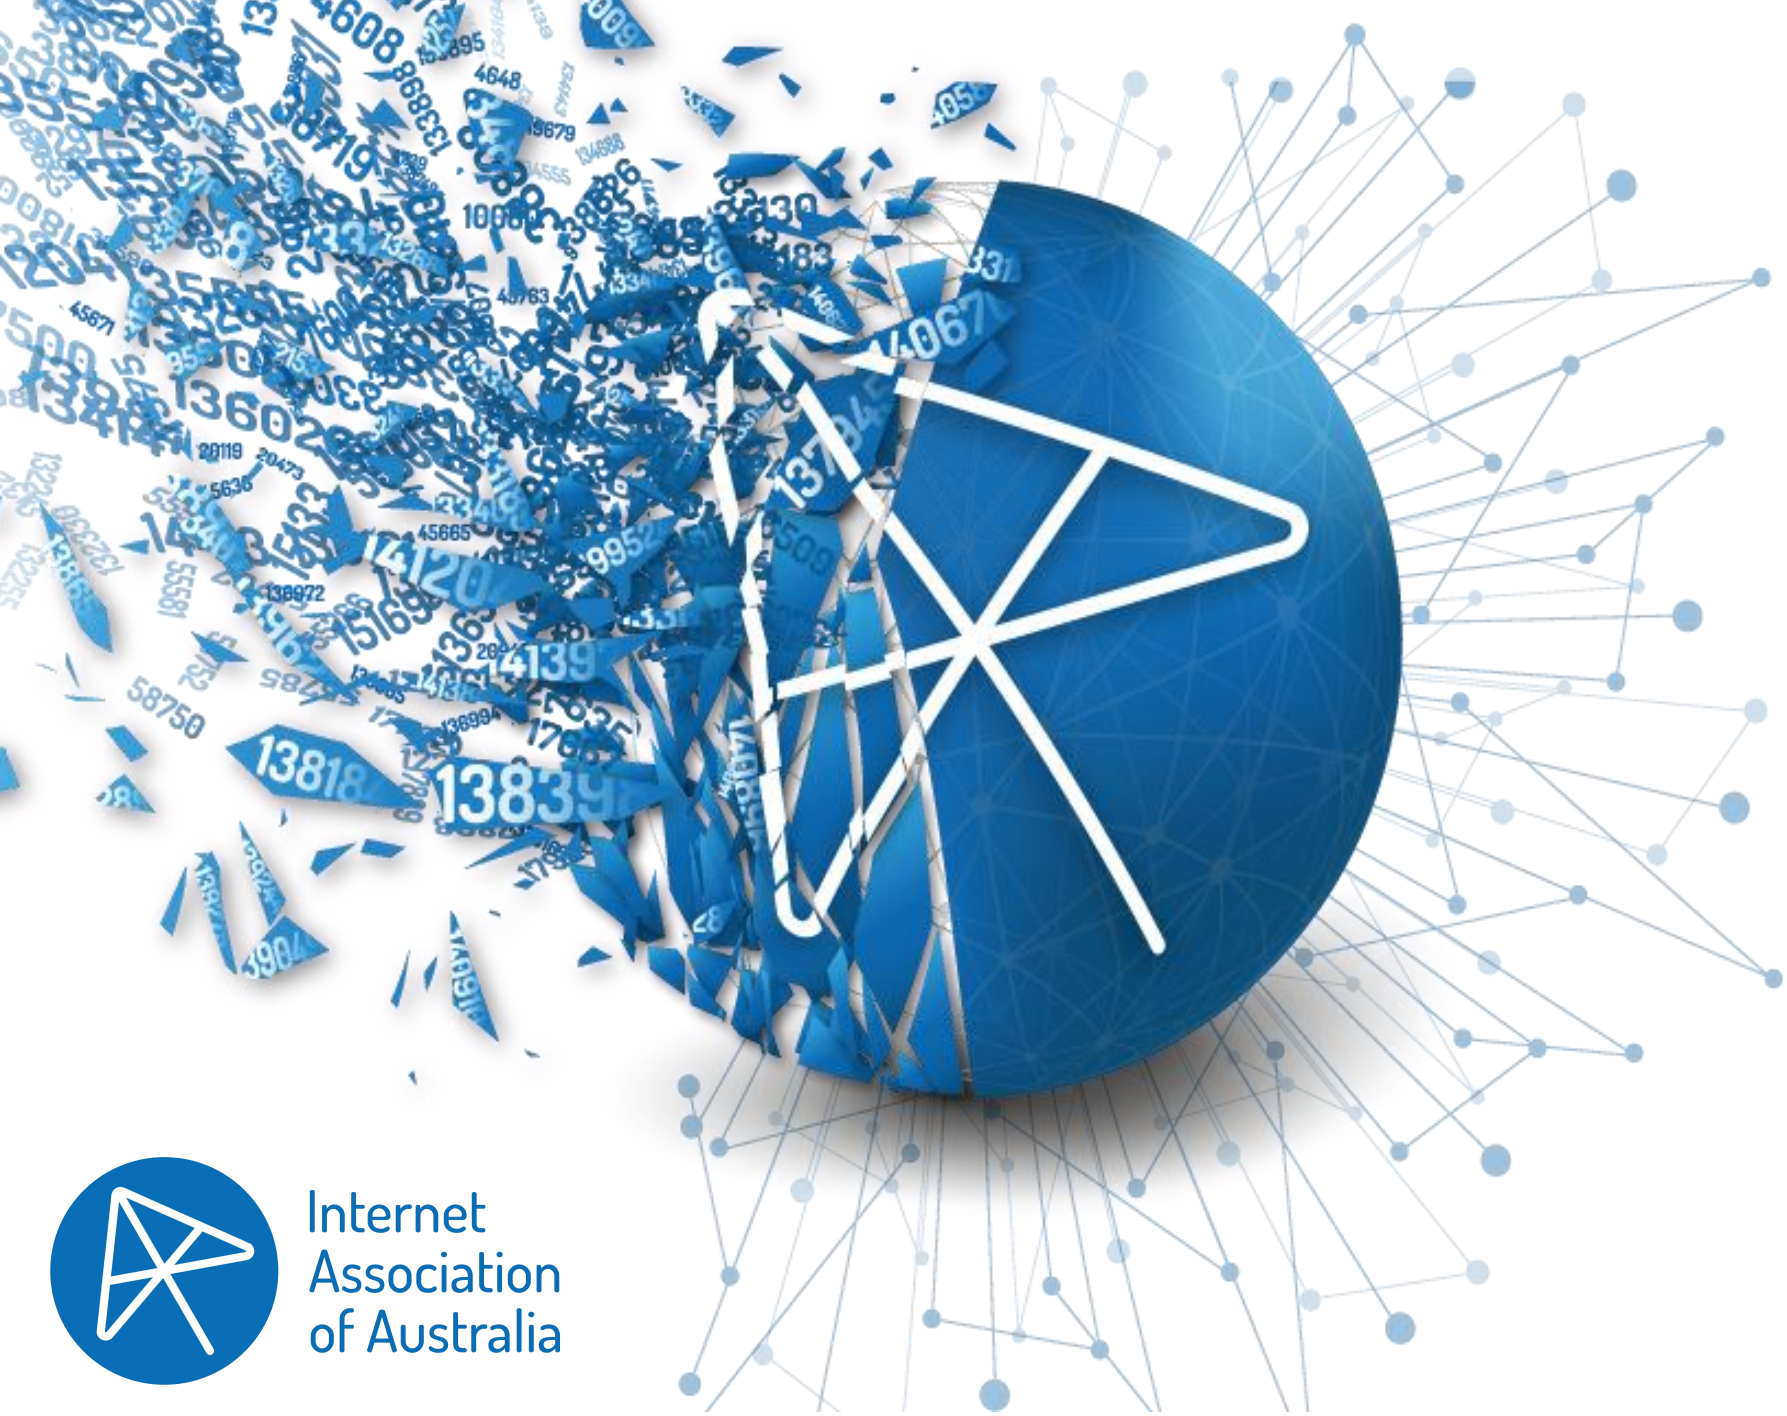

RFC 8950

An IXP Perspective

Matthew Kobayashi
Peering Engineer

Internet Association of Australia

APIX 2025

Introduction

- Peering Engineer at:
  - Internet Association of Australia (AS7606/AS10084)
  - New Zealand Internet Exchange (AS63830)
- Experience across the Australian telecommunications industry
  - Optus
  - Alcatel PABX
  - Superloop
  - Megaport
  - IAA and NZIX

Internet Exchanges today

- Layer 2 fabric with Layer 3 route servers
  - Some filters on peer ports to block problematic BUM traffic
    - ICMPv6 Router Advertisements
    - STP family
    - Etc.
- Increasingly becoming much more
  - Content delivery and caches
  - Private VLAN/VPWS services
  - Cloud on-ramps
  - ...and more

Internet Exchanges today

- Data plane is typically a L2VPN overlay
  - VPLS
  - VXLAN
  - SR-MPLS
- Control plane separation (EVPN) is becoming the de facto standard
  - Enables techniques such as ARP suppression and multi-homing
- IPv4 and IPv6 dual-stack addressing
  - Public address space used to avoid conflicts with participants' internal RFC 1918 and RFC 6598 ranges
- BGP for advertising NLRI (Network Layer Reachability Information)

What's the problem?

- Unallocated public IPv4 space is dwindling
  - IANA (Internet Assigned Numbers Authority) allocated the last /8s to RIRs (Regional Internet Registries) in February 2011
  - RIRs progressively exhausted their IPv4 pools between 2011 and 2019
- IX growth continues, but additional IPv4 space is difficult to get from RIRs
  - Layer 2 fabric and single broadcast domain requires a contiguous IPv4 block, making the problem worse for IXPs

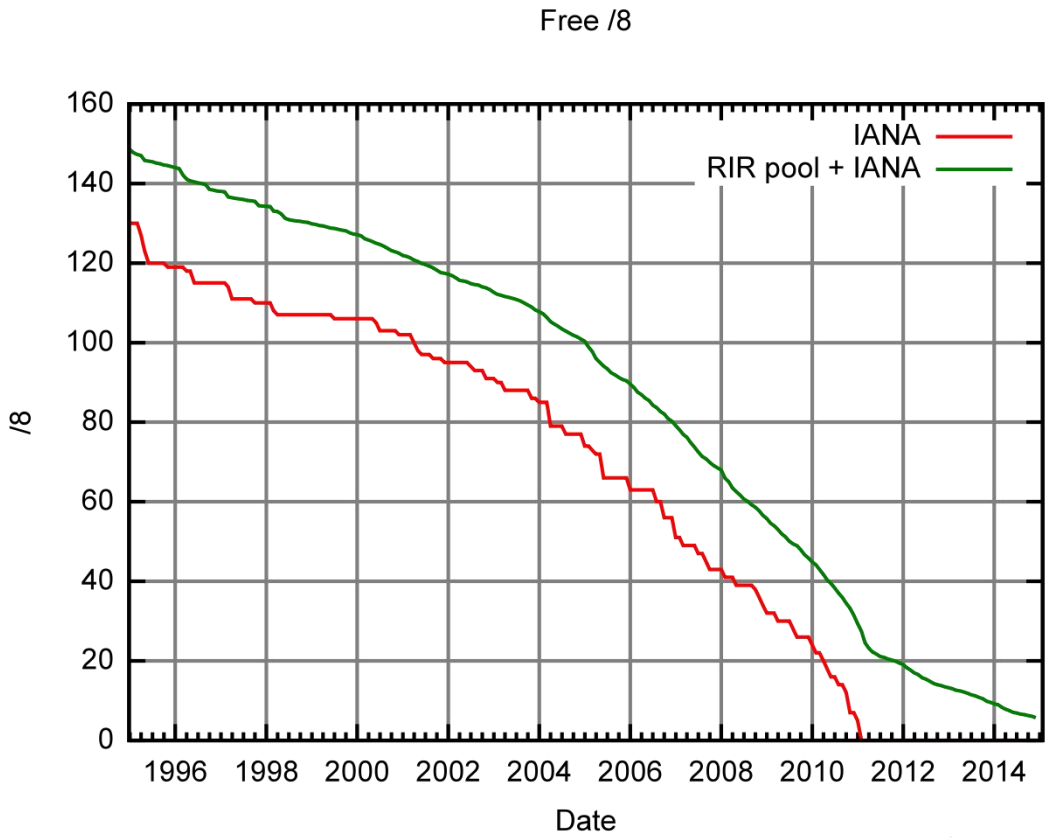

What's the problem?

- Average private sale price of IPv4 space reached US$60/address in late 2021
  - Prices have trended down since then, roughly US$25/address today
  - The market is volatile, with a price spread of US$27/address in the past month alone
  - Sale volumes of IPv4 space are steady, demand has not slowed
  - Quite expensive, and a potential heavy burden on IXP budgets

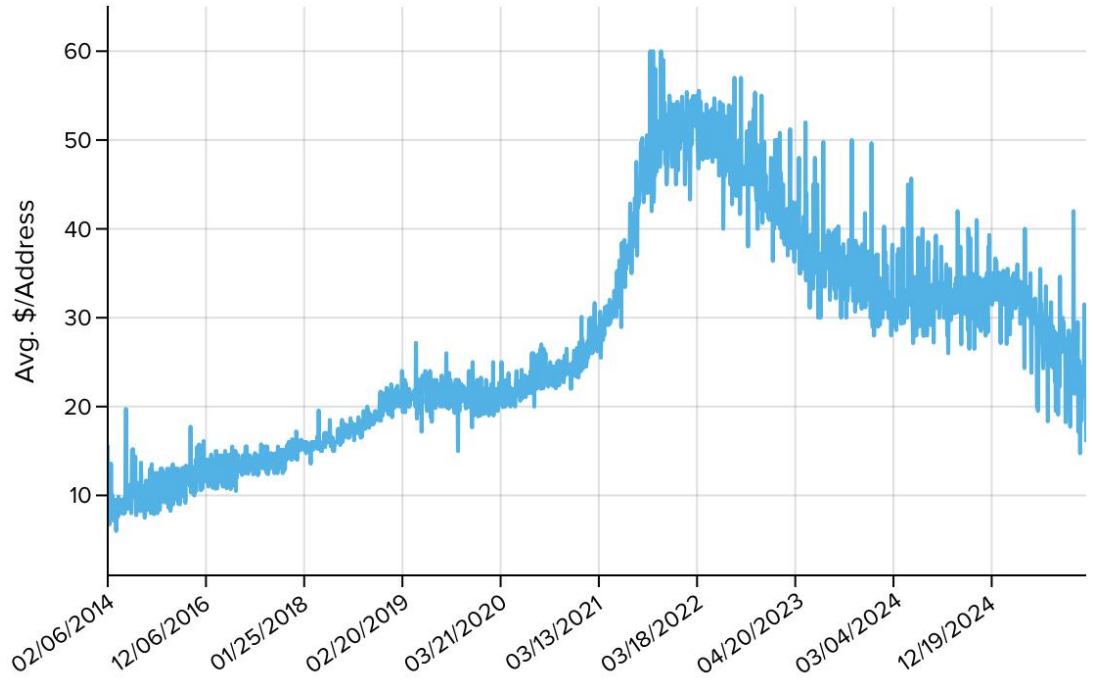

RFC 8950

- Enables IPv4 NLRI advertisements with IPv6 next-hop via BGP
  - Fully interoperable with RFC 5549
- Currently requires IPv6 GUA (globally-unique address) per RFC 2545 Section 3
  - Unfortunately, we cannot use IPv6 link-local addressing (fe80::/10)
  - A Standards Track draft produced by the IETF Interdomain Routing working group aims to enable the use of IPv6 link-local addressing for NLRI next-hops

Software Support

Software	Minimum Version	Notes
BIRD	2.0.8	- Requires Linux kernel 5.2+
OpenBGPD	8.8	
ExaBGP	4.1.0	- Cannot program Linux netlinks for RFC 5549
FRRouting	7.0.0	– Requires Linux kernel 5.2+ – Versions prior to 9.1.3 send ICMPv6 RAs if `extended-nexthop` is enabled
GoBGP	Any	

The Euro-IX RFC 8950 working group maintains a list of vendors with RFC 8950 support:
https://github.com/euro-ix/rfc8950-ixp

Vendor Support

Vendor	NOS	Minimum Version
Arista	EOS	4.22.1F
Cisco	IOS-XR	7.3.3
Juniper	JUNOS	21.2
MikroTik	RouterOS	7.20beta5
Nokia	SR-OS	20.2.R1
Nokia	SR Linux	20.06
VyOS	VyOS	1.2.2

The Euro-IX RFC 8950 working group maintains a list of vendors with RFC 8950 support:
https://github.com/euro-ix/rfc8950-ixp

Unsupported Vendors and Platforms

Vendor	NOS	Comments
Cisco	IOS-XE	
Cisco	NX-OS	Not Tested

The Euro-IX RFC 8950 working group maintains a list of vendors with RFC 8950 support:
https://github.com/euro-ix/rfc8950-ixp

Final Thoughts

- Vendor support for RFC 8950 is widespread
  - Cisco being a notable exception
- IXPs should consider changing their policies
  - Update route server software to a version with RFC 8950 support
  - Only allocate IPv4 addresses to participants upon request
  - Encourage participants to use RFC 8950 and return their allocated IPv4 address
  - Any policy changes should be weighed against participants' equipment support for RFC 8950
    - Participant MAC address OUIs can give a good passive indication
- Use of IPv6 link-local addressing for BGP NLRI next-hop may one day make IP allocation to participants unnecessary?

Questions?