

Automatic Spoof Detector  
(aka Spoofwatch)  
1/28/2002  
Advanced Network Management Lab  
Indiana University

### Background

The software agents that are covertly installed on compromised computers to generate Distributed Denial of Service (DDoS) attacks typically mask their true network identity by using an invalid source address, otherwise known as “spoofing”. Early forms of these agents, also known as zombies or bots, would choose a new random source address for each packet. Network administrators were able to combat this type of spoofing with relative ease by configuring routers to block obviously spoofed addresses. In response, zombies evolved to use more sophisticated spoofing, selecting source addresses which would be valid sources for their particular local area network (LAN).

This intelligent spoofing betrays the sub-network of the agent, but not the particular computer. Sub-networks can be quite large, sometimes including hundreds or thousands of computers, so finding the particular compromised system can be difficult. Since these agents do not answer Address Resolution Protocol (ARP) requests for their spoofed source addresses, it is possible to distinguish spoofed traffic from legitimate traffic, and in doing so collect the Ethernet address of the compromised computer. Once the compromised computer’s Ethernet address is known, it is possible to identify its location.

### Overview

Spoofwatch relies on the premise that DDoS zombies that spoof IP source addresses do not answer ARP requests for their spoofed addresses. Spoofwatch is a collection of standard linux-based utilities and some custom software that, together, has the ability to identify spoofed traffic on a locally connected LAN. Spoofwatch uses TCP-dump to collect all IP and ARP packets on its local interface. Data in the ARP packets populate a database of presumably valid IP address to Ethernet associations. At program start, a port scanning utility is used to send a single packet to all the possible valid IP addresses on the subnet, thus forcing the computer running Spoofwatch to generate an ARP request for each host. The replies are used to build Spoofwatch’s legitimate IP to Ethernet MAC address database.

When Spoofwatch captures a packet where the source IP address does not match the corresponding Ethernet address recorded in its ARP database, Spoofwatch first re-ARPs for that IP address to insure that it hasn’t been legitimately re-allocated to another computer. If the ARP database is still inconsistent, or no ARP reply is received, the packet is classified as spoofed and the IP address and Ethernet address are recorded.

### Implementation Details

Spoofwatch has to be given a list of valid networks (IP network prefixes) for the LAN it is monitoring. The current implementation is not Host Standby Router Protocol-aware (HSRP), so it may give false positives for IP addresses corresponding to the default

gateway address and its corresponding HSRP neighbor's standby address. A hack has been applied to the current implementation to ignore these addresses for the IUB campus (they always have the same host portion address per campus convention).