NEXT GENERATION INTERNET

Open Call 5

Leveraging Path Diversity to Enhance Resilience, Scalability and Energy-Efficiency with SCION

Deliverable 3: Experiment Results and Final Report

Authors	David Hausheer, hausheer@ovgu.de, OVGU Magdeburg
Due Date	30-11-2022
Submission Date	05-12-2022
Keywords	

Deliverable 3: Part I

Analysis, results, and wider impact

1 Abstract

SCION is a novel NGI architecture that has reached a level of maturity, which renders it ready today for large-scale deployment. Thus, the objective of this project is to deploy SCION on the NSF BRIDGES infrastructure over two very high-speed transatlantic links. GMU will enable and support the SCION deployment on NSF BRIDGES, with OVGU providing a European SCION testbed over GEANT. Our project demonstrates the SCION benefits by means of experiments between the US and Europe over the SCIONLab testbed to show the privacy-enhancement (i.e. by splitting traffic over multiple paths) and improved reliability (i.e. with multi-path and seamless path failover) over SCION, as well as the scalability of our SCION-based path discovery mechanisms which help to effectively reduce the network's power consumption and incentivize ISPs and transit providers to shift towards greener electricity.

2 Project Vision

SCION is a clean-slate Next-Generation Internet (NGI) architecture (cf. [1], [8]) designed to provide route control, failure isolation, and explicit trust information for end-to-end communication. SCION organizes existing Autonomous Systems (ASes) into groups of independent routing planes, called isolation domains, which interconnect to provide global connectivity. Isolation domains naturally isolate routing failures and misconfigurations, give endpoints strong control over both inbound and outbound traffic, and provide meaningful and enforceable trust. As a result, the SCION architecture provides strong resilience and security properties as an intrinsic consequence of its design.

Besides high security, SCION also provides a scalable routing infrastructure and highly efficient packet forwarding. As a path-based architecture, SCION end hosts learn about available network path segments, and combine them into end-to-end paths that are carried in packet headers. Thanks to embedded cryptographic mechanisms, path construction is constrained to the route policies of Internet Service Providers (ISPs) and receivers, offering path choice to all parties: senders, receivers, and ISPs. This approach enables path-aware communication, an emerging trend in networking. These features also enable multipath communication, which enhances privacy by obfuscating traffic over multiple paths, facilitates high availability, rapid failover in case of network failures, increased end-to-end bandwidth, dynamic traffic optimization, and resilience to DDoS attacks.

Since the start, the research on improving the SCION architecture has been complemented by implementation and deployment. Since 2014, an operational SCION network has been in operation in Switzerland. The current maturity of the software has been achieved through professional and experienced developers, which have recently built the 5th generation of the software. Moreover, this has been complemented by substantial efforts in formally modeling and verifying key parts of the SCION architecture, such that SCION has reached a level of maturity, which renders it ready today for large-scale deployment.

In order to provide a low entry bar for researchers and application developers who wish to explore and assess the capabilities of SCION, we have developed SCIONLab [2], a flexible, scalable, and extensible SCION network running at global scale. SCIONLab provides users with fast setup, enabling them to instantiate a SCION node as a VM in a few clicks, requiring little technical expertise to join the SCION network. Thereby, SCION nodes can contribute to the routing within the SCION topology and researchers can attach their own computing resources anywhere in within the SCIONLab network.

To this end, the objectives of this project are twofold: (a) To interconnect SCIONLab with the BRIDGES infrastructure over two very high-speed transatlantic links, in order to increase the path diversity and bandwidth for SCIONLab experiments and (b) to run experiments between the US and Europe over the SCIONLab testbed to demonstrate the privacy-enhancement (e.g., by splitting traffic over multiple paths) and improved reliability (e.g. with multi-path and seamless path failover) over SCION, as well as to show the scalability of our SCION-based path discovery mechanisms which help to effectively reduce the network's power consumption and incentivize ISPs and transit providers to shift towards greener electricity.

The SCIONLab testbed has currently a strong basis in Europe, with deployments, a.o., in multiple research networks including GEANT, SWITCH, DFN-GVS, SIDN, and Fed4FIRE+

Co-funded by the Horizon 2020
Framework Programme of the European Union

(VirtualWall). However, it lacks high-bandwidth connectivity and path diversity to their US counterparts, such as Internet2, FABRIC and others. By interconnecting SCIONLab with BRIDGES' infrastructure, we are able to increase the path opportunities over the Atlantic and carry out experiments at very high speeds using SCION's native path-awareness and multipath support. The enhanced SCIONLab network will be maintained even after the end of this project and offered on a continuous basis to experimenters and application developers in other NGI topics.

3 Details on participants (both EU and US)

- Prof. Dr. David Hausheer is a Professor at the Faculty of Computer Science at Otto-von-Guericke-University Magdeburg since May 2017. He holds a diploma degree in electrical engineering and a Ph.D. degree in technical sciences from ETH Zurich. From 2011 - 2017 he was an assistant professor at TU Darmstadt, Germany. Prior to that he has been employed as a senior researcher and lecturer at University of Zurich, Switzerland from 2005 - 2011, while being on leave as a visiting scholar at EECS, UC Berkeley from October 2009 to April 2011. Prof. Hausheer has been leading the project and guiding the corresponding work.
- M.Sc. Marten Gartner is a Ph.D. student in the team of Prof. Hausheer at the Faculty of Computer Science at Otto-von-Guericke-University Magdeburg since February 2021. He holds an M.Sc. degree in computer science from OVGU Magdeburg. His research interests are on the design and analysis of high-performance multipath applications in path-aware networks. He has also a strong expertise in managing SCIONLab, specifically the high-speed border router and high-speed file transfers over SCION. Marten Gartner contributed both to the setup of SCIONLab over BRIDGES as well as to the experiments.
- M.Sc. Tony John is a Ph.D. student in the team of Prof. Hausheer at the Faculty of Computer Science at Otto-von-Guericke-University Magdeburg since January 2021. He holds an M.Sc. degree in digital engineering from OVGU Magdeburg. His research interests are on secure Internet of Things based on the SCION Internet architecture. Tony John was mainly contributing to the experiments in terms of bandwidth and latency measurements of SCION over BRIDGES.
- M.Sc. Thorben Krüger is a Ph.D. student in the team of Prof. Hausheer at the Faculty of Computer Science at Otto-von-Guericke-University Magdeburg since May 2020. He holds an M.Sc. degree in system and network engineering from the University of Amsterdam. His research interests are on traffic optimization via performance-based path selection. Thorben Krüger was mainly contributing to

Co-funded by the Horizon 2020
Framework Programme of the European Union

the experiment evaluation. He will be joining the project team from October 2022.

- Jerry Sobieski is a Co-Principle Investigator of the BRIDGES Project (NSF and GMU). Prior to this, he was a CRO at NORDUnet driving advanced networking research efforts between NORDUnet, the European R&E community, and the Global R&E Networks internationally. Current areas of focus include broad scale network virtualization architecture, 5G wireless/cellular integration with core wired [R&E] networks, dynamic/automated composable [network] service provisioning, inter-domain orchestration technologies/protocols, machine learning in global networks. Jerry Sobieski has been leading the US-side of the project team and strongly supported the SCION setup over the BRIDGES infrastructure.
- Dr. Bijan Jabbari is a Professor of Electrical and Computer Engineering, College of Engineering and Computing at GMU. He holds an M.S. degree in Engineering Economics Systems and a Ph.D. degree in Electrical Engineering from Stanford University. Dr. Jabbari is a Co-Principle Investigator of the BRIDGES Project (NSF and GMU). He is tracking the contractual aspects (e.g. related to fiber crossconnects) on the US-side of the project team in support of the SCION setup over the BRIDGES infrastructure.
- Dr. Liang Zhang is a postdoctoral research fellow at GMU. He received his M.S. degree in information and communication engineering, University of Science and Technology of China (USTC), China, in 2014, and his Ph.D. degree in electrical engineering from New Jersey Institute of Technology (NJIT), USA, in 2020. Liang Zhang also supported the SCION setup over the BRIDGES infrastructure.

4 Results

The primary result of this project is the deployment of SCION over BRIDGES. To achieve this, we had to connect SCION to BRIDGES by deploying servers (including border router instances and SCION control service instances etc.) on both ends (US/EU) and establish underlying network topologies in order to implement an appropriate deployment setup of SCION links across the Atlantic.

On the European side, we have setup six dedicated SCION servers, two each in Paris, Geneva, and Frankfurt. Three of these servers are equipped with SCION high-speed border router licenses purchased from Anapaya Systems AG by OVGU.

Co-funded by the Horizon 2020 Framework Programme of the European Union

On the US side, we have setup three SCION servers (provided by CMU) within BRIDGES: two have been deployed at the BRIDGES cabinet at Equinix DC2 in Ashburn, VA, as well as one will be deployed at MAN LAN (NYC) very soon. These servers are equipped with the open source SCION border router. The BRIDGES infrastructure is connected to MARIA (Mid-Atlantic Research Infrastructure Alliance) and to Internet2. This allowed us to connect the University of Virginia (UVa) over a dedicated VLAN through MARIA to BRIDGES, which determines the first native SCION connection to a university in the US. Additionally, we will also connect Princeton via NJEdge to our SCION deployment at MAN LAN. UVa and Princeton are both new SCION nodes that we hadn't envisioned at the time of writing the proposal, but which were made possible thanks to the SCION deployment within BRIDGES. Additionally, this deployment enables many more interesting connection opportunities to academic as well as industry customers: via Internet2 exchange points in WIX (Washington IX) and MAN LAN, as well as to the Equinix IX via a local cross-connection at Equinix DC2. We will explore these options further beyond the end of the project.

In order to connect the SCION servers in GEANT with BRIDGES over dedicated L2 links at up to 2x 100Gbps speeds, we have pursued several options in parallel: In Paris, the SCION servers are in the same location as the BRIDGES servers at the Interxion data center, thus a cross-connection between the two cabinets is the most straight-forward option. However, until now there had been no clear commitment to establish this connection by the involved parties. In the meantime, we had further discussions and received commitment for a SCION deployment within SURFnet/University of Amsterdam (UvA) in Amsterdam Science Park. Since BRIDGES is also present there, this enables us to establish a connection between BRIDGES and GEANT via SURFnet/UvA. As a third option, we will also keep investigating connectivity between GEANT and BRIDGES via RENATER in Paris. To this end, we will remain in close contact with all involved parties and keep pushing beyond the end of the project. While for now we had been limited to SCION experiments within either BRIDGES or GEANT only, we expect to be able to run experiments jointly across both SCION infrastructures in the near future.

While SCION was deployed over BRIDGES, we have designed and implemented the corresponding measurement environment in order to conduct all the relevant measurements. The actual assessment includes three major experiments run across SCION, targeted at a) demonstrating the resilience enhancement over SCION, b) showing the increased scalability of our enhanced path discovery mechanisms, and c) evaluating the improved energy-efficiency over the SCION network. To this end, we have designed and implemented Hercules, a reliable, datagram oriented protocol for large-scale bulk data transfers. To scale beyond the limitations of single-path architectures, Hercules forwards traffic across multiple paths using the multipathing and path control provided by the SCION network.

Co-funded by the Horizon 2020
Framework Programme of the European Union

Hercules uses PCC [14] to regulate its sending rate in response to network congestion. Rate adjustments in PCC occur periodically as governed by the measurement interval. As a result, PCC does not need acknowledgement information until after a measurement interval when it is ready to decide whether to change its sending rate. This synergises well with the delayed acknowledgements employed in Hercules, since large files transfers are not interactive, the bulk of the data is already available for transfer, and delivering a chunk out of order does not negatively impact the overall performance. Hercules can therefore send significantly fewer acknowledgements, which reduces the load on the reverse link while effectively avoiding congestion and quickly utilising the network bandwidth.

We implemented Hercules using the Go and C languages, and evaluated it in a series of memory-to-memory (M2M) transfers. M2M transfers allow us to evaluate network utilisation without the bottlenecks of disk reads and writes.

While waiting for the SCION deployment over BRIDGES to become available, we evaluated Hercules' simultaneous use of multiple network paths by measuring its performance across 20 transfers utilising two SCIONLab paths over the transcontinental academic network GEANT, while one of the paths was congested. Hercules transferred data between Geneva, Switzerland and Amsterdam, Netherlands and was instructed to use paths through Paris, France and Hamburg, Germany. Each path had an RTT just below 25 ms; and, since each location was provisioned with a single 10 Gbps duplex link, the packets leaving Geneva on both the Paris and Hamburg paths shared Geneva's outgoing 10 Gbps capacity.

Beginning 15 s after the transmission's start, we reduced the available bandwidth of the Geneva--Hamburg--Amsterdam path by congesting the link at Hamburg for 15 s with UDP traffic. Since a single TCP flow traversing the congested link would be a poor comparison to Hercules, we instead compared it to two independent, tuned TCP flows (iperf3+) which traversed the same paths by using VLAN tags and firewall rules.

Figure 1 shows the per-path and aggregated throughputs achieved by Hercules and the two independent, tuned TCP flows over the SCIONLab infrastructure in GEANT. When faced with congestion on one path, Hercules increased transmission on the second path to make use of the available capacity of Geneva's outgoing link. This behaviour was similar to that of the two independent, tuned TCP flows, where the TCP flow via Paris reacted quickly to the bandwidth relinquished by the flow on the congested Hamburg path.

Co-funded by the Horizon 2020
Framework Programme of the European Union

The slower reaction time of Hercules is likely due to its congestion control algorithm, PCC Allegro, which is known to react slowly to changes in network conditions. Through Hercules' utilisation of both paths, it was able to maintain a total minimum throughput of 7.5 Gbps during the congested period, and an overall throughput of 8.6 Gbps -- a feat that would not be possible in the TCP/IP network without extensive, manual network configuration.

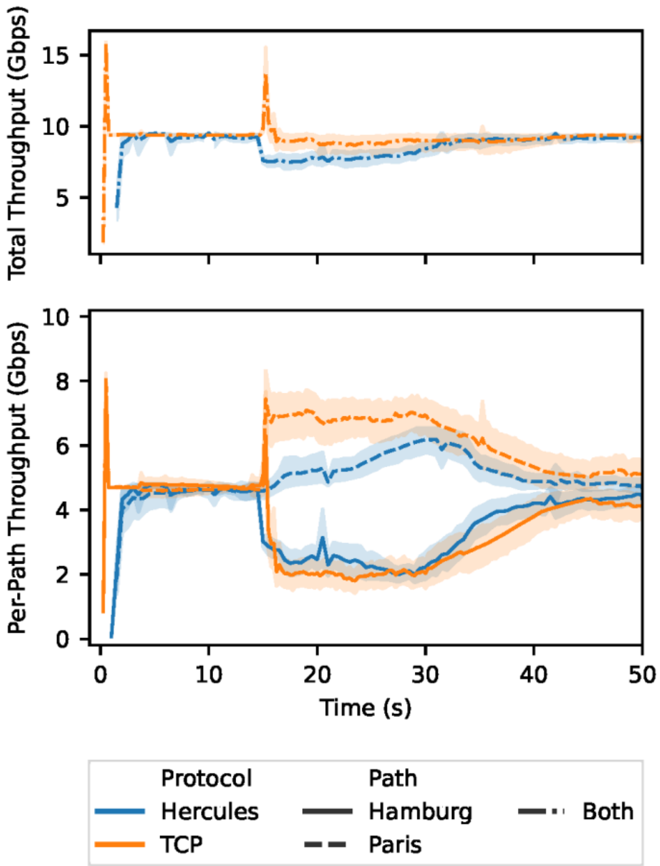

Figure 1: Per-path and total throughput of Hercules and two independent TCP flows running over GEANT.

Since SCION wasn't available in BRIDGES at the time when we needed to conduct the long-distance, intercontinental SCION experiments, we evaluated Hercules on the South Korean academic network KREONet2 [18] instead, to assess it under high bandwidth-delay-product environments, where TCP requires tuning to effectively utilise the bandwidth. KREONet2 interconnects with other research networks globally.

For comparisons with TCP, we performed similar M2M transfers using the iperf3 [19], bbcp [20], GridFTP [21] tools, as they enable high-performance M2M transfers with one (iperf3) or

Co-funded by the Horizon 2020 Framework Programme of the European Union

multiple parallel TCP flows (bbcp and GridFTP). Each transfer lasted for 45 s and was repeated at least 10 times. Throughout the evaluation, we refer to TCP as being either untuned or tuned. An experiment being run with untuned TCP refers to the system's TCP parameters remaining unmodified. This configuration may vary based on the system being evaluated, and reflects the default configuration of the operating system and its installed packages. An experiment being run with tuned TCP indicates that we manually configured the networking stack on the sending and receiving hosts to improve TCP throughput. This was done by following the guidelines set forth by the Energy Sciences Network (ESNET) [22]. Specifically, we increased the maximum send and receive TCP window sizes to 2 GiB to allow more in-flight data, switched the congestion control algorithm to H-TCP, enabled packet pacing on the sending interface, and limited bursts to 9 Gbps (on 10 Gbps links). While such changes significantly improve the performance of TCP, they require manual intervention by a system administrator and can lead to the exhaustion of system resources. GridFTP and bbcp were evaluated with tuned TCP, whereas iperf3 was evaluated with both tuned (iperf3+) and untuned TCP.

For each tool, we performed 10 transfers eastward from Hong Kong to Chicago, United States via Daejeon, South Korea; a distance of over 12 500 km with an RTT of 194 ms. Each endpoint was equipped with an Intel Xeon Silver 4114 CPU (2.20 GHz), 15.8 GiB of memory, and an Intel X710 10 Gbps Ethernet card.

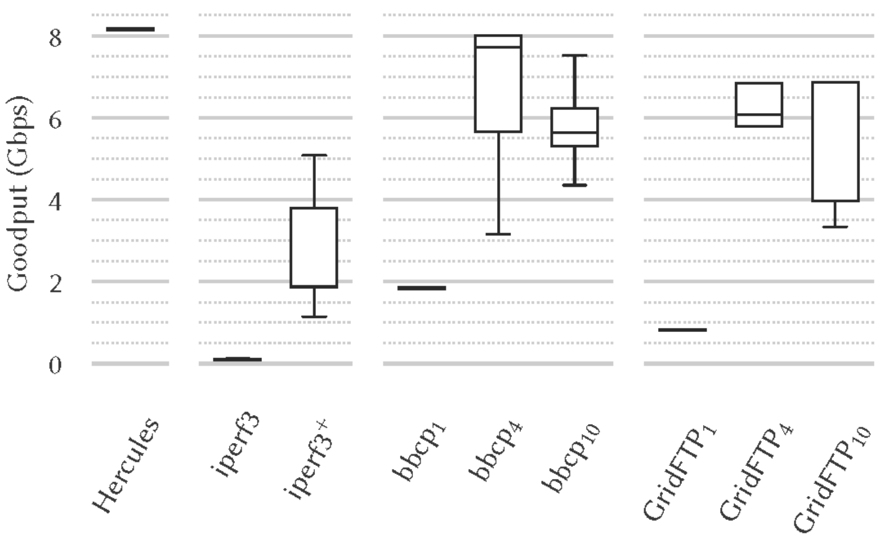

Figure 2: Goodput from Hong Kong to Chicago (10 Gbps, 194 ms RTT). Subscripts denote parallel TCP streams.

Co-funded by the Horizon 2020 Framework Programme of the European Union

Figure 2 shows the goodput in this setting. At 10 Gb*ps* and almost 200 ms RTT, KREONet2 has the highest bandwidth-delay product of our evaluation settings (around 242.5 MB) and TCP struggled to effectively utilise the network. Losses resulted in median transfer rates below 2 Gb*ps* for the single TCP connections (iperf3, bbcp, and GridFTP), thus requiring the use of multiple TCP connections ($bbcp_{4,10}$,$GridFTP_{4,10}$). In contrast, Hercules maintained above 8 Gb*ps*, thereby outperforming multiple competing TCP connections.

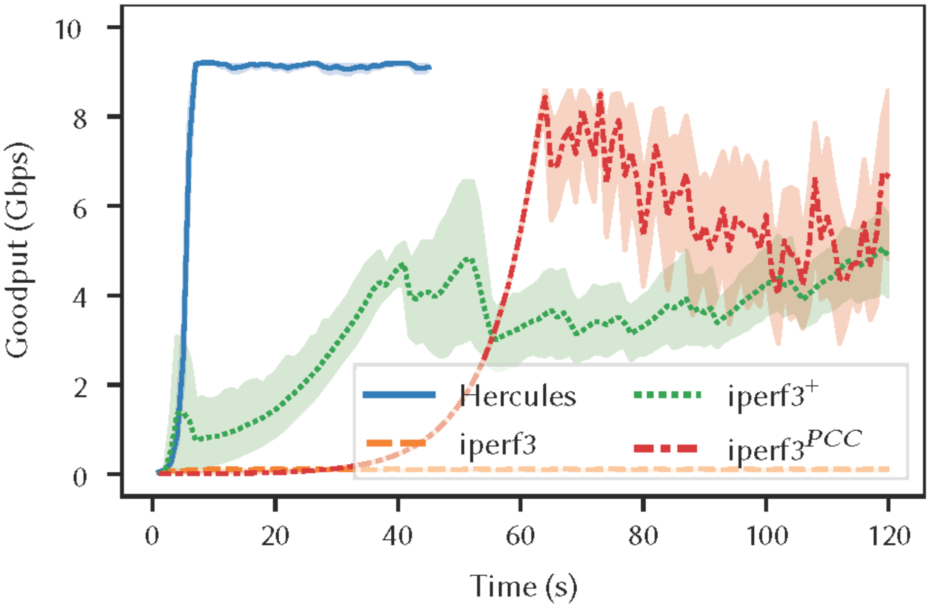

Figure 3: Intercontinental M2M goodput over time
from Hong Kong to Chicago (10 Gb*ps*, 194 ms RTT)

The change in goodput over time in Figure 3 offers an explanation for these differences. We increased the runtime for iperf3 to 120 s and incorporated iperf3 with PCC. Here we see that iperf3 with H-TCP required over 40 s to increase its transmission rate to around 5 Gb*ps*. Additionally, while iperf3 with PCC performed better, it still required over a minute to exceed 8 Gb*ps* before slowly decreasing towards 6 Gb*ps*. By contrast, within 10 s Hercules exceeded ~9 Gb*ps* and remained above it for the remainder of the transmission.

We will rerun these long-distance Hercules experiments over SCION across BRIDGES in order to confirm these results, as soon as we will have EU side connected to BRIDGES.

Co-funded by the Horizon 2020
Framework Programme of the European Union

4.1 Discussion and Analysis on Results

Hercules not only illustrates that high-speed transfers are viable on new Internet architectures, but also that they can provide new opportunities to the bulk-transfer space.
By building atop SCION, Hercules can exploit features such as the simultaneous use of multiple network paths, and can thus also quickly recover from link failures.

Our results based on real measurements provide a deep insight into the true potential of SCION in terms of privacy, reliability, and energy-efficiency enhancement by leveraging path diversity. Moreover, the results show that the user satisfaction can be further increased, while reducing the load on congested paths. In the future, the performed experiments will be complemented by long-term reliability measurements between the EU and US over the SCION network.

Based on the derived results, suitable deployment alternatives of SCION can be identified and compared, considering the network topologies both within and among different locations. This will allow us to further reduce delays and increase throughput within our SCIONLab network.

By analysing timing and throughput to different remote locations from both EU and US locations within the enhanced SCIONLab testbed, we are able to draw conclusions on the overall SCION multipath and path-awareness capabilities at a scale which would otherwise not be possible. Additionally, the resulting performance models gained from the enhanced SCIONLab network topology can be used to estimate the performance of the SCION network on a worldwide scale.

Overall, our experiments show how applications and end users are able to benefit from enhanced privacy, availability, reliability, and energy-efficiency on the SCION network. To this end, our project is also relevant with the following key NGI enabling technologies: 5G (or beyond), IoT, cybersecurity and resilience. For example, security is currently the biggest concern in adopting IoT technology (especially IoT adoption in enterprises). Threats include potential failures and attacks which hinder IoT adoption in critical infrastructures with high availability requirements (e.g. transportation or Grid infrastructures). Moreover, IoT devices often use weak authentication which may facilitate unauthorized access (e.g. smart home, healthcare devices) and use of unencrypted communication may leak privacy-sensitive data (e.g. healthcare).

Co-funded by the Horizon 2020
Framework Programme of the European Union

SCION mechanisms are able to satisfy many of those requirements: DRKey (Dynamically Recreatable Key) enables E2E encryption and supports scalable sender network authentication which prevents DDoS attacks and IP address spoofing attacks. EPIC (Every Packet is Checked) does per-packet source authentication which enforces, e.g. hidden paths to prevent exploitation of device vulnerabilities. Moreover, guaranteed low-bandwidth packet delivery with EQ and traffic filtering at high data rates with our Lightning Filter make SCION suitable for guaranteed access for command-and-control. Finally, multipath communication can complicate traffic analysis and geofencing can be done to avoid traffic leakage, which enhances privacy and compliance.

Besides, it is also possible to demonstrate how SCION can help to green the Internet by reducing the power consumption through appropriate path selection [17]. With the approach for routing current Internet traffic on the SCION network to optimize the CO2 impact, CO2 is already reduced. Even more dramatic, however, is that a virtuous cycle will arise where ISPs compete to lower their CO2 footprint and energy efficiency in order to increase traffic on their networks and consequently increase their commercial profit. We expect this virtuous cycle to reduce global CO2 footprint for networks and data centers by over 10%, which is tremendous given their power consumption exceeding 60 GW (525 TWh per year).

5 Present and Foreseen TRL

At this stage, the global SCIONLab network has been running successfully in a lab setup since 2017. SCIONLab has been deployed over multiple European-based research networks, including GEANT, SWITCH, DFN, SIDN, and FED4FIRE+ (VirtualWall), the current SCIONLab topology is depicted here: https://www.scionlab.org/topology.png. To this end, the SCIONLab network includes both native and overlay connections over the existing Internet. While the native deployments with production-grade SCION border routers (as now done in GEANT) can be considered TRL 7 or higher, the overlay-based deployments remain at TRL 4. For this reason, connections established over the SCIONLab network may still be influenced by other Internet traffic, making it difficult to perform repeatable experiments. While the European domain has already been quite well connected through dedicated L2 links over GEANT, the connection to the US domain so far lacked path diversity and high-bandwidth connections to US counterparts like Internet2, FABRIC, and others.

Therefore, the aim of this project was to interconnect SCIONLab with the BRIDGES platform. The BRIDGES facility consists of two geographically separate optical links spanning the North Atlantic and terrestrial optical links in US and Europe to form an intercontinental optical network ring topology, capable of carrying up to 200 Gbps of science data between major

Co-funded by the Horizon 2020 Framework Programme of the European Union

nodes in Washington, New York, Paris, and Amsterdam. The BRIDGES facility is a binding platform for research projects, providing a flexible research-oriented infrastructure connecting laboratories and universities in the US to their counterparts in Europe. With the deployment of SCION over BRIDGES it becomes now possible to connect the SCION ASes on the EU side (GEANT, SWITCH, SIDN, DFN, ETH, OVGU, etc.) with newly established SCION ASes on the US side (University of Virgina, Princeton, etc.) over multiple native SCION connections. These connections not only create interesting new path alternatives, but also raise the TRL across the Atlantic to a similar level as existing on the EU side.

To this end, the SCION deployment over BRIDGES significantly increases the path diversity and bandwidth between the EU domain and the US domain in SCIONLab, which is provided to experimenters by means of two geographically separate links over BRIDGES' infrastructure. Thus, the native SCION deployment across the Atlantic makes our enhanced infrastructure suitable for experimentation and application development in parts at levels beyond TRL 7.

Specifically, the transatlantic SCION network enables research on privacy and trust enhancing technologies, as well as scalable path discovery mechanisms across the network with more path opportunities between the US and Europe. Moreover, it also facilitates open Internet architecture renovation in areas that are difficult to evaluate on the current Internet such as: multipath communication, advanced and highly secure PKI systems, in-network DDoS defences, next-generation routing architecture policy definitions, path-aware applications, and path-based interdomain routing architectures.

The native SCION deployment across the Atlantic enables long-distance experiments that haven't been possible before. We expect that many more experiments will be run by researchers in the near future to evaluate those aspects across the newly extended SCION infrastructure, which will help to raise the TRL level of the overall SCION technology suite even higher.

6 Exploitation, Dissemination and Communication Status

This final report, which will be made public, details the deployment and measurement setup as well as the underlying assumptions. Furthermore, it documents the results and covers observations made based on the collected data. This report also includes a detailed documentation of the underlying deployment and measurement procedure. Based on this report, a paper presentation at an internationally renowned networking conference will be part to disseminate the results of the work further, now that the project work has officially completed.

Co-funded by the Horizon 2020 Framework Programme of the European Union

Additionally, we are also able to demonstrate the capability of the SCIONLab testbed by showing its reliability while deliberately interrupting certain connections within the SCIONLab network.

A number of dissemination activities have already taken place. Specifically, the SCION Components Analysis [9] has been presented during the PANRG meeting at IETF 114. Moreover, we gave a presentation about SCION during APNIC 54 [12]. Additionally, we have had a booth and pitched our work at a local technology transfer exhibition [13]. More recently, an update on the ongoing standardization work related to SCION was given during the RTGWG (Routing Area Working Group) and the PANRG meetings during IETF 115 [16].

We plan to disseminate the results from our experiments further in blog posts (e.g. within RIPE Labs) as well as within the future IETF and IRTF working groups (e.g. PANRG, TAPS).

7 Impacts

Our project has addressed and contributed to the impacts in relation to the NGI initiative as follows:

Impact 1: Enhanced EU – US cooperation in Next Generation Internet, including policy cooperation.

The rise of new network services and architecture proposals has led to the success of global network testbeds like PlanetLab [5] or the Peering testbed [6]. Unfortunately, current testbeds have shown shortcomings, e.g., related to the "security" of running your code or using your own configuration, requiring individual vetting by the testbed, the difficulty to evaluate DoS / security in real-world environments, and resource availability (especially computation). Moreover, these testbeds did not enable important networking aspects, including multipath routing, path-aware networking, security applications requiring per-AS certificates and cryptographic keys, secure routing, etc.

With the SCION deployment over BRIDGES and the experiments proposed in this project, we are able to significantly advance networking research through several experiments over the enhanced SCIONLab testbed across the Atlantic, which provides the possibility to run path-aware networking experiments enabled by the underlying SCION architecture.

Impact 2: Reinforced collaboration and increased synergies between the Next Generation Internet and the US Internet programmes.

Co-funded by the Horizon 2020 Framework Programme of the European Union

Our experiments meet NGI objectives such as privacy-enhancement by leveraging the SCION path diversity provided in a native fashion across the Atlantic. Additionally, experiments with our enhanced SCION-based path discovery mechanisms demonstrate the increased throughput, reliability, and energy-efficiency that can be achieved based on the SCION Internet architecture, thanks to a better utilization of network resources and through SCION's native multi-path property.

These impacts are well aligned with the objectives of our US partner in this project, GMU, who receives funding from NSF under BRIDGES - Binding Research Infrastructures for the Deployment of Global Experimental Science (Award Number: 2029221). NSF BRIDGES serves two key objectives: First, it is a prototype and demonstrator of a fully virtualized cyber-infrastructure architecture in support of future global science applications and advanced networked services such as SCION. Second, BRIDGES links European research facilities directly to US research facilities by constructing a 100 Gbps research missioned network ring spanning the North Atlantic. This BRIDGES facility has the explicit purpose to facilitate collaborative global experiments across a common, contiguous, seamless and fully federated network research infrastructure.

Impact 3: Developing interoperable solutions and joint demonstrators, contributions to standards.

There have already been a number of efforts in standardizing SCION at the IETF, e.g. within PANRG (Path Aware Networking Research Group) and TAPS (Transport Services) working groups. Specifically, Internet Drafts on the SCION overview [7], the SCION Components Analysis [9] and on DRKey (Dynamically Recreatable Key) [10] have been written up. Additionally, a presentation on PANAPI (Path Aware Networking Application Programming Interface) [11] has been given at a recent IETF TAPS meeting. The results of our experiments will also partially be fed into these ongoing standardization efforts.

More recently, SCION was also present at the IETF meeting 115 in London. Specifically, an overview on the ongoing standardization work related to SCION was given during the RTGWG (Routing Area Working Group) and the PANRG meetings and the progress with respect to these efforts was highlighted [16]. In the next steps, the received feedback will be addressed and the existing Internet Drafts on SCION will be improved. Additionally, there will be also new drafts documenting current specifications on the SCION control plane and the SCION data plane.

Co-funded by the Horizon 2020
Framework Programme of the European Union

Impact 4: An EU - US ecosystem of top researchers, hi-tech start-ups / SMEs and Internet-related communities collaborating on the evolution of the Internet

The rapid growth of the Internet is driving the emergence of various new network services (e.g., IoT) at a global scale. The limitations of the existing Internet architecture towards those new requirements has significantly increased the demand for advanced wired networking architectures overcoming those limitations.

SCION is a novel secure Internet architecture which aims to provide route control, failure isolation, and explicit trust information for end-to-end communications. As discussed in detail in the original SCION paper [4], no existing solution simultaneously provides the capabilities as SCION does, although individual aspects are covered by these efforts, which SCION builds upon. SCION has already lead to a new start-up, Anapaya Systems and a quite large SCION community, especially in Europe but also worldwide.

With our joint EU-US SCION setup and experiments over NSF BRIDGES these efforts are strengthened further. Researchers directly benefit from the major performance improvement due to the high bandwidth and low delay capabilities over BRIDGES. Furthermore, users also benefit from additional paths within SCIONLab passing across several transatlantic links. This results in a significantly improved SCIONLab network between the US and Europe to run sophisticated SCION experiments, which will give a more detailed insight into the SCION infrastructure deployed on native connections at a wide scale. To this end, we will also further disseminate the results of our experiments as outlined in Section 6 above.

8 Conclusion and Future Work

The objective of this project was to deploy SCION over the NSF BRIDGES infrastructure in order to demonstrates the SCION benefits by means of experiments across the Atlantic. To this end, the redundant BRIDGES network with dual 100G links enables researchers to select and use multiple paths across the Atlantic as facilitated by SCION. While we have successfully deployed three SCION servers within BRIDGES and six SCION servers with GEANT with native SCION links in the respective networks, the actual connection of BRIDGES with GEANT in order to exchange SCION traffic remains part of our future work, as our deployment will be further expanded. Given the huge number of connection opportunities on both sides, connecting BRIDGES with GEANT will be of great benefit for SCION research.

Co-funded by the Horizon 2020 Framework Programme of the European Union

To demonstrate these benefits, we have conducted a number of experiments with the Hercules bulk-data transfer tool over SCION. Hercules provides an easily-deployable but highly performant protocol and is useable on public networks thanks to employing existing congestion control algorithms. Furthermore, by building atop SCION, it can exploit features such as the simultaneous use of multiple network paths, and can thus also quickly recover from link failures. Despite being a user-space protocol built upon an emerging Internet architecture, Hercules achieves high transmission rates. Hercules's performance with a single network flow surpasses that 1, 4, or 10 parallel TCP flows in high BDP networks and approaches that of tuned single-flow TCP implementations in lower BDP environments, without requiring the network-stack tuning of TCP. Hercules and the SCION Internet architecture underpinning it offer a promising platform for executing bulk transfers and illustrate that high-speed bulk transfers are not only viable in new Internet architectures, but that these architectures can contribute greatly to advancements in the bulk-transfer space.

In future work, we expect to perform many more long-distance SCION experiments across the Atlantic at scales that hadn't been possible before. The new insights from the results will help to increase the performance and maturity of the overall SCION technology suite even higher than it already is today.

9 References

[1] David Barrera, Laurent Chuat, Adrian Perrig, Raphael M. Reischuk, Pawel Szalachowski. The SCION Internet Architecture. Communications of the ACM, June 2017, Vol. 60 No. 6, Pages 56-65. https://doi.org/10.1145/3085591

[2] SCIONLAB: A Next- Generation Internet Testbed, ICNP 2020. *Best Paper Award*. (https://netsec.ethz.ch/publications/papers/icnp2020_scionlab.pdf)

[3] Deployment and Scalability of an Inter-Domain Multi-Path Routing Infrastructure. CoNEXT 2021. *Best Paper Award*. (https://netsec.ethz.ch/publications/papers/2021_conext_deployment.pdf)

[4] SCION: Scalability, Control, and Isolation On Next-Generation Networks, https://www.scion-architecture.net/pdf/2011-SCION.pdf

[5] PlanetLab: http://www.planet-lab.org/

[6] PEERING: The BGP Testbed, https://peering.usc.edu/

Co-funded by the Horizon 2020
Framework Programme of the European Union

[7] Corine de Kater, Nicola Rustignoli, Adrian Perrig: SCION Overview. PANRG Internet Draft, https://datatracker.ietf.org/doc/draft-dekater-panrg-scion-overview/ Last updated 2022-08-26

[8] Chuat, L., Legner, M., Basin, D., Hausheer, D., Hitz, S., Mueller, P., and A. Perrig, "The Complete Guide to SCION", ISBN 978-3-031-05287-3, 2022, https://doi.org/10.1007/978-3-031-05288-0

[9] Nicola Rustignoli, Corine de Kater: SCION Components Analysis. PANRG Internet Draft, https://datatracker.ietf.org/doc/draft-rustignoli-panrg-scion-components/ Last updated 2022-08-26

[10] Juan A. García Pardo Giménez de los Galanes, Cyrill Krähenbühl, Benjamin Rothenberger, Adrian Perrig: Dynamically Recreatable Keys. PANRG Internet Draft, https://datatracker.ietf.org/doc/draft-garciapardo-panrg-drkey/ Last updated 2022-07-24

[11] Thorben Krüger: PANAPI, a Path-Aware Networking API. IETF TAPS Meeting, March 23, 2022 https://datatracker.ietf.org/meeting/113/materials/slides-113-taps-panapi-implementation-00

[12] David Hausheer, Nicola Rustignoli: The SCION Inter-domain Routing Architecture. APNIC 54, Sept 2022. https://conference.apnic.net/54/assets/files/APSG129/scionapnic22_1663125162.pdf

[13] David Hausheer et al. SCION-LINC: Zuverlässige Kommunikation für industrielle Steuerungssysteme (in German). https://www.transfermesse.ovgu.de/Tagesprogramm/Technologieangebote/Elektrotechnik/SCION_LINC_+Zuverl%C3%A4ssige+Kommunikation+f%C3%BCr+industrielle+Steuerungssysteme-p-9574.html, Sept 2022.

[14] Mo Dong, Qingxi Li, Doron Zarchy, P. Brighten Godfrey, Michael Schapira: PCC: Re-architecting Congestion Control for Consistent High Performance. USENIX NSDI 15.

[15] Allcock, W., Bresnahan, J., Kettimuthu, R., Link, M.: The Globus Striped GridFTP Framework and Server. SC '05: Proceedings of the 2005 ACM/IEEE Conference on Supercomputing.

[16] Corine de Kater, Nicola Rustignoli: SCION Update and Q&A. IETF 115 – RTGWG, Nov 9, 2022 https://datatracker.ietf.org/meeting/115/materials/slides-115-rtgwg-sessb-7-scion-00

[17] Seyedali Tabaeiaghdaei, Adrian Perrig: Data-Plane Energy Efficiency of a Next-Generation Internet Architecture. ISCC 2022, October 2022 https://netsec.ethz.ch/publications/papers/2022_iscc_scion_energy.pdf

[18] KREONET. Mar. 7, 2022, Available: https://www.kreonet.net/

[19] J. Dugan, S. Elliott, B. A. Mah, J. Poskanzer, and K. Prabhu: iperf3. Version 3.1.3, Jun. 8, 2016 http://software.es.net/iperf/

Co-funded by the Horizon 2020 Framework Programme of the European Union

[20] Stanford University: bbcp, version 15.02.03.00.1, Feb. 2015. https://www.slac.stanford.edu/~abh/bbcp/

[21] W. Allcock, J. Bresnahan, R. Kettimuthu, and M. Link: The Globus striped GridFTP framework and server. In SC ’05: Proceedings of the 2005 ACM/IEEE Conference on Supercomputing, Nov. 2005, pp. 54–54. doi: 10.1109/SC.2005.72.

[22] The Energy Sciences Network (ESnet). Host tuning - Linux tuning. (Sep. 3, 2021). https://fasterdata.es.net/host-tuning/linux/

10 Glossary

SCION	Scalability, Control, and Isolation On Next-Generation Networks
OVGU	Otto-von-Guericke-University
NGI	Next-Generation Internet
NSF	National Science Foundation
GMU	George Mason University
AS	Autonomous System
ISP	Internet Service Provider
DRKey	Dynamically Recreatable Key
E2E	End-to-End
EPIC	Every Packet is Checked
IP	Internet Protocol
DDoS	Distributed Denial-of-Service
BRIDGES	Binding Research Infrastructures for the Deployment of Global Experimental Science
IoT	Internet-of-Things
PANRG	Path Aware Networking Research Group
TAPS	Transport Services
RTGWG	Routing Area Working Group
PANAPI	Path Aware Networking Application Programming Interface

Co-funded by the Horizon 2020 Framework Programme of the European Union