Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 29, 2011

R. Gagliano
Cisco Systems
S. Kent
BBN Technologies
S. Turner
IECA, Inc.

February 25, 2011

Algorithm Agility Procedure for RPKI. draft-ietf-sidr-algorithm-agility-00

Abstract

This document specifies the process that Certificate Authorities (CAs) and Relying Parties (RP) participating in the Resource Public Key Infrastructure (RPKI) will need to follow to transition to a new (and probably cryptographically stronger) algorithm set. The process is expected to be completed in a time scale of months or years. Consequently, no emergency transition is specified. The transition procedure defined in this document support only a top-down migration (parent migrates before children).

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 August 29, 2011.

Table of Contents

1. Requirements notation

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 [RFC2119].

2. Introduction

The RPKI must accommodate transitions between the public keys used by CAs. Transitions of this sort are usually termed "key rollover". Planned key rollover will occur at regular intervals throughout the life of the RPKI, as each CA changes its public keys, in a non-coordinated fashion. (By non-coordinated we mean that the time at which each CA elects to change its keys is locally determined, not coordinated across the RPKI.) Moreover, because a key change might be necessitated by suspected private key compromise, one can never assume coordination of these events among all of the CAs in the RPKI. In an emergency key rollover, the old certificate is revoked and a new certificate with a new key is issued. The mechanisms to perform a key rollover in RPKI (either planned or in an emergency), while maintaining the same algorithm suite, are covered in [I-D.ietf-sidr-keyroll].

This document describes the mechanism to perform a key rollover in RPKI due to the migration to a new signature algorithm suite. A signature algorithm suite encompasses both a signature algorithm (with a specified key size range) and a one-way hash algorithm. It is anticipated that the RPKI will require the adoption of updated key sizes and/or different algorithm suites over time. This document treats the adoption of a new hash algorithm while retaining the current signature algorithm as equivalent to an algorithm migration, and requires the CA to change its key. Migration to a new algorithm suite will be required in order to maintain an acceptable level of cryptographic security and protect the integrity of certificates, CRLs and signed objects in the RPKI. All of the data structures in the RPKI explicitly identify the signature and hash algorithms being used. However, experience has demonstrated that the ability to represent algorithm IDs is not sufficient to enable migration to new algorithm suites (algorithm agility). One also must ensure that protocols, infrastructure elements, and operational procedures also accommodate migration from one algorithm suite to another. Algorithm migration is expected to be very infrequent, but it also will require support of a "current" and "next" suite for a prolonged interval, probably several years.

This document defines how entities in the RPKI execute (planned) CA key rollover when the algorithm suite changes. The description covers actions by CAs, repository operators, and RPs. It describes the behavior required of both CAs and RPs to make such key changes work in the RPKI context, including how the RPKI repository system is used to support key rollover.

This document does not specify any algorithm suite.

A failure to comply with this process during an algorithm transition MUST be considered as non-compliance with the RPKI Certificate Policy (CP) [I-D.ietf-sidr-cp].

3. Terminology

This document assumes that the reader is familiar with the terms and concepts described in "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile" [RFC5280], "X.509 Extensions for IP Addresses and AS Identifiers" [RFC3779], and "A Profile for Resource Certificate Repository Structure" [I-D.ietf-sidr-repos-struct]. Additional terms and conventions use din examples are provided below.

Algorithm migration A planned transition from one signature and hash algorithm to a new signature and hash algorithm.

Algorithm Suite A The "current" algorithm suite used for hashing and signing, in examples in this document

Algorithm Suite B The "next" algorithm suite used for hashing and signing, used in examples in this document

Algorithm Suite C The "old" algorithm suite used for hashing and signing, used in examples in this document

CA X The CA that issued CA Y's certificate (i.e., CA Y's parent), used in examples this document.

CA Y The CA that is changing keys and/or algorithm suites, used in examples this document

CA Z A CA that is a "child" of CA Y, used in examples this document

Certificate re-issuance (unilateral) a CA MAY reissue a certificate to a subordinate Subject without the involvement of the Subject. The public key, resource extensions, and most other fields (see Section X.X) are copied from the current Subject certificate into the next Subject certificate. The Issuer name MAY change, if necessary to reflect the Subject name in the CA certificate under which the reissued certificate will be validated. The validity interval also MAY be changed. This action is defined as a unilateral certificate re-issuance.

Non-Leaf CA - a CA that issues certificates to entities not under its administrative control.

PoP (proof of possession) - execution of a protocol that demonstrates to an issuer that a subject requesting a certificate possesses the private key corresponding to the public key in the certificate submitted by the subject.

Signed Product Set (or Set) - a collection of certificates, signed objects, a CRL and a manifest that are associated by virtue of being verifiable under the same parent CA certificate

4. Key Rollover steps for algorithm migration

The "current" RPKI algorithm suite (Suite A) is defined in the RPKI's CP document , by reference to [I-D.ietf-sidr-rpki-algs]. When a migration of the RPKI algorithm suite is needed, the first step MUST be an update of the [I-D.ietf-sidr-rpki-algs] document that will include all the information described in Section 4.3.

4.1. Milestones definition

CA Ready Algorithm B Date - After this date, all (non-leaf) CAs MUST be ready to process a request from a child CA to issue a certificate under the Algorithm B suite.

CA Go Algorithm B Date - After this date, all (non-leaf) CAs MUST have re-issued all of its signed product set under the Algorithm B suite.

RP Ready Algorithm B Date - After this date, all RPs MUST be prepared to process signed material issued under the Algorithm B suite.

Twilight Algorithm B - After this date, a CA MAY cease issuing signed products under the Algorithm A suite. Also, after this date, a RP MAY cease to validate signed materials issued under the Algorithm A suite.

End Of Life (EOL) Algorithm A - After this date every CA MUST NOT generate certificates, CRLs, or other RPKI signed objects under the Algorithm A suite. Also, after this date, no RP SHOULD accept as valid any certificate, CRL or signed object using the Algorithm A suite.

4.2. Process overview

The migration process described in this document involves a series of steps that MUST be executed in chronological order by CAs and RPs. The only milestone that affects both CAs and RPs, at the same moment is the EOL date. Due to the decentralized nature of the RPKI infrastructure, it is expected that the process will take several months or even years.

In order to facilitate the transition, CAs will start issuing certificates using the Algorithm B in a hierarchical top-down order. In our example, CA Y will issue certificates using the Algorithm B suite only after CA X has started to do so (CA Y Ready Algorithm B Date > CA X Ready Algorithm B Date). This ordered transition avoids the existence of mixed certificates. In RPKI an algorithm suite MUST

NOT sign a certificate carrying a subject key that corresponds to another algorithm suite.

The following figure gives an overview of the process:

Process for RPKI CAs:

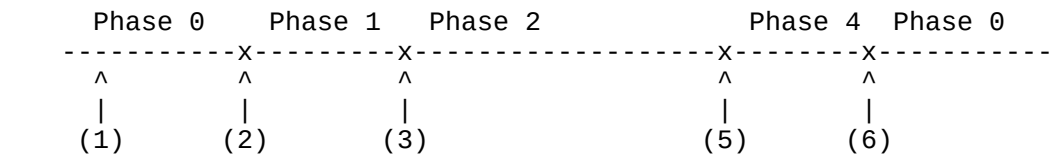

Process for RPKI RPs:

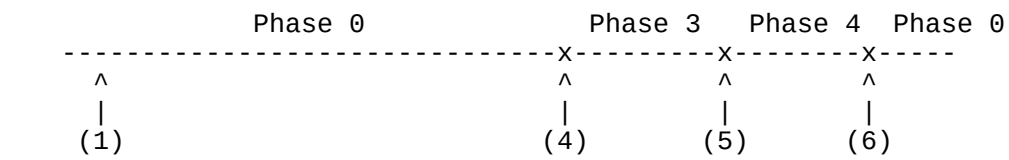

(1) RPKI's algorithm document updated.
(2) CA Ready Algorithm B Date
(3) CA Go Algorithm B Date
(4) RP Ready Algorithm B Date
(5) Twilight Date
(6) End Of Live (EOL) Date

4.3. Phase 0

Phase 0 is the initial phase of the process, during which the algorithm suite A is the only supported algorithm suite in RPKI.

The first milestone, which will initiate the migration process, is updating the [I-D.ietf-sidr-rpki-algs] document with the following definitions for the RPKI:

- Algorithm Suite A
- Algorithm Suite B
- CA Ready Algorithm B Date
- CA Go Algorithm B Date
- RP Ready Algorithm B Date

- Twilight Date
- EOL Date

All Dates MUST be represented using the local UTC date-time format specified in [RFC3339].

As an example, during Phase 0, CAs X, Y and Z are required to generate signed product sets using only the Algorithm Suite A. Also, RPs are required to validate signed product sets issued using only Algorithm Suite A.

```
CA X-Certificate-Algorithm-Suite-A (Cert-XA)
        |
        |-> CA-Y-Certificate-Algorithm-Suite-A (Cert-YA)
                |-> CA-Z-Certificate-Algorithm-Suite-A (Cert-ZA)
            |-> CA-Z-CRL-Algorithm-Suite-A (CRL-ZA)
                        |-> CA-Z-Signed-Objects-Algorithm-Suite-A
                |-> CA-Y-CRL-Algorithm-Suite-A (CRL-YA)
                |-> CA-Y-Signed-Objects-Algorithm-Suite-A
        |-> CA-X-CRL-Algorithm-Suite-A (CRL-XA)
        |-> CA-X-Signed-Objects-Algorithm-Suite-A
```

Note: Cert-XA represent the certificate for CA X, that is signed using the algorithm suite A.

4.4. Phase 1

Phase 1 starts at the CA Ready Algorithm B Date. During the Phase 1, all (non-leaf) CAs MUST be ready to process a request from a child CA to issue or revoke a certificate using the Algorithm B suite.

As the transition will happen using a (hierarchical) top-down model, a child CA will be able to issue certificates using the Algorithm B suite only after its parent CA has issued its own. The RPKI provisioning protocol can identify if a parent CA is capable of issuing certificates using the Algorithm Suite B, and can identify the corresponding algorithm suite in each Certificate Signing Request (see Section 5).

The following figure shows the status of repository entries for the three example CAs during this Phase. Two distinctive certificate chains are maintained and CA Z has not yet requested any material using the Algorithm B suite.

```
CA X-Certificate-Algorithm-Suite-A (Cert-XA)
        |
        |-> CA-Y-Certificate-Algorithm-Suite-A (Cert-YA)
                |-> CA-Z-Certificate-Algorithm-Suite-A (Cert-ZA)
            |-> CA-Z-CRL-Algorithm-Suite-A (CRL-ZA)
                        |-> CA-Z-Signed-Objects-Algorithm-Suite-A
                |-> CA-Y-CRL-Algorithm-Suite-A (CRL-YA)
                |-> CA-Y-Signed-Objects-Algorithm-Suite-A
        |-> CA-X-CRL-Algorithm-Suite-A (CRL-XA)
        |-> CA-X-Signed-Objects-Algorithm-Suite-A

CA X-Certificate-Algorithm-Suite-B (Cert-XB)
        |
        |-> CA-Y-Certificate-Algorithm-Suite-B (Cert-YB)
                |-> CA-Y-CRL-Algorithm-Suite-B (CRL-YB)
                |-> CA-Y-Signed-Objects-Algorithm-Suite-B
        |-> CA-X-CRL-Algorithm-Suite-B (CRL-XB)
        |-> CA-X-Signed-Objects-Algorithm-Suite-B
```

4.5. Phase 2

Phase 2 starts at the CA Go Algorithm B Date. During this phase all signed product sets MUST be available using both Algorithm Suite A and Algorithm Suite B. During this phase, RPs MUST be prepared to validate sets issued using Algorithm Suite A and MAY be prepared to validate sets issued using the Algorithm Suite B.

An RP that validates all signed product sets using both Algorithm Suite A or Algorithm Suite B, SHOULD expect the same results.

The following figure shows the status of the repository entries for the three example CAs during this phase, where all signed objects are available using both algorithm suites.

```
   CA X-Certificate-Algorithm-Suite-A (Cert-XA)
           |
           |-> CA-Y-Certificate-Algorithm-Suite-A (Cert-YA)
                   |-> CA-Z-Certificate-Algorithm-Suite-A (Cert-ZA)
               |-> CA-Z-CRL-Algorithm-Suite-A (CRL-ZA)
                           |-> CA-Z-Signed-Objects-Algorithm-Suite-A
                   |-> CA-Y-CRL-Algorithm-Suite-A (CRL-YA)
                   |-> CA-Y-Signed-Objects-Algorithm-Suite-A
           |-> CA-X-CRL-Algorithm-Suite-A (CRL-XA)
           |-> CA-X-Signed-Objects-Algorithm-Suite-A

   CA X-Certificate-Algorithm-Suite-B (Cert-XB)
           |
           |-> CA-Y-Certificate-Algorithm-Suite-B (Cert-YB)
                   |-> CA-Z-Certificate-Algorithm-Suite-B (Cert-ZB)
               |-> CA-Z-CRL-Algorithm-Suite-B (CRL-ZB)
                           |-> CA-Z-Signed-Objects-Algorithm-Suite-B
                   |-> CA-Y-CRL-Algorithm-Suite-B (CRL-YB)
                   |-> CA-Y-Signed-Objects-Algorithm-Suite-B
           |-> CA-X-CRL-Algorithm-Suite-B (CRL-XB)
           |-> CA-X-Signed-Objects-Algorithm-Suite-B
```

4.6. Phase 3

Phase 3 starts at the RP Ready Algorithm B Date. During this phase, all signed product sets are available using both algorithm suites and all RPs MUST be able to validate them using either suite. An object that validates using either Algorithm Suite A or Algorithm Suite B MUST be consider as valid.It is RECOMMENDED that RPs utilize only Suite B for validation during this phase, in preparation for Phase 4.

There are no changes to the CA behavior during this phase.

4.7. Phase 4

Phase 4 starts at the Algorithm B Twilight Date. At that date, the Algorithm A is labeled as "old" and the Algorithm B is labeled as "current":

```
   Before Twilight                                  -->     After Twilight

   Algorithm Suite A ("current")   -->             Algorithm Suite C ("old")
   Algorithm Suite B ("new")               -->             Algorithm Suite A ("c
urrent")
```

During this phase, all signed product sets MUST be issued using the Algorithm Suite A and MAY be issued using the Algorithm Suite C. All signed products sets issued using the Algorithm Suite A MUST be published at their corresponding publication point but signed

products sets issued using the Algorithm Suite C MAY be published at their corresponding publication points. Also, every RP MUST validate signed product sets using the Algorithm Suite A but also MAY validate signed product sets using the Algorithm Suite C.

The following figure describe a possible status for the repositories of the example CAs. In this case, CA Z no longer issues signed products using the Algorithm Suite C.

```
CA X-Certificate-Algorithm-Suite-C (Cert-XC)
        |
        |-> CA-Y-Certificate-Algorithm-Suite-C (Cert-YC)
                |-> CA-Y-CRL-Algorithm-Suite-C (CRL-YC)
                |-> CA-Y-Signed-Objects-Algorithm-Suite-C
        |-> CA-X-CRL-Algorithm-Suite-C (CRL-XC)
        |-> CA-X-Signed-Objects-Algorithm-Suite-C

CA X-Certificate-Algorithm-Suite-A (Cert-XA)
        |
        |-> CA-Y-Certificate-Algorithm-Suite-A (Cert-YA)
                |-> CA-Z-Certificate-Algorithm-Suite-A (Cert-ZA)
            |-> CA-Z-CRL-Algorithm-Suite-A (CRL-ZA)
                        |-> CA-Z-Signed-Objects-Algorithm-Suite-A
                |-> CA-Y-CRL-Algorithm-Suite-A (CRL-YA)
                |-> CA-Y-Signed-Objects-Algorithm-Suite-A
        |-> CA-X-CRL-Algorithm-Suite-A (CRL-XA)
        |-> CA-X-Signed-Objects-Algorithm-Suite-A
```

4.8. Return to Phase 0

Phase 0 starts at the EOL Algorithm Date. At this phase, ALL signed product sets using Algorithm Suite C MUST be considered invalid. CAs MUST neither issue nor publish signed products using Algorithm Suite C.

This phase closes the loop as Algorithm Suite A is the only required algorithm suite in RPKI.

5. Multi Algorithm support in the RPKI provisioning protocol

The migration described in this document is a top-down process, where two synchronization issues need to be solved between child and parent CAs:

- A child CA needs to identify which algorithm suites are supported by its parent CA
- A child CA needs to identify which algorithm suite should be used to sign a Certificate Signing Request (CSR)

The RPKI provisioning protocol [I-D.ietf-sidr-rescerts-provisioning] supports multiple algorithms suites by implementing a different resource classes for each suite.Several different resource classes also may use the same algorithm suite for different resource sets.

A child CA that wants to identify which algorithm suites are supported by its parent CA MUST perform the following tasks:

1. Establish a provisioning protocol session with its parent CA
2. Perform a "list" command as described in Section 3.1.1 of [I-D.ietf-sidr-rescerts-provisioning]
3. From the Payload in the "list response" resource class, extract the "issuer's certificate" for each class. The Algorithm Suite for each class will match the Algorithm Suite used to issue the corresponding "issuer's certificate".

A child CA that wants to specify an Algorithm Suite to its parent CA (e.g., in a certificate request) MUST perform the following tasks:

1. Perform the tasks to identify the resource class for each Algorithm Suite supported by its parent CA (as above).
2. Identify the corresponding resource class in the appropriate provisioning protocol command (e.g. "issue" or "revoke")

Upon receipt of a certificate request from a child CA, a parent CA will verify the PoP of the private key. If a child CA requests issuing a certificate using an algorithm suite that does not match a resource class, the PoP validation will fail and the request will not be performed.

6. Validation of multiple instance of signed products

During Phases 1,2,3 and 4, two algorithm suites will be simultaneously valid in RPKI. In this section, we describe the RP behavior when validating instances of the same signed product but signed with different algorithm suites. As a general rule, the validation of signed products using different algorithm suites are independent and the RP MUST NOT keep any relationship between the different hierarchies.

During Phase 1 two (corresponding) files for an object MAY be available for each signed product, one signed under Algorithm Suite A and one under Algorithm Suite B. When an RP validates these signed products, if either instance of an object validates, the product is accepted. A failure to validate one instance of a product, under either algorithm Suite MUST NOT cause the RP to reject the other instance of the product. Because both instances of such products MUST contain the same resources, relying on either instance will yield the same outcome.

During Phases 2 and 3 of this process, two (corresponding) instances of all signed products MUST be available to RPs. As in Phase 1, when an RP validates these signed products, if either instance validates, the product is accepted. A failure to validate one instance of a product, under either algorithm Suite MUST NOT cause the RP to reject the other instance of the product. Also, as above, if only one instance of a signed product can be validated, subordinate products issued under the other (non-validated) algorithm suite cannot be used, and thus need not be processed (or even retrieved).

During Phase 4 two (corresponding) files for an object MAY be available for each signed product, one signed under Algorithm Suite A and one under Algorithm Suite C. When an RP validates these signed products, if either instance of an object validates, the product is accepted. A failure to validate one instance of a product, under either algorithm Suite MUST NOT cause the RP to reject the other instance of the product. Because both instances of such products MUST contain the same resources, relying on either instance will yield the same outcome.

7. Revocations

As the algorithm migration process mandates the maintenance of two parallel certificate hierarchies, revocations requests for each algorithm suite MUST be handled independently. A Child CA MUST request revocation of a certificate relative to a specific algorithm suite.

During phase 2 and phase 3, the two parallel certificate hierarchies are designed to carry identical information. Consequently, a child CA requesting the revocation of a certificate during these two phases MUST perform that request for both algorithm suites (A and B). A non-leaf CA is NOT responsible to verify that its child CAs comply with this requirement.

8. Key rollover

Key rollover (without algorithm changes) is effected independently for each algorithm suite and MUST follow the process described in [I-D.ietf-sidr-keyroll].

9. Repository structure

The two parallel hierarchies that will exist during the transition process SHOULD have independent publications points. The repository structures for each algorithm suite are described in [I-D.ietf-sidr-repos-struct].

10. IANA Considerations

No IANA requirements

11. Security Considerations

An algorithm transition in RPKI should be a very infrequent event and it requires wide community consensus. The events that may lead to an algorithm transition may be related to a weakness of the cryptographic strength of the algorithm suite in use by RPKI, which is normal to happen over time.The procedure described in this document will take months or years to complete an algorithm transition. During that time, the RPKI system will be vulnerable to any cryptographic weakness that may have triggered this procedure.

This document does not describe an emergency mechanism for algorithm migration. Due to the distributed nature of RPKI, and the very large number of CAs and RPs, the authors do not believe it is feasible to effect an emergency algorithm migration procedure.

If a CA does not complete its migration to the new algorithm suite as described in this document (after the EOL of the "old" algorithm suite), its signed product set will not longer be valid. Consequently, the RPKI may, at the end of Phase 4, have a smaller number of valid signed products than before starting the process. Conversely, a RP that does not follow this process will lose the ability to validate signed products issued under the new algorithm suite. The resulting incomplete view of routing info from the RPKI (as a result of a failure by CAs or RPs to complete the transition) could degrade routing in the public Internet.

12. Acknowledgements

The authors would like to acknowledge the work of the SIDR working group co-chairs (Sandra Murphy and Chris Morrow) as well as the contributions given by Geoff Huston.

13. References

13.1. Normative References

[I-D.ietf-sidr-cp]
Kent, S., Kong, D., Seo, K., and R. Watro, "Certificate Policy (CP) for the Resource PKI (RPKI", draft-ietf-sidr-cp-16 (work in progress), December 2010.

[I-D.ietf-sidr-keyroll]
Huston, G., Michaelson, G., and S. Kent, "CA Key Rollover in the RPKI", draft-ietf-sidr-keyroll-05 (work in progress), December 2010.

[I-D.ietf-sidr-repos-struct]
Huston, G., Loomans, R., and G. Michaelson, "A Profile for Resource Certificate Repository Structure", draft-ietf-sidr-repos-struct-06 (work in progress), November 2010.

[I-D.ietf-sidr-res-certs]
Huston, G., Michaelson, G., and R. Loomans, "A Profile for X.509 PKIX Resource Certificates", draft-ietf-sidr-res-certs-21 (work in progress), December 2010.

[I-D.ietf-sidr-rescerts-provisioning]
Huston, G., Loomans, R., Ellacott, B., and R. Austein, "A Protocol for Provisioning Resource Certificates", draft-ietf-sidr-rescerts-provisioning-09 (work in progress), November 2010.

[I-D.ietf-sidr-rpki-algs]
Huston, G., "A Profile for Algorithms and Key Sizes for use in the Resource Public Key Infrastructure", draft-ietf-sidr-rpki-algs-04 (work in progress), November 2010.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2560] Myers, M., Ankney, R., Malpani, A., Galperin, S., and C. Adams, "X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP", RFC 2560, June 1999.

[RFC3339] Klyne, G., Ed. and C. Newman, "Date and Time on the Internet: Timestamps", RFC 3339, July 2002.

[RFC3779] Lynn, C., Kent, S., and K. Seo, "X.509 Extensions for IP Addresses and AS Identifiers", RFC 3779, June 2004.

[RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.

[RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, May 2008.

13.2. Informative References

[RFC5781] Weiler, S., Ward, D., and R. Housley, "The rsync URI Scheme", RFC 5781, February 2010.

Appendix A. Change Log

From individual submission to WG item:

1. Change form "laisez faire" to "top-down"
2. Included Multi Algorithm support in the RPKI provisioning protocol
3. Included Validation of multiple instance of signed products
4. Included Revocations
5. Included Key rollover
6. Included Repository structure
7. Included Security Considerations
8. Included Acknowledgements

Authors' Addresses

Roque Gagliano
Cisco Systems
Avenue des Uttins 5
Rolle, 1180
Switzerland

Email: rogaglia@cisco.com

Stephen Kent
BBN Technologies
10 Moulton St.
Cambridge, MA 02138
USA

Email: kent@bbn.com

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

Email: turners@ieca.com

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 12, 2011

R. Bush
Internet Initiative Japan

March 11, 2011

The RPKI Ghostbusters Record
draft-ietf-sidr-ghostbusters-03

Abstract

In the Resource Public Key Infrastructure (RPKI), resource certificates completely obscure names or any other information which might be useful for contacting responsible parties to deal with issues of certificate expiration, maintenance, roll-overs, compromises, etc. This draft describes the RPKI Ghostbusters Record containing human contact information to be signed (indirectly) by a resource-owning certificate. The data in the record are those of a severely profiled vCARD.

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 [RFC2119].

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 September 12, 2011.

Copyright Notice

Table of Contents

1. Introduction

In the operational use of the RPKI it can become necessary to contact, human to human, the party responsible for a resource-owning certificate. An important example is when the owner of a Route Origin Authorization (ROA) sees a problem, or an impending problem, with a certificate or CRL in the path between the ROA and a trust anchor. E.g., a certificate along that path has expired, is soon to expire, or a CRL associated with a CA along the path is stale, thus placing the quality of the routing of the address space described by the ROA in jeopardy.

As the names in RPKI certificates are intentionally hashes which are not meaningful to humans, see [I-D.ietf-sidr-cp], there is no way to use a certificate itself to lead to the worrisome certificate's or CRL's maintainer. So, "Who do you call?"

This document specifies the RPKI Ghostbusters Record, an object signed, indirectly via an End Entity (EE) certificate, by the certificate whose maintainer may be contacted using the human usable payload information in the Ghostbusters Record.

The Ghostbusters Record conforms to the syntax defined in [I-D.ietf-sidr-signed-object].

Note that the Ghostbusters Record is not an identity certificate, but rather an attestation to the contact data made by the issuer of the certificate signing the Ghostbusters Record.

This record is not meant to supplant or be used as resource registry whois data. It gives information about an RPKI certificate maintainer not a resource holder.

This specification has three main sections. The first, Section 4, is the format of the contact payload information, a severely profiled vCARD. The second, Section 5, profiles the packaging of the payload as a profile of the RPKI Signed Object Template specification [I-D.ietf-sidr-signed-object]. The third, Section 6, describes the proper validation of the signed Ghostbusters Record.

2. Suggested Reading

It is assumed that the reader understands the RPKI, [I-D.ietf-sidr-arch], the RPKI Repository Structure, [I-D.ietf-sidr-repos-struct], Signed RPKI Objects, [I-D.ietf-sidr-signed-object], and vCARDs [RFC2426].

3. RPKI Ghostbusters Record Payload Example

An example of an RPKI Ghostbusters Record payload with all types populated is as follows:

```
BEGIN:vCard
VERSION:3.0
FN:Human's Name
N:Name;Human's;Ms.;Dr.;OCD;ADD
ORG:Organizational Entity
ADR;TYPE=WORK:;;42 Twisty Passage;Deep Cavern; WA; 98666;U.S.A.
TEL;TYPE=VOICE,MSG,WORK:+1-666-555-1212
TEL;TYPE=FAX,WORK:+1-666-555-1213
EMAIL;TYPE=INTERNET:human@example.com
END:vCard
```

4. vCARD Profile

The goal in profiling the vCARD is not to include as much information as possible, but rather to include as few types as possible while providing the minimal necessary data to enable one to contact the maintainer of the RPKI data which threatens the ROA[s] of concern.

The Ghostbusters vCARD payload is a minimalist subset of the vCARD as described in [RFC2426].

BEGIN - pro forma packaging which MUST be the first line in the vCARD and MUST have the value "BEGIN:vCARD" as described in [RFC2426].

VERSION - pro forma packaging which MUST be the second line in the vCARD and MUST have the value "VERSION:3.0" as described in 3.6.9 of [RFC2426].

FN - the name, as described in 3.1.1 of [RFC2426], of a contactable person who responsible for the certificate.

N - the components of the name of the object the vCard represents, as described in 3.1.2 of [RFC2426].

ORG - an organization as described in 3.5.5 of [RFC2426].

ADR - a postal address as described in 3.2.1 of [RFC2426].

TEL - a voice and/or fax phone as described in 3.3.1 of [RFC2426].

EMAIL - an Email address as described in 3.3.2 of [RFC2426]

END - pro forma packaging which MUST be the last line in the vCARD and MUST have the value "END:vCARD" as described in [RFC2426].

Per [RFC2426], the BEGIN, VERSION, FN, N, and END types MUST be included in a record. To be useful, one or more of ADR, TEL, and EMAIL MUST be included. Other types MAY NOT be included.

5. CMS Packaging

The Ghostbusters Record is a CMS signed-data object conforming to the RPKI Signed Data Object Template, [I-D.ietf-sidr-signed-object].

The ContentType of a Ghostbusters Record is defined as rpkiGhostbusters, and has the numerical value of [TO BE ASSIGNED]. This OID MUST appear both within the eContentType in the encapContentInfo object as well as the ContentType signed attribute in the signerInfo object. See [I-D.ietf-sidr-signed-object].

eContent: The content of a Ghostbusters Record is described above in Section 4 above.

Similarly to a ROA, the Ghostbusters Record is verified using an EE certificate issued under the CA certificate associated with the resource-holding certificate whose maintainer is described in the vCARD.

The EE certificate used to verify the Ghostbusters Record is the one that appears in the CMS data structure that contains the payload defined above.

6. Validation

The validation procedure defined in Section 3 of [I-D.ietf-sidr-signed-object] is applied to a Ghostbusters Record. After this procedure has been performed, the Version number type within the payload is checked, and the OCTET STRING containing the vCARD data is extracted. These data are checked against the profile defined in Section 4 of this document. Only if all of these checks pass is the Ghostbusters payload deemed valid and made available to the application that requested the payload.

7. Security Considerations

Though there is no on the wire protocol in this specification, there are attacks which could abuse the data described. As the data, to be useful, need to be public, little can be done to avoid this exposure.

Phone Numbers: The vCARDs may contain real world telephone numbers which could be abused for telemarketing, abusive calls, etc.

Email Addresses: The vCARDs may contain Email addresses which could be abused for purposes of spam.

Relying parties are warned that the data in a Ghostbusters Record are self-asserted. These data have not been verified by the CA that issued a (CA) certificate to the entity that issued the EE certificate used to validate the Ghostbusters Record.

8. IANA Considerations

This document has no IANA Considerations.

9. Acknowledgments

The author wishes to thank Russ Housley, the authors of [I-D.ietf-sidr-signed-object], Stephen Kent, and Michael Elkins for their contributions.

10. References

10.1. Normative References

[I-D.ietf-sidr-signed-object] Lepinski, M., Chi, A., and S. Kent, "Signed Object Template for the Resource Public Key Infrastructure", draft-ietf-sidr-signed-object-03 (work in progress), February 2011.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2426] Dawson, F. and T. Howes, "vCard MIME Directory Profile", RFC 2426, September 1998.

10.2. Informative References

[I-D.ietf-sidr-arch]
Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", draft-ietf-sidr-arch-12 (work in progress), February 2011.

[I-D.ietf-sidr-cp]
Kent, S., Kong, D., Seo, K., and R. Watro, "Certificate Policy (CP) for the Resource PKI (RPKI", draft-ietf-sidr-cp-16 (work in progress), December 2010.

[I-D.ietf-sidr-repos-struct]
Huston, G., Loomans, R., and G. Michaelson, "A Profile for Resource Certificate Repository Structure", draft-ietf-sidr-repos-struct-07 (work in progress), February 2011.

Author’s Address

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

Phone: +1 206 780 0431 x1
Email: randy@psg.com

Network Working Group
Internet-Draft
Intended status: BCP
Expires: September 11, 2011

R. Bush
Internet Initiative Japan
March 10, 2011

RPKI-Based Origin Validation Operation
draft-ietf-sidr-origin-ops-06

Abstract

Deployment of RPKI-based BGP origin validation has many operational considerations. This document attempts to collect and present them. It is expected to evolve as RPKI-based origin validation is deployed and the dynamics are better understood.

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 [RFC2119].

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 September 11, 2011.

Copyright Notice

Table of Contents

1. Introduction

RPKI-based origin validation relies on widespread propagation of the Resource Public Key Infrastructure (RPKI) [I-D.ietf-sidr-arch]. How the RPKI is distributed and maintained globally is a serious concern from many aspects.

The global RPKI is in very initial stages of deployment, there is no root trust anchor, initial testing is being done by the IANA and some RIRs, and there is a technical testbed. It is thought that origin validation based on the RPKI will be deployed incrementally over the next year to five years.

Origin validation only need be done by an AS's border routers and is designed so that it can be used to protect announcements which are originated by large providers, upstreams and downstreams, and by small stub/enterprise/edge routers.

Origin validation has been designed to be deployed on current routers without significant hardware upgrade. It should be used by everyone from large backbones to small stub/entetprise/edge routers.

RPKI-based origin validation has been designed so that, with prudent local routing policies, there is little risk that what is seen as today's normal Internet routing is threatened by imprudent deployment of the global RPKI, see Section 5.

2. Suggested Reading

It is assumed that the reader understands BGP, [RFC4271], the RPKI, see [I-D.ietf-sidr-arch], the RPKI Repository Structure, see [I-D.ietf-sidr-repos-struct], ROAs, see [I-D.ietf-sidr-roa-format], the RPKI to Router Protocol, see [I-D.ietf-sidr-rpki-rtr], RPKI-based Prefix Validation, see [I-D.ietf-sidr-pfx-validate], and Ghostbuster Records, see [I-D.ietf-sidr-ghostbusters].

3. RPKI Distribution and Maintenance

The RPKI is a distributed database containing certificates, CRLs, manifests, ROAs, and Ghostbuster Records as described in [I-D.ietf-sidr-repos-struct]. Policies and considerations for RPKI object generation and maintenance are discussed elsewhere.

A local valid cache containing all RPKI data may be gathered from the global distributed database using the rsync protocol and a validation tool such as rcynic.

Validated caches may also be created and maintained from other validated caches. Network operators SHOULD take maximum advantage of this feature to minimize load on the global distributed RPKI database.

As RPKI-based origin validation relies on the availability of RPKI data, operators SHOULD locate caches close to routers that require these data and services. A router can peer with one or more nearby caches.

For redundancy, a router SHOULD peer with more than one cache at the same time. Peering with two or more, at least one local and others remote, is recommended.

If an operator trusts upstreams to carry their traffic, they SHOULD also trust the RPKI data those upstreams cache, and SHOULD peer with those caches. Note that this places an obligation on those upstreams to maintain fresh and reliable caches.

A transit provider or a network with peers SHOULD validate origins in announcements made by upstreams, downstreams, and peers. They still SHOULD trust the caches provided by their upstreams.

Before issuing a ROA for a block, an operator MUST ensure that any sub-allocations from that block which are announced by other ASs, e.g. customers, have correct ROAs in play. Otherwise, issuing a ROA for the super-block will cause the announcements of sub-allocations with no ROAs to be Invalid.

An environment where private address space is announced in eBGP the operator MAY have private RPKI objects which cover these private spaces. This will require a trust anchor created and owned by that environment, see [I-D.ietf-sidr-ltamgmt].

Operators issuing ROAs may have customers announce their own prefixes and ASs into global eBGP but who do not wish to go though the work to manage the relevant certificates and ROAs. The operator SHOULD provision the RPKI data for these customers just as they provision many other things for them.

4. Within a Network

Origin validation need only be done by edge routers in a network, those which border other networks/ASs.

A validating router will use the result of origin validation to influence local policy within its network, see Section 5. In

deployment this policy should fit into the AS's existing policy, preferences, etc. This allows a network to incrementally deploy validation capable border routers.

eBGP speakers which face more critical peers or up/downstreams would be candidates for the earliest deployment. Validating more critical received announcements should be considered in partial deployment.

5. Routing Policy

Origin validation based on the RPKI merely marks a received announcement as having an origin which is Valid, NotFound, or Invalid. See [I-D.ietf-sidr-pfx-validate]. How this is used in routing SHOULD be specified by the operator's local policy.

Local policy using relative preference is suggested to manage the uncertainty associated with a system in early deployment, applying local policy to eliminate the threat of unroutability of prefixes due to ill-advised certification policies and/or incorrect certification data. E.g. until the community feels comfortable relying on RPKI data, routing on Invalid origin validity, though at a low preference, MAY occur.

As origin validation will be rolled out incrementally, coverage will be incomplete for a long time. Therefore, routing on NotFound validity state SHOULD be done for a long time. As the transition moves forward, the number of BGP announcements with validation state NotFound should decrease. Hence an operator's policy SHOULD NOT be overly strict, preferring Valid announcements, attaching a lower preference to, but still using, NotFound announcements, and dropping or giving very low preference to Invalid announcements.

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

When using a metric which is also influenced by other local policy, the operator should be careful not to create privilege upgrade vulnerabilities. E.g. if Local Pref is set depending on validity state, be careful that peer community signaling MAY NOT upgrade an invalid announcement to valid or better.

Announcements with Valid origins SHOULD be preferred over those with NotFound or Invalid origins, if the latter are accepted at all.

Announcements with NotFound origins SHOULD be preferred over those with Invalid origins.

Announcements with Invalid origins MAY be used, but SHOULD be less preferred than those with Valid or NotFound.

6. Notes

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

There is some uncertainty about the origin AS of aggregates and what, if any, ROA can be used. The long range solution to this is the deprecation of AS-SETs, see [I-D.wkumari-deprecate-as-sets].

Operators which manage certificates SHOULD have RPKI Ghostbuster Records (see [I-D.ietf-sidr-ghostbusters]), signed indirectly by End Entity certificates, for those certificates on which others' routing depends for certificate and/or ROA validation.

7. Security Considerations

As the BGP origin is not signed, origin validation is open to malicious spoofing. It is only designed to deal with inadvertent mis-advertisement.

Origin validation does not address the problem of AS-Path validation. Therefore paths are open to manipulation, either malicious or accidental.

The data plane may not follow the control plane.

Be aware of the class of privilege escalation issues discussed in Section 5 above.

8. IANA Considerations

This document has no IANA Considerations.

9. Acknowledgments

The author wishes to thank Rob Austein, Steve Bellovin, Pradosh Mohapatra, Chris Morrow, Sandy Murphy, Keyur Patel, Heather and Jason Schiller, John Scudder, Maureen Stillman, and Dave Ward.

10. References

10.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[I-D.ietf-sidr-arch] Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", draft-ietf-sidr-arch-12 (work in progress), February 2011.

[I-D.ietf-sidr-repos-struct] Huston, G., Loomans, R., and G. Michaelson, "A Profile for Resource Certificate Repository Structure", draft-ietf-sidr-repos-struct-07 (work in progress), February 2011.

[I-D.ietf-sidr-roa-format] Lepinski, M., Kent, S., and D. Kong, "A Profile for Route Origin Authorizations (ROAs)", draft-ietf-sidr-roa-format-10 (work in progress), February 2011.

[I-D.ietf-sidr-rpki-rtr] Bush, R. and R. Austein, "The RPKI/Router Protocol", draft-ietf-sidr-rpki-rtr-10 (work in progress), March 2011.

[I-D.ietf-sidr-pfx-validate] Mohapatra, P., Scudder, J., Ward, D., Bush, R., and R. Austein, "BGP Prefix Origin Validation", draft-ietf-sidr-pfx-validate-01 (work in progress), February 2011.

[I-D.ietf-sidr-ghostbusters] Bush, R., "The RPKI Ghostbusters Record", draft-ietf-sidr-ghostbusters-00 (work in progress), December 2010.

[I-D.ietf-sidr-ltamgmt] Kent, S. and M. Reynolds, "Local Trust Anchor Management for the Resource Public Key Infrastructure", draft-ietf-sidr-ltamgmt-00 (work in progress), November 2010.

10.2. Informative References

[RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.

[I-D.wkumari-deprecate-as-sets] Kumari, W., "Deprecation of BGP AS_SET, AS_CONFED_SET.", draft-wkumari-deprecate-as-sets-01 (work in progress), September 2010.

Author's Address

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

Phone: +1 206 780 0431 x1
Email: randy@psg.com

Secure Inter-Domain Routing
Internet-Draft
Intended status: Standards Track
Expires: August 14, 2011

S. Kent
BBN

February 10, 2011

Threat Model for BGP Path Security
draft-kent-bgpsec-threats-01.txt

Abstract

This document describes a threat model for BGP path security (BGPSEC). BGPSEC is assumed to make use of the Resource Public Key Infrastructure (RPKI) already developed in the SIDR WG [I-D.ietf-sidr-arch], and thus threats and attacks against the RPKI are part of this model. The model assumes that BGP path security is achieved through the application of digital signatures to AS_Path Info. The document characterizes classes of potential adversaries that are considered to be threats, and examines classes of attacks that might be launched against BGPSEC. It concludes with brief discussion of residual vulnerabilities.

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 July 24, 2011.

Copyright Notice

Table of Contents

1. Introduction

This document describes the security context in which BGPSEC is intended to operate. It discusses classes of potential adversaries that are considered to be threats, and classes of attacks that might be launched against BGPSEC. Because BGPSEC depends on the Resource Public Key Infrastructure (RPKI), threats and attacks against the RPKI also are discussed.

The motivation for developing BGPSEC, i.e., residual security concerns for BGP, is well described in several documents, including "BGP Security Vulnerabilities Analysis" [RFC4272] and "Design and Analysis of the Secure Border Gateway Protocol (S-BGP)" [Kent2000]. All of these papers note that BGP does not include mechanisms that allow an Autonomous System (AS) to verify the legitimacy and authenticity of BGP route advertisements. (BGP now mandates support for mechanisms to secure peer-peer communication, i.e., for the links that connect BGP routers. There are several secure protocol options to addresses this security concern, e.g., IPsec [RFC4301] and TCP-AO [RFC5925]. This document briefly notes the need to address this aspect of BGP security, but focuses on application layer BGP security issues that are addressed by BGPSEC.)

RFC 4272 succinctly notes:

> BGP speakers themselves can inject bogus routing information, either by masquerading as any other legitimate BGP speaker, or by distributing unauthorized routing information as themselves. Historically, misconfigured and faulty routers have been responsible for widespread disruptions in the Internet. The legitimate BGP peers have the context and information to produce believable, yet bogus, routing information, and therefore have the opportunity to cause great damage. The cryptographic protections of [TCPMD5] and operational protections cannot exclude the bogus information arising from a legitimate peer. The risk of disruptions caused by legitimate BGP speakers is real and cannot be ignored.

BGPSEC is intended to address the concerns cited above, to provide significantly improved path security, and to build upon the secure route origination foundation offered by use of the RPKI. Specifically, the RPKI enables relying parties (RPs) to determine of the origin AS for a path was authorized to advertise the prefix contained in a BGP update message. This security feature is enabled by the use of two types of digitally signed data: a PKI [I-D.sidr-res-certs] that associates one or more prefixes with the public key(s) of an address space holder, and Route Origination Authorizations (ROAs) [I-D.roa-format] that allows a prefix holder to

specify the AS(es) that are authorized to originate routes for a prefix.

The security model adopted for BGPSEC does not assume an "oracle" that can see all of the BGP inputs and outputs associated with every AS or every BGP router. Instead, the model is based on a local notion of what constitutes legitimate, authorized behavior by the BGP routers associated with an AS. This is an AS-centric model of secure operation, consistent with the AS-centric model that BGP employs for routing. This model forms the basis for the discussion that follows.

This document begins with a brief set of definitions relevant to the subsequent sections. It then discusses classes of adversaries that are perceived as viable threats against routing in the public Internet. It continues to explore a range of attacks that might be effected by these adversaries, against both path security and the infrastructure upon which BGPSEC relies. It concludes with a brief review of residual vulnerabilities.

2. Terminology

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.

The following security and routing terminology definitions are employed in this document.

Adversary - An adversary is an entity (e.g., a person or an organization) perceived as malicious, relative to the security policy of a system. The decision to characterize an entity as an adversary is made by those responsible for the security of a system. Often one describes classes of adversaries with similar capabilities or motivations, rather than specific individuals or organizations.

Attack - An attack is an action that attempts to violate the security policy of a system, e.g., by exploiting a vulnerability. There is often a many to one mapping of attacks to vulnerabilities, because many different attacks may be used to exploit a vulnerability.

Autonomous System - An AS is a set of one or more IP networks operated by a single administrative entity.

AS Number (ANS) - An ASN is a 2 or 4 byte number issued by a registry to identify an AS in BGP.

Certification Authority (CA) - An entity that issues digital certificates (e.g., X.509 certificates) and vouches for the binding between the data items in a certificate.

Countermeasure - A countermeasure is a procedure or technique that thwarts an attack, preventing it from being successful. Often countermeasures are specific to attacks or classes of attacks.

Border Gateway Protocol (BGP) - A path vector protocol used to convey "reachability" information among autonomous systems, in support of inter-domain routing.

False (Route) Origination - If an ISP originates a route for a prefix that the ISP does not hold (and that it has not been authorized to originate by the prefix holder, this is termed false route origination.

Internet Service Provider (ISP) - An organization managing (and, typically, selling,) Internet services to other organizations or individuals.

Internet Number Resources (INRs) - IPv4 or IPv6 address space and ASNs

Internet Registry - An organization that manages the allocation or distribution of INRs. This encompasses the Internet Assigned Number Authority (IANA), Regional Internet Registries (RIRs), National Internet Registries (NIRs), and Local Internet registries (LIRs, ISPs).

Man in the Middle (MITM) - A MITM is an entity that is able to examine and modify traffic between two (or more) parties on a communication path

Prefix - A prefix is an IP address and a mask used to specify a set of addresses that are grouped together for purposes of routing.

Public Key Infrastructure (PKI) - A PKI is a collection of hardware, software, people, policies, and procedures used to create, manage, distribute, store, and revoke digital certificates.

Relying Parties (RPs) - An RP is an entity that makes use of signed products from a PKI, i.e., relies on signed data that is verified using certificates, and CRLs from a PKI.

RPKI Repository System - The RPKI repository system consists of a distributed set of loosely synchronized databases

Resource PKI (RPKI) - A PKI operated by the entities that manage INRs, and that issues X509 certificates (and CRLs) that attest to the holdings of INRs.

RPKI Signed Object - An RPKI signed object is a Cryptographic Message Syntax (CMS)-encapsulated data object complying with the format and semantics defined in [draft-ietf-sidr-signed-object-02.txt].

Route - In the Internet, a route is a prefix and an associated sequence of ASNs that indicates a path via which traffic destined for the prefix can be directed.

Route leak - A route leak is said to occur when AS-A advertises routes that it has received from an AS-B to AS-A's neighbors, but AS-A is not viewed as a transit provider for the prefixes in the route.

Threat - A threat is a motivated, capable adversary. An adversary that is not motivated to launch an attack is not a threat. An adversary that is motivated but not capable of launching an attack also is not a threat.

Vulnerability - A vulnerability is a flaw or weakness in a system's design, implementation, or operation and management that could be exploited to violate the security policy of a system.

3. Threat Characterization

The following classes of threats are addressed in this document.

BGP speakers - A BGP speaker, e.g., an ISP or a multi-homed non-ISP subscriber, may be a threat. (For simplicity, the remainder of this document refers to BGP speakers as ISPs.) An ISP may be motivated to cause BGP routers controlled by the ISP to emit update messages with inaccurate routing info. Such updates might cause traffic to flow via paths that would otherwise be rejected as less advantageous by other ISPs. Because the ISP controls the BGP routers that it operates, it is in a position to modify their operation. Routers operated by the ISP are vehicles for mounting MITM attacks on both control and data plane traffic. If the ISP participates in the RPKI, it will have at least CA resource certificate and may be able to generate an arbitrary number of subordinate CA certificates and ROAs. It will be authorized to populate (and may even host) its own repository publication point. If it implements BGPSEC, it will have the ability to issue certificates for its routers, and to sign updates in a fashion that will be recognized by BGPSEC-enabled ISP neighbors.

Hackers - Hackers are considered a threat. Hackers might assume control of network management computers and routers operated by ISPs, including ISPs that implement BGPSEC. In such cases, hackers would be able to act as a rogue ISP (see above). It is assumed that hackers generally do not have the capability to effect MITM attacks on most links between ISPs. Hackers might be recruited, without their knowledge, by criminals or by nations, to act on their behalf.

Criminals - Criminals may be a threat. Criminals might persuade (via threats or extortion) an ISP to act as rogue ISP (see above), and thus be able to effect a wide range of attacks. Criminals might persuade telecom staff to enable MITM attacks on links between routers. The motivation for criminals may include the ability to extort money from other ISPs or ISP clients, e.g., by adversely affecting routing for these ISPs or clients. They may wish to manipulate routing to conceal the sources of spam or of DoS attacks.

Registries - Any registry in the RPKI could be a threat. Staff at the registry are capable of manipulating repository content or mismanaging RPKI certificates. These actions could adversely affect the operation of an ISP or a client of an ISP. The staff could be motivated to do this based on political pressure from the nation in which it operates (see below).

Nations - A nation may be a threat. A nation may control one or more ISPs that operate in the nation, and thus can cause them to act as rogue ISPs. A nation may have a technical active wiretapping

capability (e.g., within its territory) that enables it to effect MITM attacks on inter-ISP traffic. It may have an ability to attack and take control of routers or management network computers of ISPs in other countries. A nation may control a registry that operates within its territory, and might force the registry to act as a rogue capacity. National threat motivations include the desire to control the flow of traffic to/from the nation or to divert traffic destined for other nations (for passive or active wiretapping, including DoS). A manifest associated with a CA's repository publication point contains a list of:

4. Attacks

This section describes classes of attacks that may be effected against Internet routing. Attacks are classified based on the target of the attack, as an element of the routing system, or the routing security infrastructure on which BGPSEC relies. In general, attacks of interest are ones that attempt to violate the integrity or authenticity of BGP traffic, or which violate the authorizations associated with entities participating in the RPKI. Attacks that violate the implied confidentiality of routing traffic are not considered significant (see 4.1 below).

4.1. Active wiretapping of links between routers

An adversary may attack the links that connect BGP routers. Passive attacks are not considered, because it is assumed that most of the info carried by BGP will otherwise be accessible to adversaries. Several classes of adversaries are assumed to be capable of MITM effecting attacks against the control plane traffic. MITM attacks may be directed against BGP, BGPSEC, or against TCP or IP. Such attacks include replay of selected BGP messages, selective modification of BGP messages, and DoS attacks against BGP routers.

4.2. Attacks on a BGP router

An adversary may attack a BGP router, whether it implements BGPSEC or not. Any adversary that controls routers legitimately, or that can assume control of a router, is assumed to be able to effect the types of attacks described below. Note that any router behavior that can be ascribed to a local routing policy decision is not considered to be an attack. This is because such behavior could be explained as a result of local policy settings, and thus is beyond the scope of what BGPSEC can detect as unauthorized behavior. Thus, for example, a router may fail to propagate some or all route withdrawals or effect "route leaks". (These behaviors are not precluded by the specification for BGP, and might be the result of a local policy that

is not publicly disclosed. As a result, they are not considered attacks.)

AS Insertion: A router might insert one or more ASNs, other than its own ASN, into an update message. This violates the BGP spec and thus is considered an attack.

False (route) Origination: A router might originate a route for a prefix, when the AS that the router represents is not authorized to originate routes for that prefix. This is an attack.

Secure Path Downgrade: A router might remove signatures from a BGPSEC update that it receives, when forwarding this update to a BGPSEC-enabled neighbor. This behavior violates the BGPSEC spec and thus is considered an attack.

Invalid Signature Insertion: A router might emit a signed update with a "bad" signature, i.e., a signature that cannot be validated by other BGPSEC routers. (This might occur due to use of a revoked or expired certificate, a computational error, or a syntactic error.) This behavior violates the BGPSEC spec and thus is considered an attack.

Stale Path Announcement: An announcement may be propagated with an origination signature segment expiry value that is not current. This behavior violates the BGPSEC spec and is considered a possible replay attack.

Premature Path Announcement Expiration: A router might emit a signed update with an origin expiry time that is very short. The BGPSEC protocol specification does not mandate a minimum expiry time. However, an immediate neighbor of a route originator should expect to see an expiry time that not substantially less than XX in the future. Later routers along a path generally cannot determine if a shorter expiry time is "suspicious" since they cannot know how long a route may have been held by an earlier AS, prior to being released. Thus this consideration applies only to an immediate neighbor of a route originator.

MITM Attack: A cryptographic key used for point-to-point security (e.g., TCP-AO or IPsec) between two BGP routers might be compromised (e.g., by extraction from a router). This would enable an adversary to effect MITM attacks on the link(s) where the key is used.

Compromised Private Key: The private key associated with an RPKI EE certificate issued to a router might be compromised by an

attack against the router. An adversary with access to this key would be able to generate updates that appear to be from this router (or from any routers that share this key and certificate). If the adversary controlled another ISP, it could use this key to forge signatures that appear to come from the router(s) in question, thus making it appear that those routers were misbehaving.

Replay Attack: An update may be signed and announced, and later withdrawn. The adversary controlling intermediate routers does not propagate the withdrawal but instead re-announces (i.e., replays) the previous announcement within its expiry time if it has not yet expired.

4.3. Attacks on ISP management computers (non-CA computers)

An adversary may choose to attack computers used by an ISP to manage its network, especially its routers. Such attacks might be effected by an adversary that has compromised the security of these computers. This might be effected via remote attacks, extortion of selected ISP staff, etc. If an adversary compromises NOC computers, it can execute any management function that authorized ISP staff would have performed. Thus the adversary could modify local routing policy to change preferences, to black-hole certain routes, etc. This type of behavior cannot be externally detected as an attack.

If the ISP participates in the RPKI, the adversary could manipulate the RP tools that extract data from the RPKI, causing the output of these tools to be corrupted in various ways. For example, an attack of this sort could cause the ISP to view valid routes as not validated, which could alter its routing behavior.

If the adversary invoked the tool used to manage the repository publication point for this ISP, it could delete any objects stored there (certificates, CRLs, manifests, ROAs, or subordinate CA certificates). This could affect the routing status of entities that have allocations/assignments from this ISP (e.g., by deleting their CA certificates).

An attacker could invoke the tool used to request certificate revocation, causing router certificates, ROAs, or subordinate CA certificates to be revoked. An attack of this sort could affect not only this ISP, but also any ISPs that receive allocations/assignments from it, e.g., because their CA certificates were revoked.

It the ISP is BGPSEC-enabled, an attack of this sort could cause the affected ISP to be viewed as not BGPSEC-enabled, possibly making routes it emits be less preferred.

If an adversary invoked a tool used to request ROAs, it could effectively re-allocate some of the prefixes allocated/assigned to the ISP (e.g., by modifying the origin AS in ROAs). This might cause other BGPSEC-enabled ISPs, and other RPKI-enabled ISPs, to view the ISP as no longer originating routes for these prefixes. Multi-homed subscribers of this ISP who received a PA allocation from the ISP might find their traffic was now routed via other connections.

If the ISP is BGPSEC-enabled, and the adversary invoked a tool used to request certificates, it could replace valid certificates for routers with ones that might be rejected by BGPSEC-enabled neighbors.

4.4. Attacks on a repository publication point

A critical element of the RPKI is the repository system. An adversary might attack a repository, or a publication point within a repository, to adversely affect routing.

This section considers only those attacks that can be launched by any adversary who controls a computer hosting one or more repository publication points, without access to the cryptographic keys needed to generate valid RPKI signed products. Such attacks might be effected by an inside or an external threat. Because all repository objects are digitally signed, attacks of this sort translate into DoS attacks against the RPKI RPs. There are a few distinct forms of such attacks, as described below.

Note first that the RPKI calls for RPs to cache the data they acquire and verify from the repository system. Attacks that delete signed products, that insert products with "bad" signatures, that tamper with object signatures, or that replace newer objects with older (valid) ones, can be detected by RPs (with a few exceptions). RPs are expected to make use of the cached repository data until attacks that violate the integrity of publication points (and which are detected) are resolved. Thus the impact of such attacks is mitigated in part by the design of the repository system.

If an adversary inserts an object into a publication point, and the object has a "bad" signature, the object will not be accepted and used by RPs.

If an adversary modifies any signed product at a publication point, the signature on the product will fail, and cause RPs to not accept it. This is equivalent to deleting the object, on many respects.

If an adversary deletes one or more CA certificates, ROAs or the CA’s

CRL at a publication point, the manifest for that publication point will allow an RP to detect this attack. (The RP would be very unhappy if there is no CRL for the CA instance anyway.) An RP can continue to use the last valid instance of the deleted object as a local policy option), thus minimizing the impact of such an attack.

If an adversary deletes a manifest (and does not replace it with an older instance), that is detectable by an RP, and should result in the CA being notified of the problem. An RP can continue to use the last valid instance of the deleted object as a local policy option), thus minimizing the impact of such an attack.

If an adversary deletes newly added CA certificates or ROAs, and replaces the current manifest with the previous manifest, the manifest (and the CRL that it matches) will be "stale" (see [ietf-sidr-manifest]). This alerts an RP that there may be a problem, and, hopefully, the CA responsible for the publication point will be asked to remedy the problem (republish the missing CA certificates and/or ROAs). An RP cannot know the content of the new certificates or ROAs that are not present, but it can continue to use what it has cached.

If a CA revokes a CA certificate or a ROA (via deleting the corresponding EE certificate), and the adversary tries to reinstate that CA certificate or ROA, the adversary would have to rollback the CRL and the manifest to undo this action by the CA. As above, this would make the CRL and manifest stale, and this is detectable by RPs. An RP cannot know which CA certificates or ROAs were deleted, and so it would use the cached instances of the affected objects. Here too one hopes that the CA will be notified of the problem and will attempt to remedy the error.

In the attack scenarios above, when a CRL or manifest is described as stale, this means that the next issue date for the CRL or manifest has passed. Until the next issue date, an RP will not be detect the attack. Thus it behooves CAs to select CRL/manifest lifetimes (the two are linked) that represent an acceptable tradeoff between risk and operational burdens.

Attacks effected by adversaries that are legitimate managers of publication points can have much greater effects, and are discussed below under attacks on or by CAs.

4.5. Attacks on an RPKI CA

Every entity to which INRs have been allocated/assigned is a CA in the RPKI. Each CA is nominally responsible for managing the repository publication point for the set of signed products that it

generates. (An INR holder may choose to outsource the operation of the RPKI CA function, and the associated publication point. In such cases, the organization operating on behalf of the INR holder becomes the CA, from an operational and security perspective. The following discussion does not distinguish outsourced CA operations.)

Note that attacks attributable to a CA may be the result of malice by the CA (i.e., the CA is the adversary) or they may result from a compromise of the CA.

All of adversaries listed in Section 2 are presumed to be capable of launching attacks against the computers used to perform CA functions. Some adversaries might effect an attack on a CA by violating personnel or physical security controls as well. The distinction between CA as adversary vs. CA as an attack victim is important. Only in the latter case should one expect the CA to remedy problems caused by a attack once the attack has been detected. Note that most of the attacks described below do not require disclosure of a CA's private key to an adversary. If the adversary can gain control of the computer used to issue certificates, it can effect these attacks, even though the private key for the CA remains "secure" (i.e., not disclosed to unauthorized parties). However, if the CA is not the adversary, and if the CA's private key is not compromised, then recovery from these attacks is much easier. This motivates use of hardware security modules to protect CA keys, at least for higher tiers in the RPKI.

An attack by a CA can result in revocation or replacement of any of the certificates that the CA issued. Revocation of a certificate should cause RPs to delete the (formerly) valid certificate (and associated signed object, in the case of a revoked EE certificate) that they have cached. This would cause repository objects (e.g., CA certificates and ROAs) that are verified under that certificate to be considered invalid, transitively. As a result, RPs would not consider as valid any ROAs or signed updates based on these certificates, which would make routes dependent on them to be less preferred. Because a CA that revokes a certificate is authorized to do so, this sort of attack cannot be detected, intrinsically, by most RPs. However, the entities affected by the revocation or replacement of CA certificates can be expected to detect the attack and contact the CA to effect remediation. If the CA was not the adversary, it should be able to issue new certificates and restore the publication point.

An adversary that controls the CA for a publication point can publish signed products that create more subtle types of DoS attacks against RPs. For example, such an attacker could create subordinate CA certificates with Subject Information Access (SIA) pointers that lead RPs on a "wild goose chase" looking for additional publication points

and signed products. An attacker could publish certificates with very brief validity intervals, or CRLs and manifests that become "stale" very quickly. This sort of attack would cause RPs to have to access repositories more frequently, and that might interfere with legitimate accesses by other RPs.

An attacker with this capability could create very large numbers of ROAs to be processed (with prefixes that are consistent with the allocation for the CA), and correspondingly large manifests. An attacker could create very deep subtrees with many ROAs per publication point, etc. All of these types of DoS attacks against RPs are feasible within the syntactic and semantic constraints established for RPKI certificates, CRLs, and signed objects.

An attack that results in revocation and replacement (e.g., key rollover or certificate renewal) of a CA certificate would cause RPs to replace the old, valid certificate with the new one. This new certificate might contain a public key that does not correspond to the private key held by the certificate subject. That would cause objects signed by that subject to be rejected as invalid, and prevent the affected subject from being able to sign new objects. As above, RPs would not consider as valid any ROAs issued under the affected CA certificate, and updates based on router certificates issued by the affected CA would be rejected. This would make routes dependent on these signed products to be less preferred. However, the constraints imposed by the use of RFC 3779 [RFC3779] extensions do prevent a compromised CA from issuing (valid) certificates with INRs outside the scope of the CA, thus limiting the impact of the attack.

An adversary that controls a CA could issue CA certificates with overlapping INRs to different entities, when no transfer of INRs is intended. This could cause confusion for RPs as conflicting ROAs could be issued by the distinct CAs.

An adversary could replace a CA certificate, use the corresponding private key to issue new signed products, and then publish them at a publication point controlled by the attacker. This would effectively transfer the affected INRs to the adversary, or to a third party of his choosing. The result would be to cause RPs to view the entity that controls the private key in question as the legitimate INR holder. Again the constraints imposed by the use of RFC 3779 extensions do prevent a compromised CA from issuing (valid) certificates with INRs outside the scope of the CA, thus limiting the impact of the attack.

Finally, an entity that manages a repository publication point can inadvertently act as an attacker (as first noted by Pogo). For example, a CA might fail to replace its own certificate in a timely

fashion (well before it expires). If might fail to issue its CRL and manifest prior to expiration, creating stale instances of these products that cause concern for RPs. A CA with many subordinate CAs (e.g., an RIR or NIR) might fail to distribute the expiration times for the CA certificates that it issues. An ISP with many ROAs might do the same for the EE certificates associated with the ROAs it generates. A CA could rollover its key, but fail to reissue subordinate CA certificates under its new key. Poor planning with regard to rekey intervals for managed CAs could impose undue burdens for RPs, despite a lack of malicious intent. All of these example of mismanagement could adversely affect RPs, despite the absence of malicious intent.

5. Residual Vulnerabilities

The RPKI, upon which BGPSEC relies, has several residual vulnerabilities that were been discussed in the preceding text (Sections 4.4 and 4.5). These vulnerabilities are of two principle forms:

- the RPKI repository system may be attacked in ways that make its contents unavailable, or not current. It is anticipated that RPs will cope with this vulnerability through local caching of repository data, and through local settings that tolerate expired or stale repository data.

- any CA in the RPKI may misbehave within the bounds of the resources allocated to it, e.g., it may issue certificates with duplicate resource allocations or revoke certificates inappropriately. This vulnerability is intrinsic in any PKI. It is anticipated that RPs will deal with this through

BGPSEC has a separate set of residual vulnerabilities:

- BGPSEC is not able to prevent what is usually referred to as route leaks, because BGP itself does not distinguish between transit and non-transit ASes- BGPSEC signatures do not protect all attributes associated with an AS_path. Some of these attributes are employed as inputs to routing decisions. Thus attacks that modify (or strip) these other attributes are not detected by BGPSEC.

6. Security Considerations

A threat model is, by definition, a security-centric document. Unlike a protocol description, a threat model does not create security

problems nor purport to address security problems. This model postulates a set of threats (i.e., motivated, capable adversaries) and examines classes of attacks that these threats are capable of effecting, based on the motivations ascribed to the threats. It describes the impact of these types of attacks on BGPSEC, including on the RPKI on which BGPSEC relies.

7. IANA Considerations

[Note to IANA, to be removed prior to publication: there are no IANA considerations stated in this version of the document.]

8. Acknowledgements

The author wishes to thank . . .

9. References

9.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

9.2. Informative References

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

[RFC4301]
Kent, S. and Seo, K., "Security Architecture for the Internet Protocol", RFC 4301, December, 2005.

[RFC3779]
Lynn, C., Kent, S., Seo, K., X.509 Extensions for IP Addresses and AS Identifiers, RFC 3779, June 2004.

[Kent2000]
Kent, S., Lynn, C., and Seo, K., "Design and Analysis of the Secure Border Gateway Protocol (S-BGP)", IEEE DISCEX Conference, January, 2000.

[RFC5925]
Touch, J., et al., "The TCP Authentication Option", RFC 5925, June 2010.

[I-D.ietf-sidr-arch]
Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", draft-ietf-sidr-arch-11.txt (work in progress), September 2010.

[I-D.sidr.signed-object]
Lepinski, M, Chi, A., and Kent, S., "Signed Object Template for the Resource Public Key Infrastructure", draft-ietf-sidr-signed-object-01.txt, (work in progress), December 2010.

[I-D.sidr-res-certs]
Huston, G., Michaelson, G., and Loomans, R. "A Profile for X.509 PKIX Resource Certificates", draft-ietf-sidr-res-certs-21.txt (work in progress), December 2010.

[I-D.roa-format]
Lepinski, M., Kent, S., and Kong, D., "A Profile for Route Origin Authorizations (ROAs)", draft-ietf-sidr-roa-format-09.txt, (work in progress), November 2010.

Author's Address

Stephen Kent BBN Technologies 10 Moulton St. Cambridge, MA 02138 USA

Email: kent@bbn.com

Network Working Group
Internet Draft
Intended status: Informational
Expires: September 7, 2011

M. Lepinski
BBN Technologies
S. Turner
IECA

March 7, 2011

An Overview of BGPSEC
draft-lepinski-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/1id-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 October 7, 2011.

Table of Contents

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.lepinski-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 accidently (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 :
    - Signature from AS 1 protecting

192.0.2/24, AS 1 and AS 2

- Signature from AS 2 protecting

  Everything AS 1's signature protected, and AS 3

- 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 BPSEC; 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 Michaelson, G., "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.lepinski-bgpsec-protocol] Lepinski, M., Ed., "BPSEC Protocol Specification", draft-lepinski-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

Authors' Addresses

Matt Lepinski
BBN Technologies
10 Moulton Street
Cambridge MA 02138

Email: mlepinski@bbn.com

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

Email: turners@ieca.com

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 8, 2011

M. Lepinski, Ed.
BBN

March 7, 2011

BGPSEC Protocol Specification
draft-lepinski-bgpsec-protocol-00.txt

Abstract

This document describes BGPSEC, a mechanism for providing path security for BGP route advertisements. BGPSEC is implemented via a new optional non-transitive BGP path attribute.

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].

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 September 8, 2011.

Copyright Notice

Table of Contents

1. Introduction

This document describes BGPSEC, a mechanism for providing path security for BGP 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 3.

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

Capability Value:

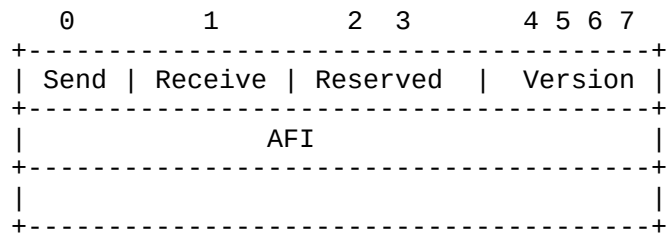

The high order bit (bit 0) of the first octet is set to 1 to indicate that the sender is able to send BGPSEC update messages, and is set to zero otherwise. The next highest order bit (bit 1) of this octet is set to 1 to indicate that the sender is able to receive BGPSEC update messages, and is set to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

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 peer, then the behavior of those peers is outside the scope of this document.)

The second two 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. BGPSEC for use with other address families may be specified in future documents. 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. Also note that a BGPSEC speaker SHOULD NOT advertise the capability

of BGPSEC support for IPv6 unless it has also advertised support for IPv6 [2].

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

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

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

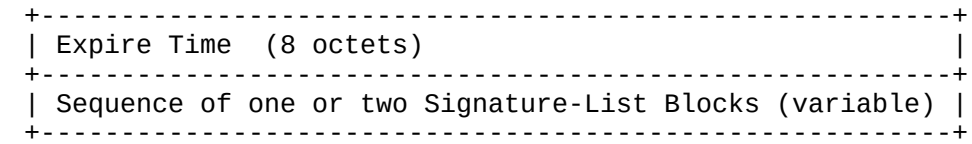

Expire Time contains a binary representation of a time as an unsigned integer number of (non-leap) seconds that have elapsed since midnight

UTC January 1, 1970. The Expire Time indicates the latest point in time that the route advertised in the update message can possibly be considered valid (see Section 5 for details on validity of BGPSEC update messages).

The BGPSEC_Path_Signatures attribute will contain one or two Signature-List Blocks, each of which corresponds to a different algorithm suite. Each of the Signature-List Blocks will contain a signature segment for each AS in the AS Path attribute. In the most common case, the BGPSEC_Path_Signatures attribute will contain only a single Signature-List Block. However, in order to enable a transition from an old algorithm suite to a new algorithm suite, it will be necessary to include two Signature-List Blocks (one for the old algorithm suite and one for the new algorithm suite) during the transition period.

Signature-List Block

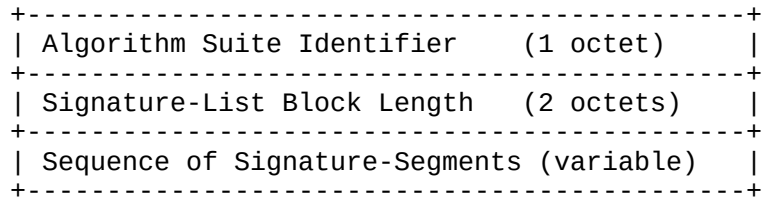

An algorithm suite consists of a digest algorithm and a signature algorithm. This version of BGPSEC only supports signature algorithms that produce a signatures of fixed length. This specification creates an IANA registry of one-octet BGPSEC algorithm suite identifiers. Additionally, this document registers a single algorithm suite which uses the digest algorithm SHA-256 and the signature algorithm RSA with 2048-bit keys [1]. The signatures produced by this algorithm suite have a length of 256 octets. Future registrations of algorithm suites for BGPSEC must specify the length of signatures produced by the algorithm suite.

BGPSEC Algorithm Suites

Algorithm Suite Identifier	Digest Algorithm	Signature Algorithm	Specification Pointer
TBA	SHA-256	RSA 2048	RFC 3447

The Signature-List Block Length is the total number of octets in all Signature-Segments (i.e., the total size of the variable-length

portion of the Signature-List block.)

A Signature-Segment has the following structure:

Signature Segments

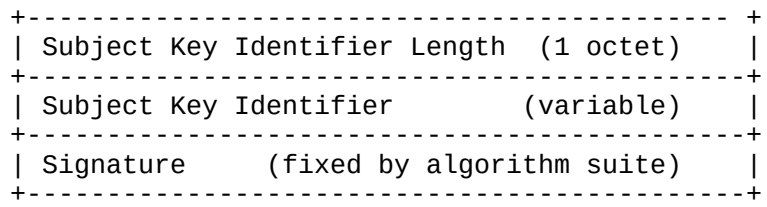

The Subject Key Identifier Length contains the size (in octets) of the value in the Subject Key Identifier field of the Signature-Segment. The Subject Key Identifier contains the value in the Subject Key Identifier extension of the RPKI end-entity certificate that is used to verify the signature (see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

The Signature contains a digital signature that protects the NLRI, the AS_Path and the BGPSEC_Path_Signatures attribute (see Sections 4 and 5 for details on generating and verifying this signature, respectively). The length of the Signature field is a function of the algorithm suite for a given Signature-List Block. The specification for each BGPSEC algorithm suite must provide the length of signatures constructed using the given algorithm suite.

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. 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 Signature-List Block for 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. 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, possibly multiple occurrences of, a single AS number), the BGPSEC speaker creates one Signature-List Block for each algorithm suite that will be used. Typically, a BGPSEC speaker will use only a single algorithm suite. 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 Signature-List Blocks for both the 'current' and the 'new' algorithm suites (see Section 6.1).

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 there exists a valid ROA authorizing the BGPSEC speaker's AS to originate routes to this prefix.

The Expire Time field is set to specify a time at which the route advertisement specified in the update message will cease to be valid. Once the Expire Time has been reached, all BGPSEC speakers who have received the advertisement will treat it as invalid. The purpose of this field is to protect the BGPSEC speaker against attacks in which the BGPSEC speaker wishes to withdraw the route, but intermediate (malicious) BGP speakers fail to propagate the withdrawal to their peers.

It is therefore necessary for the originating BGPSEC speaker to issue a new BGPSEC update prior to reaching the Expire Time. It is RECOMMENDED that a BGPSEC speaker originate a new route advertisement for a given NLRI at intervals equal to roughly one-third the validity period of the route advertisement. (Note that it is necessary to add some small amount of random jitter to the interval to avoid synchronization effects.) For instance, if a BGPSEC speaker is originating route advertisements that are valid for one day (i.e., the Expire Time is 24 hours after the generation of the update message), then it is recommended that the BGPSEC speaker re-issue new a new BGPSEC update message for advertising the given prefix roughly once every 8 hours (plus or minus a small random value).

(Editor's Note: The parameter recommendations in the previous paragraph are preliminary and may need to be updated based on further implementation and deployment experience.)

There is a natural trade-off in setting the Expire Time. Setting a later Expire Time increases the amount of time by which a malicious intermediate can delay a future route withdrawal. Similarly, setting a later Expire Time also increases the window of opportunity for malicious replay attacks in which a previous BGPSEC announcement is replayed while suppressing a more recent withdrawal for the same prefix. However, setting a sooner Expire Timed increases the frequency with which the BGPSEC speaker needs to send new announcements for the given prefix.

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

The Subject Key Identifier field (see Section 3) is populated with the identifier contained in the Subject Key Identifier extension of the RPKI end-entity certificate 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 Length field is populated with the length (in octets) of the Subject Key Identifier.

The Signature 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 Expire Time, Target AS Number, Origin AS Number, Algorithm Suite Identifier, 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.) The Origin AS number prepend to this sequence the Target AS (the AS to whom the BGPSEC speaker intends to send the update message) and the Origin AS Number refers to the AS of the BGPSEC speaker who is originating the route advertisement.

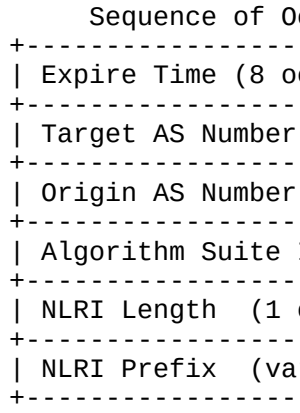

- o Apply to this octet sequence the digest algorithm (for the algorithm suite of this Signature-List) to obtain a digest value.

- 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.

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. (See Section 7 for

discussion of the security ramifications of removing 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.

To generate the BGPSEC_Path_Signatures attribute on the outgoing update message, the BGPSEC first copies the Expire Time directly from the received update message to the new update message (that it is constructing). Note that the BGPSEC speaker MUST NOT change the Expire Time as any change to Expire Time will cause the new BGPSEC update message to fail validation (see Section 5).

The BGPSEC speaker next removes from the BGPSEC_Path_Signatures attribute any Signature-List Blocks corresponding to algorithm suites that it does not support. The BGPSEC_Path_Signatures attribute for the new update message SHOULD contain a Signature-List Block for every algorithm suite that is both present in the received update message and which is supported by the BGPSEC speaker.

Note that the validation algorithm (see Section 5.1) deems a BGPSEC update message to be 'Good' if there is at least one supported algorithm suite (and corresponding Signature-List Block) that is deemed 'Good'. This means that a 'Good' BGPSEC update message may contain Signature-List Blocks which are deemed 'Not Good' (e.g., contain signatures that the BGPSEC is unable to verify). Nonetheless, such Signature-List Blocks MUST NOT be removed. (See Section 7 for a discussion of the security ramifications of this design choice.)

For each Signature-List Block corresponding to an algorithm suite that the BGPSEC speaker does support, the BGPSEC speaker then adds a new Signature-Segment to the Signature-List Block. 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 Subject Key Identifier field in the new segment is populated with the identifier contained in the Subject Key Identifier extension of the RPKI end-entity certificate 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 Length field is populated with the length (in octets) of the Subject Key Identifier.

The Signature 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:

- Construct a sequence of octets by concatenating the signature field of the most recent Signature-Segment (the one corresponding to AS from whom the BGPSEC speaker's AS received the announcement) with 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

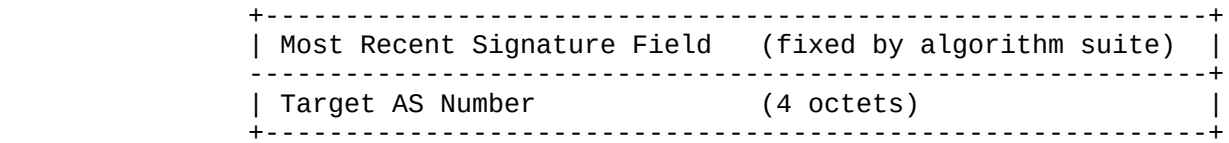

- Apply to this octet sequence the digest algorithm (for the algorithm suite of this Signature-List) to obtain a digest value.

- 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.

5. Validating a 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:

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

- 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.

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 such as setting a BGP community, or modifying a metric value such as Local_Pref or MED. 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.

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 performs the following checks:

- Check to ensure that the entire BGPSEC_Path_Signatures attribute is syntactically correct (conforms to the specification in this document).
- Check to ensure that the AS-Path attribute contains no AS-Set segments.
- Check that each Signature-List Block contains one Signature-Segment for each AS in the AS-Path attribute. (Note that the entirety of each Signature-List Block must be checked to ensure that it is well formed, even though the validation process may terminate before all signatures are cryptographically verified.)

If there are two Signature-List Blocks within the BGPSEC_Path_Signatures attribute and one of them is poorly formed (or contains the wrong number of Signature-Segments) , then the recipient should log that an error occurred, strip off that particular Signature-List Block and process the update message as though it arrived with a single Signature-List Block. If the BGPSEC_Path_Signatures attribute contains a syntax error which is not local to a single Signature-List Block, or if the AS-Path attribute contains an AS-Set segment, then the recipient should log that an error occurred, strip off the BGPSEC_Path_Signatures attribute and process the update message as though it arrived without a BGPSEC_Path_Signatures attribute.

Second, the BGPSEC speaker verifies that the update message has not yet expired. To do this, locate the Expire Time field in the BGPSEC_Path_Signatures attribute, and compare it with the current time. If the current time is later than the Expire Time, the BGPSEC update is ’Not Good’ and the validation algorithm terminates.

Third, 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.

Finally, the BGPSEC speaker examines the Signature-List Blocks in the BGPSEC_Path_Signatures attribute. Any Signature-List Block corresponding to an algorithm suite that the BGPSEC speaker does not support MUST be discarded. If all Signature-List Blocks are

discarded in this manner then the BGPSEC speaker MUST treat the update message as though it arrived without a BGPSEC_Path_Signatures attribute.

For each remaining Signature-List Block (corresponding to an algorithm suite supported by the BGPSEC speaker), the BGPSEC speaker iterates through the Signature-Segments in the Signature-List block, starting with the most recently added segment (and concluding with the least recently added segment). Note that there is a one-to-one correspondence between Signature-Segments and AS numbers in the AS-Path attribute, and the following steps make use of this correspondence.

- (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 SKI field of the Signature-Segment. If no such SKI value is found in the valid RPKI data then mark the entire Signature-List Block as 'Not Good' and proceed to the next Signature-List Block. Similarly, if the SKI exists but the AS Number associated with the SKI does NOT match the AS Number (in the AS-Path attribute) which corresponds to the current Signature-Segment, then mark the entire Signature-List Block as 'Not Good' and proceed to the next Signature-List Block.

- (Step II): Compute the digest function (for the given algorithm suite) 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

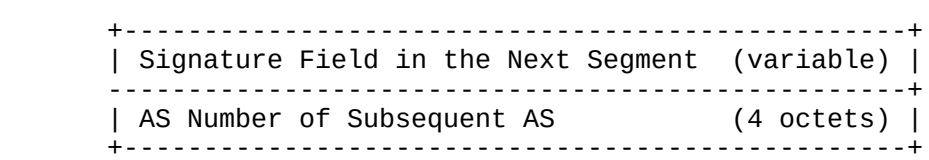

The 'Signature Field in the Next Segment' is the Signature field found in the Signature-Segment that is next to be processed (that is, the next most recently added Signature- Segment).

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 Subsequent AS' is the AS that corresponds to the Signature-Segment added immediately after the one being processed. (That is, find the AS number corresponding to the Signature-Segment currently being processed and the 'AS Number of Subsequent AS' is the next AS number that was added to the AS-Path attribute.)

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

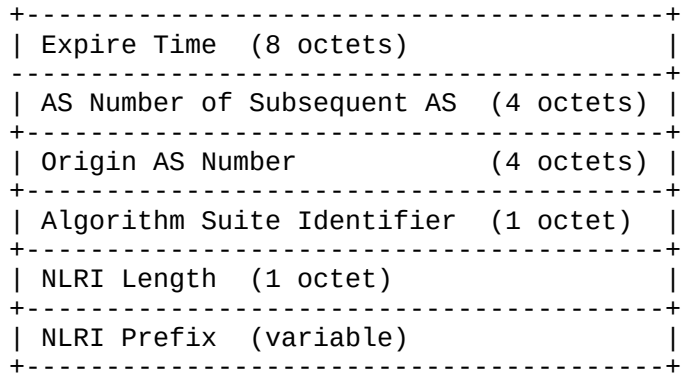

The NLRI Length, NLRI Prefix, Expire Time, and Algorithm Suite Identifier are all obtained in a straight forward manner from the NLRI of the update message or the BGPSEC_Path_Signatures attribute being validated.

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

The 'AS Number of Subsequent AS' is the AS Number added to the AS-Path immediately after the Origin AS Number. (That is, the second AS Number that was added to the AS Path.)

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, then mark the entire Signature-List Block as 'Not Good' and proceed to the next Signature-List Block. If the signature validation algorithm determines that the signature is valid, then continue processing Signature-Segments (within the current Signature-List Block).

If all Signature-Segments within a Signature-List Block pass validation (i.e., all segments are processed and the Signature-List Block has not yet been marked 'Not Good'), then the Signature-List Block is marked as 'Good'.

If at least one Signature-List Block is marked as 'Good', then the validation algorithm terminates and the BGPSEC update message is deemed to be 'Good'. (That is, if a BGPSEC update message contains two Signature-List Blocks then the update message is deemed 'Good' if the first Signature-List block is marked 'Good' OR the second Signature-List block is marked '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 single Signature-List Block (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:

- A new type of signature algorithm that produces signatures of variable length
- 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)
- 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:

- 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.

- 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.

- 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 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 both a 'Good' and a 'Not Good' Signature-List Block. 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 to each of the Signature-List Blocks using both the corresponding algorithm suite. 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' Signature-List Blocks 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' Signature-List Blocks.) 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.

8. Contributors

8.1. Authors

Rob Austein
Internet Systems Consortium
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

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
weiler@watson.org
Cobham

8.2. Acknowledgements

The authors would like to thank Sharon Goldberg, Ed Kern, Chris Morrow, Sandy Murphy, Mark Reynolds, Heather Schiller, Jason Schiller, John Scudder, and David Ward for their valuable input and review.

9. References

[1] Jonsson, J. and B. Kaliski, "PKCS #1", RFC 3447, February 2003.

[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 4760, 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", draft-ymbk-bgp-extended-messages, March 2011.

[6] Lepinski, M., Kent, S., and D. Kong, "A Profile for Route Origin Authorizations", draft-ietf-sidr-roa-format, February 2011.

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

[8] Kent, S., "Threat Model for BGP Path Security", draft-kent-bgpsec-threats, February 2011.

Author's Address

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

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

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 8, 2011

T. Manderson
ICANN
R L. Barnes
M. Lepinski
BBN

February 4, 2011

Providing first class geographical location statements for RPKI objects
draft-manderson-sidr-geo-00.txt

Abstract

This document describes the construction and use of the RPKI-GEO record. This record provides first class informational statements pertaining to the geographical attributes of the information described in RPKI objects.

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 August 8, 2011.

Copyright Notice

Table of Contents

1. Requirements Notation

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 [RFC2119].

2. Introduction

There is a constant an ongoing effort to investigate and analyse the global internet routing system from many different perspectives. One perspective is related to the geographical position of BGP [RFC4271] speakers and the terrestrial location of the route propagation. Recording of such information by passive BGP listeners in MRT format is described in the MRT BGP routing information export format with geo-location extensions [I-D.ietf-grow-geomrt]. There are of course many other efforts external to the IETF and won't be described here. Further awareness of these efforts is left to the reader.

This document describes the construction, use, and interpretation of the RPKI-GEO record. This record provides first class informational attestations pertaining to the geographical attributes relating to the information described in RPKI objects. The use of the geographical data is of an informational nature and provides a consistent and validatable approach to asserting the location properties of any item described by an RPKI object. To maintain consistency implementers and readers should condier the 9 rules in section 3 of [RFC5491].

It is not intended that the RPKI object described herein be used to directly influence routing or forwarding decisions. Its creation by any certificate maintainer is to be interpreted as informational and any replying party should only use the knowledge in the efforts of routing research or anomaly detection.

The geographic attestations made in this object are made by the certificate maintainer and their validity and accuracy is in the hands of the certificate maintainer. It is left to the relying party as how much trust is given to the geographic data provided by the certificate maintainer.

The RPKI-GEO object pertains only to the objects at the same RPKI repository publication point where it itself is published.

3. Suggested Reading

The assumption is made that the reader comprehends the RPKI, the RPKI Repository Structure, and the various RPKI objects described in the following: [I-D.ietf-sidr-arch], [I-D.ietf-sidr-res-certs], [I-D.ietf-sidr-signed-object], [I-D.ietf-sidr-roa-format], [I-D.ietf-sidr-rpki-manifests], [I-D.ietf-sidr-ghostbusters].

4. RPKI-GEO Structure

The structure of the GEO-RPKI object follows the description and the generic RPKI validation as described in Signed Object Template for the Resource Public Key Infrastructure [I-D.ietf-sidr-signed-object]

4.1. CMS Packaging

The eContentType of the RPKI-GEO object in the encapContentInfo (signed content) section of object is defined as rpkiGEO with the numerical value of TO BE ASSIGNED.

4.2. eContent

The content of a RPKI-GEO object identifies an RPKI object and the geographical coordinates associated with the item described by the RPKI object.

The ASN.1 for the RPKI-GEO object is as follows:

```
rPKIGEO ::= SEQUENCE {
        Version       [0] INTEGER DEFAULT 0,
        geoLocs SEQUENCE (SIZE(1..MAX)) OF geoOBJECTS
        }

geoObjects ::= SEQUENCE {
    objectFile          FileAndHash,
    geoAttribs          SEQUENCE (SIZE(1..MAX)) OF geoXML
    }

FileAndHash ::=     SEQUENCE {
    file            IA5String,
    hash            BIT STRING
    }

geoXML ::= SEQUENCE {
   type   INTEGER DEFAULT 0,
   xmlDoc PrintableString
}
```

4.3. RPKI-GEO data elements

4.3.1. Version

The version number of this version of the GEO-RPKI object MUST be 0.

4.3.2. geoLocs

This field is a sequence of geoObjects. Each geoObject contains a FileAndHash element and a sequence of geoXML. The geoLoc object MUST contain at least one geoXML object of type 0 for each FileAndHash element

4.3.3. FileAndHash

The single FileAndHash entry in each geoObject corresponds to each currently valid signed object that has been published by the authority (at this publication point). The description is as seen in [I-D.ietf-sidr-rpki-manifests]: Each FileAndHash is an ordered pair consisting of the name of the file in the repository publication point that contains the object in question, and a hash of the file's contents.

The publication point manifest and RPKI-GEO object's FileAndHash MUST NOT appear in a RPKI-GEO object.

4.3.4. geoXML

The geoXML contains the details of the geographical location information in an xml representation defined by the geoXML type value. The type specifies the XML schema used in the xmlDoc portion. There are 2 valid types.

```
Type 0: A GML syntax
Type 1: A Civic Address Syntax
```

geoXML schema types

4.3.4.1. Type 0

Type 0 is a constrained GML syntax [GML]. The constraints on the syntax are as follows.

Coordinate datum selection: The coordinates used in the GML will use the WGS84 datum [WGS84]. Any use of another datum specified in the GML in this object is considered illegal. This is for compatibility and uniformity.

The XML contained on the xmlDoc geoXML element for type = 0 MUST contain only one GML reference of either point or polygon representations.

```
<gml:Point srsName="urn:ogc:def:crs:EPSG::4326">
  <gml:pos>-43.5723 153.21760</gml:pos>
</gml:Point>
```

geoXML Type 0 xmlDoc GML example

4.3.4.2. Type 1

A Type 1 xmlDoc contains a Civic address representation of the location information and is defined in [RFC5139].

```
<civicAddress xml:lang="en-AU"
  xmlns="urn:ietf:params:xml:ns:pidf:geopriv10:civicAddr">
  <country>AU</country>
  <A1>NSW</A1>
  <A3>     Wollongong
  </A3><A4>North Wollongong
  </A4>
  <RD>Flinders</RD><STS>Street</STS>
  <RDBR>Campbell Street</RDBR>
  <LMK>
    Gilligan’s Island
  </LMK> <LOC>Corner</LOC>
  <NAM> Video Rental Store </NAM>
  <PC>2500</PC>
  <ROOM> Westerns and Classics </ROOM>
  <PLC>store</PLC>
  <POBOX>Private Box 15</POBOX>
</civicAddress>
```

geoXML Type 1 xmlDoc Civic address example

5. RPKI-GEO Validation

After the generic signed objects validation [I-D.ietf-sidr-signed-object] has been performed, the Version number field within the payload is checked. The payload data is checked against the profile defined in this document. All of these checks MUST pass for the RPKI-GEO payload to be considered valid and made available for use.

6. RPKI-GEO interpretation

A common sense interpretation of location data should prevail based on the type of the data that is represented in the RPKI object. For example a RPKI-GEO object that provides location information for a ROA would attest to the geographical location where the route is originated from. That may be the originating BGP speaker(s) as described in [I-D.ietf-grow-geomrt]. Similarly the location information associated with a Ghostbusters record [I-D.ietf-sidr-ghostbusters] would describe the geographical location of the entity described in the Ghostbusters VCARD.

7. IANA Considerations

This document requests IANA to add the .geo extention to the RPKI file extension namespace.

8. Security Considerations

The RPKI object described here is used in a descriptive nature and provide information that is useful in the analysis of routing systems. As such, the authors believes that it does not constitute an additional security risk. It is recommended that the issuers of the GEO-RPKI objects consider their own privacy concerns before supplying geographical coordinates in the RPKI.

9. References

9.1. Normative References

[I-D.ietf-grow-geomrt]
Manderson, T., "MRT BGP routing information export format with geo-location extensions", draft-ietf-grow-geomrt-01 (work in progress), December 2010.

[I-D.ietf-sidr-arch]
Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", draft-ietf-sidr-arch-11 (work in progress), September 2010.

[I-D.ietf-sidr-ghostbusters]
Bush, R., "The RPKI Ghostbusters Record", draft-ietf-sidr-ghostbusters-00 (work in progress), December 2010.

[I-D.ietf-sidr-res-certs]
Huston, G., Michaelson, G., and R. Loomans, "A Profile for X.509 PKIX Resource Certificates", draft-ietf-sidr-res-certs-21 (work in progress), December 2010.

[I-D.ietf-sidr-roa-format]
Lepinski, M., Kent, S., and D. Kong, "A Profile for Route Origin Authorizations (ROAs)", draft-ietf-sidr-roa-format-09 (work in progress), November 2010.

[I-D.ietf-sidr-rpki-manifests]
Austein, R., Huston, G., Kent, S., and M. Lepinski, "Manifests for the Resource Public Key Infrastructure", draft-ietf-sidr-rpki-manifests-09 (work in progress), November 2010.

[I-D.ietf-sidr-signed-object]
Lepinski, M., Chi, A., and S. Kent, "Signed Object Template for the Resource Public Key Infrastructure", draft-ietf-sidr-signed-object-02 (work in progress), December 2010.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.

[RFC5139] Thomson, M. and J. Winterbottom, "Revised Civic Location Format for Presence Information Data Format Location Object (PIDF-LO)", RFC 5139, February 2008.

[RFC5491] Winterbottom, J., Thomson, M., and H. Tschofenig, "GEOPRIV Presence Information Data Format Location Object (PIDF-LO) Usage Clarification, Considerations, and Recommendations", RFC 5491, March 2009.

9.2. Informative References

[GML] Open Geospatial Consortium, ODC., "OpenGIS Geography Markup Language (GML) Encoding Standard", December 2010, <http://portal.opengeospatial.org/files/?artifact_id=20509>.

[WGS84] Geodesy and Geophysics Department, DoD., "World Geodetic System 1984", January 2000, <http://earth-info.nga.mil/GandG/publications/tr8350.2/wgs84fin.pdf>.

Authors' Addresses

Terry Manderson
ICANN

Email: terry.manderson@icann.org

Richard L. Barnes
BBN

Email: rbarnes@bbn.com

Matt Lepinski
BBN

Email: mlepinski@bbn.com

Inter-Domain Routing
Internet-Draft
Intended status: Standards Track
Expires: September 8, 2011

A. Retana
R. Raszuk
Cisco Systems, Inc.

March 7, 2011

BGP Security State Diagnostic Message
draft-retana-bgp-security-state-diagnostic-00

Abstract

This document describes an extension to the BGP Diagnostic Message to communicate the security state of a route. An application of this extension is to propagate information about non-secure advertisements back to the eBGP peer from where the information was received.

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 September 8, 2011.

Copyright Notice

Table of Contents

1. Introduction

BGP Prefix Origin Validation [I-D.ietf-sidr-pfx-validate] defines the interaction between BGP and a database able to map prefixes to their authorized ASes. One of the potential actions resulting from an "invalid" route is to reject it.

This document describes an extension to the BGP Diagnostic Message [I-D.raszuk-bgp-diagnostic-message] and its use to communicate information about these "invalid" paths. The main motivation is to facilitate troubleshooting, monitoring, logging or even correction of the security mechanisms' operation, especially during initial deployment.

2. 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 [RFC2119].

3. The BGP Security State Diagnostic Message

The BGP Security State Diagnostic Message is a TLV to be carried in the BGP Diagnostic Message and is used to communicate the local security state of a path. It is defined as follows.

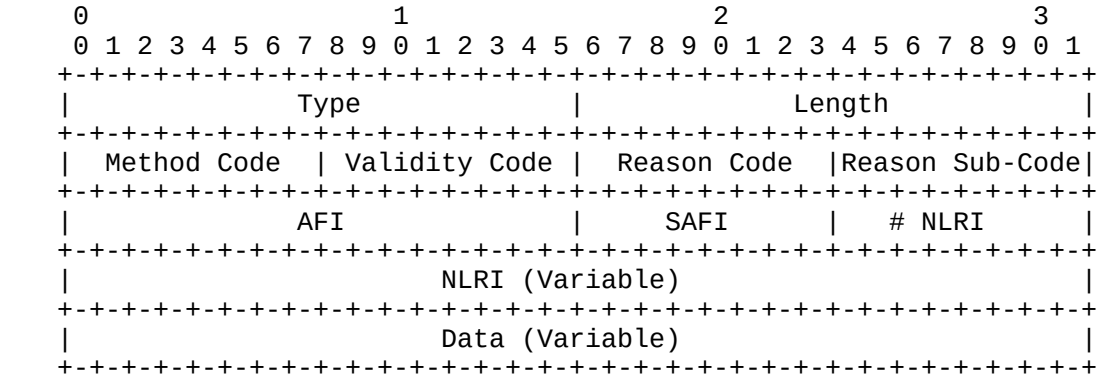

BGP Security State Diagnostic Message

Type:
Two octet field with a value TBD.

Length:
Two octet field indicating the TLV length in octets.

Method Code:
One octet field. Indicates which security mechanism was used to determine the validity of the path.

Value	Meaning
0	Reserved
1	BGP Prefix Origin Validation [I-D.ietf-sidr-pfx-validate]

Method Codes

Validity Code:
One octet field. Indicates whether the path is considered secure or not by the local AS. The values are to be interpreted relative to the Method defined above.

The following values are defined for Method Code 1:

Value	Meaning
0	Reserved
1	Not Found
2	Invalid Path

Validity Codes

Reason Code:
One octet field. Indicates the reason the security mechanism listed in the Method Code considered the path as indicated in the Validity Code. The values are specific to the Method Code used.

The following Reason Codes are defined for Method Code 1, Validity Code 2:

Value	Meaning
0	Reserved
1	Invalid Origin
2	Certificate doesn’t exist

Reason Codes

Reason Sub-Code:
One octet field. Indicates any additional information related to the Reason Code indicated for the specific Method used. At this time no specific values are defined.

AFI (Address Family Identifier):
Two octet field, encoded the same way as in RFC 4760 [RFC4760].

SAFI (Subsequent Address Family Identifier):
Two octet field, encoded the same way as in RFC 4760 [RFC4760].

NLRI (Number of Network Layer Reachability Information entries):
One octet field indicating the number of NLRI entries to follow.

NLRI:
Variable length field encoded as one or more 2-tuples of the form <length, prefix>, as described in RFC 4760 [RFC4760].

Data:
Variable length field. Indicates any additional information related to the Reason Code indicated for the specific Method used. This is an OPTIONAL field with variable length.

4. Operation

The mechanism described is intended to be primarily applied at autonomous system border routers.

When a BGP speaker receives what considers to be an invalid advertisement it MAY send a BGP Security State Diagnostic Message to the eBGP peer from where it received it. It is RECOMMENDED that a BGP speaker limit the number of messages sent to a specific peer over a given period of time and that the messages be built in such a way as to include as many NLRI as possible.

A BGP speaker SHOULD also send the BGP Security State Diagnostic Message in response to the "Prefix specific BGP query" TLV (type 17) or the "Diagnostic Message Query" TLV (type 3). The BGP Security State Diagnostic Message SHOULD NOT be sent periodically to a peer; to achieve this behavior the "Max frequency permitted" TLV (type 2) should be used to announce a value of 0.

The information contained in the BGP Security State Diagnostic Message can then be used to diagnose and correct any potential local security policy violations. Specific actions taken are outside the scope of this document, but could include withdrawing the original UPDATE or simply logging the information.

5. IANA Considerations

IANA is asked to create and maintain registries for the fields described in Section 3, and to assign the corresponding TLV type.

6. Security Considerations

The mechanism described in this document doesn't add any new security concerns.

7. Acknowledgements

The mechanism described in this document was influenced by discussions with Dacheng Zhang and Mingui Zhang.

The authors would like to thank the following people for their comments and suggestions: Bertrand Duvivier, Keyur Patel, Roque Gagliano and Russ White.

8. References

8.1. Normative References

[I-D.raszuk-bgp-diagnostic-message]
Raszuk, R., Chen, E., and B. Decraene, "BGP Diagnostic Message", draft-raszuk-bgp-diagnostic-message-00 (work in progress), October 2010.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, January 2007.

8.2. Informative References

[I-D.ietf-sidr-pfx-validate]
Mohapatra, P., Scudder, J., Ward, D., Bush, R., and R. Austein, "BGP Prefix Origin Validation", draft-ietf-sidr-pfx-validate-01 (work in progress), February 2011.

Authors' Addresses

Alvaro Retana
Cisco Systems, Inc.
7025 Kit Creek Rd.
Research Triangle Park, NC 27709
USA

Email: aretana@cisco.com

Robert Razsuk
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134
USA

Email: raszuk@cisco.com

Network Working Group
Internet-Draft
Intended status: BCP
Expires: September 16, 2011

R. Bush
Internet Initiative Japan
March 15, 2011

BGPsec Operational Considerations
draft-ymbk-bgpsec-ops-01

Abstract

Deployment of the BGPsec architecture and protocols has many operational considerations. This document attempts to collect and present them. It is expected to evolve as BGPsec is formalized and initially deployed.

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 [RFC2119].

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79. This document may not be modified, and derivative works of it may not be created, and it may not be published except as an Internet-Draft.

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 September 16, 2011.

Copyright Notice

Table of Contents

1. Introduction

BGPsec is a new protocol with many operational considerations. It is expected to be deployed incrementally over a number of years. As core BGPsec-capable routers may require large memory and crypto assist, it is thought that origin validation based on the RPKI will occur over the next two to five years and that BGPsec will start to deploy late in that window.

BGPsec relies on widespread propagation of the Resource Public Key Infrastructure (RPKI) [I-D.ietf-sidr-arch]. How the RPKI is distributed and maintained globally and within an operator's infrastructure may be different for BGPsec than for origin validation.

BGPsec need be spoken only by a AS's eBGP speaking, AKA border, routers, and is designed so that it can be used to protect announcements which are originated by small edge routers, and this has special operational considerations.

Different prefixes have different timing and replay protection considerations.

2. Suggested Reading

It is assumed that the reader understands BGP, [RFC4271], BGPsec, [I-D.lepinski-bgpsec-overview], the RPKI, see [I-D.ietf-sidr-arch], the RPKI Repository Structure, see [I-D.ietf-sidr-repos-struct], and ROAs, see [I-D.ietf-sidr-roa-format].

3. RPKI Distribution and Maintenance

The RPKI is a distributed database containing certificates, CRLs, manifests, ROAs, and Ghostbuster Records as described in [I-D.ietf-sidr-repos-struct]. Policies and considerations for RPKI object generation and maintenance are discussed elsewhere.

A local valid cache containing all RPKI data may be gathered from the global distributed database using the rsync protocol and a validation tool such as rcynic.

Validated caches may also be created and maintained from other validated caches. Network operators SHOULD take maximum advantage of this feature to minimize load on the global distributed RPKI database.

As RPKI-based origin validation relies on the availability of RPKI data, operators SHOULD locate caches close to routers that require these data and services. A router can peer with one or more nearby caches.

For redundancy, a router SHOULD peer with more than one cache at the same time. Peering with two or more, at least one local and others remote, is recommended.

If an operator trusts upstreams to carry their traffic, they SHOULD also trust the RPKI data those upstreams cache, and SHOULD peer with those caches. Note that this places an obligation on those upstreams to maintain fresh and reliable caches.

A transit provider or a network with peers SHOULD validate NLRI in announcements made by upstreams, downstreams, and peers. They still SHOULD trust the caches provided by their upstreams.

An environment where private address space is announced in eBGP the operator MAY have private RPKI objects which cover these private spaces. This will require a trust anchor created and owned by that environment, see [I-D.ietf-sidr-ltamgmt].

4. AS/Router Certificates

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

A large operator, concerned that a compromise of one router’s key would make many routers vulnerable, MAY accept a more complex certificate/key distribution burden to reduce this exposure.

On the other extreme, an edge site with one or two routers MAY use a single certificate/key.

Routers MAY be capable of generating their own keys and having their certificates signed and published in the RPKI by their NOC. This would mean that a router’s private key need never leave the router.

5. Within a Network

BGPsec is spoken by edge routers in a network, those which border other networks/ASs.

In a fully BGPsec enabled AS, Route Reflectors MUST have BGPsec

enabled if and only if there are eBGP speakers in their client cone.

A BGPsec capable router MAY use the data it receives to influence local policy within its network, see Section 8. In deployment this policy should fit into the AS's existing policy, preferences, etc. This allows a network to incrementally deploy BGPsec capable border routers.

eBGP speakers which face more critical peers or up/downstreams would be candidates for the earliest deployment. Both securing one's own announcements and validating received announcements should be considered in partial deployment.

An eBGP listener MUST NOT trust non-BGPsec markings such as communities received across a trust boundary.

6. Considerations for Edge Sites

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

BGPsec protocol capability negotiation provides for a speaker signing the data it sends but being unable to accept signed data. Thus a smallish edge router may hold only its own signing key(s) and sign it's announcement but not receive signed announcements and therefore not need to deal with the majority of the RPKI.

As the vast majority (84%) of ASs are stubs, and they announce the majority of prefixes, this allows for simpler and cheaper early incremental deployment. It may also mean that edge sites concerned with routing security will be attracted to upstreams which support BGPsec.

7. Beaconing Considerations

The BGPsec protocol attempts to reduce exposure to replay attacks by allowing the route originator to sign an announcement with a validity period and re-announce well within that period.

This re-announcement is termed 'beaconing'. All timing values are, of course, jittered.

It is only the originator of an NLRI which signs the announcement with a validity period.

To reduce vulnerability to a lost beacon announcement, a router SHOULD beacon at a rate somewhat greater than half the signature validity period it uses.

As beaconing places a load on the entire global routing system, careful thought MUST be given to any need to beacon frequently. This would be based on a conservative estimation of the vulnerability to a replay attack.

Beacon timing and signature validity periods SHOULD be as follows:

The Exemplary Citizen: Prefix originators who are not overly concerned about replay attacks might announce with a signature validity of multiple weeks and beacon one third of the validity period.

Normal Prefix: Most prefixes SHOULD announce with a signature validity of a week and beacon every three days.

Critical Prefix: Of course, we all think what we do is critical. But prefixes of top level DNS servers, and RPKI publication points are actually critical to large swaths of the Internet and are therefore tempting targets for replay attacks. It is suggested that the beaconing of these prefixes SHOULD be two to four hours, with a signature validity of six to twelve hours.

Note that this may incur route flap damping (RFD) with current default but deprecated RFD parameters, see [I-D.ymbk-rfd-usable].

8. Routing Policy

Unlike origin validation based on the RPKI, BGPsec marks a received announcement as Valid or Invalid, there is no NotFound state. How this is used in routing is up to the operator’s local policy. See [I-D.pmohapat-sidr-pfx-validate].

As BGPsec will be rolled out over years and does not allow for intermediate non-signing edge routers, coverage will be spotty for a long time. Hence a normal operator’s policy SHOULD NOT be overly strict, perhaps preferring valid announcements and giving very low preference, but still using, invalid announcements.

Local policy on the eBGP edge MAY convey the validation status of a BGP signed path through various pre-existing mechanisms, e.g. setting a BGP community, or modifying a metric value such as local-preference or MED. Some MAY choose to use the large Local-Pref hammer. Others MAY choose to let AS-Path rule and set their internal metric, which

comes after AS-Path in the BGP decision process.

A BGPsec speaker validates signed paths at the eBGP edge.

Because of possible RPKI version skew, an AS Path which does not validate at router R0 might validate at R1. Therefore, signed paths that can not be validated SHOULD have their signatures kept intact and should be signed when sent to external BGPsec speakers.

This implies that AS Paths with non-validated signatures MAY be propagated to iBGP peers. Therefore, unless local policy ensures otherwise, a signed path learned via iBGP MAY NOT have been validated. If needed, the validation state SHOULD be signaled by normal policy mechanisms such as communities or metrics.

On the other hand, local policy on the eBGP edge might preclude iBGP or eBGP announcement of signed AS Paths which are not validated.

If a BGPsec speaker receives an unsigned path, it SHOULD perform origin validation per [I-D.pmohapat-sidr-pfx-validate].

9. Notes

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

Operators which manage certificates SHOULD have RPKI Ghostbuster Records (see [I-D.ietf-sidr-ghostbusters]), signed indirectly by End Entity certificates, for those certificates on which others’ routing depends for certificate and/or ROA validation.

10. Security Considerations

BGPsec is all about security, routing security. The major security considerations for the protocol are described in [BGPsec].

11. IANA Considerations

This document has no IANA Considerations.

12. Acknowledgments

The author wishes to thank the entire BGPsec foundation team.

13. References

13.1. Normative References

[I-D.ietf-sidr-ghostbusters]
Bush, R., "The RPKI Ghostbusters Record", draft-ietf-sidr-ghostbusters-02 (work in progress), March 2011.

[I-D.ietf-sidr-roa-format]
Lepinski, M., Kent, S., and D. Kong, "A Profile for Route Origin Authorizations (ROAs)", draft-ietf-sidr-roa-format-10 (work in progress), February 2011.

[I-D.lepinski-bgpsec-overview]
Lepinski, M. and S. Turner, "An Overview of BGPSEC", draft-lepinski-bgpsec-overview-00 (work in progress), March 2011.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

13.2. Informative References

[I-D.ietf-sidr-arch]
Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", draft-ietf-sidr-arch-12 (work in progress), February 2011.

[I-D.ietf-sidr-ltamgmt]
Kent, S. and M. Reynolds, "Local Trust Anchor Management for the Resource Public Key Infrastructure", draft-ietf-sidr-ltamgmt-00 (work in progress), November 2010.

[I-D.ietf-sidr-repos-struct]
Huston, G., Loomans, R., and G. Michaelson, "A Profile for Resource Certificate Repository Structure", draft-ietf-sidr-repos-struct-07 (work in progress), February 2011.

[I-D.pmohapat-sidr-pfx-validate]

Mohapatra, P., Scudder, J., Ward, D., Bush, R., and R. Austein, "BGP Prefix Origin Validation", draft-pmohapat-sidr-pfx-validate-07 (work in progress), April 2010.

[I-D.ymbk-rfd-usable]
Pelsser, C., Bush, R., Patel, K., Mohapatra, P., and O. Maennel, "Making Route Flap Damping Usable", draft-ymbk-rfd-usable-00 (work in progress), March 2011.

[RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.

Author's Address

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

Phone: +1 206 780 0431 x1
Email: randy@psg.com

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 8, 2011

S. Bellovin
Columbia University
R. Bush
Internet Initiative Japan, Inc.
D. Ward
Juniper Networks
March 7, 2011

Security Requirements for BGP Path Validation
draft-ymbk-bgpsec-reqs-02

Abstract

This document describes requirements for a future BGP security protocol design to provide cryptographic assurance that the origin AS had the right to announce the prefix and to provide assurance of the AS Path of the announcement.

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 [RFC2119].

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79. This document may not be modified, and derivative works of it may not be created, and it may not be published except as an Internet-Draft.

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 September 8, 2011.

Copyright Notice

Table of Contents

1. Introduction

RPKI-based Origin Validation ([I-D.ietf-sidr-pfx-validate]) provides a measure of resilience to accidental mis-origination of prefixes. But it provides neither cryptographic assurance (announcements are not signed), nor assurance of the AS Path of the announcement.

This document describes requirements to be placed on a BGP security protocol, herein termed BGPsec, intended to rectify these gaps.

The threat model assumed here is documented in [RFC4593] and [I-D.kent-bgpsec-threats].

2. Recommended Reading

This document assumes knowledge of the RPKI see [I-D.ietf-sidr-arch] and the RPKI Repository Structure, see [I-D.ietf-sidr-repos-struct].

This document assumes ongoing incremental deployment of ROAs, see [I-D.ietf-sidr-roa-format], the RPKI to Router Protocol, see [I-D.ietf-sidr-rpki-rtr], and RPKI-based Prefix Validation, see [I-D.ietf-sidr-pfx-validate].

And, of course, a knowledge of BGP [RFC4271] is required.

3. General Requirements

The following are general requirements for a BGPsec protocol:

3.1 A BGPsec design must allow the receiver of a BGP announcement to determine, to a strong level of certainty, that the received PATH attribute accurately represents the sequence of eBGP exchanges that propagated the prefix from the origin AS to the receiver.

3.2 A BGPsec design MUST be amenable to incremental deployment. Any incompatible protocol capabilities MUST be negotiated.

3.3 A BGPsec design MUST provide analysis of the operational considerations for deployment and particularly of incremental deployment, e.g, contiguous islands, non-contiguous islands, universal deployment, etc..

3.4 As cryptographic payloads and memory requirements on routers are likely to increase, a BGPsec design MAY require use of new hardware. I.e. compatibility with current hardware abilities is not a requirement that this document imposes on a solution. As BGPsec will likely not be rolled out for some years, this should not be a major problem.

3.5 A BGPsec design need not prevent attacks on data plane traffic. It need not provide assurance that the data plane even follows the control plane.

3.6 A BGPsec design MUST resist attacks by an enemy who has access to the link layer, per Section 3.1.1.2 of [RFC4593]. In particular, such a design must provide mechanisms for authentication of all data, including protecting against message insertion, deletion, modification, or replay. Mechanisms that suffice include TCP sessions authenticated with IPsec [RFC4301] or TLS [RFC5246].

3.7 A BGPsec design MAY make use of a security infrastructure (e.g., a PKI) to distribute authenticated data used as input to routing decisions. Such data include information about holdings of address space and ASNs, and assertions about binding of address space to ASNs.

3.8 If message signing increases message size, the 4096 byte limit on BGP PDU size MAY be removed.

3.9 It is entirely OPTIONAL to secure AS SETs and prefix aggregation. The long range solution to this is the deprecation of AS-SETs, see [I-D.wkumari-deprecate-as-sets].

3.10 If a BGPsec design uses signed prefixes, given the difficulty of splitting a signed message while preserving the signature, it need NOT handle multiple prefixes in a single UPDATE PDU.

3.11 A BGPsec design MUST enable each BGPsec speaker to configure use of the security mechanism on a per-peer basis.

3.12 A BGPsec design MUST provide backward compatibility in the message formatting, transmission, and processing of routing information carried through a mixed security environment. Message formatting in a fully secured environment MAY be handled in a non-backward compatible manner.

3.13 While the trust level of an NLRI should be determined by the BGPsec protocol, local routing preference and policy MUST then be applied to best path and other decisions. Such mechanisms MUST conform with [I-D.ietf-sidr-ltamgmt].

3.14 If a BGPsec design makes use of a security infrastructure, that infrastructure SHOULD enable each network operator to select the entities it will trust when authenticating data in the security infrastructure. See, for example, [I-D.ietf-sidr-ltamgmt].

3.15 A BGPsec design MUST NOT require operators to reveal more than is currently revealed in the operational inter-domain routing environment, other than the inclusion of necessary security credentials to allow others to ascertain for themselves the necessary degree of assurance regarding the validity of NLRI received via BGPsec. This includes peering, customer, and provider relationships, an ISP’s internal infrastructure, etc. It is understood that some data are revealed to the savvy seeker by BGP, traceroute, etc. today.

3.16 A BGPsec design SHOULD flag security exceptions which are significant enough to be logged. The specific data to be logged are an implementation matter.

3.17 Any routing information database MAY be re-authenticated periodically or in an event-driven manner, especially in response to events such as, for example, PKI updates.

3.18 Should a BGPsec design use hashes or signatures, it should provide mechanisms for algorithm agility.

3.19 A BGPsec design SHOULD NOT presume to know the intent of the originator of a NLRI, nor that of any AS on the AS Path.

3.20 A BGP listener SHOULD NOT trust non-BGPsec markings, such as communities, across trust boundaries.

4. BGP UPDATE Security Requirements

The following requirements MUST be met in the processing of BGP UPDATE messages:

4.1 A BGPsec design MUST enable each recipient of an UPDATE to formally validate that the origin AS in the message is authorized to originate a route to the prefix(es) in the message.

4.2 A BGPsec design MUST enable the recipient of an UPDATE to formally determine that the NLRI has traversed the AS path indicated in the UPDATE. Note that this is more stringent than showing that the path is merely not impossible.

4.3 Replay of BGP UPDATE messages need not be completely prevented, but a BGPsec design MUST provide a mechanism to control the window of exposure to replay attacks.

4.4 A BGPsec design SHOULD provide some level of assurance that the origin of a prefix is still 'alive', i.e. that a monkey in the middle has not withheld a WITHDRAW message or the effects thereof.

4.5 NLRI of the UPDATE message SHOULD be able to be authenticated in real-time as the message is processed.

4.6 Normal sanity checks of received announcements MUST be done, e.g. verification that the first element of the AS_PATH list corresponds to the locally configured AS of the peer from which the UPDATE was received.

4.7 The output of a router applying BGPsec to a received signed UPDATE MUST be either Valid or Unverified. There should be no shades of grey.

5. IANA Considerations

This document asks nothing of the IANA.

6. Security Considerations

The data plane may not follow the control plane.

Security for subscriber traffic is outside the scope of this document, and of BGP security in general. IETF standards for payload data security should be employed. While adoption of BGP security measures may ameliorate some classes of attacks on traffic, these measures are not a substitute for use of subscriber-based security.

7. Acknowledgments

The author wishes to thank the authors of [I-D.ietf-rpsec-bgpsecrec] from whom we liberally stole, Russ Housley, Geoff Huston, Steve Kent, Sandy Murphy, John Scudder, Sam Weiler, and a number of others.

8. References

8.1. Normative References

[I-D.kent-bgpsec-threats]
Kent, S., "Threat Model for BGP Path Security", draft-kent-bgpsec-threats-01 (work in progress), February 2011.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC4593] Barbir, A., Murphy, S., and Y. Yang, "Generic Threats to Routing Protocols", RFC 4593, October 2006.

8.2. Informative References

[I-D.ietf-rpsec-bgpsecrec]
Christian, B. and T. Tauber, "BGP Security Requirements", draft-ietf-rpsec-bgpsecrec-10 (work in progress), November 2008.

[I-D.ietf-sidr-arch]
Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", draft-ietf-sidr-arch-12 (work in progress), February 2011.

[I-D.ietf-sidr-ltamgmt]
Kent, S. and M. Reynolds, "Local Trust Anchor Management for the Resource Public Key Infrastructure", draft-ietf-sidr-ltamgmt-00 (work in progress), November 2010.

[I-D.ietf-sidr-pfx-validate]
Mohapatra, P., Scudder, J., Ward, D., Bush, R., and R. Austein, "BGP Prefix Origin Validation", draft-ietf-sidr-pfx-validate-01 (work in progress), February 2011.

[I-D.ietf-sidr-repos-struct]
Huston, G., Loomans, R., and G. Michaelson, "A Profile for Resource Certificate Repository Structure", draft-ietf-sidr-repos-struct-07 (work in progress), February 2011.

[I-D.ietf-sidr-roa-format]
Lepinski, M., Kent, S., and D. Kong, "A Profile for Route Origin Authorizations (ROAs)",

draft-ietf-sidr-roa-format-10 (work in progress), February 2011.

[I-D.ietf-sidr-rpki-rtr] Bush, R. and R. Austein, "The RPKI/Router Protocol", draft-ietf-sidr-rpki-rtr-10 (work in progress), March 2011.

[I-D.wkumari-deprecate-as-sets] Kumari, W., "Deprecation of BGP AS_SET, AS_CONFED_SET.", draft-wkumari-deprecate-as-sets-01 (work in progress), September 2010.

[RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.

[RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.

[RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.

Authors' Addresses

Steven M. Bellovin
Columbia University
1214 Amsterdam Avenue, MC 0401
New York, New York 10027
US

Phone: +1 212 939 7149
Email: bellovin@acm.org

Randy Bush
Internet Initiative Japan, Inc.
5147 Crystal Springs
Bainbridge Island, Washington 98110
US

Phone: +1 206 780 0431 x1
Email: randy@psg.com

Dave Ward
Juniper Networks
1194 N. Mathilda Ave.
Sunnyvale, California 94089-1206
US

Phone: +1-408-745-2000
Email: dward@juniper.net

INTERNET-DRAFT
Intended Status: Proposed Standard
Expires: September 15, 2011

Mingui Zhang
Huawei
Bin Liu
Tsinghua University
Dacheng Zhang
Huawei
Beichuan Zhang
The University of Arizona

March 14, 2011

Secure Extension of BGP by Decoupling Path Propagation and Adoption
draft-zhang-idr-decoupling-02

Abstract

This draft proposes a novel mitigation scheme to protect the inter-domain data delivery during false routing announcements. A new path attribute is defined to Decouple propagation of a path and adoption of a path for data forwarding in BGP (DBGP). DBGP does not use suspicious paths for data forwarding, but still propagates them in the routing system to facilitate attack detection. It can extensively protect data delivery from routing announcements of false sub-prefixes, false origins, false nodes and false links, and works well with ongoing attack detection and prevention systems.

Status of this Memo

This Internet-Draft is submitted to IETF 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/1id-abstracts.html

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

Acknowledgements

The helpful comments of the following are hereby acknowledged, in alphabetic order: Alvaro Retana, Xiaohu Xu.

Table of Contents

1 Introduction

False routing announcements cause serious security issues to the inter-domain routing system, which can lead to widespread service disruptions. A special case is prefix hijacking, in which a network announces an IP prefix that belongs to another network. Existing works such as Pretty Good BGP (PGBGP)[PGBGP] block suspicious routing updates to protect data delivery. However, such an approach has the side effect of blocking the view of detection systems at the control plane and data plane. As a result, an attack will not be detected, and operators will not be alerted to take actions to stop the attack. This draft proposes an extension of BGP, to solve this paradox by decoupling path propagation and adoption in BGP.

In current BGP, the path a router adopts for data forwarding is the same path being propagated to neighbors. That is why upon receiving a suspicious path, a router has to either accept it (no mitigation but good detection) or reject it (good mitigation but no detection). Our idea is for a router to use trusted paths for data forwarding, but still inform its neighbors about the suspicious path. The suspicious paths will be carried in an optional transitive attribute in BGP updates, while the routers still use trusted paths for data forwarding. This way the data traffic is protected while false routing announcements are being propagated to detection systems.

2 Terminology

DBGP: Decoupling path propagation and adoption in BGP

3 False Routing Announcements

3.1 Problem

False routing announcements can be caused by either inadvertent mis-configurations or malicious attacks. For the ease of exposition, we use "attacker" to refer to the network (or Autonomous System, AS) that makes the false routing announcements regardless of the intention. Based on which part of the routing path is false, such announcements can be classified into five types, each of which has different severity:

- Prefix origin: The attacker originates someone else’s prefix. Depending on where a network is in the Internet topology, some will choose paths leading to the true origin and some will choose paths leading to the false origin.

- Intermediate node: The attacker does not forge the prefix or its origin, but inserts itself into the path as an intermediate

AS. Similar to the case of false origin, some networks will choose the correct paths and some the false paths. But usually fewer networks will choose the false path since it is longer than that of false origin attack.

- Intermediate link: The attacker forges a new link to bypass some of the ASes to get a shorter path. The shortened path is expected to attract more networks' traffic.

- Sub-prefix: The attacker originates a sub-prefix of someone else's prefix. Due to the longest match in routing lookup, a false sub-prefix will win over the original prefix. Thus, all traffic destined to the sub-prefix range will be forwarded to the attacker.

- Super-prefix: Theoretically the attacker can also announce a false super-prefix, but that will not attract any traffic unless part of the prefix range is unused by the real owner, in which case it is equivalent to announcing a false origin of the unused prefix range.

3.2 Countermeasures

The current strategies proposed for the problem of false announcements fall in three categories: prevention, detection and mitigation.

3.2.1 Prevention

Prevention schemes (e.g., [SBGP], [SoBGP], and [SPV]) use cryptographic mechanisms to protect the routing updates and let routers reject any forged announcement. Unfortunately no prevention scheme has seen much deployment on the Internet due to the lack of incentives for those first movers. Since crypto-based schemes add significant computational load to routers and require upgrade on software or hardware, individual ISPs need to see immediate benefits to justify the deployment. On the current Internet, however, the first mover's routing announcements will be accepted by other networks without authentication, and adding authentication does not bring any immediate benefit.

3.2.2 Detection

The representative detection systems proposed in recent years include [Cyclops], [PHAS], [MyASN], [IAR], [iSPY], [NWatch], [OList], and [LWeight]. Such systems are designed to detect false routing by examining routing updates, probing data paths, cross-checking with registry databases, or a combination of these techniques. Once a

false routing case is detected, the owner of the prefix will be notified, and it is expected that the owner will take actions to resolve the problem, which, in today's Internet, usually involves contacting the offending network or its upstream provider to stop the false routing announcements. This process of detection, notification and resolution takes time, ranging from an hour to a day in some past incidents and varying from network to network [NWatch][RIPE]. In the meantime, the damage to data traffic has already been made and malicious attackers may have already achieved their objectives.

3.2.3. Traditional Mitigation

A mitigation schemes attempts to protect the data traffic while an attack is going on. The common approach is to somehow identify abnormal routing updates and block them. As proposed in [PGBGP] and [PGBGP++], a router can examine the content of incoming updates. If an AS path contains unexpected prefix origin or links, it will be suppressed from propagating for a period of time to wait for the network operator's validation. The length of the time is configurable by the network operators. Instead, an alternative path (via trusted prefix origin and links) will be employed for data delivery in the suppression period. After this period, if the path is not proved to be illegal, the router will adopt and announce this new path. PGBGP gives operators certain reaction time to resolve potential false announcements while protecting data delivery in the meantime. When a real attack is detected and resolved, the corresponding false announcement will be withdrawn from the routing system, whereas legitimate announcement will stay in the routing system and eventually be accepted.

But blocking false routing announcements can get in the way of detection systems. For instance, on September 22, 2008, a Russian ISP AS8997 hijacked a large number of prefixes as it leaked an entire table [ASN8997]. However, since the upstream ISP of AS8997 blocked the routing updates, detection systems such as MyASN and IAR did not pick up this incident. The attack mainly affected ISPs and users within Russia but largely went unnoticed by prefix owners.

Each existing solution has its drawbacks, and none is sufficiently effective by itself. In future, there may be several different solutions deployed on the Internet at the same time, complementary to each other, forming a multi-line defense to protect routing and data delivery before, during, and after attacks.

3.3 Paradox of Blocking Suspicious Updates and Attack Detection

It has been realized that a mitigation mechanism and a detection system that complement each other well can be integrated into an effective routing defense solution. For instance, the mitigation mechanism can help the detection system to confine the damage caused by an attack, as the affected data traffic may be vulnerable for hours before the attack can be actually detected by the detection mechanism and be eventually stopped. Also, the mitigation mechanism can also obtain benefit from the detection system because a mitigation mechanism normally cannot identity false routing information accurately with its limited knowledge, resource and time. The output from the detection system can be used to correct the many false positives generated by the mitigation mechanism and also inform

the prefix owner to resolve the attack.

However, there is a dilemma: the mitigation mechanism tries to render the attack ineffective while the detection system needs the attack to be effective in order to detect it.

4. DBGP's Mitigation: Decoupling Path Propagation and Adoption

The suspicious paths are serially blocked hop by hop for validation in traditional mitigation schemes, which gets in the way of detection systems. In order to solve this dilemma, this document proposes a solution which decouples path propagation and path adoption. The basic idea of this solution is to extend BGP's update message with a new optional transitive path attribute so that a router can inform its neighbor routers about the suspicious path and meanwhile the router uses another trusted path for data forwarding. In order to achieve this, a new BGP attribute, DAS_PATH, is defined in Section 5. Compared to the traditional mitigation schemes, the propagation of suspicious paths through DAS_PATHs in DBGP enables the parallel validation, which accelerates the adoption process of suspected legitimate paths (the false positive).

4.1 Quarantine Time

Based on operating experiences, a false routing announcement can be detected and corrected in a certain period (e.g., one day) after it is launched, and so an announcement will be trusted if it is not withdrawn in a pre-defined period. Therefore, in mitigation schemes, when a new path is identified as a suspicious one, it will be quarantined (or blocked) from being used for a period of time, which is called the quarantine time and noted as T_q. If a new path has stayed in a router's Adj-RIBs-In for more than T_q, it will be trusted by this router. If this path is the most preferred from the Adj-RIBs-In, the router will use this path for data delivery and announce this path to its neighbors.

A router MAY determine the quarantine time itself. Assume there are two routers, R1 and R2. R1 is the downstream of R2, and the quarantine time of R1 and R2 is T1 and T2 respectively. The suspicious path is PATH at R1. There are two possibilities.

- T1 is shorter than T2. When T1 expires, PATH becomes trusted by R1 and R1 begins to use it for data delivery. R1 will announce R1-PATH as a legitimate AS path to R2. However, at this time, the path R1-PATH is still being quarantined by R2.

- T1 is longer than T2. When T2 expires, R1-PATH becomes trusted by R2. However, R2 cannot use this path for data delivery as R1

has not announced R1-PATH as an AS path. It will be cached in the Adj-RIBs-In until the downstream router R1 has announced it as an AS path when T1 expires.

5 Protocol Descriptions

Take Figure 5.1 for example. A, B, C, and O are DBGP routers residing in different ASes, , X is the attacker and p is the prefix of interest. Before the attack, the preferred path for traffic is ABCO. Here, we use the notation "R1R2...Rn-p" to denote the AS path which is destined to prefix "p" via routers R1, R2, ..., Rn. When X makes a false announcement of X-p to B, B will regard this new path as suspicious because it would divert traffic to an AS(X) that previously was not on the data path (BCO). B will store the suspicious path in its Adj-RIBs-In (The routing tables of an AS router is comprised of three sub-tables: Adj-RIBs-In, Loc-RIB and Adj-RIBs-Out [BGP4]), but keep using the existing path in its Loc-RIB for data forwarding. At the same time, B re-announces its path (BCO) in an update message to A, and encapsulates the new, suspicious path (BX) as an optional transitive attribute which is defined in Section 5.1. After receiving the update message, router A learns this suspicious path, stores it, and propagates it further to its neighbors using the optional attribute in the same way. Therefore, the suspicious path is propagated to the Internet while not adopted for data forwarding. This approach enable the detection systems to intercept the suspicious path and notify the prefix owner to take actions. Once the attack is stopped, the false announcements will be withdrawn from the routing system, i.e., deleted or replaced in the Adj-RIBs-In of the involved routers. However, a router realizes that the quarantine time has been expired and the suspicious path is still in the Adj-RIBs-In, the router regards the path as a legitimate one. Hence the DBGP router will install the path in its Loc-RIB for data forwarding and re-announce the path using the regular ASPATH attribute in the update message. For example, if BX-p has been validated as legitimate, B will announce it to A as its AS_PATH. The rest of this section will discuss the design details in new BGP attribute definition, identifying suspicious paths, choosing alternative paths for data forwarding, propagating the paths, and releasing quarantined paths.

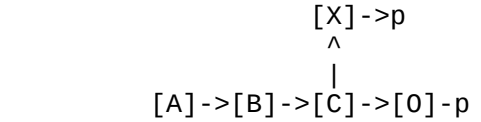

Figure 5.1: Attack Example 1

The rest of this section will discuss the design details in the

definition of the new DAS_PATH Attribute, identifying suspicious paths, choosing alternative paths for data forwarding, propagating the paths, and releasing quarantined paths.

5.1 DAS_PATH Attribute

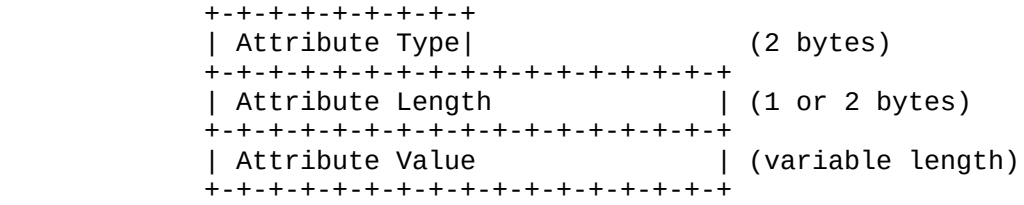

Figure 5.2: The DAS_PATH Attribute

DAS_PATH (Decoupled AS_PATH) is defined as a new optional transitive Path Attribute (Figure 5.2) to be included in BGP's UPDATE messages.

The first bit of the Attribute Type is set (1), therefore the attribute is optional. The second bit of Attribute Type is set (1), therefore the attribute is transitive and SHOULD be passed on to other BGP peers. The third and fourth bits are Partial bit and Extended Length bit respectively, which has already been defined in [BGP4]. The Attribute Type code is assigned by the IANA (Internet Assigned Numbers Authority). The definition of Attribute Length is the same as that in [BGP4].

The Attribute Value is composed of a sequence of DAS path segments. Each DAS path segment is encoded as a triple <path segment type, path segment length, path segment value>.

The path segment type is a 1-octet field with the following two allowable values:

Value	Segment	Type
1	DAS_SET	unordered set of ASs a route in the UPDATE message has traversed
2	DAS_SEQUENCE	ordered set of ASs a route in the UPDATE message has traversed

The path segment length field is 1-octet long field and contains the number of ASs in the path segment value field.

The path segment value field contains one or more AS numbers, each encoded as a 2-octets long field.

When a DAS_PATH is propagated across the network, the operations on

DAS_PATH follows the well-known AS_PATH attribute only that DAS_PATH is non-mandatory. If a router which does not deploy DBGP receives the update messages containing DAS_PATH attribute, i.e., does not understand this attribute, it will just pass it on to the next router. If its downstream router is a DBGP router, it will be able to pick up the information from this attribute and continue DBGP operations. Therefore DBGP can be incrementally deployed over the Internet.

5.2 Identification of Suspicious Paths

For a given prefix p, a path is trusted if it has been staying in the Adj-RIBs-In continuously for the required quarantine time, T_q. All the nodes, links, and origins that appear in trusted paths are trusted components, and the set of them is denoted by trusted(p). This set of trusted components is derived from current contents of all Adj-RIBs-In without using a database to store historical information like PGBGP does. Nodes, links, and origins that do not belong to trusted(p) are said to be suspicious components for this particular prefix p. A new path is suspicious if it contains any suspicious component for its prefix. However, not all suspicious paths need to be explicitly quarantined. DBGP quarantines paths that satisfy the following condition:

- A new path is quarantined if and only if it is suspicious, more preferred than other alternative paths, and contains an AS that is not in the current data forwarding path.

If the new path is not better than alternative paths, it will not be able to divert any traffic. One may suggest that the attacker can first announce a less preferred path so that DBGP routers will take it as a backup path without suspicion, and then make the primary path fail to trick the router to use the false backup path. But in this case, if the attacker has the control of the primary path, it can already get the traffic without doing this. If the attacker does not have control of the primary path, it will not know when the primary path may fail and which backup path the router will choose, thus the attack will not be effective.

If the new path does not introduce any new AS on the data path, it is not quarantined since it does not divert any traffic. In Figure 5.3, when X launches an attack by announcing X-p, this path is not quarantined by B since B already sends its traffic to X. B will accept this path and announces it to A. Assuming ABX-p is more preferred than ACO-p, A will quarantine ABX-p since this new path would divert A’s traffic to a new place, AS X, and X is a suspicious origin to A.

To reduce the potential false positives in quarantining paths, we introduce an optimization rule. If the current best path is replaced by a suspicious path (i.e., the same neighbor that sent the best path previously sends another path to replace it), then the current best path is cached for a short time period. This rule is used to accommodate some unstable prefixes, which may get announced and withdrawn or oscillate between two paths frequently. With this rule, quick re-announcement of a path will not be treated as suspicious. Previous measurement work has shown that (1) most prefixes are stable, and only a small number of prefixes are very unstable; (2) the most popular prefixes are stable; and (3) the most preferred paths are being used by routers for most of the time [BGPpop][Stable][BGPHA]. Therefore, DBGP's criterion for suspicious paths will not generate excessive false positives in reality. The majority of the rest prefixes is only suspicious infrequently.

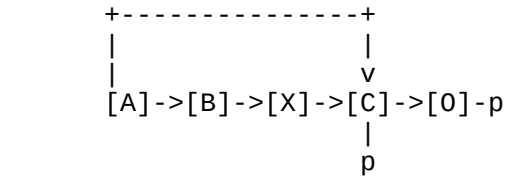

Figure 5.3: Attack Example 2

5.3 Propagating DAS_PATHs with Updates

DBGP uses the new BGP attribute, DAS_PATH, to carry a suspicious path in an update. In certain cases DBGP router may need to select a DAS_PATH from multiple ones to announce. For example, in Figure 5.4, assume C does not deploy DBGP and accepts the false announcement of X-p. When B receives BCX-p, B quarantines this new path and switches to its backup BDO-p. The update from B to A will have BDO as the AS_PATH, and BCX as the DAS_PATH attribute. However, since BDO is suspicious to A as well, A will quarantine BDO and switch to AEFO. Thus the update from A can contain either ABCX or ABDO. A router's neighbors may simultaneously send updates containing the DAS_PATH to it, which also makes the router have multiple DAS_PATHs to choose. There are two possible choices for DBGP implementations.

1. Dominate a router to select the best DAS_PATH from the multiple ones to announce, which is similar to the traditional AS_PATH selection process.

2. Encapsulate all the DAS_PATH attributes in tandem and sort them by their priorities. The priorities are determined according to the rules used in the AS_PATH selection process.

The first rule works well due to the following two facts. First, during attacks, the best DAS_PATH is most likely the bogus path aiming to attract data traffic. Second, during false positives, the best DAS_PATH will most likely become the best AS_PATH when the quarantine time ends.

The second rule allows an AS router to provide more information to its upstream, which is helpful for diagnose and correctness of attacks. Moreover, the announcement of multiple DAS_PATHs MAY help to reduce the convergence time. Take Figure 5.4 for example, A announces two DAS_PATHs, i.e., ABDO-p and ABCX-p, to its upstream node, says U. When U receives the additional DAS_APTH: ABDO-p, it will begin the validation process of this suspicious path. After A determines that BDO-p is legitimate and installs it to its Loc-RIB, the validation process MAY have been finished, therefore U can immediately start to use ABDO-p.

For the second rule, the number of DAS_PATHs in an update depends on the topology and policy of the network. Generally speaking, the number of DAS_PATHs in an update increases with the AS hop count of the AS_PATH. However, given AS paths are usually 4 to 5 hops and rarely goes to more than 10 hops, we do not expect the announcement of multiple DAS_PATHs will make DBGP message too large.

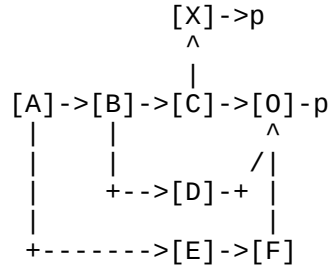

Figure 5.4: Attack Example 3

5.4 Choosing Alternative Paths

When a new path is the most preferred but suspicious, DBGP routers will use an alternative path for data delivery. The question is which alternative path to be chosen. First, if the existing path that is being used for data forwarding is still the best, then the router can stick to that path without any changes. Second, if the existing path in use will have been replaced by the suspicious path, then the router needs to pick an alternative. For example, in Figure 5.4, suppose C does not deploy DBGP and blindly accepts the false announcement X-p. B's existing path BCO-p will be replaced by a suspicious path BCX-p, therefore B needs to temporarily switch to a

backup path BDO-p from its Adj-RIBs-In. Third, if there is no alternative path or all alternative paths are labeled as suspicious, then the router err on data delivery by adopting a suspicious path to forward packets.

5.5 Releasing Quarantined Paths

If the quarantined paths are false announcements, it is likely that within T_q, the attack will be stopped and these paths being withdrawn from the routing system. In this case, there is no explicit release of the quarantined path. Just the upstream router will send an update with empty DAS_PATH attribute. If T_q has passed and the quarantined path is still in the Adj-RIBs-In, then it is more likely that this is a legitimate path. The router will treat the path as a regular path and make it go through the path selection process. If the path turns out to be the most preferred one, it will be used for data forwarding and trigger routing updates to neighbor routers.

6 Clarifications

6.1 Individual Historical Database

When DBGP is implemented in an AS router, the router does not have to purchase additional memory to store the trusted paths as that in [PGBGP]. By default, a DBGP router uses the simple rules defined in Section 5.2 to filter suspected components of an AS path based on the information stored in its Adj-RIBs-In. All a router need to do is to add a new column to its Adj-RIBs-In to record the elapse time after an AS path entered a Adj-RIB-In.

6.2 Difference from "add-path"

DBGP solution is different from the ongoing work of advertising multiple BGP paths in [add-path] where AS routers are also allowed to export multiple AS paths in one update. All advertised AS paths are available to upstream AS routers in [add-path]. Despite that DBGP allows multiple paths to be advertised in one update, except the AS path, all the other paths are actually unavailable. In other words, these paths only remain in Adj-RIBs-In of the AS router. They will not be put into either the Loc-RIB or Adj-RIBs-Out.

6.3 Detection Facilitation

Traditional mitigation mechanisms block the propagation of suspicious paths, which undermines the effectiveness of detection systems. DBGP is proposed to address this shortage. In DBGP, the data traffic is protected while the false routing announcements are spread out to be monitored by detection systems. If the first rule in Section 5.4 is

adopted, the capacity of propagating suspicious paths in DBGP is the same as that in BGP. If the second rule is adopted, this capacity is enhanced by DBGP instead.

6.4 Cooperation with Prevention

As a mitigation scheme, DBGP routers validate AS paths based on the limited information stored in local Adj-RIBs-In. This would cause some legitimate paths to be identified as suspicious and blocked from being used for data delivery (high false positive). If a down stream router would like their paths be adopted quickly rather than be suspected, it can include certificates in the update messages. For example, if AS routers adopt the solution in [pfx-va], the AS number claiming to originate an address prefix will be validated by the prefix holder. The authorized origin will not be suspected by DBGP routers. Further, if the validation can cover the whole AS path, all kinds of attacks that DBGP is trying to cope with SHOULD be prevented in the first place. In all, the deployment of DBGP actually creates the incentive for deploying prevention systems.

7 Security Considerations

The entire document is about security consideration.

8 IANA Considerations

The attribute type code of DAS_PATH should be assigned by the IANA, which identifies the attribute uniquely from all others.

9 References

9.1 Normative References

[PGBGP] J. Karlin, S. Forrest, and J. Rexford, "Pretty Good BGP: Improving BGP by Cautiously Adopting Routes," in Proceedings of IEEE ICNP, 2006.

[SBGP] S. Kent, C. Lynn, and K. Seo, "Secure Border Gateway Protocol (SBGP)," IEEE Journal on Selected Areas in Communications, vol. 18, no. 4, pp. 582-592, 2000.

[SoBGP] J. Ng, "Extensions to BGP to Support Secure Origin BGP," April 2004, ftp://ftp-eng.cisco.com/sobgp/drafts/draft-ng-sobgp-bgp-extensions-02.txt.

[SPV] Y.-C. Hu, A. Perrig, and M. Sirbu, "SPV: Secure Path Vector Routing for Securing BGP," in Proceedings of ACM

SIGCOMM, 2004.

[BGP4] J. W. Stewart, BGP4: Inter-Domain Routing in the Internet. Boston,MA, USA: Addison-Wesley Longman Publishing Co., Inc., 1998.

[pfx-va] P. Mohapatra, J. Scudder, D. Ward, R. Bush, R. Austein, "BGP Prefix Origin Validation", draft-ietf-sidr-pfx-validate-01, working in progress.

9.2 Informative References

[Cyclops] Y.-J. Chi, R. Olivera, and L. Zhang, "Cyclops: the as-level connectivity observatory," SIGCOMM Comput. Commun. Rev., vol. 38, no. 5, pp. 5-16, 2008.

[PHAS] M. Lad, D. Massey, D. Pei, B. Zhang, and L. Zhang, "PHAS: A Prefix Hijack Alert System," in 15th USENIX Security Symposium, 2006, pp.153-166.

[MyASN] "RIPE myASN System," http://www.ris.ripe.net/myasn.html.

[IAR] [Online]. Available: http://iar.cs.unm.edu/

[iSPY] Z. Zhang, Y. Zhang, Y. C. Hu, Z. M. Mao, and R. Bush, "iSPY: Detecting IP Prefix Hijacking on My Own," in Proceedings of ACM SIGCOMM, 2008.

[NWatch] G. Siganos and M. Faloutsos, "Neighborhood Watch for Internet Routing: Can We Improve the Robustness of Internet Routing Today?" in Proceedings of IEEE INFOCOM, 2007.

[Olist] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S. Wu, and L. Zhang,"Dection of Invalid Routing Announcement in the Internet," in Proceedings of the IEEE DSN, June 2002.

[LWeight] C. Zheng, L. Ji, D. Pei, J. Wang, and P. Francis, "A Light-weight Distributed Scheme for Detecting IP Prefix Hijacks in Real-time," in Proceedings of ACM SIGCOMM, 2007.

[RIPE] [Online]. Available: http://www.ripe.net/news/study-youtubehijacking.html

[ASN8997] "Prefix hijack by ASN 8997." [Online]. Available: http://www.merit.edu/mail.archives/nanog/2008-09/msg00704.html

[BGPpop] J. Rexford, J. Wang, Z. Xiao, and Y. Zhang, "BGP Routing Stability of Popular Destinations," in Proceedings of ACM IMC 2002, pp. 197-202.

[Stable] K. Butler, P. McDaniel, and W. Aiello, "Optimizing BGP Security by Exploiting Path Stability," in Proceedings of ACM CCS, Alexandria, VA, United States, 2006, pp. 298-310.

[BGPHA] R. V. Oliveira, R. Izhak-Ratzin, B. Zhang, and L. Zhang, "Measurement of Highly Active Prefixes in BGP," in Proceedings of IEEE Globecom,2005.

Appendix A: Empirical Evaluation

The following aspects of DBGP are tested on the SSFNet-2.0 simulation platform which has implemented BGP4.

- The ability to counter different types of attacks
- The ability to rectify the false positives
- The memory and message overhead

The evaluation proves that DBGP can be used to mitigate all types of attacks. Compared with previous mitigation approaches [PGBGP], it reduces the delay of legitimate announcements significantly, only incurs a small amount of messages and memory overhead.

Author's Addresses

Mingui Zhang
Huawei Technologies Co.,Ltd
HuaWei Building, No.3 Xinxi Rd., Shang-Di
Information Industry Base, Hai-Dian District,
Beijing, 100085 P.R. China

Email: mingui@huawei.com

Bin Liu
Tsinghua University
East Main Building RM9-416
Tsinghua University, Hai-Dian District,
Beijing, 100084 P.R. China

Email: lmyujie@gmail.com

Dacheng Zhang
Huawei Technologies Co.,Ltd
HuaWei Building, No.3 Xinxi Rd., Shang-Di
Information Industry Base, Hai-Dian District,
Beijing, 100085 P.R. China

Email: zhangdacheng@huawei.com

Beichuan Zhang
University of Arizona
Computer Science Department,
The University of Arizona
Tucson, AZ 85721 U.S.A.

Email: bzhang@arizona.edu