Probabilistic IP Prefix Authentication (PIPA) for Prefix Hijacking

Akmal Khan
Seoul National University (Korea)
raoakhan@mmlab.snu.ac.kr

Ted "Taekyoung" Kwon
Seoul National University (Korea)
tk@mmlab.snu.ac.kr

Hyunchul Kim
Seoul National University (Korea)
hkim@mmlab.snu.ac.kr

ABSTRACT

BGP is the most important component of Internet routing and yet it is vulnerable to many threats such as IP prefix hijacking, which has created significant problems over the decade. There have been two approaches to address the IP prefix hijacking issue: anomaly detection-based approach and cryptography-based one. Due to complexity and deployment concern of the latter, there are a lot of solutions that take the former approach. We propose a probabilistic IP prefix authentication (PIPA) scheme that leverages the existing BGP anomaly detection-based solutions as well as public internet registry information. That is, PIPA determines the authenticity of the pair (IP prefix, AS path) in BGP messages by using historical stability of the BGP information and internet registry data. We also discuss how to recover the hijacked IP prefixes in PIPA.

Categories and Subject Descriptors

C.2.0 [**Computer-Communication Networks**]: General *Security and Protection*; C.2.2 [**Computer-Communication Networks**]: Network Protocols—*Routing Protocols*; C.2.3 [**Computer-Communication Networks**]: Network Operations --*Network Monitoring*

Keywords

Inter Domain Routing, BGP Security, IP prefix hijacking

1. INTRODUCTION

The Internet is often defined as a network of networks, or a network of autonomous systems (ASs). An AS is a collection of connected IP network prefixes (or subnets) under the control of typically a single network operator. Normally, an AS is governed by a common, clearly defined routing policy. An AS should connect with other ASs to provide connectivity to end-users. The connectivity among ASs, so called AS path information, should be disseminated throughout the Internet so as to any source can send packets to any destination, which is the role of inter-domain routing. The de facto inter-domain routing protocol in the Internet is Border Gateway Protocol (BGP).

BGP is a path vector protocol to carry routing information across multiple ASs without loops. The term "path vector" comes from the fact that BGP routing information carries a sequence of AS numbers (or AS path) along the path over which an IP prefix has traversed [1]. BGP was designed with no consideration for security. Thus, any IP prefix disseminated across multiple ASs can cause a significant disruption in the Internet connectivity. As a result, there are tons of studies to address BGP security issues [2].

BGP is vulnerable to misconfigured and malicious routing information as there is no verification mechanism of the incoming routing information. One of the most notorious BGP attack is IP prefix hijacking, which occurs when a malicious or misconfigured BGP router originates an IP prefix that the router (or the AS that contains the IP subnet) does not own. IP prefix hijacking is essentially a special form of denial of service attack [6, 7]. Even though BGP operates well in practice due to simplicity and resilience, some outages may lead to significant and widespread damage. For instance, one of the early BGP hijacks happened in 1997, where traffic to be redirected to as7007 hijacked a lot of specific (or longer) IP prefixes. Some of the more recent incidents of that kind are ConEd [9] and an outage for the popular YouTube site caused by Pakistan Telecom [10]. As the number of critical applications (online banking, stock trading, and telemedicine) on the Internet grows, there will be more dependency on the underlying network infrastructure to provide reliable and secure internet connectivity [1].

That motivates us to focus on how to check the authenticity of the pair (IP prefix, AS path). Each BGP routing message carries this pair, and there are many anomaly detection-based BGP security systems that use this information. As there are many noisy BGP routing messages (e.g. even suspicious BGP messages that announces new pair of (IP prefix, AS path) may be legitimate), we take a somewhat probabilistic approach in this paper. We propose a probabilistic IP prefix authentication (PIPA) scheme that leverages some of the existing solutions to sanity-check the authenticity of AS path for an IP network prefix in BGP routing messages. One of the basic criteria in PIPA is how long the pair (IP prefix, AS path) has been announced, which is the interval since its first announcement until the announcement of new pair (the same IP prefix, new AS path). IP subprefix within the range of the same IP prefix can cause the same hijacking problem since the router prefers the longer prefix matching in making routing decisions.

PIPA will rely on several sources to check the authenticity of new pair (IP prefix/subprefix, new AS path). For instance, well-known routing monitoring systems such as RouteView and RIPE will be used. Also PIPA will refer to public internet registries like RIR and IRR. Even real time monitors like BGPmon [29] can be contacted. Furthermore, "unreachability information" available in Hubble [28] and PlanetSeer [30] will be retrieved periodically. If PIPA concludes that a newly announced pair of (IP prefix, AS path) is not credible, PIPA will take a countermeasure to remedy this hijacking incident. As for multi-origin ASes (MOASs), we

CFI'09 June 17--19, 2009, Seoul, Korea.

Table 1 : Taxonomy of Prefix Hijacking Solutions (PH: Prefix Hijacking, Y: Yes, N: No, H: History, R: Registry, Un: Unreachability, MITM: Man In The Middle)

	Detection System	Alarm Type	Prefix/Duplicate PH	Subprefix PH	Super/Indepe--ndent PH	Path Spoofing	MITM
PHAS [15]	H	Origin, Last Hop, Sub Allocation	Y	Y	N	limited	N
PG-BGP [16,17]	H	Prefix, SubPrefix	Y	Y	N	Y	limited
Jian Qiu et al. [11]	H	N	Y	Y	Y	Y	N
K.Sriram et al. [12]	H+R	N	Y	Y	N	Y	N
Nemecis [19]	R	N	Y	Y	N	N	N
Krugel et al. [26]	H	N	Y	N	N	Y	N
Hu et al. [27]	H	N	Y	Y	N	Y	N
PIPA	H+R+Un	Prefix/Sub/Independent/Path/MITM.etc.	Y	Y	Y	Y	Y

will treat each pair separately. The credibility of each pair will be handled independently.

2. BACKGROUND

Research community on inter-domain routing has worked out many protocols and technical contributions for BGP operational issues such as scalability, convergence, routing stability, and performance. However, the security aspects of BGP have not been practically solved.

A BGP attacker could be a network operator who has misconfigured its BGP routing behaviors. Also, a malicious party may gain control of a BGP speaking router on the black-market. Spammers with upstream hijacked address space might be able to operate a portion of the infrastructure [3-5]. In any case, fake routes may be fabricated by a sophisticated attacker to manipulate arbitrary address spaces so that the attacker can launch a stealthy attack or access the relevant traffic of IP prefixes [11]. There are different types of prefix hijacks and it is important to address all or the most of them. *Table 1* provides the comparison of prefix hijacking solutions and their features. Prefix hijacking can be classified into:

Prefix hijacking: an AS directly originates the route(s) of an arbitrary prefix

Sub prefix hijacking*:* an AS originates the routes of a sub-prefix (i.e. network prefix of smaller size than the hijacked prefix)

Duplicate prefix hijacking**:** an AS announces a prefix used by another AS to gain access to the traffic of the prefix.

Super prefix hijacking: an AS originates the routes of a super-prefix (i.e. network prefix of greater size than the hijacked prefix).

Independent prefix hijacking: an AS originates the routes of a prefix entirely in unused address space

Man in the middle (MITM): it allows an attacker to make traffic for certain destinations redirected to an attacker.

Over the last decade, researchers have contributed different solutions mainly based on cryptography-based solutions (S-BGP, PS-BGP [13, 14]), anomaly detection-based ones (PHAS [15], PG-BGP [16, 17]) and so on. Some of the desirable requirements for BGP security solutions are real-time, accurate, light-weight, easy deployment, incentive, and robustness. Cryptography-based solutions have been around for a while but ISPs show little interest because of their complexity and non-compatibility issues.

Anomaly detection-based systems rely on measures like the generation of alarms when anomalies are detected, access control lists (ACLs) and BGP filter lists to prevent or allow the distribution of specific IP prefixes. Anomaly detection-based solutions work by gathering BGP routing data from multiple vantage points.

Anomaly detection-based systems differ in the type(s) of data they use. Some are based on registry data from Regional Internet Registries (RIRs) and Internet Routing Registries (IRRs) - an example is the Nemecis tool [19]. Others such as the Prefix Hijack Alert System (PHAS) and the Pretty Good BGP (PG-BGP) are driven by BGP trace data. The trace data is obtained from global BGP monitoring infrastructures (e.g., RIPE-RIS, Routeviews) or a BGP speaker where the algorithm operates. There are a number of data sources of BGP routing information available for a BGP security solution to get data such as Routeviews, Reseaux Internet Protocol Europeans - Routing Information Service (RIPE-RIS), Cooperative Association for Internet Data Analysis (CAIDA). Declarative routing information is available from addressing and routing registries such as RIPE, American Registry for Internet Numbers (ARIN), Routing Assets Database (RADB). There are other BGP information sources available such as bogon lists. But correctness, freshness, and consistency of the data derived from these sources must be taken into account by any BGP Security Solution [21-25].

Anomaly detection-based solutions like PG-BGP, PHAS are very good at detecting the prefix hijacks but lack the property of timely recovery since they just rely on notifying the victim AS through email and ratio of generating false alarms is quite high.

3. PROPOSED SOLUTION

3.1. PIPA Overview

We propose PIPA by illustrating a hijacking and detecting scenario. Suppose there is an IP Prefix that has been working fine in the Internet but an attacker announces a new path for the same prefix or more specific subprefix. Then, routers may start switching to the more recently announced AS path than the old one. What happens is that all the traffic for the hijacked prefix will be black holed. In that case, unreachability to the prefix will be observed at some points in the Internet. PIPA tries to assign a probability to each IP prefix by retrieving the unreachability information (from Hubble project, PlanetSeer, etc) as well as the historical data collected from different data sources like RouteView, RIPE. If a pair of an IP prefix and it's AS path has been working well until the announcement of a new pair of the same prefix and its new AS path, the probability of hijacking is increased as the new pair causes multiple sources of unreachability reports. Based on the unreachability reports along with other analysis of data sources, PIPA can conclude whether IP prefix hijacking occurs or not.

PIPA can also initiate the recovery process once a particular prefix is concluded to be hijacked. PIPA will stop using the new AS path information of the hijacked prefix; instead, it will switch back to the old path information for the hijacked IP prefix. This fallback process is done by comparing the "hijack" probabilities of multiple AS-path entries that correspond to the same IP prefix. Hubble project can be one source of getting the unreachability information but we are also considering using distributed active probes to decide more quickly about unreachability.

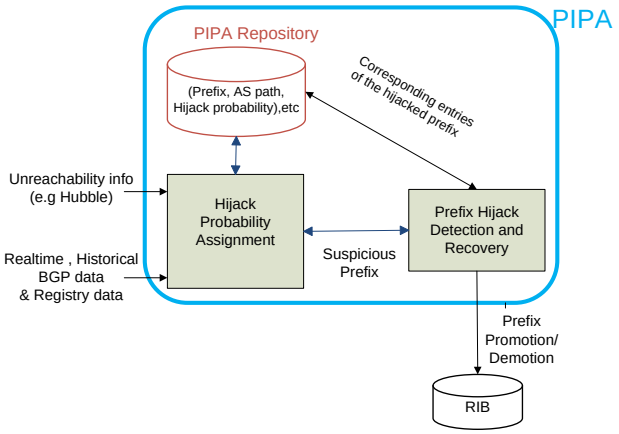

Figure 1. Functional blocks of PIPA

To check the authenticity of the AS PATH information, the routing policy information collected from RIR/IRR such as the connectivity between adjacent ASs will be retrieved. As the control plane information may not be sufficient to check authenticity, PIPA also uses the data plane information which can be collected through distributed active probing. There are a number of solutions to detect prefix hijacking; they use Route Views/RIPE-RIS and registry-based data to authenticate the AS path or prefix ownership through the history of collected for BGP traces or the current routing policy data from RIR/IRR. However, one has to consider the correctness and freshness of these data sources.

In May 2009, there are 31315 ASes in routing system, and 13327 ASes among them announce only one prefix. A single AS announces even 4303 prefixes [32].There are two main modules of PIPA namely: *Hijack Probability assignment* and *Prefix hijack detection and recovery*

3.2 Hijack Probability Assignment

Hijack Probability assignment starts by assigning an initial "hijack" probability value to every new entry of {IP Prefix, AS Path} announced in BGP data collectors. The PIPA repository in Figure 1 is used to store the {IP Prefix, AS Path, Hijack Probability value}. To determine a "hijack" probability for each entry, PIPA first checks whether an incoming BGP update message matches with the corresponding entry of BGP data collectors, RIR/IRR records and the unreachability information from the Internet. The hijack probability assignment is also dependent on (1) number of updates of a particular prefix, (2) age of an AS path and so on. If the prefix probability of a particular prefix exceeds the predetermined threshold, that prefix information is sent to the *Prefix Hijack Detection and Recovery* module, which performs the additional checks and actions below.

3.3 Prefix Hijack Detection and Recovery

Receiving the suspicious prefix information from the *Hijack Probability Assignment module*, IP prefix hijacks can be easily detected as the newly advertised suspicious prefix information must have higher hijack probability than that of the prefix information announced by the owner AS because of its limited coverage and time. A sub-prefix hijack can also be detected as PIPA performs additional checks as follows. Assume that this is not a Multi Origin AS (MOAS) problem (MOAS will be discussed later). Suppose a certain prefix entry has the lowest hijack probability value which means that it has been valid in the routing system. If there is any sub-prefix announcement from any other AS other than the owner AS of the original prefix, then this sub-prefix is suspicious and will be assigned a higher probability value. In this case, additional unreachability information can increase the hijack probability. The same rules apply to super/independent prefix hijacking.

MOAS conflicts occur when a particular prefix appears to originate from more than one AS. MOAS conflicts can be classified into OrigTranAS, SplitView, and Distinct paths [33]. For these cases, a MOAS conflict can be detected by using the up-to-date RIR/IRR routing policy data.

After concluding that a prefix is hijacked, the *Prefix hijack detection and recovery* module will trigger the recovery procedure. So far, a prevalent recovery procedure is through contacting network operators, who then manually change filters, stopping corresponding prefixes announcements or blocking malicious or misconfigured AS. Most of the BGP anomaly detection solutions generate alerts to inform network operators of their prefix hijacks attempts [15, 16], but there have been a number of false alarms generated.

Our solution to recovering from a prefix hijack is two-fold. The first mechanism is the same as providing the alert notifications to the victim ASes. The second one is based on the idea of *self healing property* of Internet. PIPA can help victim ASes recover from prefix hijack situation by suggesting which malicious or

erroneous {IP Prefix, AS Path} must be demoted and which {IP Prefix, AS Path} entry is the best or the lowest hijack probability.

4. CONCLUSION

Prefix hijacking has been a serious BGP security issue over the years. There have been a large number of proposals to the problem of detecting prefix hijacking but none has been accepted as a de facto standard. We proposed a probabilistic IP prefix authentication (PIPA) scheme by leverage the existing BGP routing information, registry data as well as prefix unreachability statistics from Internet.

5. ACKNOWLEDGEMENT

This work was supported in part by the IT RD program of MKE/IITA [2007-F-038-03, Fundamental Technologies for the Future Internet] and in part by Korea Research Council of Fundamental Science & Technology. The ICT at Seoul National University provides research facilities for this study.

6. REFERENCES

[1] http://www.ietf.org/rfc/rfc4271.txt
[2] K. Butler, T. Farley, P. McDaniel and J. Rexford, "A Survey of BGP Security," Draft August 2008
[3] Blaine Christian, Tony Tauber, "BGP Security Requirements", http://www.ietf.org/ids.by.wg/rpsec.html
[4] BGP Security Vulnerabilities Analysis, S. Murphy. RFC 2006.
[5] Nordstrom, Dovrolis,"Beware of BGP Attacks," ACM SIGCOMM Computer Communications Review 1 Volume 34, Number 2: April 2004
[6] Z. Zhang, Y. Zhang ,Y. Charlie ,Z. Morley, Randy Bush, "iSPY: Detecting IP Prefix Hijacking on My Own" SIGCOMM'08, August 17–22, 2008.
[7] Z.Zhang,Y. Zhang,Y. Charlie Hu, Z. Morley Mao, "Practical Defenses Against BGP Prefix Hijacking," CoNEXT'07, December 10-13, 2007.
[8] V.J. Bono, AS7007 explanation and apology, April 1997
[9] Renesys Blog, Con-Ed Steals the Net
[10] Renesys Blog, Pakistan Hijacks YouTube
[11] J. Qiu, L. Gao, S. Ranjan, and A. Nucci, "Detecting bogus bgp route information: Going beyond prefix hijacking," SecureComm 2007
[12] K. Sriram, O. Borchert, O. Kim, and P. Gleichmann, and D. Montgomery, "A Comparative Analysis of BGP Anomaly Detection and Robustness Algorithms," CATCH '09, Washington D.C., March 3-4, 2009.
[13] S. Kent, C. Lynn, and K. Seo, "Secure Border Gateway Protocol (S-BGP)," IEEE Journal on Selected Areas in Communications, vol. 18, no. 4, Apr. 2000.
[14] P. C. van Oorschot, T. Wang, and E. Kranakis, "On Inter-domain Routing Security and Pretty Secure BGP (psBGP)," ACM TISSEC, vol. 10, no. 3, Jul. 2007.
[15] M. Lad, D. Massey, D. Pei, Y. Wu, B. Zhang, and L. Zhang, "PHAS: A Prefix Hijack Alert System," in USENIX Security Symposium 2006.
[16] J. Karlin, S. Forrest, and J. Rexford, "Autonomous security for autonomous systems," Computer Networks, 2008
[17] J. Karlin, S. Forrest, and J. Rexford, "Pretty Good BGP: Improving BGP by Cautiously Adopting Routes," IEEE ICNP 2006, Santa Barbara, CA, USA, Nov. 2006.
[18] G. Siganos and M. Faloutsos, "A Blueprint for Improving the Robustness of Internet Routing," Security '06, 2006.
[19] G. Siganos and M. Faloutsos, "Analyzing BGP policies: methodology and tool," IEEE Infocom, 2004.
[20] BGPmon http://bgpmon.net/
[21] CAIDA http://www.caida.org/
[22] RouteViews www.routeviews.org/
[23] RIPE-RIS www.ripe.net/ris/
[24] Merit Network Routing Assets Database www.radb.net/
[25] IETF Working Group Secure Inter-Domain Routing (sidr),Routing Protocols Security(rpsec)
[26] C. Krugel, D. Mutz, W. K. Robertson, and F. Valeur, "Topology-Based Detection of Anomalous BGP Messages," in *RAID*, 2003, pp. 17–35
[27] Xin Hu and Z. Morley Mao, "Accurate Real-time Identification of IP Prefix Hijacking," *IEEE Security and Privacy*, Oakland, 2007.
[28] Studying Black holes in the Internet with Hubble http://hubble.cs.washington.edu
[29] http://bgpmon.netsec.colostate.edu/
[30] M. Zhang, C. Zhang, V. S. Pai, L. Peterson, and R. Wang. PlanetSeer, "Internet path failure monitoring and characterization in wide-area services," OSDI '04, 2004.
[31] http://www.nanog.org/mailinglist/
[32] http://www.cidr-report.org/as2.0/#General_Status
[33] X.Zhao,D.Pei.Lan Wang,D.Massey,A.Mankni,S.Felix Wu,L.Zhang ,"An Analysis of BGP multiple Origin AS(MOAS) Conflicts" IMW'01, November 2001