Source Packet Routing in Networking
Internet-Draft
Intended status: Standards Track
Expires: 30 August 2025

N. Buraglio
Energy Sciences Network
T. Mizrahi
Huawei
T. Tong
China Unicom
L. M. Contreras
Telefonica
F. Gont
SI6 Networks
26 February 2025

Segment Routing IPv6 Security Considerations
draft-ietf-spring-srv6-security-02

Abstract

SRv6 is a traffic engineering, encapsulation and steering mechanism utilizing IPv6 addresses to identify segments in a pre-defined policy. This document discusses security considerations in SRv6 networks, including the potential threats and the possible mitigation methods. The document does not define any new security protocols or extensions to existing protocols.

About This Document

This note is to be removed before publishing as an RFC.

The latest revision of this draft can be found at https://buraglio.github.io/draft-bdmgct-spring-srv6-security/draft-bdmgct-spring-srv6-security.html. Status information for this document may be found at https://datatracker.ietf.org/doc/draft-ietf-spring-srv6-security/.

Discussion of this document takes place on the Source Packet Routing in Networking Working Group mailing list (mailto:spring@ietf.org), which is archived at https://mailarchive.ietf.org/arch/browse/spring/. Subscribe at https://www.ietf.org/mailman/listinfo/spring/.

Source for this draft and an issue tracker can be found at https://github.com/buraglio/draft-bdmgct-spring-srv6-security.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Copyright Notice

Table of Contents

1. Introduction

Segment Routing (SR) [RFC8402] utilizing an IPv6 data plane is a source routing model that leverages an IPv6 underlay and an IPv6 extension header called the Segment Routing Header (SRH) [RFC8754] to signal and control the forwarding and path of packets by imposing an ordered list of segments that are processed at each hop along the signaled path. SRv6 is fundamentally bound to the IPv6 protocol and introduces a new extension header. There are security considerations which must be noted or addressed in order to operate an SRv6 network in a reliable and secure manner. Specifically, some primary properties of SRv6 that affect the security considerations are:

* SRv6 may use the SRH which is a type of Routing Extension Header defined by [RFC8754]. Security considerations of the SRH are discussed [RFC8754] section 7, and were based in part on security considerations of the deprecated routing header 0 as discussed in [RFC5095] section 5.

* SRv6 uses the IPv6 data-plane, and therefore security considerations of IPv6 are applicable to SRv6 as well. Some of these considerations are discussed in Section 10 of [RFC8200] and in [RFC9099].

* While SRv6 uses what appear to be typical IPv6 addresses, the address space is processed differently by segment endpoints. A typical IPv6 unicast address is comprised of a network prefix, host identifier. A typical SRv6 segment identifier (SID) is comprised of a locator, a function identifier, and optionally, function arguments. The locator must be routable, which enables both SRv6 capable and incapable devices to participate in forwarding, either as normal IPv6 unicast or SRv6 segment endpoints. The capability to operate in environments that may have gaps in SRv6 support allows the bridging of islands of SRv6 devices with standard IPv6 unicast routing.

This document describes various threats to SRv6 networks and also presents existing approaches to avoid or mitigate the threats.

2. Scope of this Document

The following IETF RFCs were selected for security assessment as part of this effort:

* [RFC8402] : "Segment Routing Architecture"

* [RFC8754] : "IPv6 Segment Routing Header (SRH)"

* [RFC8986] : "Segment Routing over IPv6 (SRv6) Network Programming"

* [RFC9020] : "YANG Data Model for Segment Routing"

* [RFC9256] : "Segment Routing Policy Architecture"

* [RFC9491] : "Integration of the Network Service Header (NSH) and Segment Routing for Service Function Chaining (SFC)"

* [RFC9524] : "Segment Routing Replication for Multipoint Service Delivery"

We note that SRv6 is under active development and, as such, the above documents might not cover all protocols employed in an SRv6 deployment.

3. Conventions and Definitions

3.1. 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.

3.2. Terminology

* HMAC TLV: Hashed Message Authentication Code Type Length Value [RFC8754]

* SID: Segment Identifier [RFC8402]

* SRH: Segment Routing Header [RFC8754]

* SRv6: Segment Routing over IPv6 [RFC8402]

4. Threat Model

This section introduces the threat model that is used in this document. The model is based on terminology from the Internet threat model [RFC3552], as well as some concepts from [RFC9055], [RFC7384], [RFC7835] and [RFC9416]. Details regarding inter-domain segment routing (SR) are out of scope for this document.

Internal vs. External: An internal attacker in the context of SRv6 is an attacker who is located within an SR domain. Specifically, an internal attacker either has access to a node in the SR domain, or is located on an internal path between two nodes in the SR domain. External attackers, on the other hand, are not within the SR domain.

On-path vs. Off-path: On-path attackers are located in a position that allows interception, modification or dropping of in-flight packets, as well as insertion (generation) of packets. Off-path attackers can only attack by insertion of packets.

Data plane vs. control plane vs. Management plane: Attacks can be classified based on the plane they target: data, control, or management. The distinction between on-path and off-path attackers depends on the plane where the attack occurs. For instance, an attacker might be off-path from a data plane perspective but on-path from a management plane perspective.

The following figure depicts an example of an SR domain with six attacker types, labeled 1-6. For instance, attacker 2 is located along the path between the SR ingress node and SR endpoint 1, and is therefore an on-path attacker both in the data plane and in the control plane. Thus, attacker 2 can listen, insert, delete, modify or replay data plane and/or control plane packets in transit. Off-path attackers, such as attackers 4 and 5, can insert packets, and in some cases can passively listen to some traffic, such as multicast transmissions. Attacker 3 is internal an on-path attacker in the management plane, as it is located along the path between the Network Management System (NMS) and SR endpoint 1.

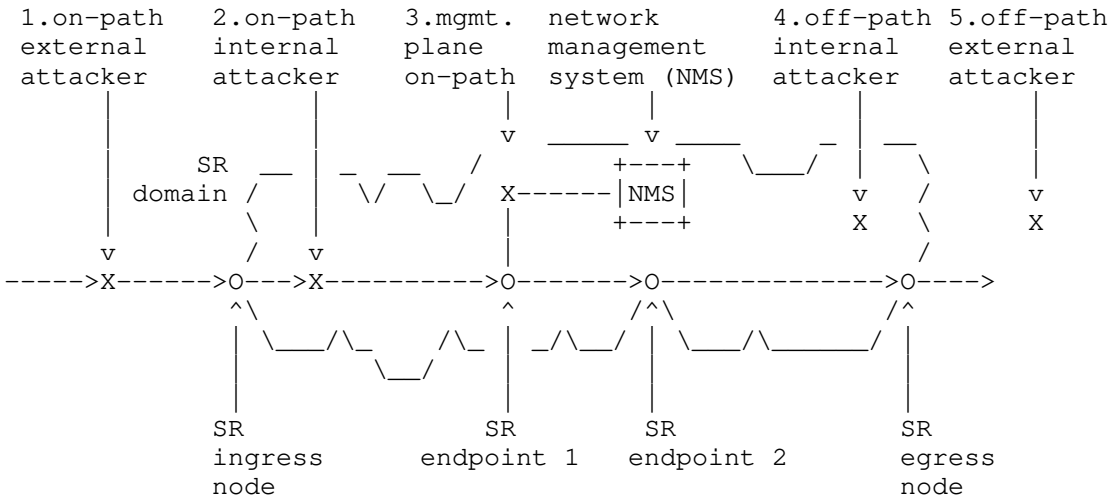

Figure 1: Threat Model Taxonomy

As defined in [RFC8402], SR operates within a "trusted domain". Therefore, in the current threat model the SR domain defines the boundary that distinguishes internal from external threats. Specifically, an attack on one domain that is invoked from within a different domain is considered an external attack in the context of the current document.

5. Impact

One of the important aspects of a threat analysis is the potential impact of each threat. SRv6 allows for the forwarding of IPv6 packets via predetermined SR policies, which determine the paths and the processing of these packets. An attack on SRv6 may cause packets to traverse arbitrary paths and to be subject to arbitrary processing by SR endpoints within an SR domain. This may allow an attacker to perform a number of attacks on the victim networks and hosts that

would be mostly unfeasible for a non-SRv6 environment.

The threat model in [ANSI-Sec] classifies threats according to their potential impact, defining six categories. For each of these categories we briefly discuss its applicability to SRv6 attacks.

* Unauthorized Access: an attack that results in unauthorized access might be achieved by having an attacker leverage SRv6 to circumvent security controls as a result of security devices being unable to enforce security policies. For example, this can occur if packets are directed through paths where packet filtering policies are not enforced, or if some security policies are not enforced in the presence of IPv6 Extension Headers.

* Masquerade: various attacks that result in spoofing or masquerading are possible in IPv6 networks. However, these attacks are not specific to SRv6, and are therefore not within the scope of this document.

* System Integrity: attacks on SRv6 can manipulate the path and the processing that the packet is subject to, thus compromising the integrity of the system. Furthermore, an attack that compromises the control plane and/or the management plane is also a means of impacting the system integrity.

* Communication Integrity: SRv6 attacks may cause packets to be forwarded through paths that the attacker controls, which may facilitate other attacks that compromise the integrity of user data. Integrity protection of user data, which is implemented in higher layers, avoids these aspects, and therefore communication integrity is not within the scope of this document.

* Confidentiality: as in communication integrity, packets forwarded through unintended paths may traverse nodes controlled by the attacker. Since eavesdropping to user data can be avoided by using encryption in higher layers, it is not within the scope of this document. However, eavesdropping to a network that uses SRv6 allows the attacker to collect information about SR endpoint addresses, SR policies, and network topologies, is a specific form of reconnaissance

* Denial of Service: the availability aspects of SRv6 include the ability of attackers to leverage SRv6 as a means for compromising the performance of a network or for causing Denial of Service (DoS). Compromising the availability of the system can be achieved by sending multiple SRv6-enabled packets to/through victim nodes, where the SRv6-enabled packets result in a negative performance impact of the victim systems (see [RFC9098] for

further details). Alternatively, an attacker might achieve attack amplification by causing packets to "bounce" multiple times between a set of victim nodes, with the goal of exhausting processing resources and/or bandwidth (see [CanSecWest2007] for a discussion of this type of attack).

Section 6 discusses specific implementations of these attacks, and possible mitigations are discussed in Section 7.

6. Attacks

6.1. Attack Abstractions

Packet manipulation and processing attacks can be implemented by performing a set of one or more basic operations. These basic operations (abstractions) are as follows:

* Passive listening: an attacker who reads packets off the network can collect information about SR endpoint addresses, SR policies and the network topology. This information can then be used to deploy other types of attacks.

* Packet replaying: in a replay attack the attacker records one or more packets and transmits them at a later point in time.

* Packet insertion: an attacker generates and injects a packet to the network. The generated packet may be maliciously crafted to include false information, including for example false addresses and SRv6-related information.

* Packet deletion: by intercepting and removing packets from the network, an attacker prevents these packets from reaching their destination. Selective removal of packets may, in some cases, cause more severe damage than random packet loss.

* Packet modification: the attacker modifies packets during transit.

This section describes attacks that are based on packet manipulation and processing, as well as attacks performed by other means. While it is possible for packet manipulation and processing attacks against all the fields of the IPv6 header and its extension headers, this document limits itself to the IPv6 header and the SRH.

6.2. Modification Attack

6.2.1. Overview

An on-path internal attacker can modify a packet while it is in transit in a way that directly affects the packet's segment list and other extension headers.

Header Modification	Impact
source address	spoof another source
destination address	modify the segment list active segment and arguments, including next segments like binding SIDs Binding SID [RFC8402] or compressed segments [I-D.ietf-spring-srv6-srh-compression]
SRH	insert or delete the SRH modifying the previous and next segments in the segment list
SRH segment list	insert, delete or modify the previous or next segments in the segment list
SRH TLV	insert, delete or modify TLVs in the SRH
SRH flags and tag	modify tags and flags

Table 1

An on-path internal attacker can also modify, insert or delete other extension headers but these are outside the scope of this document.

6.2.2. Scope

An SR modification attack can be performed by on-path attackers. If filtering is deployed at the domain boundaries as described in Section 7.1, the ability to implement SR modification attacks is limited to on-path internal attackers.

6.2.3. Impact

The SR modification attack allows an on-path internal attacker to change the segment list in the packet, i.e. the SR policy that the packet is steered through, and thus to manipulate the path and the processing that the packet is subject to.

Specifically, the SR modification attack can impact the network and the forwarding behavior of packets in one or more of the following ways:

Avoiding a specific node or path: An attacker can manipulate the DA and/or SRH in order to avoid a specific node or path. This approach can be used, for example, for bypassing the billing service or avoiding access controls and security filters.

Preferring a specific path: The packet can be manipulated to divert packets to a specific path. This attack can result in allowing various unauthorized services such as traffic acceleration. Alternatively, an attacker can divert traffic to be forwarded through a specific node that the attacker has access to, thus facilitating more complex on-path attacks such as passive listening, recon and various man-in-the-middle attacks. It is noted that the SR modification attack is performed by an on-path attacker who has access to packets in transit, and thus can implement these attacks directly. However, SR modification is relatively easy to implement and requires low processing resources by an attacker, while it facilitates more complex on-path attacks by averting the traffic to another node that the attacker has access to and has more processing resources.

Forwarding through a path that causes the packet to be discarded: SR modification may cause a packet to be forwarded to a point in the network where it can no longer be forwarded, causing the packet to be discarded.

Manipulating the SRv6 network programming: An attacker can trigger a specific endpoint behavior by modifying the destination address and/or SIDs in the segment list. This attack can be invoked in order to manipulate the path or in order to exhaust the resources of the SR endpoint.

Availability: An attacker can add SIDs to the segment list in order to increase the number hops that each packet is forwarded through and thus increase the load on the network. For example, a set of SIDs can be inserted in a way that creates a forwarding loop ([RFC8402], [RFC5095]) and thus loads the nodes along the loop. Network programming can be used in some cases to manipulate segment endpoints to perform unnecessary functions that consume processing resources. TLV fields such as the HMAC TLV can be maliciously added to the SRH in order to consume processing resources. Path inflation, malicious looping and unnecessary instructions and TLVs have a common outcome, resource exhaustion, which may in severe cases cause Denial of Service (DoS).

6.3. Passive Listening

6.3.1. Overview

An on-path internal attacker can passively listen to packets and specifically listen to the SRv6-related information that is conveyed in the IPv6 header and the SRH. This approach can be used for reconnaissance, i.e., for collecting segment lists.

6.3.2. Scope

A reconnaisance attack is limited to on-path internal attackers.

If filtering is deployed at the domain boundaries (Section 7.1), it prevents any leaks of explicit SRv6 routing information through the boundaries of the administrative domain. In this case external attackers can only collect SRv6-related data in a malfunctioning network in which SRv6-related information is leaked through the boundaries of an SR domain.

6.3.3. Impact

While the information collected in a reconnaisance attack does not compromise the confidentiality of the user data, it allows an attacker to gather information about the network which in turn can be used to enable other attacks.

6.4. Packet Insertion

6.4.1. Overview

In a packet insertion attack packets are inserted (injected) into the network with a segment list. The attack can be applied either by using synthetic packets or by replaying previously recorded packets.

6.4.2. Scope

Packet insertion can be performed by either on-path or off-path attackers. In the case of a replay attack, recording packets in-flight requires on-path access and the recorded packets can later be injected either from an on-path or an off-path location.

If filtering is deployed at the domain boundaries (Section 7.1), insertion attacks can only be implemented by internal attackers.

6.4.3. Impact

The main impact of this attack is resource exhaustion which compromises the availability of the network, as described in Section 6.2.3.

6.5. Control and Management Plane Attacks

6.5.1. Overview

Depending on the control plane protocols used in a network, it is possible to use the control plane as a way of compromising the network. For example, an attacker can advertise SIDs in order to manipulate the SR policies used in the network. Known IPv6 control plane attacks (e.g., overclaiming) are applicable to SRv6 as well.

A compromised management plane can also facilitate a wide range of attacks, including manipulating the SR policies or compromising the network availability.

6.5.2. Scope

The control plane and management plane may be either in-band or out-of-band, and thus the on-path and off-path taxonomy of Section 4 is not necessarily common between the data plane, control plane and management plane. As in the data plane, on-path attackers can be implement a wide range of attacks in order to compromise the control and/or management plane, including selectively removing legitimate messages, replaying them or passively listening to them. However, while an on-path attacker in the data plane is potentially more harmful than an off-path attacker, effective control and/or management plane attacks can be implemented off-path rather than by trying to intercept or modify traffic in-flight, for example by exchanging malicious control plane messages with legitimate routers, by spoofing an SDN (Software Defined Network) controller, or by gaining access to an NMS (Network Management System).

SRv6 domain boundary filtering can be used for mitigating potential control plane and management plane attacks from external attackers. Segment routing does not define any specific security mechanisms in existing control plane or management plane protocols. However, existing control plane and management plane protocols use authentication and security mechanisms to validate the authenticity of information.

6.5.3. Impact

A compromised control plane or management plane can impact the network in various possible ways. SR policies can be manipulated by the attacker to avoid specific paths or to prefer specific paths, as described in Section 6.2.3. Alternatively, the attacker can compromise the availability, either by defining SR policies that load the network resources, as described in Section 6.2.3, or by blackholing some or all of the SR policies. A passive attacker can use the control plane or management plane messages as a means for recon, similarly to Section 6.2.3.

6.6. Other Attacks

Various attacks which are not specific to SRv6 can be used to compromise networks that deploy SRv6. For example, spoofing is not specific to SRv6, but can be used in a network that uses SRv6. Such attacks are outside the scope of this document.

Because SRv6 is completely reliant on IPv6 for addressing, forwarding, and fundamental networking basics, it is potentially subject to any existing or emerging IPv6 vulnerabilities [RFC9099], however, this is out of scope for this document.

7. Mitigation Methods

This section presents methods that can be used to mitigate the threats and issues that were presented in previous sections. This section does not introduce new security solutions or protocols.

7.1. Trusted Domains and Filtering

7.1.1. Overview

As specified in [RFC8402]:

> By default, SR operates within a trusted domain. Traffic MUST be filtered at the domain boundaries.
> The use of best practices to reduce the risk of tampering within the trusted domain is important. Such practices are discussed in [RFC4381] and are applicable to both SR-MPLS and SRv6.

Following the spirit of [RFC8402], the current document assumes that SRv6 is deployed within a trusted domain. Traffic MUST be filtered at the domain boundaries. Thus, most of the attacks described in this document are limited to within the domain (i.e., internal attackers).

7.1.2. SRH Filtering

Filtering on presence of an SRH is possible but not useful for two reasons: 1. The SRH is optional for SID processing as described in [RFC8754] section 3.1 and 4.1. 2. A packet containing an SRH may not be destined to the SR domain, it may be simply transiting the domain.

For these reasons SRH filtering is not a useful method of mitigation, and thus filtering can only be applied based on the address range, as described below.

7.1.3. Address Range Filtering

The IPv6 destination address can be filtered at the SR ingress node and at all nodes implementing SRv6 SIDs within the SR domain in order to mitigate external attacks. Section 5.1 of [RFC8754] describes this in detail, it's summarized here as: 1. At ingress nodes, any packet entering the SR domain and destined to a SID within the SR domain is dropped. 2. At every SRv6 enabled node, any packet destined to a SID instantiated at the node from a source address outside the SR domain is dropped.

In order to apply such a filtering mechanism the SR domain needs to have an infrastructure address range for SIDs, and an infrastructure address range for source addresses, that can be detected and enforced. Some examples of an infrastructure address range for SIDs are: 1. ULA addresses 2. The prefix defined in [RFC9602]. 3. GUA addresses

Many operators reserve a /64 block for all loopback addresses and allocate /128 for each loopback interface. This simplifies the filtering of permitted source addresses.

Failure to implement address range filtering at ingress nodes is mitigated with filtering at SRv6 enabled node. Failure to implement both filtering mechanisms could result in a "fail open" scenario, where some attacks by internal attackers described in this document may be launched by external attackers.

7.2. Encapsulation of Packets

Packets steered in an SR domain are often encapsulated in an IPv6 encapsulation. This mechanism allows for encapsulation of both IPv4 and IPv6 packets. Encapsulation of packets at the SR ingress node and decapsulation at the SR egress node mitigates the ability of external attackers to attack the domain.

7.3. Hashed Message Authentication Code (HMAC)

The SRH can be secured by an HMAC TLV, as defined in [RFC8754]. The HMAC is an optional TLV that secures the segment list, the SRH flags, the SRH Last Entry field and the IPv6 source address. A pre-shared key is used in the generation and verification of the HMAC.

Using an HMAC in an SR domain can mitigate some of the SR Modification Attacks (Section 6.2). For example, the segment list is protected by the HMAC.

The following aspects of the HMAC should be considered:

* The HMAC TLV is OPTIONAL.

* While it is presumed that unique keys will be employed by each participating node, in scenarios where the network resorts to manual configuration of pre-shared keys, the same key might be reused by multiple systems as an (incorrect) shortcut to keeping the problem of pre-shared key configuration manageable.

* When the HMAC is used there is a distinction between an attacker who becomes internal by having physical access, for example by plugging into an active port of a network device, and an attacker who has full access to a legitimate network node, including for example encryption keys if the network is encrypted. The latter type of attacker is an internal attacker who can perform any of the attacks that were described in the previous section as relevant to internal attackers.

* An internal attacker who does not have access to the pre-shared key can capture legitimate packets, and later replay the SRH and HMAC from these recorded packets. This allows the attacker to insert the previously recorded SRH and HMAC into a newly injected packet. An on-path internal attacker can also replace the SRH of an in-transit packet with a different SRH that was previously captured.

8. Implications on Existing Equipment

8.1. Limitations in Filtering Capabilities

[RFC9288] provides recommendations on the filtering of IPv6 packets containing IPv6 extension headers at transit routers. However, this class of filtering is shown to not be useful and can be ignored.

Filtering on prefixes has been shown to be useful, specifically [RFC8754]'s description of packet filtering. There are no known limitations with filtering on infrastructure addresses, and [RFC9099] expands on the concept with control plane filtering.

8.2. Middlebox Filtering Issues

When an SRv6 packet is forwarded in the SRv6 domain, its destination address changes constantly, the real destination address is hidden. Security devices on SRv6 network may not learn the real destination address and fail to take access control on some SRv6 traffic.

The security devices on SRv6 networks need to take care of SRv6 packets. However, the SRv6 packets usually use loopback address of the PE device as a source address. As a result, the address information of SR packets may be asymmetric, resulting in improper filter traffic problems, which affects the effectiveness of security devices. For example, along the forwarding path in SRv6 network, the SR-aware firewall will check the association relationships of the bidirectional VPN traffic packets. And it is able to retrieve the final destination of SRv6 packet from the last entry in the SRH. When the <source, destination> tuple of the packet from PE1 to PE2 is <PE1-IP-ADDR, PE2-VPN-SID>, and the other direction is <PE2-IP-ADDR, PE1-VPN-SID>, the source address and destination address of the forward and backward VPN traffic are regarded as different flow. Eventually, the legal traffic may be blocked by the firewall.

SRv6 is commonly used as a tunneling technology in operator networks. To provide VPN service in an SRv6 network, the ingress PE encapsulates the payload with an outer IPv6 header with the SRH carrying the SR Policy segment List along with the VPN service SID. The user traffic towards SRv6 provider backbone will be encapsulated in SRv6 tunnel. When constructing an SRv6 packet, the destination address field of the SRv6 packet changes constantly and the source address field of the SRv6 packet is usually assigned using an address on the originating device, which may be a host or a network element depending on configuration. This may affect the security equipment and middle boxes in the traffic path. Because of the existence of the SRH, and the additional headers, security appliances, monitoring systems, and middle boxes could react in different ways if do not incorporate support for the supporting SRv6 mechanisms, such as the IPv6 Segment Routing Header (SRH) [RFC8754]. Additionally, implementation limitations in the processing of IPv6 packets with extension headers may result in SRv6 packets being dropped [RFC7872],[RFC9098].

8.3. Limited capability hardware

In some cases, access control lists capabilities are a resource shared with other features across a given hardware platform. Filtering capabilities should be considered along with other hardware reliant functions such as VLAN scale, route table size, MAC address table size, etc. Filtering both at the control and data plane may or may not require shared resources. For example, some platforms may require allocating resources from route table size in order to accommodate larger numbers of access lists. Hardware and software configurations should be considered when designing the filtering capabilities for an SRv6 control and data plane.

9. Security Considerations

The security considerations of SRv6 are presented throughout this document.

10. IANA Considerations

This document has no IANA actions.

11. Topics for Further Consideration

This section lists topics that will be discussed further before deciding whether they need to be included in this document, as well as some placeholders for items that need further work.

* Add tables for attack section
* The following references may be used in the future: RFC9256 [RFC8986]
* SRH compression
* Spoofing
* Path enumeration
* host to host scenario involving a WAN and/or a data center fabric.
* Terms that may be used in a future version: Locator Block, FRR, uSID
* L4 checksum: [RFC8200] specifies that when the Routing header is present the L4 checksum is computed by the originating node based on the IPv6 address of the last element of the Routing header. When compressed segment lists

[I-D.ietf-spring-srv6-srh-compression] are used, the last element of the Routing header may be different than the Destination Address as received by the final destination. Furthermore, compressed segment lists can be used in the Destination Address without the presence of a Routing header, and in this case the IPv6 Destination address can be modified along the path. As a result, some existing middleboxes which verify the L4 checksum might miscalculate the checksum. This issue is currently under discussion in the SPRING WG.

* Segment Routing Header figure: the SRv6 Segment Routing Header (SRH) is defined in [RFC8754].

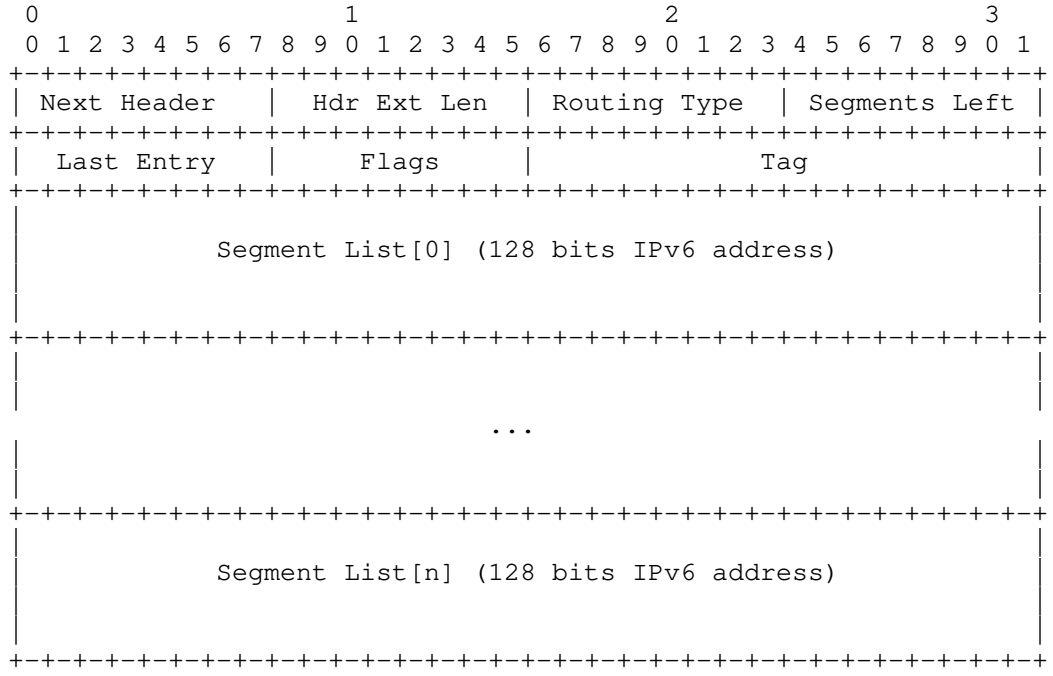

12. References

12.1. Normative References

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

[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/rfc/rfc8174>.

[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/rfc/rfc8402>.

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

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

[RFC9020] Litkowski, S., Qu, Y., Lindem, A., Sarkar, P., and J. Tantsura, "YANG Data Model for Segment Routing", RFC 9020, DOI 10.17487/RFC9020, May 2021, <https://www.rfc-editor.org/rfc/rfc9020>.

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

[RFC9491] Guichard, J., Ed. and J. Tantsura, Ed., "Integration of the Network Service Header (NSH) and Segment Routing for Service Function Chaining (SFC)", RFC 9491, DOI 10.17487/RFC9491, November 2023, <https://www.rfc-editor.org/rfc/rfc9491>.

[RFC9524] Voyer, D., Ed., Filsfils, C., Parekh, R., Bidgoli, H., and Z. Zhang, "Segment Routing Replication for Multipoint Service Delivery", RFC 9524, DOI 10.17487/RFC9524, February 2024, <https://www.rfc-editor.org/rfc/rfc9524>.

12.2. Informative References

[ANSI-Sec] "Operations, Administration, Maintenance, and Provisioning Security Requirements for the Public Telecommunications Network: A Baseline of Security Requirements for the Management Plane", 2003, <https://www.ieee802.org/1/ecsg-linksec/meetings/July03/3m150075.pdf>.

[CanSecWest2007]
"IPv6 Routing Header Security", 2007, <https://airbus-seclab.github.io/ipv6/IPv6_RH_security-csw07.pdf>.

[I-D.ietf-spring-srv6-srh-compression]
Cheng, W., Filsfils, C., Li, Z., Decraene, B., and F. Clad, "Compressed SRv6 Segment List Encoding (CSID)", Work in Progress, Internet-Draft, draft-ietf-spring-srv6-srh-compression-23, 6 February 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-spring-srv6-srh-compression-23>.

[IANAIPv6SPAR]
"IANA IPv6 Special-Purpose Address Registry", n.d., <https://www.iana.org/assignments/iana-ipv6-special-registry/iana-ipv6-special-registry.xhtml>.

[RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, DOI 10.17487/RFC3552, July 2003, <https://www.rfc-editor.org/rfc/rfc3552>.

[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/rfc/rfc5095>.

[RFC7384] Mizrahi, T., "Security Requirements of Time Protocols in Packet Switched Networks", RFC 7384, DOI 10.17487/RFC7384, October 2014, <https://www.rfc-editor.org/rfc/rfc7384>.

[RFC7855] Previdi, S., Ed., Filsfils, C., Ed., Decraene, B., Litkowski, S., Horneffer, M., and R. Shakir, "Source Packet Routing in Networking (SPRING) Problem Statement and Requirements", RFC 7855, DOI 10.17487/RFC7855, May 2016, <https://www.rfc-editor.org/rfc/rfc7855>.

[RFC7872] Gont, F., Linkova, J., Chown, T., and W. Liu, "Observations on the Dropping of Packets with IPv6 Extension Headers in the Real World", RFC 7872, DOI 10.17487/RFC7872, June 2016, <https://www.rfc-editor.org/rfc/rfc7872>.

[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/rfc/rfc8200>.

[RFC9055] Grossman, E., Ed., Mizrahi, T., and A. Hacker, "Deterministic Networking (DetNet) Security Considerations", RFC 9055, DOI 10.17487/RFC9055, June 2021, <https://www.rfc-editor.org/rfc/rfc9055>.

[RFC9098] Gont, F., Hilliard, N., Doering, G., Kumari, W., Huston, G., and W. Liu, "Operational Implications of IPv6 Packets with Extension Headers", RFC 9098, DOI 10.17487/RFC9098, September 2021, <https://www.rfc-editor.org/rfc/rfc9098>.

[RFC9099] Vyncke, Ã\211., Chittimaneni, K., Kaeo, M., and E. Rey, "Operational Security Considerations for IPv6 Networks", RFC 9099, DOI 10.17487/RFC9099, August 2021, <https://www.rfc-editor.org/rfc/rfc9099>.

[RFC9288] Gont, F. and W. Liu, "Recommendations on the Filtering of IPv6 Packets Containing IPv6 Extension Headers at Transit Routers", RFC 9288, DOI 10.17487/RFC9288, August 2022, <https://www.rfc-editor.org/rfc/rfc9288>.

[RFC9416] Gont, F. and I. Arce, "Security Considerations for Transient Numeric Identifiers Employed in Network Protocols", BCP 72, RFC 9416, DOI 10.17487/RFC9416, July 2023, <https://www.rfc-editor.org/rfc/rfc9416>.

[STRIDE] "The STRIDE Threat Model", 2018, <https://msdn.microsoft.com/en-us/library/ee823878(v=cs.20).aspx>.

Acknowledgments

The authors would like to acknowledge the valuable input and contributions from Zafar Ali, Andrew Alston, Dale Carder, Bruno Decraene, Dhruv Dhody, Joel Halpern, Bruno Hassanov, Alvaro Retana, Eric Vyncke, and Russ White.

Authors' Addresses

Nick Buraglio
Energy Sciences Network
Email: buraglio@forwardingplane.net

Tal Mizrahi
Huawei
Email: tal.mizrahi.phd@gmail.com

Tian Tong
China Unicom
Email: tongt5@chinaunicom.cn

Luis M. Contreras
Telefonica
Email: luismiguel.contrerasmurillo@telefonica.com

Fernando Gont
SI6 Networks
Email: fgont@si6networks.com