savnet
Internet-Draft
Intended status: Informational
Expires: 26 November 2026

M. Huang
Zhongguancun Laboratory
N. Geng
Huawei Technologies
D. Li
Tsinghua University

25 May 2026

Provider Interface SAV for Customer Cone Sources
draft-huang-savnet-pi-sav-for-cc-02

Abstract

Current source address validation (SAV) on an AS's provider interfaces mainly relies on loose uRPF, which cannot defend against IP spoofing using routable prefixes and is ineffective when a default route is present. This document describes a *framework* for provider interface source address validation against spoofing of customer cone source prefixes (PI-SAV for CC). Two architectural approaches are presented:

* *PI-SAV for Standalone CC*: a static solution that builds a blocklist based on topology information from BGP RIBs, RPKI ASPA, and local configuration (SLURM) etc. It requires no inter-AS communication.

* *PI-SAV for Standalone+ CC*: an enhanced solution that uses lightweight query-response coordination between the top AS and member ASes to identify sub-cones that do not actually cause traffic detours, thereby enlarging the effective blocklist.

This document is *informational* and provides only the framework, conceptual procedures, and requirements for future protocol specifications. Detailed wire protocols (e.g., for partial transit information provisioning or for query-response messaging) are out of scope and will be defined in separate documents.

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/.

Copyright Notice

Table of Contents

1. Introduction

As described in [RFC3704], [RFC8704], and [I-D.ietf-savnet-inter-domain-problem-statement], current source address validation on AS provider interfaces is mainly based on the loose uRPF mechanism. This mechanism can only check whether the source address is a valid internet prefix but cannot provide protection for routable prefixes spoofing. Furthermore, if the AS has a default route (many stub ASes configure default routes to reduce BGP processing and FIB table pressure), the protective effect of loose uRPF on routable prefixes will also be nullified.

Meanwhile, normally the vast majority of DDoS attack traffic flows into the local network through the provider interface side. Studies further indicate that reflective amplification attacks based on source address spoofing are among the most prevalent types of large-scale DDoS attacks.

The solutions described in this document, based on the customer cone (CC) source address information and its routing characteristics, generate the provider interface based source prefix blocklist for the customer cone prefixes (IBB-mode in [I-D.ietf-savnet-general-sav-capabilities], or IBA-mode combined with loose uRPF if applicable). The design goal is to prevent spoofing packets with the customer cone prefixes from flowing into the local networks, which can reduce the threat of reflective amplification DDoS attacks that leverage local servers to target local users.

This document is intentionally limited to an *informational framework*. It describes the solution architecture, operational assumptions, and requirements for both Standalone CC and Standalone+ CC. The actual protocols for obtaining partial transit information (Section 2.2) and for query-response coordination (Section 3.3) are left for future specifications. This approach allows the community to review and agree on the overall design before committing to protocol details.

1.1. Terminology

Customer Cone (CC): a set of ASes that includes an AS and its direct/indirect customer ASes, including the customer AS’s customer ASes and so on, till the stub ASes.

CC-Top-AS: the top AS for the customer cone, which means all other ASes in the CC are its direct/indirect customer ASes.

CC-Top-ASBR: ASBR (AS Border Router) with the provider interface(s) of the CC-Top-AS, which is the deployment position for the solutions in this document.

CC-Member-AS: an AS in the customer cone under the CC-Top-AS.

Partial Transit: a special transit provider service between two eBGP neighbors, a less common but operationally significant scenario. When an AS (AS-a) receives BGP update messages from its customer AS (AS-b), it disseminates the messages only to its peer AS(es) and other customer AS(es), not to its upper transit providers as usual. In this case, AS-a is partial transit provider for AS-b. Partial transit can make the customer AS (and its sub-CC) invisible for the upper transit AS.

CC-sub-Transit-AS: a CC-Member-AS that has alternative transit provider(s) outside the CC.

Sub-Transit-CC: the customer cone of a CC-sub-Transit-AS (including the CC-sub-Transit-AS and its direct/indirect customer ASes). According valley-free policy of BGP routing, traffic from the Sub-Transit-CC to other CC-Member-AS may go through the alternative transit link and provider interfaces of CC-Top-ASBR.

Standalone CC (SCC): a subset of CC, all ASes within a Standalone CC share same transit provider(s) on the CC-Top-AS (i.e. no member AS of a Standalone CC has an alternative transit provider outside the CC), all Sub-Transit-CCs in the CC must be excluded. Based on the valley-free principle of BGP routing, traffic between Standalone CC hosts will not go through the CC-Top-AS’s provider interfaces.

Standalone+ CC (SPCC): a superset of a Standalone CC that includes the Sub-Transit-CC(s) for which no detour actually occurs for intra-CC traffic. In other words, traffic between Standalone+ CC hosts will not go through the provider interfaces of the CC-Top-AS. Note that the range of Standalone+ CC may change dynamically due to traffic engineering configuration or transit link failure.

Detour Path: in the context of this document, a detour path refers to a route taken by intra-CC traffic that enters the CC via a provider interface of the CC-Top-AS (i.e., the traffic flows from an external alternative provider into the CC). Such a path makes the corresponding source prefix unsuitable for the PI-SAV blocklist, because legitimate packets with that source prefix may enter the CC through that interface.

Independence Degree: a metric that evaluates the degree of independence of a Standalone CC or Standalone+ CC within its customer cone, defined as the percentage of IP prefixes inside the CC that qualify as belonging to the Standalone CC or Standalone+ CC. A higher Independence Degree indicates that a larger fraction of the customer cone's prefixes can be protected from spoofing via the provider interface blocklist.

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. Basic Solution: PI-SAV for Standalone CC

2.1. General Procedure for SCC Solution

The PI-SAV for Standalone CC solution focuses on constructing the Standalone CC by deducting Sub-Transit-CCs from the whole CC without omission, and then generating a source prefix blocklist for the Standalone CC (after subtracting prefixes with multiple origin ASes).

Standalone CC construction requires complete knowledge of the CC-Member-ASes and whether each of them has an alternative transit provider outside the CC. The general procedure for a CC-Top-ASBR to generate the PI-SAV blocklist for Standalone CC is as follows:

1. *Initial CC-Member-AS collection*: includes (a) ASNs for all direct customer ASes of the CC-Top-AS (obtained from system configuration), and (b) partial transit deployment information for the target CC (using SLURM-based local management as suggested).

2. *CC construction*: based on the initial ASNs, find all CC-Member-ASes by searching the local-RIBs for ASes that appear under the initial ASN in the AS-PATHs. The transit relationship between member ASes must be saved for later Sub-Transit-CC construction.

3. *Sub-Transit-CC construction*: check whether any alternative transit provider outside the CC exists for each CC-Member-AS. If such a provider exists, mark the AS as a CC-sub-Transit-AS. This can be done using RPKI ASPA records [I-D.ietf-sidrops-aspa-profile], and local SLURM-based support for additional information is encouraged. If the transit provider information is not available for a CC-Member-AS, the AS MUST be marked as a CC-sub-Transit-AS. Then construct the Sub-Transit-CC for each CC-sub-Transit-AS.

4. *Standalone CC construction*: deduct all the Sub-Transit-CCs from the original CC to form the Standalone CC.

5. *Initial prefix list generation*: based on local-RIBs, collect all prefixes originated from the member ASes of the Standalone CC.

6. *Final prefix blocklist generation*: check whether any prefix in the initial list has multiple origin ASes (MOAS). If a prefix is originated not only by AS(es) inside the CC but also by AS(es) outside the CC, deduct it from the initial list. MOAS prefix identification can be based on RPKI ROA [RFC9582] and/or local ROA SLURM records. BGP adj-RIBs-in can also be used to identify external origins.

2.2. Requirements for SCC Solution

The Standalone CC solution relies on three distinct types of information to correctly construct the blocklist. Each is described below as a set of requirements, independent of any particular implementation or external database (e.g., RPKI ASPA). This separation ensures that the framework remains implementation-agnostic and can be satisfied by different mechanisms (e.g., local configuration, future protocols, or a combination).

2.2.1. Complete Member AS Discovery (Hidden ASes)

The CC-Top-AS must know the full set of CC-Member-ASes, including those that may be invisible from local-RIBs due to partial transit relationships or No-Export BGP communities.

Requirements:

* *Observability gap*: BGP local-RIBs alone is insufficient to guarantee a complete view. Hidden ASes (child ASes that do not advertise their prefixes upward) cannot be automatically inferred.

* *Explicit provisioning*: If a CC-Member-AS engages in partial transit or otherwise hides its prefixes from the CC-Top-AS, the operator of the CC-Top-AS MUST explicitly supply this information, e.g., via local SLURM [RFC8416] entries or a future declarative protocol. The CC-Top-AS cannot "assume" a hidden subtree; it can only act on explicitly known information.

* *Future specification*: A future specification document MAY define a mechanism (e.g., an extension to RPKI ASPA or a BGP attribute) for an AS to voluntarily declare its partial transit relationships, thereby making hidden ASes discoverable. Until such a mechanism exists, operators must rely on local configuration.

* *Handling missing information*: If the CC-Top-AS lacks information about whether a particular customer AS has hidden downstream ASes, it is impossible to make a safe assumption for Standalone CC - the construction of Sub-Transit-CC related to the hidden AS might be omitted. Therefore, operators MUST ensure that no partial transit information is missing, either via local configuration or a future specification; otherwise, they MUST use the Standalone+ CC solution (Section 3), which can dynamically discover hidden ASes through the diffusion-based query mechanism (R7 in Section 3.3). This requirement ensures that SCC is deployed only where topology information is complete, avoiding unsafe default assumptions.

2.2.2. Detection of Alternative Transit Providers

For each CC-Member-AS, the CC-Top-AS must determine whether that AS has any transit provider outside the CC. Such ASes are marked as CC-sub-Transit-AS, and their entire sub-cone is excluded from the Standalone CC.

Requirements:

* *Provider set completeness*: The CC-Top-AS MUST obtain the set of all transit providers for each CC-Member-AS. This can be done using RPKI ASPA records [I-D.ietf-sidrops-aspa-profile], local SLURM, or other means.

* *Safe default - assume external provider*: If the provider information for a CC-Member-AS is incomplete or unavailable (e.g., no ASPA record and no local override), the CC-Top-AS MUST assume that the AS has an external transit provider and therefore MUST mark it as a CC-sub-Transit-AS. This ensures that the blocklist never incorrectly includes prefixes that might have a feasible detour path.

2.2.3. MOAS Prefix Handling

A prefix originated by multiple ASes (MOAS) MUST be excluded from the blocklist only if at least one of its origin ASes lies outside the Standalone CC. If all origin ASes are inside the Standalone CC, the prefix MAY remain in the blocklist (internal MOAS does not create a detour path).

Detecting external origins is challenging due to two inherent limitations:

* *Best-path-only visibility*: BGP best path selection hides alternate origin ASes from the Loc-RIBs (Adj-RIBs-in may also be missing).

* *Adj-RIBs-in incompleteness*: Memory-constrained routers may not retain a full Adj-RIBs-in.

A practical deployment SHOULD combine the following information sources:

* *RPKI ROA* - Preferred when valid records exist.

* *Local SLURM* - Operator-provided overrides (especially for critical prefixes).

* *BGP Adj-RIBs-in* - Usable as a supplementary signal, with recognition that it may be incomplete.

If external origin information is incomplete, the operator MAY choose a local policy:

* *Conservative (RECOMMENDED)*: Exclude the prefix from the blocklist unless it can be confidently determined that no external origin exists. This avoids false blocking at the cost of reduced protection.

* *Optimistic*: Include the prefix in the blocklist unless an external origin is positively confirmed. This maximizes protection but risks blocking legitimate traffic if an unknown external origin exists.

Operators should weigh the impact of false blocking (denial of service for legitimate customers) versus under-protection (spoofing risk). Feedback from deployment (see Section 4.5) can help refine the information over time.

3. Improved Solution: PI-SAV for Standalone+ CC

3.1. Motivations and Operational Assumptions

Although an alternative transit link may cause traffic between CC member ASes to detour through provider interfaces of the CC-Top-AS, such detours normally only happen under specific routing policy configurations or internal transit path failures. In many cases, Sub-Transit-CCs do not actually use their external providers for intra-CC traffic.

Standalone+ CC relaxes the strict exclusion of all Sub-Transit-CCs. Instead, it verifies whether detours actually occur, and includes those Sub-Transit-CCs (or even finer-grained prefixes) that have no detour. This can significantly improve the prefix independence ratio, especially for deep customer cones.

The Standalone+ CC solution assumes a communication mechanism between the CC-Top-AS and CC-Member-ASes. This document only describes the conceptual model and the requirements for such a communication protocol; the actual protocol design is left for future Standards Track documents.

3.2. Query-Response Coordination Model

In the conceptual model, the CC-Top-AS sends *queries* to member ASes to determine whether detour paths actually exist. Two communication approaches are possible:

* *Targeted queries*: The CC-Top-AS first constructs the CC topology (using BGP RIB and possibly RPKI/ASPA) to identify which member ASes have or may have alternative transit providers. It then sends queries only to those CC-sub-Transit-ASes. This approach requires pre-existing topology knowledge and connectivity information for the targeted ASes.

* *Diffusion-based queries (recommended)*: The CC-Top-AS injects the query into BGP (e.g., as a new BGP attribute or NLRI) and relies on existing BGP propagation to deliver the query to *all* CC-Member-ASes. This approach requires no pre-computed topology and naturally reaches all ASes, including those that would otherwise be hidden due to partial transit. Moreover, it allows member ASes to respond not only with "no detour" indications but also with auxiliary information (e.g., hidden AS, their provider set, MOAS origins), thereby *replacing or complementing out-of-band registries* (RPKI ASPA, ROA and SLURM etc.) for Standalone CC.

The *diffusion-based model* is the recommended baseline for Standalone+ CC, because it:

* Eliminates the need for pre-computed CC topology before queries,

* Reaches all member ASes without requiring their explicit addresses,

* Enables incremental deployment (ASes that ignore the BGP attribute simply do not respond; the CC-Top-AS then assumes a detour),

* Provides a unified channel for both detour detection and topology/prefix information collection.

Two operational modes are defined, regardless of the underlying communication model:

* *Conservative mode*: The CC-Top-AS starts with the Standalone CC as the initial blocklist. It then sends a query to member ASes (via the chosen mechanism). If a member AS responds that no detour exists for its sub-cone (or for specific prefixes), the CC-Top-AS adds the corresponding prefixes to the Standalone+ CC blocklist. Non-responsive ASes are assumed to have a detour.

* *Aggressive mode*: The CC-Top-AS starts with the whole CC as the initial blocklist. It then sends a query. If a member AS responds that a detour does exist, the CC-Top-AS removes the corresponding prefixes from the blocklist. Non-responsive ASes are also treated as having a detour after a timeout. This mode works best when most member ASes support the protocol.

The detailed timing, retransmission, and state machines depend on the chosen communication model and are left for future specification. The remainder of this document assumes the diffusion-based model as the reference; the targeted model is considered an optional simplification (e.g., for small-scale deployments).

3.3. Requirements for a Future Communication Protocol

The SPCC solution requires a communication mechanism between the CC-Top-AS and CC-Member-ASes to exchange queries and responses. This document does not define a protocol; instead, it lists the mandatory requirements that any future Standards Track protocol MUST satisfy.

R1 - Query dissemination: The query MUST be deliverable to all CC-Member-ASes without requiring the CC-Top-AS to pre-compute their identities or addresses. A diffusion mechanism that leverages existing BGP propagation (e.g., a new BGP attribute or NLRI) is

RECOMMENDED. The query MUST include at least: - A unique query identifier, - The mode (conservative or aggressive), - A timeout value (default recommended 30 seconds), - Optionally, a list of specific ASes or prefixes to be checked.

R2 - Response collection: A CC-Member-AS that receives a query MUST be able to send a response back to the CC-Top-AS. The response MUST include: - The same query identifier, - For each checked item, an indication of whether a detour path is observed (or, in fine-grained operation, a list of prefixes with no detour).

R3 - Transport reliability: The response channel MUST provide reliable delivery (e.g., TCP). The query channel MAY use an unreliable or reliable transport; if unreliable, the protocol MUST define retransmission or fallback behavior.

R4 - Security: All communication MUST be authenticated and encrypted to prevent forgery of "no detour" responses. Use of TLS with router certificates is RECOMMENDED.

R5 - Timeout and fallback: If the CC-Top-AS does not receive a response within the timeout, it MUST assume that a detour exists for the corresponding ASes or prefixes (i.e., they are excluded from the Standalone+ CC blocklist).

R6 - Incremental updates: To avoid obsolete blocklist entries that could cause false blocking, a CC-Member-AS MUST send unsolicited update messages to the CC-Top-AS when its routing state changes (e.g., a previously direct path becomes a detour path, or a detour path reverts to direct). The CC-Top-AS MUST NOT rely on periodic polling.

R7 - Auxiliary information collection: In addition to detour indications, the protocol MAY allow a CC-Member-AS to include in its response other information that is useful for Standalone CC, such as: - Its complete set of transit providers (including those outside the CC), - Prefixes that are originated by it but hidden from the CC-Top-AS (e.g., due to No-Export), - MOAS prefixes and their origin ASes. This capability allows Standalone+ CC to gradually *replace* out-of-band registries (Partial Transit, RPKI ROA/MOAS, RPKI ASPA, SLURM etc.) for the purpose of constructing the initial Standalone CC blocklist, thereby reducing deployment dependencies.

The detailed design of message formats, state machines, and IANA registrations is left for future specification.

4. Deployment and Management Considerations

4.1. Data Completeness Requirements

Deploying Standalone CC requires different levels of completeness for different types of information:

* *Hidden ASes (complete subtree missing)* - This is the most critical gap. If a whole AS (and its downstream) is invisible due to partial transit, the operator MUST supply that information (e.g., via SLURM) or use Standalone+ CC instead. Without it, the blocklist may incorrectly retain prefixes that actually have a detour path, causing false blocking.

* *Provider set (alternative transit providers)* - For any AS where the operator provides explicit provider information (e.g., via SLURM or ASPA), that information MUST be accurate and complete for that AS. However, the operator is NOT required to provide information for every AS; for ASes without explicit data, the safe default (assume external provider) applies, which avoids false blocking at the cost of reduced protection.

* *Prefix list (originated prefixes)* - The CC-Top-AS collects prefixes from local-RIBs. Missing prefixes (e.g., due to No-Export) simply reduce protection (the blocklist is smaller) and do NOT cause false blocking. No completeness guarantee is required. However, when a hidden AS is discovered (by explicit provisioning or via Standalone+ CC), operators are encouraged to also retrieve the complete prefix list of that AS to improve protection coverage.

* *MOAS external origins* - As discussed in Section 2.2.3, incomplete external origin information may lead to false blocking if an optimistic policy is used, or under-protection if a conservative policy is used. Operators should choose a policy that matches their risk tolerance and use the feedback mechanism (Section 4.5) to refine their information.

This graduated approach allows operators to deploy Standalone CC with manageable assumptions, understanding where gaps are safe and where they are not.

4.2. Incremental Deployment and Blind Spots

The solutions support incremental deployment within a region. Not all member ASes need to support Standalone+ CC communication. For those that do not, the CC-Top-AS MUST assume that detours exist (i.e., exclude them from Standalone+ CC). Over time, as more ASes adopt the protocol, the blocklist can expand.

For Standalone CC, blind spots (e.g., hidden ASes due to partial transit) can be addressed by local SLURM configuration. Operators should prioritize configuring SLURM entries for customer ASes that are known to have complex routing policies.

4.3. Independence Degree as a Metric

The Prefix Independence Degree defined in Section 1.1 can help operators evaluate the potential benefit of PI-SAV for CC before deployment. For example: - If the Independence Degree for Standalone CC is high (e.g., >80%), the static solution already provides good coverage. - If it is low for Standalone CC but high for Standalone+ CC (e.g., >80%), deploying Standalone+ CC (with the communication protocol) is beneficial. - If both are low, the customer cone may be too intertwined with external providers, and PI-SAV for CC may offer limited protection.

These thresholds are illustrative; operators should determine their own based on network topology and risk tolerance.

4.4. Feedback-Driven Refinement

The Standalone CC and Standalone+ CC solutions rely on static or dynamically collected topology information. Incomplete information (e.g., hidden ASes due to partial transit, unknown external providers, or MOAS origins) can lead to suboptimal blocklists - either under-protection or false blocking.

To address this, an operator MAY deploy a feedback loop that uses *observed SAV validation results* to refine the information base. This applies to all types of information gaps discussed in Section 2 (hidden ASes, alternative providers, MOAS), as well as incomplete ASPA or SLURM entries.

Example feedback loop:

* *Detection*: The ASBR logs dropped packets (sampled) and their source prefixes that are blocked by the PI-SAV blocklist. If a prefix is frequently dropped but the operator later determines that the traffic is legitimate (e.g., through customer reports or traffic analysis), this indicates a possible false positive.

* *Investigation*: The operator investigates the topology around that prefix. The investigation may reveal previously missing information, such as an undisclosed external transit provider, a hidden AS, or an unknown MOAS origin.

* *Refinement*: The operator updates the local SLURM or, if using Standalone+ CC, triggers a new query to collect fresh information. The updated information is used to regenerate a more accurate blocklist.

Conversely, if spoofed packets with a CC source prefix are observed entering through a provider interface, that indicates the corresponding prefix should have been excluded from the blocklist. The operator can then similarly investigate and refine the information base.

The feedback loop can be semi-automated:

* The ASBR exports sampled drop logs to an offline collector.

* An analytics process correlates drops with known topology and flags anomalous patterns (e.g., a prefix that experiences a high drop rate but has no known external provider).

* Alerts are raised for operator review and potential SLURM updates.

Over time, this iterative process helps the system converge to a more accurate blocklist, compensating for initial information gaps.

5. Future Work: Protocol Specification Roadmap

This document intentionally does not define wire protocols for:

* Provisioning of partial transit information (see Section 2.2.1),

* Query-response coordination for Standalone+ CC (see Section 3.3).

The authors intend to collaborate with the SAVNET WG to produce one or more Standards Track documents that specify these protocols. Possible directions include:

* Extending RPKI ASPA to carry partial transit declarations,

* Defining a BGP attribute or a new NLRI for query dissemination (the diffusion-based model described in Section 3.2),

* Designing a lightweight TCP-based protocol with TLS for responses.

Among the possible communication models, the diffusion-based (BGP-integrated) approach is preferred for its scalability and ability to also collect topology information. Future specifications may also define a targeted query mode as a lightweight alternative.

These future specifications will update or extend the framework described in this document. Implementers are encouraged to experiment with the framework and provide feedback.

6. Security Considerations

The security considerations for PI-SAV for CC mainly relate to the communication mechanism for Standalone+ CC. As required in Section 3.3, any future protocol MUST provide authentication and encryption to prevent an attacker from forging "no detour" responses, which would cause the CC-Top-AS to incorrectly include a Sub-Transit-CC in the blocklist, leading to denial of service for legitimate traffic.

Additionally, the source prefixes used in response messages should not be blocked by the PI-SAV blocklist itself; otherwise, responses may be dropped. This can be avoided by ensuring that the response channel uses source addresses that are not in the blocklist (e.g., addresses from a dedicated management network).

For Standalone CC, security considerations are limited to the correctness of the input data (BGP RIB, RPKI, SLURM). Incorrect or incomplete data may result in a blocklist that either misses spoofed packets (under-protection) or blocks legitimate traffic (over-protection). Operators should validate their data sources and use the feedback mechanism (Section 4.5) to detect and correct information gaps.

7. IANA Considerations

This document includes no request to IANA.

8. References

8.1. Normative References

[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>.

[RFC8416] Ma, D., Mandelberg, D., and T. Bruijnzeels, "Simplified Local Internet Number Resource Management with the RPKI (SLURM)", RFC 8416, DOI 10.17487/RFC8416, August 2018, <https://www.rfc-editor.org/info/rfc8416>.

[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>.

8.2. Informative References

[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>.

[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>.

[RFC9582] Snijders, J., Maddison, B., Lepinski, M., Kong, D., and S. Kent, "A Profile for Route Origin Authorizations (ROAs)", RFC 9582, DOI 10.17487/RFC9582, May 2024, <https://www.rfc-editor.org/info/rfc9582>.

[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-16, 7 April 2026, <https://datatracker.ietf.org/doc/html/draft-ietf-savnet-inter-domain-problem-statement-16>.

[I-D.ietf-savnet-general-sav-capabilities] Huang, M., Cheng, W., Li, D., Geng, N., and L. Chen, "General Source Address Validation Capabilities", Work in Progress, Internet-Draft, draft-ietf-savnet-general-sav-capabilities-02, 10 October 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-savnet-general-sav-capabilities-02>.

[I-D.ietf-sidrops-aspa-profile] Snijders, J., Azimov, A., Uskov, E., Bush, R., Housley, R., and B. Maddison, "A Profile for Autonomous System Provider Authorization", Work in Progress, Internet-Draft, draft-ietf-sidrops-aspa-profile-26, 19 April 2026, <https://datatracker.ietf.org/doc/html/draft-ietf-sidrops-aspa-profile-26>.

Authors' Addresses

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

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

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