

SIDR Working Group
Internet Draft
Intended status: Informational
Expires: December 24, 2016

M. Lepinski
NCF
S. Turner
sn3rd
June 22, 2016

An Overview of BGPsec
draft-ietf-sidr-bgpsec-overview-08

Abstract

This document provides an overview of a security extension to the Border Gateway Protocol (BGP) referred to as BGPsec. BGPsec improves security for BGP routing.

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 December 24, 2016.

Copyright Notice

Copyright (c) 2016 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.	Background	3
3.	BGPsec Operation	4
3.1.	Negotiation of BGPsec	4
3.2.	Update signing and validation	5
4.	Design and Deployment Considerations	7
4.1.	Disclosure of topology information	7
4.2.	BGPsec router assumptions	7
4.3.	BGPsec and consistency of externally visible data	8
5.	Security Considerations	8
6.	IANA Considerations	9
7.	References	9
7.1.	Normative References	9
7.2.	Informative References	10
	Authors' Addresses	10

[1.](#) Introduction

BGPsec (Border Gateway Protocol Security) is an extension to the Border Gateway Protocol (BGP) that provides improved security for BGP routing [[RFC4271](#)]. This document contains a brief overview of BGPsec and its envisioned usage.

A more detailed discussion of BGPsec is provided in the following set of documents:

- * [[RFC7132](#)]:

A threat model describing the security context in which BGPsec is intended to operate.

- * [[RFC7353](#)]:

A set of requirements for BGP path security, which BGPsec is intended to satisfy.

- * [[I-D.sidr-bgpsec-protocol](#)]:

A standards track document specifying the BGPsec extension to BGP.

- * [[I-D.sidr-as-migration](#)]:

A standards track document describing how to implement an AS Number migration while using BGPsec.

- * [\[I-D.sidr-bgpsec-ops\]](#):

An informational document describing operational considerations.

- * [\[I-D.sidr-bgpsec-pki-profiles\]](#):

A standards track document specifying a profile for X.509 certificates that bind keys used in BGPsec to Autonomous System numbers, as well as associated Certificate Revocation Lists (CRLs), and certificate requests.

- * [\[I-D.sidr-bgpsec-algs\]](#)

A standards track document specifying suites of signature and digest algorithms for use in BGPsec.

In addition to this document set, some readers might be interested in [\[I-D.sriram-bgpsec-design-choices\]](#), an informational document describing the choices that were made by the design team prior to the publication of the -00 version of [draft-ietf-sidr-bgpsec-protocol](#). Discussion of design choices made since the publication of the -00 can be found in the archives of the SIDR working group mailing list.

[2. Background](#)

The motivation for developing BGPsec is that BGP does not include mechanisms that allow an Autonomous System (AS) to verify the legitimacy and authenticity of BGP route advertisements (see for example, [\[RFC4272\]](#)).

The Resource Public Key Infrastructure (RPKI), described in [\[RFC6480\]](#), provides a first step towards addressing the validation of BGP routing data. RPKI resource certificates are issued to the holders of AS number and IP address resources, providing a binding between these resources and cryptographic keys that can be used to verify digital signatures. Additionally, the RPKI architecture specifies a digitally signed object, a Route Origination Authorization (ROA), that allows holders of IP address resources to authorize specific ASes to originate routes (in BGP) to these resources. Data extracted from a valid ROA can be used by a BGP speaker to determine whether the origin AS asserted in a received route has been authorized (by the Internet Number Resource holder) to originate that route (see [\[RFC6483\]](#) and [\[RFC7115\]](#)).

By instituting a local policy that prefers routes with origins validated using RPKI data (versus routes to the same prefix that cannot be so validated) an AS can protect itself from configuration

errors by network operators and from certain mis-origination attacks. However, use of RPKI data alone provides little or no protection against a sophisticated attacker. Such an attacker could, for example, conduct a route hijacking attack by appending an authorized origin AS to an otherwise illegitimate AS path. (See [[RFC7132](#)] for a detailed discussion of the BGPsec threat model.)

BGPsec extends the RPKI by adding an additional type of certificate, referred to as a BGPsec Router Certificate, that binds an AS number to a public signature verification key. The corresponding private key is held by one or more BGP speakers within this AS. Private keys corresponding to public keys in such certificates are used within BGPsec to enable a BGP speaker to sign on behalf of its AS. The certificates thus allow a relying party to verify that a BGPsec signature was produced by a BGP speaker belonging to a given AS. The goal of BGPsec is to use such signatures to protect the AS path data in BGP update messages, so that each BGP speaker can assess the validity of this data in update messages that it receives.

[3. BGPsec Operation](#)

The core of BGPsec is a new optional (non-transitive) attribute, called BGPsec_Path. This attribute includes both AS Path data as well as a sequence of digital signatures, one for each AS in the path. (The use of this new attribute is formally specified in [[I-D.sidr-bgpsec-protocol](#)].) A new signature is added to this sequence each time an update message leaves an AS. The signature is constructed so that any tampering with the AS path data or Network Layer Reachability Information (NLRI) in the BGPsec update message can be detected by the recipient of the message.

[3.1. Negotiation of BGPsec](#)

The use of BGPsec is negotiated using BGP capability advertisements [[RFC5492](#)]. Upon opening a BGP session with a peer, BGP speakers who support (and wish to use) BGPsec include a newly-defined capability in the OPEN message [[I-D.sidr-bgpsec-protocol](#)].

The use of BGPsec is negotiated separately for each address family. This means that a BGP speaker could, for example, elect to use BGPsec for IPv6, but not for IPv4 (or vice versa) routes. Additionally, the use of BGPsec is negotiated separately in the send and receive directions. This means that a BGP speaker could, for example, indicate support for sending BGPsec update messages but require that messages it receives be traditional (non-BGPsec) update message. (To see why such a feature is useful, see [Section 4.2](#).)

If the use of BGPsec is negotiated in a BGP session (in a given

direction, for a given address family) then both BGPsec update messages (ones that contain the BGPsec_Path_Signature attribute) and traditional BGP update messages (that do not contain this attribute) can be sent within the session.

If a BGPsec-capable BGP speaker finds that its peer does not support receiving BGPsec update messages, then the BGP speaker must remove the BGPsec_Path attribute from any update messages it sends to this peer.

3.2. Update signing and validation

When a BGP speaker originates a BGPsec update message, it creates a BGPsec_Path attribute containing a single signature. The signature protects the Network Layer Reachability Information (NLRI), the AS number of the originating AS, and the AS number of the peer AS to which the update message is being sent. Note that the NLRI in a BGPsec update message is restricted to contain only a single prefix.

When a BGP speaker receives a BGPsec update message and wishes to propagate the route advertisement contained in the update to an external peer, it adds a new signature to the BGPsec_Path attribute. This signature protects everything protected by the previous signature, plus the AS number of the new peer to which the update message is being sent.

Each BGP speaker also includes a reference, called a Subject Key Identifier (SKI). The SKI identifies the BGPsec Router Certificate of the BGP speaker signing the BGPsec_Path attribute. The SKI is used by a recipient to select the public key (and associated router certificate data) needed to validate the signature.

As an example, consider the following case in which an advertisement for 192.0.2/24 is originated by AS 1, which sends the route to AS 2, which sends it to AS 3, which sends it to AS 4. When AS 4 receives a BGPsec update message for this route, it will contain the following data:

- * NLRI: 192.0.2/24
- * AS path data: 3 2 1
- * BGPsec_Path contains 3 signatures :
 - o Signature from AS 1 protecting 192.0.2/24, AS 1 and AS 2
 - o Signature from AS 2 protecting Everything AS 1's signature protected, and AS 3
 - o Signature from AS 3 protecting Everything AS 2's signature protected, and AS 4

When a BGPsec update message is received by a BGPsec speaker, the BGPsec speaker can validate the message as follows. For each signature, the BGP speaker first determines if there is a valid RPKI Router certificate matching the SKI and containing the appropriate AS number. (This would typically be done by looking up the SKI in a cache of data extracted from valid RPKI objects. A cache allows certificate validation to be handled via an asynchronous process, which might execute on another device.)

The BGPsec speaker then verifies the signature using the public key from this BGPsec router certificate. If each of the signatures can be verified in this fashion, the BGPsec speaker is assured that the update message it received was propagated via the AS path specified in the update message.

In the above example, upon receiving the BGPsec update message, a BGP speaker for AS 4 would do the following. First, it would look at the SKI for the first signature and see if this corresponds to a valid BGPsec Router certificate for AS 1. Next, it would verify the first signature using the key found in this valid certificate. Finally, it would repeat this process for the second and third signatures, checking to see that there are valid BGPsec router certificates for AS 2 and AS 3 (respectively) and that the signatures can be verified with the keys found in these certificates. Note that the BGPsec speaker for AS 4 should additionally perform origin validation as per [RFC 6483](#) [[RFC6483](#)]. However, such origin validation is independent of BGPsec.

The deployment model for BGPsec requires that all ASs in a BGPsec protected path must be BGPsec speakers. It does not permit BGPsec protection of an update that propagates through ASs that do not support BGPsec. In particular, it does not permit what is called "partial path signing", in which a BGPsec AS attaches a BGPsec_Path attribute to an unprotected update that was received from a downstream neighbor.

Partial path signing might be viewed as supplying information about a portion of a path that could be used in making better routing decisions, preferring a partially protected route. However, partial path signing implies that the entire AS path is not rigorously protected. Rigorous AS path protection is a key requirement of BGPsec [[RFC7353](#)]. Partial path signing also introduces the following attack vulnerability: If a BGPsec speaker can attach a BGPsec_Path attribute to an unprotected update, and if BGPsec protected updates would be preferred to unprotected updates, then a BGPsec speaker can manufacture any unprotected update it wants and attach a BGPsec_Path attribute to it, and thereby increase the chance that its manufactured update will be preferred. Partial path signing then

becomes a privilege elevation attack vector, that could be employed by any BGPsec AS at any point.

The need to avoid introducing that vulnerability forced the stringent deployment model.

4. Design and Deployment Considerations

In this section we provide a brief overview of several additional topics that commonly arise in the discussion of BGPsec.

4.1. Disclosure of topology information

A key requirement in the design of BGPsec was that it not disclose any new information about BGP peering topology. Since many ISPs feel peering topology data is proprietary, further disclosure of it would inhibit BGPsec adoption.

In particular, the topology information that can be inferred from BGPsec update messages is exactly the same as that which can be inferred from equivalent (non-BGPsec) BGP update messages.

4.2. BGPsec router assumptions

In order to achieve its security goals, BGPsec assumes additional capabilities in routers. In particular, BGPsec requires adding digital signatures to BGP update messages, which will significantly increase the size of these messages. Therefore, an AS that wishes to receive BGPsec update messages will require additional memory in its routers to store (e.g., in ADJ RIBs) the data conveyed in these larger update messages. Additionally, the design of BGPsec assumes that an AS that elects to receive BGPsec update messages will do some cryptographic signature verification at its edge router. This verification may require additional capability in these edge routers.

Additionally, BGPsec requires that all BGPsec speakers support 4-byte AS Numbers [[RFC6793](#)]. This is because the co-existence strategy for 4-byte AS numbers and legacy 2-byte AS speakers that gives special meaning to AS 23456 is incompatible with the security properties that BGPsec seeks to provide.

For this initial version of BGPsec, optimizations to minimize the size of BGPsec updates or the processing required in edge routers have not been considered. Such optimizations may be considered in the future.

Note also that the design of BGPsec allows an AS to send BGPsec update messages (thus obtaining protection for routes it originates)

without receiving BGPsec update messages. An AS that sends, but does not receive, BGPsec update messages, will require much less capability in its edge routers to deploy BGPsec. In particular, a router that only sends BGPsec update messages does not need additional memory to store larger updates and requires only minimal cryptographic capability (as generating one signature per outgoing update requires less computation than verifying multiple signatures on each incoming update message). See [[I-D.sidr-bgpsec-ops](#)] for further discussion related to Edge ASes that do not provide transit.

4.3. BGPsec and consistency of externally visible data

Finally note that, by design, BGPsec prevents parties that propagate route advertisements from including inconsistent or erroneous information within the AS-Path (without detection). In particular, this means that any scenarios in which a BGP speaker constructs such an inconsistent or erroneous AS Path attribute will break when BGPsec is used.

For example, when BGPsec is not used, it is possible for a single autonomous system to have one peering session where it identifies itself as AS 111 and a second peering session where it identifies itself as AS 222. In such a case, it might receive route advertisements from the first peering session (as AS 111) and then add AS 222 (but not AS 111) to the AS-Path and propagate them within the second peering session.

Such behavior may very well be innocent and performed with the consent of the legitimate holder of both AS 111 and 222. However, it is indistinguishable from the following man-in-the-middle attack performed by a malicious AS 222. First, the malicious AS 222 impersonates AS 111 in the first peering session (essentially stealing a route advertisement intended for AS 111). The malicious AS 222 then inserts itself into the AS path and propagates the update to its peers.

Therefore, when BGPsec is used, such an autonomous system would either need to assert a consistent AS number in all external peering sessions, or else it would need to add both AS 111 and AS 222 to the AS-Path (along with appropriate signatures) for route advertisements that it receives from the first peering session and propagates within the second peering session. See [[I-D.sidr-as-migration](#)] for a detailed discussion of how to reasonably manage AS number migrations while using BGPsec.

5. Security Considerations

This document provides an overview of BGPsec; it does not define the

BGPsec extension to BGP. The BGPsec extension is defined in [I-D.sidr-bgpsec-protocol]. The threat model for the BGPsec is described in [[RFC7132](#)].

6. IANA Considerations

None.

7. References

7.1. Normative References

- [RFC4271] Rekhter, Y., Li, T., and S. Hares, Eds., "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), January 2006.
- [RFC6793] Vohra, Q. and E. Chen, "BGP Support for Four-octet AS Numbers", [RFC 6793](#), December 2012.
- [RFC5492] Scudder, J. and R. Chandra, "Capabilities Advertisement with BGP-4", [RFC 5492](#), February 2009.
- [RFC6480] Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", February 2012.
- [RFC6483] Huston, G., and G. Michaelson, "Validation of Route Origination using the Resource Certificate PKI and ROAs", February 2012.
- [RFC7132] Kent, S., and A. Chi, "Threat Model for BGP Path Security", [RFC 7132](#), February 2014.
- [RFC7115] Bush, R., "RPKI-Based Origin Validation Operation", [RFC 7115](#), January 2014.
- [I-D.sidr-bgpsec-protocol] Lepinski, M., Ed., "BPSEC Protocol Specification", [draft-ietf-sidr-bgpsec-protocol](#), work-in-progress.
- [I-D.sidr-bgpsec-ops] Bush, R., "BGPsec Operational Considerations", [draft-ietf-sidr-bgpsec-ops](#), work-in-progress.
- [I-D.sidr-bgpsec-algs] Turner, S., "BGPsec Algorithms, Key Formats, & Signature Formats", [draft-ietf-sidr-bgpsec-algs](#), work-in-progress.
- [I-D.sidr-bgpsec-pki-profiles] Reynolds, M., Turner, S., and S. Kent, "A Profile for BGPsec Router Certificates, Certificate Revocation Lists, and Certification Requests", [draft-ietf-sidr-bgpsec-pki-profiles](#), work-in-progress.

[I-D.sidr-as-migration] George, W. and S. Murphy, "BGPsec Considerations for AS Migration", [draft-ietf-sidr-as-migration](#), work-in-progress.

7.2. Informative References

[RFC4272] Murphy, S., "BGP Security Vulnerabilities Analysis", [RFC 4272](#), January 2006

[I-D.sriram-bgpsec-design-choices] Sriram, K., "BGPsec Design Choices and Summary of Supporting Discussions", [draft-sriram-bgpsec-design-choices](#), work-in-progress.

[RFC7353] Bellovin, S., R. Bush, and D. Ward, "Security Requirements for BGP Path Validation", [RFC 7353](#), August 2014.

Authors' Addresses

Matt Lepinski
New College of Florida
5800 Bay Shore Road
Sarasota, FL 34243
USA

Email: mlepinski@ncf.edu

Sean Turner
sn3rd

Email: sean@sn3rd.com

