Introduction to pmacct

Paolo Lucente

pmacct

ENOG 17 – Nov 2020

whoami

Paolo Lucente

paololucente

plucente

@Paolo_Lucente

Digging data out of networks worldwide for fun and profit for more than 15 years

pmacct is open-source, free, GPL'ed software

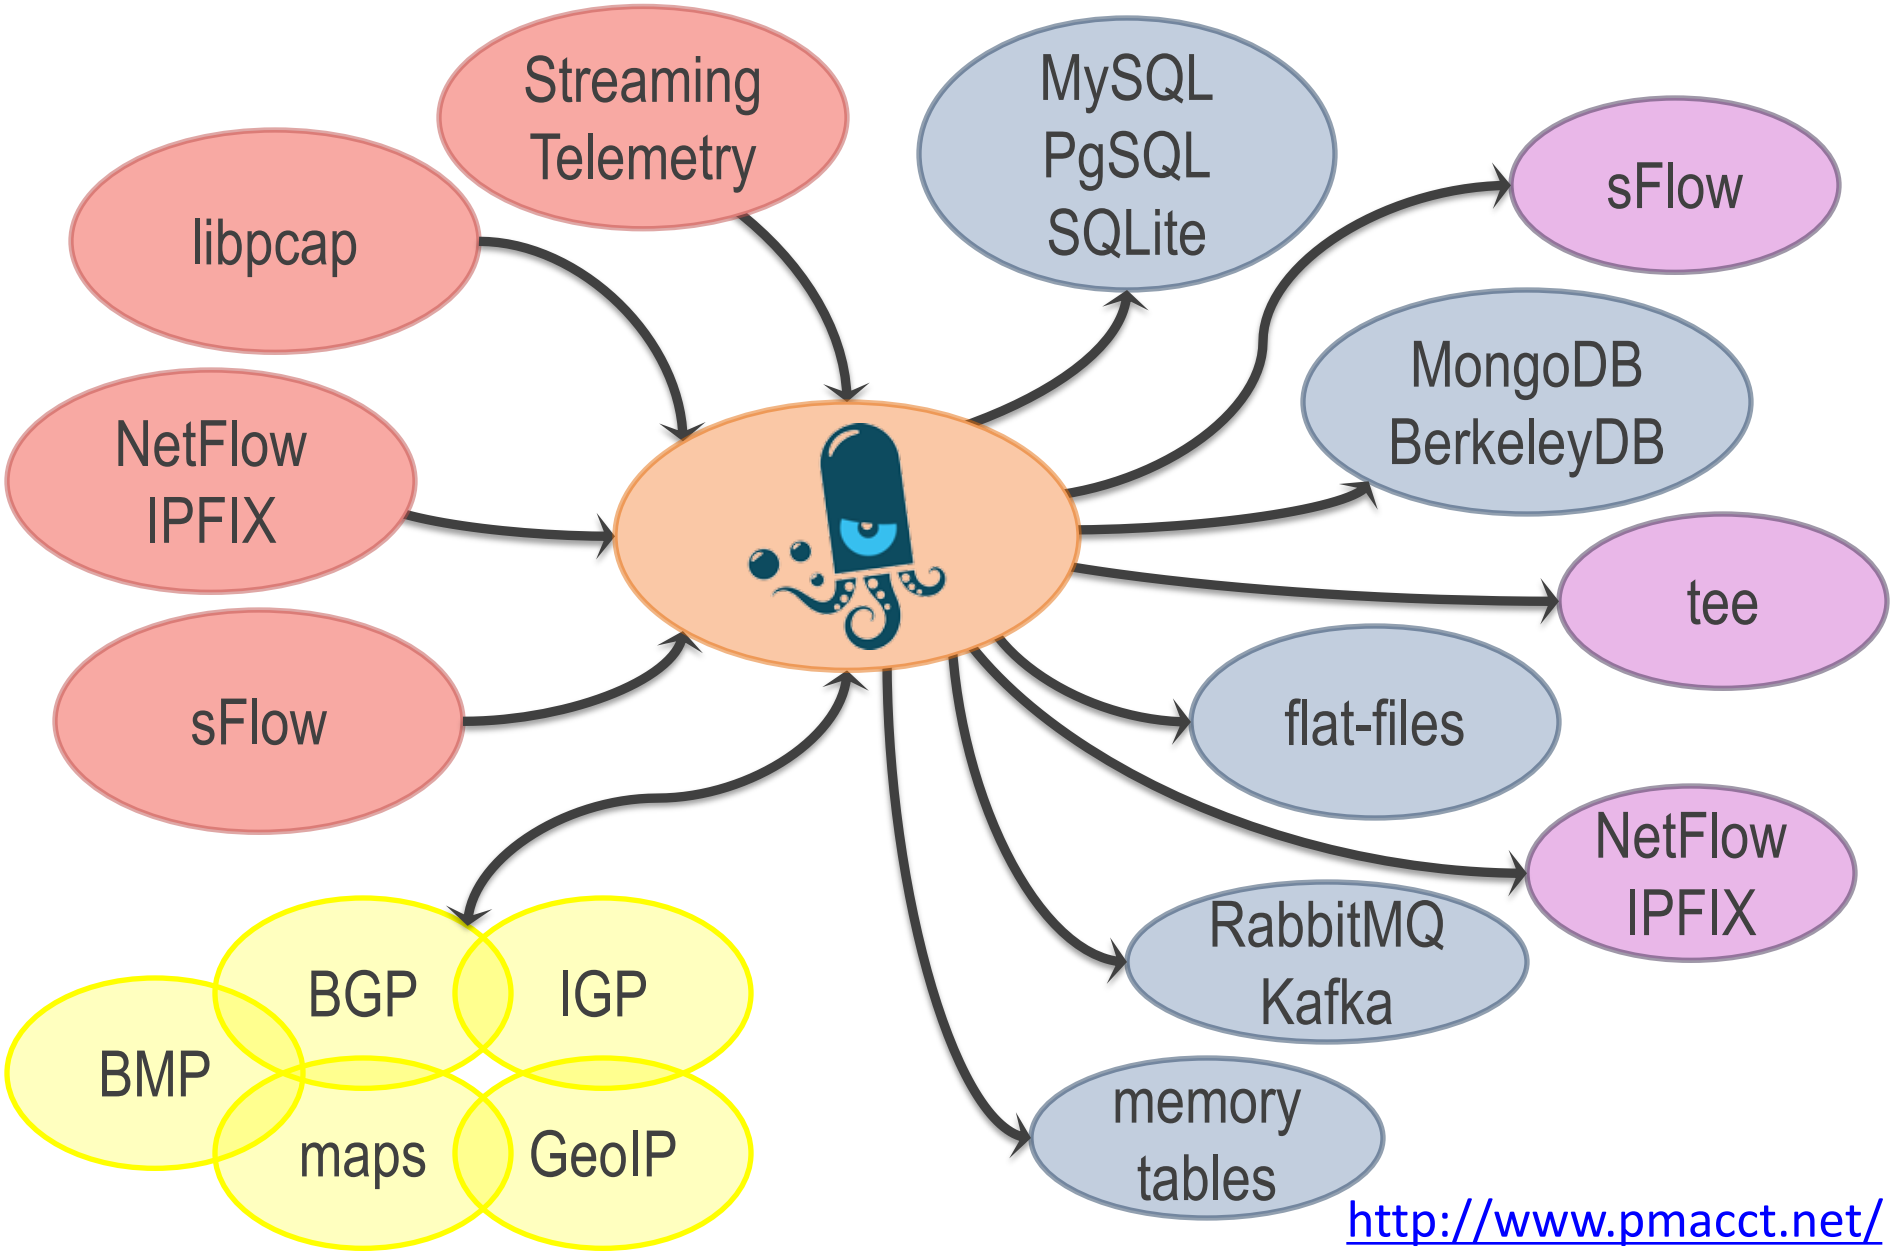

http://www.pmacct.net/

pmacct: a few simple use-cases

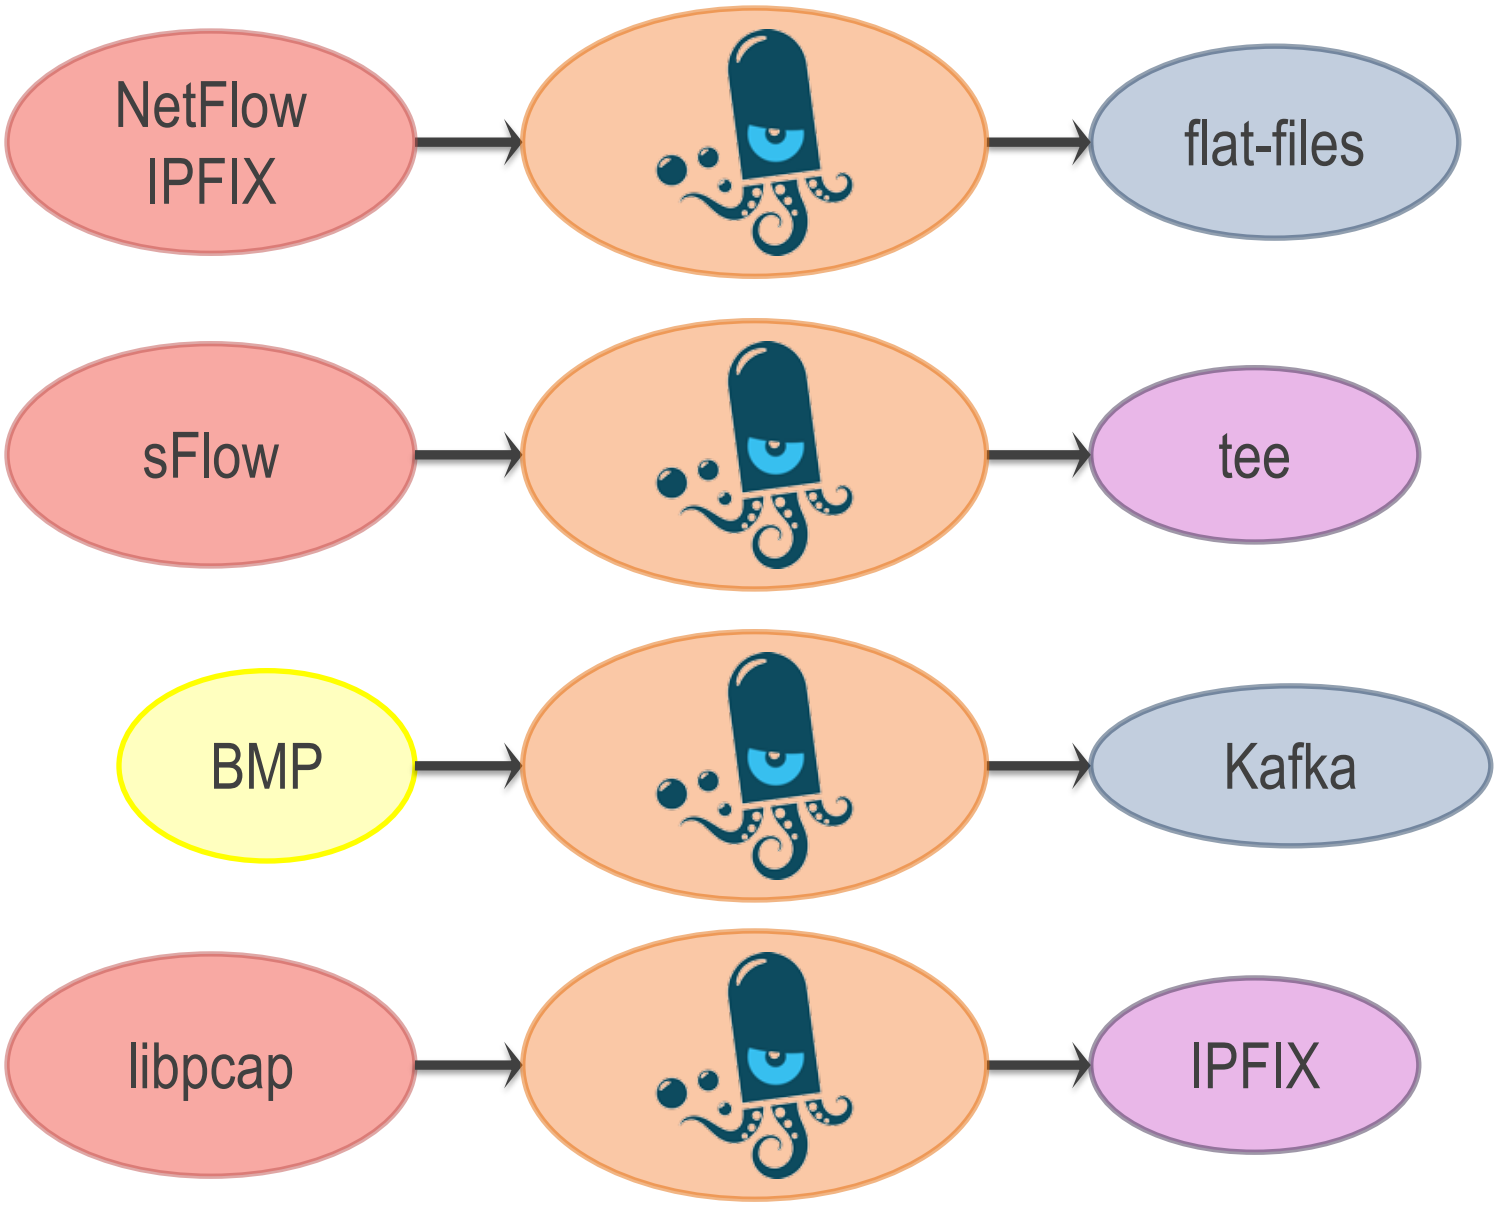

pmacct: a slightly more complex use-case

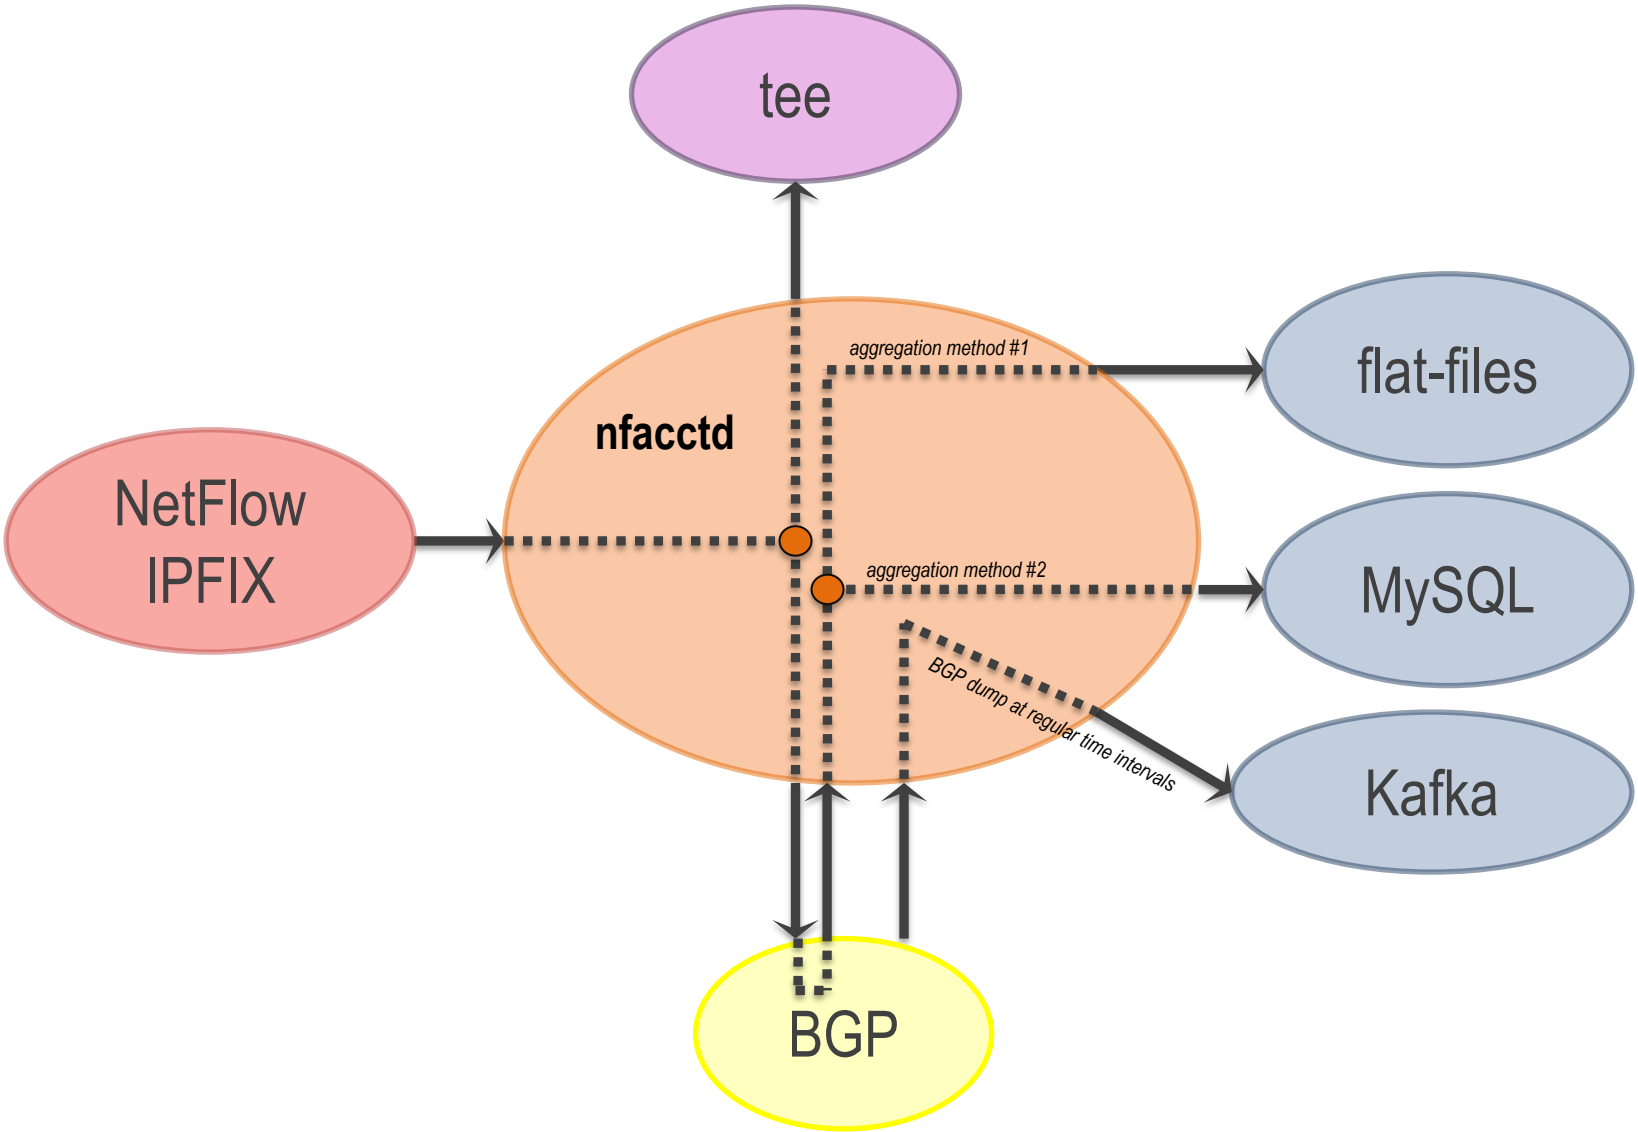

The use-case for message brokers

Sample pipeline for Network Analytics

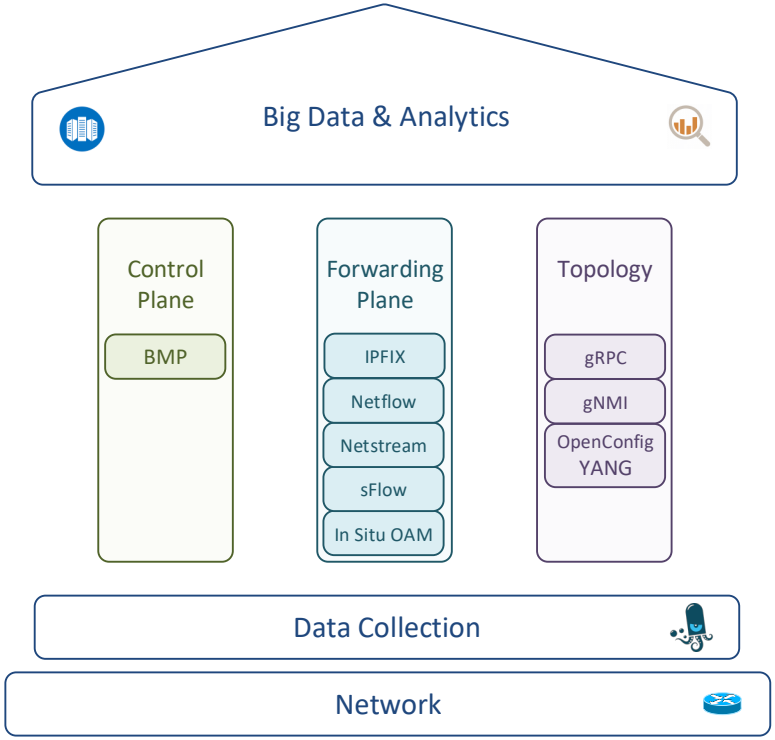

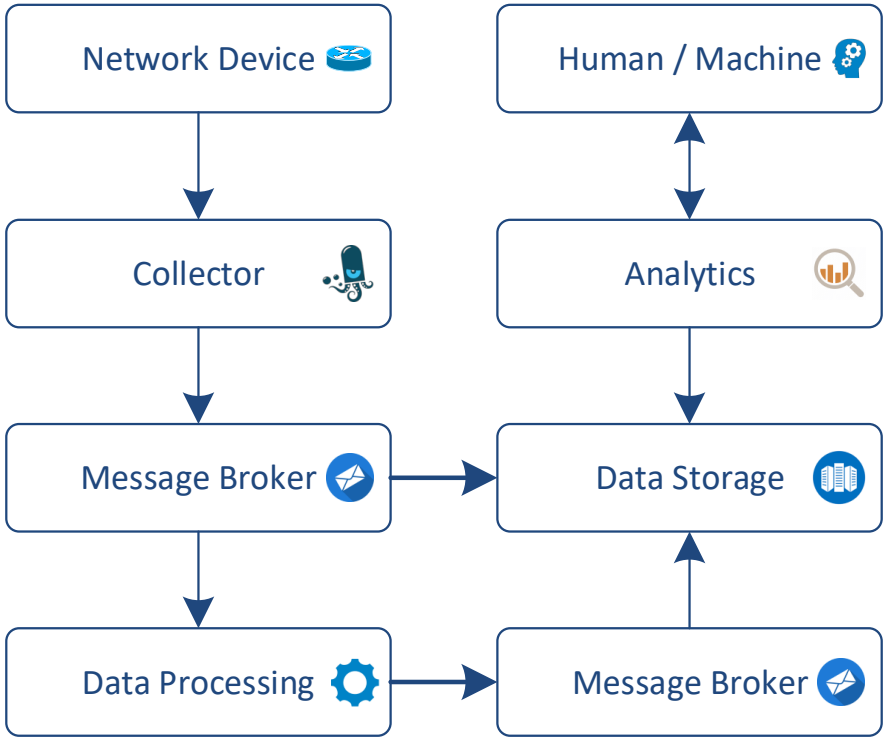

Credits to: T. Graf (Swisscom) @ UBBF 2018

Use cases by industry

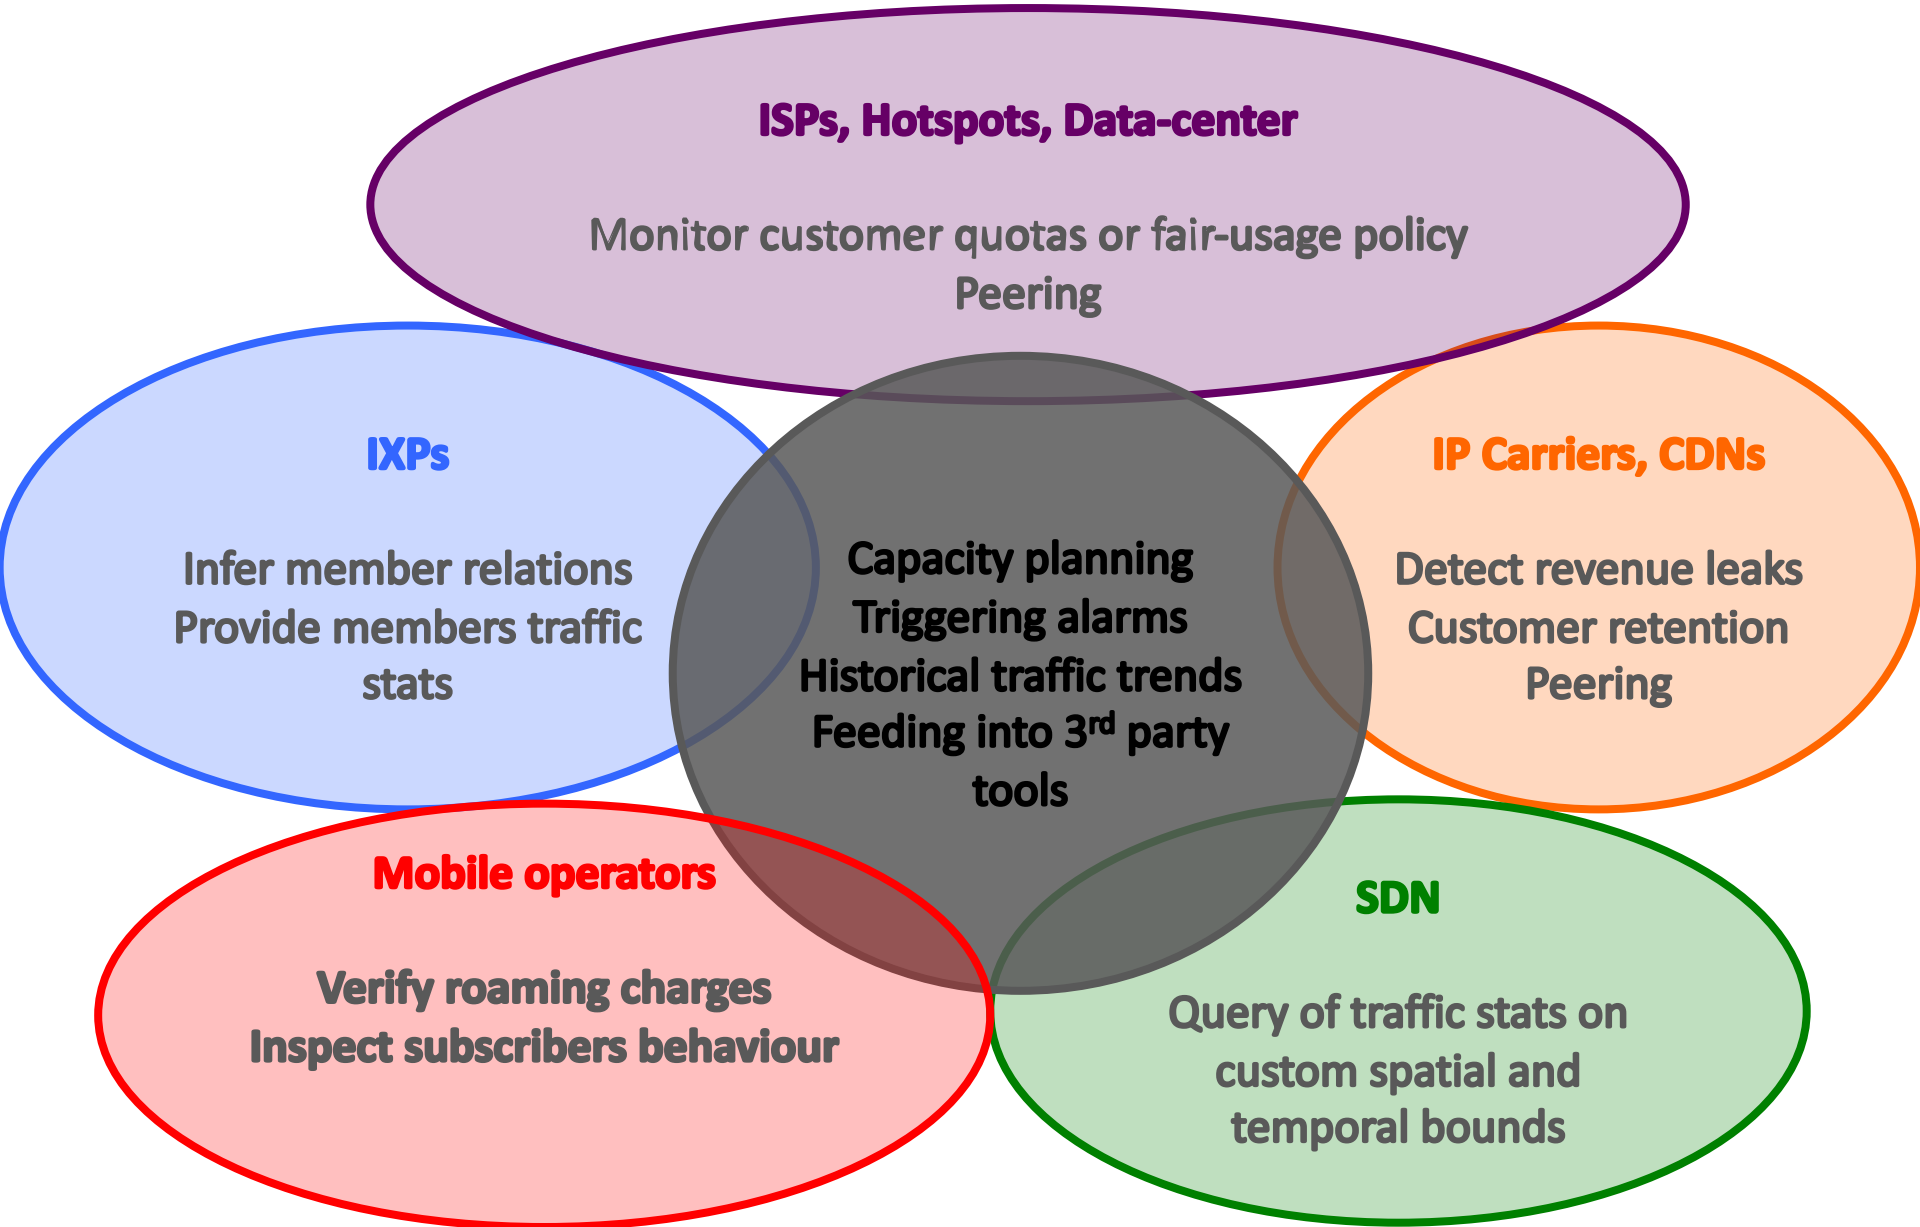

Key pmacct non-technical facts

- 15+ years old project
- Can’t spell the name after the second drink
- Free, open-source, independent
- Under active development
- Innovation being introduced
- Well deployed around, also large SPs
- SP culture and needs part of its DNA

Some technical facts (1/2)

- Pluggable architecture:
  - Can easily add support for new data sources and backends
- Correlation of data sources:
  - Natively supported data sources (ie. Flows, BGP, BMP, IGP, Streaming Telemetry)
  - External data sources via tags and labels
- Pervasive data-reduction techniques, ie.:
  - Data aggregation
  - Filtering
  - Sampling

Some technical facts (2/2)

- Build multiple views out of the very same collected network traffic dataset , ie.:
  - Unaggregated to flat-files for security and forensics; or to message brokers (RabbitMQ, Kafka) for Big Data
  - Aggregated as [<ingress router>, <ingress interface>, <BGP next-hop>, <peer destination ASN>] and sent to a SQL DB to build an internal traffic matrix for capacity planning purposes
- Enable analytics against the collected data sources (ie. BGP, BMP, Streaming Telemetry):
  - Stream real-time
  - Dump at regular time intervals (possible state compression)

Ready to scale horizontally

- Within a single system:
  - SO_REUSEPORT
- Among multiple systems:
  - ‘Tee’ (UDP replication) plugin
    - Receivers can be added/changed/removed on the fly
    - Load-balenced tee’ing (hashed or round-robin)
    - Selective tee’ing (based on telemetry exporter)
  - BGP x-connects (TCP proxying)

It all started like this ..

(easy peasy lemon squeezy)

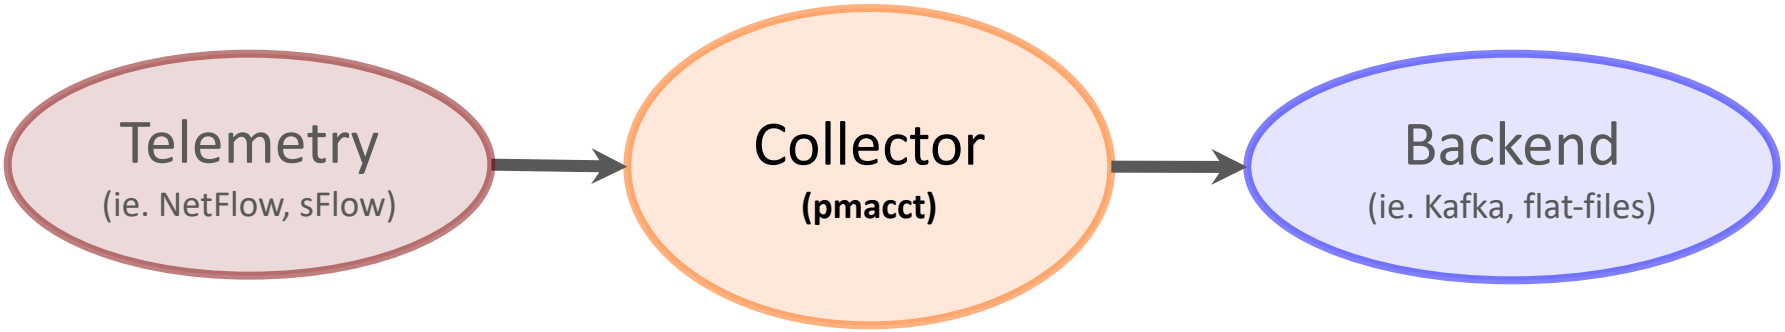

Adding events, integrating routing information

(processing increases and needs tight integration)

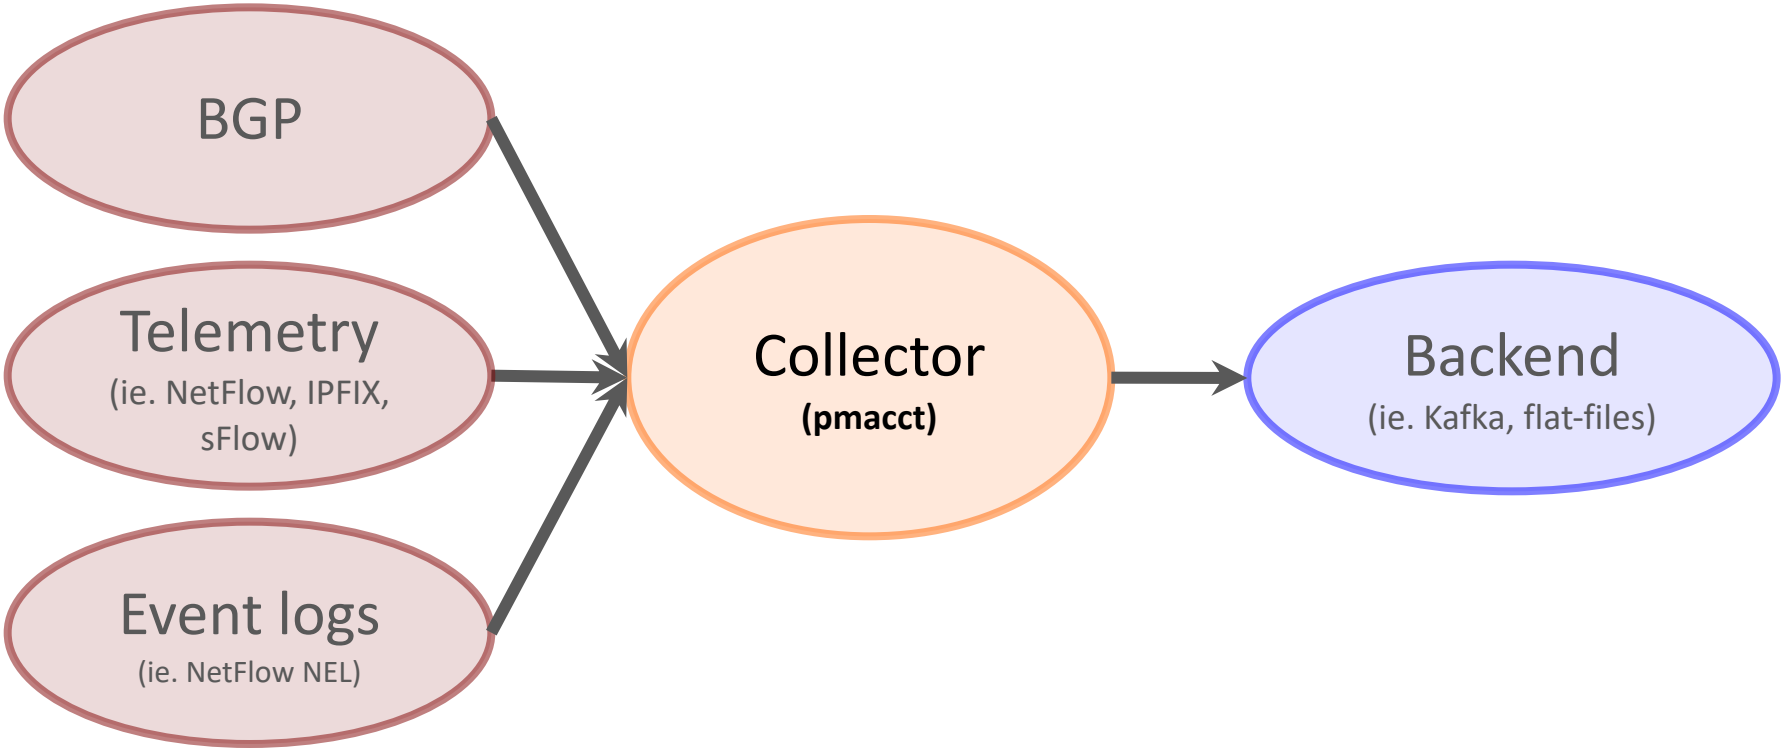

Clustering the collector

(Nice idea, but not easy to provision ..)

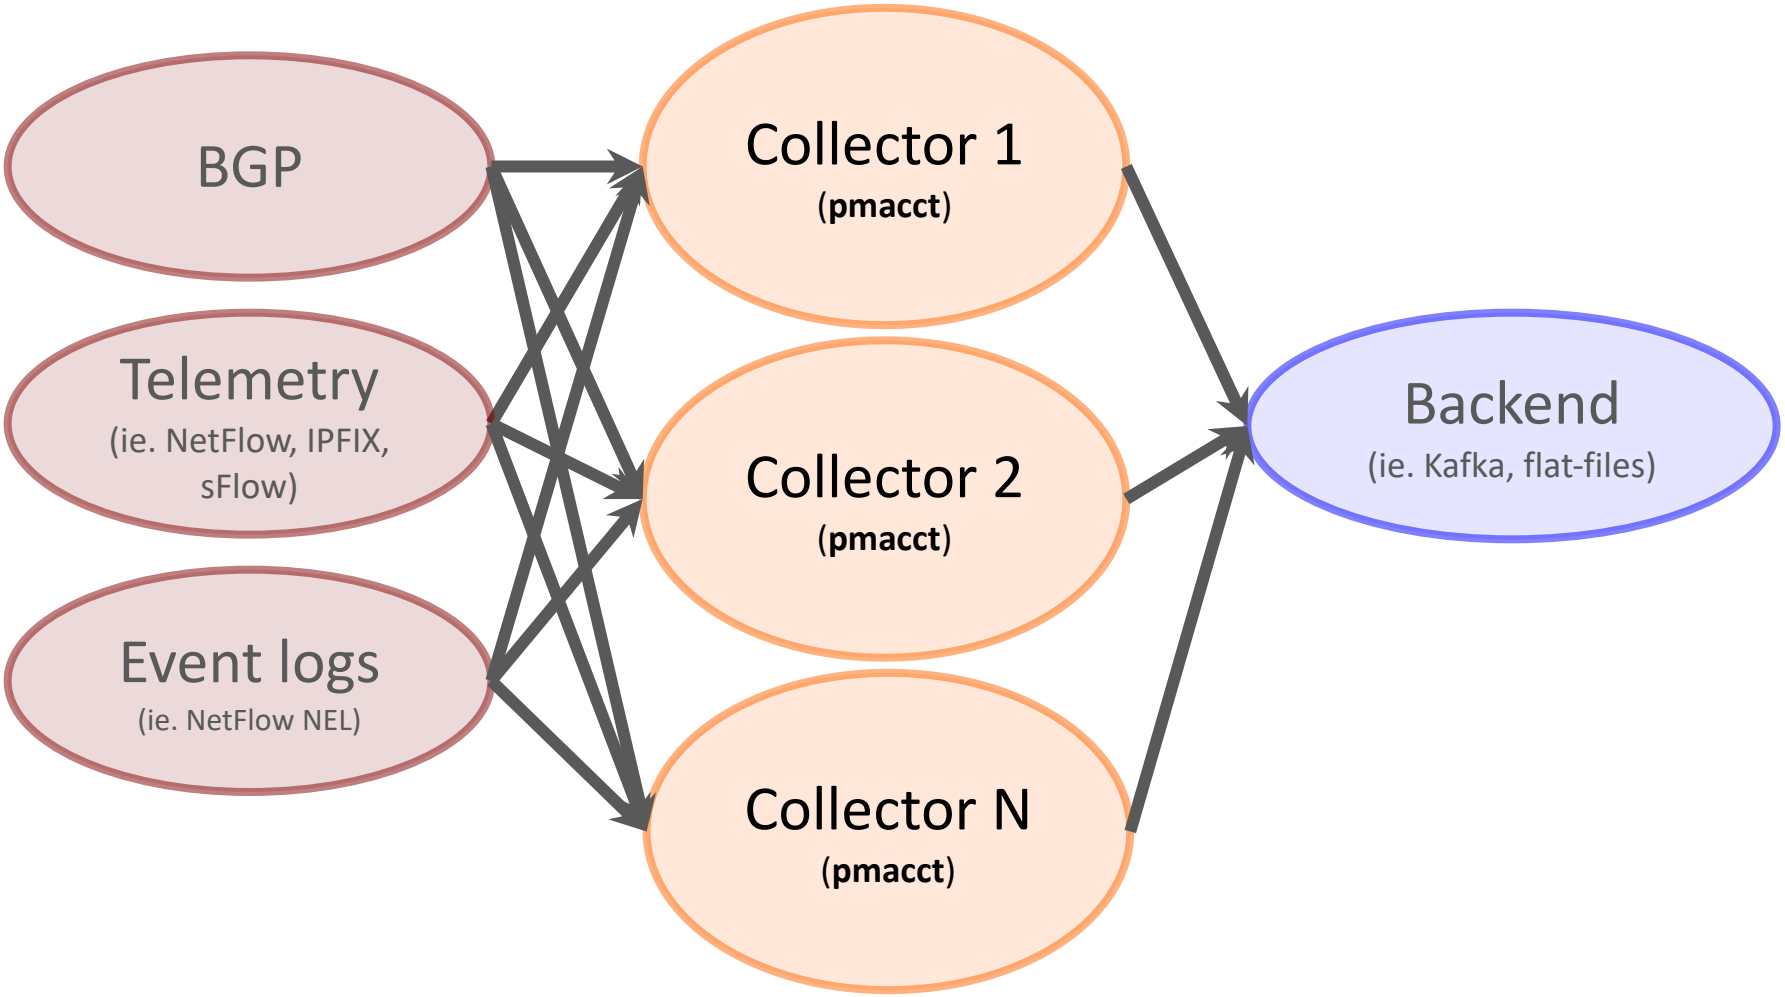

Clustering the collector

(Cluster internals segregated not exposed to the network)

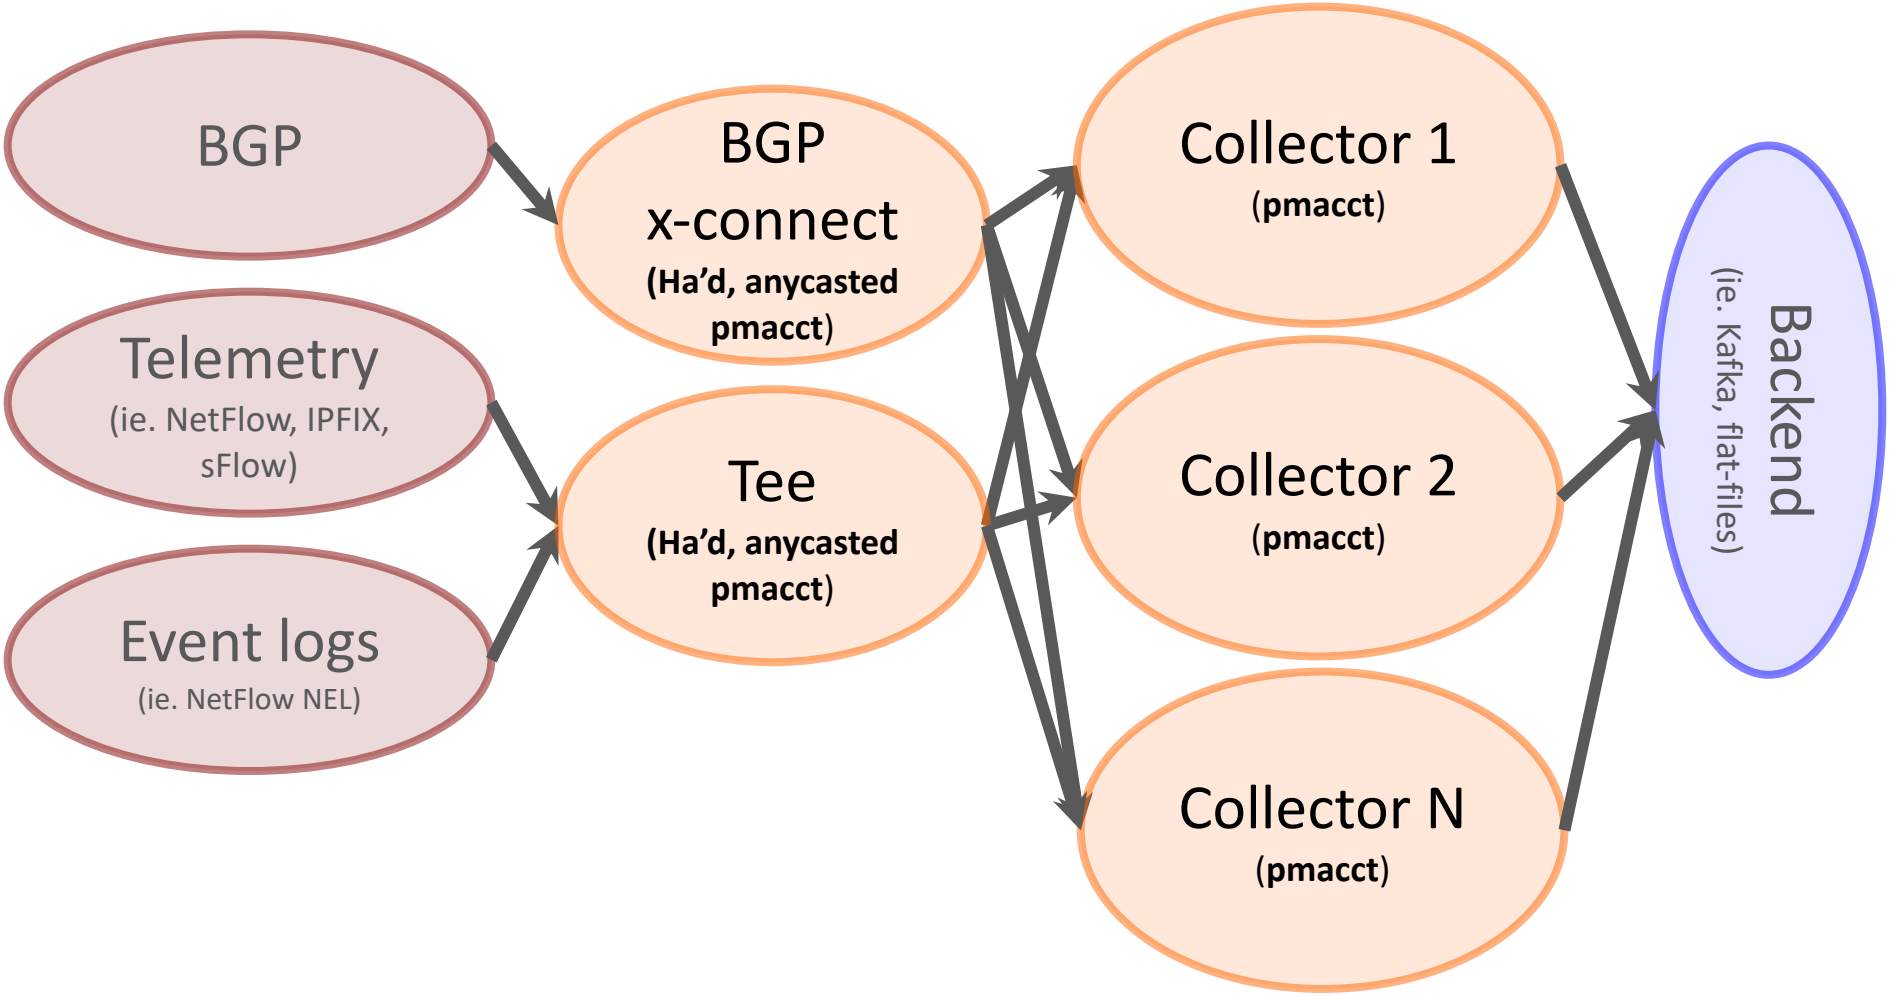

(Some of the) Works in Progress

- BGP-LS support
- Keep pace with BMP protocol development
- Make progress on IETF-proposed Streaming Telemetry (YANG Push)

Further information about pmacct

- https://github.com/pmacct/pmacct
  - Official GitHub repository, where star and watch us ☺
- http://www.pmacct.net/lucente_pmacct_uknof14.pdf
  - More about coupling telemetry and BGP
- http://ripe61.ripe.net/presentations/156-ripe61-bcp-planning-and-te.pdf
  - More about traffic matrices, capacity planning & TE
- https://github.com/pmacct/pmacct/wiki/
  - Wiki: docs, implementation notes, ecosystem, etc.

Thanks! Questions?

Paolo Lucente <paolo@pmacct.net>

http://www.pmacct.net
https://github.com/pmacct/pmacct

ENOG 17 – Nov 2020