

## RPKI-Based Policy Without Route Refresh

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. Related Work](#)
- [3. ROV Experience](#)
- [4. Keeping Partial Adj-RIB-In Data](#)
- [5. Operational Recommendations](#)
- [6. Security Considerations](#)
- [7. IANA Considerations](#)
- [8. Acknowledgements](#)
- [9. References](#)
  - [9.1. Normative References](#)
  - [9.2. Informative References](#)
- [Authors' Addresses](#)

## 1. Introduction

Memory constraints in early routers caused classic [\[RFC4271\]](#) BGP implementations to not keep a full Adj-RIB-In (Sec. 1.1). When doing RPKI-based Route Origin Validation ([\[RFC6811\]](#) and [\[RFC8481\]](#)), 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 [\[RFC7313\]](#) from its neighbors to recover the paths which 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 but mark paths affected by RPKI-based policy so Route Refresh is no longer needed.

## 2. Related Work

It is assumed that the reader understands BGP, [\[RFC4271\]](#) and Route Refresh [\[RFC7313\]](#), the RPKI [\[RFC6480\]](#), Route Origin Authorizations (ROAs), [\[RFC6482\]](#), The Resource Public Key Infrastructure (RPKI) to Router Protocol [\[I-D.ietf-sidrops-8210bis\]](#), RPKI-based Prefix Validation, [\[RFC6811\]](#), and Origin Validation Clarifications, [\[RFC8481\]](#).

### 3. ROV Experience

As Route Origin Validation dropping Invalids has deployed, some router implementations have been found which, when receiving new RPKI data (VRPs, see [[I-D.ietf-sidrops-8210bis](#)]) issue a BGP Route Refresh [[RFC7313](#)] to all sending BGP peers so that it 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; and 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 routers. As ASPA ([[I-D.ietf-sidrops-aspa-verification](#)]) becomes used, the problem will increase.

### 4. Keeping Partial Adj-RIB-In Data

Ameliorating this problem by keeping a full Adj-RIB-In can be a problem for resource constrained routers. In reality, only some data need be retained.

When RPKI data cause one or more paths to be dropped, withdrawn, or merely not chosen as best path due to RPKI-based policy (ROV, ASPA, etc.), those paths MUST be saved and marked (to not be used for best path evaluation etc.) so that later VRPs can reevaluate them against then current policy.

Policy which may drop paths due to RPKI-based checks such as ROV, ASPA, BGPsec, etc. MUST be run, and the dropped paths saved per the above paragraph, before non-RPKI policies are run, as the latter may change path attributes.

As storing these paths could cause problems in resource constrained devices, there MUST be a knob allowing operator control of this feature. Such a knob MUST NOT be per peer, as this could cause inconsistent behavior.

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 path selection and in outbound signaling.

### 5. Operational Recommendations

Routers MUST either keep the full Adj-RIB-In or implement the specification in [Section 4](#).

Operators deploying ROV and/or other RPKI based policies SHOULD ensure that the router implementation is not causing unnecessary Route Refresh requests to neighbors.

If the router does not implement these recommendations, the operator SHOULD enable the vendor's knob 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 router has insufficient resources to support this, it MUST not be used for Route Origin Validation. I.e. the knob in [Section 4](#) should only be used in very well known and controlled circumstances.

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. Pre-policy filtering such as described in [[I-D.sas-idr-maxprefix-inbound](#)] SHOULD be used to reduce this exposure.

Internet Exchange Points (IXPs) which provide [[RFC7947](#)] Route Servers 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 routers as route servers should ensure that they are not generating excessive route refresh requests.

## **6. Security Considerations**

This document describes a denial of service which 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**

None

## **8. Acknowledgements**

The authors wish to thank Ben Maddison, John Heasley, Nick Hilliard, John Scudder, Matthias Waehlich, and Saku Ytti.

## **9. References**

### **9.1. Normative References**

[[I-D.sas-idr-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://www.ietf.org/archive/id/draft-sas-idr-maxprefix-inbound-04.txt>>.

- [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>>.
- [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>>.
- [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>>.

## 9.2. Informative References

- [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-05, 22 December 2021, <<https://www.ietf.org/archive/id/draft-ietf-sidrops-8210bis-05.txt>>.
- [I-D.ietf-sidrops-aspa-verification]  
Azimov, A., Bogomazov, E., Bush, R., Patel, K., and J. Snijders, "Verification of AS\_PATH Using the Resource Certificate Public Key Infrastructure and Autonomous System Provider Authorization", Work in Progress, Internet-Draft, draft-ietf-sidrops-aspa-verification-08, 25 August 2021, <<https://www.ietf.org/archive/id/draft-ietf-sidrops-aspa-verification-08.txt>>.
- [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>>.

[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>>.

[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>>.

[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>>.

#### Authors' Addresses

Randy Bush  
IIJ Research Lab & Arrcus, Inc.  
1856 SW Edgewood Dr  
Portland, Oregon 97210  
United States of America

Email: [randy@psg.com](mailto:randy@psg.com)

Keyur Patel  
Arrcus, Inc.  
2077 Gateway Place, Suite #400  
San Jose, CA 95119  
United States of America

Email: [keyur@arrcus.com](mailto:keyur@arrcus.com)

Philip Smith  
PFS Internet Development Pty Ltd  
PO Box 1908  
Milton QLD 4064  
Australia

Email: [pfsinoz@gmail.com](mailto:pfsinoz@gmail.com)

Mark Tinka  
SEACOM  
Building 7, Design Quarter District, Leslie Avenue, Magaliessig  
Fourways, Gauteng  
2196  
South Africa

Email: [mark@tinka.africa](mailto:mark@tinka.africa)