Route Server Security & 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 closeby
- IX stories:
  - DE-CIX, AMS-IX, Seattle IX, France-IX, 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

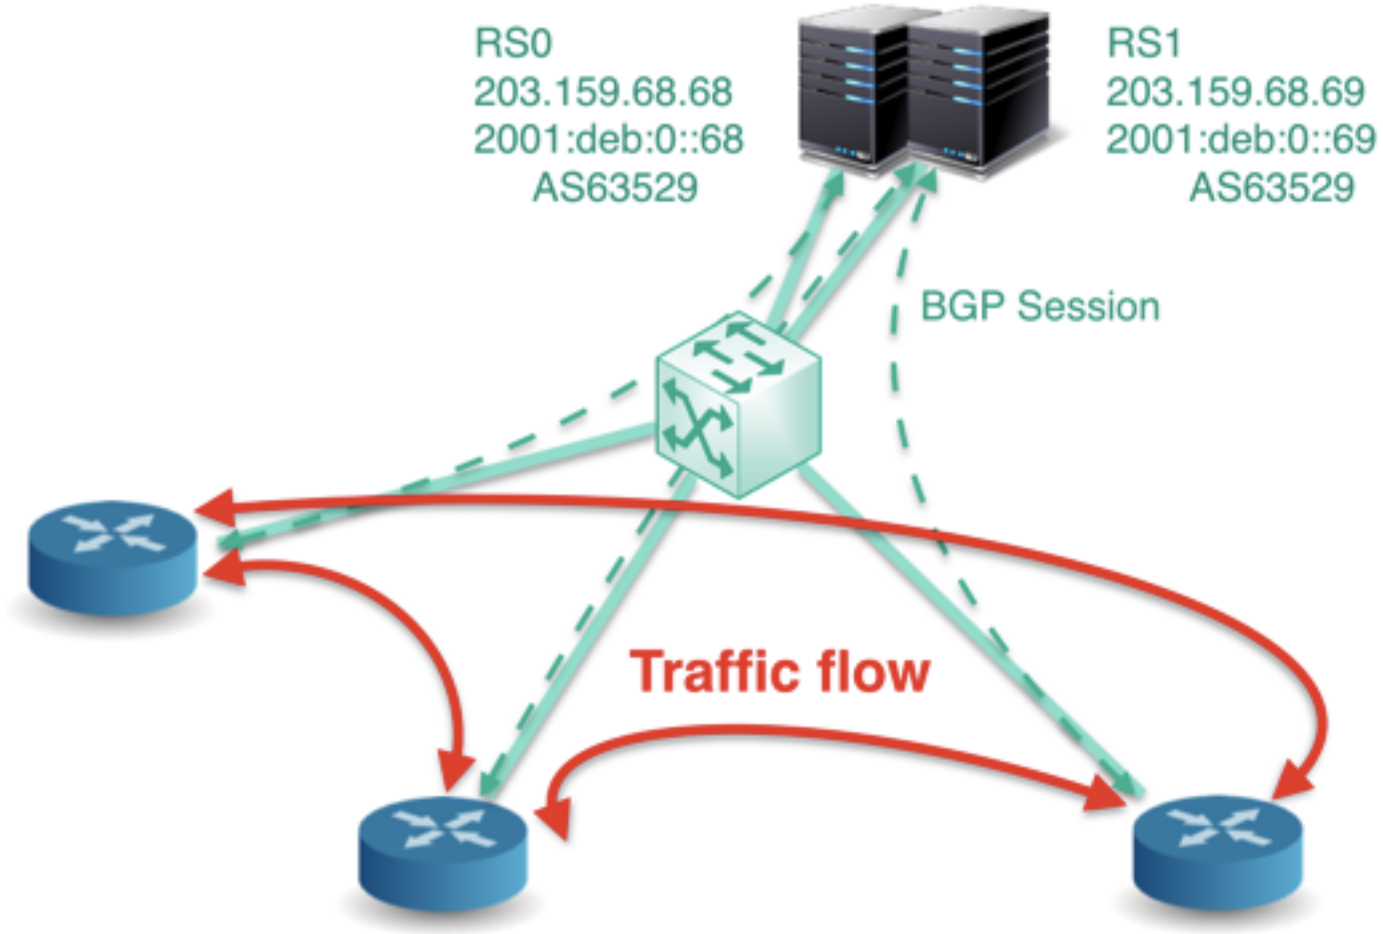

- Control-plane traffic is aggregated by the route server
- Data-plane traffic flows directly from participant to participant

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

State of route servers at top IXPs

IXP name	Route server security state
DE-CIX	Secure
AMS-IX	Secure
LINX	insecure
IX.Br	insecure
MSK-IX	Secure
DATA-IX	Secure
NL-ix	optional
Equinix	Secure
W-IX	Secure
Netnod	Secure
France-IX	Secure (per January 8th 2018)
Seattle IX	Secure
LONAP	Secure
INEX	Secure

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

IX Stories – Seattle IX

- Switched from unfiltered to secure route servers on February 23rd 2017
- Chris Caputo claims ~ 51 hours of development time
- Advice to other IXPs: stagger the work: ‘announce’ -> ‘filter rs1’ -> ‘filter rs2’
- No complaints in last 6 months, everyone appears to have updated their IRR
- Tutorial available for new folks: https://www.seattleix.net/irr-tutorial

Date	IPv4 on rs2	IPv6 on rs2	% rejected due to filter
24 jan	74066	18825	3.40% / 0.98%
23 feb	73418	19558	1.95% / 0.56%
25 feb	73945	19715	1.51% / 0.35%

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

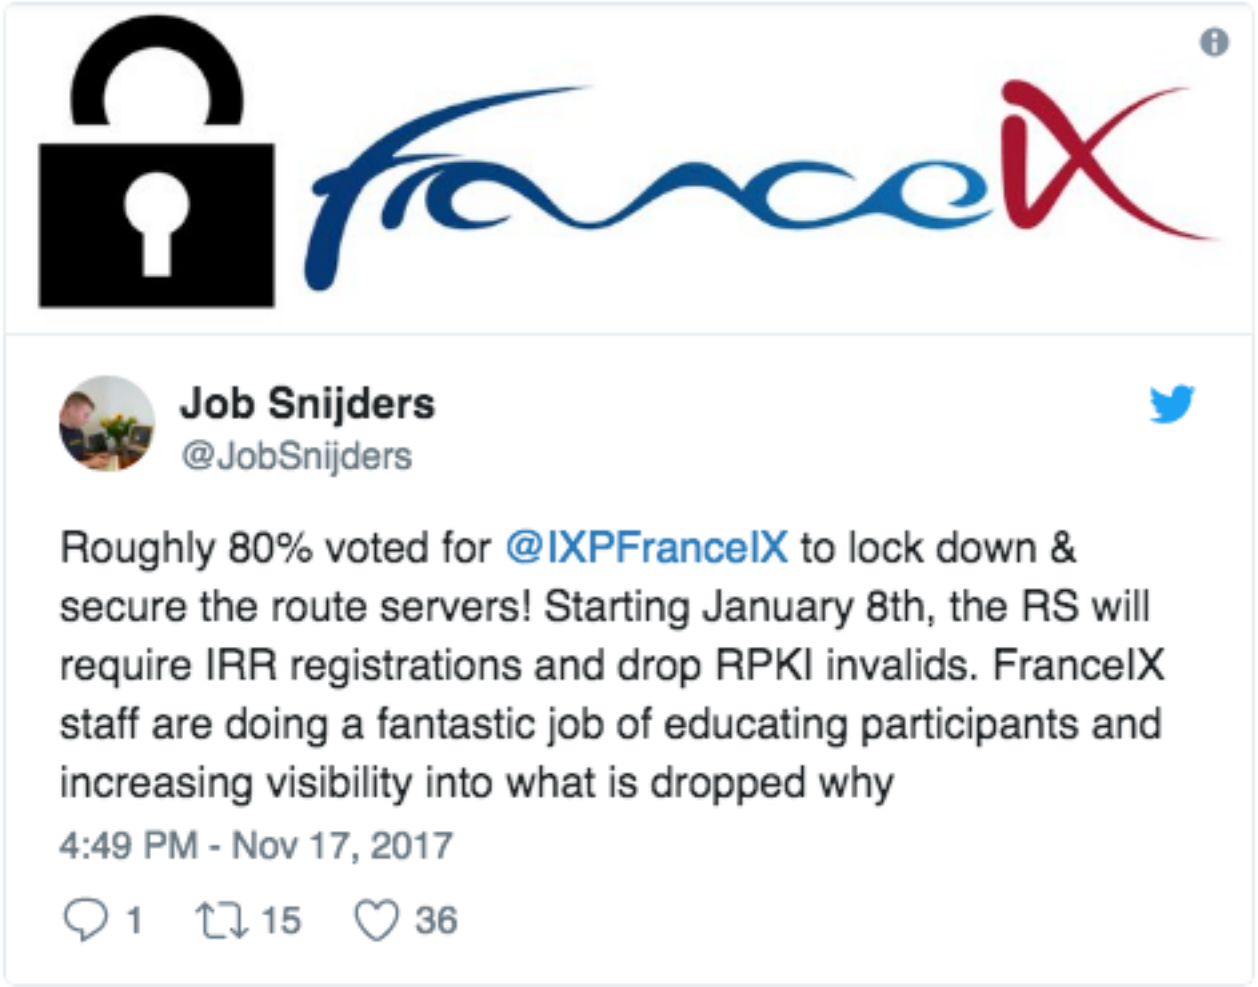

Secure | https://lg.franceix.net/welcome/RS1+RS2/ipv4

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

Number of Prefixes IPv4 and IPv6

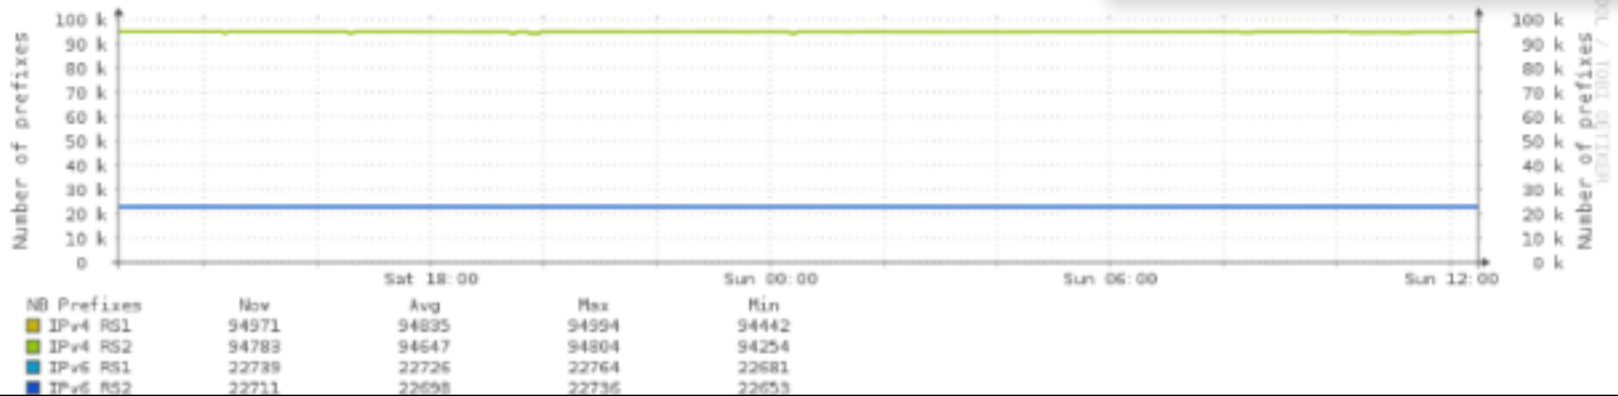

IX Stories – AMS-IX

- Years long struggle between IX participants and AMS-IX leadership
- 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 – 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.

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 LINX to use**

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

Filtering strategy recommendations

- **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)
- **Reject** announcements that contain Bogon ASNs, Bogon prefixes
- **Reject** announcements that contain 'well-known transit-free' networks anywhere in the AS_PATH: http://bgpfilterguide.nlnog.net/guides/no_transit_leaks/
- **Reject** any announcements that are classified as "RPKI Invalid"
- Generate a per-participant ***whitelist*** prefix-list of announcements using bgpq3 and **reject** any announcements for prefixes not part of that list.
- 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.
- **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.

LINX…

Customers have been asking LINX to improve their service offering for years (especially since the October 2016 ”Bofinet/Virgin” leak).

November 10th 2016 – mikeh@linx.net:

> *“We do already have the required projects in our roadmap for 2017, which will bring prefix filtering/validation to the LINX Route Servers.*
>
> *Our target dates for this is Q2 2017, but hopefully, we will be able to roll it out in different phases, with basic prefix validation and option to filter earlier in 2017.”*

LINX…

November 11th 2016 – richard@linx.net:

> *"To clarify on our position, as Mike has noted this it is on our technology roadmap for 2017 and we will indeed, review, evaluate and schedule for 2017 deployment, providing we meet all the criteria to release to the membership as a product. This includes the technical specification and equally a policy review step, so I want to set expectation that this will have a wider review with our membership before we go live. It would be great to discuss this further at LINX95. "*

LINX…

February 2nd 2017 – mikeh@linx.net:

> *"As clarified in my other response to the email from Job, the technical development of such a service is on our roadmap for 2017. But a policy review step will also be required before we can commit to launch.*
>
> *We will cover this in detail, including the request for further feedback from members during the upcoming LINX96 meeting in February [red: 2017]."*

LINX…

February 3rd 2017 – eric.loos@bics.com replies to Mike Heller:

> *"… the process you describe seems very drawn out (maybe I am wrong) and bureaucratic, so I wonder how efficient this is all going to be for what is essentially a no-brainer in terms of service improvement …*
>
> *… Just for my understanding, what arguments do you believe exist to *__not__* filter on the route-server?"*

LINX…

LINX95 November 2016 – ?

LINX96 February 2017 – Update from John Souter

LINX97 May 2017 – Update from Mike Hellers

LINX98 August 2017 – Update from Mike Hellers (based on incorrect data from FranceIX)

LINX99 November 2017 ← **WE ARE HERE NOW**

Conclusion

- Many (both large and small) IXPs have demonstrated the ability to migrate to secure route servers in weeks
- There are a number of excellent open source tools readily available
- No opposition to filtering was voiced on ops@linx.net in any of the relevant threads
- No IXP has reported issues related to “loss of traffic”

So… I’m sorry to report: ***LINX are lacking***

- LINX is the largest IXP in the world with unfiltered route servers
- Unreasonable timeline
- Lack of transparency on how filters will be generated
- Offers to assist with route server engineering have not been acted upon
- This results in a continued failure to protect its participants from misconfigurations