

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 3, 2020

R. Bush
Internet Initiative Japan & Arrcus
K. Patel
Arrcus
July 2, 2019

Origin Validation Signaling
draft-ymbk-sidrops-ov-signal-03

Abstract

Within a trust boundary, e.g. an operator's PoP, it may be useful to have only a few central devices do full Origin Validation using the Resource Public Key Infrastructure, and be able to signal to an internal sender that a received route fails Origin Validation. E.g. route reflectors could perform Origin Validation for a cluster and signal back to a sending client that it sent an invalid route. Routers capable of sending and receiving this signal can use the extended community described in [RFC8097].

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 <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 January 3, 2020.

Copyright Notice

Copyright (c) 2019 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 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.

1. Introduction

Within a routing trust boundary, e.g. an operator's Point of Presence (PoP), it may not be desirable or necessary for all routers to perform Origin Validation using the Resource Public Key Infrastructure (RPKI) per [RFC6811]. A good example is route reflectors (see [RFC4456]).

An RPKI-enabled device, an Evaluator, SHOULD signal receipt of an Invalid route back to the sender by announcing that route back to the sender marked with the BGP Prefix Origin Validation State Extended Community as defined in [RFC8097] with a last octet having the value 2, meaning "Invalid." In the rest of this document we take the liberty of calling it the "community."

We use the term "Sender" to refer to the router announcing routes to the device evaluating the Origin Validation of the announcements. Beware that the Sender receives signaling back from the Evaluator, which can be somewhat confusing.

We use the term "Evaluator" to describe the device receiving routing announcements from senders, applying RPKI-based Origin Validation, and possibly signaling route Invalidity back to the sender(s).

2. Suggested Reading

It is assumed that the reader understands BGP, [RFC4271], the RPKI, [RFC6480], RPKI-based Prefix Validation, [RFC6811], and the BGP Prefix Origin Validation State Extended Community as described in [RFC8097].

3. Trust Boundary

As a general rule, we discourage 'outsourcing trust,' i.e. letting others make security decisions for us. But there are operational environments with a somewhat wide trust boundary, a single operator's PoP for example.

This is not outsourcing trust; this is remote decision making. It is not letting a third party make the decision; it is simply doing it on a different computer. It's trust in a distributed system, where what is (sometimes) called the Policy Decision Point is not the same as the Policy Enforcement Point.

As described in [RFC7115], a PoP might have a single RPKI Cache, hence all trust is vested in it. So it is reasonable that routers in that PoP could share Origin Validation results instead of each doing full validation.

An [RFC4456] Route Reflector Cluster is an obvious candidate for this approach. The route reflector(s) would perform Origin Validation and signal an Invalid route back to the sending client.

[RFC8097] provides the obvious signaling mechanism, the BGP Prefix Origin Validation State Extended Community. The device performing OV SHOULD signal back to the sender by announcing the offending prefix marked with the extended community with the last octet having the value 2, indicating an Invalid route.

4. The OV Signaling Capability

Unfortunately, the router sending the Invalid announcement is not normally expecting to receive it back. Therefore, both parties MUST agree on this feature by using a BGP Capability [RFC5492].

To advertise the OV Signaling Capability to a peer, a BGP speaker uses BGP Capabilities Advertisement [RFC5492]. By advertising the OV Signaling Capability to a peer, a BGP speaker conveys that it is able to send, receive, and properly handle OV Signaling using the community.

A peer which does not advertise this capability MUST NOT send OV Signaling, and BGP OV Signaling MUST NOT be sent to it.

The OV Signaling Capability is a new BGP Capability defined with Capability code [TBD] and Capability length 0.

5. Recommended Action

This section assumes that the OV Signaling Capability has been negotiated by the sending and receiving routers.

An Evaluating device which performs Origin Validation on a route received from a capable sender and finds a prefix with a particular origin AS to be Invalid (in the [RFC6811] sense), MUST announce that prefix back to the sending router from which it was received with the Invalid origin AS and the addition of the community with the last octet being 2.

A sender receiving the returned prefix announcement so marked MUST treat it the way it would treat an Invalid origin that it itself detected. It should withdraw all routes it had announced to that prefix with the Invalid origin AS. This includes withdrawing any instances of additional paths with that origin AS advertised under [RFC7911].

For a sender to properly evaluate the community returned by the evaluator, the sender MUST recognize the community before loop detection. This is a change to the Phase 2 Route Selection process of [RFC4271] Section 9.1.2.

If a sender originally received the Invalid route from an evaluator within its trust boundary with which it has negotiated the OV Signaling Capability, it MAY also propagate that signal to the original sender.

6. Security Considerations

As with all communities which cause semantic change, this use of the community may be abused as an attack vector. Therefore the operator MUST configure their incoming external border to strip the community.

As the BGP sessions are already established using whatever channel security the operator chooses or not, this change specifies no additional channel or object security. Of course, the BGP transport should be protected for integrity and authentication. TCP-MD5 [RFC2385] is available on almost all platforms. If more modern methods are available, they should be used.

Outsourcing security is usually considered bad policy. Section Section 3 above discusses why that is not really the case here.

Otherwise, this document does not create security considerations beyond those of [RFC6811].

7. IANA Considerations

This document requests the IANA assign the "OV Signaling Capability" to the BGP Capabilities described in Section 2.1 in the "Capability Codes" registry's "IETF Review" range [RFC8126].. This document is the reference for the new capability.

8. Acknowledgments

Thanks to Steve Bellovin for a serious security review, and Rob Austein for a useful security snark.

9. References

9.1. 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>>.
- [RFC2385] Heffernan, A., "Protection of BGP Sessions via the TCP MD5 Signature Option", RFC 2385, DOI 10.17487/RFC2385, August 1998, <<http://www.rfc-editor.org/info/rfc2385>>.
- [RFC4456] Bates, T., Chen, E., and R. Chandra, "BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)", RFC 4456, DOI 10.17487/RFC4456, April 2006, <<http://www.rfc-editor.org/info/rfc4456>>.
- [RFC5492] Scudder, J. and R. Chandra, "Capabilities Advertisement with BGP-4", RFC 5492, DOI 10.17487/RFC5492, February 2009, <<http://www.rfc-editor.org/info/rfc5492>>.
- [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>>.
- [RFC7115] Bush, R., "Origin Validation Operation Based on the Resource Public Key Infrastructure (RPKI)", BCP 185, RFC 7115, DOI 10.17487/RFC7115, January 2014, <<http://www.rfc-editor.org/info/rfc7115>>.
- [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>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<http://www.rfc-editor.org/info/rfc8126>>.

9.2. Informative References

- [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, <<http://www.rfc-editor.org/info/rfc4271>>.
- [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>>.

Authors' Addresses

Randy Bush
Internet Initiative Japan & Arrcus
5147 Crystal Springs
Bainbridge Island, Washington 98110
US

Email: randy@psg.com

Keyur Patel
Arrcus
2077 Gateway Place, Suite #250
San Jose, CA 95119
United States of America

Email: keyur@arrcus.com