

Network Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: October 12, 2017

T. King  
D. Kopp  
DE-CIX  
A. Lambrianidis  
AMS-IX  
A. Fenoux  
France-IX  
April 10, 2017

Signaling Prefix Origin Validation Results from a Route Server to Peers  
draft-ietf-sidrops-route-server-rpki-light-02

Abstract

This document defines the usage of the BGP Prefix Origin Validation State Extended Community [RFC8097] to signal prefix origin validation results from a route server to its peers. Upon reception of prefix origin validation results peers can use this information in their local routing decision process.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in [RFC2119] only when they appear in all upper case. They may also appear in lower or mixed case as English words, without normative meaning.

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 October 12, 2017.

## Copyright Notice

Copyright (c) 2017 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 (<http://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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

1. Introduction . . . . .	2
2. BGP Prefix Origin Validation State Utilized at Route-Servers	3
3. Signaling Prefix Origin Validation Results from a Route Server to Peers . . . . .	4
4. Operational Recommendations . . . . .	4
4.1. Local Routing Decision Process . . . . .	4
4.2. Route Server Receiving the BGP Prefix Origin Validation State Extended Community . . . . .	4
4.3. Information about Validity of a BGP Prefix Origin Not Available at a Route-Server . . . . .	5
4.4. Error Handling at Peers . . . . .	5
5. IANA Considerations . . . . .	5
6. Security Considerations . . . . .	5
7. References . . . . .	6
7.1. Normative References . . . . .	6
7.2. Informative References . . . . .	6
Authors' Addresses . . . . .	7

## 1. Introduction

RPKI-based prefix origin validation [RFC6480] can be a significant operational burden for BGP peers to implement and adopt. In order to boost acceptance and usage of prefix origin validation and ultimately increase the security of the Internet routing system, IXPs may provide RPKI-based prefix origin validation at the route server [RFC7947]. The result of this prefix origin validation is signaled to peers by using the BGP Prefix Origin Validation State Extended Community as introduced in [RFC8097].

Peers receiving the prefix origin validation result from the route server(s) can use this information in their local routing decision

process for acceptance, rejection, preference, or other traffic engineering purposes of a particular route.

## 2. BGP Prefix Origin Validation State Utilized at Route-Servers

A route server that is aware of a BGP Prefix Origin Validation state for a certain route can handle this information in one of the following modes of operation:

**Simple Tagging:** The prefix origin validation state is tagged to the route as described in Section 3.

This mode of operation is like the traditional way route servers work, however, the prefix origin validation state information is additionally available for peers.

**Dropping and Tagging:** Routes for which the prefix origin validation state is "invalid" (according to [RFC6811]) are dropped by the route server. Routes which show a prefix origin validation state of "not found" and "valid" (according to [RFC6811]) are tagged accordingly to Section 3.

Security is higher rated than questionable reachability of a prefix by this mode of operation.

**Prioritizing and Tagging:** If the route server learned for a particular prefix more than one route it removes firstly the set of "invalid" routes and secondly the "not found" routes unless the set of routes is empty. Based on the set of routes left over the BGP best path section algorithm is executed. The selected route is marked accordingly to Section 3.

The BGP best path selection algorithm is changed by this mode of operation in such a way that "valid" routes are preferred even if they are unfavorable by the traditional best path selection algorithm. This puts prefix origin validation on top of the best path selection.

A route server MUST support the Simple Tagging mode of operation. Other modes of operation are OPTIONAL. The mode of operation MAY be configured by the route server operator for a route server instance or for each BGP session with a peer separately.

These mode of operations might be used in combination with [RFC7911] in order to allow a peer to receive all routes and take the routing decision by itself.

### 3. Signaling Prefix Origin Validation Results from a Route Server to Peers

The BGP Prefix Origin Validation State Extended Community (as defined in [RFC8097]) is utilized for signaling prefix origin validation result from a route server to peers.

[RFC8097] proposes an encoding of the prefix origin validation result [RFC6811] as follows:

Value	Meaning
0	Lookup result = "valid"
1	Lookup result = "not found"
2	Lookup result = "invalid"

Table 1

This encoding is re-used. Route servers providing RPKI-based prefix origin validation set the validation state according to the prefix origin validation result (see [RFC6811]).

## 4. Operational Recommendations

### 4.1. Local Routing Decision Process

A peer receiving prefix origin validation results from the route server MAY use the information in its own local routing decision process. The local routing decision process SHOULD apply to the rules as described in section 5 [RFC6811].

A peer receiving a prefix origin validation result from the route server MAY redistribute this information within its own AS.

### 4.2. Route Server Receiving the BGP Prefix Origin Validation State Extended Community

An IXP route server receiving routes from its peers containing the BGP Prefix Origin Validation State Extended Community MUST remove the extended community before the route is re-distributed to its peers. This is required regardless of whether the route server is executing prefix origin validation or not.

Failure to do so would allow opportunistic peers to advertise routes tagged with arbitrary prefix origin validation results via a route

server, influencing maliciously the decision process of other route server peers.

#### 4.3. Information about Validity of a BGP Prefix Origin Not Available at a Route-Server

In case information about the validity of a BGP prefix origin is not available at the route server (e.g., error in the ROA cache, CPU overload) the route server MUST NOT add the BGP Prefix Origin Validation State Extended Community to the route.

#### 4.4. Error Handling at Peers

A route sent by a route server SHOULD only contain none or one BGP Prefix Origin Validation State Extended Community.

A peer receiving a route from a route server containing more than one BGP Prefix Origin Validation State Extended Community SHOULD only consider the largest value (as described in Table 1) in the validation result field and disregard the other values. Values larger than two in the validation result field MUST be disregarded.

#### 5. IANA Considerations

None.

#### 6. Security Considerations

All security considerations described in RFC 6811 [RFC6811] fully apply to this document.

Additionally, threat agents polluting ROA cache server(s) run by IXP operators could cause significant operational impact, since multiple route server clients could be affected. Peers should be vigilant as to the integrity and authenticity of the origin validation results, as they are provided by a third party, namely the IXP operator hosting both the route server as well as any ROA cache server(s).

Therefore, a route server could be misused to spread malicious prefix origin validation results. However, peers already trust the route server for the collection, filtering (e.g. IRR database filtering), and redistribution of BGP routing information to other peers. So, no change in the trust level is needed for this proposal.

To facilitate trust and help with peers establishing appropriate controls in mitigating the risks mentioned above, IXPs SHOULD provide out-of-band means for peers to ensure that the ROA validation process has not been compromised or corrupted.

While being under DDoS attacks, it is a common practice for peers connected to an IXP to make use of blackholing services (see [RFC7999]). Peers are using blackholing to drop traffic, typically by announcing a more specific prefix, which is under attack. A peer SHOULD make sure that this prefix is covered by an appropriate ROA.

## 7. References

### 7.1. 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, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", RFC 4360, DOI 10.17487/RFC4360, February 2006, <<http://www.rfc-editor.org/info/rfc4360>>.
- [RFC6811] Mohapatra, P., Scudder, J., Ward, D., Bush, R., and R. Austein, "BGP Prefix Origin Validation", RFC 6811, DOI 10.17487/RFC6811, January 2013, <<http://www.rfc-editor.org/info/rfc6811>>.
- [RFC7911] Walton, D., Retana, A., Chen, E., and J. Scudder, "Advertisement of Multiple Paths in BGP", RFC 7911, DOI 10.17487/RFC7911, July 2016, <<http://www.rfc-editor.org/info/rfc7911>>.
- [RFC8097] Mohapatra, P., Patel, K., Scudder, J., Ward, D., and R. Bush, "BGP Prefix Origin Validation State Extended Community", RFC 8097, DOI 10.17487/RFC8097, March 2017, <<http://www.rfc-editor.org/info/rfc8097>>.

### 7.2. Informative References

- [RFC6480] Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", RFC 6480, DOI 10.17487/RFC6480, February 2012, <<http://www.rfc-editor.org/info/rfc6480>>.
- [RFC7947] Jasinska, E., Hilliard, N., Raszuk, R., and N. Bakker, "Internet Exchange BGP Route Server", RFC 7947, DOI 10.17487/RFC7947, September 2016, <<http://www.rfc-editor.org/info/rfc7947>>.

[RFC7999] King, T., Dietzel, C., Snijders, J., Doering, G., and G.  
Hankins, "BLACKHOLE Community", RFC 7999,  
DOI 10.17487/RFC7999, October 2016,  
<<http://www.rfc-editor.org/info/rfc7999>>.

Authors' Addresses

Thomas King  
DE-CIX Management GmbH  
Lichtstrasse 43i  
Cologne 50825  
DE

Email: [thomas.king@de-cix.net](mailto:thomas.king@de-cix.net)

Daniel Kopp  
DE-CIX Management GmbH  
Lichtstrasse 43i  
Cologne 50825  
DE

Email: [daniel.kopp@de-cix.net](mailto:daniel.kopp@de-cix.net)

Aristidis Lambrianidis  
Amsterdam Internet Exchange  
Frederiksplein 42  
Amsterdam 1017 XN  
NL

Email: [aristidis.lambrianidis@ams-ix.net](mailto:aristidis.lambrianidis@ams-ix.net)

Arnaud Fenioux  
France-IX  
88 Avenue Des Ternes  
Paris 75017  
FR

Email: [afenioux@franceix.net](mailto:afenioux@franceix.net)