How much RPKI do you want in your BGP ?

RPKI-Aware BGP Communities and related issues

Massimiliano Stucchi - DKNOG - 7th-8th March 2024

Who am I ?

- Past RIPE NCC and Internet Society
- CHIX and IXP.ge NOC
- RIPE PC, Euro-IX PC, DKNOG PC
- SwiNOG Board
- Italian FreeBSD Users Group (GUFI)
- Run AS58280

What is this talk about ?

- There is what seems to be a problem
- We are proposing a BCP Document at the IETF to moderate it


DNS IS
INNOCENT

Elements of RPKI

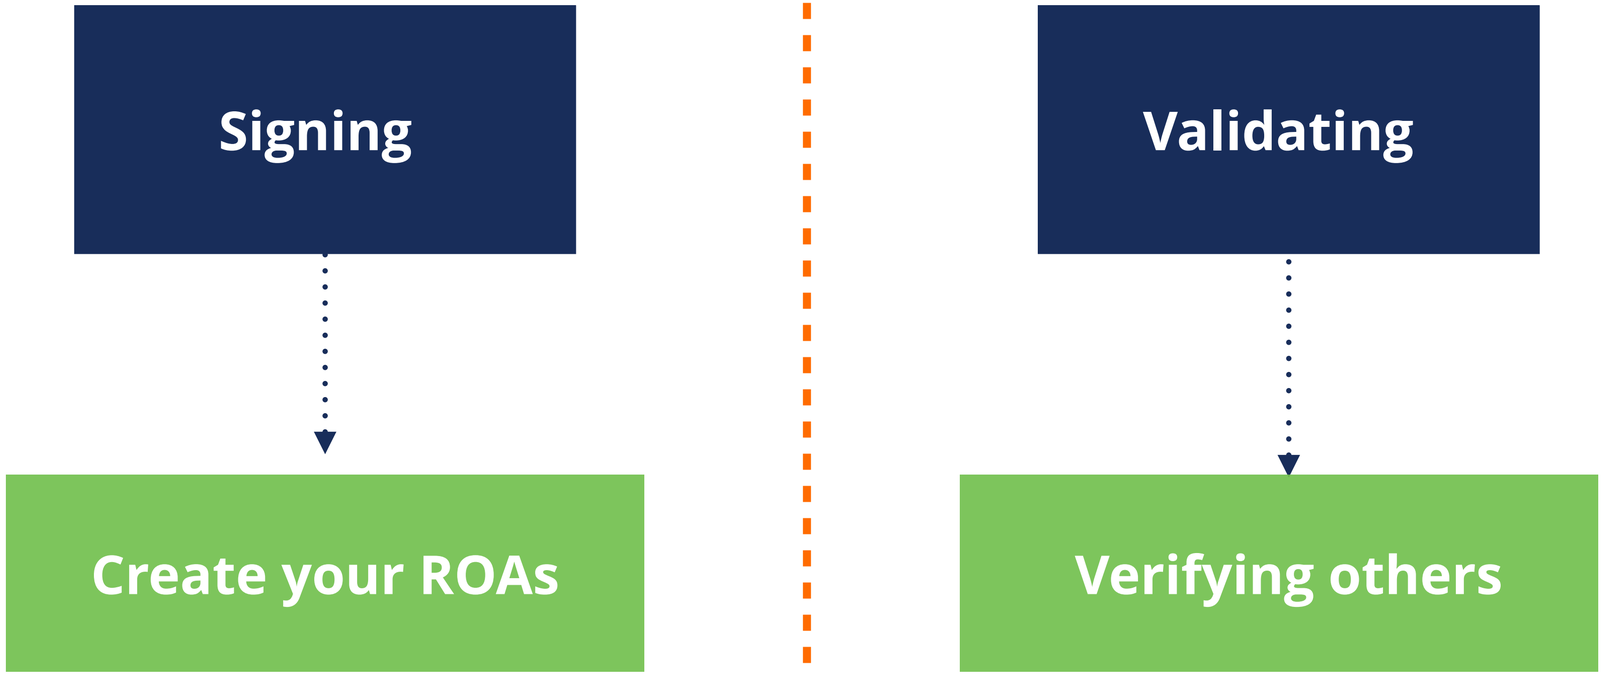

What is in a ROA ?

The network for which you are creating the ROA

The ASN supposed to originate the BGP Announcement

The maximum prefix length that ROA is authorised to advertise

RPKI - ROAs

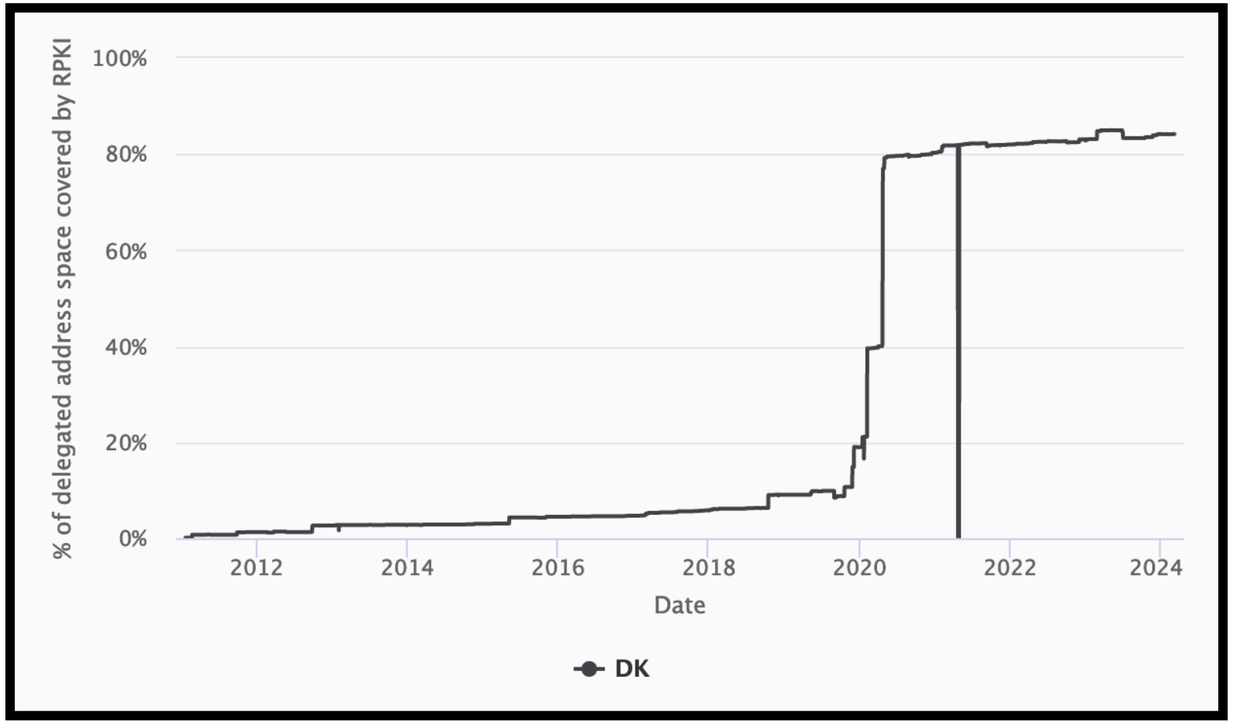

IPv4

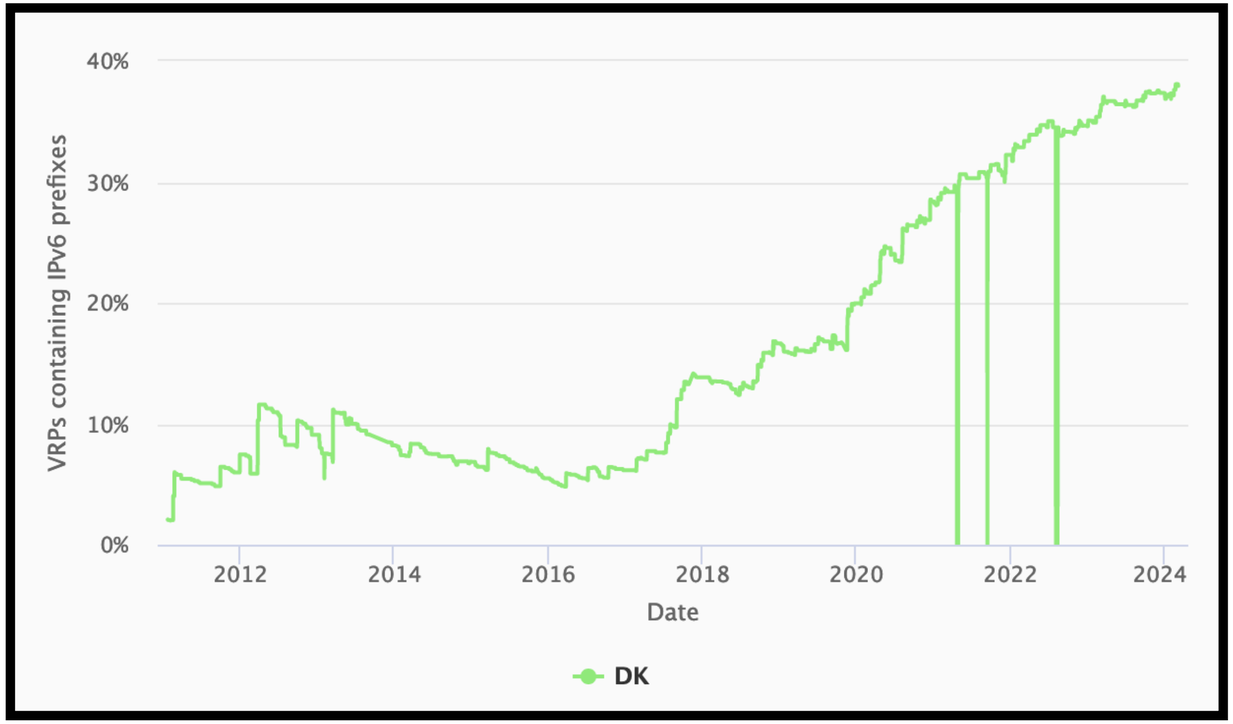

IPv6

Relying Party

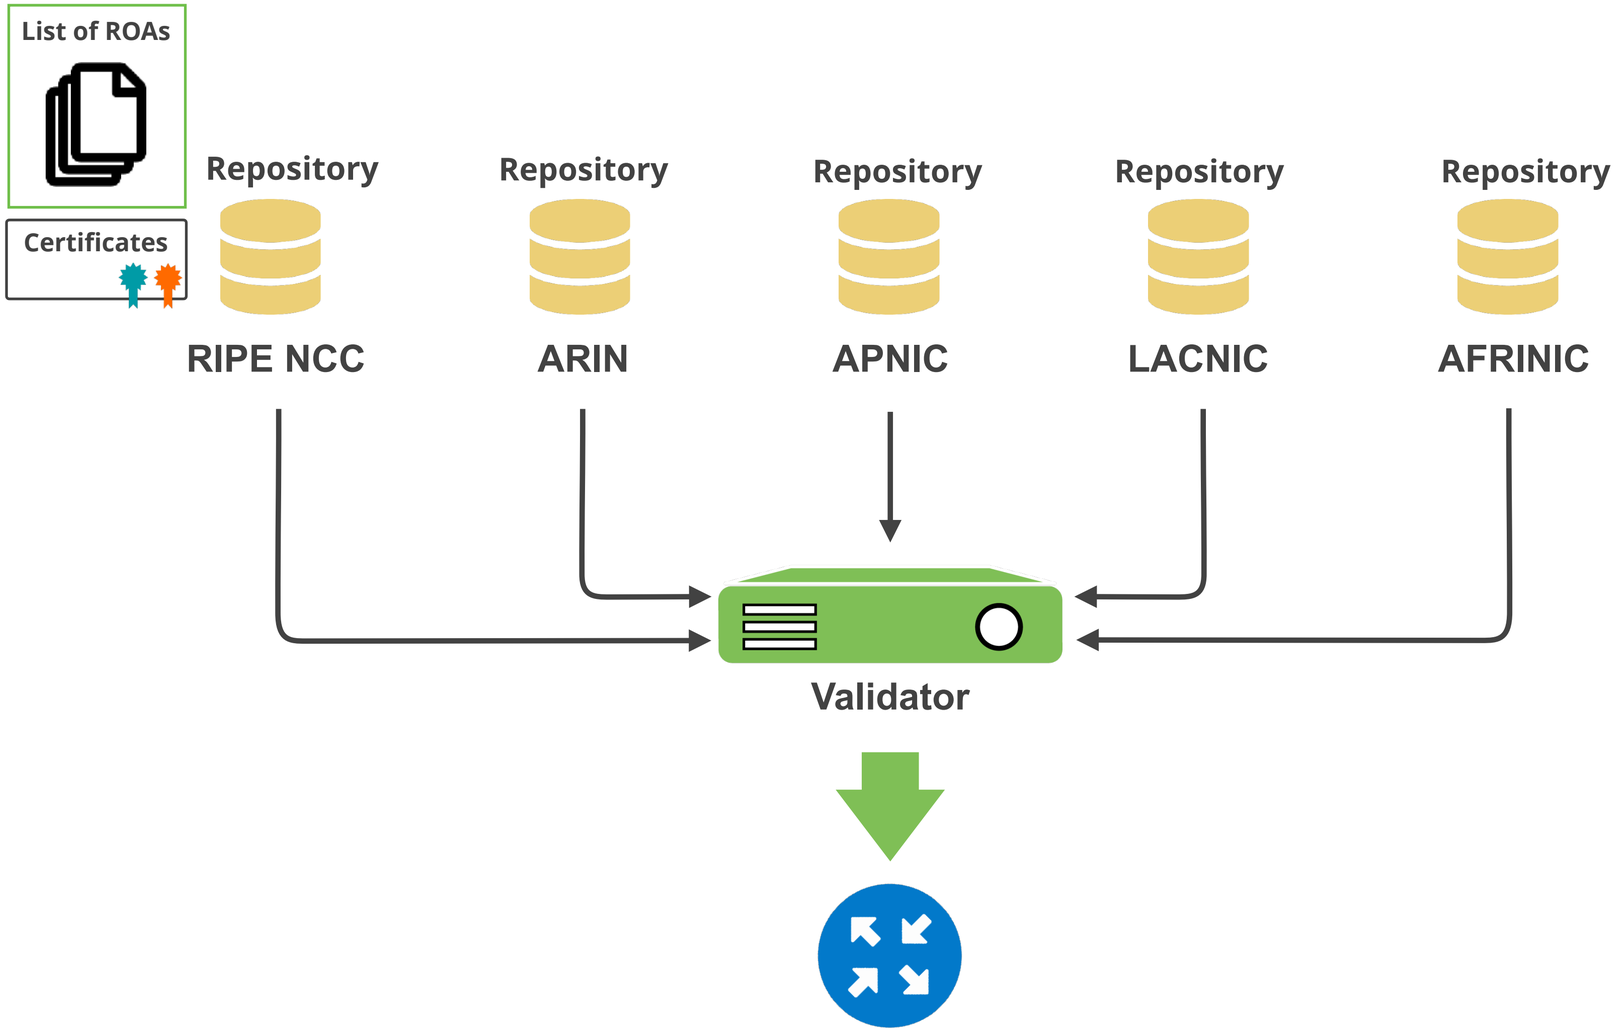

Relying Party

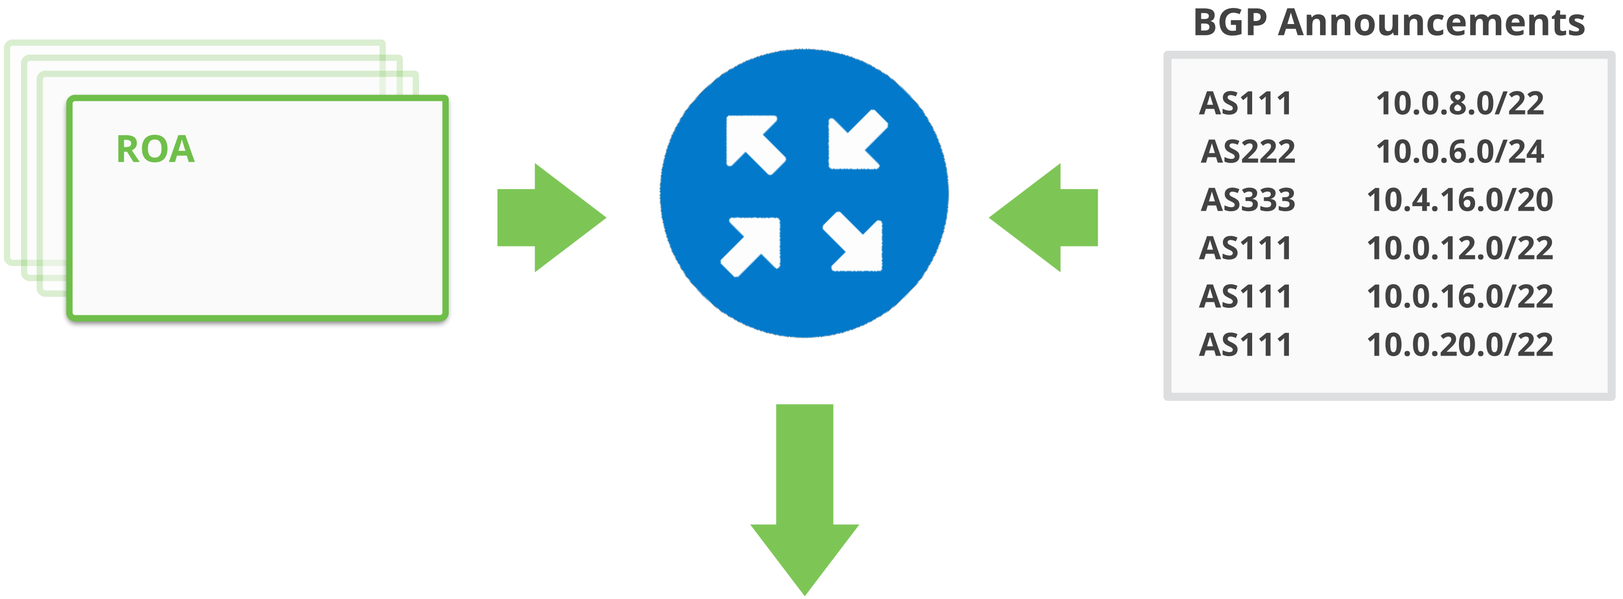

RPKI Validation States

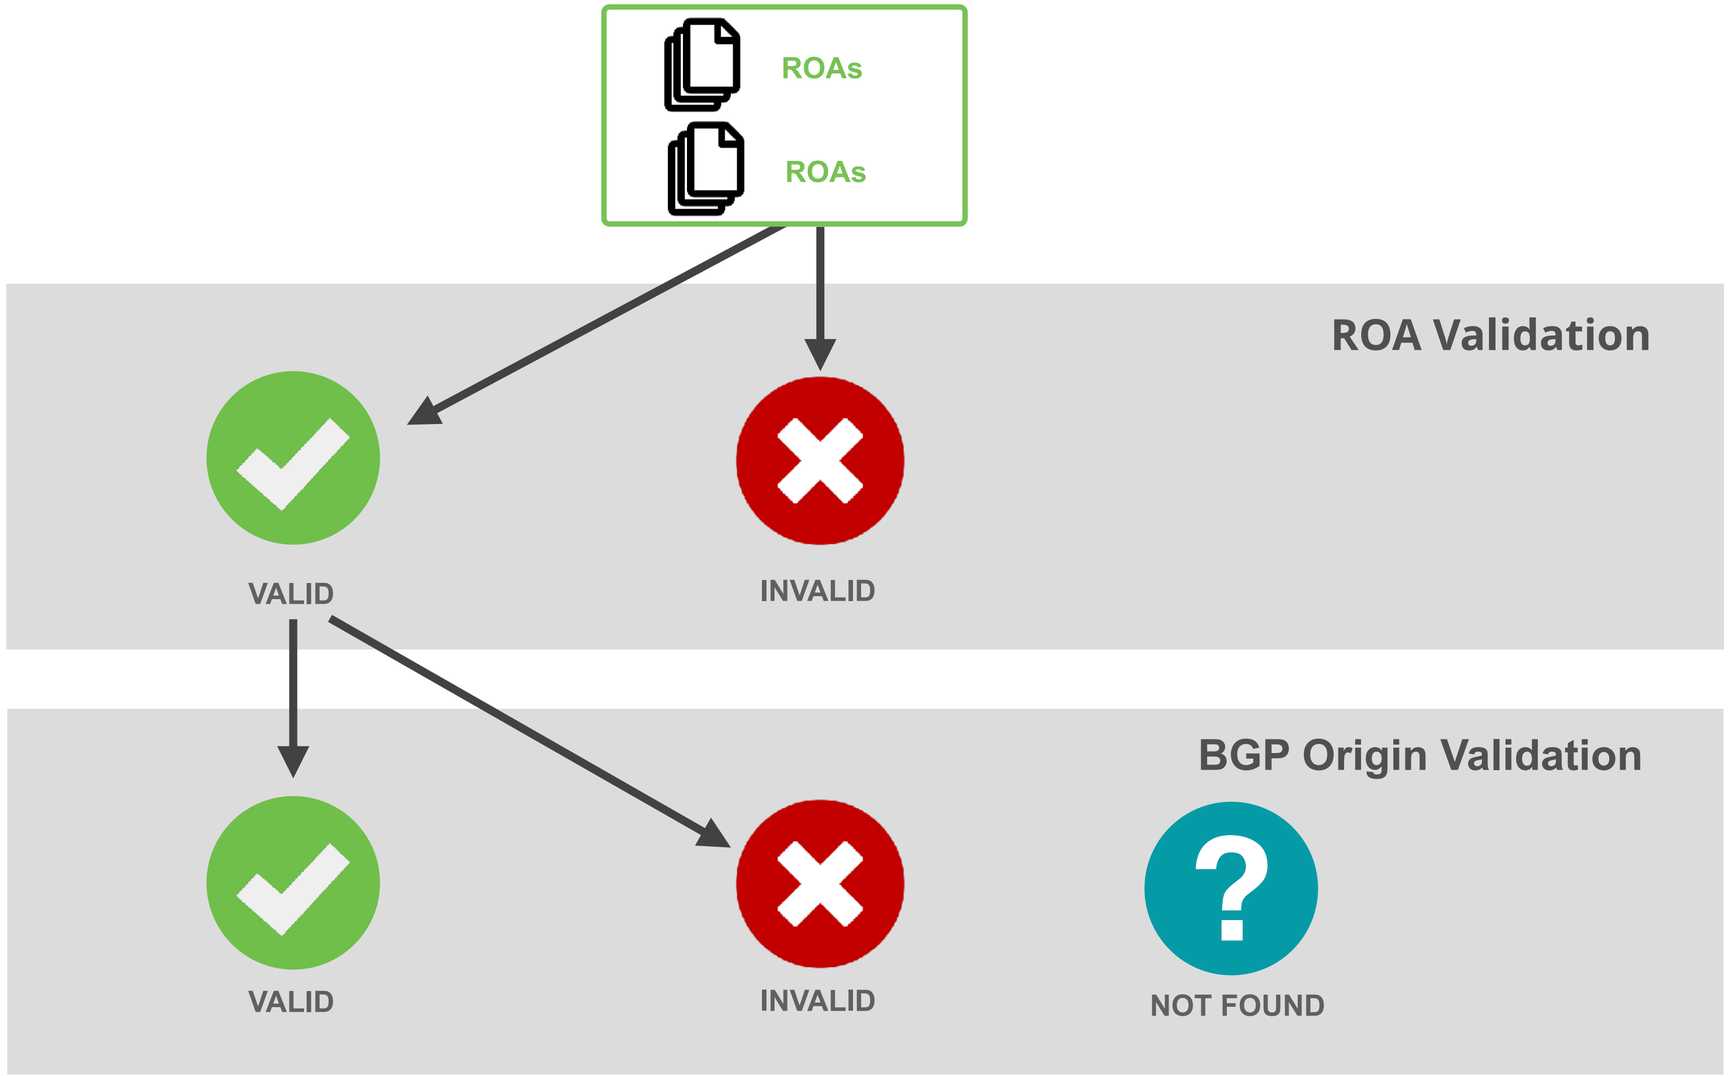

Communities

- Community is a tagging technique to mark a set of routes
  - 32-bit integer which is attached to a BGP route as an optional transitive attribute
    - as-number:community-value
  - Multiple communities can be attached to one route
  - Well-known (hard-coded) communities exist: www.iana.org/assignments/bgp-well-known-communities

Large Communities

- A simple approach continuing along the standard communities
  - 96-bit integer which is attached to a BGP route as an optional transitive attribute
    - 3x 32-bit addressing space

as-number:function:parameter

  - Larger fields, more fields, and a clean namespace separation
  - 32-bit ASN clean solution

Problem statement

- There seem to be operators using communities to carry RPKI ROV information
- Communities are a transitive attribute, so chances are they are carried along to the whole Internet
- What happens when ROV state changes for a prefix ? Can we see related updates in the DFZ ?

Outdated documentation

- There is outdated documentation out there

- Some even still recommending to install the RIPE NCC Validator
  - It was abandoned years ago

- I have approached the responsible parties of what I could find in order to update it


Although it is difficult –
stay calm and fair.

On with the data

- First stage, using BGPKit
  - Preliminary data was gathered using BGPStream

- Take a random RIB output from routeviews and check to see if some known BGP Communities carrying RPKI data come along

BGP Communities with RPKI meaning

- 1299:430 (RPKI state Valid)
- 1299:431 (RPKI state Unknown)
- 3356:901 (RPKI Valid)
- 3356:902 (RPKI Invalid)
- 3356:903 (RPKI Not Found)

Routeviews3 on 27/10/2023 at 22:00

```
Collecting data from routeviews3:
Total entries: v4 26076029 - v6 2263807
Occurrences of 3356:901 v4: 1686136 - v6: 1693
Occurrences of 3356:902 v4: 0 - v6: 0
Occurrences of 3356:903 v4: 2241052 - v6: 1199
Occurrences of 1299:430 v4: 0 - v6: 0
Occurrences of 1299:431 v4: 0 - v6: 0
```

- Good news! No 3359:902 (hence, no invalids are propagated)
- But, no visibility of anything from AS1299

Routeviews6 on 27/10/2023 at 22:00

```
Collecting data from routeviews6
Total entries: v4 0 - v6 4655520
Occurrences of 3356:901 v4: 0 - v6: 312089
Occurrences of 3356:902 v4: 0 - v6: 0
Occurrences of 3356:903 v4: 0 - v6: 279334
Occurrences of 1299:430 v4: 0 - v6: 16154
Occurrences of 1299:431 v4: 0 - v6: 16028
```

- AS3359 is more visible than AS1299 from here, but we do have some entries.

BGP Updates

- If we see these communities in the RIB, they must be also in BGP Updates
- So let’s check how many BGP updates in a 24h period have these communities
- We are going to use RIPE RIS Live using a series of route collectors around the World

BGP Communities in updates on RRC0

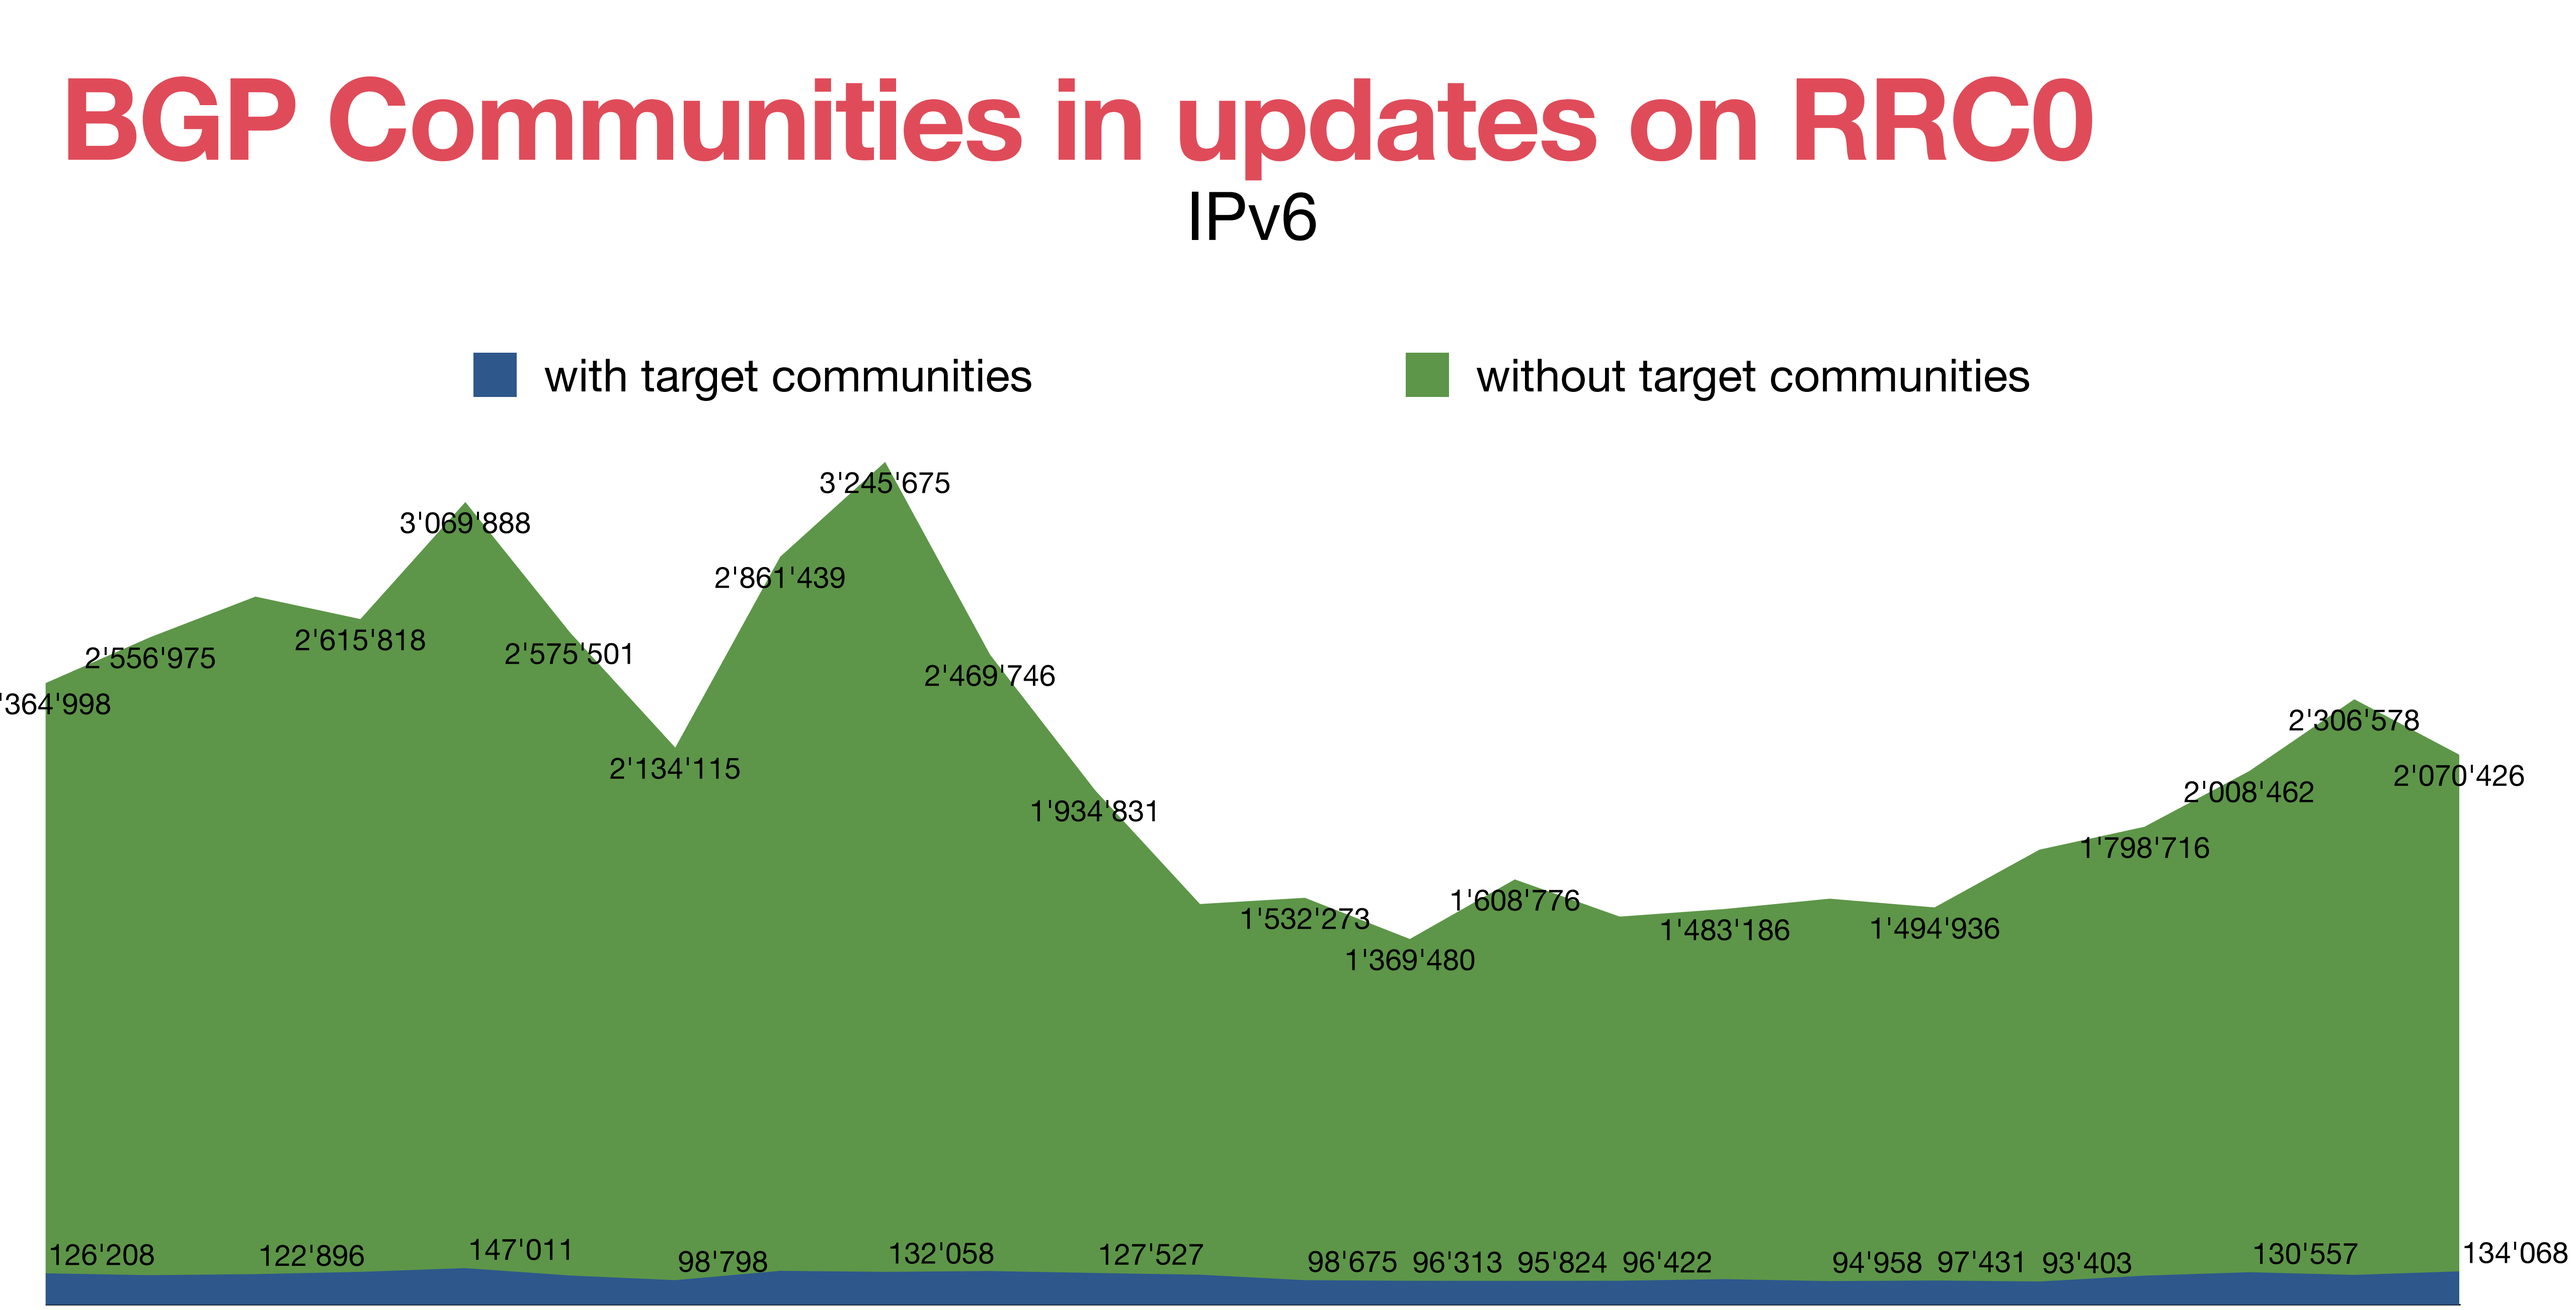

BGP Communities in updates on RRC0

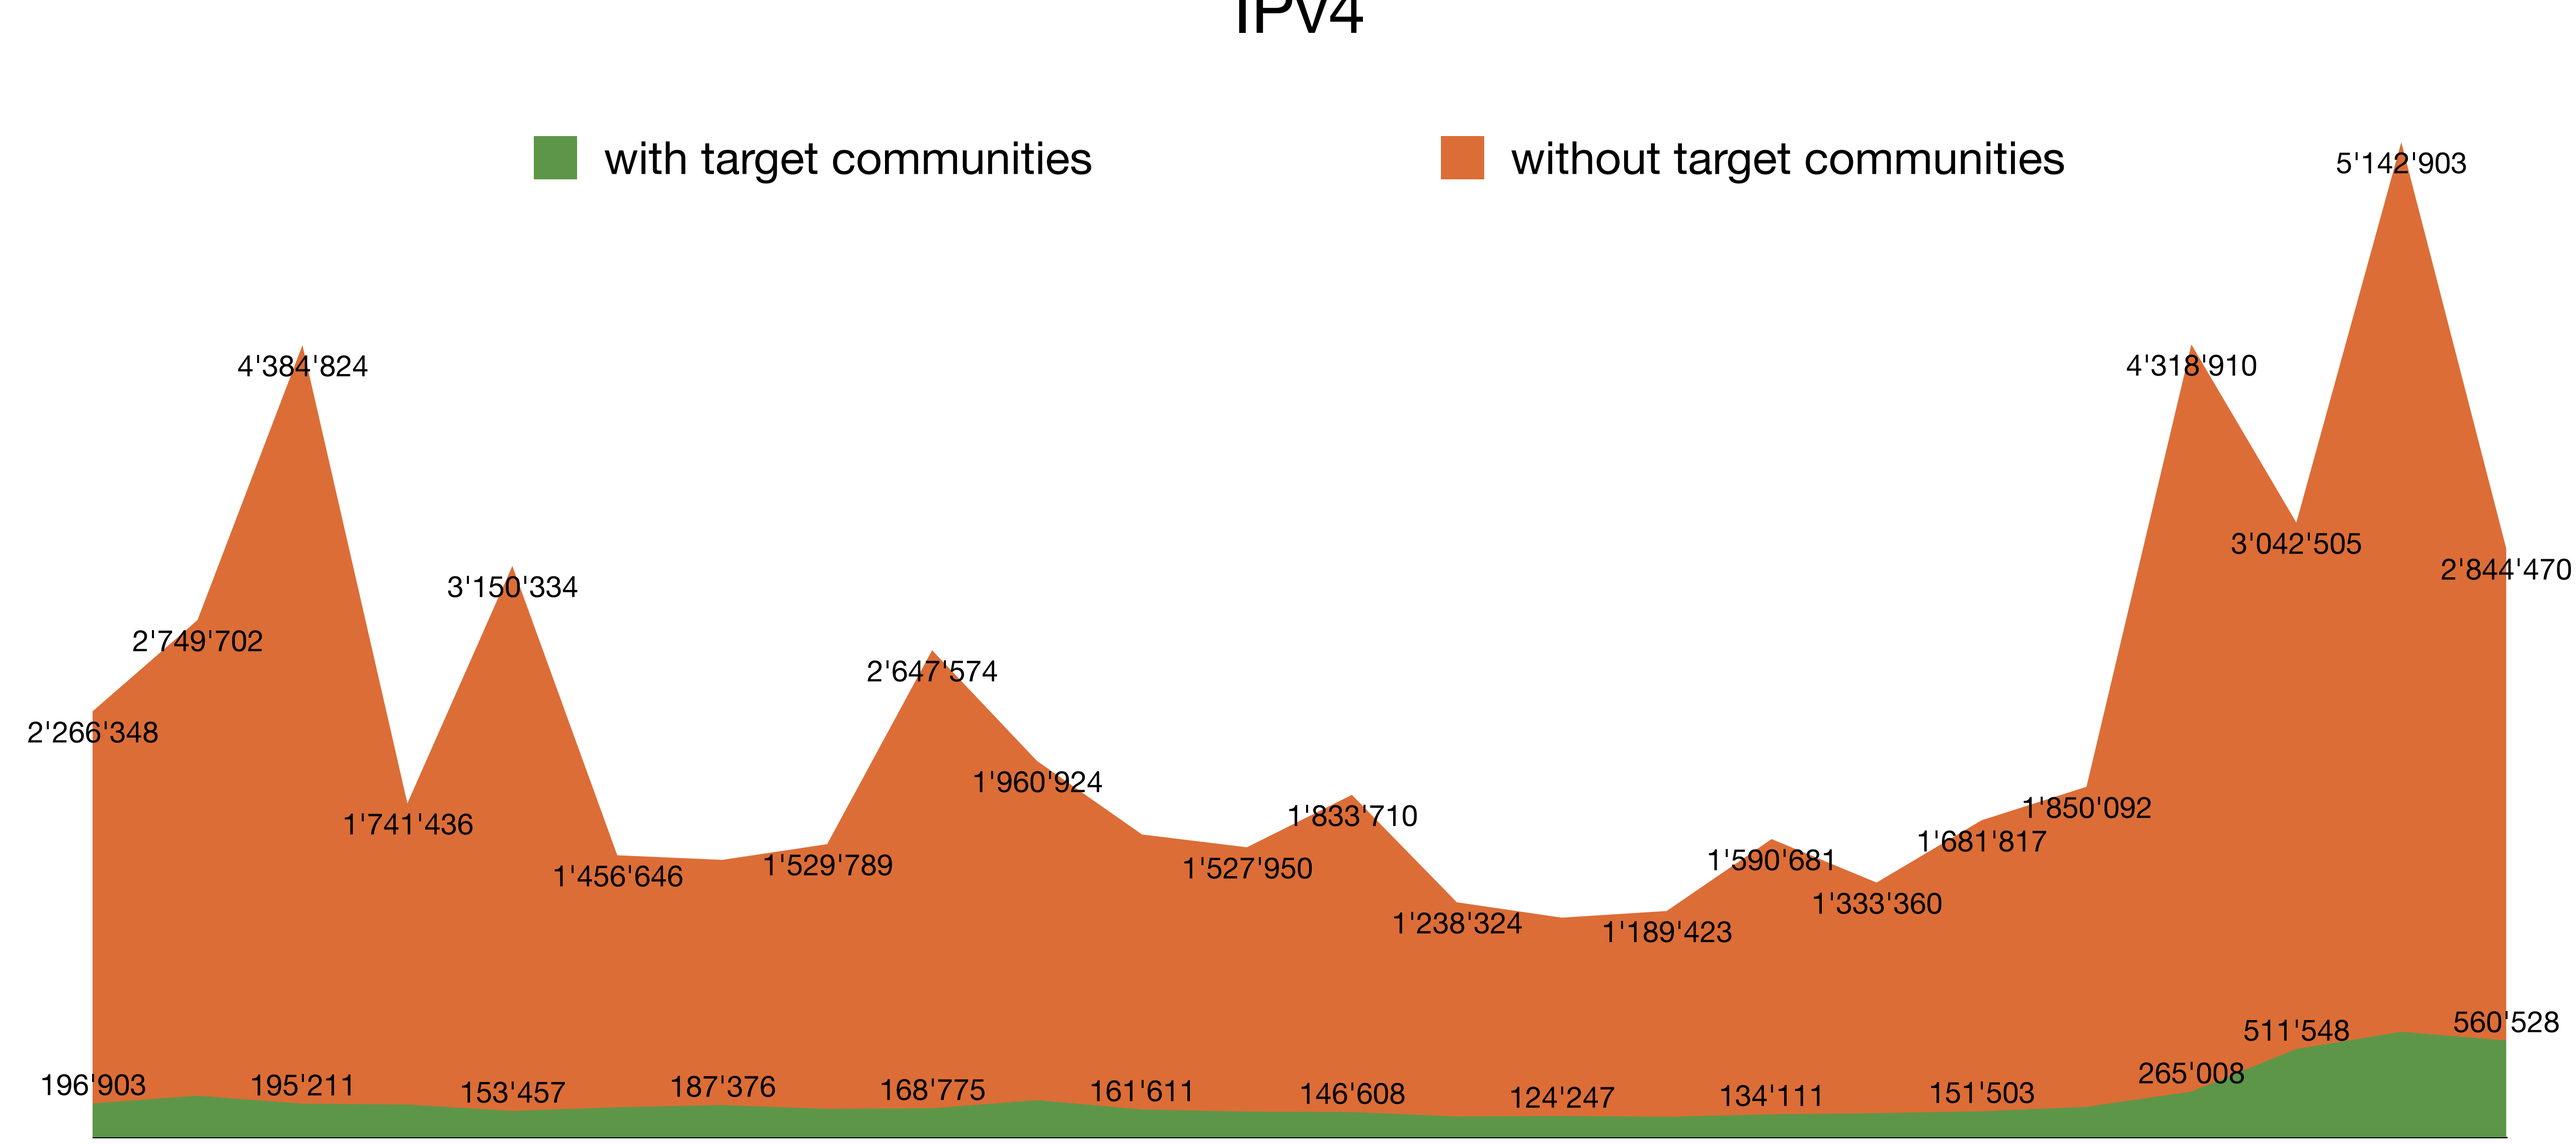

Using a test prefix

- Let’s see if announcing a prefix, and changing ROV status for it tricks updates in BGP
  - This will also help me check if other operators have special communities
- I have setup 2a0f:fd00::/29 announced by AS58280
- I have a script to update the ROA every 6 hours, either adding or removing
- And a specific filter

Results from test prefix

- Nothing visible from AS1299 or AS3356
- I spotted Coloclue having a specific community

```
Tue, Jan 30, 2024 6:05 PM - New update on RRC03 for 2a0f:fd00::/29 with communities
 [8283, 1], [8283, 101], [8283, 102], [65101, 33152], [65102, 33000], [65103, 756], [65104, 150]]
```

```
Tue, Jan 30, 2024 11:01 AM - New update on RRC03 for 2a0f:fd00::/29 with communities
 [8283, 1], [8283, 101], [65101, 33152], [65102, 33000], [65103, 756], [65104, 150]]
```

- And documenting it in RPSL

```
remarks:        ----------+------------+-------------------------------------
remarks:        8283:101  | 8283:5:1   | Accepted from peer because of valid IRR entry
remarks:        8283:102  | 8283:5:2   | Accepted from peer because of valid ROA
remarks:        8283:104  | 8283:5:4   | Accepted while RPKI invalid because it is added to our whitelist
remarks:        ----------+------------+-------------------------------------
```

Conclusions

- There is definitely a number of operators signalling RPKI ROV State in BGP Communities

- The amount of noise this generates in BGP is notable, although contained
  - WIth more and more operators enabling ROV, this noise might increase

- We need some BCP document to recommend _not_ doing it
  - Or at least, suggests removing these communities on most BGP Sessions
  - And not propagating them

BCOP Draft

- “Guidance to Avoid Carrying RPKI Validation State in Transitive BGP Path Attributes”

- https://datatracker.ietf.org/doc/draft-spaghetti-sidrops-avoid-rpki-state-in-bgp/

Future work

- Expand work to include Large Communities
  - Implemented a simple, lightweight BGP Message parser to do it on RIS Live
- A lot of manual intervention is required to check communities with the operator
  - Check their documentation or RPSL
- Repeat the prefix test but including large communities

Questions ?

max@stucchi.ch - @stucchimax@social.secret-wg.org