# BGP Security Threats and Challenges

*Geoff Huston AM*

*Chief Scientist, APNIC*

# Why do we keep seeing these headlines?

ars TECHNICA

BIZ & IT TECH SCIENCE POLICY CARS GAMING & CULTURE STORE

*THE ACCIDENTAL LEAK —*

## Google goes down after major BGP mishap routes traffic through China

Google says it doesn't believe leak was malicious despite suspicious appearances.

DAN GOODIN - 11/13/2018, 6:25 PM

Google lost control of several million of its IP addresses for more than an hour on Monday in an event that intermittently made its search and other services unavailable to many users and also caused problems for Spotify and other Google cloud customers. While Google said it had no reason to believe the mishap was a malicious hijacking attempt, the leak appeared suspicious to many, in part because it misdirected traffic to China Telecom, the Chinese government-owned provider that was recently caught improperly routing traffic belonging to a raft of Western carriers though mainland China.

FURTHER READING
Strange snafu misroutes domestic US Internet traffic through China Telecom

The leak started at 21:13 UTC when MainOne Cable Company, a small ISP in Lagos, Nigeria, suddenly updated tables in the Internet's global routing system to improperly declare that its autonomous system 37282 was the proper path to reach 212 IP prefixes belonging to Google. Within minutes, China Telecom improperly accepted the route and announced it worldwide. The move by China Telecom, aka AS4809, in turn caused Russia-based Transtelecom, aka AS20485, and other large service providers to also follow the route.

According to BGPmon on Twitter, the redirections came in five distinct waves over a 74-minute period. The redirected IP ranges transmitted some of Google's most sensitive communications, including the company's corporate WAN infrastructure and the Google VPN. This graphic from regional Internet registry RIPE NCC shows how the domino effect played out over a two-hour span. The image below shows an abbreviated version of those events.

# The Problem

We designed and built a dynamic self-learning (and self-healing) routing system that operates without any centrality of authorisation or control

This is extraordinarily flexible, robust, and scalable, as we have seen

# BGP just scales (so far!)

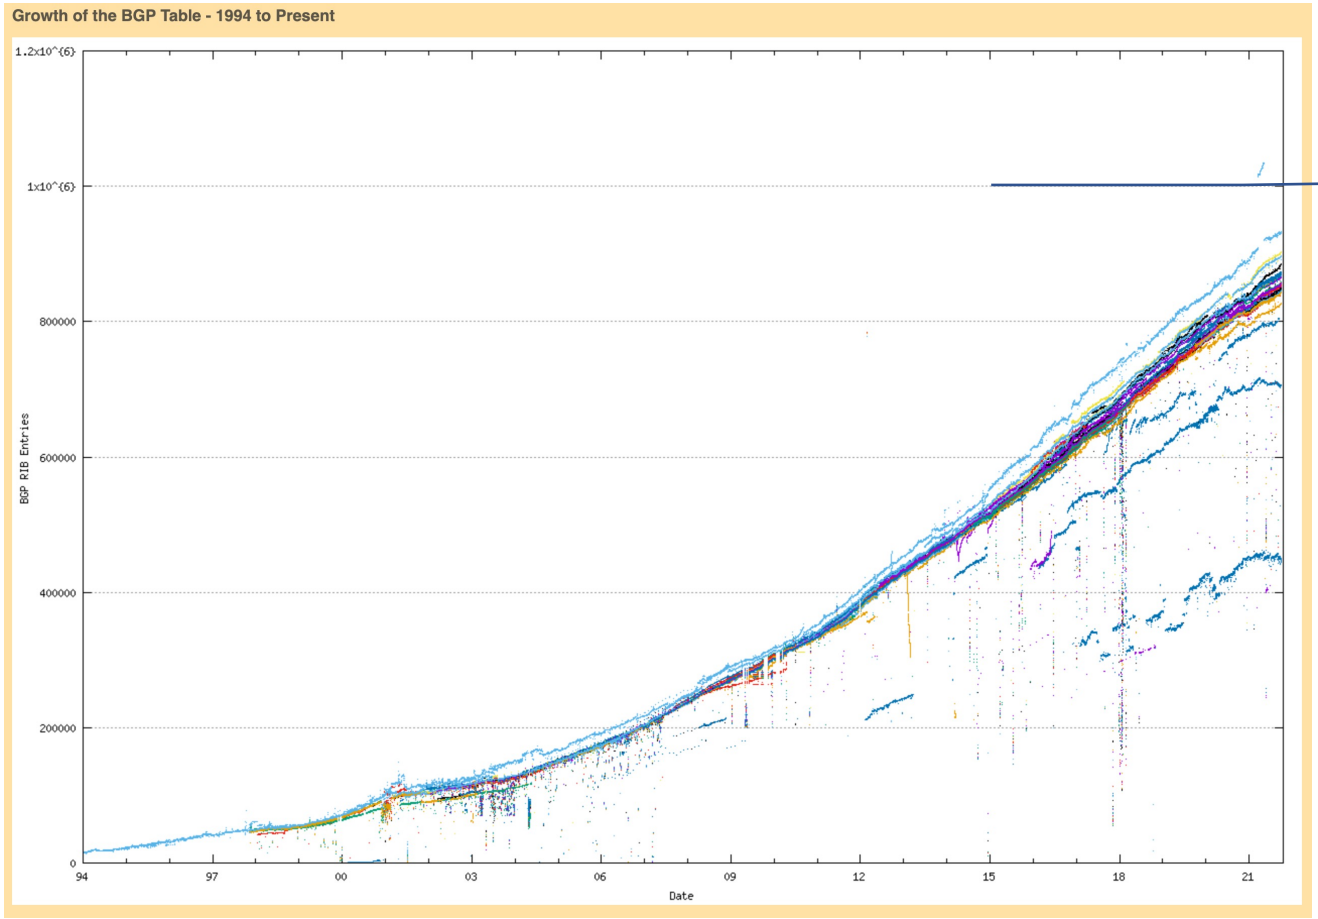


1M BGP Routes

# The Problem

We designed and built a dynamic self-learning (and self-healing) routing system that operates without any centrality of authorisation or control

This is extraordinarily flexible, robust, and scalable, as we have seen

**The problem is that nobody can tell when the routing information being passed inside this routing system is untrue**

# The Threat

Falsifying routing information  can cause:

- Traffic diversion
- Service impersonation
- Denial of service

Which can lead to outcomes of service disruption and potential theft.

Attacks to routing can be highly targeted or very broadly based

# The Threat

Route hijacking is not necessarily an end in itself, but can be a part of a larger attack

- The  April ‘18 MyEtherWallet raid was an attack involving a domain name registrar, the DNS, a susceptible certificate authority and a BGP route injection to work

Sometimes its not a hostile attack but a result of accidental self-harm

- Facebook managed to remove the IP prefixes of its own name servers from BGP on October 5th, causing a global 6-hour outage of the entire Facebook platform

# Counter Measures

- Improving the resilience of host systems cannot mitigate this threat – routing attacks rely on the host performing “normally”
- Improved application design can mitigate some of the impacts of a routing attack
  - Use of TLS provides an end-to-end encrypted session so that traffic diversion cannot reveal session contents
  - TLS also can provide server authentication so that service impersonation is more challenging when TLS is being used
- But the problem of disruption remains
  - And we can’t “fix” this at the ends
  - **We need to address this vulnerability within the routing system itself**

# The Routing Security Goal

Can we devise changes to operational practices, or operational tools or routing technologies that manage the inter-domain routing system that could prevent the propagation of false or artificial routing information across the Internet?

# This is a Very Challenging Goal

- It's a problem as old as the concept of a distributed inter-domain routing system
- Each actor applies local policy constraints on local topology knowledge to guide its local route object propagation decisions
- No single actor has sufficient "whole of system" data to determine the difference between what it should've learned and what it has learned – the routing system has no higher knowledge that it can use to filter incoming routing information to discard "lies"

# BGP has "Tunnel Vision"


What BGP “sees”


What BGP “wants”

# The Ideal

We want the interdomain routing system to advertise the **correct** reachability information for "**legitimately connected"** prefixes at all times

That means that we want to **avoid**:

- propagating reachability for bogus address prefixes
- propagating incorrect paths for reachable prefixes
- blocking paths for legitimately connected prefixes

# The Problem Space

- While we'd like to think we understand the provenance for each and every IP address, that is not exactly the case
- And even if we did, we have no precise knowledge as to which network has the authority to originate a route object for that address
- And even then, we have no exact knowledge of the inter-domain topology of the network
- And even then, we have no clear knowledge of the local policy constraints that are applied to the propagation of reachability and topology information

# The Problem

All of which means that we have no clear model of “truth” to compare to the information flow in the routing system

- Which ASes have a legitimate authority to announce a prefix into the inter-domain routing system?
- Which ASes are interconnected?
- What route policies are applied to each pair-wise AS interconnection?
- What is the “best” route to the prefix destination?

# A Weaker Routing Security Goal

Can we devise changes to operational practices, or operational tools or routing technologies that manage the inter-domain routing system that could resist attempts to inject false or contrived routing information across the Internet?

# What Data would we Like?

- An (impossible) ideal data set is the “reference set” that describes a ‘correct’ route object set that should be visible at any vantage point in the network
  - And access to a set of credentials  that support any such attestation of “correctness”
- As a compromise we could settle for a reference set that describes a ‘stable’ route object set that should be visible at any vantage point in the network

# What we want and don't want

- BGP anomaly detectors and observatories are all well and good, but they have not proved to be all that useful to the operations community
  - They are a bit like smoke alarms – they can't prevent the root cause, but simply alarm **after** it happens
- **What we would like is some form of route acceptance model that can be used as an acceptance filter for incoming route updates**

# So, we've been working on this…

We observed that we needed to improve “verifiable truth” in addressing and routing:

- We wanted to use the technology of *digital signatures* to allow receivers to validate aspects of the routing information being passed to them
- We designed a *Public Key Infrastructure* to bind public/private key pairs to ownership or IP address prefixes and/or Autonomous System Numbers
- We published tools to allow network operators to use this RPKI to generate digital authorizations
- Prefix holders could provide *verifiable “authorities”* for an AS to advertise the prefix into the inter-domain routing space (ROAs)
- Network operators could then filter routing information and discard those routes that do not match these ROAs (ROV Filtering)
- And we hoped that network operators would see the common benefit in adoption this technology

# A Word of Caution

"And we hoped that network operators would see the common benefit in adoption this technology"

- "Common Benefit" is extremely hard to identify in something as large and diverse as the Internet
- The economics of this situation work against it
  - The integrity of common infrastructure is everyone's problem which in turn quickly becomes nobody's problem
- Which implies that just identifying a benefit for all does not imply that each individual network perceives that self interest align to common interest

- **Which means that adoption is likely to be slow and not necessarily going to complete anytime soon**

# How are we doing with RPKI tool adoption?

# Lets look at some RPKI tool Adoption Measurements

Lets look at the metrics of adoption of RPKI

- Number of published ROAS
- Number of Route Objects that are validated by RAOS
- Number of clients that perform a regular fetch of RPKI material
- Adoption of Route Filtering

# Populating the RPKI - ROAs

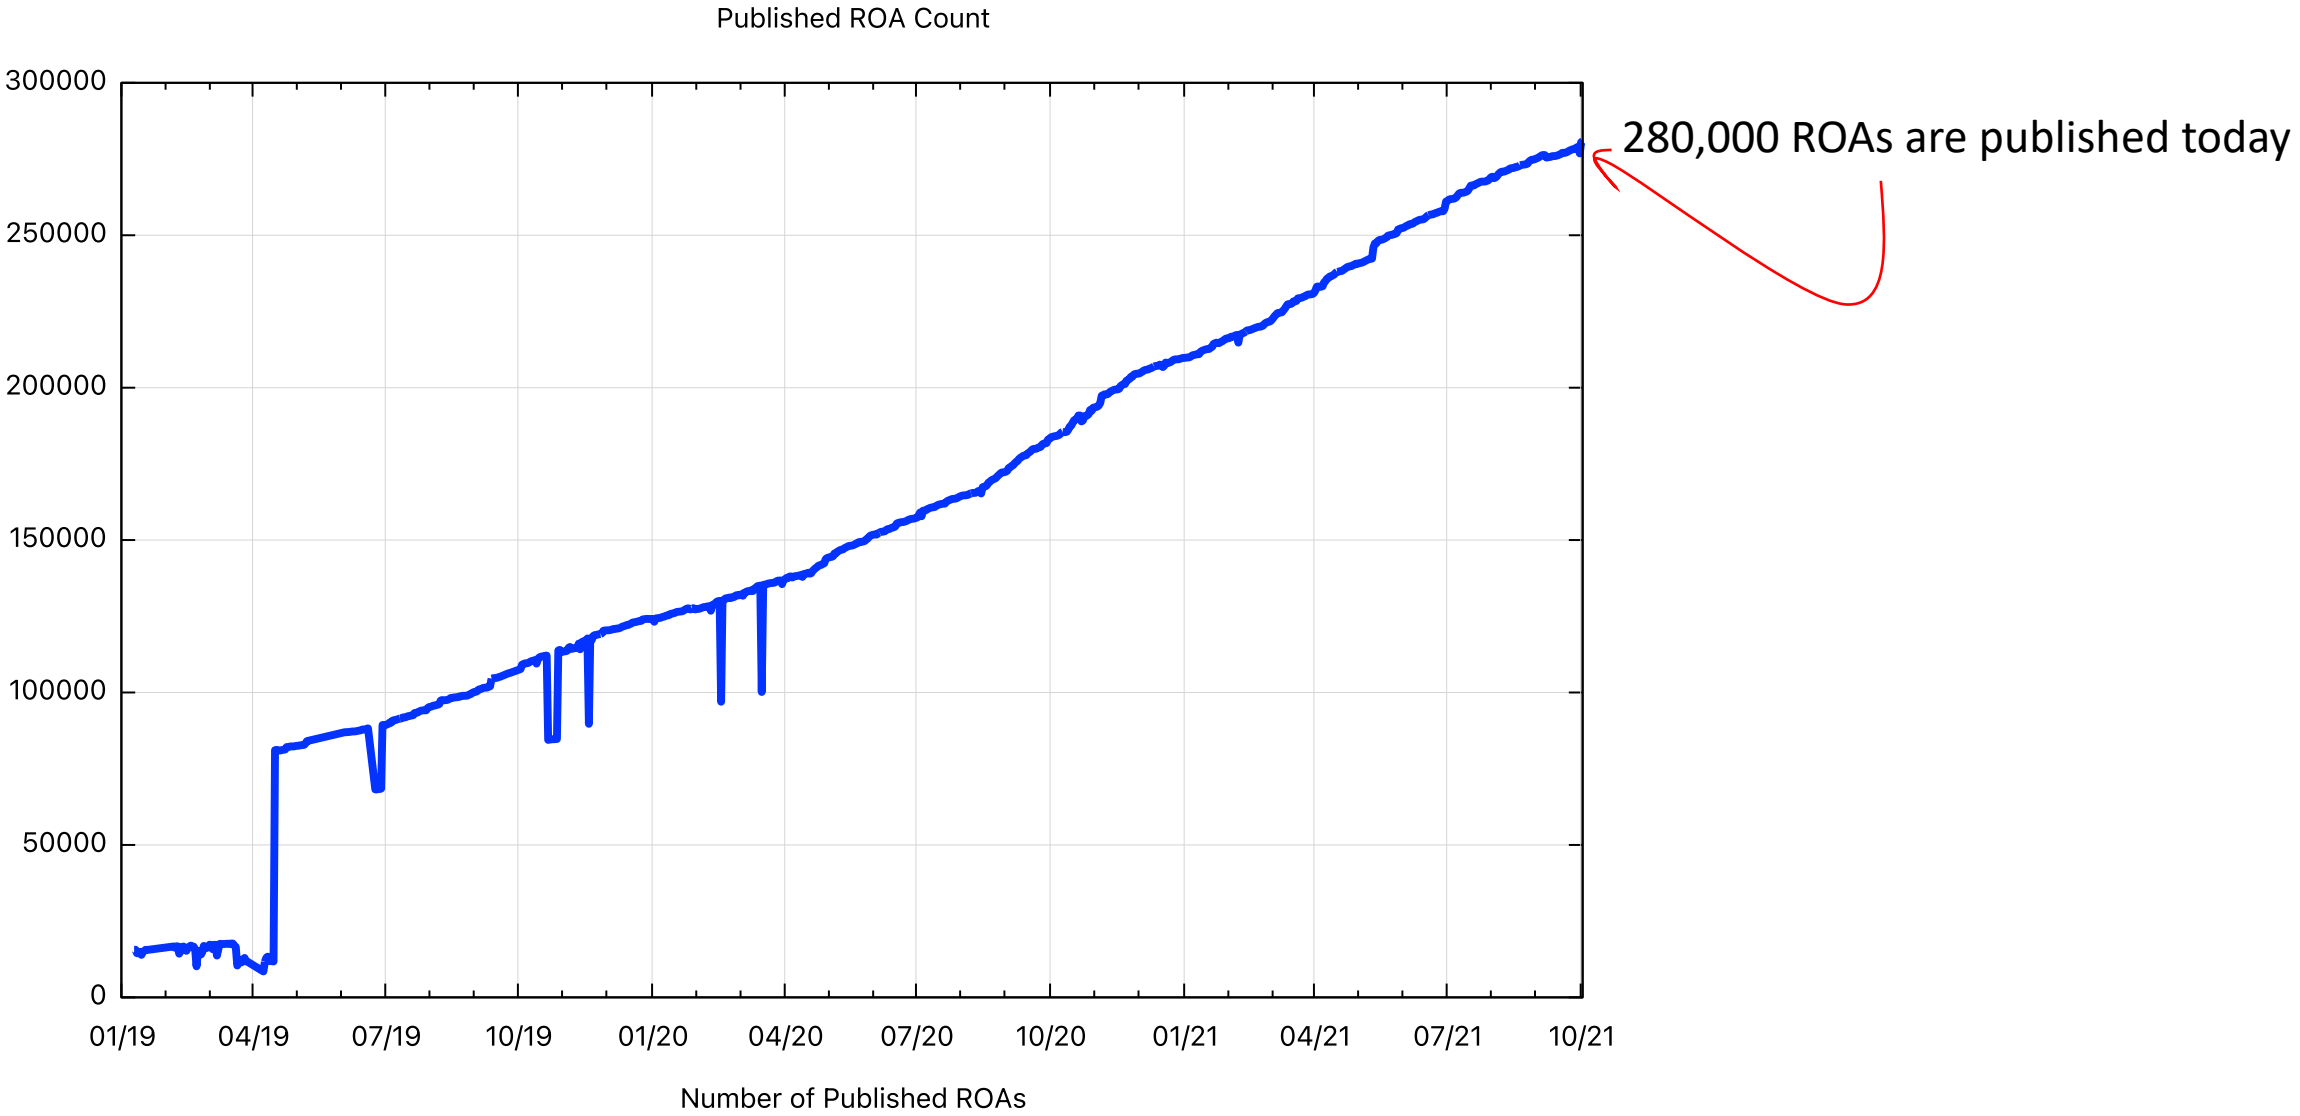

# Applying ROAs to Routes

Display: **ROAs** (Advertised ROA-Valid Route Advertisements), **IPv4**, **Count**

IPv4 Routes that are ”covered” by a ROA

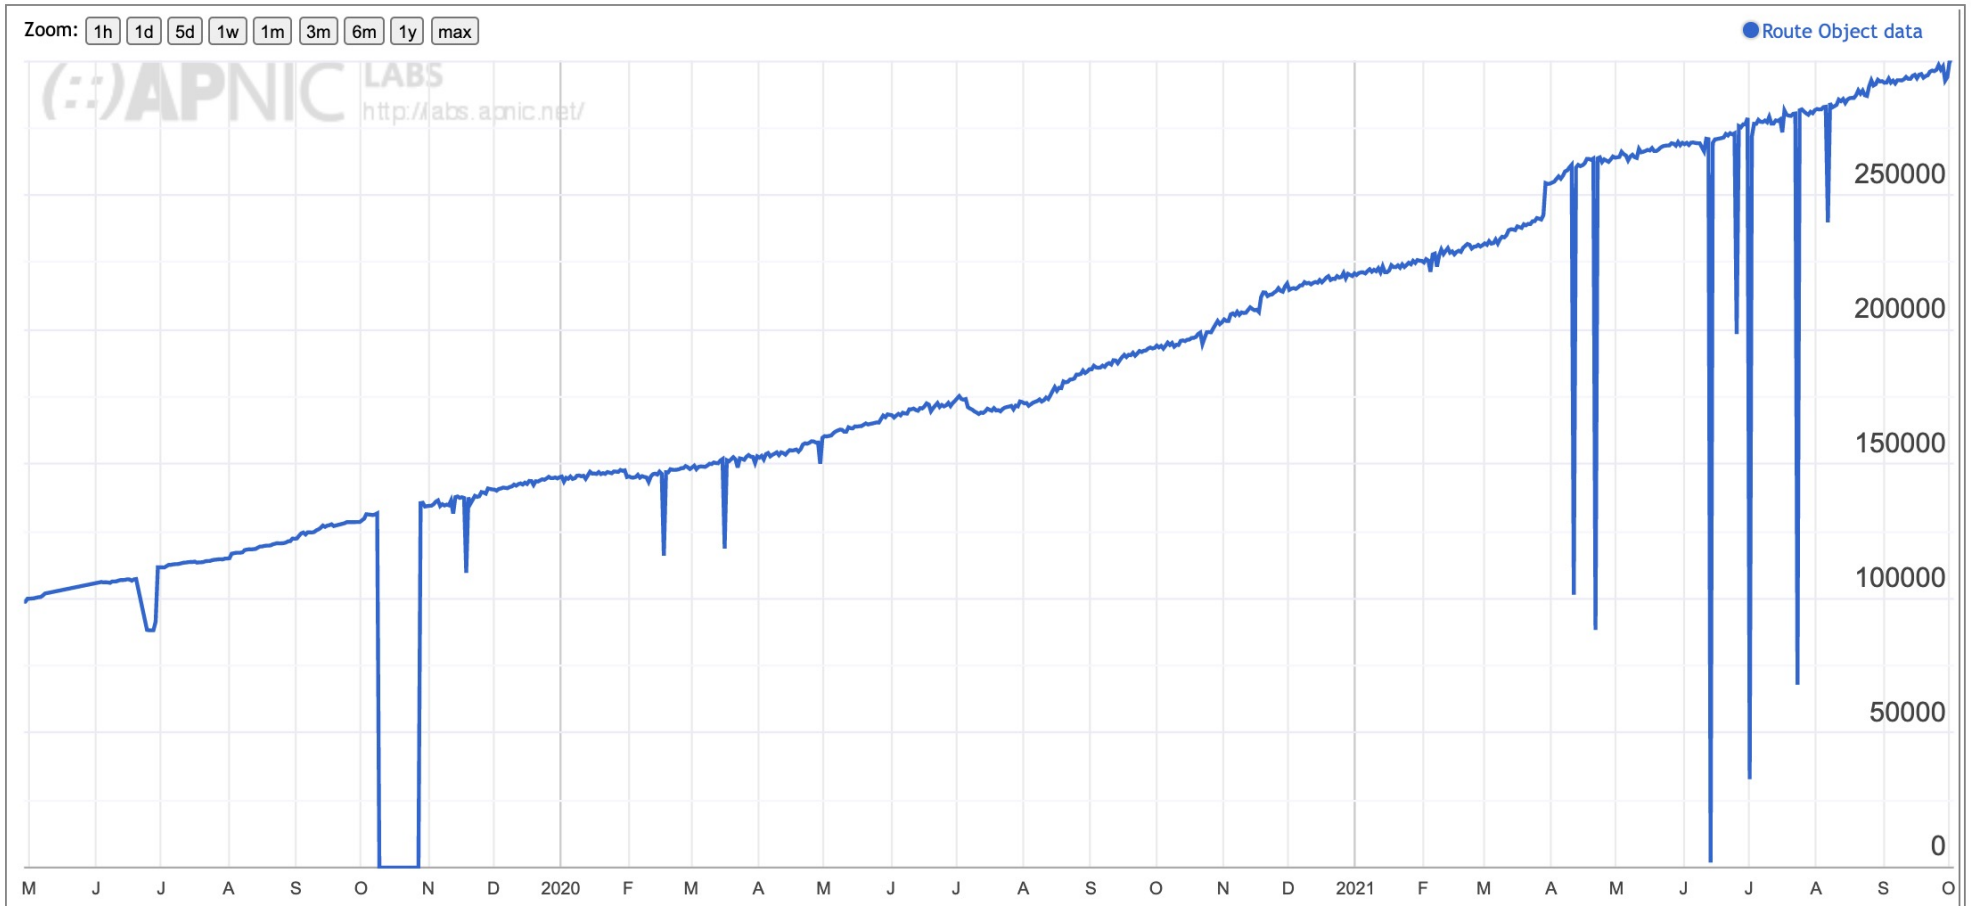

# Applying ROAs to Routes

IPv6 Routes that are ”covered” by a ROA

Display: **ROAs** (Advertised ROA-Valid Route Advertisements), **IPv6**, **Count**

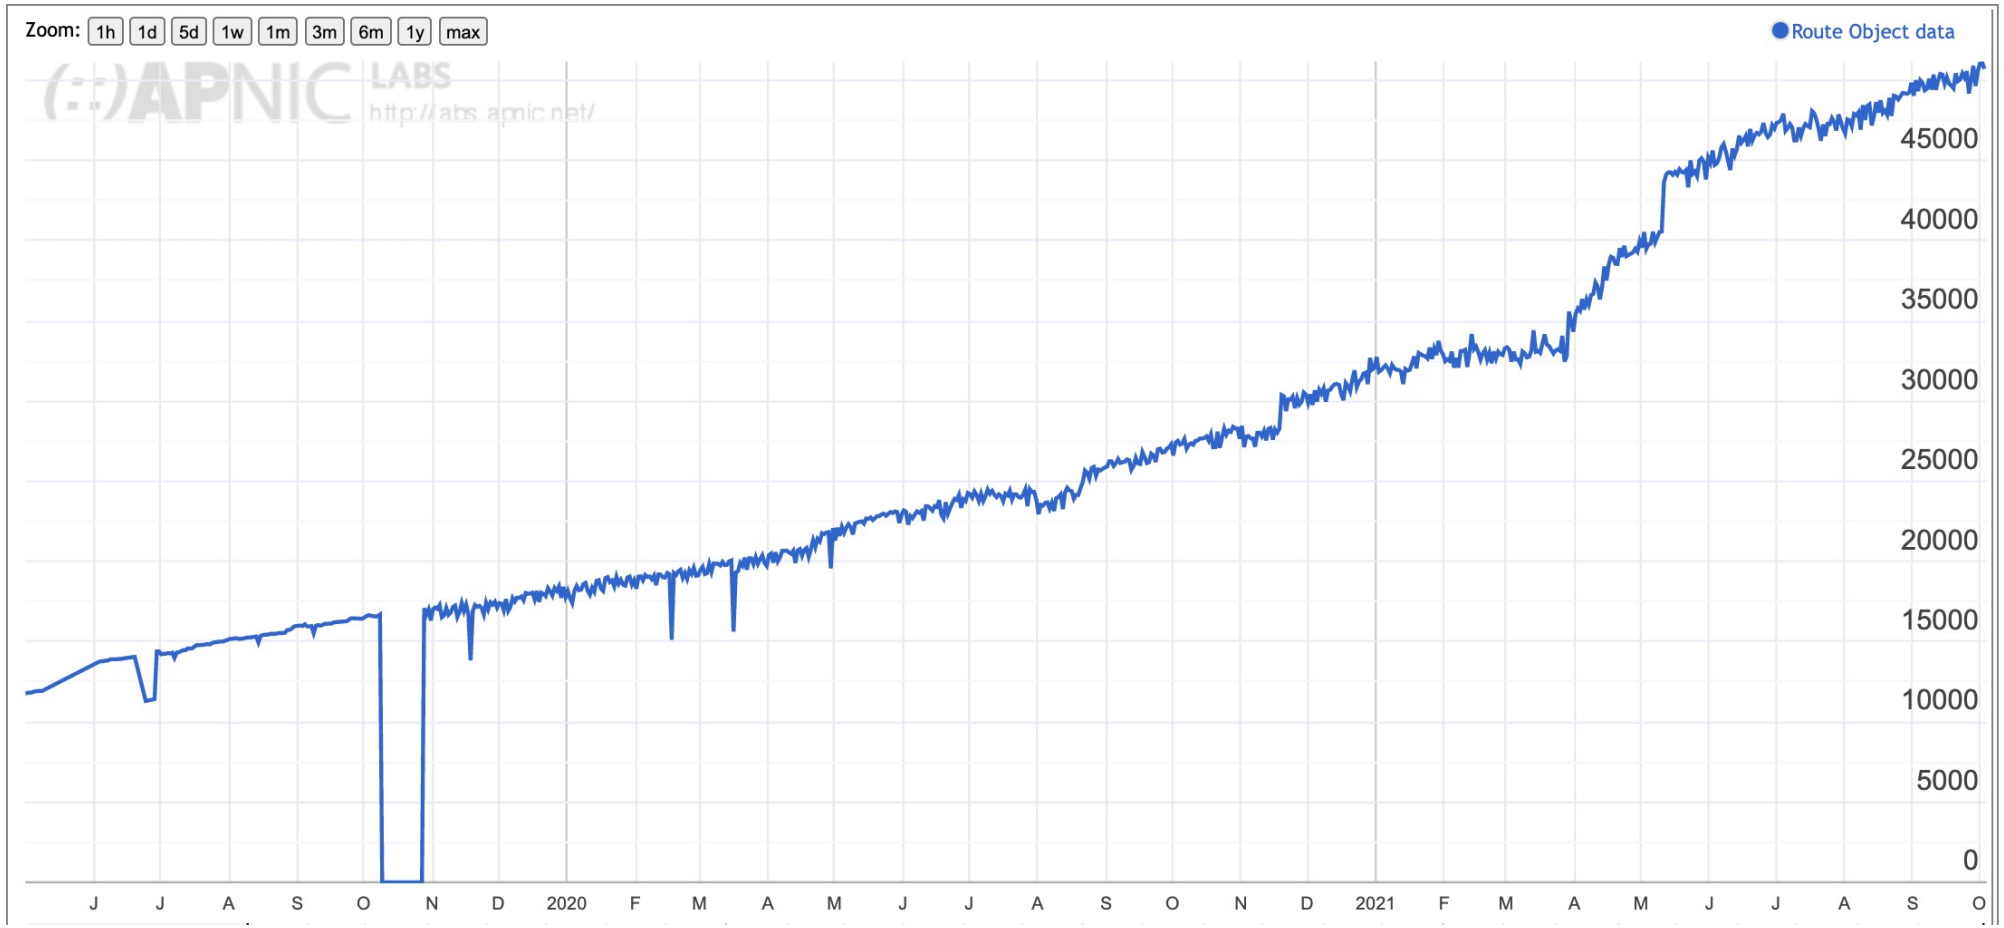

# RPKI Clients

- How many clients regularly maintain a local cache of the entire RPKI product set ?

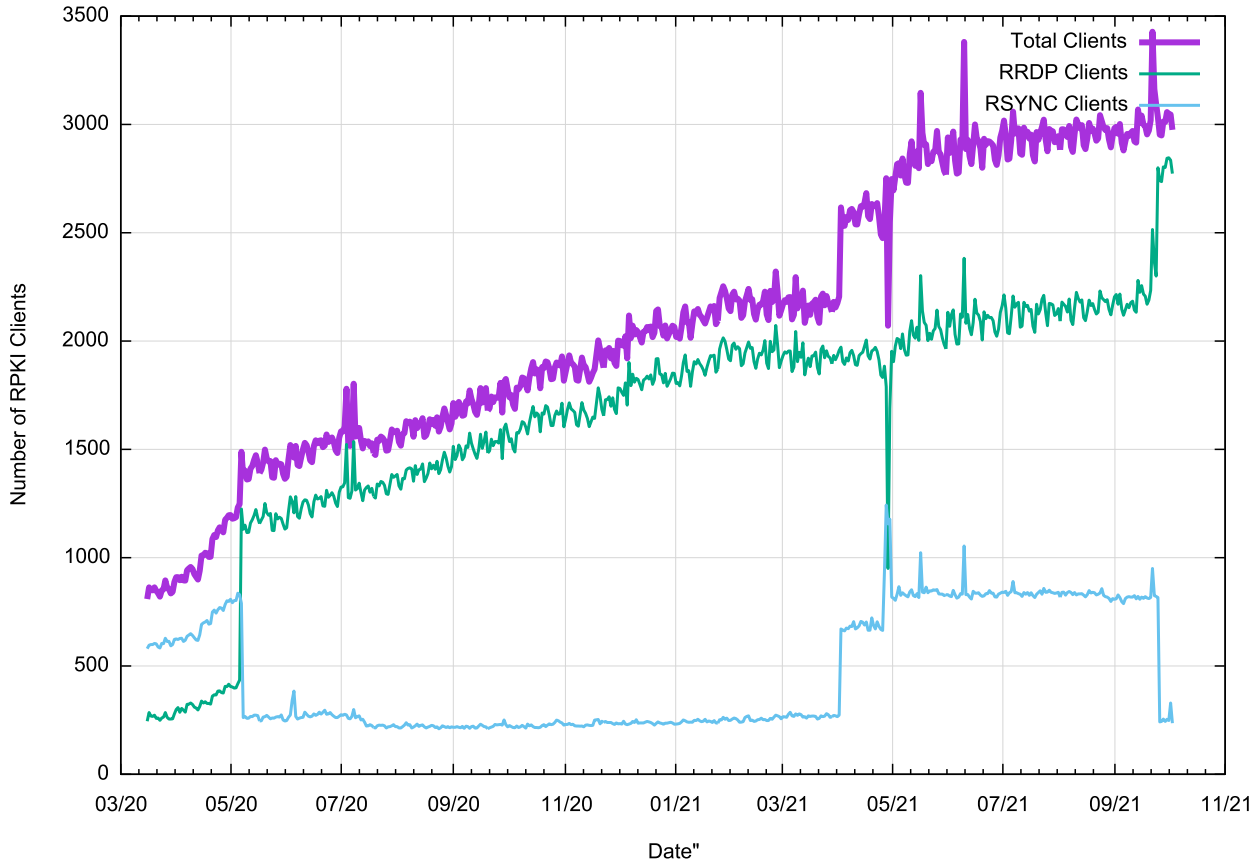

# Filtering RoV routes

What proportion of users are behind networks that filter ROV-invalid routes?

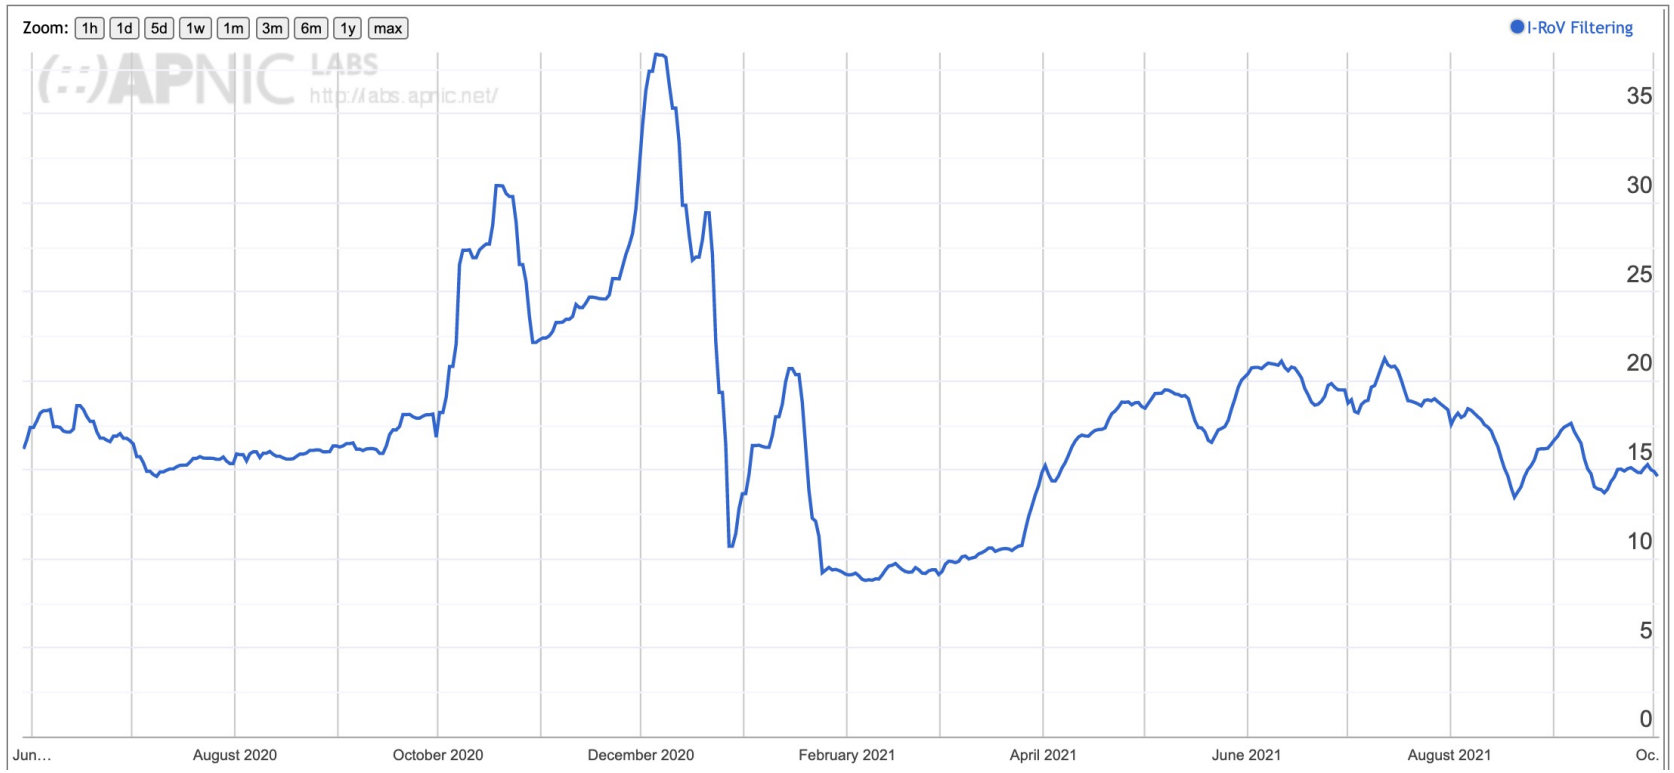

# RoV Filtering

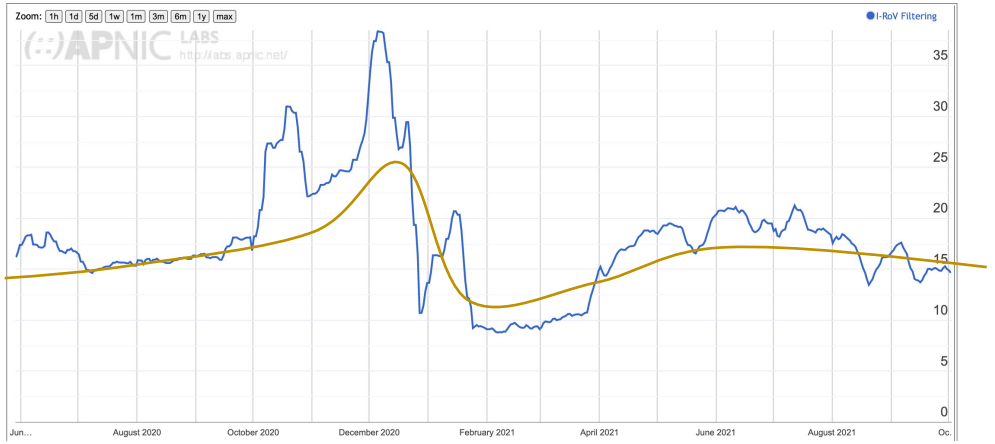


- This is an unexpected result
- Only some 15% of users cannot reach a prefix when it is advertised using a ROV-invalid prefix
- And this has been constant measurement for the past ~12 months
- It seems that few edge networks have been performing ROV dropping
- And similarly few transit networks have taken up ROV-dropping
- And there is sufficient diversity in the inter-AS topology that even if some paths are filtered out, alternate transit paths sill provide access

# Where is RoV Filtering?

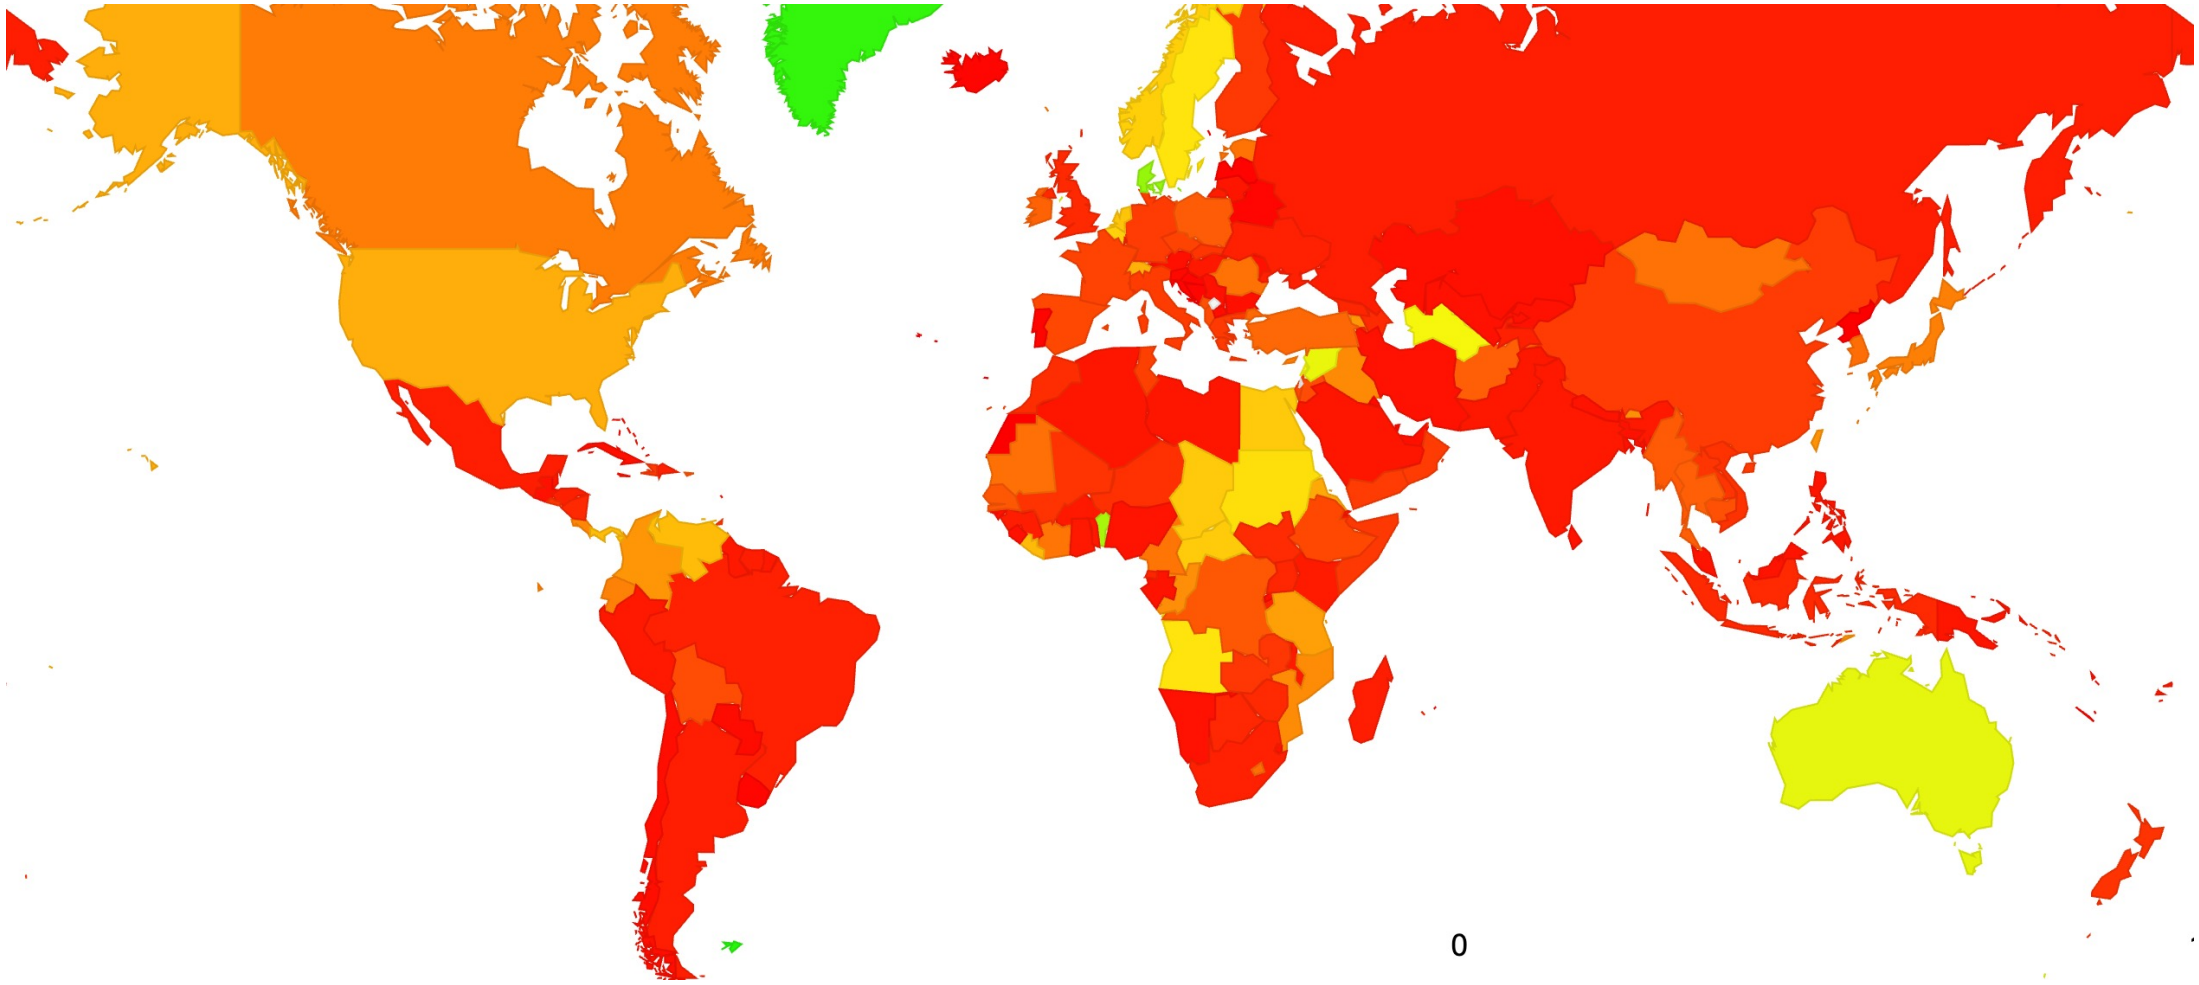

# Where is RoV Filtering?

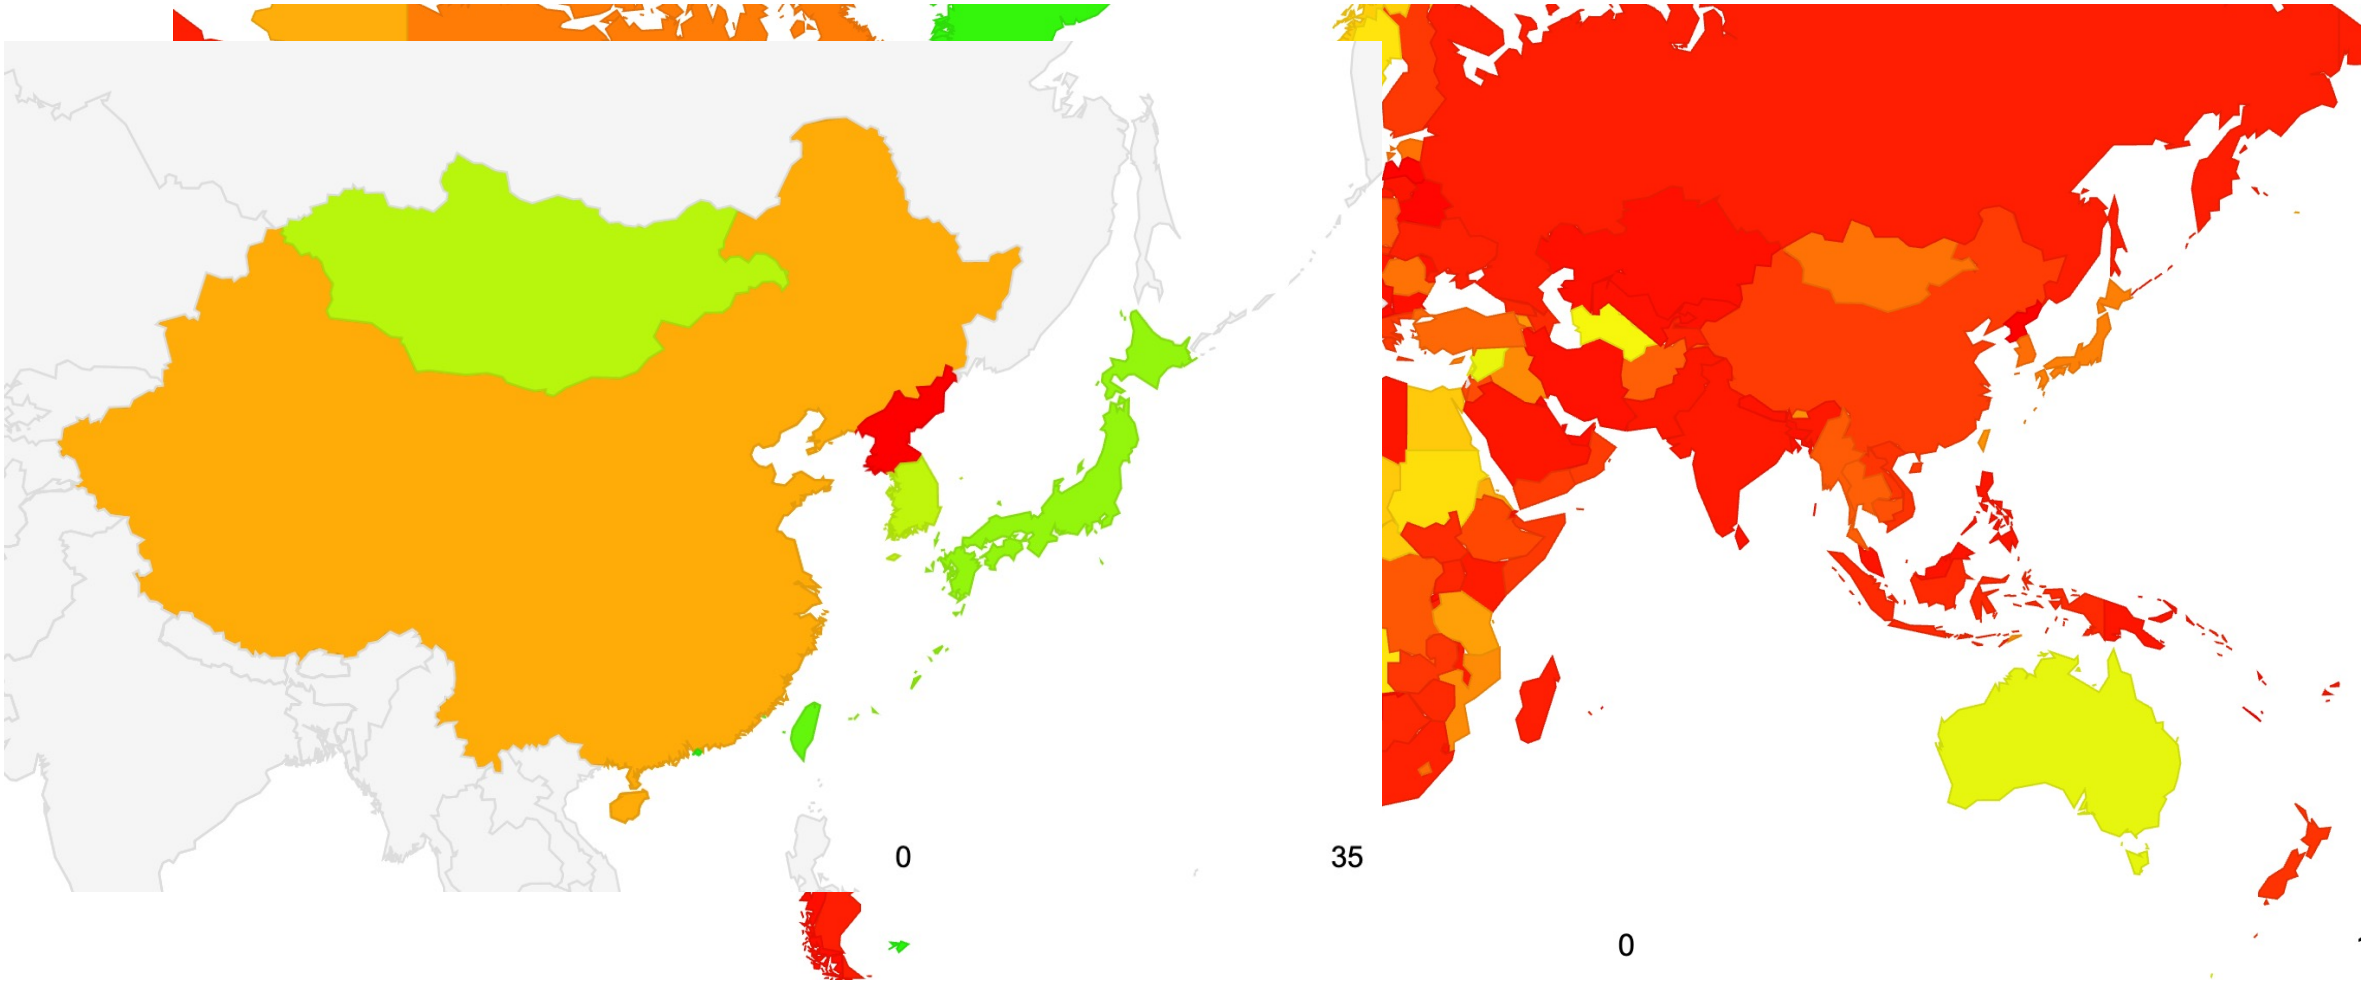

# But - all this is not enough!

This is only one part of the overall effort to secure the routing system

It is still possible to mount a routing attack by reproducing the originating AS with a faked AS Path

To complement Routing Origination Validation we need to provide a means to validate the information in the AS Path

# How about ASPA?

- AS Provider Attestations allows a network to list those networks that act as a transit provider to this network
- Its being developed in the IETF at the moment, but it looks quite promising:
  - It supports partial deployment models
  - It's a "light weight" approach compared to BGPSEC
  - It shifts the burden from the defender to the attacker
- It's likely to be entering trials and tests in the coming months

# Other Activity

- The BGP Code of Practice – ISOC's MANRS
  - Its hard to be convinced that folk will listen when nobody has been listening for the BCP38 source address filtering for more twenty years!
- BGPSEC
  - A heavyweight BGP implementation of AS Path signing developed in the IETF
  - Its hard to see this going anywhere at all! I think its DOA.
- Blockchain
  - Its currently fashionable
  - But it doesn't really address the core routing security issue
    - Its not the ledger / registry model that's the issue here – the challenge is how to integrate this information into the BGP protocol, not the derivation of trust and authority in the overall address architecture

# What can you do right now?

# What can you do right now?

1. Sign your IP address holdings

# What can you do right now?

1. Sign your IP address holdings
2. Generate ROAs to describe your announcements

# Use of ROAS in Taiwan

## Use of Route Object Validation for Taiwan (TW)

Display: **ROAs** (Advertised ROA-Valid Route Advertisements), **Total** (IPv4 + IPv6), **Percent** (of Total)

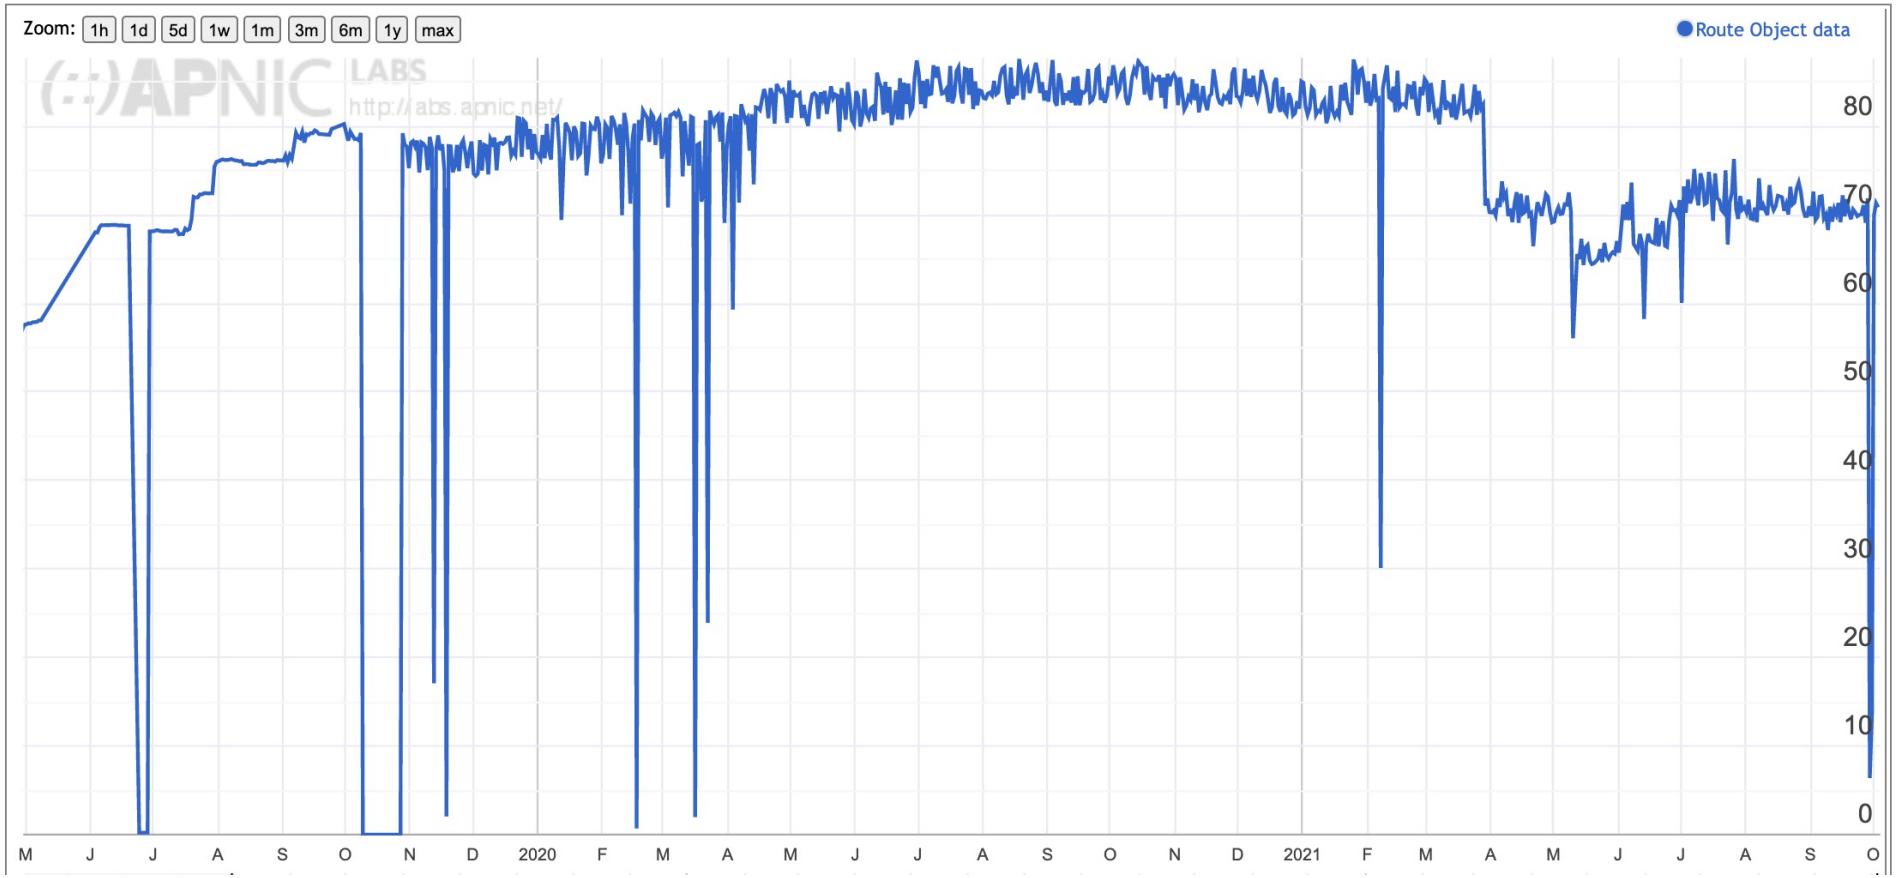

# What can you do right now?

✓ 1. Sign your IP address holdings

✓ 2. Generate ROAs to describe your announcements

# What can you do right now?

✓ 1. Sign your IP address holdings
✓ 2. Generate ROAs to describe your announcements
3. Filter invalid routes from BGP updates

# Filtering RoV routes - Taiwan

What proportion of users are behind networks that filter ROV-invalid routes?

**Use of RPKI Validation for Taiwan (TW)**

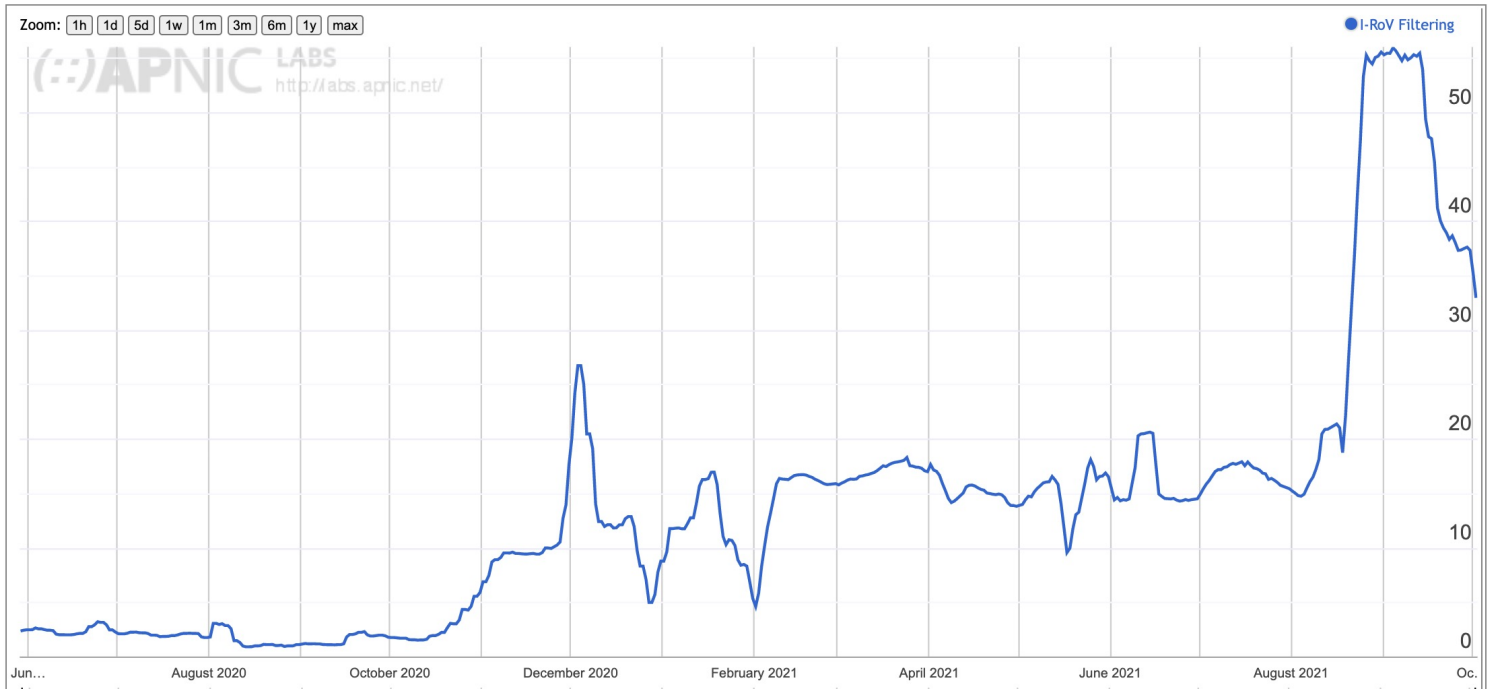

# Filtering RoV routes - Taiwan

What proportion of users are behind networks that filter ROV-invalid routes?

**Use of RPKI Validation for Taiwan (TW)**

| ASN | AS Name | RPKI Validates | Samples |
|---|---|---|---|
| AS3462 | HINET Data Communication Business Group | 33.16% | 61,459 |
| AS17421 | EMOME-NET Mobile Business Group | 53.21% | 43,653 |
| AS9674 | FET-TW Far EastTone Telecommunication Co., Ltd. | 61.89% | 28,343 |
| AS24158 | TAIWANMOBILE-AS Taiwan Mobile Co., Ltd. | 5.02% | 27,389 |
| AS24157 | VIBO-NET-AS Taiwan Star Telecom Corporation Limited.Former Vibo Telecom Inc. | 0.94% | 9,056 |
| AS131591 | AMBIT-AS-TW Ambit Microsystem Corporation | 1.41% | 4,898 |
| AS9416 | MULTIMEDIA-AS-AP Hoshin Multimedia Center Inc. | 1.01% | 4,073 |
| AS9924 | TFN-TW Taiwan Fixed Network, Telco and Network Service Provider. | 4.10% | 3,979 |
| AS131596 | TBCOM-NET TBC | 1.24% | 2,095 |
| AS38841 | KBRO-AS-TW kbro CO. Ltd. | 13.77% | 1,910 |

# What can you do?

✓ 1. Sign your IP address holdings

✓ 2. Generate ROAs to describe your announcements

✗ 3. Filter invalids from BGP updates

Could do better!

# Thank You!