Mitigating IP Spoofing by Validating BGP Routes Updates

Junaid Israr, Mouhcine Guennoun, and Hussein T. Mouftah

School of Information Technology and Engineering
University of Ottawa
800 King Edward, Ottawa, ON, Canada

Summary
IP spoofing remains a popular method to launch Distributed Denial of Service (DDOS) attacks. Several mitigation schemes have been proposed in literature to detect forged source IP addresses. Some of these solutions, like the inter domain packet filter (IDPF), construct filters based on implicit information contained in BGP route updates. The packet filters rely on the fact that BGP updates are valid and reliable. This assumption is unfortunately not true in the context of the Internet. In addition, attackers can combine control and data plane attacks to avoid detection. In this paper, we evaluate the impact of false and bogus BGP updates on the performance of packet filters. We introduce a new and easy to deploy extension to the standard BGP selection algorithm in order to detect spoofed BGP updates. The new proposal, credible BGP (CBGP), assigns credibility scores for AS prefix origination and AS path. These credibility scores are used in an extended selection algorithm to prefer valid BGP routes. Based on simulation studies, we prove that the proposed algorithm improves significantly the performance of packet filters based on BGP updates.

Key words:
BGP, IDPF, IP Spoofing.

1. Introduction

The lack of source IP address validation across multiple Autonomous Systems (ASs) in the internet makes it difficult to detect and prevent attackers from launching Distributed Denial of Service (DDoS) attacks using spoofed source IP addresses. Several popular internet sites [1] and internet infrastructure [2] have been attacked recently and such attacks have the potential to cripple the internet. Detection and prevention of these attacks is often made more complicated by attackers employing source IP address spoofing. The idea is to forge the source IP address in the "attack" packets to that of another host in the system. This allows the attacker to pose as some other host and hide its actual identity and location, making it difficult to detect the actual attacker and to protect against it. As a result, attack detection techniques that rely on source address-based filtering become less effective when source address is spoofed by the attackers.

There are several reasons why source IP address spoofing remains a popular method to launch attacks in the Internet [7]. First, when an attack is launched using source IP address spoofing, it is difficult to differentiate attack traffic from legitimate traffic. The host whose IP address has been hijacked may well be sending legitimate traffic at the same time as attack traffic is being sent from its IP address. Second, although the attack appears to be coming from a particular victim host (whose source IP address has been hijacked), it can take substantial amount of time and resources to determine that the host itself is a victim and that the true attacker still needs to be located [8,9,10]. Finally, forging of source IP addresses allows the attacker to pose as a valid host on the other end of a transaction and launch popular man-in-the-middle attacks, such as variants of TCP hijack and DNS poisoning attacks [11, 12]. Similarly, IP spoofing can be used to launch reflector-based attacks whereby an attacker uses some victim's IP address to contacts a number of hosts, resulting in the victim being flooded by replies from all these hosts [13]. These factors indicate that IP spoofing is unlikely to decrease in the near future.
Many solutions have been proposed to detect IP spoofing. Most of them are based on filtering packets based on the IP source address and the incoming interface. Indeed, if the source IP address of the packet is not expected to be received on the incoming interface then the packet is dropped. Two schemes are worth to mention: route based packet filter and inter domain packet filter (IDPF).rom this section, input the body of your manuscript according to the constitution that you had. For detailed information for authors, please refer to [1].

2. IP Spoofing Detection Techniques

The route based packet filter proposed by Park and Lee [3] relies on the basic fact that if a single-path routing scheme is assumed, there is exactly one single path p(s, d) between source node s and destination node d. Therefore, any packet with source address s and destination address d that appear in a router not in p(s, d) should be discarded. However, in order to construct a specific route-based packet filter at a node, it requires knowledge of global

Manuscript received May 5, 2009
Manuscript revised May 20, 2009

routing decisions made by all the other nodes in the network. This is impossible with the current BGP-based Internet routing infrastructure. The current Internet topology consists of more than 35,000 network domains or autonomous systems (ASs), each of which is a logical collection of networks with common administrative control. Each AS communicates with its neighbors using the Border Gateway Protocol (BGP), the de-facto inter-domain routing protocol, to exchange network layer information reachability about its own networks and others that it can reach. BGP is a *policy-based* routing protocol in that both the selection and the propagation of the best route to reach a destination at an AS are guided by some locally defined routing policies. Given the insular nature of how policies are applied at individual ASs, it is impossible for an AS to acquire the complete knowledge of routing decisions made by the other entire ASs. Hence constructing route-based packet filters as proposed in [3] is an open challenge in the current Internet routing regime.

The IDPF architecture takes advantage of the fact that while network connectivity may imply a large number of potential paths between source and destination domains, commercial relationships between ASs act to restrict to a much smaller set the number of feasible paths that can be used to carry traffic from the source to the destination [4]. IDPFs are constructed from the information implicit in BGP route updates and are deployed in network border routers. When a node receives a packet from an incoming interface, it checks if the source IP address has been advertised through this interface. The packet is discarded if the check is negative. A key feature of the scheme is that it does not require global routing information. The simulation results [4] showed that, even with partial deployment on the Internet, IDPFs can significantly limit the spoofing capability of attackers; moreover, they also help localize the actual origin of an attack packet to be within a small number of candidate networks. In addition, IDPFs also provide adequate local incentives for network operators to deploy them.

3. Security Concerns in BGP

BGP network design was undertaken in the relatively homogenous and mutually trusting environment of the early Internet. The underlying distributed distance vector computations rely heavily on informal trust models associated with information propagation to produce reliable and correct results. It can be likened to a hearsay network: information is flooded across a network as a series of point-to-point exchanges, with the information being incrementally modified each time it is exchanged between BGP speakers [14]. The approach to information exchange was not primarily designed for robustness in the face of various forms of negotiated trust or overt hostility on the part of some routing nodes in the network.

BGP has several well-known vulnerabilities. These vulnerabilities are the direct consequences of three fundamental weaknesses in the BGP and the inter-domain routing environment [5]. The first weakness is there is no mechanism to check the integrity, freshness and source authenticity of BGP messages. Also, BGP doesn't offer any mechanism to verify the authenticity of an address prefix and an AS origination of this prefix in the routing system. Last, the BGP protocol doesn't provide any way to guarantee that the attributes of a BGP UPDATE message are correct.

The lack of security concepts in BGP leaves it vulnerable to several types of control plane attacks. In addition, the IDPF scheme, which relies on BGP updates to detect and prevent source IP address spoofing, will fail if the BGP updates are not correct. The IDPF scheme assumes that BGP routing updates are secure and hence trustworthy. However, by accepting bogus BGP updates, the IDPF filters become less effective. The performance of IDPF scheme suffers when hostile nodes, which can generate non-trustable BGP updates and hence create incorrect filters, are introduced in the network (see section 5). This decline in IDPF performance can be arrested by deploying a mechanism to secure BGP. At present there are a number of practical and a number of more fundamental questions relating to securing BGP. The first is a practical question relating to the inevitable design trade-off between the level of security and the performance overheads of processing security credentials associated with BGP UPDATE messages. It is not entirely known as to what aspects of BGP performance and load are critical for the robust operation of network applications and what are not so critical. With such considerations, it is extremely important that any solution to secure BGP should try and minimize impact on current performance of BGP and should be incrementally deployable. Given this, there is a strong incentive to alter BGP such that it will provide reasonable amount of security at both control plane and data planes and will have minimal impact on BGP messaging.

4. Credible BGP

Credible BGP calls for a modification to the standard BGP route selection algorithm such that it takes into account validity state of routing updates. We define the validity state factor as the minimum of two independent scores, route origination validation score and update AS path validation score. These two scores are defined as follows:

Route origination validation score is derived based on the ability of a route receiving node to determine whether the AS originating the route actually is authorized to do so.
Route AS-Path validation score is derived based on the ability to which the node is able to determine whether the received update actually traversed the ASs listed in the AS Path.

The BGP decision process will then be modified to check the validity state of each routing update when comparing two routing updates for routing selection purposes. The validity state check must be performed before any of other prior to any of the steps defined in the decision process of [6]. The route with the highest validity state will always be preferred over other routes. In all other respects, the BGP decision process remains unchanged.

In the light of proposed changes to BGP selection algorithm, we propose to investigate its impact on the performance of IDPF scheme. The performance measurements will be analyzed to demonstrate the impact of the proposed changes on the performance of IDPF scheme when an increasing percentage of untrusted BGP routing updates are introduced in the network. The proposed changes to BGP decision process will help prevent control plane attacks in BGP networks. If an untrusted route update is accepted in the network, it can lead to black-holing of traffic. The proposed scheme will guard against such control plane attacks and will make untrusted BGP updates less acceptable.

5. Performance Metrics

False and bogus BGP updates have a significant impact on the performance of packet filters such as IDPF filter. In order to evaluate the impact of the proposed scheme on the performance of packet filters we introduce new performance metrics based on predefined set of metrics described in [3].

We define three metrics to measure the strength of the deployed solution to prevent IP spoofing attacks**.** Given the AS graph G= (V, E), we will use F to denote the sub-set F $\subseteq$ V of nodes where the new enhanced security scheme is deployed. We call $\mu=\frac{|F|}{|V|}$ the coverage ratio. We also define r as the ratio of spoofed BGP updates.

$S_{a,t}$ denotes [3] the set of nodes—more precisely, the set of IP addresses belonging to an AS node in *Sa,t*—that an attacker at AS *a* can use as spoofed source IP addresses to reach *t* without being cut-off by filters executed at autonomous systems in *T*. The larger the set *Sa,t*, the more options an attacker at *a* has in terms of forging the IP source address field with a bogus address which will go undetected. Whereas *Sa,t* is defined from the attacker's perspective, *Cs,t* captures the victim's perspective and denotes the set of nodes that could have sent an IP packet *M*(*s, t*) with spoofed source IP address *s* and destination address *t* which did not get filtered on its way. The larger *Cs,t*, the more uncertain the victim at *t* is upon receiving spoofed packet *M*(*s, t*) with respect to its true origin. If |*Cs,t*| = 1, then this means that IP address *s* cannot be used by any attacker *a* (outside of *s* itself) to mount a spoofed DoS attack aimed at *t*.

Park and Lee [3] defined three metrics to measure the strength and effectiveness of IDPF filters in limiting IP spoofing. These metrics are VictimFraction Φ, AttackFraction θ and VictimTraceFraction Ψ.

VictimFraction(τ) denotes the fraction of ASes that can be attacked with packets from at most τ ASes. Particularly, VictimFraction(1) is the fraction of ASes that are not vulnerable to IP spoofing attack. Φ is defined as:

$$\Phi(\tau)=\frac{|\{t:\forall a\in V,|S_{a,t}|<\tau\}|}{|V|} \quad (1)$$

We define Φ_μ as:

$$\Phi_\mu(r)=\frac{1}{|V|}\int_1^{|V|}\Phi_r(\tau)\mathrm{d}\tau \quad (2)$$

Φ_μ calculates the performance of IDPF in limiting the number of victims of IP spoofing in the presence of a percentage r of spoofed BGP routing updates in the network. μ is the ratio of ASes where the CBGP is deployed.

We define metric Ω to measure the average performance of IDPF in the presence of a variable rate of spoofed BGP updates. Ω is expressed as:

$$\Omega(\mu)=\int_0^1\Phi_\mu(r)dr \quad (3)$$

The enhancement of the effectiveness of IDPF in protecting ASes against spoofing based DDoS attacks is expressed as:

$$\lambda_1(\mu)=\frac{\Omega(\mu)}{\Omega(0)} \quad (4)$$

Similarly, AttackFraction denotes the fraction of ASes from which an attacker can forge addresses belonging to at most τ ASes (including the attacker's own), in attacking any other ASes in the network. Particularly, AttackFraction(1) is the fraction of ASes from which an attacker cannot spoof the IP address of any other AS to attack the network. $\theta(\tau)$ is defined in [3] as:

$$\theta(\tau) = \frac{|\{a : \forall t \in V, |S_{a,t}| < \tau\}|}{|V|} \quad (5)$$

We define θ_μ as:

$$\theta_\mu(r) = \frac{1}{|V|}\int_1^{|V|} \theta(\tau)d\tau \quad (6)$$

θ_μ calculates the strength of IDPF in limiting the number of attackers in the presence of a percentage r of spoofed BGP routing updates in the network.

We define α as a metric to measure the average strength of IDPF filters in protecting the network against attackers. α is expressed as:

$$\alpha(\mu) = \int_0^1 \theta_\mu(r)dr \quad (7)$$

The enhancement of the strength of IDPF filter in limiting the spoofing capability of an arbitrary attacker is expressed as:

$$\lambda_1(\mu) = \frac{\Omega(\mu)}{\Omega(0)} \quad (8)$$

Last, the authors of [3] define a reactive metric Ψ that measures the effectiveness of IDPF in reducing the IP trace back effort, i.e., the act of determining the true origin of spoofed packets. Ψ is defined as:

$$\Psi(\tau) = \frac{|\{t : \forall s \in V, |C_{s,t}| \leq \tau\}|}{|V|} \quad (9)$$

We define δ_μ to measure the effectiveness of IDPF filters in determining the true origin of spoofed packets in the presence of a percentage r of spoofed BGP routing updates. δ_μ is expressed as:

$$\delta_\mu(r) = \frac{1}{|V|}\int_1^{|V|} \Psi(\tau)d\tau \quad (10)$$

The average effectiveness β is defined as:

$$\beta(\mu) = \int_0^1 \delta_\mu(r)dr \quad (11)$$

The enhancement of the effectiveness of IDPF filter in determining the true origin of IP spoofing attacks is expressed as:

$$\lambda_3(\mu) = \frac{\beta(\mu)}{\beta(0)} \quad (12)$$

6. Performance of IDPF filters in the presence of BGP updates spoofing

In this section, we demonstrate how the performance of IDPF scheme declines in the face of growing number of bogus and false updates in the network. The graphs in Figure 1 to 3 demonstrate the progression of decline in the performance of IDPF scheme when an increasing percentage of untrusted BGP routing updates are introduced in the network.

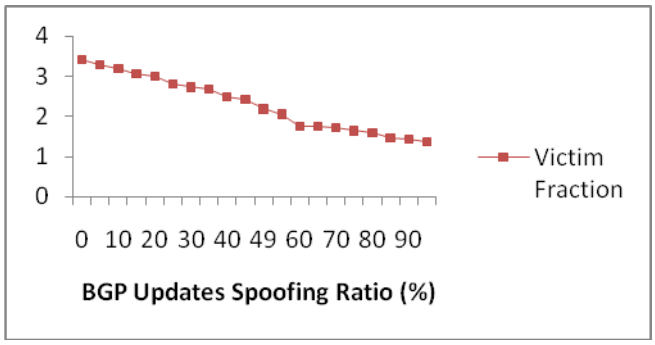

Fig. 1: Degradation of Victim Fraction Performance

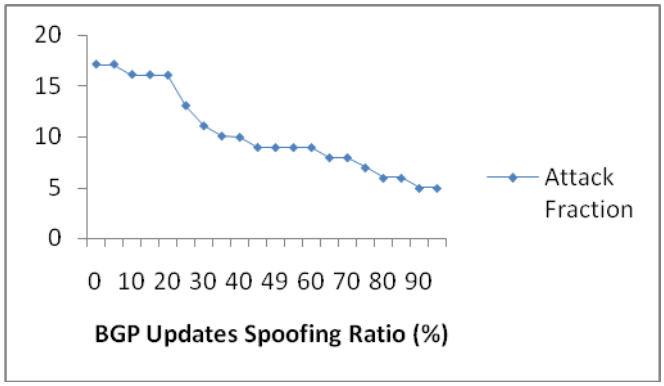

Fig. 2: Degradation of Attack Fraction Performance

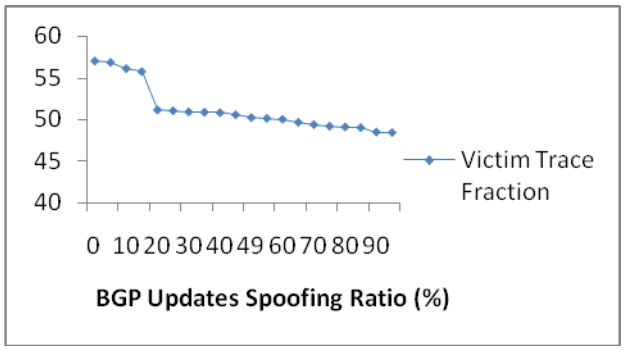

Fig. 3: Degradation of Victim Trace Fraction

The simulation results clearly demonstrate that there is a significant impact on the performance of IDPF filters when bogus and false updates are present in the network. Using this vulnerability in the IDPF scheme, attackers can combine both control plane and data plane to escape detection. It's obvious that the mitigation of IP spoofing attacks should be addressed on the control plane as well. In the next section, we will measure the enhancement of IDPF filters when Credible BGP is deployed in the network.

7. Performance Gain Ratio of IDPF filters

Results from the previous section showed that there is a need to validate the BGP updates in order to ensure proper functioning of the IDPF filters. We have deployed CBGP increasingly in the network and measured the enhancement of the strength and effectiveness of IDPF filters. The new metrics λ_1, λ_2 and λ_3 measure the overall performance enhancement of VictimFraction, AttackFraction and VictimTraceFraction metrics respectively. Figure 4 shows the simulation results with a coverage ratio μ that varies from 0 to 100%.

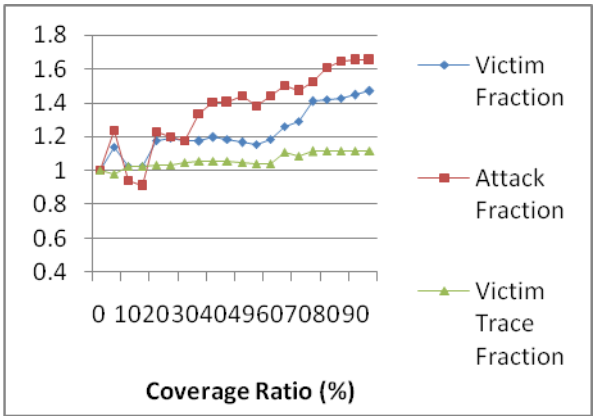

Fig. 4: Performance Gain Ratio of IDPF Filters

8. Conclusion and Future Work

In this paper we proposed an easy to deploy protocol to validate BGP routing updates. CBGP modifies the current BGP selection algorithm by adding an extra check of the validity of the origin IP prefix and the AS path. We believe that CBGP can be incrementally deployed in the Internet network without having an impact on the existing BGP infrastructure such as BGP messaging system. We proved using simulation studies that the performance of packet filters based on BGP updates is improved when CBGP is deployed in the network. In the future, we are planning to investigate the overhead and cost associated with the deployment of CBGP protocol on the current Internet infrastructure. It would be interesting to determine the impact of the proposed change in BGP selection algorithm on the control plane load in the network. Since the proposed validity state factor will override other criteria for BGP decision process, the network routing table with the validity state factor considered will appear very different from when the validity state factor is not considered.

Acknowledgments

This research was funded by a URP grant from Cisco Systems. The authors would like to thank David Ward of Cisco Systems for his support.

References

[1] Yahoo attributes a lengthy service failure to an attack. http://www.nytimes.com/library/tech/00/02/biztech/articles/08yahoo.html, February 2000.

[2] Massive DDoS attack hit DNS root servers. http://www.internetnews.com/entnews/article.php/1486981, October 2002.

[3] K. Park and H. Lee, On the effectiveness of route-based packet filtering for distributed DoS attack prevention in power-law internets, In Proc. ACM SIGCOMM, San Diego, CA, August 2001.

[4] Z. Duan et al., Constructing Inter-Domain Packet Filters to Control IP Spoofing Based on BGP Updates, http://www.cs.fsu.edu/~xyuan/paper/06infocom.pdf, 2006.

[5] S. Murphy, BGP Security Vulnerabilities Analysis, Internet Draft, draft-murphy-bgpvuln-02.txt, March 2003.

[6] Y. Rekhter, T. Li, and S. Hares, A Border Gateway Protocol 4 (BGP-4), RFC 4271, Internet Engineering Task Force, January 2006.

[7] D. Moore, C. Shannon, D. Brown, G. Voelker, and S. Savage, Inferring internet Denial-of-Service activity, ACM Transactions on Computer Systems, vol. 24, no. 2, May 2006

[8] S. Bellovin et al., ICMP Traceback Messages, February 2003, http://www3.ietf.org/proceedings/03mar/I-D/draft-ietf-itrace-04.txt

[9] S. Savage, D. Wetherall, A. Karlin, and T. Anderson, Practical network support for IP traceback, In SIGCOMM, pages 295.306, 2000.
[10] A. Snoeren, C. Partridge, L. Sanchez, C. Jones, F. Tchakountio, S. Kent, and W. Strayer, Hash-based IP traceback, In Proc. ACM SIGCOMM, 2001.
[11] M. Dalal, Improving TCP's robustness to blind in-window attacks, Internet Draft, May 2005, Work in Progress.
[12] J. Stewart, DNS cache poisoning - the next generation, Technical report, LURHQ, January 2003.
[13] V. Paxson, An analysis of using reflectors for distributed denial-of-service attacks, ACM Computer Communications Review (CCR), 31(3), July 2001.
[14] S. Murphy, BGP Security Analysis, Internet Draft, draft-murphy-bgp-secr-03.txt, June 1999.

Junaid Israr is a PhD candidate student at University of Ottawa, Canada. He has worked with Cisco systems as a software development manager in the IOS-XR MPLS Traffic Engineering (TE) and Resource Reservation Protocol (RSVP) group, and was responsible for all aspects of RSVP Traffic Engineering protocols. He completed his M. Eng in 1999 and B. Eng in 1996 from Carleton University, Ottawa, Canada.

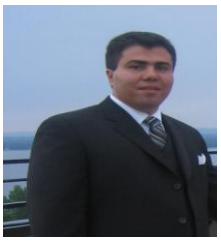

Mouhcine Guennoun graduated from the University of Ottawa, Canada and the Ecole Mohammadia d'Ingénieurs, Rabat, Maroc. He is currently a Research Assistant at the University of Oattawa. His research interests include BGP security, wireless security and detection of intrusions in Ad-Hoc wireless networks.

Dr. Hussein T. Mouftah joined the School of Information Technology and Engineering (SITE) of the University of Ottawa in 2002 as a Tier 1 Canada Research Chair Professor, where he became a *University Distinguished Professor* in 2006. He has been with the ECE Dept. at Queen's University (1979-2002), where he was prior to his departure a Full Professor and the Department Associate Head. He has six years of industrial experience mainly at Bell Northern Research of Ottawa (now Nortel Networks). He served as Editor-in-Chief of the IEEE Communications Magazine (1995-97) and IEEE ComSoc Director of Magazines (1998-99), Chair of the Awards Committee (2002-03), Director of Education (2006-07), and Member of the Board of Governors (1997-99 and 2006-07). He has been a Distinguished Speaker of the IEEE Communications Society (2000-2007). He is the author or coauthor of 6 books, 30 book chapters and more than 850 technical papers, 10 patents and 138 industrial reports. He is the joint holder of 8 Best Paper and/or Outstanding Paper Awards. He has received numerous prestigious awards, such as the 2007 Royal Society of Canada Thomas W. Eadie Medal, the 2007-2008 University of Ottawa Award for Excellence in Research, the 2008 ORION Leadership Award of Merit, the 2006 IEEE Canada McNaughton Gold Medal, the 2006 EIC Julian Smith Medal, the 2004 IEEE ComSoc Edwin Howard Armstrong Achievement Award, the 2004 George S. Glinski Award for Excellence in Research of the U of O Faculty of Engineering, the 1989 Engineering Medal for Research and Development of the Association of Professional Engineers of Ontario (PEO), and the Ontario Distinguished Researcher Award of the Ontario Innovation Trust. Dr. Mouftah is a Fellow of the IEEE (1990), the Canadian Academy of Engineering (2003), the Engineering Institute of Canada (2005) and the RSC Academy of Science (2008).