Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 6, 2014

S. Previdi, Ed.
C. Filsfils
Cisco Systems, Inc.
B. Field
Comcast
I. Leung
Rogers Communications
March 5, 2014

IPv6 Segment Routing Header (SRH)
draft-previdi-6man-segment-routing-header-00

Abstract

Segment Routing (SR) allows a node to steer a packet through a controlled set of instructions, called segments, by prepending a SR header to the packet. A segment can represent any instruction, topological or service-based. SR allows to enforce a flow through any path (topological, or application/service based) while maintaining per-flow state only at the ingress node to the SR domain.

The Segment Routing architecture can be applied to the IPv6 data plane with the addition of a new type of Routing Extension Header. This draft describes the Segment Routing Extension Header Type and how it is used by SR capable nodes.

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 September 6, 2014.

Copyright Notice

Table of Contents

1. Structure of this document

Section 3 gives an introduction on SR for IPv6 networks.

Section 4 defines the Segment Routing Header (SRH) allowing instantiation of SR over IPv6 dataplane.

Section 5 details the procedures of the Segment Routing Header.

Section 6 describes the security aspect of SR-IPv6.

2. Segment Routing Documents

Segment Routing is described in [I-D.filsfils-rtgwg-segment-routing].

Segment Routing use cases are described in [I-D.filsfils-rtgwg-segment-routing-use-cases].

Segment Routing IPv6 use cases are described in [draft-martin-spring-segment-routing-ipv6-use-cases-00].

Segment Routing protocol extensions are defined in [I-D.previdi-isis-segment-routing-extensions], and [I-D.psenak-ospf-segment-routing-ospfv3-extension].

The terminology is used in this document has been defined in [I-D.filsfils-rtgwg-segment-routing].

3. Introduction

Segment Routing (SR), defined in [I-D.filsfils-rtgwg-segment-routing], allows a node to steer a packet through a controlled set of instructions, called segments, by prepending a SR header to the packet. A segment can represent any instruction, topological or service-based. SR allows to enforce a flow through any path (topological or service/application based) while maintaining per-flow state only at the ingress node to the SR domain. Segments can be derived from different components: IGP, BGP, Services, Contexts, Locators, etc. The list of segment forming the path is called the Segment List and is encoded in the packet header.

SR allows the use of strict and loose source based routing paradigms without requiring any additional signaling protocols in the infrastructure hence delivering an excellent scalability property.

The source based routing model described in

[I-D.filsfils-rtgwg-segment-routing] is inherited from the ones proposed by [RFC1940] and [RFC2460]. The source based routing model offers the support for explicit routing capability.

3.1. Data Planes supporting Segment Routing

Segment Routing (SR), defined in [I-D.filsfils-rtgwg-segment-routing], can be instantiated over MPLS ([I-D.filsfils-spring-segment-routing-mpls]) and IPv6. This document defines its instantiation over the IPv6 data-plane.

Segment Routing for IPv6 (SR-IPv6) is required in networks where MPLS data-plane is not used or, when combined with SR-MPLS, in networks where MPLS is used in the core and IPv6 is used at the edge (home networks, datacenters).

This document defines a new type of Routing Header (originally defined in [RFC2460]) called the Segment Routing Header (SRH) in order to convey the Segment List in the packet header as defined in [I-D.filsfils-rtgwg-segment-routing]. Mechanisms through which segment are known and advertised are outside the scope of this document.

3.2. Illustrative Example

Typically, the domain ingress node obtains the path it has to use for a given packet flow through either local configuration, local computation or through an interaction with an external server such as an SDN controller.

The output of the above is a segment list: a list of IPv6 addresses (each representing a segment) that is encoded in the SRH. The Segment List represents the path of the packet.

The ingress node encodes the first segment into the Destination Address of the IPv6 header and the packet is forwarded towards the first segment endpoint.

Each segment endpoint inspects the SRH, updates the DA (with the next segment) and forwards the packet towards the next segment.

The SRH MAY be removed from the packet prior to send it to its original destination.

When traveling within a segment, a packet may traverse non-SR-capable nodes. These nodes will forward the packet based on its DA regardless the content of the SRH that, in their case, will be silently ignored as mandated by [RFC2460]. Therefore,

interoperability between SR-capable and non-SR-capable nodes being ensured, a gradual deployment of SR in existing networks is possible. The details of the procedures of SR-IPv6 are described in Section 5.

4. IPv6 Instantiation of Segment Routing

When Segment Routing is applied to IPv6, segments are encoded as 128-bit IPv6 addresses. This implies that, in the IPv6 instantiation of SR, the SID values are in fact the prefixes advertised in the IPv6 control-plane. Hence there's no need to advertise any additional specific identifier (other than IPv6 prefix) for the purpose of SR. This simplifies the introduction of IPv6 Segment Routing in existing protocols (i.e.: IS-IS, OSPF and BGP).

4.1. Segment Routing Extension Header (SRH)

A new type of the Routing Header (originally defined in [RFC2460]) is defined: the Segment Routing Header (SRH) which has a new Routing Type, to be assigned by IANA.

According to [I-D.filsfils-rtgwg-segment-routing], each segment is represented by a Segment Identifier (SID). When SR is used over IPv6 networks, the SID is an IPv6 address (or prefix) as learned by IGP, BGP or other protocols.

As an example, if an explicit path is to be constructed across a core network running ISIS or OSPF, the segment list will contain SIDs representing the nodes across the path (loose or strict) which, usually, are the IPv6 loopback interface address of each node. If the path is across service or application entities, the segment list contains the IPv6 addresses of these services or application instances.

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

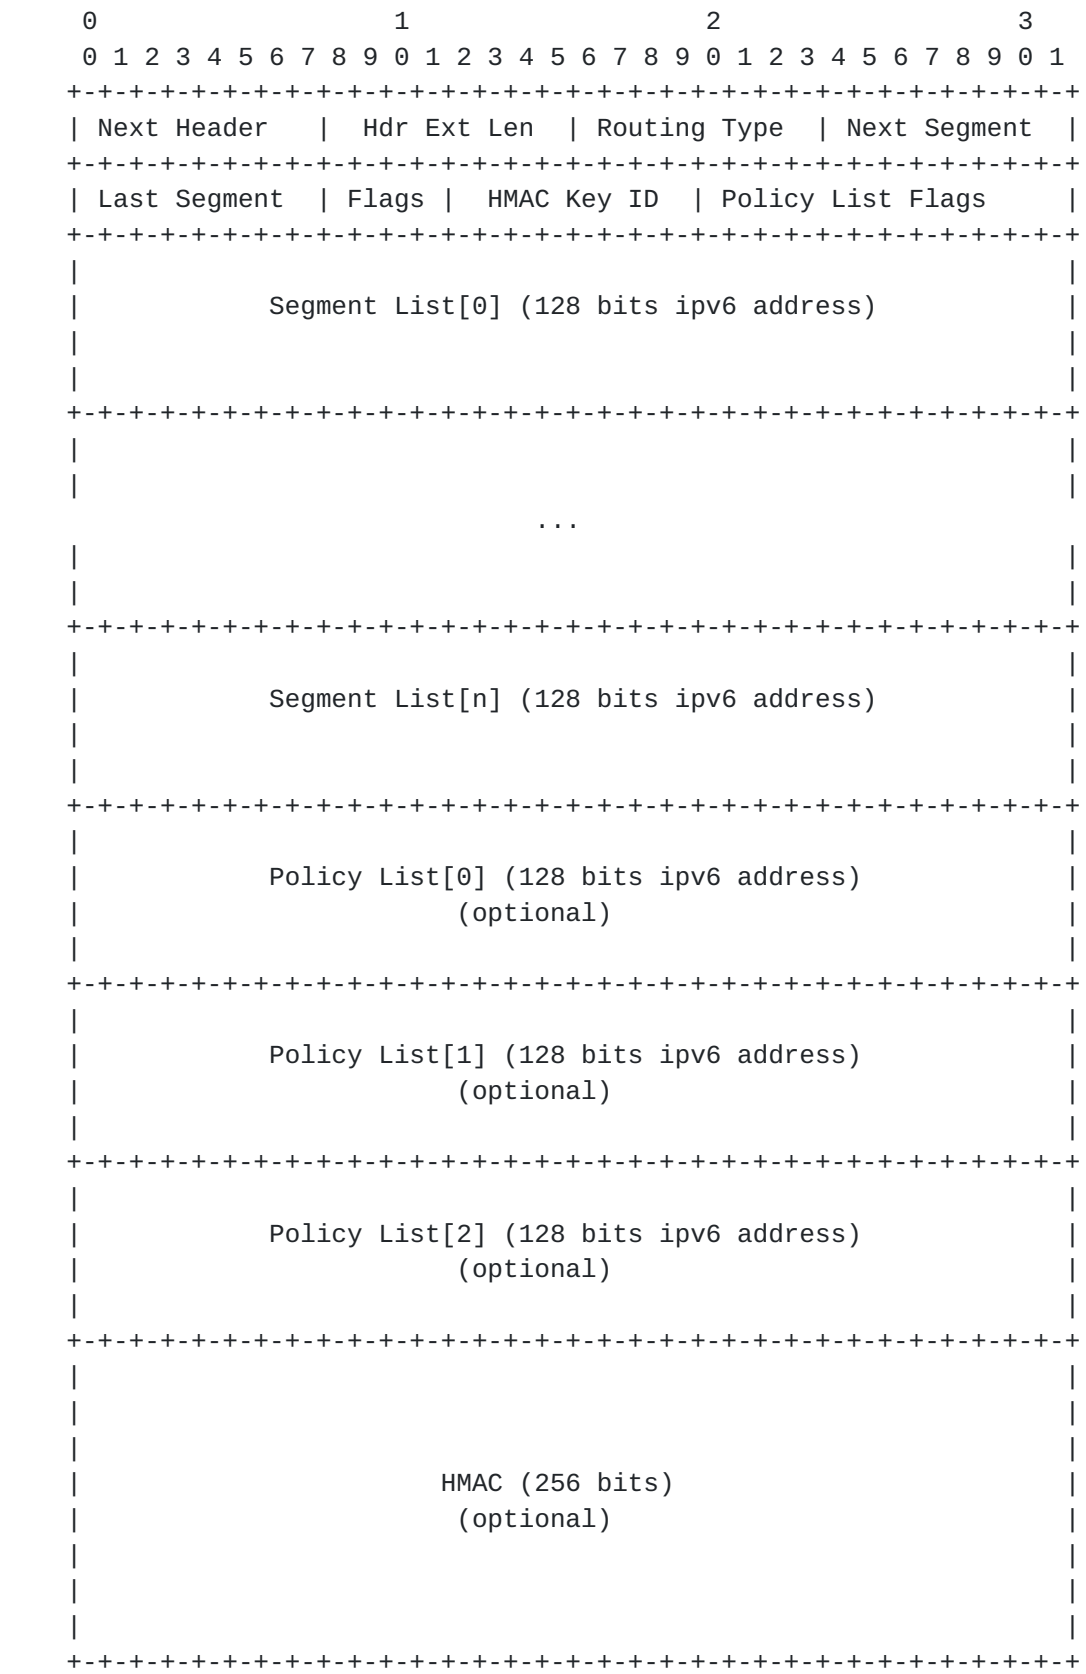
0 1 2 3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
Next Header
Hdr Ext Len
Routing Type
Next Segment
Last Segment
Flags
HMAC Key ID
Policy List Flags
Segment List[0] (128 bits ipv6 address)
...
Segment List[n] (128 bits ipv6 address)
Policy List[0] (128 bits ipv6 address)
(optional)
Policy List[1] (128 bits ipv6 address)
(optional)
Policy List[2] (128 bits ipv6 address)
(optional)
HMAC (256 bits)
(optional)

where:

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

- Hdr Ext Len: 8-bit unsigned integer, is the length of the SRH header in 8-octet units, not including the first 8 octets.

- Routing Type: TBD, to be assigned by IANA.

- Next Segment (originally defined as "Segments Left" in [RFC2460]): offset (in multiple of 8 octets not including the first 8 octets) of the next active segment (according to terminology defined in [I-D.filsfils-rtgwg-segment-routing]) in the SRH. Note that this differs from the semantic defined in the Routing Header specification ([RFC2460] defines it as "Segments Left"). Therefore, in the Segment Routing context, the "Segments Left" field is renamed as "Next Segment".

- Last Segment: offset (in multiple of 8 octets not including the first 8 octets) of the last segment of the path in the SRH.

- Flags: 4 bits of flags. Two flags are defined:

  Bit-0: Clean-up Bit. Set when the SRH has to be removed from the packet when packet reaches the last segment.

  Bit-1: Protected Bit. Set when the packet has been rerouted through FRR mechanism by a SR endpoint node. See Section 5.3 for more details.

- HMAC Key ID and HMAC field are defined in Section 6.

- Policy List flags. Define the type of the IPv6 addresses encoded into the Policy List (see below). The following have been defined:

  Bits 0-2: determine the type of the first element after the segment list.

  Bits 3-5: determine the type of the second element.

  Bits 6-8: determine the type of the third element.

  Bits 9-1: determine the type of the fourth element.

  The following values are used for the type:

0x0: Not present. If value is set to 0x0, it means the element represented by these bits is not present.

0x1: Ingress SR PE address.

0x2: Egress SR PE address.

0x3: Original Source Address.

- Segment List[n]: 128 bit IPv6 addresses representing the nth segment of the path.

- Policy List. Optional addresses representing specific nodes in the SR path such as:

  Ingress SR PE: IPv6 address representing the SR node which has imposed the SRH (SR domain ingress).

  Egress SR PE: IPv6 address representing the egress SR domain node.

  Original Source Address: IPv6 address originally present in the SA field of the packet.

  The segments in the Policy List are encoded after the segment list and they are optional. If none are in the SRH, all bits of the Policy List Flags MUST be set to 0x0.

4.1.1. SRH and RFC2460 behavior

The SRH being a new type of the Routing Header, it also has the same properties:

Can only appear once in the packet.

Only the router whose address is in the DA field of the packet header MUST inspect the SRH.

Therefore, Segment Routing in IPv6 networks implies that the segment identifier (i.e.: the IPv6 address of the segment) is moved into the DA of the packet.

The DA of the packet changes at each segment termination/completion and therefore the original DA of the packet MUST be encoded as the last segment of the path.

As illustrated in Section 3.2, nodes that are within the path of a segment will forward packets based on the DA of the packet without

inspecting the SRH. This ensures full interoperability between SR-capable and non-SR-capable nodes.

4.1.2. SRH Optimization

In order to optimize the way the SRH and, more precisely, the Segment List is processed by SR nodes, it is desirable that most of the necessary information of the SL is placed at the top of the list so to avoid reading the whole content of the SRH prior to make forwarding decisions.

With this in mind, when the SRH is created and the segment list is inserted, the order of the segments in the segment list is as follows:

- The Next Segment field points to the next segment to be examined (offset within the SRH).
- The first segment being encoded in the DA by the ingress node, it doesn't need to sit in the first position of the list.
- Hence, the first element of the segment list is the second segment of the path so that, when the packet reaches the end of the first segment, the node inspecting the SRH will find the second segment at the beginning of the segment list.
- The other segments of the path are encoded sequentially after the second segment.
- The last segment of the path is the original DA address.
- The last segment in the Segment List is used to encode the first segment. This segment will never be inspected anyway (at least not for forwarding purposes).

4.2. Segment Identifiers (SIDs)

The Segment Routing architecture described in [I-D.filsfils-rtgwg-segment-routing], defines Node-SID and Adjacency-SID. When SR is used over IPv6 data-plane the following applies.

4.2.1. Node-SID

The Node-SID identifies a node. With SR-IPv6 the Node-SID is an IPv6 prefix that the operator configured on the node and that is used as the node identifier. Typically, in case of a router, this is the IPv6 address of the node loopback interface. Therefore, SR-IPv6 does not require any additional SID advertisement. The SID is in fact the

IPv6 address of the node.

4.2.2. Adjacency-SID

The Adjacency-SID identifies a given interface. In the SR architecture a node may advertise one or more Adj-SIDs allocated to a given interface so to force the forwarding of the packet (when received with that particular Adj-SID) into the interface, regardless the routing entry for the packet destination. The same is defined for SR-IPv6: a node may advertise a given IPv6 prefix which is associated to the SR semantic of "send out the packet to the interface this prefix is allocated to". Here also, the SID is in fact the IPv6 prefix.

5. SRH Procedures

In this section we describe the different procedures on the SRH.

5.1. Segment Routing Operations

When Segment Routing is instantiated over the IPv6 data plane the following applies:

- The segment list is encoded in the SRH.
- The active segment is in the destination address of the packet.
- The Segment Routing CONTINUE operation (as described in [I-D.filsfils-rtgwg-segment-routing]) is implemented as a regular/ plain IPv6 operation consisting of DA based forwarding.
- The NEXT operation is implemented through the update of the DA with the value represented by the Next Segment field in the SRH.
- The PUSH operation is implemented through the insertion of the SRH or the insertion of additional segments in the SRH segment list.

5.2. Segment Routing Node Functions

SR packets are forwarded to segments endpoints (i.e.: nodes whose address is in the DA field of the packet). The segment endpoint, when receiving a SR packet destined to itself, does:

- Inspect the SRH.
- Determine the next segment.

- Update the SRH (or, if requested, remove the SRH from the packet).
- Update the DA.
- Send the packet to the next segment.

The procedures applied to the SRH are related to the node function. Following nodes functions are defined:

Ingress SR Node.

Transit Non-SR Node.

Transit SR Intra Segment Node.

SR Endpoint Node.

5.2.1. Ingress SR Node

Ingress Node can be a router at the edge of the SR domain or a SR-capable host. The ingress SR node obtain the segment list by either:

Local path computation.

Interaction with an SDN controller delivering the path as a complete SRH.

When creating the SRH (either at ingress node or in the SDN controller) the following is done:

Next Header and Hdr Ext Len fields are set according to [RFC2460].

Routing Type field is set as TBD (SRH).

The DA of the packet is set with the address of the FIRST segment of the path.

Next Segment field contains the offset of the SECOND segment of the path which is encoded in the FIRST position of the segment list. The segment list is encoded as follows:

The first element of the list contains the second segment (as stated above).

All subsequent segments are encoded following the second segment.

The original DA of the packet is encoded as the last segment of the path (which is NOT the last segment of the segment list).

The last segment of the segment list is the FIRST segment of the path.

Last Segment field contains the offset of the last segment of the path (i.e.: the original DA of the packet).

The packet is sent out to the first segment.

5.2.1.1. Security at Ingress

The procedures related to the Segment Routing security are detailed in Section 6.

In the case where the SR domain boundaries are not under control of the network operator (e.g.: when the SR domain edge is in a home network), it is important to authenticate and validate the content of any SRH being received by the network operator. In such case, the security procedure described in Section 6 is to be used.

The ingress node (e.g.: the host in the home network) requests the SRH to a control system (e.g.: an SDN controller) which delivers the SRH with its HMAC signature on it.

Then, the home network host can send out SR packets (with an SRH on it) that will be validated at the ingress of the network operator infrastructure.

The ingress node of the network operator infrastructure, is configured in order to validate the incoming SRH HMACs in order to allow only packets having correct SRH according to their SA/DA addresses.

5.2.2. Transit Non-SR Capable Node

SR is interoperable with plain IPv6 forwarding. Any non SR-capable node will forward SR packets solely based on the DA. There's no SRH inspection. This ensures full interoperability between SR and non-SR nodes.

5.2.3. SR Intra Segment Transit Node

Only the node whose address is in DA inspects and processes the SRH (according to [RFC2460]). An intra segment transit node is not in the DA and its forwarding is based on DA and its SR-IPv6 FIB.

5.2.4. SR Segment Endpoint Node

The SR segment endpoint node is the node whose address is in the DA. The segment endpoint node inspects the SRH and does:

```
1.  IF DA = myself (segment endpoint)
2.     IF Next Segment <> Last Segment THEN
          update DA with Next Segment
          increment Next Segment
3.     ELSE IF Last Segment <> DA THEN
          update DA with Next Segment
          IF Clean-up bit is set THEN remove the SRH
4.     ELSE give the packet to next PID (application)
            End of processing.
5.  Forward the packet out
```

5.3. FRR Flag Settings

A node supporting SR and doing Fast Reroute (as described in [I-D.filsfils-rtgwg-segment-routing-use-cases], when rerouting packets through FRR mechanisms, SHOULD inspect the rerouted packet header and look for the SRH. If the SRH is present, the rerouting node SHOULD set the Protected bit on all rerouted packets.

6. SR-IPv6 Security

This section analyses the security threat model as well as the security issues and proposed solutions related to the new routing header for segment routing (a.k.a. segment routing header SRH).

The segment routing header is simply another version of the routing header as described in [RFC2460] and is:

- inserted when entering the segment routing domain which could be done by a node or by a router;

- read and acted upon when reaching the destination of the IP header.

Routers on the path that simply forward an IPv6 packet (i.e. the IPv6 destination address is not one of theirs) will never read and process the SRH. Routers whose one interface IPv6 address equals the destination address field of the SRH will have to parse the SRH and, if supported and if the local configuration allows it, will act on the SRH.

6.1. Threat model

Using a routing extension header which is basically source routing has some well-known security issues as described in [RFC4942] section 2.1.1 and [RFC5095]:

- amplification attacks: where a packet could be forged in such a way to cause looping among a set of intermediate routers causing unnecessary traffic, hence a denial of service against bandwidth;

- reflection attack: where a hacker could force an intermediate node to appear as the immediate attacker, hence hiding the real attacker from naive forensic;

- bypass attack: where an intermediate node could be used as a stepping stone (for example in a DMZ) to attack another host (for example in the datacenter or any back-end server.

These security issues did lead to obsoleting the routing header type 0 with [RFC5095] because:

- it was assumed to be acted upon by default by each and every router on the Internet;

- it contained multiple segments in the payload.

Therefore, if intermediate nodes ONLY act on valid and authorized SRH, then there is no security threat similar to RH-0. But as SR is used for added value services, there is also a need to prevent non-participating nodes to use those services; this is called 'service stealing prevention'.

6.2. Applicability of RFC 5095 to SRH

In the segment routing architecture described in [I-D.filsfils-rtgwg-segment-routing], there are basically two kinds of nodes (routers and hosts):

- nodes within the segment routing domain, which is within one single administrative domain, i.e., where all nodes are trusted anyway else the damage caused by those nodes could be worse than amplification attacks: traffic interception and man-in-the-middle attacks, more server DoS by dropping packets, and so on.

- Nodes outside of the segment routing domain, which is outside of the administrative segment routing domain hence they cannot be trusted because there is no physical security for those nodes, i.e., they can be replaced by hostile nodes or can be coerced in

wrong behaviors.

6.3. Security fields in SRH

The security-related fields in SRH are:

- HMAC Key-id, 8bits wide, if HMAC key-id is null, then there is no HMAC field;
- HMAC, 256 bits wide.

The HMAC field is the output of the hash of the concatenation of:

- the source IPv6 address;
- last segment, clean-up bit flag, HMAC key id, all addresses in the Segment List;
- a pre-shared secret between SR nodes in the SR domain (routers, controllers, ...).

The purpose of the HMAC field is to verify the validity and authorization of the SRH itself. If an outsider of the SR domain does not have access to the pre-shared secret, then it cannot compute the right HMAC field and the first SR router on the path processing the SRH and configure to check the validity of the HMAC will simply reject the packet.

The HMAC field is located at the end of the SRH simply because only the router on the ingress of the SR domain needs to process, then all other SR nodes can ignore it (based on local policy). This is to speed up forwarding operations because some hardware platforms can only parse in hardware so many bytes.

The HMAC Key-id field allows for the simultaneous existence of several hash algorithms (SHA-128, SHA-256, ... or future ones) as well as pre-shared keys. This allows for pre-shared key roll-over when two pre-shared keys are supported for a while when all SR nodes converged to a fresher pre-shared key. The HMAC key-id is opaque, i.e., it has no syntax except as an index to the right combination of pre-shared key and hash algorithm. It also allows for interoperation among different SR domains if allowed by local policy.

How HMAC key-id and pre-shared secret are synchronized between participating nodes in the SR domain is outside of the scope of this document ([RFC6407] GDOI could be a basis).

6.4. Nodes within the SR domain

Those nodes can be trusted to generate and to process SRH received on interfaces that are part of the SR domain (AS or set of ASs under common administration where SR is enabled). These nodes MUST drop all packets received on an interface that is not part of the SR domain and containing a SRH whose HMAC field cannot be validated by local policies. This includes obviously packet with a SRH generated by a non-cooperative SR domain.

If the validation fails, then these packets MUST be dropped, ICMP error messages (parameter problem) SHOULD be generated (but rate limited) and SHOULD be logged.

6.5. Nodes outside of the SR domain

Nodes outside of the SR domain cannot be trusted for physical security; hence, they need to request by some means (outside of the scope of this document) a complete SRH for each new connection. The SRH MUST include a HMAC key-id and HMAC field which is computed correctly (see Section 6.3).

When an outside node sends a packet with an SRH and towards a SR ingress node, the packet MUST contain the HMAC field and the SR ingress node MUST be in the segment list of the SRH.

The ingress SR router, i.e., the router with an interface address equals to the destination address, MUST verify the HMAC field.

If the validation is successful, then the packet is simply forwarded as usual for a SR packet. As long as the packet travels within the SR domain, no further HMAC check is done. Subsequent routers in the SR domain MAY verify the HMAC field.

If the validation fails, then this packet MUST be dropped, an ICMP error message (parameter problem) SHOULD be generated (but rate limited) and SHOULD be logged.

7. SR and Tunneling

Encapsulation can be realized in two different ways with SR-IPv6:

- Outer encapsulation.

SRH with SA/DA original addresses.

Outer encapsulation tunneling is the traditional method where an additional IPv6 header is prepended to the packet. The original IPv6 header being encapsulated, everything is preserved and the packet is switched/routed according to the outer header (that could contain a SRH).

SRH allows encoding both original SA and DA and therefore, hence an operator may decide to change the SA/DA at ingress and restore them at egress. This can be achieved without outer encapsulation, by changing SA/DA and encoding the original values in the Segment List (the last segment of the path being the original DA) and in the Policy List (original SA).

8. Example Use Case

A more detailed description of use cases are available in [draft-martin-spring-segment-routing-ipv6-use-cases-00]. In this section, a simple SR-IPv6 example is illustrated.

In the topology described in Figure 2 it is assumed an end-to-end SR deployment. Therefore SR is supported by all nodes from A to J.

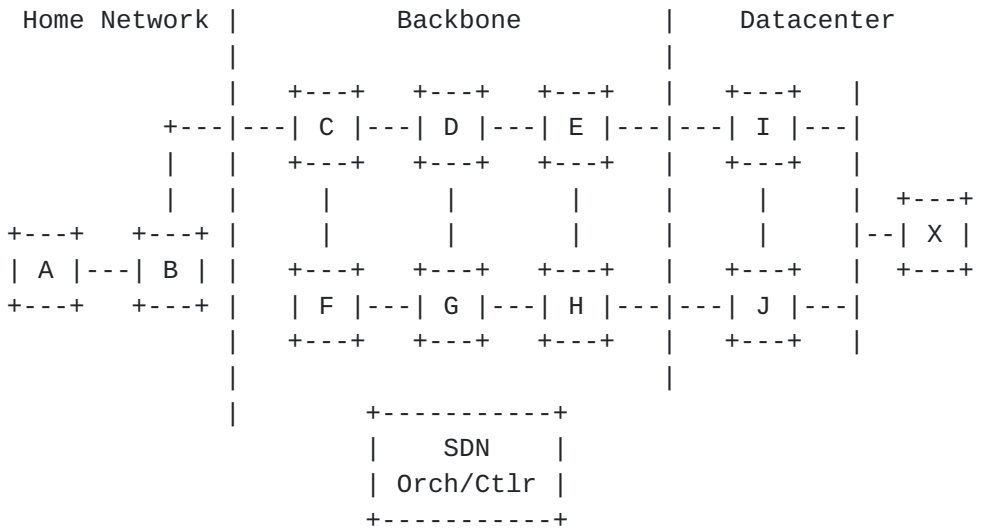

Figure 2: Sample SR topology

The following workflow applies to packets sent by host A and destined to server X.

- Host A sends a request for a path to server X to the SDN controller or orchestration system.

- The SDN controller/orchestrator builds a SRH with:
  - Segment List: C, F, J, X
  - HMAC

  that satisfies the requirements expressed in the request by host A and based on policies applicable to host A.

- Host A receives the SRH and insert it into the packet. The packet has now:
  - SA: A
  - DA: C
  - SRH with
    - SL: F,J,X,C
    - PL: C (ingress), J (egress)

    Note that X is the last segment and C is the first segment (encoded at the end of the SL).

- When packet arrives in C (first segment), C does:
  - Validate the HMAC of the SRH.

  - Update the DA with the next segment (found in SRH): DA is set to F.
  - Forward the packet to F.

- Packet arrives in F which inspects the SRH and find the next segment:
  - DA is set to J.

- Packet travels across G and H nodes which do plain IPv6 forwarding based on DA. No inspection of SRH needs to be done in these nodes. However, any SR capable node is allowed to set the Protected bit in case of FRR protection.

- Packet arrives in J where two options are available depending on the settings of the cleanup bit set in the SRH:
  - If the cleanup bit is set, then node J will strip out the SRH from the packet, set the DA as X and send the packet out.
  - If the clean-up bit is not set, the DA is set to X and the packet is sent out with the SRH.

The packet arrives in the server that may or may not support SR. The return traffic, from server to host, may be sent using the same procedures.

9. IANA Considerations

TBD

10. Manageability Considerations

TBD

11. Security Considerations

Security mechanisms applied to Segment Routing over IPv6 networks are detailed in Section 6.

12. Contributors

Eric Vynke contributed to this document through the writings of Section 6.

13. Acknowledgements

The authors would like to thank John Leddy, John Brzozowski, Mark Townsley, Christian Martin, Roberta Maglione and James Connolly for their contribution to this document.

14. References

14.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

[RFC5095] Abley, J., Savola, P., and G. Neville-Neil, "Deprecation of Type 0 Routing Headers in IPv6", RFC 5095, December 2007.

[RFC6407] Weis, B., Rowles, S., and T. Hardjono, "The Group Domain of Interpretation", RFC 6407, October 2011.

14.2. Informative References

[I-D.filsfils-rtgwg-segment-routing]
Filsfils, C., Previdi, S., Bashandy, A., Decraene, B., Litkowski, S., Horneffer, M., Milojevic, I., Shakir, R., Ytti, S., Henderickx, W., Tantsura, J., and E. Crabbe, "Segment Routing Architecture", draft-filsfils-rtgwg-segment-routing-01 (work in progress), October 2013.

[I-D.filsfils-rtgwg-segment-routing-use-cases]
Filsfils, C., Francois, P., Previdi, S., Decraene, B., Litkowski, S., Horneffer, M., Milojevic, I., Shakir, R., Ytti, S., Henderickx, W., Tantsura, J., Kini, S., and E. Crabbe, "Segment Routing Use Cases", draft-filsfils-rtgwg-segment-routing-use-cases-02 (work in progress), October 2013.

[I-D.filsfils-spring-segment-routing-mpls]
Filsfils, C., Previdi, S., Bashandy, A., Decraene, B., Litkowski, S., Horneffer, M., Milojevic, I., Shakir, R., Ytti, S., Henderickx, W., Tantsura, J., and E. Crabbe, "Segment Routing with MPLS data plane", draft-filsfils-spring-segment-routing-mpls-00 (work in progress), October 2013.

[I-D.previdi-isis-segment-routing-extensions]
Previdi, S., Filsfils, C., Bashandy, A., Gredler, H., Litkowski, S., and J. Tantsura, "IS-IS Extensions for Segment Routing", draft-previdi-isis-segment-routing-extensions-05 (work in progress), February 2014.

[I-D.psenak-ospf-segment-routing-ospfv3-extension]
Psenak, P., Previdi, S., Filsfils, C., Gredler, H., Shakir, R., and W. Henderickx, "OSPFv3 Extensions for Segment Routing", draft-psenak-ospf-segment-routing-ospfv3-extension-01 (work in progress), February 2014.

[RFC1940] Estrin, D., Li, T., Rekhter, Y., Varadhan, K., and D. Zappala, "Source Demand Routing: Packet Format and Forwarding Specification (Version 1)", RFC 1940, May 1996.

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

[draft-martin-spring-segment-routing-ipv6-use-cases-00]
Brzozowski, J., Leddy, J., Leung, I., Previdi, S.,
Townsley, M., Martin, C., Filsfils, C., and R. Maglione,
"IPv6 Segment Routing Use Cases", March 2014.

Authors' Addresses

Stefano Previdi (editor)
Cisco Systems, Inc.
Via Del Serafico, 200
Rome 00142
Italy

Email: sprevidi@cisco.com

Clarence Filsfils
Cisco Systems, Inc.
Brussels,
BE

Email: cfilsfil@cisco.com

Brian Field
Comcast
4100 East Dry Creek Road
Centennial, CO 80122
US

Email: Brian_Field@cable.comcast.com

Ida Leung
Rogers Communications
8200 Dixie Road
Brampton, ON L6T 0C1
CA

Email: Ida.Leung@rci.rogers.com