



Deployment and Evaluation of the SCION Secure Internet Architecture on Fed4FIRE+ Testbeds

**Johannes Wünsche, Fin
Christensen, Marten
Gartner, David Hausheer**
OVGU Magdeburg, Germany

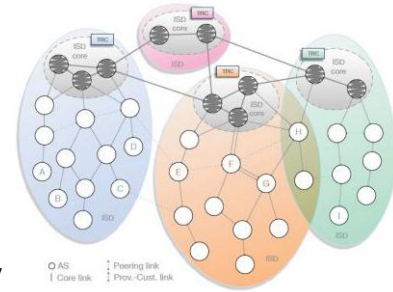
Review Open Call 5 Experiments

Virtual Meeting, 4 June 2021

Experiment Description



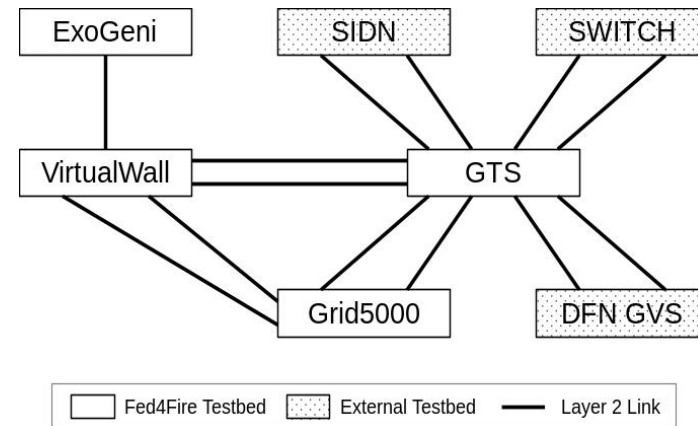
- ❖ SCION: A novel Next-Generation Internet (NGI) architecture designed to provide scalability, control, and isolation
 - Path-awareness: Control over forwarding paths
 - Multipath forwarding
 - Path optimization: high bandwidth, low latency, high availability
 - Operational SCION network since 2014 (ISPs, NRENs, banks)
 - 5th generation of software → Mature for large-scale deployment
- ❖ SCIONLab: A Next-Generation Internet testbed
 - Enables research that is difficult to evaluate on the current Internet: path-aware networking, multipath, in-network DDoS defenses etc.
 - Fast setup of a SCION AS in a VM, requiring little technical expertise, SCION AS can be attached anywhere in the SCIONLab network
 - Mostly based on overlay connections over the existing Internet



Experiment Description



- ❖ Goal: Large-scale SCION deployment and evaluation across multiple Fed4FIRE+ testbeds (GTS, VirtualWall, Grid5000, ExoGeni)
 - Analyse inter-testbed connectivity of SCION in terms of performance on top of a dedicated network unbiased by other Internet side traffic
 - Connect Fed4FIRE+ testbeds with SCIONLab to enable experiments using SCION in the interconnected testbeds
 - SCION becomes a base network substrate in Fed4FIRE+, with the possibility to establish (multipath) connections into other (yet isolated) testbeds



Experiment Description

❖ GTS

- 3x BMS (AMS, PAR, HAM)
- 10 Gbit/s internal links
- VLANs to DFN, SWITCH, SIDN

❖ VirtualWall

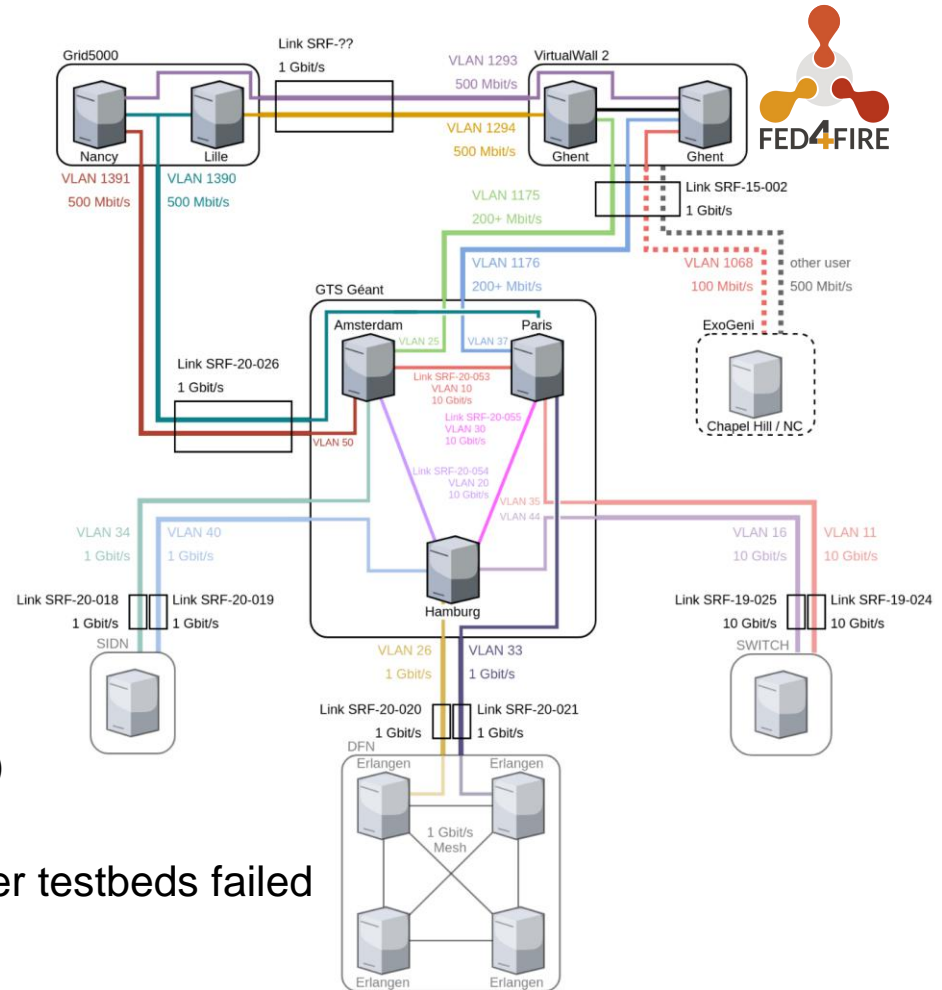
- 2x BMS (Ghent)
- 10 Gbit/s internal link
- 2x VLANs to GTS (1 Gbit/s shared)

❖ Grid5000

- 2x BMS (Nancy/Lille)
- 2x VLANs to GTS (1 Gbit/s shared)
- 2x VLANs to VWall (1 Gbit/s shared)

❖ ExoGeni

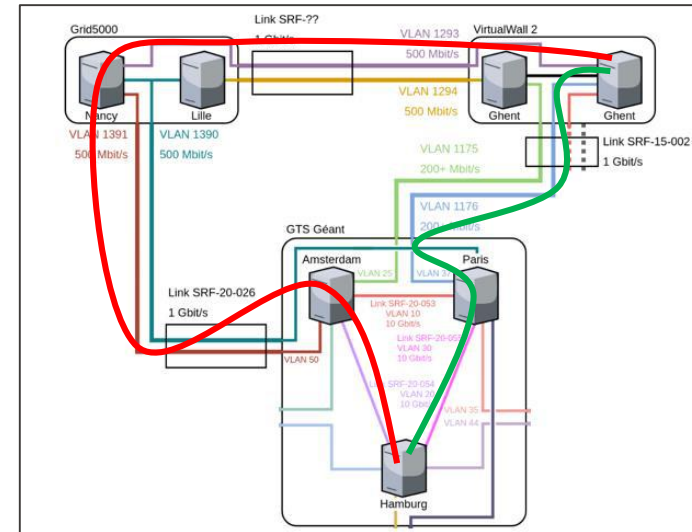
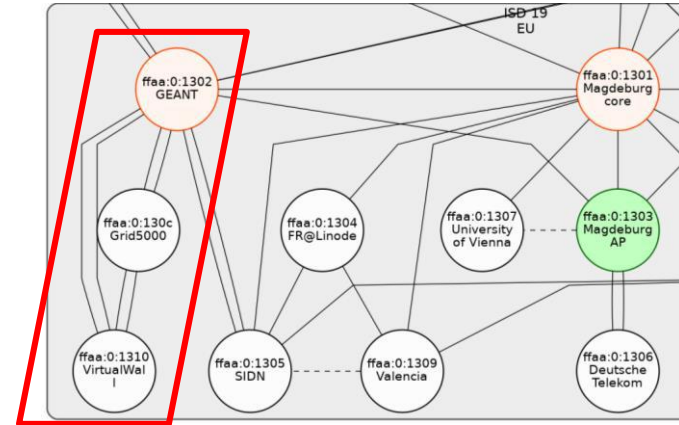
- Access problems, connection to other testbeds failed
- Support unresponsive



Experiment Description



- ❖ Every testbed forms dedicated AS
 - GTS forms a SCION core AS
 - 6 possible paths between GTS, VirtualWall, and Grid5000
 - GTS: High-speed SCION border router
 - VirtualWall, Grid5000: Open source border router
- ❖ Experiment:
 - Bandwidth aggregation over multiple paths between GTS and VirtualWall
 - Spate: A benchmarking tool to send and receive SCION traffic over multiple paths
 - Collect Prometheus metrics exposed by SCION border routers
 - Latency measurements

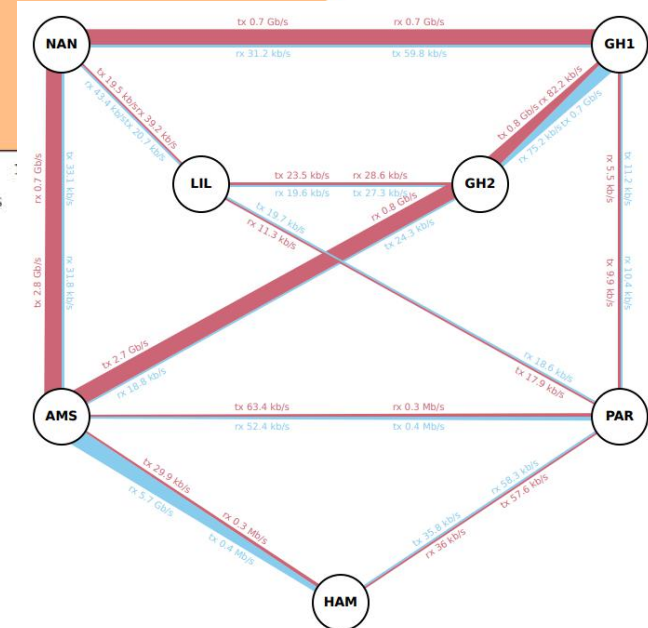
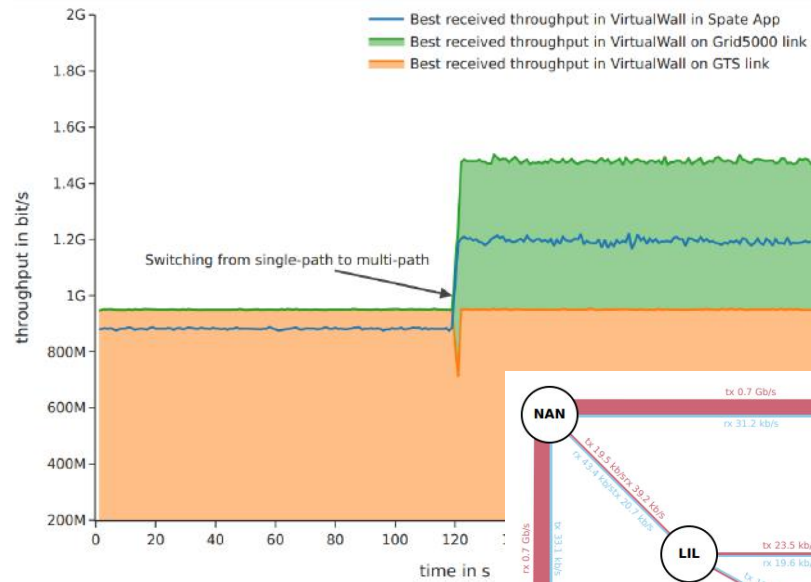


Project Results



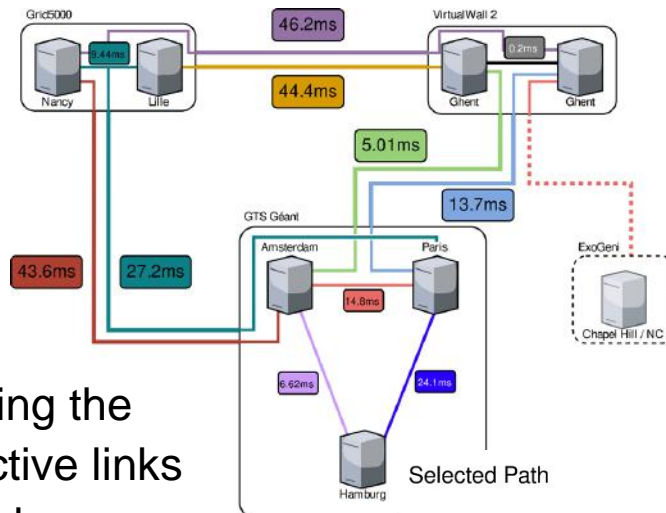
❖ Bandwidth aggregation results:

- Nearly 1 Gbit/s from Hamburg to Ghent, over direct path via AMS
- Multipath mode: Approx. 1.5 Gbit/s in total on the border routers
- Potential bottleneck related to open source border router and VLAN setup in Grid5000
- Slightly lower performance in Spate (1.2 Gbit/s) due to bottleneck in SCION endhost



Project Results

- ❖ Latency results:
 - SCION results in nearly similar latency compared to aggregating the delays over the respective links
 - Small latency overhead introduced by SCION
 - Path via Lille shows lower latency but also lower bandwidth



Round Trip Time

Hamburg	→	Amsterdam	→	Ghent 2	12.423ms						
Hamburg	→	Paris	→	Ghent 1	→	Ghent 2	38.765ms				
Hamburg	→	Paris	→	Lille	→	Ghent 2	96.421ms				
Hamburg	→	Amsterdam	→	Nancy	→	Ghent 1	→	Ghent 2	97.495ms		
Hamburg	→	Amsterdam	→	Nancy	→	Lille	→	Ghent 2	104.65ms		
Hamburg	→	Paris	→	Lille	→	Nancy	→	Ghent 1	→	Ghent 2	107.958ms

Project Results



❖ Demo Video



❖ Lessons Learned

- Grid5000: not possible for us to reserve resources on a longer term.
- Implemented Ouigo: A management script to automatically provision and migrate resources for SCION in Grid5000
- Nevertheless, experiment extension failed in many cases
- SCION configuration is challenged by heterogenous networks and interconnected research testbeds

Business Impact

- ❖ Great insight into performance of SCION across a large-scale wired network as provided by Fed4FIRE+
- ❖ Analysis of SCION latency and throughput on different paths to different destinations within Fed4FIRE+ testbeds
- ❖ Conclusions on overall SCION multipath and path-awareness capabilities can be drawn, at a scale not possible otherwise
- ❖ Results show performance potential of SCION on a real network
- ❖ Performance improvement for SCIONLab users thanks to high bandwidth and low delay capabilities over Fed4FIRE+ testbeds

Business Impact

- ❖ Users can benefit from additional paths within SCIONLab passing across Fed4FIRE+ testbeds, resulting in higher reliability
- ❖ User satisfaction can be increased, while reducing the load on congested paths
- ❖ Insight on suitable deployment alternatives of SCION in Fed4FIRE+
- ❖ Fed4FIRE+ testbeds provided excellent environment for our experiment to assess SCION's advanced wired networking capabilities

Business Impact



- ❖ Value perceived:
 - Gained knowledge about deploying and running SCION in a large-scale testbed environment
 - Learning of new competences
 - Analysis of SCION performance in its capability of filling gigabit layer links over larger time periods
 - Robustness of SCION is in volatile research testbeds
 - How to do heterogenous network deployment management
 - Experience in setting up, working with and connecting several testbeds
 - Experience in managing testbeds via different APIs

Business Impact



- ❖ Value perceived (cont'd):
 - Deployment of high-speed SCION border router in real-world testbeds revealed potential to improve deployment strategy
 - Benchmarking results show performance of SCION open-source and high-speed border routers and SCION endhosts
 - Without Fed4FIRE+, the deployment of SCION AS on a large-scale infrastructure would come along with increased complexity and effort
 - Redundant 1 Gbit/s links provided between Fed4FIRE+ testbeds are very useful
 - Implemented benchmarking tools will be used for future SCION experiments

Feedback



❖ Used resources:

Wired networking testbeds		Used?	Specify the type and amount of the resources used
	Virtual Wall (imec)	yes	2 bare metal servers, 4 Network Edges
	Geant Testbed as a Service (GTS) (Nordunet)	yes	3 bare metal servers, 7 VLANs

Cloud computing testbed		Used?	Specify the type and amount of the resources used
	Exogeni (UvA)	no	We were not able to use ExoGeni for our experiment, due to frequent crashes
	Grid5000 (Inria)	yes	2 bare metal servers, 4 KaVLANs

❖ Used tools:

Tools	Used?	Please indicate your experience with the tools. What were the positive aspects? What didn't work?
JFed	Yes	Easy to use. Problems arose with the usage of VLANs as some were outdated initially. Without further knowledge of the backbone network, the “stitched” links between different sites are quite hard to debug, more transparent error messages would have saved time here. GTS integration did not work, we switched to using the normal GTS interface outside of jFed.
GTS	Yes	Quite easy to use, but no link management was possible.
Grid5000	Yes	More complex setup with heavy focus on wiki and examples. We managed to setup the topology and the duration of it as we planned but the resource allocation policy appears not to be suitable for long-running experiments.

Feedback



❖ Procedure / Administration

- Waiting time on issues of bugs was higher than expected
- Delay due to problems regarding connections between testbeds
- Great organization of calls to manage the issues
- Implementation of found solutions took some time

❖ Setup of the experiment

- jFed integrated calling of setup scripts not used, since we needed to setup experiment configurations on the machines
- Without the activation of TLS lower than v1.3 connection was not possible
- Grid5000: high effort to implement automation script to setup machines and connect them to other testbeds
- In the end we were required to deal with each testbed separately

Feedback



❖ Fed4FIRE+ portfolio

- We only used computation and networking capabilities
- Problems in the application of network configurations to interconnect testbeds
- Connections to GTS were outdated, and indicated errors were cryptic
- Management of all resources via Fed4FIRE+ could not be done due to limitations in functionality of the provided tools
- Grid5000: stable long term experiment not possible, since policy on the usage of the testbed differs greatly from other testbeds in Fed4FIRE+
- Interoperation often failed when creating a new link between testbeds, mostly due to deprecated VLAN configurations or switch failures

❖ Documentation and support

- Requirement of third-party authentication to access the documentation

Feedback



❖ Experiment environment

- Concerns due to access tokens to machines being encoded in the URL of the VM spice interface
- In some instances, resources were not directly freed after we have ended experiments prematurely
- Resources, if scarce beforehand, were available only after the pre-allocated time of the experiment had ended

❖ Experiment execution and results

- Configuration of complete Fed4FIRE+ infrastructure not possible from jFed
- Four different specific tools and manual intervention to modify testbeds
- Major hurdles in interconnections and cross-testbed experiments: outdated configurations and non-existent stitching routes (e.g. Grid5000)

Feedback

- ❖ What is missing from your perspective?
 - A true inter-testbed connection manager
 - Tools of different testbeds had to be used
 - Not possible to depend on the information given as advertised
 - More high capability links would be nice

- ❖ Thanks to the experiment I conducted within Fed4FIRE+ ...
 - ... we discovered multiple performance bottlenecks in our tested software and gained knowledge on how to reduce the impact of those. And we worked together with the Fed4FIRE+ team to find problems in the creation of inter-testbed connections.



Co-funded by the
European Union



Co-funded by the
Swiss Confederation

This project has received funding from the European Union's Horizon 2020 research and innovation programme, which is co-funded by the European Commission and the Swiss State Secretariat for Education, Research and Innovation, under grant agreement No 732638.

**THANK YOU FOR
YOUR ATTENTION**

WWW.FED4FIRE.EU