Routing in 2016

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…

Routing the Internet

This is a view pulled together from each of the routing peers of Route-Views

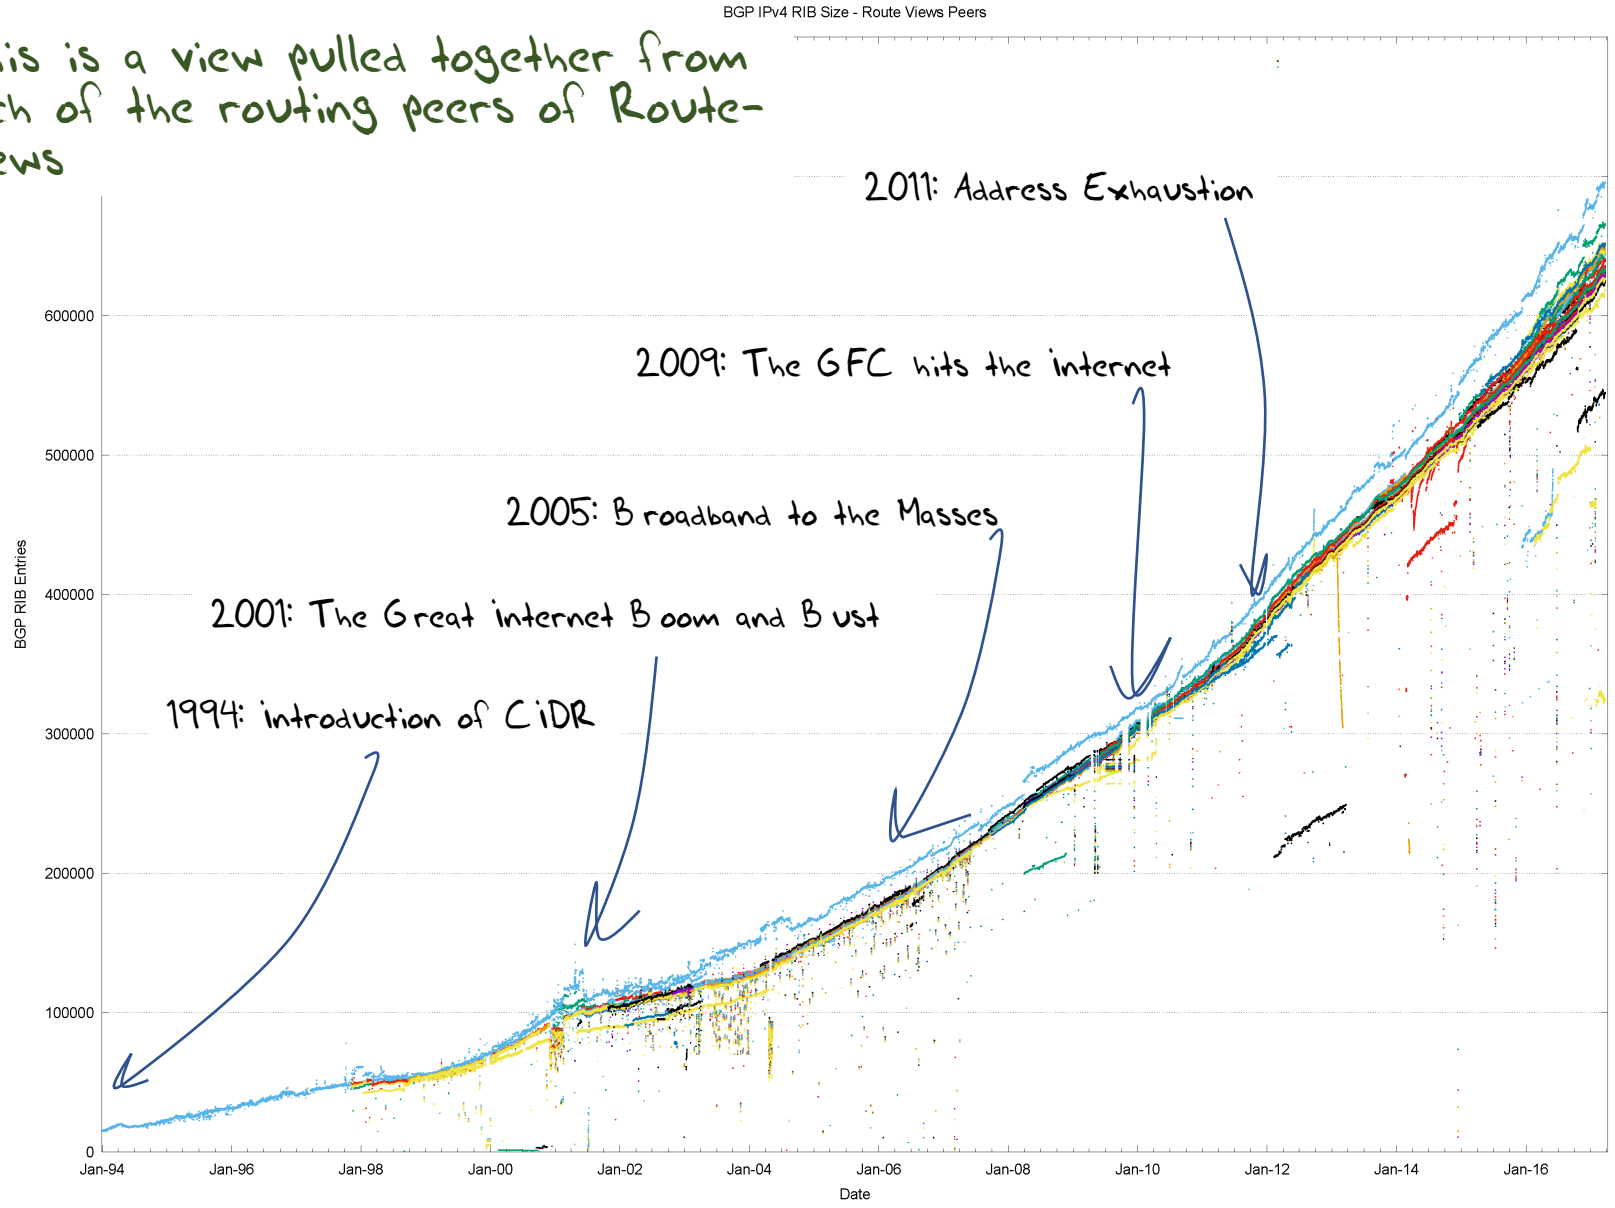

2015-2016 in detail

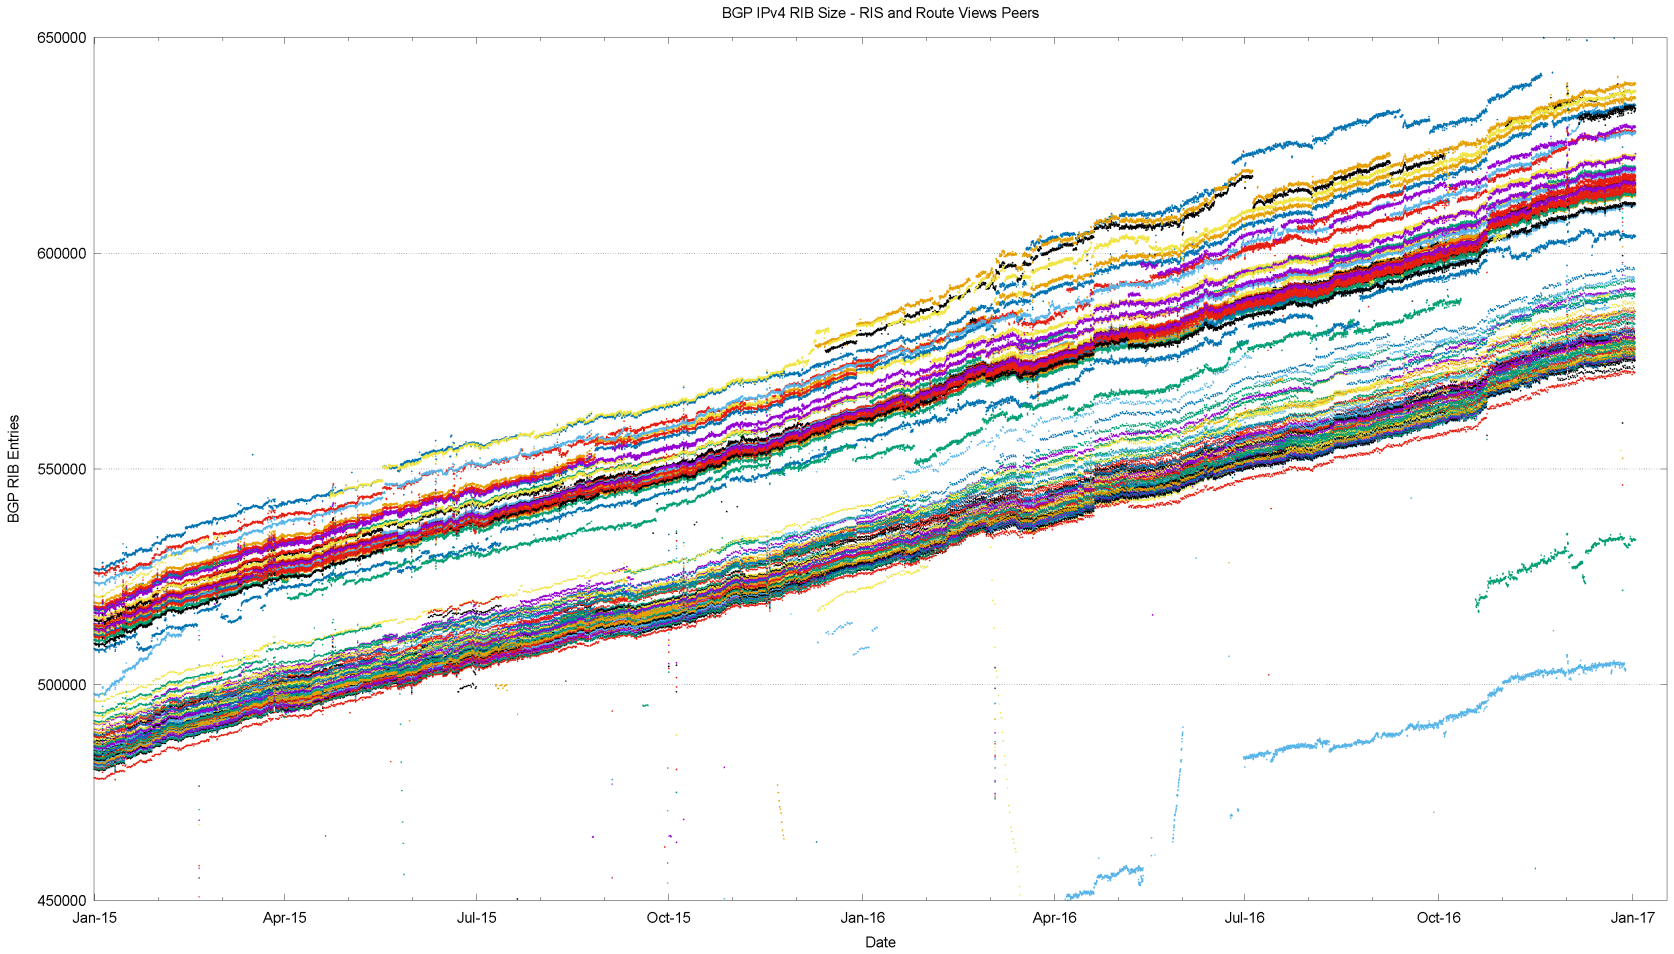

2015-2016 in detail

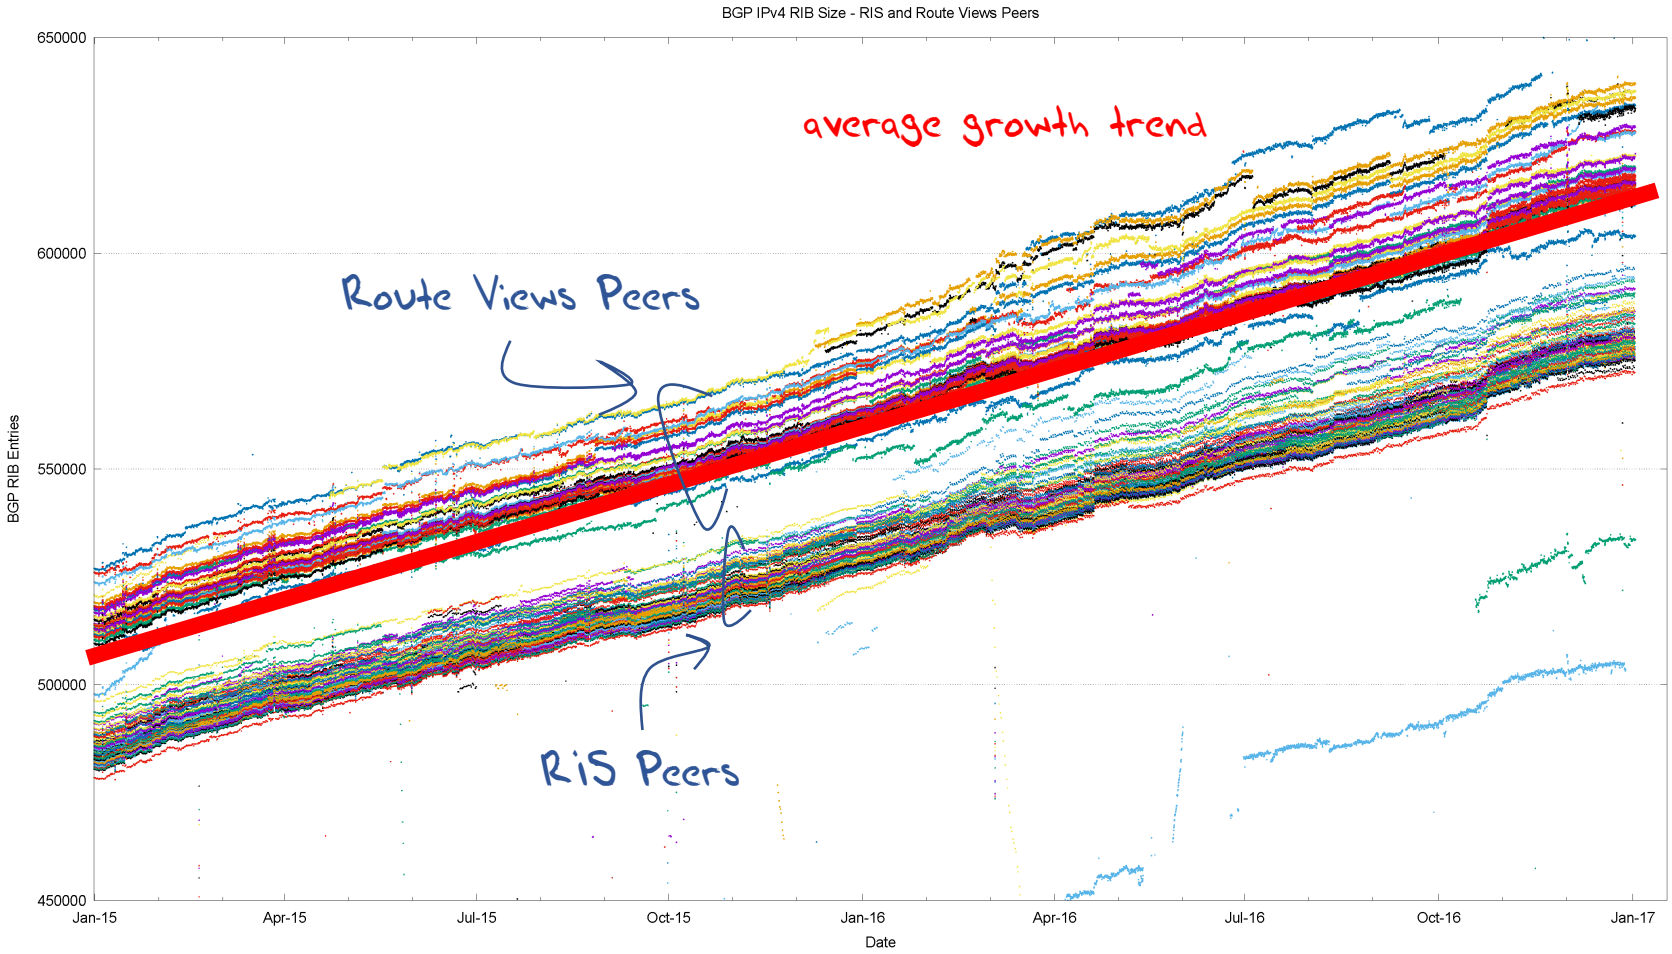

Routing Indicators for IPv4

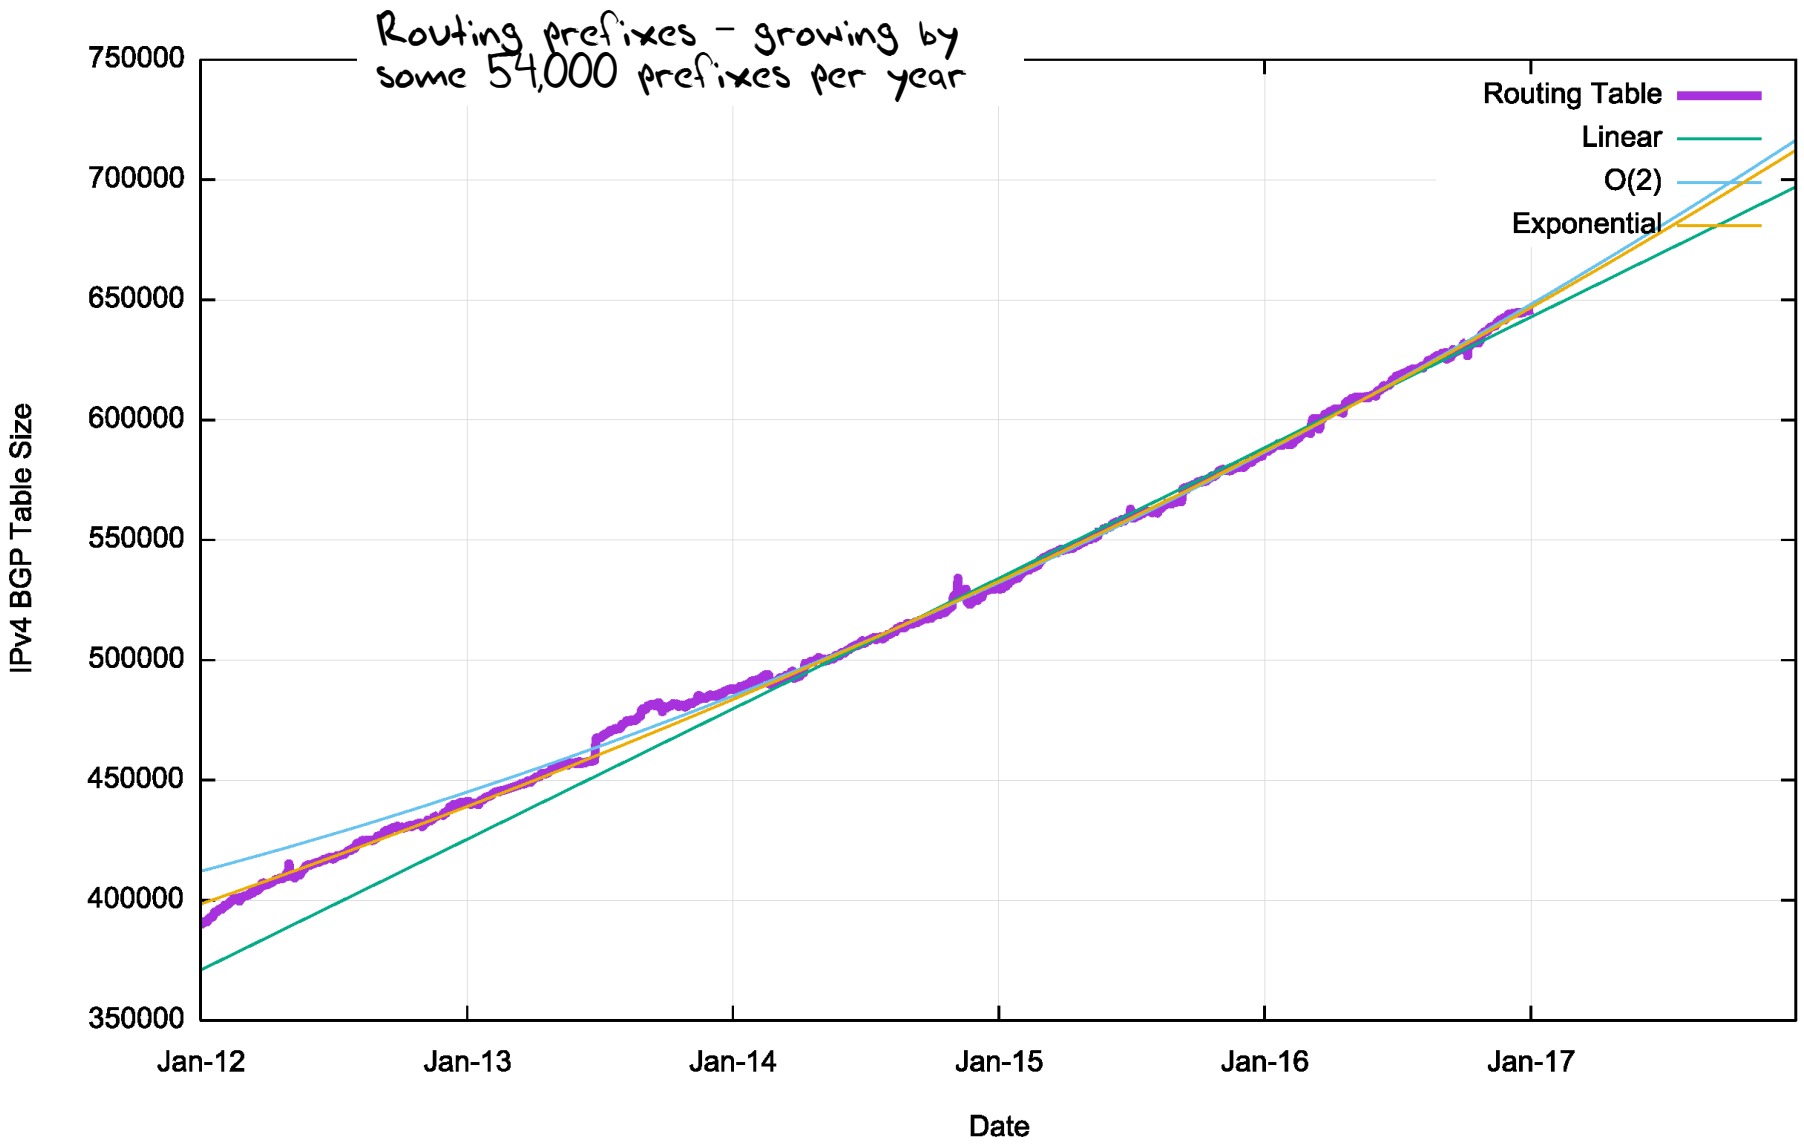

Routing Indicators for IPv4

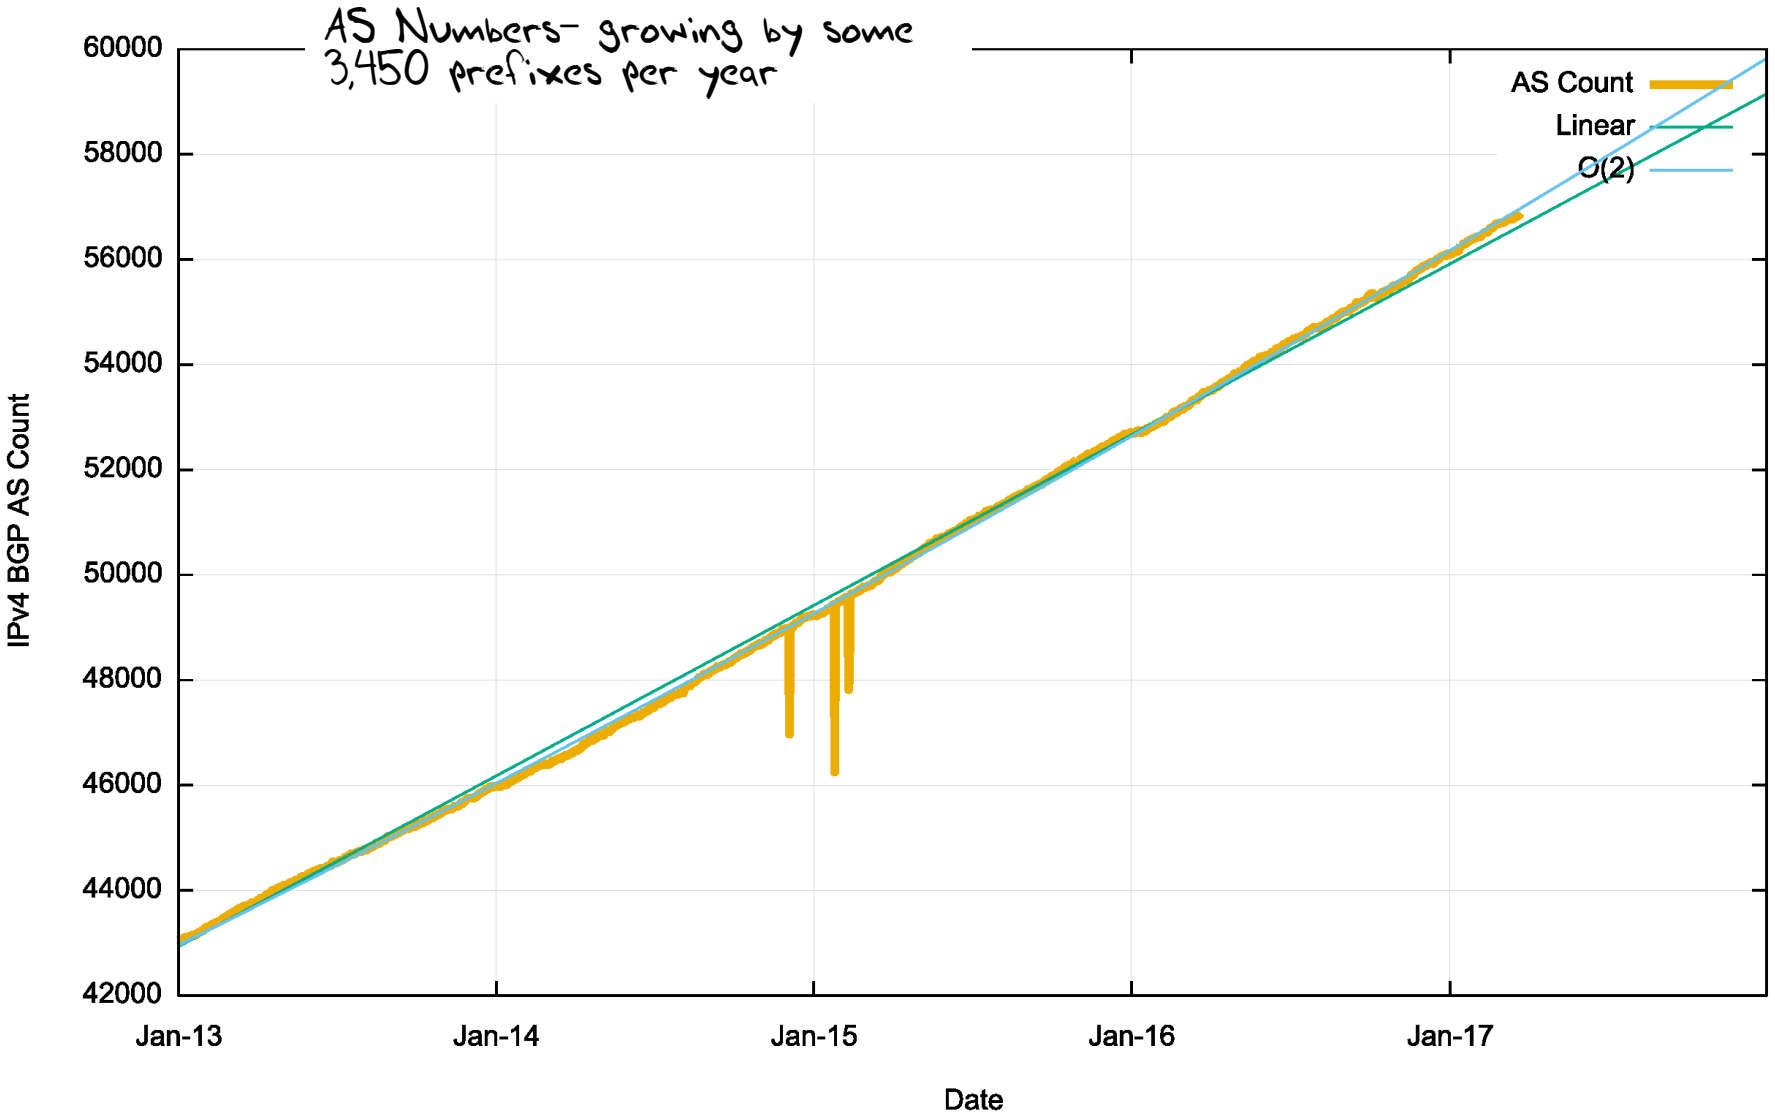

Routing Indicators for IPv4

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

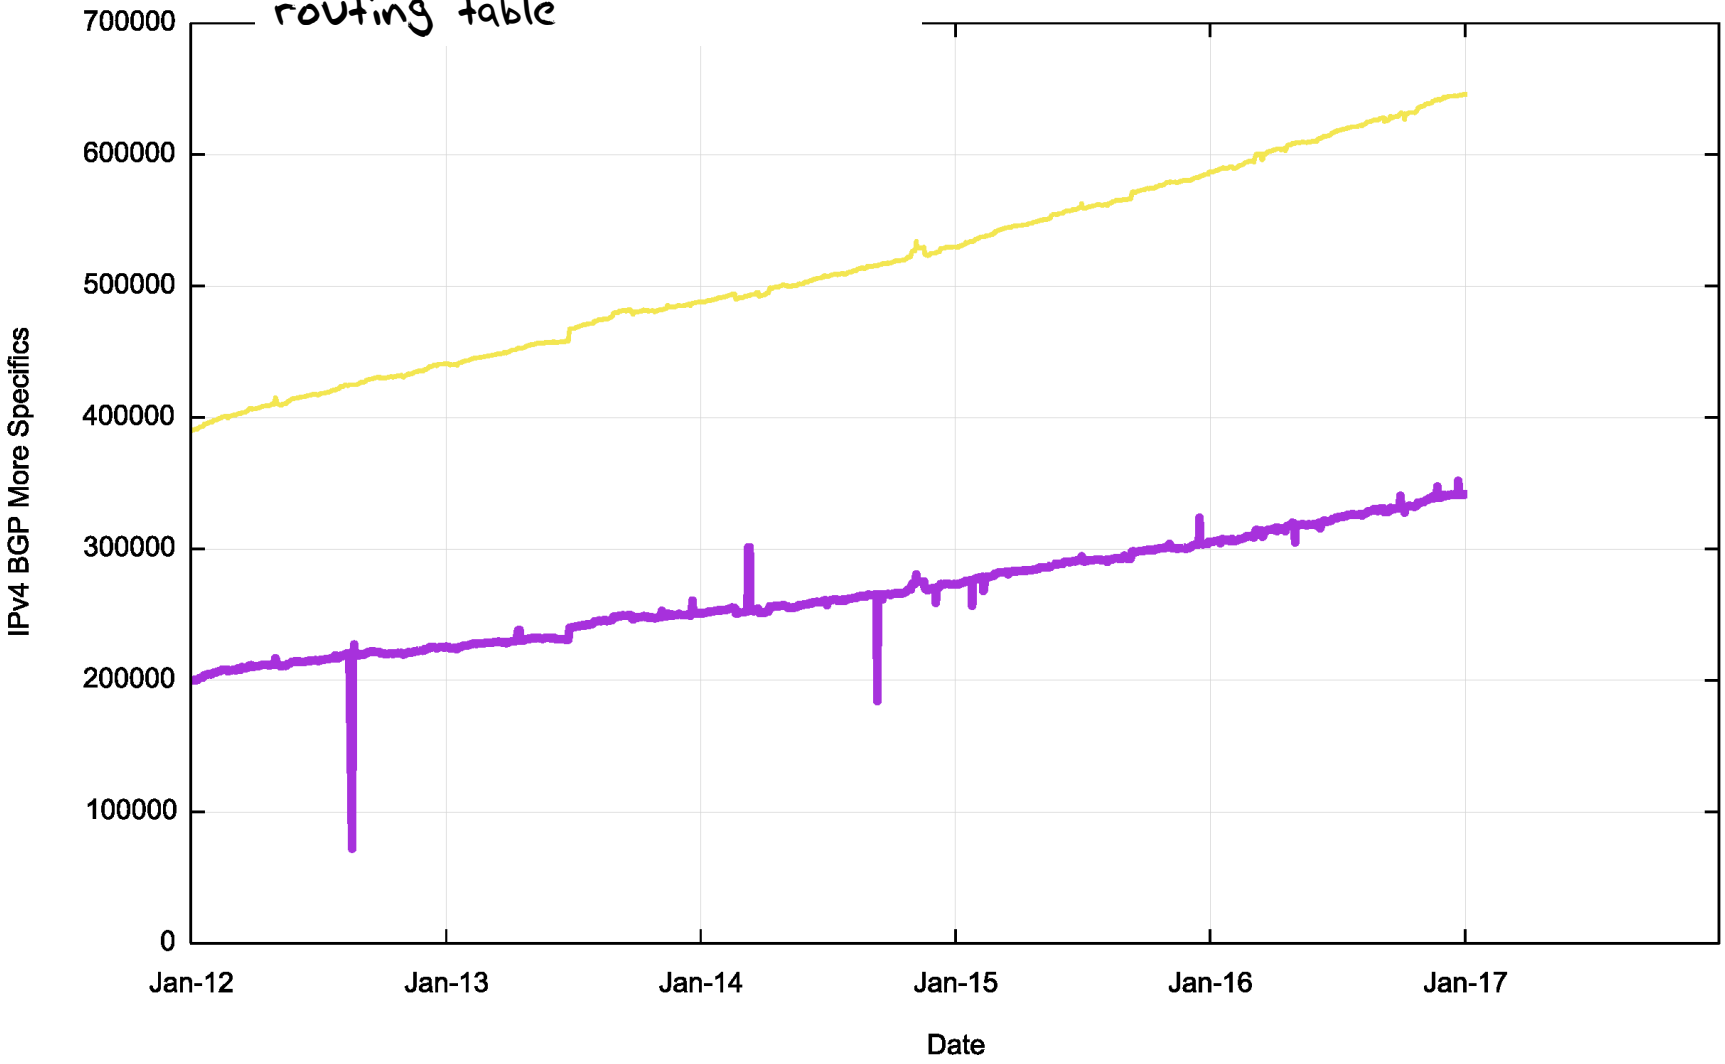

Routing Indicators for IPv4

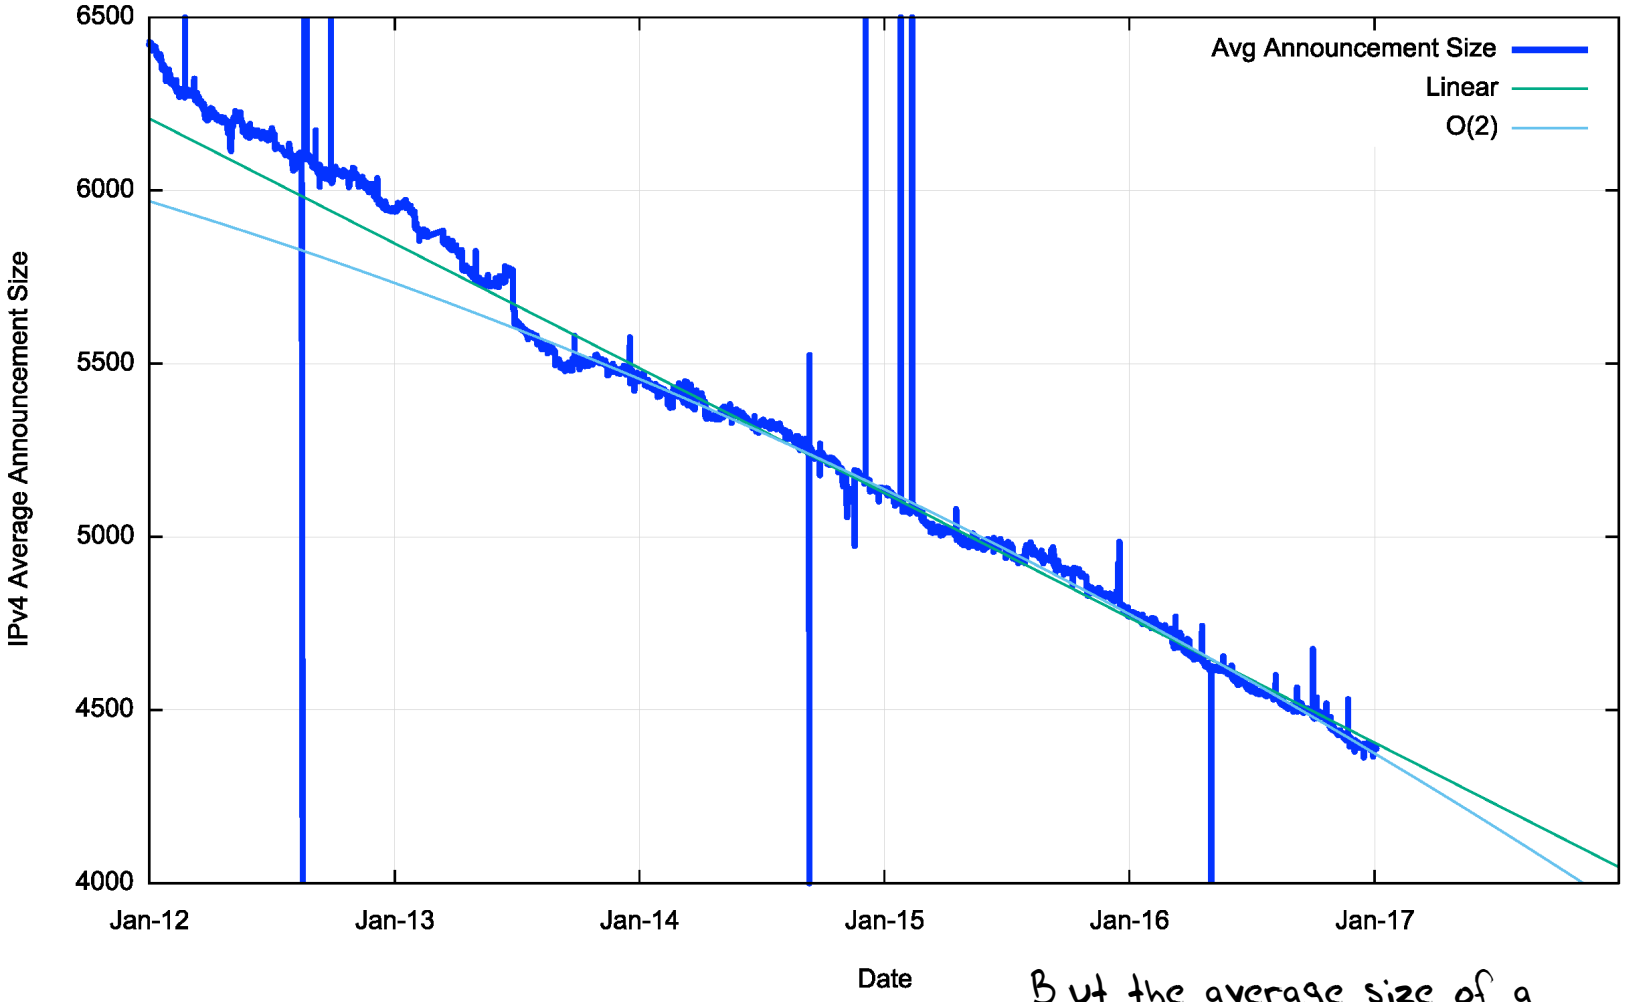

But the average size of a routing advertisement is getting smaller

Routing Indicators for IPv4

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

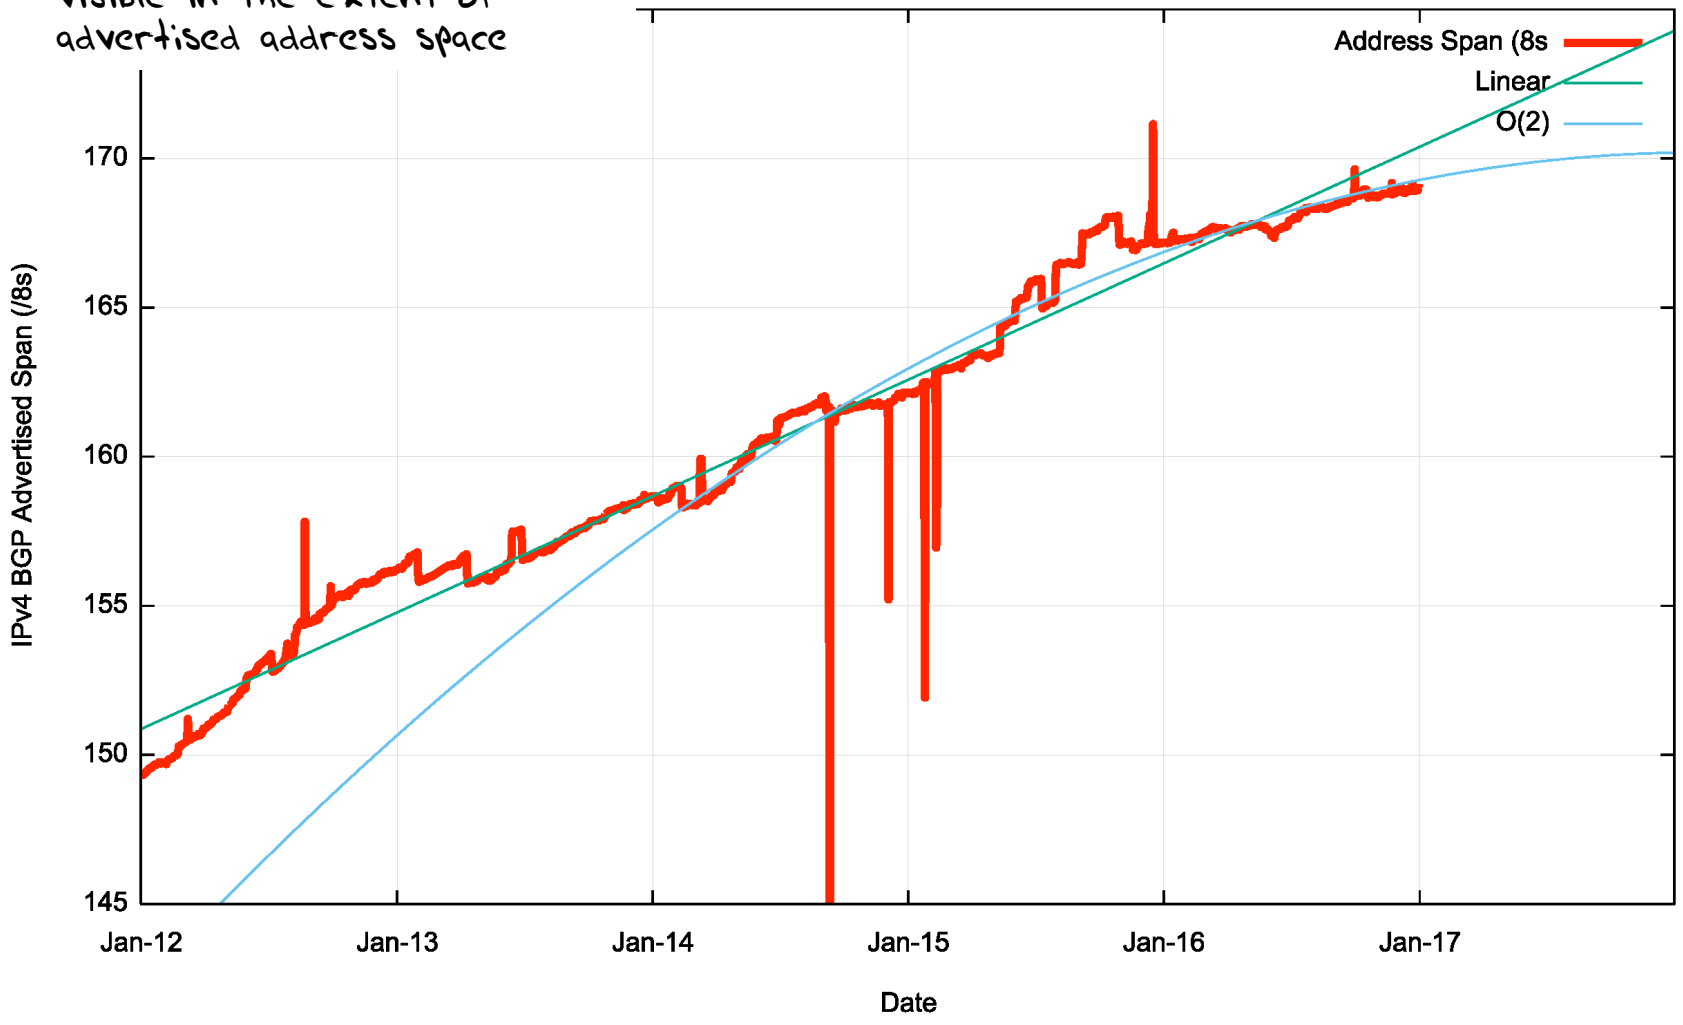

Routing Indicators for IPv4

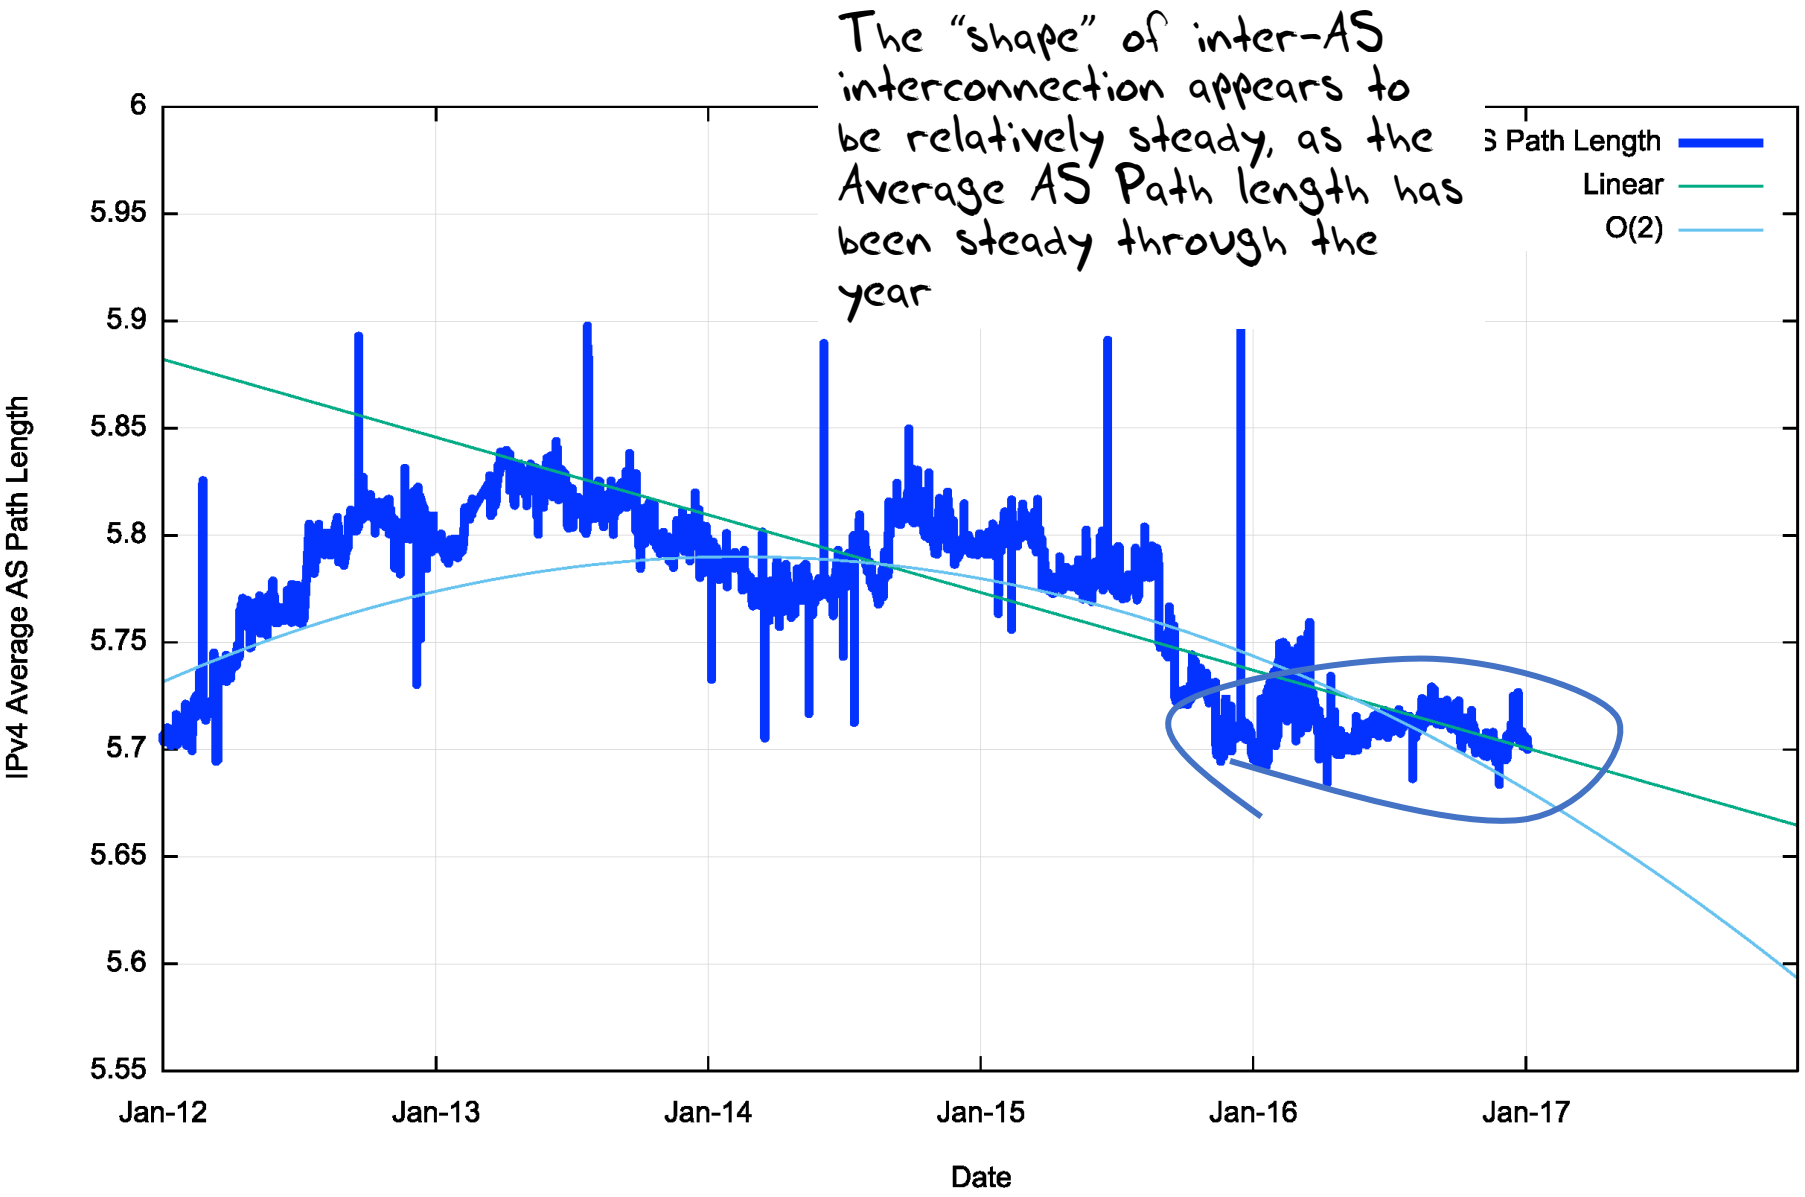

AS Adjacencies (Route-Views)

19,700 out of 57,064 ASNs have 1 or 2 AS Adjacencies (72%)

3,062 ASNs have 10 or more adjacencies

22 ASNs have >1,000 adjacencies

6,202	AS6939	HURRICANE - Hurricane Electric, Inc.,
5,069	AS174	COGENT-174 - Cogent Communication
4,767	AS3356	LEVEL3 - Level 3 Communications, In
2,632	AS3549	LVLT-3549 - Level 3 Communications,
2,397	AS7018	ATT-INTERNET4 - AT&T Services, Inc.,
1,959	AS209	CENTURYLINK-US-LEGACY-QWEST - Qw
1,953	AS57463	NETIX , BG
1,691	AS37100	SEACOM-AS, MU
1,620	AS34224	NETERRA-AS, BG

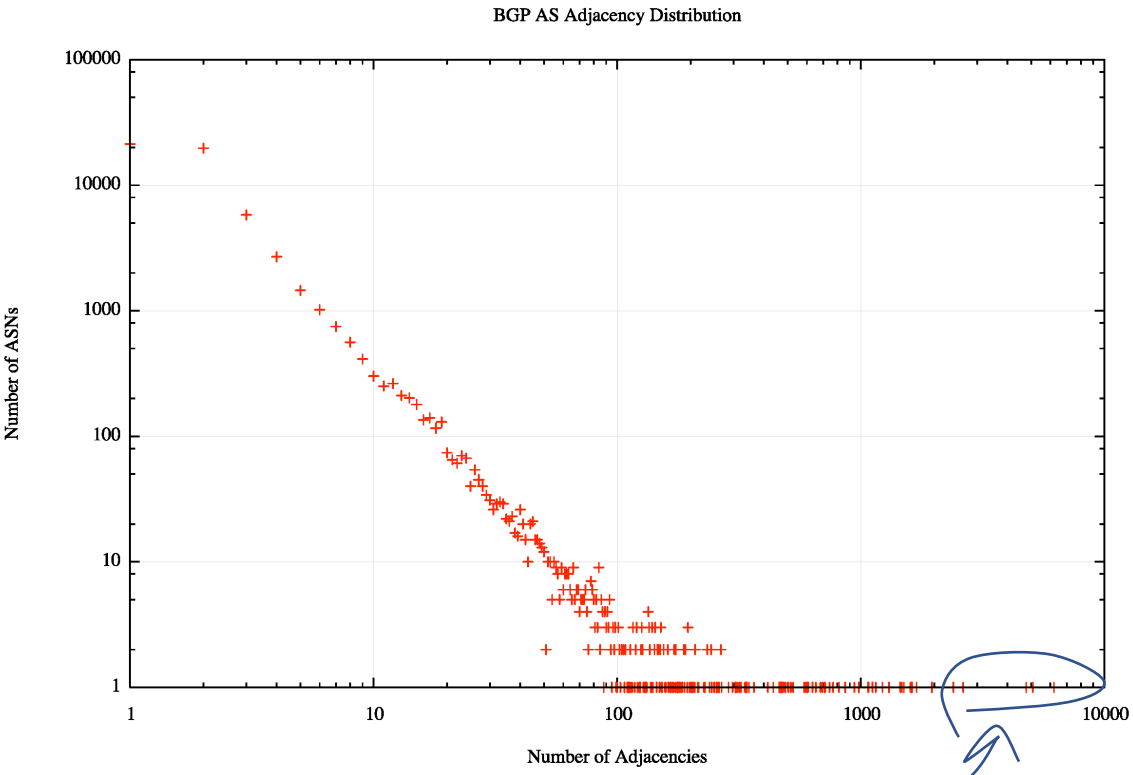

What happened in 2016 in V4?

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 is now heading to 700,000 by the end of 2017
- The pace of growth of the routing table is still relatively constant at ~54,000 new entries and 3,400 new AS's per year
    - IPv4 address exhaustion is not changing this!
    - Instead, we are advertising shorter prefixes into the routing system

How can the IPv4 network continue to grow when we are running out of IPv4 addresses?

We are now recycling old addresses back into the routing system

Some of these addresses are transferred in ways that are recorded in the registry system, while others are being “leased” without any clear registration entry that describes the lessee

IPv4 Advertised Address "Age"

2010

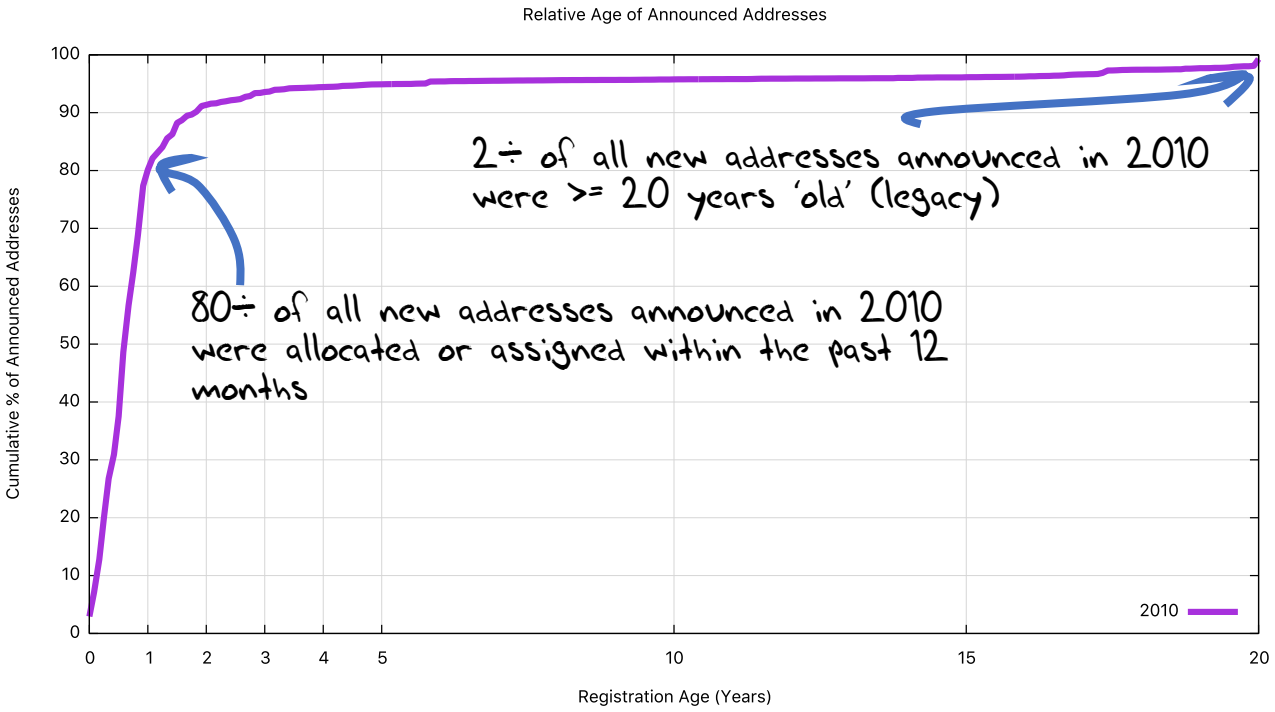

IPv4 Advertised Address "Age"

2016

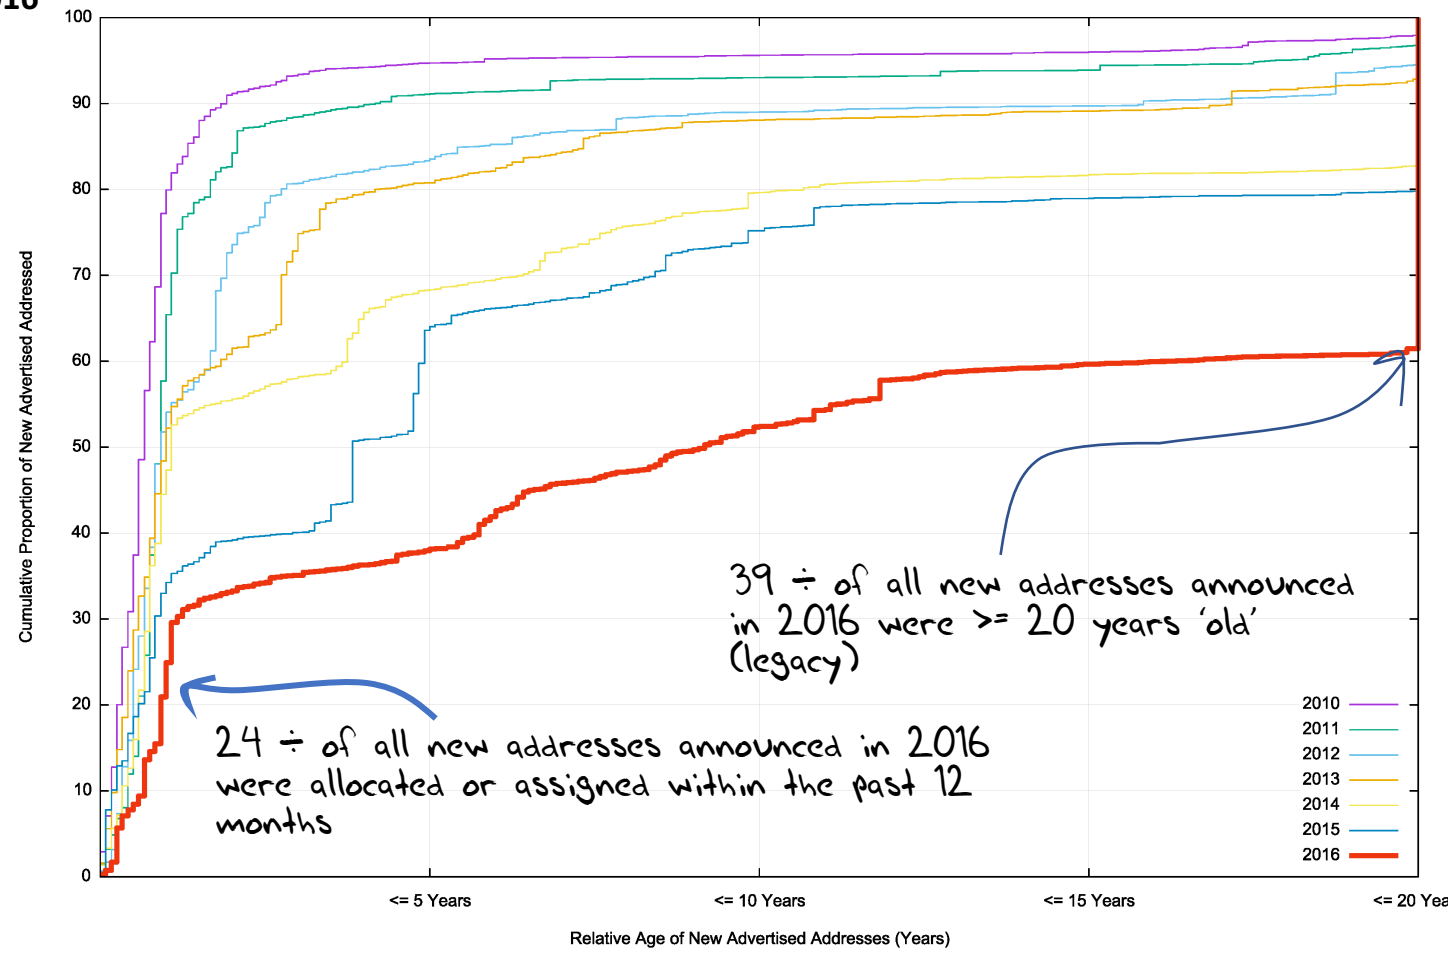

IPv4: Advertised vs Unadvertised Addresses

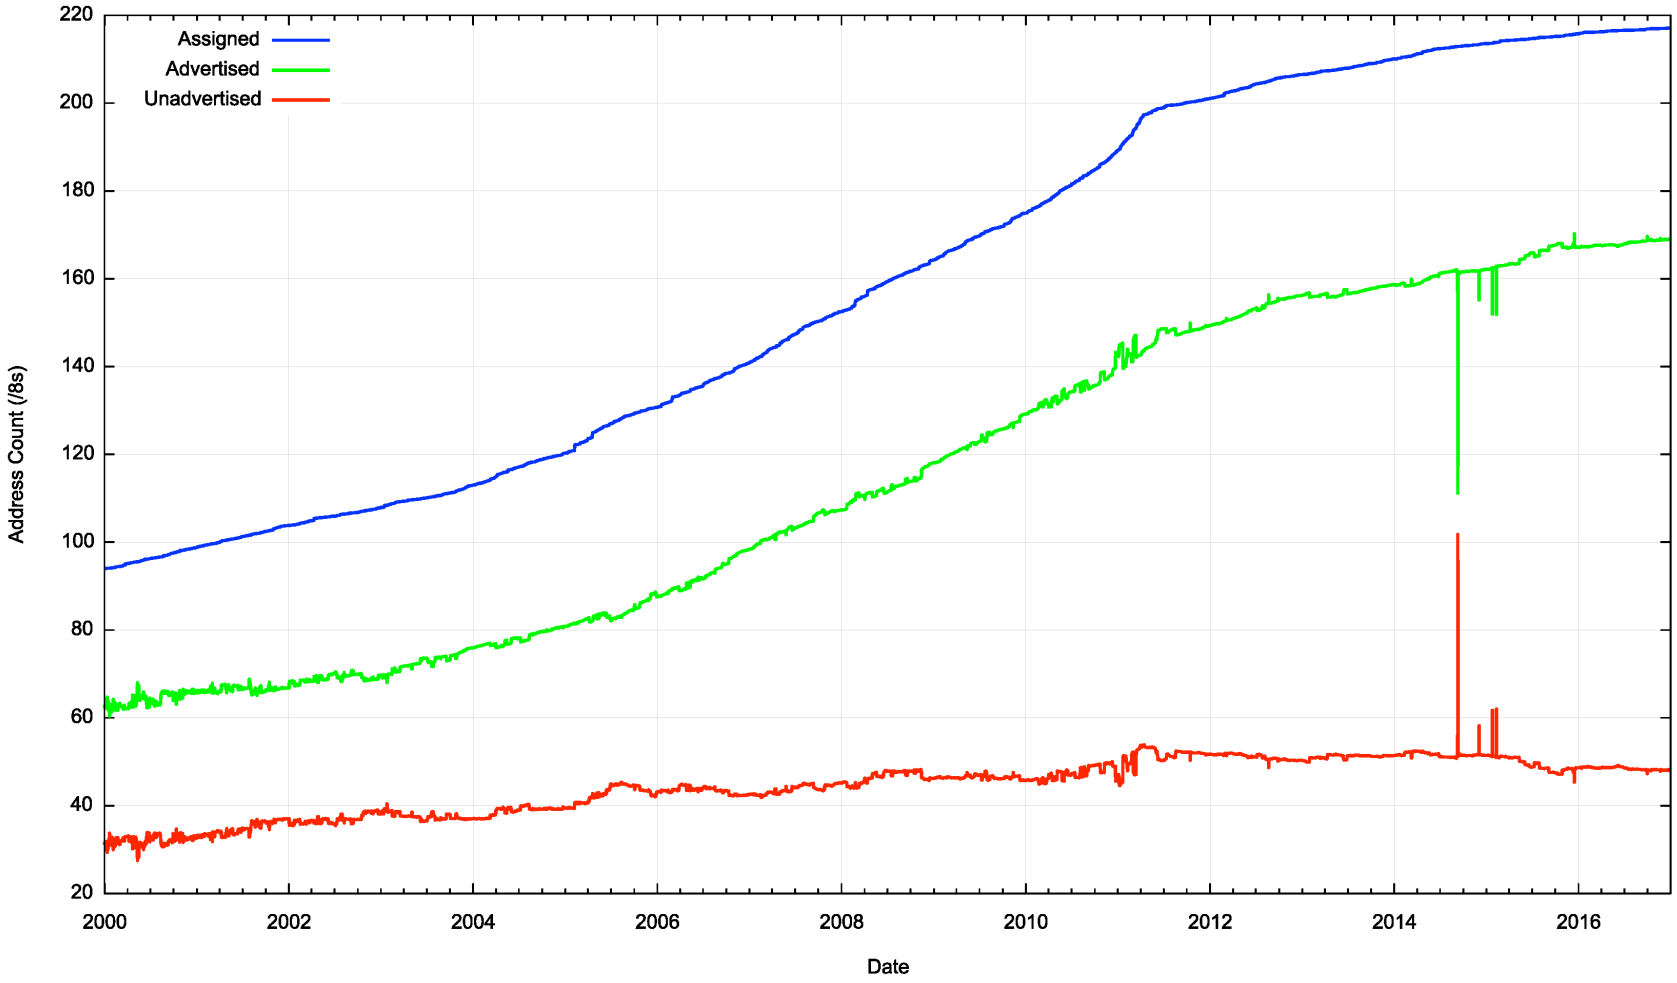

IPv4: Unadvertised Addresses

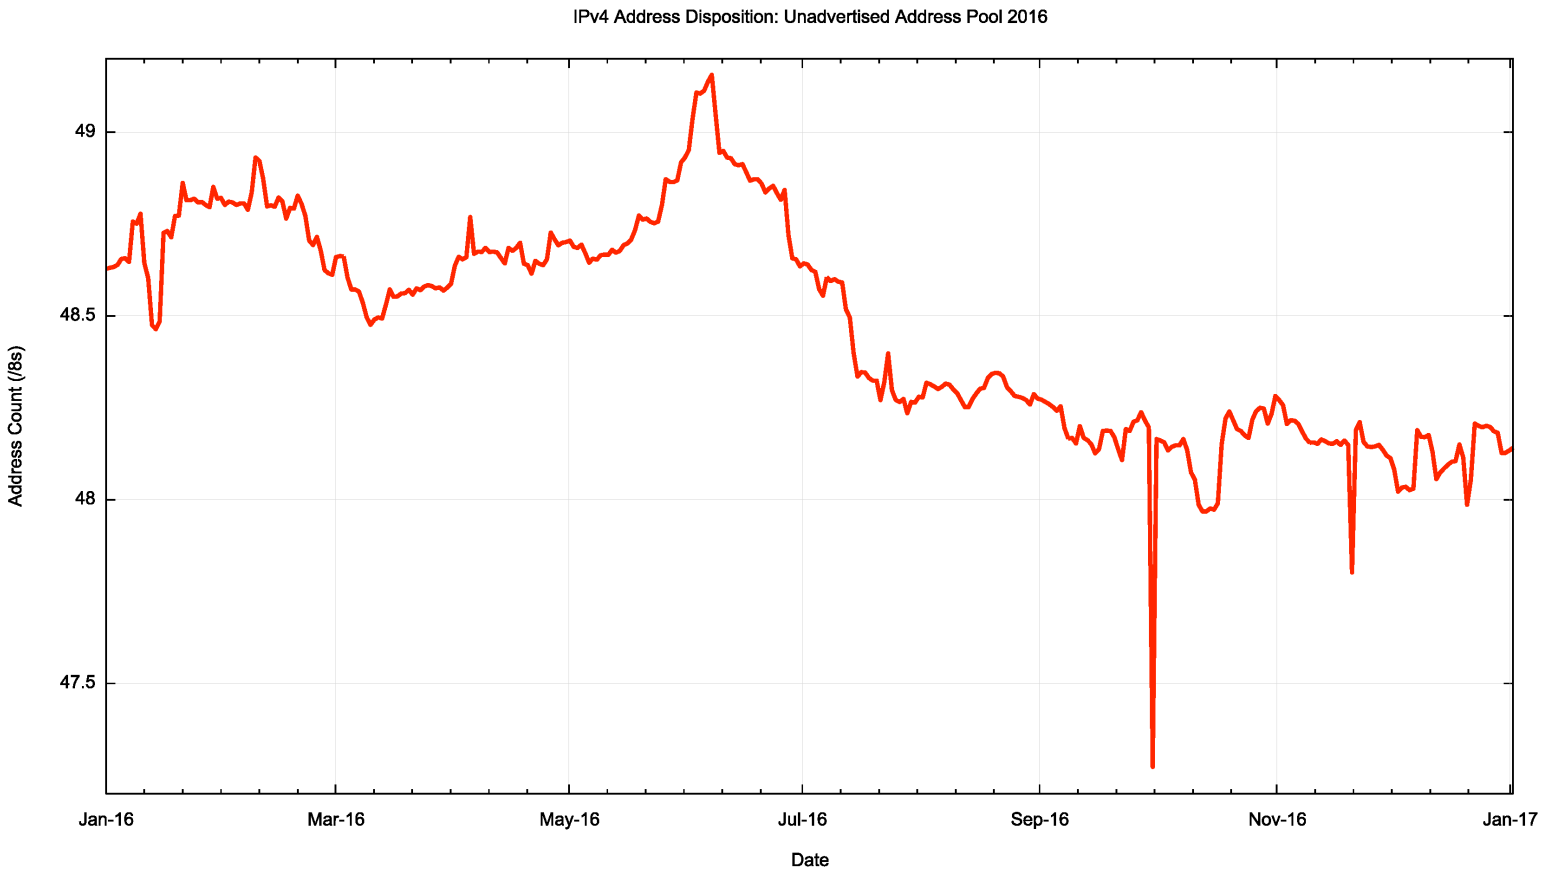

IPv4: Assigned vs Recovered

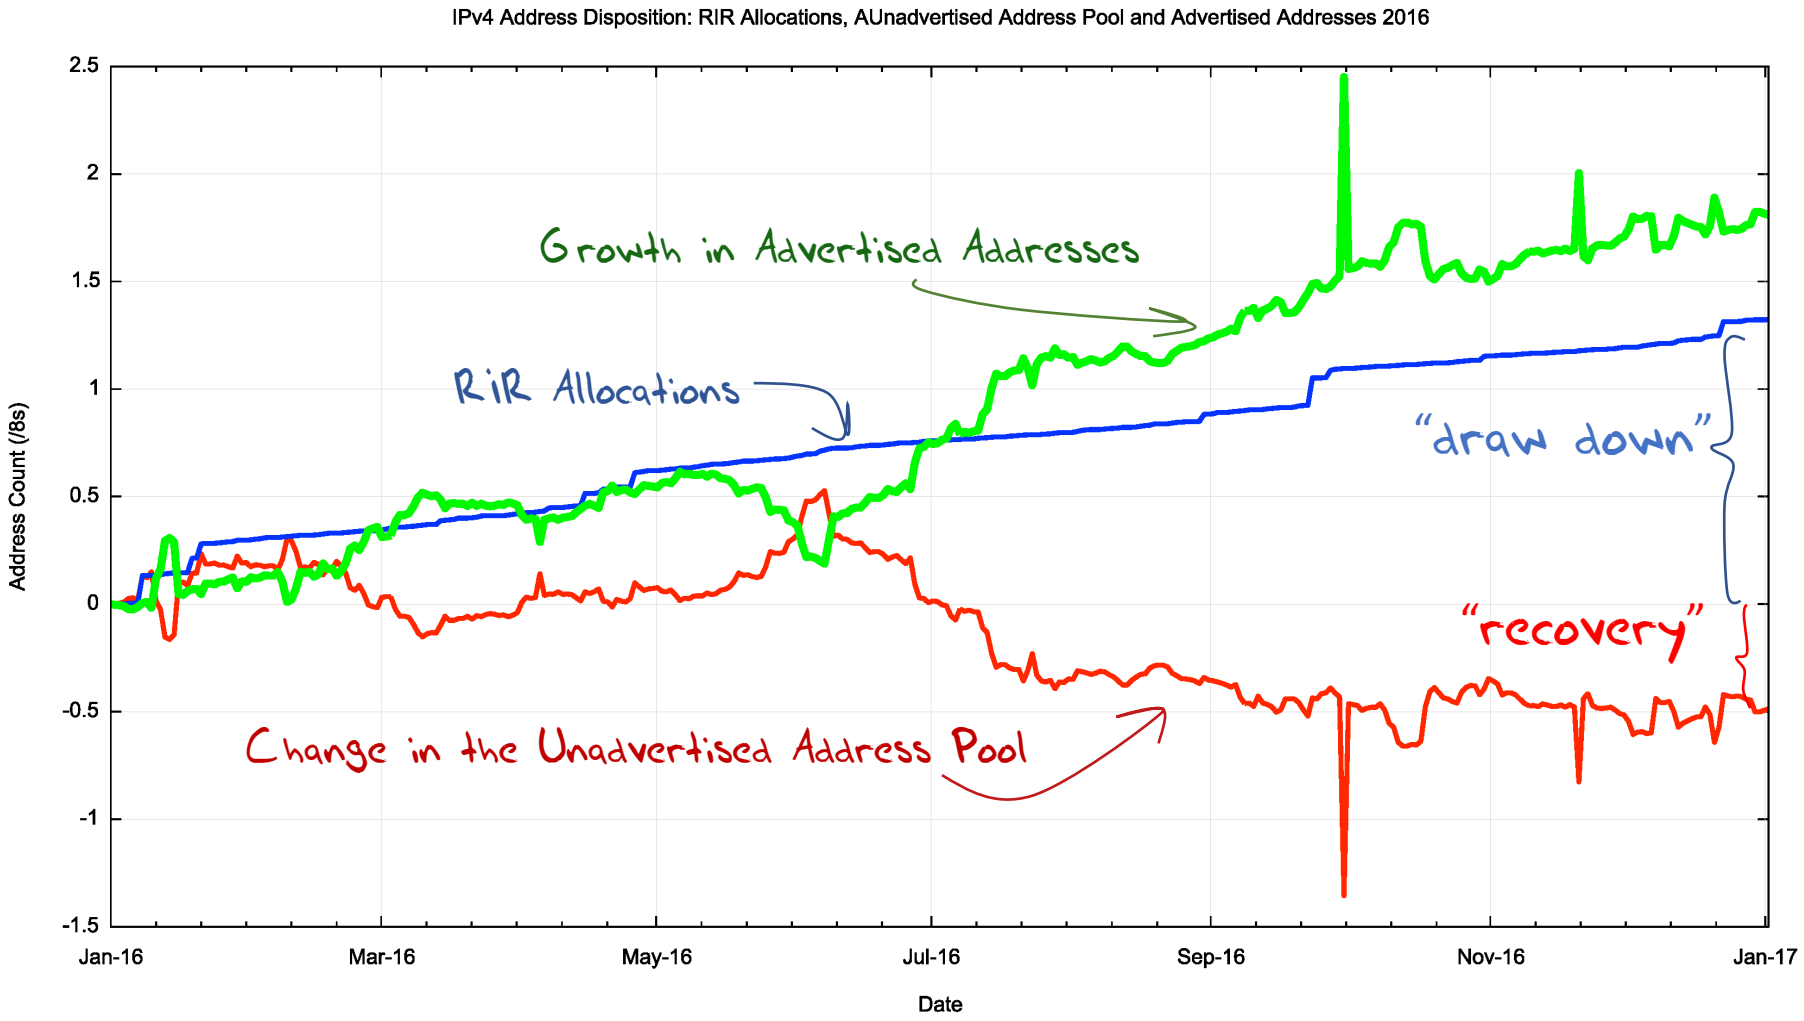

IPv4 in 2016

The equivalent of 1.8 /8s was added to the routing table across 2016

- Approximately 1.3 /8s were assigned by RIRs in 2015
  - 0.7 /8's assigned by Afrinic
  - 0.2 /8s were assigned by APNIC, RIPE NCC (Last /8 allocations)
  - 0.1 /8s were assigned by ARIN, LACNIC

- And a net of 0.5 /8's were recovered from the Unadvertised Pool

The Route-Views view of IPv6

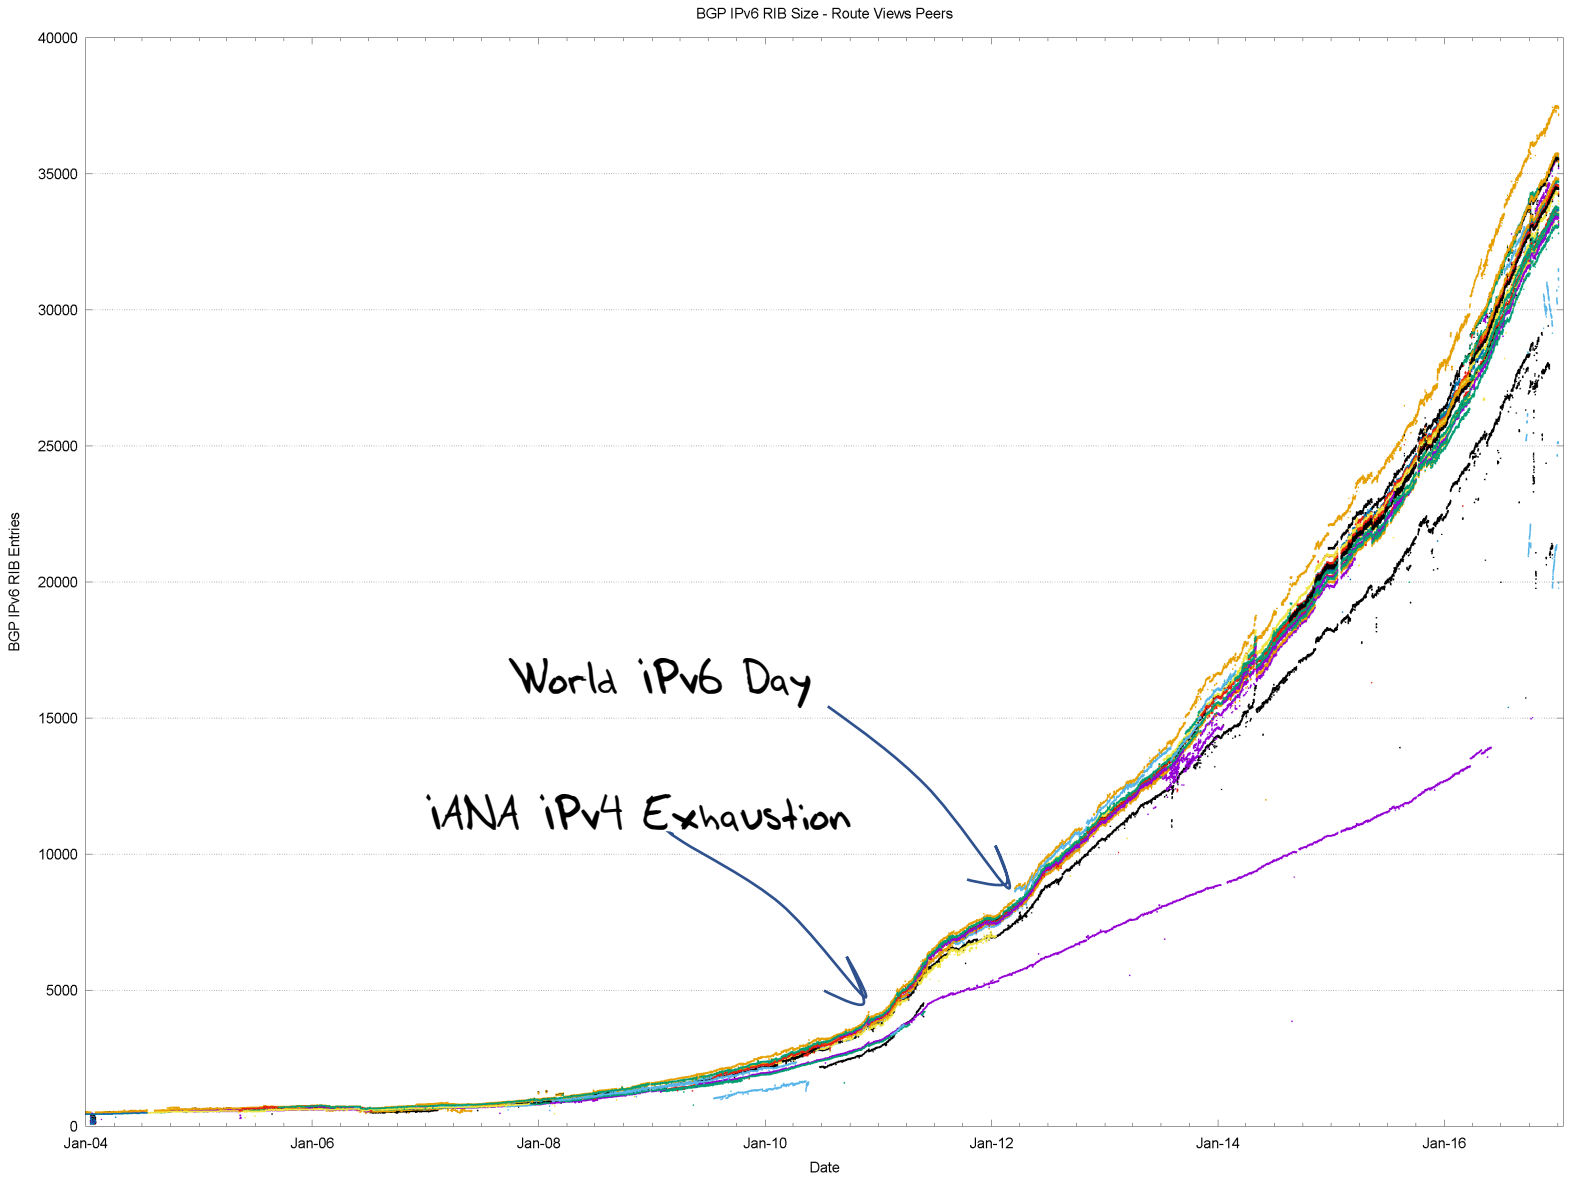

2015-2016 in detail

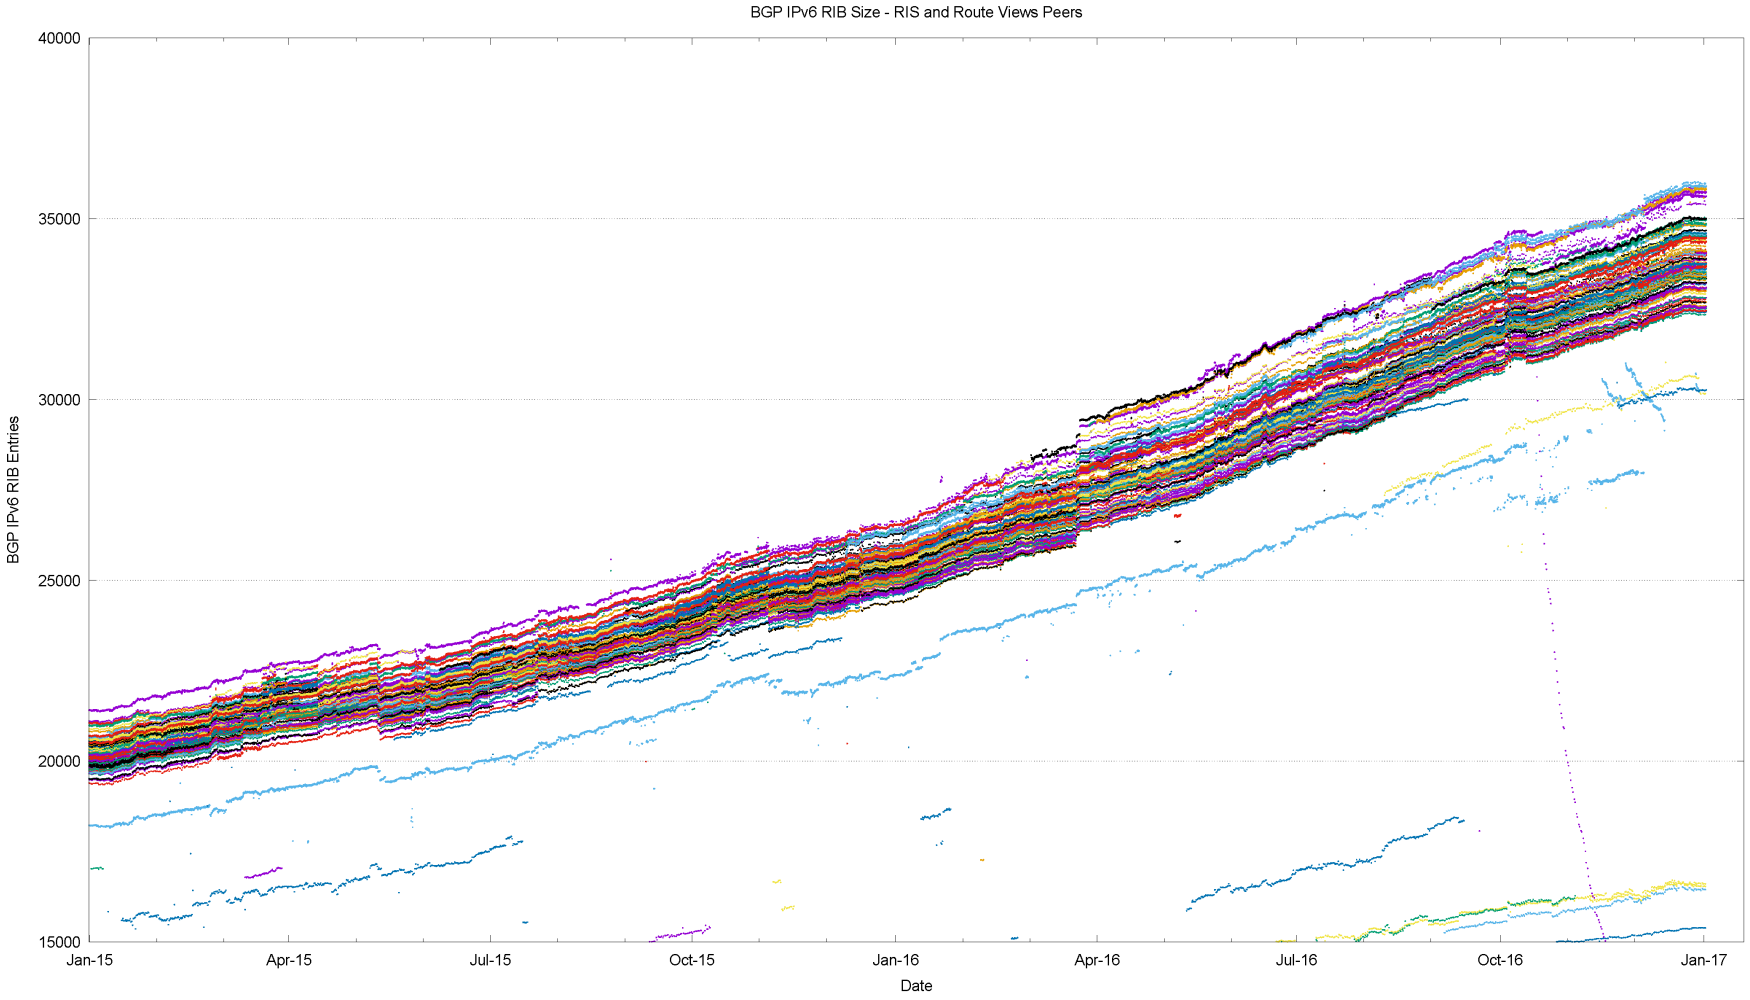

Routing Indicators for IPv6

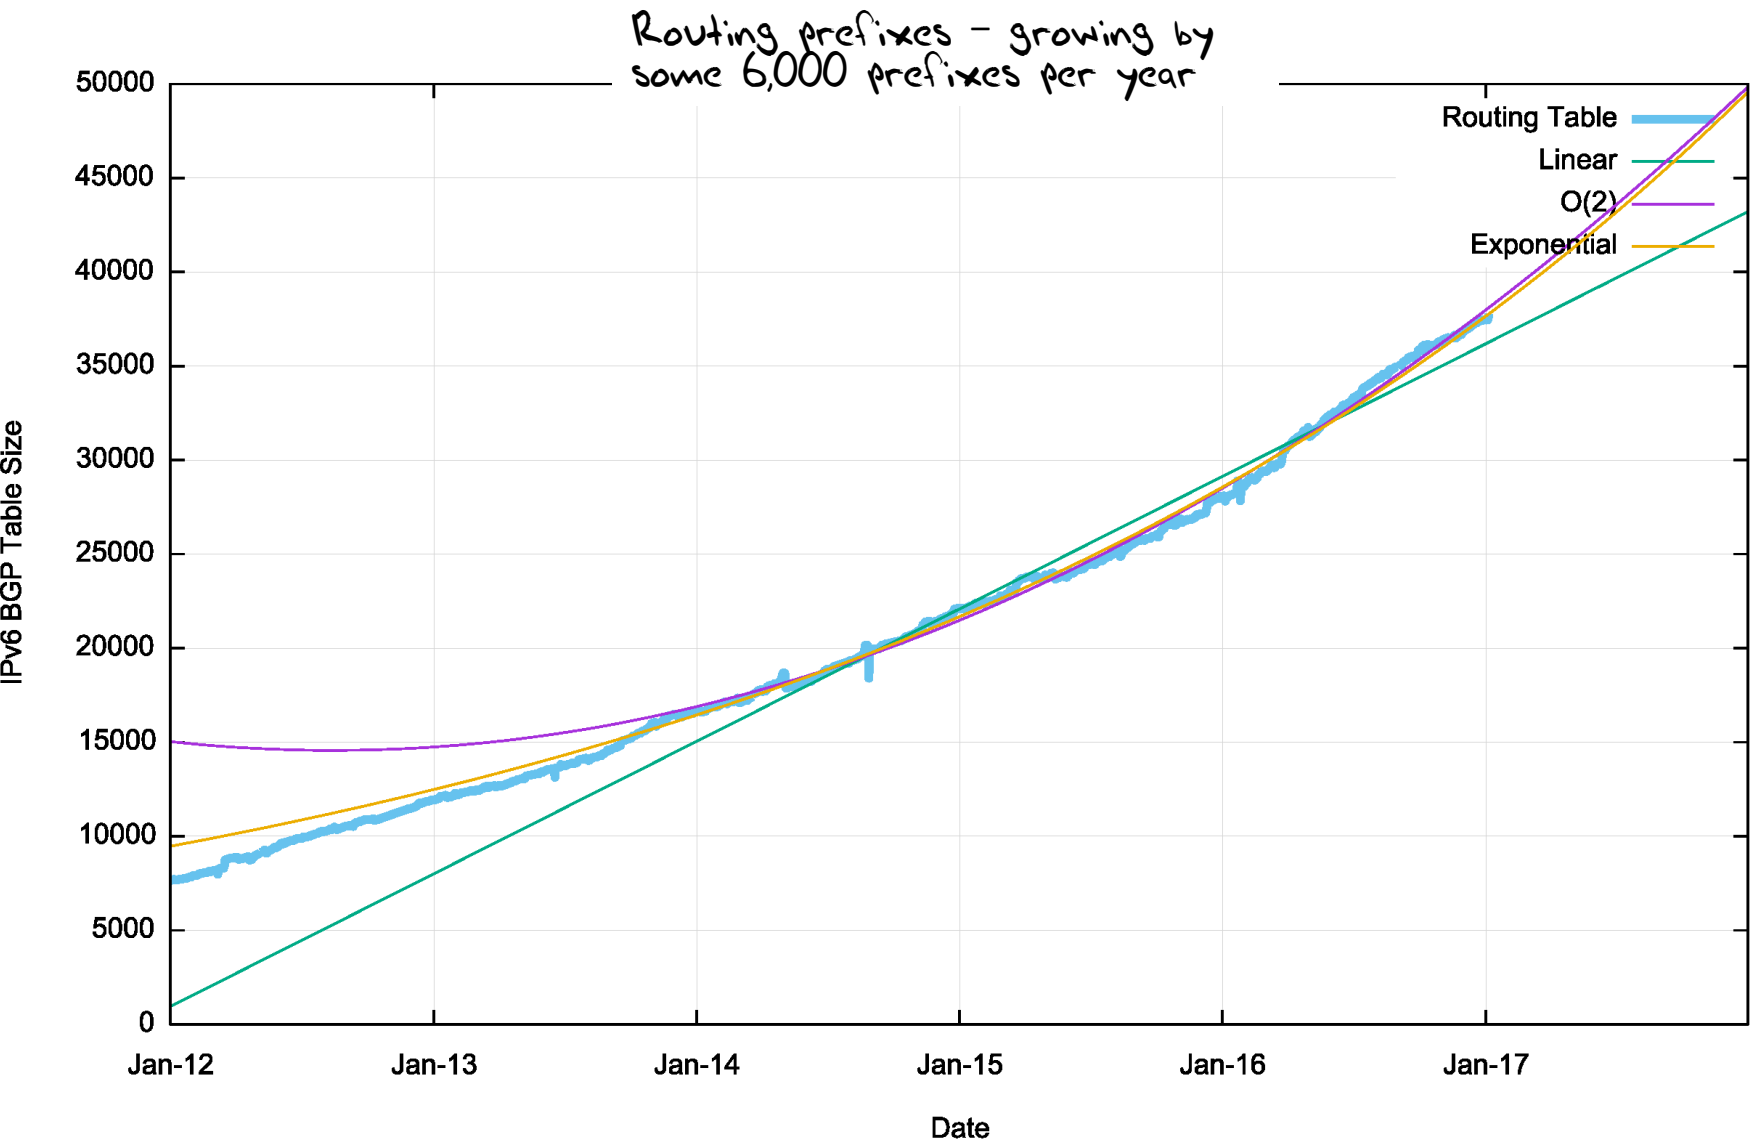

Routing Indicators for IPv6

AS Numbers- growing by some 1,700 prefixes per year (which is half the V4 growth)

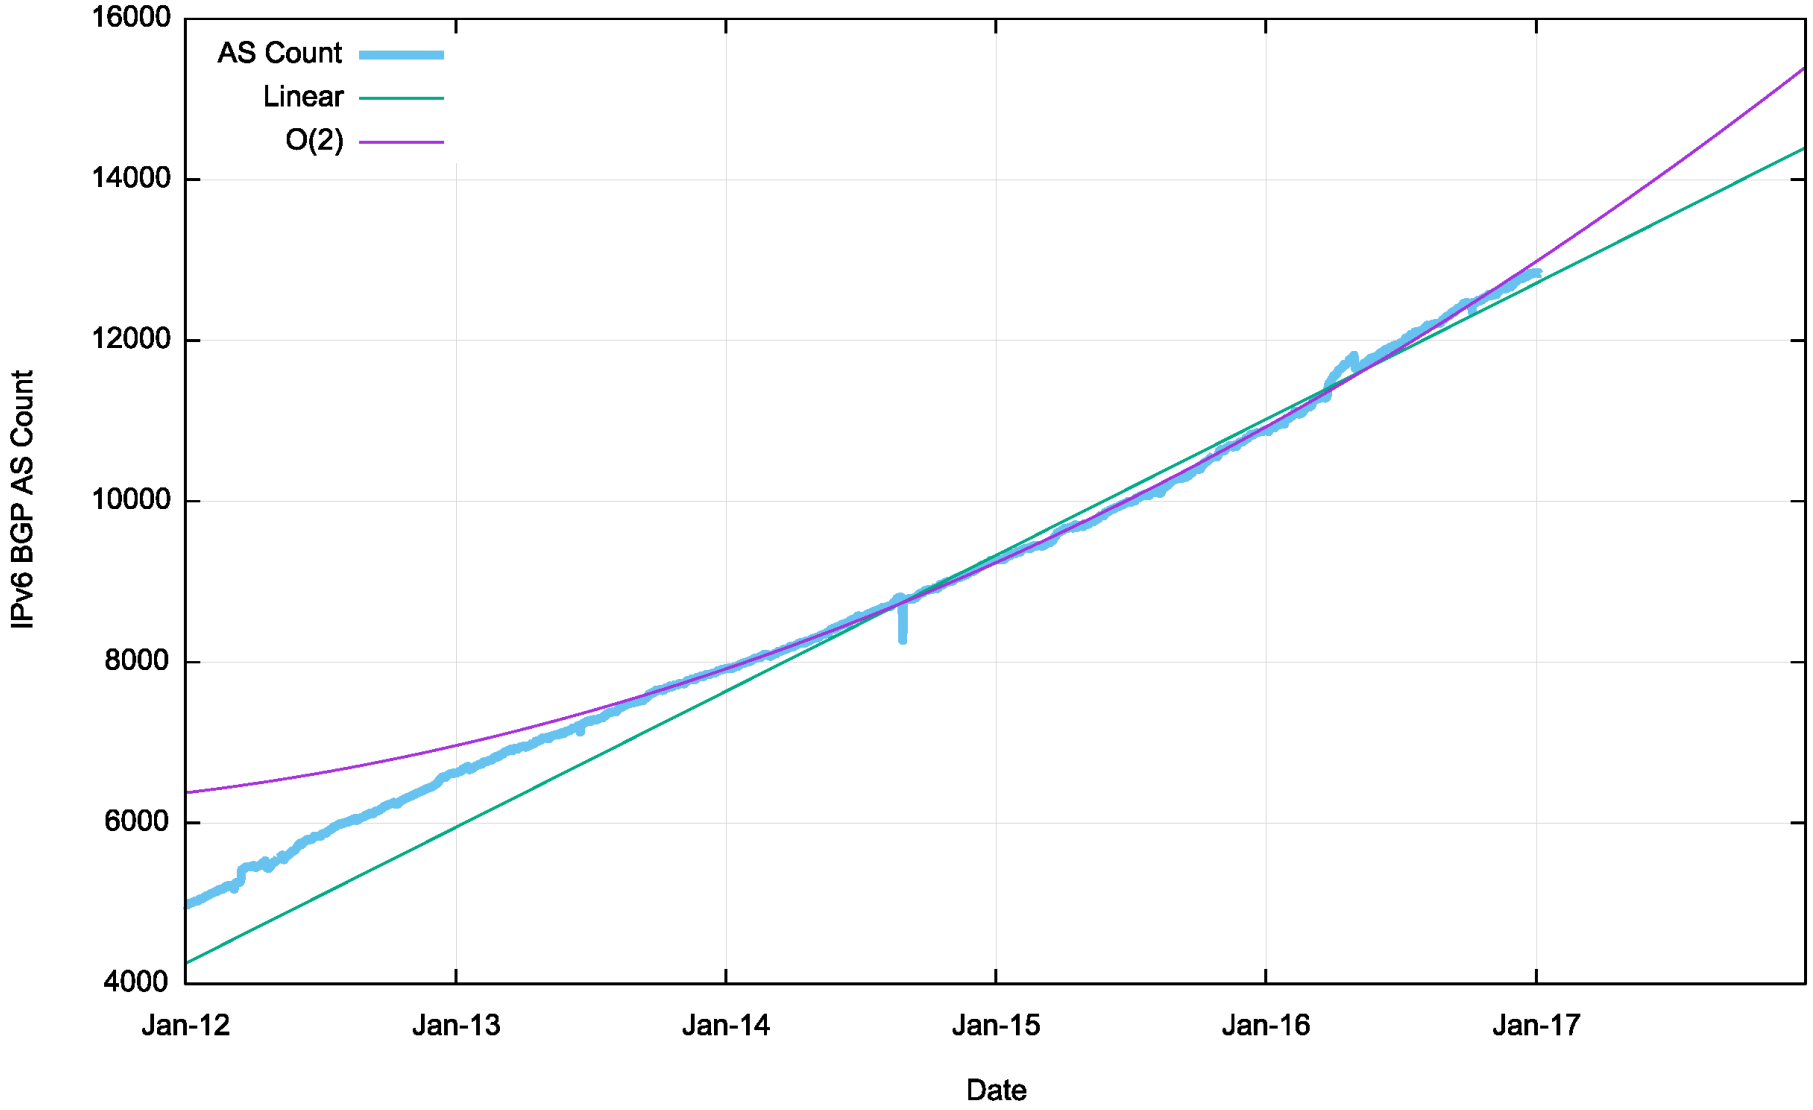

Routing Indicators for IPv6

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

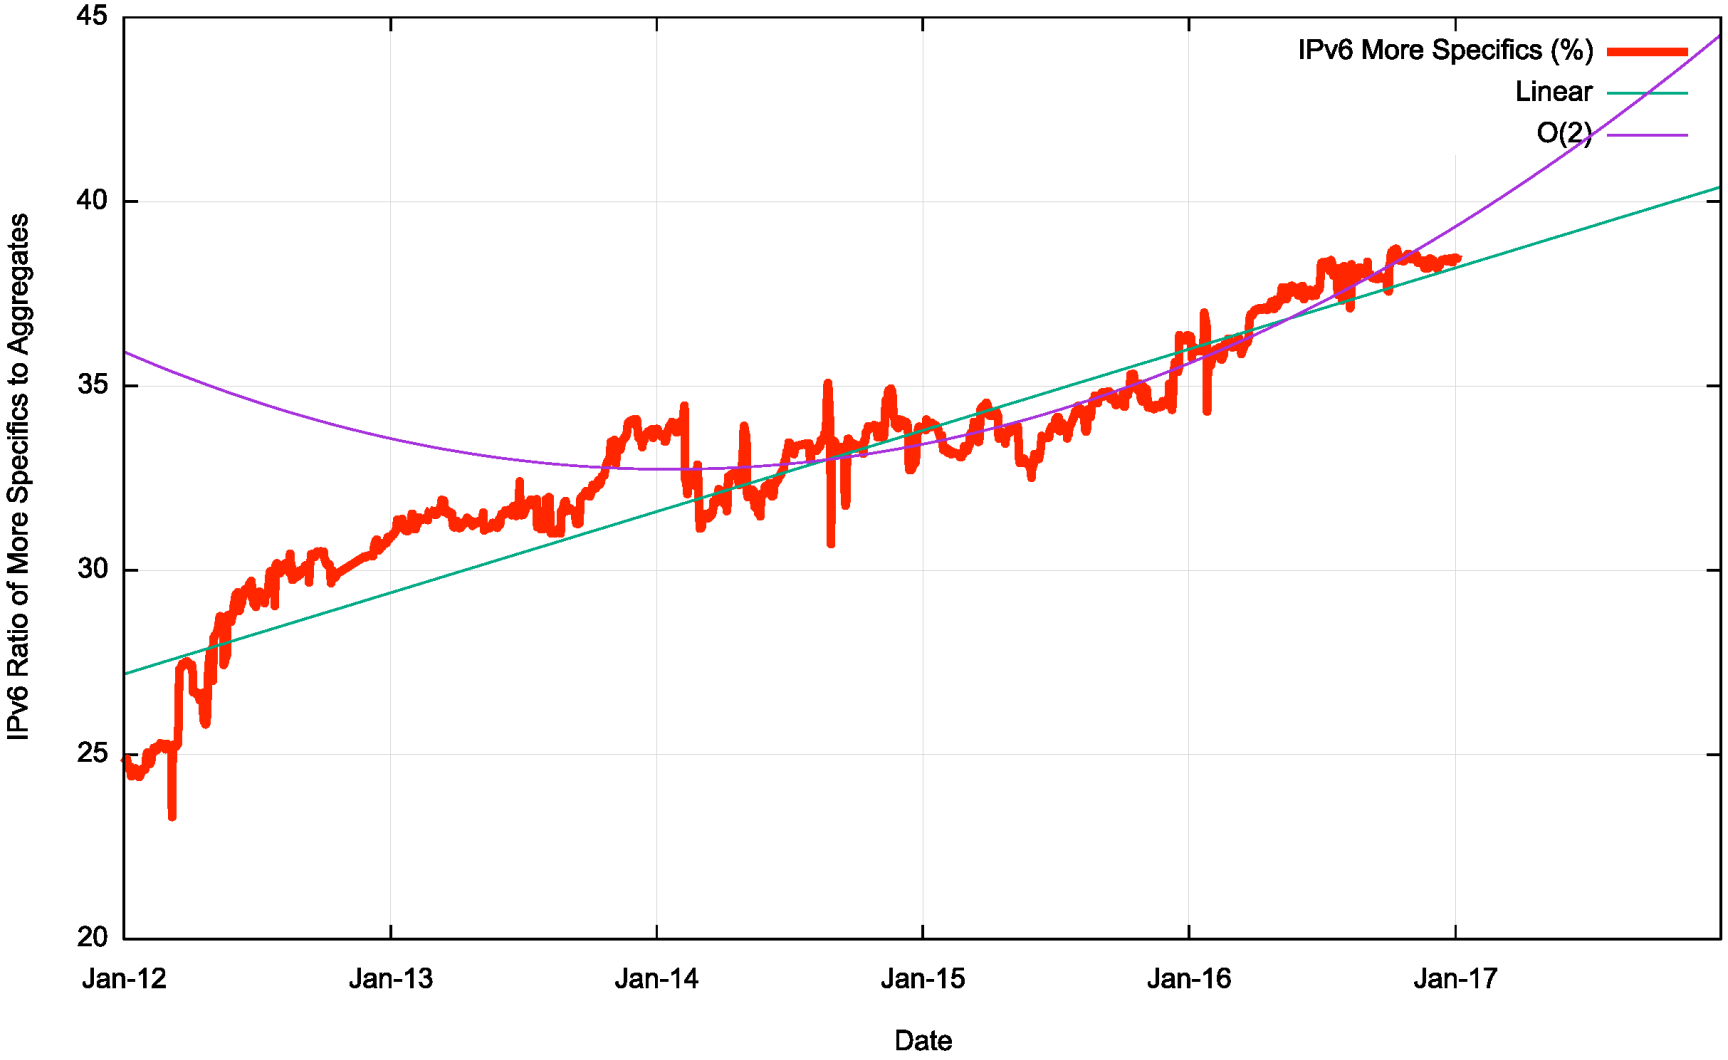

Routing Indicators for IPv6

The average size of a routing advertisement is getting smaller

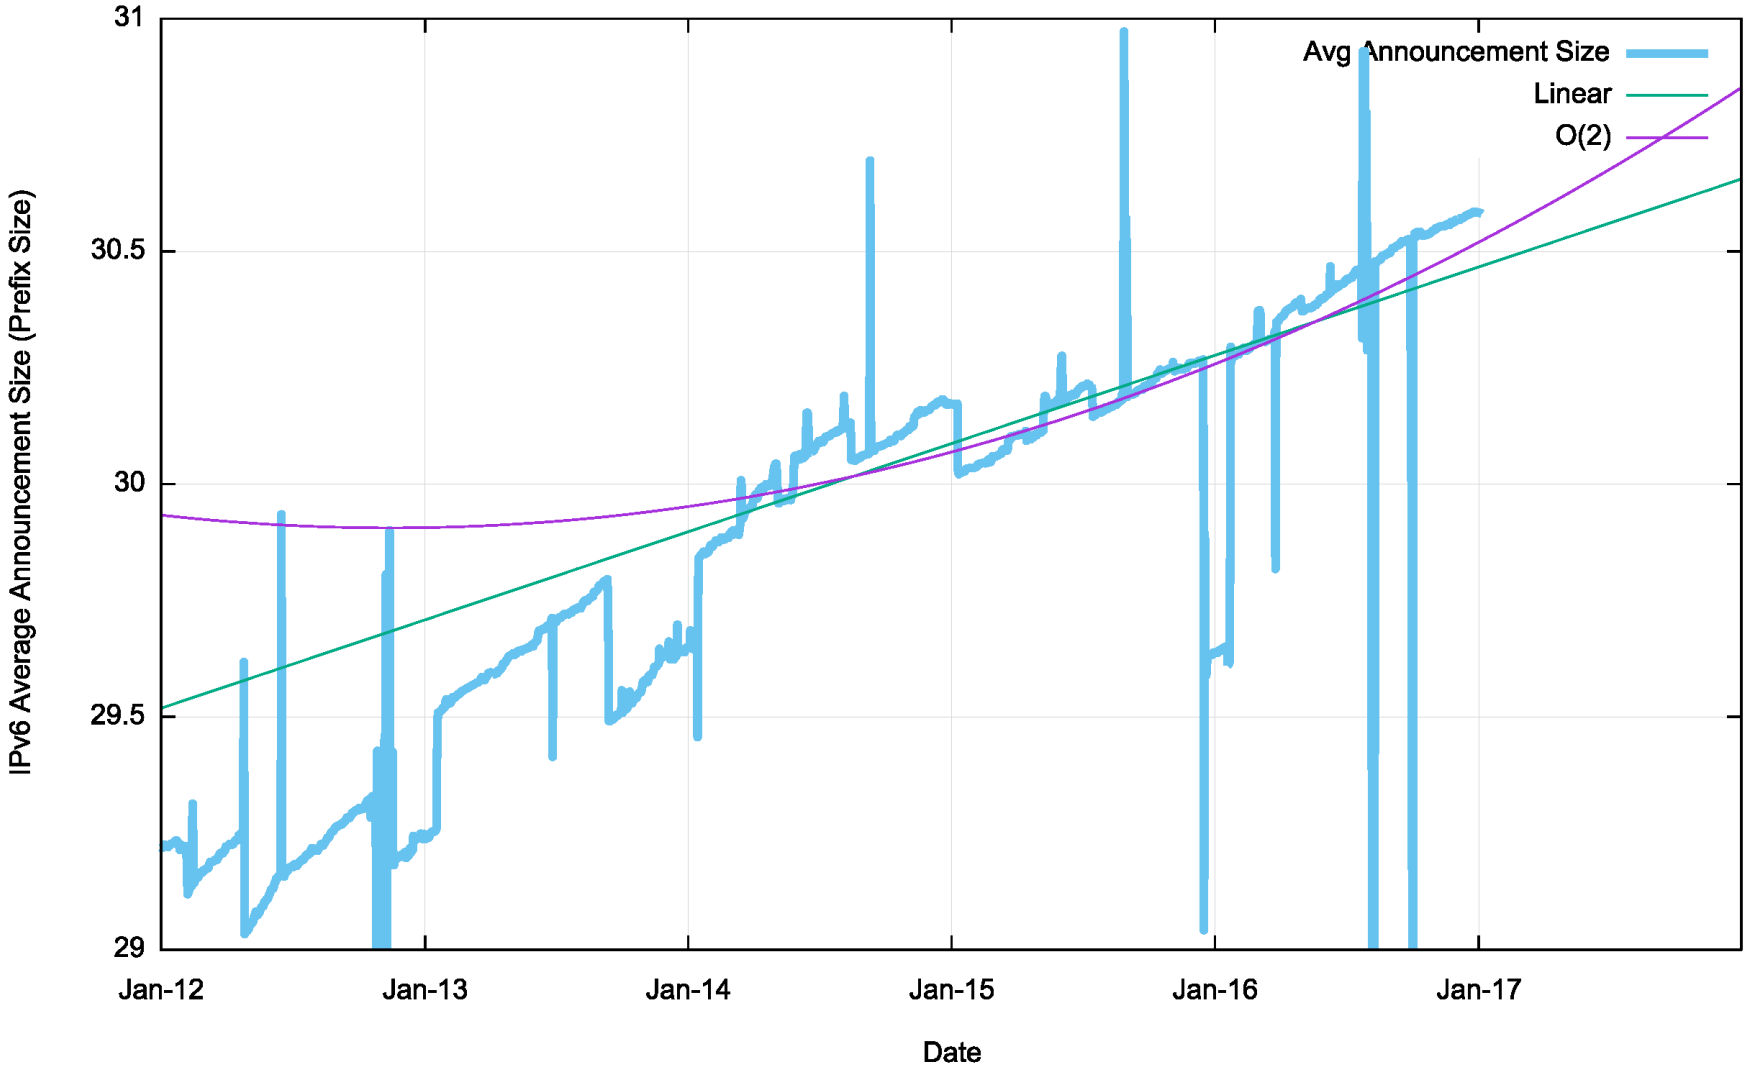

Routing Indicators for IPv6

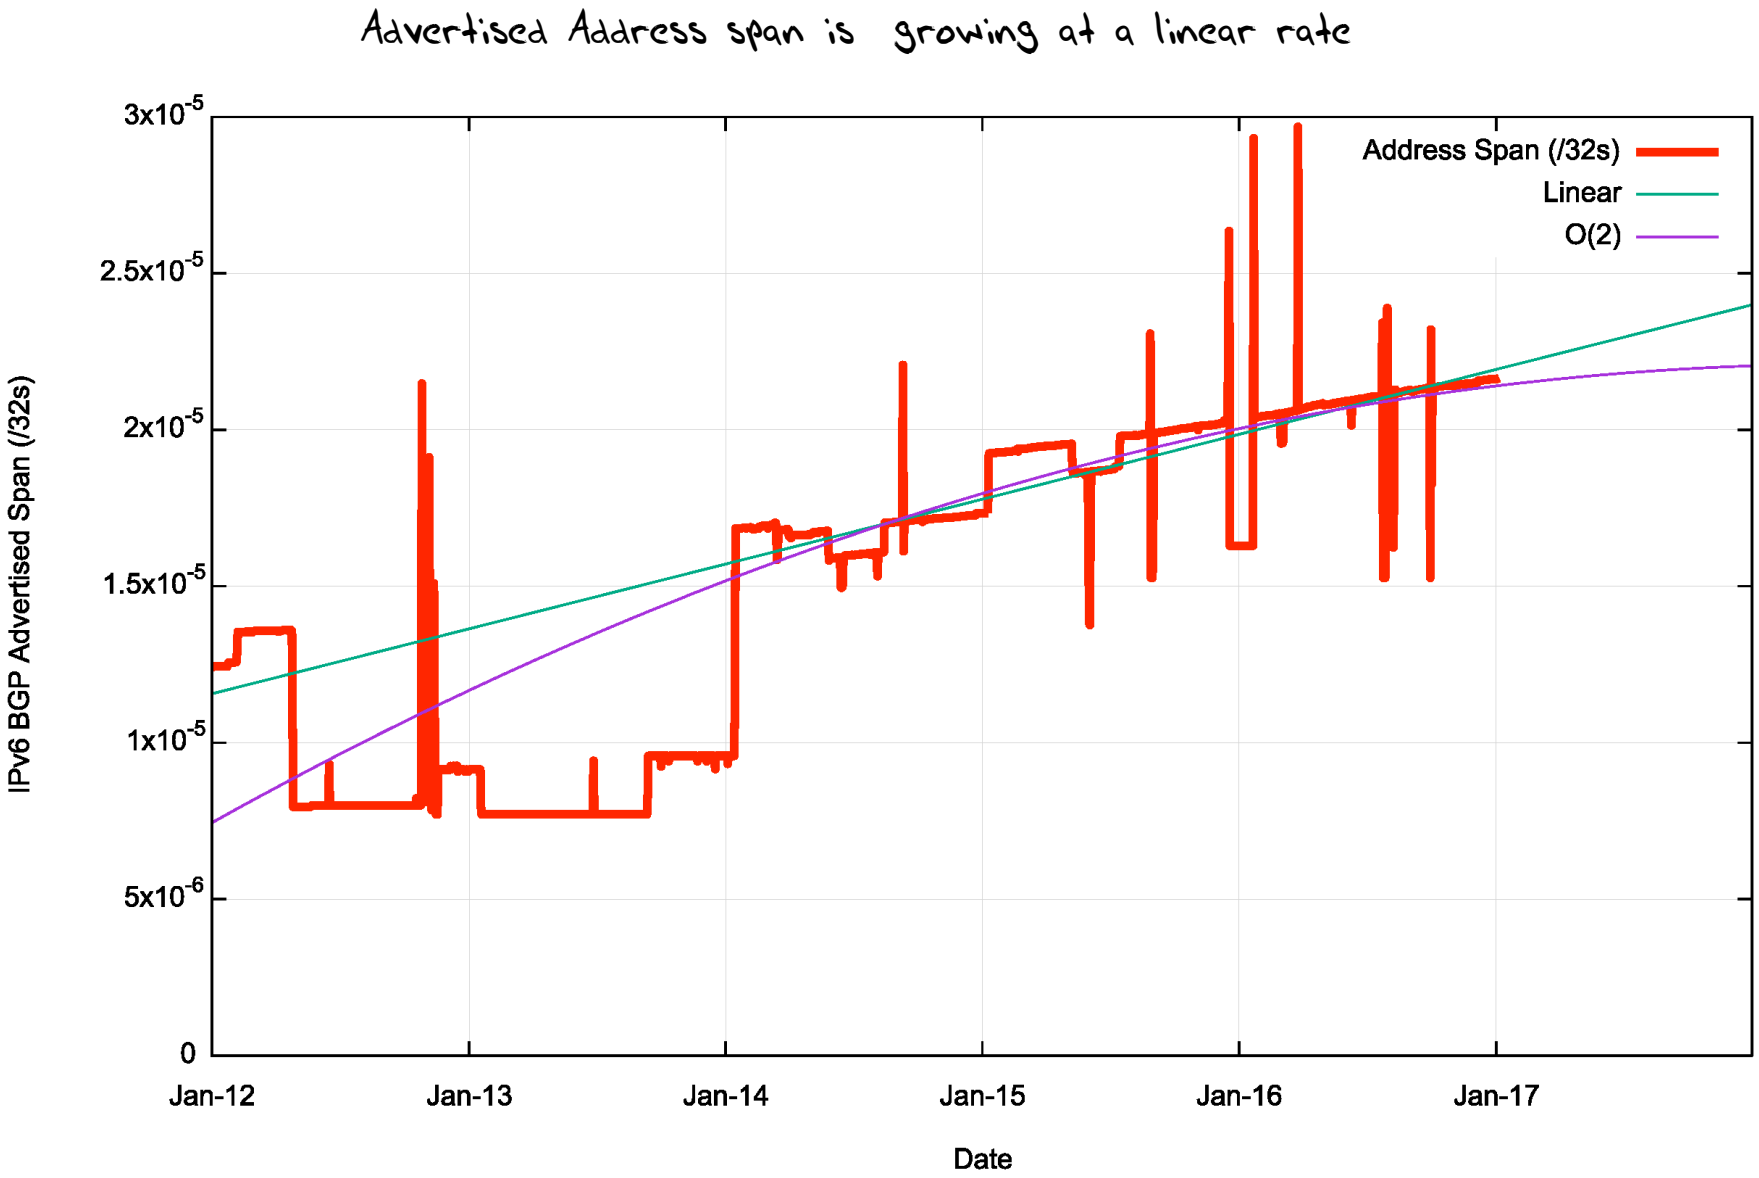

Routing Indicators for IPv6

The "shape" of inter-AS interconnection in IPv6 appears to be steady, as the Average AS Path length has been held steady

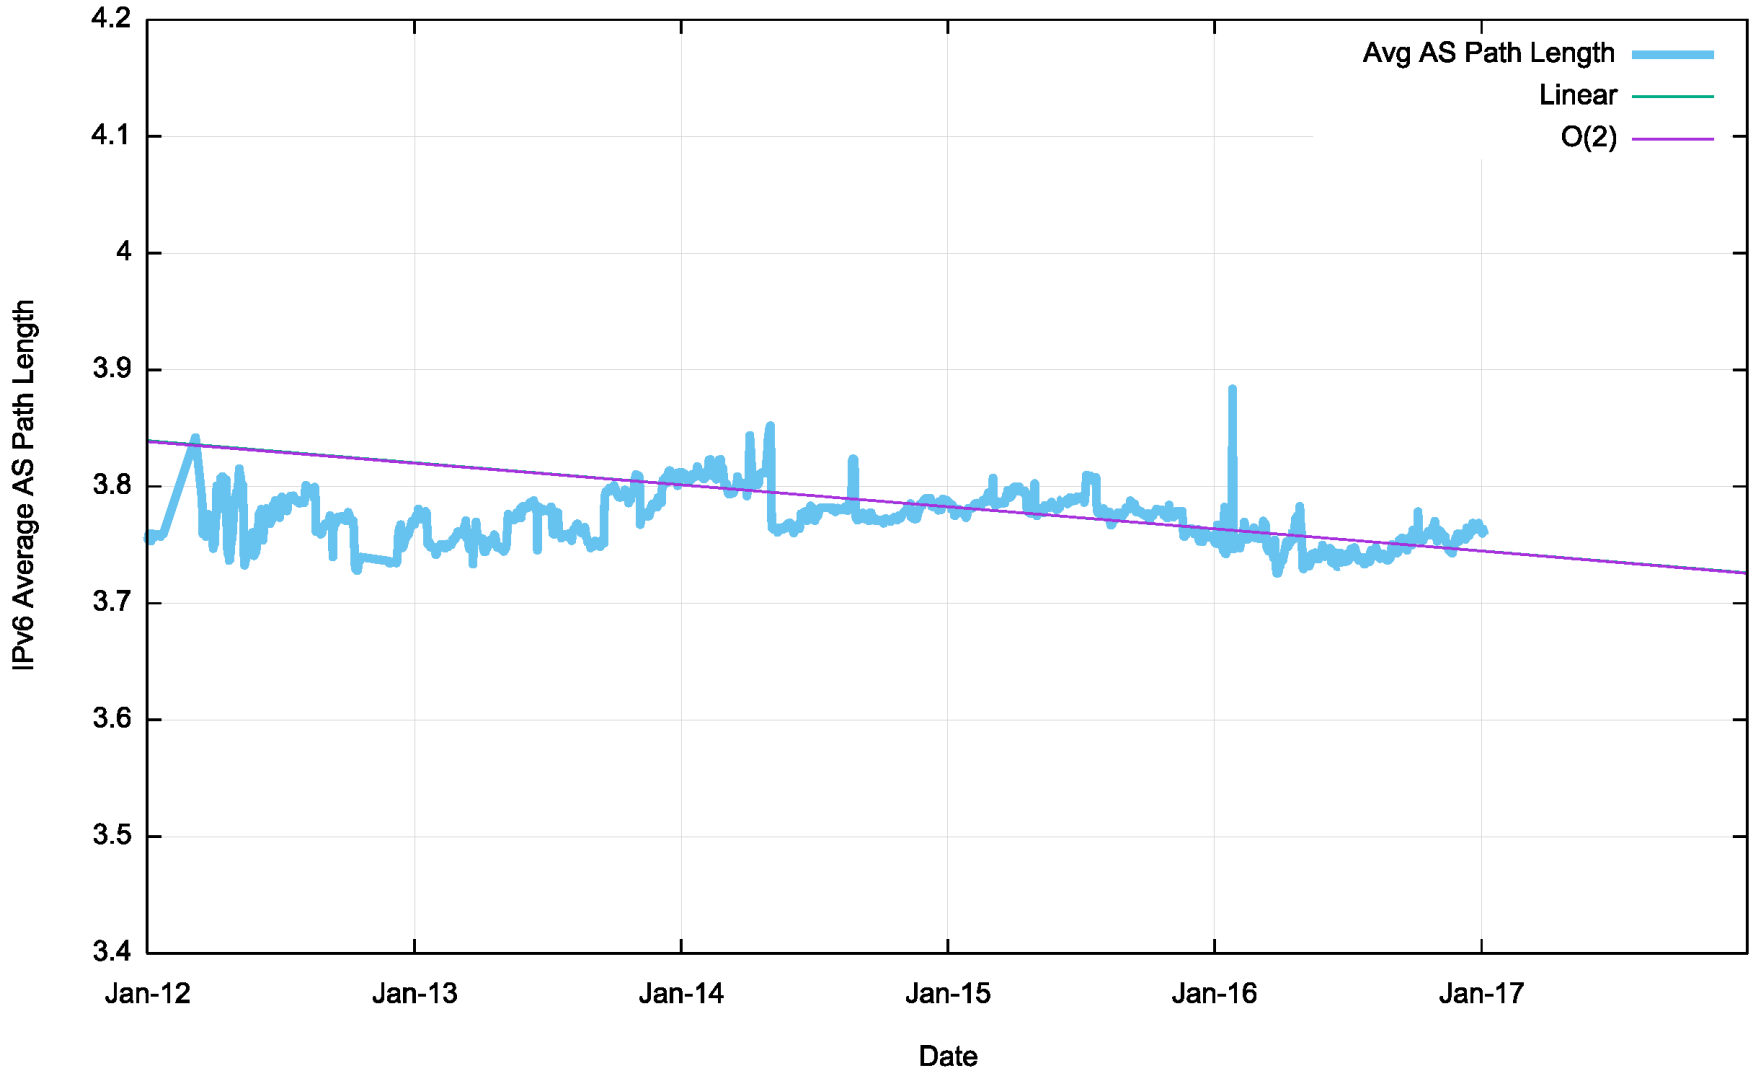

AS Adjacencies (Route Views)

9,105 out of 13,197 ASNs have 1 or 2 AS Adjacencies (69%)
917 ASNs have 10 or more adjacencies
4 ASNs have >1,000 adjacencies

3,276 AS6939 HURRICANE - Hurricane Electric, Inc., US
1,607 AS174 COGENT-174 - Cogent Communications, US
1,310 AS3356 LEVEL3 - Level 3 Communications, Inc., US
1,112 AS37100 SEACOM-AS, MU

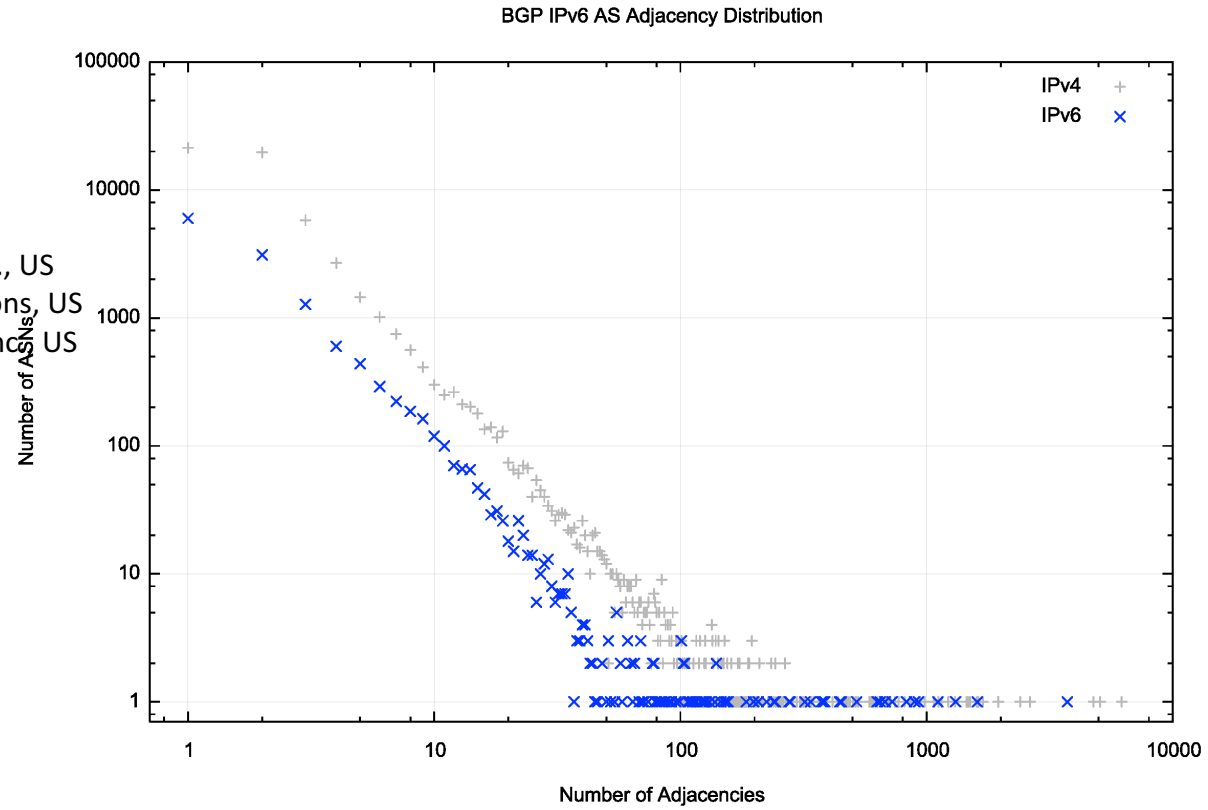

IPv6 in 2015

- Overall IPv6 Internet growth in terms of BGP is steady at some **6,000 route entries p.a.**

This is growth of BGP route objects is 1/9 of the growth rate of the IPv4 network – as compared to the AS growth rate which is 1/2 of the IPv4 AS number growth rate

What to expect

BGP Size Projections

For the Internet this is a time of extreme uncertainty

- Registry IPv4 address run out
- Uncertainty over the impacts of market-mediated movements of IPv4 on the routing table
- Uncertainty over the timing of IPv6 takeup leads to a mixed response to IPv6 so far, and no clear indicator of trigger points for change for those remaining IPv4-only networks

V4 - Daily Growth Rates

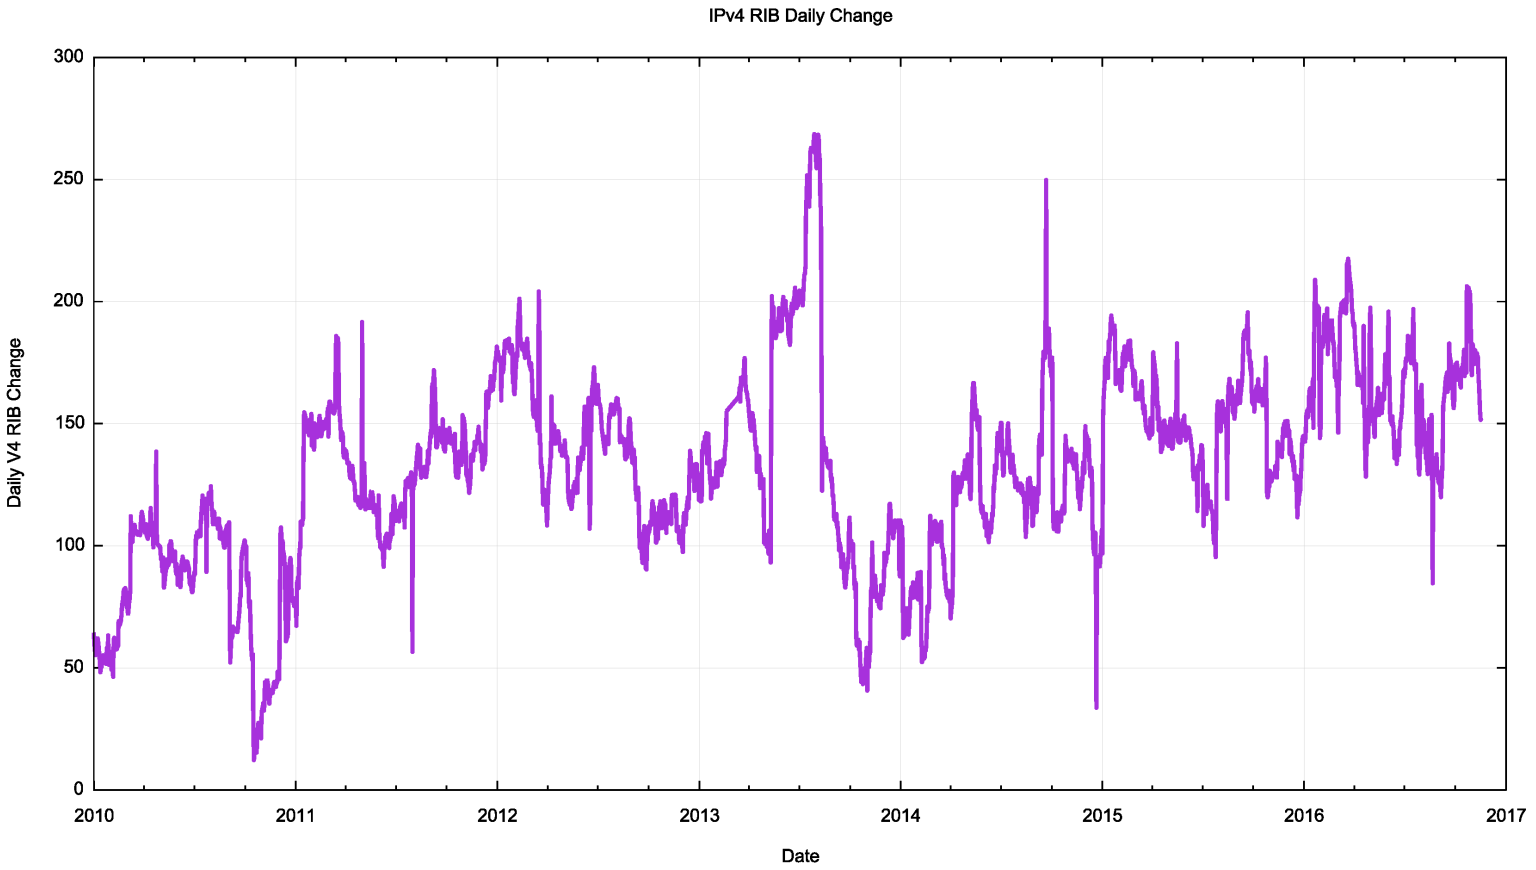

V4 - Daily Growth Rates

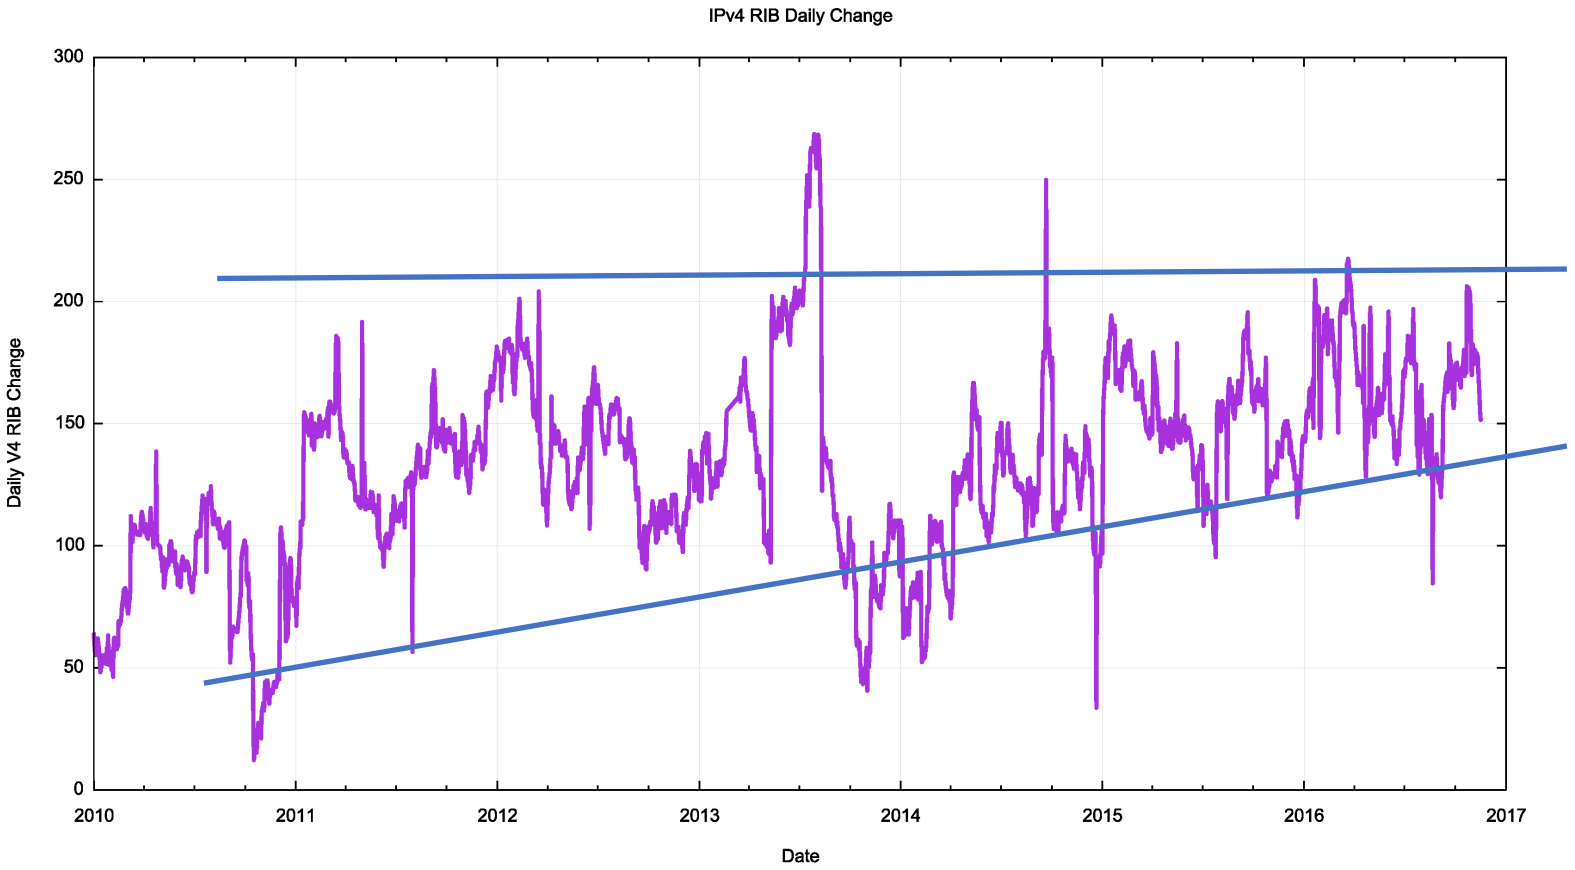

V4 - Relative Daily Growth Rates

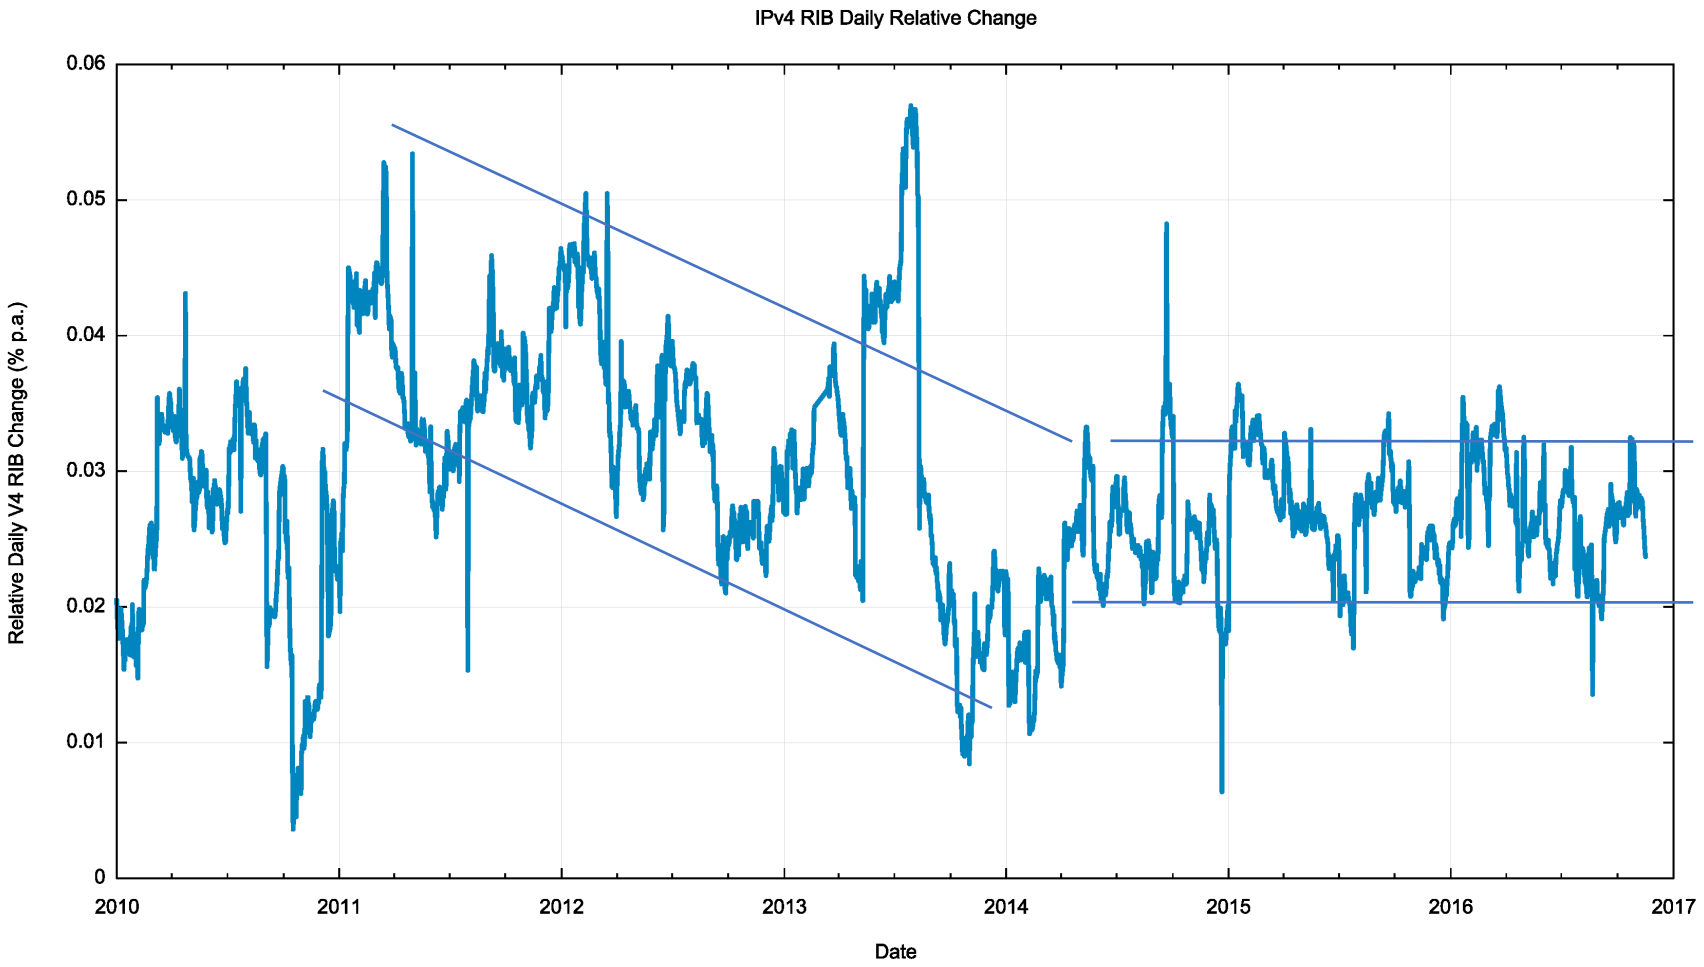

V4 - Relative Daily Growth Rates

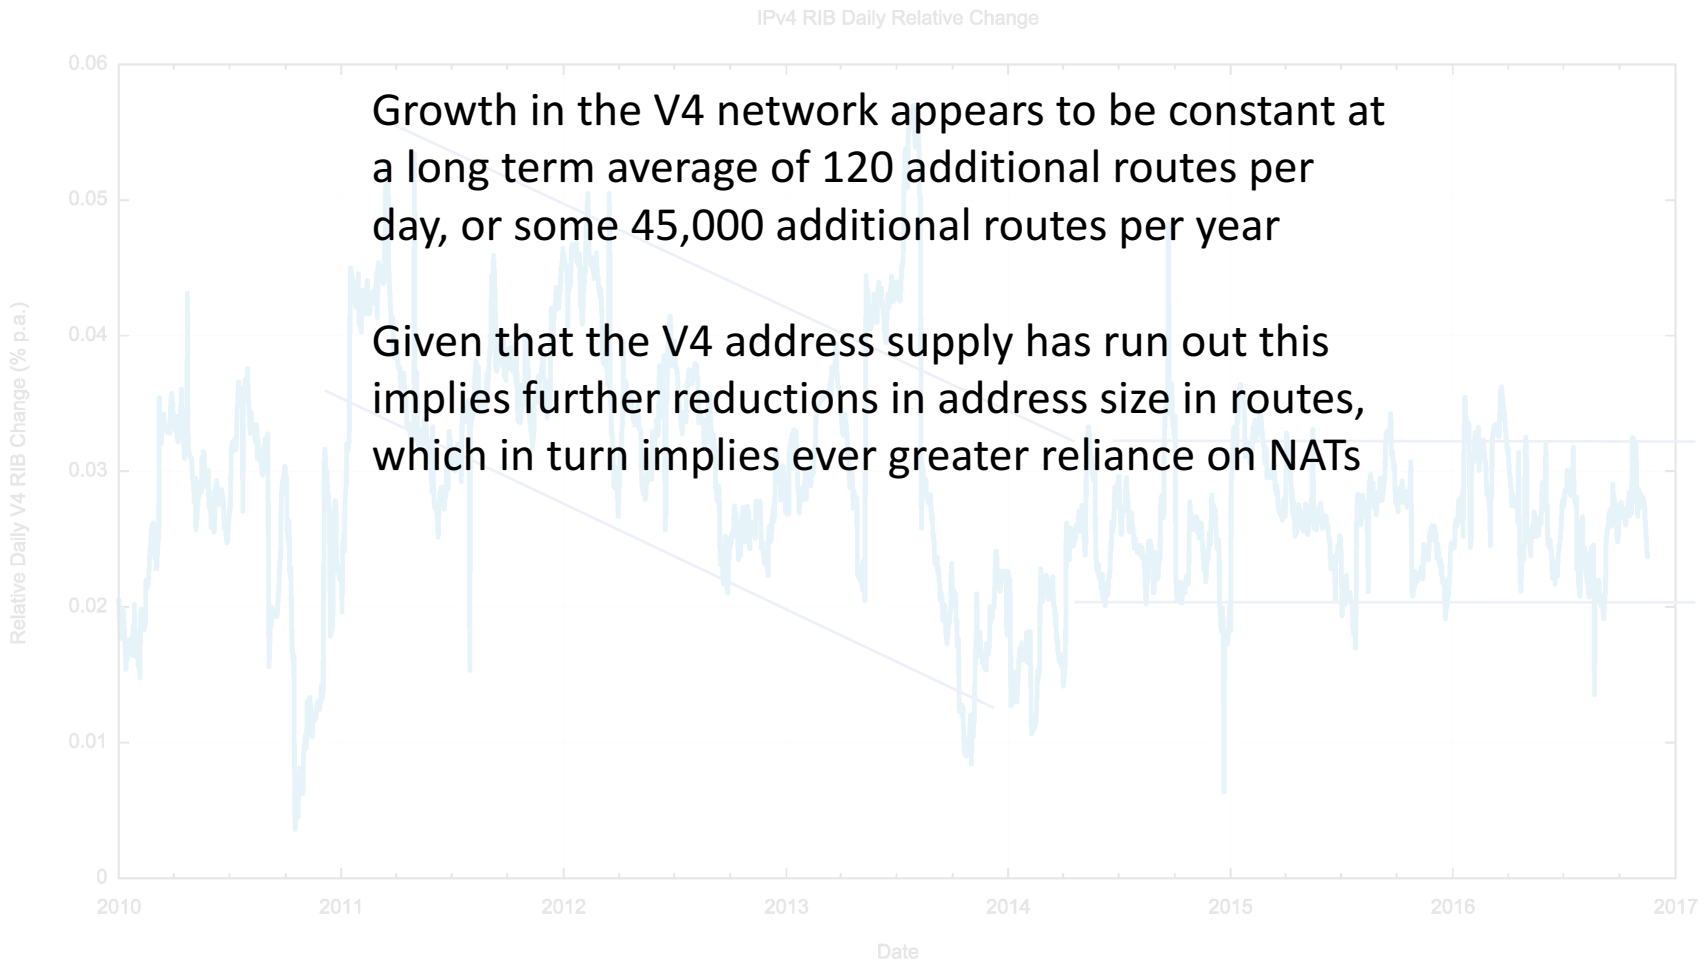

Growth in the V4 network appears to be constant at a long term average of 120 additional routes per day, or some 45,000 additional routes per year

Given that the V4 address supply has run out this implies further reductions in address size in routes, which in turn implies ever greater reliance on NATs

IPv4 BGP Table Size Predictions

	Jan 2017 PREDICTION	Jan 2016 PREDICTION	Jan 2015 PREDICTION
Jan 2013	441,000		
2014	488,000		
2015	*530,000*		
2016	*586,000*		*580,000*
2017	***646,000***	*628,000*	*620,000*
2018	*700,000*	*675,000*	*670,000*
2019	*754,000*	*722,000*	*710,000*
2020	*808,000*	*768,000*	*760,000*
2021	*862,000*	*815,000*	
2022	*916,000*		

V6 - Daily Growth Rates

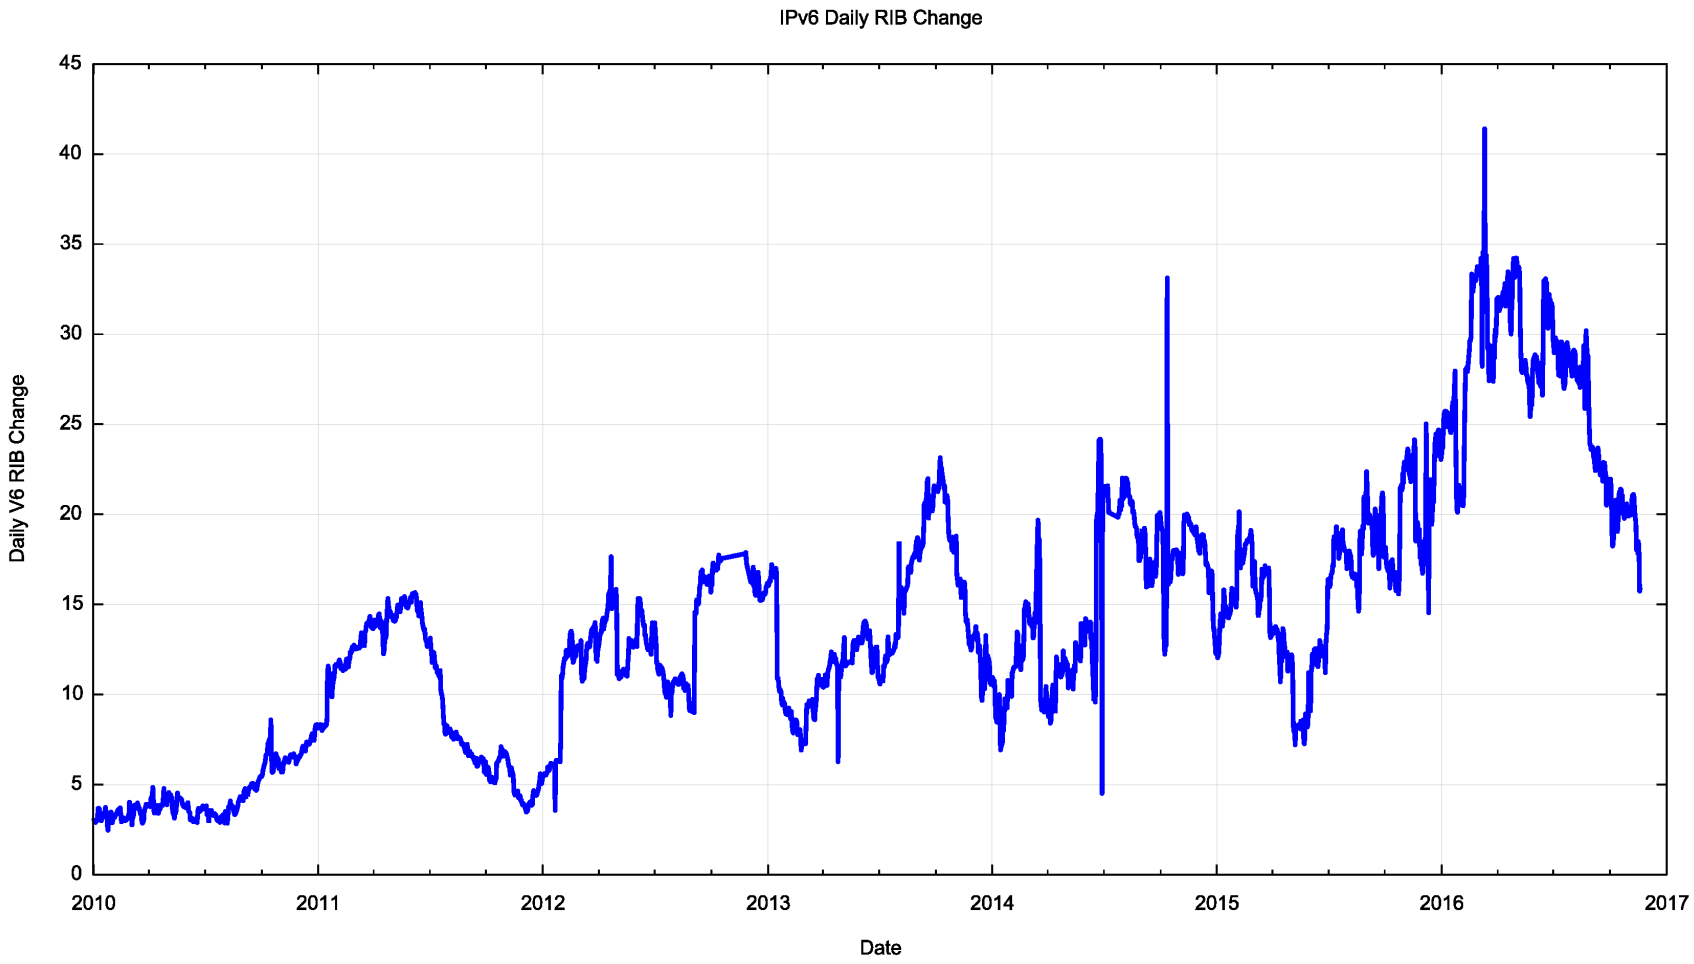

V6 - Relative Growth Rates

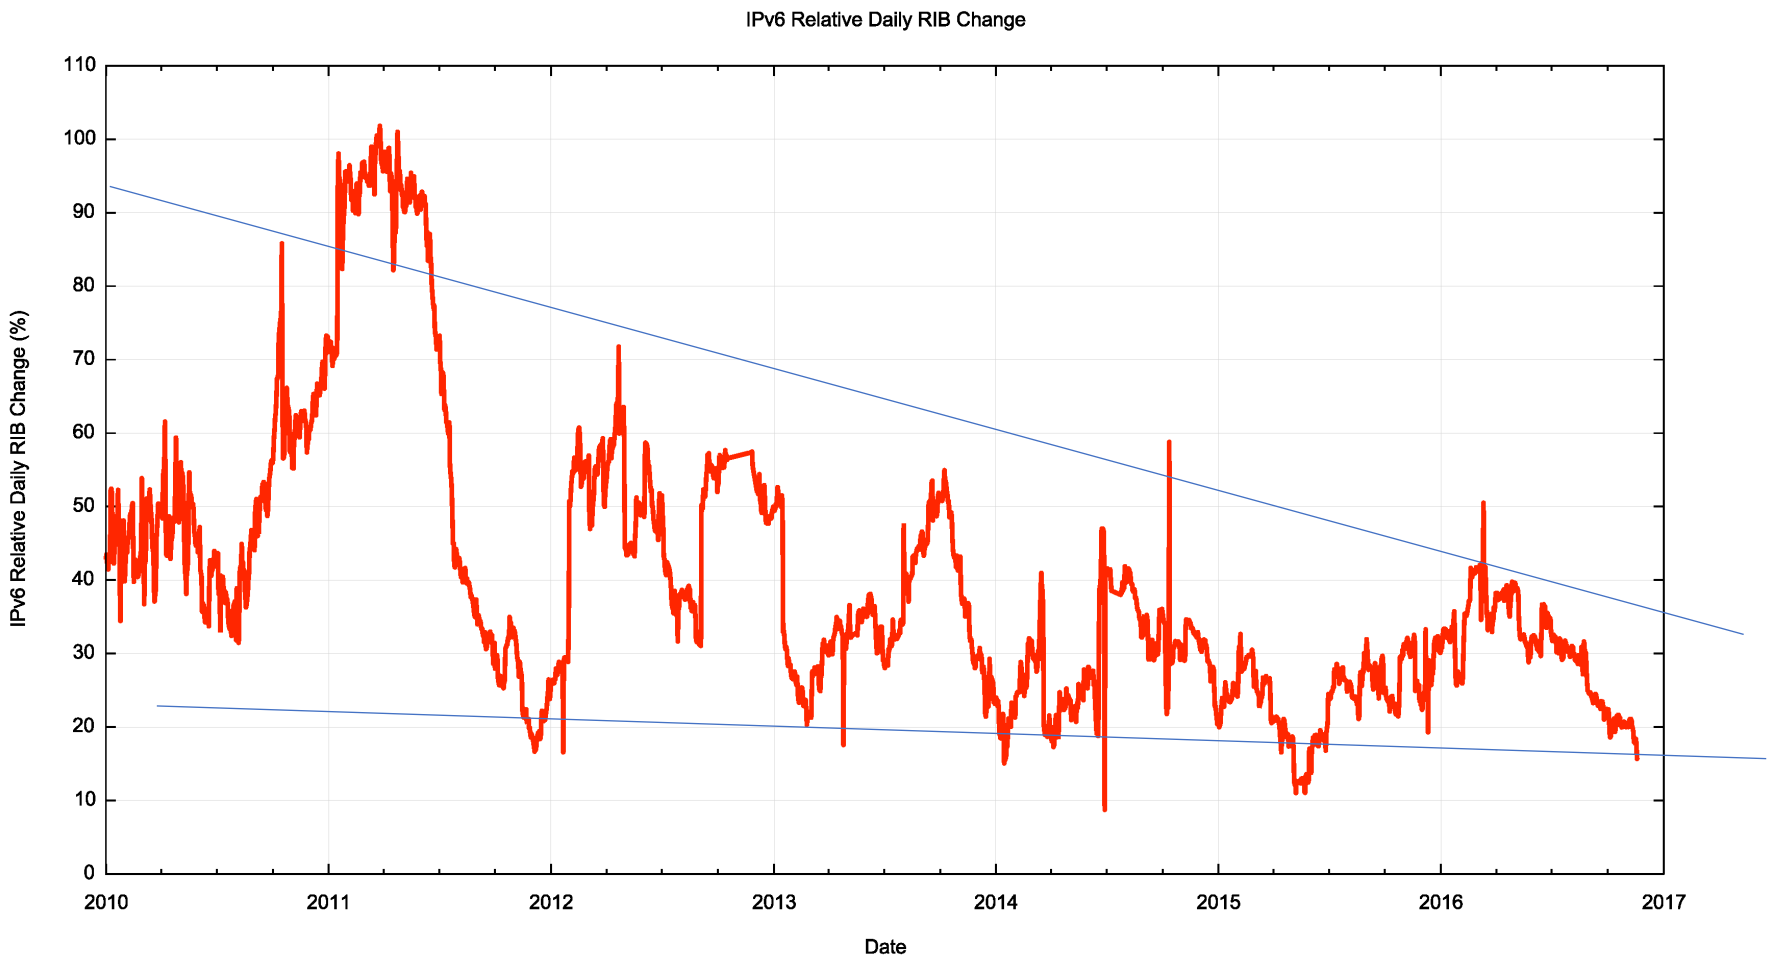

V6 - Relative Growth Rates

Growth in the V6 network appears to be increasing, but in relative terms this is slowing down.

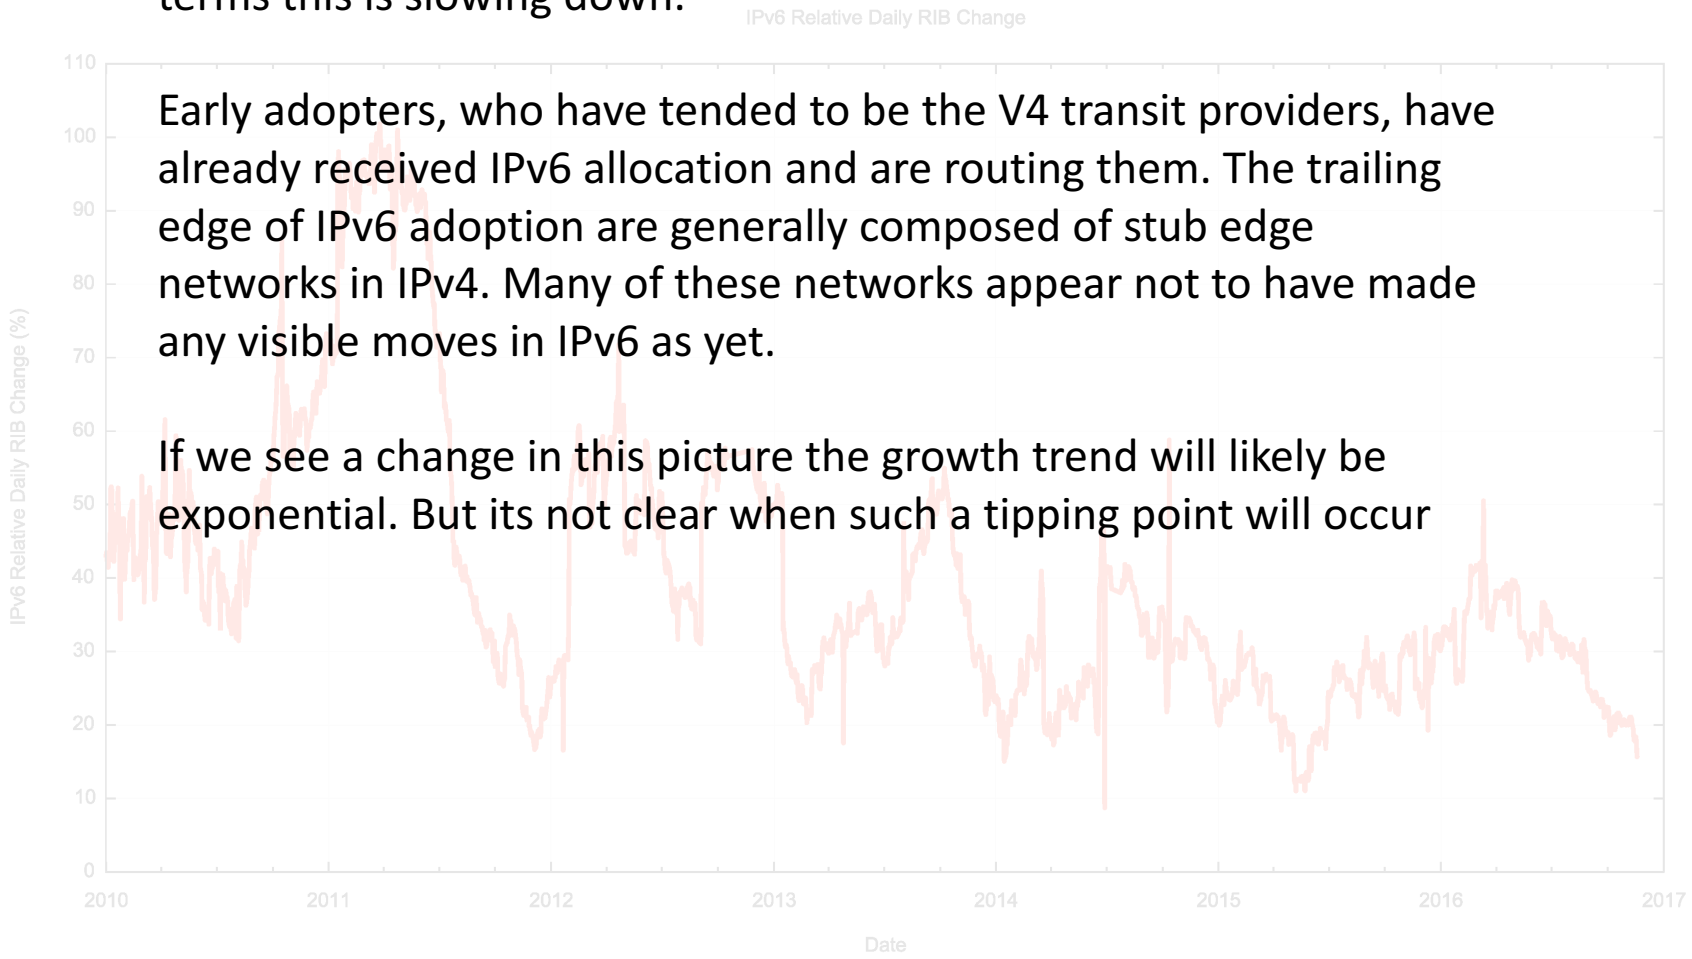

Early adopters, who have tended to be the V4 transit providers, have already received IPv6 allocation and are routing them. The trailing edge of IPv6 adoption are generally composed of stub edge networks in IPv4. Many of these networks appear not to have made any visible moves in IPv6 as yet.

If we see a change in this picture the growth trend will likely be exponential. But its not clear when such a tipping point will occur

IPv6 BGP Table Size predictions

		Exponential Model	Linear Model
Jan 2014	16,100		
2015	21,200		
2016	*27,000*		
2017	*35,000*		
2018		*50,000*	*43,000*
2019		*65,000*	*51,000*
2020		*86,000*	*59,000*
2021		*113,000*	*67,000*
2022		*150,000*	*75,000*

Range of potential outcomes

BGP Table Growth

Nothing in these figures suggests that there is cause for urgent alarm -- at present

- The overall eBGP growth rates for IPv4 are holding at a modest level, and the IPv6 table, although it is growing at a faster relative rate, is still small in size in absolute terms
- 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
- Nothing is melting in terms of the size of the routing table as yet

BGP Updates

- What about the level of updates in BGP?
- Let's look at the update load from a single eBGP feed in a DFZ context

IPv4 Announcements and Withdrawals

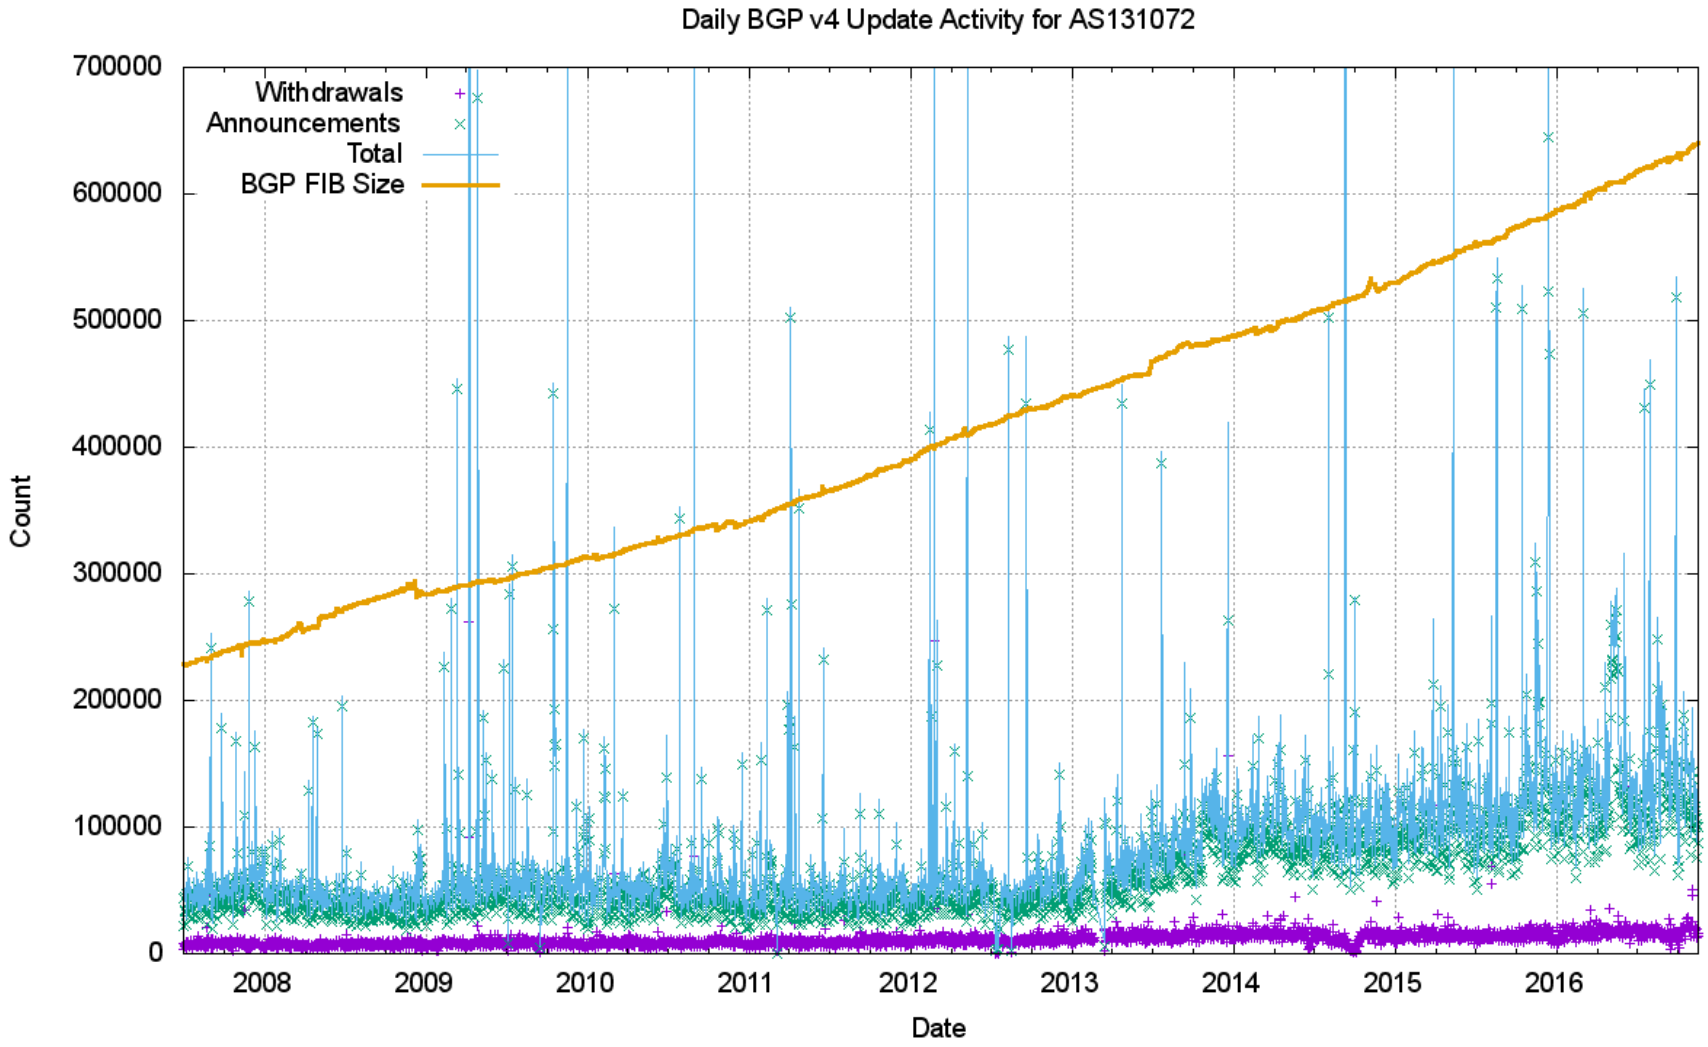

IPv4 Convergence Performance

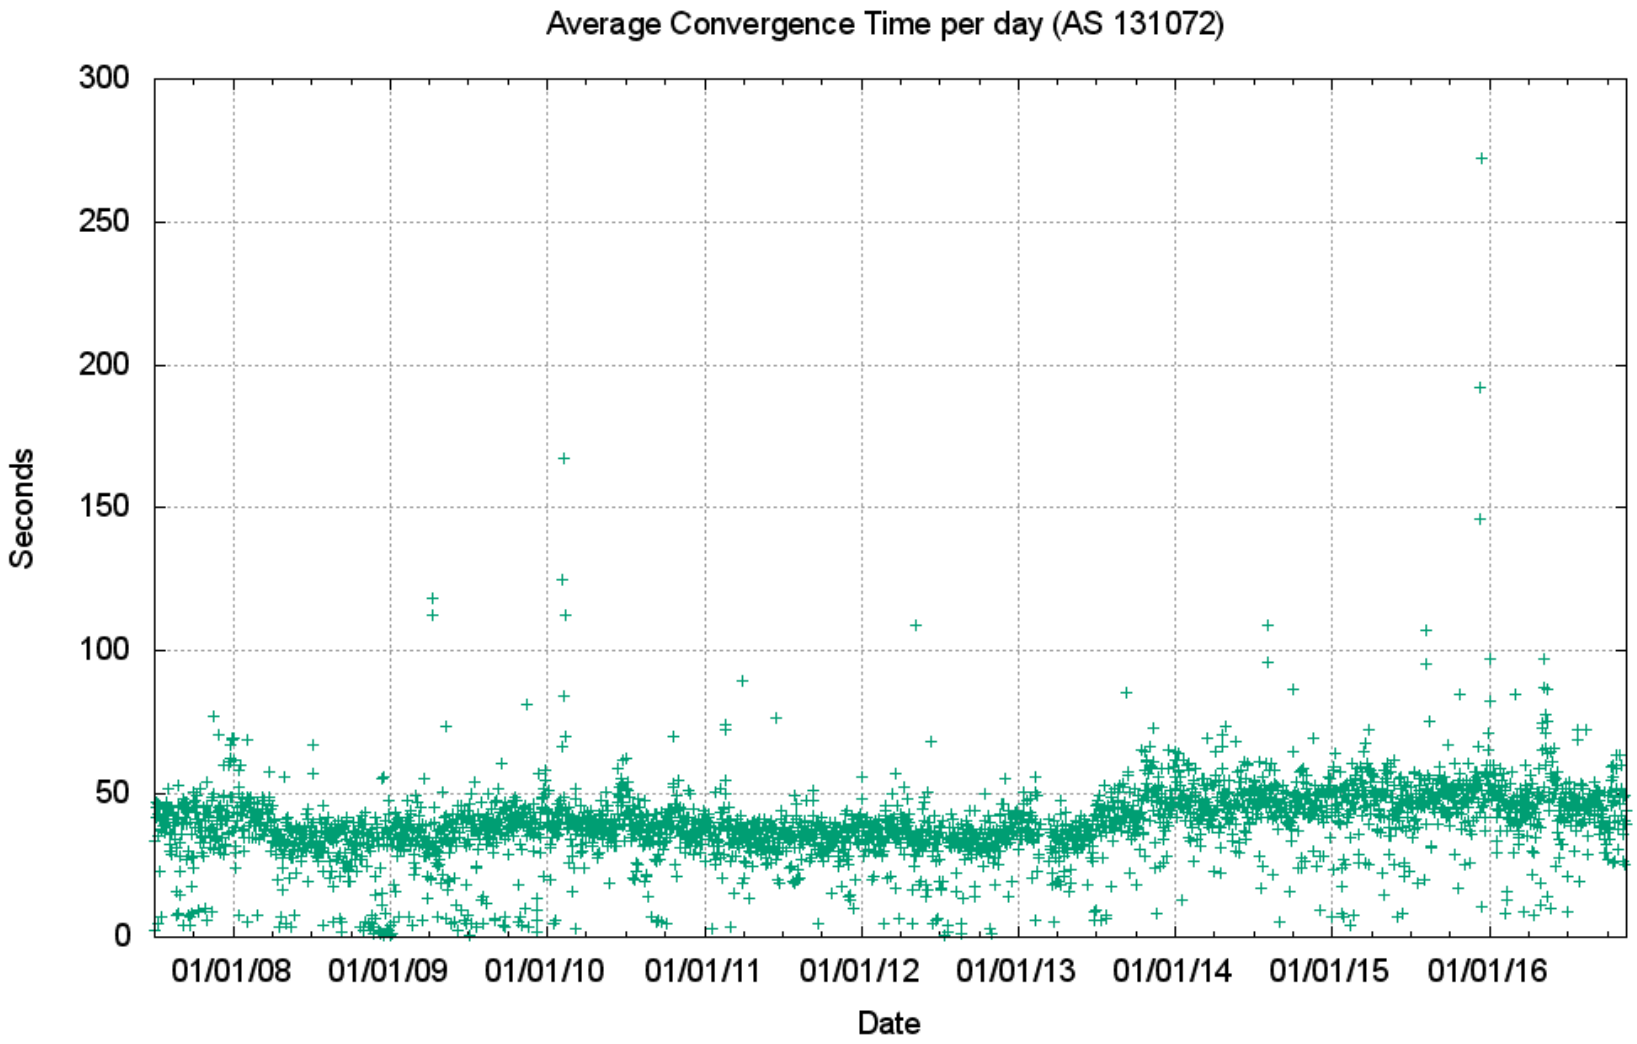

Updates in IPv4 BGP

The number of updates per instability event has been relatively constant

Which is good, but why is this happening?

Likely contributors to this outcome are the damping effect of widespread use of the MRAI interval by eBGP speakers, and the topology factors, as seen in the relatively constant V4 AS Path Length

V6 Announcements and Withdrawals

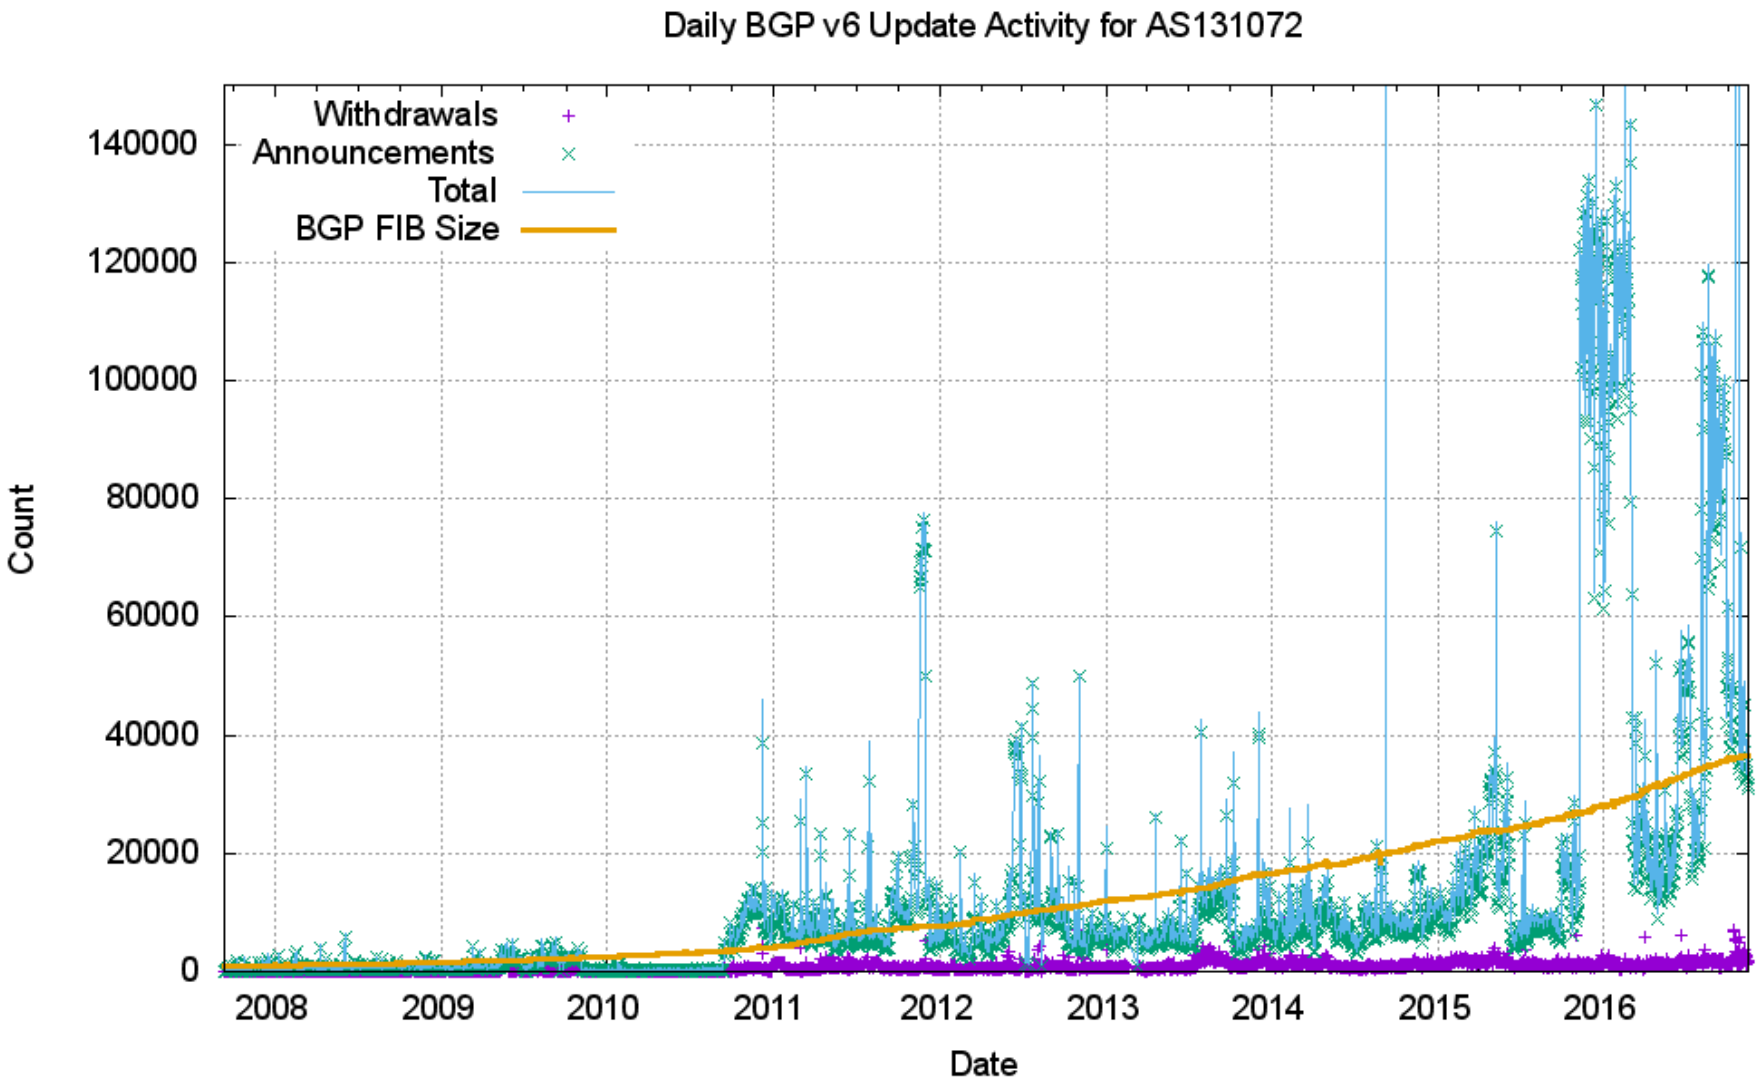

V6 Convergence Performance

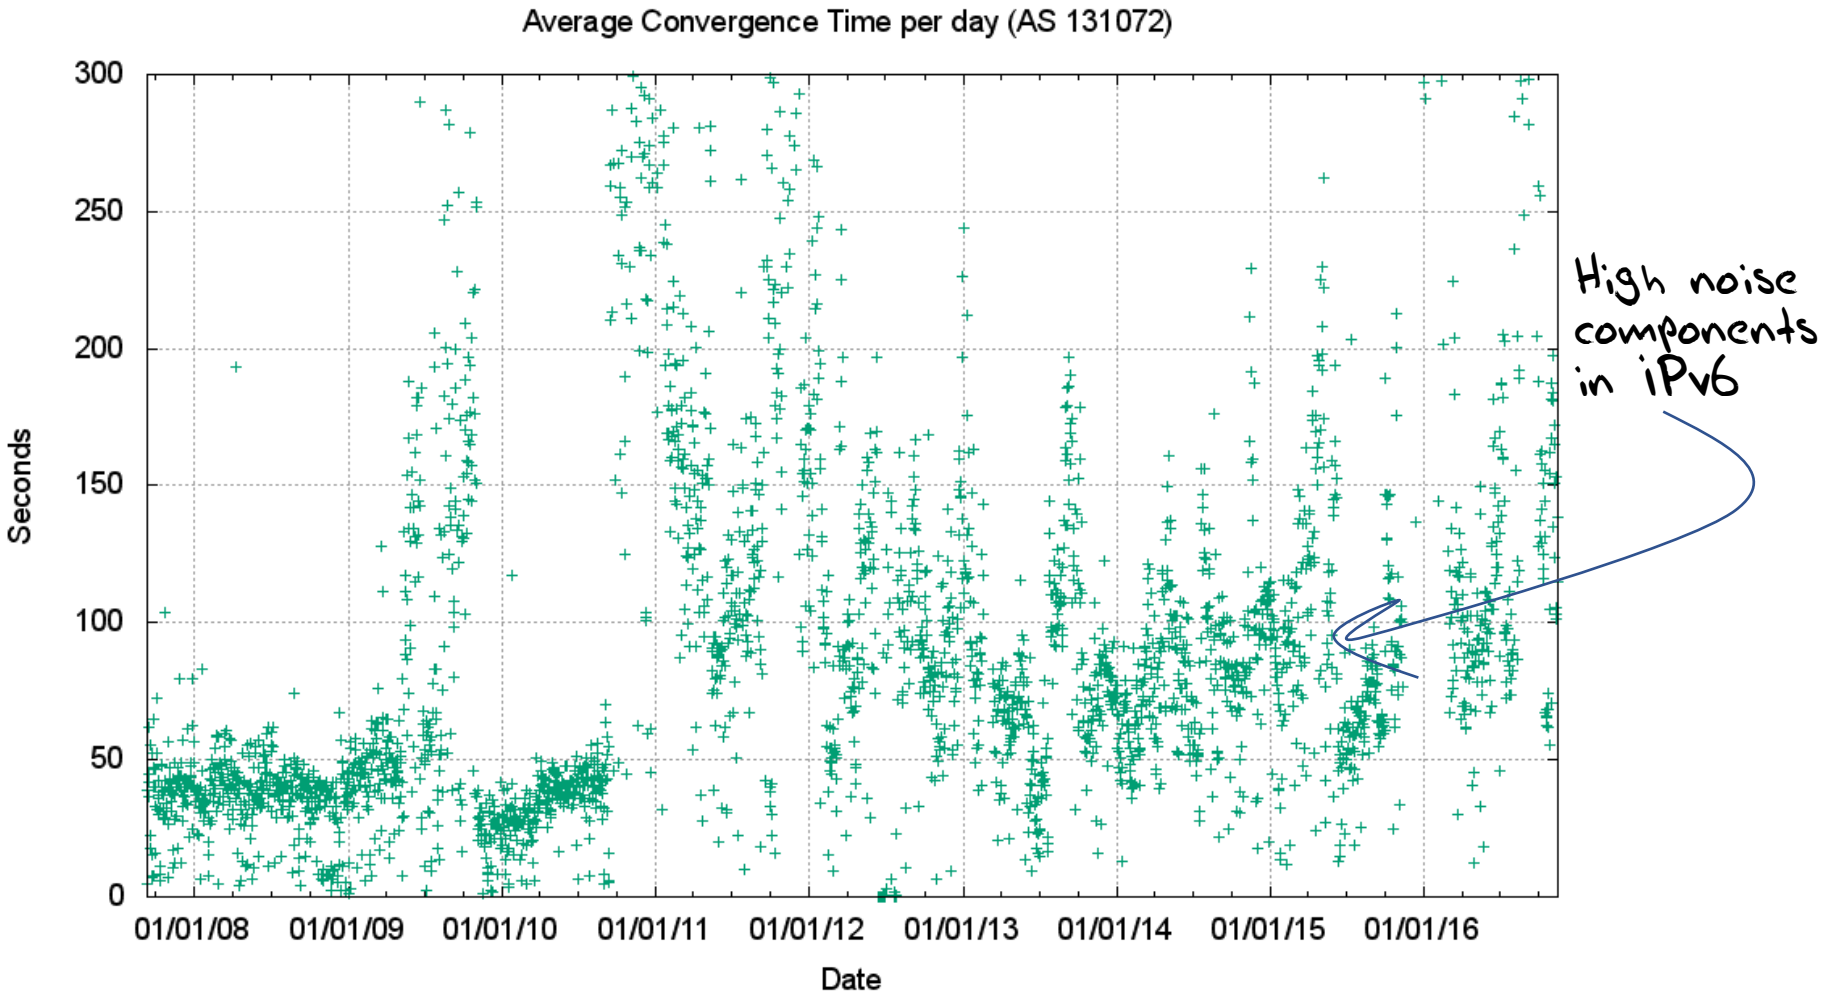

V6 Updated prefixes per day

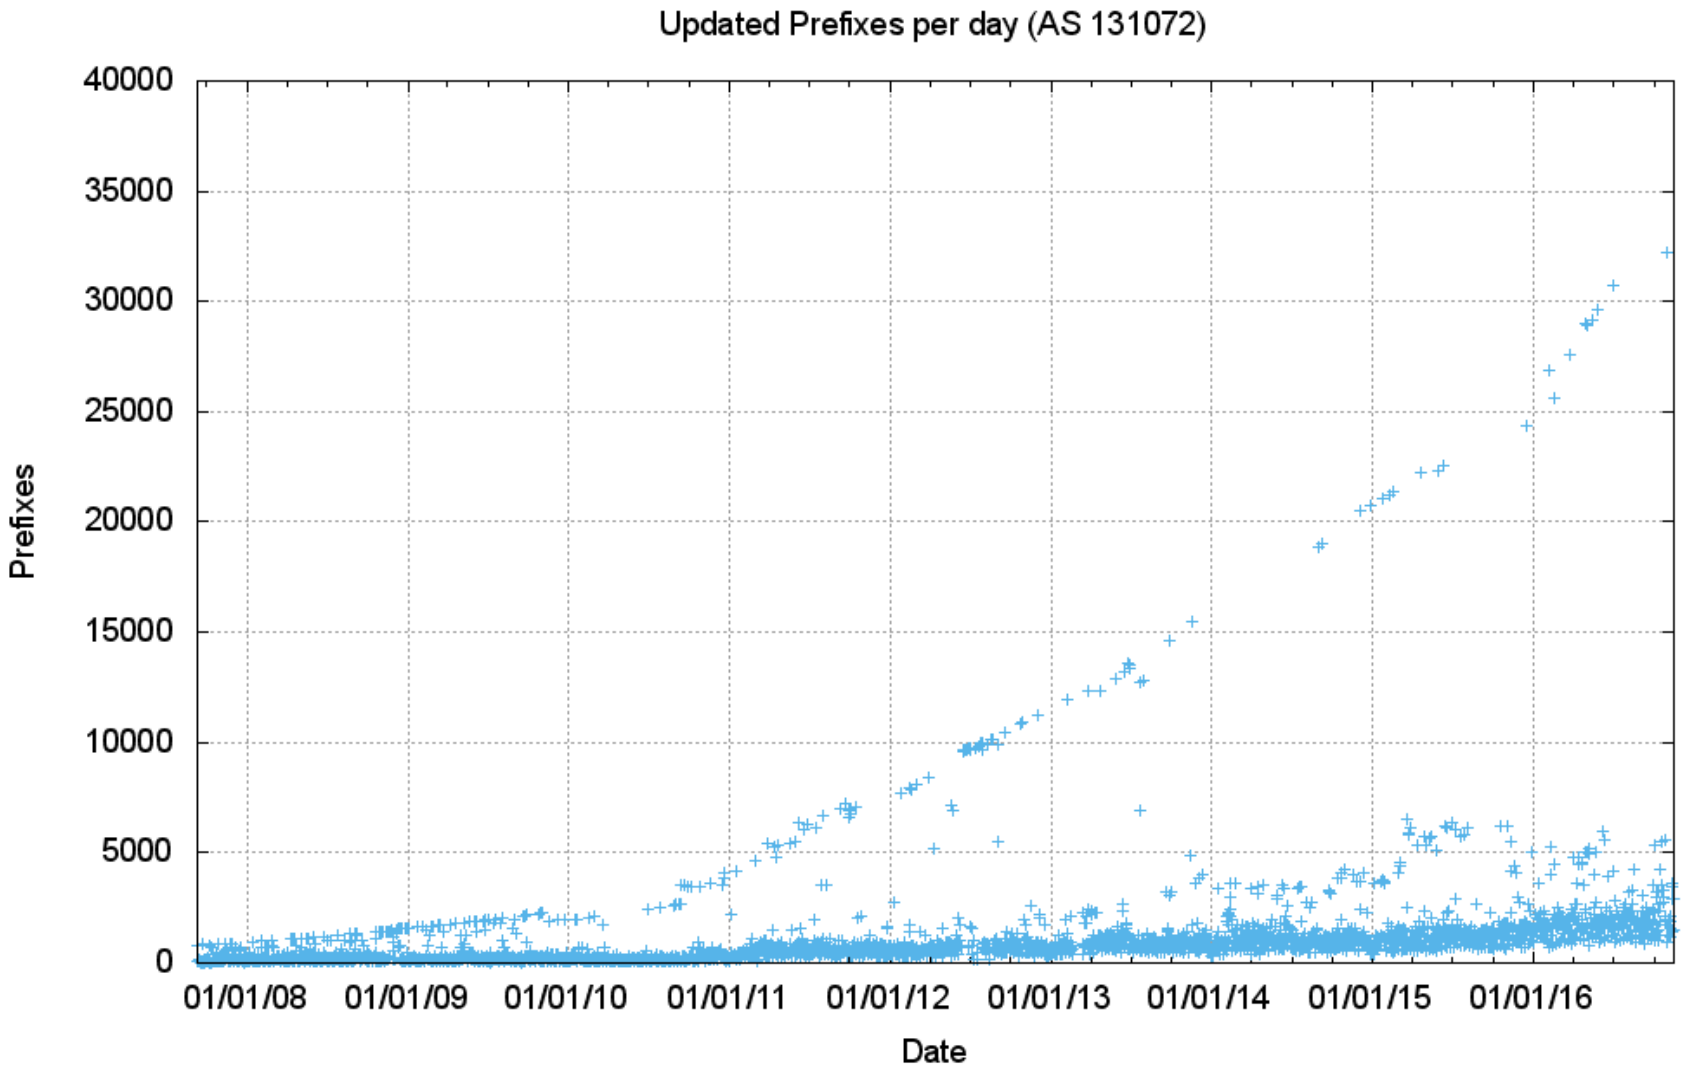

V6 Updates per event

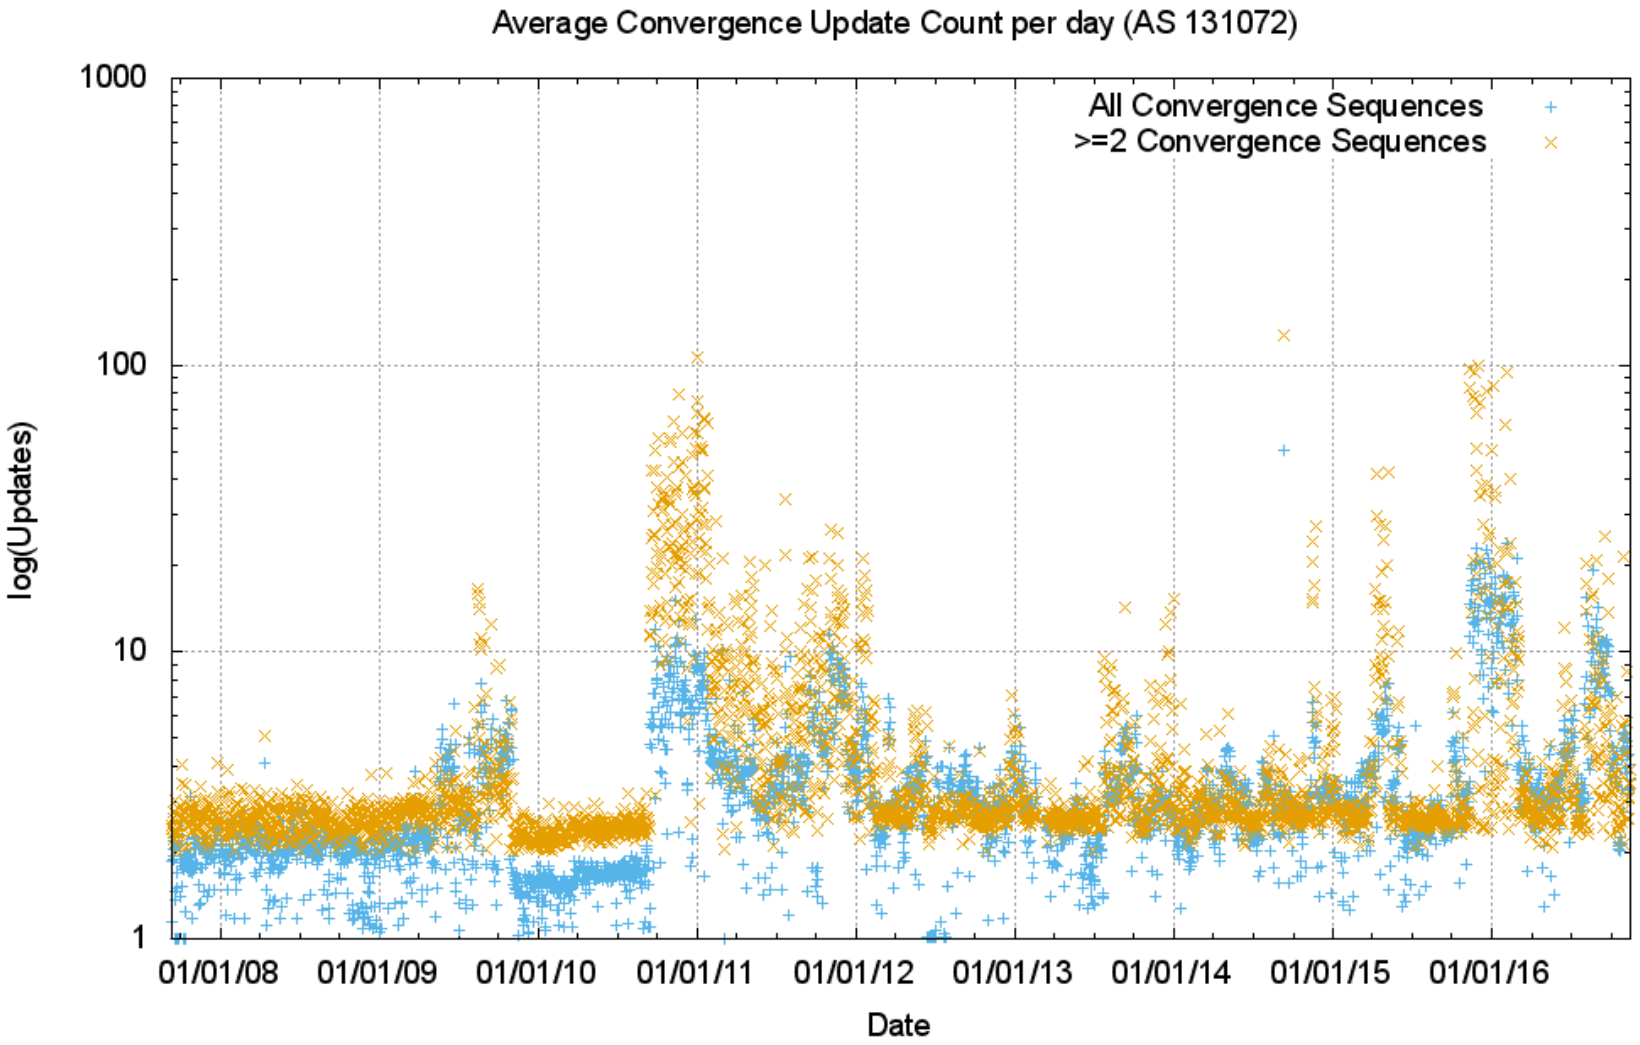

Updates in IPv6

BGP Route Updates are very unequally distributed across the prefix set – they appear to affect a very small number of prefixes which stand out well above the average

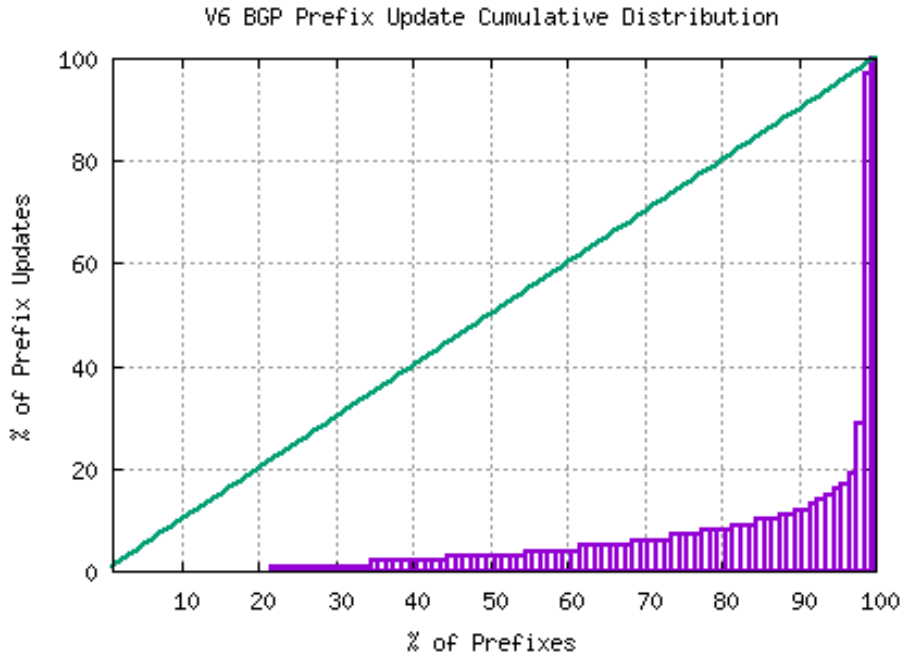

Updates in IPv6

50 Most active Prefixes for the past 31 days

RANK	PREFIX	UPDs	%	Origin AS -- AS NAME
1	2001:e00::/31	49904	2.95%	4795 -- INDOSATM2-ID INDOSATM2 ASN, ID
2	2402:f080::/48	49823	2.95%	38150 -- TELNET-AS-ID PT. TIME EXCELINDO, ID
3	2402:f080:1::/48	49821	2.95%	38150 -- TELNET-AS-ID PT. TIME EXCELINDO, ID
4	2403:ae00::/32	49801	2.95%	38766 -- BMP-AS-ID PT. Bumi Merbabu Permai, ID
5	2400:8b00:a001::/48	40699	2.41%	45727 -- THREE-AS-ID Hutchison CP Telecommunications, PT, ID
6	2001:df5:b400::/48	38068	2.25%	131735 -- IDNIC-TNC-ID PT Telemedia Network Cakrawala, ID
7	2403:8000::/32	36252	2.14%	4796 -- BANDUNG-NET-AS-AP Institute of Technology Bandung, ID
8	2804:14d:baa2::/48	29735	1.76%	28573 -- CLARO S.A., BR
9	2a00:86c0:98::/48	27722	1.64%	2906 -- AS-SSI - Netflix Streaming Services Inc., US
10	2a00:86c0:99::/48	27614	1.63%	2906 -- AS-SSI - Netflix Streaming Services Inc., US
11	2402:ab00:140::/48	27160	1.61%	24206 -- CHANNEL11-AS-ID PT Cakra Lintas Nusantara, ID
12	2620:10c:7007::/48	26703	1.58%	2906 -- AS-SSI - Netflix Streaming Services Inc., US
13	2406:9600::/32	25991	1.54%	10208 -- THENET-AS-ID-AP PT. Millenium Internetindo, ID
14	2402:a600::/32	25392	1.50%	17996 -- UIINET-ID-AP PT Global Prima Utama, ID
15	2a06:f6c0::/29	21149	1.25%	50107 -- AS_TCPCLOUD , CZ
16	2804:14d:90af::/48	19537	1.16%	28573 -- CLARO S.A., BR
17	2804:14d:9081::/48	19536	1.16%	28573 -- CLARO S.A., BR
18	2804:14d:90ae::/48	19501	1.15%	28573 -- CLARO S.A., BR
19	2804:14d:90ad::/48	19134	1.13%	28573 -- CLARO S.A., BR
20	2804:14d:90a6::/48	18939	1.12%	28573 -- CLARO S.A., BR

The busiest 50 IPv6 prefixes accounted for 1/2 of all BGP IPv6 prefix updates

http://bgpupdates.potaroo.net/instability/v6-bgpupd.html

Compared to IPv4

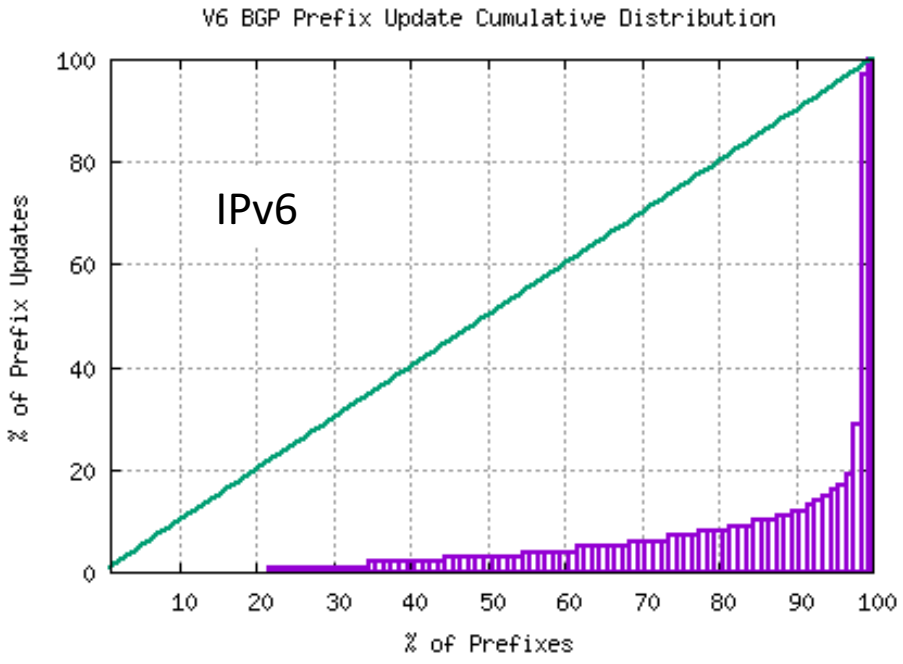

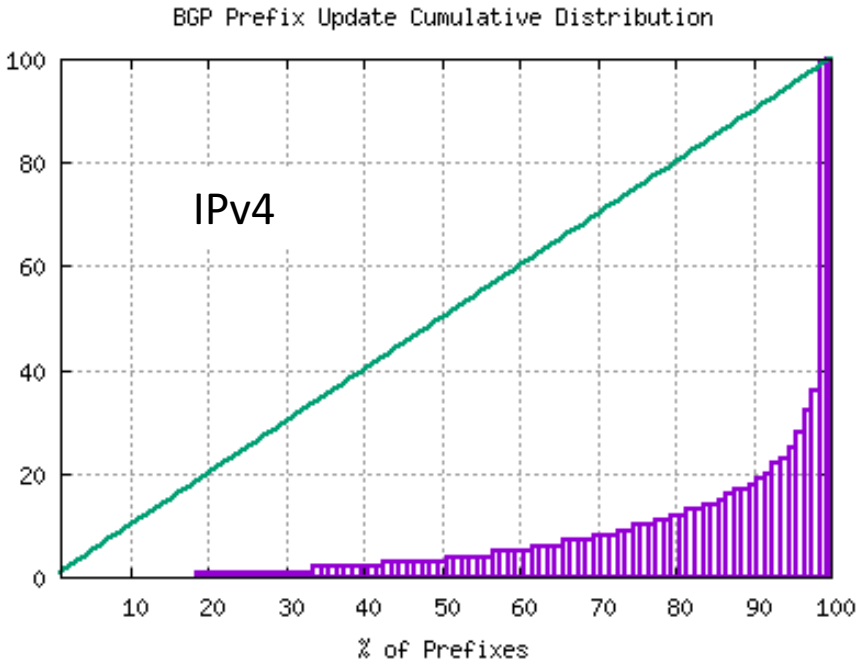

http://bgpupdates.potaroo.net/instability/v6-bgpupd.html

http://bgpupdates.potaroo.net/instability/bgpupd.html

Updates in IPv6 BGP

IPv6 routing behaviour is similar to IPv4 behaviour:

Most announced prefixes are stable all of the time

And as more prefixes are announced, most of these announced prefixes are highly stable.

But for a small number of prefixes we observe highly unstable behaviours that dominate IPv6 BGP updates which appear to be more unstable (relatively) than IPv4

Today's State of Routing

“Mostly Harmless”

The levels of growth of the tables, and the levels of growth of updates in BGP do not pose any immediate concerns

The trends are predictable and steady, so network operators can plan well in advance for the capacity of routing equipment to meet their future needs

But:

The advanced levels of instability by a small number of networks are always annoying! How can we prevent these highly unstable prefixes?

That's it!

Questions?