NPLA: Network Prefix Level Authentication

Ming Li*, Yong Cui†, Matti Siekkinen*, Antti Ylä-Jääski*

*Department of Computer Science and Engineering
Helsinki University of Technology
FI-02015, Espoo, Finland
{ming.li, matti.siekkinen, antti.yla-Jaaski}@tkk.fi

†Department of Computer Science and Technology
Tsinghua University
100084, Beijing, China
cuiyong@mail.tsinghua.edu.cn

***Abstract*—We present the design and evaluation of NPLA (Network Prefix Level Authentication), a system allowing source addresses to be validated within the network to the granularity of network prefix. Prefix routers use public key cryptography to insert NPLA headers in outgoing packets. En route entities holding the corresponding public key verify the source of a packet. NPLA provides deployment incentives because each upgraded prefix can prevent its address space from being maliciously used by other networks and its traffic is forwarded with high priority. In order to increase the scalability, NPLA does not employ PKI but leverages BGP to distribute public keys. Based on the relative damage reduction analysis, we conclude that NPLA provides more relative benefit than other approaches when they are all partially deployed. In order to decrease the overhead induced by public key cryptography, NPLA uses FPGA based hardware cryptography accelerator which has been proven to achieve several Gbps throughput on average.**

Index Terms - source spoofing, authentication, network prefix, public key cryptography

I. INTRODUCTION

The existing Internet allows packets with forged source addresses to traverse the network. Attackers leverage this vulnerability to launch spoofing attacks for anonymity and redirection of accusation. In recent years, some believe attackers no longer need to use spoofed IP addresses because of the prevalence of botnets [1] which allow the compromised hosts to use their real addresses, while recent studies [2, 3] show that IP spoofing is still a common attack vector. The quantitative results presented in [4] reveal ~31% of clients are able to spoof an arbitrary routable address.

IP addresses have an inherent hierarchical structure. According to how the granularity the addresses are authenticated, we divide the proposed solutions into three categories: Host level solution, AS (Autonomous System) level solution, and prefix level solution. Host level solutions (e.g., IPSec [5], PLA [6] and TVA [7]) provide a fine-grained source authentication scheme. They authenticate the source address of a packet to the granularity of its original host. AS level solutions (e.g., SPM [8], Passport [9] and IDPF [10]) supply a coarse-grained source authentication scheme verifying whether a packet is really from its original AS. Prefix level solutions (e.g., Ingress filtering [11]) provide a moderate-grained scheme ensuring a packet is from its original prefix network.

In this paper, we propose an alternative of the prefix level solutions, NPLA (Network Prefix Level Authentication). The basic idea is to use the public key signature as the network prefix authentication information. A packet carries a digital signature signed by a prefix border router using its private key; and the en route routers holding the corresponding public key verify the signature before forwarding the packet. We do not use PKI (Public Key Infrastructure) but the routing system (BGP [13]) to distribute the public keys among routers.

In order to decrease the overhead induced by public key cryptography, we use the ECC (Elliptic Curve Cryptography) [14] public key algorithm and a FPGA (Field Programmable Gate Array) implemented cryptography accelerator. Thus, it provides a good balance between filtering granularity, accuracy and speed. NPLA is a self-protected approach that offers deployment incentives. A prefix implementing NPLA guarantees that its addresses are not being used maliciously elsewhere on the network as long as there is at least one NPLA enabled router on the route. According to the effectiveness analysis, we conclude NPLA offers more relative benefit than other approaches.

In this paper, we assume a prefix a trust and fate-sharing unit. The network administrators may use whatever mechanisms they prefer to prevent intra-prefix spoofing. How to prevent source spoofing within a prefix is out of the scope.

The rest of the paper is organized as follows. In Section 2, we discuss the related work. Section 3 presents the design of NPLA. In Section 4, we analyze the deployment issues. In Section 5, we give the effectiveness analysis. Performance and security issues are analyzed in Section 6 and Section 7 respectively. We conclude the article in Section 8.

II. RELATED WORK

According to how the granularity the source addresses are authenticated, we divide the solutions in literature into three categories: Host level, AS level, and prefix level solutions.

Host level solutions are either implemented on the end-hosts or on both of the end-hosts and routers. End-to-end based host level solutions (e.g., IPSec [5]) do not rely upon

This work was supported in part by the Academy of Finland (no. 135230), NSFC (60911130511, 60873252), 973 Program of China (2009CB320501, 2009CB320503) and 863 Program of China (2008AA01A324).

special router functionality and are the easiest to deploy. However, they act too late to prevent the spoofed packets from consuming the network resources because the packets must reach destination hosts before being detected. Hop-by-hop based host level solutions (e.g., PLA [6] and TVA [7]) involving the support of both of hosts and routers can detect and discard spoofed packets in the middle of the network before they reach the destination hosts. Although they are the most effective solutions, they face many deployment challenges, e.g., requiring the end-hosts to upgrade.

AS level solutions (e.g., SPM [8], Passport [9] and IDPF [10]) are usually deployed in the middle of the network. The end-hosts have no idea of their existence and do not need to upgrade. Thus, as long as an AS deploys it, its addresses cannot be spoofed at other networks. The problem of these solutions is that an AS might involve millions of IP addresses. Attackers can take advantage of the intra-domain spoofed packets without being detected.

Prefix level solutions (e.g., Ingress filtering [11]) are compromising solutions considering the deployment challenges and security requirements. According to [12], each AS has around 10 prefixes advertised to other ASes. Thus, they give much smaller spoofing range the attackers can make use of compared to the AS level solutions. Additionally, they are usually deployed in the middle of the network. Thus they avoid many deployment hurdles, e.g., convincing the end-hosts to upgrade. From a practical point of view, prefix level solutions may be the most feasible solutions to mitigate the ability to spoofing.

III. General Design On NPLA

In this section, we assume all the involved entities support NPLA. In section IV, we discuss how to deal with the legacy entities during incremental deployment.

A. NPLA Architecture

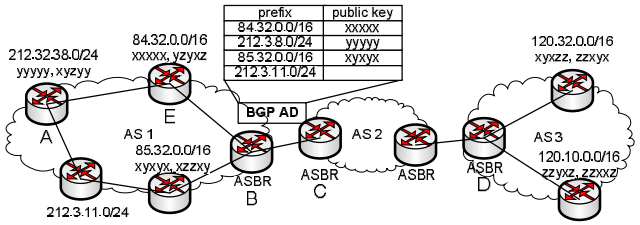

Fig. 1. A simplified NPLA architecture

Fig. 1 illustrates how NPLA works at a high level. Each network prefix that is to be advertised to other ASes through BGP holds a pair of public/private keys. In this paper, the router to which an advertised network prefix is assigned is called a prefix router, e.g., the four edge network routers in AS1 (not B) and two in AS3 (not D).

When an outgoing packet leaves a prefix router (e.g., the router A), it stamps a NPLA header (involving a digital signature) to the packet using its private key. We give the format of the NPLA header in subsection C. When the packet enters en route verifying entities, such as the AS ingress border routers like routers C and D, they verify the signature using the corresponding public key.

If the verification succeeds, it demonstrates that the packet comes from the original prefix indicated by its source address. Otherwise, the verifying router concludes the source address is spoofed. A packet with an invalid signature is discarded immediately to prevent it from consuming the resources furthermore. More detailed stamping and verification procedures are given in subsection C. Here we clearly point out which routers stamp outgoing packets and which perform the verification on incoming packets. IP addresses have hierarchical structure. In NPLA, the level of the prefix for authentication is not arbitrary. Only the prefixes an AS advertises to other ASes are responsible for the packet stamping. The packet verification is done by the AS ingress border routers and the destination prefix routers within the same AS as the source prefix routers.

B. Prefix Public Key Distribution

The objective of the public key distribution is to allow the verifying nodes to hold the public keys of network prefixes. The key distribution involves two steps: 1) Let the prefix routers within an AS share their public keys with each other; 2) Allow AS border routers to advertise their network prefixes and the corresponding public keys to other ASes.

For the first step of the key distribution, NPLA does not specify how the prefix routers within an AS exchange the public keys. There are two basic approaches: Offline manual configuration and online automated key distribution (e.g., [15]). The network administrators can use whatever they prefer to distribute the keys and they also need to explicitly configure the prefix routers according to the advertised prefixes. The specific mechanisms used for key distribution within an AS are out of the scope of this paper.

For the second step, as illustrated in Fig. 1, after obtaining the public key information of the prefix routers in the same AS, each AS border router uses BGP update messages to advertise its network prefixes and their corresponding public keys to all the other ASes. When the BGP converges, each AS obtains the public keys of all the network prefixes on the Internet. The idea of distributing key information through BGP is not new and was used earlier in Passport [9]. The example format of a BGP update message piggybacking a public key is illustrated in Fig. 2.

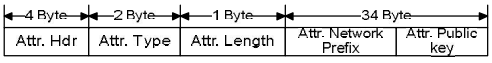

Fig. 2. BGP update message piggybacking a public key

To increase the system security, the public keys of network prefixes should update periodically, e.g., on the order of weeks or months. Simply, each prefix router periodically gets a new set of public/private keys and updates them to other routers.

During the re-keying process, each router holds two keys for a prefix: An old one and a new one. NPLA uses an alternating parity bit in the NPLA header to differentiate them. An arriving packet is authenticated using the old one if the parity bit is set and using the new one if it is unset.

The reason why we use BGP to distribute public keys is that the BGP runs a critical task for the regular Internet connections. If the assumption that the routing protocol itself is safe to depend on does not work, the Internet may collapse. However, we are not saying that the BGP is absolutely safe. BGP hijacking issue bothers the Internet for a long time. To enhance the security of NPLA, security mechanisms to protect BGP can be used, e.g., S-BGP [16] and psBGP [17]. For this paper, how to secure the routing system is out of the scope.

C. Stamping and Verification

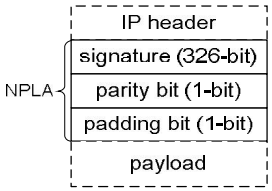

Fig. 3. NPLA header format

Before we present the stamping and verification procedures, we first give the NPLA header format as shown in Fig. 3. NPLA is inserted as a shim layer between the transport layer and network layer. It consists of two fields:

1) Signature is calculated by a prefix router using its private key over the source address, destination address, transport header and data payload.
2) Parity bit is used to differentiate the old and new public keys to authenticate a packet.

The basic stamping and verification procedure are described as follows. When a packet leaves its prefix network, the prefix router calculates a NPLA header and inserts it into the packet. When an en route AS ingress border router receives the packet, it verifies the packet using the corresponding public key: First of all, the verifying node uses the source address of the packet to lookup the corresponding network prefix, then obtains the public key of the network prefix and at last verifies the signature using the matched public key. If the verification fails, the packet is discarded. Mapping an IP to the network prefix requires the router to perform an extra prefix lookup on the source address of the packet. When the packet arrives at the destination AS or destination prefix router within the same AS as the source prefix router, the verifying node verifies it, strips off the NPLA header, and then forwards it.

IV. Deployment Consideration

In real networks, we cannot expect the full deployment of NPLA overnight. NPLA should co-exist with the existing Internet and provide deployment incentives to adopters. In this section, we discuss how NPLA deals with legacy entities and traffic, and what deployment incentives it provides.

A. Inter-operation with Legacy Entities

Among ASes, AS border routers use the optional and transitive path attributes of BGP to advertise the prefix public keys. Legacy ASes will not process the optional and transitive path attributes they received, but just include them in the routing information to propagate them to their neighboring ASes. Thus, legacy and upgraded ASes is able to co-exist.

NPLA is deployed so that the end-hosts have no idea of its existence. The source prefix router inserts a NPLA header into a packet and the destination network strips off the NPLA header before the packet reaches the destination host. However, with no full deployment, we cannot guarantee that there always exist en route NPLA enabled entities. Thus, three extra check operations are designed to prevent the NPLA packets reaching the destination hosts: 1) If the source and destination networks are in the same AS, the source prefix router sends NPLA packets only if the destination network has also upgraded. The upgraded destination prefix router strips off the NPLA header. 2) When an outgoing NPLA packet (e.g., comes from the router A in Fig. 1) arrives at the source AS border router (e.g., the router B in Fig. 1), this router checks whether there exist en route NPLA enabled ASes. If there is none, it strips off the NPLA header from the packet. This mechanism implies that within an AS, if there is at least one upgraded prefix router, the border routers of that AS also upgrade. 3) If there are en route NPLA enabled ASes, the last one of them strips off the NPLA header of the packet. These three operations work under an assumption that each AS knows the en route ASes to the destination. This assumption is feasible because an AS border router can obtain the AS path information from BGP using the AS_PATH path attribute.

In addition, if a network prefix is aggregated by its provider AS, its prefix and public key information will be lost. In this case, a prefix should rely on its provider to stamp and verify its traffic. A customer prefix desiring to stamp its own NPLA headers should require its providers not to aggregate its prefix.

B. Handling Legacy Traffic

To motivate ASes to upgrade, NPLA uses two weighted queues to handle verified and legacy traffic, allocating limited bandwidth to legacy traffic especially when a link is congested. By prioritizing verified traffic, NPLA can significantly mitigate the impact of IP spoofing. An en route NPLA enabled router will discard legacy traffic if it detects the traffic spoofs other upgraded prefixes' addresses as follows: If the source prefix of a packet has deployed NPLA (it can be confirmed by checking whether the network prefix has a public key entry), the router discards the packet because it must be a spoofed packet.

C. Incentive Deployment

NPLA provides initial benefits for early adopters. Even the first two edge deployment can gain a benefit. For example, two edge prefix networks implementing NPLA guarantees that packets originating from one of them can be verified by one

another. Moreover, if an AS detects an attack on itself, it can protect itself from spoofed packets by allowing in only packets originating from NPLA enabled networks.

NPLA also get benefits from incremental deployment during which as long as there is at least one en route NPLA enabled AS, the spoofed packets will be detected and discarded.

V. Effectiveness Analysis

In this section, we analyze the effectiveness of different approaches. Our interest is in the amount of processing effort NPLA is able to reduce as a result of spoofed traffic travelling across the Internet. Specifically, we will conduct the analysis under: 1) The ingress filtering approach [11], 2) The ingress filtering club approach, 3) The SPM approach [8], 4) The NPLA approach, and 5) Some combinations of the approaches. Under the ingress filtering club, the networks which implement ingress filtering conduct ingress filtering only to traffic destined to networks that also implement ingress filtering. Unlike ingress filtering, it provides deployment incentives to its participants. For fairness, we use the ingress filtering club to make comparison with SPM and NPLA because they all provide deployment incentives.

In this paper, we use a simplified model aiming at a clear comparison among different approaches. We assume the Internet consists of N BGP advertised prefixes denoted by INT = {1,2,…,N}. Each prefix is in charge of the traffic originating from it. We focus on demonstrating the relative benefit of NPLA with respect to the reduced spoofed traffic processing effort. To simplify the comparison, we assume all approaches are implemented on network prefix level where the prefix routers stamp the packets and the en route ASes verify them.

Consider that an attacker located in the prefix j sends spoofed traffic towards the prefix i by spoofing the source addresses belonging to the prefix k. We assume each AS as a processing unit to process the spoofed traffic. The amount of spoofed packets reaching the h^{th} en route AS is $A_{j\to i}^{(k,h)}$.

In equation (1), D presents the amount of packets with spoofed source addresses multiplied by the number of ASes traversed by those packets. It sums up the overall processing effort of forwarding the spoofed packets sent from j to i. $H_{j\to i}$ denotes the number of AS hops from j to i. In this paper, the first en route AS is the next hop AS of the source AS. Thus, h=1 presents the first en route AS.

$$D = \sum_{h=1}^{H_{j\to i}} \sum_{k\in INT} \sum_{j\in INT} \sum_{i\in INT} A_{j\to i}^{(k,h)} \quad (1)$$

The damage reduction denoted DR is the amount of filtered spoofed packets multiplied by the number of ASes. These filtered spoofed packets do not traverse the entire routes due to the deployment of defensive approaches. Thus, the damage reduction DR implies the saved processing effort of the spoofed packets sent from j to i.

Equation (2) gives the damage reduction due to the deployment of the ingress filtering club where INC denotes the set of prefixes participating in it. The ingress filtering club can stop the spoofed traffic at its original networks as long as the destination network also implements ingress filtering club. Thus, the en route ASes, from the first en route AS (denoted by h=1) to the destination AS (denoted by $h=H_{j\to i}$), do not need to consume resources to process the spoofed traffic.

$$DR = \sum_{h=1}^{H_{j\to i}} \sum_{j\in INC} \sum_{k\in INT} \sum_{i\in INC} A_{j\to i}^{(k,h)} \quad (2)$$

The damage reduction due to the deployment of the SPM is given by equation (3) where SPM denotes the prefixes that deploy the SPM. In SPM, although the destination AS can recognize and discard the spoofed traffic, the en route ASes still consume resources to process the spoofed traffic. Thus, in our evaluation method, SPM is ineffective from the saved processing effort point of view.

$$DR = \sum_{h=5}^{4} \left[\sum_{j\in INT} \sum_{k\in SPM} \sum_{i\in SPM} A_{j\to i}^{(k,h)} + \sum_{j\in SPM} \sum_{k\in (INT-SPM)} \sum_{i\in SPM} A_{j\to i}^{(k,h)} \right] = 0 \quad (3)$$

Equation (4) presents the damage reduction due to the deployment of NPLA. m ($m\in [1, H_{j\to i}]$) denotes the index of the first NPLA enabled en route AS. Thus, the saved processing effort starts from its next hop AS.

$$DR = \sum_{h=m+1}^{H_{j\to i}} \left[\sum_{j\in INT} \sum_{k\in NPLA} \sum_{i\in INT} A_{j\to i}^{(k,h)} + \sum_{j\in NPLA} \sum_{k\in (INT-NPLA)} \sum_{i\in INT} A_{j\to i}^{(k,h)} \right] \quad (4)$$

We first compare the relative damage reduction (DR/D) of different individual approaches and later some combinations of them. DR/D implies the relative benefit due to the deployment of a defensive approach. To simplify the model, we assume a constant number of AS hops from any j to any i, i.e. $H_{j\to i}=H$, and $A_{j\to i}^{(k,h)}=C$ meaning that the amount of spoofed traffic is constant for all j and i. We also assume the number of prefixes that adopt the defense approaches is P. Under the parameters assumed, we compare the DR/D as follows.

1) Under ingress filtering club from (1) and (2):
$$DR/D = P^2/N^2 \quad (5)$$
2) Under SPM from (1) and (3):
$$DR/D = 0 \quad (6)$$
3) Under NPLA from (1) and (4):
$$DR/D = (H-m)\,(2P/N - P^2/N^2)/H \quad (7)$$

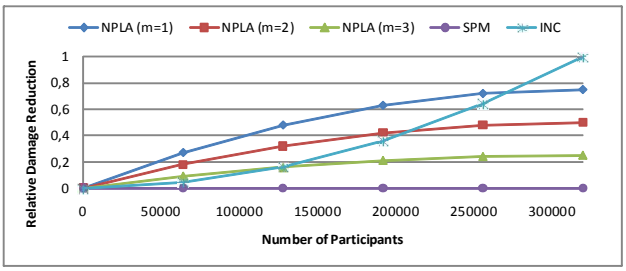

Fig. 4. The relative damage reduction of different approaches

The results are depicted in Fig. 4. We assume the number of prefixes on the Internet is 320K [12] and H equals 4 which is the average AS hops on the Internet [18]. We depict the DR/D is (Y axis) as a function of the number of prefixes (P is X axis) participating in the approaches. We do this for NPLA with m

= 1, 2, 3. The results demonstrate the following properties:

1) With ingress filtering club, the relative benefit grows slowly when the number of participants increases. It demonstrates that the scheme is not very useful when the number of participants is relatively small.
2) SPM can only detect and discard the spoofed packets at the destination network. The spoofed packets consume the same resources as the legitimate packets do. Thus, as we can see from the Fig. 4, it provides no relative benefit from the saved processing effort point of view.
3) With NPLA, the relative benefit increases quickly when the number of participants is small. It means this scheme is useful when partial ISPs implement the NPLA approach. Moreover, we note that the relative benefit increases faster when m is smaller. It implies that NPLA can provide more benefit if the first NPLA enabled AS is closer to the source NPLA enabled AS.

The reason why the ingress filtering club and NPLA behave differently as shown in Fig. 4 is the NPLA's ability to get rid of spoofed traffic as long as there's at least one NPLA enabled router on the way, while ingress filtering club requires the destination to upgrade.

For the long term consideration, combining NPLA with ingress filtering or ingress filtering club to further limit the spoofing range may be a good alternative. That is the NPLA enabled source prefix routers also enforce ingress filtering or ingress filtering club. In this case, it can discard the spoofed packets as early as possible and at the same time provide deployment incentives. The damage reduction of the combinations is given by (8) and (9) respectively where NPLAIN denotes the set of prefixes participating in NPLA and ingress filtering, and NPLAINC denotes the set of prefixes participating in NPLA and ingress filtering club.

$$DR=\sum_{h=1}^{H_{j\to i}}\sum_{j\in NPLAIN}\sum_{k\in INT}\sum_{i\in INT}A_{j\to i}^{(k,h)} + \sum_{h=m+1}^{H_{j\to i}}\sum_{j\in (INT-NPLAIN)}\sum_{k\in NPLAIN}\sum_{i\in INT}A_{j\to i}^{(k,h)} \quad (8)$$

$$DR=\sum_{h=1}^{H_{j\to i}}\sum_{j\in NPLAINC}\sum_{k\in INT}\sum_{i\in NPLAINC}A_{j\to i}^{(k,h)} + \sum_{h=m+1}^{H_{j\to i}}\sum_{j\in (INT-NPLAINC)}\sum_{k\in NPLAINC}\sum_{i\in INT}A_{j\to i}^{(k,h)} \quad (9)$$

The relative benefit due to the deployment of the combination approaches is given by equation (10) and (11).

1) Under NPLA and ingress filtering from (1) and (8):

$$DR/D=(2H-m)\,P/HN-(H-m)P^2/HN^2 \quad (10)$$

2) Under NPLA and ingress filtering club from (1) and (9):

$$DR/D=(1+m-H)\,P^2/HN^2 + (H-m)P/HN \quad (11)$$

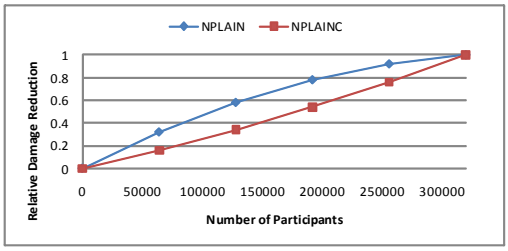

Fig. 5. The relative damage reduction of the combination of NPLA and ingress filtering, and the combination of NPLA and ingress filtering club

The results are depicted in Fig. 5 where we assume N=320K, H=4 and m=1. As shown in the figure, under the combination of NPLA and ingress filtering, the relative benefit grows faster than that of the combination of NPLA and ingress filtering club when the number of participants increases. This is also because NPLAINC requires the destination networks to implement it, while NPLAIN does not have this requirement. Thus, we suggest ISPs implement the combination of NPLA and ingress filtering.

However, NPLA does not depend on ingress filtering. Spoofed traffic is detected and discarded as long as there is at least one NPLA enabled node on the route. The combination enables a high anti-spoofing scheme by taking advantage of both of a self-defensive and a self-disciplined method.

VI. Performance and Overhead Analysis

Public key cryptography is usually criticized because of its overhead. This section analyzes the performance of a hardware cryptography accelerator to demonstrate the feasibility of NPLA to be implemented in practice.

We refer to the performance evaluation of a FPGA based hardware public cryptography accelerator. NPLA leverages ECC as public key cryptography algorithm because of its relatively small key and signature size. A 163-bit ECC key has roughly the same cryptographic strength as a 1024-bit RSA key or an 80-bit symmetric key. NPLA uses 164-bit public keys, 163-bit private keys and 326-bit signatures.

A. Header Processing Overhead and Latency

According to [19], when optimized for throughput, a FPGA implemented hardware ECC cryptography accelerator achieves 645,160 signature generations per second with a latency of 16.36 us per generation and 283,092 signature verifications per second with a latency of 24.28 us per verification. Under the assumption that the average number of AS hops is around 4 [18], the average end-to-end latency is 113.48us. This extra latency is much smaller than that of the Internet itself which is usually dozens of ms.

Based on the statistics of [20], we assume the average packet size is 6000-bit. Thus, a hardware cryptography accelerator can generate NPLA traffic at a speed of 3.87 Gbps and verify it at a speed of 1.7 Gbps on average. Prefix routers only need to insert NPLA headers for outgoing packets originated within its network not for transit and internal packets. 3.87 Gbps stamping throughput is sufficient for most edge networks. Verifying nodes need to authenticate the origin of the packets. 1.7 Gbps is fast enough for most of the edge networks. Yet it seems to be not efficient enough for the high-speed backbone routers. But we believe that in the early deployment phase legacy traffic occupies most of the traffic traversing the Internet. 1.7 Gbps may be able to handle the NPLA traffic. For the longer term consideration, according to the Moore's law, the speed of hardware public key operations

may satisfy the speed requirement in future.

B. Memory and Traffic Overhead

NPLA maintains the per-prefix public key information. According to RouteViews [12], there are less than 320K prefixes seen in BGP routing tables. The main memory overhead for each prefix is the storage of two 164-bit public keys. Thus, an upgraded border router needs around 13 MB extra memory which is not a big cost for the modern routers. The size of the NPLA header is fixed, 41 bytes. If we assume the average packet size is 750-byte (6000 bits) [20]. The average traffic overhead is about 5% for IPV4 traffic and would be smaller for IPV6 traffic. Such overhead is not critical for the rapid growth of the available bandwidth. We believe the security benefits and the saved bandwidth by discarding spoofed traffic justify this cost.

VII. Security Analysis

NPLA is a self-protective scheme. A prefix network implementing it can protect its address space from being spoofed by other networks. If an attacker within other networks wants to spoof the addresses belonging to a NPLA enabled prefix, one way to learn the private key is to launch a brute force attack. However, the 163-bit ECC public key cryptography on Koblitz Curves is the heart of NPLA. It is impossible to be cracked in short time with low cost.

The adoption of cryptography based defense systems, like NPLA, opens a door for cryptography based computation attacks where attackers may send significant amount of packets using randomly generated cryptography marks with low cost. One possible solution is to use load balancing mechanism, while it can only mitigate the effect but not solve it. The readers should know that until now there is no effective method to prevent the large scale bandwidth flooding attacks in the Internet. NPLA does not induce new security concerns.

VIII. Conclusion

In this paper, we propose NPLA an alternative method to mitigate source spoofing attacks on the network prefix level. Previously, the overhead of public key cryptography and the distribution of public keys have been identified as the major obstacles. We propose to use the FPGA implemented accelerator to decrease the overhead and leverage routing systems to distribute the public keys. NPLA is a self-protected approach. Each prefix implementing NPLA can prevent its address space from being spoofed by other networks. It also supports incremental deployment. The NPLA enabled nodes and legacy nodes can coexist in the Internet. In practice, we cannot expect an approach to be fully deployed. Under this assumption, the effectiveness of NPLA and other approaches are analyzed. We find that when partially deployed, NPLA prevents more spoofed traffic from consuming the resources of the Internet than other approaches do.

References

[1] S. Kandula, D. Katabi, M. Jacob, and A. Berger, .Botz-4-Sale: Surviving Organized DDoS Attacks that Mimic Flash Crowds,. in NSDI 2005

[2] T. Ehrenkranz and J. Li 2009. On the state of IP spoofing defense. ACM Trans. Internet Technol. 9, 2 (May. 2009), 1-29.

[3] 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.

[4] R. Beverly, A. Berger, Y. Hyun, and K. claffy. Understanding the efficacy of deployed internet source address validation filtering. In Proceedings of the 9th ACM SIGCOMM Conference on internet Measurement Conference, 2009.

[5] S. Kent and K. Seo. Security architecture for the Internet Protocol. RFC 4301, The Internet Engineering Task Force, December 2005.

[6] C. Candolin, J. Lundberg, and H. Kari. Packet level authentication in military networks. In Proc. 6th Australian Information Warfare & IT Security Conf., Geelong, Australia, Nov. 2005.

[7] X. Yang, D. Wetherall, and T. Anderson. A DoS-limiting network architecture. In Proc. ACM SIGCOMM, Aug. 2005.

[8] A. Bremler-Barr and H. Levy. Spoofing prevention method. In Proceedings of the Annual Joint Conference of the IEEE Computer and Communications Societies. InfoCom 2005.

[9] X. Liu, X. Yang, D. Wetherall, and A. Li. Passport: Secure and Adoptable Source Authentication. In Proceedings of the 5th USENIX NSDI, April 2008.

[10] Z. Duan, X. Yuan, and J. Chandrashekar. Constructing inter-domain packet filters to control IP spoofing based on BGP updates. In InfoCom 2006: Proceedings of the Annual Joint Conference of the IEEE Computer and Communications Societies.

[11] P. Ferguson and D. Senie. Network Ingress Filtering: Defeating Denial of Service Attacks which Employ IP Source Address Spoofing. RFC 2827, May 2000.

[12] RouteViewsProject. http://www.routeviews.org/.

[13] Y. Rekhter, T. Li, and S. Hares. A Border Gateway Protocol 4 (BGP-4). RFC 4271, 2006.

[14] A. Menezes. Elliptic Curve Cryptosystems. CryptoBytes, Vol.1 No.2, Summer 1995.

[15] D.Huang, A.Sinha and D.Medhi. "A key distribution scheme for double authentication in link state routing protocol". In IPCCC 2005. Proceedings of the 24th IEEE International Performance, Computing, and Communications Conference.

[16] S. Kent, C. Lynn, J. Mikkelson, and K. Seo. Secure border gateway protocol(S-BGP)-real world performance and deployment issues. In Proc. NDSS, 2000.

[17] P. V. Oorschot, T. Wan, and E. Kranakis. 2007. On interdomain routing security and pretty secure BGP (psBGP). ACM Trans. Inf. Syst. Secur. 10, 3 (Jul. 2007), 11.

[18] D. Magoni, and J. Pansiot. 2001. Analysis of the autonomous system network topology. SIGCOMM Comput. Commun. Rev. 31, 3 (Jul. 2001), 26-37.

[19] K. Javinen, and J. Skytt 6. High-Speed Elliptic Curve Cryptography Accelerator for Koblitz Curves. In FCCM 2008: Proceedings of the 16th IEEE Symposium on Fieldprogrammable Custom Computing Machines.

[20] W. John and S. Tafvelin. Analysis of internet backbone traffic and header anomalies observed. In IMC'07: Proceedings of the 7th ACM SIGCOMM conferenceon Internet measurement, pages111–116, NewYork, NY, USA, 2007. ACM