Hochschule
Bonn-Rhein-Sieg
University of Applied Sciences

RFC 9611-
PERFORMANCE TEST

IKEv2 Support for Per-Resource Child SAs

17.07.2025 • Hannes Tschofenig, Kai Jansen

AGENDA

1. The setup
2. The tests
3. The graphs
4. The observations
5. The conclusions

THE SETUP

- Basis is a bachelors thesis by Maximilian Wagner done on a Proxmox Cluster
- VMs were used
- Baremetal Hardware Intel(R) Xeon(R) w3-2423 with Mellanox „MT2910 Family“ NICs 100GB/s (thanks to Paul Wouters)
- Fedora 42 Linux
- Strongswan latest compile (Jul-25)
- Iperf3 as traffic generator
- Mpstat for cpu monitoring
- Swanctl –list-sas for SA-monitoring

THE TESTS

- 10 x 600s testruns, 1s interval
- Parallel scripts/commands
  - iperf3 -c 10.0.1.2 -P6 -t600 -fM -i1 --timestamp=%Y%m%d%H%M%S –logfile [...]
  - mpstat -P 0-5 1 1
    - cpu usage value = 100 – idle (one report per second formatted through ‚awk‘)
  - swanctl –list-sas
- All with added timestamps

THE GRAPHS

Hochschule
Bonn-Rhein-Sieg
University of Applied Sciences

- No Encryption a.k.a. strongswan service disabled => Expected results around 100Gbit/s

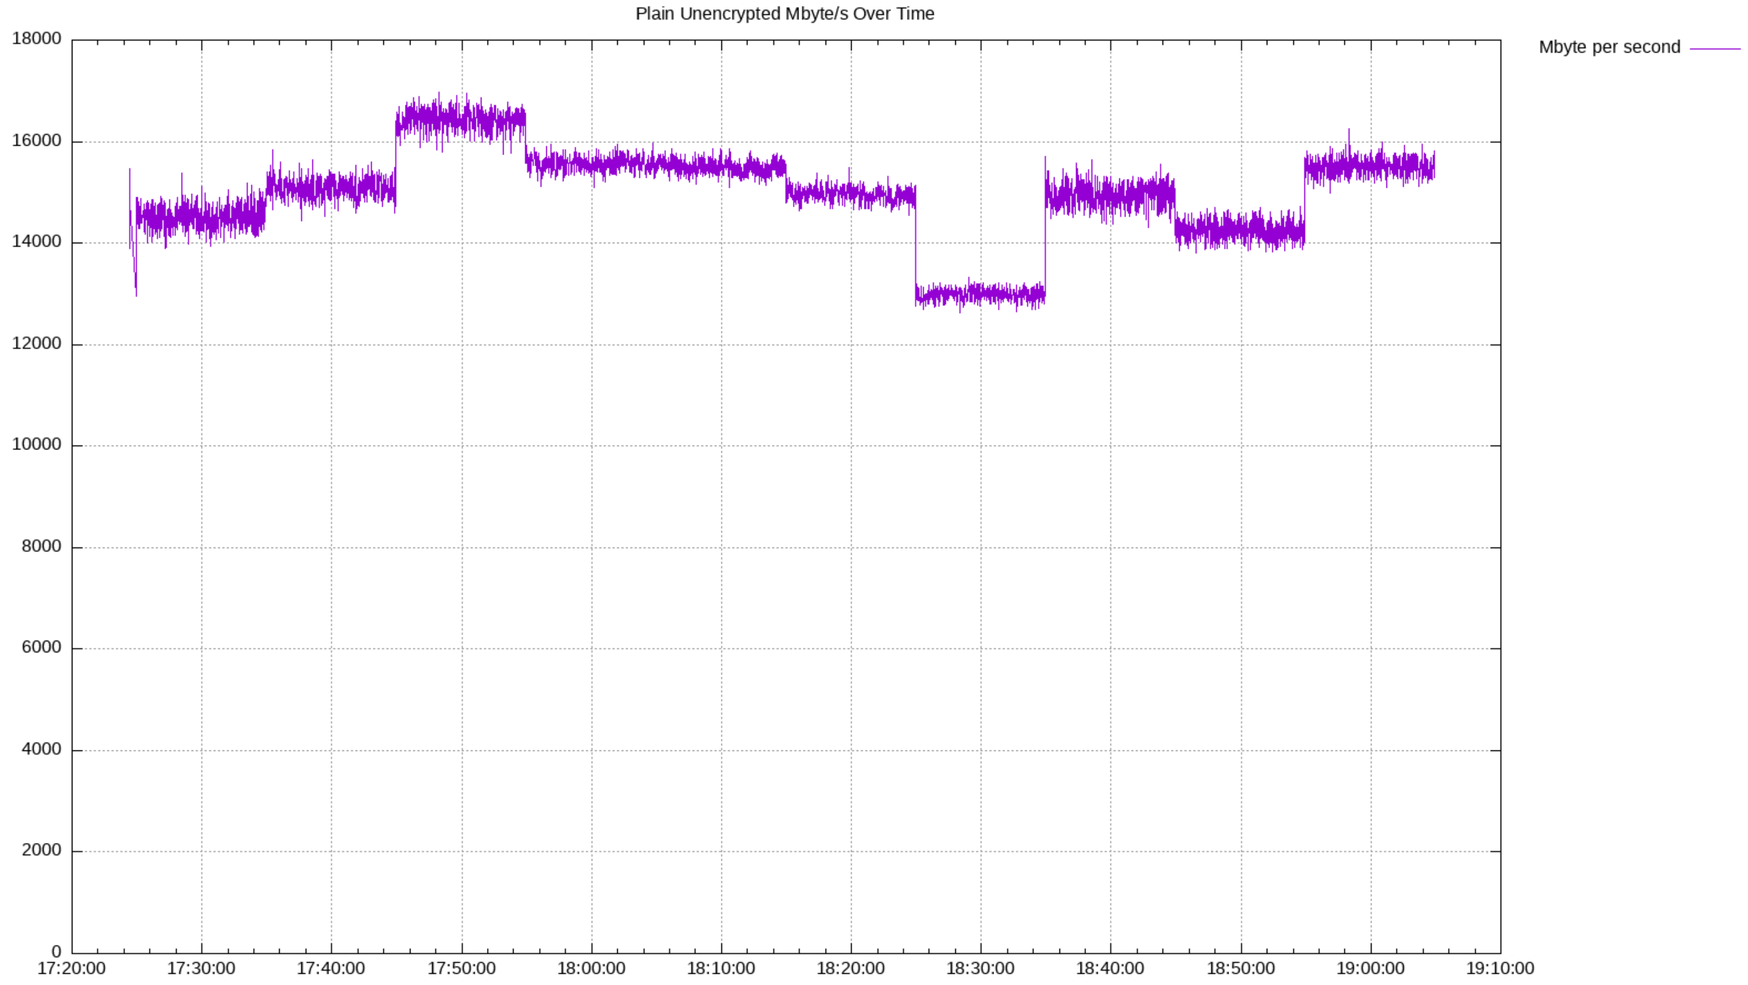

THE GRAPHS

Hochschule
Bonn-Rhein-Sieg
University of Applied Sciences

- Single SA => Consistently around 1.8 Gbit/s

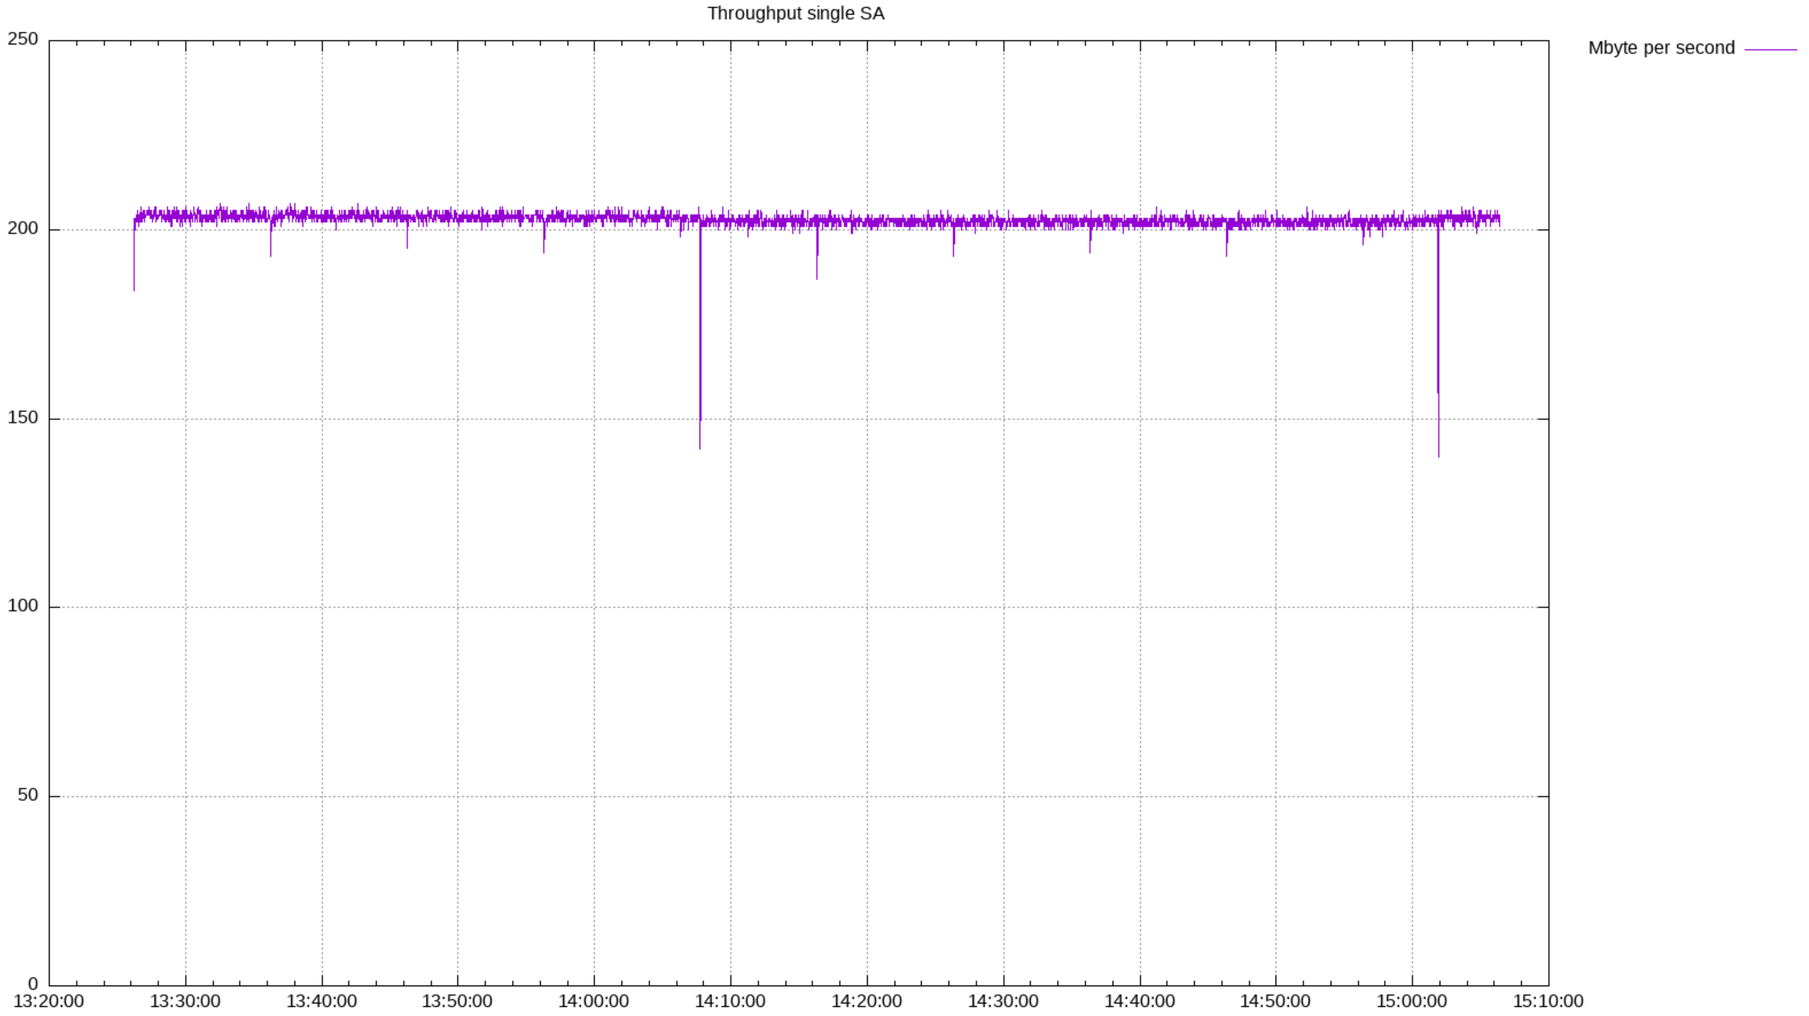

THE GRAPHS

- perCPU SA => Erratic performance between runs of iperf3 1.8Gbit/s and 6 Gbit/s

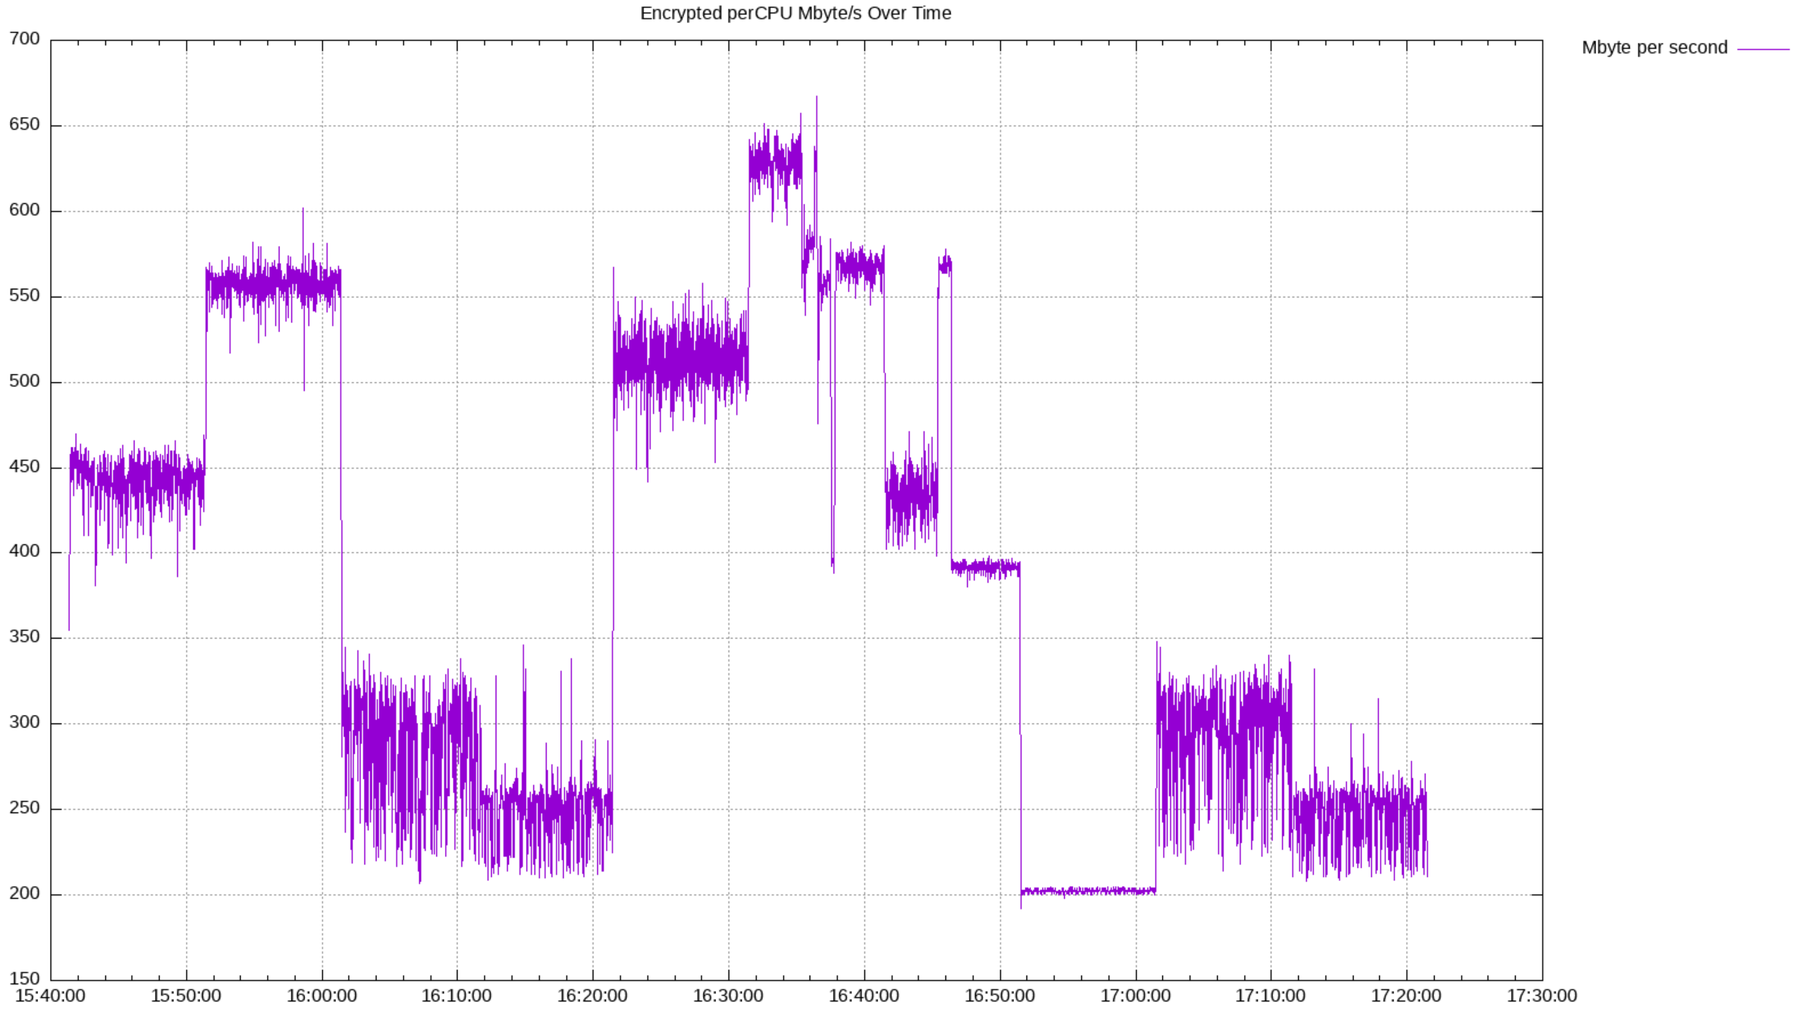

THE GRAPHS

Hochschule
Bonn-Rhein-Sieg
University of Applied Sciences

- Further Observations: SA Usage/Activity

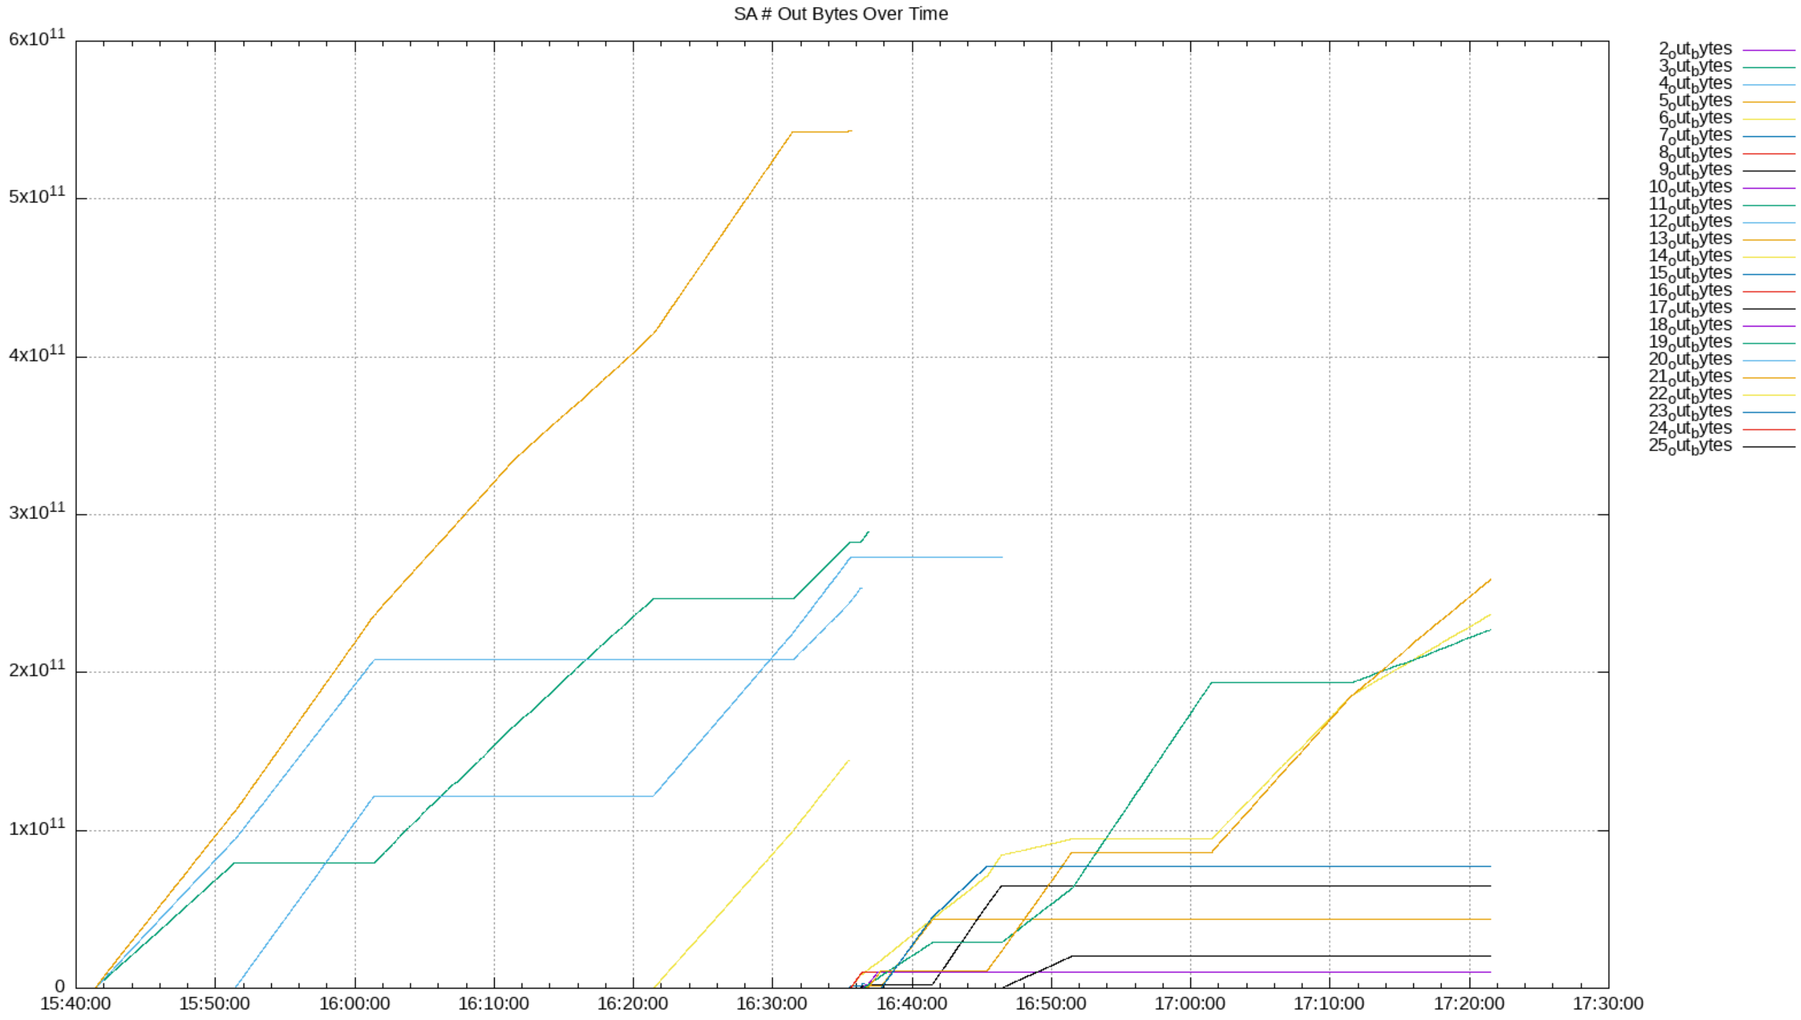

THE OBSERVATIONS

- Re-run of the same measurement → very different results
- New SAs are created over time across all consecutive runs
- SAs stop being used at some time
- Difference in available CPUs 6 (sender) → 12 (receiver) results in additional SAs being created on the sender side. BUT only on CPU 0. No other CPUs are used for the additional SAs
- Some SAs are solely being used for send OR receive, some for both simultaneously

THE CONCLUSIONS

- Traffic generation with iperf3 might be flawed.
  - Why such different results between re-runs of the same measurement?
- IPSec is hard
- „Wer misst, misst Mist“ – A naive approach to measurements is dangerous for performance conclusions
- Best practices / OOBE considerations for implementation?
- Tips for future testing? *cough*Hackathon*cough*

Hochschule
Bonn-Rhein-Sieg
University of Applied Sciences

THE THANK YOU!