Chapter 16

THE FORGER'S ARSENAL

INTRODUCTION

Communication in the Internet works by sending packets from one place to another. Each packet, like a post card, contains a source address and a destination address. The Internet fulfills the role of the post office, delivering packets to their specified destination addresses. It is interesting to notice that only the destination address is used to deliver the packet. In most cases, however, the sender wants the destination to reply. The source address is used by the destination to address the reply [1].

Unfortunately, considerable mischief can be caused by sending packets with incorrect source addresses. First, it is very likely that those sending such unwanted messages would also like to avoid being identified by the recipients. Even worse, a recipient who believes the forged source address will blame the owner of that address for the unwanted message [1].

Some of the worst attacks in the Internet today involve sending packets that cause automatic replies. Typically, in this case, neither the party that receives the original packet nor the party that receives the reply would object to a few such packets, but the attacker arranges for them to get huge numbers. Each feels like he or she is being attacked by the other. Alternatively, a large number of places are sent a smaller number of packets and the replies all converge on a victim who sees an attack that appears to come from a large number of places. Even if the attack is coming from a large number of places, that number can be made to appear much larger by reflecting the packets off many innocent intermediaries [1].

INGRESS FILTERING

The commonly recommended solution is called ingress filtering. In the post-office analogy, ingress filtering corresponds to the local post office only accepting letters with its own zip code in the return address. A potential attacker can still successfully use valid return addresses that belong to his or her neighbors, or even non-existent addresses with that

zip code. However, if these letters cause trouble, the range of possible suspects is relatively small. If the attacker sends many such letters, the local post office is likely to find out that the letters came from him or her [1].

All of the preceding statements hold for the Internet as well as the post office. The local post office corresponds to an ISP. The ISP knows that all of the packets it sends out should have source addresses in a given range. Ingress filtering simply refuses to forward those packets with source addresses outside that range [1].

> **Note:** The post office actually does have something just as good as this. The postmark shows where a letter entered the system. The recipient can compare the return address to that postmark.

There are a few problems with ingress filtering. First, as noted, it is still hard to tell when an attacker forges the address of his or her neighbor. More important, the Internet, unlike the post office, has no central authority that can force all of the ISPs to check the source addresses of outgoing packets. Finally, this extra effort does not really help the ISP that exerts the effort. It helps the rest of the world by preventing source forgery originating from that ISP, whereas the customers who pay the ISP for service get no direct benefit. In light of this, perhaps it is not so surprising that, in spite of all of the preceding recommendations (and their appeals to the spirit of cooperation), few, if any, ISPs actually do ingress filtering [1].

DISCERNING TRUE PACKET SOURCES

Unlike the post office, the Internet consists of a large number of different carriers who forward packets to each other. If things work correctly, then each forward operation brings a packet closer to its destination until it finally arrives there. When a packet arrives, each forwarder can normally tell who sent it the packet, but not where that party, in turn, got the packet, or anything more about the path the packet took through the different carriers. It is therefore recommended that each forwarder add to the packet an indication of who gave the packet to it. Each packet will then arrive with a complete forwarding path [1].

Again, this is not a new idea. The advantage of this scheme is that the path is not controlled by the sender. A forger is now in the position of writing a return address of New York when the receiver can plainly read a postmark that says Chicago! This is really a proposal to change the communication protocol used in the Internet. The new protocol is viewed as an enhancement of the Internet Protocol (IP), and is referred to as Path Enhanced IP, or PEIP [1].

The indication in the preceding of where the packet came from does not have to be an IP address. That would take much more space than necessary, space that could otherwise be used for real data. In PEIP, space is saved by encoding the paths. The encoding scheme and some additional reasons to save space are mentioned next [1].

A forwarder that can receive packets from ten different places will add to the path a number between one and ten. This, naturally, means that a receiver needs a way to decode the path. However, a separate protocol will be needed to decode the path into a sequence of IP addresses. This protocol requires each machine along the path to take the data that it (supposedly) added, and find the IP address of the neighbor for which it would have