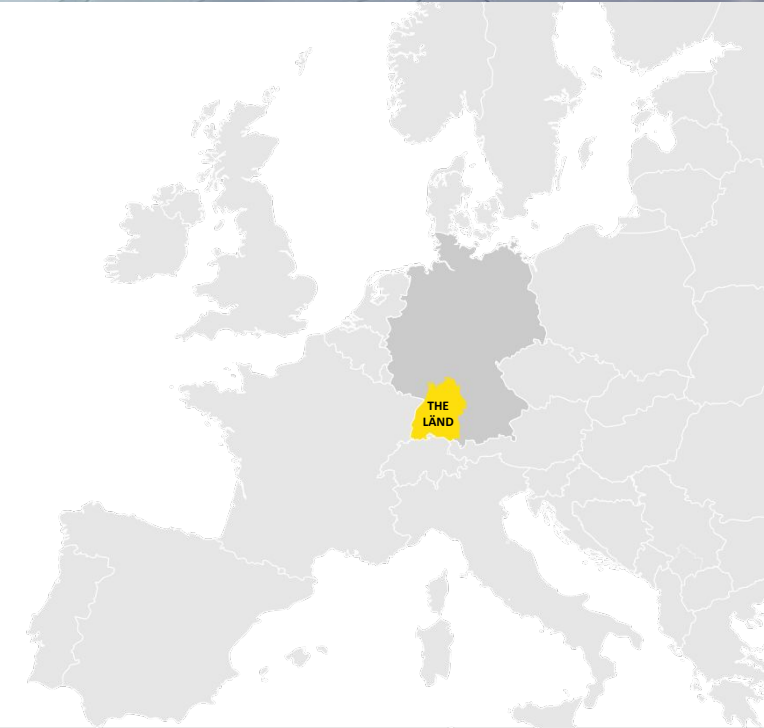


Flow Processing at BelWü

DENOG14

15th November 2022



Purposes of Flow Processing

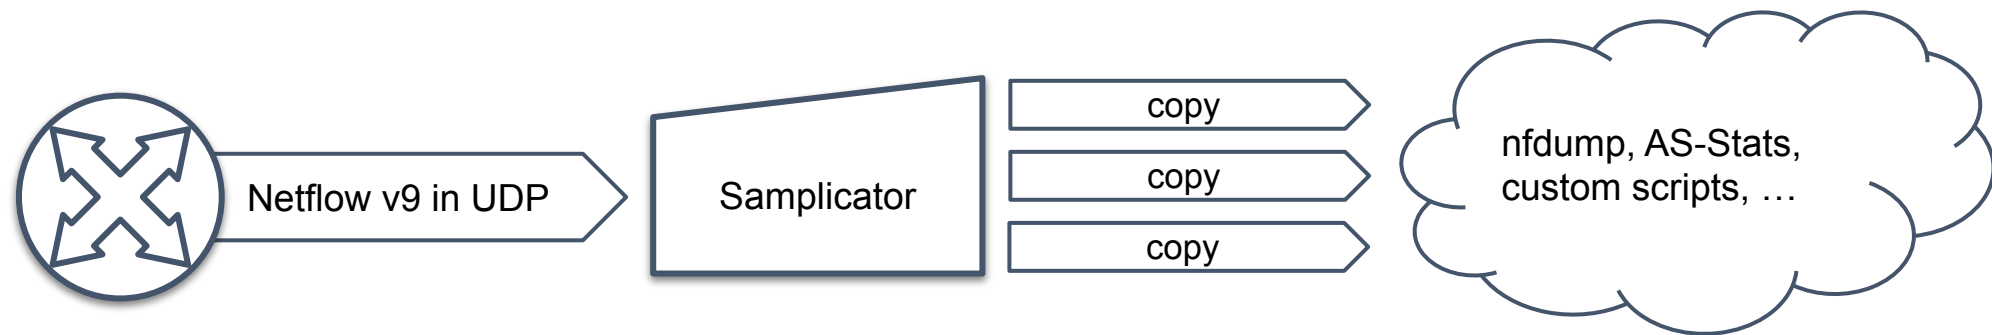
There are many, equally valid reasons to collect and process flows.

Some examples are:

- Find new, impactful peerings, e.g. regarding the utilization of some interface
- Derive flowspec rules for a customer receiving a DoS traffic
- Show awesome graphs to customers
- Detect devices that talk to known bad actors
- Answer this nagging one-off question about your traffic that someone came up with

Problems of Flow Processing

- Some applications require specific “vantage points” in a network
- Different formats and various hardware limitations
- What reality often looks like:
 - A Samplicator^[1] instance, re-sending spoofed UDP datagrams containing flows
 - Different specialized tools parse the same flows



[1] <https://github.com/sleinen/samplicator>

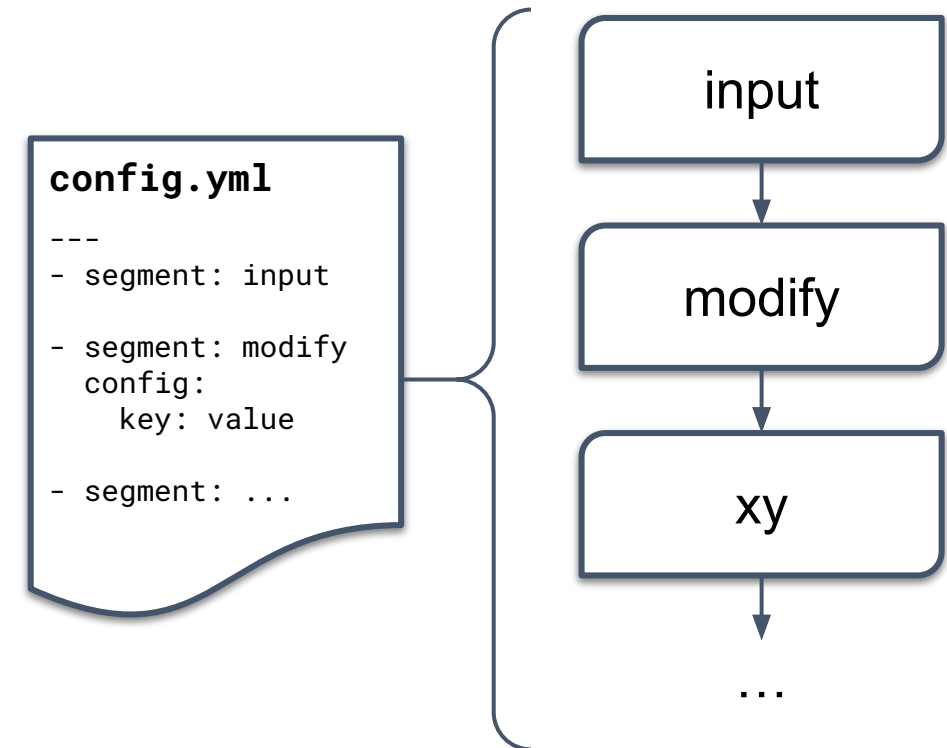
How does yet another tool solve anything?

Focus on “Application Layer” flow processing:

- Stream separation based on arbitrary criteria
- Enrich interesting flows only, drop noise early
- Full multi-tenant support for flow monitoring
- Unify flows from different sources with arbitrary granularity
- Fully reproducible yet extensible setups

flowpipeline Tooling

- Completely configuration-defined
- Single dependency-free binary
- Support for any common flow format
- Segments act on single, protobuf-encoded flow messages and pass them along
- Many different segments are available
- Open Source



Where do Flow Messages come from?

GoFlow 2

- use Goflow v2 to listen for flows in raw format
- supports network devices with sFlow, IPFIX or Netflow v9

Where do Flow Messages come from?

GoFlow 2

- use Goflow v2 to listen for flows in raw format
- supports network devices with sFlow, IPFIX or Netflow v9



- receive flows generated by another flowpipeline from a Kafka cluster
- flows can be pre-filtered or pre-enriched
- ability to output to Kafka

Where do Flow Messages come from?

GoFlow 2

- use Goflow v2 to listen for flows in raw format
- supports network devices with sFlow, IPFIX or Netflow v9

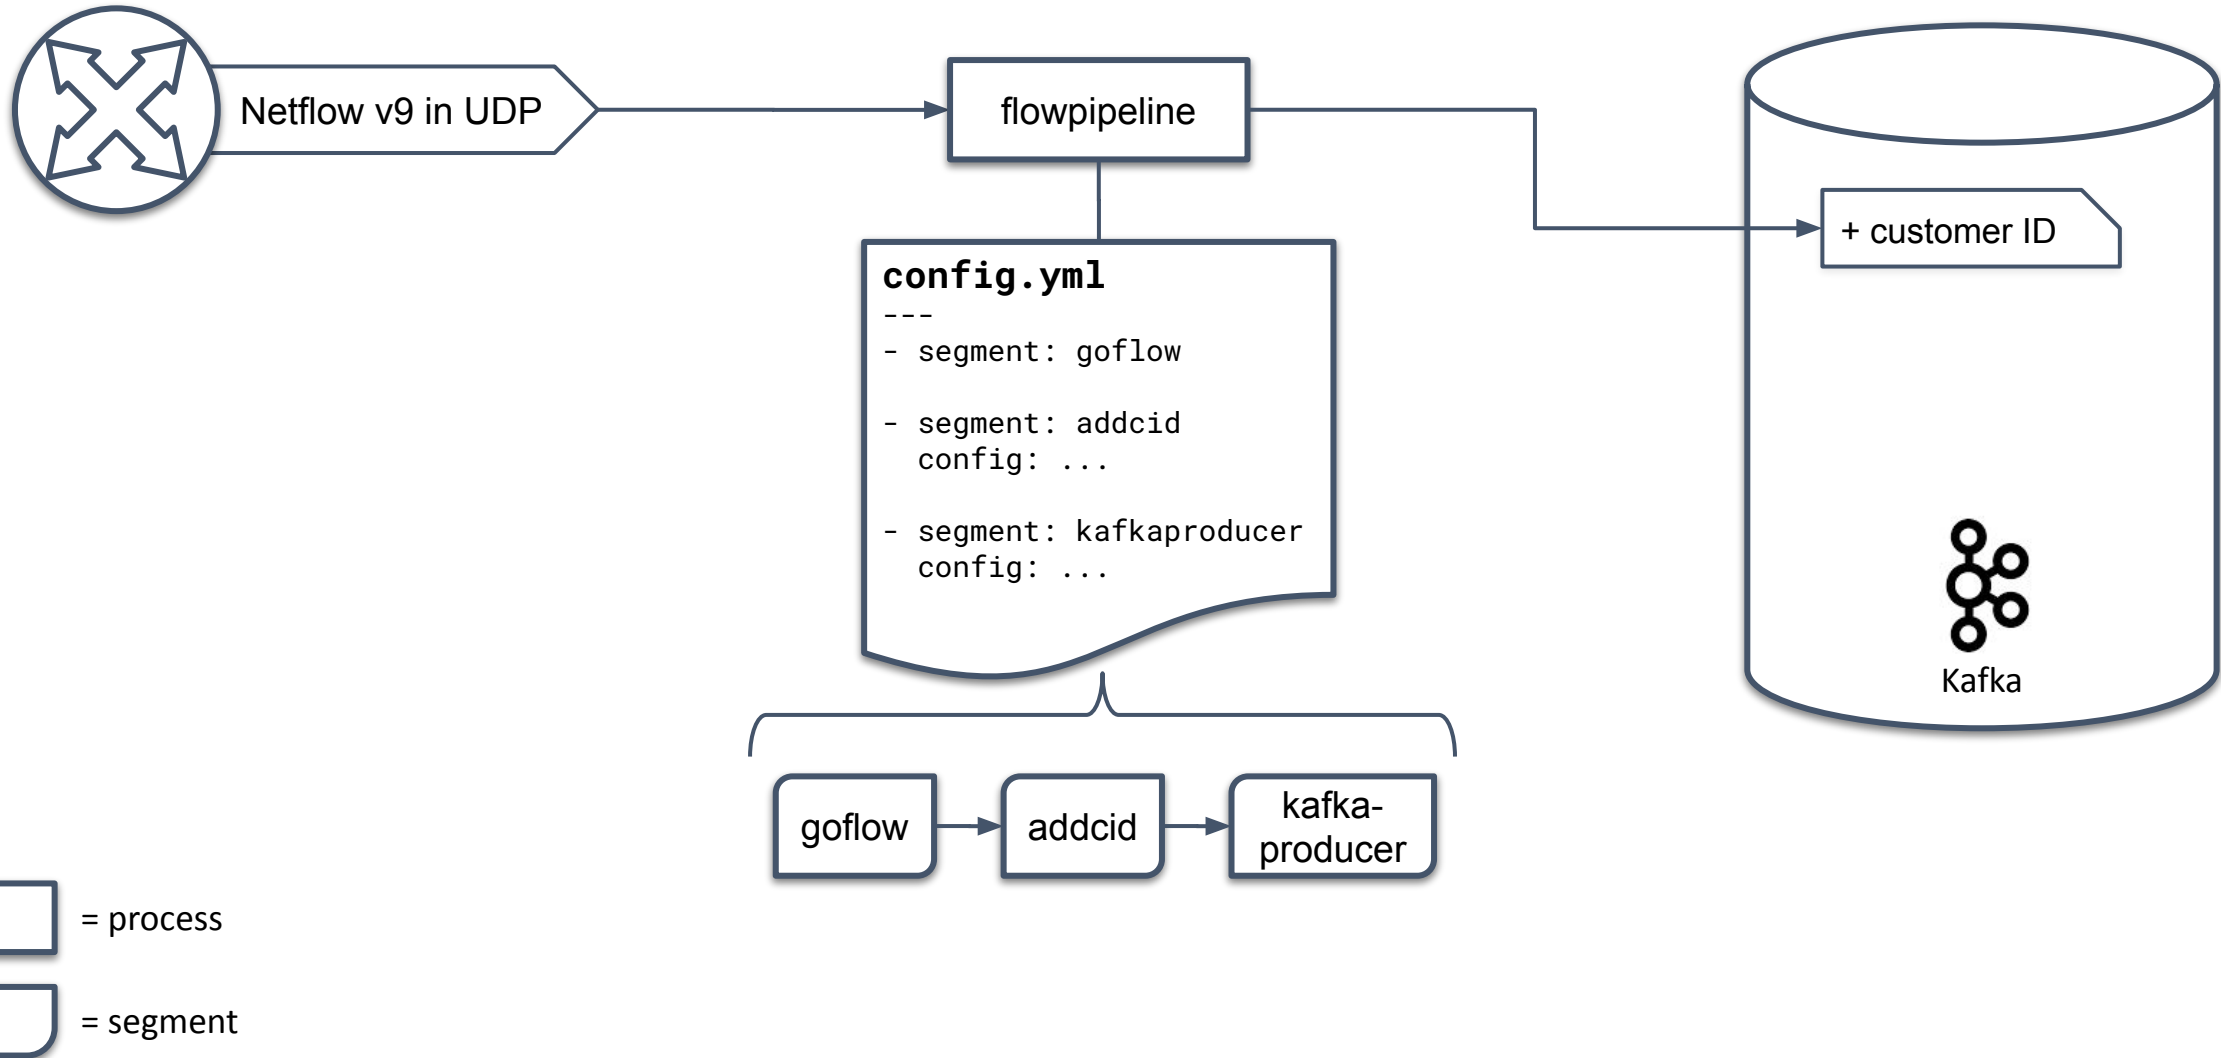


- receive flows generated by another flowpipeline from a Kafka cluster
- flows can be pre-filtered or pre-enriched
- ability to output to Kafka

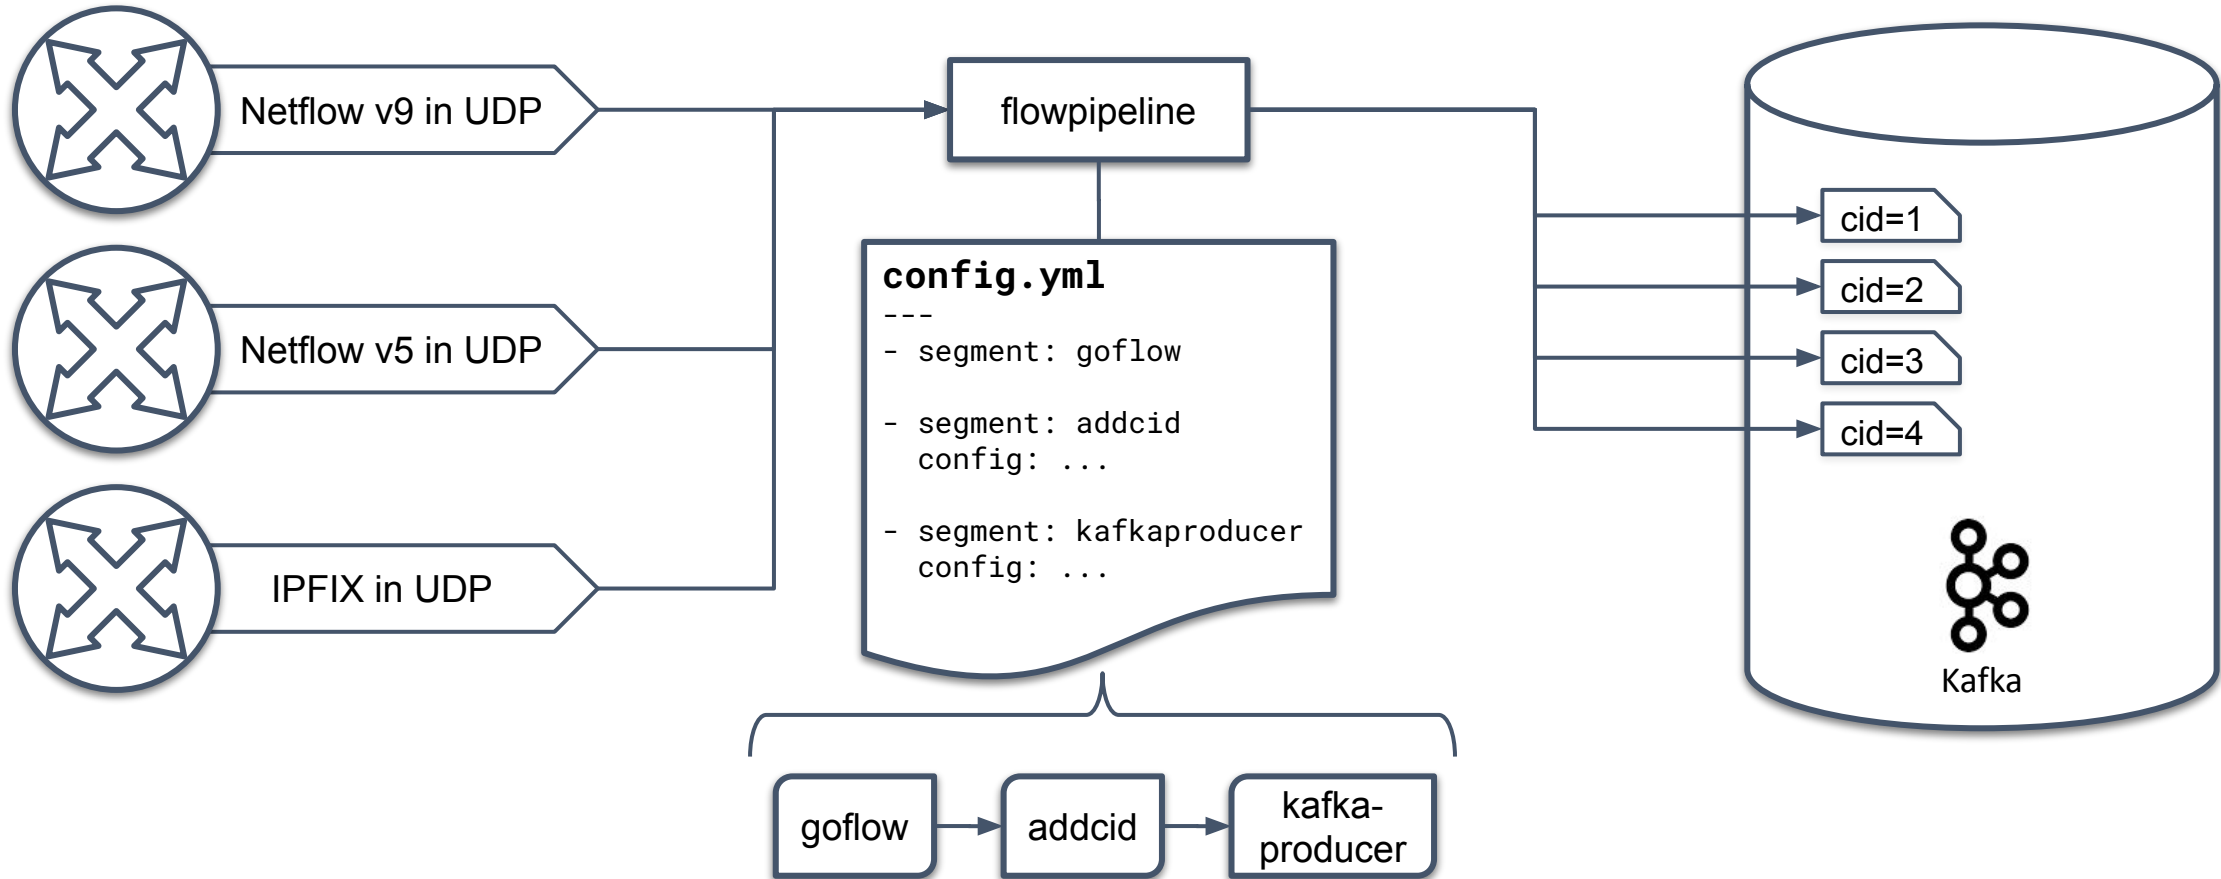


- use eBPF to dump packet headers
- match packets to flows in custom cache using 5-tuple
- additional information available (packet IAT, ...)
- working, but still WIP

Routing Flow Messages



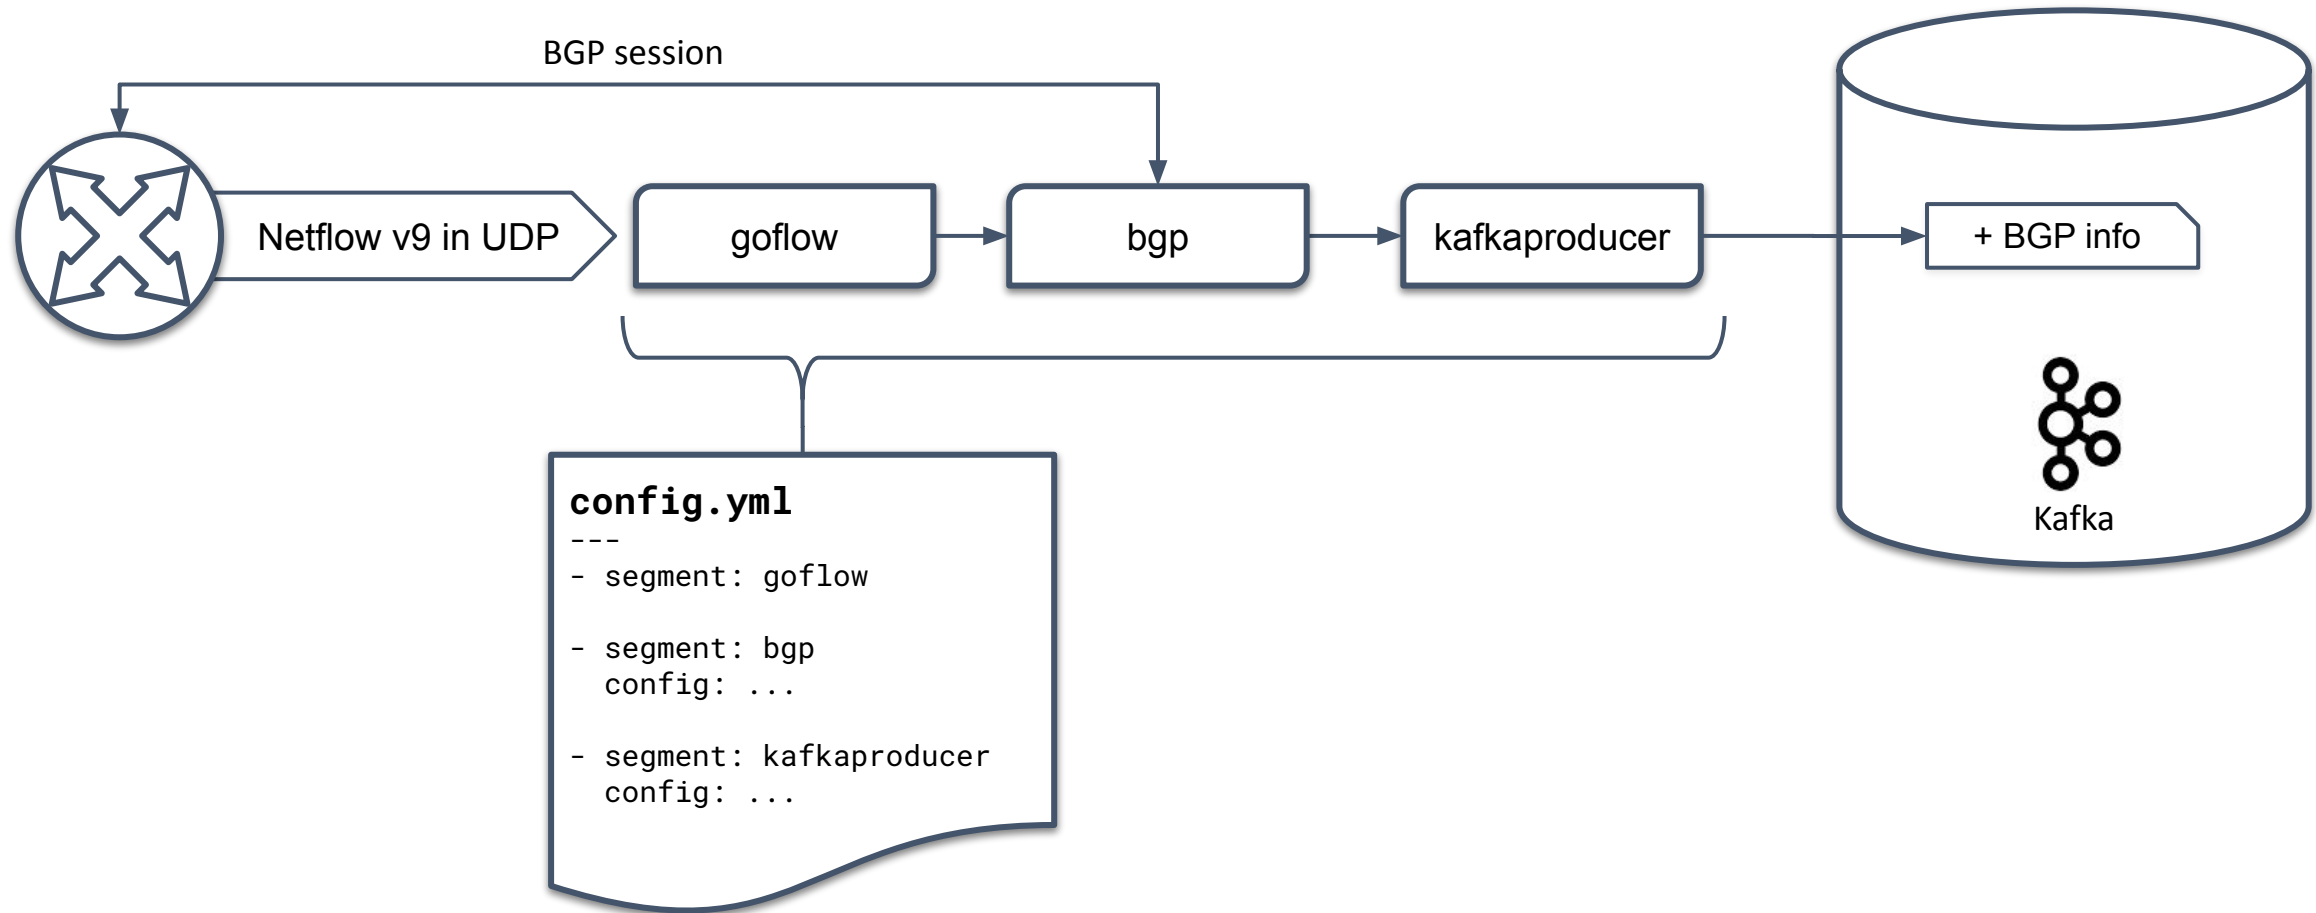
Routing Flow Messages



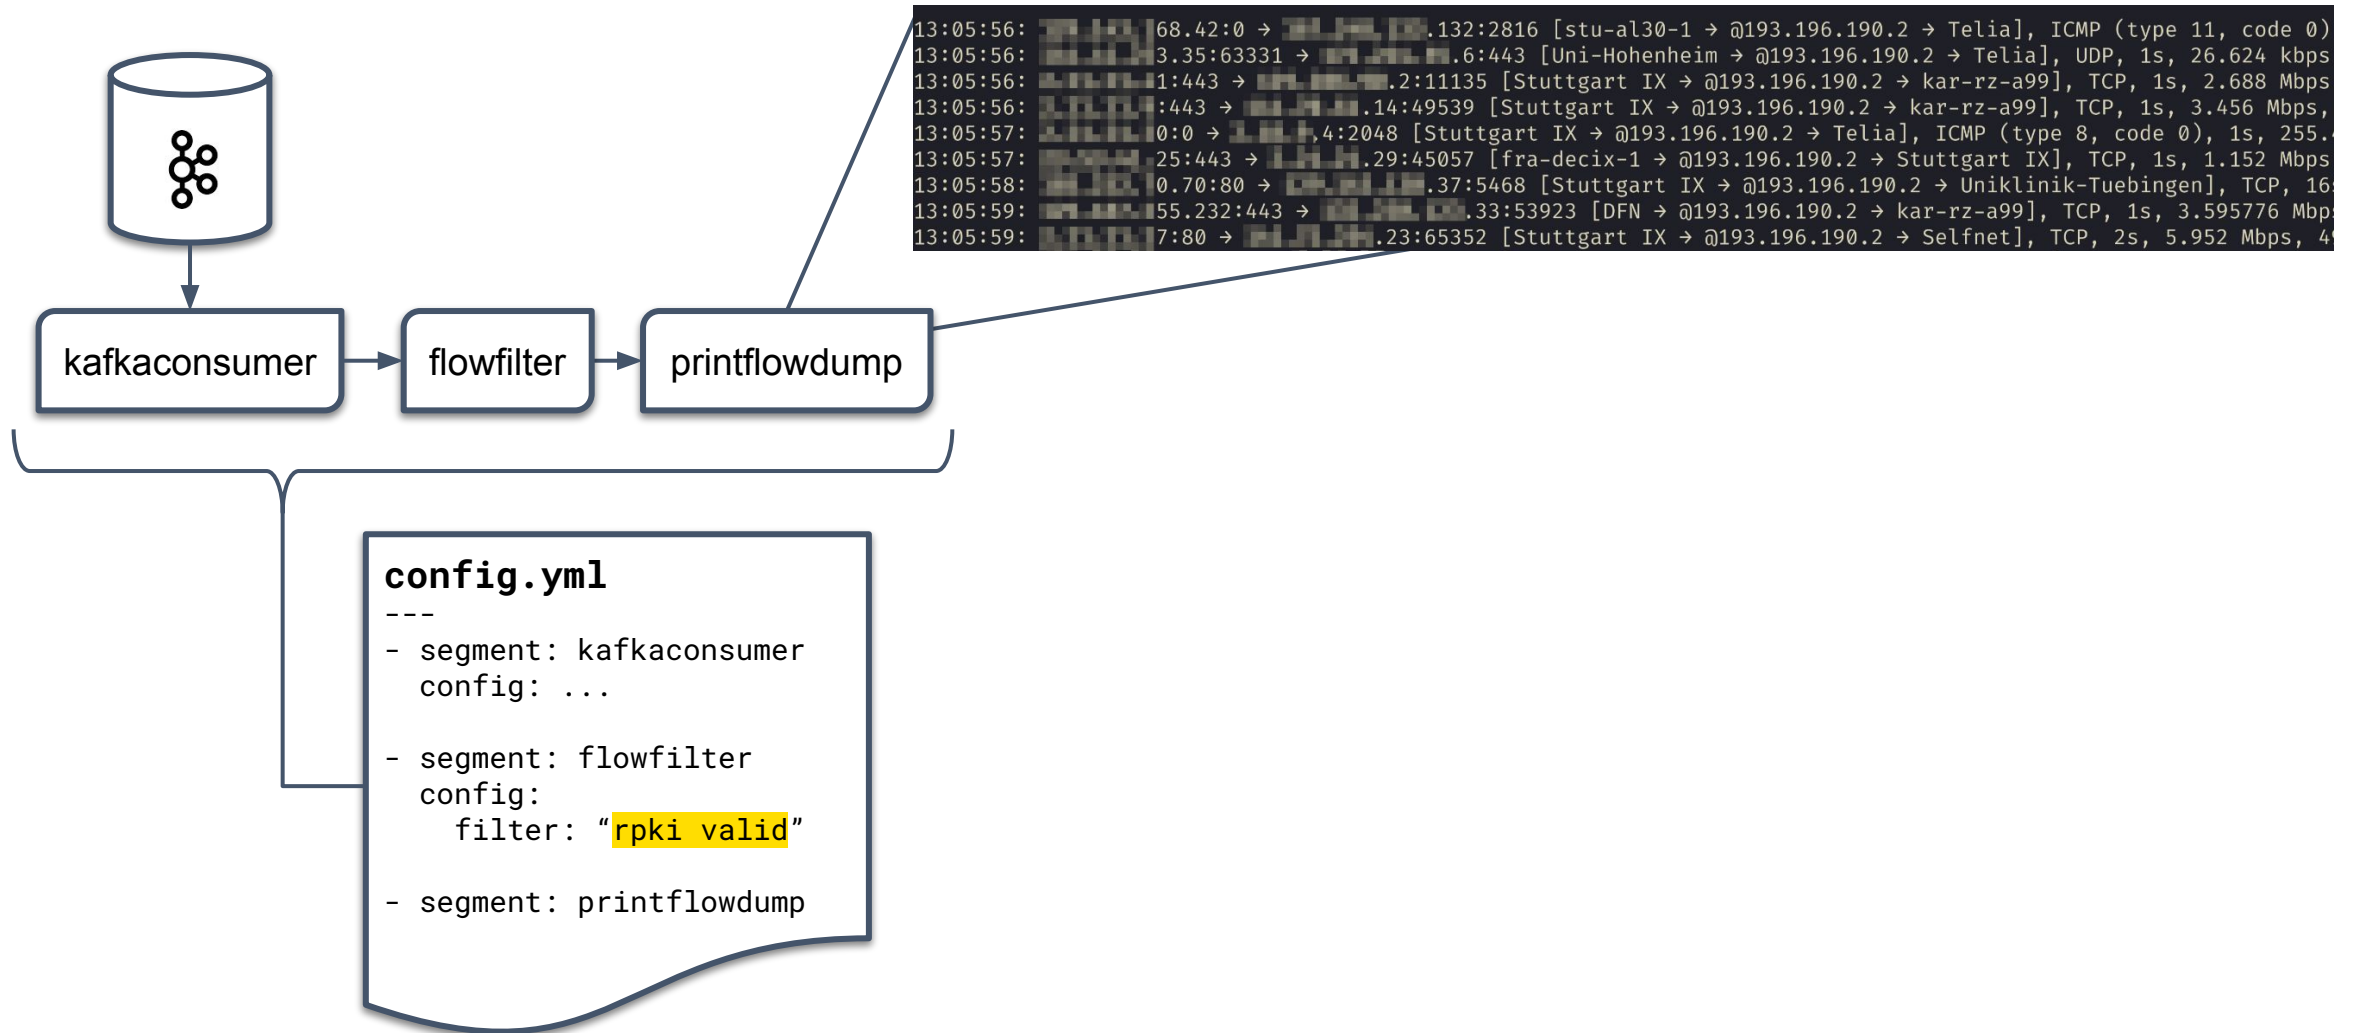
Flow Enrichment and Modification

- Built into pipeline to enable high granularity stream processing
- Options include:
 - **prefix tagging**
 - **BGP info**
 - determine remote
 - geolocation
 - DNS
 - SNMP info
 - normalization
 - **anonymization**
 - **filtering** by different means

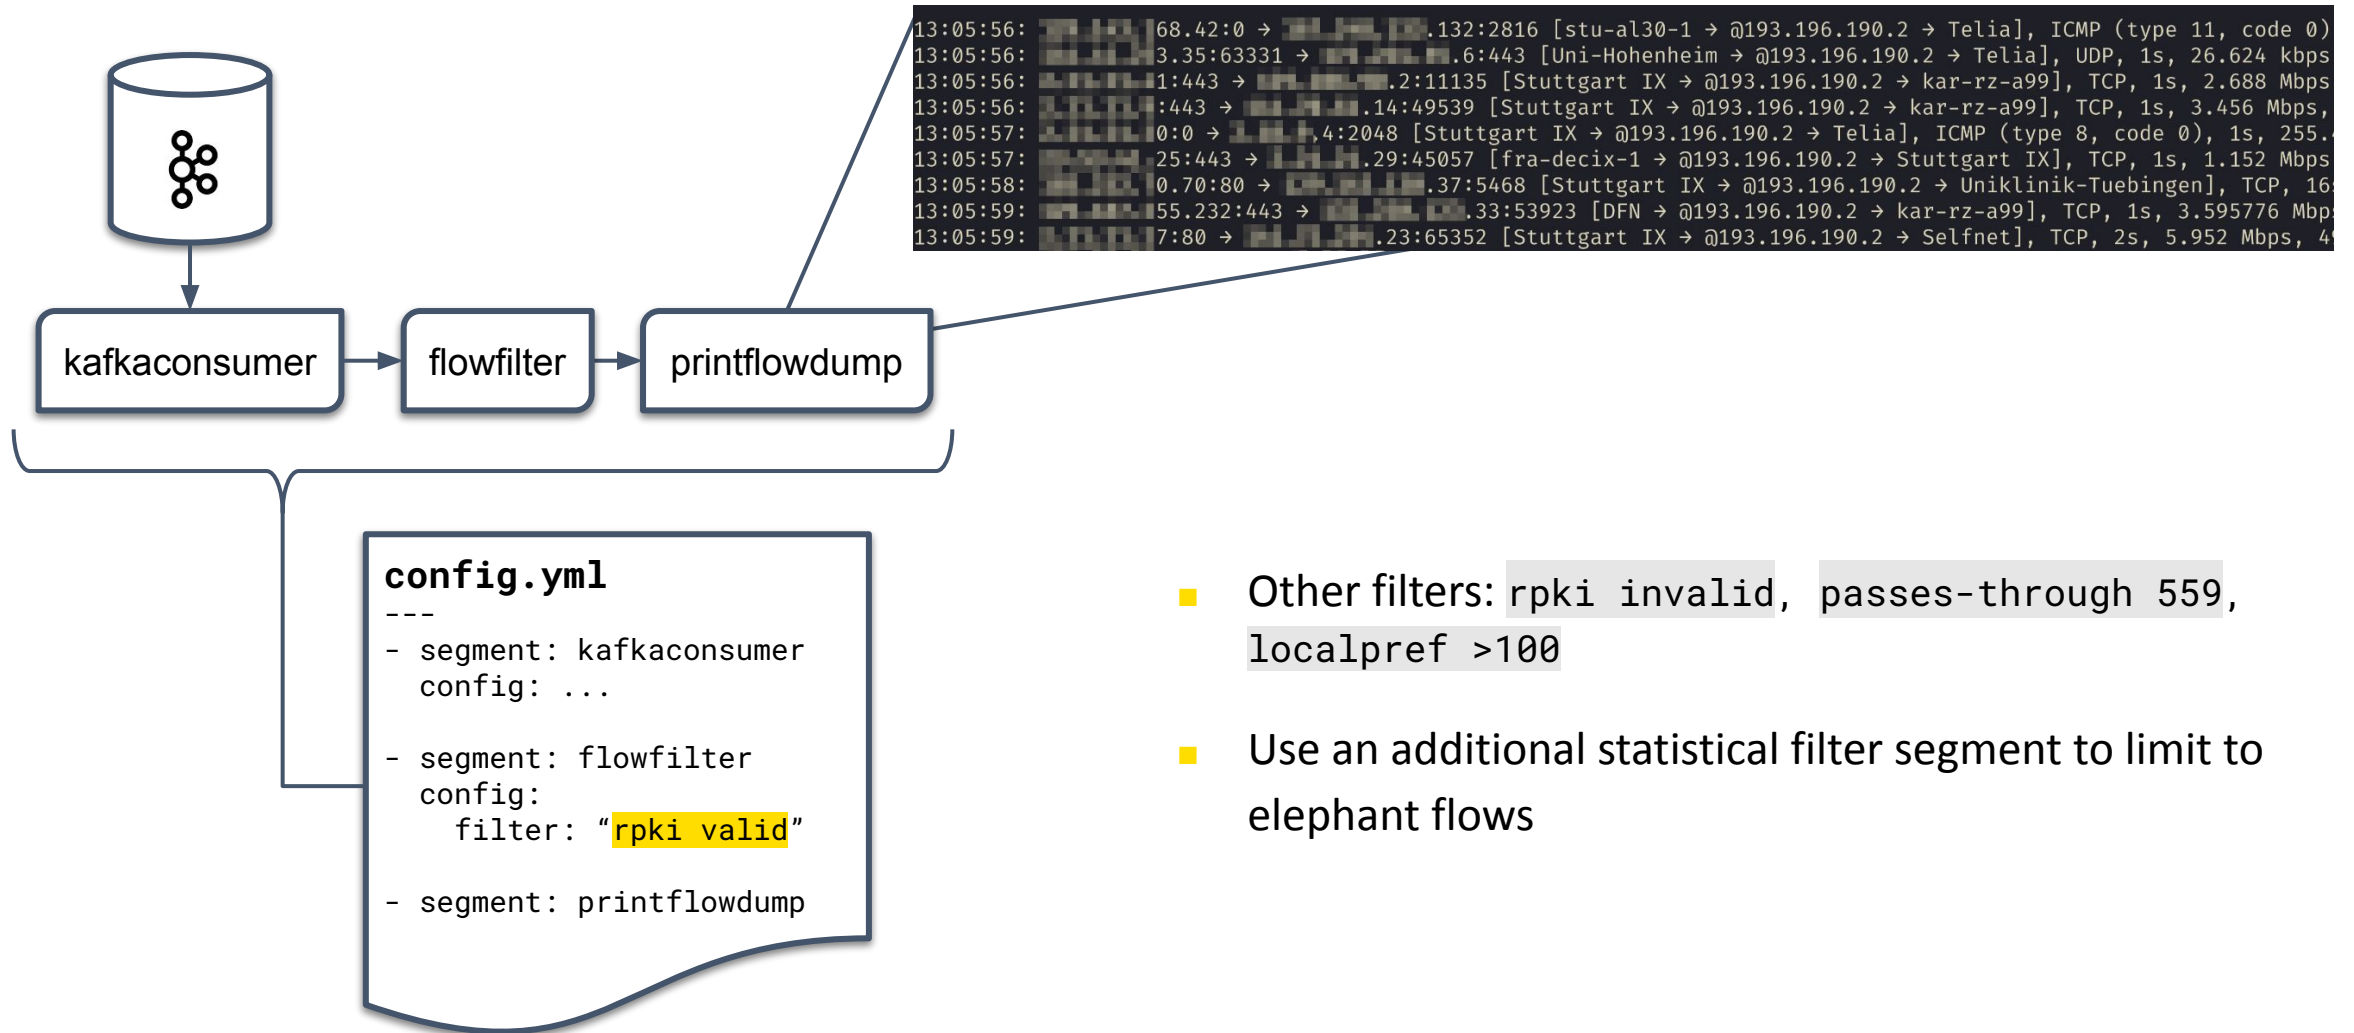
Enrichment with BGP segment



Checking for RPKI (In-)valids

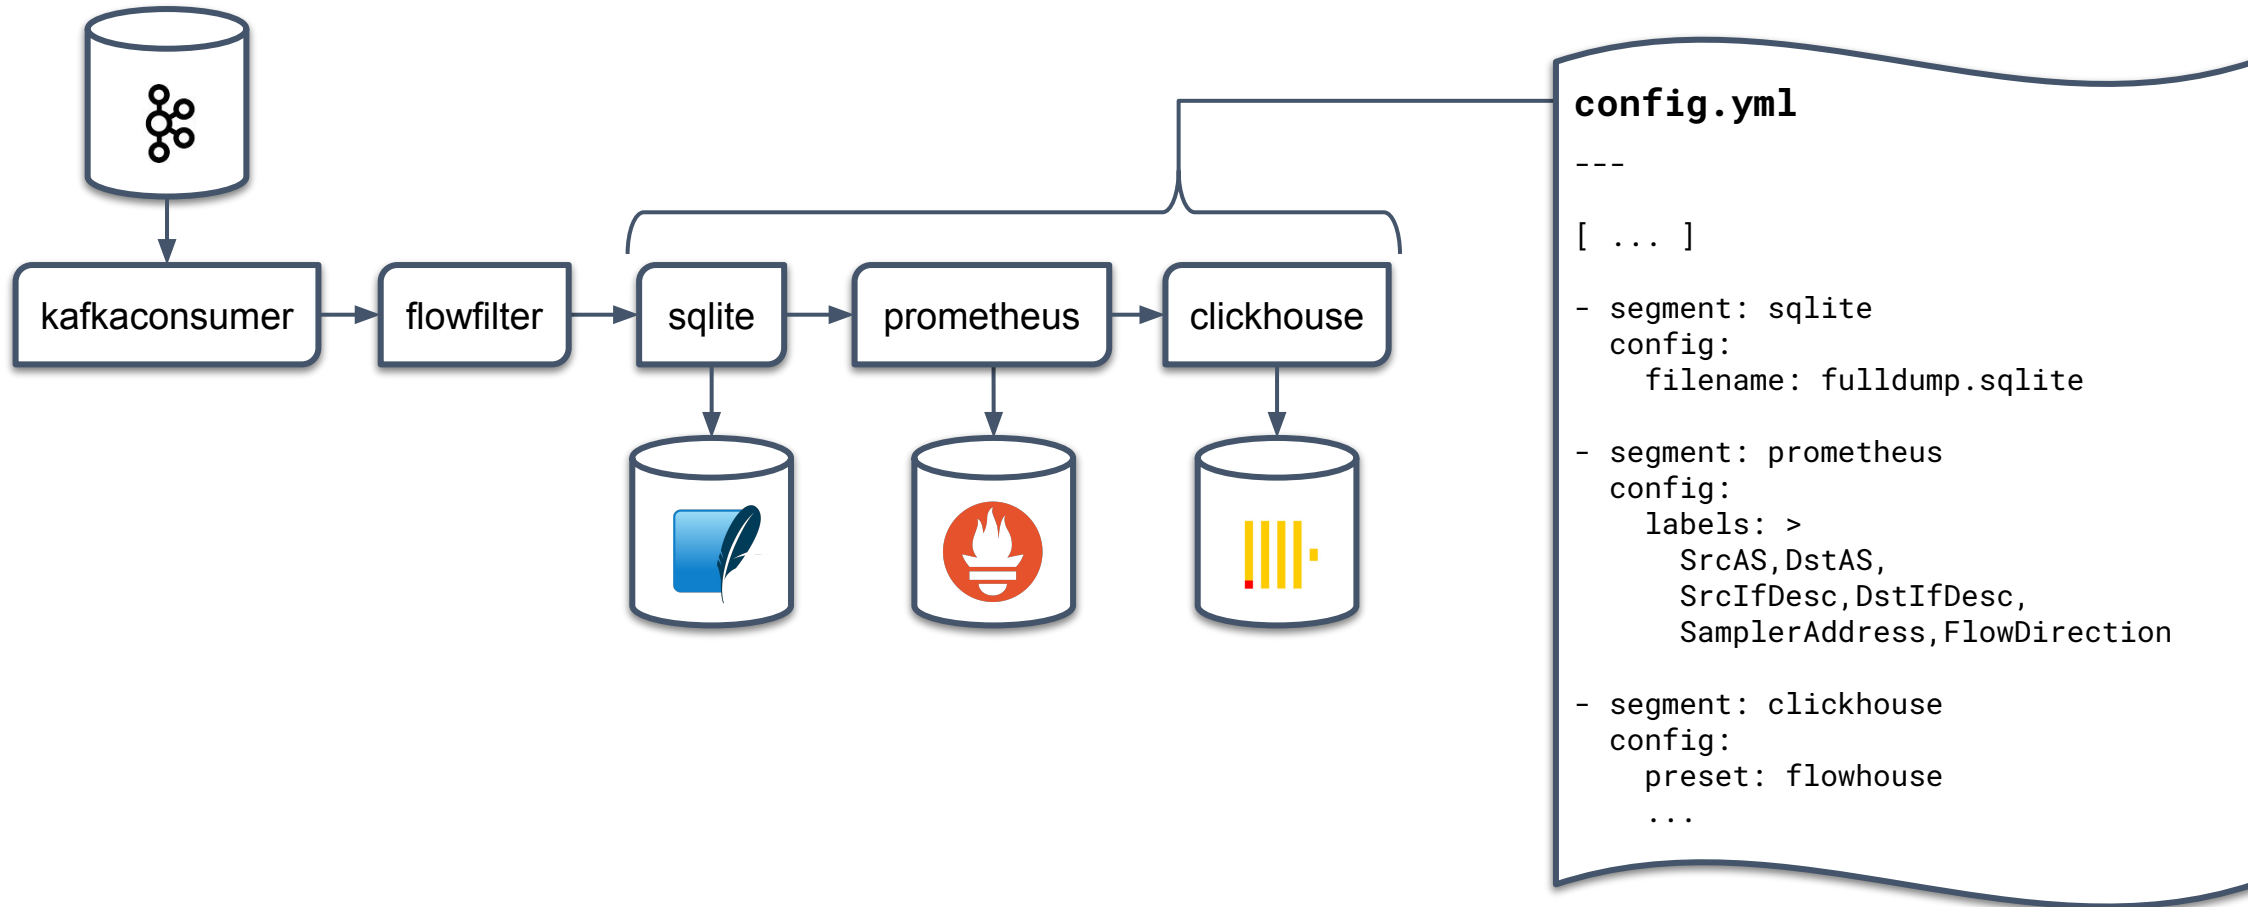


Checking for RPKI (In-)valids

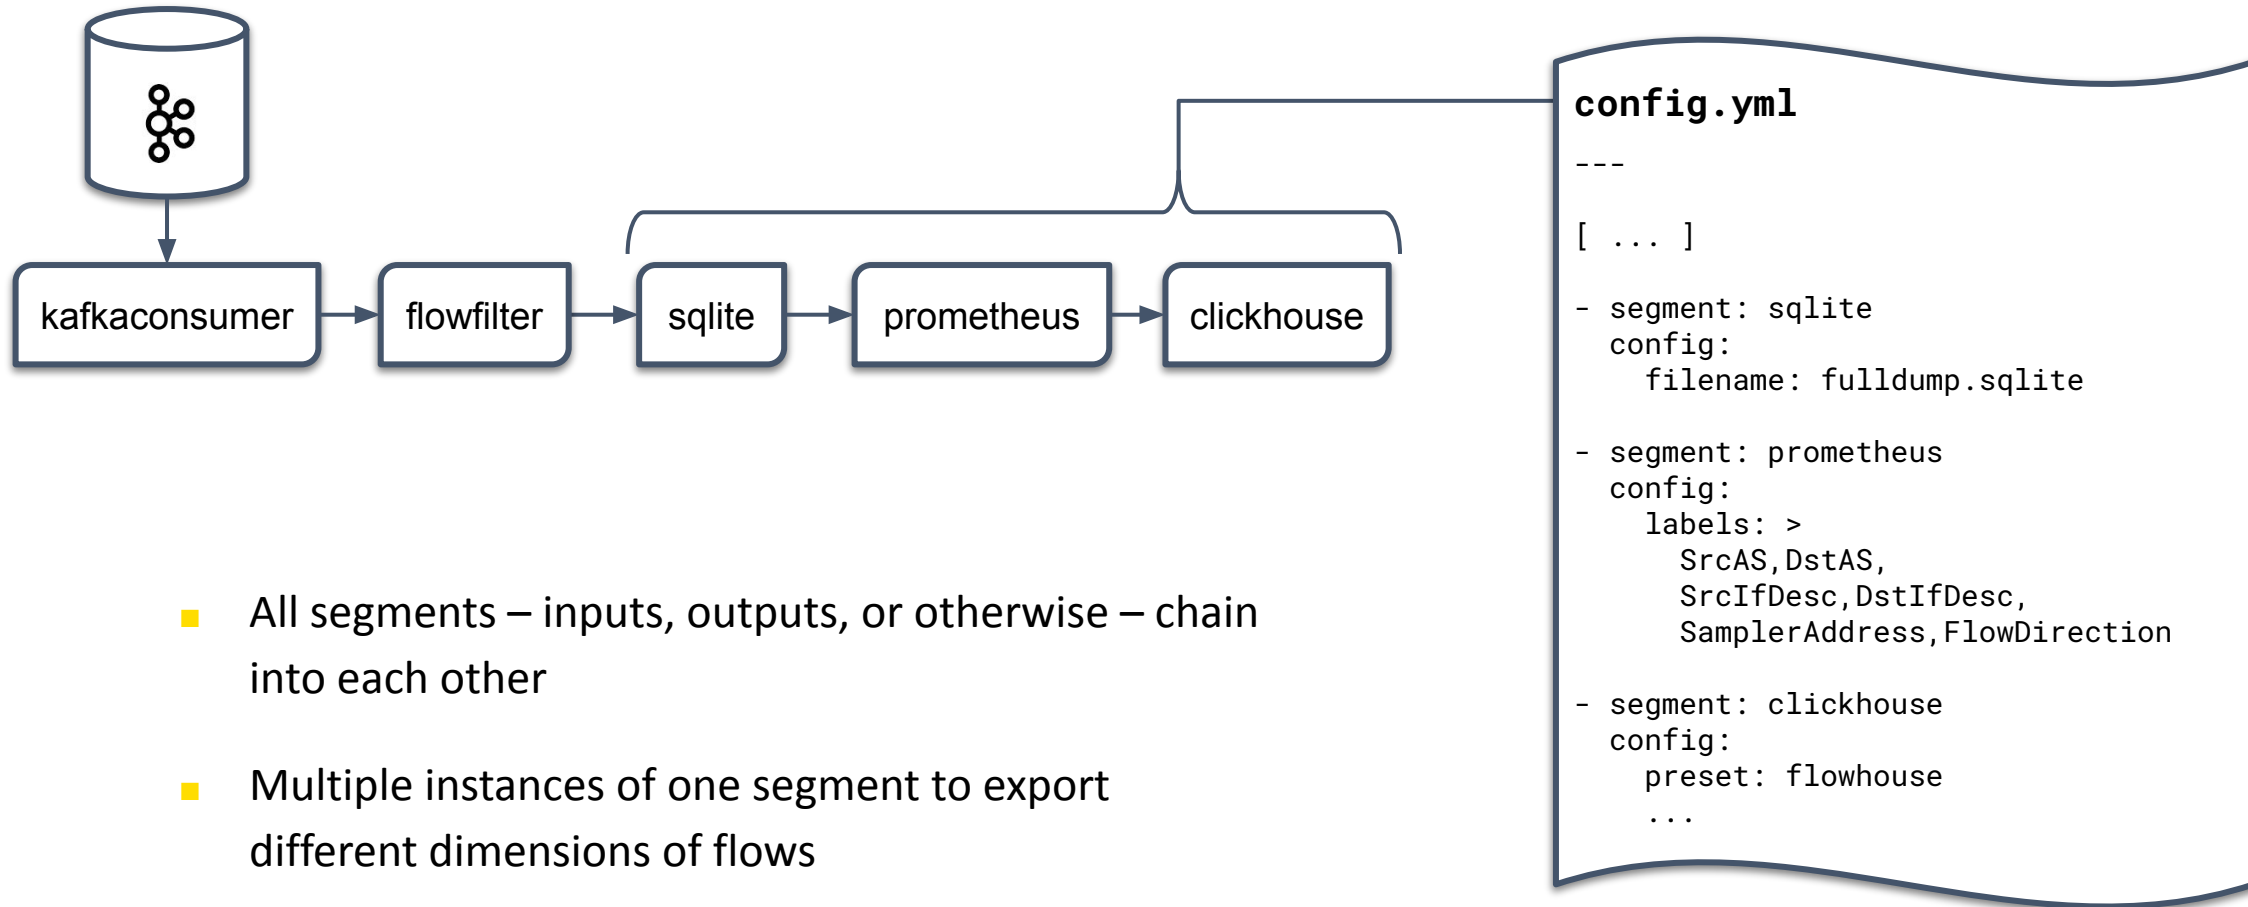


- Other filters: `rpki invalid`, `passes-through 559`, `localpref >100`
- Use an additional statistical filter segment to limit to elephant flows

Doing something useful with those Flows...

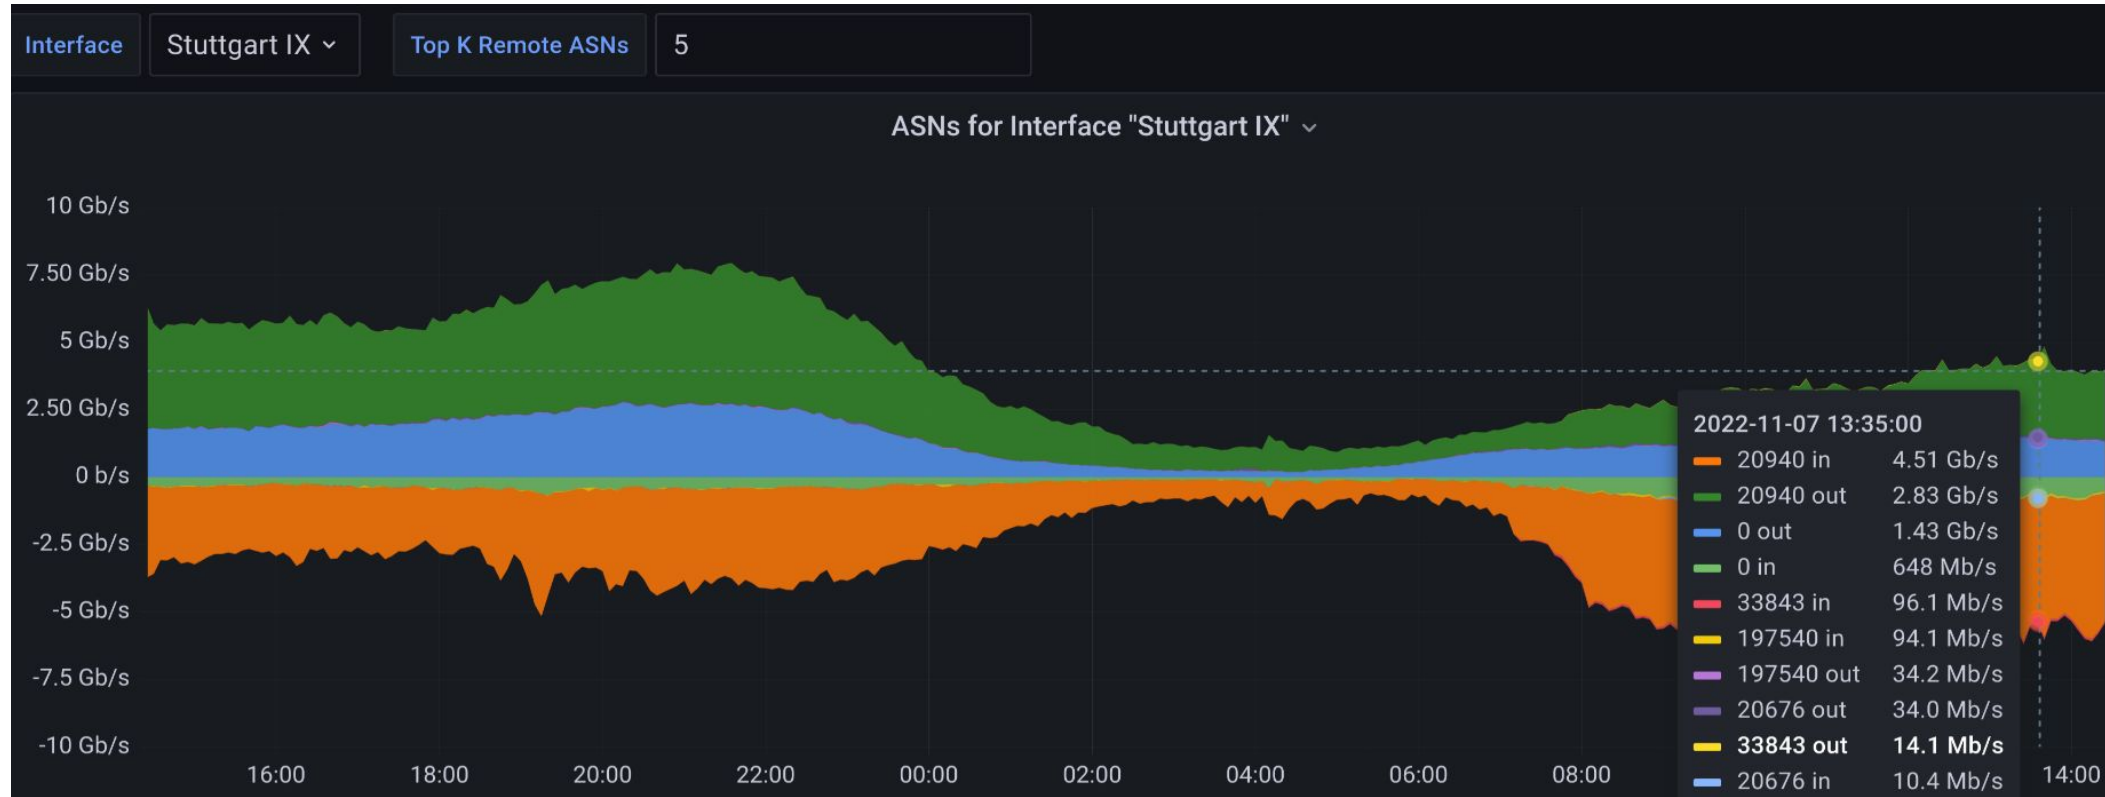


Doing something useful with those Flows...



- All segments – inputs, outputs, or otherwise – chain into each other
- Multiple instances of one segment to export different dimensions of flows

Finding Candidates for private Peering



Determine DoS Recipient Addresses

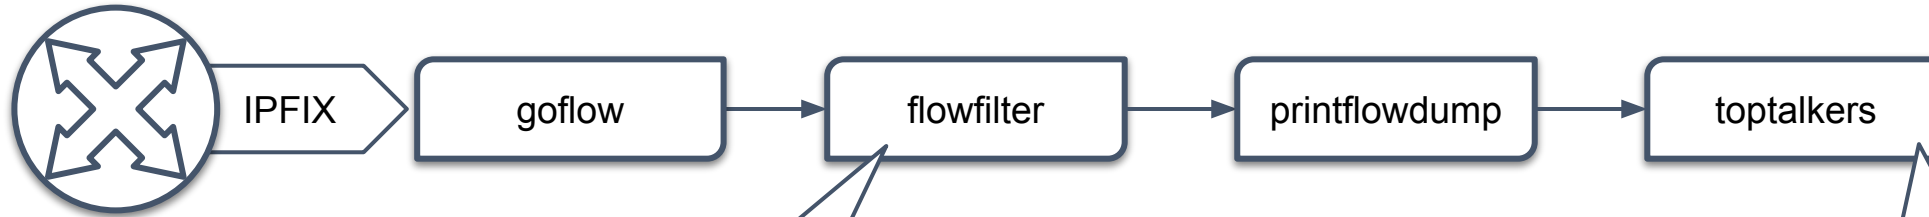


1. Start filtering for what seems to be causing problems:
proto udp and src port 123
2. Include fragmented datagrams:
... and port 0

```

193.192.168.5: 524.24576 Mbps, 350.88 kpps
193.192.168.26: 341.124978 Mbps, 473.2416 kpps
141.70.1.217: 217.036596 Mbps, 174.144 kpps
129.14.1.186: 186.202373 Mbps, 131.2 kpps
193.192.168.171: 171.416842 Mbps, 146.7584 kpps
129.14.1.166: 166.531733 Mbps, 117.386667 kpps
132.21.1.140: 140.068966 Mbps, 93.397333 kpps
192.44.1.134: 134.679093 Mbps, 94.986667 kpps
2a00:1098:1::1: 115.679232 Mbps, 80.3328 kpps
129.14.1.112: 112.629091 Mbps, 86.6944 kpps
  
```

Determine DoS Recipient Addresses

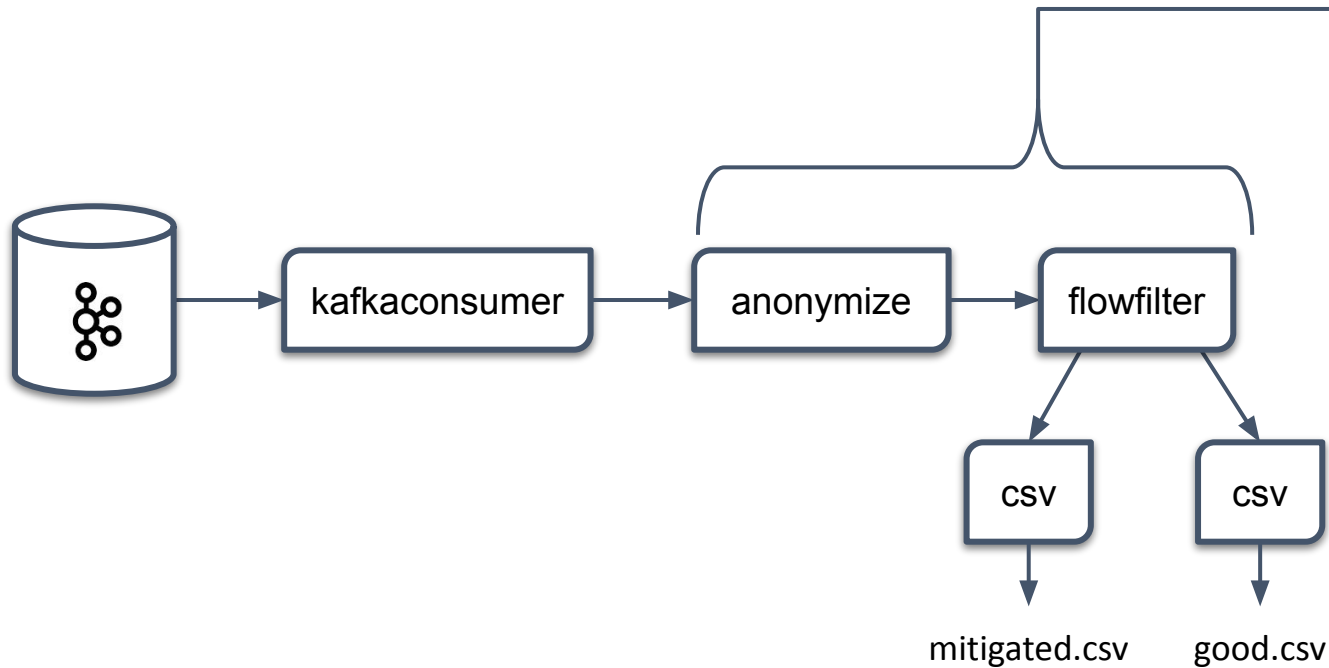


1. Limit results to SYN flooding
`tcpflags syn and not
tcpflags ack and pps >9000`
2. Check whether mitigations are working:
`... and not status dropped`

```

193.192.1.5: 524.24576 Mbps, 350.88 kpps
193.192.1.26: 341.124978 Mbps, 473.2416 kpps
141.70.1.217: 217.036596 Mbps, 174.144 kpps
129.14.1.186: 186.202373 Mbps, 131.2 kpps
193.192.1.171: 171.416842 Mbps, 146.7584 kpps
129.14.1.166: 166.531733 Mbps, 117.386667 kpps
132.21.1.140: 140.068966 Mbps, 93.397333 kpps
192.44.1.134: 134.679093 Mbps, 94.986667 kpps
2a00:1.1.1.1: 115.679232 Mbps, 80.3328 kpps
129.14.1.112: 112.629091 Mbps, 86.6944 kpps
  
```

Labeled Datasets for Researchers



config.yml

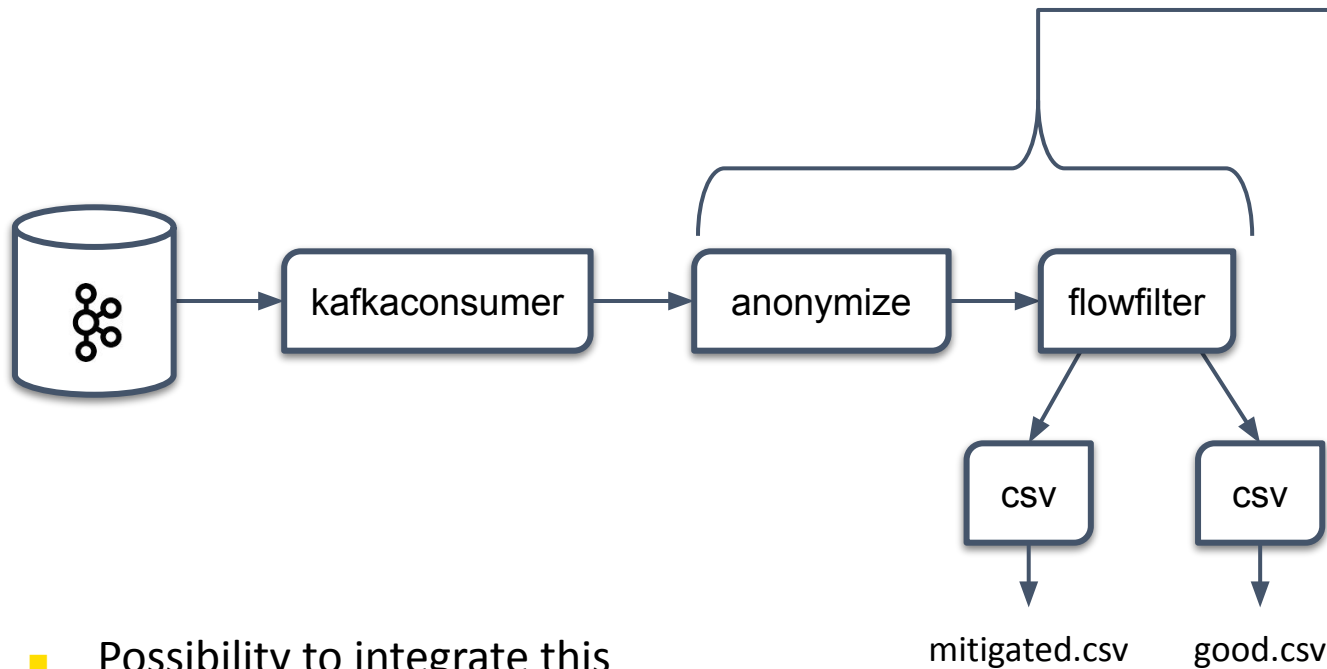
```

---
[ ... ]

- segment: anonymize
  config:
    key: "qwertyuiop"
    fields: SrcAddr,DstAddr

- segment: branch
  if:
    - segment: flowfilter
      config:
        filter: "... and status dropped"
    then:
      - segment: csv
        config:
          filename: mitigated.csv
    else:
      - segment: csv
        config:
          filename: good.csv
  
```

Labeled Datasets for Researchers



config.yml

```
---
[ ... ]

- segment: anonymize
  config:
    key: "qwertyuiop"
    fields: SrcAddr,DstAddr

- segment: branch
  if:
    - segment: flowfilter
      config:
        filter: "... and status dropped"
    then:
      - segment: csv
        config:
          filename: mitigated.csv
    else:
      - segment: csv
        config:
          filename: good.csv
```

- Possibility to integrate this during the live DoS drilldown example
- Multiple export formats at the same time

Extensibility and Integration

- Flowpipelines are extensible using custom segments, powered by the Go plugin system
- Allows implementation of **any** algorithm, data structure, or filtering
- Pluggable into existing configurations

Custom Segments

1. Code in the provided template (or copy a stock segment)
2. Compile it: `go build -buildmode=plugin ./custom.go`
3. Use the assigned name in a config and launch with `-p custom.so`

```
1 for msg := range segment.In {  
2     // [ ... ]  
3     segment.Out ← msg  
4 }
```

Thank you for your attention!

Questions?

<https://github.com/bwNetFlow/flowpipeline>



Daniel Nägele – `naegele@belwue.de` – `@debugloop` (on IRC & social)