Another year of BGP!

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

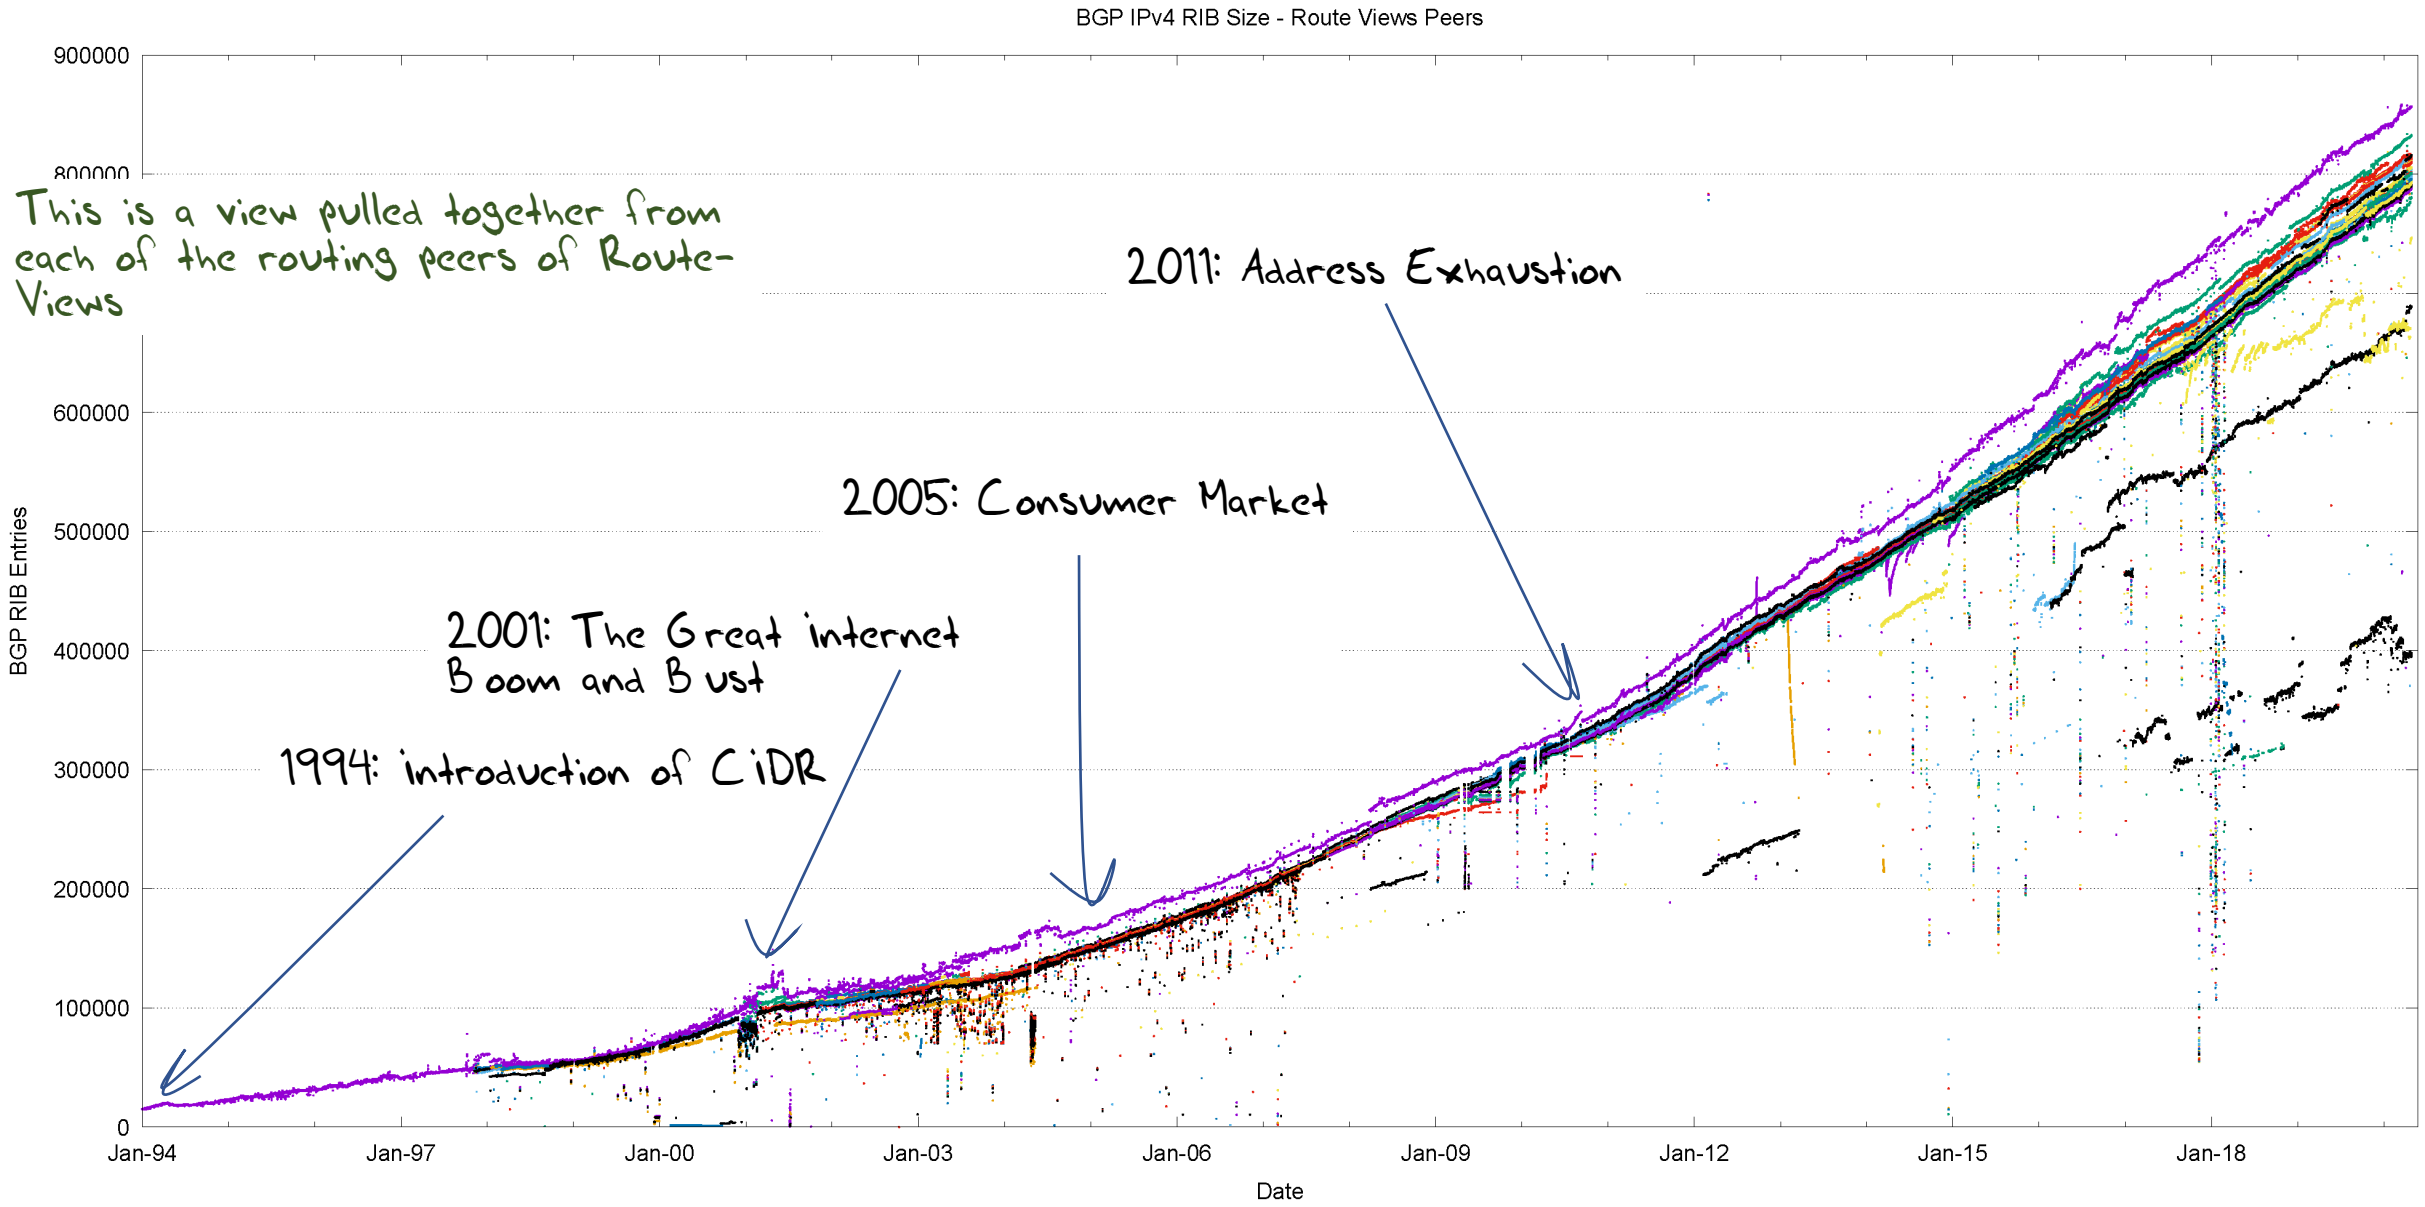

2017-2020 in detail

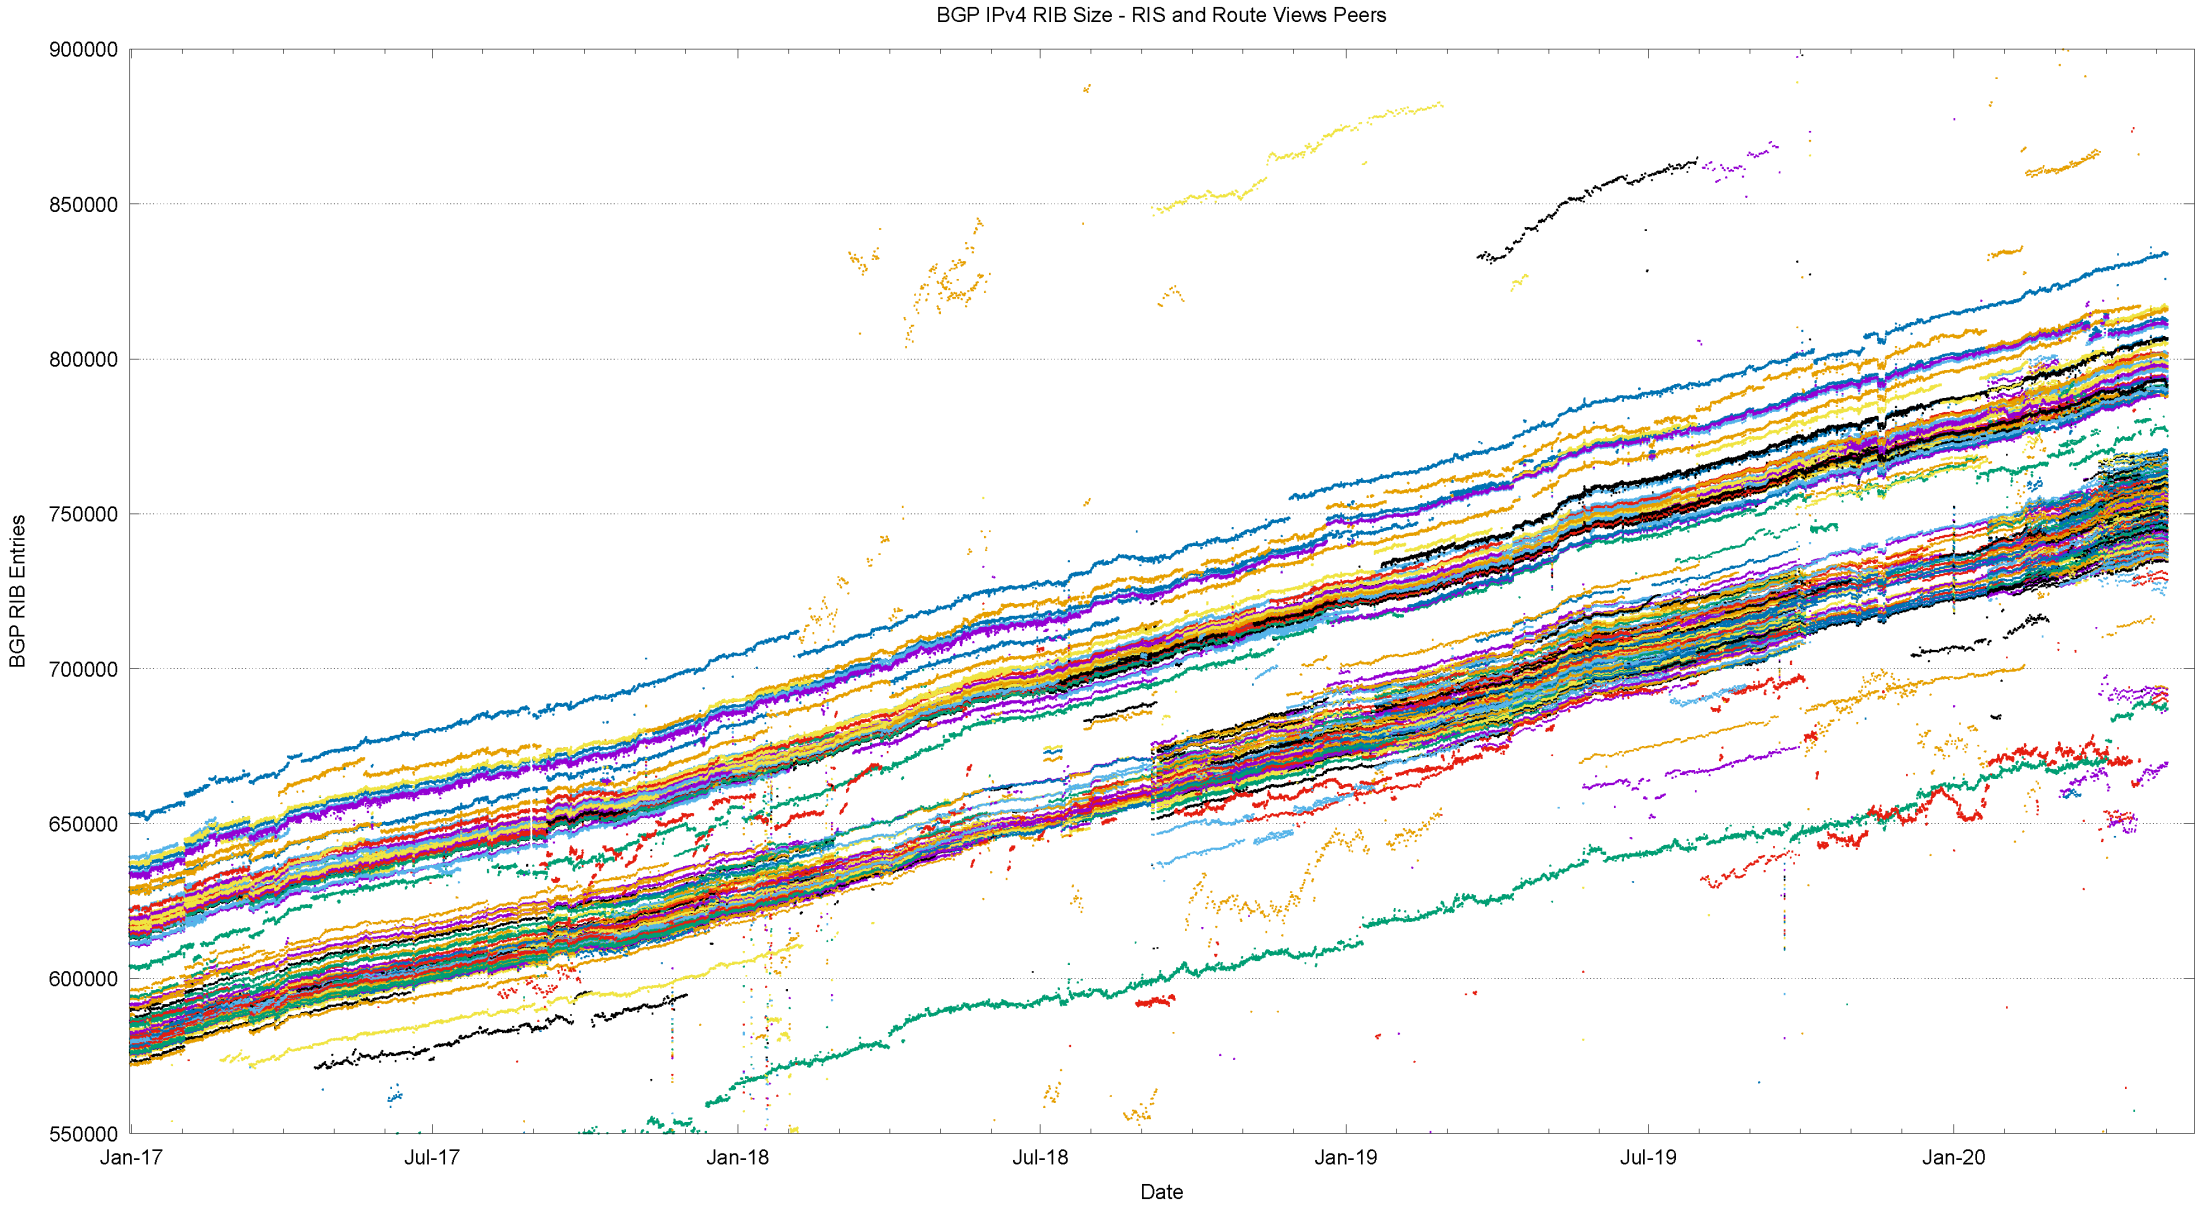

2017-2020 in detail

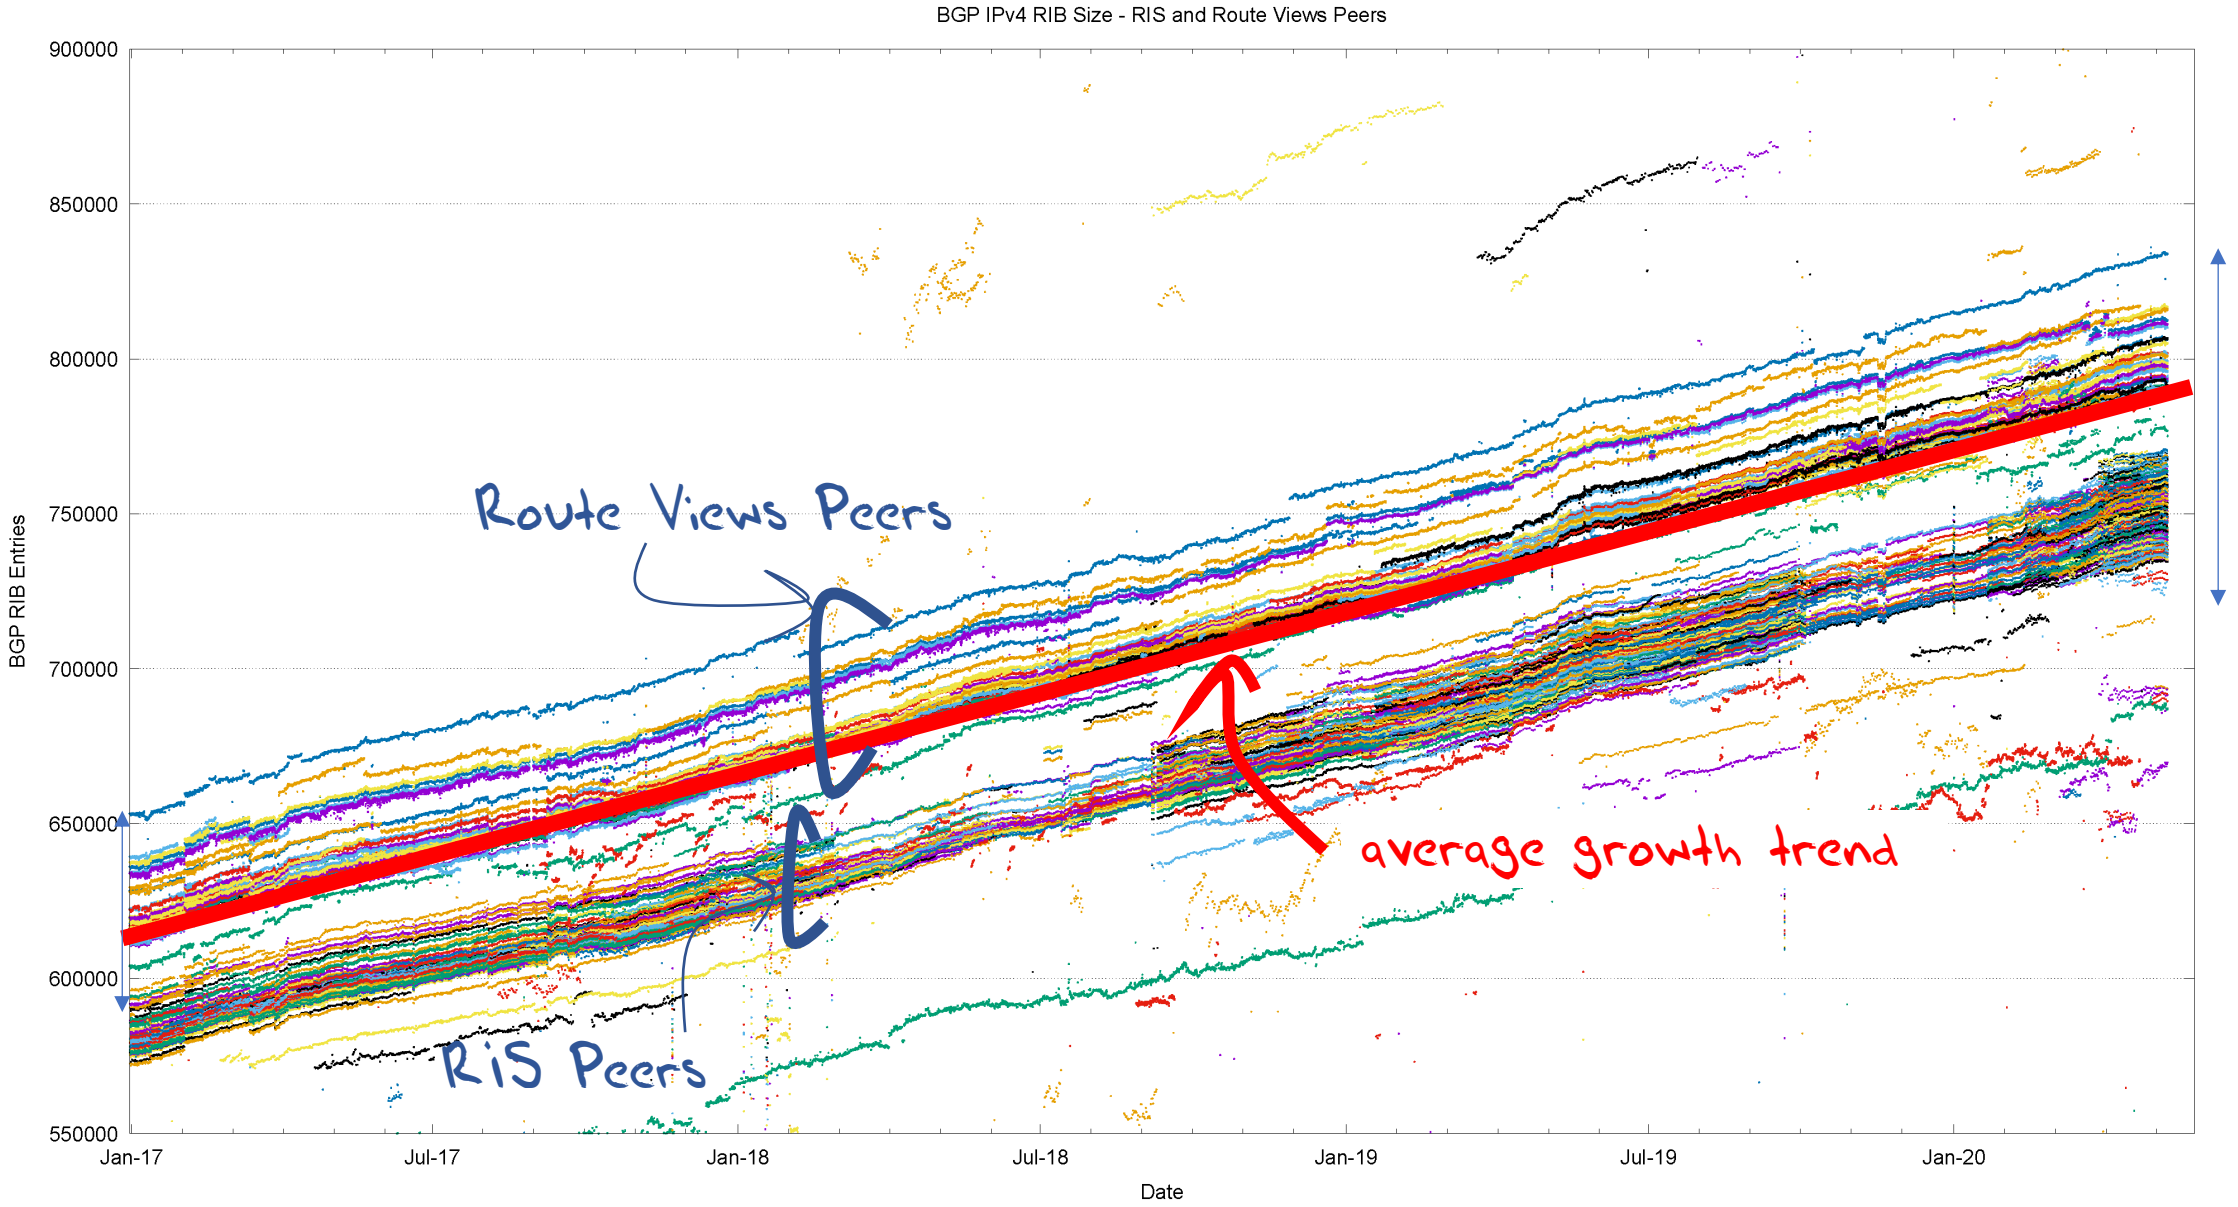

Routing Indicators for IPv4

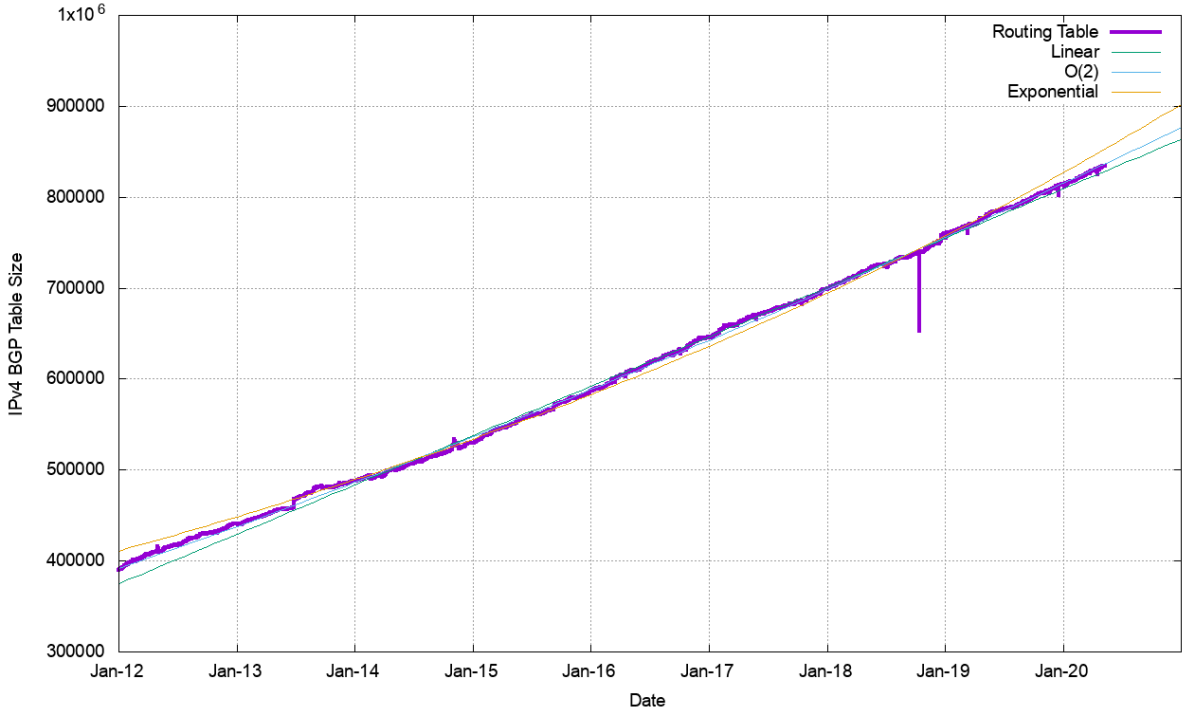

Routing prefixes – growing by some 55,000 prefixes per year

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

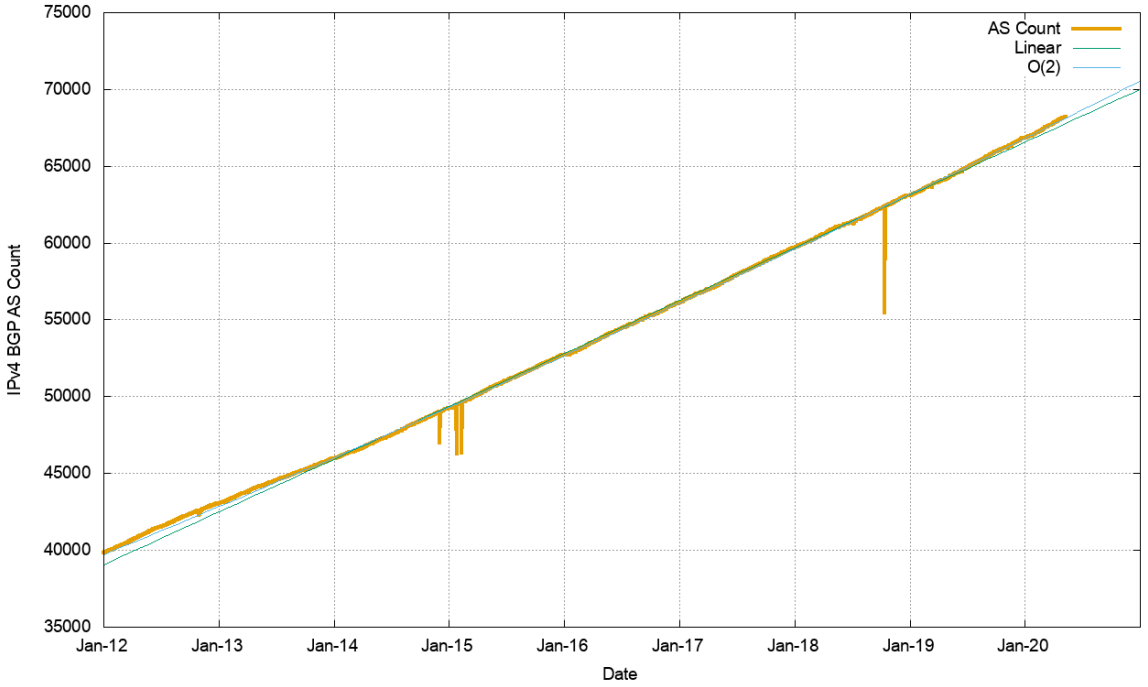

Routing Indicators for IPv4

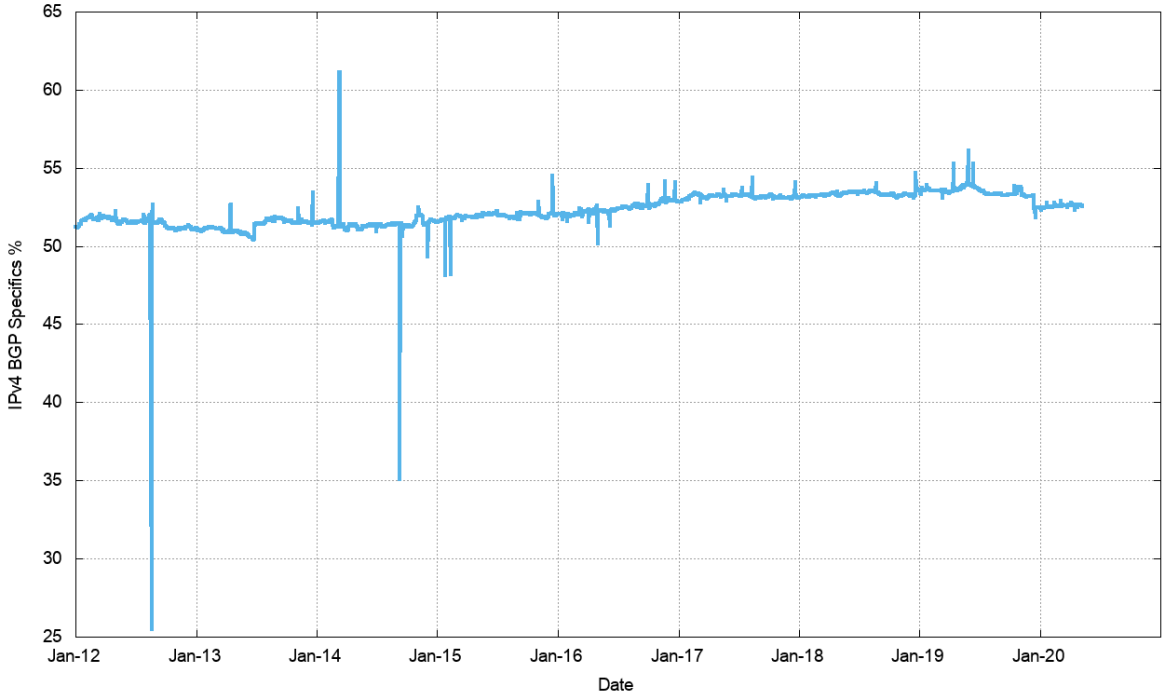

More Specifics are still taking up slightly more than one half of the routing table

But the average size of a routing advertisement continues to shrink

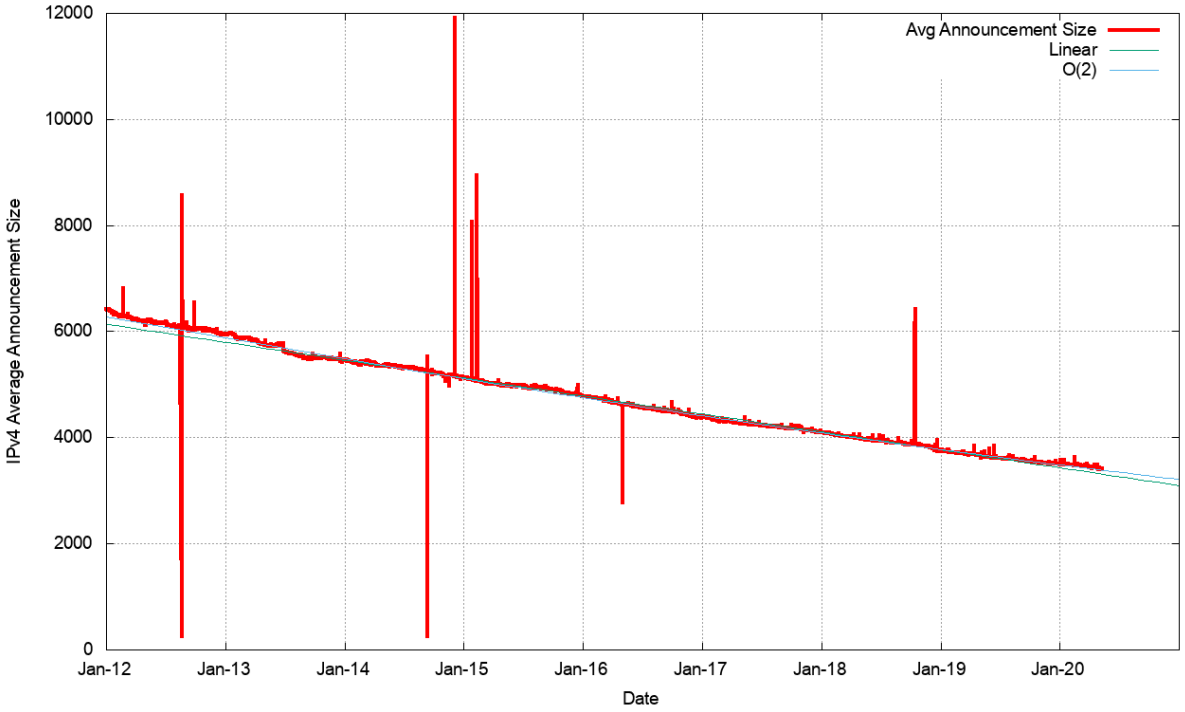

Routing Indicators for IPv4

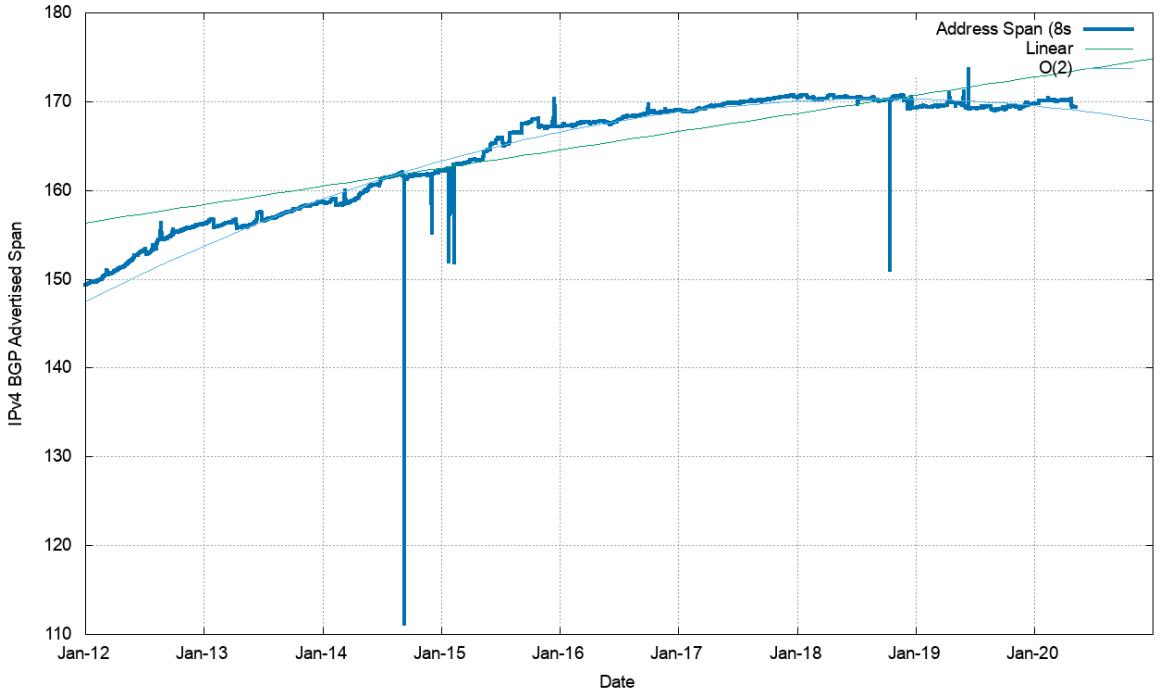

Address Exhaustion is now visible in the extent of advertised address space

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

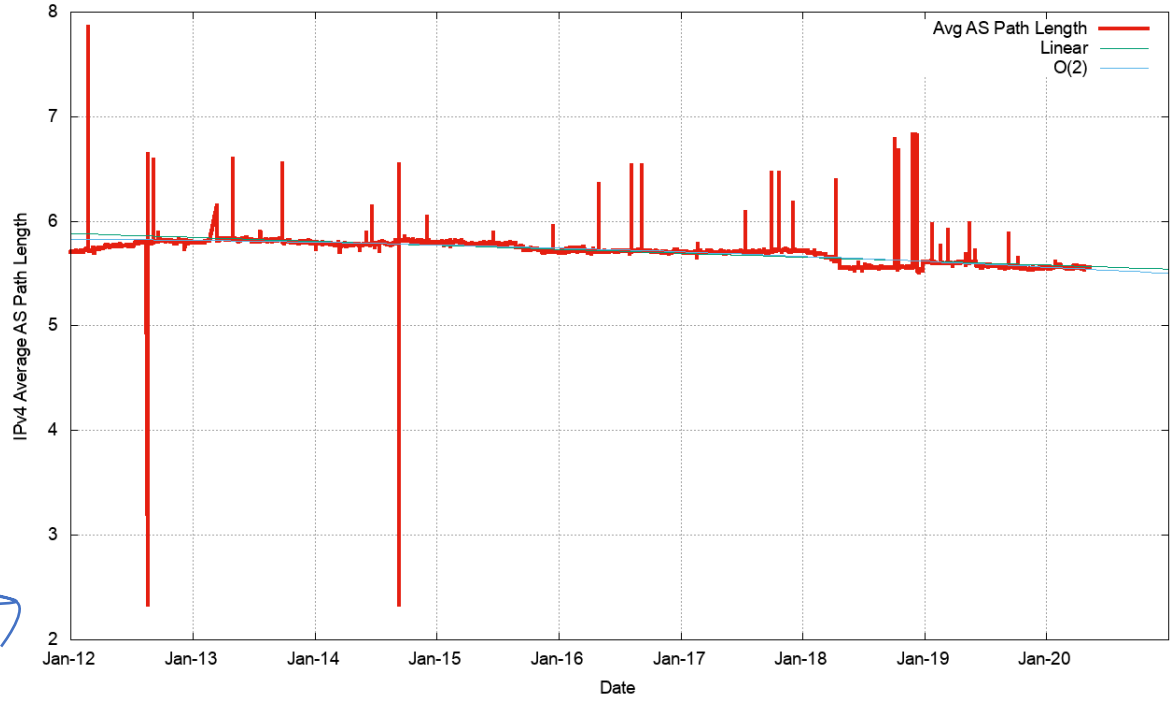

AS Adjacencies (AS131072)

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

4,788	AS6939	HURRICANE, US
4,378	AS3356	LEVEL3, US
3,870	AS174	COGENT-174, US
1,877	AS6461	ZAYO-6461, US
1,735	AS7018	ATT-INTERNET4, US
1,675	AS3257	GTT-BACKBONE GTT, DE
1,413	AS2914	NTT-COMMUNICATIONS-2914, US
1,306	AS1299	TELIANET Telia Carrier, EU
1,304	AS3549	LVLT-3549, US
1,161	AS209	CENTURYLINK-US-LEGACY-QWEST, US

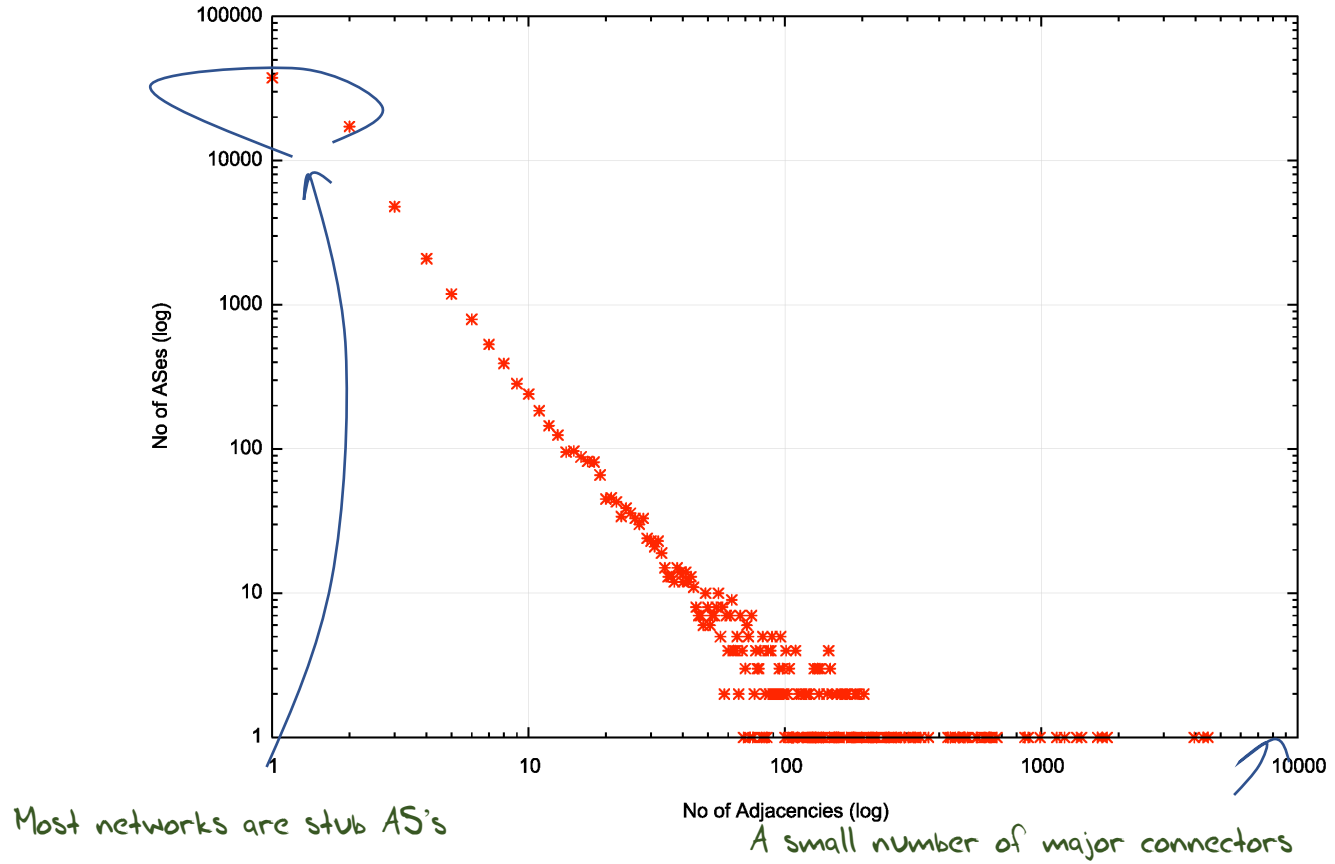

What's happening in V4?

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 has now reached 840,000
- 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?

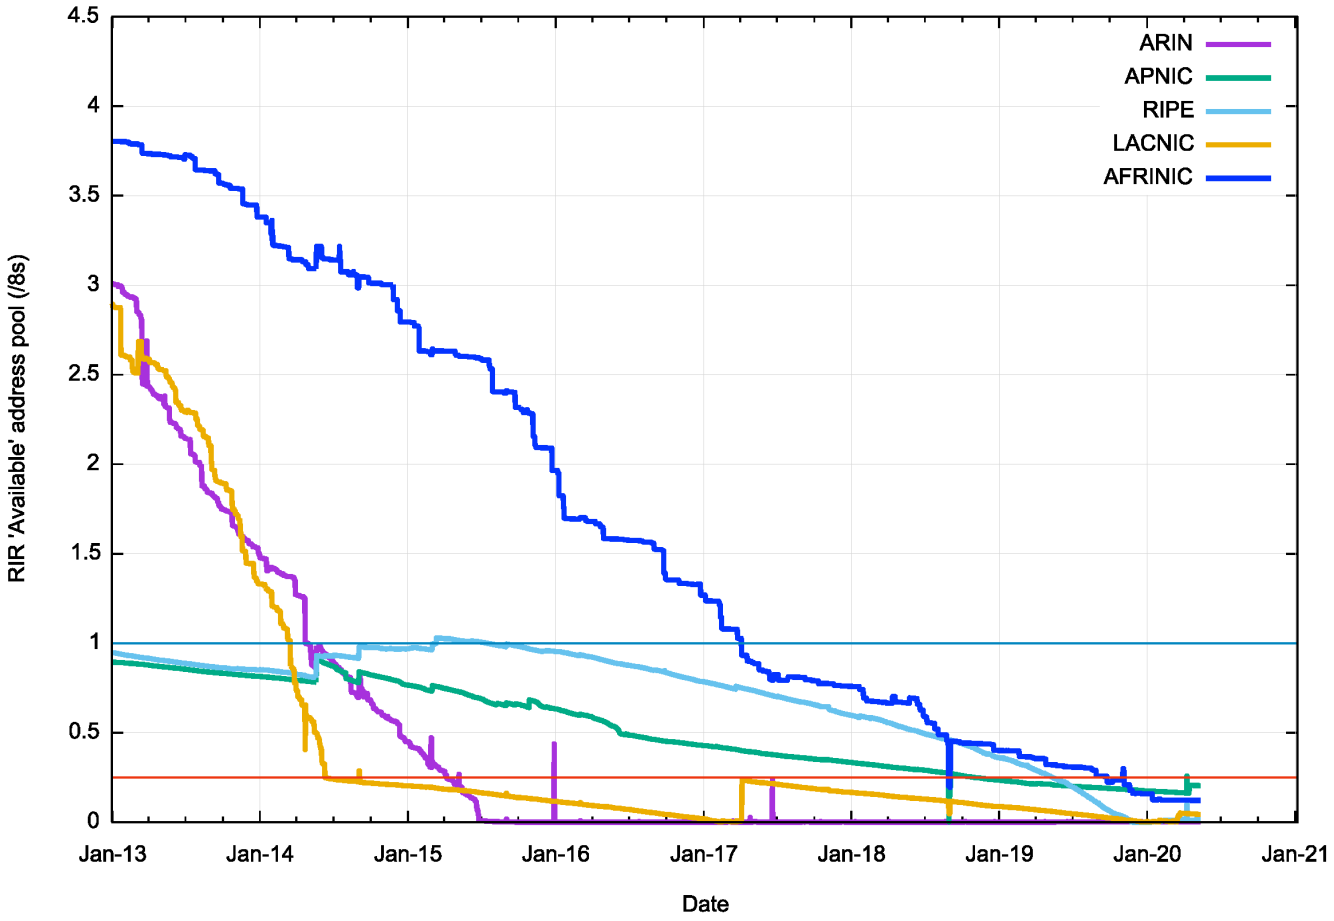

RIR Address Pool runout projections:

ARIN – no free pool left
AFRINIC – December 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

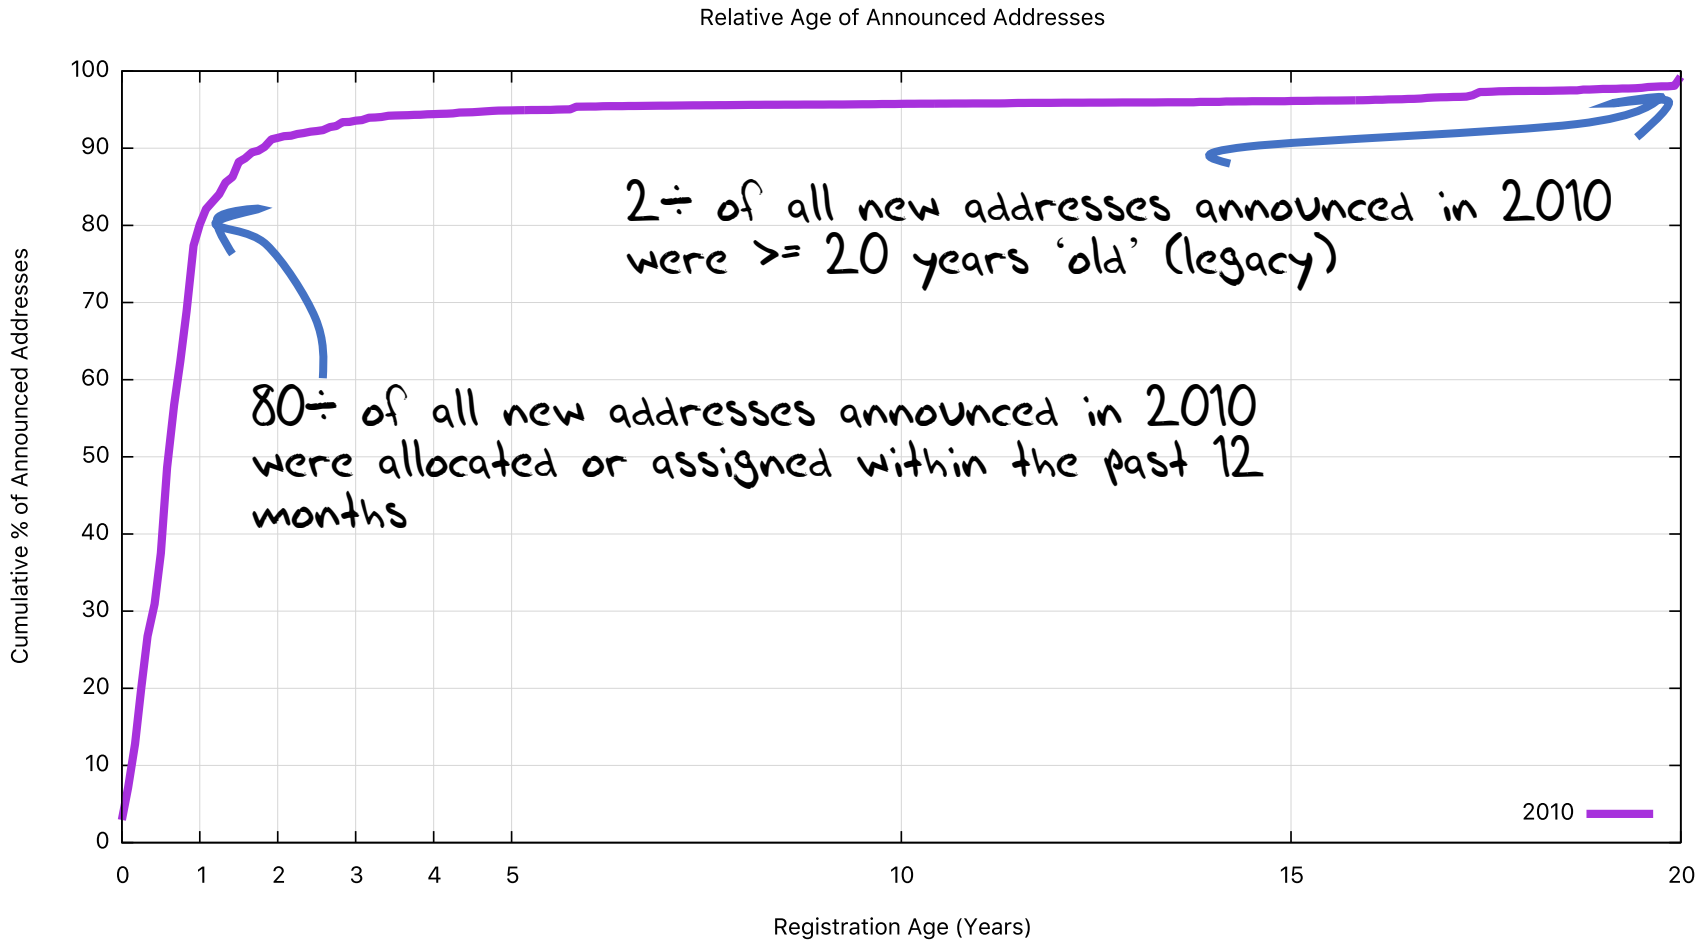

Advertised Address "Age"

2019

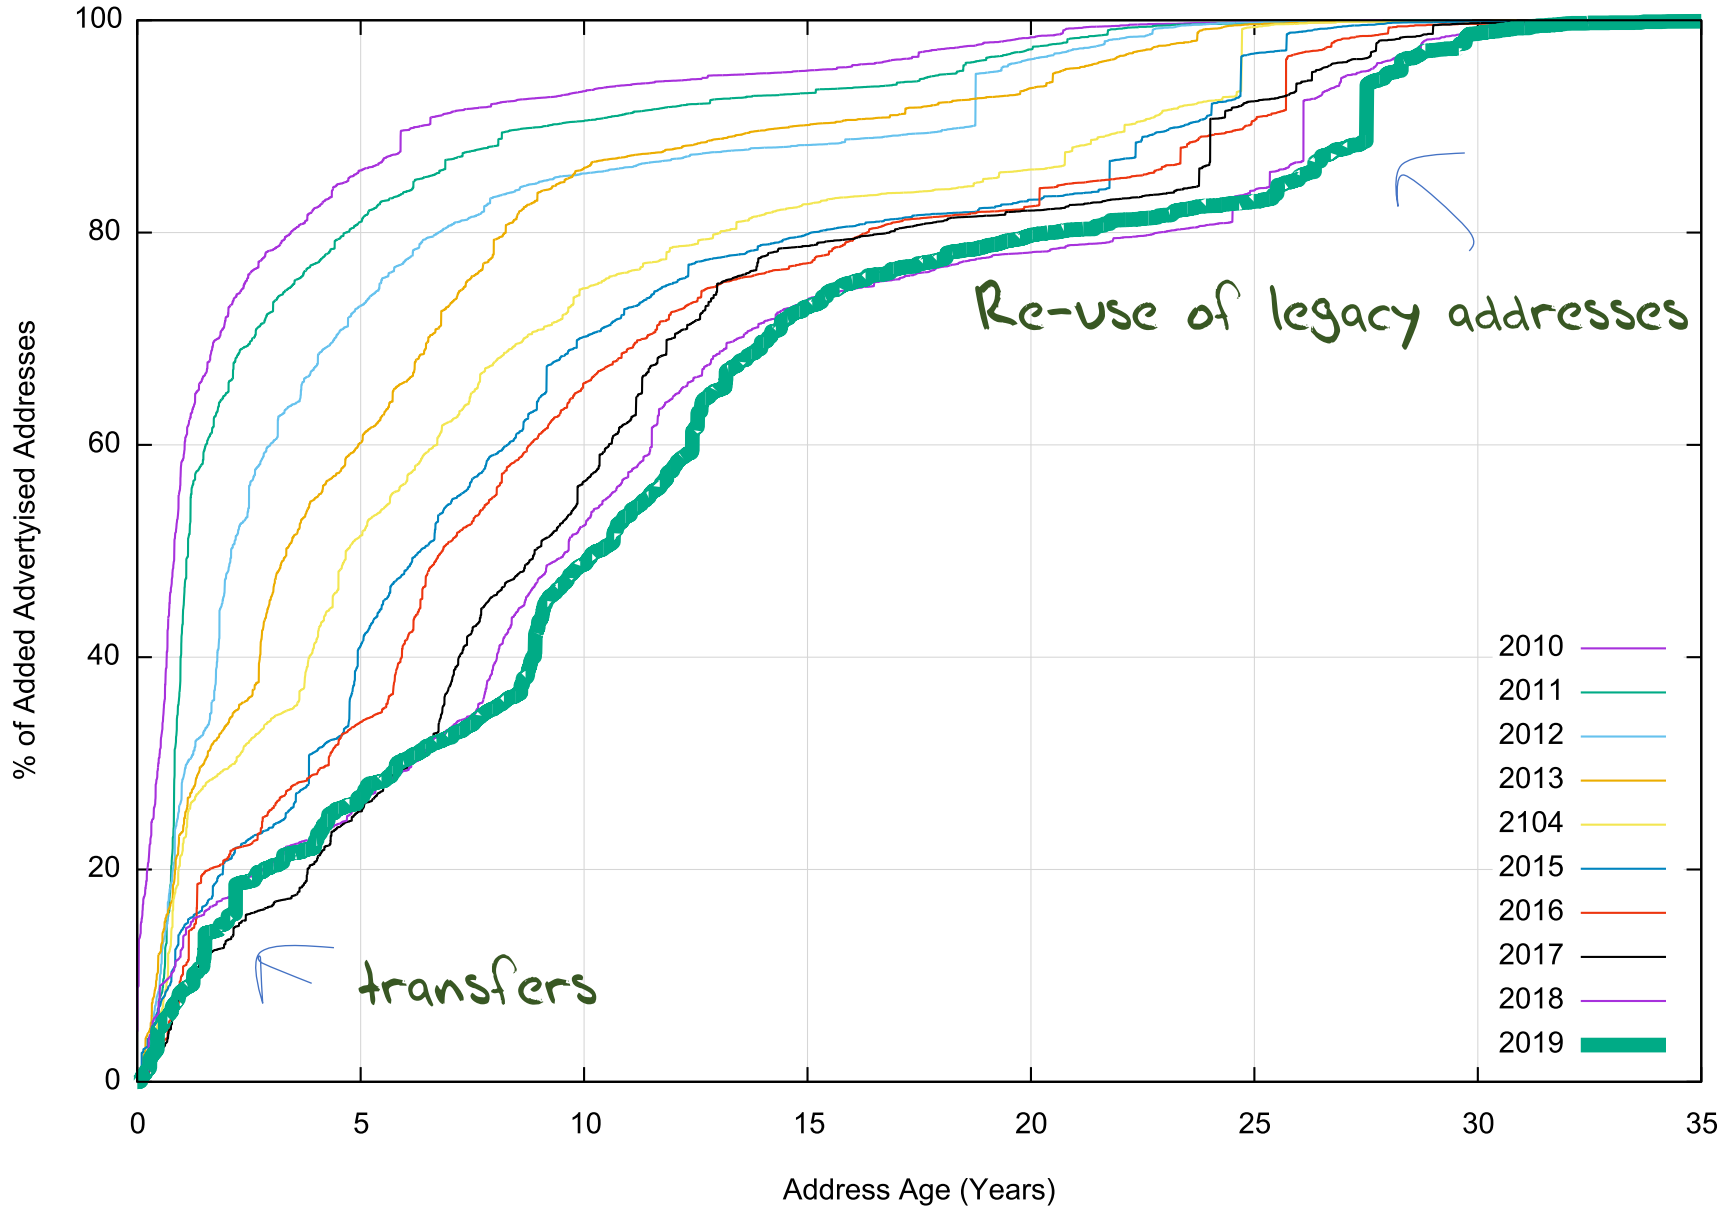

2000 - 2020: IPv4 Advertised vs Unadvertised

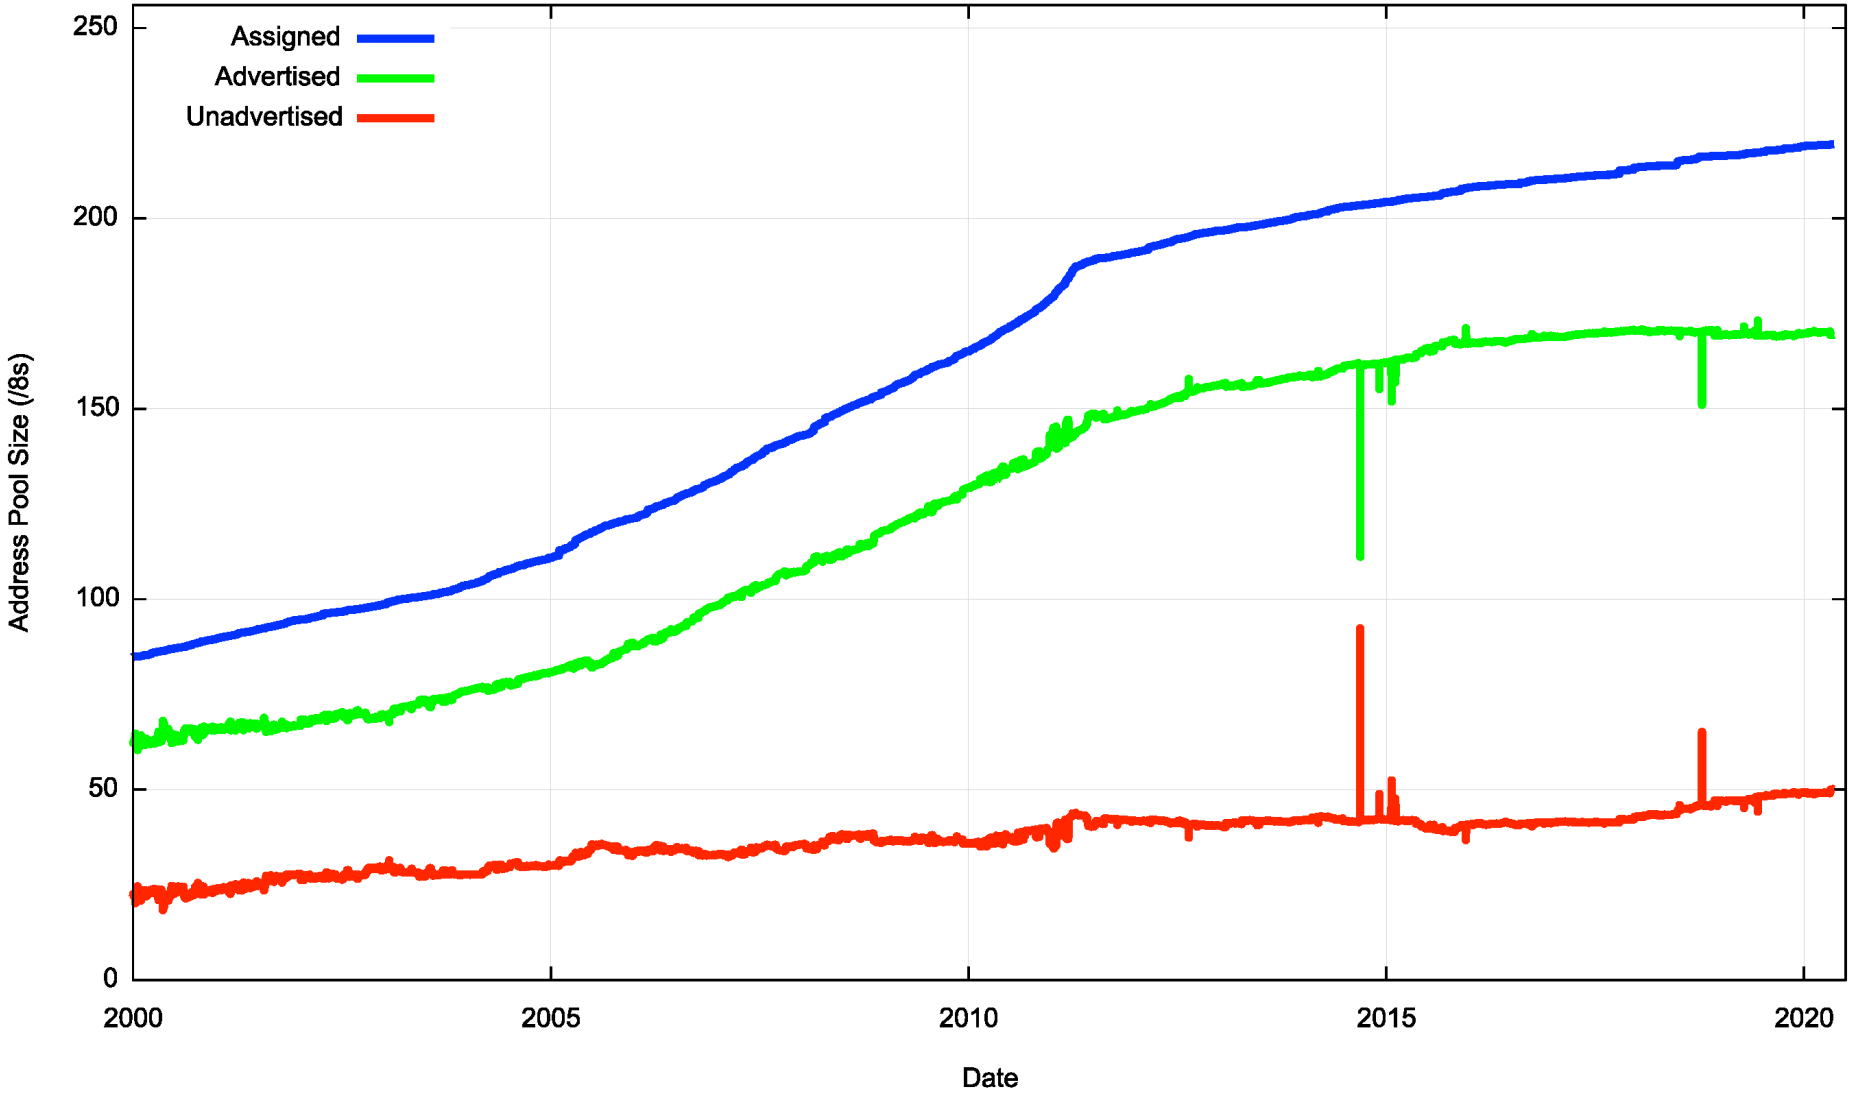

2013 - 2020: Post Free Pool Exhaustion IPv4 Advertised vs Unadvertised

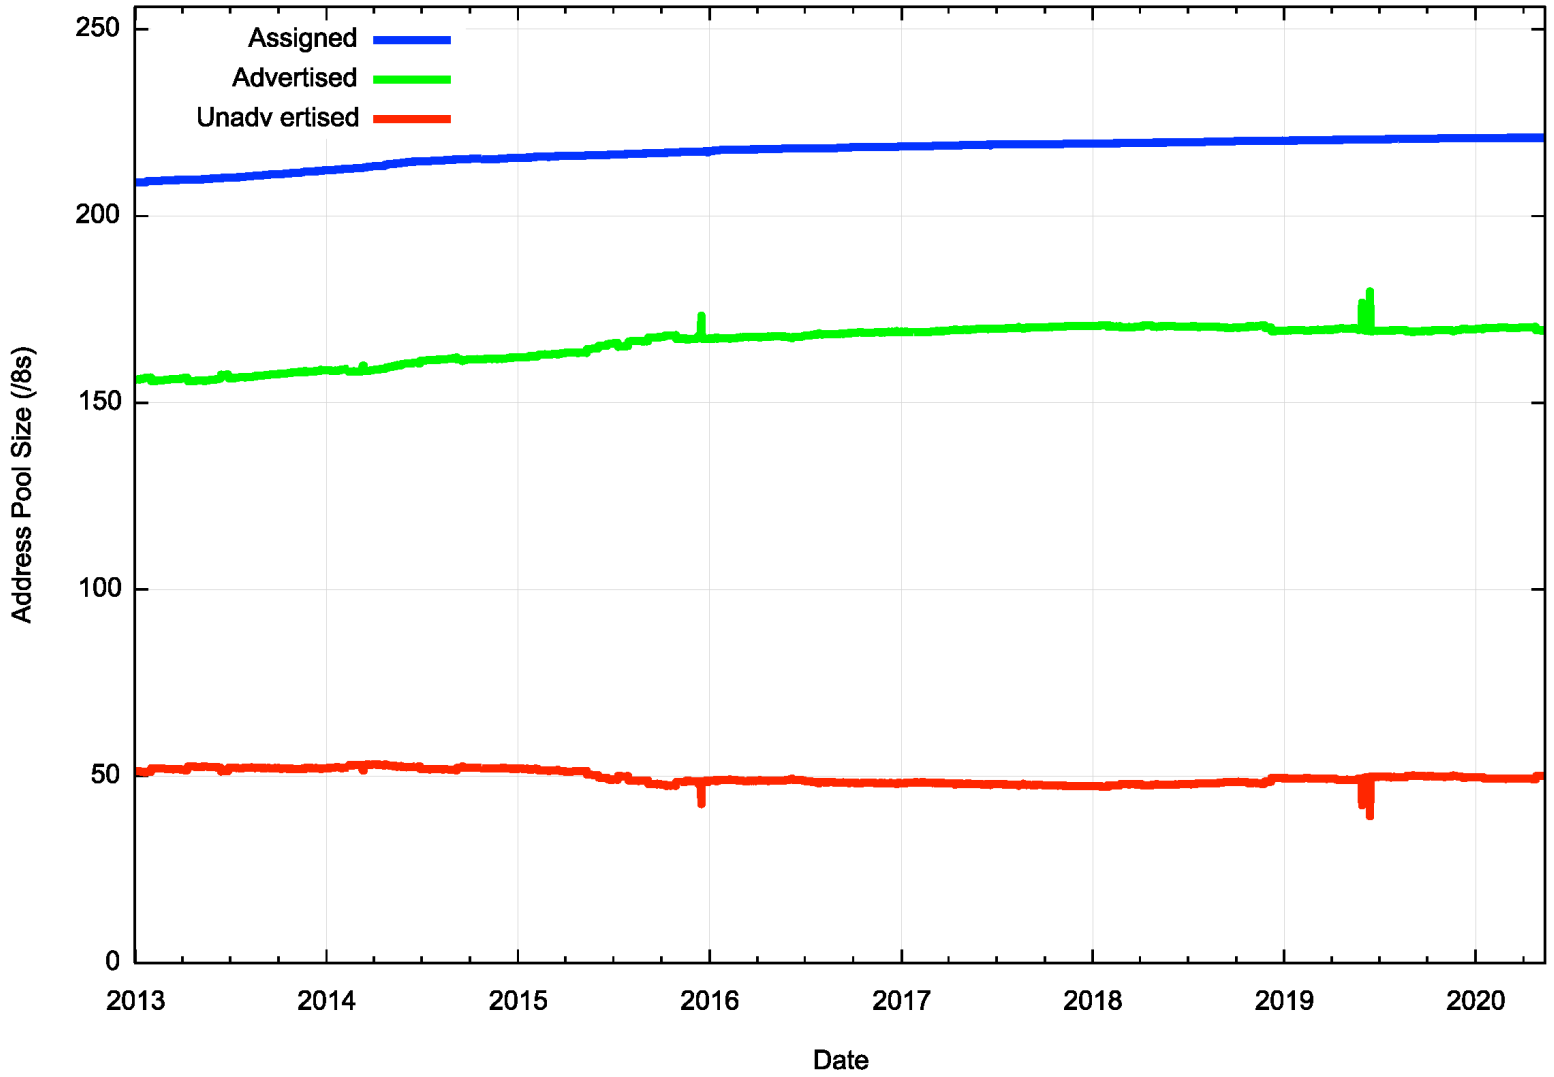

2019 - 2020: Assigned vs Unadvertised

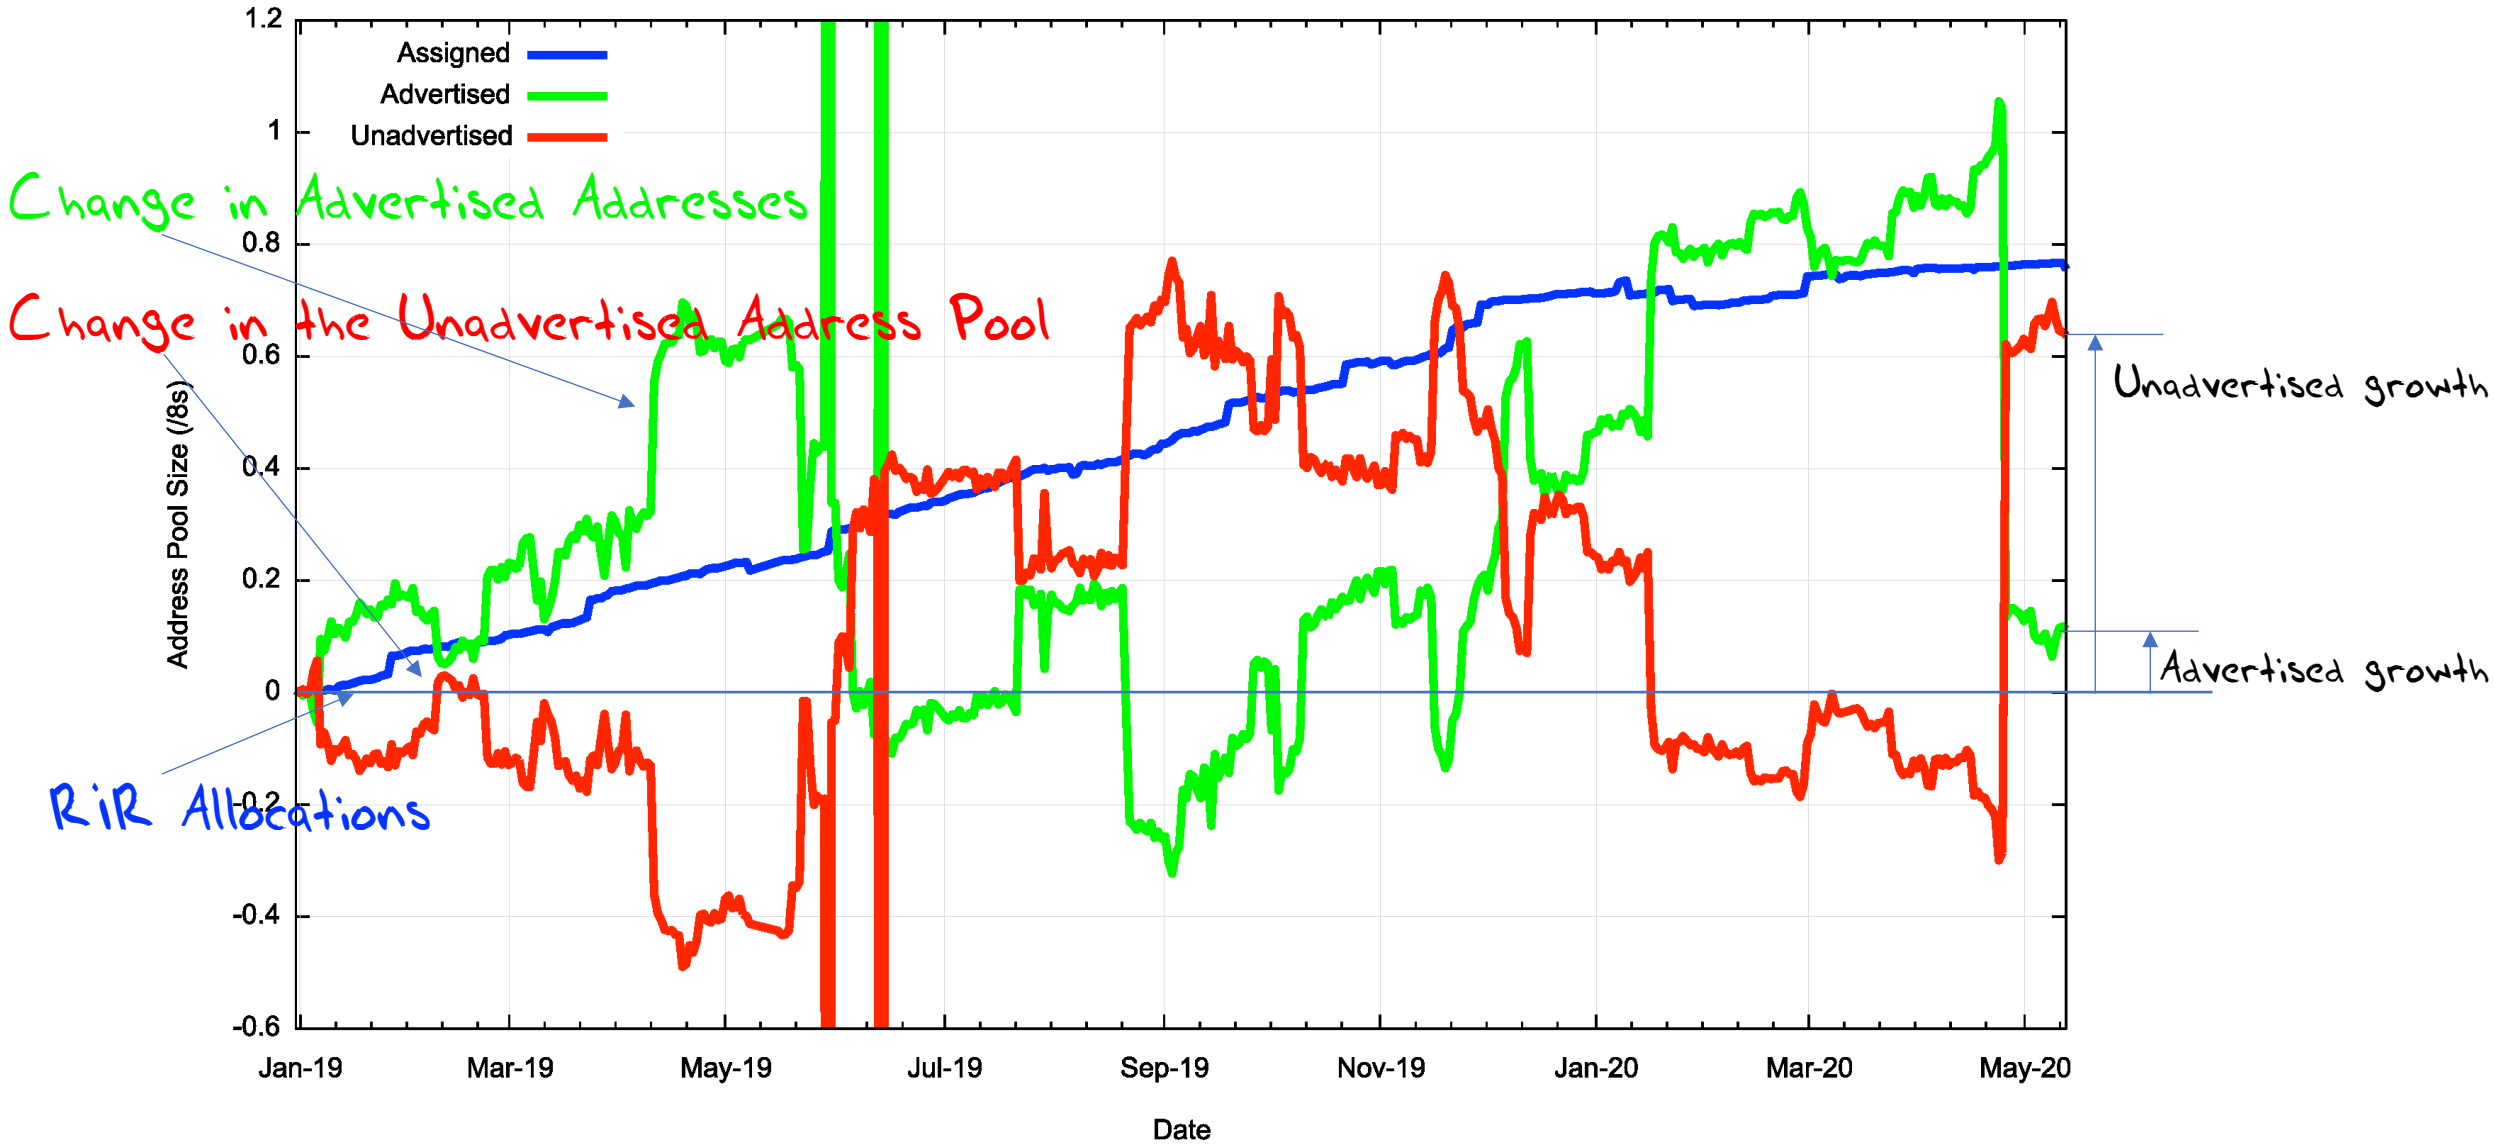

V4 for the last 16 months

- The equivalent of 0.1 /8s were **added** to the routing table
  - It would've been 1.1 /8s but on April 25 57.0.0.0/8 (SITA) was withdrawn from the routing table
- Approximately 0.8 /8s were assigned by RIRs
  - 0.35 /8s assigned by the RIPE NCC (last /8 allocations)
  - 0.24 /8's assigned by Afrinic
  - 0.10 /8s were assigned by LACNIC
  - 0.07 /8s were assigned by APNIC (last /8 allocations)
- And a net of 0.7 /8's were added to the pool of unadvertised addresses

Over this period we saw legacy blocks transferring away from ISPs / end user sites and heading towards cloud SPs.

The Route-Views View of IPv6

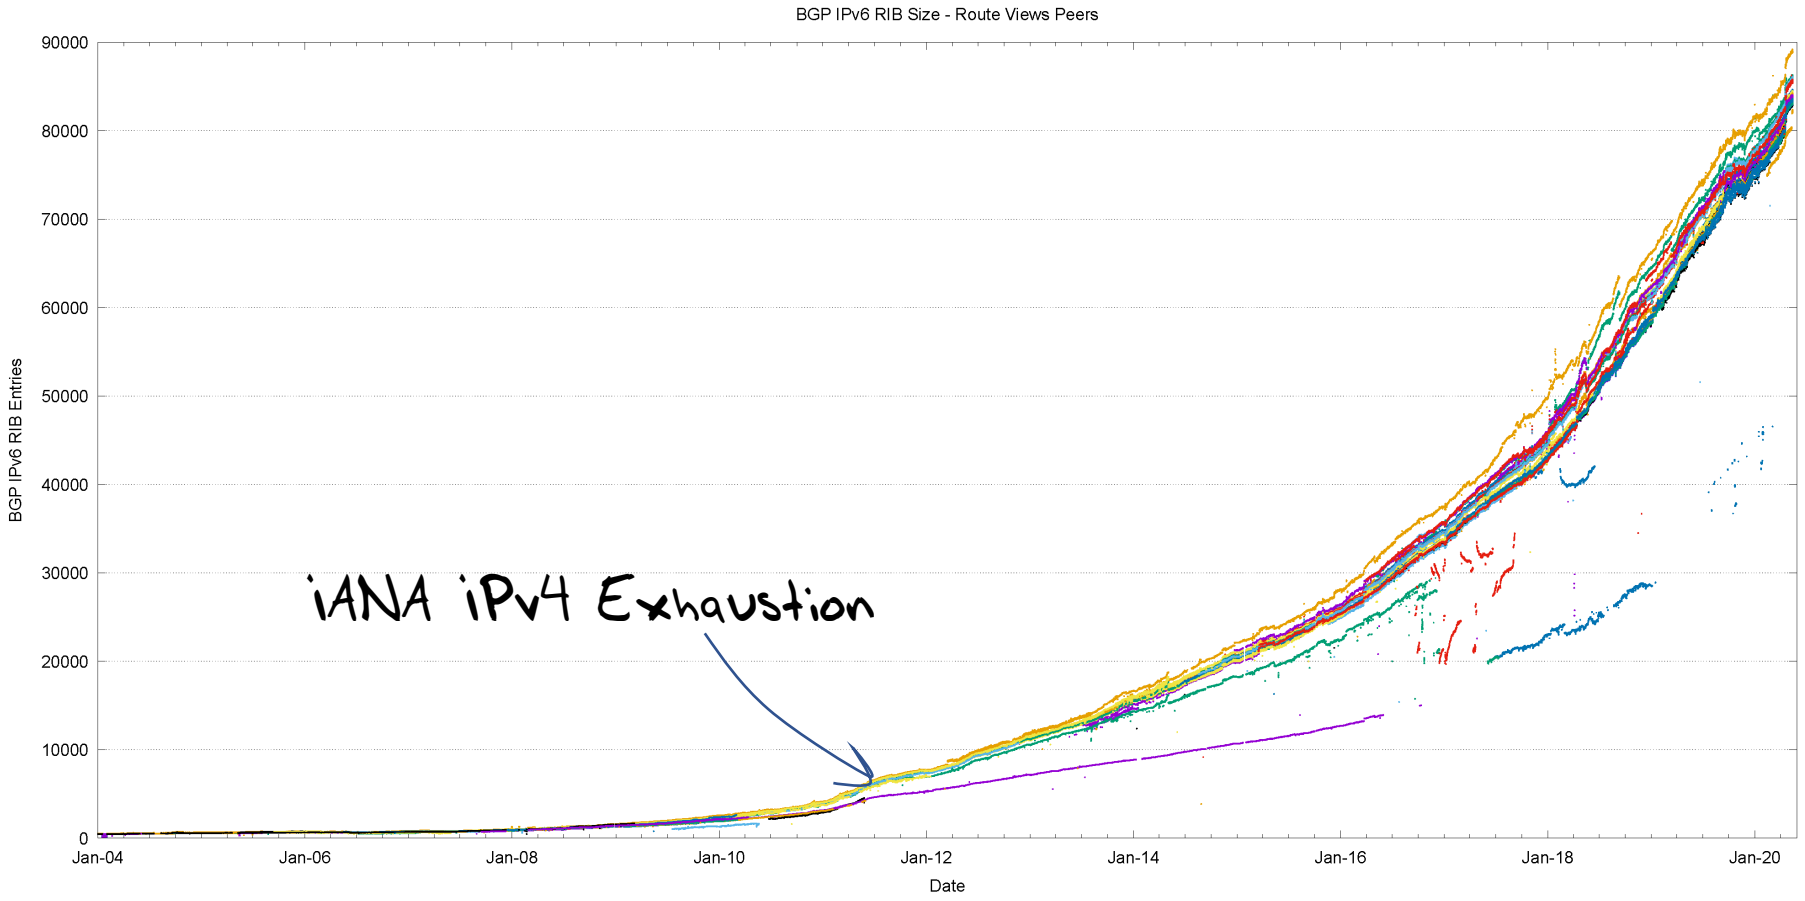

2018-2020 in Detail

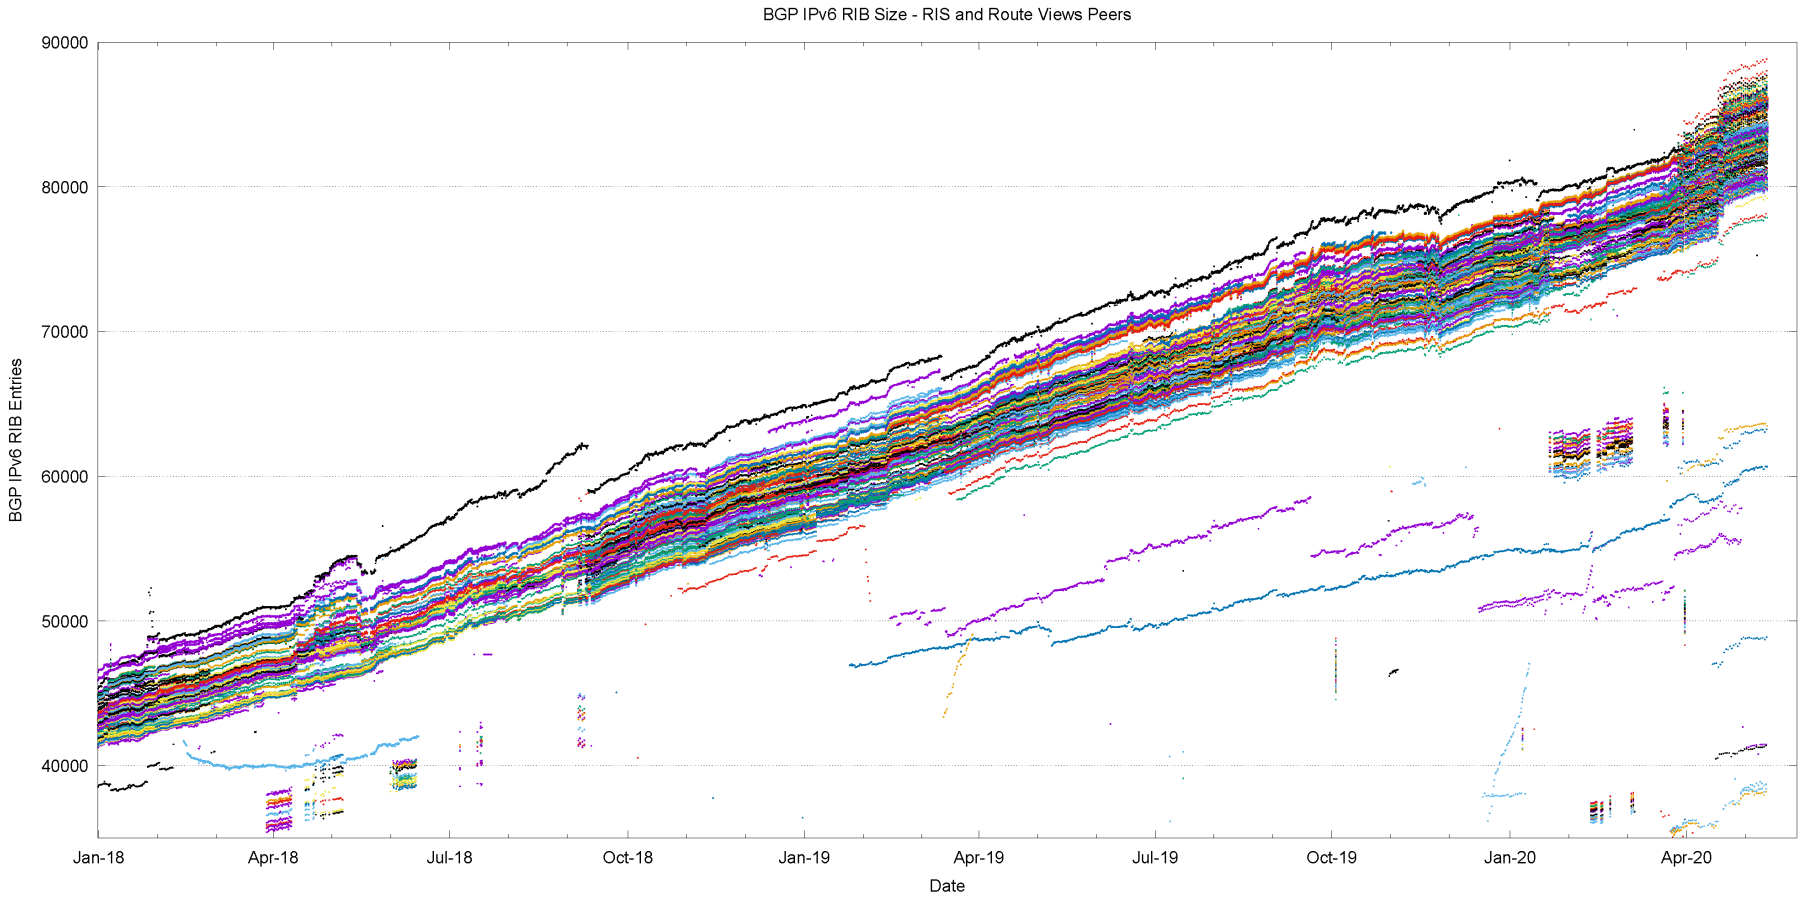

Routing Indicators for IPv6

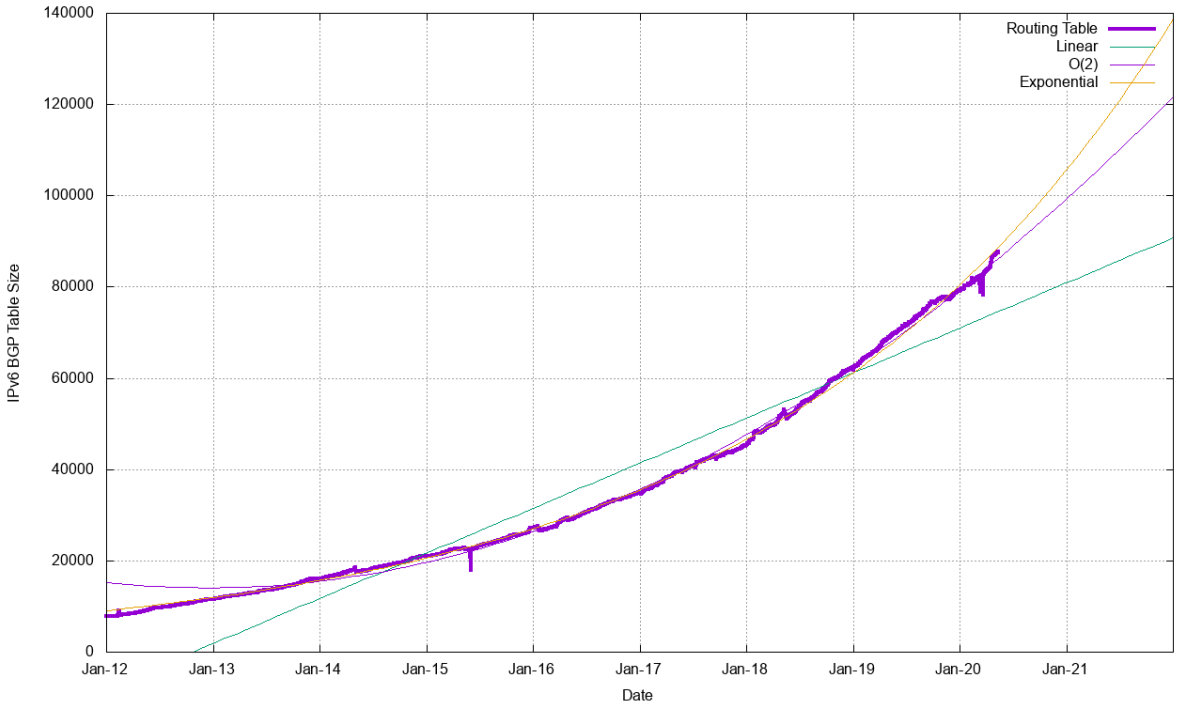

Routing prefixes – growing by some 20,000 prefixes per year

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

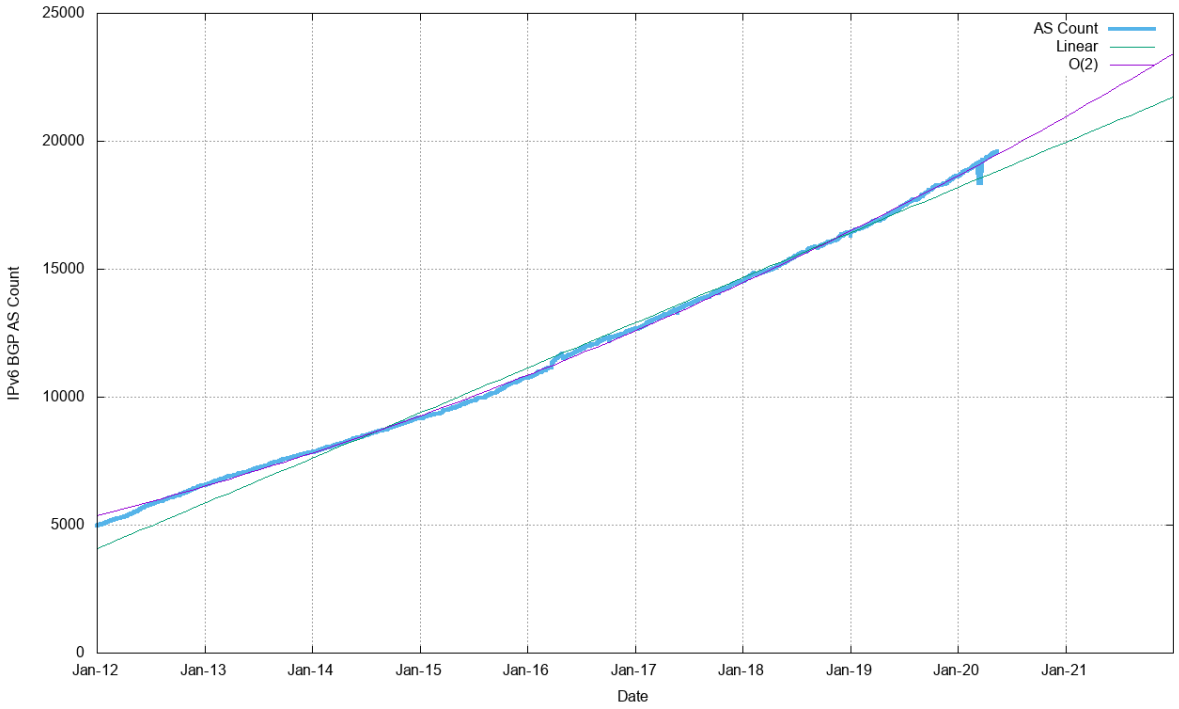

Routing Indicators for IPv6

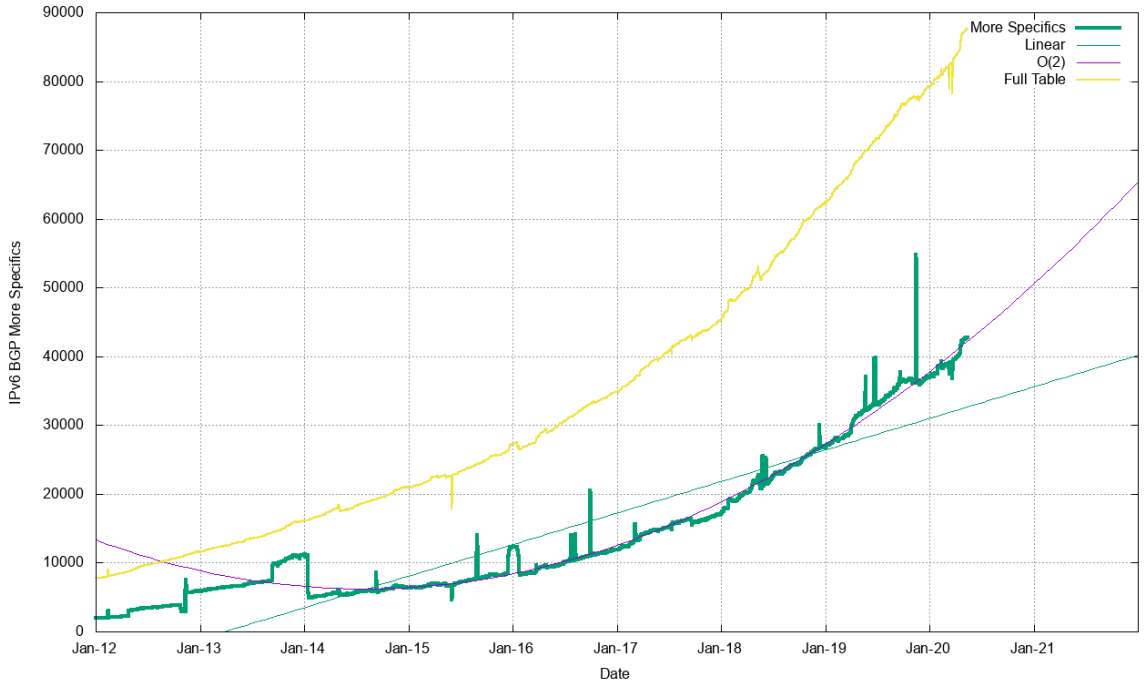

More Specifics now take up one half of the routing table

The average size of a routing advertisement is getting smaller

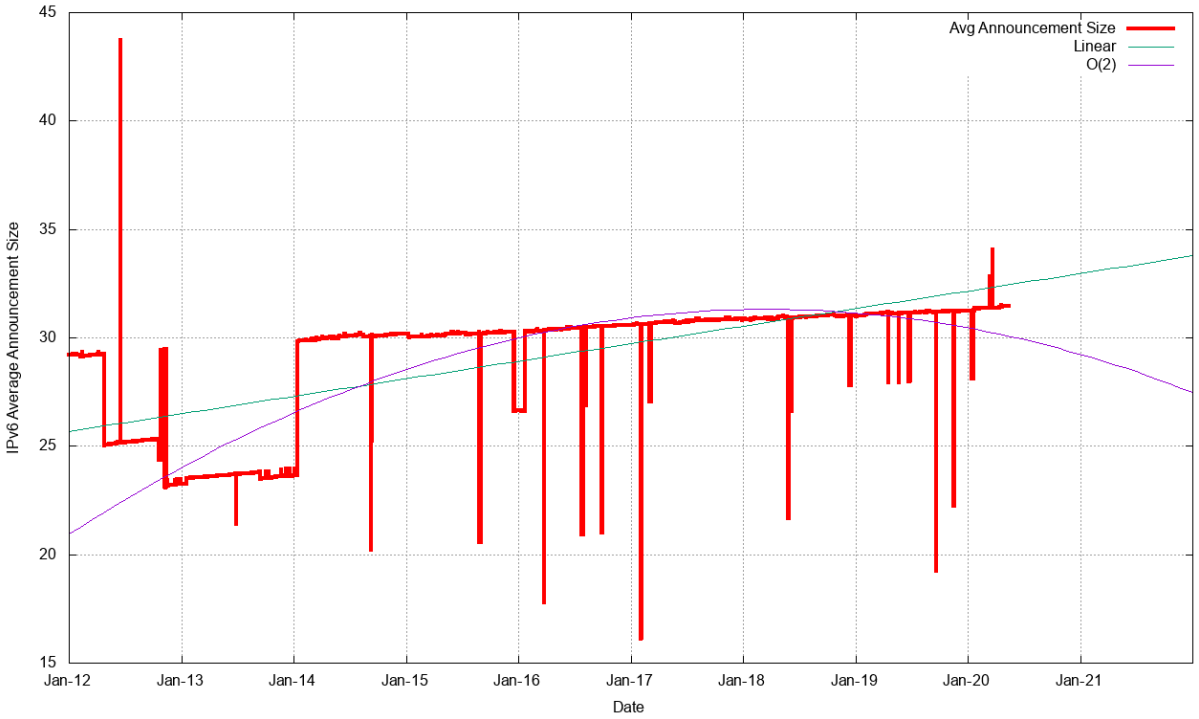

Routing Indicators for IPv6

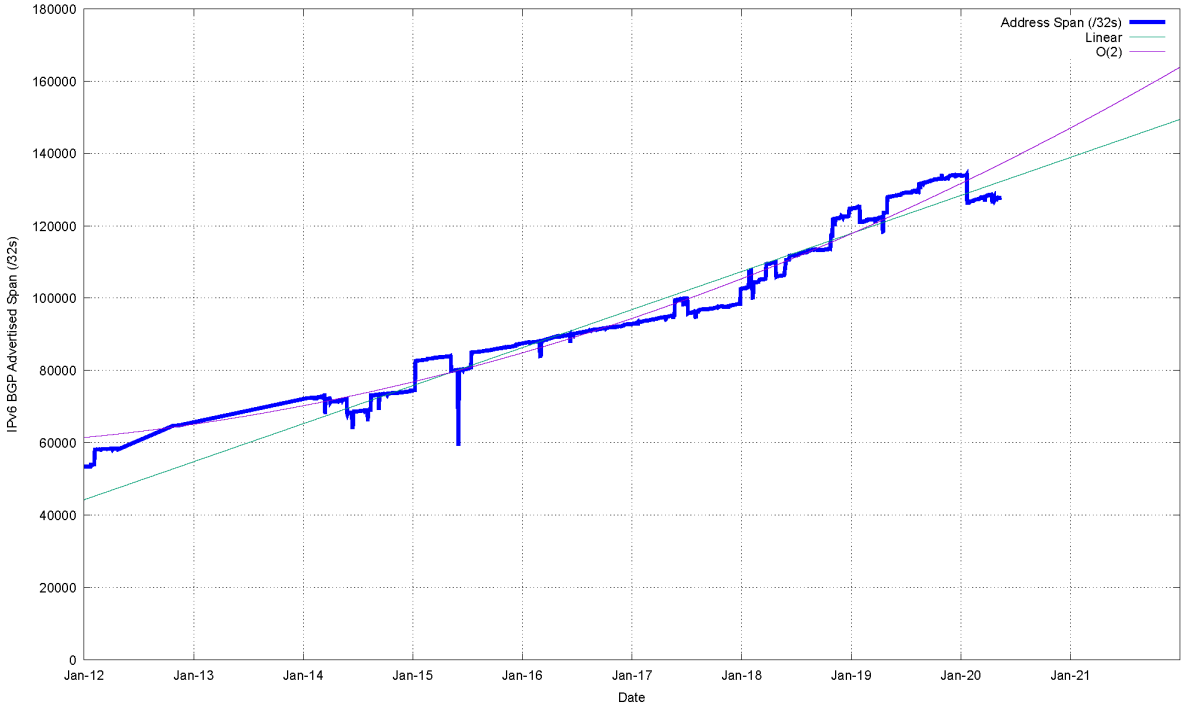

Advertised Address span is growing at an exponential rate

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

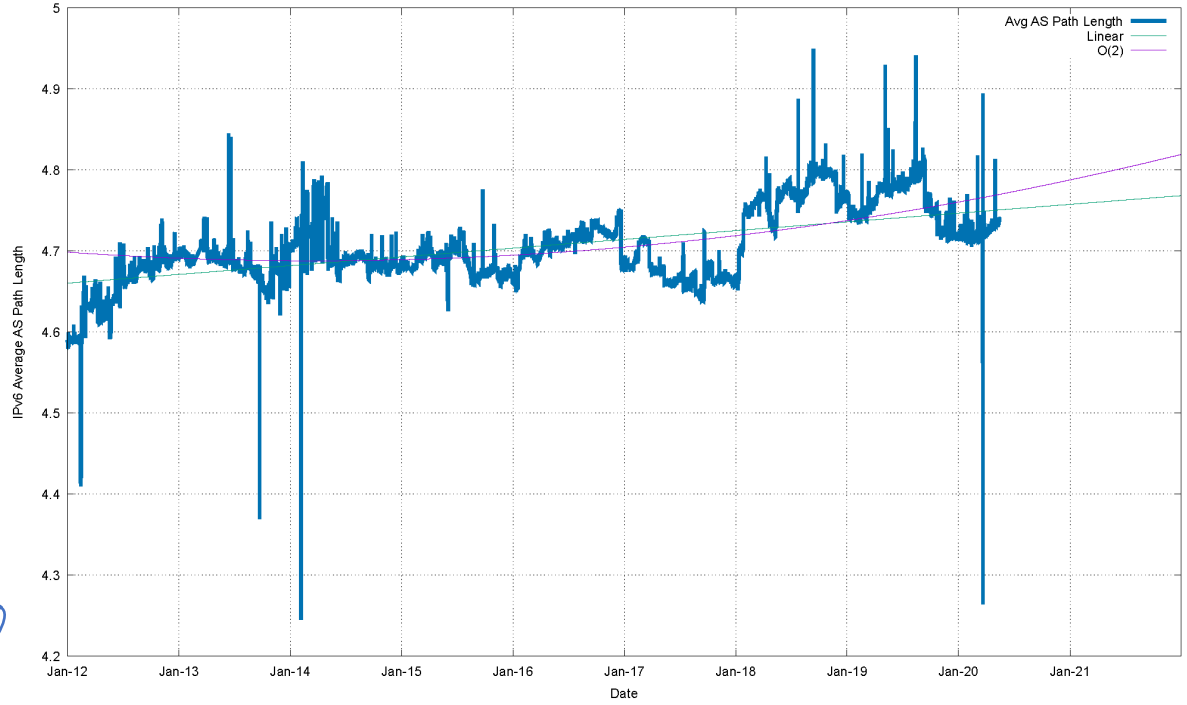

AS Adjacencies (AS131072)

15,720 out of 19,633 ASNs have 1 or 2 AS Adjacencies (80%)
698 ASNs have 10 or more adjacencies
3 ASNs have >1,000 adjacencies

4,900	AS6939	HURRICANE, US
1,093	AS3356	LEVEL3, US
1,013	AS1299	TELIANET Telia Carrier, EU
870	AS174	COGENT-174, US
830	AS2914	NTT-COMMUNICATIONS-2914, US
622	AS5539	SPACENET SpaceNET AG, DE
517	AS3257	GTT-BACKBONE GTT, DE
416	AS33891	CORE-BACKBONE GLOBAL NETWORK, DE
386	AS6461	ZAYO-6461, US
327	AS20473	AS-CHOOPA, US

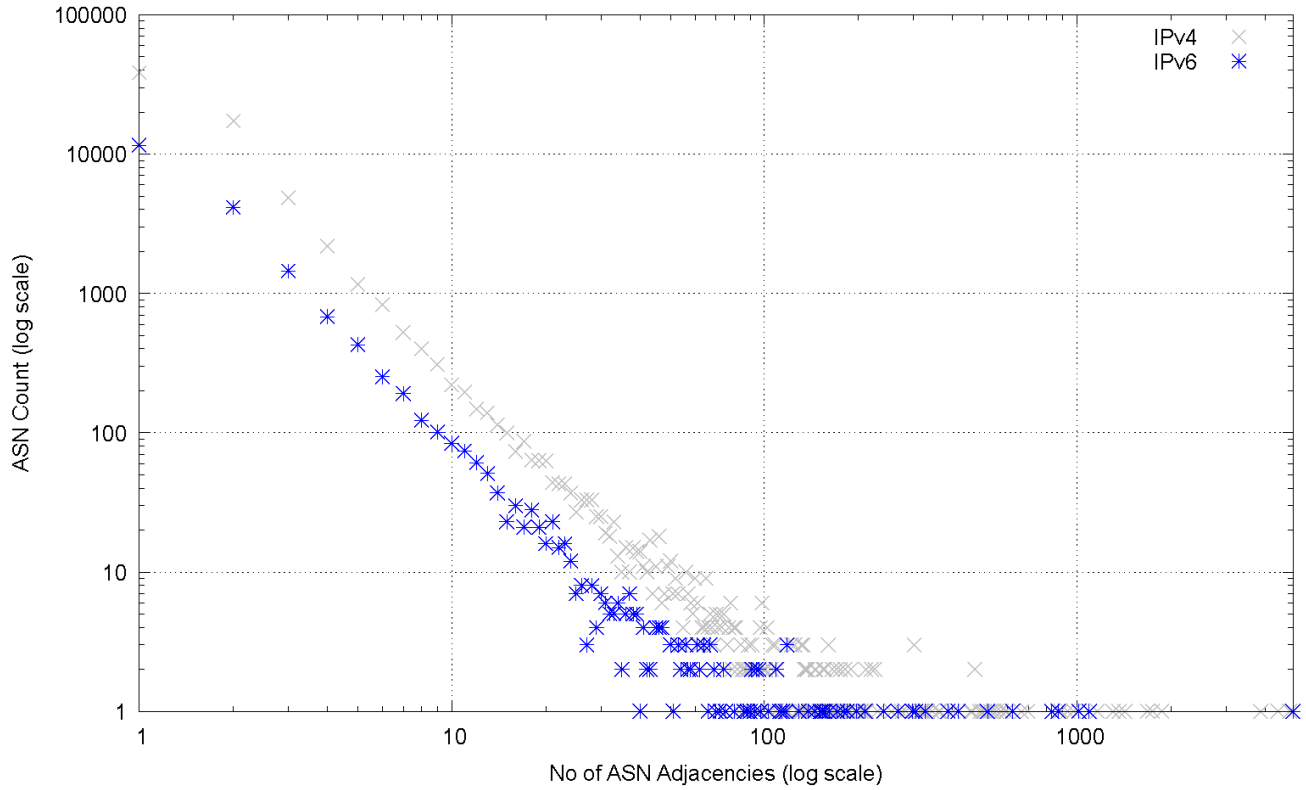

V6 in 2020

- Overall IPv6 Internet growth in terms of BGP is still increasing, and is currently at some **20,000 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

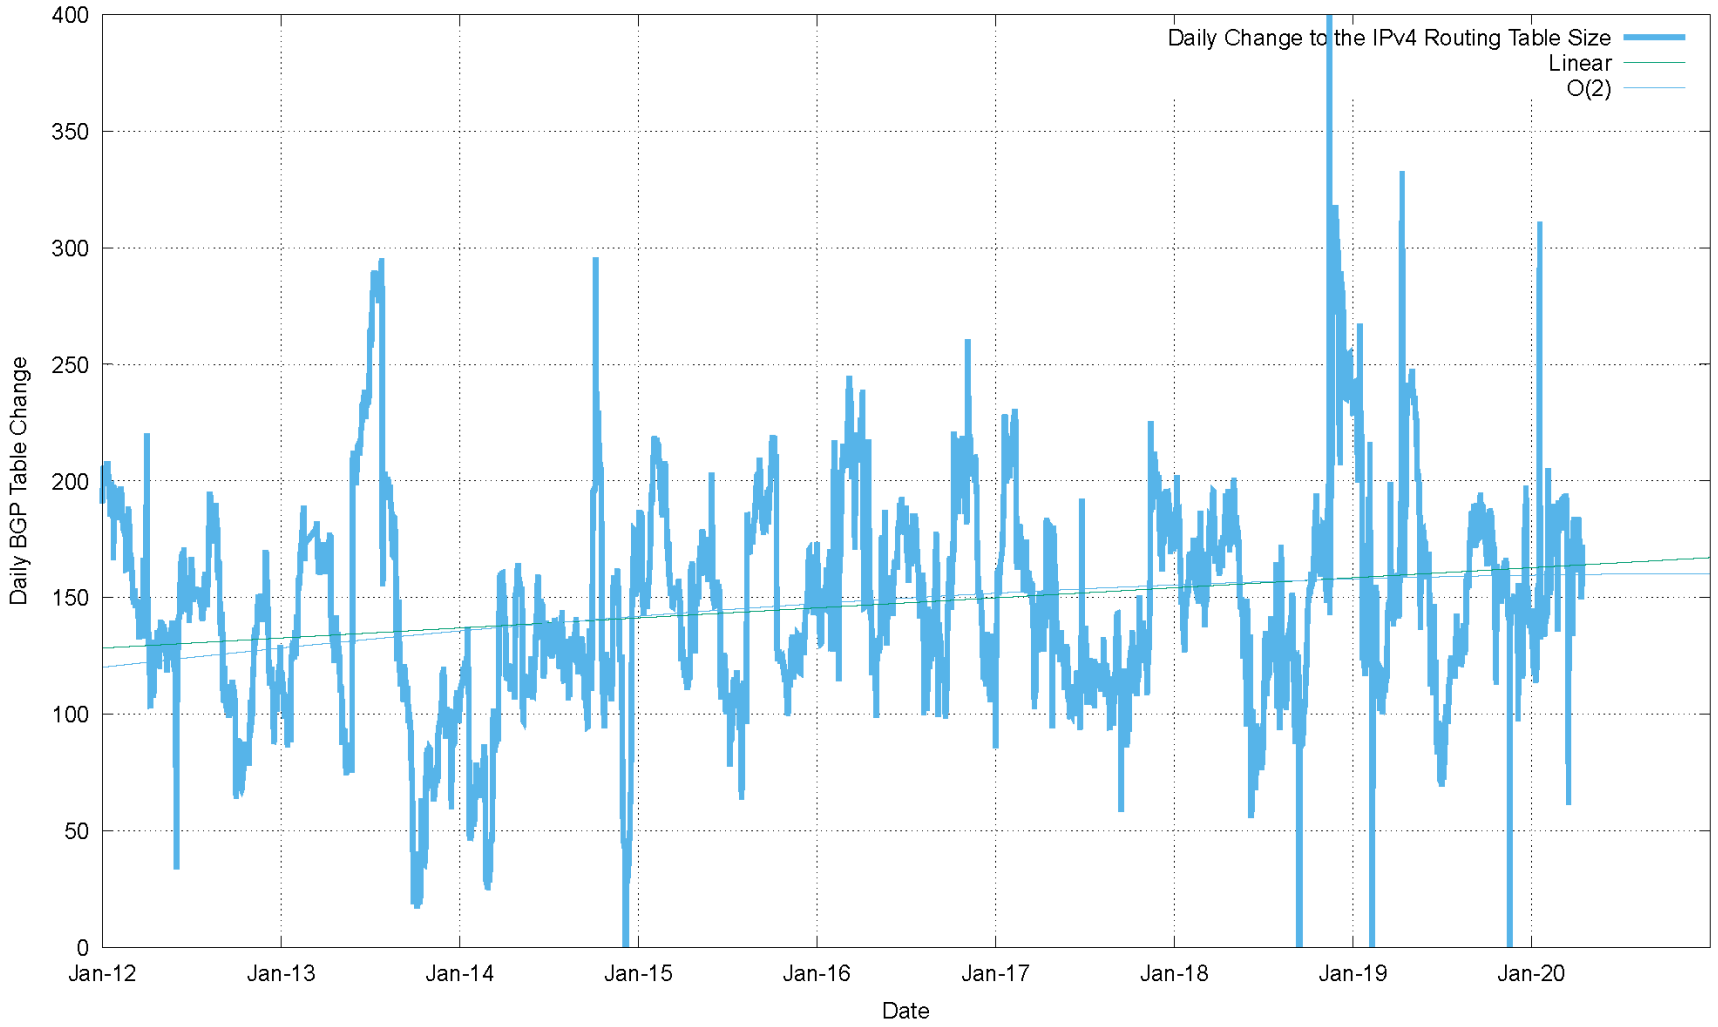

Growth in the V4 network appears to be consistent to a long-term average of 150 additional routes per day, or some 55,000 additional routes per year

V4 BGP Table Size Predictions

Jan *2017*	*646,000*
2018	*699,000*
2019	*760,000*
2020	*814,000*
2021	*864,000*
2022	*918,000*
2023	*972,000*
2024	*1,026,000*
2025	*1,081,000*
2026	*1,135,000*

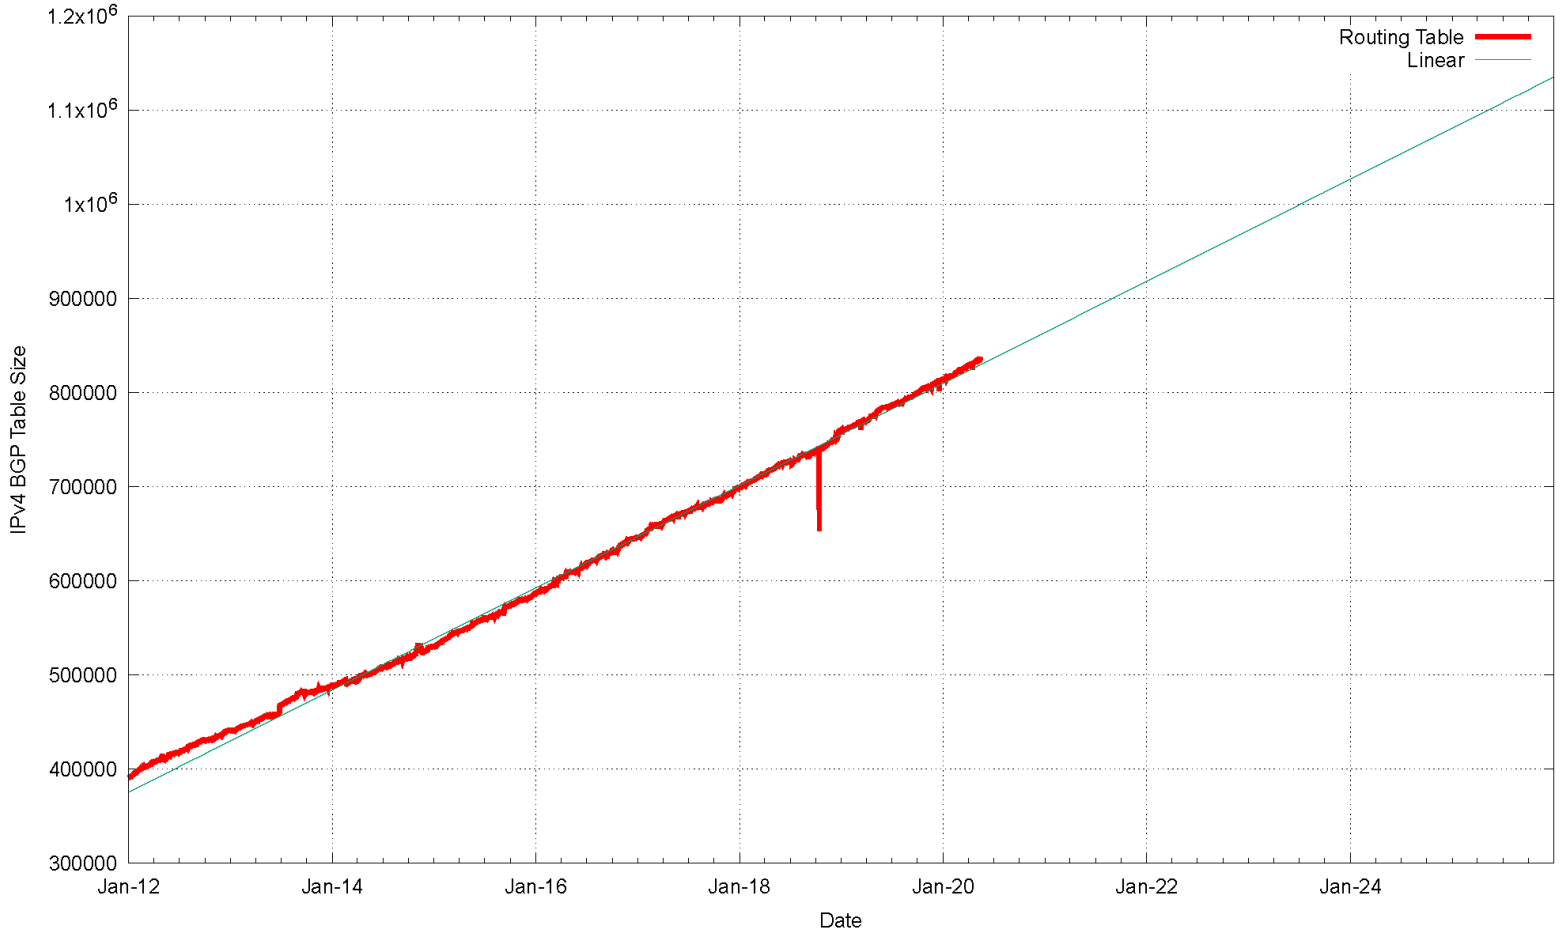

V6 - Daily Growth Rates

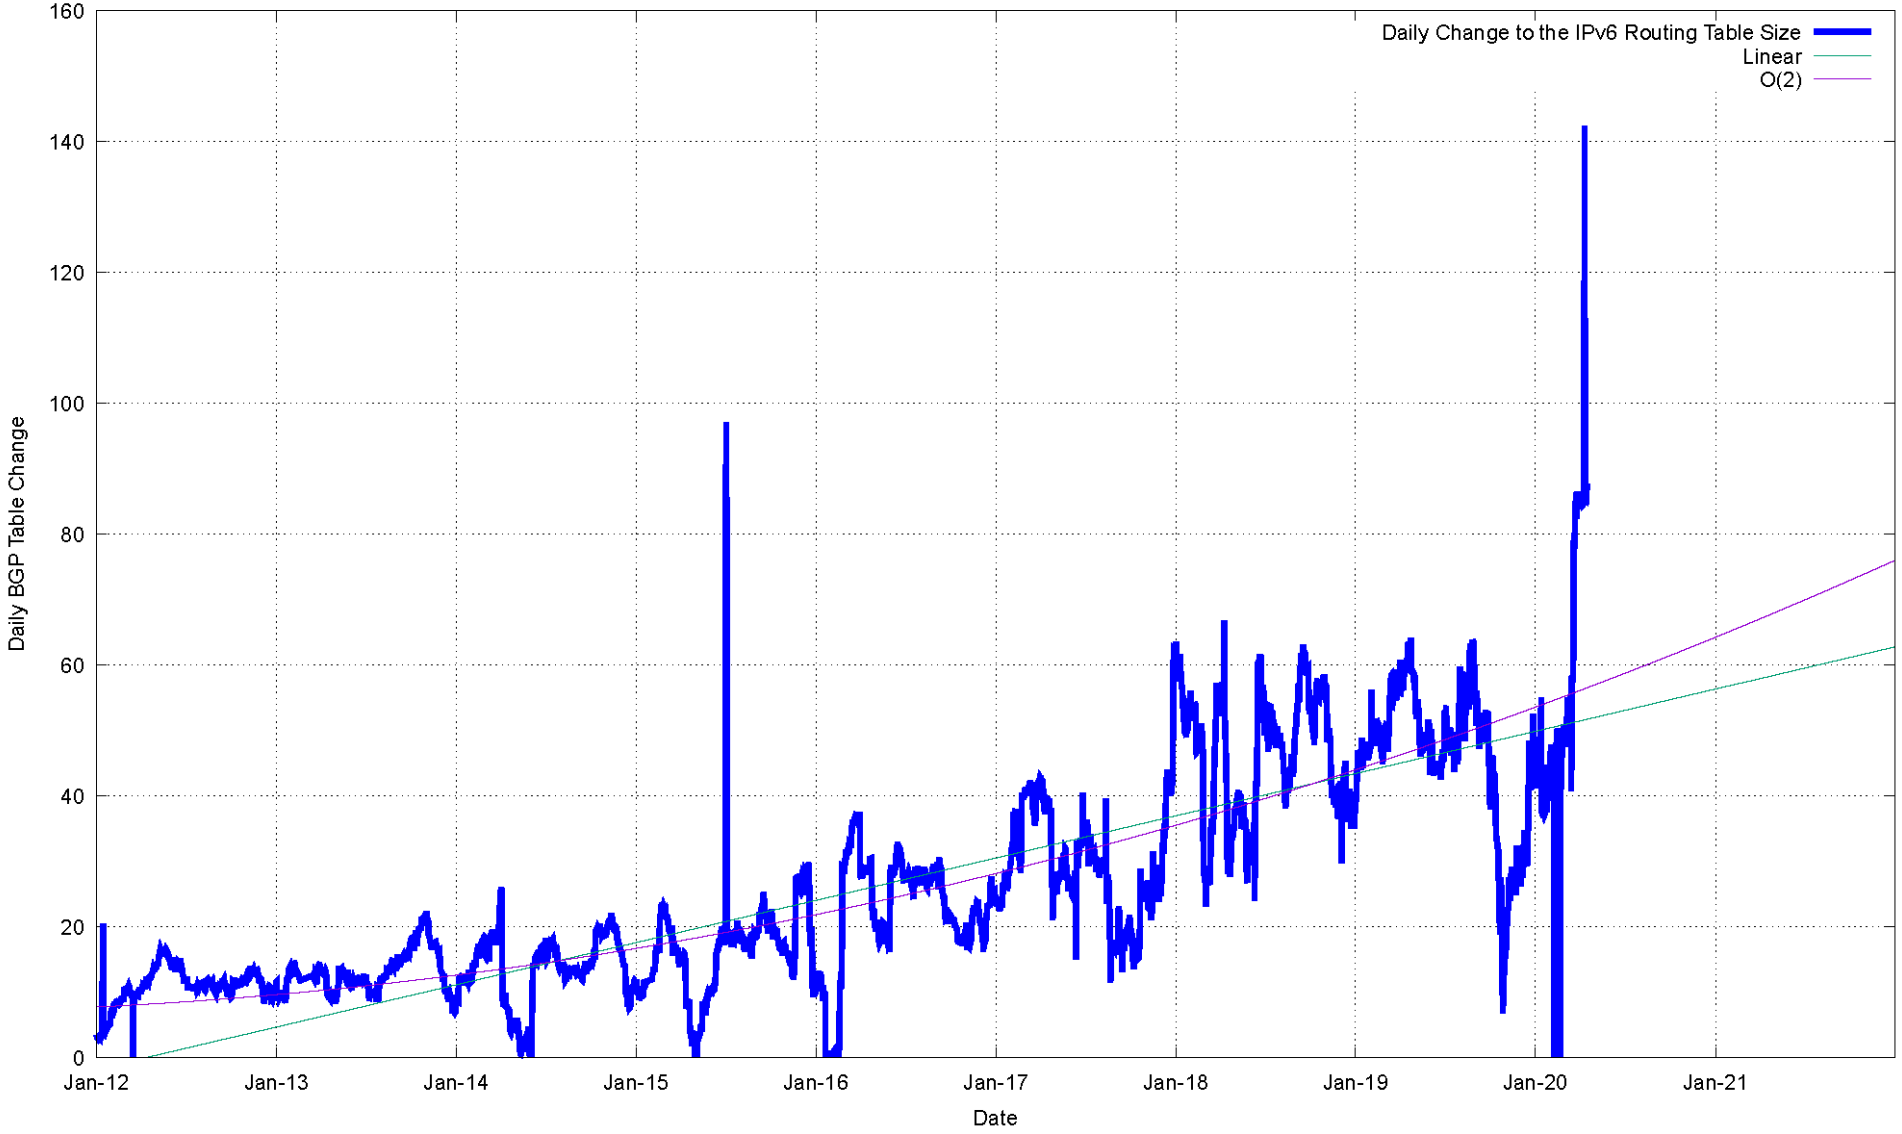

V6 BGP Table Size Predictions

	Linear	Exponential
Jan *2017*	*35,000*	
2018	*45,000*	
2019	*62,000*	
2020	*80,000*	
2021	*97,000*	*105,000*
2022	*114,000*	*137,000*
2023	*131,000*	*178,000*
2024	*148,000*	*233,000*
2025	*165,000*	*330,000*
2026	*182,000*	*395,000*

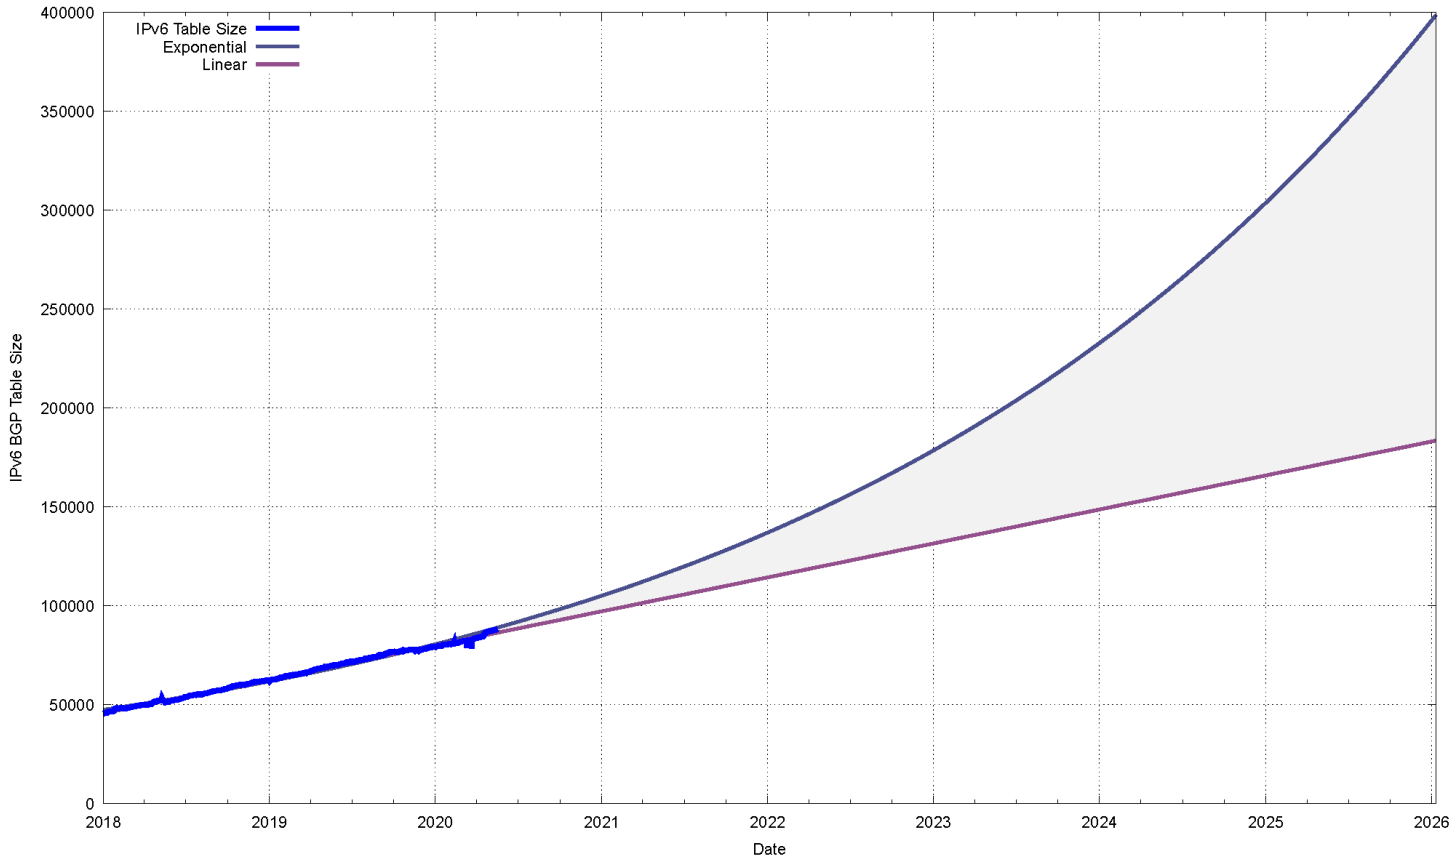

BGP Table Growth

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

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

But this is not that big a growth rate, and BGP can handle this scale of managed route objects with ease

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

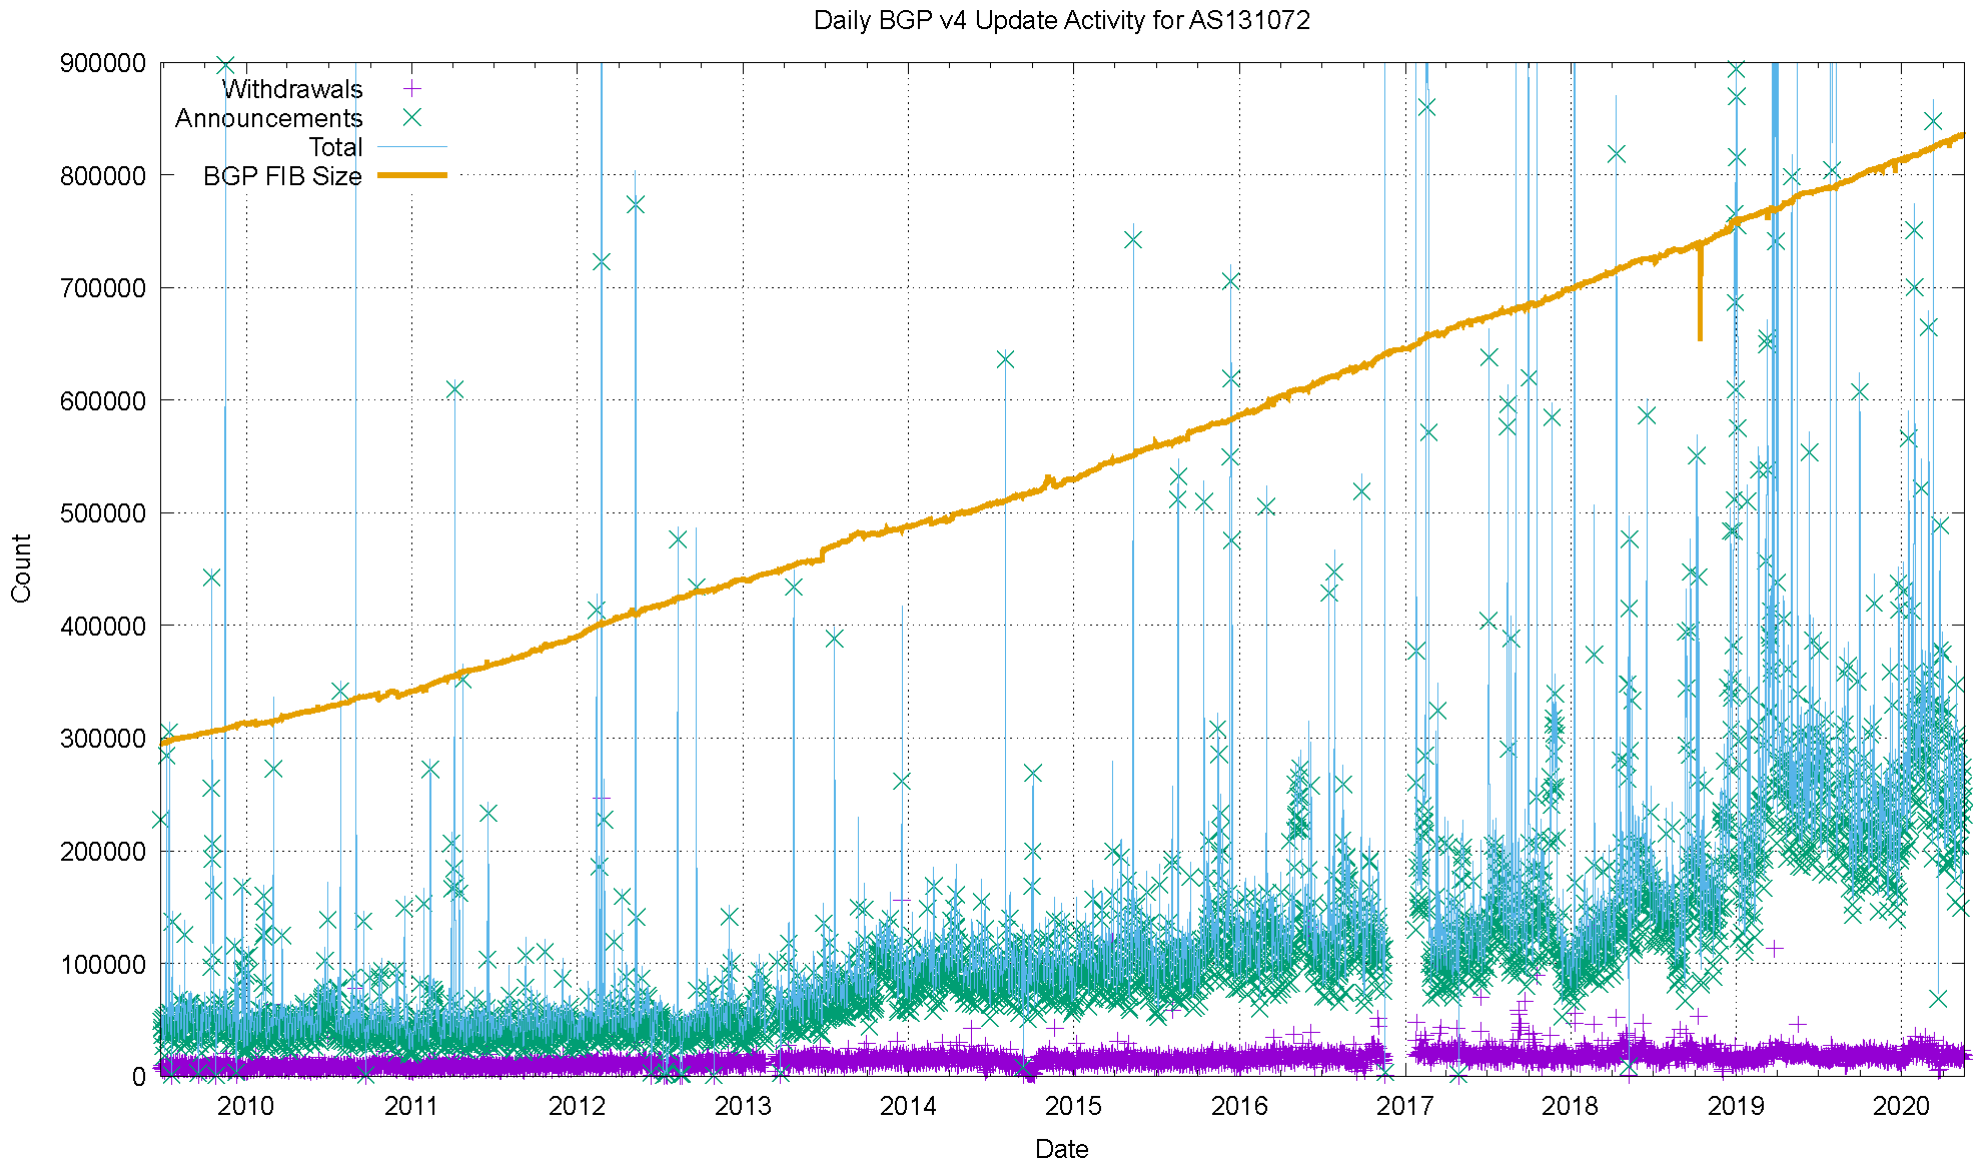

IPv4 BGP Convergence Performance

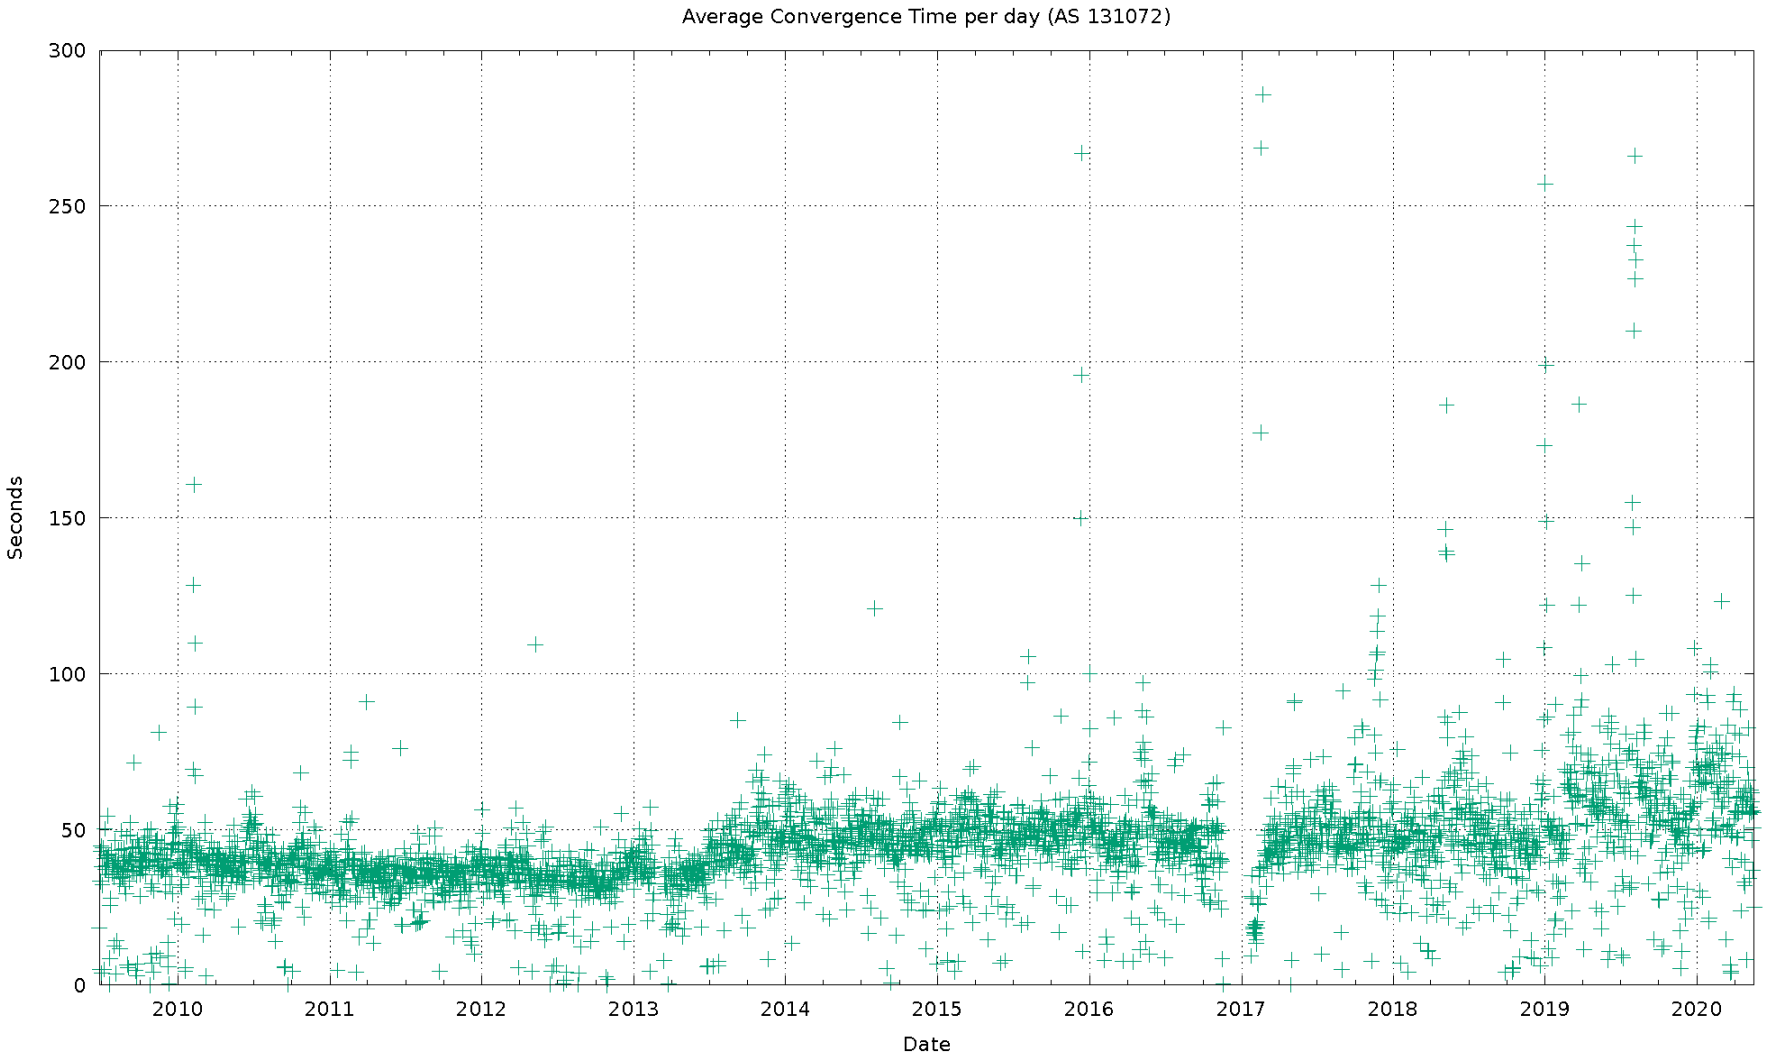

Updates in IPv4 BGP

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

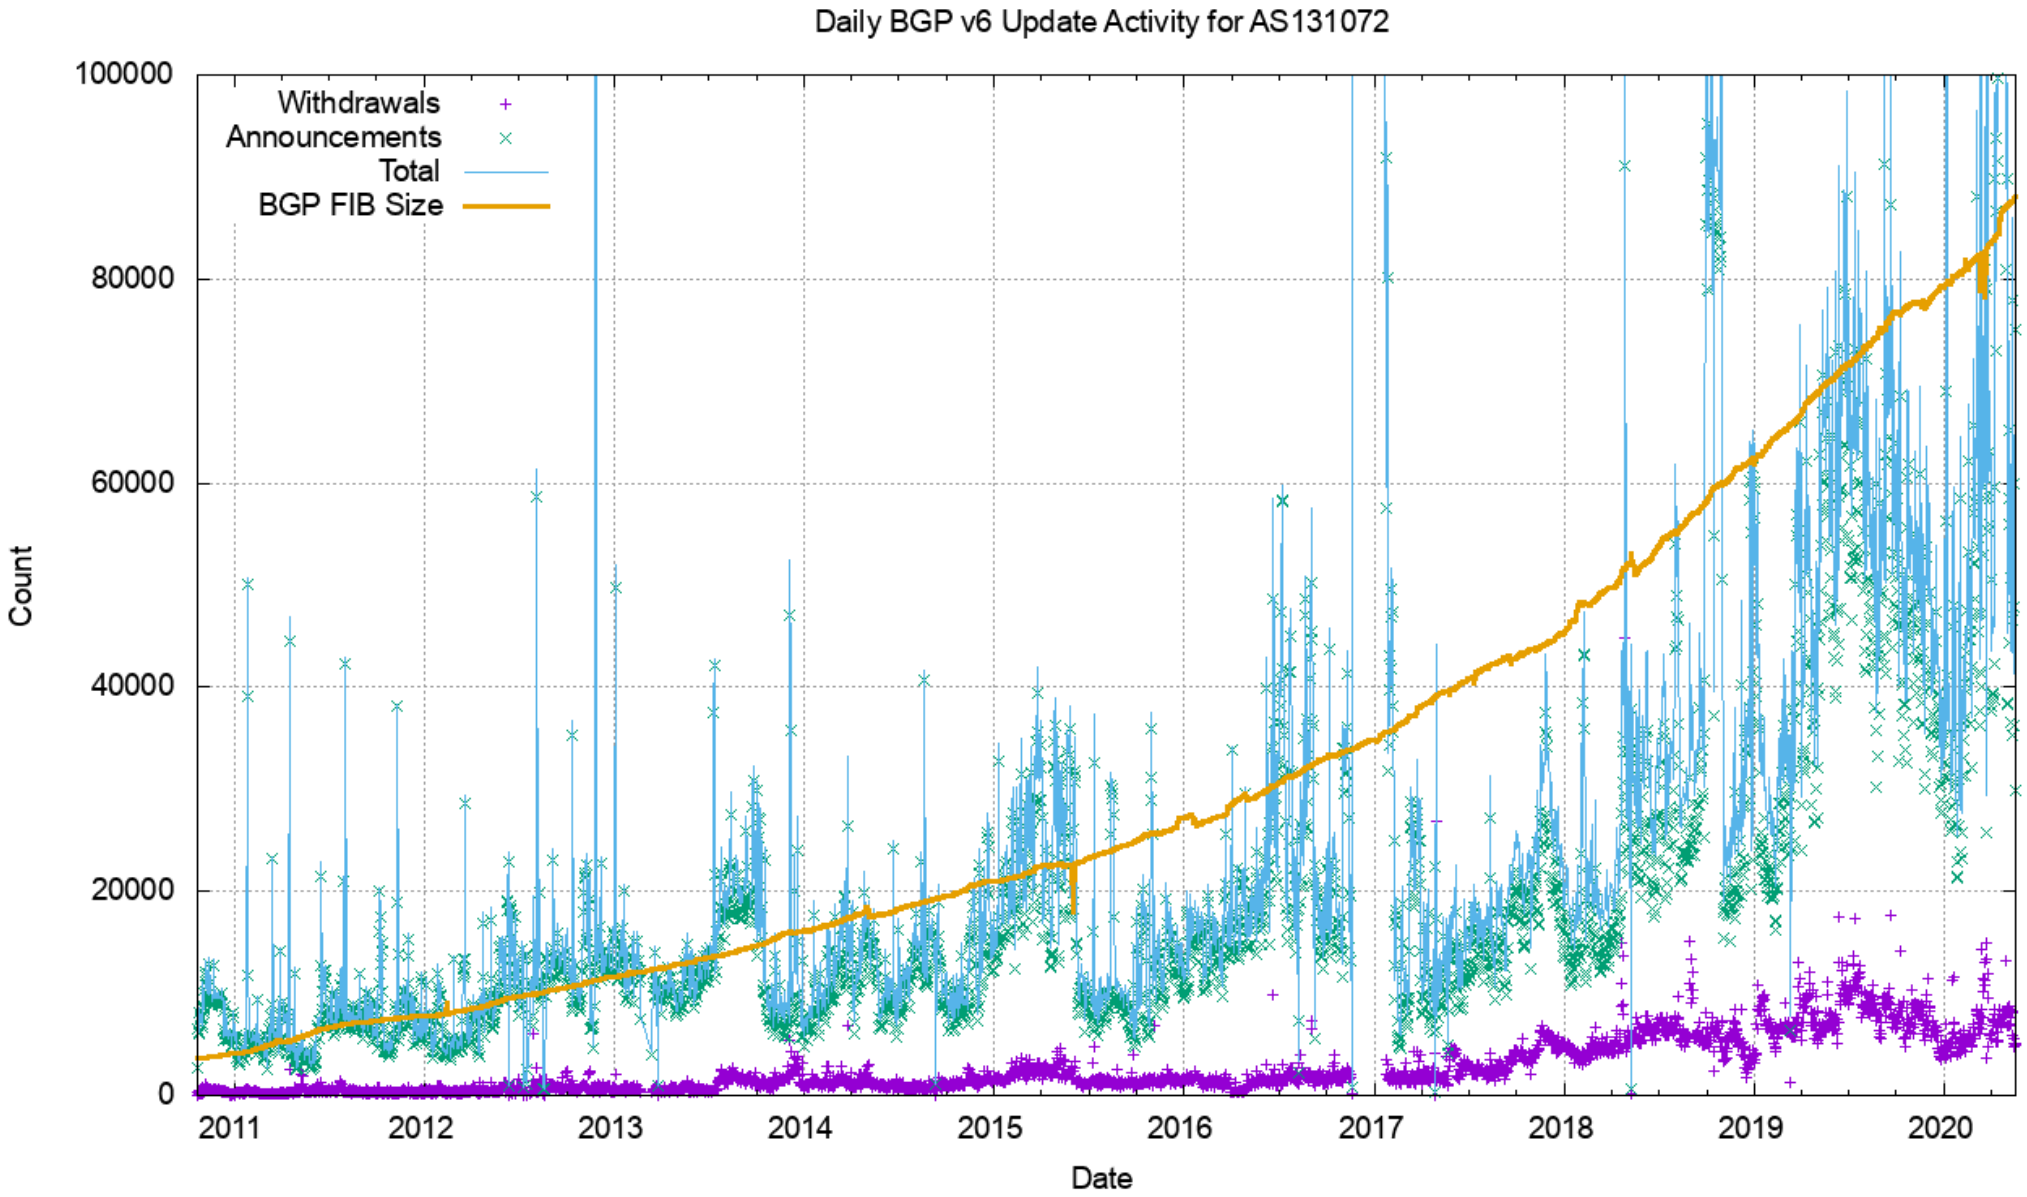

V6 Convergence Performance

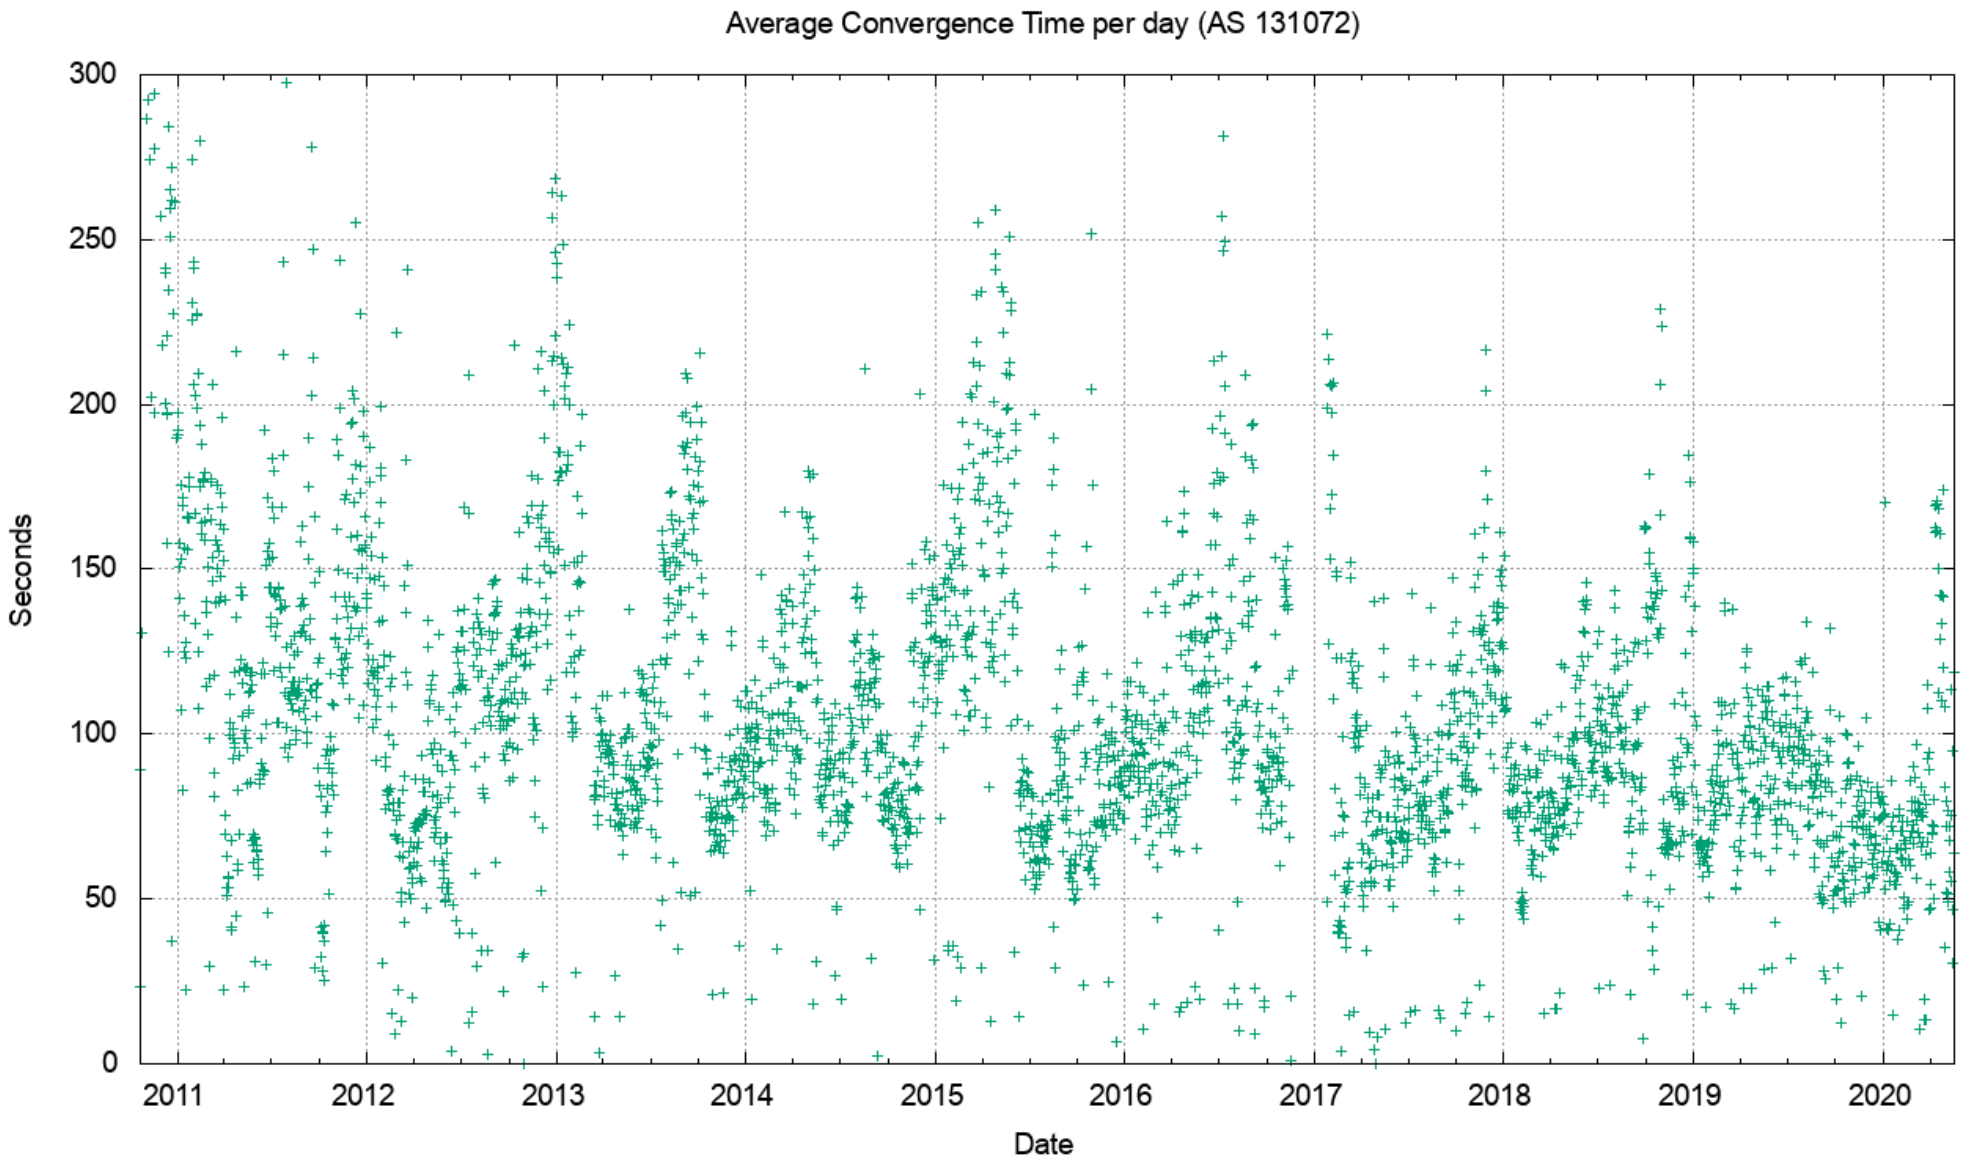

V6 Convergence Performance

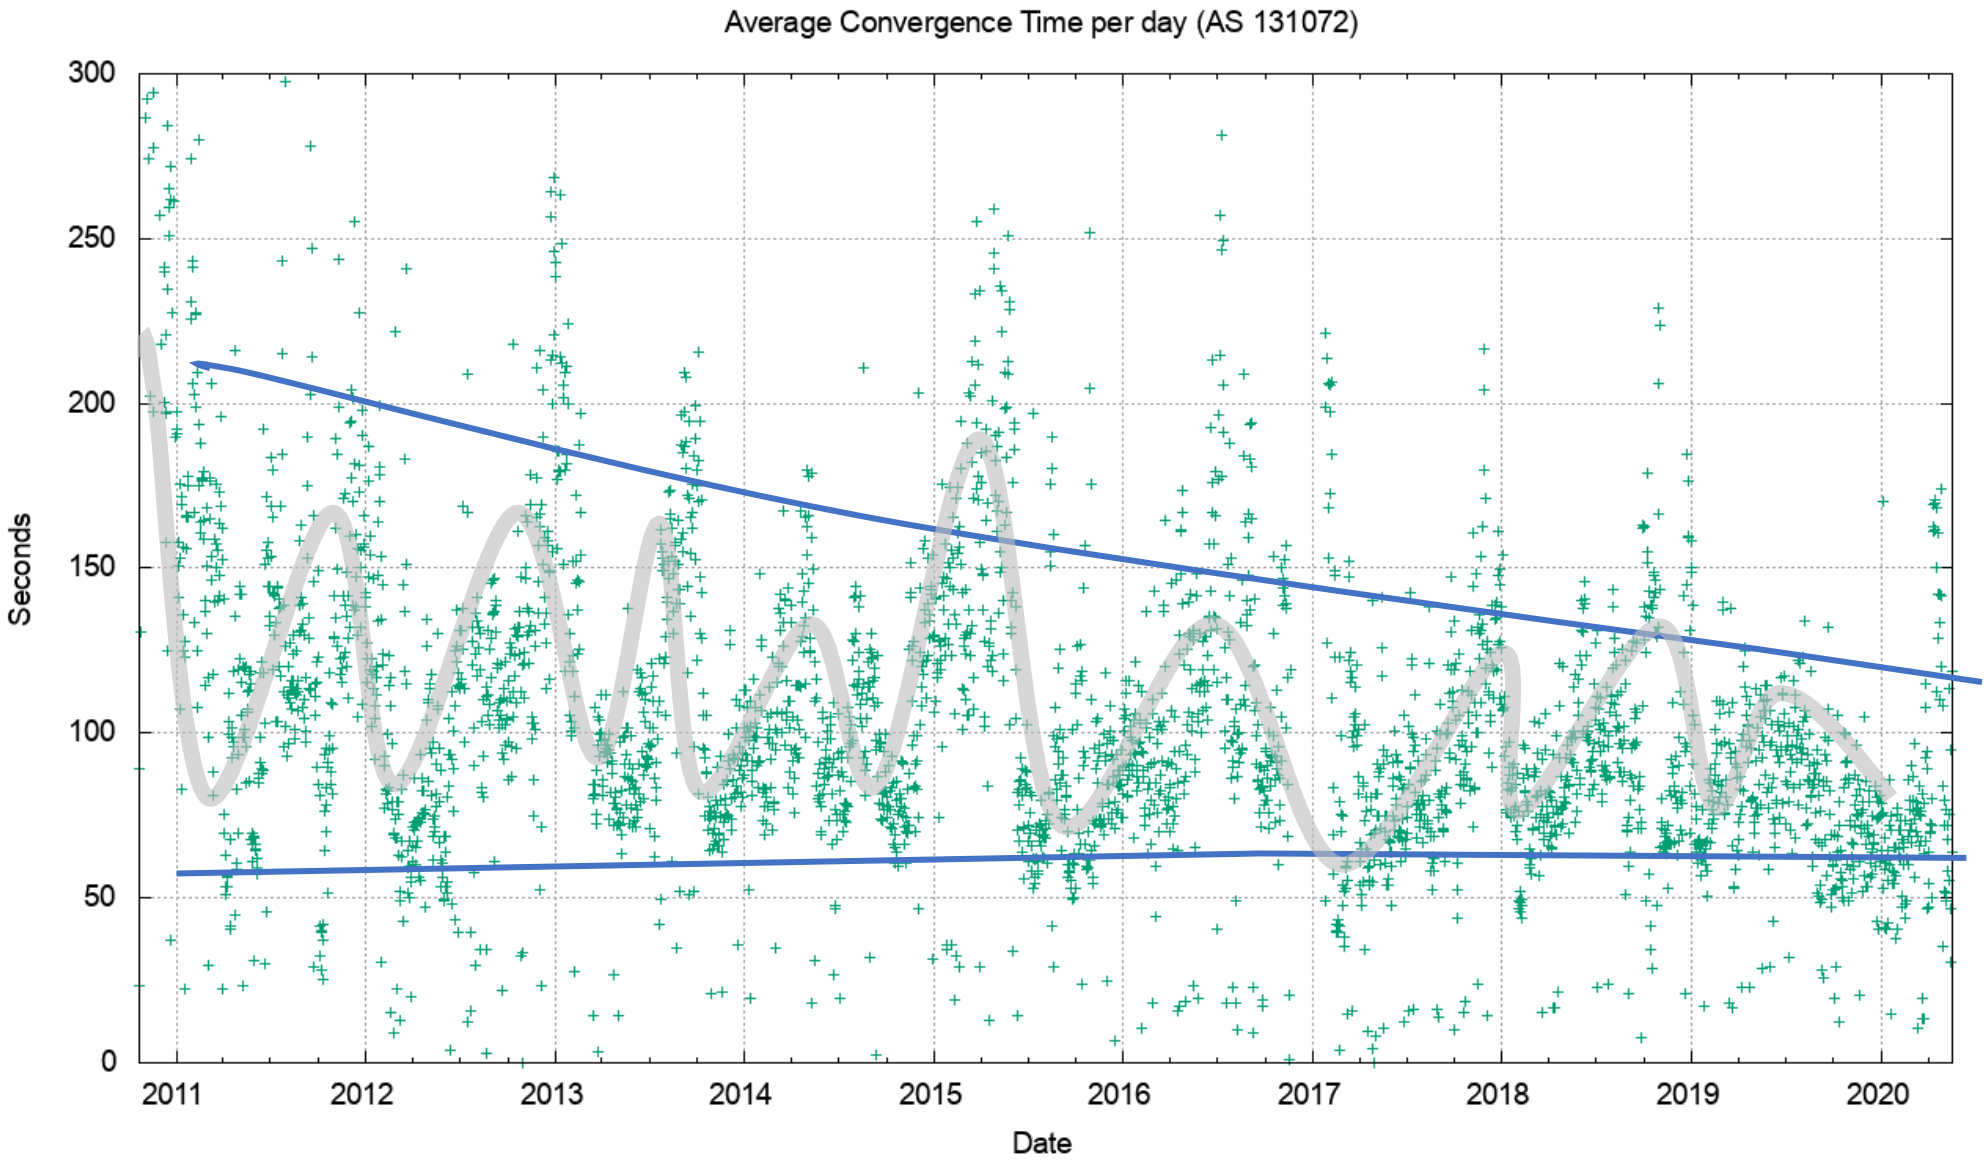

Routing Futures

- 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 memory partitioning - 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!