Zero-setup Intermediate-rate Communication Guarantees in a Global Internet

Marc Wyss and Adrian Perrig

ETH zürich

Communication availability

Objective: availability guarantees for short-lived intermediate-rate communication.
(DNS communication, accessing websites, …)

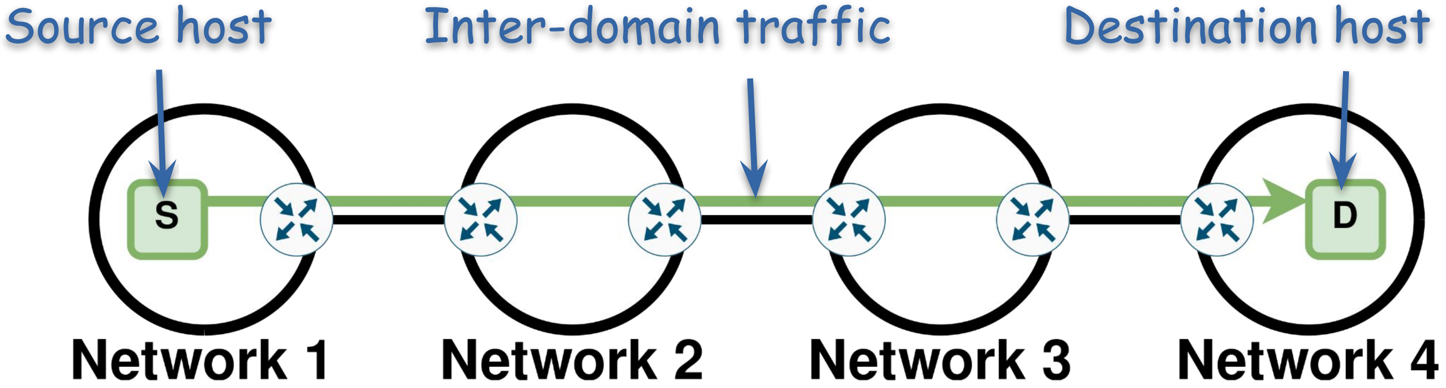

Example: *ETH Zurich* *SWITCH* *GÉANT* *OVGU*

Communication availability

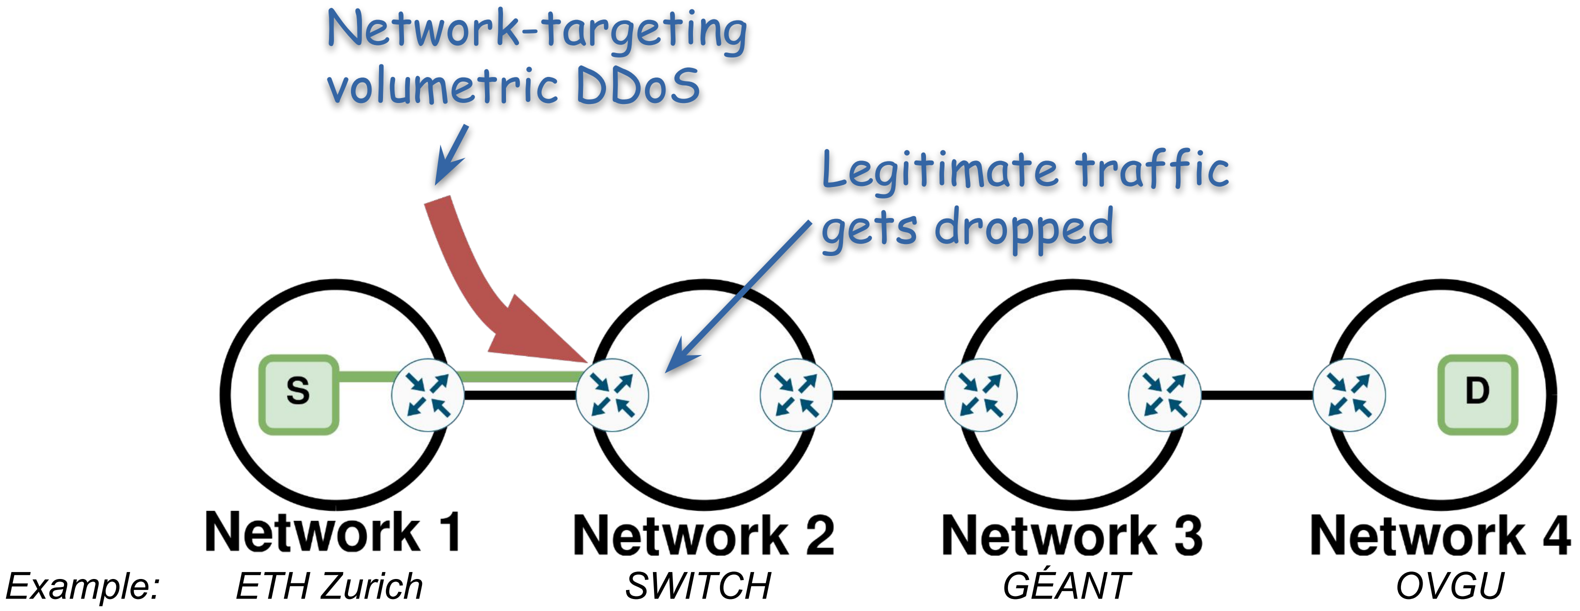

Existing solutions

- Leased lines ← $$$
- Private backbone ← $$$
- Overprovisioning ← Problem: ever higher DDoS volumes
- Rerouting & scrubbing ← Adds latency
- Per-flow fairness ← Vulnerable to spoofing
- Pushback-like ← Reactive: delay, misclassification
- Bandwidth reservations ← Substantial setup overhead
- …

Protecting short-lived intermediate-rate communication in the global Internet is challenging!

Core insights

1. ***Pre-allocate** low rate at routers to avoid setup overhead.*

Enough for intermediate-rate traffic aggregate

Allocate 1 Gbps for traffic from network 1, isolate traffic from other sources

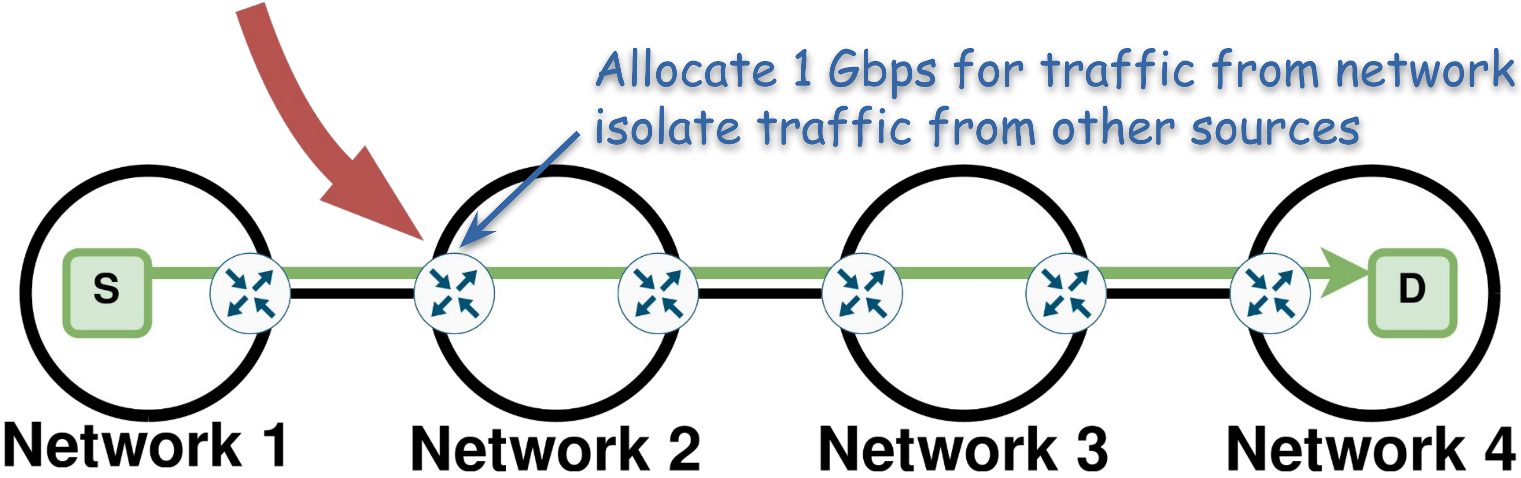

Core insights

2. *Network bandwidth isolation requires source authentication.*

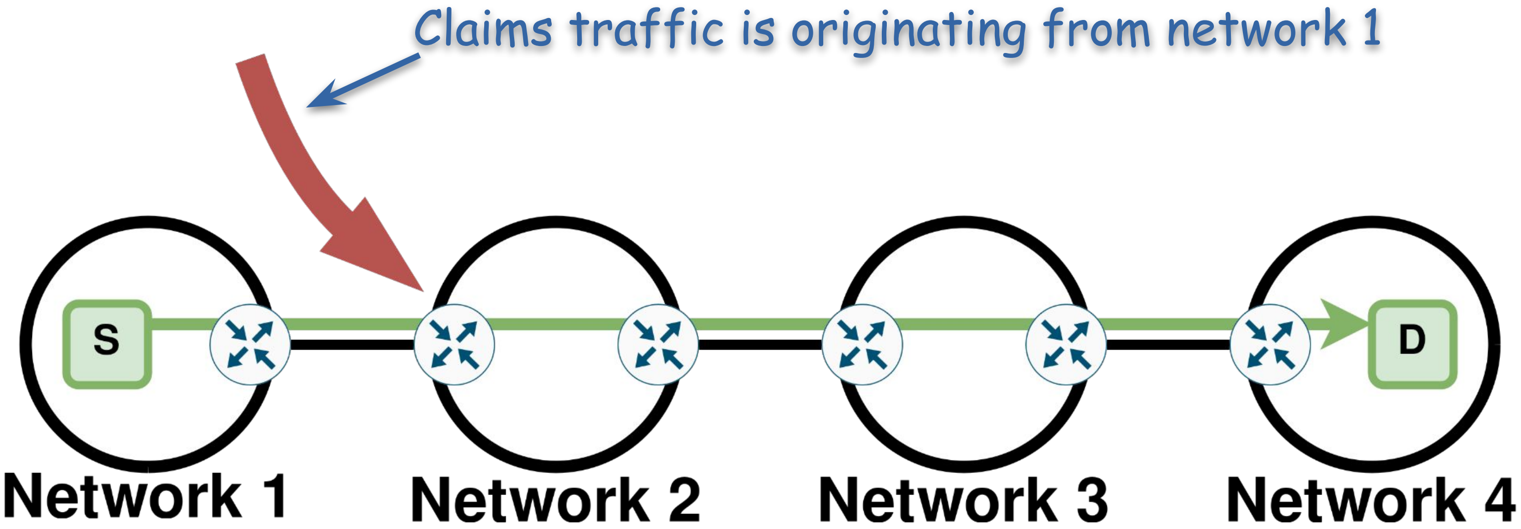
Claims traffic is originating from network 1

Core insights

2. *Network bandwidth isolation requires source authentication.*

EPIC

“Every Packet Is Checked”
[USENIX Security ‘20]

- Every router can **verify** the authenticity of every packet’s **length and origin**.
- Requires path transparency: end hosts learn the identities of on-path networks.

Core insights

3. *Secure routing is essential for communication availability.*

SCION™

- **Prevents hijacking** attacks by design.
- Provides **path transparency.**
- Isolation Domains (ISDs): Trust-based **groupings of networks**.

Z-Lane

"Z" as in "zero-setup"

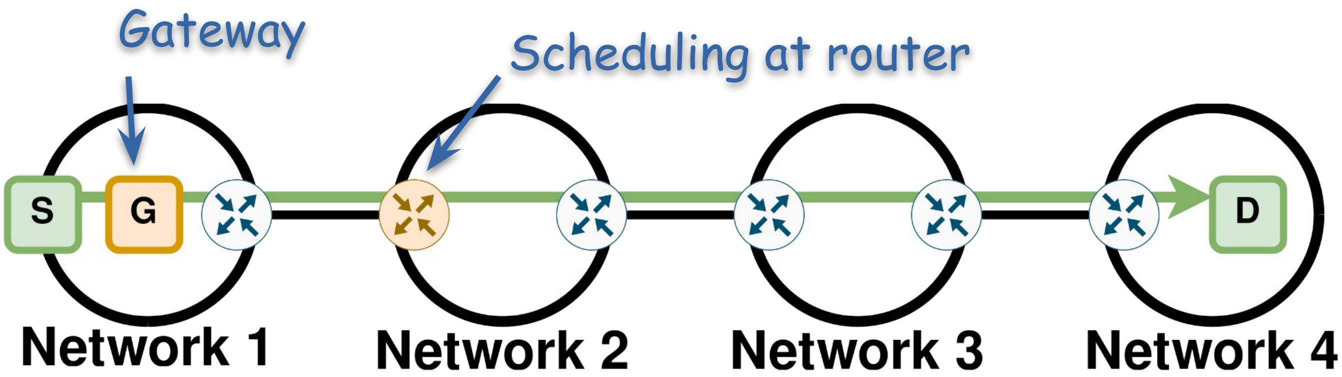

Z-Lane: router

Router implements **bandwidth isolation** of traffic from different networks.

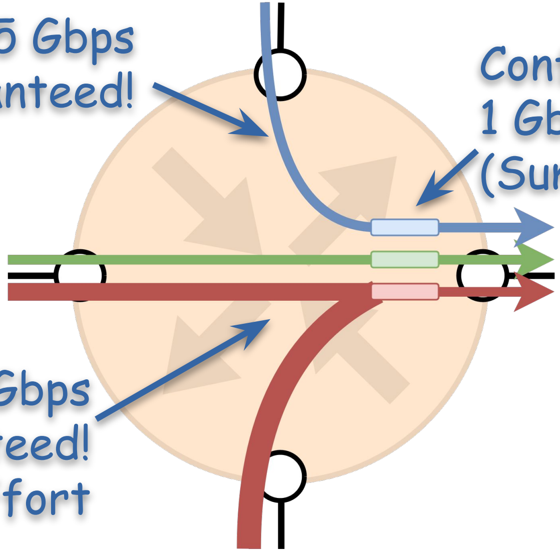

0.5 Gbps
→ guaranteed!

Configure
1 Gbps for each source network
(Sum of rates < link capacity)

3 Gbps
→ 1 Gbps guaranteed!
→ 2 Gbps best-effort

Z-Lane: router

Trivial solution for bandwidth isolation: per-network queues.

- Does not scale to size of Internet.

Z-Lane: one priority queue + per-network token buckets.

- Per token bucket, the memory cost is 20-60 bytes.
- Checking rate compliance requires tens of nanoseconds.

Optimizations:

- Memory-optimized token bucket requiring only 8 bytes of memory.
- Guarantee rates for groups of networks (SCION ISDs)
- Rates for 100’000 networks → ***5.3 kB of memory***

Z-Lane

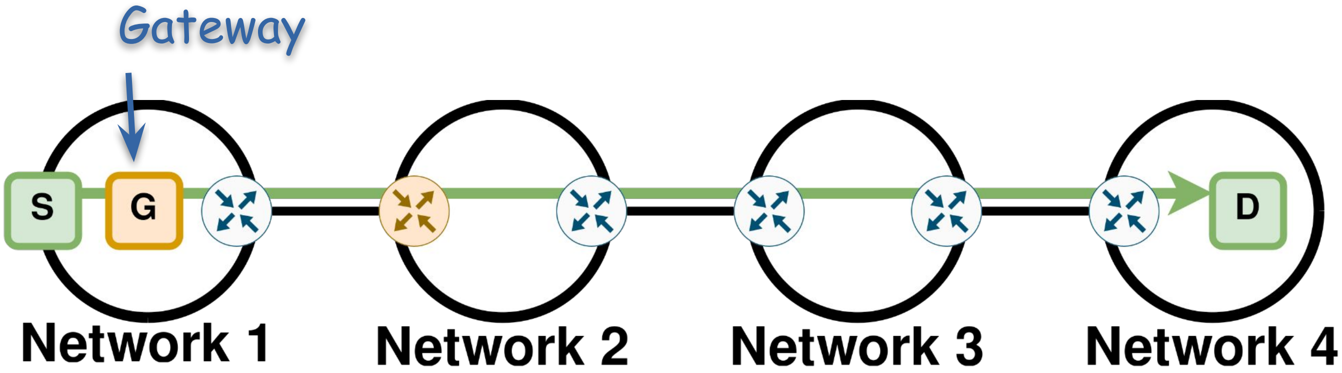

Z-Lane: gateway

… distributes network-level guaranteed rates to end hosts.
… implements bandwidth isolation for end hosts in the same network.

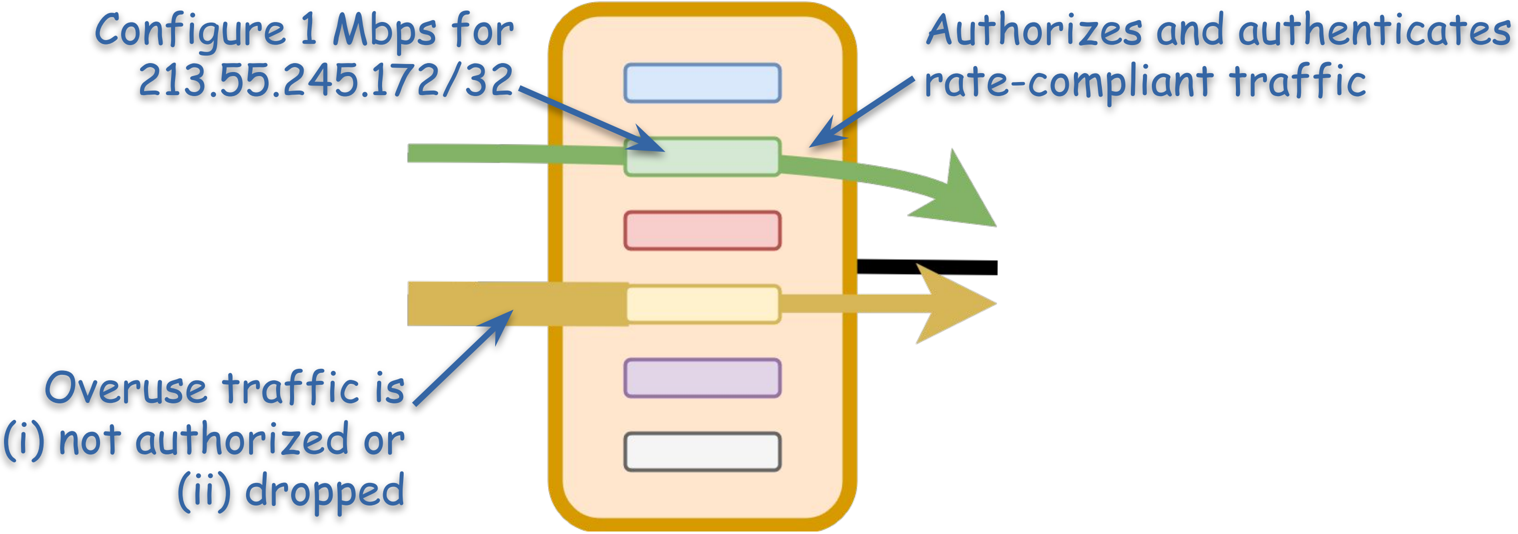

Configure 1 Mbps for 213.55.245.172/32

Authorizes and authenticates rate-compliant traffic

Overuse traffic is (i) not authorized or (ii) dropped

Z-Lane: security

Other systems are often **reactive**: try to detect malicious traffic, then block it.
Z-Lane is **proactive**: provide forwarding guarantees, works immediately.

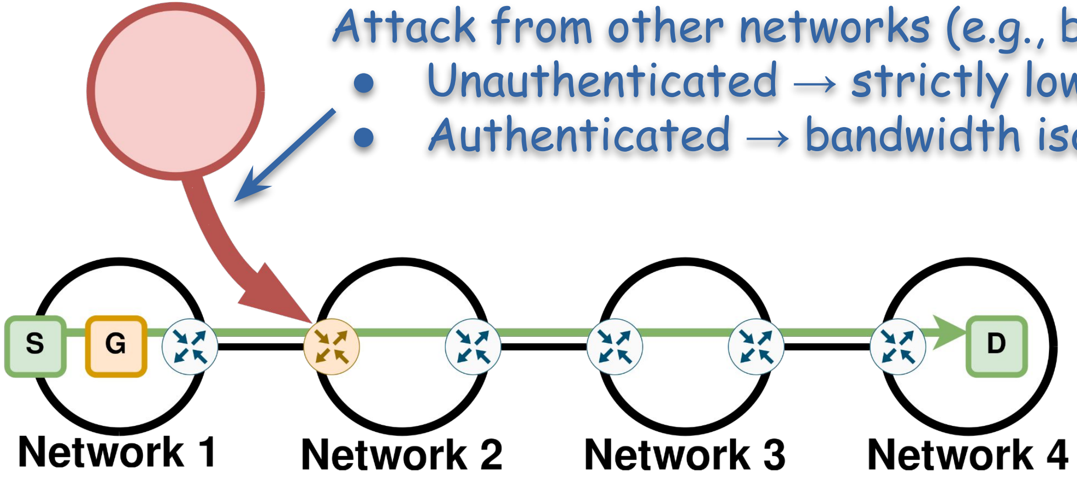

Attack from other networks (e.g., botnet):

- Unauthenticated → strictly lower priority
- Authenticated → bandwidth isolation

Z-Lane: security

Other systems are often **reactive**: try to detect malicious traffic, then block it.
Z-Lane is **proactive**: provide forwarding guarantees, works immediately.

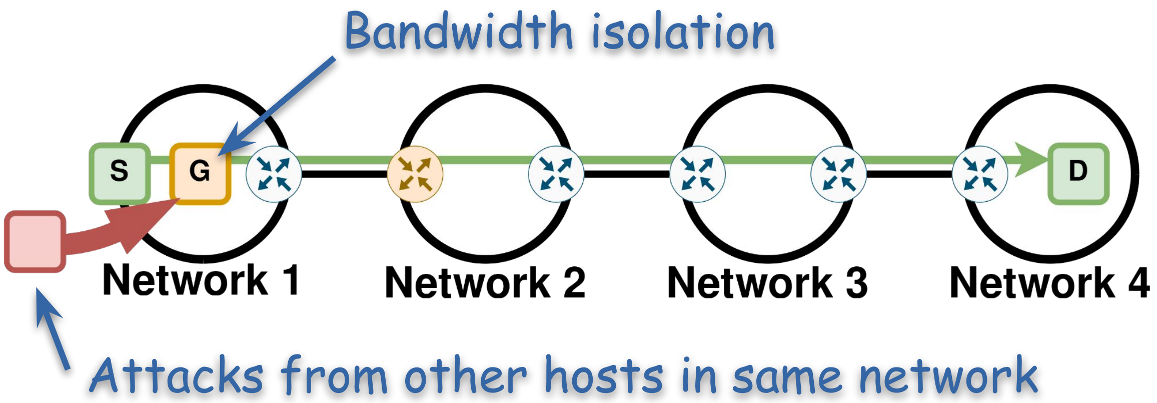

Evaluation: implementation and deployment

SCIONLab

- Global research testbed for SCION
- Seamless integration into SCION
  - No issues during three months testing period
  - Incremental deployment working

High-speed implementation

- DPDK-version of Z-Lane router and gateway
- 160 Gbps forwarding on commodity hardware
- Correct traffic scheduling (bandwidth isolation)

Evaluation: high-speed gateway

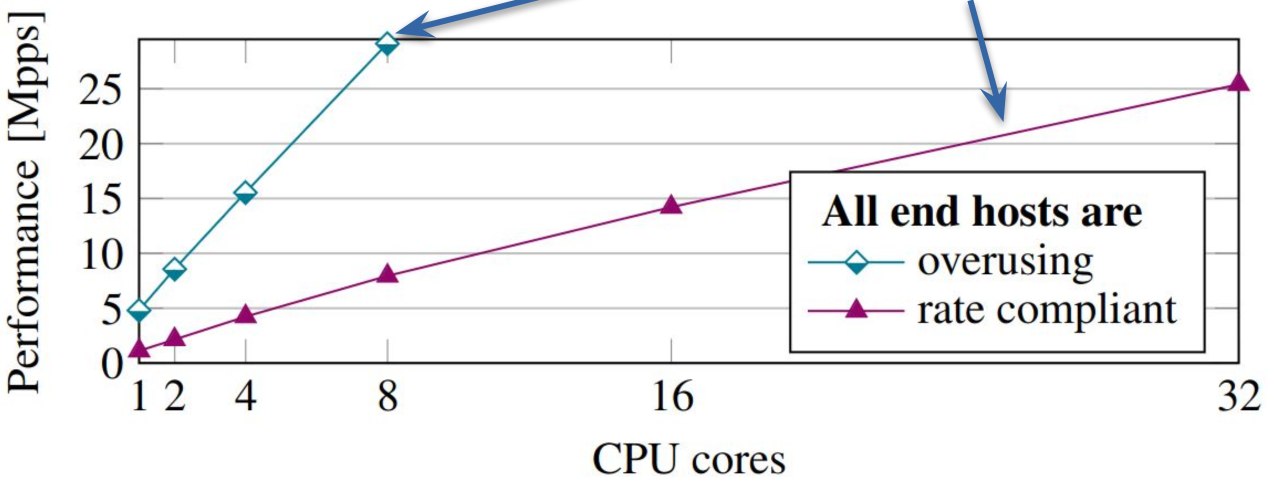

Scales linearly with #cores

… can scale performance further by deploying additional gateways.

Conclusion

- Objective: provide **communication guarantees** to **short-lived intermediate-rate traffic** despite network-targeting **volumetric DDoS** attacks.
- Our proposal: **Z-Lane**
- Can co-exist with bandwidth reservation systems.
  - Protect non-setup critical communication
- Foundation for building exciting new systems!

Thank you!

marc.wyss@inf.ethz.ch

ETH zürich

References

SCION
Laurent Chuat et al. *The Complete Guide to SCION.* Springer International Publishing, 2022.

SCIONLab
The SCIONLab research network. https://www.scionlab.org, 2024.

Pushback
Ratul Mahajan et al. *Controlling high bandwidth aggregates in the network.* SIGCOMM CCR, 2002.

EPIC
Markus Legner et al. *EPIC: Every packet is checked in the data plane of a path-aware Internet.* USENIX Security, 2020.