Workgroup: PANRG
Internet-Draft:
draft-dekater-panrg-scion-overview-03
Published: 7 March 2023
Intended Status: Informational
Expires: 8 September 2023
Authors: C. de Kater
SCION Association
N. Rustignoli
SCION Association
A. Perrig
ETH Zuerich

SCION Overview

Abstract

The Internet has been successful beyond even the most optimistic expectations and is intertwined with many aspects of our society. But although the world-wide communication system guarantees global reachability, the Internet has not primarily been built with security and high availability in mind. The next-generation inter-network architecture SCION (Scalability, Control, and Isolation On Next-generation networks) aims to address these issues. SCION was explicitly designed from the outset to offer security and availability by default. The architecture provides route control, failure isolation, and trust information for end-to-end communication. It also enables multi-path routing between hosts.

This document discusses the motivations behind the SCION architecture and gives a high-level overview of its fundamental components, including its authentication model and the setup of the control- and data plane. A more detailed analysis of relationships and dependencies between components is available in [I-D.rustignoli-scion-components]. As SCION is already in production use today, the document concludes with an overview of SCION deployments.

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-overview_I-D/draft-dekater-panrg-scion-overview.html. Status information for this document may be found at https://datatracker.ietf.org/doc/draft-dekater-panrg-scion-overview/.

Discussion of this document takes place on the WG Working Group mailing list (mailto:panrg@irtf.org), which is archived at https://datatracker.ietf.org/rg/panrg. Subscribe at https://www.ietf.org/mailman/listinfo/panrg/.

Source for this draft and an issue tracker can be found at https://github.com/scionassociation/scion-overview_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 8 September 2023.

Copyright Notice

Table of Contents

1. Introduction
  1.1. Why SCION - Motivation
    1.1.1. Scope of SCION
    1.1.2. Practical Considerations Based on Related RFCs
    1.1.3. Why Now?
  1.2. SCION Overview
    1.2.1. Network Architecture and Naming
    1.2.2. Routing
    1.2.3. Infrastructure Components
    1.2.4. Formal Verification
  1.3. Conventions and Definitions
2. Key Concepts
  2.1. Authentication
    2.1.1. The Control-Plane PKI (CP-PKI)
    2.1.2. TRC Update and Verification
    2.1.3. Dissemination of TRC Updates

2.1.4. Grace Period
2.1.5. Revocation and Recovery from a Catastrophic Event
2.2. SCION Control Plane
2.2.1. Path Exploration
2.2.2. Path Registration
2.2.3. Path Lookup
2.2.4. Link Failures
2.3. SCION Data Plane
2.3.1. Path Construction via Segment Combination
2.3.2. Path Authorization
2.3.3. Forwarding
2.3.4. Intra-AS Communication
3. Deployment
3.1. Autonomous System Deployment
3.2. Internet Exchange Points
3.3. Endpoints and Incremental Deployability
3.3.1. Native Endpoints
3.3.2. SCION to IP Gateway (SIG)
3.4. Deployment Experiences
4. IANA Considerations
5. Security Considerations
6. References
6.1. Normative References
6.2. Informative References
Acknowledgments
Authors' Addresses

1. Introduction

The Introduction section explores the motivation to develop SCION, followed by a short description of SCION's main elements. The sections after the Introduction provide further insight into SCION's key concepts and deployment scenarios. The document concludes with some concrete case studies where SCION has been successfully deployed in production.

1.1. Why SCION - Motivation

Since its inception, the Internet has continued to expand, encompassing new uses over time. The continuous expansion has brought many issues to light, including a lack of control, limitations in scalability, performance and security, occurrences of severe outages, weak fault isolation, and energy consumption. With the core focus on functionality and operation, the current Internet offers little protection against attacks such as spoofing, IP-address hijacking, denial-of-service, and combinations of these. For more background information, see [SCHUCHARD2011], [LABOVITZ2000], [GRIFFIN1999], [SAHOO2009], and [RFC4264].

There have been numerous initiatives to address the above issues. Although these initiatives have brought many improvements, concerns regarding security and scalability still remain. For more details, see, e.g., [RFC4033], [RFC6480], [RFC8205], and [RFC8446], as well as [LYCHEV2013], [LI2014], [COOPER2013], [ROTHENBERGER2017], [MORILLO2021], and [KING2022].

As a consequence, today's Internet fails to fulfil many users' requirements. This especially pertains to the demands of enterprises globally exchanging sensitive information, such as financial institutions, healthcare providers, universities, multinationals, governments, critical and transportation infrastructure operators. These users require the Internet to be highly available at all times. They expect reliable operation of the global network also in case of failures. They need availability guarantees across multiple routing domains, even in the presence of attacks. They further want to rely on an Internet that can be multilaterally governed and is free from global kill-switches.

SCION has been developed in order to meet the above-mentioned requirements. SCION aims to reach the following goals:

* Provide high-availability architecture (also in the presence of adversaries)

* Provide fast failover in the case of inter-domain link or router failures

* Prevent from IP-address hijacking, DoS, and other attacks

* Enable path transparency as well as application-specific path optimizations

* Improve the inter-domain control plane's scalability

* Prepare the Internet for tomorrow's applications, such as virtual reality, Internet of Things (IoT), and all other applications requiring high-performance connectivity.

1.1.1. Scope of SCION

The above section describes SCION's aspiration to improve *inter*-AS routing and to focus on providing end-to-end connectivity. However, SCION does not solve *intra*-AS routing issues, nor does it provide end-to-end payload encryption, and identity authentication. These topics, which are equally important for the Internet to perform well, lie outside the scope of SCION.

1.1.2. Practical Considerations Based on Related RFCs

The SCION inter-domain routing concept has initially been developed by researchers of the ETH Zuerich [PERRIG2017], and could in the meantime also gain attention and recognition in the international academic world. However, for an IT architecture to be successful, it must work well in practice, too. This section pays attention to the implementation considerations of a conceptual framework such as SCION in real life, on the basis of some RFCs exploring this topic. It also verifies in how far SCION meets the requirements mentioned and questions raised in these RFCs.

1.1.2.1. Avoiding Pitfalls

[RFC9049] describes why previous proposals to tackle some of the Internet's fundamental issues did not manage to succeed. SCION seems to avoid the pitfalls mentioned in that document. For example, SCION does not have to be adopted by the entire Internet to be effective: The routing architecture provides benefits already to early adapters. Even if only a small part of the global network works with SCION, adapters will still take advantage of using the SCION routing technology. Moreover, not only users of SCION benefit from it, also ISPs and operators benefit from deploying SCION: early deployments showed that providers can charge the use of SCION as premium connectivity, with users who are willing to pay for it. Furthermore, SCION can be installed on top of and function alongside the existing routing infrastructure and protocols--there is no need for high-impact changes in an operational network.

Another RFC that must be mentioned in this context is [RFC5218], "What Makes for a Successful Protocol?". SCION seems to fulfil most factors that contribute to the success of a protocol, as described in section 2.1 of the RFC. This includes such factors as offering a positive net value (i.e., the benefits of deploying SCION outweigh the costs), incremental deployability, and open source code availability. More importantly, SCION averts the failure criteria mentioned in section 1.4 of the RFC: SCION is already deployed and in use by many actors of the Swiss financial and academic ecosystems, and allows for multiple implementations, both open and closed source. As existing SCION implementations are easily portable, adoption in mainstream platforms is also possible.

1.1.2.2. Answering Questions

SCION can be considered a *path-aware internetworking* architecture, as described in [RFC9217]. This RFC poses (open) questions that must be answered in order to realize such a path-aware networking system. It was originally written to frame discussions in the Path Aware

Networking Research Group (PANRG), and has been published to snapshot current thinking in this space.

SCION intends to answer the questions raised in this RFC. This especially pertains to the second, third, seventh, and eighth question:

- How do endpoints and applications get access to accurate, useful, and trustworthy path properties?
- How can endpoints select paths to use for traffic in a way that can be trusted by the network, the endpoints, and the applications using them?
- How can a path-aware network in a path-aware internetwork be effectively operated, given control inputs from network administrators, application designers, and end users?
- How can the incentives of network operators and end users be aligned to realize the vision of path-aware networking, and how can the transition from current ("path-oblivious") to path-aware networking be managed?

SCION's answers to these questions can be found in Key Concepts (Section 2) and Deployments (Section 3.4), respectively.

1.1.3. Why Now?

The emergence of multiple SCION implementations and early deployments highlights the need for standardization. The time seems therefore ripe to take SCION to the IETF, also in order to contribute to the important discussion regarding path-aware networking.

1.2. SCION Overview

SCION has been designed to address the fundamental issues of today's Internet depicted in Why SCION - Motivation (Section 1.1). The following section gives a high-level overview of SCION's main elements, providing a basic understanding of this next-generation inter-network architecture.

1.2.1. Network Architecture and Naming

SCION's main goal is to offer highly available and efficient inter-domain packet delivery—even in the presence of actively malicious entities. To achieve scalability and sovereignty, SCION organizes existing ASes into groups of independent routing planes, called **Isolation Domains (ISD)**. An AS can be a member of multiple ISDs. All ASes in an ISD agree on a set of trust roots, called the **Trust Root Configuration (TRC)**. The ISD is governed by a set of **core ASes**, which

provide connectivity to other ISDs and manage the trust roots. Typically, a few distinguished ASes within an ISD form the ISD's core.

Isolation domains serve the following purposes:

- They allow SCION to support trust heterogeneity, as each ISD can independently define its roots of trust;
- They provide transparency for trust relationships;
- They isolate the routing process within an ISD from external influences such as attacks and misconfigurations; and
- They improve the scalability of the routing protocol by separating it into a process within and one between ISDs.

ISDs provide natural isolation of routing failures and misconfigurations, provide meaningful and enforceable trust, enable endpoints to optionally restrict traffic forwarding to trusted parts of the Internet infrastructure only, and enable scalable routing updates with high path-freshness.

1.2.1.1. Links

There are three types of links in SCION: core links, parent-child links, and peering links.

- A **core link** can only exist between two core ASes.
- A **parent-child link** requires that at least one of the two connected ASes is a non-core AS. ASes with a parent-child link usually belong to the same entity or have a provider-customer relationship.
- A **peering link** also includes at least one non-core AS.

See Figure 1 for a high-level overview of the SCION network structure.

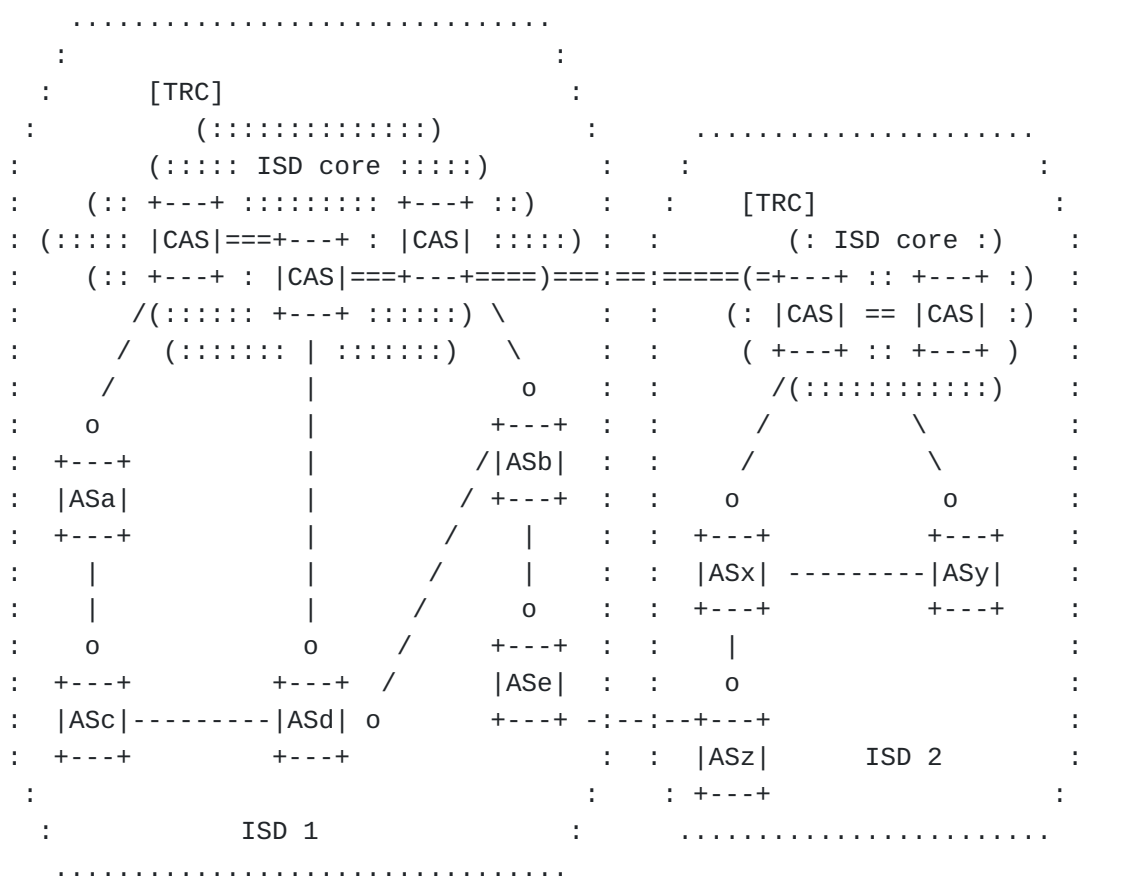

Figure 1: SCION network structure

1.2.2. Routing

SCION operates on two routing levels: intra-ISD and inter-ISD. Both levels use **path-segment construction beacons (PCBs)** to explore network paths. A PCB is initiated by a core AS and then disseminated either within an ISD (to explore intra-ISD paths) or among core ASes (to explore core paths across different ISDs). The PCBs accumulate cryptographically protected path and forwarding information on AS-level, and store this information in the form of **hop fields (HFs)**. Endpoints use information from these hop fields to create end-to-end forwarding paths for data packets, who carry this information in their packet headers. This concept is called **packet-carried forwarding state (PCFS)**. The concept also supports multi-path communication among endpoints.

The process of creating an end-to-end forwarding path consists of the following steps:

1. First, an AS discovers paths to other ASes, during the *path exploration* (or beaconing) phase.

2. The AS then selects a few PCBs according to defined policies, transforms the selected PCBs into path segments, and registers these segments with its path infrastructure, thus making them available to other ASes. This happens during the *path registration* phase.

3. During the *path resolution* phase, the actual creation of an end-to-end forwarding path to the destination takes place. For this, an endpoint performs (a) a *path lookup* step, to obtain path segments, and (b) a *path combination* step, to combine the forwarding path from the segments.

Figure 2 below shows the SCION routing process in a nutshell.

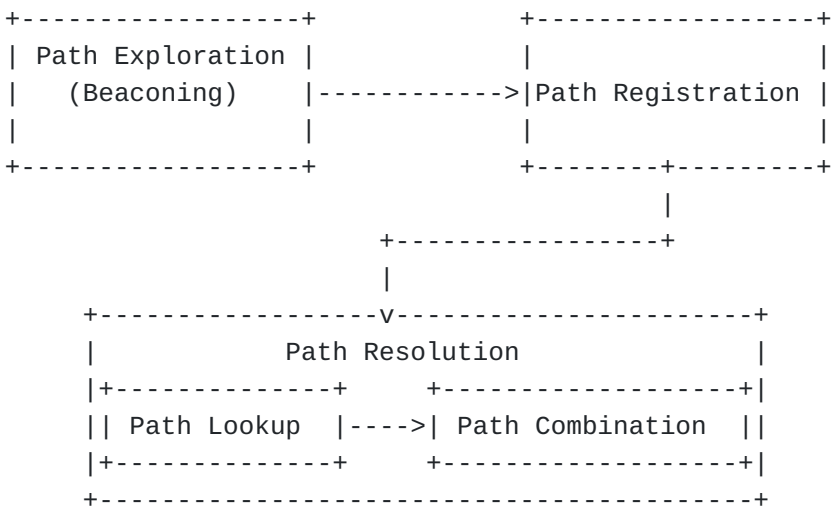

Figure 2: SCION routing in a nutshell

1.2.2.1. ISD and AS numbering

SCION decouples endpoint addressing from inter-domain routing. Routing is based on the <ISD, AS> tuple, agnostic of endpoint addressing. Existing AS numbers are inherited from the current Internet, but a 48-bit namespace allows for additional SCION AS numbers beyond the 32-bit space in use today. The endpoint local address is not used for inter-domain routing or forwarding, does not need to be globally unique, and can thus be an IPv4, IPv6, or MAC address, for example. A SCION address is therefore composed of the <ISD, AS, local address> 3-tuple.

1.2.3. Infrastructure Components

The **beacon service**, the **path service**, and the **certificate service** are the main control-plane infrastructure components within a SCION AS. Each service can be deployed redundantly, depending on the AS's size and type. Existing Internal routers are used to forward packets inside the AS, while *SCION border routers* provide interconnectivity between ASes.

- The *beacon service* discovers path information. It is responsible for generating, receiving, and propagating PCBs. Periodically, the beacon service generates a set of PCBs, which are forwarded to its child ASes or neighboring core ASes. The PCBs are flooded over policy-compliant paths to discover multiple paths between any pair of core ASes.

- The *path service* stores mappings from AS identifiers to sets of announced path segments. The path service is organized as a hierarchical caching system similar to that of DNS. Through the beacon service, ASes select the set of path segments through which they want to be reached, and they register them to the path service in the ISD core.

- The *certificate service* keeps cached copies of certificates and manages keys and certificates for securing inter-AS communication. The certificate service is queried by the beacon service when validating the authenticity of PCBs (i.e., when the beacon service lacks a certificate).

Border routers are deployed at the edge of SCION ASes. The main task of border routers is to forward packets to a neighbor border router or to the destination host within the AS. While SCION takes care of inter-domain routing, it relies on existing routing protocols (e.g., IS-IS, OSPF, SR) and communication fabric (e.g., IP, MPLS) for intra-domain forwarding. *Internal routers*, therefore, do not need to be changed to support SCION.

1.2.4. Formal Verification

An additional feature of SCION is its formal verification. The SCION network system consists of a variety of components such as routers, servers, and edge devices. Such a complex system eludes the mental capacities of human beings for considering all possible states and interactions. That is why SCION includes a formal verification framework developed by the Department of Computer Science of the ETH Zurich [KLENZE2021]. This guarantees that packet forwarding as well as SCION's authentication mechanisms and implementations are correct and consistent.

1.3. Conventions and Definitions

The key words "**MUST**", "**MUST NOT**", "**REQUIRED**", "**SHALL**", "**SHALL NOT**", "**SHOULD**", "**SHOULD NOT**", "**RECOMMENDED**", "**NOT RECOMMENDED**", "**MAY**", and "**OPTIONAL**" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Key Concepts

This section explains the SCION key concepts, including authentication, control- and data plane.

2.1. Authentication

SCION's control plane relies on a public-key infrastructure called the **control-plane PKI (CP-PKI)**, which is organized on ISD-level. Each ISD can independently build its own roots of trust, defined in a file called **trust root configuration (TRC)**.

Note: This section describes the SCION authentication concept on a very high level. A much more detailed description of SCION's authentication is available in [I-D.dekater-scion-pki].

2.1.1. The Control-Plane PKI (CP-PKI)

Trust within each isolation domain is anchored in the trust root configuration (TRC) file. Each TRC contains a collection of signed root certificates, which are used to sign CA certificates, which are in turn used to sign AS certificates. The TRC also includes ISD-policies that specify, for example, the TRC's usage, validity, and future updates. A TRC is a fundamental component of an CP-PKI.

The initial TRC in an ISD is called the **base TRC**. This base TRC constitutes the ISD's trust anchor. It is signed during a signing ceremony and then distributed throughout the ISD. All entities within the ISD obtain the initial TRC with an offline mechanism such as a USB flash drive provided by a trusted AS, like the relevant ISP, or with an online mechanism that relies on a trust-on-first-use (TOFU) approach. However, all updates to the base TRCs are performed in a straightforward process that does not require any manual or out-of-band action (such as a software update), see TRC Update and Verification (Section 2.1.2).

Figure 3 shows the TRC trust chain and associated certificates. TRC 1 is the base TRC, and TRC 2 and 3 constitute updates to this base TRC. TRC 2 must be verified using the voting certificates in TRC 1. Control-plane (CP) root certificates are used to verify other CP certificates (which are in turn used to verify path-segment construction beacons PCBs).

Each SCION AS must hold a private key (to sign PCBs) and a certificate attesting that it is the rightful owner of the corresponding public key. One of the main roles of the TRC is thus enabling the verification of **AS certificates** and PCBs.

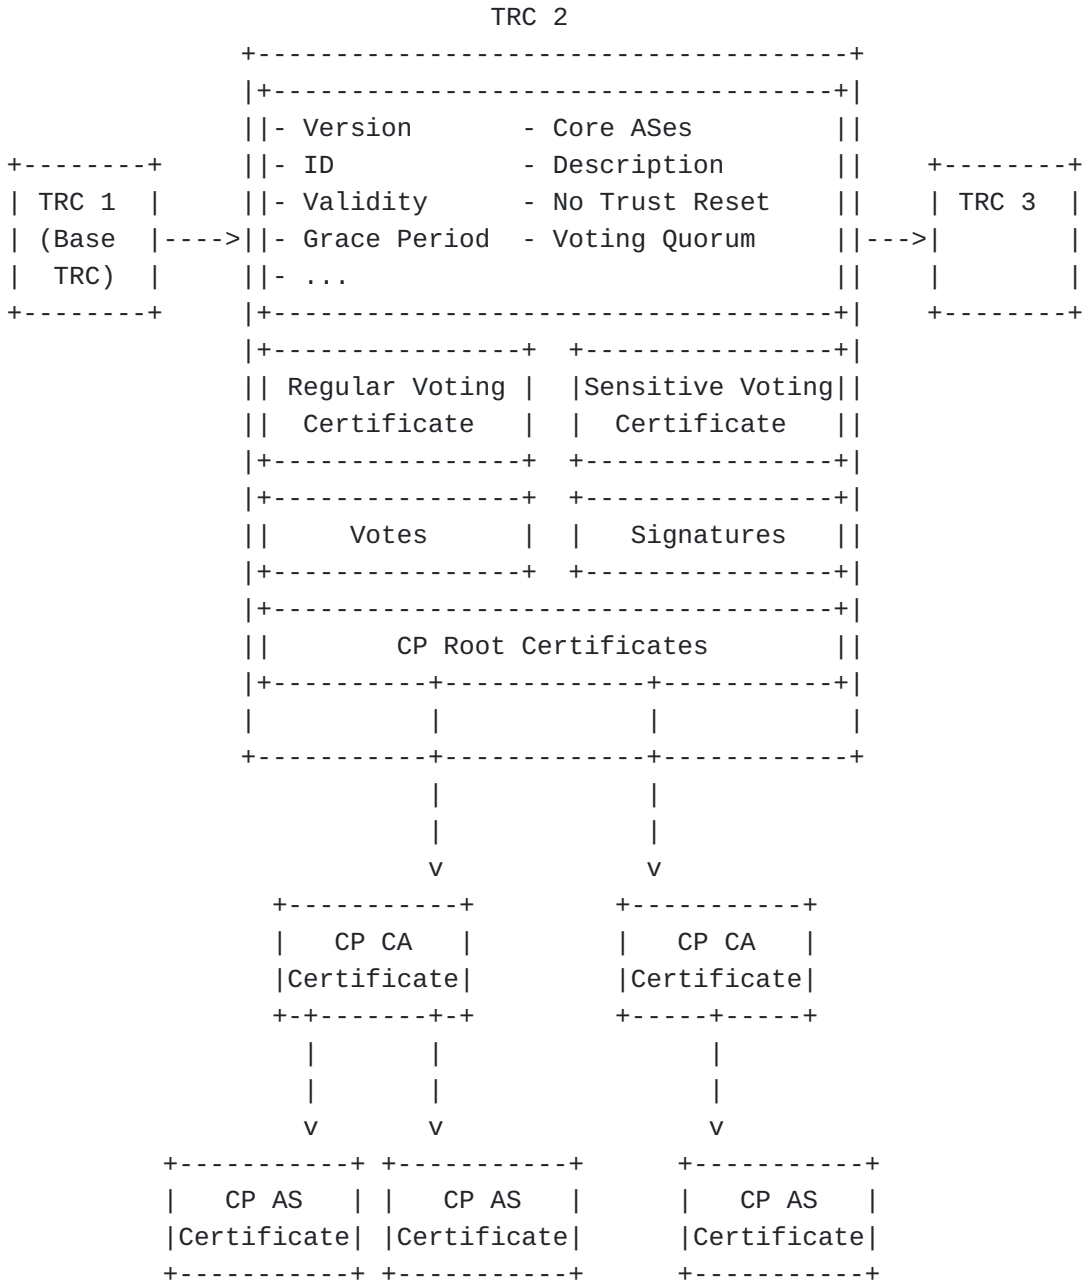

Figure 3: TRC contents and trust chain

2.1.2. TRC Update and Verification

With a base TRC as trust anchor, TRCs can be updated in a verifiable manner. There are two kinds of TRC updates: regular and sensitive updates. A *regular* TRC update happens when the TRC's validity period

expires. This period is defined by the *validity* parameter in the TRC. The default is one year. A TRC update is *sensitive* if and only if critical sections of the TRC are affected (for example, if the set of core ASes is modified). For both regular and sensitive TRC updates, a number of votes (in the form of signatures) must be cast to approve the update. This number of votes is dictated by the voting quorum and set in each TRC with the *voting quorum* parameter.

2.1.3. Dissemination of TRC Updates

Information about a TRC update is disseminated via the SCION's beaconing process, through the path-segment constructions beacons. Each PCB contains the version number of the currently active TRC. If an AS receives a PCB with a TRC version number higher than the locally stored TRC, it requests the PCB-sending AS for the new TRC. The new TRC is verified on the basis of the current one, and is accepted if it contains at least the required quorum of correct signatures by trust roots defined in the current TRC. This simple dissemination mechanism has two advantages: It is very efficient (as fresh PCBs rapidly reach all ASes), and it avoids circular dependencies with regard to the verification of PCBs and the beaconing process itself (as no server needs to be contacted over unknown paths in order to fetch the updated TRC).

2.1.4. Grace Period

At most two TRCs per ISD can be active at the same time. The TRC parameter *grace period* indicates for how long the currently active TRC can still be active after a new TRC is disseminated. This so-called **grace period** starts at the beginning of the validity period of the new TRC. An older TRC can only be active until either (1) the grace period has passed, or (2) yet a newer TRC is announced. AS certificates are validated by following the chain of trust up to an active TRC.

2.1.5. Revocation and Recovery from a Catastrophic Event

The TRC dissemination mechanism also enables rapid revocation of trust roots. When a trust root is compromised, the other trust roots can remove it from the TRC and disseminate a new TRC alongside a PCB with a new version number.

In case of a catastrophic event—such as several private root keys being disclosed due to a critical vulnerability in a cryptographic library—SCION is equipped with a recovery procedure called **trust reset**. This procedure consists of creating a new TRC with fresh, trustworthy keys (and potentially new algorithms), and redistributing this TRC out-of-band. A trust reset effectively establishes a new base TRC for the ISD. It is possible for ISDs to disable trust resets

by setting the *no trust reset* Boolean parameter to "true" in their TRC, with the effect that the entire ISD would have to be abandoned in the event of such a catastrophic compromise (this abandonment would also have to be announced out-of-band).

The partition of the SCION network into ISDs guarantees that no single entity can take down the entire network. Even if several entities formed a coalition to carry out an attack, the effects of that attack would be limited to one or a few ISDs.

2.2. SCION Control Plane

The SCION control plane is responsible for discovering path segments and making them available to endpoints. This process includes path exploration, registration, and lookup; it involves the path service, beacon service, and certificate service, both in core ASes and non-core ASes.

Note: This section describes the SCION control plane on a very high level. A much more detailed description of SCION's control plane will follow in a separate internet draft.

2.2.1. Path Exploration

In SCION, the path segment construction process is referred to as **beaconing**. The *beacon service* of each AS is responsible for the beaconing process. The beacon service generates, receives, and propagates the **path-segment construction beacons (PCBs)** on a regular basis, to iteratively construct path segments.

There are three types of path segments (note that all path segments can be used to send data traffic in both directions):

- A path segment from a non-core AS to a core AS is an *up-path segment*.
- A path segment from a core AS to a non-core AS is a *down-path segment*.
- A path segment between core ASes is a *core-path segment*.

All path segments are invertible: A core-path segment can be used bidirectional, and an up-path segment can be converted into a down-path segment, or vice versa, depending on the direction of the end-to-end path.

Path segment construction is conducted hierarchically on two levels:

- *Core beaconing* is the process of constructing path segments between core ASes. During core beaconing, the beacon service of a

core AS either initiates PCBs or propagates PCBs received from neighboring core ASes to all other neighboring core ASes. Core beaconing in SCION is similar to BGP's route-advertising process, although in SCION the process is periodic and PCBs are flooded over policy-compliant paths to discover multiple paths between any pair of core ASes.

* *Intra-ISD beaconing* creates path segments from core ASes to non-core ASes. For this, the beacon service of a core AS creates PCBs and sends them to the non-core child ASes (typically customer ASes). The beacon service of a non-core child AS receives these PCBs and forwards them to its child ASes, and so on. This procedure continues until the PCB reaches an AS without any customer (leaf AS). As a result, all ASes receive path segments to reach the core ASes of their ISD.

On its way down, a PCB accumulates cryptographically protected path- and forwarding information per traversed AS. At every AS, metadata as well as information about the AS's ingress and egress interfaces (i.e., link identifiers) is added to the PCB. The ingress and egress interface IDs identify connections to neighboring ASes. These IDs only need to be unique within each AS. Therefore, they can be chosen and encoded by each AS independently and without any need for coordination.

SCION also supports shortcuts and peering links. In a *shortcut*, a path only contains an up-path and a down-path segment, which can cross over at a non-core AS that is common to both paths. *Peering links* can be added to up- or down-path segments, resulting in an operation similar to today’s Internet.

To reduce beaconing overhead and prevent possible forwarding loops, PCBs do not traverse peering links. Instead, peering links are announced along with a regular path in a PCB. If the path segments of both ASes at the end of a peering link contain this peering link, then it is possible to use the peering link to shortcut the end-to-end path (i.e., without going through the core). SCION also supports peering links that cross ISD boundaries, according to SCION’s path transparency property.

Figure 4 shows how intra-ISD PCB propagation works, from the ISD's core AS down to child ASes. For the sake of illustration, the interfaces of each AS are numbered with integer values. In practice, each AS can choose any encoding for its interfaces; in fact, only the AS itself needs to understand its encoding. Here, AS F receives two different PCBs via two different links from core AS X. Moreover, AS F uses two different links to send two different PCBs to AS G, each containing the respective egress interfaces. AS G extends the two PCBs and forwards both of them over a single link to a child AS.

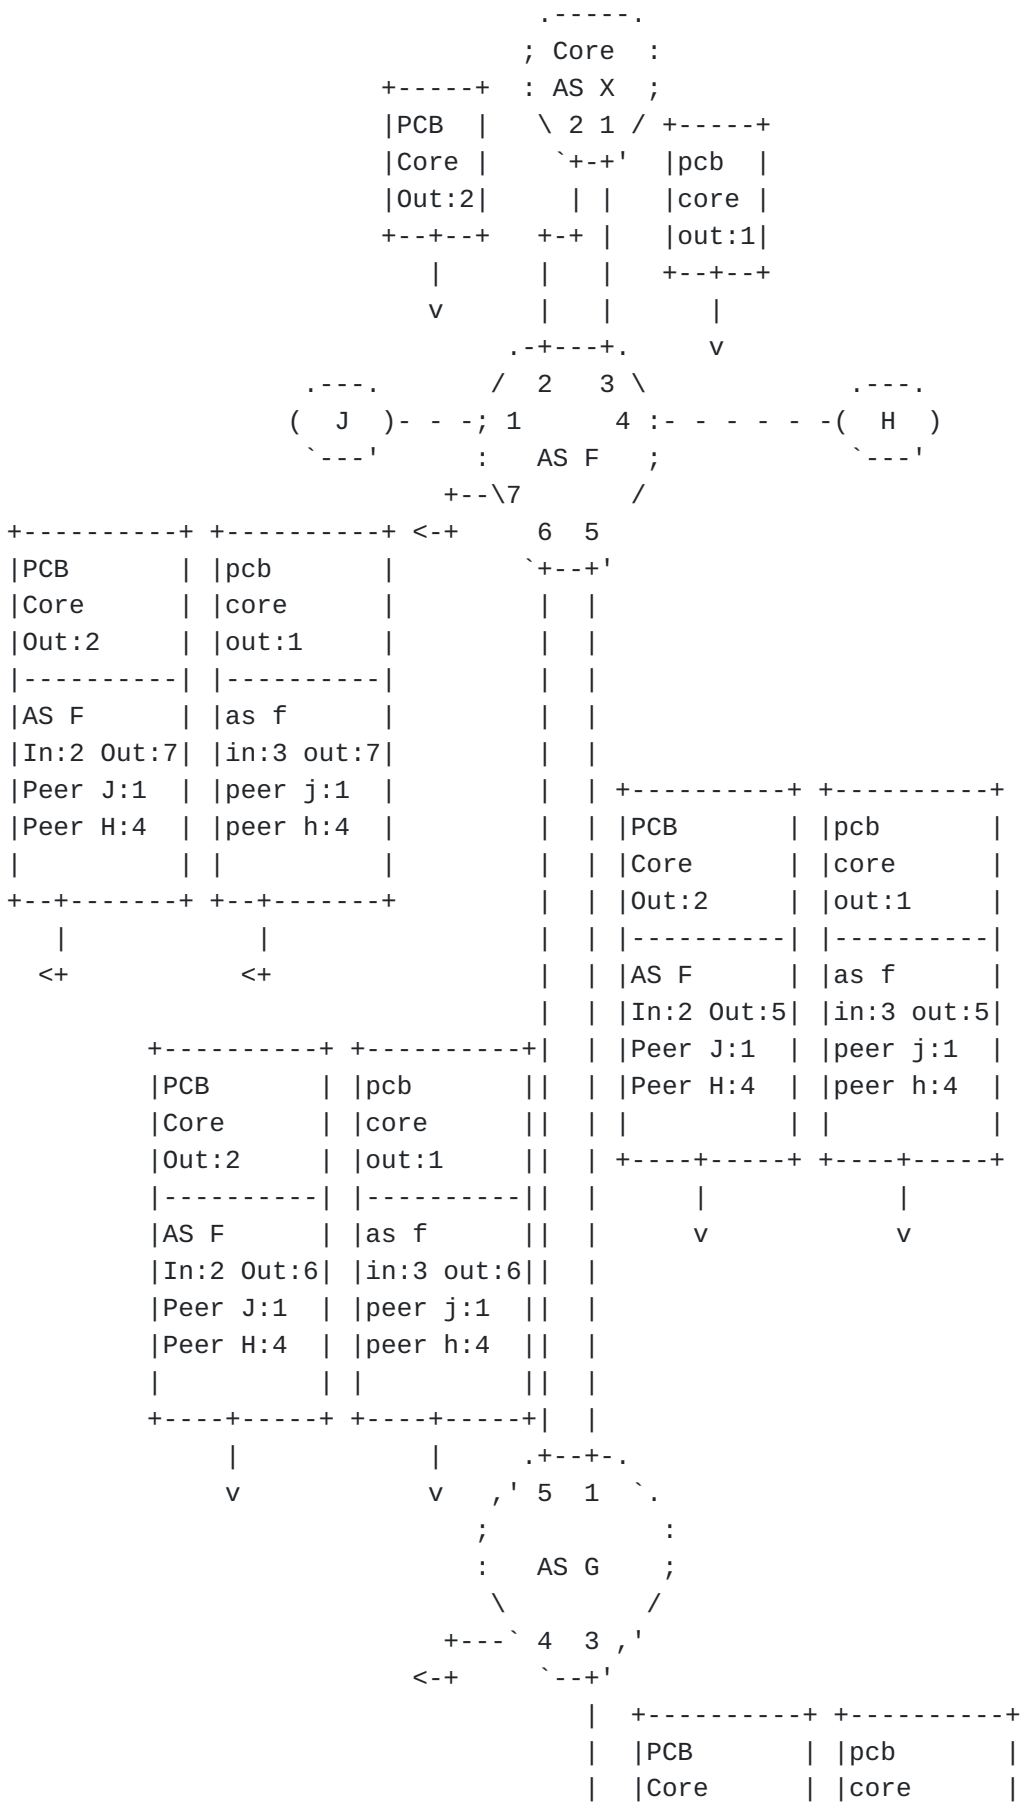
Core
AS X
PCB
Core
Out:2
pcb
core
out:1
J
H
AS F
pcb
core
out:1
as f
in:3 out:7
peer j:1
peer h:4
PCB
Core
Out:2
AS F
In:2 Out:5
Peer J:1
Peer H:4
pcb
core
out:1
as f
in:3 out:5
peer j:1
peer h:4
PCB
Core
Out:2
AS F
In:2 Out:6
Peer J:1
Peer H:4
pcb
core
out:1
as f
in:3 out:6
peer j:1
peer h:4
AS G
PCB
Core
pcb
core

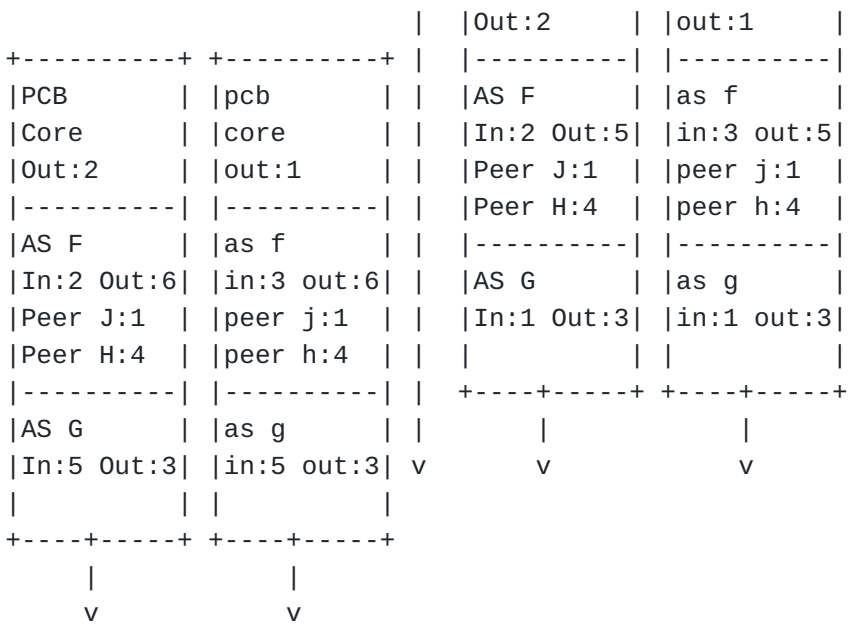

Figure 4: Intra-ISD PCB propagation from the ISD core down to child ASes

2.2.1.1. Security

Each PCB contains signatures of all on-path ASes. Every time a beacon service receives a PCB, it validates the PCB's authenticity. During this process, the beacon service can query the certificate service, for example, when it lacks an intermediate certificate.

2.2.1.2. Policies

Each AS can independently set policies dictating which PCBs are sent in which time intervals, and to which neighbors. In particular, PCBs do not need to be propagated immediately upon arrival. However, during bootstrapping and if the AS obtains a PCB containing a previously unknown path, the AS should forward the PCB immediately, to ensure quick connectivity establishment.

2.2.2. Path Registration

Both the beacon service and the path service are involved in the path registration process. A non-core AS typically receives several PCBs representing several path segments to various core ASes. Out of these PCBs, the non-core AS must select those down-path segments through which it wants to be reached. It is the task of the AS's beacon service to make this selection, according to the criteria described in Path-Segment Selection (Section 2.2.2.1). The beacon service then registers these path segments both at the local path service and at the path service of all core ASes. When links fail, segments expire, or better segments become available, the beacon service updates the down-path segments registered for its AS.

As a result, a core AS’s path service contains all intra-ISD path segments registered by the leaf ASes of its ISD. In addition, a core AS path service also stores the preferred core-path segments to other core ASes.

2.2.2.1. Path-Segment Selection

Among the received PCBs, the beacon service of an AS must choose (1) a set of PCBs to propagate further, and (2) a set of path segments to register. The selection of these PCBs and path segments is based on a path quality metric. This metric aims at identifying consistent, diverse, efficient, and policy-compliant paths:

- *Consistency* implies that at least one property along the path is uniform, such as an AS capability (e.g., high bandwidth).

- *Diversity* means that the set of paths announced over time are as path-disjoint as possible, in order to provide high-quality multipath options.

- *Efficiency* refers to the length, bandwidth, latency, utilization, and availability of a path, where more efficient paths are naturally preferred.

- *Policy compliance* implies that the path adheres to the AS's routing policy.

Based on past PCBs, the AS beacon service assigns scores to the current set of candidate path segments, and sends the best segments in the next beaconing interval.

Core beaconing operates similarly to intra-ISD beaconing, except that core PCBs only traverse core ASes. The same path selection metrics apply, where a core AS attempts to forward the set of most desirable paths to its neighbors.

2.2.3. Path Lookup

A host (source) who wants to start communication with another host (destination), requires up to three path segments: An up-path segment to reach the ISD core, a core-path segment to reach the destination ISD, and a down-path segment to reach the destination AS. The source host queries the path service in its AS for such segments. The path service has up-path segments stored in its database and furthermore checks if it has appropriate core- and down-path segments in its cache; in this case it returns them immediately.

If not, the path service in the source AS queries core path services (using locally stored up-path segments) in the source ISD for core-path segments to the destination ISD. Then, it combines up-path

segments with the newly retrieved core-path segments, and queries core path services in the remote ISD to fetch remote down-path segments. To improve overall efficiency, the local path service caches the returned path segments and uses parallelism when requesting path segments from core path services. Finally, the local path service returns all path segments to the source host.

This recursive lookup significantly simplifies the process for endpoints (which only have to send a single query, similar to stub DNS resolvers). The caching strategy ensures that path lookups are fast for frequently used destinations (similar to caching in recursive DNS resolvers).

2.2.4. Link Failures

Unlike in the current Internet, link failures are not automatically resolved by the network, but require active handling by endpoints. Since SCION forwarding paths are static, they break when one of the links fails. Link failures are handled by a two-pronged approach that typically masks link failures without any outage to the application and rapidly re-establishes fresh working paths:

* The SCION Control Message Protocol (SCMP) (the SCION equivalent of ICMP) is used for signaling connectivity problems. Instead of relying on application- or transport-layer timeouts, endpoints get immediate feedback from the network if a path stops working, and can quickly switch to an alternative path.

* SCION endpoints are encouraged to use multipath communication by default, thus masking a link failure with another working path. As multipath communication can increase availability (even in environments with very limited path choices), SCION beacon services attempt to create disjoint paths, SCION path services attempt to select and announce disjoint paths, and endpoints compose path segments to achieve maximum resilience to path failure. Consequently, most link failures in SCION remain unnoticed by the application, unlike the frequent (albeit mostly brief) outages in the current Internet. See also [ANDERSEN2001], [KATZ2012], [KUSHMAN2007], and [HITZ2021].

2.3. SCION Data Plane

While the control plane is responsible for providing end-to-end paths, the data plane ensures that packets are forwarded on the selected path. SCION border routers forward packets to the next AS based on the AS-level path in the packet header (which is extended with ingress and egress interface identifiers for each AS), without inspecting the destination address and also without consulting an inter-domain forwarding table. Only the border router at the

destination AS needs to inspect the destination address to forward it to the appropriate local endpoint.

Because SCION splits the information about the locator (the path towards the destination AS) and the identifier (the destination address), the identifier can have any format that the destination AS can interpret--only the destination needs to consider that local identifier (see also [RFC6830]). In other words, an AS can select an arbitrary addressing format for its hosts, e.g., a 4-byte IPv4, 6-byte media access control (MAC) address, 16-byte IPv6, or any other up to 16-byte addressing scheme. A valuable consequence is that hosts with different address types can directly communicate.

The next two sections describe how an endpoint combines path segments into an end-to-end forwarding path, and how border routers forward packets efficiently.

Note: This section describes the SCION data plane on a very high level. A much more detailed description of SCION's data plane will follow in a separate internet draft.

2.3.1. Path Construction via Segment Combination

Through the path lookup, the endpoint obtains path segments that must be combined into an end-to-end path. A valid SCION **forwarding path** can be created by combining up to three path segments, in the following ways:

* **Immediate combination of path segments**: The last AS on the up-path segment is also the first AS on the down-path segment. In this case, the simple combination of an up-path segment and a down-path segment creates a valid forwarding path.

* **AS shortcut**: The up-path segment and down-path segment intersect at a non-core AS. In this case, a shorter forwarding path can be created by removing the extraneous part of the path.

* **Peering shortcut**: A peering link exists between the two segments, so a shortcut via the peering link is possible. As in the AS shortcut case, the extraneous path segment is cut off. The peering link could be traversing to a different ISD.

* **Combination with a core-path segment**: The last AS on the up-path segment is different from the first AS on the down-path segment. This case requires an additional core-path segment to connect the up- and down-path segment. If the communication remains within the same ISD, a local ISD core-path segment is needed; otherwise, an inter-ISD core-path segment is required.

*   **On-path**: The destination AS is part of the up-path segment or the source AS is part of the down-path segment; in this case, a single up- or down-path segment, respectively, is sufficient to create a forwarding path.

Once a forwarding path is chosen, it is encoded in the SCION packet header. This makes inter-domain routing tables unnecessary for border routers: Both the ingress and the egress interface of each AS on the path are encoded as **packet-carried forwarding state (PCFS)** in the packet header. The destination can respond to the source by reversing the end-to-end path from the packet header, or it can perform its own path lookup and combination.

The SCION packet header contains of a sequence of **hop fields (HFs)**, one HF for each AS that is traversed on the end-to-end path. Each hop field contains the encoded numbers of the ingress and egress links, and thus defines which interfaces may be used to enter and leave an AS. In addition to the hop fields, each path segment contains an **info field (INF)** with basic information about the segment. A host can create an end-to-end forwarding path by extracting info fields and hop fields from path segments, as depicted in Figure 5. The additional meta header (META) contains pointers to the currently active INF and HF.

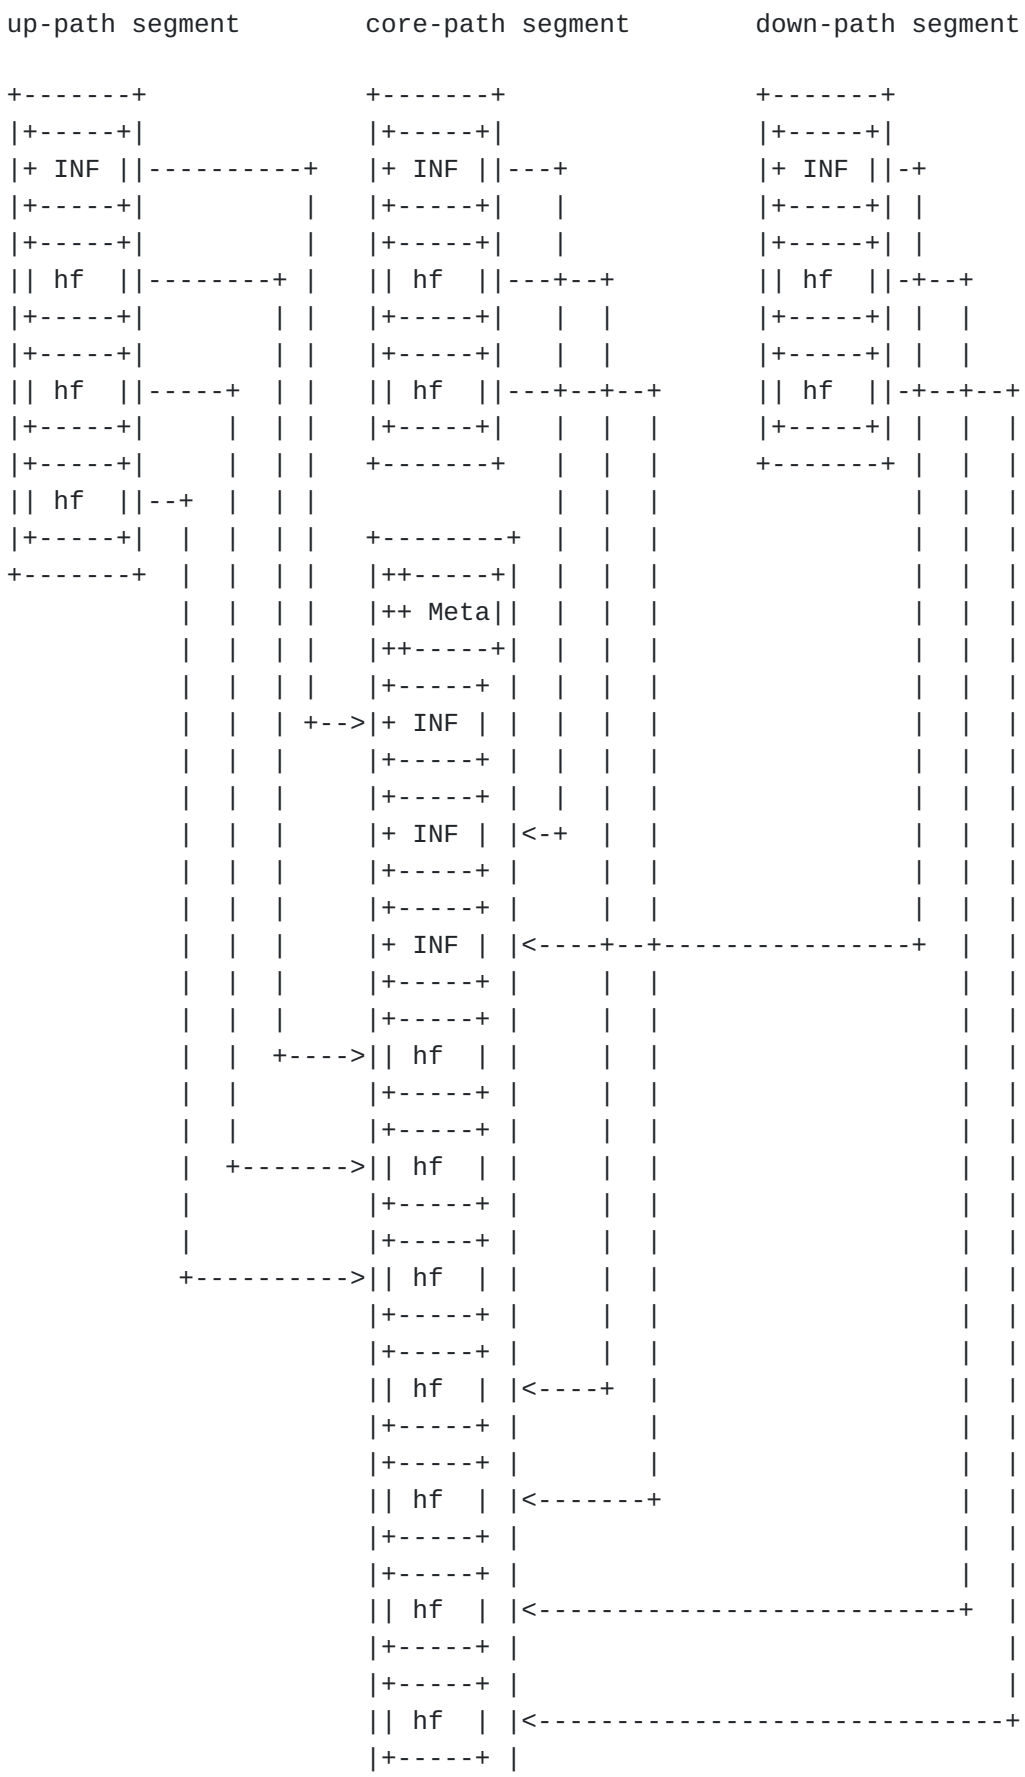
up-path segment
core-path segment
down-path segment
INF
hf
Meta

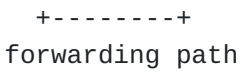

Figure 5: Combining three path segments into a forwarding path

2.3.2. Path Authorization

It is crucial for the data plane that endpoints only use paths constructed and authorized by ASes in the control plane. In particular, endpoints should not be able to craft HFs themselves, modify HFs in authorized path segments, or combine HFs of different path segments (path splicing). This property is called **path authorization** (see [KLENZE2021] and [LEGNER2020]).

SCION achieves path authorization by creating message-authentication codes (MACs) during the beaconing process. Each AS calculates these MACs using a local secret key (that is only shared between SCION infrastructure elements within the AS) and chains them to the previous HFs. The MACs are then included in the forwarding path as part of the respective HFs.

2.3.3. Forwarding

Routers can efficiently forward packets in the SCION architecture. In particular, the absence of inter-domain routing tables and of complex longest-IP-prefix matching performed by current routers enables the construction of more efficient routers.

During packet forwarding, a SCION border router at the ingress point of the AS verifies that:

- the packet entered through the correct ingress interface corresponding to the information in the HF,
- the HF is still valid, and
- the MAC in the HF is correct.

If the packet has not yet reached the destination AS, the egress interface number in the HF of the non-destination AS refers to the egress SCION border router of this AS. In this case, the packet can be sent from the ingress SCION border router to the egress SCION border router via native intra-domain forwarding (e.g., IP or MPLS). In case the packet has arrived at the destination AS, the destination AS's border router inspects the destination address and sends the packet to the corresponding host.

2.3.4. Intra-AS Communication

SCION routers use IP to communicate within an AS, therefore they rely on existing intra-domain routing protocols, such as Multiprotocol Label Switching (MPLS) or others.

3. Deployment

Adoption of a next-generation architecture is a challenging task, as it needs to be integrated with, and operate alongside existing infrastructure. SCION is designed to coexist with existing intra-domain routing infrastructure, and comprises coexistence and transition mechanisms that facilitate adoption, in accordance to principles defined in [RFC8170]. The following section discusses practical considerations for deploying SCION and briefly touches on some of the transition mechanisms, with focus on:

- Autonomous Systems (Section 3.1),
- Internet Exchange Points (Section 3.2), and
- endpoints (Section 3.3), covering both native SCION hosts and SCION to IP encapsulation.

We then describe some of the early adopters deployment experiences. A more detailed adoption plan is to be outlined in dedicated documents.

3.1. Autonomous System Deployment

A SCION AS needs to deploy the SCION infrastructure components (Section 1.2.3) and border routers. Within an AS, SCION is often deployed as an IP overlay on top of the existing network. This way SCION allows to reuse the existing intra-domain network and equipment (e.g., IP, MPLS). Customer-side SCION border routers directly connect to the provider-side border routers using last-mile connections. The SCION design assumes that AS's internal entities are considered to be trustworthy, therefore the IP overlay or the first-hop routing does not compromise or degrade any security properties SCION delivers. When it comes to inter-domain communication, an overlay deployment on top of today's Internet is not desirable, as SCION would inherit issues from its weak underlay. Thus, inter-AS SCION links are usually deployed in parallel to existing links, in order to preserve its security properties. That is, two SCION border routers from neighbour ASes are directly connected via a layer-2 cross-connection at a common point-of-presence.

All SCION AS components can be deployed on standard x86 commercial off-the-shelf servers or virtual machines. In fact, SCION border routers do not rely on forwarding tables, therefore they do not require specialized hardware. Practice shows that off-the-shelf

hardware can handle up to 100 Gbps links, while a prototype P4 implementation [DERUITER2021] showed that it is possible to forward SCION traffic even at terabit speeds.

Overall, an AS can be connected to SCION without high-impact changes to its network. In addition, use of commodity hardware for both control and data-plane components reduces initial deployment costs.

3.2. Internet Exchange Points

Internet Exchange Points (IXP) play as important a role for SCION as they do in today's Internet. SCION can be deployed at existing IXPs following a "big switch" model, where the IXP provides a large L2 switch between multiple SCION ASes. SCION has been deployed following this model at the Swiss Internet Exchange (SwissIX), currently interconnecting major SCION Swiss ISPs and enterprises through bi-lateral peering over dedicated SCION ports.

Additionally, thanks to its path-awareness, SCION offers the option of an enhanced deployment model, i.e., to expose the internal topology of an IXP within the SCION control plane. This enables IXP customers to use SCION’s multipath and fast failover capabilities to leverage the IXP’s internal links (including backup links) and to select paths depending on the application’s needs. IXPs have therefore an incentive to expose their rich internal connectivity, as the benefits from SCION’s multipath capabilities would increase their value for customers and provide them with a competitive advantage.

3.3. Endpoints and Incremental Deployability

End users can leverage SCION in two different ways: (1) using SCION-aware applications on a SCION native endpoint (Section 3.3.1), or (2) using transparent IP-to-SCION conversion (Section 3.3.2). The benefit of using SCION natively is that the full range of advantages becomes available to applications, at the cost of installing the SCION endpoint stack and making the application SCION-aware. In early deployments, the second approach is often preferred, so that no changes are needed within applications or endpoints.

3.3.1. Native Endpoints

A SCION native endpoint's stack consists of a dispatcher, which handles all incoming and outgoing SCION packets, and of a SCION daemon, which handles control-plane messages. The latter fetches paths to remote ASes and provides an API for applications and libraries to interact with the SCION control plane (i.e., for path lookup, SCION extensions). The current SCION implementation uses an UDP/IP underlay for communication between endpoints and SCION routers. This allows reuse of existing intra-domain networking infrastructure. SCION endpoints can optionally use automated

bootstrapping mechanisms to retrieve configuration from the network and establish SCION connectivity. This way, clients require no pre-existing network-specific configurations.

3.3.2. SCION to 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. A SIG can be deployed close to the end user (i.e., at branches of an enterprise, on a CPE), or within an ISP's network. In the latter case, the SIG is called carrier-grade SIG, as it serves multiple customers within the AS where it is deployed. This approach has the advantage that it does not require any changes at the customer's premises. In order to allow incremental deployability and to ease transition from legacy IP-based Internet to SCION, SIGs can be augmented with mechanisms allowing them to coordinate and automatically exchange IP prefix information. A more detailed description of the SIG and its coordination mechanisms is to be presented in dedicated documents.

3.4. Deployment Experiences

SCION has been deployed in production by multiple entities, growing its acceptance among industry. While early deployments started on academic and research networks, SCION has expanded to serve the financial industry, government, and it is being evaluated for the healthcare sector.

In 2017, SCION was evaluated for production use by a central bank, with the goal of modernising the network interconnecting banks and their branches. SCION was chosen, as it allows moving away from a dedicated private network to a reliable Internet-based solution. SCION connectivity was later extended to support system-critical applications, like the national real-time gross settlement (RTGS) system, connecting all country's banks to exchange real-time payment information. The network, called Secure Swiss Finance Network or SSFN, is implemented as a SCION ISD, where a federation of three ISPs forms the ISD core. Financial institutions are themselves SCION ASes and directly connect to one or more of the core ASes. Institutions deploy SCION–IP gateways (SIGs), transparently enabling their traditional IP-based applications to use the SCION network. The concept of the SCION ISD also provides a mechanism to implement strict governance and access control (through the issuance of AS certificates).

Besides the SSFN, SCION connectivity has also been adopted by government entities for their international communications. In addition, Swiss higher education institutions are connected thanks to the SCI-ED network.

In addition to productive deployments, SCION also comprises a global SCION research testbed called SCIONLab. It is composed of dozens of globally distributed infrastructure ASes, mostly run by academic institutions. The testbed is open to any user who can easily set up their own AS with the aid of a web-based UI, connect to the network, and run experiments. The setup has been the earliest global deployment of SCION and it has been supporting research and development of path-aware networking and SCION.

4. IANA Considerations

Currently, this document has no request for action to IANA. However, when full specification of SCION is available, requests for IANA actions are expected regarding the registration of optional packet header fields as well as the coordination of SCION ISD and AS number assignments.

5. Security Considerations

SCION has been designed from the outset to offer security by default, and thus there are manifold security considerations. As a matter of fact, SCION's protocol design has been formally verified and the open source router implementation is undergoing formal verification (see also [KLENZE2021]). Describing all security considerations here, therefore, would go beyond the scope of this document. A separate document including all security implications and considerations will follow later.

6. References

6.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/rfc/rfc2119>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/rfc/rfc8174>.

6.2. Informative References

[ANDERSEN2001] Andersen, D., Balakrishnan, H., Kaashoek, F., and R. Morris, "Resilient overlay networks", Proceedings of the eighteenth ACM symposium on Operating systems principles,

DOI 10.1145/502034.502048, October 2001, <https://doi.org/10.1145/502034.502048>.

[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>.

[DERUITER2021] de Ruiter, J. and C. Schutijser, "Next-generation internet at terabit speed: SCION in P4", Proceedings of the 17th International Conference on emerging Networking EXperiments and Technologies, DOI 10.1145/3485983.3494839, December 2021, <https://doi.org/10.1145/3485983.3494839>.

[GRIFFIN1999] Griffin, T. and G. Wilfong, "An analysis of BGP convergence properties", ACM SIGCOMM Computer Communication Review vol. 29, no. 4, pp. 277-288, DOI 10.1145/316194.316231, August 1999, <https://doi.org/10.1145/316194.316231>.

[HITZ2021] Hitz, S., "Demonstrating the reliability and resilience of Secure Swiss Finance Network", 2021, <https://perma.cc/4H3Q-WZNG>.

[I-D.dekater-scion-pki] de Kater, C. and N. Rustignoli, "SCION Control-Plane PKI", 2023, <https://datatracker.ietf.org/doc/draft-dekater-scion-pki/>.

[I-D.rustignoli-scion-components] de Kater, C. and N. Rustignoli, "SCION Components Analysis", 2023, <https://datatracker.ietf.org/doc/draft-rustignoli-panrg-scion-components/>.

[KATZ2012]
Katz-Bassett, E., Scott, C., Choffnes, D., Cunha, Í., Valancius, V., Feamster, N., Madhyastha, H., Anderson, T., and A. Krishnamurthy, "LIFEGUARD: practical repair of persistent route failures", ACM SIGCOMM Computer Communication Review vol. 42, no. 4, pp. 395-406, DOI 10.1145/2377677.2377756, August 2012, <https://doi.org/10.1145/2377677.2377756>.

[KING2022] King, D., Farrel, A., and C. Jacquenet, "Challenges for the Internet Routing Systems Introduced by Semantic

Routing", 2022, <https://datatracker.ietf.org/doc/draft-king-irtf-challenges-in-routing/>.

[KLENZE2021] Klenze, T., Sprenger, C., and D. Basin, "Formal Verification of Secure Forwarding Protocols", 2021 IEEE 34th Computer Security Foundations Symposium (CSF), DOI 10.1109/csf51468.2021.00018, June 2021, <https://doi.org/10.1109/csf51468.2021.00018>.

[KUSHMAN2007] Kushman, N., Kandula, S., and D. Katabi, "Can you hear me now?!: it must be BGP", ACM SIGCOMM Computer Communication Review vol. 37, no. 2, pp. 75-84, DOI 10.1145/1232919.1232927, March 2007, <https://doi.org/10.1145/1232919.1232927>.

[LABOVITZ2000] Labovitz, C., Ahuja, A., Bose, A., and F. Jahanian, "Delayed Internet routing convergence", Proceedings of the conference on Applications, Technologies, Architectures, and Protocols for Computer Communication, DOI 10.1145/347059.347428, August 2000, <https://doi.org/10.1145/347059.347428>.

[LEGNER2020] Legner, M., Klenze, T., Wyss, M., Sprenger, C., and A. Perrig, "EPIC: Every Packet Is Checked in the Data Plane of a Path-Aware Internet", 2020, <https://www.usenix.org/conference/usenixsecurity20/presentation/legner>.

[LI2014] Li, Q., Hu, Y., and X. Zhang, "Even Rockets Cannot Make Pigs Fly Sustainably: Can BGP be Secured with BGPsec?", Proceedings 2014 Workshop on Security of Emerging Networking Technologies, DOI 10.14722/sent.2014.23001, 2014, <https://doi.org/10.14722/sent.2014.23001>.

[LYCHEV2013] Lychev, R., Goldberg, S., and M. Schapira, "BGP security in partial deployment: is the juice worth the squeeze?", ACM SIGCOMM Computer Communication Review vol. 43, no. 4, pp. 171-182, DOI 10.1145/2534169.2486010, August 2013, <https://doi.org/10.1145/2534169.2486010>.

[MORILLO2021] Morillo, R., Furuness, J., Morris, C., Breslin, J., Herzberg, A., and B. Wang, "ROV++: Improved Deployable Defense against BGP Hijacking", Proceedings 2021 Network and Distributed System Security Symposium, DOI 10.14722/ndss.2021.24438, 2021, <https://doi.org/10.14722/ndss.2021.24438>.

[PERRIG2017] Perrig, A., Szalachowski, P., Reischuk, R., and L. Chuat, "SCION: A Secure Internet Architecture", ISBN 978-3-319-67079-9, 2017, <https://doi.org/10.1007/978-3-319-67080-5>.

[RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, DOI 10.17487/RFC4033, March 2005, <https://www.rfc-editor.org/rfc/rfc4033>.

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

[RFC5218] Thaler, D. and B. Aboba, "What Makes for a Successful Protocol?", RFC 5218, DOI 10.17487/RFC5218, July 2008, <https://www.rfc-editor.org/rfc/rfc5218>.

[RFC6480] Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", RFC 6480, DOI 10.17487/RFC6480, February 2012, <https://www.rfc-editor.org/rfc/rfc6480>.

[RFC6830] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "The Locator/ID Separation Protocol (LISP)", RFC 6830, DOI 10.17487/RFC6830, January 2013, <https://www.rfc-editor.org/rfc/rfc6830>.

[RFC8170] Thaler, D., Ed., "Planning for Protocol Adoption and Subsequent Transitions", RFC 8170, DOI 10.17487/RFC8170, May 2017, <https://www.rfc-editor.org/rfc/rfc8170>.

[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>.

[RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <https://www.rfc-editor.org/rfc/rfc8446>.

[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>.

[RFC9217] Trammell, B., "Current Open Questions in Path-Aware Networking", RFC 9217, DOI 10.17487/RFC9217, March 2022, <https://www.rfc-editor.org/rfc/rfc9217>.

[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, no. 7-10, 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, DOI 10.1145/1866307.1866411, October 2010, <https://doi.org/10.1145/1866307.1866411>.

Acknowledgments

Many thanks go to Cyrill Krähenbühl and Juan A. Garcia-Pardo for reviewing this document. We are also indebted to Laurent Chuat, Markus Legner, David Basin, David Hausheer, Samuel Hitz, and Peter Müller, for writing the book "The Complete Guide to SCION" (see [CHUAT22]), which provides the background information needed to write this informational draft.

Authors' Addresses

Corine de Kater
SCION Association

Email: cdk@scion.org

Nicola Rustignoli
SCION Association

Email: nic@scion.org

Adrian Perrig
ETH Zuerich

Email: adrian.perrig@inf.ethz.ch