

Network Working Group  
Internet Draft  
Intended status: Informational  
Expires: December 14, 2011

M. Lepinski  
BBN Technologies  
S. Turner  
IECA  
June 15, 2011

An Overview of BGPSEC  
draft-ietf-sidr-bgpsec-overview-00.txt

## 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), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

This Internet-Draft will expire on December 15, 2011.

## Copyright Notice

Copyright (c) 2011 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

<a href="#">1.</a>	Introduction.....	<a href="#">2</a>
<a href="#">2.</a>	Background.....	<a href="#">3</a>
<a href="#">3.</a>	BGPSEC Operation.....	<a href="#">4</a>
<a href="#">3.1.</a>	Negotiation of BGPSEC.....	<a href="#">4</a>
<a href="#">3.2.</a>	Update signing and validation.....	<a href="#">5</a>
<a href="#">4.</a>	Design and Deployment Considerations.....	<a href="#">7</a>
<a href="#">4.1.</a>	Disclosure of topology information.....	<a href="#">7</a>
<a href="#">4.2.</a>	BGPSEC router assumptions.....	<a href="#">7</a>
<a href="#">4.3.</a>	BGPSEC and consistency of externally visible data.....	<a href="#">8</a>
<a href="#">5.</a>	Security Considerations.....	<a href="#">8</a>
<a href="#">6.</a>	IANA Considerations.....	<a href="#">8</a>
<a href="#">7.</a>	References.....	<a href="#">9</a>
<a href="#">7.1.</a>	Normative References.....	<a href="#">9</a>
<a href="#">7.2.</a>	Informative References.....	<a href="#">9</a>

## [1.](#) Introduction

BGPSEC (Border Gateway Protocol Security) is an extension to the Border Gateway Protocol (BGP) that provides improved security for BGP routing [[RFC 4271](#)].

A comprehensive discussion of BGPSEC is provided in the following set of documents:

- [[I-D.kent-bgpsec-threats](#)]:

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

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

A standards track document specifying the BGPSEC extension to

BGP.

- . [[I-D.ymbk-bgpsec-ops](#)]:

An informational document describing operational considerations for BGPSEC deployment.

- . Certificate Profile Document (TBD)

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

- . Algorithms Document (TBD)

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

- . Design Choices Document (TBD)

An informational document describing the choices that were made in designing BGPSEC and the reasoning behind these choices.

The remainder of this document contains a brief overview of BGPSEC and envisioned usage.

## [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, [[RFC 4272](#)]).

The Resource Public Key Infrastructure (RPKI), described in [I-D.sidr-arch], 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 valid ROAs can be used by BGP speakers to determine whether a received route was originated by an AS authorized to originate that route (see [[I-D.sidr-roa-validation](#)] and [[I-D.sidr-origin-ops](#)]).

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 certain mis-

origination attacks. For example, if a BGP speaker accidentally (due to misconfiguration) originates routes to the wrong prefixes, ASes utilizing RPKI data could detect this error and decline to select these mis-originated routes. 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 [[I-D.kent-security-threats](#)] 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 of which is held by one or more BGP speakers within this AS. Private keys corresponding to public keys in such certificates can then be used within BGPSEC to enable BGP speakers to sign on behalf of their 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 signatures to protect the AS Path attribute of BGP update messages so that a BGP speaker can assess the validity of the AS Path in update messages that it receives.

### [3.](#) BGPSEC Operation

The core of BGPSEC is a new optional (non-transitive) attribute, called BGPSEC\_Path\_Signatures. This attribute consists of a sequence of digital signatures, one for each AS in the AS Path of a BGPSEC update message. (The use of this new attribute is formally specified in [[I-D.lepinski-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 or Network Layer Reachability Information (NLRI) in the BGPSEC update message will result in the recipient being able to detect that the update is

invalid.

### [3.1.](#) Negotiation of BGPSEC

The use of BGPSEC is negotiated using BGP capability advertisements [[RFC 5492](#)]. 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.

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). 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 might be 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 existing BGPSEC\_Path\_Signatures 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\_Signatures attribute containing a single signature. The signature protects the Network Layer Reachability Information (NLRI), the AS number of the originating AS, the AS number of the peer AS to whom the update message is being sent, and a few other pieces of data necessary for security guarantees. 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\_Signatures

attribute. This signature protects everything protected by the previous signature, plus the AS number of the new peer to whom the update message is being sent.

Each BGP speaker also adds a reference, called a Subject Key Identifier (SKI), to its BGPSEC Router certificate. The SKI is used by a recipient to select the public key (and selected router certificate data) needed for validation.

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 : 3 2 1

- . BGPSEC\_Path\_Signatures Attribute with 3 signatures :
  - o Signature from AS 1 protecting  
192.0.2/24, AS 1 and AS 2
  - o Signature from AS 2 protecting  
Everything AS1'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 BGP speaker, the BGP speaker can validate the message as follows. For each signature, the BGP speaker first needs to determine 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 BGP speaker then verifies the signature using the public key from this BGPSEC router certificate. If all the signatures can be verified in this fashion, the BGP speaker is assured that the update message it received actually came via the path specified in the AS\_Path attribute. Finally, the BGP speaker can check whether there exists a valid ROA in the RPKI linking the origin AS to the prefix in the NLRI. If such a valid ROA exists the BGP speaker is further assured that the AS at the beginning of the validated path was authorized to originate routes to the given prefix.

In the above example, upon receiving the BGPSEC update message, a BGP speaker for AS 4 would first check to make sure that there is a valid ROA authorizing AS 1 to originate advertisements for 192.0.2/24. It would then 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 then 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.

## [4.](#) Design and Deployment Considerations

In this section we briefly discuss 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 BGPSEC 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 involves 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 large 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 will likely require additional capability in these edge routers.

For this initial version of BGPSEC, optimizations to minimize the size of BGPSEC updates or the processing required in edge routers were NOT 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 only sends, and 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 large 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.ymbk-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 deployed 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.

## [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.lepinski-bgpsec-protocol]. The threat model for the BGPSEC is described in [[I-D.kent-bgpsec-threats](#)].

## [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.

[RFC5492] Scudder, J. and R. Chandra, "Capabilities Advertisement

with BGP-4", [RFC 5492](#), February 2009.

[I-D.sidr-arch] Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", [draft-ietf-sidr-arch](#), work-in-progress.

[I-D.sidr-roa-validation] Huston, G., and G. Michaelson, "Validation of Route Origination using the Resource Certificate PKI and ROAs", [draft-ietf-sidr-roa-validation](#), work-in-progress.

[I-D.sidr-origin-ops] Bush, R., "RPKI-Based Origin Validation Operation", [draft-ietf-sidr-origin-ops](#), work-in-progress.

[I-D.kent-bgpsec-threats] Kent, S., "Threat Model for BGP Path Security", [draft-kent-bgpsec-threats](#), work-in-progress.

[I-D.sidr-bgpsec-protocol] Lepinski, M., Ed., "BPSEC Protocol Specification", [draft-ietf-sidr-bgpsec-protocol](#), work-in-progress.

[I-D.ymbk-bgpsec-ops] Bush, R., "BGPSEC Operational Considerations", [draft-ymbk-bgpsec-ops](#), work-in-progress.

## [7.2](#). Informative References

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

Matt Lepinski  
BBN Technologies  
10 Moulton Street  
Cambridge MA 02138

Email: [mlepinski@bbn.com](mailto:mlepinski@bbn.com)

Sean Turner  
IECA, Inc.  
3057 Nutley Street, Suite 106  
Fairfax, VA 22031

Email: [turners@ieca.com](mailto:turners@ieca.com)