Inter Domain Packet Filters for IP Forging Attacks

P.V.Ravi kanth[#1]
M.Tech Pursuing
QIS College of Engg &Tech.,
Ongole, A.P, India.
pvravikanth555@gmail.com
+918019594355

M.Srinivasa Roa[#2]
M.Tech Pursuing
Kshtriya College of Engg.,
Hyderabad, A.P, India.
mr.srinu13@gmail.com
+919949965499

A. Ravi[#3]
Asst. Prof in IT
MLEC
Singarayakonda, A.P, India.
510ravi@gmail.com
+919849529095

Abstract*---The Distributed Denial-of-Service (DDoS) attack is a serious threat to the legitimate use of the Internet. Even Prevention mechanisms are attacked by the ability of attackers to forge or spoof the source addresses in IP packets. By employing IP spoofing, attackers can avoid detection and put a substantial burden on the destination network for policing attack packets. In this paper, we propose an Inter Domain Packet Filter (IDPF) architecture that can reduce the level of IP spoofing on the Internet. A key feature of our scheme is that it does not require global routing information. IDPF's are constructed from the information implicit in Border Gateway Protocol (BGP) route updates and are deployed in network border routers and IDPF's does not discard packets with valid source addresses. Here we show that, even with partial deployment on the Internet, IDPF's can proactively limit the spoofing capability of attackers. In addition, they can help localize the origin of an attack packet to a small number of candidate networks.

***Keywords*---IP spoofing, DDoS, BGP, network-level security and protection, routing protocols.**

1. INTRODUCTION

Distributed Denial of Service (DDoS) attacks pose an increasingly grave threat to the Internet, as evidenced by recent DDoS attacks mounted on both popular Internet sites and the Internet infrastructure Alarmingly, DDoS attacks are observed on a daily basis on most of the large backbone networks. One of the factors that complicate the mechanisms for policing such attacks is *IP spoofing*, the act of forging the source addresses in IP packets. By masquerading as a different host, an attacker can hide its actual identity and location, rendering source-based packet filtering less effective. It has been shown that a large part of the Internet is vulnerable to IP spoofing [3], [4].

It is our contention that IP spoofing will remain popular for a number of reasons. First, IP spoofing makes it harder to isolate attack traffic from legitimate traffic-packets with spoofed source addresses may appear to be from all around the Internet. Second, it presents the attacker with an easy way to insert a level of indirection, which shifts the burden to the victim. Substantial effort is required to localize the source of the attack traffic [2]. Finally, many popular attacks use IP spoofing and require the ability to forge source addresses. Man-in-the-middle attacks, such as variants of TCP hijack and DNS poisoning attacks are carried out by the attacker masquerading as the host at the other end of a valid transaction.

Inspired by the idea of route-based packet filters, we propose Inter-Domain Packet Filter (IDPF) architecture. The IDPF architecture takes advantage of the fact that while network connectivity may imply a large number of *potential* paths between source and destination domains, commercial relationships between ASes act to restrict to a much smaller set the number of *feasible* paths that can be used to carry traffic from the source to the destination. In this paper we focus our attention on the construction of IDPFs based solely on locally exchanged BGP updates. We will investigate how other AS relationship and routing information may help further improve the performance of IDPFs in our future work. We show that locally exchanged routing information between neighbors, i.e., BGP route updates, is sufficient to identify feasible paths and construct IDPFs. Like route-based packet filters, the proposed IDPFs cannot stop all spoofed packets. However, when spoofed packets are not filtered out, IDPFs can help localize the origin of attack packets to a small set of ASes, which can significantly improve the IP trace back situation.

We summarize the key contributions of this paper in the following:

1) We describe how to practically construct inter-domain packet filters *locally* at an AS by using only the BGP route updates being exchanged between the AS and its immediate neighbors.

2) Even with partial deployment, the architecture can proactively limit an attacker's ability to spoof packets. When a spoofed packet cannot be stopped, IDPFs can help localize the attacker to a small number of candidate ASes, reducing the effort and increasing the accuracy of IP trace back schemes.

3) We show that unlike some protection schemes that provide intangible local benefits for deployment, the IDPF architecture provides better protection against IP spoofing based DDoS attacks on local networks, which presents incentives for network operators to deploy IDPFs.

2. RELATED WORK

The idea of IDPF is motivated by the work carried out by Park and Lee [7], which was the first effort to evaluate the relationship between topology and the effectiveness of route-based packet filtering. The authors showed that packet filters that are constructed based on the *global* routing information can significantly limit IP spoofing when deployed in just a small number of ASes. In this work, we extend the idea and demonstrate that filters that are built based on *local* BGP updates can also be effective.

Unicast reverse path forwarding (uRPF) [1] requires that a packet is forwarded only when the interface that the packet arrives on is exactly the same used by the router to reach the source IP of the packet. If the interface does not match the packet is dropped. While simple, the scheme is limited given that Internet routing is inherently asymmetric, i.e., the forward and reverse paths between a pair of hosts are often quite different. In Hop-Count Filtering (HCF) [6], each end system maintains a mapping between IP address aggregates and valid hop counts from the origin to the end system. Packets that arrive with a different hop count are suspicious and are therefore discarded or marked for further processing. Bremler-Barr and Levy proposed a spoofing prevention method (SPM) [5], where packets exchanged between members of the SPM scheme carry an authentication key associated with the source and destination AS domains. Packets arriving at a destination with an invalid authentication key (w.r.t. the source) are spoofed packets and are discarded.

3. BORDER GATEWAY PROTOCOL AND AS CONNECTIONS

In this section, we briefly describe a few key aspects of BGP that are relevant to this paper. To begin with, we model the AS graph of the Internet as an *undirected* graph G = (V,E). Each node v□V corresponds to an Autonomous System (AS), and each edge e(u,v)□E represents a BGP session between two neighboring ASes u,v□V . To simplify the exposition, we assume that there is at most one edge between neighboring ASes.

Each node owns one or multiple network prefixes. Nodes exchange BGP route updates, which may be announcements or withdrawals, to learn of changes in reachability to destination network prefixes. A route withdrawal, containing a list of network prefixes, indicates that the sender of the withdrawal message can no longer reach the prefixes. In contrast, a route announcement indicates that the sender knows of a path to a network prefix. The route announcement contains a list of *route attributes* associated with the destination network.

3.1. POLICIES AND ROUTE SELECTION

Each node only selects and propagates to neighbors a single *best* route to the destination, if any. BGP is a *policy-based* routing protocol in that both the selection and the propagation of best routes are guided by locally defined routing policies. Two distinct sets of routing policies are normally employed by a node: *import* policies and *export* policies. Neighbor-specific import policies are applied upon routes learned from neighbors, whereas neighbor-specific export policies are imposed on locally-selected best routes before they are propagated to the neighbors.

In general, *import* policies can affect the .desirability. of routes by modifying route attributes. Let r be a route (to destination d) received at v from node u. We denote by import(v u)[{r}] the possibly modified route that has been *transformed* by the import policies. After the routes are passed through the import policies at node v, they are stored in v's routing table. The set of all such routes is denoted as candidate R(v,d):

Among the set of candidate routes candidate R (v,d), node v selects a single best route to reach the destination based on a well defined procedure.

$$\mathbf{candidateR}(v,d) = \{\mathrm{r} : \mathbf{import}(v \leftarrow u)[\{r\}] \neq \{\} \\ \mathrm{r.}\mathbf{prefix} = d, \forall u \in N(v)\}.$$

Here, $N(v)$ is the set of v's neighbors.

To aid in description, we shall denote the outcome of the selection procedure at node v, i.e., the best route, as best R (v,d), which reads *best route to destination* d *at node* v.

Having selected best R (v,d) from candidate R(v,d), v then exports the route to its neighbors after applying neighbor specific *export policies*. The export policies determine if a route should be forwarded to the neighbor, and if so, modify the route attributes according to the policies. We denote by export(v -> u) [{r}] the route sent to neighbor u by node v, after node v applies the export policies on route r.

3.2. AS RELATIONSHIPS AND ROUTING POLICIES

The specific routing policies that an AS employs internally are largely determined by economics: connections between ASes follow a few commercial relations. A pair of ASes can enter into one of the following arrangements:

Provider-customer: In this kind of arrangement, a customer AS pays the provider AS to carry its traffic to the rest of the Internet. This arrangement is the most common and is natural when the provider is much larger in size than the customer.

Peer-peer: In a mutual peering agreement, the ASes decide to carry traffic from each other (and their customers). This is only natural when the traffic from each other is roughly balanced. Mutual peers do not carry transit traffic for each other.

Sibling-sibling: In this type of arrangement, two ASes provide mutual transit service to each other (often as backup connectivity or for reasons of economy). Each of the two sibling ASes can be regarded as the provider of the other AS.

4. INTER DOMAIN PACKET FILTERS

In this section we discuss the intuition behind the IDPF architecture, describe how IDPFs are constructed using BGP route updates, and establish the correctness of IDPFs. After that, we discuss the case where ASes have routing policies that are less restrictive than r1-r4. We shall assume that the routing system is in the *stable routing state* in this section. We will discuss how IDPFs fare with network routing dynamics in the next section.

Let M(s,d) denote a packet whose source address is s (or more generally, the address belongs to network s), and destination address d. A packet filtering scheme decides whether a packet should be forwarded or dropped based on certain criteria. One example is the route-based packet filtering [7].

BGP is an incremental protocol: updates are generated only in response to network events. In the absence of any events, no route updates are triggered or exchanged between neighbors, and we say that the routing system is in a stable state.

Definition 1 (Stable Routing State): A routing system is in a stable state if all the nodes have selected a best route to reach other nodes and no route updates are generated (and propagated) by any node.

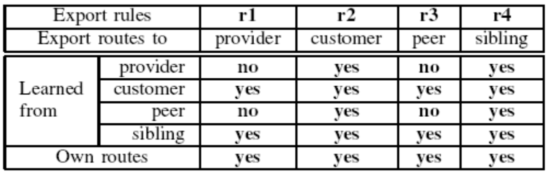

Export rules		r1	r2	r3	r4
Export routes to		provider	customer	peer	sibling
Learned from	provider	no	yes	no	yes
	customer	yes	yes	yes	yes
	peer	no	yes	no	yes
	sibling	yes	yes	yes	yes
Own routes		yes	yes	yes	yes

TABLE I
ROUTE EXPORT RULES AT AN AS

Consider the example in Fig 1, Fig 2(a) and (b) present the topological routes implied by network connectivity and feasible routes constrained by routing policies between source s and destination d, respectively. In Fig. 2(b) we assume that nodes a, b, c, and d have mutual peering relationship, and that a and b are providers to s.

We see that although there are 10 topological routes between source s and destination d, we only have 2 feasible routes that are supported by routing policies. Of more importance to IDPF is that, although network topology may imply all neighbors can forward a packet allegedly from a source to a node, feasible routes constrained by routing policies help limit the set of such neighbors.

For Example, let us consider the situation at node d. Given that only nodes a and b (but not c) are on the feasible routes from s to d as node d concerns , node d can infer that all packets forwarded by node c and allegedly from source s are spoofed and should be discarded.

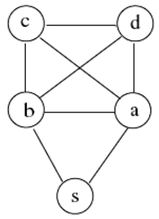

Fig. 1. An example network topology.

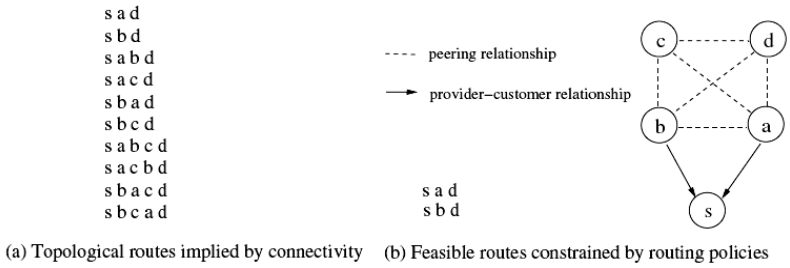

(a) Topological routes implied by connectivity (b) Feasible routes constrained by routing policies

Fig. 2. Routes between source s and destination d.

IDPF Architecture

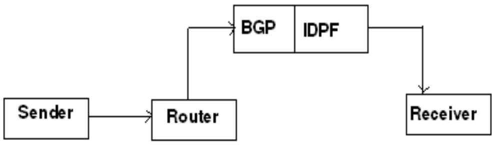

An IDPF may not be able to catch all spoofed packets forwarded by a neighbor. Note that an IDPF allows all the feasible upstream neighbors for packet M(s,d) to send the packet. However, in reality, exactly one of them will lie on best R(s,d) and forward M(s,d). On the other hand, it is worth noting that an attacker in a best upstream neighbor for packet M(s; d) can always spoof the source address s; therefore, route-based packet filters also cannot catch all spoofed packets. In the next section, we will conduct simulation studies to compare the performance of route-based packet filtering with that of the IDPF framework.

5. PERFORMANCE STUDIES

5.1 OBJECTIVES AND METRICS

We evaluate the effectiveness of IDPFs in controlling IP spoofing based DDoS attacks from two complementary perspectives.

1. We wish to understand how effective the IDPFs are in *proactively* limiting (if not preventing) the capability of an attacker to spoof addresses of ASes other than his own. Our approach does not provide complete protection and spoofed packets may still be transmitted.

2. Thus the complementary, *reactive* view is also important; we study how the deployed IDPFs can Improve IP trace back effectiveness by localizing the actual source of spoofed packets.

5.2 DATA SETS

In order to evaluate the effectiveness of IDPFs, we construct four AS graphs from the BGP data archived by the Oregon Route Views Project [33]. The rest three graphs, denotedG2003, G2004, and G2005 are constructed from single routing table snapshots (taken from the rest day in each of the years).While these provide an indication of the evolutionary trends in the growth of the Internet AS graph, they offer only a partial view of the existing connectivity [14]. In order to obtain a more comprehensive picture, similar to [34], we construct
G2004c by combining G2003 and an entire year of BGP updates between G2003 and G2004.

Results for G2004c with different IDPF node coverages

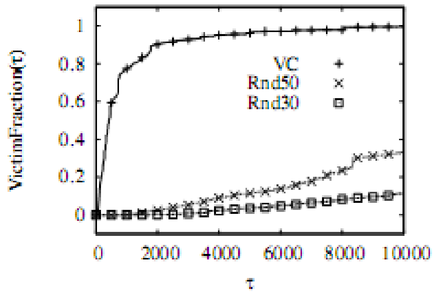

(a) $VictimFraction(\tau)$

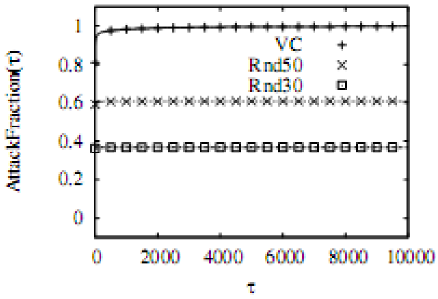

(b) $AttackFraction(\tau)$

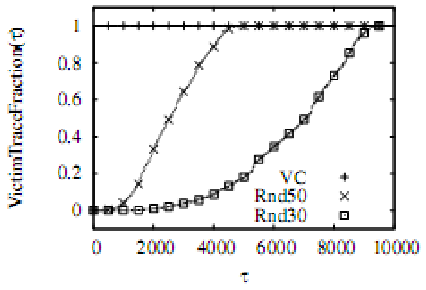

(c) $VictimTraceFraction(\tau)$

Results for G2003, G2004, G2004c, and G2005 with the VC coverage

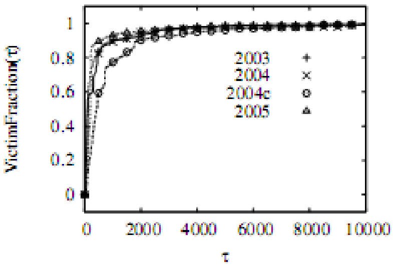

(a) $VictimFraction(\tau)$

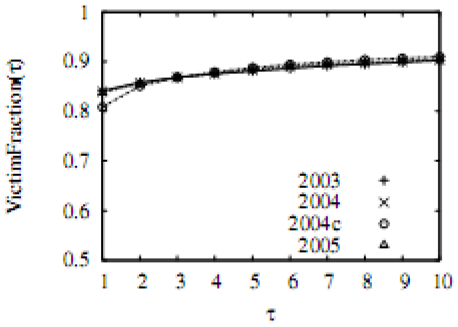

(b) $AttackFraction(\tau)$

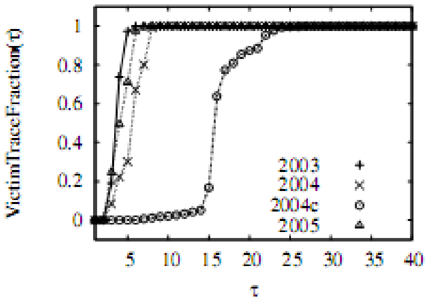

(c) $VictimTraceFraction(\tau)$

5.3 RESULTS OF PERFORMANCE STUDIES

The studies are performed with the Distributed Packet Filtering (dpf) simulation tool [12]. In addition, we also studied the impact of using BGP updates instead of precise routing information to construct packet filters, investigated the effect of overlapping prefixes in the Internet, and considered IDPFs with and without network ingress filtering. Before we describe the simulation results in detail, we briefly summarize the salient findings.

6. CONCLUSION

In this paper we proposed and studied an inter-domain packet filter (IDPF) architecture as an effective countermeasure to the IP spoofing-based DDoS attacks. IDPFs rely on BGP update messages exchanged between neighboring ASes on the Internet to infer the validity of source address of a packet forwarded by a neighbor. We showed that IDPFs can be easily deployed on the current BGP-based Internet routing architecture. Our simulation results showed that, even with partial deployment on the Internet, IDPFs can significantly limit the spoofing capability of attackers; moreover, they also help localize the actual origin of an attack packet to be within a small number of candidate networks. In addition, IDPFs also provide adequate local incentives for network operators to deploy them.

ACKNOWLEDGMENTS

This research was funded by a URP grant from Cisco Systems. The authors would like to thank David Ward of Cisco Systems for his support.

REFERENCES

1. F. Baker. Requirements for ip version 4 routers. RFC1812, June 1995.
2. S. Bellovin. ICMP trace back messages. Internet Draft,October 2001.Work in Progress.
3. R. Beverly. Spoofer project. http://momo.lcs.mit.edu/Spoofer
4. R. Beverly and S. Bauer. The Spoofer Project: Inferringthe extent of Internet source address filtering on the Internet. In *Proceedings of Usenix Steps to ReducingUnwanted Traffic on the Internet Workshop SRUTI'05* Cambridge, MA, July 2005.
5. A. Bremler-Barr and H. Levy. Spoofing prevention method. In *Proc. IEEE INFOCOM*, Florida, March,05.
6. C. Jin, H. Wang, and K. Shin. Hop-count-filtering: an Effective defense against spoofed DDoS traffic. In *Proceedings of the 10th ACM conference on Computer and communications security*, October 2003.
7. K. Park and H. Lee. On the effectiveness of route-basedPacket filtering for distributed DoS attack prevention

in Power-law Internets. In *Proc. ACM SIGCOMM*, San Diego, CA, August 2001