Mission ~~Im~~Possible

Turning IPv4 Off in an Enterprise Network

Jen Linkova, furry@google.com

Image source: https://freesvg.org/magical-unicorn

Motivation

Running out of **private** IPv4 addresses

Dogfood and testing

Dual stack is hard

Source: www.wikipedia.org

"Entities should not be multiplied without necessity."

William of Ockham

Dual-Stack (IPv4 + IPv6) Network

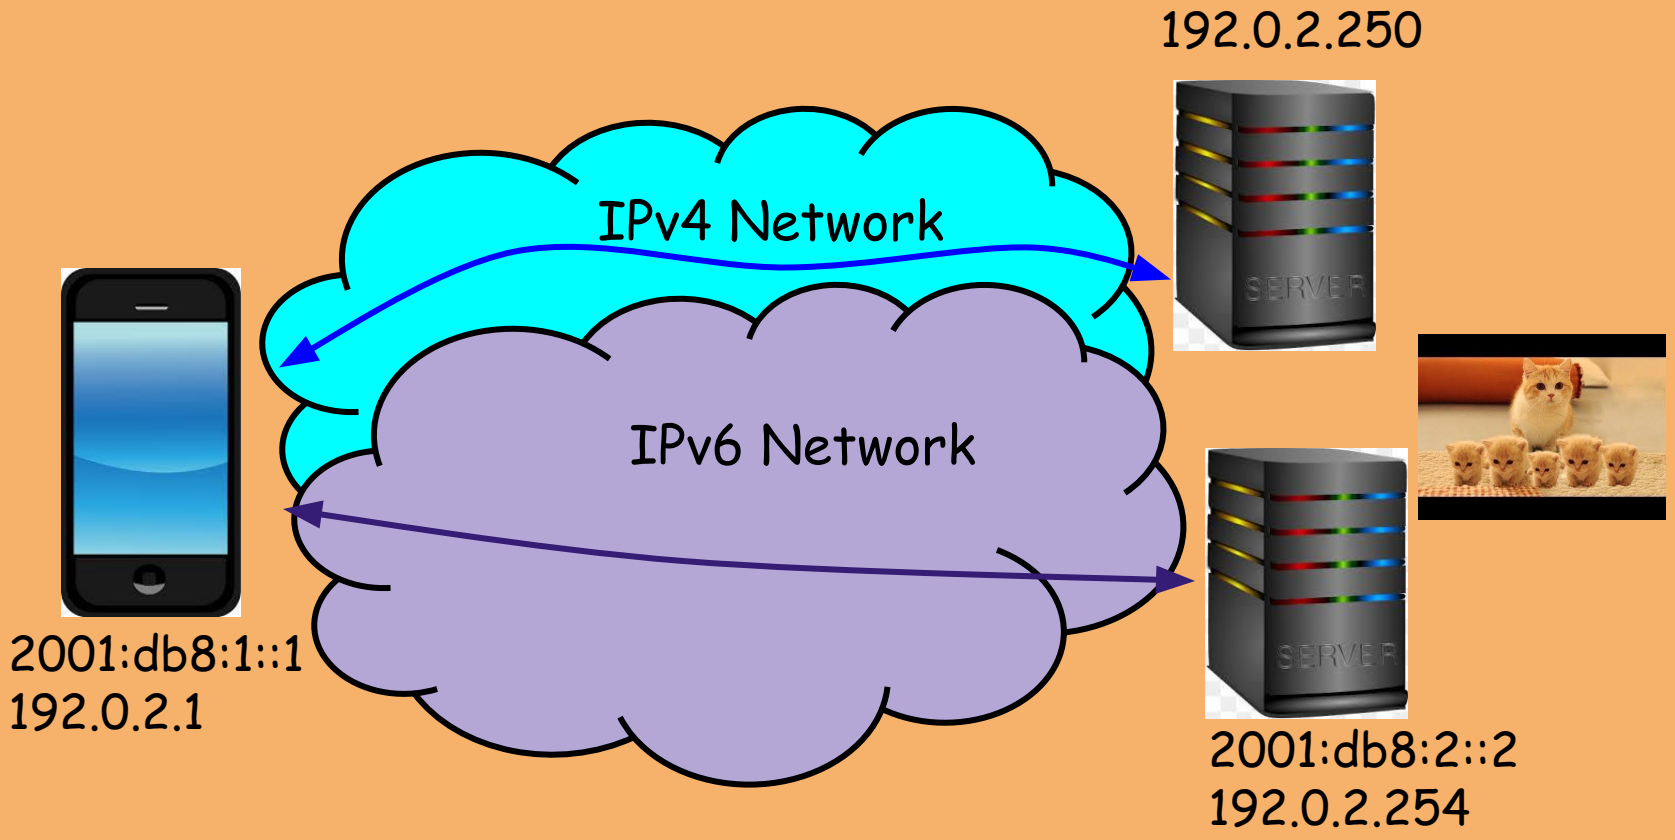

How Can We Remove IPv4?

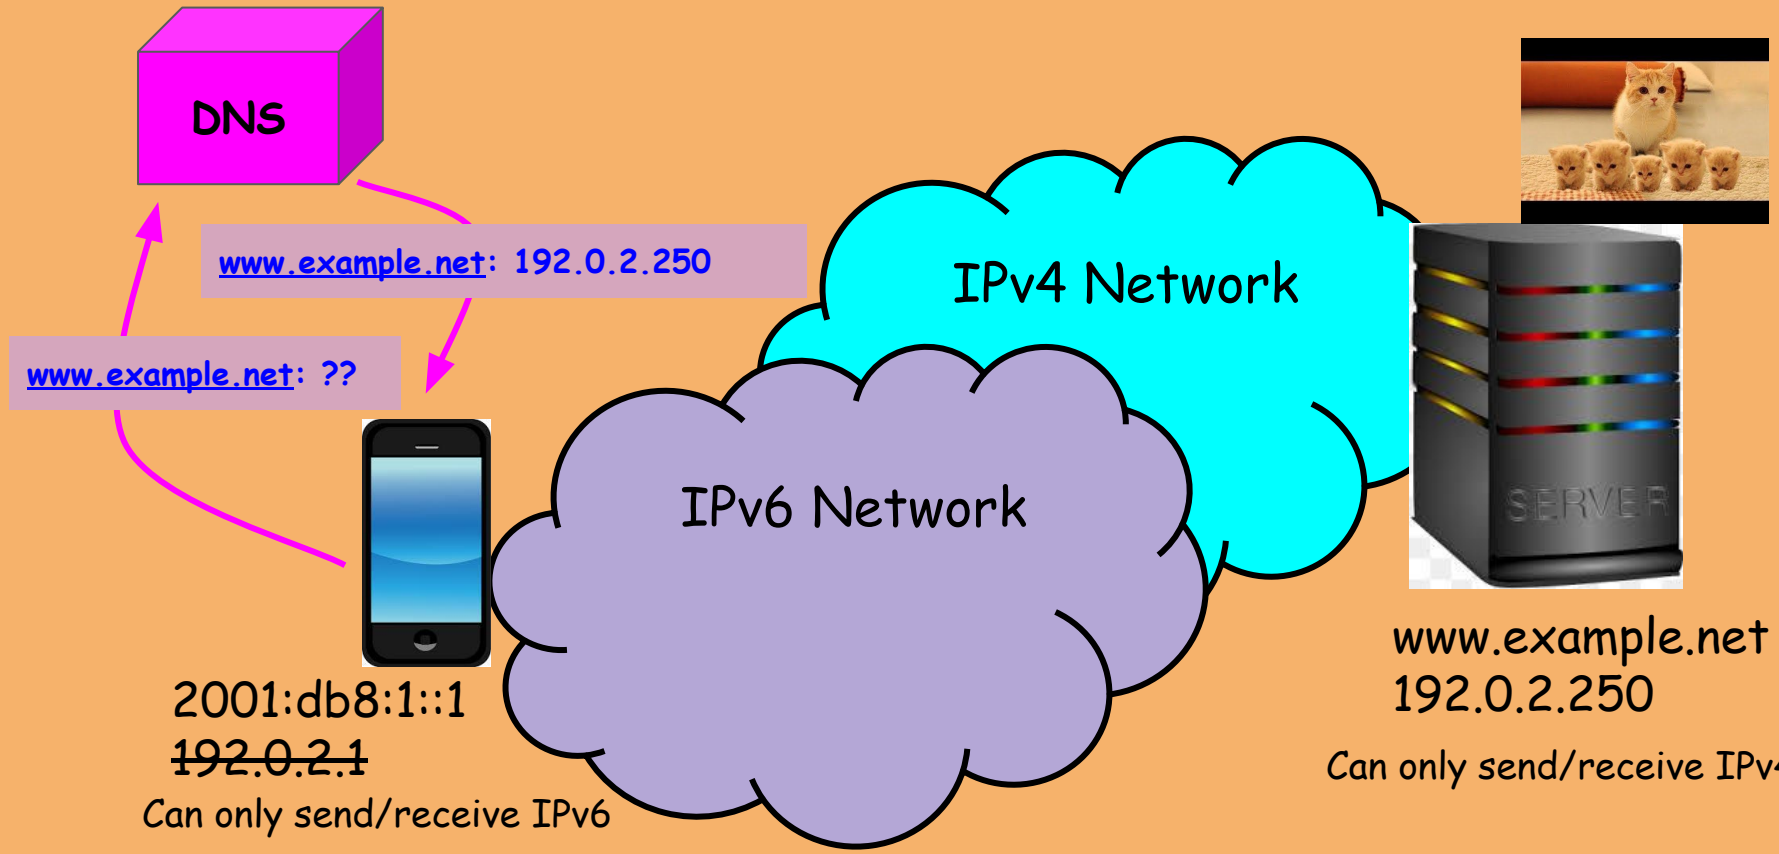

NAT64 + DNS64

IPv6 = 64:ff9b:: + IPv4

www.example.net: ??

DNS64

www.example.net: 64:ff9b::192.0.2.250

2001:db8:1::1

IPv6 packet to
64:ff9b::192.0.2.250

IPv6-only network

NAT64

IPv4 packet to
192.0.2.250

IPv4-only network

www.example.net
192.0.2.250

464XLAT (RFC6877)

DNS64 doesn't help if applications:

- Do not use DNS ("IPv4-literals)
- Only lookup IPv4 addresses
- Fail to operate w/o IPv4 address
- Uses DNSSEC

Solution: 464XLAT

- Provide applications with a private IPv4 address
- needs NAT64 only, no need for DNS64
  - DNSSEC-compatible

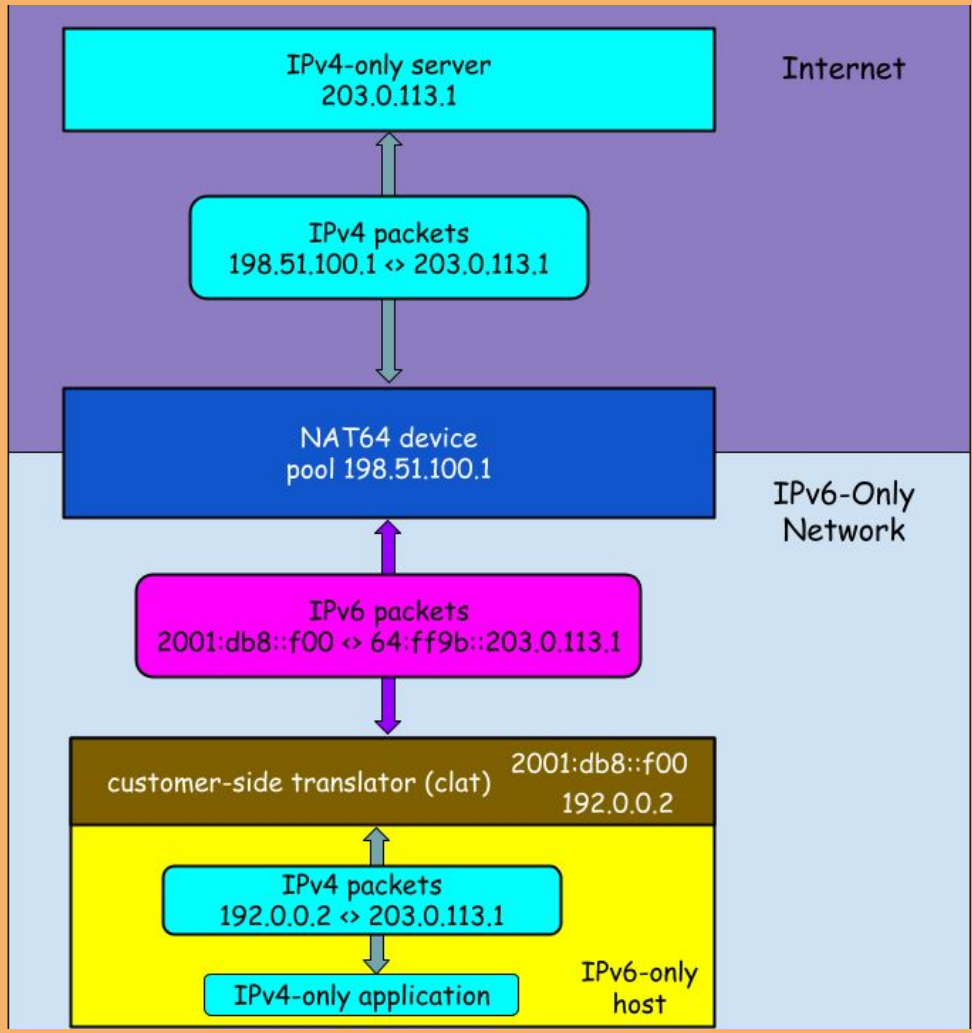

Network Overview

- SLAAC-only (no DHCPv6 for address assignment)
- NAT64/DNS64 to access IPv4-only destinations
  - NAT64 at the site edge
  - Router Advertisements options for DNS64 and PREF64
- Centralized DHCPv4 infrastructure
- Wired ports: 802.1x + dynamic vlan assignment

Stateless Address AutoConfiguration (SLAAC)

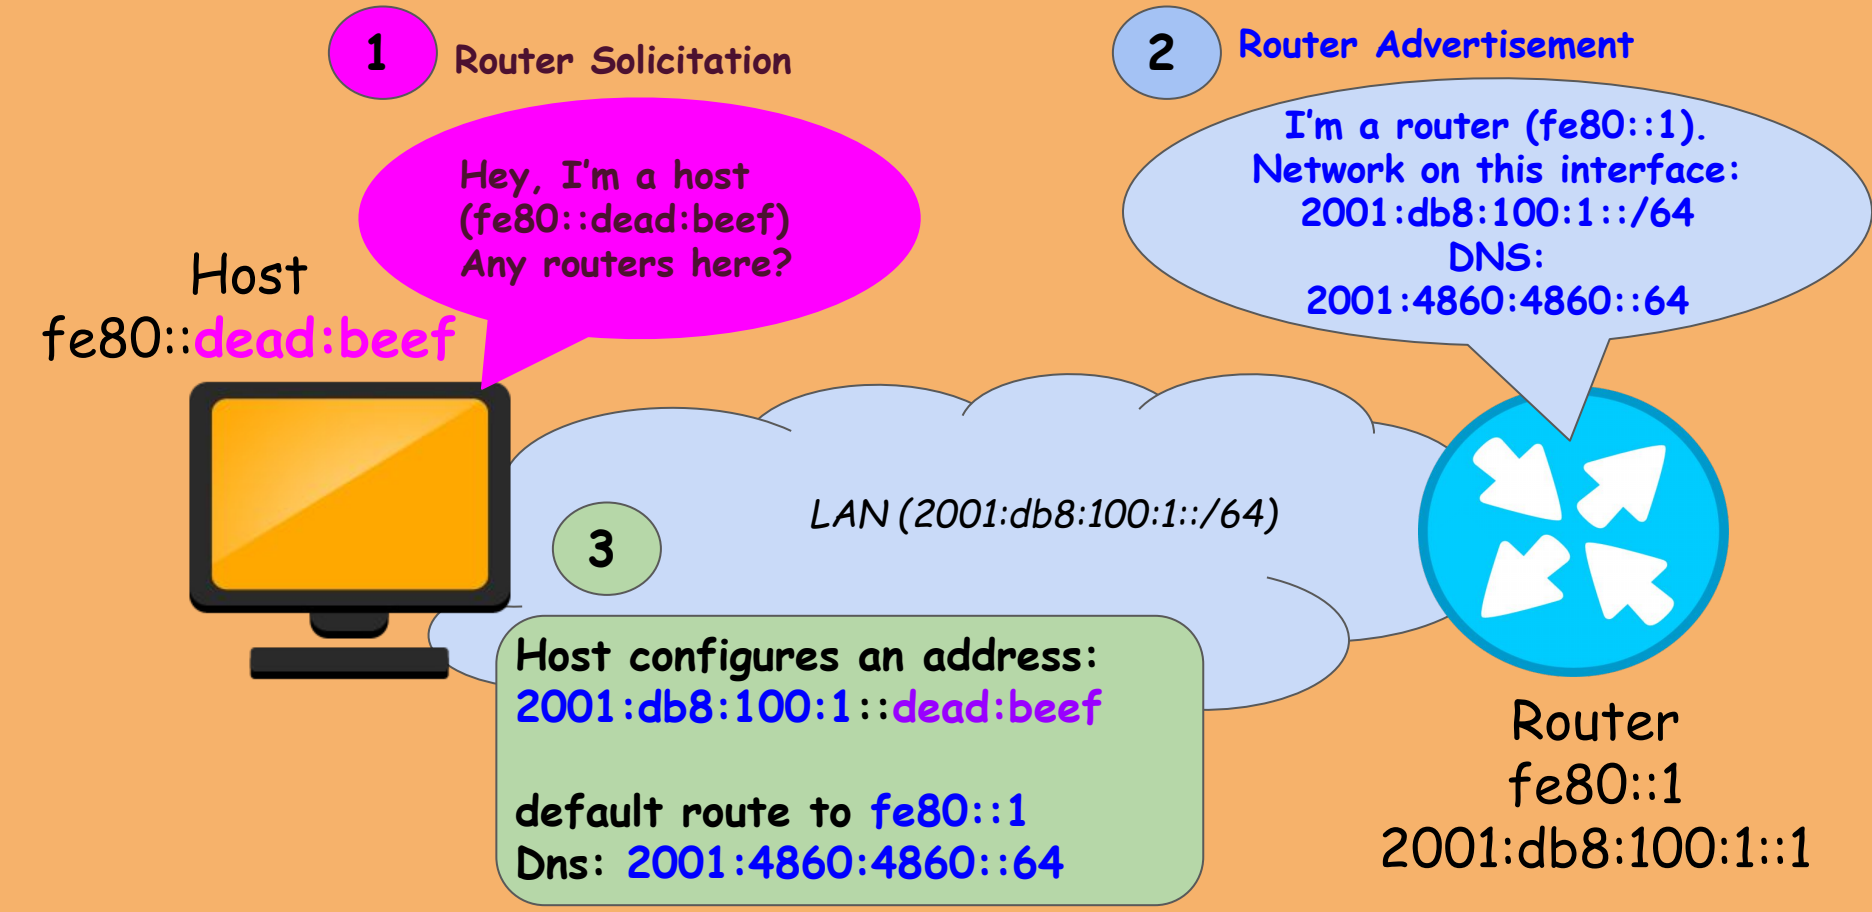

2020: First Attempt to IPv6-Only

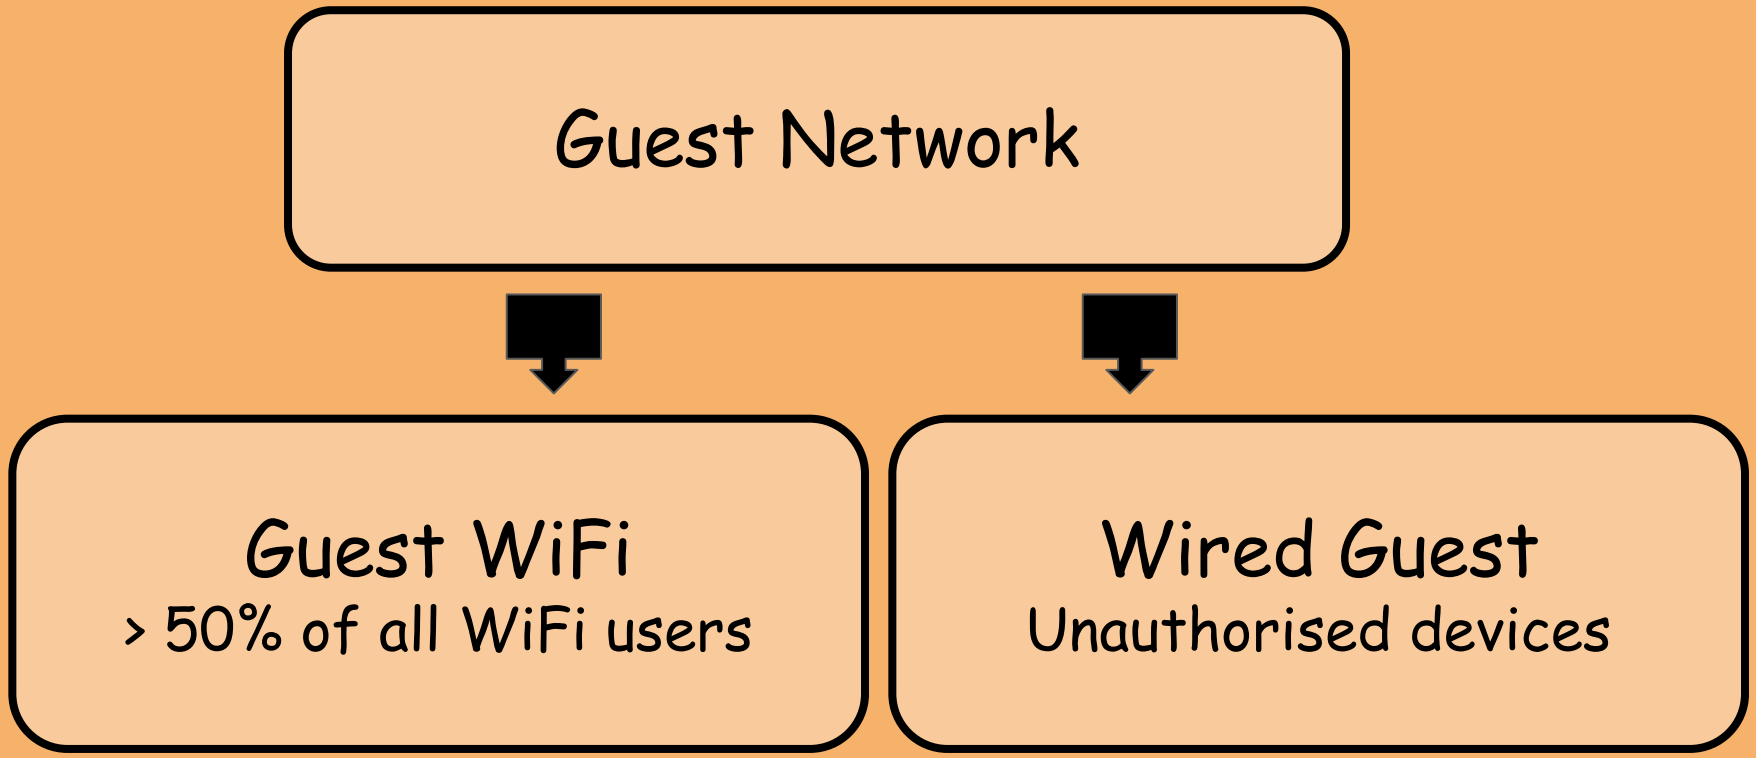

IPv6-Only Guest Overview

Two stages:

- Opt-in
  - users were invited to use IPv6-only WiFi
- Opt-out
  - the primary SSID is IPv6-only
  - a dedicated IPv4-enabled SSID for fallback

IPv6-Only Guest Overview

- Phase1: Opt-in
  - users were invited to use IPv6-only WiFi
- Phase 2: Opt-out
  - the primary SSID is IPv6-only
  - a dedicated IPv4-enabled SSID for fallback

IPv6-Only Guest Results

- Overall success
- A lot of address space reclaimed
- Many bugs/issues discovered and fixed
- IPv6-only support on endpoints
  - Mobile devices work just fine
  - Laptops/desktops: not always
    - Lack of 464xlat support
- **Need to support both IPv6-only and IPv4-enabled devices**

Dedicated SSID/VLAN: not a good idea

- Confusing for users
- Higher IPv4 consumption
- Lower visibility to issues
- Scalability concerns
- Operational complexity

We need something better!

IPv6-mostly Network

A network enabling co-existence of IPv6-only and IPv4-enabled devices

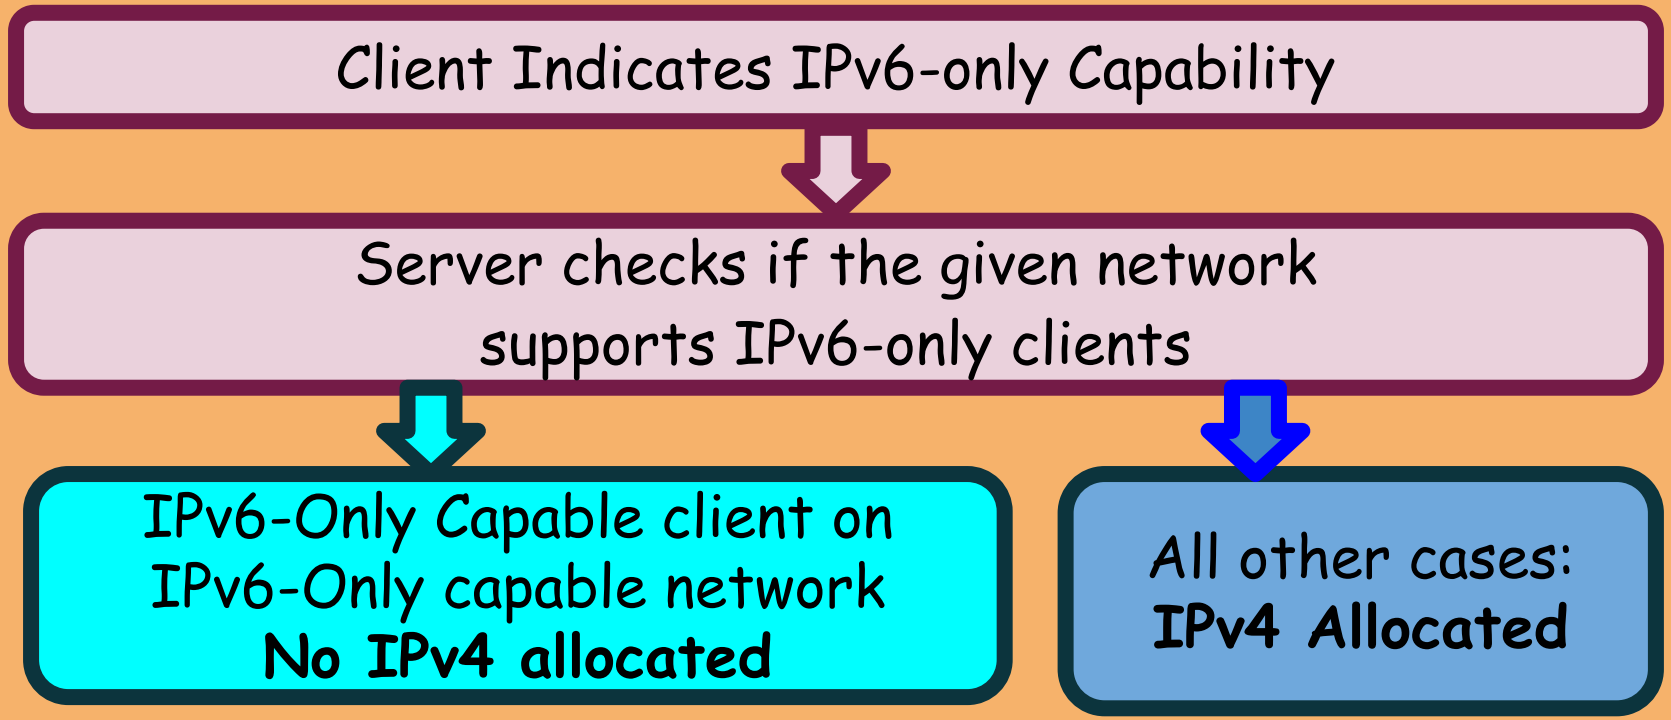

RFC8925: Use DHCPv4 to Turn IPv4 Off

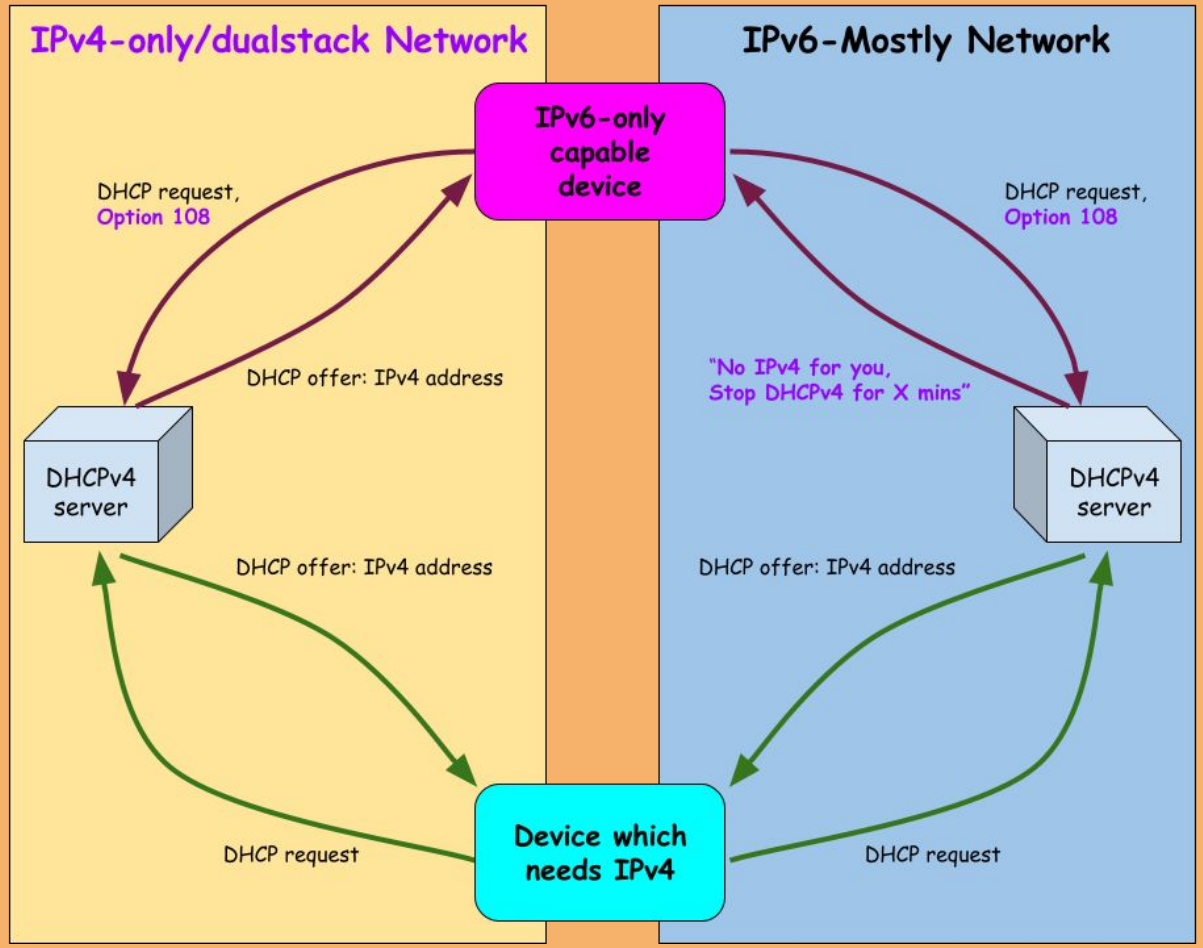

2023 Project Scope

Network Infrastructure across all offices globally:

- Corporate WiFi and IPv4-enabled (fallback) Guest WiFi
- Wired user-facing segments

Devices migrated to IPv6-Only:

- All Android, iOS (15+), MacOS 13+
  - always send DHCPv4 Option 108
  - support 464XLAT and PREF64
- Opt-in for selected ChromeOS and Linux devices

Rollout Schedule: March - Dec 2023

- Pilot in 3 locations for 2 months
- Extended pilot in 5 locations for 1 month
- “Stop the bleeding”: enable IPv6-mostly for greenfields
- Incremental rollout in 5 months, enabling Option 108 per subnet (10, 15, 25, 50, 60, 70, 80, 90, 100% of all networks)

Results

- No blocking issues found
  - A few cosmetic issues: all fixed in MacOS Sonoma
- DHCPv4 utilization dropped by 3-4 times (average) on WiFi
- Downsized subnets, reclaimed ~300K addresses

A Random Network: DHCP Utilization Drop

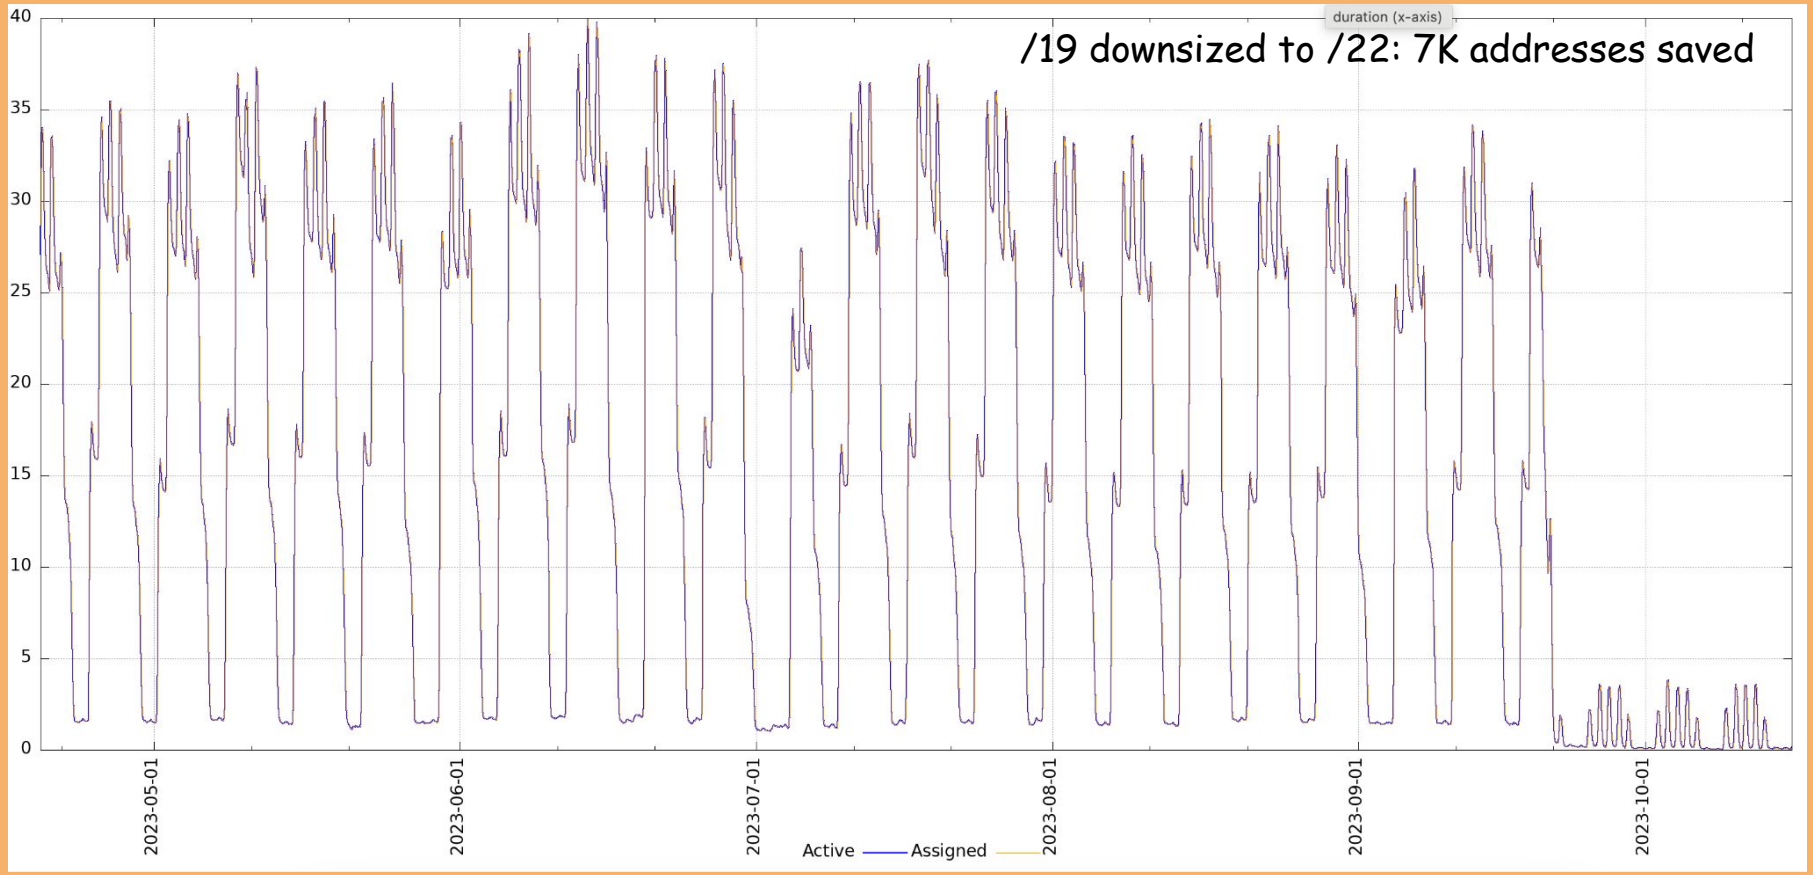

Lesson Learned #0

The only way to get IPv6 deployed:

to run out of (private) IPv4

Lesson Learned #1: “You Know Nothing, Jon Snow”

You do not really operate IPv6 until you turn IPv4 off

- Happy Eyeballs hide the problems
  - *“My workstation loses IPv6 DNS for a few mins after waking up”*
- Users do not report issues
- Issues are not getting fixed

Discovery #1: ~~Duck~~ Host Test

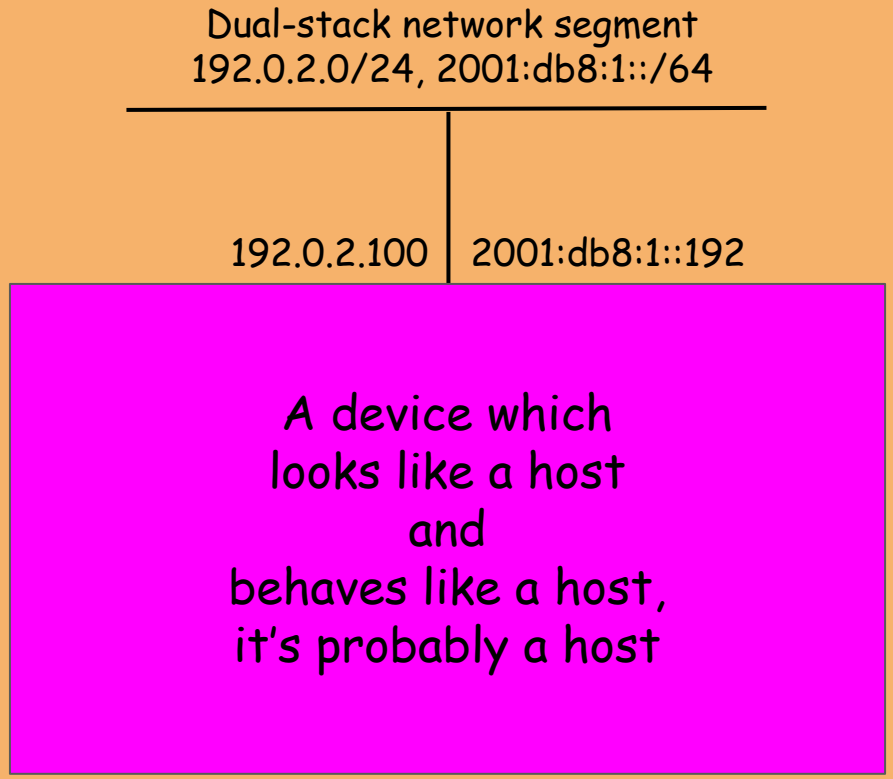

..or is it a router?

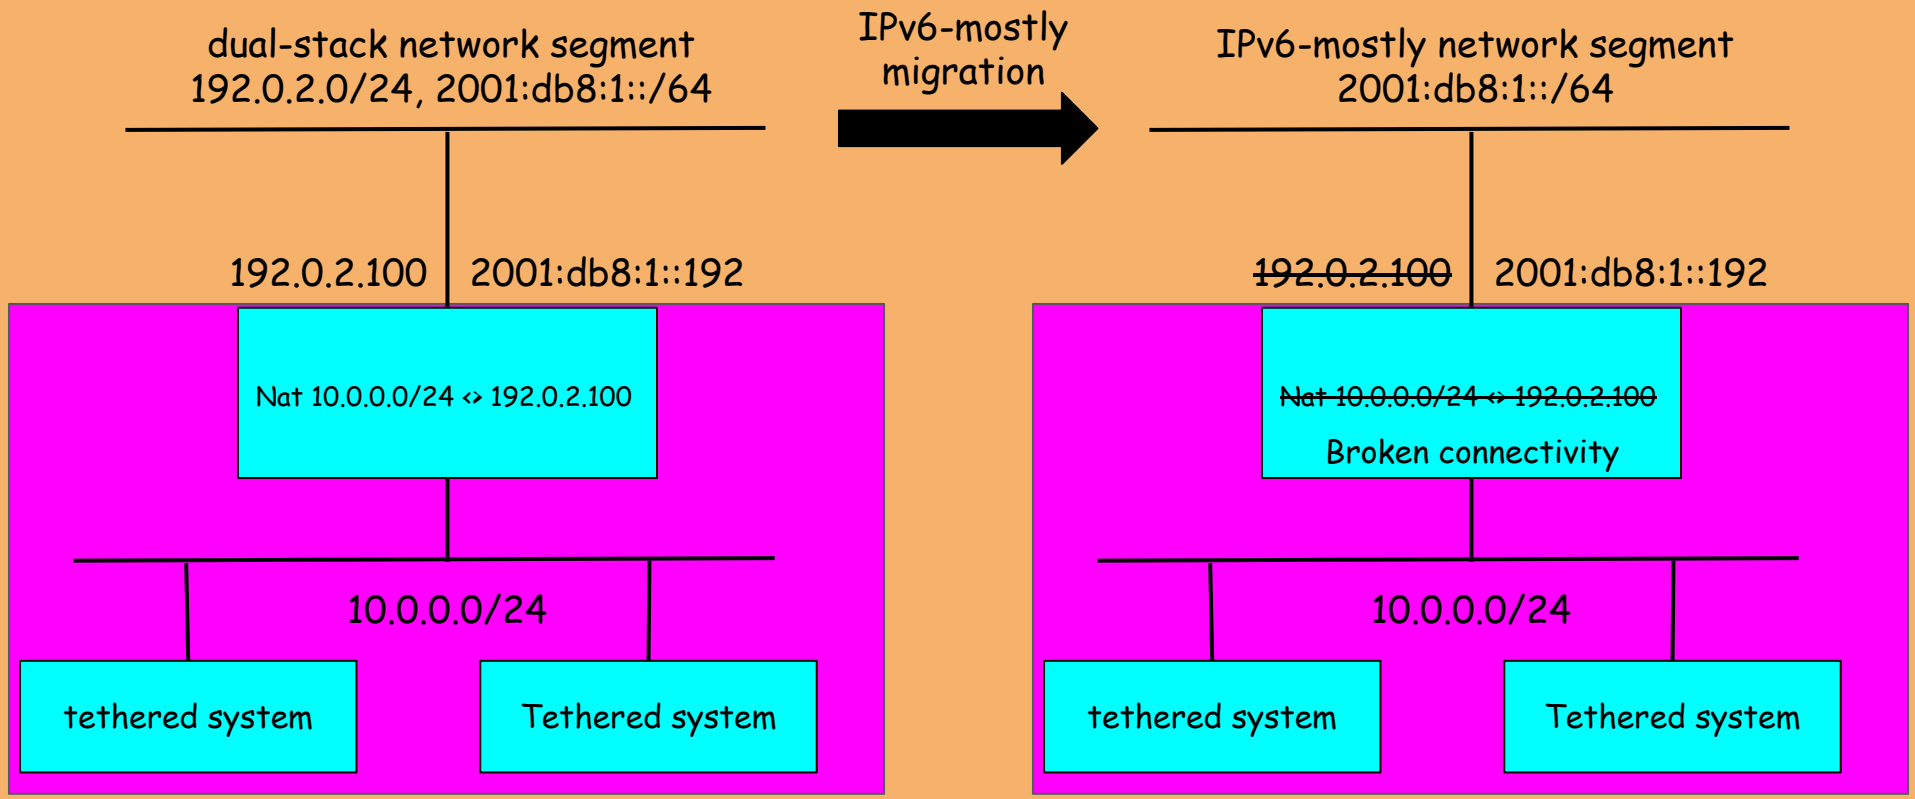

Solution: DHCPv6-PD

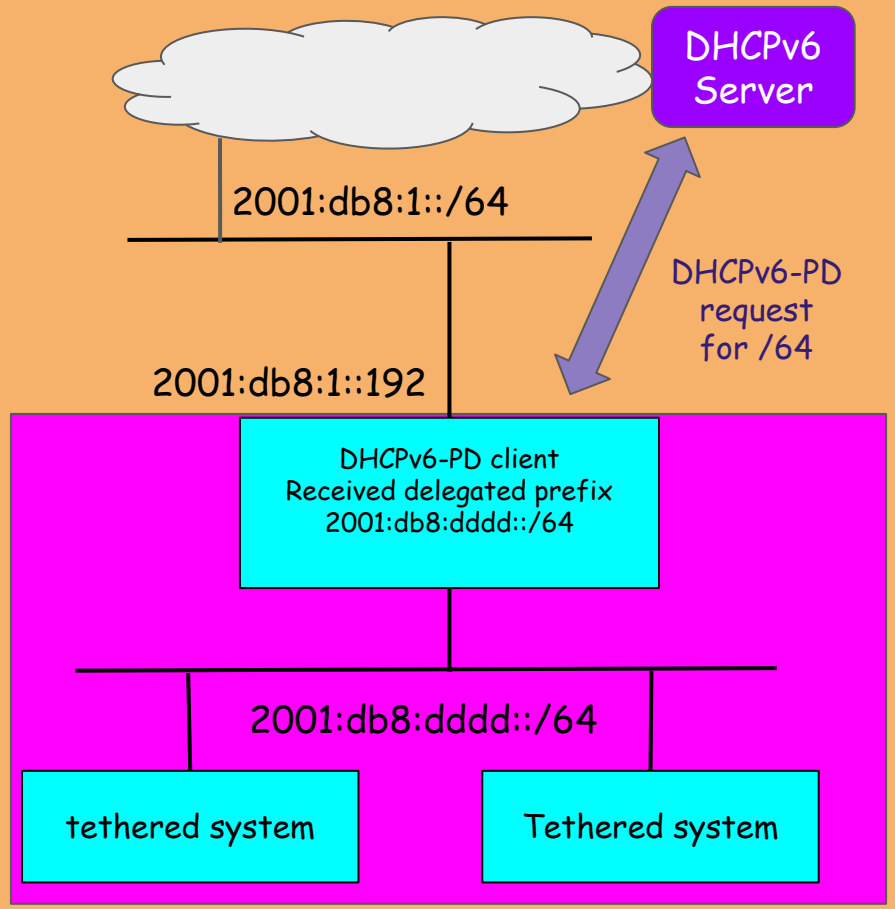

Lesson Learned #2: Extension Headers

Make sure Extension Headers are permitted

Especially

- Fragment Header
- ESP Header
  - Used by IPSec
    - VPNs
    - WiFi Calling

Discovery #3: Fragmentation Strikes Back

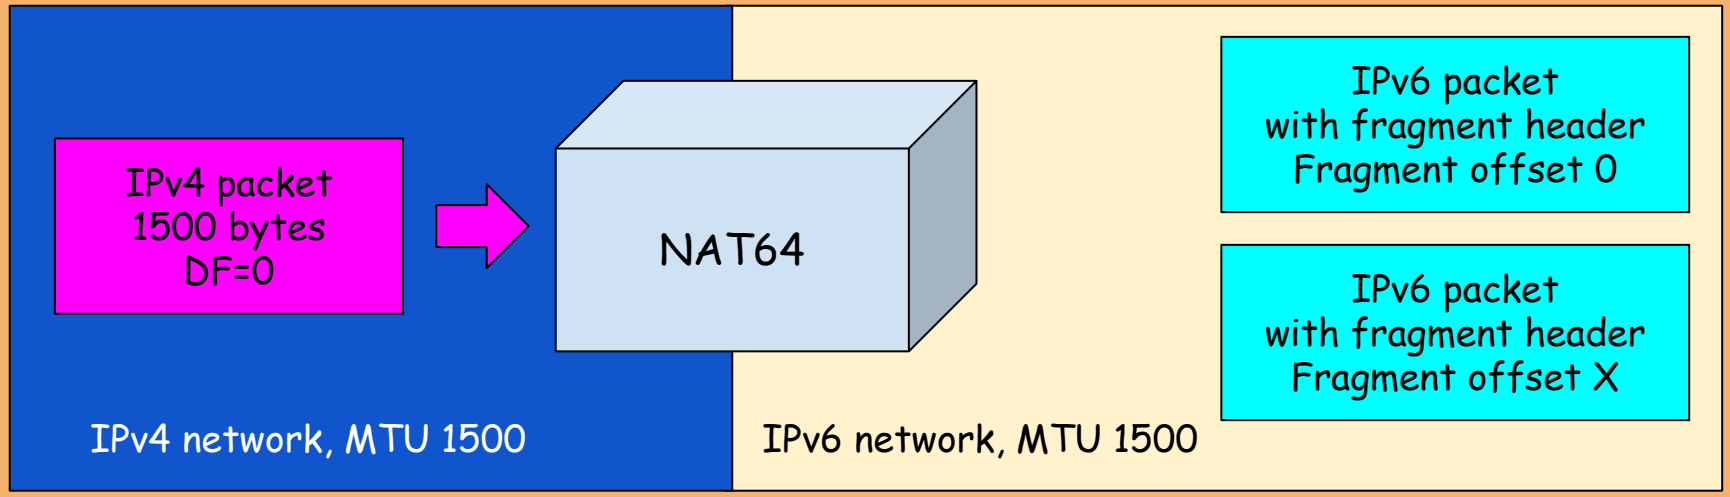

Caveats:

some NAT64 platforms use "1280" as a default size for translated packets instead of IPv6-only interface MTU.

Lesson Learned #3: Don't Disable IPv6

- "just disable IPv6 and see if it helps" wasn't a good idea.
- Had to automate enabling IPv6 on managed devices
- No way to fix it at scale for BYOD

Discovery #4: Hidden Limits

Host addresses: link-local, temporary, stable, 464XLAT

More addresses in case of virtual systems (ChromeOS: up to 20)

WiFi APs limit number of IPv6 addresses/client (limit can be as low as 7)

IPv6 addresses randomly lose connectivity

The Curious Case of Rip Van Winkle

- “My workstation loses IPv6 DNS for a few mins after waking up”
- Rootcause:
  - Router lifetime and RDNSS lifetime: 3600 secs
  - Device sleeps for > 1hr
  - A bug in the OS: DNS expires, the router is not!

RFCs Published

- RFC 8781
  - Discovering PREF64 in Router Advertisements
- RFC 8925
  - IPv6-Only Preferred Option for DHCPv4
- RFC 9131
  - Gratuitous Neighbor Discovery: Creating Neighbor Cache Entries on First-Hop Routers

IETF Work in Progress

- IPv6-Mostly Deployment Guidelines draft-link-v6ops-6mops
- Using DHCPv6-PD to Allocate Unique IPv6 Prefix per Client in Large Broadcast Networks (draft-ietf-v6ops-dhcp-pd-per-device)
- 464 Customer-side Translator (CLAT): Node Recommendations (draft-link-v6ops-claton)

Next Steps

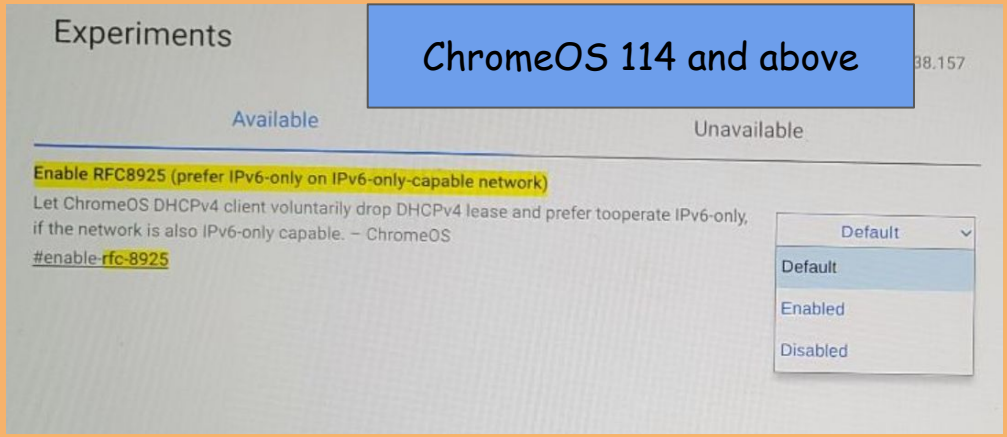

ChromeOS

- Option 108 can be enabled

Microsoft announced plans to support Option 108 + clat

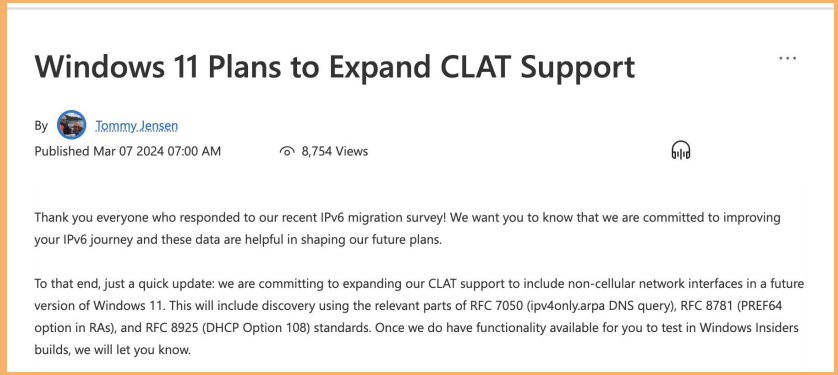

Windows 11 Plans to Expand CLAT Support

By Tommy Jensen

Published Mar 07 2024 07:00 AM 8,754 Views

Thank you everyone who responded to our recent IPv6 migration survey! We want you to know that we are committed to improving your IPv6 journey and these data are helpful in shaping our future plans.

To that end, just a quick update: we are committing to expanding our CLAT support to include non-cellular network interfaces in a future version of Windows 11. This will include discovery using the relevant parts of RFC 7050 (ipv4only.arpa DNS query), RFC 8781 (PREF64 option in RAs), and RFC 8925 (DHCP Option 108) standards. Once we do have functionality available for you to test in Windows Insiders builds, we will let you know.

QUESTIONS?