SCION
SCALABILITY, CONTROL, AND ISOLATION
ON NEXT-GENERATION NETWORKS

SCION: overview and current deployment

Adrian Perrig
ETH & Anapaya Systems

ETH zürich

SCION

Motivations for studying inter-domain routing

- Original research question: How secure can a global inter-domain routing system be?
- Security-centric design: secure all aspects, including control message protocol
- Use formal verification to ensure security
- Scalability and rapid routing convergence possible?

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, SR, BGP Multipath
- Inter-domain routing proposals: ChoiceNet, HLP, HAIR, RBF, AIP, POMO, ANA, ...
- Intra-domain / datacenter protocols: SDN, HALO, ...

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

Insight: Formal Security Verification Necessary

- To achieve strong assurance for a large-scale distributed system, formal security verification is necessary
- Performing formal verification from the beginning avoids “difficult-to-verify” components
  - Many design aspects of SCION facilitate formal verification
- Collaboration with David Basin’s and Peter Müller’s teams in the VerifiedSCION project

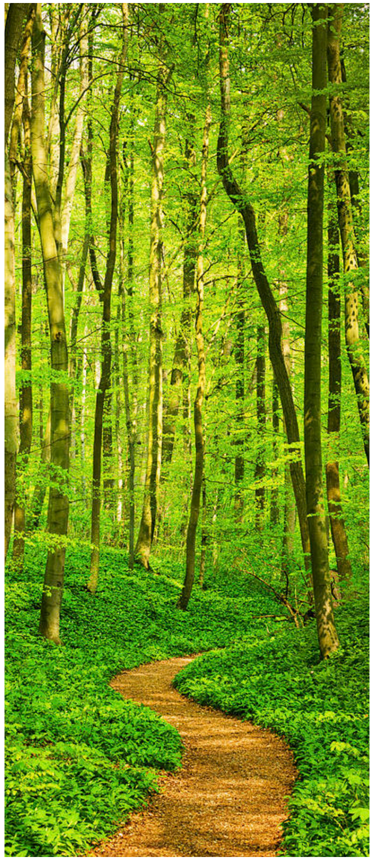

VerifiedSCION

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

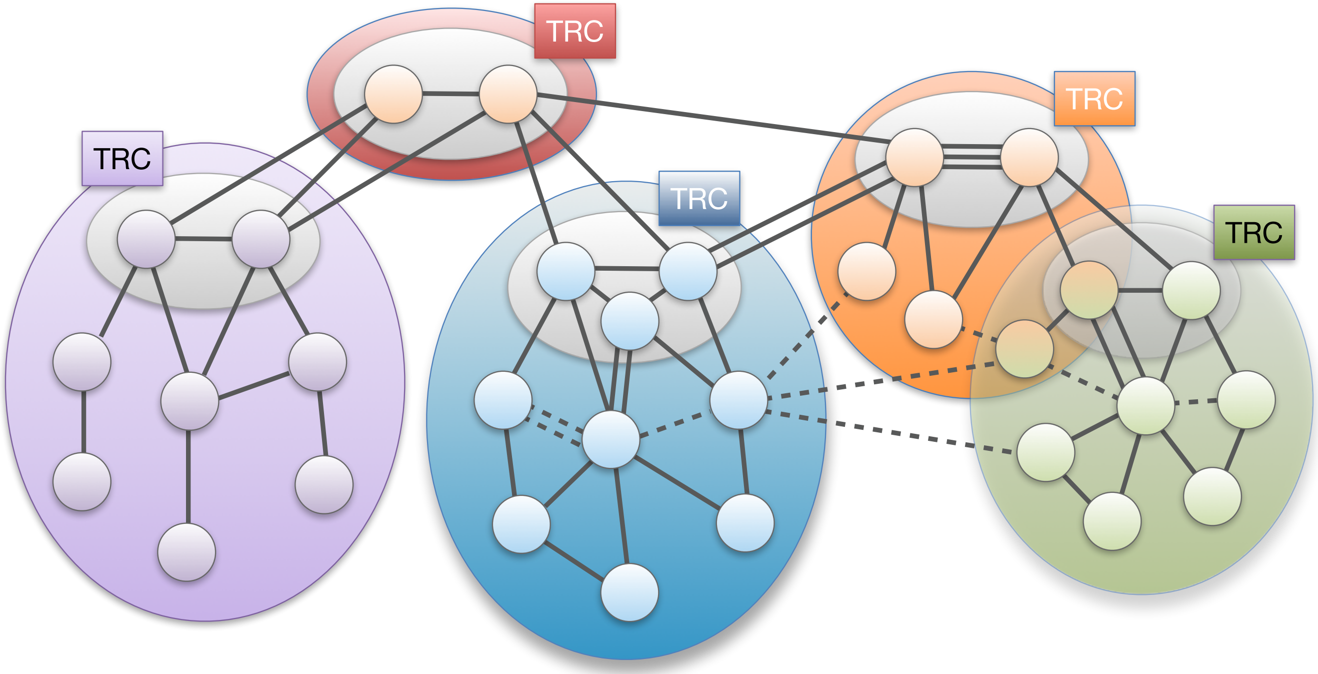

ETH zürich

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

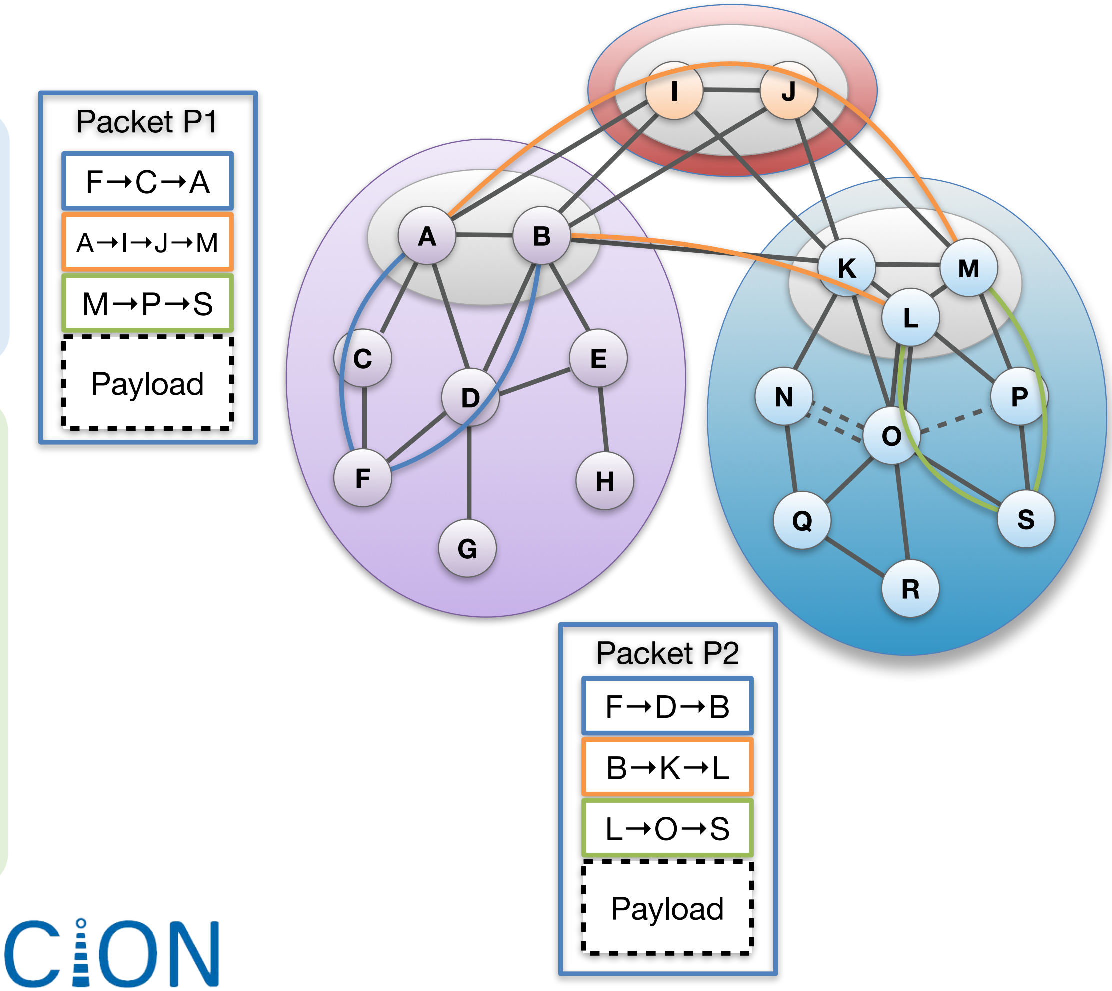

Packet P1

F→C→A

A→I→J→M

M→P→S

Payload

Packet P2

F→D→B

B→K→L

L→O→S

Payload

SCION

SCION Control and Data Plane

- Three main functions of the control plane
  1. Path exploration → path segments
  2. Path dissemination → senders requests segments
  3. Certificate dissemination/renewal
     → needed for segment verification
- Path segments contain forwarding and meta information. Meta information can include geographical location of routers, MTU, bandwidth, link latency…
- Senders extract the forwarding information from the path segments to form complete end-to-end paths
- Forwarding information is encoded in the packet header. Routers only verify the authenticity of the information
  → two AES operations replace longest-prefix match

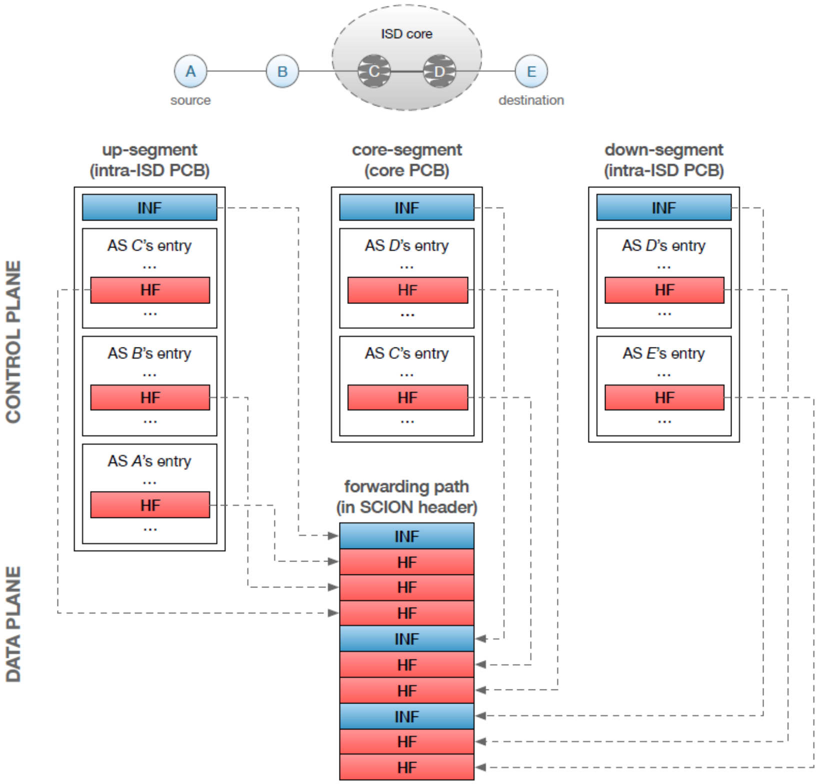

Intra-ISD Path Exploration: Beaconing

- 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

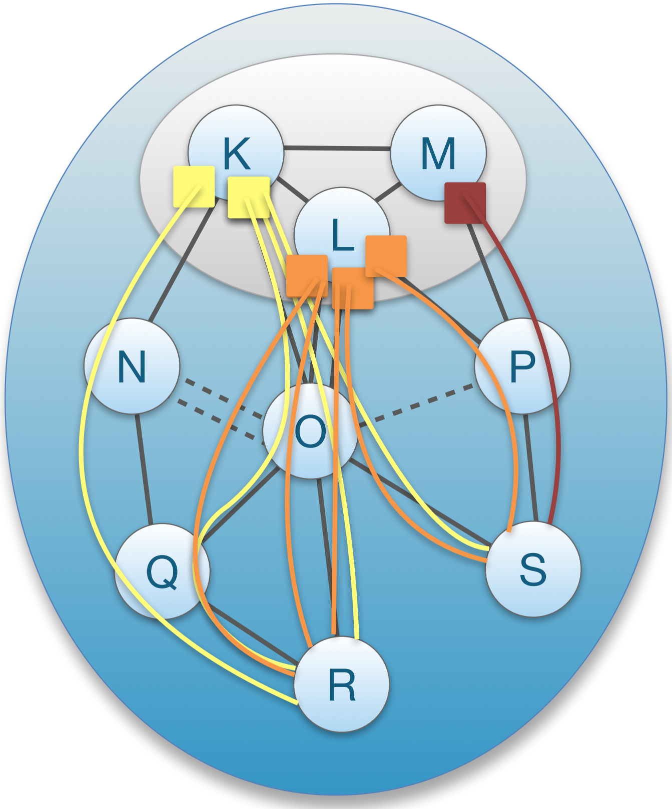

Up-Path Segment Registration

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

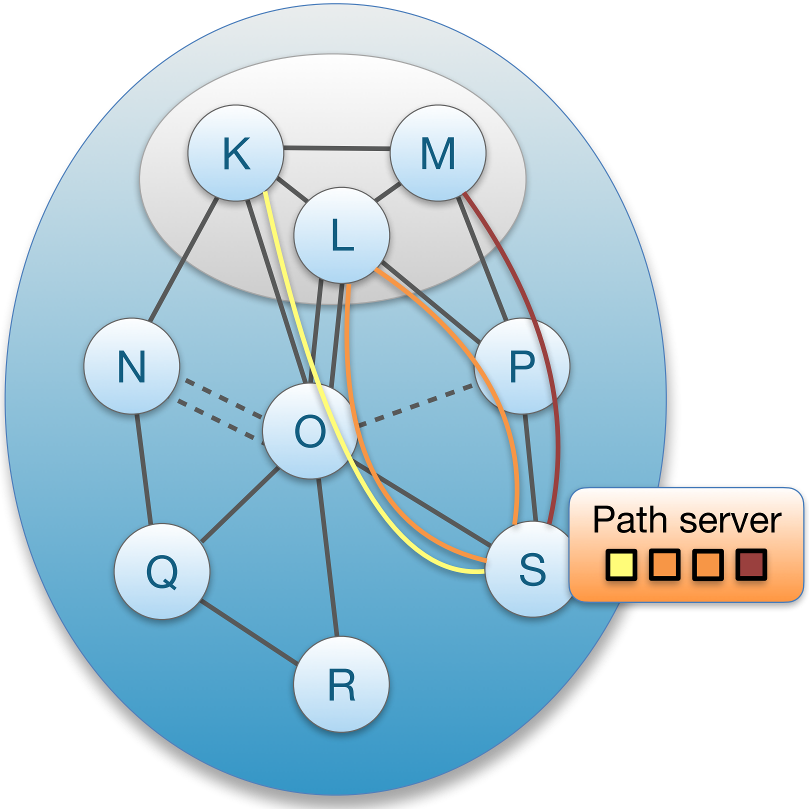

Down-Path Segment Registration

- 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

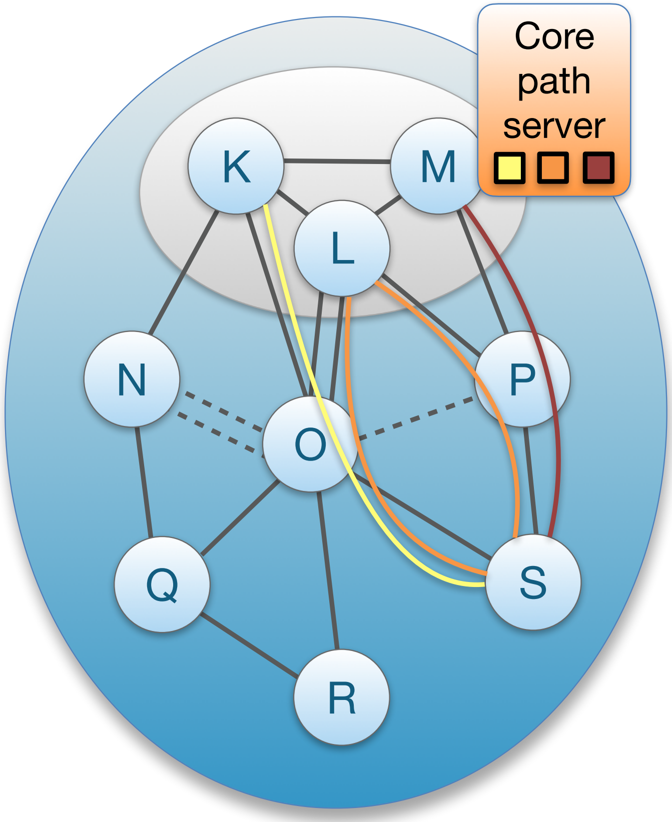

Communication within ISD

- 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)

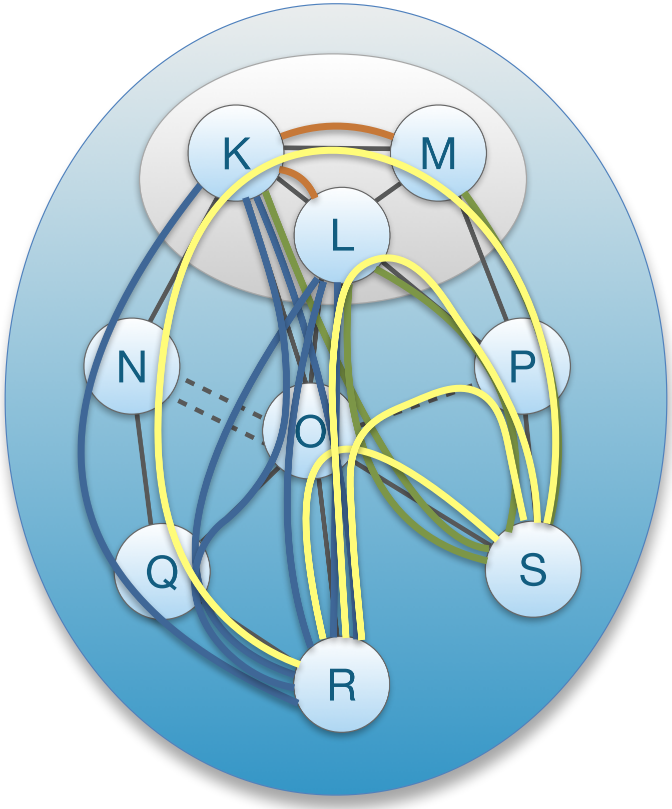

Communication to Remote ISD

- 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

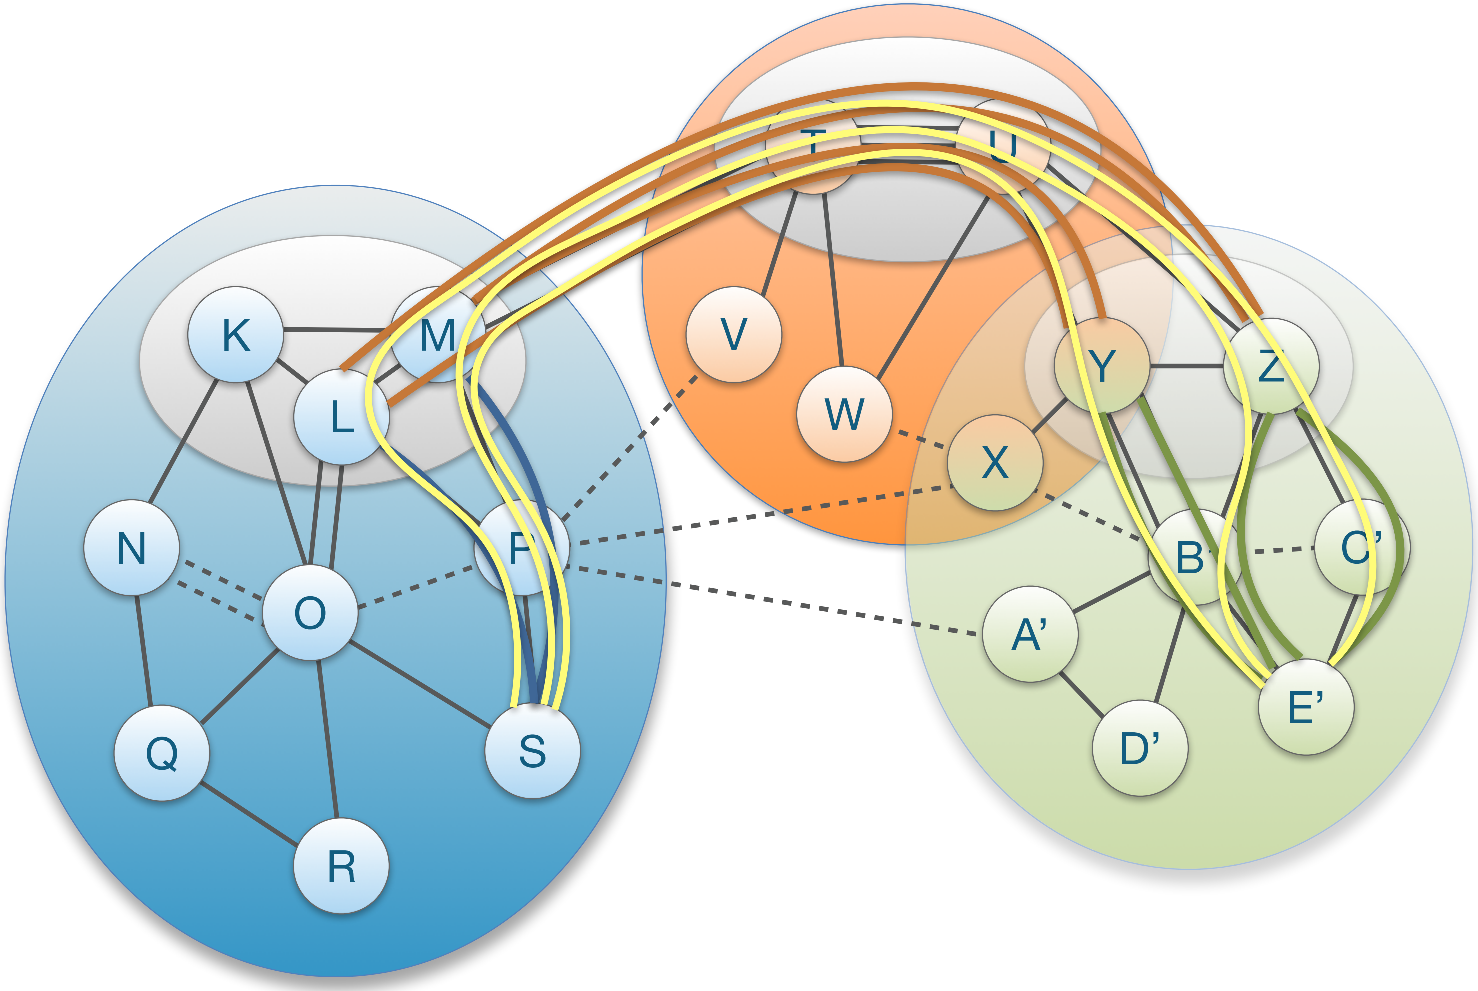

SCION Drawbacks

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

How to Deploy SCION: ISP

- CORE Routers are set up at the borders of an ISP
  - to peer with other SCION-enabled networks
  - to collect customer accesses
- No change to the internal network infrastructure of an ISP needed!

the slides, I left a few comments instead.

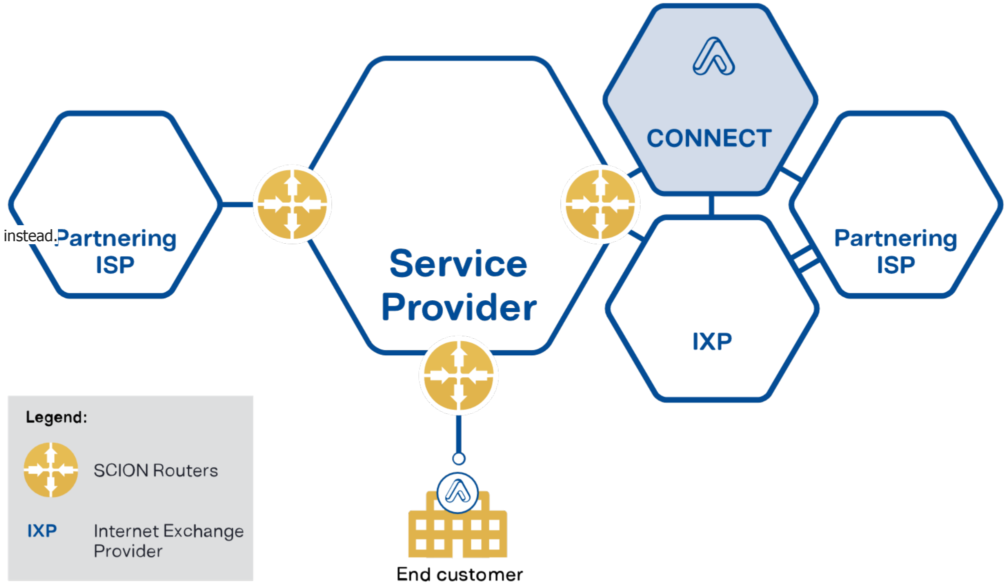

How to Deploy SCION: End Customer

- SCION IP Gateway (SIG) enables seamless integration of SCION capabilities in end-domain networks
- No upgrades of end hosts or applications needed

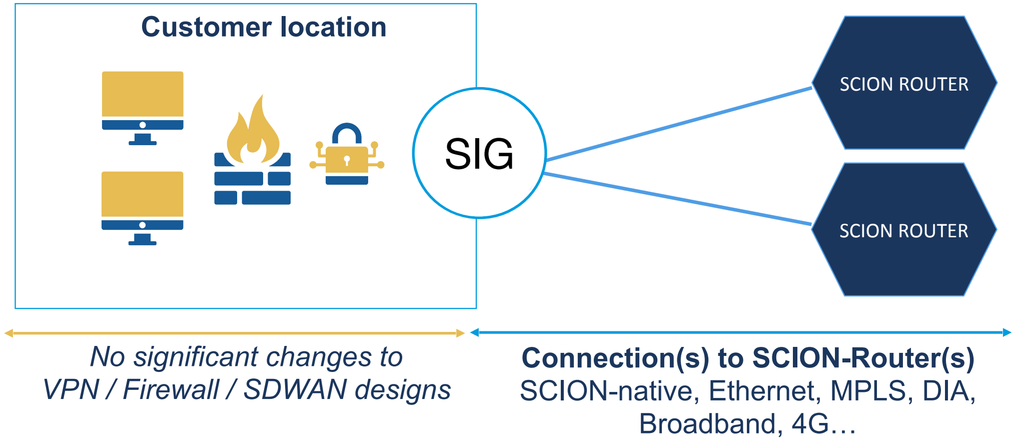

No significant changes to VPN / Firewall / SDWAN designs

Connection(s) to SCION-Router(s)
SCION-native, Ethernet, MPLS, DIA, Broadband, 4G…

Current deployments

- **Global production network (Led by Anapaya Systems)**
  - No dependence on BGP protocol
- **Three implementations**
  - Open source in Go: https://github.com/scionproto/scion
  - Vendor proprietary (high-performance) by Anapaya
  - P4 (experimental) by SIDN Labs
- Current deployment
  - ISPs: Swisscom, Sunrise, SWITCH, Telindus, CyberLink, InterCloud, …
  - IXPs: SwissIX offers SCION peering, + others joining
  - Bank deployment: Secure Swiss Finance Network

Next steps - Standardization

Success factors (RFC5218)

- Meet real need ✅
- Incremental Deployability ✅
- Open Code Availability ✅
- Freedom from Usage Restrictions ✅
- Extensible, scalable ✅
- Threats mitigated ✅
- Open Specification Availability ⏳ —> Standardization needed
- Note: side meeting at IETF 113: https://notes.ietf.org/s/LaApgxo2b

Online Resources

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

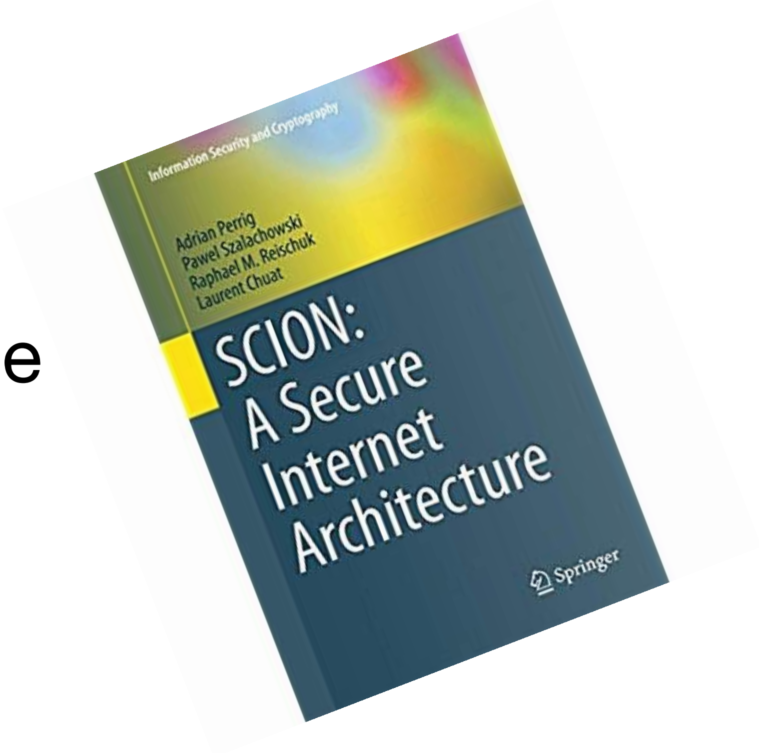

Information Security and Cryptography

Adrian Perrig
Pawel Szalachowski
Raphael M. Reischuk
Laurent Chuat

SCION: A Secure Internet Architecture

Springer

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 may not be usable by 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

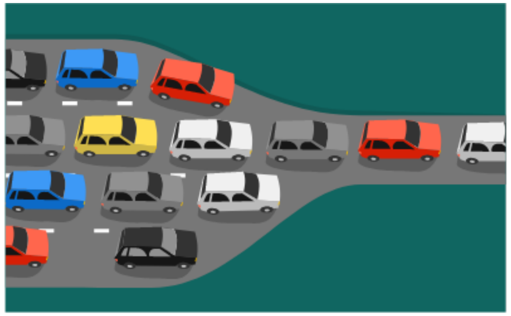

- Routing protocol switches route traversing a link with limited capacity (= bottleneck link)
- Bottleneck link traversal results in high packet loss
- Applications cannot operate and lose connectivity
- Since connectivity exists, often manual intervention needed to switch back to alternate path, outage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

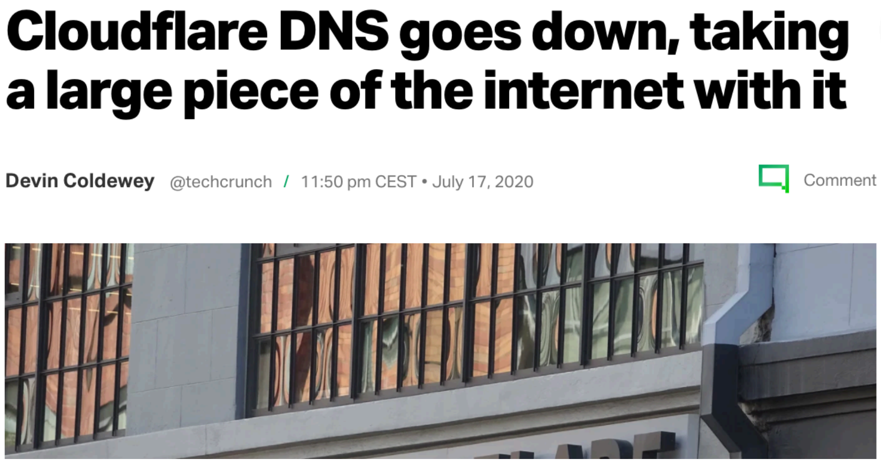

Cloudflare DNS goes down, taking a large piece of the internet with it

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

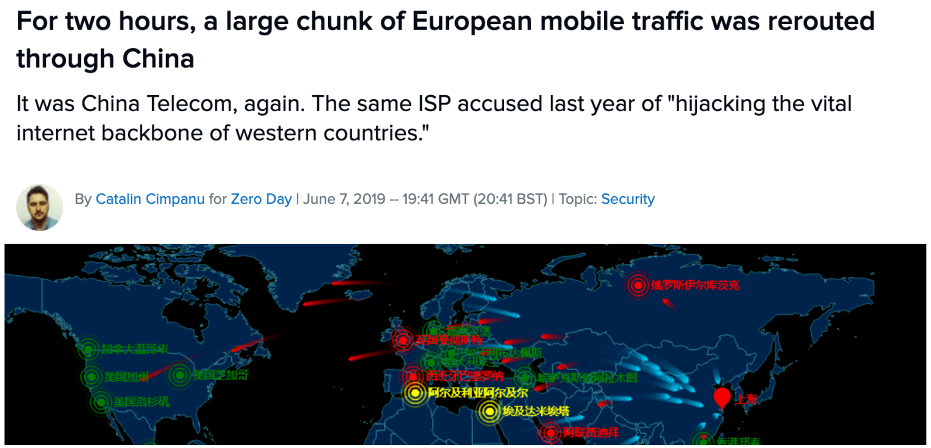

For two hours, a large chunk of European mobile traffic was rerouted through China

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”

Insight: Secure Routing Insufficient

- Secure single-path routing protocol cannot prevent outages caused by bottleneck link or continuing announcement of failed or congested routes

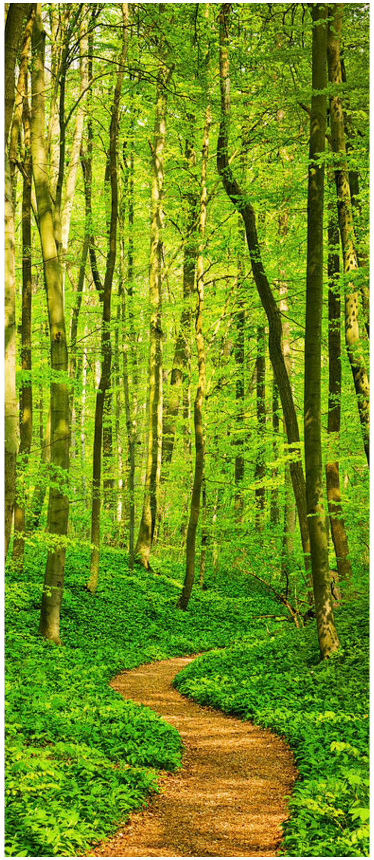

Summary

- SCION connectivity available in production from several ISPs
- 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 possible on routers
  - Formal verification of protocols and code
  - Immune against routing attacks, e.g., prefix hijacking

SCION Team