Routing 2015

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…

Through the Routing Lens

There is no Routing God!

There is no single objective “out of the system” view of the Internet’s Routing environment.

BGP distributes a routing view that is modified as it is distributed, so every eBGP speaker will see a slightly different set of prefixes, and each view is relative to a given location

So the picture I will be painting here is one that is drawn from the perspective of AS131072. This is a stub AS at edge of the Internet, and this is an eBGP view.

You may have a similar view from your network.

20 Years of Routing the Internet

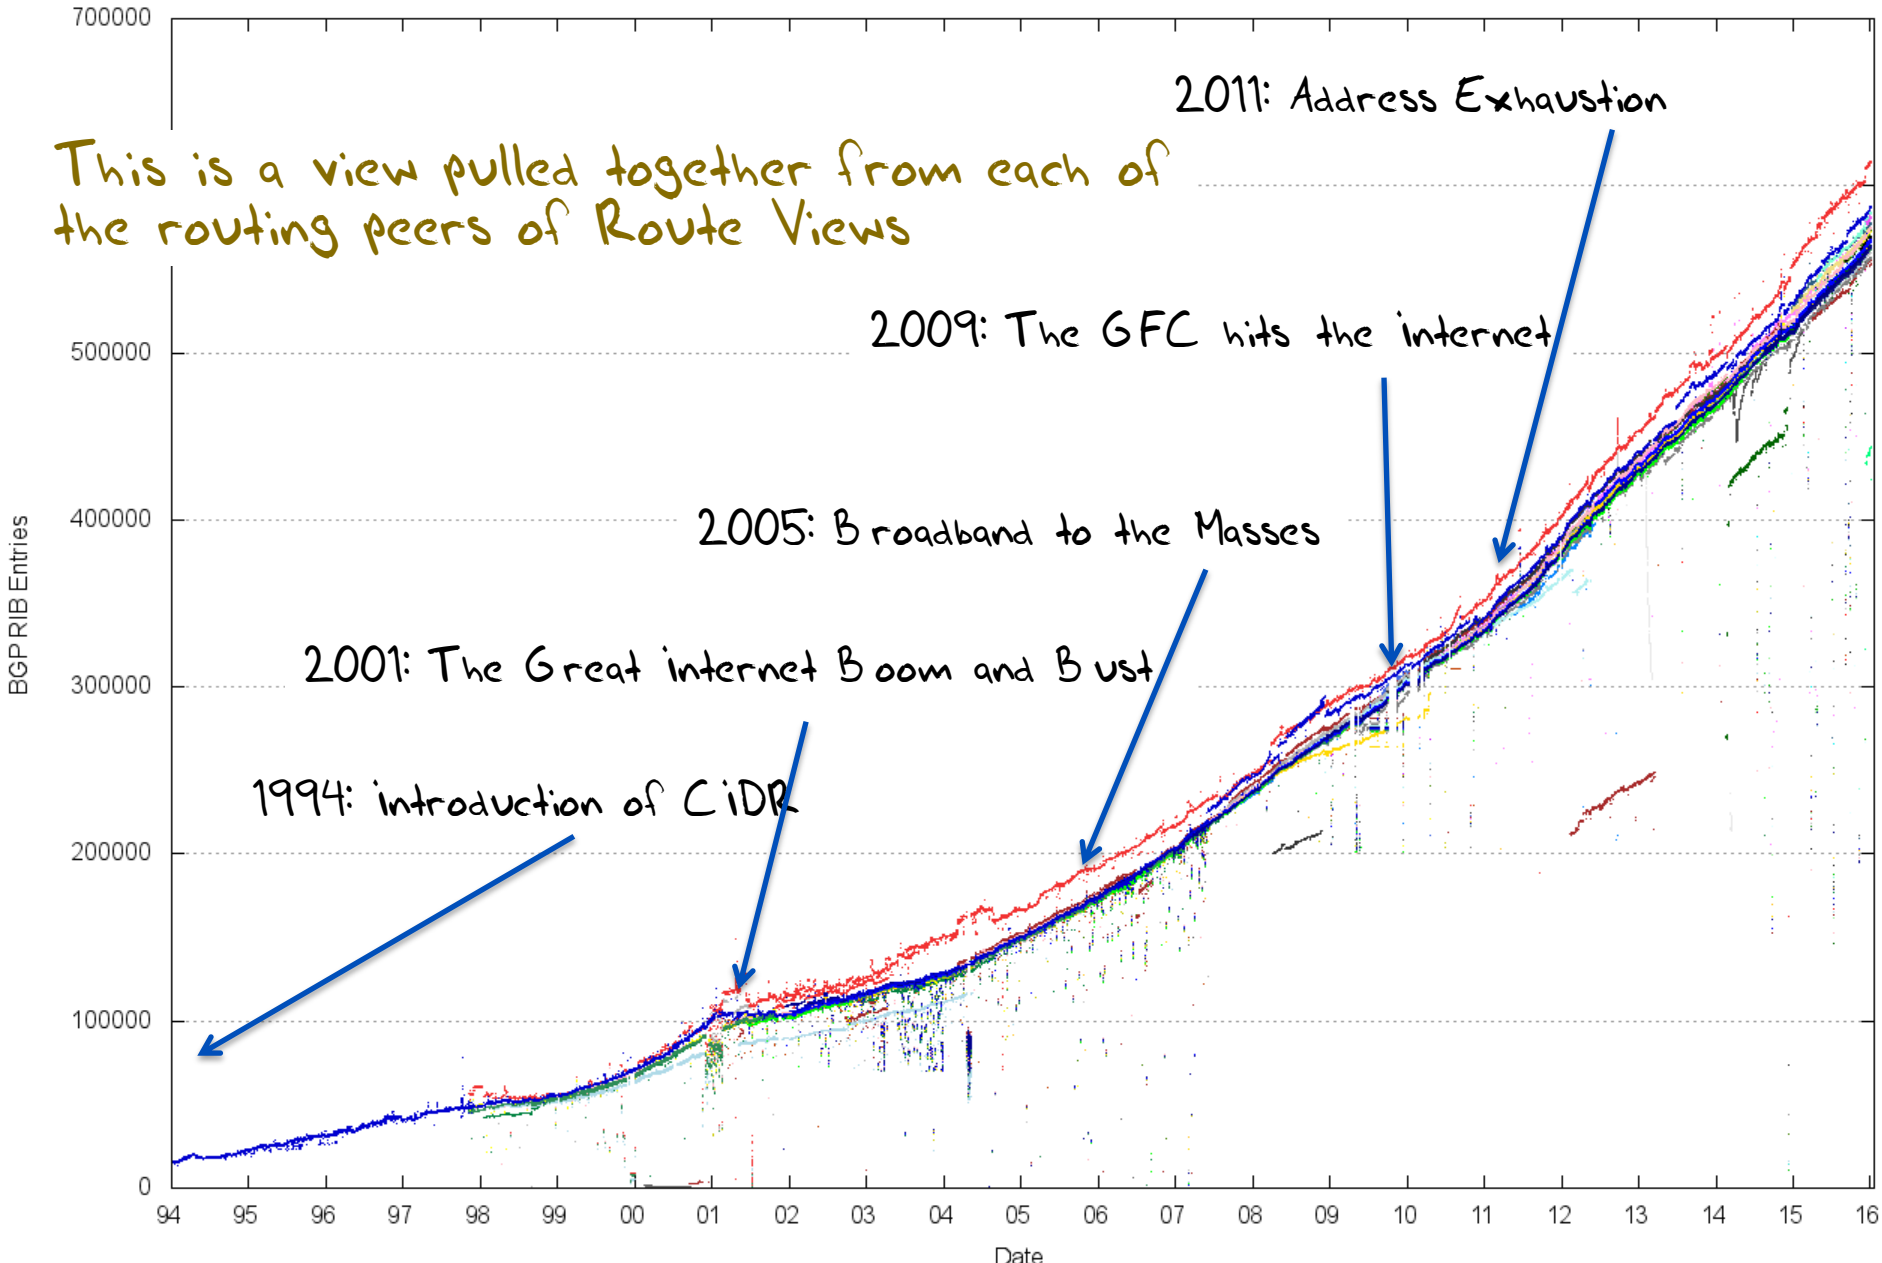

20 Years of Routing the Internet

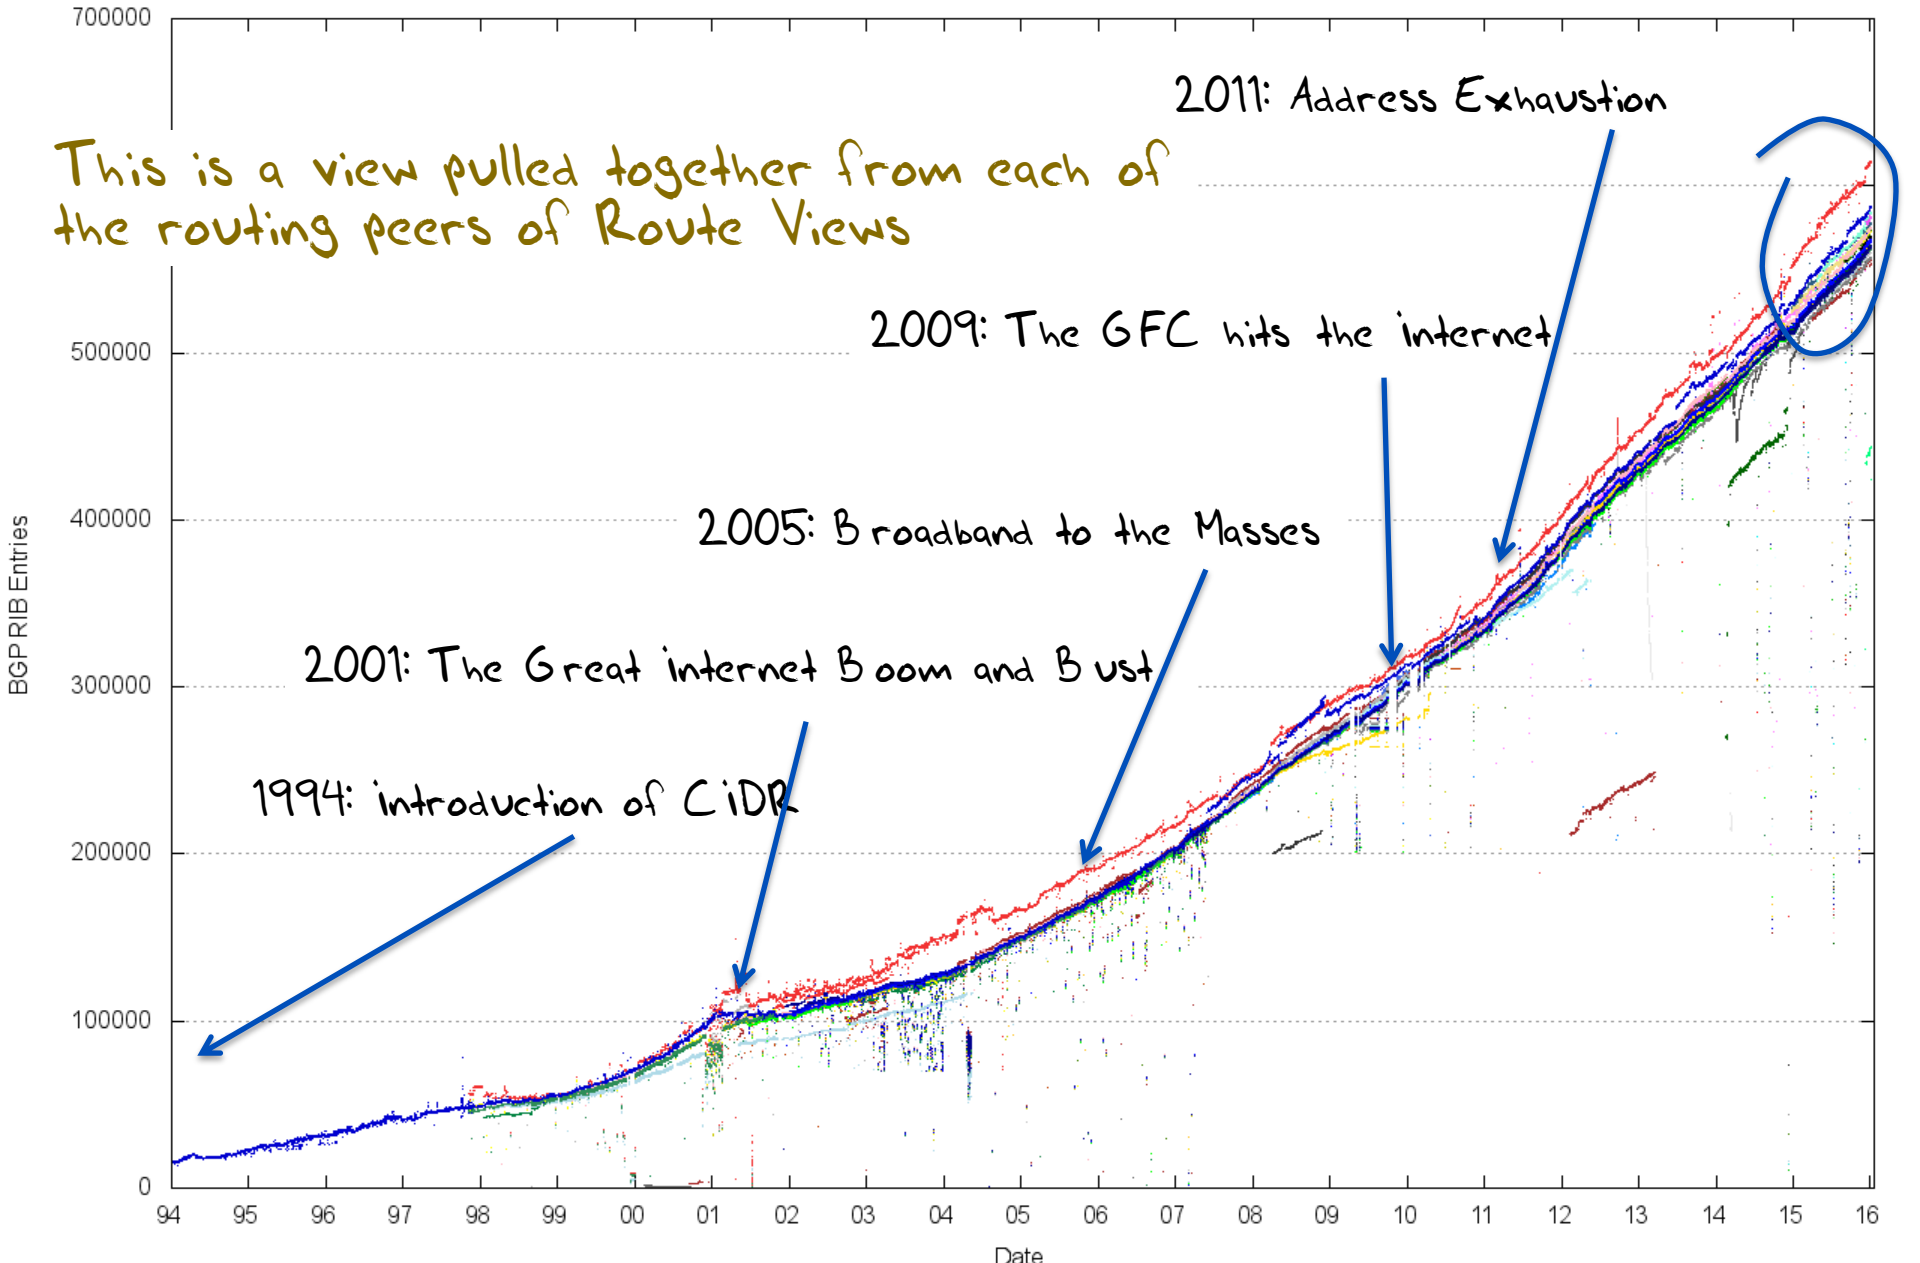

2015, as seen at Route Views

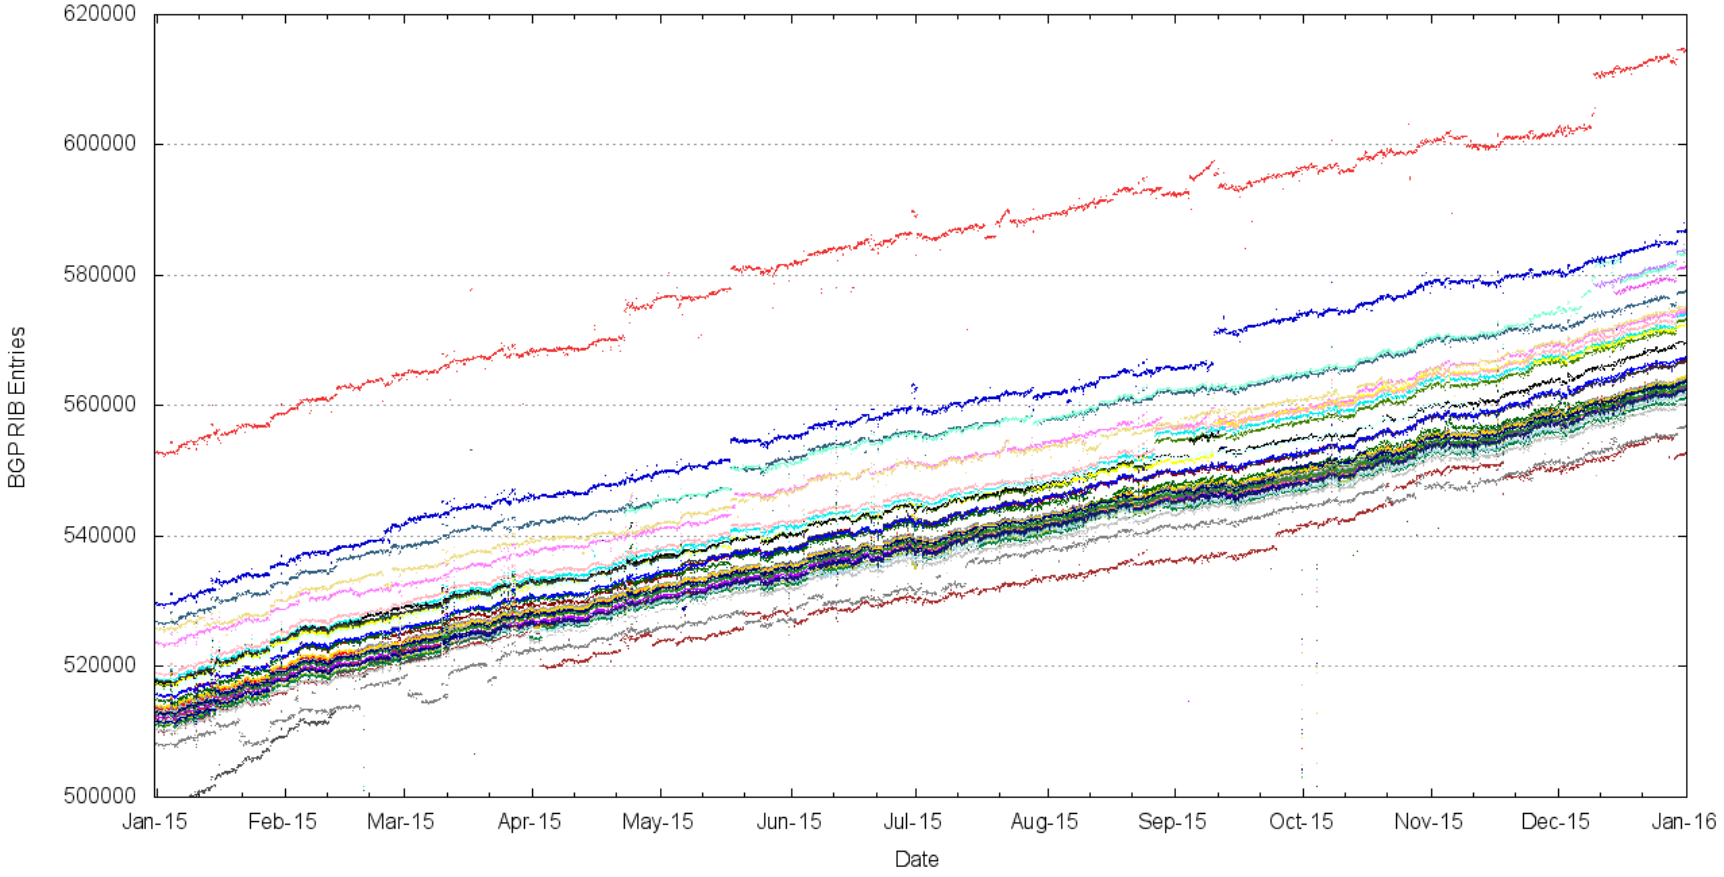

2015, as seen at Route Views

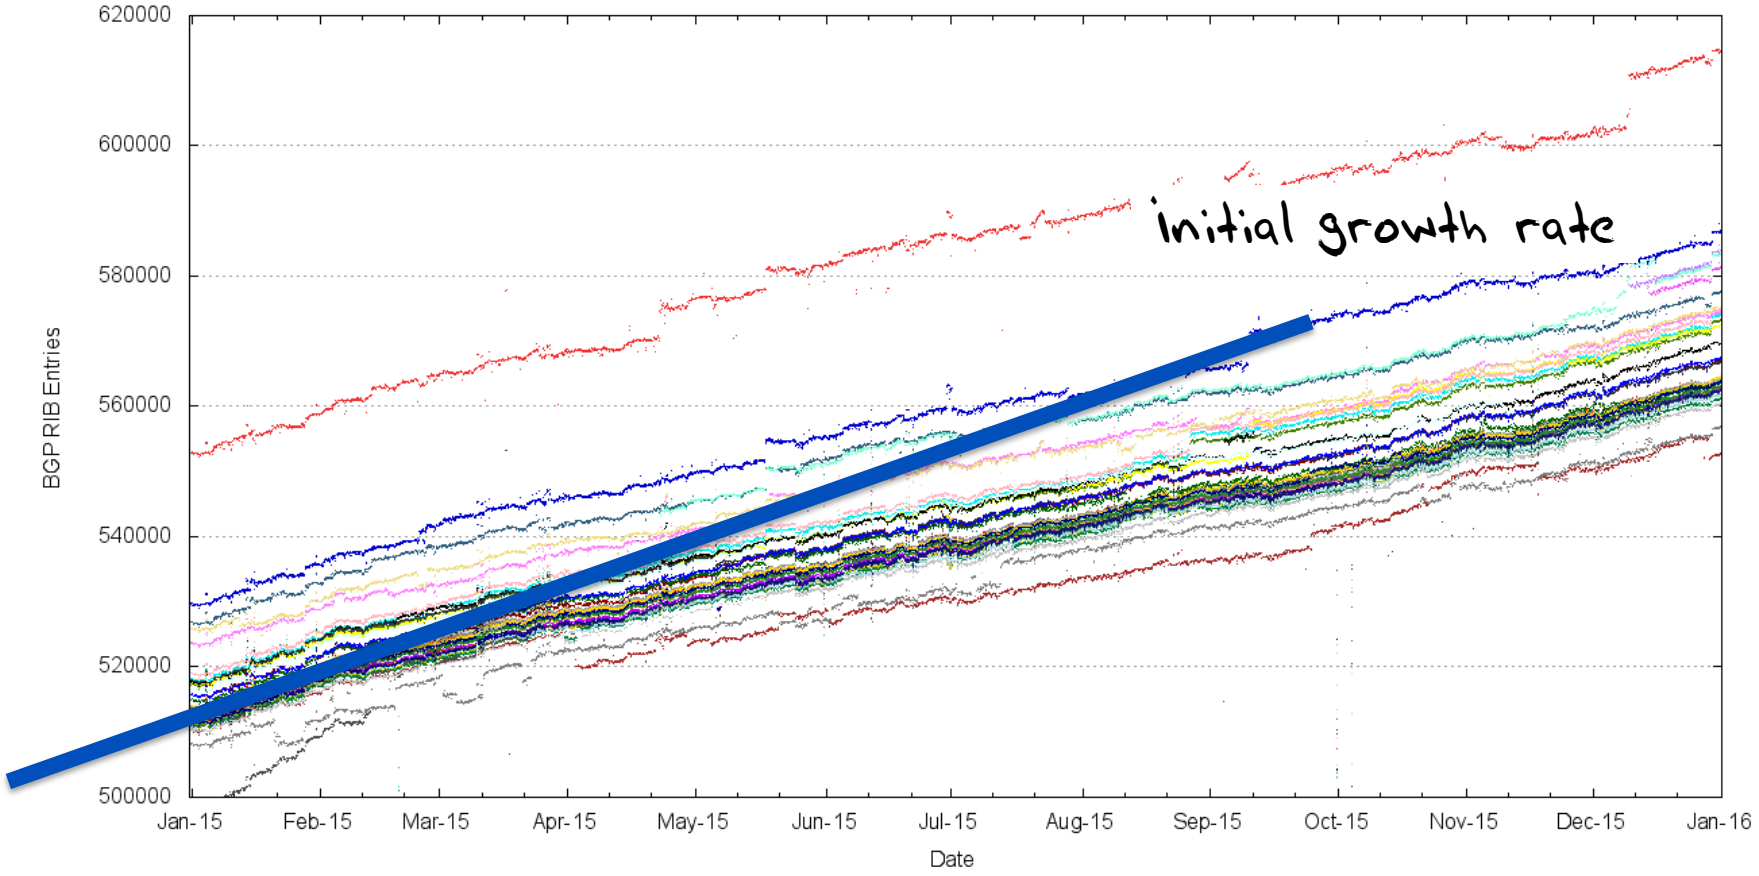

2015, as seen at Route Views

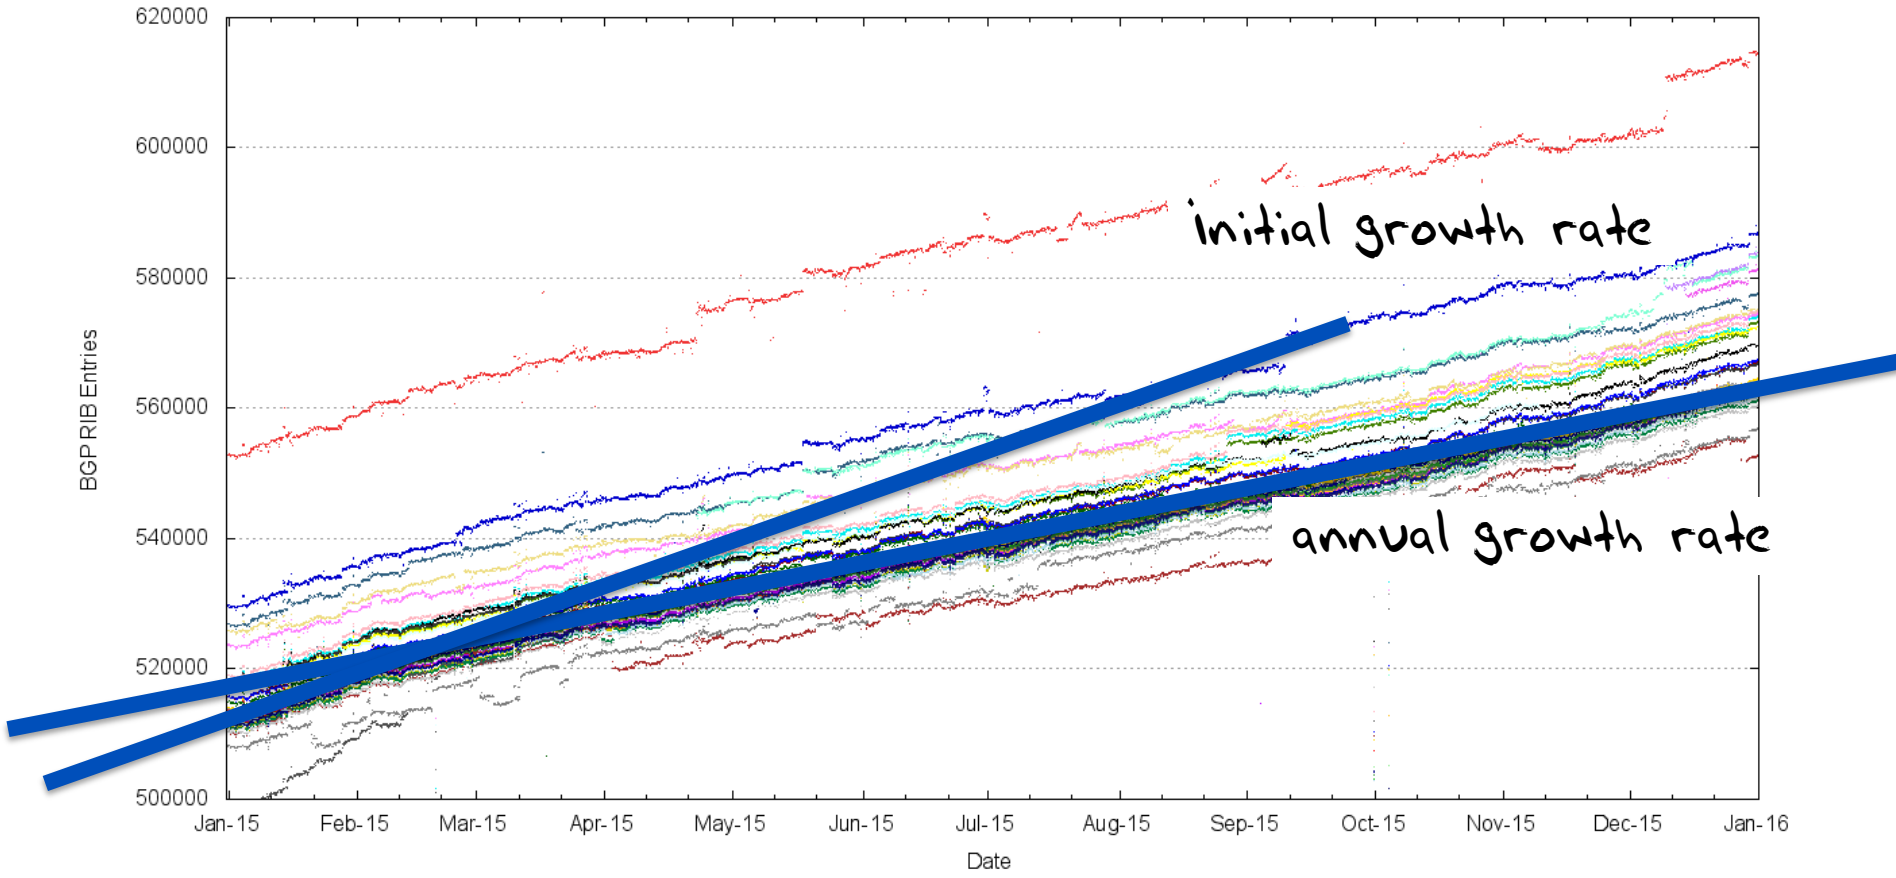

Routing Indicators for IPv4

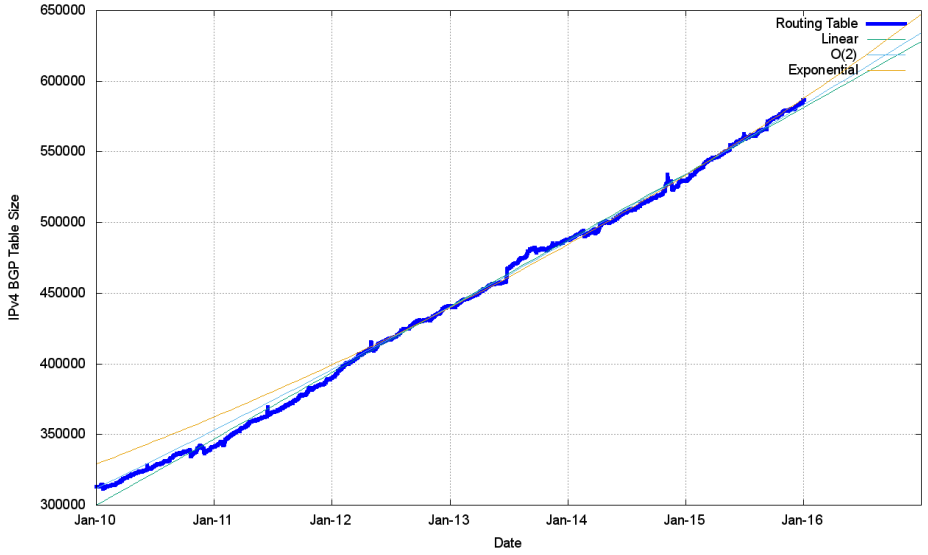

Routing prefixes – growing by some 47,000 prefixes per year

AS Numbers– growing by some 3,100 prefixes per year

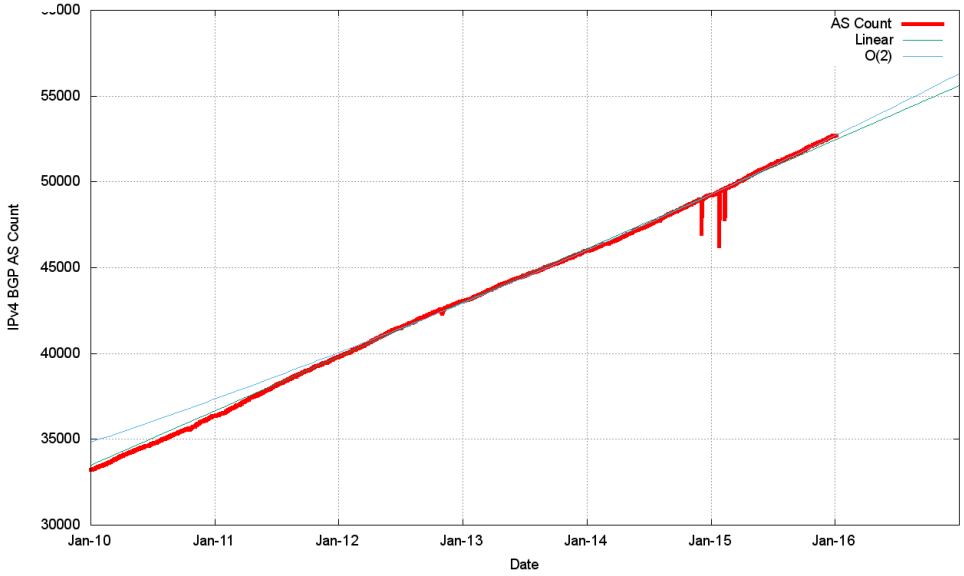

Routing Indicators for IPv4

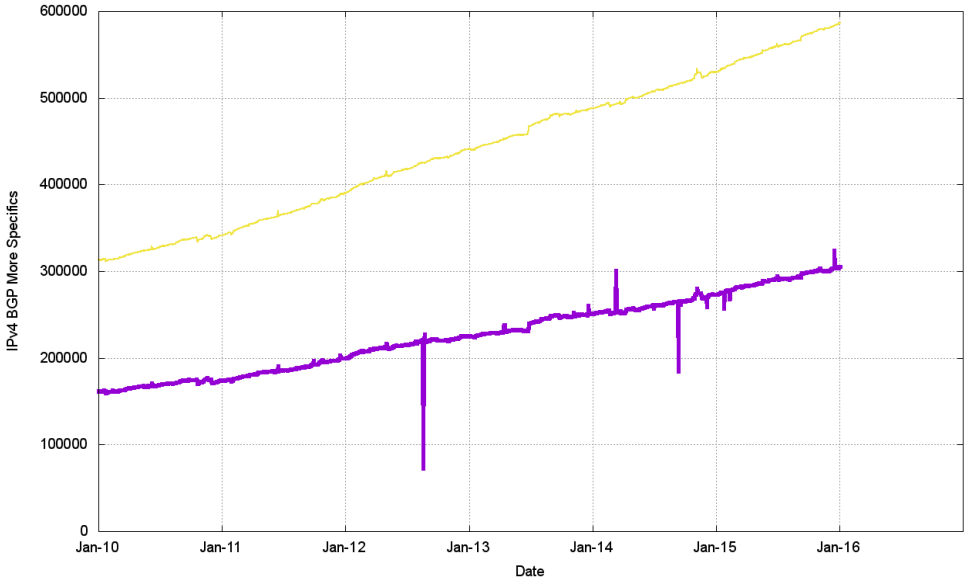

More Specifics are still taking up one half of the routing table

But the average size of a routing advertisement is getting smaller

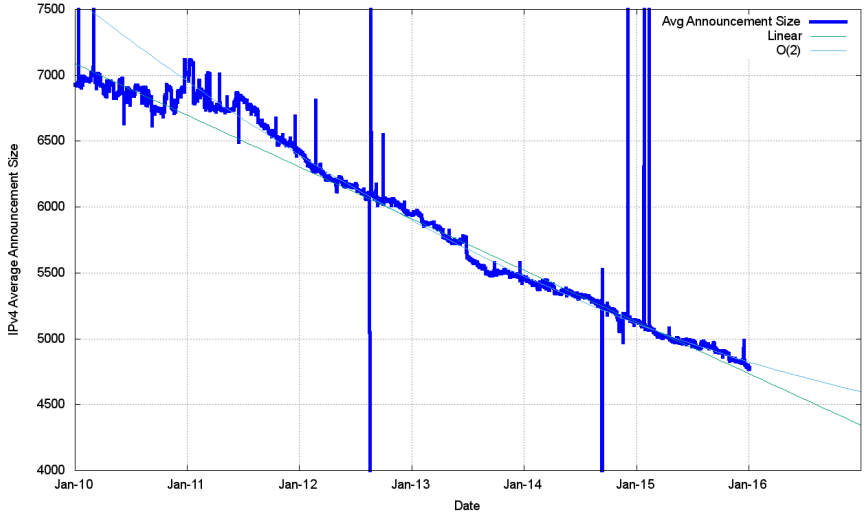

Routing Indicators for IPv4

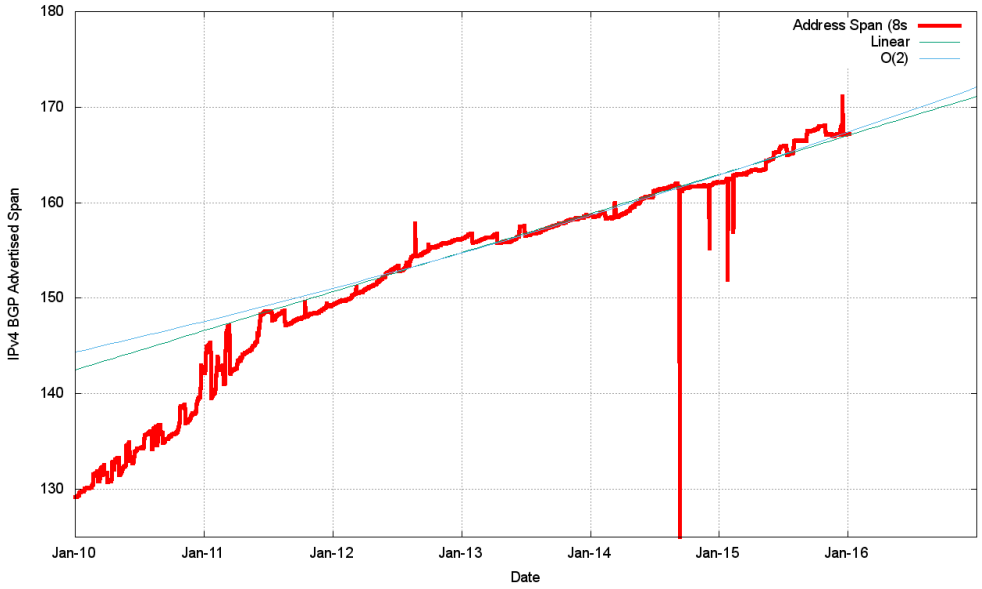

Address Exhaustion is now visible in the extent of advertised address space

The "shape" of inter-AS interconnection appears to be relatively steady, as the Average AS Path length has been steady through the year

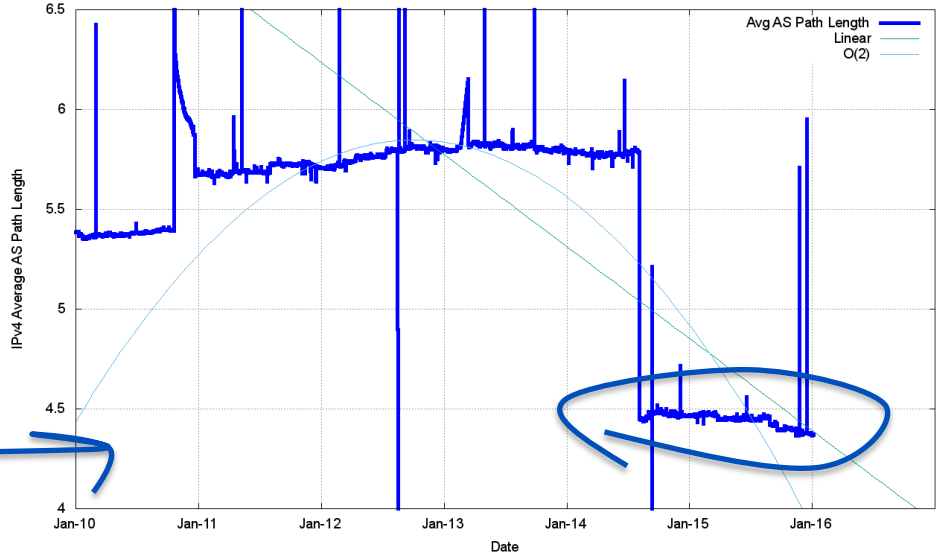

What happened in 2015 in V4?

- From the look of the growth plots, its business as usual, despite the increasing pressure on IPv4 address availability
- The number of entries in the default-free zone is now heading to 600,000
- The pace of growth of the routing table is still relatively constant at ~50,000 new entries per year
  - IPv4 address exhaustion is not changing this!

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

Address "Age"

Address “age” in 2010

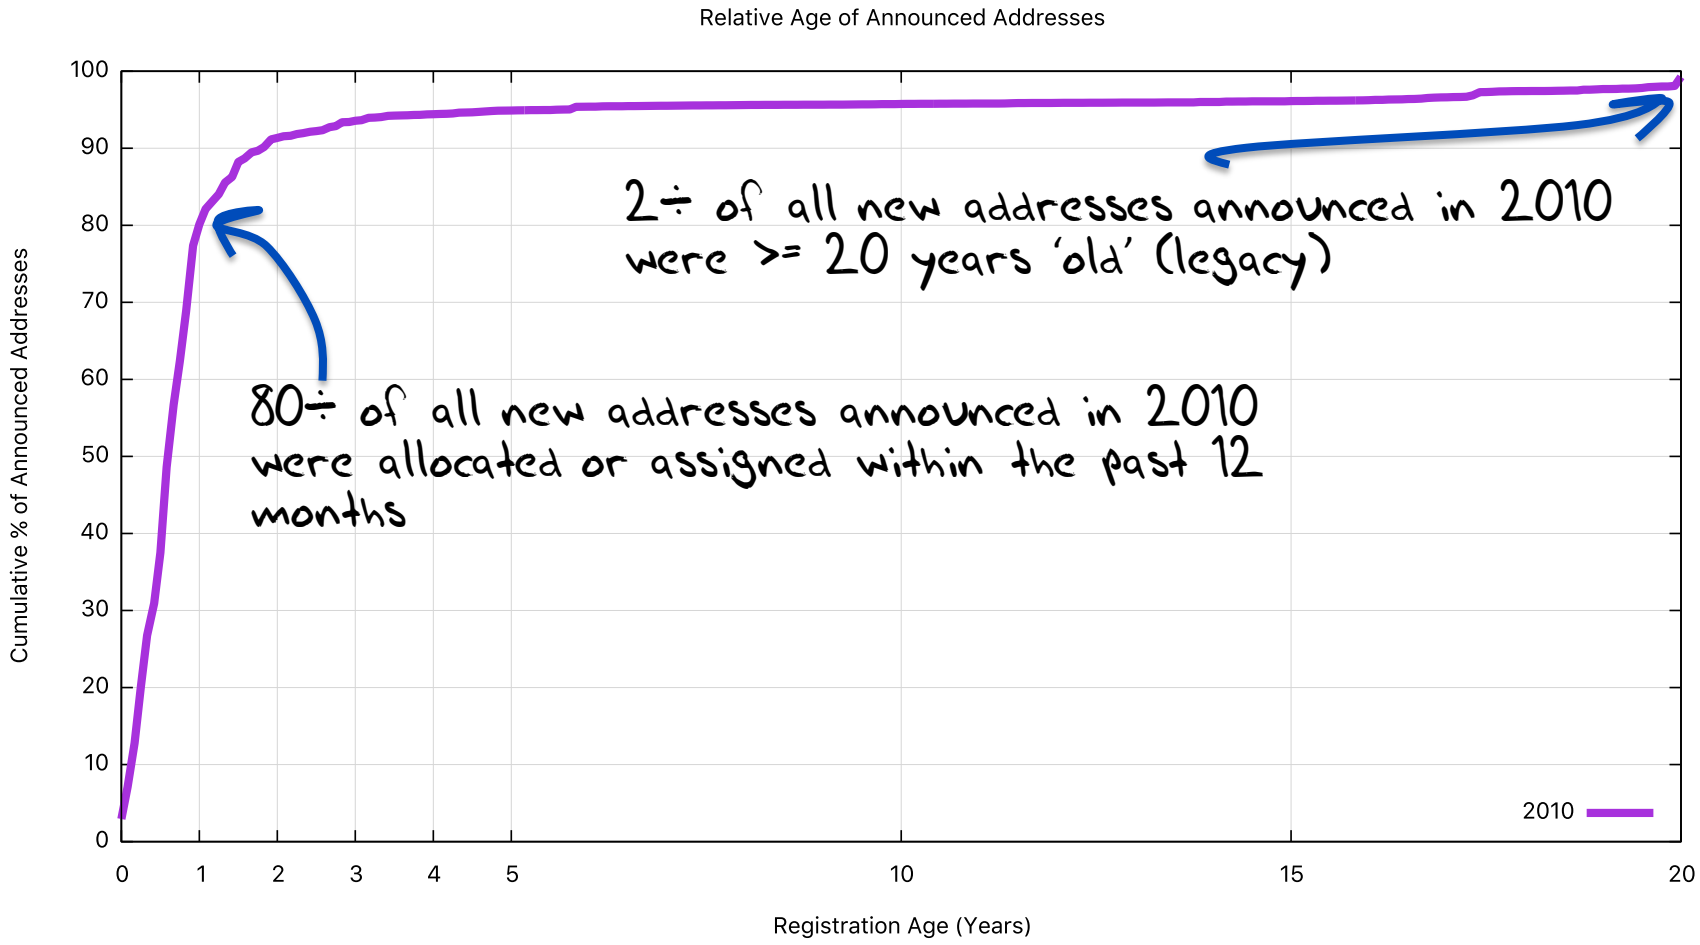

IPv4 Address Reuse

Address “age” in 2015

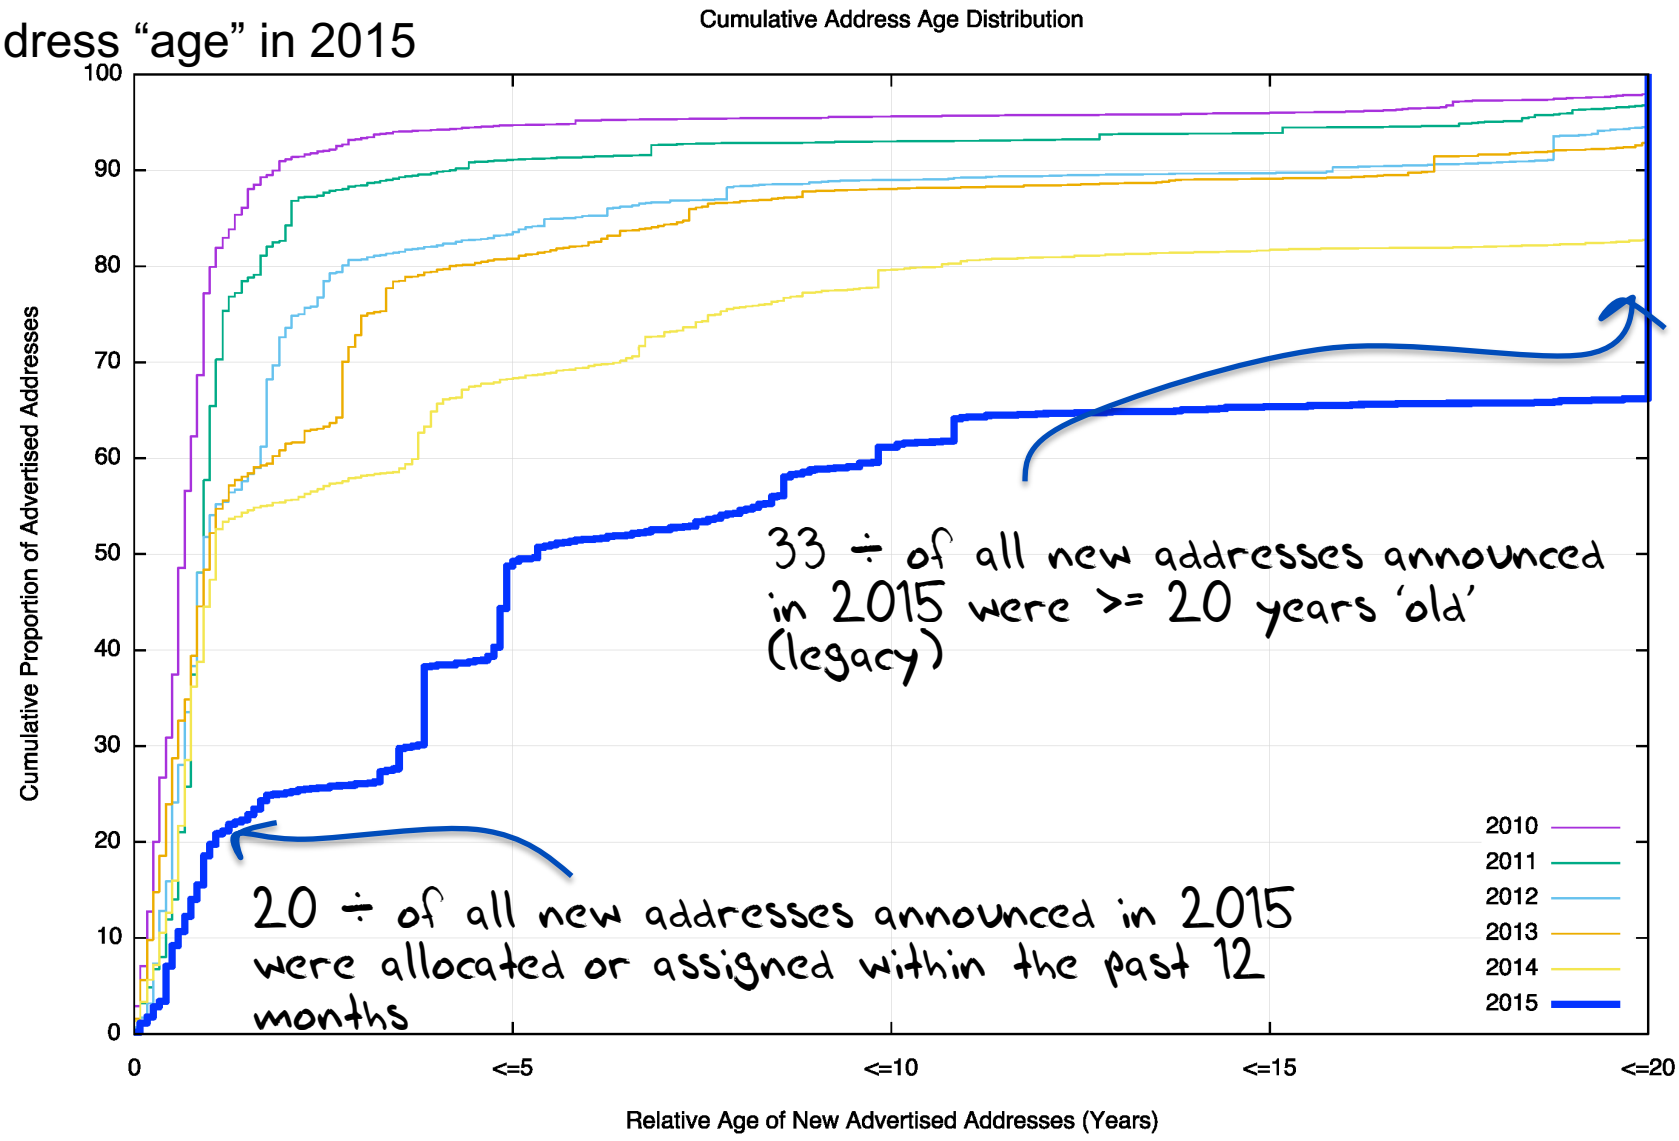

IPv4 in 2015 - Growth is Steady

- Overall IPv4 Internet growth in terms of BGP is at a rate of some **~47,000 entries p.a.**
- But we've run out of the unallocated address pools everywhere except Afrinic
- So what's driving this post-exhaustion growth?
  - Transfers?
  - Last /8 policies in RIPE and APNIC?
  - Leasing and address recovery?

IPv4: Advertised vs Unadvertised Addresses

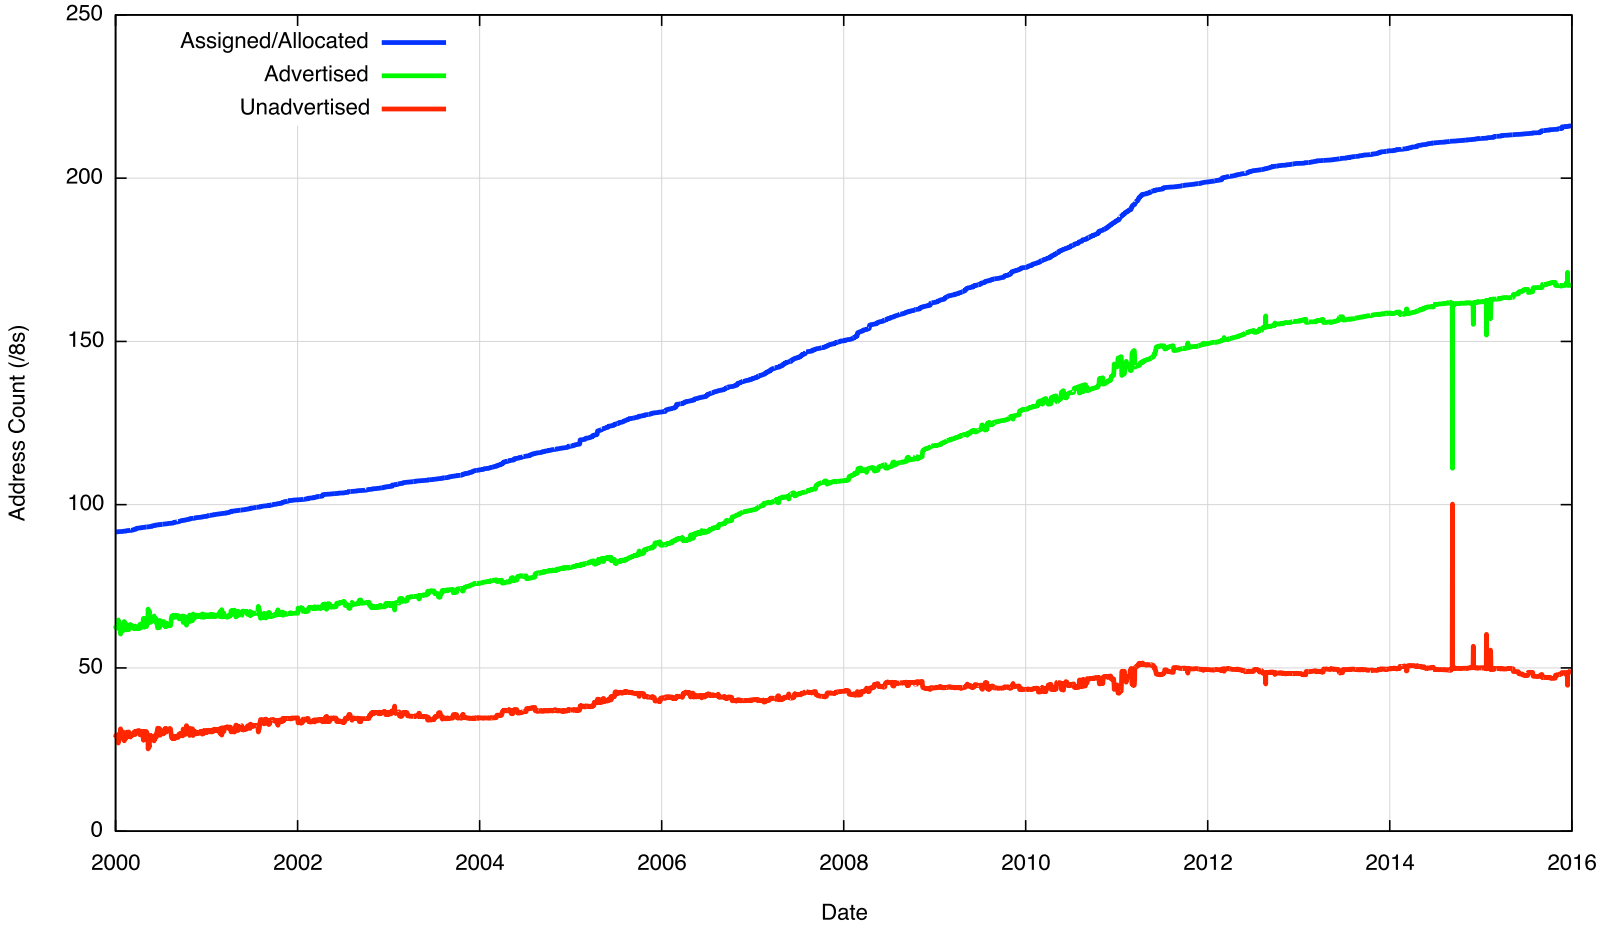

IPv4: Unadvertised Addresses

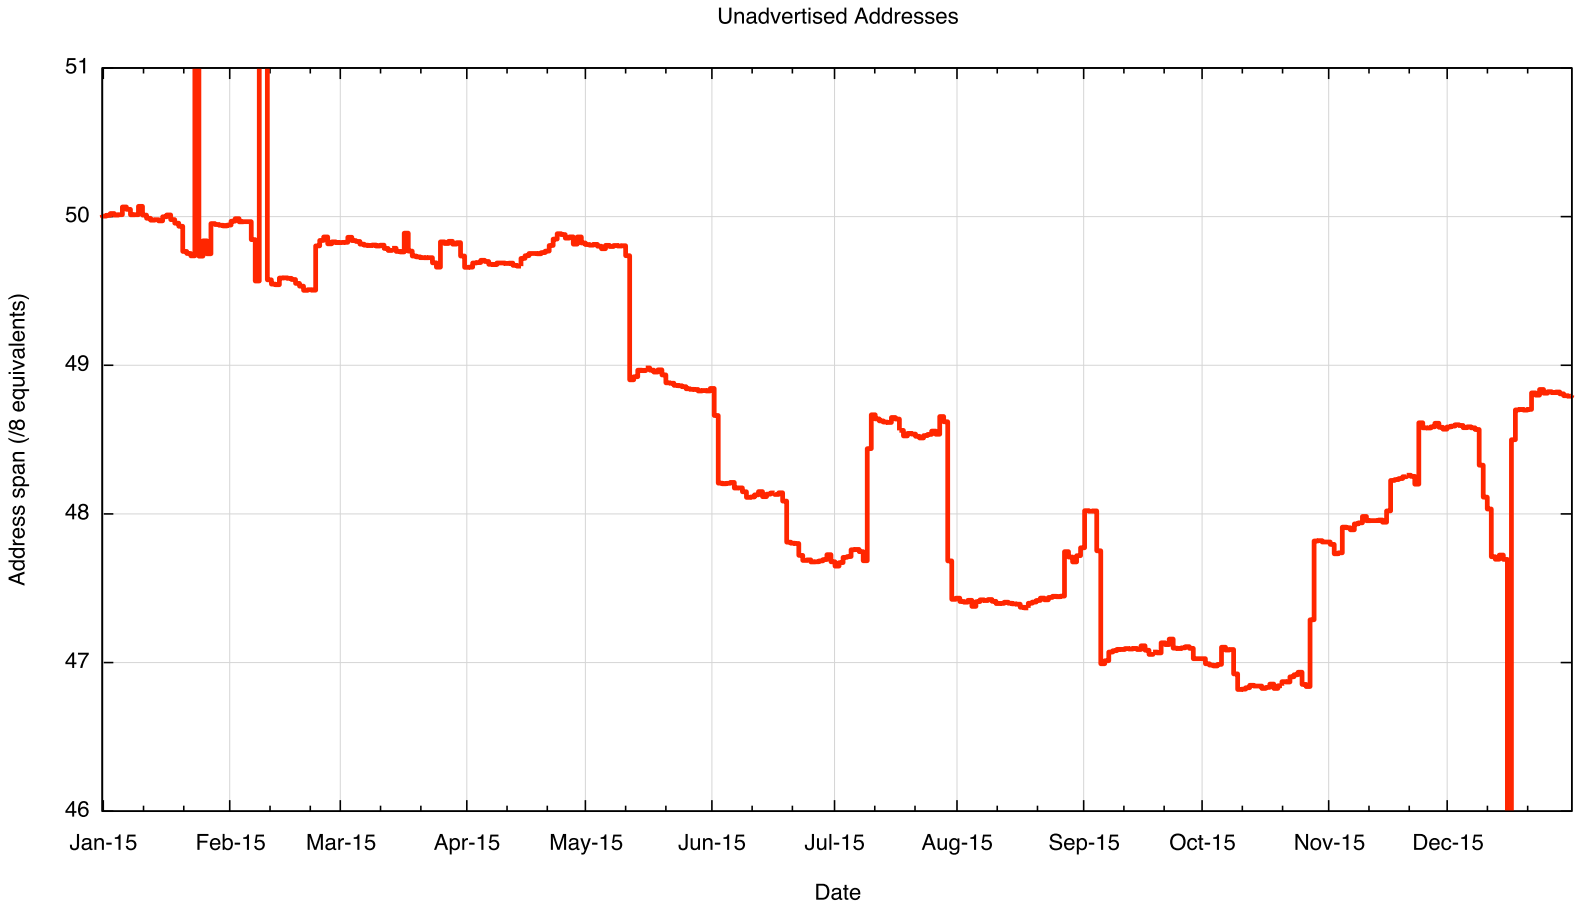

IPv4: Unadvertised Addresses

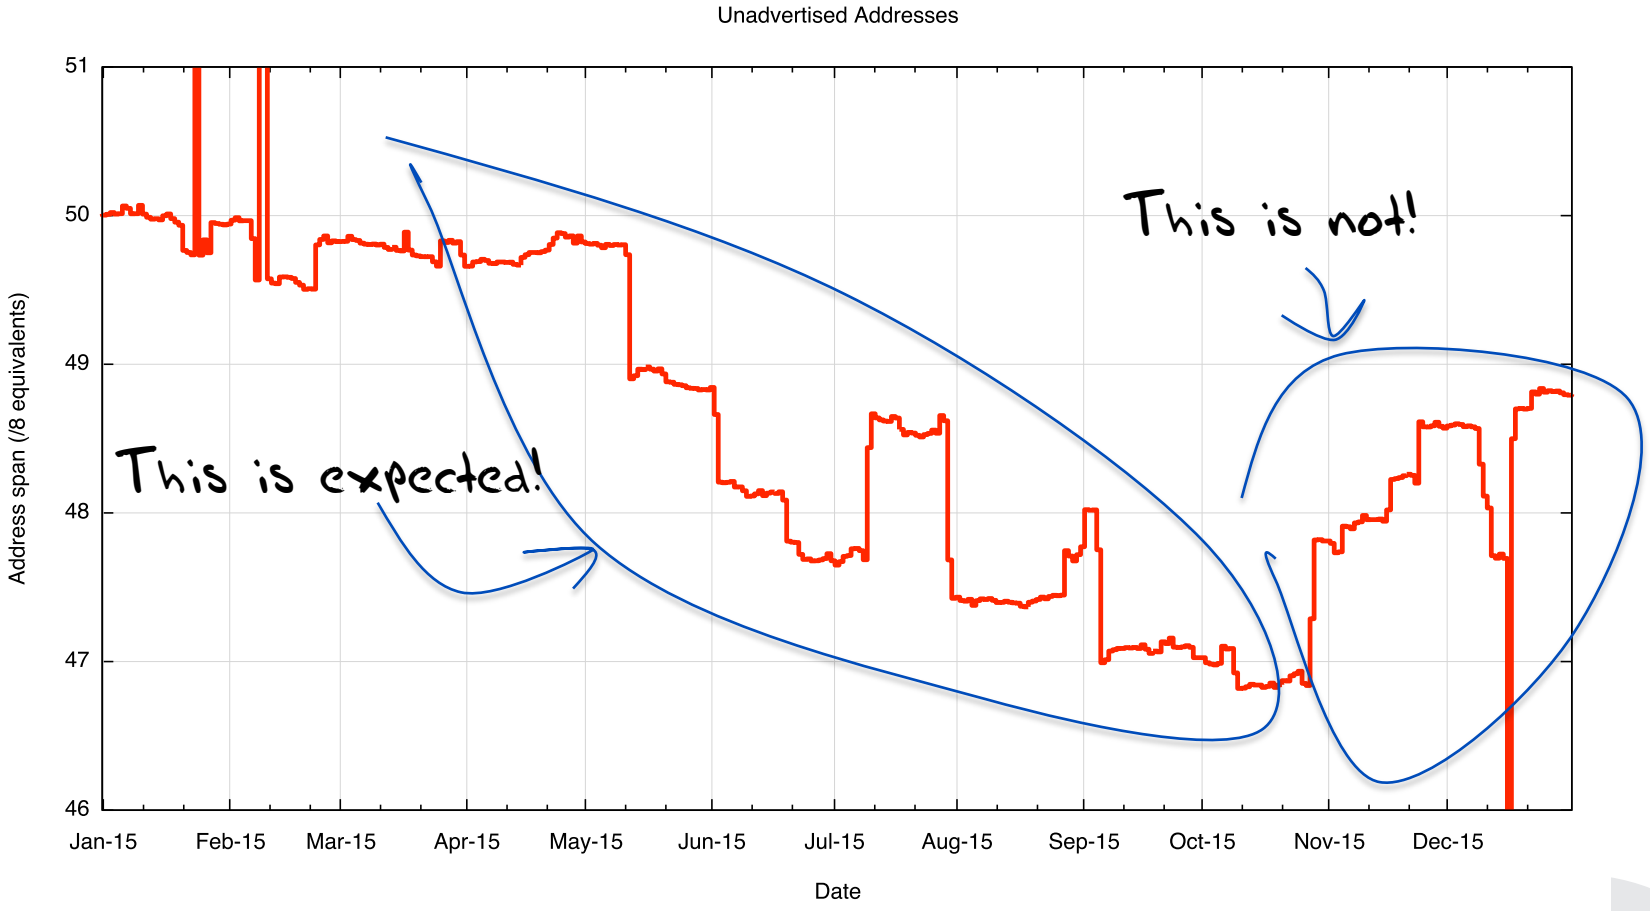

IPv4: Unadvertised Addresses

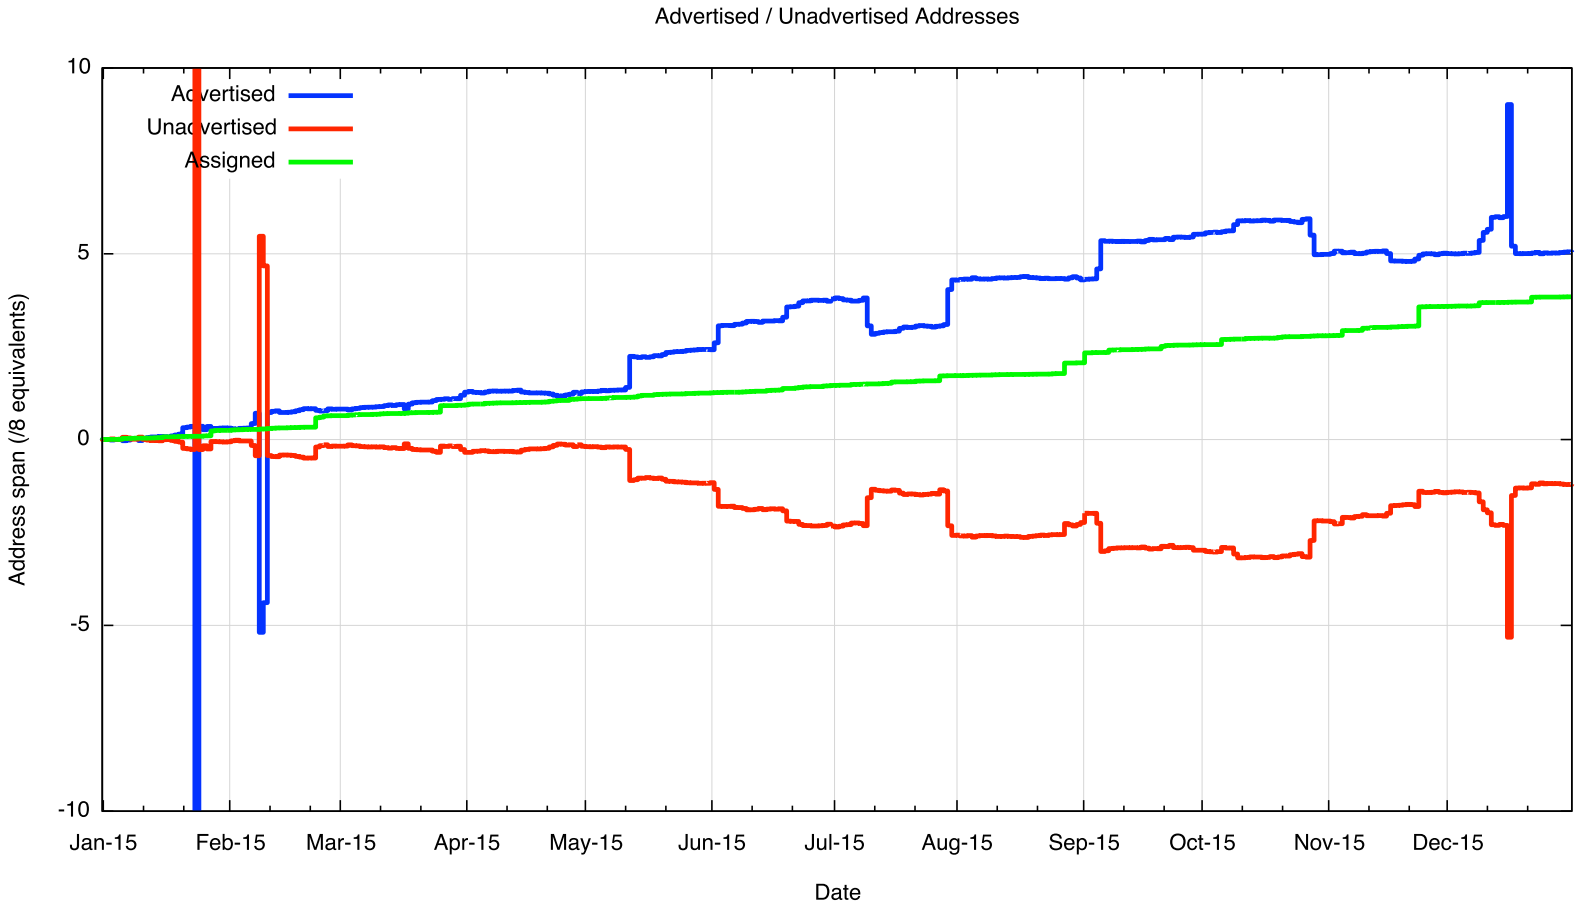

IPv4:Assigned vs Recovered

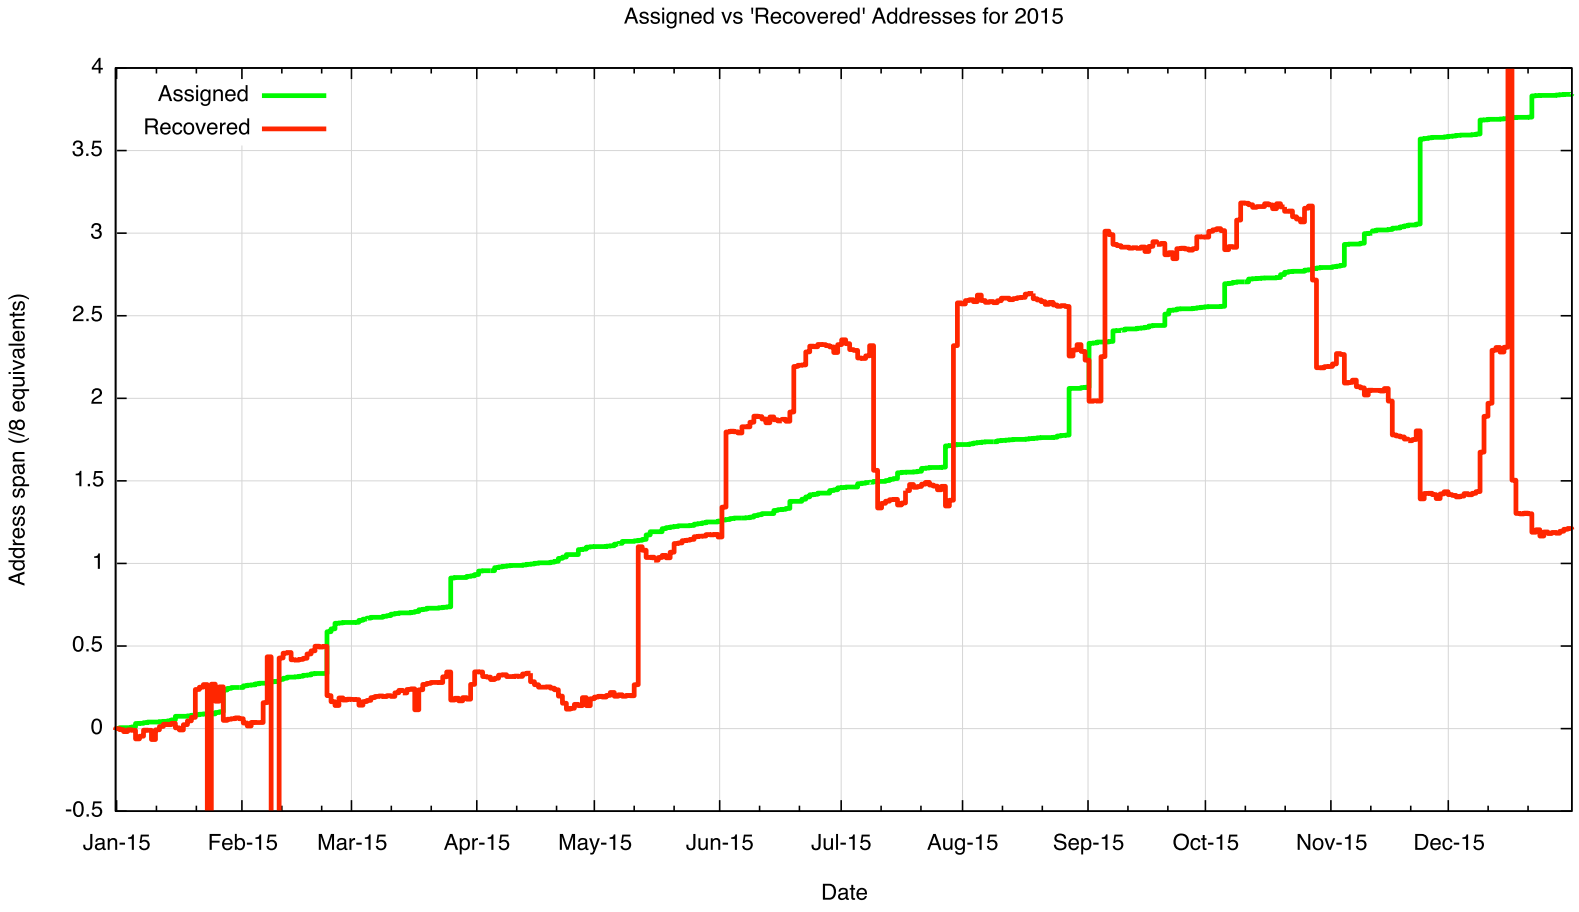

IPv4 in 2015

Approximately 4 /8s were assigned and advertised in 2015

- 2.3 /8s were assigned by ARIN
- 1 /8 assigned by AfriNIC

Up to 3 /8s were 'recovered' from the unallocated address pool and advertised during 2015

- But 2/8s of addresses were withdrawn in the last two months of the year

The Route Views view of IPv6

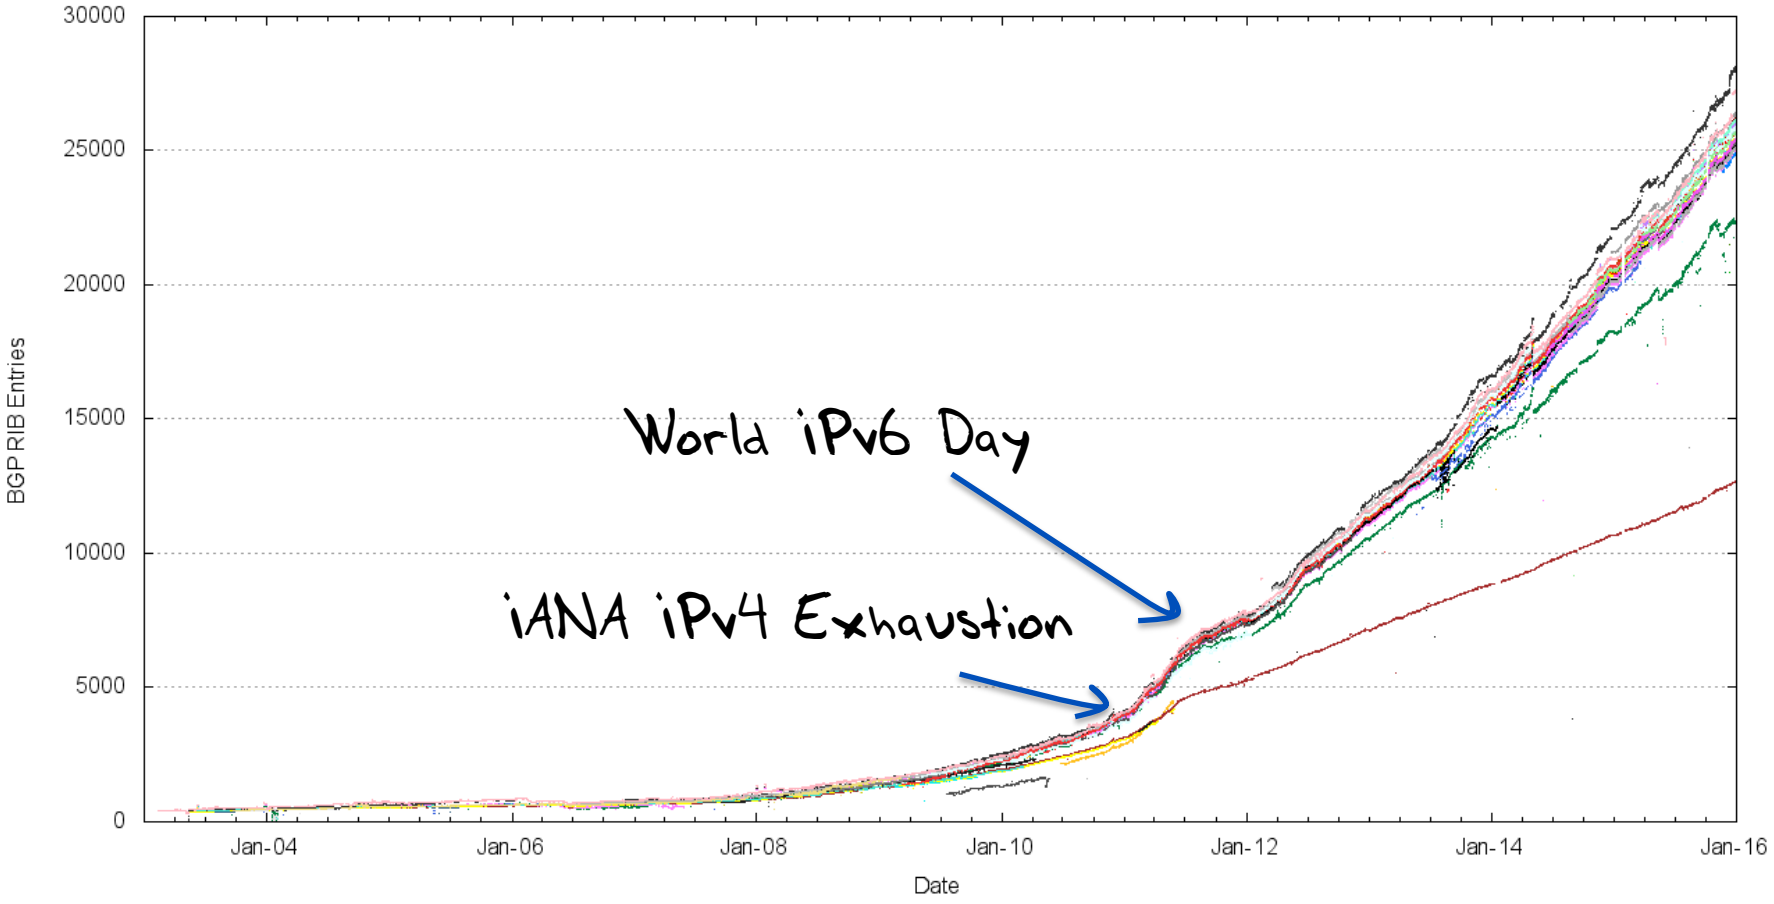

2015 for IPv6, as seen at Route Views

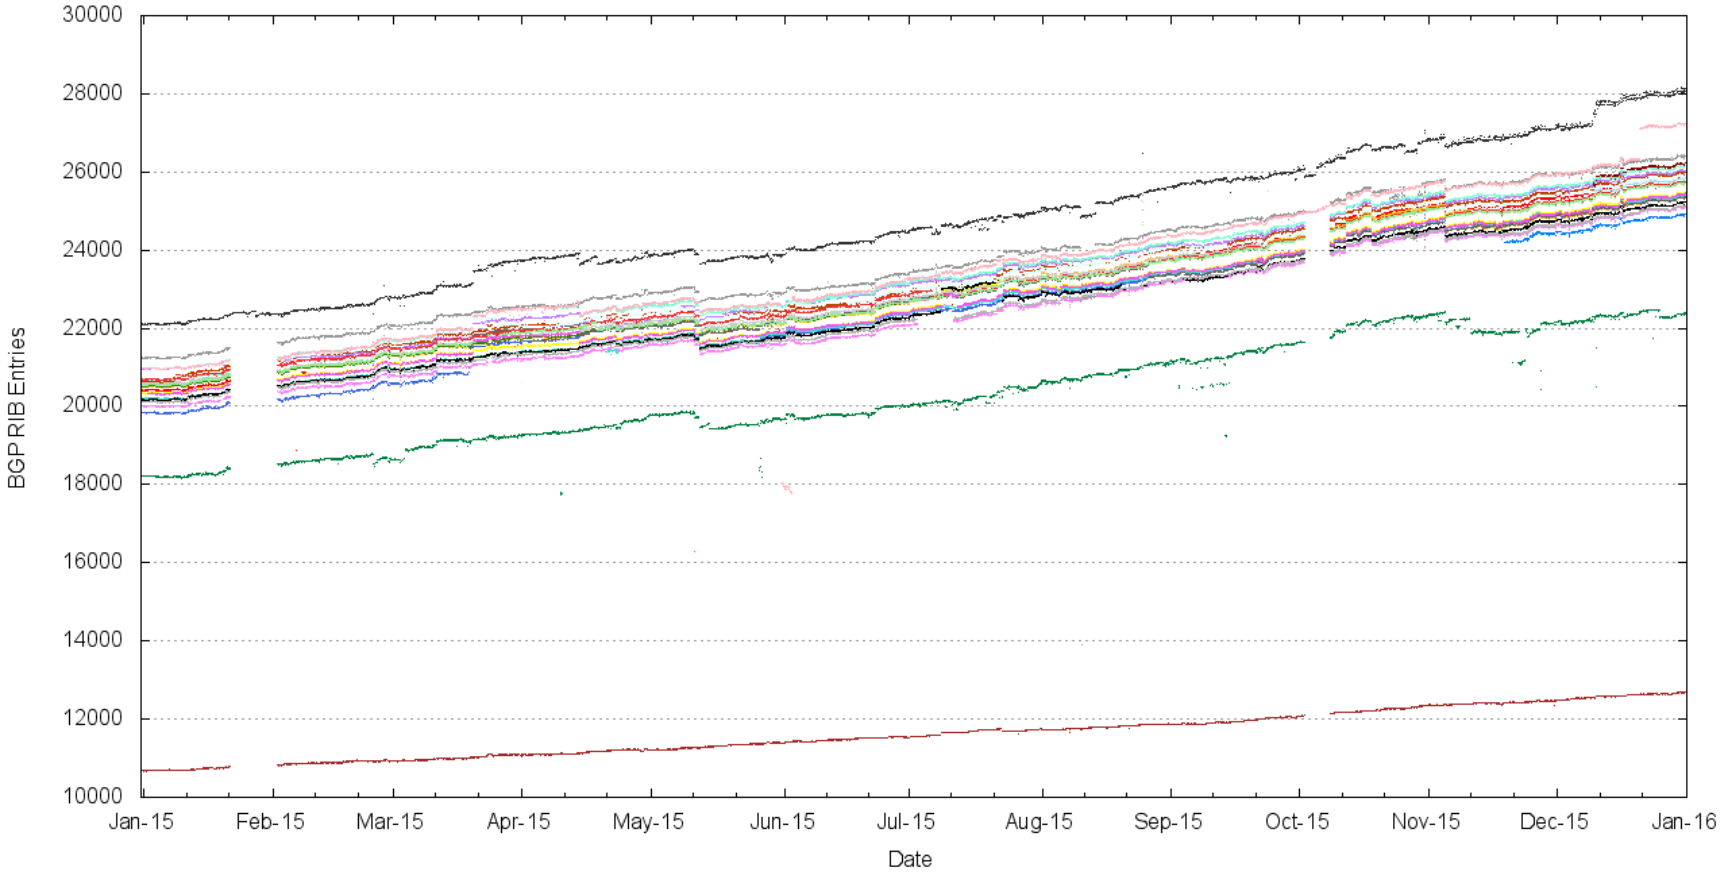

Routing Indicators for IPv6

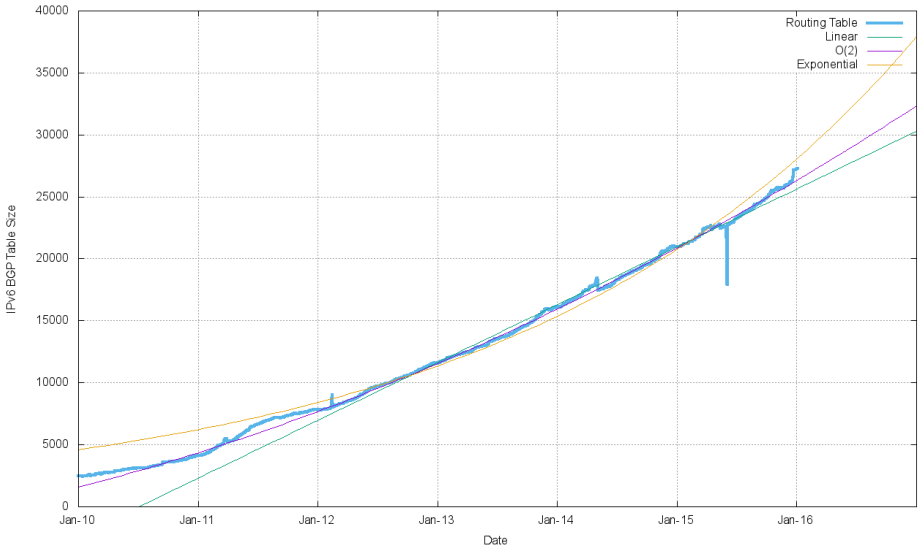

Routing prefixes – growing by some 6,000 prefixes per year

AS Numbers– growing by some 1,600 prefixes per year (which is half the V4 growth)

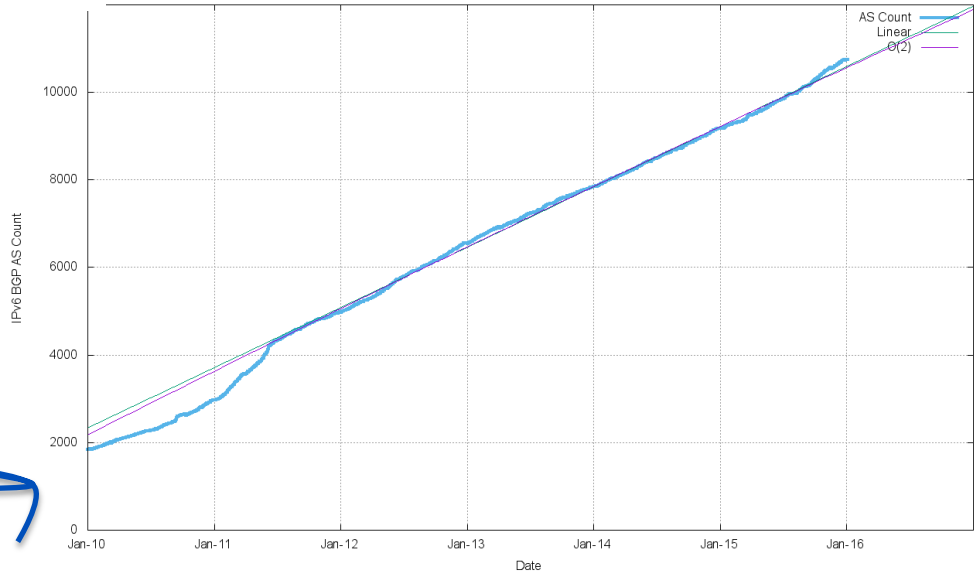

Routing Indicators for IPv6

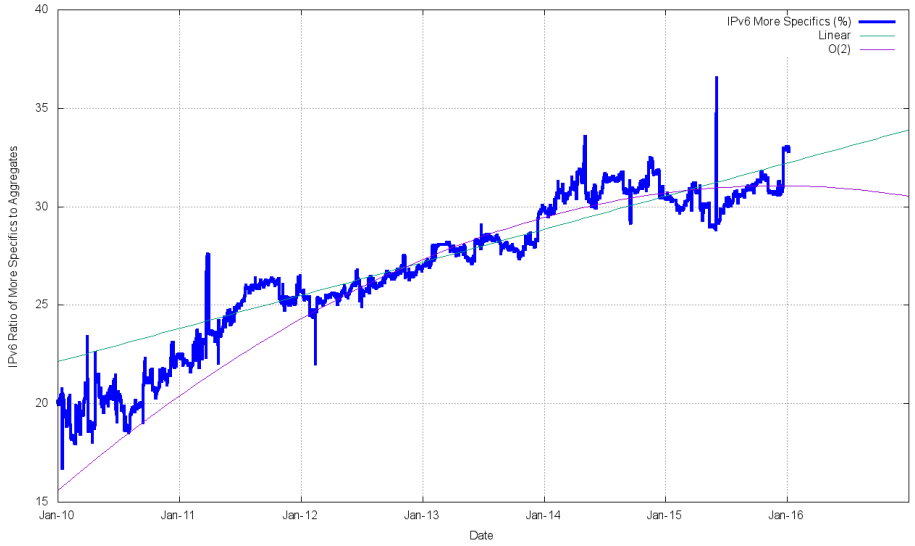

More Specifics now take up one third of the routing table

The average size of a routing advertisement is getting smaller

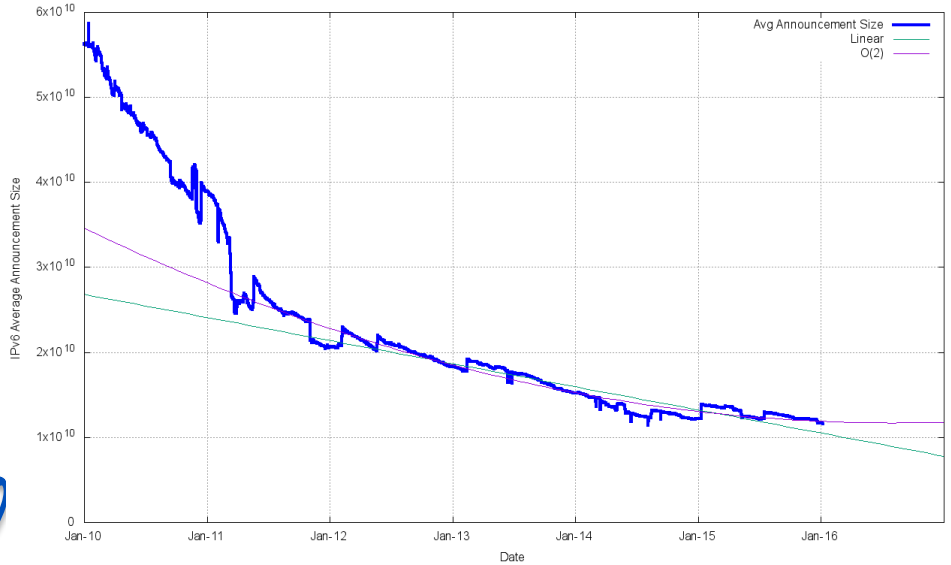

Routing Indicators for IPv6

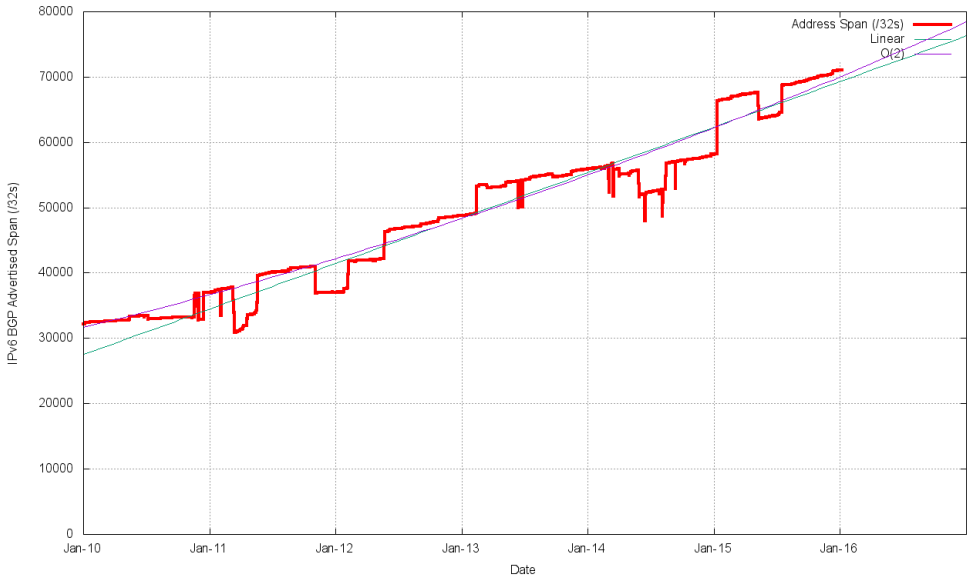

Advertised Address span is growing at a linear rate

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

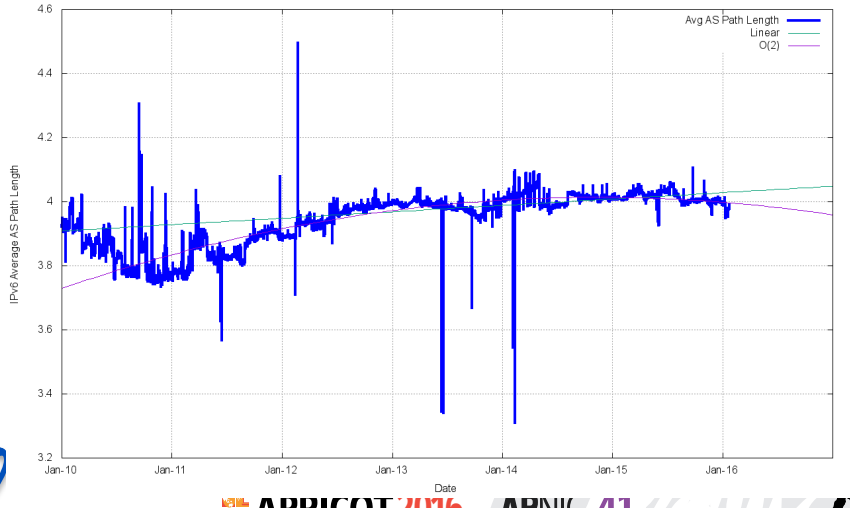

IPv6 in 2015

- 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 1/7 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 any after-market in IPv4 on the routing table
- Uncertainty over IPv6 takeup leads to a mixed response to IPv6 so far, and no clear indicator of trigger points for change

V4 - Daily Growth Rates

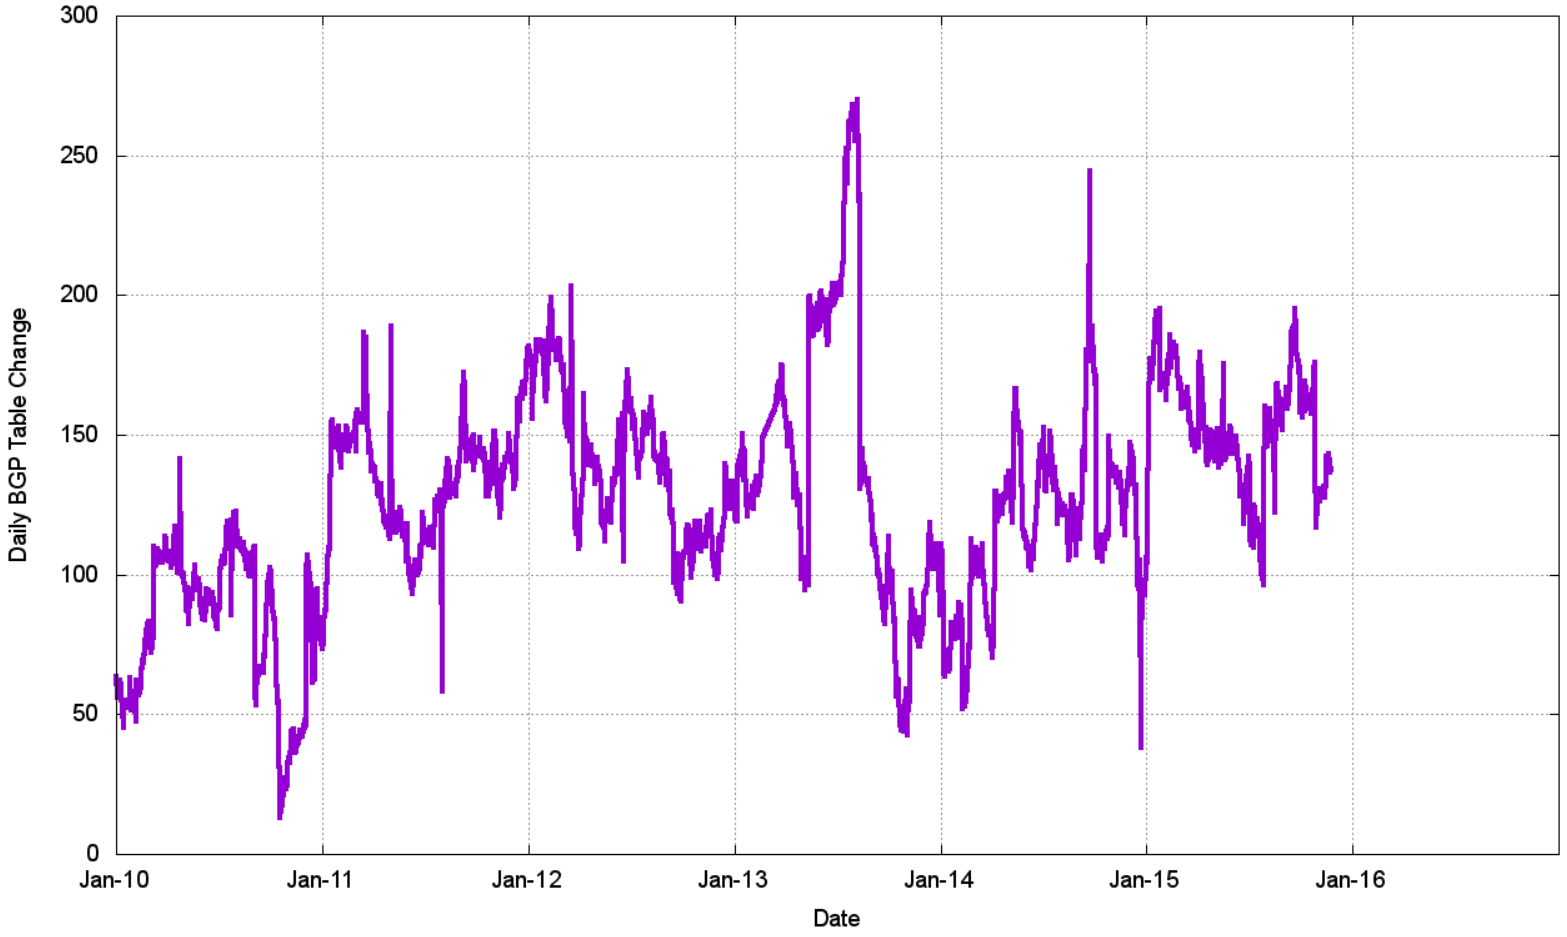

V4 - Daily Growth Rates

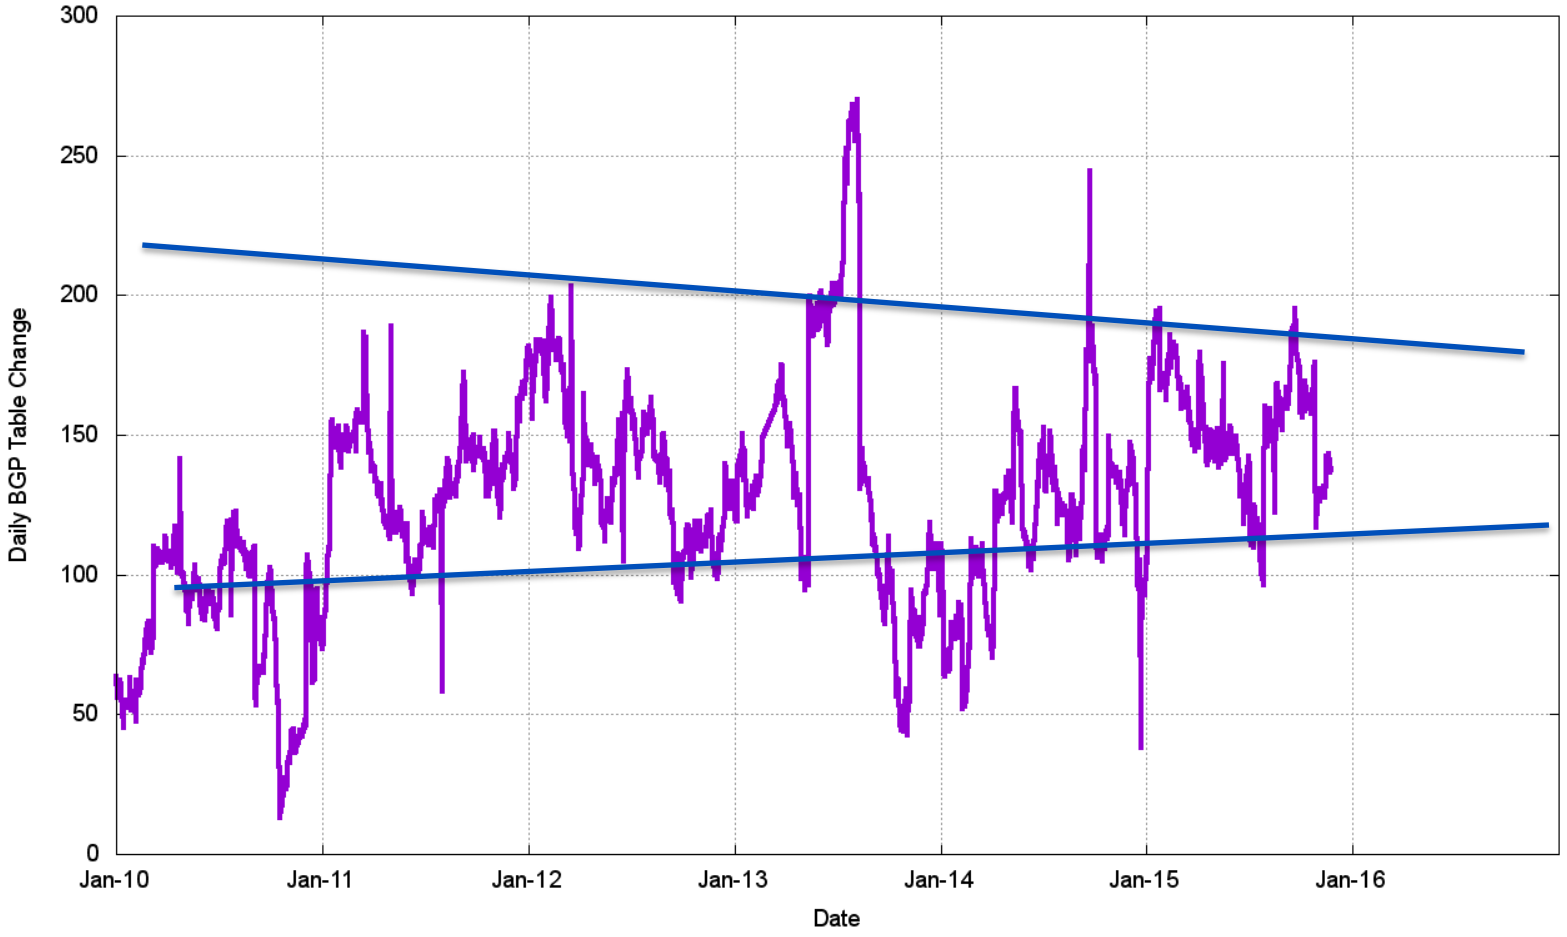

V4 - Relative Daily Growth Rates

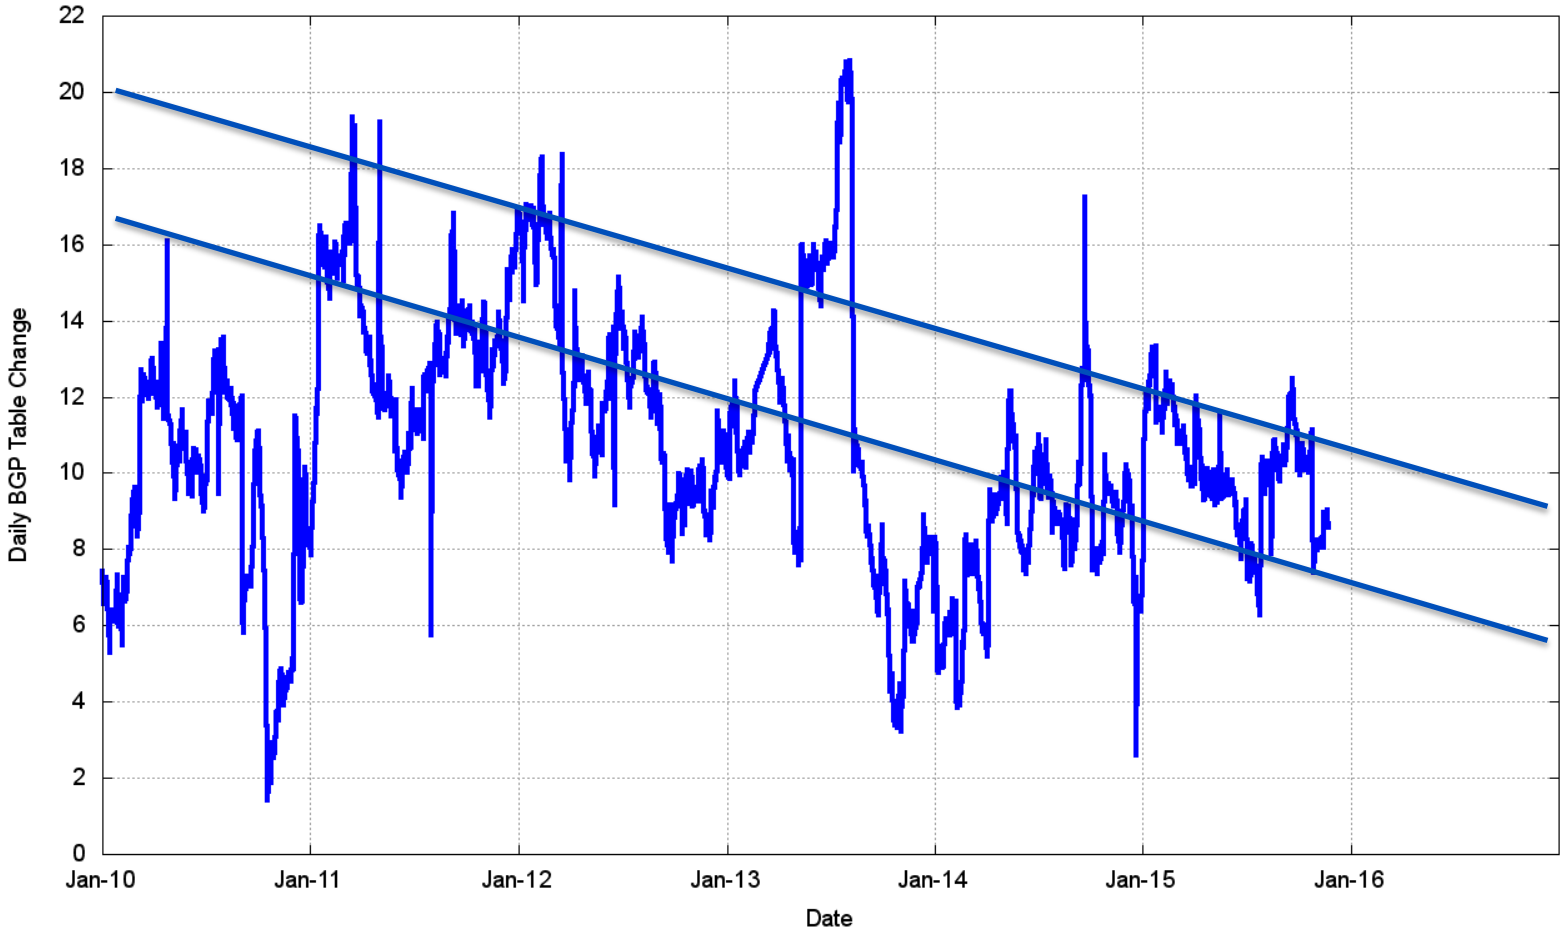

V4 - Relative Daily Growth Rates

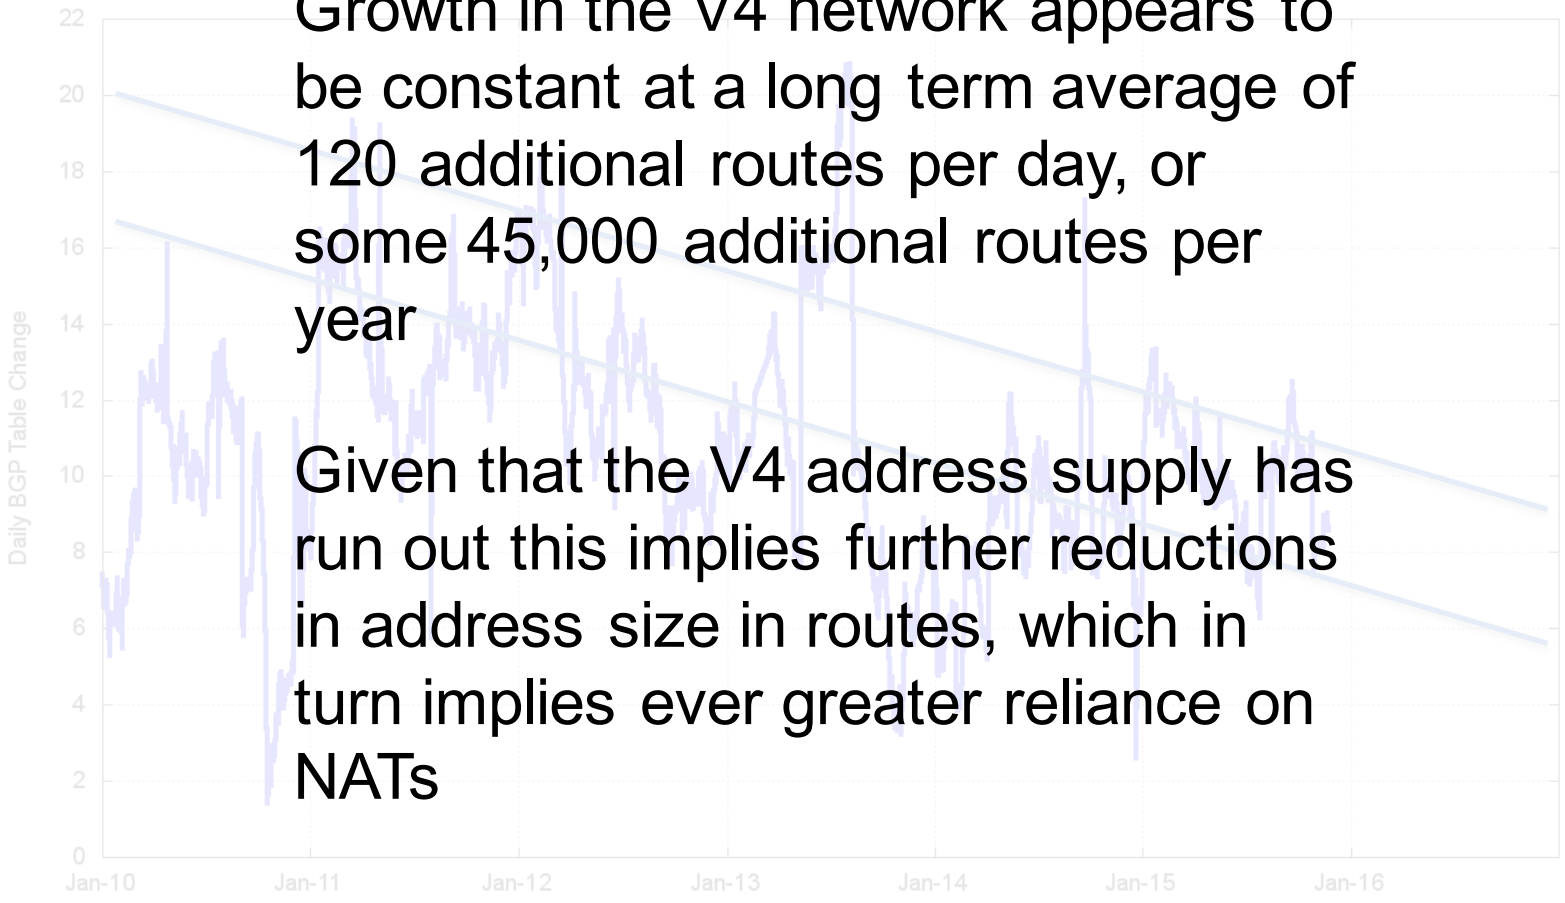

Growth in the V4 network appears to be constant at a long term average of 120 additional routes per day, or some 45,000 additional routes per year

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

Its hard to see how and why this situation will persist at its current levels over the coming 5 year horizon

IPv4 BGP Table Size predictions

		2014 PREDICTION	2013 PREDICTION
Jan 2013	441,000		
2014	488,000		
2015	*530,000*		*540,000*
2016	***586,000***	*580,000*	*590,000*
2017	*628,000*	*620,000*	*640,000*
2018	*675,000*	*670,000*	*690,000*
2019	*722,000*	*710,000*	*740,000*
2020	*768,000*	*760,000*	
2021	*815,000*		

These numbers are dubious due to uncertainties introduced by IPv4 address exhaustion pressures.

IPv6 Table Size

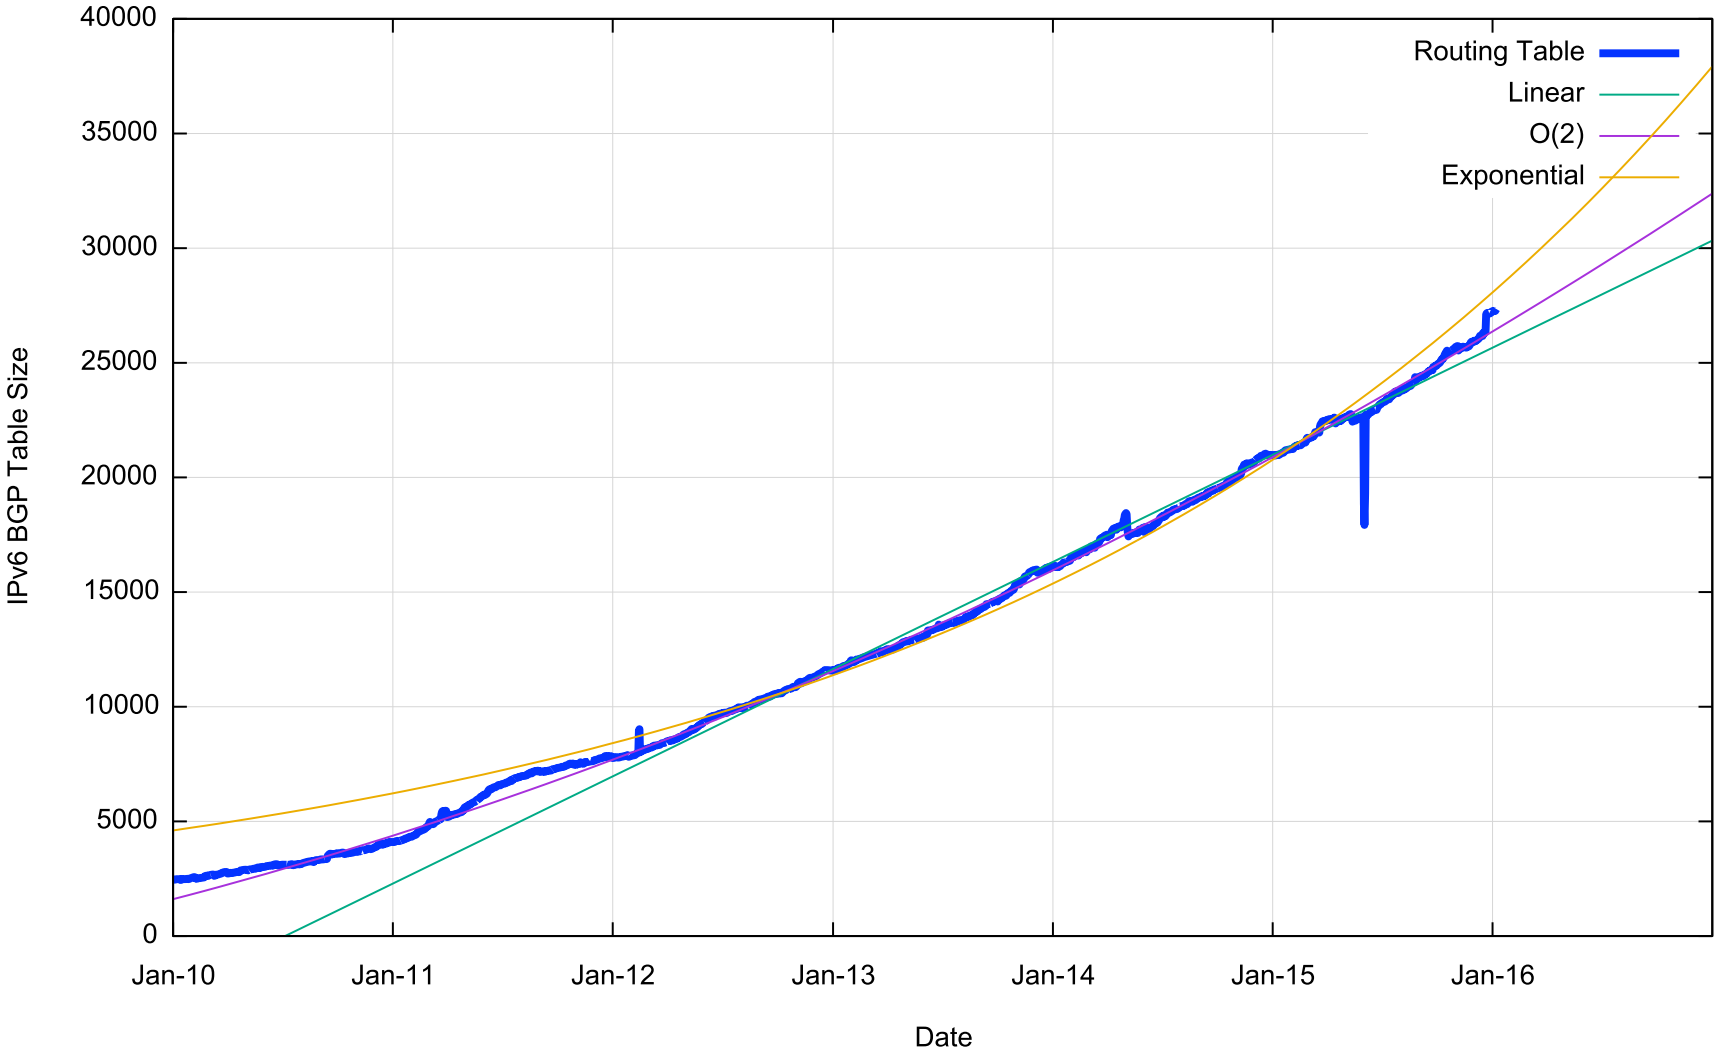

V6 - Daily Growth Rates

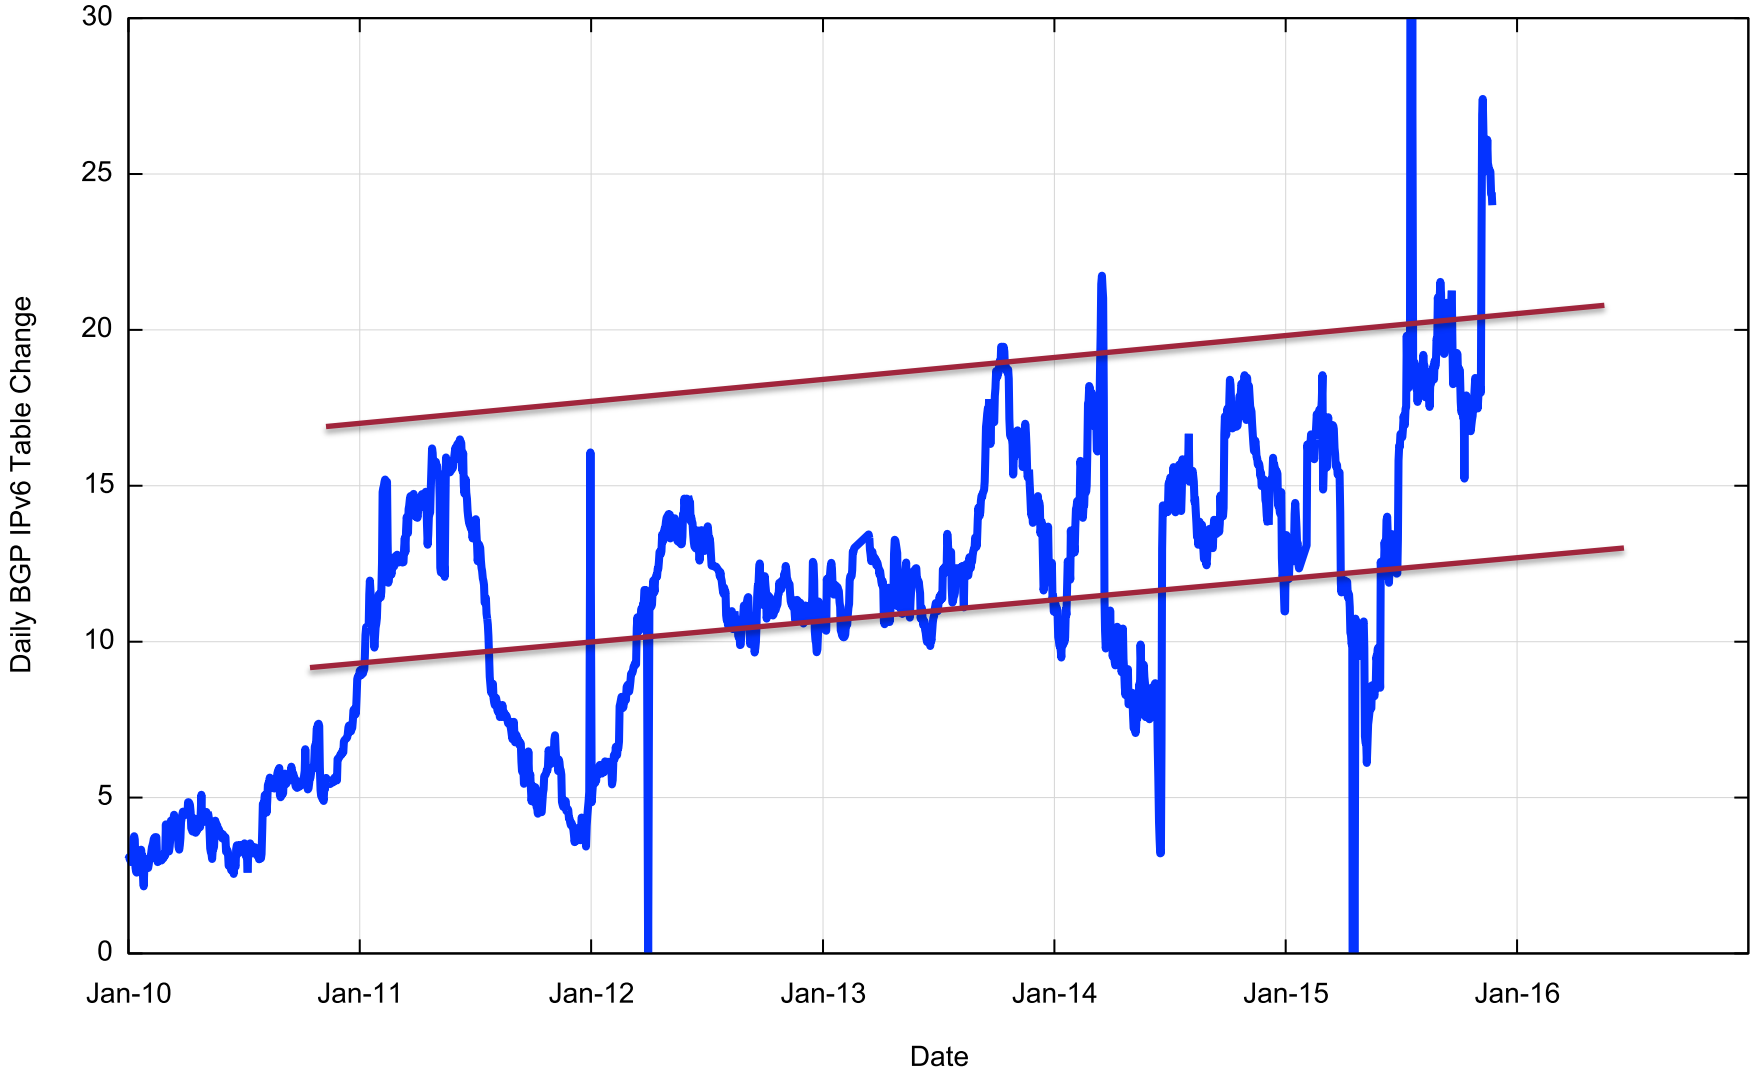

V6 - Relative Growth Rates

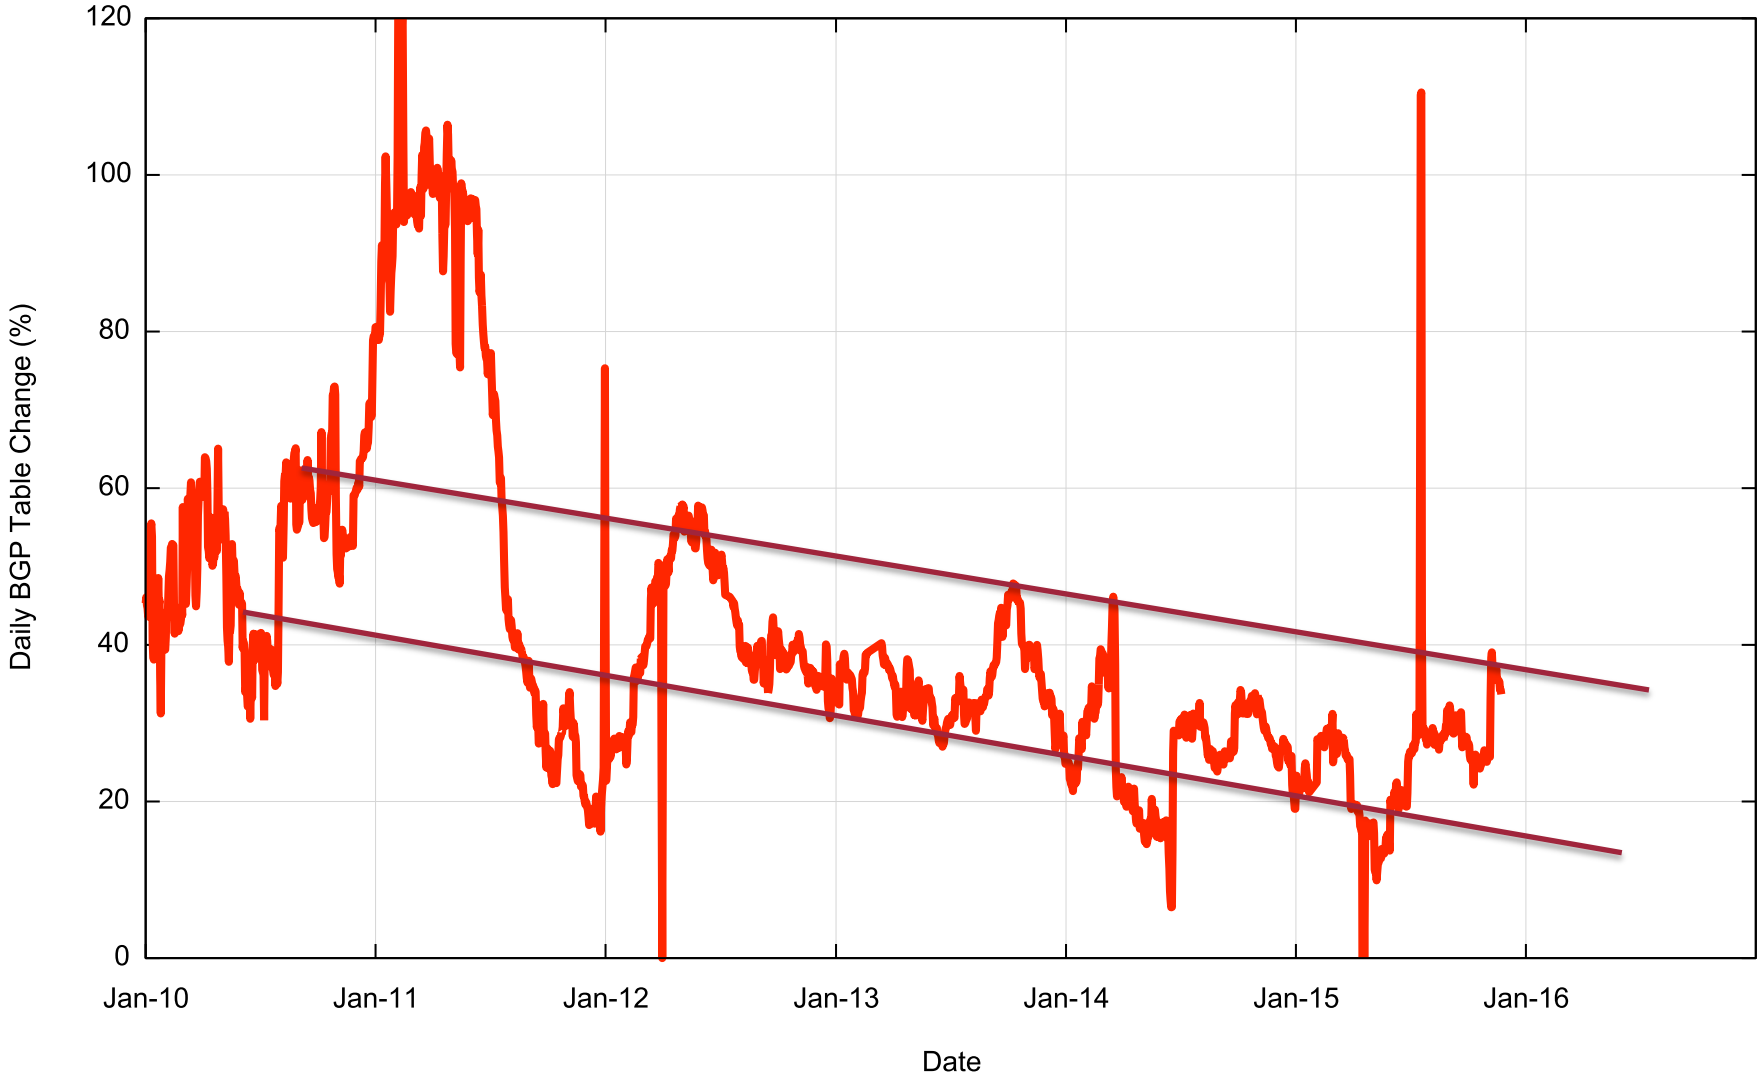

V6 - Relative Growth Rates

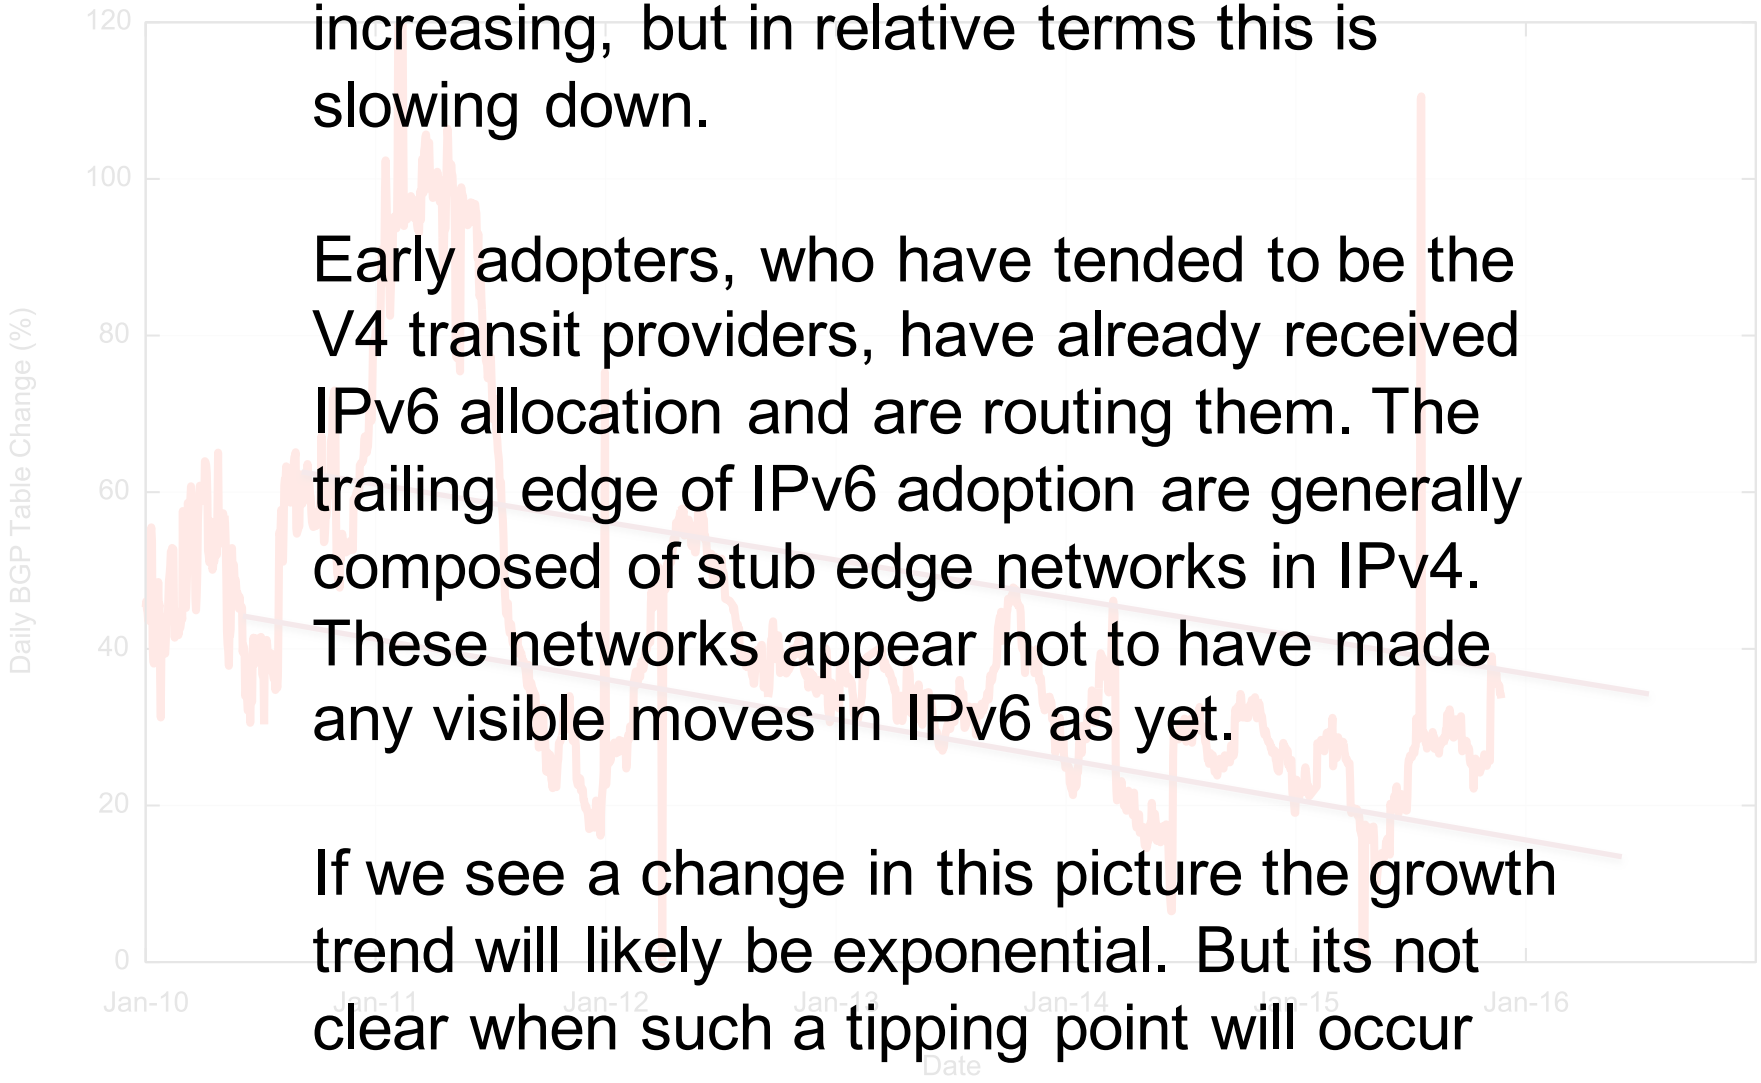

Growth in the V6 network appears to be increasing, but in relative terms this is slowing down.

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. 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 entries		
2015	21,200		
2016	*27,000*		
2017		*38,000*	*30,000*
2018		*51,000*	*35,000*
2019		*70,000*	*40,000*
2020		*94,000*	*44,000*
2021		*127,000*	*49,000*

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

Announcements and Withdrawals

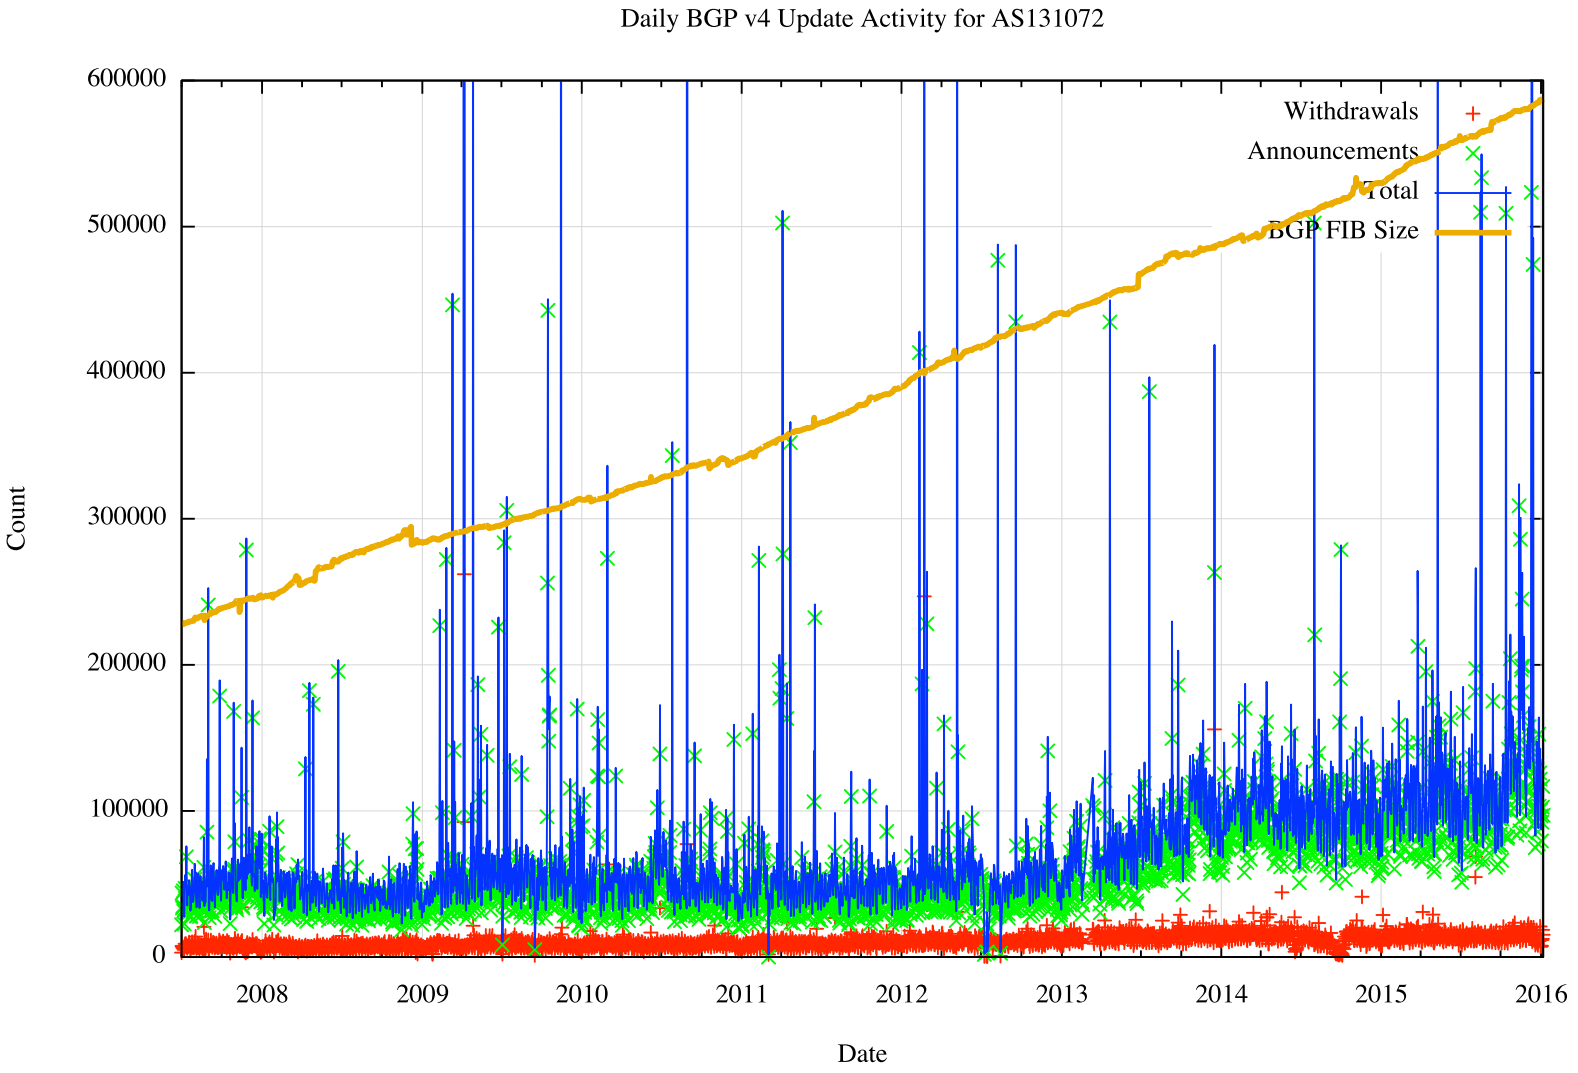

Announcements and Withdrawals

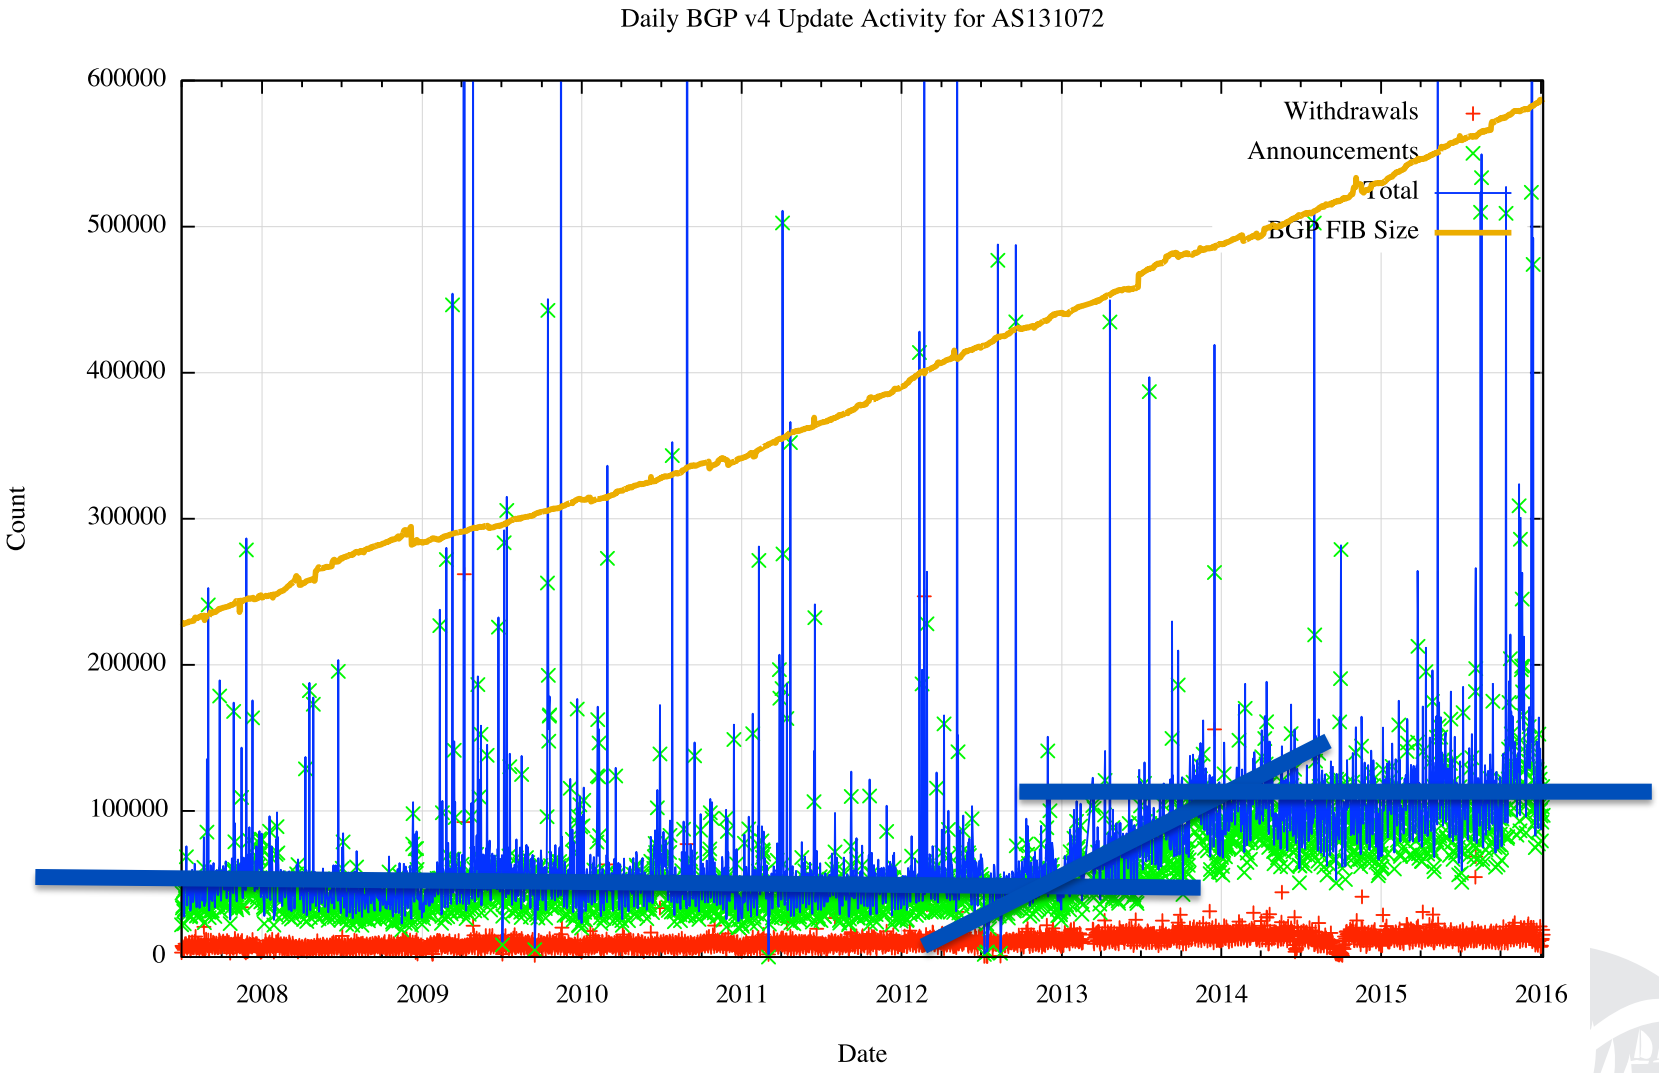

Convergence Performance

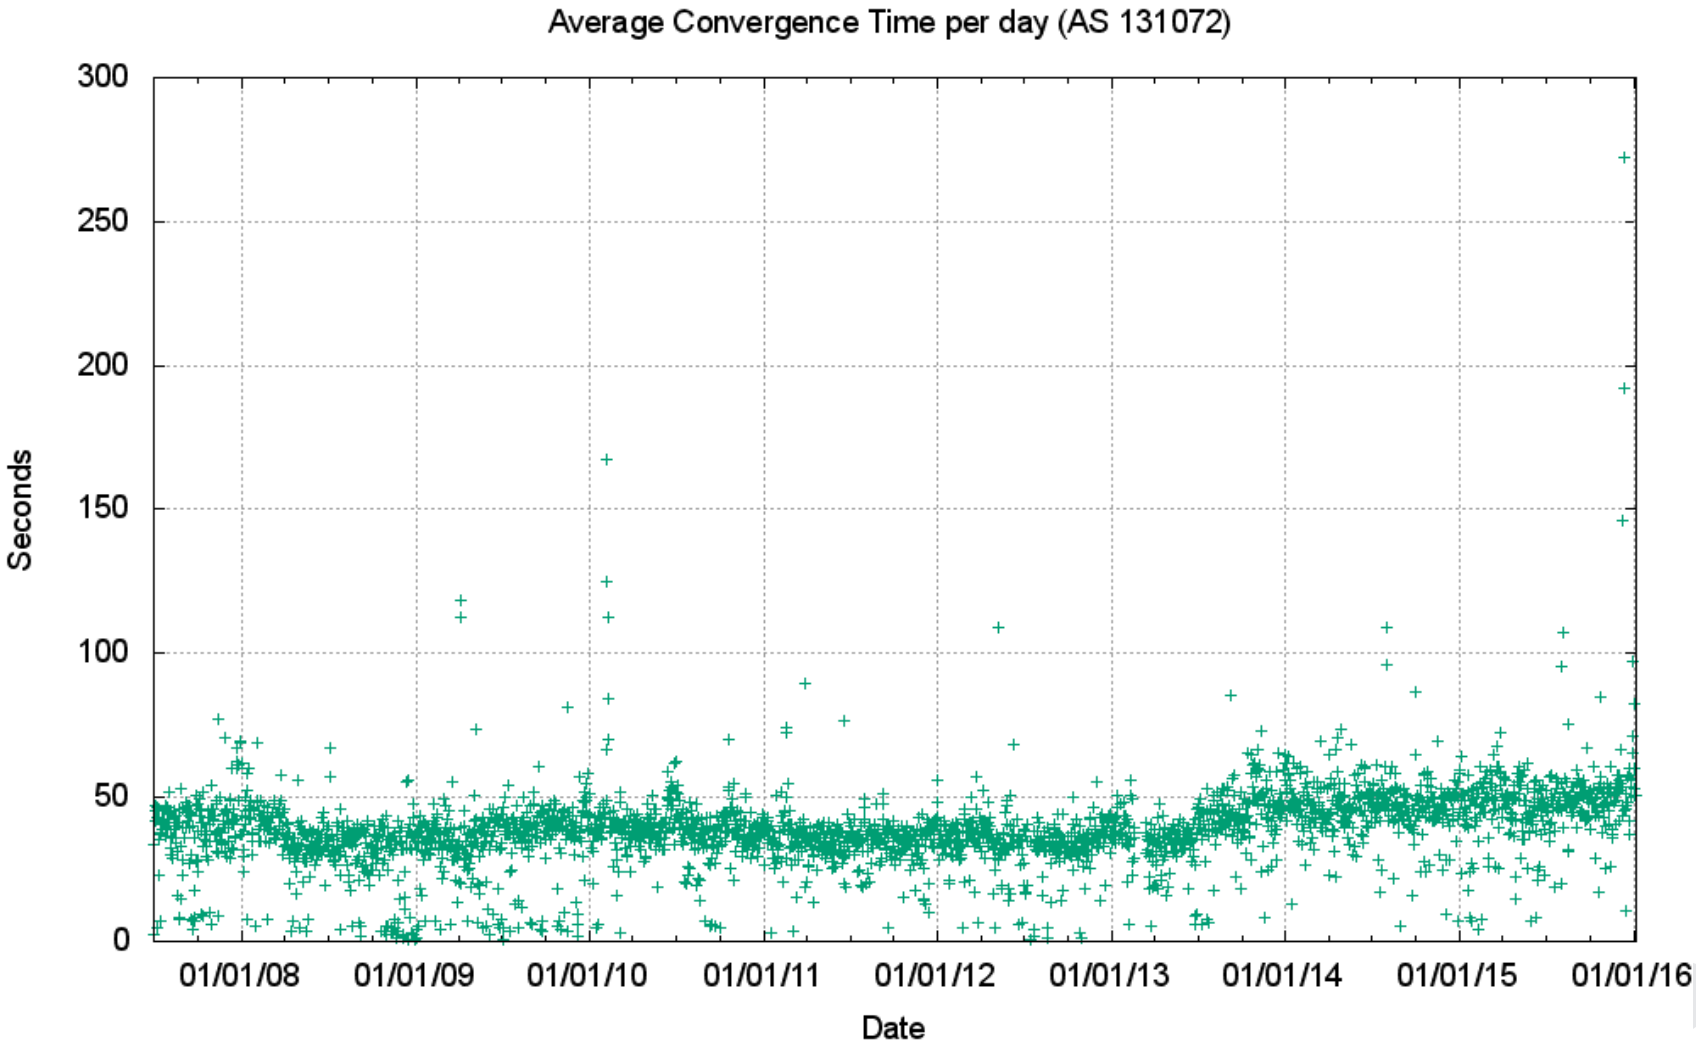

Updates in IPv4 BGP

Nothing in these figures is cause for any great level of concern …

- The number of updates per instability event has been relatively constant, which for a distance vector routing protocol is weird, and completely unanticipated. Distance Vector routing protocols should get noisier as the population of protocol speakers increases, and the increase should be multiplicative.
- But this is not happening in the Internet
- Which is good, but why is this not happening?

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

V6 Announcements and Withdrawals

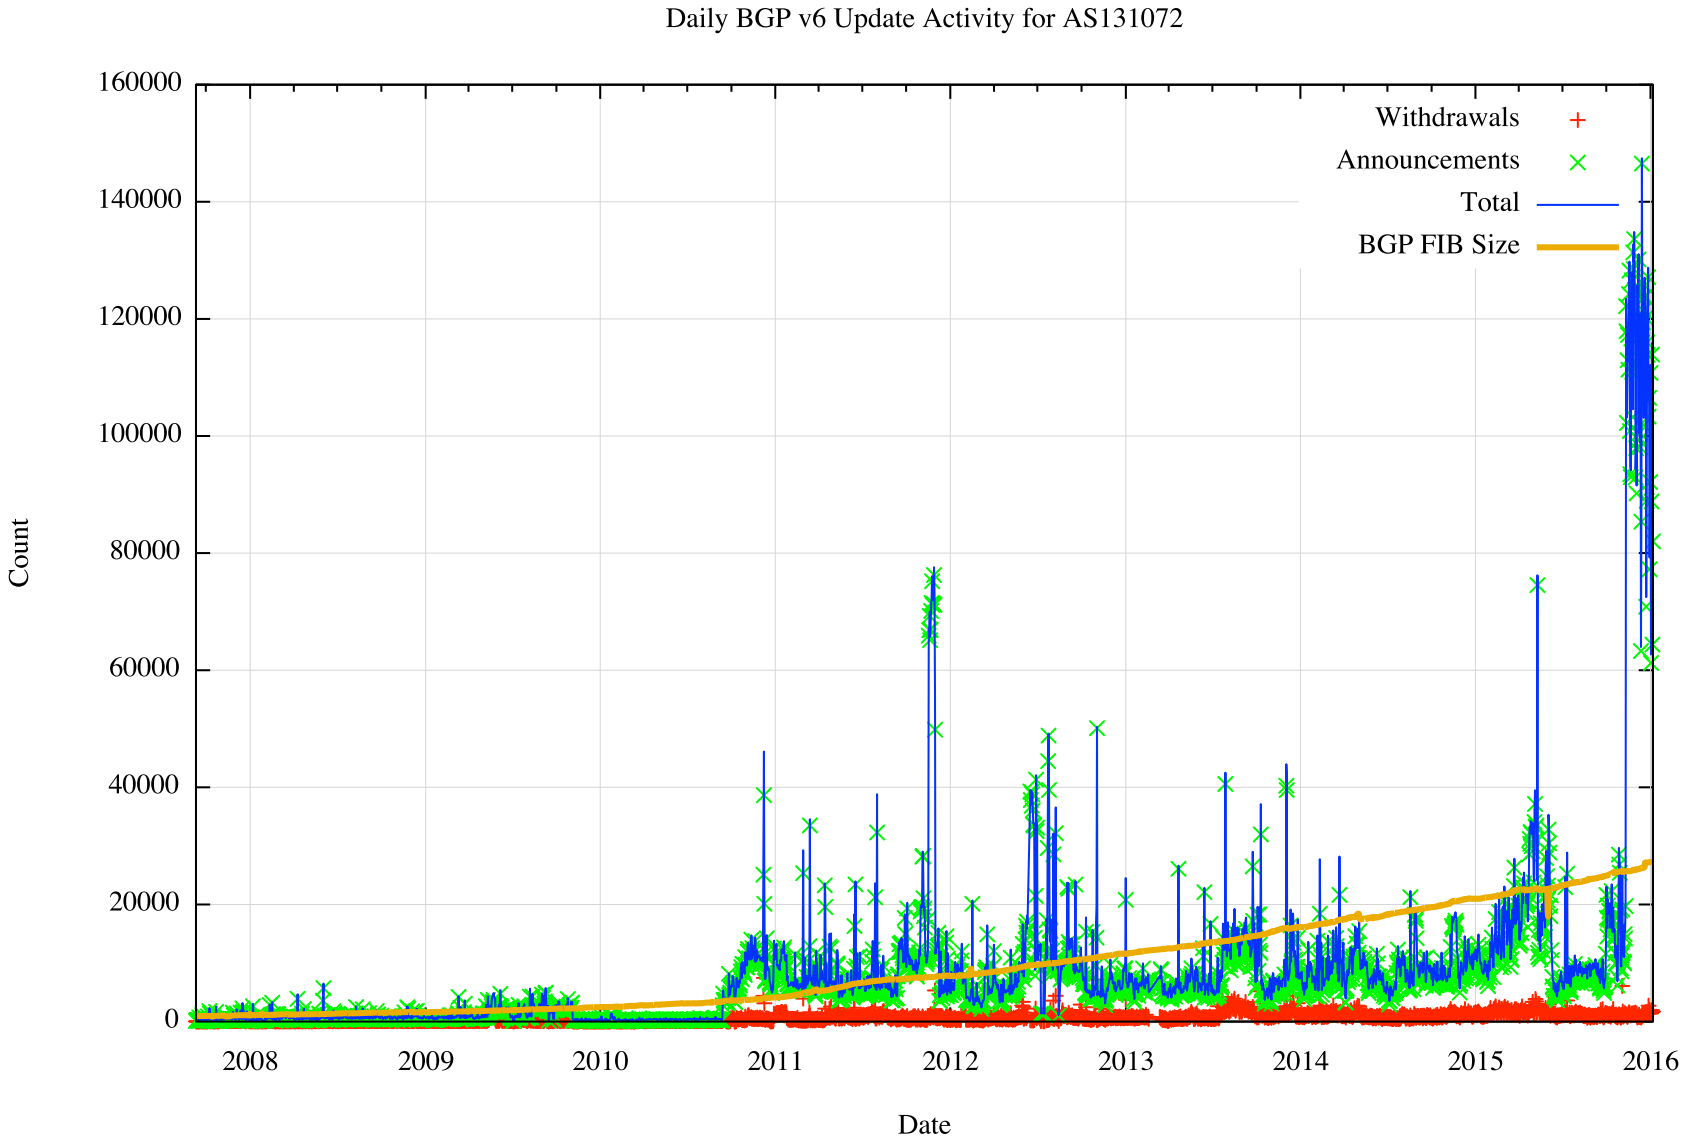

V6 Convergence Performance

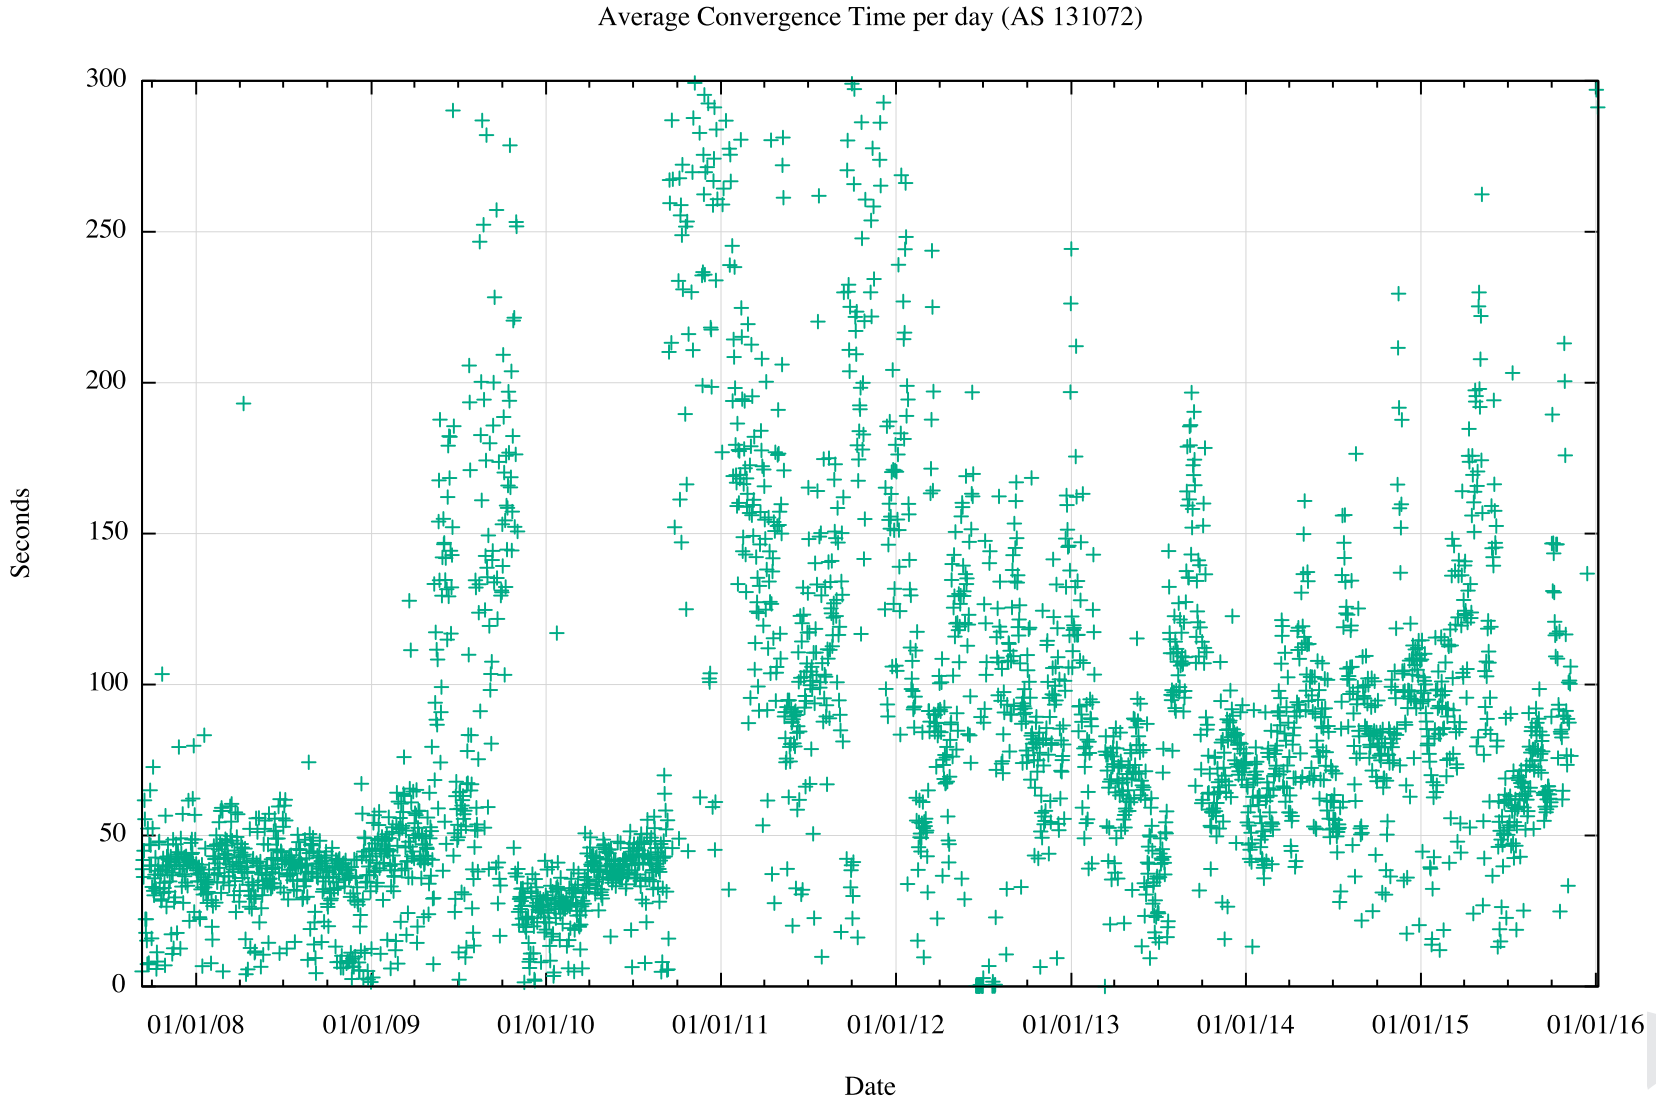

V6 Updated prefixes per day

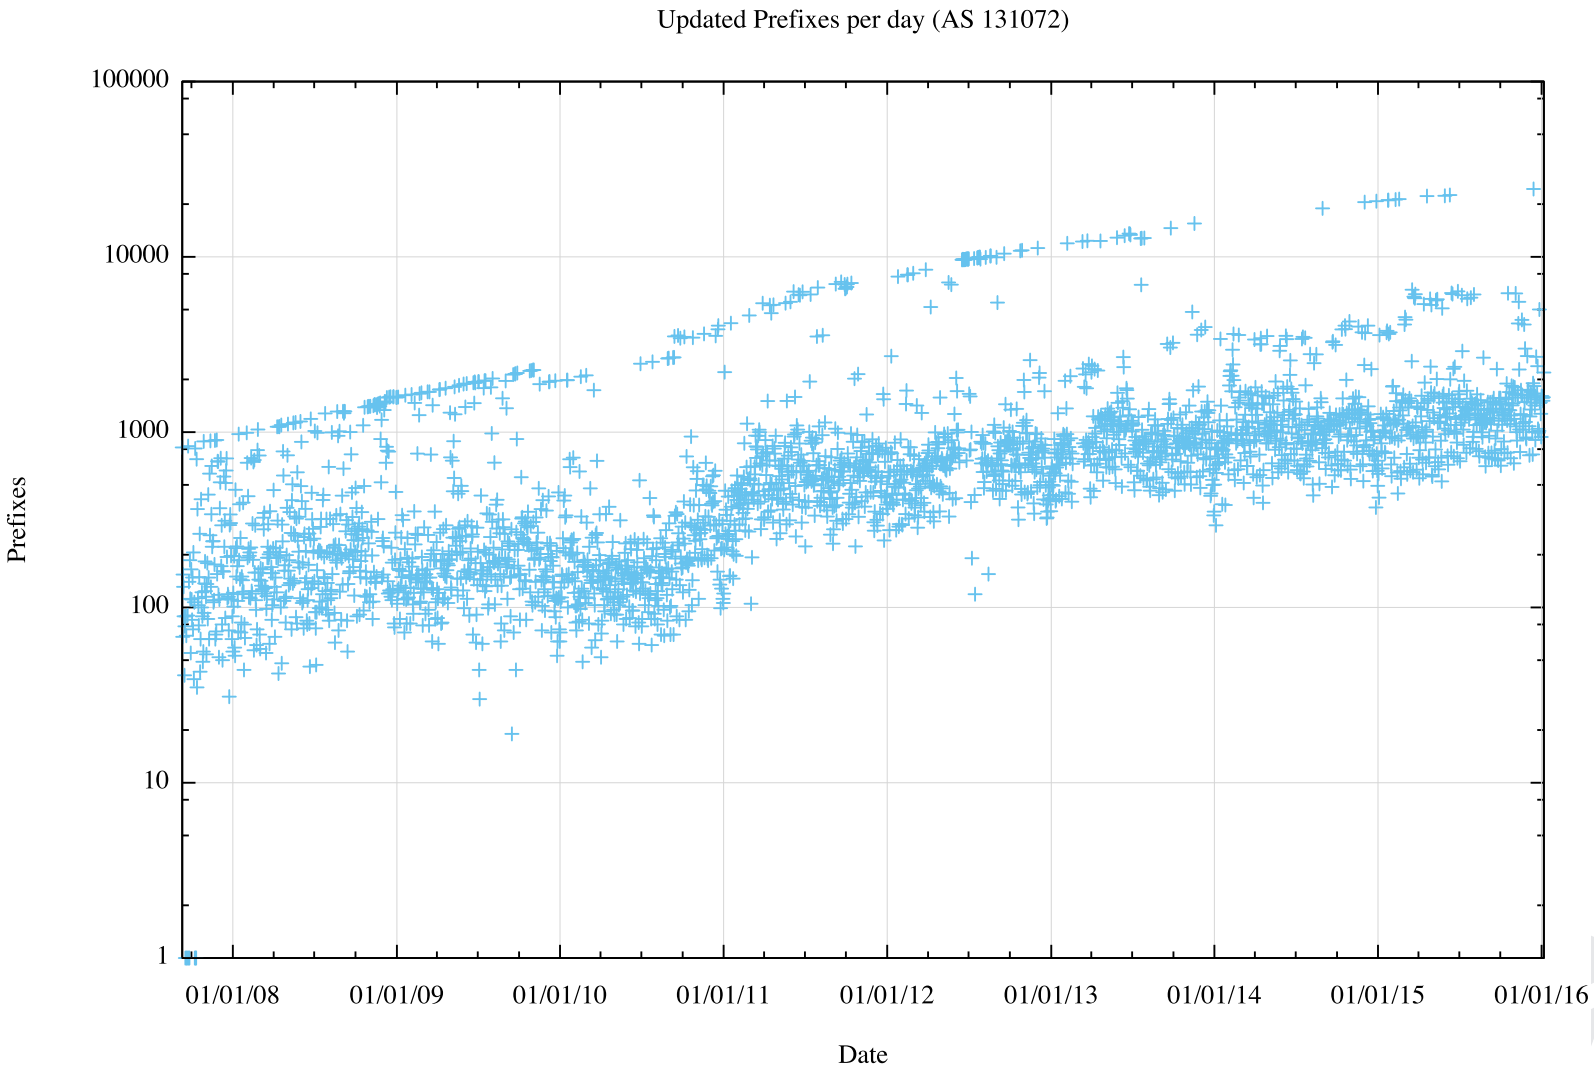

V6 Updates per event

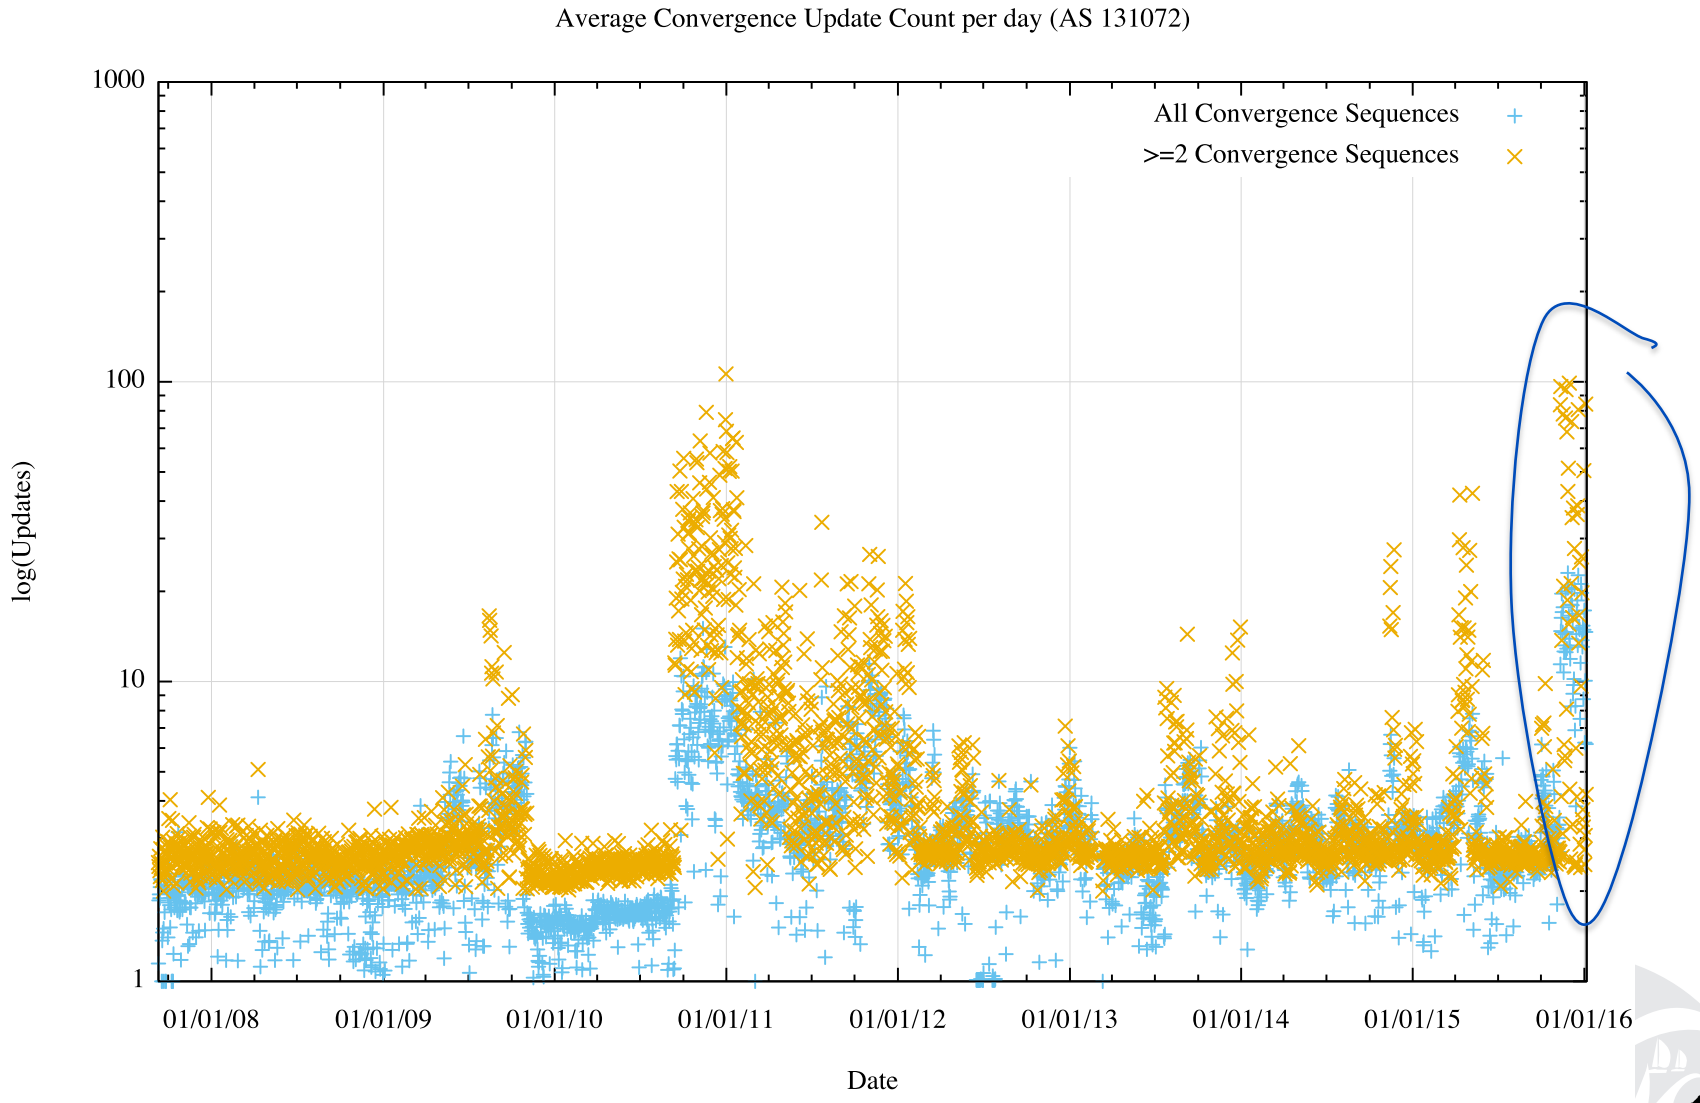

Updates in IPv6 BGP

IPv6 routing behaviour is diverging from IPv4 behaviour

The instability is greater

Its not the number of unstable prefixes, but the number of updates and elapsed time for the network to re-converge for each instability event

It this were to happen in the V4 network at the same relative scale it would be a major stability problem!

So what is going on and why has this happened?

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

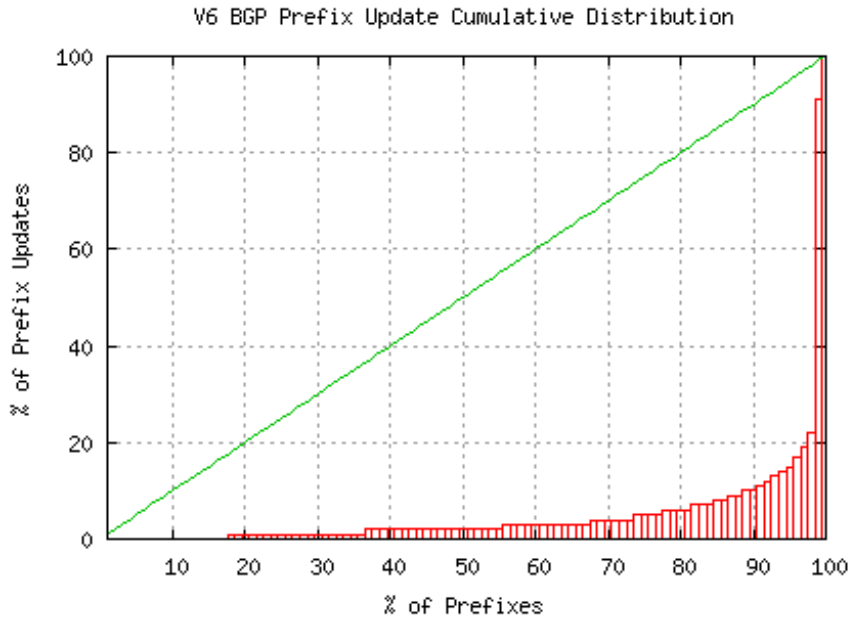

Updates in IPv6

RANK	PREFIX	UPDs	%	Origin AS -- AS NAME
1	2c0f:fe90::/32	48777	8.99%	36943 -- webafrica,ZA
2	2407:8c00:ffe1::/48	39060	7.20%	131317 -- TTSLMEIS-IN TTSL-ISP DIVISION,IN
3	2804:14d:5a82::/48	34537	6.37%	28573 -- CLARO S.A.,BR
4	2605:cf00::/32	20020	3.69%	46525 -- RURALWAVE-LTD - Rural Wave,CA
5	2403:8600:ea89::/48	19448	3.58%	131317 -- TTSLMEIS-IN TTSL-ISP DIVISION,IN
6	2001:470:e9::/48	16613	3.06%	25104 -- WEBCODE WebCode Ltd.,BG
7	2804:39c:7000::/36	12787	2.36%	28329 -- G8 NETWORKS LTDA,BR
9	2001:7fb:ff02::/48	11097	2.05%	12654 -- RIPE-NCC-RIS-AS Reseaux IP Europeens Network Coordination Centre (RIPE NCC),EU
10	2607:f870:1::/48	9530	1.76%	11992 -- CENTENNIAL-PR - Centennial de Puerto Rico,PR
11	2001:67c:370::/48	9383	1.73%	56554 -- IETF-MEETING Internet Society,CH
12	2001:67c:1230::/46	9274	1.71%	56554 -- IETF-MEETING Internet Society,CH
14	2001:df8::/32	8911	1.64%	56554 -- IETF-MEETING Internet Society,CH
15	2a02:28c8::/32	8208	1.51%	42353 -- SIMWOOD Simwood eSMS Limited,GB
16	2620:27:f::/48	6064	1.12%	10846 -- DEERE - Deere & Company,US
17	2804:14d:8085::/48	5703	1.05%	28573 -- CLARO S.A.,BR
19	2a03:4600::/32	5135	0.95%	44334 -- RTLNET-ASN RTLNET,FR
20	2804:14d:8080::/48	4992	0.92%	28573 -- CLARO S.A.,BR

The busiest 48 prefixes accounted for 2/3 of all prefix updates

That's it!

Questions?