Bhargav Bhikkaji
Balaji Venkat Venkataswami
Dell-Force10
February 29, 2012

Preventing spoofing attacks in BGP-MPLS-VPN Inter-Provider Model-C
draft-bhargav-l3vpn-inter-provider-optcsec-01

Abstract

In certain models of inter-provider Multi-Protocol-Label-Switching based Virtual Private Networks (MPLS-VPNs), spoofing attacks against VPN sites is a key concern. Unidirectional attacks towards VPN sites can compromise servers at the VPN sites and cause Denial-of-Service (DoS) situations. Currently, the inner labels associated with VPN sites are not encrypted during transmission. The Provider Edge (PE) router at the end to which the VPN customer is attached accepts any data packet with a valid label. This enables a man-in-the-middle attacker to spoof a packet to a specific site of a VPN. In this paper, we propose some secure techniques which provide security against such label-spoofing. These techniques ensure that an attacker would not be able to spoof labeled data packets. In order to make the proposed scheme robust, some additional steps are proposed over and above the initial steps specified. This makes the attacker to spend non-linear time to guess the right label for his unidirectional attacks to succeed. Our proposed technique can be applied to a specific type of inter-provider Border Gateway Protocol(BGP) based MPLS VPN and other existing variant where Multi-Protocol exterior-BGP (MP-eBGP) multi-hop is used. In future, if any other variant is proposed to use MP-eBGP multi-hop, our scheme can be used to protect against spoofing attacks.

Status of this Memo

Copyright and License Notice

Table of Contents

1 Introduction

MPLS technology helps forward data packets in the Internet using fixed-size labels [3]. By stacking labels (label-stacking technique), specific customer services such as layer 3 VPNs based on BGP extensions are widely deployed in the Internet. BGP based MPLS layer 3 VPN services are provided either on a single internet service provider (ISP) core or across multiple ISP cores. The latter cases are known as inter-provider MPLS VPNs which are broadly categorized and referred to as models: A, B and C. Model A uses back-to-back VPN Routing and Forwarding (VRF) connections between Autonomous System Border Routers (ASBRs). Model B uses eBGP redistribution of labelled VPN-IPv4 routes from AS to neighbouring AS. Model C uses multi-hop eBGP redistribution of labelled VPNIPv4 routes and eBGP redistribution of IPv4 routes from AS to neighbouring AS. Please refer [1] for more details. A. Security issues in inter-provider VPN models Model A is as secure a standard as the single-Autonomous System (AS) standard proposed in [5]. Model B can be secured well on the control plane, but on the data-plane the top-most label is not checked for validity. This weakness could be exploited to inject crafted packets from inside an MPLS core into the other. There is a work-around discussed in [1] that solves the problem. Model C can also be secured well on the control plane. But as the ASBRs do not have any VPN information and the inner-most label cannot be validated, the data-plane has architectural weakness with respect to security. This enables unidirectional packets to be sent into the VPN sites connected to the other AS, which cannot be protected against by mere configuration. Model C can therefore only be deployed where service providers trust each other. For more details refer to [1].

In [2] the authors propose encryption techniques, such as IPSec, for securing the provider edge (PE) to PE legs of the network. However the authors also highlight that the processing capacity could be over-burdened. If an attacker is located at the core of the network, or in the intermediate link or network between the providers that constitute a inter-provider MPLS VPN solution, then spoofing attacks are possible. In case an attacker spoofs the inner labels that identify packets going towards a L3 VPN site, sensitive information related to services available within the organizational servers can be compromised. A denial-of-service (DoS) attack could also be launched against these sensitive sites. As far as we know, there is no scheme adopting encryption available for installing an anti-spoofing mechanism for these VPN service labels. The proposed scheme in this document provides an alternative in case other schemes that dont adopt encryption are not suitable. The PEs at the end to which the VPN customer is attached will accept any packet with a valid label and will forward it to the VPN customer. There is no way to

ensure the veracity of a spoofed packet.

1.1 Security issue in model C

The deficiency discussed above is particularly true in the case of inter-provider BGP based MPLS VPN model C. Even though model C is highly scalable for carrying VRF routes, the vulnerability of the data-plane has rendered it unusable and the current recommendation is that model C must not be used. As discussed in [1] the insecurity for model C stems from the fact that anybody within the core of the network or at the peering points of the providers can cause DoS attacks or worm attacks. It is possible to filter all IP traffic with the exception of the required eBGP peering between the AS border routers thereby preventing a large number of potential IP traffic related attacks. Labelled packets, however, are much harder to control. In model C, there are at least two labels for each packet: the PE label, which defines the Label Switched Path (LSP) to the egress PE, and the VPN label, which defines the VPN on the PE to which the packet is associated.

1.2 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 [RFC2119].

2. Methodology of the Proposal

Consider a setup as below

The reference MPLS-eBGP based VPN network for model-C / Option-C as described in [11] is shown in Figure 1, which also shows the control plane exchanges. The near-end PE (PE_ne) and far-end PE (PE_fa) are connected through the inter-provider MPLS core. The VPN connectivity is established through a set of routers from different Autonomous Systems (AS) and their ASBRs. In the VPN, MP-eBGP updates are exchanged for a set of Forward Equivalence Classes (FECs). These FECs, which have to be protected, originate from the prefixes behind PE_ne in a VPN site or a set of VPN sites.

```
_______[AB1]________________           ___________[AB2]___________
(       /  |                 )        (           /   \          )
(      /   +-----------+      )       (          /     \         )
( Net:PE_ne        Net:PE_ne   )      (     Net:PE_ne  Net:PE_ne )
( LDP Label:POP    LDP Label:L1 )     (LDP Label:L3  LDP Label:L4)
(        |                 |    )     (  |                    |  )
(        |                 |    )     (  |                    |  )
[CE1]<-[PE_ne]__________[ASBR1]<---->[ASBR2]_________________[PE_fa]
                                                                   |
172.18.10.0/24                 NH:PE_ne          172.18.20.0/24 [CE2]
                               VPN Label: Inner Label IL1
```

For the example below we refer to PE_ne and PE_fa in the diagram as PE-1 and PE-2.

1) PE-1 and PE-2 would establish a MPeBGP-VPNV4 session between them

2) PE-1 and PE-2 will exchange a VPN-label between them which is used during forwarding

3) This provides a way for customers of PE-1 to use same address space because different routing context is provided for different customer at PE

4) This mechanism does not avoid an attacker who is trying to spoof a packet somewhere inside the core

There are 2 ways to solve this problem

2.1 Solution:1

2.1.1 Control Plane operation

1) PE-1 and PE-2 agree upon looking for "Secure Label" in addition to VPN label

2) PE-1 and PE-2 use PKI and publish their respective public keys.

```
_______[AB1]________________           ___________[AB2]___________
(       /  |                 )        (           /   \          )
(      /   +-----------+      )       (          /     \         )
( Net:PE_ne        Net:PE_ne   )      (     Net:PE_ne  Net:PE_ne )
( LDP Label:POP    LDP Label:L1 )     (LDP Label:L3  LDP Label:L4)
(        |                 |    )     (  |                |      )
(        |                 |    )     (  |                |      )
[CE1]<-[PE_ne]__________[ASBR1]<---->[ASBR2]__________[PE_fa]-->[CE2]
         ^                                                      ^
         +-------------- MP-eBGP IPVPN-V4 session --------------+

         (2) PKI exchange to exchange their respective public keys.
```

3) PE-1 and PE-2 also agree upon a universal hashing algorithm to use. It could be any of standard universal hashing algorithm which minimizes collisions to the maximum extent possible. They also agree upon a bitmask that is to be used in step 5.1. These could be exchanged using the MP-eBGP label exchange mechanism using suitable fields which will be discussed in the appendix A.1.

4) Unlike VPN label that is exchanged as part MPBGP-VPNV4 operation, "Secure Label" is generated on fly during forwarding.

```
_______[AB1]________________           ___________[AB2]___________
(       /  |                 )        (           /   \          )
(      /   +-----------+      )       (          /     \         )
( Net:PE_ne        Net:PE_ne   )      (     Net:PE_ne  Net:PE_ne )
( LDP Label:POP    LDP Label:L1 )     (LDP Label:L3  LDP Label:L4)
(        |                 |    )     (  |                |      )
(        |                 |    )     (  |                |      )
[CE1]<-[PE_ne]__________[ASBR1]<---->[ASBR2]__________[PE_fa]-->[CE2]
         ^                                                ^
         +-------------- MP-eBGP IPVPN-V4 session --------+

172.18.10.0/24    (4)   NH:PE_ne                       172.18.20.0/24
                        VPN Label: Inner Label IL1
                        Universal hashing Algorithm: UA1,
                        Bitmask: B1
```

5) Secure label is generated for prefixes reachable via PE-1 by PE-2

5.1) First take an hash on the first 256 bytes of the packet's payload received using the universal hashing algorithm agreed upon. If the PEs are from the same vendor or from a different vendor the algorithm to use is expressed in a standard identifier which is understood by both PEs.

5.2) Take any 20 bits from the hash using a bitmask that is also shared amongst the two PEs.

5.3) Take this 20 bit value and encrypt with the private key of PE-2 if traffic is flowing from PE-2 to PE-1. The inner label, universal hashing algorithm and the bitmask are sent from PE-1 to PE-2 for prefixes reachable via PE-1.

2.1.2 **Data-Plane Operation**

```
_______[AB1]________________           ___________[AB2]___________
(       /  |                 )        (         /     \           )
(      /   +------------+      )       (        /       \          )
(IL1:SL:172.18.10.1    L1:IL1:SL:172.18.10.1  /         \         )
(        |           L3:IL1:SL:172.18.10.1 <-+ L4:IL1:SL:172.18.10.1)
(        |                 |    )     (  |                 |      )
(        |                 |    )     (  |                 |      )
[CE1]<-[PE_ne]__________[ASBR1]<---->[ASBR2]__________[PE_fa]-->[CE2]

172.18.10.0/24        L2:IL1:SL:172.18.10.1            172.18.20.0/24
```

1) Let us assume a customer connected to PE-2 is sending a packet to PE-1

2) The packet could be any payload encapsulated in the MPLS header having a stack of labels for MODEL-C

3) When the packet reaches PE-2, following operation is done

3.1) Generate an hash of the received packet using the universal hashing algorithm chosen.

3.2) Take agreed 20 bits from the hash using the bitmask exchanged. (Text below talks about what 20 bits to take)

3.3) Encrypt those 20 bits with PE-2's private key

3.4) Send the packet with VPN label on top of the secure label

4) When the packet reaches PE-1, following operation is done

4.1) Generate an hash on the 256 bytes of data after the secured

label

4.2) Take agreed 20 bits using bitmask agreed upon after hash calculation in #4.1 by PE-1, call it 'K'

4.3) Decrypt received secured-label using public key of PE-1

4.4) Check if decrypted secured label and K are same.

4.5) If both are same then forward the packet otherwise drop them

2.1.3 Mapping a hash to a 20 bit value

1) The universal hashing algorithm selected generates a 128 bit value but a MPLS label is 20 bit value

2) A mechanism is necessary to map 128 bit to 20 bit

3) PE-1 and PE-2 would exchange a 128 bitmask which is used to convey what 20 bits in that 128 bits is to be used for the purposes of encryption

4) This bitmask / bitmap is generated randomly and exchanged at the time of MP-eBGP NLRI exchanges and expires every t seconds.

5) Whenever a new bitmap is generated, this would be shared between the PE's

2.2 Solution:2

2.2.1 Control Plane operation

```
_______[AB1]________________          ___________[AB2]___________
(       /  |                  )       (           /   \           )
(      /   +-----------+       )      (          /     \          )
( Net:PE_ne        Net:PE_ne   )      (     Net:PE_ne  Net:PE_ne )
( LDP Label:POP    LDP Label:L1 )     (LDP Label:L3  LDP Label:L4)
(        |                 |     )    (  |                |       )
(        |                 |     )    (  |                |       )
[CE1]<-[PE_ne]__________[ASBR1]<---->[ASBR2]__________[PE_fa]-->[CE2]
         ^                                                     ^
         +-------------- MP-eBGP IPVPN-V4 session --------------+

         (1) PKI exchange to exchange their respective public keys.
```

1) PE-1 and PE-2 use PKI and publish their respective public keys.

2) PE-1 and PE-2 agree upon looking for an digital signature below

EOS label. This would be done with an appropriate indicator in the MP-eBGP update which would tell the other PE (far-end PE) that a digital signature mechanism is to be used for purposes of protecting the payload / stream between the two PEs.

```
_______[AB1]________________           ___________[AB2]___________
(       /  |                 )        (           /   \          )
(      /   +-----------+      )       (          /     \         )
( Net:PE_ne        Net:PE_ne   )      (     Net:PE_ne  Net:PE_ne )
( LDP Label:POP    LDP Label:L1 )     (LDP Label:L3  LDP Label:L4)
(        |                 |    )     (  |                |      )
(        |                 |    )     (  |                |      )
[CE1]<-[PE_ne]__________[ASBR1]<---->[ASBR2]__________[PE_fa]-->[CE2]
         ^                                                 ^
         +-------------- MP-eBGP IPVPN-V4 session --------+

172.18.10.0/24    (4) NH:PE_ne                        172.18.20.0/24
                      VPN Label: Inner Label IL1
                      Universal hashing Algorithm: UA1,
                      Digital Signature mechanism
```

3) PE-1 and PE-2 also agree upon what universal hashing algorithm to use.

4) Digital signature is generated as follows

4.1) First take a hash on the first 256 bytes of the packet received

4.2) Encrypt with private key available for the transaction.

2.2.2 Data-Plane Operation

```
_______[AB1]________________           ___________[AB2]___________
(       /  |                 )        (          /     \          )
(      /   +-----------+      )       (         /       \         )
(IL1:DS:172.18.10.1    L1:IL1:DS:172.18.10.1   /         \        )
(        |           L3:IL1:DS:172.18.10.1 <-+ L4:IL1:DS:172.18.10.1)
(        |                 |    )     (  |                 |      )
(        |                 |    )     (  |                 |      )
[CE1]<-[PE_ne]__________[ASBR1]<---->[ASBR2]__________[PE_fa]-->[CE2]

172.18.10.0/24        L2:IL1:DS:172.18.10.1          172.18.20.0/24
```

1) Let us assume a customer connected to PE-1 is sending a packet to PE-2

2) The packet could be any payload encapsulated in the MPLS header having a stack of labels for MODEL-C

3) When the packet reaches PE-1, following operation is done

3.1) Generate a hash of the received packet involving the first 256 bytes of the packet.

3.2) Encrypt the hash with private key of PE-1.

3.3) Add this encrypted 128 bit below EOS label

3.4) Send the packet with VPN label and with encrypted digital signature below EOS label

4) When the packet reaches PE-2, following operation is done

4.1) Based on the VPN label, PE-2 knows that it needs to look for digital signature below EOS label

4.2) Remove 128 bit digital signature below EOS label

4.3) Calculate a hash after removing the digital signature, call it 'J'

4.4) Decrypt the received digital signature with public key of PE-1, Call it 'K'

4.5) compare the J and K. If they are equal, forward the packet else drops them

Note:

1) For both solutions, hash is calculated only before slapping VPN label at PE-1, ie all TTL update gets over by then

2) In case of Solution 2, to support ECMP case, we can add one nibble extra in front of digital signature based on IPV4/IPV6

2.3 Use Case:1

1) The above case talks about usage of this mechanism in VPN case

2.4 Use Case:2 (RSVP can use this during tunnel setup)

1) Head-end during tunnel setup can inform tail-end about "Secured-Label" during setup

2) We can use any of the above solutions for LSP setup as well.

2.5 Advantages

1) Any attackers packet would have to guess a 40 bit label in the case of Solution 1 and a 20 + 128-bit DS label in Solution 2.

2) Since we are using PKI, it is impossible for an attacker to create a packet with same semantics of PE, since he would have to guess the Algorithm UA(x) and the bitmask patten in Solution (1) and the PKI key as well. In Solution (2) he would have to guess the PKI key and the Algorithm UA(x).

3) Provides real security from attacker in the case of a man-in-the-middle attack say from the intervening network between the two providers.

5) The same mechanism could be used by RSVP for tunnel setup between Head-end and tail-end. Head-end during RSVP set-up can inform tail-end to use the "Secured-Label" mechanismor the DS mechanism in Solution 1 and Solution 2 respectively.

2.6 Limitations

1) An additional decryption and hashing is necessary in PE for secured labels in Solution 2 and in Solution 1 a bitmask lookup and selection is required over and above the decryption and hashing.

2) This mechanism will not work if this packet is fragmented inside the core.

3 Security Considerations

PKI is a secure mechanism as established in common security parlance. The control plane is secure in Option-C / Model-C in Inter-provider VPNs. It would take more than polynomial time complexity for an attacker to compromise the traffic using this mechanism.

4 IANA Considerations

IANA needs to assign the Type value for exchanging the additional details in the control plane as illustrated in the above two solutions.

5 References

5.1 Normative References

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

[RFC1776] Crocker, S., "The Address is the Message", RFC 1776, April 1 1995.

[TRUTHS] Callon, R., "The Twelve Networking Truths", RFC 1925, April 1 1996.

5.2 Informative References

[1] S. Alouneh, A. En-Nouaary and A. Agarwal, MPLS security: an approach for unicast and multicast environments, Annals of Telecommunications, Springer, vol. 64, no. 5, June 2009, pp. 391-400, doi:10.1007/s12243-009-0089-y.

[2] M. H. Behringer and M. J. Morrow, MPLS VPN security, Cisco Press, June 2005, ISBN-10: 1587051834.

[3] B. Daugherty and C. Metz, Multiprotocol Label Switching and IP, Part 1, MPLS VPNS over IP Tunnels, IEEE Internet Computing, May-June 2005, pp. 68-72, doi: 10.1109/MIC.2005.61.

[4] L. Fang, N. Bita, J. L. Le Roux and J. Miles, Interprovider IP-MPLS services: requirements,

implementations, and challenges, IEEE Communications Magazine, vol. 43, no. 6, June 2005, pp. 119-128, doi: 10.1109/MCOM.2005.1452840.

[5] C. Lin and W. Guowei, Security research of VPN technology based on MPLS, Proceedings of the Third International Symposium on Computer Science and Computational Technology (ISCSCT 10), August 2010, pp. 168-170, ISBN- 13:9789525726107.

[6] Y. Rekhter, B. Davie, E. Rosen, G. Swallow, D. Farinacci and D. Katz, Tag switching architecture overview, Proceedings of the IEEE, vol. 85, no. 12, December 1997, pp. 1973-1983, doi:10.1109/5.650179.

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

[8] T. H. Cormen, C. E. Leiserson, R. L. Rivest and C. Stein, Introduction to algorithms, 3rd edition, MIT Press, September 2009, ISBN-10:0262033844.

[9] C. Semeria, RFC 2547bis: BGP/MPLS VPN fundamentals, Juniper Networks white paper, March 2001.

[10] Advance MPLS VPN Security Tutorials [Online], Available:
http://etutorials.org/Networking/MPLS+VPN+security/
Part+II+Advanced+MPLS+VPN+Security+Issues, [Accessed: 10th December 2011]

[11] Inter-provider MPLS VPN models [Online], Available:
http://mpls-configuration-on-cisco-iossoftware.org.ua/1587051990/ ch07lev1sec4.html, [Accessed 10th December 2011]

[EVILBIT] Bellovin, S., "The Security Flag in the IPv4 Header", RFC 3514, April 1 2003.

[RFC5513] Farrel, A., "IANA Considerations for Three Letter Acronyms", RFC 5513, April 1 2009.

[RFC5514] Vyncke, E., "IPv6 over Social Networks", RFC 5514, April 1 2009.

Authors' Addresses

Bhargav Bhikkaji,
Dell-Force10,
350 Holger Way,
San Jose, CA 95134
U.S.A

EMail: BHARGAV_BHIKKAJI@dell.com

Balaji Venkat Venkataswami,
Dell-Force10,
Olympia Technology Park,
Fortius block, 7th & 8th Floor,
Plot No. 1, SIDCO Industrial Estate,
Guindy, Chennai - 600032.

EMail: BALAJI_VENKAT_VENKAT@dell.com