MPLS Working Group
Internet-Draft
Intended Status: Experimental RFC
Expires: February 2013

Shankar Raman
Balaji Venkat Venkataswami
Gaurav Raina
I.I.T Madras
Bhargav Bhikkaji
Dell-Force10

August 19, 2012

Avoiding Un-trusted AS thru inter-AS TE-LSPs constructed using Clipping

draft-balaji-mpls-inter-as-policy-based-te-sec-01

Abstract

For a short time sometime in the recent past , internet traffic sent between a well known site and subscribers to an internet service provider A passed through hardware belonging to a Telecom provider B other than the ISP A to which the customers were attached before reaching its final destination. Telecom Provider B was found to be many AS hops away from the well known site and ISP A. It was assumed that this was an an innocent routing error (which is the most likely explanation for the highly circuitous route that the traffic was taking), but it was troubling nonetheless. During a window that lasted 30 minutes to an hour, all unencrypted traffic passing between the victimised ISP's customers and the well known site might have been open to monitoring. Though there was no evidence any data was in fact snarfed, but it was felt that the potential for that is certainly there because the hardware belonged to the untrusted Telecom provider B.

Many such incidents have occurred in the past where the traffic has been diverted through such providers that either erroneously have let loose BGP routes or otherwise. At least one of those incidents was the result of erroneous BGP, or Border Gateway Protocol, routes that were quickly corrected. The above is a hypothetical headline that might occur in the near future if the BGP protocol is subject to such circuitous routing attacks either by mis-configuration or through purposeful intent. This is primarily owing to the fact that the BGP protocol accepts updates from providers and there exists no mechanism to figure out whether the updates for prefixes received was due to mal-intent, mis-configuration or indeed correct configuration. So there is a big blind spot that will have to be rectified. Doing the rectification through BGP would only complicate matters more.

The proposal in the scheme in this draft, warrants the use of MPLS-based inter-AS Traffic Engineered Label Switched Paths that are constructed out of a derived inter-AS topology that help to impose

policy decisions that for eg, obviate or prevent such LSPs from actually going through certain specific AS or set of ASes. Using methods like Graph construction from AS-PATH-INFO data and methods like policy based clipping of edges and nodes from such a inter-AS topology, the solution is made simple. The use of PCE (Path Computation Elements) is advised to compute such inter-AS paths that avoid ASes. Regular routing would have followed BGP updates and regular IP based forwarding. Using the TE-LSPs we can in fact set out the explicit route from AS to AS from the head-end to the tail-end avoiding specific set of ASes which dictated by policy have to be avoided.

Status of this Memo

Copyright and License Notice

Table of Contents

1 Introduction 4
1.1 Terminology 4
2. Methodology of the proposal 5
2.1 Pre-requisites for the Proposed Method 5
2.1.1 Constructing network topology using BGP strands 5
2.1.3 Explicit routing using TE-LSPs 6
2.1.4 Conclusion and Future work 7
2.1.5 ACKNOWLEDGEMENTS 8
3 Security Considerations 9
4 IANA Considerations 9
5 References 9
5.1 Normative References 9
5.2 Informative References 9
Authors' Addresses 11

1 Introduction

For a short time sometime in the recent past , internet traffic sent between a well known site and subscribers to an internet service provider A passed through hardware belonging to a Telecom provider B other than the ISP A to which the customers were attached before reaching its final destination. Telecom Provider B was found to be many AS hops away from the well known site and ISP A. It was assumed that this was an an innocent routing error (which is the most likely explanation for the highly circuitous route that the traffic was taking), but it was troubling nonetheless. During a window that lasted 30 minutes to an hour, all unencrypted traffic passing between the victimised ISP's customers and the well known site might have been open to monitoring. Though there was no evidence any data was in fact snarfed, but it was felt that the potential for that is certainly there because the hardware belonged to the untrusted Telecom provider B.

Many such incidents have occurred in the past where the traffic has been diverted through such providers that either erroneously have let loose BGP routes or otherwise. At least one of those incidents was the result of erroneous BGP, or Border Gateway Protocol, routes that were quickly corrected. The above is a hypothetical headline that might occur in the near future if the BGP protocol is subject to such circuitous routing attacks either by mis-configuration or through purposeful intent. This is primarily owing to the fact that the BGP protocol accepts updates from providers and there exists no mechanism to figure out whether the updates for prefixes received was due to mal-intent, mis-configuration or indeed correct configuration. So there is a big blind spot that will have to be rectified. Doing the rectification through BGP would only complicate matters more.

The proposal in the scheme in this draft, warrants the use of MPLS-based inter-AS Traffic Engineered Label Switched Paths that are constructed out of a derived inter-AS topology that help to impose policy decisions that for eg, obviate or prevent such LSPs from actually going through certain specific AS or set of ASes. Using methods like Graph construction from AS-PATH-INFO data and methods like policy based clipping of edges and nodes from such a inter-AS topology, the solution is made simple. The use of PCE (Path Computation Elements) is advised to compute such inter-AS paths that avoid ASes. Regular routing would have followed BGP updates and regular IP based forwarding. Using the TE-LSPs we can in fact set out the explicit route from AS to AS from the head-end to the tail-end avoiding specific set of ASes which dictated by policy have to be avoided.

1.1 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

This draft is an attempt to provide a solution. The following are the pre-requisites of the solution.

2.1 Pre-requisites for the Proposed Method

In this section we discuss the pre-requisites for the implementation of the proposed scheme.

2.1.1 Constructing network topology using BGP strands

The Inter-AS topology can be modeled as a directed graph G = (V; E; f) where the vertices (V) are mapped to AS and the edges (E) map the link that connect the neighboring AS. The direction (f) on the edge, represents the data flow from the head-end to the tail-end AS. To obtain the Inter-AS topology, the approach proposed in [5] is used. In this approach, it is shown that a sub-graph of the Internet topology, can be obtained by collecting several prefix updates in BGP. This is illustrated in Figure 1 which shows the different graph strands of AS that are recorded from the BGP packets. Figure 2 shows the strands merged together to form the topology sub-graph.

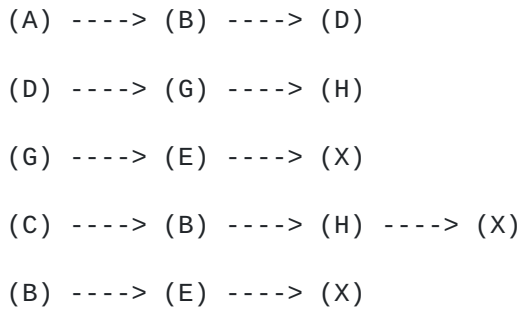

Figure 1: Different strands obtained from BGP updates, where vertices A,B,C,D and G represent the head-end AS. D,H and X form the tail-end AS. The direction of the link shows the next AS hop.

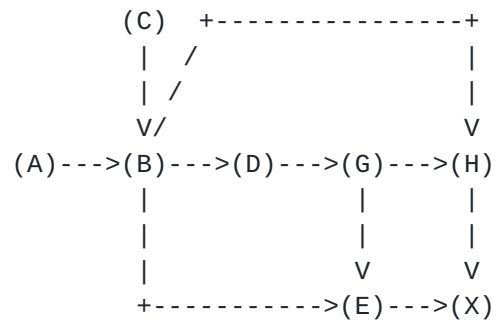

Figure 2:Combining the strands to get the topology of the Internet.

2.1.3 Explicit routing using TE-LSPs

We assume that the head-end and the tail-end may reside in different AS and the path is along multiple intervening AS. In our example the head-end is ISP providing services to its customers which is AS A and the tail-end is X which is the well known site AS, and AS H is the rogue AS that has to be avoided. The way to generate this path is by using Traffic Engineered Label Switched Paths (TE-LSPs). TE-LSPs can influence the exact path (at the AS level) that the traffic will pass through. This path can then be realized by providing these set of ASes to a protocol like Resource Reservation Protocol (RSVP). RSVP-TE then creates TE-LSPs or tunnels, using its label assigning procedure. The routers use these paths created by the explicit routing method rather than using the conventional shortest path to the destination. By this way, we can influence exclusion of a number of to-be-avoided-ASes on the way from the head-end to the tail-end AS. For example, the dotted line in Figure 5 represents the explicit route that is chosen by making use of such TE-LSPs from head-end AS A to the tail-end AS X. Note that if number of hop was the metric used by CSPF, then the route chosen is the path with 3 hops. Here the AS to be avoided is the AS H. In order to exclude the possibility of any traffic passing through H the policy is applied at the time of path computation to exclude all links to and from node H and the AS H itself. This can be used by clipping the to be excluded AS by clipping links to and from it, in this case H.

The prefixes in X and behind X need to be advertised as reachable through the TE-LSP so constructed. This way the traffic goes through trusted ASes and not into territory of ASes that are rogue and have an intent to snarf or eavesdrop on the data encrypted or non-encrypted.

The clipped topology is shown in the figure below and the path constructed after excluding AS H is shown in the figure after the one below.

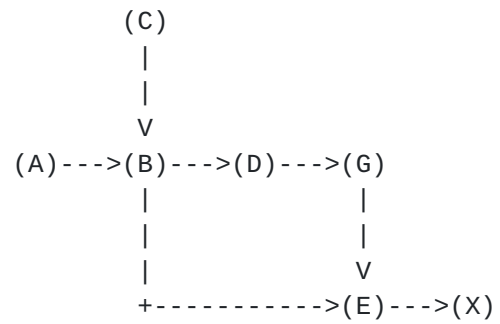

Figure 5.1: Clipped Graph excluding AS H.

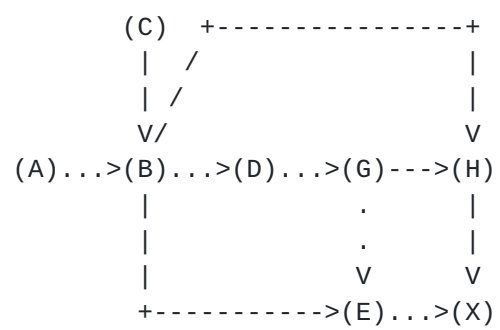

Figure 5.2:The final path is represented by the dotted lines. This path has a longer number of hops than the conventional shortest path but it avoids AS H.

It is also possible through this scheme to cut out a set of ASes rather than just one AS. Thus a gaping hole in the topology might result thus excluding one or more ASes from being considered while considering the Inter-AS TE-LSPs. This can be easily done by clipping all links to and from the graph to these set of ASes and eliminating the ASes altogether from consideration if they are not trusted by the Path Computation Element (PCE) of the TE-LSP initiating provider or AS.

This process of setting up inter-AS TE-LSPs that are passing through trusted (so called) ASes can be selectively done only for traffic heading to tail-end ASes which may be ISPs for the well-known sites or the well-known sites themselves (assuming they have an AS of their own). Such selective tunneling would take care of scalability concerns at the provider initiating these tunnels (head-ends).

2.1.4 Conclusion and Future work

Avoiding ASes and their associated links that should not be traversed towards needs be considered. One method that can be used is constructing inter-AS TE LSPs with or without bandwidth reservation to and from a head-end and a tail-end avoiding certain ASes which are

explicitly specified. Other methods are also being investigated which will be specified in due course in an updated version of this document. we show only one direction of traffic. The other direction may be suitably constructed by the tail-end AS becoming the head-end for such purposes.

2.1.5 ACKNOWLEDGEMENTS

The authors would like to acknowledge the UK EP-SRC Digital Economy Programme and the Government of India Department of Science and Technology (DST) for funding given to the IU-ATC.

3 Security Considerations

The verification of the BGP prefixes and their corresponding attributes such as AS-PATH-INFO is to be considered. Existing mechanisms can be used to verify these attributes. While constructing the strands a technique to verify these actually are indeed the strands of ASes that are traversed can be done by using a majority vote technique from among a set of strands that cover a part of the topology. In order to mount an attack with dis-information through invalid strands, the rogue AS or set of ASes have to miscommunicate the strand (AS-PATH-INFO) data in such a way that such majority vote goes in their favour.

Considering the existing BGP path selection algorithm in place today, the technique that we propose in this paper is a lot more resilient and doesnt ply itself to listening to BGP updates without some sort of verification.

4 IANA Considerations

None.

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]

[12] Davari.S et.al, Transporting PTP messages (1588) over MPLS networks, "http://datatracker.ietf.org/doc/draft-ietf-tictoc-1588overmpls/?include_text=1", Work in Progress, October 2011.

Authors' Addresses

Shankar Raman
Department of Computer Science and Engineering
I.I.T Madras,
Chennai - 600036
TamilNadu,
India.

EMail: mjsraman@cse.iitm.ac.in

Balaji Venkat Venkataswami
Department of Electrical Engineering,
I.I.T Madras,
Chennai - 600036,
TamilNadu,
India.

EMail: balajivenkat299@gmail.com

Prof.Gaurav Raina
Department of Electrical Engineering,
I.I.T Madras,
Chennai - 600036,
TamilNadu,
India.

EMail: gaurav@ee.iitm.ac.in

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

Email: Bhargav_Bhikkaji@dell.com