

Routing Security Roadmap

Job Snijders

NTT Communications

job@ntt.net

This presentation contains projections and other forward-looking statements regarding future events or our future routing performance. All statements other than present and historical facts and conditions contained in this release, including any statements regarding our future results of operations and routing positions, business strategy, plans and our objectives for future operations, are forward-looking statements (within the meaning of the Private Securities Litigation Reform Act of 1995, Section 27A of the Securities Act of 1933, as amended, and Section 21E of the Securities Exchange Act of 1934, as amended). These statements are only predictions and reflect our current beliefs and expectations with respect to future events and are based on assumptions and subject to risk and

What's this talk about?

- Differences between IRR and RPKI semantics
- What issues the industry faces
- Using the hegemonic IRR aggregator duopoly for good

Average view on routing security



Perception: it is hopeless, too many holes...



But really, there is only a **finite** amount of hurdles...



How are IRR and RPKI different?

- IRR route/route6 objects are statements:
 - About what Prefix/Origin ASN combinations can exist
 - Not necessarily made by the owner of the resource
 - Doesn't tell us anything about the validity of other route objects, or other non-matching BGP announcements
 - Unsuitable for filtering your upstream, OK-ish for peers and downstreams
 - **Not exclusive**
- RPKI on the other hand:
 - Objects are only created by resource holders
 - RFC 6811 is game changer – RPKI based BGP Origin Validation allows for non-authorized BGP announcements to be rejected
 - **Exclusive**

Exhaustive list of issues in the current ecosystem

- IRRdb / database inaccuracy (stale, autopiloted, non-validated)
- IXPs and ISPs not filtering
- Lack of Path Validation
- Lack of sufficient and good enough software

IRR – what is broken what can be fixed?

- Some IRRdbs do not perform validation
 - Meaning that virtually anyone can create virtually any route/route6 object and sneak those into the prefix-filters
- Eleven relevant IRRs not validating: RIPE, NTTCOM, RADB, ALTDB, ARIN IRR, BBOI, BELL, LEVEL3, RGNET, TC, CANARIE
- Two solutions:
 - Lock the database down (RIPE / RIPE-NONAUTH)
 - Filter on the mirror level

RIPE NWI-5 proposal & implementation

- RIPE NCC's IRR previously allowed anyone to register any non-RIPE-managed space if it had not yet been registered. *DANGER*
- The “RPSL” password & maintainer was used for this

SOLVED

Three steps were taken:

- Cannot register non-RIPE-managed space any more
- All non-RIPE space moved to separate “RIPE-NONAUTH” database
- Route/route6 ASN authorization rules have been improved

More info: <https://www.ripe.net/manage-ips-and-asns/db/impact-analysis-for-nwi-5-implementation>

OK – so current status

- Ten relevant IRRs not validating: NTTCOM, RADB, ALTDB, ARIN IRR, BBOI, BELL, LEVEL3, RGNET, TC, CANARIE
- Done: ~~RIPE~~

ARIN IRR allows anyone to register anything

```
hanna:~ job$ whois -h rr.arin.net 2001:67c:208c::
% This is the ARIN Routing Registry.
% Note: this output has been filtered.
%       To receive output for a database update, use the "-B" flag.

% Information related to '2001:67c:208c::/48AS15562'

route6:      2001:67c:208c::/48
descr:       2001:67c:208c::/48 - Job's net
remarks:     Job asked me to steal his net.  Honest!
origin:      AS15562
mnt-by:      MNT-ATTW-Z
source:      ARIN # Filtered
```

ARIN community also recognized this is an issue

- Consultation at [NANOG](#) and through [ARIN-Consult](#) mailing list
- https://www.arin.net/vault/resources/routing/2018_roadmap.html
- <https://teamarin.net/2018/07/12/the-path-forward/>

“Improve, or kill it”

ALMOST SOLVED

OK – so current status

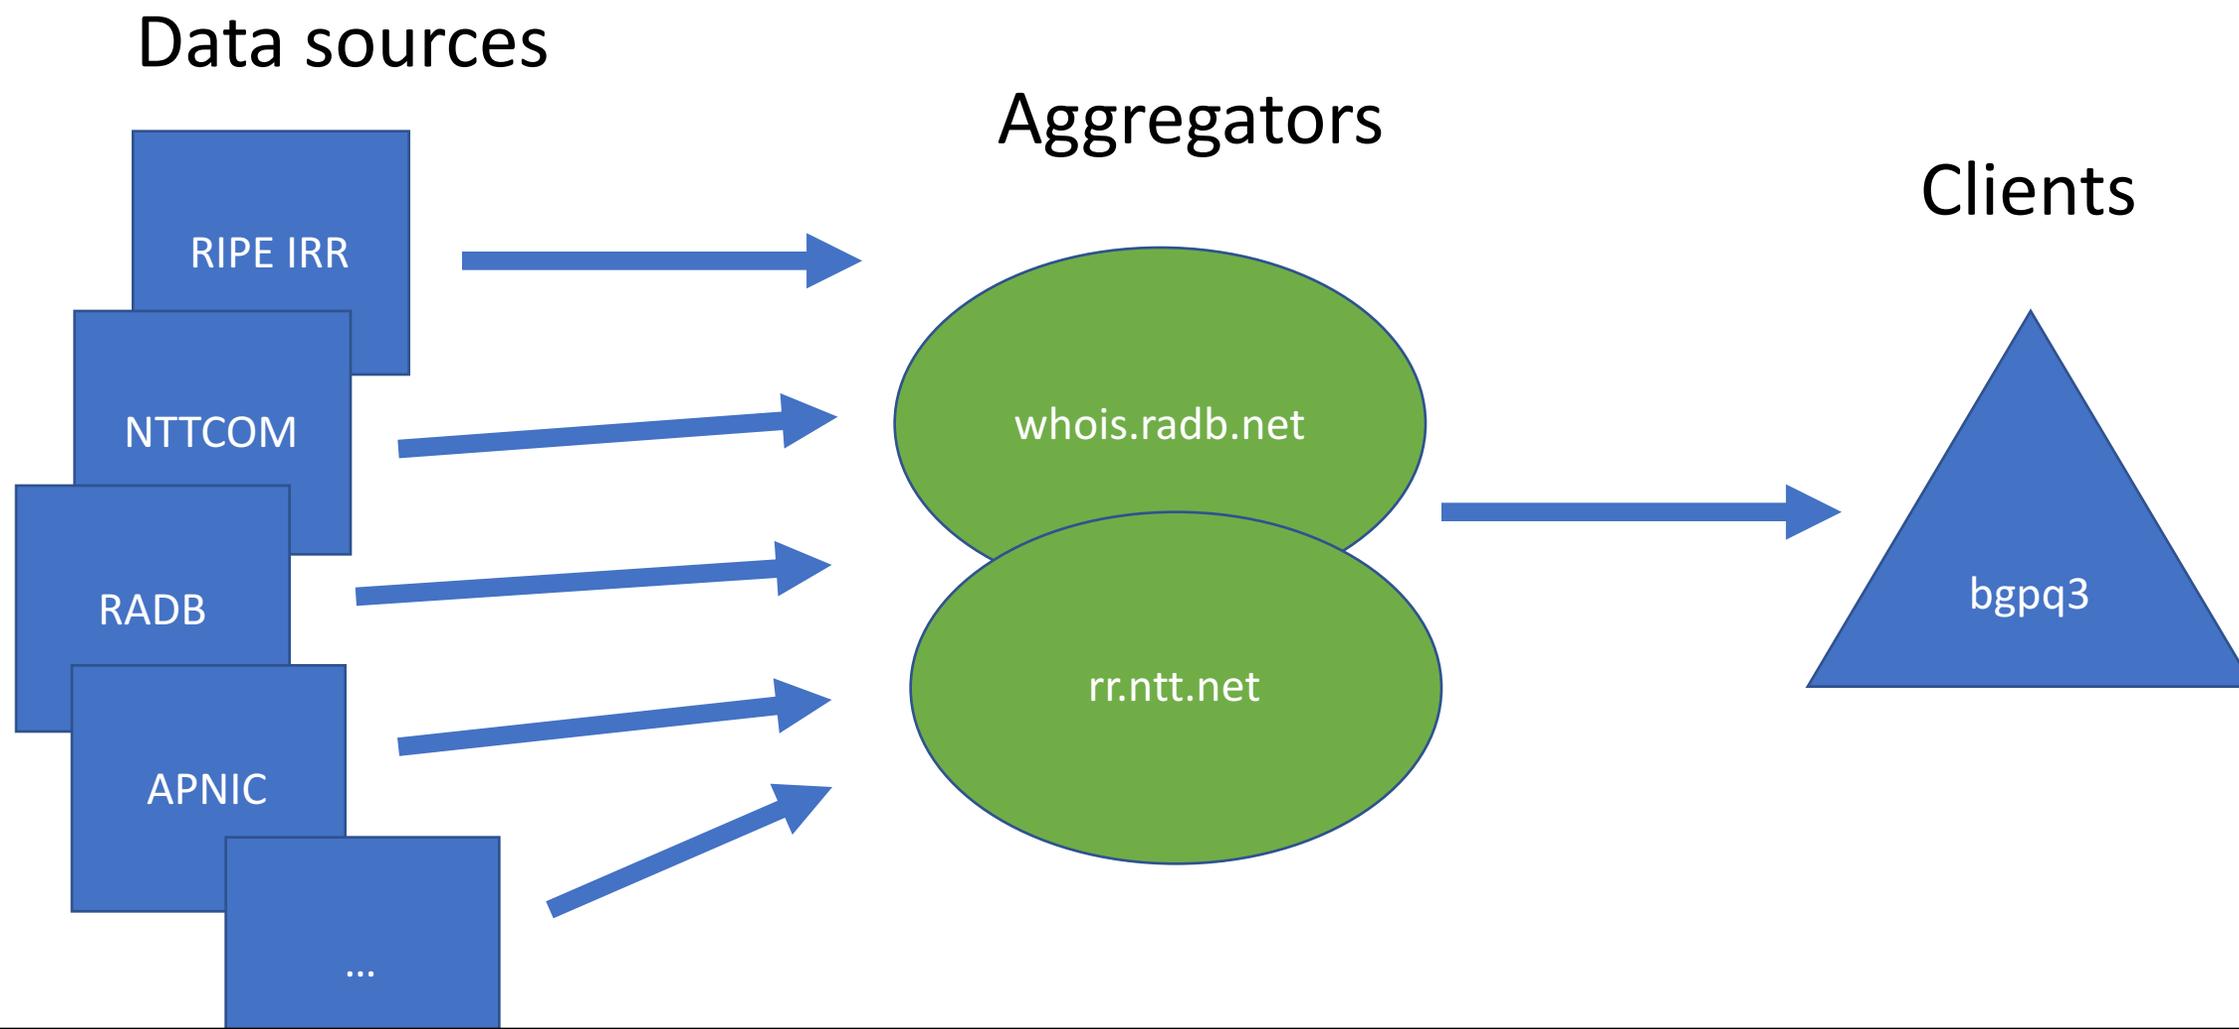
- Nine relevant IRRs not validating: NTTCOM, RADB, ALTDB, BBOI, BELL, LEVEL3, RGNET, TC, CANARIE
- Done: ~~RIPE, ARIN IRR~~
- How to deal with the remaining nine ?
- Not all of these are so easily communicated with, not all are really actively managed

The “IRR” system access

- The IRR is access through predominantly two “gateways”
 - **whois.radb.net** (the `bgpq3` and `peval` default)
 - **rr.ntt.net**
- All mirroring is essentially done with one software: [IRRd](#)

Solution: Let’s use the hegemonic duopoly for good!

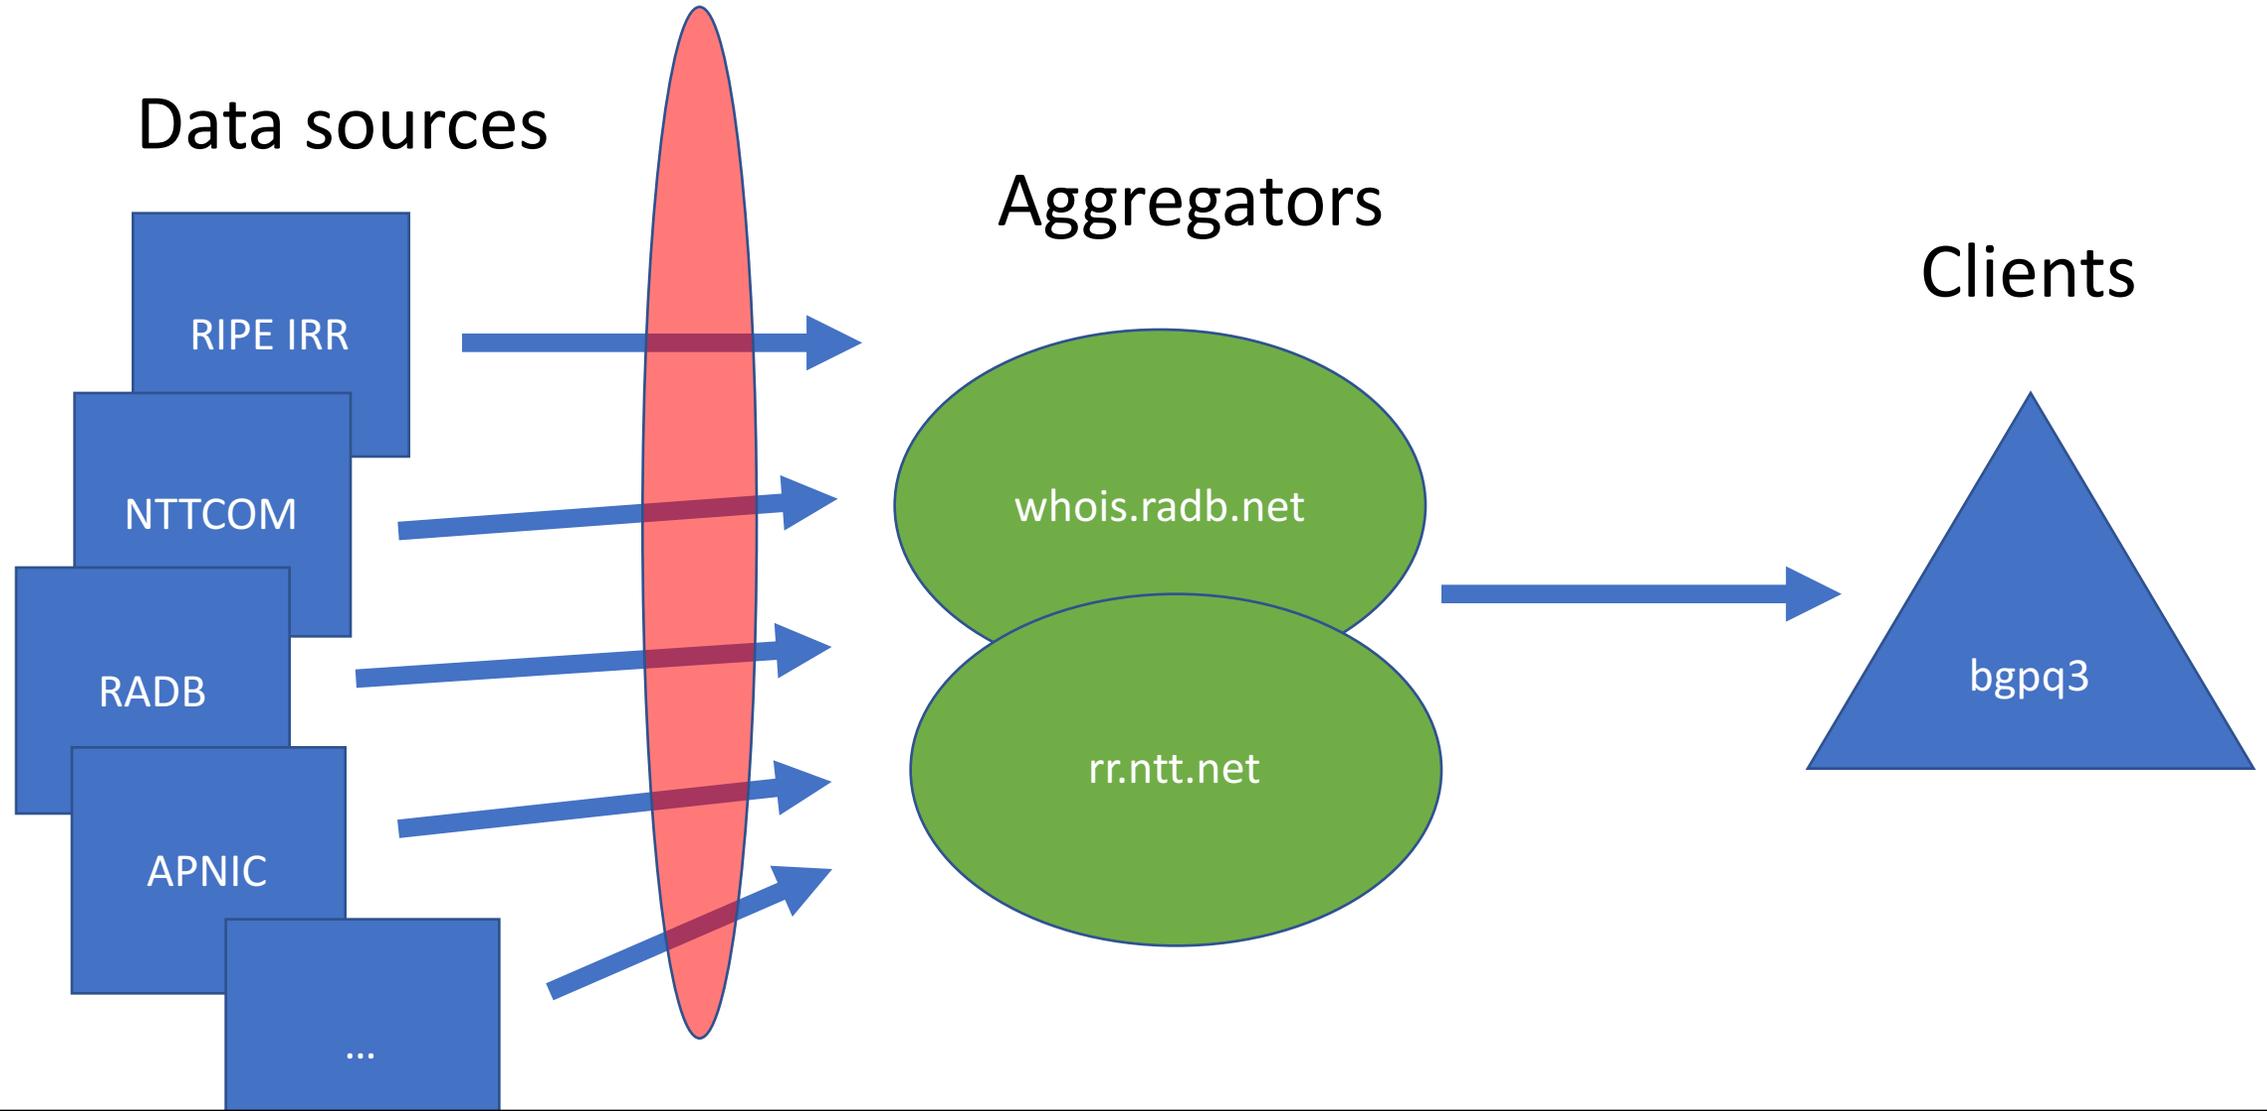
Improving security at the "aggregator"?



Proposal: Let RPKI “drown out” conflicting IRR

- RPKI can be used for *BGP Origin Validation* – but also for other things!
- A RPKI ROA is sort of a route-object
 - It has a “prefix”, “origin” and “source” (the root)
 - We can [use RPKI ROAs for provisioning BGP prefix-filters](#)
- Extend IRRd so that when IRR information is in direct conflict with a RPKI ROA – the conflicting information is suppressed ([Github](#))

RPKI filter at the aggregators



RPKI suppressing conflicting IRR advantages

- Industry-wide common method to get rid of stale proxy route objects – by creating a ROA you hide old garbage in IRRs
- By creating a ROA – you will significantly decrease the chances of people being able to use IRR to hijack your resource

This idea is also being discussed in RIPE community

OK – so current status

- IRRs not validating: no longer problematic
- Done: ~~RIPE, ARIN IRR, NTTCOM, RADB, ALTDB, BBOI, BELL, LEVEL3, RGNET, TC, CANARIE~~

SOLVED

NTT & Dashcare have started a full rewrite of IRRd to make this possible:
<https://github.com/irrdnet/irrd4>

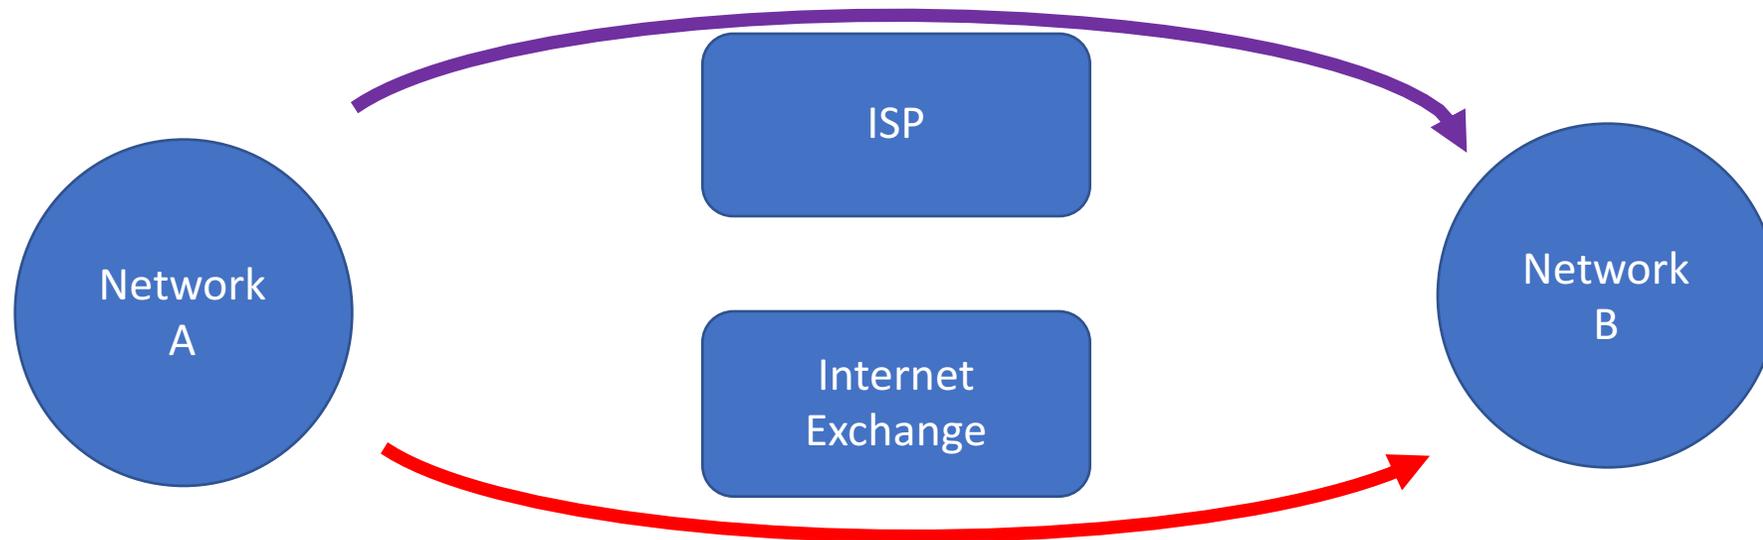
”Filtering at IXPs is hard”



- Many IXPs have come to realize their responsibilities to the Internet ecosystem and the commercial benefits of a more secure product.
- <http://peering.exposed/>
 - 9 out of top 10 IXPs are filtering, tenth will later this year. **IX.br** making good progress
- IXP filtering has become much easier, there are multiple fully featured configuration generators:
 - <https://www.ixpmanager.org/>
 - <http://arouteserver.readthedocs.io/>

Route servers must begin dropping RPKI Invalids

- Route servers *by definition* provide partial Internet tables
- No guarantees whatsoever that a given route will be available via RS
- When a route server drops a prefix, **worst case scenario is rerouting** – not an outage.



Not everyone needs to do RPKI

- Because of the centralization of the web, if a select few companies deploy RPKI Origin Validation – millions of people benefit
- (google, cloudflare, amazon, pch/quad9, facebook, akamai, fastly, liberty global, comcast, etc...)
- I think only 20 companies or so need to do Origin Validation for there to be big benefits...
- <https://dyn.com/blog/bgp-dns-hijacks-target-payment-systems/>

“RPKI Origin Validation is useless without Path Validation aka BGPSEC”

- The lack of path validation can be resolved through two methods:
 - Densely peer with each other (Example: Google & Akamai have 126+ facilities in common with each other)
 - An AS_PATH blocking mechanisms like [“peerlock”](#)
- Both effectively are “path validation for 1 hop”
- Perhaps “1 hop” already is good enough 😊

“There is no healthy software ecosystem”

- RIPE NCC Validator v3 is works and actively maintained
- NLNetlabs released their RPKI Cache Validator (Routinator 3000)
- OpenBSD is looking to fund/develop a third validator implementation

- Almost all serious routing vendors have RPKI support (Cisco, Juniper, BIRD, Nokia, FRR – and more are on the way)

- Solution: more users results in better software, start using!

SOLVED

Timeline

- IETF meetings should start **now!**
- IXPs – start doing RPKI Origin Validation on your route servers **now**
- **Quite some companies are deploying RPKI OV before the end of the year!**
- ISPs / CDNs
 - if you are pointing default somewhere and have local peering, do it **now**
- In 2019 RPKI data will be used to clean up IRR
- Hopefully the ARIN RPKI TAL situation will improve in 2019

Conclusion

