TCP Sequence Prediction Attack

A TCP sequence prediction attack is an attempt to hijack an existing TCP session by injecting packets which pretend to come from one computer involved in the TCP session.

The TCP Sequence Prediction Attack

TCP is a reliable connection-oriented layer 4 (Transport Layer) protocol. Packet transfer between hosts is accomplished by the layers below layer 4 and TCP takes responsibility to making certain the packets are delivered to higher layers in the protocol stack in the correct order. To accomplish this reordering task, TCP uses the sequence number field.

To successfully mount a TCP sequence prediction attack, you must first listen to communications between two systems, one of which is your target system. Then, you issue packets from your system to the target system with the source IP address of the trusted system that is communicating with the target system.

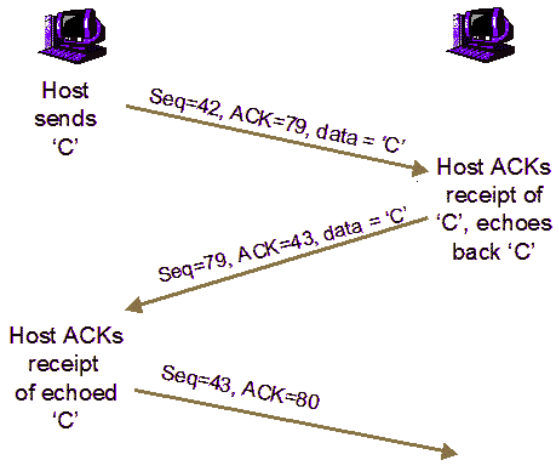

The packets you issue must have the sequence numbers that the target system is expecting. In addition, your packets must arrive before the packets from the trusted system whose connection you are hijacking. To accomplish this, it is often necessary to flood the trusted system off of the network with some form of denial of service attack. Once you have taken over the connection, you can send data to allow you to access the target host using a normal TCP/IPconnection. The most simple way to do this is:

```
echo "+ +" > /.rhosts
```

This specific technique relies upon inherent weaknesses in the BSD Unix `r` services. However, SunRPC, NFS, X-Windows, and many other services which rely upon IP address authentication can be exploited with a TCP sequence prediction attack.

Why are TCP Sequence Prediction Attacks Possible?

An excerpt from RFC 793 (Transmission Control Protocol) concerning the generation of TCP sequence numbers:

> When new connections are created, an initial sequence number (ISN) generator is employed which selects a new 32 bit ISN. The generator is bound to a (possibly fictitious) 32 bit clock whose low order bit is incremented roughly every 4 microseconds. Thus, the ISN cycles approximately every 4.55 hours. Since we assume that segments will stay in the network no more than the Maximum Segment Lifetime (MSL) and that the MSL is less than 4.55 hours we can reasonably assume that ISN's will be unique.

The developers of the BSD Unix TCP/IP stack did not follow these recommendations. TCP/IP stacks based upon BSD Unix increase the sequence number by 128,000 every second and by 64,000 for every new TCPconnection. This is significantly more predictable than the algorithm specified in the RFC.

Defending Against TCP Sequence Prediction Attacks

TCP sequence prediction attacks can be effectively stopped by any router or firewall that is configured not to allow packets from an internal IP address to originate from an external interface.
These does not fix the TCP sequence prediction vulnerability, it simply prevents TCP sequence prediction attacks from being able to reach their targets.

Diagram of the TCP Header

> TCP Header Format
>
> ----------
>
> 0 1 2 3

```
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1

+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

| Source Port | Destination Port |

+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

| Sequence Number |

+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

| Acknowledgment Number |

+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

| Data | |U|A|P|R|S|F| |

| Offset| Reserved |R|C|S|S|Y|I| Window |

| | |G|K|H|T|N|N| |

+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

| Checksum | Urgent Pointer |

+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

| Options | Padding |

+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

| data |

+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Source: http://www.tech-faq.com/tcp-sequence-prediction-attack.html