SAVNET
Internet-Draft
Intended status: Standards Track
Expires: 2 July 2026

L. Qin
Zhongguancun Laboratory
D. Li
Tsinghua University
L. Chen
L. Liu
Zhongguancun Laboratory

29 December 2025

Bicone Source Address Validation
draft-qin-savnet-bicone-sav-00

Abstract

Source address validation (SAV) aims to avoid improper blocking of legitimate traffic while maintaining directionality. Existing SAV mechanisms commonly rely on ingress allowlist filters on interfaces facing customer or lateral peer Autonomous Systems (ASes), which can result in improper blocking when the allowlist is incomplete. This document analyzes this issue and describes an alternative ingress SAV approach based on a blocklist of prefixes exclusively associated with the provider cone. Network operators may select either approach based on their operational context.

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 2 July 2026.

Copyright Notice

Table of Contents

1. Introduction

Source address spoofing remains one of the most serious security threats to today' s Internet. It is a primary attack vector for large-scale Distributed Denial-of-Service (DDoS) attacks and is widely used in reflective DDoS scenarios. To mitigate source address spoofing, a number of Source Address Validation (SAV) solutions have been proposed, including BCP38 [RFC2827] and BCP84 [RFC3704] [RFC8704]. A fundamental design objective of SAV mechanisms is to minimize improper blocking, that is, blocking legitimate traffic, while preserving directionality, as discussed in [I-D.ietf-savnet-inter-domain-problem-statement] and [RFC8704].

Existing advanced SAV mechanisms, such as EFP-uRPF [RFC8704] and BAR-SAV [I-D.ietf-sidrops-bar-sav], typically construct ingress SAV allowlist filters on interfaces facing customer or lateral peer Autonomous Systems (ASes). These allowlists are generated using information related to the customer cone of the validating AS. Under an allowlist-based approach, an interface permits incoming data packets only if their source addresses are covered by the allowlist. Consequently, the allowlist is required to include all prefixes belonging to the corresponding customer cone. If the allowlist is incomplete, legitimate traffic from the customer cone may be improperly blocked.

To address this limitation, this document explores an alternative SAV approach based on constructing ingress SAV blocklist filters. The proposed blocklist contains prefixes that exclusively belong to the provider cone, referred to as provider-cone-use-only prefixes. Unlike an allowlist, the blocklist is not required to be complete. Instead, it aims to include as many prefixes that can be confidently identified as belonging exclusively to the provider cone as possible. When a blocklist-based SAV mechanism is applied, incoming data packets with source addresses covered by the blocklist are blocked.

In practice, network operators may choose between blocklist-based and allowlist-based SAV mechanisms based on their operational requirements, deployment constraints, and risk considerations.

Readers are encouraged to be familiar with [I-D.ietf-savnet-inter-domain-problem-statement], [RFC8704], [I-D.ietf-sidrops-aspa-profile], [RFC6482], [I-D.ietf-sidrops-aspa-verification], and [I-D.qin-savnet-toa].

1.1. 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.

2. Terminology

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

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

Provider Cone: The set of ASes an AS can reach by using only Customer-to-Provider (C2P) links.

Customer Cone: The set of ASes an AS can reach by using only Provider-to-Customer (P2C) links.

3. Improper Block When the Allowlist is Incomplete

The fundamental idea of existing allowlist-based SAV solutions is to generate an ingress allowlist using information related to the customer cone of a customer or lateral peer AS. Specifically, these mechanisms identify prefixes belonging to the corresponding customer cone and permit only data packets with source addresses drawn from these prefixes on the interface facing that customer or lateral peer AS. This is based on the assumption that data packets received from a customer or lateral peer AS are expected to use source addresses belonging to the customer cone of that AS, unless a route leak occurs [RFC7908].

Limited propagation of prefixes or the presence of hidden prefixes can result in an incomplete allowlist, which may in turn lead to improper blocking (see [I-D.ietf-savnet-inter-domain-problem-statement]).

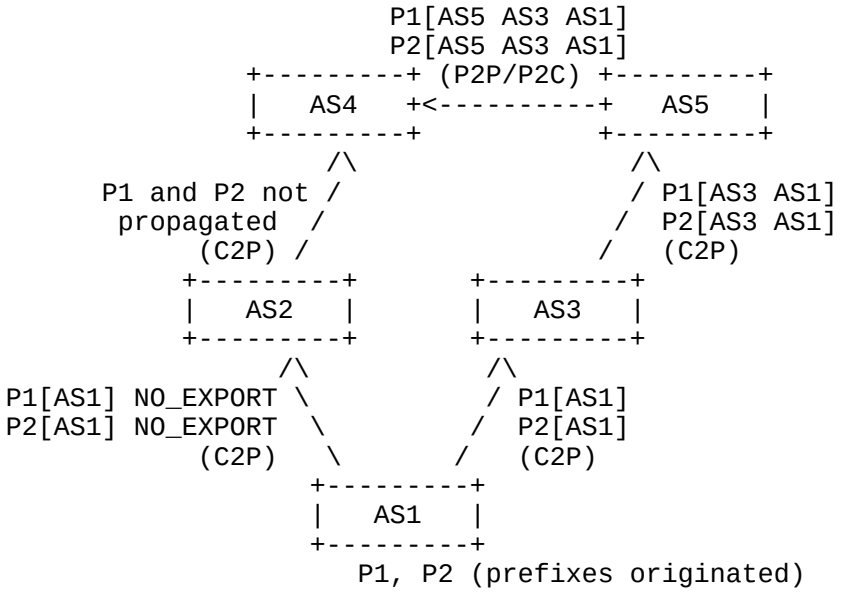

Figure 1: An example of limited propagation of prefixes in the customer cone

Figure 1 illustrates a case of limited prefix propagation within the customer cone of AS4. In the figure, arrows indicate both the propagation direction of BGP announcements and the AS relationships, namely Provider-to-Customer (P2C), Customer-to-Provider (C2P), and Peer-to-Peer (P2P), from the sending AS to the receiving AS. AS1 announces routes for prefixes P1 and P2 to its two provider ASes, AS2 and AS3. However, AS1 attaches NO_EXPORT to the BGP UPDATE message sent to AS2, preventing AS2 from propagating the routes further to AS4. As a result, AS4 receives routes to prefixes P1 and P2 only from its lateral peer or provider AS5. If AS4 applies EFP-uRPF, including Algorithm A or Algorithm B, to generate an allowlist on the AS4AS2 interface, the allowlist will not include prefixes P1 and P2. Consequently, data packets with source addresses in prefixes P1 or P2 will be improperly blocked.

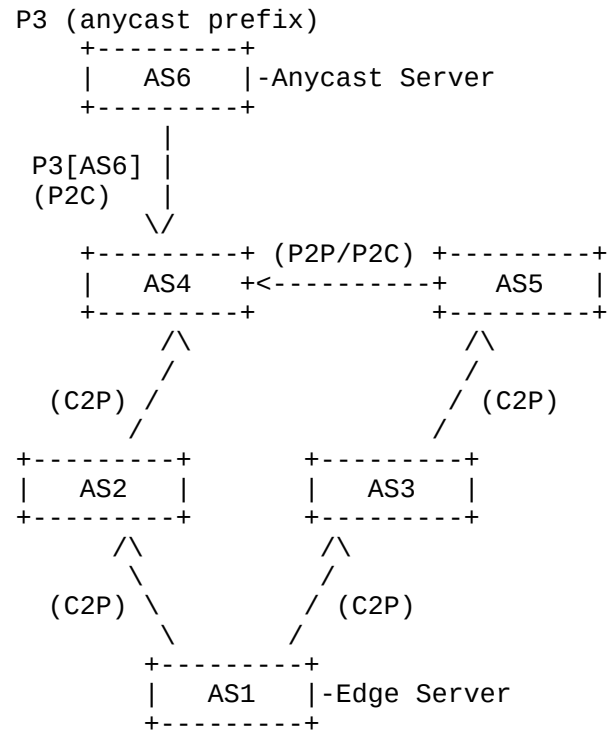

Figure 2: An example of hidden prefixes in the customer cone

Figure 2 shows an example of hidden prefixes in Content Delivery Network (CDN) and Direct Server Return (DSR) scenarios. AS6, where the anycast server is located, announces a route to the anycast prefix P3. Although AS1, where the edge server is located, is not

authorized to originate routes for prefix P3, it legitimately sends data packets using source addresses in prefix P3 as a result of DSR. If AS4 applies an allowlist on the AS4AS2 interface, the allowlist will not include prefix P3. Therefore, the allowlist filter on the AS4AS2 interface will improperly block data packets with source addresses in prefix P3.

More recent SAV mechanisms, such as BAR-SAV [I-D.ietf-sidrops-bar-sav], additionally leverage Autonomous System Provider Authorization (ASPA) [I-D.ietf-sidrops-aspa-profile] and Route Origin Authorization (ROA) [RFC6482] related to the customer cone to construct a more robust allowlist. Traffic Origin Authorization (TOA) [I-D.qin-savnet-toa] can further improve allowlist completeness in hidden-prefix scenarios. However, such authoritative information may be missing due to partial deployment, operational constraints, or incremental adoption. When some ASes or prefixes lack ASPAs, ROAs, or TOAs, the resulting allowlist may still be incomplete.

In summary, due to the inherent complexity of inter-domain routing, SAV mechanisms that rely on allowlist filters on interfaces facing customer or lateral peer ASes may fail to identify all prefixes belonging to the corresponding customer cone. In such cases, an incomplete allowlist can lead to improper blocking of legitimate traffic.

4. Allowlist-based and Blocklist-based SAV Approaches

As discussed in Section 3, allowlist-based SAV mechanisms may improperly block legitimate traffic when authoritative information is missing or incomplete. More generally, ingress SAV mechanisms can be categorized as allowlist-based or blocklist-based, depending on how authoritative information is used in filtering decisions.

An allowlist-based approach permits incoming packets only when their source addresses are explicitly authorized by available authoritative information. When the authoritative information is complete, an allowlist-based approach introduces neither improper blocking nor improper admits. However, when authoritative information is missing, legitimate traffic may be improperly blocked.

A blocklist-based approach, in contrast, permits incoming packets by default and blocks only those packets whose source addresses are explicitly identified as unauthorized by authoritative information. As a result, a blocklist-based approach does not improperly block legitimate traffic when authoritative information is missing, but may result in improper admits if unauthorized source addresses are not covered by the blocklist.

Therefore, when authoritative information is incomplete, allowlist-based and blocklist-based approaches represent different trade-offs. Allowlist-based approaches prioritize minimizing improper permits at the risk of improper blocking, whereas blocklist-based approaches prioritize avoiding improper blocking at the risk of improper permits.

5. Goals of Bicone SAV

Bicone SAV aims to provide robust ingress source address validation on interfaces facing customer or lateral peer ASes by appropriately selecting between allowlist-based and blocklist-based filtering. Its design goals are as follows:

1. Avoiding improper blocks. Bicone SAV prioritizes avoiding the blocking of legitimate data packets received from customer or lateral peer ASes. When authoritative information is insufficient to construct a complete allowlist, Bicone SAV prefers mechanisms that avoid improper blocking.

2. Maintaining directionality. Bicone SAV seeks to preserve directionality in ingress filtering in order to effectively identify source-spoofed data packets. When sufficient authoritative information is available, stricter filtering can be applied to improve spoofing detection.

6. Blocklist Generation

This section describes how to generate a blocklist using BGP UPDATE messages, ASPAs, ROAs, and TOAs associated with the provider cone.

6.1. Key Idea

The provider cone of an AS is defined as the set of ASes that the AS can reach by traversing only Customer-to-Provider (C2P) links. In the absence of route leaks [RFC7908], prefixes associated with ASes in the provider cone are not expected to be used as source addresses in data packets received from customer or lateral peer ASes. Accordingly, the blocklist consists of prefixes that belong to the provider cone.

When the blocklist is applied on an interface facing a customer or lateral peer AS, data packets received on that interface are discarded if their source addresses match any prefix in the blocklist.

To construct such a blocklist, an AS first identifies the ASes in its provider cone using ASPAs and AS-PATH information carried in BGP UPDATE messages. It then identifies prefixes associated with these ASes using ROAs and TOAs [I-D.qin-savnet-toa]. Prefixes that also belong to the AS's customer cone are subsequently removed from consideration.

Given the inherent uncertainty in determining whether a prefix belongs to the customer cone, as discussed in Section 3, a conservative strategy is to retain only prefixes that are exclusively associated with the provider cone. By blocking traffic that uses these prefixes as source addresses, the resulting blocklist-based SAV filter can mitigate the improper blocking issues observed in allowlist-based SAV filters, while still preserving directionality.

6.2. Generation Procedure

A detailed description of blocklist generation procedure is as follows:

1. Create the set of all directly connected Provider ASNs. Call it AS-set Z(1).

2. Create the set of all unique AS_PATHs in Adj-RIBs-In of all interfaces facing Providers.

3. For each unique AS_PATH with N (N>1) ASNs, i.e., [ASN_{1}, ASN_{2}, ..., ASN_{i}, ASN_{i+1}, ..., ASN_{N}] where ASN_{i} is the ith ASN in AS_PATH and the first ASN (i.e., ASN_{1}) is a directly connected Provider ASN. If all unique AS_PATHs have been processed, go to Step 8.

4. Let i = N

5. Decrement i to i-1.

6. If ASN_{i} authorizes ASN_{i+1} as a Provider in ASN_{i}'s ASPA or ASN_{i+1} is a Tier-1 AS, ASNs from ASN_{1} to ASN_{i+1} (i.e., ASN_{1}, ASN_{2}, ..., ASN_{i}, and ASN_{i+1}) are included in AS-set Z(1) and go to Step 3.

7. If i == 1, go to Step 3. Else, go to Step 5.

8. Let k = 1.

9. Increment k to k+1.

10. Create AS-set Z(k) of ASNs that are not in AS-set Z(k-1) but are authorized as Providers in ASPAs of any ASN in AS-set Z(k-1).

11. If AS-set Z(k) is null, then set k_max = k-1 and go to Step 12. Else, form the union of AS-set Z(k) and AS-set Z(k-1) as AS-set Z(k) and go to Step 9.

12. Select all ROAs and TOAs in which the authorized origin ASN is in AS-set Z(k_max). Form the union of the sets of prefixes in the selected ROAs and TOAs. Call it Prefix-set S.

13. For each unique Prefix P in Prefix-set S, check origin ASNs of Prefix P by using all ROAs and TOAs. If all unique Prefixes in Prefix-set S have been processed, go to Step 15.

14. For each prefix of Prefix P and its sub prefixes, if the prefix has at least one origin ASN not in AS-set Z(k_max), remove the prefix from Prefix-set S. Go to Step 13.

15. Apply Prefix-set S as a blocklist on interfaces facing a customer or lateral peer AS.

6.3. Incremental and Partial Deployment of ASPAs

Under incremental and partial deployment of ASPAs, an AS may be unable to fully identify all ASes in its provider cone. As a result, the resulting blocklist may not include all prefixes associated with the provider cone. Nevertheless, an incomplete blocklist does not lead to improper blocking of legitimate traffic. Instead, it can still filter source-spoofed packets whose source addresses fall within the identified subset of provider-cone prefixes. Therefore, even with partial ASPA deployment, the blocklist can provide immediate incremental benefits without introducing additional operational risk.

6.4. Incremental and Partial Deployment of ROAs and TOAs

This document does not use BGP UPDATE messages as a data source for determining the source address space associated with an AS. As discussed in Section 3, BGP information may be incomplete due to limited propagation or hidden prefixes, which can lead to improper blocking when used for SAV filtering. Instead, this document relies on Route Origin Authorizations (ROAs) [RFC6482] and Traffic Origin Authorizations (TOAs) [I-D.qin-savnet-toa] as authoritative information for identifying source address space. Because ROAs and TOAs are explicitly registered by prefix holders and are independent of BGP propagation behavior, they are not affected by the invisible scenarios described in Section 3.

Under incremental and partial deployment, ROAs and TOAs may be missing for some prefixes. If a prefix does not have any corresponding ROA or TOA, it will not be included in the generated blocklist. Consequently, missing ROAs or TOAs do not result in improper blocking of legitimate traffic, although they may reduce the effectiveness of blocking spoofed packets.

This document assumes correct operational practice by prefix holders: once ROAs or TOAs are registered for a prefix, the registration is complete and accurately reflects all legitimate route origins or traffic origins for that prefix.

7. Implementation and Operations Considerations

Network operators may choose to deploy either allowlist or blocklist filters on interfaces facing different customer or lateral peer ASes, depending on their operational requirements and deployment conditions.

7.1. Meeting the Goals

Avoiding improper blocking is a primary operational objective, as discarding legitimate traffic can cause severe service disruption. Subject to this constraint, SAV mechanisms should also minimize improper admits to improve protection against source address spoofing.

When an allowlist deployed on an interface is known to be complete, it introduces neither improper blocks nor improper admits. However, if allowlist completeness cannot be reliably ensured, for example due to hidden prefixes in the customer cone or missing authoritative information, deploying an allowlist may lead to improper blocking of legitimate traffic and thus fail to meet the primary objective. In contrast, a blocklist does not result in improper blocking caused by missing authoritative information, although it may allow some spoofed traffic when incomplete.

Accordingly, when an allowlist on an interface is known to be complete, network operators are advised to use the allowlist. Otherwise, deploying a blocklist is recommended to avoid potential improper blocking. For small ISPs with relatively small customer cones, determining allowlist completeness is generally easier, as fewer ASes are involved and routing relationships tend to be simpler. For example, operators may directly confirm with a customer or lateral peer AS whether all ASes in its customer cone have deployed the required authoritative information. In contrast, for large ISPs with extensive customer cones, determining allowlist completeness is significantly more challenging. In such cases, if the completeness of the allowlist cannot be reliably determined, deploying a blocklist is recommended.

7.2. Storage Overhead

Deploying allowlist or blocklist filters requires additional memory resources, such as ternary content-addressable memory (TCAM), on line cards. Network operators should therefore consider storage overhead when selecting between allowlists and blocklists. In general, a small ISP tends to generate a smaller allowlist and a larger blocklist, while a large ISP tends to generate a larger allowlist and a smaller blocklist.

One approach to reducing memory consumption is to maintain the original list in the control plane and install only an aggregated list in the data plane. For example, if the original list contains prefixes P1 and P2, and P1 is a less-specific prefix of P2, only P1 needs to be installed in the data plane.

7.3. Implementation and Operations Recommendations

For an interface facing a customer or lateral peer AS, the following operational guidance applies:

1. If the network operator can determine that the allowlist covers all prefixes of the corresponding customer cone, deploying an allowlist on the interface is recommended, as a complete allowlist introduces neither improper blocks nor improper permits.

2. If the network operator cannot reliably determine the completeness of the allowlist, deploying a blocklist is recommended in order to avoid improper blocking. In such cases, operators are encouraged to consider using the blocklist in conjunction with Loose uRPF [RFC3704] to improve spoofed traffic mitigation. Loose uRPF can be used to filter packets with unallocated or unroutable source addresses, while the blocklist focuses on filtering packets whose source addresses are associated with the provider cone.

Network operators may further refine the blocklist based on local knowledge. For example, operators may add special-purpose prefixes that are not expected to be used as source addresses in data packets, such as those listed in the IANA IPv4 Special-Purpose Address Registry [IANA].

8. Security Considerations

The security considerations described in [RFC8704], [I-D.ietf-sidrops-bar-sav], [I-D.ietf-sidrops-aspa-profile], [RFC6482], and [I-D.ietf-sidrops-aspa-verification] also applies to this document.

9. IANA Considerations

This document has no IANA requirements.

10. Acknowledgements

The authors would like to thank Ben Maddison, Kotikalapudi Sriram, Nan Geng, Aijun Wang, Shengnan Yue, Siyuan Teng, Igor Lubashev, Job Snijders, and many other members of the SIDROPS and SAVNET working groups for comments and discussion.

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/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>.

11.2. 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>.

[I-D.ietf-savnet-inter-domain-problem-statement] Li, D., Qin, L., Liu, L., Huang, M., and K. Sriram, "Gap Analysis, Problem Statement, and Requirements for Inter-Domain SAV", Work in Progress, Internet-Draft, draft-ietf-savnet-inter-domain-problem-statement-12, 20 October 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-savnet-inter-domain-problem-statement-12>.

[I-D.ietf-sidrops-bar-sav] Sriram, K., Lubashev, I., and D. Montgomery, "Source Address Validation Using BGP UPDATEs, ASPA, and ROA (BAR-SAV)", Work in Progress, Internet-Draft, draft-ietf-sidrops-bar-sav-08, 20 October 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-sidrops-bar-sav-08>.

[I-D.qin-savnet-toa] Qin, L., Maddison, B., Li, D., and I. Lubashev, "A Profile for Traffic Origin Authorizations (TOAs)", Work in Progress, Internet-Draft, draft-qin-savnet-toa-00, 3 November 2025, <https://datatracker.ietf.org/doc/html/draft-qin-savnet-toa-00>.

[RFC7908] Sriram, K., Montgomery, D., McPherson, D., Osterweil, E., and B. Dickson, "Problem Definition and Classification of BGP Route Leaks", RFC 7908, DOI 10.17487/RFC7908, June 2016, <https://www.rfc-editor.org/info/rfc7908>.

[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>.

[I-D.ietf-sidrops-aspa-profile]
Azimov, A., Uskov, E., Bush, R., Snijders, J., Housley, R., and B. Maddison, "A Profile for Autonomous System Provider Authorization", Work in Progress, Internet-Draft, draft-ietf-sidrops-aspa-profile-20, 18 August 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-sidrops-aspa-profile-20>.

[RFC6482] Lepinski, M., Kent, S., and D. Kong, "A Profile for Route Origin Authorizations (ROAs)", RFC 6482, DOI 10.17487/RFC6482, February 2012, <https://www.rfc-editor.org/info/rfc6482>.

[I-D.ietf-sidrops-aspa-verification]
Azimov, A., Bogomazov, E., Bush, R., Patel, K., Snijders, J., and K. Sriram, "BGP AS_PATH Verification Based on Autonomous System Provider Authorization (ASPA) Objects", Work in Progress, Internet-Draft, draft-ietf-sidrops-aspa-verification-24, 19 October 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-sidrops-aspa-verification-24>.

[IANA] "IANA IPv4 Special-Purpose Address Registry", n.d., <https://www.iana.org/assignments/iana-ipv4-special-registry/iana-ipv4-special-registry.xhtml>.

Authors' Addresses

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

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

Li Chen
Zhongguancun Laboratory
Beijing
China
Email: lichen@zgclab.edu.cn

Libin Liu
Zhongguancun Laboratory
Beijing
China
Email: liulb@zgclab.edu.cn