



Together, your Internet, even better



Arnaud Fenioux
France-IX GM-2017

FrancelIX route servers
IRR and RPKI/ROA

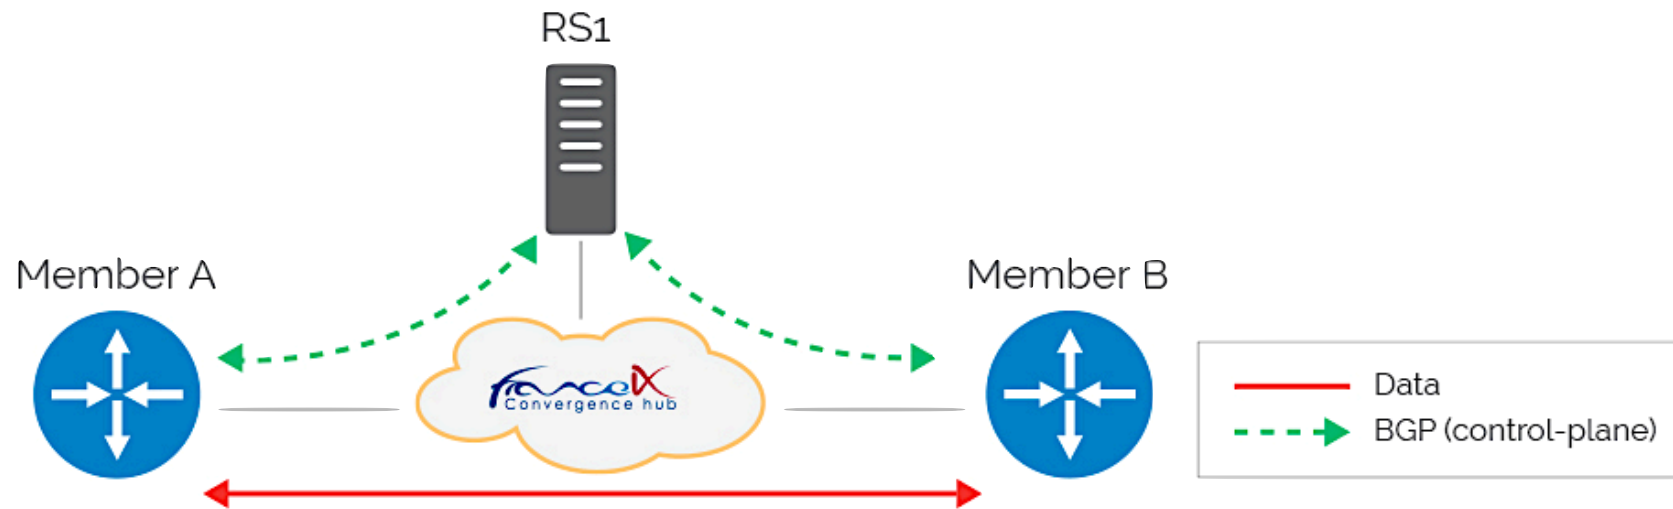
France-IX
route servers

- **Definition and benefits**
- IRR and RPKI/ROA implementation
- Facts and figures
- Use case

Definition

Route Servers are

- Multi-Lateral Peering Exchange
- Members automatically exchange routes with other members through a single BGP session
- There is no data going through the RS, it is only used to aggregate BGP information



Benefits

Routes servers are a popular service at IXPs

92% of France-IX's community use them in **Paris**

87% of France-IX's community use them in **Marseille**

Main benefits for the peers:

- Less BGP sessions to configure
- Quick and easy way to start getting lots of routes
- Easily tunable using BGP communities
- No need to make multiple peering arrangements with other members

<blink> **SAVE TIME!** </blink>

Exceptions

Can be considered as a SPOF

Some of the routing intelligence is out of the NetOps control

- ➔ Need trust into the IXP

Selective announcement may need some tweaking to keep symmetrical paths

Peers ASN will vary and increase with time :

- ➔ Adds some new destinations though the IXP

- ➔ Might not be wanted if you have strict peering policy or fine traffic tuning

- ➔ Some CDNs prefer to establish directly bilateral BGP peering

France-IX
route servers

- Definition and benefits
- **IRR and RPKI/ROA implementation**
- Facts and figures
- Use case

BGP communities

We can check transitivity with IRR and AS-SET

AS-SET -> AUT-NUM -> ROUTE(6)

ROA used to check origin AS only

Here are the communities we use to tag routes:

51706:65012 = Prefix has ROA status: **VALID**

51706:65022 = Prefix has ROA status: **INVALID**

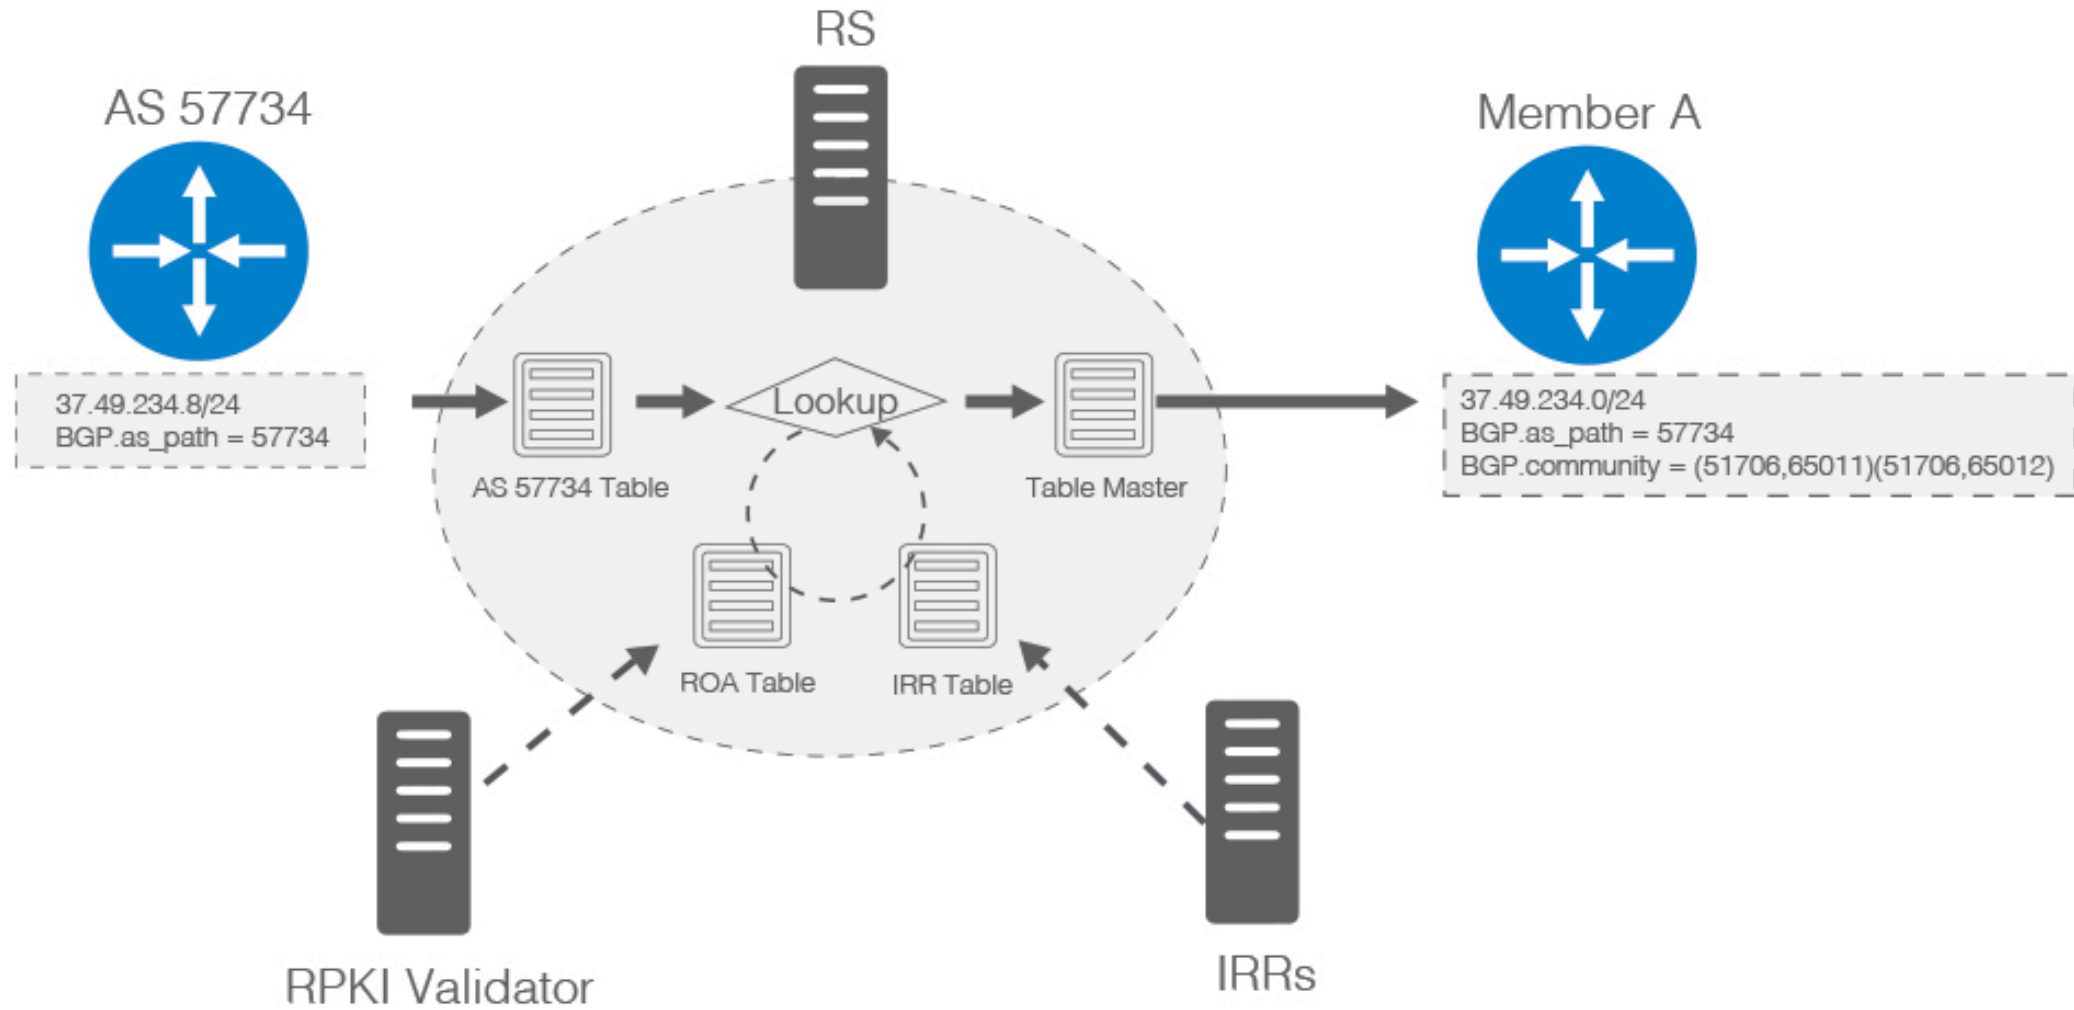
51706:65023 = Prefix has ROA status: **UNKNOWN**

51706:65011 = Prefix is **present** in an AS's announced **AS/AS-SET**

51706:65021 = Prefix is **not present** in an AS's announced **AS/AS-SET**

<https://www.franceix.net/en/technical/france-ix-route-servers/>

IRR and RPKI/ROA validation



<https://www.franceix.net/en/technical/france-ix-route-servers/>

IRR Data source

IRR Data fetched from **rr.ntt.net** with **BGPQ3**

➔ because of **IRR lockdown**

What is an IRR lockdown?

- Only honor route objects when they come from the right data source AND have been properly authenticated
- Ignore route objects covering the “locked down” IRR if they come from elsewhere

```
./bgpq3 -h rr.ntt.net -S RIPE,APNIC,AFRINIC,ARIN,NTTCOM,\  
ALTDB,BBOI,BELL,GT,JPIRR,LEVEL3,RADB,RGNET,SAVVIS,TC \  
-A -b -6 -l pfx_table_as57734 AS57734
```

```
pfx_table_as57734 = [  
  2001:7f8:54::/48,  
  2a00:a4c0::/32 ];
```



IRR Lockdown - RIPE70 - job@ntt.net

PeeringDB

AS-SET

PeeringDB Search here for a network, IX, or facility. as51706

Advanced Search

France-IX - IXP **Gold Sponsor** Edit

Organization	France-IX
Also Known As	FranceIX
Company Website	http://www.franceix.net
Primary ASN	51706
IRR Record	AS51706:AS-MEMBERS
Route Server URL	https://www.franceix.net/en/technical/france-ix-route-servers/
Looking Glass URL	https://tools.franceix.net/looking-glass
Network Type	Route Server
IPv4 Prefixes	0
IPv6 Prefixes	0
Traffic Levels	Not Disclosed
Traffic Ratios	Balanced
Geographic Scope	Regional
Protocols Supported	<input checked="" type="checkbox"/> Unicast IPv4 <input type="checkbox"/> Multicast <input checked="" type="checkbox"/> IPv6
Last Updated	2017-09-27T09:19:39Z
Notes	FranceIX is running two separated IXP based in Paris and Marseille. The IP provided on this page are these of our Route-Servers.

Your AS-SET ***ONLY***

- **No URL**
- **No text like :**
 - RIPE:: AS-SET
 - EUROPE: AS-SET
- **No bullshit**

RPKI/ROA data source

We are using the **RIPE RPKI Validator**

<https://www.ripe.net/manage-ips-and-asns/resource-management/certification/tools-and-resources>



RIPE NCC

RIPE NETWORK COORDINATION CENTRE

But the ARIN TAL is missing...
Remember to install it by yourself!

<https://www.arin.net/resources/rpki/tal.html>


France-IX route servers

- Definition and benefits
- IRR and RPKI/ROA implementation
- **Facts and figures**
- Use case

Paris

sh ip bgp sum

- **298** IPv4 BGP sessions established (+34 in 1 year)
- **204** IPv6 BGP sessions established (+20 in 1 year)
- **90 603** Unique IPv4 Routes (+16 839 in 1 year)
- **22 321** Unique IPv6 Routes (+ 5 247 in 1 year)



France-IX Services (c) 2017

Updated : 2017-09-27 15:08:01

Paris

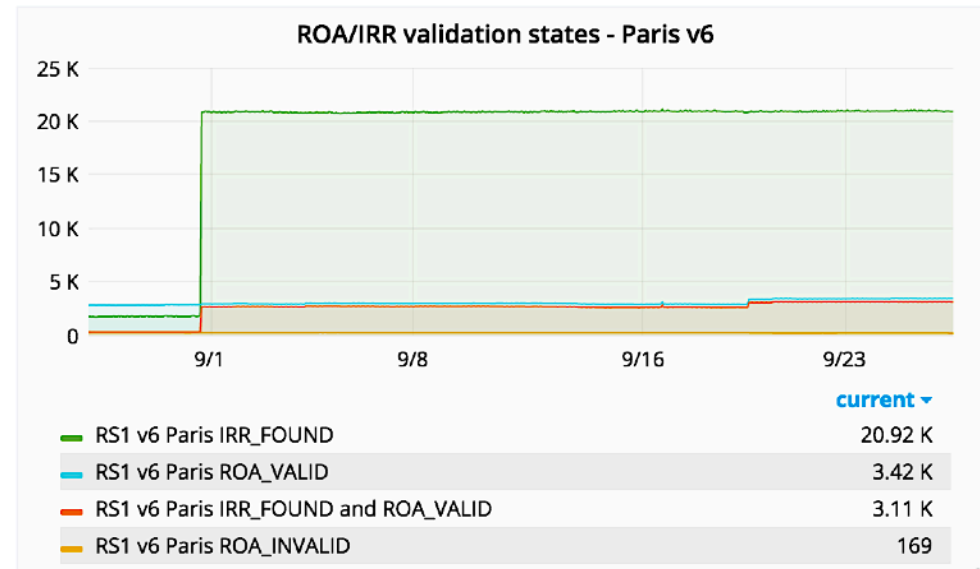
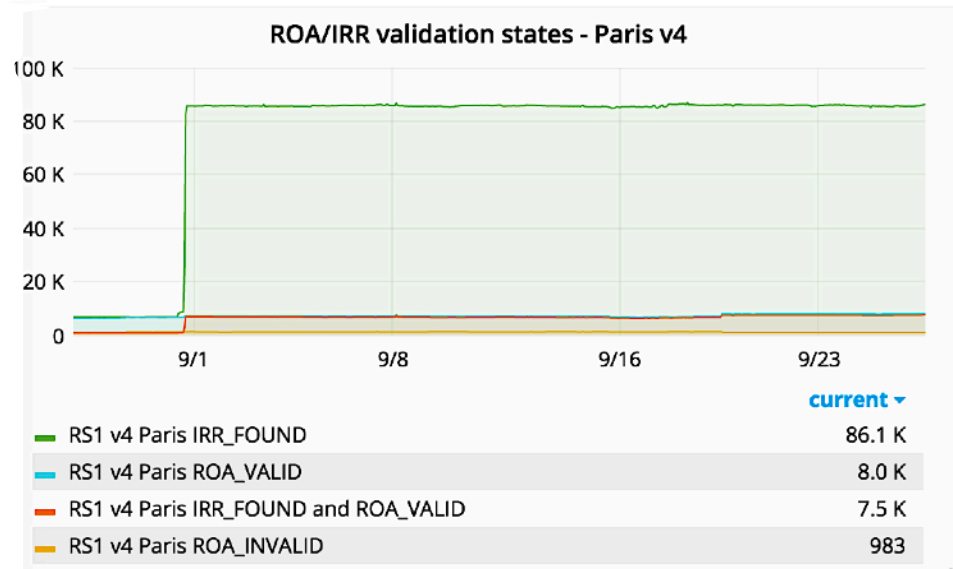
IRR and RPKI/ROA

IPv4

- 95% IRR FOUND
- 8.8% ROA VALID
- 8.3% ROA VALID + IRR FOUND
- 1.1% ROA INVALID

IPv6


- 94% IRR FOUND
- 15% ROA VALID
- 14% ROA VALID + IRR FOUND
- 0.8% ROA INVALID



Marseille

sh ip bgp sum

- **31** IPv4 BGP sessions established (+7 in 1 year)
- **25** IPv6 BGP sessions established (+5 in 1 year)
- **54 232** Unique IPv4 Routes (+11 127 in 1 year)
- **21 585** Unique IPv6 Routes (+ 5 291 in 1 year)



France-IX Services (c) 2017


Updated : 2017-09-27 15:17:01

Marseille

IRR and RPKI/ROA

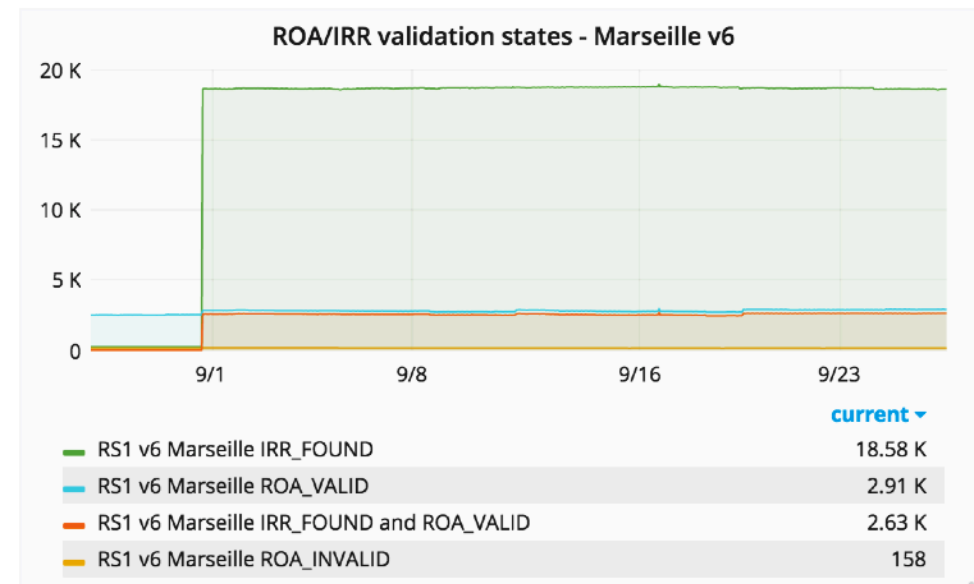
IPv4

- 90% IRR FOUND
- 5.0% ROA VALID
- 4.8% ROA VALID + IRR FOUND
- 0.8% ROA INVALID



IPv6

- 87% IRR FOUND
- 14% ROA VALID
- 12% ROA VALID + IRR FOUND
- 0.7% ROA INVALID





France-IX
route servers

- Definition and benefits
- IRR and RPKI/ROA implementation
- Facts and figures
- **Use case**

ROA and blackholing

- ROA are very easy to setup via the RIPE web interface (<https://my.ripe.net/#/rpki>)
- Beware of Max length if you plan to use the blackholing service (i.e /32 or /128)
- Blackholed prefix **MUST** pass IRR check to be accepted on RS (we do filtering ©)

The screenshot shows the RIPE RPKI Dashboard. At the top, it displays '5 CERTIFIED RESOURCES' and 'NO ALERT EMAIL CONFIGURED'. Below this, there are two main sections: '3 BGP Announcements' (with 3 Valid, 0 Invalid, and 0 Unknown) and '3 ROAs' (with 3 OK and 0 Causing problems). The dashboard is divided into three tabs: 'BGP Announcements', 'Route Origin Authorisations (ROAs)', and 'History'. The 'ROAs' tab is active, showing a table of Route Origin Authorisations. The table has columns for 'AS number', 'Prefix', 'Most specific length allowed', and 'Affects'. There are also buttons for 'Discard Changes', 'Delete ROAs', 'Causing Problems', 'Not Causing Problems', and '+ New ROA'.

AS number	Prefix	Most specific length allowed	Affects
<input type="checkbox"/> AS57734	2001:7f8:54::/48	48	1
<input type="checkbox"/> AS57734	2a00:a4c0::/32	32	1
<input type="checkbox"/> AS57734	37.49.232.0/21	21	1

Conclusion

Filtering prefixes on Route Servers :

- is “good for the internet”
- forces users to update their IRR and ROA records
- can lead to reject valid prefixes due to partial coverage (because some big ISP have too many LIR and records or they don't even know how to manage them)

IXP are working toward an effective filtering solution, enabling secure BGP announcements between members.



Merci !



Together, your Internet, even better



twitter.com/ixpfranceix



facebook.com/ixpfranceix



youtube.com/user/TheFranceIX



linkedin.com/company/france-ix

Arnaud FENIOUX

@afenioux