

AS paths: long, longer, longest

RIPE-79, 18 October 2019

Ijitsch van Beijnum

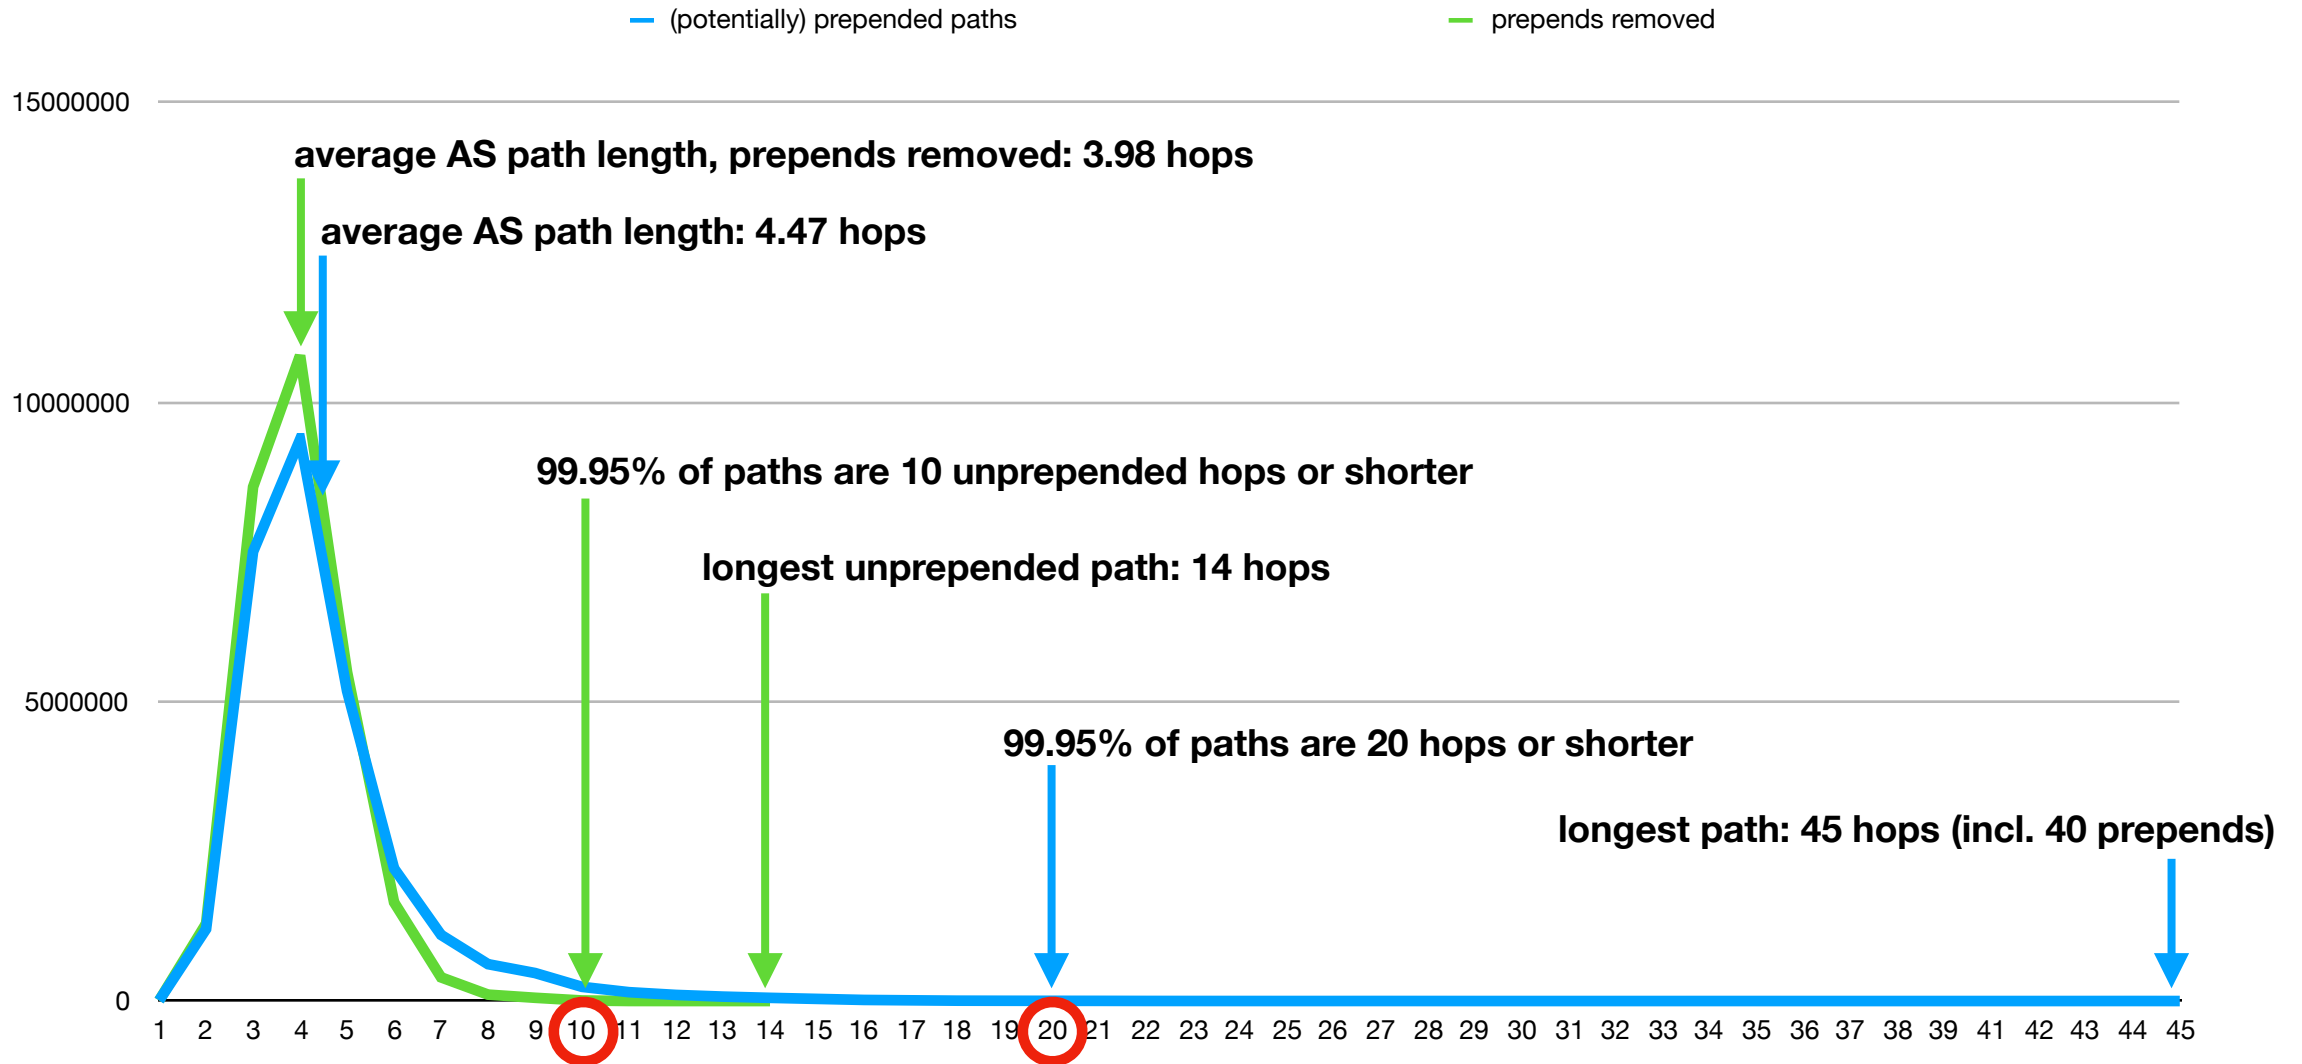
How long is too long?

- On Tuesday we heard from Doug Madory about how people prepend *wide*: towards (almost) all their neighbors
- I want to talk about how *long* the prepended paths are
- And:
 - some people filter very long AS paths
 - shouldn't we come up with a community recommendation about AS path length filtering?
 - if so: what's a reasonable limit?

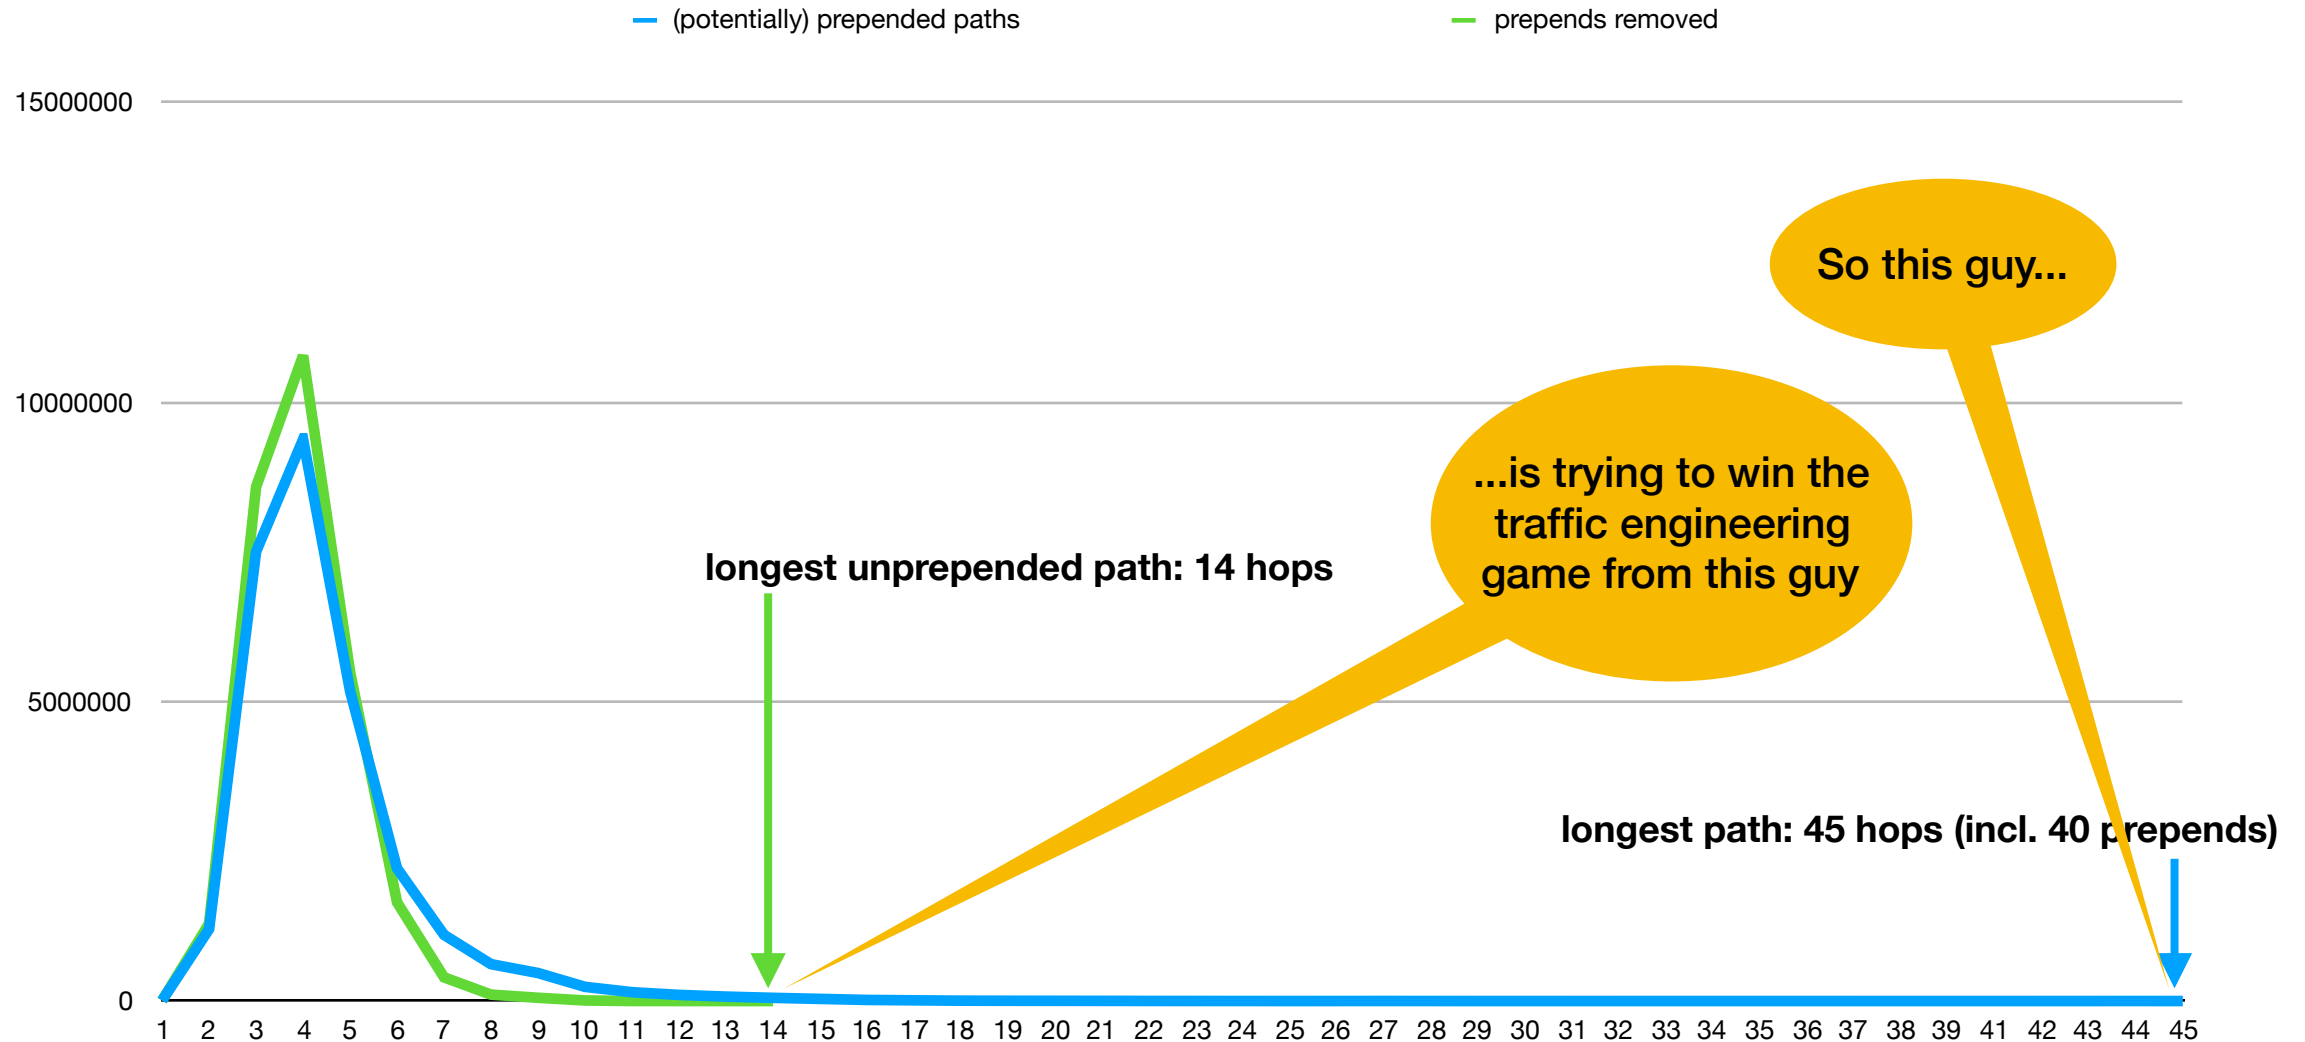
We say in Dutch: meten is weten

- First some measurements. My methodology:
 - I got an IPv4 BGP table dump from Routeviews ↓ (warning: just a single snapshot)
 - imported the whole thing into MySQL 🤪
 - removed default routes and prefixes > /24 🚫
 - removed paths from neighbor ASes with < 700k routes 🧠
 - removed duplicate ASes 🧐
 - result: 28 million prefixes from 37 neighbors
 - tried to remember what I learned in school about SQL queries... 🤔

AS paths in the wild



Overdoing things slightly



How to win TE

- The first prepend is extremely effective, often *too* effective
- But: rapidly diminishing returns
- At 7 hops, you're already longer than 98% of the unprepended internet
 - and 90% of the prepended internet
- That's at most 5 prepends
- If that's not enough, you should probably look for something other than prepending to get the results you need

Questions for you

- Do we need community guidelines for a prepending limit?
 - my answer: yes. On Tuesday, different people mentioned different AS path length limits. This means that preppers will get into trouble, but they have no idea when
- If yes, where do we publish? RIPE document, (informational) RFC?
- If yes, then what should that limit be?

AS path length filtering options

- 255 is a magic number in BGP, we'll want to stay below that
- If we see 45 in the wild, then two people with 45 hops could need 90 hops to reach each other
- If we see 14 unpreended hops in the wild, and each of those does a perfectly reasonable 3 prepends, that's 56 hops
 - $x 2 = 112...?$

My thoughts

- Ask origin networks to make sure they don't announce paths longer than what's needed for reasonable TE, say 10 hops (running total: 10)
- Ask transit networks to make sure they don't add prepends on paths that are already long, say 20 hops (running total: 20)
- Reserve hops for the other half of the path, say 10 hops (running total: 30)
- Add a margin for error, say 10 hops (running total: 40)
- Add a second margin for even more error, say another 10 hops (limit: 50)

Thanks, all!



