Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 02, 2013

R. Bush
Internet Initiative Japan
K. Patel
P. Mehta
A. Sreekantiah
Cisco Systems
L. Jalil
Verizon

October 2012

Authenticating L3VPN Origination Signaling
draft-ymbk-l3vpn-origination-02

Abstract

A BGP-signaled Layer-3 VPN’s prefix bindings sent over BGP are subject to unintentional errors, both by the legitimate originator and by non-legitimate origins. This is of special concern if the VPN traverses untrusted networks. This document describes how the sender of the Prefix/VPN binding may sign it so that recipient of the binding may authenticate it.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in RFC 2119 [RFC2119] only when they appear in all upper case. They may also appear in lower or mixed case as English words, without normative meaning.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

RFC 4364 [RFC4364] Section 7.4 describes how a Customer Edge (CE) router uses eBGP to announce to a Provider Edge (PE) router the address prefix(es) the customer provides to an L3VPN. It is possible that the originator of such an announcement could unintentionally announce prefixes they do not own.

```
      Cust(West)-CE--PE-Provider(West)--TransitA-~
            ~-TransitB--Provider(East)-PE--CE-Cust(East)
```

This document describes how the PE receiving the CE's originating announcement, West, may sign the announcement so that the PE proximal to the destination CE, East, may authenticate the NLRI see RFC 4364 [RFC4364] Section 4.3.1. Alternatively, the originating CE router may sign the announcement so that the destination CE router may authenticate the NLRI.

It is assumed that the providers already have the key creation, storage, and distribution infrastructure needed. Keys might be configured on the routers, or in some shared PKI, or, for example, the Resource Public Key Infrastructure (RPKI) could be used, see RFC 6480 [RFC6480].

A new BGP PATH Attribute, called L3VPN Origination BGP PATH Attribute (L3OPA), is created to contain the necessary keying information and signature.

2. NLRI Deaggregation

Normally, a BGP Update may contain multiple NLRI which all share the identical set of attributes. As L3OPA signalling signs over the NLRI, and NLRI can become separated as they transit the network, separation would break the signature. Therefore, a BGP announcement using L3OPA signalling MUST contain one and only one NLRI.

3. L3VPN Origination BGP Path Attribute (L3OPA)

The L3OPA is a BGP optional transitive Path Attribute RFC 4271 [RFC4271]. BGP Path Attributes are Type/Length/Value tuples.

```
    .--------------------------------------------.
    |                                            |
    |  Attribute Header, see RFC 4271 S. 4.3     |
    |                                            |
    +--------------------------------------------+
    |                                            |
    |                                            |
    |                                            |
    +----          Key Identifier           ----+
    |                                            |
    |                                            |
    |                                            |
    +--------------------------------------------+
    |    Alg   |   MUST   |                      |
    |   Suite  |    be    |   Signature Length   |
    |     1    |   Zero   |                      |
    +--------------------------------------------+
    |                                            |
    |                 Signature                  |
    |                                            |
    '--------------------------------------------'
```

The Attribute Type is two octets, the first of which, Attribute Flags, MUST have the two high order bits set to signify that attribute is optional and transitive.

The second octet of the Attribute Flags, Attribute Type, MUST be set to 0xXX, as assigned by the IANA, see Section 8, to signal that this is an L3OPA.

The Length field is one or two octets with a value of the number of octets in the entire attribute. If the length of the L3OPA is less than 256 octets, only the first octet of the length field is used. Otherwise, both octets are used to represent the Length.. See RFC 4271 [RFC4271] Section 4.2 for another explanation of this byte saving.

The Key Identifier is an eight octet value identifying the key (pair) used for the Signature. It is used when the keying is not implied by the NLRI, as it would be, for example, if the RPKI was used. It is often the VPN Identifier. If not used to identify the key, it MUST be zero.

The Algorithm Suite is a one-octet identifier specifying the digest algorithm and digital signature algorithm used to produce the Signature. The values reference the IANA registry for Algorithm Identifiers from BGPsec, see [I-D.ietf-sidr-bgpsec-algs].

The Signature Length is two octets and is the number of octets in the Signature field.

The Signature field is a digital signature that covers the NLRI and the Key Identifier.

To compute the Signature, the digest algorithm for the specified Algorithm Suite is applied to the catenation of the NLRI and the Key Identifier. This is then fed to the signature algorithm for the specified algorithm suite and the resulting value is the Signature.

```
   Signature = sign ( hash ( NLRI || Key Identifier ) )
```

4. Validation of Routes Having an L3OPA

A BGP speaker receiving routes with an L3OPA MUST perform the necessary validation if configured to do so.

The digest algorithm for the specified Algorithm Suite is applied to the catenation of the NLRI and the Key Identifier. This is then fed to the signature algorithm for the specified algorithm suite and the resulting value is compared with the Signature.

If the signature value matches the Signature in the attribute, the route MUST be marked as Valid, otherwise it MUST be marked as Invalid.

A route received without an L3OPA SHOULD be marked as having an Unknown validity state.

If L3OPA marking is disabled in the router configuration, routes are considered to have the Unknown validity state.

Configured local policy on the router may use the validity state markings to implement policy. For example, a route marked as Invalid or Unknown may be dropped or de-preferenced by appropriate use of normal BGP policy mechanisms.

Note that this is similar to annuncement marking while allowing the user to control policy as described in RPKI-Based BGP origin validation, see [I-D.ietf-sidr-pfx-validate].

5. L3VPN Deployment Scenarios

The following L3VPN deployment scenarios illustrate use of the scheme. The examples use the language of symmetric keys which have been previously agreed upon between the signer of the route and the validator. Asymmetric keying, a PKI, etc. could also be used. Signing and validation are as described above.

5.1. End CE to CE Authentication

```
CE1 ---- PE1 --------- PE2 -- CE2
                AS1

CE1 ---- PE1 --------- ASBR1 ----- ASBR2 -------- PE2 ---- CE2
                AS1                         AS2
```

CE1 and CE2 are end CEs in the same VPN. PE1, PE2, ASBR1, ASBR2 are provider PE/ASBRs which are blindly propagating the announcement with the L3OPA as generated by CE1.

As the authorization is between the originating CE1 and the terminating CE2, the keying should not be known by the provider(s). The CEs are configured with the keying information, the originating CE1 creates and signs an L3OPA for each NLRI participating in the VPN.

An update received by CE2 without an L3OPA, or having an Invalid Signature would likely be dropped. Thus the CEs are protected from incorrect prefixes originating from a provider network or unauthorized CEs.

5.2. Provider/ASBR Based Validation/Authentication

```
CE1 ---- PE1 --------- ASBR1 ----- ASBR2 -------- PE2 ---- CE2
                 AS1                        AS2
```

In the diagram, CE1 is the originating/signing CE. ASBR2 is the trusted provider with whom CE1 has collaborated. Updates generated by CE1 may be passed transparently through any number of intermediate providers, ASBR1s, which blindly propagate the L3OPA. Validation is performed when the announcement reaches the trusted validating provider, ASBR2.

Keying is agreed between CE1 and the trusted provider ASBR2, likely on per-VPN basis.

5.3. PE-PE Based Validation

```
   CE1 ------ PE1 --------- ASBR1 ---- ASBR2 -------- PE2 ------ CE2
                         AS1             |      AS2
                                         |
                      CE3 ---- ASBR3 ----+
                           AS3
```

Here PEs, possibly across ASes, agree on the keying. The Key Identifier and associated keys would normally be configured on a per VPN basis, with the PE1 signing and PE2 and PE3 validating similarly to the CEs in the previous examples.

CE1 originates an announcement, possibly with multiple NLRI, but without an L3OPA. PE1 de-aggregates the NLRI into separate announcements, signs each with the keying agreed with PE2 and PE3, and propagates them. Arbitrary providers carry the announcements toward PE2 and PE3, where the announcements have their Signatures validated, the L3OPAs removed, and are then propagated to CE2 and CE3.

6. Notes

The keying could either come from the Global RPKI or the customer or carrier running their own PKI. The keying is assumed to be asymmetric, but possibly could be symmetric. The keys can be statically configured (beware scaling and key-roll issues), dynamic, in some public or private infrastructure, etc.

If the RPKI is used, and the public key is taken from the CA certificate which owns the NLRI, the classic problem arises where all the NLRI on that certificate share fate. I.e. if one causes the need for a re-key, then all must re-key. RPKI-based origin validation solves this problem by a level of indirection, the CA certificate is used to sign an End Entity (EE) certificate which signs a Route Origin Authorization (ROA), see RFC 6480 [RFC6480] and RFC 6482 [RFC6482]. As the Key Identifier of an L3VPN signal is larger than the four octets of a ROA, a new RPKI object, for the moment let’s call it a VOA, would have to be defined and then it would have to be carried in the RPKI-Router Protocol [I-D.ietf-sidr-rpki-rtr].

If the value of the signing key, as identified by the Key Identifier, is to be rolled, in case of compromise or security policy, the technique in RFC 4808 [RFC4808] should be used.

While it is poor security practice to trust a different entity for your security/authentication/..., should a non-validating router choose to trust a validating router, they could use normal policy and signaling mechanisms, e.g. communities, to signal validation status. This page is too small to enumerate the vulnerabilities this creates.

7. Security Considerations

Signing (NLRI || Key Identifier) with the key of the NLRI-owner or some other pre-agreed key, only says that the contents were produced by the owner of the key (NLRI or other), and that no one in between has changed the (NLRI || Key Identifier). This is not protection against attacks, only configuration errors, aka 'fat fingers'. If we were trying to protect against an attacking PE replaying a signed (NLRI || Key Identifier) it has no business announcing, this design does not help.

If Key Identifier based keying is used, then the Key Identifier, and hence the signing key, MUST be unique to the VPN.

Adding a VOA which binds (NLRI || Key Identifier) still could be replayed from anywhere so really offers nothing. Like RPKI-based origin validation, this only catches fat fingers, not black hats.

8. IANA Considerations

This document requests the IANA create a new entry in the BGP Path Attributes Registry as follows:

Value	Code	Reference
TBD	L3VPN Origination	This Document

9. Acknowledgements

The authors would like to thank Eric Rosen, John Scudder, Russ Housley, and Sandy Murphy.

We note the long expired draft draft-ietf-l3vpn-auth by Ron Bonica, Yakov Rekhter, Eric Rosen, Robert Raszuk, and Dan Tappan.

10. References

10.1. Normative References

[I-D.ietf-sidr-bgpsec-algs]
Turner, S., "BGP Algorithms, Key Formats, & Signature Formats", Internet-Draft draft-ietf-sidr-bgpsec-algs-03, September 2012.

[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.

[RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, February 2006.

[RFC4808] Bellovin, S., "Key Change Strategies for TCP-MD5", RFC 4808, March 2007.

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

10.2. Informative References

[I-D.ietf-sidr-pfx-validate]
Mohapatra, P., Scudder, J., Ward, D., Bush, R. and R. Austein, "BGP Prefix Origin Validation", Internet-Draft draft-ietf-sidr-pfx-validate-10, October 2012.

[I-D.ietf-sidr-rpki-rtr]
Bush, R. and R. Austein, "The RPKI/Router Protocol", Internet-Draft draft-ietf-sidr-rpki-rtr-26, February 2012.

[RFC6482] Lepinski, M., Kent, S. and D. Kong, "A Profile for Route Origin Authorizations (ROAs)", RFC 6482, February 2012.

Authors’ Addresses

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

Email: randy@psg.com

Keyur Patel
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95134
USA

Email: keyupate@cisco.com

Pranav Mehta
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95134
USA

Email: pmehta@cisco.com

Arjun Sreekantiah
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95134
USA

Email: asreekan@cisco.com

Luay Jalil
Verizon
1201 E Arapaho Rd.
Richardson, TX 75081
USA

Email: luay.jalil@verizon.com