Routing security

Joeri de Ruiter

SIDN LABS

Operator of the .nl TLD

- Stichting Internet Domeinregistratie Nederland (SIDN)

- Critical infrastructure services
  - Lookup IP address of a domain name (almost every interaction)
  - Registration of all .nl domain names
  - Manage fault-tolerant and distributed infrastructure

- Increase the value of the Internet in the Netherlands and elsewhere
  - Enable safe and novel use of the Internet
  - Improve the security and resilience of the Internet itself

SIDN LABS

SIDN Labs

- Goal: advance operational Internet security and resilience through world-class measurement-based research and technology development

- Research challenges: core Internet systems and Internet evolution

- Daily work: help operational teams, write open source software, analyze vast amounts of data, run experiments, write academic papers, work with universities

SIDN LABS

Today's topics

- BGP
  - RPKI
  - BGPsec
- Starting from scratch: SCION

Autonomous systems

- The internet is a combination of networks
- These network are called autonomous systems (AS)
  - Controlled by a single entity
  - One or more IP prefixes
  - Identified by a unique number (ASN)
- ASes communicate routing information to their neighbours (peers)
  - Which IP prefixes can be reached through them

SIDN LABS

Border Gateway Protocol (BGP)

- BGP-4, RFC 4271
- Protocol to communicate routing information between ASes
- Announcements
  - Prefix, AS path, next hop
- Glues the Internet together
- Border routers contain forwarding tables specifying where to forward packets to depending on the prefix (using longest prefix match)

SIDN LABS

AS1103 SURFnet bv

Company Website: http://www.surf.nl/en

Country of Origin: Netherlands

Internet Exchanges: 5

Prefixes Originated (all): 97
Prefixes Originated (v4): 94
Prefixes Originated (v6): 3

Prefixes Announced (all): 214
Prefixes Announced (v4): 190
Prefixes Announced (v6): 24

BGP Peers Observed (all): 1,133
BGP Peers Observed (v4): 1,111
BGP Peers Observed (v6): 781

IPs Originated (v4): 6,194,944
AS Paths Observed (v4): 96,741
AS Paths Observed (v6): 20,522

Average AS Path Length (all): 4.225
Average AS Path Length (v4): 4.297
Average AS Path Length (v6): 3.885

BGP example

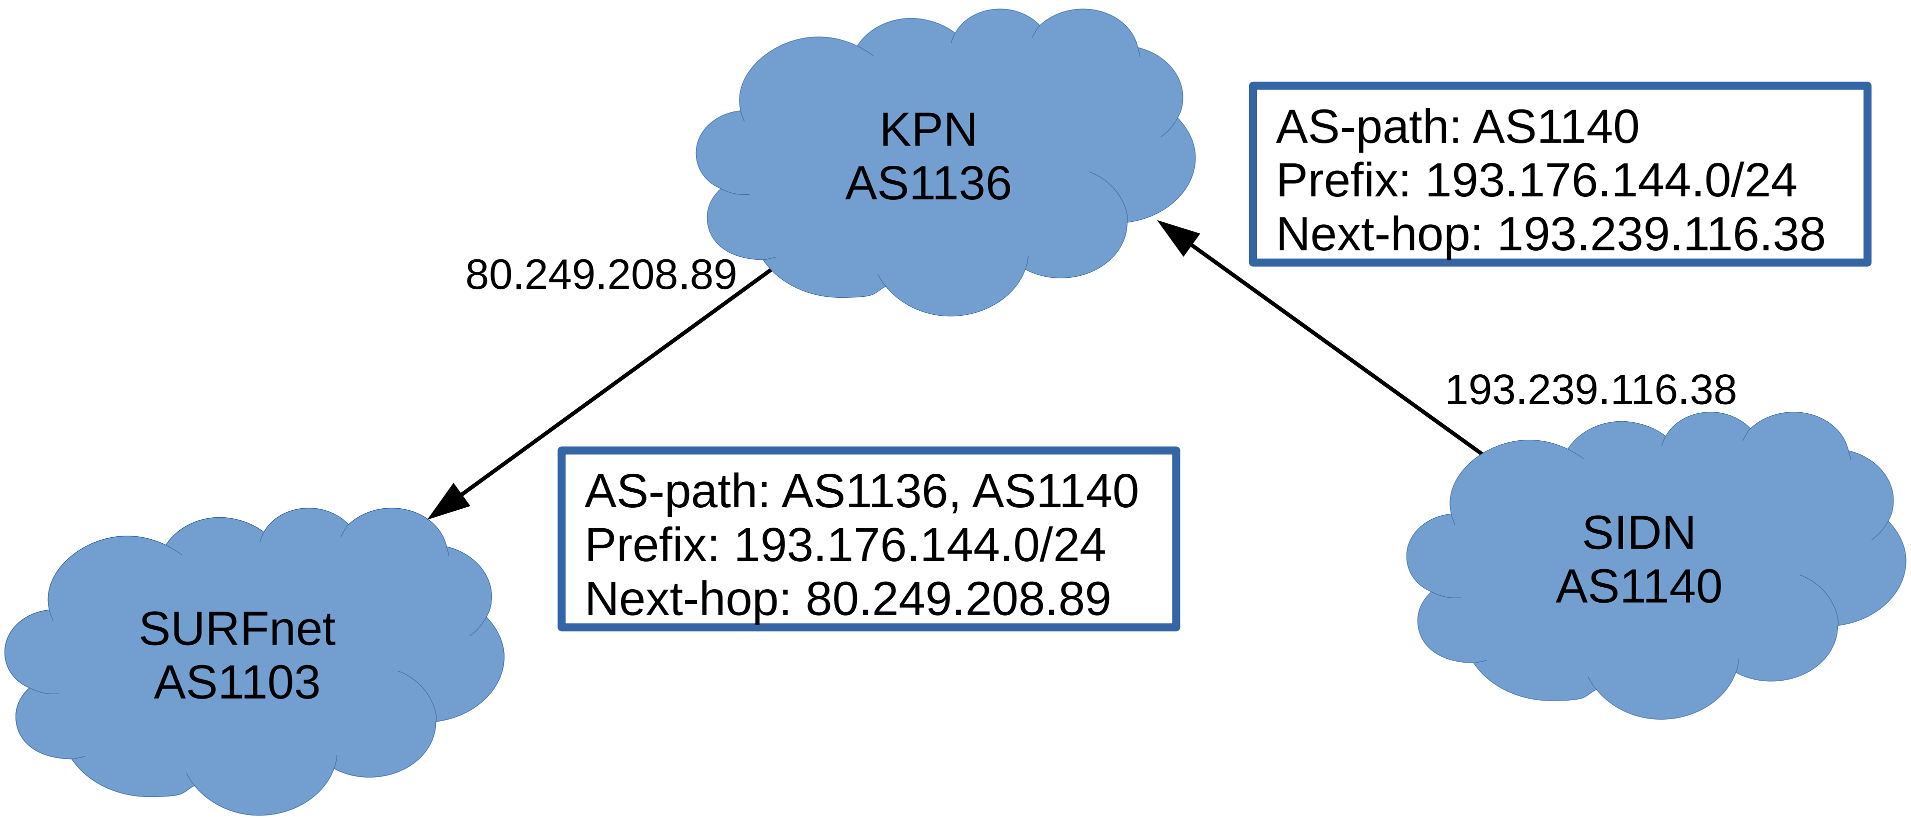
KPN
AS1136
AS-path: AS1140
Prefix: 193.176.144.0/24
Next-hop: 193.239.116.38
80.249.208.89
193.239.116.38
AS-path: AS1136, AS1140
Prefix: 193.176.144.0/24
Next-hop: 80.249.208.89
SURFnet
AS1103
SIDN
AS1140

BGP example

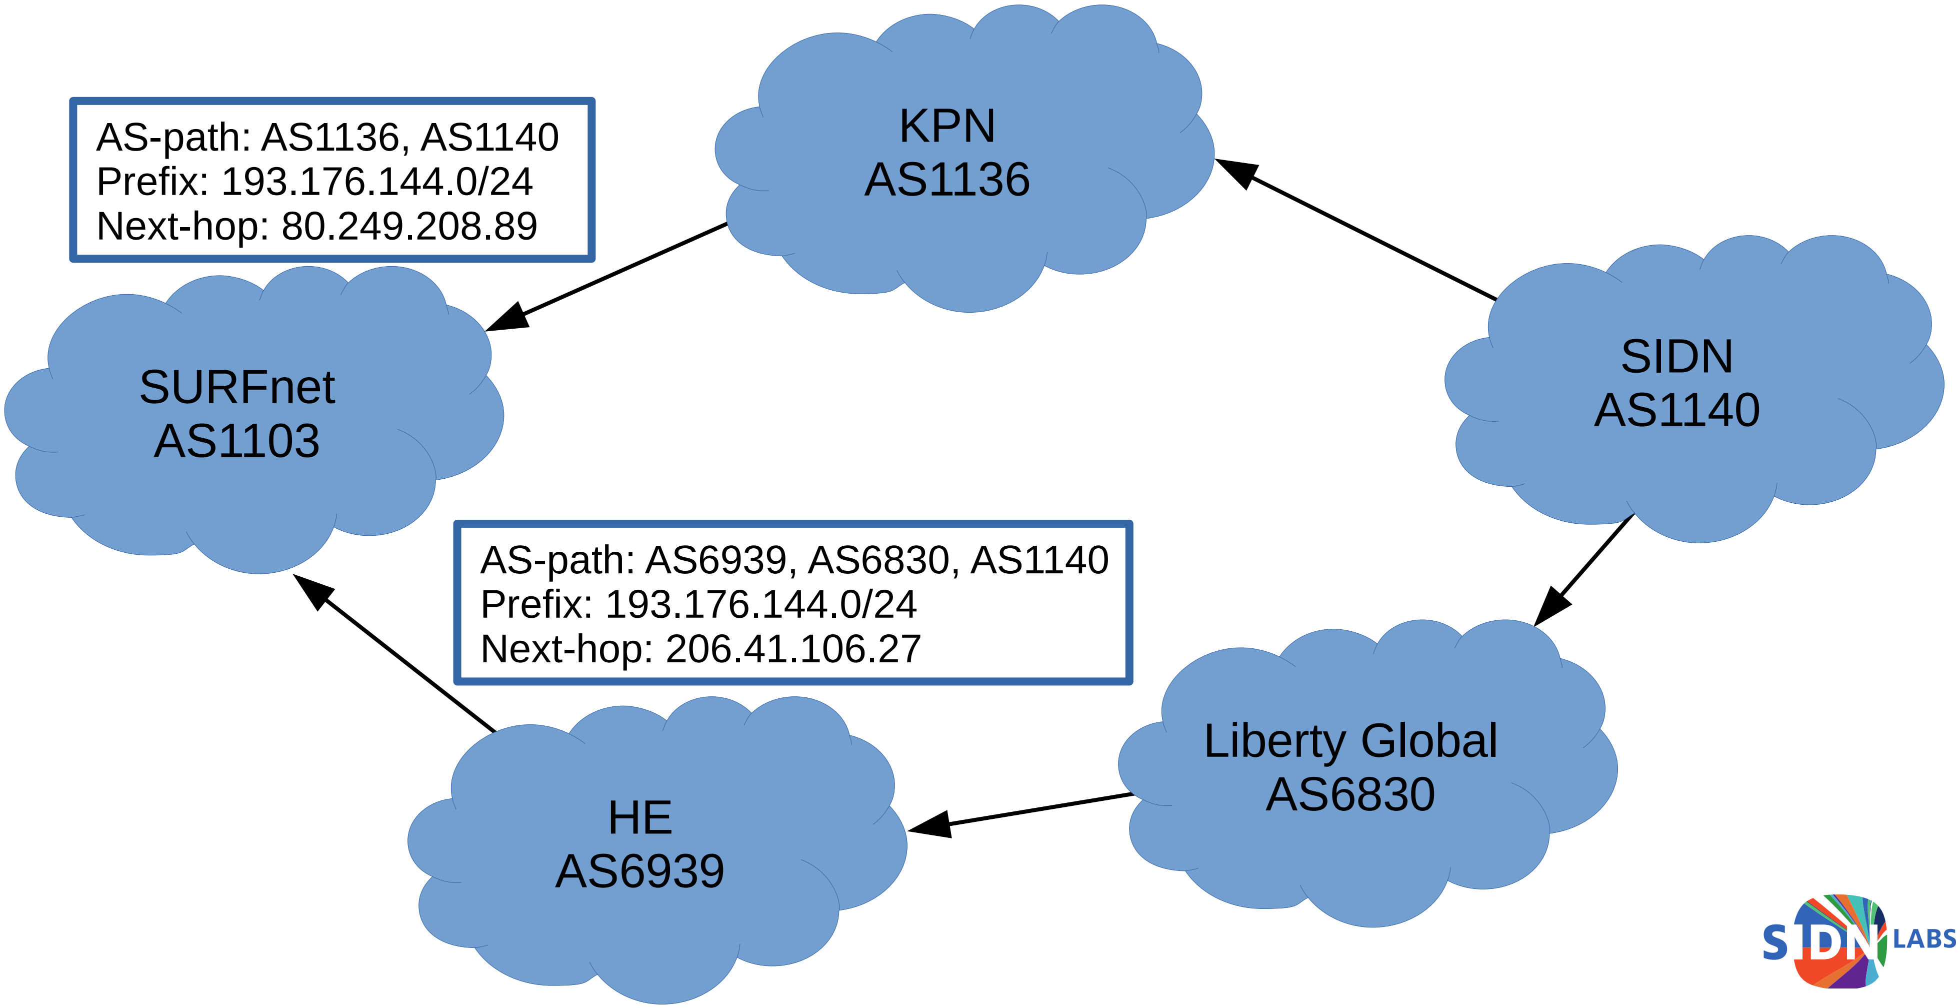
AS-path: AS1136, AS1140
Prefix: 193.176.144.0/24
Next-hop: 80.249.208.89
KPN
AS1136
SURFnet
AS1103
SIDN
AS1140
AS-path: AS6939, AS6830, AS1140
Prefix: 193.176.144.0/24
Next-hop: 206.41.106.27
Liberty Global
AS6830
HE
AS6939

BGP security

- Plaintext and unauthenticated
- Hijacking or interception of prefixes
  - Announce longer prefix or shorter path

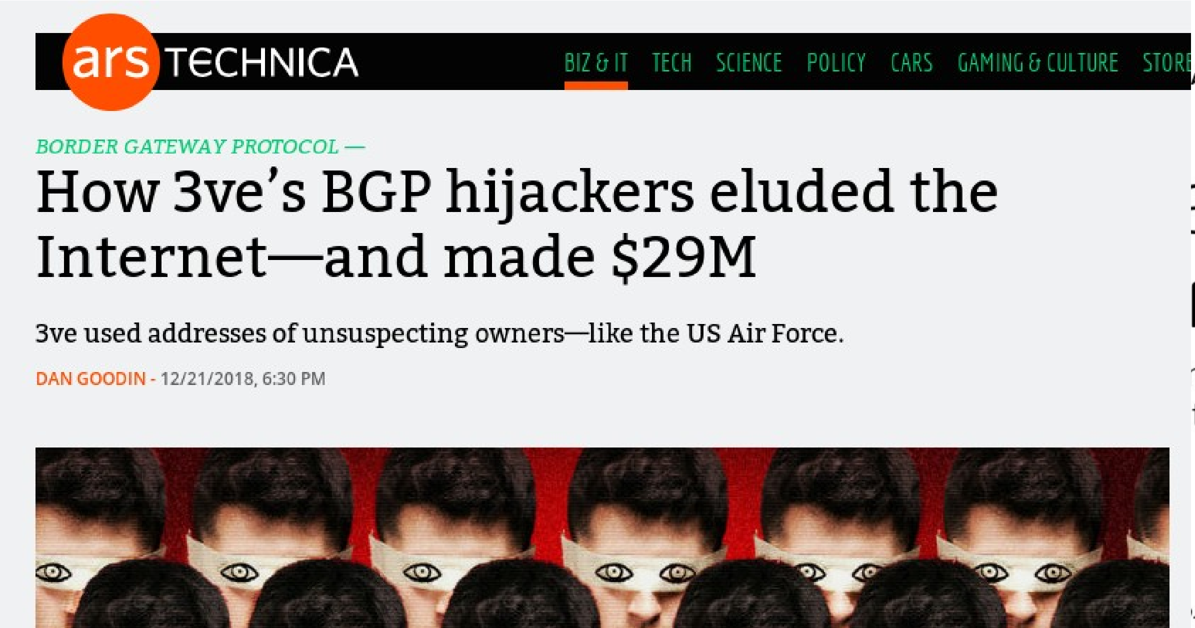
ars TECHNICA

BIZ & IT TECH SCIENCE POLICY CARS GAMING & CULTURE STORE

BORDER GATEWAY PROTOCOL —

How 3ve's BGP hijackers eluded the Internet—and made $29M

3ve used addresses of unsuspecting owners—like the US Air Force.

DAN GOODIN - 12/21/2018, 6:30 PM


ALS DOWNLOAD

ames Pakistan
r 2-hour outage

m reports that Pakistan Telecom was responsible
to erroneous Internet Protocols.

.m. to add YouTube's

Routing security

- What properties do we want?
- Origin authentication
  - You can only announce prefixes that are assigned to you
- Path authentication
  - The complete path to the origin is verifiable

Resource PKI (RPKI)

- Provides origin authentication using certificates to assign prefixes
- Deployment started in 2011 and described in RFC 6480
- Makes use of existing standards
  - E.g. X.509 certificates, extended with attributes to include IP prefixes
- Root CAs called Trust Anchor
- Leaf certificates called End-Entity Certificates
- Route Origin Authorization (ROA)
  - Bind prefix to AS
  - Signed by owner of the prefix
- One-to-one mapping between End-Entity Certificate and ROA

SIDN LABS

RPKI hierarchy

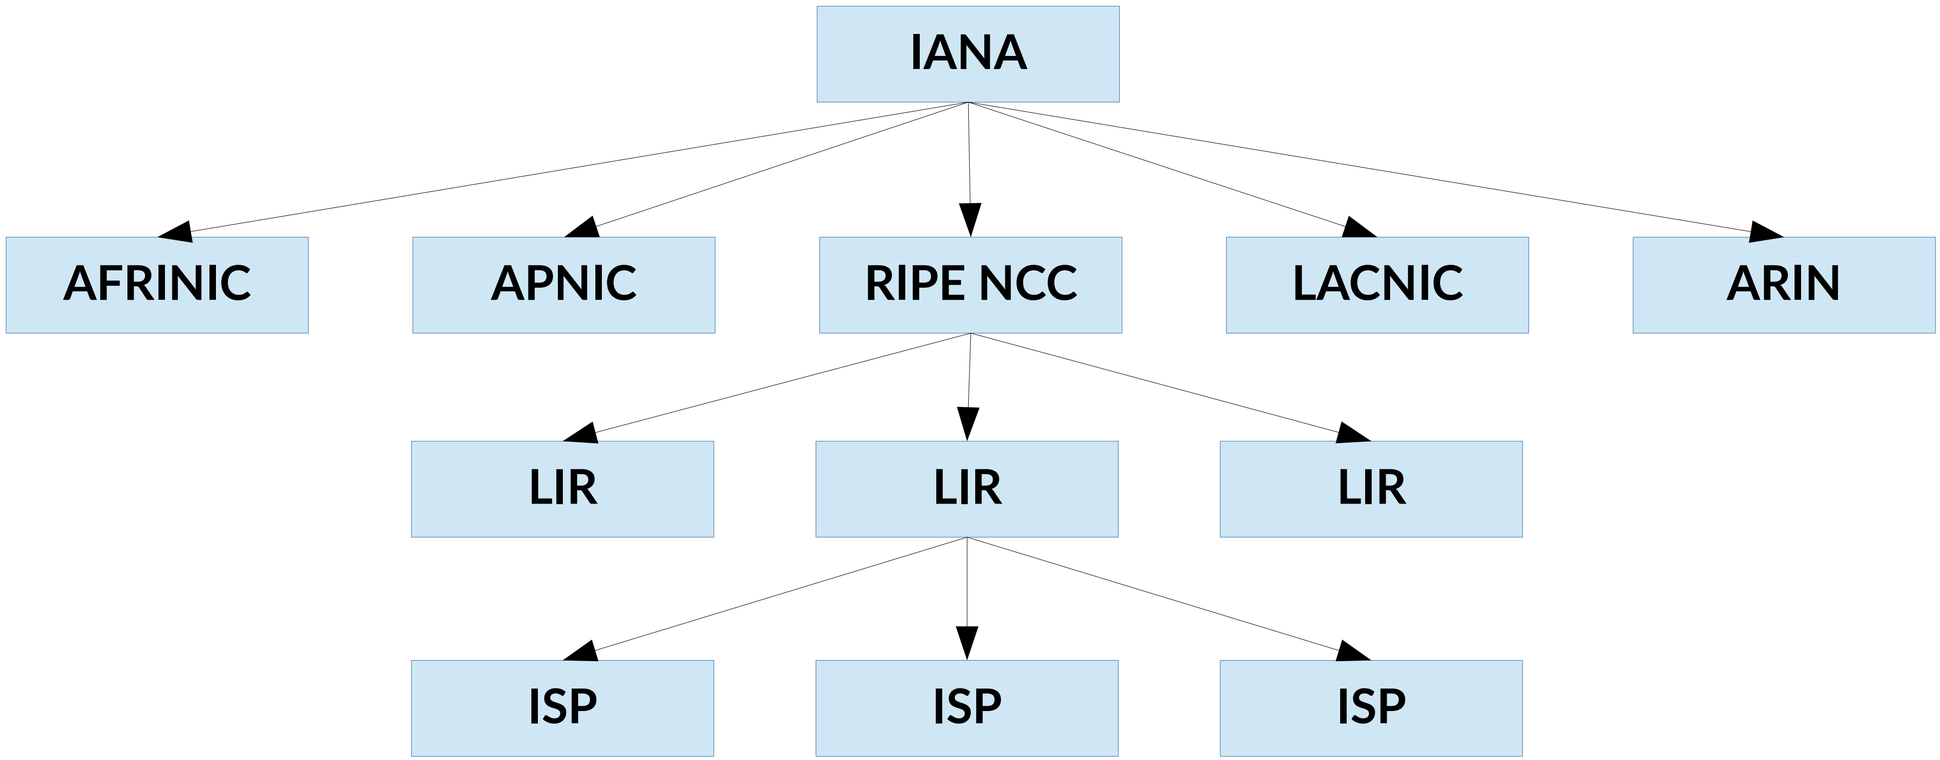

RPKI adoption – Europe

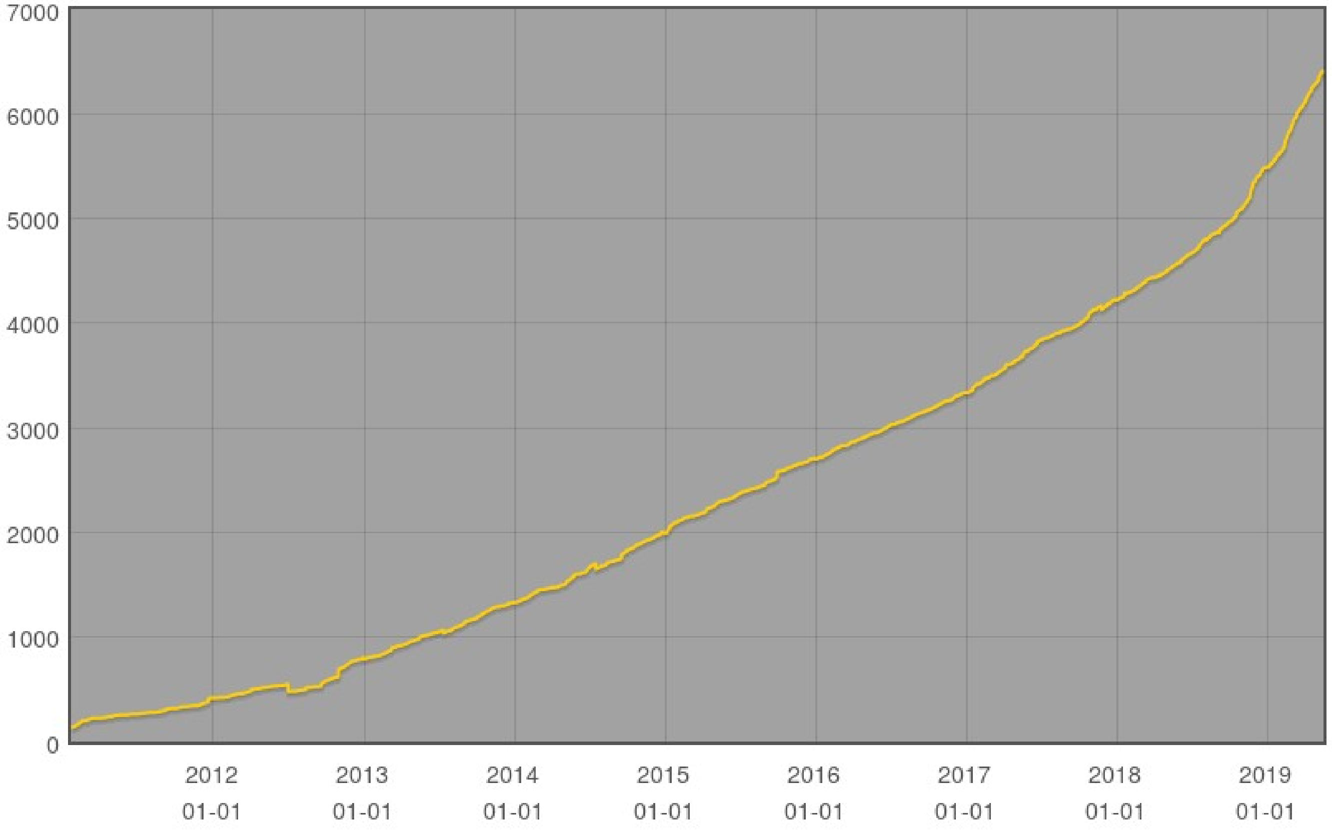

Unique ASNs in ROAs for RIPE NCC
Source: https://certification-stats.ripe.net/

Origin authentication

- Described in RFC 6493
- Cryptographic verification performed by RPKI Cache (local or at service provider)
  - Download records from repository (e.g. RIRs such as RIPE)
  - Verify chain, including assigned resources
  - Assigned resources should be a subset of the parent's resources
- Verification against BGP announcement performed by routers
  - Router retrieves stripped ROAs from RPKI Cache
  - Match BGP announcements against published ROAs
    - Valid / Invalid / NotFound
  - Verification results used in policy

BGP example

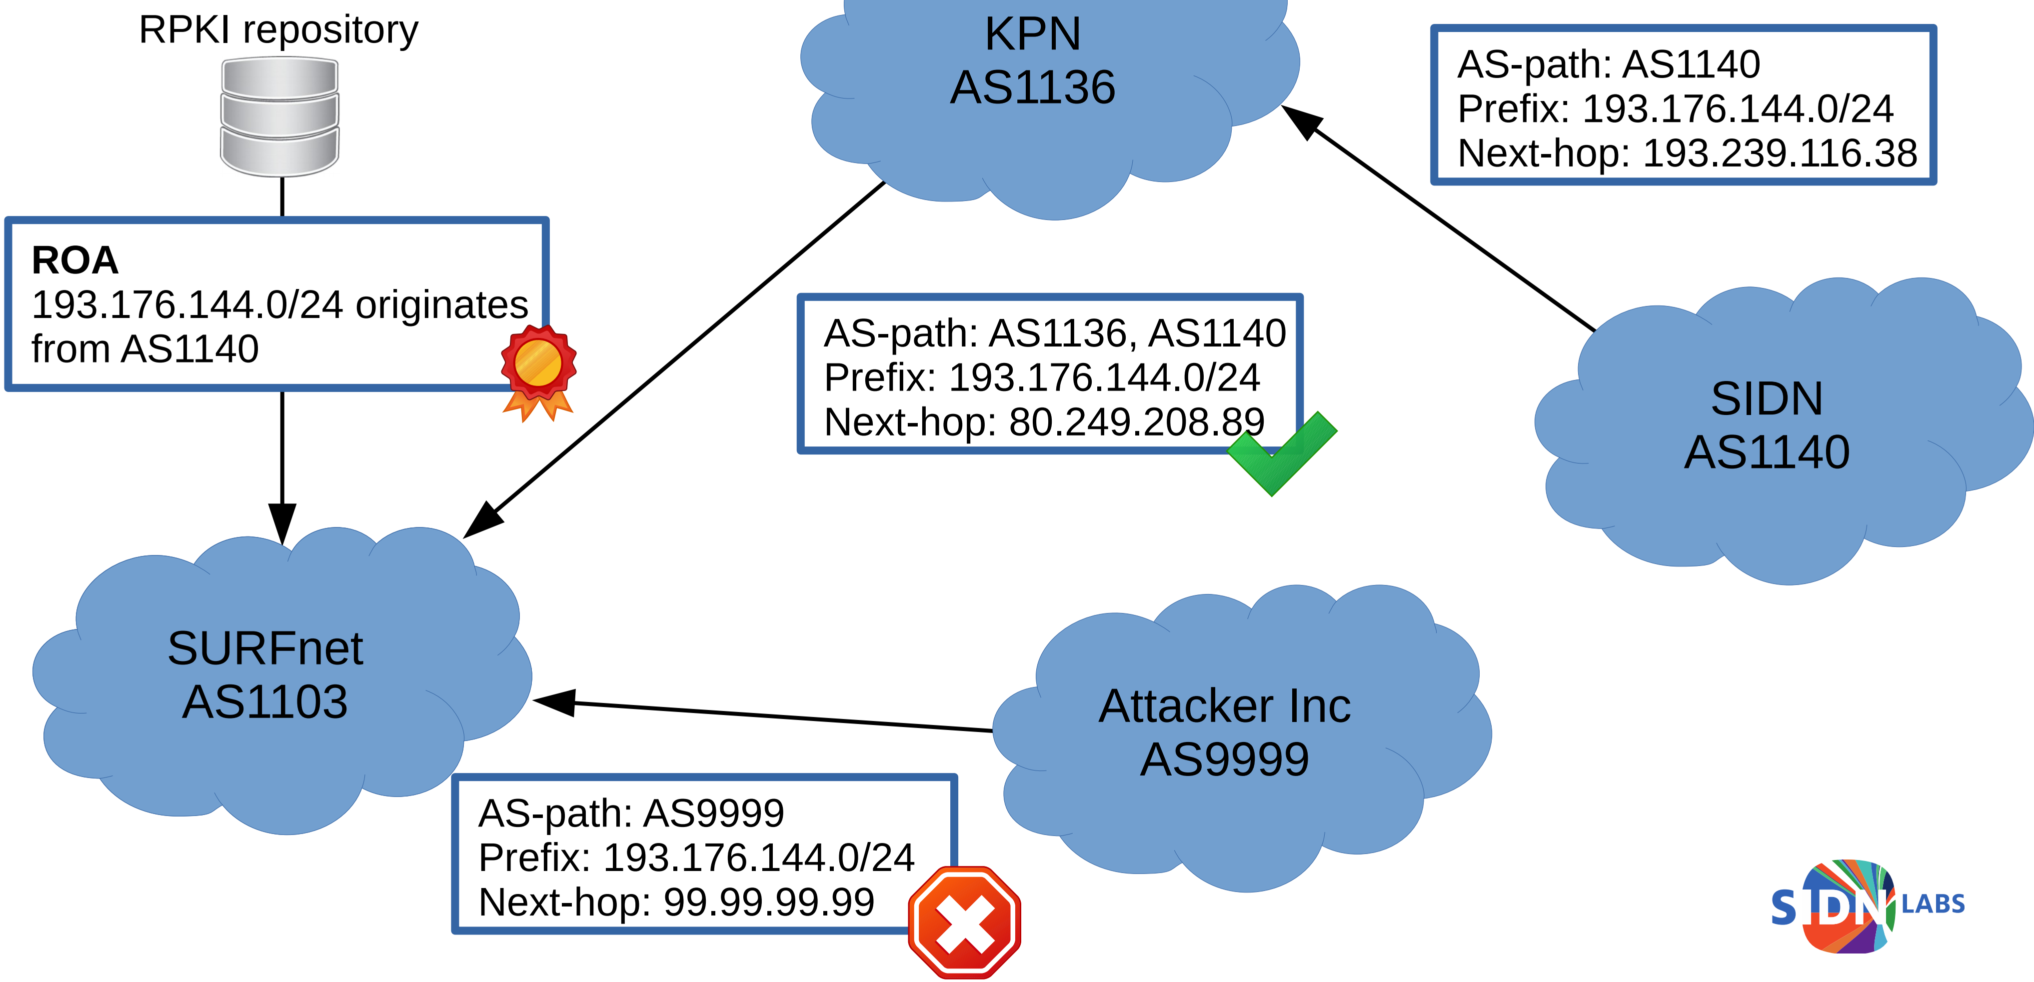

RPKI repository
ROA
193.176.144.0/24 originates from AS1140
KPN
AS1136
AS-path: AS1140
Prefix: 193.176.144.0/24
Next-hop: 193.239.116.38
AS-path: AS1136, AS1140
Prefix: 193.176.144.0/24
Next-hop: 80.249.208.89
SIDN
AS1140
SURFnet
AS1103
Attacker Inc
AS9999
AS-path: AS9999
Prefix: 193.176.144.0/24
Next-hop: 99.99.99.99

Path authentication

- BGPsec: verification of complete path in announcement
  - RFC 8205
- Uses RPKI
- AS-Path authenticated using signature in BGPsec-Path
- Every AS adds signature over previous signature and newly added path information
  - Including next AS

BGP example

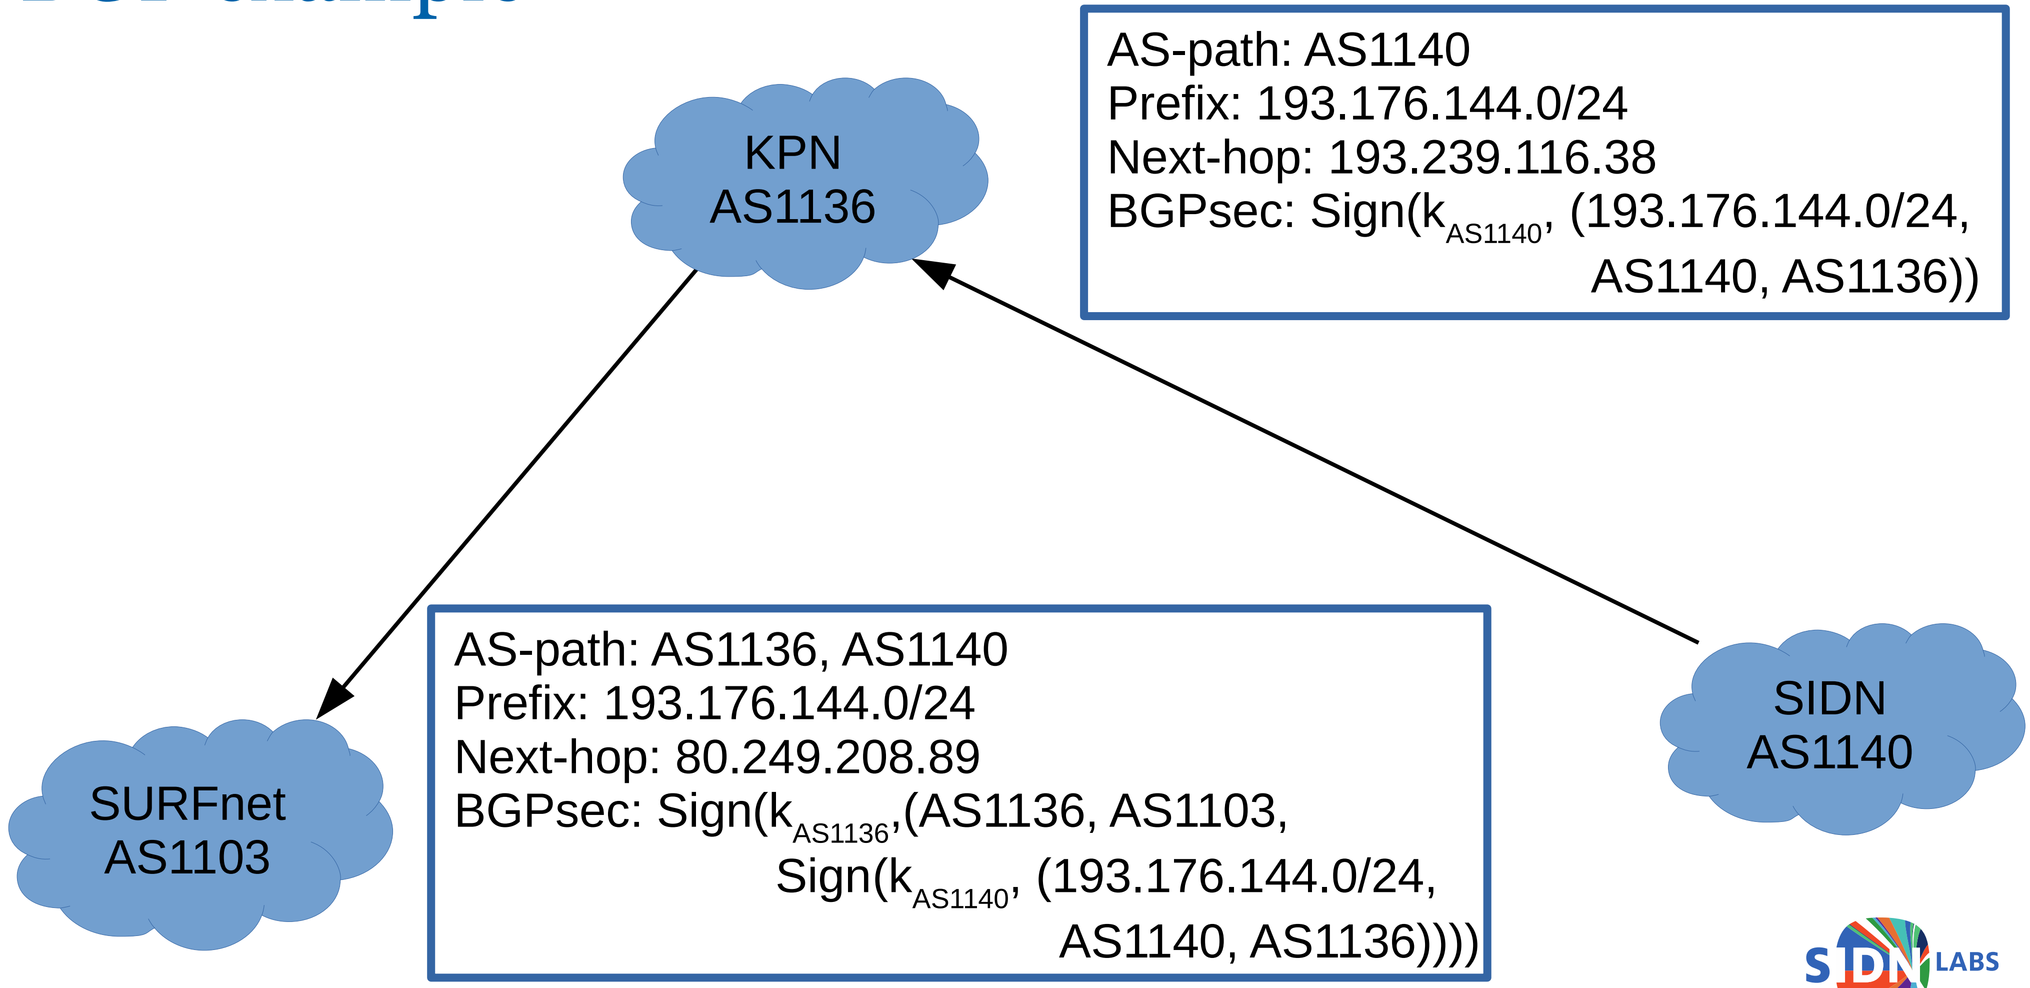

SIDN LABS

Starting from scratch

- Current Internet is a combination of patches
- Security is merely an afterthought
- Can we do better if we start (almost) from scratch?
- Scalability, Control, and Isolation On Next-generation Networks

SCION

SIDN LABS

SCION

- New internet architecture
- Research at ETH Zürich
- Scalability and security through Isolation Domains (ISDs)
  - Group of autonomous systems
  - E.g. per country or jurisdiction
- Routes authenticated both in control and data plane

SCION – Isolation Domains

- PKI organised per ISD
- ISD core: ASes managing the ISD
- Core AS: AS part of the ISD core
- Hierarchical control plane
  - Inter-ISD control plane
  - Intra-ISD control plane

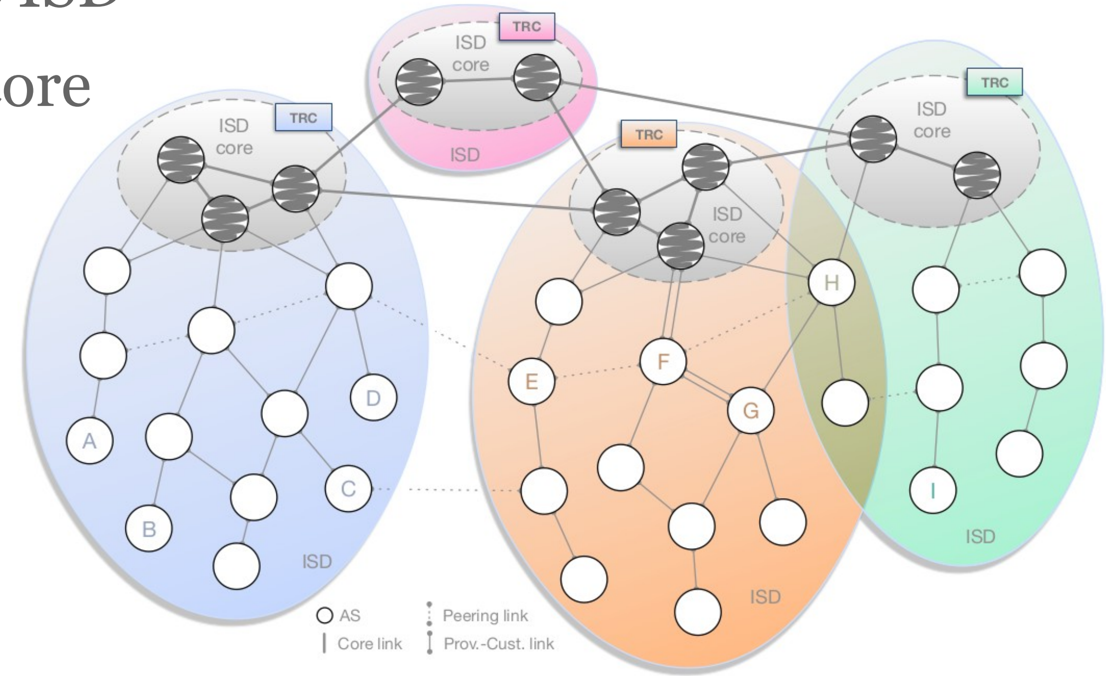

Source: The SCION Internet Architecture: An Internet Architecture for the 21st Century, Barrera et al., 2017

SIDN LABS

SCION – Autonomous systems

- Every interface that connects to neighbouring AS is assigned a unique identifier
- Several services run within AS
  - Beacon server
  - Path server
  - Certificate server

SCION – Path discovery

- Inter-ISD
  - Performed by core ASes
  - PCBs flooded similar as with BGP
  - Less ASes involved (only core)
- Intra-ISD
  - Downstream multi-path flooding

SCION – Intra-ISD path discovery

- Path Construction Beacons (PCBs) sent downstream using multi-path flooding
  - Initialised by core nodes
  - Extended and forwarded by receiving ASes
  - Add incoming and outgoing interface and optional peerings
- Eventually all nodes know how ISD core can be reached
- AS registers preferred down-segments (path from core to AS) with path server in the core
- Preferred up-segments registered with local path server

SIDN LABS

SCION – Intra-ISD path discovery

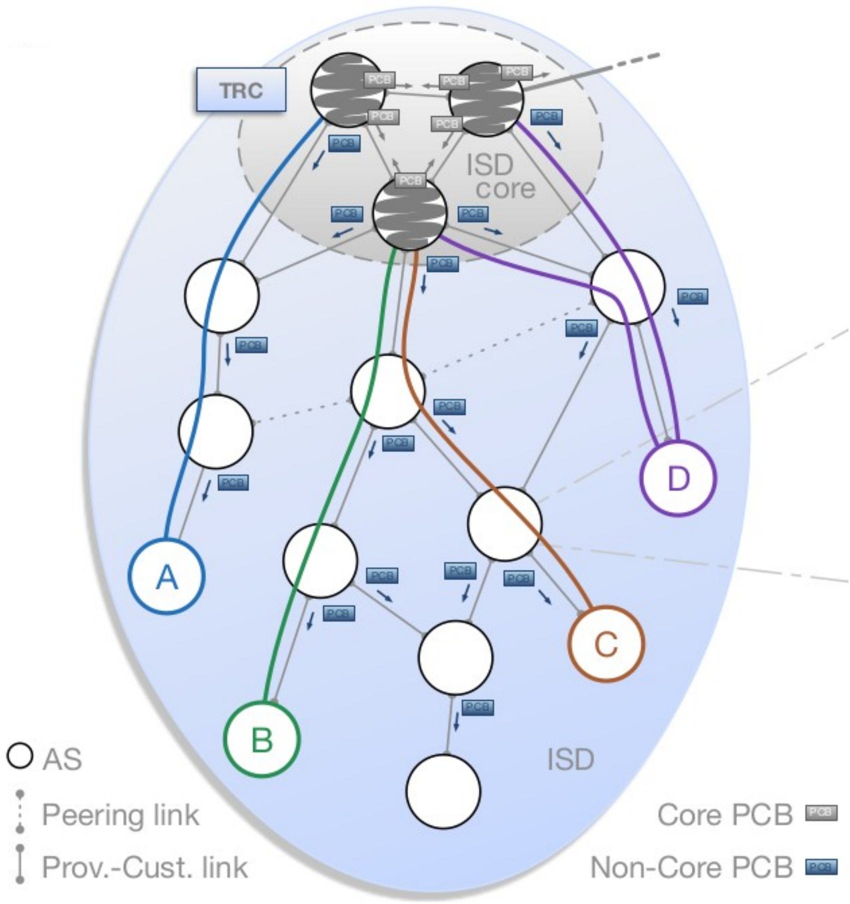

Source: The SCION Internet Architecture: An Internet Architecture for the 21st Century, Barrera et al., 2017

SCION – Intra-ISD path discovery

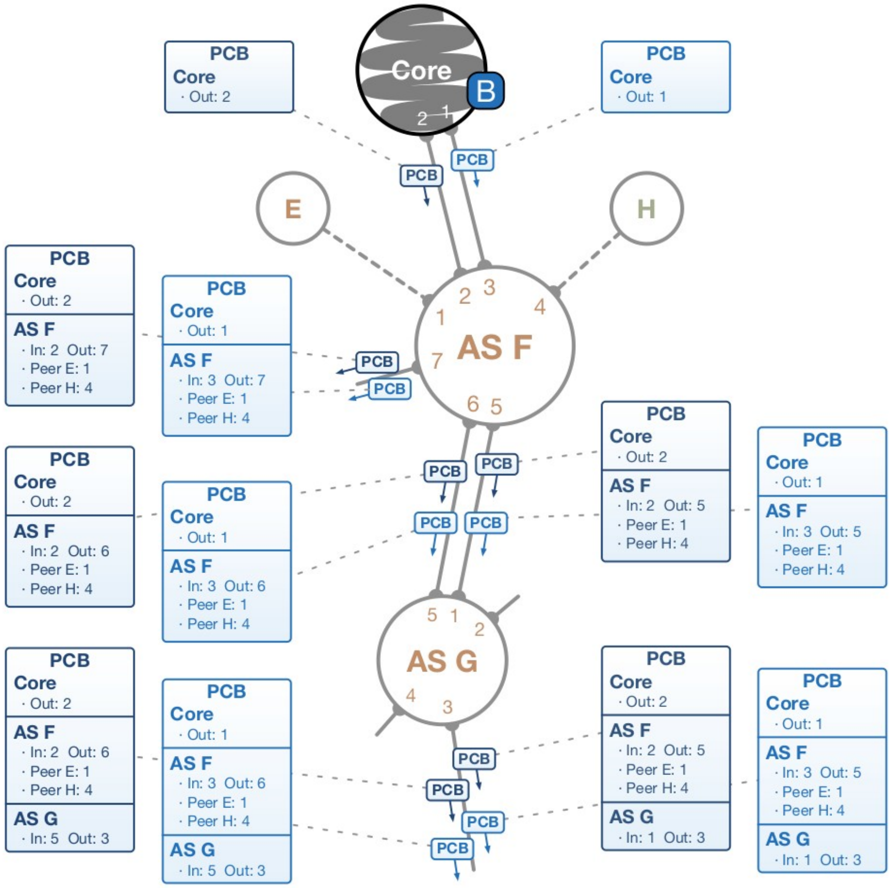

Source: The SCION Internet Architecture: An Internet Architecture for the 21st Century, Barrera et al., 2017

SIDN LABS

SCION – Path Construction Beacons

- Path Construction Beacons are signed by every AS along the path
  - Can be verified within ISD
- Hop-fields (HF) included that can be used to later select paths
  - Contain MAC computed using hop-field key
  - Only processed locally

SCION – Path Construction Beacons

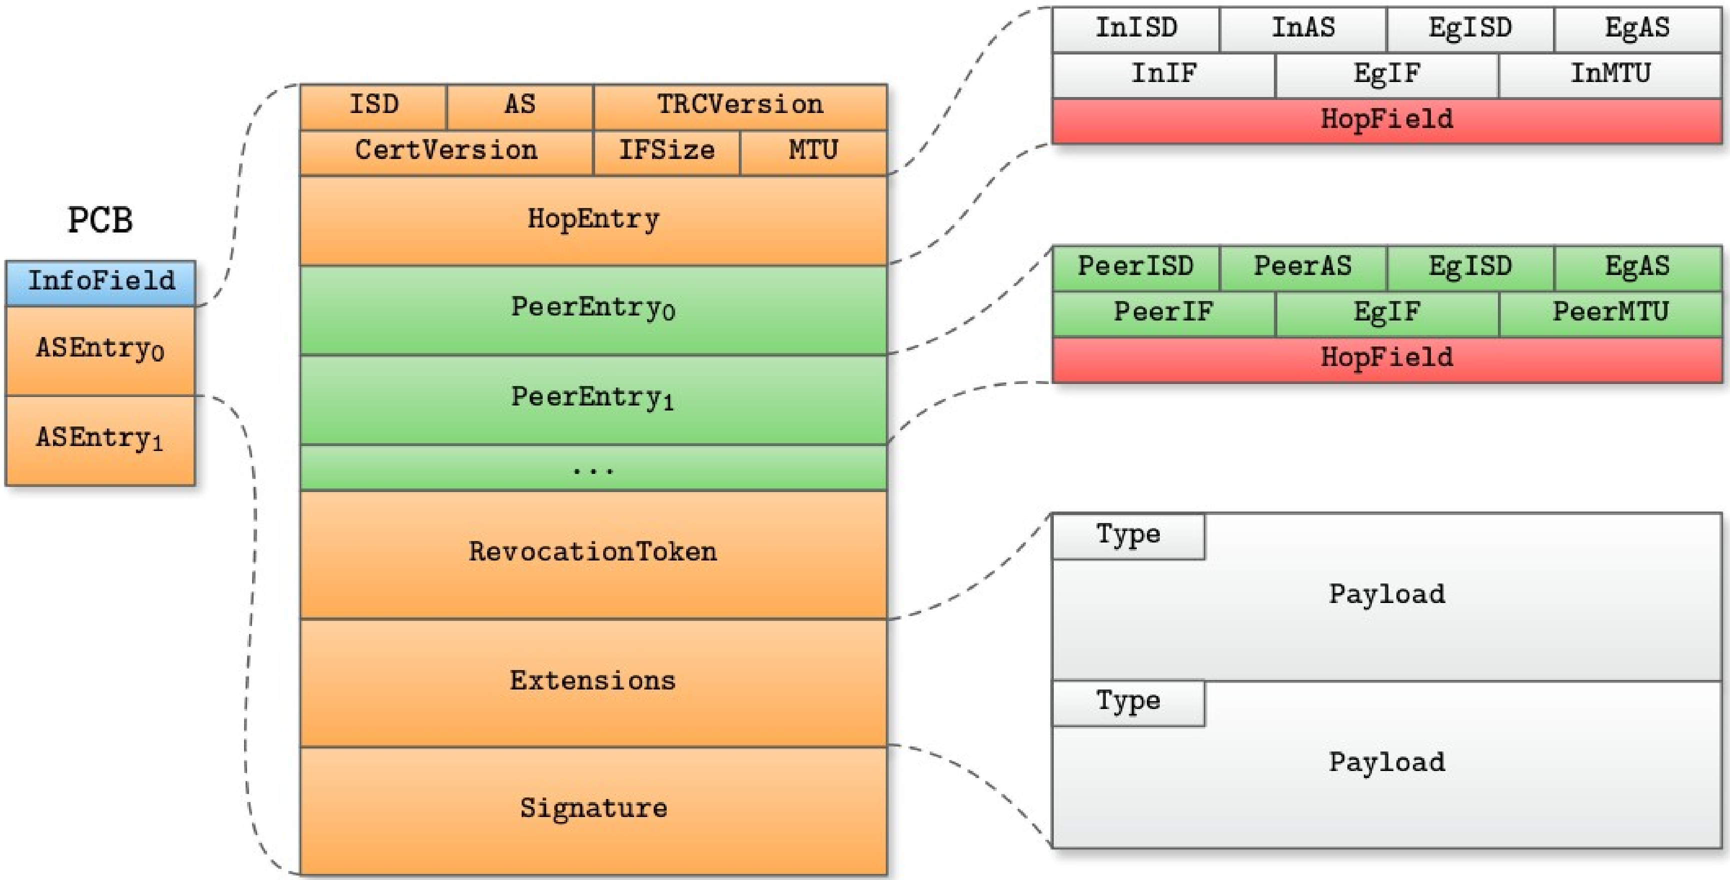

Source: SCION: A Secure Internet Architecture, Perrig et al., 2017

SCION – Path lookup

- Path construction performed by end hosts
- Request route to (ISD, AS) from local path server
- Local path server replies with
  - Up-path segments to local ISD core
  - Down-path segment in remote ISD from core to destination AS
  - Core-path segments needed to connect up-path and down-path segments
- End hosts combines segments to determine path

SCION – Path lookup

- Path server caches path segments
- If path to AS in remote ISD is not present in cache:
  - Request core- and down-path segments from local core AS
  - Core AS requests down-path segments from core AS in remote ISD
  - Up-, core- and down-segments returned to end host

SCION - Routing

- Path information included in packet headers
  - Corresponding hop-field included
  - No forwarding information necessary at routers
  - Packet-carried forwarding state (PCFS)
- Sender selects the path
  - Possible to use multiple paths
- Recipient address no longer used to route between autonomous systems
  - Only used by the destination AS

SIDN LABS

SCION - Routing

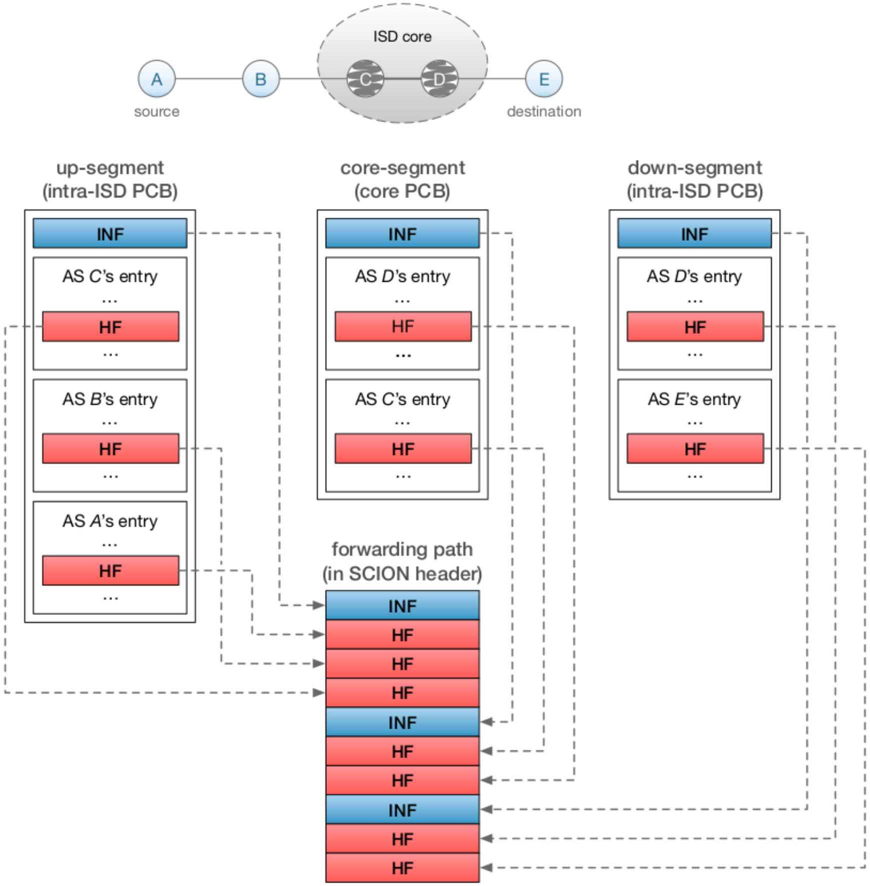

Source: SCION: A Secure Internet Architecture, Perrig et al., 2017

SCION - Security

- Trust within ISD
  - Compromise is kept locally → root key can only be used to compute certificates for local ISD
- Authenticated paths
  - Authentication in data plane
  - No path hijacking
  - No spoofing → no reflection attacks

SCION - PKI

- Control-plane
  - Comparable to RPKI
  - Short-lived certificates for ASes
- Name-resolution
  - Comparable to DNSSEC
  - Typically ISD will delegate name resolution to TLDs
- End-entity
  - Comparable to TLS
  - Certificates need to be signed by multiple CAs and registered at publicly verifiable log server

SIDN LABS

SCION – Source and path validation

- So far no validation that data was not injected and actually followed the desired path
- Extensions to SCION to achieve this:
  - OriginValidation, packet originates from source
  - PathTrace, packet followed indicated trace
  - Origin and Path Trace (OPT)

SCION - OriginValidation

- Source shares a symmetric key with every AS on the path
- Additional information in header
  - DataHash: hash over payload
  - SessionID: session identifier picked by source
  - List of OV values: MAC over DataHash with key shared between source and AS or destination
- Every intermediate AS and the destination verify its corresponding OV value
  - Overhead linear in number of ASes on the path

SCION - OriginValidation

DataHash = Hash(payload)
SessionID
OV_1 = MAC($K_{S,AS1}$, DataHash)
OV_2 = MAC($K_{S,AS2}$, DataHash)
...
OV_D = MAC($K_{S,D}$, DataHash)

SCION - PathTrace

- Source and destination share a symmetric key with every AS on the path
- Additional information in header
  - DataHash: hash over payload
  - SessionID: session identifier picked by source
  - Path Validation Field (PVF): MAC over DataHash and previous value of PVF
- Every intermediate AS updates the PVF value
  - Overhead constant
- Destination can compute MAC over data hash and final PVF for source to verify path
- Verification can be performed later: retroactive-PathTrace

SIDN LABS

SCION - PathTrace

DataHash = Hash(payload)
SessionID
PVF = MAC(K_S, DataHash)

SCION - PathTrace

DataHash = Hash(payload)
SessionID
PVF = MAC(K_{AS1}, DataHash | MAC(K_S, DataHash))

SIDN LABS

SCION in practice

- Open source implementation available
- Can be combined with existing Internet (e.g. through gateways)
- SCIONLab: international research network
  - Open for everyone to connect to
- Used in practice by banks, government and hospitals
- At SIDN
  - Permanent infrastructure node (AS) connected to SCIONLab
  - Implementation of SCION on open networking hardware

Summary

- BGP provides no secure by default
  - Hijacking and interception possible
- Origin authentication provided by RPKI and ROAs
- BGPsec introduces path authentication
- SCION introduces a new architecture that provides security by design
  - E.g. authenticated routing in data plane

SIDN LABS

Follow us

SIDN.nl

@SIDN

SIDN

Thanks for your attention!

joeri.deruiter@sidn.nl
www.sidnlabs.nl