

# Flow measurements using P4

## University of Twente

Research on Networks 2018

Luuk Hendriks

Wouter de Vries

Aiko Pras

# Measurements measurements measurements

- **History of researching and using flow measurements**
- **First NetFlow, then IPFIX**
  - IPFIX supposed to be flexible
  - True for the protocol, not for deployments per se
  - Result: vendor lock-in like restrictions, limited deployment

# Flow measurements using P4?

**P4 could become the generically applicable solution we want!**

- **Open standard**
- **Open(-ish) hardware**
- **Not restricted to a specific type of hardware e.g. dedicated flow probes / exporters**

# Research questions

- 1) How can we realise flow measurements using P4?**
- 2) To which extend can in-band telemetry be used to enhance or validate P4 flow measurements?**
- 3) How do these flow measurements compare with other technologies, e.g. NetFlow/IPFIX, qualitatively and/or quantitatively?**


# Testbed

**We focus on implementations and measurements on iron, using production traffic where possible.**

- **32x100G Barefoot switch**
- **2x10G Netronome SmartNIC**
- **2x40G Intel DumbNIC**
- **Mirror of 40G UT→ SURFnet uplink**

# Planned rack layout

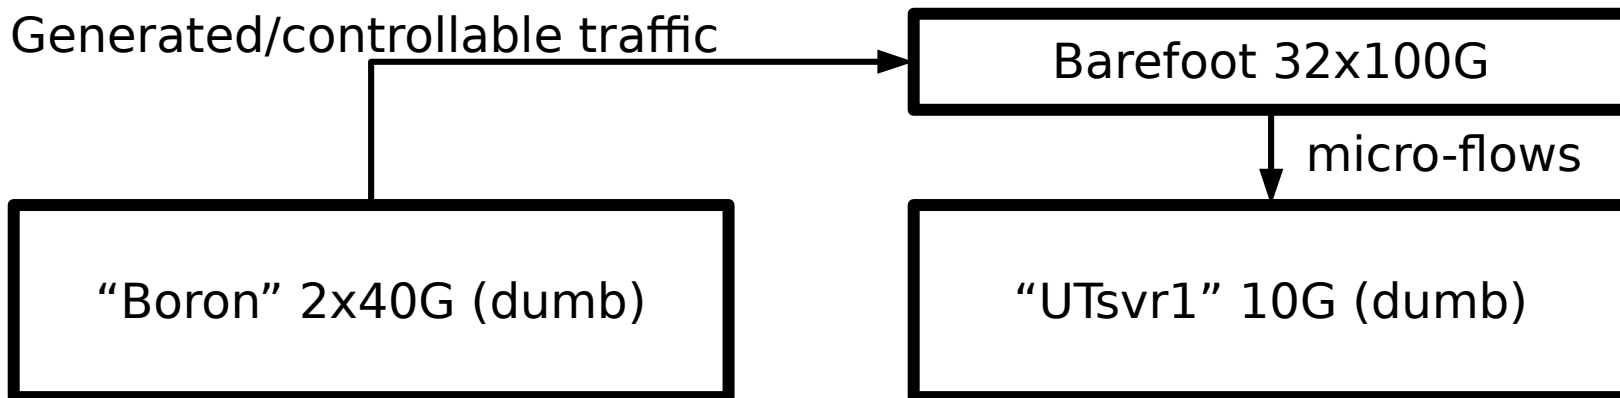
UT uplink (40G)



# Scenario #1: Flow metering

**Based on prior work (John Sonchack / Open-NFP) we might consider ‘two-stage’ flow metering:**

- 1. Use P4 pre-aggregate packets into ‘micro-flows’**
- 2. Create flow-records from these micro-flows on a separate (x86) machine**

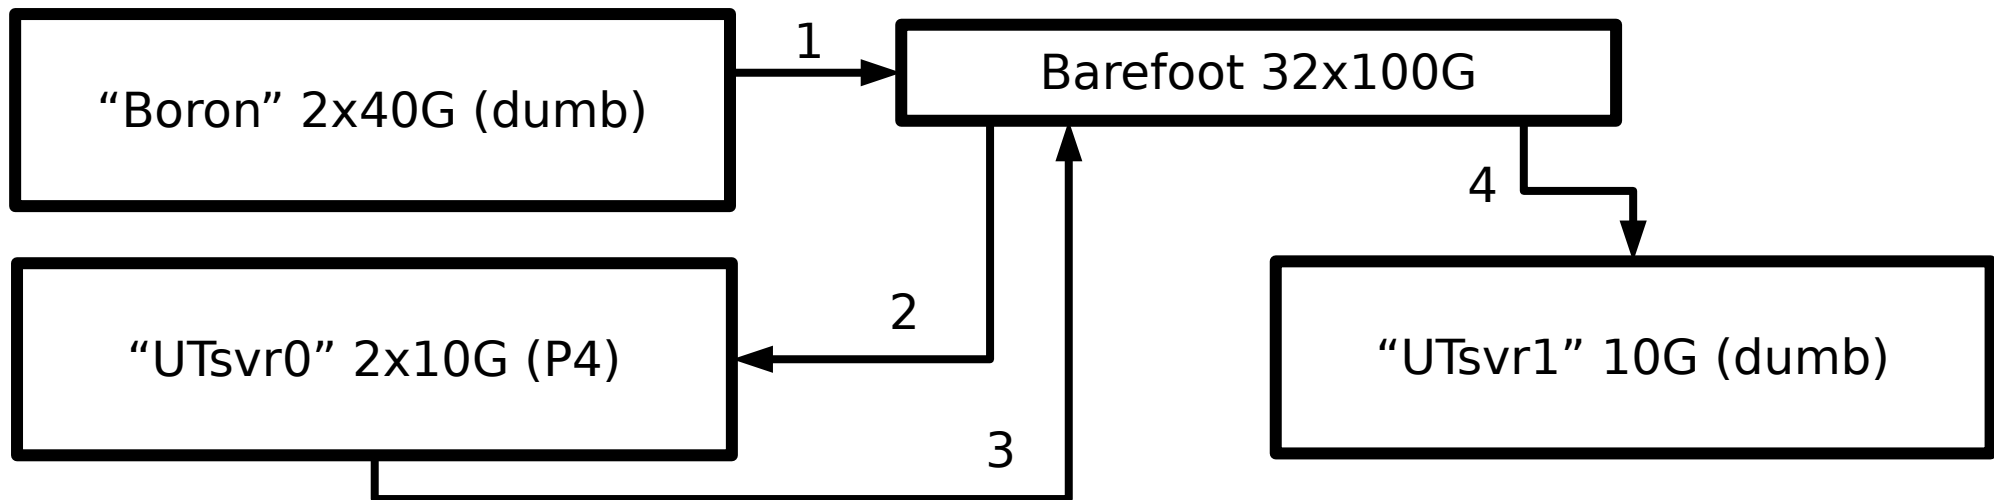


# Scenario #2: Incorporating INT

**Exporting statistics obtained via In-band Network Telemetry (INT):**

**1. P4 devices add metrics to in-band packets**

**2. The P4 exporter exports these like any other feature**





# P4 and BGP

- **To which extend can P4 be used to perform or enhance BGP operations?**
- **Can we analyse anycast performance by deploying P4 hardware in strategic places in the network?**

**Can we do actual BGP announcements from a P4 device?**

**In-band Network Telemetry to find bottlenecks?**

# Challenges, pitfalls, unknowns

**How much can we do in pure P4? To what extent do we need to fallback on (micro-)C?**

**p4\_14 vs p4\_16: wait for p4\_16 everywhere, or purposefully create a heterogeneous setup to find limitations?**