High-Performance Communication

Adrian Perrig

Network Security Group, ETH Zürich

4th SIG-NGN Meeting

ETH zürich

SCION

What used to keep me up all night …

THE VERGE TECH SCIENCE CULTURE CARS REVIEWS LONGFORM VIDEO MORE

TECH | CYBERSECURITY | ENTERPRISE

Hackers emptied Ethereum wallets by breaking the basic infrastructure of the internet

21

By Russell Brandom | @russellbrandom | Apr 24, 2018, 1:40pm EDT

SHARE

Hijack of Am
service used
hours unnot

Between 11am until 1
internet, routing you t
unknown actor.

$150K Stolen From MyEtherWallet Users in DNS Server Hijacking

363 | 7

David Floyd
Apr 24, 2018 at 16:35 UTC | Updated Apr 24, 2018 at 16:37 UTC

NEWS

Users of MyEtherWallet, a web app for storing and sending ether and ethereum-based tokens, experienced an attack Tuesday that saw users of the service lose around $152,000 worth of ether.

commercial cloud provider who count major websites such as Twitter.com as customers.

ETH zürich

What's now keeping me up all night?

ETH zürich

SCION

Internet Architecture in 21st Century

- Similar to real-world architecture, Internet Architectural trends change over time, typically not just driven by aesthetics, but also by applications
  - Early networks were circuit-switched for telephony
  - 50 years ago, packet switching started and formed the basis of today's Internet
- Recent architectural trends
  - High security and availability
  - Path-aware networking

“Self-evident” Properties of a Next-Generation Internet Architecture

- Security (broadly defined)
  - High availability even under attack
- Path awareness, path selection
- Multi-path operation
- Formal verification
- Transparency
- Sovereignty

Importance of Path Awareness & Multi-path

- Generally, two paths exist between Europe and Southeast Asia
  - High latency, high bandwidth: Western route through US, ~450ms RTT
  - Low latency, low bandwidth: Eastern route through Suez canal, ~250ms RTT
- BGP is a “money routing protocol”, traffic follows cheapest path, typically highest bandwidth path
- Depending on application, either path is preferred
- With SCION, both paths can be offered!

SCION Architecture Principles

- Near-stateless packet forwarding
- Convergence-free routing
- Path-aware networking
- Multi-path communication
- High security through design and formal verification
- Sovereignty and transparency

Vision: secure, available, and transparent global public Internet

What is SCION?

- Secure inter-domain routing architecture, to replace BGP
- Open Internet platform, open-source
- Highly efficient: enables faster communication than in current Internet
- Highly secure: attacks are either impossible by design or significantly weakened
- Verifiably secure: Security proofs through formal methods
- **Next-generation Internet: path-aware multi-path communication**

Approach for Scalability: Isolation Domain (ISD)

- Isolation Domain (ISD): grouping of ASes
- ISD core: ASes that manage the ISD and provide global connectivity
- Core AS: AS that is part of ISD core

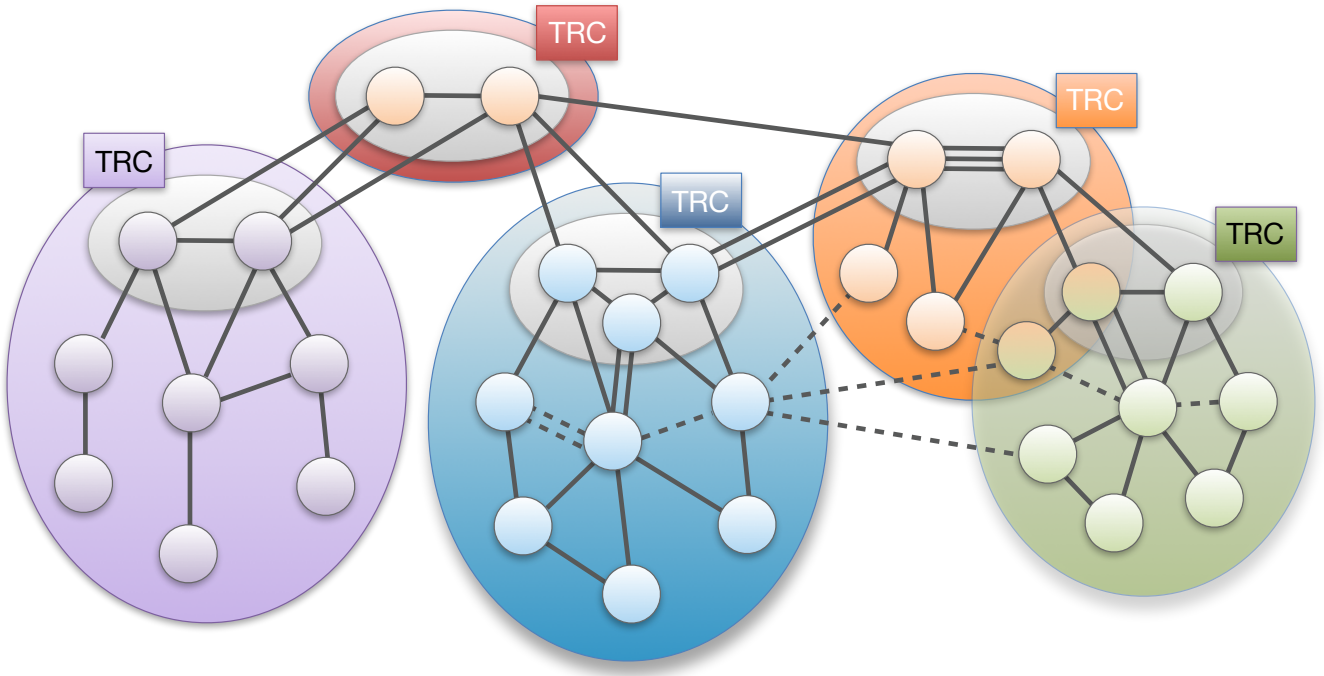

SCION Overview in One Slide

Path-aware Network Architecture

Control Plane - Routing

- Constructs and Disseminates Path Segments

Data Plane - Packet forwarding

- Combine Path Segments to Path
- Packets contain Paths
- Routers forward packets based on Path
  - Simple routers, stateless operation

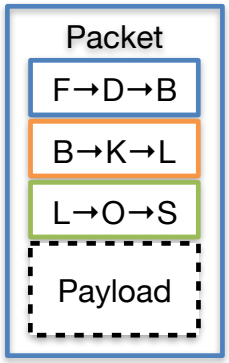

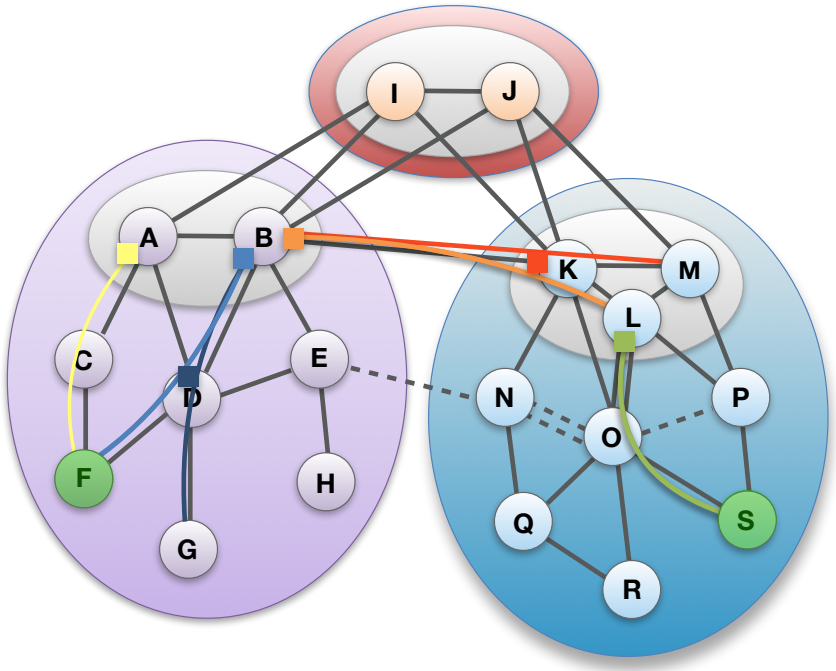

ETH zürich

SCION

Intra-ISD Path Exploration: Beaconing

- Core ASes K, L, M initiate Path-segment Construction Beacons (PCBs), or "beacons"
- PCBs traverse ISD as a flood to reach downstream ASes
- Each AS receives multiple PCBs representing path segments to a core AS

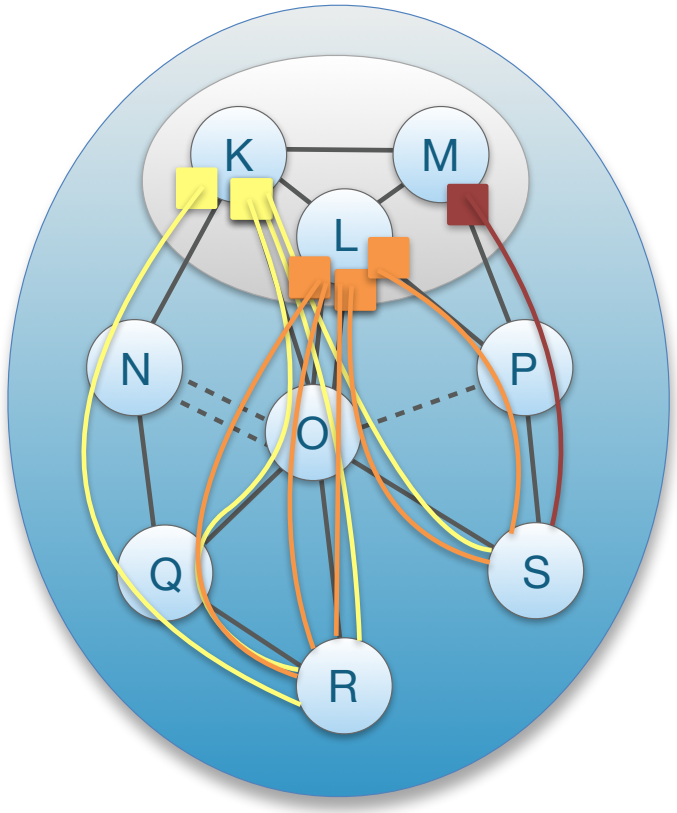

Up-Path Segment Registration

- AS selects path segments to announce as up-path segments for local hosts
- Up-path segments are registered at local path servers

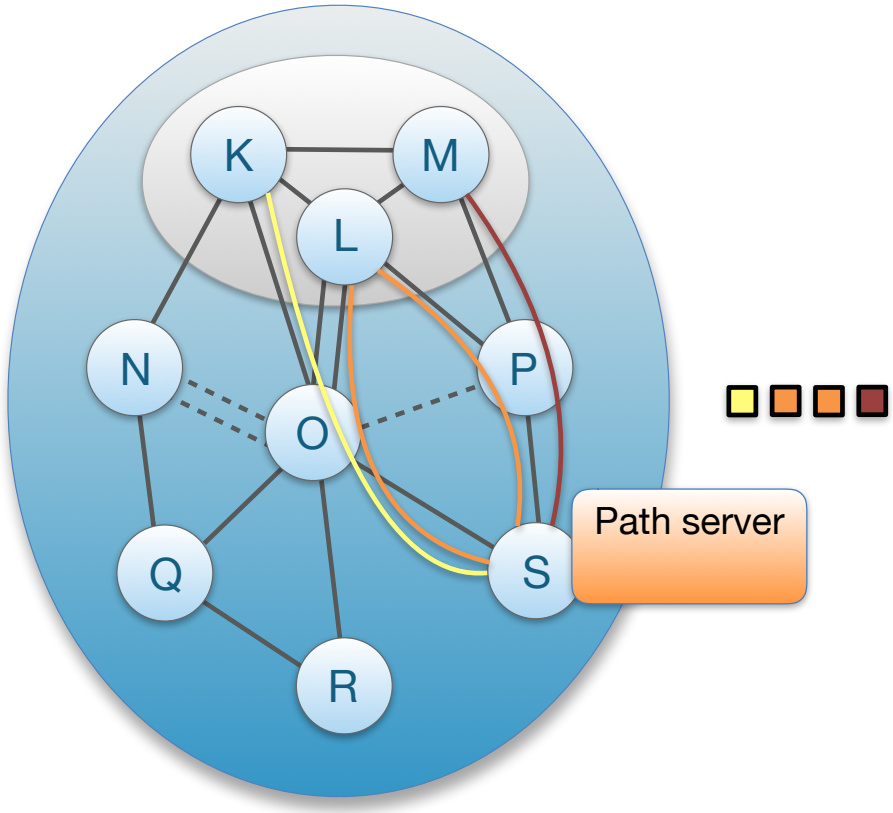

Down-Path Segment Registration

- AS selects path segments to announce as down-path segments for others to use to communicate with AS
- Down-path segments are uploaded to core path server in core AS

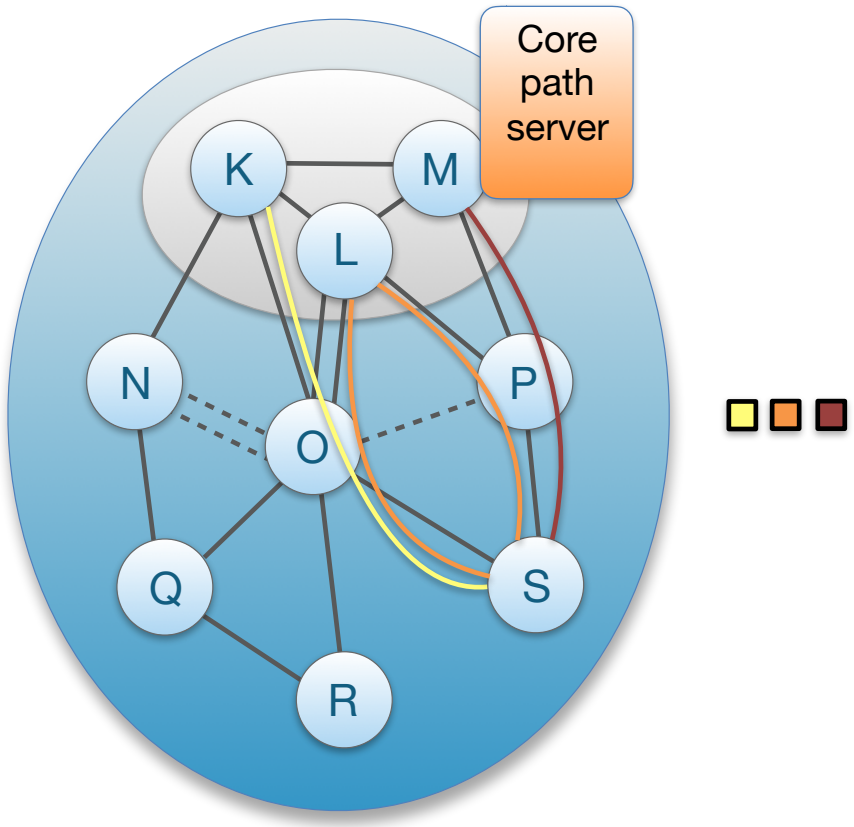

Communication within ISD

- Client obtains path segments
  - Up-path segments to local ISD core ASes (blue)
  - Down-path segments to destination (green)
  - Core-path segments as needed to connect up-path and down-path segments (orange)
- Client combines path segments to obtain end-to-end paths (yellow)

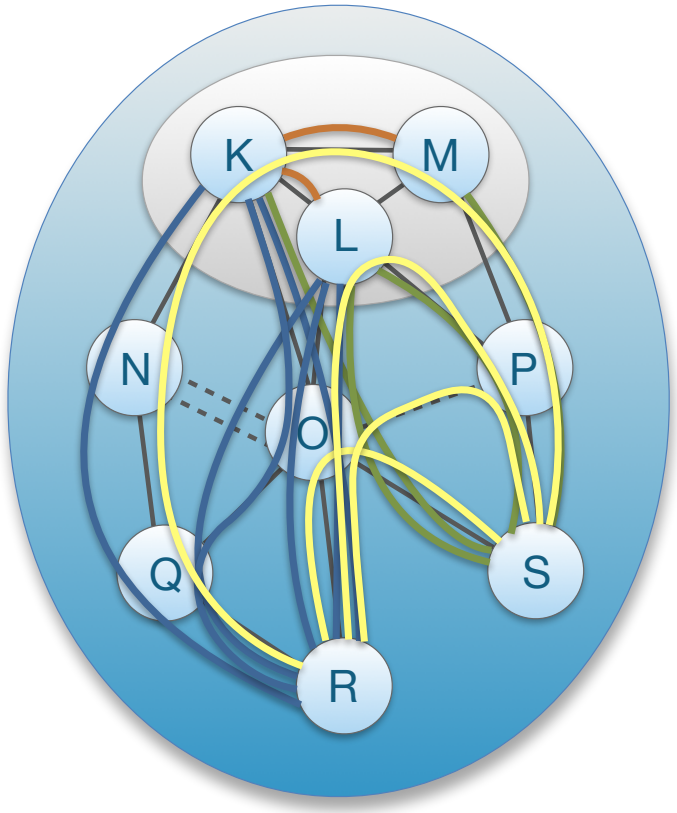

Communication to Remote ISD

- Host contacts local path server requesting <ISD, AS>
- If path segments are not cached, local path server will contact core path server
- If core path server does not have path segments cached, it will contact remote core path server
- Finally, host receives up-, core-, and down-segments

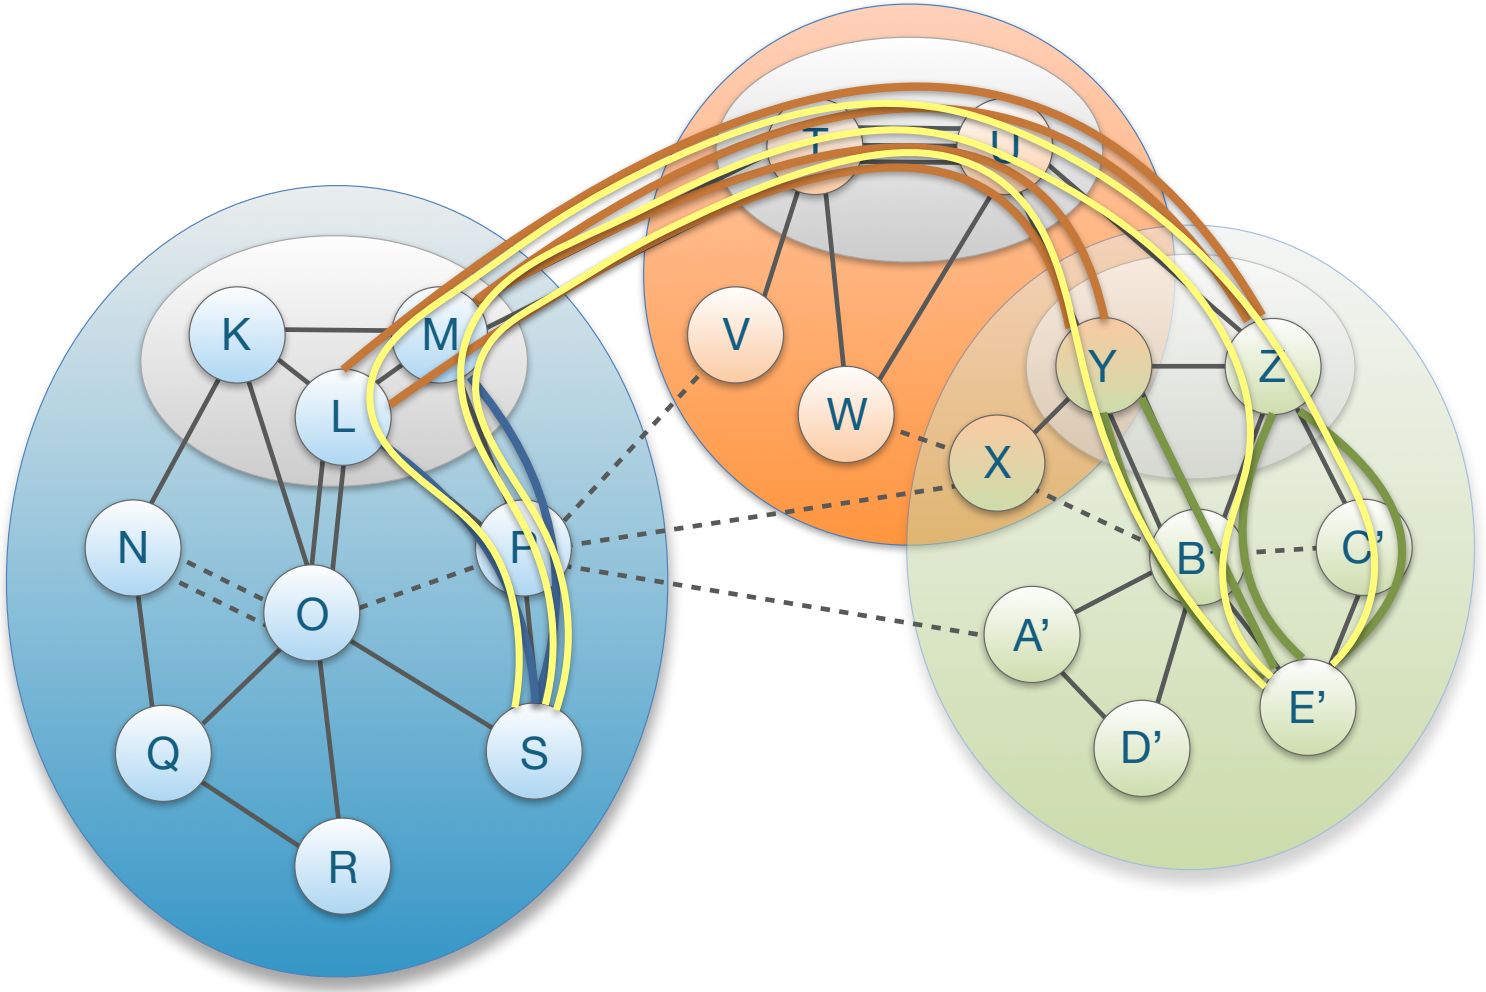

ETH zürich

SCION Drawbacks

Initial Latency Inflation

- ❖ Additional latency to obtain paths
- ✓ BUT amortized by caching & path reuse

Bandwidth Overhead

- ❖ Due to paths in the packets
- ❖ About 80 additional bytes
- ✓ Enables path control, simpler data plane, etc

Increased Complexity in Key Mgmt.

- ❖ New certificates (e.g., TRC Certificates)
- ✓ High security design

Initial Set-up Cost

- ❖ Training network operators
- ❖ Installing new infrastructures
- ✓ Offers methods to facilitate deployment

Use Case: IoT Protection through Hidden Path

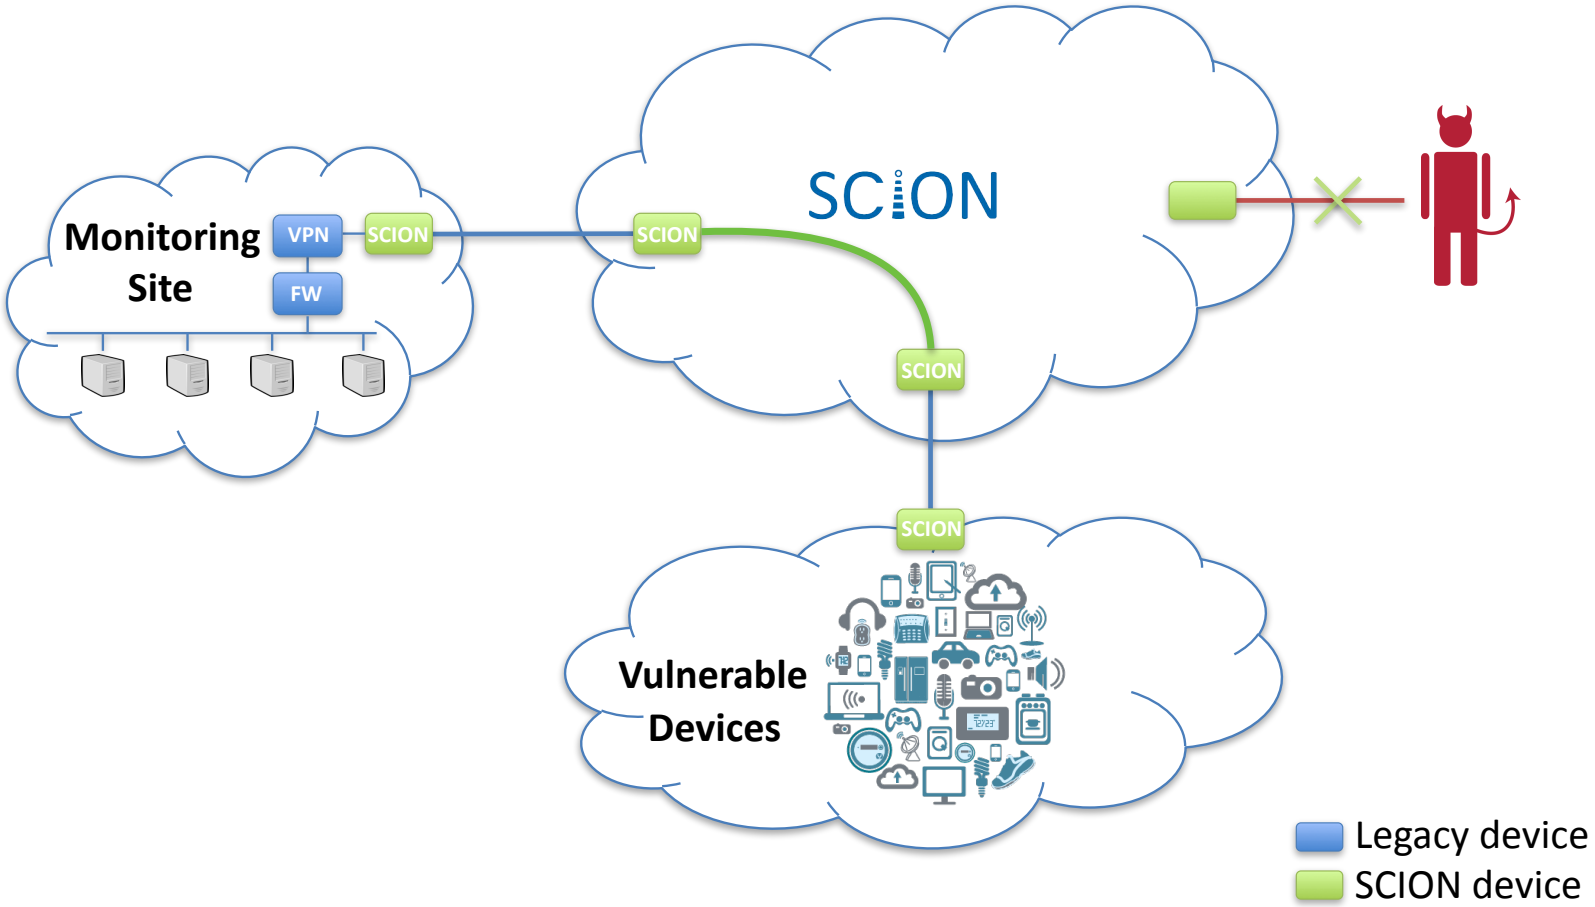

ETH zürich

SCION

Use Case: Low-Latency Connectivity

- Generally, two paths exist between Europe and Southeast Asia
  - High latency, high bandwidth: Western route through US, ~450ms RTT
  - Low latency, low bandwidth: Eastern route through Suez canal, ~250ms RTT
- BGP is a “money routing protocol”, traffic follows cheapest path, typically highest bandwidth path
- Depending on application, either path is preferred
- With SCION, both paths can be offered!

ETH zürich

Use Case: Low Earth Orbit Satellite Networks

- Previous satellite networks suffered from high latency for communication between earth and satellite
  - Geostationary satellites are at a distance of about 40’000km from earth, ~130ms latency
- New Low Earth Orbit (LEO) satellite networks are much lower and thus only require around 5ms propagation latency between earth and satellite
  - Distance about 1200km, ~4ms latency
  - Inter-Satellite Laser (ISL) links enable global communication
- Disadvantage: large number of satellites needed to provide complete coverage

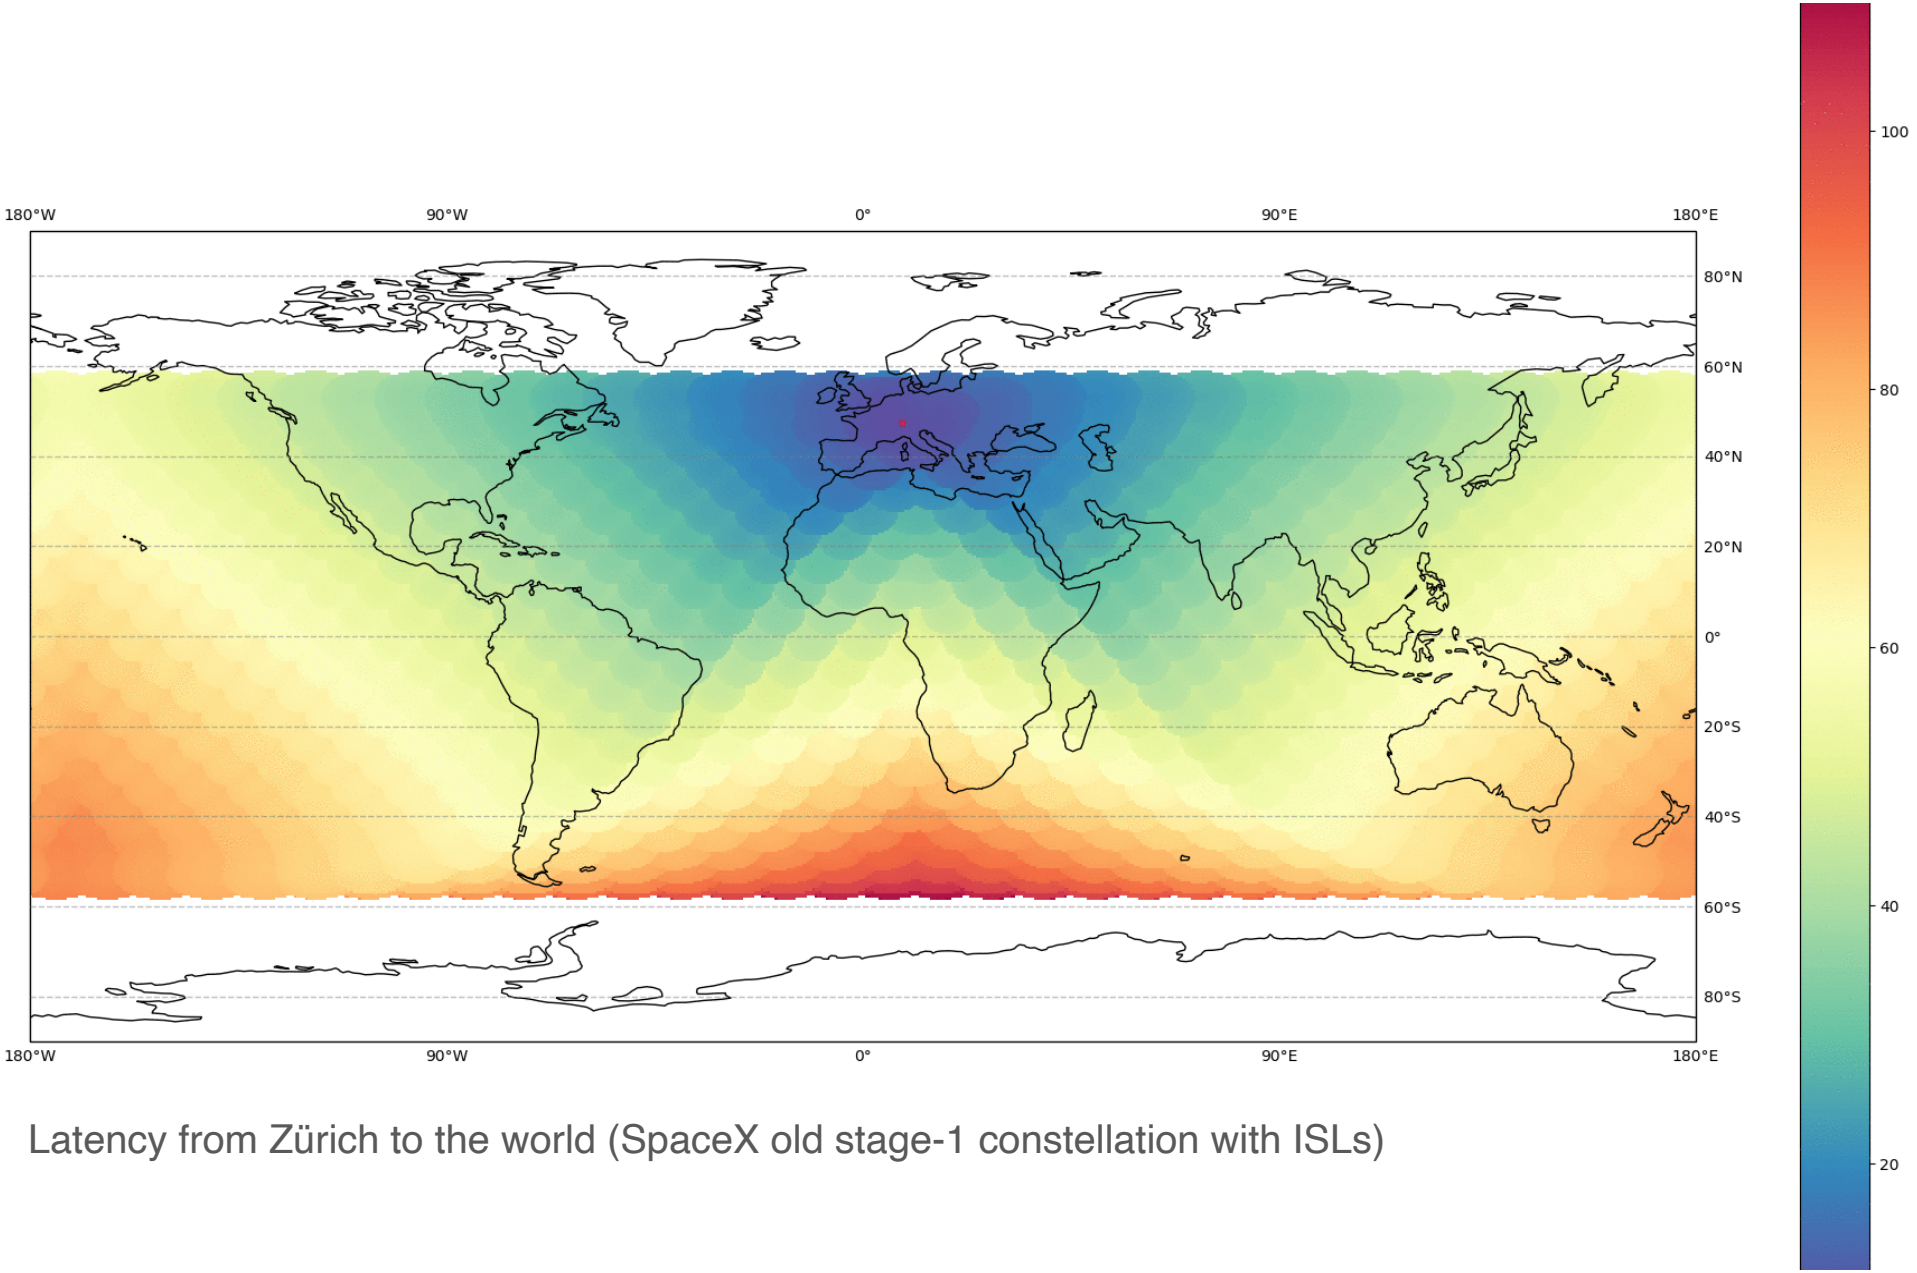

Latency from Zürich to the world (SpaceX old stage-1 constellation with ISLs)

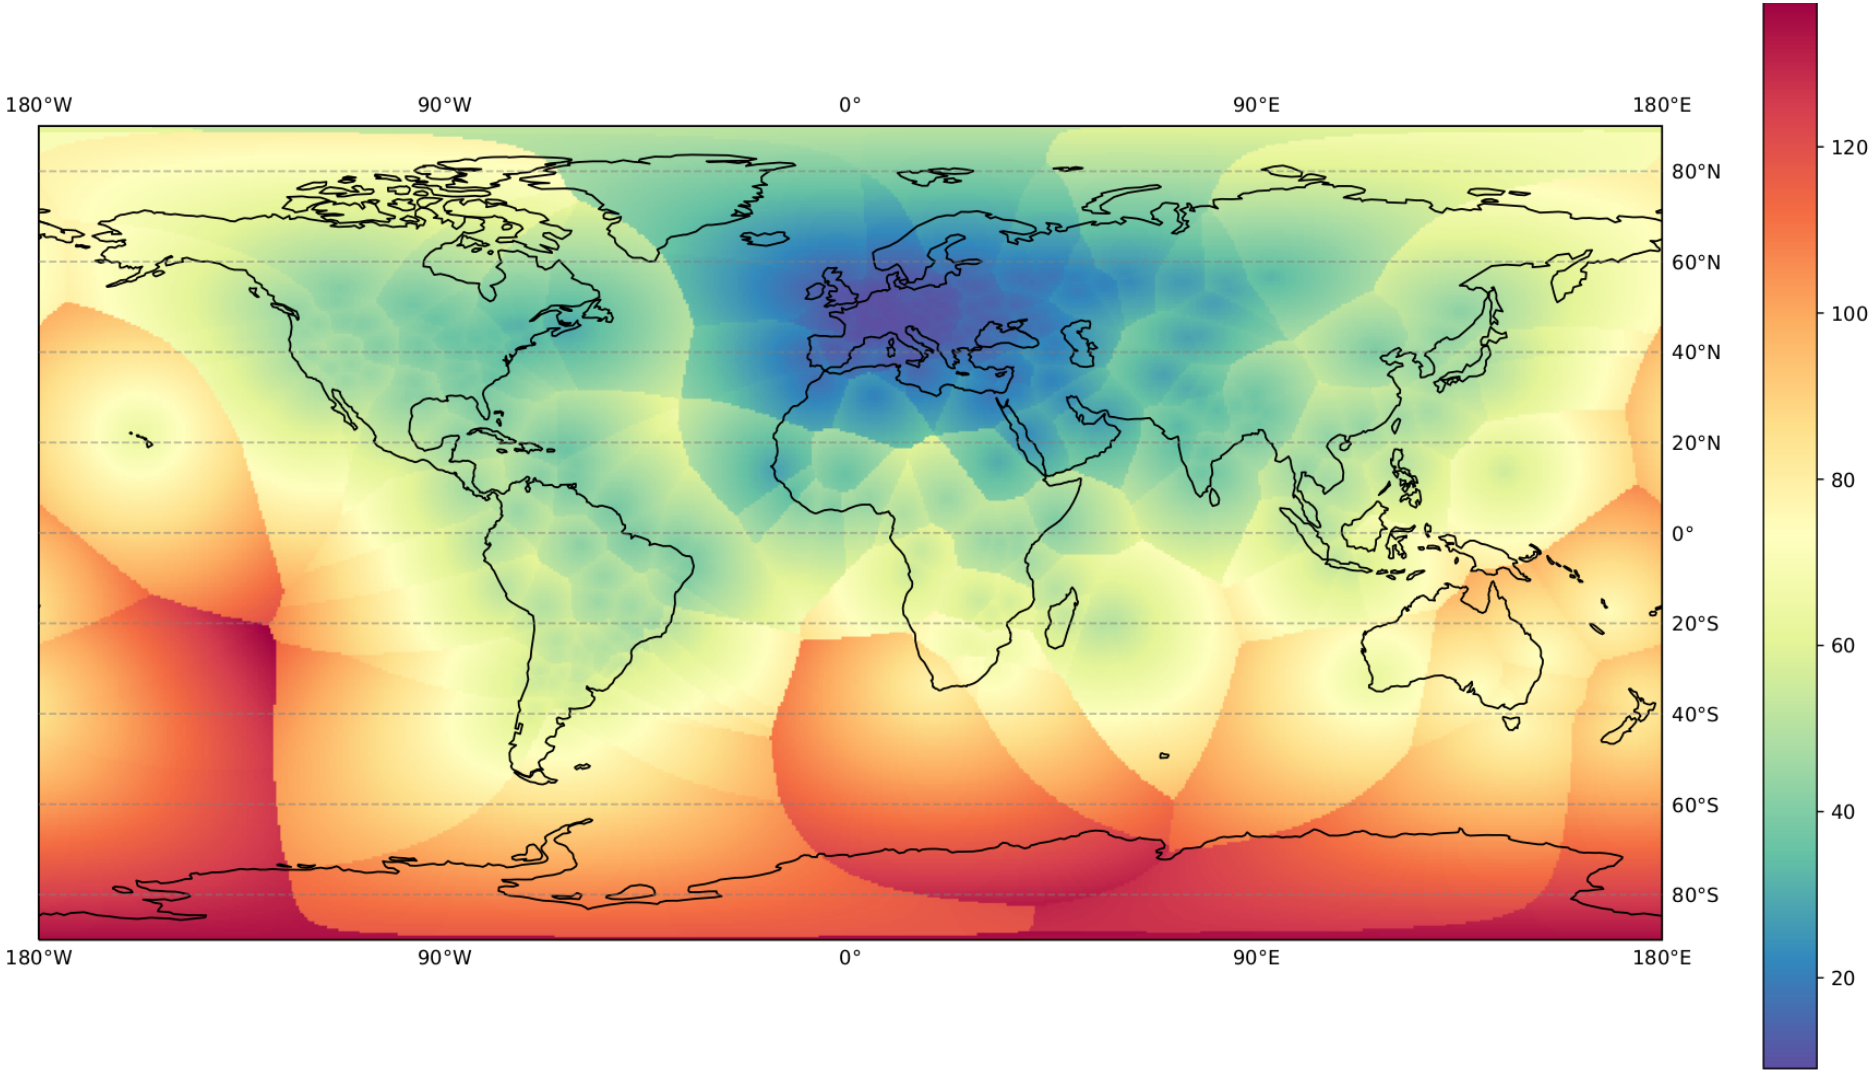

Latency from Zürich to the world, Satellite + IXP connection path

SCION Naturally Supports LEO Networks

- BGP convergence is too slow to support frequent outages / short time windows of availability for during initial deployment stages of LEO network
  - Clouds / rain can also prevent or reduce communication with satellite
- SCION can optimally integrate LEO network into Internet fabric
  - Satellite network paths can be announced next to regular Internet paths: end host can select optimal path based on bandwidth, latency, and cost
  - Beacons can be sent out before path becomes available, including start / end validity time
  - Based on weather prediction, expected bw can be added to beacon
  - End host can also select which satellite uplink station to send packets to
  - Receiver can select appropriate return link, could be terrestrial or satellite

Use Case: High-Speed Interdomain Failover

- Common failure scenarios in current Internet
  - Long-term failures (infrequent): large-scale failures require hours until BGP re-stabilizes
  - Intermediate-term failures (at each inter-domain router or link failure): 3-5 minutes until path is cleanly switched
  - Short-term failures (frequent): during BGP route change, routing loop during 5-10 seconds
- SCION: backup path is already set up and ready to be used when a link failure is observed
- Result: failover within milliseconds!

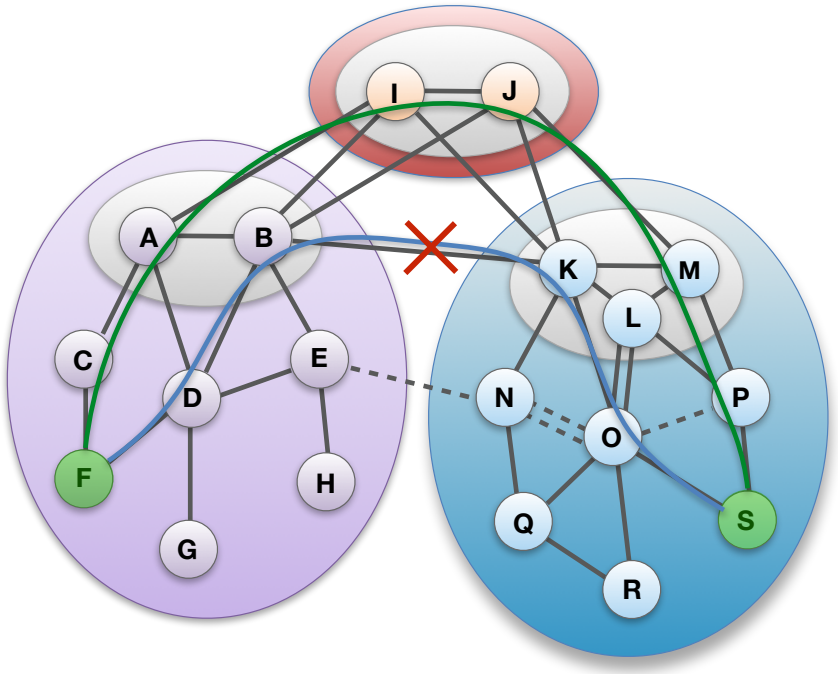

How to Deploy SCION – Core Network

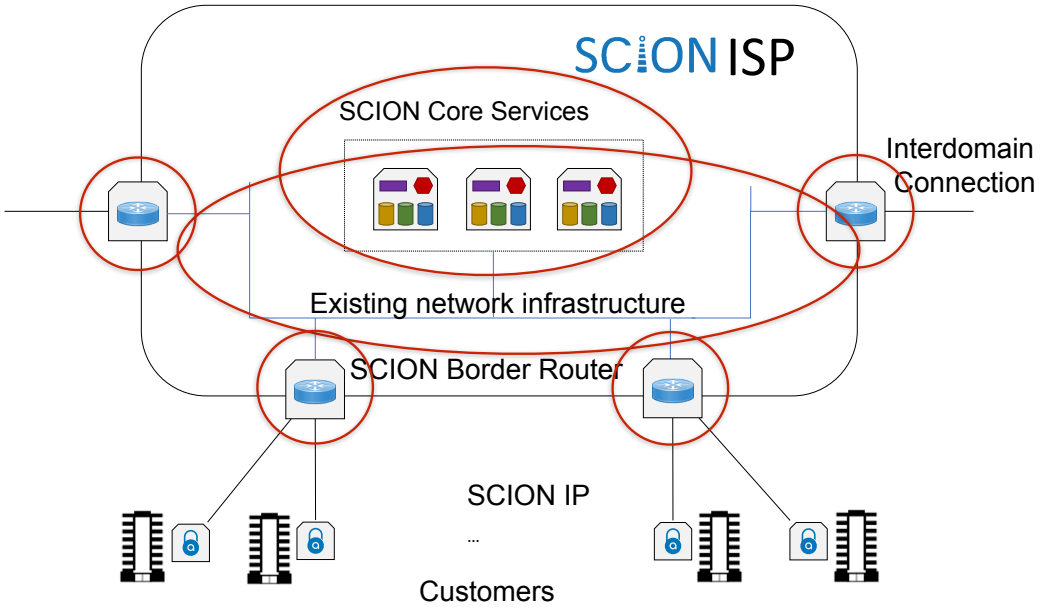

- Two components: SCION core services (control plane) and SCION border routers (data plane)
- SCION reuses existing intra-domain networking infrastructure—no need to upgrade all networking hardware

How to Deploy SCION – End Domains

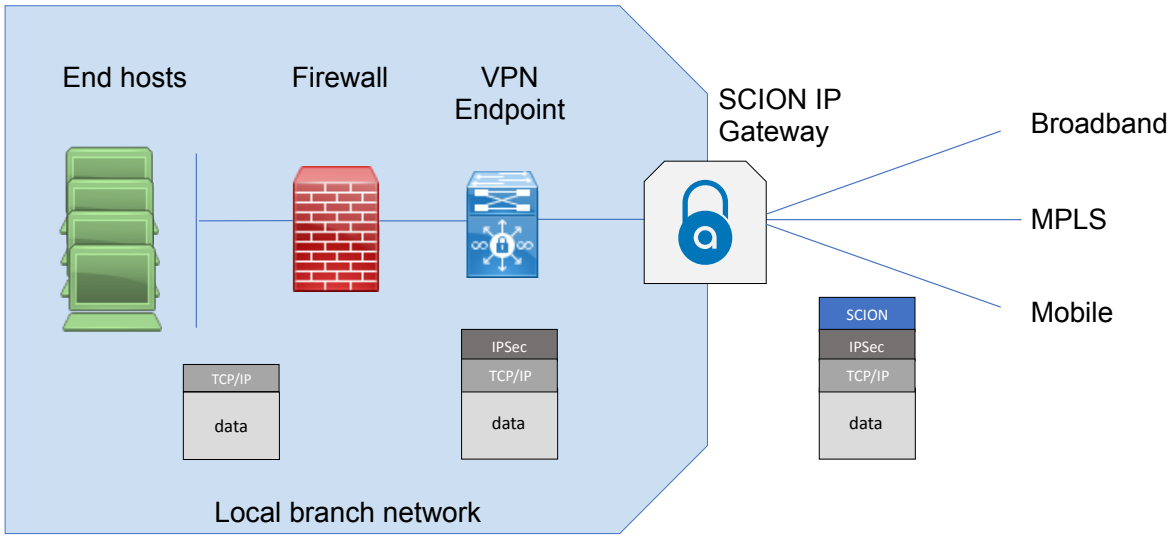

- SCION IP Gateway enables seamless integration of SCION capabilities in end-domain networks
- No upgrades of end hosts or applications needed
- SCION is transport-agnostic thus can work over many different underlaying networks

Recent Thrusts

- Main thrust: operationalize + drive deployment
- SCI-ED project
- SCIONLab
- Production network
- DRKey + control-plane PKI

SCI-ED: SCION for ETH Domain

- Goals
  - Large-scale real-world deployment: ETH, EPFL, PSI, CSCS, EMPA, EAWAG, WSL
  - Operationalize SCION in SWITCH network
  - Expand and demonstrate maturity of SCION on real-world use cases
- SCION use cases in the ETH Domain
  - High-performance data transmission
  - Secure communication of sensitive data
  - High availability for critical infrastructures
  - Platform for networking research

Approach for High-Speed Data Transmission

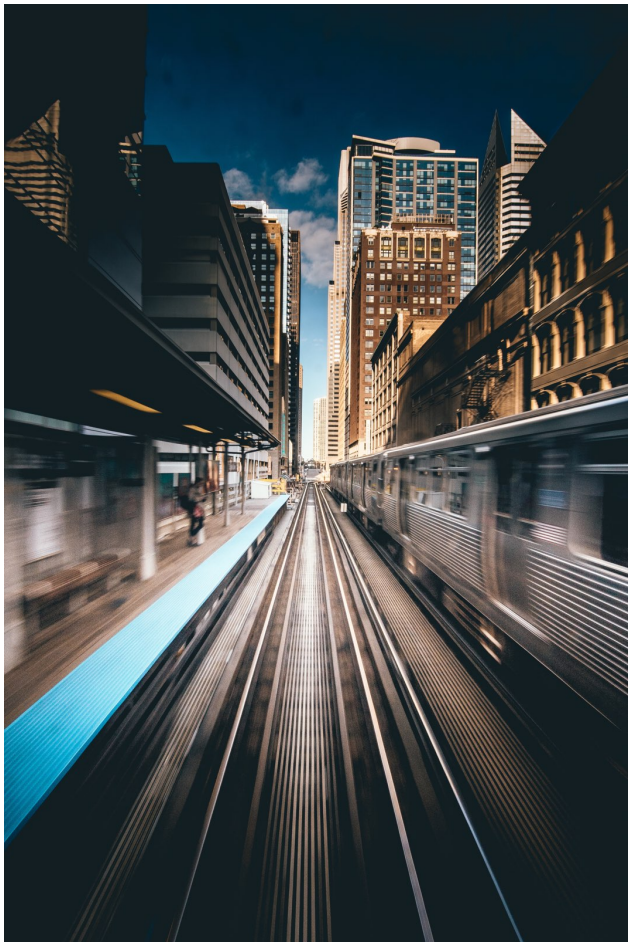

- Multipath communication, even backup links can be used simultaneously
- QUIC instead of TCP
- Firewall bypassing thanks to high-speed packet authentication
- Data transmission appliance to avoid changing end host

SCIONLab

- Global SCION research testbed
- Open to everyone: create and connect your own AS within minutes
- ISPs: Swisscom, SWITCH, KDDI, GEANT, DFN
- Korea: GLORIAD, KISTI (KREONET), KU, KAIST, ETRI
- Deployed 35+ permanent ASes worldwide, 600+ user ASes

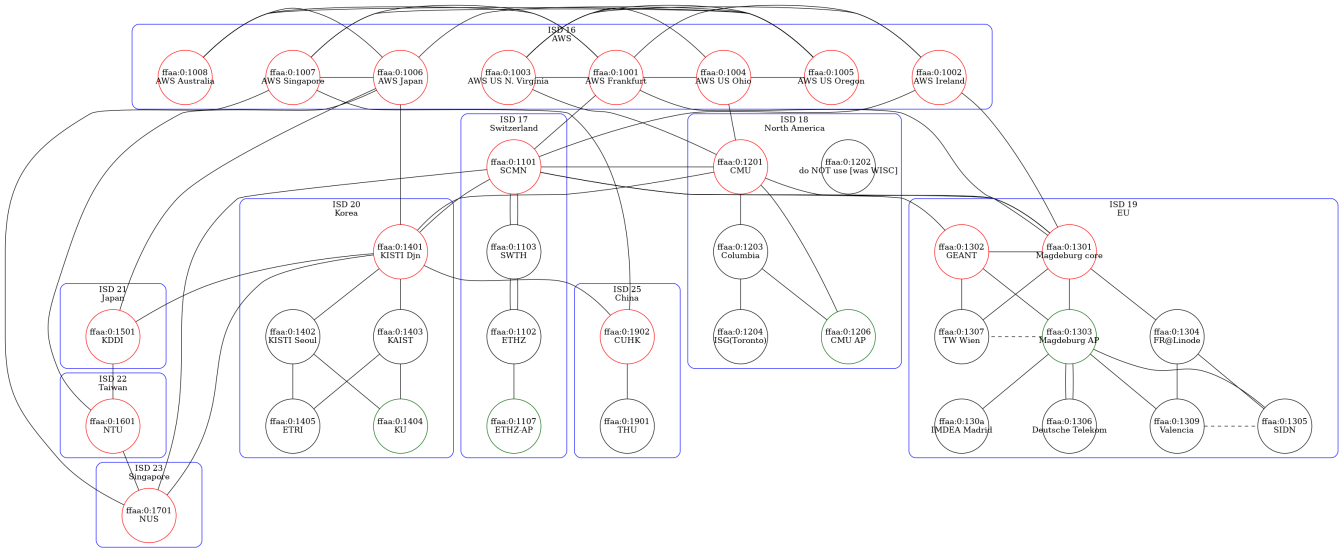

ETH zürich

SCION Production Network

- Led by Anapaya Systems 

- Important point: BGP-free global communication
  - We need failure-independence from BGP protocol
- Discussions with domestic and international ISPs
  - Goal: First **inter-continental public secure** communication network
- Construction of SCION network backbone at select locations to bootstrap adoption
- Current deployment
  - ISPs: Swisscom, Sunrise, SWITCH, +others
  - Bank deployment: 4 major Swiss banks, some in production use
  - Swiss government has SCION in production use

High-Speed Secure Communication

- Hercules: > 30Gbps file transfer using 1 core on commodity hardware
- LightningFilter: > 120Gbps firewall with per-packet cryptographic authentication on commodity hardware

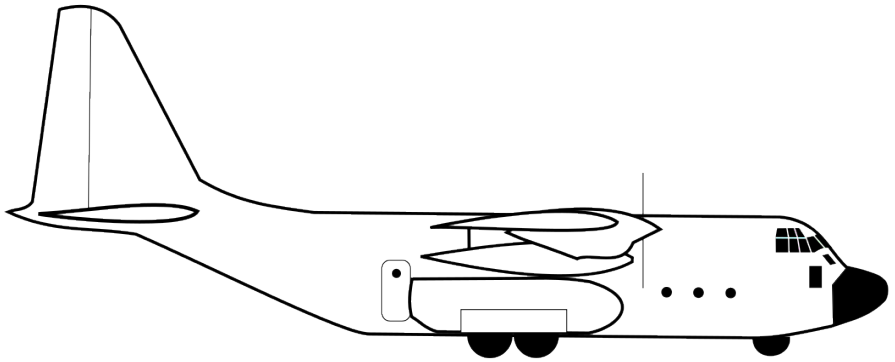

Hercules

Bulk Data Transfer over SCION

Matthias Frei and François Wirz

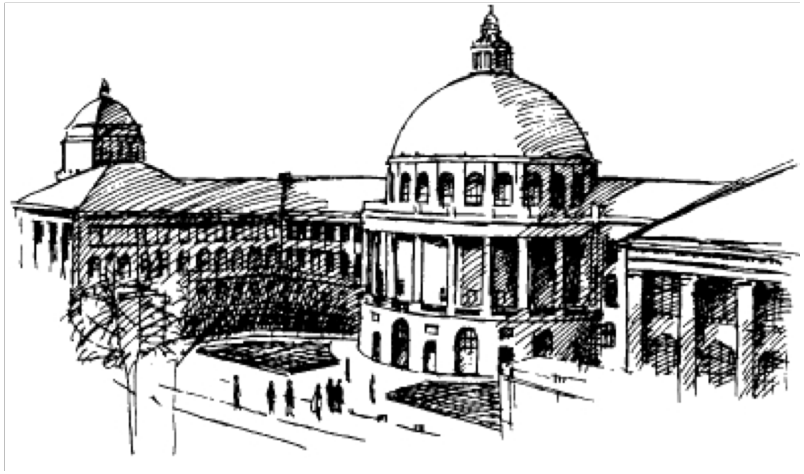

ETH zürich

SCION

Project Scope

High-speed large file transfer over Internet

- Large = Terabyte-scale data transfers

Use Cases

- Data-intensive science: healthcare, physics, big data, etc.
- Remote processing, data needs to be transmitted beforehand
- Remote backup

Traditional Approaches

- FTP over TCP/IP
  - TCP suffers from degraded performance with high latency and random losses
  - Poor multipath support
    - Open many TCP streams and hope and pray
    - Multipath TCP in the future
  - Poor utilisation of available capacity
- Science DMZ
  - Designated data transfer infrastructure, in front of enterprise firewall
  - Simple packet filter, whitelist source IPs

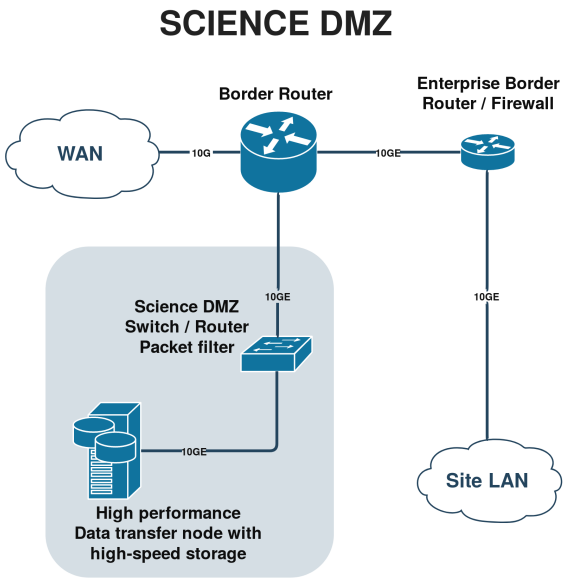

ETH zürich

SCION

How can SCION Speed up File Transfer?

- Clean multipath communication
  - Multiple disjoint paths
  - Utilize local backup links

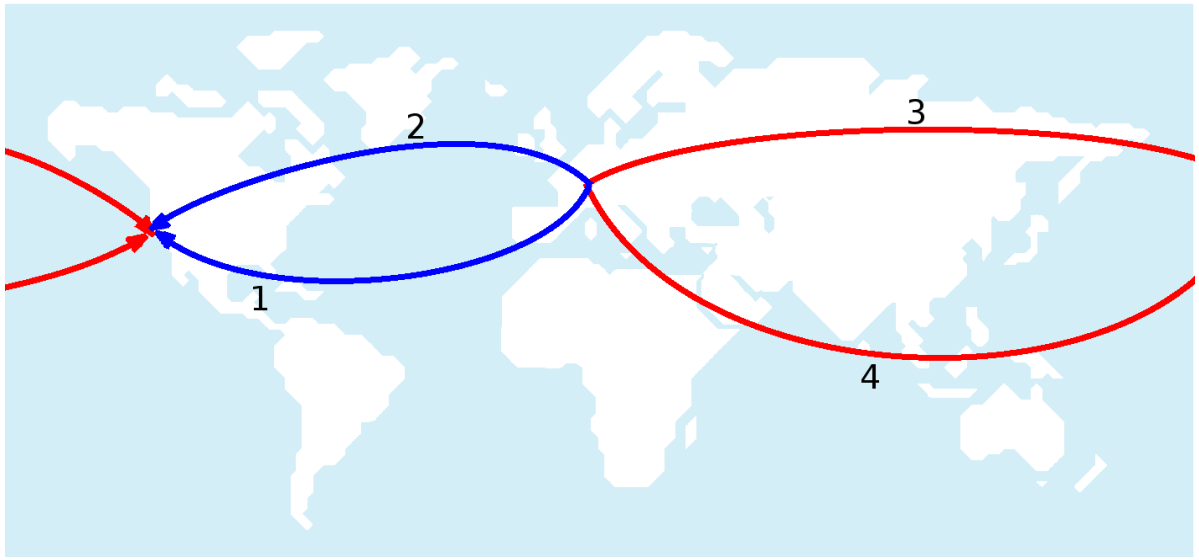

How can SCION Speed up File Transfer?

- Clean multipath communication
  - Multiple disjoint paths
  - Utilize local backup links
- Simplified congestion control & low loss thanks to COLIBRI quality-of-service system
- LightningFilter: packet filter for Science DMZ *with* strong cryptographic packet authentication

Hercules

- SCION/UDP packet blasting + retransmits
  - “Reliable Blast UDP”[1]
- Range ACKs at fixed frequency
- Performance-oriented congestion control [2]
  - Link empirical performance to actions taken

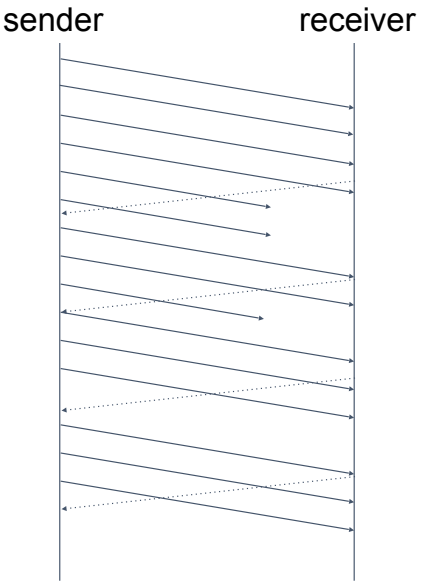

[1] "*Reliable Blast UDP : Predictable High Performance Bulk Data Transfer*", Eric He, Jason Leigh, Oliver Yu and Thomas A. DeFanti, Proceedings of IEEE Cluster Computing, Chicago, Illinois, September, 2002

[2] “PCC: Re-architecting Congestion Control for Consistent High Performance”, Mo Dong, Qingxi Li, Doron Zarchy, P. Brighten Godfrey, and Michael Schapira, 12th USENIX Symposium on Networked Systems Design and Implementation (NSDI 15)

Hercules

AF_XDP[3] for high performance SCION/UDP

- Published in December 2018
  available in Linux >= 4.18
  zero-copy mode in Linux >= 5.1
- Bypass Linux networking stack for send/receive
- Bypass SCION dispatcher

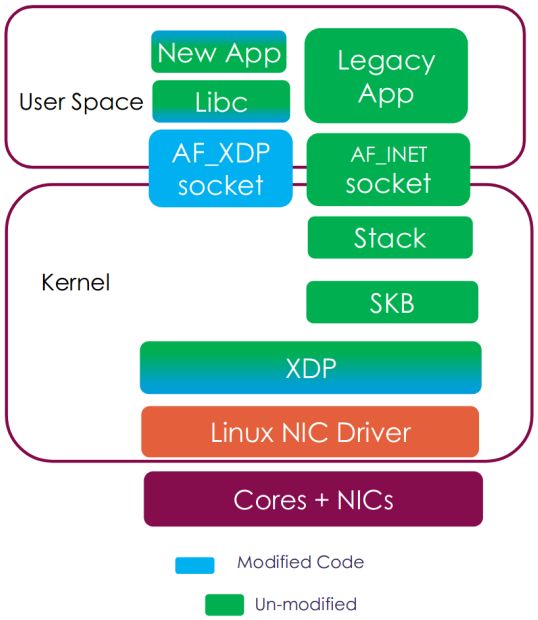

PMD for AF_XDP: Zhang Qi, Li Xiaoyun

[3] "Accelerating networking with AF_XDP", Jonathan Corbet, LWN.net, 2018

ETH zürich

SCION

Demo

- Transfer file between two SCION hosts in *same* AS
- Directly connected, 40GbE
- *Not* the target use case, but high-performance SCION links are being established

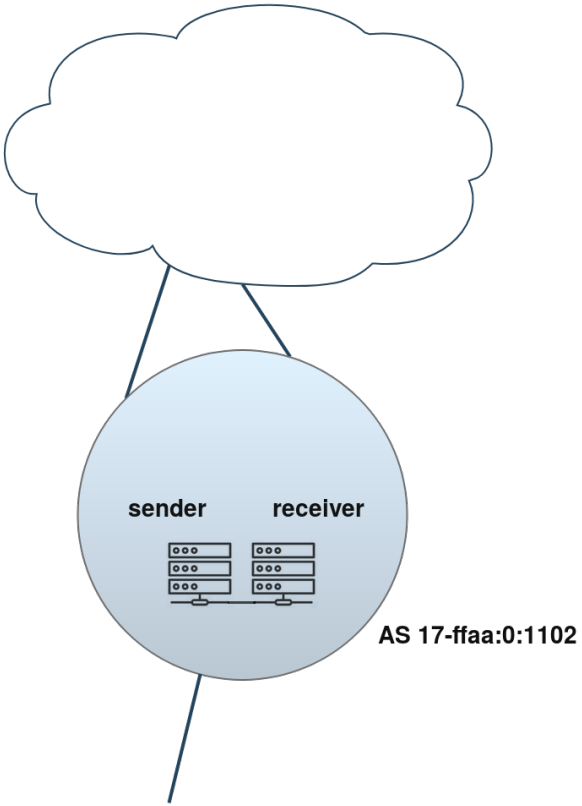

ETH zürich

SCION

Demo Summary

- Hercules achieves ~30Gbps transfer rate (using 1 core)
  - Disk I/O not included, much slower on demo host

- Comparison
  - `iperf3` with TCP achieves ~20Gbps (one thread)
  - `iperf3` with UDP, ~4Gbps
  - FTP achieves ~8Gbps

LightningFilter: Traffic Filtering at 120 Gbps

Benjamin Rothenberger

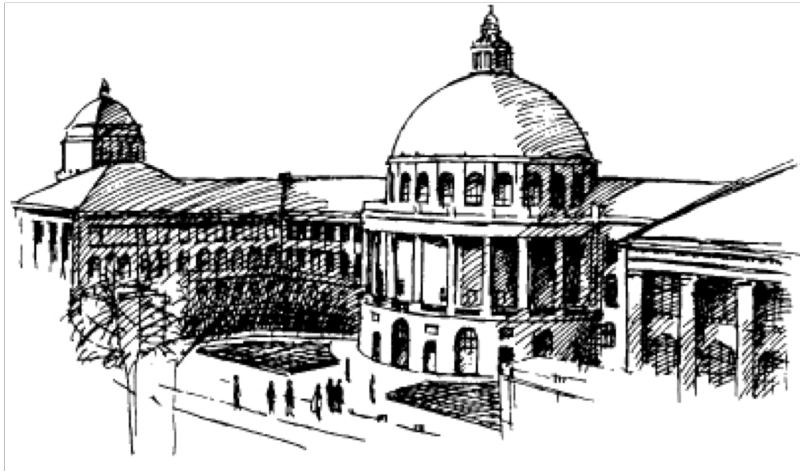

In collaboration with:

Prof. Adrian Perrig, Juan Garcìa Pardo, Dominik Roos, Jonas Gude, Pascal Sprenger, Florian Jacky

ETH zürich

SCION

Project Goals

- High-speed packet processing requires nanosecond operations
  - Example: 64-byte packets @ 100Gbps: ~5ns processing time

- Nanosecond scale key establishment
- Nanosecond scale packet authentication

- Trivia: how “long” is a nanosecond?
  - Answer: light travels about 30cm in 1ns

High-Speed Packet Processing

- Current high-speed Internet links: 400Gbit/s (Gbps)
- Arrival rate for 64-byte packets: one packet every 1.3 ns
- High-speed asymmetric signature implementation: Ed25519 SUPERCOP REF10: ~ 100μs per signature
- AES-NI instruction only requires 30 cycles: ~ 10ns
- Memory lookup from DRAM requires ~ 200 cycles: ~ 70ns
- Only symmetric crypto enables high-speed processing through parallel processing and pipelining

DRKey & Control-Plane PKI

- SCION offers a global framework for authentication and key establishment for secure network operations
- Control-pane PKI
  - Sovereign operation thanks to ISD concept
  - Every AS has a public-key certificate, enabling AS authentication
- DRKey
  - High-speed key establishment (within 20 ns), enabling powerful DDoS defense

Dynamically Recreatable Key (DRKey)

- ***Idea***: use a per-AS secret value to derive keys with an efficient Pseudo-Random Function (PRF)
- Example: AS X creates a key for AS Y using secret value SV_X
  - $K_{X \rightarrow Y} = PRF_{SVx}$ ("Y")
  - Intel AES-NI instructions enable PRF computation within 30 cycles, or 70 cycles for CMAC
    Key computation is 3-5 times faster than DRAM key lookup!
  - Any entity in AS X knowing secret value SV_X can derive $K_{X \rightarrow *}$

DRKey Performance

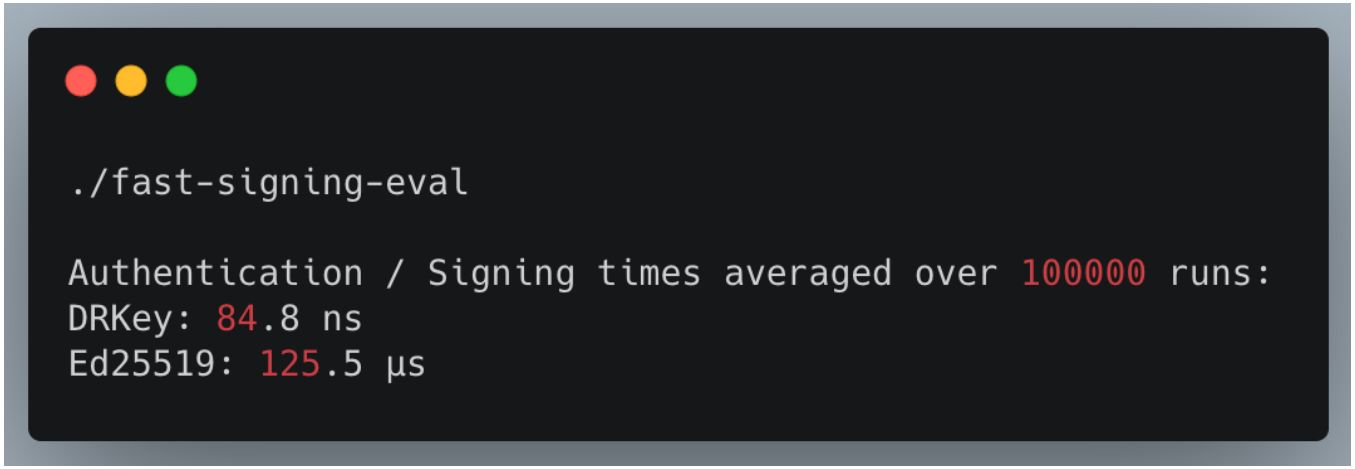

Factor:
~ 1450x

DRKey Use Case: SCMP Authentication

- Border router in AS B can derive key $K^{scmp}_{B \rightarrow A:Source}$ from SV_B
- Host “Source” can fetch key from local key server KS_A to authenticate SCMP message

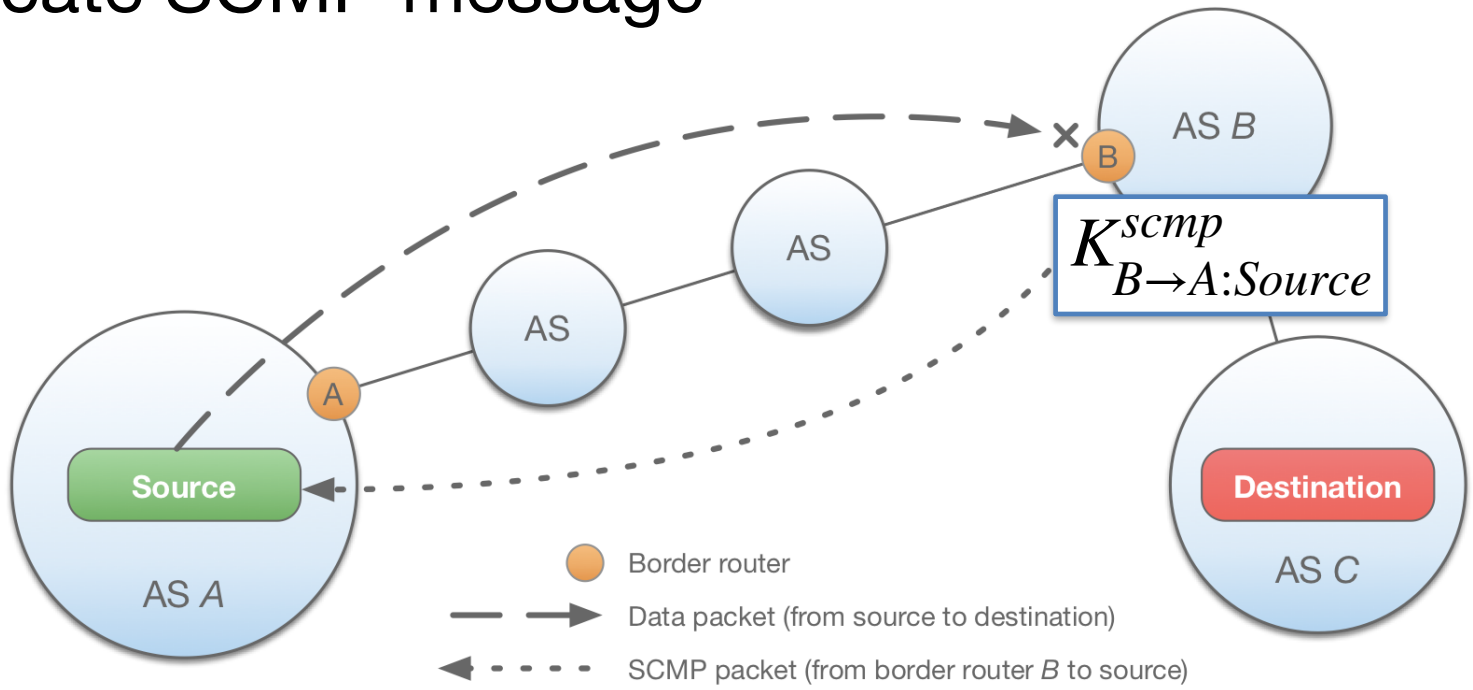

Lightning Filter

Traffic Filtering at 100 Gbps

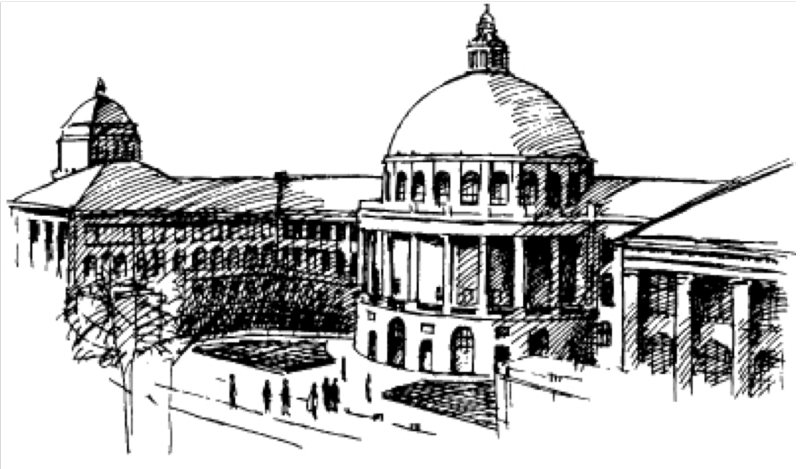

ETH zürich

SCION

Overview

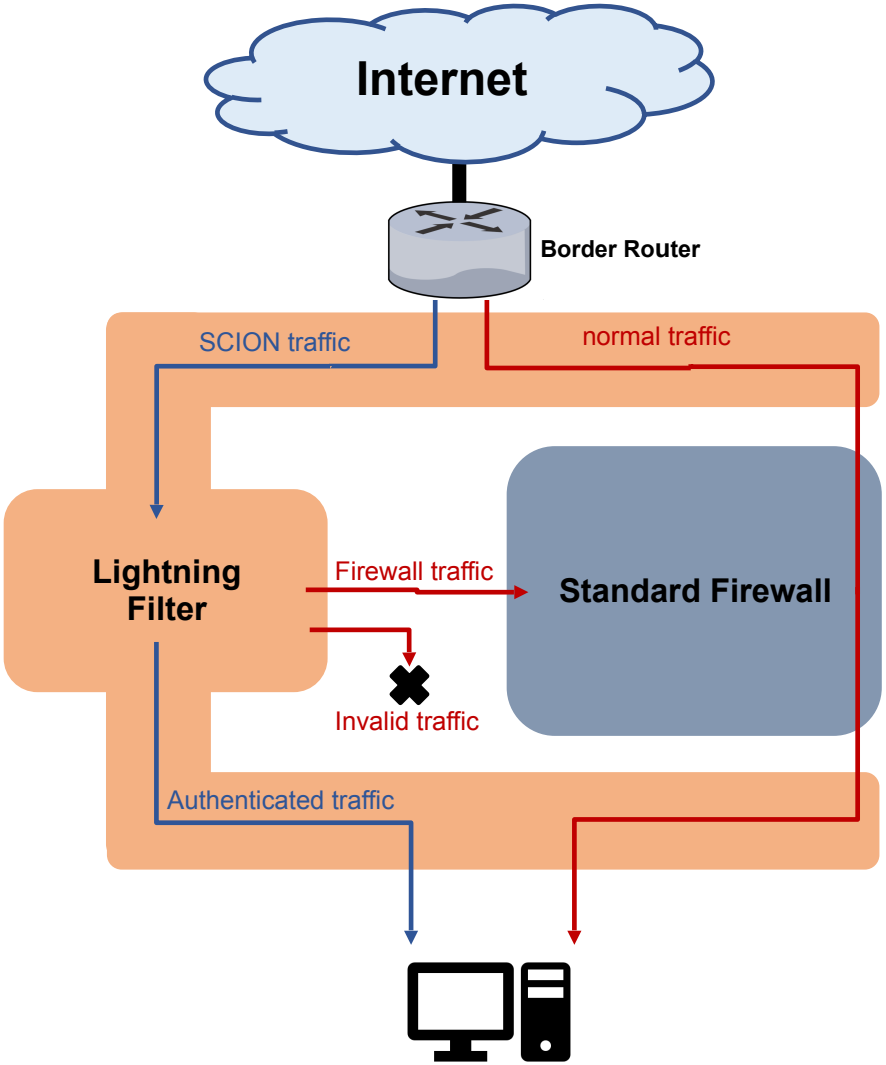

$$$

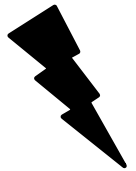

ETH zürich

SCION

System Design

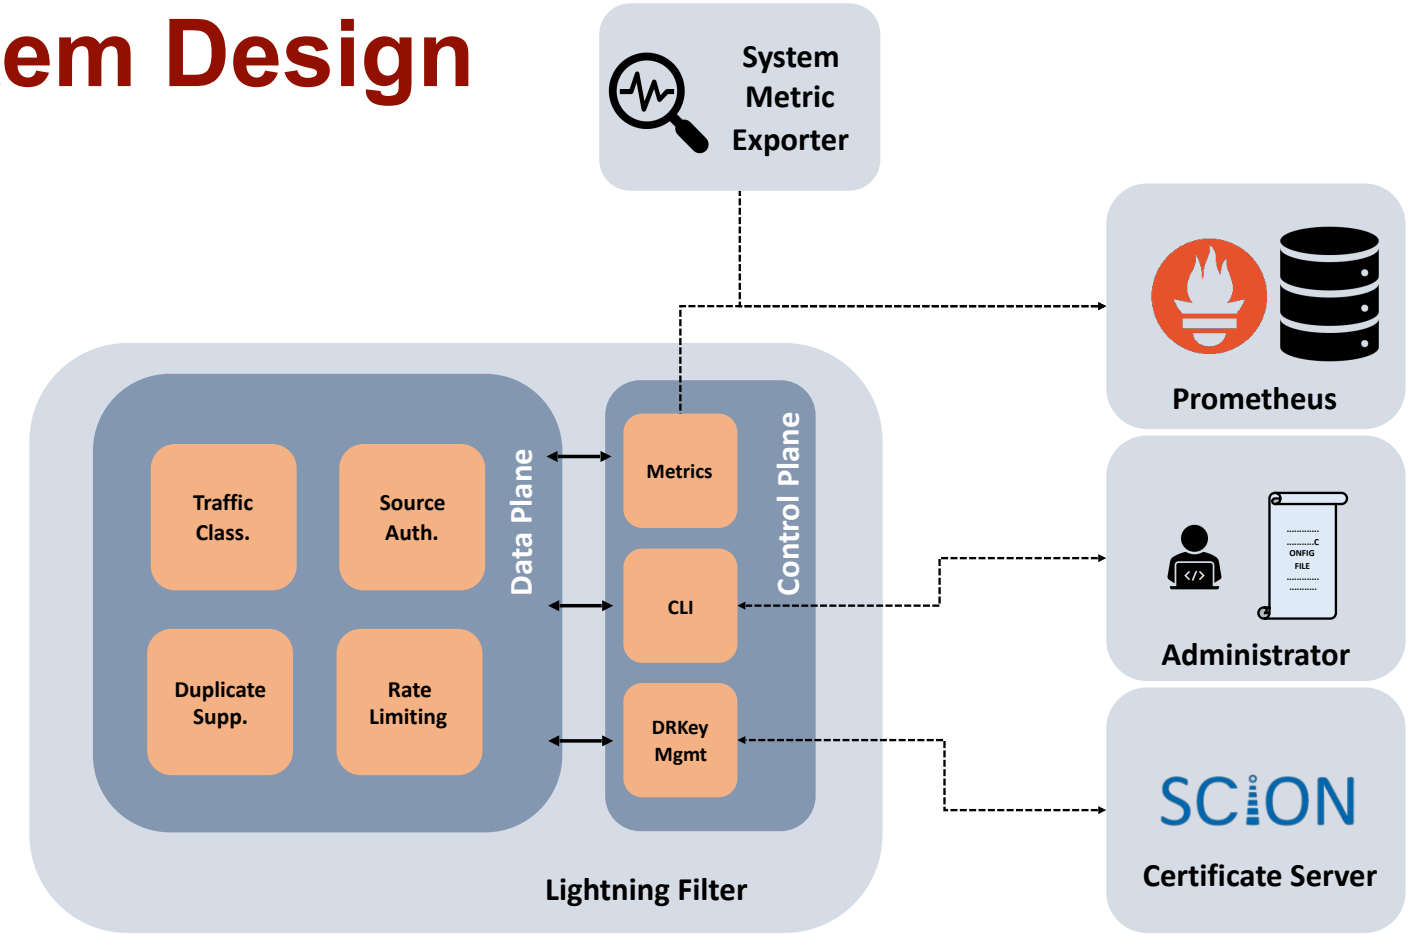

ETH zürich

SCION

Demo Outline

1. Attack scenario
   - Attacker located anywhere in Internet → Source authentication

2. Bandwidth capacity
   - 120 Gbps traffic volumne
3. Filtering based on source authentication
   - Alternate between filtering and bypass every 30s
4. Duplicate suppression
   - 80 Gbps duplicates traffic, 40 Gbps legitimate traffic

Online Resources

- https://www.scion-architecture.net
  - Book, papers, videos, tutorials
- https://www.scionlab.org
  - SCIONLab testbed infrastructure
- https://www.anapaya.net
  - SCION commercialization
- https://github.com/scionproto/scion
  - Source code

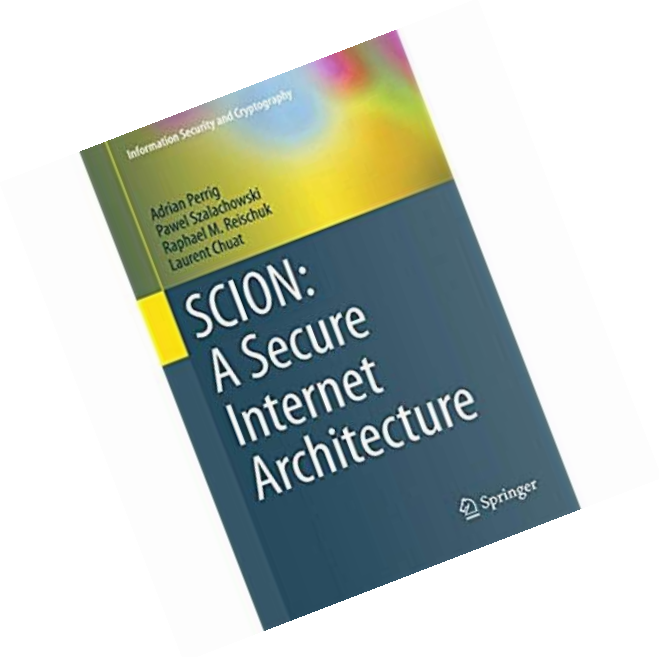

Summary

- Future Internet enables application-specific optimizations to provide enhanced efficiency
- Path-aware networking + multi-path networks are a promising direction to realize the future Internet vision
- High security and availability provide further benefits
- Join the effort, try out SCION today
  - SCIONLab research testbed
  - Production network

Thank you for your attention!