2STIC

Experimenting with the SCION Internet architecture

Caspar Schutijser, Ralph Koning (SIDN Labs)

Advanced Networking Guest Lecture, UvA/OS3

March 7, 2023

SIDN: 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

.nl = the Netherlands
17M inhabitants
6.2M domain names
3.4M DNSSEC-signed
2.5B DNS queries/day
8.6B NTP queries/day

SIDN fonds

SIDN LABS

SIDN Labs: research team of SIDN

- Goal: increase the trustworthiness (security, stability, resilience, and transparency) of our society's internet infrastructure, for .nl and the Netherlands in particular

- Strategies:
  - Applied technical research (measurements, design, prototyping, evaluation)
  - Make results publicly available and useful for various target groups
  - Work with universities, infrastructure operators, and other labs

- Three research areas: domain name security, infrastructure security (DNS, NTP, BGP), emerging internet technologies

SIDN Labs: Team

Cristian Hesselman

Directeur SIDN Labs
Leidinggevende

Marisca van der Donk

Managementassistente

Moritz Müller

Research engineer

Maarten Wullink

Research engineer

Thymen Wabeke

Research engineer

Marco Davids

Research engineer

Thijs van den Hout

Research engineer

Giovane Moura

Data scientist

Jelte Jansen

Research engineer

Caspar Schutijser

Research engineer

Ralph Koning

ResearchEngineer

Elmer Lastdrager

Research engineer

- Technical experts, diverse in seniority and nationality
- Help SIDN teams, write open-source software, analyze large amounts of data, conduct experiments, write articles, collaborate with universities
- M.Sc students help us advance specific areas

SIDN Tech Talks

- **Next**: April 12, 18:00-21:00 (twice a year)
- **Where**: Arnhem, SIDN office
- **Language**: Dutch

- **Topics**:
  - Automatic risk assessment of .nl registrations with RegCheck.
  - One-click rollout of nameservers in the .nl anycast DNS platform.
  - Drinks.

- **Link**: https://www.meetup.com/nl-NL/sidn-techtalk

2STiC program

Goal: put Dutch and European internet communities in a leading position in the field of secure, stable and transparent inter-network communication

The Internet

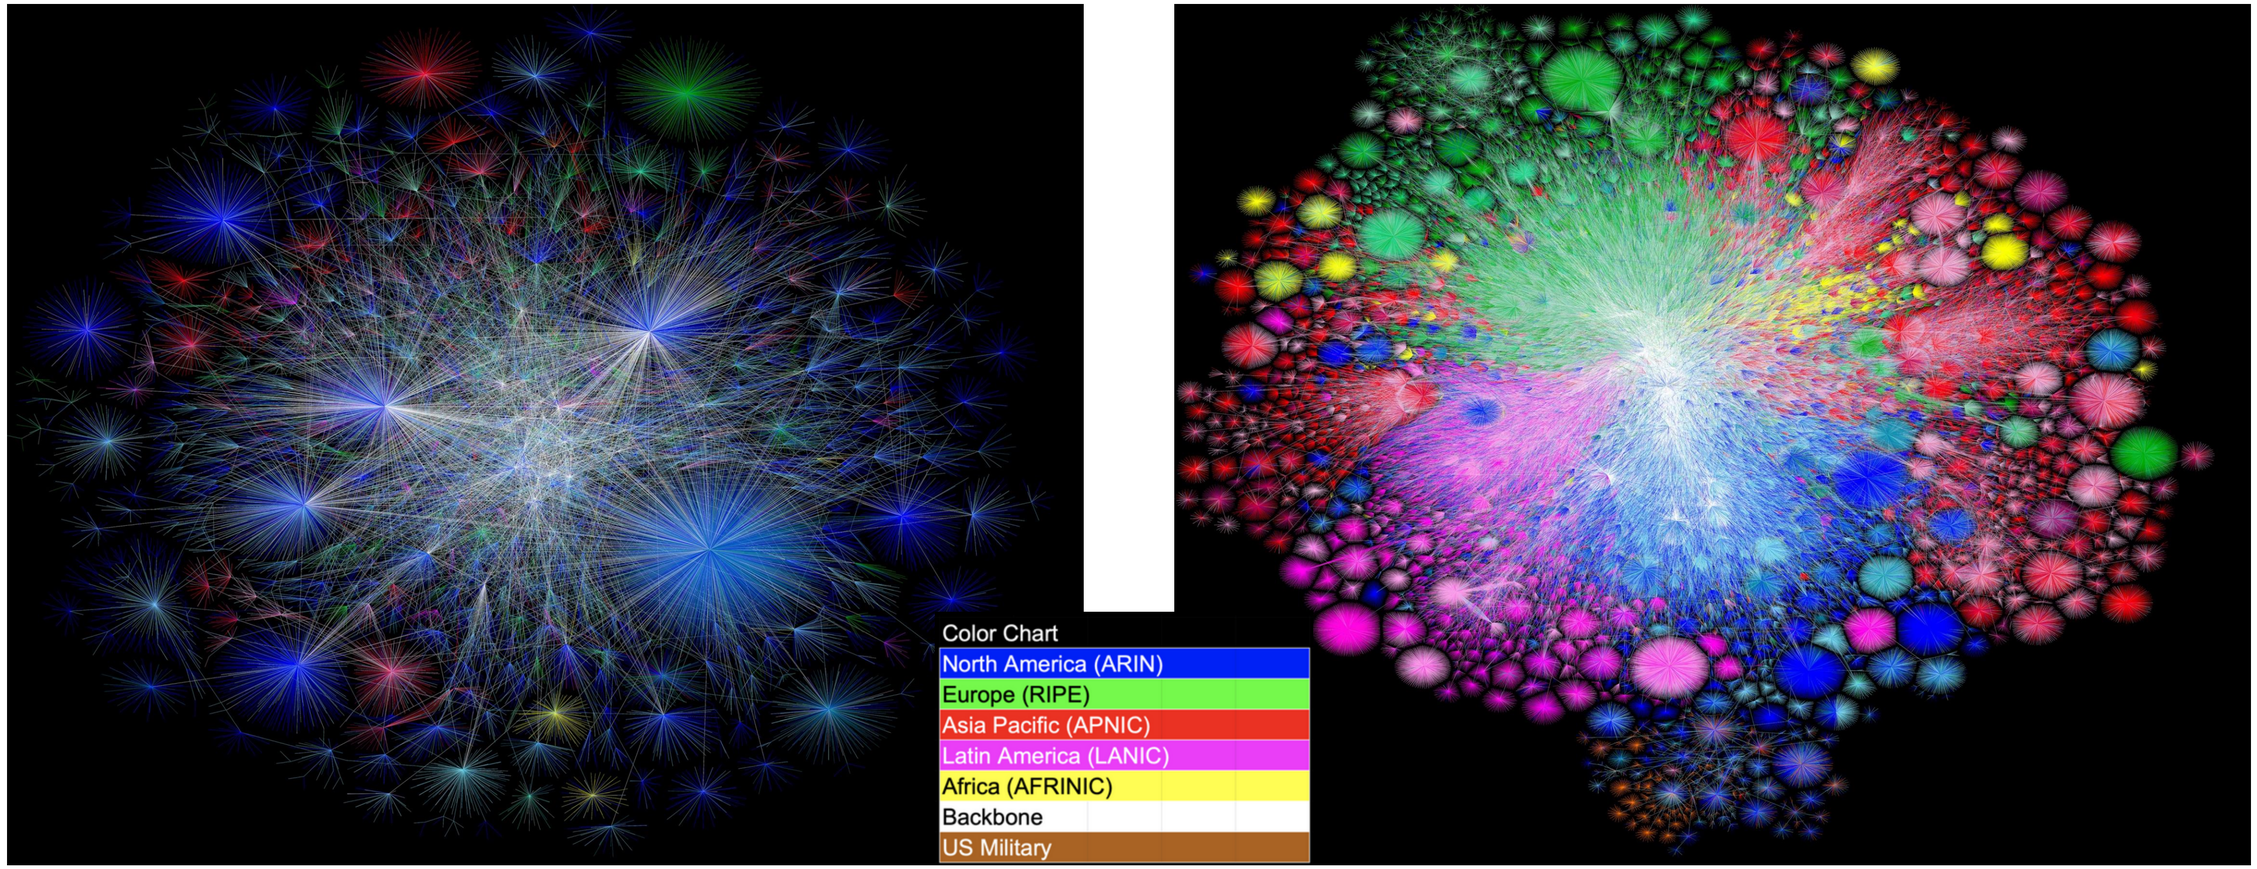

1997

2021

source: https://www.opte.org

Rate of change

Fast

Slow!

Fast

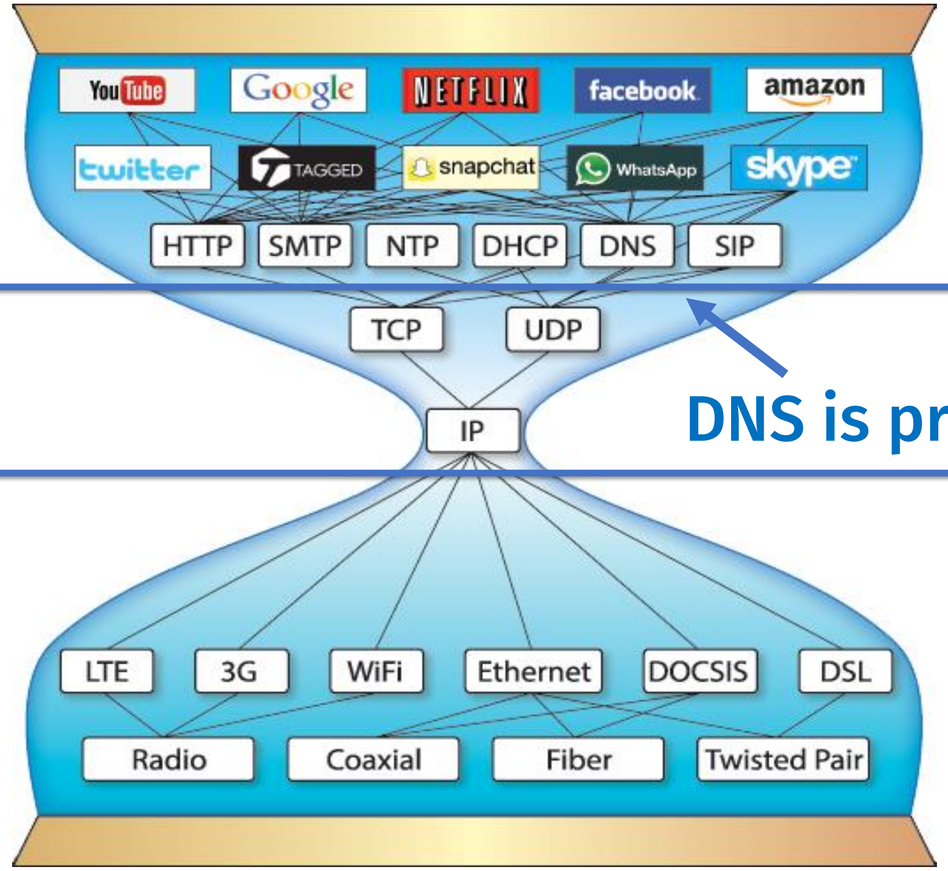

DNS is pretty slow too

New Requirements

- New applications have new **security**, **stability** and **transparency** requirements
  - More interaction with physical space (e.g., transport, smart grids, drones, remote surgery)

- To provide **trust** and and **confidence** in communication we need a **responsible** internet
  - Control over routing and verification of operational behavior

SCION
NDN
RINA
ManyNets
XIA
MobilityFirst
Nebula
Service-centric networking
FII
B4
…

Some new inter-domain networked architectures

Opening up

- Adoption of new protocols in technologies was slow, but network devices are opening up.
- (Onie) Open Network Install Environment offers OS choice on network equipment.
- OpenFlow/SDN offer control plane programmability.
- P4 provides dataplane programmability.

Potentially promising clean slate architectures

- RINA
  - Everything is IPC
  - WIP implementations: ProtoRINA, OpenIRATI
- NDN
  - Data centric
  - Stateful, lots of caching in the network
  - Implementation: named-data.net
- Extensible Internet
  - Introduces layer 3.5
  - No implementation yet
- **SCION**
  - **Path selection**
  - **Active community**
  - **Implementation: github.com/scionproto**

2STiC

SCION

2STiC

SCION

- Scalability, Control, and Isolation On Next-generation Networks
- New internet architecture
- Network Security Group, ETH Zurich
- Goal: improve security of inter-domain routing and isolation of compromise
- Scalability and security through Isolation Domains (ISDs)
  - Group of autonomous systems
  - E.g., per country or jurisdiction

SCION

- Security by design
  - Routes authenticated both in control and data plane
- Path-aware networking
  - Sender selects path
  - Enables, for example, geofencing
- Multi-path communication
  - Can be used, for example, for redundancy
- Existing application can still be used

Isolation domains

- Group of autonomous systems
  - E.g., per country or jurisdiction
- ISD core: ASes managing the ISD
- Core AS: AS part of the ISD core
- PKI organised per ISD
- Hierarchical control plane
  - Inter-ISD control plane
  - Intra-ISD control plane

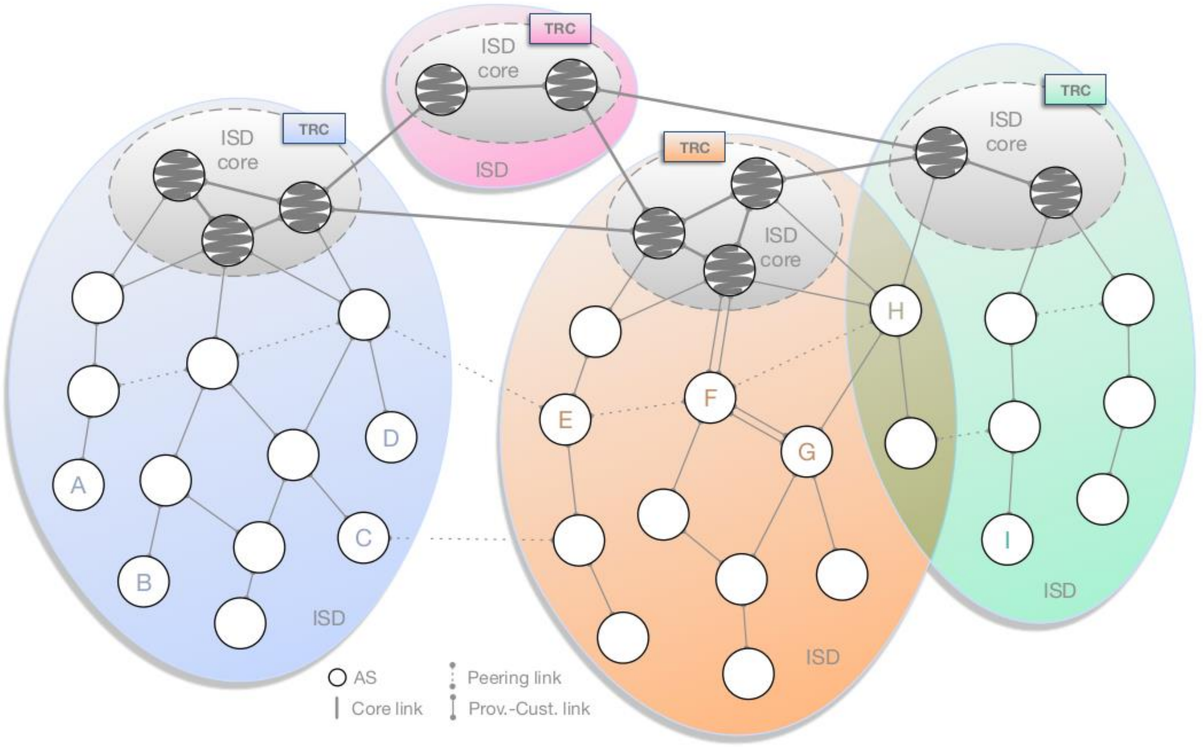

Source: The SCION Internet Architecture: An Internet Architecture for the 21st Century, Barrera et al., 2017

Routing: overview

- Control plane
  - Construct and disseminate path segments
- Data plane
  - Combine path segments to path
  - Packets contain path
  - Routers forward packets based on path (stateless)

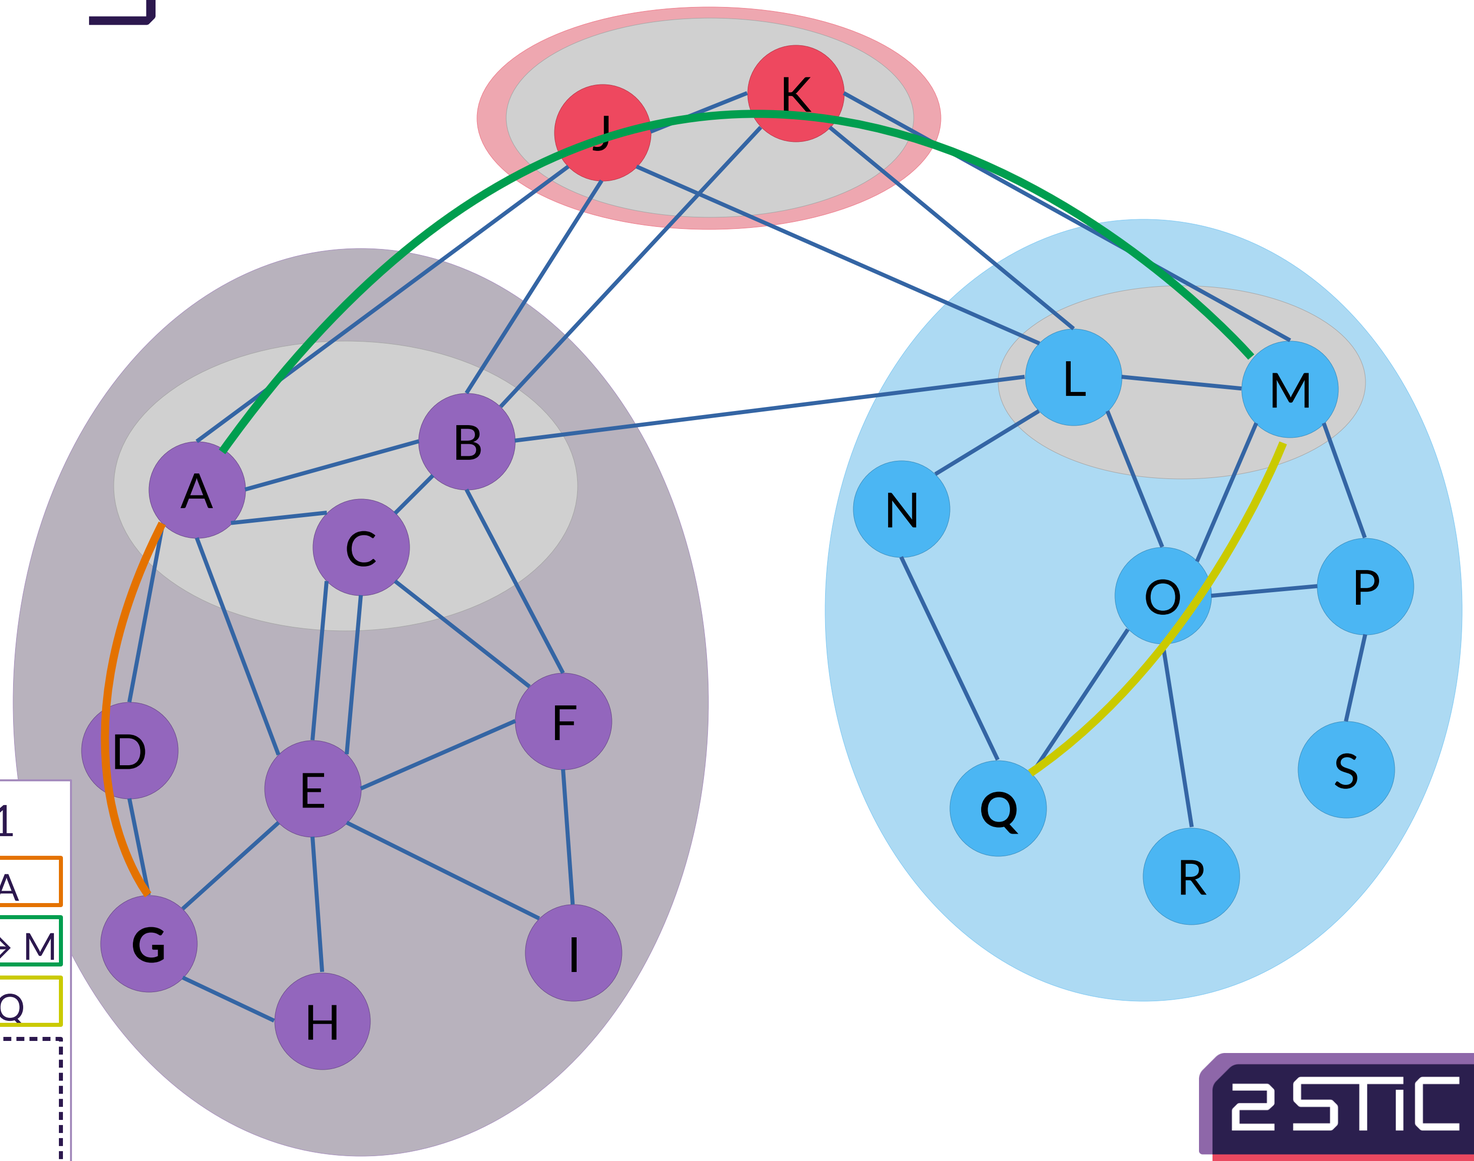

Control plane: path exploration

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

Control plane: Intra-ISD path exploration

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

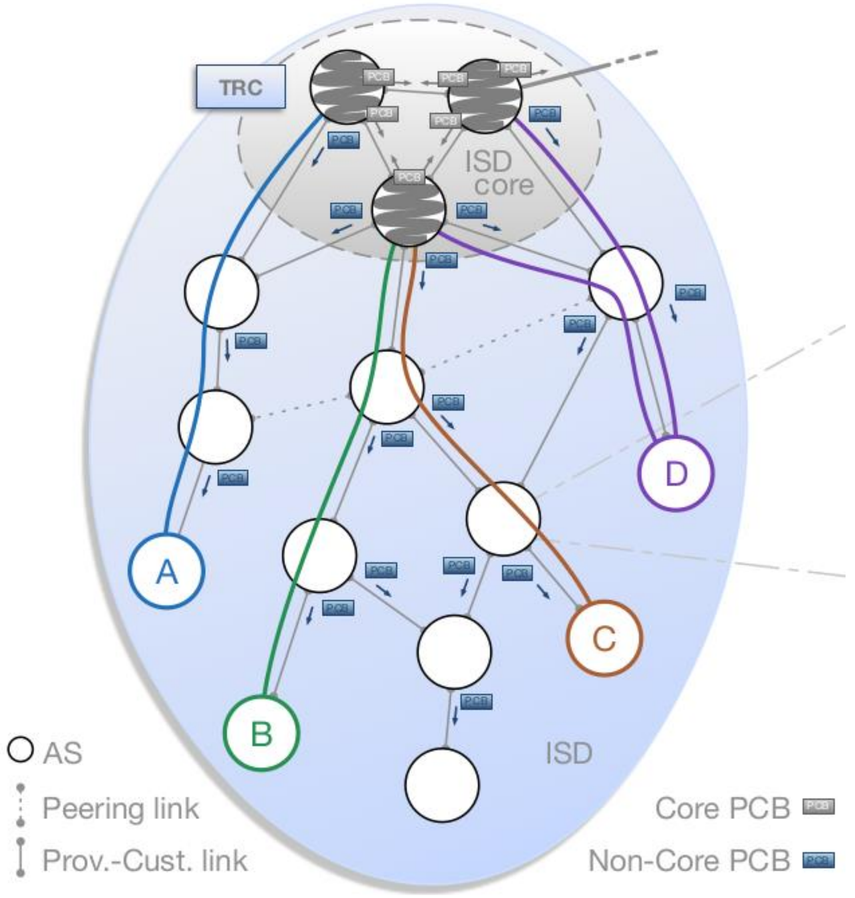

Source: The SCION Internet Architecture: An Internet Architecture for the 21st Century, Barrera et al., 2017

Control plane: Path Construction Beacons

- Path Construction Beacons are signed by every AS along the path
  - Authenticated path
- Hop fields included that can be used to later select paths
  - Contain forwarding information
  - Contain cryptographic MAC computed using hop field key
  - Only processed locally

Data plane: 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 segments in remote ISD from core to destination AS
  - Core-path segments needed to connect up-path and down-path segments
- End hosts pick and combine segments to determine path

Data plane: path combination

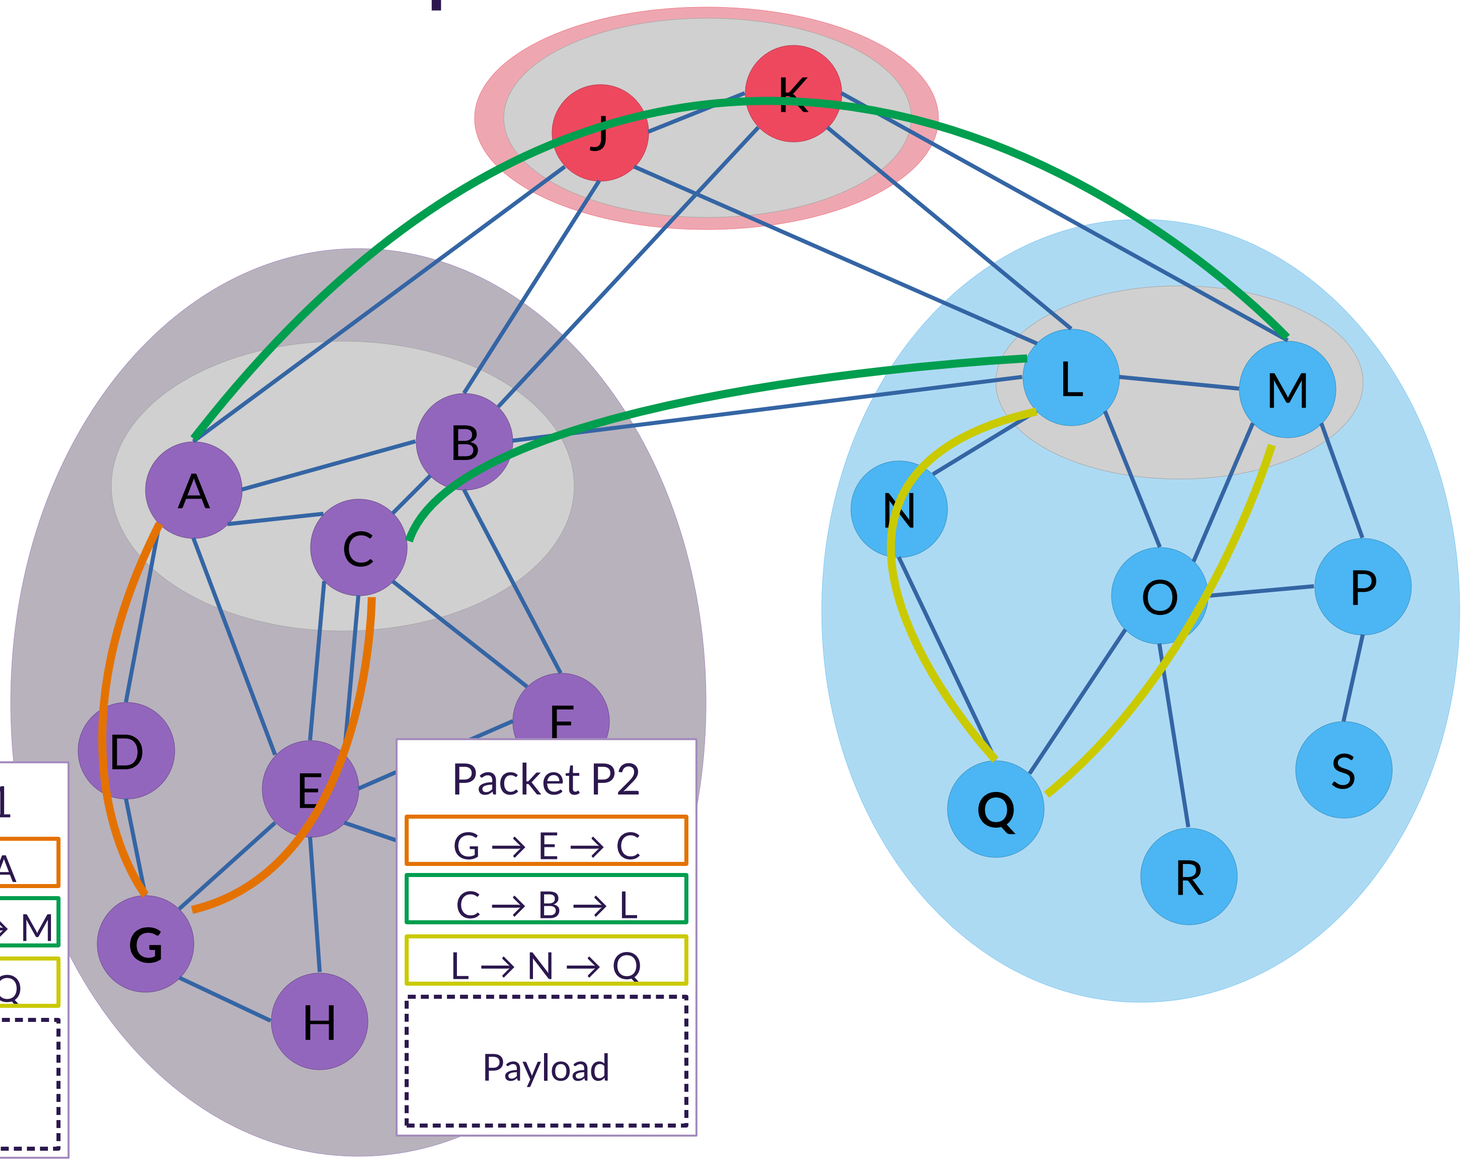

Data plane: path combination

- Possible paths determined by
  - Up-stream AS, by deciding which PCBs to forward to where
  - Core AS, by offering path segments to path server in local AS
  - Local AS, by registering down-path segments with ISD core
  - Local AS, by offering path segments to clients
  - Clients, by combining path segments offered by local path server

Routing summary

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

Security

- Path information authenticated in control plane and data plane
- Control plane
  - Beacons authenticated using digital signatures
  - No route hijacks
- Data plane
  - User selects path
  - Hop fields ensure only authorised paths possible

Security

- Address spoofing no longer possible on AS-level
  - Protects against reflection attacks
  - Reduces impact of DDoS attacks
- Hidden paths
  - Path information not published
  - Can only be used by parties that know the relevant hop fields

- EPIC (Every Packet Is Checked) is a dataplane extension that offers source authentication and path validation

Reliability and QoS

- Redundancy through use of multi-path communication
- Fast failover in case of link failure
  - No waiting for convergence
- Possible to add latency information to beacons
  - Path selection based on latency
- COLIBRI extension
  - Minimum bandwidth reservation

Deployment

- Open source implementation available
  - https://github.com/scionproto/scion
- International testbed SCIONLab
  - https://www.scionlab.org/
- Production network managed by spin-off Anapaya
- In use at banks, government and hospitals (in .ch)

- SCION Education production network in collaboration with EU NREN (Geant)

Transitioning to SCION

- Can be combined with existing applications using SCION-IP Gateway

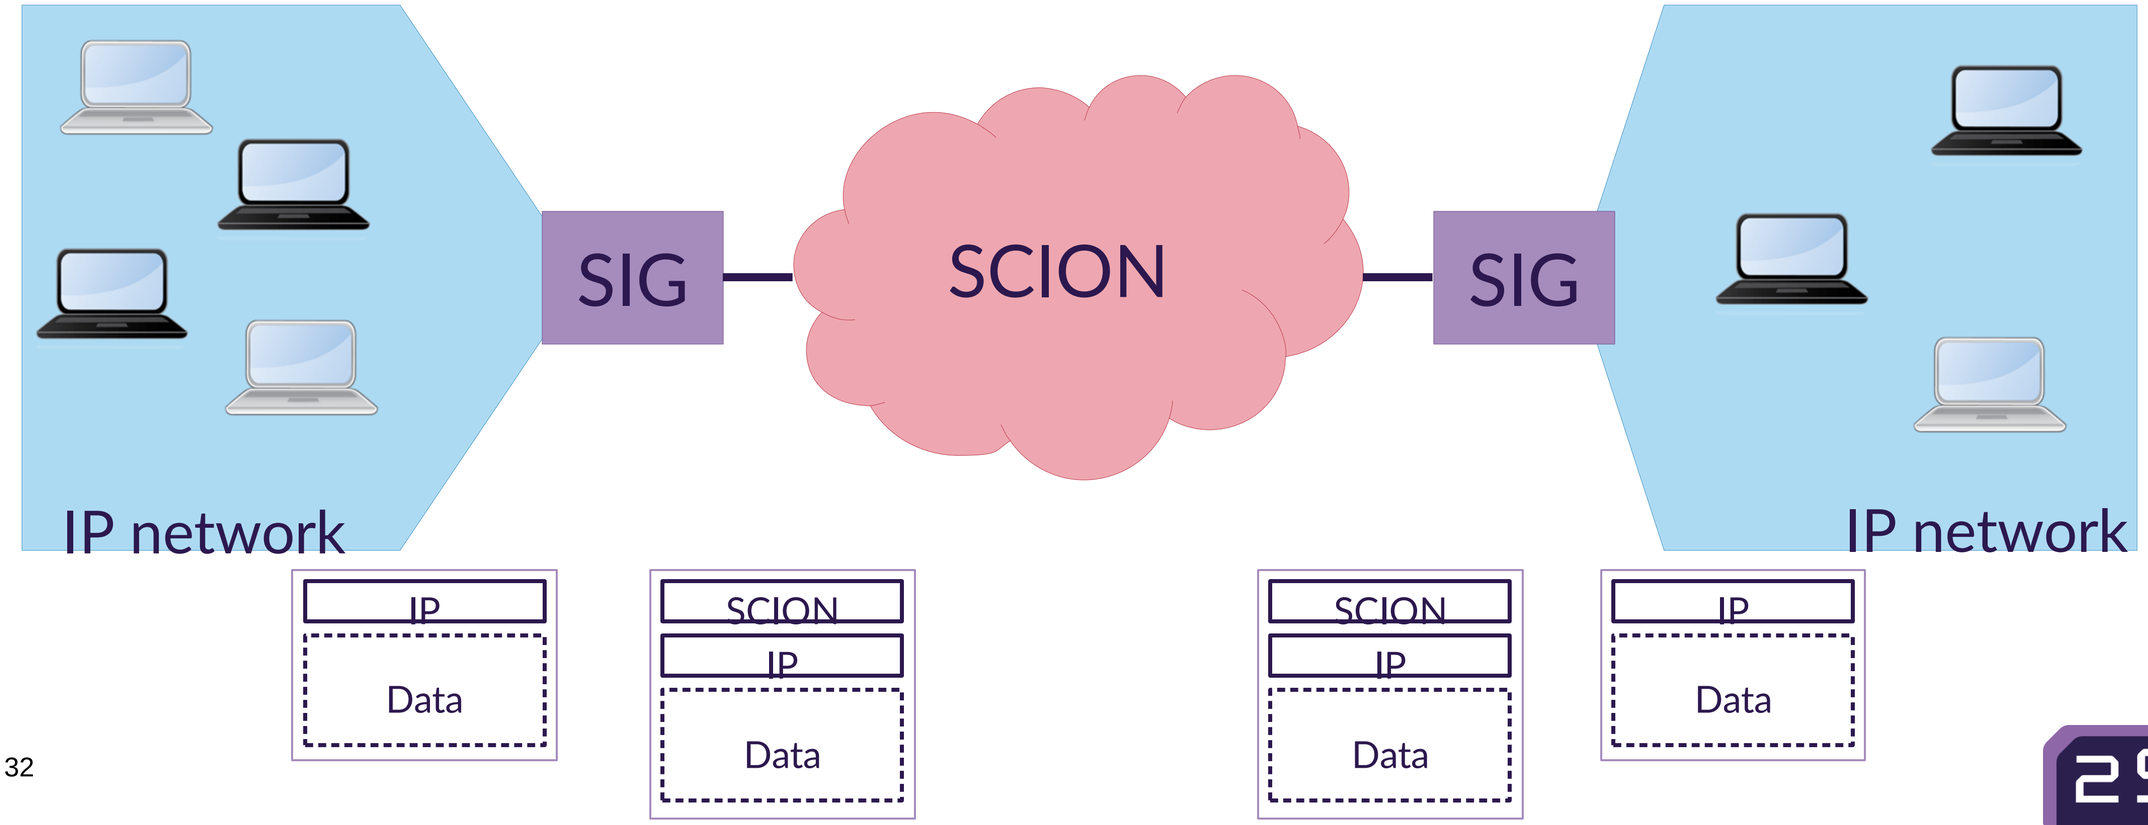

SCION in P4

A new internet architecture in P4

- We implemented the SCION internet architecture in P4 for the Intel Tofino
- Determine feasibility of running a new architecture on switch hardware and evaluate performance

P4

“Programming Protocol-independent Packet Processors (P4) is a domain-specific language for network devices, specifying how data plane devices (switches, NICs, routers, filters, etc.) process packets.”

Source: www.p4.org

Intel Tofino

- P4-programmable Ethernet switch ASIC
- Switches available with e.g., 32 or 64 100 Gbit/sec ports

Some challenges

- No support for cryptographic operations in Intel Tofino
- Protocol not designed for hardware
  - Complex headers

No cryptographic operations

- MACs verified using table containing all currently valid values
- Populated from control plane when MACs are generated
  - In the SCION control plane
  - At the switch
- Invalid entries removed

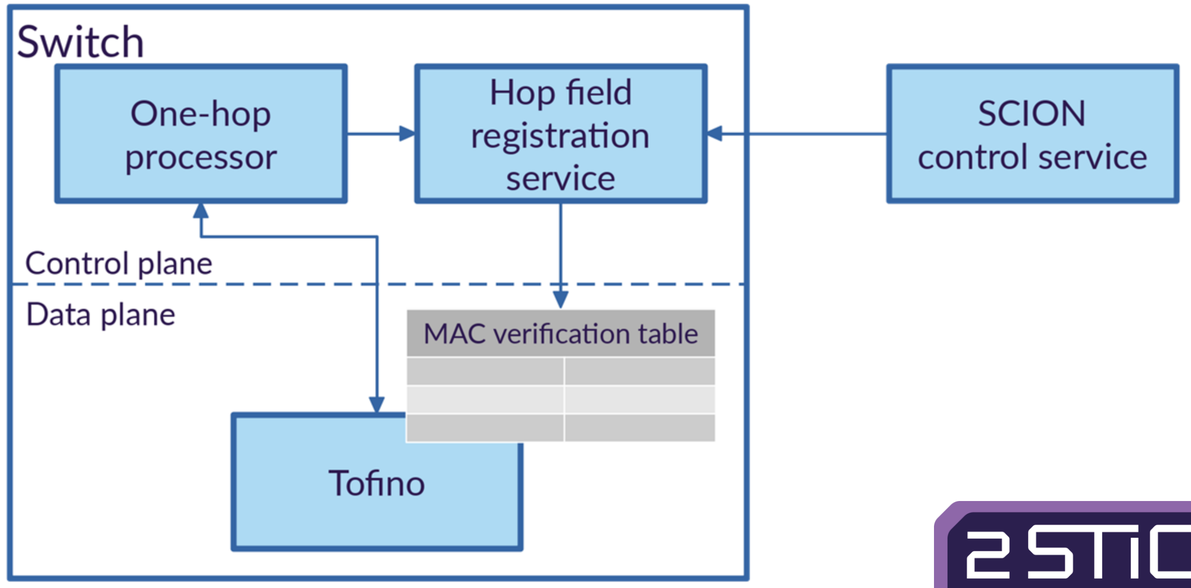

Complex header fields

- For example: forwarding path consisted of nested lists
- Flattening the structure provides for more efficient parsing

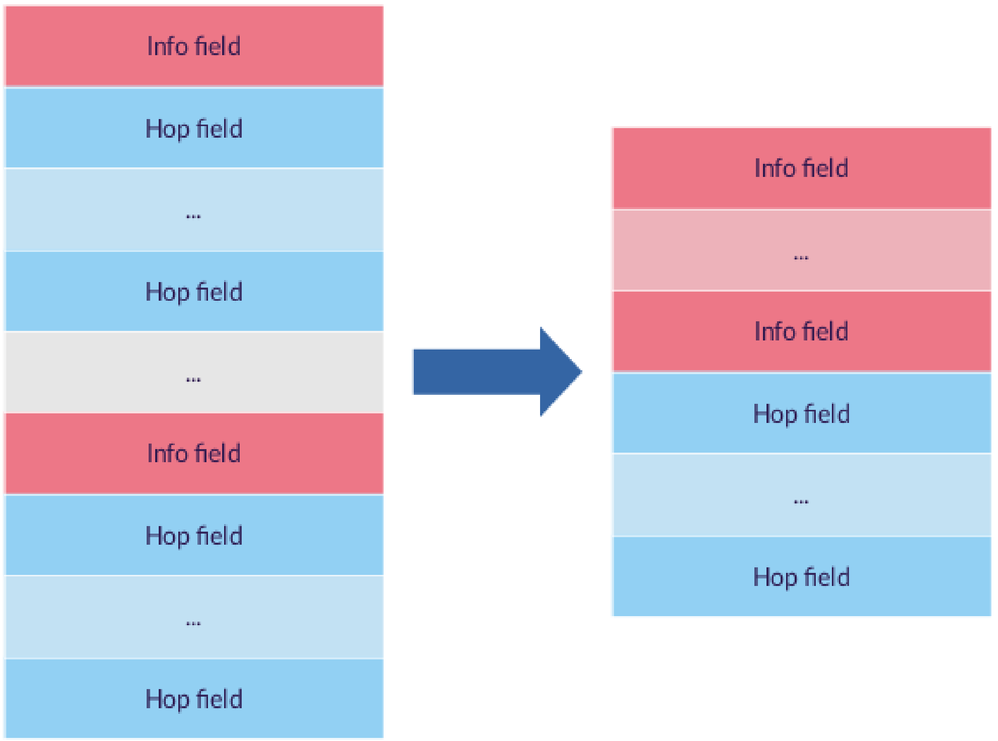

Lessons learned

- When designing a protocol with hardware in mind
  - use explicit lengths
  - do not use absolute offsets
  - limit the usage of variable length fields
  - do not use complex data structures such as nested lists

Evaluation

- Edgecore switches with 32 100 Gbps ports
- Tested functionality with topology where all border routers ran on switches
- Tested performance using packet generator for different path lengths
  - Achieved near line-rate for almost all tested path lengths

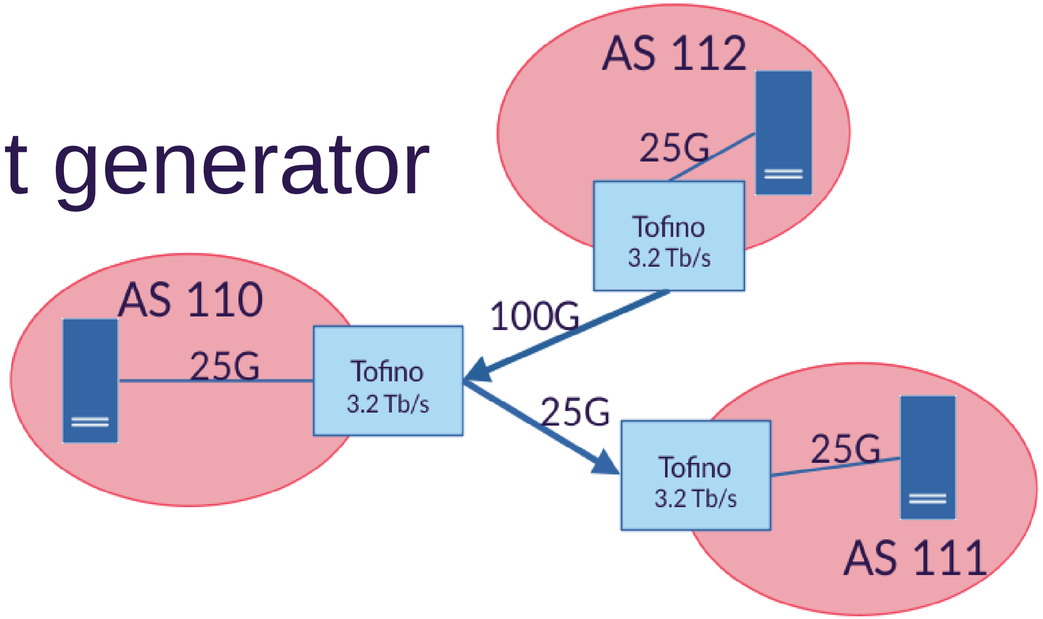

Conclusion

- SCION can be implemented for switch hardware and run on high speeds
- Several lessons learned regarding protocol design
- Future work
  - Support for protocol error handling and additional SCION-related protocols
  - More extensive performance analysis
    - Code is open source and available at github.com/SIDN/p4-scion

DISCUSSION

Discussion

To recap, today we've seen:

- Programmable network hardware;
- A new internet architecture;
- An active community.

With these developments in mind, we wonder: how do you see the future of the Internet?

Discussion

Routing transparency

- Do users need this much transparency and control?
- 3rd parties can tell a lot about how you are connected to the internet by looking at the headers in a single packet, is this desirable?

Isolation Domains

- What is a sustainable governance structure for an ISD? Should this be government controlled or not?
- Will ISDs (and extensions) violate net neutrality?

2STIC

LAB

2STIC

SCION address structure

- An AS: **`ISD-AS`**
- A host inside an AS: **`ISD-AS`**`,[`**`address`**`]`

- Examples:
  - `19-ffaa:0:1305`
  - `19-ffaa:0:1305,[127.0.0.1]`
  - `19-ffaa:0:1305,[::1]`

SCIONLab testbed

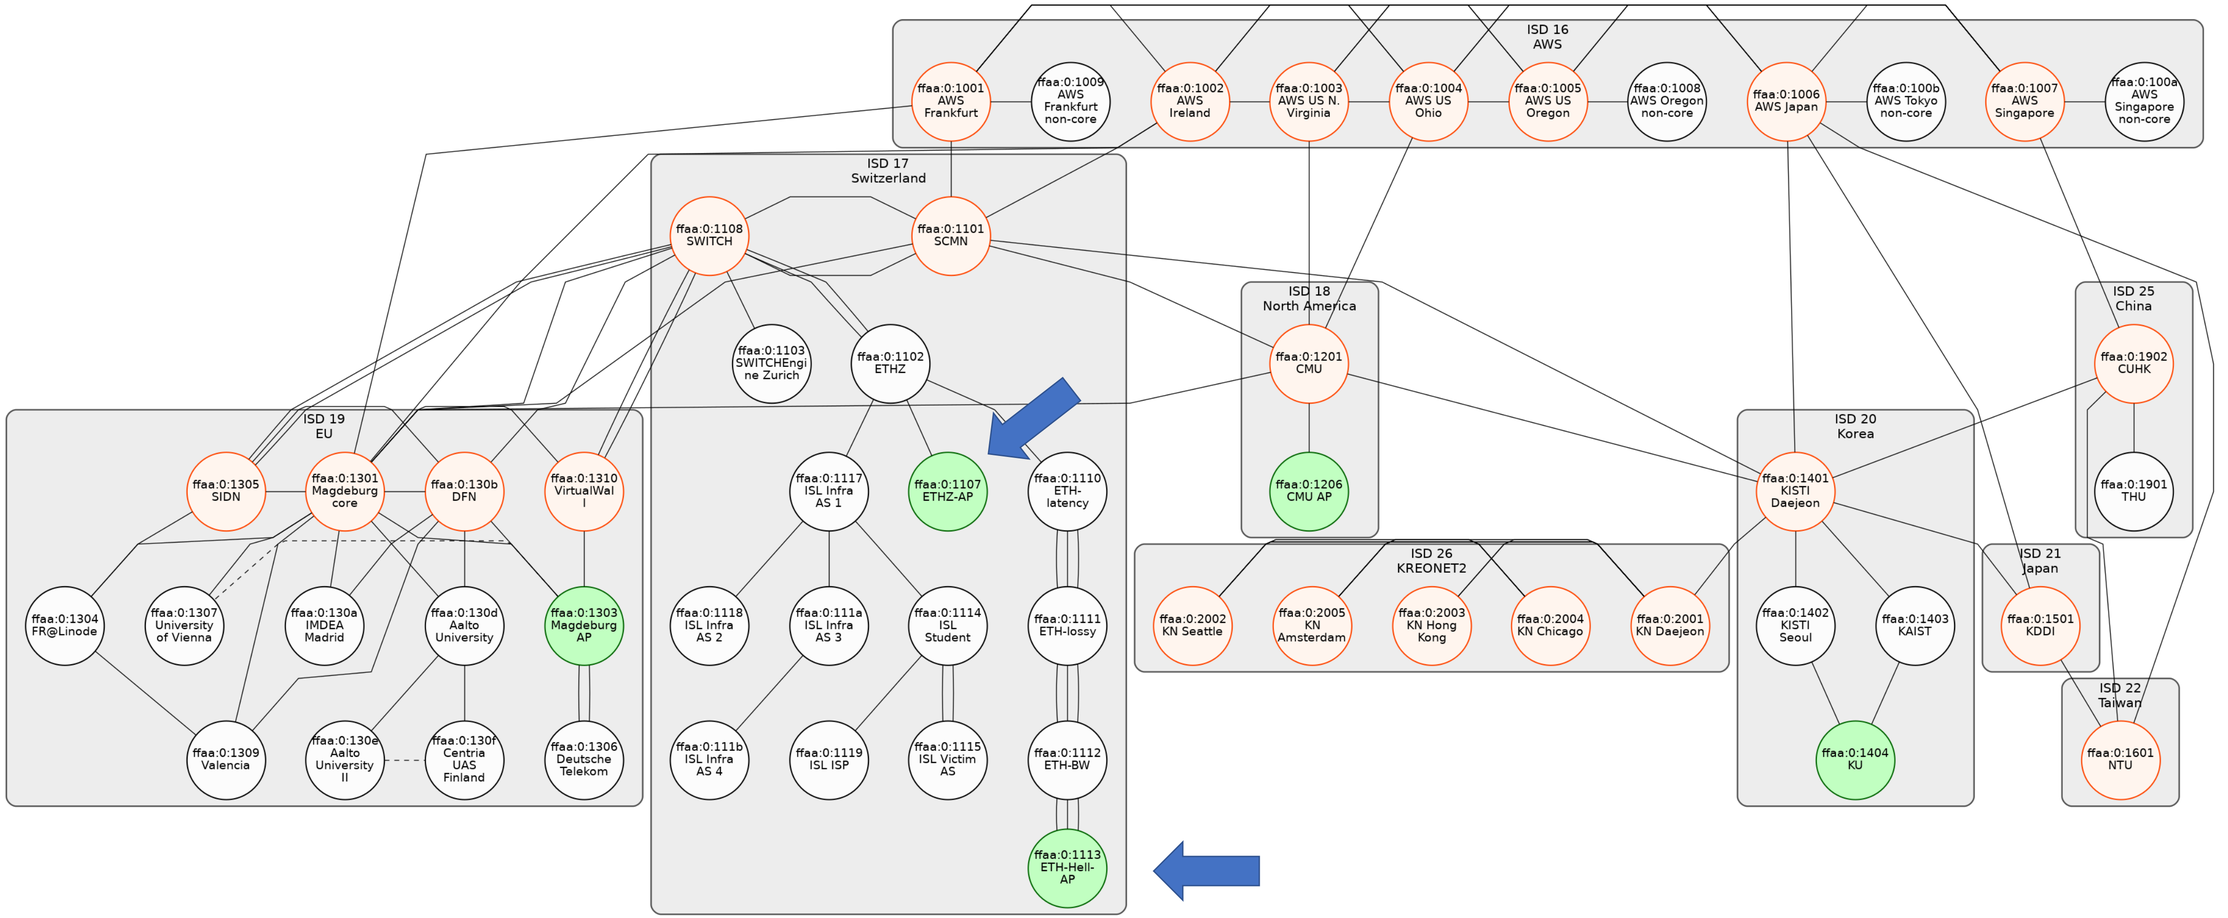

SCIONLab exercises

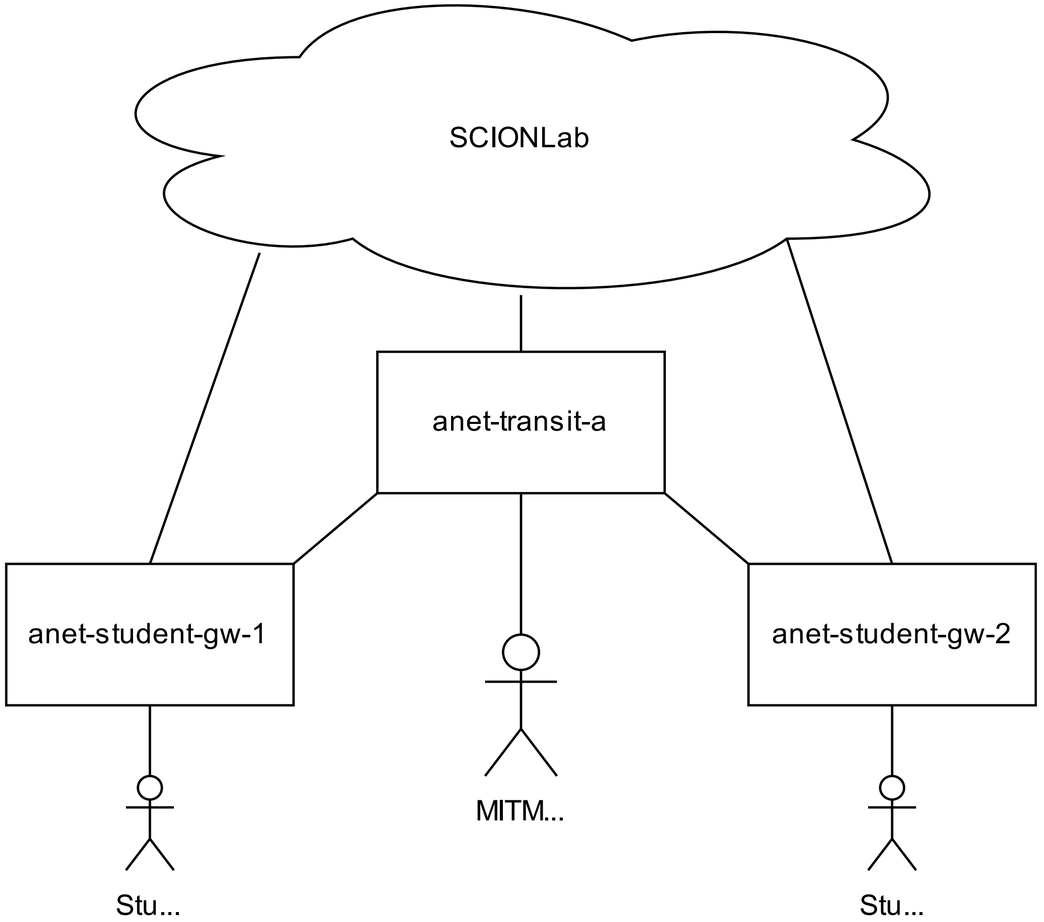

Viewer does not support full SVG 1.1

SCIONLab exercises

- Make groups of (min) 2 students.
- Instructions at https://check.sidnlabs.nl/ralph/anet-lab/
- Scion-netcat at: https://check.sidnlabs.nl/ralph/anet-lab/scion-netcat.gz
- https://www.scionlab.org

2STIC

Thanks for your attention!

Caspar Schutijser, Ralph Koning
sidnlabs.nl
2stic.nl