

Workgroup: Network Working Group
Internet-Draft:
draft-ietf-sidrops-aspa-verification-09
Published: 11 July 2022
Intended Status: Standards Track
Expires: 12 January 2023
Authors: A. Azimov (Ed.) E. Bogomazov
 Yandex Qrator Labs
 R. Bush K. Patel
 Internet Initiative Japan & Arccus Arccus, Inc.
 J. Snijders
 Fastly

BGP AS_PATH Verification Based on Resource Public Key Infrastructure (RPKI) Autonomous System Provider Authorization (ASPA) Objects

Abstract

This document defines the semantics of an Autonomous System Provider Authorization object in the Resource Public Key Infrastructure to verify the AS_PATH attribute of routes advertised in the Border Gateway Protocol. This AS_PATH verification is primarily intended for detection and mitigation of route leaks. It also provides protection against forged-origin prefix hijacks.

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.

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 12 January 2023.

Copyright Notice

Copyright (c) 2022 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- [1. Introduction](#)
- [2. Anomaly Propagation](#)
- [3. Autonomous System Provider Authorization](#)
- [4. Customer-Provider Verification Procedure](#)
- [5. AS_PATH Verification](#)
 - [5.1. Upstream Paths](#)
 - [5.2. Downstream Paths](#)
 - [5.3. Mitigation](#)
- [6. Disavowal of Provider Authorizaion](#)
- [7. Mutual Transit \(Complex Relations\)](#)
- [8. Comparison to Peerlock](#)
- [9. Security Considerations](#)
- [10. Acknowledgments](#)
- [11. References](#)
 - [11.1. Normative References](#)
 - [11.2. Informative References](#)
- [Authors' Addresses](#)

1. Introduction

The Border Gateway Protocol (BGP) was designed without mechanisms to validate BGP attributes. Two consequences are BGP Hijacks and BGP Route Leaks [[RFC7908](#)]. BGP extensions are able to partially solve these problems. For example, ROA-based Origin Validation [[RFC6483](#)] can be used to detect and filter accidental mis-originations, and [[RFC9234](#)] or [[I-D.ietf-grow-route-leak-detection-mitigation](#)] can be used to detect accidental route leaks. While these upgrades to BGP are quite useful, they still rely on transitive BGP attributes, i.e. AS_PATH, that can be manipulated by attackers.

BGPsec [[RFC8205](#)] was designed to solve the problem of AS_PATH validation using a cryptographic signatures included in the UPDATE.

Unfortunately, the cryptographic validation of path signatures results in significant computational overhead for BGP routers. More importantly, while BGPsec offers protection against path or prefix modifications, it does not protect against route leaks.

An alternative approach was introduced with soBGP [[I-D.white-sobgp-architecture](#)]. Instead of strong cryptographic AS_PATH validation, it created an AS_PATH security function based on a shared database of AS adjacencies. While such an approach has reasonable computational cost, the two-side adjacencies don't provide a way to automate anomaly detection without high adoption rate - an attacker can easily create a one-way adjacency. soBGP transported data about adjacencies in new additional BGP messages, which was recursively complex thus significantly increasing adoption complexity. In addition, the general goal of verification of all AS_PATHs was not achievable given the indirect adjacencies at Internet exchange points.

Instead of strictly checking AS_PATH correctness, this document focuses on solving real-world operational problems - automatic detection of route leaks and combined with ROA detection of malicious bgp hijacks. To achieve this, new AS_PATH verification procedures are described to automatically detect invalid (malformed) AS_PATHs in announcements that are received from customers, peers, providers, Route Servers (RSes), and RS-clients. These procedures use a shared signed database of customer-to-provider relationships using a new RPKI object - Autonomous System Provider Authorization (ASPA). This technique provides benefits for participants even during early and incremental adoption.

2. Anomaly Propagation

Both route leaks and hijacks have similar effects on ISP operations - they redirect traffic, resulting in denial of service (DoS), eavesdropping, increased latency and packet loss. But the level of risk depends significantly on the extent of propagation of the anomalies. For example, a hijack that is propagated only to customers may cause bottlenecking within a particular ISP's customer cone, but if the anomaly is propagated through peers, upstreams, or reaches Tier-1 networks, thus distributing globally, the ill effects will likely be experienced across continents.

The ability to constrain propagation of BGP anomalies to upstreams and peers, without requiring support from the source of the anomaly (which is critical if source has malicious intent), should significantly improve the security of inter-domain routing and solve the majority of problems.

3. Autonomous System Provider Authorization

As described in [[RFC6480](#)], the RPKI is based on a hierarchy of resource certificates that are aligned to the Internet Number Resource allocation structure. Resource certificates are X.509 certificates that conform to the PKIX profile [[RFC5280](#)], and to the extensions for IP addresses and AS identifiers [[RFC3779](#)]. A resource certificate is a binding by an issuer of IP address blocks and Autonomous System (AS) numbers to the subject of a certificate, identified by the unique association of the subject's private key with the public key contained in the resource certificate. The RPKI is structured so that each current resource certificate matches a current resource allocation or assignment.

ASPA is digitally signed object that bind, for a selected AFI, a Set of Provider AS numbers to a Customer AS number (in terms of BGP announcements not business), and are signed by the holder of the Customer AS. An ASPA attests that a Customer AS holder (CAS) has authorized Set of Provider ASes (SPAS) to propagate the Customer's IPv4/IPv6 announcements onward, e.g. to the Provider's upstream providers or peers. The ASPA record profile is described in [[I-D.ietf-sidrops-aspa-profile](#)]. For a selected Customer AS SHOULD exist only single ASPA object at any time. In this document we will use ASPA(AS1, AFI, [AS2, ...]) as notation to represent ASPA object for AS1 in the selected AFI.

4. Customer-Provider Verification Procedure

This section describes an abstract procedure that checks that a pair of ASNs (AS1, AS2) is included in the set of signed ASPAs. The semantics of its use is defined in next section. The procedure takes (AS1, AS2, AFI) as input parameters and returns one of three results: "Valid", "Invalid" and "Unknown".

A relying party (RP) must have access to a local cache of the complete set of cryptographically valid ASPAs when performing customer-provider verification procedure.

The following algorithm describes the customer-provider verification procedure for selected AFI:

1. Retrieve all cryptographically valid ASPAs in a selected AFI with a customer value of AS1. The union of SPAS forms the set of "Candidate Providers."
2. If the set of Candidate Providers is empty, then the procedure exits with an outcome of "Unknown."
3. If AS2 is included in the Candidate Providers, then the procedure exits with an outcome of "Valid."

4. Otherwise, the procedure exits with an outcome of "Invalid."

Since an AS1 may have different set of providers in different AFI, it should also have different SPAS in corresponding ASPAs. In this case, the output of this procedure with input (AS1, AS2, AFI) may have different output for different AFI values.

5. AS_PATH Verification

The AS_PATH attribute identifies the autonomous systems through which an UPDATE message has passed. AS_PATH may contain two types of components: AS_SEQUENCES and AS_SETs, as defined in [[RFC4271](#)].

We will use index of AS_PATH, where Seg(1) stands for the first rightmost AS in the AS_PATH. We will use Seg(I).value and Seg(I).type to represent Ith segment value and its type respectively.

We define Invalid Pair Index as a minimal I such that Seg(I).type and Seg(I+1).type equal to AS_SEQUENCE, Seg(I).value != Seg(I+1).value and customer-provider validation procedure ([Section 4](#)) with parameters (Seg(I).value, Seg(I+1).value, AFI) returns Invalid. If I index doesn't exist we put the length of AS_PATH in its value.

We define Reverse Invalid Pair Index as Invalid Pair Index calculated for a reversed AS_PATH.

We define Unknown Pair Index as a minimal I Seg(I).type and Seg(I+1).type equal to AS_SEQUENCE, Seg(I).value != Seg(I+1).value and customer-provider validation procedure ([Section 4](#)) with parameters (Seg(I).value, Seg(I+1).value, AFI) returns Unknown. If I is greater than Invalid Pair Index or I doesn't exist we equate its value to the value of Invalid Pair Index.

We define Reverse Unknown Pair Index as Unknown Pair Index calculated for a reversed AS_PATH.

The below procedures are applicable only for 32-bit AS number compatible BGP speakers.

5.1. Upstream Paths

When a route is received from a customer, a lateral peer, by a RS or RS-client at an IX, each consecutive AS_SEQUENCE pair MUST be equal (prepend policy) or belong to customer-provider or mutual transit relationship ([Section 7](#)). If there are other types of relationships, it means that the route was leaked or the AS_PATH attribute was malformed and Invalid Pair Index will be less than AS_PATH length.

If an attacker creates route leak intentionally he may try to strip his AS from the AS_PATH. To strengthen route leak detection in case of malicious activity we need to check that AS_PATH is not empty and the latest AS in the AS_PATH equals to BGP neighbour AS with the exception for routes received from transparent IXes.

At the of high adoption level there might be interest to distinguish between AS_PATHs that are Valid from AS_PATHs that can't be fully verified and may be leaked. If route is received from a customer, a lateral peer, by a RS or RS-client at an IX and Unknown Pair Index is not equal to AS_PATH length it means that there is at least one AS without ASPA record.

The goal of the procedure described below is to check the correctness of these statements.

1. If the AS_PATH has zero length then procedure halts with the outcome "Invalid";
2. If the last segment in the AS_PATH has type AS_SEQUENCE and its value isn't equal to receiver's neighbor AS and receiver is not RS-client then procedure halts with the outcome "Invalid";
3. If Invalid Pair Index is less than AS_PATH length then procedure halts with the outcome "Invalid";
4. If the AS_PATH has at least one AS_SET segment then procedure halts with the outcome "Unverifiable";
5. If Unknown Pair Index is less than AS_PATH length then procedure halts with the outcome "Unknown";
6. Otherwise, the procedure halts with an outcome of "Valid".

5.2. Downstream Paths

When a route is received from provider it may have both Upstream and Downstream fragments, where a Downstream follows an Upstream fragment. If the path differs from this rule it means that the route was leaked or the AS_PATH attribute was malformed. This statement can be transformed into the next one: if there is at least one AS between the first Upstream fragment and the last Downstream fragment it is a route leak. The length of the first Upstream segment and last Downstream segment are defined by Invalid Pair Index and Reverse Invalid Pair Index respectively. Using these indexes we can define next rule for route leak detection for routes received from providers: if sum of Invalid Pair Index and Reverse Invalid Pair Index is less than AS_PATH length, than route was leaked or the AS_PATH attribute was malformed.

Likewise we did in case of Upstream Paths, we need to check that AS_PATH is not empty and the latest AS in the AS_PATH equals to BGP neighbour AS.

Similar to route leak detection, we can distinguish the Valid AS_PATH from Unknown one by checking that sum of Unknown Pair Index and Reverse Unknown Pair Index is equal or greater than AS_PATH length.

The goal of the procedure described below is to check the correctness of these statements.

1. If the AS_PATH has zero length then procedure halts with the outcome "Invalid";
2. If a route is received from a provider and the last segment in the AS_PATH has type AS_SEQUENCE and its value isn't equal to receiver's neighbor AS, then the procedure halts with the outcome "Invalid";
3. If sum of Invalid Pair Index and Reverse Invalid Pair Index is less than AS_PATH length, then the procedure halts with the outcome "Invalid".
4. If the AS_PATH has at least one AS_SET segment then procedure halts with the outcome "Unverifiable";
5. If sum of Unknown Pair Index and Unknown Invalid Pair Index is less than AS_PATH length, then the procedure halts with the outcome "Unknown".
6. Otherwise, the procedure halts with an outcome of "Valid".

5.3. Mitigation

If the output of the AS_PATH verification procedure is "Invalid" the route MUST be rejected.

If the output of the AS_PATH verification procedure is 'Unverifiable' it means that AS_PATH can't be fully checked. Such routes should be treated with caution and SHOULD be processed the same way as "Invalid" routes. This policy goes with full correspondence to [\[I-D.kumari-deprecate-as-set-confed-set\]](#).

The above AS_PATH verification procedure is able to check routes received from customer, peers, providers, RS, and RS-clients. The ASPA mechanism combined with BGP Roles [\[RFC9234\]](#) and ROA-based Origin Validation [\[RFC6483\]](#) can provide a fully automated solution to detect and filter hijacks and route leaks, including malicious ones.

6. Disavowal of Provider Authorizaion

An ASPA is a positive attestation that an AS holder has authorized its providers to redistribute received routes to the provider's providers and peers. This does not preclude the provider ASes from redistribution to its other customers. By creating an ASPA with providers set of [0], the customer indicates that no provider should further announce its routes. Specifically, AS 0 is reserved to identify provider-free networks, Internet exchange meshes, etc.

An ASPA(AS, AFI, [0]) is a statement by the customer AS that its routes should not be received by any relying party AS from any of its customers or peers.

By convention, an ASPA(AS, AFI, [0]) should be the only ASPA issued by a given AS holder in the selected AFI; although this is not a strict requirement. An AS 0 may coexist with other provider ASes in the same ASPA (or other ASPA records in the same AFI); though in such cases, the presence or absence of the provider AS 0 in ASPA does not alter the AS_PATH verification procedure.

7. Mutual Transit (Complex Relations)

There are peering relationships which can not be described as strictly simple peer-peer or customer-provider; e.g. when both parties are intentionally sending prefixes received from each other to their peers and/or upstreams.

In this case, two corresponding records ASPA(AS1, AFI, [AS2, ...]), ASPA(AS2, AFI, [AS1, ...]) must be created by AS1 and AS2 respectively.

8. Comparison to Peerlock

ASPA has much in common with [\[Peerlock\]](#). Peerlock is a [BGP Flexsealing](#) [\[Flexsealing\]](#) protection mechanism commonly deployed by global-scale Internet carriers to protect other large-scale carriers.

Peerlock, unfortunately, depends on a laborious manual process in which operators coordinate the distribution of unstructured Provider Authorizations through out-of-band means in a many-to-many fashion. On the other hand, ASPA's use of [PKIX](#) [\[RFC5280\]](#) allows for automated, scalable, and ubiquitous deployment, making the protection mechanism available to a wider range of Internet Number Resource holders.

ASPA mechanics implemented in code instead of Peerlock AS_PATH regular expressions also provides a way to detect anomalies coming from transit providers and internet exchange route servers.

ASPA is intended to be a complete solution and replacement for existing Peerlock deployments.

9. Security Considerations

The proposed mechanism is compatible only with BGP implementations that can process 32-bit ASNs in the AS_PATH. This limitation should not have a real effect on operations - such legacy BGP routers are rare and it's highly unlikely that they support integration with the RPKI.

ASPA issuers should be aware of the validation implication in issuing an ASPA - an ASPA implicitly invalidates all routes passed to upstream providers other than the provider ASs listed in the ASPA record. It is the Customer AS's duty to maintain a correct set of providers in ASPA record(s).

While it's not restricted, but it's highly recommended maintaining for selected Customer AS a single ASPA object that covers all its providers. Such policy should prevent race conditions during ASPA updates that might affect prefix propagation. The software that provides hosting for ASPA records SHOULD support enforcement of this rule. In the case of the transition process between different CA registries, the ASPA records SHOULD be kept identical in all registries.

While the ASPA is able to detect both mistakes and malicious activity for routes received from customers, RS-clients, or peers, it provides only detection of mistakes for routes that are received from upstream providers and RS(s).

Since an upstream provider becomes a trusted point, it will be able to send hijacked prefixes of its customers or send hijacked prefixes with malformed AS_PATHs back. While it may happen in theory, it's doesn't seem to be a real scenario: normally customer and provider have a signed agreement and such policy violation should have legal consequences or customer can just drop relation with such a provider and remove the corresponding ASPA record.

10. Acknowledgments

The authors wish to thank authors of [[RFC6483](#)] since its text was used as an example while writing this document. The authors wish to thank Iljitsch van Beijnum for giving a hint about Downstream paths. Authors wish to thank Kotikalapudi Sriram for algorithm improvements and helping with text clarity in the document.

11. References

11.1. 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>>.

11.2. Informative References

[Flexsealing]

McDaniel, T., Smith, J., and M. Schuchard, "Flexsealing BGP Against Route Leaks: Peerlock Active Measurement and Analysis", November 2020, <<https://arxiv.org/pdf/2006.06576.pdf>>.

[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-00, 19 April 2019, <<http://www.ietf.org/internet-drafts/draft-ietf-grow-route-leak-detection-mitigation-00.txt>>.

[I-D.ietf-sidrops-aspa-profile]

Azimov, A., Uskov, E., Bush, R., Patel, K., Snijders, J., and R. Housley, "A Profile for Autonomous System Provider Authorization", Work in Progress, Internet-Draft, draft-ietf-sidrops-aspa-profile-00, 17 May 2019, <<http://www.ietf.org/internet-drafts/draft-ietf-sidrops-aspa-profile-00.txt>>.

[I-D.kumari-deprecate-as-set-confed-set]

Kumari, W. and K. Sriram, "Deprecation of AS_SET and AS_CONFED_SET in BGP", Work in Progress, Internet-Draft, draft-kumari-deprecate-as-set-confed-set-12, 2 July 2018, <<http://www.ietf.org/internet-drafts/draft-kumari-deprecate-as-set-confed-set-12.txt>>.

[I-D.white-sobgp-architecture]

White, R., "Architecture and Deployment Considerations for Secure Origin BGP (soBGP)", Work in Progress, Internet-Draft, draft-white-sobgp-architecture-02, 16 June 2006, <<http://www.ietf.org/internet-drafts/draft-white-sobgp-architecture-02.txt>>.

[Peerlock]

Snijders, J., "Peerlock", June 2016, <https://www.nanog.org/sites/default/files/Snijders_Everyday_Practical_Bgp.pdf>.

[RFC3779]

Lynn, C., Kent, S., and K. Seo, "X.509 Extensions for IP Addresses and AS Identifiers", RFC 3779, DOI 10.17487/RFC3779, June 2004, <<https://www.rfc-editor.org/info/rfc3779>>.

[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>>.

[RFC5280]

Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, DOI 10.17487/RFC5280, May 2008, <<https://www.rfc-editor.org/info/rfc5280>>.

[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>>.

[RFC6483]

Huston, G. and G. Michaelson, "Validation of Route Origination Using the Resource Certificate Public Key Infrastructure (PKI) and Route Origin Authorizations (ROAs)", RFC 6483, DOI 10.17487/RFC6483, February 2012, <<https://www.rfc-editor.org/info/rfc6483>>.

[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>>.

[RFC8205]

Lepinski, M., Ed. and K. Sriram, Ed., "BGPsec Protocol Specification", RFC 8205, DOI 10.17487/RFC8205, September 2017, <<https://www.rfc-editor.org/info/rfc8205>>.

[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>>.

Authors' Addresses

Alexander Azimov
Yandex

Email: a.e.azimov@gmail.com

Eugene Bogomazov

Qrator Labs

Email: eb@qrator.net

Randy Bush

Internet Initiative Japan & Arrcus

Email: randy@psg.com

Keyur Patel

Arrcus, Inc.

Email: keyur@arrcus.com

Job Snijders

Fastly

Amsterdam

Email: job@fastly.com