



# ***PCAP BGP Parser***

**RIPE 73, Madrid**

**Christoph Dietzel, Tobias Hannaske**

Research and Development, DE-CIX

## *IXPs' Route Servers*

- » They exist (yees!)
- » Process a significant amount of data
- » Crucial information for IXPs

## *Route Server as BGP Speaker at IXPs*



Customer debugging assistance

Historic analysis (new routes, new peaks)

Incidents (route hijacks, route leaks)

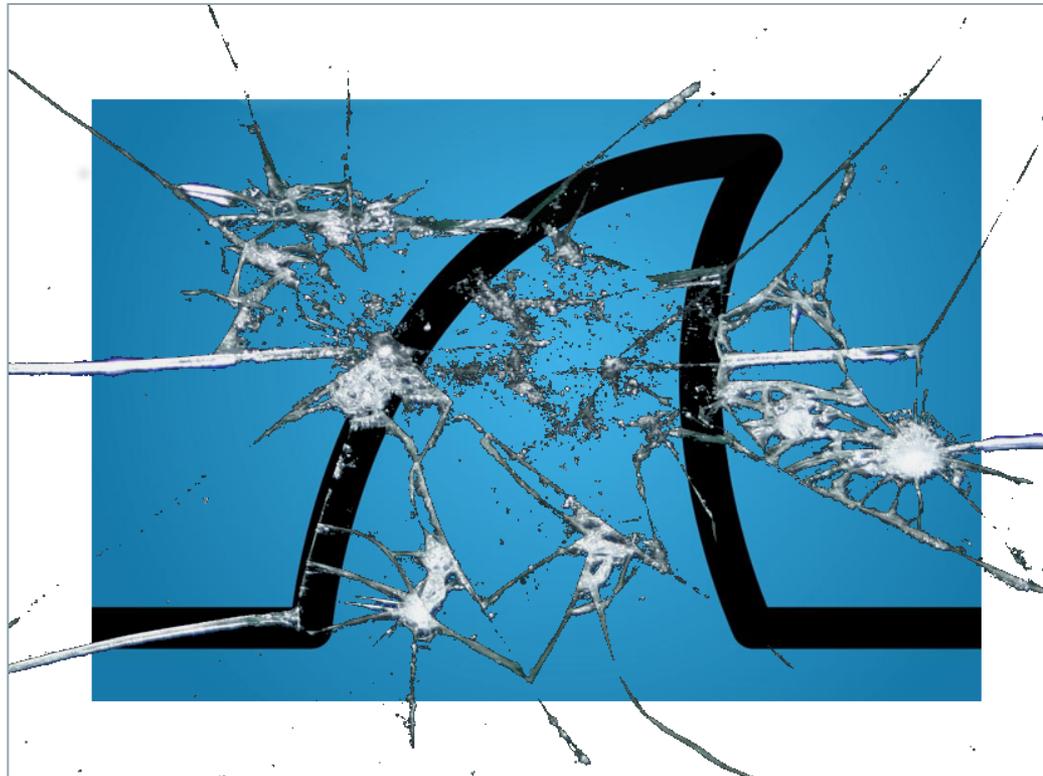
## ***BIRD's Information Export Limitations***

- » Limited long term export of BGP information
- » No continuous export of MRT for BIRD
- » No simple filtering before MRT exports
- » No insights into incoming BGP advertisements



## *Solution? - tcpdump & tshark!(!?)*

- » Complex / cumbersome
- » Output hard to process in automated fashion
- » Not build for BGP



## *PCAP BGP Parser*

- » Python 2.7 and 3.x
- » Open Source ([github.com/de-cix/pbgp-parser](https://github.com/de-cix/pbgp-parser))
- » License Apache 2.0



## Features - Input

- » Reads PCAP files (not PCAPng yet - would be easy to implement)
- » BGP parser can read from stdin (PCAP format)
- » Live reading from network interface not fully implemented yet
- » Extending is possible, as long as it relies on raw packet data

```
--interface INTERFACE      use a network interface as input source (specify  
                             interface)  
--pcap PCAP                use a pcap file as input source (specify file)  
--stdin                    use stdin as input source
```

## Features - Filtering

» Filtering before and after parsing

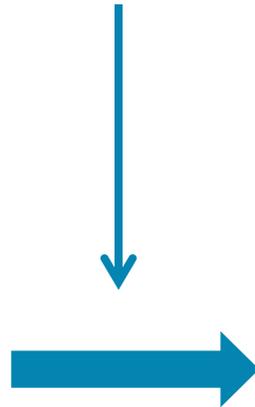


*Where  
networks  
meet*

[www.de-cix.net](http://www.de-cix.net)

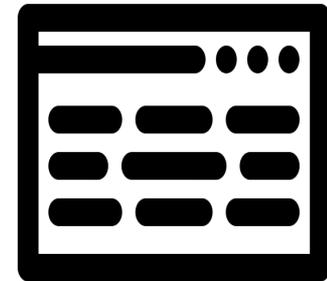
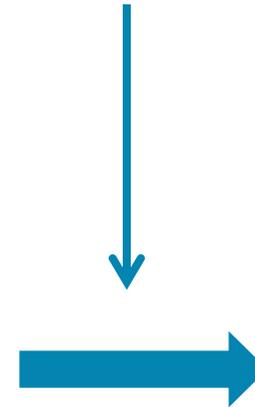
IP/MAC/port src/dst

001  
010  
011  
100



101001  
010101  
001101

other filters



## *Features – Filtering*

| Filter field        | Values; Description                                  |
|---------------------|--|
| Message type        | OPEN, UPDATE, NOTIFICATION, ROUTE-REFRESH, KEEPALIVE |
| NLRI                | Prefix, e.g., 80.81.82.0/24                          |
| Withdrawn route     | Prefix, e.g., 80.81.82.0/24                          |
| Next hop            | IP, e.g., 80.81.82.1                                 |
| ASN in AS path      | ASN, e.g., 6339                                      |
| Last ASN in AS path | ASN of the neighbor AS                               |
| Community ASN       | BGP Community, e.g., 6993:666                        |
| Source IP           | Neighbor router's IP                                 |
| Destination IP      | Neighbor router's IP                                 |
| Source MAC          | Neighbor router's MAC                                |
| Destination MAC     | Neighbor router's MAC                                |
| ...                 | ...  |

## Features - Filtering

- » Filtering to display specific BGP messages – only messages that apply are displayed
- » Combine any filters as desired
- » Different values for same filter are chained with a logical *OR*
- » Different filters are chained with a logical *AND*

```
--filter-nlri 127.0.0.0/8 --filter-nlri 192.168.1.0/32 --filter-next-hop 1.1.1.1
```

- » NLRI must contain either *127.0.0.0/8 OR 192.168.1.0/32 AND* next hop must be *1.1.1.1*

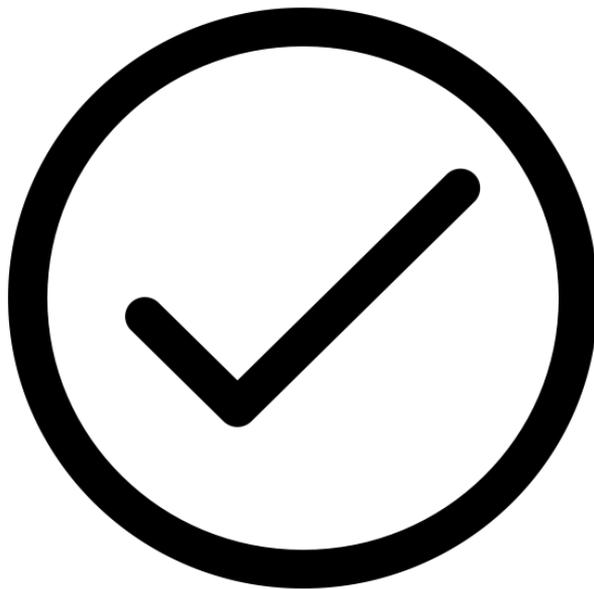
## Features - Output

```
-f {JSON,HUMAN_READABLE,LINE}, --formatter {JSON,HUMAN_READABLE,LINE}
specify data output format
```

- » Human readable
  - » Basic information about BGP msgs
  - » Easy to read
  - » Includes all important fields such as NEXT\_HOP, AS\_PATH, NLRI and/or WITHDRAWALS, etc.
- » JSON
  - » All BGP msgs + meta information (capture specific data such as timestamp, source/dest ip/mac/port)
    - » RFC 7159 (see Python internal json-package)
    - » One JSON string per line
- » Line based
  - » User can specify fields to be displayed
  - » Not all fields supported, yet
  - » Available fields for line based output are:
    - » NLRI, AS\_PATH, NEXT\_HOP, Communities, Source/Destination IP, Timestamp, Message Types

## *Evaluation Correctness*

- » Compared results of PBGP and tshark
  - » E.g., no. of packets after filtering, timestamps
  - » DE-CIX RS dump of several hours



Correct, but we keep looking

## *Limitations*

- » Kafka support with Python 2.7
- » Packet reordering issue
- » Not all features implemented yet

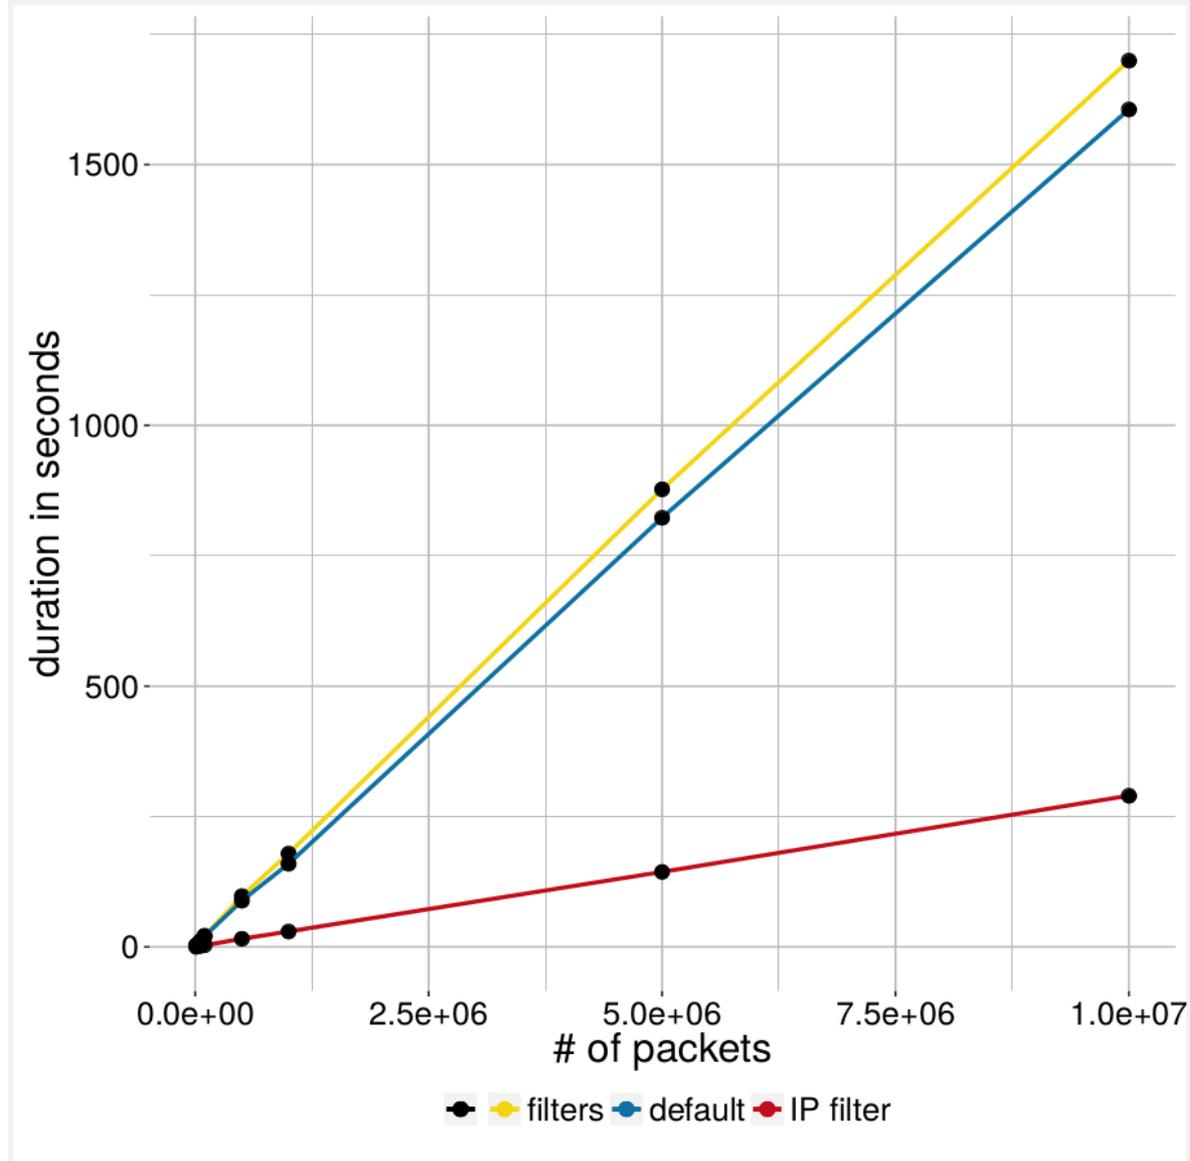


# Evaluation Performance



Where  
networks  
meet

[www.de-cix.net](http://www.de-cix.net)



## ***Conclusion / Contribution***

- » Open source PCAP BGP Parser
- » Apache 2.0 license
- » Wide range of flexible input/output parameters
- » Strong filtering capabilities
- » Nice to integrate in shell/bash/python toolchain
- » Fast enough perform “live” parsing for RS dump from large IXP

[github.com/de-cix/pbgpp-parser](https://github.com/de-cix/pbgpp-parser)



**Where  
networks  
meet**

[www.de-cix.net](http://www.de-cix.net)