ETH zürich

SCION based ScienceDMZ fast file transfer for HPCCs

François Wirz
Netsec Group, ETH Zurich

CS3

March 13th 2024

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

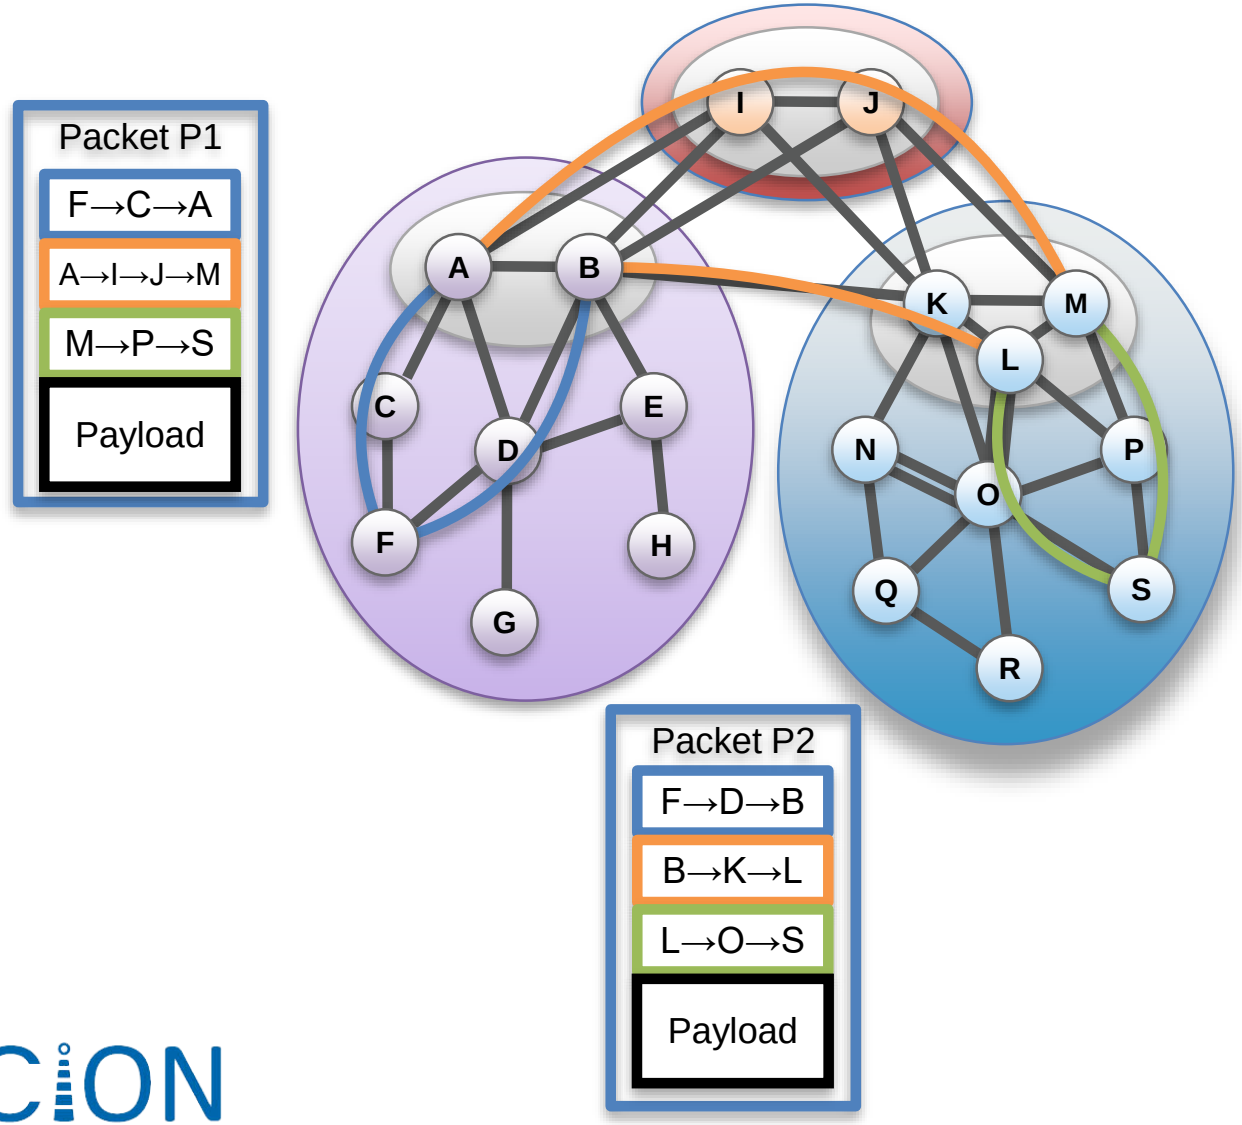

ETH zürich

SCION

SCION header overview

- Packets get routed within the AS using IPv4/v6, additional UDP header to accommodate middle boxes and OS network stacks

Link-layer/intra-domain header
SCION header
Transport-layer header

Common header
Addresses
Path
Extensions (optional)

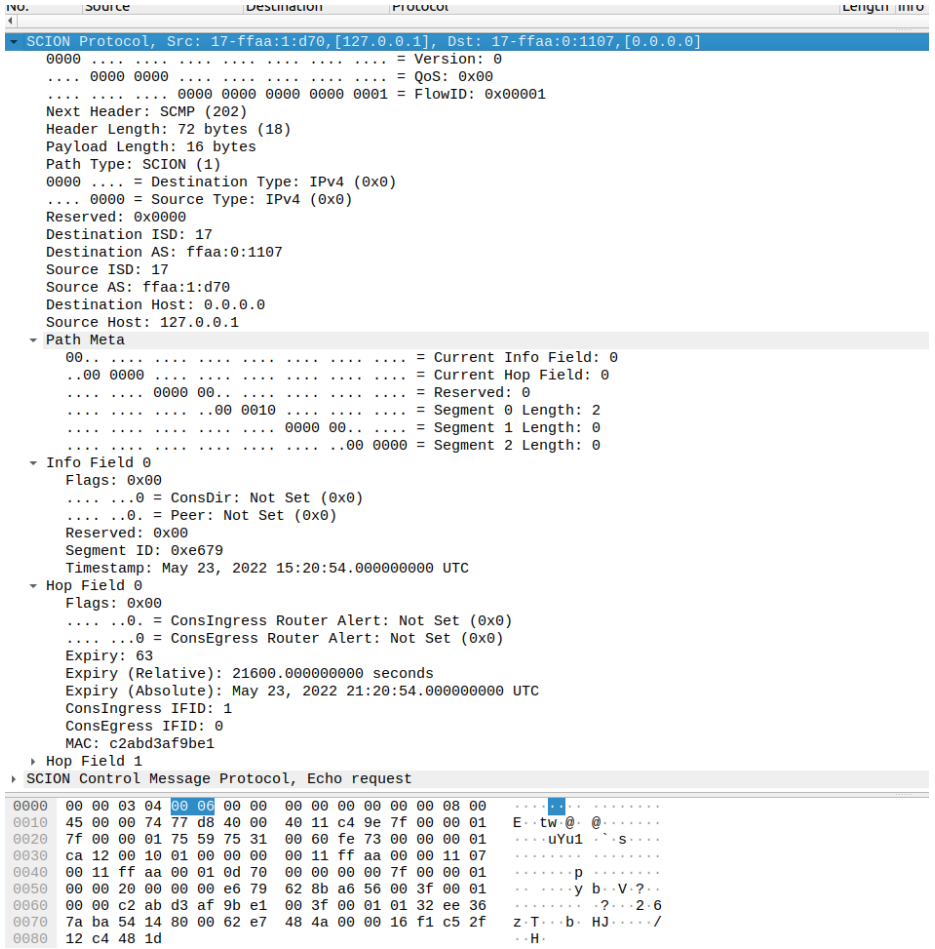

ETH zürich

SCION

Fast File Transfer on SCION

- Transmission scheme adapted to high BD product
- Bandwidth aggregation through multipathing
  - Possibility to use backup links active-active
  - Avoid in network bottlenecks
- OS network stack bypass
  - AF_XDP based solution, application traffic at line rate
- Adapted congestion control w/ PCC
  - Not affected by spurious loss, maintain fairness to competing flows
- Evaluated against existing file transfer applications
  - Compares very favourably in HBD settings vs GridFTP

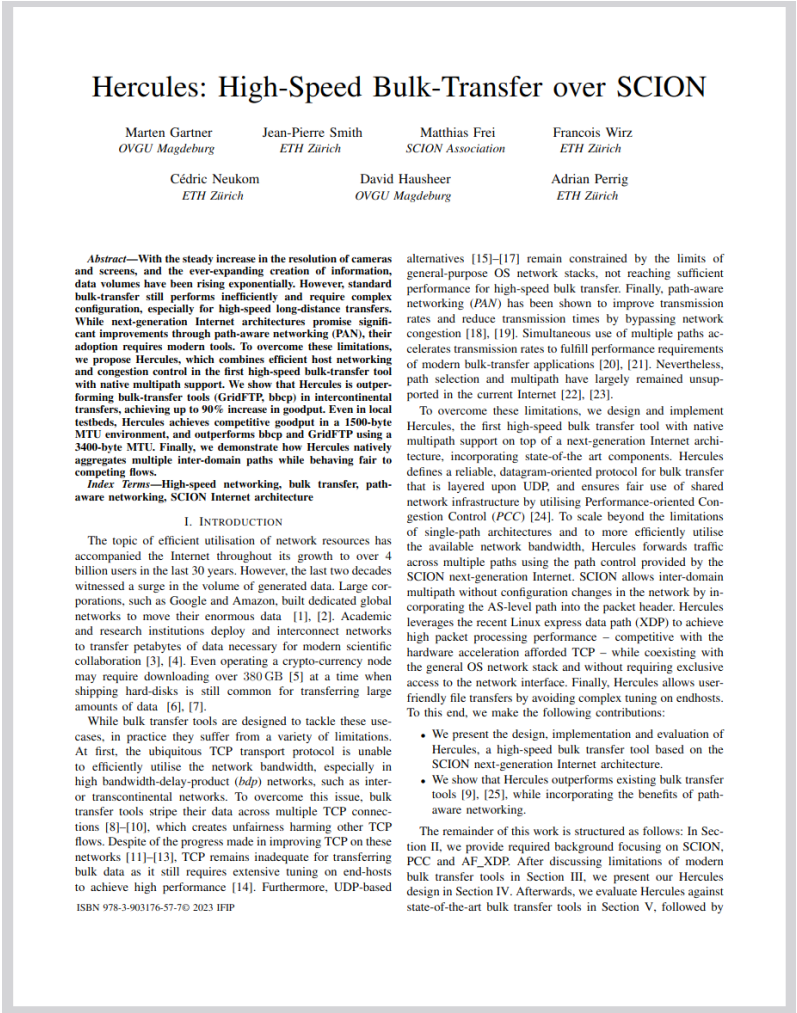

Hercules: High-Speed Bulk-Transfer over SCION

Marten Gartner, *OVGU Magdeburg*
Jean-Pierre Smith, *ETH Zürich*
Matthias Frei, *SCION Association*
Francois Wirz, *ETH Zürich*
Cédric Neukom, *ETH Zürich*
David Hausheer, *OVGU Magdeburg*
Adrian Perrig, *ETH Zürich*

***Abstract*—With the steady increase in the resolution of cameras and screens, and the ever-expanding creation of information, data volumes have been rising exponentially. However, standard bulk-transfer still performs inefficiently and require complex configuration, especially for high-speed long-distance transfers. While next-generation Internet architectures promise significant improvements through path-aware networking (PAN), their adoption requires modern tools. To overcome these limitations, we propose Hercules, which combines efficient host networking and congestion control in the first high-speed bulk-transfer tool with native multipath support. We show that Hercules is outperforming bulk-transfer tools (GridFTP, bbcp) in intercontinental transfers, achieving up to 90% increase in goodput. Even in local testbeds, Hercules achieves competitive goodput in a 1500-byte MTU environment, and outperforms bbcp and GridFTP using a 3400-byte MTU. Finally, we demonstrate how Hercules natively aggregates multiple inter-domain paths while behaving fair to competing flows.**

***Index Terms*—High-speed networking, bulk transfer, path-aware networking, SCION Internet architecture**

I. Introduction

The topic of efficient utilisation of network resources has accompanied the Internet throughout its growth to over 4 billion users in the last 30 years. However, the last two decades witnessed a surge in the volume of generated data. Large corporations, such as Google and Amazon, built dedicated global networks to move their enormous data [1], [2]. Academic and research institutions deploy and interconnect networks to transfer petabytes of data necessary for modern scientific collaboration [3], [4]. Even operating a crypto-currency node may require downloading over 380 GB [5] at a time when shipping hard-disks is still common for transferring large amounts of data [6], [7].

While bulk transfer tools are designed to tackle these use-cases, in practice they suffer from a variety of limitations. At first, the ubiquitous TCP transport protocol is unable to efficiently utilise the network bandwidth, especially in high bandwidth-delay-product (*bdp*) networks, such as inter- or transcontinental networks. To overcome this issue, bulk transfer tools stripe their data across multiple TCP connections [8]–[10], which creates unfairness harming other TCP flows. Despite of the progress made in improving TCP on these networks [11]–[13], TCP remains inadequate for transferring bulk data as it still requires extensive tuning on end-hosts to achieve high performance [14]. Furthermore, UDP-based alternatives [15]–[17] remain constrained by the limits of general-purpose OS network stacks, not reaching sufficient performance for high-speed bulk transfer. Finally, path-aware networking (*PAN*) has been shown to improve transmission rates and reduce transmission times by bypassing network congestion [18], [19]. Simultaneous use of multiple paths accelerates transmission rates to fulfill performance requirements of modern bulk-transfer applications [20], [21]. Nevertheless, path selection and multipath have largely remained unsupported in the current Internet [22], [23].

To overcome these limitations, we design and implement Hercules, the first high-speed bulk transfer tool with native multipath support on top of a next-generation Internet architecture, incorporating state-of-the art components. Hercules defines a reliable, datagram-oriented protocol for bulk transfer that is layered upon UDP, and ensures fair use of shared network infrastructure by utilising Performance-oriented Congestion Control (*PCC*) [24]. To scale beyond the limitations of single-path architectures and to more efficiently utilise the available network bandwidth, Hercules forwards traffic across multiple paths using the path control provided by the SCION next-generation Internet. SCION allows inter-domain multipath without configuration changes in the network by incorporating the AS-level path into the packet header. Hercules leverages the recent Linux express data path (XDP) to achieve high packet processing performance – competitive with the hardware acceleration afforded TCP – while coexisting with the general OS network stack and without requiring exclusive access to the network interface. Finally, Hercules allows user-friendly file transfers by avoiding complex tuning on endhosts. To this end, we make the following contributions:

- We present the design, implementation and evaluation of Hercules, a high-speed bulk transfer tool based on the SCION next-generation Internet architecture.
- We show that Hercules outperforms existing bulk transfer tools [9], [25], while incorporating the benefits of path-aware networking.

The remainder of this work is structured as follows: In Section II, we provide required background focusing on SCION, PCC and AF_XDP. After discussing limitations of modern bulk transfer tools in Section III, we present our Hercules design in Section IV. Afterwards, we evaluate Hercules against state-of-the-art bulk transfer tools in Section V, followed by

ISBN 978-3-903176-57-7© 2023 IFIP

ETH zürich

SCION

SCION ScienceDMZ

- Per packet traffic filtering leveraging the CPPKI of SCION
  - LightningFilter using DRKey
- Bandwidth aggregation through multipathing
  - Hercules using path-awareness
- Deployment model
  - Classical ScienceDMZ model with host authentication, can be deployed in parallel to existing systems, but at a fraction of the costs using COTS hardware

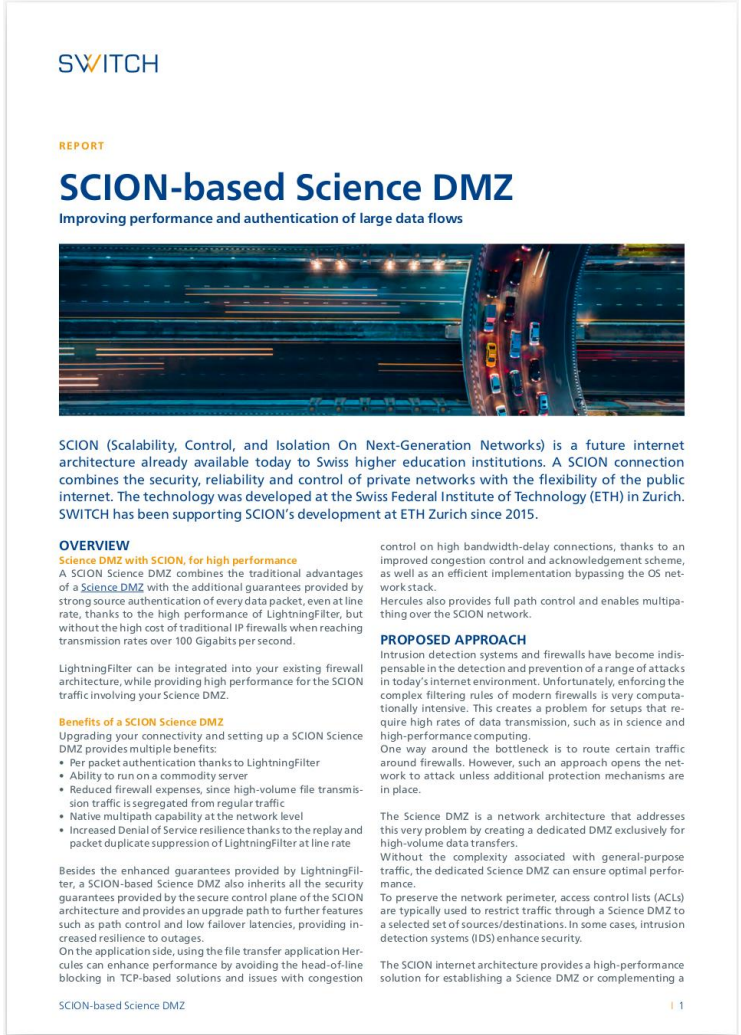

SWITCH

REPORT

SCION-based Science DMZ

Improving performance and authentication of large data flows

SCION (Scalability, Control, and Isolation On Next-Generation Networks) is a future internet architecture already available today to Swiss higher education institutions. A SCION connection combines the security, reliability and control of private networks with the flexibility of the public internet. The technology was developed at the Swiss Federal Institute of Technology (ETH) in Zurich. SWITCH has been supporting SCION's development at ETH Zurich since 2015.

OVERVIEW

Science DMZ with SCION, for high performance

A SCION Science DMZ combines the traditional advantages of a Science DMZ with the additional guarantees provided by strong source authentication of every data packet, even at line rate, thanks to the high performance of LightningFilter, but without the high cost of traditional IP firewalls when reaching transmission rates over 100 Gigabits per second.

LightningFilter can be integrated into your existing firewall architecture, while providing high performance for the SCION traffic involving your Science DMZ.

Benefits of a SCION Science DMZ

Upgrading your connectivity and setting up a SCION Science DMZ provides multiple benefits:

- Per packet authentication thanks to LightningFilter
- Ability to run on a commodity server
- Reduced firewall expenses, since high-volume file transmission traffic is segregated from regular traffic
- Native multipath capability at the network level
- Increased Denial of Service resilience thanks to the replay and packet duplicate suppression of LightningFilter at line rate

Besides the enhanced guarantees provided by LightningFilter, a SCION-based Science DMZ also inherits all the security guarantees provided by the secure control plane of the SCION architecture and provides an upgrade path to further features such as path control and low failover latencies, providing increased resilience to outages.
On the application side, using the file transfer application Hercules can enhance performance by avoiding the head-of-line blocking in TCP-based solutions and issues with congestion control on high bandwidth-delay connections, thanks to an improved congestion control and acknowledgement scheme, as well as an efficient implementation bypassing the OS network stack.
Hercules also provides full path control and enables multipathing over the SCION network.

PROPOSED APPROACH

Intrusion detection systems and firewalls have become indispensable in the detection and prevention of a range of attacks in today's internet environment. Unfortunately, enforcing the complex filtering rules of modern firewalls is very computationally intensive. This creates a problem for setups that require high rates of data transmission, such as in science and high-performance computing.
One way around the bottleneck is to route certain traffic around firewalls. However, such an approach opens the network to attack unless additional protection mechanisms are in place.

The Science DMZ is a network architecture that addresses this very problem by creating a dedicated DMZ exclusively for high-volume data transfers.
Without the complexity associated with general-purpose traffic, the dedicated Science DMZ can ensure optimal performance.
To preserve the network perimeter, access control lists (ACLs) are typically used to restrict traffic through a Science DMZ to a selected set of sources/destinations. In some cases, intrusion detection systems (IDS) enhance security.

The SCION internet architecture provides a high-performance solution for establishing a Science DMZ or complementing a

ETH zürich

SCION

First deployments and POCs

- Performance evaluation of the integrated system on the SWITCH network
  - Sustained throughput approaching link capacity achieved between Bern and Lugano during file transfer

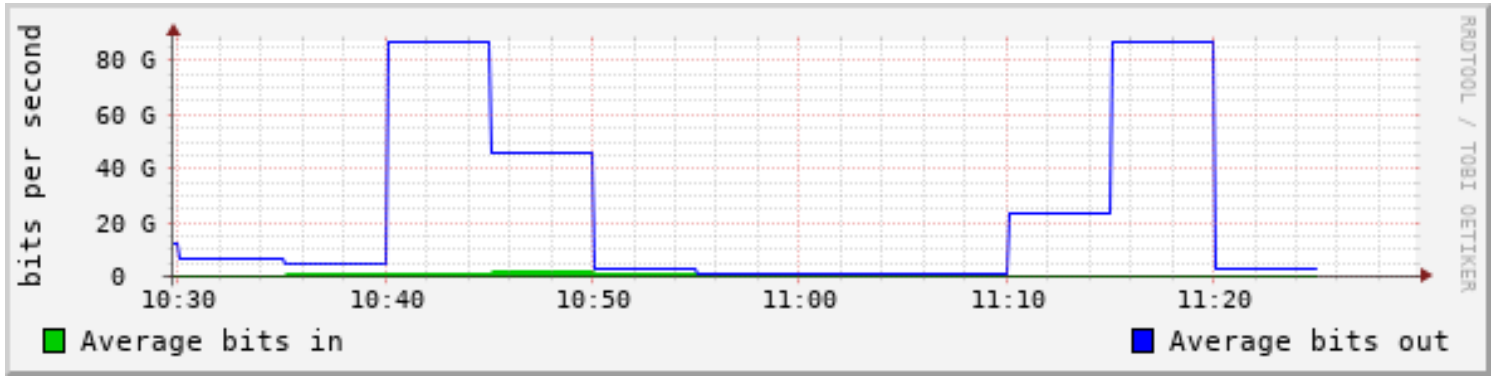

First deployments and POCs

- Deployment between CSCS and ETHZ
  - Different deployments possible depending on local network topology
  - LightningFilter can be a bump in the wire, collocated with the DTN or act as a central gateway for the ScienceDMZ

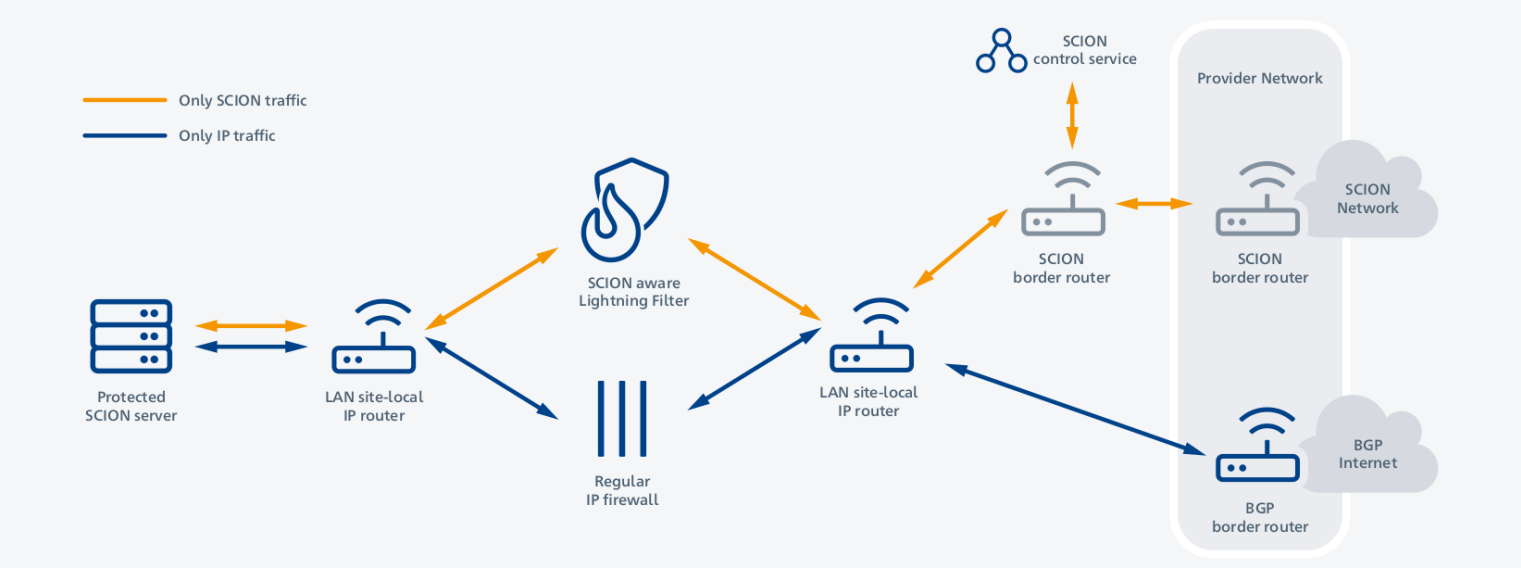

First deployments and POCs

- Deployment between CSCS and ETHZ
  - Different deployments possible depending on local network topology
  - LightningFilter can be a bump in the wire, collocated with the DTN or act as a central gateway for the ScienceDMZ
- POC between SEC Singapore and ETHZ
  - Huge improvement over existing solution, transfer time improved by two orders of magnitude
- DTN at KISTI connecting to Europe via KREONET2
  - Fully redundant links connecting the KISTI SCION AS to Switzerland

SCIERA: SCION Education, Research and Academic Network

- SCION deployment has been organically growing
  - Service offering by SWITCH in Switzerland
  - Availability within and through GEANT
  - Deployed in KREONET2
  - U.S. universities connected
  - WACREN in Africa about to join

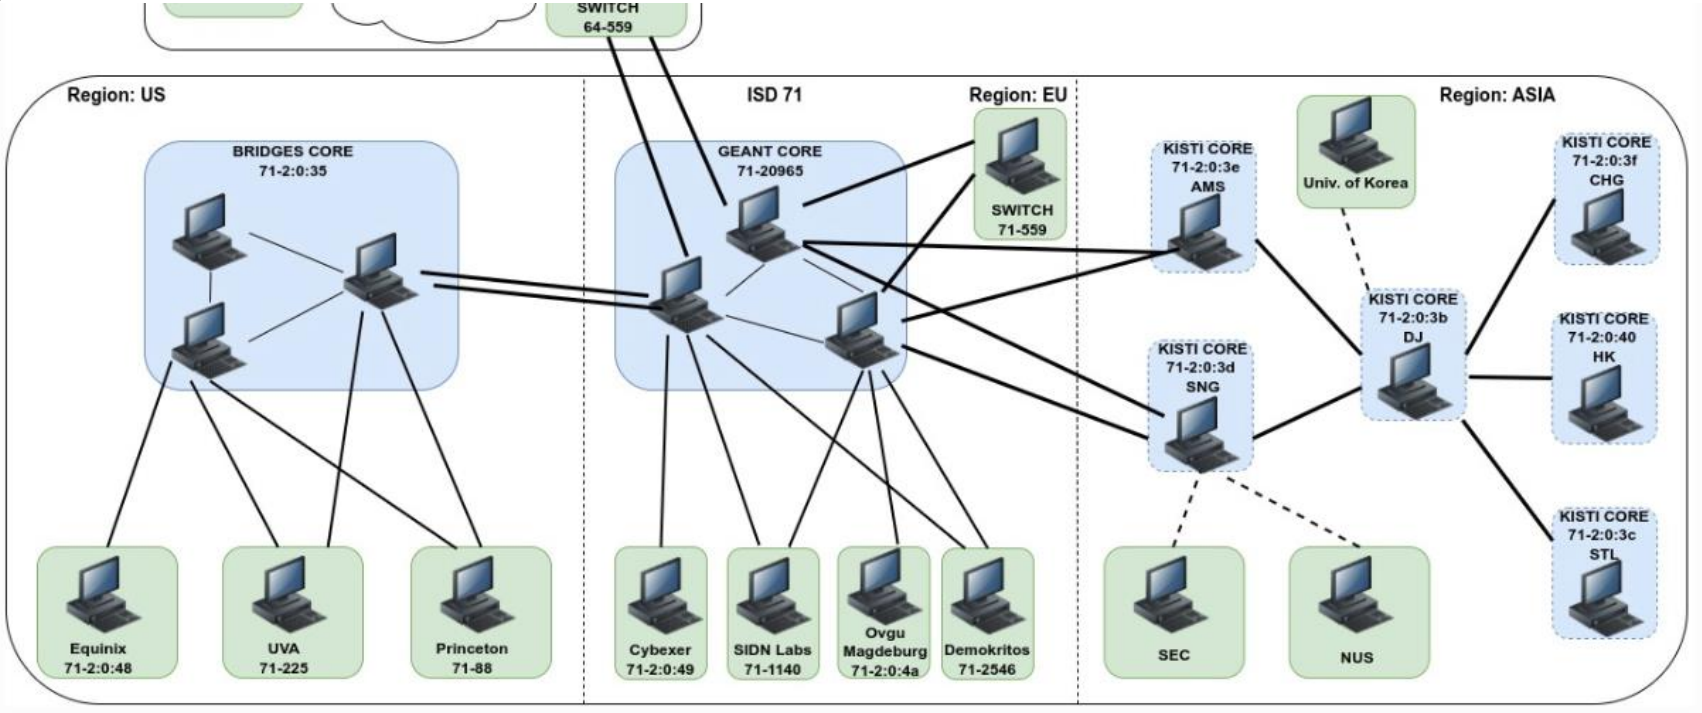

SCION Team

ETH zürich

SCION

Thank you for your attention

François Wirz

Software Engineer

Network Security Group

netsec.ethz.ch

wirzf@ethz.ch

Resources:

https://scion.org

https://scion-architecture.net

SCION Summary

- Availability, Sovereignty, and Transparency
- 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
- SCION: Next-generation Internet you can use **today**