



Conference on Networked Systems 2021
(NetSys 2021)

Towards SCION-enabled IXPs: The SCION Peering Coordinator

Lars-Christian Schulz and David Hausheer

4 pages

Towards SCION-enabled IXPs: The SCION Peering Coordinator

Lars-Christian Schulz¹ and David Hausheer²

Networks and Distributed Systems Lab
Otto-von-Guericke-University Magdeburg, Germany
Email: ¹lschulz@ovgu.de, ²hausheer@ovgu.de

Abstract: Internet eXchange Points (IXPs) around the world bring thousands of ISPs together to form dense peering fabrics. Since bilateral BGP peering sessions alone would result in large overhead, many IXPs offer route servers enabling the exchange of routing information with the entire peering population over a single multilateral BGP session. Route servers also perform RPKI validation to combat the lack of authentication and security in BGP. SCION is a novel inter-domain routing architecture addressing the security flaws of BGP by replacing it with a security and reliability centric clean-slate approach. We envision that operators of SCION ASes will be just as open to peering at Internet exchanges as they are today with BGP. Moreover, to fully utilize SCION's multipath capabilities SCION AS operators tend to deploy more different AS numbers than in BGP, further increasing the potential number of unique peering links at an IX. Since SCION has no native multilateral peering support, we propose the SCION peering coordinator, an IXP-hosted service automating SCION peering link setup based on per AS policies. As such, the SCION peering coordinator provides an open peering platform similar to BGP route servers to SCION.

Keywords: SCION, IXP, peering, route server

1 Introduction

Route servers at IXPs solve two important problems. First, they improve BGP's scalability by replacing densely meshed bilateral peering with hub-based multilateral peering. Secondly, they filter out invalid prefix announcements and perform RPKI validation to prevent network outages due to misconfigured border routers and malicious prefix hijacking. SCION, which aims to replace BGP in a future Internet, addresses the latter issue with its fully authenticated control plane, but does not provide a built-in solution to the scalability concerns of fully meshed peering. The notion of multilateral peering does not exist in SCION. Since SCION's control plane is designed to be very efficient, full mesh peering might not pose a scalability issue like it does in BGP, but multilateral peering at BGP route servers has another advantage. Route servers offer instant connectivity to a large part of the IXP's peering population to anyone joining [RSF⁺14]. SCION lacks a mechanism to get connected quickly and easily at IXPs. All peering links have to be manually negotiated and set up by human operators. The aim of our proposed peering coordinator is to provide an automatic AS configuration service to SCION ASes joining a SCION-enabled IXP. The peering coordinator accepts a set of policy rules, which are set by each connected AS, and takes care of configuring the SCION border routers to form rule conforming links.



The policies the peering coordinator supports are inspired by BGP prefix redistribution rules as made available by special BGP community attributes at some IXP route servers. With BGP communities, peers can limit the distribution of their prefixes to certain ASes or geographical regions [DEC, AMS]. Similarly, the peering coordinator can limit peering advertisements to certain ASes or administrative domains.

For IXPs, SCION with a peering coordinator offers two direct advantages over BGP. First, the need for RPKI validation at the IXP is eliminated, because SCION's control plane is authenticated by default. Second, SCION is a native multipath protocol, whereas a BGP route server advertises only a single route per destination and peer. The ability to expose multiple paths through the IXP opens up unused internal link capacity, for example from backup links, increasing the bandwidth available to SCION over BGP.

2 Background and Related Work

SCION is an inter-domain network architecture aiming to improve the Internet's scalability, reliability and security [BCP⁺17]. It achieves isolation of administrative concerns on a large scale by grouping ASes into *isolation domains* (ISDs), which fall under a common jurisdiction. An ISD is administered by its core, a consortium of densely interconnected *core ASes*. SCION's path discovery process based on flooding so called *path-segment construction beacons* (PCBs) through the network gives rise to three distinct inter-AS link types. *Core links* connect core ASes of the same or different ISDs. *Peering links* connect non-core ASes of the same or different ISDs, but in contrast to core links do not carry beacons. Finally, *provider-customer links* attach non-core ASes to their non-core transit providers and ultimately the ISD core.

To enable peering irrespective of whether an AS belongs to the core or not, the peering coordinator is able to negotiate all three types of links. *Core* and *peering* links allow ASes of the same type to peer, but are forbidden between ASes of different types by the current SCION specification. However, in the special case of peering between a core and non-core AS within the same ISD, a provider-customer link can substitute for a native peering link. Operators of core ASes connecting to a peering coordinator should keep in mind that peering in the core is transitive by default, and might want to implement a more restrictive peering policy than in non-core ASes.

Originally, the peering coordinator was designed as an extension to the SCIONLab coordinator, which is a web service maintaining a central configuration repository for the SCIONLab research network [KGL⁺20]. We decided to fork the peering coordinator as a standalone application for the following reasons: (1) We expect that IXPs want full control over their own coordinator instance. Extending the SCIONLab coordinator would put control over the peering coordinator in the hands of the SCIONLab administrators. (2) The SCIONLab coordinator manages the configuration of all AS-internal SCION services. Such deep access is unnecessary for the peering coordinator. (3) A separate peering coordinator should simplify deployment in private networks.

3 SCION Peering Coordinator

The peering coordinator is designed to alleviate the lack of multilateral peering in SCION by automating the discovery of peers. Upon joining an IXP, a connection to one or, for redundancy

reasons, multiple instances of the coordinator must be set up manually. After that, the coordinator will mediate the automatic creation of peering links between the connected ASes. AS operators can limit to which other ASes links are formed by defining a set of filtering rules.

More specifically, the peering coordinator manages two pieces of information about every connected AS: (1) The set of ASes the AS in question offers peering to. (2) The IP address and UDP port range reserved for peering links in the SCION underlay of every border router connected to the IXP. The later piece of information is required for technical reasons since SCION is currently implemented as an IP/UDP overlay. For the actual link establishment, the peering coordinator's only responsibility is to pick unused UDP ports in the underlay at the two border routers forming the link's endpoints, allowing for a fairly unsophisticated implementation.

The set of ASes to whom peering is offered is constructed from a set of peering rules installed at the coordinator by every AS. In the common case of completely open peering like at a BGP route server, the only rule an AS has to set at the start of its peering session is "peer with everyone". Peering rules in the coordinator both offer peering to other ASes and accept offers from other ASes, so an AS with a default "peer with everyone" rule will have links to as many other ASes participating at the peering coordinator as possible. In addition to open peering, the peering rules can express a limited set of selective peering policies.

Selective Peering We define four kinds of accept/reject rules, with a scope of either any AS (the default rule), all ASes of an ISD (often coincident with a geographical region), all ASes owned by the same entity (company, etc.), or individual ASes. More specific rules have precedence over less specific ones. We introduce rules based on the AS owner, because SCION AS operators are likely to employ multiple logically distinct ASes to better exploit the path control offered by SCION. Selecting potential peer ASes grouped by the underlying administrative domain seems like an obvious use case for the coordinator's filtering rules. However, SCION itself does not provide owner information, so the mapping from AS to owner must be maintained out-of-band.

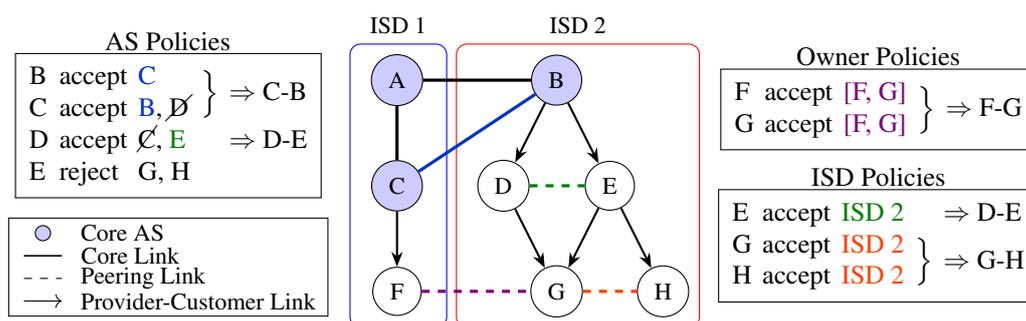


Figure 1: Example topology. ASes A to H all peer at an instance of the SCION peering coordinator. The black links are fixed (e.g., private peering or non-IXP links). Peering links negotiated by the peering coordinator and the corresponding AS-, Owner-, and ISD-level rules are colored.

As mentioned previously, open peering is achieved by setting the default rule to accept. Assuming the default rules is set to reject instead, Figure 1 illustrates a few example policies expressible in the coordinator. In the example, AS E accepts peering with everyone in its own ISD (ISD 2),



but disallows peering with G and H, which are its customers. Since AS-level rules have higher priority than less specific ISD-level policies, only the peering link to D (dashed) is created. As mentioned in [Section 2](#), due to limitations in SCION's current design, it is not possible to connect core and non-core ASes from different ISDs like C and D, consequentially such pairs are ignored.

Implementation Due to its legacy as an extension to the SCIONLab coordinator, the peering coordinator is a Django web application [[SPC](#)]. It provides a gRPC API to the peering configuration client which is installed alongside all SCION border routers connecting to the IXP. The purpose of the peering client is to perform the actual configuration changes needed to instruct the SCION infrastructure services to establish or drop links. Additionally, the coordinator serves a web interface for monitoring and management purposes.

4 Preliminary Conclusions and Future Work

The SCION peering coordinator brings the ease of public BGP peering at an IXP-hosted route server to SCION's clean-slate Internet routing approach. It supports the basic use case of open peering as well as a set of selective peering rules modeled after the prefix distribution control offered by BGP communities.

A remaining challenge with the peering coordinator is that it operates completely out-of-band. Like the SCIONLab coordinator, the peering coordinator aims at getting SCION interconnection started today, but is not integrated with the SCION control plane. In future work, we will investigate the possibility of an in-band neighbor discovery protocol in SCION which could supersede the peering coordinator. Another topic of future interest are the questions posed to the economics of public peering by SCION's inherent path control and multipath capabilities.

Bibliography

- [AMS] AMS-IX Route Servers. <https://www.ams-ix.net/ams/documentation/ams-ix-route-servers>.
- [BCP⁺17] D. Barrera, L. Chuat, A. Perrig, R. M. Reischuk, P. Szalachowski. The SCION Internet Architecture. In *Communications of the ACM*. Volume 60(6), pp. 56–65. 2017.
- [DEC] DE-CIX: Operational BGP Communities. <https://www.de-cix.net/en/resources/route-server-guides/operational-bgp-communities>.
- [KGL⁺20] J. Kwon, J. A. García-Pardo, M. Legner, F. Wirz, M. Frei, D. Hausheer, A. Perrig. SCIONLAB: A Next-Generation Internet Testbed. In *IEEE International Conference on Network Protocols*. 2020.
- [RSF⁺14] P. Richter, G. Smaragdakis, A. Feldmann, N. Chatzis, J. Boettger, W. Willinger. Peering at Peerings: On the Role of IXP Route Servers. In *Proceedings of the Internet Measurement Conference*. Pp. 31–44. ACM, 2014.
- [SPC] Peering Coordinator. <https://github.com/netsys-lab/scion-peering-coordinator>.