IXP Route Servers with RPKI and IXP Manager

APRICOT 2019, Daejeon, South Korea, Feb 2019

Barry O'Donovan
INEX & IXP Manager
@ComePeerWithMe / @barryo79

INEX
@ComePeerWithMe

Following

We demonstrated our true love for the security of INEX and the route servers, in the early hours of this morning and have given the gift 🎁 of RPKI 💝. Second half of implementation is now complete. If you love it, secure it with RPKI. #RPKILOVE ❣️

> **INEX** @ComePeerWithMe
> You'll have noticed we are big fans of RPKI at INEX and are *that* person who is so fond of something that they can't stop talking about it ... at any industry conference we can get to. Well, we are fully behind it and we have now gone "all in". #RPKIdayatINEX #Comejoinus 1/7
>
> Show this thread

8:19 am - 14 Feb 2019

INEX

RPKI at IXPs

IXP Manager

- An INEX project
- Full-stack management system for IXPs
- FOSS - GPL v2 license
- Complete route server automation
- In use at ~70 IXPs worldwide

https://www.ixpmanager.org/

github.com/inex/IXP-Manager

INEX

RPKI at IXPs

IRRDB vs. RPKI ROAs

```
route6:          2001:db8::/32
descr:           Example IPv6 route object
origin:          AS65500
created:         2006-07-12T16:11:58Z
last-modified:   2011-02-22T15:58:03Z
source:          SOME-IRRDB
```

```
route:           192.0.2.0/24
descr:           Example IPv4 route object
origin:          AS65500
created:         2004-12-06T11:43:57Z
last-modified:   2016-11-16T22:19:51Z
source:          SOME-IRRDB
```

INEX

RPKI at IXPs

ROAs - Route Origin Authorisations

- A cryptographically secure replacement for route[6] objects
- Adds maximum prefix length
- Yields route origin triplets that have been validated

```
( Origin AS, Prefix        ,  Max Length )
( AS65500,   2001:db8::/32,   /48        )
( AS65501,   192.0.2.0/24,    /24        )
```

RPKI at IXPs

Validating BGP Routing with RPKI-RTR

- A cache server *(validator)* does the cryptographic heavy lifting
- Routers receive and maintain the set of ROAs via RPKI-RTR from the cache
- RPKI gives three validation results: VALID, INVALID, UNKNOWN

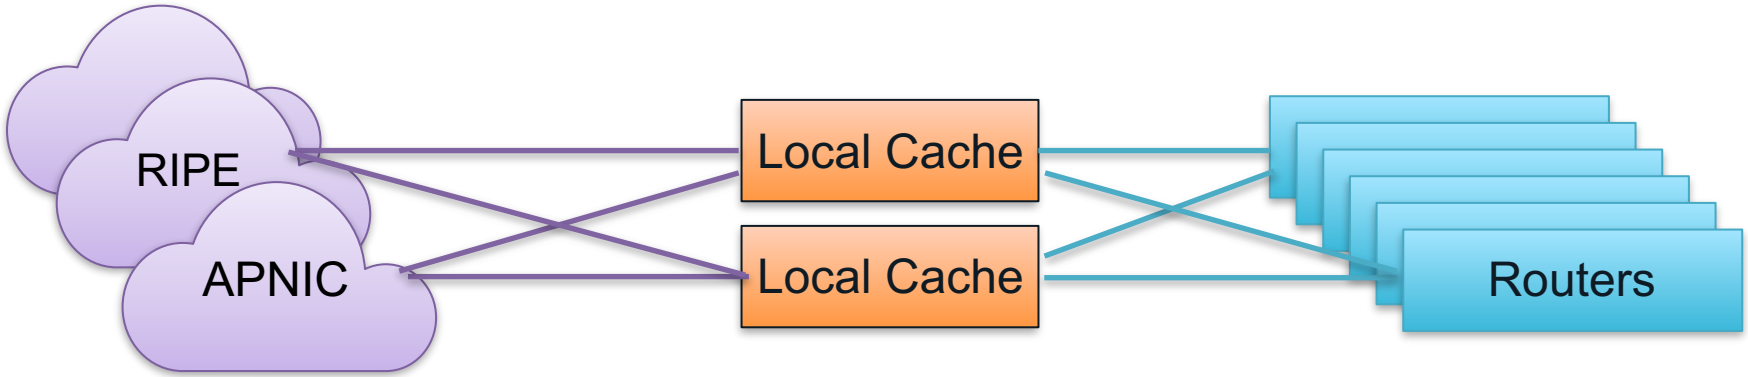

IXP Manager v5 and Route Servers

INEX

RPKI at IXPs

Route Server Refresh at INEX & IXP Manager

- RPKI just one element
- Upgrade configuration from Bird v1.6 to Bird v2.0 (retain separate protocol daemons)
- Complete rewrite of filtering workflow
  - Large communities used extensively within the route server
- Upgrade Bird's Eye[1] for Bird v2 BGP
- Overhaul IXP Manager looking glass

1. A secure micro service for querying Bird - https://github.com/inex/birdseye

INEX

RPKI at IXPs

Bird v1 to v2 Changes

- RPKI-RTR supported
- Collapsed separate daemons for IPv4 and IPv6 into a single daemon
  - master route table becomes master4 / master6
  - new protocol blocks: ipv4 { ... } / ipv6 { ... }
- Other very minor configuration changes

INEX

RPKI at IXPs

Bird v1 to v2 Changes

```
listen bgp address 192.0.2.8;

protocol bgp pb_as112_vli249_ipv4 {
        description "AS112";
        local as routerasn;
        source address 192.0.2.8;
        neighbor 192.0.2.6 as 112;
        import all;
        export none;
        table master;
}
```

```
protocol bgp pb_as112_vli249_ipv4 {
    description "AS112";
    local as routerasn;
    source address 192.0.2.8;
    strict bind yes;
    neighbor 192.0.2.6 as 112;
    ipv4 {
        import all;
        export none;
        table master4;
    };
}
```

Side note

RPKI at IXPs

Standard IX Route Server Community Filters

Description	Community	Large Community
Prevent announcement of a prefix to a certain peer	`0:peer-as`	`65500:0:peer-as`
Announce a prefix to a certain peer	`65500:peer-as`	`65500:1:peer-as`
Prevent announcement of a prefix to all peers	`0:65500`	`65500:0:0`
Announce a prefix to a all peers	`65500:65500`	`65500:1:0`

(Where we assume 65500 is the route server ASN)

INEX

RPKI at IXPs

Route Server BGP Community Usage

65500:1101:* are filtered

Side note

Description	Large Community
RPKI Valid	65500:1000:1
RPKI Unknown	65500:1000:2
IRRDB Valid	65500:1001:1
…	…

Description	Large Community
Bogon Prefix	65500:1101:3
IRRDB Invalid	65500:1101:9
RPKI Invalid	65500:1101:13
…	…

1. https://github.com/euro-ix/rs-workshop-july-2017/wiki/Route-Server-BGP-Community-usage

RPKI at IXPs

IXP Manager v5 Route Server Filtering

1. Small prefixes (default is > /24 / /48 for ipv4 / ipv6)
2. Martians / bogons
3. Ensure at least 1 ASN and <= 64 ASNs in path
4. Ensure peer AS is the same as first AS in the prefix's AS path
5. Prevent next-hop hijacking
6. Filter known transit networks
7. Ensure origin AS is in set of ASNs from member AS-SET
8. RPKI:
    - Valid -> accept
    - Invalid -> drop
9. RPKI Unknown -> revert to standard IRRDB prefix filtering

INEX

RPKI at IXPs

IXP Manager v5 Bird Topology - Import From Member

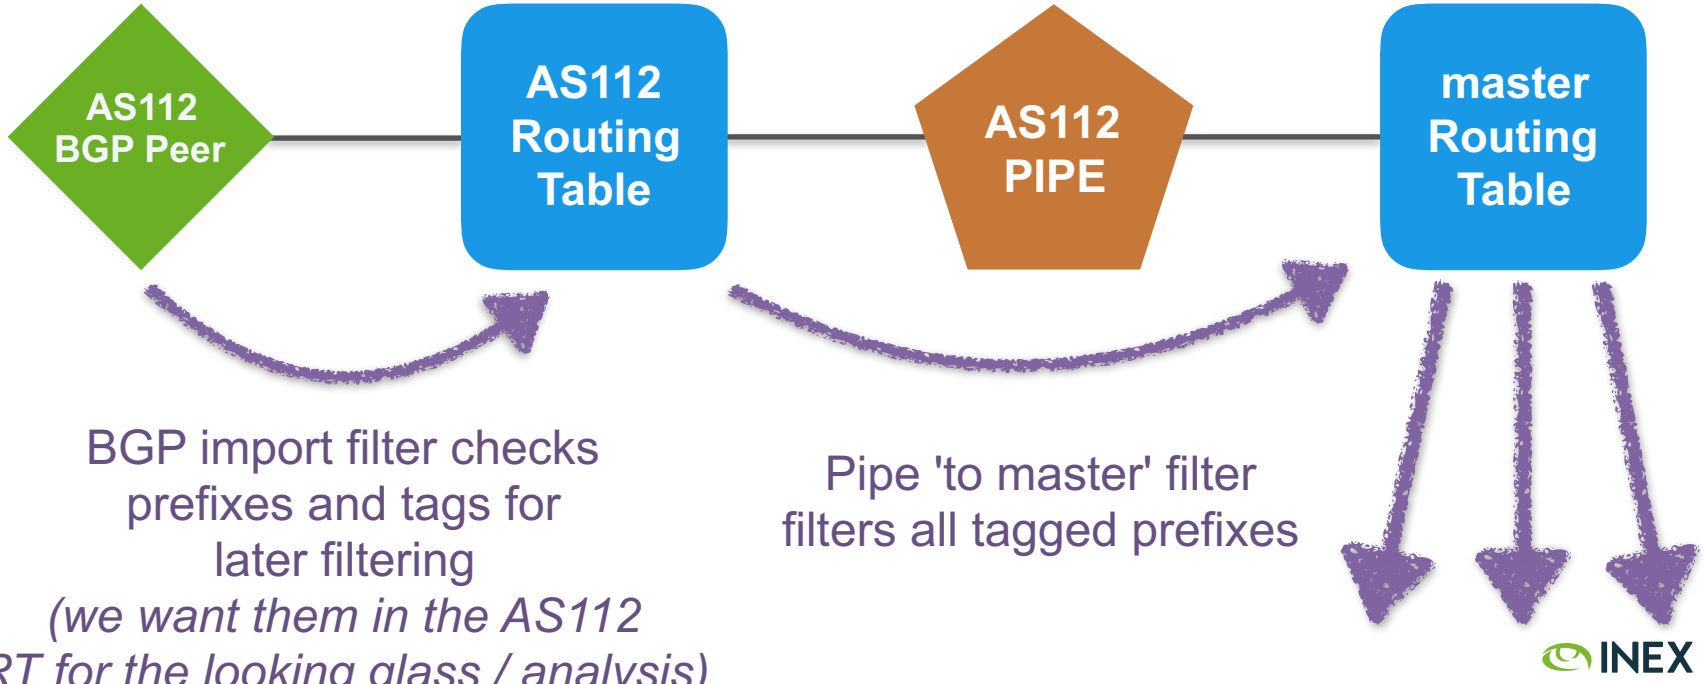

RPKI at IXPs

IXP Manager v5 Bird Topology - Export To Member

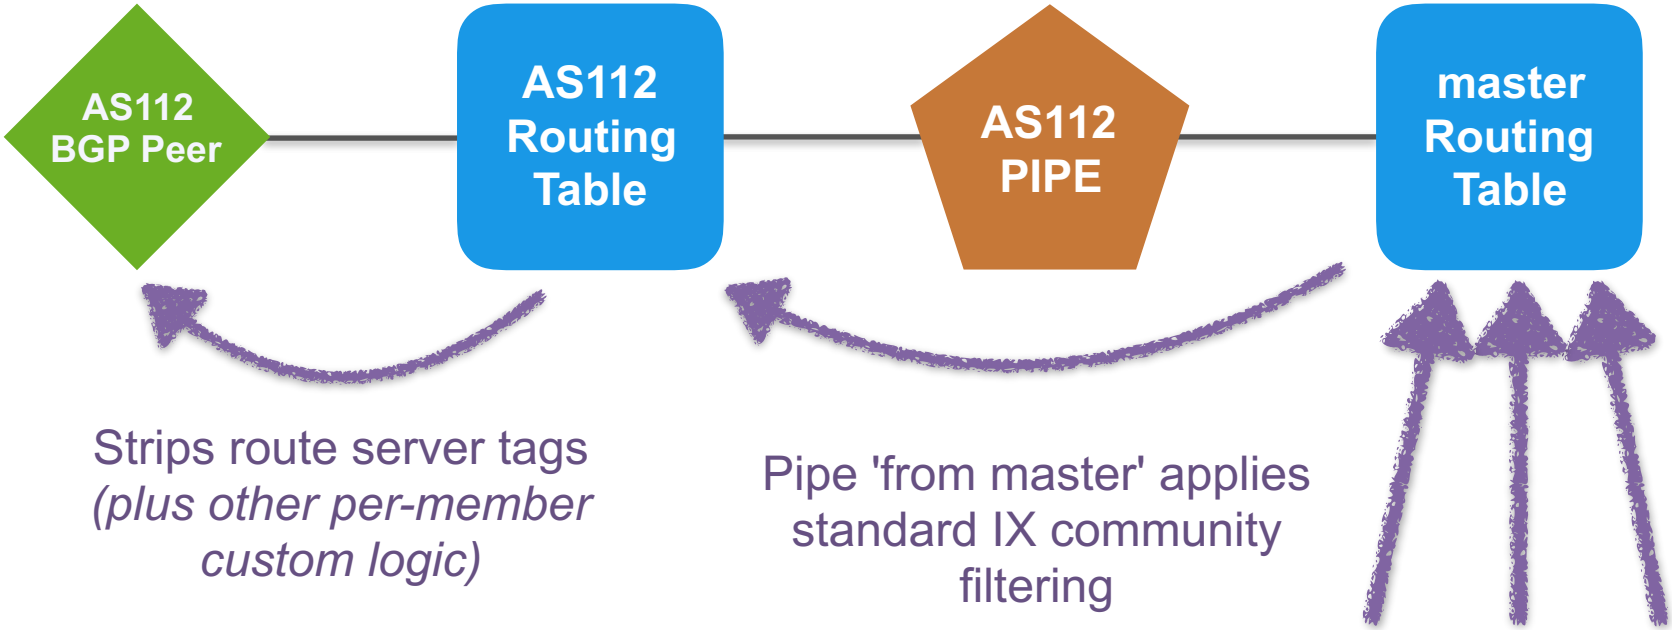

RPKI at IXPs

ROAs - RPKI Invalid Example

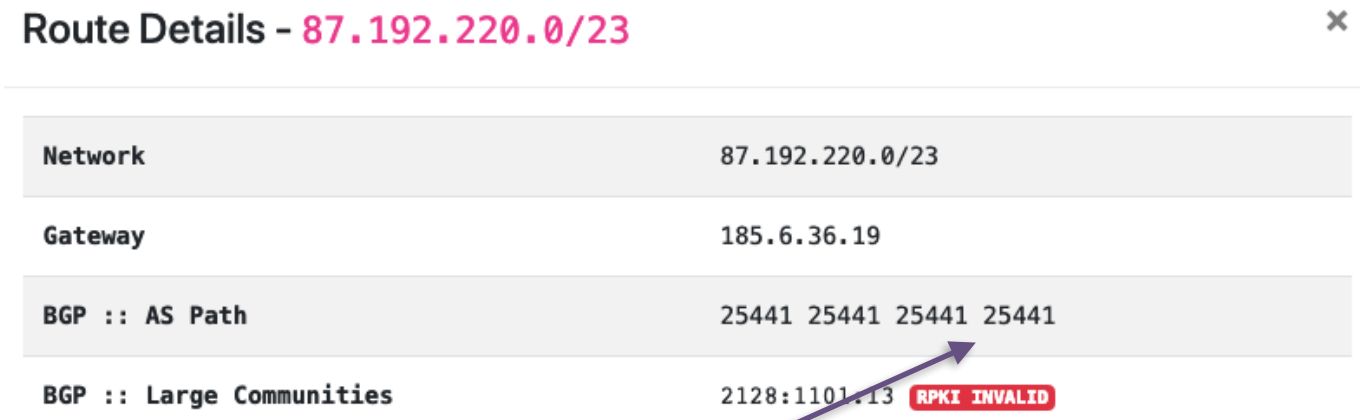


```
Table t_roa:
87.192.220.0/23-24 AS34245   [rpki2 2019-02-05 12:27:38]
                             [rpki1 2019-02-05 12:27:38]
```

INEX

Implementation Notes

INEX

RPKI at IXPs

Validator Software - RIPE NCC RPKI Validator 3

- RIPE NCC RPKI Validator 3 released in 2018
  - https://github.com/RIPE-NCC/rpki-validator-3
- Dramatically reduces installation complexity
- Modest VM requirements, runs on standard OS distributions
- Requirement to download ARIN TAL separately

```
$ wget https://ftp.ripe.net/tools/rpki/validator3/rc/generic/rpki-validator-latest-dist.tar.gz
$ wget https://ftp.ripe.net/tools/rpki/validator3/rc/generic/rpki-rtr-server-latest-dist.tar.gz
$ tar zxf rpki-validator-latest-dist.tar.gz
$ ./rpki-validator-3.0-x/rpki-validator-3.sh
$ tar zxf rpki-rtr-server-latest-dist.tar.gz
$ ./rpki-rtr-server/rpki-rtr-server-3.sh
$ open http://localhost:8080
```

INEX

RPKI at IXPs

Validator Software - Routinator 3000

- Routinator 3000 by NLnet Labs
  - https://github.com/NLnetLabs/routinator
- First impressions: low overheard, installation simplicity, stable, "just works"
- Requirement to download ARIN TAL separately

```
$ curl https://sh.rustup.rs -sSf | sh
$ source ~/.cargo/env
$ cargo install routinator
$ routinator rtrd -al 127.0.0.1:3323
```

RPKI at IXPs

Validator Software - RPKI-RTR and Bird

```
roa4 table t_roa;

protocol rpki rpki1 {

    roa4 { table t_roa; };

    remote "192.0.2.67" port 3323;

    retry keep 90;
    refresh keep 900;
    expire keep 172800;
}
```

RPKI at IXPs

Validator Software - RPKI-RTR and Bird

```
# RPKI check
rpki_result = roa_check( t_roa, net, bgp_path.last_nonaggregated );


if( rpki_result = ROA_INVALID ) then {
    ...
}


# or ROA_VALID / ROA_UNKNOWN
```

RPKI at IXPs

AS Paths

- No ability to validate AS paths in RPKI
- No ability to create AS sets in RPKI
  - draft-ietf-grow-rpki-as-cones will resolve this
- These are regressions over static IRRDB filtering
  - path validation is hard
  - AS Set / AS Cone support is critical

RPKI at IXPs

RPKI Revalidation

- Not currently supported on Bird
- Workaround:

```
10 */4 * * *   root    birdc reload in all >/dev/null
```

RPKI at IXPs

IXP Manager Development

1. Support Bird v2
2. RPKI support
3. Looking glass updates (prefix analysis)
4. Support OpenBGPd
5. Support GoBGP

RPKI at IXPs

Implementation Process at INEX

- INEX has two route servers and a route collector per LAN
- Upgrade route collector to Bird v2 + RPKI first
  - identify members who peer on the route server with RPKI invalid prefixes
  - found 4 members of ~80 with issues
    - 1 x more specific advertised than ROA allowed for
    - 1 x origin AS not matching ROA
    - 1 x member still advertising transferred space, new owners had ROAs
    - 1 x member created ROA for upstream peer-as rather than origin-as
  - members alerted to this on a "FYI basis" (i.e. non-blocking for INEX)
- Route server #1 completed Feb 7th
- Route server #2 completed Feb 14th

INEX

RPKI at IXPs

Implementation Process at INEX

- Outside of the four members with issues, no other member issues
- No production / service issues to date with Bird
- RIPE's validator has crashed twice, no issues with Routinator 3000
- There's a lot in this (Bird v2, route collector vs server, large community tagging and filtering, RPKI vs IRRDB, etc.)
    - Take the time to build internal knowledge with the operations team

INEX

RPKI at IXPs

IXP Manager

- Templates broken into units
- Templates can be skinned
- Robust production-safe config updates
- Template support for route servers, route collectors and AS112 service
- No opt-outs of RPKI when enabled
- If you use IXP Manager's route server templates, you get excellent integration for looking glasses and other tools

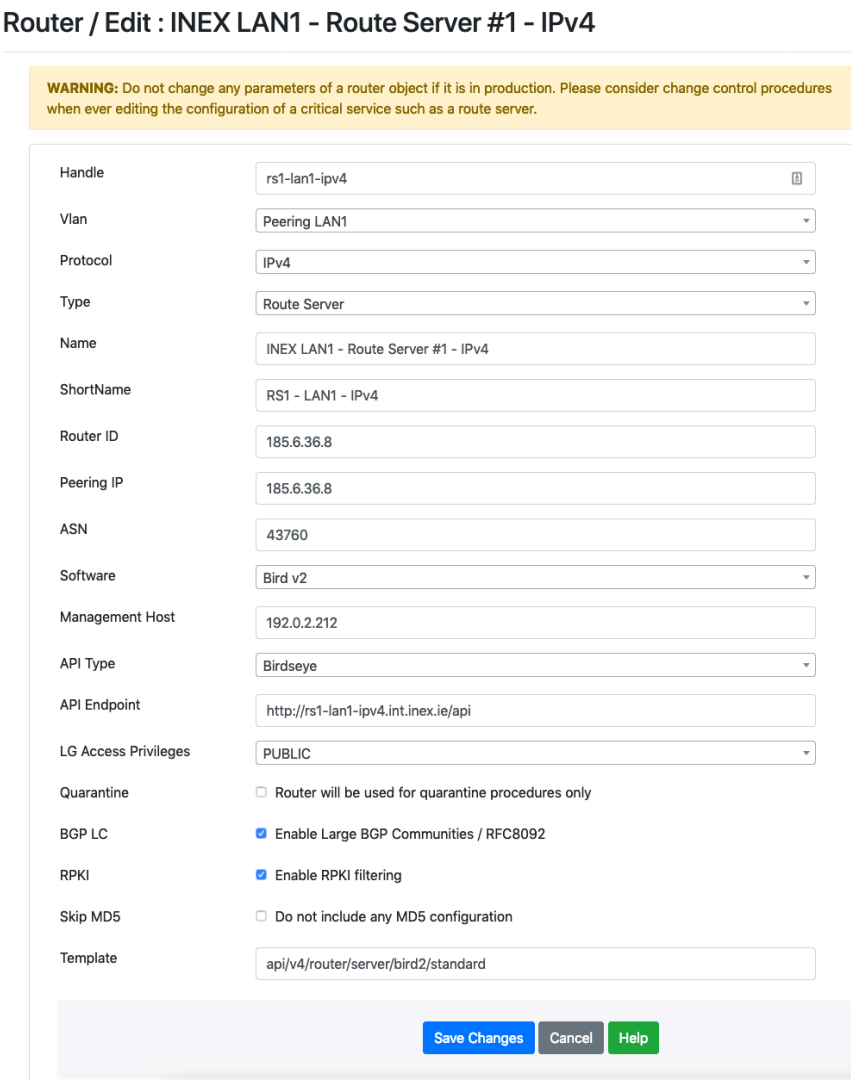

IXP MANAGER INEX Member Information Peering Statistics Support

Looking Glass INEX Cork - Route Collector - IPv4

INEX Cork - Route Collector - IPv4

This is the public looking glass. Uncached results and additional routers available when logged in.

Bird v2 2.0.3 | API: 1.2.0 | Router ID: 185.1.69.126 | Uptime: 11 days. | Last Reconfigure: 2019-02-16 15:12:02 | JSON: [status] [bgp]

Search:

Neighbor	Description	ASN	Table	PfxLimit	State/PfxRcd	PfxExp	Actions
185.1.69.6	AS112 - AS112 Reverse DNS	112	master4		2	0	Details
185.1.69.24	AS714 - Apple Distribution International	714	master4		596	0	Details
185.1.69.26	AS714 - Apple Distribution International	714	master4		597	0	Details
185.1.69.11	AS1213 - HEAnet	1213	master4		23	0	Details
185.1.69.12	AS5466 - Eir	5466	master4		77	0	Details
185.1.69.17	AS15405 - East Cork Broadband	15405	master4		5	0	Details
185.1.69.14	AS16171 - Strencom	16171	master4		4	0	Details
185.1.69.16	AS20940 - Akamai Technologies	20940	master4		1	0	Details
185.1.69.23	AS25152 - RIPE NCC k-root server	25152	master4		1	0	Details
185.1.69.10	AS31122 - Viatel	31122	master4		90	0	Details
185.1.69.19	AS41736 - Nova Telecom	41736	master4		3	0	Details
185.1.69.21	AS42090 - Rapid Broadband	42090	master4		6	0	Details

INEX

Network	Next Hop		Metric	Communities?	AS Path	
104.132.227.0/24	185.1.69.12	P	100	1 LC: 2	5466 41264	Details
109.125.0.0/18	185.1.69.12	P	100	1 LC: 2	5466 15751	Details
132.189.78.0/24	185.1.69.12	P	100	1 LC: 3 ⚠	5466 8116	Details
132.189.79.0/24	185.1.69.12	P	100	1 LC: 3 ⚠	5466 8116	Details
132.237.132.0/24	185.1.69.12	P	100	1 LC: 2	5466 30614	Details
132.237.167.0/24	185.1.69.12	P	100	1 LC: 2	5466 30614	Details
134.191.192.0/24	185.1.69.12	P	100	1 LC: 2	5466 4983	Details
134.191.216.0/22	185.1.69.12	P	100	1 LC: 2	5466 4983 4983 4983 4983 4983 4983 4983 4983 4983	Details
134.191.220.0/23	185.1.69.12	P	100	1 LC: 2	5466 4983 4983 4983 4983 4983 4983 4983 4983 4983	Details
134.191.240.0/22	185.1.69.12	P	100	1 LC: 3 ⚠	5466 4983	Details
134.191.244.0/24	185.1.69.12	P	100	1 LC: 3 ⚠	5466 4983	Details
134.191.246.0/23	185.1.69.12	P	100	1 LC: 2	5466 4983	Details
135.74.153.0/24	185.1.69.12	P	100	1 LC: 3 ⚠	5466 18676	Details
146.214.64.0/23	185.1.69.12	P	100	1 LC: 3 ⚠	5466 42213	Details

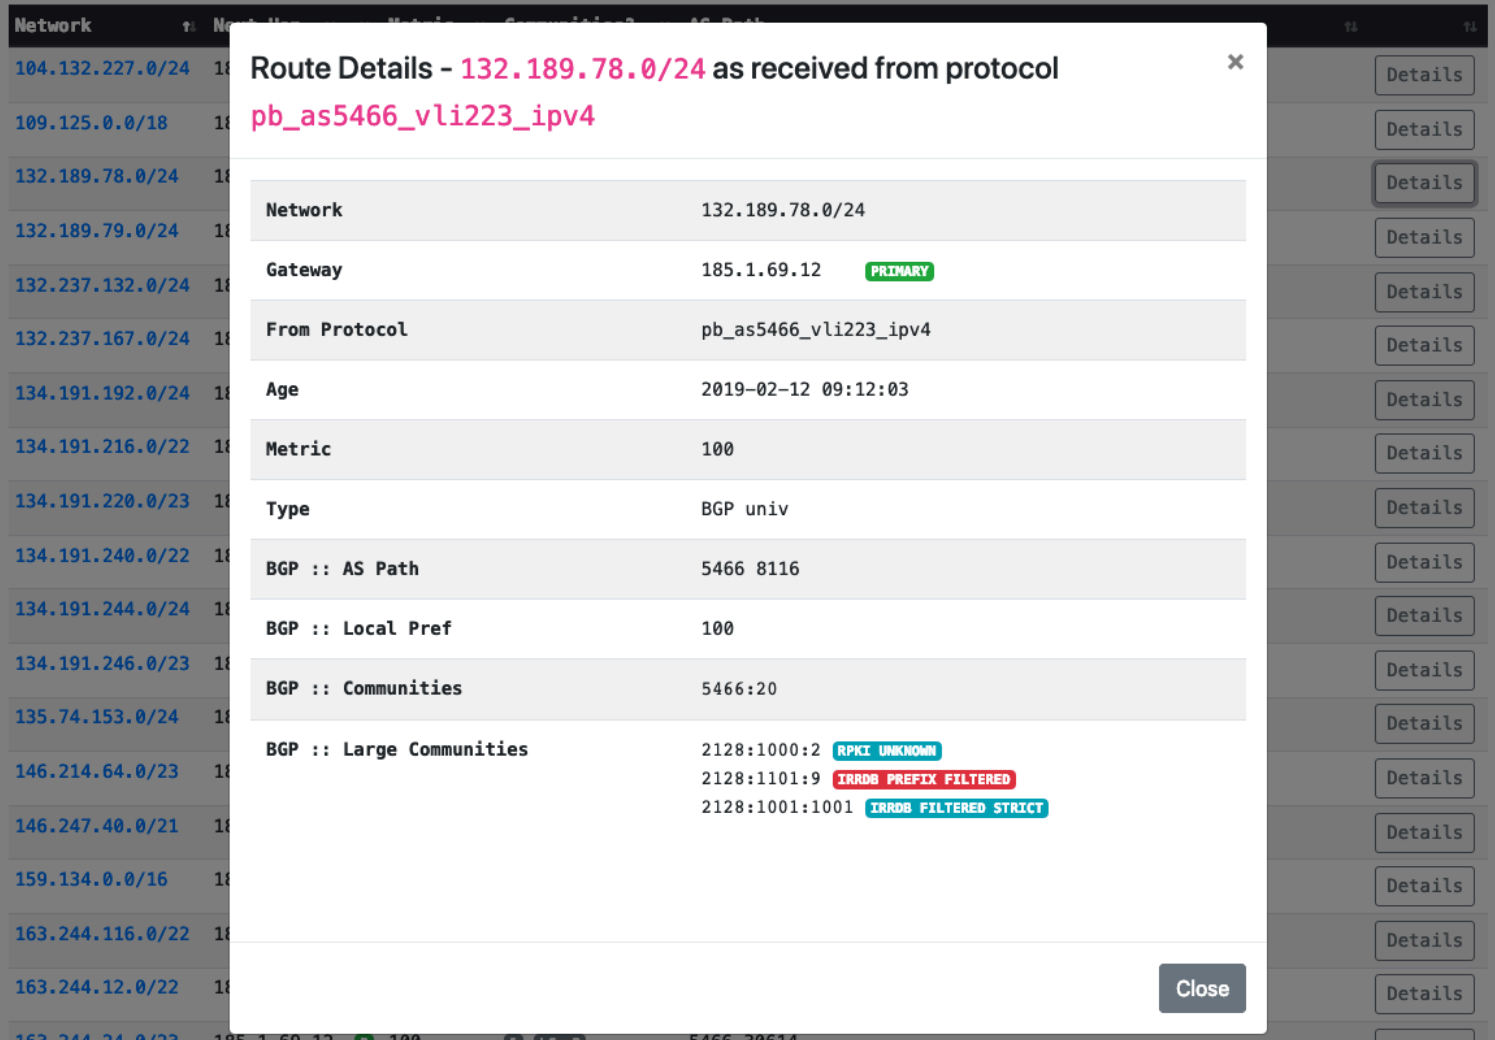
Network
104.132.227.0/24
109.125.0.0/18
132.189.78.0/24
132.189.79.0/24
132.237.132.0/24
132.237.167.0/24
134.191.192.0/24
134.191.216.0/22
134.191.220.0/23
134.191.240.0/22
134.191.244.0/24
134.191.246.0/23
135.74.153.0/24
146.214.64.0/23
146.247.40.0/21
159.134.0.0/16
163.244.116.0/22
163.244.12.0/22
163.244.24.0/23
185.1.69.12 P 100 1 LC: 2 5466 30614
Details
Route Details - 132.189.78.0/24 as received from protocol
pb_as5466_vli223_ipv4
Network 132.189.78.0/24
Gateway 185.1.69.12 PRIMARY
From Protocol pb_as5466_vli223_ipv4
Age 2019-02-12 09:12:03
Metric 100
Type BGP univ
BGP :: AS Path 5466 8116
BGP :: Local Pref 100
BGP :: Communities 5466:20
BGP :: Large Communities 2128:1000:2 RPKI UNKNOWN
2128:1101:9 IRRDB PREFIX FILTERED
2128:1001:1001 IRRDB FILTERED STRICT
Close

IXP MANAGER SPONSORS

INEX

@ComePeerWithMe

Any Questions?

Barry O'Donovan
barry.odonovan@inex.ie
@barryo79