

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 21 September 2026

C. de Kater
N. Rustignoli
SCION Association
S. Hitz
Anapaya Systems
20 March 2026

SCION Control Plane
draft-dekater-scion-controlplane-16

Abstract

This document describes the Control Plane of the path-aware, inter-domain network architecture SCION (Scalability, Control, and Isolation On Next-generation networks). A fundamental characteristic of SCION is that it gives path control to SCION-capable endpoints that can choose between multiple path options, thereby enabling the optimization of network paths. The SCION Control Plane is responsible for discovering these paths and making them available to the endpoints.

The SCION Control Plane creates and securely disseminates path segments between SCION Autonomous Systems (AS) which can then be combined into forwarding paths to transmit packets in the data plane. This document describes mechanisms of path exploration through beaconing and path registration. In addition, it describes how Endpoints construct end-to-end paths by combining path segments obtained through a path lookup process.

This document contains new approaches to secure path aware networking. It is not an Internet Standard, has not received any formal review of the IETF, nor was the work developed through the rough consensus process. The approaches in this work are offered to the community for its consideration in the further evolution of the Internet.

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-cp_I-D/draft-dekater-scion-controlplane.html. Status information for this document may be found at <https://datatracker.ietf.org/doc/draft-dekater-scion-controlplane/>.

Source for this draft and an issue tracker can be found at https://github.com/scionassociation/scion-cp_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 21 September 2026.

Copyright Notice

Copyright (c) 2026 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Table of Contents

1. Introduction	5
1.1. Terminology	6
1.2. Conventions and Definitions	8
1.3. Paths and Links	8
1.4. Routing	10
1.4.1. Path Segments	12
1.5. Addressing	13
1.5.1. ISD Numbers	13
1.5.2. SCION AS Numbers	14
1.5.3. Text Representation	15
1.6. Bootstrapping ability	16
1.7. Communication Protocol	16
2. Path Exploration or Beaconing	17
2.1. Introduction and Overview	17
2.1.1. Peering Links	18

2.1.2.	Appending Entries to a PCB	18
2.1.3.	PCB Propagation - Illustrated Examples	18
2.2.	PCB Message Format	23
2.2.1.	Segment Information	24
2.2.2.	AS Entry	25
2.2.3.	PCB Extensions	35
2.2.4.	PCB Validity	35
2.2.5.	Configuration	36
2.3.	Propagation of PCBs	37
2.3.1.	Reception of PCBs	37
2.3.2.	Storing Candidate PCBs	38
2.3.3.	PCB Selection Policies	38
2.3.4.	Propagation Interval and Best PCBs Set Size	39
2.3.5.	Propagation of Selected PCBs	40
2.4.	Distribution of Cryptographic Material	42
2.5.	Renewal of Cryptographic Material	43
3.	Registration of Path Segments	44
3.1.	Intra-ISD Path Segment Registration	45
3.1.1.	Terminating a PCB	45
3.1.2.	Transforming a PCB into an Up Segment	46
3.1.3.	Transforming a PCB into a Down Segment	47
3.2.	Core Path Segment Registration	47
3.3.	Path Segment Registration RPC API	48
3.4.	Path MTU	49
4.	Path Lookup	49
4.1.	Lookup Process	49
4.1.1.	Sequence of Lookup Requests	51
4.1.2.	Lookup Requests Message Format	51
4.1.3.	Caching	52
4.2.	Behavior of Actors in the Lookup Process	53
4.2.1.	Use of Wildcard Addresses in the Lookup Process	53
4.2.2.	Segment-Request Handler of a Non-Core Source AS	54
4.2.3.	Segment-Request Handler of a Core AS	55
5.	Control Service Discovery	55
6.	SCMP	57
6.1.	General Format	57
6.2.	Message Types	58
6.3.	Checksum Calculation	59
6.4.	Processing Rules	59
6.5.	Error Messages	60
6.5.1.	Packet Too Big	60
6.5.2.	External Interface Down	61
6.5.3.	Internal Connectivity Down	62
6.6.	Informational Messages	63
6.6.1.	Echo Request	63
6.6.2.	Echo Reply	64
6.6.3.	Traceroute Request	65
6.6.4.	Traceroute Reply	66

7.	Deployment Considerations	67
7.1.	Destination Mapping	68
7.2.	Monitoring Considerations	68
7.3.	Effects of Clock Inaccuracy	69
7.4.	Path Discovery Time and Scalability	70
7.4.1.	Intra-ISD Beaconsing	71
7.4.2.	Core Beaconsing	72
8.	Security Considerations	73
8.1.	Security Properties	74
8.2.	Manipulation of the Beaconsing Process by a Core Adversary	75
8.3.	Manipulation of the Beaconsing Process by a Non-Core Adversary	75
8.3.1.	Path Hijacking through Interposition	75
8.3.2.	Creation of Spurious ASes and ISDs	76
8.3.3.	Peering Link Misuse	76
8.3.4.	Manipulation of the Path-Selection Process	77
8.4.	Attacks on time sources	79
8.5.	Denial of Service Attacks	79
9.	IANA Considerations	80
10.	References	80
10.1.	Normative References	80
10.2.	Informative References	82
	Path-Lookup Examples	83
	Change Log	87
	draft-dekater-scion-controlplane-15	87
	draft-dekater-scion-controlplane-15	87
	draft-dekater-scion-controlplane-14	88
	draft-dekater-scion-controlplane-13	89
	draft-dekater-scion-controlplane-12	89
	draft-dekater-scion-controlplane-11	89
	draft-dekater-scion-controlplane-10	89
	draft-dekater-scion-controlplane-09	90
	draft-dekater-scion-controlplane-08	91
	draft-dekater-scion-controlplane-07	91
	draft-dekater-scion-controlplane-06	92
	draft-dekater-scion-controlplane-05	92
	draft-dekater-scion-controlplane-04	92
	Acknowledgments	93
	Authors' Addresses	93

1. Introduction

SCION (Scalability, Control, and Isolation On Next-generation networks) is a path-aware internetworking routing architecture as described in [RFC9217]. It allows endpoints and applications to select paths across the network to use for traffic, based on trustworthy path properties. SCION is an inter-domain network architecture and is therefore not concerned with intra-domain forwarding.

SCION has been developed with the following goals:

Availability - to provide highly available communication that can send traffic over paths with optimal or required characteristics, quickly handle inter-domain link or router failures (both on the last hop or anywhere along the path), and provide continuity in the presence of adversaries.

Security - to introduce a new approach to inter-domain path security that leverages path awareness in combination with a unique trust model. The goal is to provide higher levels of trustworthiness in routing information to prevent traffic hijacking, and to enable users to decide where their data travels. This routing information can be unambiguously attributed to a SCION AS, thereby ensuring that packets are only forwarded along authorized path segments (e.g., to enable geofencing). Security properties are further discussed in Section 8.1.

Scalability - to improve the scalability of the inter-domain control plane and data plane, avoiding existing limitations related to convergence and forwarding table size. The advertising of path segments is separated into a beaconing process within each Isolation Domain (ISD), and between ISDs which incurs minimal overhead and resource requirements on routers.

SCION relies on three main components:

PKI - described in [I-D.dekater-scion-pki]. To achieve scalability and trust, SCION organizes its ASes into logical groups of independent routing planes called _Isolation Domains (ISDs)_. All ASes in an ISD agree on a set of trust roots called the _Trust Root Configuration (TRC)_ which is a collection of signed root certificates in X.509 v3 format [RFC5280]. The ISD is governed by a set of _core ASes_ which typically manage the trust roots and provide connectivity to other ISDs. This is the basis of the public key infrastructure used for the authentication of messages used by the SCION Control Plane.

Control Plane - described in this document. It performs inter-domain routing by discovering and securely disseminating path information between SCION ASes. The core ASes use Path-segment Construction Beacons (PCBs) to explore intra-ISD paths, or to explore paths across different ISDs.

Data Plane - described in [I-D.dekater-scion-dataplane]. It carries out secure packet forwarding between SCION ASes over paths selected by endpoints. A SCION border router reuses existing intra-domain infrastructure to communicate to other SCION routers or SCION endpoints within its AS.

This document should be read in conjunction with the other components mentioned above. Readers are encouraged to read the introduction in this document first.

The SCION architecture was initially developed outside of the IETF by ETH Zurich with significant contributions from Anapaya Systems. It is deployed in the Swiss finance sector to provide resilient connectivity between financial institutions. The aim of this document is to document the existing protocol specification as deployed, to encourage interoperability among implementations, and to introduce new concepts that can potentially be further improved to address particular problems with the current Internet architecture. This document is not an Internet Standards Track specification; it does not have IETF consensus and it is published for informational purposes.

1.1. Terminology

***SCION Autonomous System (AS)*:** A SCION Autonomous System is a network under a common administrative control. For example, the network of a SCION service provider, company, or university can constitute an AS. While functionally similar to a BGP AS, a SCION AS operates within an Isolation Domain (ISD), utilizes a different address scheme, and serves as a locator in the addressing of end hosts. References to ASes throughout this document refer to SCION ASes.

***Beaconing*:** The Control Plane process where an AS discovers paths to other ASes.

***Control Plane*:** The SCION Control Plane is responsible for the propagation and discovery of network paths, i.e., for the exchange of routing information between SCION nodes. The Control Plane thus determines where traffic can be sent and deals with questions such as how paths are discovered, which paths exist, and how they are disseminated to endpoints, etc. Within a SCION AS, such functionalities are carried out by the Control Service whereas packet forwarding is a task carried out by the data plane.

***Control Service*:** The Control Service is the main control plane infrastructure component within a SCION AS. It is responsible for the path exploration and registration processes that take place within the Control Plane.

***Core AS*:** Each Isolation Domain (ISD) is administered by a set of distinguished SCION autonomous systems (ASes) called core ASes, which are responsible for initiating the path discovery and path construction process (called "beaconing" in SCION). Each ISD MUST have at least one Core AS.

***Endpoint*:** An endpoint is the start or the end of a SCION path, as defined in [RFC9473].

***Forwarding Path*:** A complete end-to-end path between two SCION endpoints which is used to transmit packets in the data plane. Endpoints can create paths with a combination of up to three path segments (an up segment, a core segment, and a down segment).

***Hop Field (HF)*:** As they traverse the network, Path-Segment Construction Beacons (PCBs) accumulate cryptographically protected AS-level path information in the form of Hop Fields. In the data plane, Hop Fields are used for packet forwarding: they contain the incoming and outgoing Interface IDs of the ASes on the forwarding path.

***Info Field (INF)*:** Each Path-Segment Construction Beacon (PCB) contains a single Info field, which provides basic information about the PCB. Together with Hop Fields (HFs), these are used to create forwarding paths.

***Isolation Domain (ISD)*:** SCION ASes are organized into logical groups called Isolation Domains or ISDs. Each ISD consists of ASes that span an area with a uniform trust environment (e.g., a common jurisdiction).

Message Authentication Code (MAC). In the rest of this document, "MAC" always refers to "Message Authentication Code" and never to "Medium Access Control". When "Medium Access Control address" is implied, the phrase "Link Layer Address" is used.

Path Segment: These are derived from Path-Segment Construction Beacons (PCBs). A path segment can be (1) an up segment (i.e., a path between a non-core AS and a core AS in the same ISD), (2) a down segment (i.e., the same as an up segment but in the opposite direction), or (3) a core segment (i.e., a path between core ASes, possibly traversing ISD boundaries). Endpoints use up to three path segments to create a forwarding path.

Path-Segment Construction Beacon (PCB): Core AS Control Service instances generate PCBs to explore paths within their isolation domain (ISD) and between different ISDs. ASes further propagate selected PCBs to their neighboring ASes. These PCBs traverse each AS accumulating information, including Hop Fields (HFs) which can subsequently be used by the data plane for forwarding.

SCION Control Message Protocol (SCMP): A signaling protocol analogous to the Internet Control Message Protocol (ICMP). This is described in Section 6.

Trust Root Configuration (TRC): A Trust Root Configuration or TRC is a signed collection of certificates pertaining to an isolation domain (ISD). TRCs also contain ISD-specific policies.

1.2. 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.

1.3. Paths and Links

SCION routers and endpoints connect to each other via links. A link refers to a physical or logical connection between two SCION nodes (e.g., router or endpoint). A SCION path between two endpoints traverses one or more links.

SCION ASes - each being a network under a common administrative control - are organized into logical groups called Isolation Domains (ISDs). Each ISD consists of ASes that are part of a uniform trust environment (i.e., a common jurisdiction) and is administered by a set of distinguished ASes called core ASes.

SCION distinguishes three types of links between ASes: (1) core links, (2) parent-child links, and (3) peering links.

- * _Core_ links connect two core ASes, which are either within the same or in different ISDs. Core links can exist for various reasons, including provider-customer (where the customer pays the provider for traffic) and peering relationships.
- * _Parent-child_ links create a hierarchy between the parent and the child AS within the same ISD. ASes with a parent-child link typically have a provider-customer relationship.
- * _Peering_ links exist between ASes with a peering relationship (settlement-free or paid). They can be established between any two ASes (core or non-core) and can cross ISD boundaries.

SCION paths are comprised of at most three path segments: an up segment, traversing links from child to parent, then a core segment consisting of core links, followed by a down segment traversing links from parent to child. Each path segment is established over one or more links.

SCION paths are always "valley free" whereby a child AS does not carry transit traffic from a parent AS to another parent AS. These paths can contain at most one peering link, which endpoints can use as shortcut between two path segments containing two peer ASes.

Figure 1 shows the three types of links for one small ISD with two core ASes A and C, and four non-core ASes D,E,F, and G.

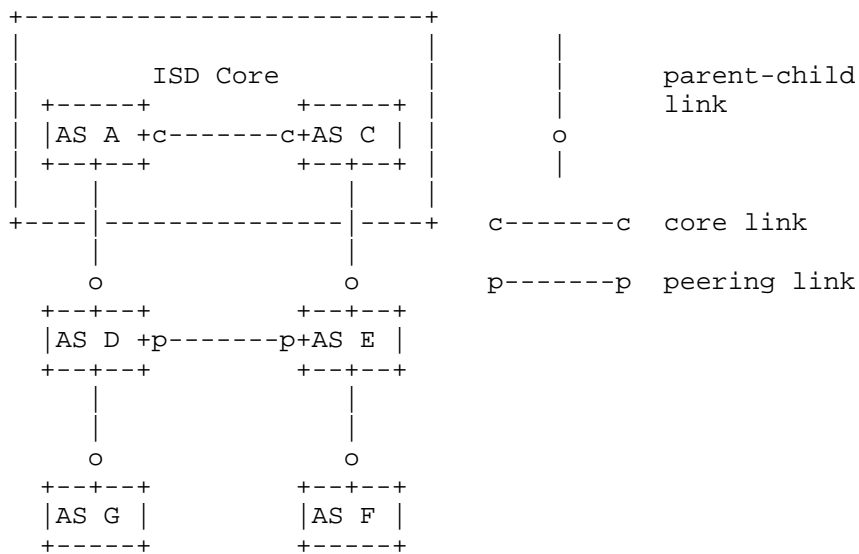


Figure 1: The three types of SCION links in one ISD. Each node in the figure is a SCION AS.

Each link connecting SCION routers is bi-directional and is identified by its corresponding egress and ingress Interface IDs. An Interface ID is a 16-bit identifier as described in [I-D.dekater-scion-dataplane] in the "Terminology" section. It is required to be unique within each AS and can therefore be chosen without any need for coordination between ASes.

1.4. Routing

SCION provides path-aware inter-domain routing between SCION ASes. The SCION Control Plane is responsible for discovering these inter-domain paths and making them available to the endpoints within the ASes.

SCION inter-domain routing operates on two levels: within an ISD which is called intra-ISD routing, and between ISDs which is called inter-ISD routing. Both levels use Path-Segment Construction Beacons (PCBs) to explore network paths. A PCB is originated 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 at an AS level and store this information in the form of Hop Fields. Endpoints use information from these Hop Fields to

create end-to-end forwarding paths for data packets that carry this information in their headers. This also supports multi-path communication among endpoints.

The creation of end-to-end forwarding paths consists of the following processes:

- * _Path exploration (or beaconing)_: This is the process where an AS Control Service discovers paths to other ASes. This is described in detail in Section 2.
- * _Path registration_: This is the process where an AS Control Service selects PCBs, according to defined policies, turns the selected PCBs into path segments, and adds these path segments to the relevant path infrastructure, thus making them available to other ASes. This is described in detail in Section 3.
- * _Path resolution_: This is the process of actually creating an end-to-end forwarding path from the source endpoint to the destination. 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. Step (a) is described in detail in Section 4 and step (b) is described in [I-D.dekater-scion-dataplane] section "Path Construction (Segment Combinations)".

All processes operate concurrently.

The **Control Service** is responsible for the path exploration and registration processes in the Control Plane. It is the main control plane infrastructure component within each SCION AS and has the following tasks:

- * Generating, receiving, and propagating PCBs. Periodically, the Control Service of a core AS generates a set of PCBs, which are forwarded to the child ASes or neighboring core ASes. In the latter case, the PCBs are sent over policy compliant paths to discover multiple paths between any pair of core ASes.
- * Selecting and registering the set of path segments via which the AS wants to be reached.

- * Distributing certificates and keys to secure inter-AS communication. Each PCB contains signatures of all on-path ASes and each time the Control Service of an AS receives a PCB, it validates the PCB's authenticity. When the Control Service lacks an intermediate certificate, it can query the Control Service of the neighboring AS that sent the PCB through the API described in Section 2.4.

Note: The Control Service of an AS is decoupled from SCION border routers and operators may deploy it anywhere within the AS.

1.4.1. Path Segments

SCION distinguishes the following types of path segments:

- * A path segment from a non-core AS to a core AS is an `_up segment_`.
- * A path segment from a core AS to a non-core AS is a `_down segment_`.
- * A path segment between core ASes is a `_core segment_`.

Each path segment starts and/or ends at a core AS. Path segments are not created between non-core ASes.

All path segments may be reversed: a core segment can be used bi-directionally, an up segment can be converted into a down segment, and a down segment can be converted into an up segment depending on the direction of the end-to-end path. This means that all path segments can be used to send data traffic in both directions.

The cryptographic protection of PCBs and path segments is based on the Control Plane PKI. The signatures are structured such that the entire message sequence constituting the path segment can be authenticated by anyone with access to this PKI.

For fast validation of the path information carried in individual packets during packet forwarding, symmetric key cryptography is used and the Hop Fields contain a MAC. These MACs are structured to allow verification of the sequence of hops, but in contrast to the PCBs can only be validated by the border routers of the respective AS.

1.5. Addressing

Inter-domain SCION routing is based on an <ISD, AS> tuple. Although a complete SCION address is composed of the <ISD, AS, endpoint address> 3-tuple, the endpoint address is not used for inter-domain routing or forwarding. The endpoint address is of variable length, does not need to be globally unique, and can thus be an IPv4, IPv6, or link layer address.

Note: As a consequence of the fact that SCION relies on existing routing protocols (e.g., IS-IS, OSPF, SR) and communication fabric (e.g., IP, MPLS) for intra-domain forwarding, existing internal routers do not need to be changed to support SCION.

1.5.1. ISD Numbers

An ISD number is the 16-bit global identifier for an ISD.

The following table gives an overview of the ISD number allocation:

ISD	Description
0	The wildcard ISD
1 - 15	Reserved for documentation and sample code (analogous to [RFC5398]).
16 - 63	Private use (analogous to [RFC6996]) - can be used for testing and private deployments
64 - 4094	Public ISDs. They MUST be globally unique.
4095 - 65535	Unallocated

Table 1: ISD number allocations

The wildcard ISD is not directly used by the control or data plane. Implementations may use it to represent any ISD, for example in path filters. ISD numbers are allocated by the SCION Association ([ISD-AS-assignments]).

1.5.2. SCION AS Numbers

A SCION AS number is the 48-bit identifier for an AS. Although they play a similar role, there is no relationship between SCION AS numbers and BGP ASNs as defined by [RFC4271]. For historical reasons some SCION Autonomous Systems use an AS number where the first 16 bits are 0 and the remaining 32 bits are identical to their BGP ASN, but there is no technical requirement for this. AS numbers of public ASes MUST be globally unique.

1.5.2.1. Wildcard Addressing

SCION endpoints use wildcard AS 0 to designate any core AS, e.g., to place requests for core segments or down segments during path lookup. These wildcard addresses are of the form I-0, to designate any AS in ISD I. For more information, see Section 4.2.1.

1.5.2.2. SCION AS numbers

AS	Size	Description
0	1	The wildcard AS
1 - 4294967295	~4.3 billion	Public SCION AS numbers
1:0:0 - 1:ffff:ffff	~4.3 billion	Public SCION AS numbers - future allocations
2:0:0 - 2:ffff:ffff	~4.3 billion	Public SCION AS numbers
3:0:0 - feff:ffff:ffff	~280 trillion	Unallocated
ff00:0:0 - ff00:0:ffff	65536	Reserved for documentation and sample code (analogous to [RFC5398])
ff00:1:0 - ffa9:ffff:ffff	~730 billion	Unallocated
ffaa:0:0 - ffaa:ff:ffff	~16.8 million	Reserved for private use (analogous to [RFC6996]) - these numbers can be used for testing and private deployments
ffaa:100:0 -	~369 billion	Unallocated

ffff:ffff:fffe		
+-----+	+-----+	+-----+
ffff:ffff:ffff	1	Reserved
+-----+	+-----+	+-----+

Table 2: AS number allocations

1.5.3. Text Representation

1.5.3.1. ISD numbers

The text representation of SCION ISD numbers MUST be its decimal ASCII representation.

1.5.3.2. AS numbers

The text representation of SCION AS numbers is as follows:

- * SCION AS numbers in the lower 32-bit range MUST be printed as decimal by implementations. Implementations may parse AS numbers in the lower 32-bit range in hexadecimal notation too (e.g., a program may accept AS number '0:1:f' for AS 65551).
- * SCION AS numbers in the higher 32-bit range MUST be printed using big-endian hexadecimal notation in 3 groups of 4, in the range 1:0:0 to ffff:ffff:ffff. Leading zeros in each group are omitted, with the exception that one zero MUST be notated if a group is entirely zeros (e.g., 1:0:1). The :: zero-compression feature of IPv6 MUST NOT be used.
- * A range of AS numbers can be shortened with a notation similar to the one used for CIDR IP ranges ([RFC4632]). In such case, hexadecimal notation MUST be used. For example, the range of the lowest 32-bit AS numbers (0-4294967295) can be represented as 0:0:0/16.

1.5.3.3. <ISD, AS> tuples

The text representation of SCION addresses MUST be <ISD>-<AS>, where <ISD> is the text representation of the ISD number, <AS> is the text representation of the AS number, and - is the literal ASCII character 0x2D. This text representation is used for the ISD-AS number attribute in the certificates (see [I-D.dekater-scion-pki]).

For example, the text representation of AS number ff00:0:1 in ISD number 15 is 15-ff00:0:1.

1.6. Bootstrapping ability

SCION uses the following mechanisms to avoid circular dependencies during bootstrapping, and to provide resiliency after systemic failures:

- * Neighbor-based path discovery: Path discovery in SCION is performed by the beaconing mechanism. In order to participate in this process, an AS Control Service only needs to be aware of its direct neighbors. As long as no path segments are available, communicating with the neighboring ASes is possible with the one-hop path type which does not rely on any path information. SCION uses these `_one-hop paths_` to propagate PCBs to neighboring ASes to which no forwarding path is available yet. The One-Hop Path Type is described in more detail in [I-D.dekater-scion-dataplane].
- * Path reversal: In SCION, every path is reversible. That is, the receiver of a packet can reverse the path in the packet header in order to produce a reply packet without having to perform a path lookup. Such a packet follows the original packet's path in the reverse direction.
- * Availability of certificates: Every entity is required to be in possession of all cryptographic material including the ISD's TRC and AS certificates, in order to verify any message it sends. This together with the path reversal means that the receiver of a message can always obtain all this material by contacting the sender.

Note: For a detailed description of a TRC and more information on the availability of certificates and TRCs, see [I-D.dekater-scion-pki].

1.7. Communication Protocol

All communication between the Control Services in different SCION ASes is expressed in terms of RPC remote procedure calls. Service interfaces and messages are defined in the Protocol Buffer "proto3" interface definition language (for details, see [proto3]).

The RPC messages are transported via [Connect]'s RPC protocol that carries messages over HTTP/3 (see [RFC9114]), which in turn uses QUIC/UDP ([RFC9000]) as a transport layer. Connect is backward compatible with [gRPC] which is supported but deprecated.

In case of failure, RPC calls return an error as specified by the RPC framework. That is, a non-zero status code and an explanatory string. Section 5 provides details about the establishment of the underlying QUIC connections.

SCION's Control Plane does not require any domain name resolution for communication.

2. Path Exploration or Beaconsing

2.1. Introduction and Overview

Path Exploration is the process where a SCION AS discovers paths to other ASes. In SCION, this process is referred to as *_beaconsing_*.

The *_Control Service_* of each SCION AS is responsible for the beaconsing process. The Control Service generates, receives, and propagates *_Path-Segment Construction Beacons (PCBs)_* on a regular basis, to iteratively construct path segments.

PCBs contain inter-domain topology and authentication information, and can include additional metadata that helps with path management and selection. The beaconsing process itself is divided into routing processes on two levels, where *_core_* or inter-ISD is based on the (selective) sending of PCBs without a defined direction, and *_intra-ISD_* beaconsing on top-to-bottom propagation. Beaconsing is initiated by core ASes, therefore each ISD MUST have at least one core AS.

- * *_Core or Inter-ISD beaconsing_* is the process of constructing path segments between core ASes in the same or in different ISDs. During core beaconsing, the Control Service of a core AS either originates PCBs or propagates PCBs received from neighboring core ASes to other neighboring core ASes.
- * *_Intra-ISD beaconsing_* creates path segments from core ASes to non-core ASes. For this, the Control Services of core ASes originates PCBs and sends them to the non-core child ASes (typically customer ASes) at regular intervals. The Control Service of a non-core child AS receives these PCBs and forwards them to its child ASes, and so on until the PCB reaches an AS without any children. As a result, all ASes within an ISD receive path segments to reach the core ASes of their ISD and register reciprocal segments with the Control Service of the associated core ASes.

On its way, 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 is added to the PCB. The full PCB message format is described in

Section 2.2. PCBs are used to construct path segments. ASes register them to make them available to other ASes, as described in Section 3.

2.1.1. Peering Links

PCBs do not traverse peering links, but peering links are instead announced along with a regular path in a PCB. If both ASes at either end of a peering link have registered path segments that include this specific peering link, then it is possible to use this during segment combination to create the end-to-end path.

2.1.2. Appending Entries to a PCB

Every propagation interval (as configured by the AS operator), the Control Service:

- * selects the best combinations of PCBs and interfaces connecting to a neighboring AS (i.e., a child AS or a core AS). This is described in Section 2.3.3.
- * propagates each selected PCB to the selected egress interface(s) associated with it. Selection policies are described in Section 2.3.5.

For every selected PCB and egress interface combination, the AS Control Service appends an `_AS entry_` to the selected PCB. This includes a Hop Field that specifies the ingress and egress interface for the packet forwarding through this AS, in the beaconing direction. The AS entry can also contain peer entries.

2.1.3. PCB Propagation - Illustrated Examples

The following three figures show how intra-ISD PCB propagation works, from the ISD's core AS down to child ASes. Interface identifiers of each AS are numbered with integer values while ASes are described with an upper case letter for the sake of illustration. Arrows represent the PCB propagation direction.




Figure 2: Intra-ISD PCB propagation from the ISD core to child ASes - Part 1

In Figure 2, core AS X sends the two different PCBs "a" and "b" via two different links to child AS Y: PCB "a" exits core AS X via egress interface "2", whereas PCB "b" is sent over egress interface "1". Core AS X adds the respective egress information to the PCBs when sending them off, as can be seen in the figure (the entries "_Core - Out:2_" and "_Core - Out:1_", respectively).




Figure 3: Intra-ISD PCB propagation from the ISD core to child ASes - Part 2

In Figure 3, AS Y receives the two PCBs "a" and "b" through two different (ingress) interfaces, namely "2" and "3", respectively. Additionally, AS Y forwards to AS Z four PCBs that were previously sent by core AS X. For this, AS Y uses the two different (egress) links "5" and "6". AS Y appends the corresponding ingress and egress interface information to the four PCBs.

AS Y also has two peering links to its neighboring peers V and W, through the interfaces "1" and "4" respectively, which is included in the information in the PCBs. Thus, each forwarded PCB accumulates path information on its way "down" from core AS X.

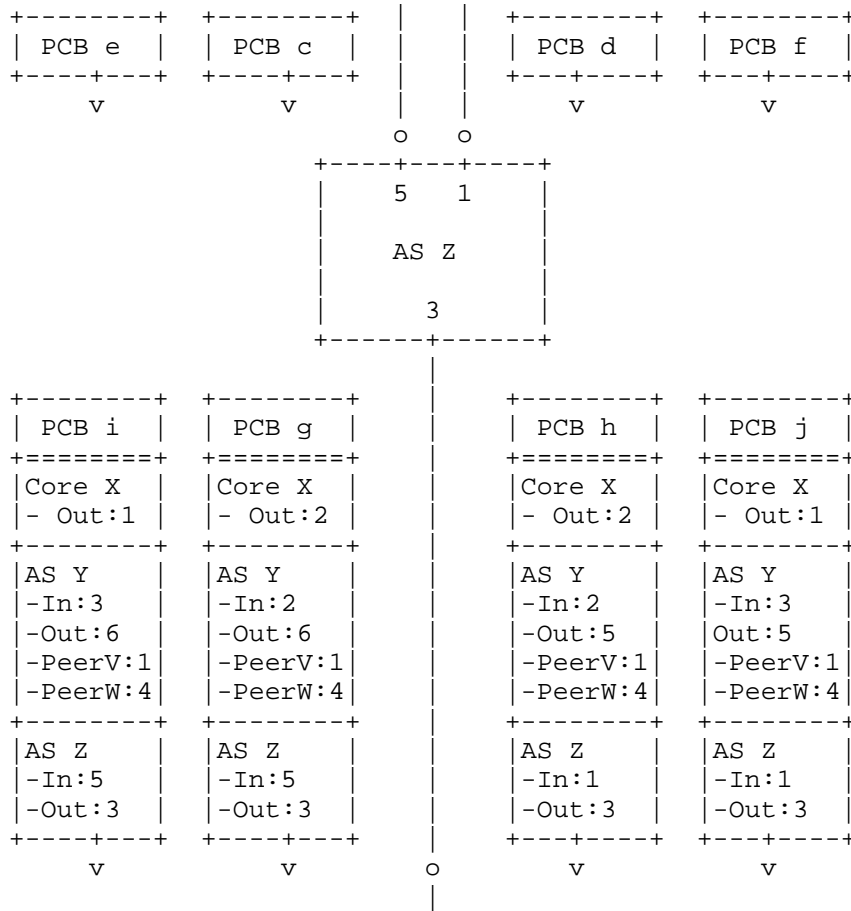


Figure 4: Intra-ISD PCB propagation from the ISD core to child ASes - Part 3

In Figure 4, the four PCBs "c", "d", "e", and "f" are coming from AS Y and are received by AS Z over two different links: PCBs "c" and "e" reach AS Z over ingress interface "5", whereas PCBs "d" and "f" enter AS Z via ingress interface "1". Additionally, AS Z propagates PCBs "g", "h", "i", and "j" further downwards over the same link (egress interface "3"), and extends the PCBs with the relevant information so that each of these includes AS hop entries from core AS X, AS Y, and AS Z.

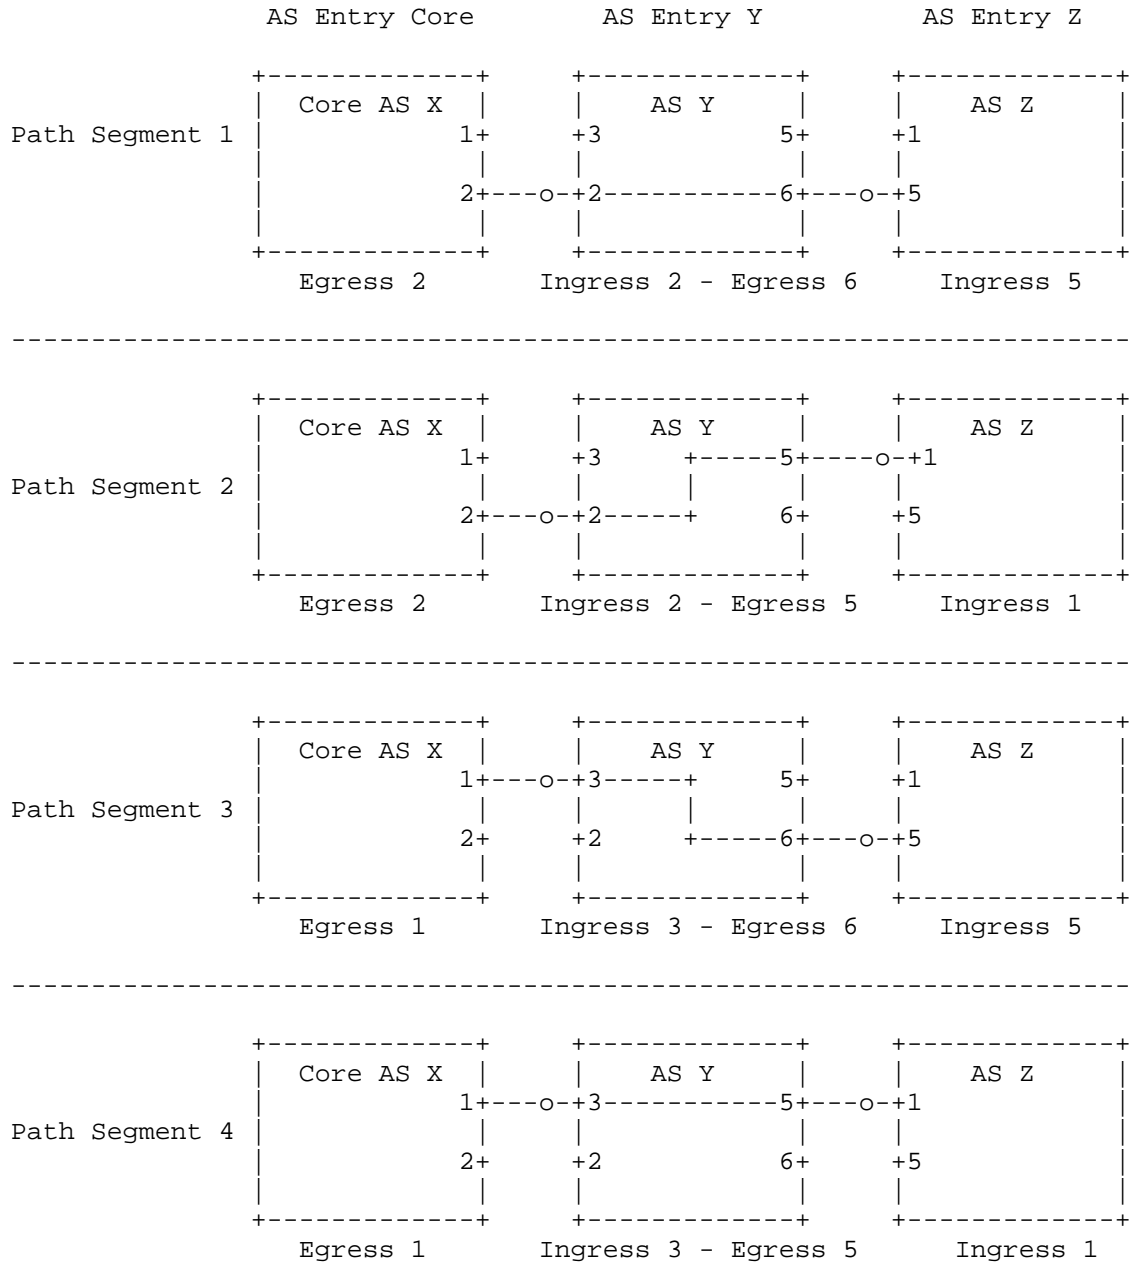


Figure 5: Possible up- or down segments for AS Z

According to Figure 2, Figure 3 and Figure 4 above, it appears that a PCB represents a single path segment. However, there is a difference between a PCB and a registered path segment as a PCB is a so-called "travelling path segment" that accumulates AS entries when traversing SCION networks. A registered path segment is instead a "snapshot" of a travelling PCB at a given time T and from the vantage point of a particular AS A. This is illustrated by Figure 5 which shows several possible path segments to reach AS Z, based on the PCBs "g", "h", "i", and "j" from Figure 4 above. It is up to AS Z to choose which of these path segments to use.

2.2. PCB Message Format

Each PCB is comprised of a message containing the following top-level fields:

```
+-----+-----+-----+-----+
|Segment Info | AS Entry 0 | AS Entry 1 | ... | AS Entry N |
+-----+-----+-----+-----+
```

Figure 6: PCB Top-Level Message Format

Each PCB MUST consist of at least:

- * An information field with an identifier and a timestamp.
- * Entries of all ASes on the path segment represented by this PCB.

The PCB top level Protobuf message format is:

```
message PathSegment {
    bytes segment_info = 1;
    repeated ASEntry as_entries = 2;
}
```

- * `segment_info`: This field is used as input for the PCB signature. It is the encoded version of the `SegmentInformation` component which provides basic information about the PCB. This component is specified in detail in Section 2.2.1.
- * `as_entries`: Contains the `ASEntry` component of all ASes on the path segment represented by this PCB. The order of the AS entries MUST correspond to the path traversal order in the PCB propagation direction.

- * ASEntry: The ASEntry component contains the complete path information of a specific AS that is part of the path segment represented by the PCB. This component is specified in detail in Section 2.2.2.

The information to be included in each of these fields is described below.

2.2.1. Segment Information




Figure 7: Segment Information Component

Each PCB MUST include a SegmentInformation message with basic information about the PCB. Its Protobuf message format is:

```

message SegmentInformation {
    int64 timestamp = 1;
    uint32 segment_id = 2;
}

```

- * timestamp: The timestamp indicates the creation time of this PCB. It is set by the originating core AS and the expiration time of each Hop Field in the PCB is computed relative to this timestamp. The timestamp is encoded as the number of seconds elapsed since the POSIX Epoch (1970-01-01 00:00:00 UTC). Note that this timestamp is encoded as a 32-bit unsigned integer in [I-D.dekater-scion-dataplane].
- * segment_id: The 16-bit identifier of this PCB and the corresponding path segment. The segment ID is REQUIRED for the computation of the message authentication code (MAC) of an AS's Hop Field. The MAC is used for Hop Field verification in the data plane and the originating core AS MUST fill this field with a cryptographically random number.

Note: See Section 2.2.2.5 for more information on the Hop Field message format. [I-D.dekater-scion-dataplane] provides a detailed description of the computation of the MAC and the verification of the Hop Field in the data plane.

2.2.2. AS Entry




Figure 8: AS Entry

Each PCB MUST also contain the entries of all ASes included in the corresponding path segment. This means that the originating core AS MUST add its AS entry to each PCB it creates, and each traversed AS MUST attach its AS entry to the PCB.

One AS entry contains the complete hop information for this specific AS in this specific path segment. It consists of a signed and an unsigned component.

The code block below defines an AS entry `ASEntry` in Protobuf message format.

```

message ASEntry {
    SignedMessage signed = 1;
    PathSegmentUnsignedExtensions unsigned = 2;
}

```

It includes the following components:

- * `SignedMessage`: The AS entry signed component.
- * `PathSegmentUnsignedExtensions`: Optional unsigned PCB extensions, further described in Section 2.2.3.




Figure 9: AS Entry Signed Component

Each AS entry of a PCB MUST include a signed component as well as a signature computed over the signed component. Each AS entry MUST be signed with the Control Plane AS Certificate (see [I-D.dekater-scion-pki]).

The signed component of an AS entry MUST include the following elements:

- * a header,
- * a body, and
- * a signature.

In the Protobuf message format, the signed component of an AS entry is specified by the SignedMessage. It consists of a header-and-body part (header_and_body) and a raw signature (signature). Their protobuf message formats are:

```
message SignedMessage {  
    bytes header_and_body = 1;  
    bytes signature = 2;  
}
```

Protobuf definition of the HeaderAndBody message used for signature computation input.

```
message HeaderAndBody {  
    bytes header = 1;  
    bytes body = 2;  
}
```

The Signed Header, Signed Body, and Signature Fields are detailed below.

2.2.2.1. AS Entry Signed Header

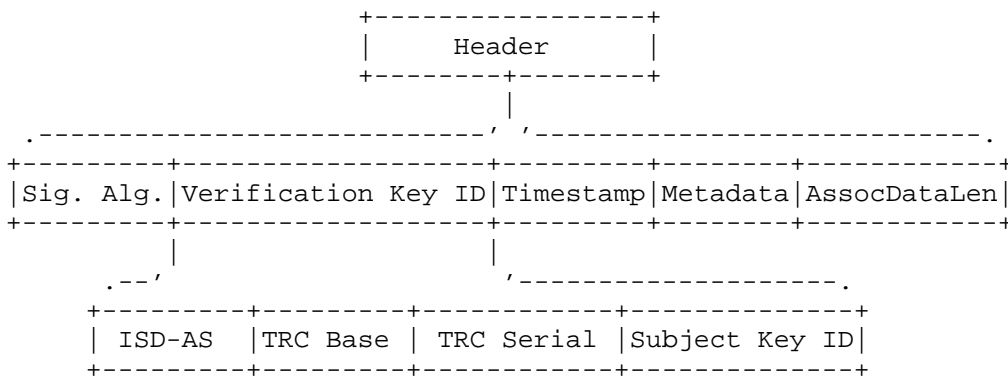


Figure 10: AS Entry Signed Header

The header part carries information that is relevant to the computation and verification of the signature. It contains the following fields:

- * `signature_algorithm`: Specifies the algorithm to compute the signature. Possible types are defined by the `SignatureAlgorithm` definition and are further discussed in [I-D.dekater-scion-pki], but an unspecified signature algorithm is never valid. Other signature algorithms or curves MAY be used in the future. This field is REQUIRED.
- * `verification_key_id`: Contains a `VerificationKeyID` message, carrying information relevant to signing and verifying PCBs and other control-plane messages. This field is REQUIRED.
- * `timestamp`: Defines the signature creation timestamp. This field is OPTIONAL.
- * `metadata`: May include metadata. While it is part of the generic Header message format, it MUST be empty in an AS entry signed header. This field is OPTIONAL.
- * `associated_data_length`: Specifies the length of the data covered by the signature but not included within the header or body. This data contains information about preceding AS entries, as described in Section 2.2.2.6. The value of this field is zero if no associated data is covered by the signature.

The Header and `SignatureAlgorithm` protobuf message formats are:

```
message Header {
    SignatureAlgorithm signature_algorithm = 1;
    bytes verification_key_id = 2;
    google.protobuf.Timestamp timestamp = 3;
    bytes metadata = 4;
    int32 associated_data_length = 5;
}

enum SignatureAlgorithm {
    SIGNATURE_ALGORITHM_UNSPECIFIED = 0;
    SIGNATURE_ALGORITHM_ECDSA_WITH_SHA256 = 1;
    SIGNATURE_ALGORITHM_ECDSA_WITH_SHA384 = 2;
    SIGNATURE_ALGORITHM_ECDSA_WITH_SHA512 = 3;
}
```

The VerificationKeyID message contains the following REQUIRED fields:

- * `isd_as`: The ISD-AS number of the current AS.
- * `subject_key_id`: Refers to the certificate that contains the public key needed to verify this PCB's signature.
- * `trc_base`: Defines the `_base_` number of the latest Trust Root Configuration (TRC) available to the signer at the time of the signature creation.
- * `trc_serial`: Defines the `_serial_` number of the latest TRC available to the signer at the time of the signature creation.

Its protobuf message format is:

```
message VerificationKeyID {
    uint64 isd_as = 1;
    bytes subject_key_id = 2;
    uint64 trc_base = 3;
    uint64 trc_serial = 4;
}
```

For more information on signing and verifying control plane messages (such as PCBs), see 'Signing and Verifying Control Plane Messages' in [I-D.dekater-scion-pki]. For more information on the TRC base and serial number, see 'Trust Root Configuration Specification' in [I-D.dekater-scion-pki].

2.2.2.2. AS Entry Signed Body




Figure 11: AS Entry Signed Body

The body of an AS entry MUST consist of the signed component `ASEntrySignedBody` of all the ASes in the path segment represented by the PCB, up until and including the current AS. Its Protobuf message format is:

```

message ASEntrySignedBody {
    uint64 isd_as = 1;
    uint64 next_isd_as = 2;
    HopEntry hop_entry = 3;
    repeated PeerEntry peer_entries = 4;
    uint32 mtu = 5;
    PathSegmentExtensions extensions = 6;
}

```

- * `isd_as`: The ISD-AS number of the AS that created this AS entry.
- * `next_isd_as`: The ISD-AS number of the downstream AS to which the PCB MUST be forwarded. The presence of this field prevents path hijacking attacks, as further discussed in Section 8.3.1.
- * `hop_entry`: The hop entry (`HopEntry`) with the information required by the data plane to forward this PCB through the current AS to the next AS. For the specification of the hop entry, see Section 2.2.2.3.
- * `peer_entries`: The list of optional peer entries (`PeerEntry`). For a specification of one peer entry, see Section 2.2.2.4.
- * `mtu`: The maximum transmission unit (MTU) in bytes that is supported by all intra-domain links within the current AS. This value is set by the Control Service when adding the AS entry to the beacon. How the Control Service obtains this information is implementation dependent, but current practice is to make it a configuration item.

- * extensions: List of signed extensions (optional). PCB extensions defined here are part of the signed AS entry. This field SHOULD therefore only contain extensions that include important metadata for which cryptographic protection is required. For more information on PCB extensions, see Section 2.2.3.

2.2.2.3. Hop Entry




Figure 12: Hop Entry

Each body of an AS entry MUST contain exactly one hop entry. It specifies forwarding information which the data plane requires to create the hop through the current AS (in the direction of the beaconing).

The HopEntry Protobuf message format is:

```

message HopEntry {
    HopField hop_field = 1;
    uint32 ingress_mtu = 2;
}

```

- * hop_field: Contains the authenticated information about the ingress and egress interfaces in the direction of beaconing. Routers need this information to forward packets through the current AS. For further specifications, see Section 2.2.2.5.
- * ingress_mtu: Specifies the maximum transmission unit (MTU) of the ingress interface (in beaconing direction) of the hop being described. The MTU of paths constructed from the containing beacon is necessarily less than or equal to this value. How the Control Service obtains the MTU of an inter-AS link is implementation dependent. It may be discovered or configured by operators, but current practice to make it a configuration item. Path MTU is further discussed in Section 3.4.

In this description, MTU and packet size are to be understood in the same sense as in [RFC1122]. That is, exclusive of any layer 2 framing or packet encapsulation (for links using an underlay network).

2.2.2.4. Peer Entry




Figure 13: Peer Entry

By means of a peer entry, an AS can announce that it has a peering link to another AS. A peer entry is an optional component of a PCB, and is only included if there is a peering link to a peer AS.

The PeerEntry Protobuf message format is:

```
message PeerEntry {
    uint64 peer_isd_as = 1;
    uint64 peer_interface = 2;
    uint32 peer_mtu = 3;
    HopField hop_field = 4;
}
```

- * `peer_isd_as`: The ISD-AS number of the peer AS. This number is used to match peering segments during path construction.
- * `peer_interface`: The 16-bit interface identifier of the peering link on the peer AS side. This identifier is used to match peering segments during path construction.
- * `peer_mtu`: Specifies the maximum transmission unit (MTU) of the peering link being described. The MTU of paths via such link is necessarily less than or equal to this value. How the Control Service obtains the MTU of an inter-AS link is implementation dependent. It may be discovered or configured, but current practice is to make it a configuration item.

- * `hop_field`: Contains authenticated information about the ingress and egress interfaces in the current AS (coming from the peering link, in the direction of beaconing - see also Figure 6). The data plane needs this information to forward packets through the current AS. For further specifications, see Section 2.2.2.5.

In this description, MTU and packet size are to be understood in the same sense as in [RFC1122]. That is, exclusive of any layer 2 framing or packet encapsulation (for links using an underlay network).

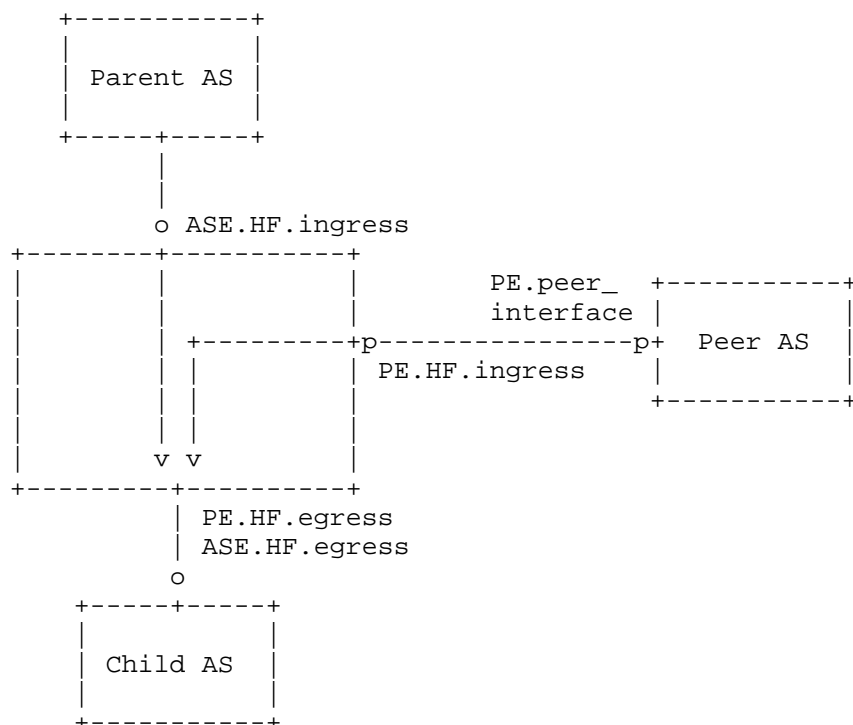


Figure 14: Peer entry information, in the direction of beaconing.
PE denotes a peer entry, ASE an AS entry, HF an hop field.

2.2.2.5. Hop Field




Figure 15: Hop Field

The Hop Field, part of both hop and peer entries, is used directly in the data plane for packet forwarding and specifies the incoming and outgoing interfaces of the ASes on the forwarding path. This information is authenticated with a Message Authentication Code (MAC) which is used by the Control Service of an AS to authenticate path segments with its border routers during packet forwarding.

The algorithm used to compute the Hop Field MAC is an AS-specific choice, although the Control Services and border routers within an AS MUST use the same algorithm. Implementations MUST also support the Default Hop Field MAC algorithm. See [I-D.dekater-scion-dataplane] section "Authorizing Segments through Chained MACs" for more information including configuration. Endpoints do not compute MACs.

The HopField Protobuf message format is:

```

message HopField {
    uint64 ingress = 1;
    uint64 egress = 2;
    uint32 exp_time = 3;
    bytes mac = 4;
}
  
```

* ingress: The 16-bit ingress interface identifier (in the direction of the path construction. That is, in the direction of beaconing through the current AS).

Note: The core AS originating a PCB MUST set the ingress interface identifier to the "unspecified" value (see [I-D.dekater-scion-dataplane] section "Terminology").

* egress: The 16-bit egress interface identifier (in the direction of beaconing).

* exp_time: The 8-bit encoded expiration time of the Hop Field, indicating its validity. This field expresses a duration in seconds according to the formula: $\text{duration} = (1 + \text{exp_time}) *$

$(24*60*60/256)$ and the minimum duration is therefore 337.5 seconds. This duration is relative to the PCB creation timestamp set in the PCB's segment information component (see also Section 2.2.1), so the absolute expiration time of the Hop Field is the sum of these two values.

- * `mac`: The message authentication code (MAC) used in the data plane to verify the Hop Field, as described in [I-D.dekater-scion-dataplane].

2.2.2.6. AS Entry Signature

Each AS entry MUST be signed with the AS certificate's private key `Ki`. The certificate MUST have a validity period that is longer than the Hop Field absolute expiration time (described in Section 2.2.2.5). The signature `Sigi` of an AS entry `ASEi` is computed over the AS entry's signed component.

This is the input for the computation of the signature:

- * The signed header and body of the current AS (`header_and_body`).
- * The `segment_info` component of the current AS. This is the encoded version of the `SegmentInformation` component containing basic information about the path segment represented by the PCB. For the specification of `SegmentInformation`, see Section 2.2.1.
- * The signed `header_and_body/signature` combination of each previous AS on this specific path segment.

The signature `Sigi` of an AS entry `ASEi` is now computed as follows. Symbol `||` represents concatenation.

$$\text{Sig}_i = K_i(\text{ASE}_i^{\text{(signed)}} || \text{SegInfo} || \text{ASE}_0^{\text{(signed)}} || \text{Sig}_0 || \dots || \text{ASE}_{(i-1)}^{\text{(signed)}} || \text{Sig}_{(i-1)})$$

The signature metadata minimally contains the ISD-AS number of the signing entity and the key identifier of the public key that other ASes can use to verify the message. For more information on signing and verifying control plane messages, see 'Signing and Verifying Control Plane Messages' in [I-D.dekater-scion-pki].

Some of the data used as an input to the signature comes from previous ASes in the path segment. This data is therefore called "associated data" and it gives the signature a nested structure. The content of associated data defined in Protobuf is:

```

input(ps, i) = signed.header_and_body || associated_data(ps, i)

associated_data(ps, i) = ps.segment_info ||
    ps.as_entries[1].signed.header_and_body ||
    ps.as_entries[1].signed.signature ||
    ...
    ps.as_entries[i-1].signed.header_and_body ||
    ps.as_entries[i-1].signed.signature

```

2.2.3. PCB Extensions

AS entries in PCBs may carry a number of optional extensions that accumulate information while traveling across ASes. Extensions can be:

- * Unsigned extensions PathSegmentUnsignedExtensions. They are part of the AS entry component (the AEntry message, see also Section 2.2.2).
- * Signed extensions PathSegmentExtensions. They are part of the signed body component of an AS entry (the AEntrySignedBody message, see also Section 2.2.2.2).

The Protobuf message format of extensions is below. As an example, it mentions the StaticInfoExtension, a signed extension that is used to carry path segment metadata, such as segment latency, bandwidth, router coordinates, link type, number of internal hops. This and other extensions are at time of writing experimental, so definitions of this message format are omitted and [PCBExtensions] should be referred to.

```

message PathSegmentUnsignedExtensions {
}

message PathSegmentExtensions {
    StaticInfoExtension static_info = 1;
}

```

If a Control Service receives an unknown PCB extension, it SHOULD skip the extension, but preserve it unmodified in case the PCB is further propagated.

2.2.4. PCB Validity

To be valid (that is, usable to construct a valid path), a PCB MUST:

- * Contain valid AS Entry signatures (Section 2.2.2.6).

- * Have a timestamp (Section 2.2.1) that is not later than the current time at the point of validation, plus an allowance for differences between the clocks of the validator and originator.
- * Contain only unexpired hop fields (Section 2.2.2.5).

It is recommend to use the hopfield expiration time (that is 337.5 seconds, see Section 2.2.2.5) as the allowance for differences between the clocks of the validator and originator.

For the purpose of validation, a hop is considered expired if its absolute expiration time (as defined in Section 2.2.2.5), is later than the current time + allowance at the point of validation.

2.2.5. Configuration

For the purpose of constructing and propagating path segments, a network operator needs to configure an AS Control Service with links to neighboring ASes. Such information may be conveyed to the Control Service in an out-of-band fashion (e.g., in a configuration file). For each link, these values MUST be configured:

- * Local Interface ID. This MUST be unique within each AS.
- * Neighbor type (core, parent, child, peer), depending on link type (see Section 1.3). Link type depends on mutual agreements between operators of the ASes at each end of each link.
- * Neighbor ISD-AS number.
- * Neighbor interface underlay address.
- * For peering links, neighbor Interface ID.

In addition, a network operator needs to configure an AS Control Service with:

- * the algorithm and forwarding key used to compute the Hop Field MAC, which are also used by routers within the AS. These are further described in [I-D.dekater-scion-dataplane].
- * registration interval (see Section 3.1).
- * propagation interval and best PCBs set size (see Section 2.3.4).
- * the maximum MTU supported by all intra-AS links may also be configured by the operator.

- * Hop Field expiration time (see Section 2.2.2.5). Current implementations default to 63, corresponding to 6 hours.

Optionally, it may configure per-link MTU (see Section 2.2.2.3) and PCB selection policies (see Section 2.3.3).

2.3. Propagation of PCBs

This section describes how PCBs are received, selected and further propagated in the path exploration process.

2.3.1. Reception of PCBs

Upon receiving a PCB, the Control Service of an AS performs the following checks:

1. PCB validity: The Control Service MUST check the validity of the PCB (see Section 2.2.4) and it MUST discard invalid PCBs. The PCB contains the version numbers of the TRC(s) and certificate(s) that MUST be used to verify its signatures which enables the Control Service to check whether it has the relevant TRC(s) and certificate(s). If not, they can be requested from the Control Service of the sending AS through the API described in Section 2.4.
2. Loop avoidance: The Control Service MUST check whether the PCB is from a core AS and whether it includes duplicate hop entries created by the core AS itself or by other ASes. If so, it MUST discard the PCB in order to avoid loops. This is necessary because core beaconing is based on propagating PCBs to all AS neighbors. Additionally, core ASes SHOULD discard PCBs that were propagated at any point by a non-core AS. Ultimately, core ASes MAY make a policy decision to propagate beacons containing path segments that traverse the same ISD more than once as this can be legitimate - e.g., if the ISD spans a large geographical area, a path between different ASes transiting another ISD may constitute a shortcut.
3. Incoming Interface: The Control Service MUST check that the last ISD-AS entry in a received PCB (in its AS Entry Signed Body) corresponds with the ISD-AS neighbor of the interface where the PCB was received. In addition, the corresponding link MUST be core or parent, otherwise the PCB MUST be discarded.
4. Continuity: When the Control Service receives a PCB containing two or more AS entries, it MUST check every AS entry except the last and discard beacons where the next_isd_as of an entry does not equal the isd_as of the next entry.

If the PCB verification is successful, the Control Service decides whether to store the PCB as a candidate for propagation based on selection criteria and policies specific for each AS.

2.3.2. Storing Candidate PCBs

An AS Control Service stores candidate PCBs in a temporary storage called the `_Beacon Store_`. The management of this storage is implementation specific, but current practice is to retain all PCBs until expired or replaced by one describing the same path with a later origination time.

2.3.3. PCB Selection Policies

The Control Service of an AS MUST select which PCBs to propagate further. The selection process can inspect and compare the properties of the candidate PCBs (e.g., length, disjointness across different paths, age, expiration time) and/or take into account which PCBs have been propagated in the past. The PCBs to select or eliminate is determined by the policy of the AS.

In order to avoid excessive overhead on the path discovery system in bigger networks, an AS Control Service should only propagate those candidate PCBs with the highest probability of meeting the needs of the endpoints that will perform path construction, in accordance with Section 2.3.4.

As SCION does not provide any in-band signal about the intentions of endpoints nor about the policies of downstream ASes, the policy will typically select a somewhat diverse set optimized for multiple, generic parameters. Selection may be based on criteria such as:

- * AS path length: from the originator core AS to the child (non-core) AS.
- * Expiration time: the maximum value for the expiration time when extending the segment.
- * ISD or AS exclusion lists: certain ASes or ISD that may not appear in a segment.
- * ISD loops: if permitted, they allow core AS to reach other core ASes in the same ISD via a third party ISDs.
- * Availability of peering links: that is the number of different peering ASes from all non-core ASes on the PCB or path segment. A greater number of peering ASes increases the likelihood of finding a shortcut on the path segment.

- * Path disjointness: Paths can be either AS disjointed or link disjointed. AS disjointed paths have no common upstream/core AS for the current AS, whereas link disjointed paths do not share any AS-to-AS link. AS disjointness allows path diversity in the event that an AS becomes unresponsive, and link disjointness provides resilience in case of link failure. Both criteria can be used depending on the objective of the AS.

The relative disjointness of two PCBs A and B may be calculated by assigning a disjointness score, calculated as the number of links in A that don't appear in B. For example, the beacon that has the highest disjointness score and is not the shortest path should be selected, but if this not better than what has already been selected, then the beacon with the shortest path yet to be selected should be chosen instead.

A PCB Selection Policy can be expressed as a stateful filter of segments, i.e., a function which indicates whether to accept or deny a given segment. This filter is stateful in that it can be updated each time its AS registers a new segment. It is expected that an AS's policy will select PCBs corresponding to paths that are commercially or otherwise operationally viable.

To ensure reachability, PCB selection policies should forward as many PCBs as possible as PCB selection is not intended as a mechanism for traffic engineering (e.g., by excluding specific PCBs for that purpose).

2.3.4. Propagation Interval and Best PCBs Set Size

PCBs are propagated in batches to each neighboring AS at a fixed frequency known as the `_propagation interval_` which happens for both intra-ISD beaconing (Section 7.4.1) and core beaconing (Section 7.4.2). At each propagation event, the AS Control Service selects a set of the best PCBs from the candidates in the Beacon Store according to the AS's selection policy.

The size of this set is called the `_best PCBs set size_`. It should be:

- * For intra-ISD beaconing (i.e., propagating to children ASes): at most 50.
- * For core beaconing (i.e., propagation between core ASes): at most 5 per immediate neighbor core AS. Current practice is for the set to contain an equal amount of PCBs from each neighbor.

These values reflect a tradeoff between scalability — limited by the computational overhead of signature verification — and the number of paths discovered. The PCBs set size should not be too low to ensure that beaconing can discover a significant number of paths. Further discussion on these trade-offs is provided in Section 7.4.

In current practice the intra-ISD set size is typically 20. Current practice also showed that in small SCION core networks, higher values of the core best PCBs set size (e.g., 20) can be used.

Depending on the selection criteria, it may be necessary to keep more candidate PCBs than the `_best PCBs set size_` in the Beacon Store in order to determine the best set of PCBs. If this is the case, an AS Control Service should have a suitable pre-selection of candidate PCBs in place in order to keep the Beacon Store capacity limited.

- * The `_propagation interval_` should be at least 5 seconds for intra-ISD beaconing and at least 60 seconds for core beaconing.

Note that to quickly (re-)establish connectivity, an AS Control Service MAY attempt to forward a PCB more frequently ("fast recovery"). Current practice is to increase the frequency of attempts if no PCB propagation is known to have succeeded within the last propagation interval:

- * because the corresponding RPC failed;
- * or because no beacon was available to propagate.

2.3.5. Propagation of Selected PCBs

The propagation process includes the following steps:

1. From the candidate PCBs stored in the Beacon Store, the Control Service of an AS selects the best PCBs to propagate to its neighboring ASes, according to the AS selection policy (see Section 2.3.3). Core ASes propagate PCBs over both core and parent-child links; additionally, they originate new PCBs over these same links. Non-core ASes propagate PCBs exclusively over parent-child links.
2. The Control Service MUST add a new AS entry (see Section 2.2.2) including any Peer Entry information (see Section 2.2.2.4) the AS is configured to advertise to every selected PCB.
3. The Control Service MUST sign each selected, extended PCB and append the computed signature.

4. As a final step, the Control Service propagates each extended PCB to the neighboring AS specified in the new AS entry by invoking the `SegmentCreationService.Beacon` remote procedure call (RPC) in the Control Services of the neighboring ASes (see also Section 2.3.5.1).

To bootstrap initial communication with a neighboring beacon service, ASes use one-hop paths. This special kind of path handles beaconing between neighboring ASes for which no forwarding path may be available yet. It is the task of beaconing to discover such forwarding paths and the purpose of one-hop paths is to break this circular dependency. The One-Hop Path Type is described in more detail in [I-D.dekater-scion-dataplane].

2.3.5.1. Propagation of PCBs in Protobuf Message Format

The last step of the above described core and intra-ISD propagation procedures is implemented as follows in Protobuf message format:

```
service SegmentCreationService {  
  rpc Beacon(BeaconRequest) returns (BeaconResponse) {}  
}  
  
message BeaconRequest {  
  PathSegment segment = 1;  
}  
  
message BeaconResponse {}
```

The propagation procedure includes the following elements:

- * `SegmentCreationService`: Specifies the service via which the extended PCB is propagated to the Control Service of the neighboring AS.
 - `Beacon`: Specifies the method that calls the Control Service at the neighboring AS in order to propagate the extended PCB.
- * `BeaconRequest`: Specifies the request message sent by the Beacon method to the Control Service of the neighboring AS. It contains the following element:
 - `PathSegment`: Specifies the path segment to propagate to the neighboring AS. For more information on the Protobuf message type `PathSegment`, see Section 2.2.
- * `BeaconResponse`: An empty message returned as an acknowledgement upon success.

2.4. Distribution of Cryptographic Material

Control Services distribute cryptographic material for the PKI (see [I-D.dekater-scion-pki]) using the following protobuf messages through the TrustMaterialService RPCs:

```
service TrustMaterialService {  
  rpc Chains(ChainsRequest) returns (ChainsResponse) {}  
  rpc TRC(TRCRequest) returns (TRCResponse) {}  
}
```

- * Chains(ChainsRequest): Returns the certificate chains that match the request.
- * TRC(TRCRequest): Returns a specific TRC that matches the request.

The corresponding protobuf message formats are:

```
message ChainsRequest {  
  uint64 isd_as = 1;  
  bytes subject_key_id = 2;  
  google.protobuf.Timestamp at_least_valid_until = 3;  
  google.protobuf.Timestamp at_least_valid_since = 4;  
}
```

```
message ChainsResponse {  
  repeated Chain chains = 1;  
}
```

```
message Chain {  
  bytes as_cert = 1;  
  bytes ca_cert = 2;  
}
```

A ChainsRequest message includes the following fields:

- * isd_as: Returns ISD-AS of Subject in the AS certificate.
- * subject_key_id: Returns SubjectKeyID in the AS certificate.
- * at_least_valid_until: Point in time at which the AS certificate must still be valid - in seconds since Epoch according to [POSIX.1-2024] Section 4.19, encoded as a 32-bit unsigned integer.
- * at_least_valid_since: Point in time at which the AS certificate must be or must have been valid - in seconds since Epoch according to [POSIX.1-2024] Section 4.19, encoded as a 32-bit unsigned integer.

A ChainsResponse includes the following fields:

- * chains: Lists the certificate chains that match the request. A Chain contains:
 - as_cert: Returns the AS certificate in the chain.
 - ca_cert: Returns the CA certificate in the chain.

Certificates are ASN.1 DER encoded.

For requesting TRCs, the protobuf messages are:

```
message TRCRequest {
  uint32 isd = 1;
  uint64 base = 2;
  uint64 serial = 3;
}
```

```
message TRCResponse {
  bytes trc = 1;
}
```

A TRCRequest includes the following fields:

- * isd: Returns the ISD number of the TRC.
- * base: Returns the base number of the TRC.
- * serial: Returns the serial number of the TRC.

The returned trc contains the ASN.1 DER encoded CMS SignedData structure containing the TRC.

2.5. Renewal of Cryptographic Material

To renew PKI cryptographic material (see [I-D.dekater-scion-pki]), Control Services MAY employ out-of-band mechanisms or utilize the ChainRenewalService RPC to exchange the Protobuf messages defined below.

```
service ChainRenewalService {
  rpc ChainRenewal(ChainRenewalRequest) returns (
    ChainRenewalResponse) {}
}
```

- * ChainRenewal(ChainRenewalRequest): returns a renewed certificate chain.

The corresponding protobuf message format for requests is:

```
message ChainRenewalRequest {  
    reserved 1;  
    bytes cms_signed_request = 2;  
}
```

A ChainRenewalRequest message includes the following fields:

- * Field number 1 is legacy and it is not in use.
- * cms_signed_request: it contains the ASN.1 DER encoded CMS SignedData structure that contains an ASN.1 DER encoded PKCS #10 request.

The protobuf message format for responses is:

```
message ChainRenewalResponse {  
    reserved 1;  
    bytes cms_signed_response = 2;  
}
```

A ChainRenewalResponse message includes the following fields:

- * Field number 1 is legacy and it is not in use.
- * cms_signed_response: it contains an ASN.1 DER encoded CMS SignedData structure containing a two-certificate chain. The chain comprises the AS certificate followed by the CA certificate, both encoded in ASN.1 DER.

3. Registration of Path Segments

Path registration is the process where a SCION AS transforms selected PCBs into path segments, and adding these segments to the relevant path databases thereby making them available to other ASes.

A non-core AS typically receives several PCBs representing several path segments to the core ASes of the ISD the AS belongs to. Out of these PCBs, the non-core AS selects those down path segments through which it wants to be reached, based on AS-specific selection criteria.

The next step is to register the selected down segments with the Control Service of the relevant core ASes in accordance with a process called `_intra-ISD path segment registration_`. In addition, each core AS Control Service also stores the preferred core path segments to other core ASes during the `_core segment registration_` process.

Both processes are described below.

3.1. Intra-ISD Path Segment Registration

Every `_registration interval_` (configured by each AS) the AS's Control Service selects two sets of PCBs to transform into two types of path segments:

- * Up segments, which allow the infrastructure entities and endpoints in this AS to communicate with core ASes; and
- * Down segments, which allow remote entities to reach this AS.

The up segments and down segments do not have to be equal as an AS may want to communicate with core ASes via one or more up segments that differ from the down segment(s) through which it wants to be reached. Therefore, an AS operator can define different selection policies for the up segment and down segment sets. In addition, the processes of transforming a PCB in an up segment or a down segment differ slightly.

3.1.1. Terminating a PCB

Both the up segments and down segments end at the AS, so by transforming a PCB into a path segment, an AS Control Service "terminates" the PCB for this AS ingress interface and at that moment in time.

The Control Service of a non-core performs the following steps to "terminate" a PCB:

1. The Control Service adds a new AS entry to the PCB which MUST be defined as follows:
 - * The next AS MUST NOT be specified.
 - In Protobuf message format, this means that the value of the `next_isd_as` field in the `ASEntrySignedBody` component MUST be "0".

- * The egress interface in the Hop Field component MUST NOT be specified.
 - In Protobuf message format, this means that the value of the egress field in the HopField component MUST be "0".
- 2. If the AS has peering links, the Control Service MAY add corresponding peer entry components to the signed body of the AS entry - one peer entry component for each peering link that the AS wants to advertise. The Control Service MUST NOT specify the egress Interface ID in the Hop Field component of each added peer entry.
 - * In Protobuf message format, this means that the value of the egress field in the HopField component MUST be "0".
- 3. The Control Service MUST sign the modified PCB and append the computed signature.

Note:

- * For more information on the signed body component of an AS entry, see Section 2.2.2.2.
- * For more information on a peer entry, see Section 2.2.2.4.
- * For more information on the Hop Field component, see Section 2.2.2.5.
- * For more information on signing an AS entry, see Section 2.2.2.6.

3.1.2. Transforming a PCB into an Up Segment

Every registration interval, the Control Service of a non-core AS performs the following steps to transform PCBs into up segments:

1. The Control Service selects the PCBs that it wants to transform into up segments from the candidate PCBs in the Beacon Store.
2. The Control Service "terminates" the selected PCBs by performing the steps described in Section 3.1.1. From this moment on, the modified PCBs are called *up segments*.
3. The Control Service adds the newly created up segments to its own path database.

Note: For more information on possible selection strategies of PCBs, see Section 2.3.3.

3.1.3. Transforming a PCB into a Down Segment

Every registration interval, the Control Service of a non-core AS performs the following steps to transform PCBs into down segments:

1. The Control Service selects the PCBs that it wants to transform into down segments from the candidate PCBs in the Beacon Store.
2. The Control Service "terminates" the selected PCBs by performing the steps described in Section 3.1.1. From this moment on, the modified PCBs are called *down segments*.
3. The Control Service registers the newly created down segments with the Control Services of the core ASes that originated the corresponding PCBs. This is done by invoking the `SegmentRegistrationService.SegmentsRegistration` remote procedure call (RPC) in the Control Services of the relevant core ASes (see also Section 3.3). The first ISD-AS entry of the path segment SHOULD be equal to the core ISD-AS where the segment is being registered, otherwise the core AS MUST reject the segment. As a result, a core AS' s Control Service contains all down segments registered by its direct or indirect customer ASes.

Note: For more information on possible selection strategies of PCBs, see Section 2.3.3.

3.2. Core Path Segment Registration

The core beaconing process creates path segments from core AS to core AS. These core segments are then added to the Control Service path database of the core AS that created the segment so that local and remote endpoints can obtain and use these core segments. In contrast to the intra-ISD registration procedure, there is no need to register core segments with other core ASes as each core AS will receive PCBs originated from every other core AS.

In every registration interval, the Control Service of a core AS performs the following operations:

1. The core Control Service selects the best PCBs towards each core AS observed so far.
2. The core Control Service "terminates" the selected PCBs by performing the steps described in Section 3.1.1. From this moment on, the modified PCBs are called *core segments*.
3. The Control Service adds the newly created core segments to its own path database.

Note: For more information on possible selection strategies of PCBs, see Section 2.3.3.

3.3. Path Segment Registration RPC API

The Control Service of a non-core AS has to register the newly created down segments with the Control Services of the core ASes that originated the corresponding PCBs. This registration step is implemented as follows in Protobuf message format:

```
enum SegmentType {
    SEGMENT_TYPE_UNSPECIFIED = 0;
    SEGMENT_TYPE_UP = 1;
    SEGMENT_TYPE_DOWN = 2;
    SEGMENT_TYPE_CORE = 3;
}

service SegmentRegistrationService {
    rpc SegmentsRegistration(SegmentsRegistrationRequest) returns (
        SegmentsRegistrationResponse) {}
}

message SegmentsRegistrationRequest {
    message Segments {
        repeated PathSegment segments = 1;
    }

    map<int32, Segments> segments = 1;
}

message SegmentsRegistrationResponse {}
```

- * SegmentType: Specifies the type of the path segment to be registered. Currently, only the following type is used:
 - SEGMENT_TYPE_DOWN: Specifies a down segment.
- * map<int32, Segments> segments: Represents a separate list of segments for each path segment type. The key is the integer representation of the corresponding SegmentType.
- * SegmentRegistrationResponse: an empty message returned as an acknowledgement upon success.

3.4. Path MTU

SCION paths represent a sequence of ASes and inter-AS links, each with possibly different MTUs. As a result, the path MTU is the minimum of the MTUs of each inter-AS link and intra-AS network it traverses. Such MTU information is disseminated during path construction:

- * The MTU of each intra-AS network traversed (represented by the MTU field of the corresponding AS Entries (Section 2.2.2.2))
- * The MTU of each inter-AS link or peering link (indicated by the `ingress_mtu` field of each Section 2.2.2.3 or the `peer_mtu` field of each Section 2.2.2.4 used)

Such information is then made available to source endpoints during the path lookup process (see Section 4). Note that the destination endpoint does not receive such information, so when using path reversibility it should use mechanisms to estimate the reverse path MTU (e.g., MTU discovery or estimate MTU from the largest packet received). SCION endpoints are oblivious to the internal topology of intermediate ASes, so when looking up a path they should assume that all hops are also constrained by the intra-AS MTU of each AS traversed.

4. Path Lookup

The `_path lookup_` is a fundamental building block of SCION's path management as it enables endpoints to obtain path segments found during path exploration and registered during path registration. This allows the endpoints to construct end-to-end paths from the set of possible path segments returned by the path lookup process. The lookup of paths still happens in the control plane, whereas the construction of the actual end-to-end paths happens in the data plane.

4.1. Lookup Process

An endpoint (source) that wants to start communication with another endpoint (destination) requires up to three path segments:

- * An up segment to reach the core of the source ISD (only if the source endpoint is a non-core AS);
- * a core segment to reach
 - another core AS in the source ISD, in case the destination AS is in the same source ISD, or;

- a core AS in a remote ISD, if the destination AS is in another ISD, and;
- * a down segment to reach the destination AS.

The actual number of required path segments depends on the location of the destination AS as well as on the availability of shortcuts and peering links. More information on combining and constructing paths is provided by [I-D.dekater-scion-dataplane].

The process to look up and fetch path segments consists of the following steps:

1. The source endpoint queries the Control Service in its own AS (i.e., the source AS) for the required segments by sending up to three SegmentsRequest messages, respectively for up, core and down segments.
2. The Control Service of the source AS answers each request with a SegmentsResponse message. Specifically, for each segment type:
 - * up segments: The local Control Service has up segments stored in the path database and returns them immediately.
 - * core segments: The local Control Service may return locally cached core segments or query the Control Services of the reachable core ASes in the source ISD. These core ASes return core segments to core ASes in the destination ISD. To reach the core Control Services, the Control Service of the source AS uses the locally stored up segments. Once obtained, it returns the core segments and adds them to the local cache for future request.
 - * down segments: The local Control Service may return locally cached core segments or query the Control Services of the remote core ASes in the destination ISD. These remote core ASes return down segments to the destination AS. To reach the remote core ASes, the Control Service of the source AS uses the previously obtained and combined up and core segments. Once obtained, it returns the down segments.
3. As the source endpoint receives each path segment, it verifies the SegmentInformation timestamp validity (see Section 2.2.4), the AS entry signature for each AS entry (see Section 2.2.2.6), and requests any missing AS or intermediate certificates from the Control Service (see Section 2.4).

4. Once it has obtained some valid path segments, the source endpoint combines them into an end-to-end path in the data plane.
5. The destination endpoint, once it receives the first packet, MAY reverse the path in the received packet in order to construct a response. This ensures that traffic flows on the same path bidirectionally.

Table 3 below shows which Control Service provides the source endpoint with which type of path segment.

Segment Type	Responsible Control Service(s)
Up-segment	Control service of the source AS
Core-segment	Control service of core ASes in source ISD
Down-segment	Control service of core ASes in destination ISD (either the local ISD or a remote ISD)

Table 3: Control services responsible for different types of path segments

4.1.1. Sequence of Lookup Requests

The overall sequence of requests to resolve a path SHOULD be as follows:

1. Request up segments for the source endpoint at the Control Service of the source AS.
2. Request core segments, which start at the core ASes that are reachable with up segments and end at the core ASes in the destination ISD. If the destination ISD coincides with the source ISD, this step requests core segments to core ASes that the source endpoint cannot directly reach with an up segment.
3. Request down segments starting at core ASes in the destination ISD.

4.1.2. Lookup Requests Message Format

Control Services provide path segments through the SegmentLookupService RPC. This API is exposed:

- * for core ASes, on the SCION dataplane to other ASes

- * for all ASes, on the intra-AS network to endpoints.

```
service SegmentLookupService {  
    rpc Segments(SegmentsRequest) returns (SegmentsResponse) {}  
}
```

They use the following protobuf messages: a `SegmentsRequest`, which includes:

- * `src_isd_as`: The source ISD-AS of the segment.
- * `dst_isd_as`: The destination ISD-AS of the segment.

The corresponding `SegmentsResponse` returns:

- * `segments`: a list of `PathSegment` matching the request.
- * a mapping from path segment type to path segments, where the key is the integer representation of the `SegmentType` enum defined in Section 3.3.

```
message SegmentsRequest {  
    uint64 src_isd_as = 1;  
    uint64 dst_isd_as = 2;  
}  
  
message SegmentsResponse {  
    message Segments {  
        repeated PathSegment segments = 1;  
    }  
    map<int32, Segments> segments = 1;  
}
```

4.1.3. Caching

For the sake of efficiency, the Control Service of the source AS SHOULD cache each returned path segment request. Caching ensures that path lookups are fast for frequently used destinations and is also essential to ensure that the path lookup process is scalable and that endpoints can perform it with low latency.

In general, to improve overall efficiency, the Control Services of all ASes SHOULD do the following:

- * Cache the returned path segments.
- * Send requests in parallel when requesting path segments from other Control Services.

4.2. Behavior of Actors in the Lookup Process

The source endpoint resolves paths with a sequence of segment requests to the Control Service of the source AS. The Control Service in the source AS either answers directly or forwards these requests to the responsible Control Services of core ASes. In SCION, the instances that handle these segment requests at the Control Services are called `_source AS segment-request handler_` and `_core AS segment-request handler_`, respectively.

This section specifies the behavior of the segment request handlers in the lookup process.

4.2.1. Use of Wildcard Addresses in the Lookup Process

Endpoints can use wildcard addresses to designate any core AS in path segment requests. The segment request handlers **MUST** expand these wildcard addresses and translate them into one or more actual addresses. Table 4 below shows who is responsible for what.

Note: For general information on the use of wildcard addresses in SCION, see Section 1.5.2.1.

Segment Request	Wildcard Represents	Expanded/ Translated By	Translated Into
Up segment	"Destination" core AS (where up segment ends)	Control service of the <code>_source AS_</code>	Actual address destination core AS in source ISD
Core segment	Source core AS (where core segment starts)^1	Control service of the <code>_source AS_</code>	Actual address source core AS in source ISD
Core segment	Destination core AS (where core segment ends)	Control service of the <code>_source core AS_</code>	Actual address destination core AS in destination ISD
Down segment	"Source" core AS (where down segment starts)^2	Control service of the <code>_source AS_</code>	Actual address source core AS in destination ISD

Table 4: Use of wildcards in path segments requests

- 1) Includes all core ASes for which an up segment from the source AS exists.
- 2) Includes all core ASes in destination ISD with a down segment to destination AS.

4.2.2. Segment-Request Handler of a Non-Core Source AS

When the segment request handler of the Control Service of a `_non-core_source` AS receives a path segment request, it MUST proceed as follows:

1. Determine the requested segment type.
2. In the case of an up segment request, look up matching up segments in the path database and return them.
3. In the case of a core segment request from a source core AS to a destination core AS:
 - * Expand the source wildcard into separate requests for each reachable core AS in the source ISD.
 - * For each core segment request;
 - If possible, return matching core segments from cache;
 - Otherwise, request the core segments from the Control Services of each reachable core AS at the source of the core segment, and then add the retrieved core segments to the cache.
4. In the case of a down segment request:
 - * Expand the source wildcard into separate requests for every core AS in the destination ISD (referring to the ISD to which the destination endpoint belongs).
 - * For each segment request;
 - If possible, return matching down segments from cache;
 - Otherwise, request the down segment from the Control Services of the core ASes at the source of the down segment. Sending the request may require looking up core segments to the source core AS of the down segment, and then adding the retrieved down segments to the cache.

4.2.3. Segment-Request Handler of a Core AS

When the segment request handler of a `_core AS_` Control Service receives a path segment request, it **MUST** proceed as follows:

1. Validate the request:
 - * The source of the path segment **MUST** be this core AS.
 - * The request **MUST** either be;
 - for a core segment to a core AS in this ISD or another ISD, or;
 - for a down segment to an AS in this ISD.
2. If the destination is a core or wildcard address, then load matching core segments from the path database and return.
3. Otherwise, load the matching down segments from the path database and return.

Appendix "Path-Lookup Examples" shows by means of an illustration how the lookup of path segments in SCION works.

5. Control Service Discovery

The Control Plane RPC APIs rely on QUIC connections over UDP/SCION (see [I-D.dekater-scion-dataplane]). Establishing such connection requires the initiator to identify the relevant peer (service resolution) and to select a path to it. Since the Control Service is itself the source of path segment information, the following bootstrapping processes apply:

- * Neighboring ASes communicate using one-hop paths, as described in [I-D.dekater-scion-dataplane]. Core ASes leverage this mechanism when originating new PCBs.
- * Paths to non-neighboring ASes are constructed from PCBs that are received from neighbors via these one-hop paths.
- * The resulting multi-hop path segments are registered with the Control Service of the origin Core AS (see Section 3.1).
- * Control Services respond to requests from remote ASes by reversing the path carried in the request packet.

Clients find the relevant Control Service at a given AS by resolving a 'service address' as follows:

1. A client sends a ServiceResolutionRequest RPC (which has no parameters) to an endpoint address in the format:

- * Common Header:

- Path type: SCION (0x01)
- DT/DL: "Service" (0b0100)

- * Address Header:

- DstHostAddr: "SVC_CS" (0x0002)

- * UDP Header:

- DstPort: 0

A ServiceResolutionRequest MUST fit within a UDP datagram, otherwise clients and servers won't be able to establish control-plane reachability.

2. The ingress border router at the destination AS resolves the service destination to an actual endpoint address. This document does not mandate any specific method for this resolution.
3. The ingress border router forwards the message to the resolved address.
4. The destination service responds to the client with a ServiceResolutionResponse. It contains one or more transport options and it MUST fit within a UDP datagram. Known transports are "QUIC". Clients MUST ignore unknown values. The response includes a Transport message containing supported addresses and port to reach the service. Supported address formats for QUIC are IPv4 and IPv6. An example of the corresponding address format is: 192.0.2.1:80 and [2001:db8::1]:80. Clients MUST treat a missing, zero or non-existent port value as an error.
5. The client uses the address and port from the "QUIC" option to establish a QUIC connection, which can then be used for other RPCs.

The service resolution API Protobuf message format is:


```
message ServiceResolutionRequest {}

message ServiceResolutionResponse {
    map<string, Transport> transports = 1;
}

message Transport {
    string address = 1;
}
```

6. SCMP

The SCION Control Message Protocol (SCMP) provides functionality for network diagnostics, such as traceroute and error messages that signal packet processing or network layer problems. SCMP is a helpful tool for network diagnostics and - in the case of External Interface Down and Internal Connectivity Down messages - a signal for endpoints to detect network failures more rapidly and fail-over to different paths. However, SCION nodes should not strictly rely on the availability of SCMP, as this protocol may not be supported by all devices and/or may be subject to rate limiting.

This document only specifies the messages used for the purposes of path diagnosis and recovery. An extended specification can be found in [SCMP]. Its security considerations are discussed in Section 8.3.4.

The logic, some message formats, and processing rules are derived from [RFC4443] and adapted for the SCION architecture. Note that there is not currently a defined mechanism for converting ICMP messages to SCMP messages, or vice-versa.

6.1. General Format

Every SCMP message is preceded by a SCION header and zero or more SCION extension headers (see [I-D.dekater-scion-dataplane] section "SCION Header Specification"). The SCMP header is identified by a NextHdr value of 202 in the immediately preceding header.

The messages have the following general format:

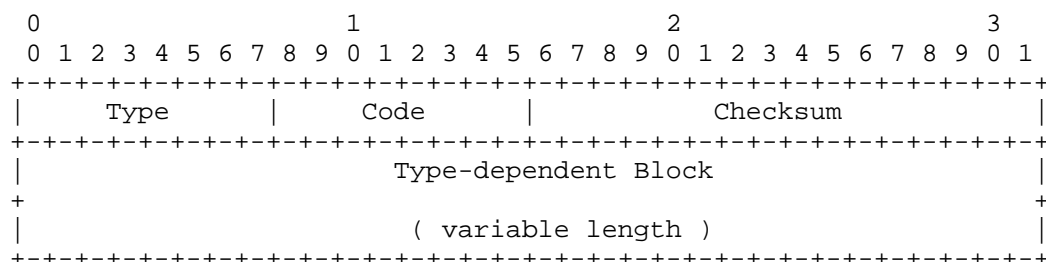


Figure 16: SCMP message format

- * Type: it indicates the type of SCMP message. Its value determines the format of the type-dependent block.
- * Code: it provides additional granularity to the SCMP type.
- * Checksum: it is used to detect data corruption.
- * Type-dependent Block: optional field of variable length which format is dependent on the message type.

6.2. Message Types

SCMP messages are grouped into two classes: error messages and informational messages. Error messages are identified by a zero in the high-order bit of the type value - i.e., error messages have a type value in the range of 0-127. Informational messages have type values in the range of 128-255.

This specification defines the message formats for the following SCMP messages:

Type	Meaning
2	Packet Too Big (Section 6.5.1)
5	External Interface Down (Section 6.5.2)
6	Internal Connectivity Down (Section 6.5.3)

Table 5: Error Message Types

Type	Meaning
128	Echo Request (Section 6.6.1)
129	Echo Reply (Section 6.6.2)
130	Traceroute Request (Section 6.6.3)
131	Traceroute Reply (Section 6.6.4)

Table 6: Informational Message Types

Further SCMP errors are covered in [SCMP].

6.3. Checksum Calculation

The checksum is calculated as the 16-bit one's complement of the one's complement sum of the entire SCMP message. This is starting with the SCMP message type field, and prepended with a "pseudo-header" consisting of the SCION address header and the Layer 4 protocol type as defined in [I-D.dekater-scion-dataplane] section "SCION Header Specification/Pseudo Header for Upper-Layer Checksum".

6.4. Processing Rules

The following rules apply to SCMP messages.

- * If an SCMP error message of unknown type is received at its destination, it MUST be passed to the upper-layer process that originated the packet that caused the error, if it can be identified.
- * If an SCMP informational message of unknown type is received, it MUST be silently dropped.
- * Every SCMP error message MUST include as much of the offending SCION packet as possible. The error message packet - including the SCION header and all extension headers MUST NOT exceed *1232 bytes* in order to fit into the minimum MTU (see [I-D.dekater-scion-dataplane] section "Deployment Considerations/MTU").
- * Implementations MUST attempt to pass SCMP error messages to the originator when possible. When using an UDP/IP underlay, implementations extract the upper-layer protocol type and port from the quoted offending packet in the body of the SCMP error

message and use it to determine the address of the originator to handle the error. Considerations related to port selection when using an UDP/IP underlay are described in [SCION-UDP]. In case of an unknown error type, implementations should assume a SCMP header length of 8 bytes, verify that the subsequent bytes represent a SCION header, and attempt to extract the offending packet. In case the origin address cannot be extracted from the SCMP error message, the SCMP message MUST be silently dropped.

- * An SCMP error message MUST NOT be originated in response to any of the following:
 - An SCMP error message.
 - A packet which source address does not uniquely identify a single node. E.g., an IPv4 or IPv6 multicast address.

[SCION-UDP] specifies the forwarding behavior of SCMP messages over an IP/UDP underlay.

The maximum size 1232 bytes is chosen so that the entire datagram, if encapsulated in UDP and IPv6, does not exceed 1280 bytes (L2 Header excluded). 1280 bytes is the minimum MTU required by IPv6 and it is assumed that this MTU can also be safely expected when using IPv4.

6.5. Error Messages

6.5.1. Packet Too Big

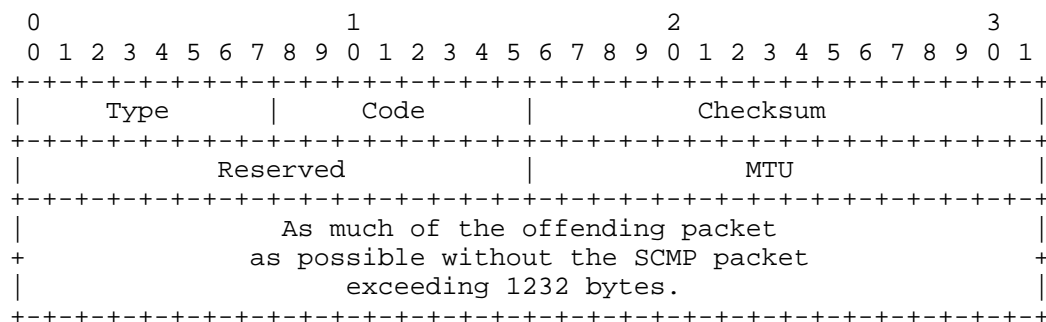


Figure 17: Packet Too Big format

+	=====	+
	Name Value	
+	=====	+
	Type 2	
+	-----	+
	Code 0	
+	-----	+
	MTU The Maximum Transmission Unit of the next-hop link.	
+	-----	+

Table 7: Error Message field values

A *Packet Too Big* message SHOULD be originated by a router in response to a packet that cannot be forwarded because the packet is larger than the MTU of the outgoing link. The router sets the MTU value to the maximum size a SCION packet can have to still fit on the next-hop link, as the sender has no knowledge of the underlay.

6.5.2. External Interface Down

0	1	2	3
0 1 2 3 4 5 6 7 8 9	0 1 2 3 4 5 6 7 8 9	0 1 2 3 4 5 6 7 8 9	0 1
+	+	+	+
	Type		Code
+	+	+	+
	ISD		Checksum
+	+	+	+
			AS
+	+	+	+
+	+	+	+
			Interface ID
+	+	+	+
+	+	+	+
			As much of the offending packet
+	+	+	as possible without the SCMP packet
			exceeding 1232 bytes.
+	+	+	+

Figure 18: External Interface Down format

Name	Value
Type	5
Code	0
ISD	The 16-bit ISD identifier of the SCMP originator
AS	The 48-bit AS identifier of the SCMP originator
Interface ID	The interface ID of the external link with connectivity issue.

Table 8: Error Message field values

A **External Interface Down** message SHOULD be originated by a router in response to a packet that cannot be forwarded because the link to an external AS is broken. The ISD and AS identifier are set to the ISD-AS of the originating router. The Interface ID identifies the link of the originating AS that is down. Recipients can use this information to route around broken data-plane links.

6.5.3. Internal Connectivity Down

0	1	2	3
0 1 2 3 4 5 6 7 8 9	0 1 2 3 4 5 6 7 8 9	0 1 2 3 4 5 6 7 8 9	0 1
Type	Code	Checksum	
ISD		AS	
Ingress Interface ID			
Egress Interface ID			
As much of the offending packet as possible without the SCMP packet exceeding 1232 bytes.			

Figure 19: Internal Connectivity Down format

Name	Value
Type	6
Code	0
ISD	The 16-bit ISD identifier of the SCMP originator
AS	The 48-bit SCION AS identifier of the SCMP originator
Ingress ID	The interface ID of the ingress link.
Egress ID	The interface ID of the egress link.

Table 9: Error Message field values

A *Internal Connectivity Down* message SHOULD be originated by a router in response to a packet that cannot be forwarded inside the AS because the connectivity between the ingress and egress routers is broken. The ISD and AS identifier are set to the ISD-AS of the originating router. The ingress Interface ID identifies the interface on which the packet enters the AS. The egress Interface ID identifies the interface on which the packet is destined to leave the AS, but the connection to which is broken.

Recipients can use this information to route around a broken data plane inside an AS.

6.6. Informational Messages

6.6.1. Echo Request

0	1	2	3
0 1 2 3 4 5 6 7 8 9	0 1 2 3 4 5 6 7 8 9	0 1 2 3 4 5 6 7 8 9	0 1
Type	Code	Checksum	
Identifier	Sequence Number		
Data (variable Len.)			

Figure 20: Echo Request format

Name	Value
Type	128
Code	0
Identifier	A 16-bit identifier to aid matching replies with requests
Sequence Nr.	A 16-bit sequence number to aid matching replies with requests
Data	Variable length of arbitrary data

Table 10: Informational Message field values

Every node SHOULD implement a SCMP Echo responder function that receives Echo Requests and originates corresponding Echo replies.

6.6.2. Echo Reply

0	1	2	3
0 1 2 3 4 5 6 7 8 9	0 1 2 3 4 5 6 7 8 9	0 1 2 3 4 5 6 7 8 9	0 1
Type	Code	Checksum	
Identifier	Sequence Number		
Data (variable Len.)			

Figure 21: Echo Reply format

Name	Value
Type	129
Code	0
Identifier	The identifier of the Echo Request
Sequence Nr.	The sequence number of the Echo Request
Data	The data of the Echo Request

Table 11: Informational Message field values

Every node SHOULD implement a SCMP Echo responder function that receives Echo Requests and originates corresponding Echo replies.

The data received in the SCMP Echo Request message MUST be returned entirely and unmodified in the SCMP Echo Reply message.

6.6.3. Traceroute Request

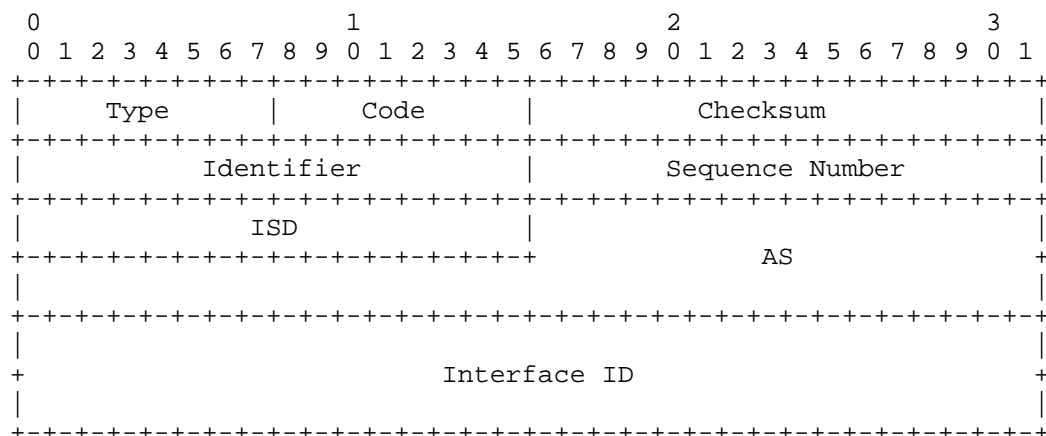


Figure 22: Traceroute Request format

Given a SCION path constituted of hop fields, traceroute allows to identify the corresponding on-path ISD-ASes.

Name	Value
Type	130
Code	0
Identifier	A 16-bit identifier to aid matching replies with requests
Sequence Nr.	A 16-bit sequence number to aid matching replies with request
ISD	Place holder set to zero by SCMP sender
AS	Place holder set to zero by SCMP sender
Interface ID	Place holder set to zero by SCMP sender

Table 12: Informational Message field values

A border router is alerted of a Traceroute Request message through the Ingress or Egress Router Alert flag set to 1 in the hop field that describes the traversal of that router in a packet's path (see [I-D.dekater-scion-dataplane] section "SCION Header Specification"). When such a packet is received, the border router SHOULD reply with a Traceroute Reply message (Section 6.6.4).

6.6.4. Traceroute Reply




Figure 23: Traceroute Reply format

Name	Value
Type	131
Code	0
Identifier	The identifier set in the Traceroute Request
Sequence Nr.	The sequence number of the Traceroute Request
ISD	The 16-bit ISD identifier of the SCMP originator
AS	The 48-bit AS identifier of the SCMP originator
Interface ID	The interface ID of the SCMP originating router

Table 13: Informational Message field values

The identifier is set to the identifier value from the Traceroute Request message (Section 6.6.3). The ISD and AS identifiers are set to the ISD-AS of the originating border router.

7. Deployment Considerations

7.1. Destination Mapping

The mechanism by which endpoints determine the destination ISD-AS corresponding to a given destination address is outside the scope of this document. One option, still experimental in existing deployments, is that SCION-aware endpoints may resolve destination SCION addresses using a naming system (e.g., DNS).

SCION-unaware endpoints may interface with a SCION network through a SCION IP Gateway (SIG), which tunnels IP traffic over SCION. In such cases, the source SIG is responsible for mapping destination IPs to the appropriate destination ISD-AS and gateway. More information can be found at [SIG].

7.2. Monitoring Considerations

In order to maintain service availability, an AS operator SHOULD monitor the following aspects when deploying the SCION control plane:

- * For routers (to enable correlation with link states): state of configured links (core, child, parent).
- * For any Control Service:
 - Fraction of path lookups served successfully (see Section 4).
 - Time synchronization offset with other ASes (see Section 7.3).
 - Fraction of ASes found in non-expired segments for which a non-expired certificate exists.
- * For a core AS:
 - Fraction of core ASes (preferably only those to which the link is up) that can be found in non-expired core segments.
 - Fraction of ASes, core or children, (preferably only those to which the link is up) to where a beacon was originated during the last propagation interval.
 - Fraction of freshly propagated beacons for which at least one corresponding down segment has been registered (see Section 3).
- * For a non-core AS:
 - Number of up segments available (may be just 0/non-0) younger than the propagation interval (or some multiple thereof).

- Fraction of up segments that were successfully registered as down segments (see Section 3).
- Fraction of children ASes (preferably only those to which the link is up) to where a beacon was propagated during the last propagation interval.

7.3. Effects of Clock Inaccuracy

A PCB originated by a given core AS Control Service is validated by all the Control Services that receive it. All have different clocks and their differences affect the validation process:

- * A fast clock at origination or a slow clock at reception will yield a lengthened expiration time for hops, and possibly an origination time in the future.
- * A slow clock at origination or a fast clock at reception will yield a shortened expiration time for hops, and possibly an expiration time in the past.

This bias comes in addition to a structural delay: PCBs are propagated at a configurable interval, typically around one minute. As a result of this and the manner in which they are iteratively constructed, PCBs with N hops may be validated up to N intervals (so maximally N minutes) after origination which creates a constraint on the expiration of hops. Hops of the minimal expiration time (337.5 seconds - see Section 2.2.2.5) would make any PCB describing a path longer than 5 hops expire. For this reason it is unadvisable to create hops with a short expiration time - i.e., they SHOULD be around 6 hours.

The Control Service and its clients authenticate each other in accordance with their respective AS's certificate. Path segments are authenticated based on the certificates of the ASes that they refer to. The RECOMMENDED expiration time of a SCION AS certificate is between 3h and 3 days, although some deployments use up to 5 days. In comparison to these time scales, clock offsets in the order of minutes are immaterial.

Each administrator of a SCION Control Service is responsible for maintaining coarse time synchronization with SCION routers within the AS, neighbor ASes Control Services, and endpoints within the AS. In typical deployments, clock deviations on the order of several minutes are acceptable.

The specific methods used to achieve this synchronization are outside the scope of this document. Security considerations on time synchronization are discussed in Section 8.4.

7.4. Path Discovery Time and Scalability

The path discovery mechanism balances the number of discovered paths and the time it takes to discover them versus resource overhead of the discovery.

The resource costs for path discovery are as follows:

- * Communication overhead is transmitting the PCBs and occasionally obtaining the required PKI material.
- * Processing overhead is validating the signatures of the AS entries, signing new AS entries, and to a lesser extent, evaluating the PCB selection policies.
- * Storage overhead is both the temporary storage of PCBs before the next propagation interval, and the storage of complete discovered path segments.

All of these are dependent on the number and length of the discovered path segments, i.e., the total number of AS entries of the discovered path segments. These in turn depend on the configured best PCBs set size (Section 2.3.4).

Relevant metrics for scalability and speed of path discovery are the time until all discoverable path segments have been discovered after a network bootstrap, and the time until a new link is usable. In general, the time until a specific PCB is built depends on its length, the propagation interval, and whether on-path ASes use "fast recovery" (see Section 2.3.4).

At each AS, the Control Service will process and propagate the PCB at the subsequent propagation event. As propagation events are not synchronized between different ASes, a PCB arrives at a random point in time during the interval and may be buffered before potentially being propagated. With a propagation interval T at each AS, the mean time until the PCB is propagated in one AS therefore is $T / 2$ and the mean total time for the propagation steps of a PCB of length L is at worst $L * T / 2$ (with a variance of $L * T^2 / 12$).

Note that link removal is not part of path discovery in SCION. For scheduled removal of links, operators let path segments expire. On link failures, endpoints route around the failed link by switching to different paths in the data plane (see [I-D.dekater-scion-dataplane] section "Handling Link Failures").

To achieve scalability, SCION ASes are partitioned into ISDs and in an ideal topology the inter-ISD core network should be kept to a moderate size. More specific observations require a distinction between intra-ISD and core beaconing.

7.4.1. Intra-ISD Beaconing

In the intra-ISD beaconing, PCBs are propagated top down along parent-child links from core to downstream ASes. Each AS discovers path segments from itself to the core ASes of its ISD.

This produces an acyclic graph which is typically narrow at the top, widens towards downstream ASes, and is relatively shallow. Intermediate provider ASes will typically have a large number of children whilst only having a small number of parents. Therefore the chain of intermediate providers from a downstream AS to a core AS is typically not long (e.g., local, regional, national provider, then core).

Each AS potentially receives PCBs for all down path segments from the core to itself. While the number of distinct provider chains to the core is typically moderate, the multiplicity of links between provider ASes has multiplicative effect on the number of PCBs. Once this number grows above the maximum recommended best PCB set size of 50, ASes SHOULD trim the set of PCBs propagated.

Ultimately, the number of PCBs received by an AS per propagation interval remains bounded by 50 for each parent link of an AS, and at most 50 PCBs per child link are propagated. The length of these PCBs and thus the number of AS entries to be processed and stored, is expected to be moderate and not grow considerably with network size. The total resource overhead for beacon propagation is easily manageable even for highly connected ASes.

To illustrate this, an AS with a rather large number of 100 parent links receives at most 5,000 PCBs during a propagation interval. Assuming a generous average length of 10 AS entries for these PCBs, this corresponds to 50,000 AS entries.

Due to the variable length fields in AS entries, the sizes for storage and transmission cannot be predicted exactly, but assume an average of 250 bytes per AS entry. At the shortest recommended

propagation interval of 5 seconds, this corresponds to an average bandwidth of around 2.5 MB/s and the processing of 10,000 signature verifications per second.

If the same AS has 1,000 child links, the propagation of the beacons will require signing one new AS entry for each of the propagated PCBs for each link (at most 50 per link) - i.e., at most 50,000 signatures per propagation event. The total bandwidth for the propagation of the PCBs for all 1,000 child links would be roughly around 25 MB/s which is manageable with even modest consumer hardware.

On a network bootstrap, path segments to each AS are discovered within a number of propagation steps proportional to the longest path. With a 5 second propagation interval and a generous longest path of length 10, all path segments are discovered after 25 seconds on average. When all ASes start propagation just after they've received the first PCBs from any of their upstreams (see "fast recovery" in Section 2.3.4), the construction of a first path to connect each AS to the ISD core is accelerated.

When a new parent-child link is added to the network, the parent AS will propagate the available PCBs in the next propagation event. If the AS on the child side of the new link has no children of its own, path discovery is complete after at most one propagation interval. Otherwise, child ASes at distance D below the new link, learn of the new link after at worst D further propagation intervals.

7.4.2. Core Beaconing

In core beaconing (typically inter-ISD), PCBs are propagated omnidirectionally along core links. Each AS discovers path segments from itself to any other core AS.

The number of distinct paths through the core network is typically very large. To keep the overhead manageable, at most 5 path segments to every destination AS are discovered and the propagation frequency is slower than in the intra-ISD beaconing (at least 60 seconds between propagation events).

Without making strong assumptions on the topology of the core network, it can be assumed that shortest paths through real world networks are relatively short - e.g., the Barabasi-Albert random graph model predicts a diameter of $\log(N)/\log(\log(N))$ for a network with N nodes [BollRio-2000] and the average distance scales in the same way. Whilst it cannot be assumed that the selected PCBs are strictly the shortest paths through the network, they are likely to be not very much longer than the shortest paths either.

With N the number of participating core ASes, an AS receives up to $5 * N$ PCBs per propagation interval per core link interface. For highly connected ASes, the number of PCBs received thus becomes rather large and in a network of 1,000 ASes, an AS with 300 core links receives up to 1.5 million PCBs per propagation interval.

Assuming an average PCB length of 6 and the shortest propagation interval of 60 seconds, this corresponds to roughly 150,000 signature validations per second or roughly 38 MB/s. For much larger, more highly connected ASes, the path discovery tasks of the Control Service can be distributed over many instances in order to increase the PCB throughput.

On a network bootstrap, full connectivity is obtained after a number of propagation steps corresponding to the diameter of the network. Assuming a network diameter of 6, this corresponds to roughly 3 minutes on average. When a new link is added to the network, it will be available to connect two ASes at distances $D1$ and $D2$ from the link respectively, at worst after a mean time $(D1+D2)*T/2$.

8. Security Considerations

As described previously, the goal of SCION's beaconing process in the control plane is to securely discover and disseminate paths between any two ASes. This section describes security considerations for SCION's Control Plane that focuses on inter-domain routing. SCION does not provide intra-domain routing, nor does it provide end-to-end payload encryption so these topics lie outside the scope of this section.

This section discusses three kinds of threats to the control plane:

1. When an adversary controls one or all core ASes of an ISD and tries to manipulate the beaconing process from the top down (see Section 8.2).
2. When "ordinary" (non-core) adversaries try to manipulate the beaconing process (see Section 8.3).
3. Denial of Services (DoS) attacks where attackers overload different parts of the infrastructure (see Section 8.5).

8.1. Security Properties

The SCION control plane provides various security properties, as discussed below. Here a SCION AS is described as 'honest' if its private keys are unknown to the attacker and it correctly performs operations in accordance with this specification (e.g., uses a unique interface identifier for each link). An honest path is one that only traverses honest ASes. An honest segment is one created by an honest AS.

Security properties are:

- * Connectivity - For every pair of honest ASes X and Y, X will eventually register enough segments to build at least one path (of any length) leading to Y.
- * Forwarding Path Consistency - For every honest path segment registered in any AS
 - its sequence of AS entries corresponds to a continuous SCION forwarding path in the network of inter-domain links
 - the inter-domain network topology remains unchanged since the segment was first generated.
- * Loop Freedom - For every honest path segment registered in any AS, its sequence of AS entries contains no duplicates, including current and next ISD-AS and Interface IDs.
- * Path Authorization - For every honest path segment registered in any AS and any AS X appearing on that segment (except for the previous one), AS X propagated a PCB corresponding to the segment portion ending in its own entry to its successor AS on the segment.

To ensure that the properties hold across the overall SCION network, these are the prerequisites to follow: - all core ASes MUST be able to reach each other with some sequence of core links; - and all non-core ASes MUST have at least one path up to a core AS.

Furthermore, to ensure that the properties hold within a single ISD, all core ASes of the ISD MUST be able to reach each other without leaving the ISD - i.e, for every pair of cores in an ISD there is a sequence of SCION links that only traverse the ISD members.

A core AS may reach other core ASes in the same ISD via other ISDs, depending on the ISD's policies.

8.2. Manipulation of the Beaconsing Process by a Core Adversary

The first risk to the beaconsing process comes from an adversary controlling one or more core ASes in an ISD. Non-core ASes depend on core ASes for beaconsing. If an adversary compromises all upstream core ASes of a target AS and suppresses PCB propagation, the discovery of new paths for will halt. Although the downstream AS will no longer receive new PCBs, previously discovered and still valid paths remain usable for data plane forwarding until they expire. This is an unlikely attack scenario, as it would require compromise of all core ASes of a given downstream AS to be effective.

8.3. Manipulation of the Beaconsing Process by a Non-Core Adversary

This section examines several possible approaches that could be taken by an "ordinary" non-core adversary to manipulate the beaconsing process in the Control Plane. For each case it shows to what extent SCION's design can prevent the corresponding attack or help mitigate it.

8.3.1. Path Hijacking through Interposition

A malicious AS M might try to manipulate the beaconsing process between two neighbor ASes A and B, with the goal to hijack traffic to flow via AS M. If AS M can interpose itself on the path between A and B, then it could attempt several potential attacks:

- * The adversary AS M could intercept and disseminate a PCB on its way from A to the neighboring AS B, and inject its own AS entry into the PCB toward downstream ASes.
- * The adversary could modify the Hop Fields of an already existing path in order to insert its own AS in the path.
- * The adversary could fully block traffic between AS A and AS B in order to force traffic redirection through an alternate path that includes its own AS.

The first type of attack is detectable and blocked by downstream ASes (e.g., AS B) because a PCB disseminated by AS A towards AS B contains the "Next ISD AS" field in the entry of AS A, pointing to AS B and protected by A's signature. If AS M manipulates the PCB while in flight from AS A to AS B, then verification of the manipulated inbound PCBs will fail at AS B, as the adversary's PCBs cannot contain AS A's correct signature (see Section 2.3.1).

The second type of attack is made impossible by the Hop Field's MAC which protects the Hop Field's integrity and chains it with the previous Hop Fields on the path.

The third type of attack generally cannot be prevented. However the alternate path would be immediately visible to endpoints as traffic MUST include Hop Fields from AS M.

8.3.2. Creation of Spurious ASes and ISDs

An alternative scenario is when an adversary tries to introduce and spoof a non-existent AS. This would enable the adversary to send traffic with the spoofed AS as a source, allowing the adversary to complicate the detection of its attack and to plausibly deny the misbehavior.

However, spoofing a new AS requires a registration of that AS with the ISD core to obtain a valid AS certificate, otherwise the adversary cannot construct valid PCBs. As this registration should include a thorough check and authentication by a CA, this cannot be done stealthily which defeats the original purpose.

Similarly to creating a fake AS, an adversary could try to introduce a new malicious ISD. This involves the generation of its own TRC, finding core ASes to peer with, and convincing other ISDs of its legitimacy to accept the new TRC. Although this setup is not entirely impossible, it requires substantial time and effort and may need the involvement of more than one malicious entity. Here the "costs" of setting up the fake ISD may outweigh the benefits.

8.3.3. Peering Link Misuse

The misuse of a peering link by an adversary represents another type of attack. Consider the case where AS A wants to share its peering link only with one of its downstream neighbors AS B, and therefore selectively includes the peering link only in PCBs sent to AS B. An adversary may now try to gain access to this peering link by prepending the relevant PCBs to its own path. For this, the adversary needs to be able to (1) eavesdrop on the link from AS A to AS B, and (2) obtain the necessary Hop Fields by querying a Control Service and extracting the Hop Fields from registered paths.

Even if an adversary succeeds in misusing a peering link as described above, SCION is able to mitigate this kind of attack. Each AS includes an egress interface as well as specific “next hop” information to the PCB before disseminating it further downstream. If a malicious entity tries to misuse a stolen PCB by adding it to its own segments, verification will fail upstream as the egress interface mismatches. Therefore, the peering link can only be used by the intended AS.

8.3.4. Manipulation of the Path-Selection Process

SCION endpoints select inter-domain forwarding paths and this section discusses some mechanisms by which an adversary can attempt to trick endpoints downstream (in the direction of beaconing) into choosing non-optimal paths. The goal of such attacks is to make paths that are controlled by the adversary more attractive than other available paths.

In SCION, overall path selection is the result of three steps. Firstly, each AS selects which PCBs are further forwarded to its neighbors. Secondly, each AS chooses the paths it wants to register at the local Control Service (as up segments) and at the core Control Service (as down segments). Thirdly, the endpoint performs path selection among all available paths resulting from a path lookup process.

These attacks are only successful if the adversary is located within the same ISD and upstream relative to the victim AS. It is not possible to attract traffic away from the core as traffic travels upstream towards the core. Furthermore, the attack may either be discovered downstream (e.g., by seeing large numbers of paths becoming available), or during path registrations. After detection, non-core ASes will be able to identify paths traversing the adversary AS and avoid these paths.

Announcing Large Numbers of Path Segments

This attack is possible if the adversary controls at least two ASes. The adversary can create a large number of links between the ASes under its control which do not necessarily correspond to physical links. This allows the adversary to multiply the number of PCBs forwarded to its downstream neighbor ASes and in turn increases the chance that one or several of these forwarded PCBs are selected by the downstream ASes.

In general, the number of PCBs that an adversary can announce this way scales exponentially with the number of consecutive ASes the adversary controls. However, this also decreases their chance of being chosen by a downstream AS for PCB dissemination or by an

endpoint for path construction as these relatively long paths have to compete with other shorter paths. Furthermore, both endpoints and downstream ASes can detect poorer quality paths in the data plane and switch to better paths.

Wormhole Attack

A malicious AS M1 can send a PCB not only to their downstream neighbor ASes, but also out-of-band to another non-neighbor colluding malicious AS M2. This creates new segments to M2 and M2's downstream neighbor ASes, simulating a link between M1 and M2 which may not correspond to an actual link in the network topology.

Similarly, two colluding ASes could announce a path through a fake peering link between them, even if in different ISDs, thus offering short paths to many destination ASes. Downstream ASes might have a policy of preferring paths with many peering links and thus are more likely to disseminate PCBs from the adversary. Endpoints are also more likely to choose short paths that make use of peering links.

In the data plane, whenever the adversary receives a packet containing a fake peering link, it can transparently exchange the fake peering Hop Fields with valid Hop Fields to the colluding AS. To avoid detection of the path alteration by the receiver, the colluding AS can replace the added Hop Fields with the fake peering link Hop Fields the sender inserted.

To defend against this attack, methods to detect the wormhole attack are needed. Per link or path latency measurements can help reveal the wormhole and render the fake peering link suspicious or unattractive. Without specific detection mechanisms these so-called wormhole attacks are unavoidable in routing.

Rogue SCMP Error Messages

SCMP External Interface Down (Section 6.5.2) and Internal Connectivity Down (Section 6.5.3) can potentially be abused by an attacker to disrupt forwarding of information and/or force the traffic through a different paths. Endpoints should therefore consider them weak hints and apply heuristics to detect fraudulent SCMP messages (e.g., by actively probing whether the affected path is actually down).

Note that this would be mitigated through authentication of SCMP messages. Authentication is not specified here since it is currently still experimental.

8.4. Attacks on time sources

Operators should maintain coarse time synchronization among Control Service instances and other system components, as discussed in Section 7.3. An adversary that significantly alters the system time of a component can disrupt SCION operations:

- * A Control Service instance: its beaconing process may halt as it cannot verify the validity of received PCBs (see Section 2.2.4) or correctly add timestamps to propagated PCBs (see Section 2.1.2).
- * An endpoint: it may fail to verify path segments during path lookup (see Section 4.1).
- * A router: packets may be dropped ahead of the Control Service intended expiration time (see Section 2.2.2.5).
- * A Certificate Authority (see [I-D.dekater-scion-pki]): it may issue AS certificates with incorrect validity periods, causing them to be rejected by verifiers.

It is therefore recommended to leverage secure time synchronization mechanisms, such as NTS [RFC8915], [BCP223], or Khronos [RFC9523], or to leverage multiple diverse time sources (e.g., GNSS and network-based).

8.5. Denial of Service Attacks

The beaconing process in the SCION Control Plane relies on control plane communication. ASes exchange control plane messages within each other when propagating PCBs to downstream neighbors, when registering PCBs as path segments, or during core path lookup. Volumetric DoS attacks, where attackers overload a link may make it difficult to exchange these messages.

SCION limits the impact of such attacks which aim to exhaust network bandwidth on links as ASes can switch to alternative paths that do not contain the congested links. Reflection-based attacks are also prevented as response packets are returned on the same path to the actual sender.

Other mechanisms are required to avoid transport protocol attacks where the attacker tries to exhaust the resources on a target server, such as for the Control Services, by opening many connections to this. The means to mitigate these kind of DoS attacks are basically the same as for the current Internet - e.g., filtering, geo-blocking or using cookies.

Thanks to its path awareness, SCION enables more fine-grained filtering mechanisms based on certain path properties. For example, control plane RPC methods that are available to endpoints within an AS are strictly separate from methods available to endpoints from other ASes. Specifically, expensive recursive path segment and trust material lookups are thus shielded from abuse by unauthorized entities.

For RPC methods exposed to other ASes, the Control Service implementation minimizes its attack surface by rejecting illegitimate callers based on ISD/AS, path type and length and any other available data points as soon as possible, i.e., immediately after determining the request type. For example:

- * `SegmentCreationService.Beacon` can only be called by direct neighbors and thus calls from peers with a path length greater than one can immediately be discarded.
- * `SegmentRegistrationService.SegmentsRegistration` can only be called from within the same ISD, thus the source address MUST match the local ISD and the number of path segments MUST be 1.

A combination of the mechanism above is used to prevent flooding attacks on the Control Service. In addition, operators SHOULD deploy the Control Service in a distributed and replicated manner so that requests can be balanced and a single instance failure does not result in a complete failure of the control plane of a SCION AS.

9. IANA Considerations

This document has no IANA actions.

The ISD and SCION AS number are SCION-specific numbers. They are allocated by the SCION Association (see [ISD-AS-assignments]).

10. References

10.1. Normative References

- [Connect] "Connect Protocol Reference", 2024,
<<https://connectrpc.com/docs/protocol/>>.
- [gRPC] "gRPC, an open-source universal RPC framework", 2023,
<<https://grpc.io/>>.
- [I-D.dekater-scion-dataplane]
de Kater, C., Rustignoli, N., Hugly, J., and S. Hitz,
"SCION Data Plane", Work in Progress, Internet-Draft,

draft-dekater-scion-dataplane-11, 6 February 2026,
<<https://datatracker.ietf.org/doc/html/draft-dekater-scion-dataplane-11>>.

[I-D.dekater-scion-pki]

de Kater, C., Rustignoli, N., and S. Hitz, "SCION Control Plane PKI", Work in Progress, Internet-Draft, draft-dekater-scion-pki-11, 16 January 2026,
<<https://datatracker.ietf.org/doc/html/draft-dekater-scion-pki-11>>.

[POSIX.1-2024]

"Standard for Information Technology--Portable Operating System Interface (POSIX) Base Specifications, Issue 8", 2024, <https://pubs.opengroup.org/onlinepubs/9799919799/basedefs/V1_chap04.html>.

[proto3]

"Protocol Buffers Language Guide version 3", 2023,
<<https://protobuf.dev/programming-guides/proto3/>>.

[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>>.

[RFC4632]

Fuller, V. and T. Li, "Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan", BCP 122, RFC 4632, DOI 10.17487/RFC4632, August 2006, <<https://www.rfc-editor.org/rfc/rfc4632>>.

[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>>.

[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>>.

[RFC9000]

Iyengar, J., Ed. and M. Thomson, Ed., "QUIC: A UDP-Based Multiplexed and Secure Transport", RFC 9000, DOI 10.17487/RFC9000, May 2021,
<<https://www.rfc-editor.org/rfc/rfc9000>>.

[RFC9114]

Bishop, M., Ed., "HTTP/3", RFC 9114, DOI 10.17487/RFC9114, June 2022, <<https://www.rfc-editor.org/rfc/rfc9114>>.

10.2. Informative References

- [BCP223] Best Current Practice 223,
<<https://www.rfc-editor.org/info/bcp223>>.
At the time of writing, this BCP comprises the following:
- Reilly, D., Stenn, H., and D. Sibold, "Network Time Protocol Best Current Practices", BCP 223, RFC 8633, DOI 10.17487/RFC8633, July 2019,
<<https://www.rfc-editor.org/info/rfc8633>>.
- [BollRio-2000]
Bollobas, B. and O. Riordan, "The diameter of a scale-free random graph", n.d.,
<https://kam.mff.cuni.cz/~ksemweb/clanky/BollobasR-scale_free_random.pdf>.
- [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>>.
- [ISD-AS-assignments]
"SCION Registry", 2026, <<http://scion.org/registry/>>.
- [PCBExtensions]
Anapaya, ETH, and SCION Association, "PCB Path Metadata Extension", 2025,
<<https://docs.scion.org/en/latest/beacon-metadata.html>>.
- [RFC1122] Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, DOI 10.17487/RFC1122, October 1989,
<<https://www.rfc-editor.org/rfc/rfc1122>>.
- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006,
<<https://www.rfc-editor.org/rfc/rfc4271>>.
- [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>>.

- [RFC5398] Huston, G., "Autonomous System (AS) Number Reservation for Documentation Use", RFC 5398, DOI 10.17487/RFC5398, December 2008, <<https://www.rfc-editor.org/rfc/rfc5398>>.
- [RFC6996] Mitchell, J., "Autonomous System (AS) Reservation for Private Use", BCP 6, RFC 6996, DOI 10.17487/RFC6996, July 2013, <<https://www.rfc-editor.org/rfc/rfc6996>>.
- [RFC8915] Franke, D., Sibold, D., Teichel, K., Dansarie, M., and R. Sundblad, "Network Time Security for the Network Time Protocol", RFC 8915, DOI 10.17487/RFC8915, September 2020, <<https://www.rfc-editor.org/rfc/rfc8915>>.
- [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>>.
- [RFC9473] Enghardt, R. and C. Krhenbhl, "A Vocabulary of Path Properties", RFC 9473, DOI 10.17487/RFC9473, September 2023, <<https://www.rfc-editor.org/rfc/rfc9473>>.
- [RFC9523] Rozen-Schiff, N., Dolev, D., Mizrahi, T., and M. Schapira, "A Secure Selection and Filtering Mechanism for the Network Time Protocol with Khronos", RFC 9523, DOI 10.17487/RFC9523, February 2024, <<https://www.rfc-editor.org/rfc/rfc9523>>.
- [SCION-UDP] Anapaya, ETH, and SCION, "SCION IP/UDP underlay", 2026, <<https://docs.scion.org/en/latest/protocols/underlay.html>>.
- [SCMP] Anapaya, ETH Zuerich, and SCION Association, "SCMP Documentation", 2025, <<https://docs.scion.org/en/latest/protocols/scmp.html>>.
- [SIG] Anapaya, ETH Zuerich, and SCION, "SCION IP Gateway Documentation", 2024, <<https://docs.scion.org/en/latest/sig.html>>.

Path-Lookup Examples

To illustrate how the path lookup works, two path-lookup examples are shown in sequence diagrams. The network topology of the examples is represented in Figure 24 below, where in both the source endpoint is in AS A.

Figure 25 shows the sequence diagram for the path lookup process in case the destination is in AS D, whereas Figure 26 shows the path lookup sequence diagram if the destination is in AS G. ASes B and C are core ASes in the source ISD, while E and F are core ASes in a remote ISD. Core AS B is a provider of the local AS, but AS C is not - i.e., there is no up-segment from A to C. "CS" stands for Control Service.





Figure 24: Topology used in the path lookup examples



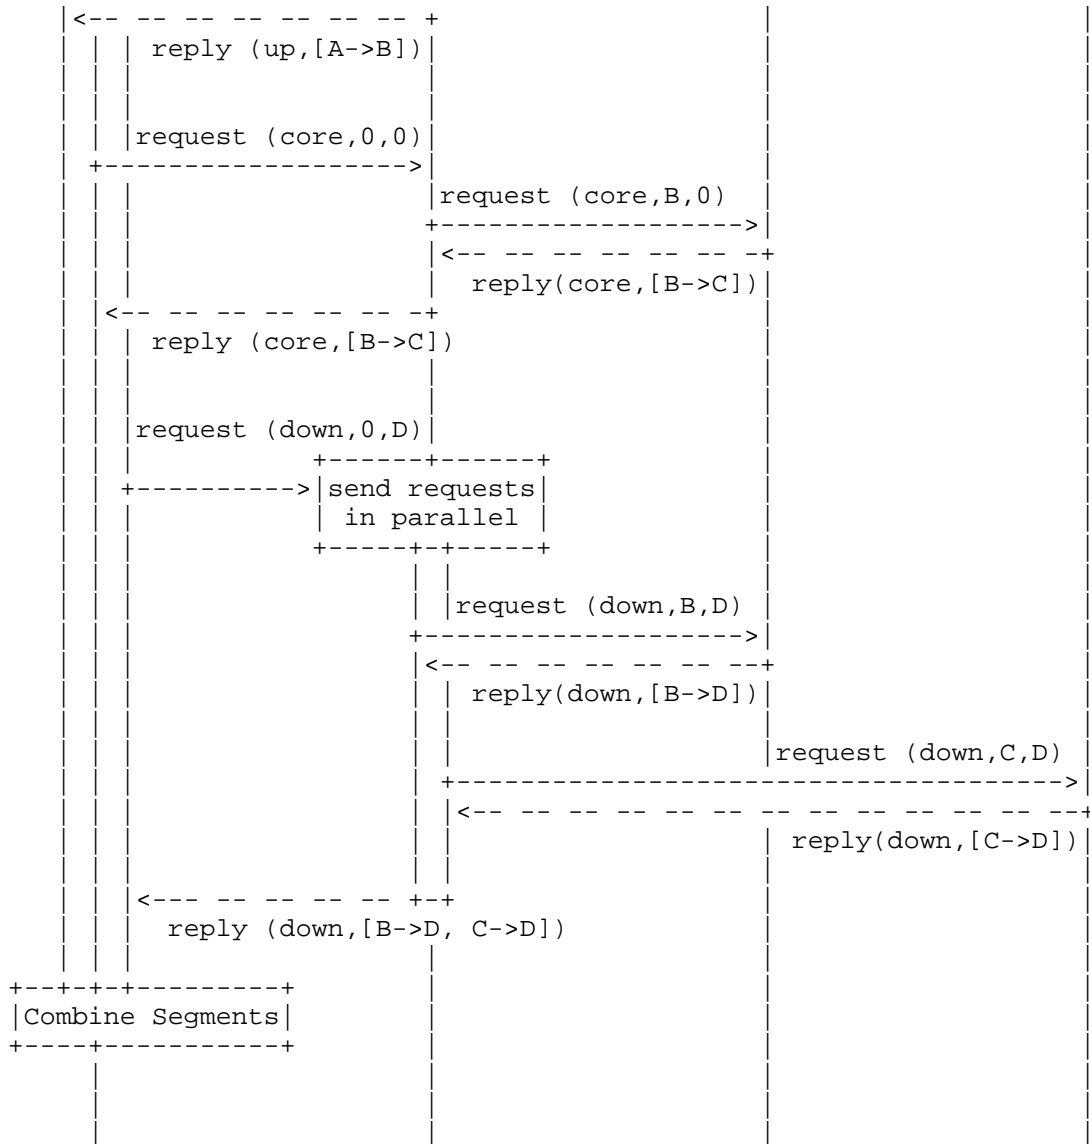


Figure 25: Sequence diagram illustrating a path lookup for a destination D in the source ISD. The request (core, 0, 0) is for all pairs of core ASes in the source ISD. Similarly, (down, 0, D) is for down segments between any core AS in the source ISD and destination D.

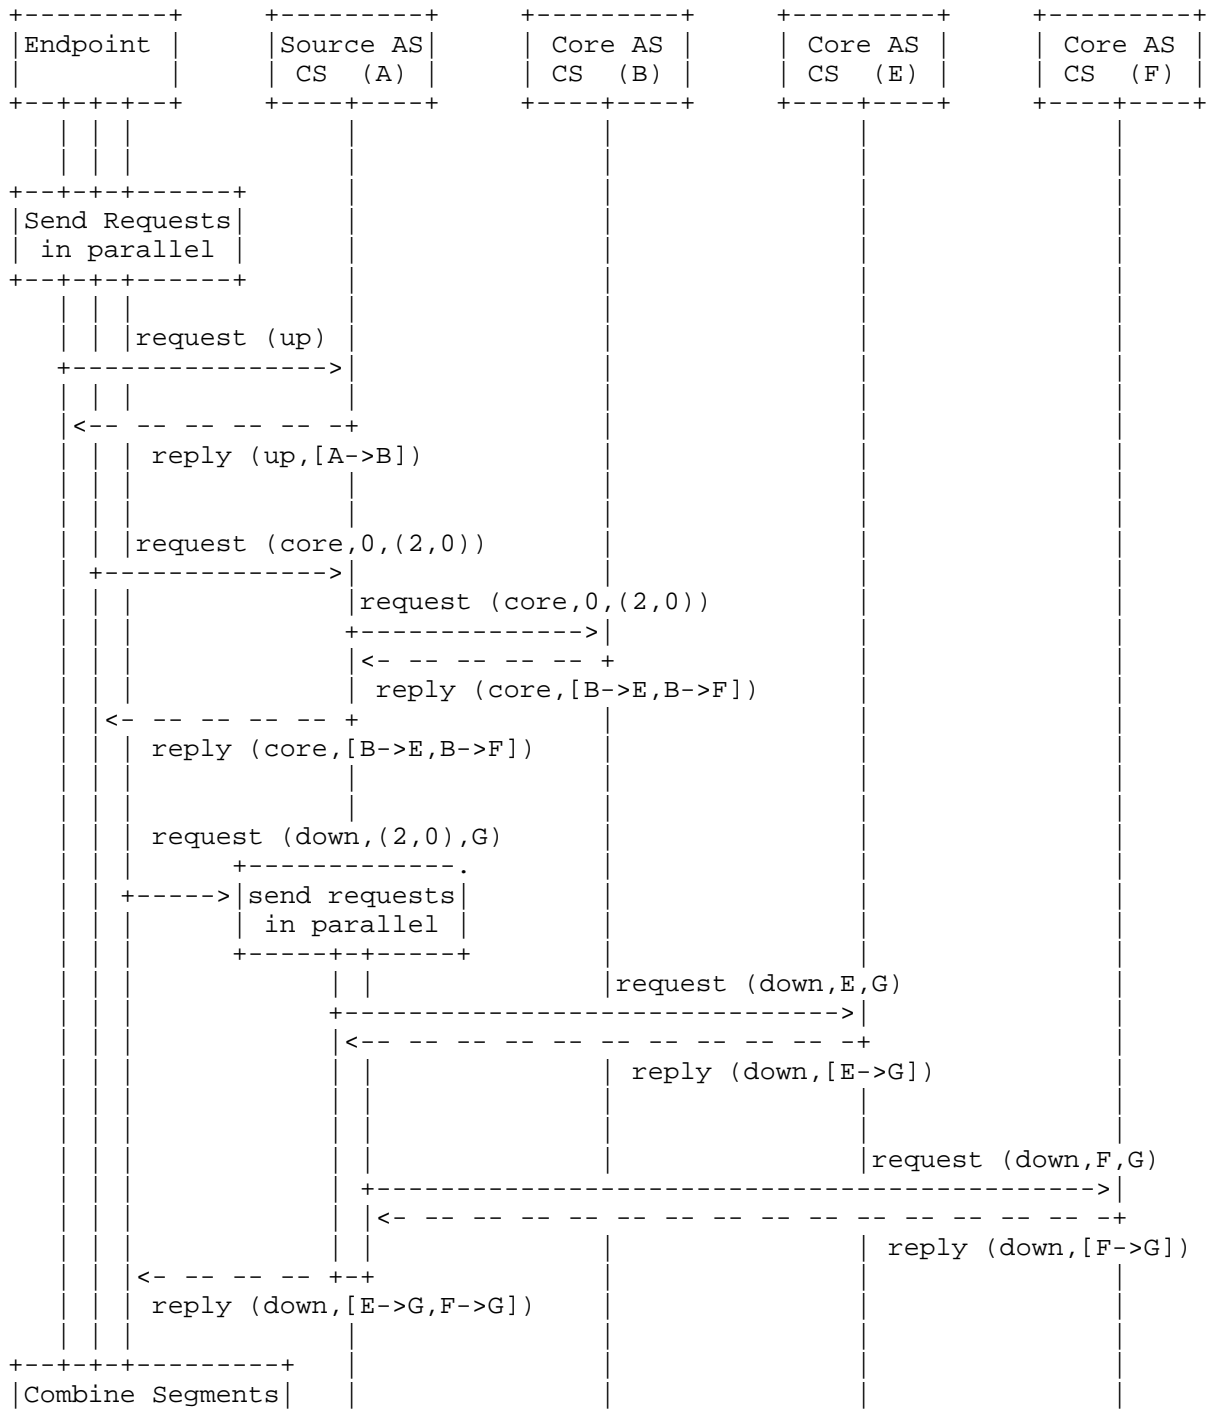




Figure 26: Sequence diagram illustrating a path lookup for a destination G in a remote ISD. The request (core, 0, (2, 0)) is for all path segments between a core AS in the source ISD and a core AS in ISD 2. Similarly, (down, (2, 0), G) is for down segments between any core AS in ISD 2 and destination G.

Change Log

Changes made to drafts since ISE submission. This section is to be removed before publication.

draft-dekater-scion-controlplane-15

- * Final read, wording
- * "originating/initiating" PCBs --> consistently use originating
- * 2.2. PCB Message Format: clarify order of as_entries
- * Section 2.3.5. Propagation of Selected PCBs: unify core and intra-ISD propagation, since steps are the same
- * Distribution of Cryptographic Material: clarify certificate encoding
- * SCMP: clarify relationship with RFC4443, clarify error processing rules and add informative reference to the SCION-UDP underlay
- * Appendix: remove SCIONLab

draft-dekater-scion-controlplane-15

- * Wording polish following ISE Editor's feedback.
- * Reduce use of passive tense and clarify subject (e.g., an AS --> An AS operator)
- * ISD and AS numbers: clarify that identifiers in public ranges must be unique
- * Remove redundant section 1.7. Resistance to partitioning

- * Section 1.7. Communication Protocol: Clarify DNS resolution is not needed
- * Figures 2, 3, 4: improve arrows in SVG version, move description text to after the figures
- * PCB Extensions: clarify behavior in case of unknown extensions
- * Configuration: ensure all items are mentioned and cross referenced.
- * Rename "registration period" to "registration interval" to ensure consistency with "propagation interval"
- * Timestamps: add normative reference to POSIX.1-2024 to clarify counting of leap seconds
- * Registration of Path Segments: clarify that a core AS has down segments registered by its direct or indirect customer ASes
- * Path Lookup Process: reformat and reword steps to clarify how an endpoint requests path segments
- * SCMP: remove experimental values from table and mention more error messages are in referenced spec
- * Move "Deployment Considerations" from section 3 to 7
- * Attacks on time sources: recommend use of secure time synchronization and mention CAs
- * Acknowledgements: ensure all reviewers are there

draft-dekater-scion-controlplane-14

- * Clarify bits in timestamps.
- * Remove informative reference to I-D.dekater-panrg-scion-overview and to Anapaya's ISD assignments, since they are taken over by SCION Association in 2026
- * Overall review and wording polish
- * Protobuf messages syntax check, add missing empty PathSegmentUnsignedExtensions message definition
- * SegmentLookupService RPC: clarify wording on API exposure

- * Peer entry figure 14 - make fields consistent with protobuf definitions

- * New section: Renewal of Cryptographic Material

draft-dekater-scion-controlplane-13

- * Clarify distinction between SCION ASes and BGP ASes through the text.

draft-dekater-scion-controlplane-12

- * Security considerations: new section "Attacks on time sources"
- * Path Lookup Process: mention checks at endpoint

draft-dekater-scion-controlplane-11

- * Clarify use of wildcard ISD and ISD-AS text representation
- * Remove redundant PCB overview figure 6 and reorganized paragraphs in 2.2. PCBs
- * Small clarifications and nits
- * Figures: small changes to use aasvg in all figures
- * Appendix "Path-Lookup Examples": use wildcard AS 0 instead of * in figures in

draft-dekater-scion-controlplane-10

Major changes:

- * New section "Distribution of Cryptographic Material" containing definitions formerly in the gRPC API appendix
- * New section "Destination Mapping" including a SIG reference
- * New section "Lookup Requests Message Format" containing definitions formerly in the gRPC API appendix
- * Move appendix "Use of the SCION Data Plane" to new section "Control Service Discovery"
- * Mention ConnectRPC as main RPC method instead of gRPC

- * Remove appendix "Full Control Service gRPC API" and move corresponding protobuf definitions in new sections mentioned above

Minor changes:

- * Rename Inter-ISD Beacons into Core Beacons for consistency
- * Clarify descriptions of fields in the HeaderAndBody message and that metadata must be empty
- * AS Entry Signature: fix order of terms in one formula, clarify validity and meaning of associated data
- * PCB Extensions: clarified text, added example of the StaticInfoExtension and informative reference
- * PCB Validity: clarify text on timestamp validity and time allowances
- * Reception of PCBs: mention that incoming link MUST be core or parent
- * PCB selection policies: discourage use for traffic engineering
- * Best PCBs Set Size: clarify tradeoffs and avoid normative language when unnecessary
- * Path reversibility: mention that destination endpoints should estimate MTU
- * Move considerations on SCMP Authentication to the security considerations section (Rogue SCMP Error Messages)
- * Security Properties: use normative language to clarify assumptions

draft-dekater-scion-controlplane-09

Major changes:

- * "SCION AS numbers": make text representation for lower 32-bit ASes consistent with PKI draft, add reference to allocation.

Minor changes:

- * Nits and wording improvements
- * Reviewed use of normative language

- * Figures: redraw and use aasvg when possible
- * "Paths and Links": clarify relationship between path segments and links
- * Ensure consistent use of example ranges

draft-dekater-scion-controlplane-08

Major changes:

- * Abstract: reword, mention goal and that document is not an Internet Standard
- * "Propagation of PCBs" section:
 - clarify checks at reception
 - introduce criteria for PCB selection policies
 - remove superfluous policy example figure
 - Propagation Interval and Best PCBs Set Size: mention tradeoff between scalability and amount of paths discovered.
 - reorganize order of paragraphs
- * New section "Security Properties" in Security considerations, based on formal model of SCION
- * New figure: Control Service RPC API - Trust Material definitions

Minor changes:

- * Moved "Special-Purpose SCION AS Numbers" table later in text
- * Split "Circular dependencies and partitioning" into two sections: "Bootstrapping ability" and "Resistance to partitioning".
- * Explain why PCBs have a next_isd_as field
- * Qualified better the choice of time allowance in the definition of segment from the future in section "PCB Validity".

draft-dekater-scion-controlplane-07

- * Moved SCMP specification from draft-dekater-scion-dataplane-03 to this document

draft-dekater-scion-controlplane-06

Major changes:

- * New section: Path MTU
- * New section: Monitoring Considerations
- * Completed description of Control Services RPC API in appendix

Minor changes:

- * Introduction: clarify goal of the document
- * Clarify typical vs recommended-limits values for best PCB set size and for certificate validity duration.
- * Clarify text representation of ISD-AS
- * General rewording
- * Added reference to SCIONLab as a testbed for implementors
- * Introduced this change log

draft-dekater-scion-controlplane-05

Minor changes:

- * Clarify beaconing fast retry at bootstrapping

draft-dekater-scion-controlplane-04

Major changes:

- * Clarified selection of MAC including a default algorithm
- * New section: PCB validity
- * New section: configuration
- * New section: Path Discovery Time and Scalability
- * New section: Effects of Clock Inaccuracy
- * New appendix: Control Service RPC API

Minor changes:

- * Introduction: Added overview of SCION components
- * Clarified path reversibility, link types, Interface IDs
- * Fixed private AS range typo
- * Clarified PCB selection policies and endpoint requirements
- * Clarified PCB propagation
- * General edits to make terminology consistent, remove duplication and rationalize text
- * Removed forward references
- * Added RFC2119 compliant terminology

Acknowledgments

Many thanks go to Alvaro Retana (Futurewei), Joel M. Halpern (Ericsson), Brian Trammel (Google), William Boye (Swiss National Bank), Matthias Frei (SCION Association), Kevin Meynell (SCION Association), Jean-Christophe Hugly (SCION Association), Juan A. Garcia Prado (ETH Zurich), Tilmann Zschke (ETH Zurich), Dominik Roos (Anapaya Systems), and Roger Lapuh (Extreme Networks) for reviewing this document. We also thank Daniel Galn Pascual and Christoph Sprenger from the Information Security Group at ETH Zurich for their inputs based on their formal verification work on SCION. We are also very grateful to Adrian Perrig (ETH Zurich), for providing guidance and feedback about every aspect of SCION. Finally, we are indebted to the SCION development teams of Anapaya and ETH Zurich for their practical knowledge and for the documentation about the SCION Control Plane, as well as to the authors of [CHUAT22] - the book is an important source of input and inspiration for this draft.

Authors' Addresses

Corine de Kater
SCION Association
Email: c_de_kater@gmx.ch

Nicola Rustignoli
SCION Association
Email: nic@scion.org

Samuel Hitz
Anapaya Systems
Email: hitz@anapaya.net