Journal of Computing and Information Technology - CIT 16, 2008, 2, 109–118
doi:10.2498/cit.1001131

Internet Deployment of DPM-based IP Traceback

Andrey Belenky and Nirwan Ansari

Advanced Networking Laboratory, ECE Department, NJIT, Newark, USA

[1] In this article, we introduce the Internet deployment guidelines for Deterministic Packet Marking (DPM) – a novel IP traceback method. Unlike other packet marking schemes, DPM cannot be deployed sporadically on the Internet. Therefore, in order to perform the traceback, a structured way of deployment is needed. Related to topology and deployment issues, discussion comparing the features of other full path schemes and ingress packet filtering to those of DPM is also presented.

Keywords: IP Traceback, DPM, topology.

1. Introduction

In recent years much interest and consideration has been paid to the topic of securing the Internet infrastructure that continues to become a medium for a broad range of transactions. A number of approaches to security have been proposed, each attempting to mitigate a specific set of concerns. After several high-profile Distributed Denial of Service (DDoS) attacks on major US web sites in 2000, numerous IP traceback approaches [1], [2] have been suggested to identify the attacker(s). The previously proposed schemes can be categorized in two broad groups. One group of the solutions relies on the routers in the network to send their identities to the destinations of certain packets, either encoding this information directly in rarely used bits of the IP header, or by generating a new packet to the same destination. The biggest limitation of this type of solutions is that they are focused only on flood-based DoS and DDoS attacks, and cannot handle attacks comprised of a small number of packets. Additionally, for the large scale DDoS attacks, these schemes are not very effective. The second group of solutions involves centralized management and logging of packet information on the network. Solutions of this type introduce a large overhead, and are complex and not scalable.

Deterministic Packet Marking (DPM), a new approach to IP traceback, was recently introduced to mitigate some of the above shortcomings. Even though DPM is classified as a scheme of the first type, the substantial differences of having only edge routers perform the marking allow DPM to perform traceback with only a few packets from the attacker and be capable of tracing thousands of attackers simultaneously.

In all previous discussions of DPM, the Internet was considered to be a homogeneous network with a single administration and perfect DPM deployment on the edges of the network. In reality, the Internet is not, and has never been, a single homogeneous network according to [3]–[5]. The Internet, from its very beginning, was an interconnection of different networks. Over time, the Internet became more structured, but its heterogeneous nature still remains. In this article, the structure of the Internet is described and the relationships among the Autonomous Systems (ASs) are explored. Then, the simple guidelines for deploying DPM are introduced and analyzed.

The rest of the paper is structured as follows. Section 2 introduces the basic DPM approach;

[1] Please address all correspondence to: Prof. Nirwan Ansari, Electrical and Computer Engineering Dept., New Jersey Inst. of Technology, Newark, NJ07102, USA. Phone/fax: +1-973-596-3670; email: ansari@njit.edu. This work was supported in part by the National Science Foundation under grant no. 0726549.

Section 3 describes the structure of the Internet and AS interrelations; Section 4 describes the perfect case of DPM deployment; Section 5 introduces the guidelines for deploying DPM on the Internet followed by an illustrative example described in Section 6; the discussion and comparison of DPM to other schemes is presented in Section 7; we conclude in Section 8.

2. DPM Principles

The basic DPM is a packet marking algorithm, which was first introduced in [6]. Subsequently, the issues of fragmentation and simultaneous multiple attackers were addressed in [7]. This section provides the general principle behind DPM.

2.1. Assumptions

The assumptions in this section were largely borrowed from [8]. The two key assumptions driving this effort are:

- An attacker may generate any packet
- Routers are both CPU and memory limited

2.2. DPM Principle

As mentioned above, DPM is a packet marking algorithm [7]. The 16-bit packet ID field and 1-bit Reserved Flag (RF) in the IP header will be used to mark packets. *Each* packet is marked when it enters the network. This mark remains unchanged for as long as the packet traverses the network. This automatically removes the issue of mark spoofing which other marking schemes have to account for. The packet is marked by the interface closest to the source of the packet on an edge ingress router, as seen in Figure 1. The routers with engraved 'DPM' have DPM enabled by configuring interfaces, and the rubber-stamps signify the interfaces on these routers that actually perform the marking. The mark is a partial address information of this interface. The interface makes a distinction between incoming and outgoing packets. Incoming packets are marked; outgoing packets are not marked. This ensures that the egress interface will not overwrite the mark in a packet placed by an ingress router.

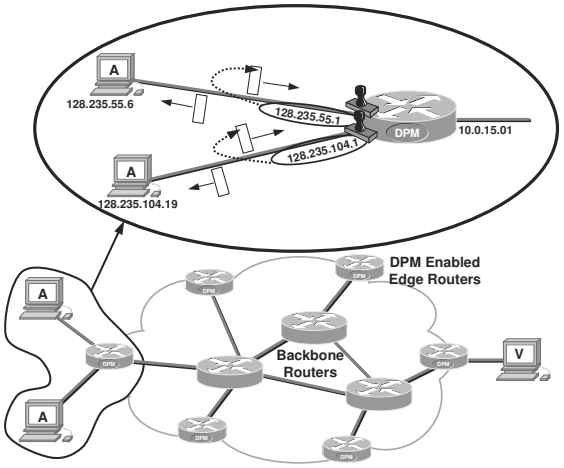

Figure 1. Deterministic Packet Marking.

In certain instances, a physical interface may be an ingress interface to a customer and at the same time it may carry traffic other than to or from this customer. In such cases, logical or virtual interfaces, such as VLANs, may be used to provide necessary granularity to DPM configuration on the DPM-enabled router.

For illustrative purposes, assume that the Internet is a network with a single administration. In this case, only interfaces closest to the customers on the edge routers will participate in packet marking. Every incoming packet will be marked. Should an attacker attempt to spoof the mark in order to deceive the victim, this spoofed mark will be overwritten with a correct mark by the very first router the packet traverses.

3. Structure of the Internet

The Internet is a hierarchical structure [3], [4]. Several ISPs, so called tier 1 ISPs, constitute the backbone of the Internet. Very few ISPs, such as UUNET, SprintLink, AboveNet, GBLX, AT&T and a few others, have the status of tier 1. The recent visualization of the Internet, where the hierarchy can be observed, is available on the web site of CAIDA [9]. These ISPs have a large geographical presence to facilitate access to a greater number of customers. Having the large presence facilitates a convenient connection of other ISPs. The ISPs of the second tier do not have the geographical presence comparable to

tier 1 ISPs. Therefore, in order to establish connectivity to the other parts of the world, tier 2 ISPs have to buy transit services from one or more tier 1 ISPs. The next tier of ISPs have even less geographical presence. Yet, these ISPs are not at the lowest levels of hierarchy. These ISPs have to buy transit services from the upper tier ISPs in order to establish global connectivity. These ISPs of medium tiers are often called regional ISPs. Finally, there are lowest tier ISPs, which sell transit services only to the retail customers such as home users and businesses. It is also suggested in [10] that the ISP's IP networks themselves have a hierarchical structure with well defined *core* links and *edge* links.

An administrative domain, such as an ISP or an enterprise network, may consist of more than one AS according to [11]. For example, when two organizations, each with its own AS, merge, the resulting network will be a single administrative domain encompassing two ASs. The connectivity among the ASs is provided by the Border Gateway Protocol (BGP), version 4 [12]. BPG is used for the exchange of routes and for the selection of routes based on the predefined policies. When using BGP, an AS advertises certain routes to its neighbors. The commercial relationships between the ASs define the rules, also called *policies*, of route advertising. In the rest of this section, relationships between ASs classified in [11], [13], which do not always fit the purely hierarchical structure will be examined. Figure 2 illustrates all of the relationships described later in this section.

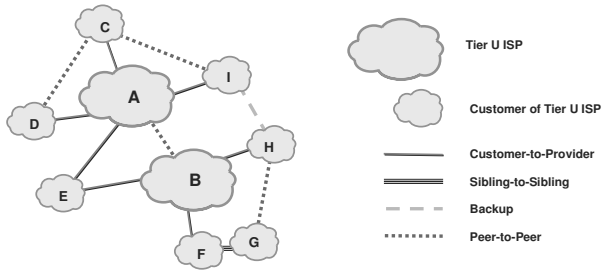

Figure 2. Illustrations of inter AS relationships.

3.1. Customer-to-provider Relationship

In the customer-provider relationship, the provider, usually an AS of a higher tier ISP, advertises all known routes to the customer, an AS of the lower tier ISP. The customer pays for this *transit service* and advertises its own routes and the routes of its customers to the provider. The provider in one relationship may be a customer in another relationship, and the customer in one relationship may be a provider in another relationship.

For increased availability, a given customer may enter a customer-provider relationship with more than one provider. Such an arrangement is called *multi-homing*. The end customer, such as a local ISP or an enterprise, that maintains customer-provider relationship with a single provider is called a *stub*. In a multi-homed arrangement, the customer does not advertise the providers' routes. Otherwise, the customer's network may be used as a transit between its providers.

It should also be noted that the customers do not necessarily have to connect only to the local, lowest tier, ISPs only. ISPs of high tiers offer retail services as well, and the end customers may purchase transit services from them directly.

In Figure 2, the ASs **C**, **D**, **E**, and **I** have the customer-to-provider relationship with AS **A**, and ASs **E**, **H**, and **F** have the customer-to-provider relationship with AS **B**. Autonomous system **E** has the customer-to-provider relationship with both ASs **A** and **B**, and, therefore, AS **E** is multi-homed.

3.2. Peer-to-peer Relationship

Two ISPs may find it mutually beneficial to exchange their routes to their customers directly. This is called a peering arrangement and is usually free of charge to both parties since it usually provides equal benefits to both ISPs. In order to establish the peering arrangement, the respective ASs have to enter a peer-to-peer relationship. In the peer-to-peer relationship, the routes to the ASs' own customers are usually advertised to the peer AS. The peering arrangements are not associative, meaning that two peers of a given AS are not peers of each other, unless they setup an explicit peer-to-peer relationship between each other.

Peer-to-peer relationships are established usually for economic reasons. Instead of sending traffic through the transit network, it is more

efficient and cost effective to send the traffic directly to the peer. This reduces the amount of traffic which is sent to the provider, as well as offers better service between the customers of the two peers.

There are several instances of peer-to-peer relationship in Figure 2. Autonomous systems **A** and **B** have a peer-to-peer relationship. Therefore, the customers of **A** would be able to communicate with the customers of **B** without using the providers of **A** and **B** (not shown in Figure 2). Autonomous system **C** has a peer-to-peer relationship established with **D** and with **I**. It means that the traffic from customers of **C** will be able to reach the customers of **D** and **I** without traversing **A**. However, the traffic from customers of AS **I** will have to traverse AS **A** in order to reach customers of AS **D**, and vice-versa. There is also a peer-to-peer relationship between ASs **G** and **H**.

3.3. Sibling-to-sibling Relationship

In the sibling-to-sibling relationship, the two ASs exchange the routes to their customers, peers and providers. In other words, the two sibling ASs share all routes. This arrangement is usually used for the ASs which belong to a single administrative domain. Also, when two stubs cannot afford a transit service by themselves, they may effectively combine their networks by entering the sibling-to-sibling relationship and use a single connection to the provider, provided the bandwidth requirements are met, in order to save money.

Autonomous systems **F** and **G** in Figure 2 have a sibling-to-sibling relationship. Since they exchange all of their available routes, customers of **G** will have access to the rest of the Internet. Similarly, although **F** does not have an explicit peer-to-peer relationship with **H**, the peer-to-peer relationship between them will exist implicitly since **G** would share routes to the **F**'s customers with **H**, and **H**'s routes with **F**.

3.4. Backup Relationship

Two ASs may have a backup relationship, given that they have different providers. If the provider of one of the ASs in the backup relationship fails, then this AS will use its backup AS as a transit network to the Internet. The purpose of this arrangement is similar to that of the multi-homing described in Section 3.1. However, instead of actually purchasing transit service from two providers, the customer purchases the transit service from one provider and enters a mutual backup relationship with another AS, which is cheaper.

Autonomous systems **H** and **I** in Figure 2 have a backup relationship. This means that if for some reasons, AS **B** would become unavailable to **H**, it would be able to maintain the connectivity to the rest of the Internet through ASs **I** and **A**. Similarly, if AS **I** would loose its connectivity to **A**, it would still maintain the connectivity to the rest of the Internet through ASs **H** and **B**. The customers do not necessarily have to connect only to the local ISPs only. In reality, as seen in Figure 3, the customers may connect directly to the ISPs of the high tiers, where **B** has two retail customers.

4. Ideal DPM Deployment

The ideal DPM deployment is a situation, where every interface connecting to the customers of all ISPs in the Internet would be DPM-enabled. This situation is depicted in Figure 3.

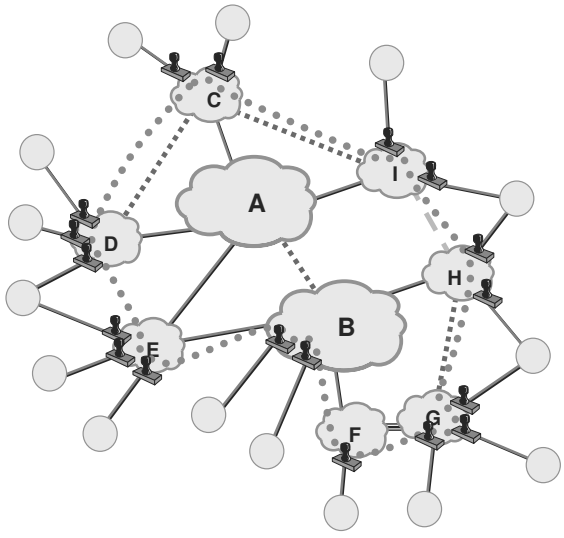

Figure 3. Ideal DPM deployment.

The collection of DPM-enabled interfaces is called a *DPM perimeter*. The perimeter must not have any *holes*, the access points through which the traffic from a customer may traverse

the Internet unmarked. The DPM perimeter is shown in Figure 3 as a thick dotted line with circular dots.

5. Guidelines for DPM Deployment

It cannot be expected that all ISPs will deploy DPM simultaneously, if at all. Also, even if that is to happen, DPM has to be disabled on some interfaces, as it is being enabled on the others to maintain the DPM perimeter. Therefore, some coordinated effort on behalf of the ISPs is needed.

Neighboring ASs are defined as two ASs, which have at least one external BGP (eBGP) session setup between them. It is possible to have more than one eBGP session between a pair of ASs. Also, it is necessary to make a natural assumption that an ISP is not acting maliciously. In other words, if the ISP claims to have deployed DPM, then it can be trusted that DPM is in fact deployed and it is deployed according to the guidelines described below.

Tier 1 ISPs must have DPM enabled on all edge interfaces of all their ASs, except those which face other ASs of tier 1 ISPs. This can be accomplished by the good will of these ISPs or by forcing them to do so. Currently, all of the tier 1 ISPs are in developed countries, thus making application of some international treaties possible.

After the Internet's core has been DPM-enabled, the rest of the ISPs can join the DPM scheme gradually. Only the ASs, whose providers have deployed DPM, may enable DPM on its edges. Once the AS of the lower tier enables DPM on its edges, this fact has to be communicated to the provider. The provider must then disable DPM on the interface to this customer.

If there are two neighboring ASs, which are in either sibling-to-sibling, peer-to-peer, or backup relationship and both of them have DPM enabled, then both should disable DPM on the interfaces facing each other. This requires that ISPs share the information about DPM deployment, but it is assumed that administrations of ASs, which have established any relationship, would have the means and incentives to communicate this information to each other.

There are rare instances when a tier-1 ISP has to use a lower tier ISP to tunnel the traffic of its customers to another tier-1 ISP. Such cases are rare, and should they occur, the lower tier ISP should preferably be included in the perimeter, or alternatively, the first tier-1 ISP should be excluded from the perimeter. In either case, the DPM perimeter will remain continuous.

6. Illustrative Example

Consider the network introduced in Figure 2 discussed so far. For illustrative purposes, it is assumed that ASs **A** and **B** are tier 1 ISPs and the rest of the ASs are tier 2. As discussed earlier, **A** and **B** have to deploy DPM before the rest of the ASs; then the following sequence of AS deployment is considered: **F**, **H**, **C**, **E**, **D**, **I**, and **G**.

Autonomous systems belonging to the tier 1 ISPs must deploy DPM first on their edge interfaces. Since AS **A** and AS **B** have a peer-to-peer relationship and both of them have DPM enabled, the DPM must be disabled on the interfaces between them, as can be seen in Figure 4.

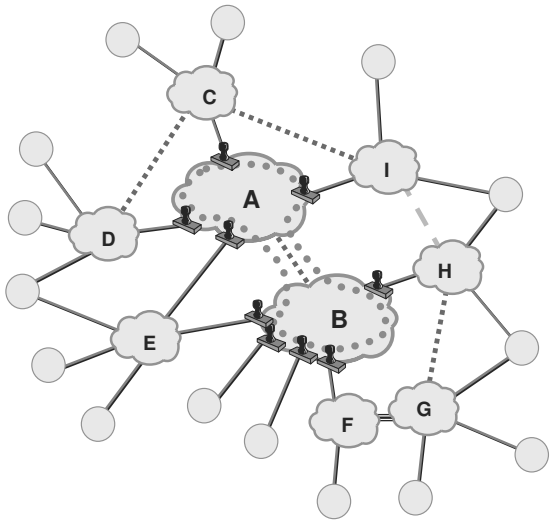

Figure 4. DPM deployment example, DPM enabled on tier 1 ISPs **A** and **B**.

Next, DPM is deployed on the edge interfaces of the first tier 2 ISP **F** as seen in Figure 5. Autonomous system **F** has a single interface to **B** and since both of them will have DPM enabled, **B** has to disable DPM on its interface to **F**. If **B** does not disable DPM on the interface to **F**, then the attack traffic from the **F**'s customers will be

marked by the edge interface of **F**, and then those marks will be overwritten by the marks of **B**. So, effectively, the traceback will be possible only to the edge of **B**. Autonomous system **F** has a sibling-to-sibling relationship with AS **G**. Since **G**, at this moment, had not deployed DPM on its edges, **F** has to deploy DPM on the interface to **G**.

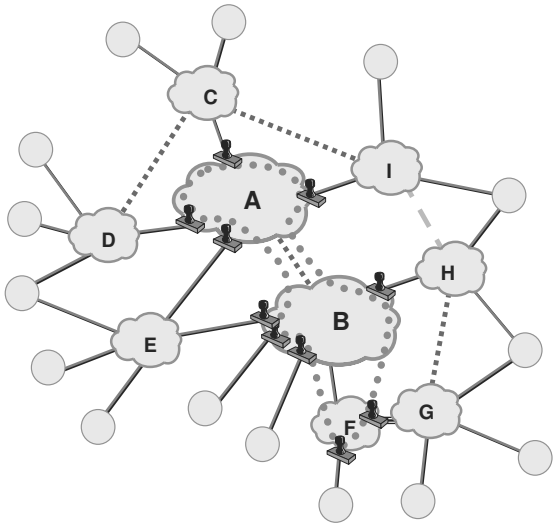

Figure 5. DPM deployment example, DPM enabled on AS **F**.

Deployment of ASs **H**, **C**, and **E** happens in exactly the same way as deployment of AS **F** by enabling DPM on their edges and disabling DPM on the interfaces between them and tier 1 ISPs. Figure 6 depicts the network after DPM on these three ISPs has been deployed.

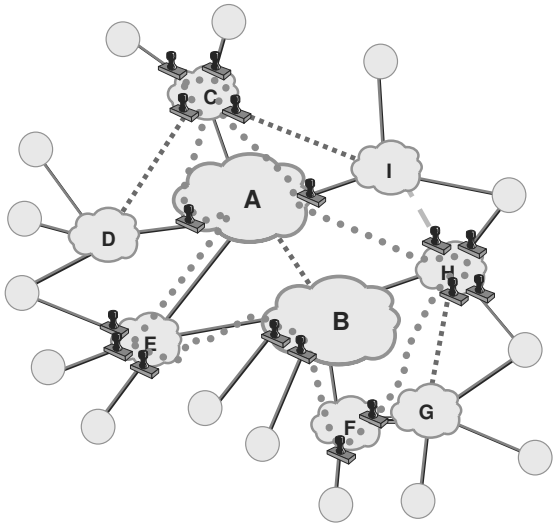

Figure 6. DPM deployment example, DPM enabled on AS **E**.

When AS **D** enables DPM on its edges, not only AS **A**, but also its peer, AS **C**, has to disable DPM on the interfaces to **D** because at that point both **C** and **D** will have DPM deployed as shown in Figure 7.

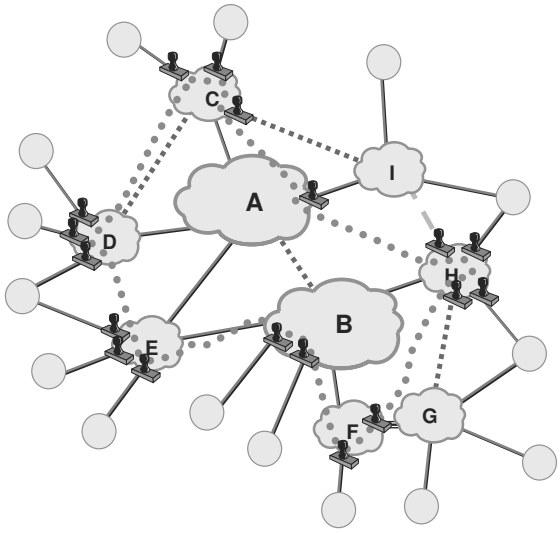

Figure 7. DPM deployment example, DPM enabled on AS **D**.

Similarly, when AS **I** enables DPM on its edge interfaces, the interfaces to **C** and **H** would not need to have DPM enabled on them since the respective ASs have DPM enabled. Also, the DPM would have to be disabled on **C**'s and **H**'s interfaces to **I**. Finally, when DPM is enabled on the edges of AS **G**, it has to be disabled on the interfaces of its peer **H** and of its sibling **F**, which provides connectivity to **G**.

If it is assumed for illustrative purposes that the network described so far is a complete Internet, then ideal DPM deployment, shown in Figure 3 in Section 4, will result after AS **G** deploys DPM on its edges.

7. Discussion of DPM

Traditionally, all IP traceback schemes perform what is known as the full path traceback, where a complete path of the attack packets through the Internet is determined. Deterministic Packet Marking does not perform full path traceback. Only the closest to the source interface which belongs to the DPM perimeter is determined

during the traceback. In this section, the advantages of the DPM traceback versus the full path traceback will be addressed.

It can also be argued that the ingress address filtering, described in [14], is as effective as DPM if deployed around the same perimeter. This argument is addressed in this section as well.

Finally, some concluding remarks on DPM and traceback in general are made.

7.1. Comparison of DPM to Full Path Traceback Schemes

Deterministic Packet Marking can be viewed as a special case of Probabilistic Packet Marking (PPM) described in [8]. The differences are that marking happens only at the edges of the collection of the deployed networks and the probability of marking is 100%. So what is lost and what is gained by these changes?

Deterministic Packet Marking must be deployed according to the guidelines outlined in Section 5. This requires the synchronization of efforts on behalf of the ISPs. Probabilistic Packet Marking, on the other hand, can be deployed independently on every ISP. Consider the situation depicted in Figure 8, where ISP **O** of tier U does not deploy a traceback scheme, and ISPs **P** and **R** of tier $U + 1$ do deploy the traceback scheme. Also, the attack path from the attacker **A** to the victim **V** traverses the path **P-O-R**.

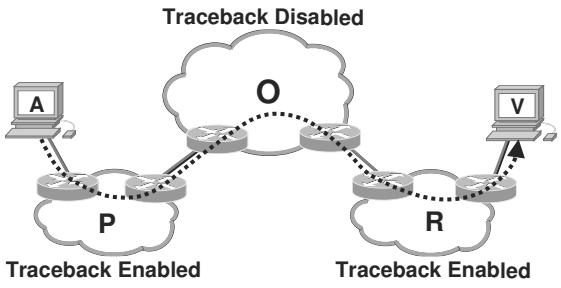

Figure 8. Example of partial traceback deployment.

If a full-path traceback scheme is deployed, the partial path encompassing the routers in **P** and **R** will be reconstructed. If DPM is deployed, the address of the **R**'s interface to **O** will be reconstructed. The marks of the ingress interface of **P** inserted in the packets will be overwritten by the marks of the **R**'s interface to **O**. That is exactly why this situation is not possible in DPM, if the deployment guidelines described in Section 5 are followed.

Deterministic Packet Marking cannot trace the attacks which were initiated from inside of the DPM perimeter, situation depicted in Figure 9. There is an unsubstantiated opinion that the attackers subvert one or more routers as a part of most attacks. In reality, subverting a router is a difficult task, usually possible only as a result of an improper router configuration. To get a feel for how vulnerable is the network equipment compared to the workstations, the Computer Emergency Response Team (CERT) vulnerability notes database [15] was examined. A vulnerability is a flaw in the system that can be used to take full or partial control over the system, or just bring it down. Vulnerability notes database is a collection of known vulnerabilities, which have been reported so far. For example, as of May 12th, 2003, there were 825 known vulnerabilities in the database. Only 31 (less than 4%) of them were attributed to the ISP grade network equipment, the rest to various software packages of different platforms, and some to home office or Local Area Network (LAN) network equipment. Of these 31, only a handful were severe enough to allow an attacker to take control over the device. The low percentage of ISP grade network equipment vulnerabilities may serve as an indication on how difficult it is to accomplish the situation depicted in Figure 9. It can also be concluded that in the vast majority of the attacks, the attack packets are generated by the workstations, and therefore are traceable by DPM.

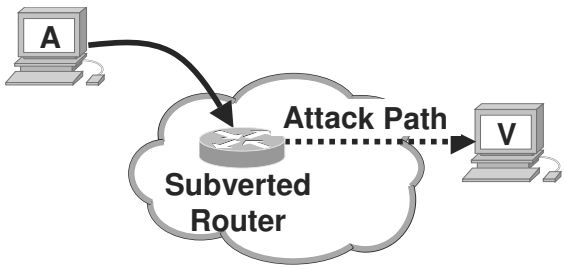

Figure 9. Situation of the subverted router inside an ISP.

There are many advantages of DPM over the full-path marking schemes. Security issues of PPM-like schemes arise from the fact that an attacker can inject a packet, which is marked with erroneous information. Such behavior is called

mark spoofing. Prevention of such behavior is accomplished by special coding techniques, and is not 100% proof. Deterministic packet marking ensures that every packet which arrives to the victim is correctly marked, and thus the need in those complex and processor intensive encoding techniques is unnecessary.

The following observation about all full-path traceback schemes is made: in a datagram packet network, the full-path traceback is as good as the address of an ingress point in terms of identifying the attacker. By definition, each packet in a datagram network is individually routed. Since every packet may take a different path from the source to the destination, only the ingress interface on the router closest to the source must be the same. Packets may take different routes even if their source and destination are identical. This may happen for two reasons: due to unwanted oscillation of the network routing, or due to desired bandwidth management techniques such as load balancing. The changes in the route between the source of the attack packets and the victim will be detrimental for the full-path traceback since more than one path for the single source would be reconstructed. This, however, does not affect DPM.

Internet service providers may only use public addresses for interfaces to customers and other networks, and use private addressing plans within their own networks. In this case, the usefulness of the full-path traceback becomes very low since the information produced for the most part cannot tell the victim much more than a few IP addresses on the borders between ISPs. Even if this is not the case, and public addressing is used within ISPs' networks, ISPs generally feel reluctant to disclose their topologies. Full-path traceback schemes reveal topology of all networks by design. To limit this undesirable behavior, only routers, whose addresses are already known, should implement such schemes.

7.2. Comparison of DPM to Ingress Address Filtering

Source address filtering is a mechanism of ensuring that only the packets with the valid SAs are entering the Internet. The range of valid addresses is set up manually on the ingress interface and usually corresponds to the range of IP addresses of the hosts, which are expected to connect to the Internet through this interface.

Ingress address filtering is usually associated with high processing overhead. It is usually a subset of a large filtering mechanism, which enables filtering of packets by many other fields of layer 3 and 4 headers in the packets. To perform this filtering, every packet has to be taken off the fast switching hardware-based path and be analyzed by software, thus drastically increasing the processing overhead incurred by the router for every packet. The processing overhead, however, is not what precludes ingress address filtering from becoming an effective protection against the DDoS attacks. Hardware-based mechanism, which would filter packets based on their SA only, would not be difficult to implement. In fact, the processing overhead incurred for every packet on the routers would be comparable to DPM, and none of the victim's processing would be necessary.

Unfortunately, ingress filtering is effective in preventing DDoS attacks only if it is performed on close to 100% of interfaces in the Internet according to [16]. In addition, the ingress filtering has to be constantly managed if it is not deployed on the edges of the network, and its effectiveness is severely degraded. Consider Figure 10 where ISP **O** of tier U does implement ingress address filtering, and ISPs **P** and **R** of tier $U + 1$ do not. The customer network connects to ISP **P**. Assume also that at some point in time ISP **O** is aware of the address ranges used by both lower tier ISPs and the customer. If the customer changes its provider from **P** to **R**, then the ranges on the interface performing the ingress filtering on the edges of **O** have to be reconfigured to allow the traffic from the customer through **O**. In case the customer

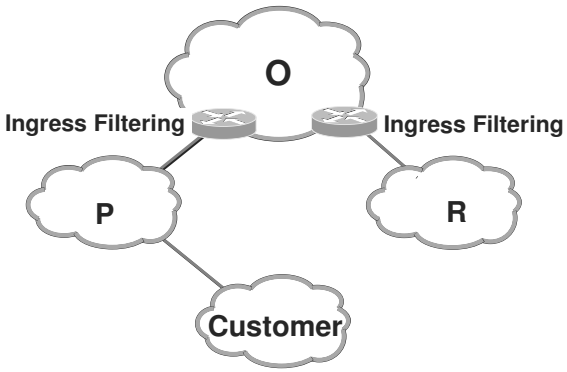

Figure 10. Limitations of ingress address filtering.

changes its IP address range, similar reconfiguration tasks would have to be performed. If the customer is multihomed, then **O** would have to accept traffic from customers with SAs of the interfaces from **R** and **P**. In other words, any workstation could spoof the source address of their packets to the ones of this customer, and this traffic would be allowed to pass.

Once the attacker finds a way for his or her packets with the spoofed SAs to pass through the ingress filtering, the attack cannot be stopped. DPM, however, even if deployed on the same network as the ingress address filtering, would provide the victim with the concrete IP address, where the attack traffic entered the DPM perimeter.

8. Conclusions

In this article, the structure of the Internet and AS interrelations were discussed. The simple guidelines for DPM deployment were introduced. Following those guidelines would ensure that DPM produces the best results possible for the given collection of ISPs deploying DPM. It was concluded that by following those deployment guidelines, DPM can perform the traceback as well as any other full-path traceback schemes.

References

[1] Z. GAO, N. ANSARI, Tracing Cyber Attacks from the Practical Perspective. *IEEE Communications Magazine*, **43**(5) (May 2005), pp. 123–131.

[2] A. BELENKY, N. ANSARI, On IP Traceback. *IEEE Communications Magazine*, **41**(7) (July 2003), pp. 142–153.

[3] G. HUSTON, Interconnection, Peering and Settlements – Part I. *The Internet Protocol Journal*, **2**(1) (March 1999), 2–17.

[4] ——, Interconnection, Peering and Settlements – Part II. *The Internet Protocol Journal*, **2**(2) (June 1999), 2–24.

[5] L. SUBRAMANIAN, S. AGARWAL, J. REXFORD, R. H. KATZ, Characterizing the internet hierarchy from multiple vantage points. In *Proceedings of INFOCOM 2002 Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies*, **2** (June 2002), pp. 618–627.

[6] A. BELENKY, N. ANSARI, IP traceback with deterministic packet marking. *IEEE Commun. Lett.*, **7**(4) (April 2003), pp. 162–164.

[7] ——, On deterministic packet marking. *Computer Networks*, **51**(10), (July 11, 2007), 2677–2700.

[8] S. SAVAGE, D. WETHERALL, A. KARLIN, T. ANDERSON, Network support for IP traceback. *IEEE/ACM Trans. Networking*, **9**(3) (June 2001), pp. 226–237.

[9] CAIDA, AS internet graphy. Apr. 2002. [Online]. Available: `http://www.caida.org/analysis/topology/as_core_network/pics/ascoreApr2002.gif`

[10] A. FELDMANN, J. REXFORD, IP network configuration for intradomain traffic engineering. *IEEE Network*, **15**(5) (Sept./Oct. 2001), pp. 46–57.

[11] L. GAO, On inferring autonomous system relationships in the internet. *IEEE/ACM Trans. Networking*, **9**(6) (December 2001), pp. 733–745.

[12] Y. REKHTER, T. LI, Border gateway protocol 4 (BGP-4). RFC 1771, March 1995.

[13] L. GAO, J. REXFORD, Stable internet routing without global coordination. *IEEE/ACM Trans. Networking*, **9**(6) (December 2001), pp. 681–692.

[14] P. FERGUSON, D. SENIE, Network ingress filtering: defeating denial of service attacks which employ IP source address spoofing. RFC 2827, May 2000.

[15] VULNERABILITY DATABASE, Available on-line, Cert, May 2003. [Online]. Available: `http://www.cert.com/`

[16] K. PARK, H. LEE, On the Effectiveness of Route-based Packet Filtering for Distributed DoS Attack Prevention in Power-law Internets. In *Proc. of the 2001 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, (August 2001) pp. 15–26.

Received: October, 2007
Accepted: November,2007

Contact addresses:
Andrey Belenky and Nirwan Ansari
Advanced Networking Laboratory
ECE Department, NJIT
Newark, NJ 07102
USA
Phone/Fax: +1-973-596-3670
e-mail: Nirwan.Ansari@NJIT.EDU

ANDREY BELENKY received the B.S.Comp.E. degree (summa cum laude) and the M.S. degree in telecommunication networks from Polytechnic University, Brooklyn, NY in 1998, and the Ph.D. degree from New Jersey Institute of Technology (NJIT), Newark, NJ in 2003. The main focus of his research was distributed denial of service attacks and IP Traceback. Prior to receiving the Ph.D. degree, he worked in Telcordia Technologies (formerly Bellcore) as a network engineer. He is currently a licensed patent attorney in NY.

NIRWAN ANSARI received the B.S.E.E. (summa cum laude) from the New Jersey Institute of Technology (NJIT), Newark, in 1982, the M.S.E.E. degree from University of Michigan, Ann Arbor, in 1983, and the Ph.D. degree from Purdue University, West Lafayette, IN, in 1988. He joined NJIT's Department of Electrical and Computer Engineering as an Assistant Professor in 1988, and has been a Full Professor since 1997. He has also assumed various administrative positions at NJIT. He authored Computational Intelligence for Optimization (Springer, 1997, translated into Chinese in 2000) with E.S.H. Hou, and edited Neural Networks in Telecommunications (Springer, 1994) with B. Yuhas. His current research focuses on various aspects of broadband networks and multimedia communications. He has also contributed over 300 technical papers, of which over one third in refereed journals and magazines. He is a Senior Technical Editor of the IEEE Communications Magazine, and also serves on the editorial board of Computer Communications, the ETRI Journal, and the Journal of Computing and Information Technology. He was the founding general chair of the First IEEE International Conference on Information Technology: Research and Education (ITRE2003), was instrumental, while serving as its Chapter Chair, in rejuvenating the North Jersey Chapter of the IEEE Communications Society which received the 1996 Chapter of the Year Award and a 2003 Chapter Achievement Award, served as Chair of the IEEE North Jersey Section and in the IEEE Region 1 Board of Governors during 2001-2002. He has also been serving in various IEEE committees such as Chair of IEEE COMSOC Technical Committee on Ad Hoc and Sensor Networks, and (TPC) Chair/Vice-chair of several conferences/symposia. His awards and recognitions include the NJIT Excellence Teaching Award in Graduate Instruction (1998), IEEE Region 1 Award (1999), an IEEE Leadership Award (2007, from IEEE Princeton/Central Jersey Section), and the designation as an IEEE Communications Society Distinguished Lecturer.