New Implementation of IP Traceback

Raghavendra.V*
Arokiaraj Jovith.A*
Godwin Ponsam.J*

Abstract

Distributed denial of service attacks currently represents a serious threat to the appropriate operation of Internet services. To overcome this threat we are proposing an overlay network with the help of ip tracebacking of autonomous systems. Our proposed system will not require a prior knowledge of the network for that incrementally deploying the system to different autonomous systems. In previous paper their contribution for two techniques those are new extension for the BGP update message and sequence marking process for the packets but in our proposed for the update message community attribute is used and marking process with the help of hashmapping. The false positives are reduced. In previous work only detection but in our proposed blocking and also if attacker is perform spoofing also blocking the system after tracebacking. We also tested with the help of java programming in incremental basis installing in the systems and tracing the system that is performing the attack. The main conclusion is that the proposed system is suitable for large-scale networks such as the Internet because it provides efficient traceback and allows incremental deployment.

Keywords: Ddos, BGP, Internet Services

1. Introduction

Identifying the sources of large-scale distributed denial of service (DDoS) attacks [1] is a challenging task because:

- IP routing is based solely on the destination IP address carried by each packet.
- IP packets are not authenticated at the moment they are forwarded, enabling spoofed source-IP addresses to be used in DDoS attacks.
- Attacker packets can be sent by zombie hosts (remotely controlled by an attacker), whose owners are unaware that they are participating in a DDoS attack.
- Usually, no information about packet forwarding is kept at the intermediate routers due to scalability restrictions.
- Identifying the sources of an attack does not imply that the attackers were found because they could be behind a firewall or protected by a private IP address, and thus the traceback indicates only the network middleboxes from where the packets are coming.

The combination of these issues contributes to the current scenario in which attackers have virtual guarantee of anonymity. In general, the defense against attacks requires three different steps:

- Intrusion detection, usually performed by intrusion

Department of Information Technology, SRM University, Chennai, Tamil Nadu, India.

detection and prevention systems
- The identification, at least partially, of the route(s) of attacker packets
- The filtering or blocking of attacker packets at key points along the routes.

We emphasize that our proposed system focuses on the Second step, that is, on the discovery of the route(s) followed by attacker packets. It is exactly this functionality that IP traceback systems are intended to provide. In contrast, we propose an IP-traceback system that could be partially deployed in large-scale networks such as the Internet. The traceback system operates on border routers of an autonomous system (AS), which after exchanging information carried by the Border Gateway Protocol (BGP), builds an AS-level overlay network for IP traceback. Our first contribution is thus a new extension to the BGP update-message community attribute that enables information to be passed across ASs that are not necessarily involved in the overlay network. The second contribution is a new sequence-marking process to remove ambiguities in the traceback path. Through a simulation study, we show that our system can be partially and incrementally deployed over the network, and it still provides good performance. This feature enables ASs willing to collaborate with the traceback effort to join other participating ASs, thus leading to an increased efficiency in IP traceback. Moreover, our findings indicate that having a relatively low number of Ass using the system as long as these ASs are selected strategically is sufficient to provide an efficient IP traceback at the AS level.

2. Related Work of IP Traceback

Several IP-traceback systems were proposed in recent years [2]. Snoeren et al. [3] introduce the source path isolation engine (SPIE), a log-based IP-traceback system that stores digests of packets inside Bloom filters [4] when packets are forwarded by routers. Bloom filters are space-efficient probabilistic data structures based on hash functions that typically are used to test whether an element is a member of a set. This system, based on the use of Bloom filters, can traceback a single IP packet. Laufer et al. [5] propose another use of Bloom filters; when packets cross a router, a mark is inserted into a generalized Bloom filter (GBF) [6] located in the IP header of each packet. Compared with a conventional Bloom filter, GBF limits the false positive probability at the expense of introducing the possibility of false negative events. The main advantage of GBF is that both the false positive and the false negative ratios are upper bounded. Furthermore, these upper bounds depend only on the chosen parameters of the filter and not on its initial condition. Therefore, when a packet reaches its destination, the packet carries the mark of all traversed routers in the GBF. To traceback a packet, the victim verifies which of its neighbor routers has its mark in the resulting GBF and then sends a reconstruction packet to it. This procedure is repeated at each router until the reconstruction packet arrives at the router from where the attacker packets are coming. Finally, to finish the traceback, this router sends a message back to the victim with the discovered route.

Analyzing IP-traceback systems proposed so far, we observe that none of them can be adopted effectively in a large-scale network such as the Internet because of the:
- Increased network overhead both on the intermediate routers and on the victim at the moment the traceback process takes place
- Limited scalability
- Requirement to purchase new specialized devices
- Requirement of authentication mechanisms
- Requirement of previous knowledge about network topology

Moreover, we emphasize that such systems require the deployment of their respective solutions in all routers of the monitored network domain, thus contributing to the unlikely deployment in the Internet and limiting their efficacy against large-scale DDoS attacks. We believe that deploying a traceback system in all routers is not required to enable an efficient IP traceback. Rather, it suffices to identify some key points in the path where attacker packets are being forwarded to enable efficient countermeasures to be taken in a distributed way to block the ongoing attack (e.g., at the closest traceback-collaborative ASs with respect to the sources of a DDoS attack, or at the AS that forwards more traffic). This kind of IP-traceback system should typically operate at the AS level. Indeed, recent work included placing monitors in the Internet, specifically within ASs [7]. Durresi et al. [8] propose an AS-based traceback system using a probabilistic packet marking (PPM) technique. Korkmaz et al. [9] propose an AS-level single-packet traceback (AS-SPT) mechanism that could operate in a partial deployment scenario.

Although the proposed architecture allows partial deployment, it requires previous knowledge of the network topology. In a previous work [10], we evaluated the accuracy of deploying a traceback system partially, however, without building an overlay network. We build upon the promising results in this previous work and propose the AS-level IP-traceback system presented in this article, thus contributing the mechanism to form the AS-level overlay network for IP traceback and a new sequence-marking process to remove ambiguities in the traceback path. In contrast with other recent proposals, our proposed AS-level IP-traceback system does not require previous knowledge of the network topology and allows single-packet traceback and incremental deployment. Previous work on IP traceback typically requires complete deployment over the network, that is, the system must operate on all routers in the monitored network to traceback an ongoing DDoS attack properly. This requirement comes from the way traceback is performed on those systems, with a focus on rebuilding the complete path taken by attacker packets. We believe this constraint limits the possibility of such IP-traceback systems to deal with large-scale DDoS attacks that are present in the Internet today.

3. Overlay Network for IP Traceback

In our proposed IP-traceback system, packet marking is performed similarly to the way it is performed by Laufer et al. [5]. The data inserted into the GBF of an IP packet carries the marks of the routers that the packets traverse. Our proposed system, however, uses BGP as a vehicle for information exchange among ASs that are participating in the scheme. This feature allows the establishment of an overlay network at the AS level for IP traceback. Moreover, its usage enables the discovery of the next hop in the reverse path followed by attacker packets to reach the

victim. Therefore, our system is capable of working at the AS level, eliminating the requirement of being deployed sequentially in all routers of the monitored network. In other words, the system could be partially and incrementally deployed over the Internet at the AS level.

A. Role of BGP in Proposed System

BGP routers periodically use update messages to exchange routing information with each other. An update message has a field named Path Attributes, which is actually a collection of attributes associated to a given route that may influence the route selection process. One of these attributes, called Community Attribute, is used to group destinations that share common characteristics. The community attribute is highly flexible and is indeed used for many different purposes, such as multi-home routing, traffic engineering, support of virtual private networks (VPNs), and mobile honeypot systems.In our IP-traceback system, we propose the creation of a new IP-traceback community comprising information about the presence of our traceback system on collaborating ASs. This indicates that these ASs then could create and participate on the AS-level overlay network for IP traceback. An important characteristic of a community attribute is that it is an optional and transitive attribute of BGP. This means that if the BGP running in a border router does not recognize an attribute present in the update message, a verification of whether the transitive flag is set or not is made. In the case of when the flag is set, the attribute is forwarded in a new update message by the AS border router to its peers. This feature enables the information about the IP-traceback community also to be forwarded by ASs not directly participating in the scheme, and so, information eventually reaches Ass that implement the proposed IP-traceback system. After a sequence of update messages is forwarded, the overlay network for traceback is established or updated. At this time, each IP-traceback AS has an overlay table that contains a list of all other IP-traceback ASs known by the owner of the table. Each entry in the overlay table contains two basic pieces of information: the identification of the neighbor AS in the overlay network that has the IP-traceback system installed and the AS responsible for the BGP update message that indicated this overlay neighbor in the topology. In other words, this table contains all AS-level network overlay neighbors for IP traceback. The size of these overlay tables at each collaborating AS is proportional to the number of neighbors it has in the proposed AS-level overlay network. The update of the overlay table occurs in a similar way as the updating of the BGP routing table because the overlay information is carried on BGP messages.

B. Building An Overlay Network for Ip Traceback

The overlay network enables IP traceback among participating AS routers (they are not required to be adjacent routers at the routing or AS level). The IP traceback is performed hop by hop in the AS-level overlay network. This feature eliminates the requirement adopted by several previous approaches that the traceback system should be deployed in all routers of the monitored network. Figure 1a illustrates the construction of the overlay network; ASs with a flag have the traceback system running. At the beginning, overlay tables are empty. When AS1 sends an update message to its peers (step 1), AS3 receives the message and updates its table by registering AS1 as its neighbor in the overlay network. On the other hand, since AS2 does not have the system deployed, AS2 just generates a new update message, sending in a transparent way, the information previously received about the IP-traceback community to its neighbors AS4 and AS3 (step 2). This happens because this information is set as transitive in the update message received from AS1. AS3 receives the update message coming from AS2 and simply ignores it because AS3 already received the information about AS1. In its turn, AS4 inserts AS1 as its neighbor in the overlay table. When AS3 and AS4 create a new update message, they insert their information about the IP traceback community and send the message to their neighbors (steps 3 and 4, respectively). After receiving the update message from each other, AS3 and AS4 insert each other as neighbors in their overlay tables. Similar procedures are repeated by all ASs in the network until all the participating ASs are reached and know about each other, thus forming the overlay network. The resulting overlay network is illustrated in Fig. 1b where thick lines represent the connections of the overlay level.

We emphasize that the exchanged information about IP traceback community incurs no significant additional overhead to the network because such information is carried inside update messages that are native and exchanged periodically by BGP routers.

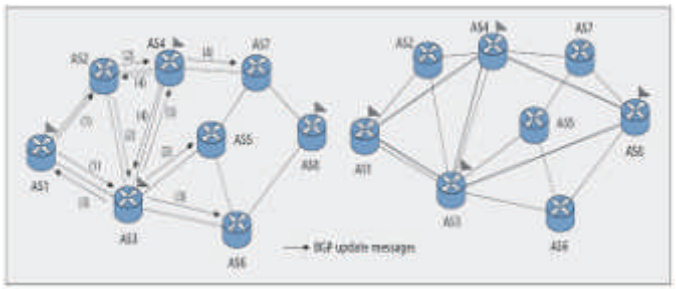

Fig 1: Building the as-level overlay network for ip traceback

C. Packet Marking

The packet-marking process originally proposed by Laufer etal. [5] was modified to properly operate with our proposed system. In the original process, whenever a packet traverses a router, the router inserts a mark into the GBF so that when the packet reaches the destination, the GBF contains marks of all traversed routers. The inserted filter may be carried as an option of the IP header. However, when a router is performing the traceback process, marks from more than one neighbor router could be found in the GBF. This problem is illustrated at the router level by the scenario presented in Fig. 2 black arrows indicate the attacker path where router RT1 verifies that routers RT2 and RT3 have their marks in the GBF because both form part of the route taken by attacker packets. In this case, RT1 does not have a way to distinguish which router, RT2 or RT3, is the direct next hop in the reverse path. Therefore, if router RT1 sends the reconstruction path packet to RT2, the traceback tends to finish without problems. However, if the reconstruction path packet is sent to RT3, a problem could arise because RT3 finds marks belonging to routers RT2 and RT5 in the GBF. If RT3 sends the reconstruction path packet to RT2, the traceback could finish unexpectedly without being accomplished or unnecessarily generate repeated messages in the network. Note that similar issues can arise if tracebacking at the AS level is performed.

To prevent the problem mentioned above in our AS-level IP-traceback system, we propose to identify the AS sequence in the marking process as follows. When an edge router with the traceback system deployed receives a packet, it first makes an exclusive or (XOR) operation between two 16-bit numbers: the first one is the autonomous system number (ASN) itself (the 16-bit number that identifies each AS), and the other one is formed by the time-to-live (TTL) of the packet at that moment. Since TTL is eight bits long, we complete the remaining eight most significant bits with 0. The actual marking process is performed just afterwards, that is, it is the result of the XOR operation that is submitted to a hash to generate the mark that should be inserted into the GBF. This procedure (illustrated in Fig. 3) is performed by all IP-traceback ASs. Thus, when the packet reaches its destination, it still has the marks of all participating ASs from where the packet traverses. Moreover, this method of packet marking is required for correct reverse path reconstruction, as described in further detail later in this section. Considering a possible deployment of the system, it is worth noting that the marking process does not demand a high computation effort because it can be performed through basic logical operations: the previously detailed XOR operation to avoid ambiguities in the path reconstruction in addition to an OR and another AND operation required to update the GBF of each packet [5, 6]. The size of the GBF filter to be adopted actually depends on a trade-off between the targeted maximum false-positive probability, the number of considered hash functions, and the expected number of elements to be represented in the filter (for a full analysis of this trade-off, see Laufer et al. [5, 6]). For example, following this analysis and considering 16 hash functions (eight to set and eight to reset bits in the GBF), a filter of 64 bits would have as false positive probabilities 0.09 percent, 0.5 percent, and 3 percent, after inserting three, five, and eight entries in the filter, respectively. It should be noted that Siganos et al. [11] observe that the average number of ASs that a packet traverses from its source to its destination is at most five for more than 95 percent of flows. Therefore, we may expect a relatively small number of entries in a typical filter because only ASs that take part of the overlay network insert their mark into the filter. The possibility of single-packet traceback without requiring per-packet state in the network core is achieved with additional per-packet overhead to carry the AS-level path information in the filter.

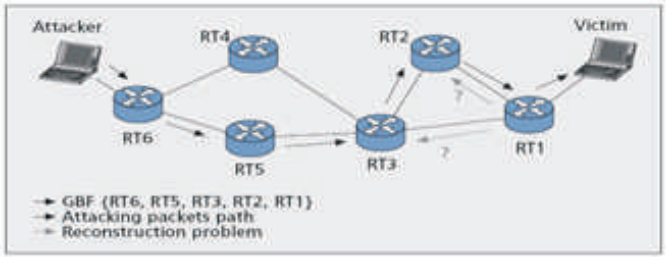

Fig 2: Packet marking problem

Fig 3: As sequence marking process to solve the problem

D. Tracebacking System

The traceback process can be started by an AS administrator or by a detection system installed therein. It initiates the reverse procedure of packet marking as explained in the previous section. The traceback process is illustrated in Fig. 4. Continuous arrows show the attacker path, and numbers indicate the TTL of the packet at that point. The victim's AS (AS8) starts the traceback by checking its overlay table. From this table, it searches for GBF marks belonging to either AS3 or AS4, that is, its neighbors in the overlay network (as also illustrated in Fig. 1b). To check where attacker packets come from (AS3 or AS4), AS8 proceeds as follows. First, an XOR operation is performed between the ASN of AS3 and the TTL of the packet increased by one. Note that TTL at AS8 is 251; then the TTL at AS3 must be 252 or greater. The result of the XOR operation is hashed and compared against the GBF of the packet; because there is no match, the procedure is now performed using the ASN of AS4. As the result is negative for both, the TTL is increased to 253, and the procedure is repeated until a match is found. In this case, the check is positive for AS4. Therefore, AS8 sends a reconstruction path packet to AS4 (step 1), which in turn increases the TTL (254) and repeats the process looking for marks belonging to either AS1 or AS3. The result is positive for AS3 (step 2). Then, the same procedure is repeated at AS3, and it finishes when the reconstruction path packet reaches AS1 (step 3) — in this case, because the source of the attack is behind AS1. Note that the traceback process actually could be finished in two ways: when the TTL reaches 256 or when an AS cannot find marks of any other neighbor in the GBF, thus concluding it is the closest AS to the source of the attack. When the traceback process finishes, ASs belonging to the overlay network that compose the path are aware of it, including the closest traceback-collaborative AS to the source of attacker packets. In the case of DDoS attacks, the traceback process in the overlay network is performed for different filters carried by the attacker packets, thus indicating different attacker paths from distinct sources. This results in a distributed identification of the closest traceback-collaborative ASs to the different attacker sources, thus allowing distributed countermeasures to be taken. These closest ASs could start packet filtering procedures to block the attack as close as possible to the source(s) of the attack, then preventing the attack and the network overhead incurred due to the forwarding of attacker packets to the victim. The filtering method to be adopted, however, is beyond the scope of this article, which is focused on IP-traceback mechanisms.

4. System Evaluation

In this section, we investigate the impact on the performance of our AS-level IP-traceback system resulting from the adopted strategy to deploy it in the network. We thus evaluate two ways to place our IP-traceback system in the network:

- Strategic placement, where the most connected ASs have a traceback system deployed first
- Random placement, where ASs are selected randomly to have the traceback system installed

The rationale behind strategic placement is based on several recent analyses about Internet topology [11, 12]. There is a clear tendency of new nodes in the Internet to connect to others that are highly connected a feature also known as preferential

attachment. This tendency helps to explain why typically a few nodes are highly connected whereas most nodes have only a few connections in the Internet topology. Our results confirm that strategically placing the proposed system in a limited portion of ASs is enough to enable an efficient IP-traceback system.

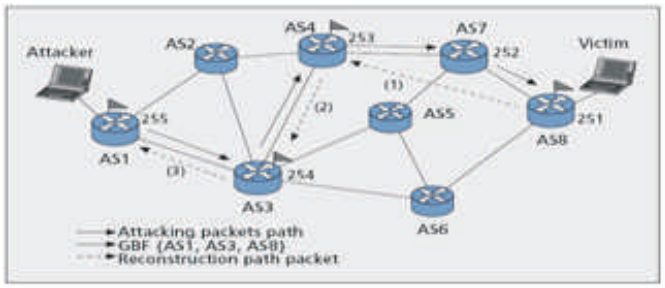

Fig 4: Tracebacking path with the help of GBF

5. Implementation

We adopted java technology for performing traceback process to identifying the attacker node for that purpose by using javafx concept designing the network and in which systems to perform the tracebacking process select those systems as with traceback facility and remaining are with- out traceback after that when neighbours are connected and identify it is having traceback means that entry is stored in the overlay table otherwise storing that entry into the normal table by using these entries tracebacking is very easy. We are performing the attack if it is exceeded the limit of the threshold then attack is happened when sending the information to other neighbours each and every packet is marked with their TTL value of packet and ASN number hashing is performed then it is stored into the GBF filter after the attack is happened the traceback process will start from which intermediate systems is coming then detect the source node that entire information is send to the destination node from this source node attack is happened and also if the source node is spoofed also we can detect how means we are using the TTL value of the packet and their autonomous system number after that blocking that source to avoiding the attack for other systems. We can reduce the false positives also by using the hash-mapping technique in java.

We are tested this for the 30, 60, 90 systems also process is fine we can able to detect the source node who is performing the attack it is shown in the fig 5.

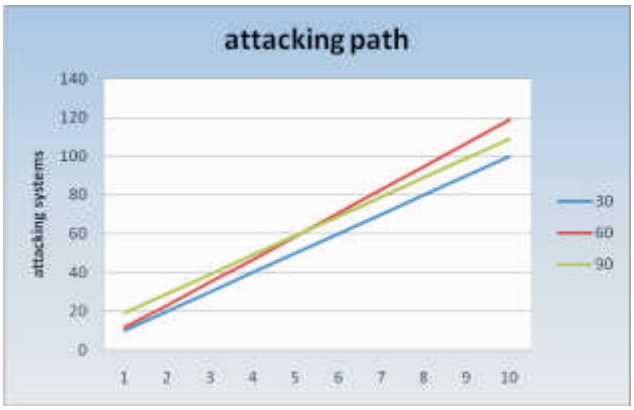

Fig 5: Efficiency discovered of attacking paths

6. Conclusion

We propose an AS-level IP-traceback system that takes advantage of some characteristics of BGP to build an AS-level overlay network for IP traceback. To establish and maintain this AS-level overlay network, we define a new community attribute for BGP that is responsible for disseminating the information about the ASs that have deployed the traceback system. This feature allows our proposed system to be partially and incrementally deployed in the network regardless of the topology, eliminating the requirement of being deployed in all network routers, which was an usual requirement in previous work. Furthermore, we introduce a method to identify the AS sequence in the marking process, avoiding possible problems of the identification of the correct sequence of traceback-collaborative ASs taking part in the attacker path. Our proposed AS-level traceback system depends on some collaboration among ASs to be efficient. We cite at least two incentives for ASs to adopt the proposed system and then participate in the AS-level traceback overlay:

- If a collaborative AS is in the path of an ongoing attack, in addition to helping other ASs in the traceback task, it can perform filtering of the attacker packets, thus reducing the consumption of network resources in its own domain because these resources are being used by attacker flows.

An AS with the traceback system installed can initiate its own traceback request when an attack is detected within its domain. Implementation results suggest that if a relatively small amount of ASs take part in the AS-level traceback overlay network given that they are strategically placed - a large portion of the AS-level route(s) taken by the attacker packets can be identified.

This enables countermeasures to be taken in a highly distributed way, that is, close to the attacker sources reducing the effect of a large-scale DDoS attack before attacker packets reach the victim's AS. Comparing the strategic and random placement strategies, we note that previous knowledge of the network topology contributes to better results. Although previous knowledge of the network topology directly influences the traceback accuracy, this is not a mandatory requirement for the operation of our proposed traceback system, as it is in [9].

7. Future Work

In future work, we intend to investigate the feasibility of integrating our AS-level traceback system with a router-level traceback system. Therefore, we intend to perform the traceback at two levels: first, the traceback can discover ASs from where packets are sent and second, the traceback can be performed inside these ASs, increasing the chance of getting closer to the attacker sources and thereby performing more efficient filtering.

8. References

1. *D. Moore et al., "Inferring Internet Denial-of-Service Activity," ACM Trans.Comp. Sys., vol. 24, no. 2, May 2006, pp. 115–39.*
2. *A. Belenky and N. Ansari, "On IP Traceback," IEEE Commun. Mag., vol. 41,no. 7, July 2003, pp. 142–53.*
3. *A. C. Snoeren et al., "Single-Packet IP Traceback," IEEE/ACM*

Trans. Net., vol. 10, no. 6, Dec. 2002, pp. 721–34.

4. *B. Bloom, "Space/Time Trade-Offs in Hash Coding with Allowable Errors," Commun. ACM, vol. 13, no. 7, July 1970, pp. 422–26.*
5. *R. P. Laufer et al., "Towards Stateless Single-Packet IP Traceback," 32nd IEEE Conf. Local Comp. Net., Dublin, Ireland, Oct. 2007.*
6. *R. P. Laufer, P. B. Velloso, and O. C. M. B. Duarte, "Generalized Bloom Filters," COPPE/UFRJ, Tech. Rep. GTA-05-43, Sept. 2005; http://www.cs.ucla. edu/_rlaufer/publications/gbf.pdf*
7. *A. W. Jackson et al., "A Topological Analysis of Monitor Placement," Proc. 6th IEEE Int'l. Symp. Net. Comp. and Applications, July 2007, pp. 169–78.*
8. *A. Durresi et al., "Efficient and Secure Autonomous System Based Traceback," J. Interconnection Netw., vol. 5, no. 2, June 2004, pp. 151–64.*
9. *T. Korkmaz et al., "Single Packet IP Traceback in AS-level Partial Deployment Scenario," Int'l. J. Security and Netw., vol. 2, no. 1–2, 2007, pp. 95–108.*
10. *A. O. Castelucio, R. M. Salles, and A. Ziviani, "Evaluating the Partial Deployment of an AS-level IP Traceback System," Proc. ACM Symp. Applied Comp. '08, Fortaleza, Brazil, Mar. 2008.*
11. *G. Siganos et al., "Power Laws and the AS-Level Internet Topology*
12. *A. Medina, I. Matta, and J. Byers, "On the Origin of Power Laws in Internet Topologies,"*
13. *"An AS-Level Overlay Network for IP Traceback" by André Castelucio and Artur Ziviani,*