Multi-core IPsec tunnels

Daniel Xu, Vlad Dumitrescu, Antony Antony

Additional credits

- Kernel work ~3-4 years ago @Netdev 0x13 in Prague.
- Many thanks to Steffen Klassert for xfrm patches
- Tobias Brunner : strongSwan support
- Paul Wouters : IETF standardization, and testing using libreswan
- Sowmini Varadhan : initial use case
- Benedict Wong and Tuomo Soini : hacking and testing
- Jonathan Lemon : ENA driver XDP multi-buffer support

Background

- IPsec tunnels have well known scalability limitations
  - Crypto state, counters, and sequence numbers cannot be efficiently shared across cores
- Link speeds vastly outpacing single tunnel performance improvements
- Would like to take advantage of modern multi-core systems

Typical topology

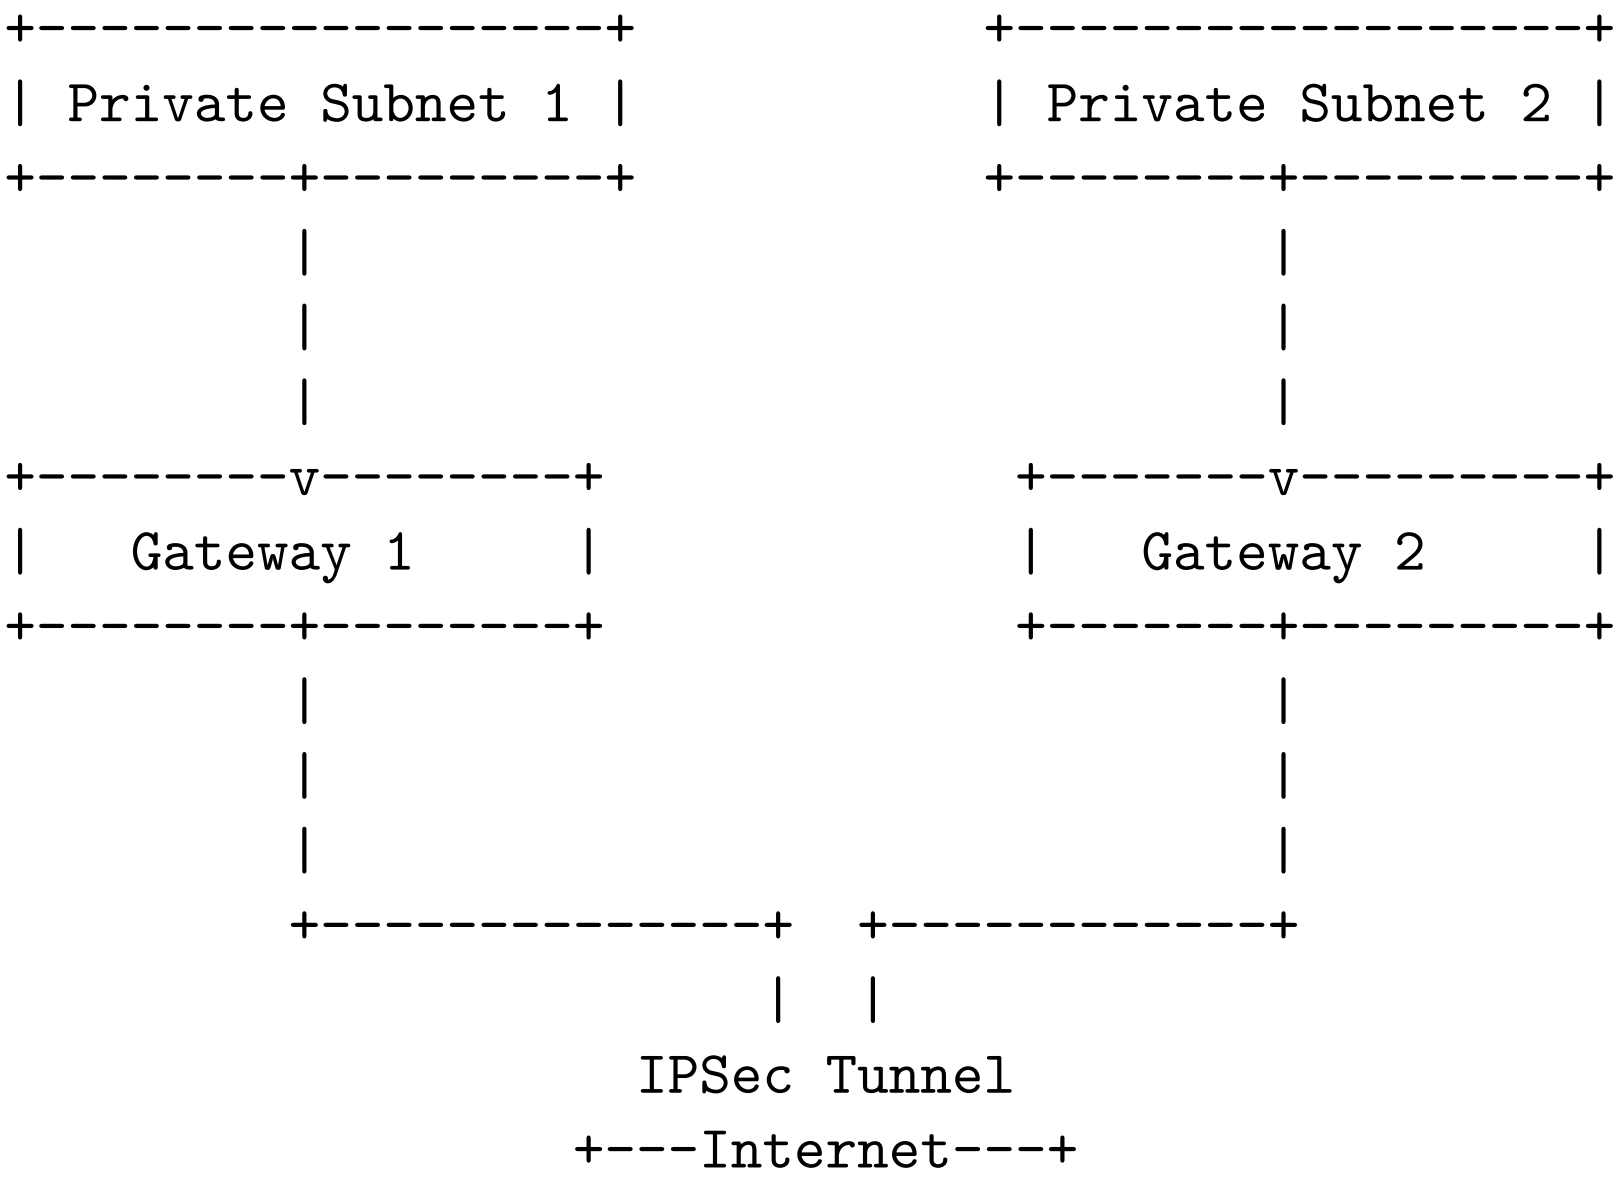

Aviatrix topology

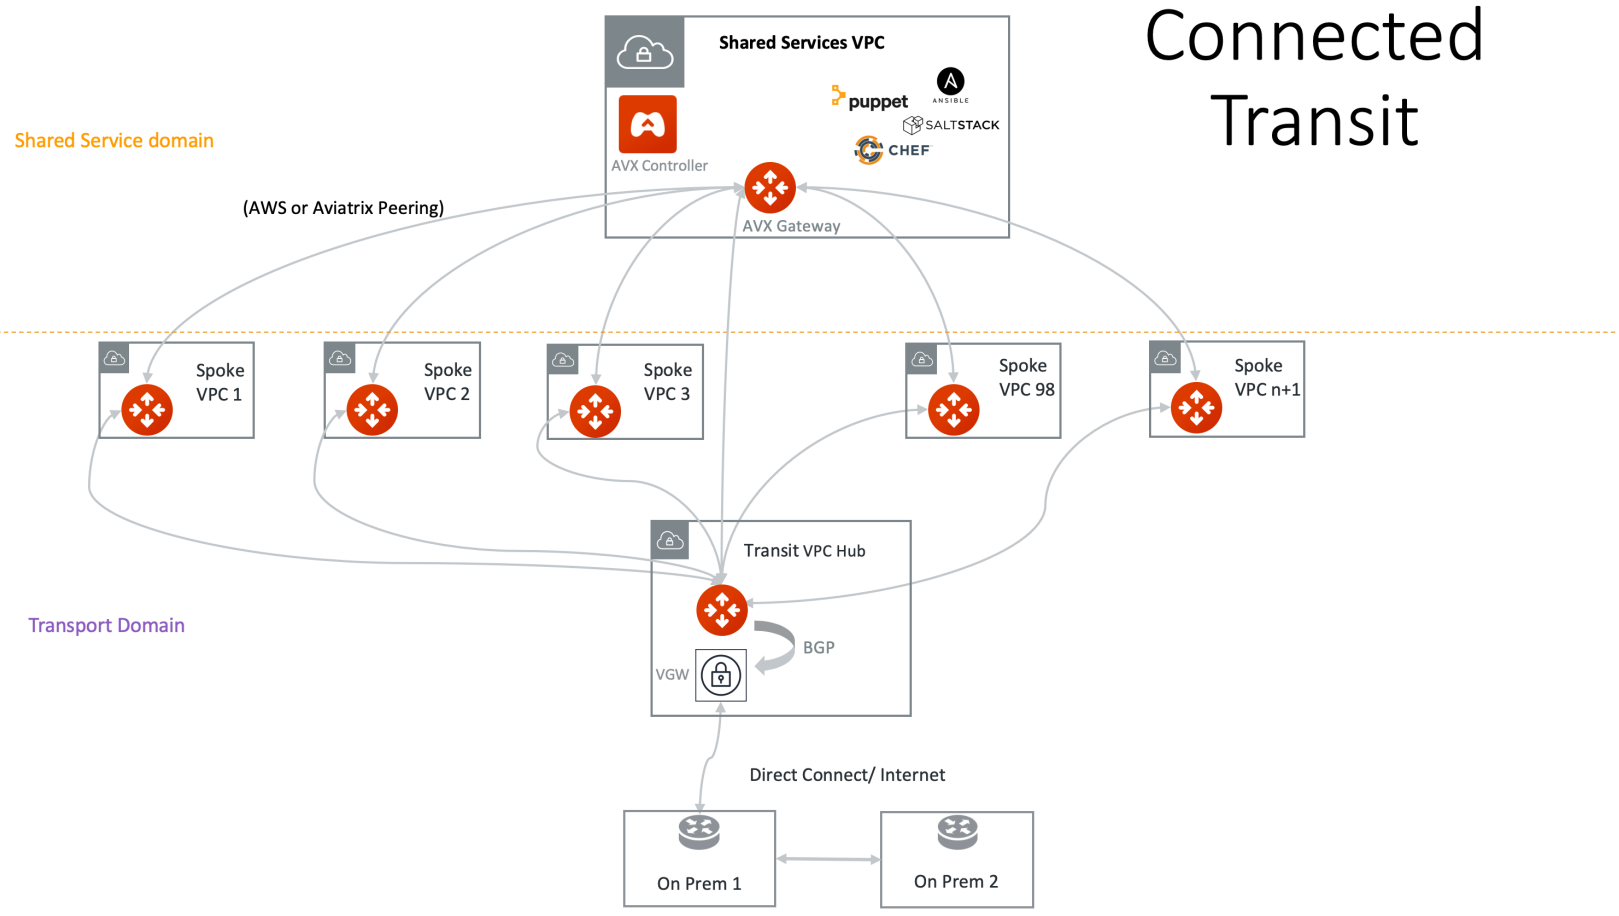

pCPU IPsec tunnel design

- Negotiate a pair of SAs for each CPU
  - On demand and sender driven
- On TX, the pCPU SA is chosen based on current CPU
- On RX, expect a given pCPU SPI to always land on same CPU
  - Hardware RSS or software RSS (XDP_REDIRECT)
- If RX and TX constraints are met: lockless operation and linear scaling

End to end sequence

dxuuu.xyz/r/ipsec-pcpu.png

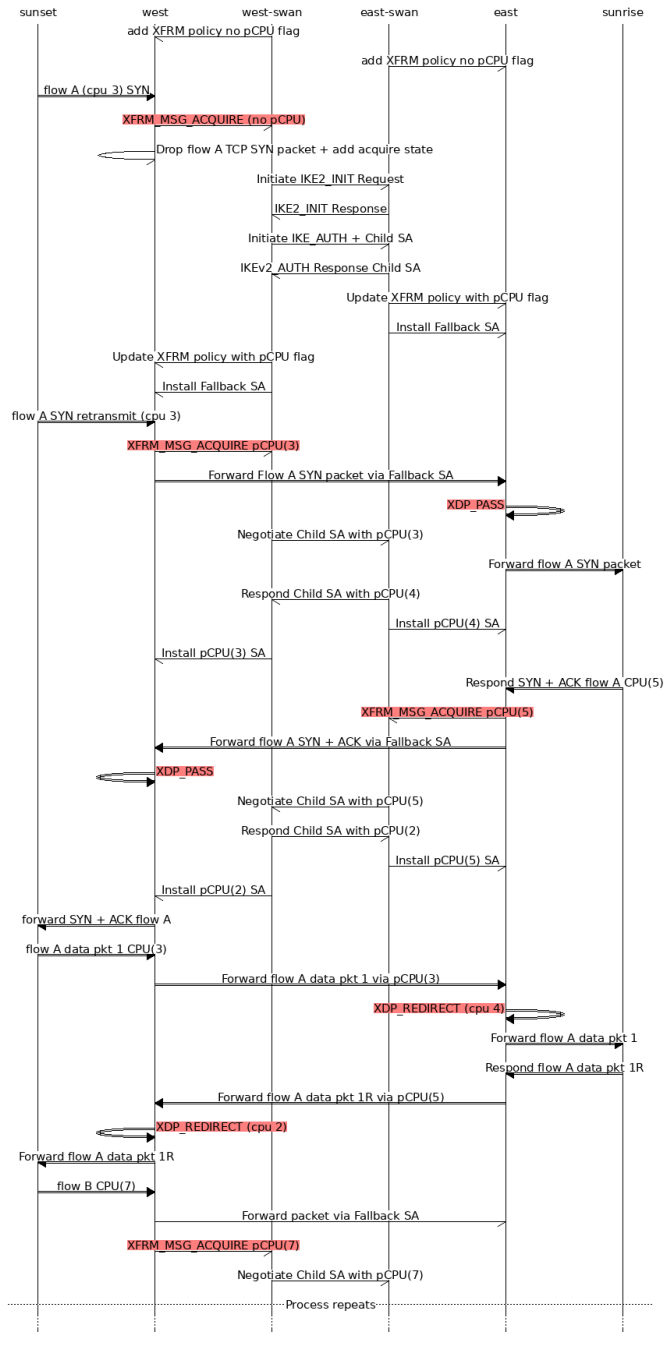

Experimental setup

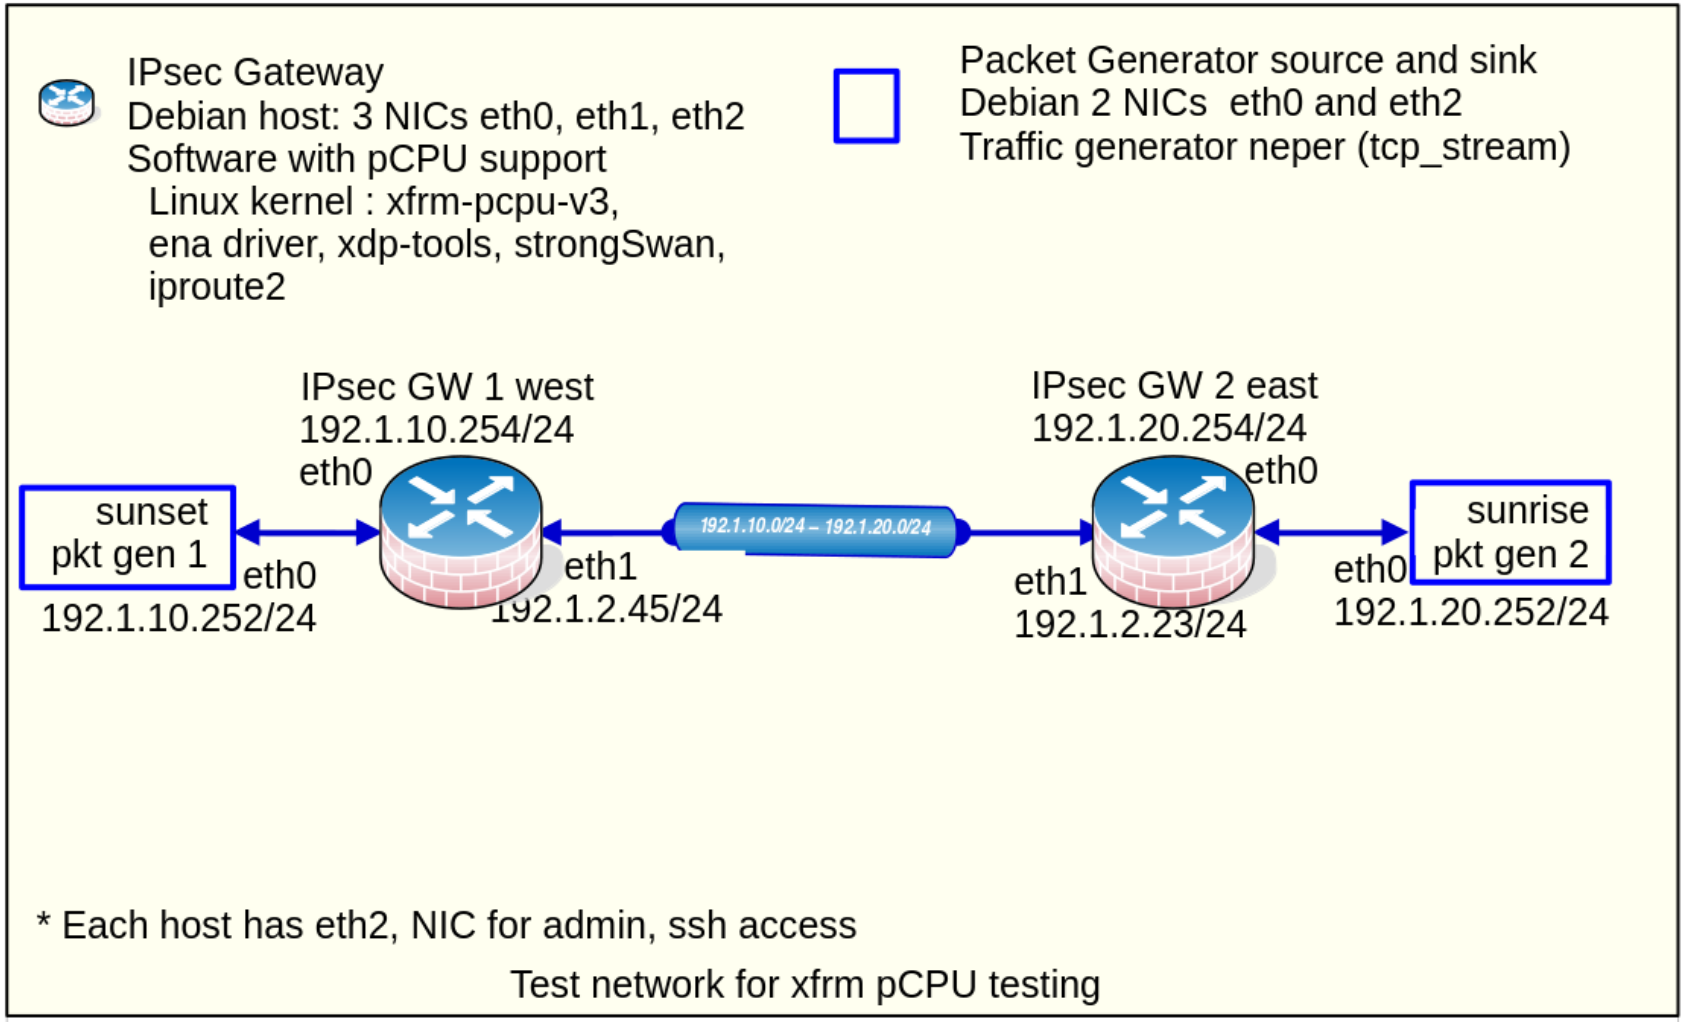

Experimental setup details

- All hosts are EC2 `c6in.16xlarge`
  - 32 physical cores (hyperthreading disabled)
  - 100 Gbps instance bandwidth
  - 10 Gbps single flow limit
  - 16 combined rx/tx queues
- xfrm pcpu patches applied to 6.5.6 Debian sid kernel
- ENA patches applied to prevent XDP queue halving and for jumbo frames
- GRO disabled on all dataplane interfaces (more on this later)
- XDP_REDIRECT used for steering
  - SPI for rx
  - sport/dport for plaintext tx
- `neper` (`tcp_stream`) used for traffic
- UDP encap used to overcome single flow limit
  - AWS does not differentiate ESP flows based on SPI :(

Results (1/3)

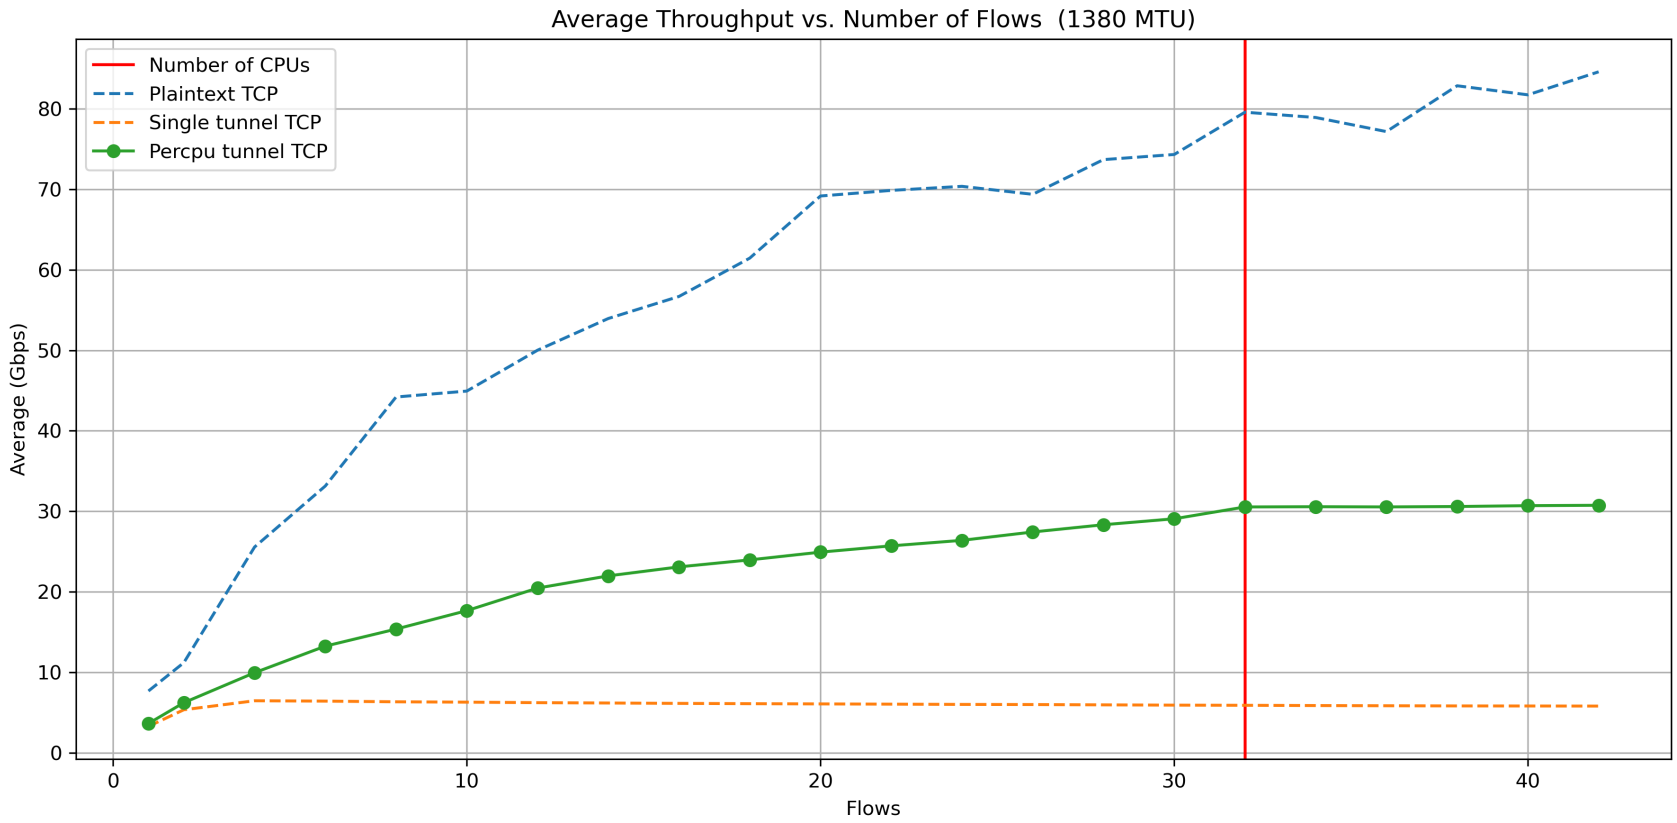
Average Throughput vs. Number of Flows (1380 MTU)
Number of CPUs
Plaintext TCP
Single tunnel TCP
Percpu tunnel TCP
Average (Gbps)
Flows
0
10
20
30
40
0
10
20
30
40
50
60
70
80

Results (2/3)

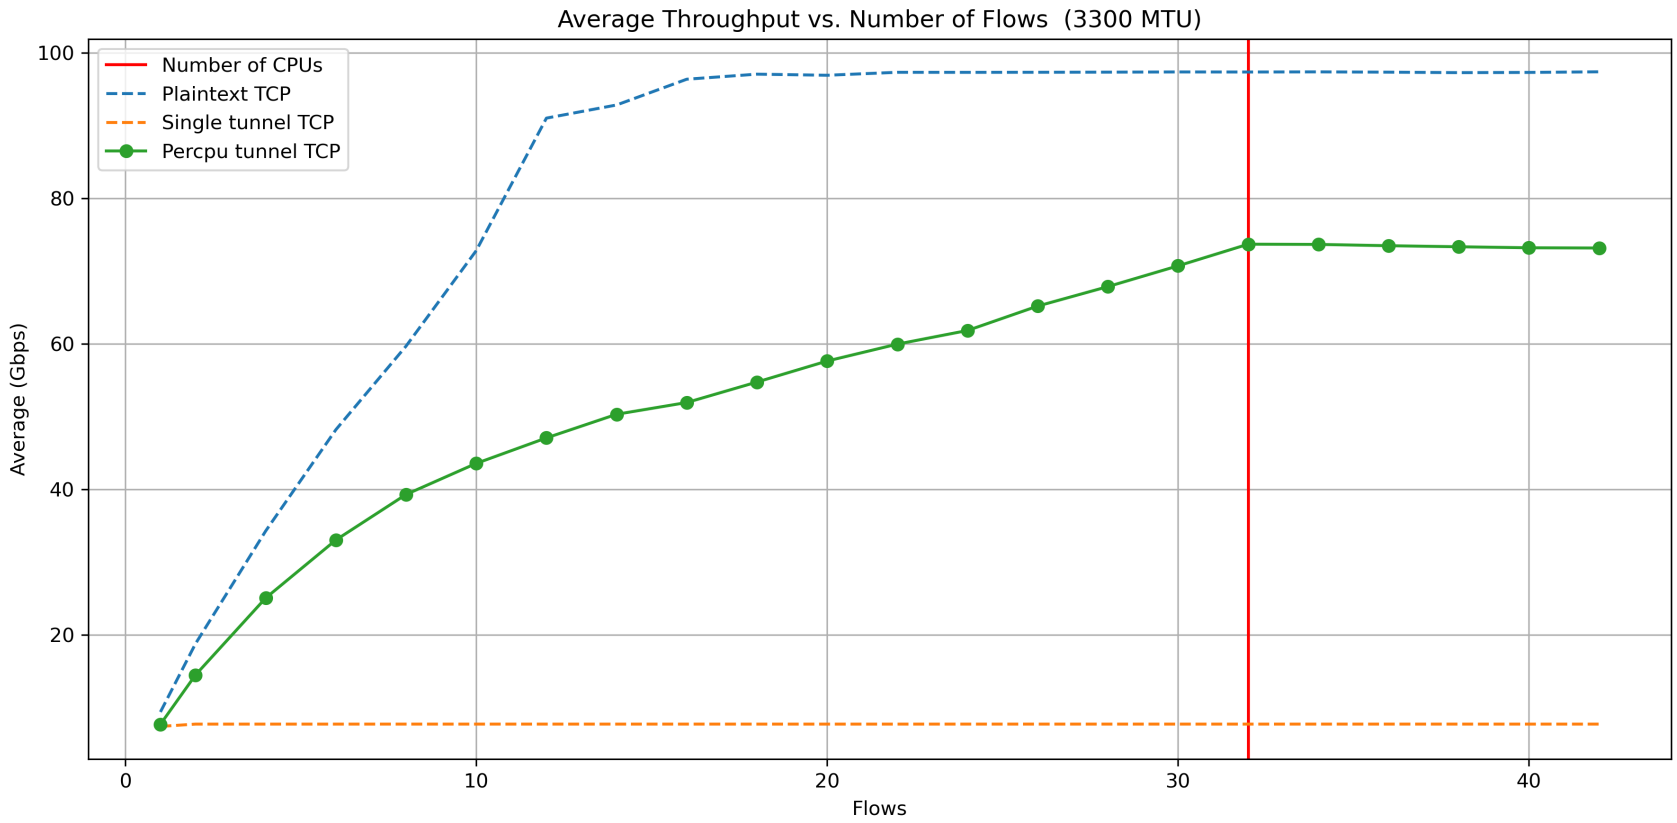
Average Throughput vs. Number of Flows (3300 MTU)
Number of CPUs
Plaintext TCP
Single tunnel TCP
Percpu tunnel TCP
Average (Gbps)
Flows
100
80
60
40
20
0
10
20
30
40

Results (3/3)

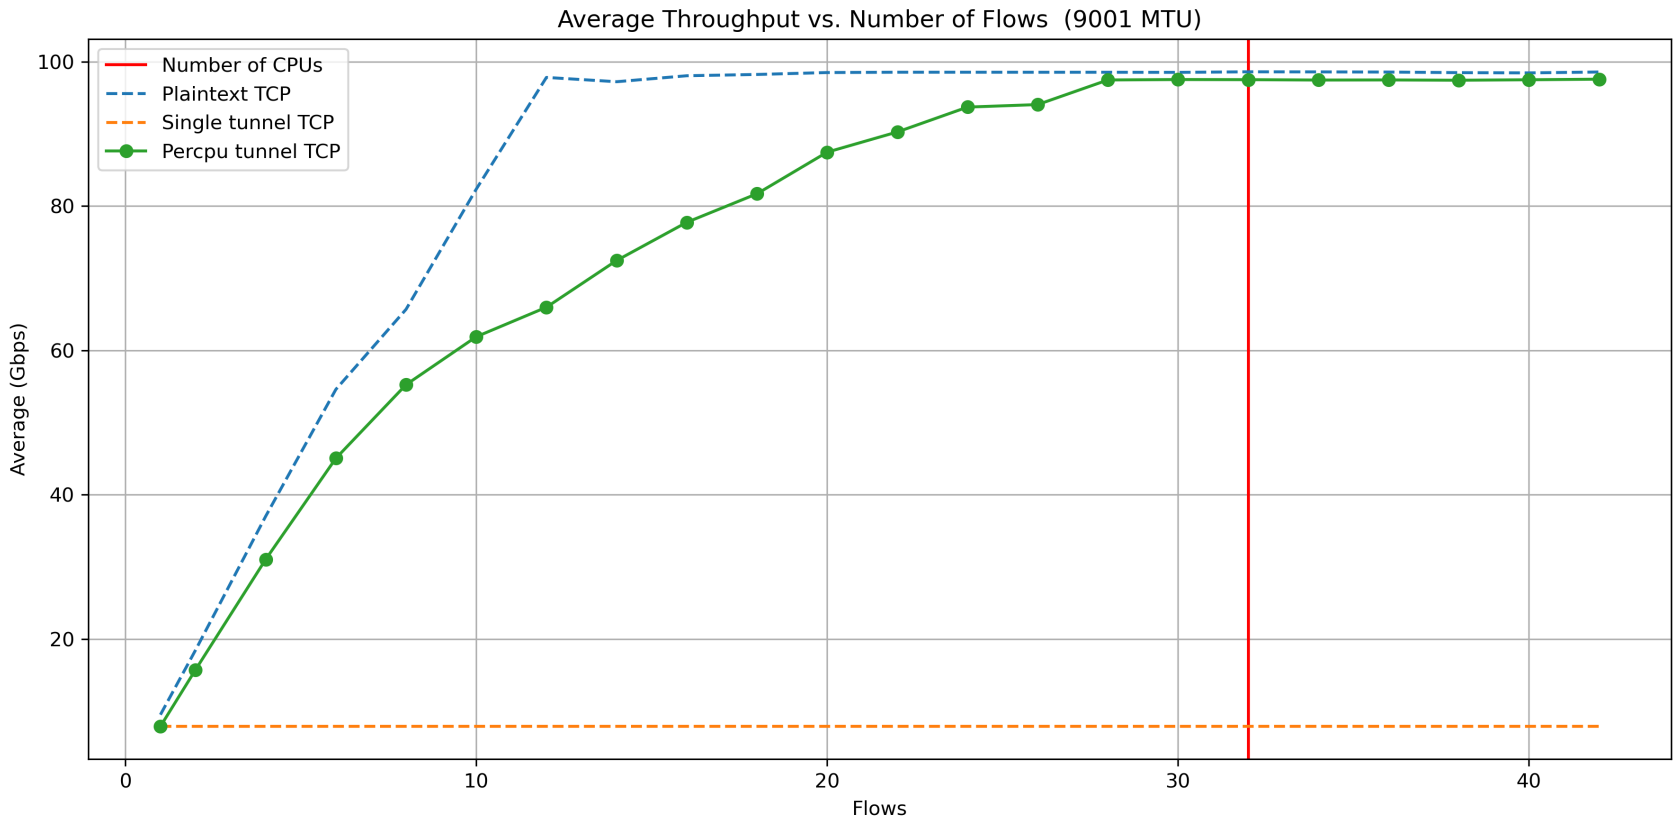
Average Throughput vs. Number of Flows (9001 MTU)
Number of CPUs
Plaintext TCP
Single tunnel TCP
Percpu tunnel TCP
Average (Gbps)
Flows
100
80
60
40
20
0
10
20
30
40

Near-term improvements

- XDP cpumap GRO support
  - To help batch up plaintext flows to hand to xfrm
  - Big expected win here
  - Patches already exist!
- xfrm pcpu tx contention
  - Unexpected appearanes in cpu profile:
    - `xfrm_resolve_and_create_bundle()`
    - `xfrm_state_find()`

It takes a village

Changes in:

- kernel xfrm
- kernel bpf
- Amazon ENA driver
- strongSwan
- xdp-tools
- iproute2
- IETF