Network Working Group M. Lepinski, Ed.
Internet-Draft BBN
Intended status: Standards Track March 12, 2012
Expires: September 13, 2012

BGPSEC Protocol Specification
draft-ietf-sidr-bgpsec-protocol-02

Abstract

This document describes BGPSEC, an extension to the Border Gateway Protocol (BGP) that provides security for the AS-PATH attribute in BGP update messages. BGPSEC is implemented via a new optional non-transitive BGP path attribute that carries a digital signature produced by each autonomous system on the AS-PATH.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [4].

Table of Contents

1. Introduction

This document describes BGPSEC, a mechanism for providing path security for Border Gateway Protocol (BGP) [1] route advertisements. That is, a BGP speaker who receives a valid BGPSEC update has cryptographic assurance that the advertised route has the following two properties:

1. The route was originated by an AS that has been explicitly authorized by the holder of the IP address prefix to originate route advertisements for that prefix.

2. Every AS listed in the AS_Path attribute of the update explicitly authorized the advertisement of the route to the subsequent AS in the AS_Path.

This document specifies a new optional (non-transitive) BGP path attribute, BGPSEC_Path_Signatures. It also describes how a BGPSEC-compliant BGP speaker (referred to hereafter as a BGPSEC speaker) can generate, propagate, and validate BGP update messages containing this attribute to obtain the above assurances.

BGPSEC relies on the Resource Public Key Infrastructure (RPKI) certificates that attest to the allocation of AS number and IP address resources. (For more information on the RPKI, see [7] and the documents referenced therein.) Any BGPSEC speaker who wishes to send BGP update messages to external peers (eBGP) containing the BGPSEC_Path_Signatures must have an RPKI end-entity certificate (as well as the associated private signing key) corresponding to the BGPSEC speaker's AS number. Note, however, that a BGPSEC speaker does not require such a certificate in order to validate update messages containing the BGPSEC_Path_Signatures attribute.

2. BGPSEC Negotiation

This document defines a new BGP capability [3]that allows a BGP speaker to advertise to its neighbors the ability to send and/or receive BGPSEC update messages (i.e., update messages containing the BGPSEC_Path_Signatures attribute).

This capability has capability code : TBD

The capability length for this capability MUST be set to 5.

The three octets of the capability value are specified as follows.

Capability Value:

```
  0        1        2  3       4 5 6 7
+----------------------------------------+
| Send | Receive | Reserved  |  Version |
+----------------------------------------+
|               AFI                      |
+----------------------------------------+
|                                        |
+----------------------------------------+
|             Reserved                   |
+----------------------------------------+
|               SAFI                     |
+----------------------------------------+
```

The high order bit (bit 0) of the first octet is set to 1 to indicate that the sender is able to send BGPSEC update messages, and is set to zero otherwise. The next highest order bit (bit 1) of this octet is set to 1 to indicate that the sender is able to receive BGPSEC update messages, and is set to zero otherwise. The next two bits of the capability value (bits 2 and 3) are reserved for future use. These reserved bits should be set to zero by the sender and ignored by the receiver.

The four low order bits (4, 5, 6 and 7) of the first octet indicate the version of BGPSEC for which the BGP speaker is advertising support. This document defines only BGPSEC version 0 (all four bits set to zero). Other versions of BGPSEC may be defined in future documents. A BGPSEC speaker MAY advertise support for multiple versions of BGPSEC by including multiple versions of the BGPSEC capability in its BGP OPEN message.

If there does not exist at least one version of BGPSEC that is supported by both peers in a BGP session, then the use of BGPSEC has not been negotiated. (That is, in such a case, messages containing the BGPSEC_Path_Signatures MUST NOT be sent.)

If version 0 is the only version of BGPSEC for which both peers (in a BGP session) advertise support, then the use of BGPSEC has been negotiated and the BGPSEC peers MUST adhere to the specification of BGPSEC provided in this document. (If there are multiple versions of BGPSEC which are supported by both peers, then the behavior of those peers is outside the scope of this document.)

The second and third octets contain the 16-bit Address Family Identifier (AFI) which indicates the address family for which the BGPSEC speaker is advertising support for BGPSEC. This document only

specifies BGPSEC for use with two address families, IPv4 and IPv6, AFI values 1 and 2 respectively. BGPSEC for use with other address families may be specified in future documents.

The fourth octet in the capability is reserved. It is anticipated that this octet will not be used until such a time as the reserved octet in the Multi-protocol extensions capability advertisement [2] is specified for use. The reserved octet should be set to zero by the sender and ignored by the receiver.

The fifth octet in the capability contains the 8-bit Subsequent Address Family Identifier (SAFI). This value is encoded as in the BGP multiprotocol extensions [2].

Note that if the BGPSEC speaker wishes to use BGPSEC with two different address families (i.e., IPv4 and IPv6) over the same BGP session, then the speaker must include two instances of this capability (one for each address family) in the BGP OPEN message. A BGPSEC speaker SHOULD NOT advertise the capability of BGPSEC support for any <AFI, SAFI> combination unless it has also includes the multiprotocol extension capability for the same <AFI, SAFI> combination [2].

By indicating support for receiving BGPSEC update messages, a BGP speaker is, in particular, indicating that the following are true:

- o The BGP speaker understands the BGPSEC_Path_Signatures attribute (see Section 3).

- o The BGP speaker supports 4-byte AS numbers (see RFC 4893).

Note that BGPSEC update messages can be quite large, therefore any BGPSEC speaker announcing the capability to receive BGPSEC messages SHOULD also announce support for the capability to receive BGP extended messages [5].

A BGP speaker MUST NOT send an update message containing the BGPSEC_Path_Signatures attribute within a given BGP session unless both of the following are true:

- o The BGP speaker indicated support for sending BGPSEC update messages in its open message.

- o The peer of the BGP speaker indicated support for receiving BGPSEC update messages in its open message.

3. The BGPSEC_Path_Signatures Attribute

The BGPSEC_Path_Signatures attribute is a new optional (non-transitive) BGP path attribute.

This document registers a new attribute type code for this attribute : TBD

The BGPSEC_Path_Signatures attribute has the following structure:

```
                  BGPSEC_Path_Signatures Attribute
     +---------------------------------------------------------+
     | Flags Octet                     (1 octet)               |
     +---------------------------------------------------------+
     | Algorithm Suite Identifier 1    (1 octet)               |
     +---------------------------------------------------------+
     | Algorithm Suite Identifier 2    (1 octet)               |
     +---------------------------------------------------------+
     | Reserved                       (8 octets)               |
     +---------------------------------------------------------+
     | Sequence of Signature-Segments (variable)               |
     +---------------------------------------------------------+
```

The flags octet is an unsigned octet that contains flags to aid in receiver processing.

```
              Flags Octet in Path_Signatures Attribute

                  0               1  2  3  4  5  6  7
            +---------------------------------------------+
            | Two Algorithms  |        Reserved           |
            +---------------------------------------------+
```

The first bit in the Flags octet is set to zero in the common case that each Signature-Segment contains a single signature. The first bit of the Flags octet is set to one in the case that each Signature-Segment contains two signatures, produced by two different algorithm suites. (Note that this second case is necessary to support a transition between two algorithm suites, see Section 8.) The remaining 7 bits of the Flags octet are reserved for future use. These bits should be set to zero by the sender and ignored by the receiver.

Algorithm Suite Identifier 1 contains a one-octet identifier specifying the digest algorithm and digital signature algorithm used to produce the first signature in each Signature-Segment. An IANA

registry of algorithm identifiers for use in BGPSEC is created in the BGPSEC algorithms document[10].

Algorithm Suite Identifier 2 contains a one-octet identifier specifying the digest algorithm and digital signature algorithm used to produce the second signature in each Signature-Segment. This field is ignored by the receiver if the first bit in the Flags octet is set to zero (indicating that only one signature algorithm is used in this BGPSEC update). An IANA registry of algorithm identifiers for use in BGPSEC is created in the BGPSEC algorithms document[10].

There are eight octets reserved for future use. These octets are digitally signed (see Section 4 below).

EDITOR'S NOTE: In a previous version of this document there was an Expire Time that was used to provide protection against replay of old (stale) digital signatures or failure to propagate a withdrawal message. This mechanism was removed from the current version of the document. Please see the SIDR mailing list for discussions related to protection against replay attacks. Depending on the result of discussions within the SIDR working group this reserved field could at some future point be used to re-introduce Expire Time, or some other octets used in a future replay protection mechanism.

The BGPSEC_Path_Signatures attribute contains one Signature-Segment for each AS along the path of the route advertisement in this update message. (For a detailed explanation of how an AS processes a BGPSEC update message and adds a new Signature_Segment, see Section 4.) A Signature-Segment has the following structure:

Signature Segments

```
+-------------------------------------------- +
| AS Number                        (4 octets)  |
+---------------------------------------------+
| pCount                            (1 octet)  |
+---------------------------------------------+
| Subject Key Identifier 1 Length  (1 octet)  |
+---------------------------------------------+
| Subject Key Identifier 1        (variable)  |
+---------------------------------------------+
| Signature 1 Length                (1 octet)  |
+---------------------------------------------+
| Signature 1                     (variable)  |
+---------------------------------------------+
| Subject Key Identifier 2 Length  (1 octet)  |
+---------------------------------------------+
| Subject Key Identifier 2        (variable)  |
+---------------------------------------------+
| Signature Length 2                (1 octet)  |
+---------------------------------------------+
| Signature 2                     (variable)  |
+---------------------------------------------+
```

The AS Number is the Autonomous System Number of the BGPSEC speaker that produced the digital signature(s) in this Signature Segment.

The pCount field contains an unsigned integer indicating the number of repetitions of the associated autonomous system number that the signature covers. This field enables a BGPSEC speaker to mimic the semantics of adding multiple copies of their AS to the AS-PATH without requiring the speaker to generate multiple signatures.

The Subject Key Identifier 1 Length field contains the size (in octets) of the value in the Subject Key Identifier 1 field of the Signature-Segment. The Subject Key Identifier 1 field contains the value in the Subject Key Identifier extension of the RPKI end-entity certificate that is used to verify the first signature in the Signature-Segment (see Section 5 for details on validity of BGPSEC update messages).

The Signature 1 Length field contains the size (in octets) of the value in the Signature 1 field. The Signature 1 field contains a digital signature that protects the NLRI and the BGPSEC_Path_Signatures attribute (see Sections 4 and 5 for details on generating and verifying this signature, respectively).

The Subject Key Identifier 2 Length field contains the size (in octets) of the value in the Subject Key Identifier 2 field of the Signature-Segment. This length field SHOULD be zero if the first bit in the Flags octet is zero (indicating that only one algorithm suite is being used to generate signatures for this update message). The Subject Key Identifier 2 field contains the value in the Subject Key Identifier extension of the RPKI end-entity certificate that is used to verify the second signature in the Signature-Segment (see Section 5 for details on validity of BGPSEC update messages). This field is ignored by the receiver when the first bit in the Flags octet is zero (indicating that only one algorithm suite is being used to generate signatures for this update message).

The Signature 2 Length field contains the size (in octets) of the value in the Signature 2 field. This length field SHOULD be zero if the first bit in the Flags octet is zero (indicating that only one algorithm suite is being used to generate signatures for this update message). The Signature 2 field contains a digital signature that protects the NLRI and the BGPSEC_Path_Signatures attribute (see Sections 4 and 5 for details on generating and verifying this signature, respectively). This field is ignored by the receiver when the first bit in the Flags octet is zero (indicating that only one algorithm suite is being used to generate signatures for this update message).

4. Generating a BGPSEC Update

Sections 4.1 and 4.2 cover two cases in which a BGPSEC speaker may generate an update message containing the BGPSEC_Path_Signatures attribute. The first case is that in which the BGPSEC speaker originates a new route advertisement (Section 4.1). That is, the BGPSEC speaker is constructing an update message in which the only AS to appear in the AS_PATH attribute is the speaker's own AS (normally appears once but may appear multiple times if AS prepending is applied). The second case is that in which the BGPSEC speaker receives a route advertisement from a peer and then decides to propagate the route advertisement to an external (eBGP) peer (Section 4.2). That is, the BGPSEC speaker has received a BGPSEC update message and is constructing a new update message for the same NLRI in which the AS_PATH attribute will contain AS number(s) other than the speaker's own AS.

In the remaining case where the BGPSEC speaker is sending the update message to an internal (iBGP) peer, the BGPSEC speaker populates the BGPSEC_Path_Signatures attribute by copying the BGPSEC_Path_Signatures attribute from the received update message. That is, the BGPSEC_Path_Signatures attribute is copied verbatim.

Note that in the case that a BGPSEC speaker chooses to forward to an iBGP peer a BGPSEC update message that has not been successfully validated (see Section 5), the BGPSEC_Path_Signatures attribute SHOULD NOT be removed. (See Section 7 for the security ramifications of removing BGPSEC signatures.)

The information protected by the signature on a BGPSEC update message includes the AS number of the peer to whom the update message is being sent. Therefore, if a BGPSEC speaker wishes to send a BGPSEC update to multiple BGP peers, it MUST generate a separate BGPSEC update message for each unique peer AS to which the update message is sent.

A BGPSEC update message MUST advertise a route to only a single NLRI. This is because a BGPSEC speaker receiving an update message with multiple NLRI is unable to construct a valid BGPSEC update message (i.e., valid path signatures) containing a subset of the NLRI in the received update. If a BGPSEC speaker wishes to advertise routes to multiple NLRI, then it MUST generate a separate BGPSEC update message for each NLRI.

Note that in order to create or add a new signature to a BGPSEC update message with a given algorithm suite, the BGPSEC speaker must possess a private key suitable for generating signatures for this algorithm suite. Additionally, this private key must correspond to the public key in a valid Resource PKI end-entity certificate whose AS number resource extension includes the BGPSEC speaker's AS number [11]. Note also new signatures are only added to a BGPSEC update message when a BGPSEC speaker is generating an update message to send to an external peer (i.e., when the AS number of the peer is not equal to the BGPSEC speaker's own AS number). Therefore, a BGPSEC speaker who only sends BGPSEC update messages to peers within its own AS, it does not need to possess any private signature keys.

4.1. Originating a New BGPSEC Update

In an update message that originates a new route advertisement (i.e., an update whose AS_Path contains a single AS number), a BGPSEC speaker will use only a single algorithm suite. That is, the BGPSEC speaker will set the Two_Algorithms flag to 0 in the BGPSEC_Path_Signatures attribute and include only a single signature in the Signature-Segment (setting the Signature 2 Length and Subject Key Identifier 2 Lengths to zero). However, to ensure backwards compatibility during a period of transition from a 'current' algorithm suite to a 'new' algorithm suite, it will be necessary to originate update messages containing both the 'current' and the 'new' algorithm suites (see Section 6.1). In such a case the BGPSEC speaker will set the Two_Algorithms flag to 1 in the

BGPSEC_Path_Signatures attribute and include two separate digital signatures (one for each algorithm suite). For the remainder of this section we describe the common case where the Two_Algorithms flag is set to one. However, the construction of the second signature is completely analogous (the only change is the replacement of 1 by 2 in the field names corresponding to the second signature).

The Resource PKI enables the legitimate holder of IP address prefix(es) to issue a signed object, called a Route Origination Authorization (ROA), that authorizes a given AS to originate routes to a given set of prefixes (see [6]). Note that validation of a BGPSEC update message will fail (i.e., the validation algorithm, specified in Section 5.1, returns 'Not Good') unless there exists a valid ROA authorizing the first AS in the AS PATH attribute to originate routes to the prefix being advertised. Therefore, a BGPSEC speaker SHOULD NOT originate a BGPSEC update advertising a route for a given prefix unless a ROA has previously been created (and published in the repository system) that authorizing the BGPSEC speaker's AS to originate routes to this prefix.

EDITOR'S NOTE: In a previous version of this document there was a description here of a mechanism that used that used periodic repetition of update messages (aka "beaconing") to protect against replay of old (stale) digital signatures or failure to propagate a withdrawal message. This mechanism was removed from the current version of the document. Please see the SIDR mailing list for discussions related to protection against replay attacks. Depending on the result of discussions within the SIDR working group a mechanism for protection against replay of digital signatures may be re-introduced into BGPSEC in the future.

When originating a new route advertisement, the BGPSEC_Path_Signatures attribute MUST contain a single Signature-Segment. The following describes how the BGPSEC speaker populates the fields of the Signature-Segment (see Section 3 for more information on the syntax of the Signature-Segment).

The AS field is set to the AS number of the BGPSEC speaker. That is, the AS number that the BGPSEC speaker advertised in the Open message of the current BGP session.

The pCount field is typically set to the value 1. However, a BGPSEC speaker may set the pCount field to a value greater than 1. Setting the pCount field to a value greater than one has the same semantics as repeating an AS number multiple times in the AS_PATH of a non-BGPSEC update message (e.g., for traffic engineering purposes). Setting the pCount field to a value greater than one permits this repetition without requiring a separate digital signature for each

repetition.

The Subject Key Identifier 1 field (see Section 3) is populated with the identifier contained in the Subject Key Identifier extension of the RPKI end-entity certificate (containing keys suitable for use with Algorithm Suite 1) used by the BGPSEC speaker. This Subject Key Identifier will be used by recipients of the route advertisement to identify the proper certificate to use in verifying the signature.

The Subject Key Identifier 1 Length field is populated with the length (in octets) of the Subject Key Identifier 1 field.

The Signature 1 field contains a digital signature that binds the NLRI, AS_Path attribute and BGPSEC_Path_Signatures attribute to the RPKI end-entity certificate used by the BGPSEC speaker. The digital signature is computed as follows:

- Construct a sequence of octets by concatenating the Target AS Number, AS Number (from the Signature_Segment), pCount, Algorithm Suite Identifier 1, Reserved field of the BGPSEC_Path_Signatures attribute and NLRI. The Target AS Number is the AS to whom the BGPSEC speaker intends to send the update message. (Note that the Target AS number is the AS number announced by the peer in the OPEN message of the BGP session within which the update is sent.)

```
                Sequence of Octets to be Signed
           +-----------------------------------------+
           | Target AS Number (4 octets)             |
           +-----------------------------------------+
           | AS Number        (4 octets)             |
           +-----------------------------------------+
           | pCount          (1 octet)               |
           +-----------------------------------------+
           | Algorithm Suite Identifier 1 (1 octet)  |
           +-----------------------------------------+
           | Expire Time     (8 octets)              |
           +-----------------------------------------+
           | NLRI Length     (1 octet)               |
           +-----------------------------------------+
           | NLRI Prefix    (variable)               |
           +-----------------------------------------+
```

- Apply to this octet sequence the digest algorithm (for Algorithm Suite 1) to obtain a digest value.

- Apply to this digest value the signature algorithm, (for Algorithm Suite 1) to obtain the digital signature. Then populate the Signature 1 field with this digital signature.

The Signature 1 Length field is populated with the length (in octets) of the Signature 1 field.

4.2. Propagating a Route Advertisement

When a BGPSEC speaker receives a BGPSEC update message containing a BGPSEC_Path_Signatures algorithm (with one or more signatures) from a (internal or external) peer, it may choose to propagate the route advertisement by sending to its (internal or external) peers by creating a new BGPSEC advertisement for the same prefix.

A BGPSEC speaker MUST NOT generate an update message containing the BGPSEC_Path_Signatures attribute unless it has selected, as the best route to the given prefix, a route that it received in an update message containing the BGPSEC_Path_Signatures attribute. In particular, this means that whenever a BGPSEC speaker generates an update message with a BGPSEC_Path_Signatures attribute that it will possess a received update message for the same prefix that also contains a BGPSEC_Path_Signatures attribute.

Additionally, whenever a BGPSEC speaker selects as the best route to a given prefix a route that it received in an update message containing the BGPSEC_Path_Signatures attribute, it is RECOMMENDED that if the BGPSEC speaker chooses to propagate the route that it generate an update message containing the BGPSEC_Path_Signatures attribute. However, a BGPSEC speaker MAY propagate a route advertisement by generating a (non-BGPSEC) update message that does not contain the BGPSEC_Path_Signatures attribute. Note that if a BGPSEC speaker receives a route advertisement containing the BGPSEC_Path_Signatures attribute and chooses for any reason (e.g., its peer is a non-BGPSEC speaker) to propagate the route advertisement as a non-BGPSEC update message without the BGPSEC_Path_Signatures attribute, then it MUST follow the instructions in Section 4.2.1.

The Subject Key Identifier 1 field (see Section 3) is populated with the identifier contained in the Subject Key Identifier extension of the RPKI end-entity certificate (containing keys suitable for use with Algorithm Suite 1) used by the BGPSEC speaker. This Subject Key Identifier will be used by recipients of the route advertisement to identify the proper certificate to use in verifying the signature.

The Subject Key Identifier 1 Length field is populated with the length (in octets) of the Subject Key Identifier 1 field.

Note that removing BGPSEC signatures (i.e., propagating a route advertisement without the BGPSEC_Path_Signatures attribute) has significant security ramifications. (See Section 7 for discussion of

the security ramifications of removing BGPSEC signatures.) Therefore, when a route advertisement is received via a BGPSEC update message, propagating the route advertisement without the BGPSEC_Path_Signatures attribute is NOT RECOMMENDED. Furthermore, note that when a BGPSEC speaker propagates a route advertisement with the BGPSEC_Path_Signatures attribute it is attesting to the fact that: (1) it received a BGPSEC update message that advertised this route; and (2) it chose this route as its best path to the given prefix. That is, the BGPSEC speaker is not attesting to the validation state of the update message it received. (See Section 7 for more discussion of the security semantics of BGPSEC signatures.)

If the BGPSEC speaker is producing an update message which contains an AS-SET (e.g., the BGPSEC speaker is performing proxy aggregation), then the BGPSEC speaker MUST NOT include the BGPSEC_Path_Signatures attribute. In such a case, the BGPSEC speaker must remove any existing BGPSEC_Path_Signatures in the received advertisement(s) for this prefix and produce a standard (non-BGPSEC) update message.

If the received BGPSEC update message uses two algorithm suites (i.e., the Two_Algorithms flag is set to 1) and the BGPSEC speaker supports both of the corresponding algorithms suites, then the BGPSEC speaker SHOULD generate a new update message that uses both algorithm suites (i.e., set the Two_Algorithms flag to 1). If the received BGPSEC update message that uses two algorithm suites and the BGPSEC speaker does not support the second algorithm suite, then the BGPSEC speaker MUST set the Two_Algorithms flag to 1 and remove the Signature 2 and Subject Key Identifier 2 fields from each Signature-Segment in the BGPSEC_Path_Signatures attribute (and set the corresponding lengths to zero). Note that this case can happen during an algorithm transition when the BGPSEC speaker has not yet been updated to support the new algorithm, see Section 6 for more details. If the BGPSEC speaker does not support the first algorithm suite in a BGPSEC update message, then the BGPSEC speaker MUST NOT propagate the route advertisement with the BGPSEC_Path_Signatures attribute. (Note that if this case occurs, something has gone wrong, as algorithm transitions are designed to never produce this case.)

The Reserved field from the BGPSEC_Path_Signatures attribute is copied directly from the Reserved field in the received update message.

The BGPSEC speaker then creates a new Signature-Segment. This Signature-Segment is prepended to the list of Signature-Segments (placed in the first position) so that the list of Signature-Segments appears in the same order as the corresponding AS numbers in the AS_PATH attribute. The BGPSEC speaker populates the fields of this new Signature-Segment as follows.

The AS field is set to the AS number of the BGPSEC speaker. That is, the AS number that the BGPSEC speaker advertised in the Open message of the current BGP session.

The pCount is typically set to the value 1. A BGPSEC speaker may set the pCount field to a value greater than 1. (See Section 4.1 for a discussion of setting pCount to a value greater than 1.) A route server that participates in the BGP control path, but does not act as a transit AS in the data plane, may choose to set pCount to 0. This option enables the route server to participate in BGPSEC and obtain the associated security guarantees without increasing the effective length of the AS_PATH. (Note that the Signature_Segmenet still contains the AS Number of the route server as this information is necessary for signature verification.) Note that the option of setting pCount to 0 is intended only for use by route servers that desire not to increase the effective AS-PATH length of routes they advertise. The pCount field SHOULD NOT be set to 0 in other circumstances. BGPSEC speakers SHOULD drop incoming update messages with pCount set to zero in cases where the BGPSEC speaker does not expect its peer to set pCount to zero (i.e., cases where the peer is not acting as a route server).

The Subject Key Identifier 1 field (see Section 3) is populated with the identifier contained in the Subject Key Identifier extension of the RPKI end-entity certificate (containing keys suitable for use with Algorithm Suite 1) used by the BGPSEC speaker. This Subject Key Identifier will be used by recipients of the route advertisement to identify the proper certificate to use in verifying the signature.

The Subject Key Identifier 1 Length field is populated with the length (in octets) of the Subject Key Identifier 1 field.

The Signature 1 field in the new segment contains a digital signature that binds the NLRI, AS_Path attribute and BGPSEC_Path_Signatures attribute to the RPKI end-entity certificate used by the BGPSEC speaker. The digital signature is computed as follows:

- o Construct a sequence of octets by concatenating the Signature 1 Length and Signature 1 fields of the most recent Signature-Segment (the one corresponding to AS from whom the BGPSEC speaker's AS received the announcement) with the pCount field inserted by the signer, and the Target AS (the AS to whom the BGPSEC speaker intends to send the update message). Note that the Target AS number is the AS number announced by the peer in the OPEN message of the BGP session within which the BGPSEC update message is sent.

Sequence of Octets to be Signed

```
+-------------------------------------------------------------+
| Most Recent Signature 1 Length Field      (1 octet)         |
+-------------------------------------------------------------+
| Most Recent Signature 1 Field             (variable)        |
+-------------------------------------------------------------+
| pCount Field of Signer        (1 octet)                     |
+-------------------------------------------------------------+
| Target AS Number              (4 octets)                    |
+-------------------------------------------------------------+
```

- o Apply to this octet sequence the digest algorithm (for the algorithm suite of this Signature-List) to obtain a digest value.

- o Apply to this digest value the signature algorithm, (for the algorithm suite of this Signature-List) to obtain the digital signature. Then populate the Signature Field with this digital signature.

The Subject Key Identifier 1 Length field is populated with the length (in octets) of the Subject Key Identifier 1 field.

5. Processing a Received BGPSEC Update

Validation of a BGPSEC update messages makes use of data from RPKI certificates and signed Route Origination Authorizations (ROA). In particular, to validate update messages containing the BGPSEC_Path_Signatures attribute, it is necessary that the recipient have access to the following data obtained from valid RPKI certificates and ROAs:

- o For each valid RPKI end-entity certificate containing an AS Number extension, the AS Number, Public Key and Subject Key Identifier are required

- o For each valid ROA, the AS Number and the list of IP address prefixes

Note that the BGPSEC speaker could perform the validation of RPKI certificates and ROAs on its own and extract the required data, or it could receive the same data from a trusted cache that performs RPKI validation on behalf of (some set of) BGPSEC speakers. (The latter case in analogous to the use of the RPKI-RTR protocol [12] for origin validation.)

To validate a BGPSEC update message containing the

BGPSEC_Path_Signatures attribute, the recipient performs the validation steps specified in Section 5.1. The validation procedure results in one of two states: 'Good' and 'Not Good'.

It is expected that the output of the validation procedure will be used as an input to BGP route selection. However, BGP route selection and thus the handling of the two validation states is a matter of local policy, and shall be handled using existing local policy mechanisms. It is expected that BGP peers will generally prefer routes received via 'Good' BGPSEC update messages over routes received via 'Not Good' BGPSEC update messages as well as routes received via update messages that do not contain the BGPSEC_Path_Signatures attribute. However, BGPSEC specifies no changes to the BGP decision process and leaves to the operator the selection of an appropriate policy mechanism to achieve the operator's desired results within the BGP decision process.

BGPSEC validation need only be performed at eBGP edge. The validation status of a BGP signed/unsigned update MAY be conveyed via iBGP from an ingress edge router to an egress edge router. Local policy in the AS determines the specific means for conveying the validation status through various pre-existing mechanisms (e.g., modifying an attribute). As discussed in Section 4, when a BGPSEC speaker chooses to forward a (syntactically correct) BGPSEC update message, it SHOULD be forwarded with its BGPSEC_Path_Signatures attribute intact (regardless of the validation state of the update message). Based entirely on local policy settings, an egress router MAY trust the validation status conveyed by an ingress router or it MAY perform its own validation.

EDITOR'S NOTE: Text will be inserted here for dealing with the AS_PATH attribute. Note that the BGPGSEC_Path_Signatures attribute now contains all of the information needed to construct the AS_PATH attribute. Therefore, there seem to be two options. One option the BGPSEC speaker checks the AS_PATH attribute against the information in the BGPSEC_Path_Signatures attribute and returns "Not Good" if the two do not match. The other option is that the BGPSEC speaker discards anything in the AS_PATH attribute and reconstructs the AS_PATH from the data in the BGPSEC_Path_Signatures attribute. I believe that there are no interoperability problems if the choice between these two options is left up to the BGPSEC speaker.

5.1. Validation Algorithm

This section specifies an algorithm for validation of BGPSEC update messages. A conformant implementation MUST include an BGPSEC update validation algorithm that is functionally equivalent to the external behavior of this algorithm.

First, the recipient of a BGPSEC update message performs a check to ensure that the message is properly formed. Specifically, the recipient checks that the BGPSEC_Path_Signatures attribute is properly formed (as specified in Section 3). If the BGPSEC_Path_Signatures attribute is not properly formed, then the recipient should log that an error occurred and drop the update message containing the error.

Second, the BGPSEC speaker verifies that the origin AS is authorized to advertise the prefix in question. To do this, consult the valid ROA data to obtain a list of AS numbers that are associated with the given IP address prefix in the update message. Then locate the last (least recently added) AS number in the AS-Path. If the origin AS in the AS-Path is not in the set of AS numbers associated with the given prefix, then BGPSEC update message is 'Not Good' and the validation algorithm terminates.

Third, the BGPSEC speaker examines the Algorithm Suite identifiers and the Two-Algorithms flag in the BGPSEC_Path_Signatures attribute. If the BGPSEC speaker does not support the first Algorithm Suite, then the BGPSEC speaker MUST treat the update message in the same manner that the BGPSEC speaker would treat an update message that arrived without a BGPSEC_Path_Signatures attribute. (Note that algorithm transitions are designed so that this case will never happen, therefore if this case occurs the BGPSEC speaker SHOULD log an error message.) If the Two-Algorithms flag is set to 1 and the BGPSEC speaker supports only the first algorithm suite then it follows the instructions below to validate the signatures using the first algorithm suite, and ignore Signature 2 in each Signature-Segment. If the Two-Algorithms flag is set to 1 and the BGPSEC speaker supports both algorithm suites, then the BGPSEC speaker follows the instructions below to validate the signatures using the first algorithm suite. The BGPSEC speaker MAY then analogously validate the second set of signatures using Algorithm Suite 2. If the BGPSEC speaker chooses to validate both sets of signatures, it returns "Good" if either the first or the second set of signatures successfully validate.

- o (Step I): Locate the public key needed to verify the signature (in the current Signature-Segment). To do this, consult the valid RPKI end-entity certificate data and look for an SKI that matches the value in the Subject Key Identifier 1 field of the Signature-Segment. If no such SKI value is found in the valid RPKI data then validation fails and returns "Not Good". Similarly, if the SKI exists but the AS Number associated with the SKI does NOT match the AS Number in the Signature-Segment, then validation fails and returns "Not Good".

o (Step II): Compute the digest function (for Algorithm Suite 1) on the appropriate data. If the segment is not the (least recently added) segment corresponding to the origin AS, then the digest function should be computed on the following sequence of octets:

Sequence of Octets to be Hashed

```
+--------------------------------------------------------+
| Signature 1 Length Field in the Next Segment  (1 octet) |
+--------------------------------------------------------+
| Signature 1 Field in the Next Segment         (variable) |
+--------------------------------------------------------+
| pCount Field in the Current Segment           (1 octet)  |
+--------------------------------------------------------+
| AS Number of Previous AS                      (4 octets) |
+--------------------------------------------------------+
```

The 'Signature 1 Field in the Next Segment' and 'Signature 1 Length Field in Next Segment' are the Signature 1 field and Signature 1 Length fields found in the Signature-Segment that is next to be processed (that is, the next most recently added Signature- Segment). The 'pCount Field in the Current Segment' is the pCount field found in the Signature-Segment that is currently being processed.

For the first segment to be processed (the most recently added segment), the 'AS Number of Subsequent AS' is the AS number of the BGPSEC speaker validating the update message. Note that if a BGPSEC speaker uses multiple AS Numbers (e.g., the BGPSEC speaker is a member of a confederation), the AS number used here MUST be the AS number announced in the OPEN message for the BGP session over which the BGPSEC update was received.

For each other Signature-Segment, the 'AS Number of Previous AS' is the AS number in the Signature-Segment that was most recently processed.

Alternatively, if the segment being processed corresponds to the origin AS, then the digest function should be computed on the following sequence of octets:

```
                 Sequence of Octets to be Hashed
           -------------------------------------------+
           | AS Number of Previous AS    (4 octets)   |
           +------------------------------------------+
           | Origin AS Number            (4 octets)   |
           +------------------------------------------+
           | Algorithm Suite 1 Identifier  (1 octet)  |
           +------------------------------------------+
           | pCount       (1 octet)                   |
           +------------------------------------------+
           | NLRI Length  (1 octet)                   |
           +------------------------------------------+
           | NLRI Prefix  (variable)                  |
           +------------------------------------------+
```

The NLRI Length, NLRI Prefix, Expire Time, and Algorithm Suite Identifier are all obtained in a straight forward manner from the NLRI of the update message or the BGPSEC_Path_Signatures attribute being validated. The pCount field is taken from the Signature-Segment currently being processed.

The Origin AS Number is the same Origin AS Number that was located in Step I above. (That is, the AS number in the least recently added Signature-Segment.)

The 'AS Number of Previous AS' is the AS Number in the Signature-Segment that was most recently processed (i.e., processed before the current segment).

o (Step III): Use the signature validation algorithm (for the given algorithm suite) to verify the signature in the current segment. That is, invoke the signature validation algorithm on the following three inputs: the value of the Signature field in the current segment; the digest value computed in Step II above; and the public key obtained from the valid RPKI data in Step I above. If the signature validation algorithm determines that the signature is invalid, validation has failed and return 'Not Good'. If the signature validation algorithm determines that the signature is valid, then continue processing Signature-Segments.

If all Signature-Segments pass validation (i.e., all segments are processed and the algorithm has not yet returned 'Not Good'), then validation succeeds and returns 'Good'.

6. Algorithms and Extensibility

6.1. Algorithm Suite Considerations

Note that there is currently no support for bilateral negotiation between BGPSEC peers to use of a particular (digest and signature) algorithm suite using BGP capabilities. This is because the algorithm suite used by the sender of a BGPSEC update message must be understood not only by the peer to whom he is directly sending the message, but also by all BGPSEC speakers to whom the route advertisement is eventually propagated. Therefore, selection of an algorithm suite cannot be a local matter negotiated by BGP peers, but instead must be coordinated throughout the Internet.

To this end, a mandatory algorithm suites document will be created which specifies a mandatory-to-use 'current' algorithm suite for use by all BGPSEC speakers. Additionally, the document specifies an additional 'new' algorithm suite that is recommended to implement.

It is anticipated that in the future the mandatory algorithm suites document will be updated to specify a transition from the 'current' algorithm suite to the 'new' algorithm suite. During the period of transition (likely a small number of years), all BGPSEC update messages SHOULD simultaneously use both the 'current' algorithm suite and the 'new' algorithm suite. (Note that Sections 3 and 4 specify how the BGPSEC_Path_Signatures attribute can contain signatures, in parallel, for two algorithm suites.) Once the transition is complete, use of the old 'current' algorithm will be deprecated, use of the 'new' algorithm will be mandatory, and a subsequent 'even newer' algorithm suite may be specified as recommend to implement. Once the transition has successfully been completed in this manner, BGPSEC speakers SHOULD include only a signatures corresponding to the 'new' algorithm.

6.2. Extensibility Considerations

This section discusses potential changes to BGPSEC that would require substantial changes to the processing of the BGPSEC_Path_Signatures and thus necessitate a new version of BGPSEC. Examples of such changes include

- o A new type of signature algorithm for which the number of signatures in the Signature-List Block is not equal to the number of ASes in the AS_PATH (e.g., aggregate signatures)

- o Changes to the data that is protected by the BGPSEC signatures (e.g., protection of attributes other than AS_PATH)

In the case that such a change to BGPSEC were deemed desirable, it is expected that a subsequent version of BGPSEC would be created and

that this version of BGPSEC would specify a new BGP Path Attribute, let's call it BGPSEC_PATH_SIG_TWO, which is designed to accommodate the desired changes to BGPSEC. In such a case, the mandatory algorithm suites document would be updated to specify algorithm suites appropriate for the new version of BGPSEC.

At this point a transition would begin which is analogous to the algorithm transition discussed in Section 6.2. During the transition period all BGPSEC speakers SHOULD simultaneously include both the BGPSEC_PATH_SIGNATURES attribute and the new BGPSEC_PATH_SIG_TWO attribute. Once the transition is complete, the use of BGPSEC_PATH_SIGNATURES could then be deprecated, at which point BGPSEC speakers SHOULD include only the new BGPSEC_PATH_SIG_TWO attribute. Such a process could facilitate a transition to a new BGPSEC semantics in a backwards compatible fashion.

7. Security Considerations

For discussion of the BGPSEC threat model and related security considerations, please see [8].

A BGPSEC speaker who receives a valid BGPSEC update message, containing a route advertisement for a given prefix, is provided with the following security guarantees:

o The origin AS number corresponds to an autonomous system that has been authorized by the IP address space holder to originate route advertisements for the given prefix.

o For each subsequent AS number in the AS-Path, a BGPSEC speaker authorized by the holder of the AS number selected the given route as the best route to the given prefix.

o For each AS number in the AS Path, a BGPSEC speaker authorized by the holder of the AS number intentionally propagated the route advertisement to the next AS in the AS-Path.

That is, the recipient of a valid BGPSEC Update message is assured that the AS-Path corresponds to a sequence of autonomous systems who have all agreed in principle to forward packets to the given prefix along the indicated path. (It should be noted BGPSEC does not offer a precise guarantee that the data packets would propagate along the indicated path; it only guarantees that the BGP update conveying the path indeed propagated along the indicated path.) Furthermore, the recipient is assured that this path terminates in an autonomous system that has been authorized by the IP address space holder as a legitimate destination for traffic to the given prefix.

Note that although BGPSEC provides a mechanism for an AS to validate that a received update message has certain security properties, the use of such a mechanism to influence route selection is completely a matter of local policy. Therefore, a BGPSEC speaker can make no assumptions about the validity of a route received from an external BGPSEC peer. That is, a compliant BGPSEC peer may (depending on the local policy of the peer) send update messages that fail the validity test in Section 5. Thus, a BGPSEC speaker MUST completely validate all BGPSEC update messages received from external peers. (Validation of update messages received from internal peers is a matter of local policy, see Section 5).

Note that there may be cases where a BGPSEC speaker deems 'Good' (as per the validation algorithm in Section 5.1) a BGPSEC update message that contains two sets of signatures, one 'Good' and one 'Not Good'. That is, the update message contains two sets of signatures corresponding to two algorithm suites, and one set of signatures verifies correctly and the other set of signatures fails to verify. In this case, the protocol specifies that if the BGPSEC speaker propagates the route advertisement received in such an update message then the BGPSEC speaker SHOULD add its signature using both the algorithm suites. Thus the BGPSEC speaker creates a signature using both algorithm suites and creates a new update message that contains both the 'Good' and the 'Not Good' set of signatures (from its own vantage point).

To understand the reason for such a design decision consider the case where the BGPSEC speaker receives an update message with both a set of algorithm A signatures which are 'Good' and a set of algorithm B signatures which are 'Not Good'. In such a case it is possible (perhaps even quite likely) that some of the BGPSEC speaker's peers (or other entities further 'downstream' in the BGP topology) do not support algorithm A. Therefore, if the BGPSEC speaker were to remove the 'Not Good' set of signatures corresponding to algorithm B, such entities would treat the message as though it were unsigned. By including the 'Not Good' set of signatures when propagating a route advertisement, the BGPSEC speaker ensures that 'downstream' entities have as much information as possible to make an informed opinion about the validation status of a BGPSEC update.

Note also that during a period of partial BGPSEC deployment, a 'downstream' entity might reasonably treat unsigned messages different from BGPSEC updates that contain a single set of 'Not Good' signatures. That is, by removing the set of 'Not Good' signatures the BGPSEC speaker might actually cause a downstream entity to 'upgrade' the status of a route advertisement from 'Not Good' to unsigned. Finally, note that in the above scenario, the BGPSEC speaker might have deemed algorithm A signatures 'Good' only because

of some issue with RPKI state local to his AS (for example, his AS might not yet have obtained a CRL indicating that a key used to verify an algorithm A signature belongs to a newly revoked certificate). In such a case, it is highly desirable for a downstream entity to treat the update as 'Not Good' (due to the revocation) and not as 'unsigned' (which would happen if the 'Not Good' signatures were removed).

A similar argument applies to the case where a BGPSEC speaker (for some reason such as lack of viable alternatives) selects as his best route to a given prefix a route obtained via a 'Not Good' BGPSEC update message. (That is, a BGPSEC update containing only 'Not Good' signatures.) In such a case, the BGPSEC speaker should propagate a signed BGPSEC update message, adding his signature to the 'Not Good' signatures that already exist. Again, this is to ensure that 'downstream' entities are able to make an informed decision and not erroneously treat the route as unsigned. It may also be noted here that due to possible differences in RPKI data at different vantage points in the network, a BGPSEC update that was deemed 'Not Good' at an upstream BGPSEC speaker may indeed be deemed 'Good' at another BGP speaker downstream.

Therefore, it is important to note that when a BGPSEC speaker signs an outgoing update message, it is not attesting to a belief that all signatures prior to its are valid. Instead it is merely asserting that:

1. The BGPSEC speaker received the given route advertisement with the indicated NLRI and AS Path;

2. The BGPSEC speaker selected this route as the best route to the given prefix; and

3. The BGPSEC speaker chose to propagate an advertisement for this route to the peer (implicitly) indicated by the 'Target AS'

The BGPSEC update validation procedure is a potential target for denial of service attacks against a BGPSEC speaker. To mitigate the effectiveness of such denial of service attacks, BGPSEC speakers should implement an update validation algorithm that performs expensive checks (e.g., signature verification) after less expensive checks (e.g., syntax checks). The validation algorithm specified in Section 5.1 was chosen so as to perform checks which are likely to be expensive after checks that are likely to be inexpensive. However, the relative cost of performing required validation steps may vary between implementations, and thus the algorithm specified in Section 5.1 may not provide the best denial of service protection for all implementations.

Finally, the mechanism of setting the pCount field to zero is included in this specification to enable route servers in the control path to participate in BGPSEC without increasing the effective length of the AS_PATH. However, entities other than route servers could conceivably use this mechanism (set the pCount to zero) to attract traffic (by reducing the effective length of the AS_PATH) illegitimately. This risk is largely mitigated if every BGPSEC speaker drops incoming update messages that set pCount to zero but come from a peer that is not a route server. However, note that a recipient of a BGPSEC update message in which an upstream entity that is two or more hops away set pCount to zero is unable to verify for themselves whether pCount was set to zero legitimately.

8. Contributors

8.1. Authors

Rob Austein
Dragon Research Labs
sra@hactrn.net

Steven Bellovin
Columbia University
smb@cs.columbia.edu

Randy Bush
Internet Initiative Japan
randy@psg.com

Russ Housley
Vigil Security
housley@vigilsec.com

Matt Lepinski
BBN Technologies
lepinski@bbn.com

Stephen Kent
BBN Technologies
kent@bbn.com

Warren Kumari
Google
warren@kumari.net

Doug Montgomery
USA National Institute of Standards and Technology
dougm@nist.gov

Kotikalapudi Sriram
USA National Institute of Standards and Technology
kotikalapudi.sriram@nist.gov

Samuel Weiler
Cobham
weiler+ietf@watson.org

8.2. Acknowledgements

The authors would like to thank Luke Berndt, Wes George, Sharon Goldberg, Ed Kern, Chris Morrow, Doug Maughan, Pradosh Mohapatra, Russ Mundy, Sandy Murphy, Keyur Patel, Mark Reynolds, Heather Schiller, Jason Schiller, John Scudder, Ruediger Volk and David Ward for their valuable input and review.

9. Normative References

[1] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4", RFC 4271, January 2006.

[2] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, January 2007.

[3] Scudder, J. and R. Chandra, "Capabilities Advertisement with BGP-4", RFC 5492, February 2009.

[4] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[5] Patel, K., Ward, D., and R. Bush, "Extended Message support for BGP", March 2011.

[6] Lepinski, M., Kent, S., and D. Kong, "A Profile for Route Origin Authorizations", February 2011.

[7] Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", February 2011.

[8] Kent, S., "Threat Model for BGP Path Security", June 2011.

[9] Bush, R., "BGPsec Operational Considerations", October 2011.

[10] Turner, S., "BGP Algorithms, Key Formats, & Signature Formats", December 2011.

[11] Reynolds, M., Turner, S., and S. Kent, "A Profile for BGPSEC Router Certificates, Certificate Revocation Lists, and Certification Requests", December 2011.

[12] Bush, R. and R. Austein, "The RPKI/Router Protocol", October 2011.

Author's Address

Matthew Lepinski (editor)
BBN
10 Moulton St
Cambridge, MA 55409
US

Phone: +1 617 873 5939
Email: mlepinski@bbn.com