Opsec Working Group
Internet-Draft
Intended status: Best Current Practice
Expires: November 4, 2017

K. Sriram
NIST
D. Montgomery
US NIST
May 3, 2017

Enhanced Feasible-Path Unicast Reverse Path Filtering
draft-sriram-opsec-urpf-improvements-01

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 is expected to alleviate ISPs' concerns about the possibility of disrupting service for their customers, and encourage greater deployment of uRPF.

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 http://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 November 4, 2017.

Copyright Notice

Table of Contents

1. Introduction

This internet draft identifies a need for improvement of the unicast Reverse Path Filtering techniques (uRPF) [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 denying 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 an edge router), then data packets with source addresses in any of those prefixes are allowed to be received on any of those interfaces. This technique is expected

to add greater operational logic and efficacy to uRPF, and alleviate ISPs' concerns about the possibility of disrupting service for their customers. It should encourage greater deployment of uRPF to realize its DDoS prevention benefits network wide.

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 deterrence against DDoS attacks with spoofed addresses [RFC2827] [RFC3704]. There are also some techniques used for prevention of reflection-amplification attacks [RRL] [TA14-017A], which are used in achieving greater 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 the necessary protections 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/outgoing 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. Hence, this method may be operationally difficult or infeasible in dynamic environments such as when a customer network is multihomed, has address space allocations from multiple ISPs, or dynamically varies its BGP announcements (i.e. routing) for traffic engineering purposes.

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 on an interface at the border router is accepted only if the Forwarding Information Base (FIB) contains a prefix that encompasses the source address and packet forwarding for that prefix points to said interface. In other words, the best path for routing to that source address (if it were used as a destination address) should point to said interface. It is well known that this method has limitations when a network or autonomous system is 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.

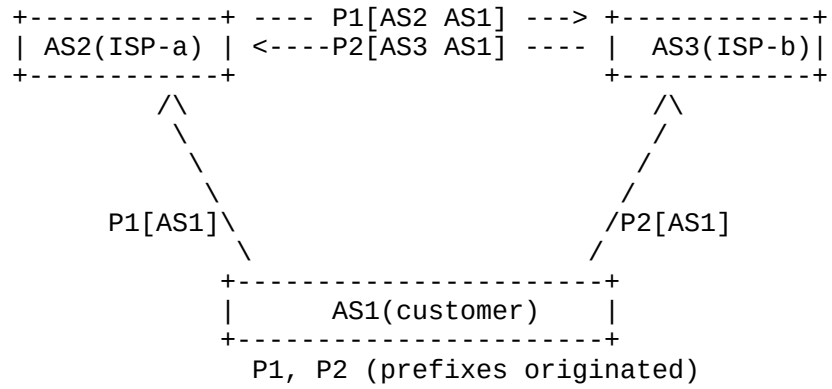

Figure 1: Scenario 1 for illustration of efficacy of uRPF schemes.

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 the difference is that instead of inserting one best route in the FIB (or an equivalent RPF table), alternative routes are also added there. This method relies on announcements for the same prefixes (albeit some may be prepended to effect lower preference) propagating to all the routers performing feasible-path uRPF check. So 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 route is preferred at AS2 and AS3 over the shorter path.

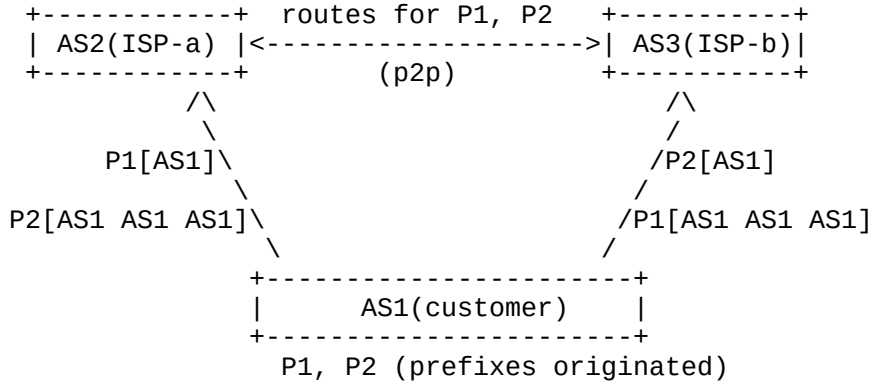

Consider data packets received at AS2 via AS3
that originated from AS1 and have source address in P1:
* Feasible-path uRPF works (if customer route preferred at AS3 over shorter path)
* Feasible-path uRPF fails (if shorter path preferred at AS3 over customer route)
* Loose uRPF works (but not desirable)
* 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 heeded. 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. In most cases, this method is not useful for prevention of address spoofing. It only drops packets if the spoofed address is non-routable (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 logic and efficacy to existing uRPF methods discussed in Section 2. It can be best explained with an example. 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 belongs 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 prefixes passed route filtering and/or origin validation (i.e. the origin AS-x is deemed 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 for those scenarios. 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. Applying the principle of enhanced feasible-path uRPF, given the commonality of the origin AS across the above-mentioned routes for P1 and P2, AS4 permits the SA in data packets to belong in P1 or P2 on any of the three interfaces (from AS2, AS3, and AS5).

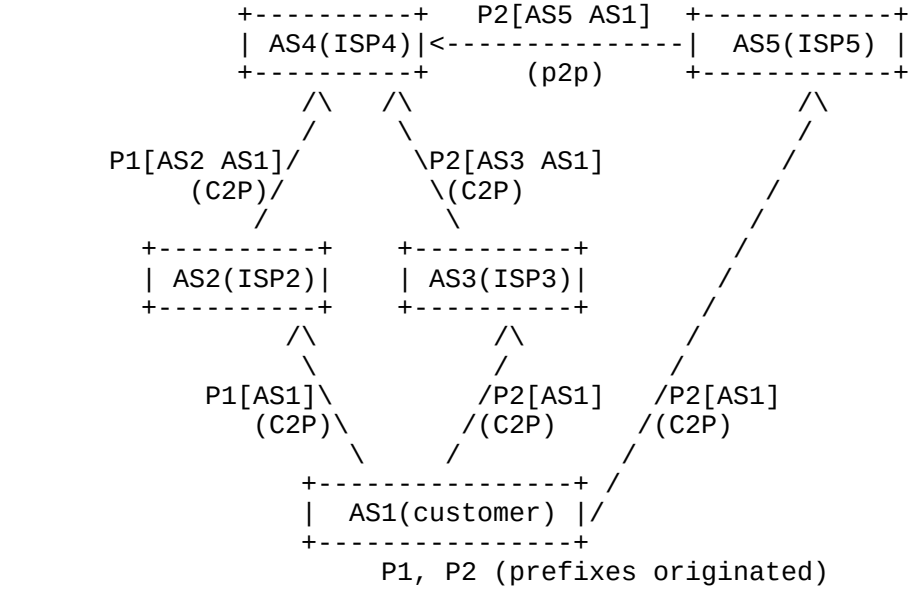

Consider that data packets (sourced from AS1) may be received at AS4 with source address in P1 or P2 via any of the neighbors (AS2, AS3, AS5):
* Feasible-path uRPF fails
* Loose uRPF works (but not desirable)
* Enhanced Feasible-path uRPF works best

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

Based on the above, it can be possibly rationalized that the proposed enhanced feasible-path uRPF method would help alleviate ISP concerns about possible service disruption for their customers and encourage greater adoption of uRPF.

3.2. Operational Recommendations

The following operational recommendations if followed will make robust the desired operation of the enhanced feasible-path uRPF proposed here.

For multi-homed stub AS:

- A multi-homed stub AS SHOULD announce at least one of its origination prefixes to each transit provider AS.

For non-stub AS:

- A non-stub AS SHOULD announce at least one of its origination prefixes to each transit provider AS.

- Additionally, from the routes it has learned from customers, a non-stub AS SHOULD announce at least one route for each unique {prefix, origin AS} pair to each transit provider AS.

(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.)

Also, it should be observed that in the absence of ASes adhering the above recommendations, the following type of 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. But this can be clearly avoided if the above recommendations for stub and non-stub ASes are followed. In this example, this would mean that the NO_EXPORT is avoided and instead AS prepending is used (to depref routes) on the AS1-AS2 peering session.

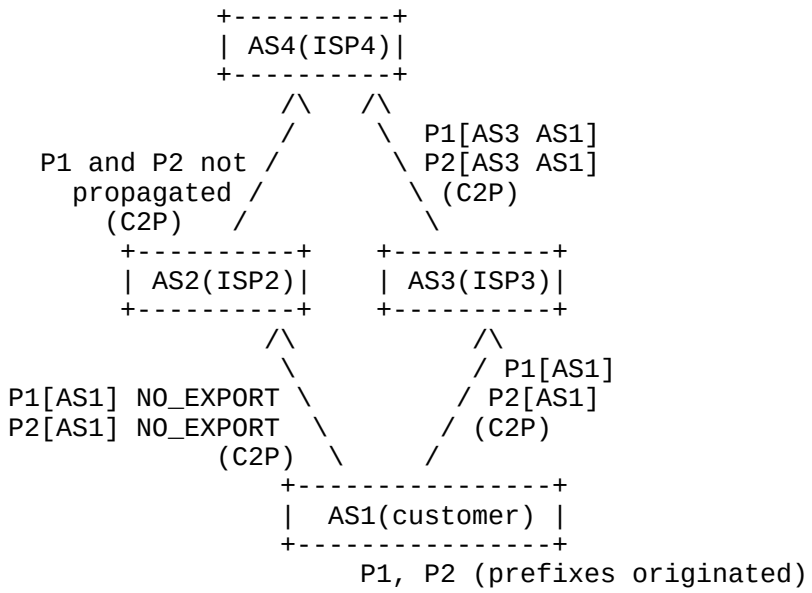

Figure 4: Illustration of a challenging scenario.

3.3. Customer Cone Consideration

An additional degree of flexibility that can be incorporated in the enhanced feasible-path uRPF can be described as follows. Let I = {I1, I2, ..., In} represent the set of all directly-connected customer interfaces at customer-facing edge routers in a transit provider's AS. Let P = {P1, P2, ..., Pm} represent the set of all prefixes for which routes have been received over the interfaces in set I. Then, over all interfaces in the set I, the edge router SHOULD permit ingress data packets with SA in any of the prefixes in the set P.

3.4. Implementation Consideration

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 tied to 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).

4. Security Considerations

This document offers a technique to improve the security features of uRPF. 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 Jeff Haas, Job Snijders, Marco Marzetti, Marco d'Itri, Nick Hilliard, Gert Doering, Barry Greene, and Joel Jaeggli for comments and suggestions.

7. Informative References

[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, <http://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, <http://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, <http://www.rfc-editor.org/info/rfc3704>.

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

[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
NIST
100 Bureau Drive
Gaithersburg MD 20899
USA

Email: ksriram@nist.gov

Doug Montgomery
US NIST
100 Bureau Drive
Gaithersburg MD 20899
USA

Email: dougm@nist.gov