RIPE NCC
RIPE NETWORK COORDINATION CENTRE

Deploying IPv6-mostly access networks

IPv6-only and dual stack in one network

Ondřej Caletka | 12 December 2024 | INEX Member Meeting

The endless transition to IPv6

- IPv6 is *slowly* being deployed
- IPv4 is still the protocol of the Internet
- There are simply **not enough IPv4 addresses**
  - repurposing 240/4, 127/8 or 0/8 will not help
- There are many transition mechanisms, two of which are special:
  - **Dual stack:** Running both protocols at the same time
  - **NAT64:** Allowing limited access from unmodified IPv6-only hosts to IPv4 resources

The best transition mechanism

- IPv4-only and IPv6-only resources **directly accessible**
- IPv6 preferred for dual-stack resources
- Problems with IPv6 **masked** by Happy Eyeballs algorithm
- But it **does not address IPv4 scarcity**

Dual Stack

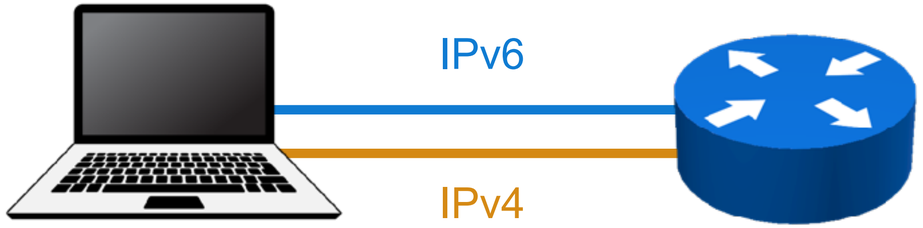

NAT64 allows IPv6-only networks

- IPv6 accessible natively
- IPv4 is translated into part of IPv6 address space
- Together with **DNS64**, everything seems to be **accessible over IPv6**
- **But sometimes you run into…**
  - IPv4 literals
  - Legacy software opening IPv4-only sockets
  - Dual-stack servers with broken IPv6

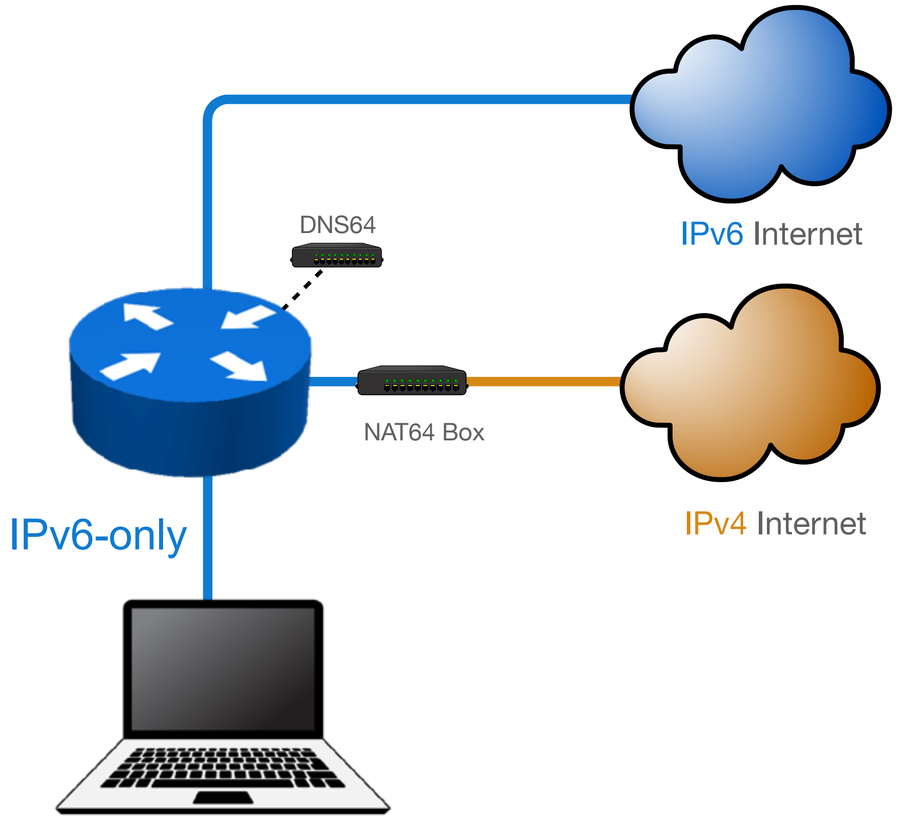

Mobiles are ready

- Apple forces all iOS apps to work well on IPv6-only networks with NAT64
- There is **Happy Eyeballs 2.0** for IPv4 literals or broken IPv6 on dual stack servers
- Finally **CLAT** is used for tethering to a computer
- Android uses just CLAT (464XLAT)
  - so IPv4 is accessible via two translations

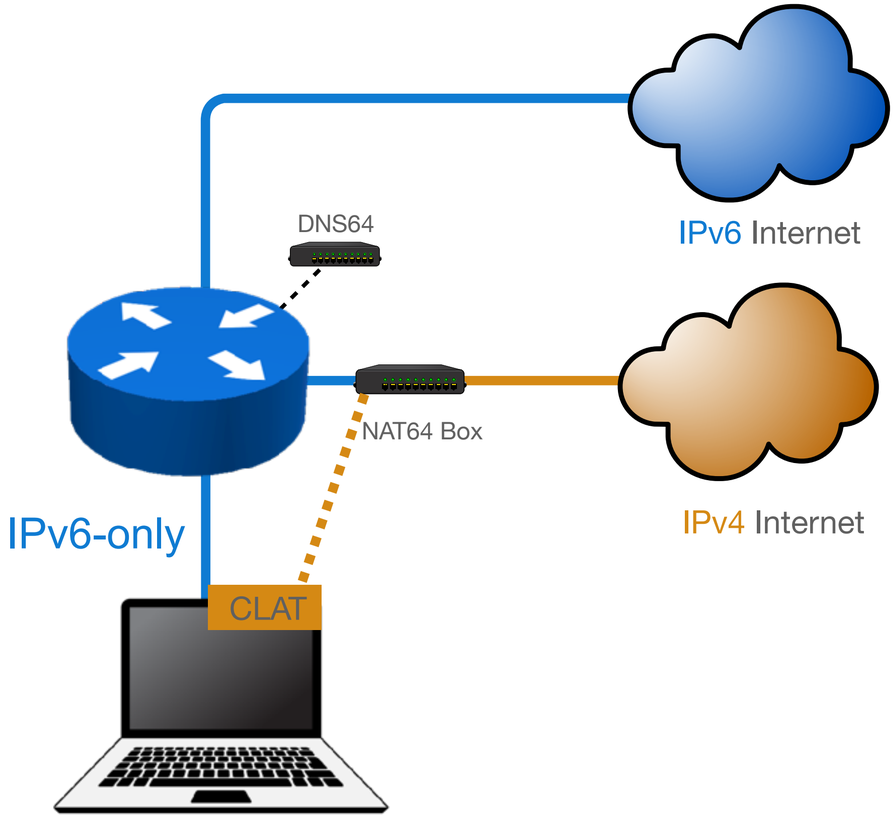

Desktops suffer on IPv6-only

- No Happy Eyeballs 2.0 implementation outside Apple
  - and even on Apple, only high-level APIs support it *(eg. Safari, not Firefox)*
  - Chrome has “Use NAT64 translation for IPv4 literals” feature
- **No CLAT** in Windows, Linux or ChromeOS
- Well known small problems:
  - Legacy applications using IPv4-only sockets
  - IPv4 literals do not work *(except Chrome)*
  - Dual-stack servers where IPv6 is broken do not work
  - Legacy Happy Eyeballs **doesn’t help** since there's no IPv4 to fall back to
  - Most corporate VPNs do not work (often *just* a configuration issue)

Can my device work on IPv6-only?

Fully

Android

iOS

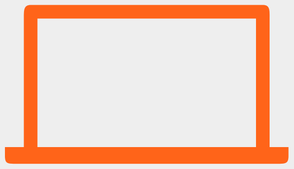

macOS

- CLAT is present*
- Some mobile networks run billions of IPv6-only phones for years already

Mostly

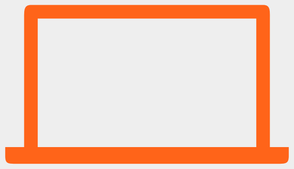

Windows

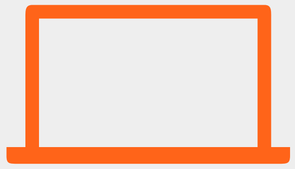

Linux

- No CLAT*
- Applications relying on IPv4 are **broken**

No way!

IoT

Smart home

- No IPv6 support*
- Native IPv4 **required**

***) some statements are simplified**

Can we do IPv6-only?

At least for a subset of devices?

IPv6-only Preferred option of DHCP

(RFC 8925)

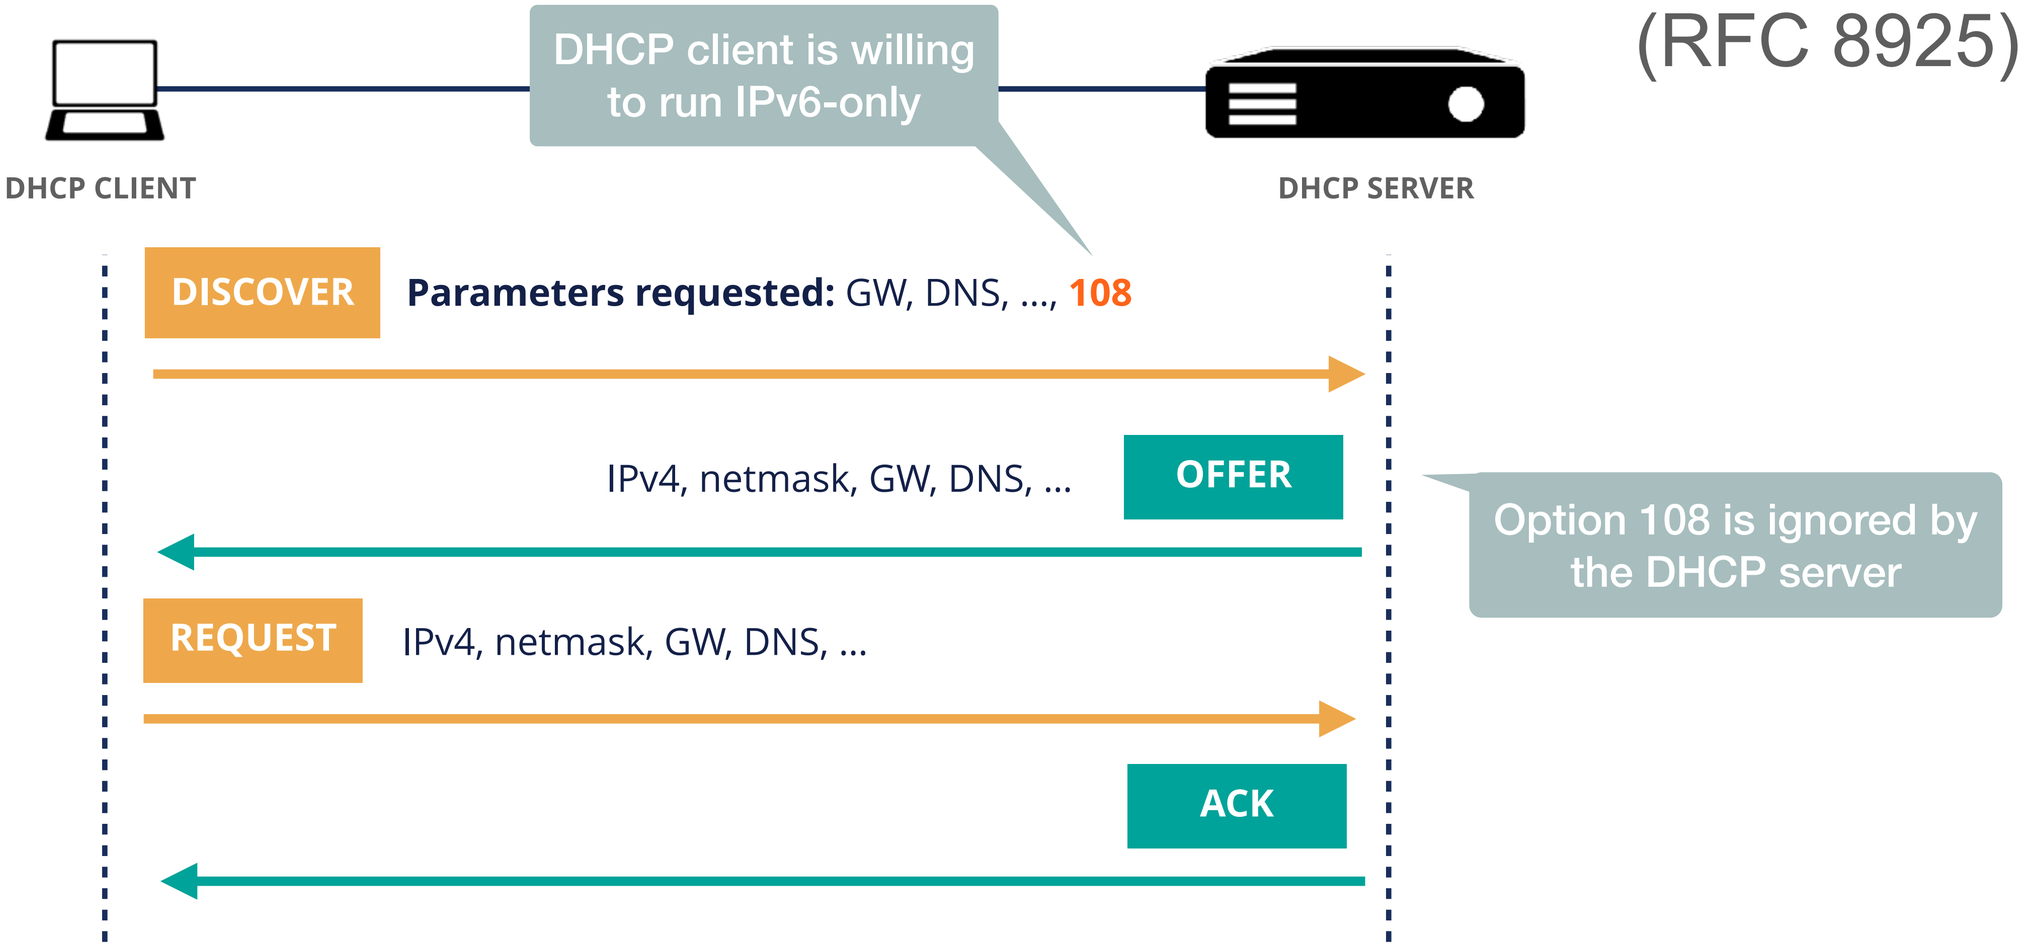

Using DHCP to turn IPv4 off

(RFC 8925)

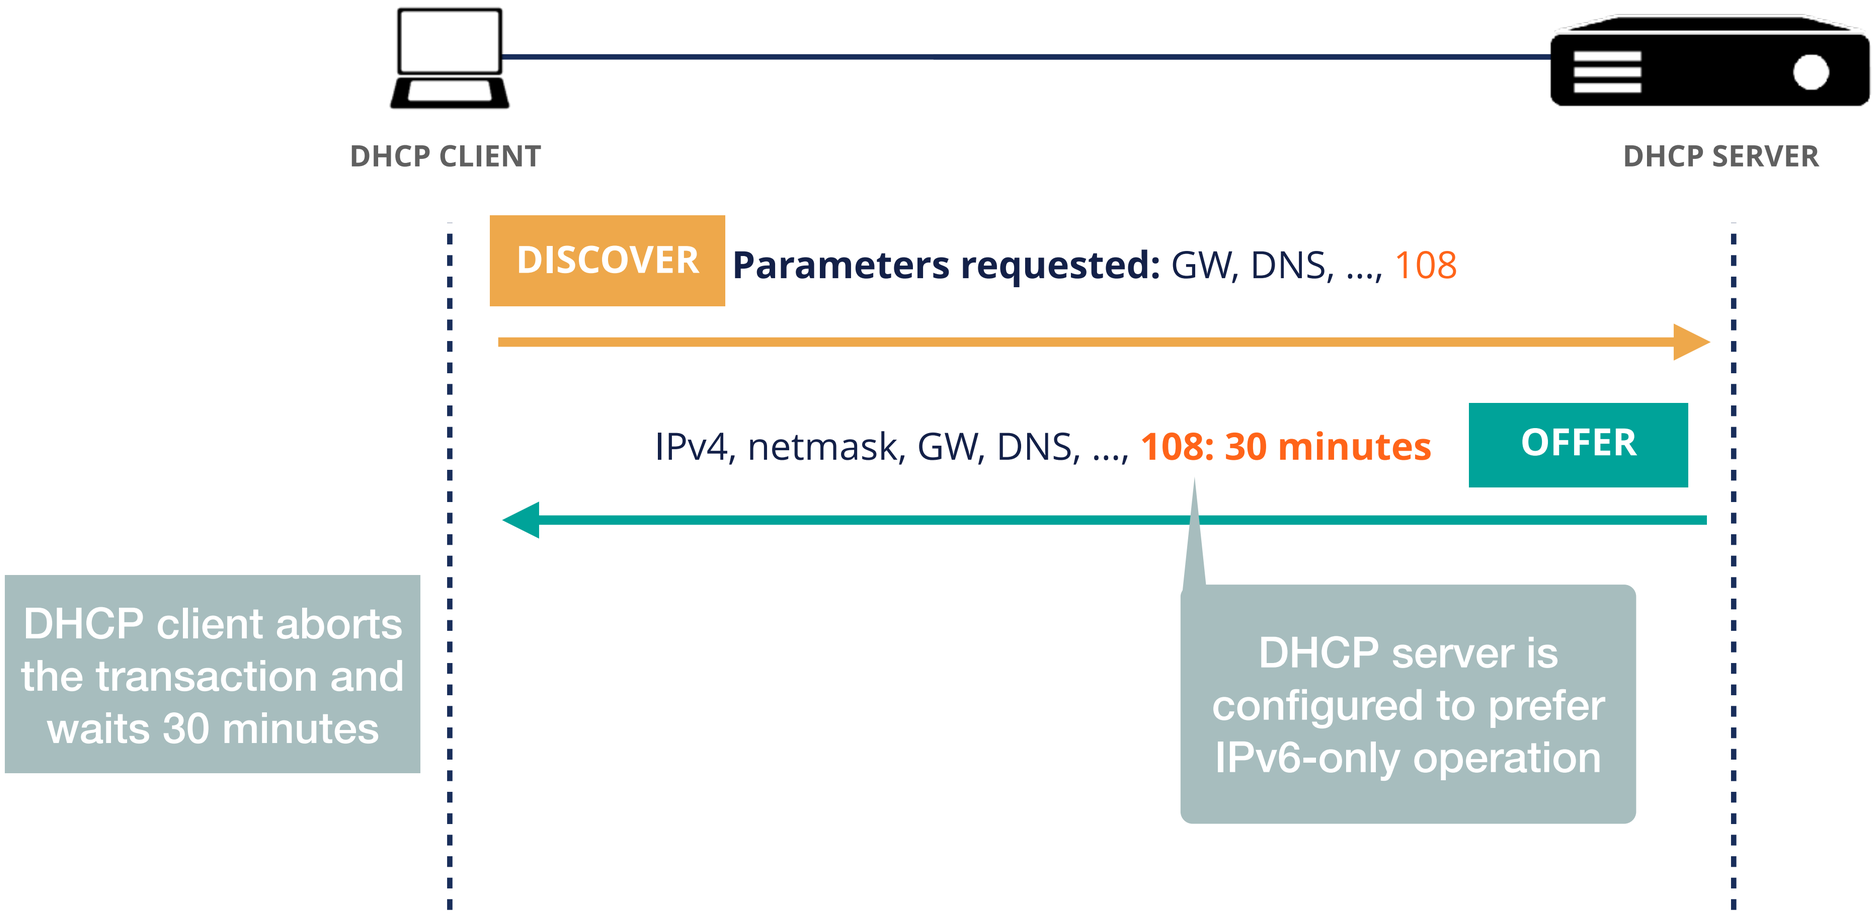

Is DHCP option 108 already deployed?

You bet! Option 108 is requested by recent:

Android	iOS	**macOS**

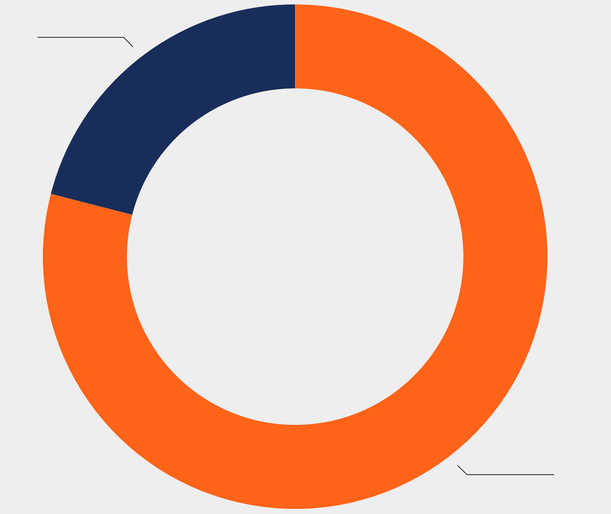

Devices are ready, networks are lagging behind.

PREF64 RA Option

- A Router Advertisement option **carrying NAT64 prefix**
- Needed for **CLAT configuration**, local DNS64 or Happy Eyeballs 2.0 *(dealing with IPv4 literals)*
- **Shares fate** with other configuration parameters
  - can be trusted **a bit more** than DNS64
- Supported by recent Android, iOS and macOS

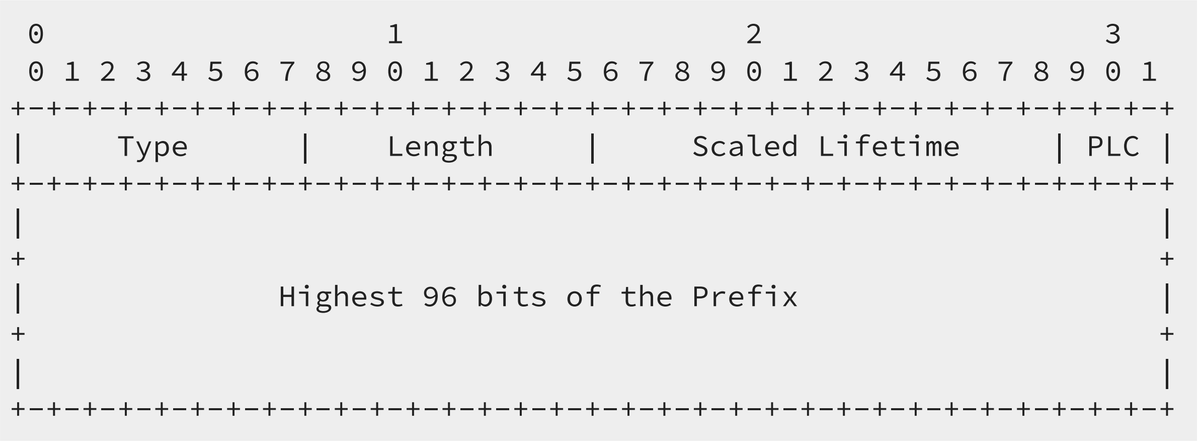

Running IPv6-mostly

DHCP option 108 is easy

- **Native support** in the latest Kea
- Most DHCP servers support defining **custom options**
  - for instance: `dnsmasq -O 108,0:0:1:2c`
  - the option value represents duration for which the IPv4 stack should be disabled
- **No special processing** on the DHCP server side is *required*
- But there **have to be free addresses** in the IPv4 address pool
  - Otherwise the DHCP server will not respond

PREF64 RA option is harder

- No **custom RA option** support in routers
  - We already **had this issue** with Recursive DNS Server option, now we **have it again**
  - Router vendors should really implement **custom options** similar to DHCP
- Adoption is *slowly* increasing:
  - `radvd` (merged but unreleased)
  - `odhcpd` (used in OpenWRT)
  - rad (part of OpenBSD)
  - MikroTik RouterOS v7.8 beta2
  - BIRD 2.14 (as a part of **custom router advertisement** option)

Summary

Pros

- **Only one network** to join
- **No waste of IPv4** addresses for every single device
  - Cool if you **don't use NAT**
- Lower reliance on DHCP

Cons

- Deployment is **still dual-stack**
- NAT64 is **needed**
- **Problematic** interoperability between dual-stack and IPv6-only hosts within the network
  - Setting up a Chromecast from an Android phone is *impossible*

When to consider IPv6-mostly

- You **don't use NAT** and your DHCP pool is filling up
- You do use NAT but are **running out** of private addresses
- There are mostly **mobile** or **Apple devices** in your network
- You **already have NAT64** in place and want to gradually undeploy IPv4

An RFC draft is **being written**: draft-ietf-v6ops-6mops-00

Phased IPv6 transition

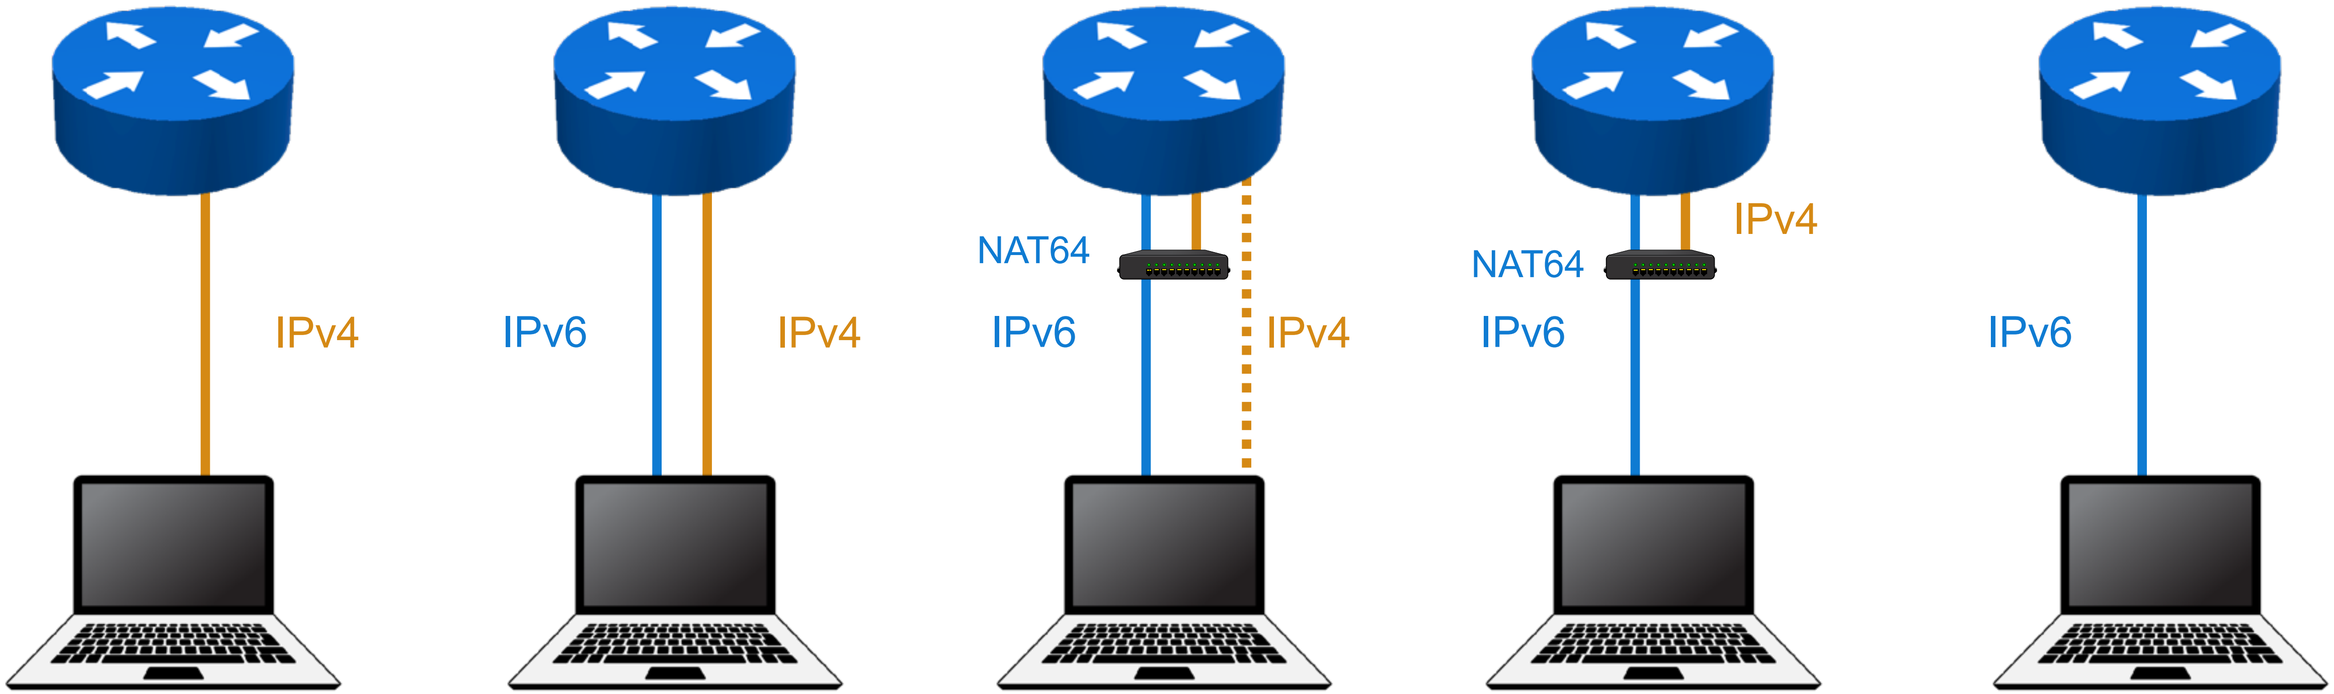

IPv6-only support on other platforms

- **Windows 11** pledged to implement CLAT and DHCP Option 108
- Linux has **some support** for DHCP Option 108
  - NetworkManager
  - systemd-networkd
  - dhcpcd
- Proper CLAT for Linux is **still missing**
  - Not supported by the kernel itself
  - eBPF might be the way forward

Questions

Ondrej.Caletka@ripe.net
https://Ondřej.Caletka.nl