RIPE NCC
RIPE NETWORK COORDINATION CENTRE

Basic IPv6

NOG.HR - Tutorial

October 2023

Overview

- Introduction
- IPv6 Address Basics
- Getting it
- IPv6 Protocol Basics
- Addressing Plan
- IPv6 Packets
- Deploying IPv6
- Tips

IPv4?

Section 1

IANA IPv4 Pool

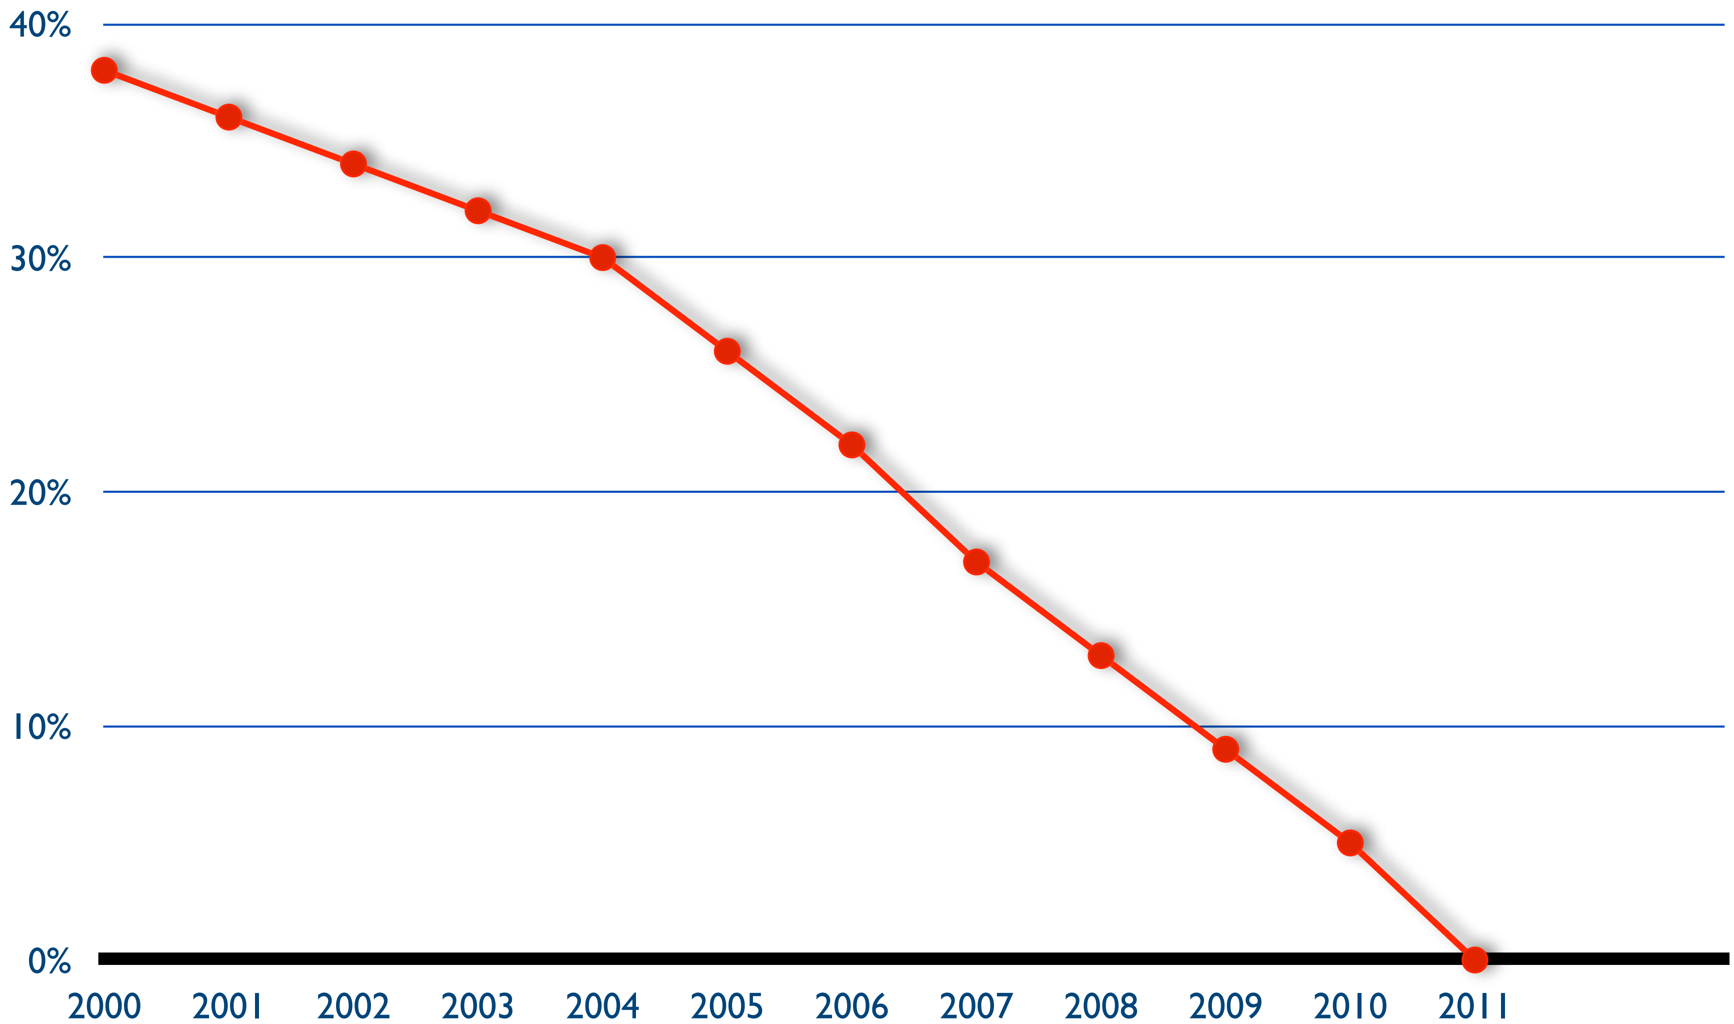

IPv4 run-out

"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

IPv6 Address Basics

Section 2

IP Address Distribution

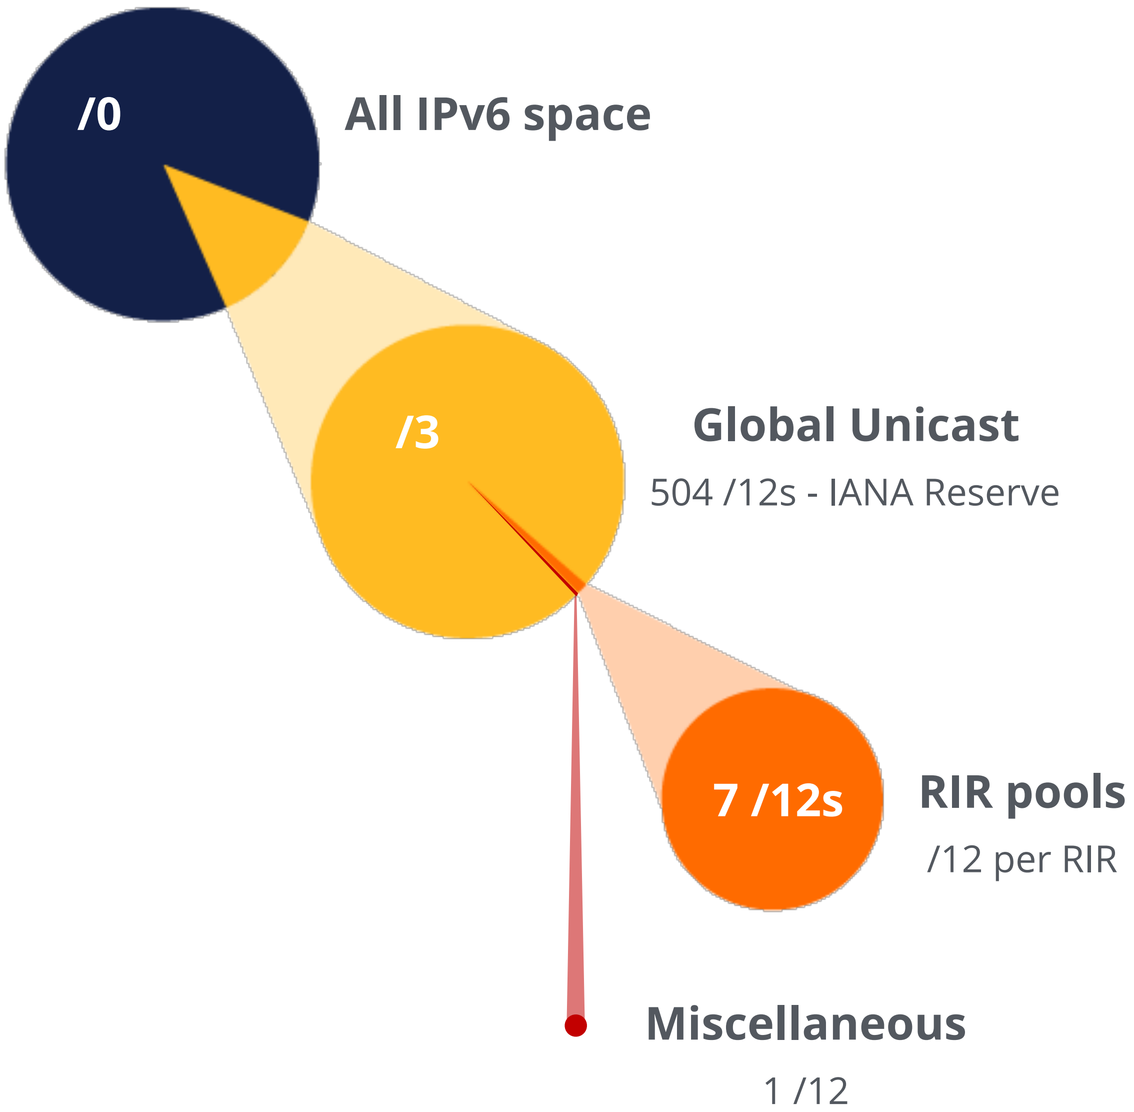

RIR Pools

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

IP Address Distribution

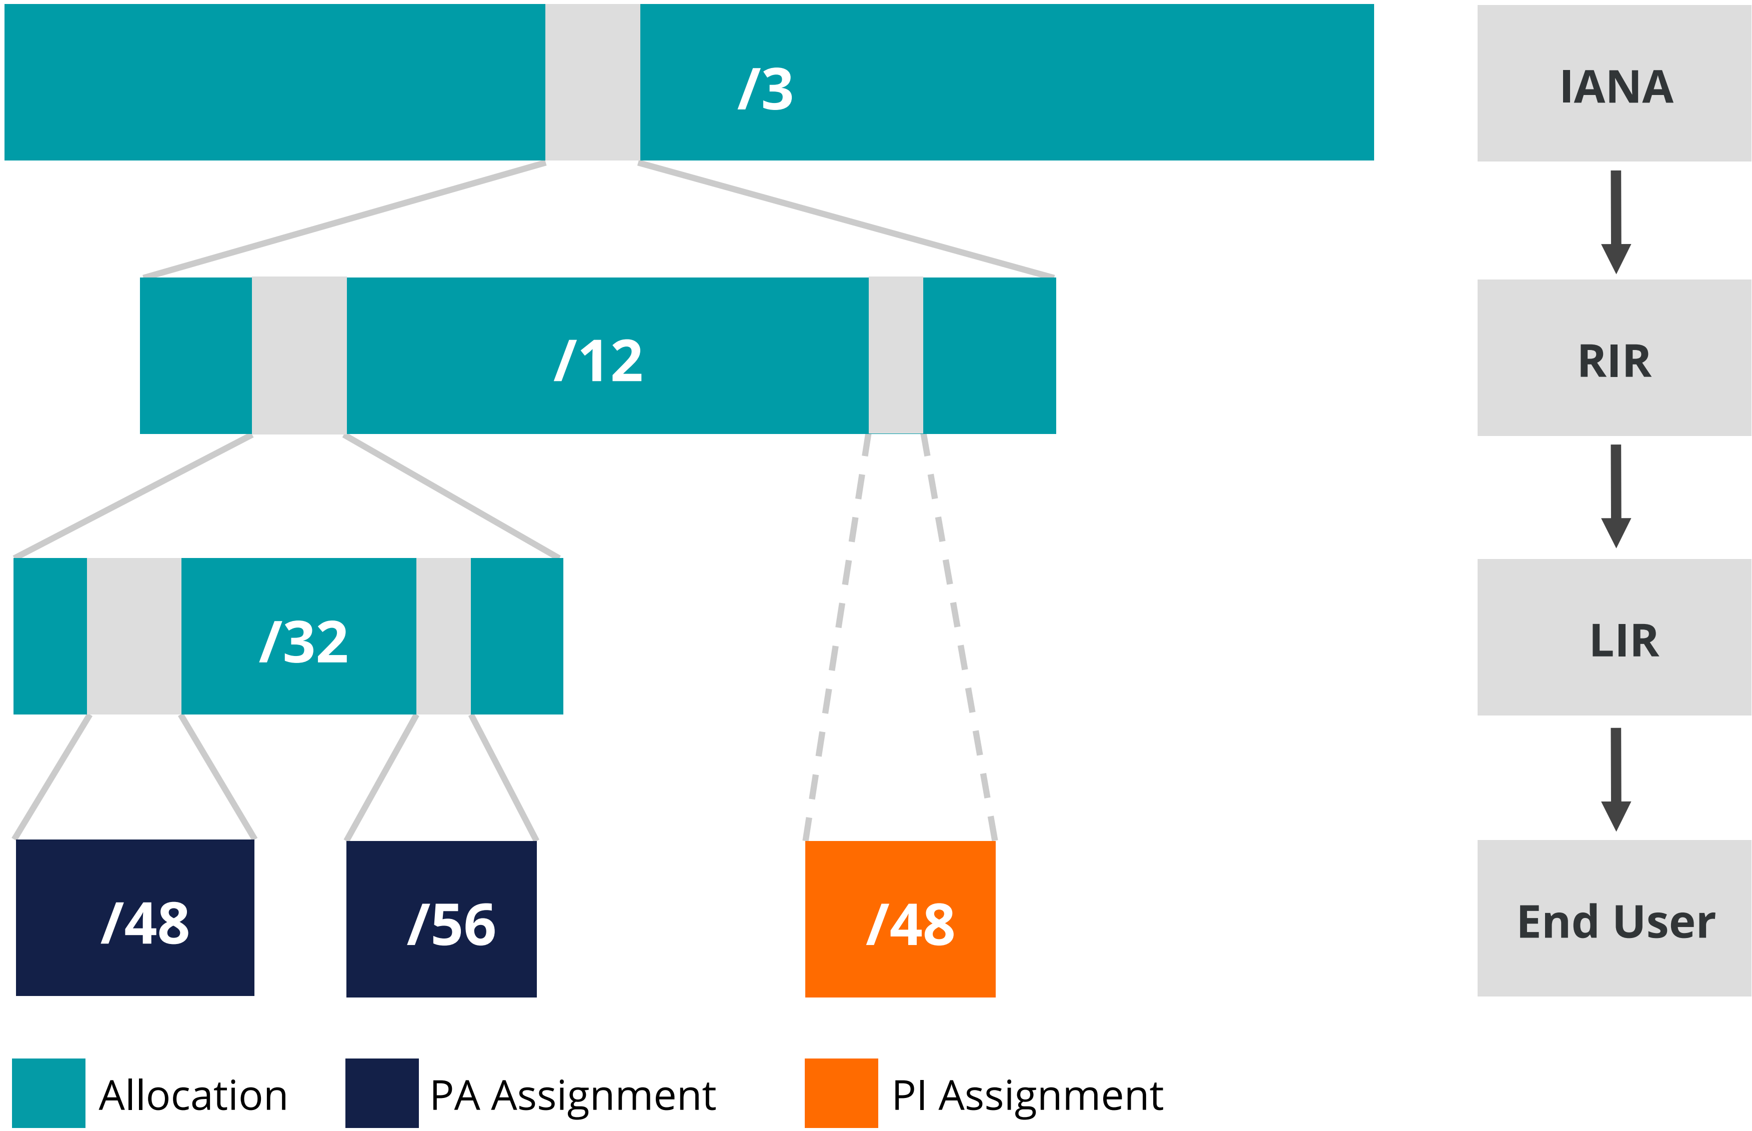

IPv6 Address Basics

- 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

1	1	1	0	1	1	1	1	0	0	0	1	0	0	0	1

IPv6 Subnetting

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

64 bits interface ID

/64

/60 = 16 x /64

/56 = 256 x /64

/52 = 4096 x /64

/48 = 65536 x /64

/32 = 65536 x /48

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

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
- For more than /48 - you need to document it well!
- 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
	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

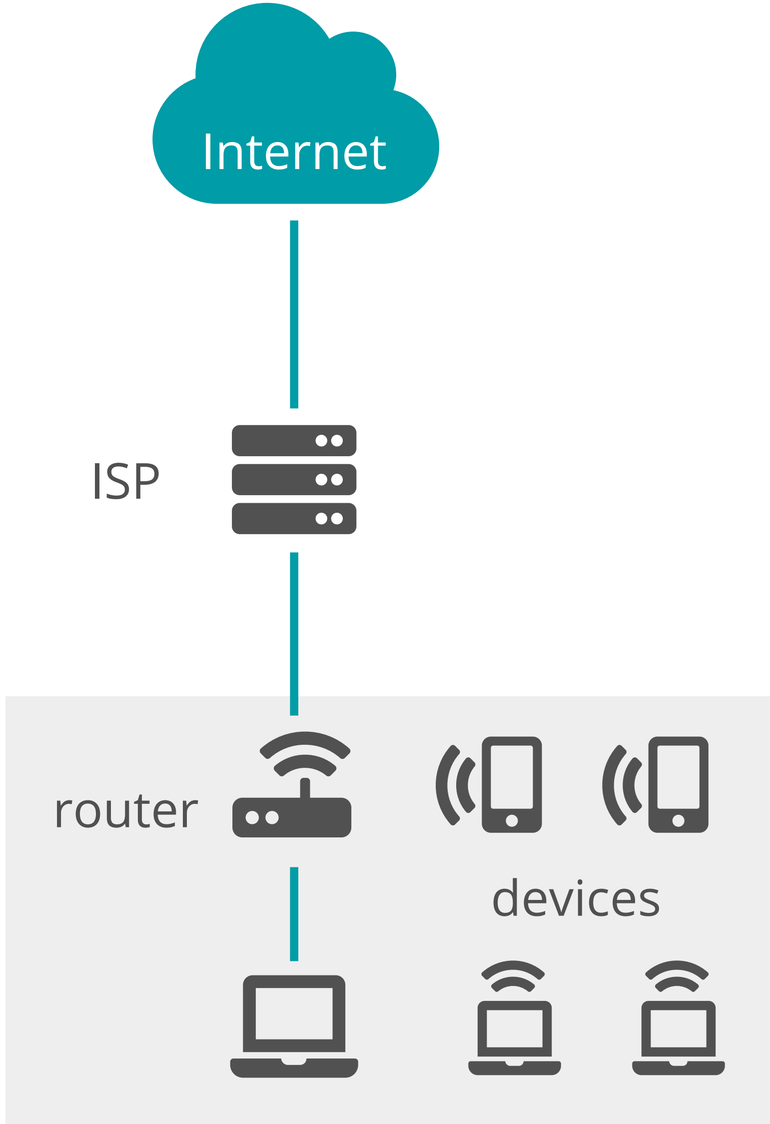

One assignment
= ‘ASSIGNED’

Using ASSIGNED

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

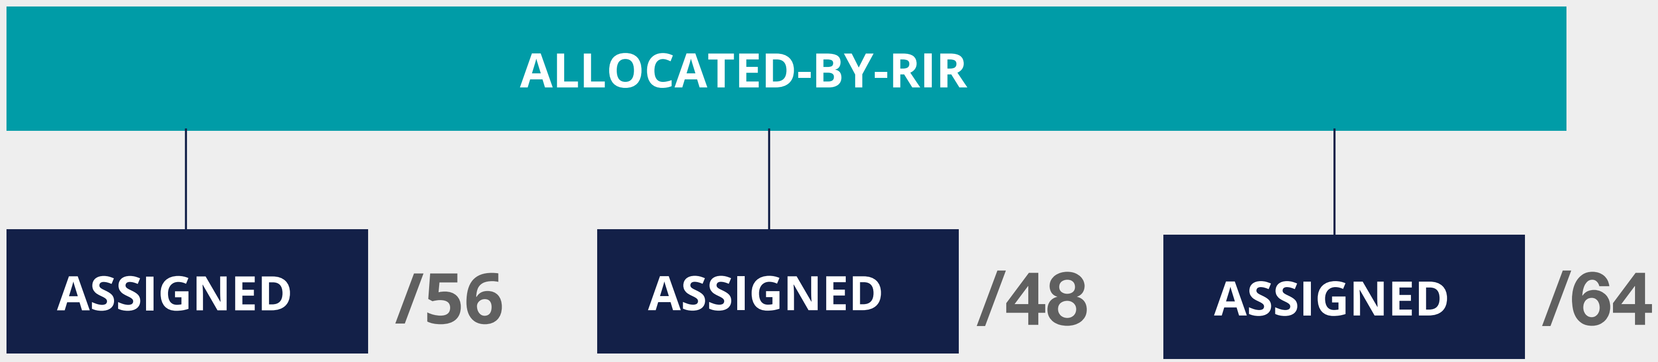

Using ASSIGNED - Example Object

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

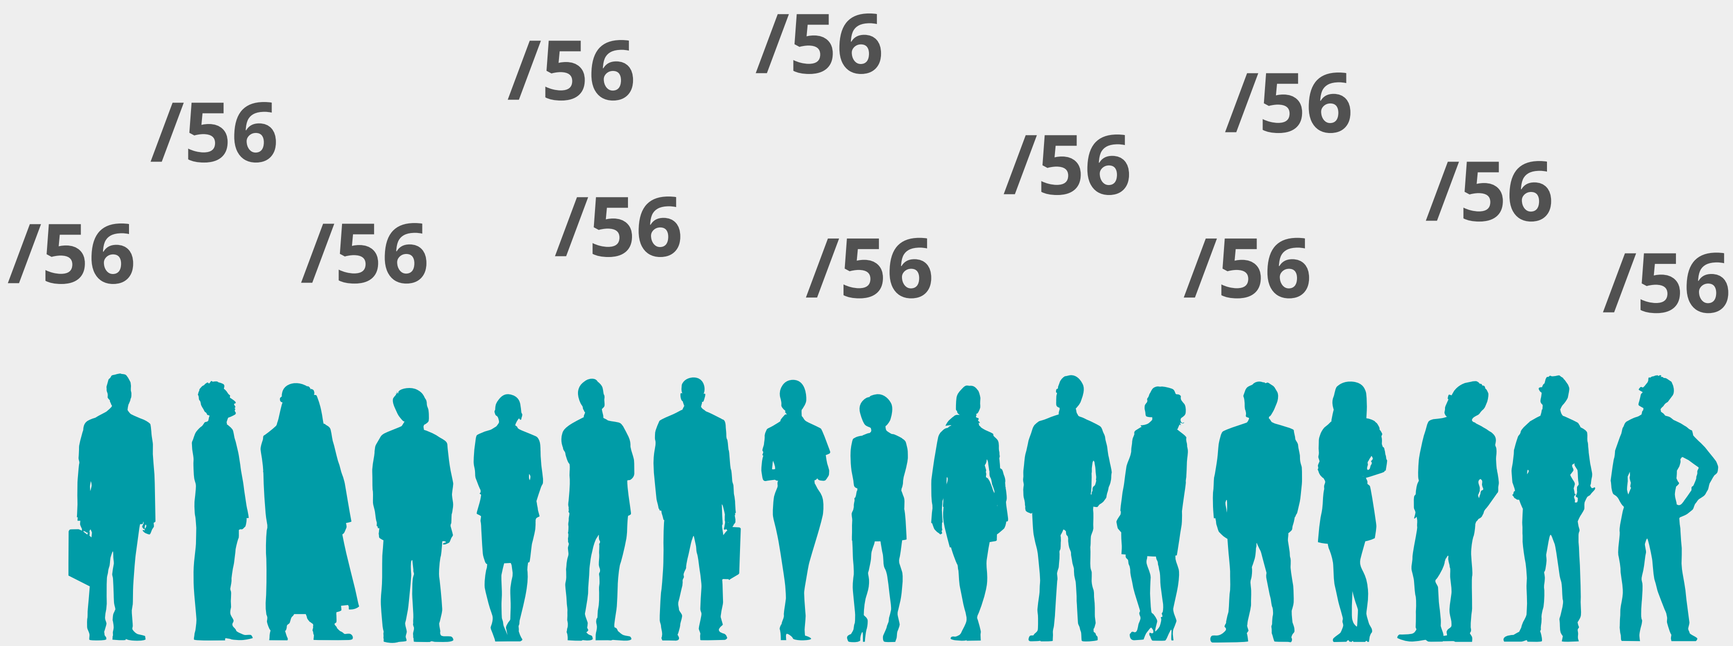

Using AGGREGATED-BY-LIR

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

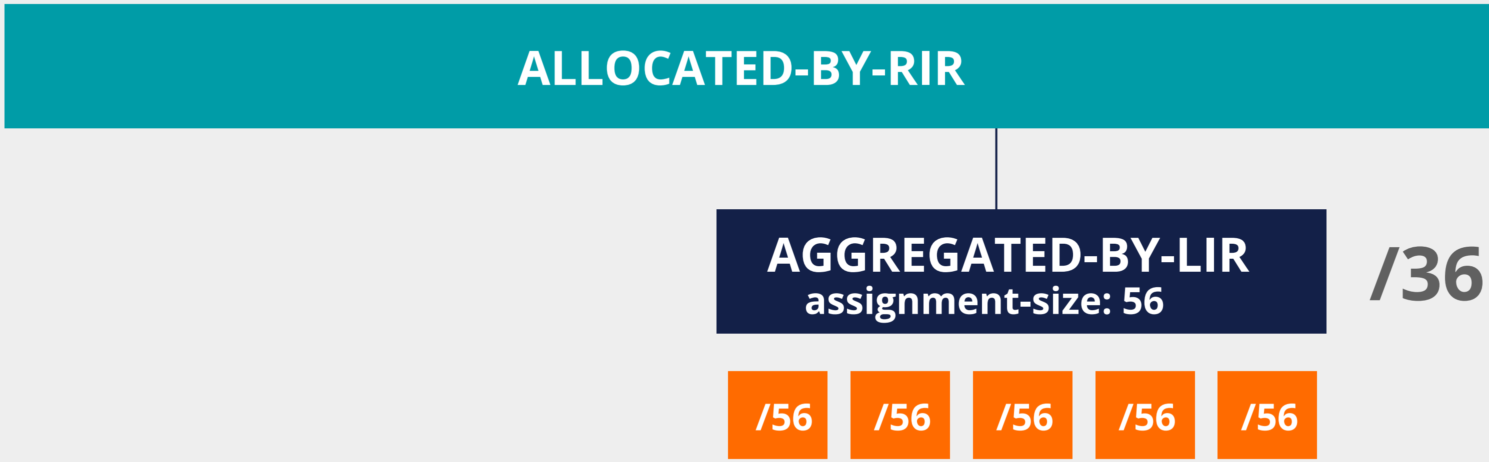

Using AGGREGATED-BY-LIR - Example

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

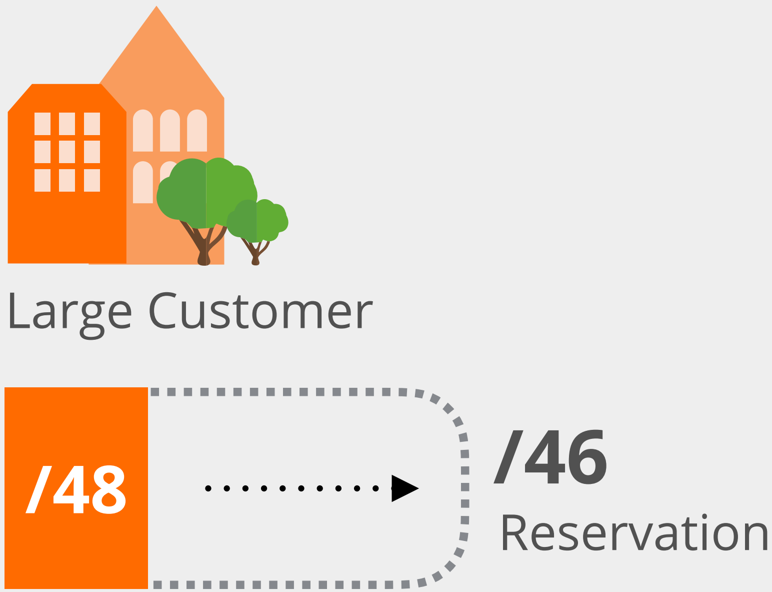

Branch office or department

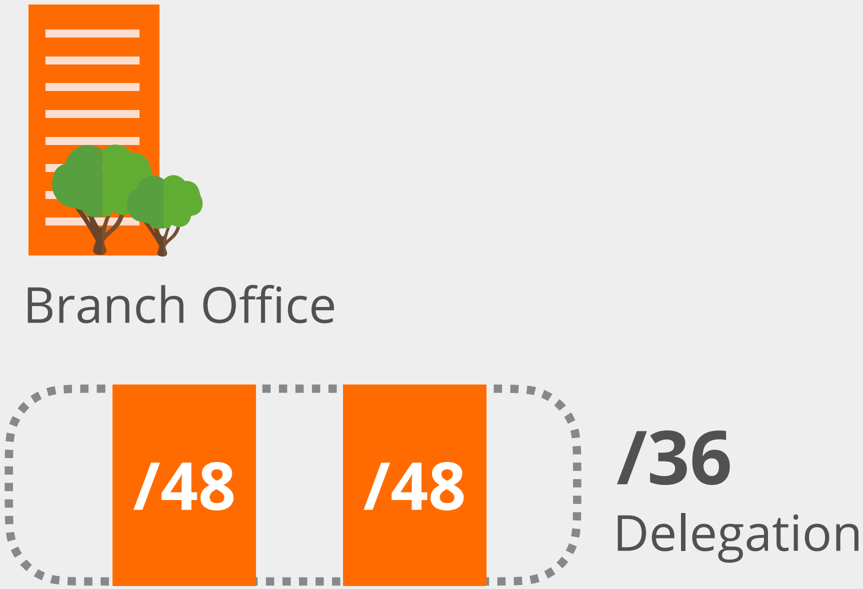

Using ALLOCATED-BY-LIR

Can be used for customers with **potential for growth**

- Or for your own infrastructure
- Or to delegate address space to a downstream ISP

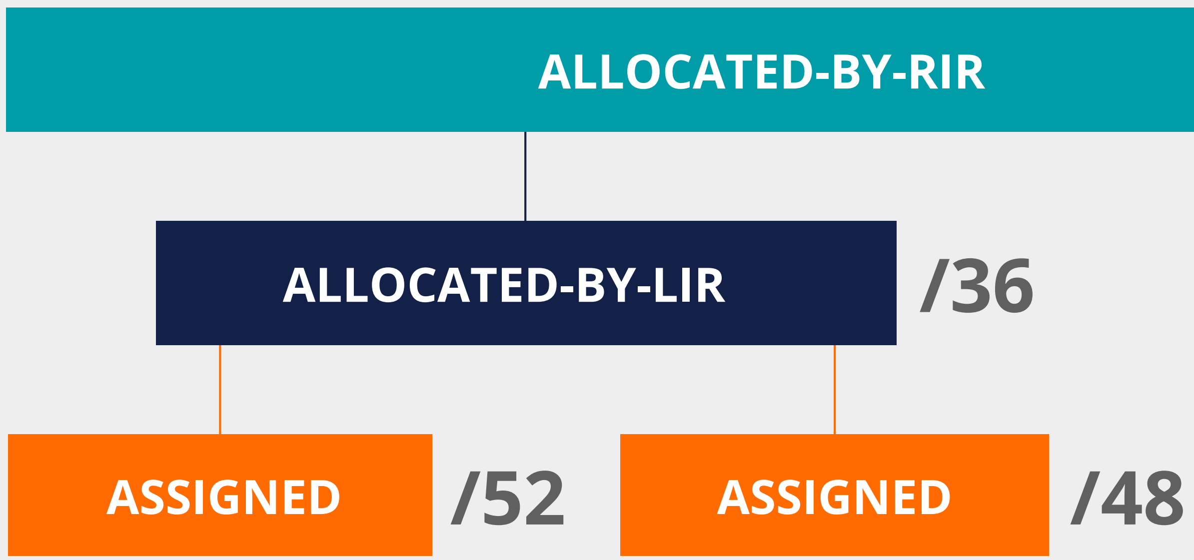

Using ALLOCATED-BY-LIR - Example

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

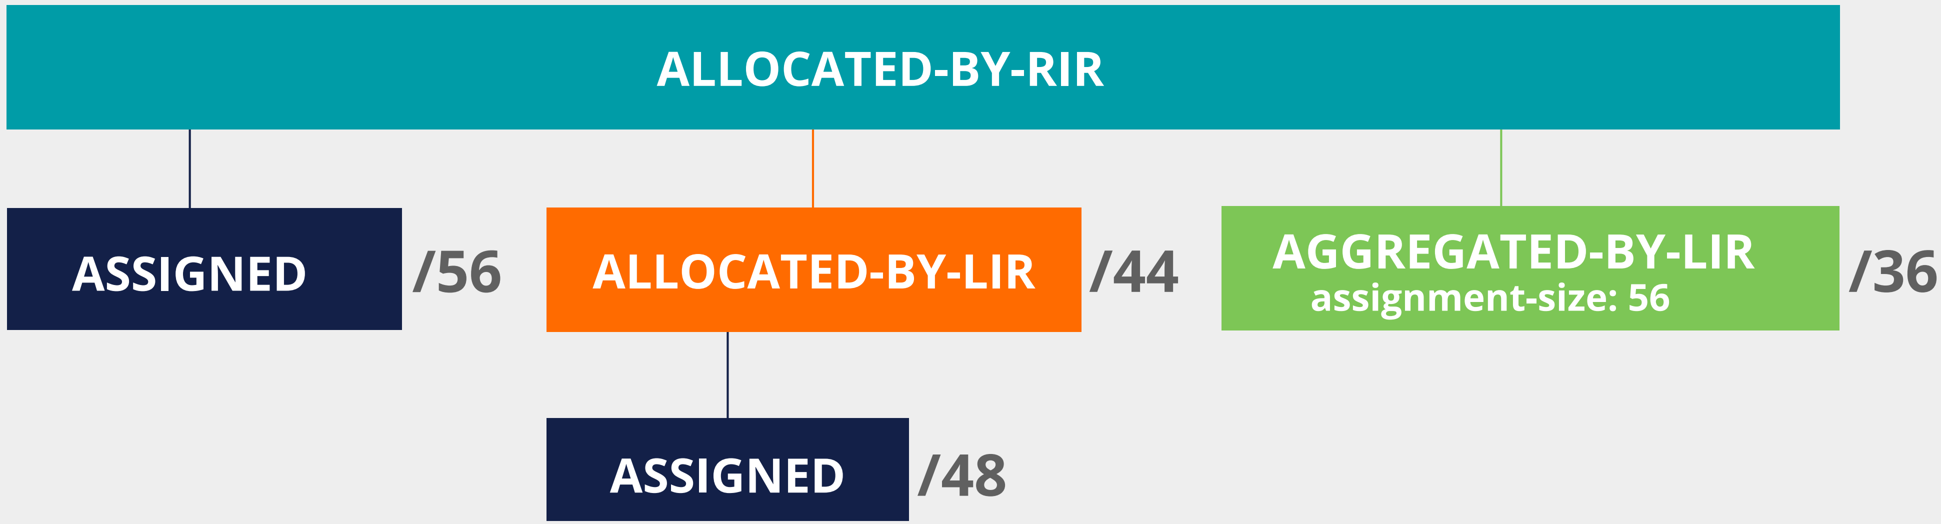
ALLOCATED-BY-RIR
ASSIGNED
/56
ALLOCATED-BY-LIR
/44
AGGREGATED-BY-LIR
assignment-size: 56
/36
ASSIGNED
/48

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 **cannot** be used for sub-assignments

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

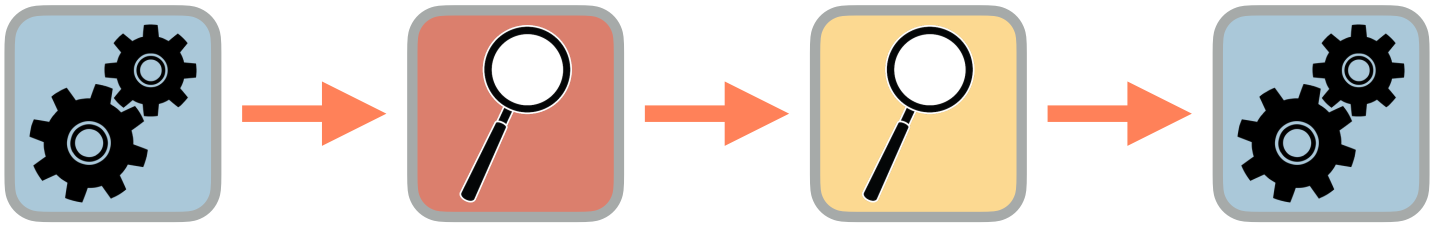

Making a Link-Local Address

- Interface ID is made from the MAC address

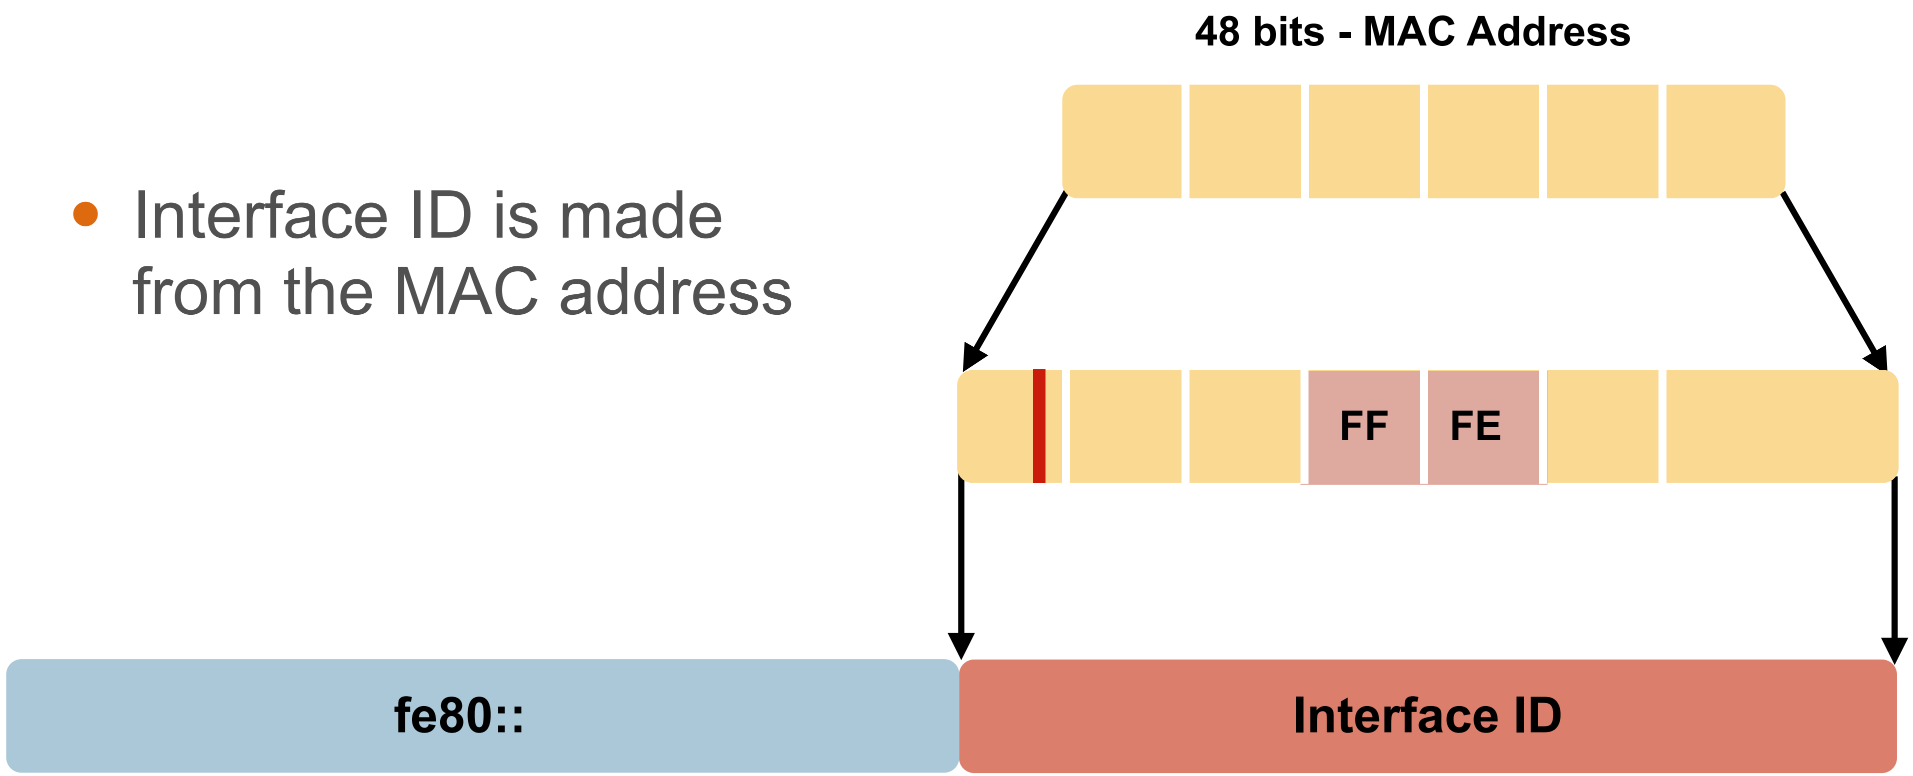

- fe80:: + Interface ID = Link-Local address for the host

Checking for Duplicates

Neighbor Solicitation

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

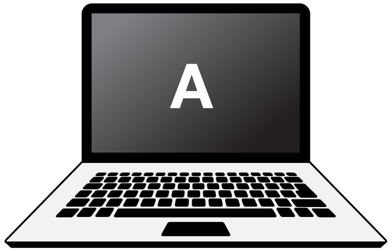

Neighbor Advertisement

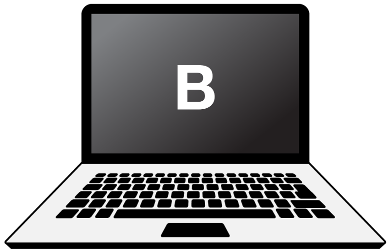

Hello! Yes, I'm using that IPv6 address.
My MAC address is 72:D6:0C:2F:FC:01

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

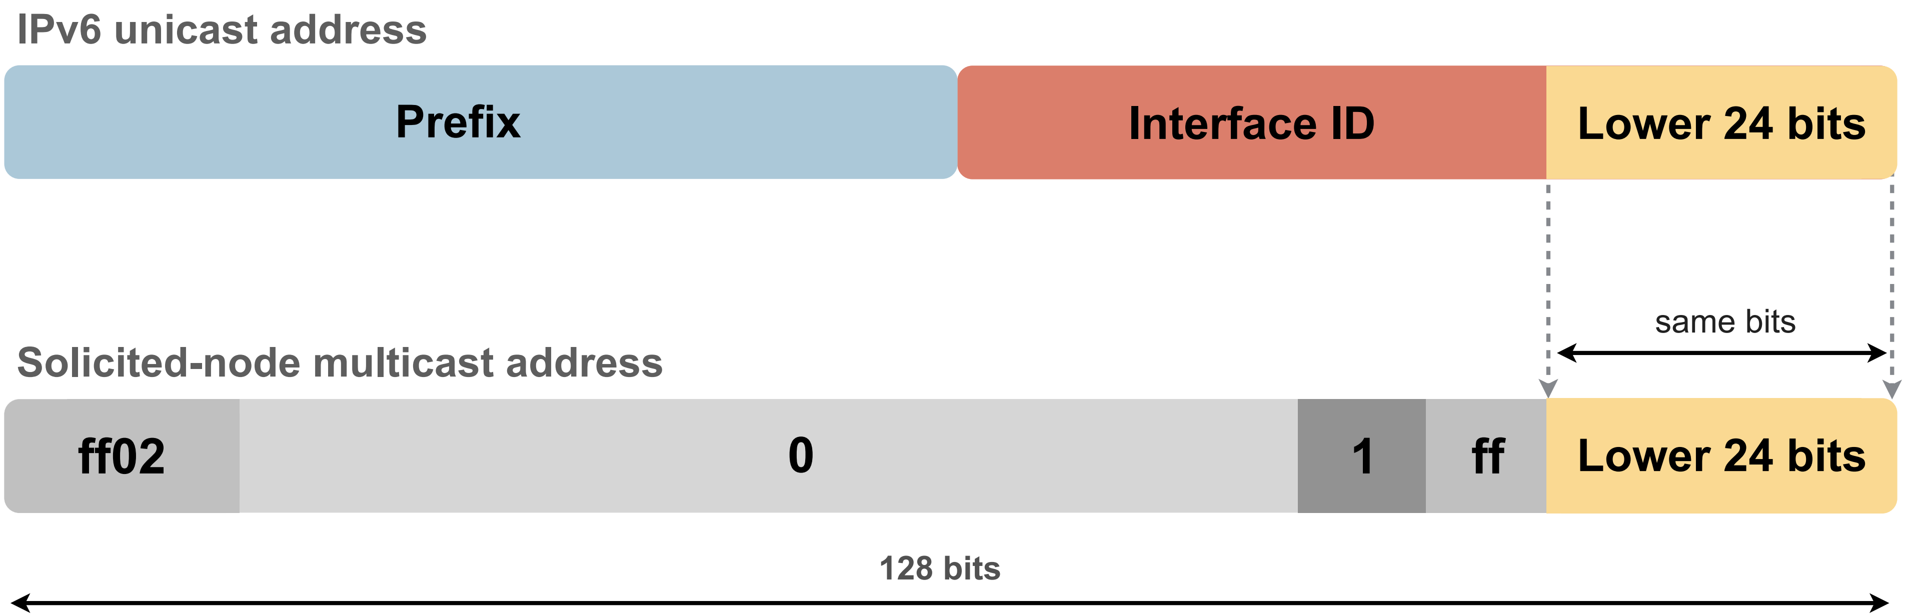

Solicited Node Multicast Address

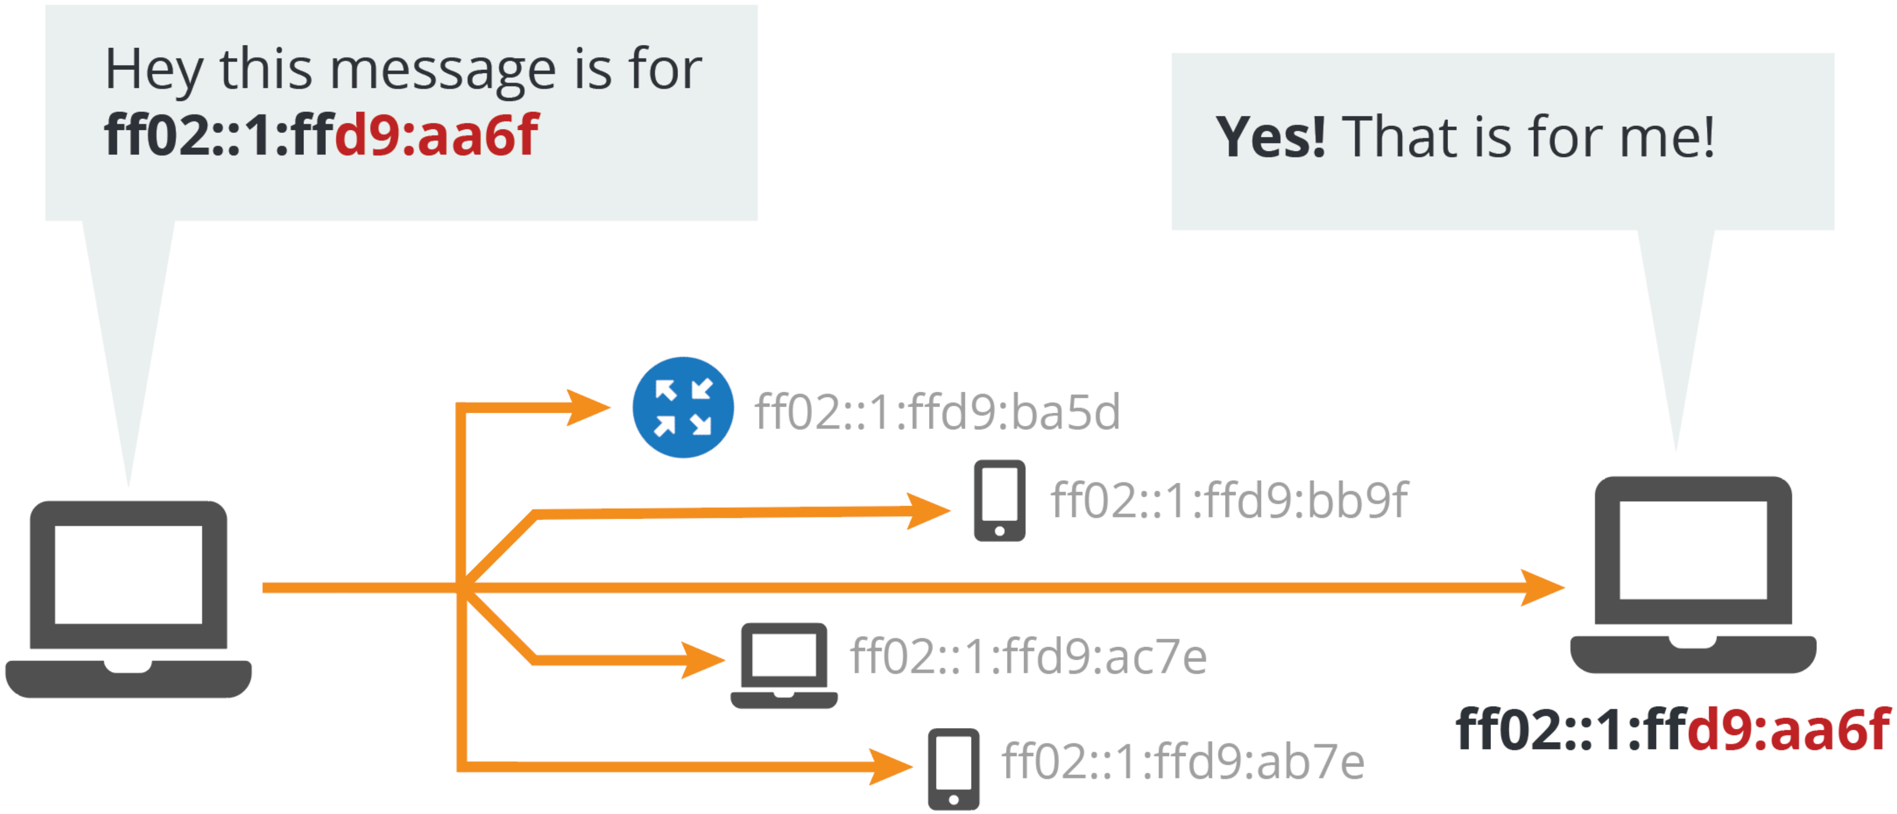

Searching for Routers

Router Solicitation

Hello! Is there a router out there?

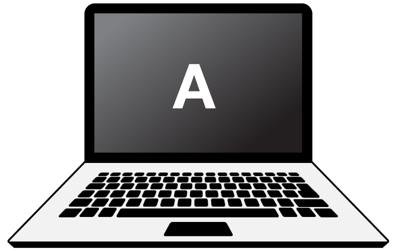

Router Advertisement

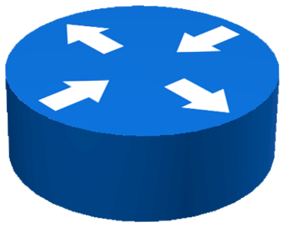

Hello! I'm a router and I have some information for you…

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?

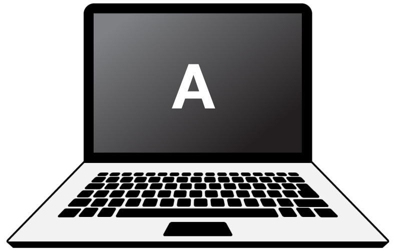

Neighbor Advertisement

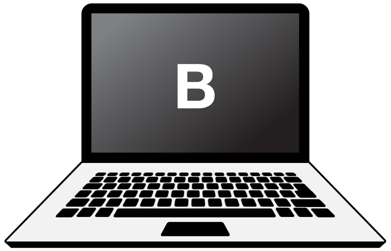

Hello! Yes, I'm still online.
My MAC address is 72:D6:0C:2F:FC:01

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

Redirects

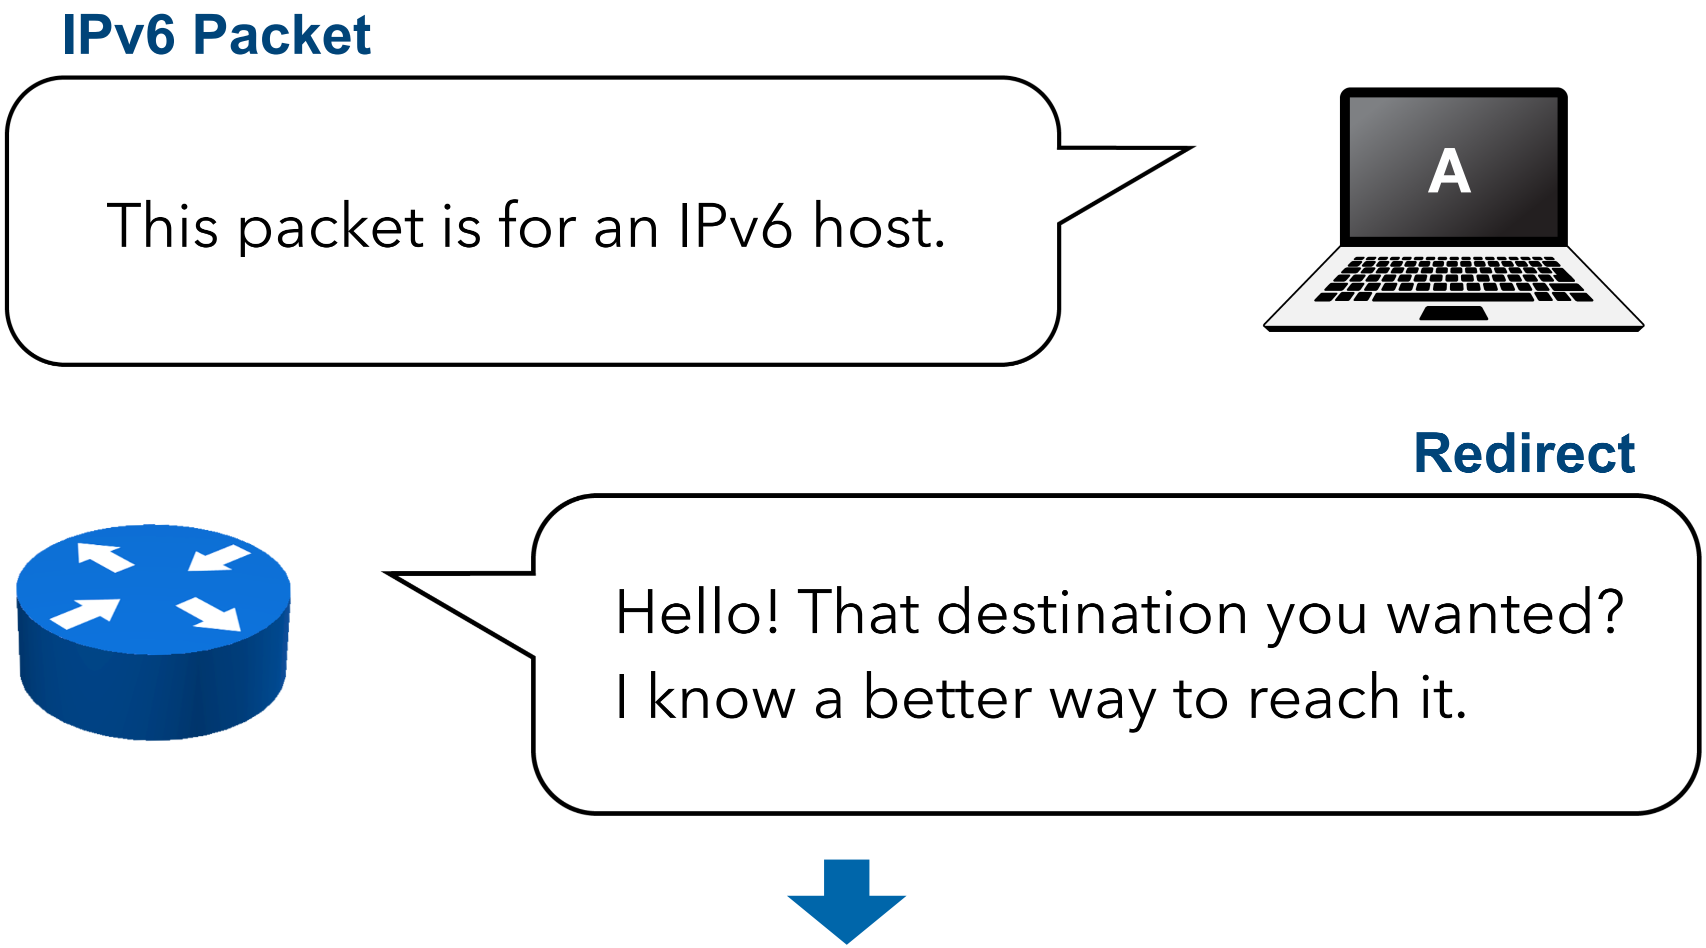

- Hosts can be redirected to a better first-hop router
- They can also be informed that the destination is a neighbor on the link

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 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

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

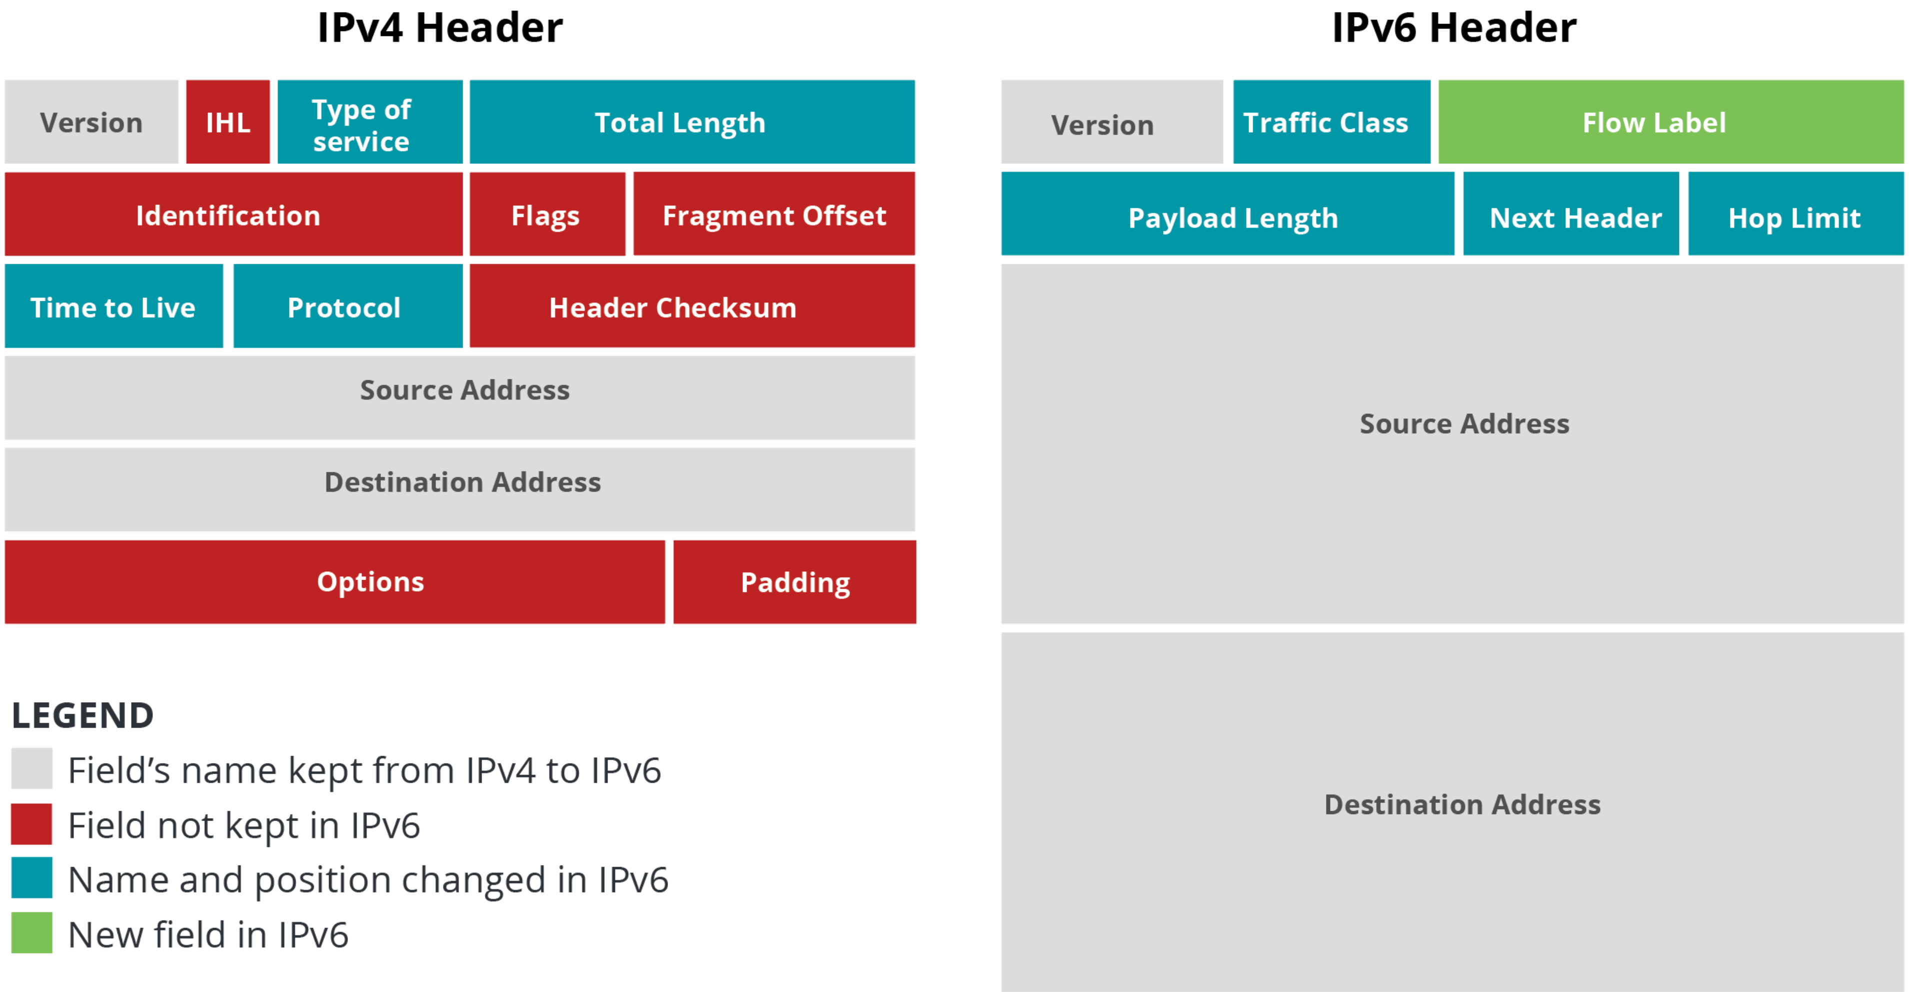

IPv4 Header
Version
IHL
Type of service
Total Length
Identification
Flags
Fragment Offset
Time to Live
Protocol
Header Checksum
Source Address
Destination Address
Options
Padding
IPv6 Header
Version
Traffic Class
Flow Label
Payload Length
Next Header
Hop Limit
Source Address
Destination Address
LEGEND
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

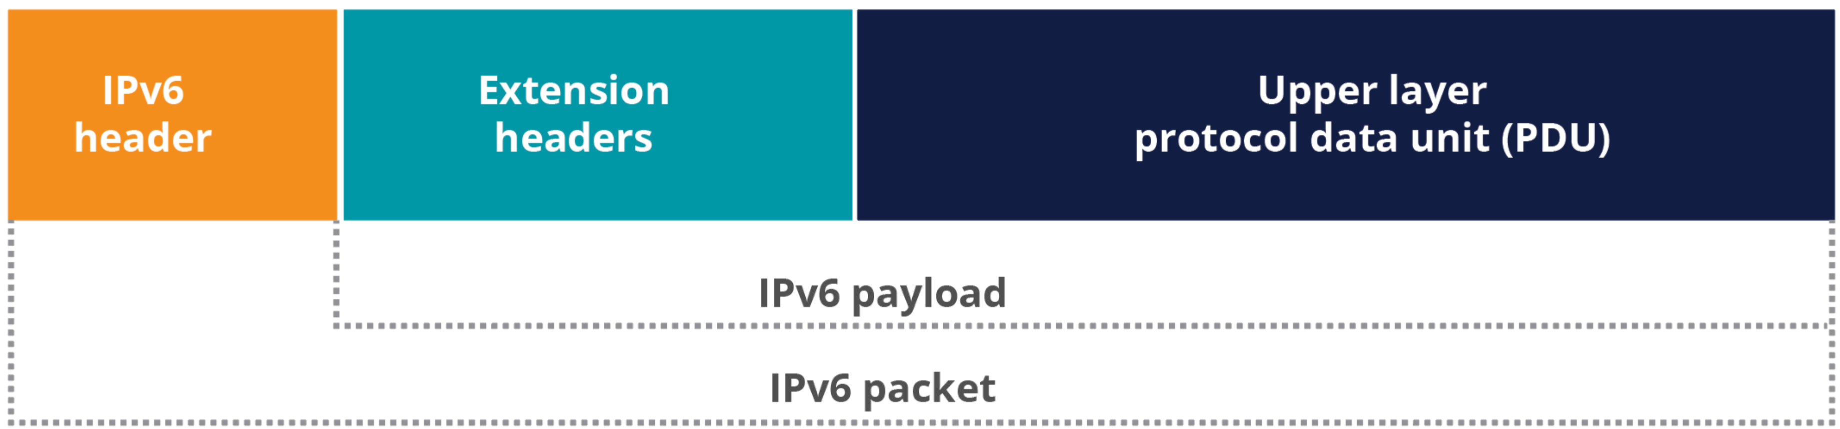

IPv6 Header

- Daisy-chained after the main header

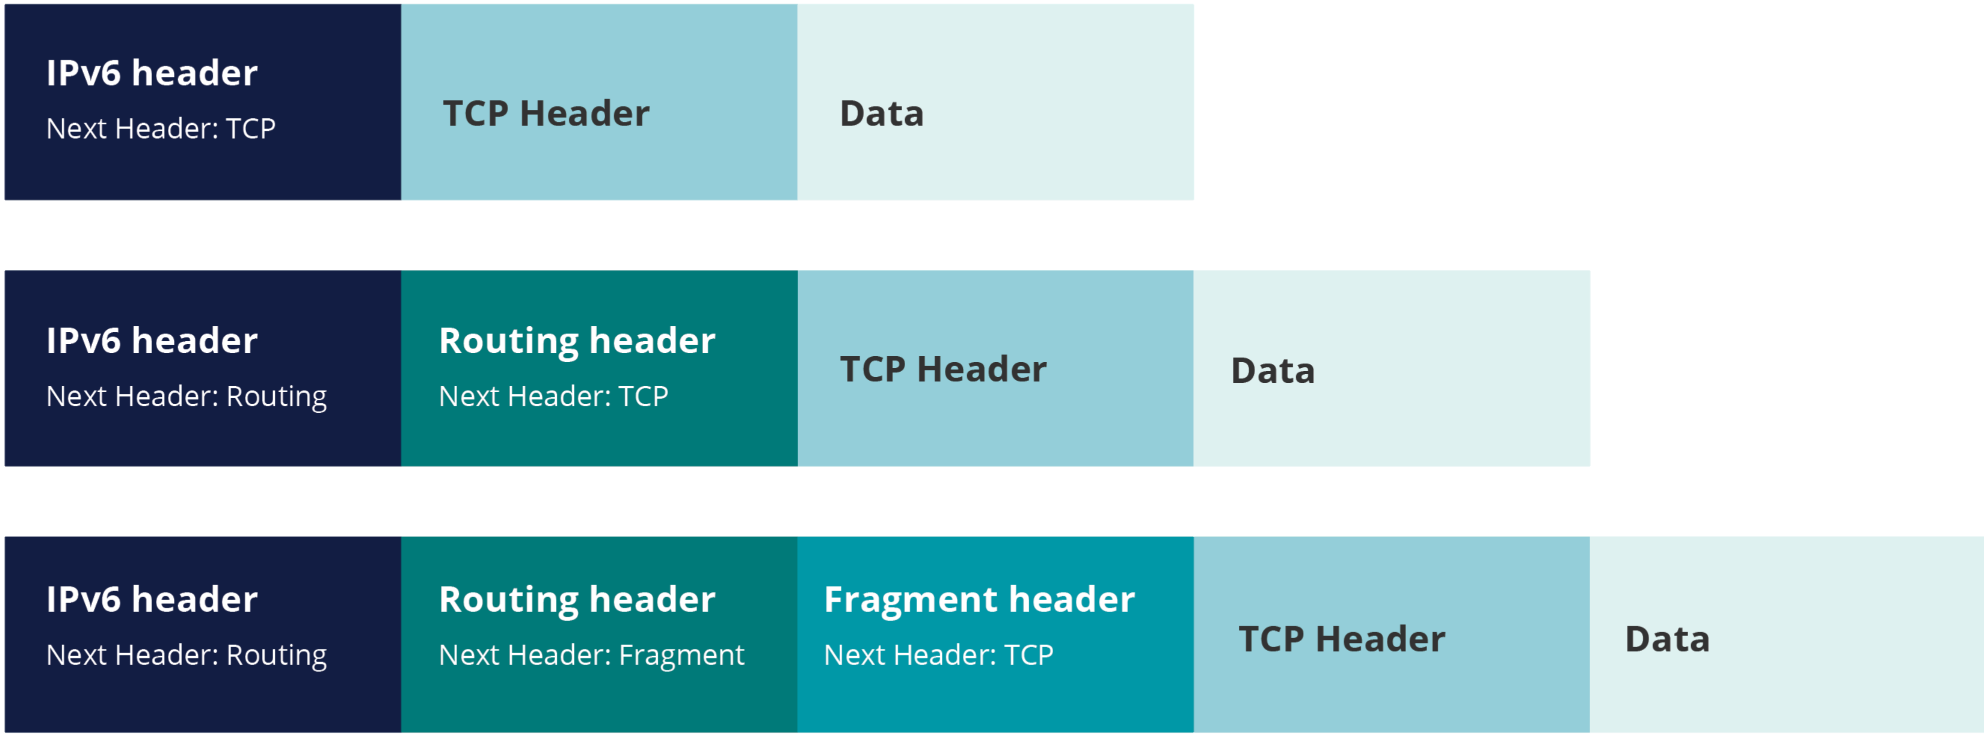

Fragmentation

- 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

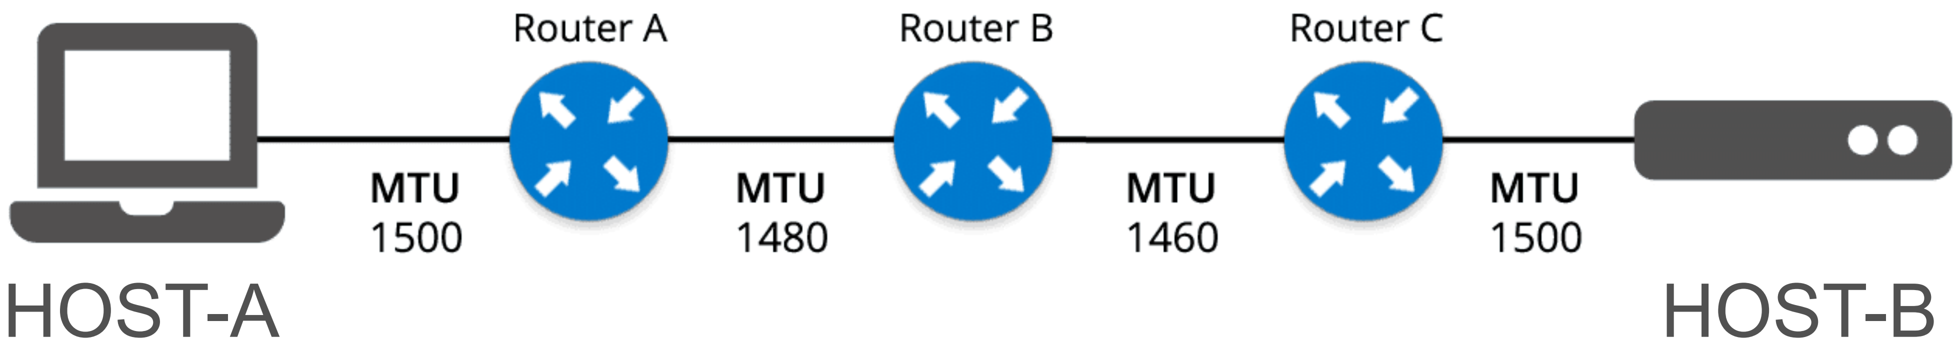

Ordering of Headers

- 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

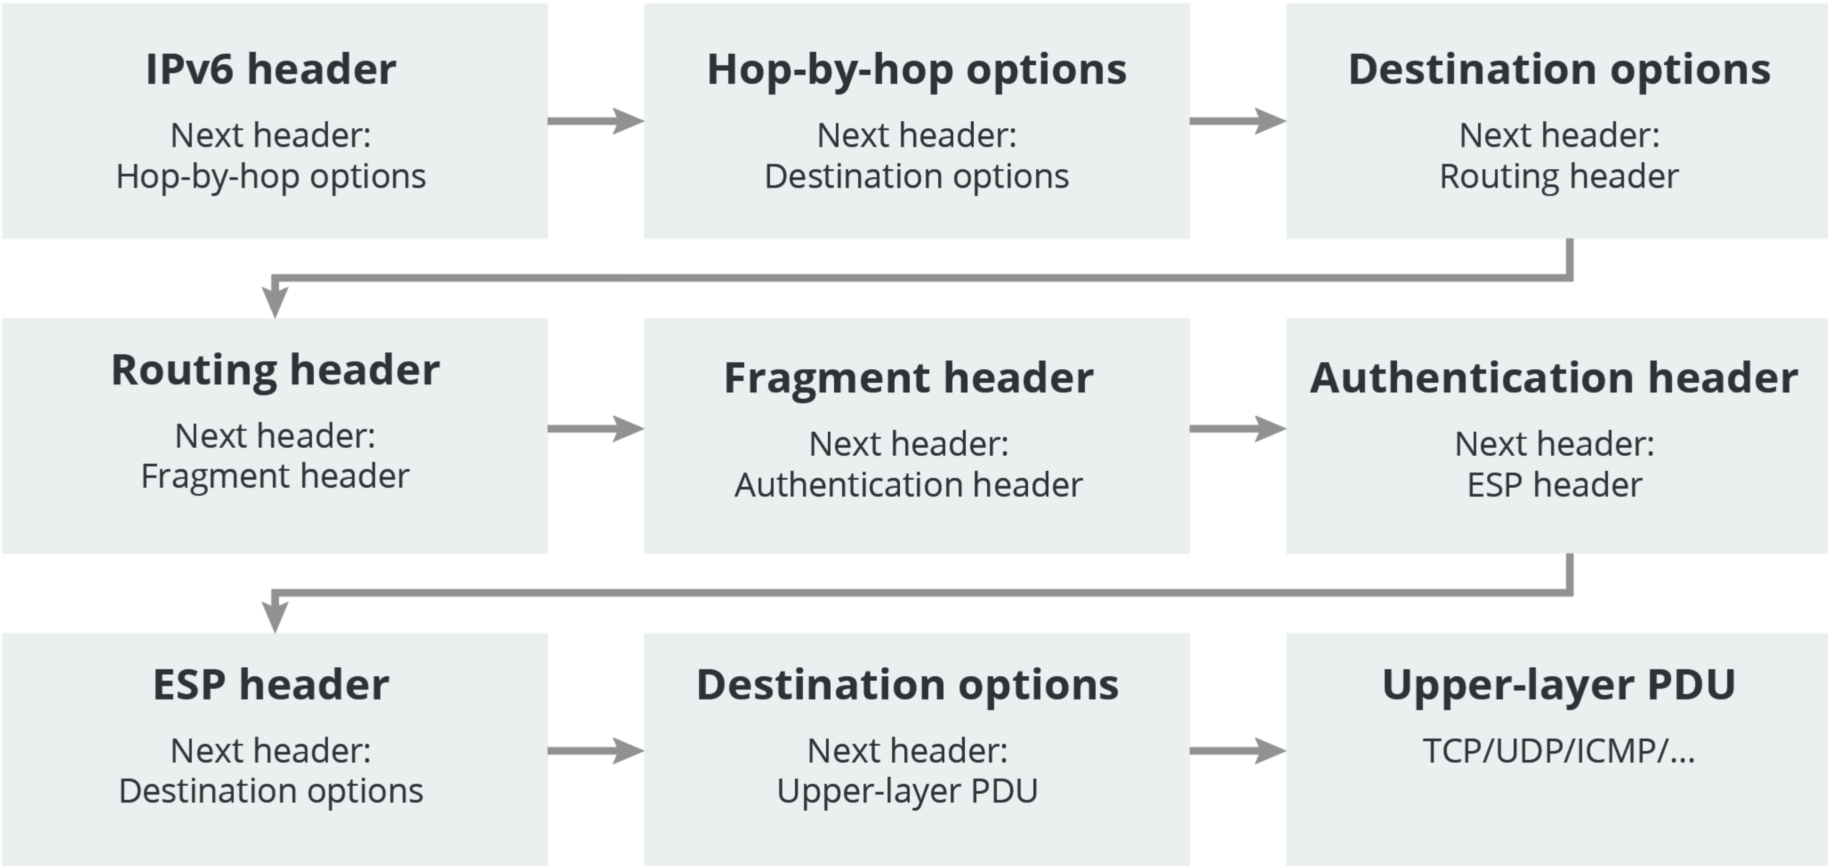

Broadcast

- 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

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)

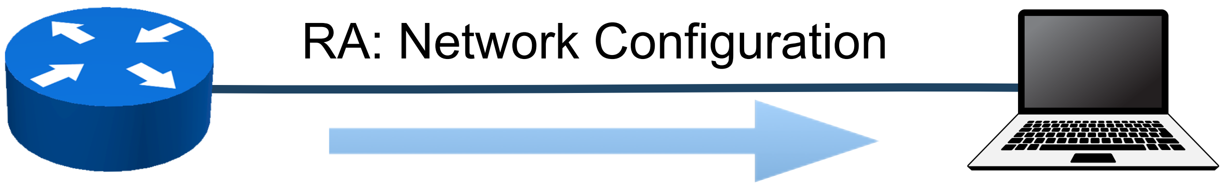

SLAAC IID Generation Options

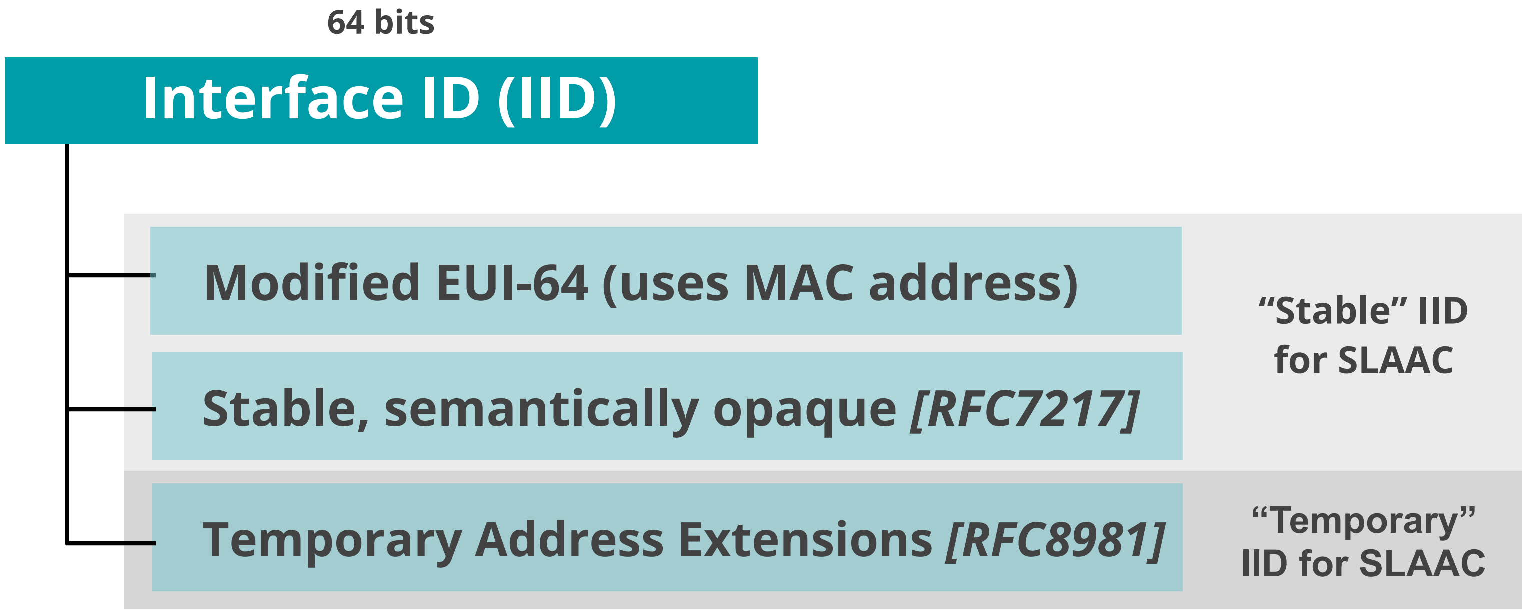

DHCPv6

- Used to give additional information like DNS servers or to manage the address pool

- Router Advertisement message contains hints
  - If “managed” flag = ‘1’ ⇒ 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)

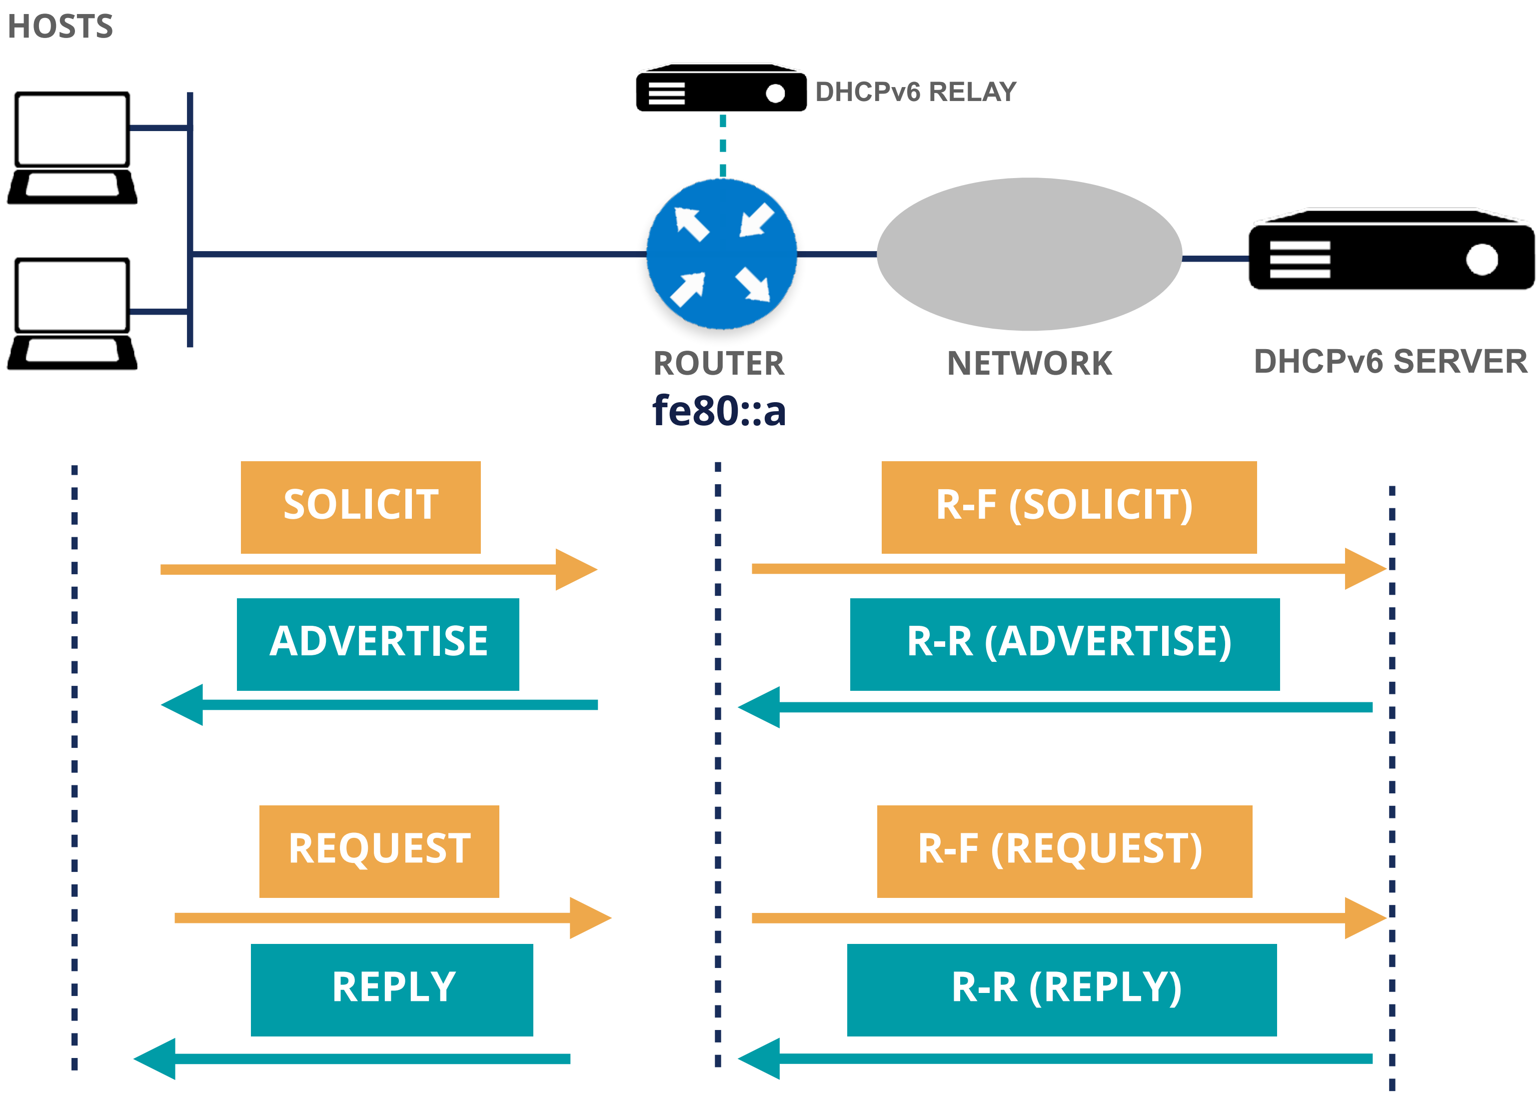
HOSTS
DHCPv6 RELAY
ROUTER
fe80::a
NETWORK
DHCPv6 SERVER
SOLICIT
ADVERTISE
REQUEST
REPLY
R-F (SOLICIT)
R-R (ADVERTISE)
R-F (REQUEST)
R-R (REPLY)

DHCPv6 (M=0, O=1)

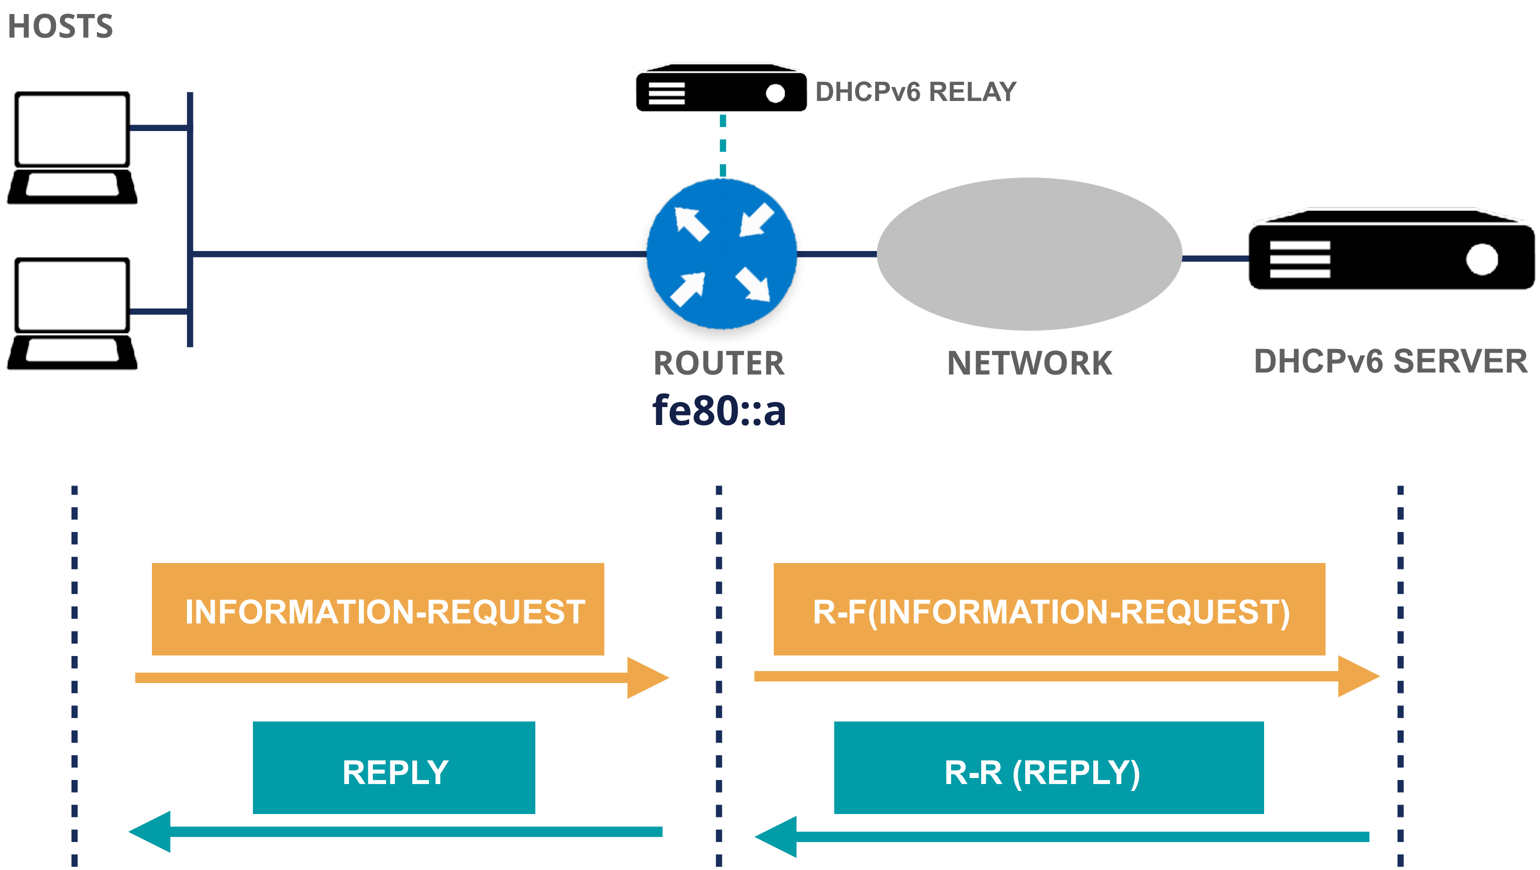

HOSTS
DHCPv6 RELAY
ROUTER
fe80::a
NETWORK
DHCPv6 SERVER
INFORMATION-REQUEST
R-F(INFORMATION-REQUEST)
REPLY
R-R (REPLY)

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

Tips

Section 8

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

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!


RIPE NCC
Academy
Learn something new today!
academy.ripe.net


RIPE NCC Certified Professionals
RIPE NCC Certified Professionals
IPv6 Fundamentals Analyst
RIPE NCC Certified Professionals
RIPE Database Associate
RIPE NCC Certified Professionals
BGP Security Associate
RIPE NCC Certified Professionals
IPv6 Security Expert
https://getcertified.ripe.net/

Questions


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