

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

Why are we doing any of this?

- Creating filters based on public data, forces malicious actors to leave a trail in IRR, WHOIS or other data sources: **auditability**
- **Bugs happen!** – your router may suddenly ignore parts of your configuration, you'll then rely on your EBGP peer's filters
- **Everyone makes mistakes** – a typo is easily made

Average view on routing security



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



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



Exhaustive list of issues in the current ecosystem

- IRRdb / database inaccuracy (stale, autopiloted, non-validated)
- IXPs 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



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

Using RPKI to clean up the IRR




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!

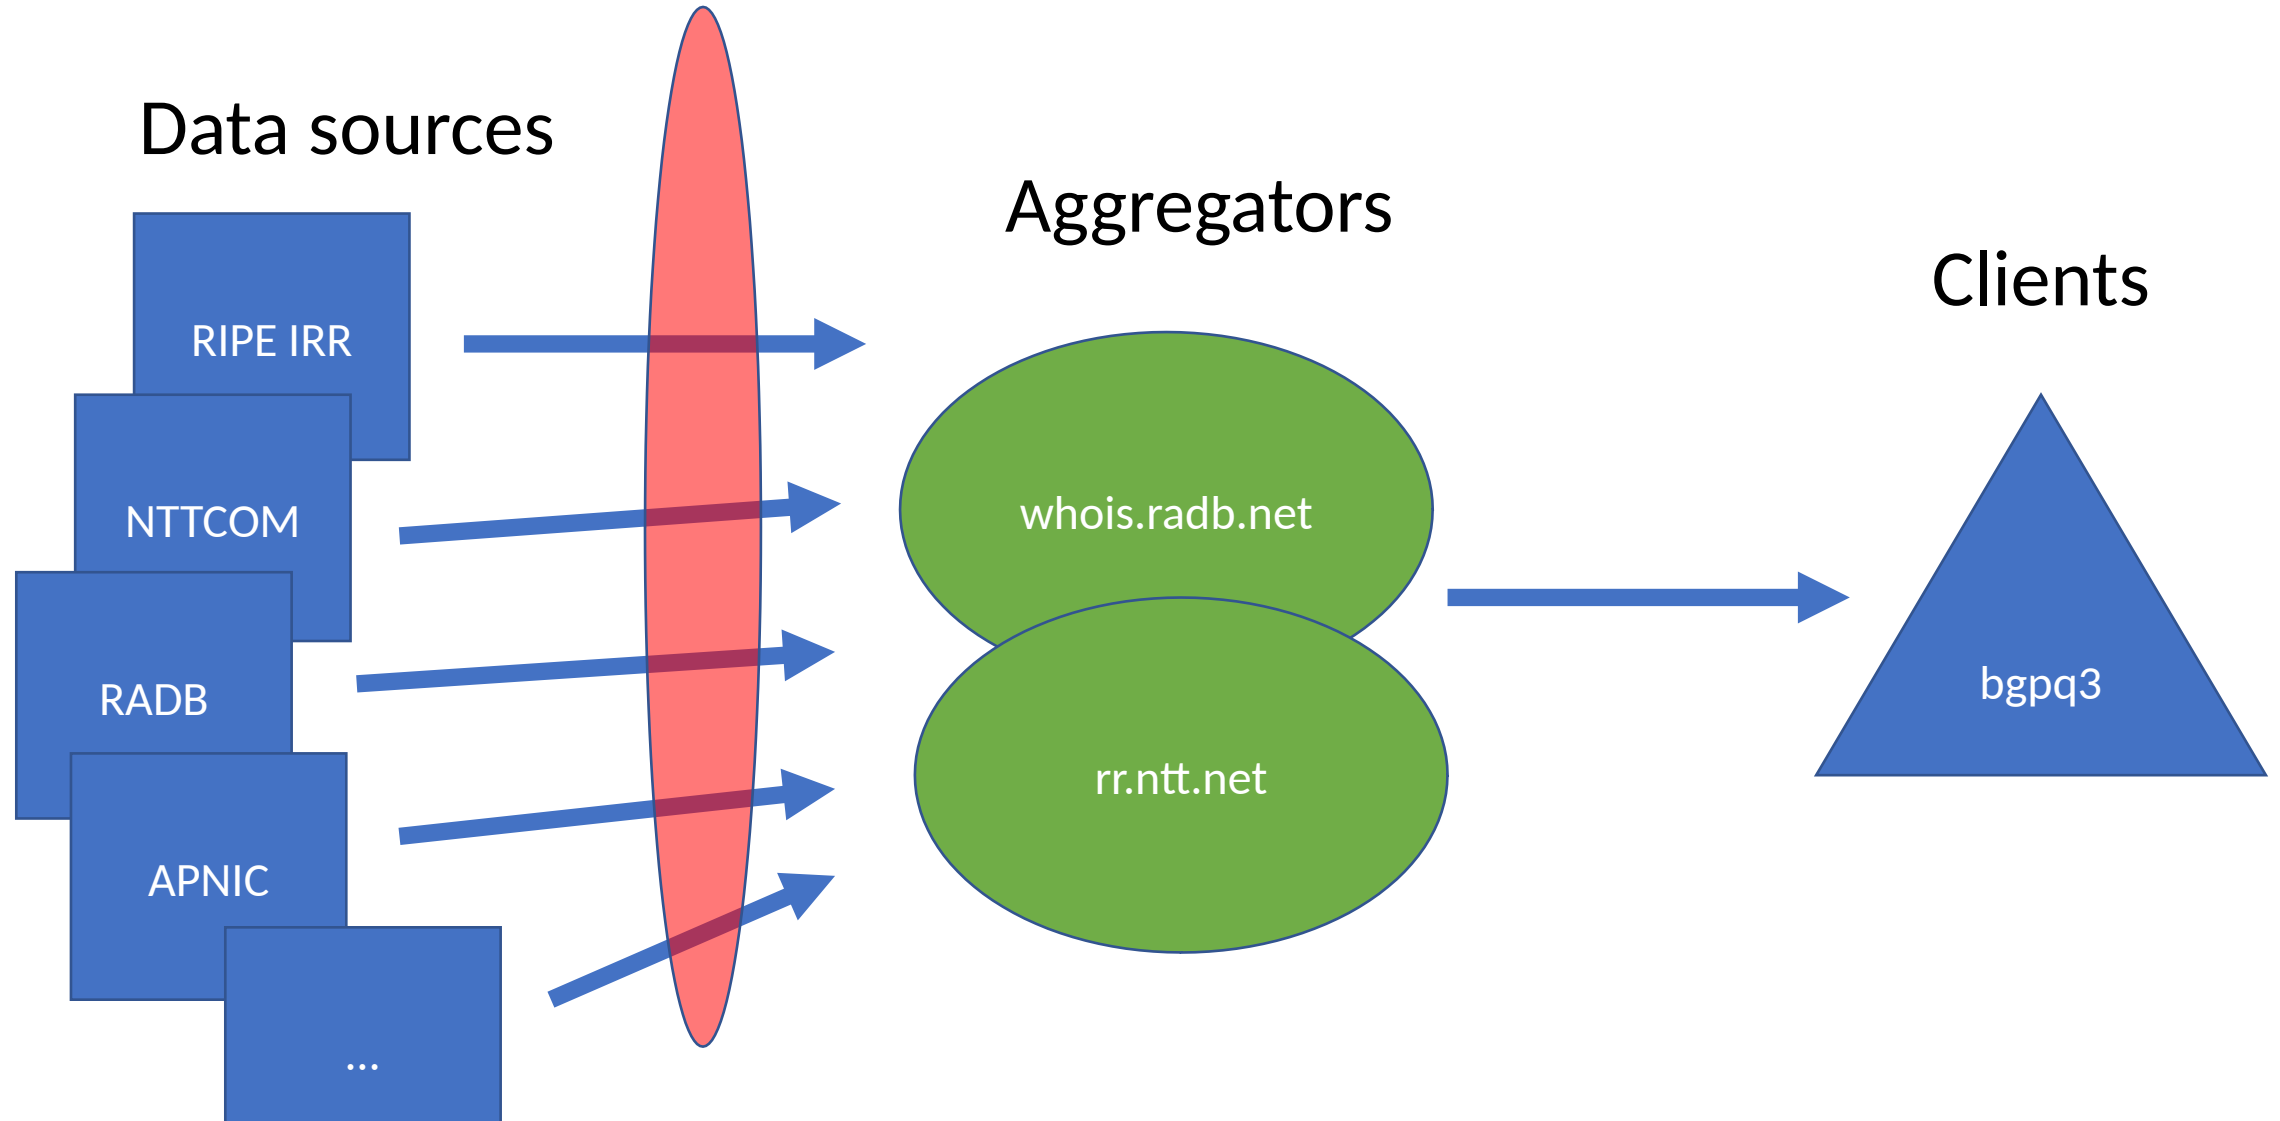
Improving security at the "aggregator"?



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



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**

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

OK – so current status

SOLVED

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



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/> version 5 has RPKI support!
 - <http://arouteserver.readthedocs.io/>

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/>

Issue: "RPKI Origin Validation is useless without Path Validation aka BGPSEC...."



The internet keeps connecting directly

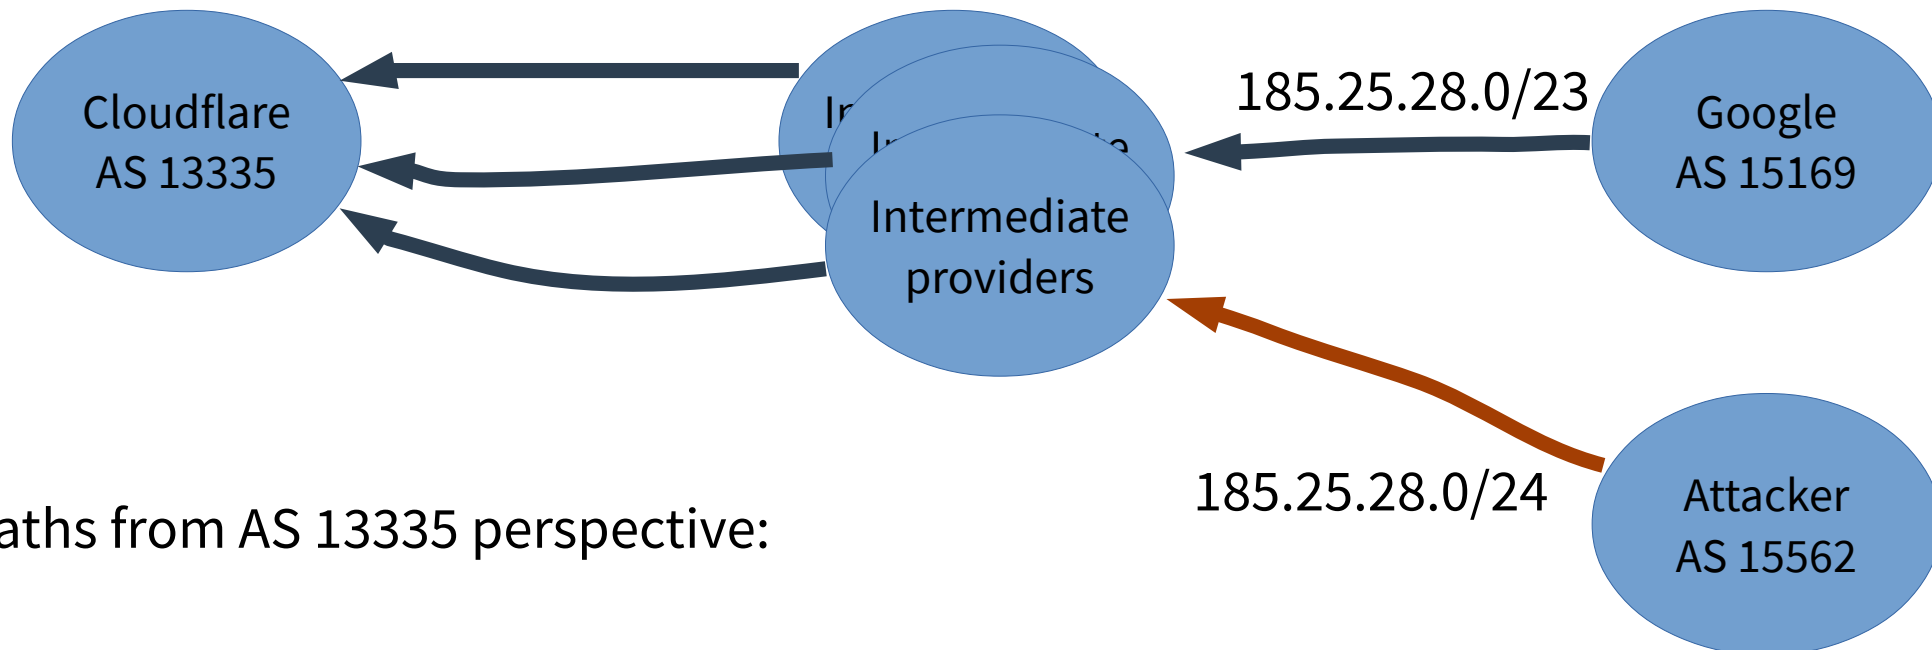
Average AS Path Length



2012 Source:

<https://labs.ripe.net/Members/mirjam/update-on-as-path-lengths-over-time>

Hijack / misconfiguration scenario




Paths from AS 13335 perspective:

185.25.28.0/23 13335_XXX_15169

185.25.28.0/23 13335_YYY_15169

185.25.28.0/24 13335_ZZZ_15562 (wins)

Hijack / misconfiguration scenario – direct peering




Enter RPKI ROAs

| | |
|---------------------|----------------------------------|
| Prefix: | 185.25.28.0/23 |
| Prefix description: | Google |
| Country code: | CH |
| Origin AS: | 15169 |
| Origin AS Name: | G00GLE - Google LLC, US |
| RPKI status: | ROA validation successful |
| MaxLength: | 23 |
| First seen: | 2016-01-08 |
| Last seen: | 2019-02-26 |
| Seen by #peers: | 40 |

Hijack / misconfiguration scenario – RPKI ROA

Cloudflare applying “invalid == reject”




Paths from AS 13335 perspective:

185.25.28.0/23 13335_15169 (wins)

185.25.28.0/24 13335_15562 (rejected, wrong prefix length)

Change of tactics: announce same prefix

Cloudflare applying “invalid == reject”



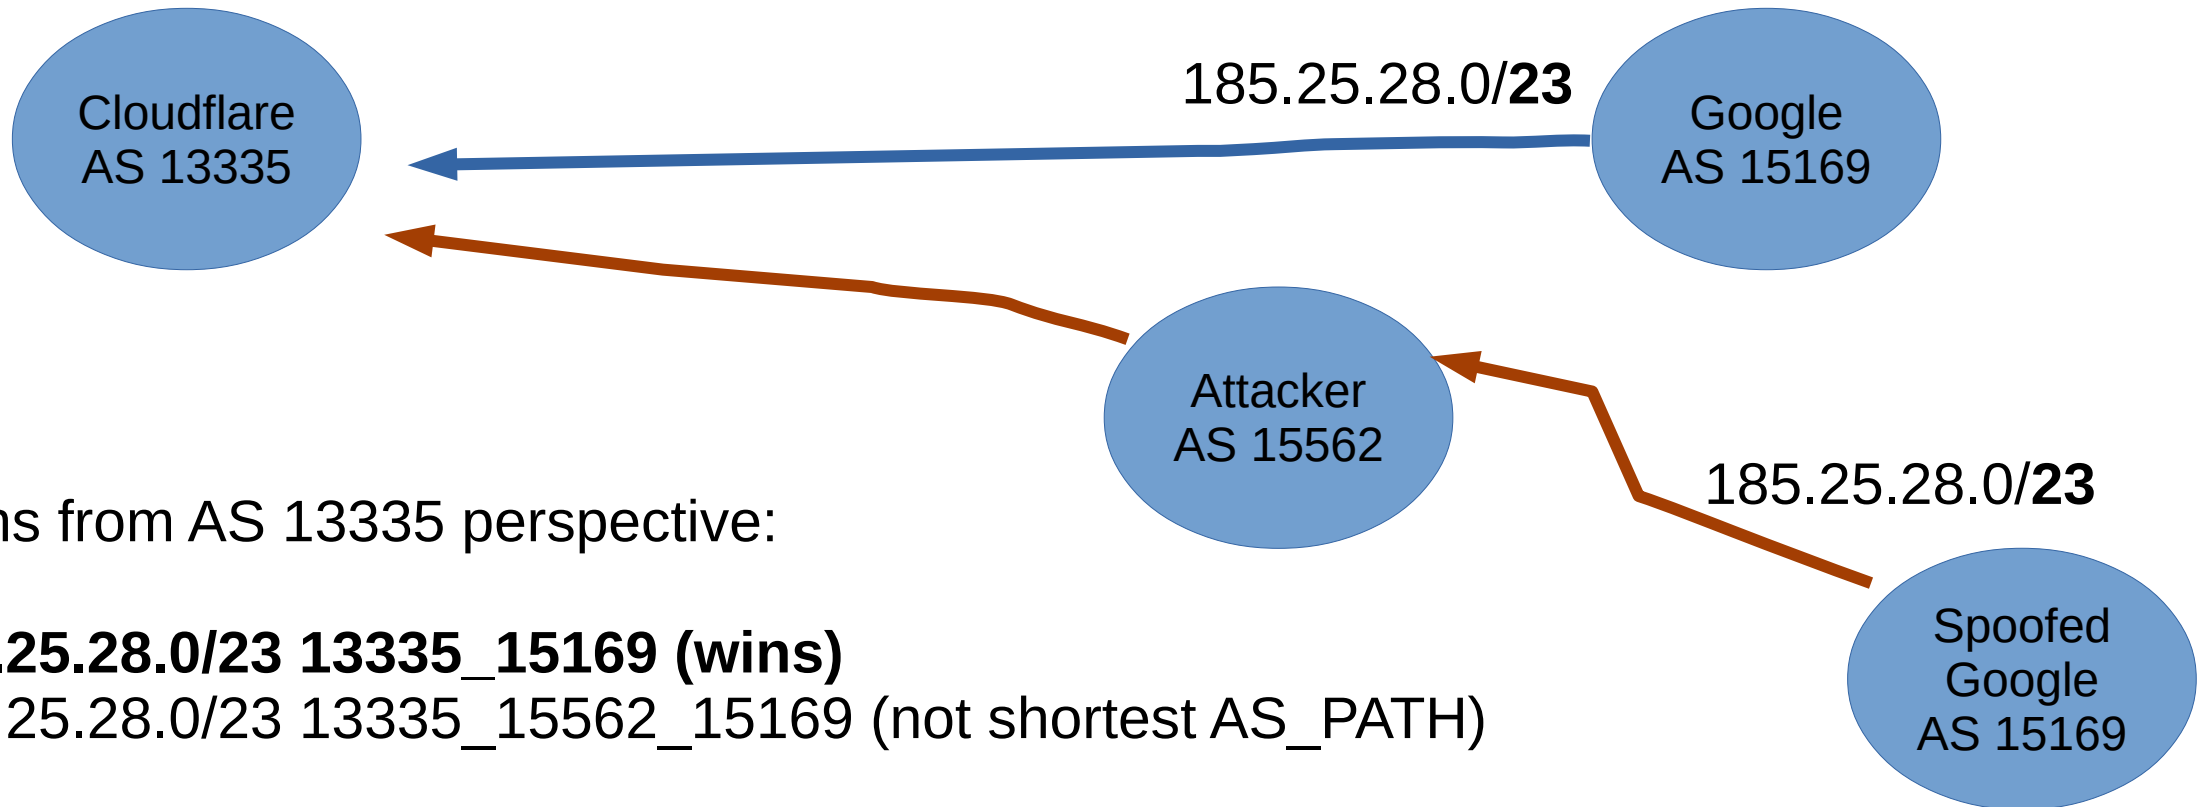
Paths from AS 13335 perspective:

185.25.28.0/23 13335_15169 (wins)

185.25.28.0/23 13335_15562 (rejected, wrong Origin ASN)

Change of tactics: spoof origin – NOT EFFECTIVE!

Cloudflare applying “invalid == reject”




“There is no healthy software ecosystem”



- RIPE NCC Validator v3 is works and actively maintained
- NLNetlabs is released a RPKI Cache Validator (Routinator 3000)
- Dragon Research RPKI Toolkit
- RPSTIR
- OpenBSD close to releasing a new RPKI 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!


RPKI based traffic analysis with pmacct

pmacct: one slightly more complex use-case



Pmacct RPKI capabilities

- RFC 6811 Origin Validation procedure is applied
- Mark traffic based on ROA status, without deploying RPKI in your network
- This helps us understand the effects of dropping RPKI Invalid announcements
- Version 1.7.3 - <https://github.com/pmacct/pmacct>




Timeline

- All ISPs, create [RPKI ROAs](#) now - it's easy!
- IXPs – start doing RPKI Origin Validation on your route servers **now**
- In 2019 RPKI data will be used to clean up IRR
- Hopefully the ARIN RPKI TAL situation will improve

Deployment update

- Cloudflare
- YYCIX


**YYCIX**
Canada
Internet Exchange
@yycix

Following

Good news! On Oct 7, the [#YYCIX](#) route servers started filtering prefixes which are RPKI ROA invalid. We are among the leaders in performing this validation -- probably the first [#IXP](#) in North America!

8:58 PM - 7 Oct 2018

11 Retweets 18 Likes



**Jerome Fleury**
@Jerome_UZ

Following

As of today, 75% of the [@cloudflare](#) PoPs (116/155) have RPKI strict validation enabled on all peering sessions. That's about 17,000 RPKI enabled peerings. Great work from [@lpoinsig](#)!



RPKI and BGP: our path to securing Internet Routing

This article will talk about our approach to network security using technologies like RPKI to sign Internet routes and protect our users and customers from route hijack...
blog.cloudflare.com

2:35 AM - 20 Dec 2018 from [San Francisco, CA](#)

48 Retweets 141 Likes



RPKI Deployment

- AT&T rejects invalids on peering sessions
- Nordunet rejects invalids on all EBGP sessions
- KPN / 286 rejects invalids on customer sessions
- Seacomm & Workonline will drop invalids per April 2019
- INEX
- AMS-IX
- DE-CIX
- France-IX
- Netnod (**soon!**)

Conclusion

