Route Server Security and the Role of IXPs

Job Snijders

NTT Communications

job@ntt.net

Agenda

- Advantages of route servers
- Why security matters
- State of route servers around the world and close-by
- IX stories:
  - DE-CIX, AMS-IX, Seattle IX, France-IX, NL-IX, LINX, YYCIX
- Open source software:
  - IXP Manager, arouteserver, bgpq3, irrexplorer.nlnog.net
- Conclusion

Advantages of route servers

- Low maintenance aggregation point sessions
- Immediate value for newcomers.
- Debugging tool to have a sense what's going on at the IX

Further reading:

"Peering at Peerings: on the role of IXP route servers"
https://people.csail.mit.edu/richterp/imc238-richterA.pdf
"Internet Exchange BGP Route Server" – RFC 7947
"Internet Exchange BGP Route Server Operations" – RFC 7948

How a route server works

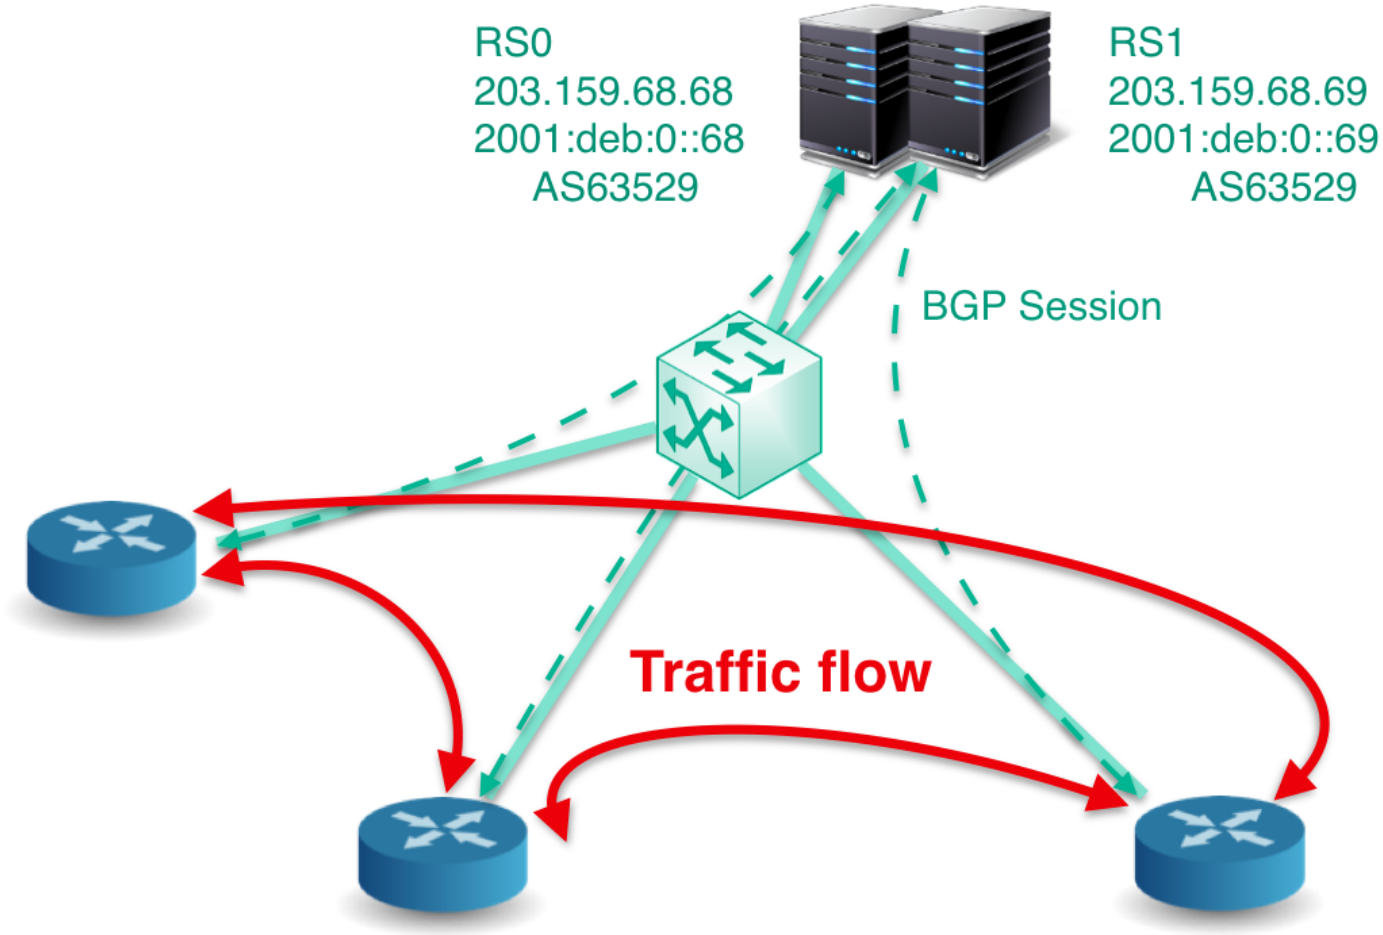

- Control-plane traffic is aggregated by the route server
- Data-plane traffic flows directly from participant to participant

Image created by bknix.co.th

Why security matters, for everyone

- Forcing malicious actors to leave a trail in the IRR helps fight crime
- Enforce basic hygiene: scrub bogon ASNs, bogon prefixes, etc
- Non-RS-participants can be affected: if someone leaks NTT prefixes to the Route Server, I won't be happy
- Level playing field between IXPs, internet is as strong as the weakest link, everyone benefits if everyone who can filter; filters.
- Bugs happen, BGP implementations may suddenly ignore filters
- Misconfigurations are easy to make, everyone has made typos

An IX's value increases as their trustworthiness increases

Why security matters, for everyone

Even if your policy is not to peer with Route Servers – insecure route servers can negatively affect your business operation!

- What if a mutual customer leaks your full table to the IXP RS?
- What if a malicious party plans a hijack and uses the route servers to obfuscate the trail?

As a customer or member of the IXP you are in an excellent position to provide your IXP with feedback about their route server product. When everyone asks their IXPs for security – we all benefit.

State of route servers at top largest IXPs

IXP name	Route server security state
DE-CIX	Secure
AMS-IX	Secure
LINX Lon/Man/Cardiff/Nova	Secure
IX.Br	Insecure (but working on it!)
MSK-IX	Secure
DATA-IX	Secure
NL-ix	Secure
Equinix	Secure
W-IX	Secure
Netnod	Secure
France-IX	Secure
Seattle IX	Secure
LONAP	Secure
INEX	Secure

More extensive overview: http://peering.exposed/

http://peering.exposed/ April 2018 status

- ~ 50 IXPs indicated they have per-customer filters
- ~ 15 IXPs are known not to have filters
- ~ 60 IXPs we lack data and don't know if they filter or not

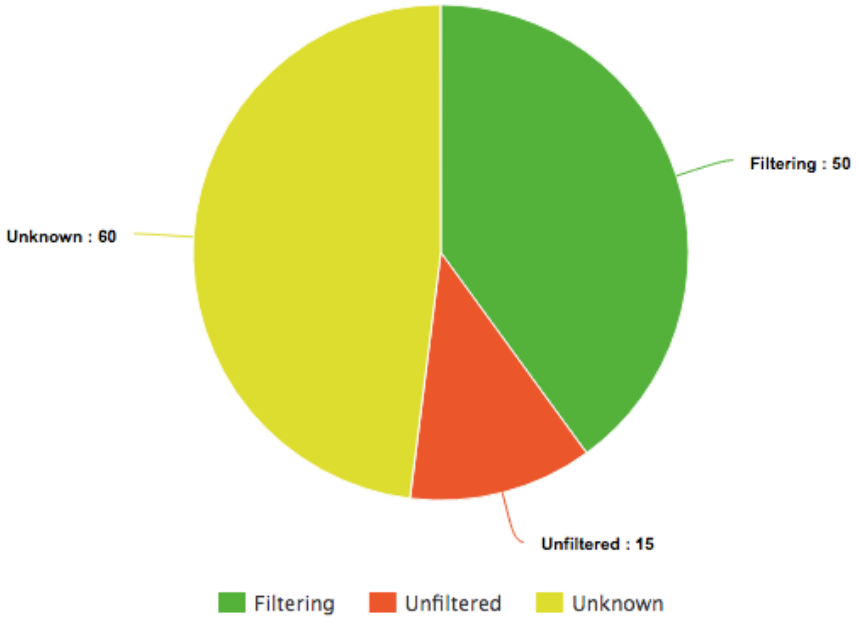

IX Stories – DE-CIX

- Started filtering in 2001 (16 years ago!)
  - Arnold Nipper wrote some sed, awk & /bin/sh to build per customer filters on zebra
- Now have sophisticated toolchain, and have open sourced parts of it:
  - bgperf RS performance measurement tool
  - Pbgpp (PCAP BGP parser)

Advice to other IXPs:

"Help your customers / participants to make effective and efficient use of the route servers. Ask them what they want and need. Whatever helps your participants to make a more sophisticated decision where to route traffic to the better."

IX Stories – FranceIX

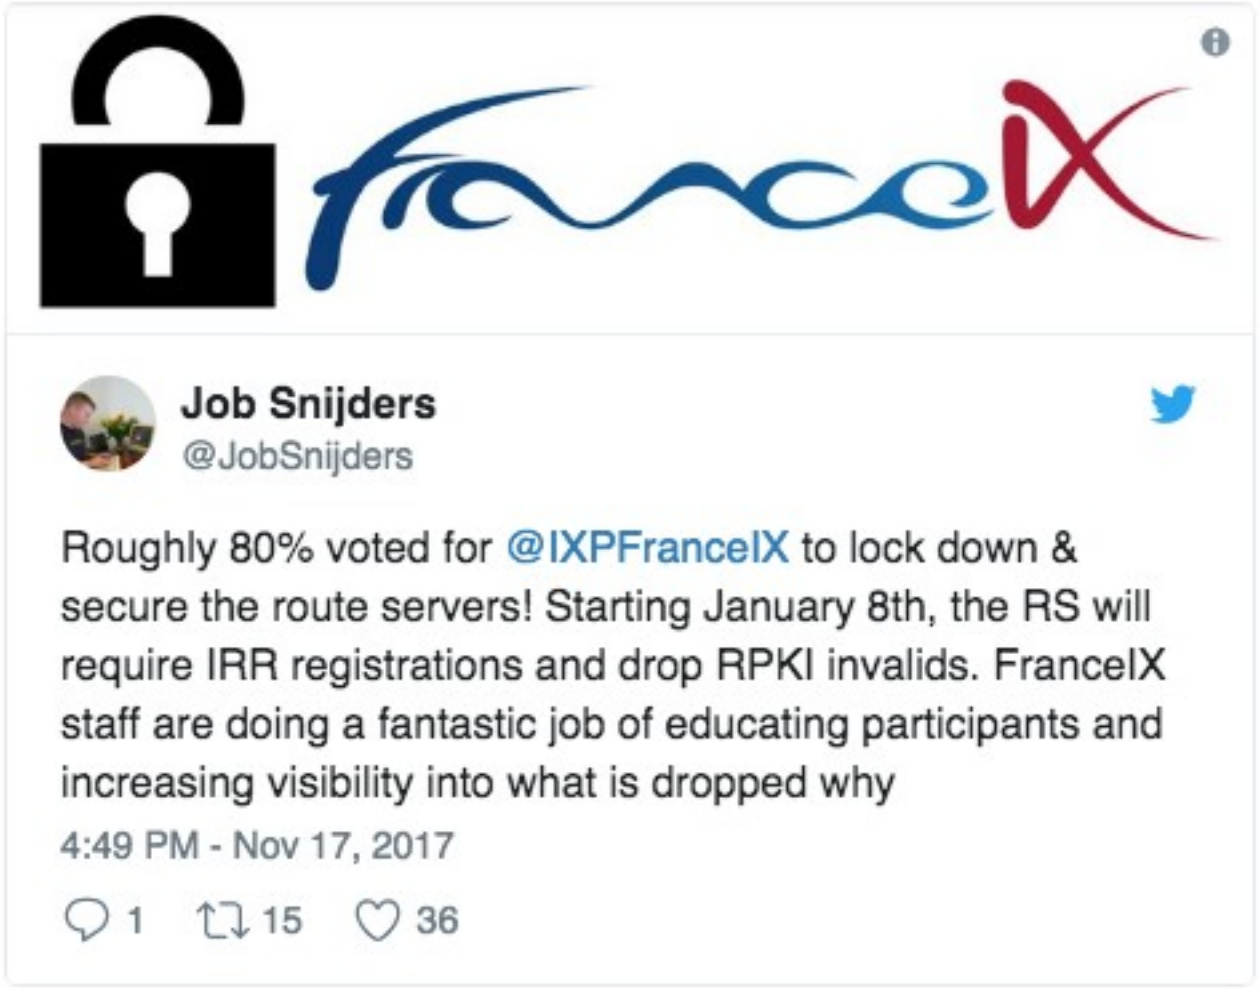

Job Snijders
@JobSnijders

Roughly 80% voted for @IXPFranceIX to lock down & secure the route servers! Starting January 8th, the RS will require IRR registrations and drop RPKI invalids. FranceIX staff are doing a fantastic job of educating participants and increasing visibility into what is dropped why

4:49 PM - Nov 17, 2017

1 15 36

Secure | https://lg.franceix.net/welcome/RS1+RS2/ipv4

Franceix.net - LG | Nodes: all RS1 RS2 | Protocols: ipv4 ipv6 | show informations

- show informations
- show route ...
- show route ... all
- show route ... (bgpmap)
- show route ROA Valid for ASN ...
- show route ROA Invalid for ASN ...
- show route ROA Unknown for ASN ...
- show route IRR Found for ASN ...
- show route IRR NOT Found for ASN ...
- show route IRR Found from ASN ...
- show route IRR NOT Found from ASN ...
- show all routes ROA Invalid
- show all routes IRR NOT Found

RS1 Statistics show route count

Routes ipv4	105826
Unique Routes ipv4	94975
Routes ipv6	25306
Unique Routes ipv6	22738

RS2 Statis

Route
Unique Route
Route
Unique Route

Graphics

This graphics represent the number of uniques prefixes advertised by FranceIX only)

Day

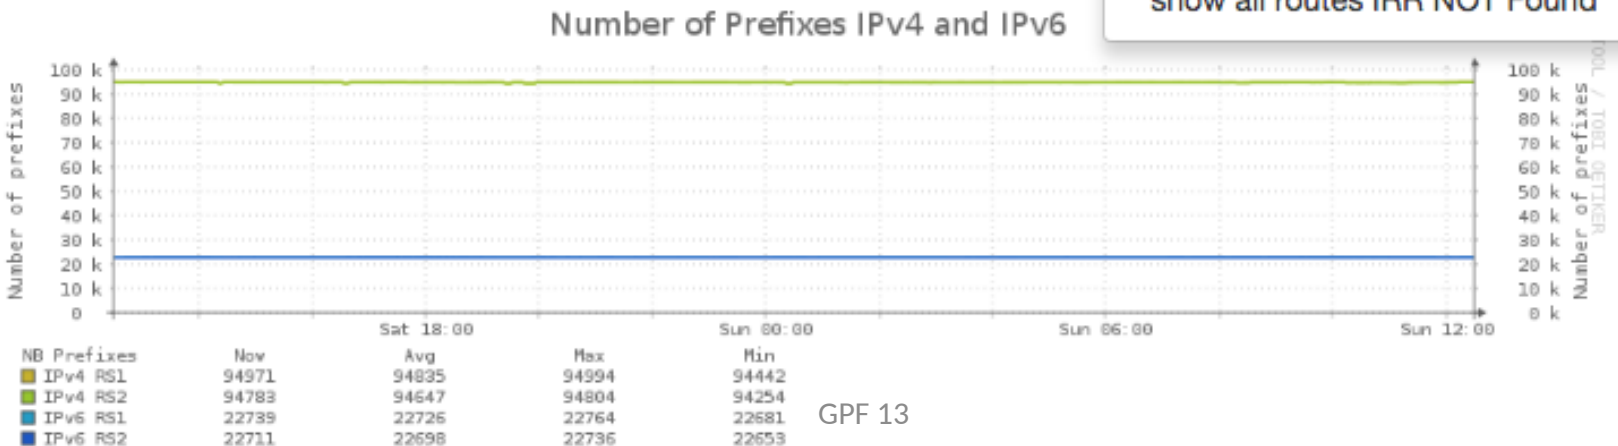

IX Stories – FranceIX

- December 2017/ January 2018: a looking-glass was implemented to provide insight
- December 2017/ January 2018: a testing environment was implemented to validate the application of the strict filtering and confirm there is no side-effect.
- 8 February 2018: maintenance for the implementation of the strict filtering on RS-MRS-1 (Marseille) and RS-PAR-1 (Paris).
- 15 February 2018: maintenance for the implementation of the strict filtering on RS-MRS-2 (Marseille) and RS-PAR-2 (Paris).

Full time line: https://blog.franceix.net/route-servers-filtering-policy-current-status-and-next-steps/

IX Stories – AMS-IX

- Converted from ‘insecure’ to ’secure by default’ in October 2017
- Participants can choose between four modes via a webportal:
  - “IRR + RPKI filtered”, “IRR filtered”, “RPKI filtered”, “No filter, only BGP community tagging (aka poison mode)”

Leadership worried about “traffic loss”, however:

*“**No traffic loss detected**, although advertised prefixes (with IRR+RPKI filtering) went from **~165K to ~68K**.”*

“we were quite surprised ourselves by the non-linear relation between prefixes and traffic”

Source: https://mailarchive.ietf.org/arch/msg/sidrops/Vf6r7EoRYkIbHOwjx1x_IQobq_I

IX Stories – LINX

- Converted from ‘insecure’ to ’secure’ on all its IXPs throughout December 2017 - February 2018 after extensive community consultation
- Datapoint from IXManchester: “The total IPv4 prefix count has dropped, as expected, by around 20% from previously 55,000 prefixes to now 42,000 prefixes.”
  - “A large part of the prefixes are learned from a single member and we are working on implementing additional validation criteria to improve in those cases.” (meaning RPKI and ARIN WHOIS)
- Traffic impact: No significant impact to the exchanges as a whole

IX Stories – YYCIX

YYCIX
Calgary
Internet Exchange

- Calgary, Canada, famous for the security research (OpenBSD, OpenSSH.. ;-)
- Runs route servers on OpenBGPD (the other IXPs mentioned use BIRD)
- 2 weeks to get IRR updated, project done in October 2017
- Lockstep migration: first migrate rs1 ⎵ help everyone based on rs1 data ⎵ flip the switch on rs2
- ~ 900 emails spent helping peers
- **No traffic loss**
- AS-SETs come from PeeringDB, Routing statements from IRR, RPKI & WHOIS
- Positive reactions from participants
- 3 fat finger routing errors, 2 redundancy issues diagnosed in first month

Open Source software – IXP Manager

IXP Manager is a full stack management system for Internet eXchange Points (IXPs) which includes an administration and customer portal; provides end to end provisioning; and both teaches and implements best practice. Maintained by the excellent INEX folks.

Produces: simple BIRD configurations, comes with full IXP management tool.

Downside: no support for RPKI, Registro.br, WHOIS

https://www.barryodonovan.com/2016/09/19/a-brief-history-of-ixp-manager

Open Source software - Arouteserver

Arouteserver is a Python tool to automatically build (and test) feature-rich configurations for BGP route servers. Written by Pier Carlo Chiodi.

Produces:

- Very feature rich BIRD *and* OpenBGPD configurations
- Parity between classic & large communities
- IRR, RPKI, ARIN WHOIS as whitelist (let customers choose where and how to register)
- fetches AS-SETs from PeeringDB (and/or from local database)
- easy to plug into existing portals / customer lists / management systems
- YYCIX is used as real-world test platform
- Active development, very reliable quality due to extensive regression testing
- **Arouteserver is what I would recommend people to use**

https://blog.apnic.net/2017/03/17/ixp-automation-made-easy-new-open-source-tool/

Filtering strategy recommendations

1. **Use PeeringDB** to find what AS-SET to use for what ASN (also show what is used and allow an override through a web portal)
2. **Reject** announcements that contain Bogon ASNs, Bogon prefixes
3. **Reject** announcements that contain ‘well-known transit-free’ networks anywhere in the AS_PATH: http://bgpfilterguide.nlnog.net/guides/no_transit_leaks/
4. **Reject** any announcements that are classified as “RPKI Invalid”
5. Generate a per-participant ***whitelist*** prefix-list of announcements using bgpq3 and **reject** any announcements for prefixes not part of that list.
6. Generate a per-participant ***whitelist*** as-path-filter based on the AS-SET using bgpq3, and **reject** any announcement originated by an ASN which is not part of the participant’s AS-SET.
7. **Visibility**: show in a web portal what announcements are rejected, and use BGP communities to attach a rejection reason to each such announcement for easy debugging.
8. **Never** make any ”they are too big to be filtered”-exception for any of the peers

Potential Future work

1. Enhance PeeringDB with some “**Never-Via-Routeservers**” flag, so networks can self-declare their ASN should never appear in AS_PATHS distributed by route servers? (as an alternative approach to ”block transit free networks”)
2. **Standardize what a ‘secure route server’ is in IETF BCP / RFC ?**
3. Promote the idea to **remove ability to receive unfiltered feeds** (aka filters should not be opt-in/opt-out – but always on)

Conclusion

- Many (both large and small) IXPs have demonstrated the ability to migrate to secure route servers in a matter of weeks
- There are a number of excellent open source tools readily available
- IXPs in Europe and US report no “loss of traffic”

If you need help converting your IXP RS to ‘secure’ – ask me!