

# NO SHORTCUTS TO LONG PREFIXES

As we squeeze IPv4 address space ever further, we ought to revisit the question of how prefixes longer than a /24 propagate across the internet. The obvious risk with longer prefixes is the potential for an explosion in routing table sizes. At the same time, ARIN has an address block -- 23.128.0.0/10 -- from which it will allocate blocks as small as a /28.

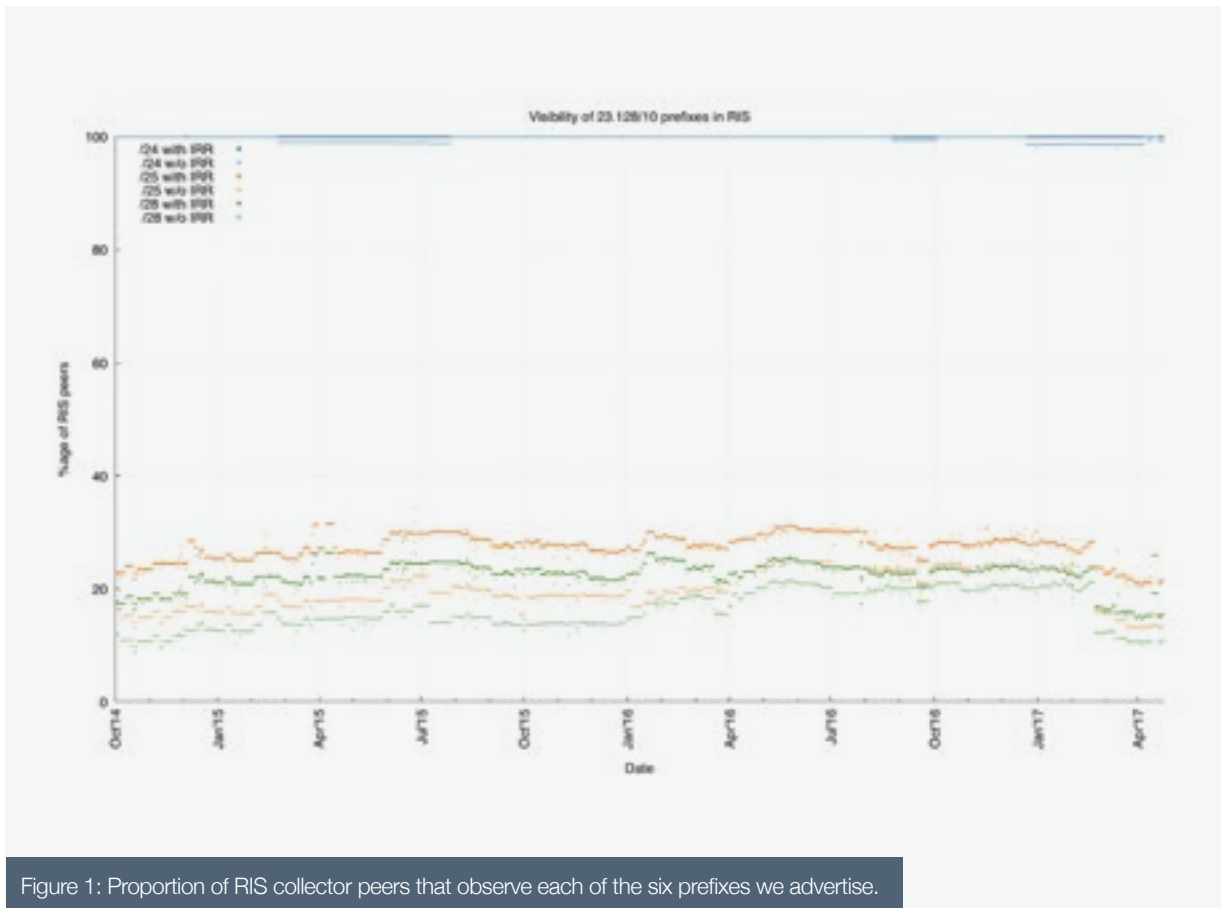


Figure 1: Proportion of RIS collector peers that observe each of the six prefixes we advertise.

Since October 2014, the RIPE NCC has been advertising six IPv4 prefixes drawn from this range (two /24s, two /25s, and two /28s) to better understand how they propagate across the network, to determine how they behave in the wild, and to assist decision-making for policy on smaller allocations. Currently, less than 1% of this /10 has been allocated

and none of those allocations were smaller than a /24. If smaller allocations were made, how useful would they be for public routing?

Fig. 1 shows the fraction of RIS (RIPE Routing Information Service) peers that observe the six prefixes. The first thing to note is that visibility of the longer prefixes hasn't changed much

over the years: at best, we might expect a /25 to be visible by 20-25% of our peers. Next, we see a drop in visibility between 21 and 22 of February 2017. This drop stems from stronger filtering after switching to new route servers at ASN3333. The trend, in any case, is definitely not increasing.

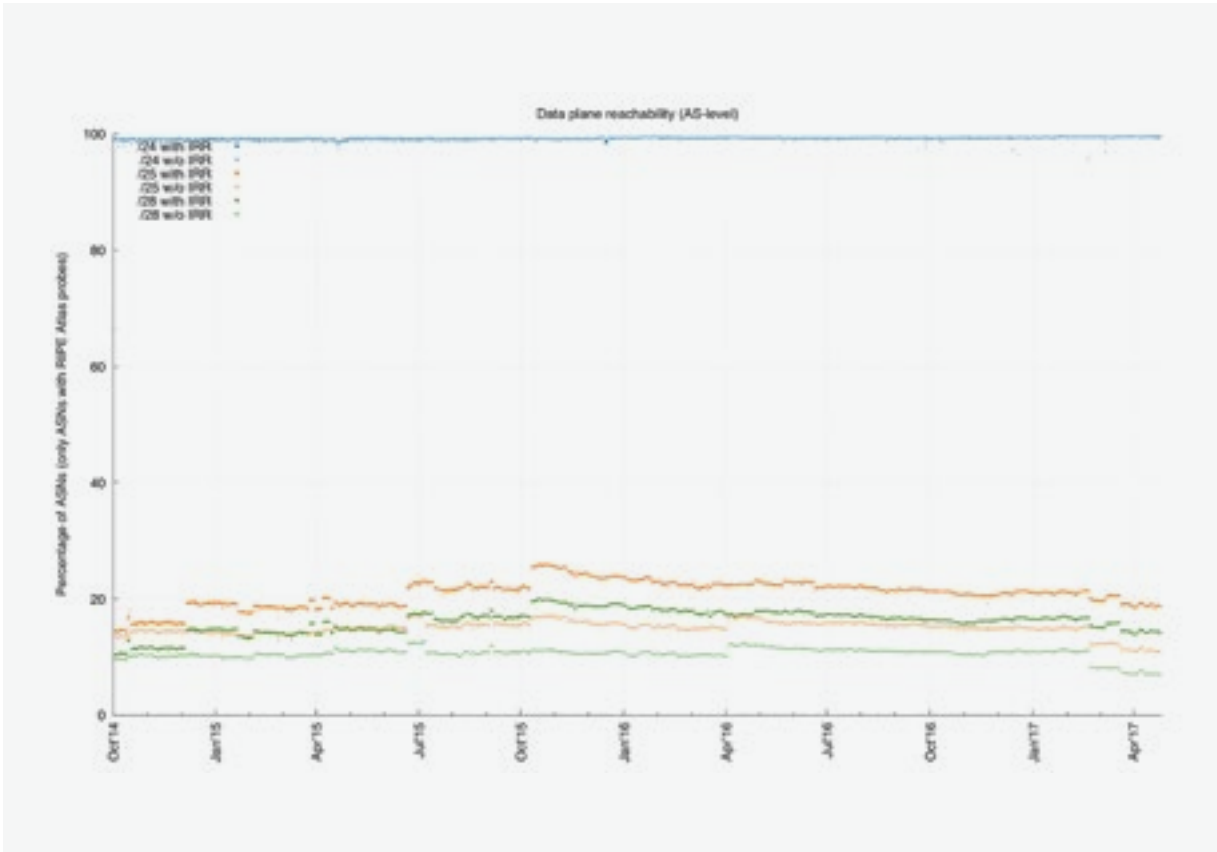


Figure 2: Traceroute reachability of all six prefixes from all RIPE Atlas probes

We are able to corroborate this with other datasets:

- The Route Views project is another source of archived BGP data, offering additional vantage points from which to observe prefix propagation. On 1 May 2017, Route Views collectors observe similar visibility to RIS collectors: /24s propagate widely, but the /25s and /28s have limited visibility. The longer prefixes that have route objects in the routing registry propagate further than those that don't, a slight advantage for networks when they want to advertise a longer prefix onto the public network.
- Packet Clearing House (PCH) provides BGP table dumps for route collectors located in IXPs around the world. Many of these are not full tables: 89 out of 135 observe fewer than 100,000 unique prefixes; 8 see more than 300,000 unique prefixes, and only one appears to have a full table. Given this, on 1 May the PCH dataset shows the /24 with the route object in 11 of the BGP table dumps, but seven see the longer prefixes with route objects, and only 1 BGP table contains all six. So the pattern is similar, in that there

are common cases where a route object will assist the propagation of a long prefix, even if they will not propagate as widely as a /24.

We are also able to use RIPE Atlas to conduct long-running active measurements. Within each BGP advertisement, we have one address which is responsive for exactly this purpose, giving us insight into reachability from a broad sample of the public network. Fig. 2 shows the data-plane reachability from RIPE Atlas probes. What we observe matches the above: most probes cannot reach addresses within the longer prefix.

The received wisdom is that we should not advertise IPv4 prefixes longer than a /24 to the global network because they won't achieve full coverage. The data bears this out. Further, the lack of path diversity makes reachability less stable. From this, it's clear a /25 is still not very useful if you want global reachability. But there are some differences, and it's clear that an appropriate route object in the IRR helps, and perhaps strong RPKI and IRR filters as requirements for long prefixes in specific ranges (such as this one) may be the way towards permitting widespread advertisements longer than a /24.

For more details, please see BGP Even More Specifics in 2017 on RIPE Labs: [https://labs.ripe.net/Members/stephen\\_strowes/bgp-even-more-specifics-in-2017](https://labs.ripe.net/Members/stephen_strowes/bgp-even-more-specifics-in-2017)

## About the RIPE NCC

The RIPE NCC is the Regional Internet Registry for Europe, the Middle East and parts of Central Asia. As such, we allocate and register blocks of Internet number resources to Internet service providers (ISPs) and other organisations. We're a not-for-profit organisation that works to support the RIPE (Réseaux IP Européens) community and the wider Internet community. The RIPE NCC membership consists mainly of Internet service providers, telecommunication organisations and large corporations.