Global Information Assurance Certification Paper

Firewalk : Can Attackers See Through Your Firewall?

David Irby
SANS GIAC Level One Security Essentials Practical
MBUS 511

Introduction

In order to complete an information security attack there are certain steps or phases an attacker must complete [1]. The first phase is information gathering wherein the attacker tries to gain enough information about the target systems to enable an actual attack. The second phase is the actual exploitation of a target system. The third and final phase (metastasis) involves hiding the traces of the initial exploit and then installing tools to allow continued access as well as gathering more information to be used in expanding the attack. In the 'information gathering' phase the attacker will generally attempt to ascertain the identities of interesting hosts and determine which services might be available on those hosts. Any information about the target network topology and operating systems being attacked is also valuable. I have chosen to report on a methodology and supporting tools which aid the information gathering functions of service availability and network topology. The Firewalk methodology as devised by David Goldsmith and Michael Schiffman [2] uses a traceroute-like method to map out the services allowed through a firewall or other access controlled device. A firewall is generally expected to hide the details of the protected network from the outside world. The Firewalk tool shows that this is not always the case and attackers may be able to learn more about your systems than you expect. In order to understand the method used by the Firewalk tool we need to first understand the workings of the traceroute command.

Traceroute

Traceroute is a networking utility designed to list the routers involved in making a connection from one host to another across a network. It lists the number of hops the packets take and the IP addresses of each router along the way. In order to determine this information traceroute relies on the IP time to live (TTL) feature [3]. The time to live feature was implemented in IP to prevent packets from looping indefinitely in the network. As each device receives a packet it decrements the time to live counter and if the counter is less than or equal to zero the packet is dropped and an ICMP "TTL Exceeded in Transit" error message is generated and returned to the originator. This error message will contain the IP address of the router dropping the packet as the originator. Traceroute uses this behavior and manipulates the TTL counter so that each router on the way to the target host will generate the error message and thus reveal its IP address. The Windows version (tracert.exe) uses pings (ICMP Echo) as the packets being sent while Unix versions of traceroute generally use UDP datagrams. The datagrams are sent to port 33434 by default and the port number is incremented for each successive packet. It is common for traceroute to send 3 packets (to successive ports) with the same TTL value to guard against packet loss. Below is a sample of the output from the Windows tracrt.exe program:

```
C:\WINDOWS>tracert quote.yahoo.com

Tracing route to finance.yahoo.com [204.71.203.155]
over a maximum of 30 hops:
```

```
 1   99 ms   100 ms   119 ms  tnt3.culpeper.va.da.uu.net [206.115.221.174]
 2   99 ms   119 ms   115 ms  206.115.233.205
 3  106 ms   104 ms   102 ms  Fddi0-0.HR1.DCA1.ALTER.NET [137.39.33.130]
 4  112 ms    95 ms   113 ms  102.ATM3-0.XR1.DCA1.ALTER.NET [146.188.160.254]
 5  103 ms    98 ms   104 ms  195.at-7-2-0.XR1.DCA8.ALTER.NET [146.188.163.6]
 6   98 ms   111 ms   111 ms  POS6-0.BR3.DCA8.ALTER.NET [152.63.36.5]
 7  110 ms   102 ms   104 ms  137.39.52.18
 8  106 ms   104 ms   112 ms  pos2-0-155M.cr1.WDC2.gblx.net [208.178.174.53]
 9  172 ms   180 ms   167 ms  pos7-0-2488M.cr2.SNV.gblx.net [208.50.169.86]
10  168 ms   165 ms   167 ms  ge1-0-1000M.hr8.SNV.gblx.net [206.132.254.41]
11  168 ms   174 ms   165 ms  bas1r-ge3-0-hr8.snv.yahoo.com [208.178.103.62]
12  176 ms   169 ms   175 ms  finance.yahoo.com [204.71.203.155]

Trace complete.
```

Many firewalls are configured to block traceroute and ping traffic from the outside to prevent attackers from learning the details of the internal networks and hosts. The following example shows the tracert.exe output when a firewall or router access control list blocks the ping traffic:

```
C:\WINDOWS>tracert vanguard.com

Tracing route to vanguard.com [192.175.182.6]
over a maximum of 30 hops:

 1  103 ms    98 ms    97 ms  tnt3.culpeper.va.da.uu.net [206.115.221.174]
 2  105 ms   104 ms   104 ms  206.115.233.205
 3  103 ms    97 ms   104 ms  Fddi0-0.HR1.DCA1.ALTER.NET [137.39.33.130]
 4  101 ms   736 ms   103 ms  102.ATM2-0.XR2.DCA1.ALTER.NET [146.188.160.250]
 5  105 ms   105 ms   103 ms  294.at-7-2-0.XR2.DCA8.ALTER.NET [146.188.163.30]
 6  100 ms   104 ms   118 ms  POS7-0.BR2.DCA8.ALTER.NET [152.63.35.193]
 7  107 ms   105 ms   106 ms  uu-gw.wswdc.ip.att.net [192.205.32.133]
 8  103 ms   104 ms   103 ms  gbr4-p50.wswdc.ip.att.net [12.123.9.54]
 9  100 ms   102 ms    98 ms  gbr1-p60.wswdc.ip.att.net [12.122.1.221]
10  101 ms   117 ms   126 ms  ar1-a3120s4.wswdc.ip.att.net [12.123.8.45]
11  118 ms   103 ms   104 ms  12.127.47.50
12    *        *        *     Request timed out.
13    *        *        *     Request timed out.
```

As you can see we are unable to complete the trace and begin receiving timeout messages at the host which drops the ping packets. We are unable to determine any information beyond this system.

Firewalking

The traceroute program will show the hosts up to and including the system which is dropping the packets. The firewall stopping the flow of traffic will still respond to the traceroute packet directed at itself but will not allow further packets to pass on their way to the target system. Since firewalls are installed for a useful purpose there must be some sort of traffic allowed through, even though the packets used by traceroute are blocked. The firewalking methodology is based on determining what traffic types are allowed and then using those packet types as the basis for further traceroute type scanning. A common firewall implementation might be to only allow DNS queries (UDP port 53).

Thus, if we can send traffic to UDP port 53 with the next TTL value it will pass through the initial firewall and return information about the next host in line. Since the traceroute functionality is built on the handling of the TTL field at the IP level any of the various transport mechanisms (UDP, TCP or ICMP) can be used and thus any service based on those protocols may be spoofed.

Once a firewall has been identified along the path to the target host scanning that system with the firewalking methodology will reveal the open ports on that system. These ports will be known even if the next system down the line refuses to pass information on the target port. This information can be used to map the overall access control lists for each of the firewalls along the way. If each host along the path is not inspected there is the possibility of falsely reporting ports closed when the traffic is actually blocked by some intermediate system. The following diagram illustrates how this might happen.

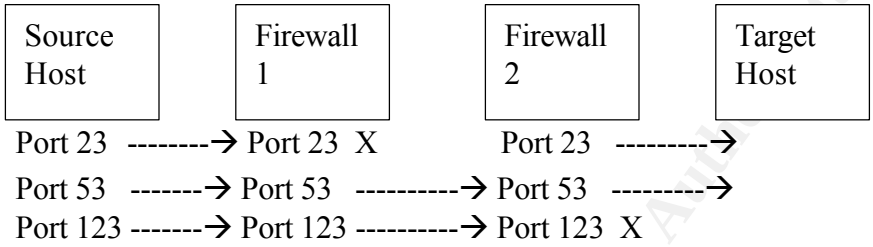

In the above scenario a firewalking scan of "Firewall 2" towards the "Target Host" leads to the conclusion that port 23 is closed on "Firewall 2" when the traffic is actually blocked by "Firewall 1" and was never received at "Firewall 2". To avoid false negatives all hosts must be scanned along the way to the target. This process is of course much slower than starting the scan at the furthest detected firewall system.

One of the greatest threats posed by the firewalking scan is that most firewalls do not log traffic on allowed ports. A careful and patient attacker could easily collect a wealth of information about the systems inside your firewall leaving no traces of their presence in the firewall logs.

Firewalk – The Firewalking Tool

The authors of the Firewalking paper [2] have also developed a proof of concept tool named Firewalk which has become quite popular in the security community [4]. The Firewalk tool implements the firewalking strategy and includes the full scanning of all intermediate hops across the network to the target. This prevents the false negative reports as described above. Firewalk is currently available on Linux and has recently been upgraded with a graphical user interface based on the GTK toolkit. The Firewalk tool allows any ports to be scanned but does not attempt to actually spoof the service being attempted. Because of this some firewalls which inspect the actual packet contents may stop Firewalk scans even on ports which are allowed. Intrusion Detection tools can also use this behavior to detect Firewalk scans. The Firewalk tool is available from http://www.packetfactory.net/firewalk/.

The Hping2 tool is another popular network security tool and has also implemented a firewalking type scan. It is available at http://www.kyuzz.org/antirez/hping/.

Conclusion

Firewalking can be stopped by blocking all outgoing TTL Exceeded in Transit packets in the firewall or by using Network Address Translation to hide the addresses on your internal networks. If a host on the other side of the firewall cannot be targeted then firewalking will not be successful.

One of the most important points to take away from this report is that a single layer of defense is never enough. There are many very clever people in the world and new ideas are springing up daily in the race to obtain protected data. It is certainly clear to me now that I cannot trust firewalls as my sole source of security and I cannot expect a firewall to prevent attackers from learning about my network and systems. The defense in depth strategy in which even hosts protected by firewalls implement strong security measures and host based intrusion detection seems more sensible than ever. One of the major points in the training so far has been never to place too much trust in any one form of security. Learning about the firewalking tool has convinced me that firewalls cannot be trusted to hide your systems details from the outside and that all systems must be protected with multiple layers of overlapping security.

References

[1] Stewart, Andrew J. "Distributed Metastasis : A Computer Network Penetration Methodology". August 12, 1999. URL: http://www.packetfactory.net/Papers/index.html (9 December 2000)

[2] Goldsmith, David and Schiffman, Michael. "Firewalking : A Traceroute-Like Analysis of IP Packet Responses to Determine Gateway Access Control Lists". October 1998. URL: http://www.packetfactory.net/firewalk/firewalk-final.html (1 December 2000)

[3] Graham, Robert. "FAQ: Firewalls: What am I seeing?". January 15, 2000. URL: http://www.robertgraham.com/pubs/firewall-seen.html (4 December 2000)

[4] Insecure.Org. "Top 50 Security Tools". August 19, 2000. URL: http://www.insecure.com/tools.html (1 December 2000)

[5] Lynn, Karl. "Strategic Scanning and Assessment of Remote Hosts (SSARH)". June 14, 1999. URL: http://www.attrition.org/security/newbie/pen/ssarh.html (6 December 2000)

[6] "Firewalking, a new method to gather information on a remote host." October 29, 1998. URL: http://www.securiteam.com/unixfocus/Firewalking__a_new_method_to_gather_information_on_a_remote_host.html (1 December 2000)