Introduction

Spoofing can take on many forms in the computer world, all of which involve some type fraudulent representation of information. There are a variety of methods and types of spoofing. I would like to introduce and explain four in this paper:

IP
ARP
Web
DNS

There are no legal or constructive uses for implementing spoofing of any type. Some of the outcomes might be sport, theft, vindication or some other malicious goal. The gravity of these attacks can be very severe, can cost us millions of dollars and should not be overlooked by the Internet security community.
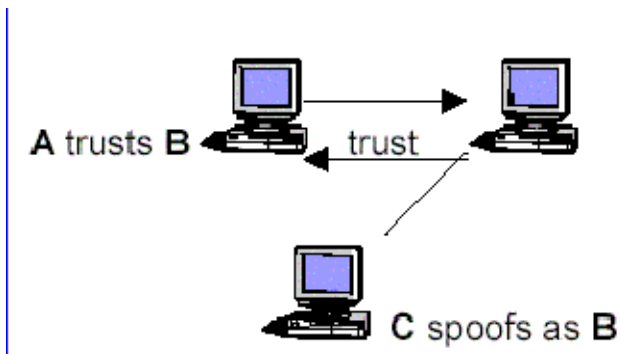
IP Spoofing

IP spoofing is used to gain unauthorized access to a computer. The attacker forwards packets to a computer with a source address indicating that the packet is coming from a trusted port or system. Attackers must go through some complicated steps to accomplish the task. They must:

Acquire a target
Acquire an IP address of a trusted machine
Disable communication of the trusted machine (e.g. syn flooding)
Sample a communication between the target and trusted hosts
Guess the sequence numbers of the trusted machine
Modify the packet headers so that it appears that the packets are coming from the trusted host
Attempt connection to an address authenticated service or port.
If successful, the attacker will plant some kind of backdoor access for future reference.

System A impersonates system B by sending B's address instead of its own. The reason for doing this is that systems tend to function within groups of other ``trusted'' systems. This trust is implemented in a one- to-one fashion; system A trusts system B. IP spoofing occurs in the following manner: if system A trusts system B and system C spoofs system B, then system C can gain otherwise denied access to system A. This is all made possible by means of IP address authentication, and if the packets are coming from external sources- poorly configured routers.

One of the major drawbacks with IP spoofing is that C never "sees" the responses from A. This is completely blind attack, much experience and knowledge of what to expect from the target?s responses is needed to successfully carry out his attack.

Some of the most common ways to avoid this type of attack are to disable source-routed packets and to disable all external incoming packets with the same source address as a local host.

For an excellent example of what a spoofing attack might look like visit: http://www.blinky-lights.org/shimomura-25jan95.html. This includes packet captures and contains explanations as to what the

attacker is trying to accomplish as he is doing it.

ARP Spoofing

ARP (Address Resolution Protocol) is used to map IP addresses to hardware addresses. A table, usually called the ARP cache, is used to maintain a correlation between each MAC address and its corresponding IP address. ARP provides the protocol rules for making this correlation and providing address resolution in both directions. When an incoming packet sent to a host machine on a network arrives at a router, it asks the ARP program to find a MAC address that matches the IP address. The ARP program looks in the ARP cache and, if it finds the address, provides it so that the packet can be converted to the right packet length and format and sent to the machine. If no entry is found for the IP address, ARP broadcasts a request packet in a special format to all the machines on the network to determine if any machine knows who has that IP address. A machine that recognizes the IP address as its own returns a reply so indicating. ARP updates the ARP cache for future reference and then sends the packet to the MAC address that replied. Here is a sample ARP broadcast query:
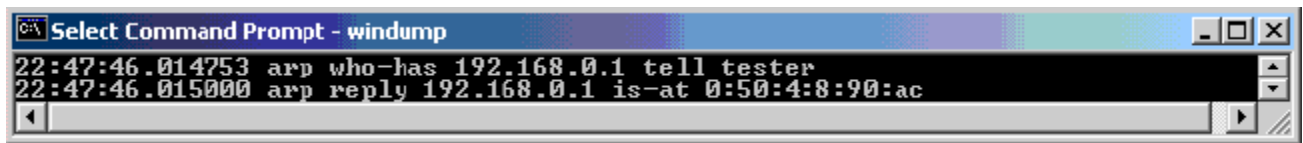
One might deduct that this addressing scheme could also be spoofed to provide a host with incorrect information "ARP Spoofing involves constructing forged ARP request and reply packets. By sending forged ARP replies, a target computer could be convinced to send frames destined for computer A to instead go to computer B.". This referred to as ARP poisoning.

There are currently programs that automate the process of ARP poisoning ? ARPoison, Ettercap, and Parasite. All three have the capability to provide spoofed ARP packets and therefore redirect transmission, intercept packets, and/or perform some type of man in the middle attack.

Either enabling MAC binding at a switch or implementing static ARP tables achieves prevention of ARP spoofing. MAC binding makes it so that once an address is assigned to an adapter; it cannot be changed without authorization. Static ARP management is only realistically achieved in a very small network. In a large dynamic network, it would be impossible to manage the task of keeping the entries updated. ARPWATCH, for Unix based systems, monitors changes to the ARP cache and alerts administrator as to the changes.

Web Spoofing

As with the other forms of spoofing Web or Hyperlink spoofing provides victims with false information. Web Spoofing is an attack that allows someone to view and modify all web pages sent to a victim's machine. They are able to observe any information that is entered into forms by the victim. This can be of particular danger due to the nature of information entered into forms, such as addresses, credit card numbers, bank account numbers, and the passwords that access these accounts.

Web Spoofing works on both Internet Explorer and Netscape and is not necessarily prevented by secure connections. This is due the way that the SSL protocol uses certificates to authenticate websites. The attacker can observe and modify all web pages and form submissions, even when the browser is indicating that there is a secure connection. The user usually never sees anything that is out of the ordinary (O?Dwyer, Frank, http://www.brd.ie/papers/sslpaper/sslpaper.html)

The attack can be implemented using JavaScript and Web server plug-ins, and works in two parts. First, the attacker causes a browser window to be created on the victim's machine, with some of the normal status and menu information replaced by identical-looking components supplied by the attacker. Then, the attacker causes all Web pages destined for the victim's machine to be routed through the attacker's server. On the attacker's server, the pages are rewritten in such a way that their appearance does not change at all, but any actions taken by the victim (such as clicking on a link) would be logged by the attacker. In addition, any attempt by the victim to load a new page would cause the newly loaded page to be routed through the

attacker's server, so the attack would continue on the new page. The attack is initiated when the victim visits a malicious Web page, or receives a malicious email message (if the victim uses an HTML- enabled email reader).

Current browsers do not completely prevent Web Spoofing, and there seems to be little movement in the direction of addressing this problem. I believe that there can be no fully secure electronic commerce on the Web until the Spoofing vulnerability has been addressed.

DNS Spoofing

DNS spoofing essentially occurs in three ways.

An attacker compromises a DNS server and changes the hostname to IP address mappings. When someone requests the URL of one if the altered mappings that person is sent to a machine that is under the complete control of the attacker.
An attacker can spoof the responses from DNS server before the real one comes. Guessing the connectionless, UDP sequence numbers does this. Sequence numbers are incremented by one, which makes this is relatively easy to accomplish. If the attacker can sniff a transaction from the DNS server, her or she will then be able to accurately guess the sequence numbering.
The DNS Cache can be "poisoned" simply sending false replies with a high TTL (so that the poisoning will last) to requesting the DNS server. A remote name server that is controlled by attacker, that is set up with false name resolutions is called upon, and then cached by the requester. (Erdfelt, Johannes http://www.the-project.org/admins/0897-1097/0381.html)

DNS spoofing can be utilized in some of the same ways as web spoofing. If someone is directed through a compromised server, any data that is entered into a website or online order form can be viewed. Much personal data can be gathered by this means. Email can also be redirected through a maligned server, which could be especially dangerous for companies who share proprietary information via email.

Some of the precautions to take would involve using the most up to date patches for all versions and implementations of DNS.

Conclusion

With the current implementations of spoofing, the network security community needs to be aware of the gravity and potential cost of these types of attacks. People can effectively maintain patching and monitoring of logs to minimize the potential damage.

Neil B. Riser
July 1, 2001