Global Communication Guarantees in the Presence of Adversaries

Adrian Perrig, Network Security Group

ETH zürich

SCION

The Internet is on Fire!

- Lack of sovereignty
- Frequent outages
  - https://downdetector.com
- Constant DDoS attacks
  - https://www.digitalattackmap.com
- Frequent routing attacks
  - https://bgpstream.com
- Lack of communication guarantees
- Expensive maintenance

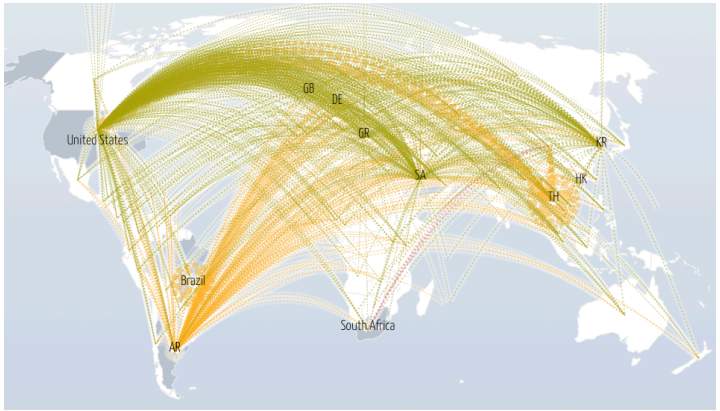

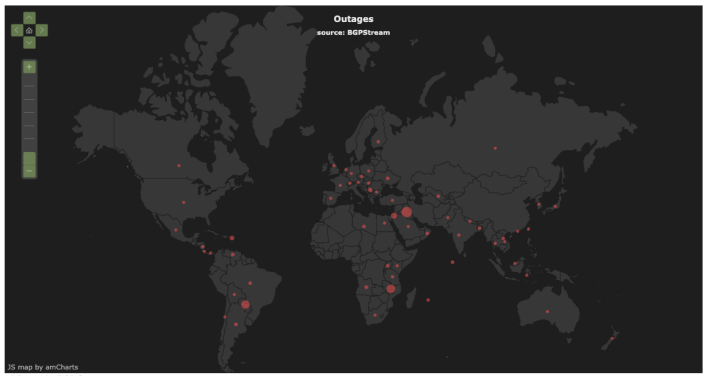

All Events for BGP Stream.

Event type	Country	ASN	Start time (UTC)	End time (UTC)	More info
Possible Hijack		*Expected Origin AS:* ZOHO-EU, NL (AS 205111) *Detected Origin AS:* LVLT-3549, US (AS 3549)	2020-10-06 01:01:28		More detail
Possible Hijack		*Expected Origin AS:* ZOHO-EU, NL (AS 205111) *Detected Origin AS:* LVLT-3549, US (AS 3549)	2020-10-06 01:01:28		More detail
Outage		SWIFTNETBROADBAND-AS SWIFTNET BROADBAND PRIVATE LIMITED, IN (AS 133713)	2020-10-05 22:18:00	2020-10-05 22:22:00	More detail
Outage		U-LAN-AS, RU (AS 48128)	2020-10-05 21:24:00		More detail
Outage		TPODLASIE, PL (AS 39375)	2020-10-05 20:00:00	2020-10-05 20:52:00	More detail

ETH zürich

SCION

Inspirations for a New Beginning

- Many exciting next-generation Internet projects over the past 25 years
- General Future Internet Architectures (FIA)
  - XIA: enhance flexibility to accommodate future needs
  - MobilityFirst: empower rapid mobility
  - Nebula (ICING, SERVAL): support cloud computing
  - NIMROD: improved scale and flexibility
  - NewArch (FARA, NIRA, XCP)
  - RINA: clean API abstractions simplify architecture
- Content-centric FIAs: NDN, CCNx, PSIRP, SAIL / NETINF
- Routing security: BGPSEC, S-BGP, soBGP, psBGP, SPV, PGBGP, H-NPBR
- Path control: MIRO, Deflection, Path splicing, Pathlet, I3
- Inter-domain routing proposals: ChoiceNet, HLP, HAIR, RBF, AIP, POMO, ANA, ...
- Intra-domain / datacenter protocols: SDN, HALO, ...

Why attempt redesigning Internet Architecture?

- We started our expedition asking the question:
  How secure can a global Internet be?
  - Answer: global communication guarantees can be achieved as long as a path of benign domain exists
- During our journey we discovered that path-aware networking and multi-path communication are powerful concepts that can provide higher efficiency than single-path Internet
  - Enables path optimization depending on application needs
  - Simultaneous use of several paths unlocks additional bandwidth
- Explore new networking concepts without the constraints imposed by current infrastructure!

SCION Ambition: A Global Next-Generation Public Internet

SCION Architecture Principles

- Stateless packet forwarding (no inconsistent forwarding state)
- “Instant convergence” routing
- Path-aware networking
- Multi-path communication
- High security through design and formal verification
- Sovereignty and transparency for trust roots

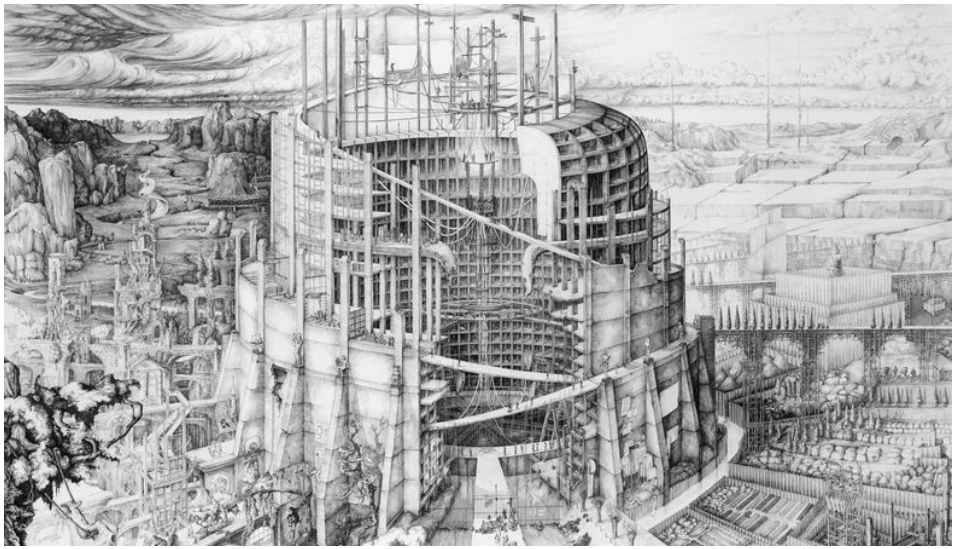

ETH zürich

SCION

Multi-path Communication is a Necessity Not a Luxury

- Necessary for high availability
  - Rapid failover without routing system convergence
- Enables higher network capacity
  - No more passive links for redundancy, all links can be active
  - Simultaneous use of several links
- Enables higher communication efficiency
  - Latency- vs. bandwidth optimal paths can be chosen
- Helps defend against DoS attacks, as adversary needs to congest all links
- QoS needs multi-path, as several alternatives need to be available to attempt resource reservations

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!

Approach for Scalability: Isolation Domain (ISD)

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

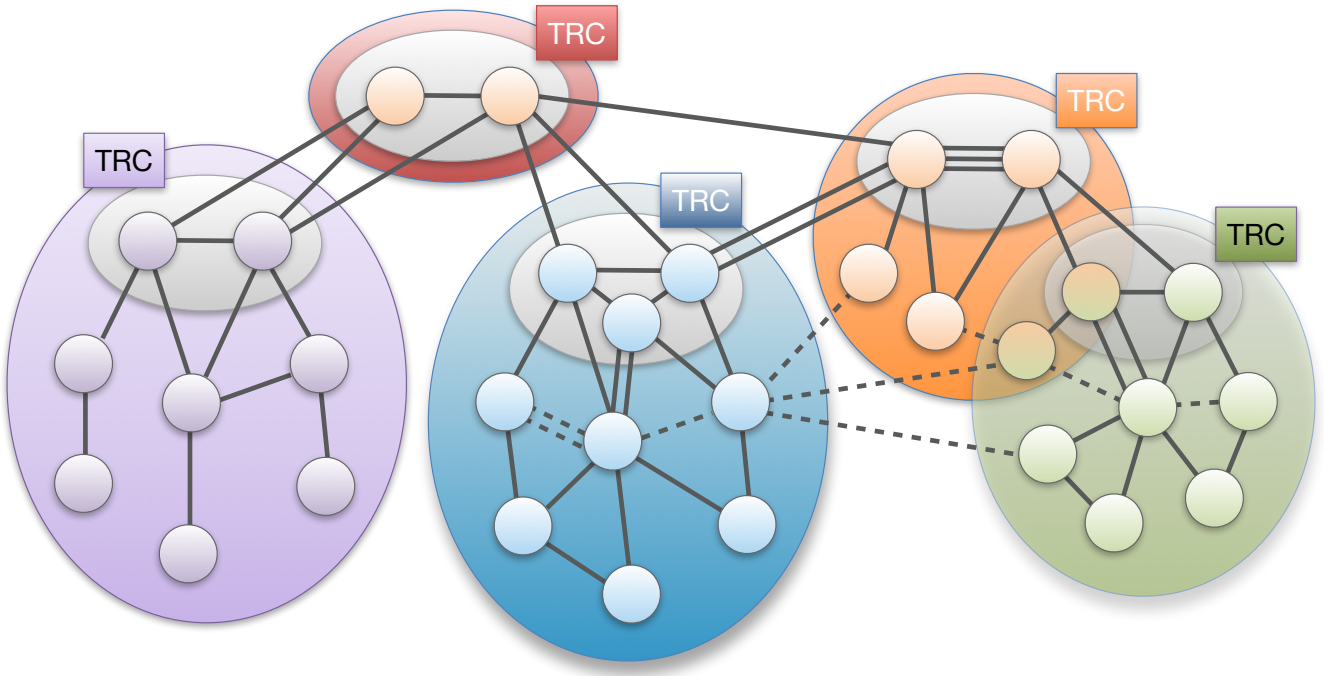

SCION Overview in One Slide

Path-based Network Architecture

Control Plane - Routing

- Constructs and Disseminates Path Segments

Data Plane - Packet forwarding

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

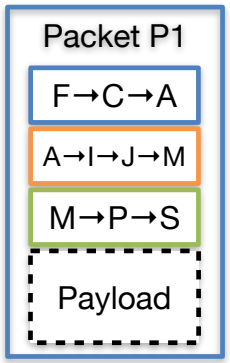

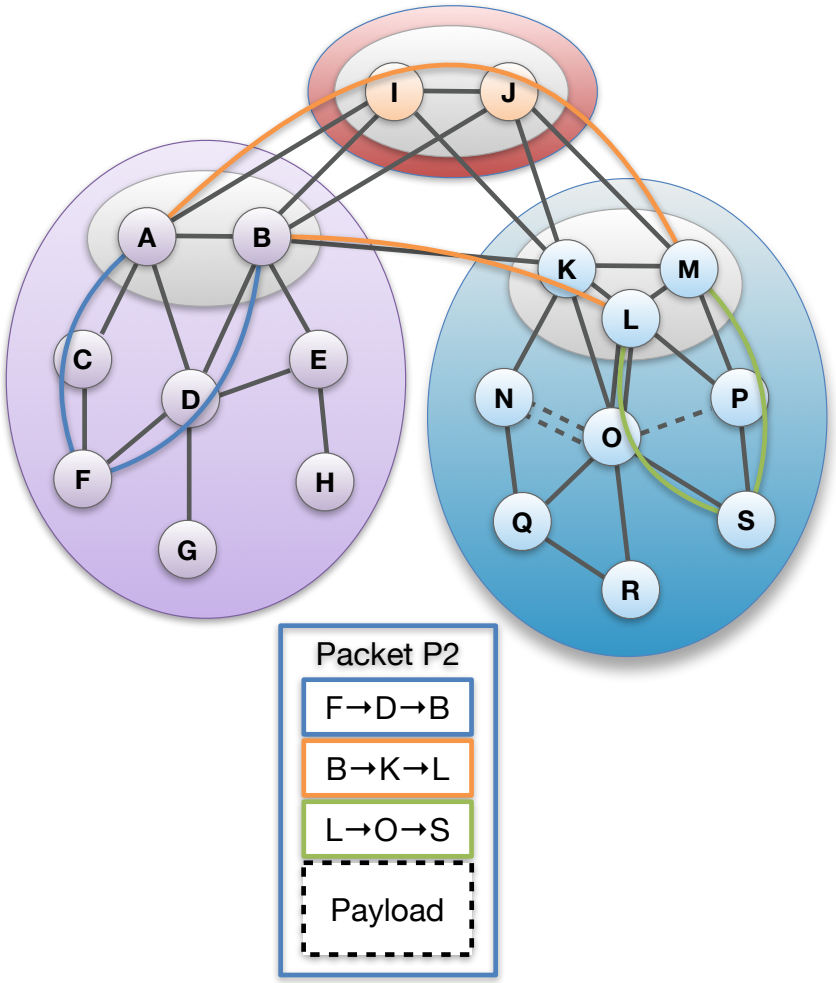

SCION

How to Deploy SCION – Core Network

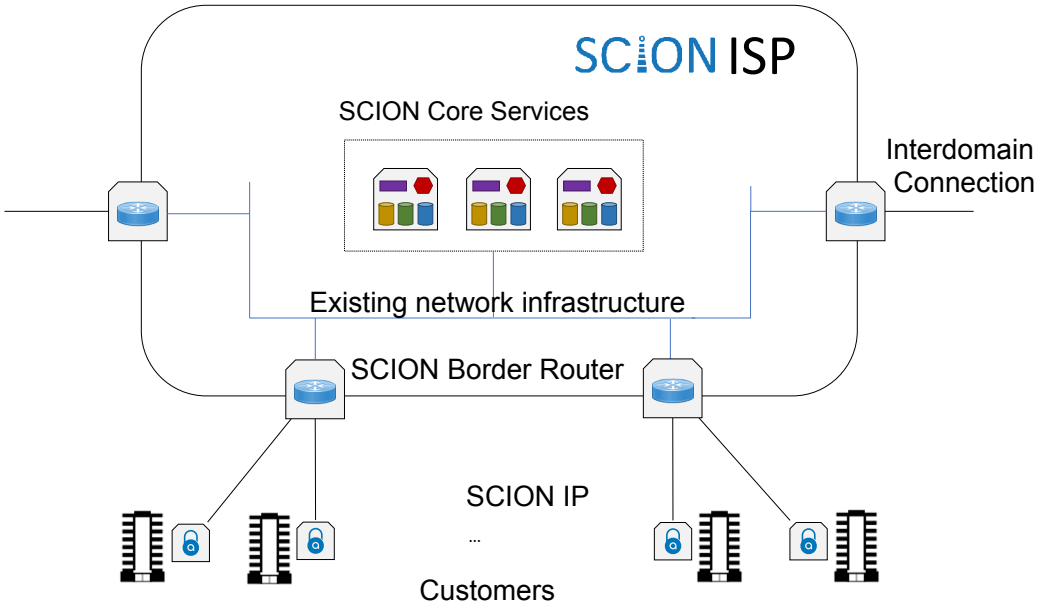

- Two components: SCION core services (control plane) and SCION border routers (data plane)
- SCION reuses existing intra-domain networking infrastructure—no need to upgrade all networking hardware

How to Deploy SCION – End Domains

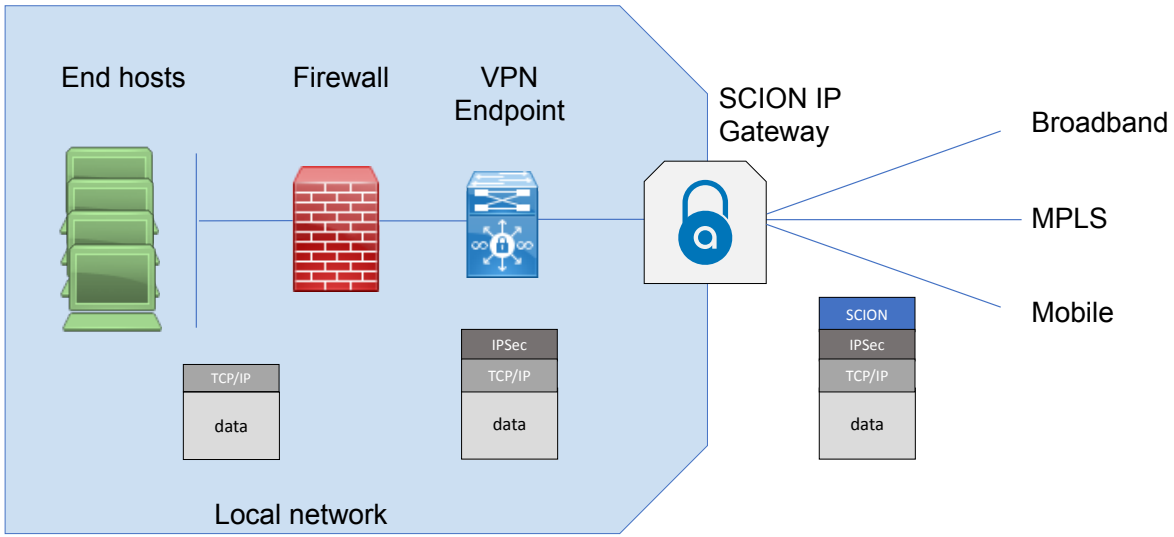

- SCION IP Gateway enables seamless integration of SCION capabilities in end-domain networks
- No upgrades of end hosts or applications needed
- SCION is transport-agnostic thus can work over many different underlaying networks

SCIONLab

- Global SCION research testbed: https://www.scionlab.org
- Open to everyone: create and connect your own AS within minutes
- ISPs: Swisscom, SWITCH, KDDI, GEANT, DFN
- Deployed 35+ permanent ASes worldwide, 600+ user ASes

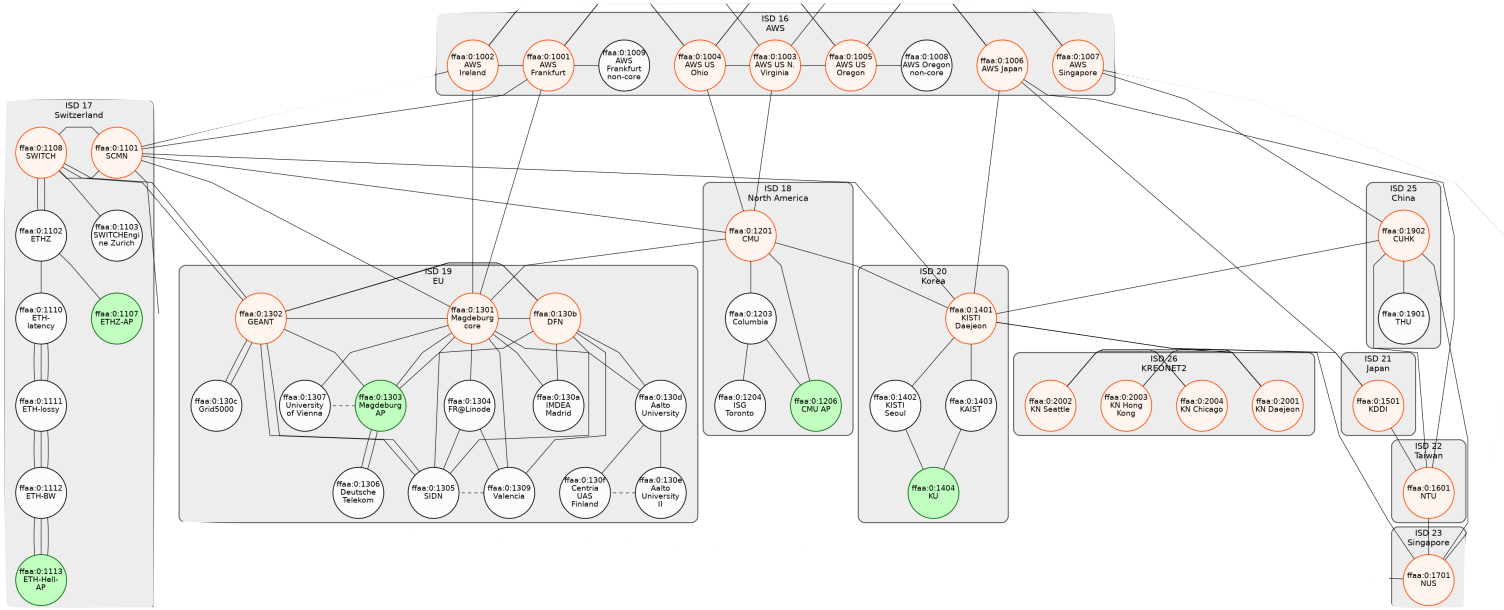

SCION Production Network

- Led by Anapaya Systems 

- **BGP-free global communication**
  - Fault independent from BGP protocol

- Deployment with domestic and international ISPs
  - Goal: First **global public secure** communication network
- Construction of SCION network backbone at select locations to bootstrap adoption
- Current deployment
  - ISPs: Swisscom, Sunrise, SWITCH, +others underway
  - Bank deployment: 4 major Swiss banks, some in production use

ETH zürich

SCION

Three SCION Project Thrusts

- **Thrust I: Security**
  - Sovereignty
  - Transparency
  - Routing security
  - DDoS resilience
  - Secure web connectivity (PKI)
  - Formal verification of protocols and code

ETH zürich

- **Thrust II: Efficiency**
  - Higher network capacity
  - Low-latency paths
  - High-bandwidth paths
  - Simultaneous use of multiple links
  - Fast failover
  - High-speed firewall

SCION

- **Thrust III: Green net**
  - Energy reduction vs. current Internet
  - More efficient forwarding
  - Use idle backup links
  - Improved network utilization
  - QoS savings (zero-loss, limited ACKs)

SCION Summary

- SCION: Next-generation Internet you can use today!
- High-performance
  - Path-aware network enables application-specific optimizations to provide enhanced efficiency
  - Multi-path communication enables simultaneous use of multiple paths, increasing available bandwidth
- Secure, high assurance, high availability
  - Per-packet authentication verification possible on routers
  - Formal verification of protocols and code
  - Immune against routing attacks, e.g., BGP prefix hijacking

ETH zürich

SCION

Global Communication Guarantees in the Presence of Adversaries

- Goal: If (routing policy compliant) path of benign ASes exists (with operational infrastructure), sender can find, use, and achieve minimum bandwidth guarantee on that path
- Challenges
  - Network routing instabilities, misconfigurations, etc.
  - DoS attacks at various levels (control plane, data plane, end host)

Observation: Stable Forwarding + Multi-path Necessary

- Single-path forwarding cannot achieve strong availability guarantees
  - During routing protocol convergence, no path may be available
  - Equipment failure on path will result in unavailability until routing protocol updates and forwarding tables are adjusted
  - If forwarding path experiences high packet loss, then path is not usable for practical applications
- Approaches
  - Stable forwarding: packet-carried forwarding state protects forwarding from routing instabilities
  - Multi-path ensures presence of several paths, so as long as a single path works, end-to-end connectivity is assured

Bottleneck Routing Disrupts Availability

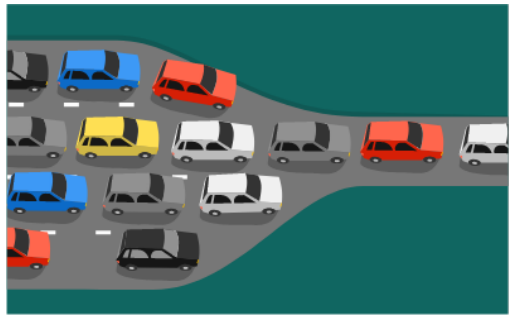

- Routing protocol switches route traversing a link with limited capacity (= bottleneck link)
- Bottleneck link traversal results in high packet loss
- Applications cannot operate and lose connectivity
- Since connectivity exists, often manual intervention needed to switch back to alternate path, outage typically persists for 30+ minutes
- Frequent reason for outage, caused by misconfiguration or attack

Cloudflare DNS goes down, taking a large piece of the internet with it

Devin Coldewey @techcrunch / 11:50 pm CEST • July 17, 2020 Comment

For two hours, a large chunk of European mobile traffic was rerouted through China

It was China Telecom, again. The same ISP accused last year of "hijacking the vital internet backbone of western countries."

By Catalin Cimpanu for Zero Day | June 7, 2019 -- 19:41 GMT (20:41 BST) | Topic: Security

Announcement of Failed Routes

- In some cases, networks continue to announce routes that failed
- Example: August 30 CenturyLink/Level(3) Outage
  https://blog.cloudflare.com/analysis-of-todays-centurylink-level-3-outage
  "CenturyLink/Level(3)'s network was not honoring route withdrawals and continued to advertise routes to networks like Cloudflare's even after they'd been withdrawn"

Multi-path Routing Provides Alternate Paths

- Important observation: Even secure routing protocol cannot prevent these issues, as announcements are often legitimate
- A multi-path routing approach provides alternate paths that can be used by end hosts

Global Communication Guarantees in the Presence of Adversaries

- Goal: If (routing policy compliant) path of benign ASes exists (with operational infrastructure), sender can find, use, and achieve minimum bandwidth guarantee on that path
- Challenges
  - Network routing instabilities, misconfigurations, etc.
  - DoS attacks at various levels (control plane, data plane, end host)

Availability in a public Internet is threatened by different types of DoS attacks

Link-flooding attacks	Attacker floods network links with excessive amount of traffic
	Can target **access links** (last mile) or **core links** in the network
	Often executed using botnets and/or amplification techniques
End-system attacks	Attacker exhausts **computational** or **memory resources** of victim
	Often possible due to other defense mechanisms such as firewalls
	Examples: **state exhaustion**, **signature flooding**
Control-plane attacks	Attacker disrupts important **control-plane mechanisms** or **access to services**
	Services are essential for a functioning network
	Examples in SCION: beacon server, path server, certificate server

SCION is an Internet architecture with both *strong security* properties and *high availability*

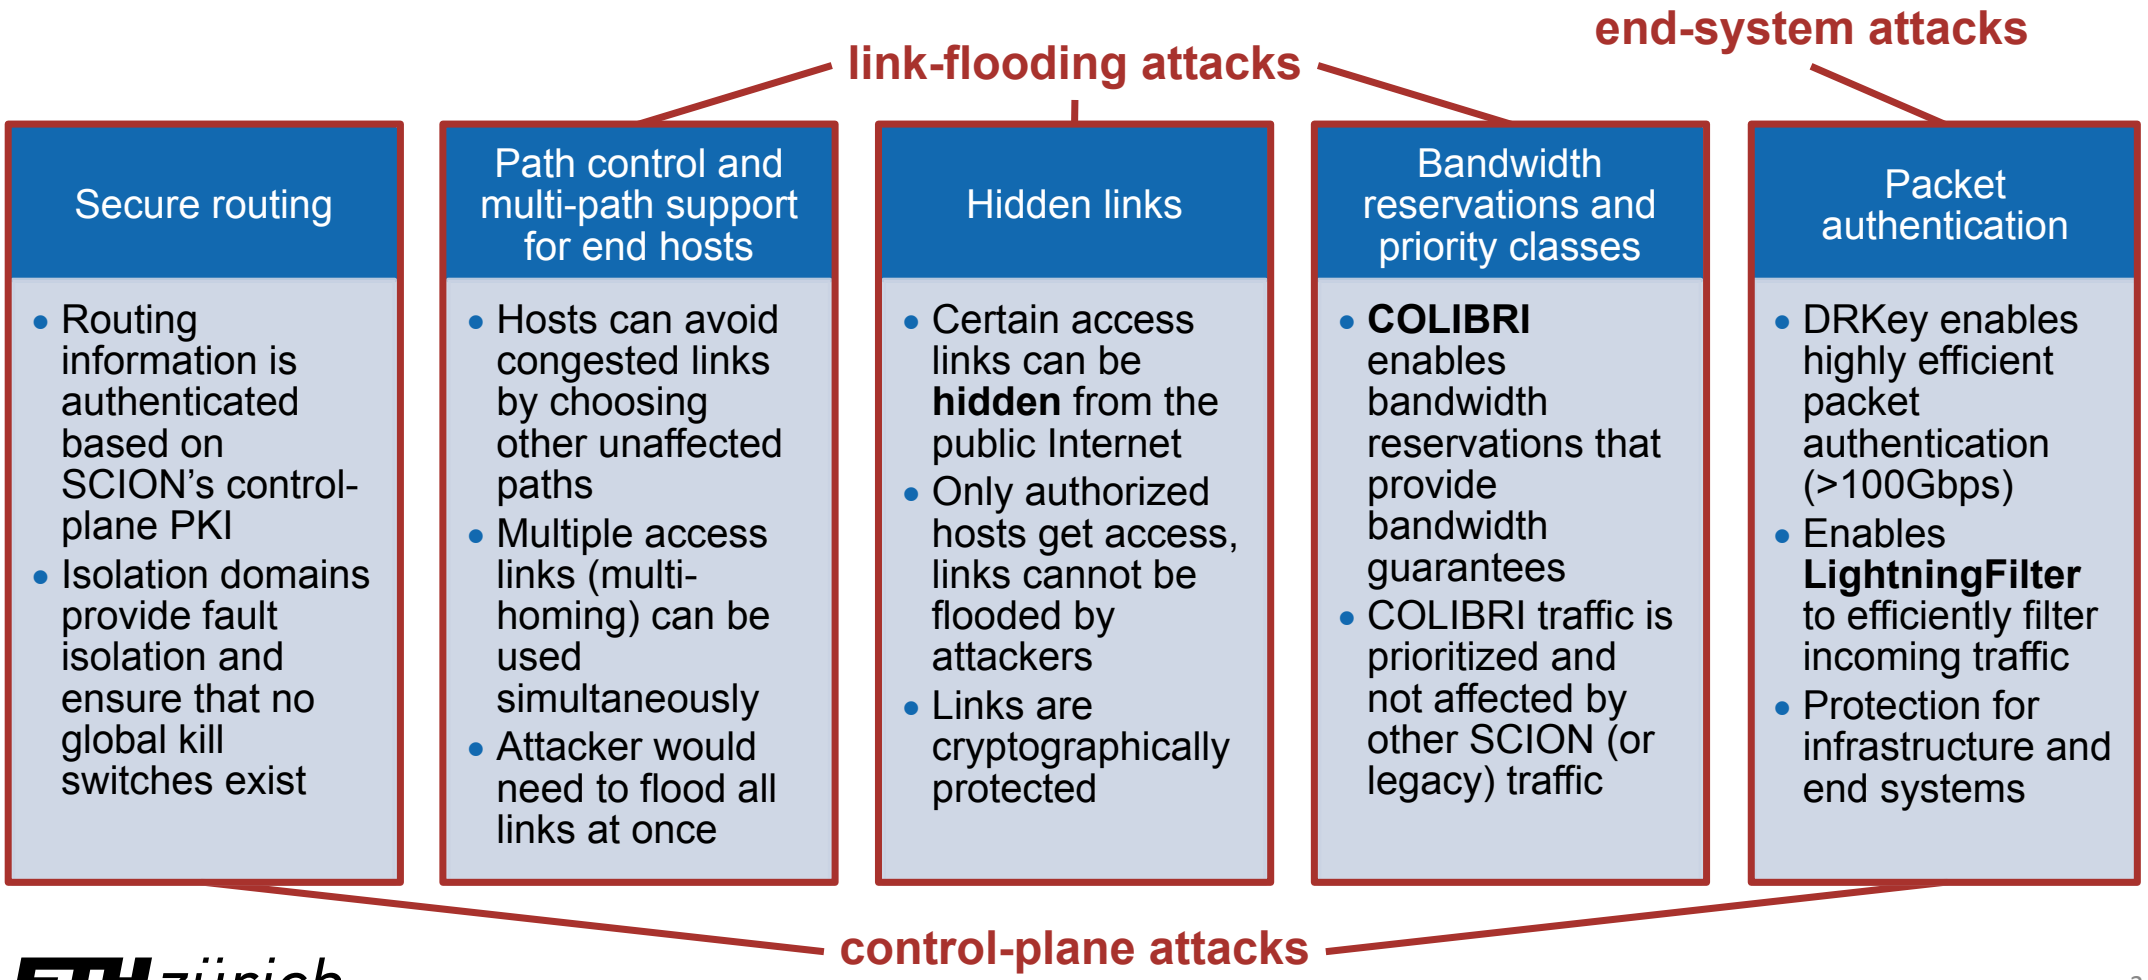

ETH zürich

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 mechanisms
- PISKES: Pragmatic Internet-Scale Key-Establishment System, Rothenberger et al., ACM Asia Conference on Computer and Communications Security (ASIACCS) 2020

Avoid Asymmetric Crypto for High Performance

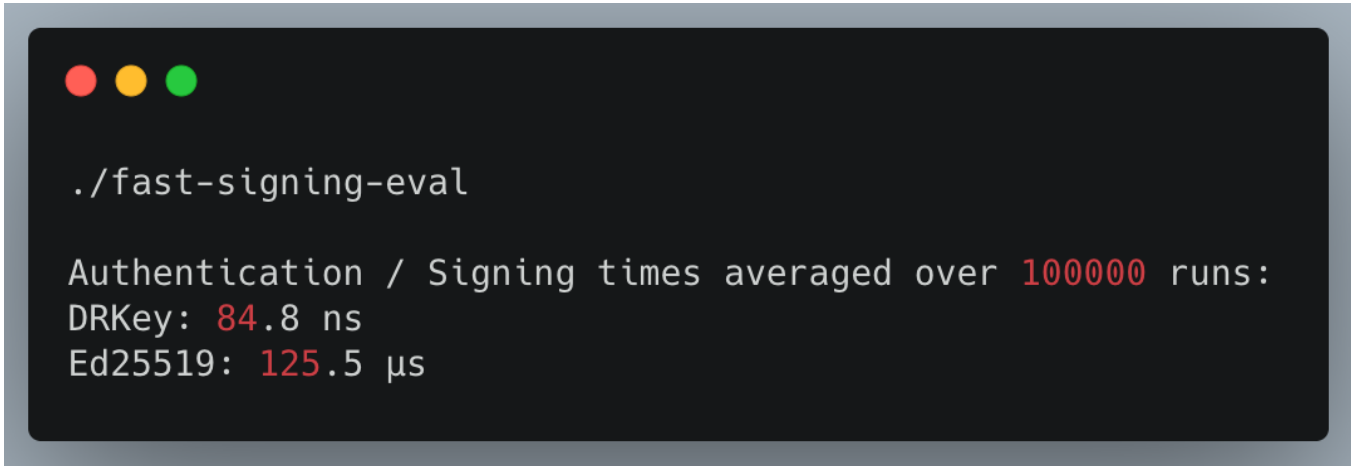

Factor:
~ 1450x

ETH zürich

SCION

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 $K_{X \rightarrow *}$

Coremelt Attack [Studer, Perrig, Esorics 2009]

- Adversary controls many bots distributed across the Internet
- Bots send traffic between each other, thus all traffic is desired by destination
  - Traffic is not sent to victim as in regular DDoS attacks
- Adversary can exhaust bandwidth on victim link
- Result: attack traffic exhausts bandwidth in per-flow fair sharing systems

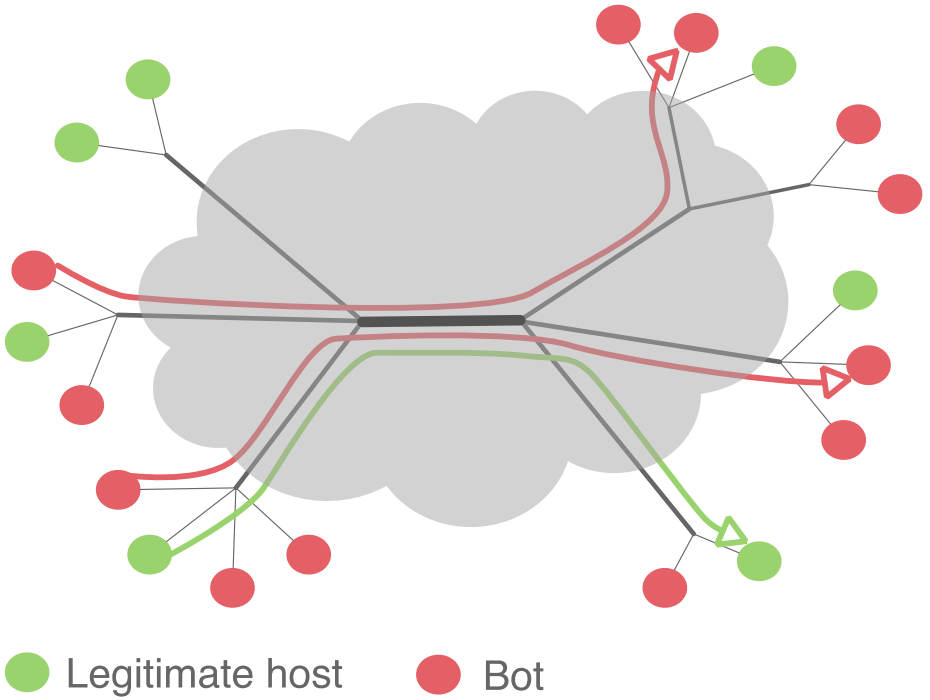

Crossfire Attack [Kang, Lee, Gligor, IEEE S&P 2013]

- Adversary controls distributed bot army
- Observation: due to route optimization, few links are actually used to connect a target region to rest of Internet
- Adversary can contact selected servers to overload target links
- Result: disconnect target region from remainder of Internet

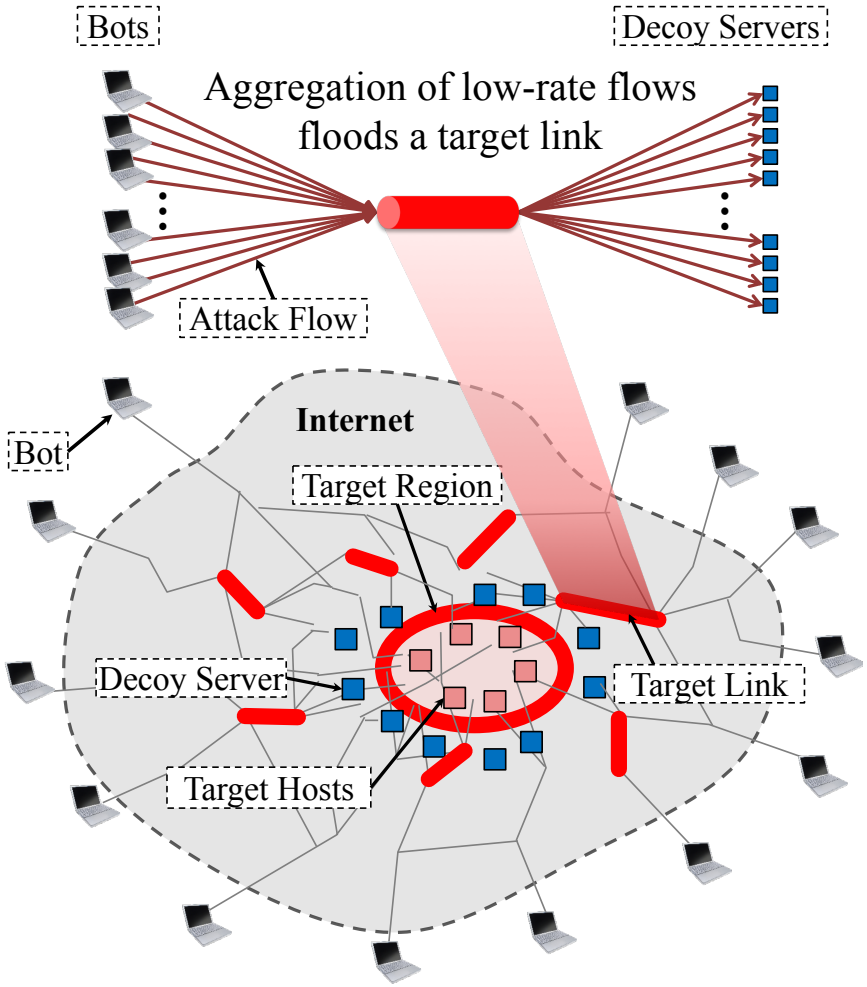

ETH zürich

SCION

Volumetric DDoS Attacks

- Attacker overloads network link to induce congestion
- Defense requires sophisticated approaches
  - EPIC dynamic hop field computation
  - COLIBRI global resource allocation and reservation

EPIC: Every Packet Is Checked

- Goals
  - Per-packet source authentication by every router and destination
  - Per-packet-unique hop fields
  - Path validation by destination
- Assumption: global time synchronization (± 100ms)
- Attacks prevented
  - Malicious router replays packets or increases packet size
  - Hop field MAC is brute forced and destination attacked until expiration time
- EPIC: Every Packet Is Checked in the Data Plane of a Path-Aware Internet, Legner et al., USENIX Security Symposium 2020

ETH zürich SCION

EPIC Level 1

- Current SCION Hop field MAC: $MAC_{Ki}(TS \| IgIF \| EgIF \| ExpT \| \sigma_{i-1})$
  - TS: time beacon creation, ExpT: Expiration time, IgIF/EgIF: Ingress/ Egress Interface, T: packet creation time, σ_{i-1} : previous hop field MAC
- One additional level of indirection:
  $K_H = MAC_{Ki}(TS \| IgIF \| EgIF \| ExpT \| S_{i-1})$
  $S_i = H(K_H)$
  Hop field MAC: $MAC_{K_H} (T \| H(P) \| len(P))$
- Hop key K_H is distributed via beacon (PCB)
  and S_i is included in packet
- Result: every packet has unique hop field MAC, foiling brute force attacks on hop field MAC

ETH zürich SCION

EPIC Level 2

- Goal: **line-speed source authentication for every packet on every router**
- Approach: include DRKey $K_{X \rightarrow Y:H}$ in hop field MAC
- $K_H = MAC_{Ki}(TS \| IgIF \| EgIF \| ExpT \| S_{i-1})$

  $S_i = H(K_H)$

  Hop field MAC: $MAC_{K_{X \rightarrow E:e}} (T \| H(P) \| len(P) \| K_H)$

  For host e in AS E, traversing AS X
- Router in AS X can efficiently derive $K_{X \rightarrow E:e}$ (2 AES operations)
- Host e needs to fetch one key per AS traversed from local certificate server
- Result: efficient per-packet per-hop source authentication!
  (5 AES op)

COLIBRI: Global QoS System

Designed to scale to the Internet

Admission Control

- Source Authentication
  - ✓ Eliminates free-riders
  - ✓ Basis to achieve fairness
  - Per-packet authentication based on DRKey

Resource Allocation

- Two-tier reservation
  - Tier-1: between ASes
  - Tier-2: between end hosts
- Fairness
  - Per-neighbor AS fairness

Traffic Policing

- Designed for scalability
- Two layer model
  - @ Edge: host-based traffic shaping
  - @ Core: neighbor-AS-based traffic shaping

Traffic Monitoring

- Designed for scalability
- Two layer model
  - @ Edge: stateful per-flow monitoring
  - @ Core: per-flow-stateless probabilistic monitoring

Admission Algorithm with Per-Neighbor Fairness

- Each AS defines neighbor-to-neighbor minimum bandwidth guarantees
- For any path, AS-to-AS minimum bandwidth guarantee can be computed, regardless of other demands
- Algorithm guarantees that no set of ASes can reserve a disproportionate amount of bandwidth through any link

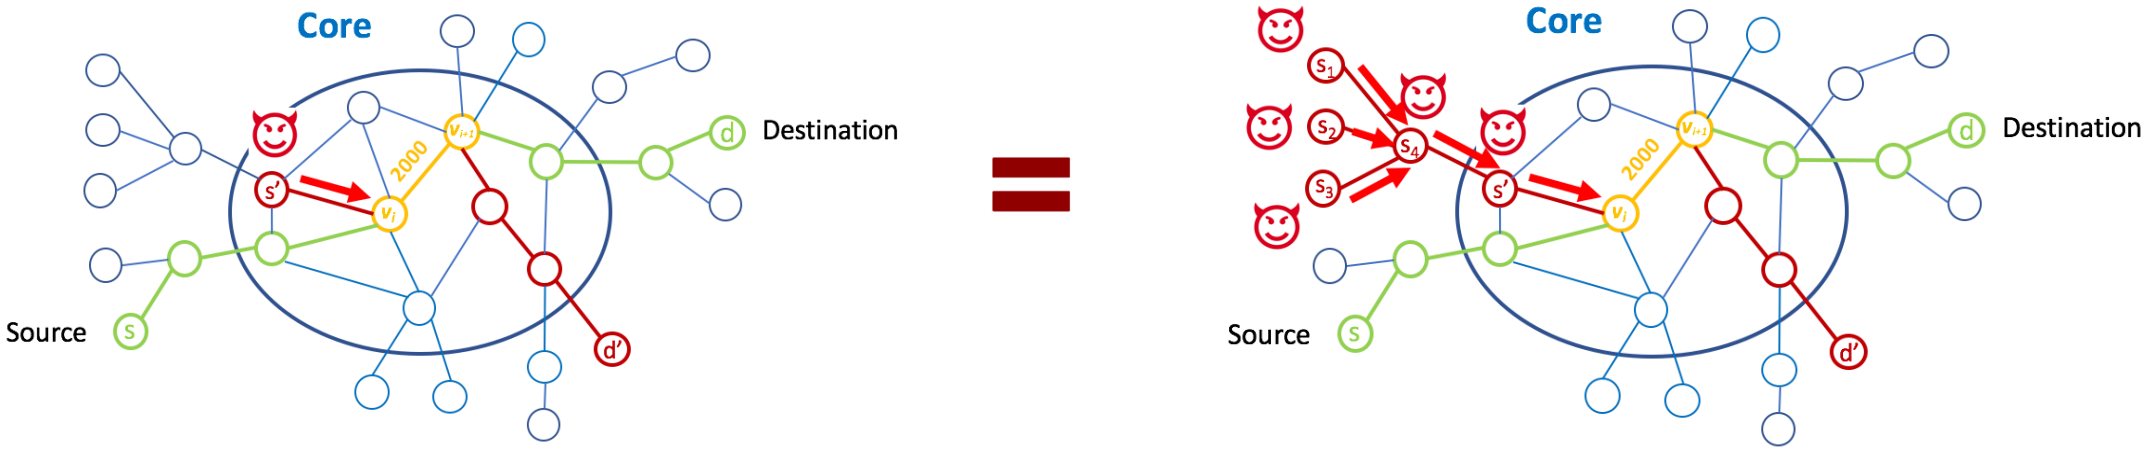

ETH zürich

SCION

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

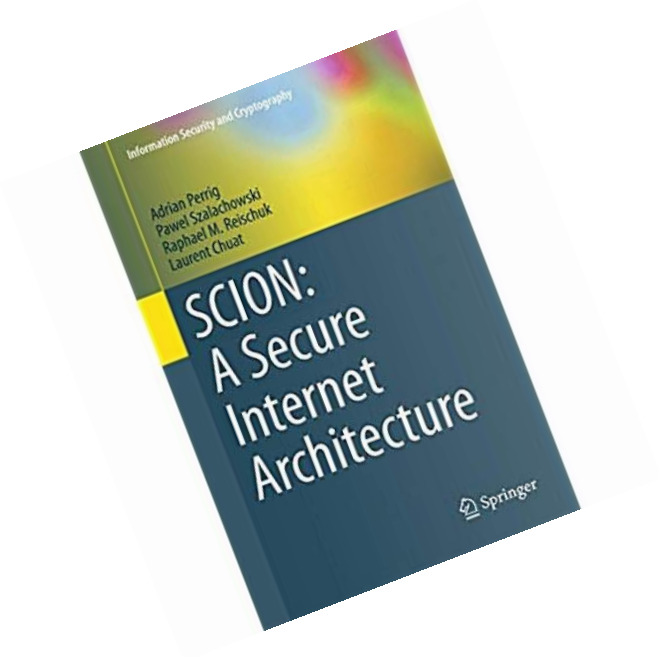

Lessons Learnt

- A global high-security public Internet is possible
  - Sovereign operation, yet global connectivity is feasible
  - Global available communication is possible on public networks
  - Guaranteed DDoS resilience is possible
  - Protocol and code verification are necessary to obtain strong properties for large-scale distributed systems
- Static paths + multi-path routing enables powerful concepts
  - Fast failover: do not rely on network-based active failover but on redundancy with simultaneous use
  - Possibility to unlock additional network capacity

SCION Team (2019)