savnet
Internet-Draft
Intended status: Informational
Expires: 15 August 2026

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

11 February 2026

Provider Interface SAV for Customer Cone Sources
draft-huang-savnet-pi-sav-for-cc-01

Abstract

Current source address validation (SAV) on AS’s provider interfaces primarily relies on loose uRPF, which can protect against IP spoofing based on unannounced internet prefixes but lacks the ability to defend against spoofing based on valid internet prefixes. Furthermore, if a default route is present, this protection is effectively nullified.

Based on the traffic flow characteristics between the ASes in a customer cone, this document describes a general solution architecture -- provider interface SAV for customer cone source spoofing (PI-SAV for CC). This mechanism can be applied on the provider interfaces of an AS to prevent spoofing traffic with the customer cone source addresses from flowing into the local AS and its customer cone. Specially, this mechanism offers protection against reflective amplification DDoS attacks that leverage local servers to target victims within the customer cone.

Table of Contents

1. Introduction

As described in [RFC3704][RFC8704][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, generates 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.

In contrast to loose uRPF, the key advancement of the PI-SAV for CC mechanism is its ability to create a precise blocklist for _routable_ source prefixes belonging to the local customer cone, thereby preventing spoofing even when a default route is used.

Section 4 introduces the basic solution: PI-SAV for Standalone CC. This solution focuses on how to construct a Standalone CC -- a subset of CC, in which all member ASes share transit providers on the top AS of the CC, no any other alternative transit provider. Under the valley-free principle of BGP routing, the traffic between the hosts in the Standalone CC must not traverse the provider interfaces of the CC-Top-AS, so a PI-SAV source prefix blocklist can be generated based on this Standalone CC.

Section 5 introduces the improved solution: PI-SAV for Standalone+ CC. This solution does not simply assume alternative transit provider will cause CC traffic detour through the provider interfaces as the PI-SAV for Standalone CC solution does, but to check whether there is route policy configuration or link failure causing so. If no detour exists, a Standalone CC+ can be formed by adding the corresponding ASes (the sub-CC under the alternative transit provider) to the Standalone CC. By doing so, a more complete PI-SAV prefix blocklist can be formed based on the Standalone+ CC. The mechanism for Standalone+ CC solution can also be used to enhance Standalone CC solution, working as an in-band channel to get provider information of CC-Member-AS in stead of relying on out-band RPKI ASPA or SLRUM records.

Note that this document mainly elaborates the basic principles and general procedures of the solutions, the details for additional information provisioning (e.g. Partial-Transit object management) or communication protocol between ASes for Standalone+ CC is not included.

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 Boarder 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, which is rarely deployed in the wild. When an AS (AS-a) receives the 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, the 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: a subset of CC, all ASes within a Standalone CC share same transit provider(s) on the CC-Top-AS, i.e. the member AS of a Standalone CC has no alternative transit provider outside the CC, all Sub-Transit-CCes in the CC must be excluded. Based on the valley-free principle of BGP routing, the traffic flow between the Standalone CC hosts will not go through the CC-Top-AS's provider interfaces.

Standalone+ CC: a superset of a Standalone CC that includes the Sub-Transit-CC(s) which has no detour happened for the CC traffic, i.e., 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.

Independence Degree: a metric that evaluates the degree of independence of a Standalone CC or Standalone+ CC within its customer cone. It is defined as the percentage of ASes or prefixes inside the CC that qualify as Standalone/Standalone+ CC. Two variations are considered:

* AS Independence Degree: The ratio of Standalone/Standalone+ ASes within the CC to the total ASes in the CC.

* Prefix Independence Degree: The ratio of Standalone/Standalone+ prefixes within the CC to the total prefixes in the CC.

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

The PI-SAV for Standalone CC solution mainly focuses on solving two key problems: 1) constructing the Standalone CC by deducting the Sub-Transit-CC(es) from the whole CC without any omission, and 2) generating source prefix blocklist for the Standalone CC by subtracting the prefix(es) with multi-origin AS.

Standalone CC construction requires the complete knowledge of the CC-Member-AS and whether each of them has alternative transit provider outside the CC. Basically, CC-Top-ASBR can get the complete list of CC member AS based on BGP local-RIBs if there is no Partial Transit implemented in the CC, that is, ASBR needs to know whether there is any Partial Transit applied in the CC to guarantee the completeness of CC member AS list. Note that, although RPKI ASPA [I-D.ietf-sidrops-aspa-profile] can cover the knowledge of partial transit, and itself can also be used to find the list of CC member AS, it is not recommended to primarily rely on RPKI ASPA for generating CC member list. 1) It will be a long run for RPKI ASPA registration getting ubiquitous, and more importantly, ASPA record themselves cannot indicate whether all ASes have registered. 2)

Timeliness due to RPKI record update latency. It is normal situation that the RPKI record updates lay behind the real network changes. 3) It needs to search for all the ASPA records per customer layer, and carefully deal with the race condition between RPKI ASPA records.

While Standalone CC gets constructed, the initial prefix blocklist for PI-SAV can be generated based on local-RIBs. However, prefix in this list may get involved with multiple origin, that is, the same prefix may also be announced by CC outside network. Furthermore, different ASes in the CC may choose different origin AS path for a specific MOAS prefix, in this case, legitimate traffic with this source prefix may flow in the CC through provider interfaces of the CC, so the CC source prefixes with outside origin AS should be deducted from the initial blocklist. The knowledge of MOAS prefix may be covered by RPKI ROA [RFC9582] and/or RPKI ROA SLURM.

Based on the above discussion, the solution proposed in this document constructs the PI-SAV blocklist primarily based on local-RIBs information and additional supports from local SLURM and RPKI records. The general procedure for CC-Top-ASBR generating PI-SAV blocklist for Standalone CC listed as below:

1. Initial CC-Member-AS collection. Includes 2 parts: a) ASNs for all direct customer ASes of the CC-Top-AS. The CC-Top-ASBR should have this information in its system. b) Get partial transit deployment information (ASN of CC-Member-AS which only partial transit with upper ASes in the CC) for the target CC. SLURM based local management for partial transit information is suggested.

2. CC construction. Based on the initial ASNs collected above, find all CC-Member-ASes by searching the local-RIBs for which is under the initial ASN in the AS-paths. The transit relationship between member ASes must be saved for further Sub-Transit-CC construction.

3. Sub-Transit-CC construction. Firstly, check whether there is any alternative transit provider outside the CC for each CC-Member-AS, mark the CC-Member-AS as a CC-sub-Transit-AS if there is. This can be done by RPKI ASPA records, and SLURM based local support for additional information is encouraged. Note, if the transit provider information is not available for any CC-Member-AS, the AS must be marked as CC-sub-Transit-AS. Then construct Sub-Transit-CC for each CC-sub-Transit-AS.

4. Standalone CC construction. Deduct all the Sub-Transit-CCes from the CC to form Standalone CC.

5. Generating the initial prefix list for the Standalone CC. Based on BGP local-RIBs, collect all prefixes initiated from the member ASes of the Standalone CC.

6. Final prefix blocklist generation. Check whether any prefix in the initial list get involved with multiple origin AS (MOAS), that is, the prefix is originated not only by the AS(es) inside the CC but also by the AS(es) outside the CC, deduct the prefix with MOAS from the initial list. MOAS prefix indentification can be based on RPKI ROA and/or local ROA SLURM records. Local-RIBs-in can also be used to indentify whether there is any external origin AS for a CC source prefix. Note, router implementation in current practice may not save all local-RIBs-in records, so it needs further consideration if we want find all CC prefixes with external origin AS.

Note: the prefix blocklist must avoid improper block while generating the prefix list, which means while a prefix is not sure and hard to determine whether it should be included in the blocklist, just do not include it. The hidden prefix described in [I-D.ietf-savnet-inter-domain-problem-statement] (No-Export and DSR etc.) might be the cases.

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

Although alternative transit link may cause the traffic between CC member ASes detouring through provider interfaces of the CC, however, the detour path, which normally is longer path compared to inner CC path, only happens while specific routing policy configuration or internal transit path failure making so. That is, traffic detour through provider interface for CC communication is not the normal case.

On the basis of the previous PI-SAV for Standalone CC solution, the initial goal of the PI-SAV for Standalone+ CC solution is to identify as many prefixes as possible for the PI-SAV blocklist where the Sub-Transit-CC they belongs actually has no detour happen. This solution will be more meaningful if Standalone CC can only count a small portion of whole CC.

Another benefit of the improved solution is that it can enhance PI-SAV for Standalone CC solution by providing additional information support via in-band channel (communication between CC member ASes), rather than heavily relying on out-band source (RPKI and/or local SLURM), where the former can provide more accurate, complete, and timely information support.

There are two modes described in this document:

1) Conservative Mode: This mode starts from the Standalone CC, forming the Standalone+ CC by adding a Sub-Transit-CC back when it is confirmed that no detour happened for this Sub-Transit-CC.

2) Aggressive Mode: This mode starts from the whole CC, forming the Standalone+ CC by pruning a Sub-Transit-CC out when it is confirmed that detour happened for this Sub-Transit-CC. The aggressive mode requires all CC-sub-Transit-ASes support cooperation with the CC-Top-AS based on the PI-SAV for Standalone+ CC solution.

3.1. Conservative Mode

The general procedures for PI-SAV for Standalone+ CC conservative mode:

1) CC-Top-ASBR enables PI-SAV for Standalone+ CC conservative mode. The system sets the Standalone CC as the initial Standalone+ CC. And then it sends a notification message to all CC-Member-ASes. The information in the message includes:

-- The mode of PI-SAV for Standalone+ CC.

-- CC-Top-ASBR information, including the router identification and information about how it can be connected for member AS to response.

-- The CC-Member-AS and CC-sub-Transit-AS information.

After the notification message, update message must be sent if relevant information changes.

2) CC-Member-AS receives the notification/update message. Based on the BGP local-RIBs, CC-Member-AS verifies whether it has preferred an external transit provider for the BGP updates originated from CC member ASes. Then send response message back to CC-Top-ASBR.

Since then, response update message must be sent while related route preferences change (due to routing policy or transit link status change etc.).

3) CC-Top-ASBR receives the response message. It will add the corresponding CC-sub-Transit-AS and its Sub-Transit-CCs into the Standalone+ CC while the message shows no detour through alternative provider.

4) CC-Top-ASBR generates blocklist for Standalone+ CC like Standalone CC does, including deduction of MOAS prefixes.

Note that, in conservative mode, in case no response message is received from a sub-transit-AS, it will be assuming that there is external detour paths exist for it.

Based on the communication mechanism between CC-Top-AS and CC-Member-AS, the procedure can also be used for enhancing PI-SAV for Standalone CC solution. It can provide more accurate, complete and real-time feedback for additional information required, for example, whether a member AS has partial transit implemented, has external provider existence, or a CC prefix with external origin AS, rather than just relying on out-band pre-registration (RPKI or local SLURM).

3.2. Aggressive Mode

The general procedures for PI-SAV for Standalone+ CC aggressive mode:

1) When the CC-Top-ASBR enables PI-SAV for Standalone+ CC aggressive mode, the system sets the whole CC as the initial Standalone+ CC. Then it sends a notification message to all CC-Member-ASes, the information in the message is similar with conservative mode but with aggressive mode identifier.

After that, update message need to be sent if relevant information changes.

2) The CC-Member-AS receives the notification/update message, based on the BGP Local-RIBs, verifying whether it has preferred an external transit provider for the BGP updates originated from CC member ASes. Then send response message back to CC-Top-ASBR.

Since then, response update message must be sent while related route preference change (due to routing policy or transit link status change etc.).

3) After CC-Top-ASBR receives the response message, it will prune the corresponding CC-sub-Transit-AS and its sub-standalone-CCs from the initial Standalone+ CC while the message shows detour through alternative provider.

4) Only after CC-Top-AS gets response from all CC-sub-Transit-ASes, generates blocklist for Standalone+ CC like Standalone CC does, including deduction of MOAS prefixes.

Same as the previous conservative mode, this mechanism can also be used for PI-SAV for Standalone CC solution enhancement.

4. Considerations

This document focuses on the general solution architecture for PI-SAV for CC, including the communication mechanism between CC-Top-AS and CC-sub-Transit-AS for PI-SAV for Standalone+ CC solution, the details of protocol design is not included. Some considerations:

4.1. Deployment Considerations.

The deployment position for the solutions is suggested for a region center AS, rather than a top/high tier AS with complex connection across regions. This means the ASes in the CC are normally/mainly connected to outside internet through the region top AS, and it would be easier and more efficient for them to implement comprehensive PI-SAV for CC solution.

Meanwhile, the solution should allow the incremental deployment in this region/the CC, which means some member ASes may not support the solution or some information (e.g. ASPA records for some ASes) may not be available. However, 1) The partial transit information must be fully covered to avoid the hidden AS in the CC. 2) MOAS prefixes must be fully identified to avoid related prefix improper block.

4.2. Security Considerations.

This is mainly related to the communication mechanism between ASes in the Standalone+ scheme. Generally (not base on the detailed protocol design), the communication mechanism must be equipped with the reliability and security features inherent in conventional transport protocols. Additionally, to avoid MITM (Man-in-the-middle) attacks, the protocol message can be encrypted and authenticated based on router-keys.

Note: the source prefix of response message should not be in the blocklist in case the message needs to detour back through provider interfaces, this can be done at early normal situation.

4.3. Operation and Management Considerations

The deployment of PI-SAV for CC solutions involves operational decisions that can be guided by the Independence Degree metric defined in this document. Operators may evaluate the feasibility and benefit of deploying PI-SAV for CC based on the following aspects:

1) Using Independence Degree for Deployment Decision

Before enabling PI-SAV for CC, the operator of the CC-Top-AS may compute the AS Independence Degree or Prefix Independence Degree for the Standalone CC and Standalone+ CC. These metrics indicate the coverage and effectiveness that can be achieved by the Standalone/ Standalone+ CC solution.

* If both degrees are high (e.g., above 80%), deploying the Standalone CC solution can provide substantial protection with relatively simple operations.

* If the degree for Standalone CC is moderate (e.g., 40%~80%) while the degree for Standalone+ CC is higher (e.g., above 80%), operators may consider deploying the Standalone+ CC solution to extend coverage, especially if there is evidence that many Sub-Transit-CCs do not actually detour through external providers.

* If the degrees are low (e.g., below 40%), the benefit of deploying PI-SAV for CC may be limited unless the Standalone+ CC solution with aggressive mode is supported by most member ASes.

2) Monitoring and Reporting

It is recommended that implementations provide monitoring capabilities for Independence Degree over time. Operators should track:

* Changes in AS and prefix counts within the CC, Standalone CC, and Standalone+ CC.

* Reasons for exclusion of ASes or prefixes (e.g., due to alternative transit, MOAS, partial transit).

* Alerting when Independence Degree falls below/increases over a configured threshold.

Such monitoring helps operators understand the dynamics of their customer cone and adjust PI-SAV policies accordingly.

3) Incremental Deployment and Cooperation

Even if the Independence Degree is initially low, incremental deployment remains possible. Operators may start with a pilot deployment among a subset of cooperative member ASes (e.g., those without alternative transit). The Standalone+ CC solution’s communication mechanism can be used to gradually expand coverage as more member ASes participate.

Operators should also consider establishing agreements or incentives for CC-Member-ASes to share accurate provider and MOAS information, thereby improving the accuracy and completeness of Independence Degree calculation and PI-SAV blocklist generation.

5. IANA Considerations

This document includes no request to IANA.

6. References

6.1. 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>.

6.2. 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-12, 20 October 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-savnet-inter-domain-problem-statement-12>.

[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]
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-22, 6 February 2026, <https://datatracker.ietf.org/doc/html/draft-ietf-sidrops-aspa-profile-22>.

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