

IETF 119 – SPRING WG

Segment Routing over IPv6 (SRv6) Proof of Transit

draft-iannone-spring-srv6-pot-00


L. Iannone, A. Fressancourt

IETF 119 – Brisbane

Path Verification Problem Statement

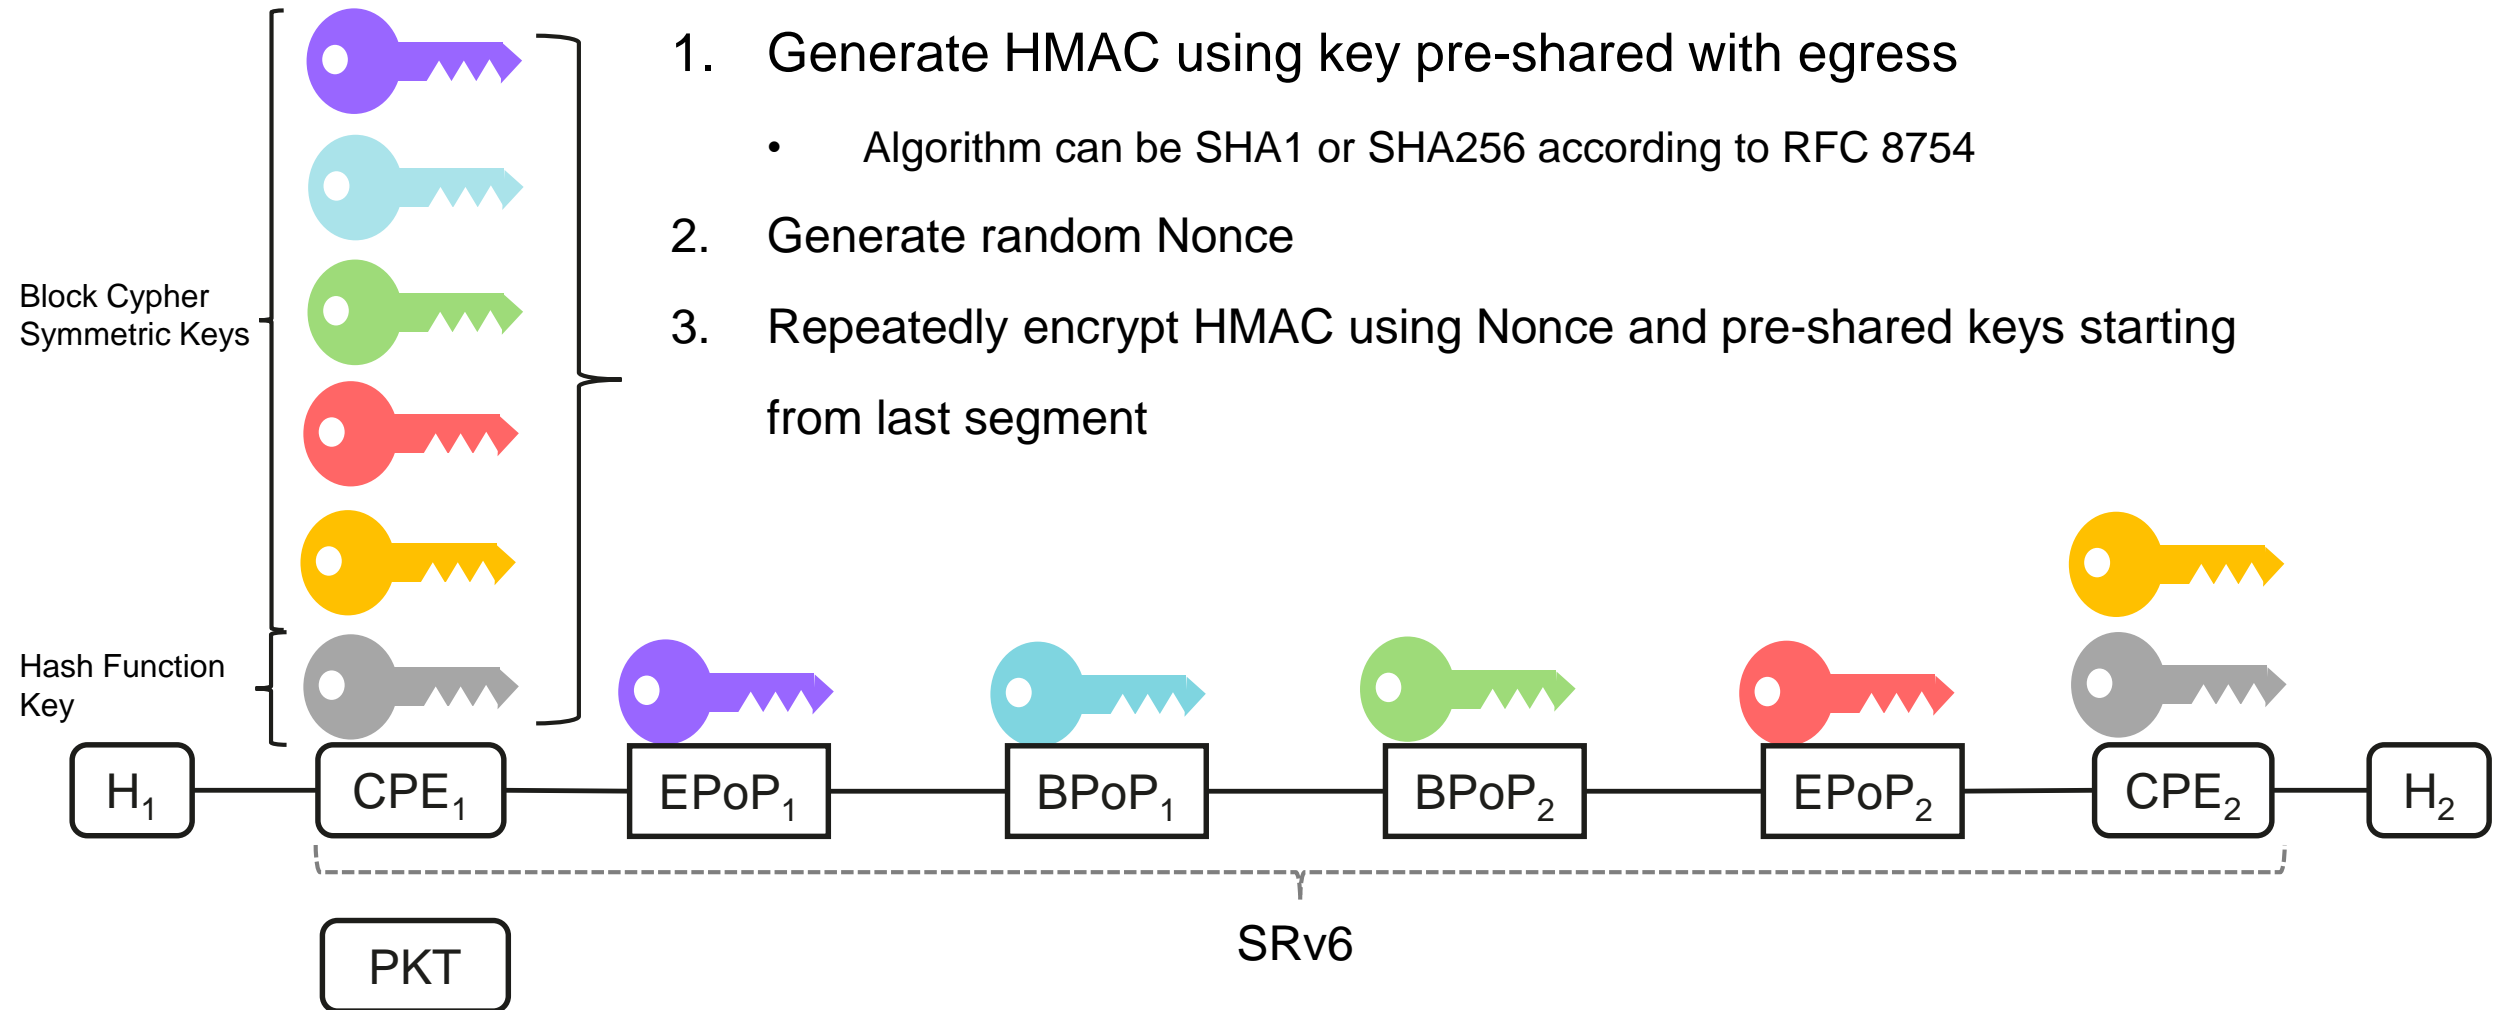
(or should we say segment verification...)

- Path Validation via Proof of Transit makes sure the path we chose is the actual path the traffic travels
- How to make sure that **DESIRED** path equals **ACTUAL** path?



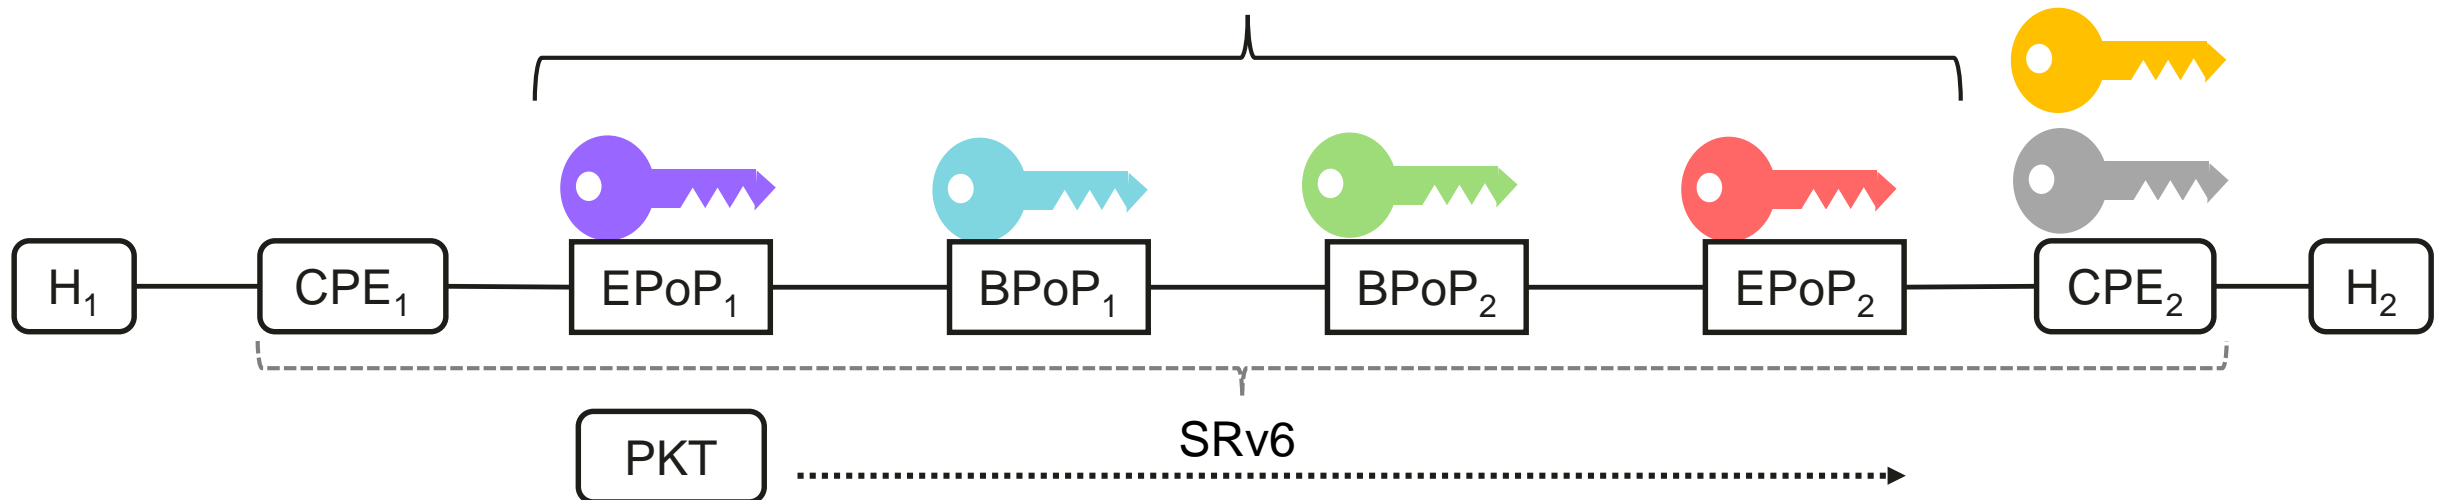
- Growing interest Proof of Transit (POT) for:
 - Service Function Chaining (SFC)
 - Workload identity
 - Traffic path compliance
 - uRPF validation
 - Segment Routing Security

POT in Practice: Ingress Node Operations



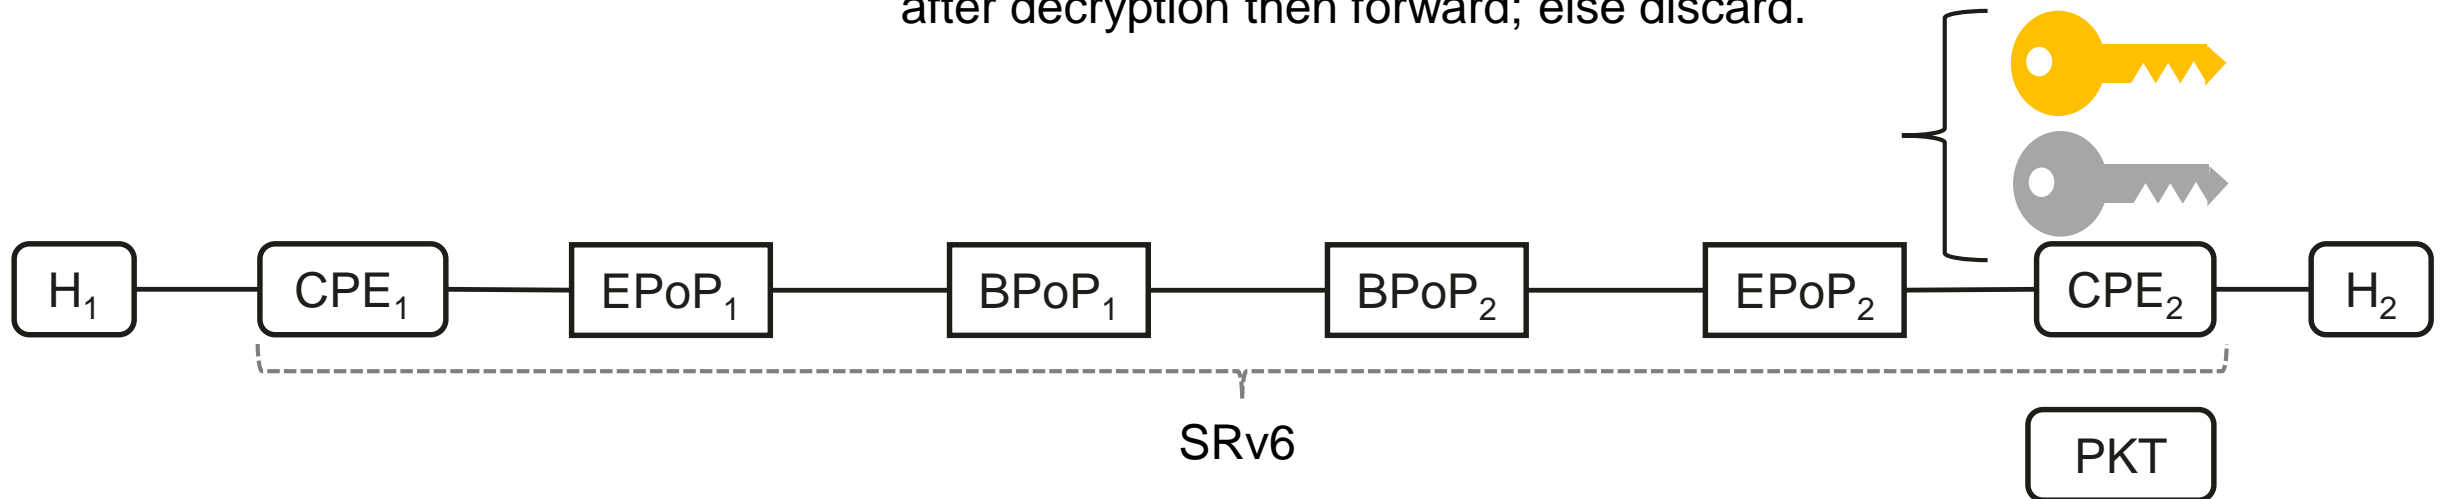
POT in Practice: Middle Segments Operations

- For each packet containing PoT TLV:
 1. Retrieve pre-shared key
 2. Decrypt HMAC once
 3. Forward Packet to next segment



POT in Practice: Last Segment Operations

- For each packet containing PoT TLV:
 1. Retrieve pre-shared key
 2. Decrypt HMAC once
 3. Compute packet HMAC
 - If newly compute HMAC == HMAC obtained from the packet after decryption then forward; else discard.



Recursive Cryptography as POT

1. Packet creation:

Given the set of keys

$$[k_1, k_2, k_3, k_4, k_D, k_{HSD}]$$

Identified by a **Key Set ID** and a randomly generated value **N**, the following Path Verification Factor (PVF) is computed:

$$PVF_{HSD} = HMAC_{k_{HSD}}(SID-list, \dots)$$

$$PVF_1 = Enc_{k_1}(Enc_{k_2}(Enc_{k_3}(Enc_{k_4}(Enc_{k_D}(PVF_{HSD}), N), N), N), N), N)$$

Then along with the SRH the following is sent:

$$M = Key\ Set\ ID\ | N\ | PVF_1$$

2. Packet processing

Each intermediate node i receives a packet

$$M = Key\ Set\ ID\ | N\ | PVF_{i-1}$$

and computes:

$$PVF_i = Dec_{k_i}(PVF_{i-1}, N)$$

Then it sends:

$$M = Key\ Set\ ID\ | N\ | PVF_i$$

3. Packet verification:

The destination receives a packet:

$$M = Key\ Set\ ID\ | N\ | PVF_{i-1,t}$$


and computes:

$$PVF_{HSD} = Dec_{k_D}(PVF_{D-1}, N)$$

To verify the HMAC, it computes:

$$PVF_{HSD}' = HMAC_{k_{HSD}}(SID-list, \dots)$$

If $PVF_{HSD} = PVF_{HSD}'$, then SID-list and the path has been verified



Proof of Transit Segment Routing Header TLV

0								1								2								3							
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
Type								Length								Reserved (8 bits)								Nonce Length (8 bits)							
Key Set ID (32 bits)																															
Nonce (variable length)																															
Encrypted HMAC (variable length)																															

- Improvements compared to SRv6 HMAC:
 - Sequential encryption and decryption of the HMAC → verification that the routers listed in the SID list have been crossed
- Drawbacks compared to SRv6 HMAC:
 - Intermediate routers are unable to verify the SID list's integrity.

Linux Kernel Implementation

- WSL Kernel 6.1.21.1
 - Header (TLV) Operations
 - Crypto: AES/XTS
 - Sysctl API
 - CLI Configuration
 - Support in iproute2 (v6.6.0)

```
#include <linux/types.h>
#include <linux/seg6.h>

#define SEG6_PVAL_FIELD_LEN 32


#define SEG6_PVAL_ALGO_NAME_MAX_LENTH 64
#define SEG6_PVAL_SECRET_MAX_LENTH 64
#define SEG6_HMAC_SECRET_MAX_LENTH 64

/* These should go in
 * include/uapi/linux/seg6.h
```

```
ggx@DebianNoVPN:~/linux-kernel$ sudo sysctl net.ipv6 |grep seg6
net.ipv6.conf.all.seg6_enabled = 0
net.ipv6.conf.all.seg6_require_hmac = 0
net.ipv6.conf.all.seg6_require_pval = 0
net.ipv6.conf.default.seg6_enabled = 0
net.ipv6.conf.default.seg6_require_hmac = 0
net.ipv6.conf.default.seg6_require_pval = 0
net.ipv6.conf.docker0.seg6_enabled = 0
net.ipv6.conf.docker0.seg6_require_hmac = 0
net.ipv6.conf.docker0.seg6_require_pval = 0
net.ipv6.conf.eth0.seg6_enabled = 0
net.ipv6.conf.eth0.seg6_require_hmac = 0
net.ipv6.conf.eth0.seg6_require_pval = 0
net.ipv6.conf.lo.seg6_enabled = 0
```

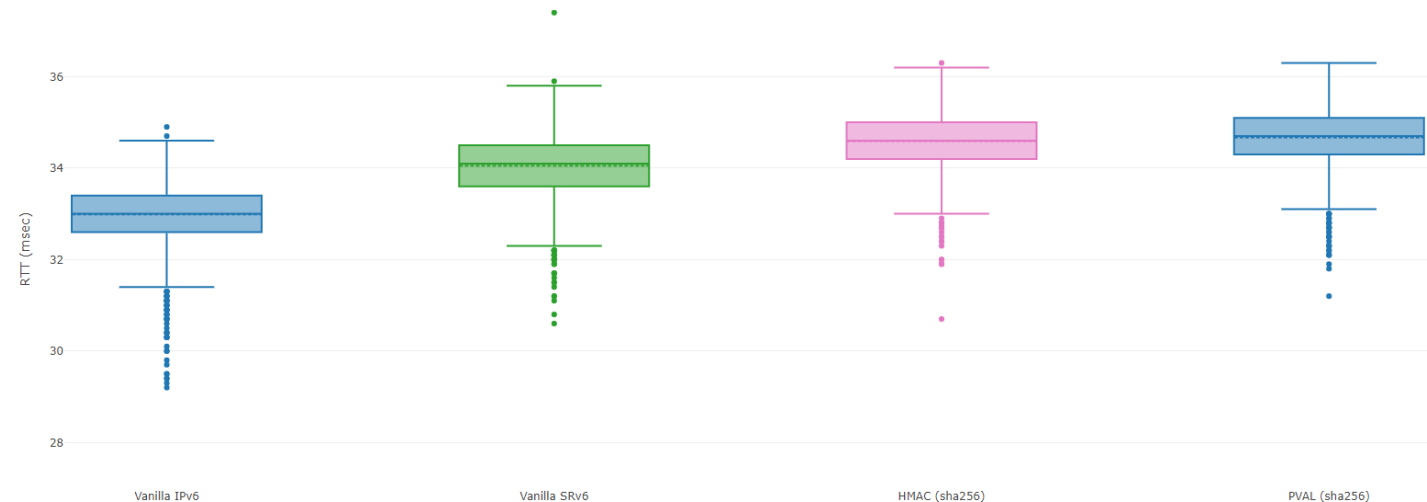
Sysctl API

```
ggx@DebianNoVPN:~/linux-kernel/iproute2-main$ sudo ./ip/ip sr pval set 338 aes
Enter secrets for Path Validation (blank to delete):
ggx@DebianNoVPN:~/linux-kernel/iproute2-main$ sudo ./ip/ip sr pval set 666 chacha20
Enter secrets for Path Validation (blank to delete):
ggx@DebianNoVPN:~/linux-kernel/iproute2-main$ sudo ./ip/ip sr pval show
pval: id 338 algo aes secret "thisisatest"
pval: id 666 algo chacha20 secret "thisisasecret"
```



Conclusion

- Initial evaluation quite promising
 - Cost in terms of delay similar to Segment-by-Segment HMAC
 - (in the specific setup)



- Happy to discuss more and show you a demo!

