RIPE

How to Secure Routing Header for Segment Routing?

Eric Vyncke, Distinguished Engineer
evyncke@cisco.com
@evyncke

Agenda

- Special use case
- Security of Routing Header & RFC 5095
- Segment Routing Security
- Packets with Extension Headers are lost?

Special Use Case

SR-IPv6 Capable Service Chaining

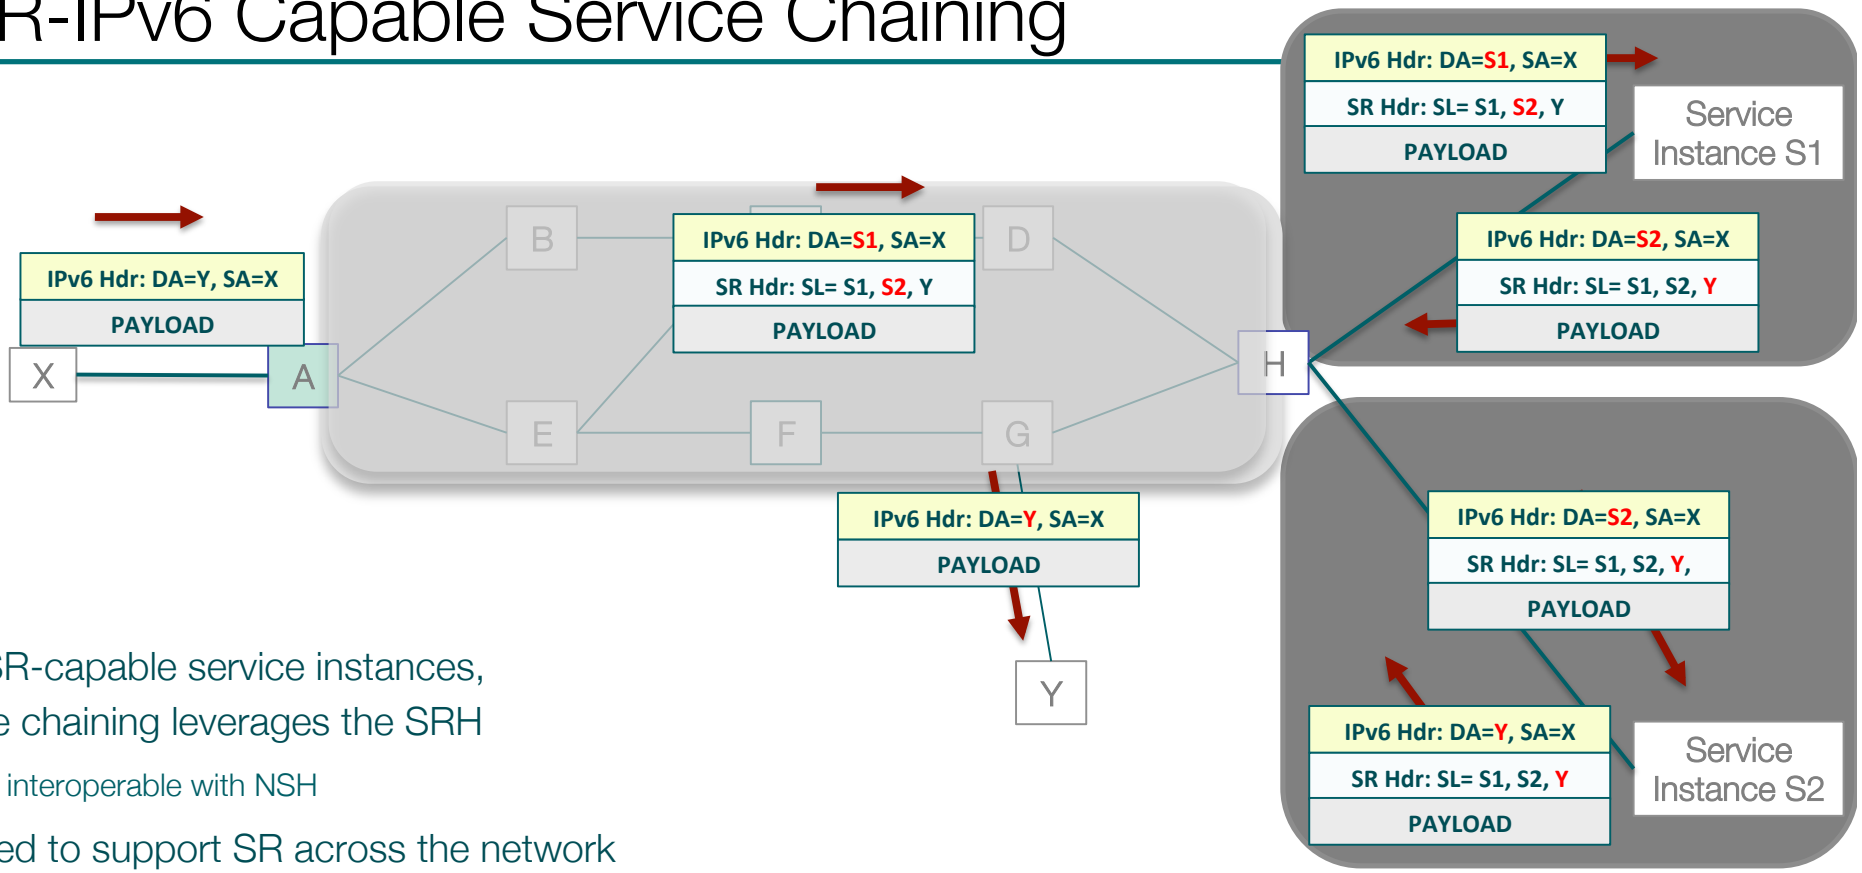

- With SR-capable service instances, service chaining leverages the SRH
  - Still interoperable with NSH
- No need to support SR across the network
  - Transparent to network infrastructure
- Next Step: allow SR service chaining with non-SR applications…
  - Work in progress

“Extreme Traffic Engineering” from CPE/Set-up Box?

- What about mobile node away from SP network?

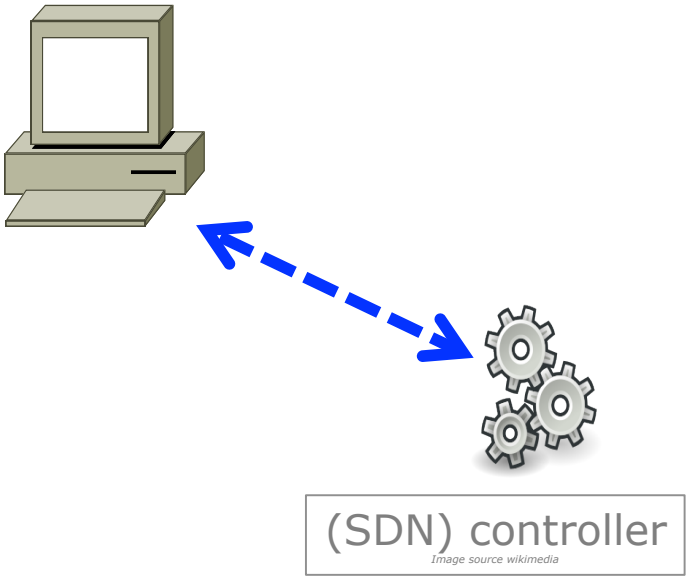

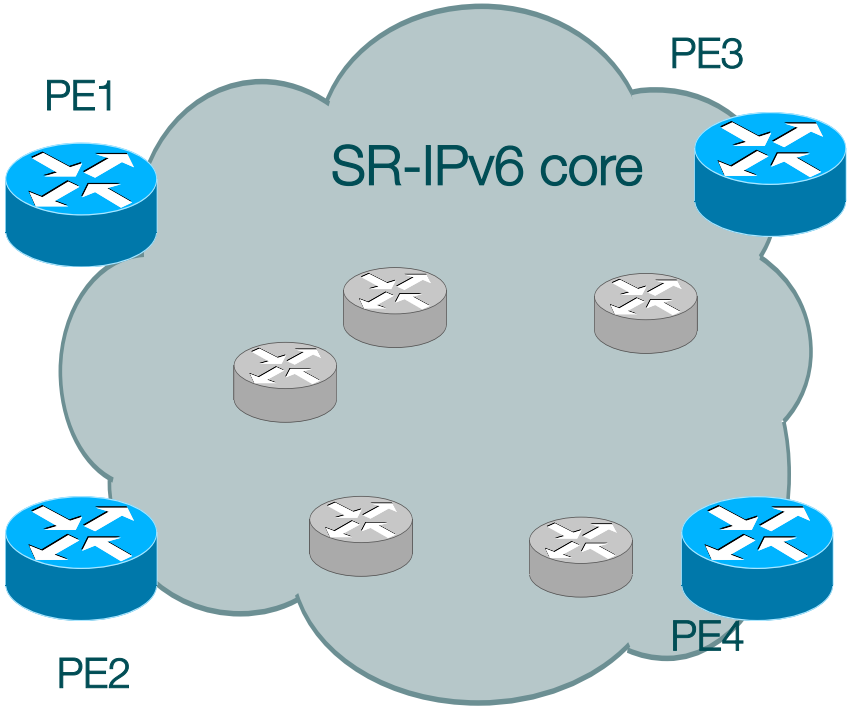

“Extreme Traffic Engineering” from CPE/Set-up Box?

- What about mobile node away from SP network?

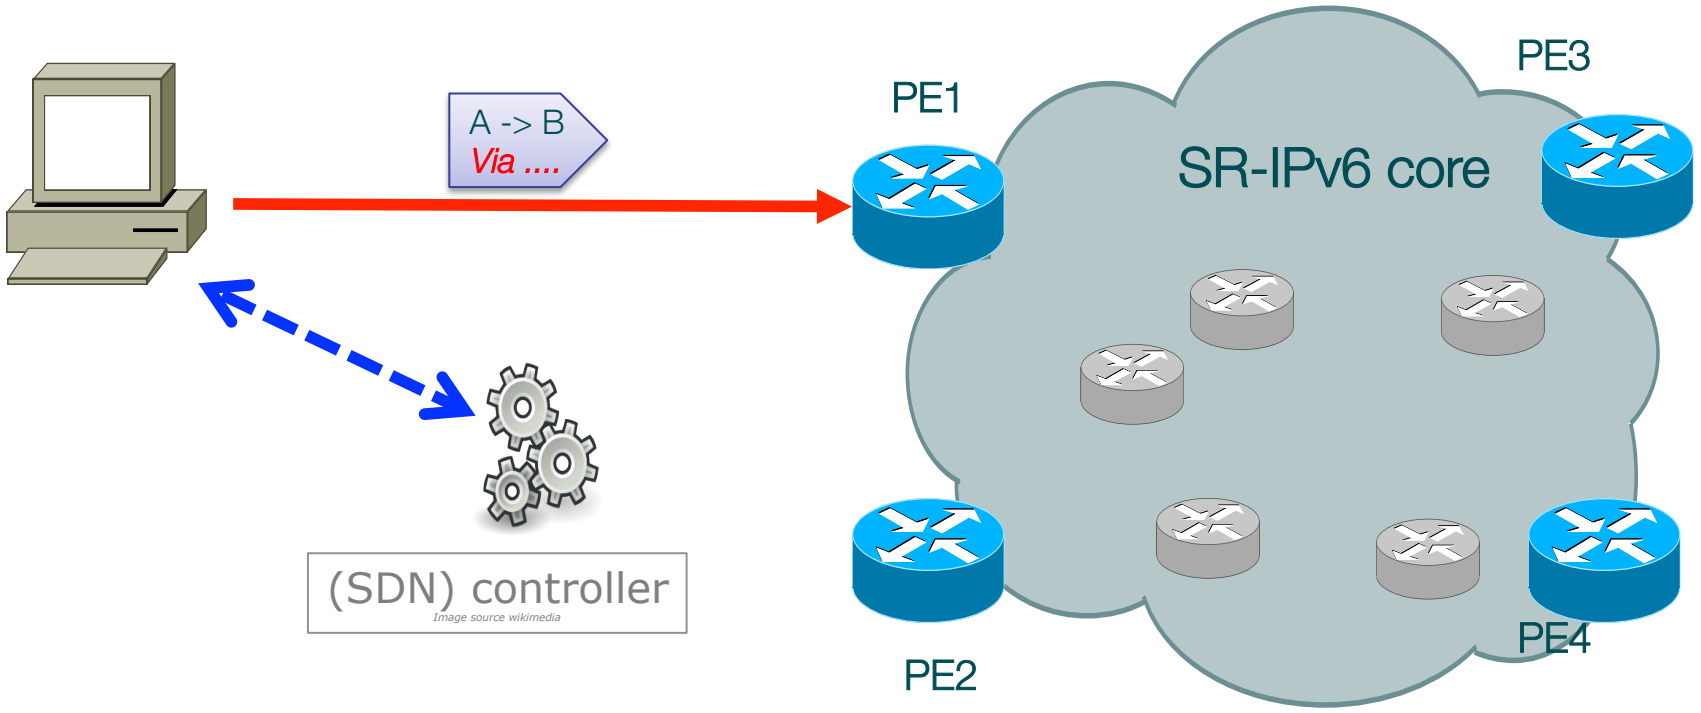

“Extreme Traffic Engineering” from CPE/Set-up Box?

- What about mobile node away from SP network?

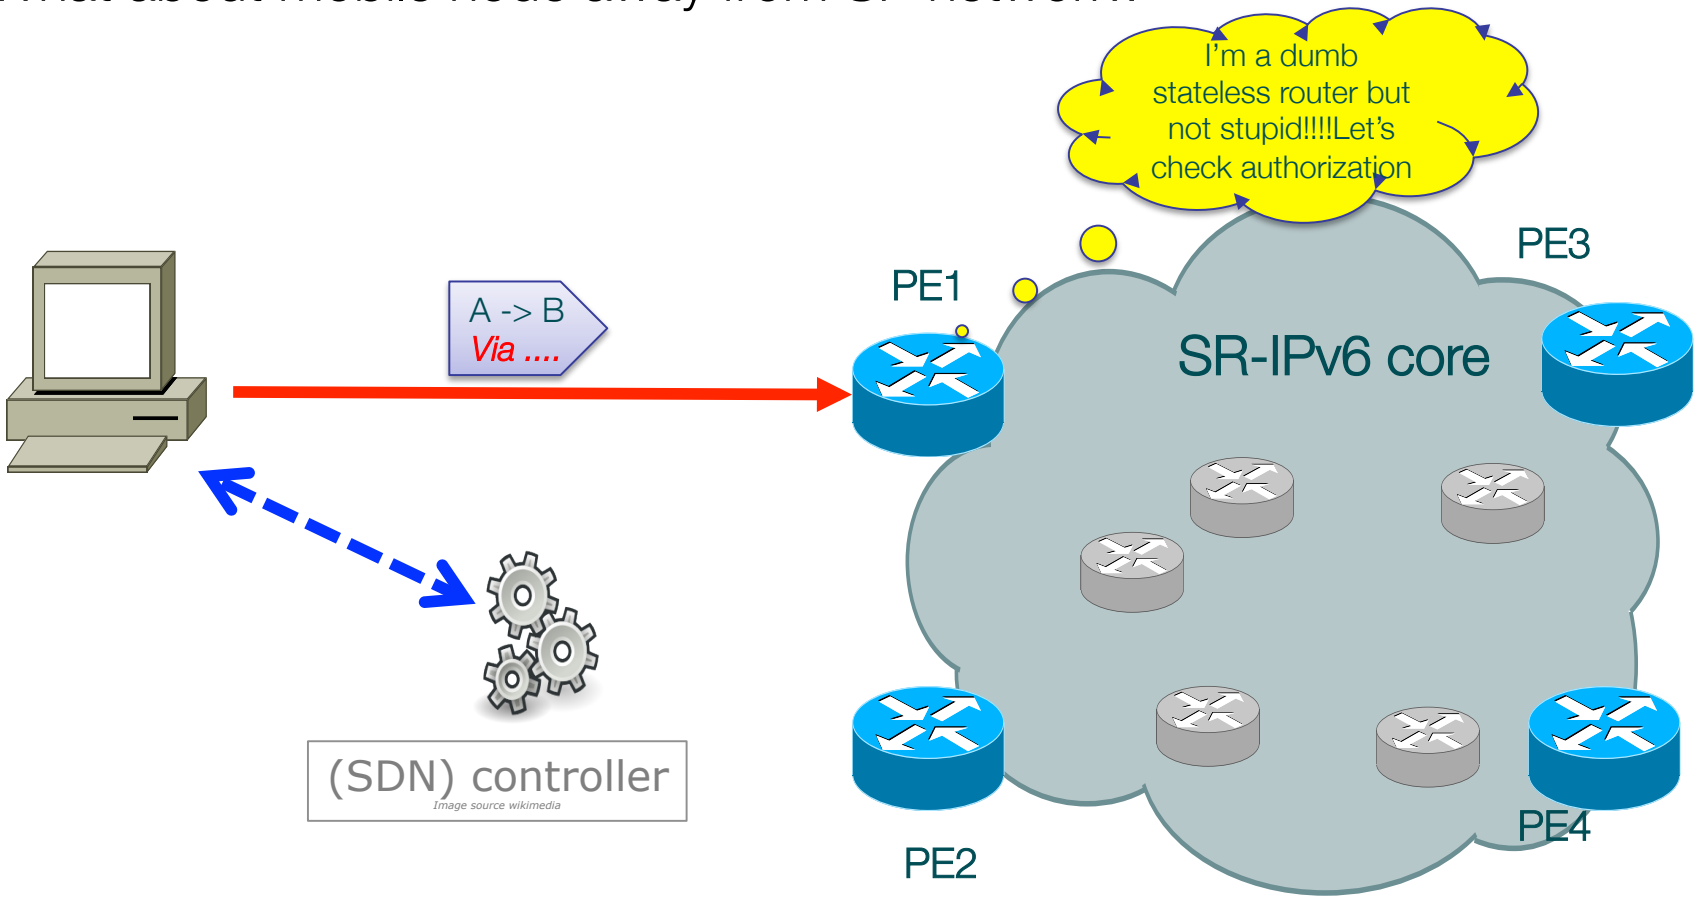

“Extreme Traffic Engineering” from CPE/Set-up Box?

- What about mobile node away from SP network?

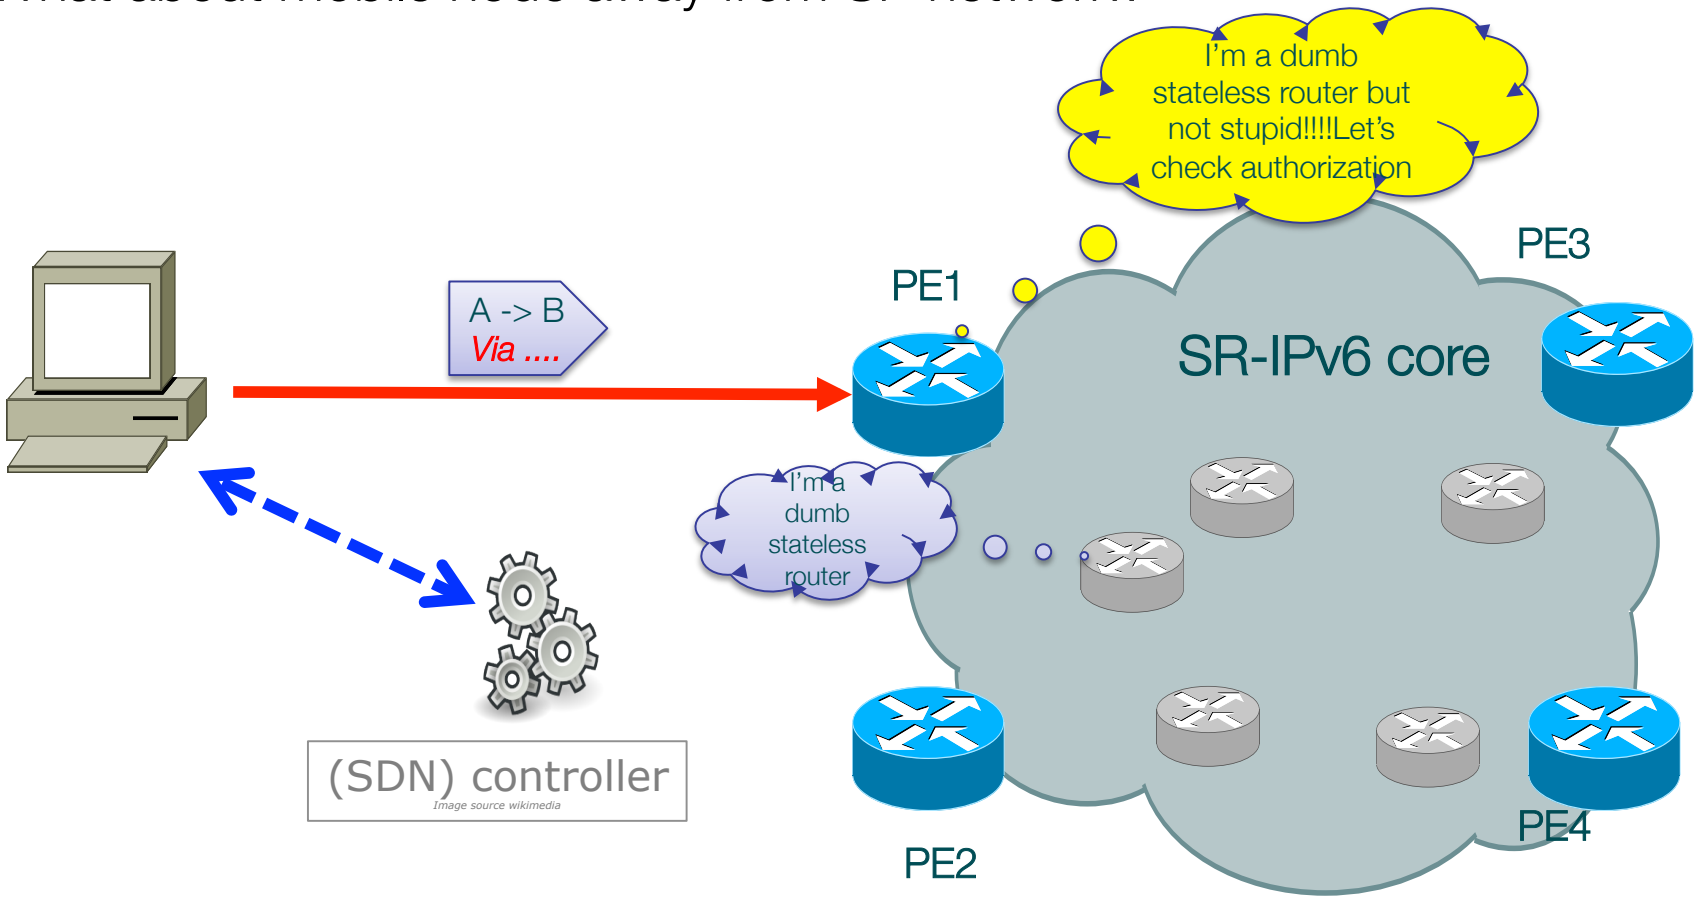

“Extreme Traffic Engineering” from CPE/Set-up Box?

- What about mobile node away from SP network?

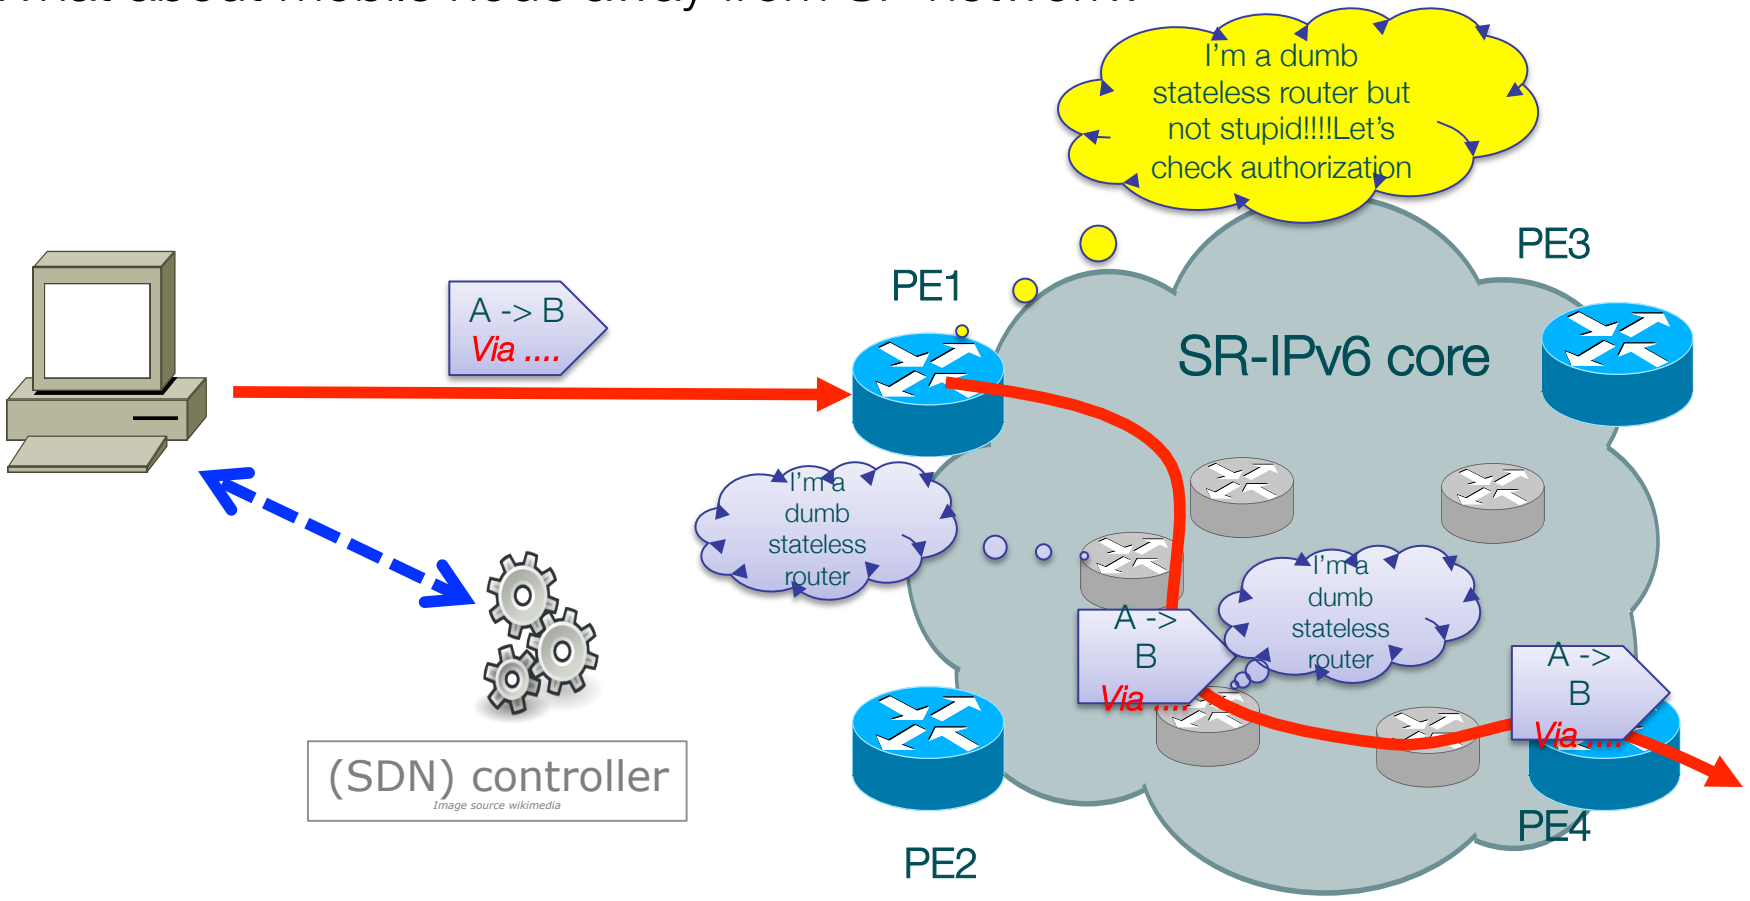

Huh??? Source Routing Security? What about RFC 5095?

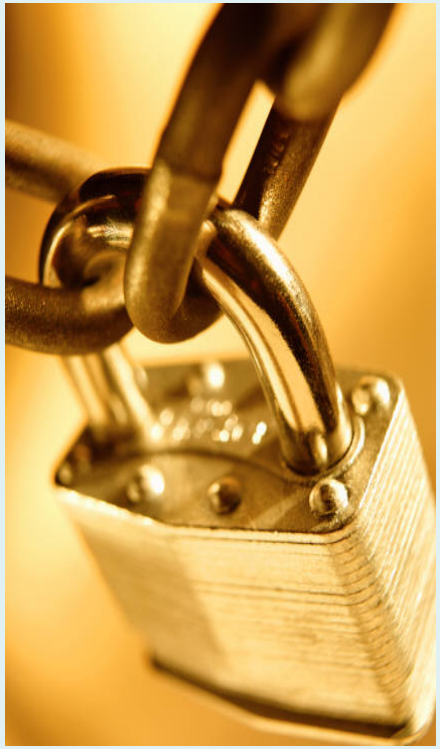

IPv6 Routing Header

- An extension header, processed by intermediate routers
- Three types
  - Type 0: similar to IPv4 source routing (multiple intermediate routers)
  - Type 2: used for mobile IPv6
  - Type 3: RPL (Routing Protocol for Low-Power and Lossy Networks)

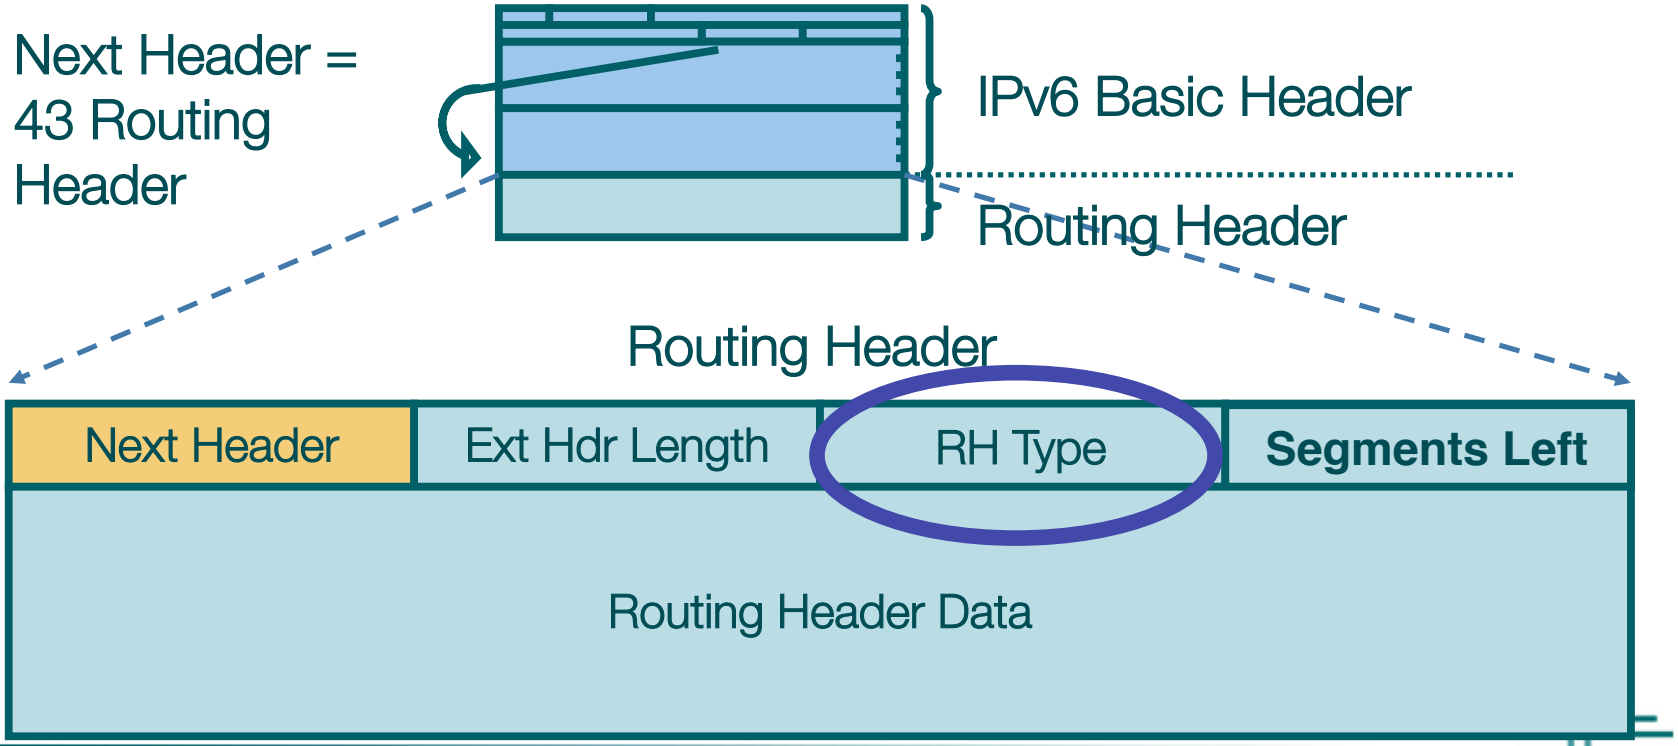

RH0: Amplification Attack

- What if attacker sends a packet with RH containing
  - A -> B -> A -> B -> A -> B -> A -> B -> A ....
- Packet will loop multiple time on the link A-B
- An amplification attack!

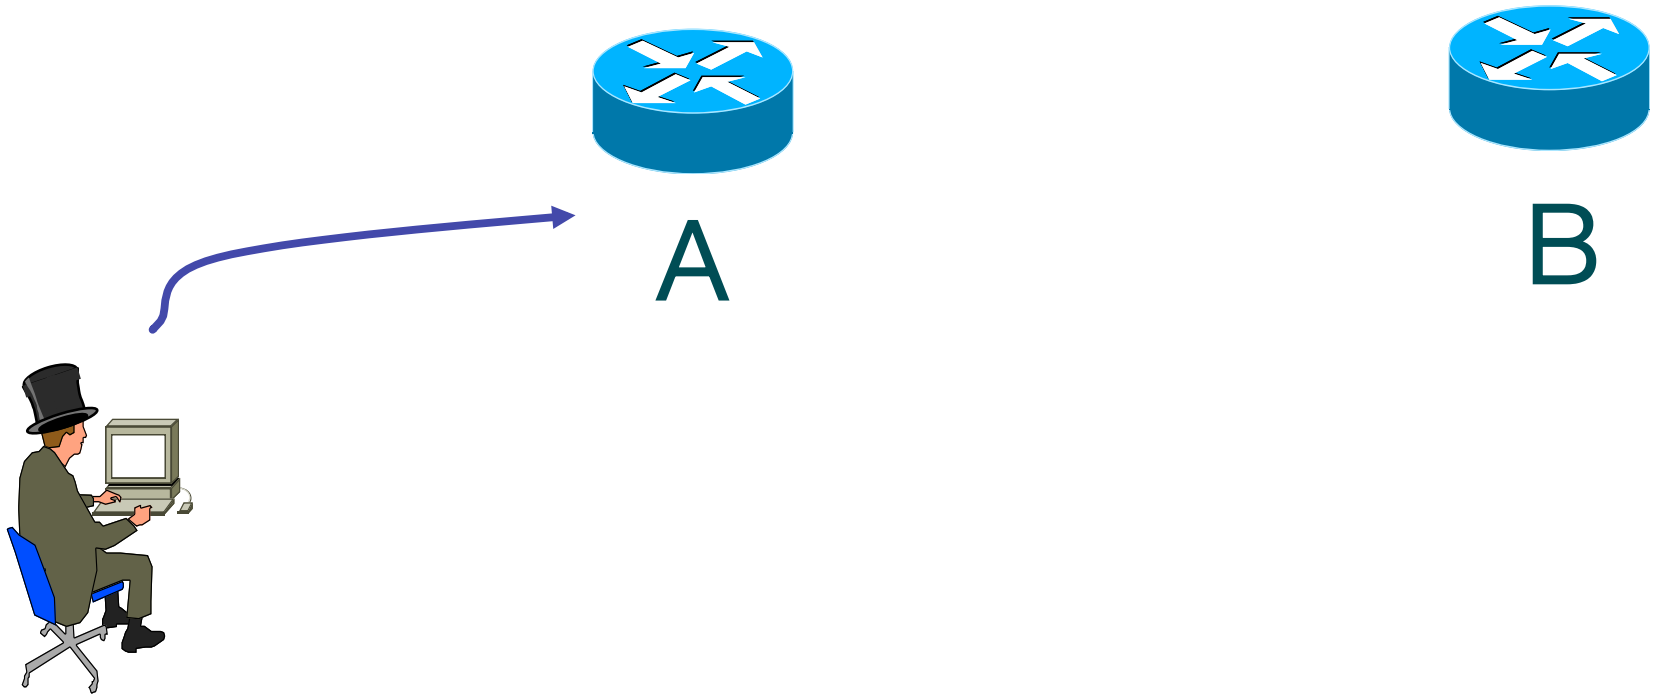

RH0: Amplification Attack

- What if attacker sends a packet with RH containing
  - A -> B -> A -> B -> A -> B -> A -> B -> A ....
- Packet will loop multiple time on the link A-B
- An amplification attack!

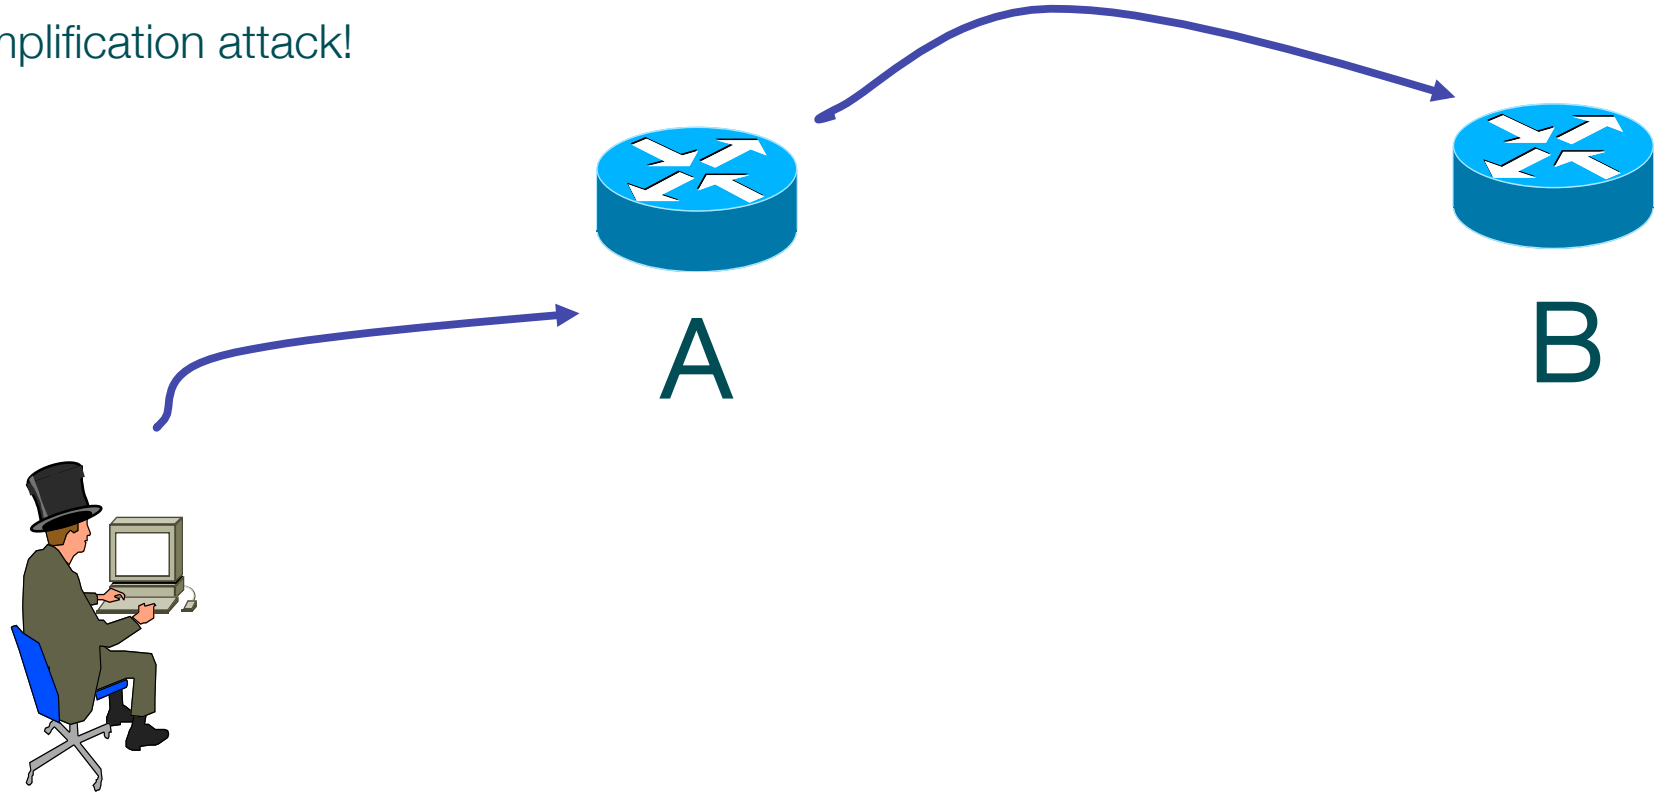

RH0: Amplification Attack

- What if attacker sends a packet with RH containing
  - A -> B -> A -> B -> A -> B -> A -> B -> A ....
- Packet will loop multiple time on the link A-B
- An amplification attack!

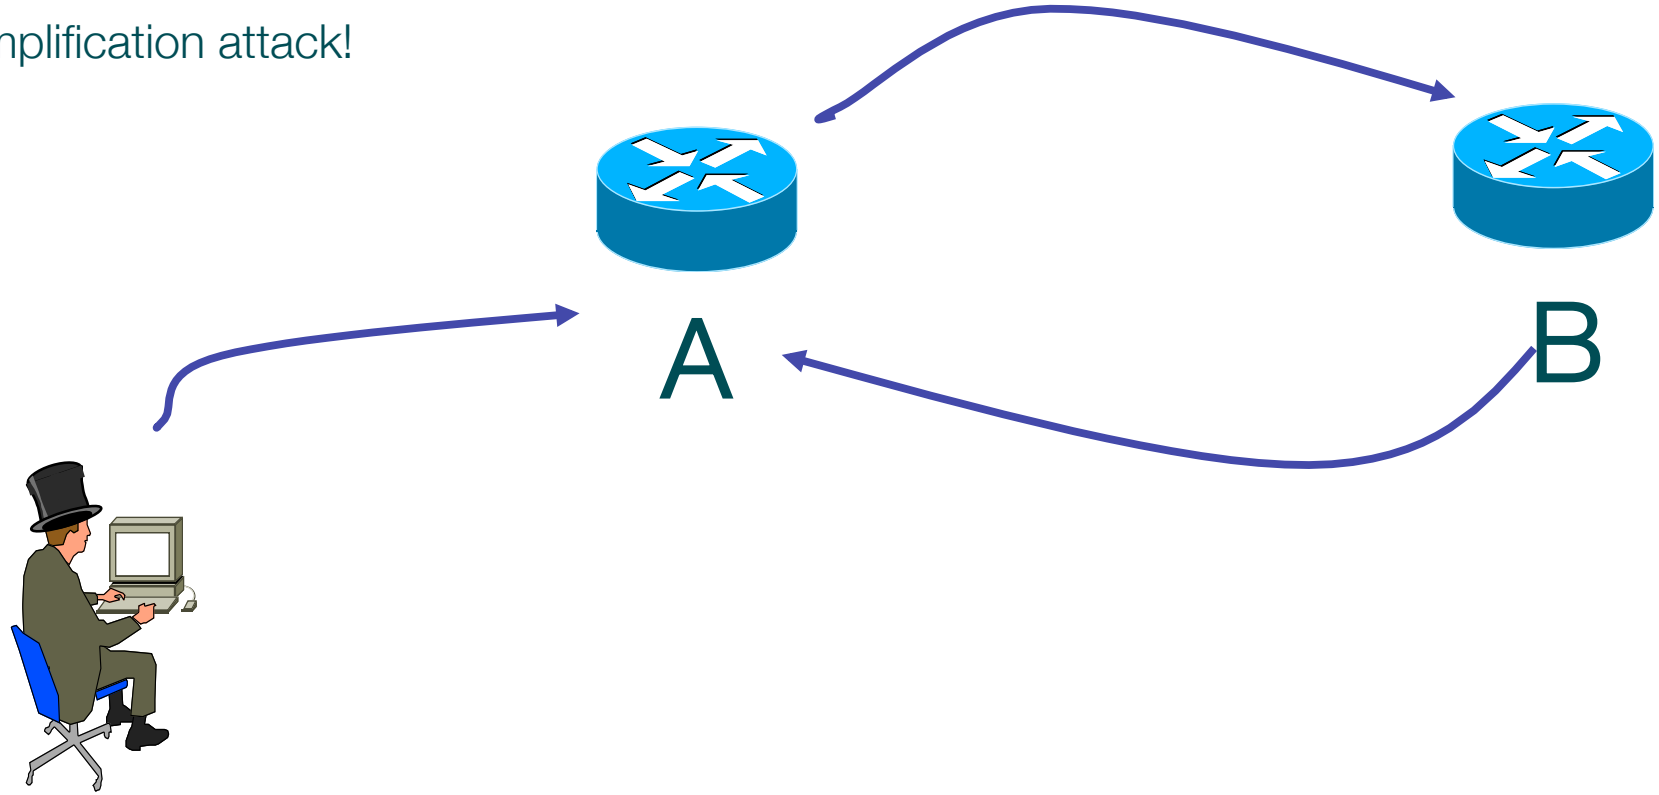

RH0: Amplification Attack

- What if attacker sends a packet with RH containing
  - A -> B -> A -> B -> A -> B -> A -> B -> A ....
- Packet will loop multiple time on the link A-B
- An amplification attack!

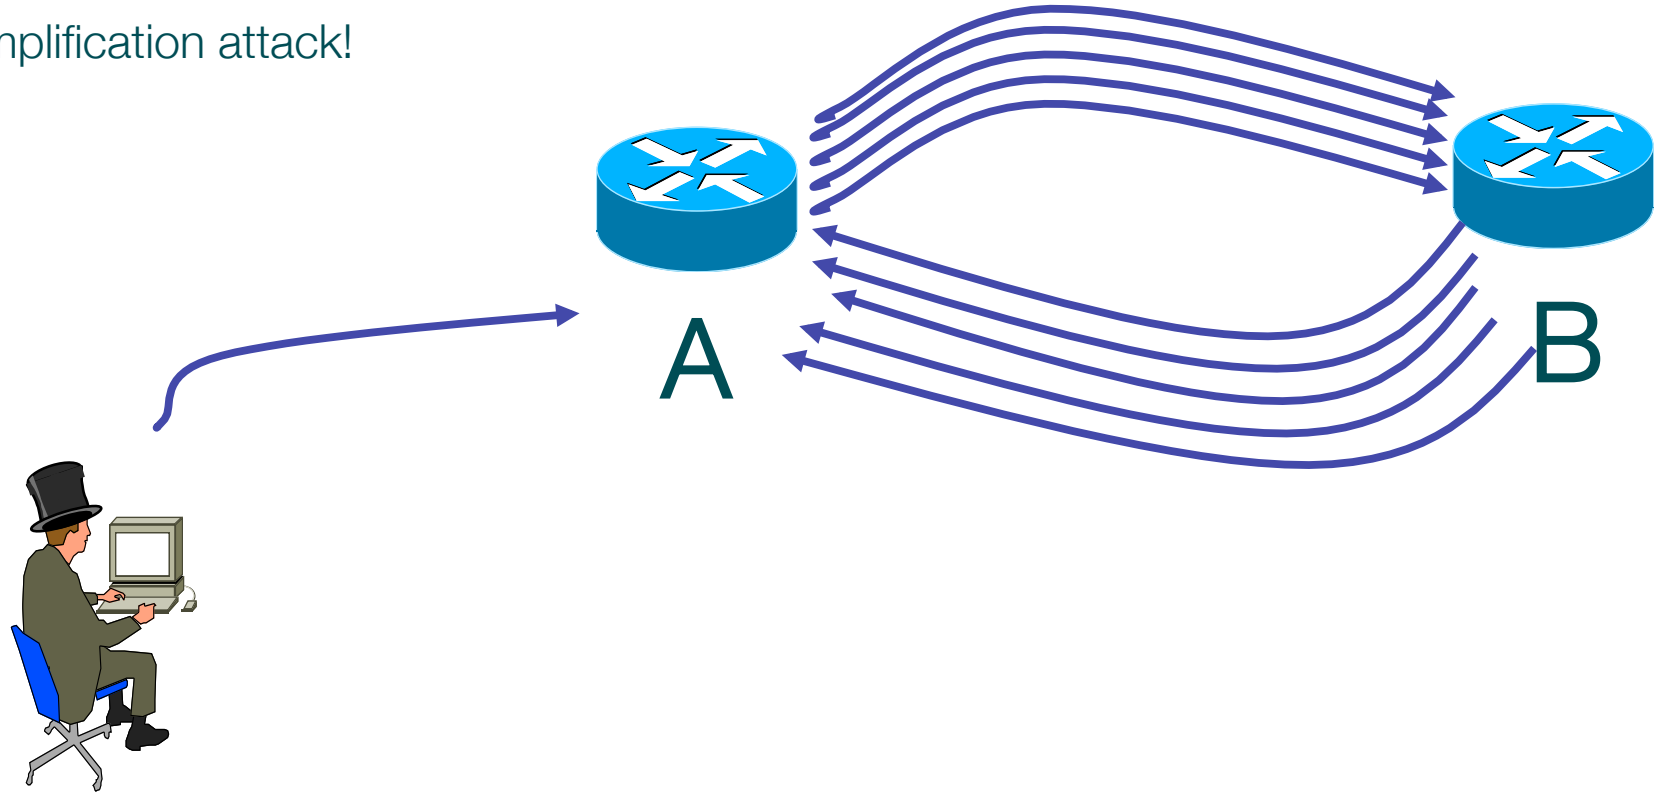

What RFC 5095 Says

J. Abley
Afilias
P. Savola
CSC/FUNET
G. Neville-
Neil
Neville-Neil Consulting

December 2007

Deprecation of Type 0 Routing Headers in IPv6
RFC 5095

"The severity of this threat is considered to be sufficient to warrant deprecation of RH0 entirely. A side effect is that this also eliminates benign RH0 use-cases; however, such applications may be facilitated by future Routing Header specifications."

Type 1: NIMROD

- A 1994 project funded by DARPA
  - Mobility
  - Hierarchy of routing (kind of LISP)
- Type 1 was deprecated in 2009
  - not because of security
  - but project was defunct and AFAIK not a single NIMROD packet was sent over IPv6...

Source: Clipartpanda.com

IPv6 Type 2 Routing Header: no problem

- Rebound/amplification attacks impossible
  - Only one intermediate router: the mobile node home address

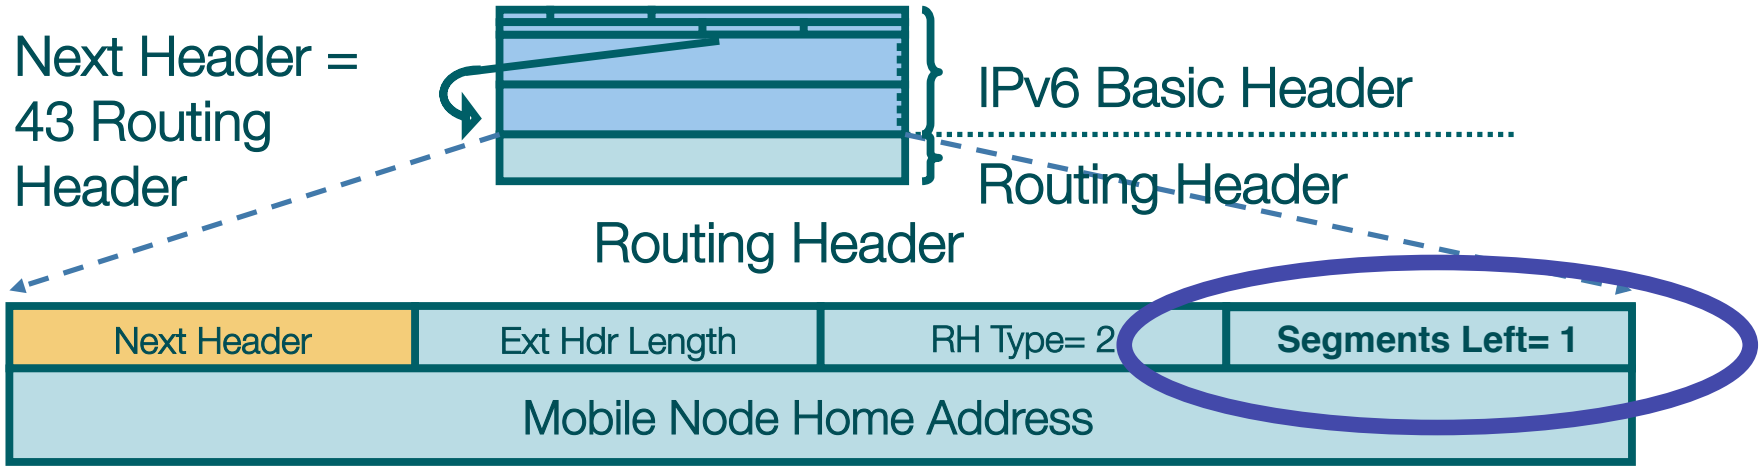

RH-3 for RPL: no problem

- Used by Routing Protocol for Low-Power and Lossy Networks
- But only within a single trusted network (strong authentication of node), never over a public untrusted network
  - Damage is limited to this RPL network
  - If attacker was inside the RPL network, then he/she could do more damage anyway

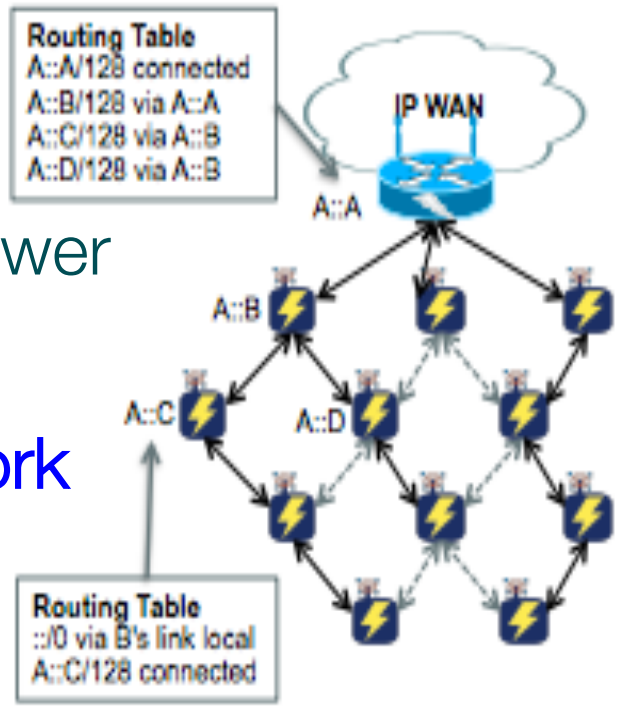

Segment Routing Security

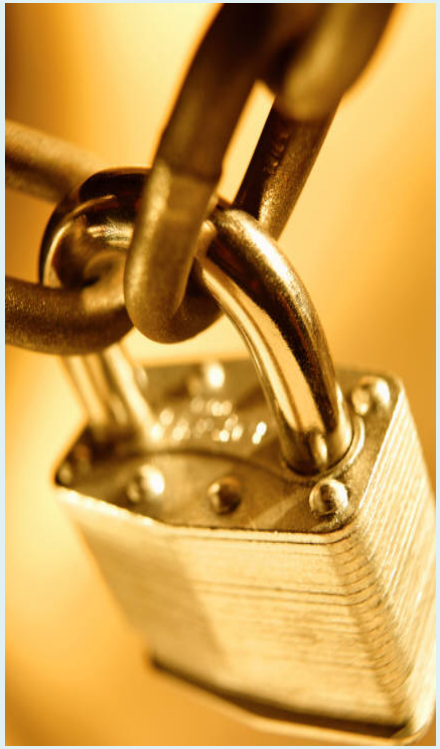

Segment Routing Security

- Addresses concerns of RFC5095
  - HMAC field to be used at ingress of a SR domain in order to validate/authorize the SRH
  - Inside SR domain, each node trust its brothers (RPL model)

- HMAC requires a shared secret (SDN & SR ingress routers)
  - Outside of current discussions
  - Pretty much similar to BGP session security or OSPFv3 security

SRH: identical to RFC 2460

- Next Header: 8-bit selector. Identifies the type of header immediately following the SRH
- Hdr Ext Len: 8-bit unsigned integer. Defines the length of the SRH header in 8-octet units, not including the first 8 octets
- Routing Type: TBD by IANA (SRH)
- Segment Left: index, in the Segment List, of the current active segment in the SRH. Decremented at each segment endpoint.

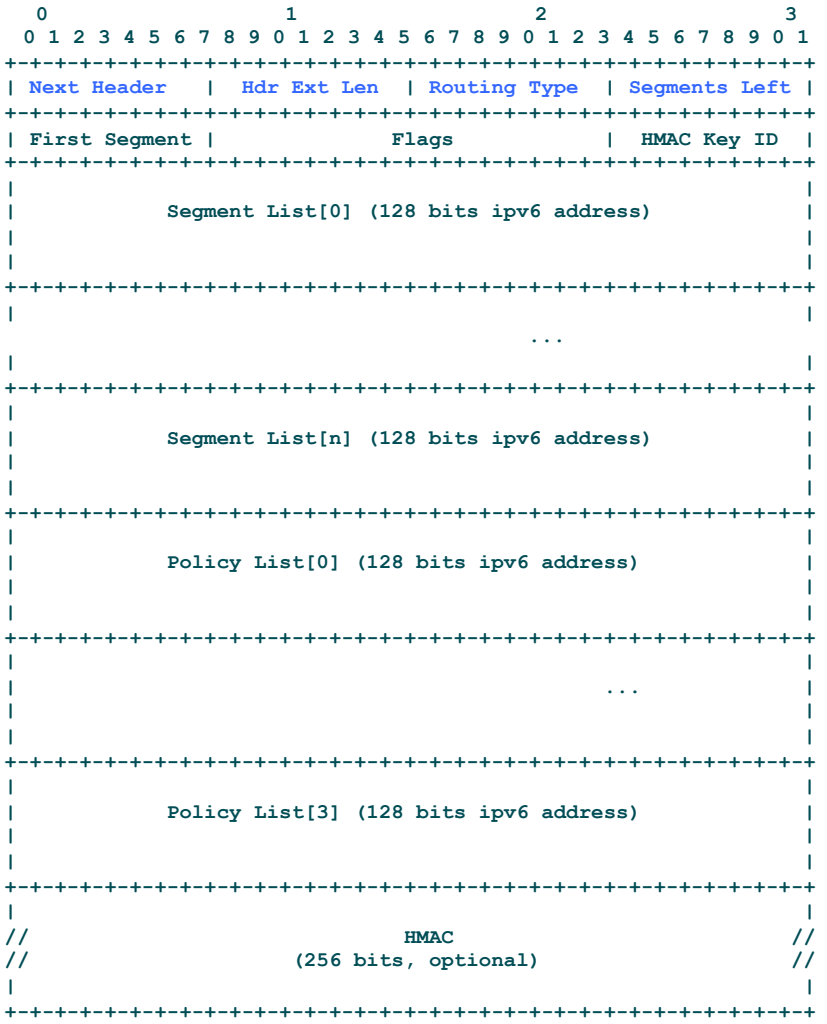

SRH: New

- First Segment: offset in the SRH, not including the first 8 octets and expressed in 16-octet units, pointing to the last element of the Segment List
- Flags:
  - bit-0: cleanup
  - bit-1: rerouted packet
  - bits 2 and 3: reserved
  - bits 4 to 15: policy flags
- Segment List[n]: 128 bit IPv6 addresses representing each segment of the path. The segment list is encoded in the reverse order of the path: the last segment is in the first position of the list and the first segment is in the last position
- Policy List[n] (optional): to mark ingress/ ingress SR address, to remember original source address

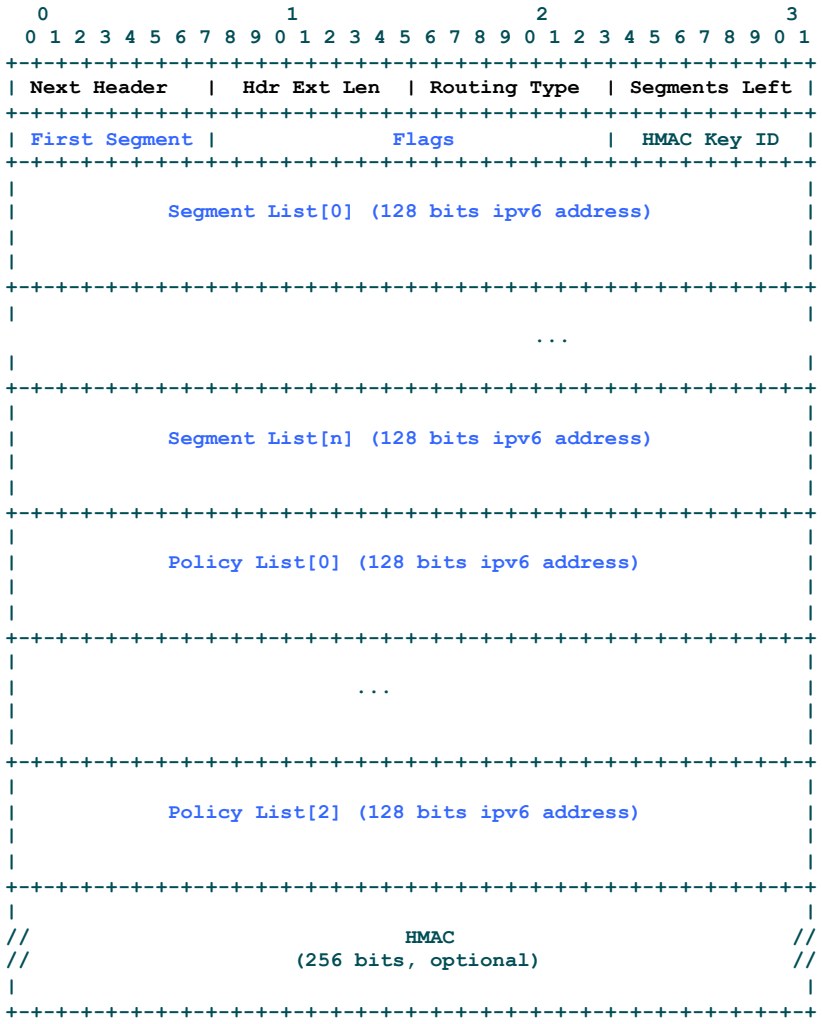

SRH: New for Security

- HMAC Key-id: identifies the shared secret used by HMAC. If 0, HMAC field is not present
- HMAC: SRH authentication (optional)

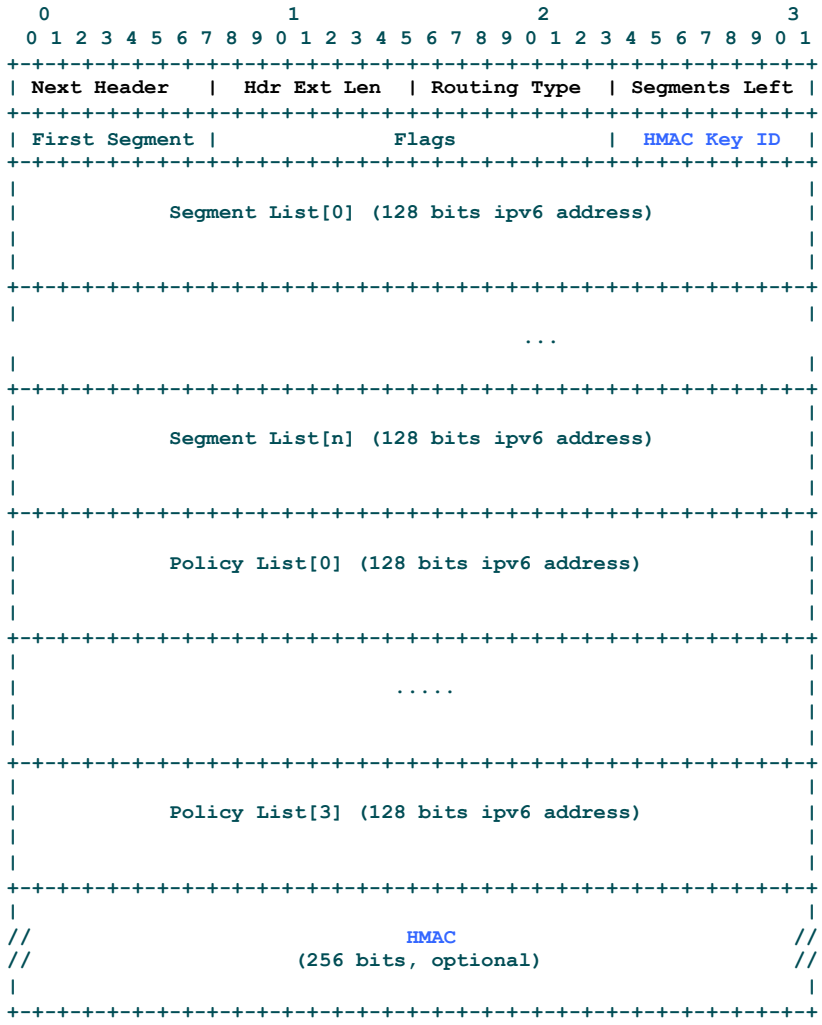

SRH: HMAC Coverage

- Source Address (not shown): as it is immutable and to prevent SR service stealing
- First Segment: offset in the SRH, not including the first 8 octets and expressed in 16-octet units, pointing to the last element of the Segment List
- Flags:
  - bit-0: cleanup
  - bit-1: rerouted packet
  - bits 2 and 3: reserved
  - bits 4 to 15: policy flags
- HMAC Key ID:
- Segment List[n]: all segments

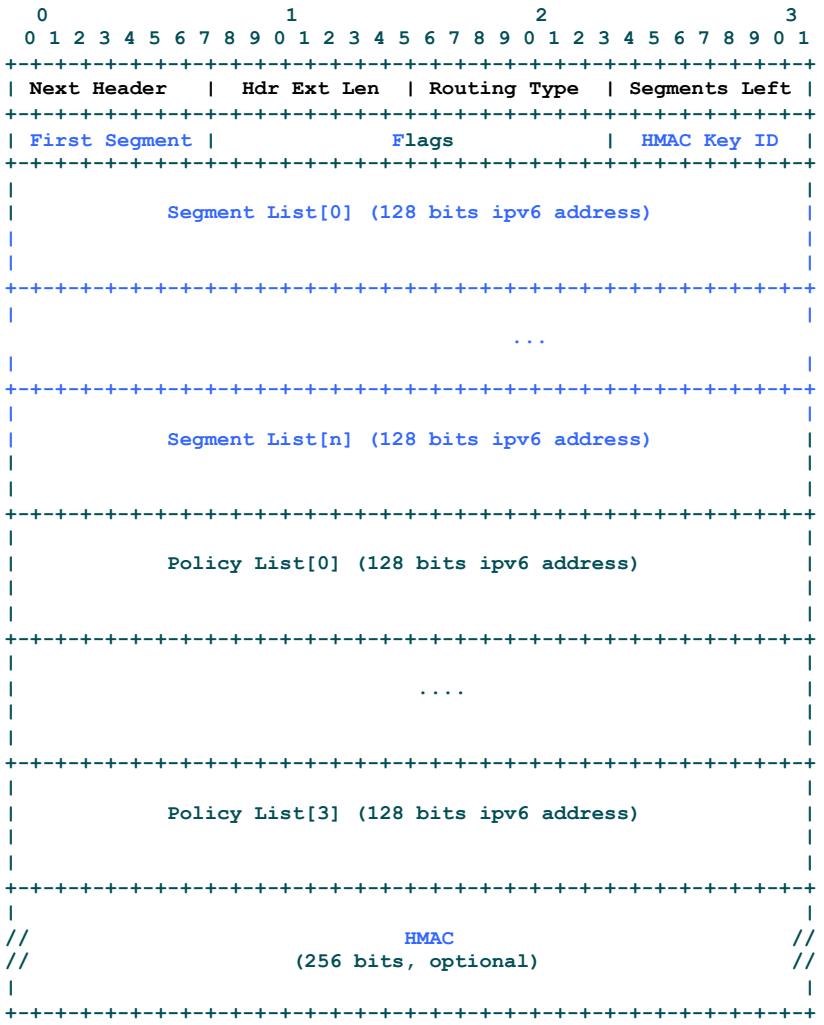

Lost of Packets with Extension Headers

Issue: Ext Hdr are dropped on the Internet

- draft-gont-v6ops-ipv6-ehs-in-real-world
  - About 20-40% of packets with Ext Hdr are dropped over the Internet
- SRH works only within one administrative domain
  - => not an issue as operator set the security/drop policy
- Test on your own: http://www.vyncke.org/sr.php
  - And let us know !

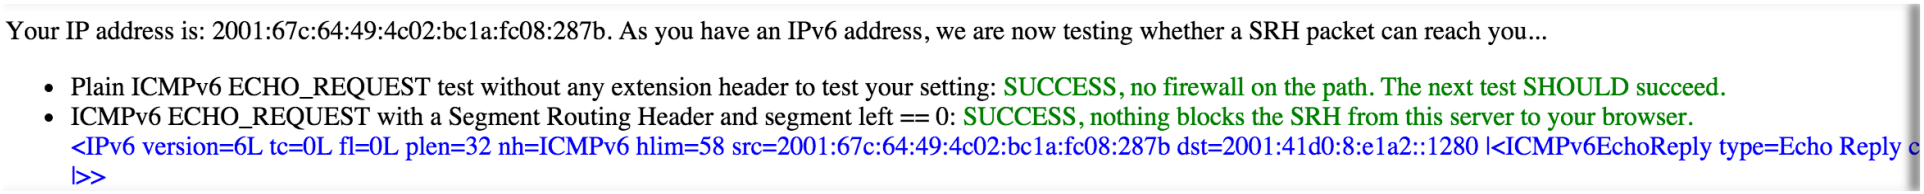
Your IP address is: 2001:67c:64:49:4c02:bc1a:fc08:287b. As you have an IPv6 address, we are now testing whether a SRH packet can reach you...

- Plain ICMPv6 ECHO_REQUEST test without any extension header to test your setting: SUCCESS, no firewall on the path. The next test SHOULD succeed.
- ICMPv6 ECHO_REQUEST with a Segment Routing Header and segment left == 0: SUCCESS, nothing blocks the SRH from this server to your browser.
  <IPv6 version=6L tc=0L fl=0L plen=32 nh=ICMPv6 hlim=58 src=2001:67c:64:49:4c02:bc1a:fc08:287b dst=2001:41d0:8:e1a2::1280 |<ICMPv6EchoReply type=Echo Reply c
  |>>

Another View of Packet Drops

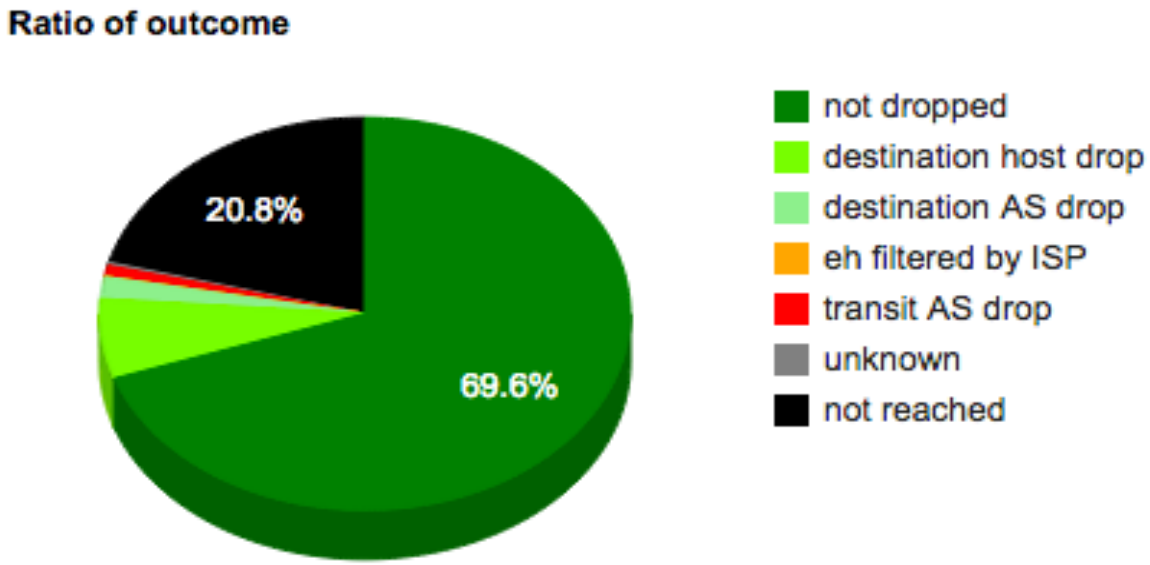

- Current research by Polytechnique Paris (Mehdi Kouhen) and Cisco (Eric Vyncke)
- And VM provided by Sander Steffann
- http://btv6.vyncke.org/exthdr/index.php?ds=bgp&t=rh4 (work in progress!)

Questions?