



**RIPE** Réseaux IP  
Européens

# IPV6-MOSTLY EXPERIENCES AT 2 INDUSTRY EVENTS

Éric Vyncke, Distinguished Engineer, Cisco Belgium

Warren Kumari, Google, IETF NOC member

Slide content by Joe Clarke, Cisco, Cisco Live & IETF NOC member

GenAI graphics by Gemini


# IPV6-MOSTLY ?

In IPv6-Mostly networks, 2 classes of hosts


- IPv6-Only capable: do not use IPV4 at all
- Not capable: use dual-stack or IPv4-only

Goal:

- Foster IPv6 adoption
- Lower IPv4 address consumption
- Until 100% are IPv6-only capable



# IPV6-MOSTLY IN | SLIDE



- IPv6-only capable hosts
  - Rely on client-level address translation (CLAT), IPv6 prefix from the PREF64 RA to do translation of IPv4 addresses to v6
  - Decline to request IPv4 address if DHCPv4 server replies to DHCPv4 option 108 RFC 8925 “IPv6-Only Preferred Option for DHCPv4”
- <https://datatracker.ietf.org/doc/draft-ietf-v6ops-6mops/> by Nick Buraglio , Ondřej Caletka , Jen Linkova





# THE EXPLORERS

- Several other explorers of this route
  - Jen Linkova @ Google
  - Ondřej Caletka for RIPE network since RIPE-85 October 2022 !
  - David Stockdale, Imperial College London with 20.000 students, see [https://www.ipv6.org.uk/wp-content/uploads/2024/10/08\\_IPv6-Mostly\\_at\\_Imperial.pdf](https://www.ipv6.org.uk/wp-content/uploads/2024/10/08_IPv6-Mostly_at_Imperial.pdf)



# CISCO LIVE

- Amsterdam, February 2025
- 20.000 attendees
- Technical sessions and exhibition by Cisco for its customers
- Cisco NOC has months to prepare and do very careful planning (including contingency plans)
- **Default SSID was IPv6-mostly**
- **Public IPv6 space, private IPv4**



# Wired Core Team



Lionel  
Hercot



James Botting




Andrew  
Yourtchenko




Martin Kriska



Ivan Ivanov



# EDGE DESIGN

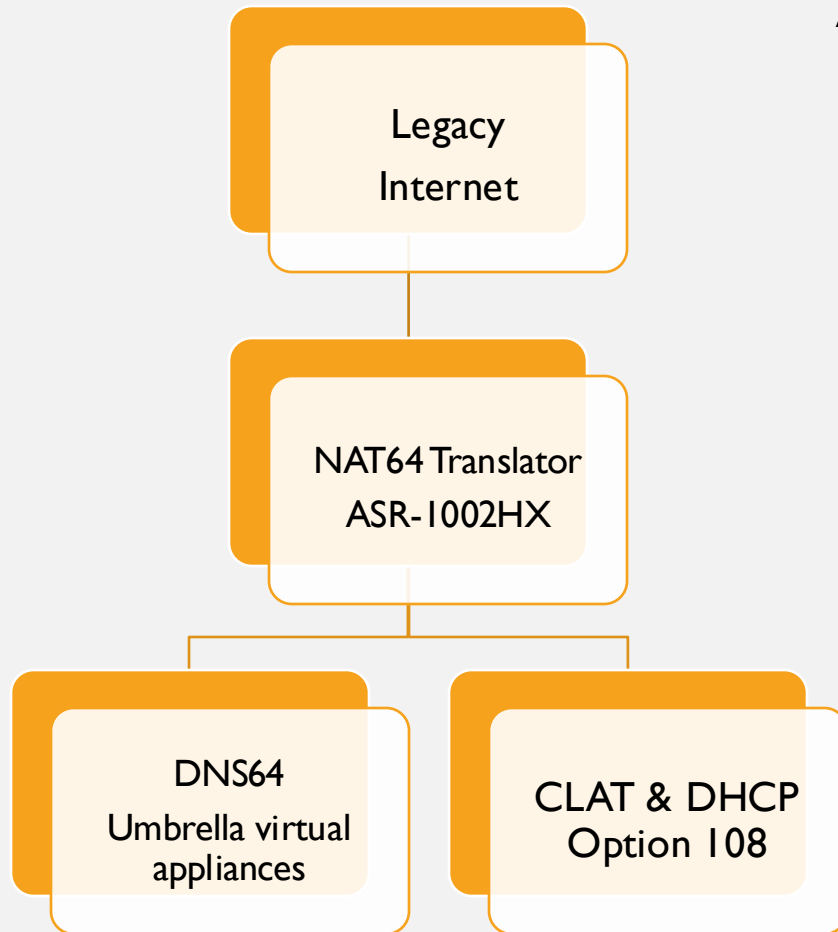


2x ASR 1002-HX for NAT44 and Native IPv6  
 2x ASR 1002-HX for NAT64 only  
 /26 of V4 public per NAT pool

Each unit BGP peered to Venue  
 1x 10Gb/s connectivity to Venue  
 2x 40Gb/s to core

# NAT64 DESIGN

Aspiring for 100% IPv6 inside the network



IPv4 Traffic  
83.97.13.0/24

IPv6 Traffic  
64:ff9b::/96



# NAT64 MORE POPULAR THAN EXPECTED

## SEGMENTING YOUR IPV4 SPACE BETWEEN NAT44 AND NAT64 IS TRICKY

- Didn't really know how popular NAT64 would be over NAT44 edge.
- Ultimately, not enough addresses in NAT64 pool and too many in NAT44
- Changed translation expiry timers to resolve, added more IP's out of hours.
- *Additionally, Cisco Secure Client split tunnel VPNs didn't like CLAT*

i	Time	Event
>	2/11/25 2:23:22.000 PM	%NAT64-6-ADDR_ALLOC_FAILURE: Address allocation failed; pool 1 may be exhaustedQFP:0.1 Thread:013 TS:00001045967706653903 host = corel-nat64 : source = sc4s : sourcetype = cisco:ios
>	2/11/25 2:22:23.000 PM	%NAT64-6-ADDR_ALLOC_FAILURE: Address allocation failed; pool 1 may be exhaustedQFP:0.0 Thread:041 TS:00001045907705145316 host = corel-nat64 : source = sc4s : sourcetype = cisco:ios
>	2/11/25 2:20:22.000 PM	%NAT64-6-ADDR_ALLOC_FAILURE: Address allocation failed; pool 1 may be exhaustedQFP:0.1 Thread:062 TS:00001045787702266807 host = corel-nat64 : source = sc4s : sourcetype = cisco:ios
>	2/11/25 2:19:22.000 PM	%NAT64-6-ADDR_ALLOC_FAILURE: Address allocation failed; pool 1 may be exhaustedQFP:0.0 Thread:247 TS:00001045727701264995 host = corel-nat64 : source = sc4s : sourcetype = cisco:ios
>	2/11/25 2:18:22.000 PM	%NAT64-6-ADDR_ALLOC_FAILURE: Address allocation failed; pool 1 may be exhaustedQFP:0.0 Thread:030 TS:00001045667699795864 host = corel-nat64 : source = sc4s : sourcetype = cisco:ios
>	2/11/25 2:17:21.000 PM	%NAT64-6-ADDR_ALLOC_FAILURE: Address allocation failed; pool 1 may be exhaustedQFP:0.1 Thread:073 TS:00001045607699166472 host = corel-nat64 : source = sc4s : sourcetype = cisco:ios
>	2/11/25 2:16:22.000 PM	%NAT64-6-ADDR_ALLOC_FAILURE: Address allocation failed; pool 1 may be exhaustedQFP:0.1 Thread:175 TS:00001045547695930319 host = corel-nat64 : source = sc4s : sourcetype = cisco:ios
>	2/11/25 2:15:22.000 PM	%NAT64-6-ADDR_ALLOC_FAILURE: Address allocation failed; pool 1 may be exhaustedQFP:0.1 Thread:217 TS:00001045487695689167 host = corel-nat64 : source = sc4s : sourcetype = cisco:ios
>	2/11/25 2:14:22.000 PM	%NAT64-6-ADDR_ALLOC_FAILURE: Address allocation failed; pool 1 may be exhaustedQFP:0.0 Thread:220 TS:00001045427695452648 host = corel-nat64 : source = sc4s : sourcetype = cisco:ios
>	2/11/25 2:13:22.000 PM	%NAT64-6-ADDR_ALLOC_FAILURE: Address allocation failed; pool 1 may be exhaustedQFP:0.0 Thread:167 TS:00001045367693901283 host = corel-nat64 : source = sc4s : sourcetype = cisco:ios

Clients that support IPv6 Mostly will request DHCPv4 option 108.

If available, those clients will prefer v6 and do CLAT in order

Number of Clients That Support IPv6 Mostly 📊

Number of Clients That DO NOT Support IPv6 Mostly 📊

20401

24005

MAIN SSID V6 MOSTLY SUPPORT

Total Internet Volume since Core Power-on

90.26 TB

TOTAL TRAFFIC TO THE  
INTERNET...

20,000 attendees over five days




# IETF 122

- Bangkok, March 2025
- 764 onsite registered participants
- Typical geek audience
- **Default SSID is dual-stack**
  - *Plan for IETF-123 IPv6-Mostly ;-)*
- **"ietf-ipv6-mostly" widely advertised and used by many**
- **Public IPv4 and IPv6 space**



# IPV6-MOSTLY COMPONENTS



ALL WAS GOOD FOR MACOS AND  
PHONES, BUT...

## Sunday Evening...

Bill mentions, I can't join ietf-ipv6-mostly...



ietf-ipv6-mostly

In the process of debugging, we discovered IPv6  
Mostly was IPv6 [NAT64] Only

- DHCP discovers weren't being relayed to Kea
- Kea wasn't configured for Option 108 🤪



# After Much Troubleshooting

10

```
static {
  rule-set NAT64 {
    from zone [ junos-host nat64 ];
    rule NAT64 {
      match {
        source-address 2001:67c:370:1998::/64; <===== "the line"
        destination-address 64:ff9b::/96;
      }
      then {
        static-nat {
          inet;
        }
      }
    }
  }
}
```






## SRX

- Without **The Line™**, DHCP replies were being eaten by the SRX
- No one knows why this line is needed as it's not documented in examples for NAT64
- In addition, Kea needed option 108 (set to 900), and Unbound needed to hand out DNS64\*

\* Config was finalized, but too late. So DNS64 over v6 and DNS over v4

# Overall, People Seem Happy

13

Name	Last * 	Max	Mean
 SSID: ietf	274	582	201
 SSID: ietf-ipv6-mostly	100	138	29
 SSID: eduroam	18	52	19
 SSID: ietf-legacy122	0	1	0

And 371 “unique” clients total



## REAL LIFE EXPERIMENT WITH ENGINEERS

- Discovery: MacOS SSH forced to an IPv4 host address. This didn't work. But Tommy Pauly @ Apple was told and hacked up a fix while at the meeting 😊
- *“Great environment to deploy such features because when things go wrong, the right people are in the room to fix the problems !”* (Quote: Joe Clarke, IETF NOC)



# IPV6-MOSTLY SUMMARY

- If you pay attention, it is not hard at all
- Network operations and debugging are different => ! Learning curve
- Issues & limitations:
  - Split-tunnel VPN (esp. MacOS)
  - Slack weirdness
  - Some corporate MDM (for DNS64)
  - Microsoft Windows current lack of support for CLAT & RFC 8925 are still dual-stack hosts

