What's hiding behind IPv6 extension headers?

MAT-WG, RIPE73, Madrid, Spain

Luuk Hendriks
Design and Analysis of Communication Systems

UNIVERSITY OF TWENTE.

SURF NET

What's wrong?

Flow-based measurements are based on a *key*, made from specific header fields.

Classic 5-tuple:
L3 src/dst, L4 sport/dport, proto

The classic tuple

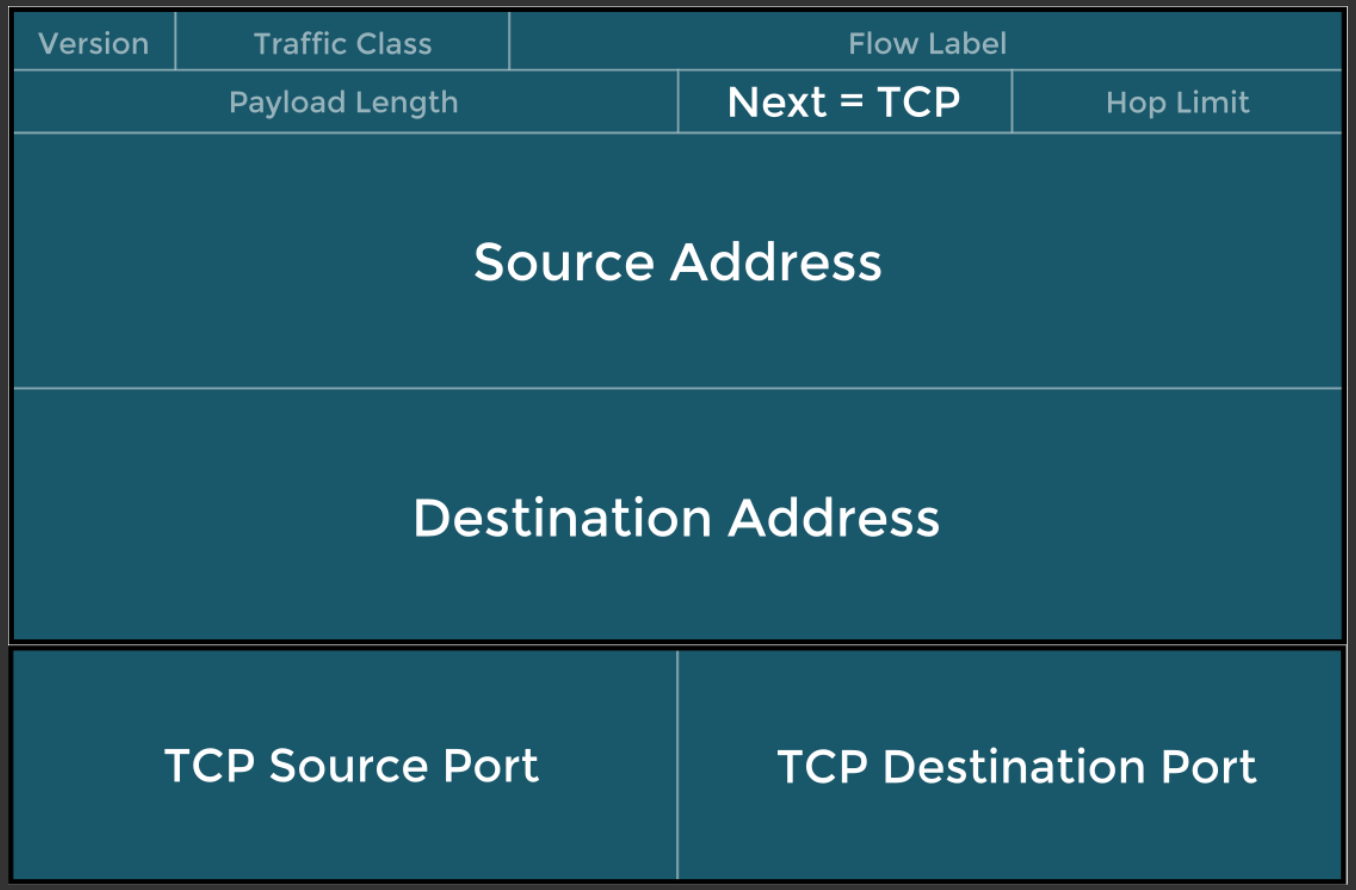

Enter Extension Headers

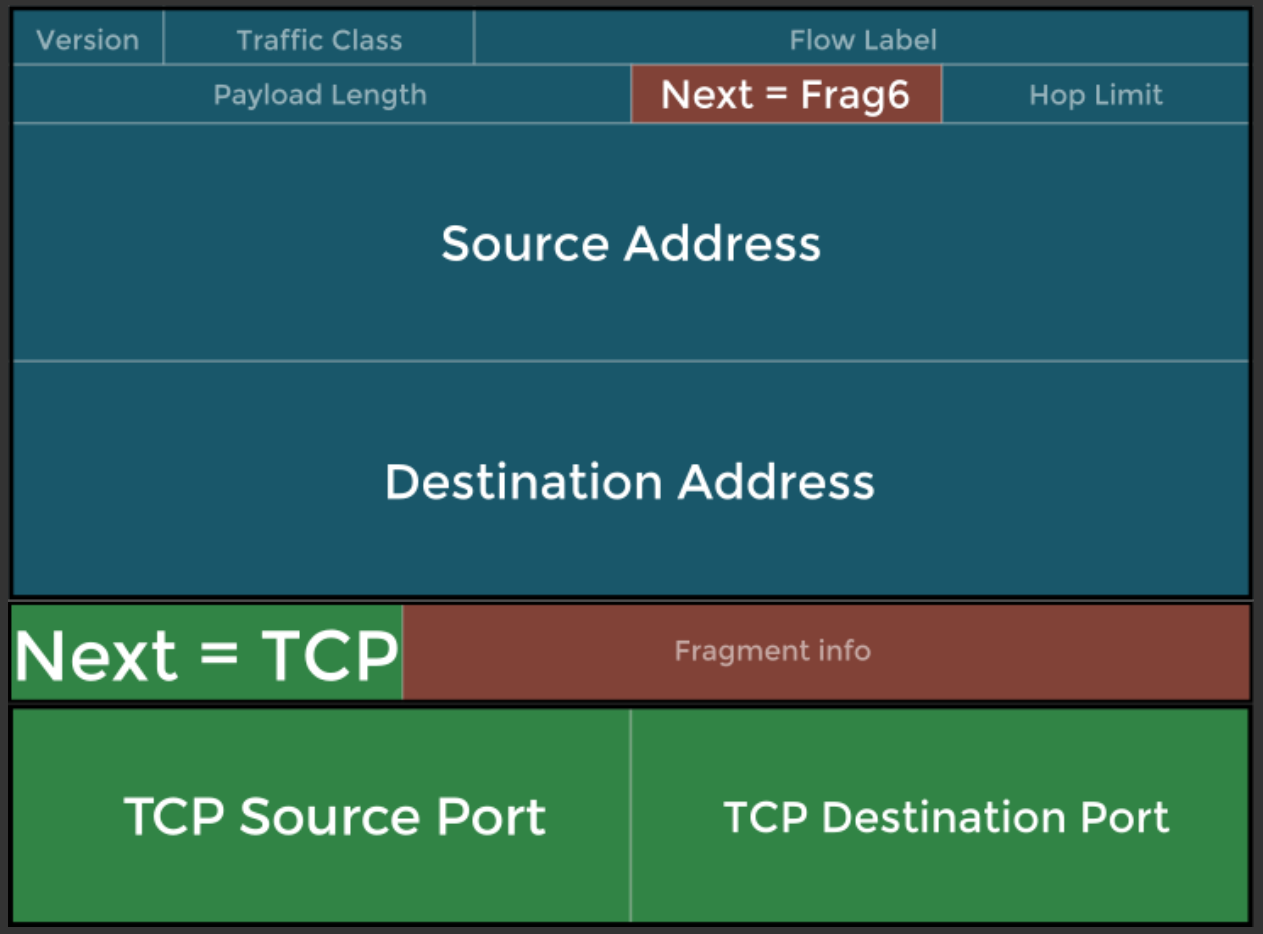

When using flow-based measurements,

Extension Headers in IPv6
are **hiding information** on
the **actual upper layer**.

```
Proto          Source address                port
Ipv6-Frag      2001:db8:1:0:4777::140        0
               Destination Address           port
               2001:db8:db8:a120::17         0

pkt    bytes       flows
8       9792        1
```

```
Proto          Source address                port
Ipv6-Frag      2001:db8:1:0:4777::140        0
               Destination Address           port
               2001:db8:db8:a120::17         0

pkt   bytes      flows
8      9792       1
```

```
Proto            Source address                port
Ipv6-Frag        2001:db8:1:0:4777::140        0
                 Destination Address           port
                 2001:db8:db8:a120::17         0

pkt    bytes        flows
8        9792         1
```

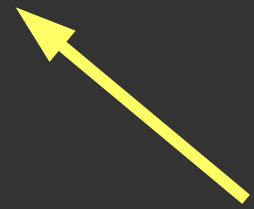

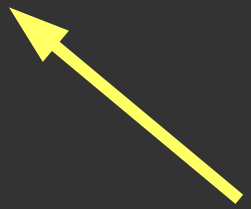

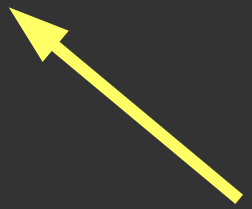

?

?

What's hidden then?

- Actual upper layer proto
- Actual upper layer sport/dport
- All extension headers after the first one

Furthermore,

- Wrongful aggregation
  hides actual byte/packet/flow counts

Challenges in flow-land

- How can we get the hidden information?
  - Export new fields! But what fields?
- How can we fix the wrongful aggregation?
  - Use a different cache key! But what fields?
- Any collector-side changes?

Implementation

We implented a Flowmon plugin to export

- Upper Proto/~sport/~dport
- Extension header list/~total size

Adapted cache key to include upperProto, upperSport, upperDport

Measurement at CESNET

- May 2016
- 10 links, our plugin on FlowMon probes
- IPv6 flows only
- Unsampled
- Anonymized IP addresses
- 1 single collector

~4000M IPv6 flows

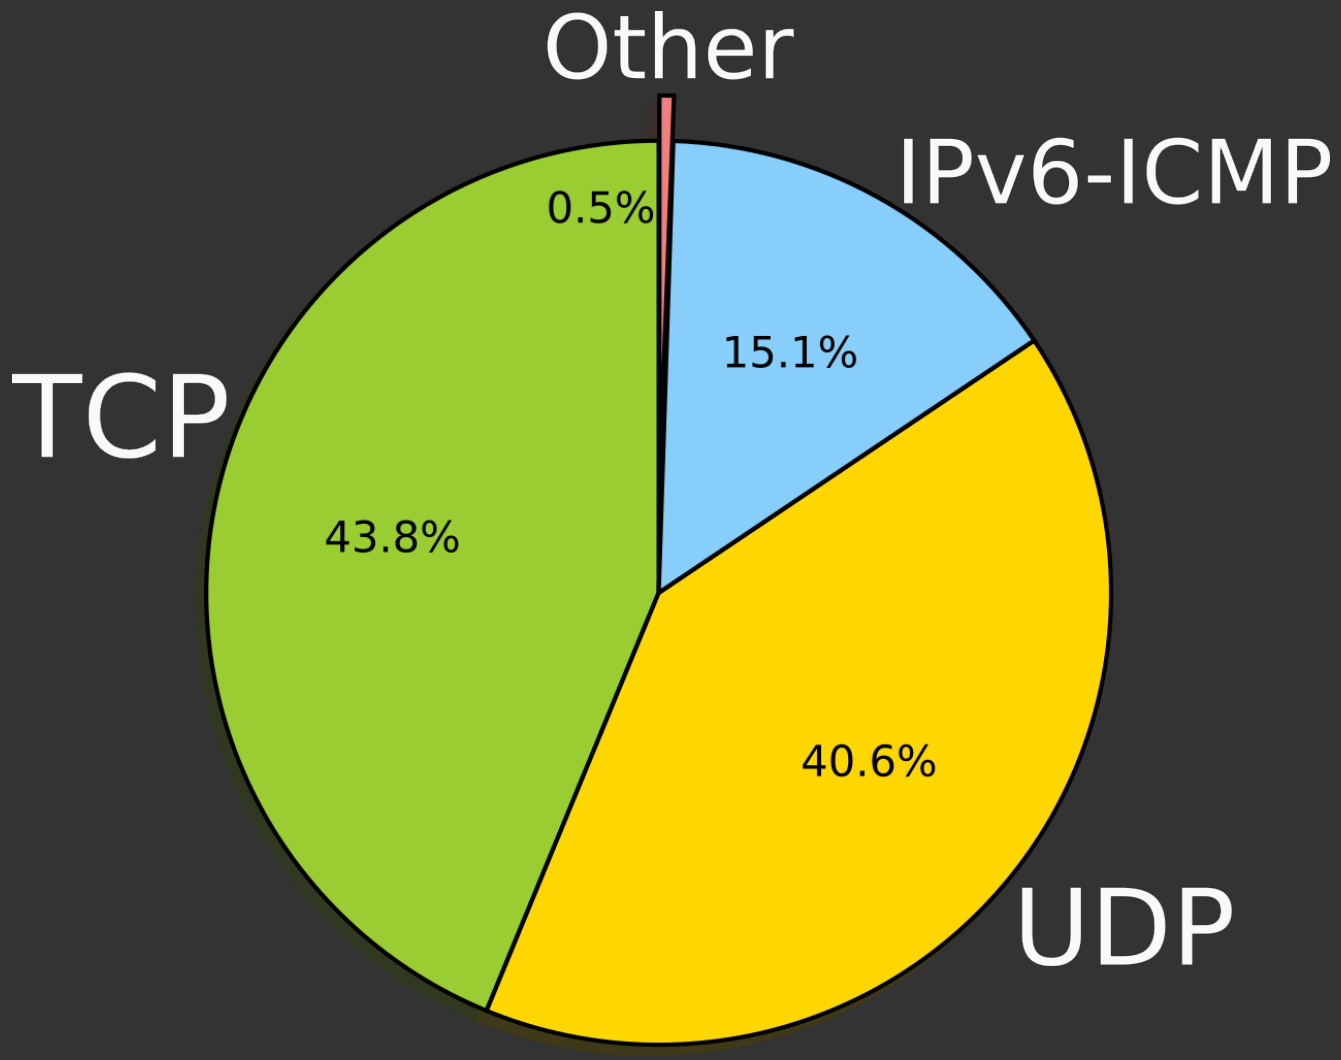

Flows with an Upper Proto

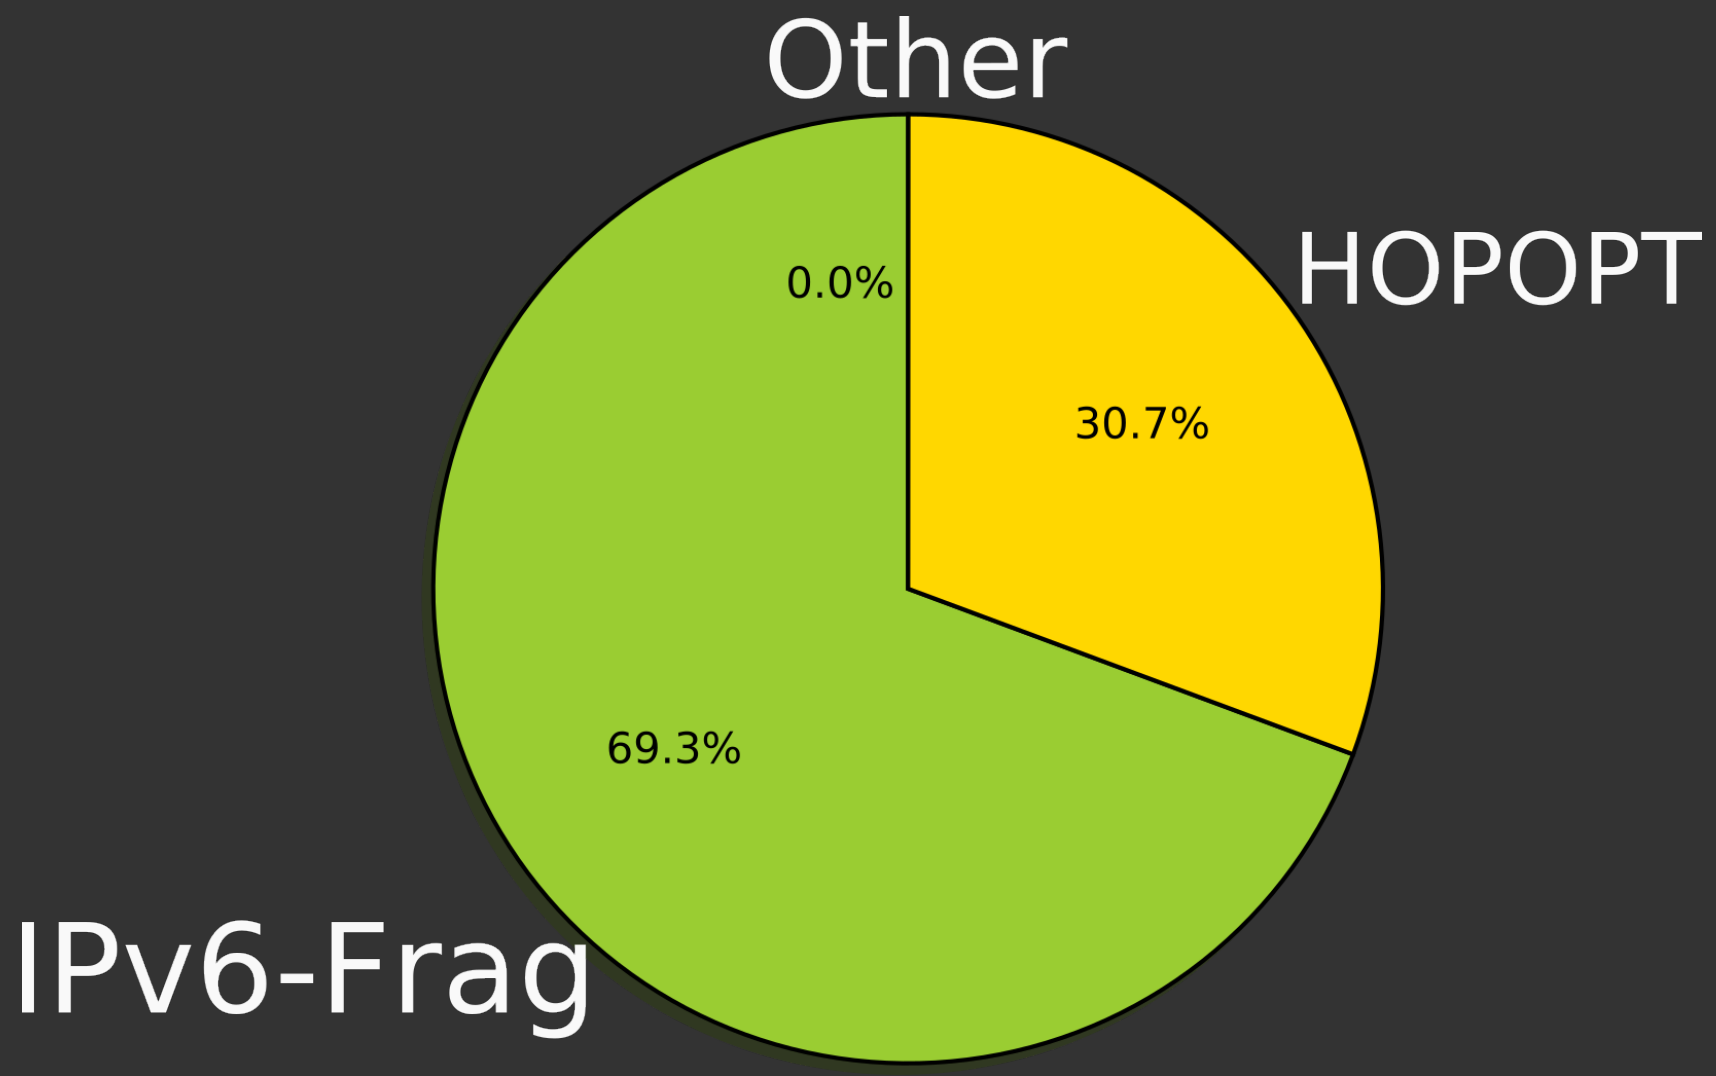

Distribution of Upper TCP ports

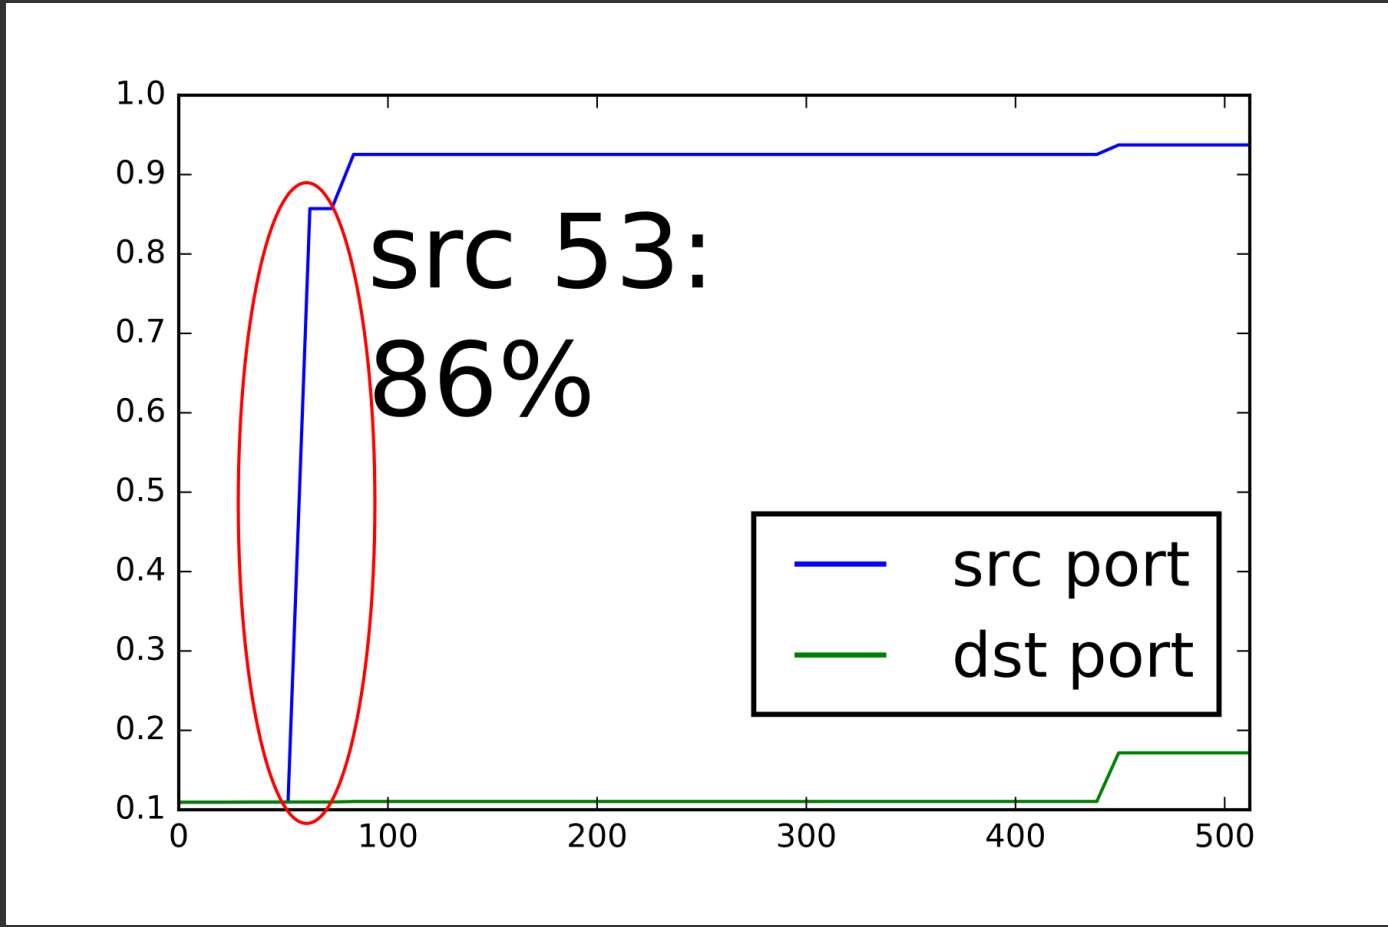

Concluding, ...

Share of flows with EHs, is not that big.
However, actual higher layer payload is often **important for (end-user) QoE**, e.g. DNS.

Measurement technologies need to **traverse the Extension Header chain**, in order to give **correct** and realistic **results**.

Thank you

Petr Vlan (CESNET)

for support in both plugin development and deployment

Open source software:
Ipfixcol / fbitdump

What's hiding behind IPv6 extension headers?

As presented at MAT-WG,
RIPE73, Madrid, Spain

Luuk Hendriks
luuk.hendriks@utwente.nl

UNIVERSITY OF TWENTE.

SURF NET