

Global Information Assurance Certification Paper

Copyright SANS Institute Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permited without express written permission.

Interested in learning more?

Check out the list of upcoming events offering "Security Essentials: Network, Endpoint, and Cloud (Security 401)" at http://www.giac.org/registration/gsec

Inverse mapping using disguised TCP resets

Minna Kangasluoma 13 April 2001 version 1.2c

1. Introduction

Today's Internet is full of scans. These scans are sometimes targeted at specific networks, sometimes they are completely random, searching for vulnerable hosts to use in attacks into other, better protected computers, or as slaves in distributed Denial-of-Service attacks. In response, many organizations seek to protect their internal networls using different filtering devices, e.g., firewalls, to limit the traffic allowed into the network. The goal is to deny the attacker knowledge of the machines and configuration behind the firewall, as well as the services which might be vulnerable to an attack; this knowledge would allow him to target specific attacks to those machines or services.

When gathering information preparatory to an attack, the attacker needs first identify specific machines as potential targets. This can be accomplished by many means, including searching public databases (whois, dns) or simply scanning to see which IP addresses are in use. These scans come in many forms. Simple pings can be used to find live hosts. Unfortunately for the attacker, many firewalls nowadays block ping traffic, forcing the attacker to use more sophisticated means. Some of these scans use the inverse scanning method.

2. Inverse scanning

In inverse scanning, the attacker sends a packet addressed to a host located in a network segment protected by a firewall. If no reaction ensues, either because the packet reached the host in question or because it was dropped by a filter, then it can be tentatively concluded that the host may exist. But if the firewall or a router sends back an ICMP 'host unreachable' message, the host in question does not exist. Then the attacker can concentrate on those hosts tentatively presumed to exist, and leave the non-existent hosts in peace. Note that this scan does not tell which hosts exist, only those that do not. See Picture 1 for a representation of the packets exchanged. [1]

TCP resets are often used for these attacks, since TCP resets are ubiquitous in the Internet, and few current IDS systems bother to log them. This will change in the future, as IDS systems evolve. Resets aimed at non-existing hosts will certainly become one of the targets of analysis, for in normal use they are only seen as results of error conditions. Thus an abnormally large amount of resets aimed at non-existing hosts is a clear indication of a scan in progress.

Picture 1.

The problem for the attacker is that in order to gain information from the scan, he has to provide at least one genuine source address where he can study the returned packets. Many scanning took, e.g. nmap[2], provide a way of sending decoy packets[5] from several forged source addresses to confuse the IDS systems. But always there is one genuine address among the rest, making the tracing of the attacker possible if not probable. This paper presents a method for disguising the origin of the scan, as well as the limitations of and countermeasures for such a scan.

3. Disguised TCP resets

In order to disguise the origin of the scan, the attacker may use other machines to echo the scan packets, thus concealing his origins. A machine controlled by the attacker sends a TCP packet to an unwitting accomplice (UA) with a forged source address of the target machine, using a packet-generating tool, e.g. hping[3]. The packet has the ACK bit set, so the accomplice assumes it refers to an existing connection. Since the accomplice knows of no such connection, it will generate a RST packet and send it to the forged source address, i.e., the target machine.

When the packet reaches the firewall or router, it will be either dropped or passed depending on the firewall rules. [The attacker can choose the source port, so he can pick either a port that is likely to pass the firewall, e.g. HTTP port 80, or one that is likely to be dropped, e.g. one of the unpriviledged ports.] In case the target host does not exist, the firewall or the router who knows this will send back an ICMP host unreachable packet. If the attacker positions himself somewhere on the line between the UA and the target network and listens to the traffic, he will see any ICMP packets sent by the firewall/router. If the host exist, no reply is sent.



Picture 2.

The attacker can use any machine in the Internet as the unwitting accomplice. The only requirement is an open TCP port that will reset a non-existent connection. The most useful accomplices are well-known and often used hosts and ports such as web-servers. Even if the scan is detected, the victim may hesitate to contact the owner of a well-known web service to complain about the scan.

To further confuse the issue, the attacker can use several hosts as accomplices, perhaps even a different host for each address. This kind of spread of source IPs makes the scan almost impossible to detect.

4. Limitations

The greatest limitation with this scanning technique is the requirement for a listening host somewhere between the accomplice and the target network. A simple place would be a host situated in the same network segment as the firewall or router protecting the inner network. In this case, the attacker could use any Internet host as an accomplice.

If the listening point is farther from the target, the attacker is limited to those accomplices whose traffic with the target passes through the point he listens at. This limits the available accomplices, but not significantly. Gaining access to such a host might be difficult, though, as the network segments between company LANs tend to be better watched.

Another, more probable way would be to use a single compromised host in a local network and bounce the packets off a machine in the same local net, thus hiding the actual location of the compromised machine. If multiple accomplices were used, the target site might easily conclude that he is seeing residue from a scan aimed at the accomplice site, where his site had been used as a decoy (see [4]).

5. Countermeasures

To counter these kinds of s cans, the defender uses all the same techniques as for countering normal TCP reset scans. A firewall may block ICMP host unreachable packet originating from the inner network. It may also act pre-emptively and simply deny or drop all packets destined to non-existent hosts.

Stateful inspections at the firewall also foil the scan. Since the reset packets do not belong to any existing connection, the firewall simply drops them, thus denying the attacker any useful knowledge of the inner network.

Network Address Translation (NAT) is another excellent way of confusing this scan, depending on the type of the translation. Statically mapping private IP addresses to real ones does not help much, but almost any kind of dynamism is enough to render the results irrelevant. When mapping several private addresses to a single real IP, especially if the port bindings are dynamic, the results of the scan are not very useful.

Detection of this scan is difficult at best. Resets abound in the Internet, and sometimes the packets are only second order effects of decoys sent towards another site. In this case, the attacker has forged the target site's IP addresses for his decoy scans. The analysis of the ports used together with the site acting as a relay may yield some indication on which site was the actual target.

These scans may appear on an ID system, if resets aimed at non-existent hosts are studied over a period of time. If a trend of resets from the same source or a couple of sources show up, and the sources are reasonably well-known hosts with slight chance of compromise, the defender may conclude that the source host is being used as an accomplice. This is probable especially in the cases where the source port does not vary. If different source ports appear only once, the packets more likely are residue from a scan targeted at the source site.

Tracing this kind of scan is theoretically hard, but practically almost trivial. In theory, the listening machine could be located anywhere between the common path of accomplices to the target network. In practice, the location must be either very close to the target network, i.e., just outside the firewall/router, or very close to the accomplices, probably on the same local network. To gain access to a machine elsewhere on the path would be much more difficult, and allow the attacker to do much more damaging things than simple scanning.

A second route is to trace the forged packets arriving at the accomplice hosts. Note that this host may or may not be the same as the listening host. The methods with any likelihood of success would require either cooperative traffic analysis at each network node between target and accomplice [6], or some sort of IP tracing [7-9]. Although theoretical work in this area has been done, the results are not readily suited for tracing single scans. All the tracing methods probabilistic, and require much more traffic for analysis than is produced by a single scan. Cooperative traffic analysis is also more suited for bigger amount of traffic. Also this sort of massive effort is unlikely to happen for a simple scan.

6. Conclusions

Disguised scan uses other, innocent hosts as accomplices to bounce the TCP reset scan to the intended target. The attacker listens somewhere between the accomplice and the

target for the ICMP messages identifying non-existing hosts. The main limitation for this scan is the requirement for this listening host, which the attacker must gain control of first. For this reason, the scan type described above may prove to be more of an academic interest than a practical application.

On the defensive side, this scan does not work against networks properly shielded from straight-forward TCP reset scans. The only bonus for the attacker is the difficulty in detecting this kind of attack, since it is quite easy to disguise it as innocent echoes from scans targeted at the accomplices. On the negative side, gaining access to a host on a route that will allow effective disguising, i.e. multiple distant accomplices, may offset the advantages.

I do not think this kind of scanning will be common at any time, for there are simpler methods to disguise the scan origins. The main benefit of this type of scan as opposed to using decoys and compromised hosts as scan sources is the confusion created. Depending on the accomplices used, the scan can be easily mistaken for either second-order effects of a scan on the accomplice, or a compromised machine at the accomplice's site. Neither conclusion will help track the attacker down.

7. Bibliography

[1] Ofir Arkin. "Network Scanning Techniques". November 1999.
URL: <u>http://www.sys-security.com/archive/papers/Network Scanning Techniques.pdf</u> (28 Mar. 2001)

[2] "NMAP -- The Network Mapper". 10 March 2001. URL: <u>http://www.insecure.org/nmap/</u> (8 Apr. 2001)

[3] "Official Hping homepage". 16 March 2001. URL: <u>http://www.kyuzz.org/antirez/hping.html</u> (8 Apr. 2001)

[4] Green, John, Marchette, David, Northcutt, Stephen and Ralph, Bill. "Analysis Techniques for Detecting Coordinated Attacks and Probes". In *Proceedings of the 1st Work shop on Intrusion Detection and Network Monitoring*, Santa Clara, California, USA, April 1999.

URL:

http://www.usenix.org/publications/library/proceedings/detection99/full_papers/green/green/green/html/ (28 Mar. 2001)

[5] Graham, Robert. "FAQ: Firewalls: What am I seeing?" URL: <u>http://packetstorm.securify.com/papers/firewall/firewall-seen.htm</u> (8 Apr. 2001)

[6] Mansfield, Glenn, Ohta, Kohei, Takei, Y., Kato, N., Nemoto, Y. "Towards trapping wily intruders in the large." In *Proceedings of Second International Workshop on Recent Advances in Intrusion Detection (RAID99)*, West Lafayette, Indiana, USA. September 1999.

URL: <u>http://www.cerias.purdue.edu/raid/proceedings/1999/more/mansfiel.pdf</u> (28 Mar. 2001)

[7] Bellovin, Steven M. "ICMP traceback messages." Internet draft bellovin-trace-00. March 2000.

[8] Doeppner, Thomas W., Klein, Philip N., Koyfman, Andrew. "Using router stamping to identify the source of IP packets." In *Proceedings of the 7th ACM Conference on Computer and Communications Security (CCS 2000)*, pages 184-189, Athens Greece, November 2000.

[9] Savage, Stefan, Wetherall, David, Karlin, Anna, Anderson, Tom. "Practical network support for IP traceback." In *Proceedings of ACM SIGCOMM 2000*, pages 295-306, Stockholm, Sweden, August 2000.

© SANS Institute 2000 - 2002