

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: July 10, 2016

C. Jacquenet
M. Boucadair
Orange
January 7, 2016

An IPv6 Extension Header for Service Function Chaining (SFC)
draft-jacquenet-sfc-ipv6-eh-01

Abstract

This document specifies an IPv6 extension header for Service Function Chaining (SFC) purposes. This extension header is used by SFC data plane elements to make forwarding decisions in an IPv6-enabled SFC domain and it conveys metadata that are processed by SFC-aware nodes.

This extension is intended to be used within a single administrative domain.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

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 <http://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 July 10, 2016.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

Internet-Draft

[draft-jacquenet-ipv6-eh-sfc-00.txt](#)

January 2016

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
2.	Overview	3
3.	SFC IPv6 Extension Header Format	3
3.1.	Source Routed SFP	6
4.	Operation	7
4.1.	Generic Considerations	7
4.2.	Generating the SFC Extension Header	8
4.3.	Processing the SFC Extension Header	8
4.3.1.	Processing Source-Routed SFP Information	8
5.	IANA Considerations	9
6.	Security Considerations	9
6.1.	Privacy	10
6.2.	Invalid Context Information	10
6.3.	Forwarding Threats	10
7.	Acknowledgements	11
8.	References	11
8.1.	Normative References	11
8.2.	Informative References	11
	Authors' Addresses	12

[1.](#) Introduction

Service Function Chaining (SFC) can be seen as a technique that facilitates the dynamic enforcement of differentiated traffic forwarding policies within an SFC-enabled domain. SFC operation assumes the manipulation of some information that is typically carried by packets within an SFC-enabled domain. In particular, this information is meant to assist Service Function Forwarders (SFFs) in making forwarding decisions within the SFC-enabled domain.

The overall SFC problem space is discussed in [[RFC7498](#)], while a data

plane architecture is documented in [[RFC7665](#)].

Several options can be used to carry SFC-specific information. Some of them can take advantage of various existing tools such as encapsulation schemes (e.g., IP-in-IP), or specific fields in an IP

header. This document specifies an IPv6 Extension Header ([RFC6564](#)) to carry SFC-related information.

The SFC extension header is intended to be used within a single administrative domain.

The reader should be familiar with the terms defined in [[RFC7498](#)] and [[RFC7665](#)].

2. Overview

Unlike some other solutions that require the use of yet another shim layer to carry SFC information, the use of an IPv6 Extension Header (EH) in IPv6-enabled SFC domains has the advantage to get rid of any specific transport encapsulation scheme when forwarding packets between nodes that are connected to the same subnet. Figure 1 shows the case of a packet that carries the SFC EH and which is forwarded to the SFC Next Hop that is connected to the same subnet.

Figure 2 shows a packet that is encapsulated in an IPv6 packet that contains the SFC EH. Such encapsulation scheme can also be used to carry IPv4 packets within an IPv6-enabled SFC domain.

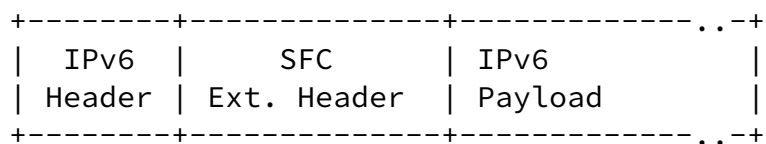
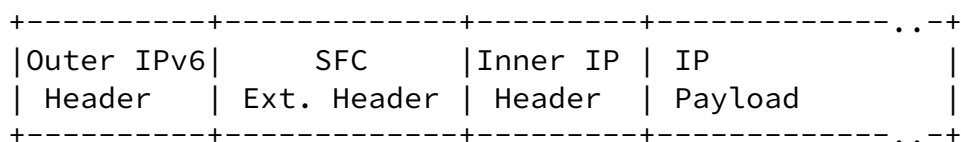


Figure 1



|--Original IP packet-----|

|-----Encapsulated IPv6 packet-----|

Figure 2

3. SFC IPv6 Extension Header Format

The IPv6 Extension Header to carry SFC metadata has format shown in Figure 3.

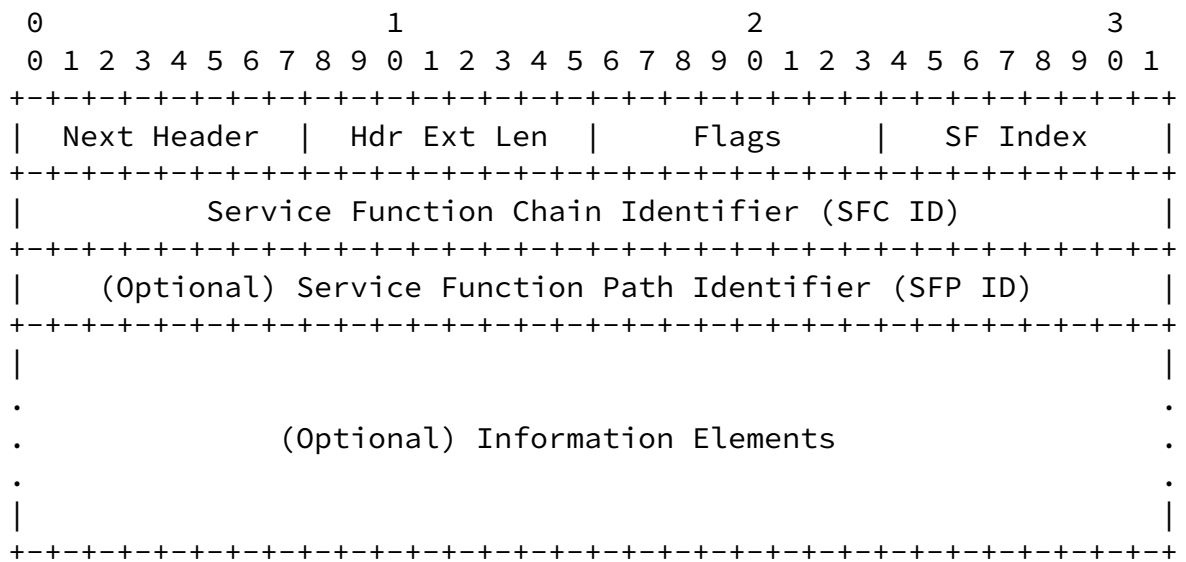


Figure 3

The description of the fields is as follows:

Next Header: 8-bit selector. Identifies the type of header immediately following the extension header.

Hdr Ext Len: 8-bit unsigned integer. Length of the extension header in 8-octet units, excluding the first 8 octets.

Flags: The Flags field comprises a set of 8 flags:

```

+---+---+---+---+
|P|E|r|r|r|r|r|M|
+---+---+---+---+

```

where "rrrrrr" are reserved bits for future assignment as additional flag bits. r bits MUST each be sent as zero and MUST be ignored on receipt.

When set, the P-flag indicates that a Service Function Path Identifier (SFP ID) field is present in the SFC EH. This flag is set to 0 by default: this means that there is no SFP ID information carried in the SFC EH.

When set, the E-flag indicates that a source routed SFP field is present in the SFC EH. This flag is set by default to 0, meaning there is no source routed SFP field present in the SFC EH.

When set, the M-flag indicates that an extended set of a 32-bit encoded Flags field is present in the SFC EH. The default value

of the M flag is 0. This feature allows to extend the SFC EH with new flags while ensuring backward compatibility. When present, the extended flag field MUST be positioned right after the SFC ID field.

SF Index: 8-bit unsigned integer. This field is decremented by 1 and used to detect SFC loops.

Service Function Chain Identifier (SFC ID): 8-bit unsigned integer. Identifies the Service Function Chain that is associated to the IPv6 packet.

(Optional) Service Function Path Identifier (SFP ID): This field MUST be supplied only if 'P' flag is set. This field is used to convey an identifier of a path that is bound to a given service function chain. A null value of this field means that no specific constraint is to be applied when forwarding this packet with this service function chain. It is RECOMMENDED to use this field only if non-null identifiers are to be carried in the SFC EH. A null value with P-flag set leads to the same behavior as with P-flag set to 0.

(Optional) Information Elements: Conveys one or multiple optional data that may be supplied within an SFC-enabled domain. The format of an optional Information Element can either be associated with the definition of a new flag or encoded according to the following TLV format Figure 4.

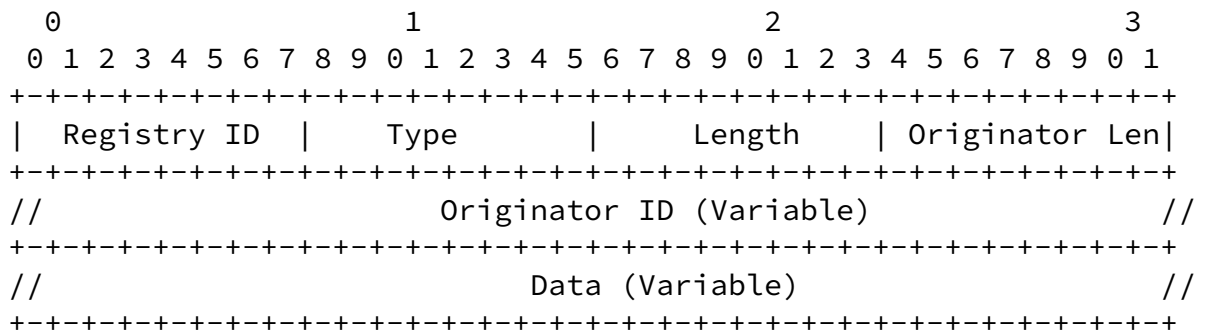


Figure 4

The description of the fields is as follows:

Registry ID: In order to foster service innovation, this field allows to inherit from existing code point registries that are likely to be useful in a SFC context. The following value is reserved by this specification:

1: IPFIX [[IPFIX](#)].

Type: Indicates the code point of the information element. If Registry ID is set, then the interpretation of this field must conform to the one defined for that specific registry.

Length: Indicates, in octets, the length of the data carried in the Information Element (including the "Originator Len" and "Originator ID" fields).

Originator Len: Indicates, in octets, the length of the "Originator ID" field.

Originator ID: Provides the identifier of the entity that injected this Information Element in the SFC Extension Header.

Originator ID: Conveys the identifier of the entity that injected

this Information Element in the SFC header (e.g., a service function identifier, a classifier, etc.). This document does not make any assumption about the structure of the information carried in this field because this is deployment-specific. This information is used by SFC-aware elements to enforce policies such as: process a context information if and only if it was supplied by a given entity. This information can be used as a safeguard against misbehaving nodes that inject illegitimate data in the SFC EH.

Data (Variable): The semantics of this field depend on the "Registry ID" and "Type" fields.

3.1. Source Routed SFP

If the E-flag is set, a "Source Routed SFP" field MUST be present in the SFC Extension Header. This field MUST be positioned right after the "Service Function Path Identifier (SFP ID)" field and "Extended Flag bits", if P-flag or M-flag are set. It MUST be positioned right after the "Service Function Chain Identifier (SFC ID)" field if P-flag and M-flag are both set to 0.

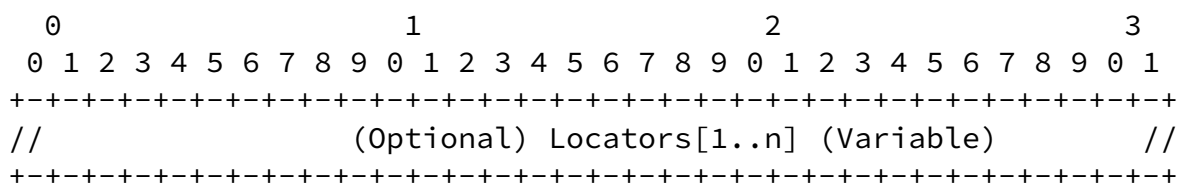
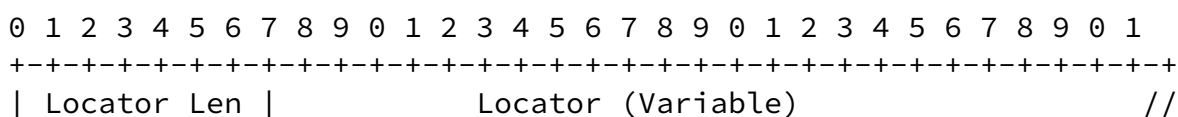


Figure 5

The optional "Source Routed SFP" field is structured as a vector of SF locators (whereas a SF locator could be an IP address, a MAC address or a FQDN, for example).

A SFP Source Route is said to be "strict" when the exact set of all the SF instances that need to be invoked along the SFP is explicitly and exhaustively mentioned in the field. Each locator in the list is encoded as follows:



A classifier typically inserts the SFC EH to every incoming packet that matches one of the entries of the SFC classification policy table maintained by the classifier. The classifier is supposed to be configured with the set of context information that must be supplied for each service function. If no such instruction is provided, the classifier inserts by default an identifier of the service chain and optionally an identifier of the service function path.

The classifier may also inject a source SFP as part of the SFC EH that will be injected in packet matching its classification policies.

The SFC EH is inserted either following the format in Figure 1 (if the next hop is within the same subnet as the classifier) or as shown in Figure 2 otherwise. When an encapsulation is required, the destination IP address is set to the IPv6 address of the first hop in the chain.

SFC-aware nodes that are configured to inject context information for a given service function chain can update the context of an SFC EH.

[4.3.](#) Processing the SFC Extension Header

Upon receipt of an IPv6 packet that carries the SFC EH, a SFF must, eventually decapsulate the packet, and process the metadata information carried in the SFC EH: typically, the SF node that embeds the SFF capability will use these metadata to (1) position itself in the forwarding path, (2) determine which SF instance(s) need to be invoked next and (3) make its forwarding decision according to the SFC instructions carried in the SFC EH and as per the matching entry of its SFC Forwarding Policy Table.

Once the packet is processed by the corresponding SF, SF Index is decremented by 1.

An SFC-aware node MUST discard packets with an "SF Index" equal to 0. This event must be logged locally.

[4.3.1.](#) Processing Source-Routed SFP Information

If the SFC EH carried in the incoming IPv6 packet contains Source-Routed SFP ([Section 3.1](#)), the SFF will forward the packet according to the instructions carried in the corresponding field: if this is a Strict Source Route, the SFF will forward the packet towards the next SF node that embeds the SF instance identified by the SF Locator carried in the Source-Routed SFP field, possibly upon completion of

some SF operation, depending on the nature of the chain and its corresponding instructions.

If the explicit route happens to be a loose source route:

1. If the next SF instance that needs to be invoked is explicitly identified by its Locator, then the SFF forwards the packet accordingly: the next SF to be invoked can be reached according to the corresponding entry of its SFC Forwarding Policy Table.
2. If the next SF instance that needs to be invoked is valued to "ANY", then the SFF forwards the packet according to the best matching entry of its SFC Forwarding Policy Table, as per the SFC ID and Hop Index carried in the SFC EH.

By default, packets destined outside the SFC-enabled domain MUST be strip any SFC EH that is carried in the packet.

When a node receives an IPv6 packet with a "ICMPv6 Destination Unreachable Code, "Error in Source Routing Header"xxx

[5.](#) IANA Considerations

This document requests IANA to assign the following values from the register in <http://www.iana.org/assignments/ipv6-parameters/ipv6-parameters.xhtml#extension-header>:

Protocol Number	Description	Reference
TBD	SFC Extension Header	[This document]

Also, this document requests IANA to assign a sub-registry for Registry ID. The following value is reserved by this specification:

1: IPFIX

[6.](#) Security Considerations

Security considerations discussed in [RFC7045] and [RFC7112] apply. Additional considerations are discussed in the following sub-sections.

[6.1.](#) Privacy

Because context headers may reveal privacy information (e.g., IMSI, user name, user profile, location, etc.). SFC Extension Headers MUST NOT be exposed outside the SFC domain. Also, means to protect context headers from eavesdroppers SHOULD be enforced.

[6.2.](#) Invalid Context Information

In order to control the information that can be supplied by a SFC-aware node, and therefore influence the behavior of an SFC-aware node within the SFC-enabled domain, the Originator ID field can be used as a first safeguard to check that the node is entitled to supply such information. If so, the Originator ID field can also be used to check whether the supplied information can be processed as part of the instructions that pertain to a given service function chain.

An SFC-aware node can be provided with the appropriate SFC instructions by the SFC control plane or by configuration.

[6.3.](#) Forwarding Threats

This specification is not subject to infinite forwarding loops because a loop can be detected by an SF Index equal to 0.

Several attacks (e.g., evade access controls based on destination addresses, amplification attacks) have been identified in [[RFC4942](#)]. Such attacks can be prevented in the SFC context by the enforcement of adequate policies at the boundaries of the SFC domain. Typically, SFC border nodes of a SFC-enabled domain can be configured to discard any SFC EH that may be present in a packet that enters the domain, and strip the SFC EH when the packet is forwarded outside of the SFC-enabled domain, so that the information carried by the SFC EH is not leaked outside the domain when the packet exits the SFC-enabled domain.

Nevertheless, a node of a SFC-enabled domain may alter the contents of the SFC EH, thereby possibly distorting the SFP. Misbehaving nodes can be detected and countermeasures applied, if adequate

monitoring is enforced. Also, means to protect traffic against illegitimate SFs/SFFs that do not belong to the SFC-enabled domain must be enabled. Such means should typically be defined in service function discovery specifications.

[7.](#) Acknowledgements

TBD

[8.](#) References

[8.1.](#) Normative References

- [IPFIX] International Organization for Standardization, "IP Flow Information Export (IPFIX) Entities", 1992, <<http://www.iana.org/assignments/ipfix/ipfix.xhtml>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC6564] Krishnan, S., Woodyatt, J., Kline, E., Hoagland, J., and M. Bhatia, "A Uniform Format for IPv6 Extension Headers", [RFC 6564](#), DOI 10.17487/RFC6564, April 2012, <<http://www.rfc-editor.org/info/rfc6564>>.
- [RFC7045] Carpenter, B. and S. Jiang, "Transmission and Processing of IPv6 Extension Headers", [RFC 7045](#), DOI 10.17487/RFC7045, December 2013, <<http://www.rfc-editor.org/info/rfc7045>>.

[8.2.](#) Informative References

- [RFC4942] Davies, E., Krishnan, S., and P. Savola, "IPv6 Transition/Co-existence Security Considerations", [RFC 4942](#), DOI 10.17487/RFC4942, September 2007,

<<http://www.rfc-editor.org/info/rfc4942>>.

- [RFC7112] Gont, F., Manral, V., and R. Bonica, "Implications of Oversized IPv6 Header Chains", [RFC 7112](#), DOI 10.17487/RFC7112, January 2014, <<http://www.rfc-editor.org/info/rfc7112>>.
- [RFC7498] Quinn, P., Ed. and T. Nadeau, Ed., "Problem Statement for Service Function Chaining", [RFC 7498](#), DOI 10.17487/RFC7498, April 2015, <<http://www.rfc-editor.org/info/rfc7498>>.
- [RFC7665] Halpern, J., Ed. and C. Pignataro, Ed., "Service Function Chaining (SFC) Architecture", [RFC 7665](#), DOI 10.17487/RFC7665, October 2015, <<http://www.rfc-editor.org/info/rfc7665>>.

Jacquenet & Boucadair

Expires July 10, 2016

[Page 11]

Internet-Draft

[draft-jacquenet-ipv6-eh-sfc-00.txt](#)

January 2016

Authors' Addresses

Christian Jacquenet
Orange

Email: christian.jacquenet@orange.com

Mohamed Boucadair
Orange

Email: mohamed.boucadair@orange.com

