Authors: K. Sriram   I. Lubashev   D. Montgomery
         USA NIST   Akamai   USA NIST

Source Address Validation Using BGP UPDATEs, ASPA, and ROA (BAR-SAV)

Abstract

Designing an efficient source address validation (SAV) filter
requires minimizing false positives (i.e., avoiding blocking
legitimate traffic) while maintaining directionality (see RFC8704).
This document advances the technology for SAV filter design through
a method that makes use of BGP UPDATE messages, Autonomous System
Provider Authorization (ASPA), and Route Origin Authorization (ROA).
The proposed method's name is abbreviated as BAR-SAV. BAR-SAV can be
used by network operators to derive more robust SAV filters and thus
improve network resilience. This document updates RFC8704.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

Spoofed source addresses are often used in Denial of Service (DoS)
and Distributed DoS (DDoS) attacks. Source address validation (SAV)
filtering is used to drop packets with spoofed source addresses (see
BCP 84 [RFC3704] [RFC8704]). A detailed review of unicast Reverse
Path Forwarding (uRPF) techniques for SAV is provided in [RFC8704]).
Also, [RFC8704] describes enhanced feasible-path uRPF (EFP-uRPF)
methods that aim to minimize false positives (i.e., avoid blocking
legitimate traffic) while maintaining directionality (see
definitions in [RFC3704]).

New technology for securing the Border Gateway Protocol (BGP)
[RFC4271] using Resource Public Key Infrastructure (RPKI) [RFC6480]
is seeing increasing adoption. Two of the currently existing or
proposed types of signed objects in the RPKI can be leveraged for a
more accurate SAV filter design as well. These are the Route Origin
Authorization (ROA) and the Autonomous System Provider

Authorizations (ASPA) objects. A ROA is a cryptographically signed attestation by an IP address-resource holder listing their prefixes that are authorized to be originated in BGP by a specific autonomous system (AS) [RFC6482]. ROAs are currently used for RPKI-based Route Origin Validation (RPKI-ROV) [RFC6811] [RFC9319]. An ASPA is a cryptographically signed attestation by an AS listing its transit provider AS numbers (ASNs) [I-D.ietf-sidrops-aspa-profile]. The ASPA data is designed to be used for a form of AS path validation that can detect and mitigate route leaks [I-D.ietf-sidrops-aspa-verification] [sriram1]. See [RFC7908] for the definition of route leaks.

This document advances the technology for SAV filter design using methods that make use of ASPA, ROA, and/or BGP UPDATE data. A method is presented in Section 3 that makes use of only ASPA and ROA data to design the SAV filter. This method is for use in the future when the adoption of ROA and ASPA is considered to be ubiquitous. However, for use in the period before that, another method for SAV is presented in Section 4 that makes complementary use of BGP UPDATE messages along with ASPA and ROA data. Accordingly, the latter method's name is abbreviated as BAR-SAV. It is hoped that just as the adoption of ROAs is growing at present [Monitor], the adoption of ASPA will also gain momentum in the near future. The BAR-SAV method additionally incorporates a refined version of Algorithm A of the EFP-uRPF technique (Section 3.1 of [RFC8704]). BAR-SAV can be used by network operators to derive more robust SAV filters and thus improve network resilience.

The focus of this document is on the design of ingress SAV allowlist filters for an interface facing a customer or lateral peer AS. The same procedure applies in both cases (Section 2).

Throughout this document, ROA and ASPA data mean the payload data in cryptographically valid ROA and ASPA objects (see Section 4 in [RFC6482] and Section 4 in [I-D.ietf-sidrops-aspa-profile]).

The reader is encouraged to be familiar with [RFC8704], [RFC6482], [RFC6811], [I-D.ietf-sidrops-aspa-profile], and [I-D.ietf-sidrops-aspa-verification].

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. **Same Procedure Applies to Customers and Lateral Peers**

    The same procedure applies for the construction of a permissible
    ingress SAV filter for a customer or lateral peer interface, because
    the data packets received from a customer or lateral peer should
    have source addresses belonging only to the prefixes in the customer
    cone (CC) of said customer or lateral peer. The focus, therefore, is
    only on the CC of the neighbor in each case. Note that the CC
    includes the AS belonging to the customer or lateral peer.

3. **SAV Using ASPA and ROA (Procedure X)**

    The procedure (called Procedure X) described in this section is for
    future scenarios when ASPA and ROA adoption is ubiquitous. In that
    scenario, robust SAV filters can be generated from the RPKI
    information (ASPA and ROA data) alone. The procedure is applicable
    for ingress SAV filter design for customer and lateral peer
    interfaces. An ISP may use Procedure X on a customer (or lateral
    peer) interface if it expects full adoption of ROAs and ASPAs in the
    CC of the neighbor.

    A description of Procedure X (one that makes use of only ASPA and
    ROA data):

      *Step A: Compute the set of ASNs in the Customer's or Lateral
       Peer's customer cone using ASPA data.

      *Step B: Compute from ROA data the set of prefixes authorized to
       be announced by the ASNs found in Step A. Keep only the unique
       prefixes. This set is the permissible prefix list for SAV for the
       interface in consideration.

    A detailed description of Procedure X is as follows:

      1. Let the Customer or Lateral Peer ASN be denoted as AS-k.

      2. Let i = 1. Initialize: AS-set S(1) = {AS-k}.

      3. Increment i to i+1.

      4. Create AS-set S(i) of all ASNs whose ASPA data declares at
         least one ASN in AS-set S(i-1) as a Provider.

      5. If AS-set S(i) is null, then set i_max = i - 1 and go to Step
         6. Else, go to Step 3.

      6. Form the union of the sets, S(i), i = 1, 2, ..., i_max, and
         name this union as AS-set A.

7. Select all ROAs in which the authorized origin ASN is equal to
      any ASN in AS-set A. Form the union of the sets of prefixes
      listed in the selected ROAs. Name this union set of prefixes as
      P-set.

   8. Apply P-set as the list of permissible prefixes for SAV.

4. **SAV using BGP UPDATE Messages, ASPA, and ROA (BAR-SAV)**

   SAV using BGP UPDATE Messages, ASPA, ROA (BAR-SAV) as well as ACLs
   is described in this section and is meant for the period when there
   is a partial deployment of ROAs and ASPAs. To compensate for
   incomplete RPKI information, BAR-SAV augments ASPA data with BGP
   UPDATE AS_PATH data (and ASN ACLs) for discovering CC ASes, and it
   augments ROA data with BGP UPDATE data (and Prefix ACLs) for
   discovering all prefixes associated with ASes in the CC. The details
   of this procedure are described below.

   BAR-SAV additionally incorporates a refined version of Algorithm A
   of EFP-uRPF (Section 3.1 of [RFC8704]). Algorithm A in [RFC8704]
   picked only the originating ASes from AS_PATHs received on the
   customer or lateral peer interface in consideration and included
   them for SAV filter computation. The variant of Algorithm A in
   [RFC8704] used here includes all ASes in the AS_PATHs for the SAV
   filter computation. Unless there is a route leak [RFC7908], each AS
   is a customer of the AS added next in AS_PATHs of BGP UPDATE
   messages received from a customer or lateral peer. Further customer-
   provider AS relations within the CC are discovered by examining all
   unique ASes in the AS_PATHs in BGP UPDATEs received on all
   interfaces (from transit providers, customers, lateral peers, and
   IBGP peers). This is described in the step-by-step procedure later
   in this section.

   Note that if a multi-homed AS is present in an above-mentioned
   AS_PATH and did not originate any prefix in the CC in consideration
   but originated a prefix into an overlapping neighboring CC, then the
   AS and prefix will still be detected and included in the design of
   the SAV filter. This improves the accuracy of the SAV filter in the
   BAR-SAV method in comparison to Algorithm A in [RFC8704].

   One should not compute a customer cone by separately processing ASPA
   data and AS_PATH data and then merging the two sets of ASes at the
   end. Doing so is likely to miss ASes from the customer cone.
   Instead, both ASPAs and AS_PATHs should be used to iteratively
   expand the discovered customer cone. When new ASes are discovered,
   both ASPA and AS_PATH data should be used to discover customers of
   those ASes. This process is repeated for newly discovered customer
   ASes until there are no new ASes to be found.

As a measure of security, validation of the AS_PATH data in Adj-RIBs-In [RFC4271] SHOULD be performed using the procedures in [I-D.ietf-sidrops-aspa-verification] and any Invalid AS_PATHs must be excluded from inputs to the BAR-SAV procedure. This ensures that BGP UPDATEs containing route leaks are not considered for BAR-SAV filter design. Please see additional discussion about route leaks in Section 9.

As a further measure of security, validation of BGP routes in Adj-RIBs-In MUST be performed by applying RPKI-ROV [RFC6811] and any Invalid routes must be excluded from inputs to the BAR-SAV procedure. Please see additional discussion about prefix/route filtering in Section 9.

A detailed description of the BAR-SAV procedure is as follows:

1. Let the Customer or Lateral Peer ASN be denoted as AS-k.

2. Let i = 1. Initialize: AS-set Z(1) = {AS-k}.

3. Extend AS-set Z(1) to include ASNs from any ASN ACL configured for this interface.

4. Increment i to i+1.

5. Create AS-set A(i) of all ASNs whose ASPA data declares at least one ASN in AS-set Z(i-1) as a Provider.

6. Create AS-set B(i) of all customer ASNs each of which is a customer of at least one ASN in AS-set Z(i-1) according to unique AS_PATHs in Adj-RIBs-In of all interfaces at the BGP speaker computing the SAV filter.

7. Form the union of AS-sets A(i) and B(i) and call it AS-set C. From AS-set C, remove any ASNs that are present in Z(j), for j=1 to j=(i-1). Call the resulting set Z(i).

8. If AS-set Z(i) is null, then set i_max = i - 1 and go to Step 9. Else, go to Step 4.

9. Form the union of the AS-sets, Z(i), i = 1, 2, ..., i_max, and name this union as AS-set D.

10. Select all ROAs in which the authorized origin ASN is in AS-set D. Form the union of the sets of prefixes listed in the selected ROAs. Name this union set of prefixes as Prefix-set P1.

11. Using the routes in Adj-RIBs-In of all interfaces, create a
    list of all prefixes originated by any ASN in AS-set D. Name
    this set of prefixes as Prefix-set P2.

12. Form the union of Prefix-set P1, Prefix-set P2, and any Prefix
    ACL configured for this interface. Apply this union set as the
    list of permissible prefixes for SAV.

5. Operational Recommendations

Network operators SHOULD implement the BAR-SAV method (Section 4)
for computing the permissible ingress prefix list for SAV on
interfaces facing customers and lateral peers. BAR-SAV offers
immediate incremental benefits to early adopters.

The operational recommendations provided in Section 3.2 of [RFC8704]
are applicable and helpful for BAR-SAV (Section 4). Since Procedure
X (Section 3) and the BAR-SAV procedure (Section 4) benefit from the
registration of ROAs, network operators are RECOMMENDED to register
ROAs and enable RPKI-ROV in their ASes. When ASPA registration
becomes available, network operators are also RECOMMENDED to
register ASPAs at that time.

The registration of ROAs and ASPAs helps with the detection and
inclusion of otherwise hidden prefixes in the permissible list for
SAV. As mentioned earlier, prefixes hidden in other SAV techniques
often arise from the use of multi-homing in conjunction with limited
propagation of prefixes in a given CC (for example, by attaching
NO_EXPORT to all prefixes announced from a customer AS to a transit
provider AS). In these situations, the registration of ASPAs and
ROAs helps improve the accuracy of SAV.

5.1. Considerations for the CDN and DSR Scenario

Direct Server Return (DSR) is a common asymmetric routing scenario
that is not supported by existing BCP-84 uRPF [RFC3704] and EFP-uRPF
[RFC8704] SAV methods. DSR is commonly used by Content Delivery
Networks (CDNs) that wish to use anycast service addresses but
deliver data from edge locations that do not announce anycast
addresses.

For example, in Figure 1, the CDN announces an anycast prefix P3
(from AS3) from a well-connected location with CDN control
infrastructure. When a User from prefix P1 (AS1) establishes a
connection to the anycast address and requests an object, an Anycast
Server at the CDN may determine that the best location to serve the
object is an Edge Server in a location close to the User. The Edge
Server is reachable only via prefix P2 (AS2). The Anycast Server can
forward packets arriving from the User to the Edge Server (via IP-IP
tunneling or similar means), but the bulk data transmission would

need to happen directly from the Edge Server to the User with an
anycast source address (a P3 address).
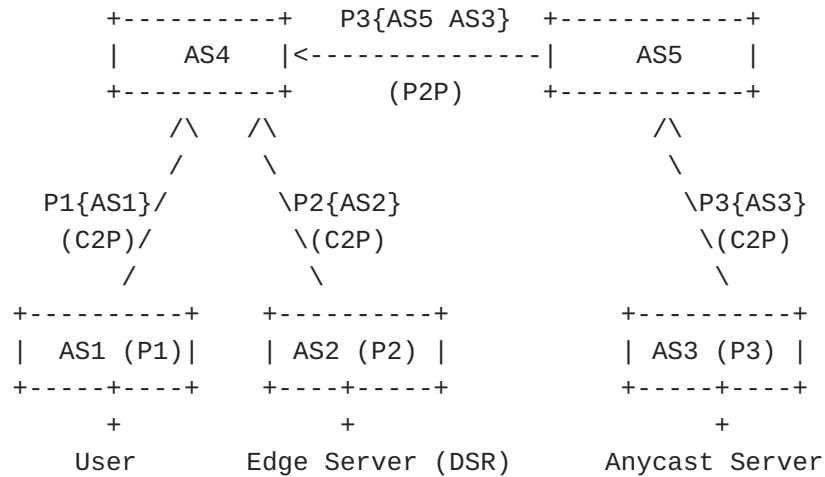
```
         +----------+   P3{AS5 AS3}  +------------+
         |    AS4   |<---------------|    AS5     |
         +----------+      (P2P)     +------------+
            /\    /\                      /\
           /      \                        \
      P1{AS1}/      \P2{AS2}                \P3{AS3}
       (C2P)/        \(C2P)                  \(C2P)
         /             \                       \
    +----------+    +----------+         +----------+
    | AS1 (P1)|    | AS2 (P2) |         | AS3 (P3) |
    +-----+----+    +----+-----+         +-----+----+
          +              +                     +
        User        Edge Server (DSR)     Anycast Server

        Consider AS4 generating SAV list for interface to AS2:
        CDN's ROAs: {P3, AS3}, {P3, AS2}, {P2, AS2}
        AS2 should not/does not announce P3
        With the SAV methods in this document,
           AS4 correctly includes P2 and P3 in the SAV list
```

Figure 1: Illustration of how the solution functions for the CDN/DSR
scenario.

Existing SAV methods of [RFC3704] and EFP-uRPF [RFC8704] would not
allow AS4 to include P3 as a legitimate SA prefix on the interface
to AS2. However, if the CDN (owner of prefix P3) registers a ROA
object authorizing AS2 to originate P3, and AS4 uses an SAV
procedure specified in this document (Section 4), then AS4 will use
that ROA object to include P3 as a valid source prefix for the AS2
customer interface. The CDN may never want to announce a route to P3
from AS2, but the existence of this ROA would result in the
construction of an SAV filter that would permit AS2 to send data
packets with source addresses belonging to P3.

The CDN example above is just one DSR scenario. There are other
cloud-based DSR scenarios that include low-latency gaming, mobile
roaming, corporate networks of global enterprises, and others.

Recommendation: In a DSR scenario, a network operator MUST register
ROAs that bind the edge server ASes with the anycast service prefix.
This is in addition to registering a ROA authorizing the anycast
server AS to announce the anycast prefix.

6. Operations and Management Considerations

This section highlights some important operations and management considerations and was motivated in part to address the comments received from the SIDROPS working group members.

6.1. Applicability of ASPA and ROA

A transit provider is a network that (a) offers its customers outbound (customer to Internet) data traffic connectivity and/or (b) further propagates in all directions (towards providers, lateral peers, and other customers) any BGP Updates that the customer may send [I-D.ietf-sidrops-aspa-profile]. In the latter case, it also provides transport for inbound data traffic. In all cases, the customer AS SHOULD follow the specification in [I-D.ietf-sidrops-aspa-profile] and include the transit provider AS in its ASPA. Registering an ASPA prevents forged-origin hijacks for the customer AS and its prefixes, prevents route leaks involving the customer AS, and facilitates BAR-SAV.

If a prefix is used for source addresses for hosts attached at an AS but not announced in BGP from that AS (e.g., the DSR scenario in Section 5.1), a ROA MUST be registered binding the prefix and the AS. This ROA registration assists in preventing hijacking of the prefix and helps facilitate BAR-SAV. The risk of this ROA registration enabling a forged-origin prefix hijack for the prefix is minimal since the ASPA-based path verification [I-D.ietf-sidrops-aspa-verification] prevents forged-origin attacks. It may be noted that a similar usage of ROA is made in the context of DDoS mitigation (see Section 5.1 in [RFC9319]), where hypothetically the prefix may never need to be originated by the AS of the DDoS mitigation provider.

6.2. BAR-SAV and Routing Policy

BAR-SAV identifies all ASes in a customer's (or lateral peer's) customer cone (CC), and then it discovers all prefixes that could plausibly be used as source addresses in data traffic originated from the ASes in the CC. If ASPA and ROA have been adopted by all ASes and prefix owners, respectively, in the CC of interest, then the list of plausible source address prefixes will be complete with no improper block (i.e., traffic with legitimate source addresses is not blocked). Further, deploying BAR-SAV by all ASes within the CC ensures no improper permit (i.e., traffic with spoofed source address is not admitted). Note that routing policies of ASes may be such that some of the discovered prefixes may never be used as source addresses on a given customer interface of interest, but this does not impact BAR-SAV's accuracy.

6.3. Where to Deploy BAR-SAV

The discussion in Section 3.6.1 of [RFC8704] of the Forwarding
Information Base (FIB) size estimates and the networks where SAV
would be most effective are applicable to BAR-SAV as well. Smaller
ISPs (and possibly some midsize and regional ISPs) are expected to
implement the BAR-SAV method, since SAV in general is most effective
closer to the edges of the Internet. For such networks, the
conservatively estimated SAV filter list size is only a small
fraction of the anticipated FIB memory size (see details in Section
3.6.1 of [RFC8704]).

6.4. Automation is the Key

SAV done manually, e.g., using ACLs, usually does not get much
adoption because of operational costs, susceptibility to human
errors, and tendency of SAV filters to get out of date due to the
need for any changes by customers or peers to be coordinated with
multiple parties (providers and peers). Automated uRPF technique,
such as BAR-SAV, however, allow for easy, accurate, and cost
effective deployments. The BAR-SAV method makes it possible to
automate the construction of SAV filter lists aiming for no improper
block and a minimal probability of improper permit of data traffic.
As ASPA adoption picks up alongside the ongoing ROA adoption, BAR-
SAV's accuracy of discovering all possible source addresses
(prefixes) for the customer cone of interest improves even further
in complex scenarios.

6.5. Implementation Guidelines

When a SAV filter is used to police data traffic, and an incomplete
SAV filter list could cause legitimate traffic to be blocked, the
use of robust implementation practices for RPKI data retrieval and
cache management practices become paramount. Some of such
recommended practices are discussed in this section.

6.5.1. Management of Local RPKI Repository Caches

RPKI infrastructure does not guarantee continuous availability of
RPKI repositories. Local caches of RPKI signed objects, manifest
files (MFTs), and certificate revocation lists (CRLs) are already
maintained for managing ROA objects and router certificates
[RFC8210]. That is being extended to ASPA objects as well
[I-D.ietf-sidrops-8210bis]. The cache refresh frequency currently
used for RPKI data should be sufficient for BAR-SAV purposes as
well. If an RPKI repository publication point is unavailable, or
there is any other failure in fetching its objects, the latest
cached version of the objects associated with the repository MUST
continue to be used, as described in [RFC9286].

If the local cache of some repository objects required for BAR-SAV computation is unavailable (for example, due to a filesystem failure) and/or the RPKI data cannot be fetched from the repository publication point, the SAV system SHOULD "fail open" and downgrade the SAV function on a given interface to "loose uRPF" described in [RFC3704] and [RFC8704]. This downgrade is better than suspending SAV entirely since at least source addresses in unallocated and bogon space are rejected.

7. Provider Interfaces

For most networks, "loose uRPF" SAV mode described in [RFC3704] and [RFC8704] is the current best practice for provider interfaces to ensure no improper block. It may be possible to use BGP UPDATE as well as RPKI ROA and ASPA data to compute a list of prefixes that originate exclusively within the customer cone and should only be received by the AS on customer interfaces. A set of such prefixes can be used in conjuction with "loose uRPF" SAV mode as a blocklist on provider interfaces.

8. IANA Considerations

This document includes no request to IANA.

9. Security Considerations

The security considerations described in [RFC8704], [RFC6811], [I-D.ietf-sidrops-aspa-profile], and [I-D.ietf-sidrops-aspa-verification] also apply to this document.

The security and robustness of BAR-SAV are strengthened by supporting mechanisms for detecting and dropping BGP routes that are misoriginations or leaks. Section 4 stated the requirement of validating BGP route origins using RPKI-ROV [RFC6811]. It further helps if route origin validation using trusted IRR route objects and prefix filtering are also deployed (see [RFC7454] [NIST-800-189]). It is also advised that one or more of the available methods to prevent, detect, and mitigate route leaks are deployed (e.g., [RFC9234] [I-D.ietf-grow-route-leak-detection-mitigation] [I-D.ietf-sidrops-aspa-verification] [sriram1]).

10. References

10.1. 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>.

[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>.

[RFC4271]   Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A
                 Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI
                 10.17487/RFC4271, January 2006, <https://www.rfc-
                 editor.org/info/rfc4271>.

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

[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>.

[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>.

[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>.

10.2. Informative References

[RFC6480]   Lepinski, M. and S. Kent, "An Infrastructure to Support
                 Secure Internet Routing", RFC 6480, DOI 10.17487/RFC6480,
                 February 2012, <https://www.rfc-editor.org/info/rfc6480>.

[RFC7454]   Durand, J., Pepelnjak, I., and G. Doering, "BGP
                 Operations and Security", BCP 194, RFC 7454, DOI
                 10.17487/RFC7454, February 2015, <https://www.rfc-
                 editor.org/info/rfc7454>.

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

[RFC8210]   Bush, R. and R. Austein, "The Resource Public Key
                 Infrastructure (RPKI) to Router Protocol, Version 1", RFC
                 8210, DOI 10.17487/RFC8210, September 2017, <https://
                 www.rfc-editor.org/info/rfc8210>.

[RFC9234]
Azimov, A., Bogomazov, E., Bush, R., Patel, K., and K.
Sriram, "Route Leak Prevention and Detection Using Roles
in UPDATE and OPEN Messages", RFC 9234, DOI 10.17487/
RFC9234, May 2022, <https://www.rfc-editor.org/info/
rfc9234>.

[RFC9286] Austein, R., Huston, G., Kent, S., and M. Lepinski,
"Manifests for the Resource Public Key Infrastructure
(RPKI)", RFC 9286, DOI 10.17487/RFC9286, June 2022,
<https://www.rfc-editor.org/info/rfc9286>.

[RFC9319] Gilad, Y., Goldberg, S., Sriram, K., Snijders, J., and B.
Maddison, "The Use of maxLength in the Resource Public
Key Infrastructure (RPKI)", BCP 185, RFC 9319, DOI
10.17487/RFC9319, October 2022, <https://www.rfc-
editor.org/info/rfc9319>.

[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-12, 29 January
2023, <https://datatracker.ietf.org/doc/html/draft-ietf-
sidrops-aspa-profile-12>.

[I-D.ietf-sidrops-aspa-verification]
Azimov, A., Bogomazov, E., Bush, R., Patel, K., Snijders,
J., and K. Sriram, "BGP AS_PATH Verification Based on
Autonomous System Provider Authorization (ASPA) Objects",
Work in Progress, Internet-Draft, draft-ietf-sidrops-
aspa-verification-12, 6 March 2023, <https://
datatracker.ietf.org/doc/html/draft-ietf-sidrops-aspa-
verification-12>.

[I-D.ietf-grow-route-leak-detection-mitigation] Sriram, K. and A.
Azimov, "Methods for Detection and Mitigation of BGP
Route Leaks", Work in Progress, Internet-Draft, draft-
ietf-grow-route-leak-detection-mitigation-08, 24 October
2022, <https://datatracker.ietf.org/doc/html/draft-ietf-
grow-route-leak-detection-mitigation-08>.

[I-D.ietf-sidrops-8210bis] Bush, R. and R. Austein, "The Resource
Public Key Infrastructure (RPKI) to Router Protocol,
Version 2", Work in Progress, Internet-Draft, draft-ietf-
sidrops-8210bis-10, 16 June 2022, <https://

datatracker.ietf.org/doc/html/draft-ietf-sidrops-8210bis-10>.

[sriram1]   Sriram, K. and J. Heitz, "On the Accuracy of Algorithms
             for ASPA Based Route Leak Detection", IETF SIDROPS
             Meeting, Proceedings of the IETF 110, March 2021,
             <https://datatracker.ietf.org/meeting/110/materials/
             slides-110-sidrops-sriram-aspa-alg-accuracy-01>.

[Monitor]   "NIST RPKI Monitor", National Institute of Standards and
             Technology, accessed June 2022, <https://rpki-
             monitor.antd.nist.gov/>.

[NIST-800-189] Sriram, K. and D. Montgomery, "Resilient Interdomain
             Traffic Exchange: BGP Security and DDoS Mitigation", NIST
             Special Publication, NIST SP 800-189, December 2019,
             <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/
             NIST.SP.800-189.pdf>.

Acknowledgements

Authors' Addresses

Kotikalapudi Sriram
USA National Institute of Standards and Technology
100 Bureau Drive
Gaithersburg, MD 20899
United States of America

Email: ksriram@nist.gov

Igor Lubashev
Akamai Technologies
145 Broadway
Cambridge, MA 02142
United States of America

Email: ilubashe@akamai.com

Doug Montgomery
USA National Institute of Standards and Technology
100 Bureau Drive
Gaithersburg, MD 20899
United States of America

Email: dougm@nist.gov