Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 4, 2015

E. Vyncke
S. Previdi
Cisco Systems, Inc.
B. Field
Comcast
I. Leung
Rogers Communications

July 3, 2014

IPv6 Segment Routing Header (SRH) Security Considerations
draft-vyncke-6man-segment-routing-security-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.

Segment Routing can be applied to the IPv6 data plane with the addition of a new type of Routing Extension Header. This draft analyses the security aspects the Segment Routing Extension Header Type and how it is used by SR capable nodes to deliver a secure service.

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 will expire on January 4, 2015.

Copyright Notice

Table of Contents

1. Segment Routing Documents

Segment Routing terminology is defined in [I-D.filsfils-spring-segment-routing].

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

Segment Routing IPv6 use cases are described in [I-D.ietf-spring-ipv6-use-cases].

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

2. Introduction

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.

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

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

- o inspected and acted upon when reaching the destination address of the IP header.

Routers on the path that simply forward an IPv6 packet (i.e. the IPv6 destination address is none of theirs) will never inspect 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.

3. Threat model

3.1. Source routing threat

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

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

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

- o 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, RH-0, with [RFC5095] because:

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

- o 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.

3.2. Applicability of RFC 5095 to SRH

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

- o 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.

- o 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.

3.3. Service stealing threat

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'.

3.4. Topology disclosure

The SRH also contains all IPv6 addresses of intermediate SR-nodes, this obviously reveals those addresses to the potentially hostile attackers if those attackers are on the path.

4. Security fields in SRH

This section summarizes the use of specific fields in the SRH; they are integral part of [I-D.previdi-6man-segment-routing-header] and they are again described here for reader's sake.

The security-related fields in SRH are:

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

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

- o the source IPv6 address;
- o last segment field, an octet whose bit-0 is the clean-up bit flag and others are 0, HMAC key-id, all addresses in the Segment List;
- o a pre-shared secret between SR nodes in the SR domain (routers, controllers, ...);
- o if required by the hash algorithm a pad field filled with 0.

The purpose of the HMAC field is to verify the validity, the integrity and the authorization of the SRH itself. If an outsider of the SR domain does not have access to a current pre-shared secret, then it cannot compute the right HMAC field and the first SR router on the path processing the SRH and configured 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 it, then all other SR nodes can ignore it (based on local policy) because they can trust the upstream router. 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-256, SHA3-256 ... 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.

When a specific SRH is linked to a time-related service (such as turbo-QoS for a 1-hour period) where the DA, SID are identical, then it is important to refresh the shared-secret frequently as the HMAC validity period expires only when the HMAC key-id and its associated shared-secret expires. 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).

4.1. Selecting a hash algorithm

The HMAC field in the SRH is 256 bit wide. Therefore, the HMAC MUST be based on a hash function whose output is at least 256 bits. If the output of the hash function is 256, then this output is simply inserted in the HMAC field. If the output of the hash function is larger than 256 bits, then the output value is truncated to 256 by taking the least-significant 256 bits and inserting them in the HMAC field.

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

4.2. Performance impact of HMAC

While adding a HMAC to each and every SR packet increases the security, it has a performance impact. Nevertheless, it must be noted that:

- the HMAC field is used only when SRH is inserted by a device (such as a home set-up box) which is outside of the segment routing domain. If the SRH is added by a router in the trusted segment routing domain, then, there is no need for a HMAC field, hence no performance impact.

- when present, the HMAC field MUST only be checked and validated by the first router of the segment routing domain, this router is named 'validating router'. Downstream routers SHOULD NOT inspect the HMAC field.

- this validating router can also have a cache of <IPv6 header + SRH, HMAC field value> to improve the performance. It is not the same use case as in IPsec where HMAC value was unique per packet, in SRH, the HMAC value is unique per flow.

- Last point, hash functions such as SHA-2 have been optimized for security and performance and there are multiple implementations with good performance.

With the above points in mind, the performance impact of using HMAC is minimized.

4.3. Pre-shared key management

The field HMAC key-id allows for:

- o key roll-over: when there is a need to change the key (the hash pre-shared secret), then multiple pre-shared keys can be used simultaneously. The validating routing can have a table of <key-id, pre-shared secret> for the current and future keys.

- o different algorithm: by extending the previous table to <key-id, hash function, pre-shared secret>, the validating router can also support simultaneously several hash algorithm (see section Section 4.1)

The pre-shared secret distribution can be done:

- o in the configuration of the validating routers, either by static configuration or any SDN oriented approach;

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

NOTE: this section needs more work but the intent is NOT to define yet-another-key-distribution-protocol.

5. Deployment Models

5.1. Nodes within the SR domain

Those nodes can be trusted to generate SRH and to process SRH received on interfaces that are part of the SR domain. 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.

5.2. 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 (i.e.

new destination address). The SRH MUST include a HMAC key-id and HMAC field which is computed correctly (see Section 4).

When an outside node sends a packet with an SRH and towards a SR ingress node, the packet MUST contain the HMAC key-id and HMAC field and the SR ingress node MUST be the destination address.

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

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 needs to be done. Subsequent routers in the SR domain MAY verify the HMAC field when they process the SRH (i.e. when they are the destination).

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.

5.3. SR path exposure

As the intermediate SR nodes addresses appears in the SRH, if this SRH is visible to an outside then he/she could reuse this knowledge to launch an attack on the intermediate SR nodes or get some insider knowledge on the topology. This is especially applicable when the path between the source node and the first SR-node in the domain is on the public Internet.

The first remark is to state that 'security by obscurity' is never enough; in other words, the security policy of the SR domain MUST assume that the internal topology and addressing is known by the attacker. A simple traceroute will also give the same information (with even more information as all intermediate nodes between SID will also be exposed). IPsec Encapsulating Security Payload (RFC 4303) cannot be use to protect the SRH as per RFC 4303 the ESP header must appear after any routing header (including SRH).

To prevent a user to leverage the gained knowledge by intercepting SRH, it it recommended to apply an infrastructure Access Control List (iACL) at the edge of the SR domain. This iACL will drop all packets from outside the SR-domain whose destination is any address of any router inside the domain. This security policy should be tuned for local operations.

6. IANA Considerations

There are no IANA request or impact in this document.

7. Manageability Considerations

TBD

8. Security Considerations

This document describes the security mechanisms applied to the Segment Routing Header defined in [I-D.previdi-6man-segment-routing-header]

9. Acknowledgements

The authors would like to thank Dave Barach for his contribution to this document.

10. References

10.1. 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, 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.

10.2. Informative References

[I-D.filsfils-spring-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-spring-segment-routing-03 (work in progress), June 2014.

[I-D.filsfils-spring-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-spring-segment-routing-use-cases-00 (work in progress), March 2014.

[I-D.ietf-isis-segment-routing-extensions]
Previdi, S., Filsfils, C., Bashandy, A., Gredler, H., Litkowski, S., Decraene, B., and J. Tantsura, "IS-IS Extensions for Segment Routing", draft-ietf-isis-segment-routing-extensions-02 (work in progress), June 2014.

[I-D.ietf-spring-ipv6-use-cases]
Brzozowski, J., Leddy, J., Leung, I., Previdi, S., Townsley, W., Martin, C., Filsfils, C., and R. Maglione, "IPv6 SPRING Use Cases", draft-ietf-spring-ipv6-use-cases-00 (work in progress), May 2014.

[I-D.previdi-6man-segment-routing-header]
Previdi, S., Filsfils, C., Field, B., and I. Leung, "IPv6 Segment Routing Header (SRH)", draft-previdi-6man-segment-routing-header-01 (work in progress), June 2014.

[I-D.psenak-ospf-segment-routing-ospfv3-extension]
Psenak, P., Previdi, S., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPFv3 Extensions for Segment Routing", draft-psenak-ospf-segment-routing-ospfv3-extension-02 (work in progress), July 2014.

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

Authors' Addresses

Eric Vyncke
Cisco Systems, Inc.
De Kleetlaann 6A
Diegem 1831
Belgium

Email: evyncke@cisco.com

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

Email: sprevidi@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