

SAVNET
Internet-Draft
Intended status: Informational
Expires: 27 February 2026

D. Li
J. Wu
Tsinghua University
L. Qin
M. Huang
Zhongguancun Laboratory
N. Geng
Huawei
26 August 2025

Source Address Validation in Intra-domain Networks Gap Analysis, Problem
Statement, and Requirements
draft-ietf-savnet-intra-domain-problem-statement-18

Abstract

This document provides a gap analysis of existing intra-domain source address validation mechanisms, describes the fundamental problems, and defines the basic requirements for technical improvements.

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 27 February 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	2
1.1. Terminology	5
1.2. Requirements Language	5
2. Current Operational Intra-domain SAV Mechanisms	5
3. Gap Analysis	6
3.1. Intra-domain SAV for Traffic from Sub Networks or Directly-Connected Hosts	6
3.1.1. Asymmetric Routing	7
3.1.2. Hidden Prefix	9
3.2. Intra-domain SAV for Traffic from the Rest of the Internet	11
4. Problem Statement	12
5. Requirements for New SAV Mechanisms	13
5.1. Accurate Validation	13
5.2. Automatic Update	13
5.3. Working in Incremental Deployment	13
5.4. Fast Convergence	14
5.5. Security	14
6. Security Considerations	14
7. IANA Considerations	14
8. Acknowledgements	14
9. References	14
9.1. Normative References	15
9.2. Informative References	15
Authors' Addresses	16

1. Introduction

Source Address Validation (SAV) defends against source address spoofing. Network operators can enforce SAV at the following levels (see [RFC5210]):

- * Within in the access network
- * Within in the domain (i.e., the autonomous system)
- * Between domains (i.e., autonomous systems)

Access networks have already deployed SAV mechanisms. These mechanisms typically are deployed on switches and prevent hosts from using the source address of another host on the Internet. Mechanisms include:


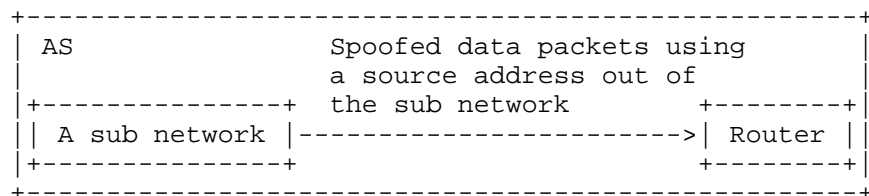
- * Source Address Validation Improvement (SAVI) Solution for DHCP [RFC7513]
- * IP Source Guard (IPSG) based on DHCP snooping [IPSG]
- * Cable Source-Verify [cable-verify]

A domain can trust a sub network or an access network if that network reliably enforces SAV. Sadly, SAV mechanisms are not universally deployed and many domains cannot trust the sub networks or hosts to which they are attached. Therefore, intra-domain SAV mechanisms are required.

This document analyzes intra-domain (i.e., intra-AS) SAV. As illustrated in Figure 1, intra-domain SAV has the following goals:

- * To prevent sub-networks from injecting packets with spoofed source addresses outside their own sub-network, and to prevent hosts from injecting packets with spoofed source addresses outside their own network segment into the domain
- * To prevent other domains (i.e., autonomous systems) from injecting packets with spoofed internal-use-only source addresses into the domain

Goal i: To prevent sub-networks from injecting packets with spoofed source addresses outside their own sub-network, and to prevent hosts from injecting packets with spoofed source addresses outside their own network segment into the domain



Goal ii: To prevent other domains (i.e., autonomous systems) from injecting packets with spoofed internal-use-only source addresses into the domain

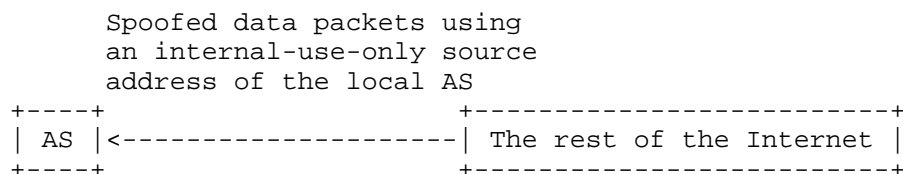


Figure 1: Two Goals of intra-domain SAV

Building on the second goal of intra-domain SAV, inter-domain (i.e., inter-AS) SAV further prevents other domains from injecting packets with other spoofed source addresses into the domain.

This document provides a gap analysis of the current operational intra-domain SAV mechanisms, identifies key problems to solve, and proposes basic requirements for solutions.

1.1. Terminology

Sub Network: A sub network may operate its own internal routing protocols (e.g., a separate IGP), but it is considered part of the AS in the global routing system. It is connected to one or more routers of the AS for Internet connectivity. It originates traffic but does not transit traffic for other networks.

SAV Rule: The rule in a router that describes the mapping relationship between a source address (prefix) and the valid incoming interface(s). It is used by a router to make SAV decisions.

Improper Block: The validation results that the packets with legitimate source addresses are blocked improperly due to inaccurate SAV rules.

Improper Permit: The validation results that the packets with spoofed source addresses are permitted improperly due to inaccurate SAV rules.

SAV-specific Information: The information specialized for SAV rule generation.

1.2. 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.

2. Current Operational Intra-domain SAV Mechanisms

Although BCP38 [RFC2827] and BCP84 [RFC3704] provides multiple ingress filtering methods which are typically used for inter-domain SAV, some of them have also been used for intra-domain SAV in practice. This section introduces the current operational mechanisms used to implement intra-domain SAV.

- * Access Control List (ACL) [RFC2827] is a SAV filter that checks the source address of every data packet against a list of acceptable or unacceptable prefixes. By performing the ACL rule at a router interface, every data packet received at this interface that does not match the ACL rule will be blocked. It can be used at router interfaces facing a sub network or a set of hosts, only allowing data packets using an acceptable source address. It is also usually used on router interfaces facing the rest of the Internet, blocking data packets using an unacceptable

source address, such as internal source addresses owned by the local AS [nist-rec]. The ACL rule is typically configured and updated manually, so it is critical to update ACL rules in time when the set of prefixes changes.

- * Strict uRPF [RFC3704] implements a SAV filter in an automatic way by checking the source address of every data packet against the router's local Forwarding Information Base (FIB). The router deploying strict uRPF accepts a data packet only when i) the local FIB contains a prefix covering the packet's source address and ii) the corresponding outgoing interface for the prefix in the FIB matches the packet's incoming interface. Otherwise, the packet will be blocked. Strict uRPF is often used to block spoofed data packets originated from a sub network or a directly-connected host.
- * Loose uRPF [RFC3704] also uses the local FIB to implement SAV but checks only for the existence of the prefix. Loose uRPF accepts a data packet if the router's local FIB contains a prefix covering the packet's source address regardless of the interface from which the packet is received. Routers can use loose uRPF to block spoofed data packets using a non-global or non-routed source address [nist-rec].

EFP-uRPF [RFC8704] is another advanced SAV mechanism but it is specifically designed for inter-domain SAV. It implements a SAV filter on router interfaces facing a customer AS by using BGP data provided by other ASes. This document does not analyze EFP-uRPF because it is out of the scope.

3. Gap Analysis

This section elaborates the gap scenarios and key problems of the current operational intra-domain SAV mechanisms.

3.1. Intra-domain SAV for Traffic from Sub Networks or Directly-Connected Hosts

Towards Goal i described in Section 1 and shown in Figure 1, the AS operator can use ACL rules or strict uRPF on appropriate routers to implement intra-domain SAV for traffic originated from a sub network or a directly-connected host.

For example, the AS operator can configure an ACL rule on router interfaces facing a sub network or a directly-connected host, containing a set of prefixes which can be used as the source address. By using this ACL rule, the router will block any data packet using a source address out of the set of prefixes. However, the key problem

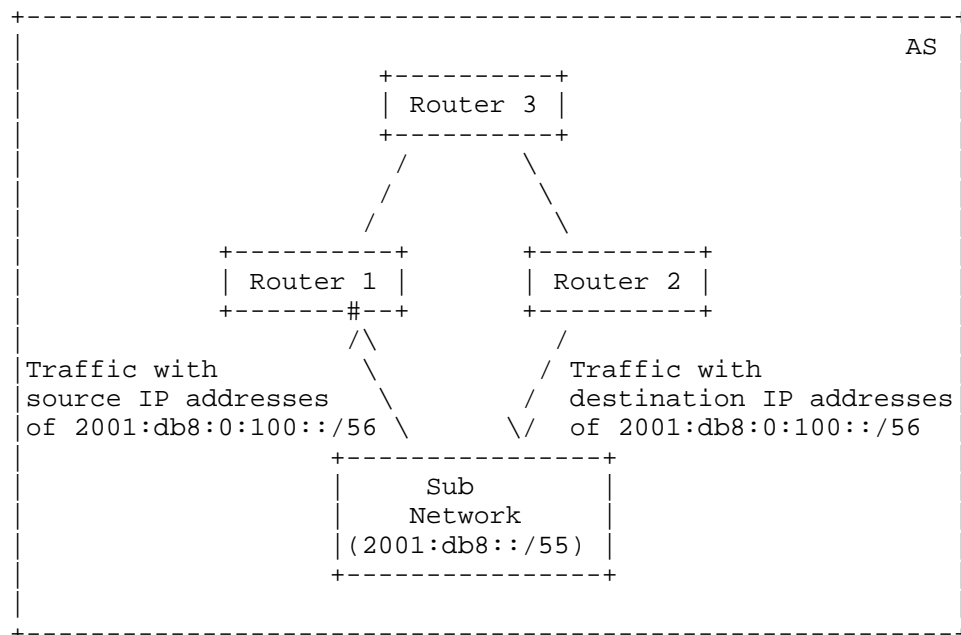
of ACL-based SAV is the high operational overhead. Since the ACL rule is typically maintained manually, the AS operator must update the ACL rule according to prefix changes or topology changes timely. If the AS operator forgets to update the ACL rule when the set of prefixes changes, the outdated ACL rule may improperly block legitimate data packets.

Strict uRPF can generate and update SAV rules in an automatic way but it will improperly block legitimate data packets in the scenario of asymmetric routing or hidden prefix. In the following, this section introduces two gap scenarios when using strict uRPF to implement intra-domain SAV.

3.1.1. Asymmetric Routing

Asymmetric routing means a packet traverses from a source to a destination in one path and takes a different path when it returns to the source. Asymmetric routing can occur within an AS due to routing policy, traffic engineering, etc. For example, a sub network connected to multiple routers of the AS may need to perform load balancing on incoming traffic, thereby resulting in asymmetric routing.

Figure 2 shows an example of asymmetric routing. The sub network owns prefix 2001:db8::/55 [RFC6890] and is connected to two routers of the AS, i.e., Router 1 and Router 2. Router 1, Router 2, and Router 3 exchange routing information through the intra-domain routing protocol. For load balancing of traffic flowing to the sub network, the sub network expects the incoming traffic destined for prefix 2001:db8:0::/56 to come from Router 1 and the incoming traffic destined for prefix 2001:db8:0:100::/56 to come from Router 2. To this end, it requires that Router 1 advertises the route information of prefix 2001:db8:0::/56 and Router 2 advertises the routing information of prefix 2001:db8:0:100::/56 through the intra-domain routing protocol. Figure 2 shows the FIB entries of Router 1 and Router 2 associated with the two prefixes.



FIB of Router 1

Dest	Next_hop
2001:db8:0::/56	Sub
2001:db8:0:100::/56	Router 3

FIB of Router 2

Dest	Next_hop
2001:db8:0:100::/56	Sub
2001:db8:0::/56	Router 3

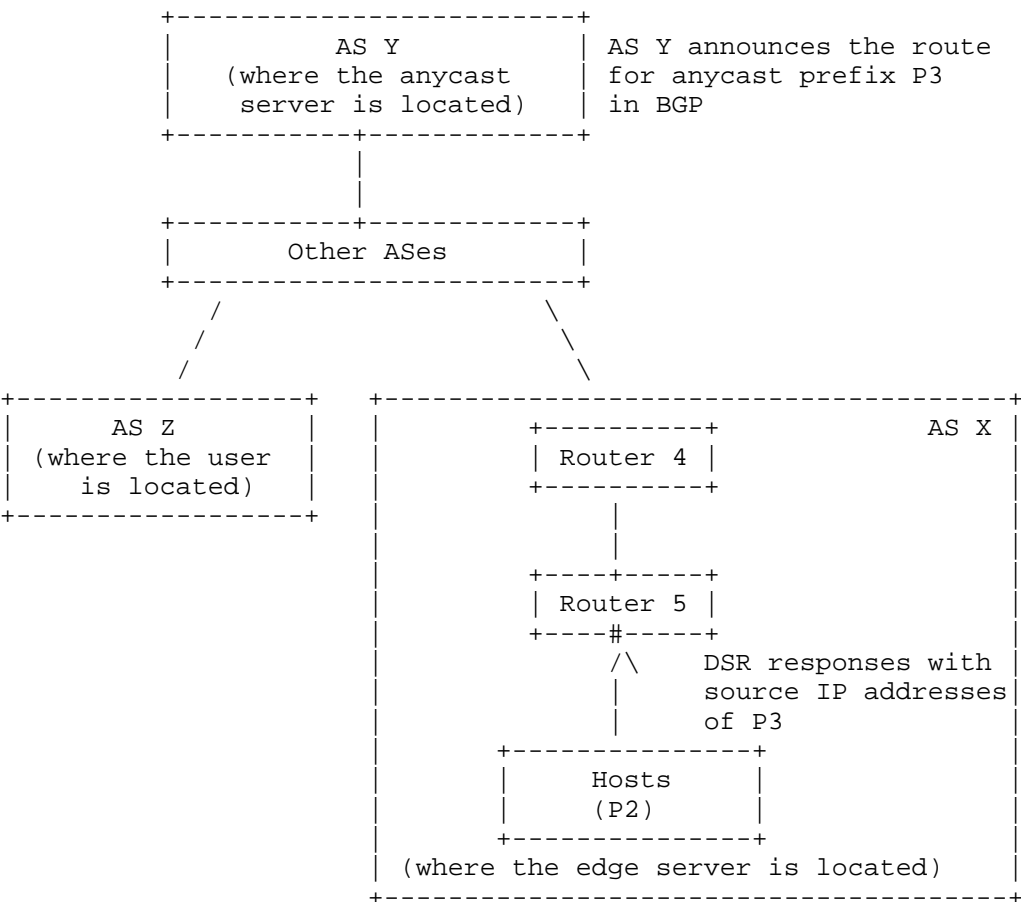
The legitimate traffic originated from sub network with source addresses in 2001:db8:0:100::/56 will be improperly blocked by strict uRPF on Router 1.

Figure 2: An example of asymmetric routing

While the sub network does not expect traffic destined for prefix 2001:db8:0:100::/56 to come from Router 1, it can send traffic with source addresses of prefix 2001:db8:0:100::/56 to Router 1. As a result, there is asymmetric routing of data packets between the sub network and Router 1. Arrows in the figure indicate the flowing direction of traffic. If Router 1 adopts strict uRPF, by checking the FIB entry that matches prefix 2001:db8:0:100::/56, the SAV rule is that Router 1 only accepts data packets with a source address of 2001:db8:0:100::/56 from Router 3. Therefore, when the sub network sends data packets with a source address of 2001:db8:0:100::/56 to Router 1, strict uRPF on Router 1 will improperly block these legitimate data packets. Similarly, if Router 2 adopts strict uRPF, it will improperly block legitimate data packets from the sub network that use a source address of prefix 2001:db8:0::/56.

3.1.2. Hidden Prefix

In the hidden prefix scenario, a host originates data packets using a source address that is hidden or invisible to the intra-domain routing protocol and intra-domain routers. The Content Delivery Networks (CDN) and Direct Server Return (DSR) technology is a representative example of hidden prefix scenario.



DSR response packets from edge server with source IP addresses of P3 (i.e., the anycast prefix) will be improperly blocked by Router 5 if Router 5 uses strict uRPF.

Figure 3: Hidden prefix in CDN and DSR scenario

For example, in Figure 3, the anycast server in AS Y is reachable via prefix P3, while the edge server in AS X is reachable only via prefix P2. When a user in AS Z sends a request to the anycast server in AS Y, the request is received by the anycast server and then tunneled to the edge server in AS X. The edge server subsequently generates a DSR response packet, but uses the anycast server’s address (i.e., a P3 address) as the source. Since Router 5 does not learn the route to prefix P3 from the interface facing the edge server, the response packets from the edge server would be improperly dropped if Router 5 applies strict uRPF for intra-domain SAV.

Specifically, if Router 5 adopts strict uRPF, the SAV rule is that Router 5 only accepts packets with a source address of prefix P2 from the edge server. As a result, when the edge server returns DSR response packets using the source address of prefix P3, DSR response packets will be improperly blocked. In addition, even if Router 5 adopts loose uRPF, it will also improperly block these DSR response packets because prefix P3 does not exist in the FIB of Router 5.

3.2. Intra-domain SAV for Traffic from the Rest of the Internet

Towards Goal ii described in Section 1 and shown in Figure 1, intra-domain SAV is typically deployed on router interfaces facing the rest of the Internet to block data packets using an internal source address (see Figure 4). ACL-based SAV is often used for this purpose. The AS operator can configure an ACL rule which contains a set of unacceptable prefixes (e.g., internal prefixes) to block any data packet using a source address of these prefixes. However, the operational overhead of maintaining ACL rules will be extremely high, especially when there are multiple router interfaces that need to configure the ACL rule as shown in Figure 4.




Figure 4: Intra-domain SAV for Traffic from the Rest of the Internet

In addition, loose uRPF can be used in this case to block data packets from the rest of the Internet using a non-global or non-routed source address. But it will improperly permit spoofed data packets using an internal source address because internal prefixes/addresses exist in the router's local FIB.

4. Problem Statement

As analyzed above, the current operational intra-domain SAV mechanisms have significant limitations in terms of automatic update or accurate validation.

ACL-based SAV entirely relies on manual maintenance and thus requires high operational overhead in dynamic networks. To guarantee accuracy of ACL-based SAV, AS operators have to manually update the ACL rule in time when the prefix or topology changes. Otherwise, improper block or improper permit problems may appear.

Strict uRPF can automatically update SAV rules, but may improperly block legitimate traffic under asymmetric routing scenario or hidden prefix scenario. As analyzed in Section 3.1, it may mistakenly consider a valid incoming interface as invalid, resulting in legitimate data packets being blocked (i.e., improper block problem).

Loose uRPF is also an automated SAV mechanism but its SAV rules are overly loose. As analyzed in Section 3.2, any spoofed data packet using a source address covered by the FIB will be accepted by loose uRPF (i.e., improper permit problem).

In summary, the AS operator has a comprehensive perspective, so it can configure the correct ACL rules. However, manual maintenance can lead to high operational overhead and improper blocks may occur due to the negligence of the AS operator. uRPF cannot guarantee the accuracy of SAV because it solely uses the router's local FIB information to determine SAV rules, which may not match the incoming interfaces of legitimate data packets from the source. As a result, strict uRPF will improperly block legitimate data packets in the case of asymmetric routing and hidden prefix, while loose uRPF has limited effect on preventing source address spoofing because it only blocks non-global or non-routed addresses.

5. Requirements for New SAV Mechanisms

This section lists five general requirements for technical improvements that should be considered when designing the future intra-domain SAV architecture and solution. These informational requirements can not be used to initiate standards-track protocol changes. Any protocol changes would require a standards-track requirements document, not a non-normative reference to this informational document.

5.1. Accurate Validation

The new intra-domain SAV mechanism **MUST** improve the accuracy upon existing intra-domain SAV mechanisms. It **MUST** achieve the two goals described in Section 1 to block those spoofing traffic from sub networks, hosts, and the rest of the Internet. Meanwhile, it **MUST** avoid blocking legitimate data packets, especially when there are prefix changes, asymmetric routes, or hidden prefixes. To overcome the improper block problems, routers may need to use more information (e.g., SAV-specific information) besides the local FIB information to determine SAV decisions. By integrating SAV-specific information, routers may learn asymmetric routes or hidden prefixes, resulting in more accurate SAV rules.

5.2. Automatic Update

The new intra-domain SAV mechanism **MUST** be able to automatically generate and update SAV rules on routers, rather than relying entirely on manual updates like ACL-based SAV. Even if some necessary initial configurations may be needed to improve the accuracy of SAV, automation helps reduce subsequent maintenance overhead of the AS operator.

5.3. Working in Incremental Deployment

The new intra-domain SAV mechanism **MUST** specify the deployment scope (i.e., which routers the mechanism is used on) and **MUST** provide incremental benefits when incrementally deployed within the specified deployment scope. That is, it **MUST NOT** be effective only when fully deployed within the deployment scope. In the incremental deployment scenario, it **MUST** be able to fulfill or partially fulfill the goals described in Section 1 and **MUST** avoid improper blocks.

5.4. Fast Convergence

The new intra-domain SAV mechanism MUST be able to update SAV rules in time when prefix changes, route changes, or topology changes occur in an AS. Two considerations must be taken into account if SAV-specific information is communicated through a protocol. First, the mechanism MUST allow routers to learn the updated SAV-specific information in a timely manner. Second, the mechanism MUST NOT communicate too much SAV-specific information for the SAV function, because this may greatly increase the burden on the control plane of routers and even compromise the performance of the current protocols.

5.5. Security

The new intra-domain SAV mechanisms MUST NOT introduce additional security vulnerabilities or confusion to the existing intra-domain architectures or protocols. Section 6 details the security scope and security considerations for the new intra-domain SAV mechanism.

6. Security Considerations

Similar to the security scope of intra-domain routing protocols, intra-domain SAV mechanisms can ensure integrity and authentication of protocol messages that deliver the required SAV-specific information, and consider avoiding unintentional misconfiguration. It is not necessary to provide protection against compromised or malicious intra-domain routers which poison existing control or management plane protocols. Compromised or malicious intra-domain routers may not only affect SAV, but also disrupt the whole intra-domain routing domain. Security mechanisms to prevent these attacks are beyond the capability of intra-domain SAV.

7. IANA Considerations

This document does not request any IANA allocations.

8. Acknowledgements

Many thanks to the valuable comments from: Jared Mauch, Barry Greene, Fang Gao, Kotikalapudi Sriram, Anthony Somerset, Yuanyuan Zhang, Igor Lubashev, Alvaro Retana, Joel Halpern, Aijun Wang, Michael Richardson, Li Chen, Gert Doering, Mingxing Liu, Libin Liu, John O'Brien, Roland Dobbins, Xiangqing Chang, Tony Przygienda, Yingzhen Qu, Changwang Lin, James Guichard, Linda Dunbar, Robert Sparks, Yu Fu, Stephen Farrel etc.

9. References

9.1. 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/info/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/info/rfc8174>>.

9.2. Informative References

- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, DOI 10.17487/RFC2827, May 2000, <<https://www.rfc-editor.org/info/rfc2827>>.
- [RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, DOI 10.17487/RFC3704, March 2004, <<https://www.rfc-editor.org/info/rfc3704>>.
- [RFC5210] Wu, J., Bi, J., Li, X., Ren, G., Xu, K., and M. Williams, "A Source Address Validation Architecture (SAVA) Testbed and Deployment Experience", RFC 5210, DOI 10.17487/RFC5210, June 2008, <<https://www.rfc-editor.org/info/rfc5210>>.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, DOI 10.17487/RFC4301, December 2005, <<https://www.rfc-editor.org/info/rfc4301>>.
- [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/info/rfc9256>>.
- [RFC2784] Farinacci, D., Li, T., Hanks, S., Meyer, D., and P. Traina, "Generic Routing Encapsulation (GRE)", RFC 2784, DOI 10.17487/RFC2784, March 2000, <<https://www.rfc-editor.org/info/rfc2784>>.
- [RFC8704] Sriram, K., Montgomery, D., and J. Haas, "Enhanced Feasible-Path Unicast Reverse Path Forwarding", BCP 84, RFC 8704, DOI 10.17487/RFC8704, February 2020, <<https://www.rfc-editor.org/info/rfc8704>>.

- [RFC7513] Bi, J., Wu, J., Yao, G., and F. Baker, "Source Address Validation Improvement (SAVI) Solution for DHCP", RFC 7513, DOI 10.17487/RFC7513, May 2015, <<https://www.rfc-editor.org/info/rfc7513>>.
- [cable-verify] "Cable Source-Verify and IP Address Security", January 2021, <<https://www.cisco.com/c/en/us/support/docs/broadband-cable/cable-security/20691-source-verify.html>>.
- [IPSG] "Configuring DHCP Features and IP Source Guard", January 2016, <https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2960/software/release/12-2_53_se/configuration/guide/2960scg/swdhcp82.html>.
- [RFC7039] Wu, J., Bi, J., Bagnulo, M., Baker, F., and C. Vogt, Ed., "Source Address Validation Improvement (SAVI) Framework", RFC 7039, DOI 10.17487/RFC7039, October 2013, <<https://www.rfc-editor.org/info/rfc7039>>.
- [RFC6890] Cotton, M., Vegoda, L., Bonica, R., Ed., and B. Haberman, "Special-Purpose IP Address Registries", BCP 153, RFC 6890, DOI 10.17487/RFC6890, April 2013, <<https://www.rfc-editor.org/info/rfc6890>>.
- [nist-rec] "Resilient Interdomain Traffic Exchange - BGP Security and DDoS Mitigation", January 2019, <<https://www.nist.gov/publications/resilient-interdomain-traffic-exchange-bgp-security-and-ddos-mitigation>>.

Authors' Addresses

Dan Li
Tsinghua University
Beijing
China
Email: tolidan@tsinghua.edu.cn

Jianping Wu
Tsinghua University
Beijing
China
Email: jianping@cernet.edu.cn

Lancheng Qin
Zhongguancun Laboratory
Beijing
China
Email: qinlc@mail.zgclab.edu.cn

Mingqing Huang
Zhongguancun Laboratory
Beijing
China
Email: huangmq@mail.zgclab.edu.cn

Nan Geng
Huawei
Beijing
China
Email: gengnan@huawei.com