Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: August 5, 2013

W. George
Time Warner Cable
S. Murphy
SPARTA, Inc., a Parsons Company
February 1, 2013

BGPSec Considerations for AS Migration
draft-george-sidr-as-migration-01

Abstract

This draft discusses considerations and methods for supporting and securing a common method for AS-Migration within the BGPSec protocol.

Table of Contents

1. Introduction

There is a method of managing an ASN migration using some BGP knobs that while commonly-used are not formally part of the BGP4 [RFC4271] protocol specification and may be vendor-specific in exact implementation. In order to ensure that this behavior is understood and considered for future modifications to the BGP4 protocol specification, especially as it concerns the handling of AS_PATH attributes, the behavior and process has been defined in draft-ga-idr-as-migration [I-D.ga-idr-as-migration]. Accordingly, it is necessary to discuss this de facto standard to ensure that the process and features are properly supported in BGPSec [I-D.ietf-sidr-bgpsec-protocol], because BGPSec is explicitly designed to protect against changes in the BGP AS_PATH, whether by choice, by misconfiguration, or by malicious intent. It is critical that the BGPSec protocol framework is able to support this operationally necessary tool without creating an unacceptable security risk or exploit in the process.

1.1. 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].

2. General Scenario

This draft assumes that the reader has read and understood the ASN migration method discussed in draft-ga-idr-as-migration [I-D.ga-idr-as-migration] including its examples, as they will be heavily referenced here. The use case being discussed in the referenced draft is as follows: For whatever the reason, a provider is in the process of merging two or more ASNs, where eventually one subsumes the other(s). Confederations RFC 5065 [RFC5065] are *not* being implemented between the ASNs, but vendor-specific configuration knobs are being used to allow the migrating PE to masquerade as the old ASN for the PE-CE eBGP session, or to manipulate the AS_PATH, or both. While BGPSec [I-D.ietf-sidr-bgpsec-protocol] does have a case to handle standard confederation implementations, it may not be applicable in this exact case. The reason that this may drive a slightly different solution in BGPSec than a standard confederation is that unlike in a confederation, eBGP peers may not be peering with the "correct" external ASN, and the forward-signed updates are for a public ASN, rather than a private one, so there is no expectation that the BGP speaker should strip the updates before propagating the route to its eBGP neighbors.

In the following examples, AS300 is being subsumed by AS200, and both ASNs represent a Service Provider (SP) network. AS100 and 400 represent end customer networks. References to PE, CE, and P routers mirror the diagrams and references in the above draft.

3. RPKI Considerations

Since the methods and implementation discussed in draft-ga-idr-as-migration [I-D.ga-idr-as-migration] are not technically a part of the BGP4 protocol implementation, but rather a vendor-specific optimization, BGPSec is not technically required to ensure that it continues functioning as it does today. However, this is widely used during network integrations resulting from mergers and acquisitions, as well as network redesigns, and therefore it is not feasible to simply eliminate this capability on any BGPSec-enabled routers/ASNs. What follows is a discussion of the potential issues to be considered regarding how ASN-migration and BGPSec [I-D.ietf-sidr-bgpsec-protocol] validation might interact.

One of the primary considerations for this draft and migration is that companies rarely stop with one merger/acquisition/divestiture, and end up accumulating several legacy ASNs over time. Since they are using methods to migrate that do not require coordination with customers, they do not have a great deal of control over the length of the transition period as they might with something completely under their administrative control like a key roll. This leaves many SPs with multiple legacy ASNs which don't go away very quickly, if at all. As solutions were being proposed for RPKI implementations to solve this transition case, operational complexity and hardware scaling considerations associated with maintaining multiple legacy ASN keys on routers throughout the combined network have been carefully considered. While part of the recommendation may be "SPs SHOULD NOT remain in this transition phase indefinitely because of the operational complexity and scaling considerations associated with maintaining multiple legacy ASN keys on routers throughout the combined network", this is of limited utility as a solution, and so every effort has been made to keep the additional complexity during the transition period to a minimum, on the assumption that it will likely be protracted.

3.1. Origin Validation

Origin Validation does not need a unique solution to enable migration, as the existing protocol and procedure allows for a solution. In the scenario discussed, AS300 is being replaced by AS200. If there are any existing routes originated by AS300 on the router being moved into the new ASN, this simply requires generating

new ROAs for the routes with the new ASN and treating them as new routes to be added to AS200. However, we also need to consider the situation where one or more other PEs are still in AS300, and are originating one or more routes that may be distinct from any that the router under migration is originating. When those routes arrive at PE1, which is now a part of AS200 and instructed to use replace-as to remove AS300 from the path, PE1 needs to be able to handle routes originated from AS300. If the route now shows up as originating from AS200, any downstream peers' validation check will fail unless a ROA is *also* available for AS200 as the origin ASN, meaning that there will be overlapping ROAs until all routers originating prefixes from AS300 are migrated to AS200. Overlapping ROAs are permissible perRFC 6480 [RFC6480] section 3.2, and so managing origin validation during a migration like this is merely applying the defined case where a set of prefixes are originated from more than one ASN. Therefore, for each ROA that authorizes AS300 to originate a prefix, a new ROA SHOULD also be created that authorizes AS200 to originate the same prefix.

3.2. Path Validation

BGPSec Path Validation requires that each router in the AS_PATH cryptographically sign its update to assert that "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." Since this migration technique is explicitly modifying the AS_PATH between two eBGP peers who are not coordinating with one another (are not in the same administrative domain), no level of trust can be assumed, and therefore it may be difficult to identify legitimate manipulation of the AS_PATH for migration activities when compared to manipulation due to misconfiguration or malicious intent.

3.2.1. Outbound announcements (PE-->CE)

When PE1 is moved from AS300 to AS200, it will be provisioned with the appropriate keys for AS200 so that it can begin forward-signing routes using AS200. However, there is currently no guidance in the BGPSec protocol specification on whether or not the forward-signed ASN value MUST match the configured "remote-as" to validate properly. That is, if CE1's BGP session is configured as "remote-as 300", the presence of "local-as 300" on PE1 will ensure that there is no ASN mismatch on the BGP session itself, but if CE1 receives updates from its remote neighbor (PE1) forward-signed from AS200, should the BGPSec validator on CE1 still consider those valid by default? If it does, is there any potential attack vector to consider? RFC4271 [RFC4271] section 6.3 mentions this match between the ASN of the peer and the AS_PATH data, but it is listed as an optional validation, rather than a requirement.

3.2.2. Inbound announcements (CE-->PE)

Inbound is more complicated, because the CE doesn't know that PE1 has changed ASNs, so it is forward-signing all of its routes with AS300, not AS200. The BGPSec speaker cannot manipulate previous signatures, and therefore cannot manipulate the previous AS_Path without causing a mismatch that will invalidate the route. If the updates are simply left intact, the ISP would still need to publish and maintain valid and active public-keys for AS 300 if it is to appear in the BGPSec_Path_Signature in order that receivers can validate the BGPSEC_Path_Signature arrived intact/whole. However, if the updates are left intact, this will cause the AS_PATH length to be increased, which as previously stated is undesirable.

4. Requirements

These requirements are written under the assumption that the currently vendor-specific implementations will be standardized via draft-ga-idr-as-migration [I-D.ga-idr-as-migration], as it makes little sense to build support into a standard for something that is not actually a standard itself. However, should IETF choose not to standardize the discussed method of AS migration, it is possible that this draft could be considered implementation guidance for those vendors that have support for this method of AS migration and wish to support it in their BGPSec implementation. Any solution to the described problem needs to consider the following requirements, listed in no particular order:

- BGPSec MUST support AS Migration for both inbound and outbound route announcements (see Section 3.2.1 and 3.2.2). It SHOULD do this without reducing BGPSec's protections for route path

- MUST NOT require any reconfiguration on the remote eBGP neighbor (CE)

- SHOULD confine configuration changes to the migrating PEs e.g. can't require global configuration changes to support migration

- MUST NOT lengthen AS Path during migration

- MUST operate within existing trust boundaries e.g. can't expect remote side to accept pcount=0 from untrusted/non-confed neighbor

5. Solution

As noted in [I-D.ietf-sidr-bgpsec-protocol], section 4.2, BGPSec already has a solution for hiding ASNs where increasing the AS_PATH length is undesirable. So one might think that a simple solution would be to retain the keys for AS300 on PE1, and forward-sign towards CE1 with AS300 and Pcount=0. However, this would mean passing a pcount=0 between two ASNs that are in different administrative and trust domains such that it could represent a significant attack vector to manipulate BGPSec-signed paths. The expectation for legitimate instances of Pcount=0 (to make a route-server that is not part of the transit path invisible) is that there is some sort of existing trust relationship between the operators of the route-server and the downstream peers such that the peers could be explicitly configured by policy to permit PCount=0 announcements only on the sessions where they are expected, and otherwise reject them. For the same reason that things like local-as are used for ASN migration without end customer coordination, it is unrealistic to assume any sort of coordination between the SP and the administrators of CE1 to ensure that they will by policy accept PCount=0 signatures during the transition period, and therefore this is not a workable solution.

However, a better solution presents itself when considering how to handle routes coming from the CE toward the PE, where the routes are forward-signed to AS300, but will eventually need to show AS200 in the outbound route announcement. Because both AS200 and AS300 are in the same administrative domain, a signature from AS300 forward-signed to AS200 with Pcount=0 would be acceptable as it would be within the appropriate trust boundary so that each BGP speaker could be explicitly configured to accept Pcount=0 where appropriate between the two ASNs. At the very simplest, this could potentially be used at the eBGP boundary between the two ASNs during migration. But since the AS_PATH manipulation described above usually happens at the PE router on a per-session basis, and does not happen network-wide simultaneously, it is not generally appropriate to apply this AS hiding technique across all routes exchanged between the two ASNs, and may result in routing loops and other undesirable behavior. Therefore the most appropriate place to implement this is on the local PE that still has eBGP sessions associated with AS300 (using the transition knobs detailed in the companion draft). Since that PE has been moved to AS200, it is not possible for it to forward-sign AS300 with Pcount=0 without some minor changes to the BGPSec implementation to address this use case.

AS migration is using AS_PATH and remote-AS manipulation to act as if a PE under migration exists simultaneously in both ASNs even though it is only configured with one global ASN. This draft proposes

applying a similar technique to the BGPSec signatures generated for routing updates processed through this migration machinery. Each routing update that is received from or destined to an eBGP neighbor that is still using the old ASN (300) will be signed twice, once with the ASN to be hidden and once with the ASN that will remain visible. In essence, we are treating the update as if the PE had an internal BGP hop and the update was passed across an eBGP session between AS200 and AS300, configured to use and accept Pcount=0, while eliminating the processing and storage overhead of actually creating an actual eBGP session within the PE router. This will result in a properly secured AS_PATH attribute, because the PE router will be provisioned with valid keys for both AS200 and AS300. The procedure is slightly different depending on whether the PE under migration is receiving the routes from one of its eBGP peers ("inbound" as in section 3.2.2) or destined toward the eBGP peers ("outbound" as in section 3.2.1).

5.1. Outbound (PE->CE)

When a PE router receives an update destined for an eBGP neighbor that is locally configured with AS-migration knobs as discussed in draft-ga-idr-as-migration [I-D.ga-idr-as-migration] to facilitate a move from an old ASN to a new one, it MUST generate a valid BGPSec signature as defined in [I-D.ietf-sidr-bgpsec-protocol] for _both_ configured ASNs. It MUST generate a signature from the new (global) ASN forward signing to the old (local) ASN with Pcount=0, and then it MUST generate a forward signature from the old (local) ASN to the target eBGP ASN with Pcount=1 as normal.

5.2. Inbound (CE->PE)

When a PE router receives an update from an eBGP neighbor that is locally configured with AS-migration knobs (i.e. the opposite direction of the previous route flow), it MUST generate a signature from the old (local) ASN forward signing to the new (global) ASN with PCount=0. It is not necessary to generate the second signature from the new (global) ASN because the ASBR will generate that when it forward signs towards its eBGP peers as defined in normal BGPSec operation. This is a deviation from standard BGPSec behavior in that typically a signature is not added when a routing update is sent across an iBGP session, and the next signature is added by the ASBR when it forward-signs toward its eBGP peer as the routing update exits the ASN.

5.3. Other considerations

In this case, the PE is adding BGPSec attributes to routes received from or destined to an iBGP neighbor, and using PCount=0 to mask

them. While this is not prohibited by the current BGPSec specification, routers that receive updates from iBGP neighbors MUST NOT reject updates with new (valid) BGPSec attributes, including the presence of PCount=0 on a previous signature, or they will interfere with this implementation. In similar fashion, any route-reflectors in the path of these updates MUST reflect them transparently to their clients.

In order to secure this set of signatures, the PE router MUST be provisioned with valid keys for _both_ configured ASNs (old and new), and the key for the old ASN MUST be kept valid until all eBGP sessions are migrated to the new ASN. Downstream neighbors will see this as a valid BGPSec path, as they will simply trust that their upstream neighbor accepted Pcount=0 because it was explicitly configured to do so based on a trust relationship and business relationship between the upstream and its neighbor (the old and new ASNs).

5.4. Example

The following example will illustrate the method being used above. As with previous examples, PE1 is the router being migrated, AS300 is the old AS, which is being subsumed by AS200, the "keep" AS. Some additional notation is used to delineate the details of each signature as follows:

The origin BGPSEC signature attribute takes the form: sig(<Target ASN>, Origin ASN, pcount, NLRI Prefix) key

Intermediate BGPSEC signature attributes take the form: sig(<Target ASN>, Signer ASN, pcount, <most recent sig field>) key

Before Merger

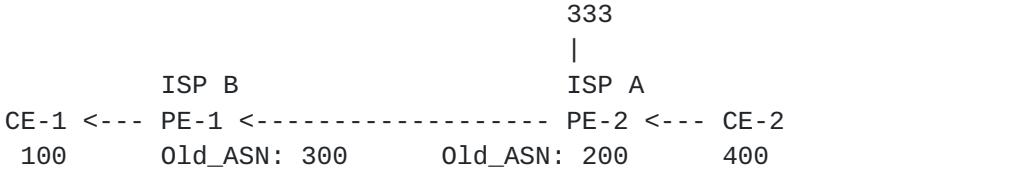

CE-2 to PE-2: sig(<200>, O=400, pcount=1, N)K_400-CE2 [sig1]
AS_PATH=(400)
length=sum(pcount)=1

PE-2 to 333: sig(<333>, 200, pcount=1, <sig1>)K_200-PE2 [sig2]
sig(<200>, 400, pcount=1, N)K_400-CE2 [sig1]
AS_PATH=(200,400)
length=sum(pcount)=2

PE-2 to PE-1: sig(<300>, 200, pcount=1, <sig1>)K_200-PE2 [sig3]
sig(<200>, 400, pcount=1, N)K_400-CE2 [sig1]
AS_PATH=(200,400)
length=sum(pcount)=2

PE-1 to CE-1: sig(<100>, 300, pcount=1, <sig3>)K_300-PE1 [sig4]
sig(<300>, 200, pcount=1, <sig1>)K_200-PE2 [sig3]
sig(<200>, 400, pcount=1, N)K_400-CE2 [sig1]
AS_PATH = (300,200,400)
length=sum(pcount)=3

```
   Migrating, route flow outbound PE-1 to CE-1
                                      333
                                      |
             ISP A'                   ISP A'
   CE-1 <--- PE-1 <------------------ PE-2 <--- CE-2
    100      Old_ASN: 300      Old_ASN: 200      400
             New_ASN: 200      New_ASN: 200


   CE-2 to PE-2:  sig(<200>, 400, pcount=1, N)K_400-CE2  [sig11]
                  AS_PATH=(400)
                  length=sum(pcount)=1

   PE-2 to 333:   sig(<333>, 200, pcount=1, <sig11>)K_200-PE2  [sig12]
                  sig(<200>, 400, pcount=1, N)K_400-CE2  [sig11]
                  AS_PATH=(200,400)
                  length=sum(pcount)=2

   PE-2 to PE-1:  [sig11]

   PE-1 to CE-1:  sig(<100>, 300, pcount=1, <sig13>)K_300-PE1  [sig14]
                  sig(<300>, 200, pcount=0, <sig11>)K_200-PE2  [sig13]
                  sig(<200>, 400, pcount=1, N)K_400-CE2  [sig11]
                  AS_PATH=(300,400)
                  length=sum(pcount)=2 (length is NOT 3)
   #PE1 adds [sig13] acting as AS200
   #PE1 accepts [sig13] with PCount=0 acting as AS300,
   #as it would if it received sig13 from an eBGP peer
```

```
Migrating, route flow inbound CE-1 to PE-1
                                    333
                                    |
          ISP A'                    ISP A'
CE-1 ---> PE-1 -------------------> PE-2 ---> CE-2
 100      Old_ASN: 300      Old_ASN: 200      400
          New_ASN: 200      New_ASN: 200


CE-1 to PE-1:  sig(<300>, 100, pcount=1, N)K_100-CE1   [sig21]
               AS_PATH=(100)
               length=sum(pcount)=1

PE-1 to PE-2:  sig(<200>, 300, pcount=0, <sig21>)K_300-PE1  [sig22]
               sig(<300>, 100, pcount=1, N)K_100-CE1   [sig21]
               AS_PATH=(100)
               length=sum(pcount)=1 (length is NOT 2)
#PE1 adds [sig22] acting as AS300
#PE1 accepts [sig22] with PCount=0 acting as AS200,
#as it would if it received sig22 from an eBGP peer

PE-2 to 333:   sig(<333>, 200, pcount=1, <sig22>)K_200-PE2  [sig23]
               sig(<200>, 300, pcount=0, <sig21>)K_300-PE1  [sig22]
               sig(<300>, 100, pcount=1, N)K_100-CE1   [sig21]
               AS_PATH=(200,100)
               length=sum(pcount)=2 (length is NOT 3)

PE-2 to CE-2:  sig(<400>, 200, pcount=1, <sig22>)K_200-PE2  [sig24]
               sig(<200>, 300, pcount=0, <sig21>)K_300-PE1  [sig22]
               sig(<300>, 100, pcount=1, N)K_100-CE1   [sig21]
               AS_PATH=(200,100)
               length=sum(pcount)=2 (length is NOT 3)
```

6. Acknowledgements

Thanks to Kotikalapudi Sriram and Shane Amante for their review comments.

Additionally, the solution presented in this draft is an amalgam of several SIDR interim meeting discussions plus a discussion at IETF85, collected and articulated thanks to Sandy Murphy.

7. IANA Considerations

This memo includes no request to IANA.

8. Security Considerations

This draft discusses a process by which one ASN is migrated into and subsumed by another. Because this involves manipulating the AS_Path to make it deviate from the actual path that it took through the network, it is in some ways attempting to do exactly what BGPSec is working to prevent. The BGPSec implementation MUST be able to manage this legitimate use of AS_Path manipulation without generating a vulnerability in the RPKI route security infrastructure that can be exploited by a malicious actor.

The solution discussed above is considered to be reasonably secure from exploitation by a malicious actor because it requires both signatures to be secured as if they were forward-signed between two eBGP neighbors. This requires any router using this solution to be provisioned with valid keys for both the migrated and subsumed ASN so that it can generate valid signatures for each of the two ASNs it is adding to the path. If the AS's keys are compromised, or zero-length keys are permitted, this does potentially enable an AS_PATH shortening attack, but this is not fundamentally altering the existing security risks for BGPSec.

9. References

9.1. Normative References

[I-D.ga-idr-as-migration]
George, W. and S. Amante, "Autonomous System (AS) Migration Features and Their Effects on the BGP AS_PATH Attribute", draft-ga-idr-as-migration-00 (work in progress), September 2012.

[I-D.ietf-sidr-bgpsec-protocol]
Lepinski, M., "BGPSEC Protocol Specification", draft-ietf-sidr-bgpsec-protocol-06 (work in progress), October 2012.

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

9.2. Informative References

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

[RFC5065] Traina, P., McPherson, D., and J. Scudder, "Autonomous System Confederations for BGP", RFC 5065, August 2007.

[RFC6480] Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", RFC 6480, February 2012.

Authors' Addresses

Wesley George
Time Warner Cable
13820 Sunrise Valley Drive
Herndon, VA 20171
US

Phone: +1 703-561-2540
Email: wesley.george@twcable.com

Sandy Murphy
SPARTA, Inc., a Parsons Company
7110 Samuel Morse Drive
Columbia, MD 21046
US

Phone: +1 443-430-8000
Email: sandy@tislabs.com