Measuring RPKI

Geoff Huston

APNIC

comments

- Draft pack for the PC – there is more material to come but I wanted to show the PC the general intent of the presentation and the issues I want to raise in it.

- I expect to add slides of a few countries (Bhutan and China are \good examples) and networks

- I also want to explore next steps in this measurement

- Geoff

Objective

- The measure the “impact” of invalid route filtering on users
- The question we want to answer here is user-centric:
  - What proportion of users can’t reach a destination when the destination route is invalid according to ROV?
- We’d like this as a long term whole-of-Internet measurement

Methodology

- Set a prefix and AS in a delegated RPKI repository
- Regularly revoke and re-issue ROAs that flip the validity state between valid and invalid states
- Anycast the prefix and AS pair in a number of locations across the Internet
- Load the URL of the destination into a measurement script
- Feed the script into the advertising systems
- Collect and analyse data

Flipping ROA states

- What's a good frequency to flip states?
- What's the re-query interval for clients of a RPKI publication point?
  - There is no standard-defined re-query interval so implementors exercised their creativity

Re-Query Intervals

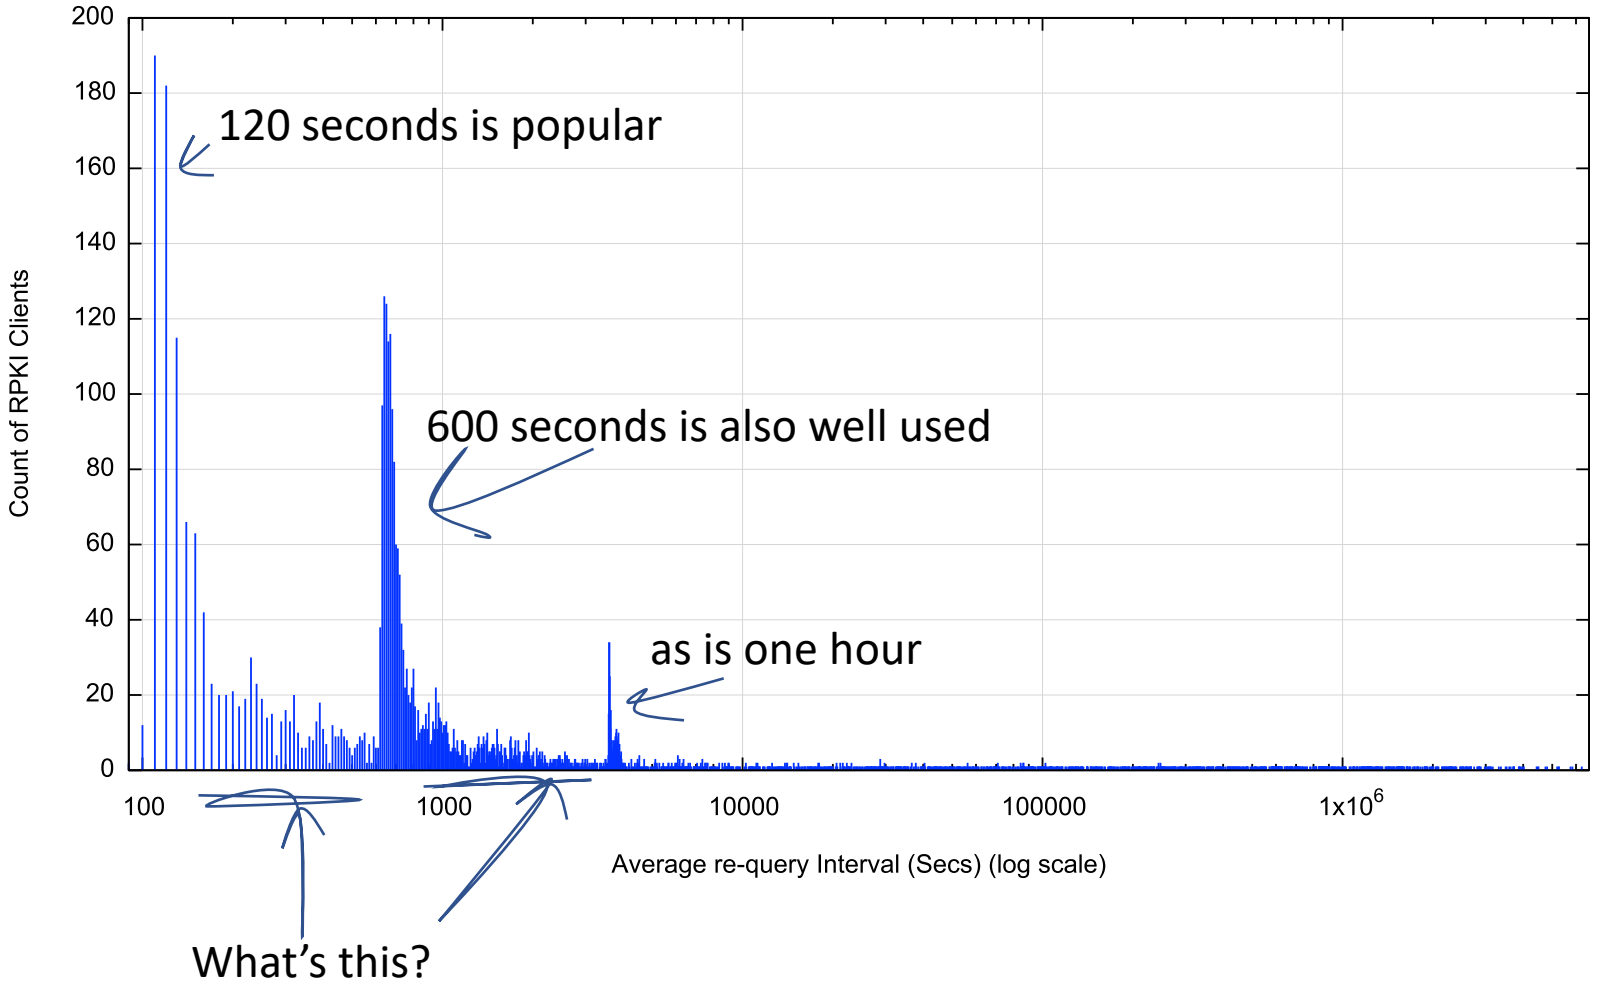

Re-Query Intervals (first hour)

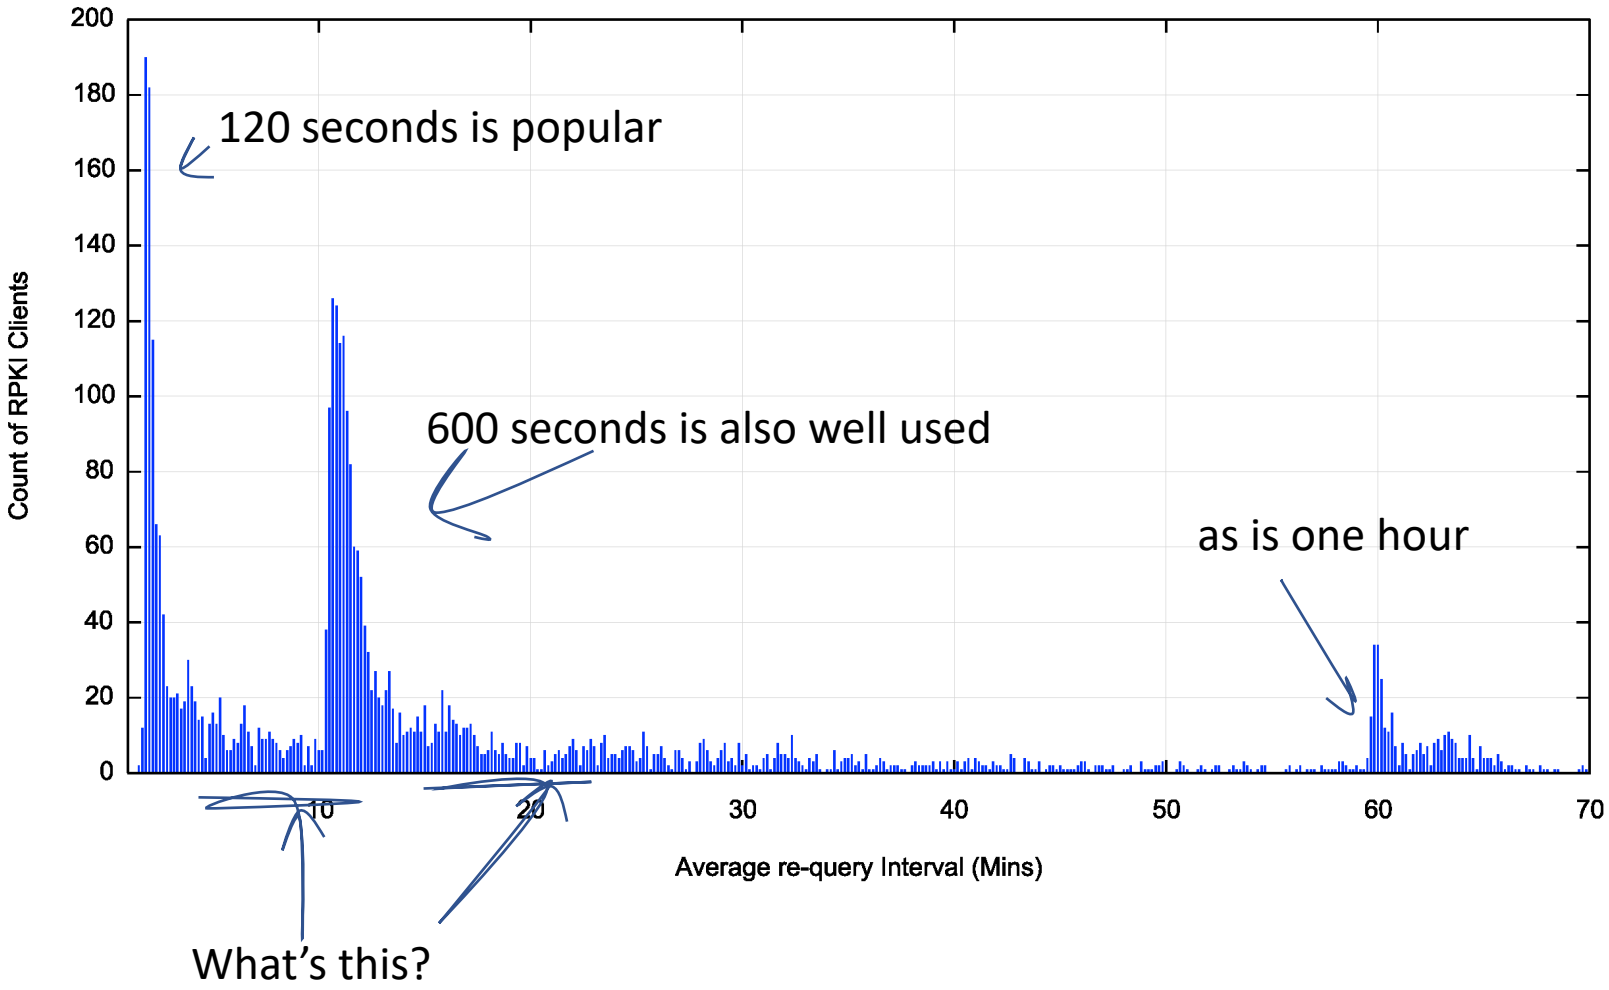

Re-Query – Cumulative Distribution

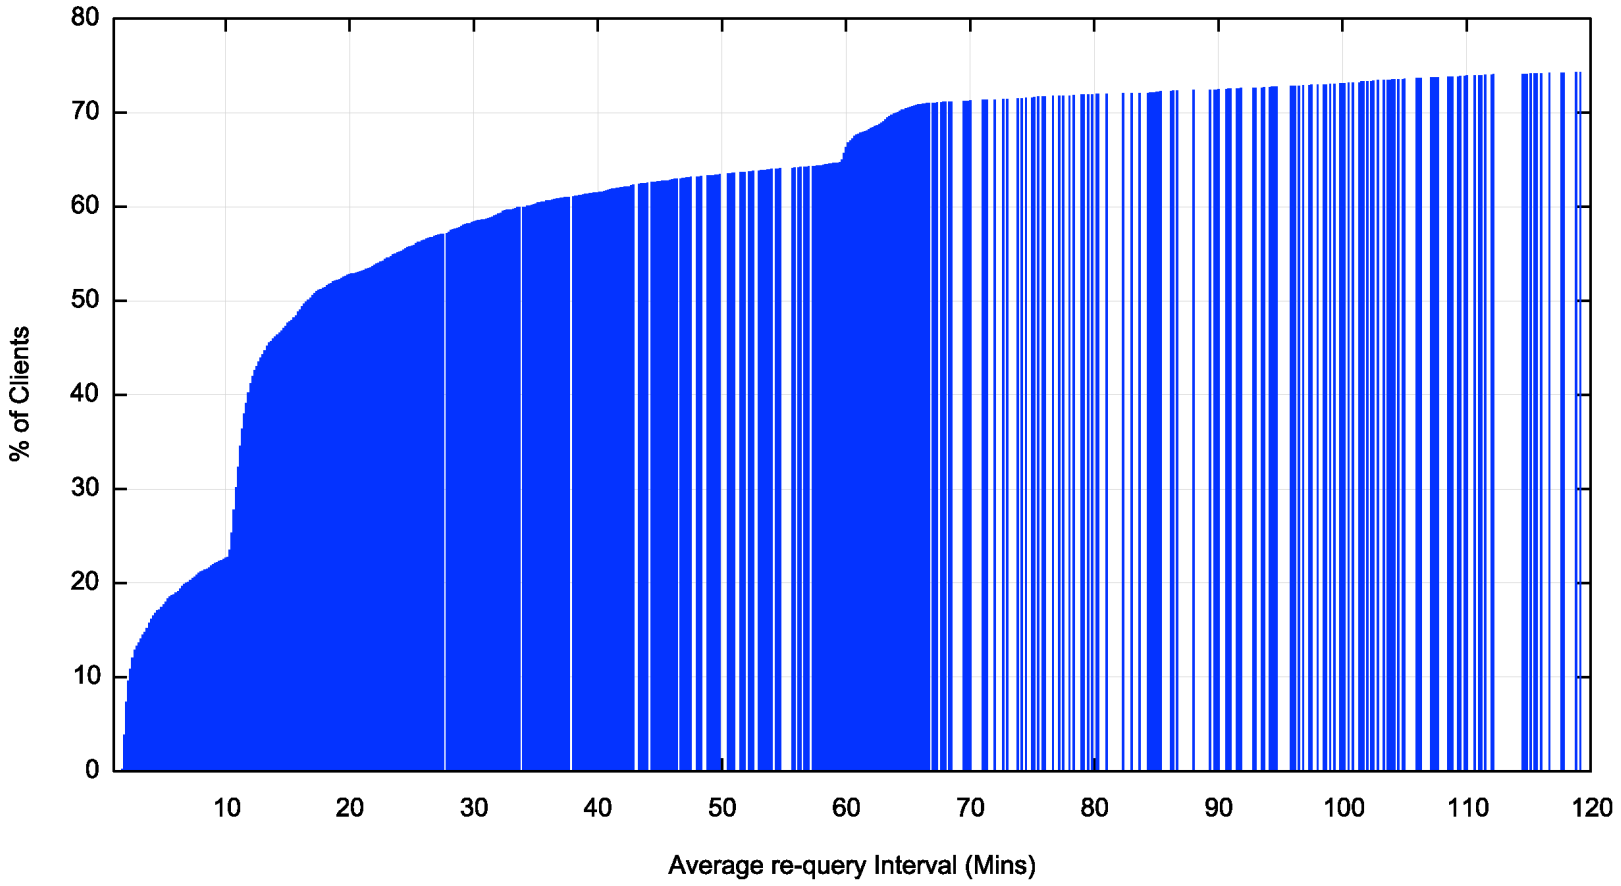

At 2 hours we see 75% of clients perform a requery

Why the lag?

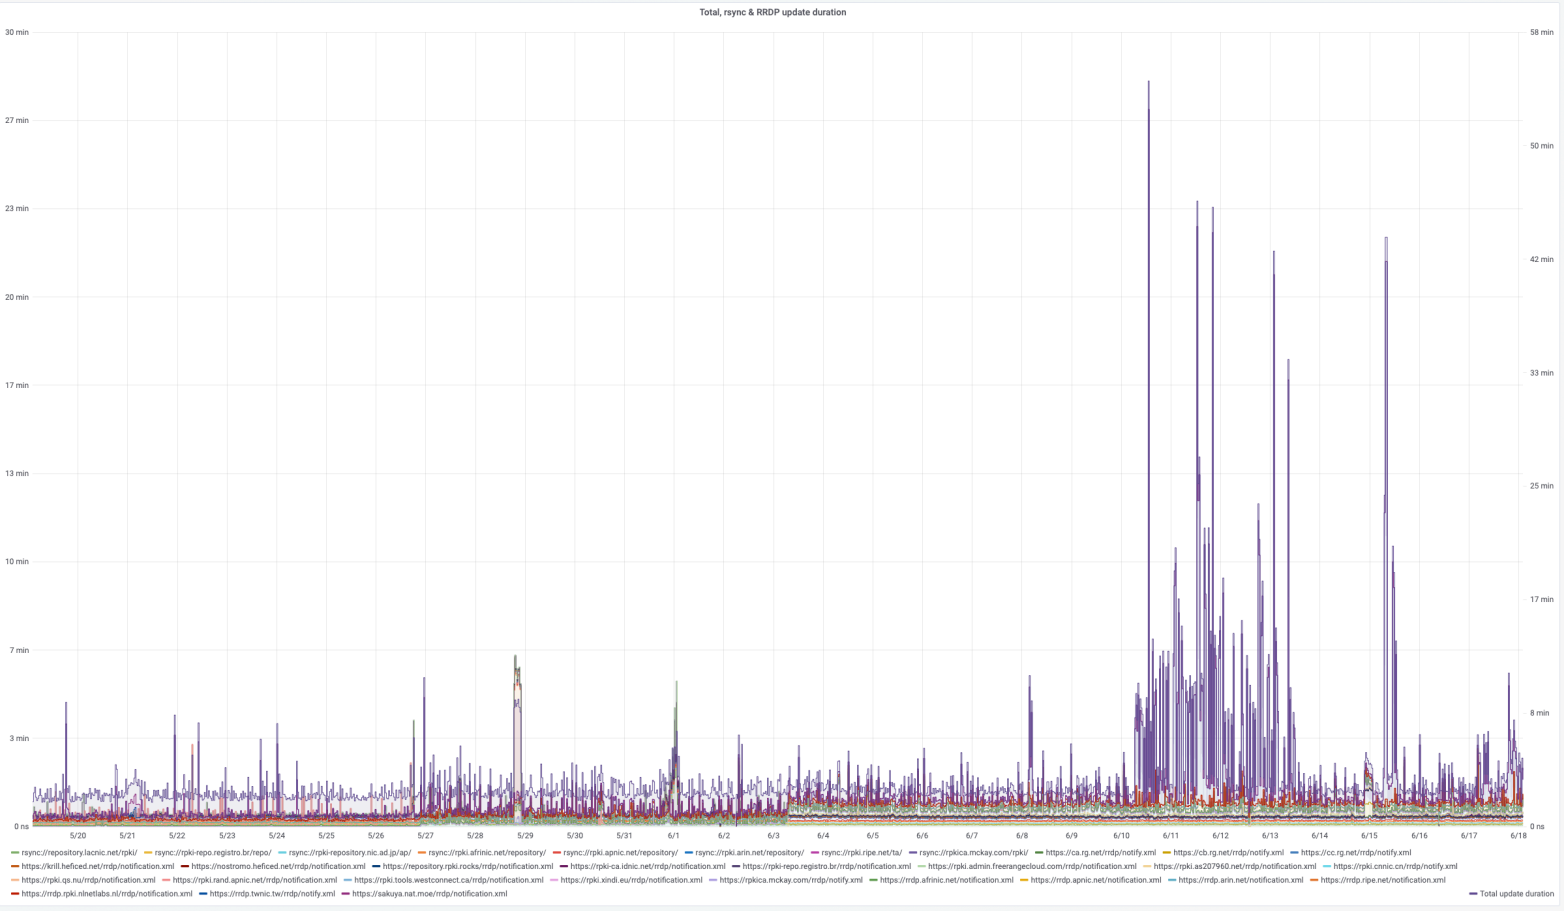

https://grafana.wikimedia.org/d/UwUa77GZk/rpki?panelId=59&fullscreen&orgId=1&from=now-30d&to=now

Clients can take a significant amount of time to complete a pass through the entire RPKI distributed repository set, which makes the entire system sluggish to respond to changes

We used 12 and 36 hour flips

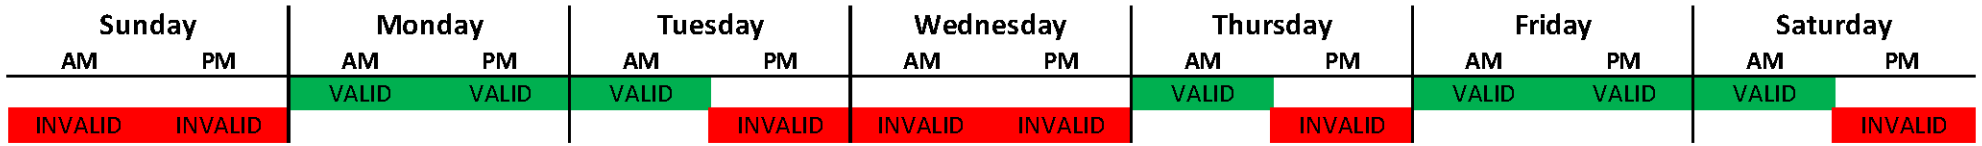

We used 12 and 36 hour flips

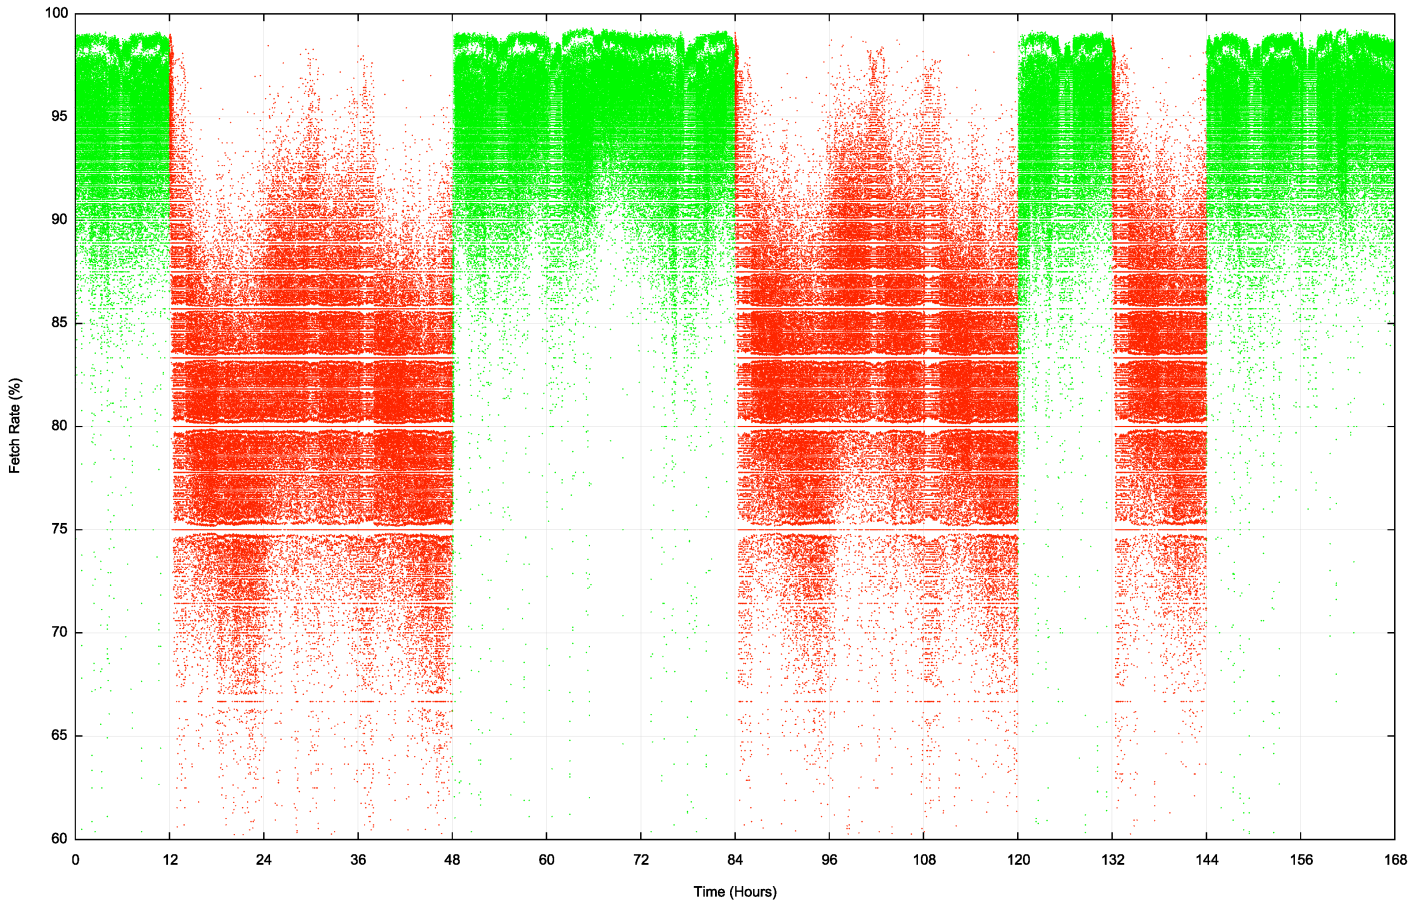

This shows the per-second fetch rate when the route is valid (green) and invalid (red) over a 7 day window

The route validity switches are clearly visible

Transition – Valid to Invalid

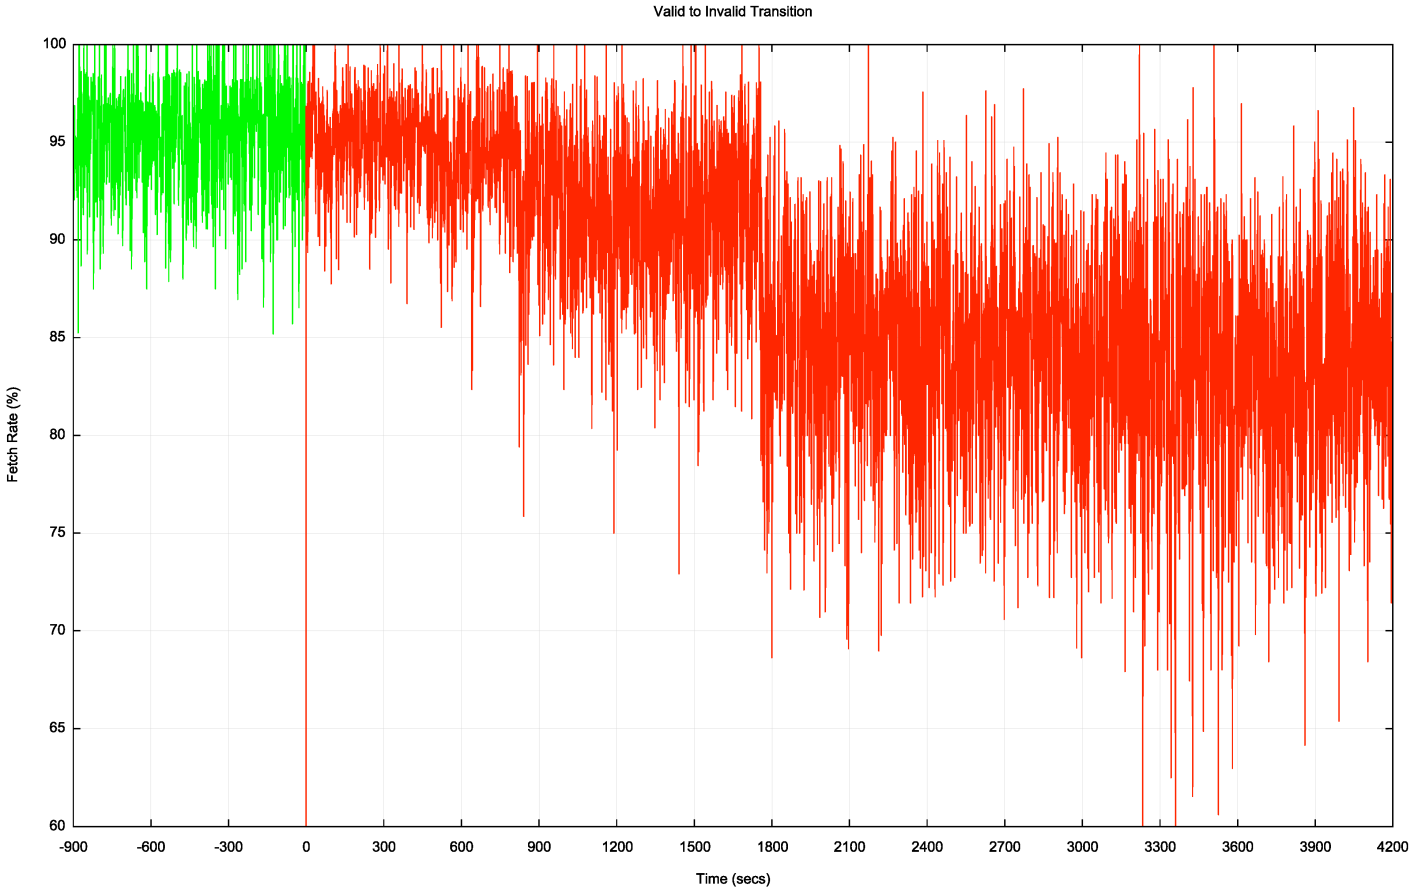

It takes some 30 minutes for the valid to invalid transition to take effect in this measurement

It appears that this is a combination of slow re-query rates at the RPKI publication point and some delays in making changes to the filters being fed into the routers

This system is dependant on the last transit ISP to withdraw

Transition – Invalid to Valid

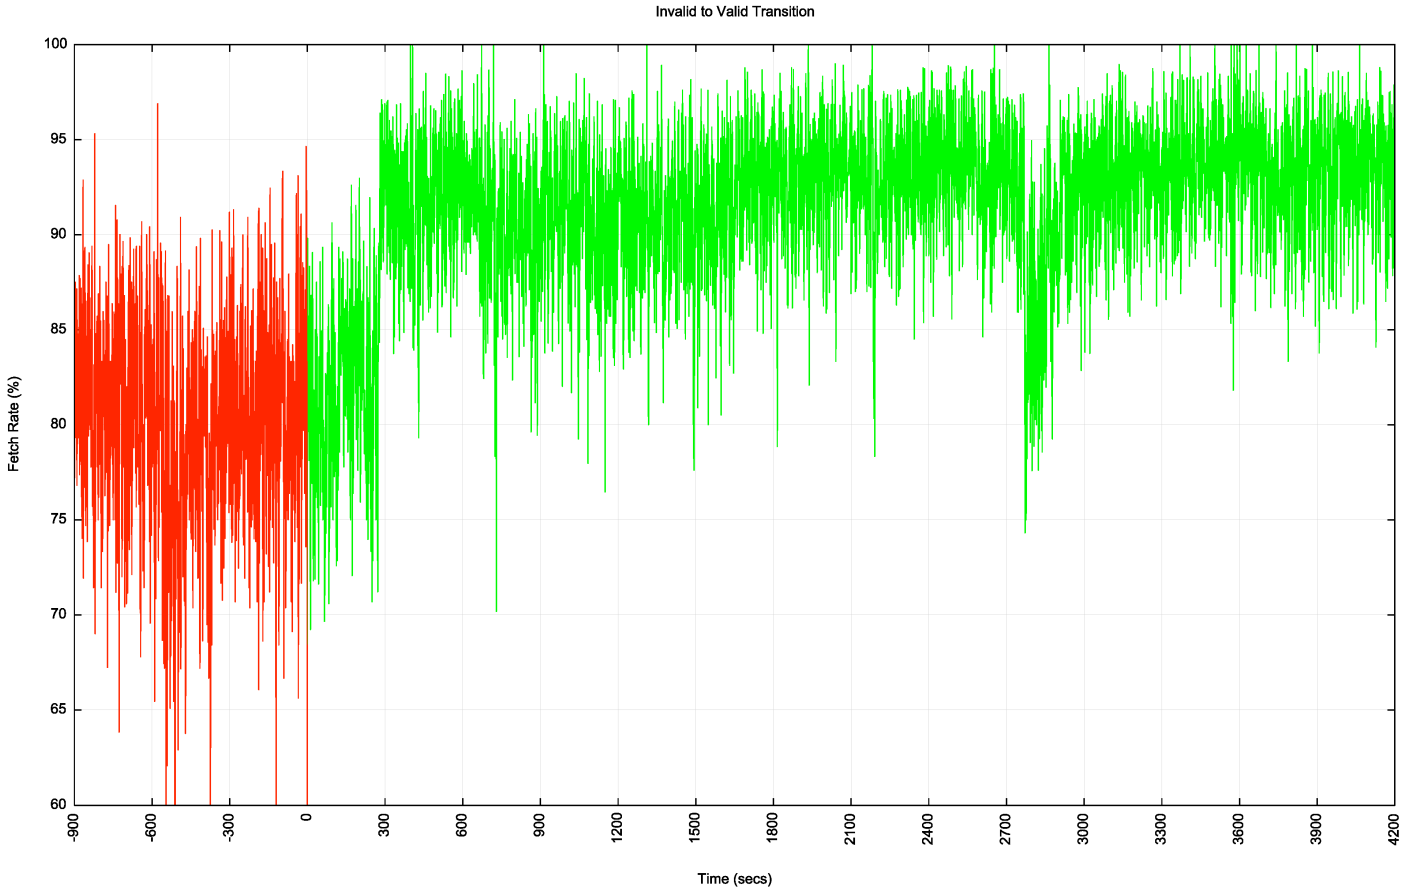

It takes some 5 minutes for the invalid to valid transition to take effect in this measurement

This system is dependant on the first transit ISP to announce, so it tracks the fastest system to react

RPKI “sweep” software

- There is a mix of 2, 10 and 60 minute timers being used
- 2 minutes seems like a lot of thrashing with little in the way of outcome – the system is held back by those clients using longer re-query timers
- 60 minutes seems too slow
- I’d go with a 10 minute query timer as a compromise here

User impact of RPKI filtering

Use of RPKI Validation for World (XA)

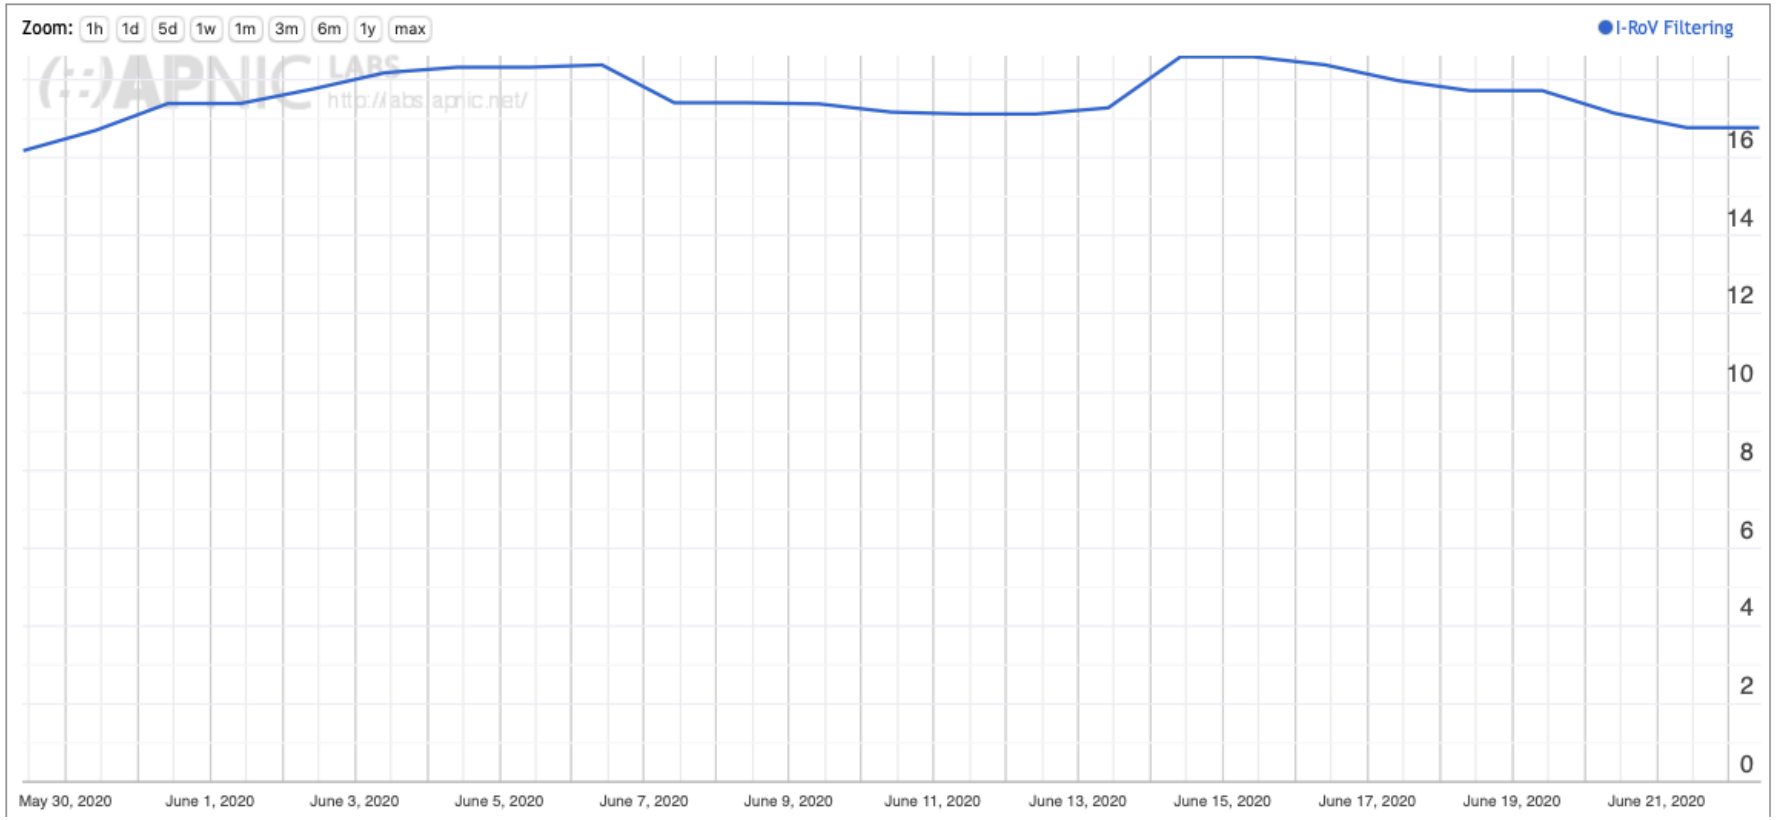

At 16% of users that's surprisingly large for a very recent technology

User impact of RPKI filtering

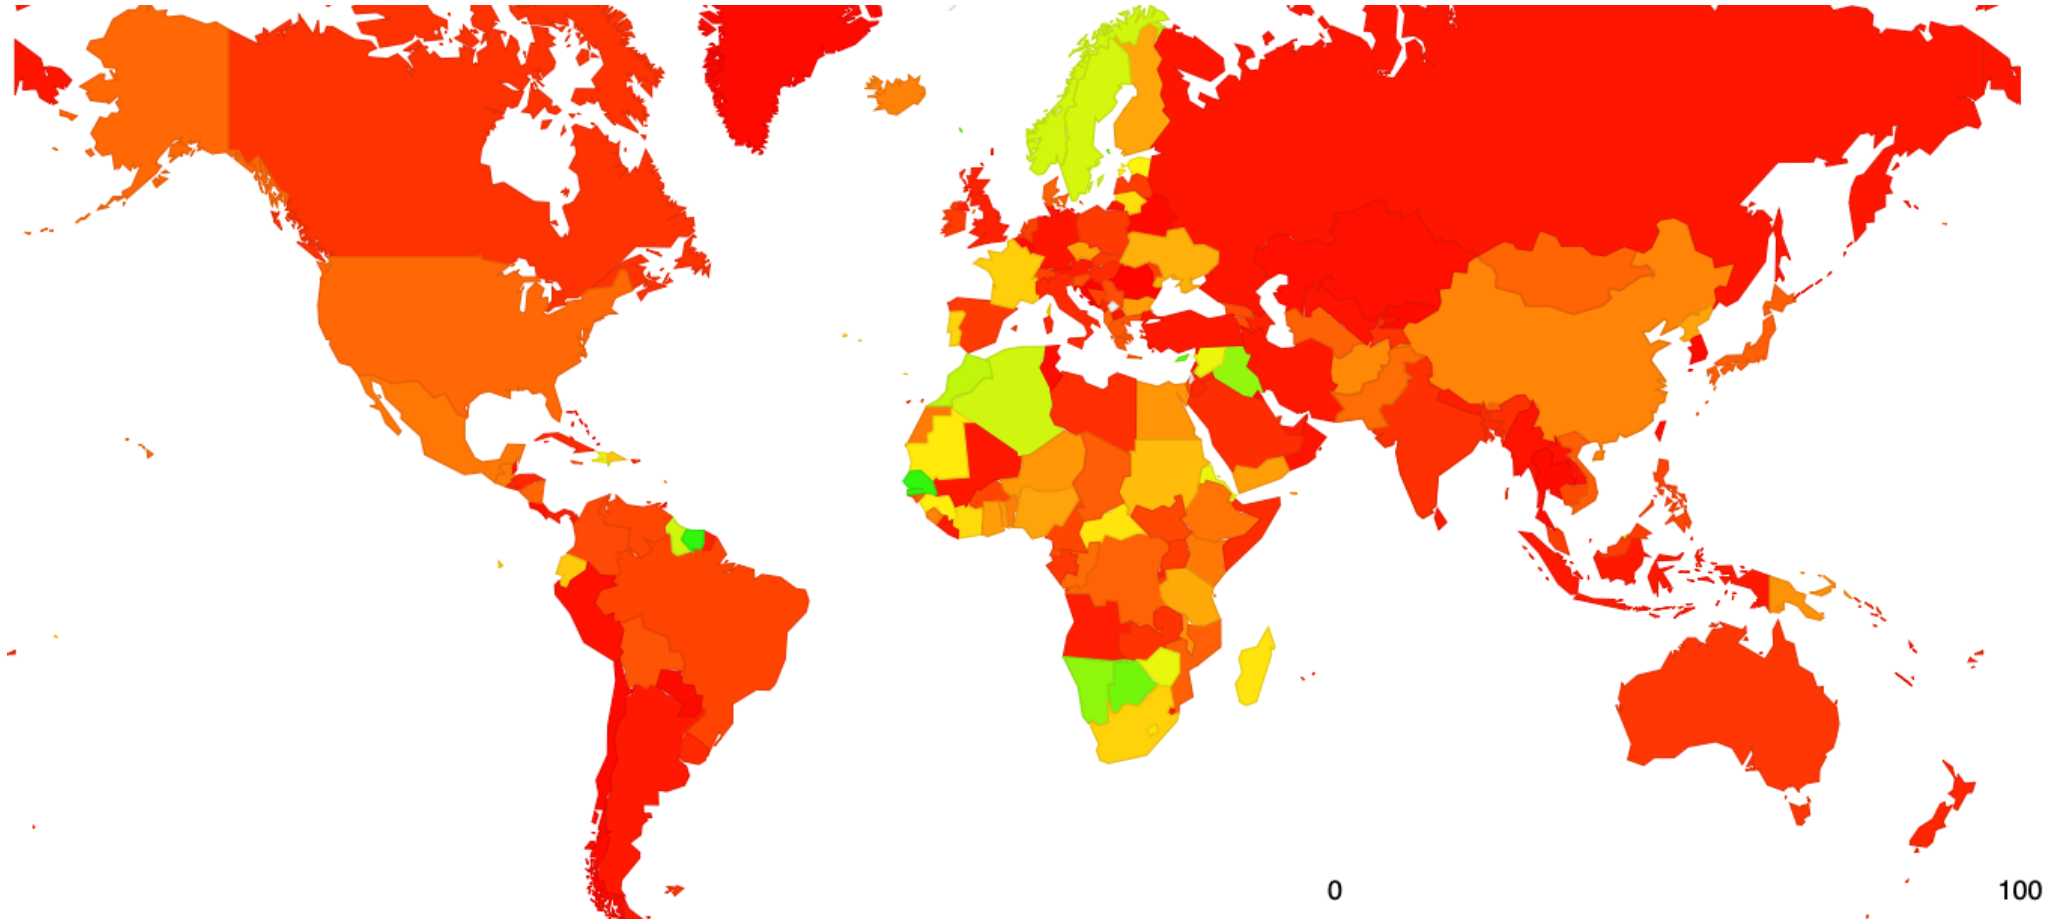

Why?

- This map is a mix of two factors
  - Networks that perform invalid route filtering

RPKI I-ROV Per-Country filtering for AS37100: SEACOM-AS, South Africa (ZA)

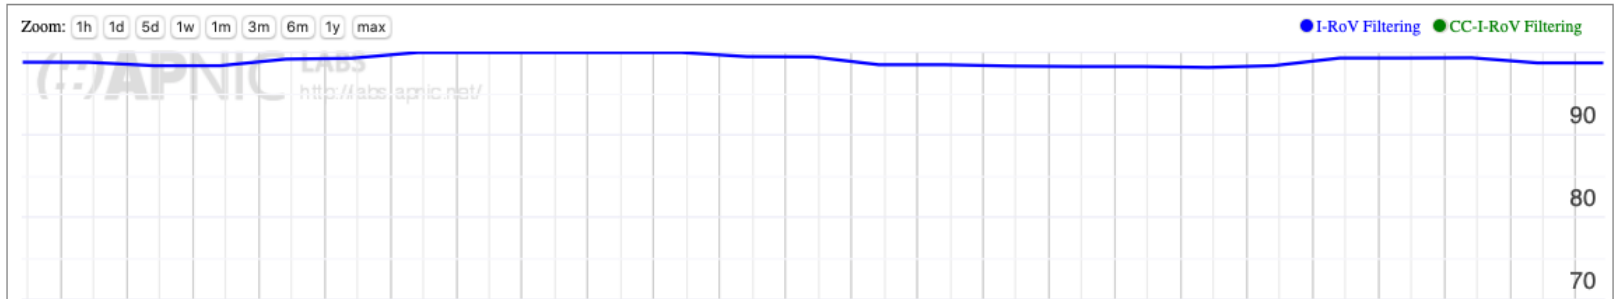

RPKI I-ROV Per-Country filtering for AS7018: ATT-INTERNET4, United States of America (US)

Why?

- This map is a mix of two factors
  - Networks that perform invalid route filtering
  - Network that do not filter themselves by are customers of transit providers who filter
- In either case the objective is achieved, in that end users and their networks are not exposed to invalid route objects

Questions?

- Is it **necessary** for every AS to operate RPKI ROV infrastructure and filter invalid routes?
- If not, what's the minimal set of filtering networks that could provide similar levels of filtering
- What's more important: protecting others from your operational mishaps or protecting yourself from the mishaps of others?