Internet routing security, a distributed database problem

Job Snijders

job@bsd.nl

What is routing security?

Security generally refers to protection from **potential harm** or **unwanted change**.

It is a complex concept that relates to the stability and resilience of Internet networks and societies.

20 years ago

My first experience with routing security

20 years ago

My first experience with routing security

a German person on the phone

yelling at me

20 years ago

My first experience with routing security

a German person on the phone

yelling at me

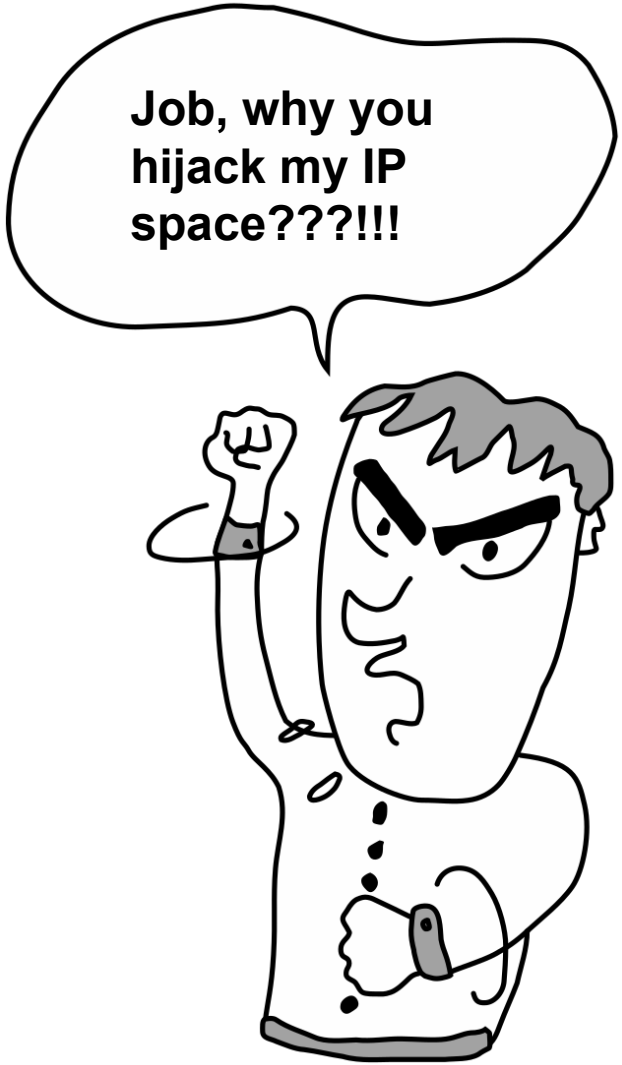

What happened?

A core problem: keys are very close to each other!

Keys?

- Keys on the keyboard
- IP network addresses
- Autonomous System Numbers

It’s all the same

What to do about it?

How to grow the Internet? How to scale the global routing system?

An approach: safety through isolation

https://commons.wikimedia.org/wiki/File:Great_Wall_of_China_around_1900.jpg

But, there is endless change to the shape of the isolated resources

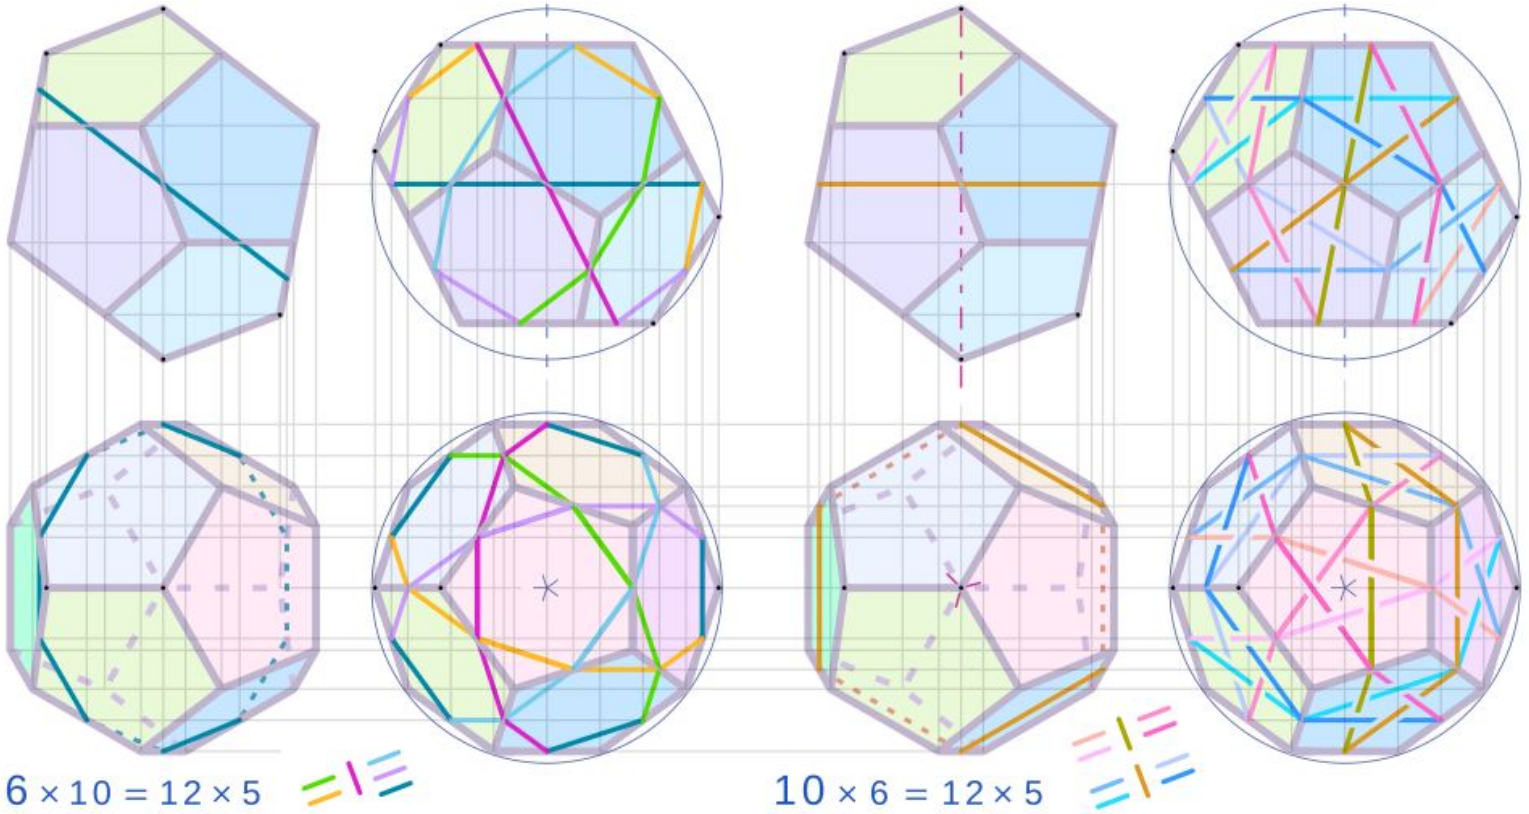

https://commons.wikimedia.org/wiki/File:4_pairs_of_views_of_the_regular_equatorial_cross_sections_of_a_Platonic_dodecahedron.svg

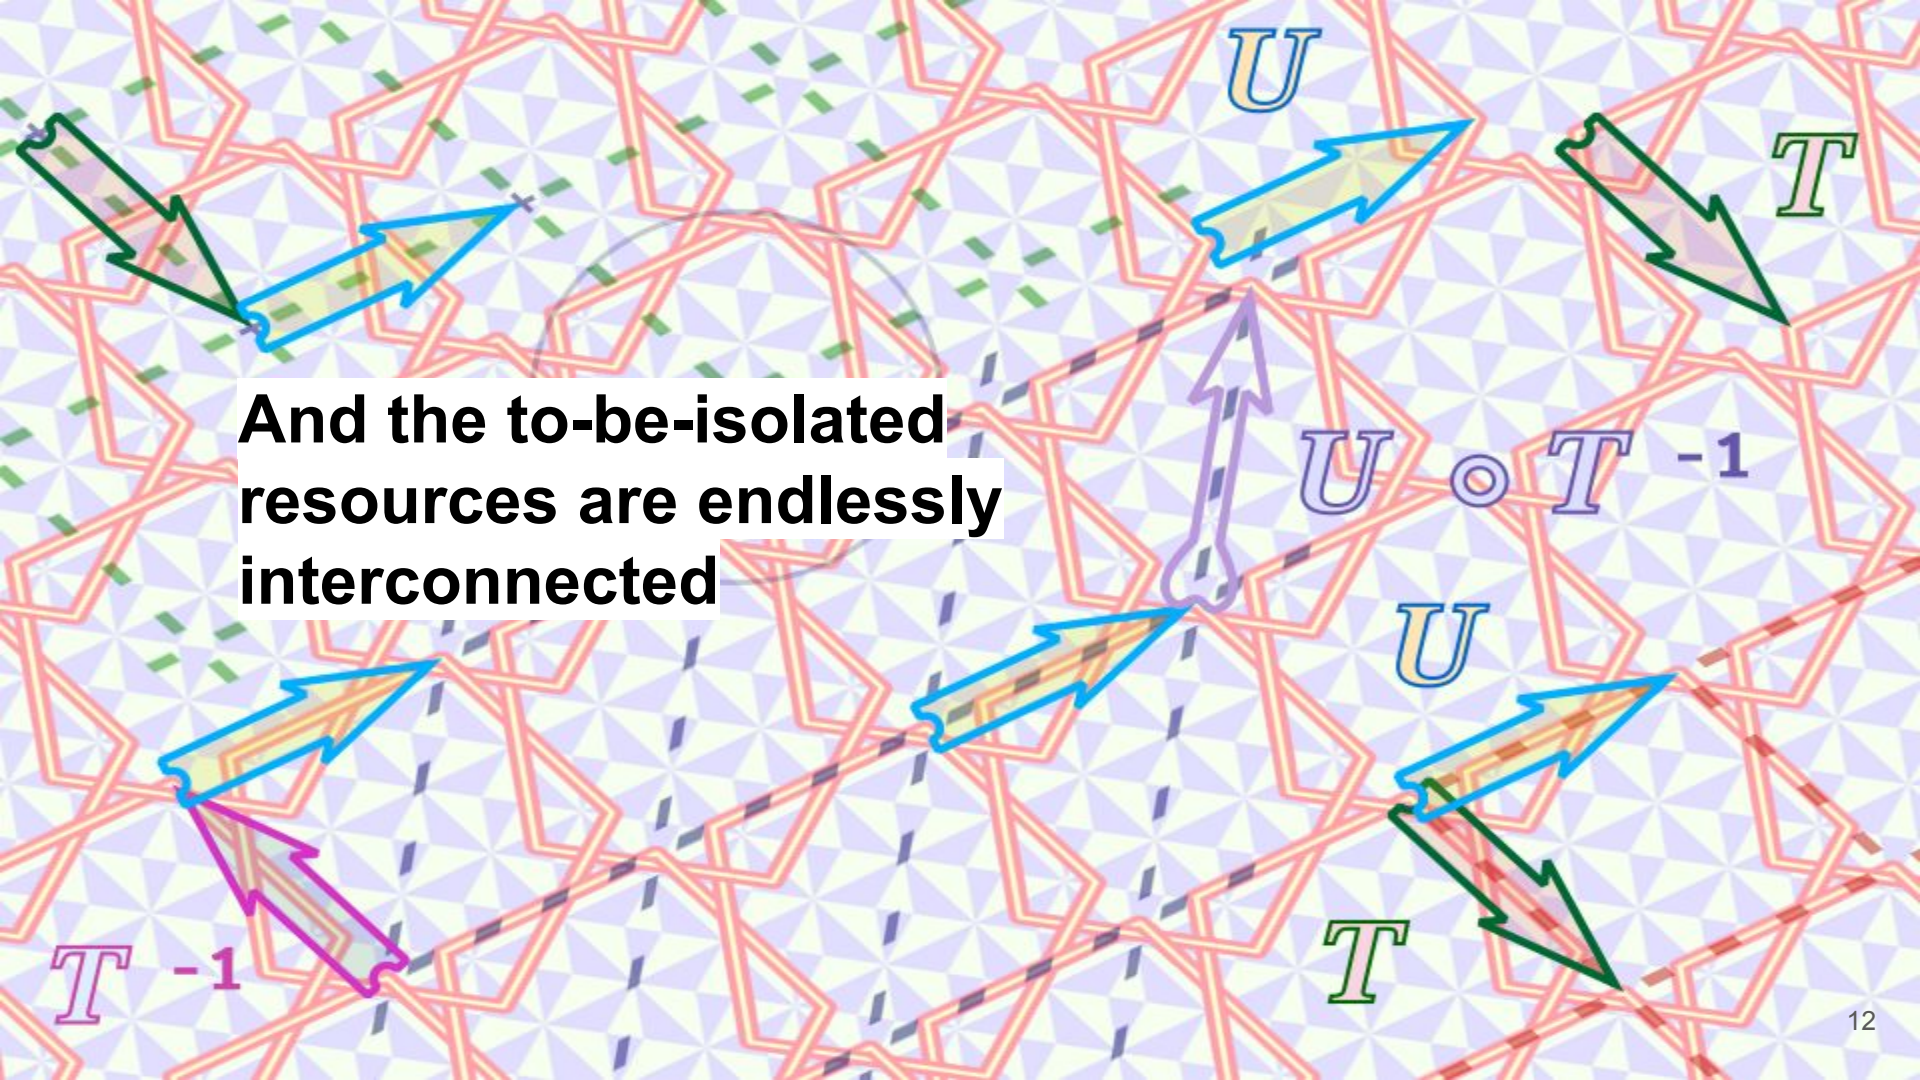

And the to-be-isolated resources are endlessly interconnected

A core problem: what's where? who operates what?

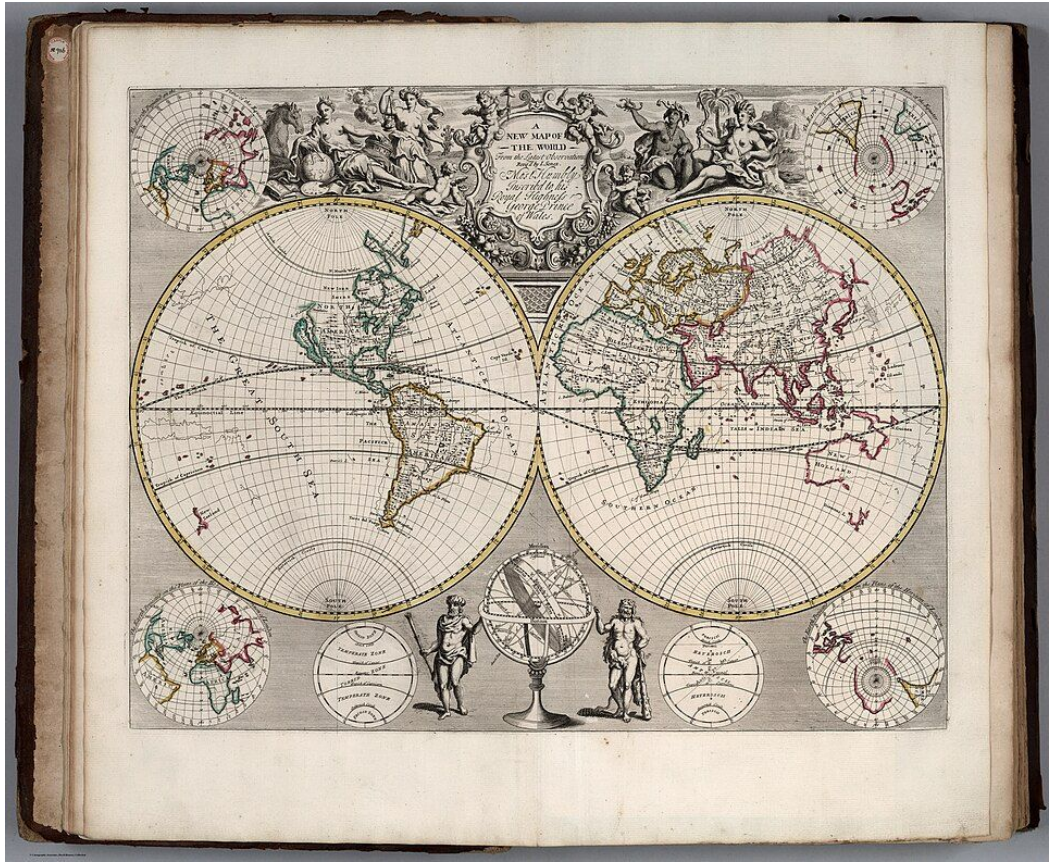

An approach: hierarchical delegation of authority

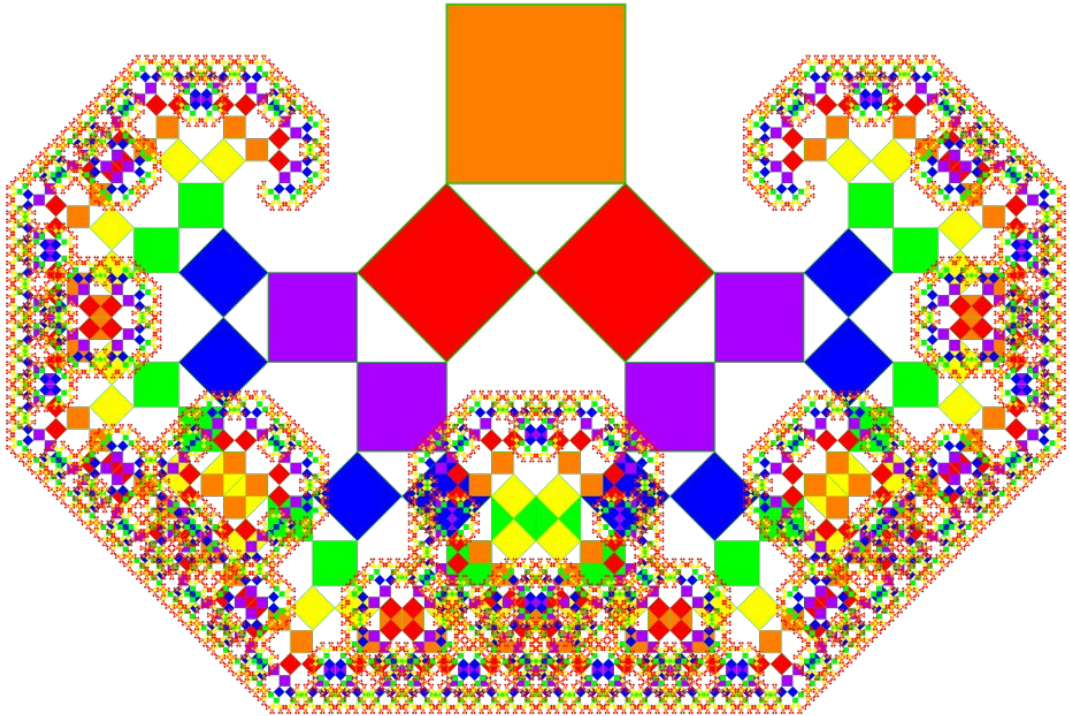

https://commons.wikimedia.org/wiki/File:Pythagoras_tree_1_1_13_Summer.svg

An approach: a PKI, *the RPKI*

So what was the routing security plan?

*A global distributed database for hierarchical delegations of authority to **isolate and interconnect** resources.*

War story: RPKI is working as intended

Job Snijders
Principal Software Engineer,
Fastly

November 05, 2024

CDN & Delivery | Security

Read the full text here:
https://www.fastly.com/blog/war-story-rpki-is-working-as-intended

To be very forward, this really is a story about something that turned out to be no problem at all. But sometimes boring stories deserve to be told. To provide context for this one, we have to go back to February 2008. Back then - through no fault of their own - one of the world's most popular video-sharing platforms suffered a disastrous multi-hour outage, interrupting millions of video viewings. The impact was so significant that even mainstream media reported extensively on what


SOCIETY WITH FULL RPKI DEPLOYMENT
imgflip.com

A perspective on the timeline

2007-2012 – IETF community writes & publishes RPKI RFCs

2012-2017 – Global network community largely unaware of RPKI

2019-2020 – Operators deploy Route Origin Validation at global scale

2020-2024 – Lots of bug fixing in RPKI standards & implementations

2024- … – Optimizing for efficiency, auditability, *performance*

Where are we now?

- *Most* of the Internet routing table covered by RPKI ROAs
- *Most* of the traffic flows to RPKI ROA covered destinations
- Filtering (rejecting) invalid route announcements happens where it matters: invalids don’t propagate far
- There currently are ~ 5000 RPKI-ROV invalid routes out there, *these are considered “background noise”*
- *Keep in mind:* Routing security is more than RPKI, RFC 9234 “BGP OPEN Roles” is a great example of this

The RPKI is a distributed database

The RPKI is a distributed database

The RPKI is ***not*** a routing protocol

This means… distributed database challenges!

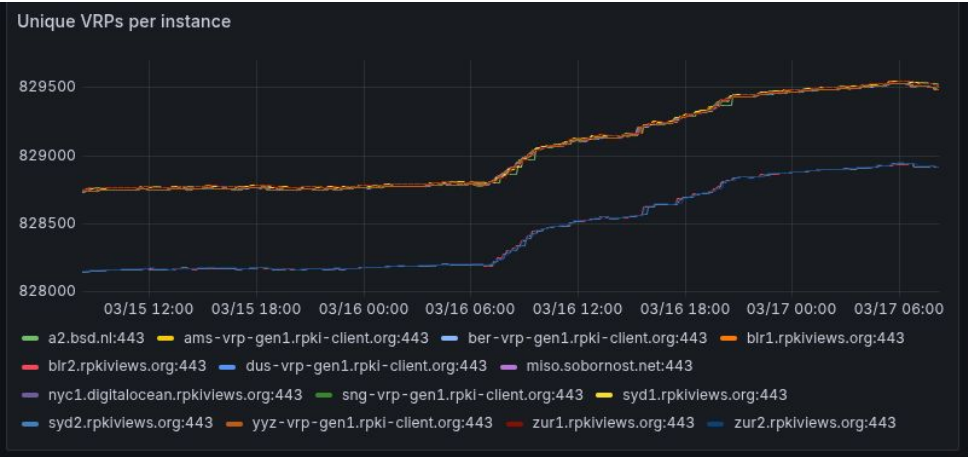

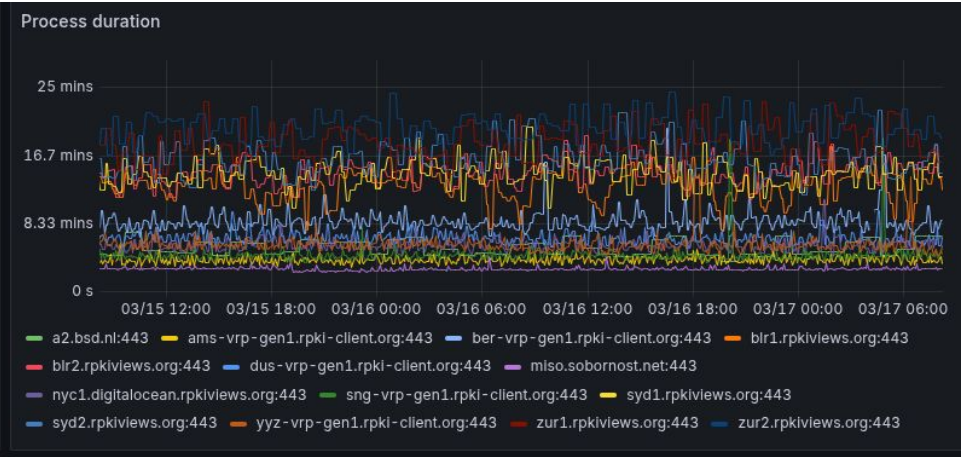

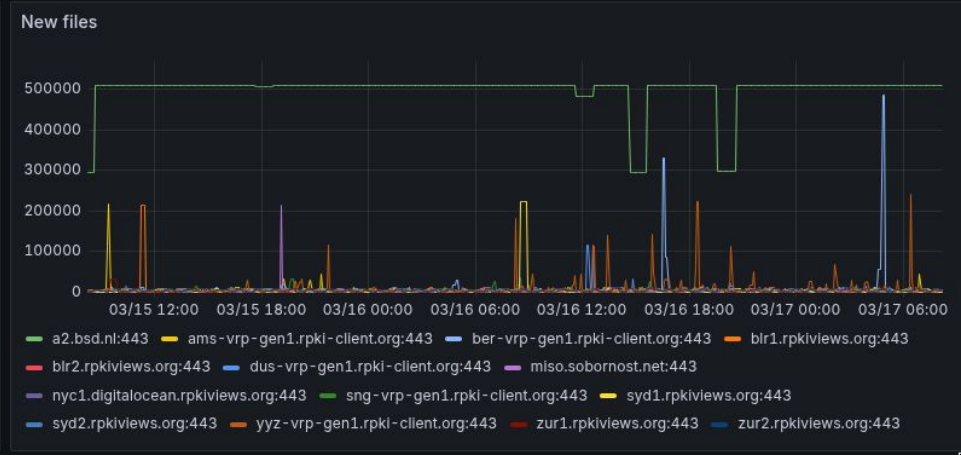

Distributed database systems challenges

data consistency, partition tolerance, network latency, replication, security and authorization, scaling query processing

Some RPKI statistics

- Every second, two new RPKI objects appear somewhere on the planet
- Every two minutes, someone somewhere adds/deletes a Route Origin Authorization (ROA)
- Every 26 minutes, someone somewhere adds/deletes an ASPA
- Every 24 hours, 33% of the database changes
  (180,000 out of 520,000 objects are ‘upserted’)
- Total database size currently is around 1 gigabyte (compressed 500 MB)
- The RPKI currently has ~ 50,000 participating Certification Authorities
- The RPKI currently might have ~ 5,000 relying parties?
  (caveat: external visibility is murky!!!)

How do I know this?


Radio Shack
HI OUT
SPEED
LO IN
RPKIViews.org
SELECTOR
OFF
TAPE
INPUT
STEREO
ADD
MONO
LISTENING VOLUME
FAST WINDING
PLAY OR RECORD
STOP
RECORDING VOLUME
RADIO/PHONO
MICROPHONE
NORMAL

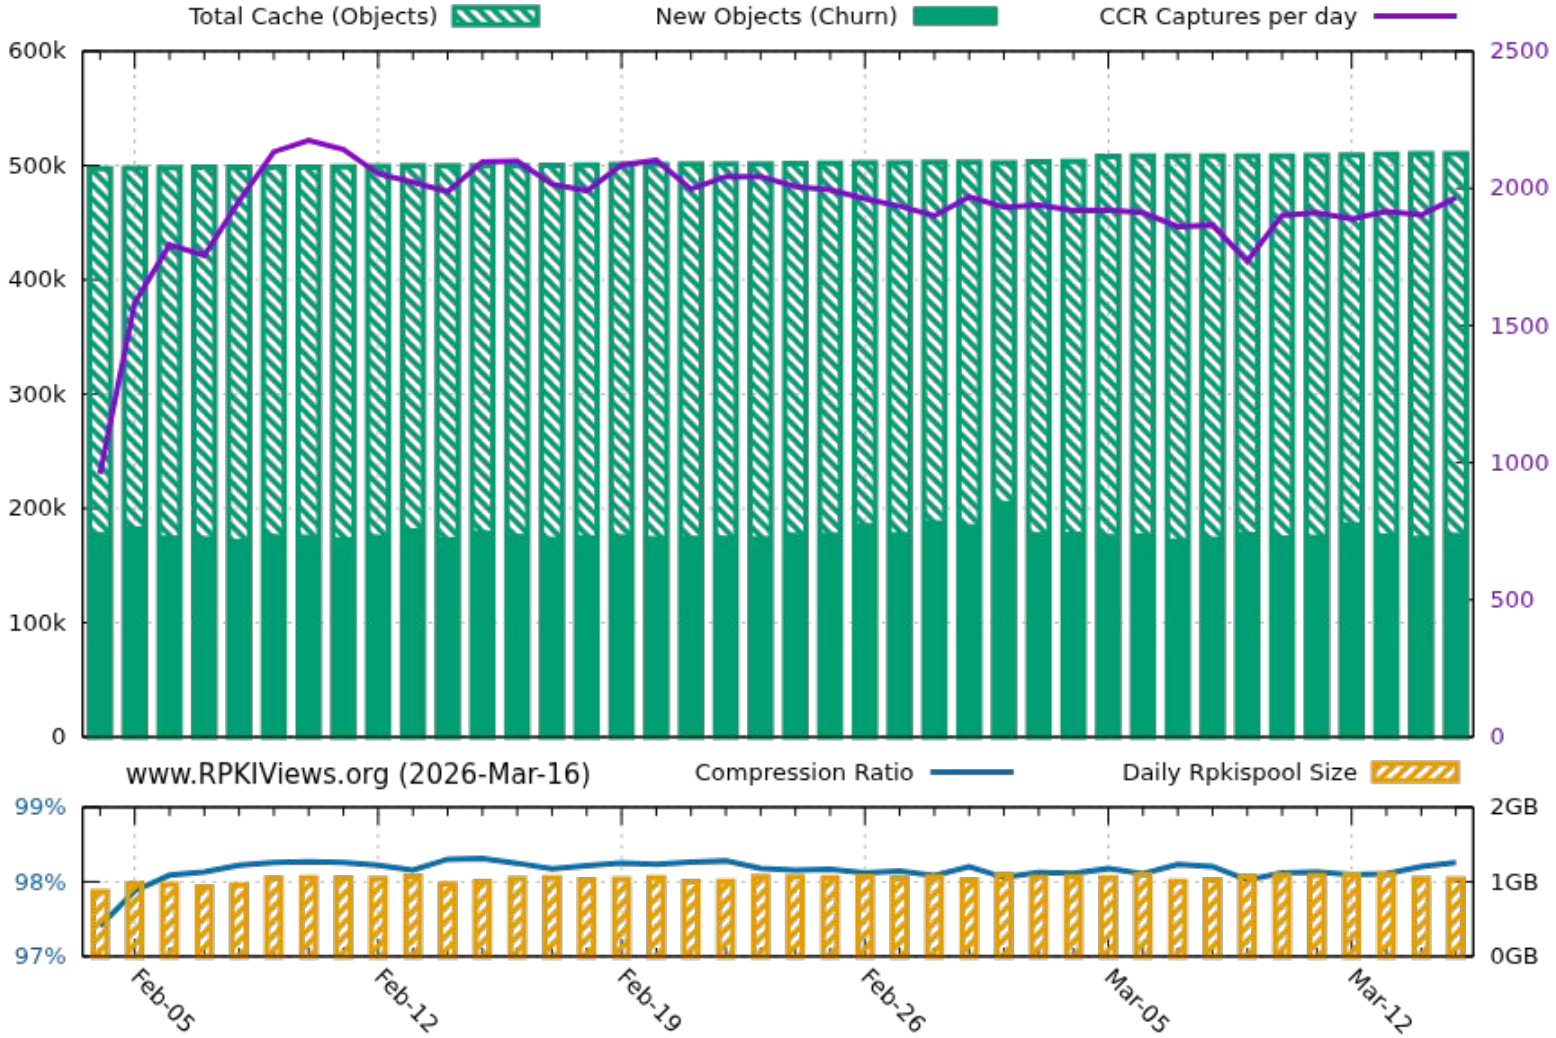
Total Cache (Objects)
New Objects (Churn)
CCR Captures per day
600k
500k
400k
300k
200k
100k
0
2500
2000
1500
1000
500
0
www.RPKIViews.org (2026-Mar-16)
Compression Ratio
Daily Rpkispool Size
99%
98%
97%
2GB
1GB
0GB
Feb-05
Feb-12
Feb-19
Feb-26
Mar-05
Mar-12

Maturing the RPKI ecosystem: incremental (re)building

- Development of new applications: e.g. ASPA (draft), RSC (RFC 9323)
- Auditing & debugging: CCR, RPKISPOOL
  - Analogy: “what MRT is for BGP”, “what DNSTAP is for DNS”
- Improving performance: Erik Synchronization Protocol
  - The *faster* the new RPKI data is distributed globally, the *better* for everyone!
- Open datasets: RPKIViews.org
- Educate next generations on how the global RPKI works!

Some lessons:

- Staying the course in multi-decade projects is very hard
  - Patience, perseverance, compassion, and trusting each other
- System designers, implementers, and operators need to jointly write the standards.
  - When only two out of three are present it will result in problems. *Working Groups must require running code.*
- Using off-the-shelf “prefab” components can come at a cost.
  - Prefab never is quite optimized for your use-case, and replacing components in a running system can be costly!