PIONEERING TECHNOLOGIES
FOR A BETTER INTERNET

Cs3, Inc.

5777 W. Century Blvd.
Suite 1185
Los Angeles, CA 90045-5600

Phone: 310-337-3013
Fax: 310-337-3012
E-mail: info@cs3-inc.com

Attack Attribution in Non-Cooperative Networks

Donald Cohen* & K. Narayanaswamy*, Member, IEEE

I. BACKGROUND

This paper reports on preliminary research concepts in attack attribution that have been developed in Cs3's project being conducted for Advanced Research and Development Activity (ARDA)[1]. The ARDA BAA [1] identified 4 levels of attribution**:**

- Level 1: Attribution to the specific hosts involved in the attack;
- Level 2: Attribution to the primary controlling host;
- Level 3: Attribution to the actual human actor;
- Level 4: Attribution to an organization with the specific intent to attack.

Cs3's research specifically focuses on attribution in situations where universal cooperation is not available for the attribution effort. This paper concentrates only on research concepts that show promise in resolving the Level 1 attribution problem. The name of the project is **S**ystematically **T**racking **A**ttackers through **R**outing **D**ata, **E**vents, and **C**ommunication **K**nowledge (**STARDECK**).

II. PROBLEM

Our model is that an IP packet, P, is generated by a machine, G, forwarded by a sequence of IP routers, and finally, if not dropped along the way, delivered to a recipient machine. The goal of Level 1 attribution is, given P, to identify G. In general, such identification might be requested by the recipient machine, by any of the forwarding routers, or by any other machine to which those routers communicate information about P.

Note that the source address field of every IP packet is supposed to contain the IP address of the machine that generated it. If this requirement were actually enforced on the Internet, then Level 1 attribution would be trivial. The Level 1 attribution problem is more commonly referred to as "*traceback*" in the literature. We use the terms interchangeably for the rest of this paper.

Ideally, the tracker would like to identify the physical machine (G) that generated the packet **P**. In practice, he will be satisfied with an IP address that is unique within the network visible to him. Traceback from a single packet is deemed to be a critical requirement of the problem in order to attribute all attacks, and not just those that feature a high volume of attack packets (such as DoS floods).

III. ASSESSMENT OF RELATED WORK

We are not aware of other work specifically intended to identify sets of possible origins given partial cooperation. We have, however, found two general categories of related work, which we call link identification and filtering.

Link identification techniques call for routers or other machines in the network to monitor or mark traffic so that the forwarding path of a packet can be "*traced*" back toward its origin [3 - 12]. These techniques can determine the origin of any packet

[1] This work is supported by ARDA under contract NBCHC030115. The views expressed in this document are those of the author(s) and Cs3 Inc., and not those of ARDA or its representatives. * Computer Scientists at Cs3 Inc., Los Angeles, California.

given universal cooperation. However, in the case of partial cooperation the result is a partial path, which only slightly restricts the origin. Unfortunately, the methods cited above only identify routers on the forwarding paths of floods and do not work from a single packet.

Link identification techniques call for routers or other machines in the network to monitor or mark traffic so that the forwarding path of a packet can be "traced" back toward its origin [3 - 12]. These techniques can determine the origin of any packet given universal cooperation. However, in the case of partial cooperation the result is a partial path, which only slightly restricts the origin. Unfortunately, the methods cited above only identify routers on the forwarding paths of floods and do not work from a single packet.

Filtering techniques call upon routers to use routing knowledge to recognize that certain packets have incorrect source addresses. The routers simply drop such packets [13 - 15]. Unlike the work in link identification, there are known results that relate the amount of cooperation available in the network infrastructure to the effectiveness of filtering. In the case of ingress filtering [13,14], it is well known that the set of possible origins is the union of the network containing the source address and all networks that do not cooperate (i.e., that do not do ingress filtering).

[15] describes a generalization of ingress filtering. Packets are filtered if the route from their alleged origin to their destination does not pass through the link on which the router receives them. [15] explicitly addresses the question of how the quality of the result (in this case, what packets with false source addresses go unfiltered) is related to the number and placement of the cooperating machines in the network.

IV. STARDECK ATTACK ATTRIBUTION

The STARDECK approach combines two kinds of data. One is link identification data. In general, a cooperating machine that is able to observe the traffic on some link in the network allows the tracker to determine that a given packet either did or did not traverse that link. Therefore, a set of cooperating machines provides the tracker with a set of links known to have been traversed by the packet, which we call the *positive links* for that packet, and a set of links known not to have been traversed by the packet, which we call the *negative links* for that packet.

The second kind of data is routing data. In abstract terms, routing data relates origins, destinations, and forwarding paths. If the tracker had access to complete routing data for the network, he could compute, for every origin x, and destination y, all of the possible forwarding paths[2] [footnote:.] from x to y.

Given a packet, P, the tracker uses his cooperating machines to compute sets of positive and negative links. The forwarding path of P

- must be one of the forwarding paths determined by routing data
- must end at the destination address of P
- must be consistent with the observed link identification data:
  – every positive link must be in the path
  – no negative link can be in the path.

The set of possible origins of P is the origins of the paths that satisfy all of the requirements above.

The biggest problem is that the tracker generally does not have complete routing data. This is also the problem with route based filtering [15]. Our ongoing work includes finding more ways to get routing data, finding more ways to use routing data to compute what is needed (or at least useful) for the approach above, estimating errors in approximations that can be computed from available data and relating those errors in routing data to errors in the attribution result.

V. STARDECK INNOVATIVE CONTRIBUTIONS

STARDECK combines any data available from any of the previously described link identification methods. That is, one method might be used to identify one link as positive or negative, while another method is used on another link. Some of the link identification methods do not apply to individual packets but only to aggregates. Our approach can still use the data from those methods to identify the possible sources of the aggregates.

[2] A path is defined as a sequence of links.

It may not be obvious that the STARDECK approach can also make use of filtering methods. More precisely, the knowledge that certain packets are filtered at certain places amounts to negative link data. If the tracker knows, for instance, that a router at a given link filters packets with some given source address, then he can infer that a packet with that source address must not have traversed the link. This rules out as possible origins all origins from which the packet would have traversed that link. Filtering routers in some sense cooperate with everyone. A current problem remaining for the tracker is to find out where filtering routers are installed, and what exactly they filter.

VI. REFERENCES

[1] http://www.nbc.go /pip.cfm

[2] RFC 1812, F. Baker, Editor Requirements for IP Version 4 Routers, http://www.ietf.org/rfc/rfc1812.txt, June 1995

[3] Characterizing and Packet Floods Using Cisco Routers, http://www.cisco.com/warp/public/707/22.html

[4] Tracing Anonymous Packets to their Approximate Sources, Hal Burch and Bill Cheswick, In Usenix Lisa, December 2000

[5] ICMP Traceback Messages; Steven Bello in, http://www.research.att.com/%7Esmb/papers/draft-belloin-itrace-00.txt, Internet draft, March 2000

[6] Practical Network Support for IP Traceback; Stefan Savage, David Wetherall, Anna Karlin, Tom Anderson; In ACM SIGCOMM, August 2000

[7] Single Packet IP Traceback; A. C. Snoeren, C. Partridge, L. A. Sanchez, C. E. Jones, F. Tchakountio, B. Schwartz, S. T. Kent, and W. T. Strayer; IEEE/ACM Transactions on Networking (ToN), 10(6), Dec. 2002

[8] Providing for responsibility in a Global Information Infrastructure; Fred Cohen, http://www.all.net/journal/ntb/responsible.html

[9] Changing IP to Eliminate Source Forgery; Don Cohen, K. Narayanaswamy, Fred Cohen; http://www.cs3-inc.com/pubs/eliminating-source-forgery.pdf

[10] A Fair Service Approach to Defenses Against Packet Flooding Attacks; Don Cohen and K. Narayanaswamy; http://www.cs3-inc.com/pubs/fair-service-ddos-defense.pdf

[11] PI : A New Defense Mechanism Against IP Spoofing and DoS Attacks; A. Yaar, A. Perrig, D. Song; IEEE Symposium on Security, 2003