Routing in 2018

Geoff Huston
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

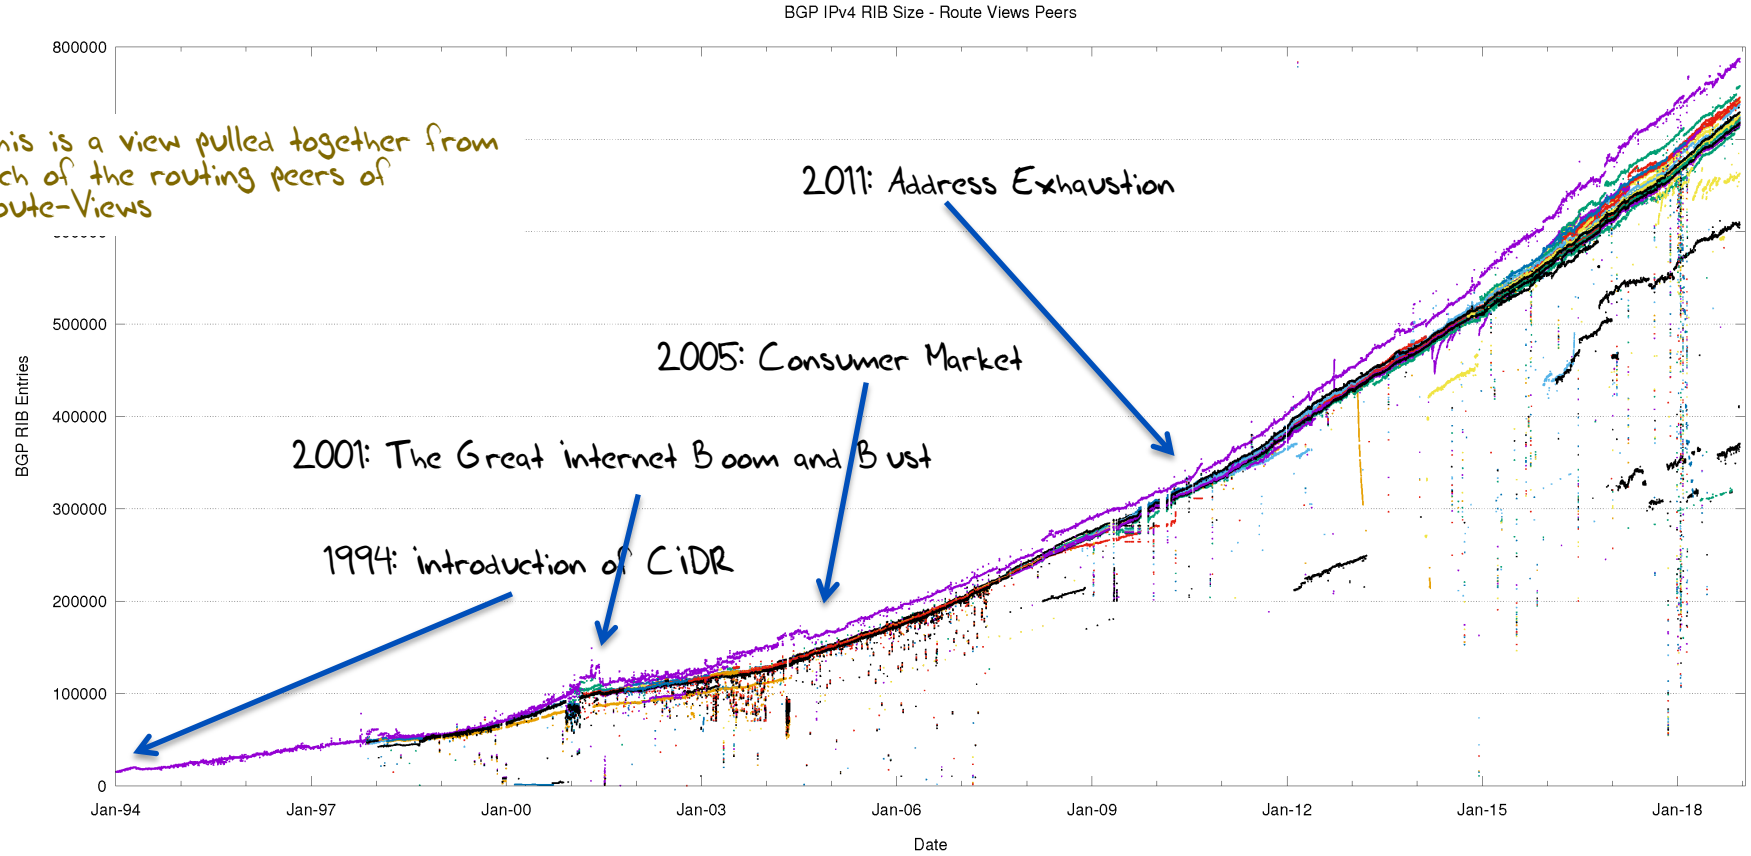

2016-2018 in detail

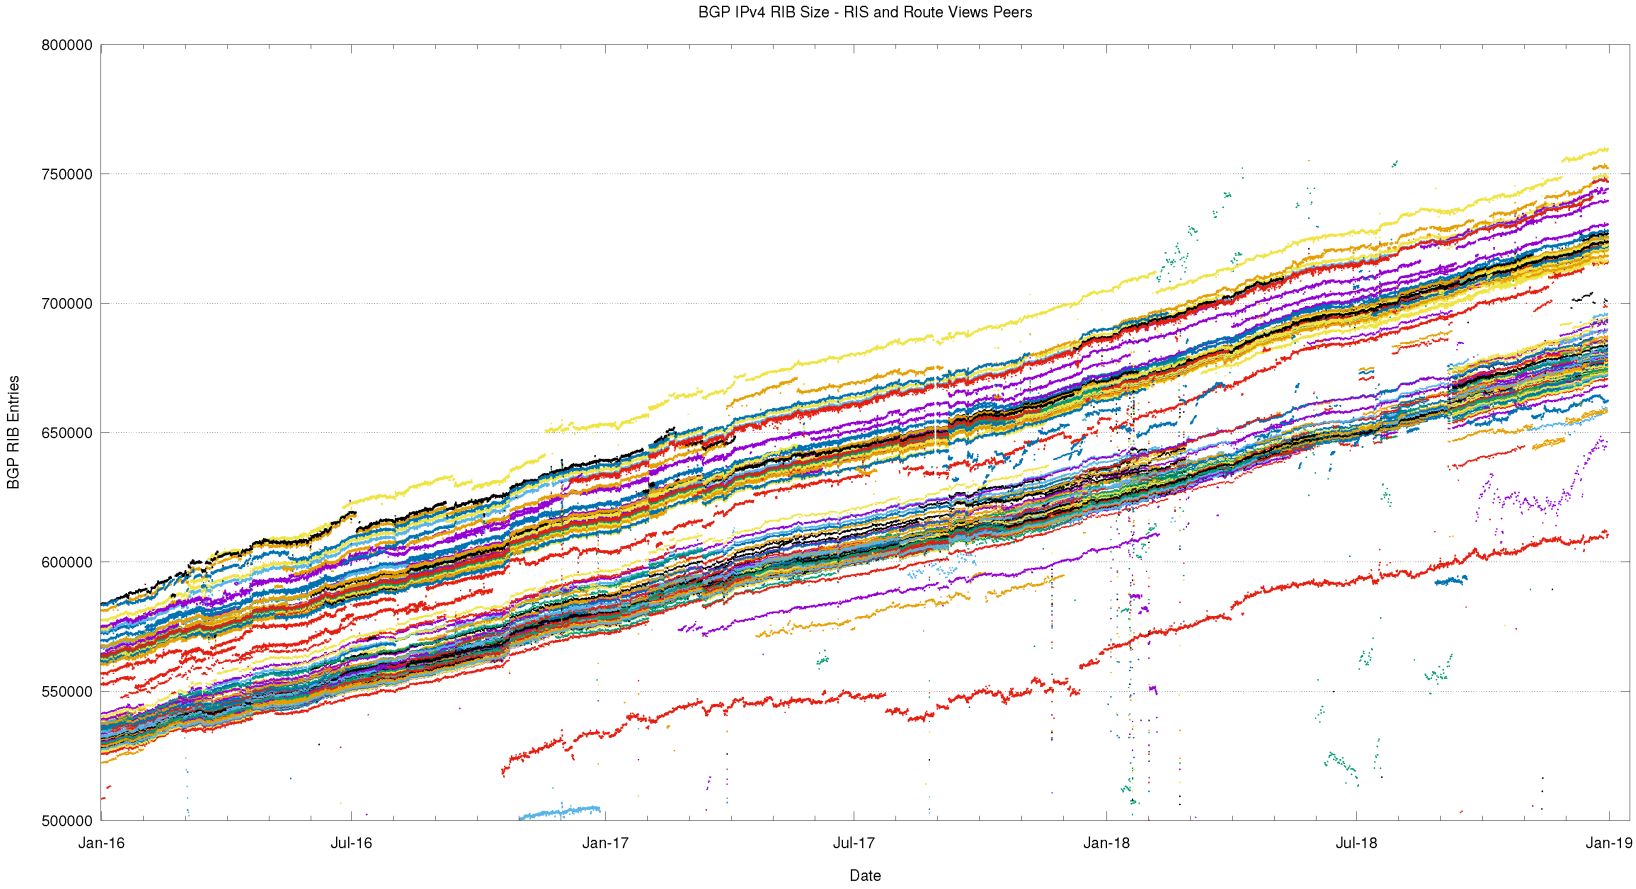

2016-2018 in detail

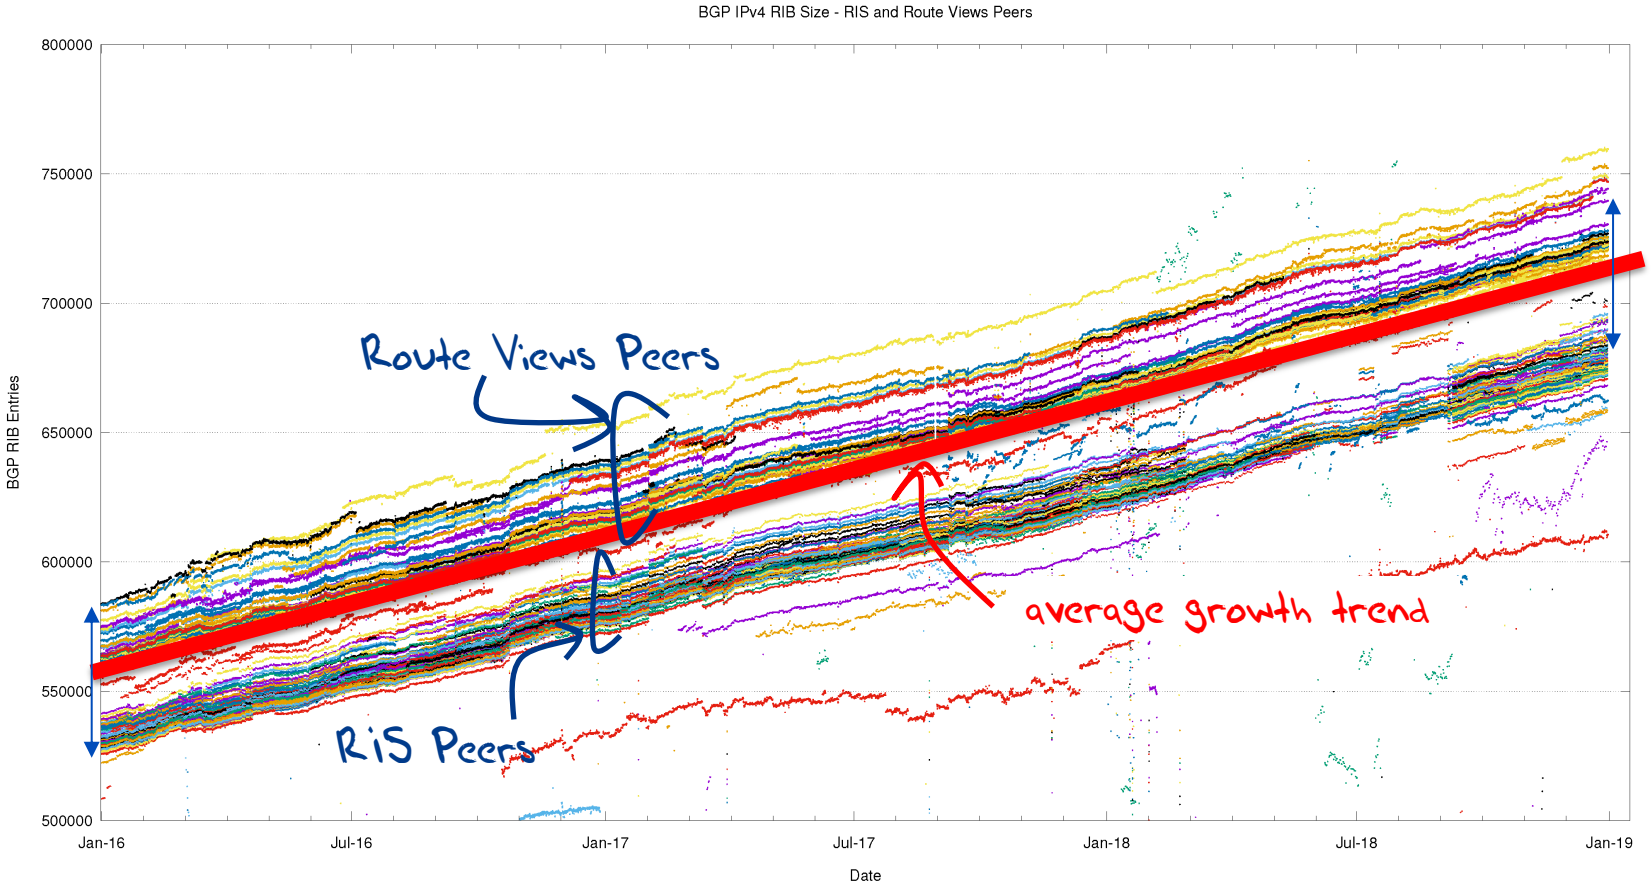

Routing Indicators for IPv4

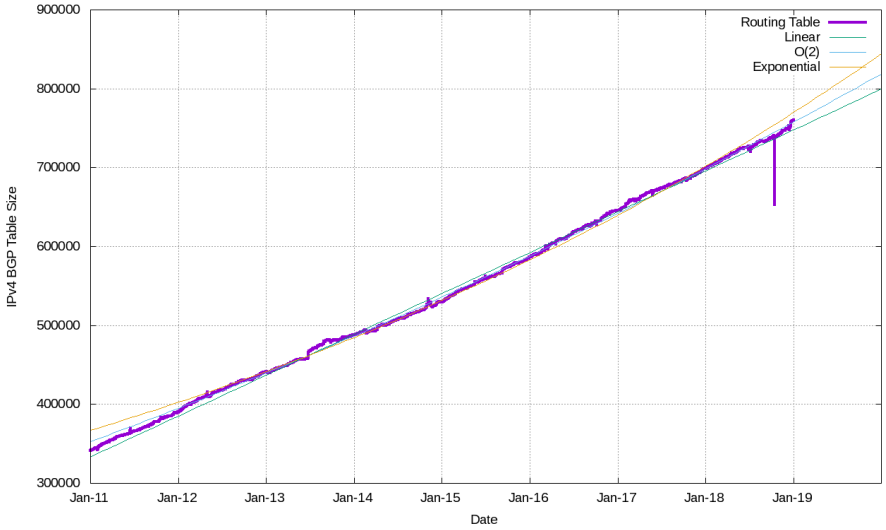

Routing prefixes – growing by some 52,000 prefixes per year

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

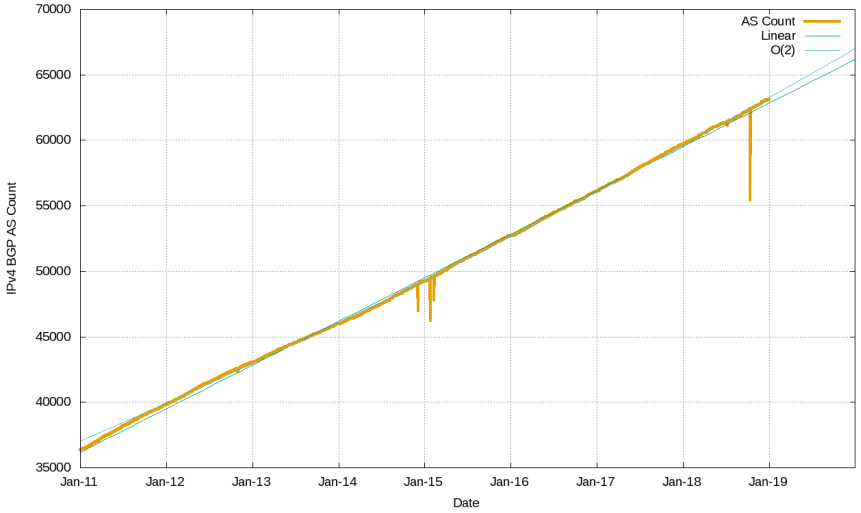

Routing Indicators for IPv4

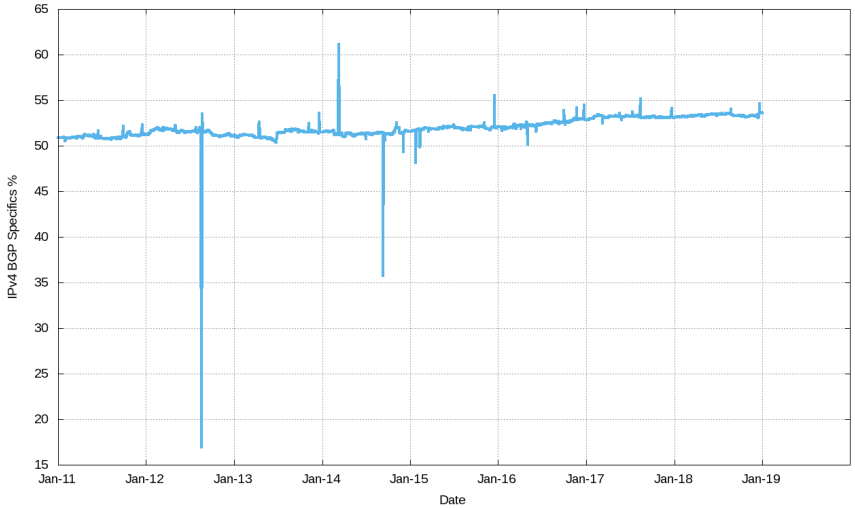

More Specifics are still taking up slightly more than one half of the routing table

But the average size of a routing advertisement continues to shrink

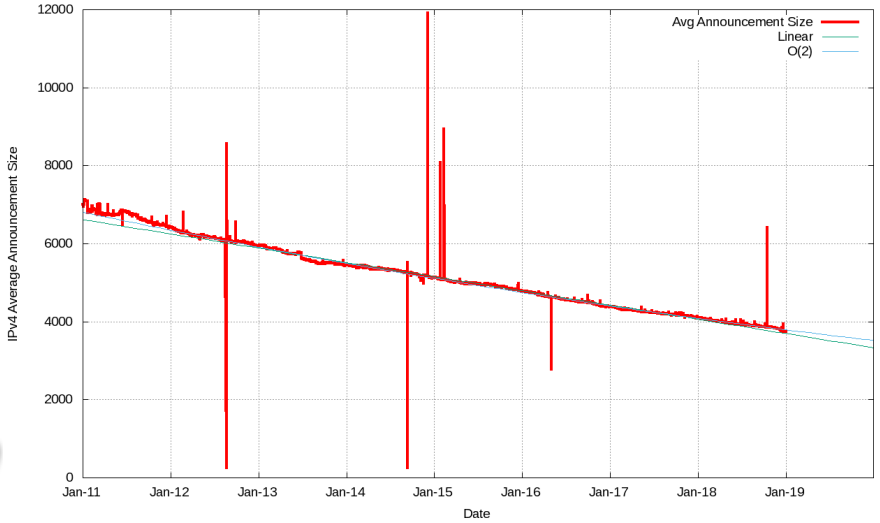

Routing Indicators for IPv4

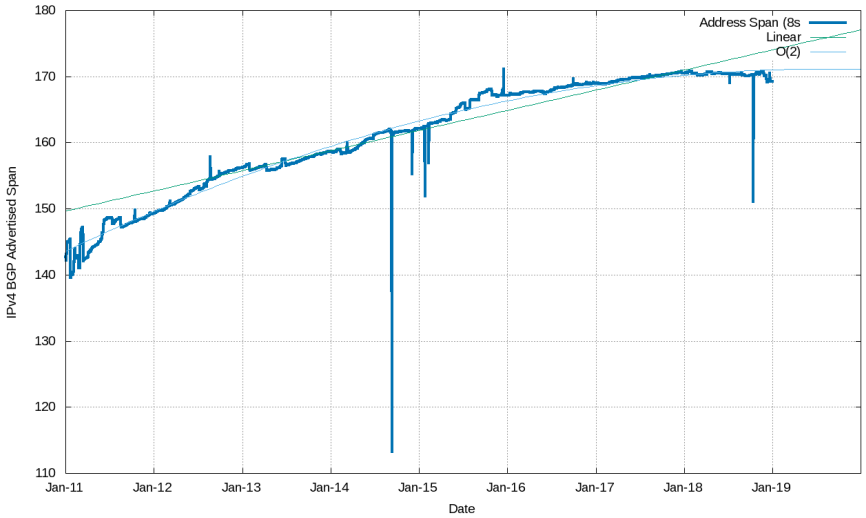

Address Exhaustion is now visible in the extent of advertised address space

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

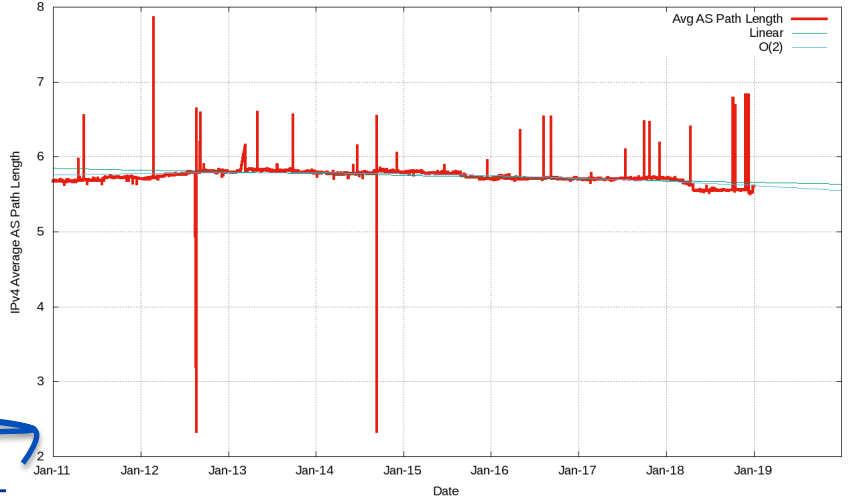

AS Adjacencies (AS131072)

51,613 out of 63,080 ASNs have 1 or 2 AS Adjacencies (82%)

1,803 ASNs have 10 or more adjacencies

9 ASNs have >1,000 adjacencies

4,144	AS6939	HURRICANE - Hurricane Electric, Inc., US
4.032	AS3356	LEVEL3 - Level 3 Communications, Inc., US
3,702	AS174	COGENT-174 - Cogent Communications, US
1,724	AS6461	ZAYO Bandwidth, US
1,646	AS7018	ATT-INTERNET4 - AT&T Services, Inc., US
1,618	AS3549	LVLT – Level 3 Parent, US
1,428	AS3257	GTT-Backbone, DE
1,377	AS2914	NTT America, US
1,208	AS209	CENTURYLINK, US
957	AS701	Verizon Business, US

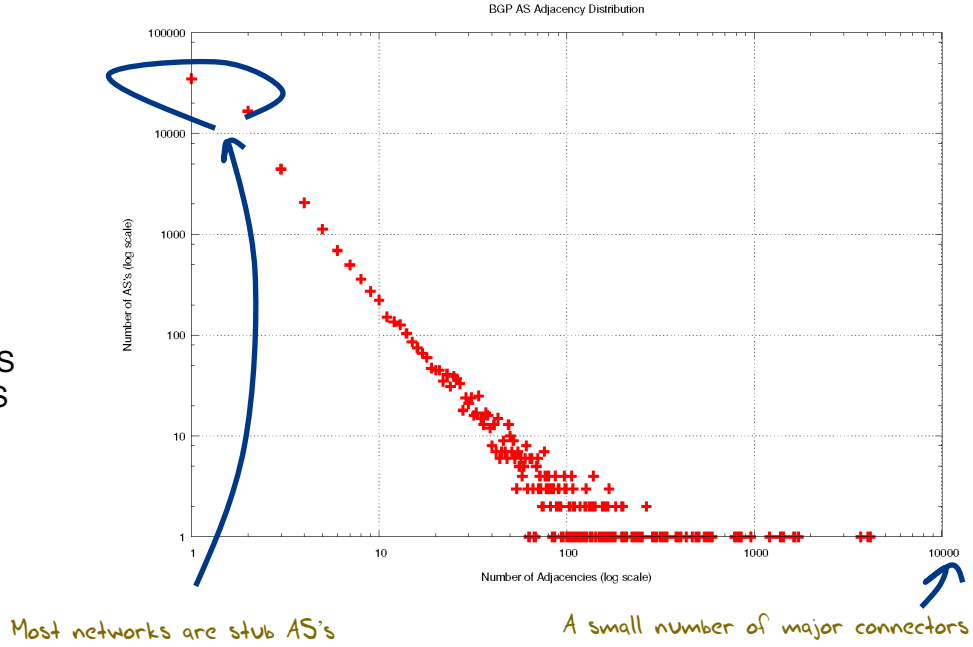

What happened in 2018 in V4?

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 750,000 by the end of 2018
- The pace of growth of the routing table is still relatively constant at ~52,000 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?

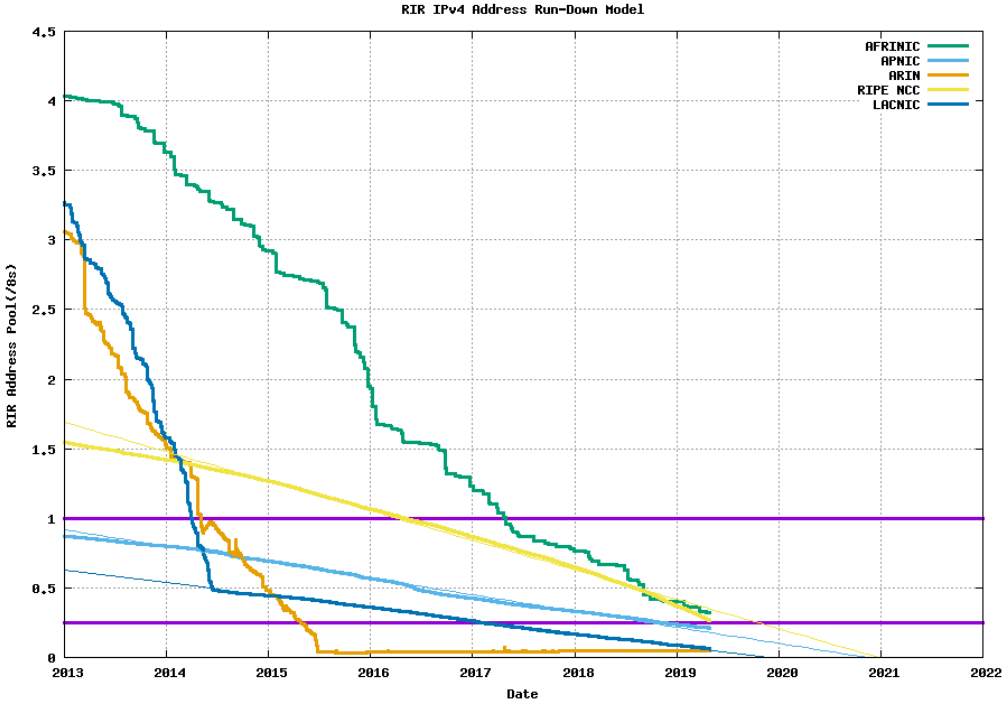

RIR Address Pool runout projections (as of April 2019):

ARIN	– no free pool left
AFRINIC	– May 2020
LACNIC	– November 2019
APNIC	– November 2020
RIPE NCC	– January 2020

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

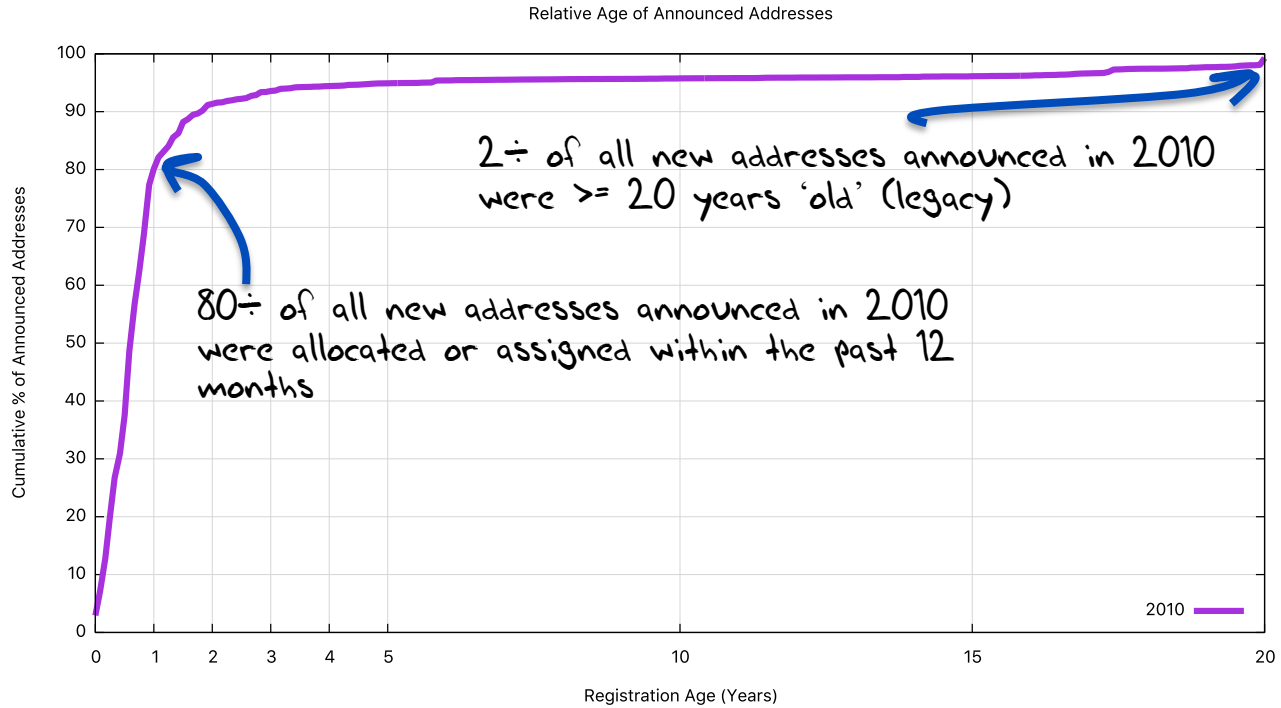

Advertised Address "Age"

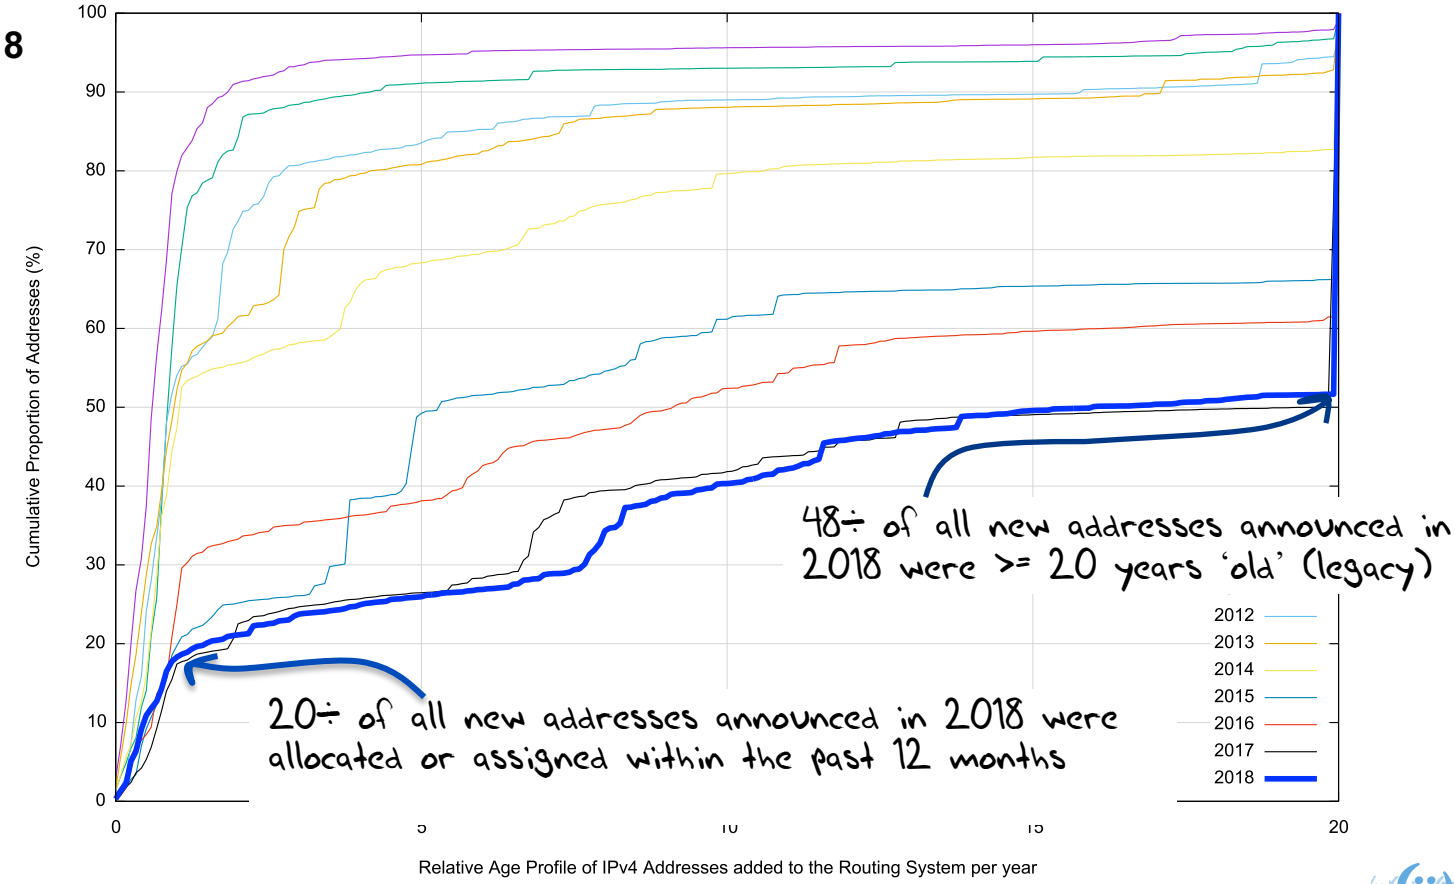

2000 - 2018: IPv4 Advertised vs Unadvertised

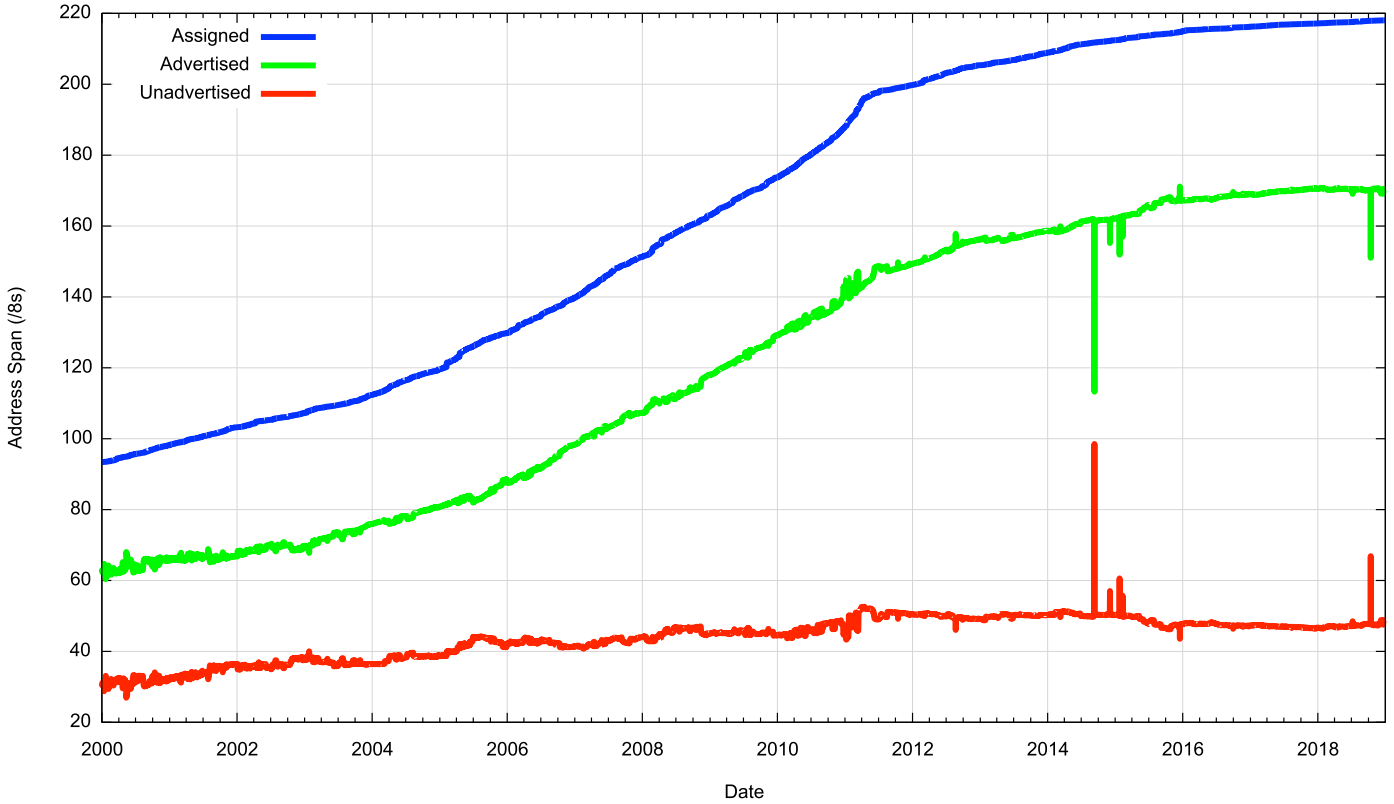

2000 - 2018: Unadvertised Addresses

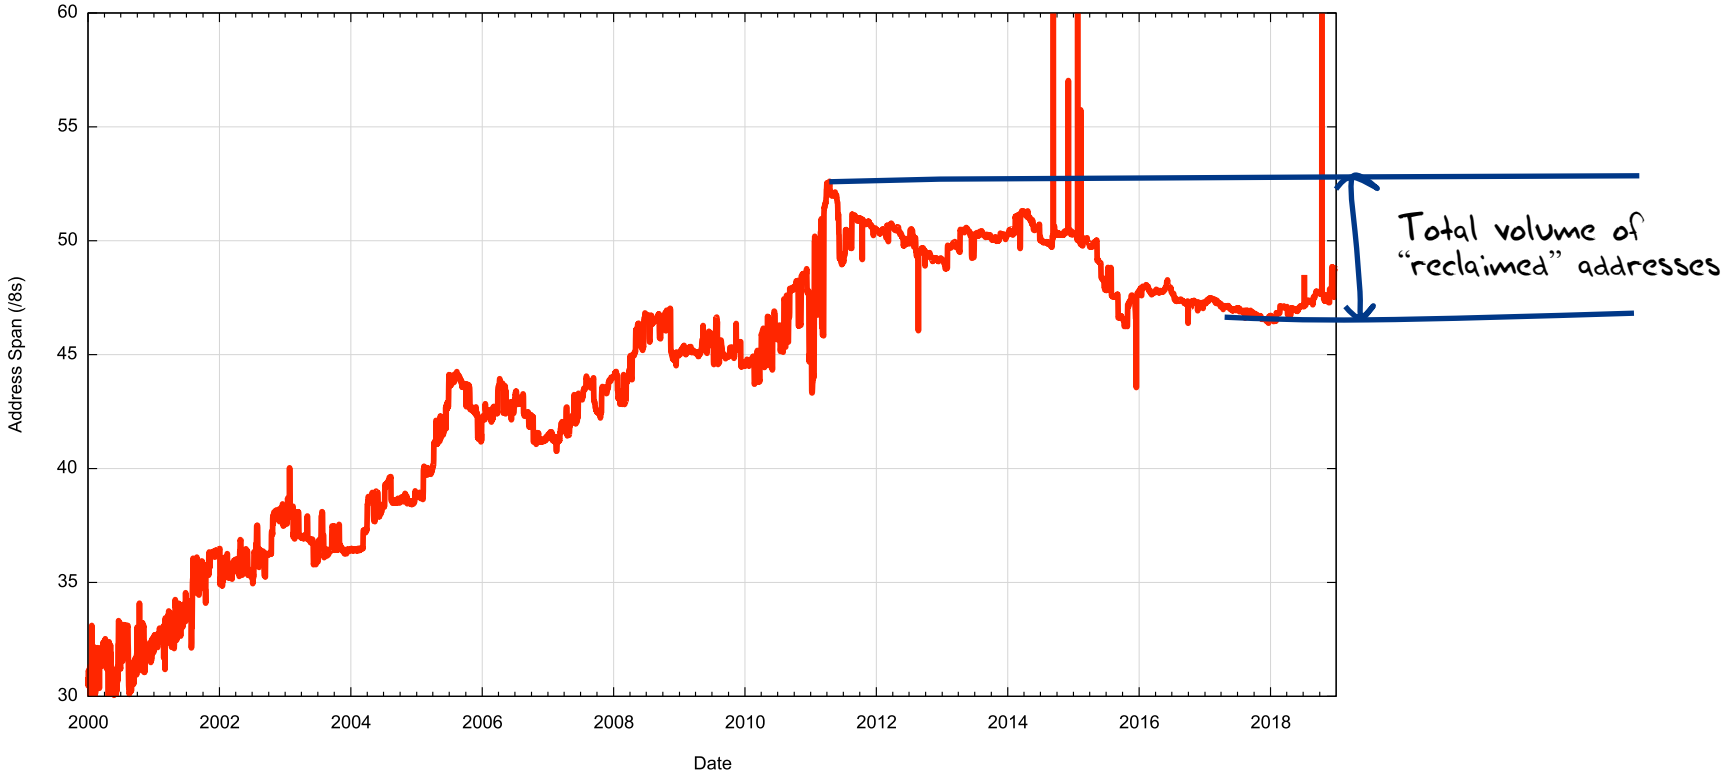

2018: Assigned vs Recovered

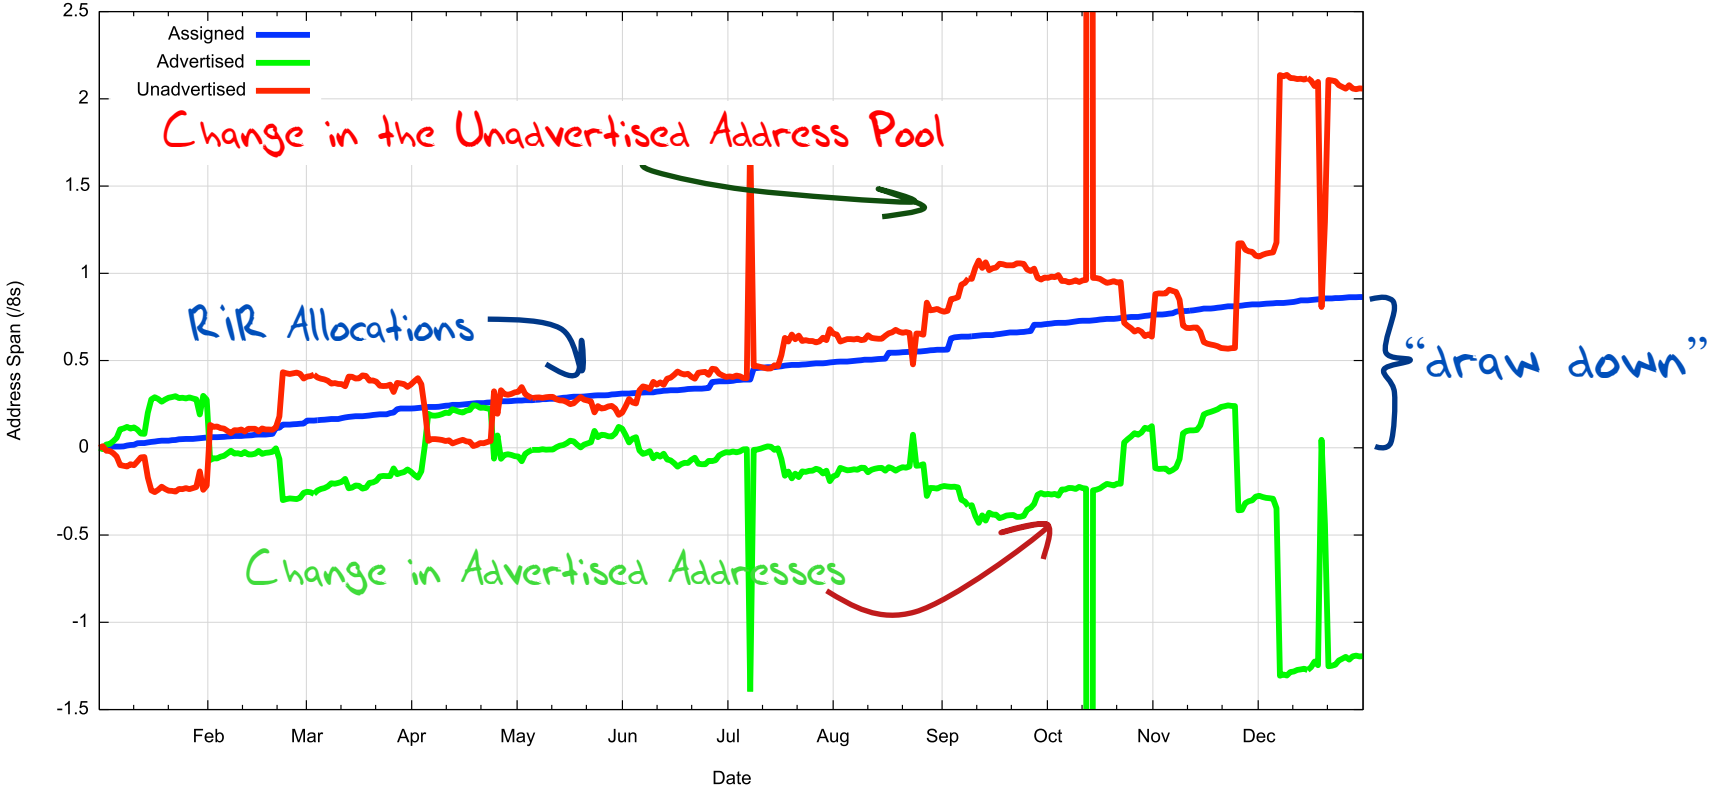

V4 in 2018

- The equivalent of 1.4 /8s were **removed** from the routing table across 2018
- Approximately 0.86 /8s were assigned by RIRs in 2015
  - 0.37 /8's assigned by Afrinic
  - 0.28 /8s assigned by the RIPE NCC (last /8 allocations)
  - 0.10 /8s were assigned by APNIC (last /8 allocations)

- And a net of 2.1 /8's were added to the pool of unadvertised addresses

In 2018 we saw legacy blocks transferring away from ISPs / end user sites and heading towards cloud SPs.

The Route-Views View of IPv6

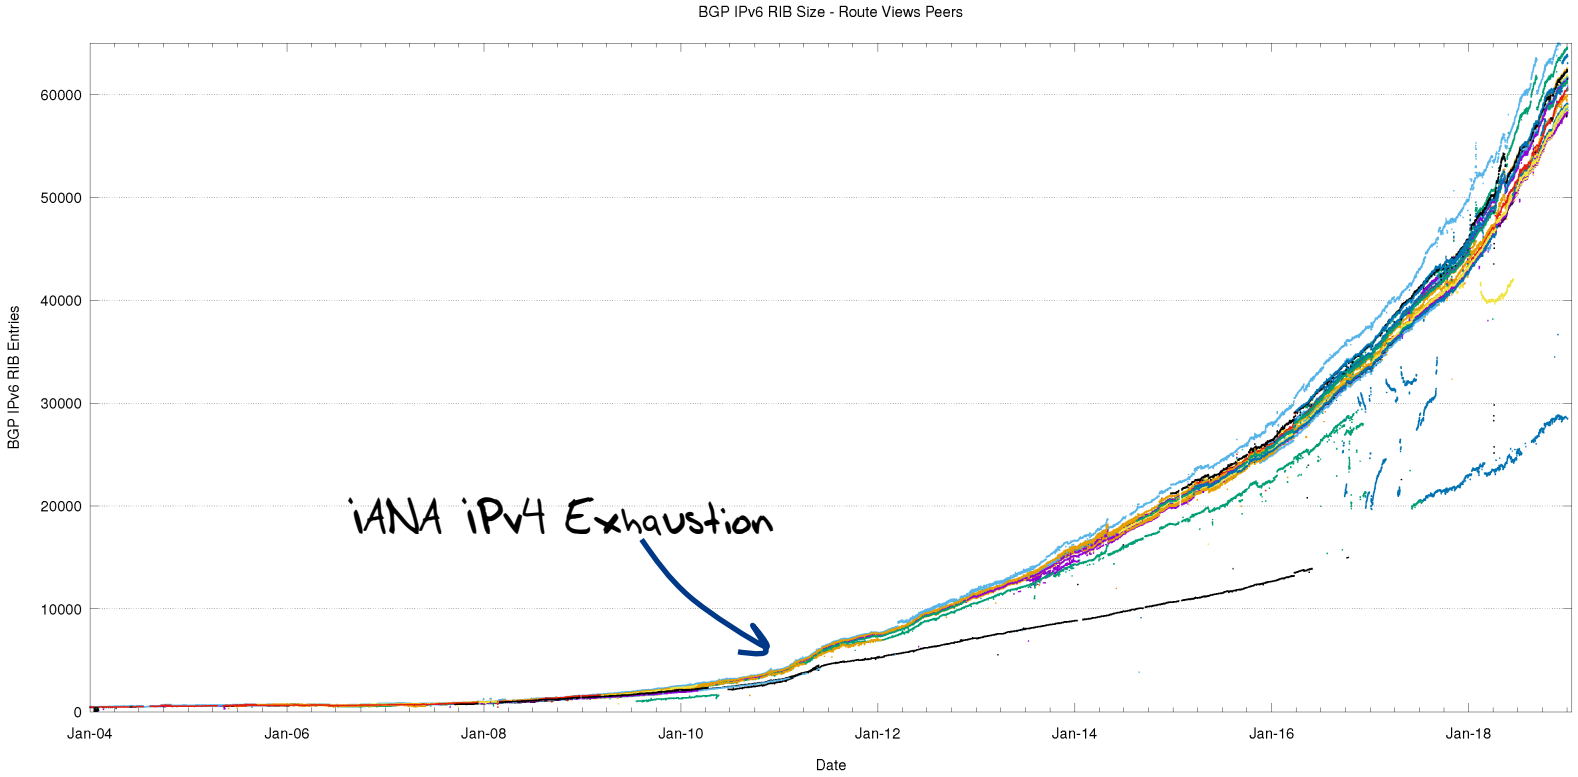

2017-2018 in Detail

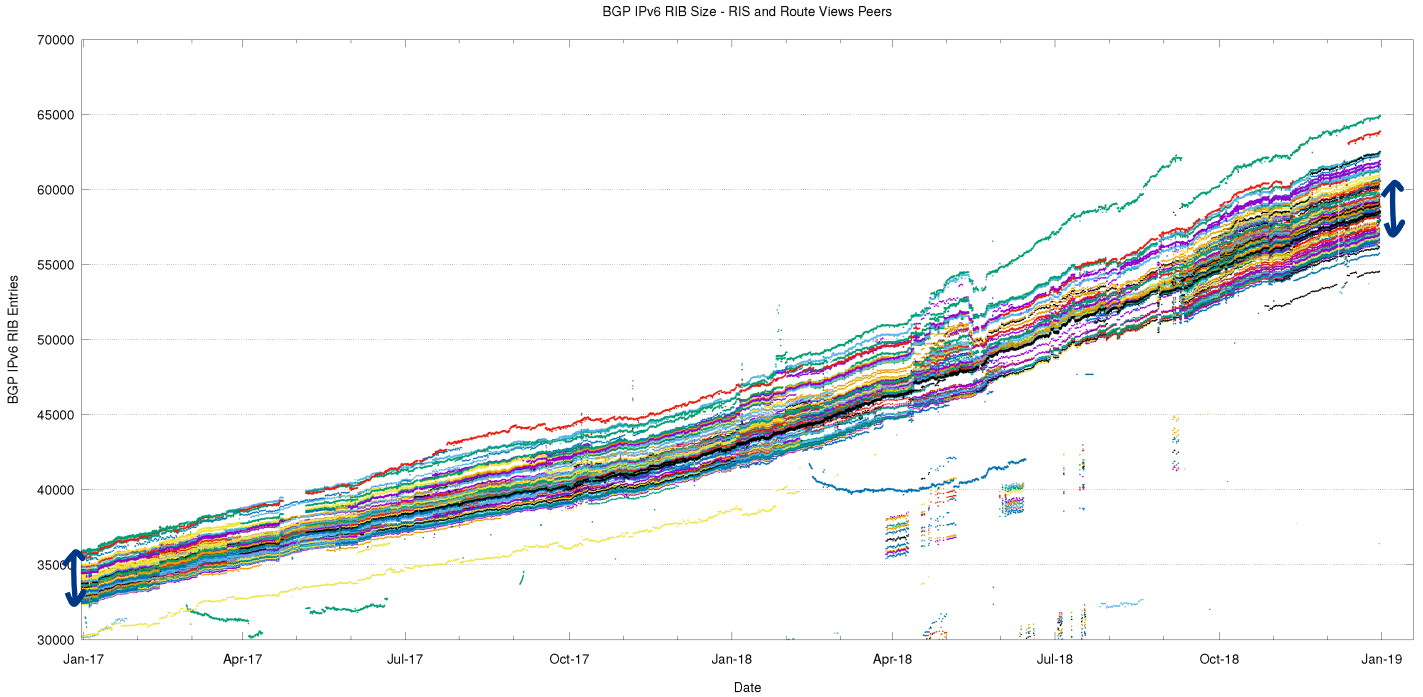

Routing Indicators for IPv6

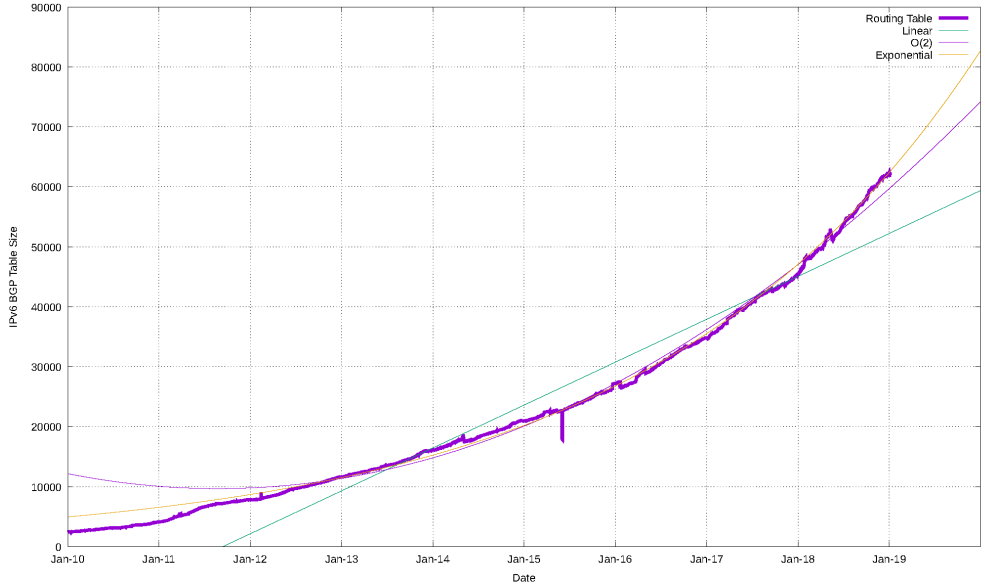

Routing prefixes – growing by some 15,000 prefixes per year

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

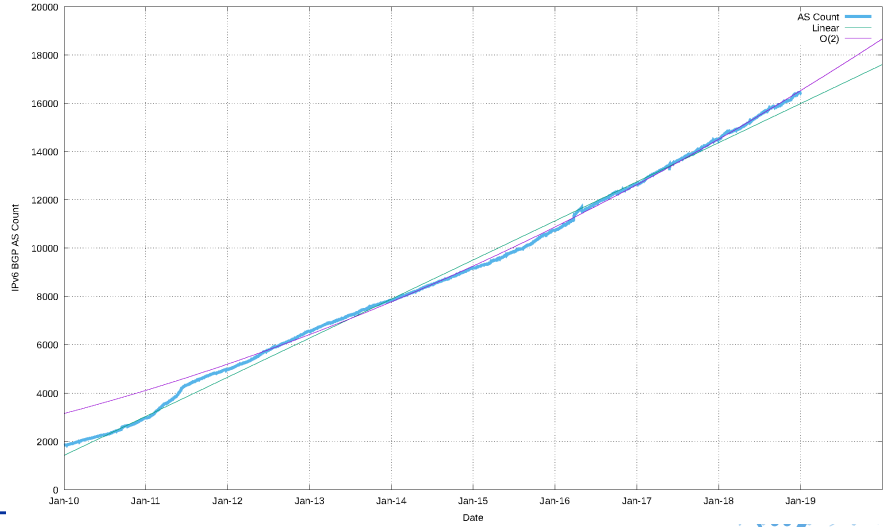

Routing Indicators for IPv6

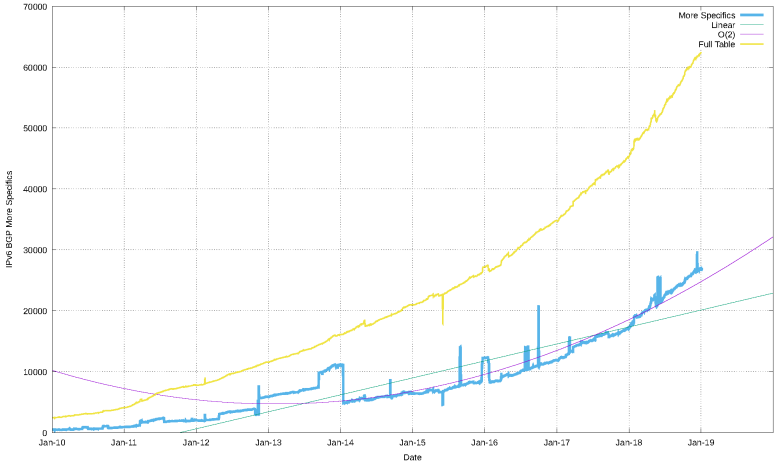

More Specifics now take up more than one third of the routing table

The average size of a routing advertisement is getting smaller

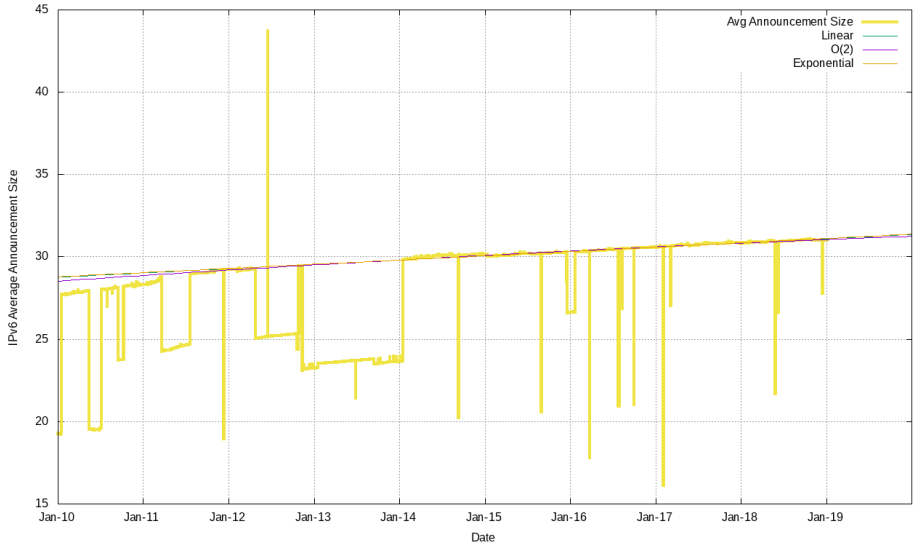

Routing Indicators for IPv6

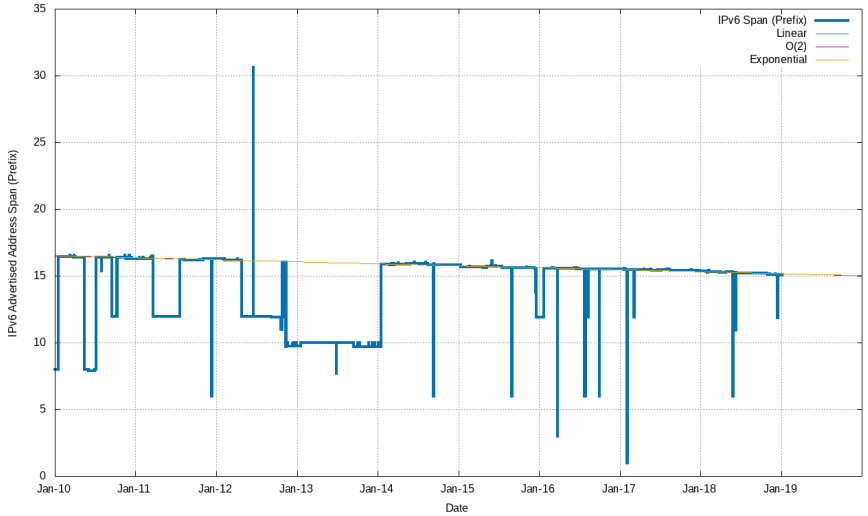

Advertised Address span is growing at an exponential rate

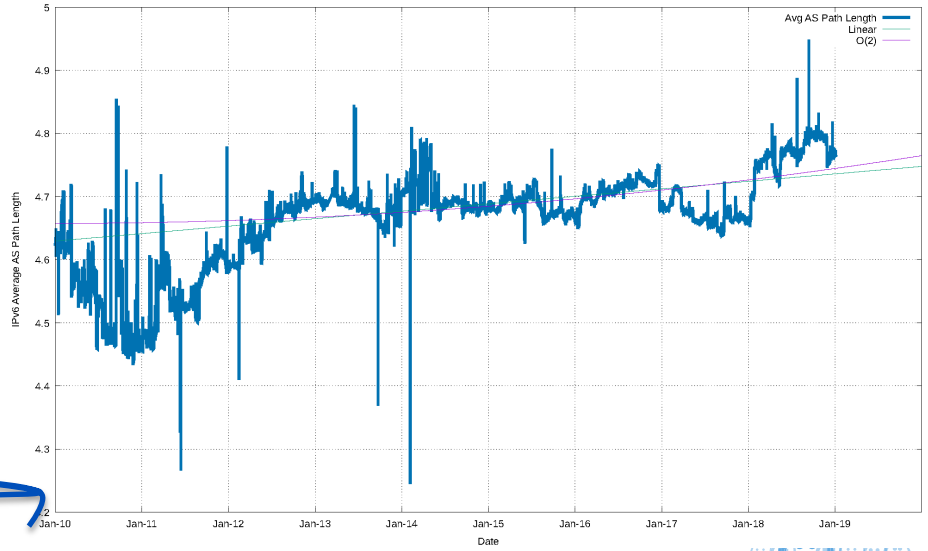

The "shape" of inter-AS interconnection in IPv6 is rising slightly. Local connections appear to be replacing overlay trunk transits

AS Adjacencies (AS131072)

13,095 out of 16,465 ASNs have 1 or 2 AS Adjacencies (79%)
573 ASNs have 10 or more adjacencies
2 ASNs have >1,000 adjacencies

4,295	AS6939	HURRICANE - Hurricane Electric, Inc., US
1,049	AS3356	LEVEL3 - Level 3 Communications, Inc., US
749	AS174	COGENT-174 - Cogent Communications, US
719	AS2915	NTT America, US
632	AS1299	Telia Carrier, SE

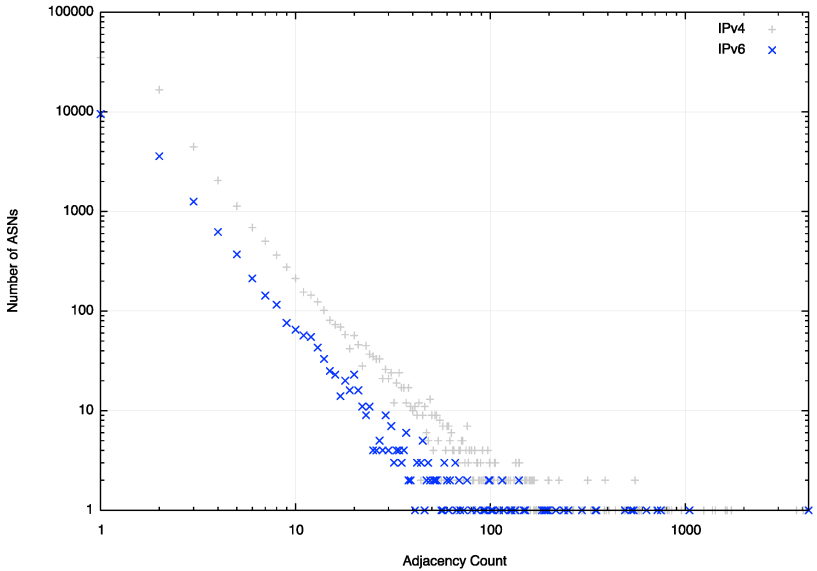

V6 in 2018

- Overall IPv6 Internet growth in terms of BGP is still increasing, and is currently at some **15,000 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

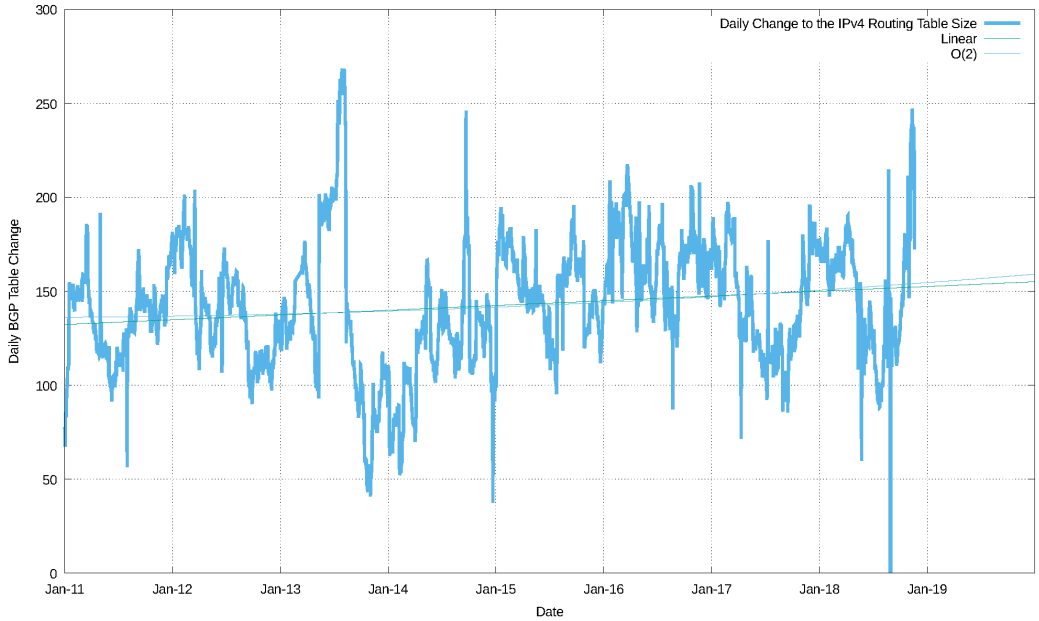

Growth in the V4 network appears to be constant at a long term average of 140 additional routes per day, or some 52,000 additional routes per year

V4 BGP Table Size Predictions

Jan *2017*	*646,000*
2018	*699,000*
2019	***755,000***
2020	*807,000*
2021	*859,000*
2022	*911,000*
2023	*963,000*
2024	*1,015,000*

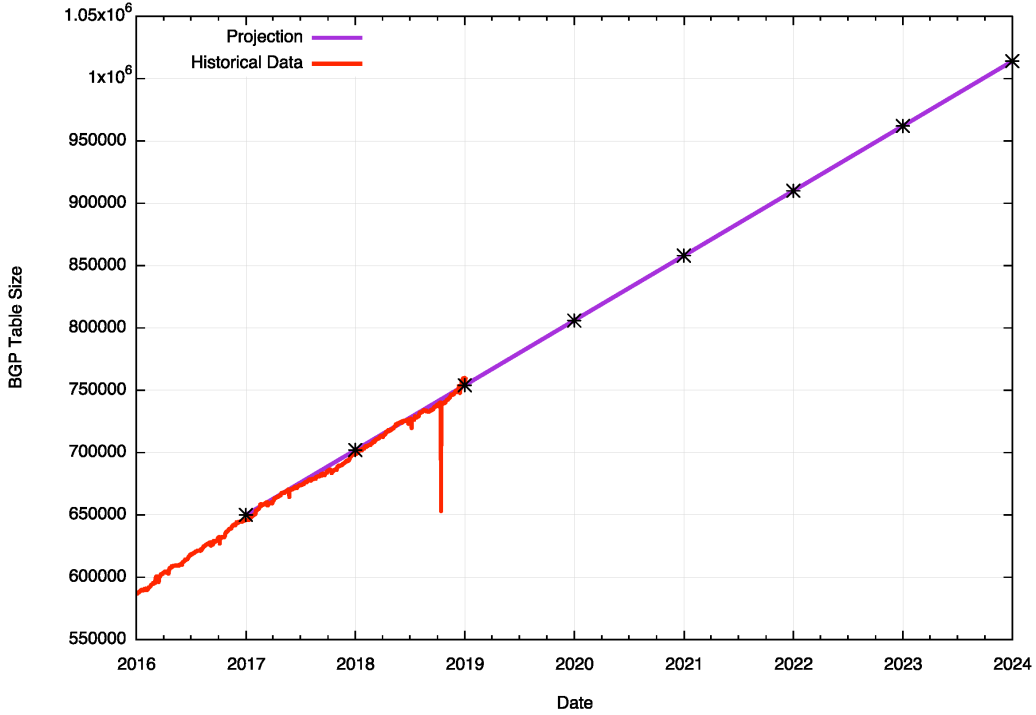

V6 - Daily Growth Rates

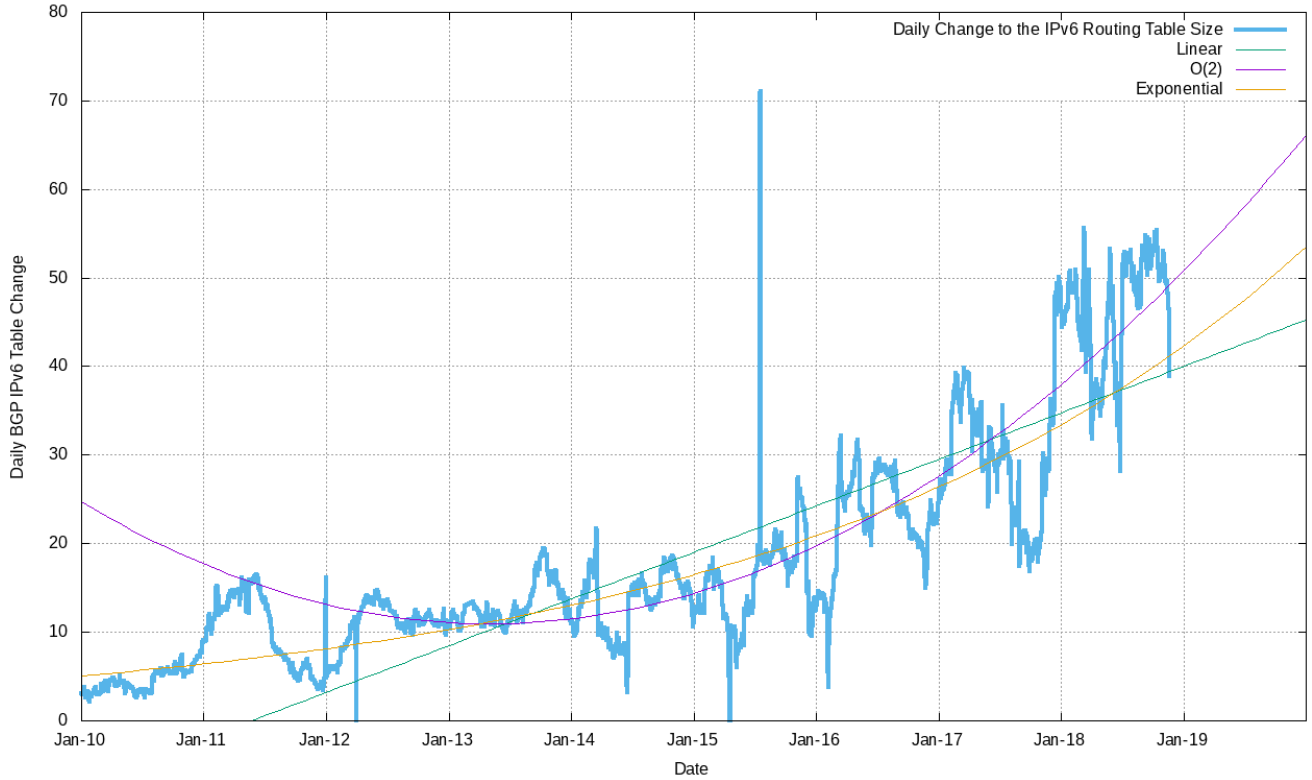

V6 BGP Table Size Predictions

	Linear	Exponential
Jan *2017*	*35,000*	*36,000*
2018	*49,000*	*47,000*
2019	***62,000***	***62,000***
2020	*75,000*	*83,000*
2021	*89,000*	*109,000*
2022	*102,000*	*145,000*
2023	*116,000*	*192,000*
2024	*130,000*	*255,000*

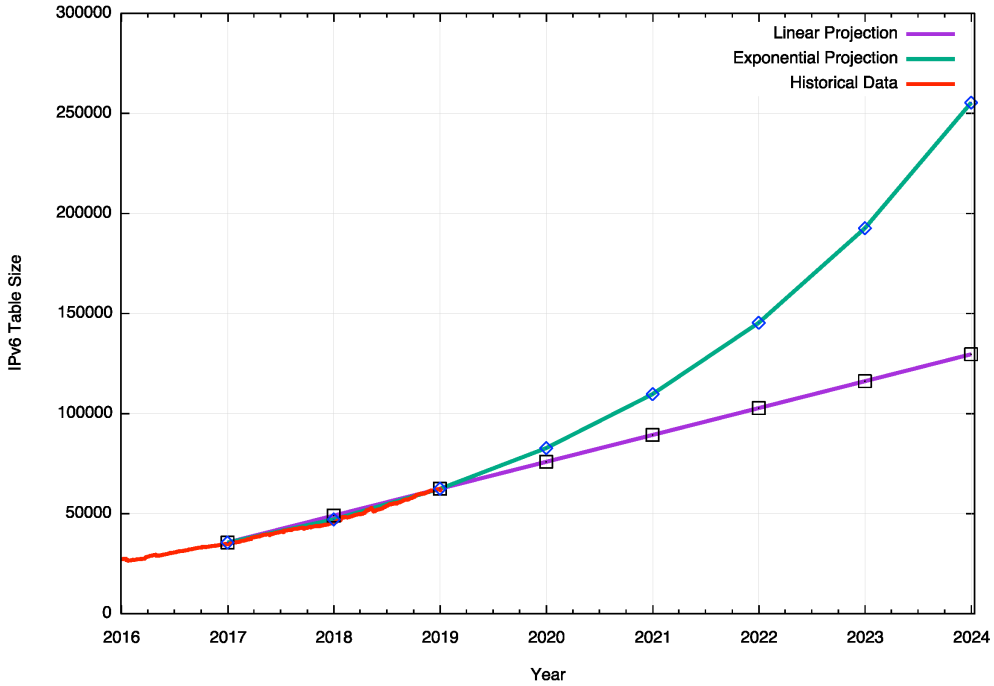

BGP Table Growth

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 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

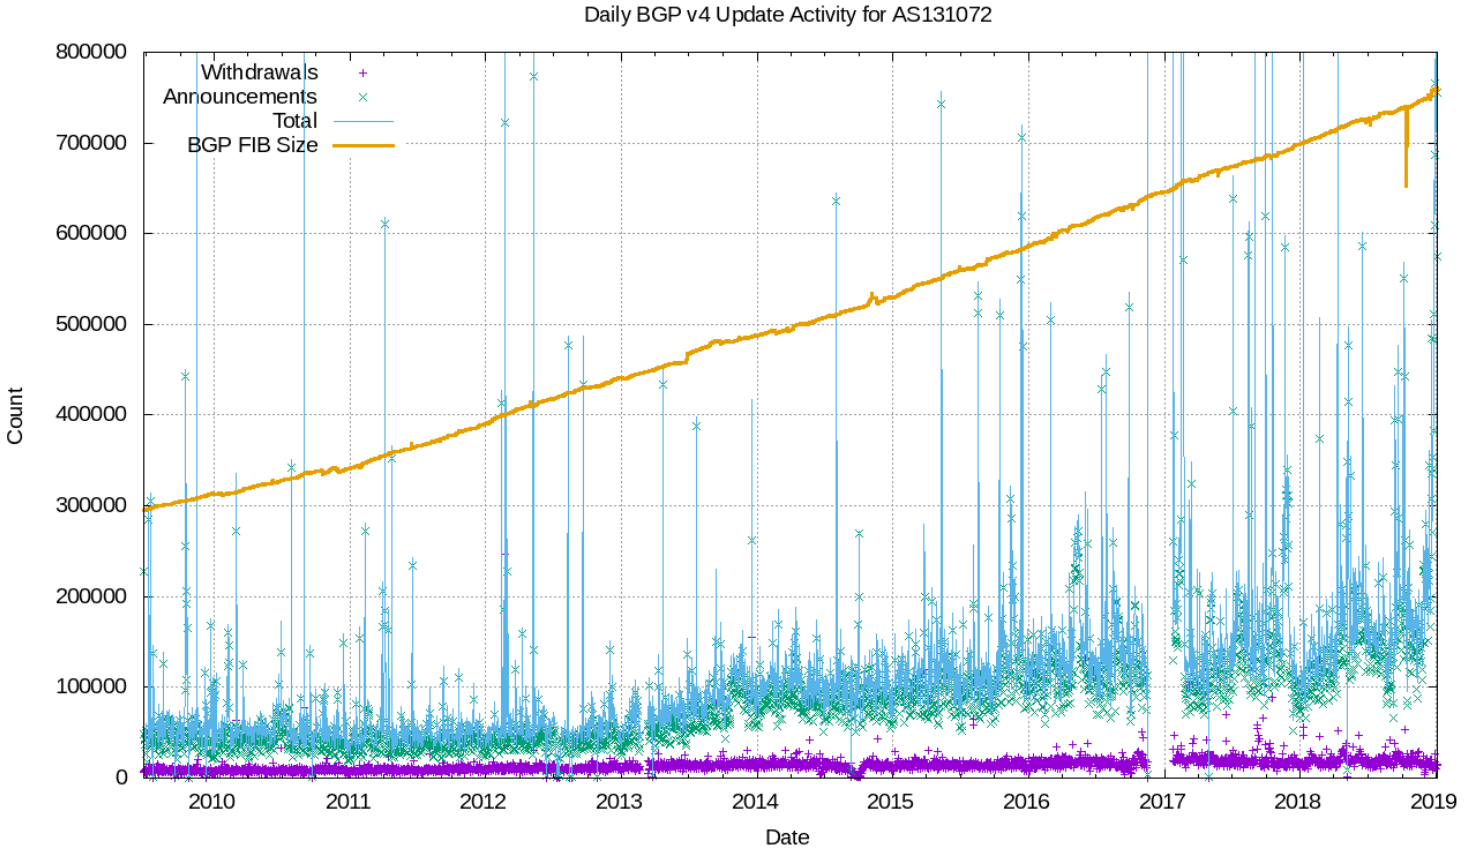

IPv4 BGP Convergence Performance

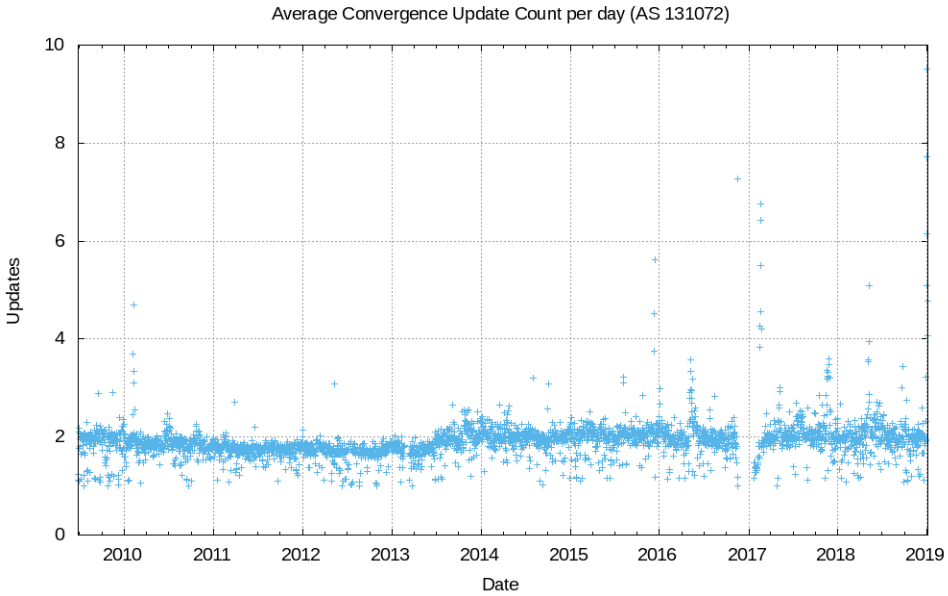

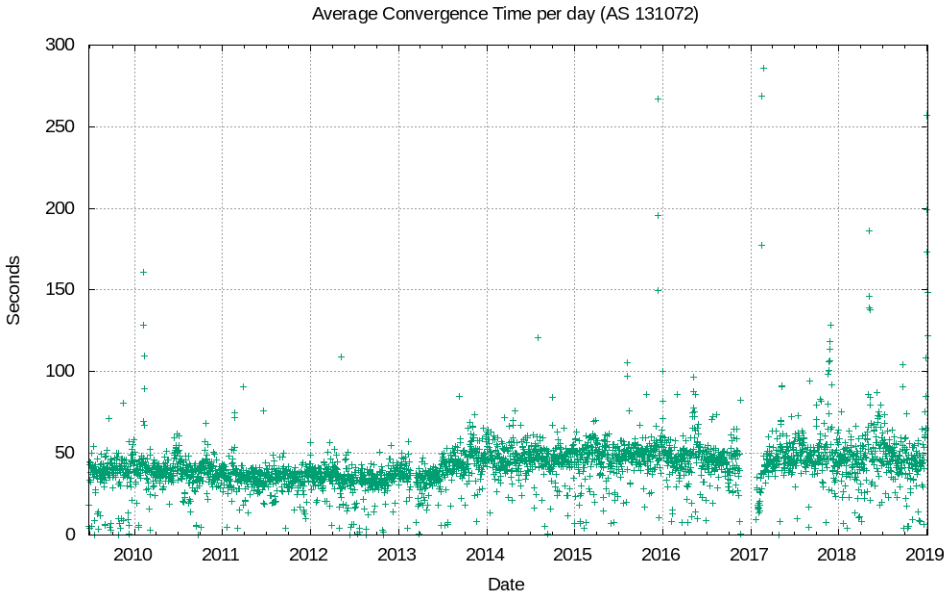

Updates in IPv4 BGP

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

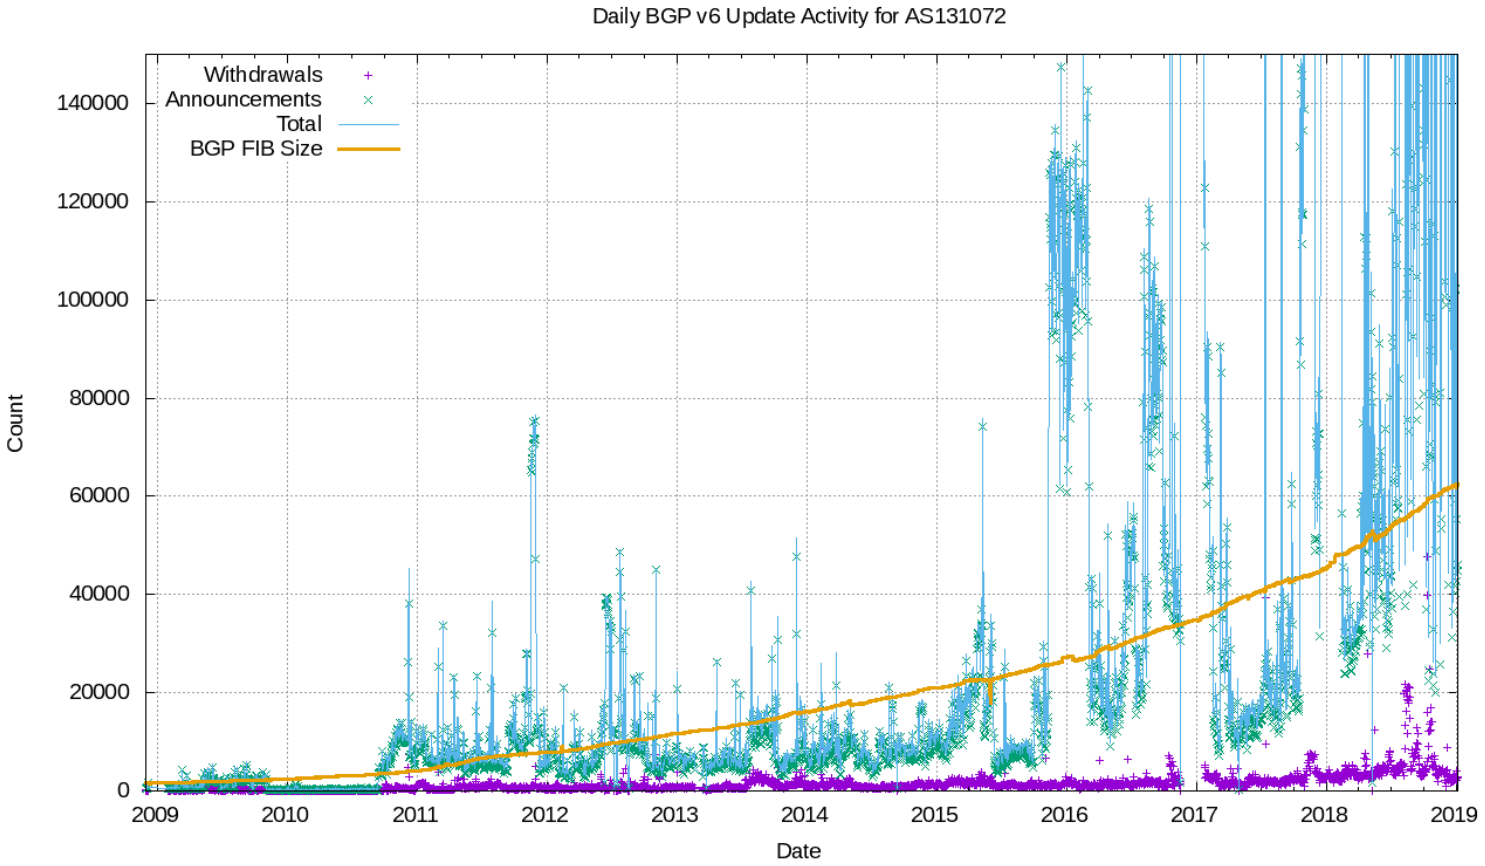

V6 Convergence Performance

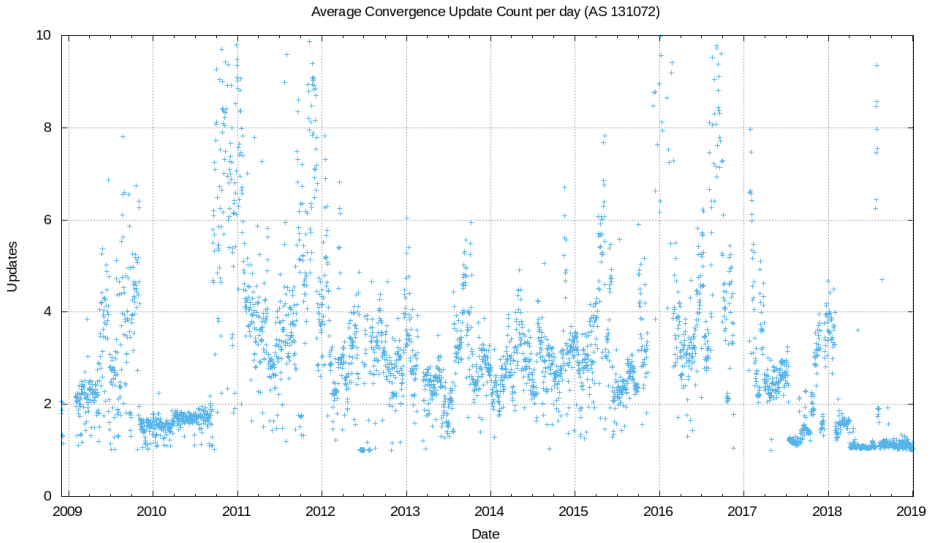

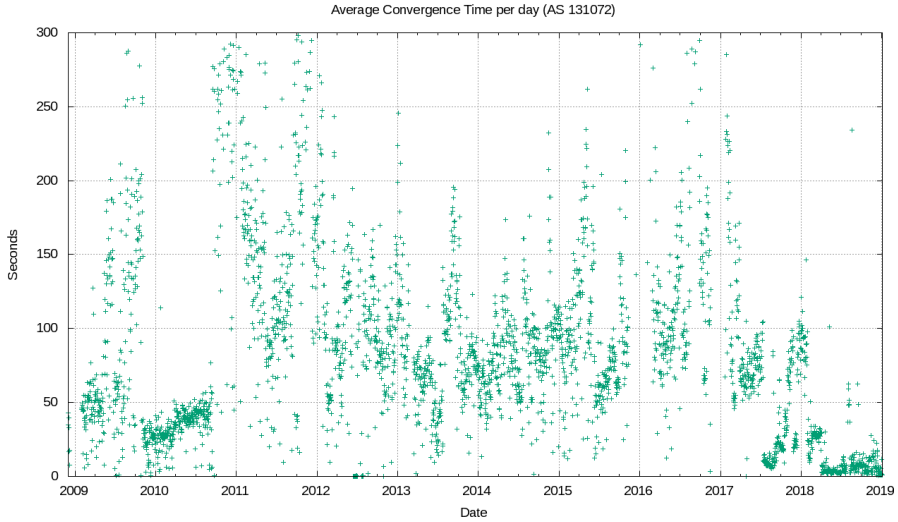

Routing Futures

- 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 speed 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 900,000 IPv4 slots and 110,000 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

That's it!