Internet Engineering Task Force (IETF)
Request for Comments: 9324
Updates: 8481
Category: Standards Track
ISSN: 2070-1721

R. Bush
IIJ Research Lab & Arrcus, Inc.
K. Patel
Arrcus, Inc.
P. Smith
PFS Internet Development Pty Ltd
M. Tinka
SEACOM

December 2022

Policy Based on the Resource Public Key Infrastructure (RPKI) without Route Refresh

Abstract

A BGP speaker performing policy based on the Resource Public Key Infrastructure (RPKI) should not issue route refresh to its neighbors because it has received new RPKI data. This document updates RFC 8481 by describing how to avoid doing so by either keeping a full Adj-RIB-In or saving paths dropped due to ROV (Route Origin Validation) so they may be reevaluated with respect to new RPKI data.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at https://www.rfc-editor.org/info/rfc9324.

Table of Contents

1. Introduction
   1.1. Requirements Language
2. Related Work
3. ROV Experience
4. Keeping Partial Adj-RIB-In Data
5. Operational Recommendations
6. Security Considerations
7. IANA Considerations

8. References
8.1. Normative References
8.2. Informative References
Acknowledgements
Authors' Addresses

1. Introduction

Memory constraints in early BGP speakers caused classic BGP implementations [RFC4271] to not keep a full Adj-RIB-In (Section 1.1 of [RFC4271]). When doing RPKI-based Route Origin Validation (ROV) [RFC6811] [RFC8481] and similar RPKI-based policy, if such a BGP speaker receives new RPKI data, it might not have kept paths previously marked as Invalid, etc. Such an implementation must then request a route refresh [RFC2918] [RFC7313] from its neighbors to recover the paths that might be covered by these new RPKI data. This will be perceived as rude by those neighbors as it passes a serious resource burden on to them. This document recommends implementations keep and mark paths affected by RPKI-based policy, so route refresh is no longer needed.

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. Related Work

It is assumed that the reader understands BGP [RFC4271], route refresh [RFC7313], the RPKI [RFC6480], Route Origin Authorizations (ROAs) [RFC6482], the Resource Public Key Infrastructure (RPKI) to Router Protocol [RPKI-ROUTER-PROT-v2], RPKI-Based Prefix Validation [RFC6811], and Origin Validation Clarifications [RFC8481].

Note that the term "RPKI-based Route Origin Validation" in this document means the same as the term "Prefix Origin Validation" used in [RFC6811].

3. ROV Experience

As Route Origin Validation dropping Invalids has deployed, some BGP speaker implementations have been found that, when receiving new RPKI data (Validated ROA Payloads (VRPs) [RPKI-ROUTER-PROT-v2]), issue a BGP route refresh [RFC7313] to all sending BGP peers so that they can reevaluate the received paths against the new data.

In actual deployment, this has been found to be very destructive, transferring a serious resource burden to the unsuspecting peers. In reaction, RPKI-based Route Origin Validation (ROV) has been turned off. There have been actual de-peerings.

As RPKI registration and ROA creation have steadily increased, this problem has increased, not just proportionally, but on the order of the in-degree of ROV implementing BGP speakers. As Autonomous System Provider Authorization (ASPA) [AS_PATH-VER] becomes used, the problem will increase.

Other mechanisms, such as automated policy provisioning, which have flux rates similar to ROV (i.e., on the order of minutes), could very well cause similar problems.

Therefore, this document updates [RFC8481] by describing how to avoid this problem.

4. Keeping Partial Adj-RIB-In Data

If new RPKI data arrive that cause operator policy to invalidate the best route and the BGP speaker did not keep the dropped routes, then the BGP speaker would issue a route refresh, which this feature aims to prevent.

A route that is dropped by operator policy due to ROV is, by nature, considered ineligible to compete for the best route and MUST be kept in the Adj-RIB-In for potential future evaluation.

Ameliorating the route refresh problem by keeping a full Adj-RIB-In can be a problem for resource-constrained BGP speakers. In reality, only some data need be retained. If an implementation chooses not to retain the full Adj-RIB-In, it MUST retain at least routes dropped due to ROV for potential future evaluation.

As storing these routes could cause problems in resource-constrained devices, there MUST be a global operation, CLI, YANG, or other mechanism that allows the operator to enable this feature and store the dropped routes. Such an operator control MUST NOT be per peer, as this could cause inconsistent behavior.

As a side note, policy that may drop routes due to RPKI-based checks such as ROV (and ASPA, BGPsec [RFC8205], etc., in the future) MUST be run and the dropped routes saved per this section, before non-RPKI policies are run, as the latter may change path attributes.

5. Operational Recommendations

Operators deploying ROV and/or other RPKI-based policies should ensure that the BGP speaker implementation is not causing route refresh requests to neighbors.

BGP speakers MUST either keep the full Adj-RIB-In or implement the specification in Section 4. Conformance to this behavior is an additional, mandatory capability for BGP speakers performing ROV.

If the BGP speaker does not implement these recommendations, the operator should enable the vendor’s control to keep the full Adj-RIB-In, sometimes referred to as "soft reconfiguration inbound". The operator should then measure to ensure that there are no unnecessary route refresh requests sent to neighbors.

If the BGP speaker’s equipment has insufficient resources to support either of the two proposed options (keeping a full AdjRibIn or at least the dropped routes), the equipment SHOULD either be replaced with capable equipment or SHOULD NOT be used for ROV.

The configuration setting in Section 4 should only be used in very well-known and controlled circumstances where the scaling issues are well understood and anticipated.

Operators using the specification in Section 4 should be aware that a misconfigured neighbor might erroneously send a massive number of paths, thus consuming a lot of memory. Hence, pre-policy filtering such as described in [MAXPREFIX-INBOUND] could be used to reduce this exposure.

If route refresh has been issued toward more than one peer, the order of receipt of the refresh data can cause churn in both best route selection and outbound signaling.

Internet Exchange Points (IXPs) that provide route servers [RFC7947] should be aware that some members could be causing an undue route

refresh load on the route servers and take appropriate administrative and/or technical measures. IXPs using BGP speakers as route servers should ensure that they are not generating excessive route refresh requests.

6. Security Considerations

This document describes a denial of service that Route Origin Validation or other RPKI policy may place on a BGP neighbor and describes how it may be ameliorated.

Otherwise, this document adds no additional security considerations to those already described by the referenced documents.

7. IANA Considerations

This document has no IANA actions.

8. References

8.1. 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>.

[RFC2918] Chen, E., "Route Refresh Capability for BGP-4", RFC 2918, DOI 10.17487/RFC2918, September 2000, <https://www.rfc-editor.org/info/rfc2918>.

[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>.

[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>.

[RFC7313] Patel, K., Chen, E., and B. Venkatachalapathy, "Enhanced Route Refresh Capability for BGP-4", RFC 7313, DOI 10.17487/RFC7313, July 2014, <https://www.rfc-editor.org/info/rfc7313>.

[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>.

[RFC8481] Bush, R., "Clarifications to BGP Origin Validation Based on Resource Public Key Infrastructure (RPKI)", RFC 8481, DOI 10.17487/RFC8481, September 2018, <https://www.rfc-editor.org/info/rfc8481>.

8.2. Informative References

[AS_PATH-VER] Azimov, A., Bogomazov, E., Bush, R., Patel, K., Snijders, J., and K. Sriram, "BGP AS_PATH Verification Based on Resource Public Key Infrastructure (RPKI) Autonomous System Provider Authorization (ASPA) Objects", Work in Progress, Internet-Draft, draft-ietf-sidrops-aspa-verification-11, 24 October 2022, <https://datatracker.ietf.org/doc/html/draft-ietf-sidrops-aspa-verification-11>.

[MAXPREFIX-INBOUND]
Aelmans, M., Stucchi, M., and J. Snijders, "BGP Maximum Prefix Limits Inbound", Work in Progress, Internet-Draft, draft-sas-idr-maxprefix-inbound-04, 19 January 2022, <https://datatracker.ietf.org/doc/html/draft-sas-idr-maxprefix-inbound-04>.

[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>.

[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>.

[RFC7947] Jasinska, E., Hilliard, N., Raszuk, R., and N. Bakker, "Internet Exchange BGP Route Server", RFC 7947, DOI 10.17487/RFC7947, September 2016, <https://www.rfc-editor.org/info/rfc7947>.

[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>.

[RPKI-ROUTER-PROT-v2]
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>.

Acknowledgements

The authors wish to thank Alvaro Retana, Ben Maddison, Derek Yeung, John Heasley, John Scudder, Matthias Waehlisch, Nick Hilliard, Saku Ytti, and Ties de Kock.

Authors’ Addresses

Randy Bush
IIJ Research Lab & Arrcus, Inc.
1856 SW Edgewood Dr
Portland, OR 97210
United States of America
Email: randy@psg.com

Keyur Patel
Arrcus, Inc.
2077 Gateway Place, Suite #400
San Jose, CA 95119
United States of America
Email: keyur@arrcus.com

Philip Smith
PFS Internet Development Pty Ltd
PO Box 1908
Milton QLD 4064
Australia
Email: pfsinoz@gmail.com

Mark Tinka

SEACOM
Building 7, Design Quarter District
Leslie Avenue, Magaliessig
Fourways, Gauteng
2196
South Africa
Email: mark@tinka.africa