

# Karsolink IPv6 Deploy

## The post IPv4-runout ISP experience

AS204471

# WHO are WE?

- ▶ The network KARSOLINK is run by 2S Computers, founded in 2014 by two partners.
- ▶ Three owners & 2 employers
- ▶ We operate mostly in Trieste province in the north-est of Italy on the border with Slovenija.
- ▶ RIPE Members as a LIR since 2021
- ▶ AS204471

# How we started in 2018?

- ▶ Small WISP with less than 300 users
- ▶ Minimal budget
- ▶ Running on Mikrotik ROS6
- ▶ FTTH Uplink - 500Mbps and a /28
- ▶ NATed IP to customers
- ▶ Selling just FWA access

# NOW, 2024

- ▶ ISP/WISP with 1000 customers
- ▶ Bigger budget (but still not that big :-)
- ▶ Mikrotik ROS7
- ▶ Dual-stack backbone with a virtual overlay
- ▶ Redundant 10G Uplinks and own ASN & resources
- ▶ Multiple transit providers, three exchanges, private peerings
- ▶ NATed, public IP, public networks, backup links, multi-homed links, and IPv6 networks to customers
- ▶ Selling FWA and FTTC/FTTH
- ▶ RPKI - ROA & Prefix validation
- ▶ MANRS Members

# WHY IMPLEMENTING IPv6?

## THE MOTIVATION

- ▶ **Why not?**
- ▶ Lack of IPv4 addresses
  - ▶ Should we **invest** and buy IPv4 resources? - What will the value of a /24 be in 10 years?
- ▶ Offering customers free **public** IP assignments
- ▶ Get rid of NAT
  - ▶ Improving user experience
  - ▶ Cheaper network operations (LOGS)
  - ▶ Faster CPE, without NAT sessions table limits
- ▶ Get rid of private addresses
- ▶ End-to-end connectivity
- ▶ Routing only network
- ▶ **Natural network evolution path**

# IPv6, FIRST STEPS

- ▶ Check if all the equipment supports IPv6
- ▶ Check if router software modules support IPv6
  - ▶ IPv6 PD over PPPoE is not supported by PPPoE Server in ROS6
- ▶ IPv6 address plan -> brainstorming and RIPE Academy IPv6 course
- ▶ Backbone: The easy part of the work!
  - ▶ Initial setup and internal routing (OSPFv3, BGP, Route reflector cluster)
  - ▶ Basic security (Firewall, PtP / 126, Spoofing guard,...)
  - ▶ Announce the network, filtering
- ▶ DNS
  - ▶ Dual-stack forwarders
  - ▶ Dual-stack authoritative name servers
- ▶ RADIUS
  - ▶ Dual-stack user portal
  - ▶ DNS filtering / Parental control

# What to do with all those networks?

Only the second or third attempt of an IPv6 address plan will be proper!  
Take your time and scratch it down.

## ▶ 2a12:d8c0::/29

- ▶ 2a12:d8c0::/32 - Backbone
- ▶ 2a12:d8c1::/32 - Customers
- ▶ 2a12:d8c2::/32 - Future use
- ▶ ...
- ▶ 2a12:d8c7::/32 - Future use

**WE NEED AN  
ADDRESS PLAN!!**

# IPv6 address PLAN

## Backbone / Anycast / Loopback

- ▶ 2a12:d8c0::/29
  - ▶ 2a12:d8c0::/32
    - ▶ 2a12:d8c0:1000::/36
      - ▶ 2a12:d8c0:1000::/44 not used
      - ▶ 2a12:d8c0:1010::/44 Service Zone 1
      - ▶ 2a12:d8c0:1020::/44 Reserved
      - ▶ 2a12:d8c0:1030::/44 Reserved
      - ▶ 2a12:d8c0:1040::/44 Reserved
      - ▶ 2a12:d8c0:1050::/44 Service Zone 2
        - ▶ 2a12:d8c0:1050::/48
        - ▶ ...
        - ▶ 2a12:d8c0:105a::/48 Local Services
        - ▶ ...
        - ▶ 2a12:d8c0:105f::/48 PtP links /64
      - ▶ 2a12:d8c0:1060::/44 Reserved
      - ▶ 2a12:d8c0:1070::/44 Reserved
- ▶ 2a12:d8c0::/32
  - ▶ 2a12:d8c0:0000:0000::/64 - Anycast services
    - ▶ 2a12:d8c0::a/128 - DNS
    - ▶ 2a12:d8c0::5/128 - NTP
    - ▶ ...
  - ▶ 2a12:d8c0:ffff:0000::/64 - Loopbacks addresses
    - ▶ 2a12:d8c0:ffff::201/128
    - ▶ 2a12:d8c0:ffff::202/128
    - ▶ ...
- ▶ Multiple instances of the same service are distributed with OSPF or BGP
- ▶ Easy to deliver: one IP per service



# IPv6 address PLAN for Customers

## RIPE 690 - 4.2.1. /48 for everybody

- ▶ In this case, we replicate the running IPv4 topology -> Ease of use for all the involved personnel
- ▶ Routing optimization: only /40 are injected in the internal routing table, and exceptions, of course
- ▶ Zero networks are reserved or not used for practical reasons
- ▶ F networks are primarily used for point-to-point links
  - ▶ RIPE-690 - 4.1. Numbering the WAN link
- ▶ In case: 65.024 /48 networks for each /32.
  
- ▶ 10.30.0.0/16 - Customers space
  - ▶ 10.30.1.0/24 - Zone 1
    - ▶ 10.30.1.11/32 - Customer 11
    - ▶ 10.30.1.201/32 - Customer 201
    - ▶ 10.30.1.zzz/32 - Customer zzz
  - ▶ 10.30.11.0/24 - Zone 11
    - ▶ 10.30.11.11/32 - Customer 11
    - ▶ 10.30.11.zzz/32 - Customer zzz
  
- ▶ 2a12:d8c1::/32 - Customers space
  - ▶ 2a12:d8c1:0-F::/36 - Geographic Zone
    - ▶ 2a12:d8c1:11zz::/40 - GEO 1 / Zone 1
      - ▶ 2a12:d8c1:110A::/48 - Customer 11
      - ▶ 2a12:d8c1:11C9::/48 - Customer 201
      - ▶ 2a12:d8c1:11zz::/48 - Customer zzz
  - ▶ 2a12:d8c1:BBzz::/40 - GEO B / Zone 11
    - ▶ 2a12:d8c1:BB0A::/48 - Customer 11
    - ▶ 2a12:d8c1:BBzz::/48 - Customer zzz

# IPv6 address PLAN

## Business customers reservations

- ▶ 2a12:d8c0::/29
  - ▶ 2a12:d8c1::/32
    - ▶ 2a12:d8c1:0300::/44 - Business customers network
      - ▶ 2a12:d8c1:03[0-3]::/48 - not used
      - ▶ 2a12:d8c1:0304::/48 - Customer 1 /48
      - ▶ 2a12:d8c1:0305::/48 - Reserved /47
      - ▶ 2a12:d8c1:0306::/48 - Reserved /46
      - ▶ 2a12:d8c1:0307::/48 - Reserved /46
      - ▶ 2a12:d8c1:0308::/48 - Customer 2 /48
      - ▶ 2a12:d8c1:0309::/48 - Reserved /47
      - ▶ 2a12:d8c1:030A::/48 - Reserved /46
      - ▶ 2a12:d8c1:030B::/48 - Reserved /46
    - ▶ Each business customer, upon request, is given its network from a separate pool
    - ▶ For each /48 assigned a /46 is reserved for future use
    - ▶ Upon request a contiguous /47 or /46 can be assigned to the customer

# MAIN PROBLEMS in MASS DEPLOY

Mainly, all the problems reside on the customers' side


- ▶ Most CPEs have IPv6 disabled by default
- ▶ Not all CPEs support PD /56 and /48 -> Change CPE or upgrade the firmware
- ▶ Assign static IPv6 delegations to all RADIUS customers
  - ▶ RIPE-690 - 5.3. Why persistent prefix assignments are recommended
- ▶ Security -> direct access to CPE management web page
- ▶ IPv6 was disabled on many hosts (They thought that was causing problems...)
- ▶ Some customers are concerned about IPv6 security, mainly because a friend told them so. Many need to see the benefits of implementing it, and it could be costly.

# VENDOR SPECIFIC PROBLEMS


- ▶ Some parts of the backbone cannot yet be migrated to IPv6 or dual-stack
  - ▶ MPLS/LDP/VPLS/BGP-VPLS/TE are not yet ready in ROS7
  - ▶ Now, our network overlay runs on IPv6 ONLY VxLANs
- ▶ IPv6 PD over PPPoE is not supported by PPPoE Server in ROS6
  - ▶ IPv6 PD over PPPoE runs perfectly on ROS7
- ▶ IPv6 accounting on Mikrotik testing from ROS 7.15
  - ▶ Need to modify the RADIUS server and the customer portal
  - ▶ Helped the static delegations

# IPv4 vs IPv6 Traffic comparison

- ▶ The current total IPv6 adoption rate is 49%
  - ▶ It mainly depends on when the customer joined the network and the CPE that is using
- ▶ The overall IPv6 traffic varies from 10% to 30%
- ▶ Most OTT traffic will pass on IPv6 (Google, AWS, META, Netflix, Apple, Microsoft)



ETTYPE	MIN	MAX	AVERAGE	~95TH
IPv4	1.51Gbps	15.74Gbps	5.94Gbps	12.88Gbps
IPv6	215.67Mbps	4.92Gbps	1.29Gbps	4.37Gbps



SRC AS	MIN
32590: Valve Corporation	16.00bps
15169: Google	16.48Mbps

2024-04-10 23:40:00	
32590: Valve Corporation	0.00
15169: Google	371.61M
32934: Meta	3.48G
204471: Karsolink	279.25M
16509: Amazon.com	95.20M
41327: Fiber Telecom (FT)	120.60M
20940: Akamai Technologies	107.61M
714: Apple	88.62M
6185: Apple CDN	46.30M
54113: Fastly	849.75K
Other	54.73M

# IPv6 DEPLOY MAIN BENEFITS

- ▶ Reduced IPv4 traffic from 20% to 30%
  - ▶ Reduced traffic logs: 1000 customers generate 70 GB of logs annually, and the traffic is increasing. Logs must be stored using a compliant and “expensive” method.
  - ▶ Reduced load on ISP NAT devices
- ▶ Increased performance: traffic in IPv6 is routed ONLY!
- ▶ Greater efficiency and failover
  - ▶ We can use active/active uplinks
- ▶ Easier debugging
- ▶ No more ISP NAT & CPE NAT!!
- ▶ It could become a business case: help business customers to adopt IPv6

# NEXT IDEAS

- ▶ Migrate internal backbone to IPv6 ONLY
  - ▶ Most of the routers will be stateless
  - ▶ IPv4 with IPv6 next hop (Static routes only in ROS7 so far)
- ▶ Become an IPv6 beta tester for our VoIP provider
- ▶ IPv6 DAY
  - ▶ Invite the customers with their CPEs
  - ▶ Change or upgrade and configure IPv6 on the CPEs
  - ▶ Explain at best to the customers what it is all about
  - ▶ EAT & DRINK!

# THANK YOU!

vasja@karsolink.com

Let us know if we should do it differently. 😊

[test-ipv6.karsolink.com](http://test-ipv6.karsolink.com)