

BGP in 2023

Geoff Huston AM
APNIC Labs




The Highlights


- IPv4 Summary
- IPv6 Summary
- FIB Projections
- Churn
- Conclusions

30 Years of Routing the IPv4 Internet

BGP IPv4 RIB Size - RIS and Route Views Peers




2023 in detail




2023 in detail

BGP IPv4 RIB Size - RIS and Route Views Peers



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 Addresses

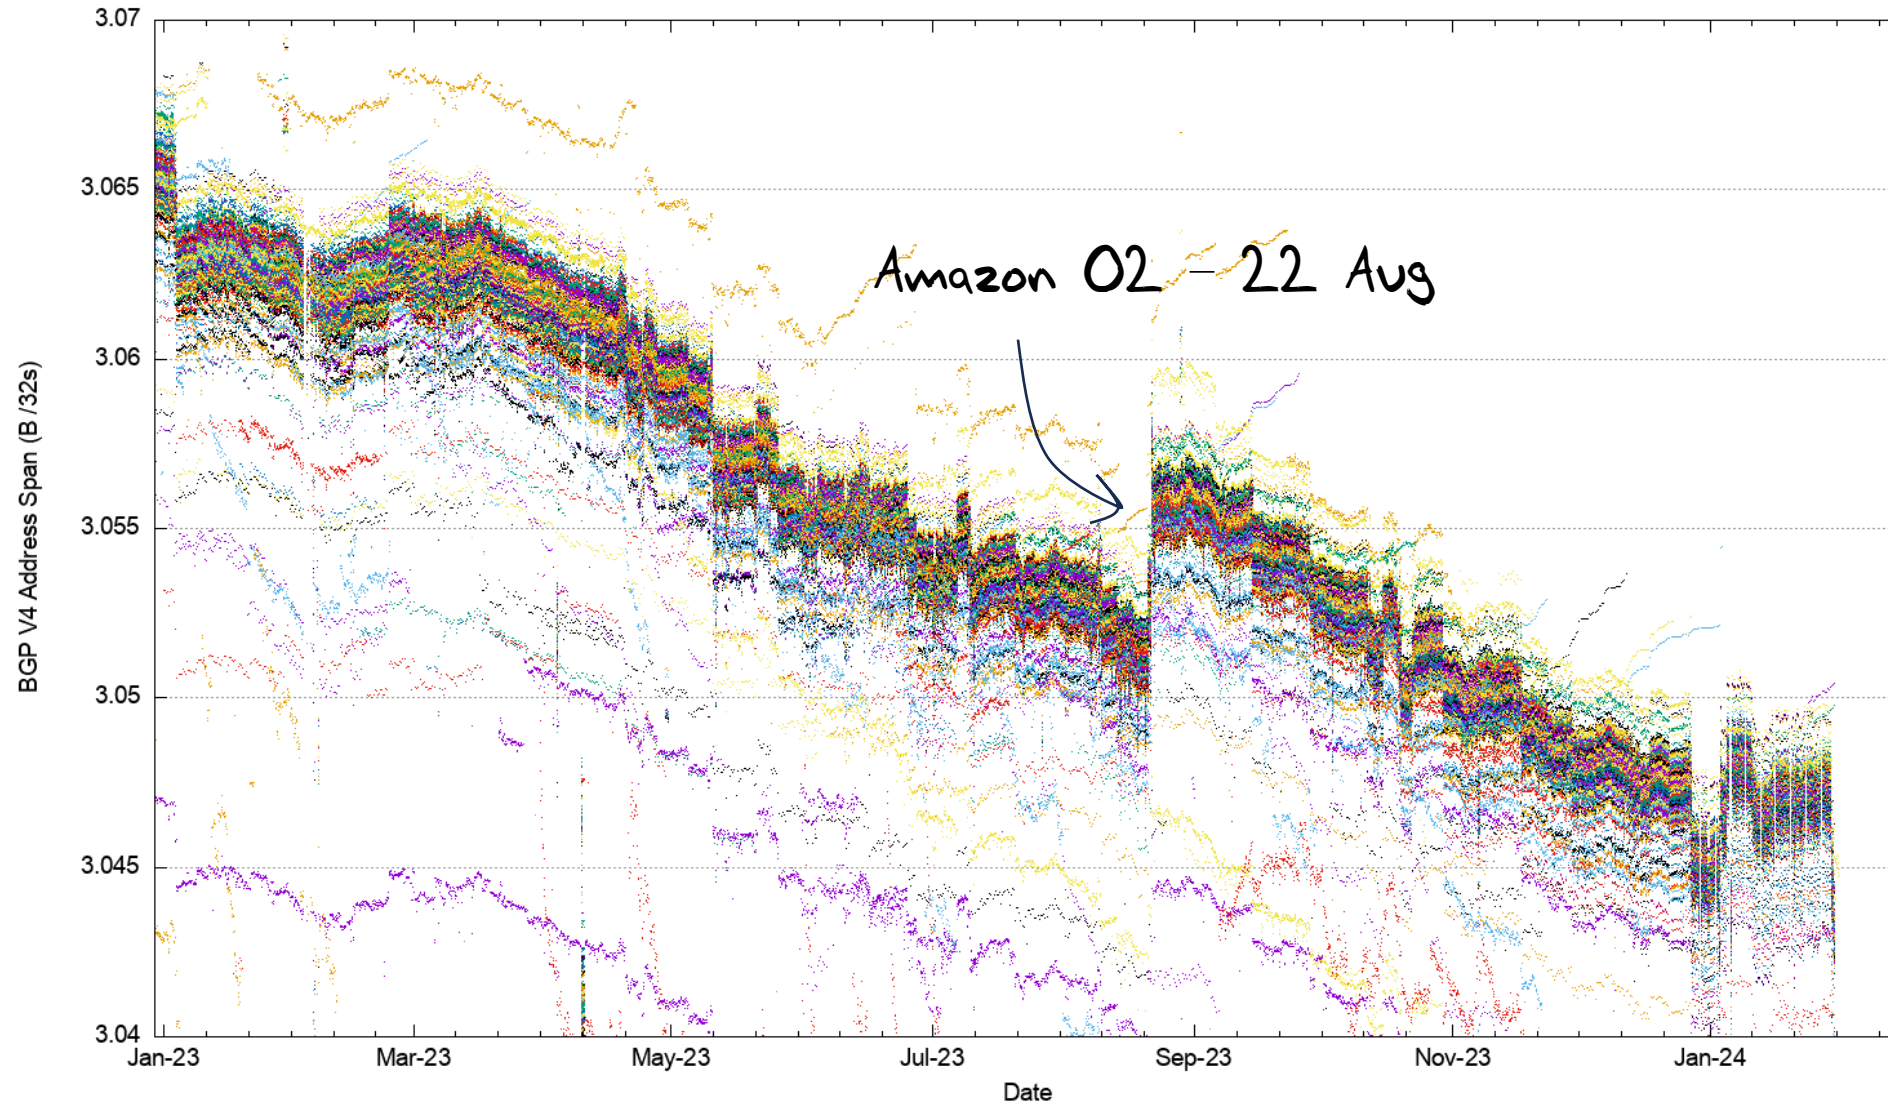
BGP IPv4 Routed Address Span - RIS and Route Views Peers

This is a view pulled together from each of the routing peers of RouteViews and RIPE RIS

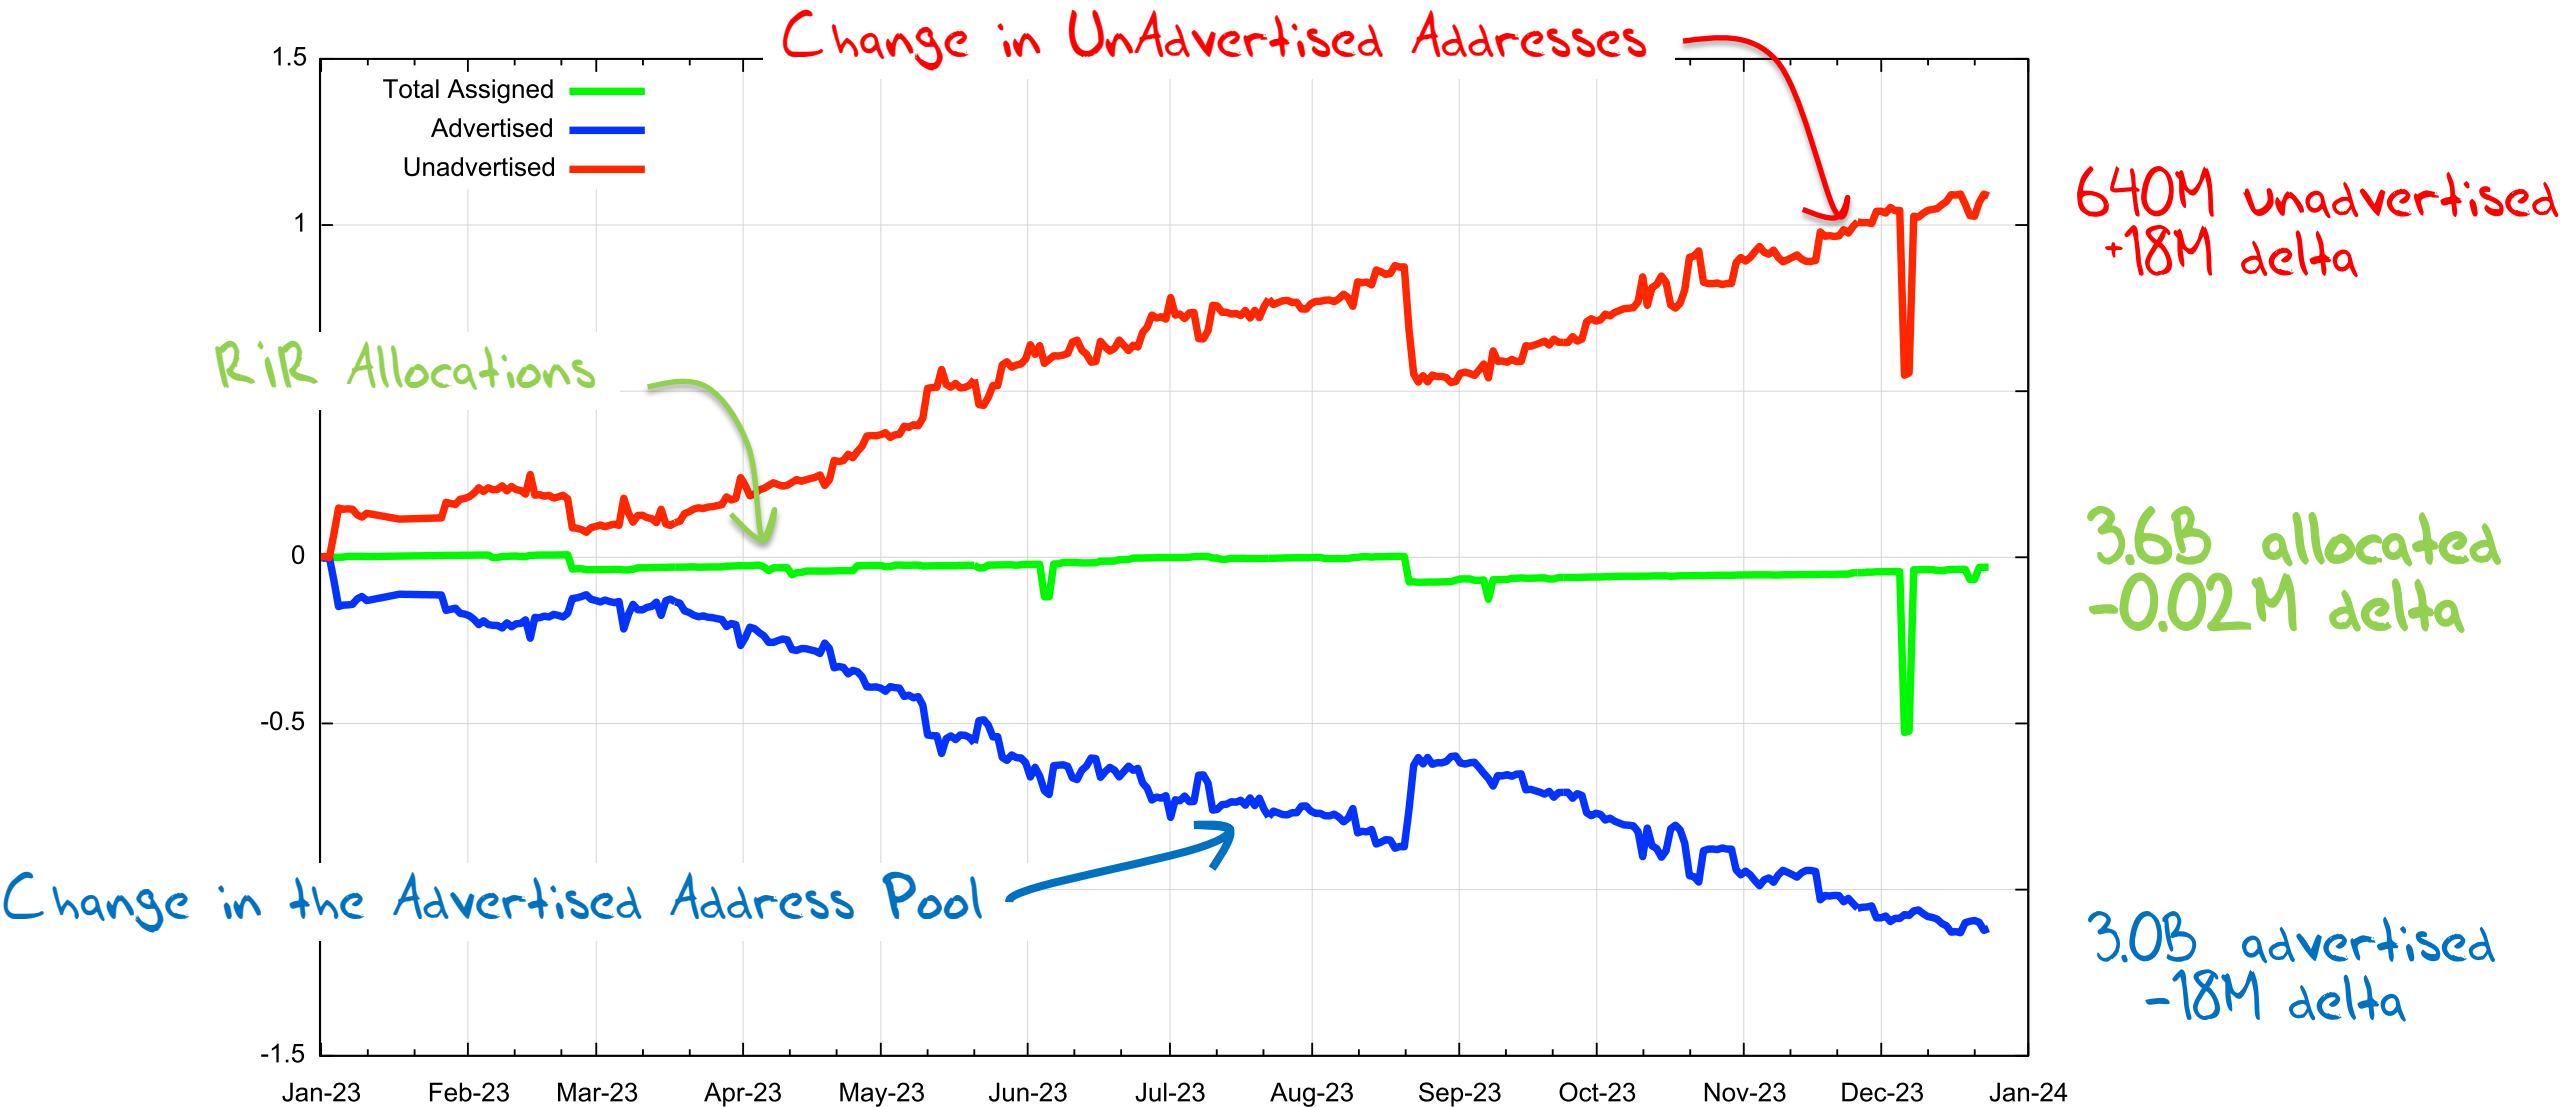


2023 in Detail

BGP IPv4 Routed Address Span - RIS and Route Views Peers



2023: Assigned vs Recovered



2023: Assigned vs Recovered



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A total of 18M addresses were withdrawn from the advertised network, and shifted to the unadvertised pool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

unadvertised delta

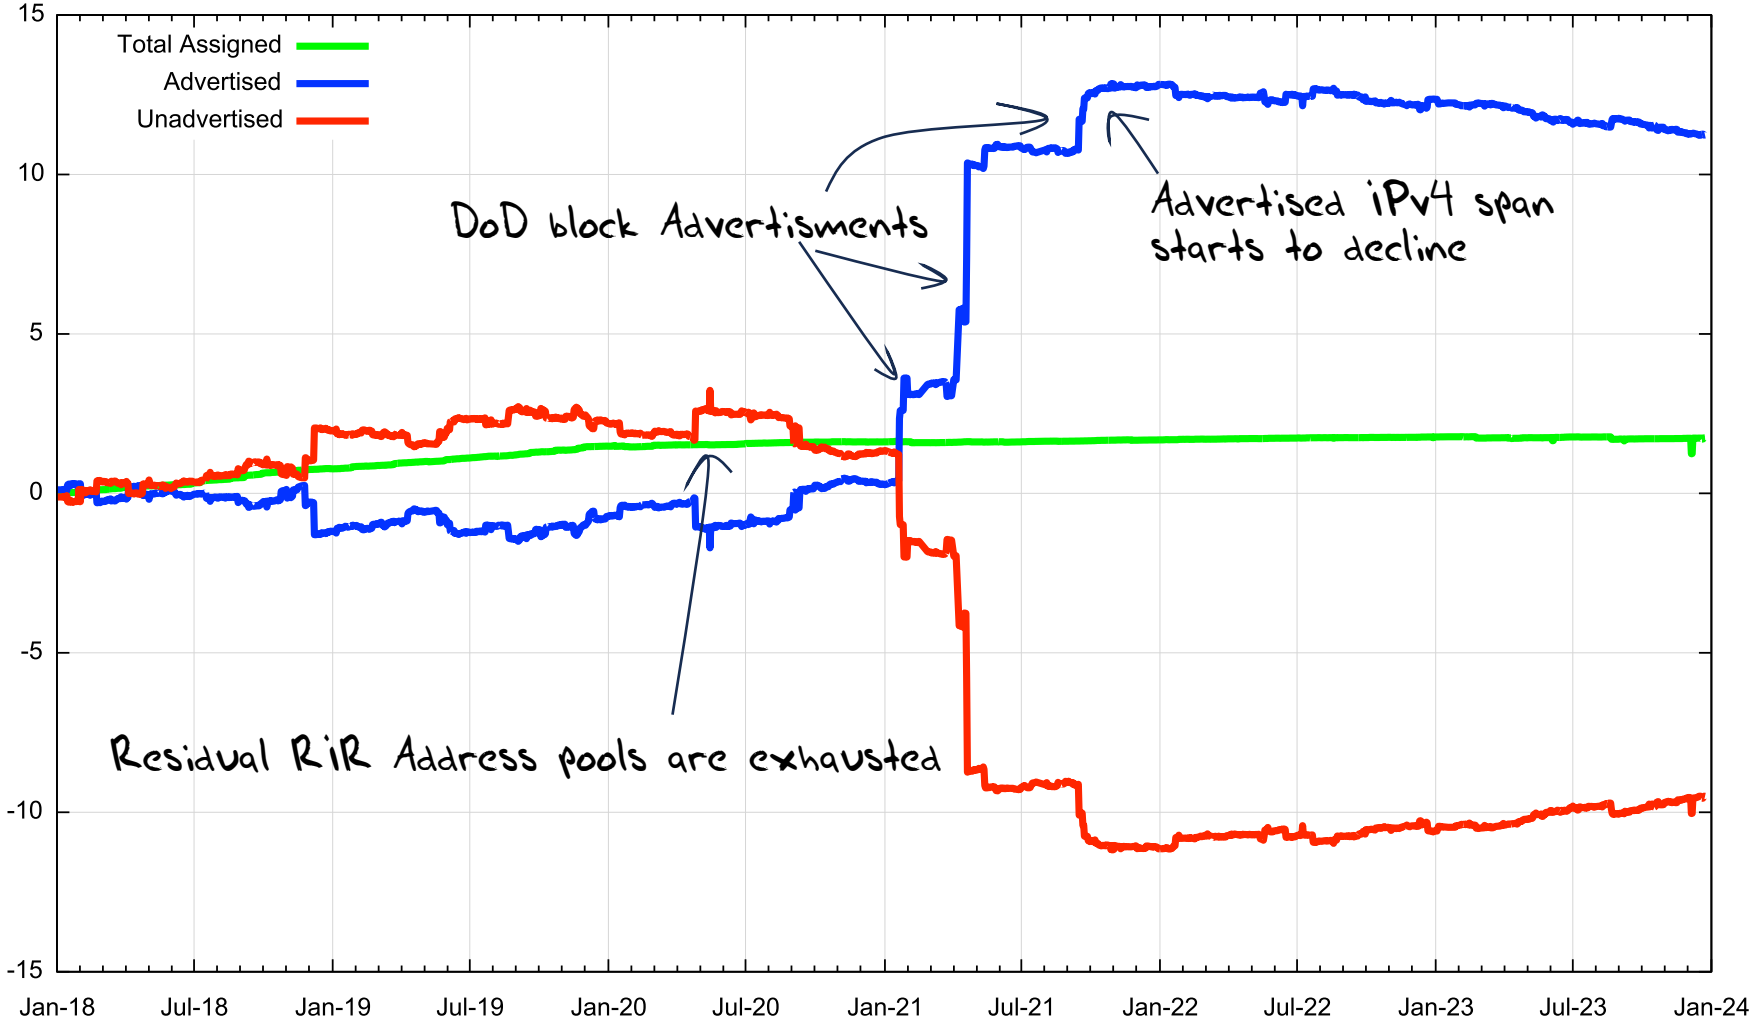
allocated delta

Change in the Advertised Address Pool

3.0 advertised
-18M delta



2018-2023: 6 Year Assigned vs Recovered



Advertised Span per 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 route set **reached a peak** at the end of 2023
- The percentage of new ASes added in 2023 was **only lower** than the rolling 5-year average (was 36,000 in 2021)
- The number of new ASes added in 2023 was slightly lower with **1,100 new AS's advertised in 2023** (was 1,400 in 2022)
- 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**

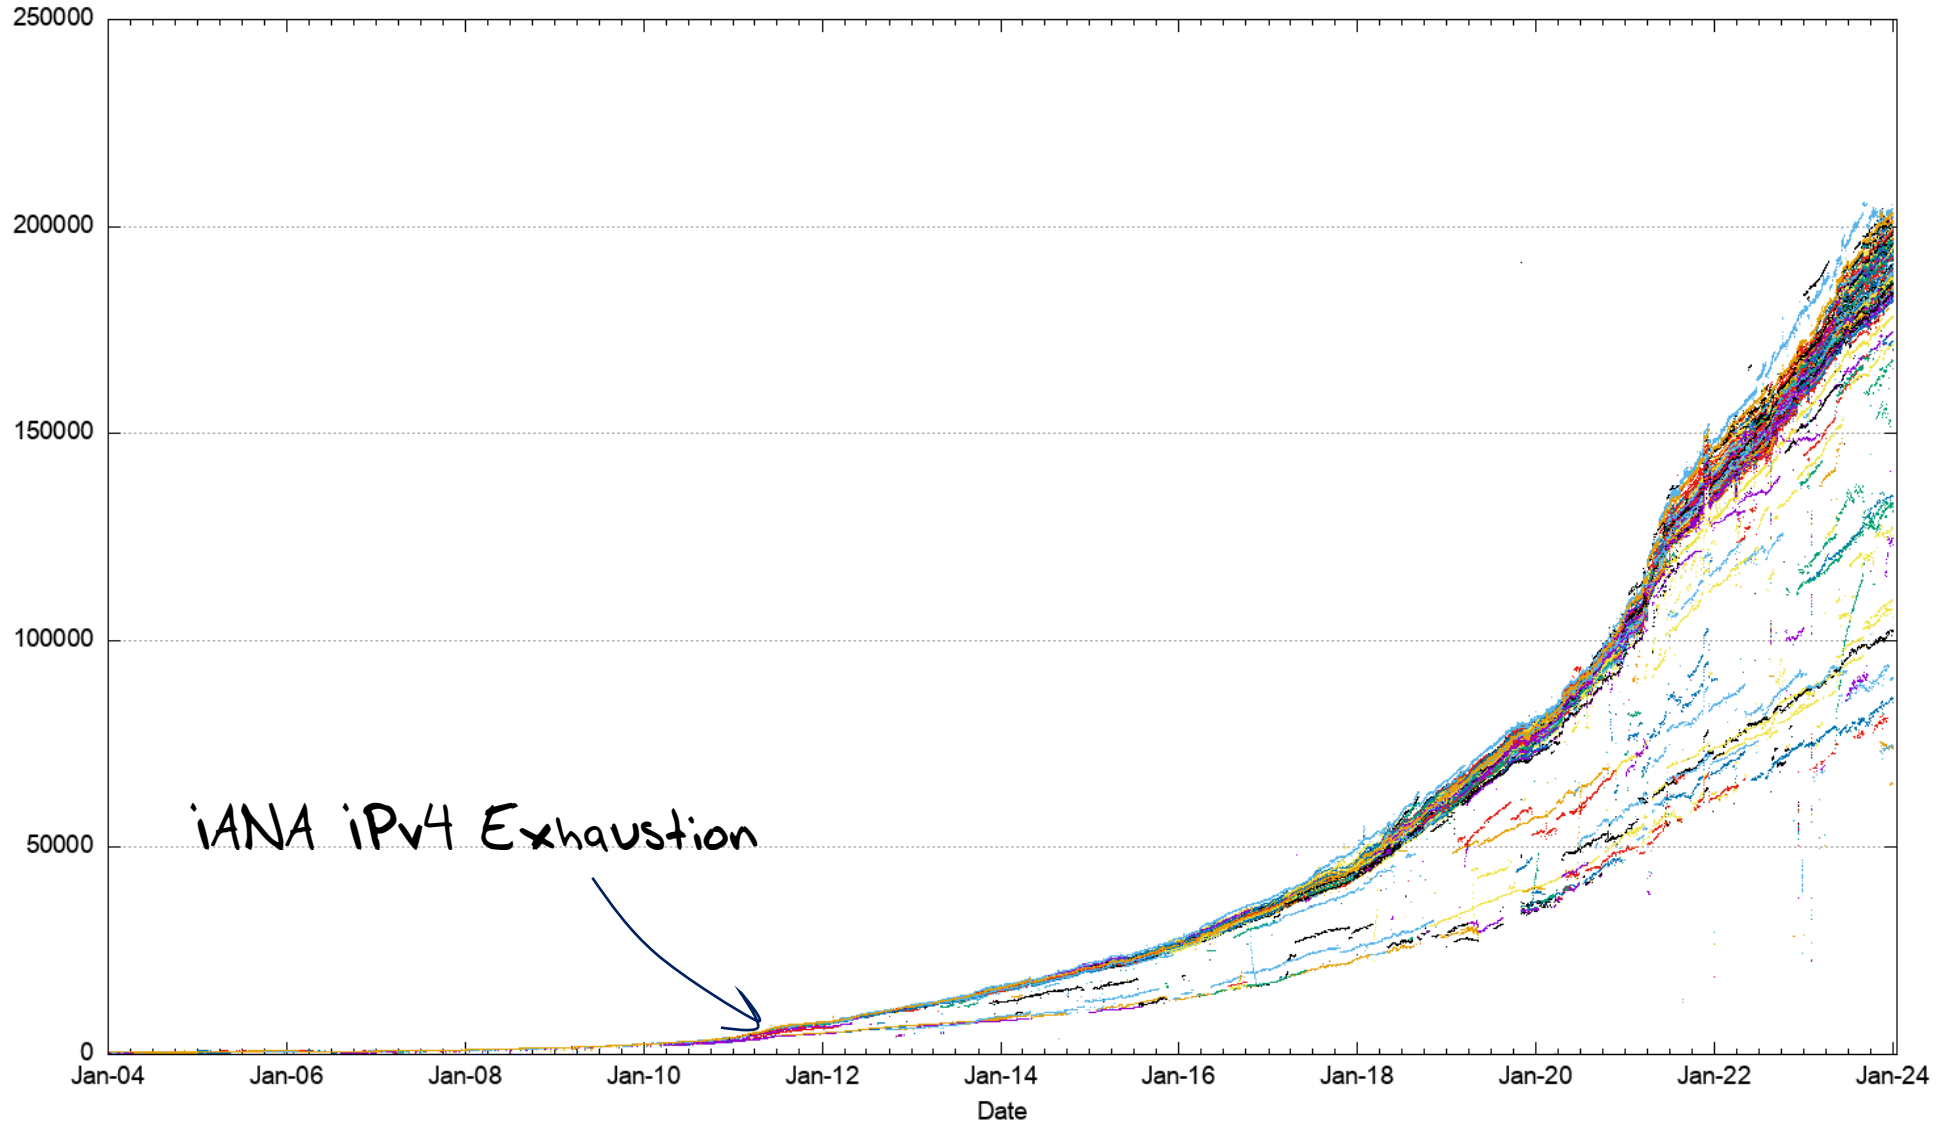
Have we reached Peak IPv4?

The Highlights


- IPv4 Summary
- IPv6 Summary
- FIB Projections
- Churn
- Conclusions

The 20-Year View of IPv6

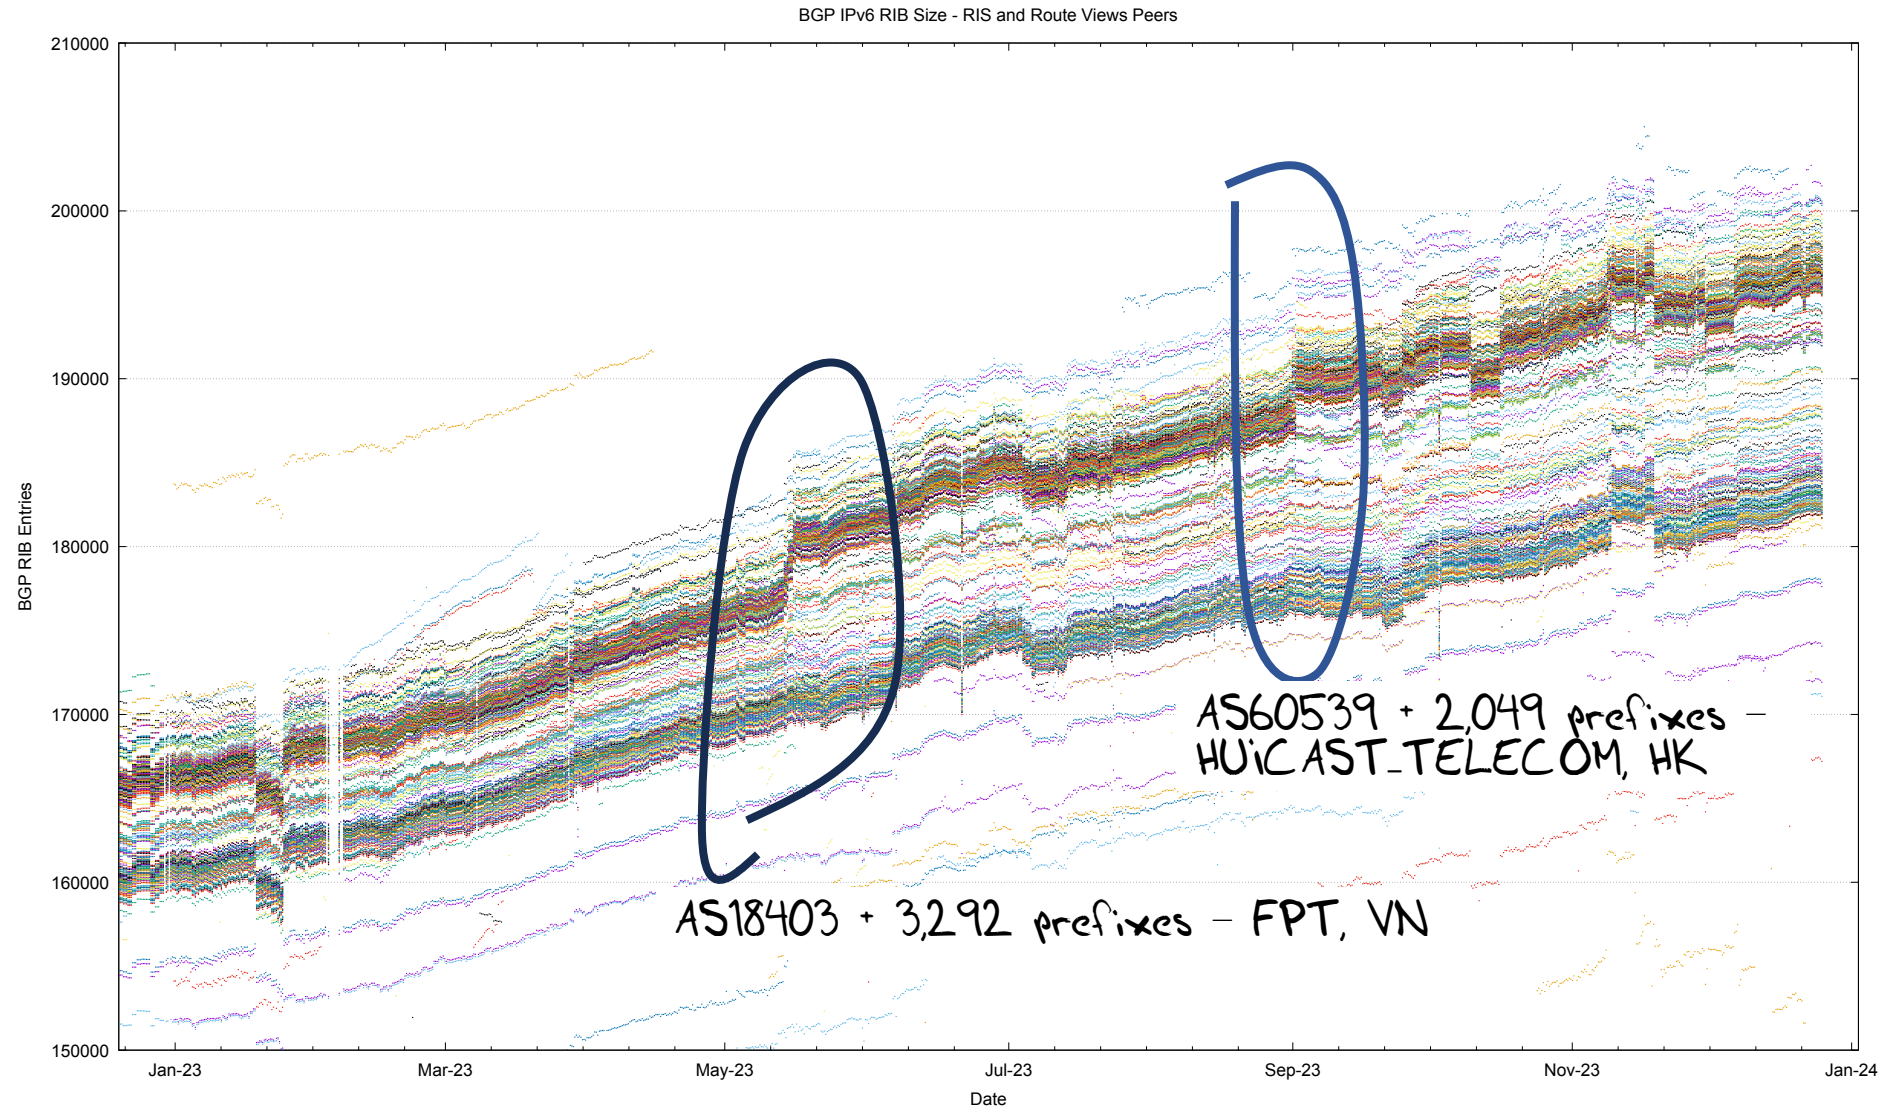
BGP IPv6 RIB Size - RIS and Route Views Peers




2023 in Detail



2023 in Detail



More Specifics in IPv6




45% of all IPv6 prefixes are /48's


IPv6

IPv4

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
- Conclusions

V4 BGP Table Size Predictions

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



V4 BGP Table Size Predictions

Date RIB Size Prediction

Jan 2019 760,000

2020 814,000

2021

2022

2023

2024

2025

2026

2027

2028

2029


958,000

958,000


947,000

928,000

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



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
- Conclusions


IPv4 BGP Updates - Daily Updates

Daily BGP v4 Update Activity for AS131072



IPv4 BGP Updates - Daily Updates


Daily BGP v4 Update Activity for AS131072



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

BGP v4 Daily Unstable Prefix Count




IPv4 Unstable Prefixes per Day


BGP v4 Daily Unstable Prefix Count



IPv4 BGP Convergence Performance



IPv4 BGP Convergence Performance




Updates in IPv4 BGP

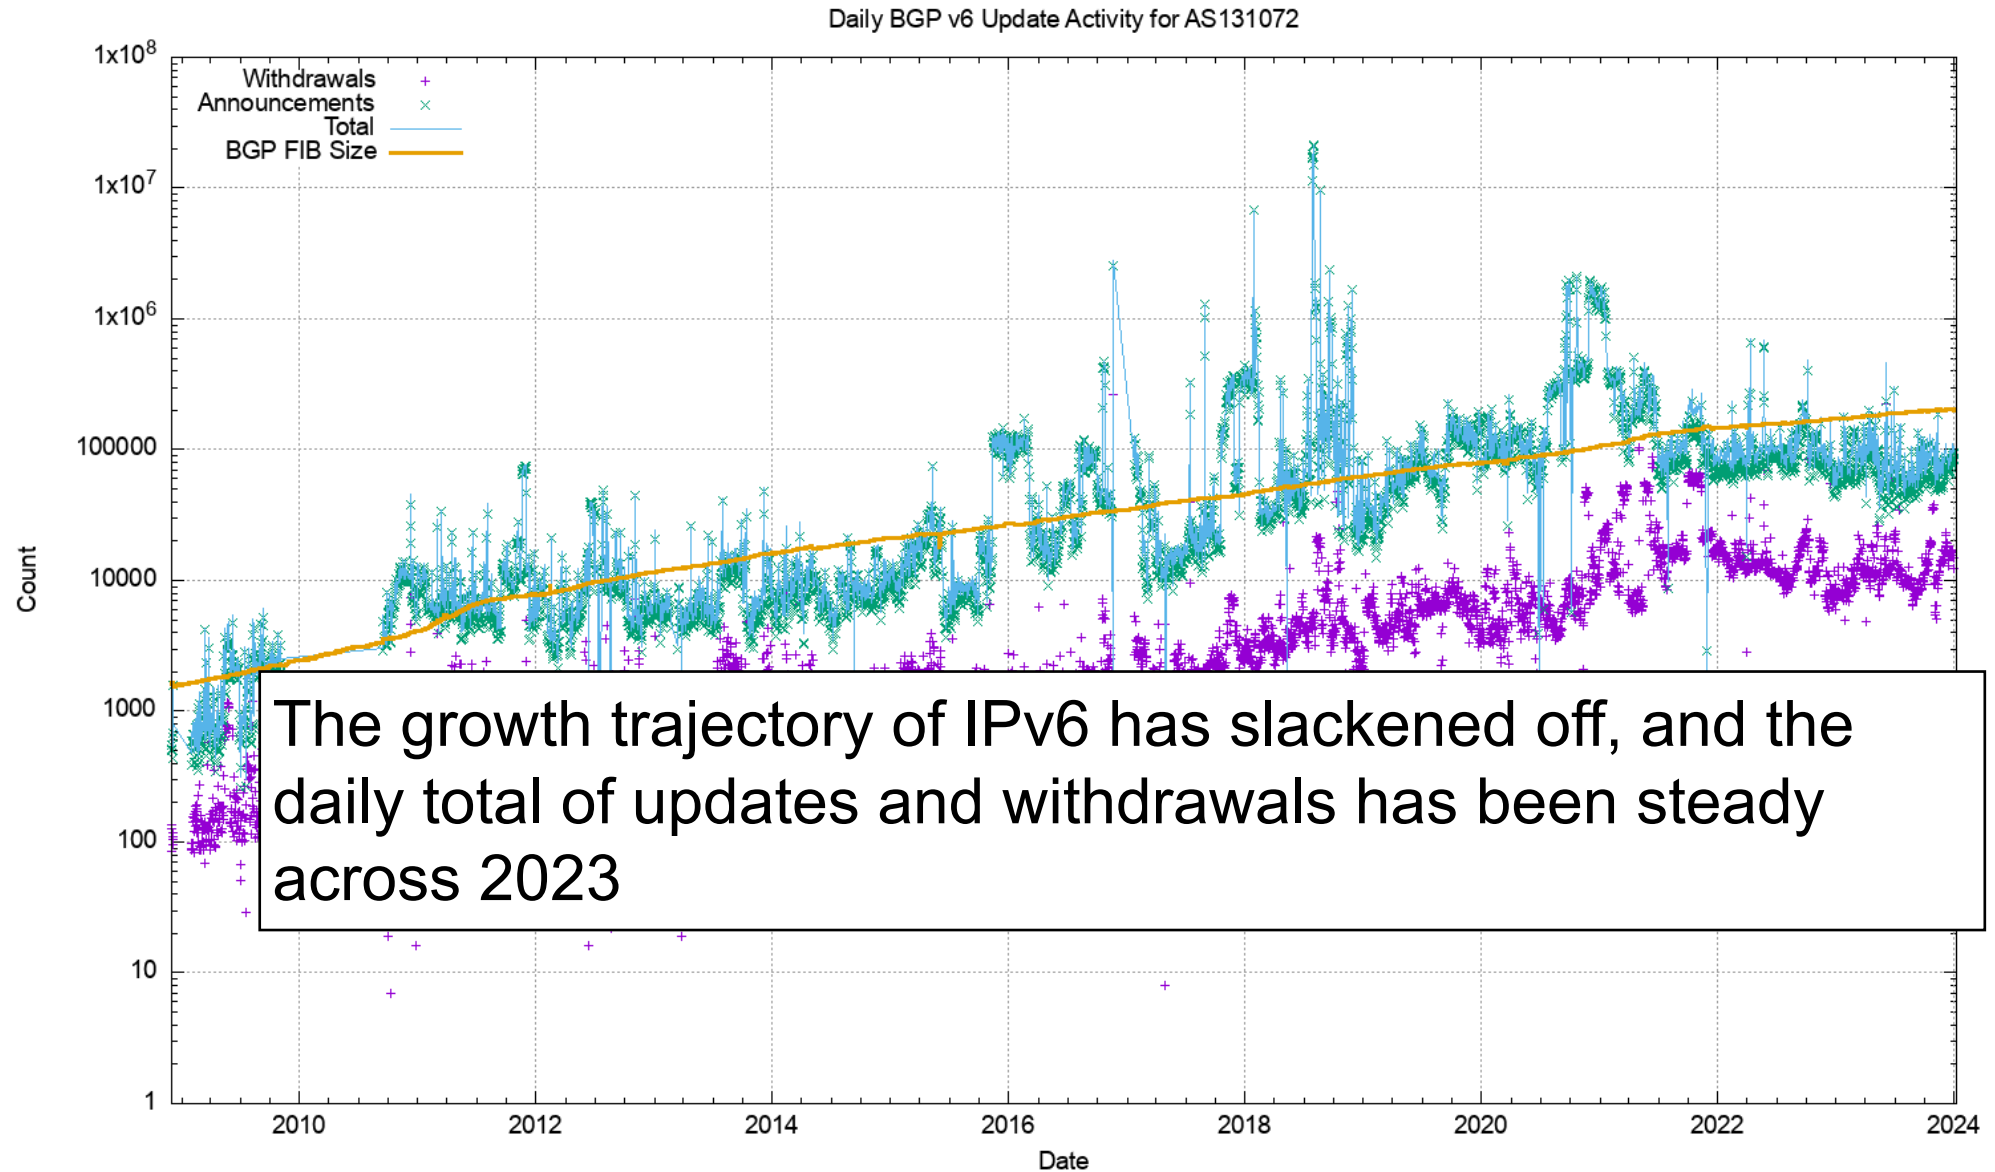
The IPv4 inter-domain routing system is still highly stable ...

- 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




V6 Unstable Prefixes




V6 Unstable Prefixes



V6 Convergence Performance




V6 Convergence Performance



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 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
- **Conclusions**

Routing Futures

- There is still little in the way of scaling pressure from BGP as a routing protocol – the relatively compressed inter-AS topology and stability of the infrastructure links tend to ensure that BGP remains effective in routing the internet.
- Instability levels are rising, generally driven by a small set of highly unstable “super generators”

Routing Futures

- The frenetic pace of expansion in the routing infrastructure of the Internet has slowed down across 2023. The IPv4 BGP network shrank slightly in 2023, and the pace of growth the IPv6 network also slackened off across the year.
- The drivers for growth were the product of the population of networks with discrete routing policies and the need to balance incoming traffic across multiple paths (traffic engineering) – the rapid increase in the use of CDN platforms has reduced the dependence on transit routes to provide content and services to end users.
- Much of the overall volume of content traffic has shifted across into privately operated platforms, and the demands on the public routed common infrastructure have started to decline

Some Practical Suggestions

For those network operators whose service needs are dependent on the BGP network to some extent (which is mostly everyone)

Know your network's limits:

- Understand your routing FIB capacity in the default-free parts of your network
- There may be some default max prefix setting of ~1M FIB entries which will cause a BGP session shutdown when tripped
 - AS 4804 (Optus) in Australia appeared to encounter this situation in October 2023
- Use a max prefix exceeded switch which avoids sessions shutdown where possible

Some Practical Suggestions

Know your network's limits

Review your routers' settings

- Review your IPv4 / IPv6 portioning in the FIB tables - a dual-stack eBGP router will conservatively need a 1.2 M 32-bit IPv4 slots and 320K 128-bit IPv6 slots for a full eBGP routing table in line cards by 2026 if they are using a full eBGP FIB load (plus internal routes of course). That's roughly the same memory footprint for IPv4 and IPv6!

Some Practical Suggestions

Know your network's limits

Review your routers' settings

Default routes can be helpful

- Judicious use of **default** routes in your internal network may allow you drop this high speed line card memory requirement significantly

Some Practical Suggestions

Know your network's limits

Review your routers' settings

Default routes can be helpful

Time for hot caching in line card FIBs?

- Using a hot cache for line card FIB cache would reduce the high-speed TCAM memory requirement significantly without visible performance cost

Some Practical Suggestions

Know your network's limits

Review your routers' settings

Default routes can be helpful

Time for hot caching in line card FIBs?

That's it!

Questions?