The Routing and Addressing in the Internet - 2019 in Review

Geoff Huston
Chief Scientist, APNIC

Through the Routing Lens …

There are very few ways to assemble a single view of the entire Internet

The lens of routing is one of the ways in which information relating to the entire reachable Internet is bought together

Even so, its not a perfect lens, but it can provide some useful insights about the entire scope of the Internet

25 Years of Routing the Internet

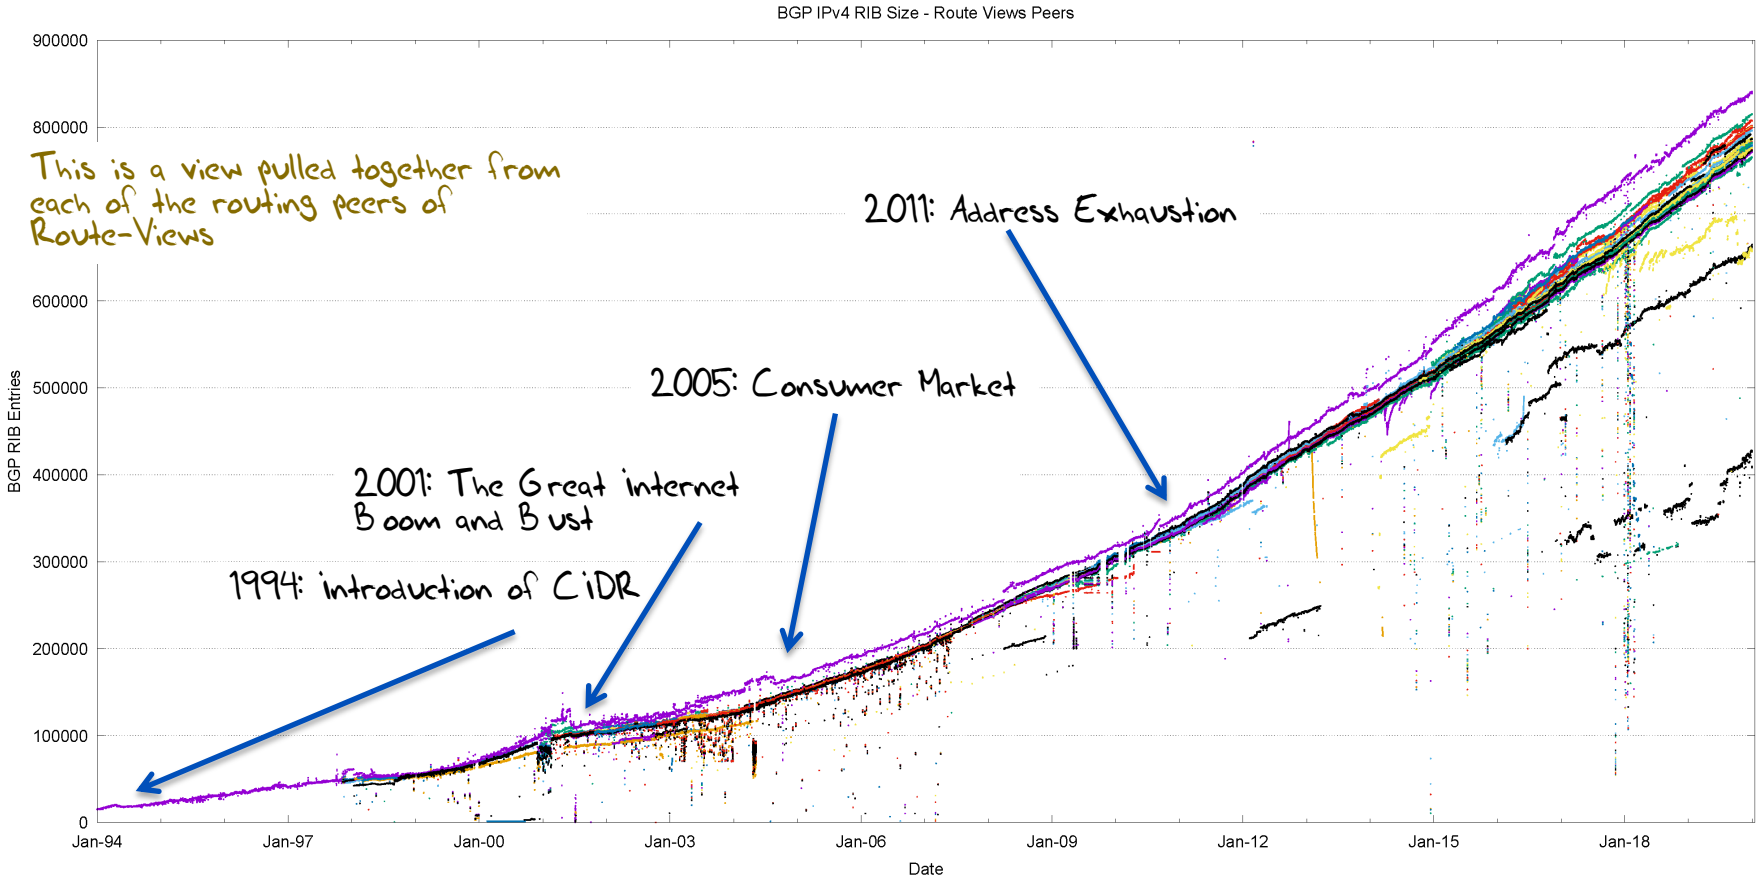

2017-2019 in detail

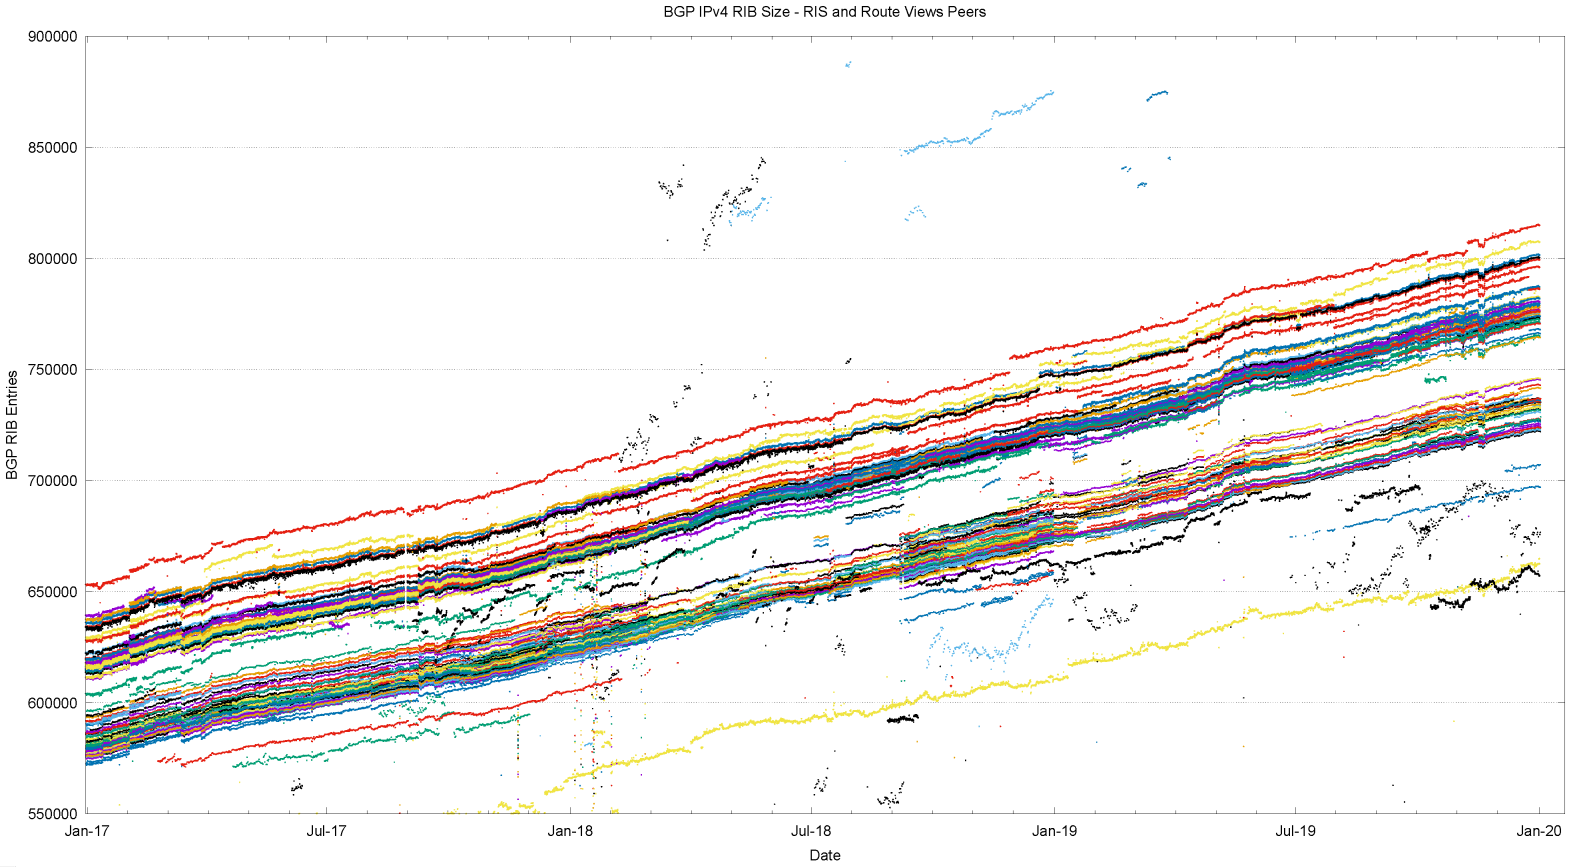

2017-2019 in detail

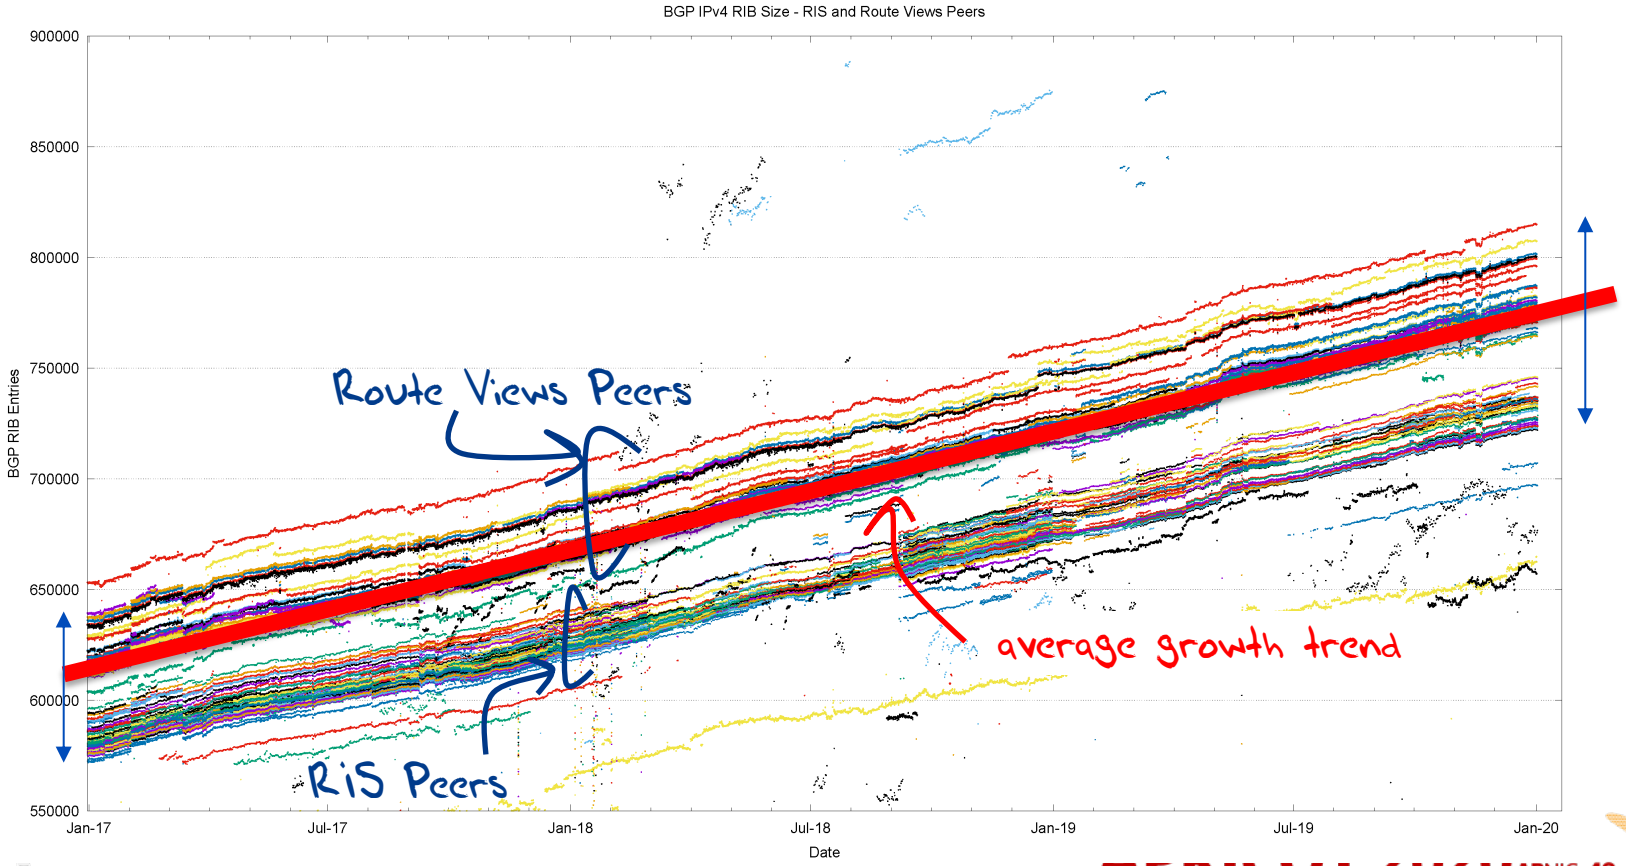

Routing Indicators for IPv4

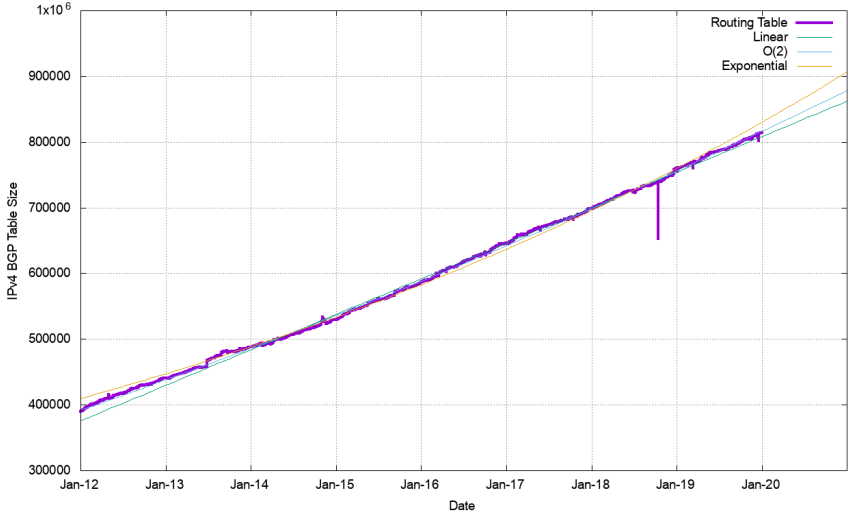

Routing prefixes – growing by some 51,000 prefixes per year

AS Numbers – growing by some 3,400 prefixes per year

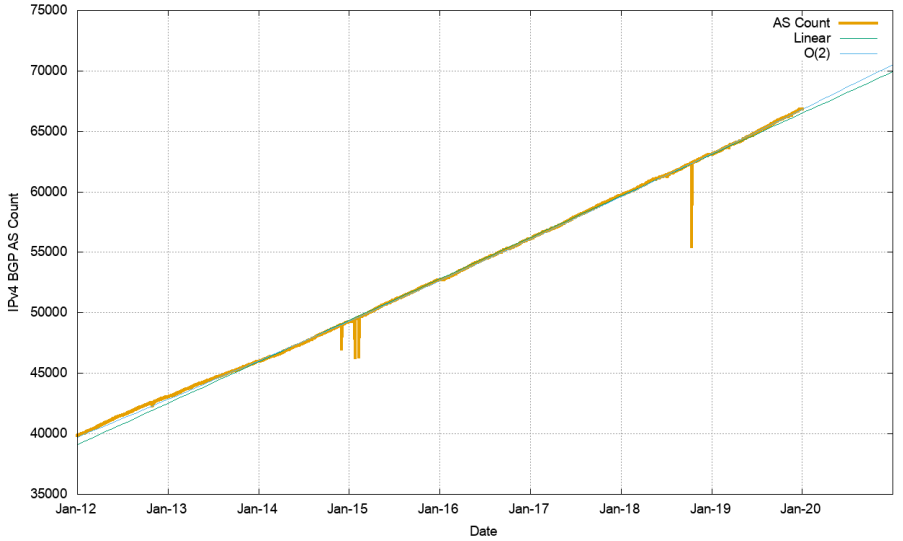

Routing Indicators for IPv4

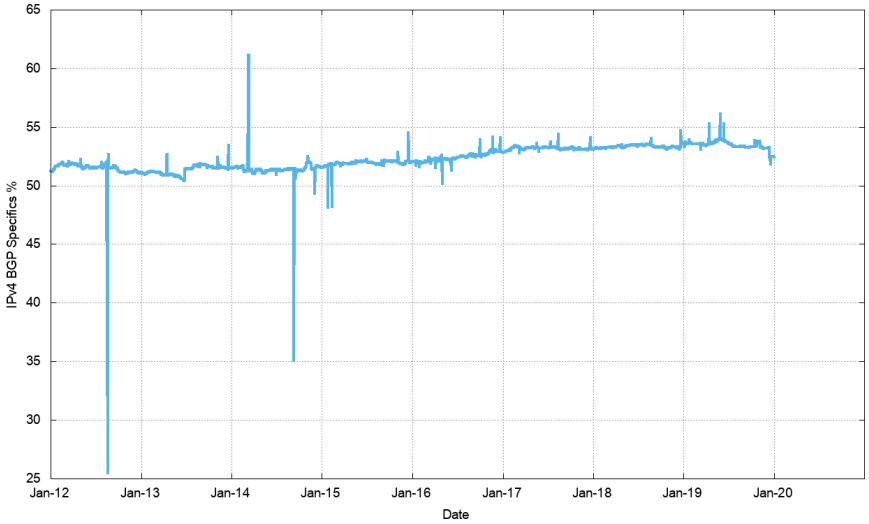

More Specifics are still taking up slightly more than one half of the routing table

But the average size of a routing advertisement continues to shrink

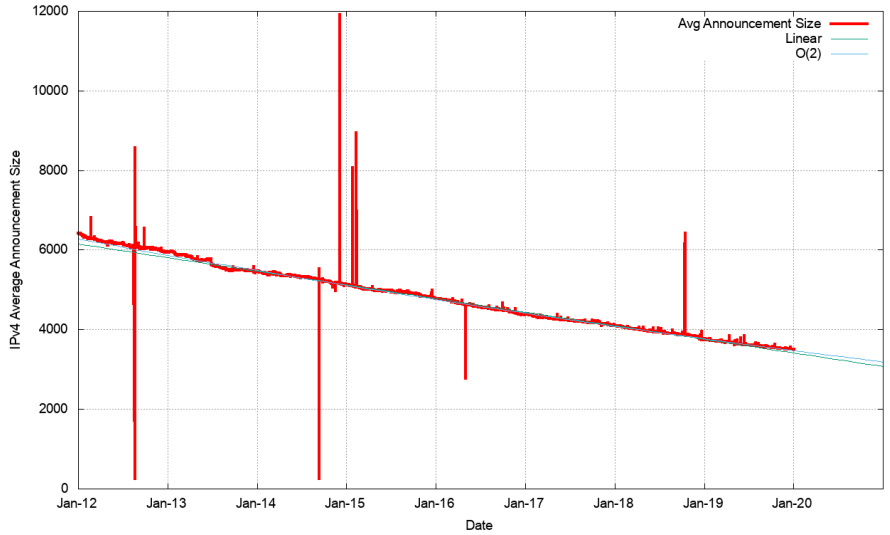

Routing Indicators for IPv4

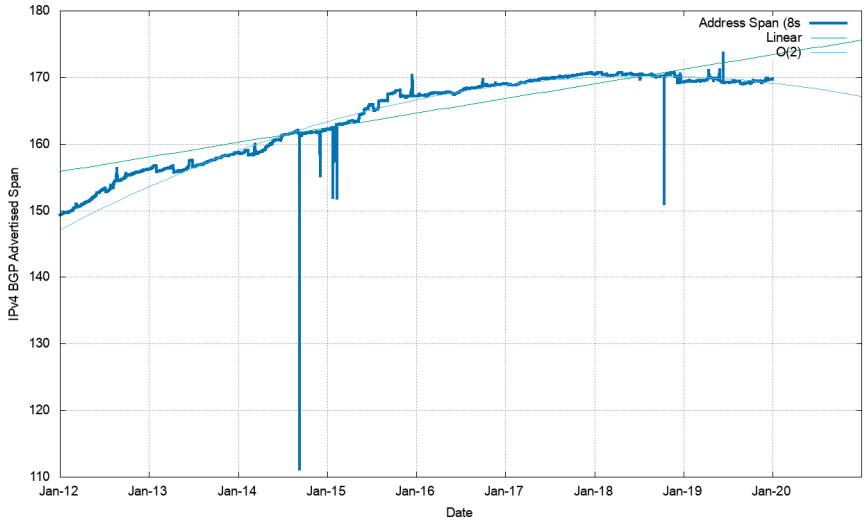

Address Exhaustion is now visible in the extent of advertised address space

The "shape" of inter-AS interconnection appears to be relatively steady

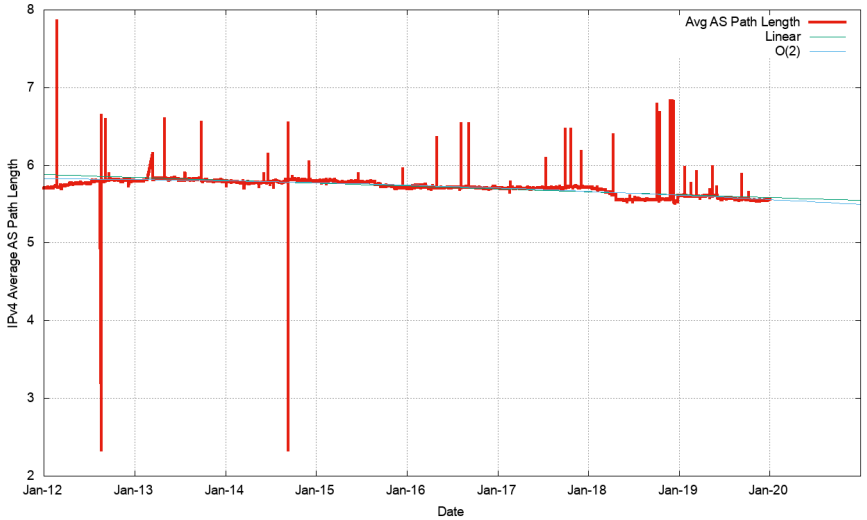

AS Adjacencies (AS131072)

54,697 out of 66,928 ASNs have 1 or 2 AS Adjacencies (82%)

2,195 ASNs have 10 or more adjacencies

10 ASNs have >1,000 adjacencies

4,498	AS6939	HURRICANE - Hurricane Electric, Inc., US
4,291	AS3356	LEVEL3 - Level 3 Communications, Inc., US
3,947	AS174	COGENT-174 - Cogent Communications, US
1,803	AS6461	ZAYO Bandwidth, US
1,722	AS3257	GTT-Backbone, DE
1,649	AS7018	ATT-INTERNET4 - AT&T Services, Inc., US
1,422	AS2914	NTT America, US
1,377	AS3549	LVLT – Level 3 Parent, US
1228	AS1299	TELIANET Telia Carrier, SE
1,148	AS209	CENTURYLINK, US

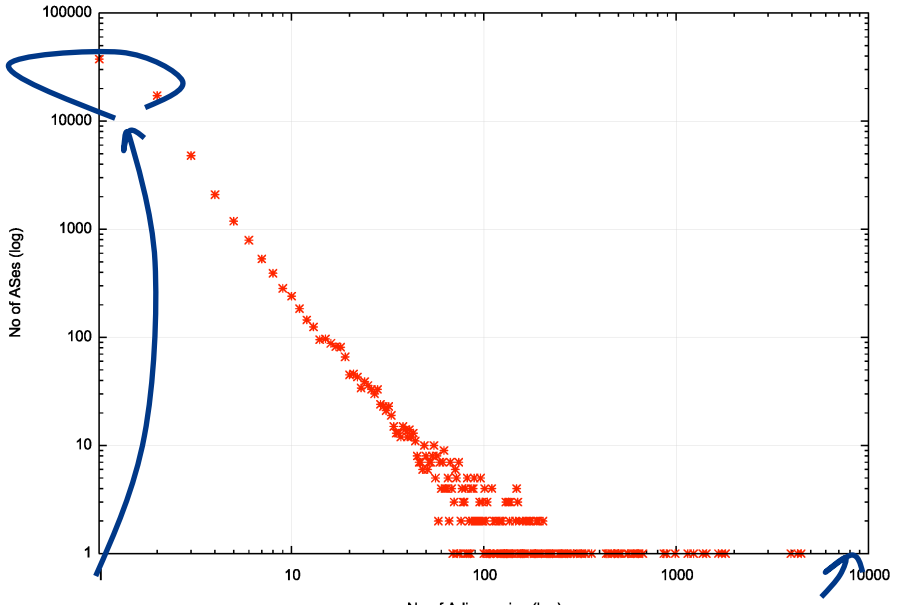

What happened in 2019 in V4?

Routing Business as usual – despite IPv4 address exhaustion!

- From the look of the growth plots, its business as usual, despite the increasing pressures on IPv4 address availability
- The number of entries in the IPv4 default-free zone reached 800,000 by the end of 2019
- The pace of growth of the routing table is still relatively constant at ~51,000 new entries and 3,400 new AS's per year
  - IPv4 address exhaustion is not changing this!
  - Instead, we appear to be advertising shorter prefixes into the routing system

What about IPv4 Address Exhaustion?

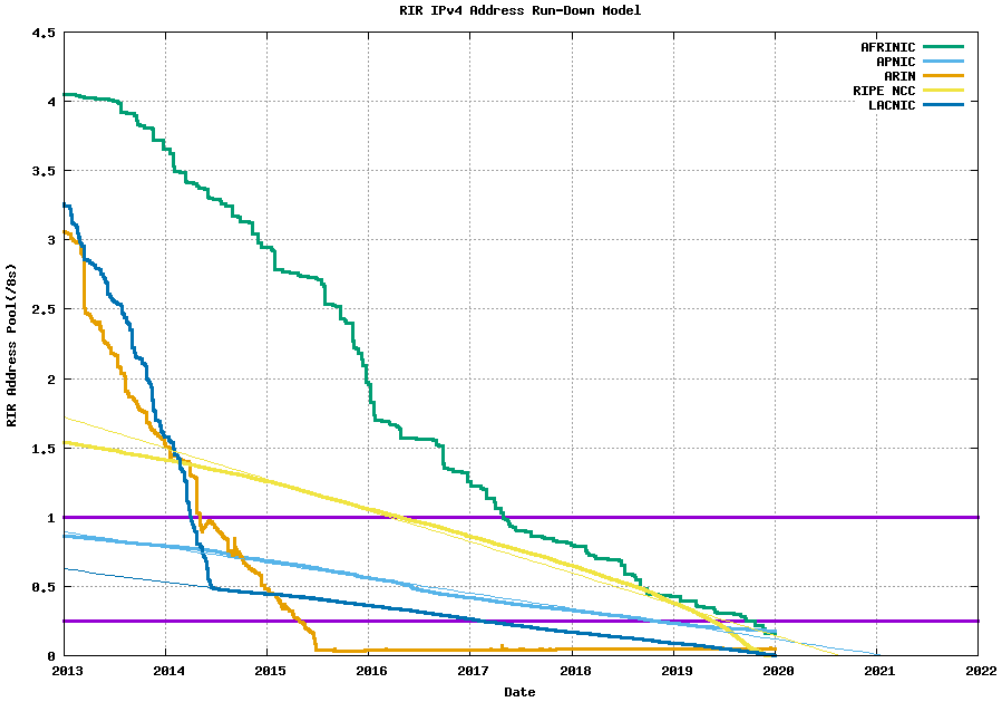

RIR Address Pool runout projections as of the start of 2020:

ARIN	– no free pool left
AFRINIC	– July 2020
LACNIC	– no free pool left
APNIC	– January 2021
RIPE NCC	– no free pool left

Post-Exhaustion Routing Growth

- What's driving this post-exhaustion growth?
  - Transfers?
  - Last /8 policies in RIPE and APNIC?
  - Leasing and address recovery?

Advertised Address "Age"

2010

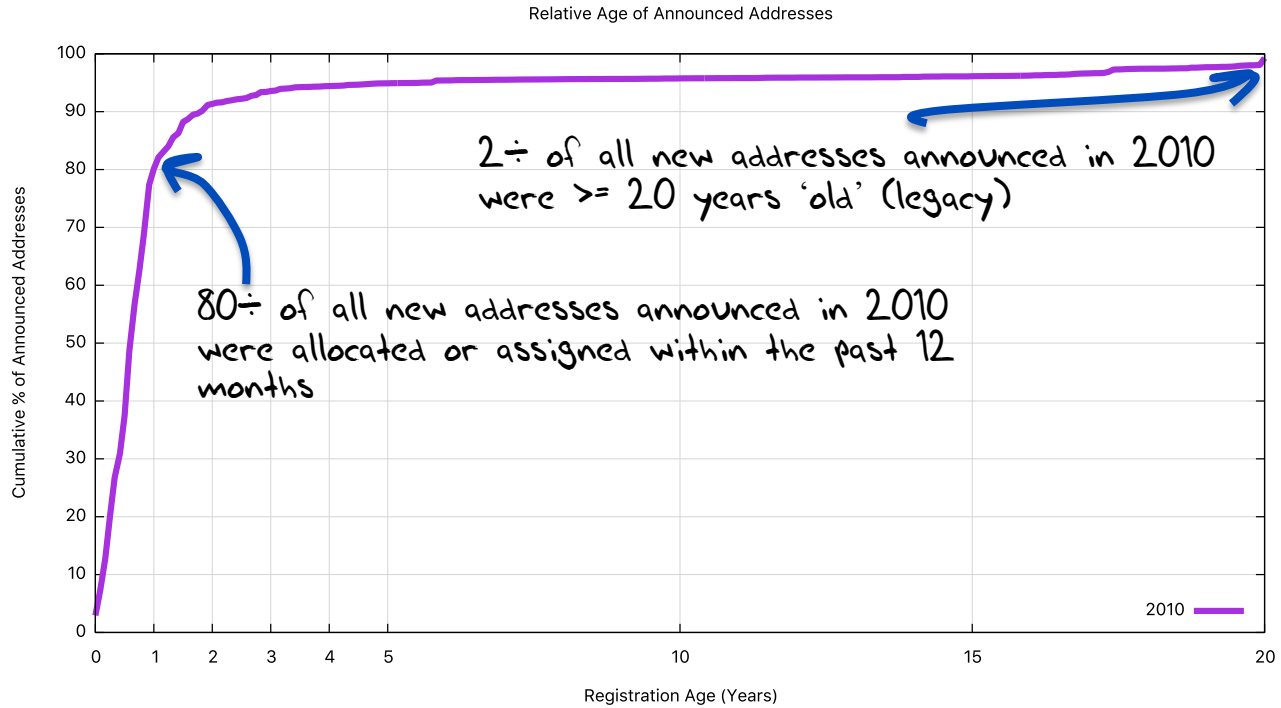

Advertised Address "Age"

2019

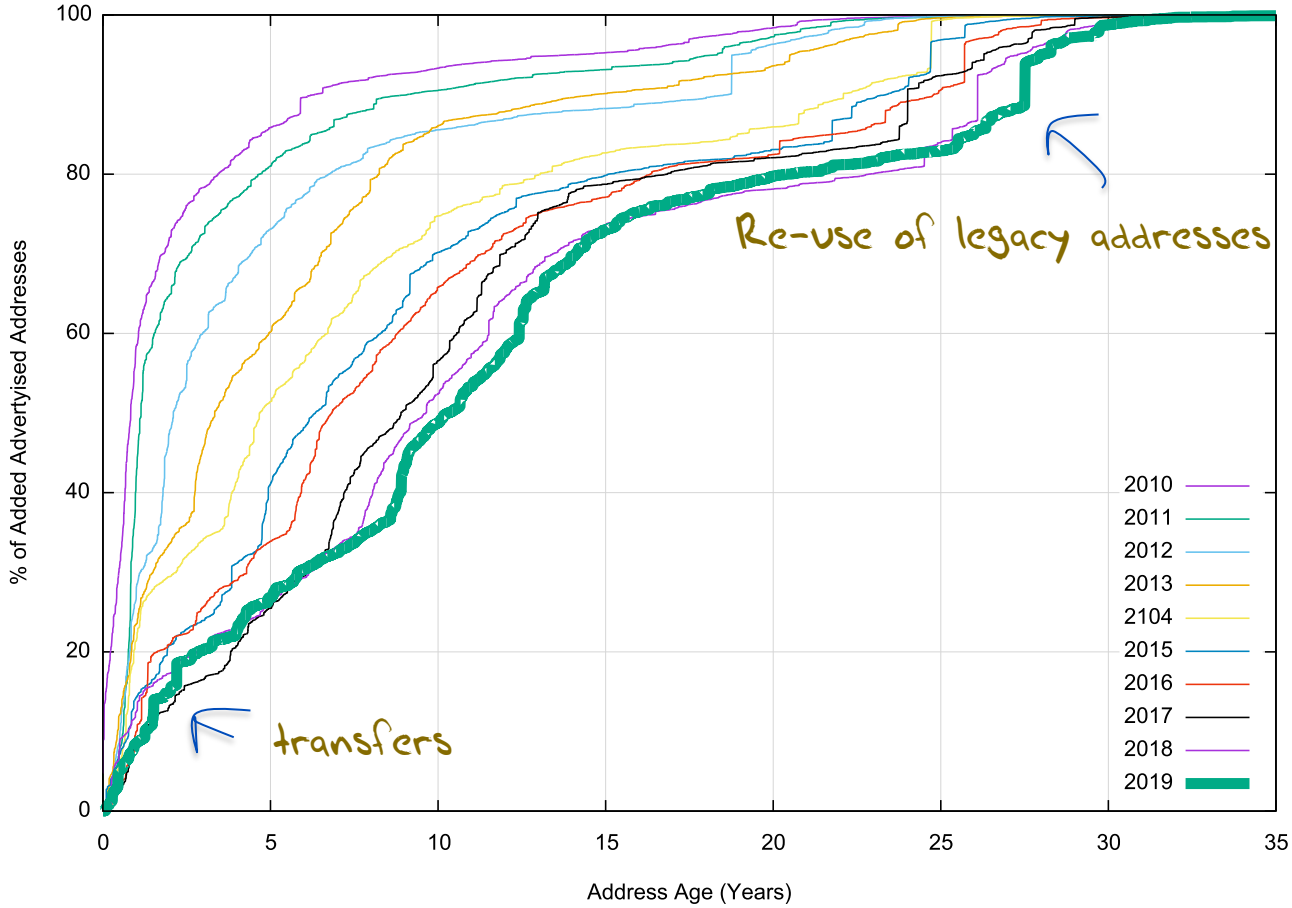

2000 - 2019: IPv4 Advertised vs Unadvertised

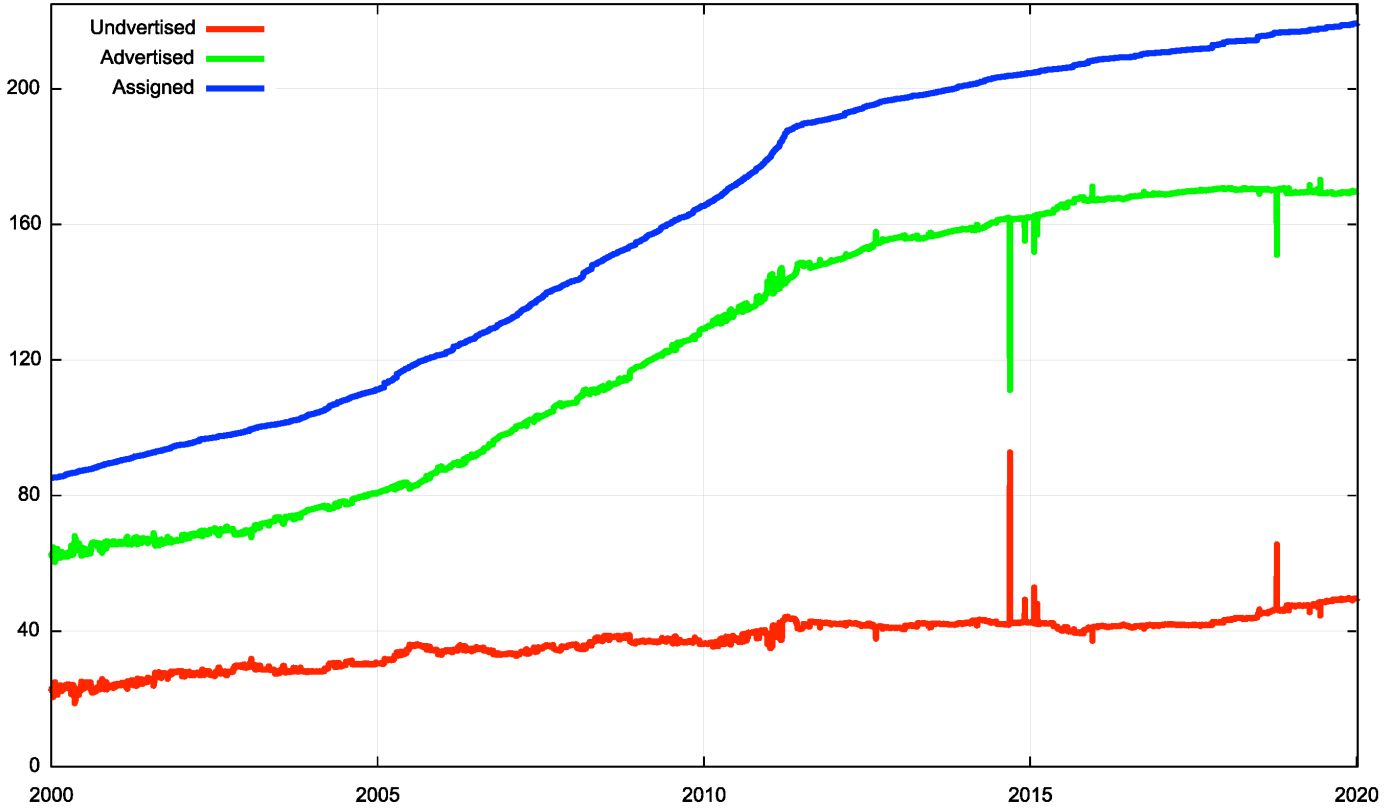

2005 - 2020: Unadvertised Addresses

2019: Assigned vs Recovered

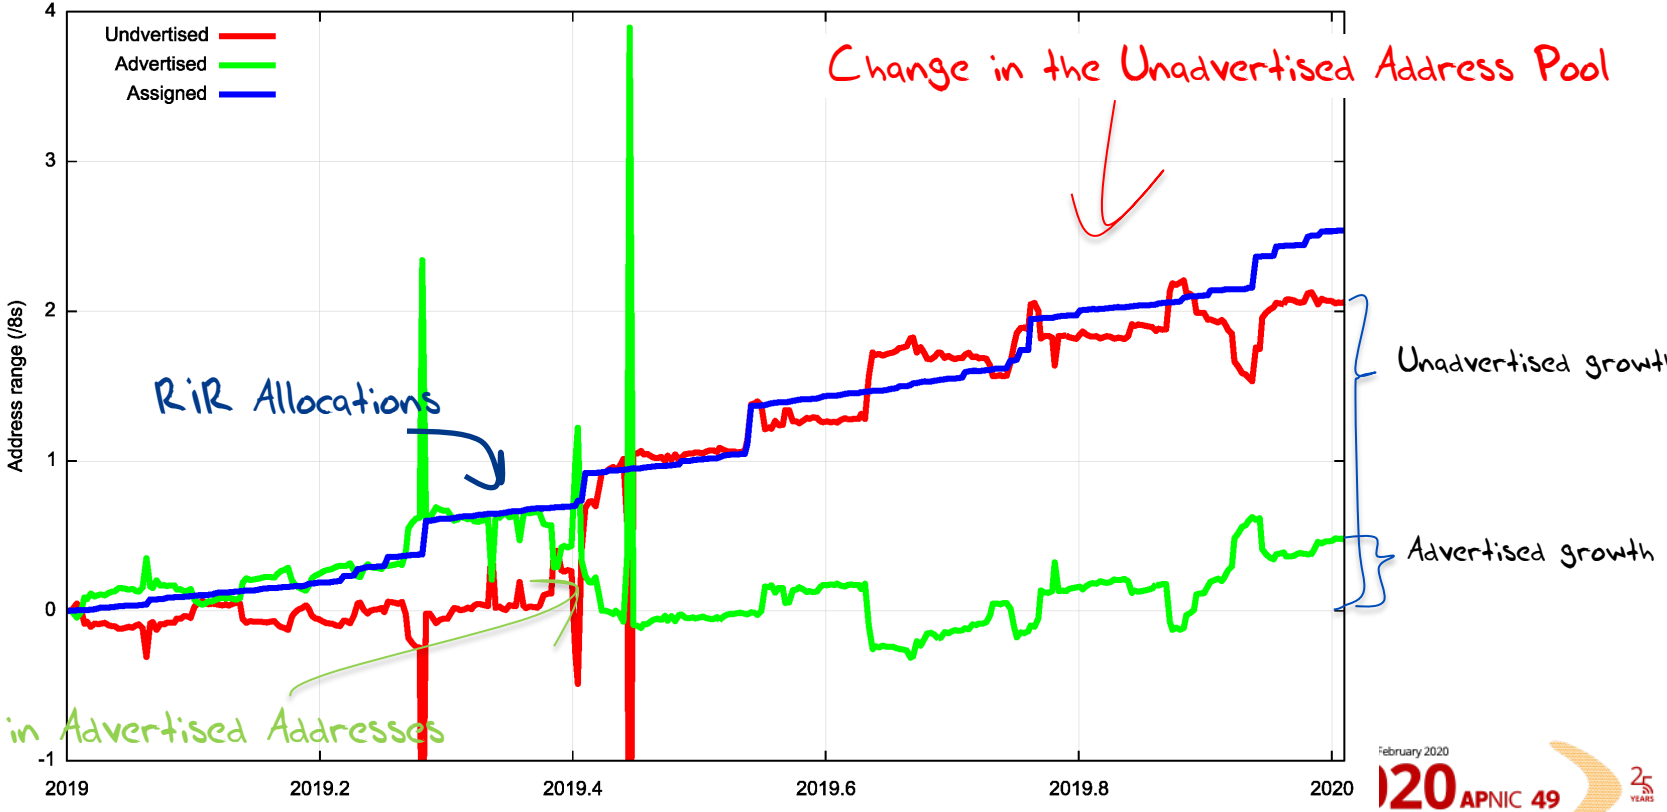

V4 in 2019

- The equivalent of 0.4 /8s were **added** to the routing table across 2019
- Approximately 2.5 /8s were assigned by RIRs in 2019
  - 0.38 /8s assigned by the RIPE NCC (last /8 allocations)
  - 0.27 /8's assigned by Afrinic
  - 0.09 /8s were assigned by LACNIC
  - 0.06 /8s were assigned by APNIC (last /8 allocations)
  - 1.7 /8s assigned by ARIN (transfers)
- And a net of 2.1 /8's were added to the pool of unadvertised addresses

In 2019 we saw legacy blocks transferring away from ISPs / end user sites and heading towards cloud SPs.

The Route-Views View of IPv6

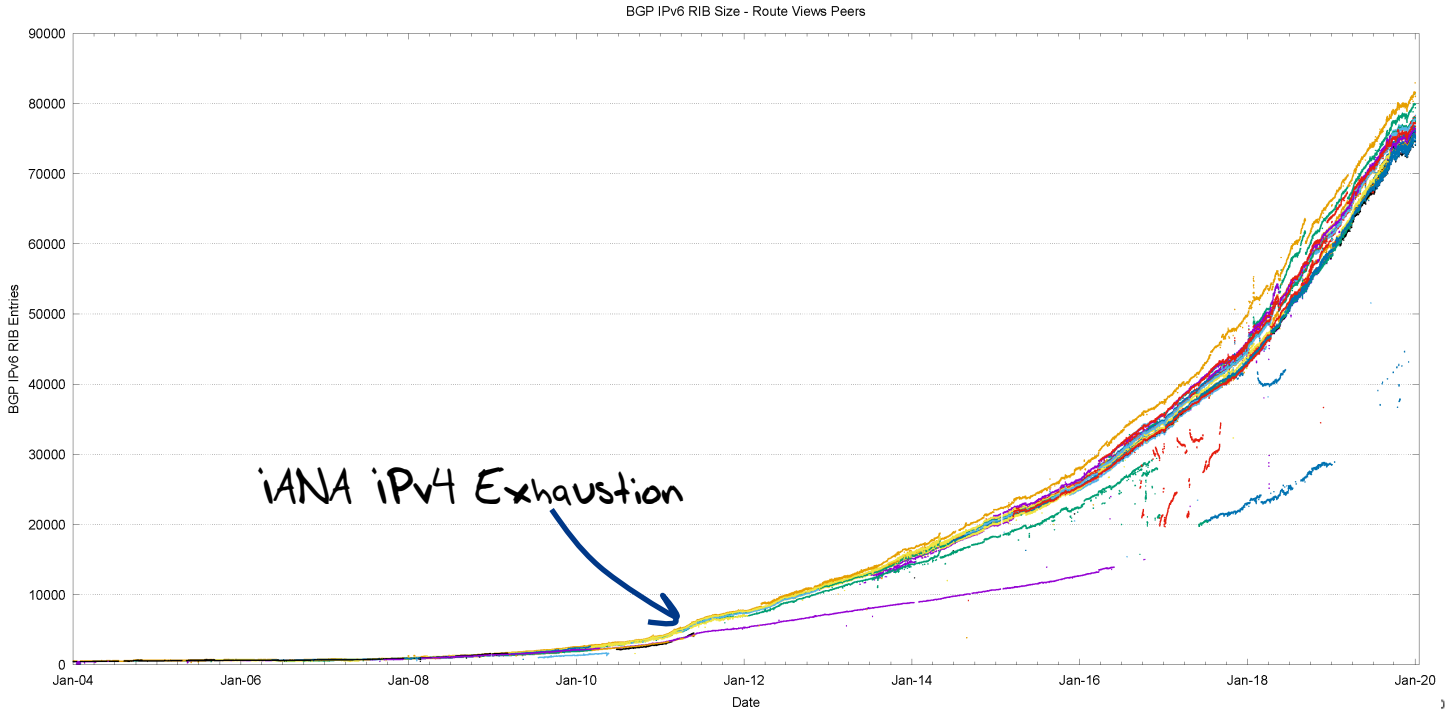

2018-2019 in Detail

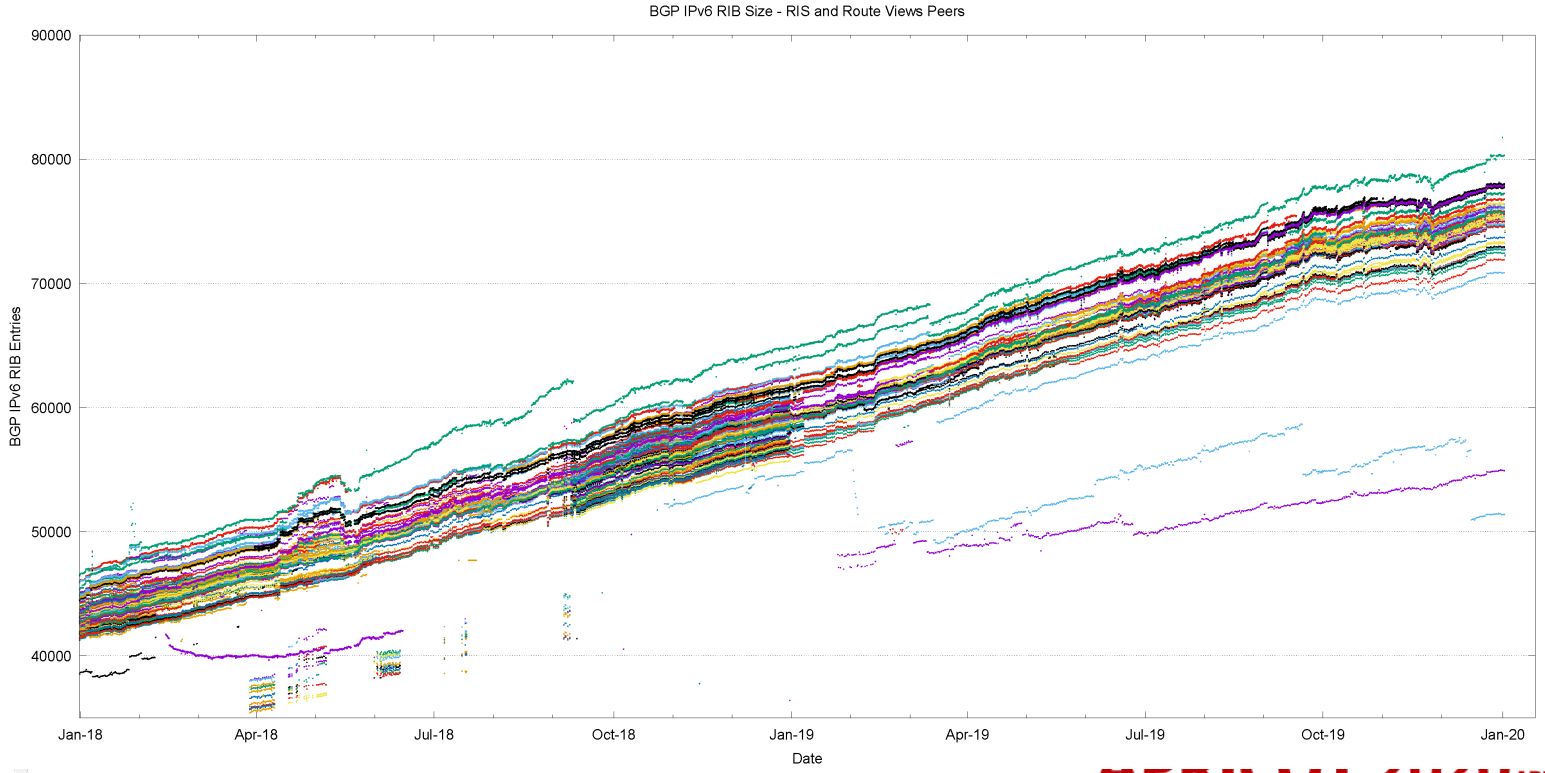

Routing Indicators for IPv6

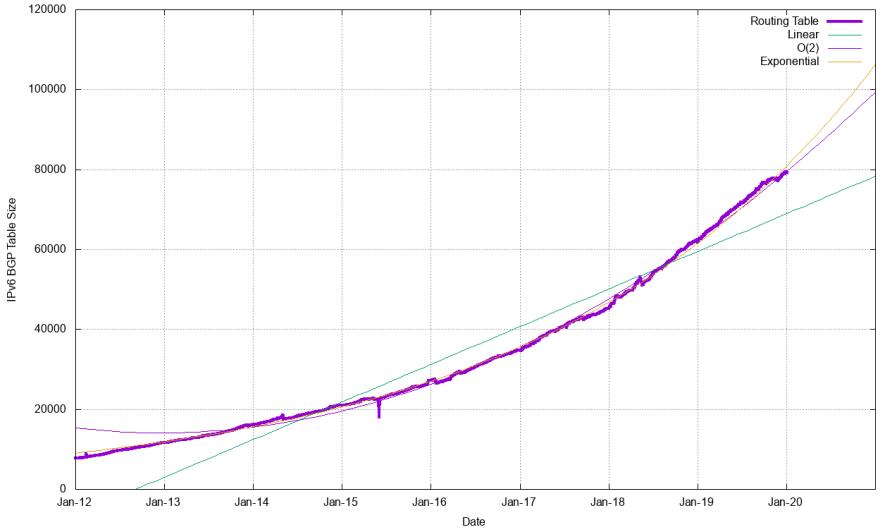

Routing prefixes – growing by some 17,000 prefixes per year

AS Numbers – growing by some 2,000 ASNs per year (which is 60÷ the V4 growth)

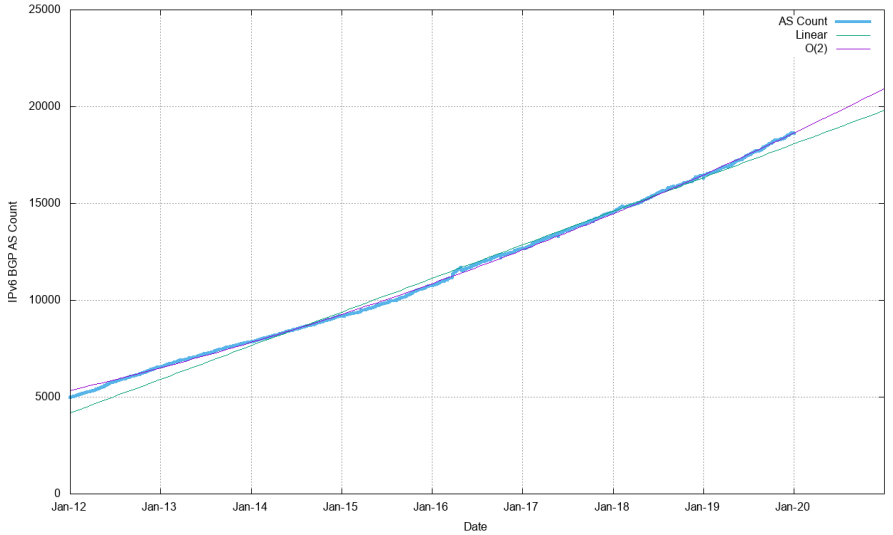

Routing Indicators for IPv6

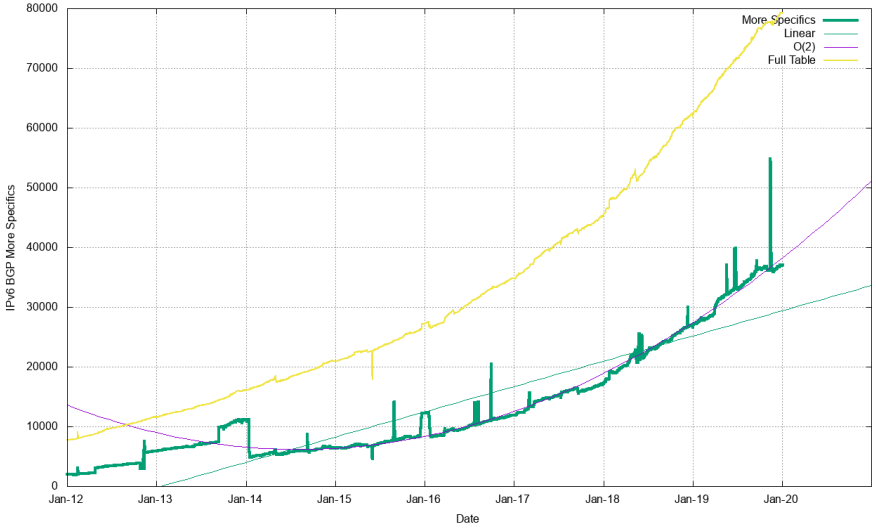

More Specifics now take up one half of the routing table

The average size of a routing advertisement is getting smaller

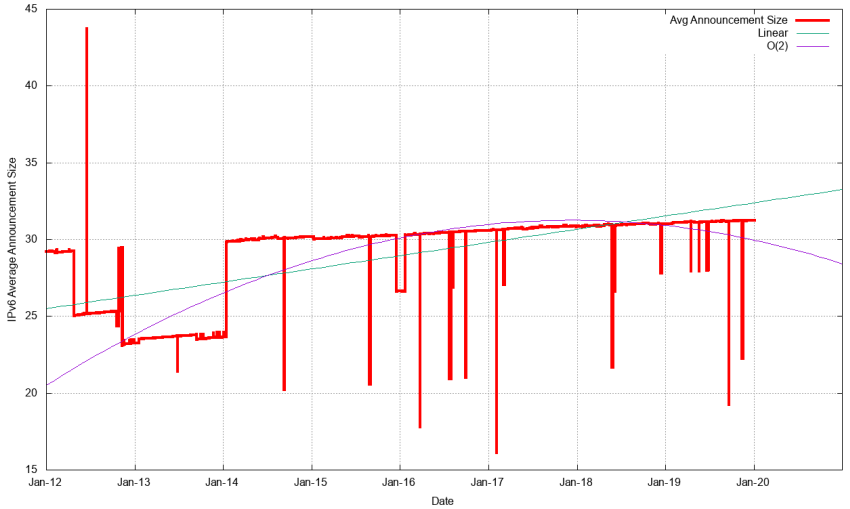

Routing Indicators for IPv6

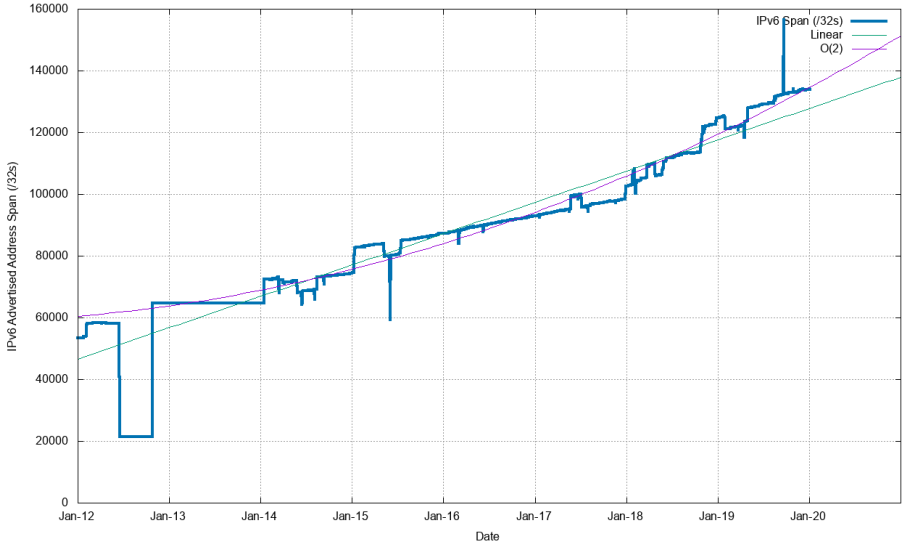

Advertised Address span is growing at an exponential rate

The "shape" of inter-AS interconnection in IPv6 is rising slightly. Local connections appear to be replacing overlay trunk transits

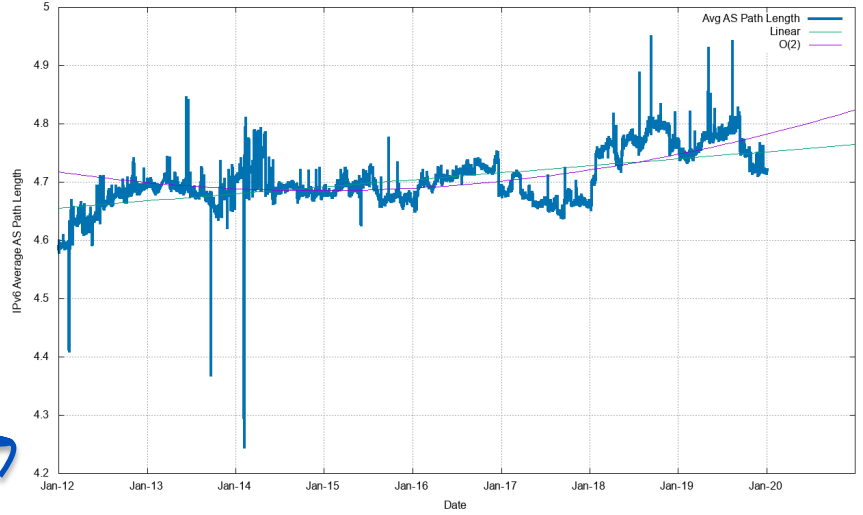

AS Adjacencies (AS131072)

14,997 out of 18,720 ASNs have 1 or 2 AS Adjacencies (80%)
654 ASNs have 10 or more adjacencies
2 ASNs have >1,000 adjacencies

4,728	AS6939	HURRICANE - Hurricane Electric, Inc., US
1,011	AS3356	LEVEL3 - Level 3 Communications, Inc., US
955	AS174	COGENT-174 - Cogent Communications, US
948	AS1299	Telia Carrier, SE
818	AS2914	NTT America, US

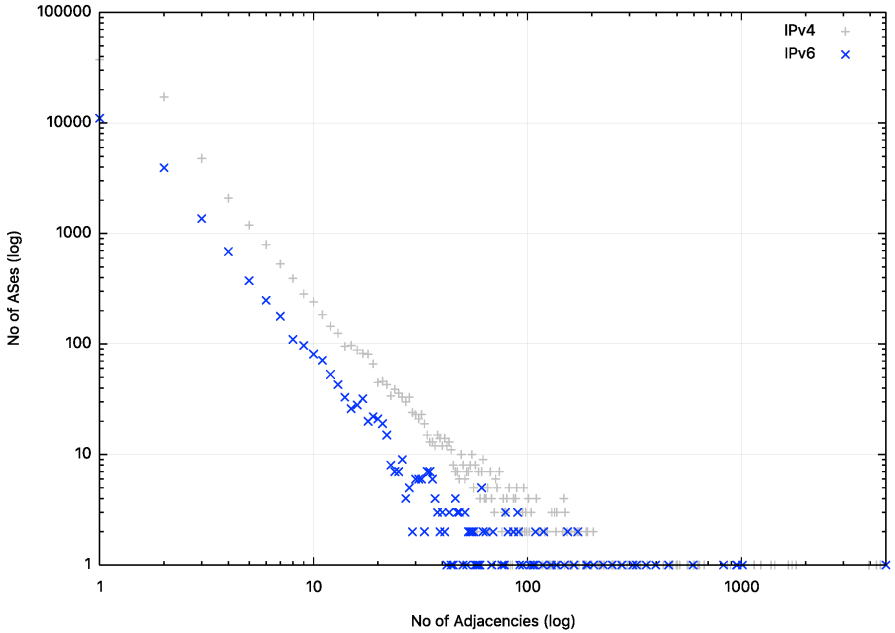

V6 in 2018

- Overall IPv6 Internet growth in terms of BGP is still increasing, and is currently at some **17,000 route entries p.a.**

What to expect

BGP Size Projections

How quickly is the routing space growing?

What are the projections of future BGP FIB size?

V4 - Daily Growth Rates

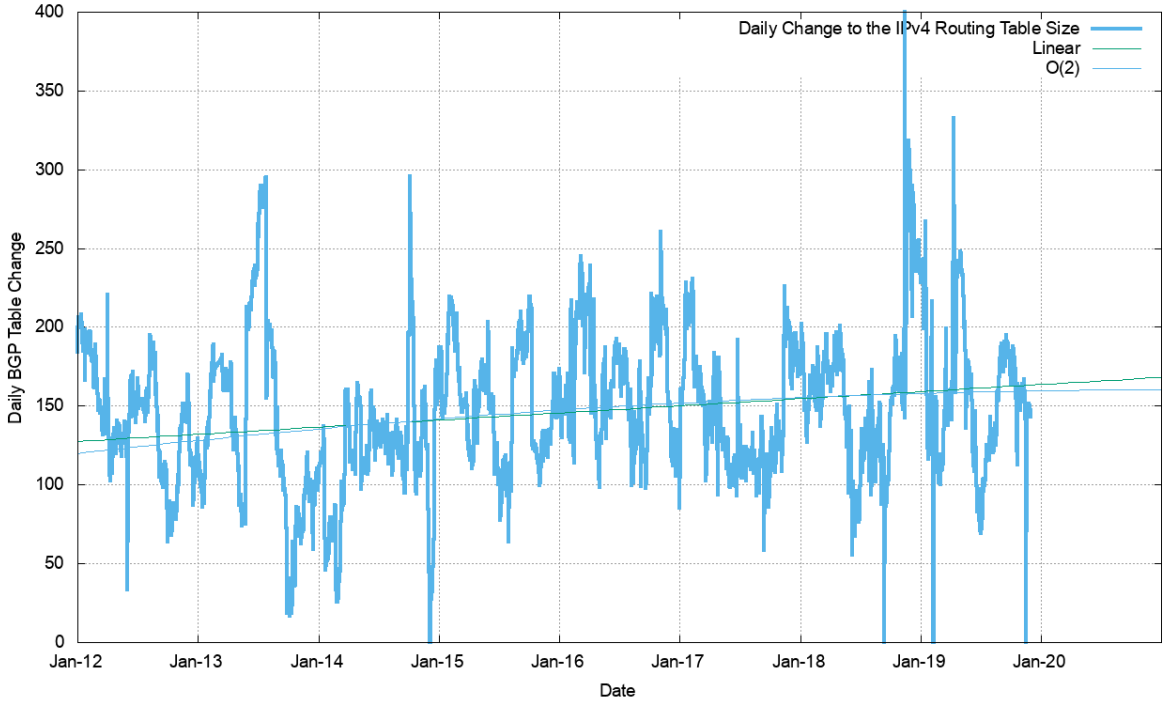

Growth in the V4 network appears to be constant at a long term average of 150 additional routes per day, or some 51,000 additional routes per year

V4 BGP Table Size Predictions

Jan	2017	646,000
	2018	699,000
	2019	760,000
	2020	814,000
	2021	862,000
	2022	916,000
	2023	970,000
	2024	1,024,000
	2025	1,079,000

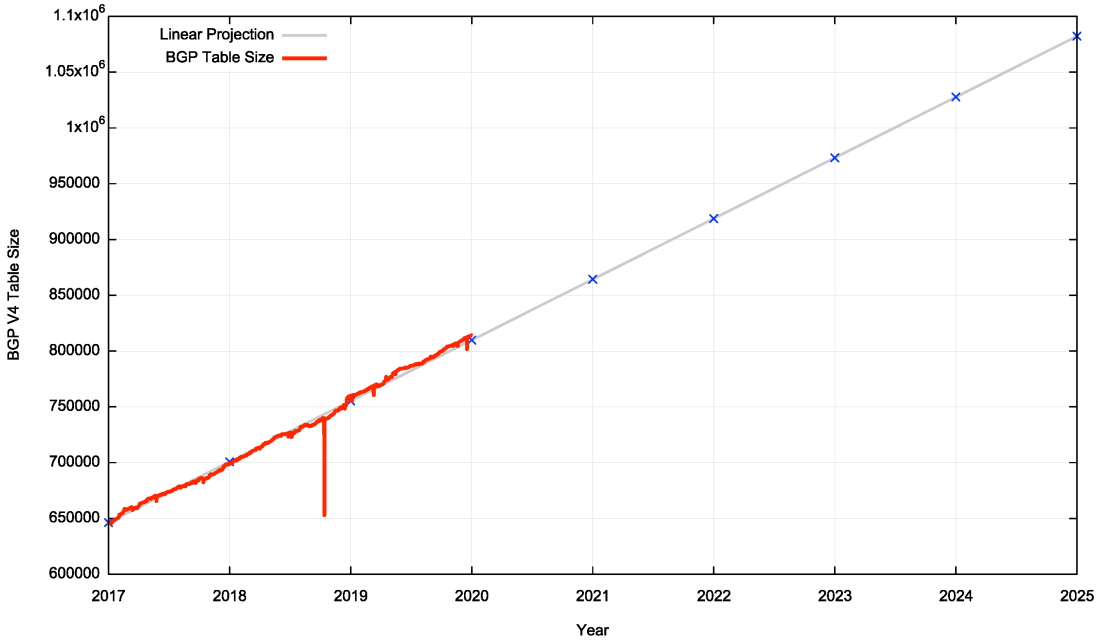

V6 - Daily Growth Rates

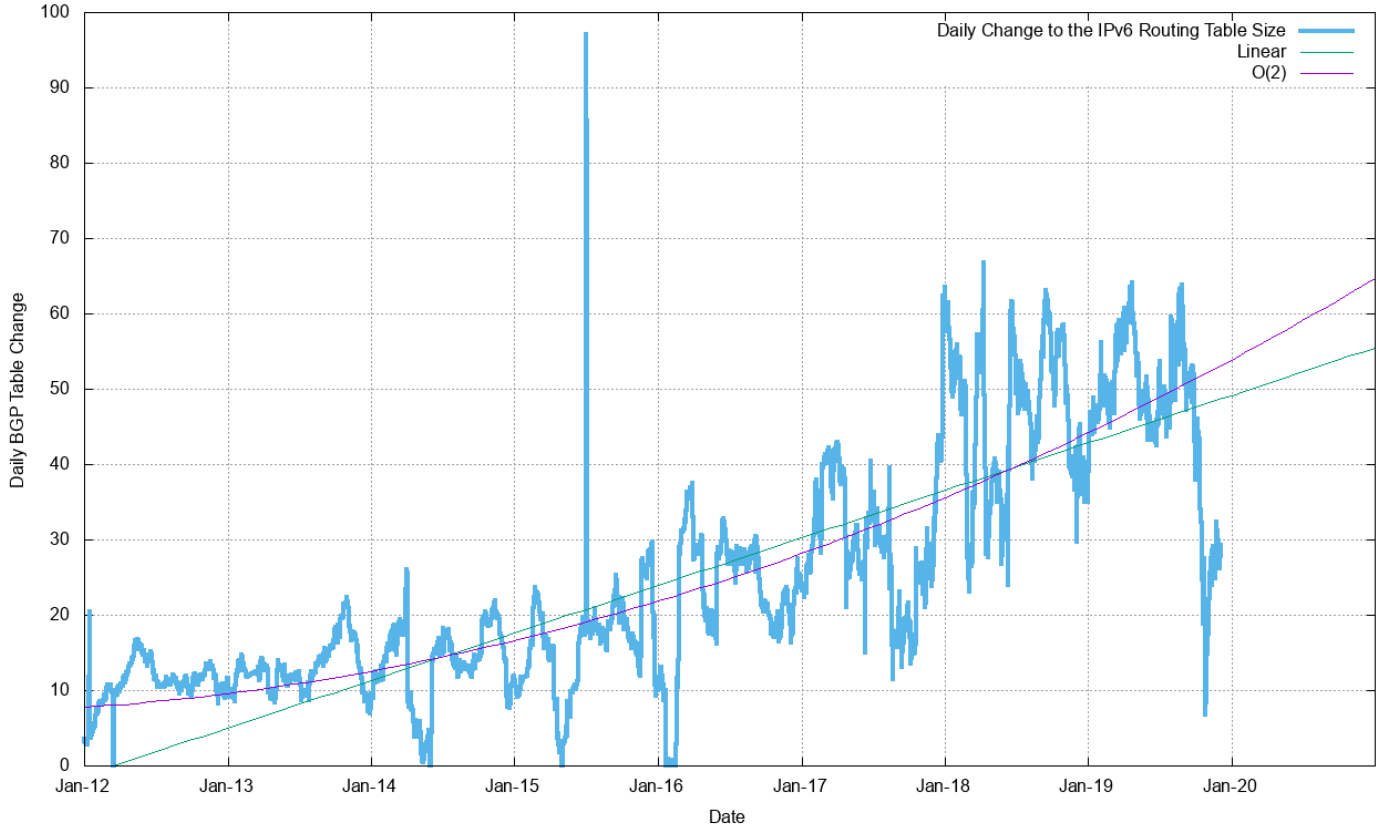

V6 BGP Table Size Predictions

	Linear	Exponential
Jan *2017*	*35,000*	
2018	*45,000*	
2019	*62,000*	
2020	*75,000*	
2021	*96,000*	*106,000*
2022	*112,000*	*140,000*
2023	*128,000*	*184,000*
2024	*144,000*	*242,000*
2025	*160,000*	*318,000*

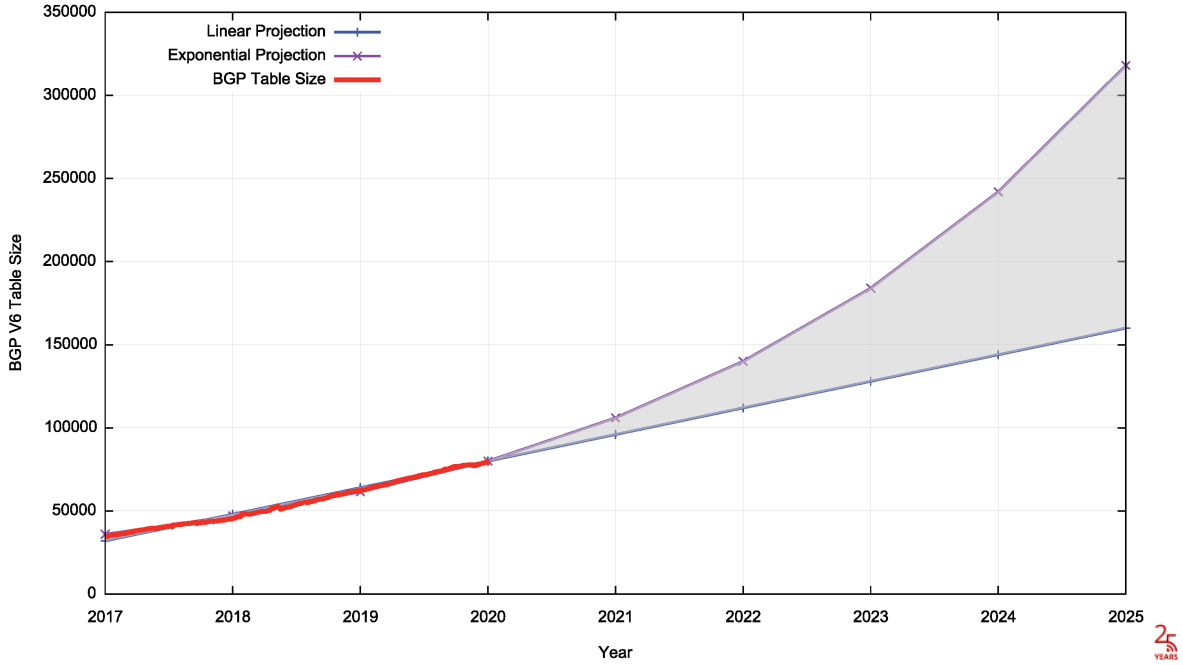

BGP Table Growth

The absolute size of the IPv6 routing table is growing much faster than the IPv4 table

They will require the same memory size in around 5 years time, given that each IPv6 entry is 4 times the memory size of an IPv4 entry

As long as we are prepared to live within the technical constraints of the current routing paradigm, the Internet's use of BGP will continue to be viable for some time yet

BGP Updates

- What about the level of updates in BGP?

IPv4 BGP Updates

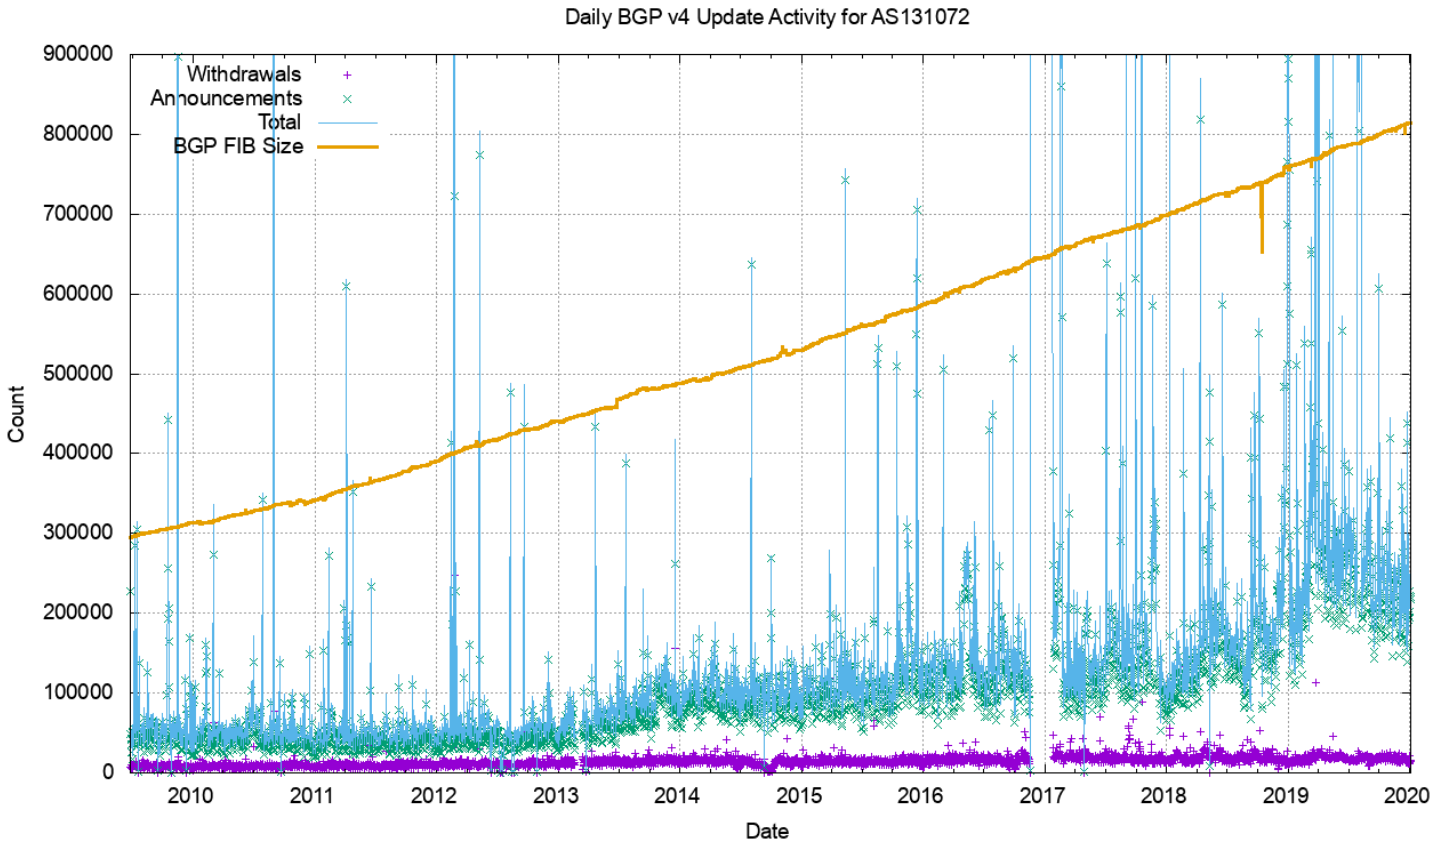

IPv4 BGP Convergence Performance

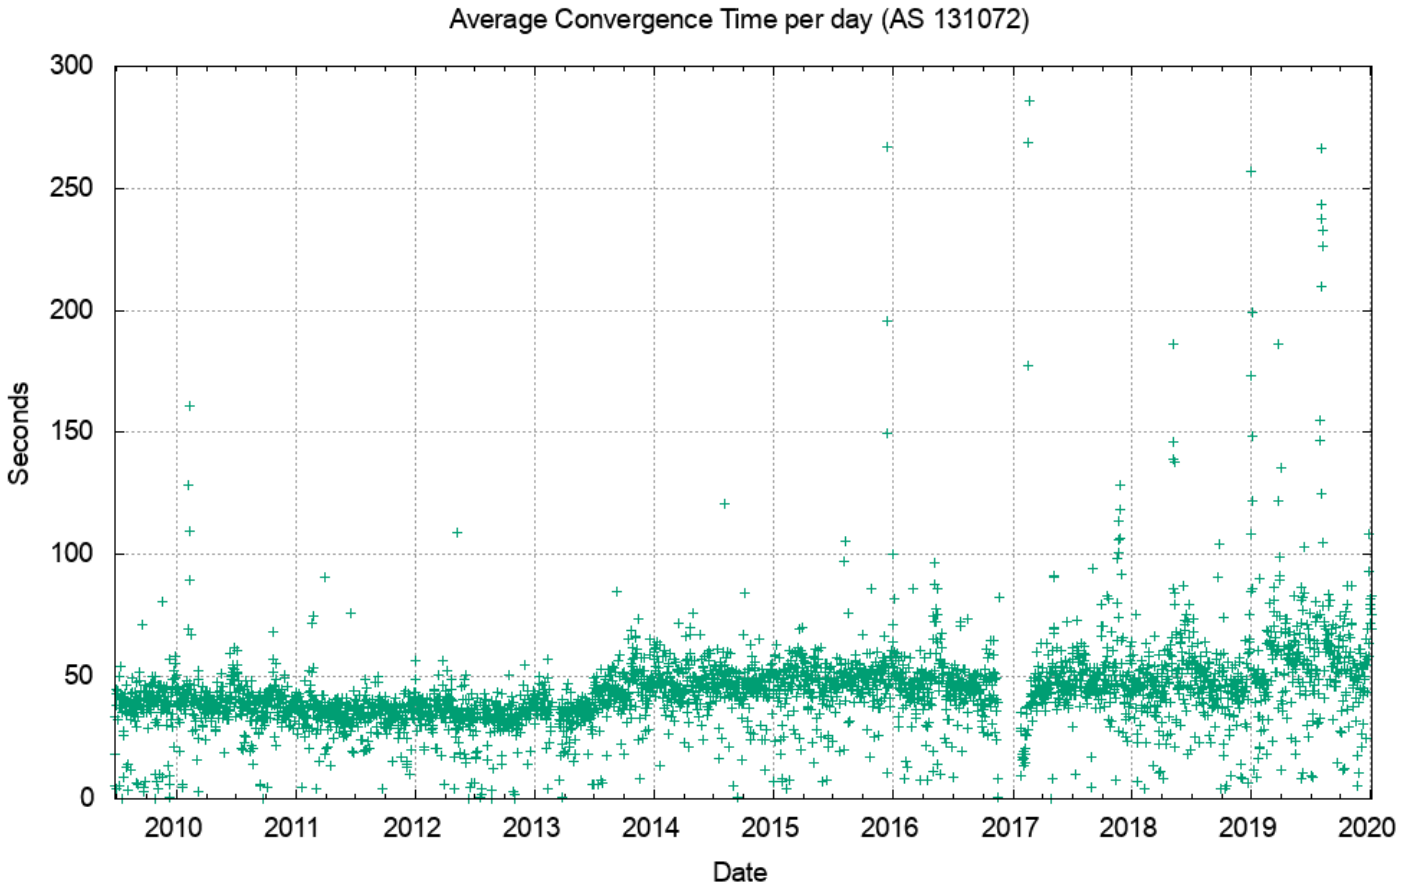

Updates in IPv4 BGP

Still no great level of concern …

- The number of updates per instability event and the time to converge has been relatively constant
- Likely contributors to this outcome are the damping effect of widespread use of the MRAI interval by eBGP speakers, and the compressed topology factor, as seen in the relatively constant AS Path Length

V6 BGP Updates

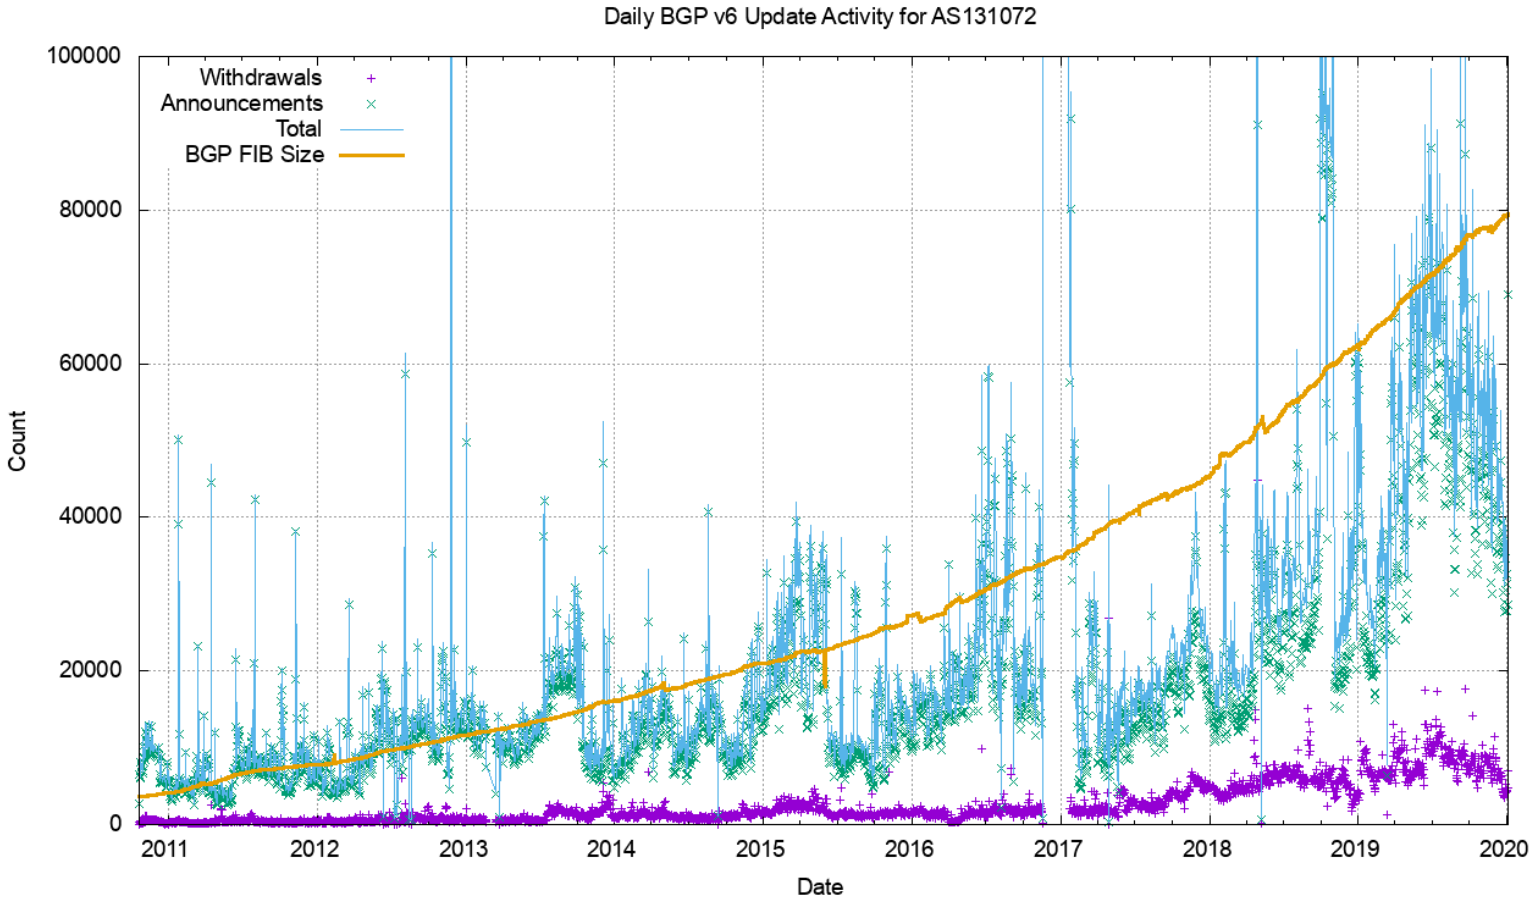

V6 Convergence Performance

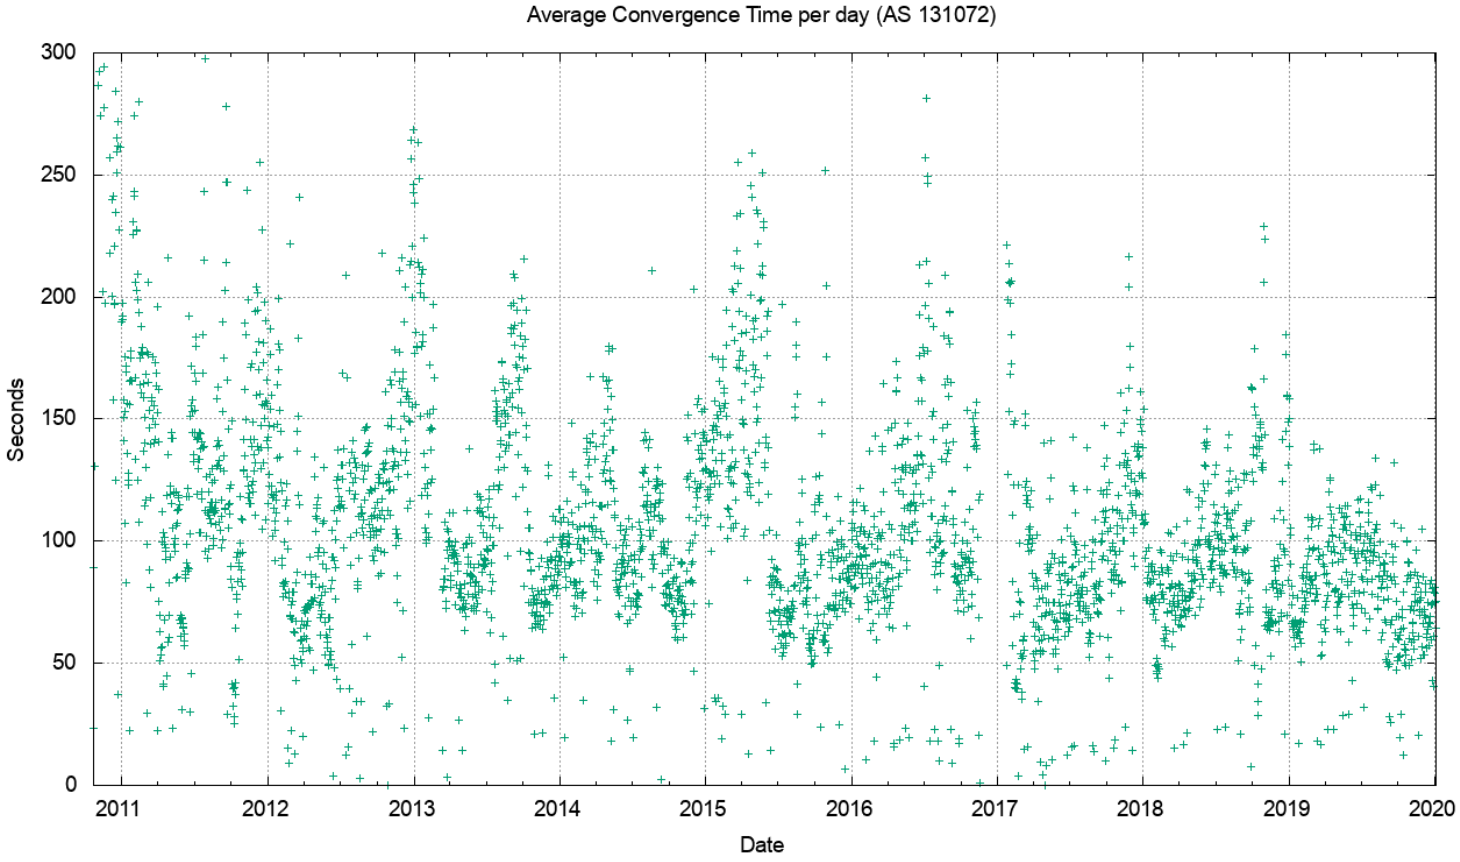

V6 Convergence Performance

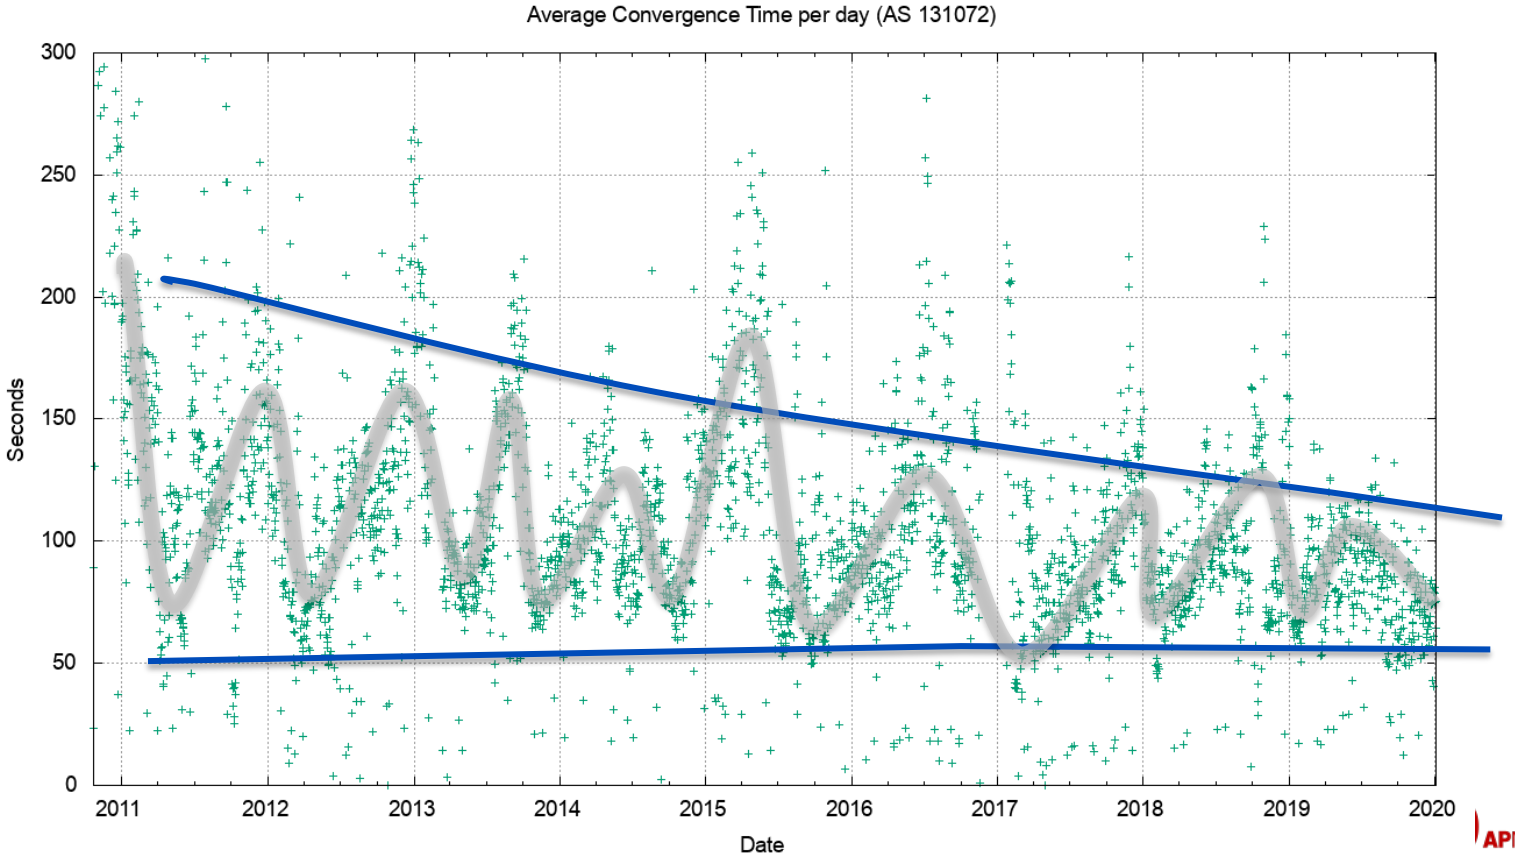

Routing Futures

- There is little in the way of scaling pressure from BGP as a routing protocol – the relatively compressed topology and stability of the infrastructure links tend to ensure that BGP remains effective in routing the internet

- The issues of FIB size, line speeds and equipment cost of line cards represent a more significant issue for hardware suppliers – we can expect cheaper line cards to to use far smaller LRU cache local FIBs in the high-speed switches and push less-used routes to a slower / cheaper lookup path. This approach may also become common in very high-capacity line cards

Some Practical Suggestions

- Understand your hardware's high speed FIB capacity in the default-free parts of your network
- Review your IPv4 / IPv6 portioning - a dual-stack eBGP router will need 920,000 IPv4 slots and 140,000 IPv6 slots for a full eBGP routing table in line cards over the coming 24 months if they are using a full FIB load
- Judicious use of default routes in your internal network may allow you drop this requirement significantly
- Using a hot cache for line card FIB cache would reduce the memory requirement significantly without visible performance cost

That's it!


25 YEARS
2020
APRICOT
APNIC 49
MELBOURNE
AUSTRALIA
12 – 21 February 2020
#apricot2020