

# Network Scanning Techniques

Understanding how it is done

**Ofir Arkin**

[ofir@publicom.co.il](mailto:ofir@publicom.co.il)  
<http://www.publicom.co.il>  
<http://www.sys-security.com>

November, 1999



Copyright © Ofir Arkin, 1999

# 1.0 Introduction

## 1.1 Introduction to Intelligence Gathering Techniques

Imagine the following scenario:

A military target is to be attacked. What's the first step considered? – Gathering Intelligence, naturally. To do so, a satellite will photo the target zone and a special recon unit will patrol the area with maximum caution to eliminate the possibility of detection. After enough information has been gathered, a wing of stealth bombers will bomb the target – Mission accomplished.

Gathering intelligence is extremely important. If the amount of data collected is not sufficient, or alternatively, if the target is tightly defended, no attack will be launched.

The same is true of computer hacking.

An intelligent hacker will conduct a lot of research before attempting to gain privileged access to your systems.

If the intelligence gathered shows a poorly defended computer system, an attack will be launched, and unauthorized access will be gained.

However, if the target is highly protected, the hacker will think twice before attempting to break in. It will be dependent upon the tools and systems that protect the target.

Again, the key here is the amount of information he has gathered beforehand.

In the computer hacking world, intelligence gathering can be roughly divided into three major steps:

- Foot printing
- Scanning
- Enumeration

**Foot Printing** - The information collected by the hacker makes a unique *footprint* or a *profile* of an organization security posture.

With *foot printing*, using rather simple tools, we gather information such as:

1. Administrative, technical, and billing contacts, which include employee names, email addresses, and phone & fax numbers.
2. IP address range
3. DNS servers
4. Mail servers

And we can also identify some of the systems that are directly connected to the Internet.

Most of the information here can be freely accessed on the Internet.

**Scanning** - The art of detecting which systems are alive and reachable via the Internet, and what services they offer, using techniques such as *ping sweeps*, *port scans*, and *operating system identification*, is called *scanning*.

The kind of information collected here has to do with the following:

1. TCP/UDP services running on each system identified.
2. System architecture (Sparc, Alpha, x86).
3. Specific IP addresses of systems reachable via the Internet.
4. Operating system type.

Scanning can be compared with a thief checking all the doors and windows of a house he wants to break into.

**Enumeration** - *Enumeration* is the process of extracting valid accounts or exported resource names from systems. The information is gathered using active connections to systems and queries, which is more intrusive in nature than foot printing and scanning.

The techniques are mostly operating system specific, and can gather information such as:

1. User & group names.
2. System banners
3. Routing tables
4. SNMP information

This article will focus on scanning, normally the second phase of computer intelligence gathering technique.

## 1.2 Introduction to scanning

Today the number of automated scanners is constantly increasing, and as a result, more and more attacks are successfully initiated.

In order to be better prepared, we need to fully understand the scanning tools and the methods that these tools are using against us.

The questions we need to ask ourselves are:<sup>1</sup>

- What are scanners doing?
- What do they look like (signature)?
- How they operate in order to accomplish their tasks?
- What kind of information is collected?
- How serious is the threat?

We need to identify the intruder's behavior and understand the scanning techniques. If we have an intrusion detection system, or planning on implementing one in the future, finding scanning patterns in our log (manually, or automatically by the IDs) will give us an indication of a probable upcoming attempt to gain unprivileged access to our systems.

Only after we understand scanning techniques we can try to protect ourselves against them.

---

<sup>1</sup> John Green, NSW Shadow Team, Identifying Scanners in the Wild.

## 2.0 PING Sweeps

### 2.1 ICMP sweeps (ICMP ECHO requests)

We can use ICMP packets to determine whether a target IP address is alive or not, by simply sending an *ICMP ECHO request* (ICMP type 8) packets to the targeted system and waiting to see if an *ICMP ECHO reply* (ICMP type 0) is received. If an ICMP ECHO reply is received, it means that the target is alive; No response means the target is down.

Querying multiple hosts using this method is referred to as *Ping Sweep*. Ping Sweeps is the most basic step in mapping out a network.

This is an older approach to mapping, and the scan is fairly slow.

Some of the tools used for this kind of scan include –

UNIX:

- *fping* & *gping*<sup>2</sup>
- *nmap*<sup>3</sup>

Windows:

- *Pinger* from Rhino9<sup>4</sup>

*Pinger* is one of the fastest ICMP sweep scanners. Its advantage lies in its ability to send multiple ICMP ECHO packets concurrently and wait for the response. It also allows you to resolve host names and save the output to a file.

Blocking ICMP sweeps is rather easy, simply by not allowing ICMP ECHO requests into your network from the void.

If you are still not convinced that you should block ICMP ECHO requests, bear in mind that you can also perform *Broadcast ICMP's*.

### 2.2 Broadcast ICMP

Sending ICMP ECHO request to the network or/and broadcast addresses will produce all the information you need for mapping a targeted network in even a simpler way.

The request will be broadcast to all alive hosts on the target network, and they will send ICMP ECHO reply to the attacker source IP after only one or two packets have been sent by him.

Here we can first distinguish between Unix and Windows machines. While Unix machines often still answers to requests directed to the network address (the answer will be the fully qualified network address), Windows machines will ignore it.

---

<sup>2</sup> [ftp.tamu.edu/pub/Unix/src/](http://ftp.tamu.edu/pub/Unix/src/)

<sup>3</sup> <http://www.insecure.org>

<sup>4</sup> <http://207.98.195.250/software>

## 2.3 Non-ECHO ICMP

Blocking incoming ICMP ECHO requests is not enough. We can use non-ECHO ICMP protocols for gathering various information about a system.

Good examples are ICMP type 13 messages (*TIMESTAMP*), and ICMP type 17 messages (*ADDRESS MASK REQUEST*).

ICMP timestamp request and reply allow a system to query another for the current time.

The ICMP address mask request (and reply) is intended for diskless systems to obtain its subnet mask at bootstrap time. We can use it to request the netmask of a particular device.

We can use the *icmpush* & *icmpquery* tools to perform this kind of scanning.

Many firewalls are configured to block only ICMP ECHO traffic, and in this case it makes the non-ECHO requests a valid form of host identification.

Even if ICMP traffic is blocked on the border router or firewall, there are additional techniques that can be used to determine which systems are actually alive, although these techniques are not as accurate as a normal ICMP Sweep.

## 2.4 TCP Sweeps

The TCP connection establishment process is called "*the three way handshake*", and is combined of three segments.

1. A client sends a SYN segment specifying the port number of a server that the client wants to connect to, and the client *initial sequence number*.
2. If the server's service (or port) is active the server will respond with its own SYN segment containing the server's initial sequence number. The server will also acknowledge the client's SYN by ACKing the client's SYN+1.

If the port is not active, the server will send a RESET segment, which will reset the connection.

3. The client will acknowledge the server's SYN by ACKing the servers ISN+1.

When will a RESET be sent? – Whenever an arriving segment does not appear correct to the *referenced connection*. Referenced connection means the connection specified by the destination IP address and port number, and the source IP address and the port number <sup>5</sup>.

With the TCP Sweep technique, instead of sending ICMP ECHO request packets we send TCP ACK or TCK SYN packets (depending if we have root access or not) to the target network. The port number can be selected to meet our needs. Usually a good pick would be one of the following ports – 21 / 22 / 23 / 25 / 80 (especially if a firewall is protecting the targeted network). Receiving a response is a good indication that something is up there. The response depends on the target's operating system, the nature of the packet sent and any firewalls, routers or packet-filtering devices used.

Bear in mind that firewalls can spoof a RESET packet for an IP address, so TCP Sweeps may not be reliable.

---

<sup>5</sup> RFC 793, <http://www.ietf.org/rfc/rfc0793.txt>

*Nmap* and *Hping*<sup>6</sup> are tools that support TCP Sweep, both for the Unix platform. Hping even adds an additional option to *fragment packets*, which allows the TCP packet to pass through certain access control devices.

An example with nmap:

```
[root@mia /root] ./nmap -sP -PT80 192.168.2.0/24
```

TCP probe port is 80

Starting nmap V. 2.2-BETA4 by Fyodor (fyodor@dhp.com, www.insecure.org/nmap/)

Host host1.MyDomain.com (192.168.2.0) appears to be up.

Host host2.MyDomain.com (192.168.2.1) appears to be up.

Host host3.MyDomain.com (192.168.2.2) appears to be up.

Host host4.MyDomain.com (192.168.2.3) appears to be up.

Host host5.MyDomain.com (192.168.2.4) appears to be up.

Host host6.MyDomain.com (192.168.2.5) appears to be up.

...

Host host254.MyDomain.com (192.168.2.254) appears to be up.

Nmap run completed -- 32 IP addresses (13 hosts up) scanned in 12 seconds

## 2.5 UDP Sweeps (Also known as UDP Scans)

This method relies on the ICMP PORT UNREACHABLE message, initiated by a closed UDP port. If no ICMP PORT UNREACHABLE message is received after sending a UDP data gram to a UDP port that we wish to examine on a targeted system, we may assume the port is opened.

UDP scanning is unreliable because of a number of reasons<sup>7</sup>:

- Routers can drop UDP packets as they cross the Internet.
- Many UDP services do not respond when correctly probed.
- Firewalls are usually configured to drop UDP packets (except for DNS).
- UDP sweep relies on the fact that a non-active UDP port will respond with an ICMP PORT UNREACHABLE message.

## 3.0 Port Scanning

Ping Sweeps help us identify which systems are alive. The next step is trying to determine what services (if any) are running or in a LISTENING state on the targeted system, by connecting to the TCP and UDP ports of that system. This is called – Port Scanning.

For the hacker it is critical to identify listening ports, because it helps him identify the operating system and application in use.

The services detected as listening may suffer from vulnerabilities which may result from two reasons:

- Misconfiguration of the service
- The version of the software is known to have security flaws

If identified, these vulnerabilities can lead to unprivileged access gained by the attacker.

We will further discuss port scanning types, techniques, and tools.

---

<sup>6</sup> <http://www.kyuzz.org/antirez>

<sup>7</sup> Ron Gula, How to Handle and Identify Network Probes, Netowrk Defense Consulting.

## 3.1 Port Scanning Types

### 3.1.1 TCP connect() scan

With this type of scan we use the basic TCP connection establishment mechanism. To open a connection to an interesting port on the targeted machine:

1. A SYN packet is sent to the target's system interesting port.
2. Now we wait to see what type of packet is sent back from the target.
  - If a SYN/ACK packet is received it usually means the port is in a LISTENING state.
  - If a RST/ACK packet is received, it usually means the port is not LISTENING and the connection will RESET.
3. We finish the three-way handshake (if SYN/ACK packet was received) by sending an ACK.
4. A connection is terminated after the full connection establishment process has been completed.

This kind of scan is easily detected. Inspecting the target system log will show a number of connections and error messages immediately after each one of them was initiated.

### 3.1.2 TCP SYN Scan (half open scanning)

This type of scan differs from TCP connect() scan because we do not open a full TCP connection. You send a SYN packet to initiate the three-way handshake and wait for a response. If we receive an SYN/ACK it indicates the port is LISTENING. If we receive an RST/ACK it indicates a non-LISTENING port. If we do receive a SYN/ACK packet we immediately tear down the connection by sending a RESET.

Because the TCP three-way handshake was not completed some of the sites will probably not log these scanning attempts.

### 3.1.3 Stealth Scan <sup>8</sup>

Chris Klaus was one of the first people to write a paper about stealth scans. In a paper called "stealth scanning – Bypassing Firewalls/SATAN Detectors" <sup>9</sup>, he describes a technique which intentionally violates the TCP three-way handshake. This technique is what people refer to today as "half-open" scanning.

Today, some people use the "stealth" term to mean NULL flags (no flags or code bits set).

"Stealth" can also be defined as a scanning technique family, doing one of the following:

- Pass through filtering rules.
- Not to be logged by the targeted system logging mechanisms.
- Try to hide themselves at the usual site / network traffic.

---

<sup>8</sup> Stephen Northcutt, Network Intrusion Detection an Analyst's Handbook, New – Riders, 1999.

<sup>9</sup> <http://www.netsys.com/firewalls/firewalls-9512/0085.html>

### 3.1.3.1 Explicit Stealth Mapping Techniques

The frequently used explicit mapping techniques are:

- SYN/ACK scan.
- FIN scans.
- XMAS scans.
- NULL scans.

According to RFC 793<sup>10</sup> closed ports are required to reply with a RESET packet to our probe packets, while open ports must ignore any packet in question.

This applies to all scanning methods described here.

#### 3.1.3.1.1 SYN/ACK

This scan intentionally disregards the TCP three-way handshake. We send a SYN/ACK packet, which is step two in the TCP three-way handshake, while there is no SYN packet sent for step one.

Sending SYN/ACK packet to a closed port:

Because TCP is stateful, it knows no SYN has been sent, which is the first step in the three-way TCP handshake. TCP figures this packet must be a mistake and sends a RESET to tear down the connection. This is what we wished for – any kind of response to give away the existence of the system and the fact that the probed port is closed.

If we send the SYN/ACK to an open port, it will ignore any such packet.

#### 3.1.3.1.2 FIN

The FIN scan uses the FIN packet as the probe. The hacker sends a FIN packet to a targeted port on a probed system and waits for a response<sup>11</sup>.

#### 3.1.3.1.3 XMAS (Christmas Tree)

XMAS is a scanning type, which sends a TCP packet with the URG, ACK, PST, RST, SYN and FIN flags set. All the TCP flags are set.

#### 3.1.3.1.4 NULL

Null scan is a scanning type, which sends a TCP packet that turns off all flags. The probed system should send back a RESET to all closed ports.

According to RFC 793 this should work against every implementation of TCP regardless of the operating system it runs on. Life is not always simple. Windows, CISCO, BSDI, HP/UX, MVS & IRIX have a broken TCP implementation – they send RESETs to open ports as well.

One method suggested by Fyodor<sup>12</sup> (nmap tool programmer) is – if you scan using FIN, XMAS or NULL and the results show all ports are closed, and then you SYN scan and find out

---

<sup>10</sup> <http://www.ietf.org/rfc/rfc0793.txt>

<sup>11</sup> Uriel Maimon, Port Scanning without the SYN flag, Phrack 49, Volume Seven, Issue Forty Nine.

<sup>12</sup> Fyodor, The art of port scanning, Phrack 51



open ports it probably means the probed system belong to the “weird” group of operating systems.

### **3.1.3.2 Inverse Mapping**

A group of techniques, which gathers information about hosts or networks that are unreachable. By having a clue about what is not there we make assumptions about where things are.

#### **3.1.3.2.1 RESET Scans** <sup>13</sup>

Routers will give away internal architecture information of a network, even if the question they were asked does not make any sense.

A router looks at the IP address and makes decisions based on that. The RESET, which is out of place, is not taken into any consideration. The router will report addresses that do not exist with host (or net) unreachable or time exceeded message.

We can probe the router for a list of IPs. The answers we receive from the router will help us map the infrastructure of the targeted network IPs - because if we don't receive the HOST UNREACHABLE or TIME EXCEEDED replies about a certain IP, we can assume the IP exists.

In the past, RESET scans used a fixed ACK number in the ACK field which was easily detected. Recent RESET scans use random acknowledgment numbers, and if you generate decoys with the RESET scan it makes it even harder to detect.

#### **3.1.3.2.2 Domain Query Answers**

The probing computer sends answers to domain questions that were never asked. The goal here is to get ICMP unreachable messages for a host or a subnet that do not exist.

### **3.1.4 Proxy Scanning / FTP Bounce Scanning**

The FTP protocol supports the following scenario:

Attacker.com connects to an FTP server, which has a world writable directory, and establishes a control communication connection. The attacker can then ask the FTP server to initiate an active server data transfer process and send a file anywhere on the Internet, presumably to a user data transfer process.

Hobbit wrote a post to Bugtrack on July 12 1995 where he described this kind of attack<sup>14</sup>. Under the “other possibilities” section, he wrote that this protocol flaw “can be used to post virtually untraceable mail and news, hammer on services of various sites, fill up disks, try to hop firewalls, and generally be annoying and hard to track down at the same time”.

This method can be used to scan TCP ports from a “proxy” FTP server.

---

<sup>13</sup> Analysis Techniques for Detecting Coordinated Attacks and Probes, John Green, David Merchette, Stephen Northcutt and Bill Ralph, Proceedings of the USENIX Workshop on Intrusion detection and Network Monitoring.

<sup>14</sup> [www.geek-girl/bugtrack/1995\\_3/0047.html](http://www.geek-girl/bugtrack/1995_3/0047.html)

An example can be scanning behind a firewall – connect to an FTP server behind a firewall, and then try to scan ports that the firewall blocks. If a directory is writable for the account you are using on the FTP server, you can also send data to the ports you find open<sup>15</sup>.

Nmap supports this kind of scan and uses the PORT FTP command to declare that our passive user data transfer process is listening on the target box at a certain port number. We then use the LIST FTP command to try to list the current directory. The result is sent over the server data transfer process channel.

If the transfer is successful (150 and 226 response), the target host is listening on the specified port scanned. Otherwise, a “425 Can’t build data connection: Connection refused” message will be received.

By using this method we can scan all the target’s ports simply by issuing PORT commands one after another.

This scan is rather slow. Some FTP servers disable the “Proxy” feature, but there are still many who do not, making this kind of scanning still available.

### **3.1.5 TCP Reverse Ident Scanning**

The ident protocol (RFC 1413<sup>16</sup>) is used to determine the owner username of a particular TCP connection by communicating with port 113.

This involves opening a full TCP connection to the target machine port.

The scanning method gives the attacker information that helps determine which server to attack.

If a certain server is running as “root”, and one of its services was compromised, attackers can gain instant root access. So it is more common that in a group of services that can be attacked, the services which are at the highest privilege will be attacked first.

## **3.2 Port Scanning Techniques**

### **3.2.1 “Random” Port Scan**

Many commercial intrusion detection systems and firewalls are looking for sequential connection attempts. When the pattern is matched a port scan is reported.

Randomizing the sequence of ports probed may prevent detection.

### **3.2.2 Slow Scan**

Intrusion detection systems can determine if a specific IP tries to port scan the network they are defending. It is done by analyzing the network traffic over a certain amount of time.

The amount of time is called the site detection threshold.

Some hackers are very patient and can use network scanners that spread out the scan over a long period of time.

The scan rate can be, for example, as low as 2 packets per day per target site.

---

<sup>15</sup> Fyodor, The Art of Port Scanning, Phrack 51

<sup>16</sup> [www.ietf.org/rfc/rfc1413.txt](http://www.ietf.org/rfc/rfc1413.txt)

If the attacker can guess the detection threshold of its target, he can reduce the chances of detection to a minimum or even to no detection at all, as long as he doesn't include a signature with his packet that alerts the intrusion detection system in other way.

### 3.2.3 Fragmentation Scanning

All IP packets that carry data can be fragmented.

If we need to send information using TCP and to fragment our data, normally the destination port and source port along with the flags field will be "traveling" at the first packet sent.

RFC 791<sup>17</sup> defines the minimum and maximum sizes of fragments.

In the case of TCP the 8 octets of data (minimum fragment size) are enough to contain the source and destination port numbers. This will force the TCP flags field into the second fragment<sup>18</sup>.

Some filtering devices and intrusion detection systems may incorrectly reassemble or completely miss portions of the scan. They may assume that this was just another segment of traffic that has already passed through their access list.

Filtering devices that queue all IP fragments can handle this method. Linux is a good example with the CONFIG\_IP\_ALWAYS\_DEFRAG kernel option. Some networks cannot afford the performance hit this causes and disable this feature<sup>19</sup>.

This kind of scan has been fixed in most vendors' products.

### 3.2.4 Decoy

Some network scanners include options for Decoys or spoofed addresses in their attacks.

It would appear to the attacked network / host that the host(s) you specified as decoys are scanning them as well. This will drive intrusion detection systems into thinking that the target network is being port scanned by all the hosts, and determining who the real attacker is, will be nearly impossible.

One way that helped intrusion detection systems detect the decoy hosts in the past was the TTL (Time to Live) field values in the scanned packets. If all the incoming packets TTL values have the same value, it is likely that they were generated in the same "factory".

If the attacker is using nmap intrusion detection systems cannot use this method since nmap generates random TTL fields between 51-65.

Another way of detecting the real scan among the decoys is to try to traceroute the source IP. If the attacker used non-routable IPs or the IPs belongs to a host that is not up something is suspicious. This can also result in a denial of service.

Attackers solved this by using spoofed IPs of hosts that were up.

Another aspect of detection is using IPs rather than names, which is a really smart move, because the decoys network will not see the attackers IP in their name server's log<sup>20</sup>.

---

<sup>17</sup> <http://www.ietf.org/rfc/rfc0791.txt>

<sup>18</sup> RFC 1858 <http://www.ietf.org/rfc/rfc1858.txt>

<sup>19</sup> nmap network security scanner man page, [http://insecure.org/nmap/nmap\\_manpage.html](http://insecure.org/nmap/nmap_manpage.html)

<sup>20</sup> Ron Gula, How to Handle and Identify Network Probes, Network Defense Consulting.

### 3.2.5 Coordinated Scans

Probing a few target systems from a single IP within a certain amount of time will usually turn on the alarm of the intrusion detection systems.

We have already discussed a way to try to bypass this – using slow scans. But even a slow scan can sometimes be detected.

When a group of attackers are working together to achieve a common goal, trying to get unauthorized access on a targeted network for example, we call this – coordinated attacks.

Coordinated attacks can be used to target a single host or even an entire network.

If multiple IPs probe a target network, each one of them probes for a certain service on a certain machine in a different time period, and therefore it would be nearly impossible to detect these scans.

When coordinated attack efforts are aimed at a certain host, the detection (if any at all) is dependent on the time period these probes take place<sup>21</sup>.

Coordinated attacks, when intelligently launched, are the most discrete way of probing a target. Most of them are still undetected.

## 4.0 Operating System Detection

Because many security holes are operating system dependent, identifying which operating system runs on the target host / machine is of major importance.

An attacker can scan a range of IPs for opened TCP and UDP ports and the operating system type. Then he can put the list aside for a while.

When a security hole that is operating system dependent is discovered, all the attacker has to do is pull the list and look for a matching information.

### 4.1 Banner Grabbing

Some services can be used to identify an operating system. The TELNET service is the most notable example. If a system provides the TELNET service, just by telneting to the box and looking at the welcome banner we can, in most cases, identify the operating system:

```
[root@pooh] # telnet 192.168.1.13
Debian GNU/Linux 2.1 target.domain.com
target login:
```

Other services have banners that give the same information, for example the mail server:

```
220 target.domain.com ESMTP Sendmail 8.9.3/8.9.3/Debian/GNU; Thu, 28 Oct 1999 09:56:32 +0200
```

---

<sup>21</sup> Analysis Techniques for Detecting Coordinated Attacks and Probes, John Green, David Merchette, Stephen Northcutt and Bill Ralph, Proceedings of the USENIX Workshop on Intrusion detection and Network Monitoring.

Today, an increasing number of systems keep their banners turned off or make their banners display forged information.

Some applications leak information. A good example is the SYST command on FTP servers and IIS server:

```
[root@pooh] # telnet 192.168.1.17
```

```
get
HTTP/1.1 400 Bad Request
Server: Microsoft-IIS/4.0
Date: Thu, 28 Oct 1999 08:29:46 GMT
Content-Type: text/html
Content-Length: 87
```

```
<html><head><title>Error</title></head><body>The parameter is incorrect. </body></html>
```

## 4.2 DNS HINFO Record

The Host information record is a pair of strings identifying the host's hardware type and the operating system<sup>22</sup>.

```
www    IN      HINFO   "Sparc Ultra 5"    "Solaris 2.6"
```

This is an old technique that is rarely effective today because administrators avoid using this record.

## 4.3 TCP/IP Stack Fingerprinting

Stack fingerprinting is a technique that uses distinct variations in TCP stack implementation to determine the type of a remote operating system.

The idea is to send "specific" TCP packets to the target IP and observe the response. The response will be unique to a certain group or individual operating system(s). The response varies because one vendor's IP stack implementation is not the same as another. This comes from different interpretation of specific RFC guidelines when vendors wrote their TCP/IP stack.

The tools often used for stack fingerprinting are Queso<sup>23</sup> written by Savage, and Nmap. Nmap has a larger database of responses of operating systems than any other tool. In order to get maximum reliability, Nmap needs at least one port opened at the target host.

The definite information source of the subject is Fyodor's article "Remote OS Detection via TCP/IP Stack Fingerprinting".<sup>24</sup>

### 4.3.1 Types of Probes Used to Determine the Operating System Type

#### 1. The FIN Probe

A FIN packet is sent to an open port. RFC 793 states that the correct behavior is not to respond to the FIN packet.

Many stack implementations will respond with a RESET. This group includes: Windows, BSDI, CISCO, HP-UX, MVS, and IRIX with a RESET.

---

<sup>22</sup> Paul Albitz & Cricket Liu, DNS & BIND, third edition, O'reilly 1998.

<sup>23</sup> Queso, <http://apostols.org/projectz/queso>

<sup>24</sup> [www.insecure.org/nmap/nmap-fingerprintng-article.html](http://www.insecure.org/nmap/nmap-fingerprintng-article.html)

## 2. The Bogus Flag Probe

Here, a SYN packet with an undefined flag set is sent to the targeted host.

Machines running the Linux operating system with kernel prior to 2.0.35 will keep the flag set in their response.

Some operating systems will RESET the connection when getting this kind of probe.

## 3. TCP Initial sequence number sampling

Finding patterns of the initial sequence numbers chosen by the TCP implementations when responding to a connection request.

We can divide the answers into four groups:

- Traditional 64k (older UNIXs)
- Random increment (FreeBSD, DG-UX, IRIX, new versions of Solaris)
- True “random” (Linux)
- Time dependent modules (MS Windows)

## 4. “Don’t Fragment Bit”

To enhance performance some operating systems set the “Don’t fragment bit”. Monitoring this behavior can give the attacker more information about the target operating system.

## 5. TCP Initial Window

Some stack implementations have a unique initial window size on their returned packets.

AIX for example is the only operating system using the 0x3F25 value. OpenBSD and FreeBSD use 0x402E.

## 6. ACK Value

In some cases IP stacks even differ in the value they use for the ACK field.

For example, sending a FIN|PSH|URG to a closed port. Most implementations will set the acknowledgement number in the returned packet to be the same as the sequence number received. Windows responds with an ACK field that is the sequence number+1.

## 7. ICMP error Message Quenching

RFC 1812<sup>25</sup> suggests limiting the rate at which various error messages are sent. Only a few operating systems are known to follow this RFC.

An attacker can use this to send UDP packets to a random, high UDP port and count the number of unreachable messages received within a given amount of time.

---

<sup>25</sup> <http://www.ietf.org/rfc/rfc1812.txt>

## 8. ICMP Message Quoting

ICMP error messages should quote a small amount of information from the ICMP message that caused the error.

The information is quoted when the PORT UNREACHABLE message is received in the IP header + 8 bytes, with almost all the implementations. Solaris sends more information than is needed and Linux even more.

## 9. ICMP Error Message Echoing Integrity

When sending back an ICMP error message, some stack implementations may alter the IP header.

If an attacker examines the types of alternation that have been made to the headers, he may be able to make certain assumptions about the target operating system.

## 10. Type of Service (TOS)

When an ICMP PORT UNREACHABLE message is sent, an attacker can examine the type of service field.

Nearly all implementations use 0 for this value, Linux uses 0xC0.

## 11. Fragmentation Handling

Different stack implementations handle overlapping fragments differently. This was pointed out by Thomas Ptacek and Tim Newsham in their paper "Insertion, Evasion, and Denial of Service: Eluding Network Intrusion Detection".

Some implementations will either overwrite the old data portions with the new data or vice versa, when the fragments are reassembled.

## 12. TCP Options

RFC 793 defines the TCP options. RFC 1323<sup>26</sup> defines the more advanced TCP options.

- Not all hosts implement TCP options
- When sending a query with an option set to a targeted host, the target host will set the option in the reply only if it supports it.
- We can test all the options at the same time if we send one packet that includes all the options.

When you examine the response packet, you look at the Options field for Options that were set. These are the supported options.

Some operating systems support all the advanced options while others support very few.

---

<sup>26</sup> <http://www.ietf.org/rfc/rfc1323.txt>

## 5.0 Firewalking

Firewalking<sup>27</sup> is a technique used to gather information about a remote network protected by a firewall.

The technique is being used for two purposes:

- Determining the rule set or ACL of a firewall or other packet-filtering device (mapping open ports on a firewall).
- Mapping a network behind a firewall. When a firewall's policy is to drop ICMP ECHO Request/reply this technique is very effective.

### How does Firewalking work?

It is using a traceroute-like packet filtering to determine whether or not a particular packet can pass through a packet-filtering device.

Since traceroute is dependent on the IP layer (TTL field), any transport protocol can be used the same way (TCP, UDP, and ICMP).

For Firewalking we need two pieces of information in advance:

- The IP address of the last known gateway before the firewalking takes place
- The IP address of a host located behind the firewall.

The first IP address serves as our waypoint. The second IP address is used as a destination to direct the packet flow.

If we try to traceroute the machine behind a firewall and get blocked by an ACL filter that prohibits the probe, we can only determine what the last gateway which responded was – probably the firewall. The firewall then becomes a waypoint for further investigations. We try again to traceroute the same machine, this time we use a different traceroute-like probe using a different transport protocol. If we get a response we can conclude the following:

- That particular traffic is allowed by the firewall
- We know a host behind the firewall

If we are continuously kept blocked by the ACL filters at our waypoint, we know that this kind of traffic is blocked.

Trying to pass packets on all ports and protocols through the firewall and monitor the response, will produce the ACL.

Sending packets to every host behind the packet-filtering device can generate an accurate map of a network's topology.

---

<sup>27</sup> Firewalking - A Traceroute-Like Analysis of IP Packet Responses to Determine Gateway Access Control Lists, <http://www.packetfactory.net/firewalk/firewalk-final.html>



## 6.0 Conclusion

We have reviewed some scanning types combined with hard-to-detect or even non-detectable scanning techniques.

Understanding the importance of detecting these scans can prevent, in some cases, an intrusion to the systems that were scanned. The detection can be partly achieved by implementing an Intrusion Detection system<sup>28</sup>. Tuning it right is another issue of major importance. The second part is maintenance of the system, getting information on new and wicked scanning types and techniques, understanding their signatures, and implementing new filters to detect them.

The number of simple automated scans is constantly on the rise, and consequently the number of attempts to get intelligence and eventually trying to hack into sites.

Tightening your security to the maximum can drive some attackers away.

Some attackers will take the challenge.

Identifying these probing attempts will give you an indication that an upcoming attack might be on its way.

---

<sup>28</sup> I recommend you check out Network Flight Recorder at <http://www.nfr.net>.