



FAKULTÄT FÜR
INFORMATIK



SCION Edge Router for Legacy IP Applications based on Intel Tofino

M.Sc. Lars-Christian Schulz

M.Sc. Robin Wehner

Prof. Dr. David Hausheer

Structure

1. Background – SCION
2. Motivation for a SCION IP Translator
3. Concepts – Translating IP addresses to SCION
4. P4 Implementation on Tofino
5. Deployment in SCIERA
6. Conclusion & Future Work

1. Background: SCION Architecture


Path-based Network Architecture

Control Plane - Routing

- ❖ **Constructs** and **Disseminates** Path Segments

Data Plane - Packet forwarding

- ❖ **Combine** Path Segments to Path
- ❖ Packets contain Path
- ❖ Routers forward packets based on Path
 - ▶ Simple routers, stateless operation



1. Background: SCION Architecture





Path-based Network Architecture

Control Plane - Routing

- ❖ **Constructs** and **Disseminates** Path Segments

Data Plane - Packet forwarding

- ❖ **Combine** Path Segments to Path
- ❖ Packets contain Path
- ❖ Routers forward packets based on Path
 - ▶ Simple routers, stateless operation



1. Background: SCION Architecture





Path-based Network Architecture

Control Plane - Routing

- ❖ **Constructs** and **Disseminates** Path Segments

Data Plane - Packet forwarding

- ❖ **Combine** Path Segments to Path
- ❖ Packets contain Path
- ❖ Routers forward packets based on Path
 - ▶ Simple routers, stateless operation



1. Background: SCION Architecture



Path-based Network Architecture

Control Plane - Routing

- ❖ **Constructs** and **Disseminates** Path Segments

Data Plane - Packet forwarding

- ❖ **Combine** Path Segments to Path
- ❖ Packets contain Path
- ❖ Routers forward packets based on Path
 - ▶ Simple routers, stateless operation



1. Background: SCION Architecture




Path-based Network Architecture

Control Plane - Routing

- ❖ **Constructs** and **Disseminates** Path Segments


Data Plane - Packet forwarding

- ❖ **Combine** Path Segments to Path
- ❖ Packets contain Path
- ❖ Routers forward packets based on Path
 - ▶ Simple routers, stateless operation



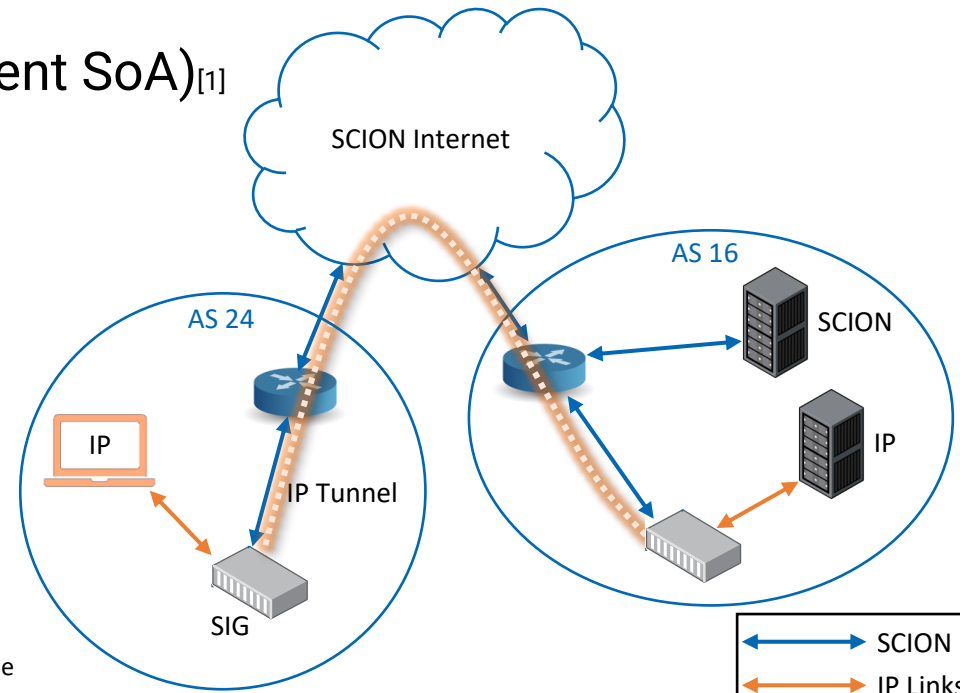
2. Motivation

- ❖ SCION solves many issues of today's Internet
- ❖ How can legacy IP applications use the SCION network?



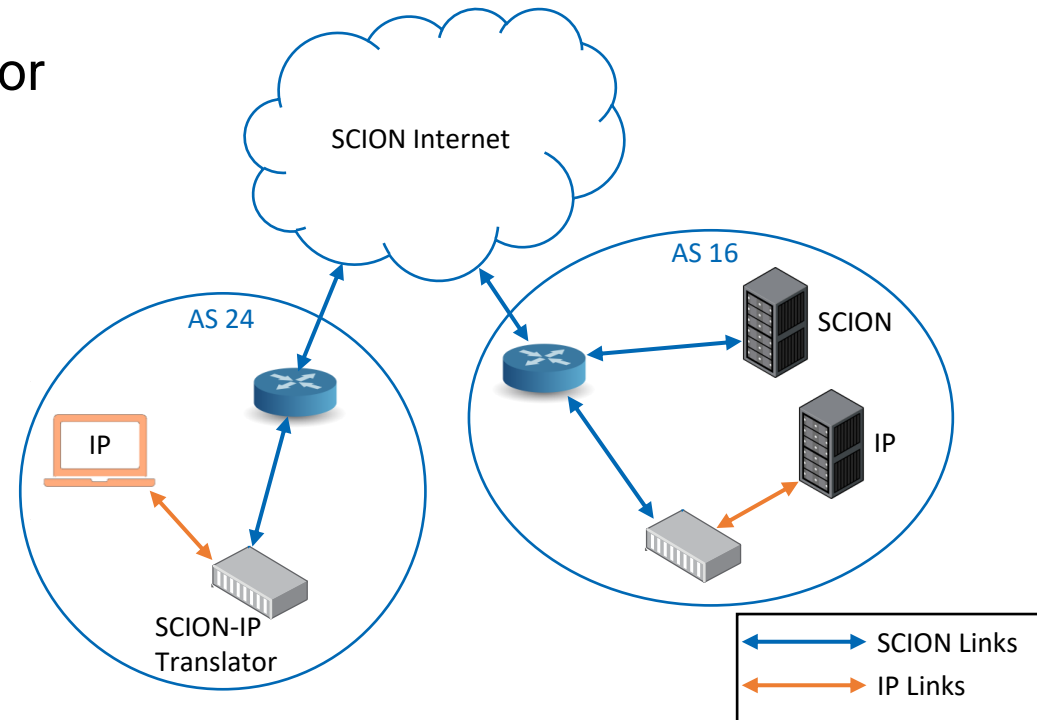
2. Motivation

- ❖ SCION solves many issues of today's Internet
- ❖ How can legacy IP applications use the SCION network?
 - Use SCION-IP-Gateway (SIG) (Current SoA)^[1]
 - Tunnels IP over SCION
 - **Complex encapsulation**
 - **Have to exchange IP routes with each other**
 - **No communication between SCION and IP apps**



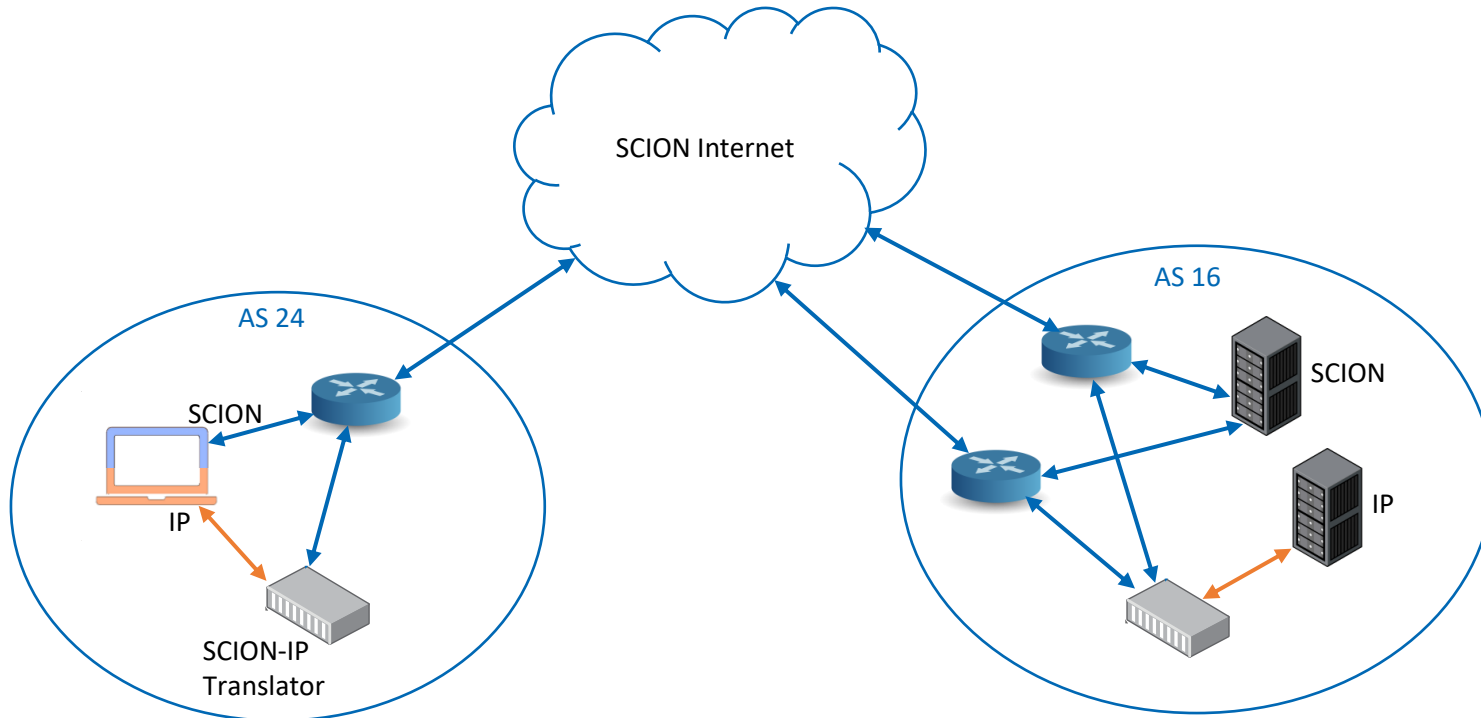
2. Motivation

- ❖ SCION solves many issues of today's Internet
- ❖ How can legacy IP applications use the SCION network?
 - We propose the SCION IP Translator
 - Allows IP applications to use the SCION network
 - Allows communication between SCION and IP implementations of an app




3. Concepts - Deployment

❖ Independent SCION-IP translator



3. Concepts - Deployment

❖ Integrated SCION-IP translator



3. Concepts - Translating IP addresses to SCION

- ❖ SCION addressing: **ISD** – **AS** – Host Address
Example server: **1** **18** **fd00::0d93**
- ❖ Encode SCION addresses in IPv6 (“SCION-mapped IP”)

IPv6:

128-bit IPv6 address					
8 bit	12 bit	20 bit	24 - m bit	m bit	64 bit
global routing prefix				subnet ID	interface ID
SCION prefix	ISD	ASN	local routing prefix	subnet ID	interface ID

IPv4:

128-bit IPv6 address					
8 bit	12 bit	20 bit	24 bit	32 bit	32 bit
global routing prefix				interface ID	
SCION prefix	ISD	ASN	0	0x0000ffff	IPv4 address

3. Concepts - Translating IP addresses to SCION

- ❖ SCION addressing: **ISD** – **AS** – Host Address
Example server: **1** **18** **fd00::0d93**
- ❖ Encode SCION addresses in IPv6 (“SCION-mapped IP”)
fc
 - 8 Bit global SCION prefix

IPv6:

128-bit IPv6 address					
8 bit	12 bit	20 bit	24 - m bit	m bit	64 bit
	global routing prefix			subnet ID	interface ID
SCION prefix	ISD	ASN	local routing prefix	subnet ID	interface ID

IPv4:

128-bit IPv6 address					
8 bit	12 bit	20 bit	24 bit	32 bit	32 bit
global routing prefix				interface ID	
SCION prefix	ISD	ASN	0	0x0000ffff	IPv4 address

3. Concepts - Translating IP addresses to SCION

- ❖ SCION addressing: **ISD** – **AS** – Host Address

Example server: **1** **18** **fd00::0d93**

- ❖ Encode SCION addresses in IPv6 (“SCION-mapped IP”)

fc00:1

- 12 Bit ISD number because currently assigned from range 64-4094

IPv6:

128-bit IPv6 address					
8 bit	12 bit	20 bit	24 - m bit	m bit	64 bit
	global routing prefix			subnet ID	interface ID
SCION prefix	ISD	ASN	local routing prefix	subnet ID	interface ID

IPv4:

128-bit IPv6 address					
8 bit	12 bit	20 bit	24 bit	32 bit	32 bit
global routing prefix			interface ID		
SCION prefix	ISD	ASN	0	0x0000ffff	IPv4 address

3. Concepts - Translating IP addresses to SCION

❖ SCION addressing: **ISD** – **AS** – Host Address

Example server: **1** **18** **fd00::0d93**

❖ Encode SCION addresses in IPv6 (“SCION-mapped IP”)

fc00:1000:18

➤ 20 Bit ASN because current BGP allocation from block below 2^{19}

IPv6:

128-bit IPv6 address					
8 bit	12 bit	20 bit	24 - m bit	m bit	64 bit
global routing prefix			subnet ID	interface ID	
SCION prefix	ISD	ASN	local routing prefix	subnet ID	interface ID

IPv4:

128-bit IPv6 address					
8 bit	12 bit	20 bit	24 bit	32 bit	32 bit
global routing prefix			interface ID		
SCION prefix	ISD	ASN	0	0x0000ffff	IPv4 address

3. Concepts - Translating IP addresses to SCION

- ❖ SCION addressing: **ISD** – **AS** – Host Address

Example server: **1** **18** **fd00::0d93**

- ❖ Encode SCION addresses in IPv6 (“SCION-mapped IP”)

fc00:1000:18xx:xxxx::0d93

- 64 Bit Interface ID for simple host address assignment

IPv6:

128-bit IPv6 address					
8 bit	12 bit	20 bit	24 - m bit	m bit	64 bit
global routing prefix				subnet ID	interface ID
SCION prefix	ISD	ASN	local routing prefix	subnet ID	interface ID

IPv4:

128-bit IPv6 address					
8 bit	12 bit	20 bit	24 bit	32 bit	32 bit
global routing prefix				interface ID	
SCION prefix	ISD	ASN	0	0x0000ffff	IPv4 address

3. Concepts - Translating IP addresses to SCION

- ❖ SCION addressing: **ISD** – **AS** – Host Address


Example server: **1** **18** **fd00::0d93**

- ❖ Encode SCION addresses in IPv6 (“SCION-mapped IP”)

fc00:1000:18xx:xxxx::0d93

- ❖ All hosts need a SCION-mapped IP additionally to its existing IPv6

- Legacy SCION traffic to **fd00::1**
- IP traffic to SCION mapped IPv6
fc00:1000:18xx:xxxx::0d93



3. Concepts – ICMP-SCMP Translation

- ❖ Support for conversion of SCMP and ICMP
- ❖ **MTU exceeded** to facilitate path MTU discovery (also when SCION-IP conversion exceeds MTU)
- ❖ Path failures are not translated, but used to inform path selection
 - Choose another path

Code	ICMP	SCMP	Meaning
1	yes	yes	Destination unreachable
2	yes	yes	Packet too big
3	yes	no	Time exceeded
4	yes	yes	Parameter problem
5	no	yes	External interface down
6	no	yes	Internal connectivity down
128	yes	yes	Echo request
129	yes	yes	Echo reply
130	no	yes	Traceroute Request
131	no	yes	Traceroute reply

3. Concepts – ICMP-SCMP Translation

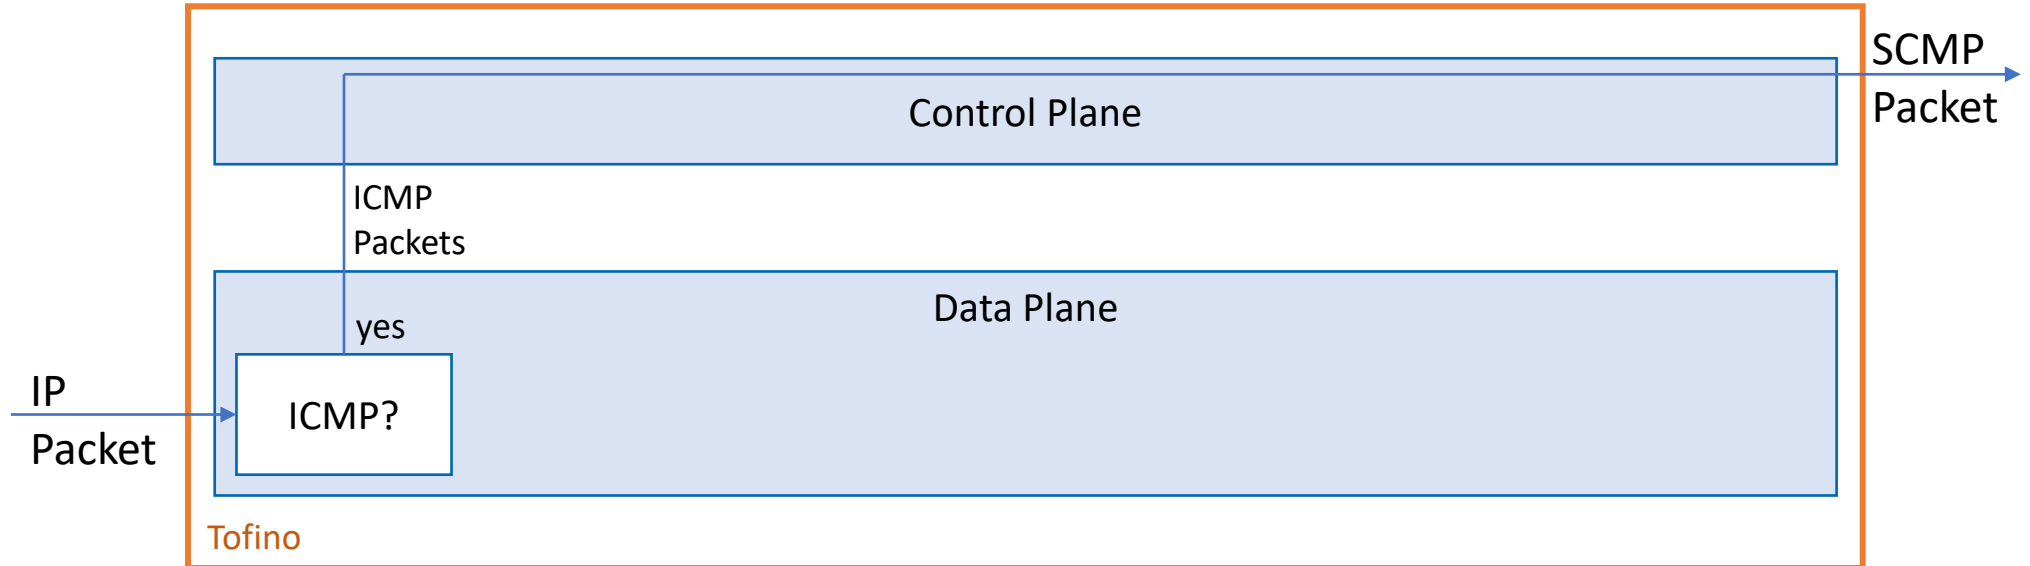
- ❖ Support for conversion of SCMP and ICMP
- ❖ MTU exceeded to facilitate path MTU discovery (also when SCION-IP conversion exceeds MTU)
- ❖ **Path failures** are not translated, but used to inform path selection
 - Choose another path

Code	ICMP	SCMP	Meaning
1	yes	yes	Destination unreachable
2	yes	yes	Packet too big
3	yes	no	Time exceeded
4	yes	yes	Parameter problem
5	no	yes	External interface down
6	no	yes	Internal connectivity down
128	yes	yes	Echo request
129	yes	yes	Echo reply
130	no	yes	Traceroute Request
131	no	yes	Traceroute reply

4. Implementation

- ❖ P4 implementation targeting Tofino 1 & Tofino 2
- ❖ Translate IP to SCION

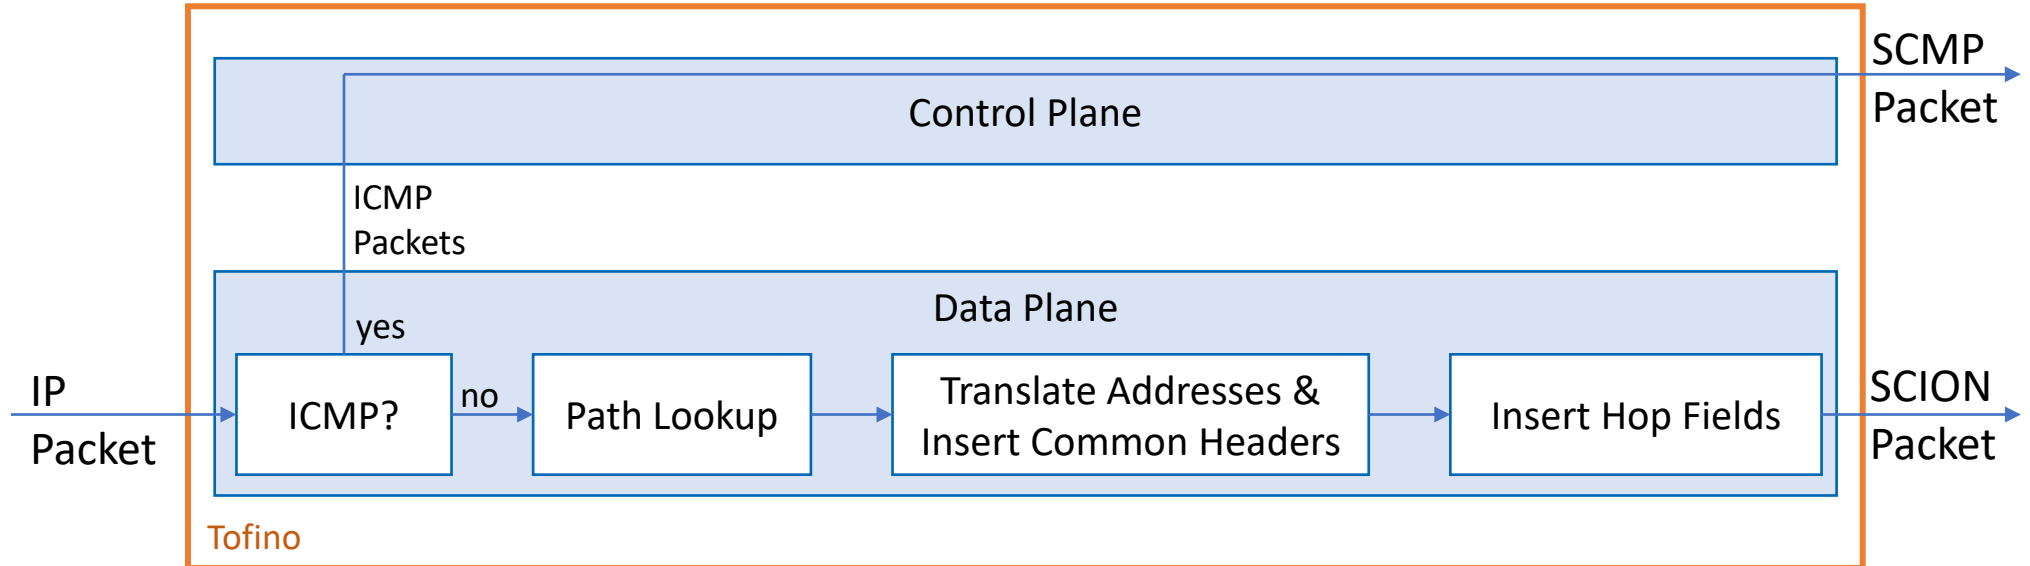
SCION Control Service



4. Implementation

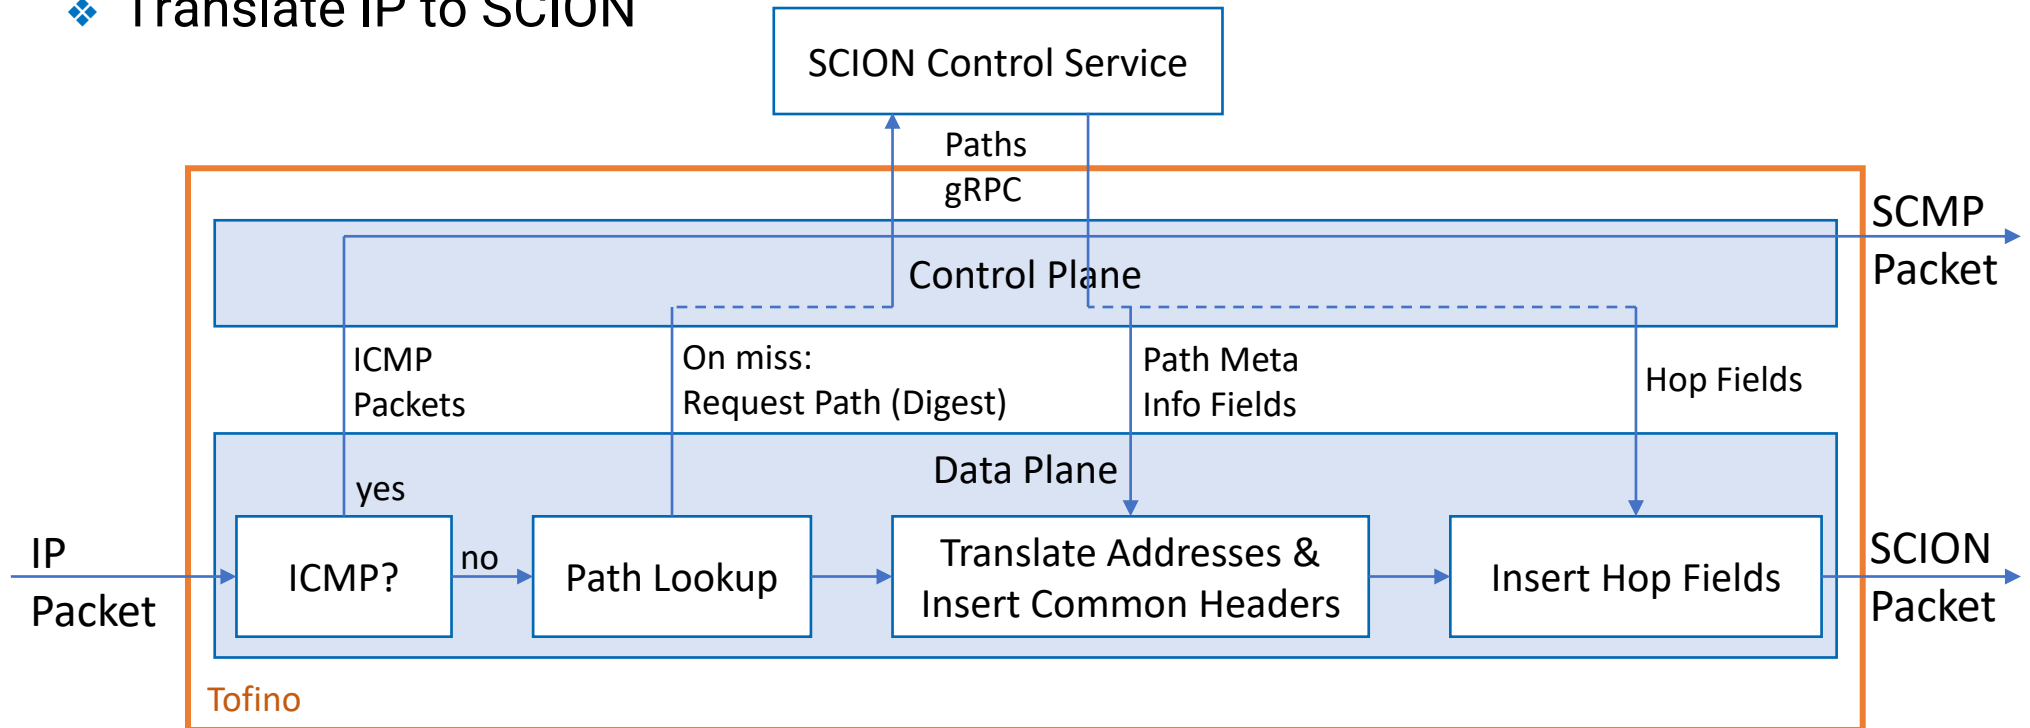
- ❖ P4 implementation targeting Tofino 1 & Tofino 2
- ❖ Translate IP to SCION

SCION Control Service



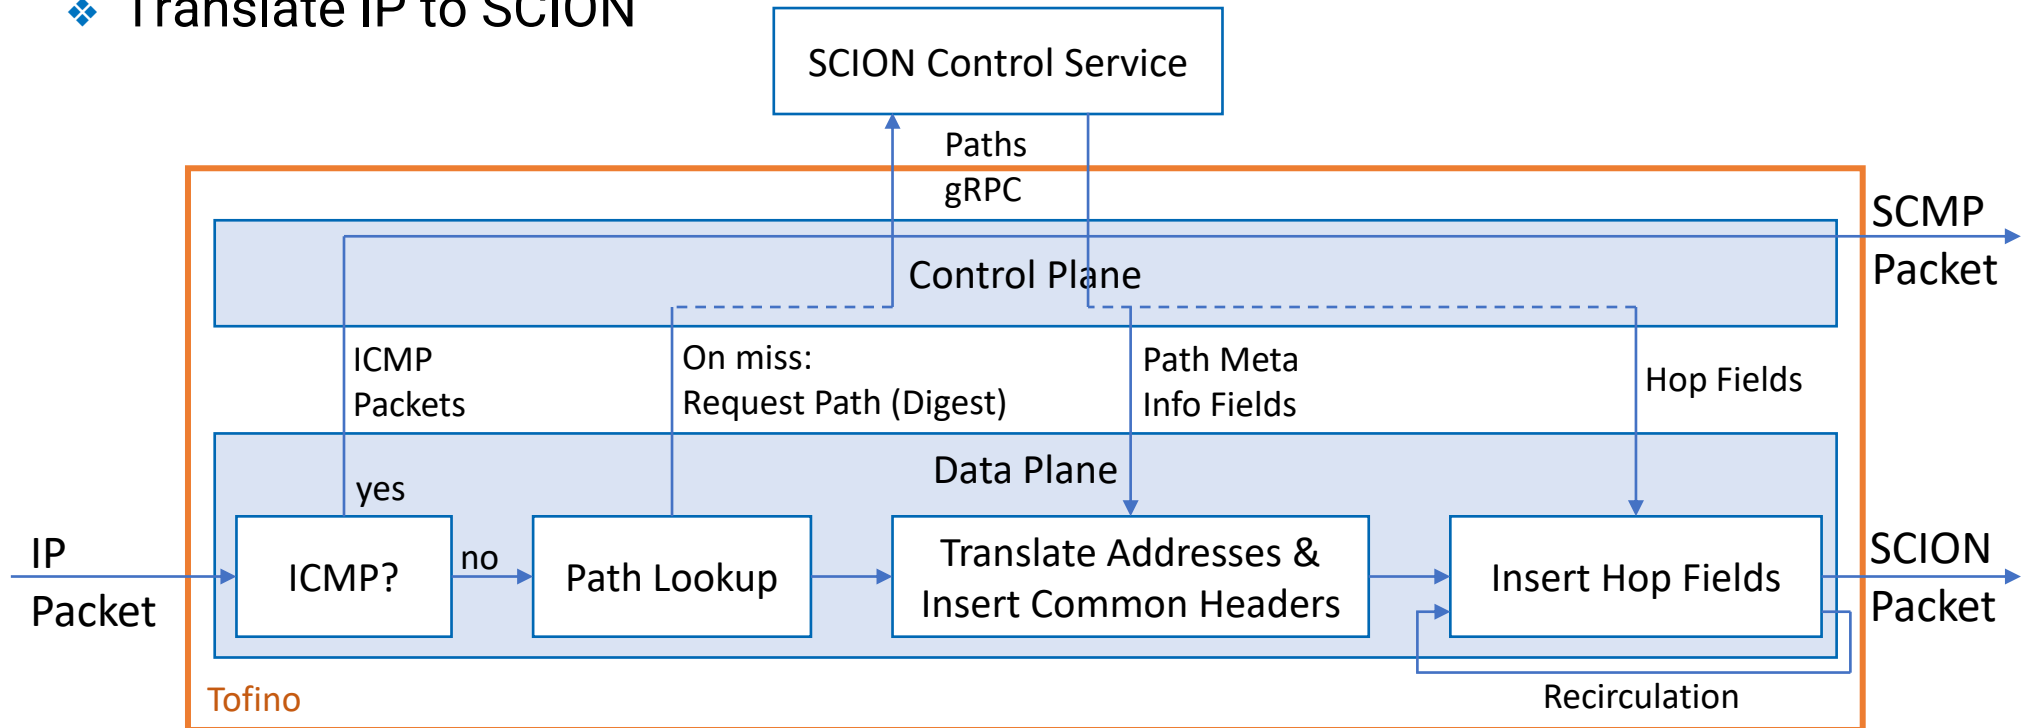
4. Implementation

- ❖ P4 implementation targeting Tofino 1 & Tofino 2
- ❖ Translate IP to SCION



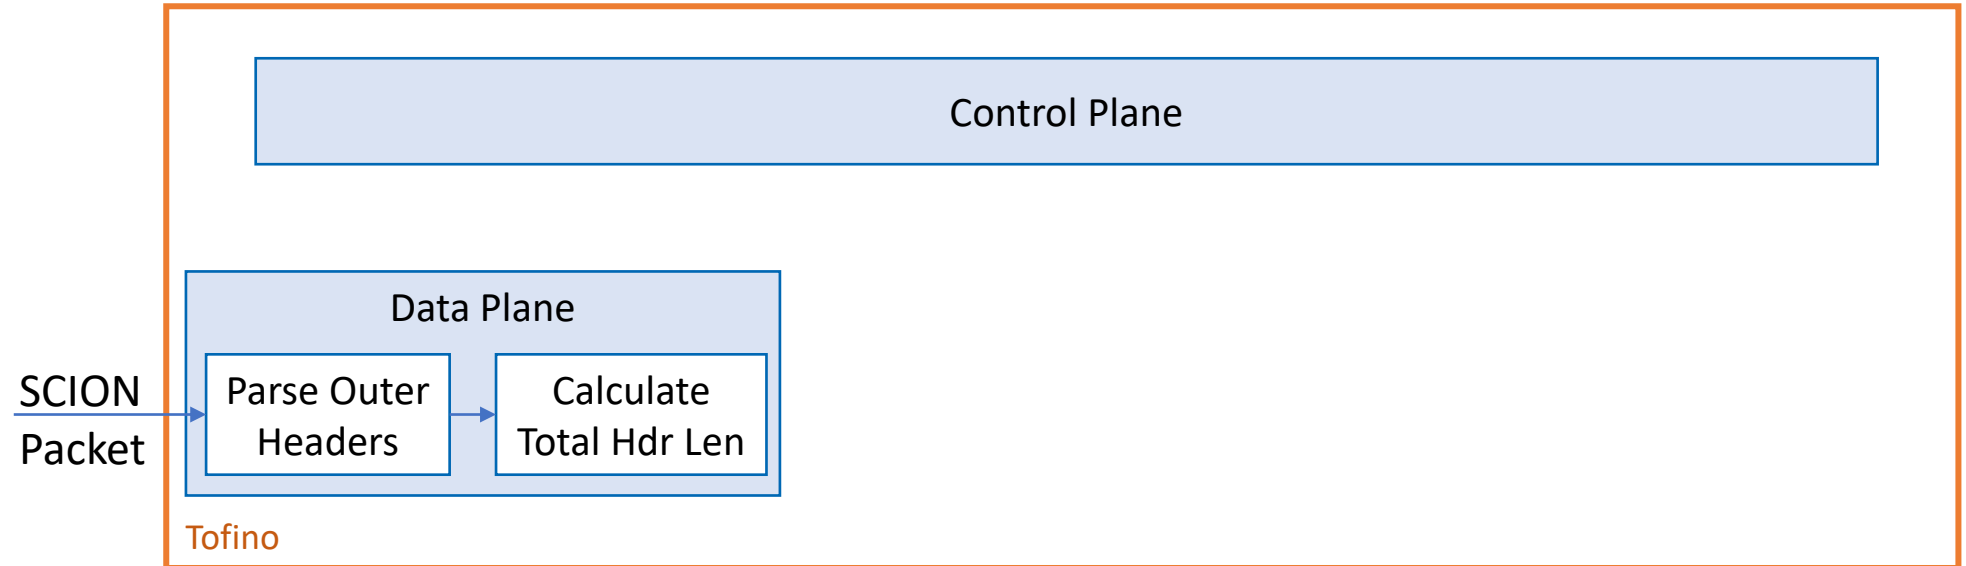
4. Implementation

- ❖ P4 implementation targeting Tofino 1 & Tofino 2
- ❖ Translate IP to SCION



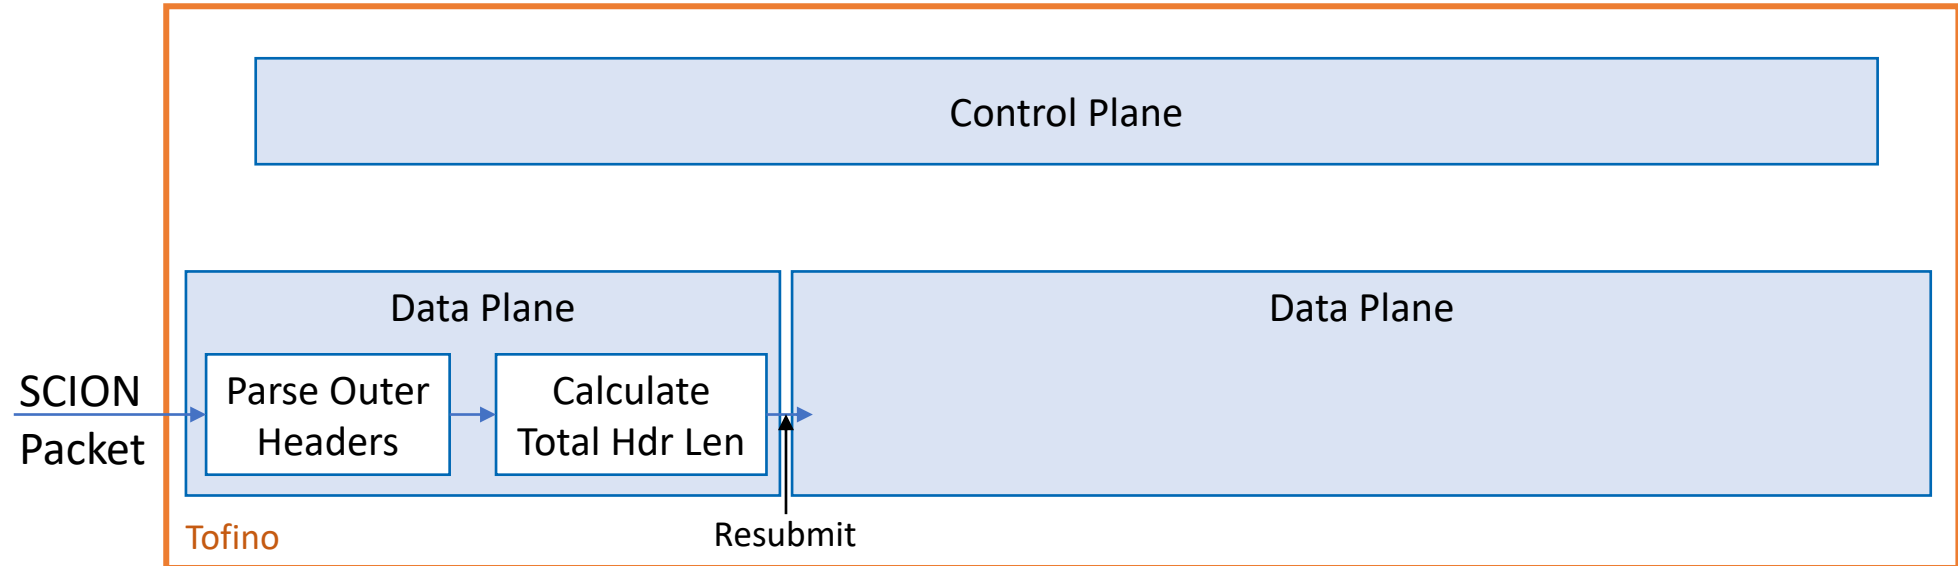
4. Implementation

- ❖ P4 implementation targeting Tofino 1 & Tofino 2
- ❖ Translate SCION to IP



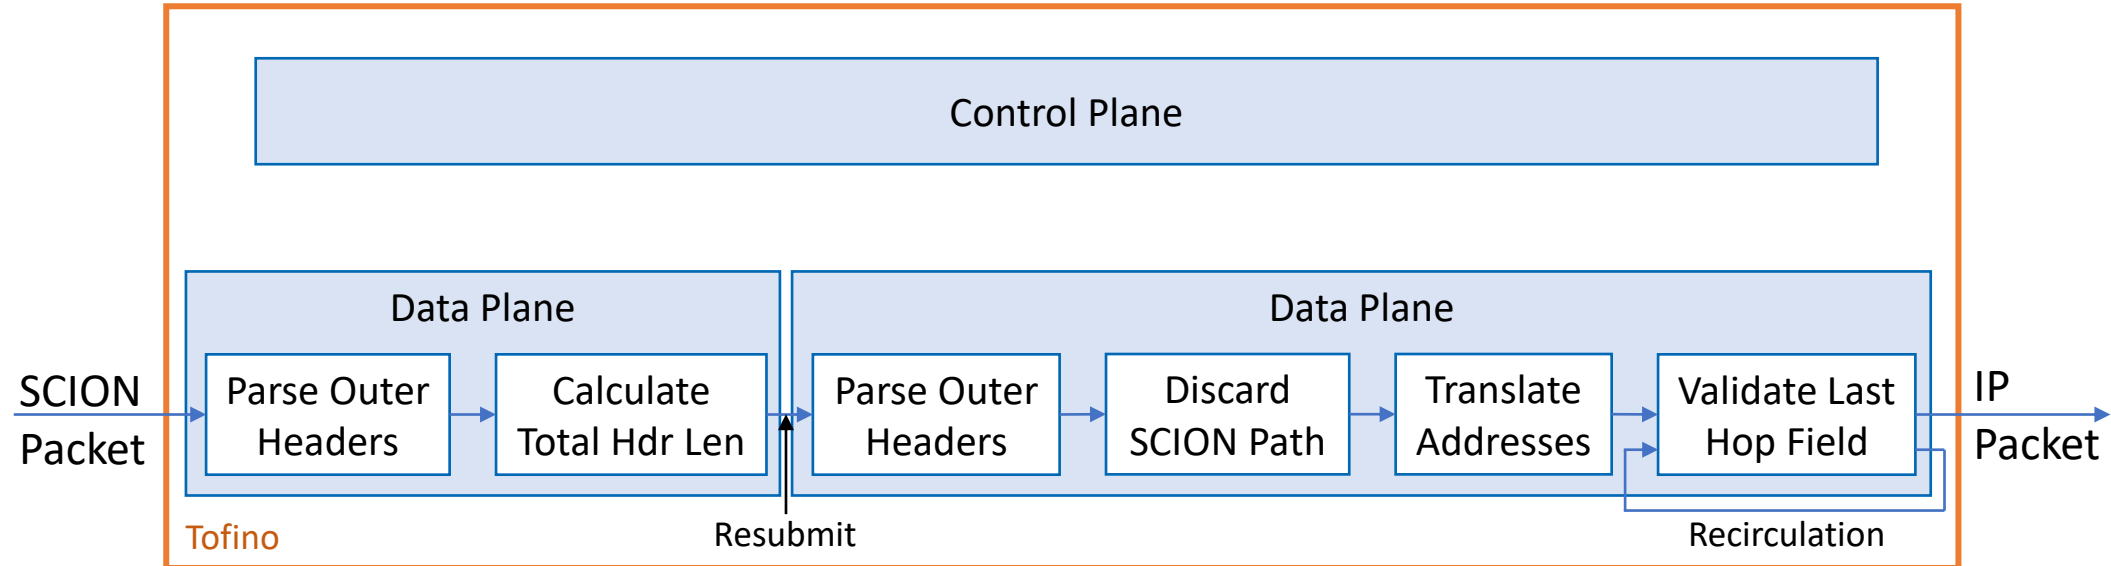
4. Implementation

- ❖ P4 implementation targeting Tofino 1 & Tofino 2
- ❖ Translate SCION to IP



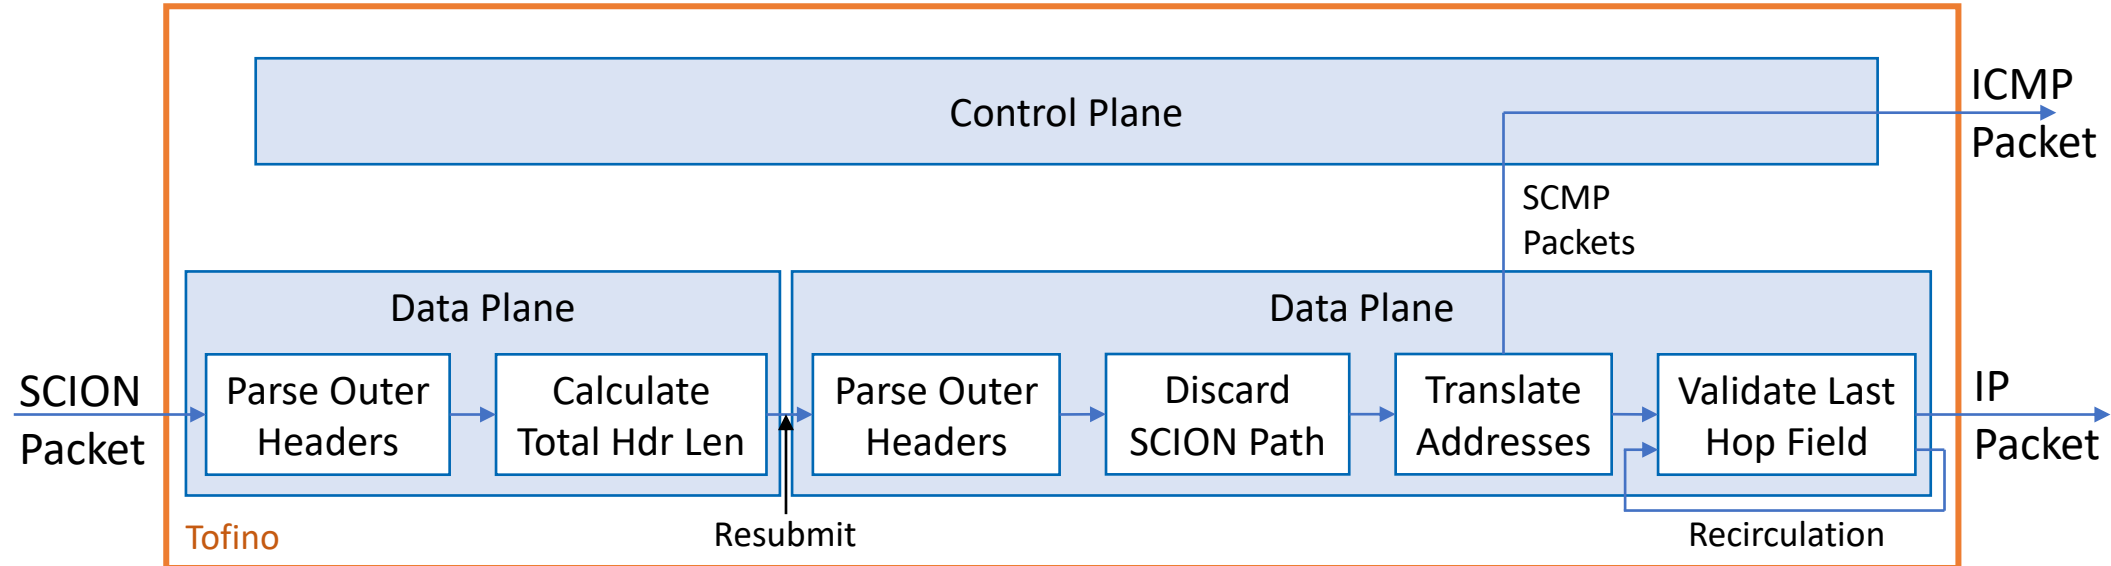
4. Implementation

- ❖ P4 implementation targeting Tofino 1 & Tofino 2
- ❖ Translate SCION to IP



4. Implementation

- ❖ P4 implementation targeting Tofino 1 & Tofino 2
- ❖ Translate SCION to IP



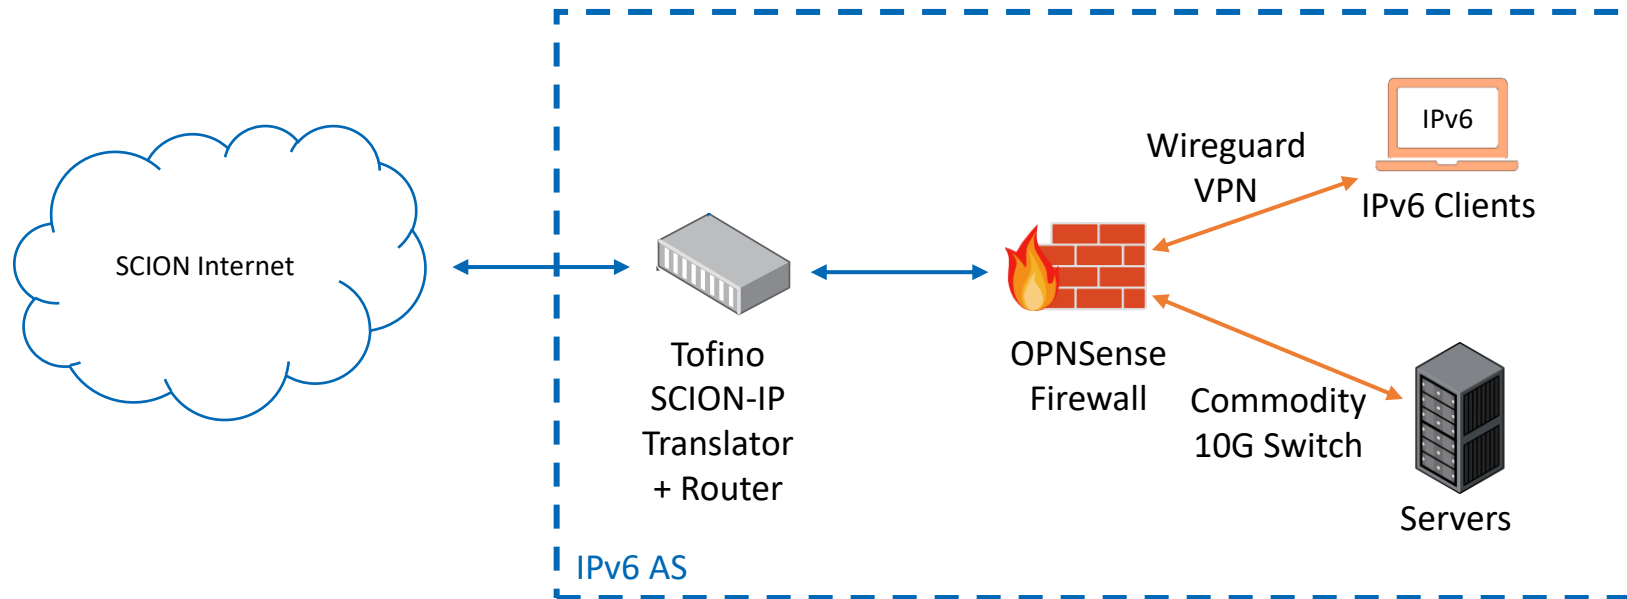
5. Deployment

- ❖ Resource Consumption in Tofino
- ❖ Only restricted by PHV registers, as every inserted hop field requires 12 byte of space

Switch Generation	Tofino 1	Tofino 2
Inserted hop fields per pass	16	24
Recirculations for MAC validation	5	2

5. Deployment

- ❖ Router is deployed in SCIERA (SCION Education, Research and Academic Network)



6. Conclusion and Future Work

The SCION-IP translator...

- ❖ Enables legacy IP applications over SCION
- ❖ Allows interconnection between IP and SCION implementations of an application
- ❖ Is easier to implement than the SIG (or rewriting applications to SCION)

We plan to...

- ❖ Improve path selection to specific applications to improve connections

Thanks for your attention!

E-Mails: lschulz@ovgu.de
robin.wehner@ovgu.de