ETH zürich
Network Security Group

OTTO VON GUERICKE UNIVERSITÄT MAGDEBURG

ANAPAYA

Deployment and Scalability of an Inter-Domain Multi-Path Routing Infrastructure

Cyrill Krähenbühl[*], Seyedali Tabaeiaghdaei[*], Christelle Gloor[*], Jonghoon Kwon[*], Adrian Perrig[*‡], David Hausheer[†], Dominik Roos[‡]
[*]ETH Zürich, [†]OVGU Magdeburg, [‡]Anapaya Systems

CoNEXT 2021

A New Era for Internet Technologies

- AR / VR with haptic feedback
- Interactive telepresence
- 3D hologram imaging
- Immersive games
- …

Image source: Siemens, A new era for the Internet of Things

WHO BOUGHT THE FIRST FAX MACHINE?

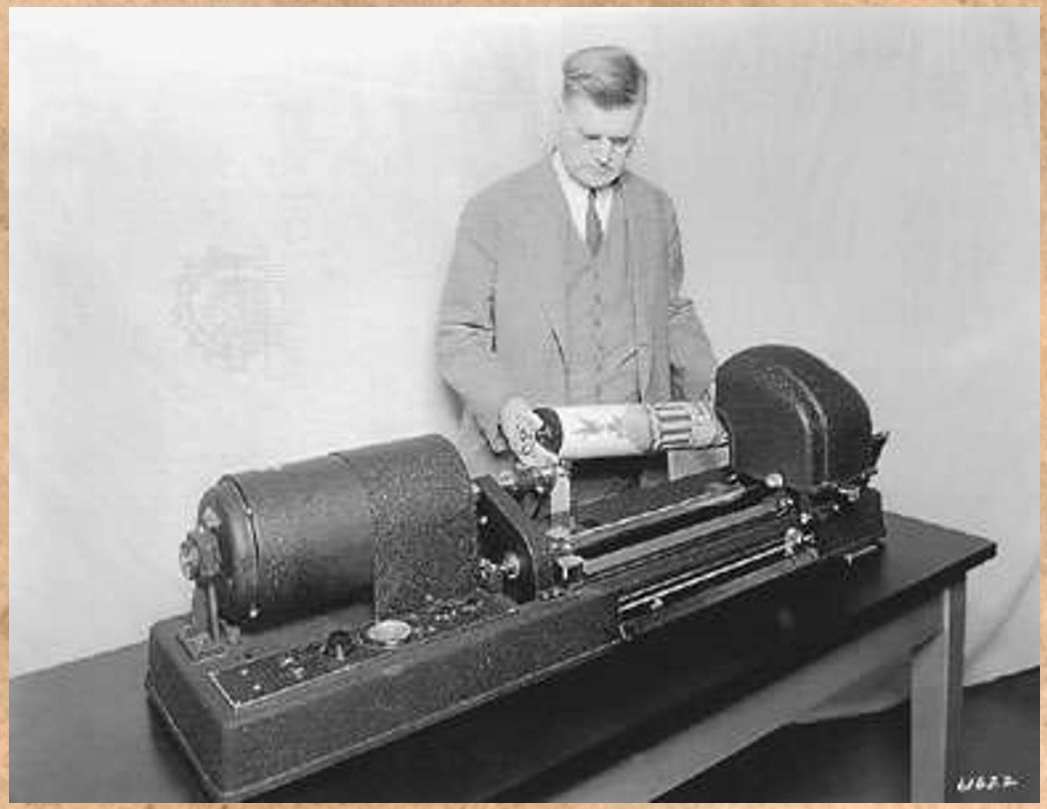

Difficulties in Deployment

New technologies lacking early adopters' incentives …

Egress filtering

BGPsec

DNSsec

… experience limited deployment in real-world

In this presentation...

We explore a path towards a next-generation Internet

SCION: Next-Generation Internet

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

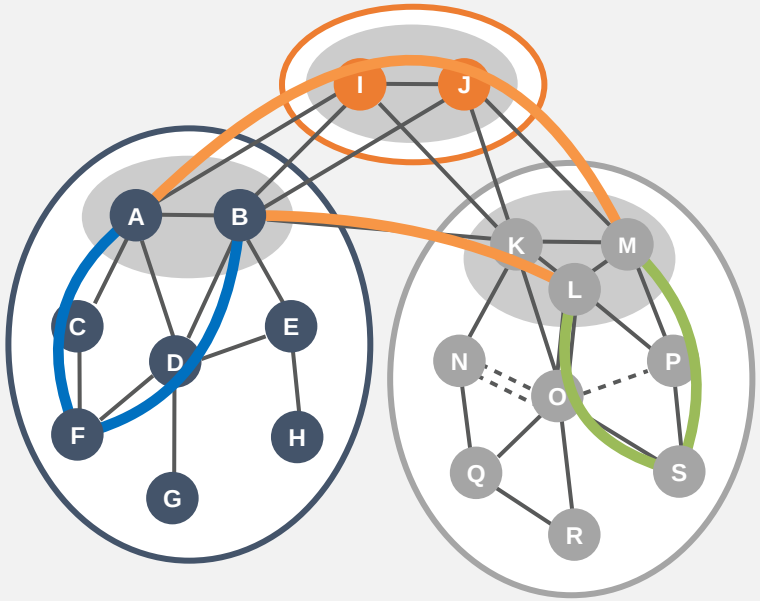

Packet P1

- F→C→A
- A→I→J→M
- M→P→S

Packet P2

- F→D→B
- B→K→L
- L→O→S

SCION: Next-Generation Internet

Path-aware Internet architecture which enables endpoint path control

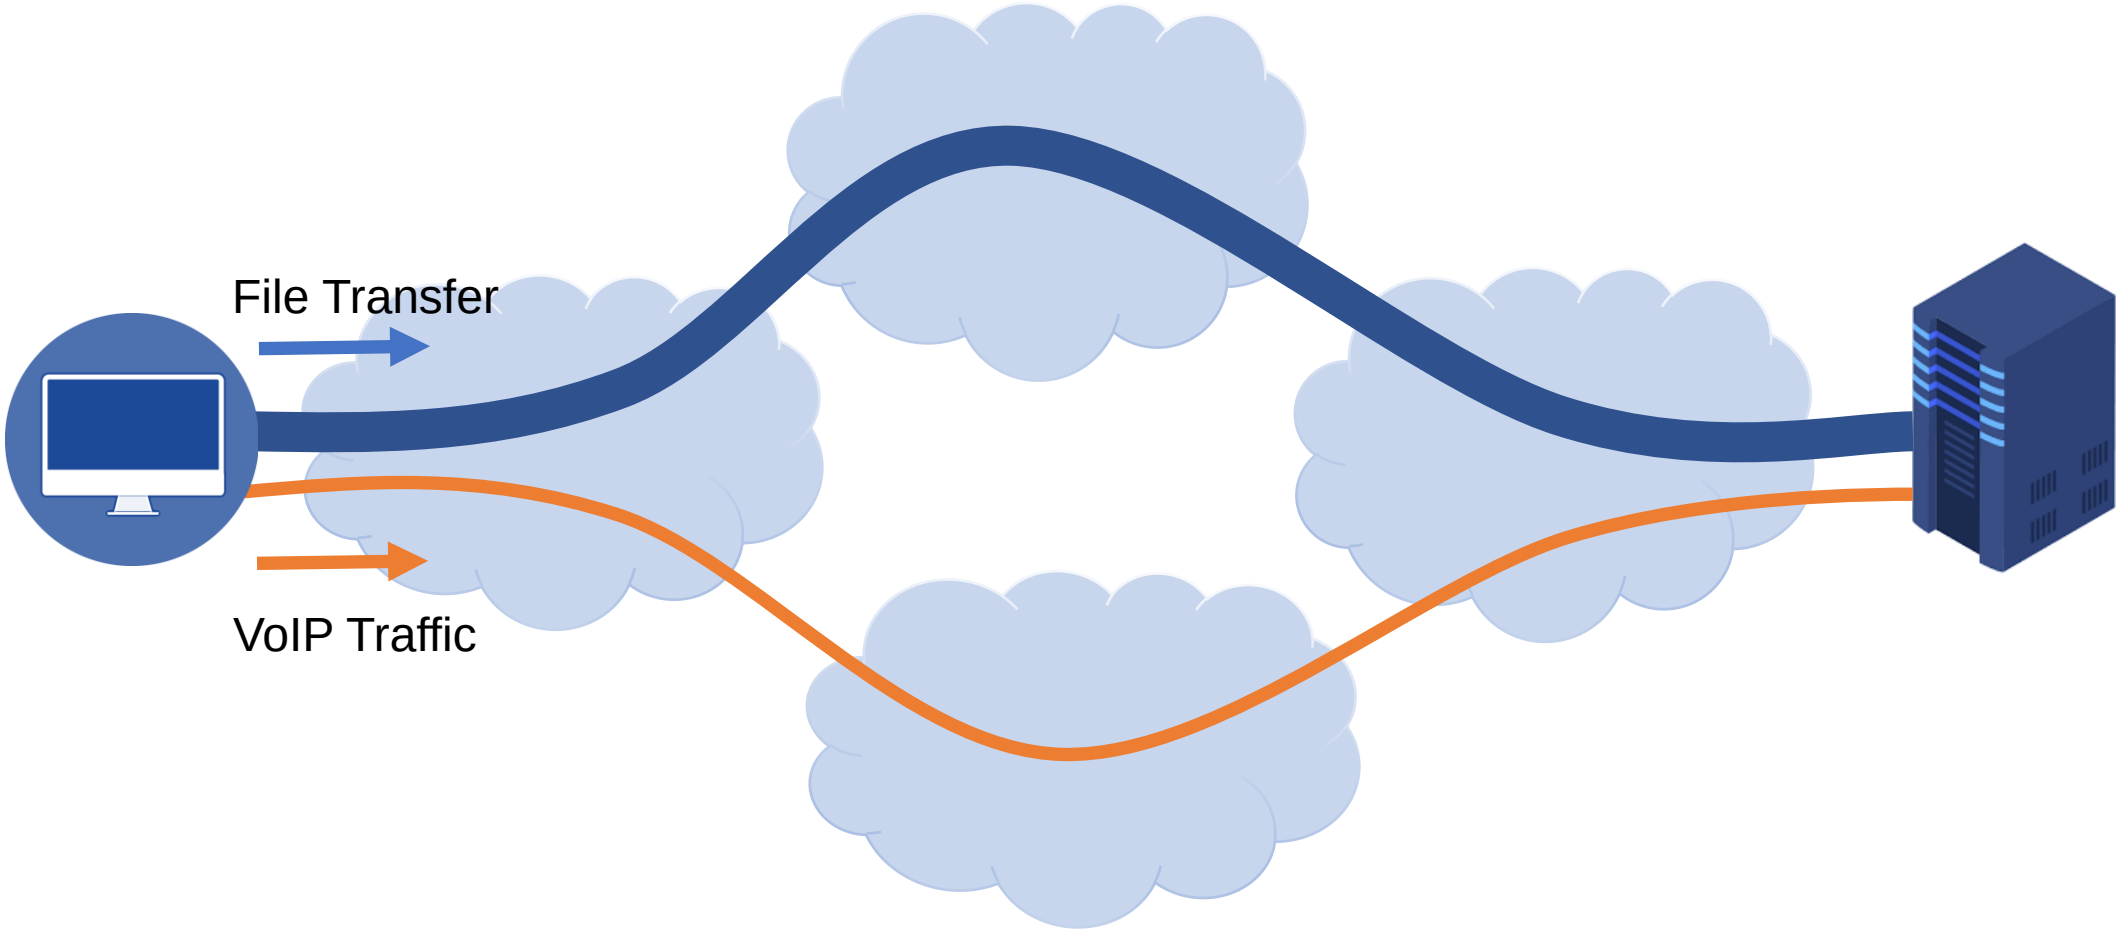

SCION: Next-Generation Internet

Provides in-network multi-path

→ even if single-homed

→ further benefits from multi-homing

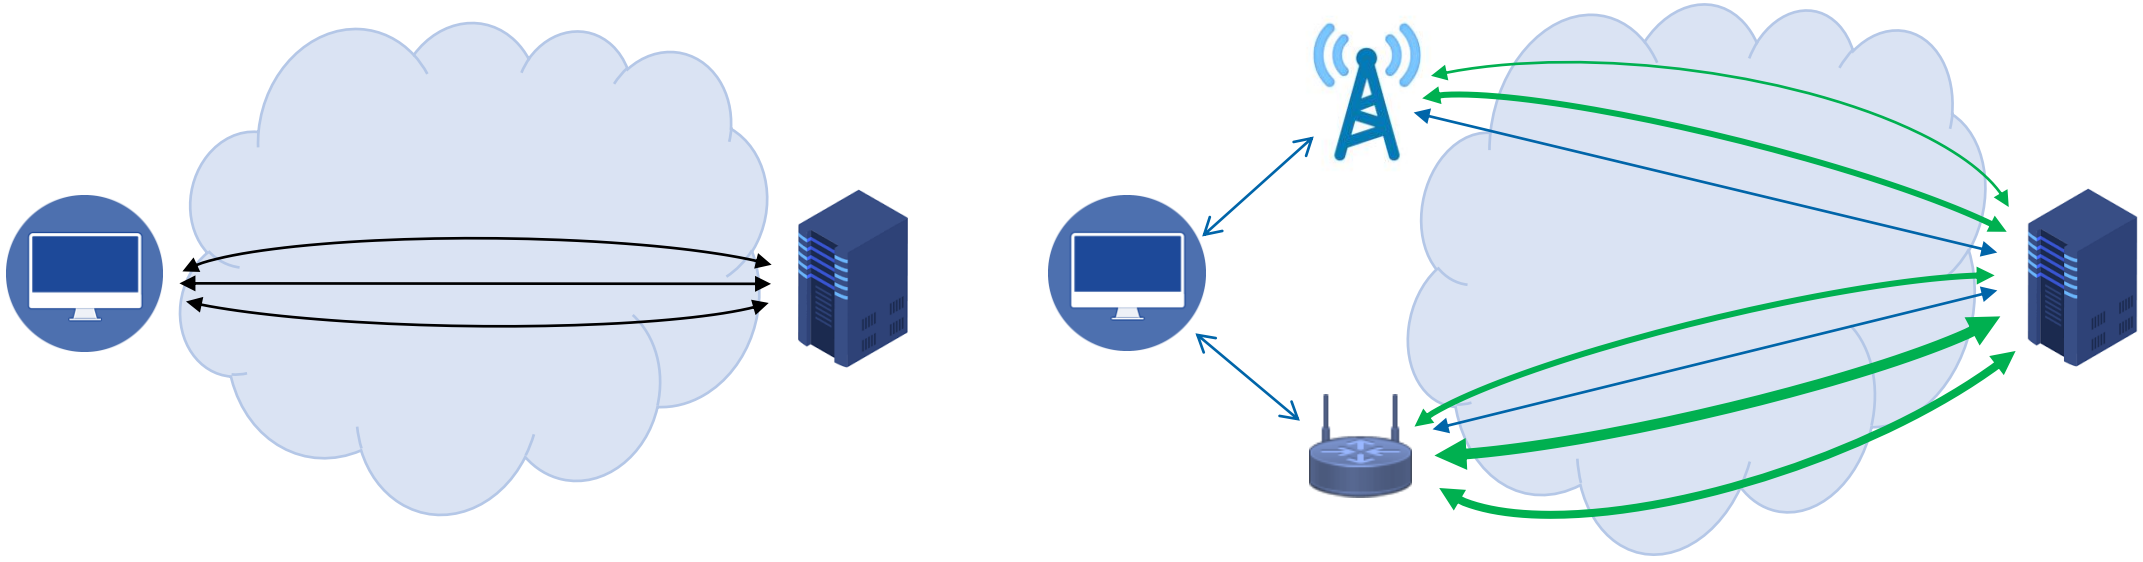

This seems impossible to deploy!

Need early adopters deploying the architecture

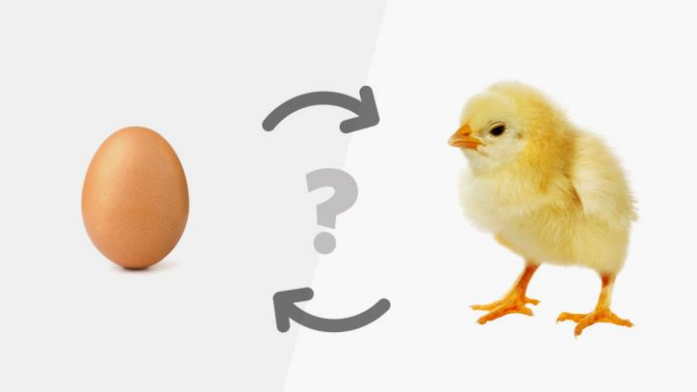

Need initial deployment to convince early adopters

Need long-term incentives for providers

- ✓ New opportunities to generate revenue
- ✓ Reduce operational network cost

ISPs Offer Official Products (2021)

A quantum leap in cybersecurity

The Swiss education, research and innovation community depends on secure data transfer. The new SCION architecture is now available to all universities via SWITCH's network.

Text: **Daniel Bertolo**, published on 15.07.2021

SWITCH is making secure data transfer with the SCION internet
universities. Developed at ETH Zurich, SCION offers universities
improved security, availability and performance.

swisscom Offer Insights and Success Stories

Business customers > Enterprise Customers > Offer > Network, Internet & Landline

SCION – transporting sensitive data securely

The secure, high-speed Internet under your control

Contact an expert

Sunrise | Medium & Large Business EN

SCION - the innovative solution for secure multi-site connectivity.

SCION stands for, **Scalability, Control and Isolation On next-generation Networks**. SCION is a new, state of the art technology developed in Switzerland's ETH Zurich, in collaboration with Anapaya Systems. The SCION protocol enables a fundamentally secure multi-site connectivity option, offering path control, fault isolation and explicit trust information for end-to-end communication.

↓ Go to Offer

SCION is federated

SCION is federated and offered as a product deployed by several ISPs

Any ISP can join SCION and benefit from its properties

SCION Deployment is overlay-free

Building on weak foundation can be problematic

No reliance on BGP (i.e., overlay-free) **in production network**

How did we get here?

Deployment Milestones

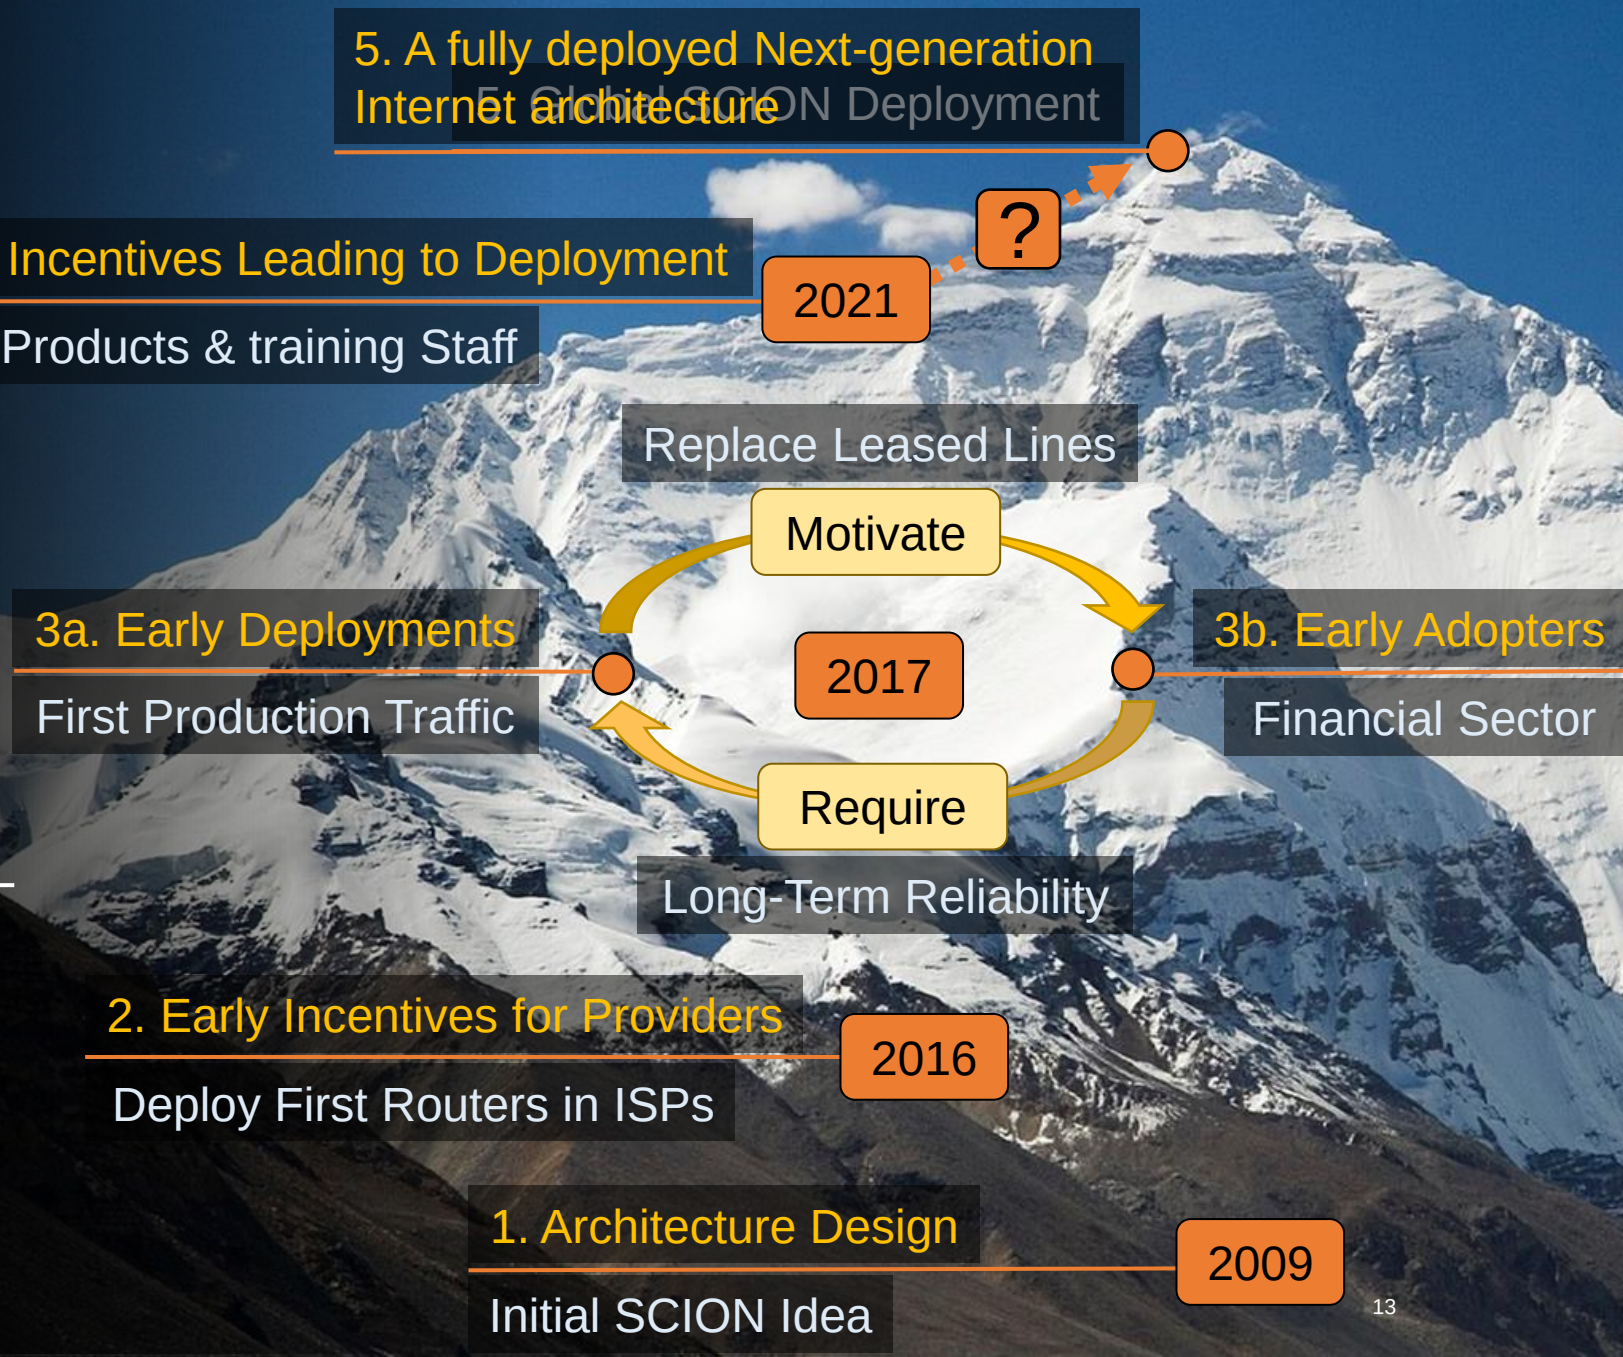

Global SCION Backbone Service

Anapaya Connect:

- ✓ SCION-transit service
- ✓ Native SCION connectivity
- ✓ BGP free routing
- ✓ Over 100 data centers
- ✓ 10+ countries

https://www.anapaya.net/anapaya-connect

End-Domain Deployment

Native deployment provides the biggest SCION benefits

SCION IP Gateway (SIG) simplifies deployment

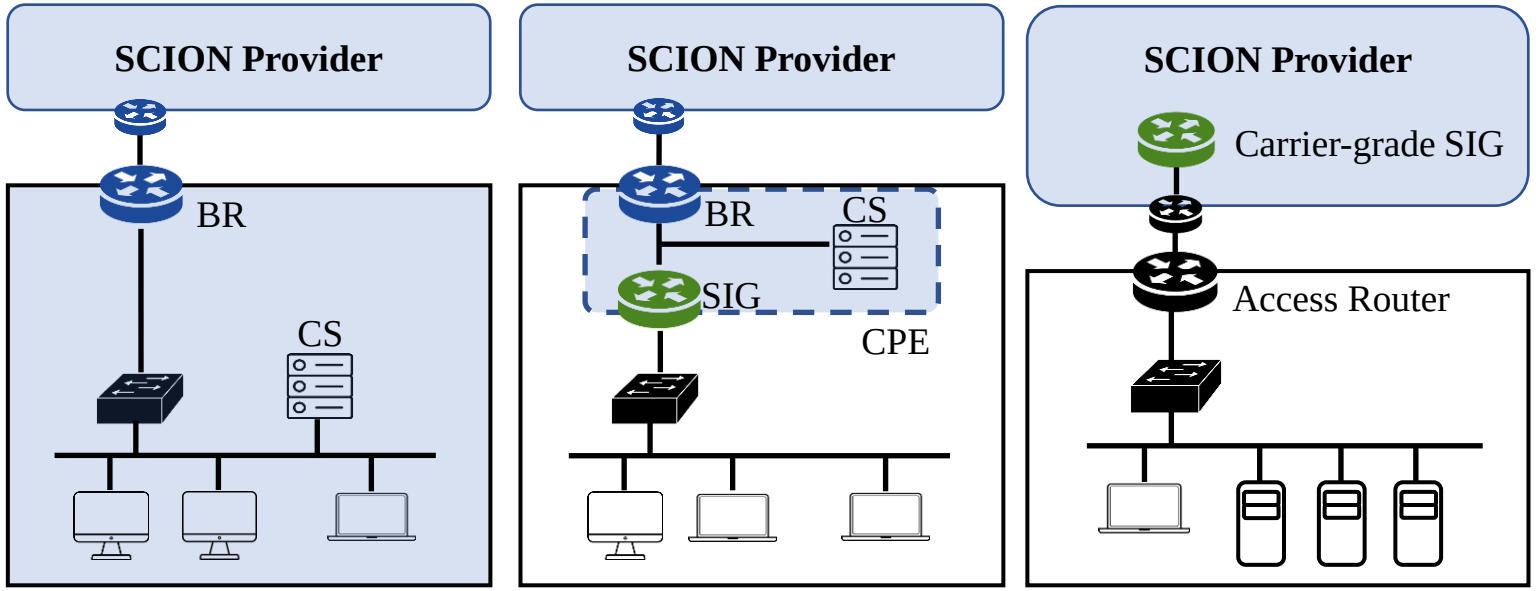

IXP Deployment

Two IXP Deployment Models:

- An IXP is treated as a large layer 2 switch between its customer ASes
- An IXP exposes its internal structure by modelling each site as an individual AS

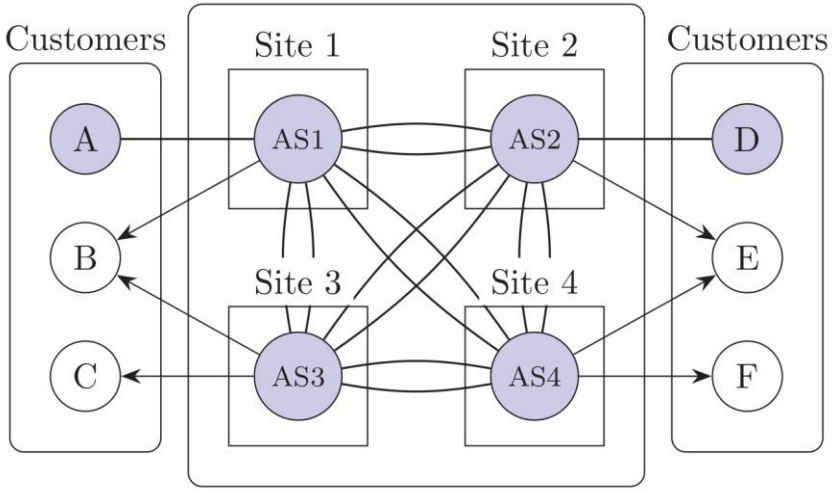

SwissIX SCION Peering Mesh: www.swissix.ch/participation/participants-scion

Preparing to reach the summit

What if we reach our goal of a global SCION deployment?

Over 70’000 ASes in the Internet and still growing!

Must ensure scalability of the system

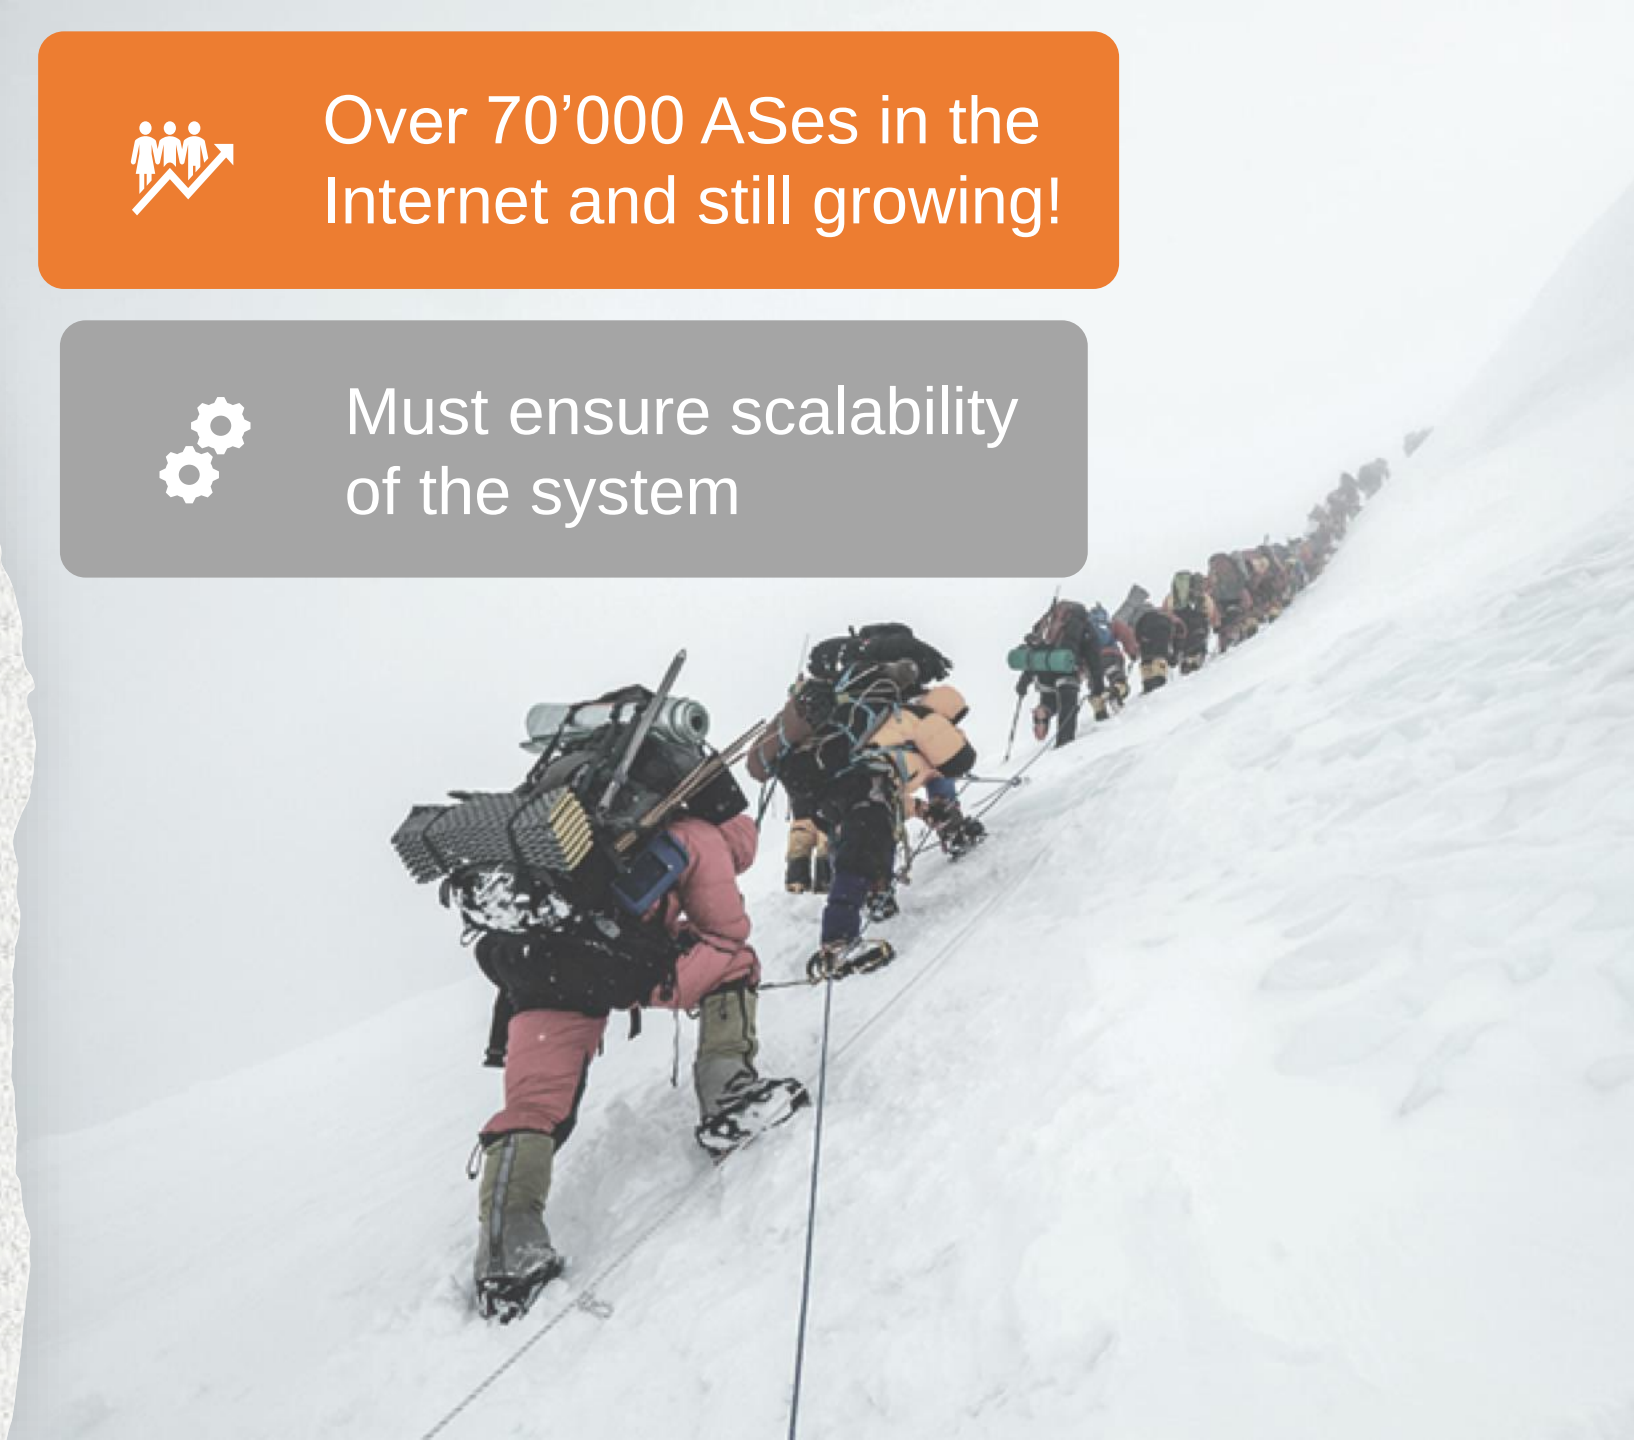

Scalability

- Grouping ASes into ISDs
- AS-level routing
- No inter-domain forwarding tables on routers
- No control-plane operations on routers
- Hierarchical beaconing
- Push-based connectivity establishment with pull-based path lookup

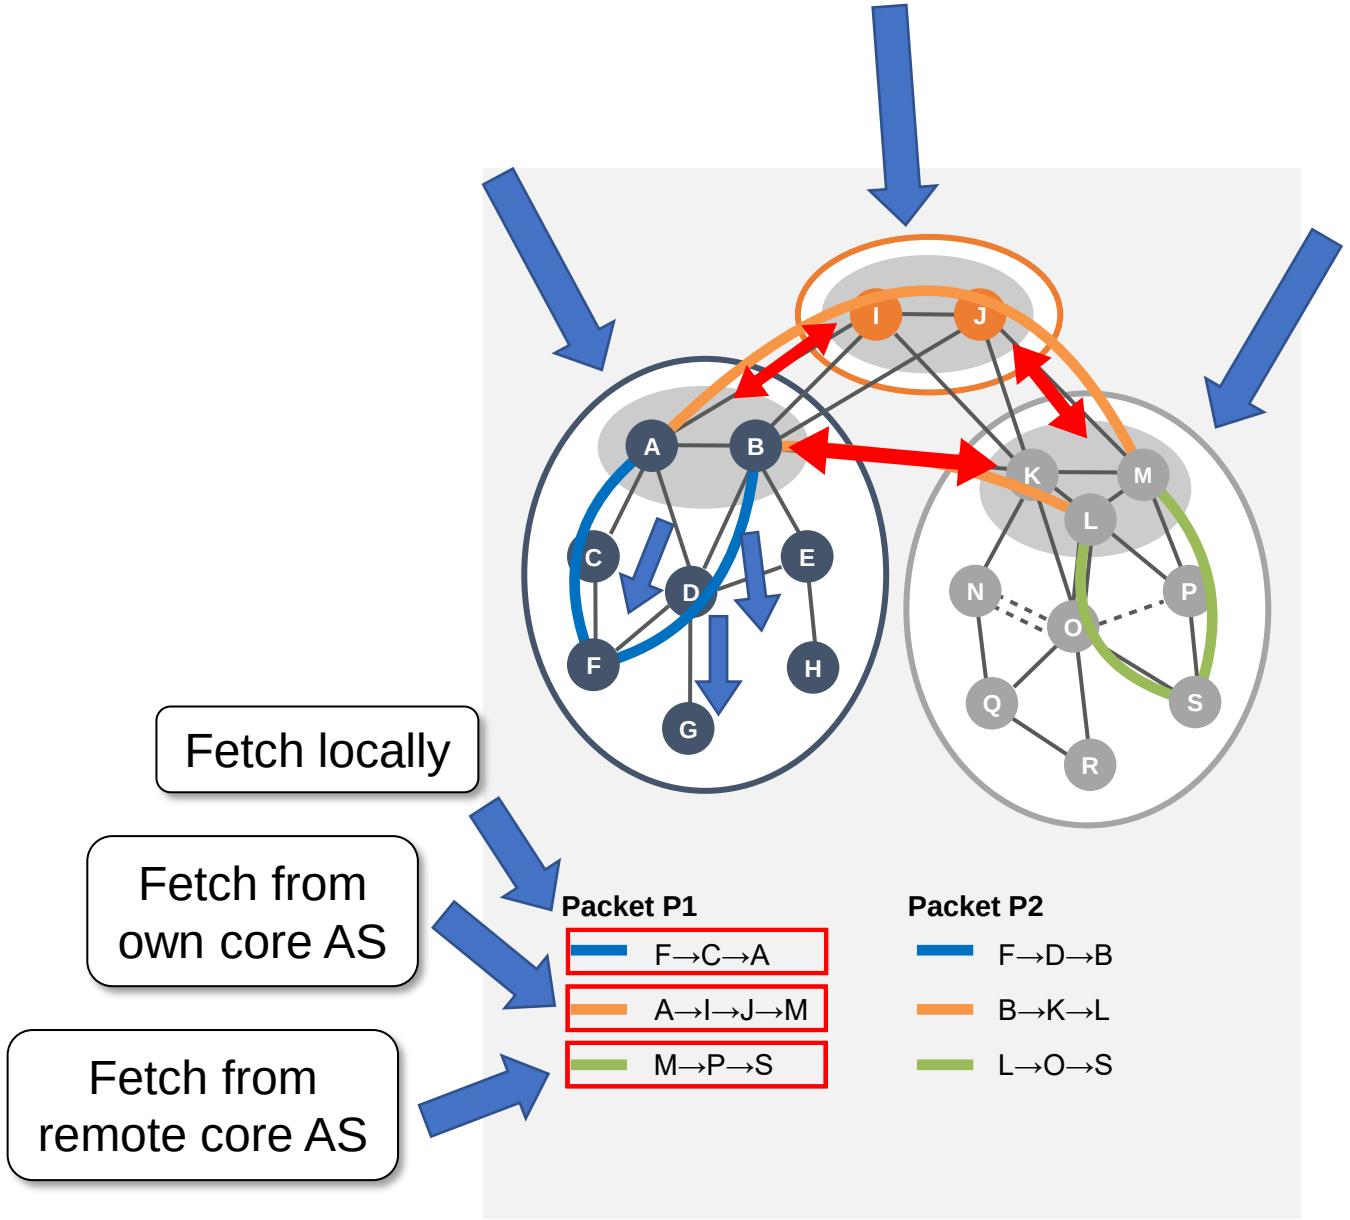

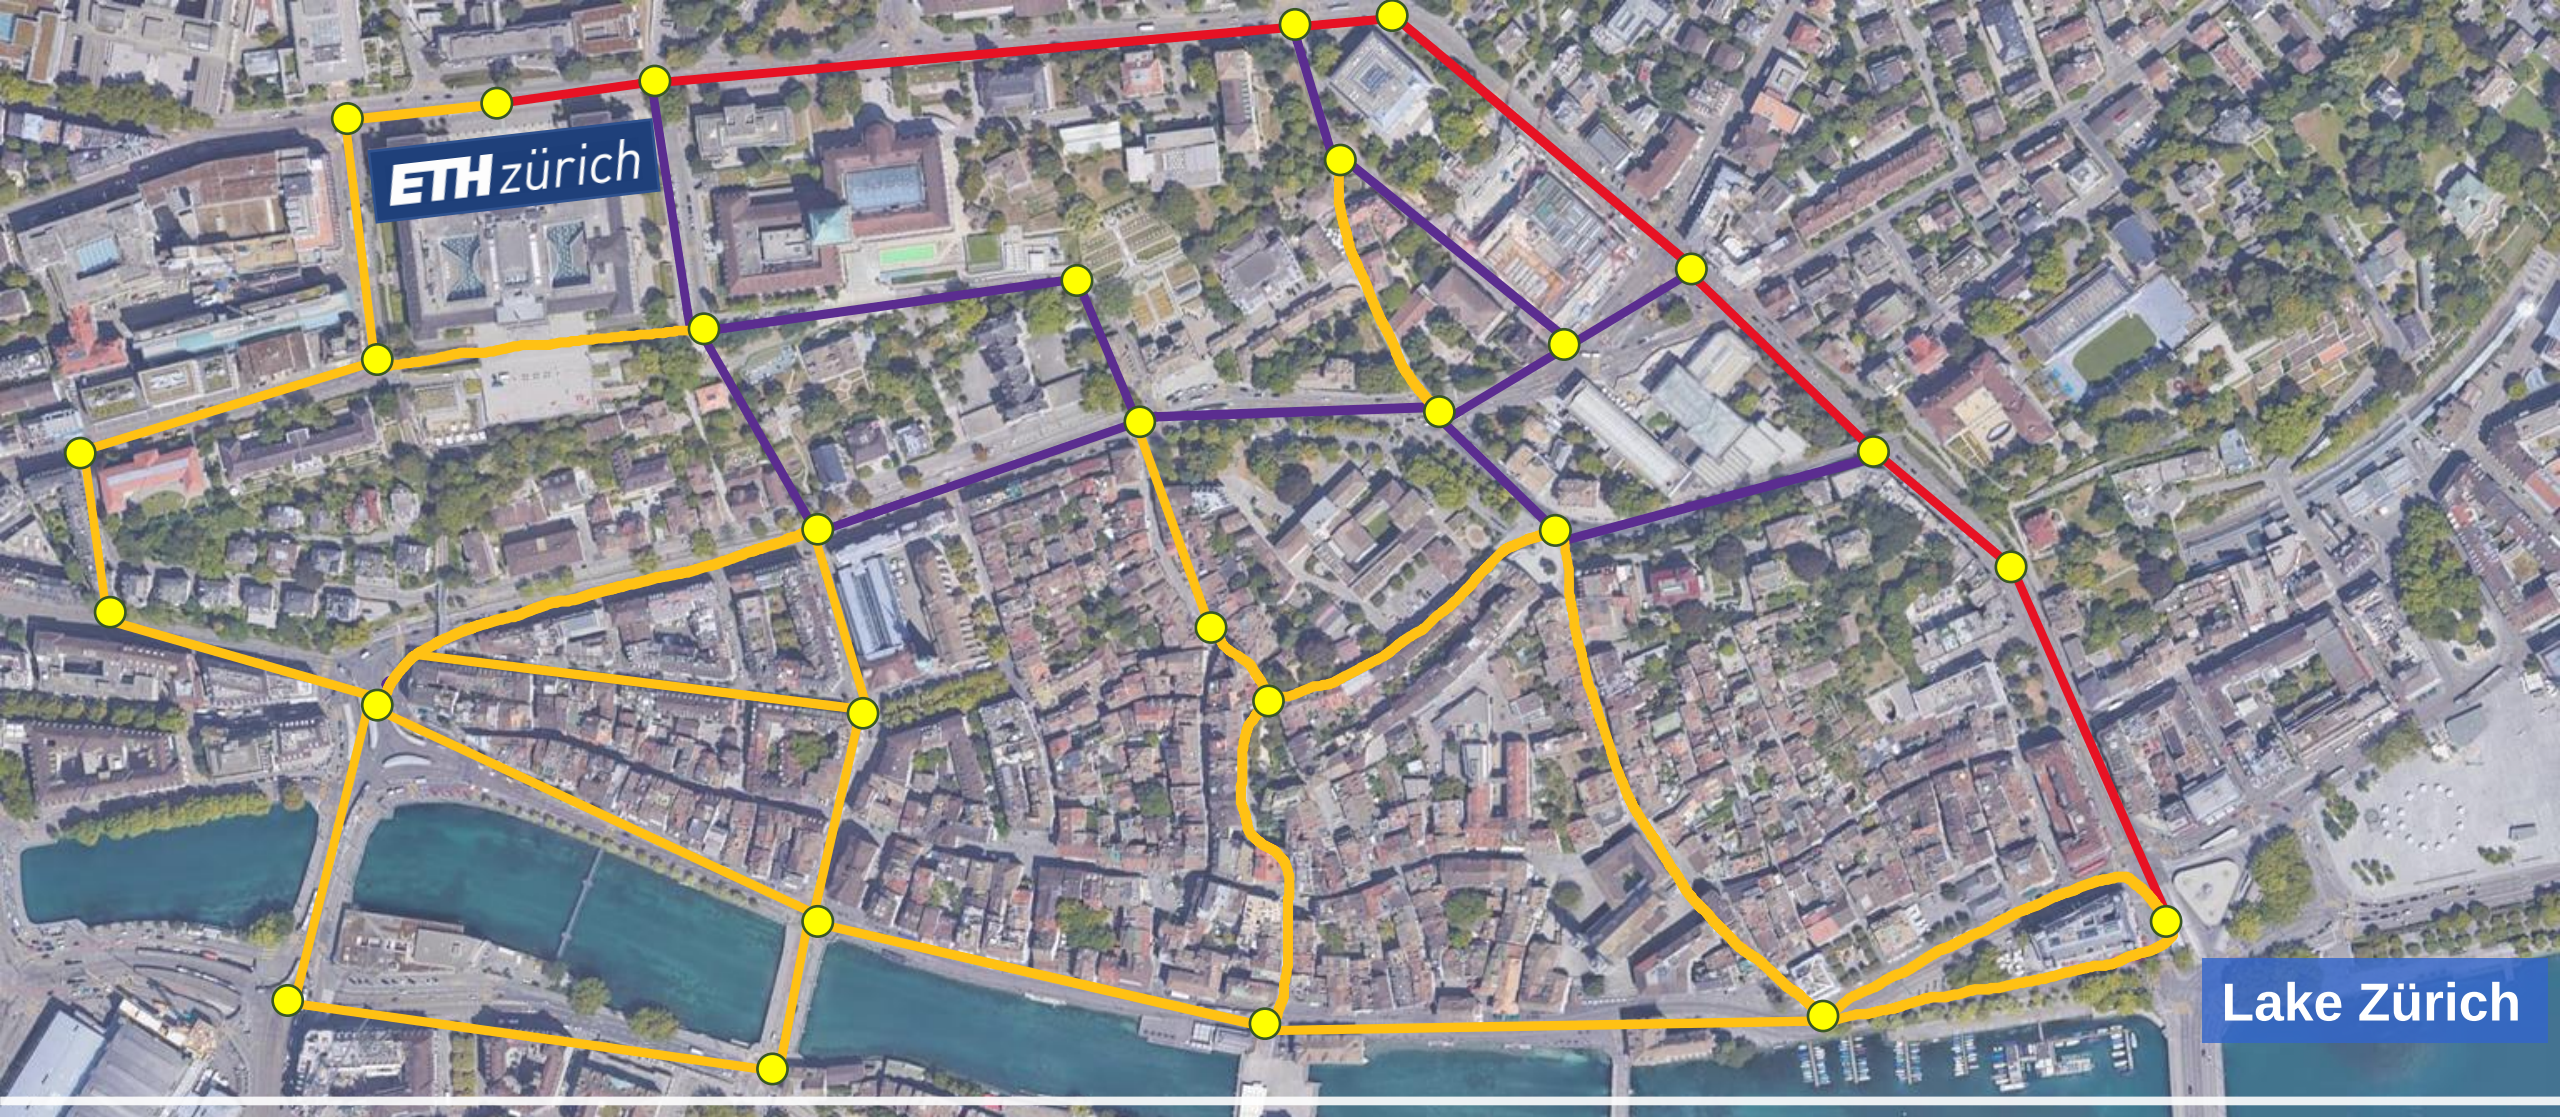

Path-Diversity-Aware Beaconing

Experimental Setup

- BGP data from RouteViews
- BGPsec and SCION data from simulations
- BGPsec topology (12000 ASes)
  - Based on CAIDA AS dataset and extrapolated to Internet scale
- Core topology (2000 ASes)
  - Based on CAIDA AS dataset by pruning lowest degree Ases
- Intra-ISD topology (~7000 ASes)
  - Find largest customer cone topology in CAIDA AS dataset

Control Plane Scalability Analysis

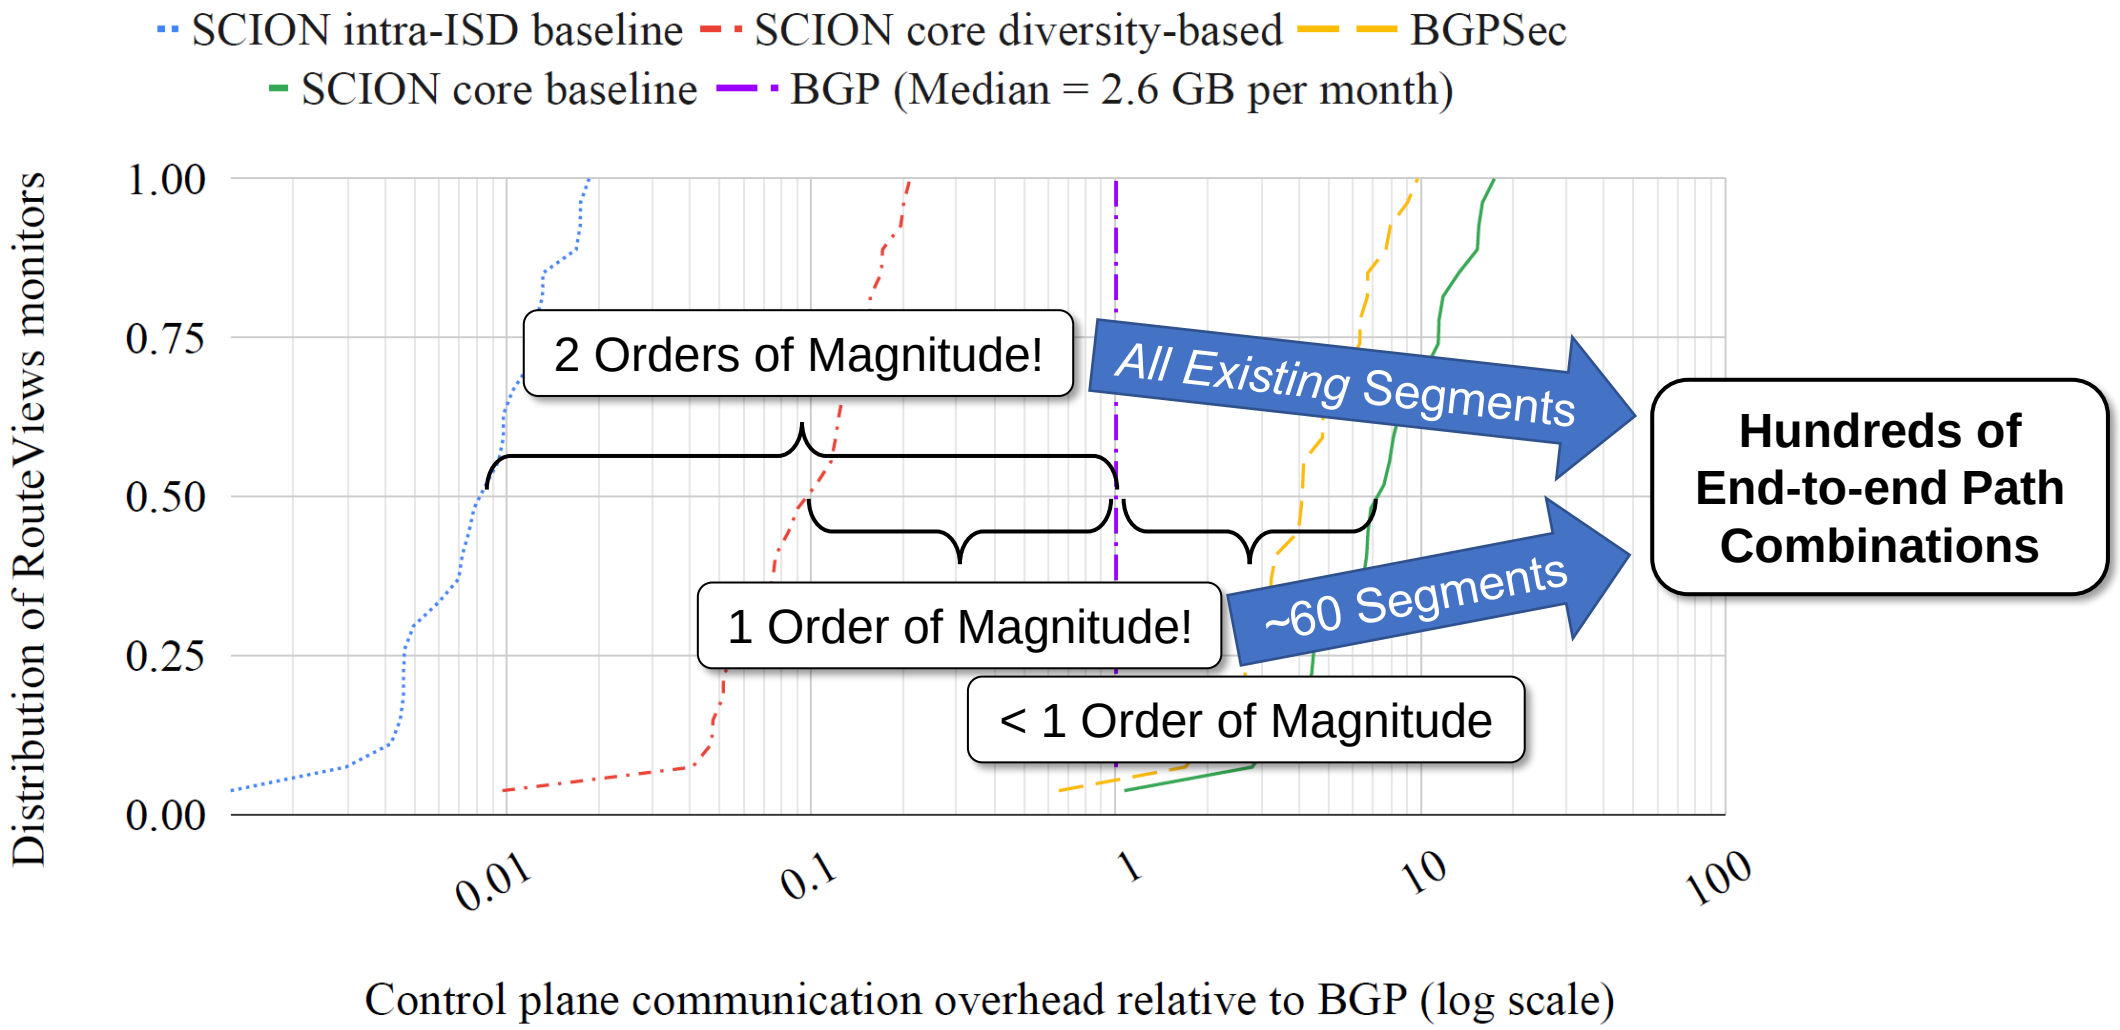

Deployment and Scalability of an Inter-Domain Multi-Path Routing Infrastructure

Cyrill Krähenbühl, ETH Zürich
Seyedali Tabaeiaghdaei, ETH Zürich
Christelle Gloor, ETH Zürich
Jonghoon Kwon, ETH Zürich
Adrian Perrig, Anapaya Systems, ETH Zürich
David Hausheer, OVGU Magdeburg
Dominik Roos, Anapaya Systems

ABSTRACT

Path aware networking (PAN) is a promising approach that enables endpoints to participate in end-to-end path selection. PAN unlocks numerous benefits, such as fast failover after link failures, application-based path selection and optimization, and native inter-domain multi-path. The utility of PAN hinges on the availability of a large number of high-quality path options. In an inter-domain context, two core questions arise. Can we deploy such an architecture *natively* in today's Internet infrastructure without creating an overlay relying on BGP? Can we build a scalable multi-path routing system that provides a large number of high-quality paths?

We first report on the real-world native deployment of the SCION next-generation architecture, providing a usable PAN infrastructure operating in parallel to today's Internet. We then analyze the scalability of the architecture in an Internet-scale topology. Finally, we introduce a new routing approach to further improve scalability.

CCS CONCEPTS

• **Networks** → **Network design principles**; **Control path algorithms**; **Network simulations**; **Routing protocols**.

KEYWORDS

SCION, Next-generation Internet Architecture, BGP, Deployment, Scalability, Control-Plane Algorithm Design, Multi-Path, Inter-Domain Routing, Network Simulation

ACM Reference Format:
Cyrill Krähenbühl, Seyedali Tabaeiaghdaei, Christelle Gloor, Jonghoon Kwon, Adrian Perrig, David Hausheer, and Dominik Roos. 2021. Deployment and Scalability of an Inter-Domain Multi-Path Routing Infrastructure. In *The 17th International Conference on emerging Networking EXperiments and Technologies (CoNEXT '21), December 7–10, 2021, Virtual Event, Germany*. ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3485983.3494862

about the paths their packets traverse. This is unlike the current Internet, where path selection is performed implicitly within the network. PAN enables exciting opportunities to evolve the Internet: multi-path communication can harness inherent Internet path diversity, application-based path selection allows endpoints to influence routing and choose optimal paths, and rapid failover can mask link failures.

Although intra-domain multi-path techniques are already in use, a large body of inter-domain multi-path research [5, 18, 19, 25, 30, 33, 39, 40, 59–62] has not seen large-scale deployment, mainly due to the challenges of deploying inter-domain protocols. Deploying a new inter-domain multi-path architecture requires a sizable nancial investment, and the architecture can only gain real-wo traction with tangible incentives for early adopters [55]. The Internet architecture SCION is the first inter-domain multi architecture in practical use. In this work, we revisit the deplo concept that underlies the SCION production network, al the incentive scheme it provides to early adopters. In th we will answer the following questions:

(1) What are the (early adopters') incentives for
(2) How could the deployment of SCION start the pre-eminent BGP-based infrastructur
(3) Can SCION scale to the size of the Inte

To understand why SCION was deployed cessful in the future, we must answer th incentives for initial and future SCIO discuss both short- and long-term i early deployment.

To answer the second questi SCION has been deployed free manner, co-existin on it. Furthermore,

More in the Paper

Case Study on Deployment
ISP, IXP, End-domain Deployment

Scalable Multi-Path Discovery Algorithm
Quality-aware Path Construction

Extensive Evaluation
Scalability analysis and comparison with BGP and BGPsec

Conclusion

- BGP(overlay)-free inter-domain multi-path infrastructure
- Operating on publicly available federated networks
- Shown to be scalable to a global deployment
- Innovation in inter-domain routing is possible

Thank you for your attention!

Cyrill Krähenbühl
Network Security Group
ETH Zürich
cyrill.kraehenbuehl@inf.ethz.ch