Export of Segment Routing IPv6 Information in IPFIX

draft-tgraf-opsawg-ipfix-srv6-srh

Enabling insights in SRv6 forwarding plane
by adding Segment Routing dimensions

thomas.graf@swisscom.com
benoit.claise@huawei.com
16. March 2022

MPLS-SR @ IPFIX

Adressed with RFC 9160 @ OPSAWG

- In MPLS-SR the data-plane is still the same as in MPLS. **Only the routing protocol providing the label changes.**
- IE70 mplsTopLabelStackSection is the top label FEC used to forward. Each following label in the label stack is **decomposed in IE71-79 separately.**

```
 0                   1                   2
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                Label                  | Exp |S|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

Label:  Label Value, 20 bits
Exp:    Experimental Use, 3 bits
S:      Bottom of Stack, 1 bit
```

- IE47 mplsTopLabelIPv4Address is the top label IP address where the traffic is forwarded to.
- IE46 mplsTopLabelType describes from which routing protocol the top label IP address and label is coming from. **Updated with RFC 9160 to cover MPLS-SR routing protocols.**

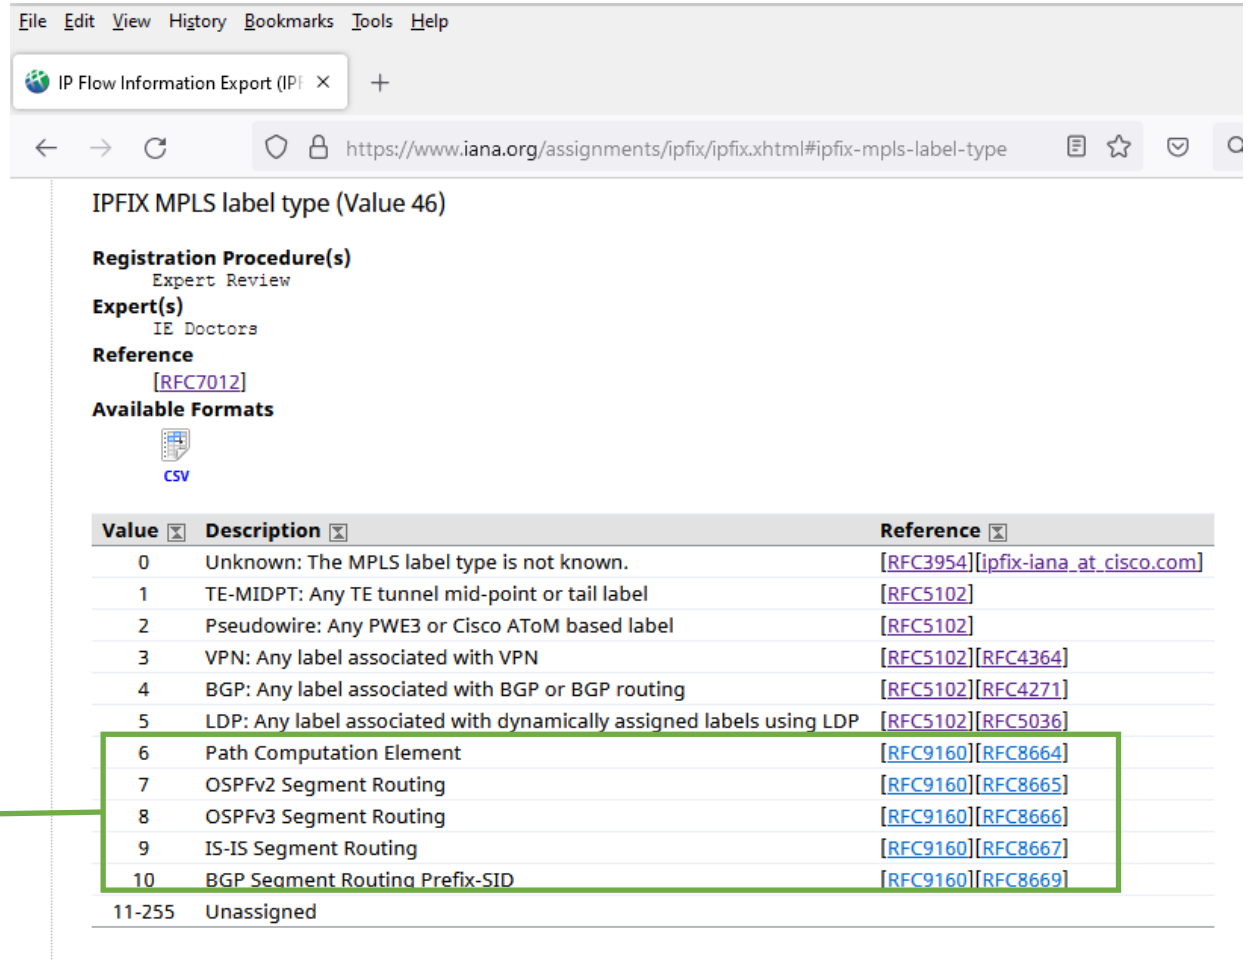

Value	Description	Reference
0	Unknown: The MPLS label type is not known.	[RFC3954][ipfix-iana_at_cisco.com]
1	TE-MIDPT: Any TE tunnel mid-point or tail label	[RFC5102]
2	Pseudowire: Any PWE3 or Cisco AToM based label	[RFC5102]
3	VPN: Any label associated with VPN	[RFC5102][RFC4364]
4	BGP: Any label associated with BGP or BGP routing	[RFC5102][RFC4271]
5	LDP: Any label associated with dynamically assigned labels using LDP	[RFC5102][RFC5036]
6	Path Computation Element	[RFC9160][RFC8664]
7	OSPFv2 Segment Routing	[RFC9160][RFC8665]
8	OSPFv3 Segment Routing	[RFC9160][RFC8666]
9	IS-IS Segment Routing	[RFC9160][RFC8667]
10	BGP Segment Routing Prefix-SID	[RFC9160][RFC8669]
11-255	Unassigned	

SRv6 @ IPFIX

Data-Plane visibility is missing in SRv6

- SRv6 is already deployed at network operators (draft-matsushima-spring-srv6-deployment-status). If you know any other network operator which migrated from MPLS to SRv6 yet.

  -> Feedback welcome

- Data-Plane visibility is missing in SRv6. Unable to see how much traffic is being forwarded or dropped with which SID. **Network operators flying blind.**
- Segment Routing Header is defined in Section 2 of RFC 8754.

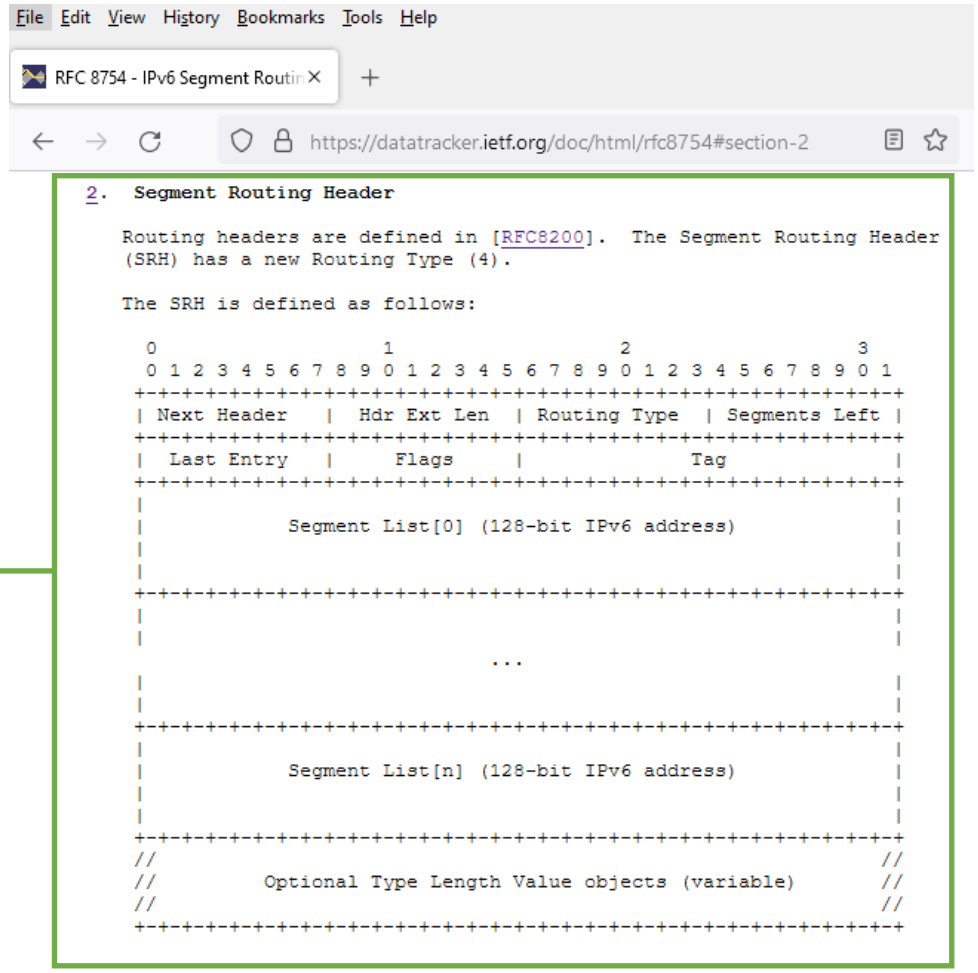

SRv6 @ IPFIX

IPFIX entities in context of the SRH (1)

- **ipv6SRHSegmentsLeft**

  8-bit unsigned integer defining the number of route segments remaining to reach the end of the segment list.

- **ipv6SRHTag**

  16-bit tag field defined in the SRH that marks a packet as part of a class or group of packets sharing the same set of properties.

- **ipv6SRHFlags**

  8-bit flags defined in the SRH.

- **ipv6SRHSegmentType**

  Name of the routing protocol or PCEP extension from where the active SRv6 segment has been learned from.

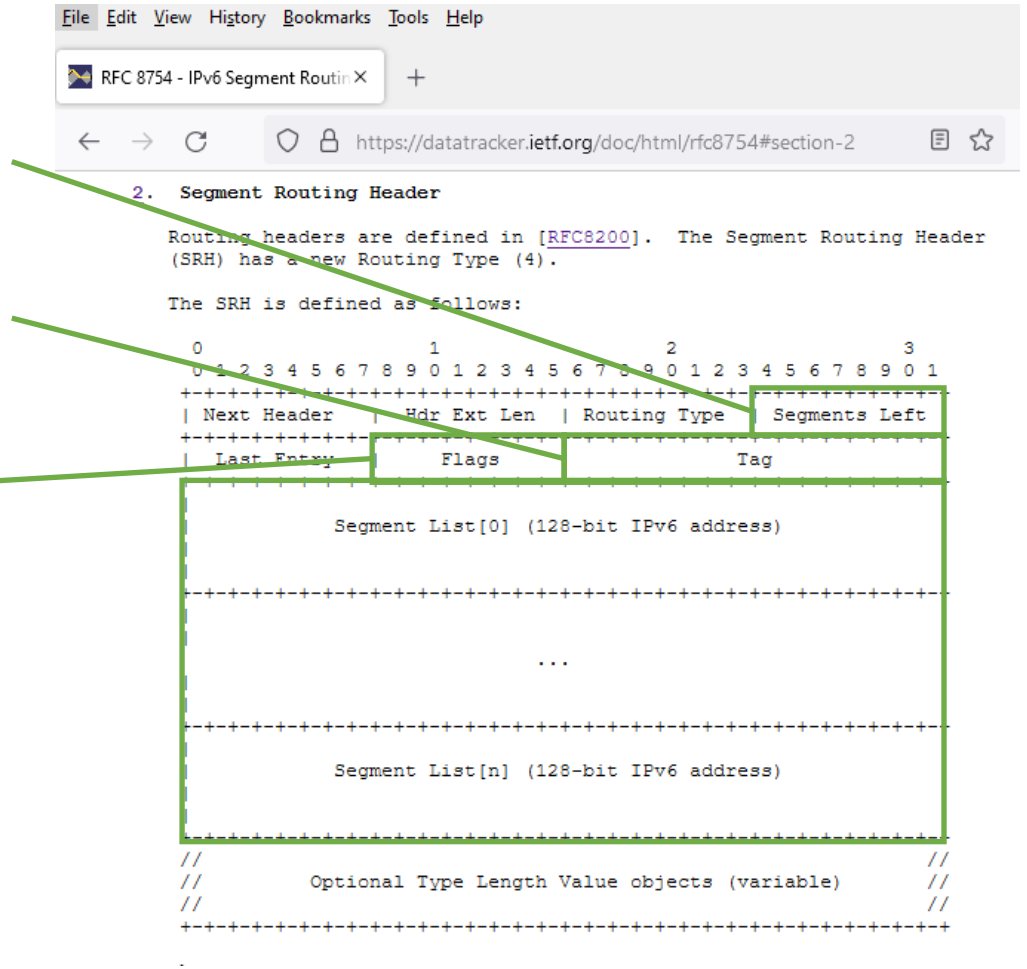

SRv6 @ IPFIX

IPFIX entities in context of the SRH (2)

- **ipv6SRHSection**

  Exposes the SRH and its TLV's as defined in section 2 of [RFC8754] as series of n octets.

- **ipv6SRHSegmentListSection**

  Exposes the SRH Segment List as defined in section 2 of [RFC8754] as series of n octets.

- **ipv6SRHSegment**

  128-bit IPv6 address that represents an SRv6 segment.

- **ipv6SRHSegmentBasicList**

  Ordered basicList [RFC6313] of zero or more 128-bit IPv6 addresses in the SRH that represents the SRv6 segment list. The Segment List is encoded starting from the active segment of the SR Policy.

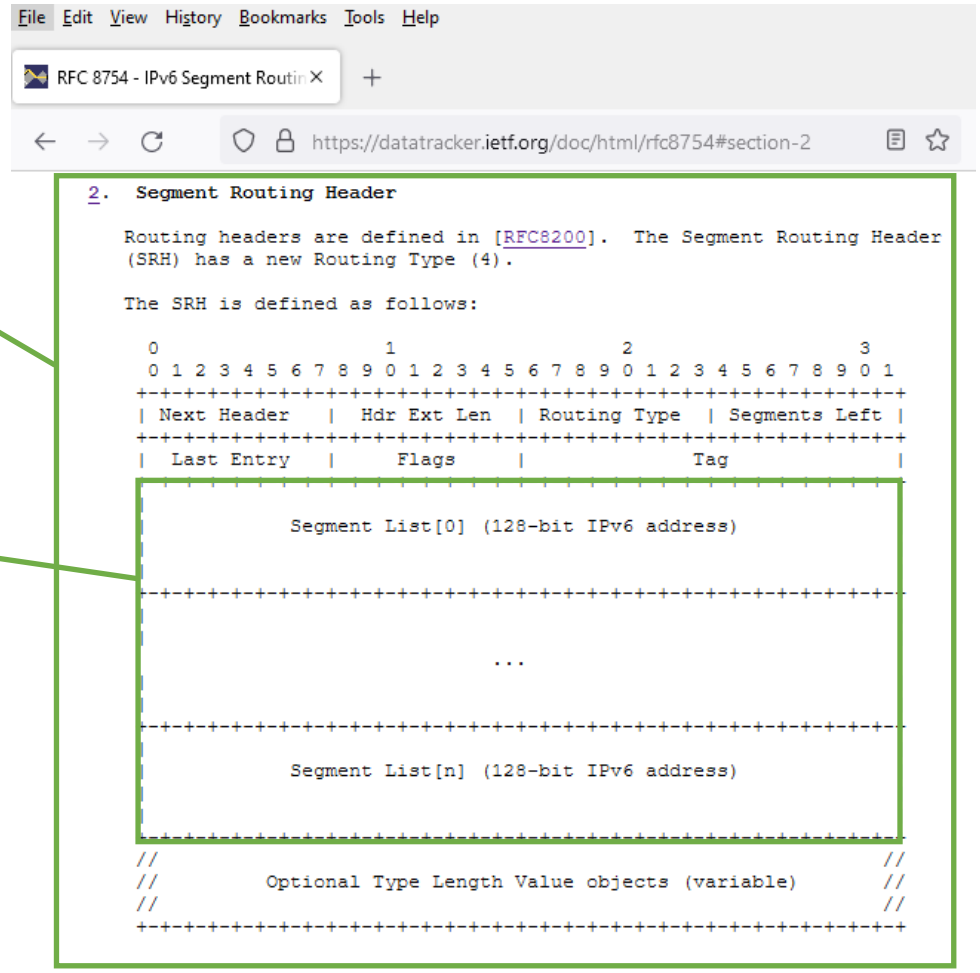

SRv6 @ IPFIX

Operational Considerations

- **ipv6SRHSegmentBasicList**
  Encodes the SID list of IPv6 addresses with a basicList, specified in the IPFIX Structured Data [RFC6313]. This encoding offers the advantage to the data collection that the different IPv6 addresses are already structured as a list, without the need of post processing. However, **this method requires some extra processing on the exporter,** to realize the BasicList data mapping.
- **ipv6SRHSegmentListSection**
  On the other hand, encodes the list of IPv6 addresses as an octetArray. This doesn't impose any data flow manipulation on the exporter, facilitating the immediate export. However, **the data collection must be able to decode** the IPv6 addresses according the SR specifications.

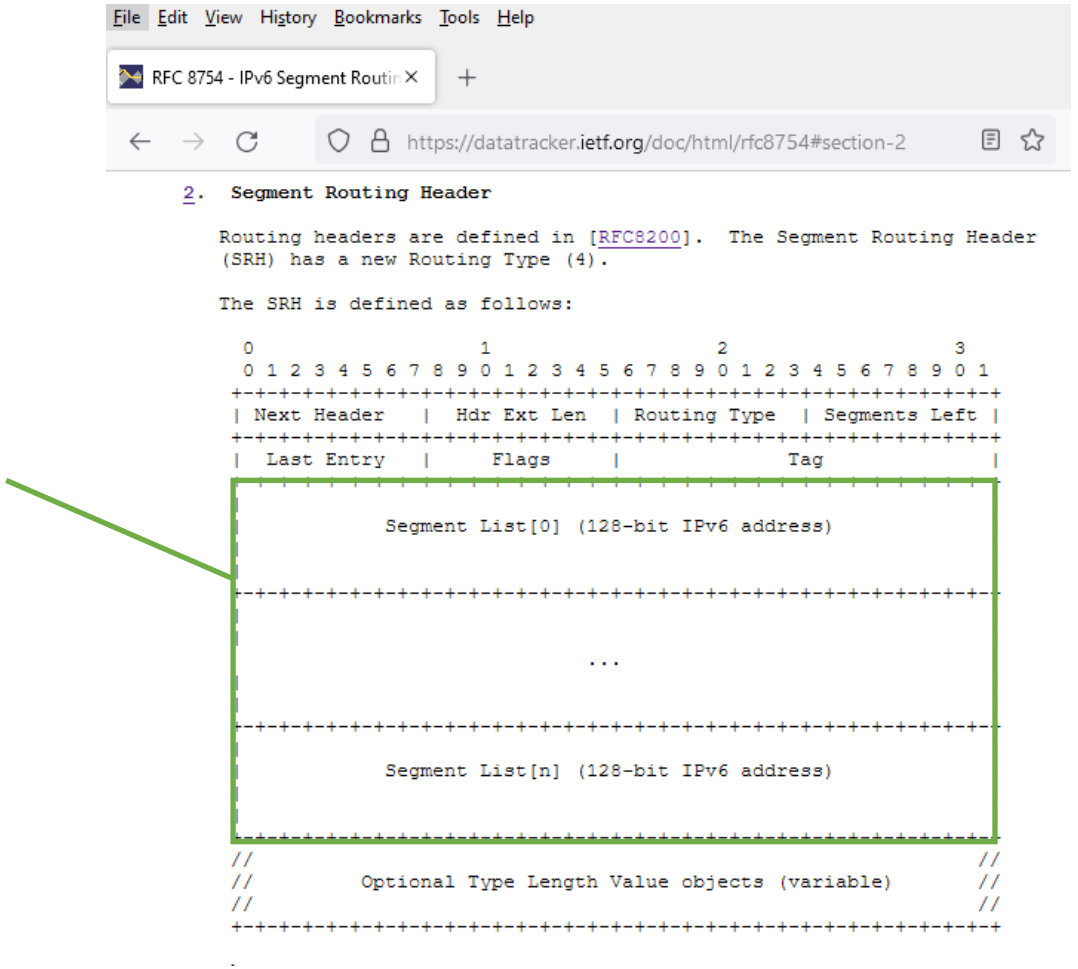

SRv6 @ IPFIX

Compressed-SID container

- As described in draft-ietf-spring-srv6-srh-compression-00.html#section-2, the compressed-SID container (C-SID container) is 128-bit long and contains a sequence of C-SIDs. Therefore, the ipv6SRHSegmentList **contains either a list of IPv6 SID's, a list of C-SID containers or both.** They are not mutually exclusive.
- It probably makes sense to add an operational consideration section to how an IPFIX data-collection could distinct between a list of IPv6 SID's and a list of C-SID containers.
- From what we understood it does not bring much added value to decompose the C-SID container into Compressed-SID (C-SID) in IPFIX.

  -> Feedback welcome

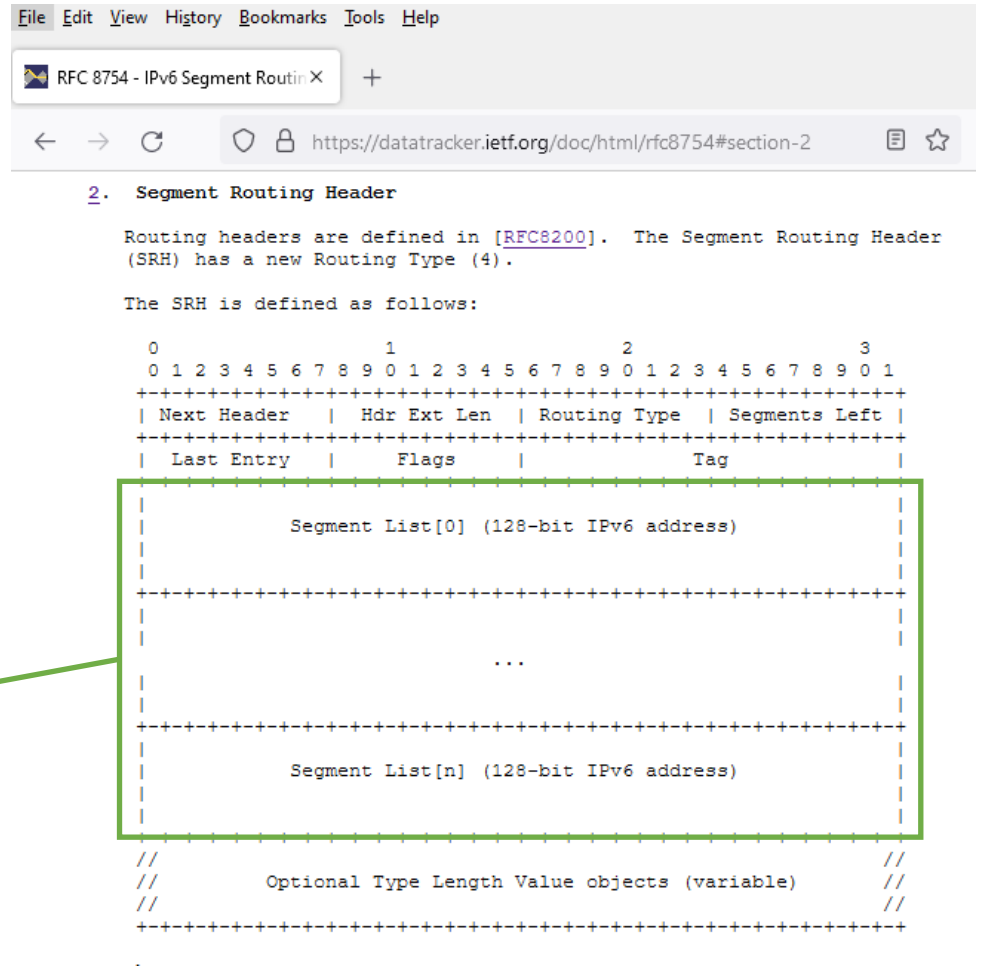

SRv6 @ IPFIX

Draft Status

- Feedback collected from SPRING, OPSAWG and IPFIX doctor.
- ipv6SRHSection and ipv6SRHSegmentListSection added to allow export of entire SRH and Segment List in one IPFIX entity.
- ipv6SRHSegmentsLeft added to express at which position of the Segment List the forwarding happens. Useful for detecting forwarding loops.
- Added operational considerations section to describe when ipv6SRHSection and ipv6SRHSegmentListSection makes sense.
- Updated IANA considerations to be in line with RFC 8126.
- The document doesn't introduce any new protocols. It is for documentation purposes. However, because new IPFIX registries are introduced, we specified the document as Internet standard, to be aligned with previous similar IPFIX RFCs (7133, 8158)

SRv6 @ IPFIX

Next Steps

- Data-Plane visibility is missing in SRv6. **Do you recognize the problem statement?**
- Authors believe that document should progress quickly through IETF to avoid private enterprise code points being used in SRv6 deployments.

-> Call for adoption at OPSAWG at IETF 113

thomas.graf@swisscom.com
benoit.claise@huawei.com
16. March 2022