Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 25, 2019

C. Filsfils, Ed.
Cisco Systems, Inc.
S. Previdi
Huawei
J. Leddy
Individual
S. Matsushima
Softbank
D. Voyer, Ed.
Bell Canada

October 22, 2018

IPv6 Segment Routing Header (SRH)
draft-ietf-6man-segment-routing-header-15

Abstract

Segment Routing can be applied to the IPv6 data plane using a new type of Routing Extension Header. This document describes the Segment Routing Extension Header and how it is used by Segment Routing capable nodes.

Requirements Language

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.

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 April 25, 2019.

Table of Contents

1. Introduction

Segment Routing can be applied to the IPv6 data plane using a new type of Routing Extension Header (SRH). This document describes the Segment Routing Extension Header and how it is used by Segment Routing capable nodes.

The Segment Routing Architecture [RFC8402] describes Segment Routing and its instantiation in two data planes MPLS and IPv6.

SR with the MPLS data plane is defined in [I-D.ietf-spring-segment-routing-mpls].

SR with the IPv6 data plane is defined in [I-D.filsfils-spring-srv6-network-programming].

The encoding of MPLS labels and label stacking are defined in [RFC3032].

The encoding of IPv6 segments in the Segment Routing Extension Header is defined in this document.

Terminology used within this document is defined in detail in [RFC8402]. Specifically, these terms: Segment Routing, SR Domain, SRv6, Segment ID (SID), SRv6 SID, Active Segment, and SR Policy.

2. Segment Routing Extension Header

Routing Headers are defined in [RFC8200]. The Segment Routing Header has a new Routing Type (suggested value 4) to be assigned by IANA.

The Segment Routing Header (SRH) is defined as follows:

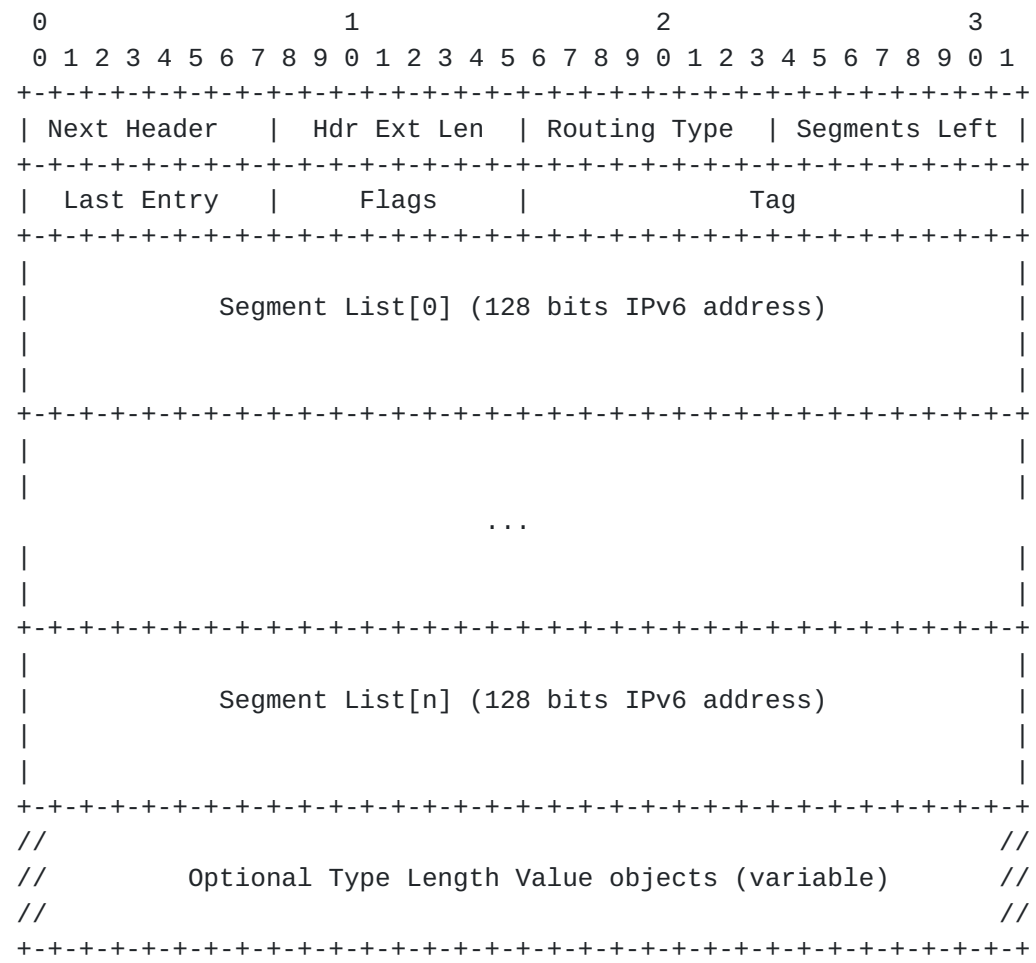

where:

- Next Header: Defined in [RFC8200]
- Hdr Ext Len: Defined in [RFC8200]
- Routing Type: TBD, to be assigned by IANA (suggested value: 4).
- Segments Left: Defined in [RFC8200]
- Last Entry: contains the index (zero based), in the Segment List, of the last element of the Segment List.

- Flags: 8 bits of flags. Following flags are defined:

```
 0 1 2 3 4 5 6 7
+-+-+-+-+-+-+-+-+
|U U U U U U U U|
+-+-+-+-+-+-+-+-+
```

  U: Unused and for future use. MUST be 0 on transmission and ignored on receipt.

- Tag: tag a packet as part of a class or group of packets, e.g., packets sharing the same set of properties. When tag is not used at source it MUST be set to zero on transmission. When tag is not used during SRH Processing it SHOULD be ignored. The allocation and use of tag is outside the scope of this document.

- Segment List[n]: 128 bit IPv6 addresses representing the nth segment in the Segment List. The Segment List is encoded starting from the last segment of the SR Policy. I.e., 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.

- Type Length Value (TLV) are described in Section 2.1.

2.1. SRH TLVs

This section defines TLVs of the Segment Routing Header.

```
 0                   1
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-----------------------
|     Type      |    Length     | Variable length data
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-----------------------
```

Type: An 8 bit value. Unrecognized Types MUST be ignored on receipt.

Length: The length of the Variable length data. It is RECOMMENDED that the total length of new TLVs be multiple of 8 bytes to avoid the use of Padding TLVs.

Variable length data: Length bytes of data that is specific to the Type.

Type Length Value (TLV) contain OPTIONAL information that may be used by the node identified in the Destination Address (DA) of the packet.

Each TLV has its own length, format and semantic. The code-point allocated (by IANA) to each TLV Type defines both the format and the semantic of the information carried in the TLV. Multiple TLVs may be encoded in the same SRH.

TLVs may change en route at each segment. To identify when a TLV type may change en route the most significant bit of the Type has the following significance:

0: TLV data does not change en route

1: TLV data does change en route

Identifying which TLVs change en route, without having to understand the Type, is required for Authentication Header Integrity Check Value (ICV) computation. Any TLV that changes en route is considered mutable for the purpose of ICV computation, the Type Length and Variable Length Data is ignored for the purpose of ICV Computation as defined in [RFC4302].

The "Length" field of the TLV is used to skip the TLV while inspecting the SRH in case the node doesn't support or recognize the Type. The "Length" defines the TLV length in octets, not including the "Type" and "Length" fields.

The following TLVs are defined in this document:

Padding TLV

HMAC TLV

Additional TLVs may be defined in the future.

2.1.1. Padding TLVs

There are two types of padding TLVs, pad0 and padN, the following applies to both:

Padding TLVs are used to pad the TLVs to a multiple of 8 octets.

More than one Padding TLV MUST NOT appear in the SRH.

The Padding TLVs are used to align the SRH total length on the 8 octet boundary.

When present, a single Pad0 or PadN TLV MUST appear as the last TLV.

When present, a PadN TLV MUST have a length from 0 to 5 in order to align the SRH total length on a 8-octet boundary.

Padding TLVs are ignored by a node processing the SRH TLV, even if more than one is present.

Padding TLVs are ignored during ICV calculation.

2.1.1.1. PAD0

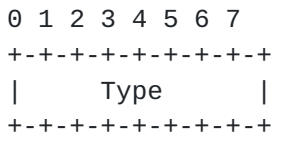

Type: to be assigned by IANA (Suggested value 128)

A single Pad0 TLV MUST be used when a single byte of padding is required. If more than one byte of padding is required a Pad0 TLV MUST NOT be used, the PadN TLV MUST be used.

2.1.1.2. PADN

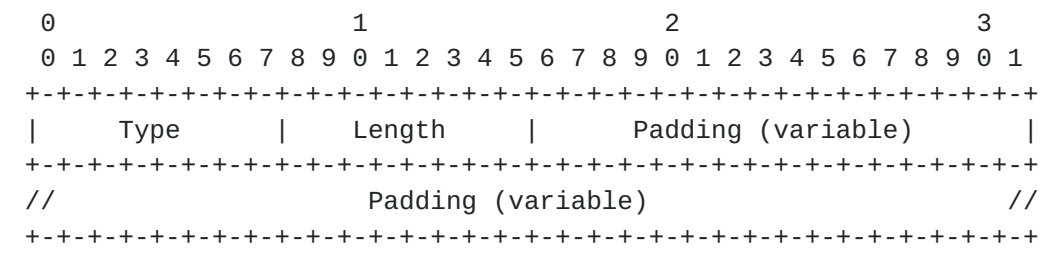

Type: to be assigned by IANA (suggested value 129).

Length: 0 to 5

Padding: Length octets of padding. Padding bits have no semantics. They MUST be set to 0 on transmission and ignored on receipt.

The PadN TLV MUST be used when more than one byte of padding is required.

2.1.2. HMAC TLV

The keyed Hashed Message Authentication Code (HMAC) TLV is OPTIONAL and has the following format:

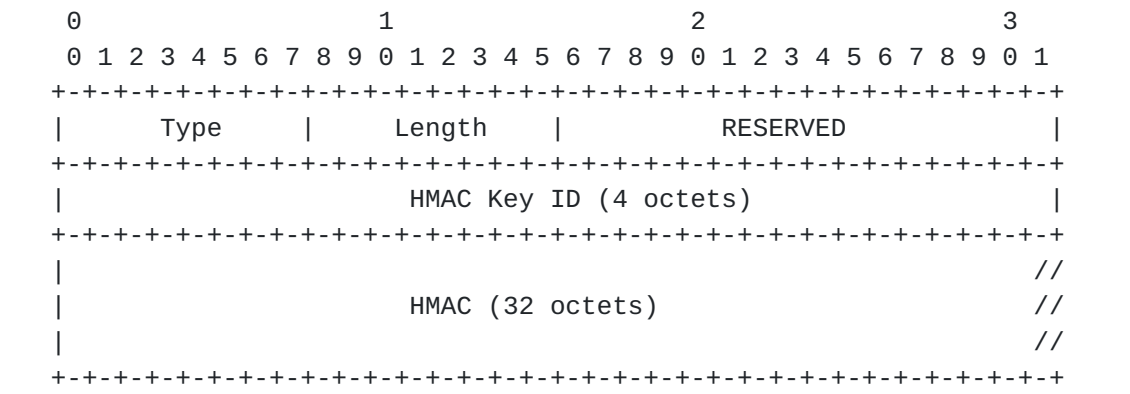

where:

- Type: to be assigned by IANA (suggested value 5).
- Length: 38.
- RESERVED: 2 octets. MUST be 0 on transmission and ignored on receipt.
- HMAC Key ID: A 4 octet opaque number which uniquely identifies the pre-shared key and algorithm used to generate the HMAC. If 0, the HMAC is not included.
- HMAC: 32 octets of keyed HMAC, not present if Key ID is 0.

The HMAC TLV is used to verify the source of a packet is permitted to use the current segment in the destination address of the packet, and ensure the segment list is not modified in transit.

2.1.2.1. HMAC generation

The HMAC field is the output of the HMAC computation as defined in [RFC2104], using:

- key: the pre-shared key identified by HMAC Key ID
- HMAC algorithm: identified by the HMAC Key ID
- Text: a concatenation of the following fields from the IPv6 header and the SRH, as it would be received at the node verifying the HMAC:
  - IPv6 header: source address (16 octets)
  - IPv6 header: destination address (16 octets)

* SRH: Segments Left (1 octet)

* SRH: Last Entry (1 octet)

* SRH: Flags (1 octet)

* SRH: HMAC Key-id (4 octets)

* SRH: all addresses in the Segment List (variable octets)

The HMAC digest is truncated to 32 octets and placed in the HMAC field of the HMAC TLV.

For HMAC algorithms producing digests less than 32 octets, the digest is placed in the lowest order octets of the HMAC field. Remaining octets MUST be set to zero.

2.1.2.2. HMAC Verification

Local policy determines when to check for an HMAC and potentially a requirement on where the HMAC TLV must appear (e.g. first TLV). This local policy is outside the scope of this document. It may be based on the active segment at an SR Segment endpoint node, the result of an ACL that considers incoming interface, or other packet fields.

If HMAC verification is successful, the packet is forwarded to the next segment.

If HMAC verification fails, an ICMP error message (parameter problem, error code 0, pointing to the HMAC TLV) SHOULD be generated (but rate limited) and SHOULD be logged.

2.1.2.3. HMAC Pre-Shared Key Algorithm

The HMAC Key ID field allows for the simultaneous existence of several hash algorithms (SHA-256, SHA3-256 ... or future ones) as well as pre-shared keys.

The HMAC Key ID field is opaque, i.e., it has neither syntax nor semantic except as an identifier of the right combination of pre-shared key and hash algorithm, and except that a value of 0 means that there is no HMAC field.

At the HMAC TLV verification node the Key ID uniquely identifies the pre-shared key and HMAC algorithm.

At the HMAC TLV generating node the Key ID and destination address uniquely identify the pre-shared key and HMAC algorithm. Utilizing the destination address with the Key ID allows for overlapping key IDs amongst different HMAC verification nodes. The Text for the HMAC computation is set to the IPv6 header fields and SRH fields as they would appear at the verification node, not necessarily the same as the source node sending a packet with the HMAC TLV.

Pre-shared key roll-over is supported by having two key IDs in use while the HMAC TLV generating node and verifying node converge to a new key.

SRH implementations can support multiple hash functions but MUST implement SHA-2 [FIPS180-4] in its SHA-256 variant.

The selection of pre-shared key and algorithm, and their distribution is outside the scope of this document, some options may include:

- in the configuration of the HMAC generating or verifying nodes, either by static configuration or any SDN oriented approach

- dynamically using a trusted key distribution protocol such as [RFC6407]

3. SR Nodes

There are different types of nodes that may be involved in segment routing networks: source SR nodes originate packets with a segment in the destination address of the IPv6 header, transit nodes that forward packets destined to a remote segment, and SR segment endpoint nodes that process a local segment in the destination address of an IPv6 header.

3.1. Source SR Node

A Source SR Node is any node that originates an IPv6 packet with a segment (i.e. SRv6 SID) in the destination address of the IPv6 header. The packet leaving the source SR Node may or may not contain an SRH. This includes either:

A host originating an IPv6 packet.

An SR domain ingress router encapsulating a received packet in an outer IPv6 header, followed by an optional SRH.

The mechanism through which a segment in the destination address of the IPv6 header and the Segment List in the SRH, is derived is outside the scope of this document.

3.2. Transit Node

A transit node is any node forwarding an IPv6 packet where the destination address of that packet is not locally configured as a segment nor a local interface. A transit node is not required to be capable of processing a segment nor SRH.

3.3. SR Segment Endpoint Node

A SR segment endpoint node is any node receiving an IPv6 packet where the destination address of that packet is locally configured as a segment or local interface.

4. Packet Processing

This section describes SRv6 packet processing at the SR source, Transit and SR segment endpoint nodes.

4.1. Source SR Node

A Source node steers a packet into an SR Policy. If the SR Policy results in a segment list containing a single segment, and there is no need to add information to SRH flag or TLV, the DA is set to the single segment list entry and the SRH MAY be omitted.

When needed, the SRH is created as follows:

Next Header and Hdr Ext Len fields are set as specified in [RFC8200].

Routing Type field is set as TBD (to be allocated by IANA, suggested value 4).

The DA of the packet is set with the value of the first segment.

The first element of the SRH Segment List is the ultimate segment. The second element is the penultimate segment and so on.

The Segments Left field is set to n-1 where n is the number of elements in the SR Policy.

The Last Entry field is set to n-1 where n is the number of elements in the SR Policy.

HMAC TLV may be set according to Section 7.

The packet is forwarded toward the packet's Destination Address (the first segment).

4.1.1. Reduced SRH

When a source does not require the entire SID list to be preserved in the SRH, a reduced SRH may be used.

A reduced SRH does not contain the first segment of the related SR Policy (the first segment is the one already in the DA of the IPv6 header), and the Last Entry field is set to n-2 where n is the number of elements in the SR Policy.

4.2. Transit Node

As specified in [RFC8200], the only node allowed to inspect the Routing Extension Header (and therefore the SRH), is the node corresponding to the DA of the packet. Any other transit node MUST NOT inspect the underneath routing header and MUST forward the packet toward the DA according to its IPv6 routing table.

When a SID is in the destination address of an IPv6 header of a packet, it's routed through an IPv6 network as an IPv6 address. SIDs, or the prefix(es) covering SIDs, and their reachability may be distributed by means outside the scope of this document. For example, [RFC5308] or [RFC5340] may be used to advertise a prefix covering the SIDs on a node.

4.3. SR Segment Endpoint Node

Without constraining the details of an implementation, the SR segment endpoint node creates Forwarding Information Base (FIB) entries for its local SIDs.

When an SRv6-capable node receives an IPv6 packet, it performs a longest-prefix-match lookup on the packets destination address. This lookup can return any of the following:

A FIB entry that represents a locally instantiated SRv6 SID
A FIB entry that represents a local interface, not locally instantiated as an SRv6 SID
A FIB entry that represents a non-local route
No Match

4.3.1. FIB Entry Is Locally Instantiated SRv6 END SID

This document, and section, defines a single SRv6 SID called END. Future documents may define additional SRv6 SIDs. In which case, the entire content of this section will be defined in that document.

If the FIB entry represents a locally instantiated SRv6 SID, process the next header of the IPv6 header as defined in section 4 of [RFC8200]

The following sections describe the actions to take while processing next header fields.

4.3.1.1. SRH Processing

```
When an SRH is processed {
  If Segments Left is equal to zero {
    Proceed to process the next header in the packet, whose type
    is identified by the Next Header field in the Routing header.
  }
  Else {
    If local policy requires TLV processing {
      Perform TLV processing (see TLV Processing)
    }
    max_last_entry  =  ( Hdr Ext Len /  2 ) - 1

    If  ((Last Entry > max_last_entry) or
         (Segments Left is greater than (Last Entry+1)) {
      Send an ICMP Parameter Problem, Code 0, message to the
      Source Address, pointing to the Segments Left field, and
      discard the packet.
    }
    Else {
      Decrement Segments Left by 1.
      Copy Segment List[Segments Left] from the SRH to the
      destination address of the IPv6 header.
      If the IPv6 Hop Limit is less than or equal to 1 {
        Send an ICMP Time Exceeded -- Hop Limit Exceeded in
        Transit message to the Source Address and discard
        the packet.
      }
      Else {
        Decrement the Hop Limit by 1
        Resubmit the packet to the IPv6 module for transmission
        to the new destination.
      }
    }
  }
}
```

4.3.1.1.1. TLV Processing

Local policy determines how TLV's are to be processed when the Active Segment is a local END SID. The definition of local policy is outside the scope of this document.

For illustration purpose only, two example local policies that may be associated with an END SID are provided below.

```
Example 1:
For any packet received from interface I2
  Skip TLV processing

Example 2:
For any packet received from interface I1
  If first TLV is HMAC {
    Process the HMAC TLV
  }
  Else {
    Discard the packet
  }
```

4.3.1.2. Upper-layer Header or No Next Header

Send an ICMP parameter problem message to the Source Address and discard the packet. Error code (TBD by IANA) "SR Upper-layer Header Error", pointer set to the offset of the upper-layer header.

A unique error code allows an SR Source node to recognize an error in SID processing at an endpoint.

4.3.2. FIB Entry is a Local Interface

If the FIB entry represents a local interface, not locally instantiated as an SRv6 SID, the SRH is processed as follows:

> If Segments Left is zero, the node must ignore the Routing header and proceed to process the next header in the packet, whose type is identified by the Next Header field in the Routing Header.
>
> If Segments Left is non-zero, the node must discard the packet and send an ICMP Parameter Problem, Code 0, message to the packet's Source Address, pointing to the unrecognized Routing Type.

4.3.3. FIB Entry Is A Non-Local Route

Processing is not changed by this document.

4.3.4. FIB Entry Is A No Match

Processing is not changed by this document.

4.3.5. Load Balancing and ECMP

Within an SR domain, an SR source node encapsulates a packet in an outer IPv6 header for transport to an endpoint. The SR source node MUST impose a flow label computed based on the inner packet. The computation of the flow label is as recommended in [RFC6438] for the sending Tunnel End Point.

At any transit node within an SR domain, the flow label MUST be used as defined in [RFC6438] to calculate the ECMP hash toward the destination address. If flow label is not used, the transit node may hash all packets between a pair of SR Edge nodes to the same link.

At an SR segment endpoint node, the flow label MUST be used as defined in [RFC6438] to calculate any ECMP hash used to forward the processed packet to the next segment.

5. Illustrations

This section provides illustrations of SRv6 packet processing at SR source, transit and SR segment endpoint nodes.

5.1. Abstract Representation of an SRH

For a node k, its IPv6 address is represented as Ak, its SRv6 SID is represented as Sk.

IPv6 headers are represented as the tuple of (source, destination). For example, a packet with source address A1 and destination address A2 is represented as (A1,A2). The payload of the packet is omitted.

An SR Policy is a list of segments. A list of segments is represented as <S1,S2,S3> where S1 is the first SID to visit, S2 is the second SID to visit and S3 is the last SID to visit.

(SA,DA) (S3, S2, S1; SL) represents an IPv6 packet with:

- Source Address is SA, Destination Addresses is DA, and next-header is SRH.

- SRH with SID list <S1, S2, S3> with SegmentsLeft = SL.

- Note the difference between the <> and () symbols. <S1, S2, S3> represents a SID list where the leftmost segment is the first segment. Whereas, (S3, S2, S1; SL) represents the same SID list but encoded in the SRH Segment List format where the leftmost segment is the last segment. When referring to an SR policy in a high-level use-case, it is simpler to use the <S1, S2, S3> notation. When referring to an illustration of detailed behavior, the (S3, S2, S1; SL) notation is more convenient.

At its SR Policy headend, the Segment List <S1,S2,S3> results in SRH (S3,S2,S1; SL=2) represented fully as:

```
Segments Left=2
Last Entry=2
Flags=0
Tag=0
Segment List[0]=S3
Segment List[1]=S2
Segment List[2]=S1
```

5.2. Example Topology

The following topology is used in examples below:

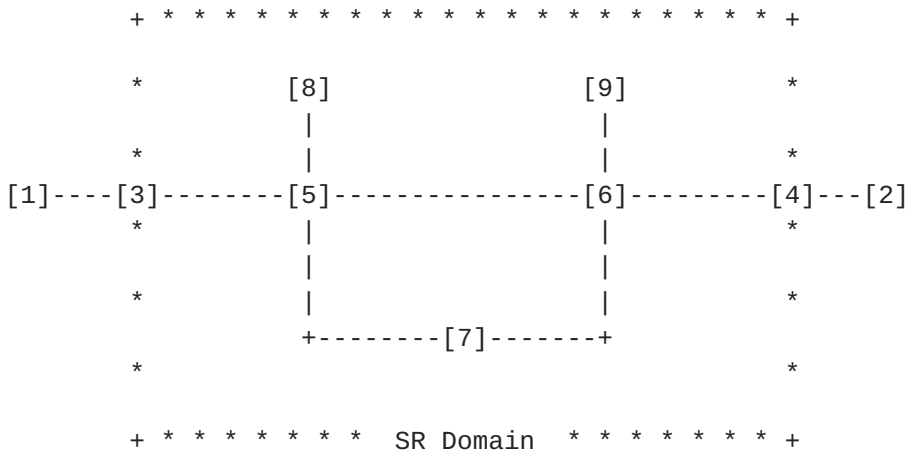

Figure 3

- 3 and 4 are SR Domain edge routers

- 5, 6, and 7 are all SR Domain routers

- 8 and 9 are hosts within the SR Domain

- 1 and 2 are hosts outside the SR Domain

5.3. Source SR Node

5.3.1. Intra SR Domain Packet

When host 8 sends a packet to host 9 via an SR Policy <S7,A9> the packet is

P1: (A8,S7)(A9,S7; SL=1)

5.3.1.1. Reduced Variant

When host 8 sends a packet to host 9 via an SR Policy <S7,A9> and it wants to use a reduced SRH, the packet is

P2: (A8,S7)(A9; SL=1)

5.3.2. Transit Packet Through SR Domain

When host 1 sends a packet to host 2, the packet is

P3: (A1,A2)

The SR Domain ingress router 3 receives P3 and steers it to SR Domain egress router 4 via an SR Policy <S7, S4>. Router 3 encapsulates the received packet P3 in an outer header with an SRH. The packet is

P4: (A3, S7)(S4, S7; SL=1)(A1, A2)

If the SR Policy contains only one segment (the egress router 4), the ingress Router 3 encapsulates P3 into an outer header (A3, S4). The packet is

P5: (A3, S4)(A1, A2)

5.3.2.1. Reduced Variant

The SR Domain ingress router 3 receives P3 and steers it to SR Domain egress router 4 via an SR Policy <S7, S4>. If router 3 wants to use a reduced SRH, Router 3 encapsulates the received packet P3 in an outer header with a reduced SRH. The packet is

P6: (A3, S7)(S4; SL=1)(A1, A2)

5.4. Transit Node

Nodes 5 acts as transit nodes for packet P1, and sends packet

P1: (A8,S7)(A9,S7;SL=1)

on the interface toward node 7.

5.5. SR Segment Endpoint Node

Node 7 receives packet P1 and, using the logic in section 4.3.1, sends packet

P7: (A8,A9)(A9,S7; SL=0)

on the interface toward router 6.

6. Deployment Models

6.1. Nodes Within the SR domain

SR Source Nodes within an SR Domain are trusted to generate IPv6 packets with SRH. SR segment endpoint nodes receiving packets on interface that are part of the SR Domain may process any packet destined to a local segment, containing an SRH.

A SR Source Node connected to the SR Domain via a secure tunnel, e.g. IPSec tunnel mode [RFC4303] or Ethernet pseudowire [RFC4448], may be considered trusted and directly connected. Some types of tunnels may result in additional processing overhead that should be considered in a deployment.

6.2. Nodes Outside the SR Domain

Nodes outside the SR Domain cannot be trusted. SR Domain Ingress routers SHOULD discard packets destined to SIDs within the SR Domain (regardless of the presence of an SRH) to avoid attacks on the SR Domain as described and referenced in [RFC5095]. As an additional layer of protection, SR Segment Endpoint nodes SHOULD discard packets destined to local SIDs from source addresses not part of the SR Domain.

For example, using the example topology from section 5, all SIDs in the SR Domain (SIDS S1-S9) are assigned within a single IPv6 prefix, Prefix-S. All SIDs assigned to a node k are assigned within a single IPv6 prefix Prefix-Sk, all addresses permitted to source packets destined to SIDs in the SR Domain are assigned within a single IPv6 prefix Prefix-A.

An Infrastructure Access List (IACL), applied to the external interfaces of SR Domain ingress nodes 3 and 4, that discards packets destined to a SID covered by Prefix-S is used to discard packets destined to SIDs within the SR Domain.

An IACL, applied to each interface of SR Segment Endpoint Nodes k, that discards packets destined to a SID covered by Prefix-Sk with a source address not covered by Prefix-A.

Failure to implement a method of ingress filtering, as defined above, exposes the SR domain to source routing attacks from nodes outside the SR Domain, as described and referenced in [RFC5095].

6.2.1. SR Source Nodes Not Directly Connected

Nodes outside the SR Domain may request, by some trusted means outside the scope of this document, a complete SRH including an HMAC TLV which is computed correctly for the SRH.

SR Domain ingress routers permit traffic destined to select SIDs with local policy requiring HMAC TLV processing for those select SIDs, i.e. those SIDs provide a gateway to the SR Domain for a set of segment lists.

If HMAC verification is successful, the packet is forwarded to the next segment. Within the SR Domain no further HMAC check need be performed.

If HMAC verification fails, an ICMP error message (parameter problem, error code 0, pointing to the HMAC TLV) SHOULD be generated (but rate limited) and SHOULD be logged.

For example, extending the topology defined in Figure 3, consider node 3 offering access to a premium SLA service to node 20. Node 20 is a trusted SR Source not directly connected to the SR Domain.

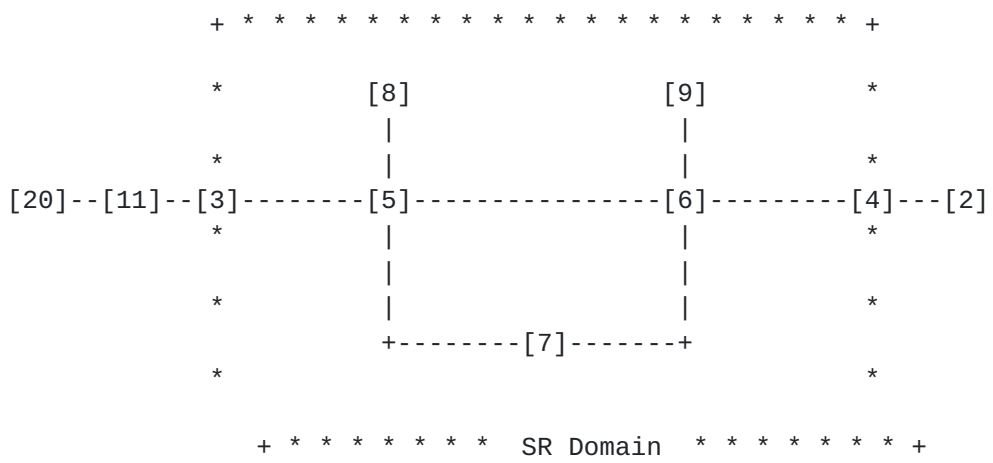

In order to access the SLA service, node 20 must be able to access segments within the SR Domain. To provide a secure entry point for the SLA service, SIDs with local policy requiring HMAC verification at node k are defined as Hk and assigned from a prefix Prefix-H. Prefix-H is disjoint with Prefix-S and Prefix-A defined earlier.

Prefix-H is not part of the IACLs applied at the external facing interfaces of node 3 and 4, allowing external nodes access to it.

SID H3 is a SID covered by Prefix-H at node 3.

Node 20 requests the premium SLA service to node 2 and is provided a pre-computed SRH and HMAC with destination address H3.

Node 20 sends a packet with destination addresses set to H2, SRH and HMAC TLV are as provided for the premium SLA service.

Node 3 receives the packet and verifies the HMAC as defined in section 4.3, forwarding the packet to the next segment in the segment list or dropping it based on the HMAC result.

This use of an HMAC is particularly valuable within an enterprise based SR Domain to authenticate a host which is using SRv6 segment routing as documented in [SRN]. In that example, the HMAC is used to validate a source node is using a permitted segment list.

7. Security Considerations

This section reviews security considerations related to the SRH, given the SRH processing and deployment models discussed in this document.

As describe in Section 6, it is necessary to filter packets ingress to the SR Domain destined to segments within the SR Domain. This

ingress filtering is via an IACL at SR Domain ingress border nodes. Additional protection is applied via an IACL at each SR Segment Endpoint node, filtering packets not from within the SR Domain, destined to SIDs in the SR Domain. ACLs are easily supported for small numbers of prefixes, making summarization important, and when the prefixes requiring filtering is kept to a seldom changing set.

Additionally, ingress filtering of IPv6 source addresses as recommended in BCP38 SHOULD be used.

SR Source Nodes not directly connected to the SR Domain may access specific sets of segments within the SR Domain when secured with the SRH HMAC TLV. The SRH HMAC TLV provides a means of verifying the validity of ingress packets SRH, limiting access to the segments in the SR Domain to only those source nodes with permission.

7.1. Source Routing Attacks

[RFC5095] deprecates the Type 0 Routing header due to a number of significant attacks that are referenced in that document. Such attacks include bypassing filtering devices, reaching otherwise unreachable Internet systems, network topology discovery, bandwidth exhaustion, and defeating anycast.

Because this document specifies that the SRH is for use within an SR domain protected by ingress filtering via IACLs, and by cryptographically authenticated SR source nodes not directly connected to the SR Domain; such attacks cannot be mounted from outside an SR Domain. As specified in this document, SR Domain ingress edge nodes drop packets entering the SR Domain destined to segments within the SR Domain.

Aditionally, this document specifies the use of IACL on SR Segment Endpoint nodes within the SR Domain to limit the source addresses permitted to send packets to a SID in the SR Domain.

Such attacks may, however, be mounted from within the SR Domain, from nodes permitted to source traffic to SIDs in the domain. As such, these attacks and other known attacks on an IP network (e.g. DOS/ DDOS, topology discovery, man-in-the-middle, traffic interception/ siphoning), can occur from compromised nodes within an SR Domain.

7.2. Service Theft

Service theft is defined as the use of a service offered by the SR Domain by a node not authorized to use the service.

Service theft is not a concern within the SR Domain as all SR Source nodes and SR segment endpoint nodes within the domain are able to utilizing the services of the Domain. If a node outside the SR Domain learns of segments or a topological service within the SR domain, IACL filtering denies access to those segments.

Nodes outside the SR Domain, capable of intercepting packets from SR Source nodes not directly connected to the SR Domain utilizing the SRH HMAC, may steel the outer IP header SRH and HMAC TLV. If such an attacker is capable of spoofing the source address of the original sender it may use the IP header and HMAC to access services of the SR Domain intended for the original SR Source node.

Frequent rekeying of the HMAC TLV helps mitigate against this attack but cannot prevent it.

However, as described in Section 6.2.1, there exist use cases where the risk of service threat is of minimum concern and the HMAC TLV is used primarily to validate that the source is permitted to use the segment list in the SRH.

7.3. Topology Disclosure

The SRH may contains SIDs of some intermediate SR-nodes in the path towards the destination, this reveals those addresses to attackers if they are able to intercept packets containing SRH.

This is applicable within an SR Domain but the disclosure is less relevant as an attacker has other means of learning topology.

For an SR Source node not directly connected to the SR Domain this disclosure is applicable. While the segments within the SR domain disclosed in SRH are protected by ingress filtering, they may be learned by an attacker external to the SR Domain.

As described in Section 6.2.1, there exist use cases where the risk of topology disclosure is of minimum concern when the HMAC TLV is used primarily to validate that the source is permitted to use the segment list in the SRH.

7.4. ICMP Generation

The generation of ICMPv6 error messages may be used to attempt denial-of-service attacks by sending an error-causing destination address or SRH in back-to-back packets. An implementation that correctly follows Section 2.4 of [RFC4443] would be protected by the ICMPv6 rate-limiting mechanism.

8. IANA Considerations

This document makes the following registrations in the Internet Protocol Version 6 (IPv6) Parameters "Routing Type" registry maintained by IANA:

Suggested Value	Description	Reference
4	Segment Routing Header (SRH)	This document

This document request IANA to create and maintain a new Registry: "Segment Routing Header TLVs"

8.1. Segment Routing Header Flags Register

This document requests the creation of a new IANA managed registry to identify SRH Flags Bits. The registration procedure is "Expert Review" as defined in [RFC8126]. Suggested registry name is "Segment Routing Header Flags". Flags is 8 bits, the following bits are defined in this document:

Suggested Bit	Description	Reference
4	HMAC	This document

8.2. Segment Routing Header TLVs Register

This document requests the creation of a new IANA managed registry to identify SRH TLVs. The registration procedure is "Expert Review" as defined in [RFC8126]. Suggested registry name is "Segment Routing Header TLVs". A TLV is identified through an unsigned 8 bit codepoint value. The following codepoints are defined in this document:

Suggested Value	Description	Reference
5	HMAC TLV	This document
128	Pad0 TLV	This document
129	PadN TLV	This document

9. Implementation Status

This section is to be removed prior to publishing as an RFC.

9.1. Linux

Name: Linux Kernel v4.14

Status: Production

Implementation: adds SRH, performs END processing, supports HMAC TLV

Details: https://irtf.org/anrw/2017/anrw17-final3.pdf and [I-D.filsfils-spring-srv6-interop]

9.2. Cisco Systems

Name: IOS XR and IOS XE

Status: Pre-production

Implementation: adds SRH, performs END processing, no TLV processing

Details: [I-D.filsfils-spring-srv6-interop]

9.3. FD.io

Name: VPP/Segment Routing for IPv6

Status: Production

Implementation: adds SRH, performs END processing, no TLV processing

Details: https://wiki.fd.io/view/VPP/Segment_Routing_for_IPv6 and [I-D.filsfils-spring-srv6-interop]

9.4. Barefoot

Name: Barefoot Networks Tofino NPU

Status: Prototype

Implementation: performs END processing, no TLV processing

Details: [I-D.filsfils-spring-srv6-interop]

9.5. Juniper

Name: Juniper Networks Trio and vTrio NPU's

Status: Prototype & Experimental

Implementation: SRH insertion mode, Process SID where SID is an interface address, no TLV processing

9.6. Huawei

Name: Huawei Systems VRP Platform

Status: Production

Implementation: adds SRH, performs END processing, no TLV processing

10. Contributors

Kamran Raza, Darren Dukes, Brian Field, Daniel Bernier, Ida Leung, Jen Linkova, Ebben Aries, Tomoya Kosugi, Eric Vyncke, David Lebrun, Dirk Steinberg, Robert Raszuk, Dave Barach, John Brzozowski, Pierre Francois, Nagendra Kumar, Mark Townsley, Christian Martin, Roberta Maglione, James Connolly, Aloys Augustin contributed to the content of this document.

11. Acknowledgements

The authors would like to thank Ole Troan, Bob Hinden, Ron Bonica, Fred Baker, Brian Carpenter, Alexandru Petrescu, Punit Kumar Jaiswal, and David Lebrun for their comments to this document.

12. References

12.1. Normative References

[FIPS180-4] National Institute of Standards and Technology, "FIPS 180-4 Secure Hash Standard (SHS)", March 2012, <http://csrc.nist.gov/publications/fips/fips180-4/fips-180-4.pdf>.

[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>.

[RFC5095] Abley, J., Savola, P., and G. Neville-Neil, "Deprecation of Type 0 Routing Headers in IPv6", RFC 5095, DOI 10.17487/RFC5095, December 2007, <https://www.rfc-editor.org/info/rfc5095>.

[RFC6407] Weis, B., Rowles, S., and T. Hardjono, "The Group Domain of Interpretation", RFC 6407, DOI 10.17487/RFC6407, October 2011, <https://www.rfc-editor.org/info/rfc6407>.

[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>.

[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>.

12.2. Informative References

[I-D.filsfils-spring-srv6-interop]
Filsfils, C., Clad, F., Camarillo, P., Abdelsalam, A., Salsano, S., Bonaventure, O., Horn, J., and J. Liste, "SRv6 interoperability report", draft-filsfils-spring-srv6-interop-01 (work in progress), September 2018.

[I-D.filsfils-spring-srv6-network-programming]
Filsfils, C., Camarillo, P., Leddy, J., daniel.voyer@bell.ca, d., Matsushima, S., and Z. Li, "SRv6 Network Programming", draft-filsfils-spring-srv6-network-programming-05 (work in progress), July 2018.

[I-D.ietf-spring-segment-routing-mpls]
Bashandy, A., Filsfils, C., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing with MPLS data plane", draft-ietf-spring-segment-routing-mpls-14 (work in progress), June 2018.

[RFC2104] Krawczyk, H., Bellare, M., and R. Canetti, "HMAC: Keyed-Hashing for Message Authentication", RFC 2104, DOI 10.17487/RFC2104, February 1997, <https://www.rfc-editor.org/info/rfc2104>.

[RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, DOI 10.17487/RFC3032, January 2001, <https://www.rfc-editor.org/info/rfc3032>.

[RFC4302] Kent, S., "IP Authentication Header", RFC 4302, DOI 10.17487/RFC4302, December 2005, <https://www.rfc-editor.org/info/rfc4302>.

[RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, DOI 10.17487/RFC4303, December 2005, <https://www.rfc-editor.org/info/rfc4303>.

[RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", STD 89, RFC 4443, DOI 10.17487/RFC4443, March 2006, <https://www.rfc-editor.org/info/rfc4443>.

[RFC4448] Martini, L., Ed., Rosen, E., El-Aawar, N., and G. Heron, "Encapsulation Methods for Transport of Ethernet over MPLS Networks", RFC 4448, DOI 10.17487/RFC4448, April 2006, <https://www.rfc-editor.org/info/rfc4448>.

[RFC5308] Hopps, C., "Routing IPv6 with IS-IS", RFC 5308, DOI 10.17487/RFC5308, October 2008, <https://www.rfc-editor.org/info/rfc5308>.

[RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, DOI 10.17487/RFC5340, July 2008, <https://www.rfc-editor.org/info/rfc5340>.

[RFC6438] Carpenter, B. and S. Amante, "Using the IPv6 Flow Label for Equal Cost Multipath Routing and Link Aggregation in Tunnels", RFC 6438, DOI 10.17487/RFC6438, November 2011, <https://www.rfc-editor.org/info/rfc6438>.

[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>.

[SRN] and , "Software Resolved Networks: Rethinking Enterprise Networks with IPv6 Segment Routing", 2018, <https://inl.info.ucl.ac.be/system/files/sosr18-final15-embedfonts.pdf>.

Authors' Addresses

Clarence Filsfils (editor)
Cisco Systems, Inc.
Brussels
BE

Email: cfilsfil@cisco.com

Stefano Previdi
Huawei
Italy

Email: stefano@previdi.net

John Leddy
Individual
US

Email: john@leddy.net

Satoru Matsushima
Softbank

Email: satoru.matsushima@g.softbank.co.jp

Daniel Voyer (editor)
Bell Canada

Email: daniel.voyer@bell.ca