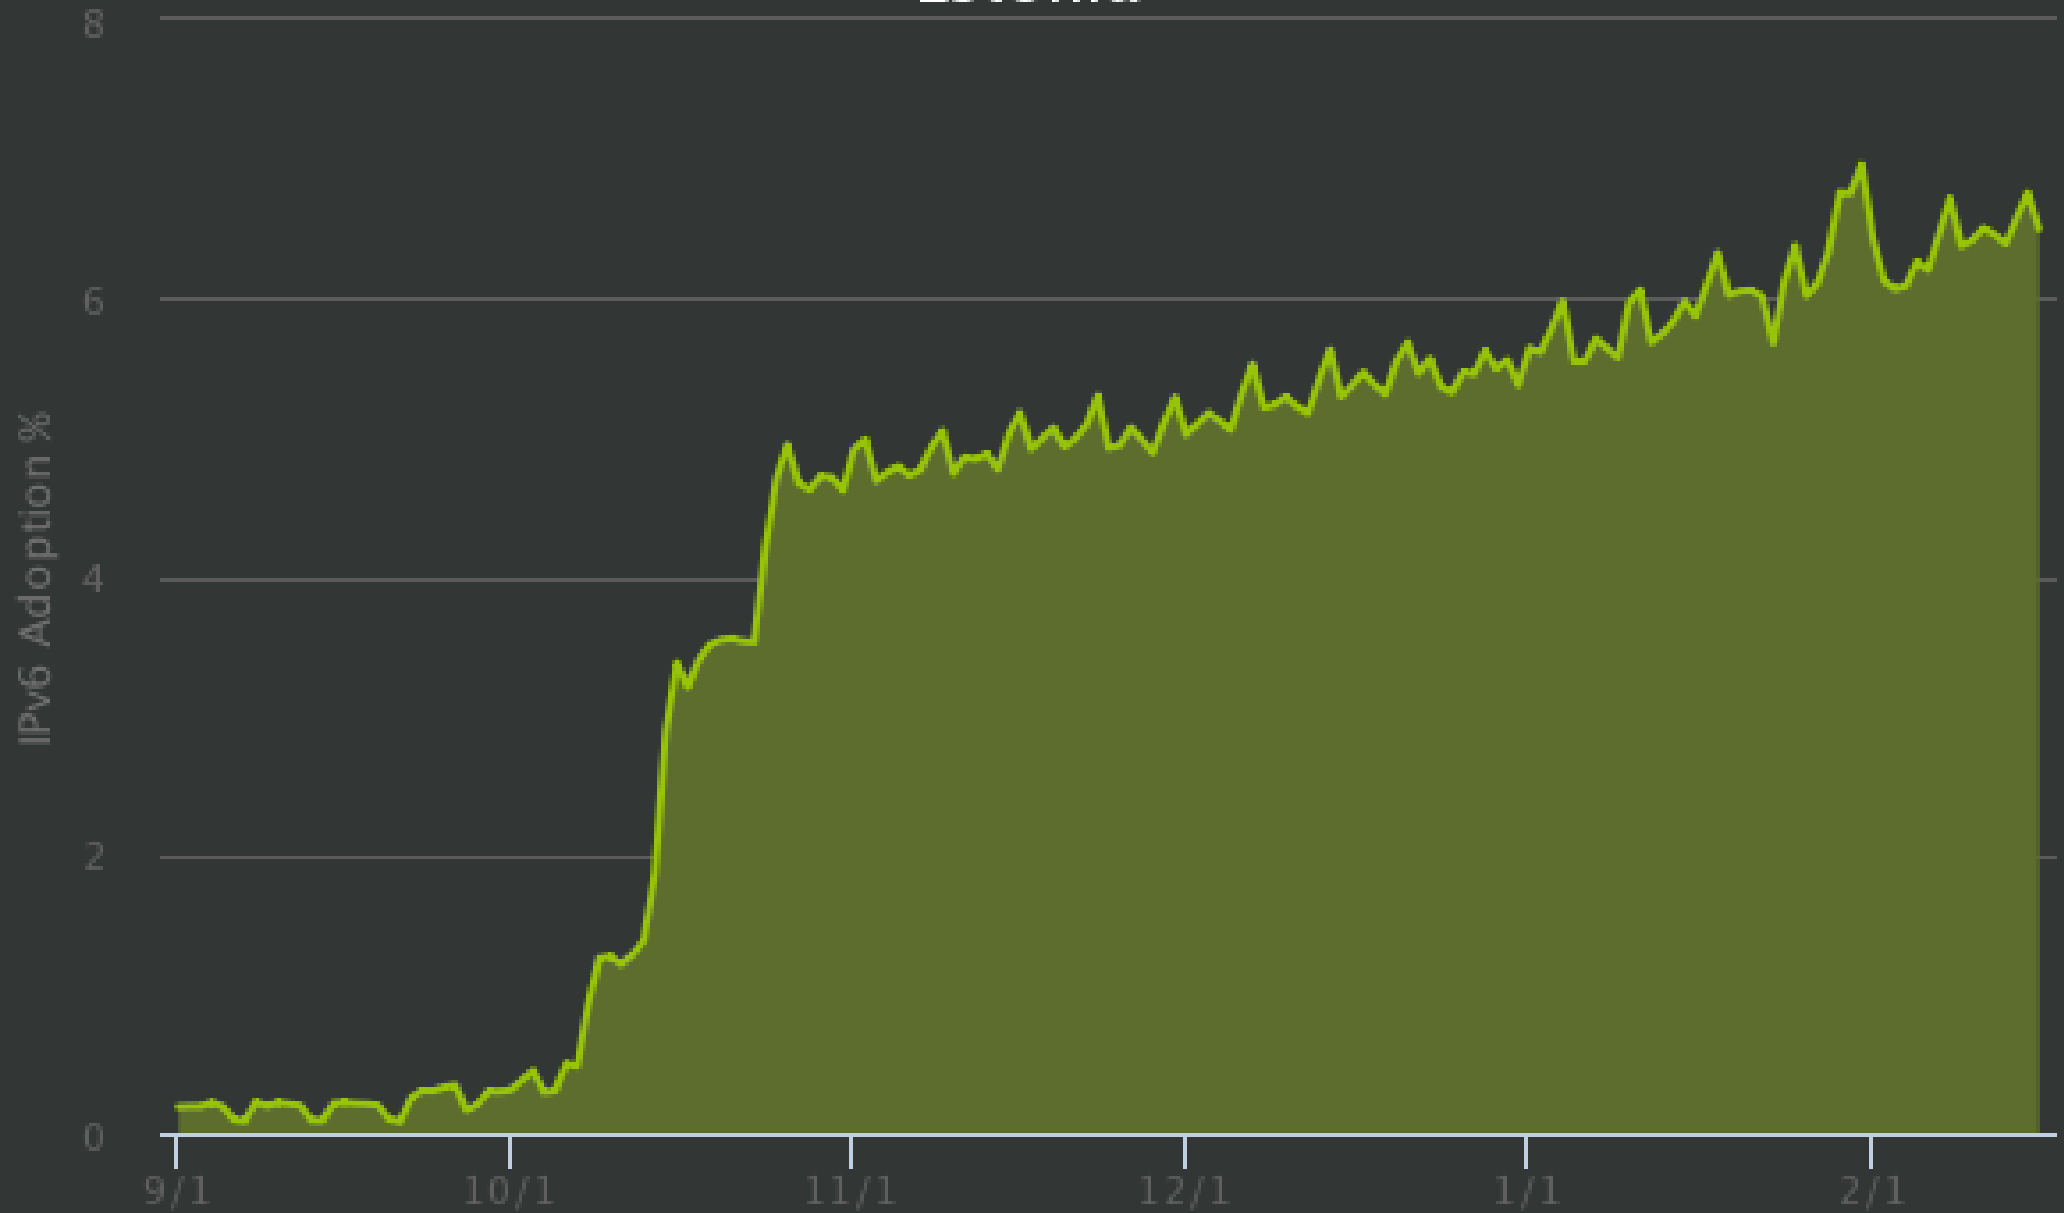


Enabling and Securing IPv6 in Service Provider Networks

Tarko Tikan

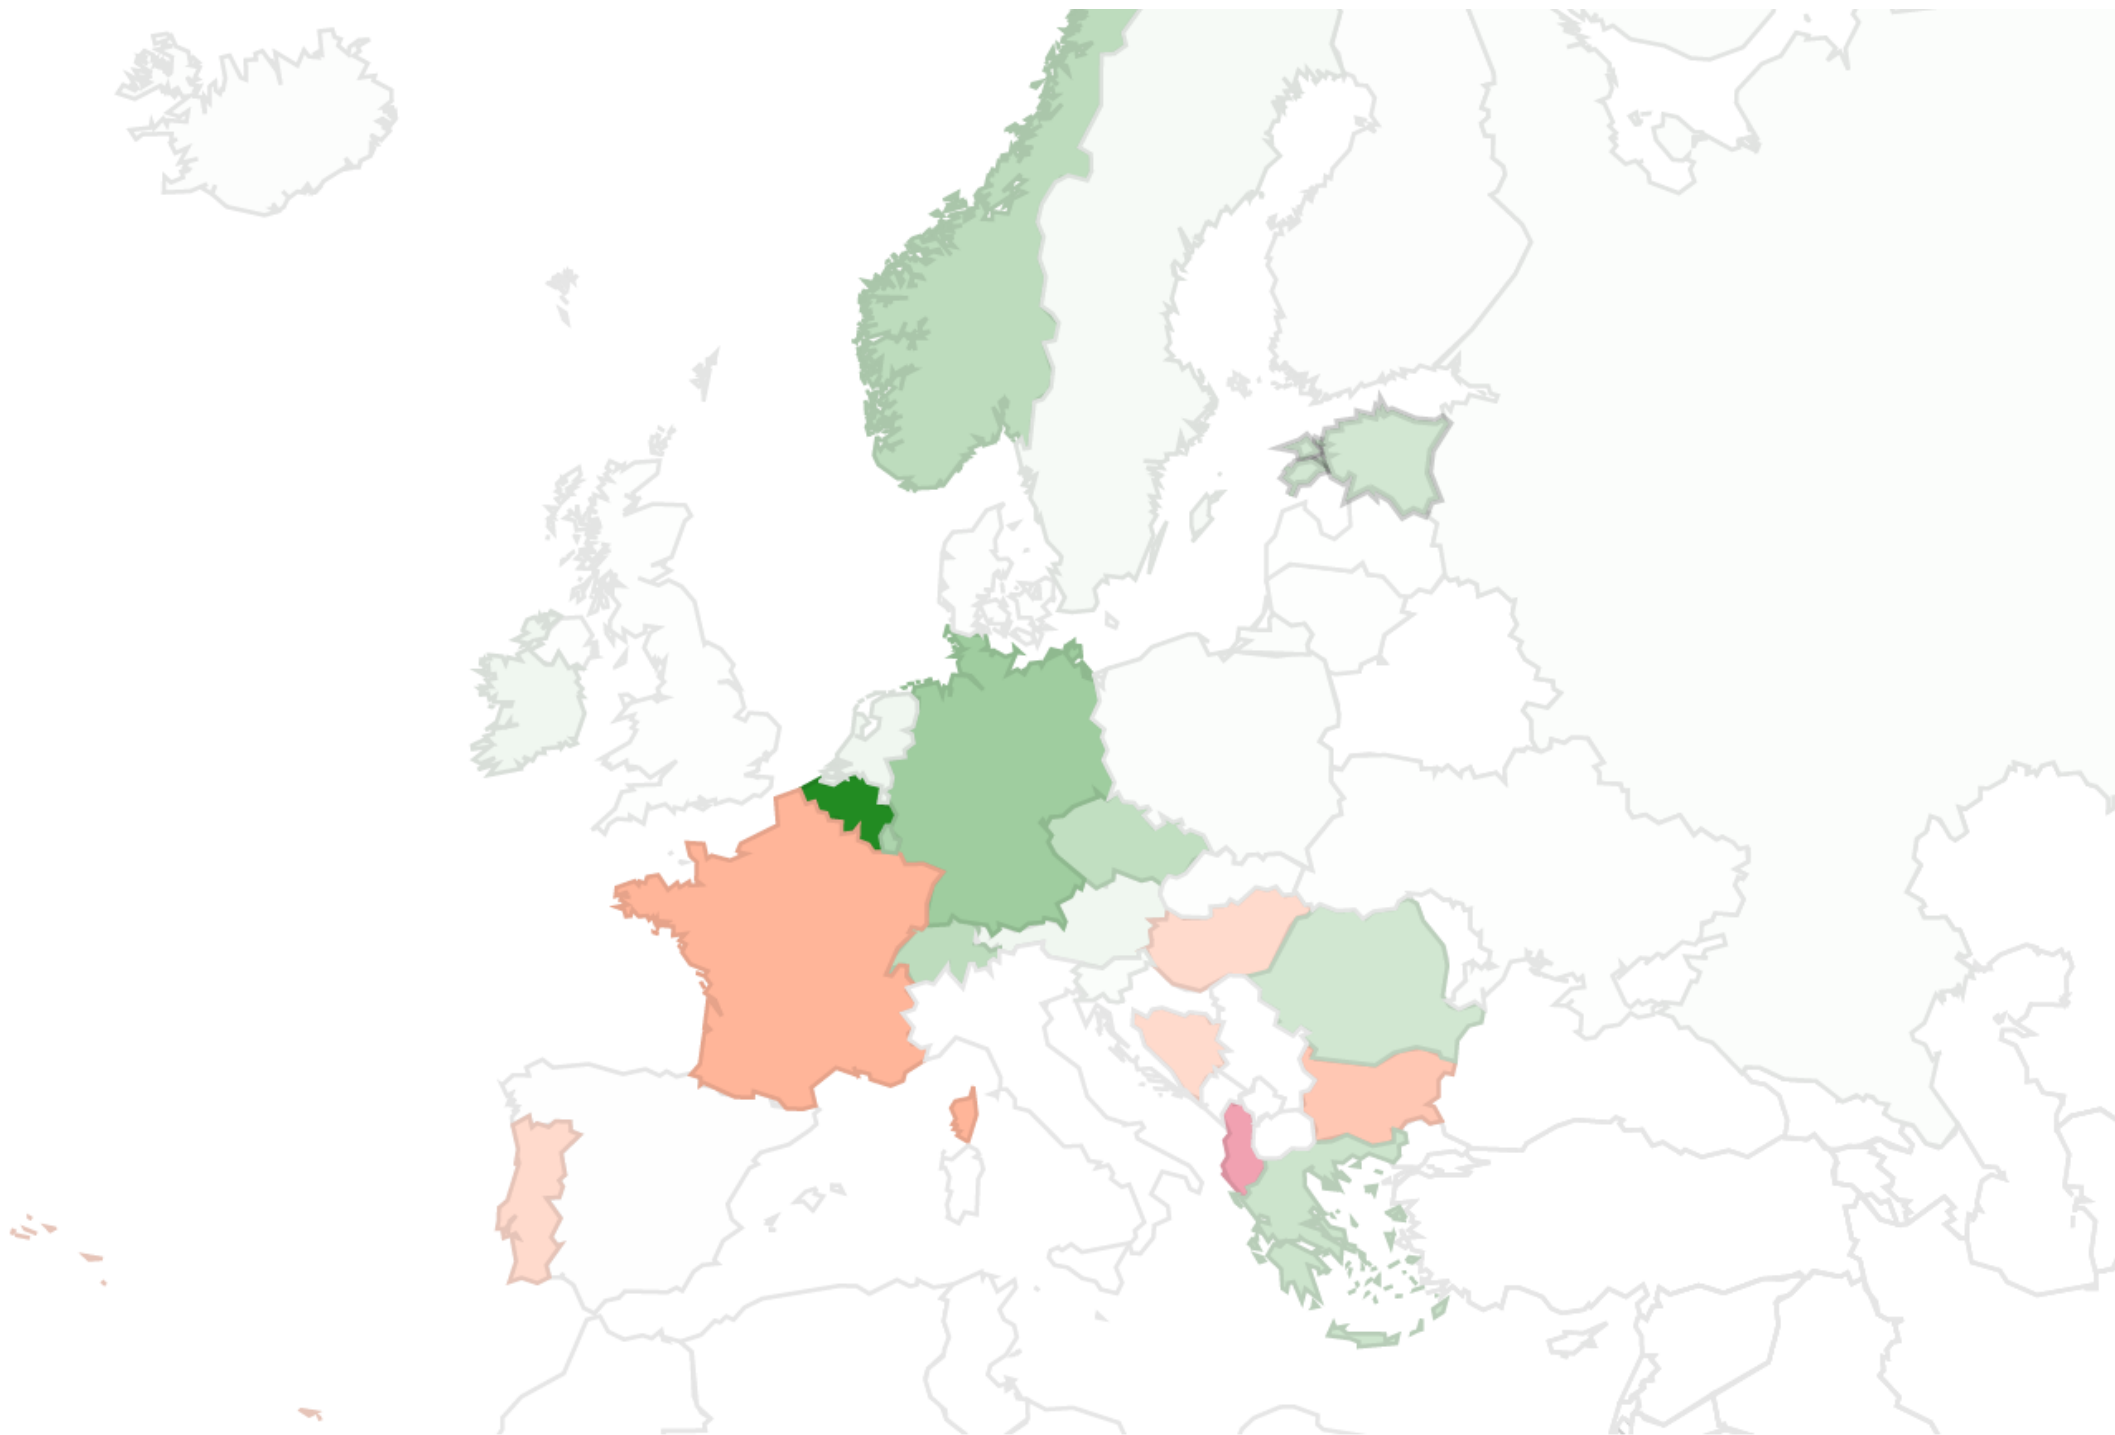
firstn . lastn @ telekom . ee

Estonia


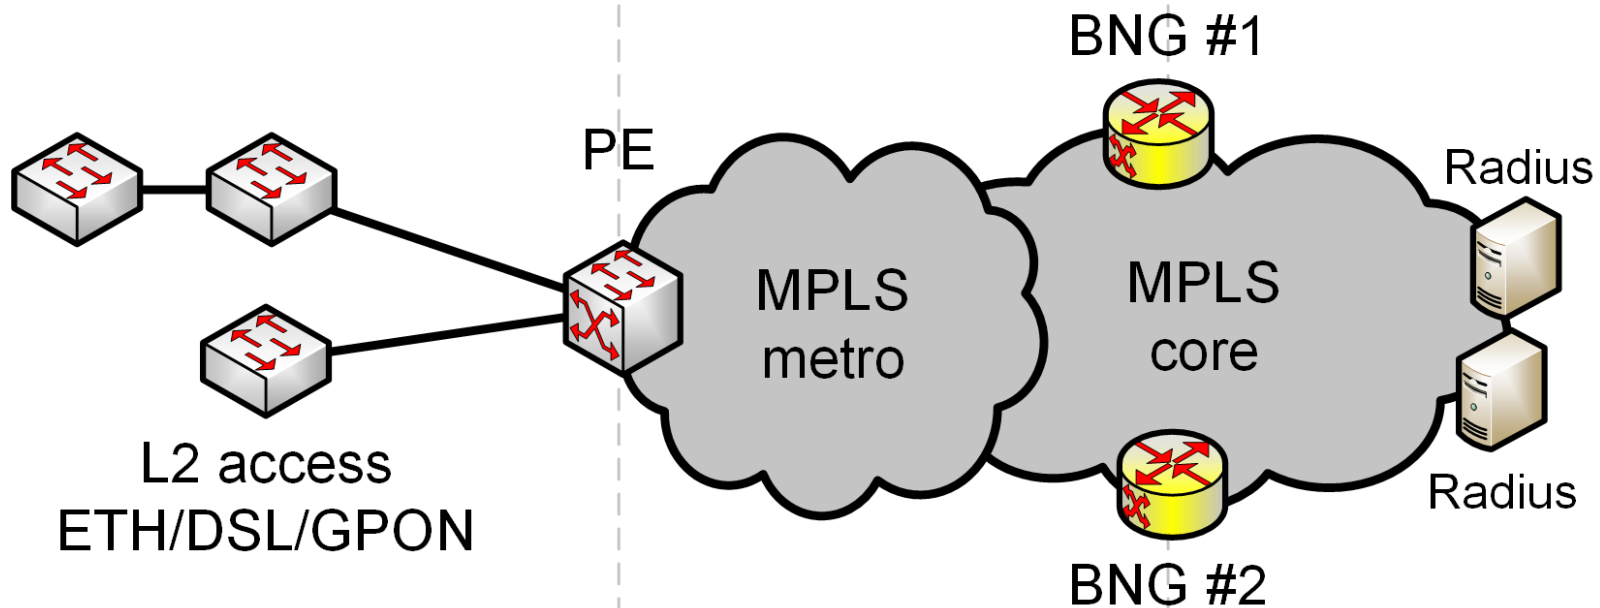


Select with your mouse to zoom-in on graph

Source: Akamai State of the Internet Report



- Have been working on it for 3+ years
- BNG replacement was the main driver
 - Proper, native, IPv6
- No replacements in aggregation or access
 - \$\$\$ and more importantly time
 - No tradeoffs in security
- New Linux (OpenWrt) based CPE
- Happy Eyeballs
- Deprecated PPPoE



Enablers

- MPLS
 - You get “free” L2 redundancy with EoMPLS
 - 6PE core (boring) (LDPv6/SR is finally shipping)
- BNG (aka subscriber management)
 - Centralised subscriber state
 - BringYourOwnRadius (SDN!)
 - Address usage logging
 - Unicast RA hack

Enablers

- Split horizon (aka private-vlan)
 - No traffic sent directly between subscribers
- Still need to protect against L2 threats
 - Limit number of MACs per AN port
 - Make sure BNG MACs couldn't be learned from AN ports
 - Limit unknown unicast, multicast, broadcast
- Need proxy-ARP for v4
 - No hacks needed for v6, routed via BNG

Bootup

- DHCPv4 & ARP
 - Authenticated by DHCP option 82
- RAs starting, unicasted to CPE MAC
 - Unsolicited
 - Horrible but required hack for ::/0
 - No on-link prefixes, only M-bit
 - Virtual link-local address at BNG
 - CPE still sends to multicast but only reaches BNG

Bootup

- DHCPv6
 - Delayed, IA_PD only, no IA_NA
 - ND info learned from Solicit
 - No LDRA support in ANs
 - Authenticated using cached v4 data from radius DB
 - Linked using MAC + port (L2 domain)
 - Radius req. rejected if v4 session not found
 - Possible DOS vector
- Shared shapers with v4
- V6 DNS information in DHCPv6

CPE

- /56 per subscriber
 - Cached for up to 24h
 - Take care to sync RA and DHCP lifetimes
- Divided to:
 - 1st /64 – loopback for CPE management
 - 2nd /64 – LAN (RA on-link, RDNSS, stateless DHCPv6)
 - 3rd /64 – public WIFI (if enabled)
- Ingress firewall configurable by customer
- ND logging for helpdesk

Can you do separate v6 BNG?

- Absolutely
- Shapers will not be shared
- But radius database will be
- Have to protect two sets of BNG MACs
- Can you do normal v4 and v6 BNG?
 - Probably but I wouldn't
 - BNG is not for everyone

What about the 4 weeks?

- We cheated
 - Deployed all v6 config during BNG change
 - Filter all 0x86DD at BNG
 - Waited until all CPE bugs are fixed
 - Replaced filter with one that allows v6 for certain MACs
 - Monitor carefully (IPFIX)
- **0** customer-affecting problems

Future

- Support other CPEs
- Static prefix from radius
- Static routes in LAN
 - address hints needed
- Hierarchical DHCPv6 PD
- VOIP and IPTV to IPv6
- Mobile network (VoLTE)

Stats

- 38000+ v6 enabled subscribers
- 81% of them have have at least one IPv6 enabled device in the LAN (70% have more than one)
- Content:
 - Zone.ee
 - Google + Youtube + GA
 - Akamai, Cloudflare, Limelight, Edgecast
 - Facebook
 - VKontakte, Yandex
 - Torrent & Hamachi

...