Roadmap for a more secure global Internet Routing System 2023-2028

Job Snijders
IEPG @ IETF 117

job@fastly.com

Fastly, OpenBSD, RIPE NCC Exec Board, GROW, RSSF, PeeringDB

Agenda

- Recap: where are we now?
- What’s problem space
- The holy trifecta: ROV + ASPA/OPEN + BGPsec
- Status update on where we are in our journey towards the holy trifecta
- Q & A

The last 20 years

2004 - RFC 3779 published: *X.509 Extensions for IP Addresses and AS Identifiers*
2006 - Department of Homeland Security publishes a road map for Secure Protocols for the Routing Infrastructure Initiative.
2006 - IETF Secure Inter-Domain Routing (SIDR) working group started
2012 - RFC 6480 published: *An Infrastructure to Support Secure Internet Routing*
2013 - RFC 6811 published: *BGP Prefix Origin Validation*
2016 - Peerlock popularized and deployed at various global Carriers
2017 - RFC 8212 published: *EBGP policies 'secure by default'*
2017 - RFC 8205 published: *BGPsec Protocol Specification*
2018 - RIPE community successfully closes password 'RPSL' loophole with completion of NWI-5
2019 - RIPE community successfully use RPKI data to clean-up stale IRR data as a continuous process: RIPE-731.
2020 - Global launch of RPKI Origin Validation (Telia, NTT, LINX, Telstra, HK-IX, GTT, Cogent, Amazon, many others)

https://bgpsec.net/history.html

Re-cap: where are we now? (compared to 2019 roadmap)

✓ - Globally available hierarchical public key infrastructure for IPs & ASNs

✓ - First “RPKI Application” launched: ROAs for Route Origin Validation (RPKI-ROV)

✓ - IRR cleanup: RIPE-NONAUTH (with RPKI filtering), ARIN-NONAUTH (deprecated)

✓ - IRRd version 4 has RPKI-ROV filtering for route:/route6: objects

✓ - RPKI-ROV Origin Validation: deployed by major ISP providers & IXPs

In fact, most todo items from the previous Roadmap (2019) have been cleared:

https://events.dknog.dk/event/4/contributions/37/attachments/15/25/DKNOG9_Snijders_Routing_security_roadmap1.pdf

What’s the problem space

Mis originations: the keys are so close to each other

Image source: https://commons.wikimedia.org/wiki/File:Typo2_iPhone_Keyboard_(16862375121).jpg

Route leaks

I didn’t mean to send you that…

https://www.kentik.com/blog/a-brief-history-of-the-internets-biggest-bgp-incidents/

Image source: https://commons.wikimedia.org/wiki/File:Fire_Hydrant_with_Defective_Temporary_Meter.jpg

Route Leak? AS11845 / Vox Telecom Ltd

Phil Lavin phil.lavin at vonage.com
Wed Jun 21 13:13:58 UTC 2023

- Previous message (by thread): Call for Nominations for Board of Trustees and Advisory Council
- Next message (by thread): Route Leak? AS11845 / Vox Telecom Ltd
- **Messages sorted by:** [date] [thread] [subject] [author]

Hi Folks,

Seeing traffic from AWS use1 (216.147.2.235) to UK (62.3.100.19) has been attempting to transit via AS11845 for the last few hours. Looks like a route leak from Vox->AWS at DE-CIX.

```
traceroute to 62.3.100.19 (62.3.100.19), 30 hops max, 60 byte packets
 1  ec2-3-236-61-157.compute-1.amazonaws.com (3.236.61.157)  2.505 ms  2.001
ms  1.988 ms
 2  240.0.228.64 (240.0.228.64)  0.348 ms  0.443 ms  0.440 ms
 3  240.0.228.89 (240.0.228.89)  0.334 ms  0.331 ms  0.430 ms
 4  240.192.54.146 (240.192.54.146)  0.260 ms  0.310 ms  0.307 ms
```

AS spoofing: will the real slim shady, please stand up?

Image source: still from Emimem - The Real Slim Shady video

AMERICAN HORROR STORY

AS Impersonation attacks are real!

<< tell a wild story >>

- Social engineering is used to get direct BGP connectivity
  - (often facilitated by IRR / fraudulent domain name registrations)
- BGP hijacks start
- Everyone ultra confused / super hard to debug

How to solve these things?

Origin Validation

ASPA +
RFC 9234

BGPsec

Image source: We Live - Oscar Lukefahr CM (2010)

Executive briefing on ROA / RPKI-ROV

Safety mechanism against fat-finger typo mis-configurations

- All relevant BGP implementations support RTR v0 and RPKI-ROV
  - 100s of bugs in BGP implementations got fixed:
- All major tier-1 implemented RPKI-ROV (invalid == reject) except for 3
  - Sparkle + Zayo + Deutsche Telekom still missing?
- All major IXPs do RPKI-ROV on their Route Servers
- Global Community is doing excellent job for outreach & education
- 2020 was the year RPKI-ROV became big

Todo: ROA duplication across RIRs (inter-RIR certification services)

Executive briefing on RFC 9234 - BGP OPEN Policy

Safety mechanism against leaking

- Doesn’t use the RPKI at all!
- Configure for each BGP session if you are the customer/peer/provider role
- A BGP Path Attribute is set based on the above
- OpenBGPD + BIRD + FRRouting support this (COTS BGP to come later)

```
neighbor "CustomerA" {
        remote-as 15562
        role customer
        neighbor 212.114.120.72
}
```

This is an OpenBGPD example, lot’s of magic and implied RPL in the above!

Executive briefing on ASPA - the timeline

SIDROPS has been working hard firing on all cylinders to produce ASPA

2023: first wave of OSS projects: rpki-client, routinator, OpenBGPD, StayRTR
2024: ASPA RFCs likely to be published. Start putting it in RFPs.
2025: (all?) RIRs start supporting creation of ASPA objects in their Web portals
2026: General availability of ASPA-support in COTS implementations
Many IXPs deploy ASPA-verification on their Route Servers
QA testing starts for large ISP deployments
2027: Nationwide ISPs / Tier-1s deploy ASPA verification

https://www.manrs.org/2023/05/estimating-the-timeline-for-aspa-deployment/

http://bgpsec.net/

Executive briefing on BGPsec benefits

- To combat spoofing, the concept of BGP in-band signatures is **unavoidable**
- Incremental deployment: first protect the most valuable sessions
  - (sessions where you assign routes a high LOCAL_PREF)
- Signing & Validation decoupled in BGPsec: you can opt to only sign
  - Helps reduce resource consumption in a few places
- Benefit in BGPsec for peering, as transit is the last-resort is trash anyway…
- BGPsec essentially is fully automated KYC (Know Your Customer)
  - Can view it as RPKI-assisted TCP-MD5 dance
- Perfect for Closed User Groups (CUGs)
  - IXP Route Servers
  - Blackhole servers

Myths and concerns about BGPsec

- Requires everyone to participate (not true!)
- Too slow (who cares: put it in the cloud, or only enable it on important peers)
- Doesn’t scale (prove it! What even is scale? HTTPS worked out in the end!)
- If you uppref an unsigned path, the unsigned path is used (yes!)
- Downgrade attack? (Just configure your router not to downgrade!)

I am *very* interested to make BGPsec a reality - I think it’s worth doing!

BGPsec deployment plan, we are at 5%

- ✓ RPKI Validators (rpki-client, Routinator)
- ✓ RPKI-To-Router implementation (StayRTR)
- ✓ Proof-of-concept BGP implementation (NIST-SRx)
- ✓ Signers: Krill, Dragon Labs rpki.net
- ❌ OpenBGPD (**NEED FUNDING!**), BIRD, FRRouting
- ❌ First real-world test deployments on Private Peering
- ❌ Commercial-of-the-shelf BGP implementations
- ❌ Deployment at IXP Route Servers / CUGs like Blackhole route servers

Overall status: Internet Routing Security is at 42%

The toolbox for routing is going to be:

- ASPA
  - coarse filter optimized for the general use-case
- BGP Open Policy (Only-to-customer attribute) - RFC 9234
  - extra per-route per-session-specific signal (if you also peer with your customers)
- BGPsec
- RPKI-ROV

Other useful RPKI applications:

- RPKI-signed Geofeed files
- RPKI-signed challenge/responses (RSC)
- GhostBuster Records (Point-of-Contacts for RPKI)

Todo list:

- Some IRR features need to be ported to RPKI (ASGroups/ASCones/ASSets, Prefixlist objects)
- **inter-RIR certification services for redundancy**

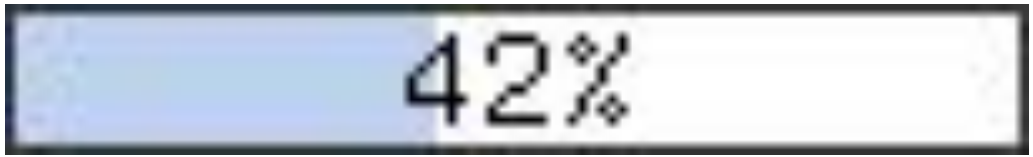

Origin Validation

ASPA +
RFC 9234

BGPsec

Image source: We Live - Oscar Lukefahr CM (2010)