

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 21 September 2022

T. Graf
Swisscom
B. Claise
Huawei
20 March 2022

Export of Segment Routing IPv6 Information in
IP Flow Information Export (IPFIX)
draft-tgraf-opsawg-ipfix-srv6-srh-03

Abstract

This document introduces new IP Flow Information Export (IPFIX) information elements to identify the SRv6 Segment Routing Header dimensions and SRv6 Control Plane Protocol that traffic is being forwarded with.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 21 September 2022.

Copyright Notice

Copyright (c) 2022 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Revised BSD License.

Table of Contents

1.	Introduction	2
2.	IPFIX Information Elements	2
3.	Use Cases	3
4.	IANA Considerations	3
5.	Operational Considerations	6
6.	Security Considerations	7
7.	Acknowledgements	7
8.	References	7
8.1.	Normative References	7
8.2.	Informative References	8
Appendix A.	IPFIX Encoding Example	8
	Authors' Addresses	10

[1.](#) Introduction

A new type of Routing Extension Header called Segment Routing Header (SRH) is defined by [\[RFC8754\]](#) which is used for applying Segment Routing (SR) on the IPv6 data plane.

Three routing protocol extensions, OSPFv3 Extensions [\[I-D.li-lsr-ospfv3-srv6-extensions\]](#), IS-IS Extensions [\[I-D.ietf-lsr-isis-srv6-extensions\]](#), BGP Prefix Segment Identifiers (Prefix-SIDs) [\[I-D.ietf-bess-srv6-services\]](#) and one Path Computation Element Communication Protocol (PCEP) Extension [\[I-D.ietf-pce-segment-routing-ipv6\]](#) have been defined to propagate Segment Identifiers (SIDs) for the IPv6 data plane.

This document defines eight new IPFIX Information Elements (IEs) and two new subregistries within the "IPFIX Information Elements" registry [\[RFC7012\]](#), respectively for the new SRH dimensions and routing protocol and PCEP extensions.

[2.](#) IPFIX Information Elements

This section defines and describes the new IPFIX IEs.

ipv6SRHFlags

8-bit flags defined in the SRH.

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.

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.

ipv6SRHSegmentListSection

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

ipv6SRHSegmentsLeft

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

ipv6SRHSection

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

ipv6SRHSegmentType

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

[3.](#) Use Cases

By using ipv6SRHSegmentBasicList(TBD4), ipv6SRHSegmentsLeft (TBD6), ipv6SRHSegmentType(TBD8) and forwardingStatus(89) it is possible to identify

- * how many packets are forwarded or dropped
- * if dropped, for which reasons,
- * identify the control plane protocol which defined the active segment,
- * the SRv6 segment list and

* how many SRv6 segments are left.

4. IANA Considerations

This document requests IANA to create new IEs (see table1) and two new subregistries called "IPFIX IPv6 SRH Flags" (table 2) and "IPFIX IPv6 SRH Segment type" (table 3) under the "IPFIX Information Elements" registry [[RFC7012](#)] available at [[IANA-IPFIX](#)] and assign the following code initial points.

Element ID	Name
TBD1	ipv6SRHFlags
TBD2	ipv6SRHTag
TBD3	ipv6SRHSegment
TBD4	ipv6SRHSegmentBasicList
TBD5	ipv6SRHSegmentListSection
TBD6	ipv6SRHSegmentsLeft
TBD7	ipv6SRHSection
TBD8	ipv6SRHSegmentType

Table 1: Creates IEs in the "IPFIX Information Elements" registry

Note to the RFC-Editor:

- * Please replace TBD1 - TBD12 with the values allocated by IANA
- * Please replace the [RFC-to-be] with the RFC number assigned to this document

[4.1.](#) ipv6SRHFlags

Name: ipv6SRHFlags ElementID: TBD1 Description: This Information Element identifies the 8-bit flags defined in the SRH. Values for this Information Element are listed in the "IPFIX IPv6 SRH Flags" registry, see [[IANA-IPFIX](#)]. Initial values in the registry are defined by the table below. New assignments of values will be administered by IANA and are subject to Expert Review Abstract Data Type: unsigned8 Data Type Semantics: flags Reference: [RFC-to-be], [RFC8754](#)[[RFC8126](#)]. Experts need to check definitions of new values for completeness, accuracy, and redundancy.

Graf & Claise

Expires 21 September 2022

[Page 4]

Internet-Draft IPFIX Segment Routing IPv6 Information

March 2022

Value	Description	Reference
0-7	Unassigned	

Table 2: "IPFIX IPv6 SRH Flags" registry

[4.2.](#) ipv6SRHTag

Name: ipv6SRHTag ElementID: TBD2 Description: This Information Element identifies the 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. Abstract Data Type: unsigned16 Data Type Semantics: identifier Reference: [RFC-to-be], [RFC8754](#)

[4.3.](#) ipv6SRHSegment

Name: ipv6SRHSegment ElementID: TBD3 Description: This Information Element identifies the 128-bit IPv6 address that represents an SRv6 segment. Abstract Data Type: ipv6address Data Type Semantics:

default Reference: [RFC-to-be], [RFC8754](#)

[4.4.](#) ipv6SRHSegmentBasicList

Name: ipv6SRHSegmentBasicList ElementID: TBD4 Description: This Information Element identifies the 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. Abstract Data Type: basicList Data Type Semantics: list Reference: [RFC-to-be], [RFC8754](#)

[4.5.](#) ipv6SRHSegmentListSection

Name: ipv6SRHSegmentListSection ElementID: TBD5 Description: Exposes the SRH Segment List as defined in [section 2](#) of Abstract Data Type: octetArray Data Type Semantics: default Reference: [RFC-to-be], [RFC8754](#)[[RFC8754](#)] as series of n octets.

[4.6.](#) ipv6SRHSegmentsLeft

Name: ipv6SRHSegmentsLeft ElementID: TBD6 Description: This Information Element identifies the 8-bit unsigned integer defining the number of route segments remaining to reach the end of the segment list. Abstract Data Type: unsigned8 Data Type Semantics: quantity Reference: [RFC-to-be], [RFC8754](#)

[4.7.](#) ipv6SRHSection

Name: ipv6SRHSection ElementID: TBD7 Description: This Information Element exposes the SRH and its TLV's as defined in [section 2](#) of Abstract Data Type: octetArray Data Type Semantics: default Reference: [RFC-to-be], [RFC8754](#)[[RFC8754](#)] as series of n octets.

[4.8.](#) ipv6SRHSegmentType

Name: ipv6SRHSegmentType ElementID: TBD8 Description: This Information Element identifies the name of the routing protocol or PCEP extension from where the active SRv6 segment has been learned from. Values for this Information Element are listed in the "IPFIX IPv6 SRH Segment type" registry, see [[IANA-IPFIX](#)]. Initial values in the registry are defined by the table below. New assignments of

values will be administered by IANA and are subject to Expert Review
 Abstract Data Type: unsigned8 Data Type Semantics: identifier
 Reference: [RFC-to-be][RFC8126]. Experts need to check definitions of new values for completeness, accuracy, and redundancy.

Value	Description	Reference
TBD9	Unknown	[RFC-to-be]
TBD10	Path Computation Element	[RFC-to-be], draft-ietf-pce-segment-routing-ipv6
TB11	OSPFv3 Segment Routing	[RFC-to-be], draft-li-ospf-ospfv3-srv6-extensions
TBD12	IS-IS Segment Routing	[RFC-to-be] draft-ietf-lsr-isis-srv6-extensions
TBD13	BGP Segment Routing Prefix-SID	[RFC-to-be], draft-ietf-bess-srv6-services

Table 3: "IPFIX IPv6 SRH Segment type" subregistry

5. Operational Considerations

The zero or more 128-bit IPv6 addresses in the SRH [RFC8754] can be exported in two different ways, with two different IPFIX IEs:

- * ipv6SRHSegmentBasicList

- * ipv6SRHSegmentListSection

The 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.

The `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.

It is not expected that an exporter would support both `ipv6SRHSegmentBasicList` and `ipv6SRHSegmentListSection` at the same time.

[6.](#) Security Considerations

There exists no significant extra security considerations regarding the allocation of these new IPFIX IEs compared to [\[RFC7012\]](#).

[7.](#) Acknowledgements

I would like to thank Pierre Francois, Yao Liu and Paolo Lucente for their review and valuable comments.

[8.](#) References

[8.1.](#) Normative References

- [RFC6313] Claise, B., Dhandapani, G., Aitken, P., and S. Yates, "Export of Structured Data in IP Flow Information Export (IPFIX)", [RFC 6313](#), DOI 10.17487/RFC6313, July 2011, <<https://www.rfc-editor.org/info/rfc6313>>.
- [RFC7012] Claise, B., Ed. and B. Trammell, Ed., "Information Model for IP Flow Information Export (IPFIX)", [RFC 7012](#), DOI 10.17487/RFC7012, September 2013, <<https://www.rfc-editor.org/info/rfc7012>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 8126](#), DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.

[8.2.](#) Informative References

[I-D.ietf-bess-srv6-services]

Dawra, G., Filsfils, C., Talaulikar, K., Raszuk, R., Decraene, B., Zhuang, S., and J. Rabadan, "SRv6 BGP based Overlay Services", Work in Progress, Internet-Draft, [draft-ietf-bess-srv6-services-13](https://www.ietf.org/archive/id/draft-ietf-bess-srv6-services-13), 19 March 2022, <<https://www.ietf.org/archive/id/draft-ietf-bess-srv6-services-13.txt>>.

[I-D.ietf-lsr-isis-srv6-extensions]

Psenak, P., Filsfils, C., Bashandy, A., Decraene, B., and Z. Hu, "IS-IS Extensions to Support Segment Routing over IPv6 Dataplane", Work in Progress, Internet-Draft, [draft-ietf-lsr-isis-srv6-extensions-18](https://www.ietf.org/archive/id/draft-ietf-lsr-isis-srv6-extensions-18), 20 October 2021, <<https://www.ietf.org/archive/id/draft-ietf-lsr-isis-srv6-extensions-18.txt>>.

[I-D.ietf-pce-segment-routing-ipv6]

Li, C., Negi, M., Sivabalan, S., Koldychev, M., Kaladharan, P., and Y. Zhu, "PCEP Extensions for Segment Routing leveraging the IPv6 data plane", Work in Progress, Internet-Draft, [draft-ietf-pce-segment-routing-ipv6-12](https://www.ietf.org/internet-drafts/draft-ietf-pce-segment-routing-ipv6-12), 6 March 2022, <<https://www.ietf.org/internet-drafts/draft-ietf-pce-segment-routing-ipv6-12.txt>>.

[I-D.li-lsr-ospfv3-srv6-extensions]

Li, Z., Hu, Z., Cheng, D., Talaulikar, K., and P. Psenak, "OSPFv3 Extensions for SRv6", Work in Progress, Internet-Draft, [draft-li-lsr-ospfv3-srv6-extensions-00](https://www.ietf.org/archive/id/draft-li-lsr-ospfv3-srv6-extensions-00), 15 January 2020, <<https://www.ietf.org/archive/id/draft-li-lsr-ospfv3-srv6-extensions-00.txt>>.

[IANA-IPFIX]

"IANA, "IP Flow Information Export (IPFIX) Entities"", <<https://www.iana.org/assignments/ipfix/ipfix.xhtml>>.

[RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", [RFC 8754](https://www.rfc-editor.org/info/rfc8754), DOI 10.17487/RFC8754, March 2020, <<https://www.rfc-editor.org/info/rfc8754>>.

[Appendix A](#). IPFIX Encoding Example

In this section an example is provided to show the encoding format for the newly introduced IEs.

SRH Nr	SRH Flags	SRH Tag	Segment Type	Segment List
1	0	123	IS-IS	2001:db8::1, 2001:db8::2, 2001:db8::3
2	0	456	IS-IS	2001:db8::4, 2001:db8::5
3	0	789	IS-IS	2001:db8::6

Table 4: 3 observed SRH headers and their routing protocol

A.1. Template Record

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
+--+																																							

Table 4: Template Record Encoding Format

In this example, the Template ID is 256, which will be used in the Data Record. The field length for ipv6SRHSegmentBasicList is 0xFFFF, which means the length of this IE is variable, and the actual length of this IE is indicated by the List Length field in the basicList format as per [RFC6313].

A.2. Data Set

The data set is represented as follows:

```

0          1          2          3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          SET ID = 256          |          Length = 136          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| ipv6SRHFlags |          ipv6SRHTag = 123          | ipv6SRHSegment |
| = 0          |          |          Type = TBD12          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          255          |          List Length = 53          | semantic=
|          |          |          ordered          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          ipv6SRHSegment = TBD4          |          Field Length = 16          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          Segment List[0] = 2001:db8::1 (16 bytes)          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          Segment List[1] = 2001:db8::2 (16 bytes)          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          Segment List[2] = 2001:db8::3 (16 bytes)          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| ipv6SRHFlags |          ipv6SRHTag = 456          | ipv6SRHSegment |
| = 0          |          |          Type = TBD12          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          255          |          List Length = 37          | semantic=
|          |          |          ordered          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          ipv6SRHSegment = TBD4          |          Field Length = 16          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          Segment List[0] = 2001:db8::4 (16 bytes)          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          Segment List[1] = 2001:db8::5 (16 bytes)          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| ipv6SRHFlags |          ipv6SRHTag = 789          | ipv6SRHSegment |
| = 0          |          |          Type = TBD12          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          255          |          List Length = 21          | semantic=
|          |          |          ordered          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          ipv6SRHSegment = TBD4          |          Field Length = 16          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```


