

Internet Engineering Task Force (IETF)
Request for Comments: 8207
BCP: 211
Category: Best Current Practice
ISSN: 2070-1721

R. Bush
Internet Initiative Japan
September 2017

BGPsec Operational Considerations

Abstract

Deployment of the BGPsec architecture and protocols has many operational considerations. This document attempts to collect and present the most critical and universal. Operational practices are expected to evolve as BGPsec is formalized and initially deployed.

Status of This Memo

This memo documents an Internet Best Current Practice.

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 BCPS 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 <http://www.rfc-editor.org/info/rfc8207>.

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
- 1.1. Requirements Language 3
- 2. Suggested Reading 3
- 3. RPKI Distribution and Maintenance 3
- 4. AS/Router Certificates 3
- 5. Within a Network 4
- 6. Considerations for Edge Sites 4
- 7. Routing Policy 5
- 8. Notes 7
- 9. Security Considerations 7
- 10. IANA Considerations 7
- 11. References 8
- 11.1. Normative References 8
- 11.2. Informative References 8
- Acknowledgements 10
- Author’s Address 10

1. Introduction

Origin validation based on the Resource Public Key Infrastructure (RPKI) [RFC6811] is in its early phases. As BGPsec [RFC8205] may require larger memory and/or more modern CPUs, it is expected to be deployed incrementally over a longer time span. BGPsec is a new protocol with many operational considerations that this document attempts to describe. As with most operational practices, they will likely change over time.

BGPsec relies on widespread propagation of the RPKI [RFC6480]. How the RPKI is distributed and maintained globally and within an operator’s infrastructure may be different for BGPsec than for origin validation.

BGPsec needs to be spoken only by an Autonomous System’s (AS’s) eBGP-speaking border routers. It is designed so that it can be used to protect announcements that are originated by resource-constrained edge routers. This has special operational considerations, see Section 6.

Different prefixes may have different timing and replay protection considerations.

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. Suggested Reading

It is assumed that the reader understands BGP [RFC4271], BGPsec [RFC8205], the RPKI [RFC6480], the RPKI Repository Structure [RFC6481], and Route Origin Authorizations (ROAs) [RFC6482].

3. RPKI Distribution and Maintenance

The considerations for RPKI objects (Certificates, Certificate Revocation Lists (CRLs), manifests [RFC6481], and Ghostbusters Records [RFC6493]), Trust Anchor Locators (TALs) [RFC7730], cache behaviors of synchronization, and validation from the section on RPKI Distribution and Maintenance of [RFC7115] apply. Specific considerations relating to ROA objects do not apply to this document.

4. AS/Router Certificates

As described in [KEY], BGPsec-speaking routers are capable of generating their own public/private key-pairs and having their certificates signed and published in the RPKI by the RPKI Certification Authority (CA) system, and/or are given public/private key-pairs by the operator.

A site/operator may use a single certificate/key in all their routers, one certificate/key per router, or any granularity in between.

A large operator, concerned that a compromise of one router's key would make other routers vulnerable, may deploy a more complex certificate/key distribution burden to reduce this exposure.

At the other end of the spectrum, an edge site with one or two routers may choose to use a single certificate/key.

In anticipation of possible key compromise, a prudent operator SHOULD pre-provision each router's 'next' key in the RPKI so that there is no propagation delay for provisioning the new key.

5. Within a Network

BGPsec is spoken by edge routers in a network, specifically those that border other networks/ASes.

In an AS where edge routers speak BGPsec and, therefore, inject BGPsec paths into the iBGP (internal BGP), Route Reflectors (RRs) MUST have BGPsec enabled if and only if there are eBGP (external BGP) speakers in their client cone, i.e., an RR client or the transitive closure of a client's customers.

A BGPsec-capable router MAY use the data it receives to influence local policy within its network, see Section 7. In deployment, this policy should fit into the AS's existing policy, preferences, etc. This allows a network to incrementally deploy BGPsec-enabled border routers.

eBGP speakers that face more critical peers or upstreams or downstreams would be candidates for early deployment. Both securing one's own announcements and validating received announcements should be considered in partial deployment.

An operator should be aware that BGPsec, as any other policy change, can cause traffic shifts in their network. And, as with normal policy shift practice, a prudent operator has the tools and methods to predict, measure, modify, etc.

On the other hand, an operator wanting to monitor router loading, shifts in traffic, etc., might deploy incrementally while watching those and similar effects.

BGPsec does not sign over communities, so they are not formally trustable. Additionally, outsourcing verification is not a prudent security practice. Therefore, an eBGP listener SHOULD NOT strongly trust unsigned security signaling, such as communities, received across a trust boundary.

6. Considerations for Edge Sites

An edge site that does not provide transit and trusts its upstream(s) may only originate a signed prefix announcement and not validate received announcements.

An operator might need to use hardware with limited resources. In such cases, BGPsec protocol capability negotiation allows for a resource-constrained edge router to hold only its own signing key(s) and sign its announcements, but not receive signed announcements.

Therefore, the router would not have to deal with the majority of the RPKI, potentially saving the need for additional hardware.

As the vast majority of ASes are stubs, and they announce the majority of prefixes, this allows for simpler and less expensive incremental deployment. It may also mean that edge sites concerned with routing security will be attracted to upstreams that support BGPsec.

7. Routing Policy

As BGPsec-signed paths cannot traverse non-BGPsec topology, partial BGPsec deployment forms islands of assured paths. As islands grow to touch each other, they become larger islands.

Unlike origin validation based on the RPKI, BGPsec marks a received announcement as Valid or Not Valid, there is no explicit NotFound state. In some sense, an unsigned BGP4 path is the equivalent of NotFound. How this is used in routing is up to the operator's local policy, similar to origin validation as in [RFC6811].

As BGPsec will be rolled out over years and does not allow for intermediate non-signing edge routers, coverage will be spotty for a long time. This presents a dilemma; should a router evaluating an inbound BGPsec_PATH as Not Valid be very strict and discard it? On the other hand, it might be the only path to that prefix, and a very low local-preference would cause it to be used and propagated only if there was no alternative. Either choice is reasonable, but we recommend dropping because of the next point.

Operators should be aware that accepting Not Valid announcements, no matter the local preference, will often be the equivalent of treating them as fully Valid. Local preference affects only routes to the same set of destinations. Consider having a Valid announcement from neighbor V for prefix 10.0.0.0/16 and a Not Valid announcement for 10.0.666.0/24 from neighbor I. If local policy on the router is not configured to discard the Not Valid announcement from I, then the longest match forwarding will send packets to neighbor I no matter the value of local preference.

Validation of signed paths is usually deployed at the eBGP edge.

Local policy on the eBGP edge MAY convey the validation state of a BGP-signed path through normal local policy mechanisms, e.g., setting a BGP community for internal use, or modifying a metric value such as local-preference or Multi-Exit Discriminator (MED). Some may choose

to use the large Local-Pref hammer. Others may choose to let AS path rule and set their internal metric, which comes after AS path in the BGP decision process.

As the mildly stochastic timing of RPKI propagation may cause version skew across routers, an AS Path that does not validate at router R0 might validate at R1. Therefore, signed paths that are Not Valid and yet propagated (because they are chosen as best path) MUST NOT have signatures stripped and MUST be signed if sent to external BGPsec speakers.

This implies that updates which a speaker judges to be Not Valid MAY be propagated to iBGP peers. Therefore, unless local policy ensures otherwise, a signed path learned via iBGP may be Not Valid. If needed, the validation state should be signaled by normal local policy mechanisms such as communities or metrics.

On the other hand, local policy on the eBGP edge might preclude iBGP or eBGP announcement of signed AS Paths that are Not Valid.

A BGPsec speaker receiving a path SHOULD perform origin validation per [RFC6811] and [RFC7115].

A route server is usually 'transparent', i.e., does not insert an AS into the path so as not to increase the AS hop count, and thereby affect downstream path choices. But, with BGPsec, a client router R needs to be able to validate paths that are forward signed to R. But the sending router cannot generate signatures to all the possible clients. Therefore, a BGPsec-aware route server needs to validate the incoming BGPsec_PATH and to forward updates that can be validated by clients that must, therefore, know the route server's AS. This implies that the route server creates signatures per client including its own AS in the BGPsec_PATH, forward signing to each client AS, see [RFC8205]. The route server uses a pCount of 0 to not increase the effective AS hop count, thereby retaining the intent of 'transparency'.

If it is known that a BGPsec neighbor is a transparent route server, or otherwise may validly use a pCount of 0 (e.g., see [RFC8206]), the router SHOULD be configured to accept and process a received pCount of 0. Routers MUST disallow a pCount of 0 by default.

To prevent exposure of the internals of the BGP confederations [RFC5065], a BGPsec speaker exporting to a non-member removes all intra-confederation Secure_Path Segments. Therefore, signing within the confederation will not cause external confusion even if non-unique private ASes are used.

8. Notes

For protection from attacks replaying BGP data on the order of a day or longer old, rekeying routers with new keys (previously) provisioned in the RPKI is sufficient. For one approach, see [ROLLOVER].

A router that once negotiated (and/or sent) BGPsec should not be expected to always do so.

Like the DNS, the Global RPKI presents only a loosely consistent view, depending on timing, updating, fetching, etc. Thus, one cache or router may have different data about a particular prefix or router than another cache or router. There is no 'fix' for this, it is the nature of distributed data with distributed caches.

Operators who manage certificates SHOULD have RPKI Ghostbuster Records (see [RFC6493]), signed indirectly by end entity certificates, for those certificates on which others' routing depends for certificate and/or ROA validation.

Operators should be aware of impending algorithm transitions, which will be rare and slow-paced, see [RFC6916]. They should work with their vendors to ensure support for new algorithms.

As a router must evaluate certificates and ROAs that are time dependent, routers' clocks MUST be correct to a tolerance of approximately an hour. The common approach is for operators to deploy servers that provide time service, such as [RFC5905], to client routers.

If a router has reason to believe its clock is seriously incorrect, e.g., it has a time earlier than 2011, it SHOULD NOT attempt to validate incoming updates. It SHOULD defer validation until it believes it is within reasonable time tolerance.

9. Security Considerations

This document describes operational considerations for the deployment of BGPsec. The security considerations for BGPsec are described in [RFC8205].

10. IANA Considerations

This document does not require any IANA actions.

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>>.
- [RFC6493] Bush, R., "The Resource Public Key Infrastructure (RPKI) Ghostbusters Record", RFC 6493, DOI 10.17487/RFC6493, February 2012, <<https://www.rfc-editor.org/info/rfc6493>>.
- [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>>.
- [RFC7115] Bush, R., "Origin Validation Operation Based on the Resource Public Key Infrastructure (RPKI)", BCP 185, RFC 7115, DOI 10.17487/RFC7115, January 2014, <<https://www.rfc-editor.org/info/rfc7115>>.
- [RFC7730] Huston, G., Weiler, S., Michaelson, G., and S. Kent, "Resource Public Key Infrastructure (RPKI) Trust Anchor Locator", RFC 7730, DOI 10.17487/RFC7730, January 2016, <<https://www.rfc-editor.org/info/rfc7730>>.
- [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>>.
- [RFC8205] Lepinski, M., Ed. and K. Sriram, Ed., "BGPsec Protocol Specification", RFC 8205, DOI 10.17487/RFC8205, September 2017, <<http://www.rfc-editor.org/info/rfc8205>>.

11.2. Informative References

- [KEY] Bush, R., Turner, S., and K. Patel, "Router Keying for BGPsec", Work in Progress, draft-ietf-sidr-rtr-keying-13, April 2017.
- [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>>.

- [RFC5065] Traina, P., McPherson, D., and J. Scudder, "Autonomous System Confederations for BGP", RFC 5065, DOI 10.17487/RFC5065, August 2007, <<https://www.rfc-editor.org/info/rfc5065>>.
- [RFC5905] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, DOI 10.17487/RFC5905, June 2010, <<https://www.rfc-editor.org/info/rfc5905>>.
- [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>>.
- [RFC6481] Huston, G., Loomans, R., and G. Michaelson, "A Profile for Resource Certificate Repository Structure", RFC 6481, DOI 10.17487/RFC6481, February 2012, <<https://www.rfc-editor.org/info/rfc6481>>.
- [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>>.
- [RFC6916] Gagliano, R., Kent, S., and S. Turner, "Algorithm Agility Procedure for the Resource Public Key Infrastructure (RPKI)", BCP 182, RFC 6916, DOI 10.17487/RFC6916, April 2013, <<https://www.rfc-editor.org/info/rfc6916>>.
- [RFC8206] George, W. and S. Murphy, "BGPsec Considerations for Autonomous System (AS) Migration", RFC 8206, DOI 10.17487/RFC8206, September 2017, <<http://www.rfc-editor.org/info/rfc8206>>.
- [ROLLOVER] Weis, B., Gagliano, R., and K. Patel, "BGPsec Router Certificate Rollover", Work in Progress, draft-ietf-sidrps-bgpsec-rollover-02, August 2017.

Acknowledgements

The author wishes to thank Thomas King, Arnold Nipper, Alvaro Retana, and the BGPsec design group.

Author's Address

Randy Bush
Internet Initiative Japan
5147 Crystal Springs
Bainbridge Island, Washington 98110
United States of America

Email: randy@psg.com