Workgroup: Path Aware Networking RG
Internet-Draft:
draft-rustignoli-panrg-scion-components-00
Published: 11 July 2022
Intended Status: Informational
Expires: 12 January 2023
Authors: N. Rustignoli C. de Kater
ETH Zuerich ETH Zuerich

SCION Components Analysis

Abstract

SCION is a future Internet architecture that focuses on security and availability. Its fundamental functions are carried out by a number of components.

This document illustrates the dependencies between its core components and extensions. It also discusses the relationship between SCION and existing protocols, with focus on illustrating which existing protocols are reused or extended. Additionally, it describes the motivations behind cases where a greenfield approach is needed, and the properties that can be achieved thanks to it. It then briefly touches on the maturity level of components.

About This Document

This note is to be removed before publishing as an RFC.

The latest revision of this draft can be found at https://scionassociation.github.io/scion-components_I-D/draft-rustignoli-panrg-scion-components.html. Status information for this document may be found at https://datatracker.ietf.org/doc/draft-rustignoli-panrg-scion-components/.

Discussion of this document takes place on the Path Aware Networking RG Research Group mailing list (mailto:panrg@irtf.org), which is archived at https://www.ietf.org/mail-archive/web/panrg/.

Source for this draft and an issue tracker can be found at https://github.com/scionassociation/scion-components_I-D.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute

working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 12 January 2023.

Table of Contents

1. Introduction

While SCION was initially developed in academia, the architecture has now "slipped out of the lab" and counts its early productive

deployments (including the Swiss inter-banking network SSFN). The architecture is composed of a system of related components, some of which are essential to set up end-to-end SCION connectivity. Add-ons provide additional functionality, security, or backwards compatibility. Discussions at PANRG [PANRG-INTERIM-Min] showed the need to describe the relationships between SCION's core components. This document, therefore focuses on each component, describing its functionality, properties, dependencies and relationships to existing protocols. The goal is not to describe each component's specification, but to illustrate the engineering decisions that made SCION what it is and to provide a basis for further discussions.

Before reading this document, please refer to [I-D.dekater-scion-overview] for a generic overview of SCION and its components, the problems it solves, and existing deployments. For an in-depth description of SCION, refer to [CHUAT22].

1.1. Design Goals

SCION was created from the start with the intention to provide the following properties for inter-domain communication.

* *Availability*. SCION aims to provide highly available communication. Its focus is not only on handling failures (both on the last hop or anywhere along the path), but also on allowing communication in the presence of adversaries. Availability is fundamental as applications move to cloud data centers, and enterprises increasingly rely on the Internet for mission-critical communication. For example, as highlighted in [I-D.rtgwg-net2cloud-problem-statement], achieving reliable inter-domain Internet connectivity remains an open challenge for cloud providers.

* *Security*. SCION comes with an arsenal of mechanisms, designed by security researchers with the goal of making most network-based and routing attacks either impossible or easy to mitigate. The relevance of Internet's routing security issues is testified by the fact that these issues now have the attention of policymakers, while previously they were only well known in industry and academia. One example is the 2022 FCC inquiry on routing security [FCC2022]. SCION strongly focuses on preventing routing attacks, IP prefix hijackings, DoS, providing stronger guarantees than the existing Internet. Security is tightly related to trust. SCION therefore offers end-hosts transparency and control over forwarding paths. In addition, SCION's design starts from the assumption that any two entities on the global Internet do not mutually trust each other. SCION therefore enables trust agility, allowing its users to decide the roots of trust they wish to rely upon.

*_Scalability_. Security and high availability should not result in compromises on scalability. At the same time, a next-generation Internet architecture should not suffer from scalability issues due to network growth or forwarding table size. The S in SCION, indeed, stands for scalability. The architecture proposes a design that is scalable both in the control plane and in the data plane (making secure forwarding efficient).

Many research efforts have analysed whether such properties could be achieved by extending the existing Internet architecture. But as described in Section 2.2.1, tradeoffs between properties would be unavoidable when exclusively relying on or extending existing protocols.

The following paragraphs describe the key properties of SCION's core components. They then describe the components' mutual dependencies and their relation with existing protocols.

2. Minimal Stack - Core Components

In order to establish end-to-end connectivity, SCION relies on three main components. SCION's data plane carries out secure path-aware forwarding. Its control plane performs routing and provides a selection of path segments. The Control Plane PKI then handles cryptographic material.

The control plane is responsible for discovering and disseminating routing information. Path discovery is performed by each autonomous system (AS) thanks to an authenticated path-exploration mechanism called beaconing. SCION end hosts query their respective AS control plane and obtain authenticated and authorized network paths, in the form of path segments. End hosts select one or more of the end-to-end network paths, based on the application requirements (i.e., latency). End hosts then craft SCION packets containing the end-to-end path to the destination. The data plane is responsible for forwarding SCION packets while authenticating them at each hop.

Both the control and data plane rely on the Control-Plane PKI (CP-PKI) for authentication. SCION's authentication mechanisms aim at protecting the whole end-to-end path at each hop. SCION Autonomous Systems are organised in Isolation Domains (ISDs), that independently define their own roots of trust. ISD members share a uniform trust environment (i.e., a common jurisdiction). They can transparently define trust relationships between parts of the network by deciding whether to trust other ISDs. SCION therefore relies on a unique trust model, which differs from other PKIs. The motivation behind this design choice is clarified in Section 2.3.

All above mentioned core components are deployed in production (e.g., they are in use within the SSFN, the Swiss Finance Network). There are commercial implementations of all core components (including a high performance data-plane).

2.1. Routing - Control Plane

The SCION control plane's main purpose is to discover and disseminate routing information, in the form of path segments. Path exploration is based on path-segment construction beacons (PCBs), which are initiated by a subset of ASes and accumulate cryptographically protected path forwarding information. Each AS selects a few PCBs and makes them available to end hosts via its path service. End hosts query the control plane for path segments, and combine them into forwarding paths to transmit packets in the data plane. For an overview of the process to create and disseminate path information, refer to [I-D.dekater-scion-overview], section 1.2.2.

2.1.1. Key Properties in Relationship to Existing Protocols

On first sight, it might seem that the SCION control plane takes care of similar duties as BGP. While both focus on disseminating routing information, there are substantial differences in their mechanisms and properties offered. This section describes the core properties provided by the SCION control plane, and its relationships with existing protocols.

* *Host addressing.* SCION decouples routing from end-host addressing: inter-domain routing is based on ISD-AS tuples rather than on end-host addresses, making SCION agnostic to end-host addressing. This design decision has two outcomes: First of all, SCION can reuse existing host addressing schemes, as IPv6, IPv4, or others. Secondly, its control plane does not carry prefix information, avoiding known issues of using routing tables (i.e., scalability, the need for dedicated hardware).

* *Multipath.* SCION ASes can select PCBs according to their policies, and register the corresponding path segments, making them available to other ASes and end hosts. SCION hosts can leverage a wide range of inter-domain paths, selecting them at each hop based on application requirements or path conditions. One existing mechanism is BGP ADD-PATH [RFC7911], focusing on advertising multiple paths for the same prefix in order to provide a backup path. However, BGP multipath does not allow end hosts to select the whole end-to-end paths, therefore traffic cannot be routed based on application requirements. In addition, it faces scalability concerns typical for BGP (i.e., increased resource requirement on routers), as discussed in the above-

mentioned RFC. Similarly to BGP multipath, other approaches based on BGP either are only able to provide backup paths that can solely be activated in case of failure (i.e., "Diverse BGP Paths" [RFC6774]), or they face scalability limitations. Such concerns motivate an alternative approach, such as SCION.

* *Hop-by-hop path authorization.* SCION packets can only be forwarded along authorized path segments. This is achieved thanks to message authentication codes (MACs) within each hop field. During beaconing, each AS's control plane creates MACs, which are then verified during forwarding. This gives end hosts strong guarantees about the path where the data is routed. Other approaches, such as BGPSec ([RFC8205]), suffer from challenges with scalability, introduce circular dependencies [COOPER2013] and global kill switches [ROTHENBERGER2017]. Giving end hosts guarantees about the full inter-domain path is important in order to avoid traffic interception, and to enable geofencing (i.e., keeping data in transit within a well-defined trusted area of the global Internet).

* *Scalability.* The SCION's beaconing algorithm is around two orders of magnitude more efficient than BGP due to the following reasons: The routing process is divided in a process within each ISD (intra-ISD) and one between ISDs (inter-ISD), SCION beaconing does not need to iteratively converge, and SCION makes AS-based announcements instead of BGP's IP prefix-based announcements. Scalability of the routing process is fundamental not only in order to support network size growth, but also in order to quickly react to failures. Refer to [KRAHENBUHL2022] for an in-depth study of SCION's scalability in comparison to BGP.

* *Convergence time.* Since routing decisions are decoupled from the dissemination of path information, SCION features faster convergence times than path-vector protocols such as BGP. Path information is propagated across the network by PCBs in times that are within the same order of magnitude of network round trip time. In addition, the division of the beaconing process into intra- and inter-ISD helps in speeding up global distribution of routing information. This means that SCION has the capability to restore global reachability, even after catastrophic failures, within tens of seconds. This is in contrast to BGP, which in certain situations will never converge to a stable state, or converge only non-deterministically (see [GRIFFIN1999] and [RFC4264]). Convergence under BGP may also simply take too much time [SAHOO2009].

* *Transparency.* SCION end hosts have full visibility about the inter-domain path where their data is forwarded. This is a property that is missing in traditional IP networks, where

routing decisions are made by each hop, therefore end hosts have no visibility nor guarantees on where their traffic is going. Additionally, SCION users have visibility on the roots of trust that are used to forward traffic. SCION therefore makes it harder to redirect traffic through an adversary's vantage point. Moreover, SCION gives end users the ability to select which parts of the Internet to trust. This is particularly relevant for workloads that currently use segregated networks.

* *Fault isolation.* As the SCION routing process is hierarchically divided into intra-ISD and inter-ISD, faults have a generally limited and localized impact. Misconfigurations, such as an erroneous path policy, may suppress some paths. However, as long as an alternative path exists, communication is possible. In addition, while the control plane is responsible for creating new paths, it does not invalidate existing paths. The latter function is handled by end hosts upon detecting failures or eventually receiving a SCMP message from the data plane. This separation of control and data plane prevents the control plane from cutting off an existing communication.

* *Authenticated control messages.* BGP has no built-in security mechanisms and does not provide any tools for ASes to authenticate the information they receive through BGP update messages. This opens up a multitude of attack opportunities. SCION control-plane messages, instead, are all authenticated, avoiding pitfalls that could possibly prevent deployment, as discussed in [RFC9049]. In addition, currently the Internet Control Message Protocol (ICMP) lacks authentication support, see [RFC4443] and [RFC0791]. Unauthenticated ICMP messages can potentially be used to affect or even prevent traffic forwarding. SCION therefore provides the SCION Control Message Protocol (SCMP), which is analogous to ICMP. It provides functionality for network diagnostics, such as ping and traceroute, and error messages that signal packet processing or network layer problems. SCMP is the first control message protocol that supports the authentication of network control messages.

Additionally, the SCION control plane design takes into account some of the lessons learned discussed in [RFC9049]: It does not try to outperform end-to-end mechanisms, as path selection is performed by end hosts. SCION, therefore, can leverage existing end-to-end mechanisms to switch paths, rather than competing with them. In addition, there is no component in the architecture that needs to keep connection state, as this task is pushed to end hosts.

Overall, several of the SCION control plane properties and key mechanisms depend on the fact that SCION ASes are grouped into Isolation Domains (ISDs). For example, ISDs are fundamental to

achieve transparency, routing scalability, fault isolation, and fast propagation of routing information. The SCION control plane therefore is built around the concept of ISDs, and relies on the SCION Control-Plane PKI (see Section 2.3) for authenticating control information.

2.2. Forwarding - Data Plane

SCION is an inter-domain network architecture and as such does not interfere with intra-domain forwarding. This corresponds to the practice today where BGP is used for inter-domain routing, while ASes use an intra-domain protocol of their choice (i.e., OSPF, IS-IS, MPLS, ...). SCION therefore re-uses the intra-domain network fabric to provide connectivity among its infrastructure services, border routers, and end hosts, minimising changes to the internal infrastructure.

SCION routers are deployed at the network edge. They receive and validate SCION packets from neighbours, then they use their intra-domain forwarding information to transmit packets to the next border router or SCION end host.

SCION packets are at the network layer (layer-3), and the SCION header sits in between the transport and link layer. The header contains a variable type and length end-host address, and can therefore carry any address (IPv4, IPv6, ...). In addition, end-host addresses only need to be unique within an AS, and can be, in principle, reused. In early deployments, intra-AS SCION packets are sometimes encapsulated into an IP packet, for backwards compatibility.

2.2.1. Key Properties in Relationship to Existing Protocols

Thanks to its data plane, SCION achieves properties that are difficult to achieve when exclusively extending existing protocols.

* *Path selection.* In SCION, end hosts select inter-domain network paths, rather than routers. The end hosts are empowered to make end-to-end path choices based on application requirements. This means that routers do not carry the burden of making enhanced routing or forwarding decisions.

* *Scalability.* SCION routers can efficiently forward packets without the need to look up forwarding tables or keeping per-connection state. Routers only need to verify MACs in hop fields. This operation is based on modern block ciphers such as AES, can be computed faster than performing a memory lookup and is widely supported in modern CPUs. Routers, therefore, do not require expensive and energy-intensive dedicated hardware, and can be deployed on off-the-shelf hardware. Lack of forwarding tables

also implies that the growing size of forwarding tables is of no concern to SCION. Additionally, routers that keep state of network information can suffer from denial-of-service (DoS) attacks exhausting the router's state [SCHUCHARD2011], which is less of a problem to SCION.

* *Recovery from failures.* SCION hosts usually receive more than one path to a given destination. Each host can select (potentially disjoint) backup paths that are available in case of failure. In contrast to the IP-based Internet, SCION packets are not dynamically rerouted by the network in case of failures. Routers use BFD [RFC5880] to detect link failures, and in case they cannot forward a packet, they send an authenticated SCMP message triggering path revocation. End hosts can use this information, or alternatively perform active monitoring, to quickly reroute traffic in case of failures. There is therefore no need to wait for inter-domain routing protocol convergence.

* *Extensibility.* SCION, similarly to IPv6, supports extensions in its header. Such extensions can be hop-by-hop (and are processed at each hop), or end-to-end.

* *Path validation.* SCION routers validate network paths in packets at each hop, so that they are only forwarded along paths that were authorized by all on-path ASes in the control plane. Thanks to a system of nested message authentication codes, traffic hijackings attacks are avoided.

In conclusion, in comparison to today's Internet, the SCION's data plane pushes some of the responsibilities away from routers onto end hosts (such as selecting paths or reacting to failures). This contributes to creating a data plane that is more efficient and scalable, and that does not require routers with specialised routing table lookup hardware. Routers validate network paths so that packets are only forwarded on previously authorized packets.

2.3. Authentication - SCION PKI

SCION's control plane messages are all authenticated. The verification of those messages relies on a public-key infrastructure (PKI) called the Control-Plane PKI or CP-PKI. It consists of a set of mechanisms, roles, and policies related to the management and usage of certificates, which enables the verification of signatures of, e.g., path-segment construction beacons (PCBs).

2.3.1. Key Properties

One might ask why SCION requires its own PKI, rather than reusing some of the existing PKI architectures. There are several properties

that distinguish the CP-PKI from others, and motivate SCION's distinct approach.

- *Unique decentralised trust model.* SCION is designed to enable global secure connectivity, where ASes do not necessarily share mutual trust. This requires a trust model that is different from existing ones that are behind commonly deployed PKIs in today's Internet. In a monopolistic model, all entities trust one or a small number of roots of trust. In an oligopolistic model, there are multiple equally trusted roots (e.g., in the Web PKI). In both models, some or all certification authorities are omnipotent. If their key is compromised, then the security of the entire system collapses. Both models do not scale well to a global environment, because mutually distrustful entities cannot agree on a single root of trust (monopoly) and because in the oligopoly model, the security is as strong as its weakest root. The SCION trust model differs from classic PKIs in two ways. First, no entity is omnipotent, as Isolation Domains elect their own root of trust, and the capabilities of each ISD (authentication-wise) are limited to communication channels in which they are involved. Second, the trust roots of each ISD are located in a single file, the TRC, which is co-signed by multiple entities in a process called voting.

- *Resilience to compromised entities and keys.* Compromised or malicious trust roots outside an ISD cannot affect operations that stay within that ISD. Moreover, each ISD can be configured to withstand the compromise of any single voting key.

- *Trust flexibility.* Each ISD can define its own trust policy. ASes must accept the trust policy of the ISD(s) in which they participate, but they can decide which ISDs they want to join, and they can participate in multiple ISDs.

- *Scalability.* The authentication infrastructure scales to the size of the Internet and is adapted to the heterogeneity of today's Internet constituents.

- *A basis for authentication of data-plane messages.* Authentication based on digital signatures works well for the relatively low message rates in the control plane, but it does not meet the performance requirements for the high message rate of the data plane. The authentication of data-plane traffic and control messages requires a highly efficient and ideally stateless system to achieve high bandwidths on commodity hardware, and to avoid creating opportunities for DoS attacks. SCION comprises a component called DRKey, which enables high-speed data-plane elements, like border routers, to derive symmetric cryptographic keys from local secrets only. This DRKey component is used to

authenticate SCMP messages. Today's Internet also lacks a fundamental mechanism to share a secret key between two end hosts for secure end-to-end communication. Existing approaches (i.e., SSH) resort to trust-on-first-use (TOFU), where a host's initial public key is accepted without verification. DRKey addresses this issue as well. For more information, refer to the draft [I-D.garciapardo-drkey].

The CP-PKI is based on certificates that follow the X.509v3 standard [RFC5280]. There are already several professional industry-grade implementations. Trust within an ISD is normally bootstrapped with an initial ceremony. Subsequent updates to the root of trust are handled automatically.

SCION is built around a unique trust model, allowing mutually distrustful entities to communicate. This justifies the existence of the CP-PKI, which differs from existing PKI architectures. Thanks to the CP-PKI, control and data plane packets are authenticated. This helps avoiding some of the obstacles to deployment mentioned in [RFC9049], where several path-aware methods failed to achieve deployment because of lack of authentication or lack of mutual trust between end hosts and the intermediate network.

3. Additional Components

This document mainly focuses on describing the fundamental components needed to run a minimal SCION network. Beyond that, SCION comprises a number of extensions and transition mechanisms that provide additional properties, as improved incremental deployability, security, additional features. For the sake of completeness, this paragraph briefly mentions some of these transition mechanisms and extensions.

3.1. Transition Mechanisms

As presented in [I-D.dekater-scion-overview], incremental deployability is a focus area of SCION's design. It comprises transition mechanisms that allow partial deployment and coexistence with existing protocols. These mechanisms require different levels of changes in existing systems, and have different maturity levels (from research to production). Rather than describing how each mechanism works, this document provides a short summary of each approach, focusing on its functions and properties, as well as on how it reuses, extends or interacts with existing protocols.

* *SCION-IP-Gateway (SIG).* A SCION-IP-Gateway (SIG) encapsulates regular IP packets into SCION packets with a corresponding SIG at the destination that performs the decapsulation. This mechanism enables IP end hosts to benefit from a SCION deployment by

transparently obtaining improved security and availability properties. SCION routing policies can be configured on SIGs, in order to select appropriate SCION paths based on application requirements. SIGs have the ability to dynamically exchange prefix information, currently using their own encapsulation and prefix exchange protocol. This does not exclude reusing existing protocols in the future. SIGs are deployed in production SCION networks, and there are commercial implementations.

* *SIAM.* To make SIGs a viable transition mechanism in an Internet-scale network with tens of thousands of ASes, an automatic configuration system is required. SIAM creates mappings between IP prefixes and SCION addresses, relying on the authorisations in the Resource Public Key Infrastructure (RPKI). SIAM is currently a research prototype, further described in [SUPRAJA2021].

* *SBAS* is an experimental architecture aiming at extending the benefits of SCION (in terms of performance and routing security) to potentially any IP host on the Internet. SBAS consists of a federated backbone of entities. SBAS appears on the outside Internet as a regular BGP-speaking AS. Customers of SBAS can leverage the system to route traffic across the SCION network according to their requirements (i.e., latency, geography, ...). SBAS contains globally distributed PoPs that advertise its customer's announcements. SBAS relies on RPKI to validate IP prefix authorization. Traffic is therefore routed as close as possible to the source onto the SCION network. The system is further described in chapter 13 of [CHUAT22].

3.2. Extensions and Other Components

In addition to the components mentioned above, there are others that aim at facilitating deployment or at better integrating SCION with existing networks. As an example, PANAPI (Path-Aware Networking API) [slides-113-taps-panapi] aims at making path-awareness and multipath to the transport layer at end hosts. DRKey [I-D.garciapardo-drkey] is a SCION extension that provides an Internet-wide key-establishment system allowing any two hosts to efficiently derive a symmetric key. This extension can be leveraged by other components to provide additional security properties. For example, LightningFilter [slides-111-panrg-lightning-filter] leverages DRKey to provide high-speed packet filtering between trusted SCION ASes. The SCION Control Message Protocol (SCMP) provides authenticated error messages and network diagnostics. COLIBRI [GIULIARI2021] is SCION's inter-domain bandwidth reservation system. RHINE (Robust and High-performance Internet Naming for End-to-end security, formerly RAINS) is a secure-by-design naming system that provides a set of desired security, reliability, and performance properties beyond

what the DNS security infrastructure offers today. It is further described in chapter 19 of [CHUAT22].

These additional components are briefly mentioned here in order to provide additional context. As extensions, they build upon the three SCION core components described earlier in this document. They are therefore unlikely to be the first components being standardised.

4. Related Work

A question that is often asked is whether SCION could simply reuse or extend existing protocols. This section tries to clarify this question, giving an overview of the relationships between SCION and other approaches. This section discusses what properties can be achieved by extending existing protocols, already deployed in the wild, and what properties can only be achieved with an approach like SCION.

4.1. SCION and RPKI

One might ask why SCION could not just rely on RPKI. Summarising the points discussed in this document, the CP-PKI distinguishes itself because of its trust model, which comprises independent trust roots that are a fundamental building block for SCION's Isolation Domains. RPKI's trust model follows the same structure as the IP allocation hierarchy, where the five RIRs run a CA. This clashes with the trust model required for SCION's Isolation Domains, therefore the SCION control plane would not be able to leverage RPKI instead of the CP-PKI. In addition, RPKI is only meant to provide authorisation, but not authentication. SCION indeed does not provide, by design, IP authorisation. Rather, one of IP-to-SCION's coexistence mechanisms mentioned earlier (SIAM) relies on RPKI for IP origin attestation.

4.2. SCION and Segment Routing

Given its path-aware properties, some of SCION's characteristics might seem similar to the ones provided by Segment Routing (SR) [RFC8402]. There are, however, fundamental differences that distinguish and motivate SCION. The most salient one is that Segment Routing is designed to be deployed across a single trusted domain. SR therefore does not focus on security, which remains an open question, as outlined in [I-D.spring-srv6-security-consideration]. SCION, instead, is designed from the start to facilitate inter-domain communication between (potentially mutually distrustful) entities. It comes, therefore, with built-in security measures to prevent attacks (i.e., authenticating all control-plane messages and all critical fields in the data-plane header). Rather than competing, SCION and SR complement each other. SCION relies on existing intra-domain routing protocols, therefore SR can be one of

the possible intra-domain forwarding mechanisms. A possible integration of its path-aware properties remains for now an open question.

4.3. SCION and Other Routing Approaches

There is an increasing motivation to extend inter-domain routing beyond mere reachability. See for example [I-D.trossen-routing-beyond-reachability], which provides a summary of some of the existing methods, and states that wider architectural approaches are needed. One proposed approach is semantic routing [I-D.irtf-introduction-to-semantic-routing], which adds support for advanced routing and forwarding into packets and into the data-plane. SCION takes a different approach: Path selection is carried out by end hosts, which have the ability to select network paths based on application requirements. This means that there is no need to include semantics in packets. This comes with the benefit that the SCION data plane can provide advanced routing without increased complexity or strain on routers. BGP ADD-PATH [RFC7911] extends BGP to allow additional path announcements for a certain prefix, without implicitly revoking the existing path. However, when additional paths are advertised for a large number of prefixes, router memory consumption is significantly increased. Furthermore, the additional path diversity is not exposed to end-hosts, therefore the additional paths can only be used for redundancy.

5. Dependency Analysis

This section briefly discusses dependencies between SCION's core components, with the goal of facilitating a discussion on whether it is possible to implement each of SCION's core components on its own, independently from other core components.

* *Control-plane PKI.* The CP PKI enables the verification of signatures, e.g., on path-segment construction beacons (PCBs). As discussed in Section 2.3, it is built on top of a peculiar trust model, where entities are able to select their roots of trust. Overall, it constitutes the most independent and self-contained building block, as it could potentially be leveraged by SCION or other protocols. The PKI itself does not have significant dependencies on other SCION components, therefore it could represent a good starting point for standardisation. Its unique trust properties, interfaces, and processes (as voting), could be a good candidate for a first draft.

* *Control plane.* The SCION control plane is built around the concept of Isolation Domains, being the routing process divided into an intra- and inter-ISD one. It heavily relies on the CP-PKI for beaconing (i.e., for authenticating routing information).

Each Isolation Domain requires its own root of trust in order to carry out path exploration and dissemination. Decoupling the control plane from the CP-PKI would severely affect the properties and guarantees that can be provided by the control plane. The control plane could, therefore, be specified in parallel with the CP-PKI. The control plane is internally formed by multiple sub-components (as the beacon service, responsible for path discovery, and the path service, responsible for path dissemination). Processes and interfaces between these sub-components could be topic for one or multiple drafts.

* *Data plane.* In order to be able to transmit data, end hosts need to fetch path information from their AS control plane, as discussed in Section 2.2. In addition, the SCION data plane requires that hosts validate paths, and that routers authenticate path information at each hop. This authentication mechanism relies on the control-plane PKI. It is what allows SCION to distinguish itself from other proposals, gaining many of the security and availability proprieties discussed earlier. The data plane, therefore, relies on both the control plane and the control-plane PKI in order to function. Should the data plane be used independently, without end-to-end path validation, SCION would loose many of its security properties, which are fundamental in an inter-domain scenario where entities are mutually distrustful. As discussed in [RFC9049], lack of authentication has often been the cause for path-aware protocols never being adopted because of security concerns. SCION should avoid such pitfalls and therefore its data plane should rely on the corresponding control plane and control-plane PKI.

6. Conclusions

This document describes the three fundamental SCION core components, together with their properties and dependencies. It highlights how such components allow SCION to provide unique properties. It then discusses how the main components are interlinked, with the goal of fostering a discussion on the standardisation of key components. As this document is an early draft, the authors welcome feedback from the IETF community for future iterations.

7. Informative References

[CHUAT22] 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>.

[COOPER2013] Cooper, D., Heilman, E., Brogle, K., Reyzin, L., and S. Goldberg, "On the risk of misbehaving RPKI authorities",

Proceedings of the Twelfth ACM Workshop on Hot Topics in Networks, DOI 10.1145/2535771.2535787, November 2013, <https://doi.org/10.1145/2535771.2535787>.

[FCC2022] Federal Communications Commission, "Notice of Inquiry on Secure Internet Routing", 2022, <https://www.fcc.gov/document/fcc-launches-inquiry-internet-routing-vulnerabilities>.

[GIULIARI2021] Giuliari, G., Roos, D., Wyss, M., García-Pardo, J., Legner, M., and A. Perrig, "Colibri: A Cooperative Lightweight Inter-domain Bandwidth-Reservation Infrastructure", 2022, <https://netsec.ethz.ch/publications/papers/2021_conext_colibri.pdf>.

[GRIFFIN1999] Griffin, T. and G. Wilfong, "An analysis of BGP convergence properties", ACM SIGCOMM Computer Communication Review Vol. 29, pp. 277-288, DOI 10.1145/316194.316231, October 1999, <https://doi.org/10.1145/316194.316231>.

[I-D.dekater-scion-overview] de Kater, C., Rustignoli, N., and A. Perrig, "SCION Overview", 2022, <https://datatracker.ietf.org/doc/draft-dekater-panrg-scion-overview/>.

[I-D.garciapardo-drkey] Pardo, J., Krähenbühl, C., Rothenberger, B., and A. Perrig, "Dynamically Recreatable Keys", 2022, <https://datatracker.ietf.org/doc/draft-garciapardo-panrg-drkey/>.

[I-D.irtf-introduction-to-semantic-routing] Farrel, A. and D. King, "An Introduction to Semantic Routing", 2022, <https://datatracker.ietf.org/doc/draft-li-spring-srv6-security-consideration/>.

[I-D.rtgwg-net2cloud-problem-statement] Dunbar, L., Malis, A., Jacquenet, C., Toy, M., and K. Majumdar, "SCION Overview", 2022, <https://datatracker.ietf.org/doc/draft-ietf-rtgwg-net2cloud-problem-statement/>.

[I-D.spring-srv6-security-consideration] Li, C., Li, Z., Xie, C., Tian, H., and J. Mao, "Security Considerations for SRv6 Networks", 2022, <https://datatracker.ietf.org/doc/draft-li-spring-srv6-security-consideration/>.

[I-D.trossen-routing-beyond-reachability] Trossen, D., Lou, D., and S. Jiang, "Continuing to Evolve Internet Routing Beyond 'Mere' Reachability", 2022, <https://

datatracker.ietf.org/doc/draft-trossen-rtgwg-routing-beyond-reachability/>.

[KRAHENBUHL2022] Kraћenbuћl, C., Tabaeiaghdaei, S., Gloor, C., Kwon, J., Perrig, A., Hausheer, D., and D. Roos, "Deployment and Scalability of an Inter-Domain Multi-Path Routing Infrastructure", 2022, <https://netsec.ethz.ch/publications/papers/2021_conext_deployment.pdf>.

[PANRG-INTERIM-Min] "Path Aware Networking Research Group - Interim 106 Minutes", June 2022, <https://datatracker.ietf.org/meeting/interim-2022-panrg-01/materials/minutes-interim-2022-panrg-01-202206011700-00>.

[RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <https://www.rfc-editor.org/rfc/rfc791>.

[RFC4264] Griffin, T. and G. Huston, "BGP Wedgies", RFC 4264, DOI 10.17487/RFC4264, November 2005, <https://www.rfc-editor.org/rfc/rfc4264>.

[RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", STD 89, RFC 4443, DOI 10.17487/RFC4443, March 2006, <https://www.rfc-editor.org/rfc/rfc4443>.

[RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, DOI 10.17487/RFC5280, May 2008, <https://www.rfc-editor.org/rfc/rfc5280>.

[RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, DOI 10.17487/RFC5880, June 2010, <https://www.rfc-editor.org/rfc/rfc5880>.

[RFC6774] Raszuk, R., Ed., Fernando, R., Patel, K., McPherson, D., and K. Kumaki, "Distribution of Diverse BGP Paths", RFC 6774, DOI 10.17487/RFC6774, November 2012, <https://www.rfc-editor.org/rfc/rfc6774>.

[RFC7911] Walton, D., Retana, A., Chen, E., and J. Scudder, "Advertisement of Multiple Paths in BGP", RFC 7911, DOI

10.17487/RFC7911, July 2016, <https://www.rfc-editor.org/rfc/rfc7911>.

[RFC8205] Lepinski, M., Ed. and K. Sriram, Ed., "BGPsec Protocol Specification", RFC 8205, DOI 10.17487/RFC8205, September 2017, <https://www.rfc-editor.org/rfc/rfc8205>.

[RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <https://www.rfc-editor.org/rfc/rfc8402>.

[RFC9049] Dawkins, S., Ed., "Path Aware Networking: Obstacles to Deployment (A Bestiary of Roads Not Taken)", RFC 9049, DOI 10.17487/RFC9049, June 2021, <https://www.rfc-editor.org/rfc/rfc9049>.

[ROTHENBERGER2017] Rothenberger, B., Asoni, D., Barrera, D., and A. Perrig, "Internet Kill Switches Demystified", Proceedings of the 10th European Workshop on Systems Security, DOI 10.1145/3065913.3065922, April 2017, <https://doi.org/10.1145/3065913.3065922>.

[SAHOO2009] Sahoo, A., Kant, K., and P. Mohapatra, "BGP convergence delay after multiple simultaneous router failures: Characterization and solutions", Computer Communications Vol. 32, pp. 1207-1218, DOI 10.1016/j.comcom.2009.03.009, May 2009, <https://doi.org/10.1016/j.comcom.2009.03.009>.

[SCHUCHARD2011] Schuchard, M., Mohaisen, A., Foo Kune, D., Hopper, N., Kim, Y., and E. Vasserman, "Losing control of the internet: using the data plane to attack the control plane", Proceedings of the 17th ACM conference on Computer and communications security - CCS '10, DOI 10.1145/1866307.1866411, 2010, <https://doi.org/10.1145/1866307.1866411>.

[slides-111-panrg-lightning-filter]
Garcia Pardo, J. A., "Lightning Filter: High-Speed Traffic Filtering based on DRKey", 2021, <https://datatracker.ietf.org/meeting/111/materials/slides-111-panrg-lightning-filter-high-speed-traffic-filtering-based-on-drkey-00.pdf>.

[slides-113-taps-panapi] Kruğer, T., "PANAPI, a Path-Aware Networking API", 2022, <https://datatracker.ietf.org/

meeting/113/materials/slides-113-taps-panapi-implementation-00.pdf>.

[SUPRAJA2021] Supraja, S., Wirz, F., de Ruiter, J., Schutijser, C., Legner, M., and A. Perrig, "Global Distributed Secure Mapping of Network Addresses", 2021, <https://netsec.ethz.ch/publications/papers/sridhara_taurin2021_siam.pdf>.

Acknowledgments

The authors are indebted to Adrian Perrig, Laurent Chuat, Markus Legner, David Basin, David Hausheer, Samuel Hitz, and Peter Mueller, for writing the book "The Complete Guide to SCION" [CHUAT22], which provides the background information needed to write this document.

Authors' Addresses

Nicola Rustignoli
ETH Zuerich

Email: nicola.rustignoli@inf.ethz.ch

Corine de Kater
ETH Zuerich

Email: corine.dekatermuehlhaeuser@inf.ethz.ch