

# Measuring the route collector eco system

---

Franziska Lichtblau<sup>1</sup>, Florian Streibelt<sup>1</sup>, Lars Prehn<sup>1</sup>, Cristel Pelsser<sup>3</sup>,  
Randy Bush<sup>5</sup>, Anja Feldmann<sup>1</sup>

October 15, 2019

<sup>1</sup> Max Planck Institute for Informatics (MPII), <sup>3</sup> University of Strasbourg, <sup>5</sup> Internet Initiative Japan (IIJ) & Arrcus, inc

# Beware!!

*Originally this presentation should have been on the RIPE79 plenary. Due to short notice data inconsistency we withdrew the presentation. This talk is a summary of the project idea, some numbers and results which are preliminary and lack accuracy.*

## What's interesting about route collectors?

- Multiple projects: RIPE/RIS, RouteViews, Isolario, PCH, ...
- Multiple collectors in different geographical regions
- Different peers at each collector
- Different peering strategies per project  
(e.g. multi-hop BGP sessions vs. directly attached peers)

*Our question: Does the selection of collectors impact the result of an analysis?*

**Side effect: Survey of the state of the route collector ecosystem...**

## **Findings: Route Collector Projects**

---

## Differences between collector projects

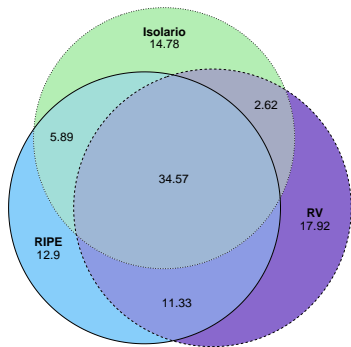
- Do all collectors resp. projects share the same view?
  - In terms of IP address space
  - In terms of inter AS links
- Does geographical location of collectors change the view?
- How do features like multi-hop BGP sessions or BGP ADD-PATH change the picture?

## Visibility: ASes and IP prefixes


| Project  | #AS   | #Prefixes<br>IPv4 | #Prefixes<br>IPv6 |
|----------|-------|-------------------|-------------------|
| RIPE     | 63575 | 966K              | 83K               |
| RV       | 63400 | 1,5M              | 72K               |
| Isolario | 63798 | 1.1M              | 87K               |
| Union    | 64108 | 1.8M              | 101K              |

Dataset based on a union of all announcements as of May 2018

# Project Comparison: AS links<sup>1</sup>



**Figure 1:** % AS links as seen per route collector project (IPv4).



**Figure 2:** % AS links as seen per route collector project (IPv6).

<sup>1</sup>100% = Links discovered by RIPE+RV+Isolario

# Regional Comparison: AS links<sup>2</sup>

Is it worth the effort to place additional route collectors to less covered regions?

## South America

- 4 collectors
- Additional AS links:  
51.4K (14.8%)

## Asia

- 4 collectors
- Additional AS links:  
4.1K ( 1.19%)

## Africa

- 4 collectors
- Additional AS links:  
10.6K ( 3%)

## Australia

- 2 collectors
- Additional AS links:  
4K (1.17%)

---

<sup>2</sup>100% = Links discovered by all non-multiform collectors



## Misc: BGP ADD-PATH offered by Isolario

- Sends all available paths
- Uncovers additional paths as currently "inactive" paths are visible even if they are never the best path
- ADDPATH announcement: Impossible to tell what the best path of an AS currently is

**40.6% of all Isolario's unique paths are learned only via add-path**

## Are partial feeds of any use?

- Even though route collector projects are asking for full BGP views some ASes only send partial feeds
- About 45% of all observed feeds send only a partial view
- 5.7% (26K) of all uncovered AS links are only visible due to partial feeds

## Concluding Remarks

---

## Recommendations for data users

- Expect all kinds of weirdness in real life BGP data
- Pay attention to effects which can skew your results: multi-hop collectors, BGP ADD-PATH, operators feeding weird stuff into the collectors for \$reason
- Check data quality and expect unavailable time-frames
- Don't hesitate to discuss results with operators - they're a friendly crowd

## Recommendations for route collector operators and peers<sup>3</sup>

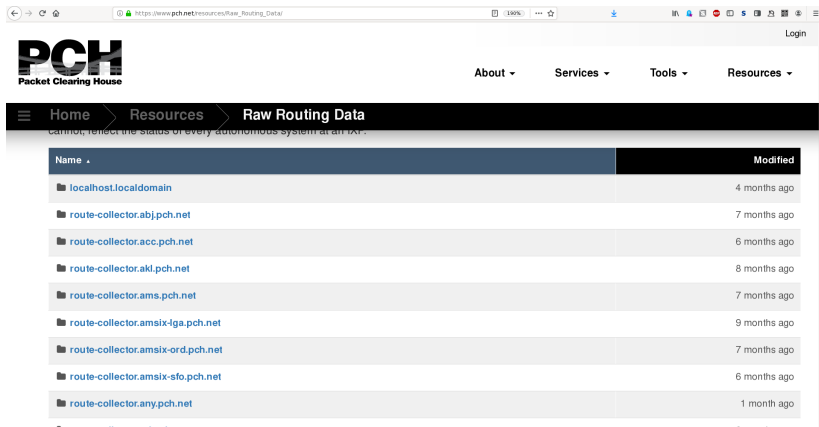
Please document your collectors! Data users care about meta data, e.g., location, peering options, number of peers,...

---

<sup>3</sup>Humble wishes from a sometimes humble data user...

# Recommendations for route collector operators and peers<sup>3</sup>

Please document your collectors! Data users care about meta data, e.g., location, peering options, number of peers,...



The screenshot shows a web browser displaying the Packet Clearing House (PCH) website. The URL is [https://www.pch.net/resources/Raw\\_Routing\\_Data/](https://www.pch.net/resources/Raw_Routing_Data/). The page features a navigation menu with 'About', 'Services', 'Tools', and 'Resources'. The main content area is titled 'Raw Routing Data' and contains a table of route collector entries. The table has two columns: 'Name' and 'Modified'. The entries are as follows:

| Name                              | Modified     |
|-----------------------------------|--------------|
| localhost.localdomain             | 4 months ago |
| route-collector.abj.pch.net       | 7 months ago |
| route-collector.acc.pch.net       | 6 months ago |
| route-collector.akl.pch.net       | 8 months ago |
| route-collector.ams.pch.net       | 7 months ago |
| route-collector.amsix-iga.pch.net | 9 months ago |
| route-collector.amsix-ord.pch.net | 7 months ago |
| route-collector.amsix-sfo.pch.net | 6 months ago |
| route-collector.any.pch.net       | 1 month ago  |

<sup>3</sup>Humble wishes from a sometimes humble data user...

# Recommendations for route collector operators and peers<sup>3</sup>

Provide BGP Updates and RIBs.

Some long lived routes really get missed when only updates are considered!

---

<sup>3</sup>Humble wishes from a sometimes humble data user...

Document collector outages!

It is really hard to guess what a bunch of missing files mean.

---

<sup>3</sup>Humble wishes from a sometimes humble data user...



Place collectors in sparsely covered regions. In case you are a BGP speaker in a sparsely covered region, try reaching out to collector projects - your feeds matter.

It really changes the picture we draw of the Internet!

---

<sup>3</sup>Humble wishes from a sometimes humble data user...