

Source Packet Routing in Networking
Internet-Draft
Intended status: Standards Track
Expires: 22 March 2026

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

Segment Routing IPv6 Security Considerations
draft-ietf-spring-srv6-security-07

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://github.com/buraglio/draft-bdmgct-spring-srv6-security>.
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.

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 22 March 2026.

Copyright Notice

Copyright (c) 2025 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	3
2. Scope of this Document	4
3. Conventions and Definitions	4
3.1. Requirements Language	5
3.2. Terminology	5
4. Threat Terminology	5
5. Effect	6
6. Attacks	8
6.1. Attack Abstractions	8
6.2. Data Plane Attacks	9
6.2.1. Modification Attack	9
6.2.2. Passive Listening	11
6.2.3. Packet Insertion	11
6.2.4. Other Attacks	12
6.3. Control Plane Attacks	12
6.3.1. Overview	12
6.3.2. Routing Protocol Attacks	12
6.3.3. OAM Attacks	14
6.3.4. Central Control Plane Attacks	15
6.4. Management Plane Attacks	15

6.4.1. Overview	16
6.5. Attacks - Summary	17
7. Mitigation Methods	18
7.1. Trusted Domains and Filtering	18
7.1.1. Overview	18
7.1.2. SRH Filtering	19
7.1.3. Address Range Filtering	19
7.2. Encapsulation of Packets	20
7.3. Hashed Message Authentication Code (HMAC)	20
7.4. Control Plane Mitigation Methods	21
7.5. Management Plane Mitigation Methods	21
8. Implications on Existing Equipment	22
8.1. Middlebox Filtering Issues	22
8.2. Limited capability hardware	23
9. Security Considerations	23
10. IANA Considerations	23
11. References	23
11.1. Normative References	24
11.2. Informative References	25
Acknowledgments	29
Authors' Addresses	29

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]. This header is used 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 the aforementioned 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 in [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 and a 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 Terminology

This section introduces the threat taxonomy that is used in this document. This taxonomy 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 within the premises of 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 control plane perspective.

The following figure depicts an example of an SR domain with five attacker types, labeled 1-5. As an example, 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. In this example a Path Computation Element as a Central Controller (PCECC) [RFC9050] is used as part of the control plane. Thus, attacker 3 is an internal on-path attacker in the control plane, as it is located along the path between the PCECC 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. Effect

One of the important aspects of threat analysis is assessing the potential effect or outcome 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 effect, 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 that are unable to enforce security policies for SRv6. 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 affecting the system integrity. Specific SRv6-targeted attack may cause one or more of the following outcomes:
 - Avoiding a specific node or path: when an SRv6 policy is manipulated, specific nodes or paths may be bypassed, for example in order to avoid the billing service or circumvent access controls and security filters.
 - Preferring a specific path: packets can be manipulated so that they are diverted to a specific path. This 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, which facilitates more complex on-path attacks such as passive listening, recon and various man-in-the-middle attacks.
 - Causing header modifications: SRv6 network programming determines the SR endpoint behavior, including potential header modifications. Thus, one of the potential outcomes of an attack is unwanted header modifications.
- * 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 of user data can be avoided by using encryption in higher layers, it is not within the scope of this document. However, eavesdropping of a network that uses SRv6 is a specific form of reconnaissance. This reconnaissance allows the attacker to collect information about SR endpoint addresses, SR policies, and network topologies.
- * 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), including:
 - Resource exhaustion: compromising the availability of the system can be achieved by sending SRv6-enabled packets to/through victim nodes in a way that results in a negative performance impact of the victim systems (e.g., [RFC9098]). For example, network programming can be used in some cases to manipulate segment endpoints to perform unnecessary functions that consume processing resources. Resource exhaustion may in severe cases cause Denial of Service (DoS).
 - Forwarding loops: an attacker might achieve attack amplification by increasing the number hops that each packet is forwarded through and thus increase the load on the network. For instance, a set of SIDs can be inserted in a way that creates a forwarding loop ([RFC8402], [RFC5095], [CanSecWest2007]) and thus loads the nodes along the loop.
 - Causing packets to be discarded: an attacker 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.

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. This could lead to using more resources or security devices being unable to track connections correctly.
- * Packet insertion: an attacker generates and injects a packet to the network. The generated packet may be maliciously crafted to include false information; including false addresses, SRv6-related information, or other intentionally incorrect 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. Data Plane Attacks

6.2.1. Modification Attack

6.2.1.1. 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.

A modification attack can be performed in one or more of the following ways:

- * SID list: the SRH can be manipulated by adding or removing SIDs, or by modifying existing SIDs.
- * IPv6 Destination Address (DA): when an SRH is present modifying the destination address (DA) of the IPv6 header affects the active segment. However, DA modification can affect the SR policy even in the absence of an SRH. One example is modifying a DA which is used as a Binding SID [RFC8402]. Another example is modifying a

DA which represents a compressed segment list [RFC9800]. SRH compression allows encoding multiple compressed SIDs within a single 128-bit SID, and thus modifying the DA can affect one or more hops in the SR policy.

- * Add/remove SRH: an attacker can insert or remove an SRH.
- * SRH TLV: adding, removing or modifying TLV fields in the SRH.

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 redirecting the traffic to another node that the attacker has access to with more processing resources.

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

6.2.1.2. 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.

6.2.1.3. Effect

SR modification attacks, including adding or removing an SRH, modifying the SID list, and modifying the IPv6 DA, can have one or more of the following outcomes, which are described in Section 5.

- * Unauthorized access
- * Avoiding a specific node or path
- * Preferring a specific path
- * Causing header modifications
- * Causing packets to be discarded
- * Resource exhaustion
- * Forwarding loops

Maliciously adding unnecessary TLV fields can cause further resource exhaustion.

6.2.2. Passive Listening

6.2.2.1. 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.

6.2.2.2. Scope

A reconnaissance 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.

6.2.2.3. Effect

While the information collected in a reconnaissance 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.

6.2.3. Packet Insertion

6.2.3.1. 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.

6.2.3.2. 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.

6.2.3.3. Effect

The main effect of this attack is resource exhaustion, which compromises the availability of the network, as described in Section 6.2.1.3.

6.2.4. 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]. This, however, is out of scope for this document.

6.3. Control Plane Attacks

6.3.1. Overview

The SRv6 control plane leverages existing control plane protocols, such as BGP, IS-IS, OSPF and PCEP. Consequently, any security attacks that can potentially compromise these protocols are also applicable to SRv6 deployments utilizing them. Therefore, this document does not provide an exhaustive list of the potential control plane attacks. Instead, it highlights key categories of attacks, focusing on three primary areas: attacks targeting routing protocols, centralized control plane infrastructures, and OAM protocols. In this document, the term OAM refers specifically to Operations, Administration, and Maintenance, in alignment with the definition provided in [RFC6291]. As such, it explicitly excludes management-related functions. Security considerations pertaining to the management plane are addressed in Section 6.4.

6.3.2. Routing Protocol Attacks

6.3.2.1. Overview

Generic threats applicable to routing protocols are discussed in [RFC4593]. Similar to data plane attacks, the abstractions outlined in Section 6.1 are also applicable to control plane traffic. These include passive eavesdropping, message injection, replay, deletion, and modification.

Passive listening enables an attacker to intercept routing protocol messages as they traverse the network. This form of attack does not alter the content of the messages but allows the adversary to analyze routing information, infer network topology, and gather intelligence on routing behavior.

Active attacks involve the unauthorized injection or alteration of control plane messages. Such attacks can compromise routing integrity by introducing falsified information, modifying legitimate routing data, or triggering incorrect forwarding decisions. These disruptions may result in denial-of-service conditions or traffic misdirection.

For example, an attacker may advertise falsified SIDs to manipulate SR policies. Another example in the context of SRv6 is the advertisement of an incorrect Maximum SID Depth (MSD) value [RFC8476]. If the advertised MSD is lower than the actual capability, path computation may fail to compute a viable path. Conversely, if the value is higher than supported, an attempt to instantiate a path that can't be supported by the head-end (the node performing the SID imposition) may occur.

An additional case could be the manipulation of backup paths [RFC8355], where the attacker could alter the SIDs defining such backup path then directing traffic over suboptimal or compromised paths, enabling eavesdropping, traffic analysis, or selective denial of service, compromising the service integrity and confidentiality if traffic is diverted to unauthorized nodes or paths.

Finally, in situations of interworking with other domains, as for BGP Egress Peer Engineering (BGP-EPE) [RFC9087] an attacker injecting malicious BGP-EPE policies may steer traffic through unauthorized peers or paths. This facilitates interception, traffic analysis, or denial of service. Attackers gaining access to the BGP-EPE controller can manipulate SRv6 route selection and segment lists, compromising network integrity and confidentiality.

6.3.2.2. Scope

The location of an attacker in the network significantly affects the scope of potential attacks. Off-path attackers are generally limited to injecting malicious routing messages, while on-path attackers can perform a broader range of attacks, including active modification, or passive listening.

6.3.2.3. Effect

Attacks targeting the routing protocol can have diverse impacts on network operation, including the aspects described in Section 5. These impacts may include incorrect SR policies or the degradation of network availability, potentially resulting in service disruption or denial of service.

6.3.3. OAM Attacks

6.3.3.1. Overview

Since SRv6 operates over an IPv6 infrastructure, existing OAM protocols designed for IPv6 networks are applicable to SRv6 as well. Consequently, the security considerations associated with conventional IPv6 OAM protocols are also relevant to SRv6 environments. As noted in [RFC7276], successful attacks on OAM protocols can mislead operators by simulating non-existent failures or by concealing actual network issues. SRv6-specific OAM aspects are specified in [RFC9259].

The O-flag in the SRH serves as a marking bit in user packets to trigger telemetry data collection and export at the segment endpoints. An attacker may exploit this mechanism by setting the O-flag in transit packets, thereby overloading the control plane and degrading system availability. Additionally, an on-path attacker may passively intercept OAM data exported to external analyzers, potentially gaining unauthorized insight into network topology and behavior.

6.3.3.2. Scope

Off-path attackers may attempt to degrade system availability by injecting fabricated OAM messages or SRv6 packets with the O-bit set, thereby triggering unnecessary telemetry processing. They may also probe SRv6 nodes to infer information about network state and performance characteristics.

On-path attackers possess enhanced capabilities due to their position within the traffic path. These include passive interception of OAM data, unauthorized modification of the O-bit in transit packets, and tampering with legitimate OAM messages to mislead network monitoring systems or conceal operational issues.

6.3.3.3. Effect

Attacks targeting OAM protocols may impact network availability or facilitate unauthorized information gathering. Such attacks can disrupt normal operations or expose sensitive details about network topology, performance, or state.

6.3.4. Central Control Plane Attacks

6.3.4.1. Overview

Centralized control plane architectures, such as those based on the Path Computation Element (PCE) [RFC4655] and PCE as a Central Controller (PCECC) [RFC8283], inherently introduce a single point of failure. This centralization may present a security vulnerability, particularly with respect to denial-of-service (DoS) attacks targeting the controller. Furthermore, the central controller becomes a focal point for potential interception or manipulation of control messages exchanged with individual Network Elements (NEs), thereby increasing the risk of compromise to the overall network control infrastructure.

6.3.4.2. Scope

As with other control plane attacks, an off-path attacker may attempt to inject forged control messages or impersonate a legitimate controller. On-path attackers, by virtue of their position within the communication path, possess additional capabilities such as passive interception of control traffic and in-transit modification of messages exchanged between the controller and Network Elements (NEs).

For example, an attacker may manipulate SR policies instantiated via the central controller (using protocols like PCEP or BGP) at the head end, thereby altering both the paths of the SR policy and the traffic steered over it. Additionally, PCECC enables manipulation of SID allocation and distribution.

6.3.4.3. Effect

A successful attack may result in any of the adverse effects described in Section 5, potentially impacting availability and operational correctness.

6.4. Management Plane Attacks

6.4.1. Overview

Similar to the control plane, a compromised management plane can enable a broad range of attacks, including unauthorized manipulation of SR policies and disruption of network availability. The specific threats and their potential impact are influenced by the management protocols in use.

As with centralized control systems, a centralized management infrastructure may introduce a single point of failure, rendering it susceptible to denial-of-service (DoS) attacks or making it a target for eavesdropping and message tampering.

Unauthorized access in a network management system can enable attackers or unprivileged users to gain control over network devices and alter configurations. In SRv6-enabled environments, this can result in the manipulation of segment routing policies or cause denial-of-service (DoS) conditions by disrupting traffic or tampering with forwarding behavior.

Management functionality is often defined using YANG data models, such as those specified in [RFC9020], [I-D.ietf-lsr-isis-srv6-yang] and [I-D.ietf-lsr-ospf-srv6-yang]. As with any YANG module, data nodes marked as writable, creatable, or deletable may be considered sensitive in certain operational environments. Unauthorized or unprotected write operations (e.g., via edit-config) targeting these nodes can adversely affect network operations. Some of the readable data nodes in a YANG module may also be considered sensitive or vulnerable in some network environments.

6.4.1.1. Scope

As with control plane attacks, an off-path attacker may attempt to inject forged management messages or impersonate a legitimate network management system. On-path attackers, due to their privileged position within the communication path, have additional capabilities such as passive interception of management traffic and unauthorized modification of messages in transit. An attacker with unauthorized access to a management system can cause significant damage, depending on the scope of the system and the strength of the access control mechanisms in place.

6.4.1.2. Effect

A successful attack may result in any of the adverse effects described in Section 5, potentially impacting availability and operational correctness.

6.5. Attacks - Summary

The following table summarizes the attacks that were described in the previous subsections, and the corresponding effect of each of the attacks. Details about the effect are described in Section 5.

Attack	Details	Effect
Modification	Modification of: <ul style="list-style-type: none"> * SID list * IPv6 DA Add/remove/modify: <ul style="list-style-type: none"> * SRH * SRH TLV 	<ul style="list-style-type: none"> * Unauthorized access * Avoiding a specific node or path * Preferring a specific path * Causing header modifications * Causing packets to be discarded * Resource exhaustion * Forwarding loops
Passive listening	Passively listen to SRv6-related information	* Reconnaissance
Packet insertion	Maliciously inject packets with a segment list	<ul style="list-style-type: none"> * Resource exhaustion * Security tooling confusion
Control plane attacks	<ul style="list-style-type: none"> * Routing protocol attacks * OAM attacks * Central control plane attacks 	<ul style="list-style-type: none"> * Unauthorized access * Avoiding a specific node or path * Preferring a specific path * Causing header modifications
Management plane attacks	<ul style="list-style-type: none"> * Centralized management attacks * Unauthorized access to the management system 	<ul style="list-style-type: none"> * Causing packets to be discarded * Resource exhaustion * Forwarding loops

Figure 2: Attack Summary

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).

Such an approach has been commonly referred to as the concept of "fail-open", a state of which the attributes are frequently described as containing inherently more risk than fail-closed methodologies. The reliance of perfectly crafted filters on on all edges of the trusted domain pose a demonstrable risk of inbound or outbound leaks if the filters are removed or adjusted in an erroneous manner. It is also important to note that some filtering implementations have limits on the size, complexity, or protocol support that can be applied, which may prevent the filter adjustments or creation required to properly secure the trusted domain for a new protocol such as SRv6.

Practically speaking, this means successfully enforcing a "Trusted Domain" may be operationally difficult and error-prone in practice, and that attacks that are expected to be unfeasible from outside the trusted domain may actually become feasible when any of the involved systems fails to enforce the filtering policy that is required to define the Trusted Domain.

7.1.2. SRH Filtering

Filtering can be performed based on the presence of an SRH. More generally, [RFC9288] provides recommendations on the filtering of IPv6 packets containing IPv6 extension headers at transit routers. However, filtering based on the presence of an SRH is not necessarily 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 necessarily a useful method of mitigation.

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 and a summary is presented here: 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 nodes. 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.

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.

7.2. Encapsulation of Packets

Packets steered in an SR domain are often contained in an IPv6 encapsulation. 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 and also allows for encapsulation of both IPv4 and IPv6 packets.

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 6.2.1).

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.

These considerations limit the extent to which HMAC TLV can be relied upon as a security mechanism that could readily mitigate threats associated with spoofing and tampering protection for the IPv6 SRH.

7.4. Control Plane Mitigation Methods

Mitigation strategies for control plane attacks depend heavily on the specific protocols in use. Since these protocols are not exclusive to SRv6, this section does not attempt to provide an exhaustive list of mitigation techniques. Instead, it is focused on considerations particularly relevant to SRv6 deployments.

Routing protocols can employ authentication and/or encryption to protect against modification, injection, and replay attacks, as outlined in [RFC6518]. These mechanisms are essential for maintaining the integrity and authenticity of control plane communications.

In centralized SRv6 control plane architectures, such as those described in [I-D.ietf-pce-segment-routing-policy-cp], it is RECOMMENDED that communication between PCEs and PCCs be secured using authenticated and encrypted sessions. This is typically achieved using Transport Layer Security (TLS), following the guidance in [RFC8253] and best practices in [RFC9325].

When the O-flag is used for Operations, Administration, and Maintenance (OAM) functions, as defined in [RFC9259], implementations should enforce rate limiting to mitigate potential denial-of-service (DoS) attacks triggered by excessive control plane signaling.

The control plane should be confined to a trusted administrative domain. As specified in [I-D.ietf-idr-bgp-ls-sr-policy], SR Policy information advertised via BGP should be restricted to authorized nodes, controllers, and applications within this domain. Similarly, the use of the O-flag is assumed to occur only within such a trusted environment, where the risk of abuse is minimized.

7.5. Management Plane Mitigation Methods

Mitigating attacks on the management plane, much like in the control plane, depends on the specific protocols and interfaces employed.

Management protocols such as NETCONF and RESTCONF are commonly used to configure and monitor SRv6-enabled devices. These protocols must be secured to prevent unauthorized access, configuration tampering, or information leakage.

The lowest NETCONF layer is the secure transport layer, and the mandatory-to-implement secure transport is Secure Shell (SSH) [RFC6242]. The lowest RESTCONF layer is HTTPS, and the mandatory-to-implement secure transport is TLS [RFC8446].

The Network Configuration Access Control Model (NACM) [RFC8341] provides the means to restrict access for particular NETCONF or RESTCONF users to a pre-configured subset of all available NETCONF or RESTCONF protocol operations and content.

SRv6-specific YANG modules should be designed with the same security considerations applied to all YANG-based models. Writable nodes must be protected using access control mechanisms such as NACM and secured transport protocols like SSH or TLS to prevent unauthorized configuration changes. Readable nodes that expose sensitive operational data should be access-controlled and transmitted only over encrypted channels to mitigate the risk of information leakage.

8. Implications on Existing Equipment

8.1. Middlebox Filtering Issues

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

The security devices on SRv6 networks need to take care of SRv6 packets. However, SRv6 packets are often encapsulated by an SR ingress device with an IPv6 encapsulation that has the loopback address of the SR ingress device as a source address. As a result, the address information of SR packets may be asymmetric, resulting in improper traffic filter 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. It is therefore able to retrieve the final destination of an SRv6 packet from the last entry in the SRH. When the <source, destination> tuple of the packet from PE1 (Provider Edge 1) 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 traffic are regarded as different flows. Thus, legitimate traffic may be blocked by the firewall.

Forwarding SRv6 traffic through devices that are not SRv6-aware might in some cases lead to unpredictable behavior. Because of the existence of the SRH, and the additional headers, security appliances, monitoring systems, and middle boxes could react in different ways if they 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].

Upper-layer checksum calculations rely on a pseudo-header that includes the IPv6 Destination Address. [RFC8200] specifies that when the Routing header is present the upper-layer checksum is computed by the originating node based on the IPv6 address of the last element of the Routing header. When compressed segment lists [RFC9800] 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 defined in [RFC9800], the Destination Address used in the upper-layer checksum calculation is the address as expected to be received by the ultimate destination. As a result, some existing middleboxes which verify the upper-layer checksum might miscalculate the checksum.

8.2. Limited capability hardware

In some cases, access control list 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. References

11.1. 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>>.

11.2. 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-idr-bgp-ls-sr-policy] Previdi, S., Talaulikar, K., Dong, J., Gredler, H., and J. Tantsura, "Advertisement of Segment Routing Policies using BGP Link-State", Work in Progress, Internet-Draft, draft-ietf-idr-bgp-ls-sr-policy-17, 6 March 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-ls-sr-policy-17>>.
- [I-D.ietf-lsr-isis-srv6-yang] Hu, Z., Ye, D., Qu, Y., Geng, X., and Q. Ma, "YANG Data Model for IS-IS SRv6", Work in Progress, Internet-Draft, draft-ietf-lsr-isis-srv6-yang-08, 2 September 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-lsr-isis-srv6-yang-08>>.
- [I-D.ietf-lsr-ospf-srv6-yang] Qu, Y., Hu, Z., Geng, X., Raza, S. K., and A. Lindem, "YANG Data Model for OSPF SRv6", Work in Progress, Internet-Draft, draft-ietf-lsr-ospf-srv6-yang-08, 2 September 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-lsr-ospf-srv6-yang-08>>.
- [I-D.ietf-pce-segment-routing-policy-cp] Koldychev, M., Sivabalan, S., Sidor, S., Barth, C., Peng, S., and H. Bidgoli, "Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing (SR) Policy Candidate Paths", Work in Progress, Internet-Draft, draft-ietf-pce-segment-routing-policy-cp-27, 4 April 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-pce-segment-routing-policy-cp-27>>.
- [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>>.

- [RFC4381] Behringer, M., "Analysis of the Security of BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4381, DOI 10.17487/RFC4381, February 2006, <<https://www.rfc-editor.org/rfc/rfc4381>>.
- [RFC4593] Barbir, A., Murphy, S., and Y. Yang, "Generic Threats to Routing Protocols", RFC 4593, DOI 10.17487/RFC4593, October 2006, <<https://www.rfc-editor.org/rfc/rfc4593>>.
- [RFC4655] Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, DOI 10.17487/RFC4655, August 2006, <<https://www.rfc-editor.org/rfc/rfc4655>>.
- [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>>.
- [RFC6242] Wasserman, M., "Using the NETCONF Protocol over Secure Shell (SSH)", RFC 6242, DOI 10.17487/RFC6242, June 2011, <<https://www.rfc-editor.org/rfc/rfc6242>>.
- [RFC6291] Andersson, L., van Helvoort, H., Bonica, R., Romascanu, D., and S. Mansfield, "Guidelines for the Use of the "OAM" Acronym in the IETF", BCP 161, RFC 6291, DOI 10.17487/RFC6291, June 2011, <<https://www.rfc-editor.org/rfc/rfc6291>>.
- [RFC6518] Lebovitz, G. and M. Bhatia, "Keying and Authentication for Routing Protocols (KARP) Design Guidelines", RFC 6518, DOI 10.17487/RFC6518, February 2012, <<https://www.rfc-editor.org/rfc/rfc6518>>.
- [RFC7276] Mizrahi, T., Sprecher, N., Bellagamba, E., and Y. Weingarten, "An Overview of Operations, Administration, and Maintenance (OAM) Tools", RFC 7276, DOI 10.17487/RFC7276, June 2014, <<https://www.rfc-editor.org/rfc/rfc7276>>.
- [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>>.
- [RFC7835] Saucez, D., Iannone, L., and O. Bonaventure, "Locator/ID Separation Protocol (LISP) Threat Analysis", RFC 7835, DOI 10.17487/RFC7835, April 2016, <<https://www.rfc-editor.org/rfc/rfc7835>>.

- [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>>.
- [RFC8253] Lopez, D., Gonzalez de Dios, O., Wu, Q., and D. Dhody,
"PCEPS: Usage of TLS to Provide a Secure Transport for the
Path Computation Element Communication Protocol (PCEP)",
RFC 8253, DOI 10.17487/RFC8253, October 2017,
<<https://www.rfc-editor.org/rfc/rfc8253>>.
- [RFC8283] Farrel, A., Ed., Zhao, Q., Ed., Li, Z., and C. Zhou, "An
Architecture for Use of PCE and the PCE Communication
Protocol (PCEP) in a Network with Central Control",
RFC 8283, DOI 10.17487/RFC8283, December 2017,
<<https://www.rfc-editor.org/rfc/rfc8283>>.
- [RFC8341] Bierman, A. and M. Bjorklund, "Network Configuration
Access Control Model", STD 91, RFC 8341,
DOI 10.17487/RFC8341, March 2018,
<<https://www.rfc-editor.org/rfc/rfc8341>>.
- [RFC8355] Filsfils, C., Ed., Previdi, S., Ed., Decraene, B., and R.
Shakir, "Resiliency Use Cases in Source Packet Routing in
Networking (SPRING) Networks", RFC 8355,
DOI 10.17487/RFC8355, March 2018,
<<https://www.rfc-editor.org/rfc/rfc8355>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol
Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018,
<<https://www.rfc-editor.org/rfc/rfc8446>>.
- [RFC8476] Tantsura, J., Chunduri, U., Aldrin, S., and P. Psenak,
"Signaling Maximum SID Depth (MSD) Using OSPF", RFC 8476,
DOI 10.17487/RFC8476, December 2018,
<<https://www.rfc-editor.org/rfc/rfc8476>>.

- [RFC9050] Li, Z., Peng, S., Negi, M., Zhao, Q., and C. Zhou, "Path Computation Element Communication Protocol (PCEP) Procedures and Extensions for Using the PCE as a Central Controller (PCECC) of LSPs", RFC 9050, DOI 10.17487/RFC9050, July 2021, <<https://www.rfc-editor.org/rfc/rfc9050>>.
- [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>>.
- [RFC9087] Filssils, C., Ed., Previdi, S., Dawra, G., Ed., Aries, E., and D. Afanasiev, "Segment Routing Centralized BGP Egress Peer Engineering", RFC 9087, DOI 10.17487/RFC9087, August 2021, <<https://www.rfc-editor.org/rfc/rfc9087>>.
- [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, ., 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>>.
- [RFC9259] Ali, Z., Filssils, C., Matsushima, S., Voyer, D., and M. Chen, "Operations, Administration, and Maintenance (OAM) in Segment Routing over IPv6 (SRv6)", RFC 9259, DOI 10.17487/RFC9259, June 2022, <<https://www.rfc-editor.org/rfc/rfc9259>>.
- [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>>.
- [RFC9325] Sheffer, Y., Saint-Andre, P., and T. Fossati, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", BCP 195, RFC 9325, DOI 10.17487/RFC9325, November 2022, <<https://www.rfc-editor.org/rfc/rfc9325>>.
- [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>>.

- [RFC9602] Krishnan, S., "Segment Routing over IPv6 (SRv6) Segment Identifiers in the IPv6 Addressing Architecture", RFC 9602, DOI 10.17487/RFC9602, October 2024, <<https://www.rfc-editor.org/rfc/rfc9602>>.
- [RFC9800] Cheng, W., Ed., Filsfils, C., Li, Z., Decraene, B., and F. Clad, Ed., "Compressed SRv6 Segment List Encoding", RFC 9800, DOI 10.17487/RFC9800, June 2025, <<https://www.rfc-editor.org/rfc/rfc9800>>.

Acknowledgments

The authors would like to acknowledge the valuable input and contributions from Zafar Ali, Andrew Alston, Dale Carder, Bruno Decraene, Dhruv Dhody, Mike Dopheide, Darren Dukes, Joel Halpern, Boris Hassanov, Sam Oehlert, 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