www.ijecs.in
International Journal Of Engineering And Computer Science ISSN:2319-7242
Volume2 Issue 8 August, 2013 Page No. 2565-2568

Internet Routing With Lightweight Route Attestation

M.V.R Jyothisree [*] V.Ramakrishna[2] Dr.A.V.Krishna Prasad [3]

[*] M.Tech Student, CVSR College of Engineering, Department of Computer Science, A.P. India

jyothisree.manne@gmail.com

[2] Assistant Professor of Computer Science and Engineering, Anurag group of Institutions, A.P. India

mekrishnait2111@gmail.com

[3] Associate Professor, Department of IT, MVSR Engineering College, A.P. India

kpvambati_it@mvsrec.edu.in

ABSTRACT

The Internet routing system plays an essential role of gluing together tens of thousands of individual networks to create a global data delivery substrate. Over the years many efforts have been devoted to securing the routing system, the weak trust model in Border Gateway Protocol (BGP) introduces severe vulnerabilities for Internet routing including active malicious attacks and unintended misconfigurations. Although various secure BGP solutions have been proposed, the complexity of security enforcement and data-plane attacks still remain open problems. We propose TBGP, a trusted BGP scheme aiming to achieve high authenticity of Internet routing with a simple and lightweight attestation mechanism. TBGP introduces a set of route update and withdrawal rules that, if correctly enforced by each router, can guarantee the authenticity and integrity of route information that is announced to other routers in the Internet. To verify this enforcement, an attestation service running on each router provides interfaces for a neighbouring router to challenge the integrity of its routing stack, enforced rules, and the attestation service itself. If this attestation succeeds, the neighbouring router updates its routing table or announces the route to its neighbours, following the same rules. Thus, a router on a routing path only needs to verify one neighbour's routing status to ensure that the route information is valid. Through this, TBGP builds a transitive trust relationship among all routers on a routing path. We implement a prototype of TBGP to investigate its practicality.

Keywords: Internet Routing, Border Gate way Protocol, Security Attacks

I. INTRODUCTION

As the Internet penetrates into every corner of the human society ranging from daily life information search to transactions in financial sector to management of critical infrastructure such as power supply systems, securing the global routing system also becomes of paramount importance. All applications of the Internet depend on the reliable functioning of the routing system to deliver their data to the right destinations. A routing system failure can lead to the failure of all applications, and a routing fault can result in denial of services to applications, or even compromises of applications security. BGP is a protocol based on trust that does not authenticate route update messages. The Border Gateway Protocol (BGP) is the only widely deployed inter domain routing protocol connecting different IP networks or autonomous systems (ASes) to construct the whole Internet [5]. In ordinary BGP, every AS announces its route information with different prefixes. However, its neighbouring ASes cannot validate this route information, but rather directly propagate it across the Internet. Obviously, this weak trust model allows forged route announcement propagations, which is a fundamental security weakness of BGP. Forged routes,

which can be generated by configuration errors or malicious attacks, can cause large-scale network connectivity problems. The situation could be worse if forged routes are generated by remote attacks [3].

In order to effectively eliminate false announcements and improve the security of BGP, several security-enhanced BGP solutions have been proposed. They generally can be classified into two categories: cryptography-based prevention and anomaly detection. Cryptographic approaches, such as SBGP and SoBGP, use a centralized routing registration authority and public key infrastructure (PKI) to ensure the authentication of routing announcements. These solutions are not sufficient to prevent data-plane attacks, where an AS can announce a route not adopted by itself [12]. Moreover, they usually consume a significant amount of extra router resources including computation and storage, and exacerbate the routing convergence performance. It is obvious that pure cryptography-based solutions are not cost-efficient to defend against routing attacks, and this impedes their deployment on the Internet. On the other hand, anomaly detection approaches aim to discover underlying hijacks in BGP announcements, e.g., by comparing BGP announcements with out-of-band information and querying third-party routing services [10]. However, most of the anomaly detection solutions raise false positives and require network operators to take actions in order to block detected anomalous routes [9]–[11].

II. IMPLEMENTATIONS

We propose a trusted BGP scheme called TBGP, which aims to *use minimal computation cost to achieve BGP security goals*. Unlike existing cryptography-based approaches, we do not solely rely on cryptography mechanisms to secure routing. Instead, we propose a set of well-defined route update and withdrawal rules that are enforced by the filters of each BGP router along a routing path. These rules guarantee that route announcements comply with the BGP specification [1]. Thus, the enforcement of these rules provides automatic route authenticity in each router and prevents the spread of forged routes over the Internet.

Our main goal in this paper is to show that FIs are no more vulnerable than traditional communication networks (such as IP networks) that do not export control on forwarding.

1. Forward Infrastructure (FI) achieves certain specific security properties, the essential features and efficiency for Network Path and Data Router.
2. Our main defense technique, which is based on light-weight cryptographic constraints on forwarding entries, prevents several attacks including eavesdropping, loops, and traffic amplification.
3. TBGP a trusted BGP scheme aiming to achieve high authenticity of
4. Internet routing with a simple and lightweight attestation mechanism.
5. From earlier work, we leverage some techniques, such as challenge-responses and erasure-coding, to other attacks.
6. Asymmetric to construct a consistent view of the network topology to secure the Network Path.

BGP is essential to the operation of the Internet, but is vulnerable to both accidental failures and malicious attacks. We propose a new protocol that works in concert with BGP, which Autonomous Systems will use to help detect and mitigate accidentally or maliciously introduced faulty routing information. The protocol differs from previous efforts at securing BGP in that it is receiver-driven, meaning that there is a mechanism for recipients of BGP UPDATE messages to corroborate the information they receive and to provide feedback. We argue that our new protocol can be adopted incrementally, and we show that there is incentive for network operators to do so. We also describe our prototype implementation. There are tens of routing protocols; they can be broadly split into two categories: *intra* domain, or internal, routing protocols, and *inter* domain, or external, routing protocols. Organizations under cohesive administrative control (companies, universities, Internet service providers) use intra domain routing protocols to exchange information about how to reach machines within their own purview. Inter-domain routing protocols are used to exchange and propagate reachability information *between* such organizations. This split reflects the coarse structure of the Internet: many networks connected to each other. It also reflects the different needs and requirements for routing protocols for use in intra- versus inter domain routing. While there are several internal routing protocols in use today, there is only one inter domain routing protocol: the Border Gateway Protocol (BGP) [8,9].

BGP views the Internet as a collection of interconnected *Autonomous Systems*. An Autonomous System (AS) is a portion of the network under single administrative control (at least as far as routing is concerned). Each AS connects to other ASes; the routers in each AS that connect to their counterpart in other ASes are called *border* routers. These neighbouring border routers connect *directly* to each other, that is, there are no routers between them. (This is not strictly true, nor is the assertion that only neighbouring routers speak BGP to each other, but the details are beyond the scope of this paper.) Over this direct connection, border routers establish *BGP sessions*; there may be many BGP sessions over each link, but there are (almost) never BGP sessions between non-neighbouring routers. BGP sessions are used to exchange network reachability information— each router tells its neighbour what address ranges (also known as address prefixes, or just prefixes) it knows how to route to, along with ancillary information that issued to make the decision of whether this router will actually be used to route that part of the address space.

As BGP provides information for controlling the flow of packets between ASes, the protocol plays a critical role in Internet efficiency, reliability, and security. The Internet can be severely impacted by BGP failures. Accidental misconfigurations have resulted in serious routing problems and loss of service [13]. However, failures are not always

accidental—attacks intended to cause widespread outage on the Internet will (and do) target BGP [15, 16].
Denial of service is not the only concern; an attacker might redirect the flow of some traffic through his network so that he can eavesdrop on it. BGP has several well-known vulnerabilities. Neither the originating announcement of a route, nor the information attached to it as it traverses ASes is guaranteed to be correct. Moreover, BGP does not provide any way of identifying the source of bad data. Hence, misconfigured or malicious routers can, among others things, force other ASes to accept bad or inefficient routes, hijack address ranges, or simply flood the network with useless route information. The security limitations of BGP are compounded by the fact that the protocol itself does not always converge [12,17]. Because BGP is potentially unstable at any time, it is particularly difficult to analyze. Complexity is always at odds with security. Getting the routing system to work at an acceptable level has taken huge effort in terms of designing, implementing, and deploying protocols. Moreover, as the nature of the Internet changes, these protocols have been required to provide functionality not originally envisioned. It comes as no surprise that security has not been the first priority of designers, implementers, or even operators; it is this lack of security that makes the routing system, and hence the entire Internet, susceptible to an increasing number of both accidental failures and malicious attacks.

Attacks against Internet routing are increasing in number and severity. Contributing greatly to these attacks is the absence of origin authentication: there is no way to validate claims of address ownership or location. The lack of such services enables not only attacks by malicious entities, but indirectly allows seemingly inconsequential misconfigurations to disrupt large portions of the Internet. Consider the semantics, design, and costs of origin authentication in interdomain routing. We formalize the semantics of address delegation and use on the Internet, and develop and characterize broad classes of origin authentication proof systems. Routing in the Internet dictates the path that IP packets take to get from their source to their destination. In its most general form, this path, called the route, is a sequence of routers and the links between them. To compute such paths, routers use a routing protocol to exchange reachability data, and perform computations on these data to compute the desired routes. The Border Gateway Protocol is the interdomain routing protocol used on the Internet. BGP routing domains, called Autonomous Systems (ASes) announce IP address ranges called prefixes to its neighbouring ASes. Each AS also announces the pre-fixes that it learns from each of its neighbours to its other neighbours. The design of BGP reflects its egalitarian origins: ASes are trusted to behave per specification and to perform due diligence in providing timely and accurate routing information. In other words, BGP does not currently provide security. The need for security in interdomain routing has been widely acknowledged and evaluated, and interim and long-term solutions are seeking broad adoption.

BIND: A Fine-grained Attestation Service for Secure Distributed Systems[18].

The term BIND is also used in Domain Name Service (DNS) terminology to stand for the Berkeley Internet Name Daemon.

Securing distributed systems continues to be an important research challenge. One hard problem in securing a distributed system arises from the fact that a remote software platform may be compromised and running malicious code.

In TBGP, a set of route attestation rules is strictly enforced in each router to simplify route attestations and build a trusted Internet routing infrastructure, and thus aggregated signatures are eliminated without sacrificing the security of BGP. Our prototype leverages the trusted computing (TC) technology to build transitive trust relationships between BGP speakers, and the identity-based signature (IBS) algorithm to sign/verify BGP routes and reduce the complexity of security operations in existing secure BGP solutions. Our security analysis and performance study shows that TBGP meets the security goals of BGP with significantly better convergence performance and lower resource cost than traditional solutions.

The lack of security in interdomain routing protocols is increasingly recognized as an important problem. An important aspect of any comprehensive approach is the means by which it performs origin authentication. An origin authentication service traces and validates the delegation of address usage from authorities to organizations, and ultimately to the ASes which originate them. Previous works have identified simple solutions, but no work has defined and generalized origin authentication or evaluated solutions using a complete picture of delegation on the Internet.

This work is composed of three serial efforts: formalization, modelling, and simulation. We initially formalized the semantics of address advertisements and proofs of delegation. Broad classes of origin authentication services are defined by extending existing cryptographic proof systems.

Securing the current interdomain routing infrastructure is likely to be a lengthy process. The security and networking communities must continually re-evaluate the assumptions and environments upon which the solutions are based. Work such as this serve as important contributions to this process. A thorough understanding of the trade-offs inherent to these services is essential.

III. CONCLUSION AND FUTURE WORK

The security of the routing system remains an open challenge in today's Internet. In this paper we assess the state of the art in proposed solutions. Our examination over the advantages and disadvantages of different classes of solutions suggests that detect-and-react type of solutions are most promising: they do not require any change in the protocol and they are the only solutions currently used by network operators to battle against attacks and faults in the BGP routing systems. As code attestation technology receives increasing attention in the

research community, we are interested in addressing the following questions: 1) what are the desired properties we would ultimately like to achieve out of attestation? 2) Suppose we were able to build a perfect attestation service with all of the desired properties, and make it available on every platform, how can it aid us in designing secure distributed systems in general? 3) how far are we from the perfect attestation service and how far can we push our limits toward this goal using currently available TCG and microprocessor technology? We propose BIND, a fine-grained attestation service that ties the proof of what code has executed to the data the code has produced. By attesting to the critical code immediately before it executes, we narrow the gap between time-of -use and time-of-attestation. BIND is useful for establishing a trusted environment for distributed systems, and greatly simplifies the design of secure distributed systems. For future work, we want to investigate the feasibility of a hardware based design for BIND. The current version of BIND runs in the Secure Kernel and assumes that the Secure Kernel is trustworthy, which is a hybrid hardware and software solution. However, it will be desirable to place trust only on hardware and no software components at all. The two main mechanisms we need to secure BGP is to verify the correctness of the origin of the prefix (to prevent prefix theft), and to prevent a malicious AS from altering the ASPATH in any other way than appending its own ASN to the path.

IV. REFERENCES

[1]. G. Goodell, W. Aiello, T. Griffin, J. Ioannidis, P. McDaniel, and A. Rubin, "Working around BGP: An incremental approach to improving security and accuracy of interdomain routing," in *Proc. ISOC NDSS*, 2003, pp. 75–85.

[2]. E. Keller, M. Yu, M. Caesar, and J. Rexford, "Virtually eliminating router bugs" in *Proc. ACM CoNext*, 2009.

[3]. R. Gummadi, H. Balakrishnan, P. Maniatis, and S. Ratnasamy. Not-a-bot: Improving service availability in the face of botnet attacks. In *Proc. of the NSDI*, 2009.

[4]. J. Caballero, T. Kampouris, D. Song, and J. Wang, "Would diversity really increase the robustness of the routing infrastructure against software defects?" in *Proc. ISOC NDSS*, 2008.

[5]. J. Karlin, S. Forrest, and J. Rexford, "Autonomous security for autonomous systems," *Computer Networks*, vol. 52, pp. 2908–2923, 2008.

[6]. L. Subramanian, V. Roth, I. Stoica, S. Shenker, and R. Katz, "Listen and whisper: Security mechanisms for BGP," in *Proc. NSDI*, 2004.

[7]. N. Aaraj, A. Raghunathan, and N. K. Jha, "Analysis and design of a hardware/software trusted platform module for embedded systems", *ACM Trans. Embedded Computing Syst.*, vol. 8, no. 1, 2008.

[8]. P. McDaniel, W. Aiello, K. R. B. Butler, and J. Ioannidis, "Origin authentication in interdomain routing," *Computer Networks*, vol. 50, no. 16, pp. 2953–2980, 2006.

[9]. P. van Oorschot, T. Wan, and E. Kranakis, "On inter-domain routing security and pretty secure BGP (psBGP)," *Proc. ACM TISSEC*, vol. 10, no. 3, pp. 1–41, 2007.

[10]. Youtube Hijacking: RIPE NCC RIS Case Study [Online]. Available:http://www.ripe.net/news/study-youtube-hijacking.html

[11]. Ricardo Oliveira, Mohit Lad, Lixia Zhang, "Understanding the Challenges in Securing Internet Routing".

[12] Y. Hu, A. Perrig, and M. Sirbu. SPV: Secure path vector routing for securing bgp. In *Proc. of the ACM SIGCOMM*, pages 179–192, 2004.

[13]. T. G. Griffin, F. B. Shepherd, and G. Wilfong. The stable paths problem and interdomain routing. *IEEE/ACM Transactions on Networking*, 10(2):232–243, 2002.

[14]. S. Goldberg, S. Halevi, A. D. Jaggard, V. Ramachandran, and R.N. Wright, "Rationality and traffic attraction: Incentives for honest path announcements in BGP," in *Proc. ACM SIGCOMM*, 2008, pp. 267–278.

[15]. S. Kent, C. Lynn, and K. Seo, "Secure border gateway protocol," *IEEE J. SAC*, vol. 18, no. 4, pp. 582–592, Apr. 2000.

[16] Qi L, Mingwei Xu, Jianping Wu, Xinwen Zhang, Patrick P. C. Lee, "Enhancing the Trust of Internet Routing with Lightweight Route Attestation" , IEEE Transactions on Information Forensics and Security, Vol.7, No.2, April 2012.

[17]. P. Reynolds, O. Kennedy, E. G. Sirer, and F. B. Schneider. Securing BGP using external security monitors. *Cornell University, Computing and Information Science, Technical Report TR2006-2065*, 2006.

[18]. E. Shi, A. Perrig, and L. van Doorn, "BIND: A fine-grained attestation service for secure distributed systems," in *Proc. IEEE Symp. Security Privacy*, 2005, pp. 154–168.