

Opsec Working Group
Internet-Draft
Intended status: Best Current Practice
Expires: May 3, 2018

K. Sriram
D. Montgomery
US NIST
J. Haas
Juniper Networks, Inc.
October 30, 2017

Enhanced Feasible-Path Unicast Reverse Path Filtering
draft-sriram-opsec-urpf-improvements-02

Abstract

This document identifies a need for improvement of the unicast Reverse Path Filtering techniques (uRPF) [BCP84] for source address validation (SAV) [BCP38]. The strict uRPF is inflexible about directionality, the loose uRPF is oblivious to directionality, and the current feasible-path uRPF attempts to strike a balance between the two [BCP84]. However, as shown in this draft, the existing feasible-path uRPF still has short comings. This document proposes an enhanced feasible-path uRPF technique, which aims to be more flexible (in a meaningful way) about directionality than the feasible-path uRPF. It can potentially alleviate ISPs' concerns about the possibility of disrupting service for their customers, and encourage greater deployment of uRPF techniques.

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 May 3, 2018.

Copyright Notice

Copyright (c) 2017 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
1.1. Requirements Language	3
2. Review of Existing Source Address Validation Techniques . . .	3
2.1. SAV using Access Control List	4
2.2. SAV using Strict Unicast Reverse Path Filtering	4
2.3. SAV using Feasible-Path Unicast Reverse Path Filtering .	5
2.4. SAV using Loose Unicast Reverse Path Filtering	6
3. Proposed New Technique: SAV using Enhanced Feasible-Path uRPF	7
3.1. Description of the Method	7
3.2. Operational Recommendations	8
3.3. A Challenging Scenario	9
3.4. Overcoming the Above Challenge: Algorithm with Full Flexibility Across Customer Cone	10
3.5. Implementation Considerations	11
3.5.1. Impact on FIB Memory Size Requirement	11
4. Security Considerations	12
5. IANA Considerations	12
6. Acknowledgements	13
7. Informative References	13
Authors' Addresses	14

1. Introduction

This internet draft identifies a need for improvement of the unicast Reverse Path Filtering (uRPF) techniques [RFC2827] for source address validation (SAV) [RFC3704]. The strict uRPF is inflexible about directionality, the loose uRPF is oblivious to directionality, and the current feasible-path uRPF attempts to strike a balance between the two [RFC3704]. However, as shown in this draft, the existing feasible-path uRPF still has short comings. Even with the feasible-path uRPF, ISPs are often apprehensive that they may be dropping customers' data packets with legitimate source addresses.

This document proposes an enhanced feasible-path uRPF technique, which aims to be more flexible (in a meaningful way) about directionality than the feasible-path uRPF. It is based on the

principle that if BGP updates for multiple prefixes with the same origin AS were received on different interfaces (at border routers), then incoming data packets with source addresses in any of those prefixes should be accepted on any of those interfaces (described in Section 3.1). For some challenging ISP-customer scenarios (see Section 3.3), we further propose (a) Forming a list of all unique prefixes in the collection of routes received on all customer interfaces; and (b) Including that list in the RPF list of each customer interface (described in Section 3.4). Implementation considerations are discussed in Section 3.5.

Note: Definition of Reverse Path Filtering (RPF) list: The list of permissible source address prefixes for incoming data packets on a given interface.

The proposed techniques are expected to add greater operational robustness and efficacy to uRPF, while minimizing ISPs' concerns about accidental service disruption for their customers. It is expected that this will encourage more deployment of uRPF so as to realize its DDoS prevention benefits networkwide.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2. Review of Existing Source Address Validation Techniques

There are various existing techniques for mitigation against DDoS attacks with spoofed addresses [RFC2827] [RFC3704]. There are also some techniques used for mitigating reflection attacks [RRL] [TA14-017A], which are used to amplify the impact in DDoS attacks. Employing a combination of these preventive techniques in enterprise and ISP border routers, DNS servers, broadband and wireless access networks, and data centers provides reasonably effective protection against DDoS attacks.

Source address validation (SAV) is performed in network edge devices such as border routers, Cable Modem Termination Systems (CMTS), Digital Subscriber Line Access Multiplexers (DSLAM), and Packet Data Network (PDN) gateways in mobile networks. Ingress Access Control List (ACL) and unicast Reverse Path Filtering (uRPF) are techniques employed for implementing SAV [RFC2827] [RFC3704] [ISOC].


2.1. SAV using Access Control List

Ingress/egress Access Control Lists (ACLs) are maintained which list acceptable (or alternatively, unacceptable) prefixes for the source addresses in the incoming Internet Protocol (IP) packets. Any packet with a source address that does not match the filter is dropped. The ACLs for the ingress/egress filters need to be maintained to keep them up to date. Updating the ACLs is an operator driven manual process, and hence operationally difficult or infeasible.

Typically, the egress ACLs in access aggregation devices (e.g. CMTS, DSLAM) permit source addresses only from the address spaces (prefixes) that are associated with the interface on which the customer network is connected. Ingress ACLs are typically deployed on border routers, and drop ingress packets when the source address is spoofed (i.e. belongs to obviously disallowed prefix blocks, RFC 1918 prefixes, or provider's own prefixes).

2.2. SAV using Strict Unicast Reverse Path Filtering

In the strict unicast Reverse Path Filtering (uRPF) method, an ingress packet at border router is accepted only if the Forwarding Information Base (FIB) contains a prefix that encompasses the source address and forwarding information for that destination prefix points back to the interface over which the packet was received. In other words, the reverse path for routing to that source address (if it were used as a destination address) should use the same interface over which the packet was received. It is well known that this method has limitations when networks are multi-homed and there is asymmetric routing of packets. Asymmetric routing occurs (see Figure 1) when a customer AS announces one prefix (P1) to one transit provider (ISP-a) and a different prefix (P2) to another transit provider (ISP-b), but routes data packets with source addresses in the second prefix (P2) to the first transit provider (ISP-a) or vice versa.



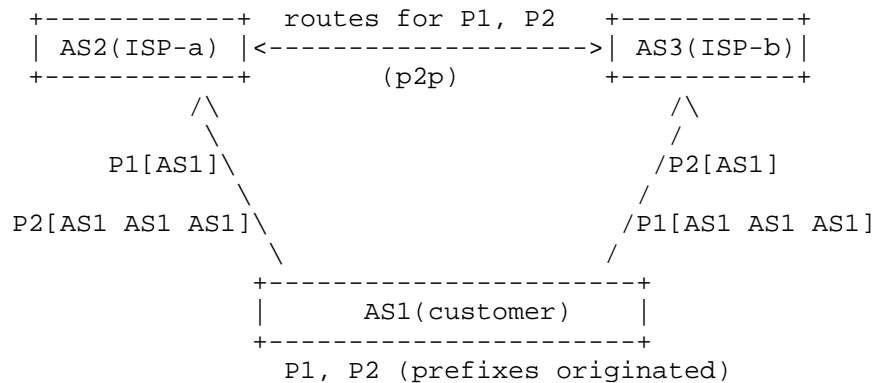
Consider data packets received at AS2
 (1) from AS1 with source address in P2, or
 (2) from AS3 that originated from AS1
 with source address in P1:

- * Strict uRPF fails
- * Feasible-path uRPF fails
- * Loose uRPF works (but ineffective in IPv4)
- * Enhanced Feasible-path uRPF works best

Figure 1: Scenario 1 for illustration of efficacy of uRPF schemes.

2.3. SAV using Feasible-Path Unicast Reverse Path Filtering

The feasible-path uRPF helps partially overcome the problem identified with the strict uRPF in the multi-homing case. The feasible-path uRPF is similar to the strict uRPF, but in addition to inserting the best-path prefix, additional prefixes from alternative announced routes are also included in the RPF table. This method relies on announcements for the same prefixes (albeit some may be prepended to effect lower preference) propagating to all routers performing feasible-path uRPF checks. Therefore, in the multi-homing scenario, if the customer AS announces routes for both prefixes (P1, P2) to both transit providers (with suitable prepends if needed for traffic engineering), then the feasible-path uRPF method works (see Figure 2). It should be mentioned that the feasible-path uRPF works in this scenario only if customer routes are preferred at AS2 and AS3 over a shorter non-customer route.



Consider data packets received at AS2 via AS3 that originated from AS1 and have source address in P1:

- * Feasible-path uRPF works (if customer route to P1 is preferred at AS3 over shorter path)
- * Feasible-path uRPF fails (if shorter path to P1 is preferred at AS3 over customer route)
- * Loose uRPF works (but ineffective in IPv4)
- * Enhanced Feasible-path uRPF works best

Figure 2: Scenario 2 for illustration of efficacy of uRPF schemes.

However, the feasible-path uRPF method has limitations as well. One form of limitation naturally occurs when the recommendation of propagating the same prefixes to all routers is not followed. Another form of limitation can be described as follows. In Scenario 2 (described above, illustrated in Figure 2), it is possible that the second transit provider (ISP-b or AS3) does not propagate the prepended route for prefix P1 to the first transit provider (ISP-a or AS2). This is because AS3's decision policy permits giving priority to a shorter route to prefix P1 via a peer (AS2) over a longer route learned directly from the customer (AS1). In such a scenario, AS3 would not send any route announcement for prefix P1 to AS2. Then a data packet with source address in prefix P1 that originates from AS1 and traverses via AS3 to AS2 will get dropped at AS2.

2.4. SAV using Loose Unicast Reverse Path Filtering

In the loose unicast Reverse Path Filtering (uRPF) method, an ingress packet at the border router is accepted only if the FIB has one or more prefixes that encompass the source address. That is, a packet is dropped if no route exists in the FIB for the source address. Loose uRPF sacrifices directionality. This method is not effective for prevention of address spoofing since there is little unrouted address space in IPv4. It only drops packets if the spoofed address

is unreachable in the current FIB (e.g. RFC 1918, unallocated, allocated but currently not routed).

3. Proposed New Technique: SAV using Enhanced Feasible-Path uRPF

3.1. Description of the Method

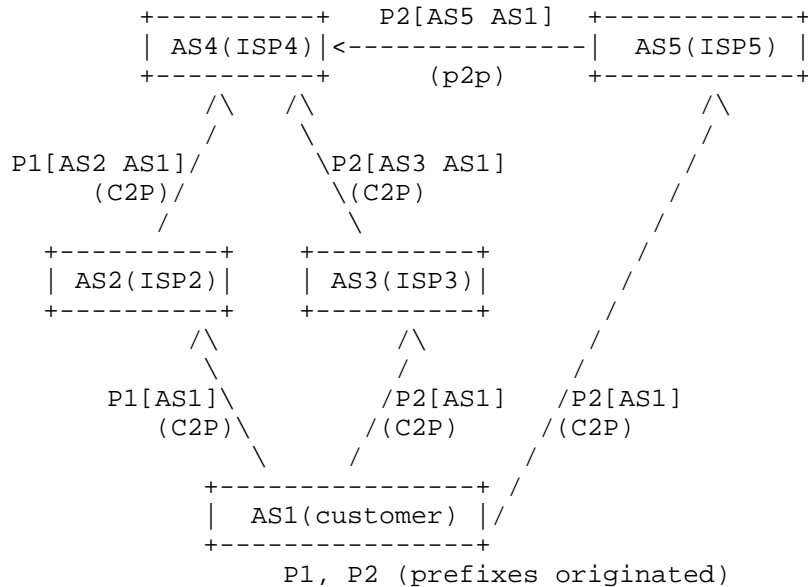
Enhanced feasible-path uRPF adds greater operational robustness and efficacy to existing uRPF methods discussed in Section 2. The proposed technique is based on the principle that if BGP updates for multiple prefixes with the same origin AS were received on different interfaces (at border routers), then incoming data packets with source addresses in any of those prefixes should be accepted on any of those interfaces. It can be best explained with an example as follows:

Let us say, a border router of ISP-A has in its Adj-RIB-in the set of prefixes {Q1, Q2, Q3} each of which has AS-x as its origin and AS-x is in ISP-A's customer cone. Further, the border router received a route for prefix Q1 over a customer facing interface, while it learned routes for prefixes Q2 and Q3 from a lateral peer and an upstream transit provider, respectively. All these routes passed route filtering and/or origin validation (i.e. the origin AS-x is deemed legitimate). That is, the route announcements are considered legitimate. In this example scenario, the enhanced feasible-path uRPF method allows source addresses to belong in {Q1, Q2, Q3} on any of the three specific interfaces in question (customer, peer, provider) on which the three routes were learned.

Thus, enhanced feasible-path uRPF defines feasible paths in a more generalized but precise way (as compared to feasible-path uRPF). In the above example, routes for prefixes Q2 and Q3 were not received on a customer facing interface at the border router, yet data packets with source addresses in Q2 or Q3 are accepted by the router if they come in on the same customer interface on which the route for prefix Q1 was received (based on these prefix routes having the same origin AS).

Looking back at Scenarios 1 and 2 (Figure 1 and Figure 2), the enhanced feasible-path uRPF provides comparable or better performance than the other uRPF methods. Scenario 3 (Figure 3) further illustrates the enhanced feasible-path uRPF method with a more concrete example. In this scenario, the focus is on operation of the feasible-path uRPF at ISP4 (AS4). ISP4 learns a route for prefix P1 via a customer-to-provider (C2P) interface from customer ISP2 (AS2). This route for P1 has origin AS1. ISP4 also learns a route for P2 via another C2P interface from customer ISP3 (AS3). Additionally, AS4 learns an alternate route for P2 via a peer-to-peer (p2p)

interface from ISP5 (AS5). Both routes for P2 have the same origin AS (i.e. AS1) as does the route for P1. Using the proposed enhanced feasible-path uRPF scheme, given the commonality of the origin AS across the above-mentioned routes for P1 and P2, AS4 would permit source addresses belonging to either P1 or P2 in data packets received on any of the three interfaces (from AS2, AS3, and AS5).



Consider that data packets (sourced from AS1) may be received at AS4 with source address in P1 or P2 via any of the neighbors (AS2, AS3, AS5):

- * Feasible-path uRPF fails
- * Loose uRPF works (but not desirable)
- * Enhanced Feasible-path uRPF works best

Figure 3: Scenario 3 for illustration of efficacy of uRPF schemes.

Based on the above, the proposed enhanced feasible-path uRPF method would reduce ISP concerns about possible service disruption affecting their customers and encourage greater adoption of uRPF.

3.2. Operational Recommendations

The following operational recommendations will make the operation of the proposed enhanced feasible-path uRPF robust:

For multi-homed stub AS:

- o A multi-homed stub AS SHOULD announce at least one of the prefixes it originates to each of its transit provider ASes.

For non-stub AS:

- o A non-stub AS SHOULD also announce at least one of the prefixes it originates to each of its transit provider ASes.
- o Additionally, from the routes it has learned from customers, a non-stub AS SHOULD announce at least one route per origin AS to each of its transit provider ASes.

(Note: It is worth noting that in the above recommendations if "at least one" is replaced with "all", then even traditional feasible-path uRPF will work as desired.)

3.3. A Challenging Scenario

It should be observed that in the absence of ASes adhering the above recommendations, the following example scenarios may be constructed which pose a challenge for the enhanced feasible-path uRPF (as well as for traditional feasible-path uRPF). In the scenario illustrated in Figure 4, since routes for neither P1 nor P2 are propagated on the AS2-AS4 interface, the enhanced feasible-path uRPF at AS4 will reject data packets received on that interface with source addresses in P1 or P2.




Figure 4: Illustration of a challenging scenario.

3.4. Overcoming the Above Challenge: Algorithm with Full Flexibility Across Customer Cone

Adding further flexibility to the enhanced feasible-path uRPF method can help address the potential limitation identified above using the scenario in Figure 4 (Section 3.3). In the following, "route" refers to a route currently existing in the Adj-RIB-in. Including the additional degree of flexibility, the modified algorithm can be described as follows:

- o Let $I = \{I_1, I_2, \dots, I_n\}$ represent the set of all directly-connected customer interfaces at customer-facing edge routers in a transit provider's AS.
- o Let $P = \{P_1, P_2, \dots, P_m\}$ represent the set of all unique prefixes for which routes were received over the interfaces in Set I .
- o Let $A = \{AS_1, AS_2, \dots, AS_k\}$ represent the set of all unique origin ASes seen in the routes that were received over the interfaces in Set I .
- o Let $Q = \{Q_1, Q_2, \dots, Q_j\}$ represent the set of all unique prefixes for which routes were received over peer or provider interfaces such that each of the routes has its origin AS belonging in Set A .

- o Then, Set Z = Union(P,Q) represents the RPF list for each customer-facing edge router in the AS in question. That is, over each interface in Set I, the edge router SHOULD permit only those ingress data packets that have SA in any of the prefixes in Set Z.

When this algorithmic flexibility is incorporated, then the type of limitation identified in Figure 4 (Section 3.3) goes away. This should significantly reduce the possibility of blocking legitimate customer-data packets in uRPF implementations.

3.5. Implementation Considerations

The existing RPF checks in edge routers take advantage of existing line card implementations to perform the RPF functions. For implementation of the proposed technique, the general necessary feature would be to extend the line cards to take arbitrary RPF lists that are not necessarily the same as the existing FIB contents. For example, in the proposed method, the RPF lists are constructed by applying a set of rules to all received BGP routes (not just those selected as best path and installed in FIB).

3.5.1. Impact on FIB Memory Size Requirement

The proposed technique requires that there should be FIB memory (i.e., TCAM) available to store the RPF lists in line cards. For an ISP's AS, the RPF list size for each line card will roughly and conservatively equal the total number of prefixes in its customer cone (assuming the algorithm in Section 3.4 is used). The following table shows the measured customer cone sizes for various types of ISPs [sriram-ripe63]:

Type of ISP	Measured Customer Cone Size in # Prefixes (in turn this is an estimate for RPF list size on line card)
Very Large Global ISP	32392
Very Large Global ISP	29528
Large Global ISP	20038
Mid-size Global ISP	8661
Regional ISP (in Asia)	1101

Table 1: Customer cone sizes (# prefixes) for various types of ISPs.

For some super large global ISPs that are at the core of the Internet, the customer cone size (# prefixes) can be as high as a few hundred thousand [caida]. But uRPF is most effective when deployed at ASes at the edges of the Internet where the customer cone sizes are smaller as shown in Table 1.

A very large global ISP's router line card is likely to have a FIB size large enough to accommodate 2 to 6 million routes [cisco1]. Similarly, the line cards in routers corresponding to a large global ISP, a mid-size global ISP, and a regional ISP are likely to have FIB sizes large enough to accommodate about 1 million, 0.5 million, and 100K routes, respectively [cisco2]. Comparing these FIB size numbers with the corresponding RPF list size numbers in Table 1, it can be surmised that the conservatively estimated RPF list size is only a small fraction of the anticipated FIB memory size under various ISP scenarios.

4. Security Considerations

This document offers a technique to improve the robustness features of uRPF and thus improve the security of the Internet as a whole. The proposed technique does not warrant any additional security considerations.

5. IANA Considerations

This document does not request new capabilities or attributes. It does not create any new IANA registries.

6. Acknowledgements

The authors would like to thank Job Snijders, Marco Marzetti, Marco d'Itri, Nick Hilliard, Gert Doering, Igor Gashinsky, Barry Greene, and Joel Jaeggli for comments and suggestions.

7. Informative References

- [caida] "Information for AS 174 (COGENT-174)", CAIDA Spoofer Project , <<https://spoofer.caida.org/as.php?asn=174>>.
- [cisco1] "Internet Routing Table Growth Causes ROUTING-FIB-4-RSRC_LOW Message on Trident-Based Line Cards", Cisco Trouble-shooting Tech-notes , January 2014, <<https://www.cisco.com/c/en/us/support/docs/routers/asr-9000-series-aggregation-services-routers/116999-problem-line-card-00.html>>.
- [cisco2] "Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 5.x (Chapter: Managing the Unicast RIB and FIB)", Cisco Configuration Guides , June 2017, <https://www.cisco.com/c/en/us/td/docs/switches/data-center/sw/5_x/nx-os/unicast/configuration/guide/13_cli_nxos/13_manage-routes.html#22859>.
- [ISOC] Vixie (Ed.), P., "Addressing the challenge of IP spoofing", ISOC report , September 2015, <<https://www.us-cert.gov/ncas/alerts/TA14-017A>>.
- [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>>.
- [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>>.
- [RFC6811] Mohapatra, P., Scudder, J., Ward, D., Bush, R., and R. Austein, "BGP Prefix Origin Validation", RFC 6811, DOI 10.17487/RFC6811, January 2013, <<https://www.rfc-editor.org/info/rfc6811>>.

[RRL] "Response Rate Limiting in the Domain Name System",
Redbarn blog , <<http://www.redbarn.org/dns/ratelimits>>.

[sriram-ripe63]
Sriram, K. and R. Bush, "Estimating CPU Cost of BGPSEC on
a Router", Presented at RIPE-63; also at IETF-83 SIDR WG
Meeting, March 2012,
<[http://www.ietf.org/proceedings/83/slides/
slides-83-sidr-7.pdf](http://www.ietf.org/proceedings/83/slides/slides-83-sidr-7.pdf)>.

[TA14-017A]
"UDP-Based Amplification Attacks", US-CERT alert
TA14-017A , January 2014,
<<https://www.us-cert.gov/ncas/alerts/TA14-017A>>.

Authors' Addresses

Kotikalapudi Sriram
US NIST
100 Bureau Drive
Gaithersburg MD 20899
USA

Email: ksriram@nist.gov

Doug Montgomery
US NIST
100 Bureau Drive
Gaithersburg MD 20899
USA

Email: dougmon@nist.gov

Jeffrey Haas
Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale CA 94089
USA

Email: jhaas@juniper.net