On the Deployability of Inter-AS Spoofing Defenses

Bingyang Liu and Jun Bi

Abstract

IP spoofing makes network attacks more destructive and harder to prevent. AS spoofing defenses mitigate these attacks by enforcing AS-level source address validity. Although many inter-AS defenses have been proposed, none is sufficiently deployed by Internet service providers. In this article we investigate the deployability of the defenses by evaluating their deployment benefit, deployment cost, and operational risk. Evaluation results reveal the technical features of highly deployable defenses, which can help improve existing defenses and the design of future defenses.

Distributed Denial of Service (DDoS) attacks have long been the top operational threat to ISPs. Using IP spoofing, attackers can hide the origin of attack traffic and amplify attack volume. There are two types of spoofing based DDoS attacks: d-DDoS and s-DDoS.[1] In d-DDoS, an attacker controls agents (also known as botnets) to send flooding packets to the victim with arbitrary source addresses for anonymity. In s-DDoS, agents use the victim's source address to send requests to innocent destination hosts for reflection, as the responses will be sent to the victim. Since responses are usually much larger than the requests in size, s-DDoS can be used for traffic amplification.

As the best current practices (BCPs) for spoofing defense, ingress/egress filtering (IEF) [1] and variants of unicast reverse path forwarding (uRPF) [2] are widely implemented in routers. Measurement reveals that about 60.5 percent of ASes on the Internet are unspoofable due to the adoption of IEF and uRPF [3]. However, in recent years spoofing attacks still frequently occur and repeatedly break the world record on network attack volume [4, 5], which indicates that the current deployment of BCPs is insufficient. Furthermore, measurement also shows that the deployment ratio has not been improved in four years, because many ISPs do not have incentives to deploy them [6]. So far, many new defenses have been proposed, but none of them can attract broad deployment in the real world. This background motivates us to study the deployability of spoofing defenses.

An inter-AS spoofing defense is deployed on AS border routers to filter spoofing traffic by enforcing AS-granularity source address validity. In this article we focus on inter-AS spoofing defenses for two reasons. First, despite being coarse-grained, AS-granularity source address validity is cost-effective in preventing attacks [7]. Second, ASes are autonomous units on the Internet, which allows us to analyze their autonomous economic interest and deployment behavior. According to our survey with the operators of ISPs, ASes are concerned about three properties of a defense in a real-world deployment [8]:

- Deployment benefit: by deploying a defense, an AS should gain significant additional protection, that is, the additional reduction of the spoofing traffic attacking the AS should be significant.
- Deployment cost: the deployment investment should be low.
- Operational risk: a defense should avoid dropping genuine packets, that is, the false positive rate (FPR) should be minimized.

In this article we survey existing inter-AS spoofing defenses and evaluate them against these metrics. We observe several technical features of highly deployable defenses based on the evaluation results, and then verify the observations with an example. Finally, we summarize the conclusions and discuss future work.

Inter-AS Spoofing Defenses

We first summarize basic filtering principles and inter-AS collaboration modes on which existing defenses are based. Then we briefly introduce seven major inter-AS spoofing defenses.

Filtering Principles

There are three filtering principles: end based, end-to-end (e2e) based, and path based. A defense can be based on one or multiple filtering principles.

End: End based filtering regards the Internet as two parts: the local end (AS) and the outside world (Fig. 1b). IEF is an end based defense. It drops the inbound packets whose source addresses belong to the local AS (ingress filtering), and the outbound packets whose source addresses do not belong to it (egress filtering).

E2e: E2e based filtering regards the Internet as individual ends (ASes) and ignores the connections between ends. E2e

The authors are with Institute for Network Sciences and Cyberspace, Tsinghua University, and Tsinghua National Laboratory for Information Science and Technology (TNList). The corresponding author is Jun Bi.

The authors would like to acknowledge the support from the National High-tech R&D Program ("863" Program) of China (No.2013AA013505) and the National Science Foundation of China (No.61472213).

[1] "d-" indicates that the victim is destination addresses of spoofing packets; "s-" indicates that the victim is the source addresses.

0890-8044/15/$25.00 © 2015 IEEE

based defenses require collaboration between ASes. Conceptually, each ordered pair of ASes s and d shares a secret key. The source AS s stamps a tag generated with the key into every outbound packet from s to d, that is, (s, d), and the destination AS d verifies the tag in the inbound packet (s, d). For example, in Fig. 1c assume that d, s, w, and y have deployed an e2e based defense, but x and z have not. From d's view, s, w, and y are distinct ends, since they are associated with distinct keys. The other ASes are unknown ends, since d has no information about them. Therefore, a deployer AS cannot be spoofed by other ASes if the destination is d.

Path: Path based filtering identifies a spoofing packet by verifying the packet's forwarding path. Although an attacker can forge a packet's source address, they cannot forge its forwarding path. Particularly, for a same packet (s, d), when it is originated from different locations, the forwarding paths would also be different. Denoted by a: (s, d) the packet (s, d) originated from a, a: (s, d) is a genuine packet if s is equal to a; otherwise it is a spoofing packet. For example, in Fig. 1d a spoofing packet x: (s, d) follows a different path from the genuine packet s: (s, d). From d's view, the genuine packet should come from the upstream AS z. Thus, the spoofing packet x: (s, d), which comes from an invalid upstream AS, can be identified and dropped by d.

Inter-AS Collaboration

In some defenses, ASes collaboratively share filtering information or help each other filter spoofing traffic. For example, e2e based defenses share secret keys, and path based defenses can share forwarding path information, so that deployers can construct more accurate filtering tables and enhance filtering ability. There are two collaboration modes: *source mode* and *destination mode*. In source/destination mode, a deployer has greater ability to identify a spoofing packet if its source/destination address belongs to another deployer, or drops an identified spoofing packet only if its source/destination address belongs to another deployer. We will show that collaborative defenses offer better protection to deployers than non-deployers.

Brief Introduction to Inter-AS Spoofing Defenses

We will introduce seven major inter-AS spoofing defenses. Some of them, such as IEF and uRPF, were not designed to work in an inter-AS scenario, but their methodology can be easily applied to this scenario. Some well known spoofing defenses are not discussed in this article because their filtering effectiveness is tightly bound to IP-level properties such as IP hop counts, which makes it difficult to accurately analyze them at the AS-level granularity. The properties of these defenses are summarized in Table 1.

IEF: IEF is an end based defense. It is applied on AS border routers to block the inbound packets whose source addresses belong to the local AS, and the outbound packets whose source addresses do not belong to it. To apply IEF, network operators only need to know the IP prefixes owned by the local AS and the incoming interfaces of routers for inbound and outbound traffic. Hence, IEF does not require inter-AS collaboration.

DPF: Distributed packet filtering (DPF) is a path based defense [7]. DPF introduces a path based filter, which associates each source AS s to a set of valid upstream ASes, that is, the ASes from which genuine packets sent by s can be received. The authors prove that if the filter is complete and accurate (i.e. a perfect filter), DPF can be very effective in spoofing defenses. However, how to construct perfect filters is not specified, which is extensively studied by other path based defense proposals.

uRPF: uRPF reversely uses a router's forwarding table as the path based filter, that is, it takes the next hop ASes towards AS s in the forwarding table as the valid upstream ASes associated with s. The advantages of uRPF are that it does not require a separate filtering table and the filter can be locally constructed. However, the disadvantage is that it may drop genuine packets when forwarding paths are asymmetric.

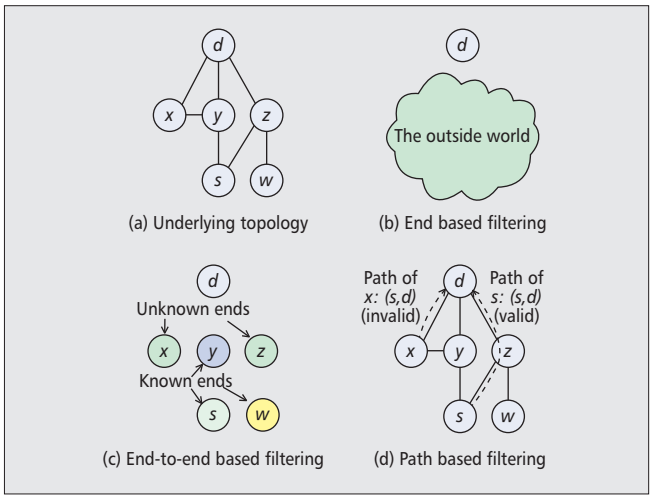

Figure 1. Given the same underlying topology, different filtering principles have different views of the Internet for spoofing identification.

IDPF: Inter-domain packet filter (IDPF) [9] constructs imperfect filters by inferring feasible paths for every source AS, where a feasible path is a path that conforms to a set of route export rules. The authors prove that u is a feasible upstream AS of v for source s, if and only if u has exported to v via Border Gateway Protocol (BGP) its best route towards s. Thus, an IDPF deployer can independently infer feasible paths by monitoring BGP updates. However, feasible paths may involve some paths that are not taken by genuine packets. For example, in Fig. 1d both $s \rightarrow d$ and $s \rightarrow z \rightarrow d$ are feasible paths, but only $s \rightarrow z \rightarrow d$ is taken by the genuine packets. Therefore, the spoofing packet y: (s, d) cannot be identified by d.

SAVE: Source Address Validity Enforcement (SAVE) [10] is a protocol that enables ASes to collaboratively construct perfect path based filters. A SAVE-enabled source AS s periodically generates updates toward all the other ASes. Since the updates follow the same paths as ordinary data packets, all the SAVE deployers along the paths can learn the valid upstream ASes for s. Only if s deploys SAVE can other deployers learn the filter entries for the source s. Thus, the collaboration is in source mode.

SPM: Spoofing prevention method (SPM) [11] employs e2e based verification. Each ordered pair of SPM deployers, s and d, shares a secret tag tag(s, d). An outbound packet s: (s, d) is stamped with tag(s, d) by the border router of s, and this tag is verified by the border router of d. If the tag is valid, the packet is forwarded; otherwise it is dropped. Since a destination deployer knows how to filter for a source AS only if the source AS also deploys SPM and shares tags with it, SPM employs source-mode collaboration. For example, in Fig. 1c d can filter out w: (s, d), since the packet does not carry tag(s, d). SPM deployers also enforce collaborative egress filtering (end based). Only if the destination of an outbound packet is a deployer will egress filtering be enforced on the packet to prevent the destination deployer from receiving spoofing packets. Thus, SPM also employs destination-mode collaboration.

Passport: Packet passport [12] deployers collaborate to filter spoofing packets in "destination + source" mode using the end based egress filtering and the e2e tag verification similar

Defense	Filtering principle	Collaboration mode
IEF	End	n/a
DPF	Path	n/a
uRPF	Path	n/a
IDPF	Path	n/a
SAVE	Path	Source
SPM	End + e2e	Destination + source
Passport	End + e2e	Destination + source

Table 1. Defenses' filtering principles and inter-AS collaboration modes.

to SPM, while the main difference is that Passport uses per packet cryptographic tags generated with the shared secret keys. A drawback of e2e based defenses is that some spoofing packets can only be identified when they reach the destination. On the path before reaching the destination, a spoofing packet consumes the same network resources as a genuine packet. Passport seeks to enable intermediate ASes on the path to identify spoofing packets, but rather than dropping the spoofing packets, intermediate ASes only demote them in a low-priority queue to avoid false positive (FP).

Deployability Evaluation

This section presents an evaluation of the defenses based on deployment benefit, deployment cost, and operational risk. Some evaluation results on deployment cost and operational risk have been published in [13]; we summarize and analyze them here. For the evaluation of deployment benefit, we perform comprehensive Internet-scale simulations, and will describe them in more detail.

Deployment Benefit

To incentivize a non-deployer AS to deploy a defense, the defense must provide the AS with sufficient additional protection, that is, the additional reduction of the spoofing traffic attacking the AS. Suppose that when the set of deployers is D the amount of reduced spoofing traffic attacking a non-deployer n is $red(D, n)$, and the amount becomes $red(D \cup \{n\}, n)$ when n becomes a deployer. Thus, the deployment benefit of n can be calculated: $ben(D, n) = red(D \cup \{n\}, n) - red(D, n)$. Since there is no analytical solution for every defense, we use simulations to compute their deployment benefits.

Simulation Settings: We obtain the real Internet AS-level topology from the UCLA Internet topology collection project and use C-BGP to construct inter-AS routing. Then we implement the filtering algorithms of the seven defenses. Given n as the victim, we generate 250000 random spoofing flows, and calculate $ben(D, n)$ as the ratio of the flows that can be filtered by $D \cup \{n\}$ but cannot be filtered by D with the filtering algorithms. For each flow a: (s, d), we set that all IP addresses have the uniform probability to be a, s, and d (unless s or d is the victim). Under another setting where the probability uniformly distributes among all ASes, the results are similar to the former setting. For space reasons, we only present the results of the former setting.

In each simulation we initially set $D = \phi$, and in each round we expand D by selecting some new deployers and compute $ben(D)$ as the average $ben(D, n)$ across all $n \notin D$. Some defenses may have optimal deployment strategies, where the optimal set of new deployers is selected in each round to maximize $ben(D)$. However, since the optimal strategies for some defenses are unavailable or computationally infeasible, we randomly select new deployers for a fair comparison between defenses.

As the number of deployers increases, $ben(D)$ changes, so we can draw a curve representing how the deployment benefit changes with the deployment ratio. We run the simulation 50 times and draw the averaged curves. Figures 2 and 3 show the evaluation results on deployment benefit for d-DDoS and s-DDoS, respectively. Since Passport has the same packet dropping behavior as SPM, they have the same curves. Thus, we omit Passport in these figures.

Deployment Benefit against d-DDoS: IEF's deployment benefit against d-DDoS is close to 0 since its ability to filter inbound spoofing traffic is very weak; indeed, only the packets whose source addresses belong to the local AS can be identified. Because the local address space is small compared with the global address space, the gained additional protection is marginal.

The benefit curve of DPF starts high since an AS can independently filter a lot of inbound spoofing traffic using DPF's perfect filter. However, as the deployment ratio grows, more and more spoofing traffic is filtered out by existing deployers, making it less and less attractive for a non-deployer to deploy DPF on its own since the additional reduction becomes less.

uRPF's curve is similar to DPF but is higher, which is unexpected since DPF employs perfect filters while uRPF infers imperfect filters. The reason is that uRPF achieves high benefit at the risk of high FPR. As a path based defense, DPF cannot identify the spoofing packets that are on the same path of the genuine packets. However, if the path is asymmetric, uRPF will drop all the packets, including the spoofing ones as well as the genuine ones. Since the deployment benefit criterion only assesses spoofing packets, the curve of uRPF is higher than DPF.

The curve of IDPF is similar to DPF but is much lower, because some invalid paths, which can be identified by DPF's perfect filters, are treated as feasible paths by IDPF.

Unlike other path based defenses, SAVE's curve starts from 0 and grows at the early deployment stage, because SAVE's filter is constructed collaboratively, that is, only if a source AS deploys SAVE can other deployers learn how to filter for it. The more deployers there are, the more inbound spoofing traffic can be filtered out. However, the curve stops growing at the later deployment stage because more and more spoofing traffic is filtered out by existing deployers.

Similar to SAVE, SPM's filter is also constructed collaboratively and the benefit curve also starts from 0. The main difference is that SPM employs destination-mode collaboration. Only if the destination AS of an outbound packet is a deployer does the source deployer enforce egress filtering on the packet. Thus, the more deployers there are, the less spoofing traffic is received by a deployer. On the other hand, non-deployers cannot benefit from deployers' egress filtering. Thus, if a non-deployer becomes a deployer, the additional protection is significant, and the benefit monotonically grows with the deployment ratio.

Deployment Benefit against s-DDoS: The deployment benefit of IEF, DPF, uRPF, and IDPF is close to 0 because these defenses are non-collaborative. In s-DDoS, spoofing packets are sent from attackers to innocent destinations, while the victim AS has a very low chance to be on the packets' forwarding paths to filter them. Thus, without other ASes' collaboration,

an AS cannot improve its protection against s-DDoS much by deploying these defenses.

The benefit curve of SAVE monotonically increases since it employs source-mode collaboration. If an AS does not deploy SAVE, other deployers cannot learn how to filter for its source address space. If it becomes a deployer, all other deployers can learn the filter for it and will filter out the spoofing packets whose source addresses belong to it. Thus, the more deployers there are, the higher deployment benefit is produced.

The benefit curve of SPM grows at the early deployment stage due to its source-mode collaboration. If an AS deploys SPM and shares tags with other deployers, other deployers will learn how to filter for its source address space. The more deployers there are, the more additional protection can be gained. However, at the later deployment stage, most spoofing traffic is filtered out by the egress filtering at source ASes, so it becomes less necessary for a non-deployer to deploy SPM on its own.

Recap: Based on the simulation results, we assess the deployment benefit of IEF as marginal since its benefit is close to 0 for both d-DDoS and s-DDoS. DPF, uRPF, and IDPF are assessed as low since they provide some benefit for d-DDoS but little benefit for s-DDoS. SAVE is assessed as medium since it provides low benefit for d-DDoS but high benefit for s-DDoS. SPM is assessed as medium-high since it provides high benefit for d-DDoS and some benefit for s-DDoS.

Deployment Cost

Deployment cost is the expense of deploying a defense. Previous work has investigated the computation, storage, and network overhead of the defenses [13]. However, to convince ISPs to deploy a defense, one must convert technical overhead metrics into economic expenses. In this article we propose to classify the defenses' deployment cost into three levels according to the required update or upgrade on routers. At the low level, only configuration updates and maintenance are required to implement the defense. At the medium level, the software or firmware of routers needs to be upgraded to implement new functionalities required by the defense. At the high level, the hardware of the router needs to be upgraded or replaced to fulfill the defense's resource demands. The three levels can be easily mapped to economic expense levels. Configuration only involves ordinary operation and maintenance. The medium level requires vendors to develop new software or firmware and ISPs to upgrade and reboot routers. The high level requires ISPs to purchase new hardware, which is generally expensive. Next, we analyze the defenses' deployment cost based on the evaluation results of the previous work.

IEF is recommended as BCP by the Internet Engineering Task Force (IETF). To apply IEF, an AS only needs to maintain the local IP prefixes, configure access control lists (ACLs) on routers, and apply the ACLs on proper interfaces of the routers. Hence the deployment cost of IEF is low.

The cost of DPF is unknown since it does not specify how to construct and store the filter in routers.

The deployment cost of uRPF is low since its functionality is widely implemented in routers. Some variant of uRPF, for example, feasible RPF, is implemented using equal-cost multipath (ECMP), an existing function on modern routers. Thus, only configuration is required to apply uRPF.

IDPF maps a source prefix to a set of feasible previous hops. The mapping is learned from BGP updates, which requires software updates of routers. However, the bottleneck is on storage. The authors in [13] show that if the average

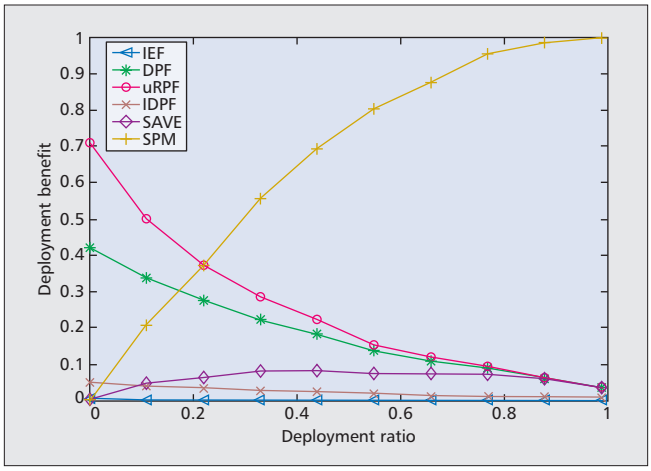

Figure 2. Defenses' deployment benefit against d-DDoS.

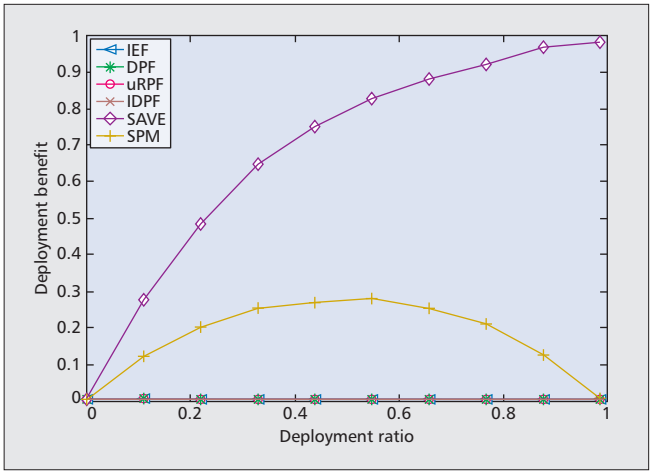

Figure 3. Defenses' deployment benefit against s-DDoS.

number of AS neighbors is 280, and each neighbor's feasibility is indicated by a single bit, the required space for a source prefix is 35 bytes. However, in the worst case, the largest AS on the Internet (Level 3) has 3621 neighbors, which requires 453 bytes for each prefix. For line-speed packet processing, the filter must be implemented in the data plane, and typically stored in static random access memory. Without further optimization or approximation, the table size will exceed the storage resources of existing routers, which means that hardware upgrade is required. Thus, the cost of IDPF is at the high level.

SAVE introduces a new protocol to learn filtering tables. Software must be updated to support SAVE. By optimizing data structures and algorithms, the authors prove that the computation, storage, and network overhead can be handled by existing routers. Thus, SAVE's deployment cost is medium.

SPM has a passive way to learn tags, which does not consume additional network bandwidth. Packet sizes are not enlarged since SPM tags are embedded into the IPID field of IP packets. A following work of SPM [14] implements the mechanism in existing routers by merely upgrading the firmware. Thus, the cost of SPM is medium.

Passport requires per packet cryptographic computation, making it impossible to perform high-speed packet processing without new hardware. Therefore, the deployment cost of Passport is high.

Defense	Deployment benefit	Deployment cost	Operational risk
IEF	Marginal	Low	Minimal
DPF	Low	Unknown	Unknown
uRPF	Low	Low	High
IDPF	Low	High	Medium
SAVE	Medium	Medium	Medium
SPM	Medium-high	Medium	Minimal
Passport	Medium-high	High	Minimal

Table 2. Defenses' deployability evaluation results.

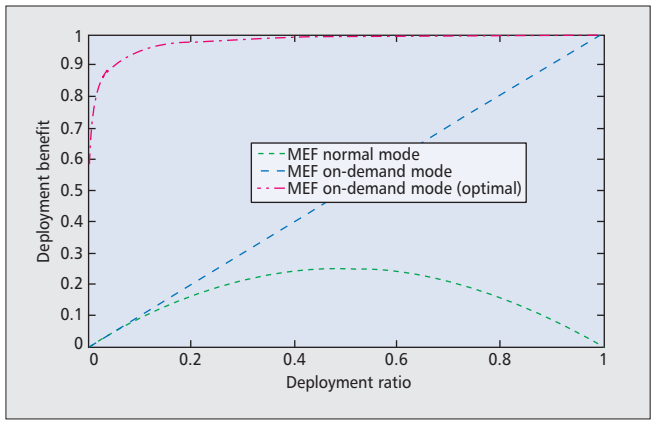

Figure 4. MEF's deployment benefit (d-DDoS and s-DDoS).

Operational Risk

The defenses prevent spoofing based attacks by dropping spoofing packets. If a defense improperly drops genuine packets, customers' network services may be degraded, and ISPs may have the risk to incur economic penalty. Next, we analyze the defenses' FPR.

To correctly operate IEF, operators need to maintain the local IP prefixes and the router interfaces for inbound and outbound traffic, and keep them up-to-date. The authors in [13] show that the FPR of IEF can be eliminated with careful operation. Thus, the risk of IEF is minimal.

Since DPF does not specify how to construct the filters, the FPR cannot be determined. Thus, the risk of DPF is unknown.

uRPF may drop genuine packets under path asymmetry. For example, in Fig. 1d, if y is the next hop toward s in d's forwarding table, d will deny the genuine packet s: (s, d) from z. Since path asymmetry is common on the current Internet, the risk of uRPF is high.

IDPF assumes that BGP updates comply with a set of route export rules. However, it is observed that at least 4 percent of the updates violate the rules. If the rules are violated, a source AS may send traffic via a path without announcing BGP updates via it. As a result, downstream ASes cannot learn this path for the source AS, which causes FP. Since not all violations cause FP, the FPR should be lower than 4 percent. Thus, the risk of IDPF is medium.

When SAVE is partially deployed, protocol updates will not be generated upon a routing change if the change occurs at a legacy AS. Under this condition, a new forwarding path may be used without updating the filtering tables of downstream ASes, so that genuine packets on the new path are dropped. The FPR can be high at the early deployment stage when legacy ASes dominate, but will decrease with the incremental deployment. Thus, the risk of SAVE is assessed as medium.

SPM employs egress filtering and e2e tag verification. The FPR of egress filtering is minimal. Without implementation and operation flaws, the FPR of e2e tag verification should be 0 [13]. Thus, the risk of SPM is minimal.

Passport has the same packet dropping behavior as SPM, while the path based verification only degrades the packets' priority. Thus, the risk of Passport should be equal to SPM, that is, minimal.

Evaluation Results Summary

The evaluation results are summarized in Table 2. It can be seen that collaborative defenses (SAVE, SPM, and Passport) have higher deployment benefit than the non-collaborative defenses. There are two reasons for this phenomenon. First, in collaborative defenses, deployers share the information for identifying spoofing traffic, so their filtering ability becomes higher. Second, collaborative defenses have more information to filter the traffic attacking deployers than non-deployers, or even filter only for deployers, which makes deployers gain much more protection than non-deployers and increases deployment benefit.

IEF and uRPF have low deployment cost since they are widely implemented in routers. IDPF and Passport have high cost since their storage and computation demands exceed the capacity of current routers, respectively. The cost of SAVE and SPM is medium since only software or firmware upgrades are required.

End or e2e based defenses have low operational risk. All the information required for filtering is locally available or shared across deployers in an e2e manner, so that its correctness does not depend on other non-deployers. On the other hand, path based defenses have higher risk. The required filtering information either is inferred (uRPF and IDPF) or depends on intermediate non-deployers (SAVE), which makes the risk hard to control.

Improving the Deployability of the Defenses

Table 2 shows that none of the studied defenses have satisfactory results for all the three evaluation metrics. To improve their deployability, one can seek to:

- Promote inter-AS collaboration to enhance deployment benefit.
- Use existing techniques on routers to reduce deployment cost.
- Use end and e2e based filtering principles to minimize operational risk.

As the first attempt toward this, mutual egress filtering (MEF) applies "destination + source" mode inter-AS collaboration to egress filtering to enhance the deployment benefit [15], that is, a spoofing packet identified by egress filtering is dropped only if its source or destination address belongs to a deployer. MEF also has a fine-grained on-demand filtering mode, in which filtering is enforced only on the traffic related to the victim. In this way, the operational risk (if any) can be restricted within the traffic related to the victim, while unconcerned traffic remains unaffected. Like IEF, MEF is implemented with ACLs on routers, so its cost is low.

Figure 4 shows the benefit curves of MEF. We already know that the benefit of IEF is close to 0. With inter-AS collaboration, the curve of MEF normal mode is much higher. The benefit of the on-demand mode monotonically increases, and the curve grows very fast at the early deployment stage under the optimal deployment strategy, where the non-

deployer who sends the most spoofing traffic is selected to be the new deployer in each round.

MEF verifies that employing inter-AS collaboration is an efficient approach to improve deployment benefit, and fine-grained collaboration can restrict operational risk.

Conclusion and Future Work

In this article we have surveyed seven inter-AS spoofing defenses, and analyzed their deployability by evaluating them against the metrics related to deployers' economic interest, including deployment benefit, deployment cost, and operational risk. The results have shown that collaborative defenses have higher deployment benefit, and end and e2e based defenses have lower operational risk. We believe that these observations can aid the design of future defenses. As a verification, we have shown that the deployability of egress filtering can be significantly enhanced by the employment of inter-AS collaboration.

With the aid of these observations, future work can exploit the potential room for improvement in deployability of existing defenses, or design new defenses with high deployability. Beyond the scope of this article, future work can study the policy concerns in inter-AS collaboration, as well as the performance and security issues introduced by the defenses.

References

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

[2] F. Baker and P. Savola, "Ingress Filtering for Multihomed Networks," RFC 3704, Mar. 2004.

[3] R. Beverly *et al.*, "Spoofer Project," retrieved Apr. 2014; http://spoofer.cmand.org/

[4] M. Prince, "The DDoS that Almost Broke the Internet," Mar. 2013; http://blog.cloudflare.com/the-ddos-that-almost-broke-the-internet

[5] B. Zigterman, "Largest DDoS Attack Ever Nails Company that Protects Against DDoS Attacks," Feb. 2014; http://bgr.com/2014/02/11/largest-ddos-attack/

[6] R. Beverly, A. Berger, and Y. Hyun, "Understanding the Efficacy of Deployed Internet Source Address Validation Filtering," *Proc. ACM SIGCOMM IMC*, Nov. 2009, pp. 356–69.

[7] K. Park and H. Lee, "On the Effectiveness of Route-based Packet Filtering for Distributed DoS Attack Prevention in Power-law Internets," *Proc. ACM SIGCOMM*, Aug. 2001, pp. 15–26.

[8] J. Bi and B. Liu, "Problem Statement of SAVI beyond the First Hop," Internet Draft, Nov. 2012.

[9] Z. Duan, X. Yuan, and J. Chandrashekar, "Controlling IP Spoofing through Interdomain Packet Filters," *IEEE Trans. Dependable Secure Computing*, vol. 5, no. 1, 2008, pp. 22–36.

[10] J. Li *et al.*, "SAVE: Source Address Validity Enforcement Protocol," *Proc. INFOCOM*, Jun. 2002, pp. 1557–66.

[11] A. Bremler-Barr and H. Levy, "Spoofing Prevention Method," *Proc. INFOCOM*, Mar. 2005, pp. 536–47.

[12] X. Liu *et al.*, "Passport: Secure and Adoptable Source Authentication," *Proc. NSDI*, Apr. 2008, pp. 365–78.

[13] J. Mirkovic and E. Kissel, "Comparative Evaluation of Spoofing Defenses," *IEEE Trans. Dependable Secure Computing*, vol. 8, no. 2, 2011, pp. 218–32.

[14] B. Liu, J. Bi, and Y. Zhu, "Deployable Approach for Inter-AS Anti-spoofing," *Proc. ICNP*, Nov. 2011, pp. 19–24.

[15] B. Liu, J. Bi, and A. Vasilakos, "Towards Incentivizing Anti-spoofing Deployment," *IEEE Trans. Inf. Forens. Security*, vol. 9, no. 3, 2014, pp. 436–50.

Biographies

BINGYANG LIU (liubingyang@tsinghua.edu.cn) received a B.S. degree in computer software from Tsinghua University, China. He was a joint Ph.D. student in the Department of Computer Science, Duke University. He received a Ph.D. degree in computer science from Tsinghua University, China. He is now a postdoctoral researcher at the Institute for Network Sciences and Cyberspace, Tsinghua University. His research fields include Internet architecture, DDoS defense, and software-defined networking (SDN).

JUN BI (junbi@tsinghua.edu.cn) received B.S., M.S., and Ph.D. degrees in computer science from Tsinghua University, China. He was a postdoctoral scholar at Bell Laboratories Research and a research scientist at Bell Labs, USA. Currently he is a full professor and director of the Network Architecture Research Division, Institute for Network Sciences and Cyberspace at Tsinghua University, and a key member of Tsinghua National Laboratory for Information Science and Technology (TNList). His research interests include Internet architecture and protocols. He has successfully lead tens of government supported or international collaboration research projects, published more than 100 research papers and 20 Internet RFCs or drafts (four of them were approved), owned 20 innovation patents, and received national science and technology advancement prizes. He is co-chair of the AsiaFI (Asia Future Internet Forum) Steering Group, and co-founder of the China SDN Commission and serves as executive chair. He served as co-chairs of Workshops/Tracks at INFOCOM, ICNP, Mobihoc, ICCCN, etc., and served on the Organization Committees or Technical Program Committees at SIGCOMM, ICNP, CoNEXT, SOSR/HotSDN, etc. He is a senior member of IEEE, a senior member of ACM, and a distinguished member of CCF (China Computer Federation).