BGP in 2023

Geoff Huston

A super-brief summary of BGP!

- BGP is the distributed routing protocol that manages the reachability of address prefixes across the Internet
  - Its objective is to distribute to all BGP speakers a current list of reachable address prefixes and the local next-hop interface to use to forward packets to each address prefix
- This is an instance of a Bellman-Ford Distance Vector routing protocol
  - BGP uses a “Path Vector” to prevent loop formation and aid in the comparison of routes

For example:…

```
bgpd# show ip bgp
BGP table version is 0, local router ID is 203.133.248.2
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
              i internal, r RIB-failure, S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*  1.0.0.0/24       202.12.28.1                              0 4777 13335 i
*                   203.119.104.2                            0 4608 13335 i
*>                  203.119.104.1                            0 4608 13335 i
*  1.0.4.0/22       202.12.28.1                              0 4777 2500 7660 4635 7545 2764 38803 i
*                   203.119.104.2                            0 4608 1221 2764 38803 i
*>                  203.119.104.1                            0 4608 1221 2764 38803 i
*  1.0.5.0/24       202.12.28.1                              0 4777 2500 7660 4635 7545 2764 38803 i
*>                  203.119.104.2                            0 4608 1221 2764 38803 i
*                   203.119.104.1                            0 4608 1221 2764 38803 i
*  1.0.16.0/24      202.12.28.1                              0 4777 2516 3356 2519 i
*                   203.119.104.2                            0 4608 1221 3356 2519 i
*>                  203.119.104.1                            0 4608 1221 3356 2519 i
```

Why look at BGP Tables?

- BGP is one of the few tools that lets us see all of the Internet at once.
- The functional objective of BGP is to flood reachability information for all reachable address prefixes to all BGP speakers
- Which means that looking at BGP tables can be related to looking at the entire Internet!

The Highlights

- IPv4 Summary
- IPv6 Summary
- FIB Projections
- Churn
- Directions

30 Years of Routing the IPv4 Internet

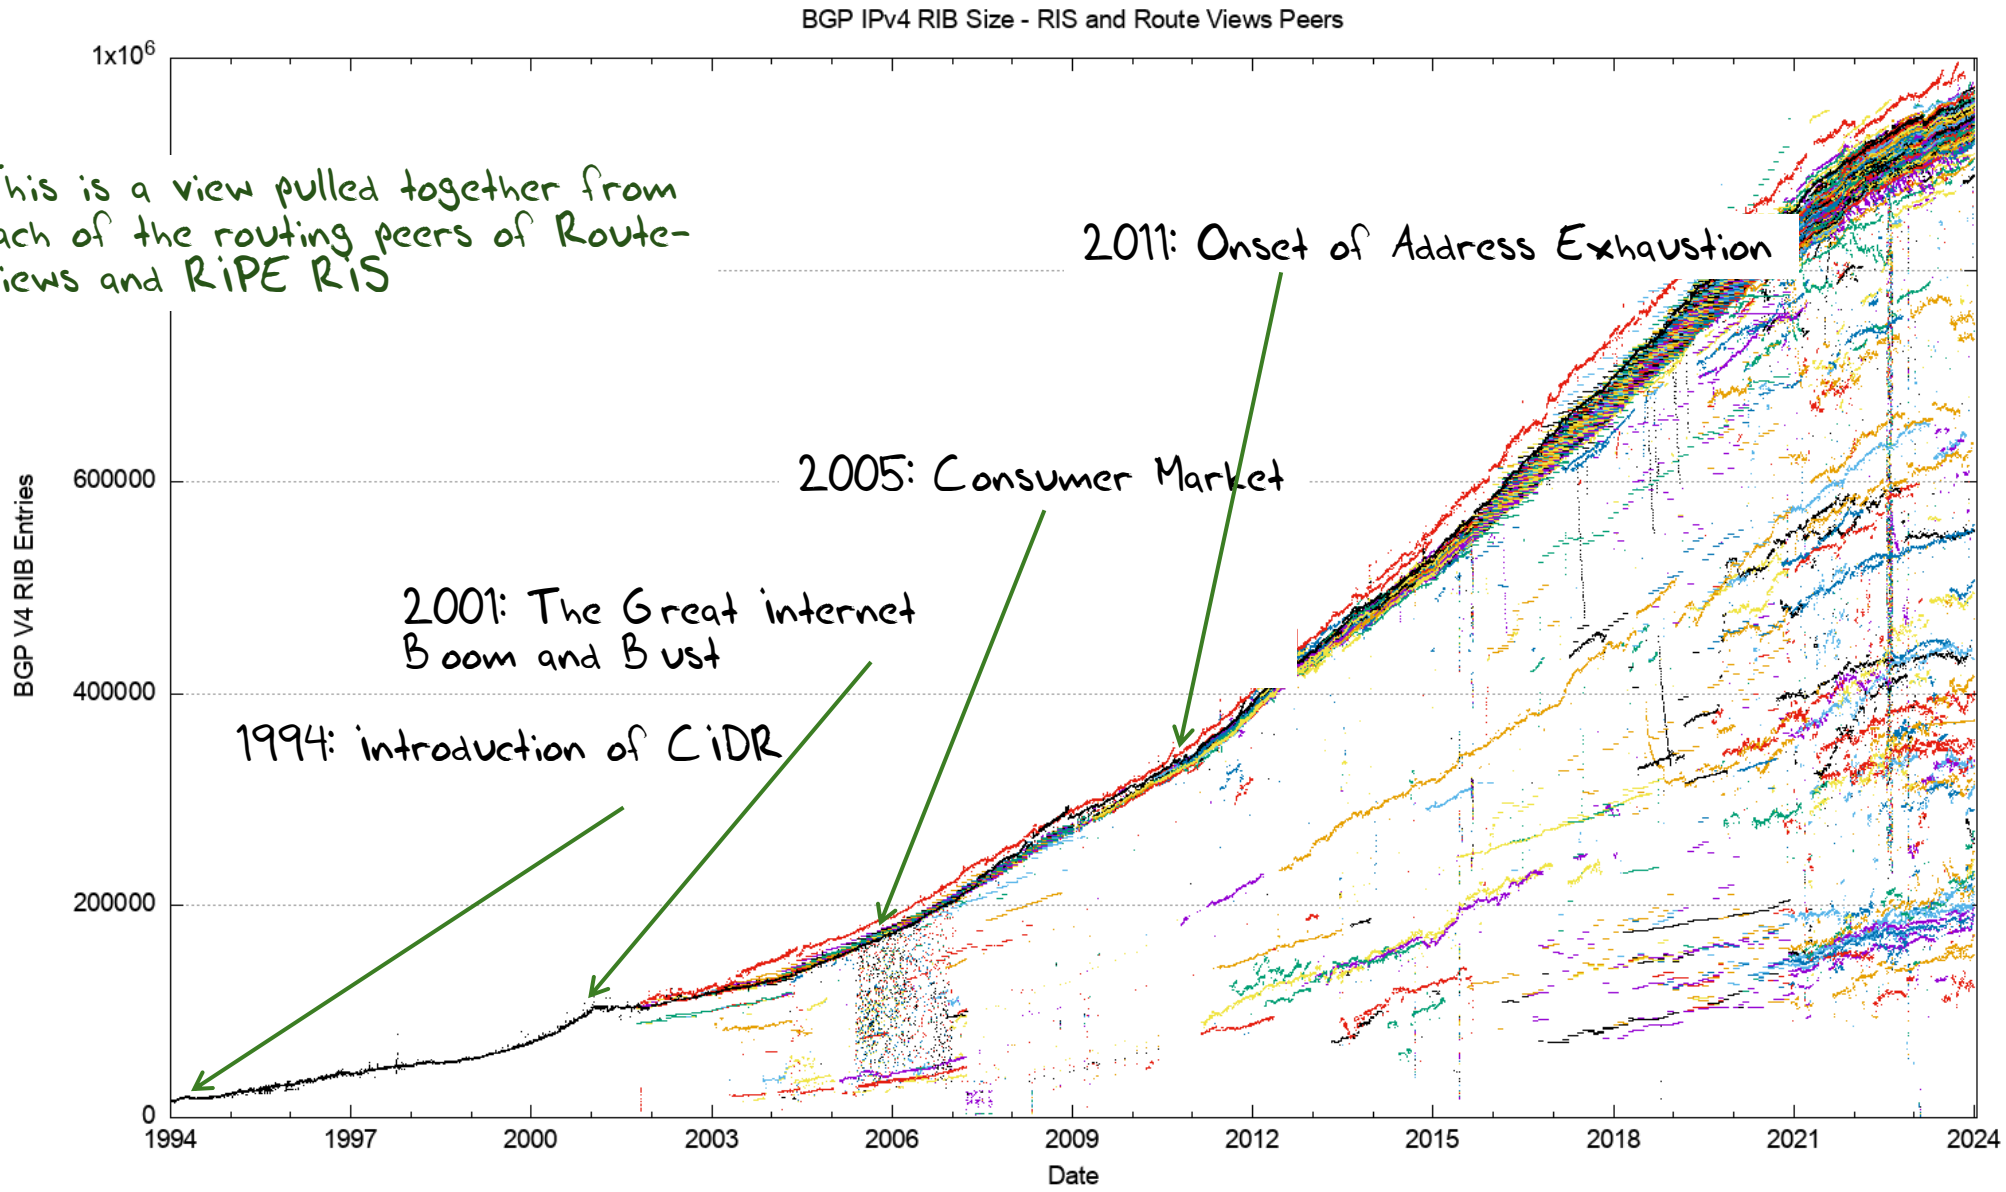

IPv4 in 2023

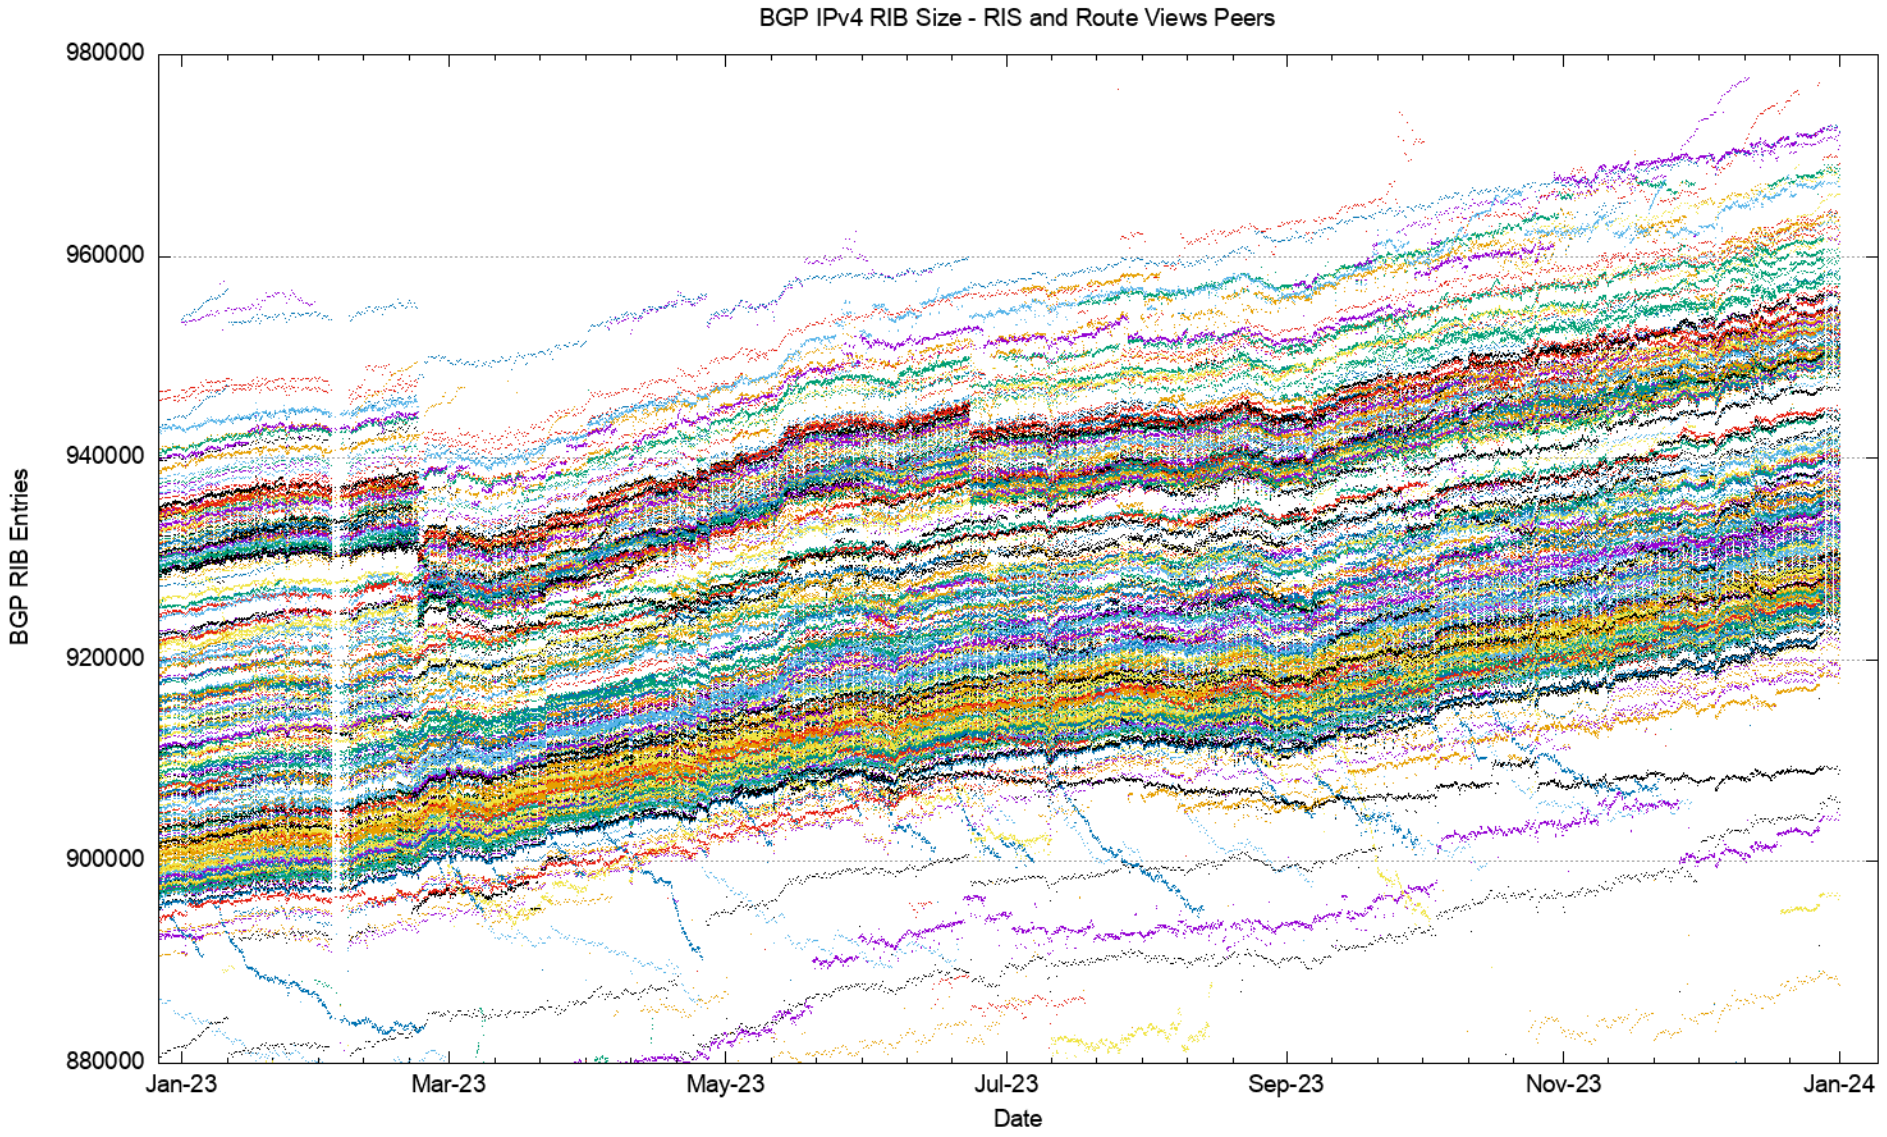
BGP IPv4 RIB Size - RIS and Route Views Peers
BGP RIB Entries
980000
960000
940000
920000
900000
880000
Jan-23
Mar-23
May-23
Jul-23
Sep-23
Nov-23
Jan-24
Date

IPv4 in 2023

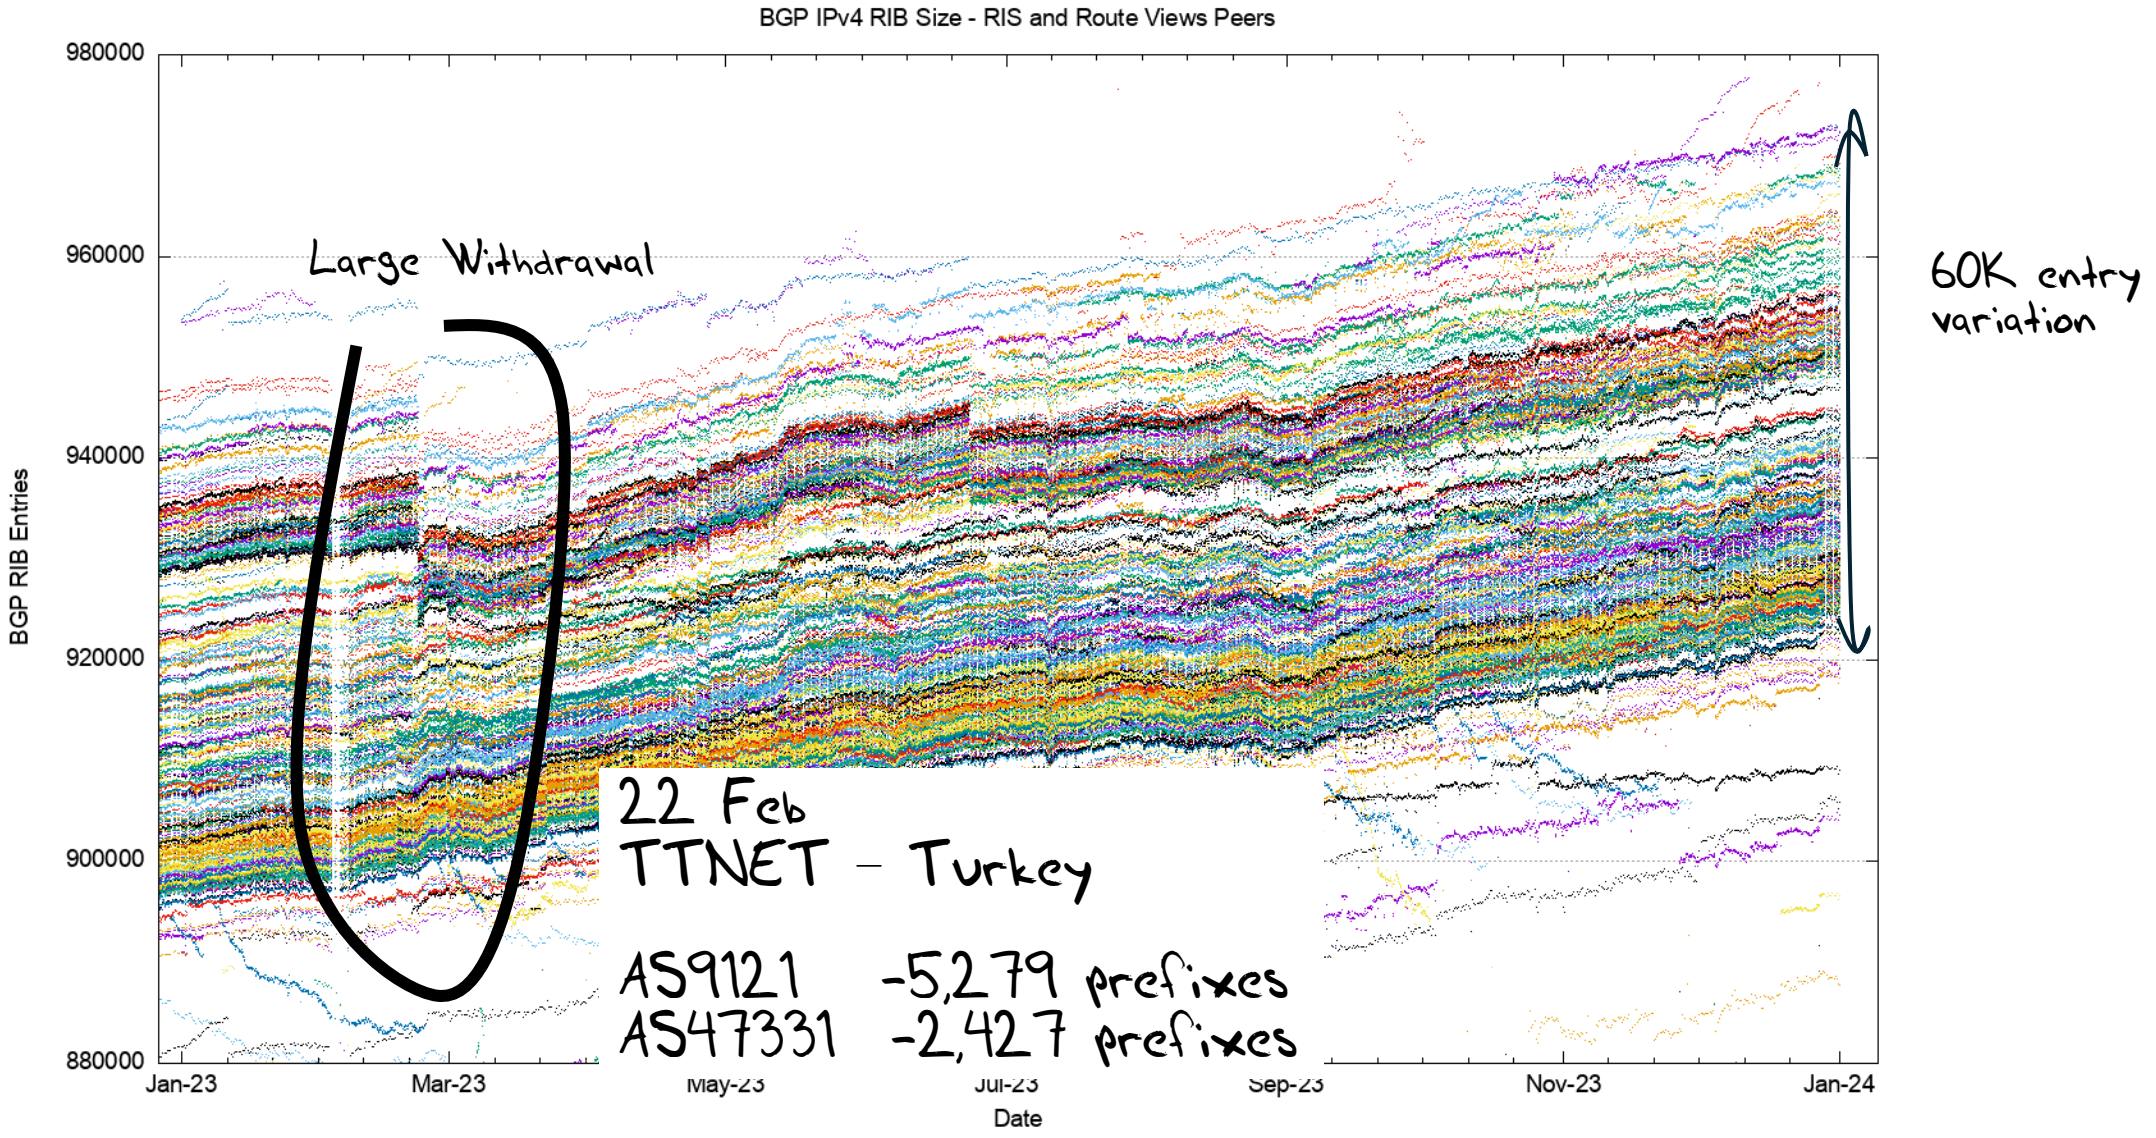
BGP IPv4 RIB Size - RIS and Route Views Peers
BGP RIB Entries
980000
960000
940000
920000
900000
880000
Jan-23
Mar-23
May-23
Jul-23
Sep-23
Nov-23
Jan-24
Date
Large Withdrawal
22 Feb
TTNET - Turkey
AS9121 -5,279 prefixes
AS47331 -2,427 prefixes
60K entry variation

Aside: What happened to TTNET?

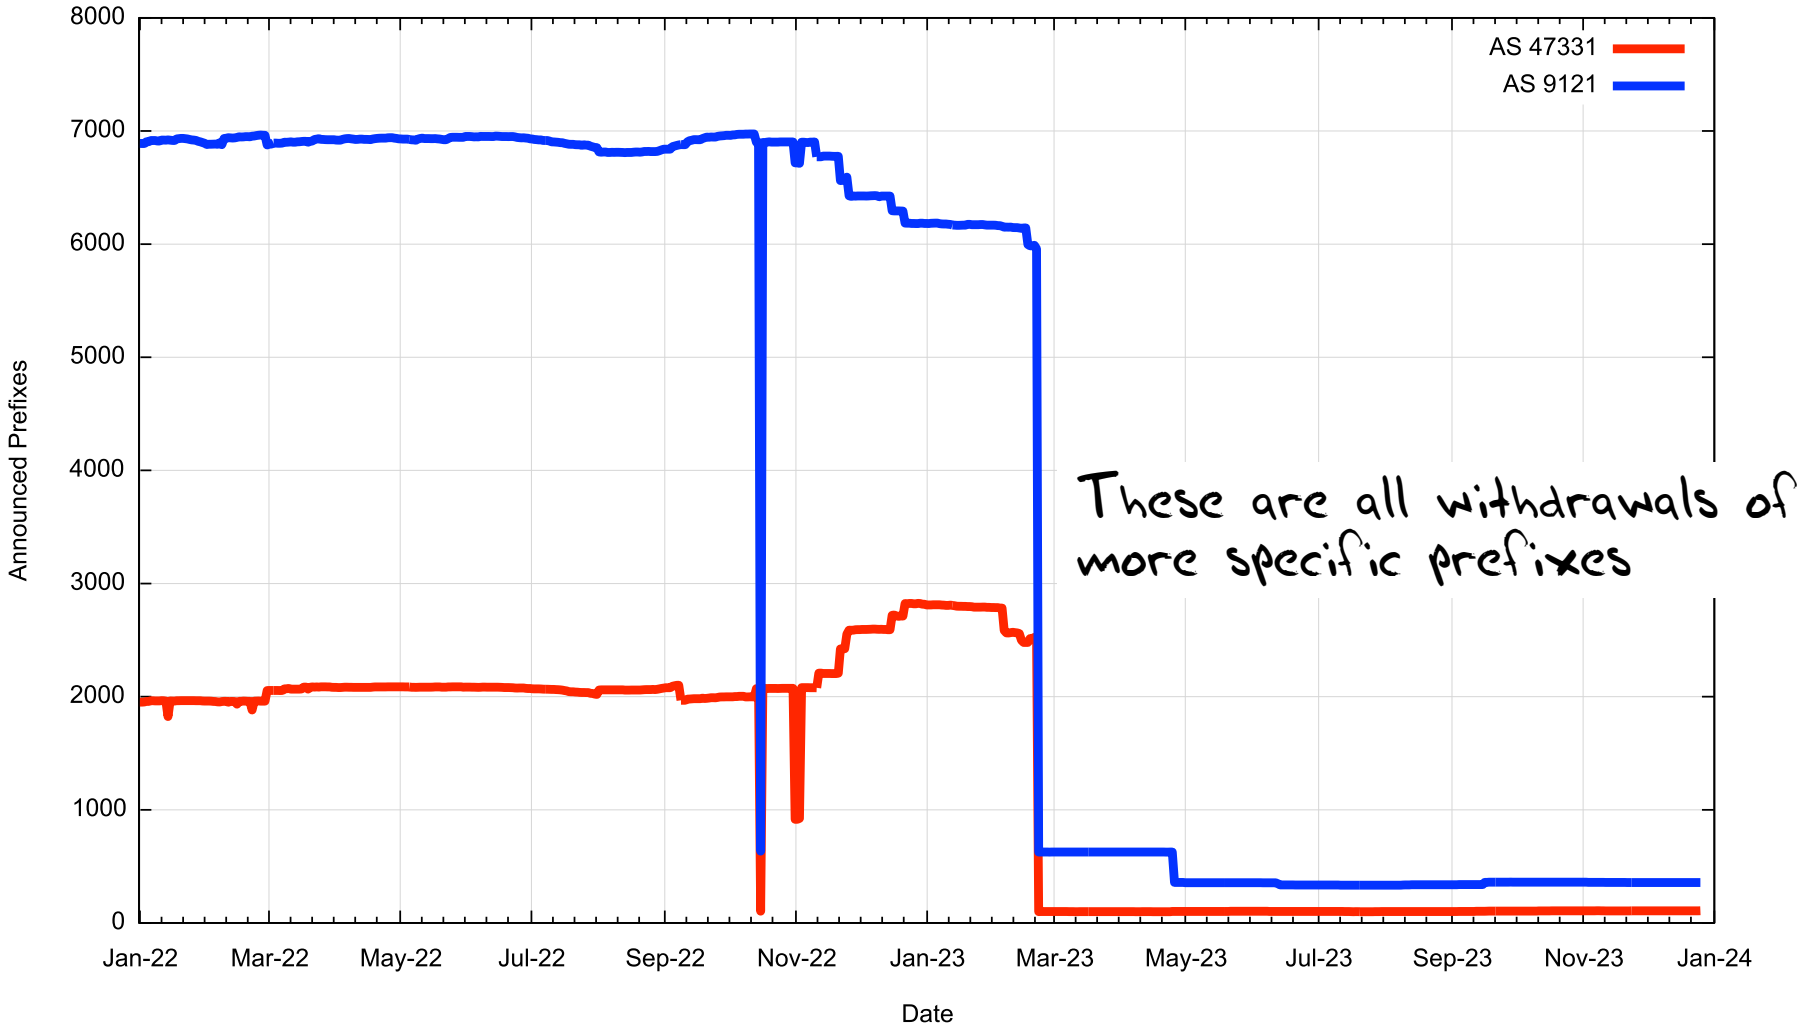

AS Prefix Count over 2023

Dropped Prefixes

AS Num	Change	Jan-23	Dec-23	Name. CC
AS1291	**-5,826**	6,183	357	TTNET, TR
AS47331	**-2,703**	2,811	108	TTNET, TR
AS6849	**-1,183**	2,251	1,068	UKRTELNET, UA
AS1239	**-784**	1,204	420	SPRINTLINK, US
AS209	**-631**	2,343	1,712	CENTURYLINK, US
AS1289	**-629**	692	63	HOTNET, IL
AS9394	**-568**	1,052	484	CTT, CN
AS135887	**-568**	1028	460	Telstra Belong, AU
AS35908	**-482**	770	288	VPLSNET, US
AS40676	**-418**	806	388	Psychz, US

Added Prefixes

AS Num	Change	Jan-23	Dec-23	Name. CC
AS18403	**3,034**	1,499	4,533	FPT Telecom, VN
AS16509	**1,709**	7,761	9,470	Amazon-O2, US
AS367	**1,403**	1,558	2,961	DNIC, US
AS44477	**1,376**	77	1,453	STARK, GB
AS8151	**1,221**	1,939	3,160	UNINET, MX
AS3737	**1,122**	26	1,148	PTD, US
AS140292	**1,079**	1,258	2,337	China Telecom, Jiangsu, CN
AS207990	**1,012**	116	1,128	HR, IN
AS9009	**857**	2,650	3,507	M247, RO
AS4155	**846**	-	846	USDA-1, US

30 Years of IPv4 Advertised Address Span

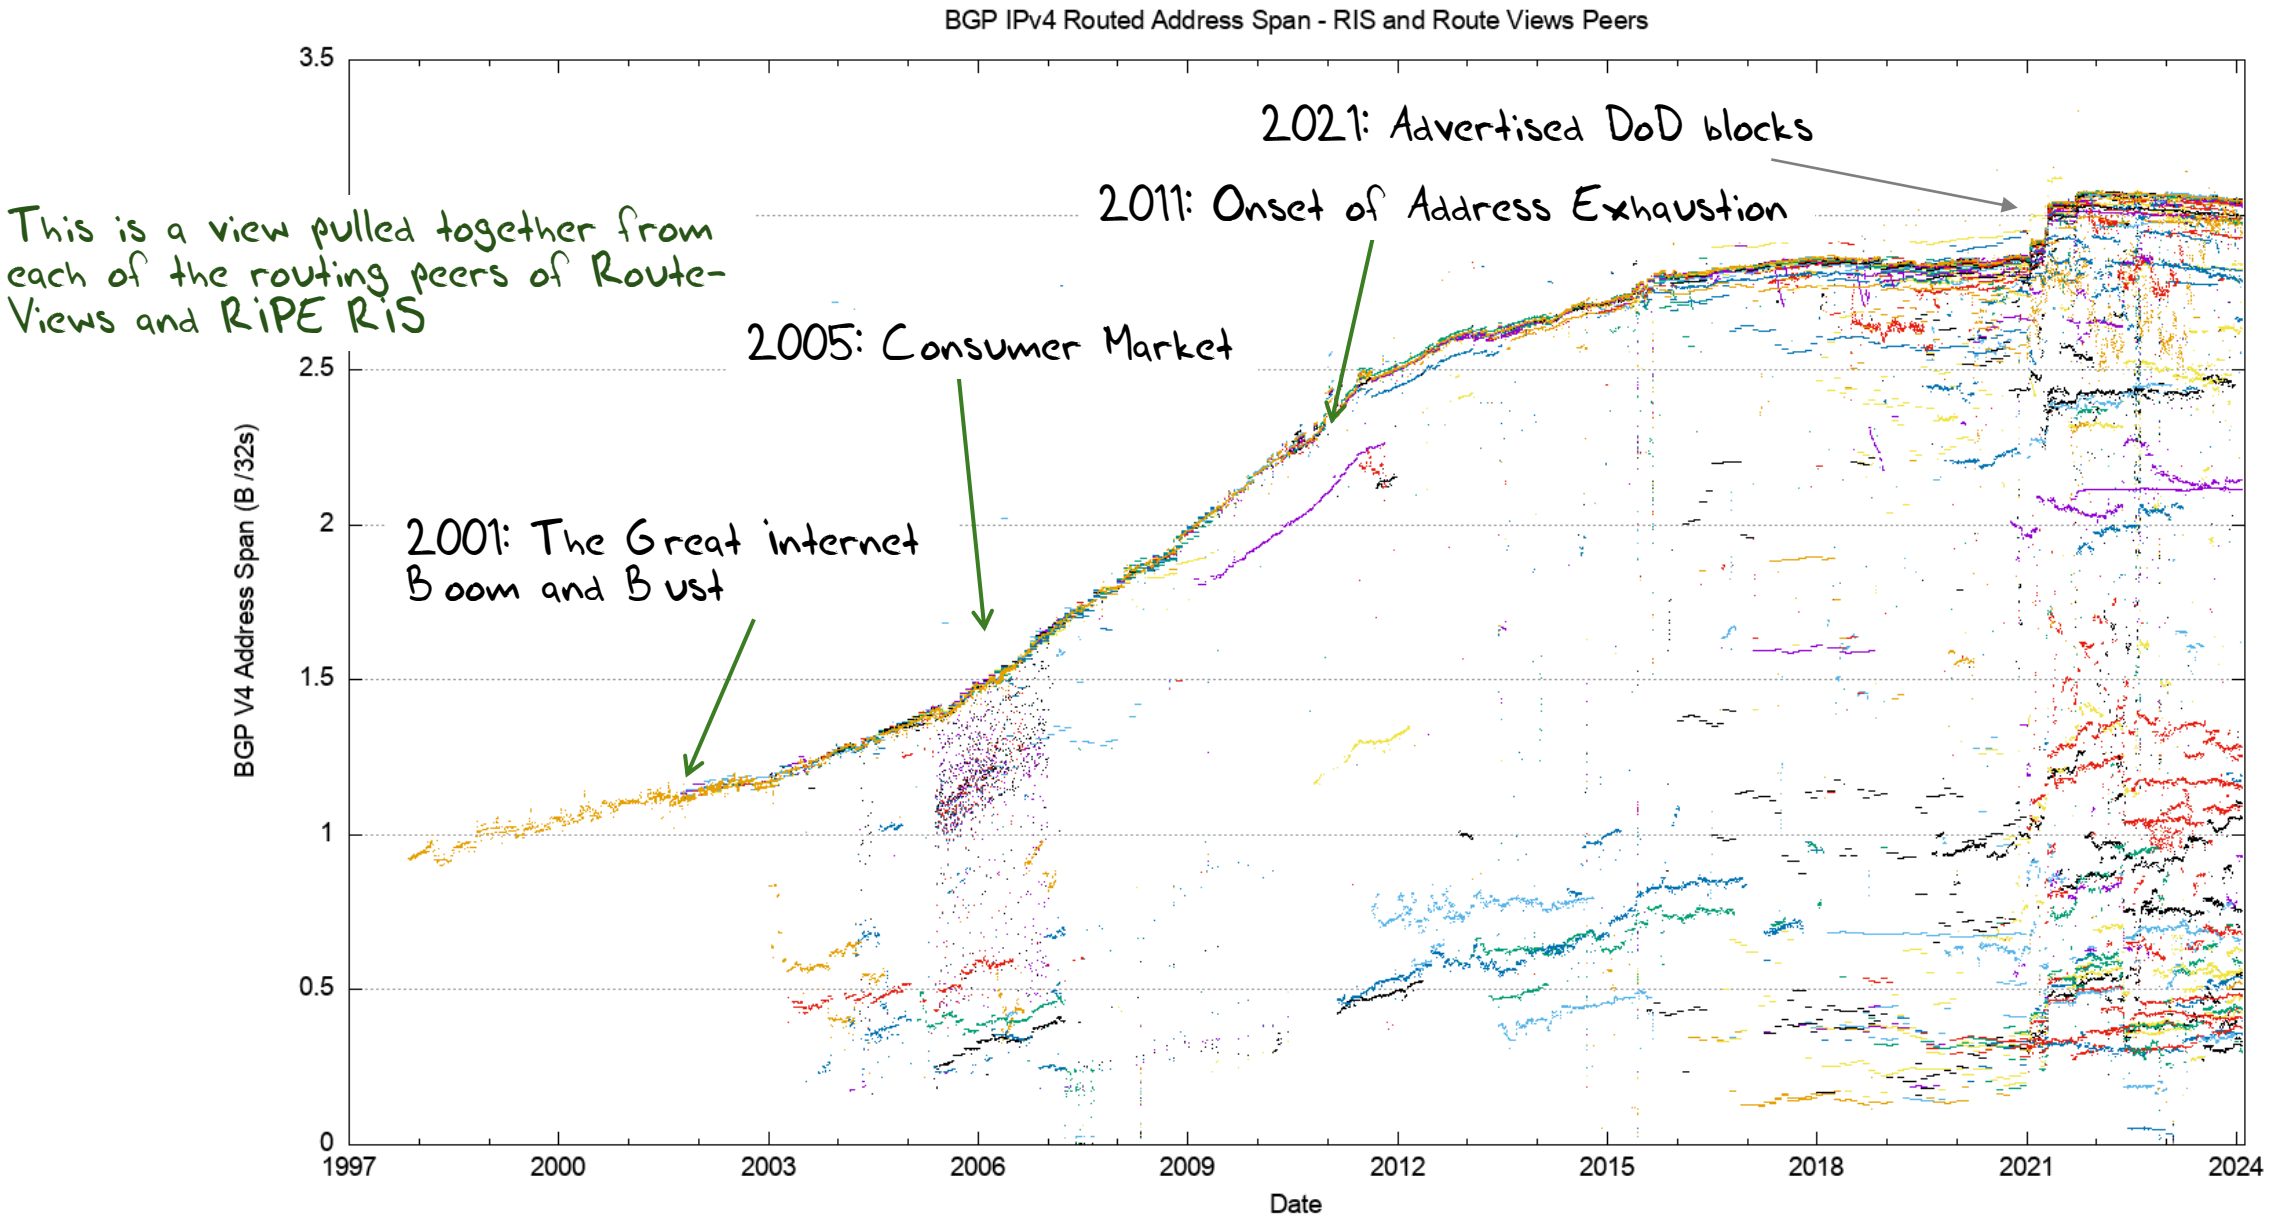

2023 in Detail

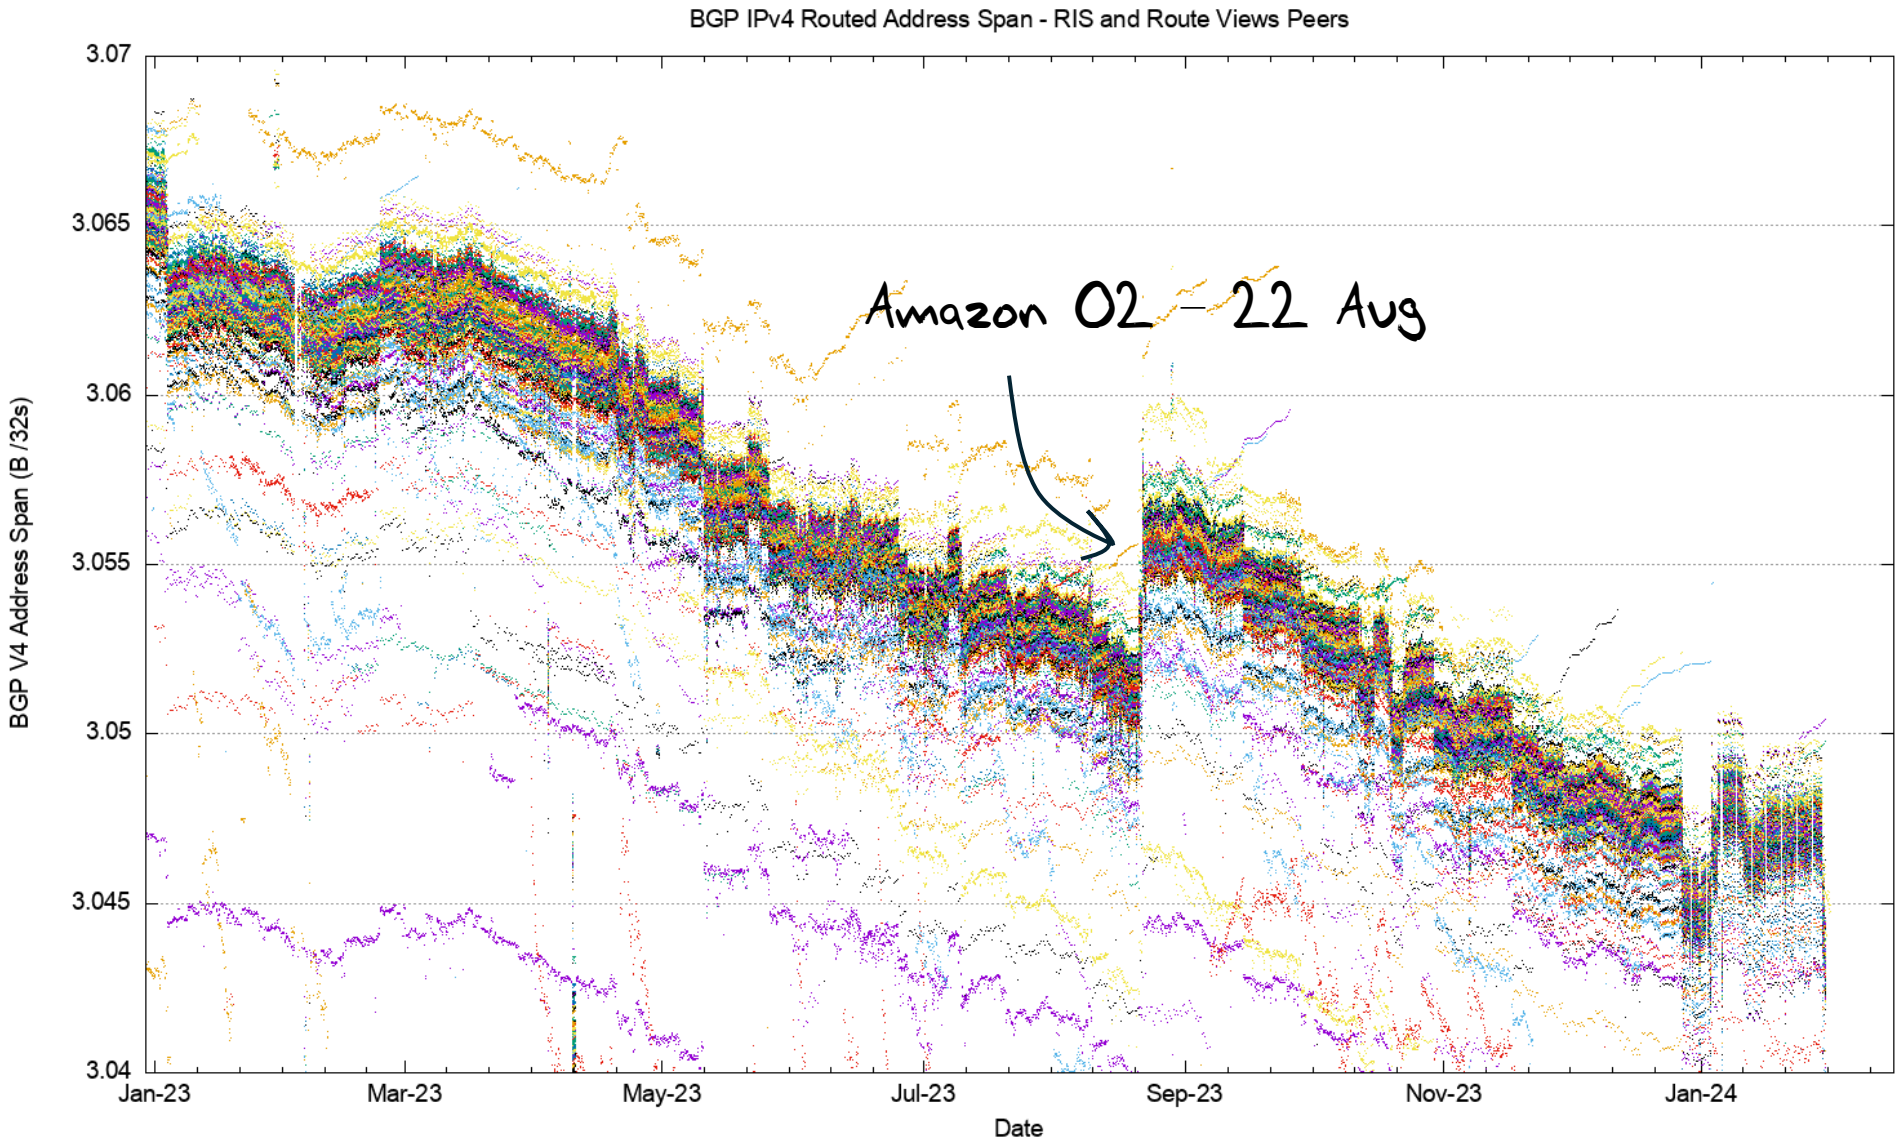

2023: Assigned vs Recovered Addresses

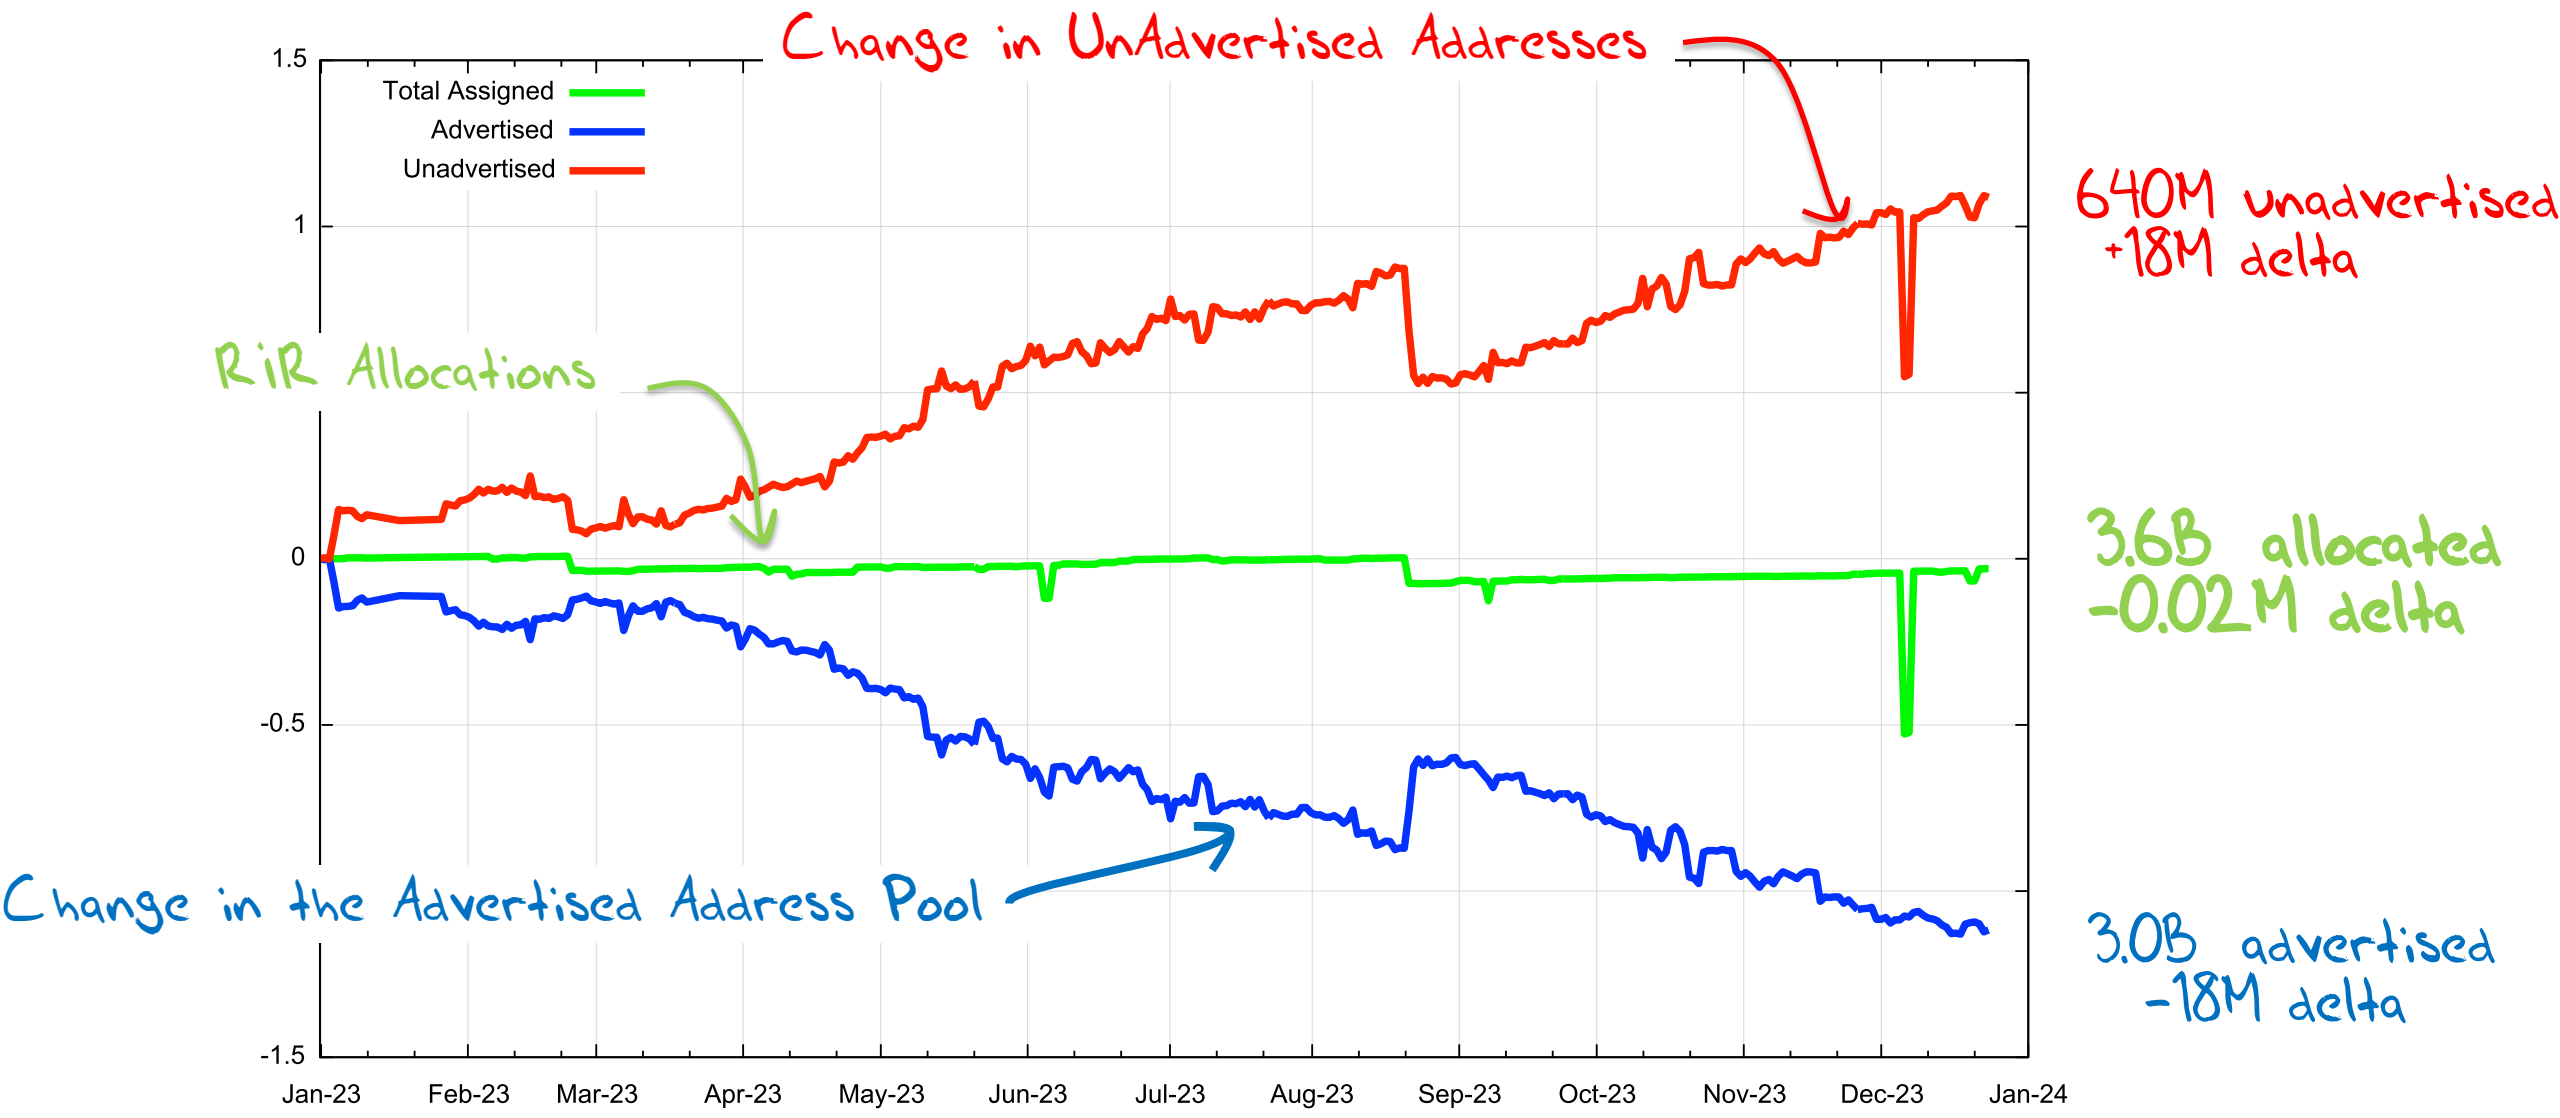

2023: Assigned vs Recovered Addresses

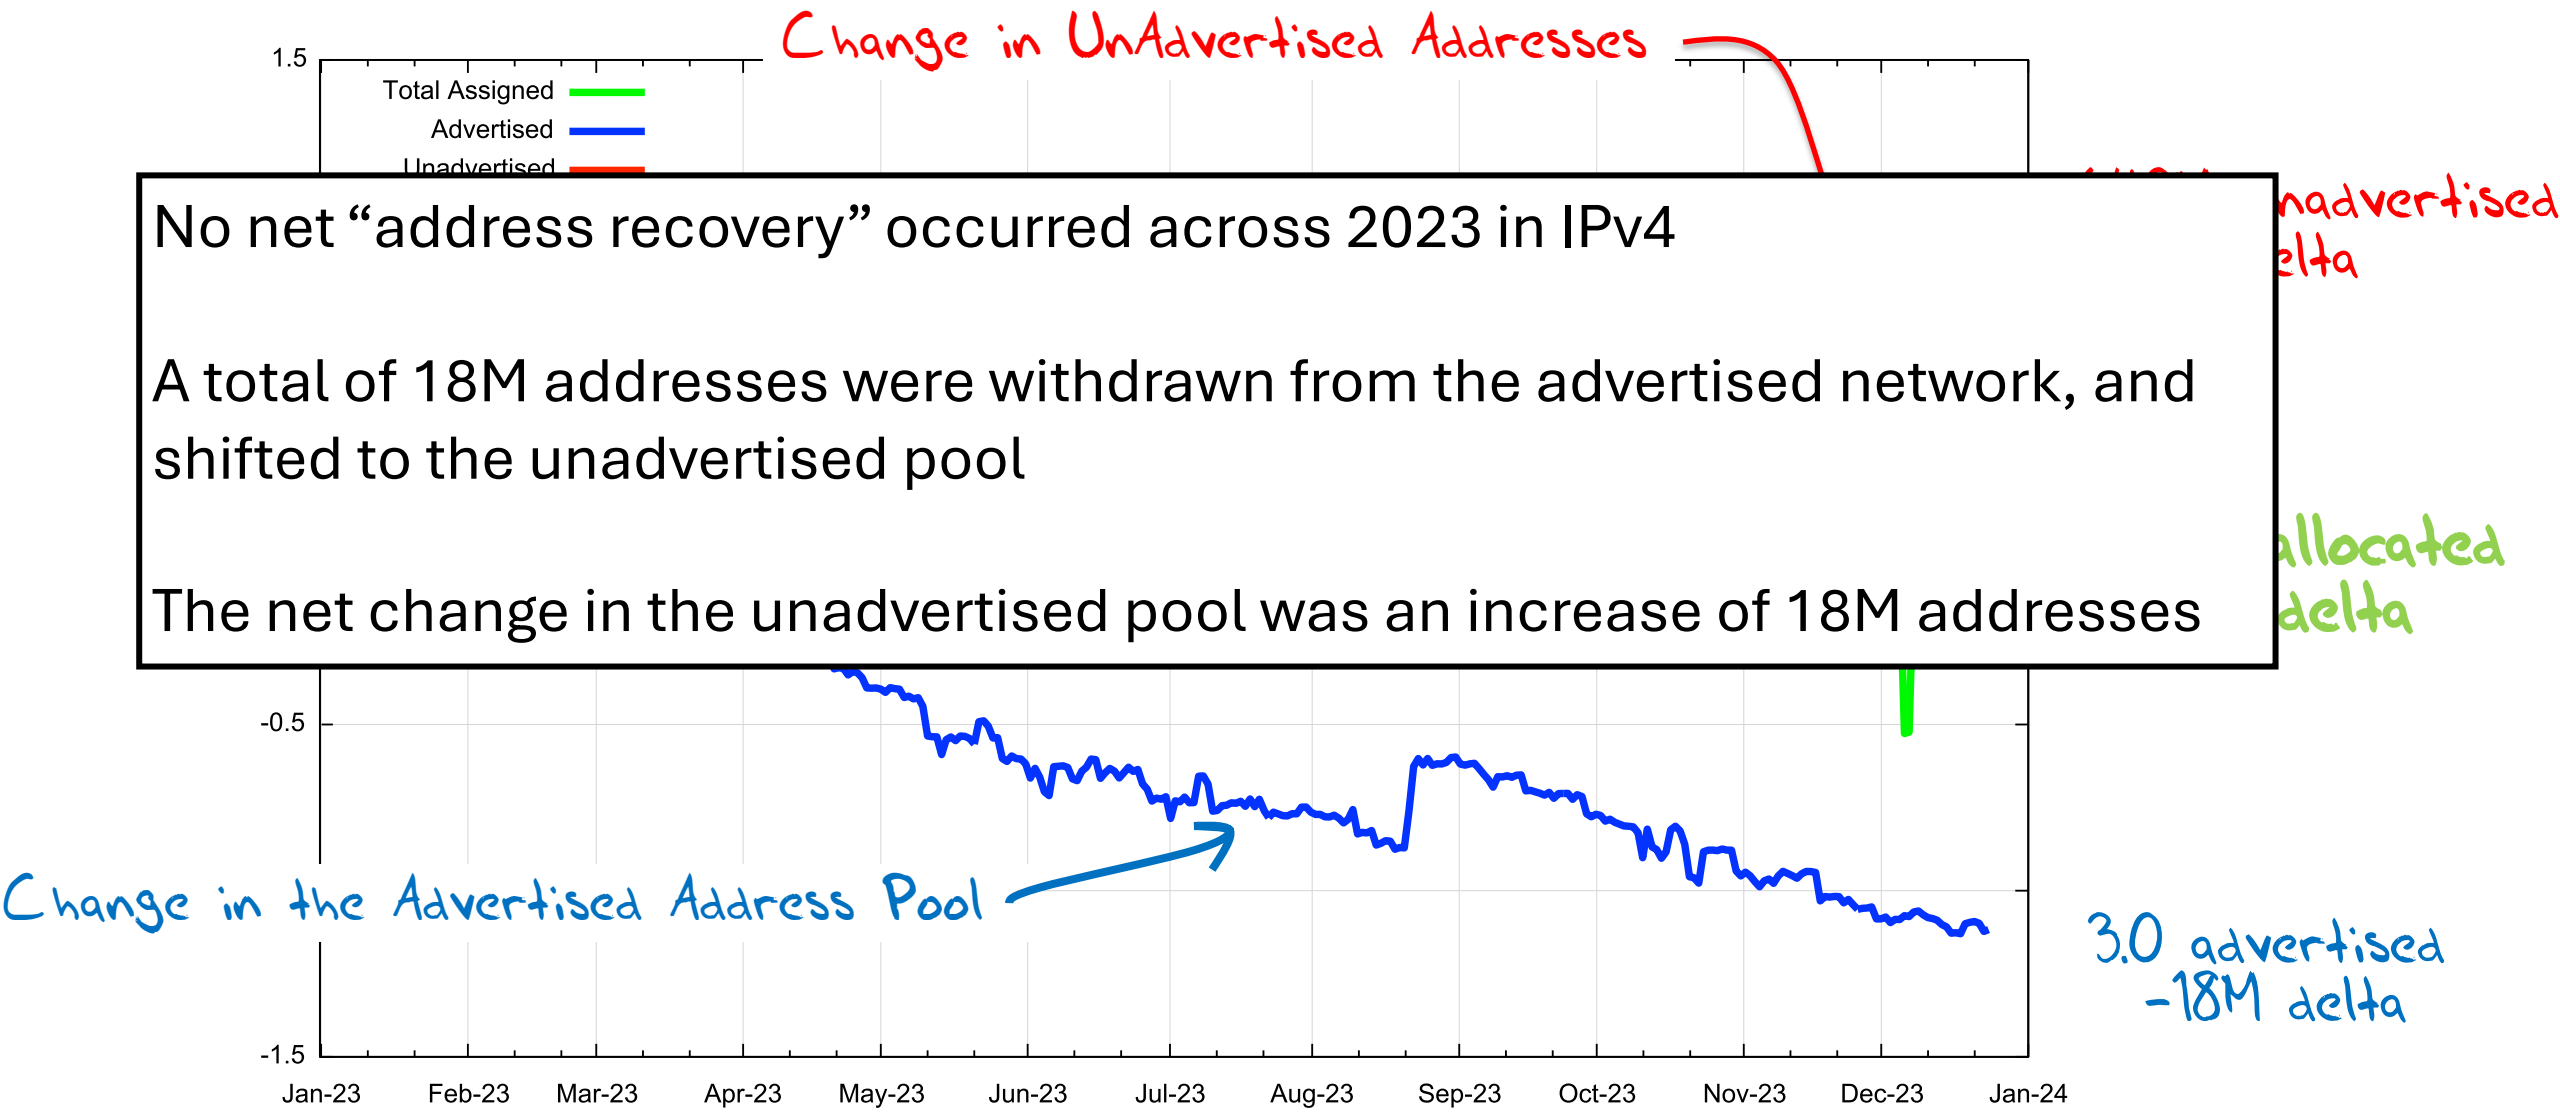

No net “address recovery” occurred across 2023 in IPv4

A total of 18M addresses were withdrawn from the advertised network, and shifted to the unadvertised pool

The net change in the unadvertised pool was an increase of 18M addresses

2018-2023: 6 Year Assigned vs Recovered Addresses

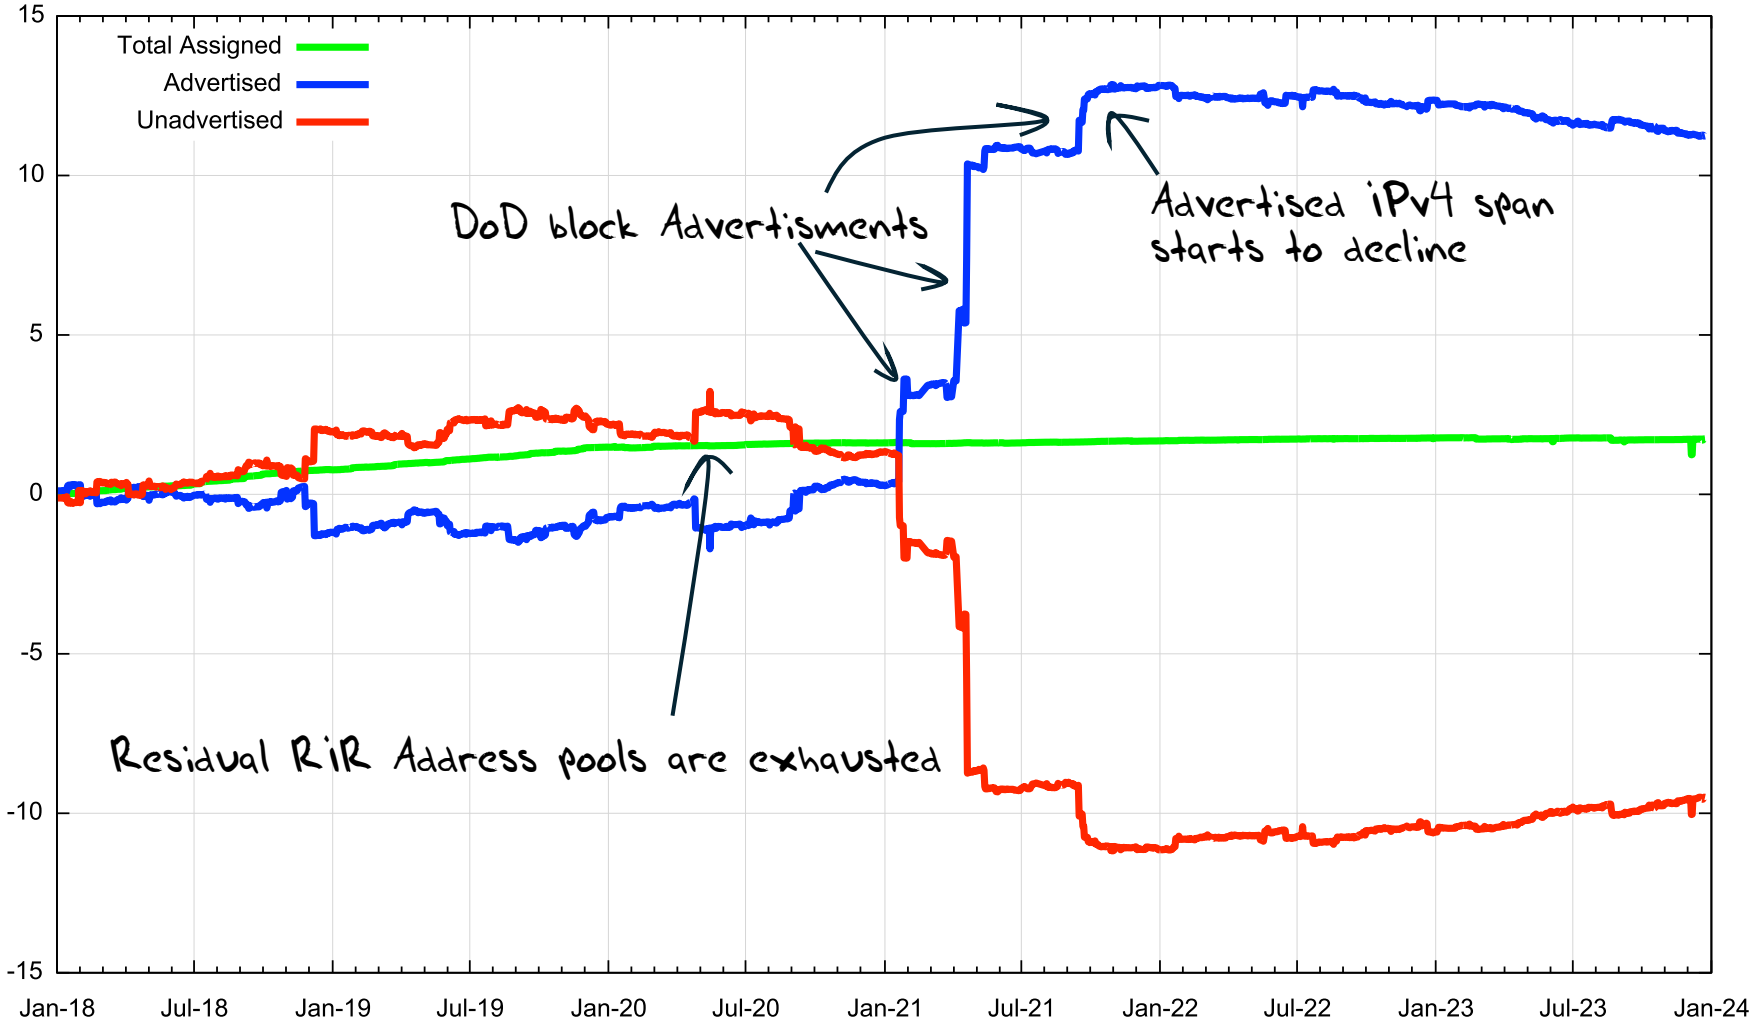

Advertised Span per Origin AS over 2023

Reduced Advertised Address Span

AS Num	Change	Jan-23	Dec-23	Name. CC
1239	**-5,958,912**	18,021,376	12,062,464	SPRINTLINK, US
9394	**-2,751,232**	19,795,968	17,044,736	CTTNET, CN
10455	**-2,163,456**	4,590,336	2,426,880	LUCENT, US
7018	**-2,095,360**	98,855,168	96,759,808	ATT-INTERNET4, US
4249	**-1,900,544**	8,585,216	6,684,672	LILLY-AS, US
47331	**-1,550,336**	1,687,808	137,472	TTNET, TR
16625	**-1,514,496**	7,378,432	5,863,936	AKAMAI-AS, US
9105	**-1,490,944**	2,868,480	1,377,536	TISCALI, GB
15169	**-1,236,992**	10,250,752	9,013,760	GOOGLE, US
7922	**-1,152,512**	71,294,720	70,142,208	COMCAST, US

Increased Advertised Address Span

AS Num	Change	Jan-23	Dec-23	Name, CC
749	**17,830,400**	207,162,880	224,993,280	DNIC, US
367	**9,184,256**	6,606,592	15790848	DNIC, US
11003	**4,165,888**	458,752	4,624,640	PANDG, US
16509	**2,304,512**	43,574,272	45,878,784	AMAZON-02, US
19901	**2,237,696**	-	2,237,696	BRSPD, US
3257	**2,133,504**	4,558,080	6,691,584	GTT, US
6167	**1,391,104**	11,270,144	12,661,248	CELLCO, US
6306	**1,317,888**	623,616	1,941,504	TELEFONICA, VE
984	**1,180,160**	4,352	1,184,512	OWS, US
29447	**1,048,576**	458,752	1,507,328	Iliad, FR

What happened in 2023 in V4?

- From the look of the routing growth plots, the growth of the size of the IPv4 network **is slowing down**
- The number of entries in the IPv4 default-free zone reached 920K – 960K by the end of 2023
- The pace of growth of the routing table was slightly lower than the rolling 5-year average, with **20,000 new entries in 2023** (was 36,000 in 2022)
- The AS position was slightly lower with **1,100 new AS's advertised in 2023** (was 1,400 in 2021)
- Transit relationships have not changed materially over 2022 for most networks
- The address range spanned by the advertised route set **declined** in 2023 by the equivalent of 1 /8
- **The overall IPv4 routing growth trends slowed down or even reversed through 2023**

What happened in 2023 in V4?

- From the look of the routing growth plots, the growth of the size of the network **is slowing down**
- The number of entries in the IPv4 default-free table by the end of 2023
- The pace of growth was slightly lower than the rolling 5-year ... 2023 (was 36,000 in 2021)
- The number of ASs was slightly lower with **1,100 new AS's advertised in 2023** (was ...)
- Transit relationships have not changed materially over 2022 for most networks
- The address range spanned by the advertised route set **declined** in 2023 by the equivalent of 1 /8
- **The overall IPv4 routing growth trends slowed down or even reversed through 2023**

Have we reached Peak IPv4?

The Highlights

- IPv4 Summary
- IPv6 Summary
- FIB Projections
- Churn
- Directions

The 20-Year View of IPv6

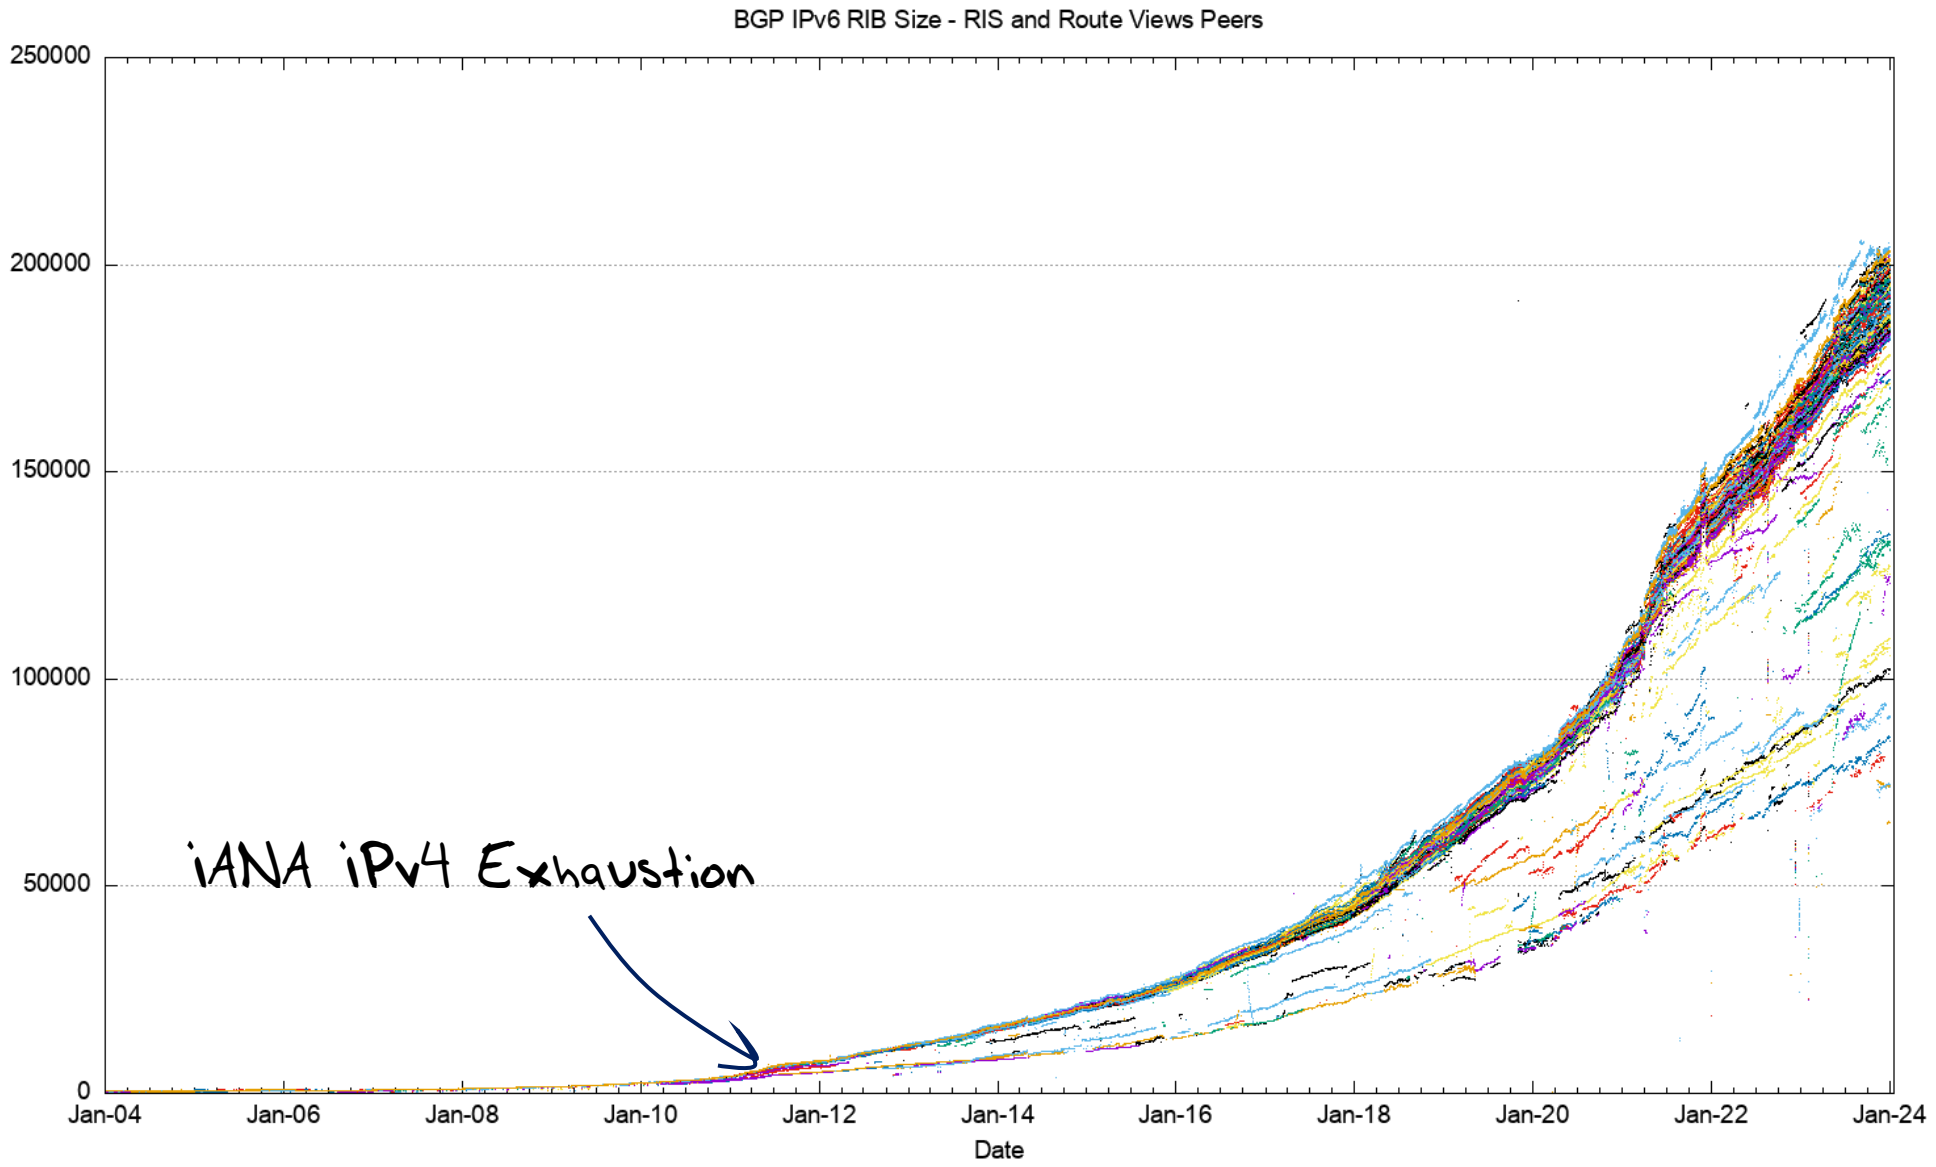

2023 IPv6 FIB in Detail

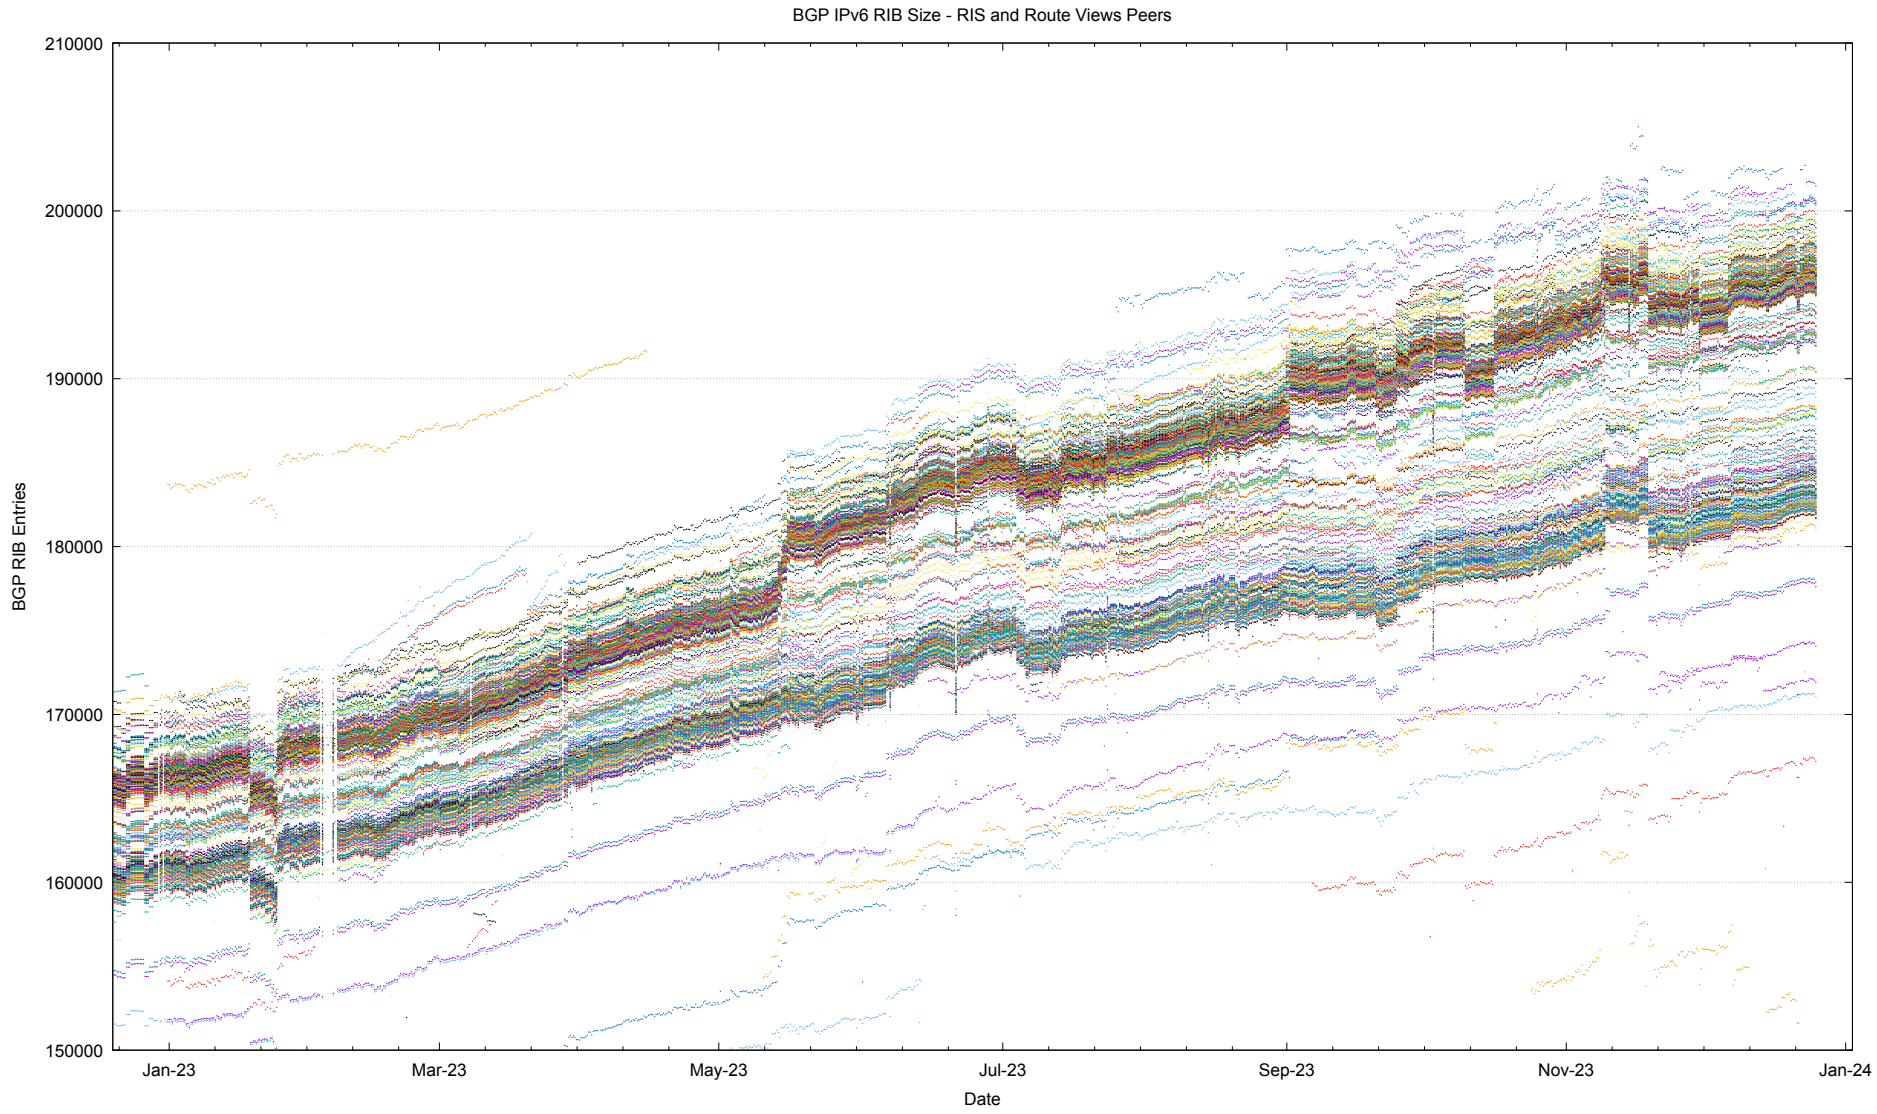

2023 IPv6 FIB in Detail

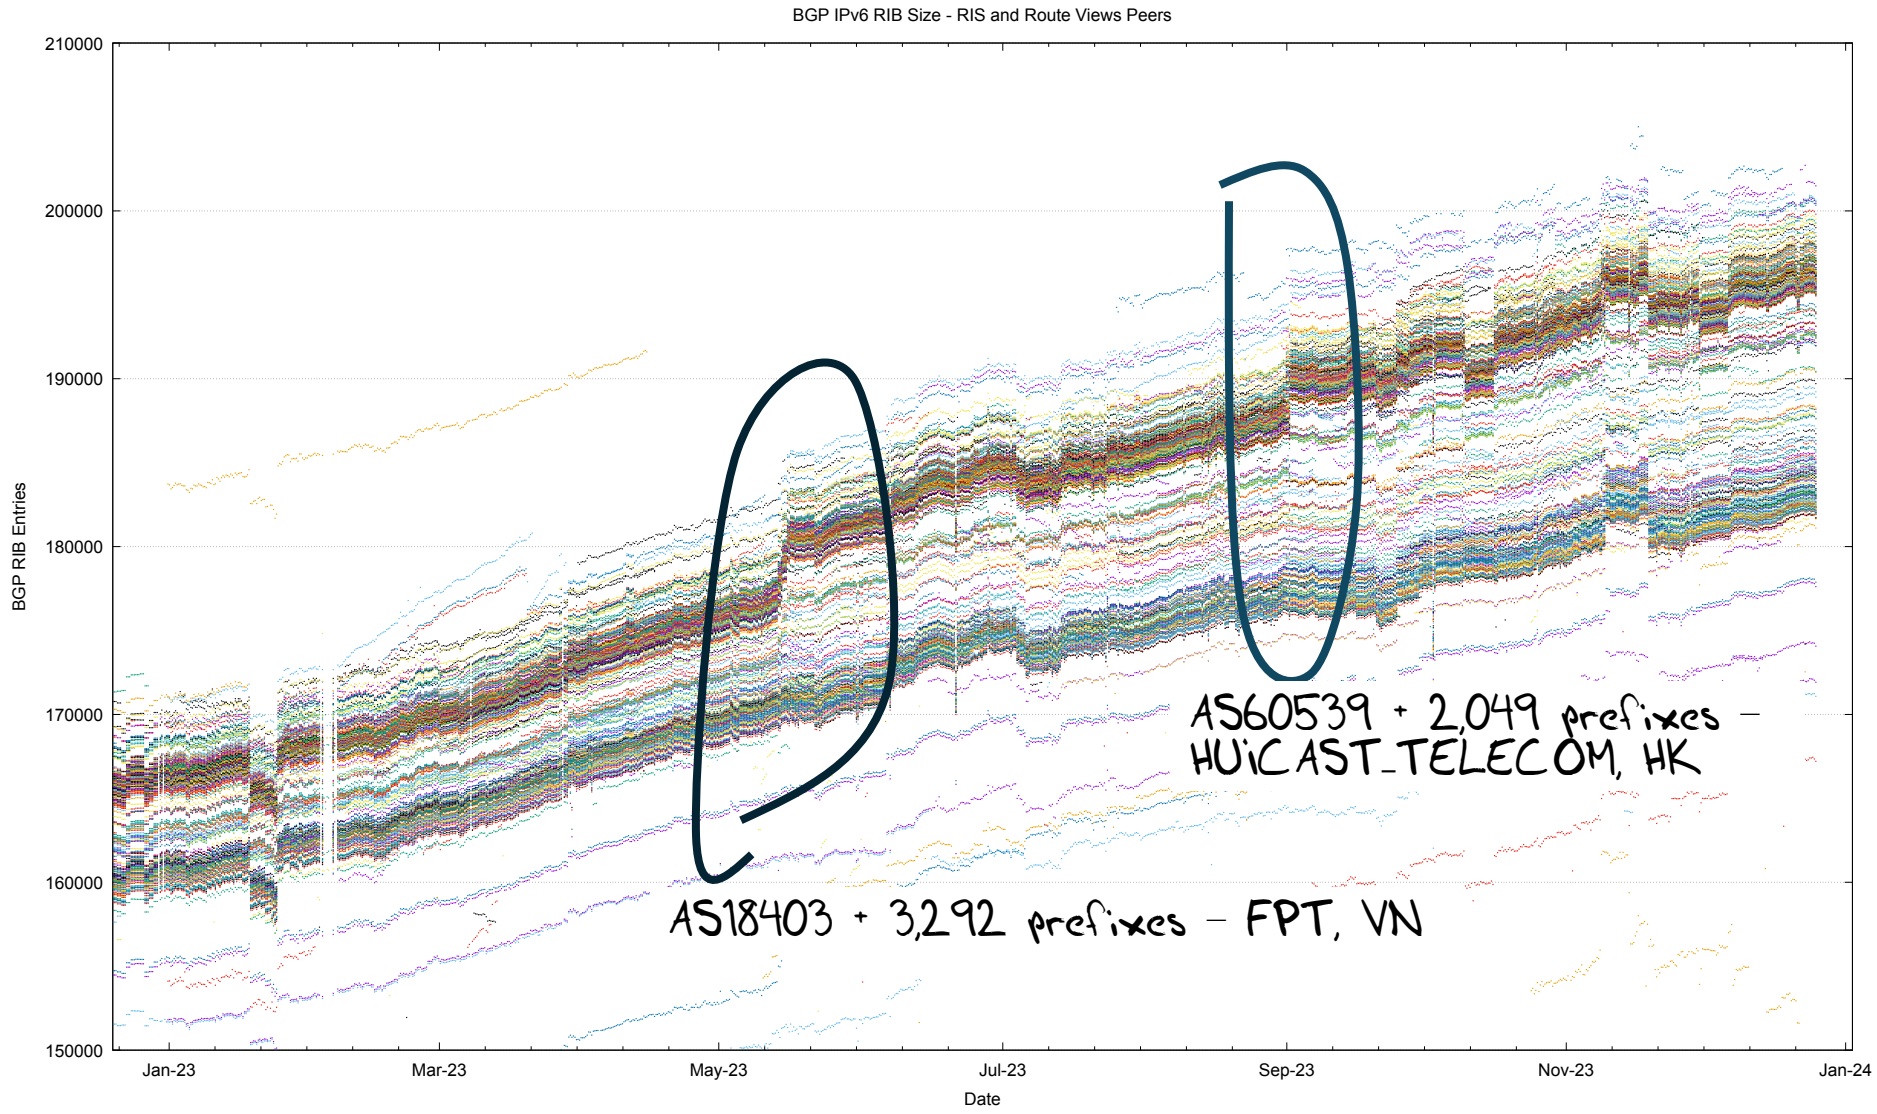

More Specifics in IPv6

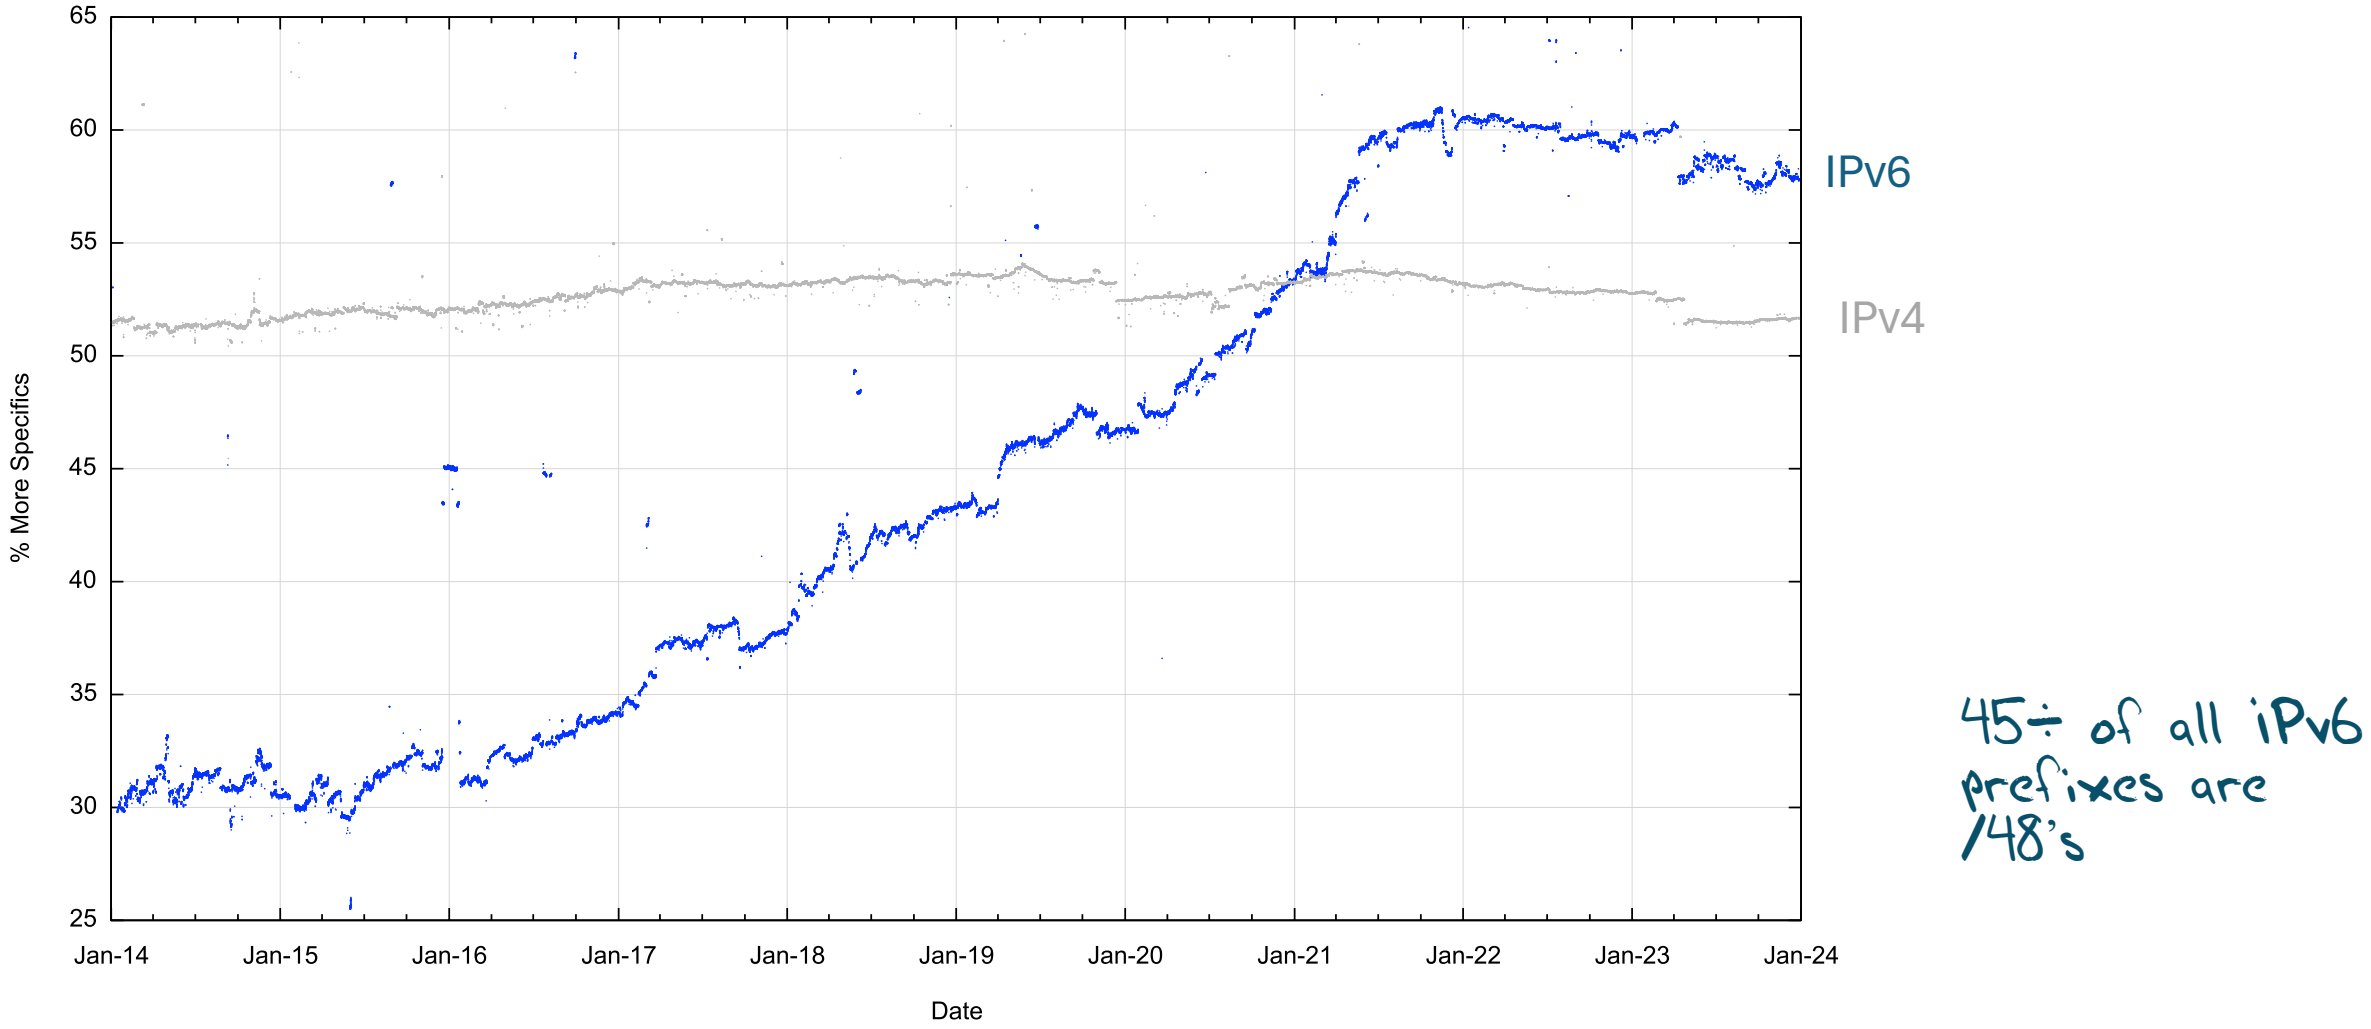

45÷ of all iPv6 prefixes are /48's

20-Year IPv6 Advertised Address Span

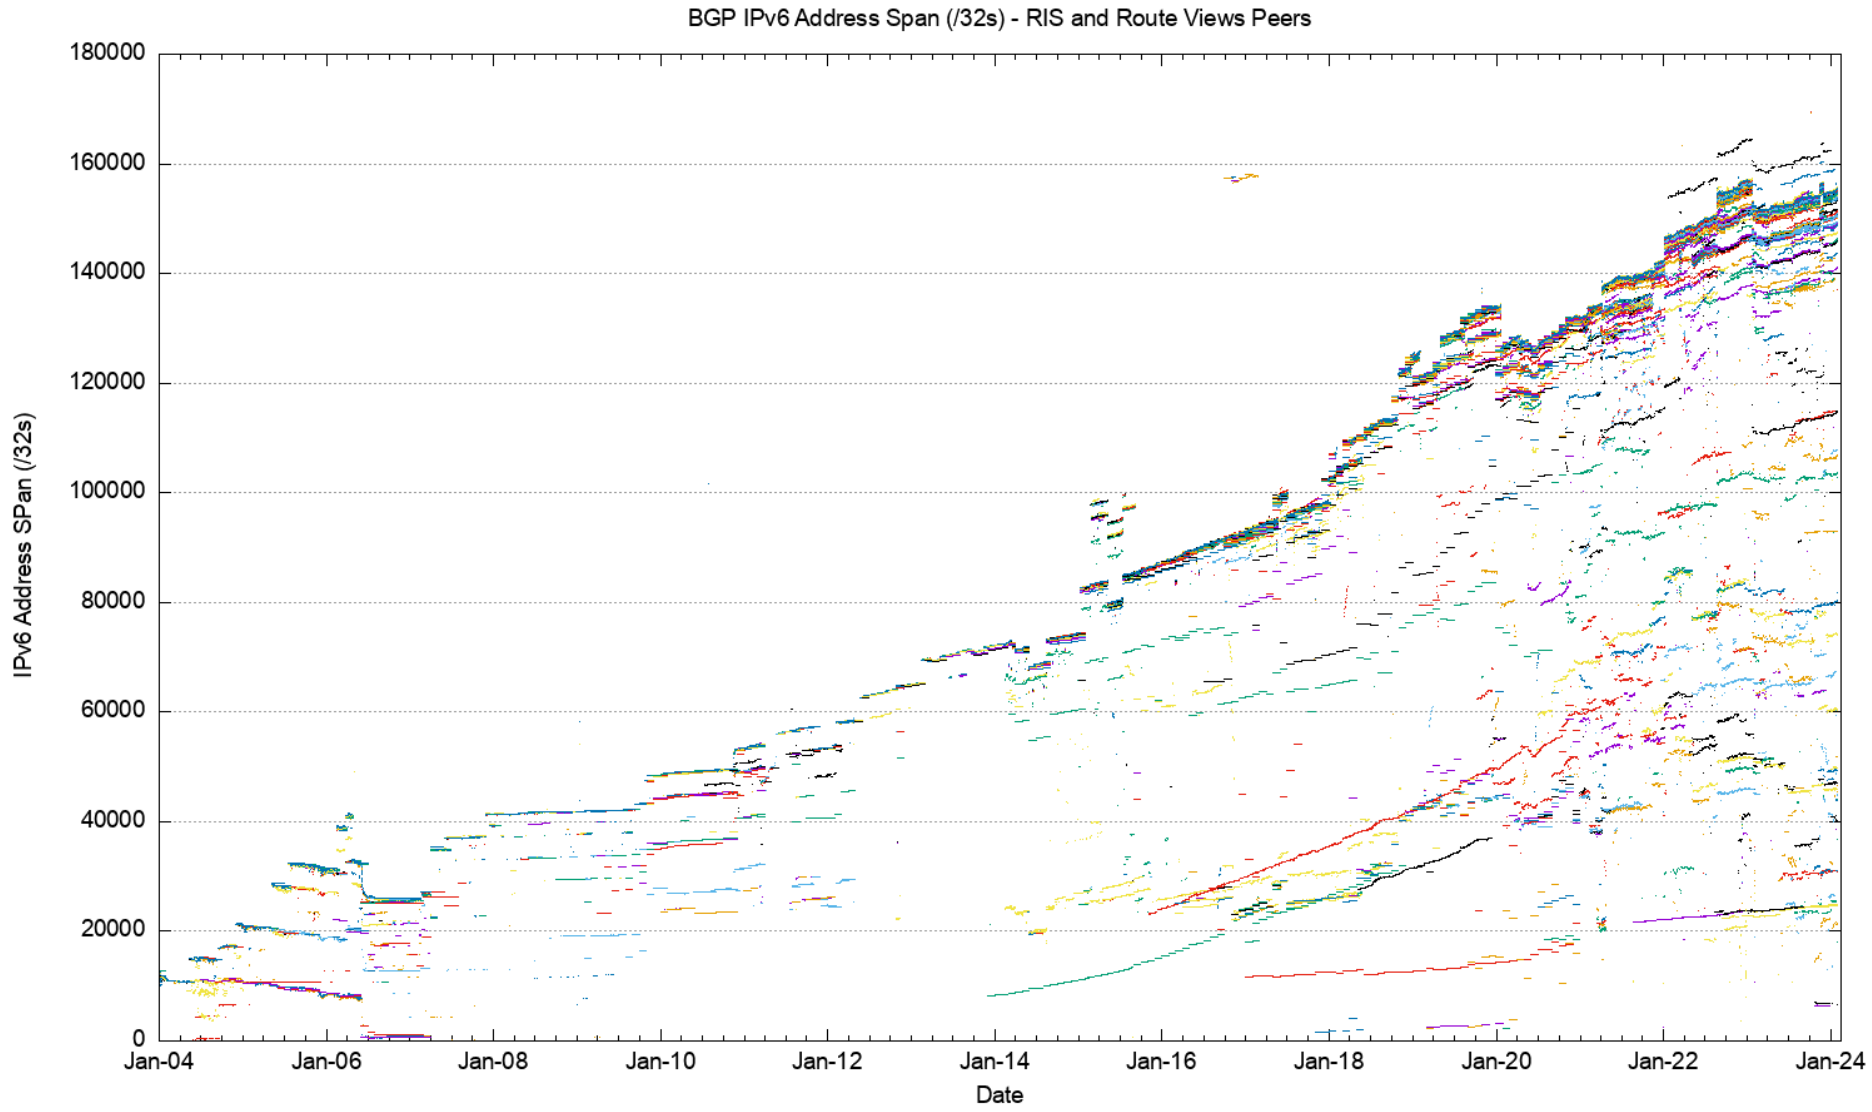

IPv6 Advertised Address Span in 2023

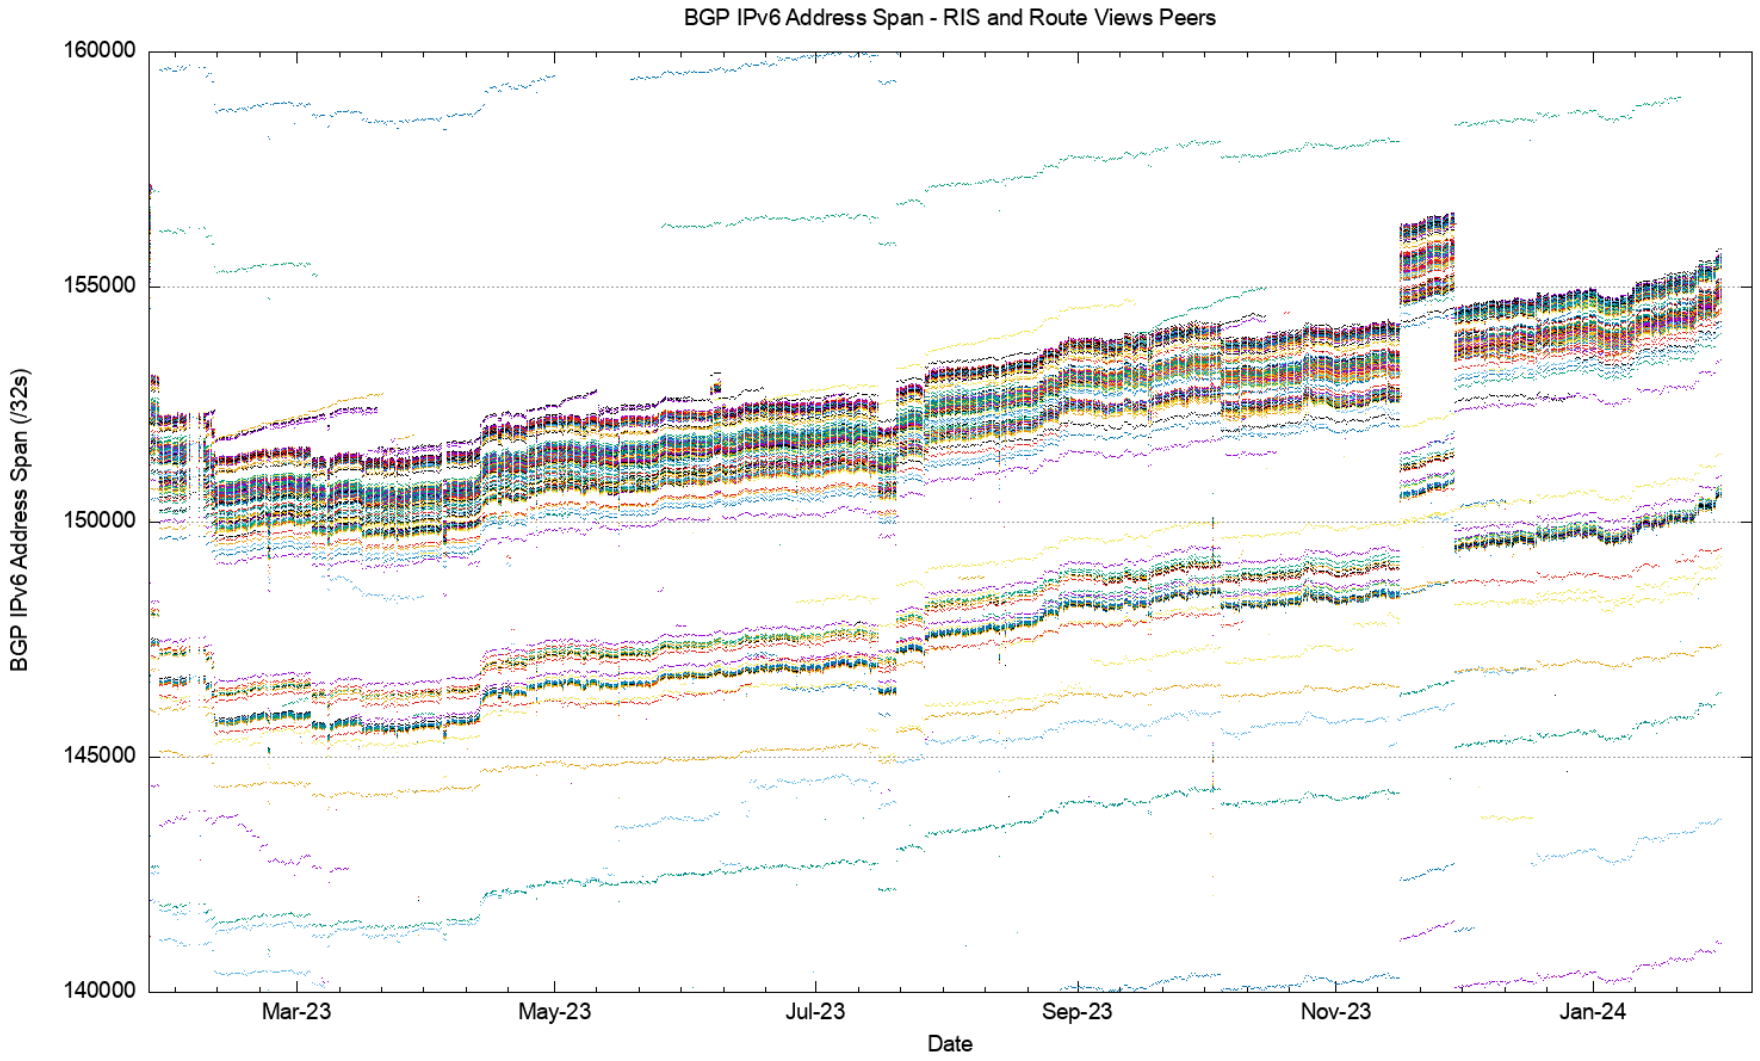

V6 in 2023

- Overall IPv6 Internet growth in terms of BGP is still increasing, and is currently at some **30,000 route entries p.a.** (17%)

  - Predominate use of /48 more specifics
  - 2,000 more AS's advertising IPv6 prefixes
  - Growth of 2,500 /32 equivalents in the advertised address span (1.6%)
  - Growth rates across 2023 are lower than 2018 – 2020 annual rates

The Highlights

- IPv4 Summary
- IPv6 Summary
- FIB Projections
- Churn
- Directions

V4 BGP Table Size Predictions

Date	RIB Size	Prediction
Jan 2019	760,000	
2020	814,000	
2021	856,000	
2022	906,000	
2023	942,000	
2024	944,0000	949,000
2025		958,000
2026		958,000
2027		947,000
2028		928,000
2029		898,000

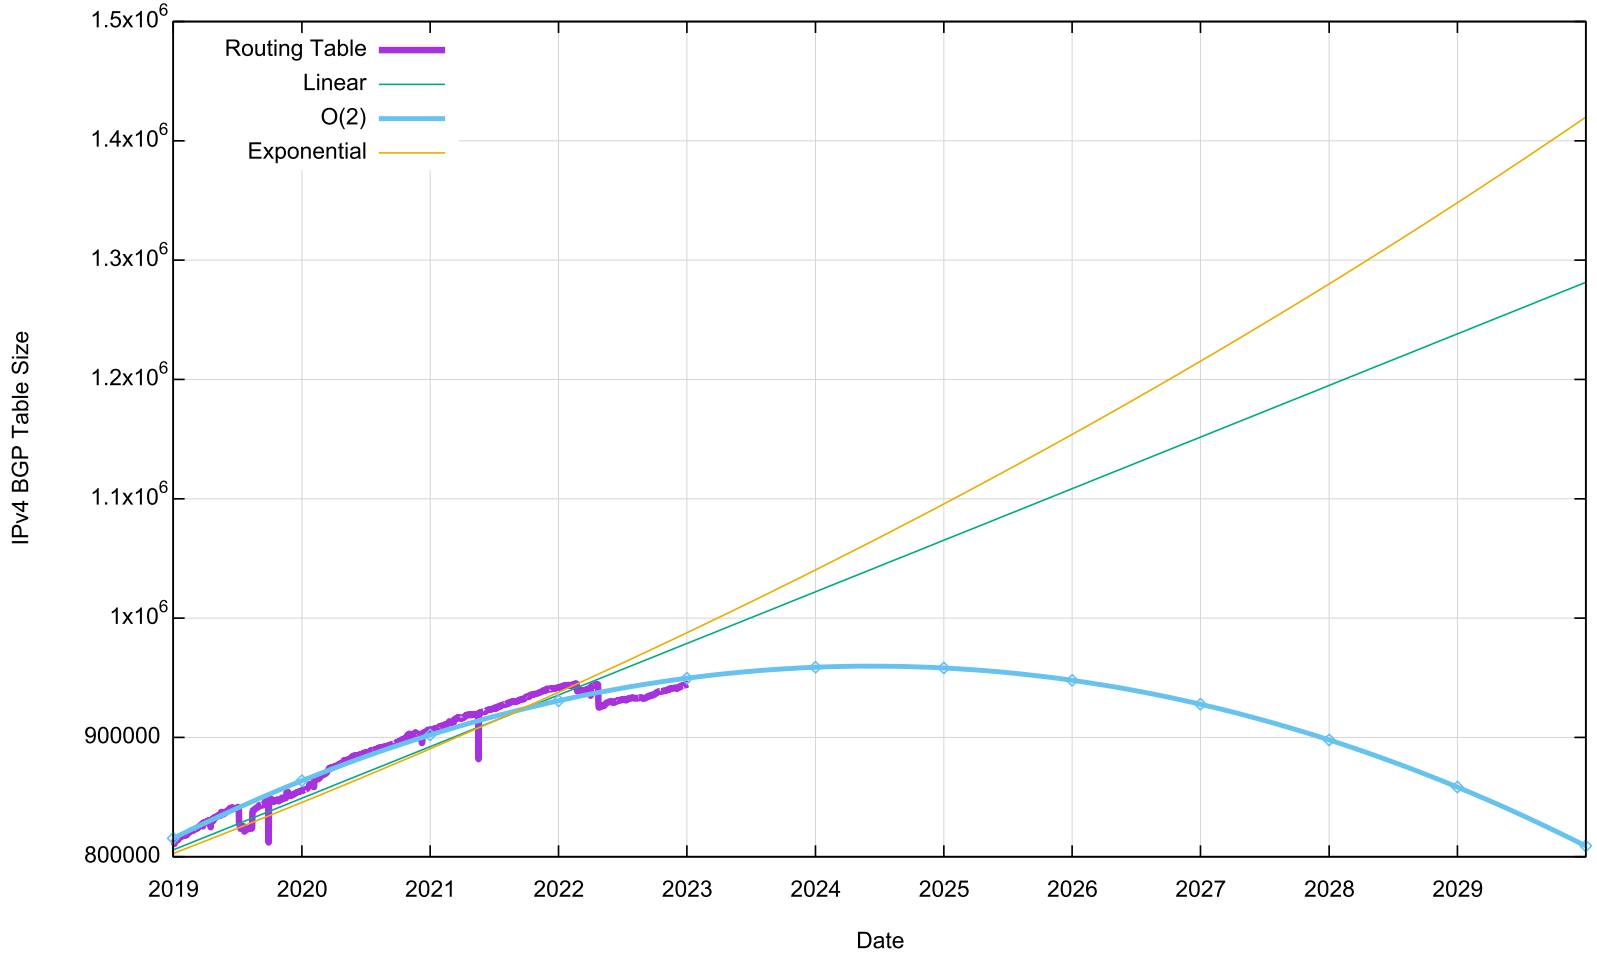

V4 BGP Table Size Predictions

Date	RIB Size	Prediction
Jan 2019	760,000	
2020	814,000	
2021		
2022		
2023		
2024		
2025		958,000
2026		958,000
2027		947,000
2028		928,000
2029		898,000

The best fit projection model to the past 4 years of FIB data is a polynomial function that peaks in 2024 at just under 1M entries

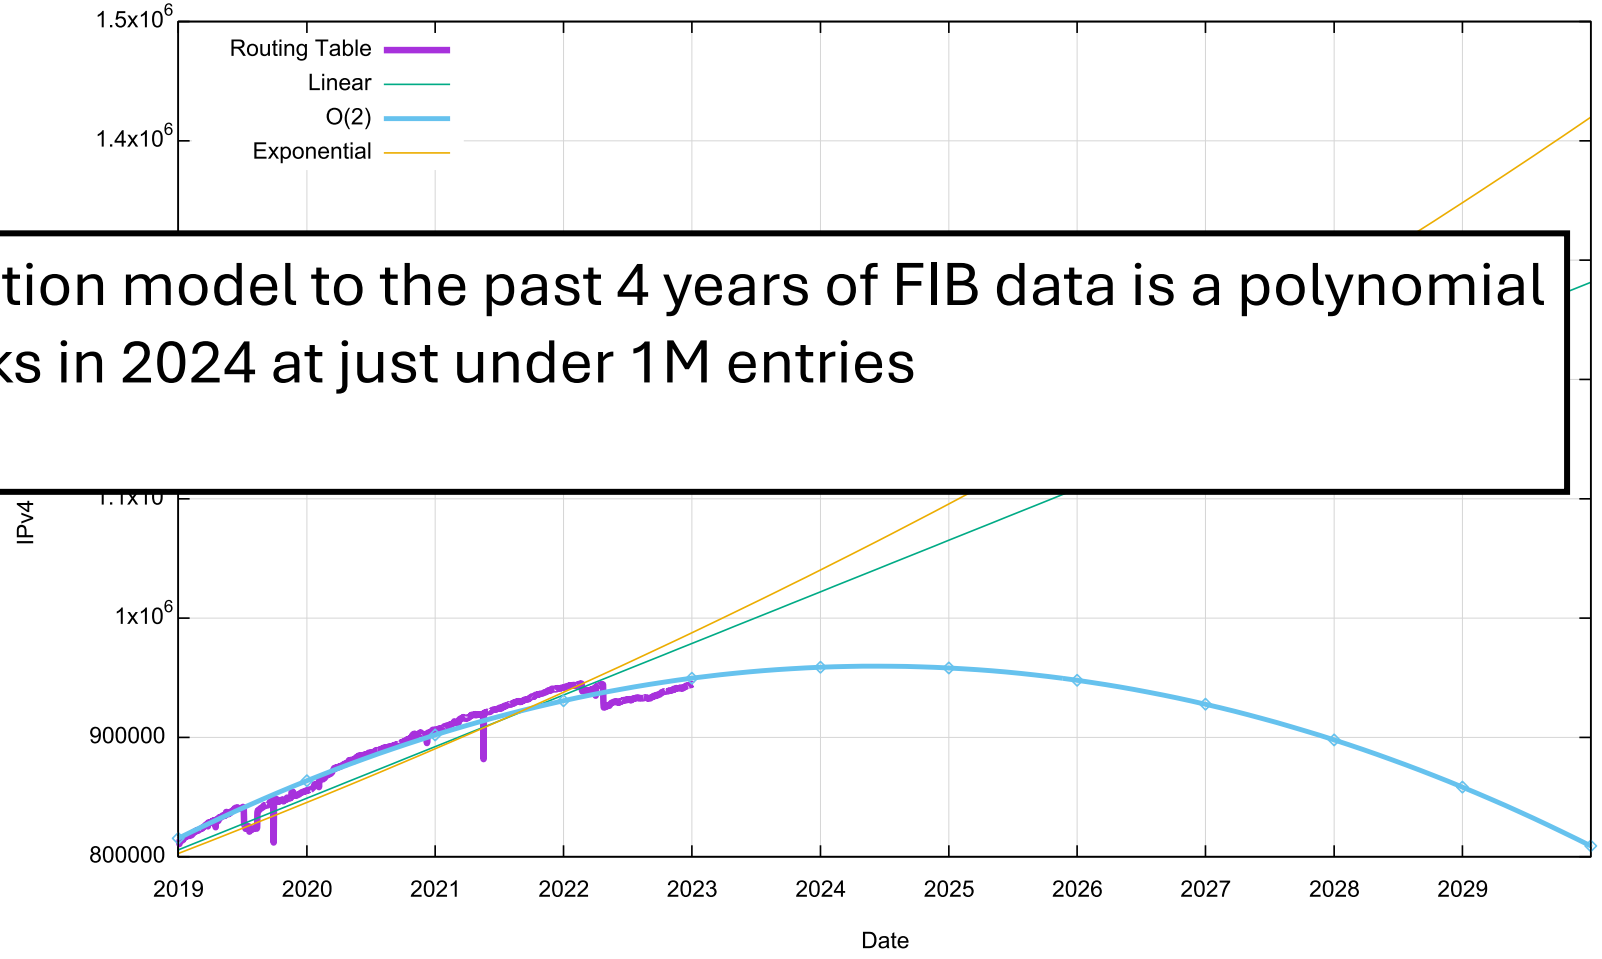

V6 BGP Table Size Predictions

		Linear	Exponential
Jan	2019	62,000	
	2020	79,000	
	2021	106,000	
	2022	147,000	
	2023	172,000	
	2024	**201,000**	198,000	227,000
	2025		226,000	293,000
	2026		254,000	380,000
	2027		282,000	491,000
	2028		310,000	635,000
	2029		337,000	822,000

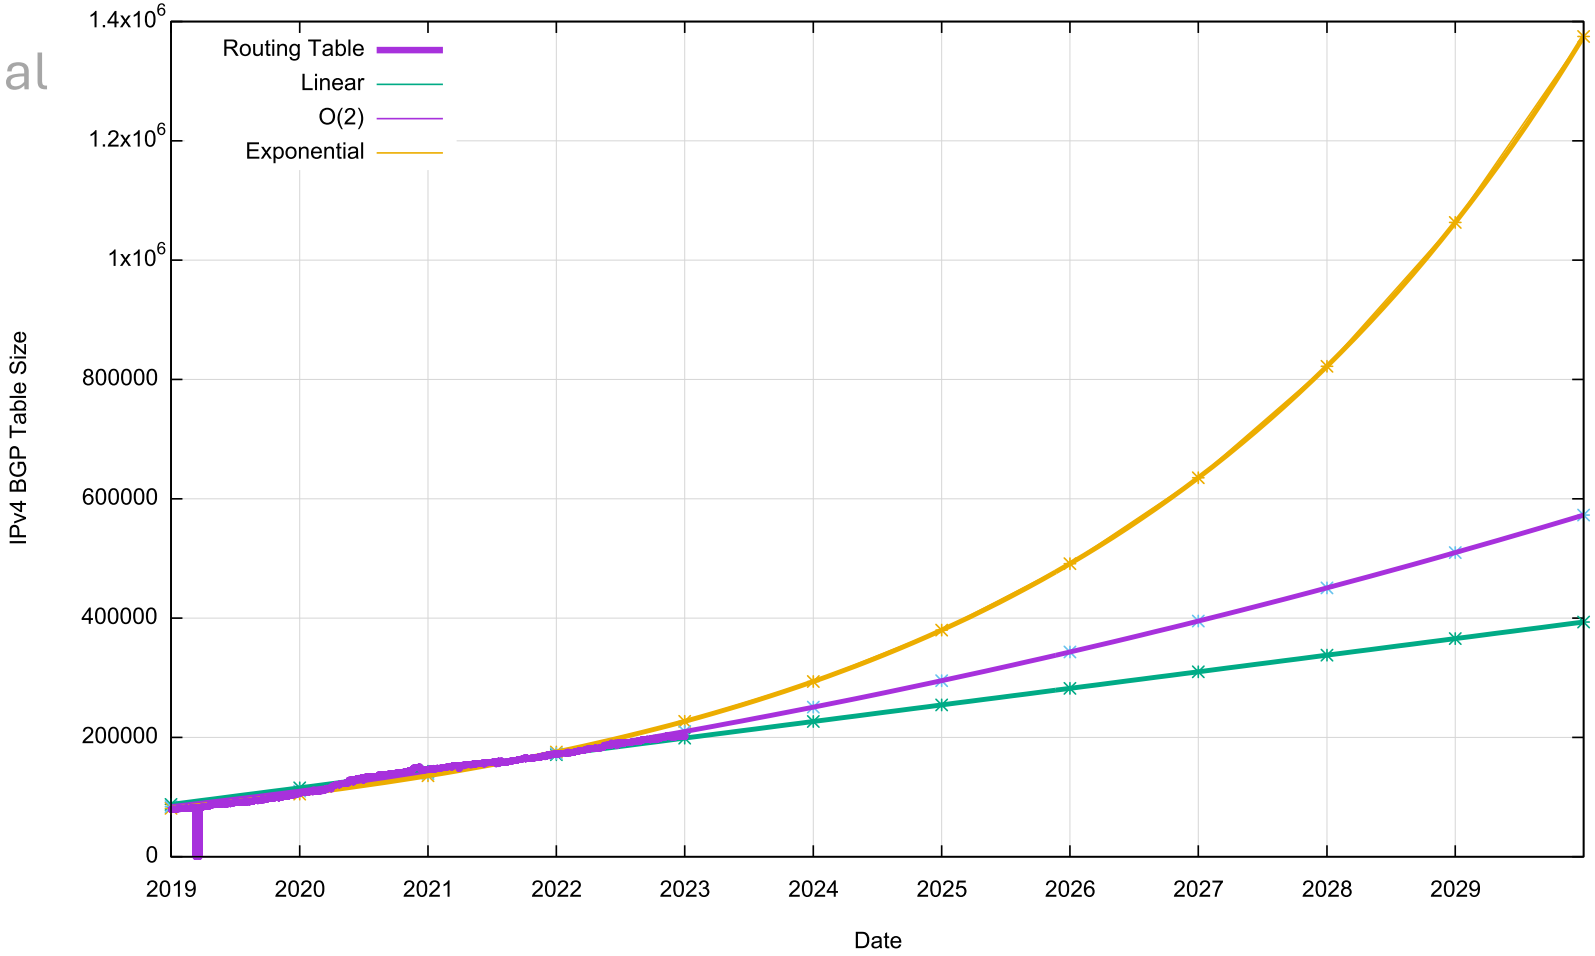

Note that the IPv6 tables are 128bits wide – i.e. 4x the size of the IPv4 tables!

V6 BGP Table Size Predictions

		Linear	Exponential
Jan	2019	62,000	
	2020	79,000	
	2021	106,000	
	2022	1	
	2023	1	
	2024	**20**	
	2025		
	2026		
	2027		
	2028	310,000	635,000
	2029	337,000	822,000

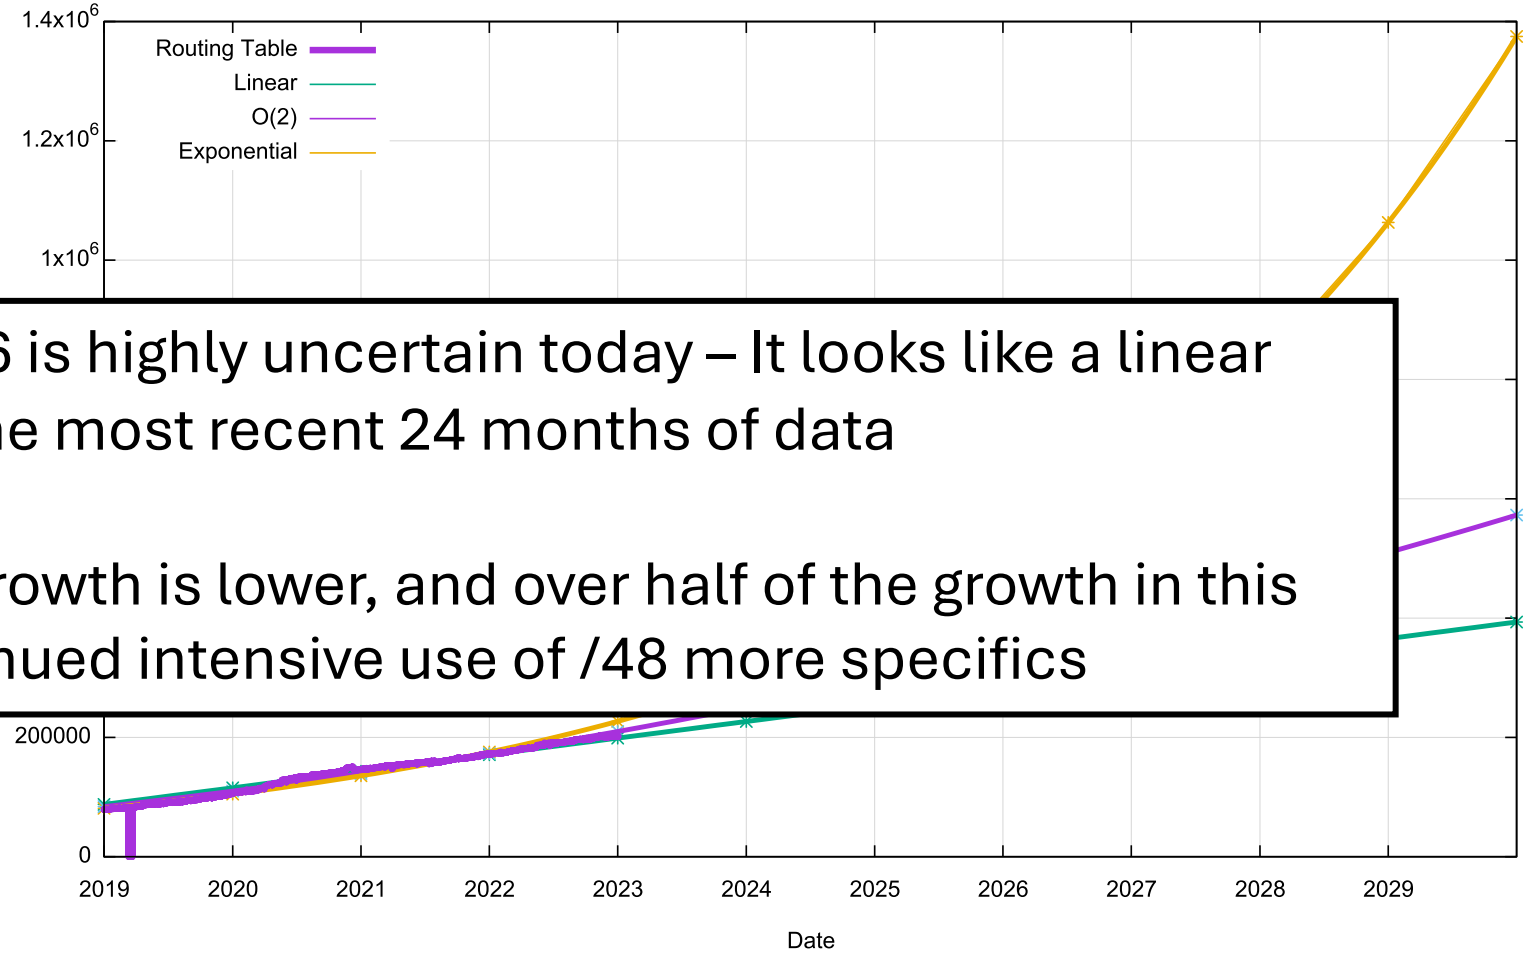

The growth model for IPv6 is highly uncertain today – It looks like a linear growth model matches the most recent 24 months of data

The underlying network growth is lower, and over half of the growth in this model is due to the continued intensive use of /48 more specifics

Note that the IPv6 tables are 128bits wide – i.e. 4x the size of the IPv4 tables!

BGP Table Growth

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

These two tables will require the same storage/lookup size in around 1 year from now, given that each IPv6 entry is 4 times the bit size of an IPv4 entry

The good news ...

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

The Highlights

- IPv4 FIB Summary
- IPv6 FIB Summary
- FIB Projections
- Churn
- Directrions

IPv4 BGP Updates - Daily Updates

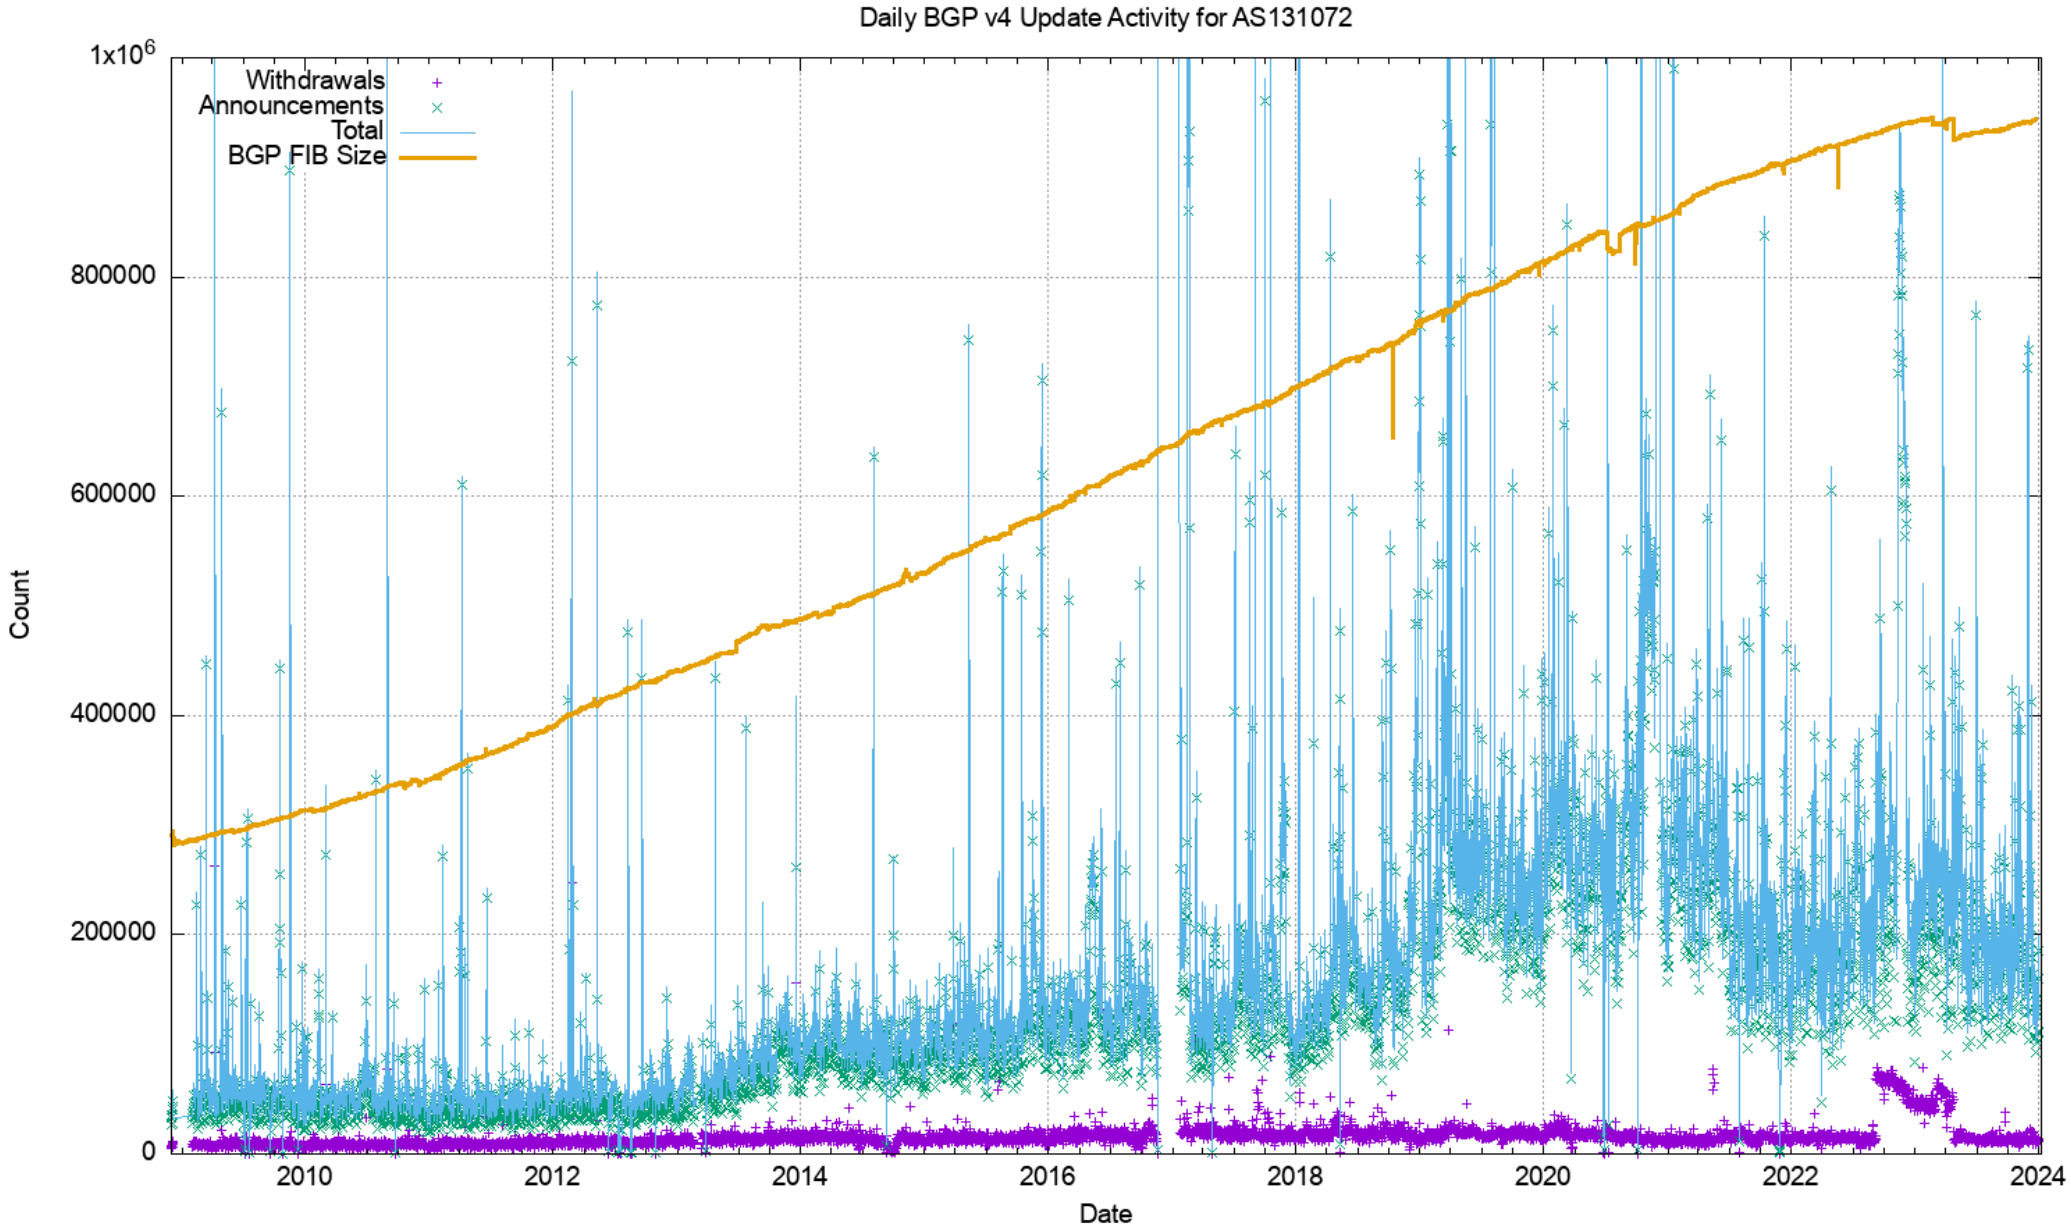

IPv4 BGP Updates - Daily Updates

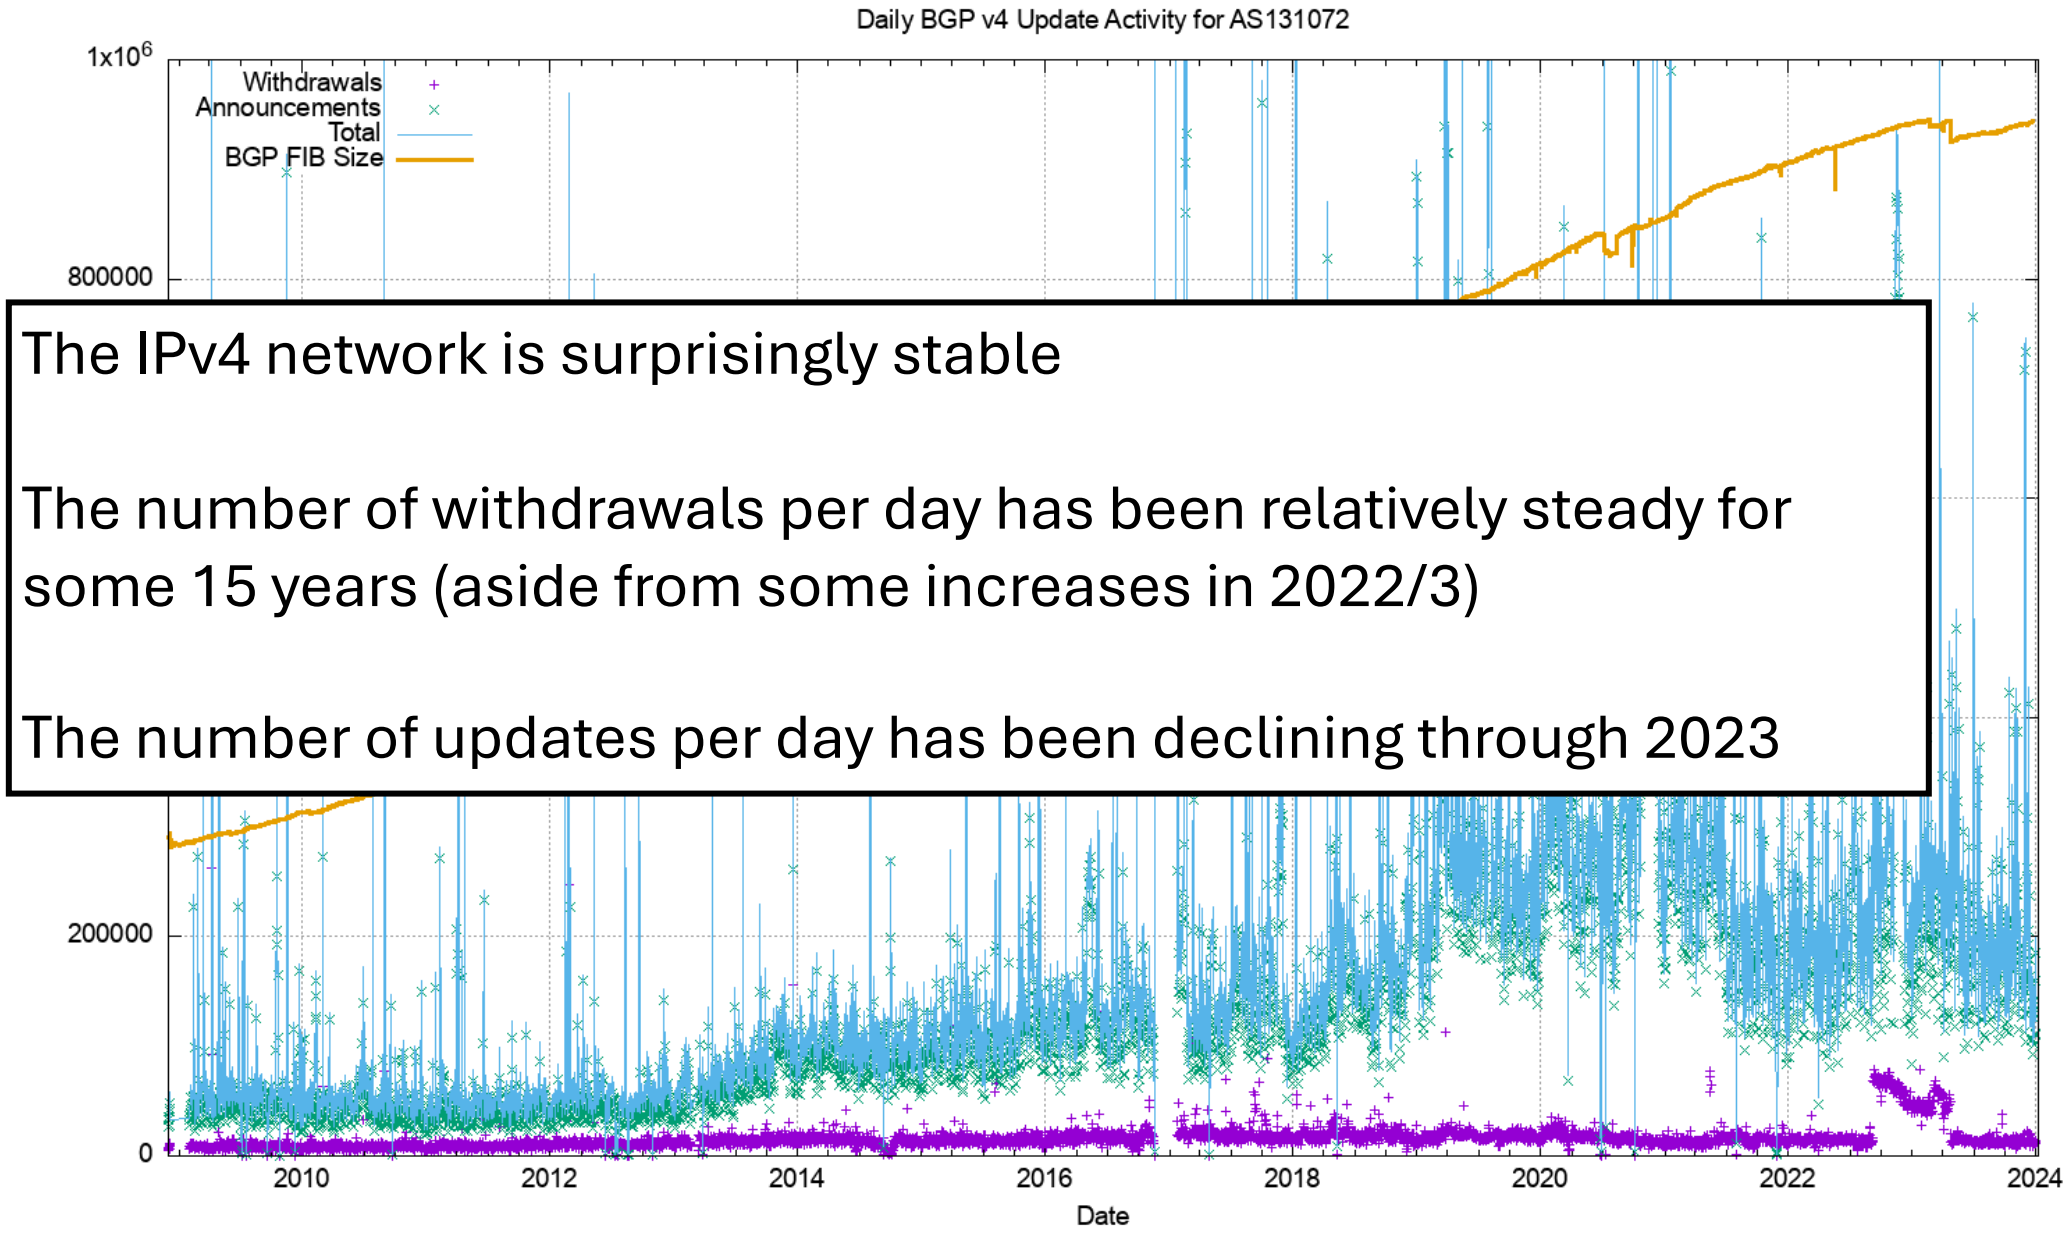

The IPv4 network is surprisingly stable

The number of withdrawals per day has been relatively steady for some 15 years (aside from some increases in 2022/3)

The number of updates per day has been declining through 2023

IPv4 Unstable Prefixes per Day

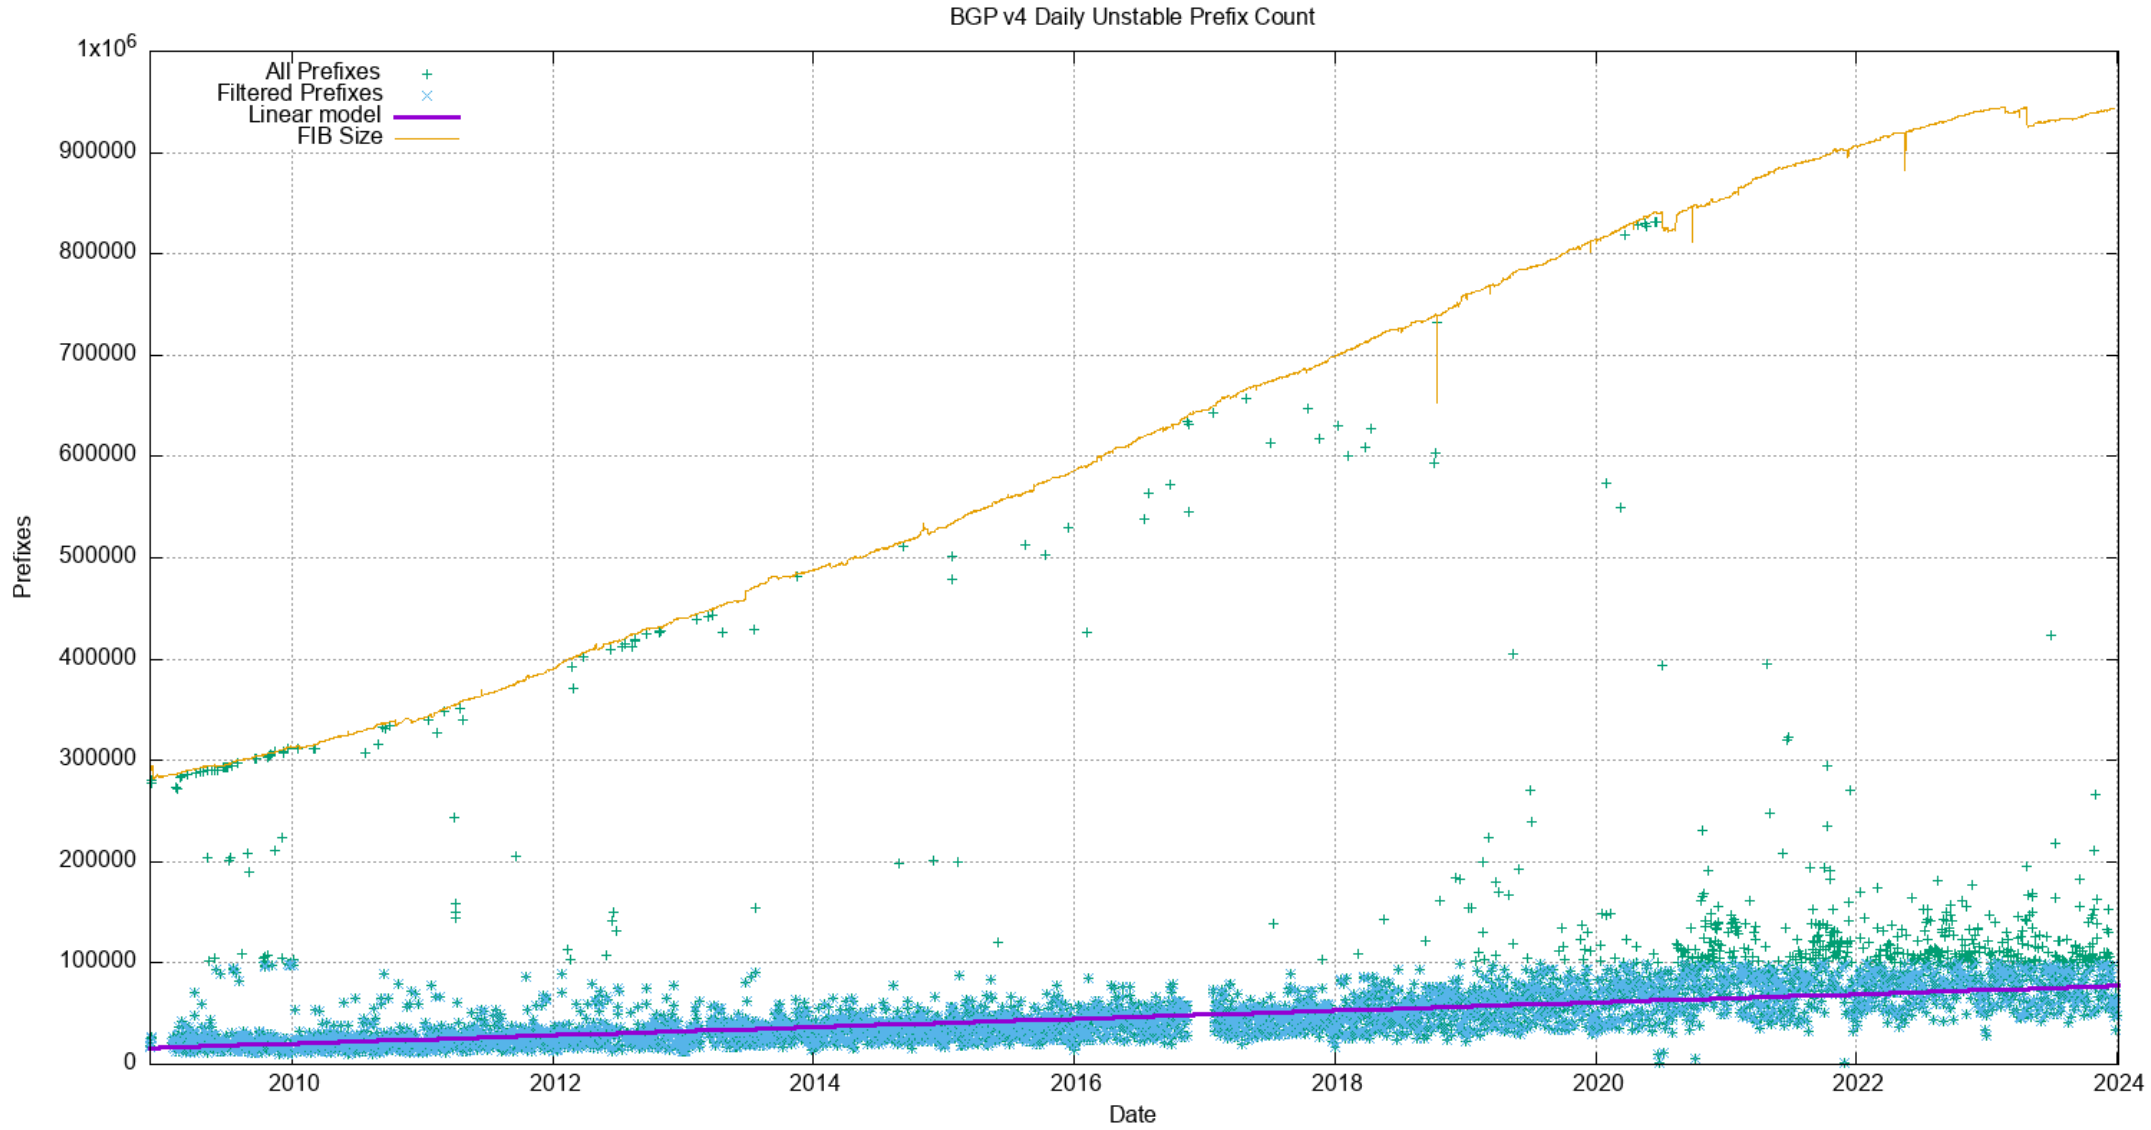

IPv4 Unstable Prefixes per Day

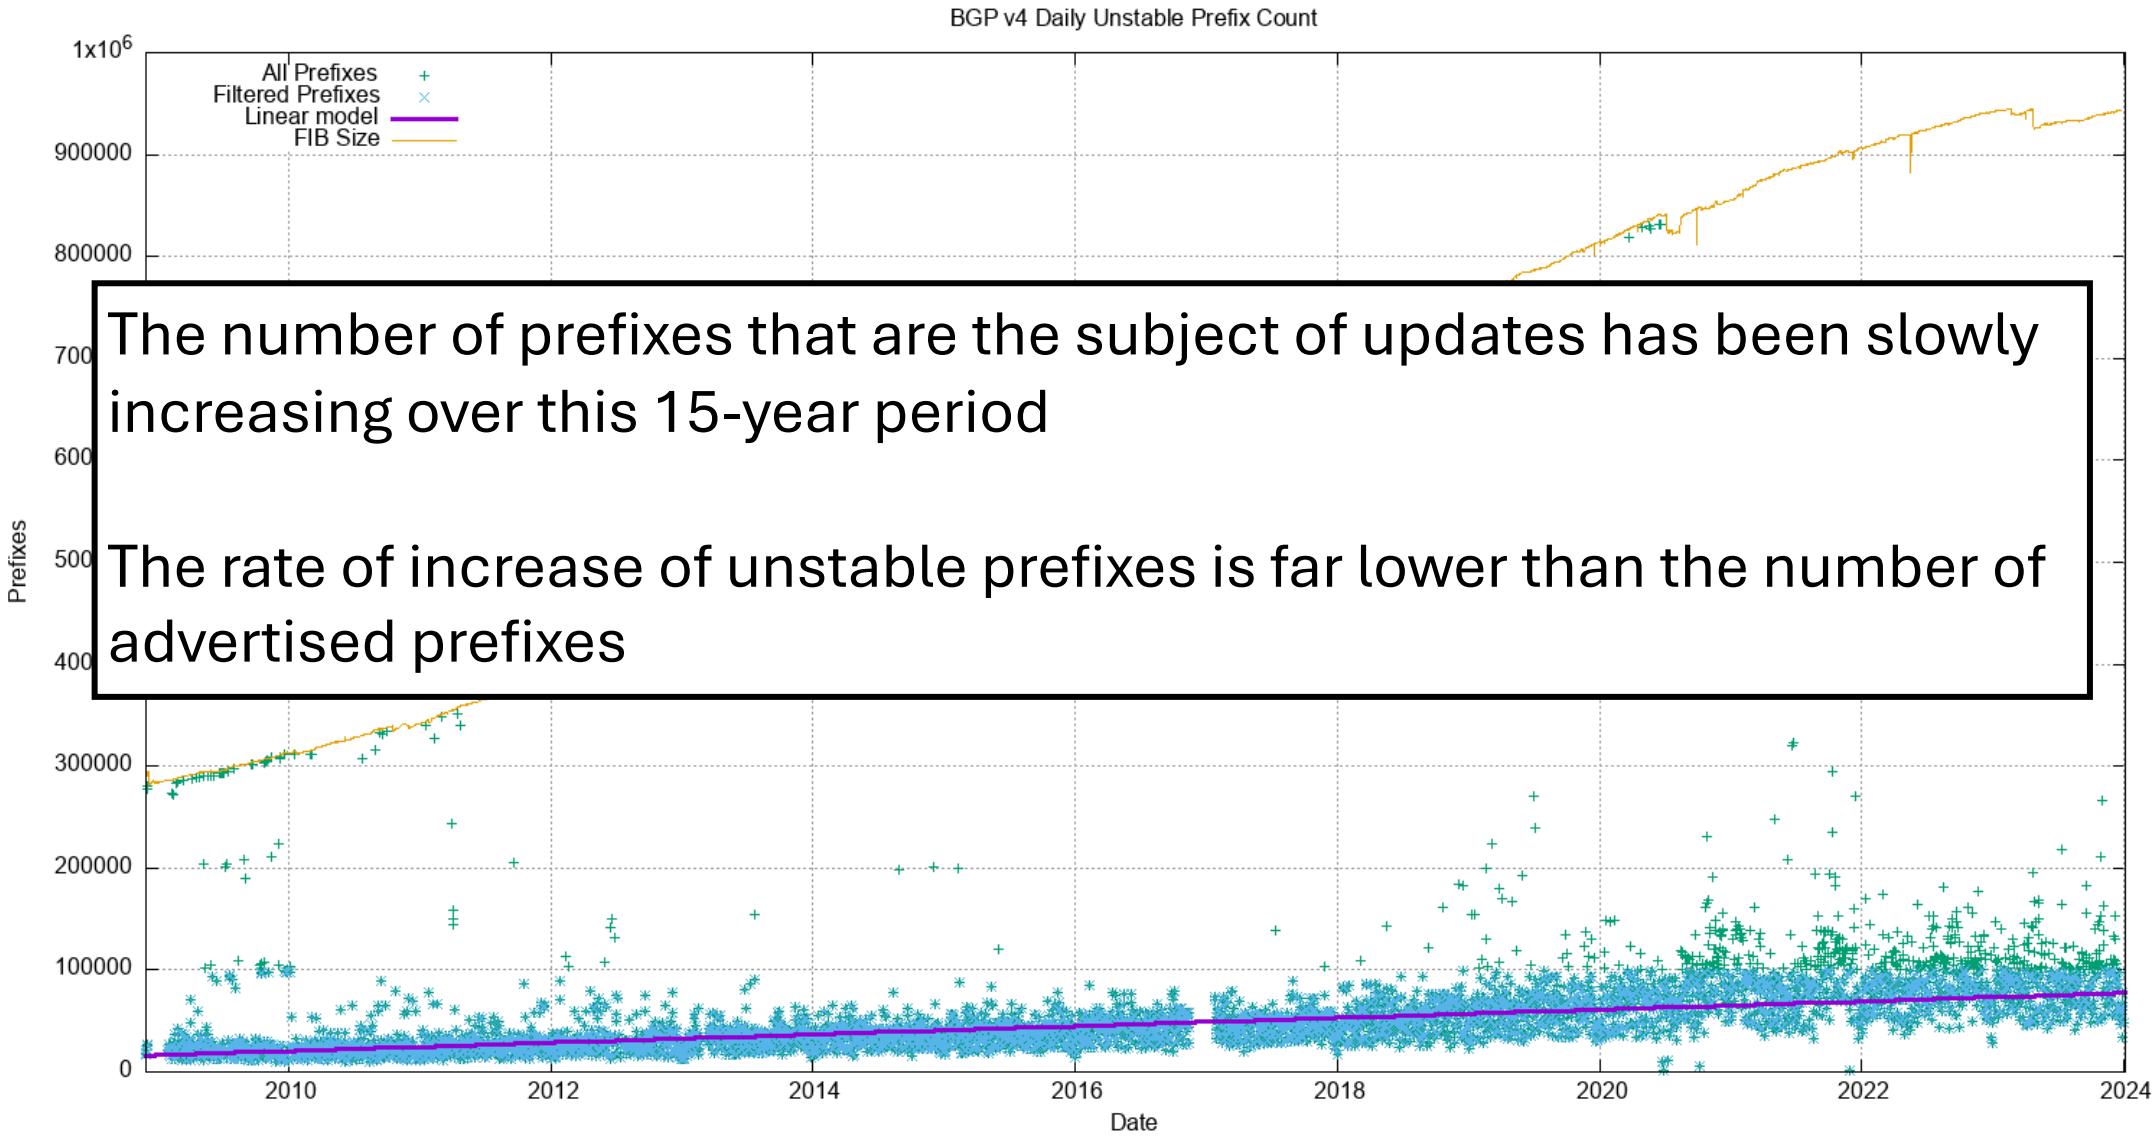

The number of prefixes that are the subject of updates has been slowly increasing over this 15-year period

The rate of increase of unstable prefixes is far lower than the number of advertised prefixes

IPv4 BGP Convergence Performance

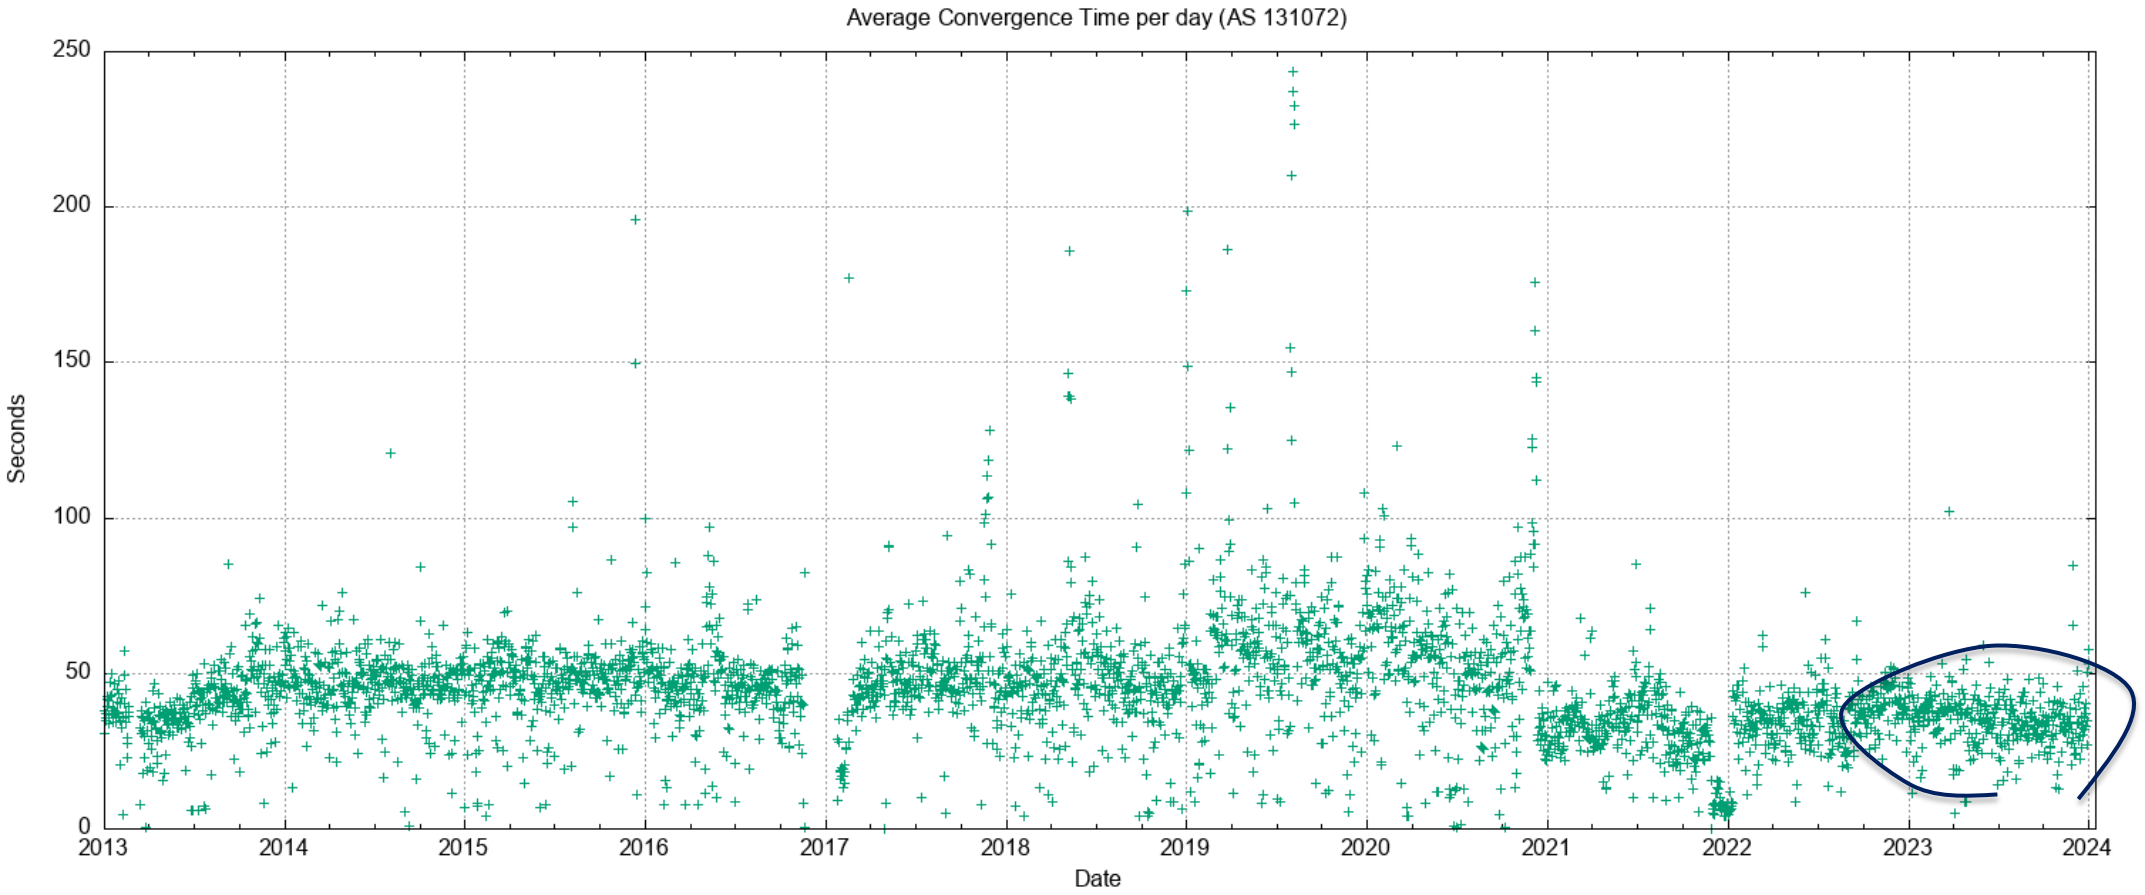

IPv4 BGP Convergence Performance

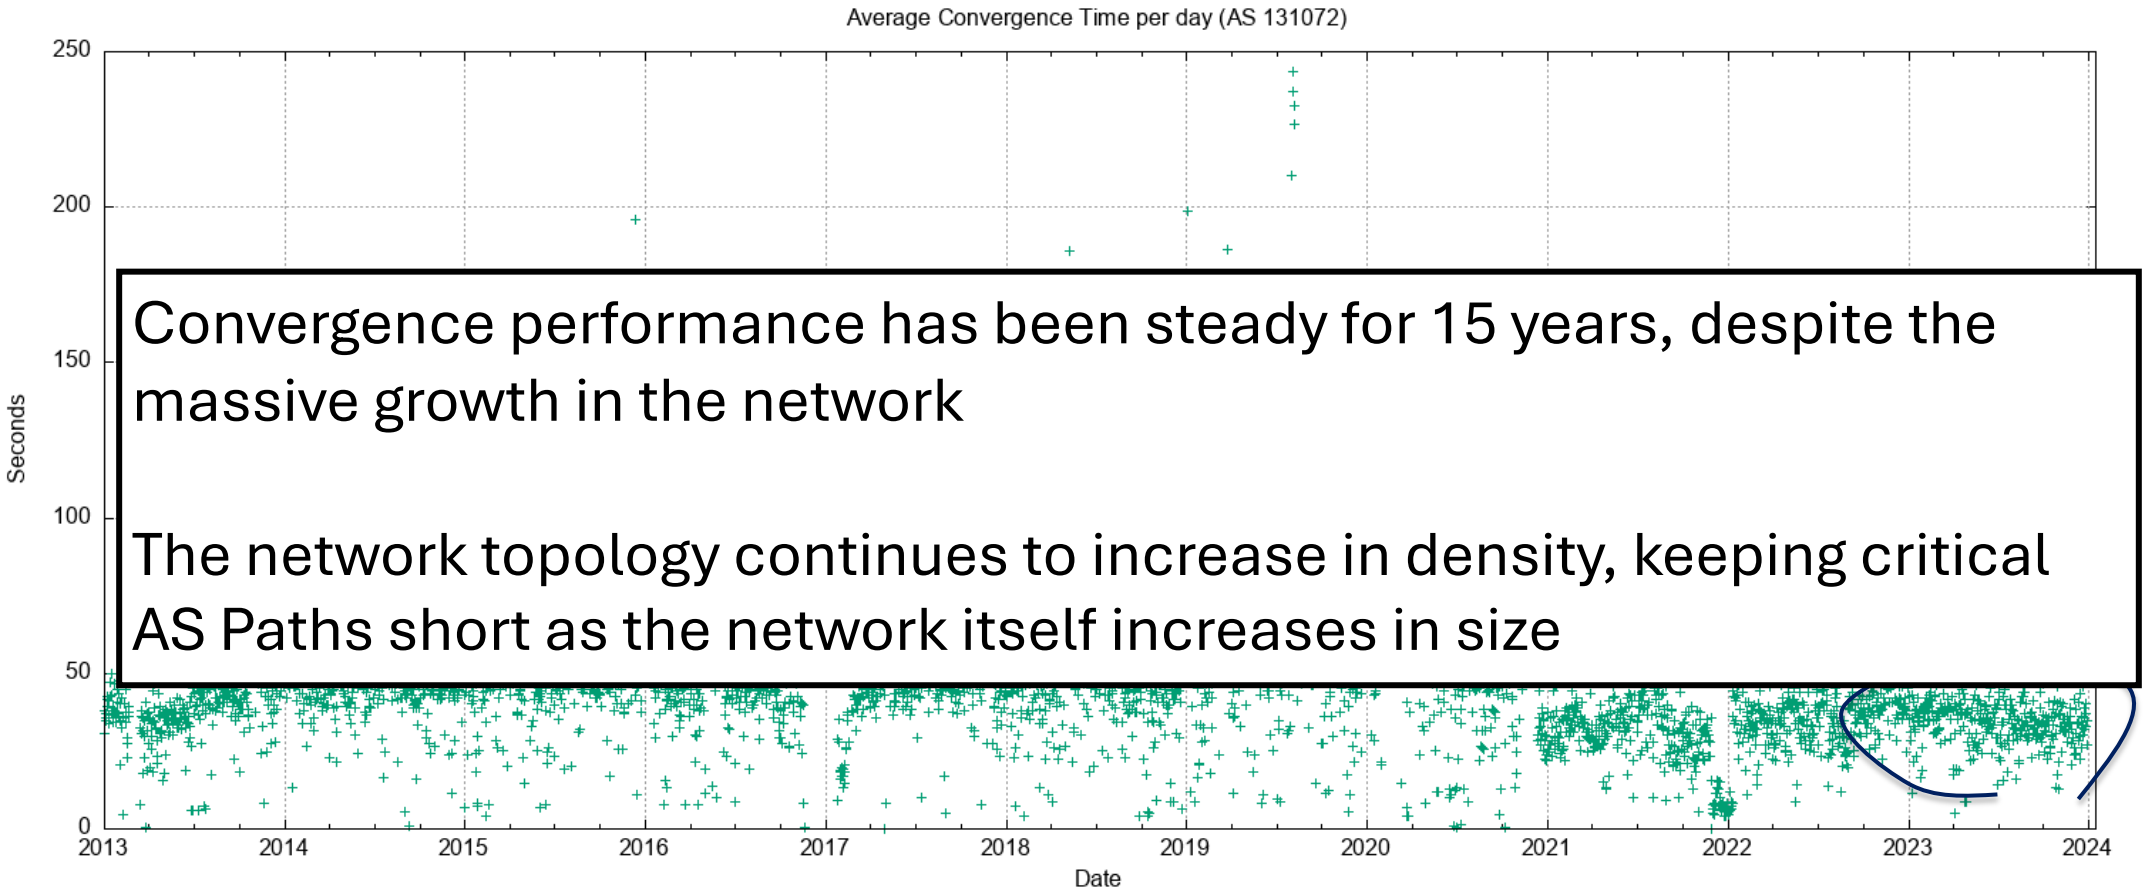

Convergence performance has been steady for 15 years, despite the massive growth in the network

The network topology continues to increase in density, keeping critical AS Paths short as the network itself increases in size

Updates in IPv4 BGP

The IPv4 inter-domain routing system is still highly stable

- The number of updates per instability event and the time to converge to a stable forwarding state has been relatively constant for many years - it rose in 2019 - 2020 and has declined again in 2021, and stabilized in 2022
- 20% of prefixes generate 80% of all updates. Less than 5% of all origin networks are linked to 80% of all updates. **Routing instability is concentrated in a small number of highly unstable cases.**

V6 BGP Updates

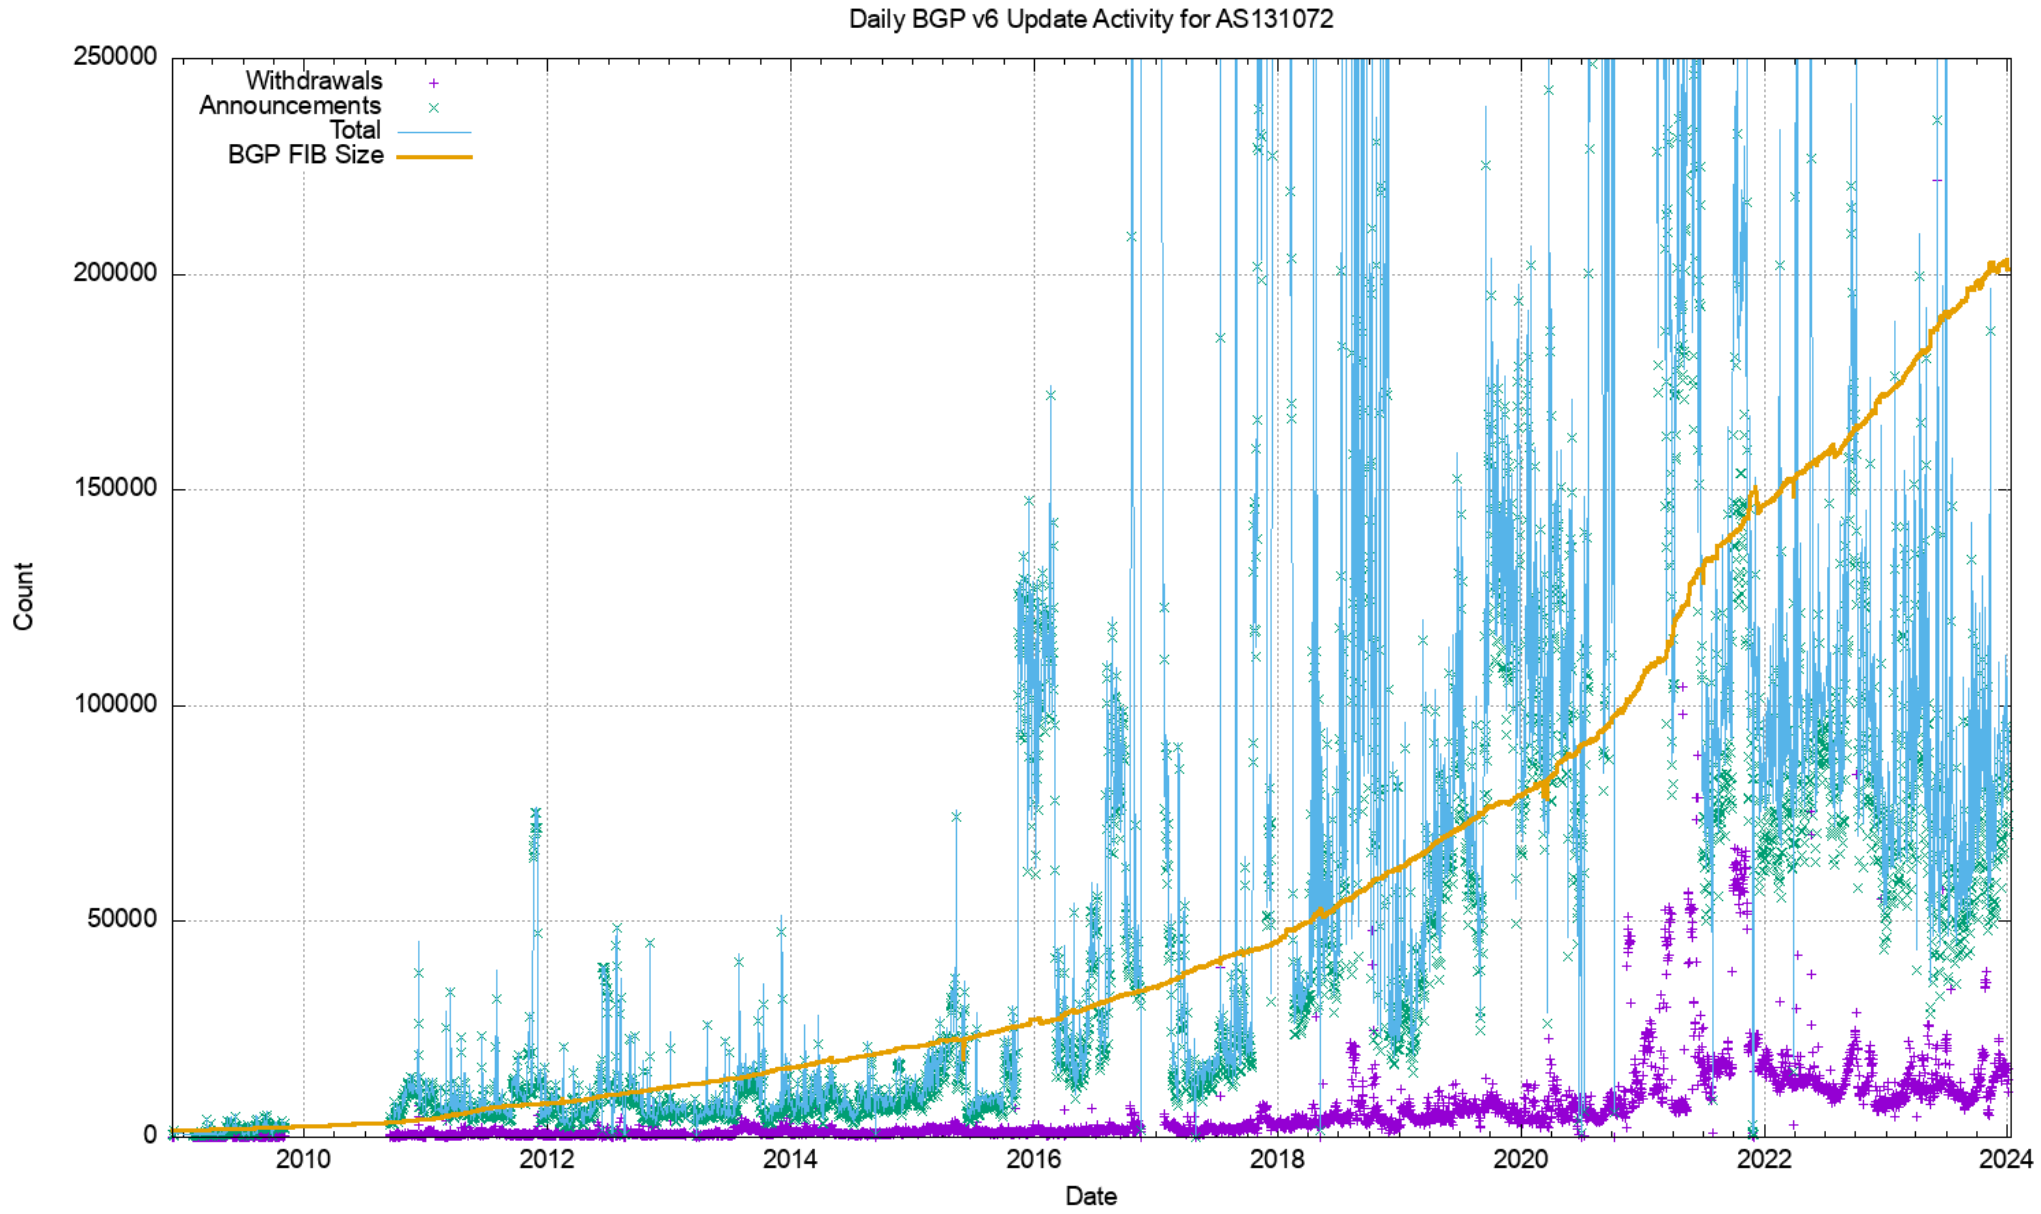

V6 BGP Updates

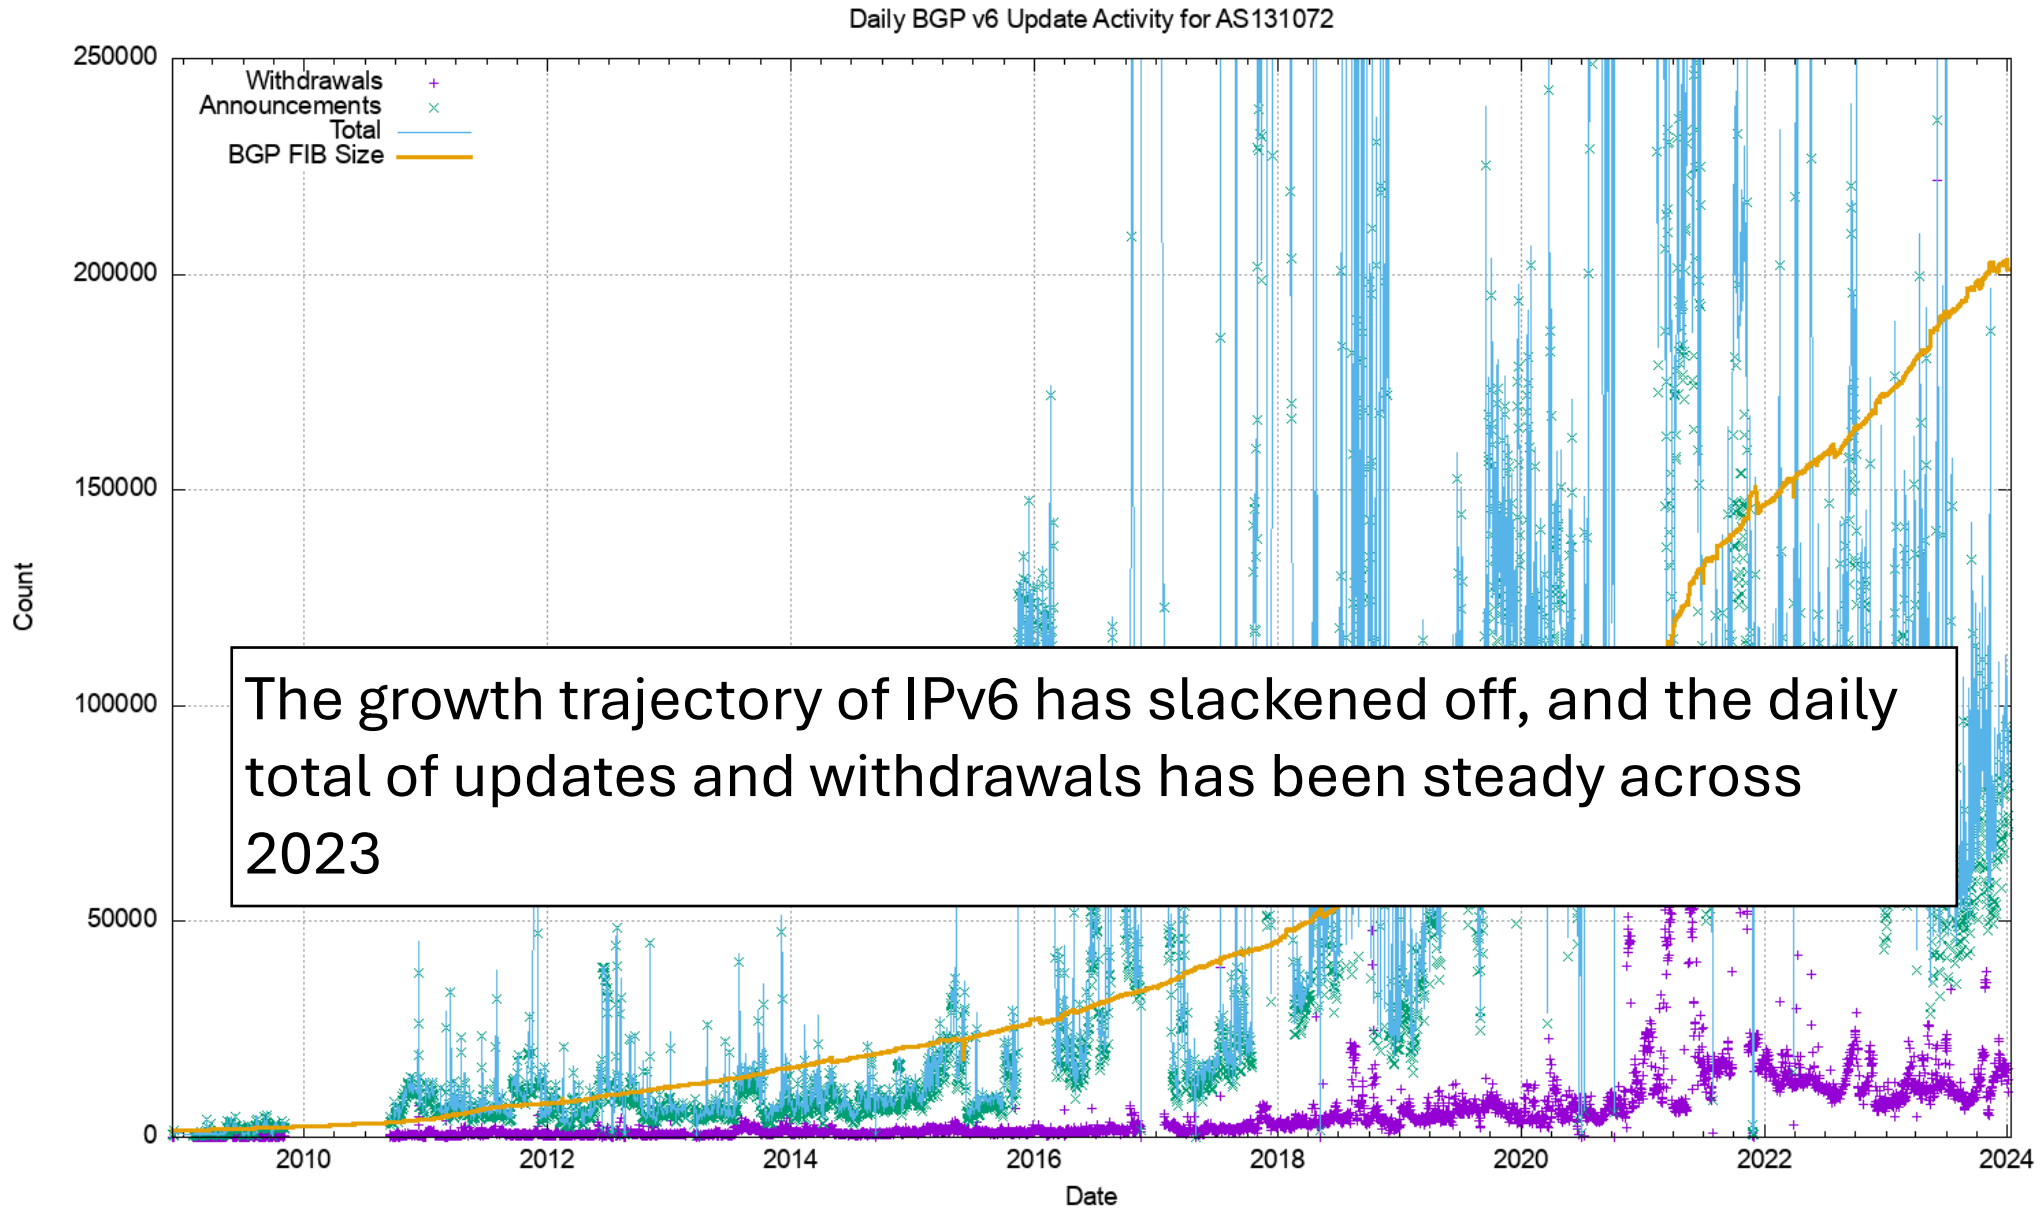

The growth trajectory of IPv6 has slackened off, and the daily total of updates and withdrawals has been steady across 2023

V6 Unstable Prefixes

LOG Scale!

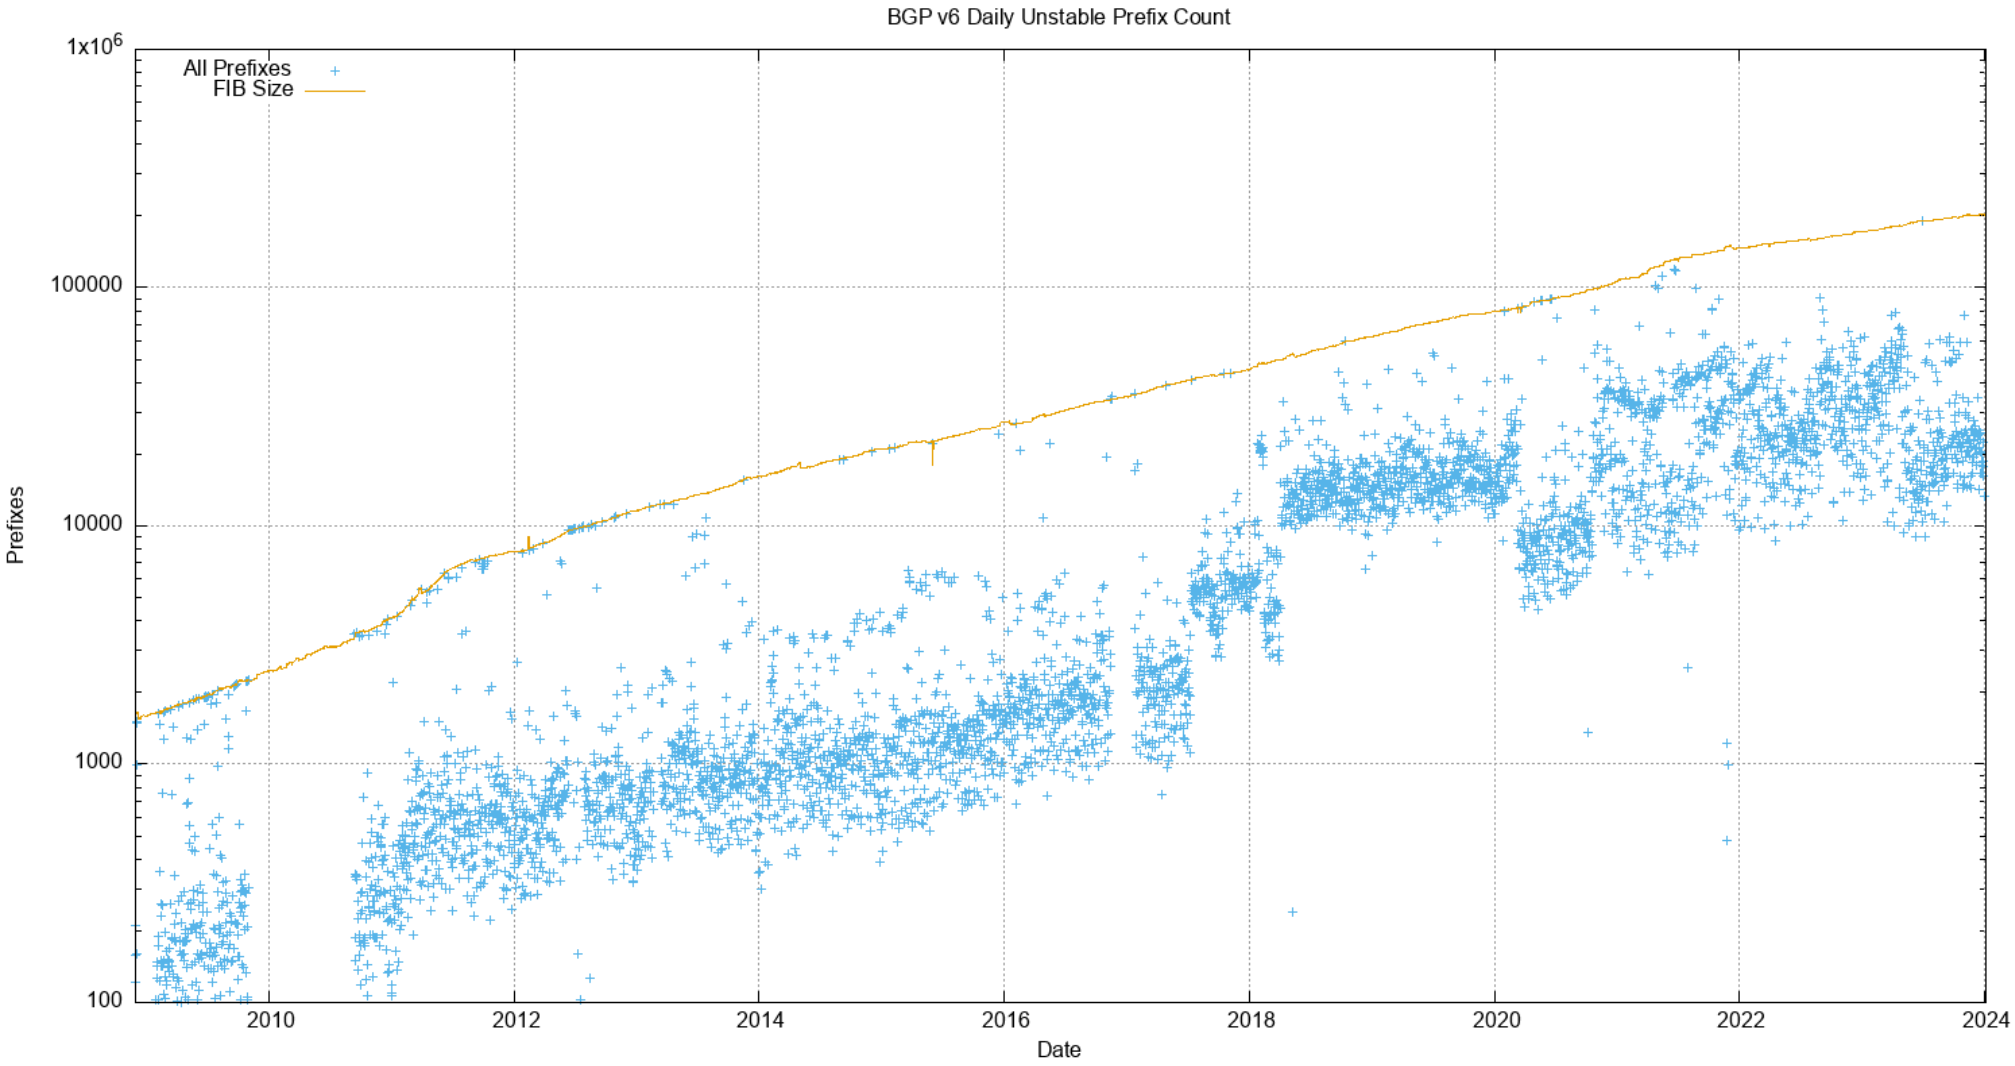

V6 Unstable Prefixes

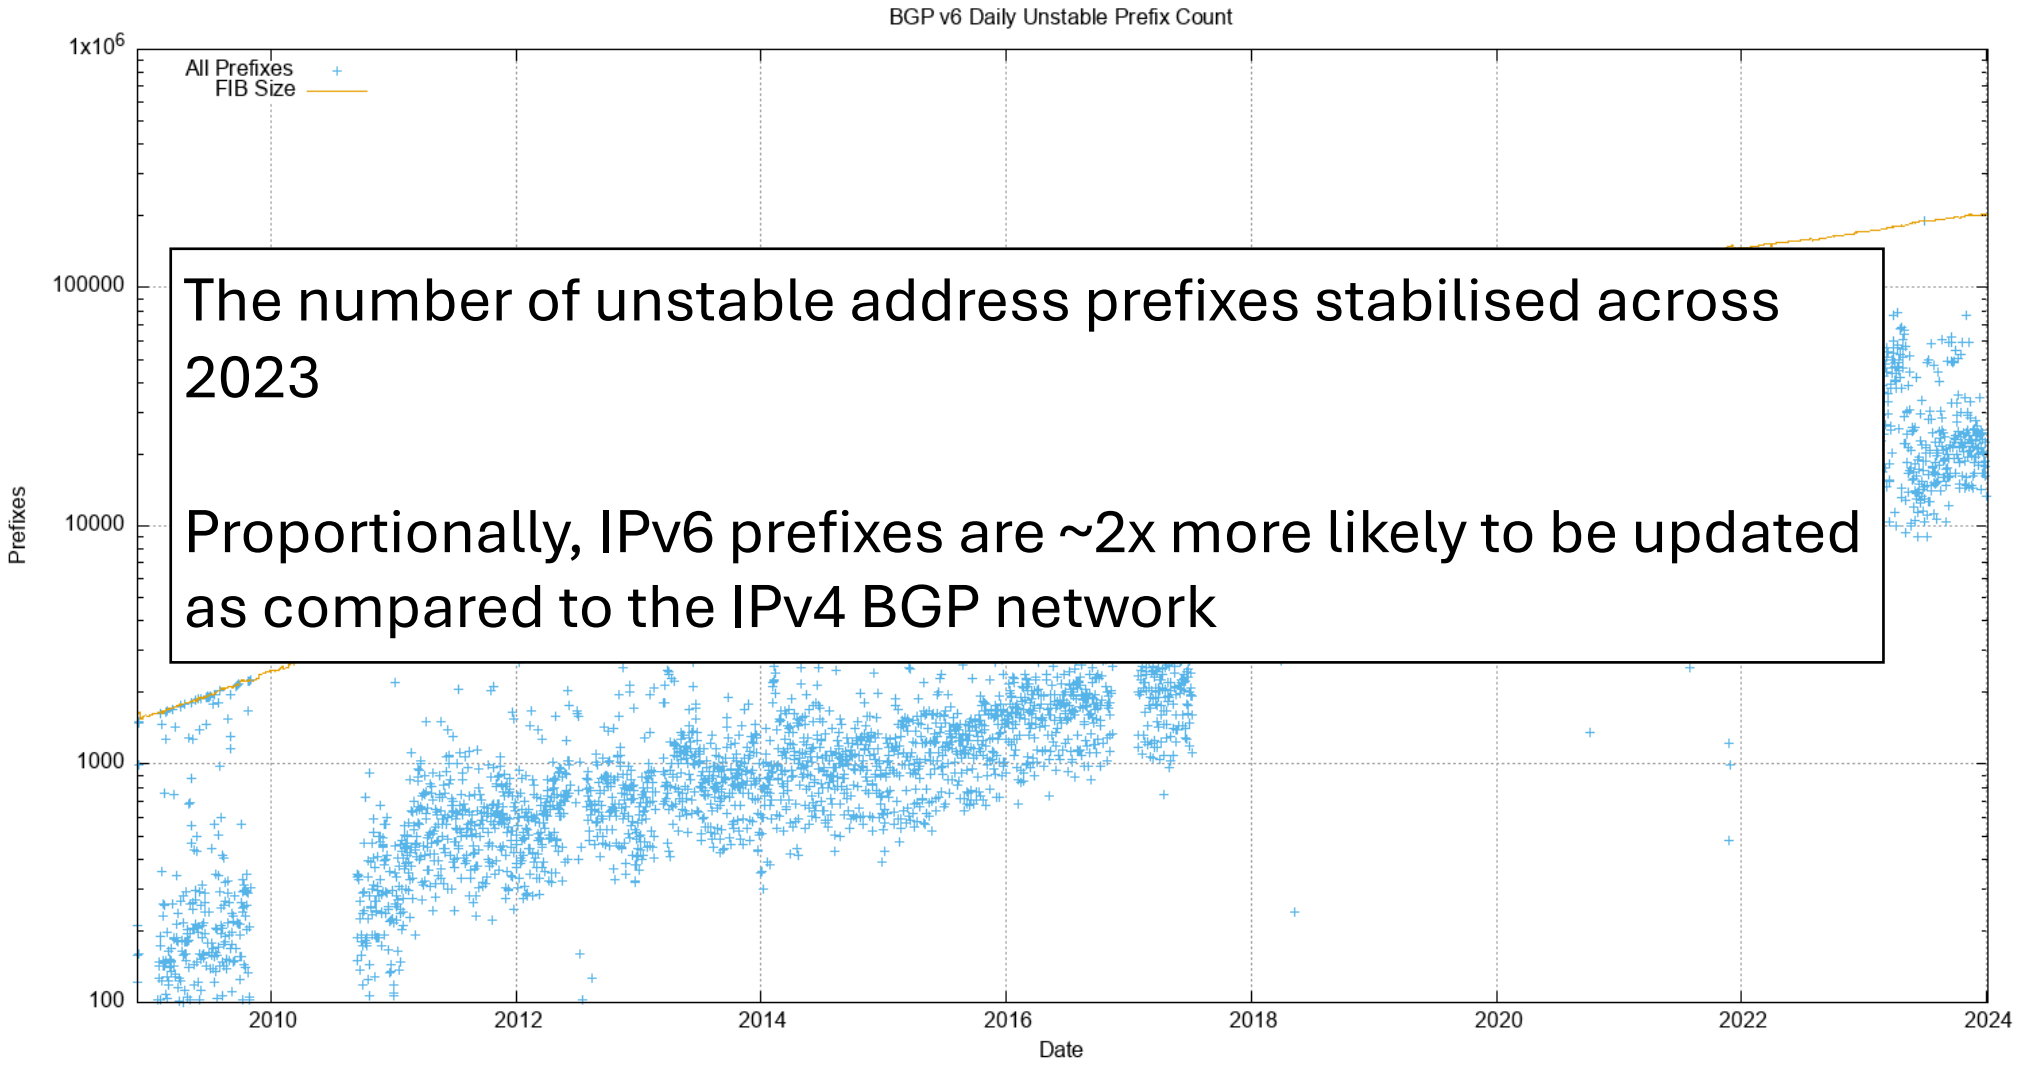

LOG Scale!

The number of unstable address prefixes stabilised across 2023

Proportionally, IPv6 prefixes are ~2x more likely to be updated as compared to the IPv4 BGP network

V6 Convergence Performance

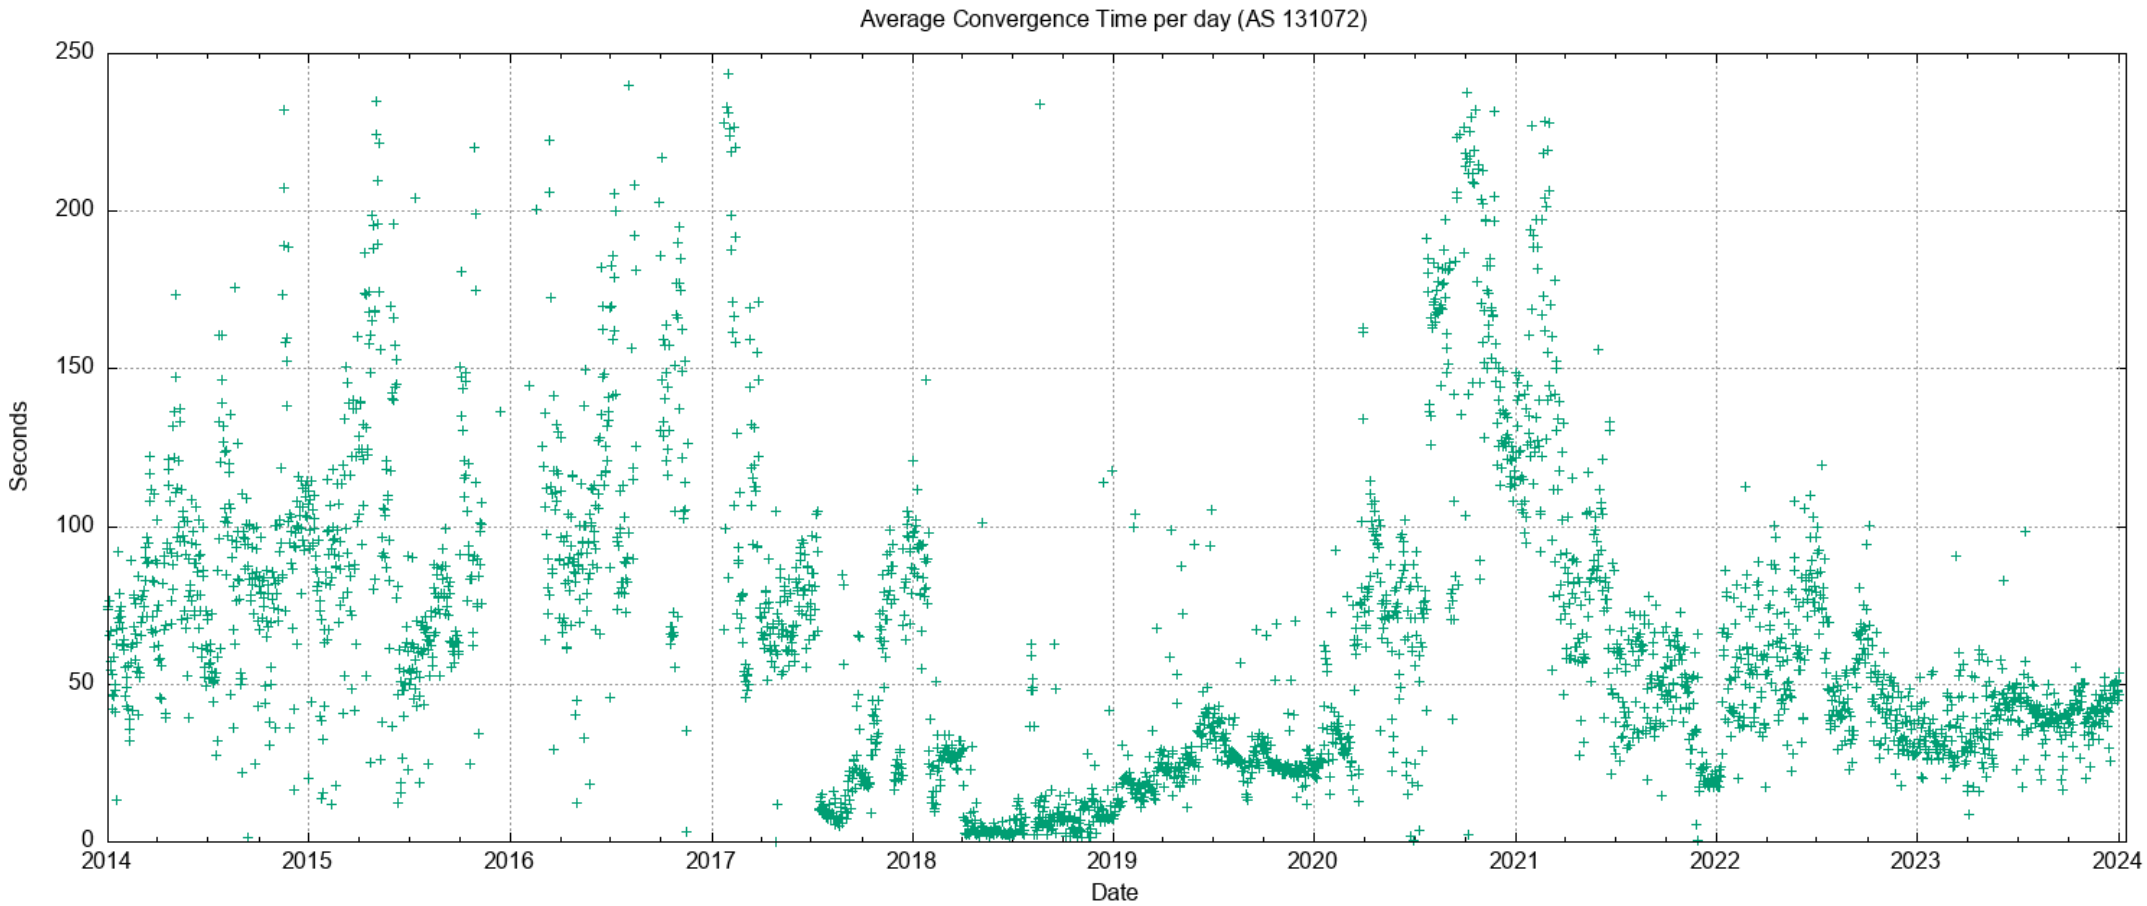

V6 Convergence Performance

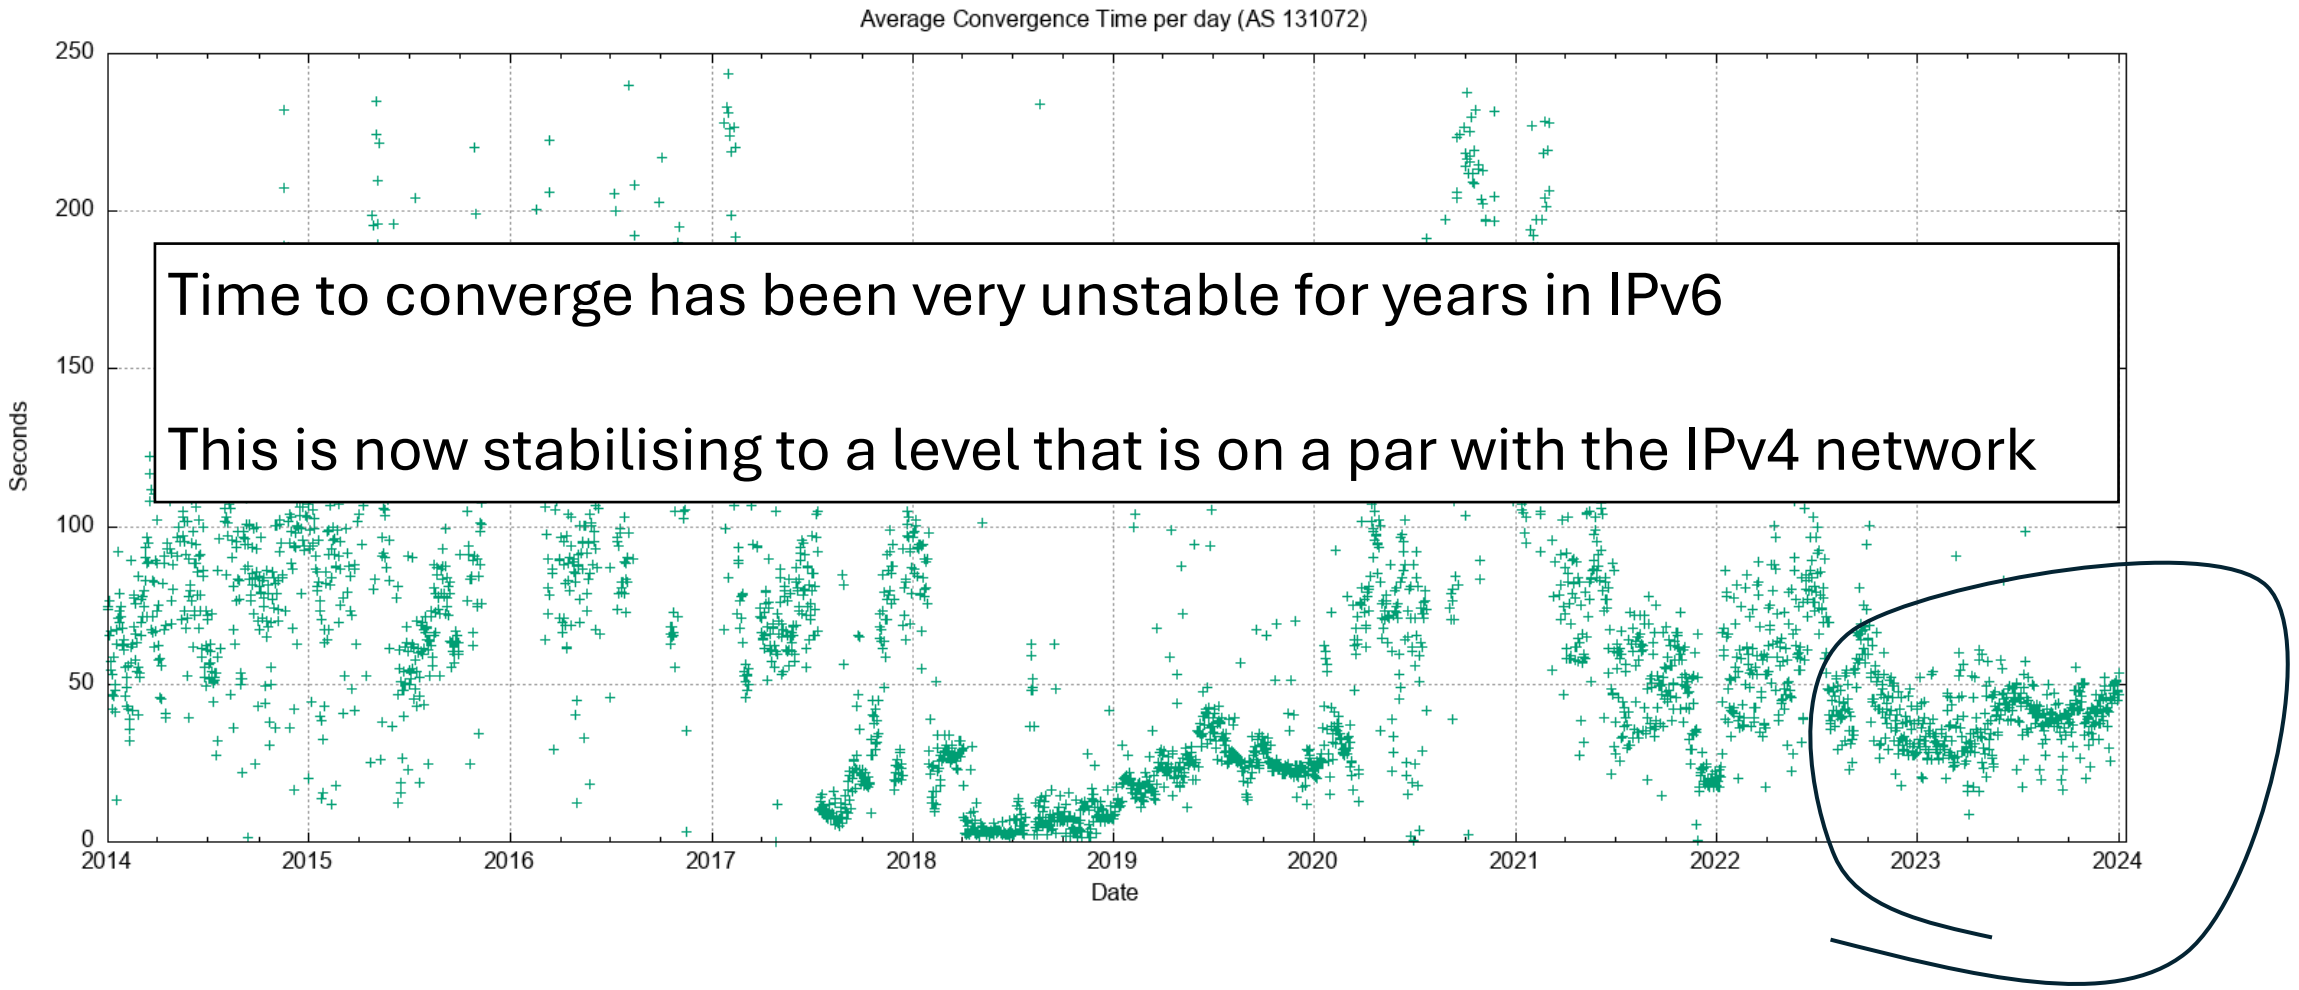

Time to converge has been very unstable for years in IPv6

This is now stabilising to a level that is on a par with the IPv4 network

Updates in IPv6 BGP

It's improving …

- Compared to IPv4, the IPv6 network has exhibited a high level of skew of routing instability, where a small number of networks contribute disproportionately to the overall level of BGP updates in IPv6.
- Just 2 AS's generated 50% of the BGP IPv6 update load in the last 2 weeks of 2023. IPv6 routing instability is still concentrated in a small number of pathologically unstable cases.

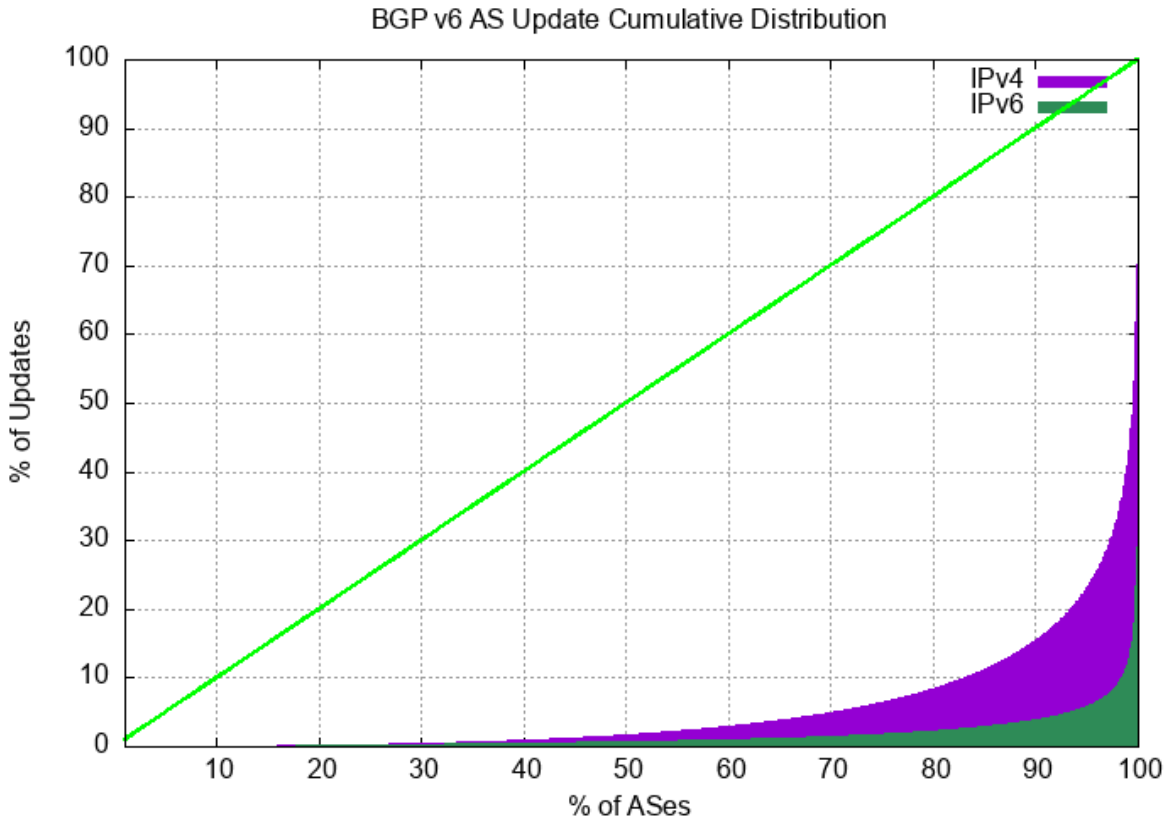

The Highlights

- IPv4 FIB Summary
- IPv6 FIB Summary
- FIB Projections
- Churn
- Directions

Directions: Securing Routing

The plan of study in securing the routing system is to:

- Characterize the current behaviour of the Internet's inter-domain routing system
- Understand the bounds of "normal" behavior
- Characterize the ways in which routing information can be manipulated
- Ways to identify anomalous routing information
- Augment the routing environment with means to perform this detection

Some Open Questions

- Should a useful secured routing system be reactive or preventative?
- In a market-drive economy who pays for security?
- Is the most effective response technical or regulatory?
- Where / how should we respond?
  - Is this a network issue with a network solution?
  - Or should applications protect themselves?
- What is the nature of “trust” in a networked environment?
- How can crypto help?

That's it!

Questions?