object



inet6num: 2001:db8:1000::/48

netname: CUSTOMER-NET

country: NL

admin-c: ADM321-RIPE

tech-c: NOC123-RIPE

status: ASSIGNED

mnt-by: LIR-MNT


created: 2015-05-31T08:23:35Z

last-modified: 2015-05-31T08:23:35Z



Examples AGGREGATED-BY-LIR

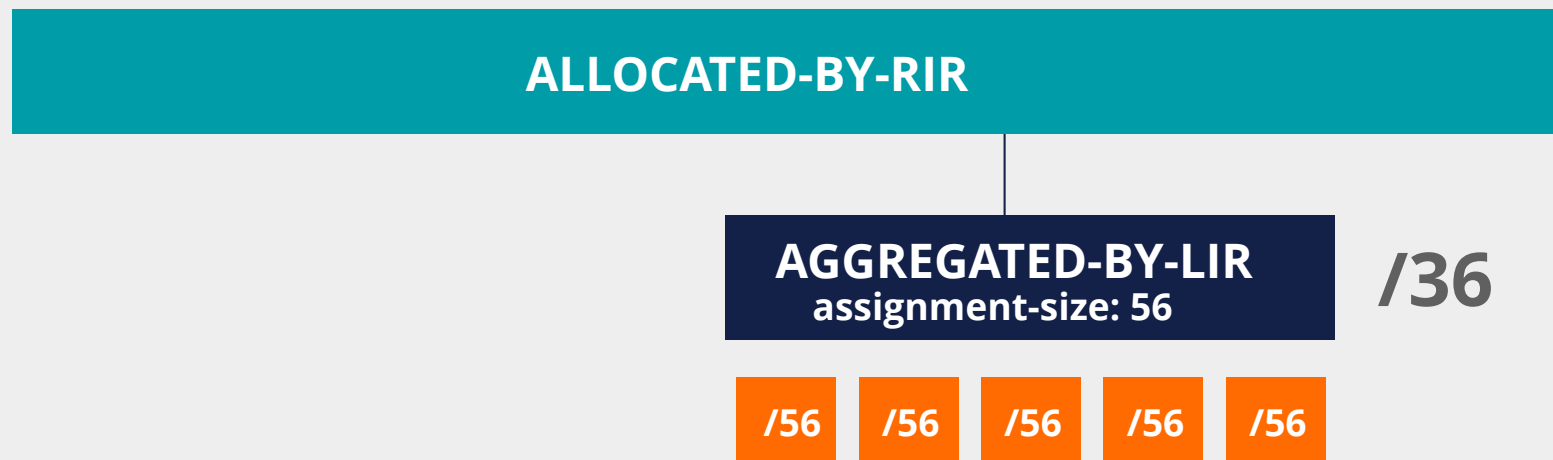
- Group of customers
- Same assignment size





Using AGGREGATED-BY-LIR

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



Using AGGREGATED-BY-LIR - Example



inet6num: 2001:db8:1000::/36

netname: DSL-Broadband-Pool

country: NL

admin-c: ADM321-RIPE

tech-c: NOC123-RIPE

status: AGGREGATED-BY-LIR

assignment-size: 56

mnt-by: LIR-MNT

notify: noc@example.net

created: 2015-05-31T08:23:35Z


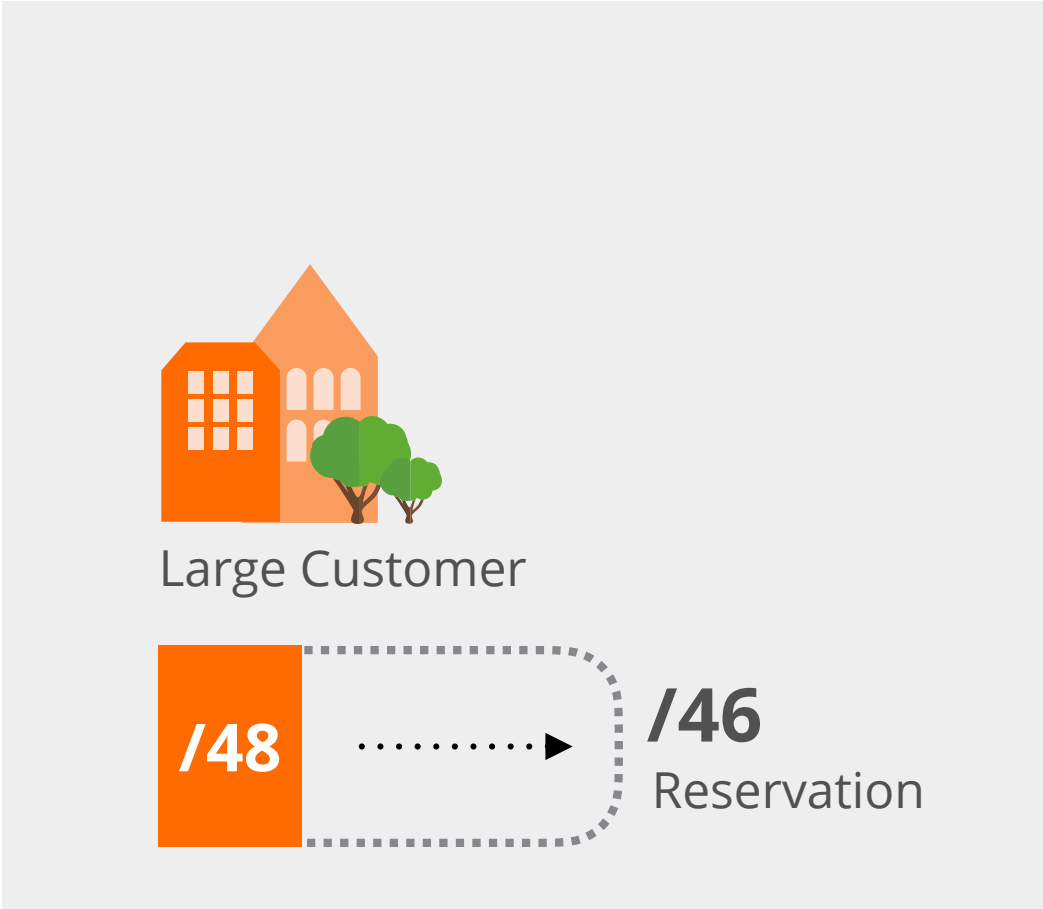
last-modified: 2015-05-31T08:23:35Z

source: RIPE



Examples ALLOCATED-BY-LIR

Reservation for a large customer






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



Using ALLOCATED-BY-LIR - Example



inet6num: 2001:db8:50::/44

netname: Branch-Office-Network

country: NL

admin-c: ADM321-RIPE

tech-c: NOC123-RIPE

status: ALLOCATED-BY-LIR

mnt-by: LIR-MNT

mnt-lower: BRANCH-OFFICE-MNT

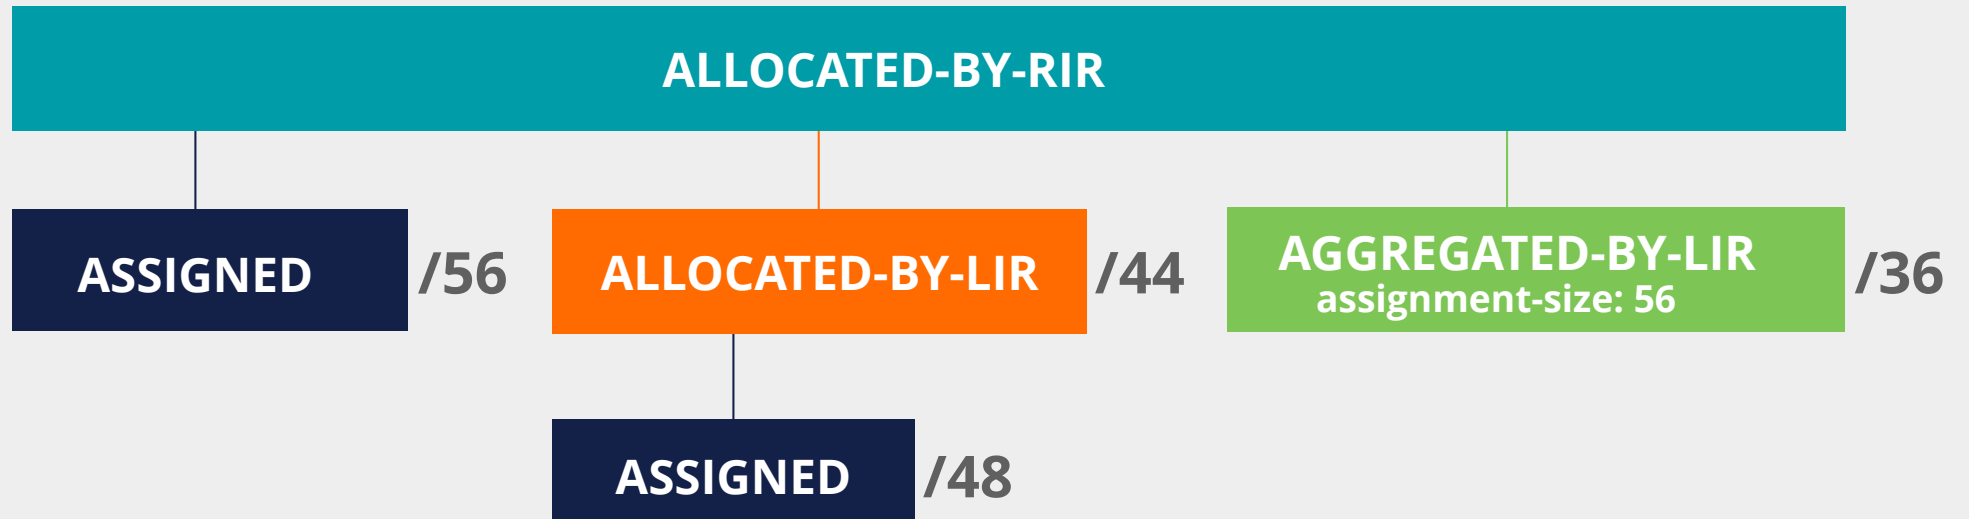
notify: noc@example.net

created: 2015-05-31T08:23:35Z

last-modified: 2015-05-31T08:23:35Z

source: RIPE

Overview





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 **cannot** be used for sub-assignments



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`



Making Assignments

Exercise

Create assignments for a smart city!





Context

- You work for the LIR: **nl.ripencc-ts**
- Your LIR has a /32 allocation: **2001:db8::/32**
- Your customer Future Casa is working on a project called “Smart Home 6”
- They need IPv6 addresses from your address space
- Future Casa wants to connect **1 million** Smart Homes




Product Description

- Each home will be equipped with a 4G-enabled base unit
- The base unit will be the central gateway for smart services inside the house
- Each smart service runs on a **dedicated subnet**
- Services can be enabled or disabled at any point from a user's smartphone app
- Future Casa will be rolling out **new services in the future**

Smart Home 6 Network Diagram



IPv6 Internet





Calculations...

- **/64 = 1 subnet**
 - Not enough. We need one subnet alone for the p2p conn.
- **/63 = 2 subnets**
 - Not enough subnets.
 - Not on the 4-bit boundary!
- **/60 = 16 subnets**
 - Is it enough to meet the future needs?
 - You want to avoid having to renumber!



Calculations...

- **/56 = 256 subnets**

- Sounds reasonable. How many subnets can a house need?

- **/52 = 4096 subnets**

- More than enough.

- **/48 = 65K subnets**

- Definitely more than enough.



Calculations...

One million smart homes

X

/56 per home

=

/36



Possible options for /36 subnets

2001:db8::/32

/32

/36 /36 /36 /36 /36 /36 /36 /36 /36 /36 /36 /36 /36 /36 /36

/36

2001:db8:0000::/36

2001:db8:1000::/36

2001:db8:2000::/36

2001:db8:3000::/36

2001:db8:4000::/36

2001:db8:5000::/36

2001:db8:6000::/36

2001:db8:7000::/36

2001:db8:8000::/36

2001:db8:9000::/36

2001:db8:a000::/36

2001:db8:b000::/36

2001:db8:c000::/36

2001:db8:d000::/36

2001:db8:e000::/36

2001:db8:f000::/36

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

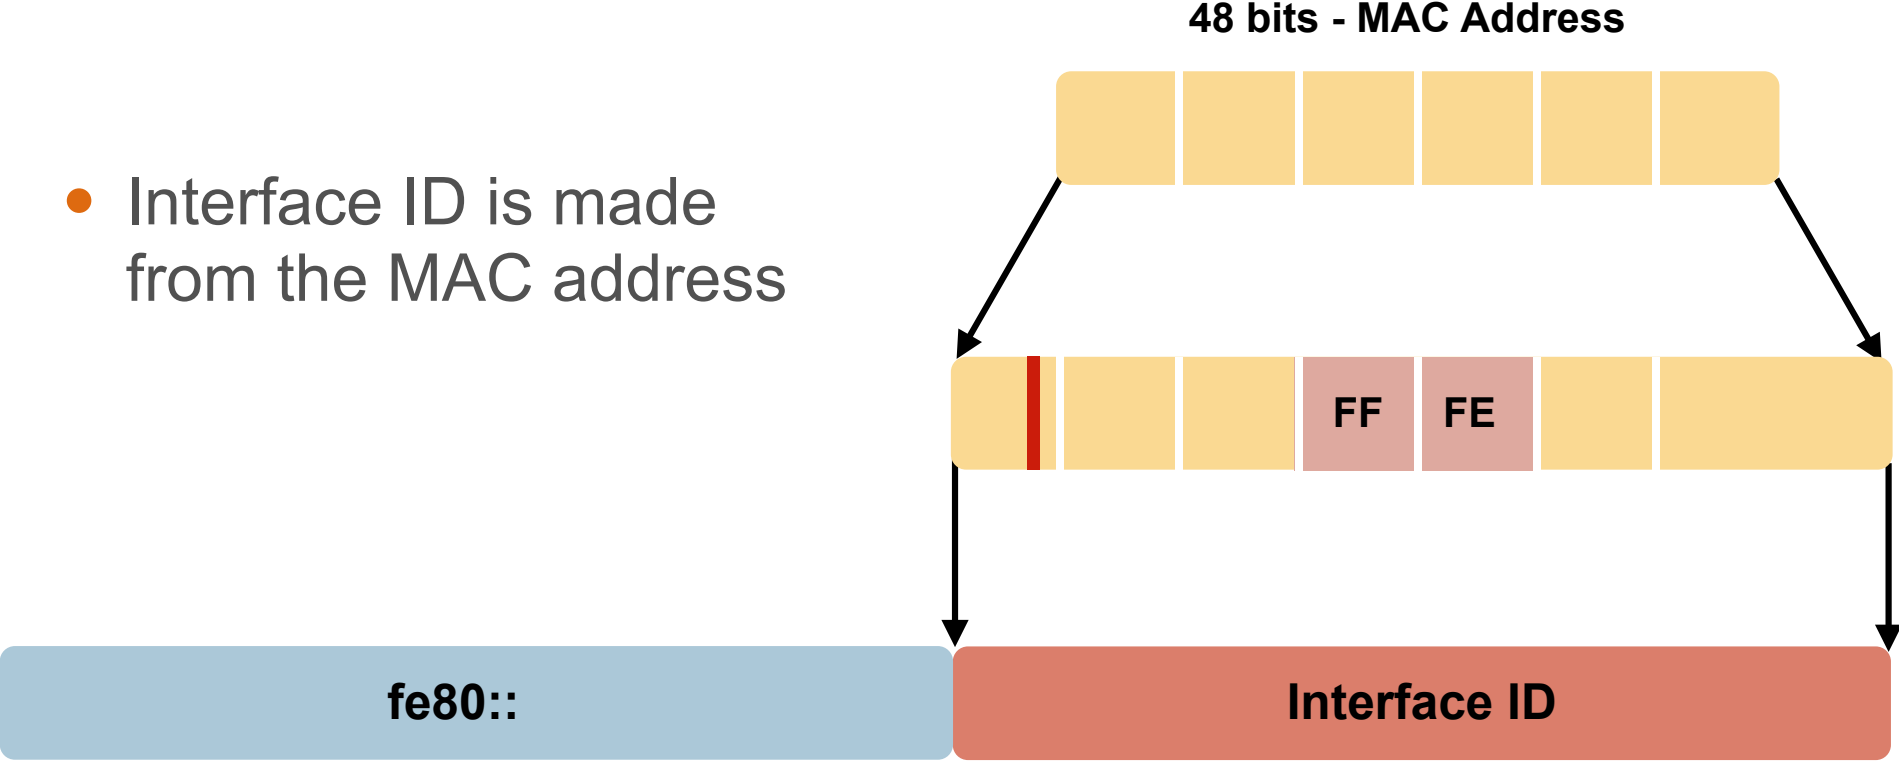
1. Make a Link-Local address
2. Check for duplicates on the link
3. Search for a router
4. Make a Global Unicast address





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

Solicited Node Multicast Address



- Used in Neighbor Discovery Protocol for obtaining the layer 2 link-layer (MAC) addresses

IPv6 unicast address



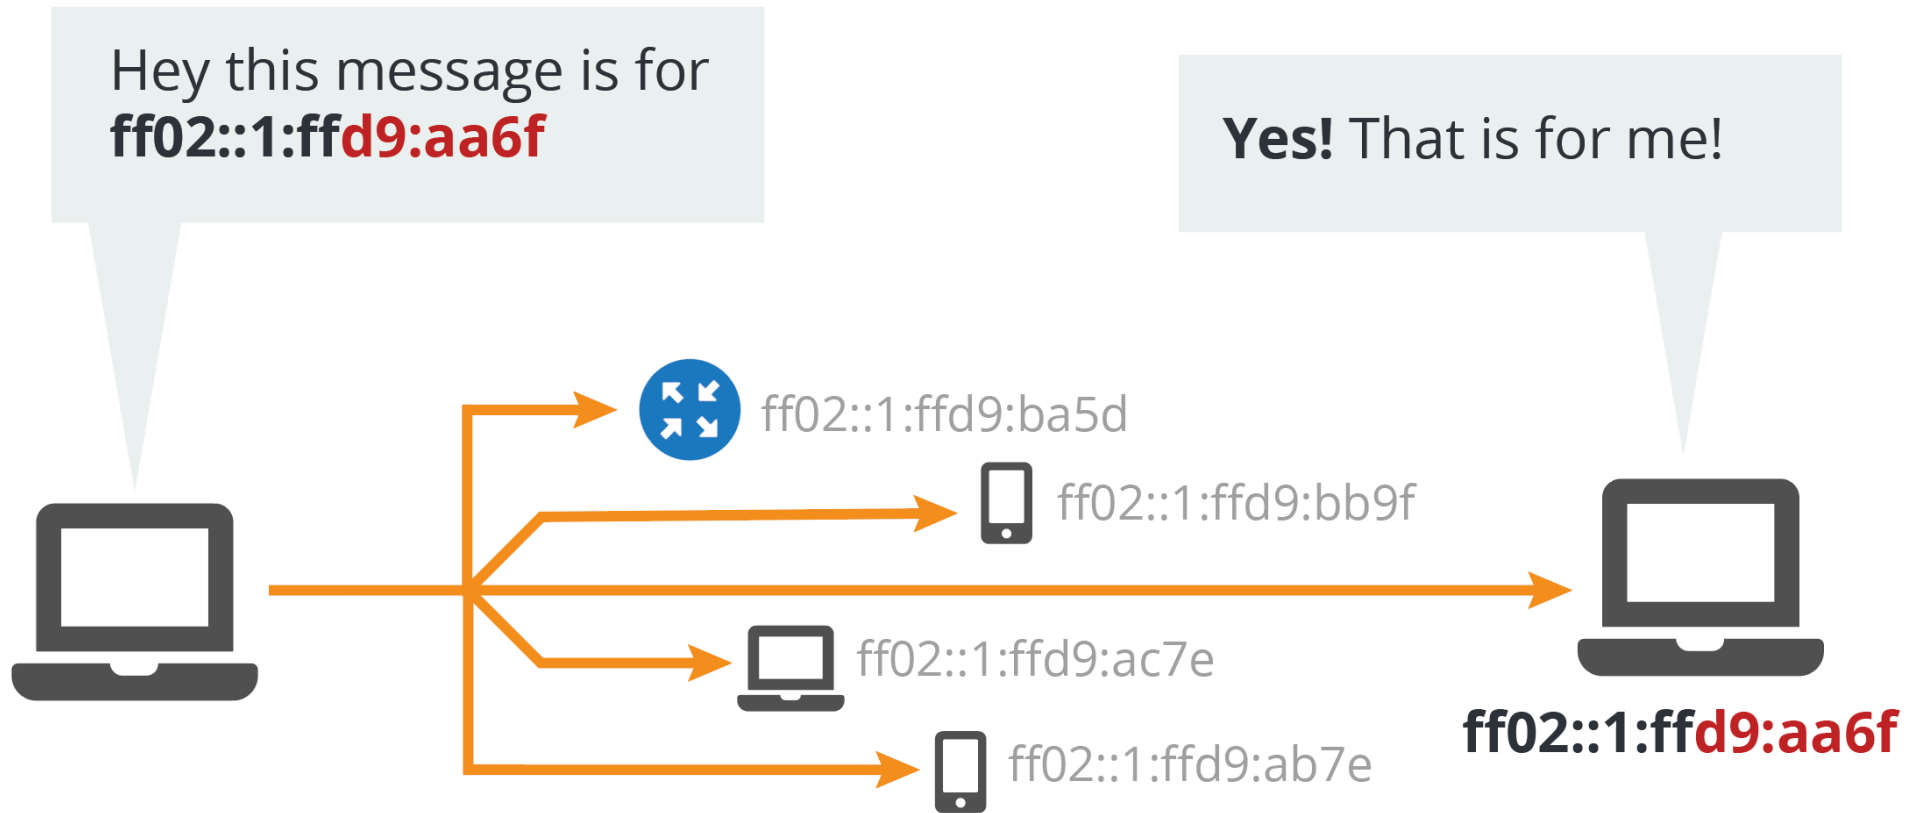
Solicited-node multicast address



128 bits



Solicited Node Multicast 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 Addressing Plan?



- **Benefits of 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 /64s in a /48
 - There are 65.536 /48s in a /32
 - There are 524.288 /48s in a /29
 - There are **16.777.216** /56s in a /32
 - There are **134.217.728** /56s in a /29
- Find a suitable IPAM solution



Addressing Plan

Exercise

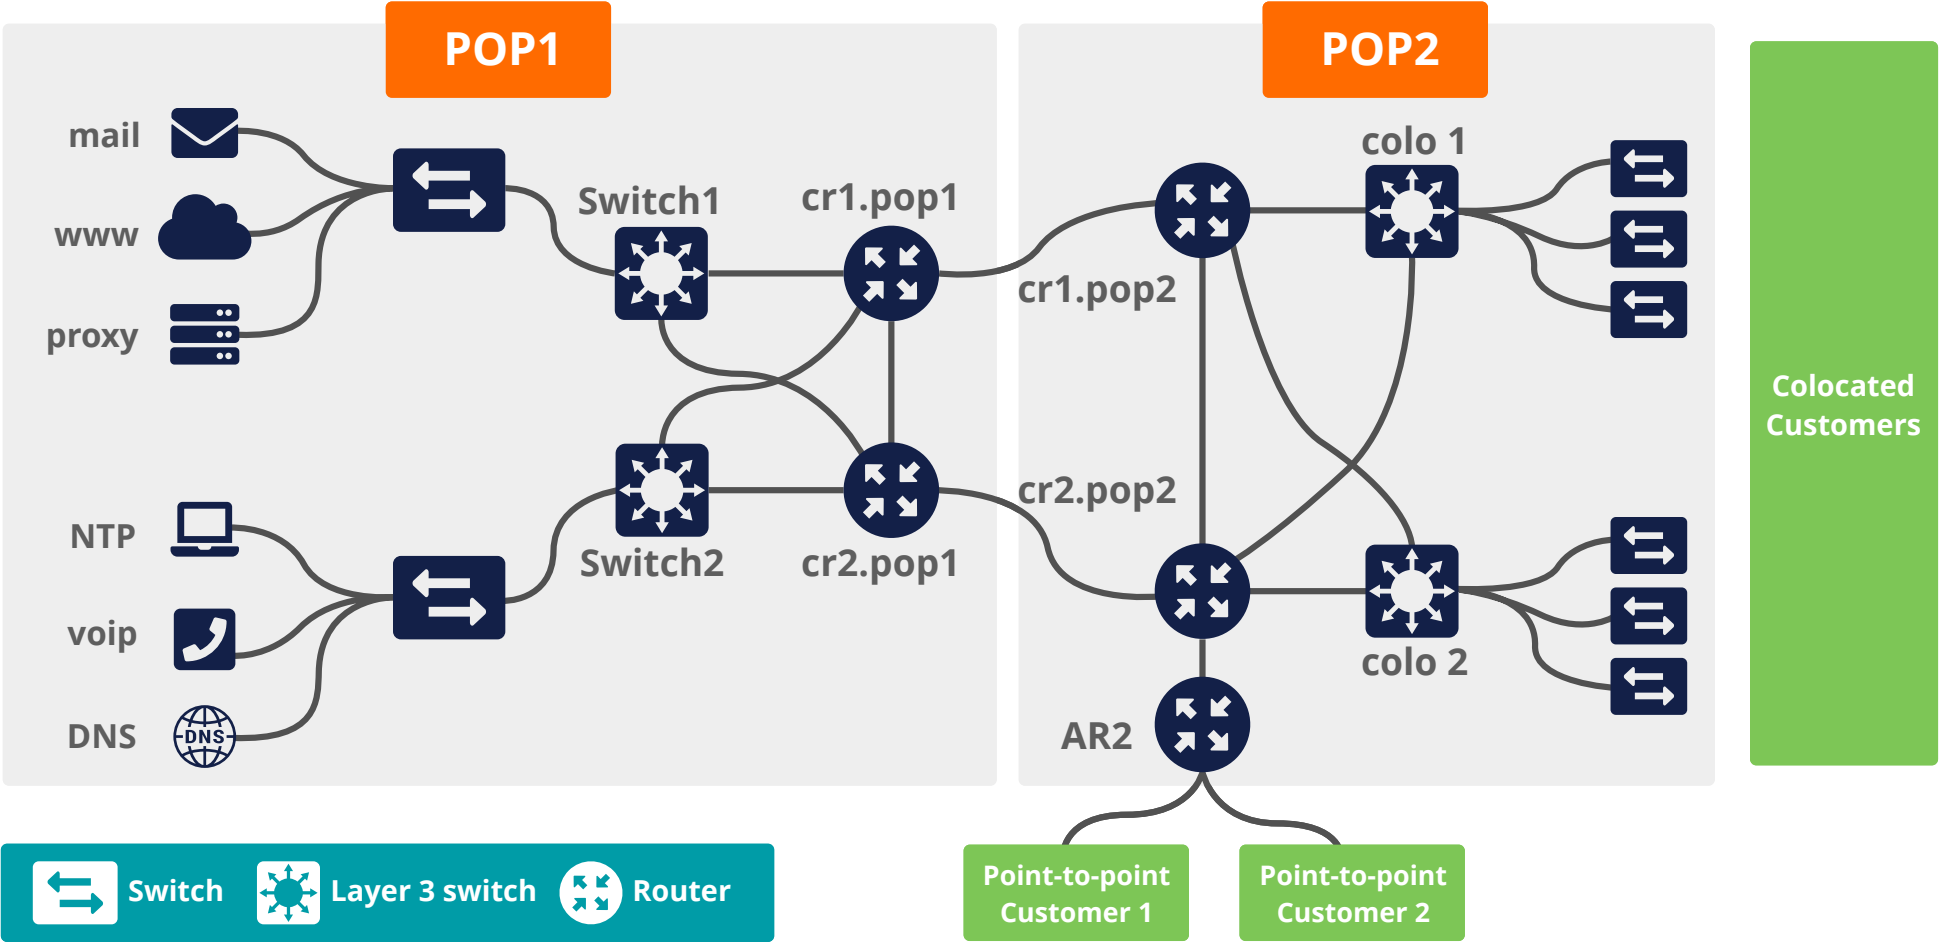


Addressing Plan Exercise

- 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

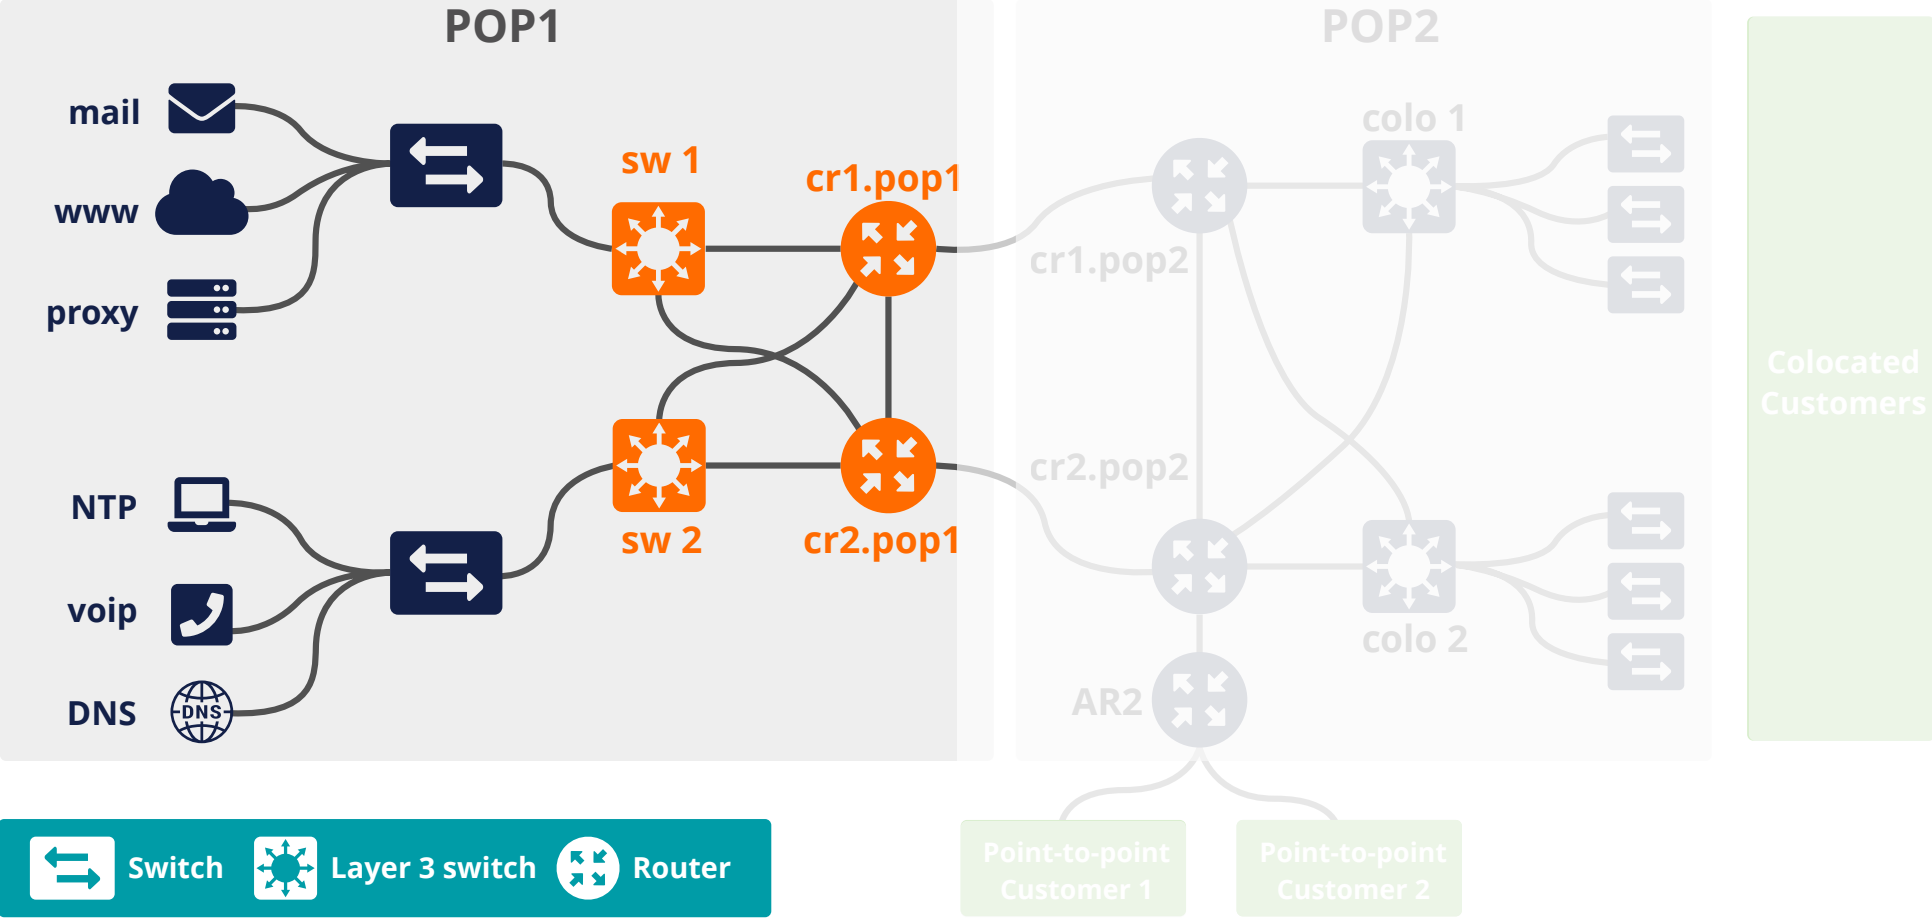


Network Diagram - POPs



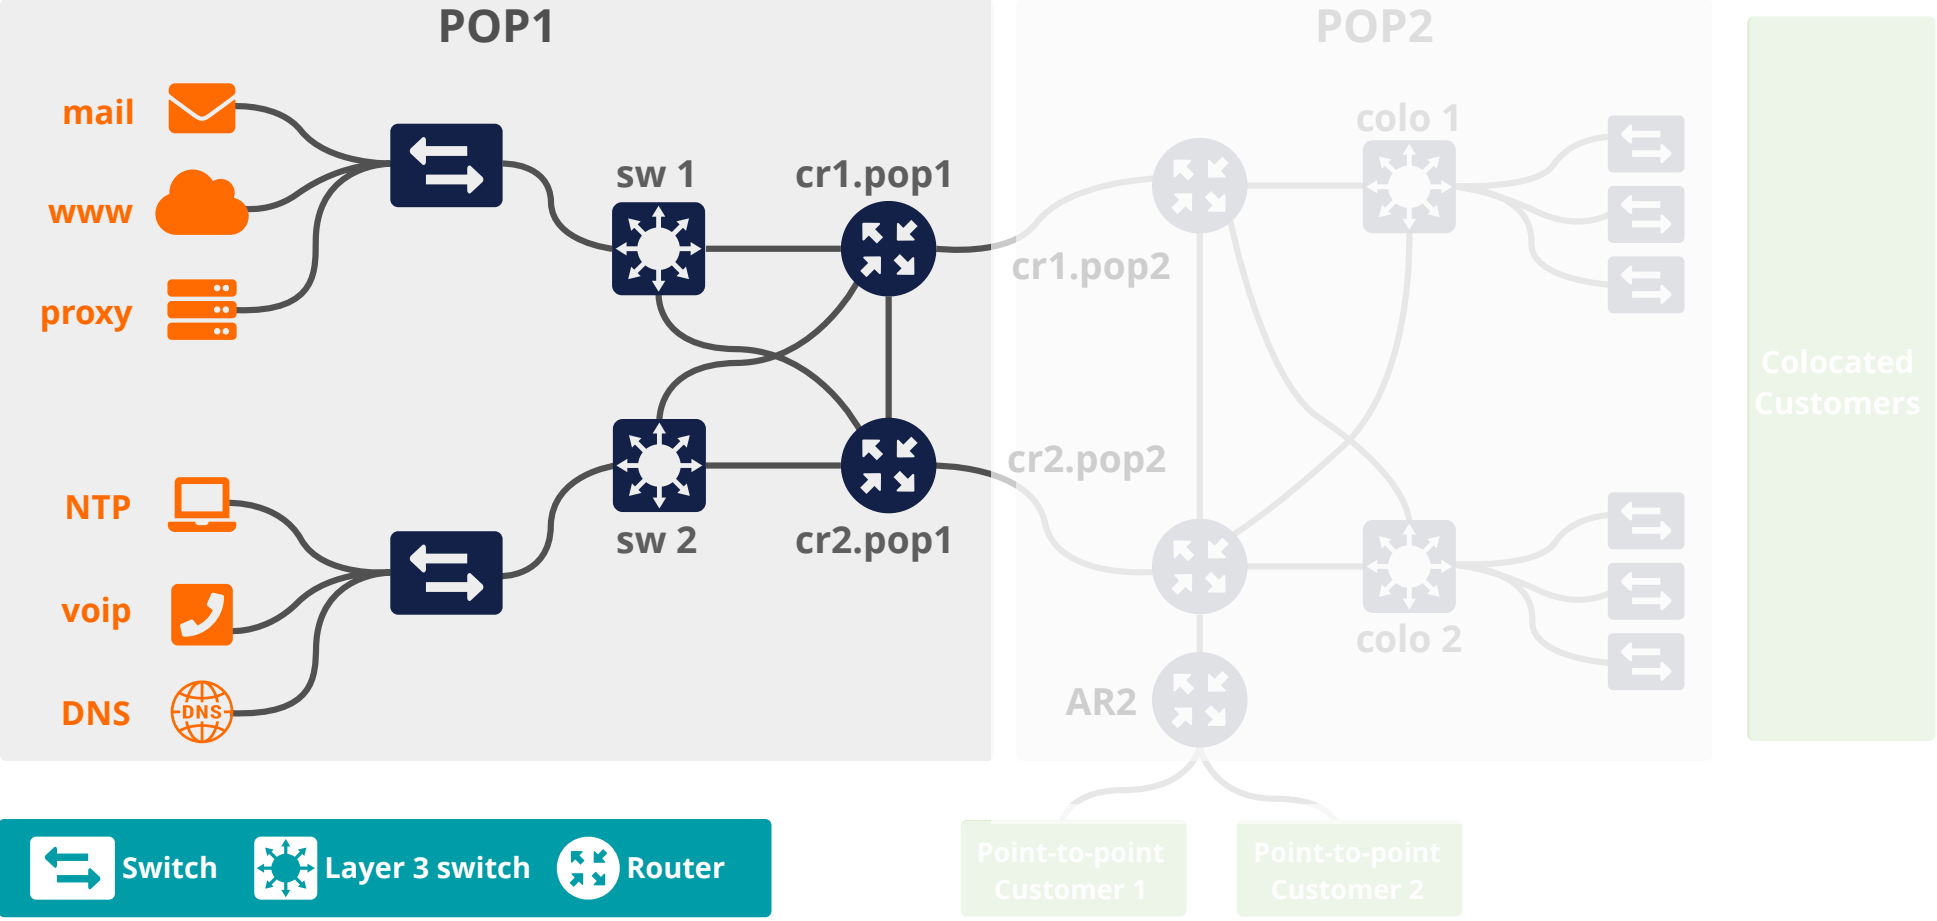


Network Diagram - POP1

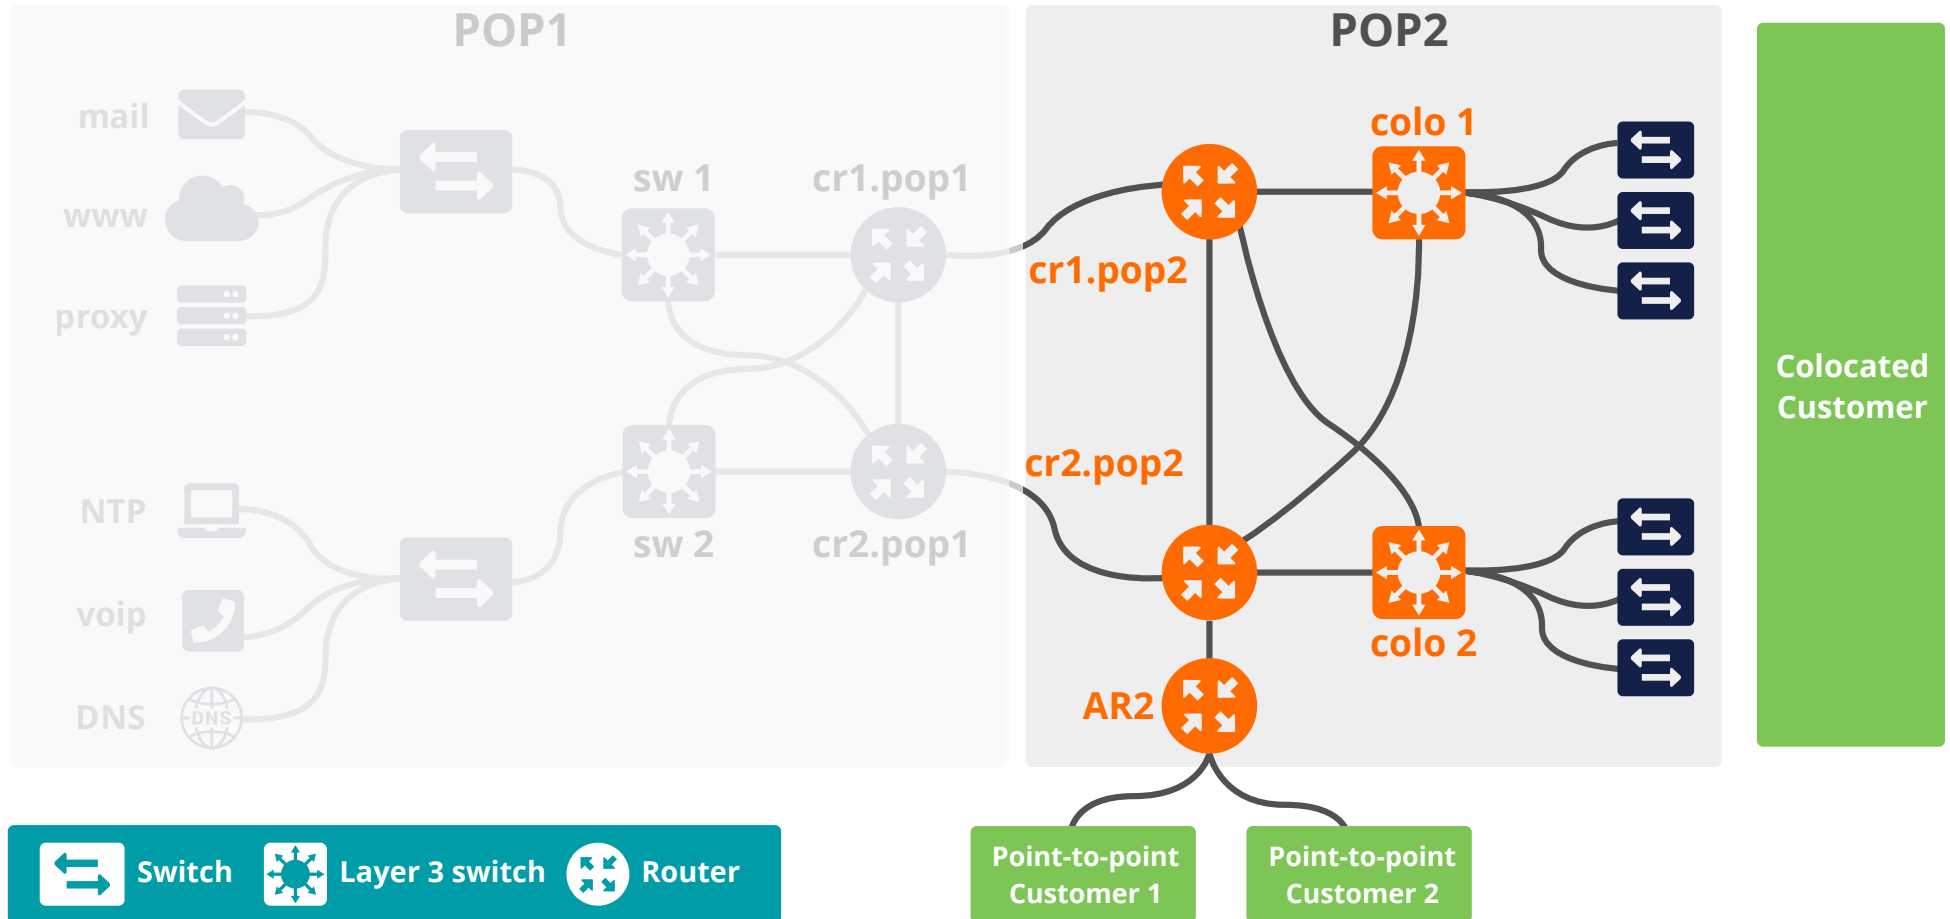




Network Diagram - POP1

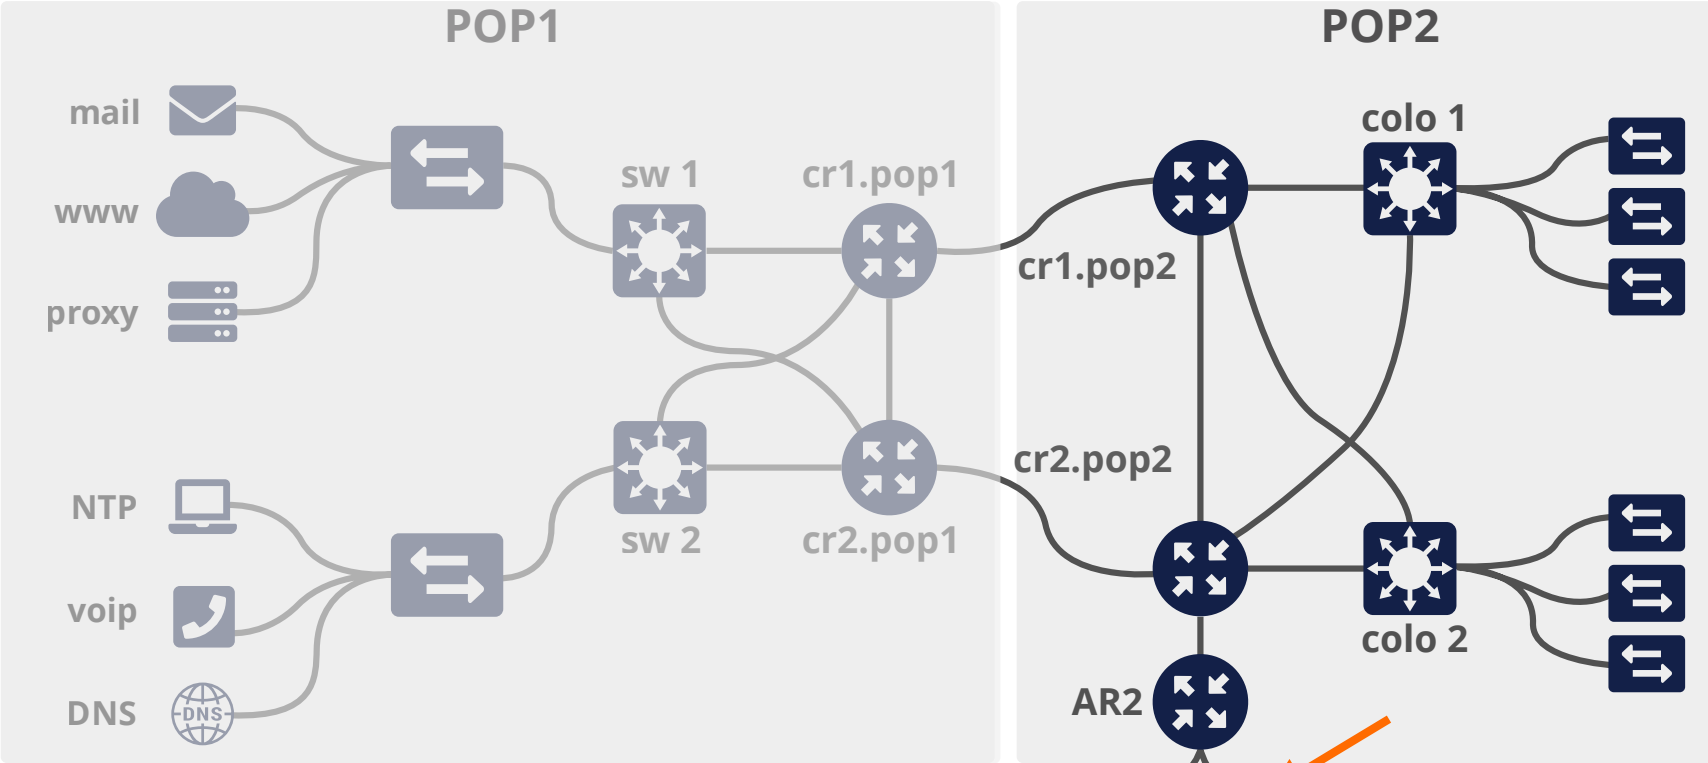


Network Diagram - POP2





Network Diagram - POP2



Colocated Customer

Point-to-point Customer 1
Point-to-point Customer 2



Addressing plans

- /64 for each subnet
- 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



The /64 story

- “Every interface ID must be a /64” (RFC 4291)
- Because of SLAAC
- Other RFCs followed this

- The **only** exception is a /127 for point-to-point links



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





IPv6 Packets

Section 6



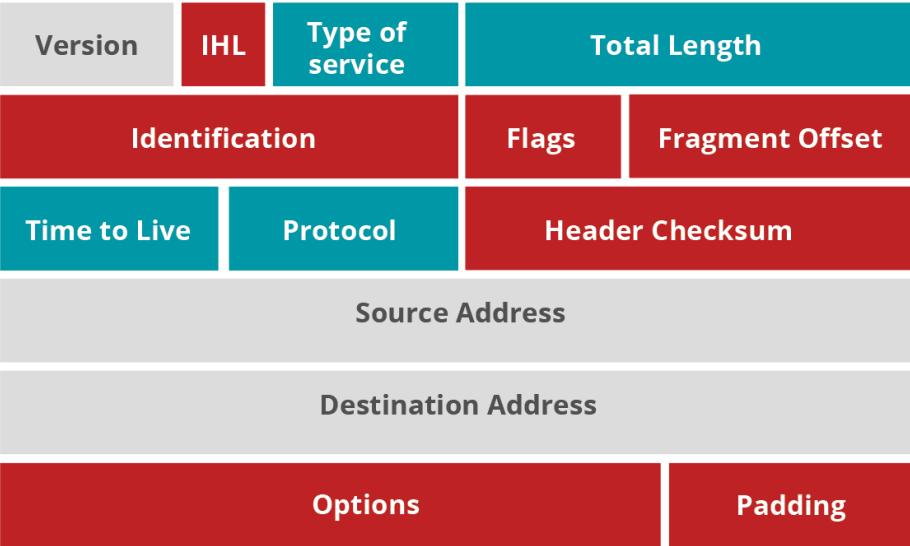
IPv6 Header Format

- Fixed length
 - Optional headers are daisy-chained
- IPv6 header is twice as long (40 bytes) as IPv4 header without options (20 bytes)

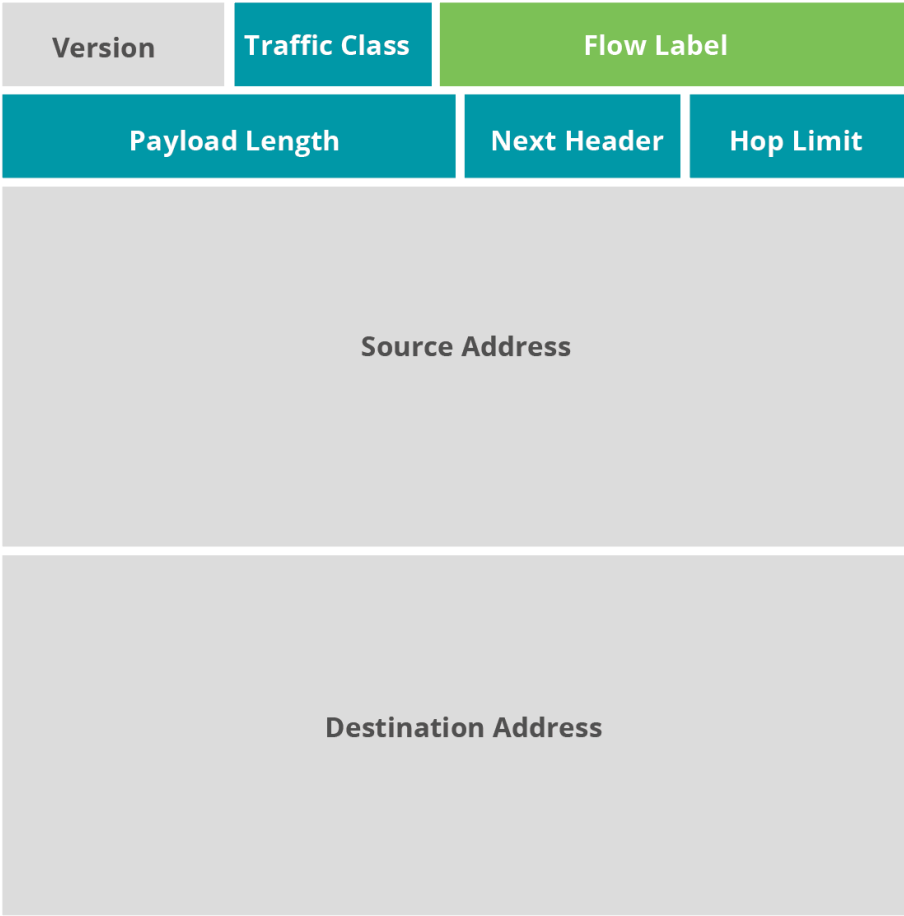


IPv6 Header

IPv4 Header



IPv6 Header




LEGEND

- Field's name kept from IPv4 to IPv6
- Field not kept in IPv6
- Name and position changed in IPv6
- New field in IPv6



IPv6 Header

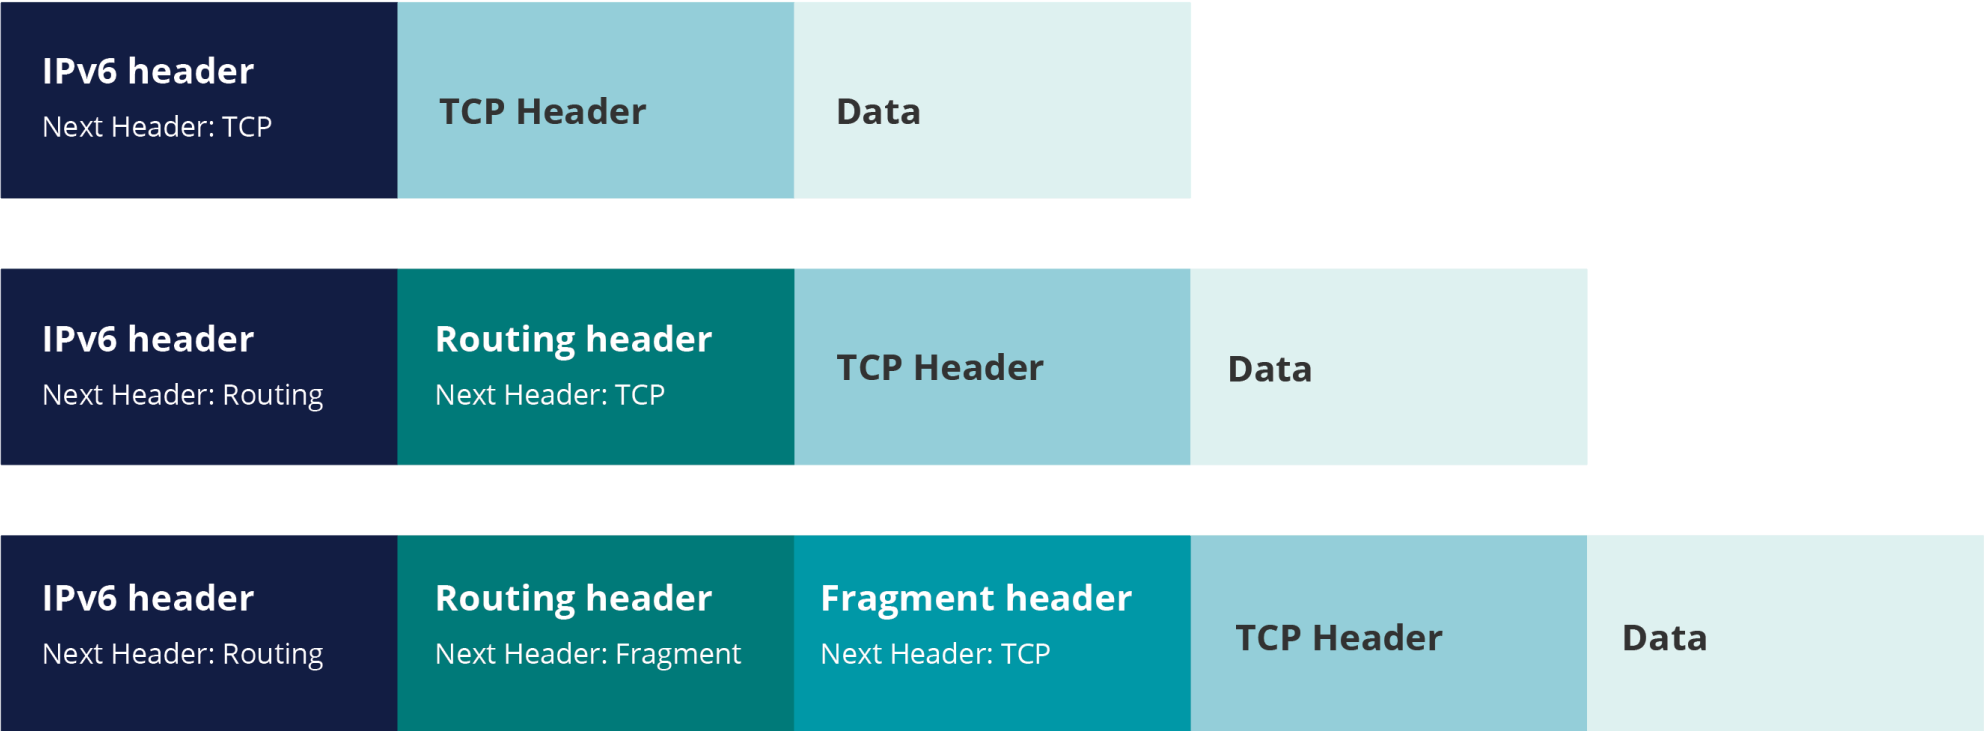
- Optional fields go into extension headers





IPv6 Header

- Daisy-chained after the main header





Common Headers

- Common values of Next Header Fields:
 - 0 Hop-by-hop option (extension)
 - 6 TCP (payload)
 - 17 UDP (payload)
 - 43 Routing (extension)
 - 44 Fragmentation (extension)
 - 50 Encrypted Security Payload (extension)
 - 58 ICMPv6




Fragmentation

- Routers don't fragment packets with IPv6
 - More efficient handling of packets in the core
 - Fragmentation is being done by host
- If a packet is too big for next hop:
 - "Packet too big" error message
 - This is an ICMPv6 message
 - Filtering ICMPv6 causes problems



Path MTU Discovery

- A sender who gets this “message-too-big” ICMPv6 error tries again with a smaller packet
 - A hint of size is in the error message
 - This is called Path MTU Discovery

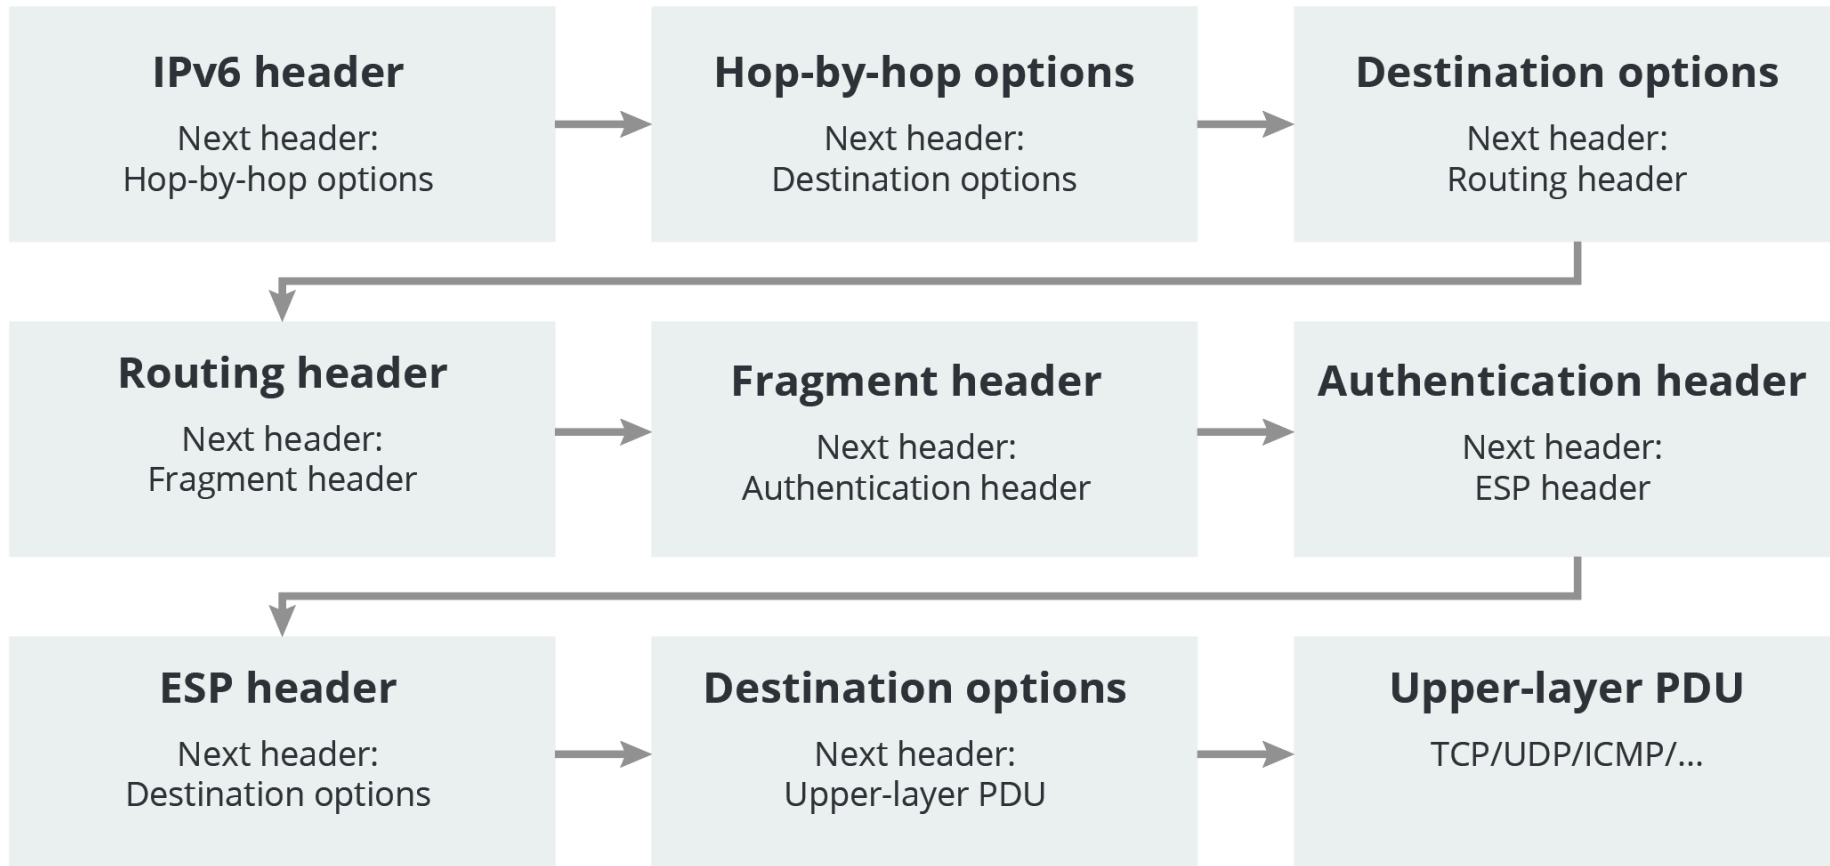




Ordering of Headers

- Order is important:
 - Only hop-by-hop header has to be processed by every node
 - Routing header needs to be processed by every router
 - Fragmentation has to be processed before others at the destination

Ordering of Headers





Broadcast

- IPv6 has no broadcast
- There is an “all nodes” multicast group
 - ff02::1
- Disadvantages of broadcast:
 - It wakes up all nodes
 - Only a few devices are involved
 - Can create broadcast storms



Neighbor Discovery

- IPv6 has no ARP
- Replacement is called Neighbor Discovery
 - Uses ICMPv6
 - Uses Multicast
- Every ARP request wakes up every node
- Each ND request only wakes up a few nodes



Neighbor Discovery

- ND is used by nodes:
 - For address resolution
 - To find neighboring routers
 - To track address changes
 - To check neighbor reachability
 - To do Duplicate Address Detection

- ND uses 5 different ICMPv6 packet types



Questions





Deploying IPv6

Section 7




Assigning Addresses

- Routers influence how hosts connect to network
- Several options:
 - Manual configuration
 - Router Advertisement only (SLAAC)
 - RA + DHCPv6 ('M' flag on)
 - RA + DHCPv6 ('O' flag on)
 - RA ('A' flag off) + DHCPv6 ('M' flag on)
- Gateway is always provided by the RA

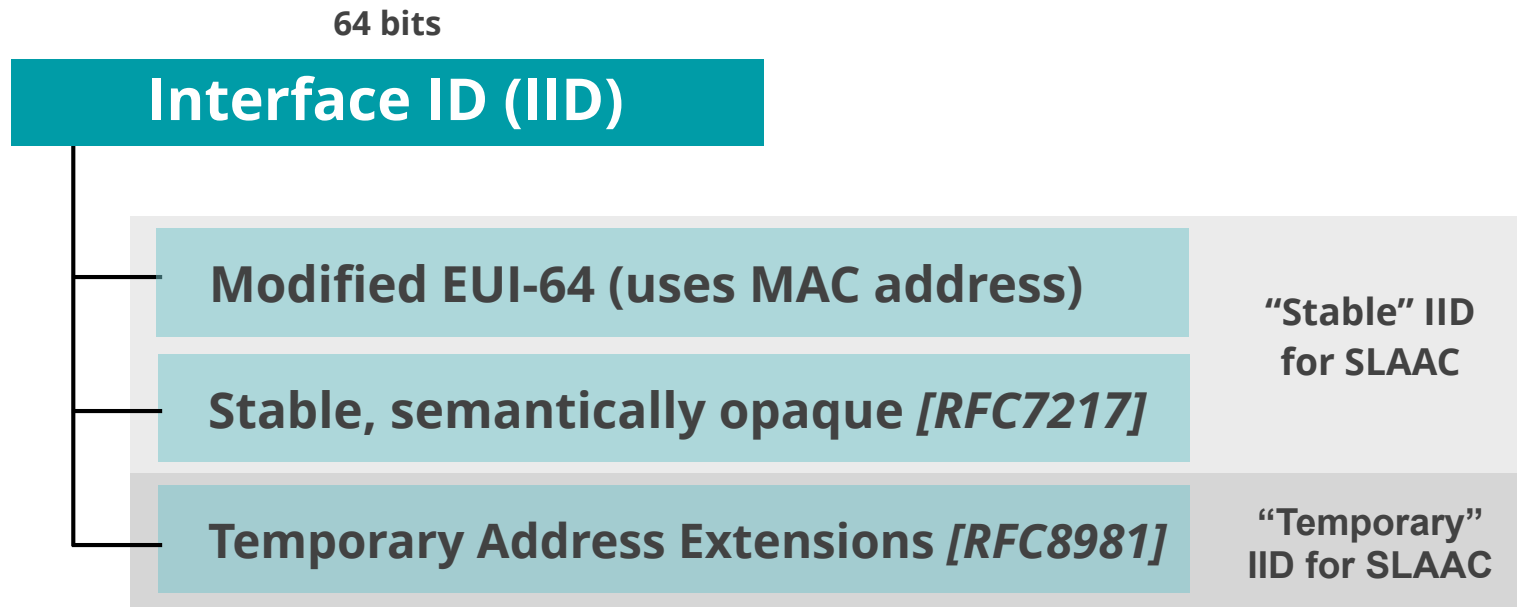


Router Advertisement Options

- RA message is used to provide configuration info
 - Default gateway address
 - Which prefix(es) to use on the link? Prefix length?
 - Is SLAAC allowed?
 - Is DHCPv6 available? For address/options? Only options?
 - What is the preference of a router on the link?
 - DNS servers / Domain (optional)
 - MTU size (optional)



SLAAC IID Generation Options





Privacy Extensions for SLAAC

- Provides privacy for users
- Privacy Extensions changes the interface ID randomly over time
- Duplicate Address Detection ensures uniqueness
- In case of collision, a new address should be generated



Stable, Semantically Opaque IID



- Consider IID bits “**opaque**”, no value or meaning *[RFC7136]*

How to generate IIDs *[RFC7217]*

Different for each interface in the same network prefix

Not related to any fixed interface identifier

Always the same when same interface connected to same network

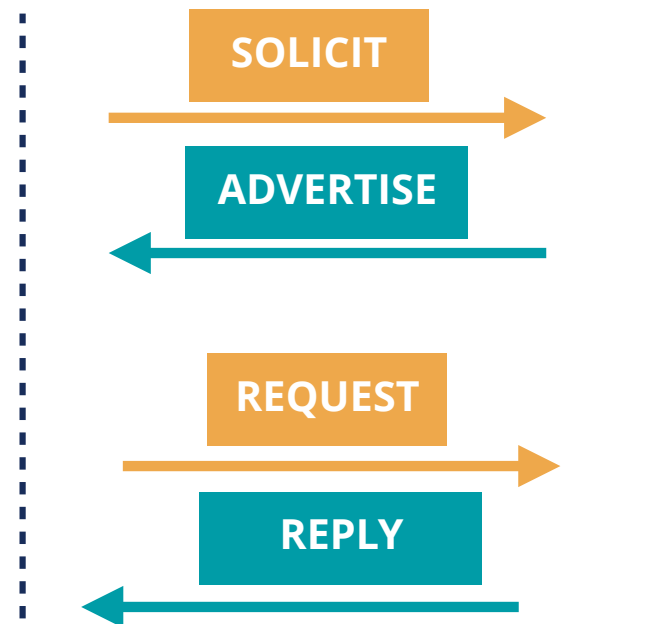

- Widely used and standardised for “stable” addresses *[RFC8064]*



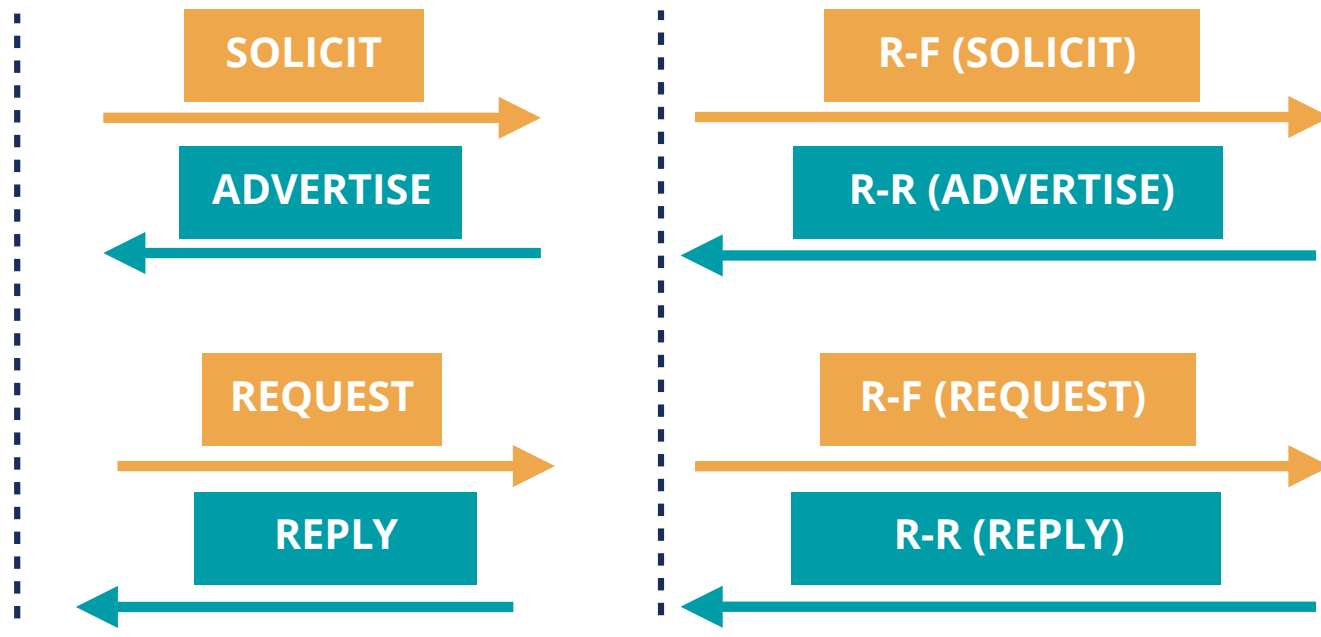
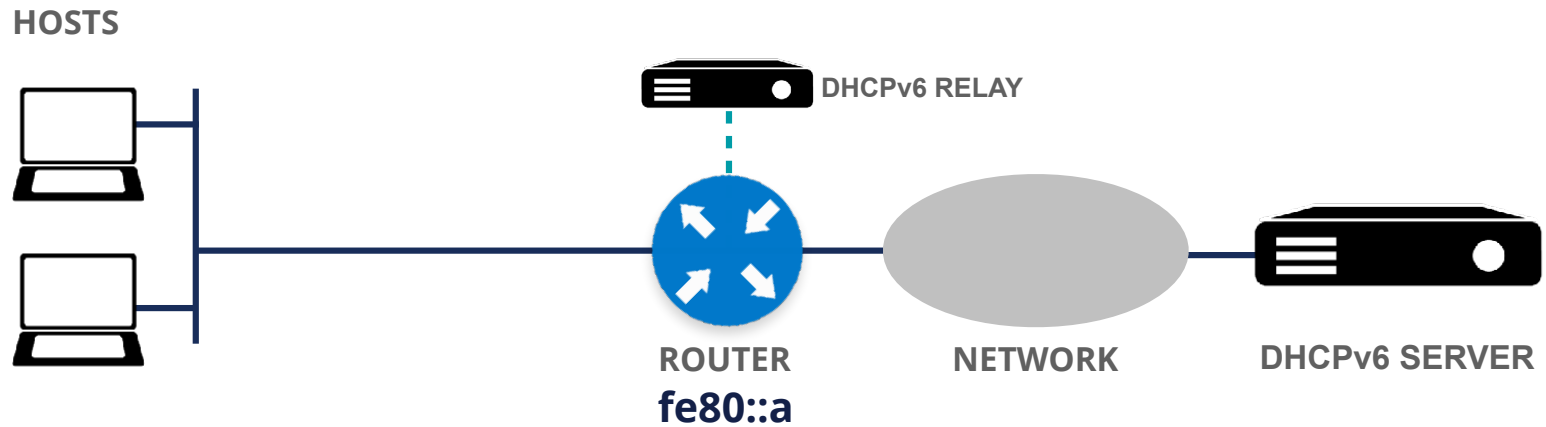
DHCPv6

- Used to give additional information like DNS servers or to manage the address pool
- Router Advertisement message contains hints
 - If “managed” flag = ‘1’ \Rightarrow can use DHCPv6 to get an address
 - Optionally provide the address of a DNS server (RFC 8106)
- Using additional flags, the network admin can disable SLAAC and force DHCPv6



DHCPv6 (M=1)



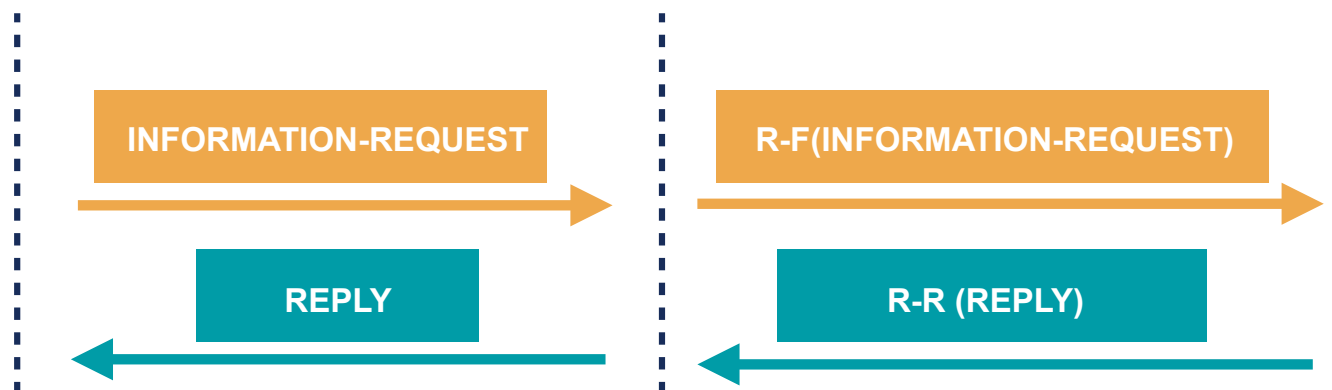
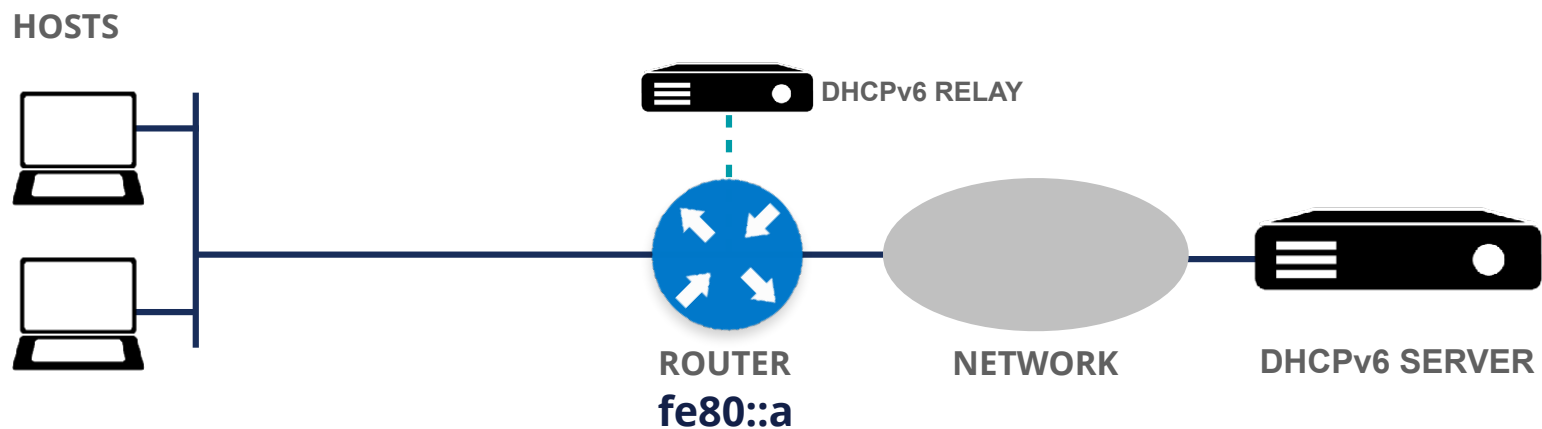
DHCPv6 (M=1)



DHCPv6 (M=0, O=1)



DHCPv6 (M=0, O=1)





MLD

- Multicast Listener Discovery (MLD) is an important component of IPv6
- IPv6 routers use MLD to discover multicast listeners on a directly attached link, similar to IGMP in IPv4
- MLD is embedded in ICMPv6. Two versions exist:
 - MLDv1 similar to IGMPv2
 - MLDv2 similar to IGMPv3

MLD



- 3 types of messages: Query, Report, Done

MLD	IGMP	Message Type	ICMPv6 Type	Function
MLDv1 (RFC2710)	IGMPv2	Listener Query	130	Discover multicast listeners
		Listener Report	131	Response to a Query, joins a group
		Listener Done	132	Node reports that it has stopped listening
MLDv2 (RFC3810)	IGMPv3	Listener Query	130	Discover multicast listeners
		Listener Report	143	Current multicast listening state, or changes



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.0.0.0.0.0.0.0.1.1.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.0.0.0.0.0.0.0.1.1.f.e.e.3.0.0.8.b.d.0.1.0.0.2.ip6.arpa. PTR yourname.domain.tld.



IPv6 and Domain Objects

- IPv6 prefix: **2001:db8::/32**
- Domain object:
 - domain: **8.b.d.0.1.0.0.2.ip6.arpa**
 - descr: rDNS for my whole IPv6 network
 - admin-c: NOC12-RIPE
 - tech-c: NOC12-RIPE
 - zone-c: NOC12-RIPE
 - nserver: pri.example.net
 - nserver: sns.company.org
 - ds-rdata: 45062 8 2 275d9acbf3d3fec11b6d6...
 - mnt-by: EXAMPLE-LIR-MNT
 - created: 2015-01-21T13:52:29Z
 - last-modified: 2016-02-07T15:09:46Z
 - source: RIPE



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 router 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 two IP protocols



Configuring IPv6

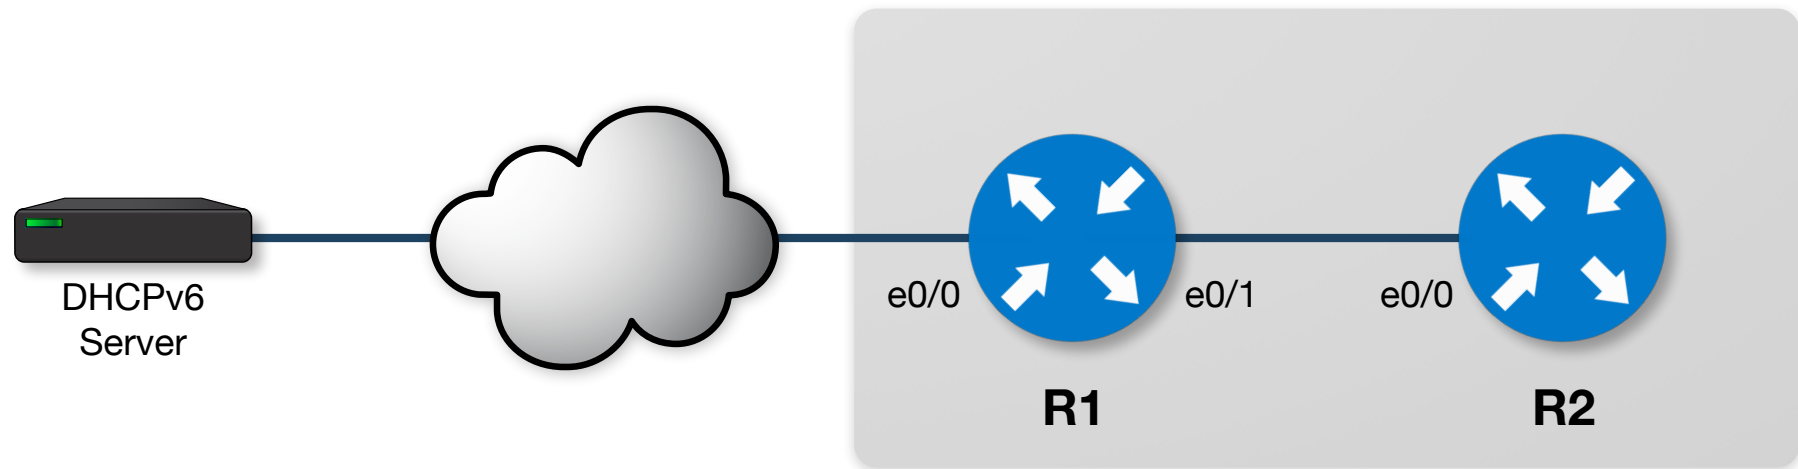
Exercise



Assigning Addresses

- **R1** will send the RAs and act as DHCPv6 Relay Agent
- **R2** will get IPv6 configuration info in three ways:
 - RA + SLAAC only
 - RA + SLAAC + 'O' flag (DHCPv6 Other Configuration)
 - RA + **no SLAAC** + 'M' flag (DHCPv6 Managed)
- The DHCPv6 server is already configured

Network Diagram



Router roles:

R1: Default gateway router
DHCPv6 relay agent

R2: Client device
SLAAC
DHCPv6 client



Exercise: Configuring IPv6

- Make sure you have connectivity
- Go to: workbench.ripe.net
- Choose the lab (ask the trainers)
- Your login is your number on participants list
- The trainers will provide the password

- Choose “**RA and DHCPv6**” from the menu



Check R2

- Verify that the interface e0/0 has no address yet

```
show ipv6 interface brief
```



Basic IPv6 Settings

- Before configuring IPv6 on your router interfaces, the basic IPv6 settings must be enabled
- On both **R1** and **R2**

```
configure terminal
```

```
ipv6 unicast-routing  
ipv6 cef
```




1st Case: SLAAC only (Router)

- On **R1** we will configure an IPv6 address from a /64 prefix on interface e0/1

```
interface e0/1
  ipv6 address 2001:ffxx:1::a/64
```

Where **xx** is your given number for the LAB!

1 = 01

2 = 02

10 = 10

11 = 11

1st Case: SLAAC only (Client)



- On **R2** we will configure SLAAC on the interface e0/0

```
interface e0/0
  ipv6 address autoconfig default
```



Check R2

- Verify that the interface e0/0 has an IPv6 address

```
end      (exits config mode)
```

```
show ipv6 interface e0/0
```

- And a default route

```
show ipv6 route
```



Check R2

- Unfortunately, **R2** has no DNS name servers

```
show ip dns view
```

- This information was not provided in the RA from **R1**

2nd Case: SLAAC + O flag (Router)



- On **R1** we will configure the 'O' flag for the RAs on interface e0/1

```
interface e0/1
  ipv6 nd other-config-flag
```



2nd Case: SLAAC + O flag (Client)

- On **R2** we will first bring down the interface e0/0

```
configure terminal  
interface e0/0  
shutdown
```

- And then bring it back up...

```
no shutdown
```

2nd Case: SLAAC + O flag (Client)



- Verify that the interface e0/0 has an IPv6 address and other configuration

```
end      (exits config mode)
```

```
show ipv6 interface e0/0
```

```
show ip dns view
```

```
show ipv6 dhcp interface e0/0
```



3rd Case: RA + M flag (Router)

- On **R1** we will configure the 'M' flag for the RAs on interface e0/1

```
interface e0/1
  no ipv6 nd other-config-flag
  ipv6 nd managed-config-flag
```




3rd Case: RA + M flag (Client)

- On **R2** we will first bring down the interface e0/0

```
configure terminal  
interface e0/0  
shutdown
```

- Remove the SLAAC configuration

```
no ipv6 address autoconfig default
```



3rd Case: RA + M flag (Client)

- On R2, configure the DHCP client

```
ipv6 address dhcp  
ipv6 enable  
ipv6 nd autoconfig default-route
```

- And then bring the interface back up...

```
no shutdown
```



3rd Case: RA + M flag (Client)

- Verify that the interface e0/0 has an IPv6 address and other configuration

```
end      (exits config mode)
```

```
show ipv6 interface e0/0
```

```
show ipv6 dhcp interface e0/0
```



Questions





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



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



RIPE-772 Document

- “Requirements for IPv6 in ICT Equipment”
 - Best Current Practice describing what to ask for when requesting IPv6 Support
 - Useful for tenders and RFPs
 - Original version was ripe-554
 - Ripe-554 Originated by the Slovenian Government
 - Adopted by various others (Germany, Sweden)

Link to the document:

<https://www.ripe.net/publications/docs/ripe-772>



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>





Customers And Their /48

- Customers have no idea how to handle **65,536 subnets!**
- Provide them with information!



Link to the document:

<https://www.ripe.net/support/training/material/basicipv6-addressing-plan-howto.pdf>



Also useful

- Websites
 - <http://www.getipv6.info>
 - <http://www.ipv6actnow.org>
 - <http://datatracker.ietf.org/wg/v6ops/>
 - <https://www.ripe.net/publications/docs/ripe-772>
- Mailing lists
 - <http://lists.cluenet.de/mailman/listinfo/ipv6-ops>
 - <http://www.ripe.net/mailman/listinfo/ipv6-wg>



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!



Questions





What's Next in IPv6



Webinars

Attend another webinar live wherever you are.

- ❖ Introduction to IPv6 (2 hrs)
- ❖ IPv6 Host Configuration (2 hrs)
- ❖ IPv6 Addressing Plan (1 hr)
- ❖ Basic IPv6 Protocol Security (2 hrs)
- ❖ IPv6 Associated Protocols (2 hrs)
- ❖ IPv6 Security Myths, Filtering and Tips (2 hrs)

↓ For more info
click the link
below



learning.ripe.net



Face-to-face

Meet us at a location near you for a training session delivered in person.

- ❖ Basic IPv6 (8.5 hrs)
- ❖ Advanced IPv6 (17 hrs)
- ❖ IPv6 Security (8.5 hrs)



E-learning

Learn at your own pace at our online Academy.

- ❖ IPv6 Fundamentals (15 hrs)
- ❖ IPv6 Security (24 hrs)

↓ For more info
click the link
below



academy.ripe.net



Examinations

Learnt everything you needed? Get certified!

- ❖ IPv6 Fundamentals - Analyst
- ❖ IPv6 Security - Expert

↓ For more info
click the link
below



getcertified.ripe.net



Learn something new today!
academy.ripe.net






RIPE NCC Certified Professionals



<https://getcertified.ripe.net/>





We want your feedback!

What did you think about this session?

Take our survey at:

<https://www.ripe.net/feedback/bv6/>



Copyright Statement

[...]

The RIPE NCC Materials may be used for **private purposes, for public non-commercial purpose, for research, for educational or demonstration purposes**, or if the materials in question specifically state that use of the material is permissible, and provided the RIPE NCC Materials are not modified and are properly identified as RIPE NCC documents. Unless authorised by the RIPE NCC in writing, any use of the RIPE NCC Materials for advertising or marketing purposes is strictly forbidden and may be prosecuted. The RIPE NCC should be notified of any such activities or suspicions thereof.

[...]

Find the full copyright statement here:

<https://www.ripe.net/about-us/legal/copyright-statement>



Änn Соңы An Críoch پايان Y Diwedd
Vége Endir Finvezh Ende Koniec
Son დასასრული უტრეზ Kинецъ Finis
Lõpp Amaia תסוה Tmiem Krai
Sfârșit Loppu Slutt Liðugt Krai
Kraj النهاية Конецъ Fund
Fine Fin Fí Konec Τέλος
Einde Край Pabaiga
Slut Beigas
Fim

E₁ N₁ D₂



ANNEX



Transition Mechanisms

Annex 1

Transitioning: Solving Two Problems



- a.** Maintaining connectivity to IPv4 hosts by sharing IPv4 addresses between clients
 - a.** Extending the address space with NAT/CGN/LSN
 - b.** Translating between IPv6 and IPv4


- b.** Provide a mechanism to connect to the emerging IPv6-only networks
 - a.** Tunnelling IPv6 packets over IPv4-only networks

6in4



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

6in4






6RD

- a.** Encodes the IPv4 address in the IPv6 prefix
- b.** Uses address space assigned to the operator
- c.** The operator has full control over the relay
- d.** Traffic is symmetric across a relay
 - a.** Or at least stays in your domain
- e.** Can work with both public and private IPv4 space
- f.** Needs additional software for signalling

6RD






NAT64 / DNS64

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

NAT64 / DNS64



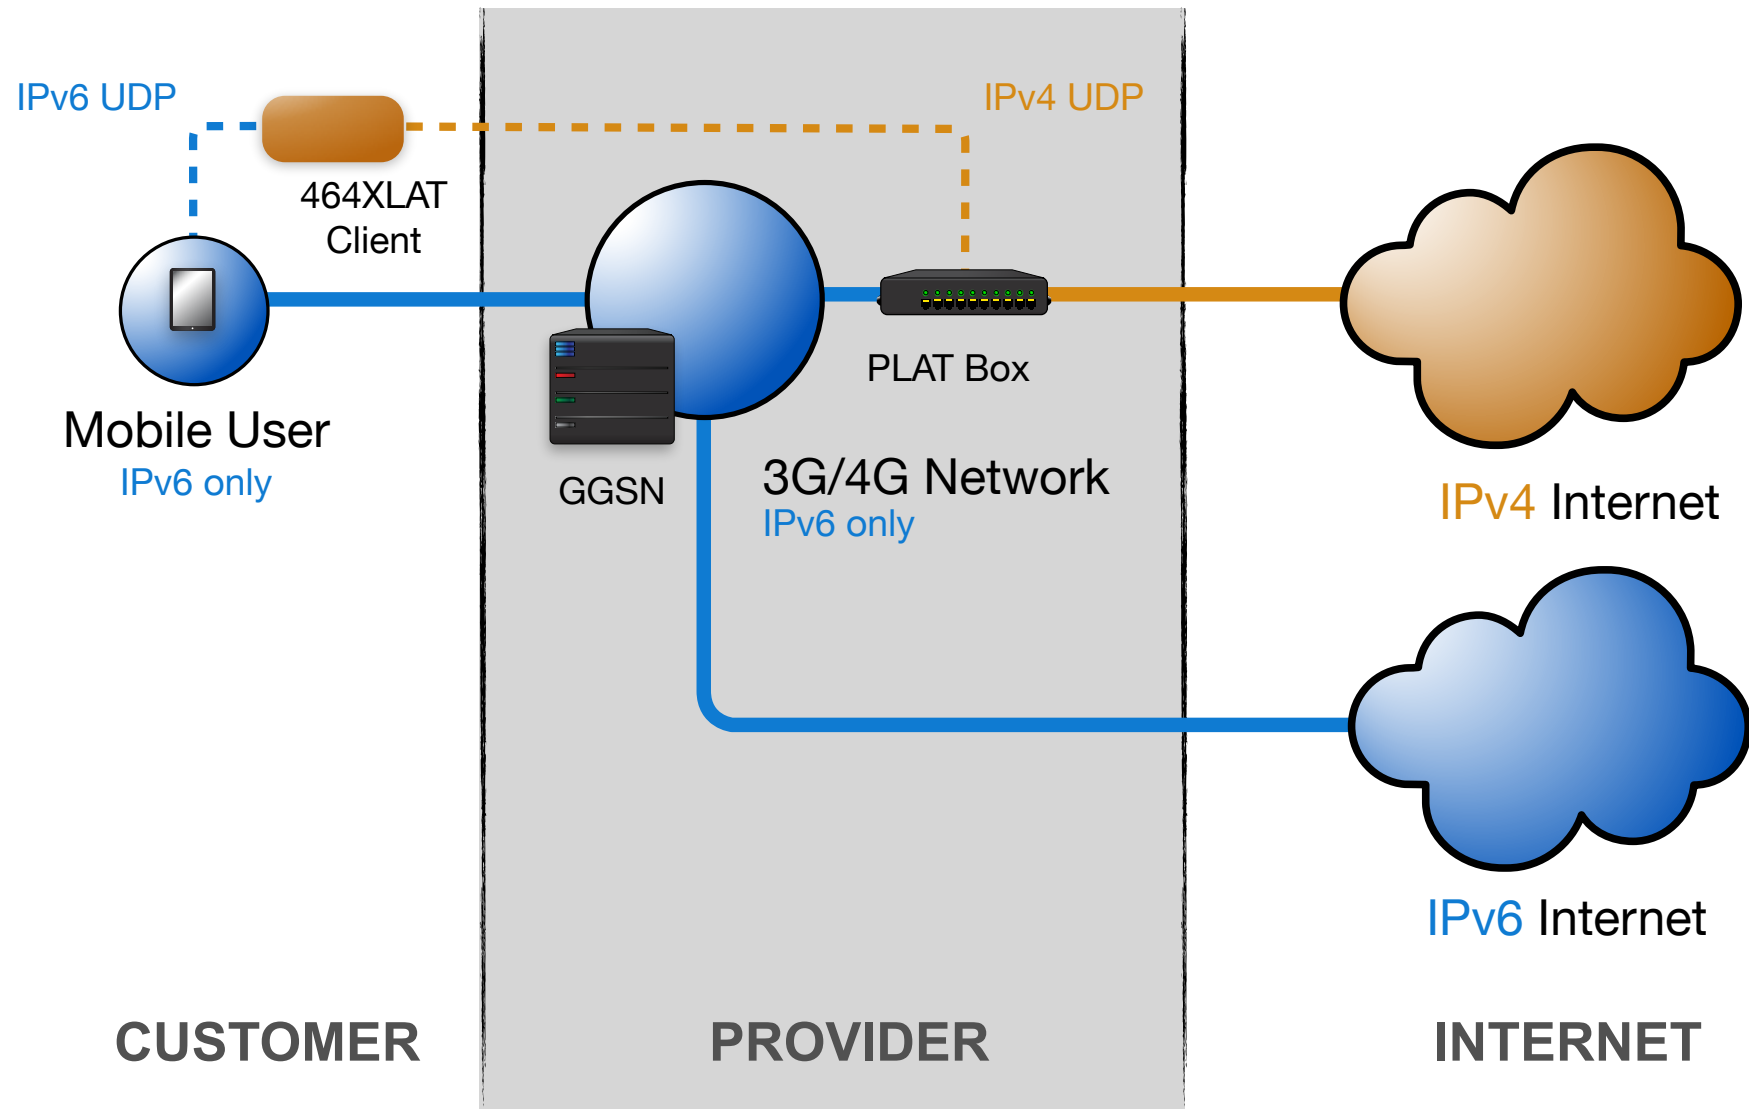


464XLAT

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

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

464XLAT

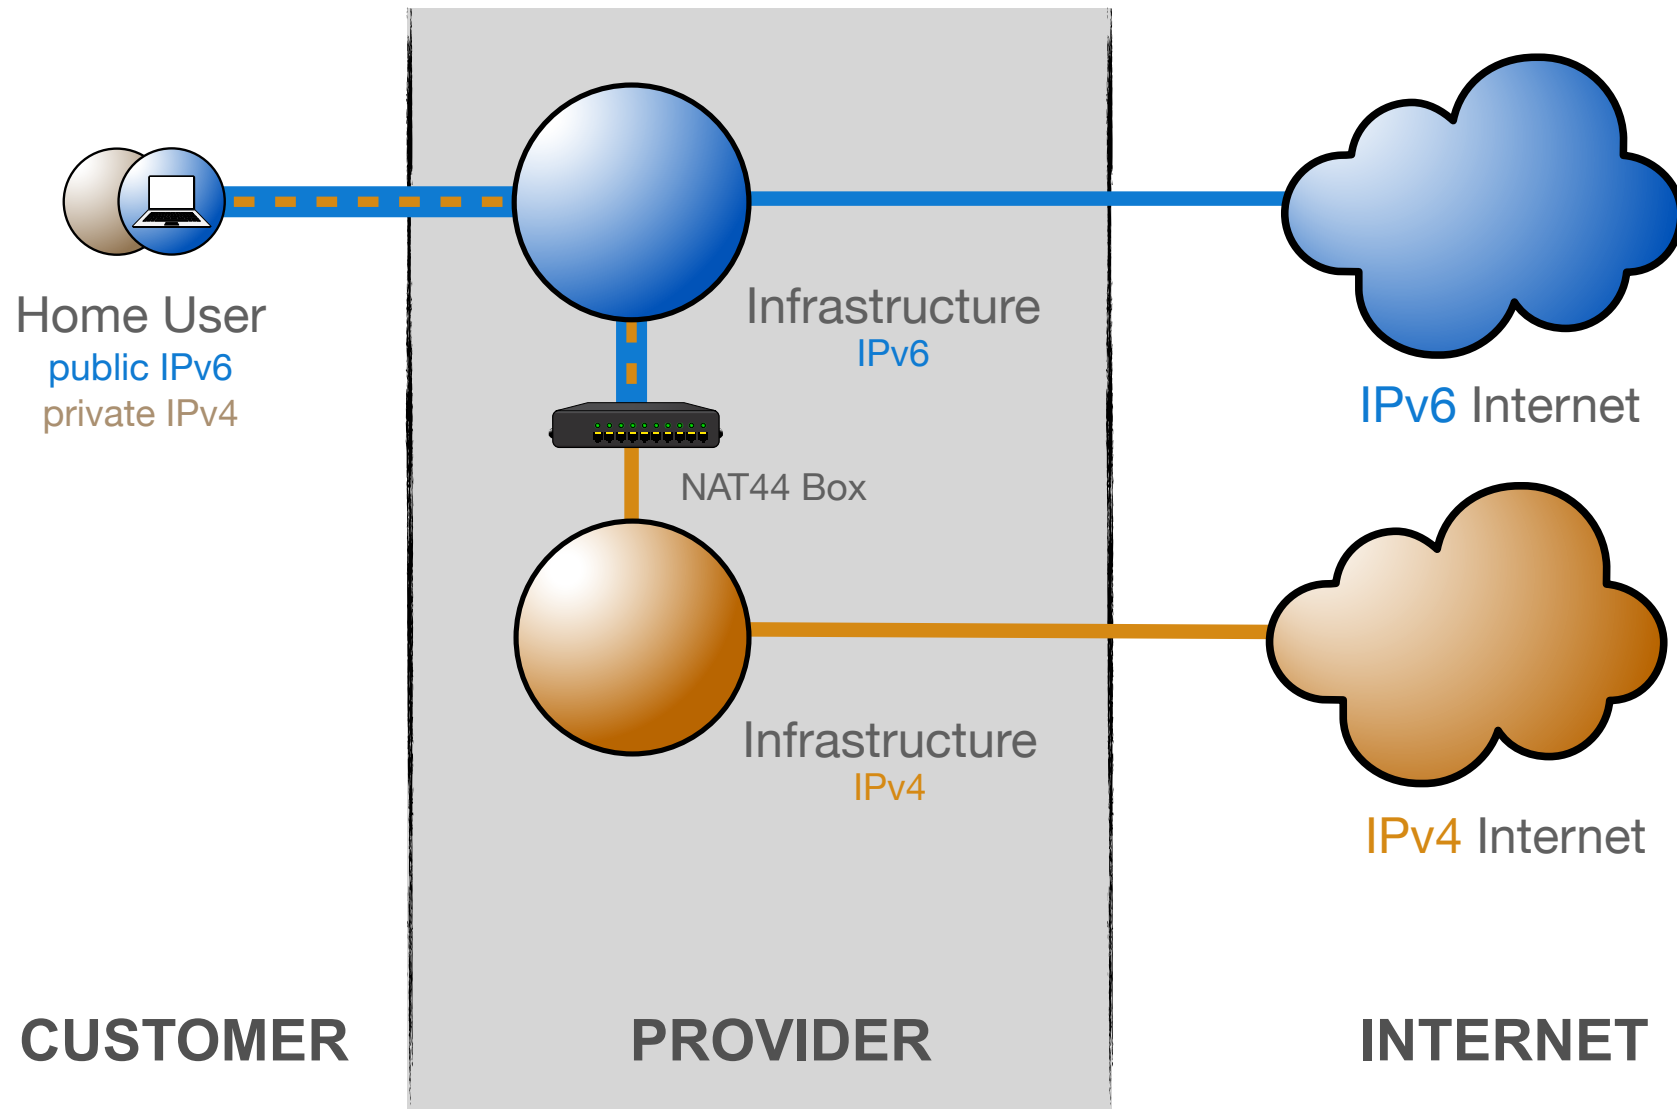




DS-lite

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

DS-lite






MAP-E / MAP-T

- a.** IPv4 over IPv6 - Encapsulated or Translated
- b.** Clients get private IPv4 and public IPv6
- c.** IPv4 address/port mapped into IPv6 address
- d.** Stateless NAT44 allows traffic to flow asymmetrically in and out of MAP domain




MAP-E / MAP-T



Best Transition Mechanism?



Dual Stack





RIPE NCC

Academy

Graduate to the next level!

<http://academy.ripe.net>

Follow us!



twitter

@TrainingRIPENCC

Feedback!



<https://www.ripe.net/training/basic-ipv6/survey>