IPv6 addressing inside VPN tunnel between endpoints with rotating prefixes

RIPE 84
May 19th, 2022
Matthias Scheer

avm.de

$ whois

About us...

- AVM is best known for its series of residential /SOHO gateways
- Linux based OS with custom IP routing stack
- Early deployment of IPv6 in the consumer field (2009)

About me...

- Software / Network engineer 10+ years with AVM
- Current focus: overhaul of gateway-terminated VPN

Current IPv6 deployment – our perspective

Customer feedback – as of April 2022 (representative sample):

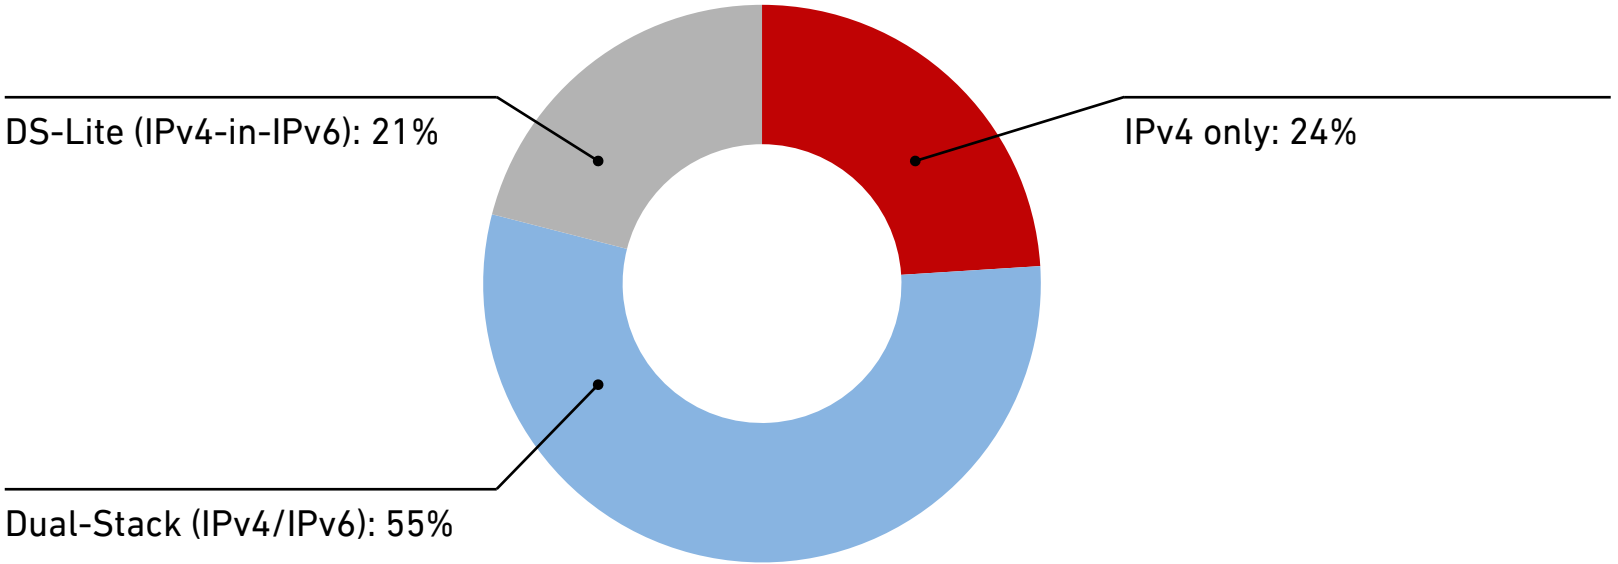

Current IPv6 deployment – our conclusion

- ~ 3 out of 4 residential / SOHO connections with active IPv6 already
  - Time to switch default configuration from opt-in to opt-out

- Next major OS update will enable IPv6 on all connections
  - Already deployed in our public beta program („Labor“)
  - Fail-over if no IPv6 is supported by the ISP

IPv6-enabled VPN

- Upcoming release also introduces IPv6 connectivity between VPN endpoints
- IPv6 is preferred during connection establishment, endpoints are expected to behave similarly
- But no IPv6 connectivity inside VPN tunnel yet (IPv4-in-IPv6)
- Why…? Rotating LAN prefixes at the endpoints

Rotating prefixes

- Residential / SOHO IPv6 delegated prefixes change on every connection dial-in with most ISP (sync loss, power outage, ...)
- Forced disconnect after 24h with some ISP

➢ Routing and name resolution need to know the changed opposite endpoint adresses

➢ DNS resolution inside the tunnel is essential

IPv6 in the tunnel

Site-to-Site vs. Road Warrior

- Road Warrior usually has a ::/0 into the VPN tunnel
- Only DNS server needs to be known at the Road Warrior side

- Endpoint in Site-to-Site configuration must be able to reach resources in opposite endpoint‘s current IPv6 address space
- Endpoints need to be statically configured with DNS domains of opposite endpoint as well as their DNS servers

IPv6 in the tunnel

IPSec vs. WireGuard

- Versatility vs. Simplicity
- IPSec contains dedicated configuration exchanges, which allow to transmit dynamically changed parameters
- WireGuard uses static endpoint configuration
  (no dynamic routing parameters)
- WireGuard clients prefer IPv4 despite A and AAAA records

➢ Different VPN protocols require different solutions

IPv6 in the tunnel

GUA-based Approach (for IPSec)

- Dynamic configuration during IKE exchange using dedicated payloads:
  - LAN prefix announced by INTERNAL_IP6_SUBNET
  - DNS servers announced by INTERNAL_IP6_DNS
- Traffic selectors embedded in IPSec SA must match the prefix announced by INTERNAL_IP6_SUBNET
- Routes and DNS servers are adjusted accordingly at both endpoints after the connection has been successfully established
- Caveat: WAN adresss of gateway must not be part of LAN prefix

IPv6 in the tunnel

GUA-based Approach (for IPSec)

Left VPN configuration:
Domain = right.fritz.box
DNS = 2001:db8:a:b::1
Route = 2001:db8:a:b::/64

INTERNAL_IP6_DNS = 2001:db8:a:b::1
INTERNAL_IP6_SUBNET = 2001:db8:a:b::/64

Right VPN configuration:
Domain = left.fritz.box
DNS = 2001:db8:1:2::1
Route = 2001:db8:1:2::/64

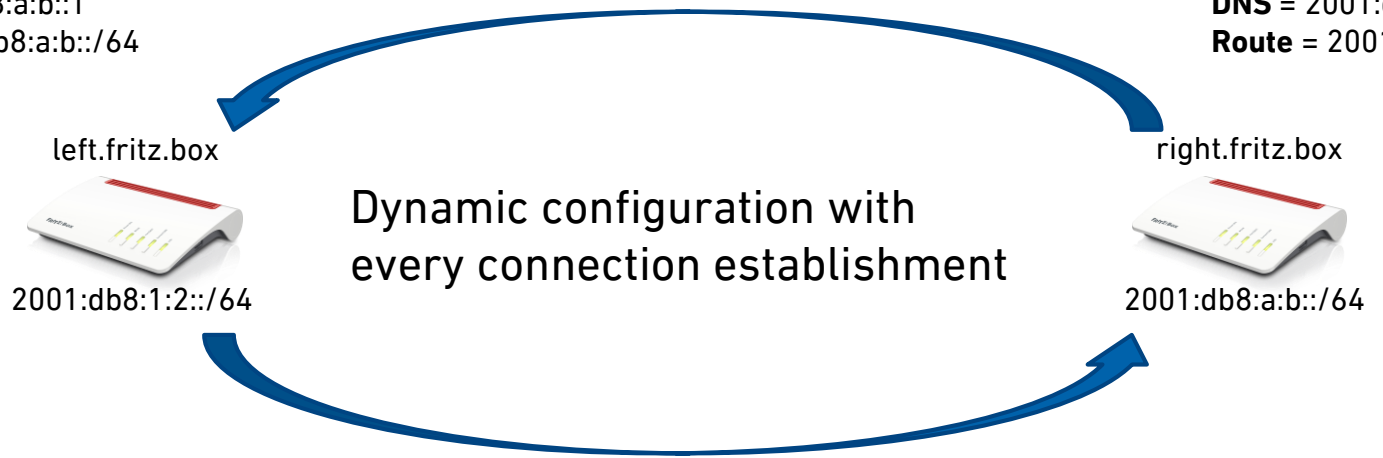

INTERNAL_IP6_DNS = 2001:db8:1:2::1
INTERNAL_IP6_SUBNET = 2001:db8:1:2::/64

IPv6 in the tunnel

ULA-based Approach (for WireGuard)

- RFC 4193 explicitly suggests the use of ULAs for site-to-site-VPN
- ULA configuration on both sides necessary
- Static WireGuard configuration uses ULA of opposite endpoint:
  - ULA prefix defines route into VPN tunnel
  - ULA of opposite gateway defines DNS server
- Caveat: IPv6 source address selection (RFC6724) must choose ULA

IPv6 in the tunnel

ULA-based Approach (for WireGuard)

Left VPN configuration:
Domain = right.fritz.box
DNS = fd00:a:b:c::1
Route = fd00:a:b:c::/64

Right VPN configuration:
Domain = left.fritz.box
DNS = fd00:1:2:3::1
Route = fd00:1:2:3::/64

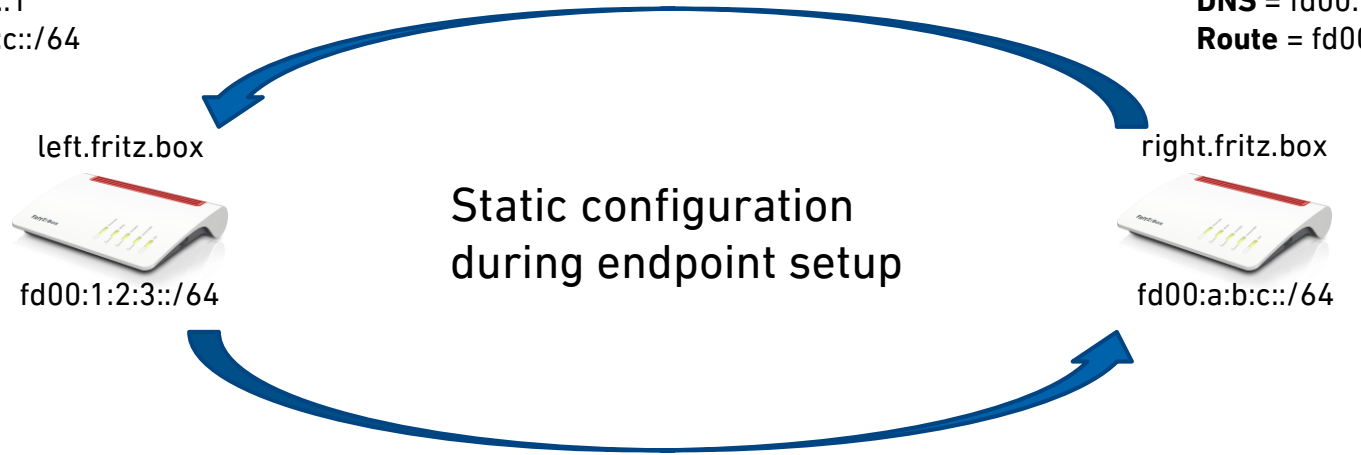

Questions, Suggestions, ...

Thank you for your attention!

Feel free to come talk to us after the session!