Workgroup: Network Working Group
Internet-Draft:
draft-ietf-opsawg-ipfix-srv6-srh-04
Published: 11 November 2022
Intended Status: Standards Track
Expires: 15 May 2023
Authors: T. Graf B. Claise P. Francois
Swisscom Huawei INSA-Lyon

Export of Segment Routing over IPv6 Information in IP Flow Information Export (IPFIX)

Abstract

This document introduces new IP Flow Information Export (IPFIX) Information Elements to identify a set of Segment Routing over IPv6 (SRv6) related information such as data contained in a Segment Routing Header (SRH), the SRv6 control plane, and the SRv6 behavior 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 15 May 2023.

Table of Contents

1. Introduction

A dedicated Routing Extension Header, called Segment Routing Header (SRH), is defined in [RFC8754] for use of Segment Routing (SR) over IPv6 data plane.

Also, three routing protocol extensions, OSPFv3 [I-D.ietf-lsr-ospfv3-srv6-extensions], IS-IS [I-D.ietf-lsr-isis-srv6-extensions] and BGP Prefix Segment Identifiers(Prefix-SIDs) [RFC9252] and Path Computation Element Communication Protocol (PCEP) Extension [I-D.ietf-pce-segment-routing-ipv6] and Segment Routing Policy [RFC9256] are defined to propagate Segment Identifiers (SIDs).

SRv6 segment endpoint behaviors are defined in [RFC8986] and describe how packets should be processed by SRv6 Segment Endpoint Nodes.

This document specifies eleven new IPFIX Information Elements (IEs) and three new subregistries within the "IPFIX Information Elements" registry [RFC7012], for SRv6 purposes.

These IEs are used to export SRv6 active segment and its control plane protocol, the SRv6 segment list, the next SRv6 node and its type, and then numbers of SRv6 segments left.

Some examples are provided in Appendix A.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

This document makes use of the terms defined in [RFC7011], [RFC8402] and [RFC8754].

The following terms are used as defined in [RFC7011].

- *IPFIX
- *IPFIX Information Elements
- *Template
- *Template Record
- *Options Template
- *Options Template Record
- *Data Record
- *Data Set

The following terms are used as defined in [RFC8402].

*Segment Routing (SR)

*Segment

*Segment List

*Active Segment

*Segment Identifier (SID)

*SRv6

*SRv6 SID

The following terms are used as defined in [RFC8754].

*Segment Routing Header (SRH)

*SR Source Node

*Transit Node

*SR Segment Endpoint Node

*Reduced SRH

*Segments Left

*Last Entry

3. IPFIX Information Elements

This section specifies the new SRv6 IPFIX IEs.

srhFlagsIPv6
8-bit flags defined in the SRH (Section 2 of [RFC8754]).

srhTagIPv6
16-bit tag field defined in the SRH (Section 2 of [RFC8754]). A tag is used to mark a packet as part of a class or group of packets sharing the same set of properties.

srhSegmentIPv6
128-bit IPv6 address that represents an SRv6 segment.

srhActiveSegmentIPv6
128-bit IPv6 address that represents the active SRv6 segment.

srhSegmentIPv6BasicList
Ordered basicList [RFC6313] of zero or more 128-bit IPv6 addresses in the SRH that represents the SRv6 segment list. As specified in Section 2 of [RFC8754], the Segment List is encoded starting from the last segment of the SR Policy. That is, the first element of the Segment List (Segment List[0]) contains the last segment of the SR Policy, the second element contains the penultimate segment of the SR Policy, and so on.

srhSegmentIPv6ListSection
Exposes the SRH Segment List as defined in Section 2 of [RFC8754] as series of n octets in IPFIX.

srhSegmentIPv6sLeft
8-bit unsigned integer defining the number of segments remaining to reach the end of the segment list.

srhSectionIPv6
Exposes the SRH and its TLVs as specified in Section 2 of [RFC8754] as series of n octets in IPFIX.

srhActiveSegmentIPv6Type
Name of the routing protocol or PCEP extension from where the active SRv6 segment has been learned from.

srhSegmentLocatorLength
The number of significant bits. Together with srhSegmentIPv6 it enables the calculation of the SRv6 Locator.

srhSegmentEndpointBehavior
16-bit unsigned integer that represents a SRv6 Endpoint behavior as per Section 4 of [RFC8986].

Note that the srhSegmentIPv6, srhSegmentLocatorLength, and srhSegmentEndpointBehavior IPFIX IEs are generic fields, to be used

in the context of IPFIX Options Templates or IPFIX Structured Data [RFC6313].

4. Sample Use Cases

The IPFIX IEs srhSegmentIPv6BasicList(TBD5) or srhSegmentIPv6ListSection (TBD6), srhActiveSegmentIPv6 (TBD4), srhSegmentIPv6sLeft (TBD7), srhActiveSegmentIPv6Type(TBD9), forwardingStatus(89) [RFC7270] [IANA-IPFIX], and some existing counters information [IANA-IPFIX] allow to provide answers to the following questions (amongst others):

* How many packets steered with a SR policy are forwarded or dropped using SRv6 in a network?

* If dropped, for which reasons?

* What is the current active segment and its associated control plane protocol?

* What is the SRv6 segment list?

* What is the next SRv6 node and its type?

* How many SRv6 segments are left?

5. IANA Considerations

This document requests IANA to create new IEs (see table 1) and three new subregistries called "IPFIX IPv6 SRH Flags" (table 2), "IPFIX IPv6 SRH Segment Type" (table 3), and "IPFIX SRv6 Endpoint Behavior" (table 4) under the "IPFIX Information Elements" registry [RFC7012] available at [IANA-IPFIX].

Element ID	Name
TBD1	srhFlagsIPv6
TBD2	srhTagIPv6
TBD3	srhSegmentIPv6
TBD4	srhActiveSegmentIPv6
TBD5	srhSegmentIPv6BasicList
TBD6	srhSegmentIPv6ListSection
TBD7	srhSegmentIPv6sLeft
TBD8	srhSectionIPv6
TBD9	srhActiveSegmentIPv6Type
TBD10	srhSegmentLocatorLength
TBD11	srhSegmentEndpointBehavior

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

Note to the RFC-Editor:

*Please replace TBD1 - TBD17 with the values allocated by IANA

*Please replace the [RFC-to-be] with the RFC number assigned to this document

5.1. srhFlagsIPv6

Name:
srhFlagsIPv6

ElementID: TBD1

Description: The 8-bit flags defined in the SRH. Assigned flags and their meanings are provided in the "Segment Routing Header Flags" IANA registry.

Abstract Data Type: unsigned8

Data Type Semantics: flags

Additional Information: RFC8754. See the assignments in the "Segment Routing Header Flags" IANA registry at https:// www.iana.org/assignments/ipv6- parameters/ipv6- parameters.xhtml#segment-routing-header-flags

Reference: [RFC-to-be]

5.2. srhTagIPv6

Name: srhTagIPv6

ElementID: TBD2

Description: 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

Additional Information: RFC8754

Reference: [RFC-to-be]

5.3. srhSegmentIPv6

Name: srhSegmentIPv6

ElementID: TBD3

Description: The 128-bit IPv6 address that represents an SRv6 segment.

Abstract Data Type: ipv6address

Data Type Semantics: default

Additional Information:
RFC8754

Reference: [RFC-to-be]

5.4. srhActiveSegmentIPv6

Name: srhActiveSegmentIPv6

ElementID: TBD4

Description: The 128-bit IPv6 address that represents the active SRv6 segment.

Abstract Data Type: ipv6address

Data Type Semantics: default

Additional Information: RFC8754

Reference: [RFC-to-be]

5.5. srhSegmentIPv6BasicList

Name: srhSegmentIPv6BasicList

ElementID: TBD5

Description: The Ordered basicList [RFC6313] of zero or more 128-bit IPv6 addresses in the SRH that represents the SRv6 segment list. As described in section 2 of [RFC8754], the Segment List is encoded starting from the last segment of the SR Policy. That is, the first element of the Segment List (Segment List[0]) contains the last segment of the SR Policy, the second element contains the penultimate segment of the SR Policy, and so on.

Abstract Data Type: basicList

Data Type Semantics: list

Additional Information: RFC8754

Reference: [RFC-to-be]

5.6. srhSegmentIPv6ListSection

Name: srhSegmentIPv6ListSection

ElementID: TBD6

Description:

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

Abstract Data Type: octetArray

Data Type Semantics: default

Additional Information: RFC8754

Reference: [RFC-to-be]

5.7. srhSegmentIPv6sLeft

Name: srhSegmentIPv6sLeft

ElementID: TBD7

Description: 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

Additional Information: RFC8754

Reference: [RFC-to-be]

5.8. srhSectionIPv6

Name: srhSectionIPv6

Description: The SRH and its TLVs as defined in Section 2 of [RFC8754] as series of n octets.

Abstract Data Type: octetArray

Data Type Semantics: default

Additional Information: RFC8754

Reference: [RFC-to-be]

5.9. srhActiveSegmentIPv6Type

Name: srhActiveSegmentIPv6Type

ElementID: TBD9

Description: 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" subregistry, see [IANA-IPFIX].

Abstract Data Type: unsigned8

Data Type Semantics: identifier

Additional Information: [RFC-to-be]

Reference: [RFC-to-be]

5.9.1. IPFIX IPv6 SRH Segment Type Subregistry

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 [RFC8126]. Experts need to check definitions of new values for completeness, accuracy, and redundancy.

Value	Description	Additional Information
TBD12	Unknown	[RFC-to-be]
TBD13	Segment Routing Policy	[RFC-to-be], <xref target="RFC9256"/>
TBD14	Path Computation Element	[RFC-to-be],<xref target="I-D.ietf-pce-segment-routing-ipv6"/>
TBD15	OSPFv3 Segment Routing	[RFC-to-be],<xref target="I-D.ietf-lsr-ospfv3-srv6-extensions"/>
TBD16	IS-IS Segment Routing	[RFC-to-be],<xref target="I-D.ietf-lsr-isis-srv6-extensions"/>
TBD17	BGP Segment Routing Prefix-SID	[RFC-to-be], <xref target="RFC9252"/>

Table 3: "IPFIX IPv6 SRH Segment type" subregistry

5.10. srhSegmentLocatorLength

Name: srhSegmentLocatorLength

ElementID: TBD10

Description:
The number of significant bits and together with srhSegmentIPv6 enables the calculation of the SRv6 Locator.

Abstract Data Type: unsigned8

Data Type Semantics: default

Additional Information: RFC8986 Section 3.1

Reference: [RFC-to-be]

5.11. srhSegmentEndpointBehavior

Name: srhSegmentEndpointBehavior

ElementID: TBD11

Description: The 16-bit SRv6 Endpoint behavior. Assigned values and their meanings are provided in the "SRV6 Endpoint Behavior" IANA registry [IANA-IPFIX].

Abstract Data Type: unsigned16

Data Type Semantics: identifier

Additional Information: RFC8986 Section 4. See the assignments in the "SRV6 Endpoint Behavior" IANA registry at https://www.iana.org/assignments/segment-routing/segment-routing.xhtml#srv6-endpoint-behaviors

Reference: [RFC-to-be]

6. Operational Considerations

6.1. SRv6 Segment List

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

- *srhSegmentIPv6BasicList
- *srhSegmentIPv6ListSection

The srhSegmentIPv6BasicList encodes the SRv6 segment list 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 srhSegmentIPv6ListSection, 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. Compared to the srhSegmentIPv6BasicList, the srhSegmentIPv6ListSection flow records length is slightly reduced.

It is not expected that an exporter would support both srhSegmentIPv6BasicList and srhSegmentIPv6ListSection at the same time.

6.2. Compressed SRv6 Segment List Decomposition

The SRv6 segment list in the IPFIX IEs srhSegmentIPv6BasicList and srhSegmentIPv6ListSection could contain compressed-SID containers as described in [I-D.ietf-spring-srv6-srh-compression]. The SID endpoint behaviors described in section 4 of [I-D.ietf-spring-srv6-srh-compression] determine wherever the segment list is compressed or not. The SID Locator as described in section 3.1 [RFC8986], determines the common most significant bits.

6.3. Multiple Segment Routing Headers

[RFC8200] describes the support of multiple extension headers in one IPv6 packet. Allowing the use of multiple SRH per SRv6 packet. The export of the same IE multiple times in one data record and related template is supported and the order within the packet SHOULD be preserved in the IPFIX export according to Section 8 of [RFC7011]. If the network node is not capable to export IPFIX for more than one SRH, it MUST export IPFIX for the active SRH.

7. Implementation Status

Note to the RFC-Editor: Please remove this section before publishing.

7.1. FD.io VPP

INSA Lyon implemented the following IEs as part of a prototype in the FD.io VPP (Vector Packet Processing) platform:

* srhActiveSegmentIPv6

* srhSegmentEndpointBehavior

* srhSegmentIPv6sLeft

*srhFlagsIPv6

*srhTagIPv6

*srhSegmentIPv6ListSection

*srhSegmentIPv6BasicList

The open source code can be obtained here: [INSA-Lyon-VPP] and was validated at the IETF 115 hackathon.

7.2. Huawei VRP

Huawei implemented the following IEs as part of a a production implementation in the VRP platform:

*srhActiveSegmentIPv6

*srhSegmentIPv6sLeft

*srhFlagsIPv6

*srhTagIPv6

*srhSegmentIPv6ListSection

The implementation was validated at the IETF 115 hackathon.

7.3. Pmacct Data Collection

Paolo Lucente implemented the IE srhSegmentIPv6ListSection decomposition as part of a a production implementation in the open source Network Telemetry data collection project pmacct.

The source code can be obtained here: [Paolo-Lucente-Pmacct] and was validated at the IETF 115 hackathon.

8. Acknowledgements

The authors would like to thank Yao Liu, Eduard Vasilenko, Bruno Decraene and Mohamed Boucadair for their review and valuable comments, and Paolo Lucente and Alex Huang Feng for the implementation and validation.

9. Security Considerations

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

10. References

10.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[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>.

[RFC7011] Claise, B., Ed., Trammell, B., Ed., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, RFC 7011, DOI 10.17487/RFC7011, September 2013, <https://www.rfc-editor.org/info/rfc7011>.

[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>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

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

10.2. Informative References

[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, 20 October 2021,

<https://www.ietf.org/archive/id/draft-ietf-lsr-isis-srv6-extensions-18.txt>.

[I-D.ietf-lsr-ospfv3-srv6-extensions] Li, Z., Hu, Z., Talaulikar, K., and P. Psenak, "OSPFv3 Extensions for SRv6", Work in Progress, Internet-Draft, draft-ietf-lsr-ospfv3-srv6-extensions-08, 14 September 2022, <https://www.ietf.org/archive/id/draft-ietf-lsr-ospfv3-srv6-extensions-08.txt>.

[I-D.ietf-pce-segment-routing-ipv6] Li, C., Negi, M. S., Sivabalan, S., Koldychev, M., Kaladharan, P., and Y. Zhu, "Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing leveraging the IPv6 dataplane", Work in Progress, Internet-Draft, draft-ietf-pce-segment-routing-ipv6-15, 23 October 2022, <https://www.ietf.org/archive/id/draft-ietf-pce-segment-routing-ipv6-15.txt>.

[I-D.ietf-spring-srv6-srh-compression] Cheng, W., Filsfils, C., Li, Z., Decraene, B., Cai, D., Voyer, D., Clad, F., Zadok, S., Guichard, N., Aihua, L., Raszuk, R., and C. Li, "Compressed SRv6 Segment List Encoding in SRH", Work in Progress, Internet-Draft, draft-ietf-spring-srv6-srh-compression-02, 11 July 2022, <https://www.ietf.org/archive/id/draft-ietf-spring-srv6-srh-compression-02.txt>.

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

[INSA-Lyon-VPP] "INSA Lyon, FD.io VPP implementation", <https://github.com/insa-unyte/vpp>.

[Paolo-Lucente-Pmacct] "Paolo Lucente, Pmacct open source Network Telemetry Data Collection", <https://github.com/pmacct/pmacct>.

[RFC7270] Yourtchenko, A., Aitken, P., and B. Claise, "Cisco-Specific Information Elements Reused in IP Flow Information Export (IPFIX)", RFC 7270, DOI 10.17487/RFC7270, June 2014, <https://www.rfc-editor.org/info/rfc7270>.

[RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <https://www.rfc-editor.org/info/rfc8200>.

[RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment

Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <https://www.rfc-editor.org/info/rfc8402>.

[RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <https://www.rfc-editor.org/info/rfc8986>.

[RFC9252] Dawra, G., Ed., Talaulikar, K., Ed., Raszuk, R., Decraene, B., Zhuang, S., and J. Rabadan, "BGP Overlay Services Based on Segment Routing over IPv6 (SRv6)", RFC 9252, DOI 10.17487/RFC9252, July 2022, <https://www.rfc-editor.org/info/rfc9252>.

[RFC9256] Filsfils, C., Talaulikar, K., Ed., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", RFC 9256, DOI 10.17487/RFC9256, July 2022, <https://www.rfc-editor.org/info/rfc9256>.

Appendix A. IPFIX Encoding Examples

This appendix represents three different encodings for the newly introduced IEs, for the example values in Table 5. The three different encodings use the following IEs, respectively: srhSegmentIPv6BasicList, srhSegmentIPv6ListSection, and srhSectionIPv6.

SRH Nr	SRH Flags	SRH Tag	Active 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 5: three observed SRH headers and their routing protocol

A.1. Three Observed SRH Headers and their Routing Protocol

A.1.1. Template Record and Data Set with Segment Basic List

With encoding in Figure 1, the examples in Table 5 are represented with the following IEs: "=>" is used to indicate which IE is mapped to a given information.

*SR Flags => srhFlagsIPv6 (TBD1)

*SR Tag => srhTagIPv6 (TBD2)

*Active Segment Type => srhActiveSegmentIPv6Type (TBD9)

*Segment List => srhSegmentIPv6BasicList (TBD5)

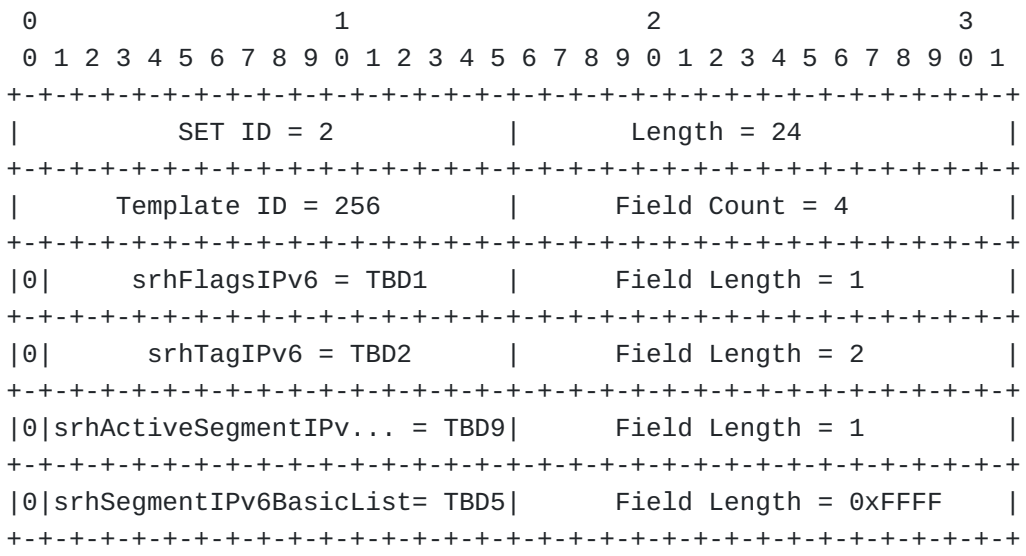

Figure 1: Template Record with Basic List Encoding Format

In this example, the Template ID is 256, which will be used in the Data Record. The field length for srhSegmentIPv6BasicList 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].

The data set is represented as follows:

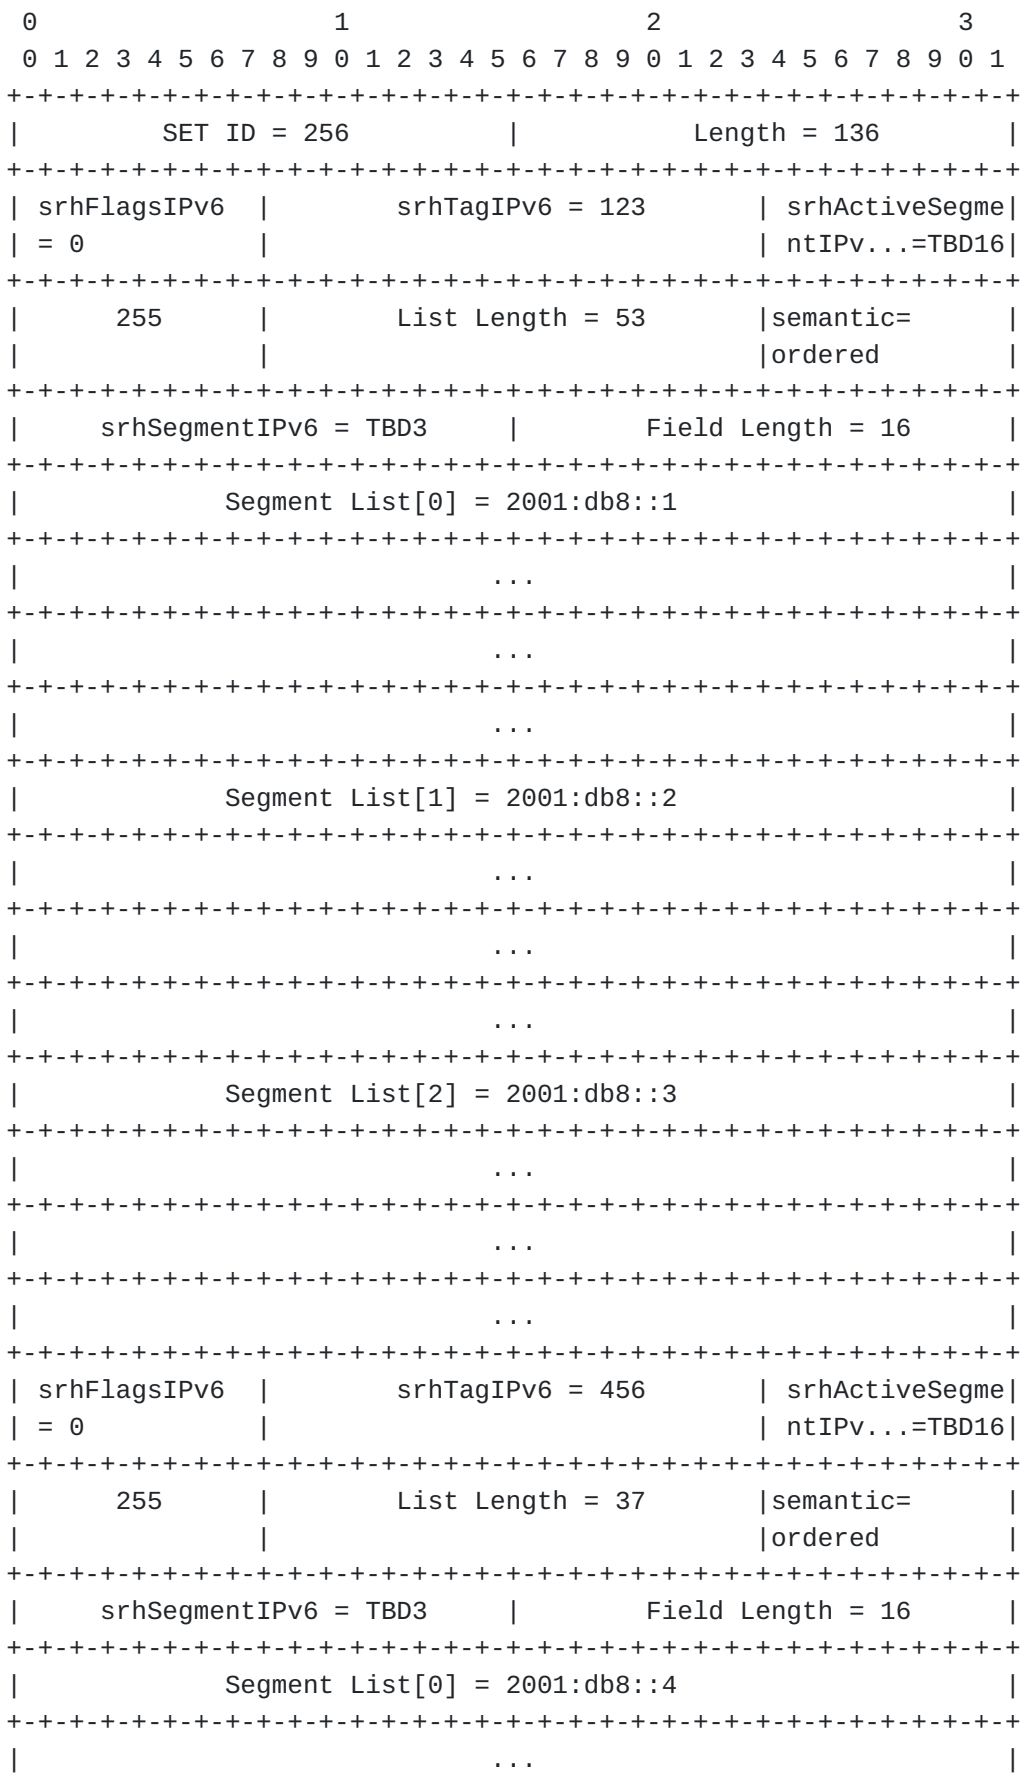
0 1 2 3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
SET ID = 256
Length = 136
srhFlagsIPv6 = 0
srhTagIPv6 = 123
srhActiveSegmentIPv6...=TBD16
255
List Length = 53
semantic=ordered
srhSegmentIPv6 = TBD3
Field Length = 16
Segment List[0] = 2001:db8::1
...
...
...
Segment List[1] = 2001:db8::2
...
...
...
Segment List[2] = 2001:db8::3
...
...
...
srhFlagsIPv6 = 0
srhTagIPv6 = 456
srhActiveSegmentIPv6...=TBD16
255
List Length = 37
semantic=ordered
srhSegmentIPv6 = TBD3
Field Length = 16
Segment List[0] = 2001:db8::4
...

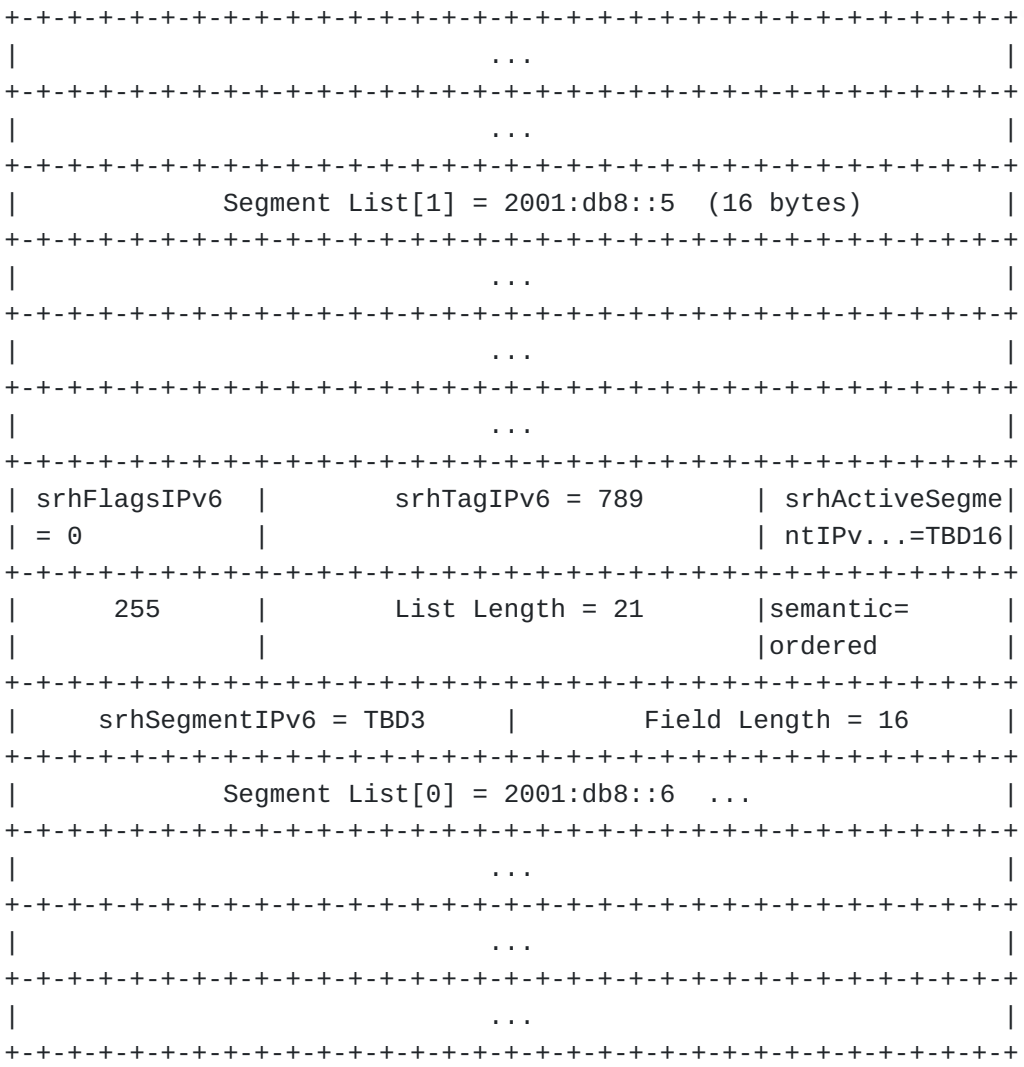

Figure 2: Data Set Encoding Format for Basic List

A.1.2. Template Record and Data Set with Segment List Section

With encoding in Figure 3, the examples in Table 5 are represented with the following IEs: "=>" is used to indicate which IE is mapped to a given information.

*SR Flags => srhFlagsIPv6 (TBD1)

*SR Tag => srhTagIPv6 (TBD2)

*Active Segment Type => srhActiveSegmentIPv6Type (TBD9)

*Segment List => srhSegmentIPv6ListSection (TBD6)

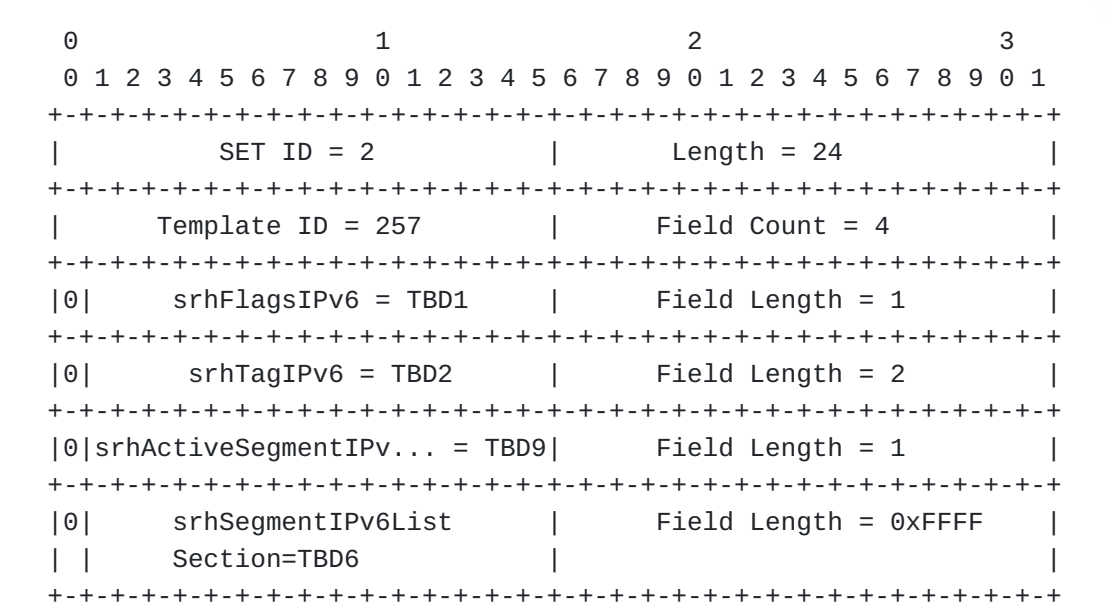

Figure 3: Template Record with Segment List Section Encoding Format

In this example, the Template ID is 257, which will be used in the Data Record. The field length for srhSegmentIPv6ListSection in the Template Record is 0xFFFF, which means that the length of this IE is variable: its actual length is encoded in the Data Set. Note that, with an actual length inferior to 255 in the Data Record example, the length field is encoded in 8 bits (see https://www.rfc-editor.org/rfc/rfc7011.html#section-7)

The data can be represented as follows:

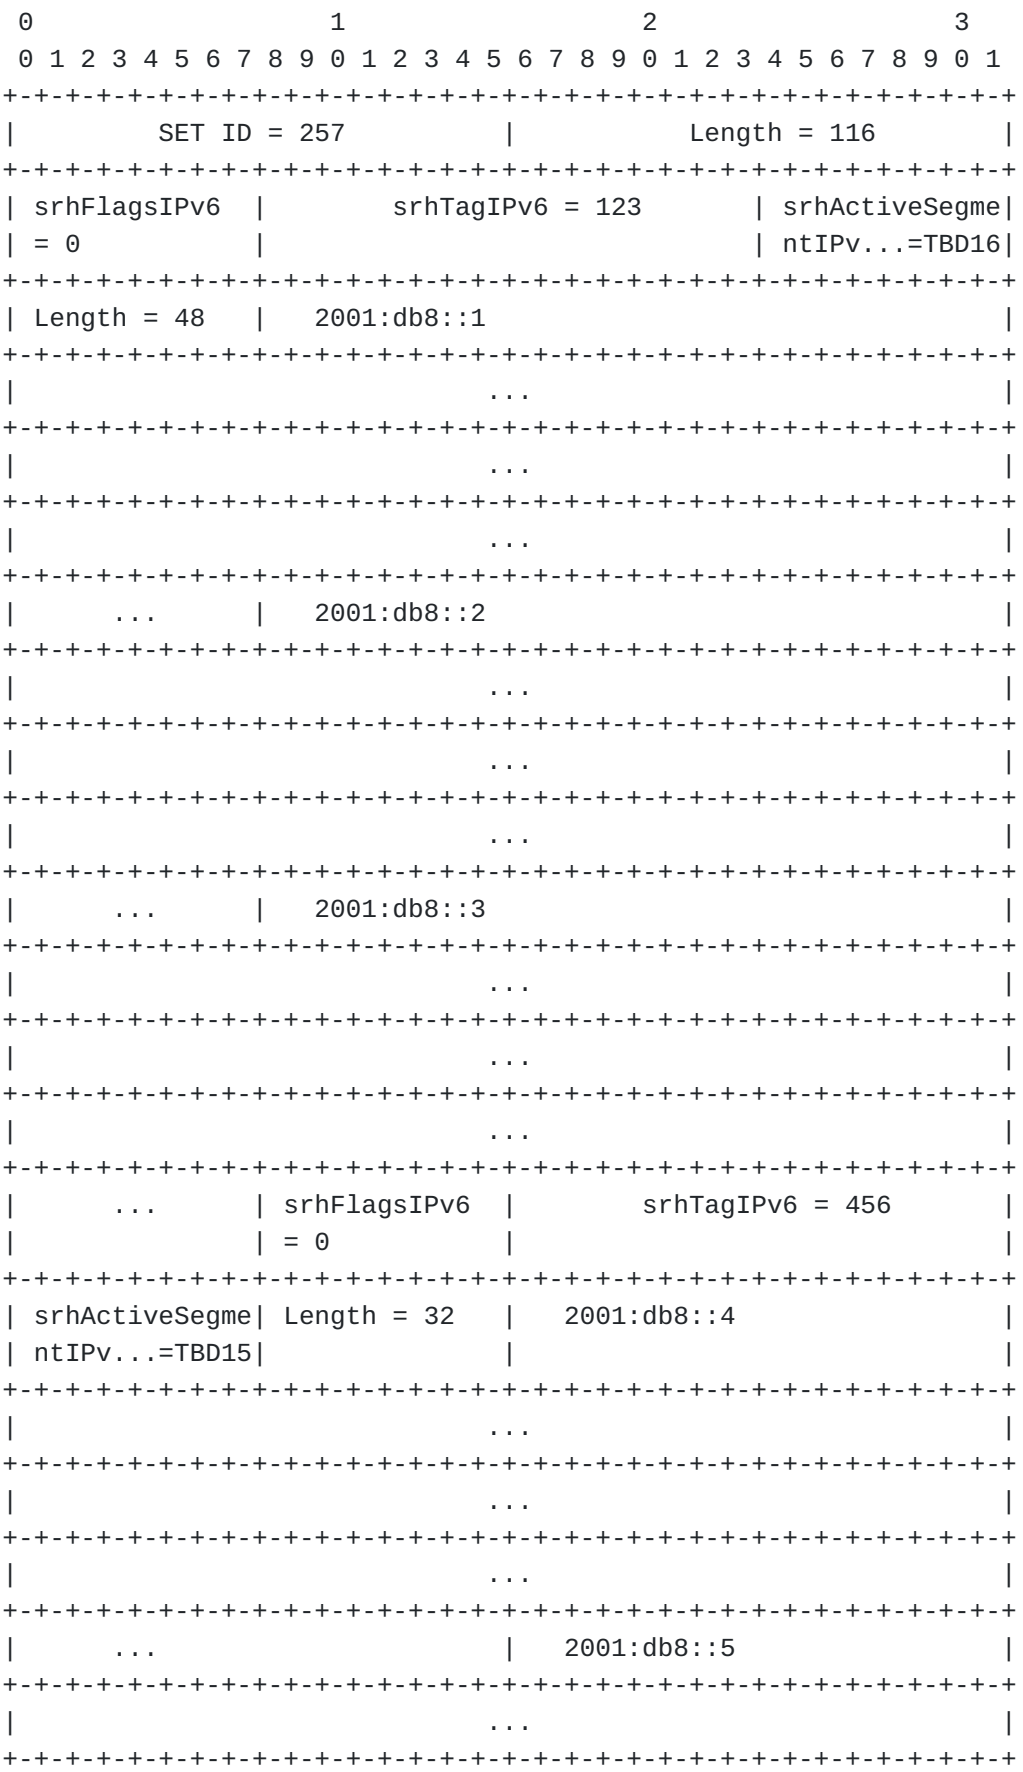
0 1 2 3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
SET ID = 257
Length = 116
srhFlagsIPv6 = 0
srhTagIPv6 = 123
srhActiveSegmentIPv...=TBD16
Length = 48
2001:db8::1
2001:db8::2
2001:db8::3
srhFlagsIPv6 = 0
srhTagIPv6 = 456
srhActiveSegmentIPv...=TBD15
Length = 32
2001:db8::4
2001:db8::5

Figure 4: Data Set Encoding Format for Segment List Section

A.1.3. Template Record and Data Set with SRH Section

With encoding in Figure 5, the examples in Table 5 are represented with the following IEs: "=>" is used to indicate which IE is mapped to a given information.

*SR Flags + SR Tag + Segment List => srhSectionIPv6 (TBD8)

*Active Segment Type => srhActiveSegmentIPv6Type (TBD9)

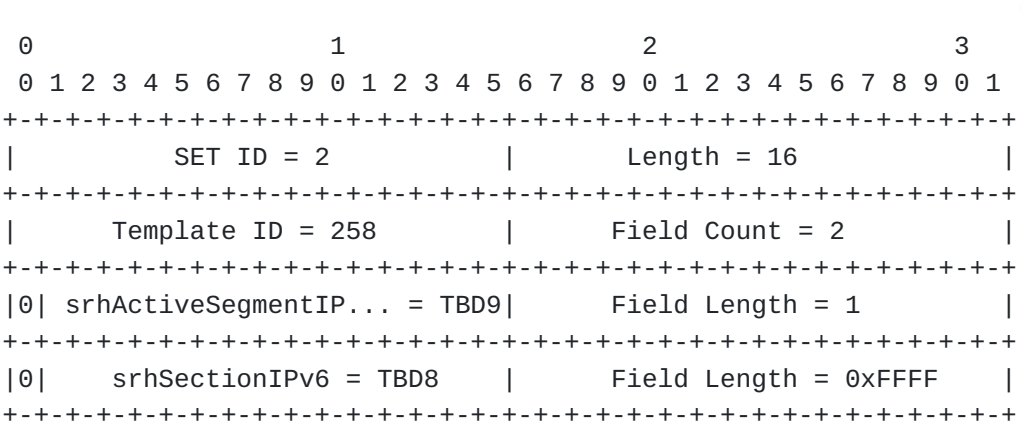

Figure 5: Template Record with SRH Section Encoding Format

In this example, the Template ID is 258, which will be used in the Data Record. The field length for srhSectionIPv6 in the Template Record is 0xFFFF, which means that the length of this IE is variable: its actual length is encoded in the Data Set. Note that, with an actual length inferior to 255 in the Data Record example, the length field is encoded in 8 bits (see https://www.rfc-editor.org/rfc/rfc7011.html#section-7)

The data can be represented as follows:

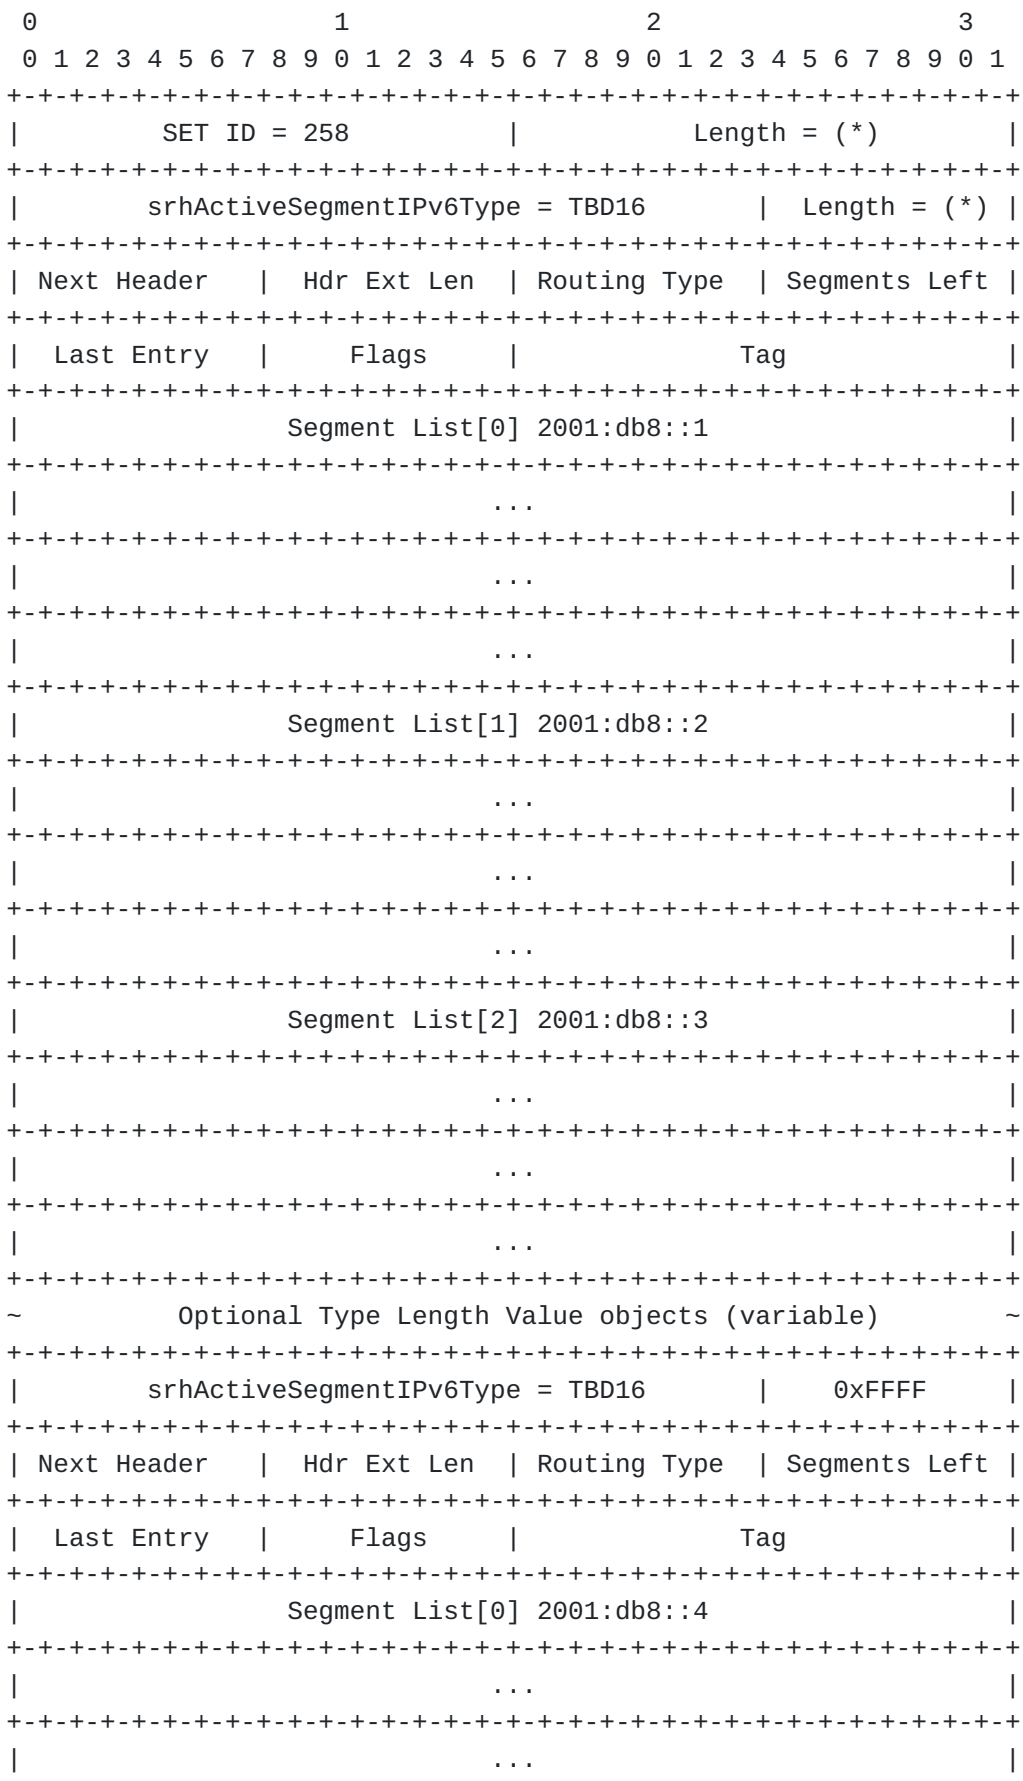
0 1 2 3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
SET ID = 258
Length = (*)
srhActiveSegmentIPv6Type = TBD16
Length = (*)
Next Header
Hdr Ext Len
Routing Type
Segments Left
Last Entry
Flags
Tag
Segment List[0] 2001:db8::1
...
...
...
Segment List[1] 2001:db8::2
...
...
...
Segment List[2] 2001:db8::3
...
...
...
Optional Type Length Value objects (variable)
srhActiveSegmentIPv6Type = TBD16
0xFFFF
Next Header
Hdr Ext Len
Routing Type
Segments Left
Last Entry
Flags
Tag
Segment List[0] 2001:db8::4
...
...

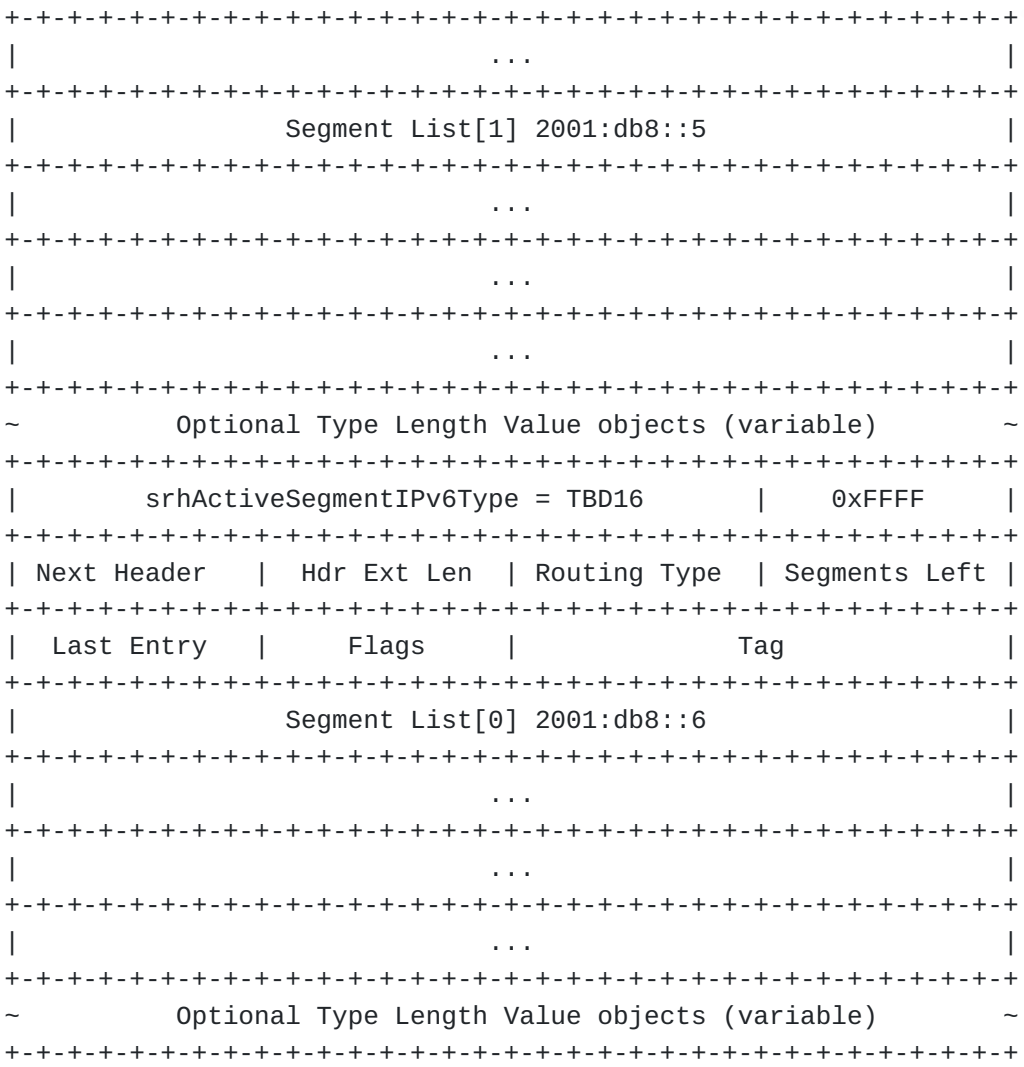

Figure 6: Data Set Encoding Format for SRH Section

(*) The Length MUST be calculated to include the optional Type Length Value objects.

A.2. Options Template Record and Data Set for SRv6 Segment End Point behavior and Locator Length

This appendix provides an SRv6 EndPoint Bevahior Options Template example, for the values presented in Table 12. In the Options Template case, the srhActiveSegmentIPv6 Information Element is a Scope field

Entry Nr	SRH End Point IPv6	SRH End Point Behavior	SRH Segment Locator Length
1	2001:db8::1	End [1]	48
2	2001:db8::4	End with NEXT-CSID [43]	48
3	2001:db8::6	End.DX6 [16]	48

Figure 7: three observed SRv6 Segment End Point Behaviors

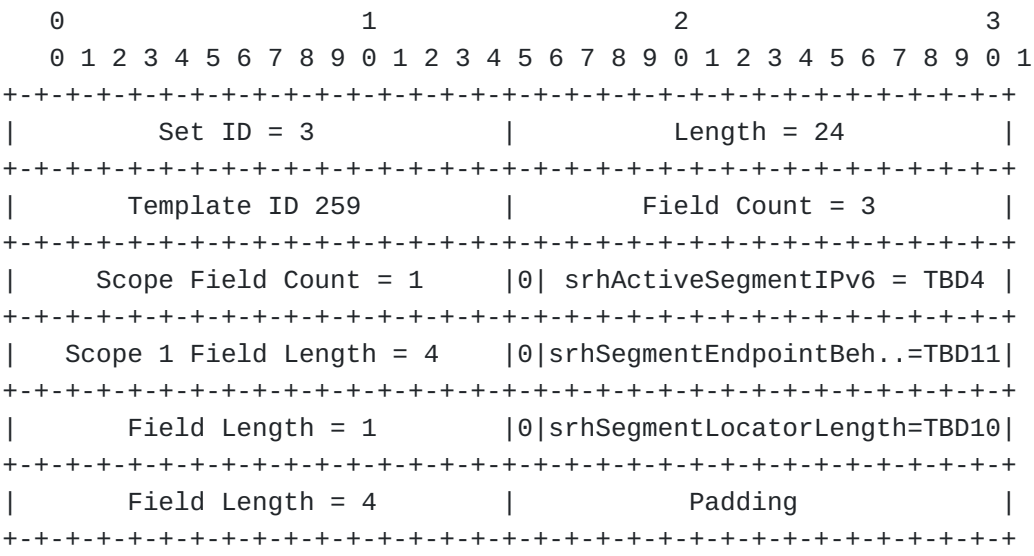

Figure 8: Template Record with SRH Section Encoding Format

In this example, the Template ID is 259, which will be used in the Data Record.

The data set is represented as follows:

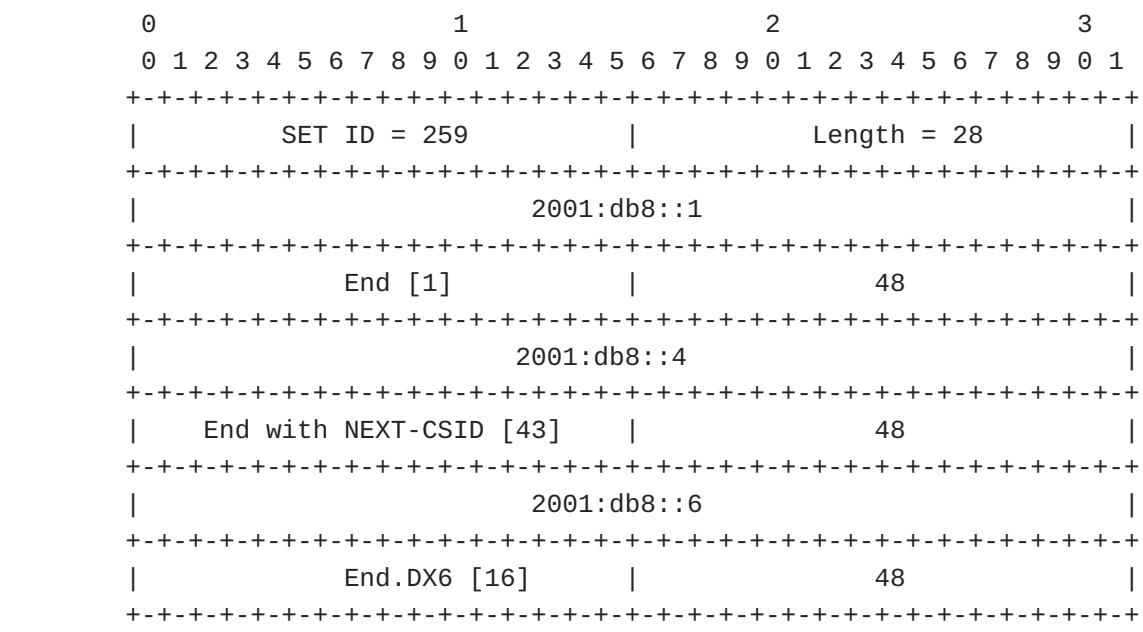

Figure 9: Data Set Encoding Format for SRH Section

(*) The Length MUST be calculated to include the optional Type Length Value objects.

Authors' Addresses

Thomas Graf
Swisscom
Binzring 17
CH-8045 Zurich
Switzerland

Email: thomas.graf@swisscom.com

Benoit Claise
Huawei

Email: benoit.claise@huawei.com

Pierre Francois
INSA-Lyon
Lyon
France

Email: pierre.francois@insa-lyon.fr