

[IP-spoofing Demystified] (Trust-Relationship Exploitation)

by daemon9 / route / infinity
for Phrack Magazine
June 1996 Guild Productions, kid

comments to route@infonexus.com

The purpose of this paper is to explain IP-spoofing to the masses. It assumes little more than a working knowledge of Unix and TCP/IP. Oh, and that yur not a moron...

IP-spoofing is complex technical attack that is made up of several components. (In actuality, IP-spoofing is not the attack, but a step in the attack. The attack is actually trust-relationship exploitation. However, in this paper, IP-spoofing will refer to the whole attack.) In this paper, I will explain the attack in detail, including the relevant operating system and networking information.

SECTION I. BACKGROUND INFORMATION]

--[The Players]--

A: Target host
B: Trusted host
X: Unreachable host
Z: Attacking host
1)2: Host 1 masquerading as host 2

--[The Figures]--

There are several figures in the paper and they are to be interpreted as per the following example:

```
ick host a control host b  
1 A ---SYN---> B
```

tick: A tick of time. There is no distinction made as to *how* much time passes between ticks, just that time passes. It's generally not a great deal.

host a: A machine participating in a TCP-based conversation.

control: This field shows any relevant control bits set in the TCP header and the direction the data is flowing

host b: A machine participating in a TCP-based conversation.

In this case, at the first referenced point in time host a is sending a TCP segment to host b with the SYN bit on. Unless stated, we are generally not concerned with the data portion of the TCP segment.

--[Trust Relationships]--

In the Unix world, trust can be given all too easily. Say you have an account on machine A, and on machine B. To facilitate going betwixt the two with a minimum amount of hassle, you want to setup a full-duplex trust relationship between them. In your home directory at A you create a .rhosts file: ``echo "B username" > ~/.rhosts`` In your home directory at B you create a .rhosts file: ``echo "A`

[IP-spoofing Demystified] (Trust-Relationship Exploitation)

by daemon9 / route / infinity
for Phrack Magazine
June 1996 Guild Productions, kid

username" > ~/.rhosts` (Alternately, root can setup similar rules in /etc/hosts.equiv, the difference being that the rules are hostwide, rather than just on an individual basis.) Now, you can use any of the r* commands without that annoying hassle of password authentication. These commands will allow address-based authentication, which will grant or deny access based off of the IP address of the service requestor.

--[Rlogin]--

Rlogin is a simple client-server based protocol that uses TCP as its transport. Rlogin allows a user to login remotely from one host to another, and, if the target machine trusts the other, rlogin will allow the convenience of not prompting for a password. It will instead have authenticated the client via the source IP address. So, from our example above, we can use rlogin to remotely login to A from B (or vice-versa) and not be prompted for a password.

--[Internet Protocol]--

IP is the connectionless, unreliable network protocol in the TCP/IP suite. It has two 32-bit header fields to hold address information. IP is also the busiest of all the TCP/IP protocols as almost all TCP/IP traffic is encapsulated in IP datagrams. IP's job is to route packets around the network. It provides no mechanism for reliability or accountability, for that, it relies on the upper layers. IP simply sends out datagrams and hopes they make it intact. If they don't, IP can try to send an ICMP error message back to the source, however this packet can get lost as well. (ICMP is Internet Control Message Protocol and it is used to relay network conditions and different errors to IP and the other layers.) IP has no means to guarantee delivery. Since IP is connectionless, it does not maintain any connection state information. Each IP datagram is sent out without regard to the last one or the next one. This, along with the fact that it is trivial to modify the IP stack to allow an arbitrarily chosen IP address in the source (and destination) fields make IP easily subvertable.

--[Transmission Control Protocol]--

TCP is the connection-oriented, reliable transport protocol in the TCP/IP suite. Connection-oriented simply means that the two hosts participating in a discussion must first establish a connection before data may change hands. Reliability is provided in a number of ways but the only two we are concerned with are data sequencing and acknowledgement. TCP assigns sequence numbers to every segment and acknowledges any and all data segments received from the other end. (ACK's consume a sequence number, but are not themselves ACK'd.) This reliability makes TCP harder to fool than IP.

--[Sequence Numbers, Acknowledgements and other flags]--

Since TCP is reliable, it must be able to recover from lost, duplicated, or out-of-order data. By assigning a sequence number to every byte transferred, and requiring an acknowledgement from the other end upon receipt, TCP can guarantee reliable delivery. The receiving end uses the sequence numbers to ensure proper ordering of the data and to eliminate duplicate data bytes.

TCP sequence numbers can simply be thought of as 32-bit counters. They range from 0 to 4,294,967,295. Every byte of data exchanged across a TCP connection (along with certain flags) is

[IP-spoofing Demystified] (Trust-Relationship Exploitation)

by daemon9 / route / infinity
for Phrack Magazine
June 1996 Guild Productions, kid

sequenced. The sequence number field in the TCP header will contain the sequence number of the *first* byte of data in the TCP segment. The acknowledgement number field in the TCP header holds the value of next *expected* sequence number, and also acknowledges *all* data up through this ACK number minus one.

TCP uses the concept of window advertisement for flow control. It uses a sliding window to tell the other end how much data it can buffer. Since the window size is 16-bits a receiving TCP can advertise up to a maximum of 65535 bytes. Window advertisement can be thought of an advertisement from one TCP to the other of how high acceptable sequence numbers can be.

Other TCP header flags of note are RST (reset), PSH (push) and FIN (finish). If a RST is received, the connection is immediately torn down. RSTs are normally sent when one end receives a segment that just doesn't jive with current connection (we will encounter an example below). The PSH flag tells the receiver to pass all the data it has queued to the application, as soon as possible. The FIN flag is the way an application begins a graceful close of a connection (connection termination is a 4-way process). When one end receives a FIN, it ACKs it, and does not expect to receive any more data (sending is still possible, however).

--[TCP Connection Establishment]--

In order to exchange data using TCP, hosts must establish a connection. TCP establishes a connection in a 3 step process called the 3-way handshake. If machine A is running an rlogin client and wishes to connect to an rlogin daemon on machine B, the process is as follows:

fig(1)

```
1  A   ---SYN--->  B
2  A  <---SYN/ACK--- B
3  A   ---ACK--->  B
```

At (1) the client is telling the server that it wants a connection. This is the SYN flag's only purpose. The client is telling the server that the sequence number field is valid, and should be checked. The client will set the sequence number field in the TCP header to its ISN (initial sequence number). The server, upon receiving this segment (2) will respond with its own ISN (therefore the SYN flag is on) and an ACKnowledgement of the clients first segment (which is the client's ISN+1). The client then ACK's the server's ISN (3). Now, data transfer may take place.

--[The ISN and Sequence Number Incrementation]--

It is important to understand how sequence numbers are initially chosen, and how they change with respect to time. The initial sequence number when a host is bootstrapped is initialized to 1. (TCP actually calls this variable 'tcp_iss' as it is the initial *send* sequence number. The other sequence number variable, 'tcp_irs' is the initial *receive* sequence number and is learned during the 3-way connection establishment. We are not going to worry about the distinction.) This practice is wrong, and is acknowledged as so in a comment the tcp_init() function where it appears. The ISN is incremented by 128,000 every second, which causes the 32-bit ISN counter to wrap every 9.32 hours

[IP-spoofing Demystified] (Trust-Relationship Exploitation)

by daemon9 / route / infinity
for Phrack Magazine
June 1996 Guild Productions, kid

if no connections occur. However, each time a connect() is issued, the counter is incremented by 64,000.

One important reason behind this predictability is to minimize the chance that data from an older stale incarnation (that is, from the same 4-tuple of the local and remote IP-addresses TCP ports) of the current connection could arrive and foul things up. The concept of the 2MSL wait time applies here, but is beyond the scope of this paper. If sequence numbers were chosen at random when a connection arrived, no guarantees could be made that the sequence numbers would be different from a previous incarnation. If some data that was stuck in a routing loop somewhere finally freed itself and wandered into the new incarnation of its old connection, it could really foul things up.

--[Ports]--

To grant simultaneous access to the TCP module, TCP provides a user interface called a port. Ports are used by the kernel to identify network processes. These are strictly transport layer entities (that is to say that IP could care less about them). Together with an IP address, a TCP port provides provides an endpoint for network communications. In fact, at any given moment *all* Internet connections can be described by 4 numbers: the source IP address and source port and the destination IP address and destination port. Servers are bound to 'well-known' ports so that they may be located on a standard port on different systems. For example, the rlogin daemon sits on TCP port 513.

[SECTION II. THE ATTACK]

...The devil finds work for idle hands....

--[Briefly...]--

IP-spoofing consists of several steps, which I will briefly outline here, then explain in detail. First, the target host is chosen. Next, a pattern of trust is discovered, along with a trusted host. The trusted host is then disabled, and the target's TCP sequence numbers are sampled. The trusted host is impersonated, the sequence numbers guessed, and a connection attempt is made to a service that only requires address-based authentication. If successful, the attacker executes a simple command to leave a backdoor.

--[Needful Things]--

There are a couple of things one needs to wage this attack:

- (1) brain, mind, or other thinking device
- (1) target host
- (1) trusted host
- (1) attacking host (with root access)
- (1) IP-spoofing software

Generally the attack is made from the root account on the attacking host against the root account on the target. If the attacker is going to all this trouble, it would be stupid not to go for root. (Since root access is needed to wage the attack, this should not be an issue.)

[IP-spoofing Demystified] (Trust-Relationship Exploitation)

by daemon9 / route / infinity
for Phrack Magazine
June 1996 Guild Productions, kid

--[IP-Spoofing is a 'Blind Attack']--

One often overlooked, but critical factor in IP-spoofing is the fact that the attack is blind. The attacker is going to be taking over the identity of a trusted host in order to subvert the security of the target host. The trusted host is disabled using the method described below. As far as the target knows, it is carrying on a conversation with a trusted pal. In reality, the attacker is sitting off in some dark corner of the Internet, forging packets purportedly from this trusted host while it is locked up in a denial of service battle. The IP datagrams sent with the forged IP-address reach the target fine (recall that IP is a connectionless-oriented protocol-- each datagram is sent without regard for the other end) but the datagrams the target sends back (destined for the trusted host) end up in the bit-bucket. The attacker never sees them. The intervening routers know where the datagrams are supposed to go. They are supposed to go the trusted host. As far as the network layer is concerned, this is where they originally came from, and this is where responses should go. Of course once the datagrams are routed there, and the information is demultiplexed up the protocol stack, and reaches TCP, it is discarded (the trusted host's TCP cannot respond-- see below). So the attacker has to be smart and *know* what was sent, and *know* what response the server is looking for. The attacker cannot see what the target host sends, but she can *predict* what it will send; that coupled with the knowledge of what it *will* send, allows the attacker to work around this blindness.

--[Patterns of Trust]--

After a target is chosen the attacker must determine the patterns of trust (for the sake of argument, we are going to assume the target host *does* in fact trust somebody. If it didn't, the attack would end here). Figuring out who a host trusts may or may not be easy. A 'showmount -e' may show where filesystems are exported, and rpcinfo can give out valuable information as well. If enough background information is known about the host, it should not be too difficult. If all else fails, trying neighboring IP addresses in a brute force effort may be a viable option.

--[Trusted Host Disabling Using the Flood of Sins]--

Once the trusted host is found, it must be disabled. Since the attacker is going to impersonate it, she must make sure this host cannot receive any network traffic and foul things up. There are many ways of doing this, the one I am going to discuss is TCP SYN flooding.

A TCP connection is initiated with a client issuing a request to a server with the SYN flag on in the TCP header. Normally the server will issue a SYN/ACK back to the client identified by the 32-bit source address in the IP header. The client will then send an ACK to the server (as we saw in figure 1 above) and data transfer can commence. There is an upper limit of how many concurrent SYN requests TCP can process for a given socket, however. This limit is called the backlog, and it is the length of the queue where incoming (as yet incomplete) connections are kept. This queue limit applies to both the number of incomplete connections (the 3-way handshake is not complete) and the number of completed connections that have not been pulled from the queue by the application by way of the accept() system call. If this backlog limit is reached, TCP will silently discard all incoming SYN requests until the pending connections can be dealt with. Therein lies the attack.

[IP-spoofing Demystified] (Trust-Relationship Exploitation)

by daemon9 / route / infinity
for Phrack Magazine
June 1996 Guild Productions, kid

The attacking host sends several SYN requests to the TCP port she desires disabled. The attacking host also must make sure that the source IP-address is spoofed to be that of another, currently unreachable host (the target TCP will be sending it's response to this address. (IP may inform TCP that the host is unreachable, but TCP considers these errors to be transient and leaves the resolution of them up to IP (reroute the packets, etc) effectively ignoring them.) The IP-address must be unreachable because the attacker does not want any host to receive the SYN/ACKs that will be coming from the target TCP (this would result in a RST being sent to the target TCP, which would foil our attack). The process is as follows:

fig(2)

```
1  Z(x)  ---SYN--->  B
    Z(x)  ---SYN--->  B
    Z(x)  ---SYN--->  B
    Z(x)  ---SYN--->  B
    Z(x)  ---SYN--->  B
    ...
2  X  <---SYN/ACK---  B
    X  <---SYN/ACK---  B
    ...
3  X  <---RST---  B
```

At (1) the attacking host sends a multitude of SYN requests to the target (remember the target in this phase of the attack is the trusted host) to fill its backlog queue with pending connections. (2) The target responds with SYN/ACKs to what it believes is the source of the incoming SYNs. During this time all further requests to this TCP port will be ignored.

Different TCP implementations have different backlog sizes. BSD generally has a backlog of 5 (Linux has a backlog of 6). There is also a 'grace' margin of 3/2. That is, TCP will allow up to backlog*3/2+1 connections. This will allow a socket one connection even if it calls listen with a backlog of 0.

AuthNote: [For a much more in-depth treatment of TCP SYN flooding, see my definitive paper on the subject. It covers the whole process in detail, in both theory, and practice. There is robust working code, a statistical analysis, and a lengthy paper. Look for it in issue 49 of Phrack. -daemon9 6/96]

--[Sequence Number Sampling and Prediction]--

Now the attacker needs to get an idea of where in the 32-bit sequence number space the target's TCP is. The attacker connects to a TCP port on the target (SMTP is a good choice) just prior to launching the attack and completes the three-way handshake. The process is exactly the same as fig(1), except that the attacker will save the value of the ISN sent by the target host. Often times, this process is

[IP-spoofing Demystified] (Trust-Relationship Exploitation)

by daemon9 / route / infinity
for Phrack Magazine
June 1996 Guild Productions, kid

repeated several times and the final ISN sent is stored. The attacker needs to get an idea of what the RTT (round-trip time) from the target to her host is like. (The process can be repeated several times, and an average of the RTT's is calculated.) The RTT is necessary in being able to accurately predict the next ISN. The attacker has the baseline (the last ISN sent) and knows how the sequence numbers are incremented (128,000/second and 64,000 per connect) and now has a good idea of how long it will take an IP datagram to travel across the Internet to reach the target (approximately half the RTT, as most times the routes are symmetrical). After the attacker has this information, she immediately proceeds to the next phase of the attack (if another TCP connection were to arrive on any port of the target before the attacker was able to continue the attack, the ISN predicted by the attacker would be off by 64,000 of what was predicted).

When the spoofed segment makes its way to the target, several different things may happen depending on the accuracy of the attacker's prediction:

- If the sequence number is EXACTly where the receiving TCP expects it to be, the incoming data will be placed on the next available position in the receive buffer.
- If the sequence number is LESS than the expected value the data byte is considered a retransmission, and is discarded.
- If the sequence number is GREATER than the expected value but still within the bounds of the receive window, the data byte is considered to be a future byte, and is held by TCP, pending the arrival of the other missing bytes. If a segment arrives with a sequence number GREATER than the expected value and NOT within the bounds of the receive window the segment is dropped, and TCP will send a segment back with the *expected* sequence number.

--[Subversion...]--

Here is where the main thrust of the attack begins:

fig(3)

```
1  Z(b)  ---SYN---->  A
2  B     <---SYN/ACK--- A
3  Z(b)  ---ACK---->  A
4  Z(b)  ---PSH---->  A
```

[...]

The attacking host spoofs her IP address to be that of the trusted host (which should still be in the death-throes of the D.O.S. attack) and sends its connection request to port 513 on the target (1). At (2), the target responds to the spoofed connection request with a SYN/ACK, which will make its way to the trusted host (which, if it *could* process the incoming TCP segment, it would consider it an error, and immediately send a RST to the target). If everything goes according to plan, the SYN/ACK will be dropped by the gagged trusted host. After (1), the attacker must back off for a bit to give the target ample time to send the SYN/ACK (the attacker cannot see this segment). Then, at (3) the attacker sends an ACK to the target with the predicted sequence number (plus one, because we're

[IP-spoofing Demystified] (Trust-Relationship Exploitation)

by daemon9 / route / infinity
for Phrack Magazine
June 1996 Guild Productions, kid

ACKing it). If the attacker is correct in her prediction, the target will accept the ACK. The target is compromised and data transfer can commence (4).

Generally, after compromise, the attacker will insert a backdoor into the system that will allow a simpler way of intrusion. (Often a ``cat + + >> ~/.rhosts`` is done. This is a good idea for several reasons: it is quick, allows for simple re-entry, and is not interactive. Remember the attacker cannot see any traffic coming from the target, so any responses are sent off into oblivion.)

--[Why it Works]--

IP-Spoofing works because trusted services only rely on network address based authentication. Since IP is easily duped, address forgery is not difficult. The hardest part of the attack is in the sequence number prediction, because that is where the guesswork comes into play. Reduce unknowns and guesswork to a minimum, and the attack has a better chance of succeeding. Even a machine that wraps all its incoming TCP bound connections with Wietse Venema's TCP wrappers, is still vulnerable to the attack. TCP wrappers rely on a hostname or an IP address for authentication...

[SECTION III. PREVENTITIVE MEASURES]

...A stitch in time, saves nine...

--[Be Un-trusting and Un-trustworthy]--

One easy solution to prevent this attack is not to rely on address-based authentication. Disable all the `r*` commands, remove all `.rhosts` files and empty out the `/etc/hosts.equiv` file. This will force all users to use other means of remote access (telnet, ssh, skey, etc).

--[Packet Filtering]--

If your site has a direct connect to the Internet, you can use your router to help you out. First make sure only hosts on your internal LAN can participate in trust-relationships (no internal host should trust a host outside the LAN). Then simply filter out `*all*` traffic from the outside (the Internet) that purports to come from the inside (the LAN).

--[Cryptographic Methods]--

An obvious method to deter IP-spoofing is to require all network traffic to be encrypted and/or authenticated. While several solutions exist, it will be a while before such measures are deployed as defacto standards.

--[Initial Sequence Number Randomizing]--

Since the sequence numbers are not chosen randomly (or incremented randomly) this attack works. Bellovin describes a fix for TCP that involves partitioning the sequence number space. Each connection would have its own separate sequence number space. The sequence numbers would still be incremented as before, however, there would be no obvious or implied relationship between the numbering in these spaces. Suggested is the following formula:

[IP-spoofing Demystified] (Trust-Relationship Exploitation)

by daemon9 / route / infinity
for Phrack Magazine
June 1996 Guild Productions, kid

ISN=M+F(localhost,localport,remotehost,remotepoort)

Where M is the 4 microsecond timer and F is a cryptographic hash. F must not be computable from the outside or the attacker could still guess sequence numbers. Bellovin suggests F be a hash of the connection-id and a secret vector (a random number, or a host related secret combined with the machine's boot time).

[SECTION IV. SOURCES]

- Books: TCP/IP Illustrated vols. I, II & III
- RFCs: 793, 1825, 1948
- People: Richard W. Stevens, and the users of the Information Nexus for proofreading
- Sourcecode: rbone, mendax, SYNflood

This paper made possible by a grant from the Guild Corporation.