



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 permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Network Monitoring and Threat Detection In-Depth (Security 503)"
at <http://www.giac.org/registration/gcia>

SANS GCIA Practical

SANS 2001, New Orleans

Tom Chmielarski

© SANS Institute 2000 - 2002, Author retains full rights.

Assignment 1 – Analysis of 5 detects

The first part of the SANS GCIA practical is to analyze five attack detects. The detects analyzed are as follows:

- Detect 1 – DNS Cache Corruption, from the corporate intranet (internal).
- Detect 2 – Web Server Folder Traversal, from a corporate web server (external).
- Detect 3 – SOCKS Proxy probe, from the corporate intranet (internal).
- Detect 4 – RPC Port Probe, from home computer connected to a cable modem.
- Detect 5 – DNS scan, from unknown network as reported to SANS

Detect # 1 – DNS Cache Corruption

Detect Packet Capture

```
15:06:46.364907 0:30:65:66:f6:5c ff:ff:ff:ff:ff:ff 0800 60: MY.NET.195.96.49173 >
MY.NET.195.255.137:
>>> NBT UDP PACKET(137): OPUNKNOWN; REQUEST; UNICAST
TrnID=0x6865
OpCode=13
NmFlags=0x46
Rcode=12
QueryCount=28460
AnswerCount=8311
AuthorityCount=28530
AddressRecCount=27748
Corrupt packet??
(ttl 64, id 1118)
0x0000 4500 0028 045e 0000 4011 dbc2 89a2 c360 E..(^..@.....`
0x0010 89a2 c3ff c015 0089 0014 6436 6865 6c6c .....d6hell
0x0020 6f2c 2077 6f72 6c64 5555 5555 5555 o,.worldUUUUUU

15:06:46.364907 0:30:65:66:f6:5c ff:ff:ff:ff:ff:ff 0800 60: MY.NET 195.96.49173 >
MY.NET.195.255.137:
>>> NBT UDP PACKET(137): OPUNKNOWN; REQUEST; UNICAST
TrnID=0x6865
OpCode=13
NmFlags=0x46
Rcode=12
QueryCount=28460
AnswerCount=8311
AuthorityCount=28530
AddressRecCount=27748
Corrupt packet??
(ttl 64, id 1118)
0x0000 4500 0028 045e 0000 4011 dbc2 89a2 c360 E..(^..@.....`
0x0010 89a2 c3ff c015 0089 0014 6436 6865 6c6c .....d6hell
0x0020 6f2c 2077 6f72 6c64 5555 5555 5555 o,.worldUUUUUU

15:06:46.364907 0:30:65:66:f6:5c ff:ff:ff:ff:ff:ff 0800 60: MY.NET.195.96.49173 >
MY.NET.195.255.137:
>>> NBT UDP PACKET(137): OPUNKNOWN; REQUEST; UNICAST
TrnID=0x6865
OpCode=13
NmFlags=0x46
Rcode=12
QueryCount=28460
AnswerCount=8311
AuthorityCount=28530
AddressRecCount=27748
```

```

Corrupt packet??
(ttl 64, id 1118)
0x0000  4500 0028 045e 0000 4011 dbc2 89a2 c360  E..(.^...@.....`
0x0010  89a2 c3ff c015 0089 0014 6436 6865 6c6c  .....d6hell
0x0020  6f2c 2077 6f72 6c64 5555 5555 5555  o,.worldUUUUUU

```

1. Source Of Attack

The attack originated from a system on the internal network, on the same physical subnet as the target systems.

2. Detect was Generated By

This detect was generated by Black ICE, Enterprise version, and was flagged as a "DNS Cache Corruption". Black ICE is primarily a host-based IDS that combines IDS and firewall functionality. The 'Evidence' file, a tcpdump-style network capture, was retrieved from the sensor and provided the trace above. The output above was created, by processing the Black ICE created file with Windump 2.1 beta.

3. Probability Source IP was Spoofed

Minimal. The originating Ethernet frame header from the attack was compared against the machine at the packet's source IP address and found to be matching.

4. Description of Attack

A malformed name resolution packet was broadcast to the Windows Name Service Port - UDP 137 - to all systems on a single subnet. The packet is illegal in that it contained both a name query and a name response. This is roughly similar to CVE-1999-0288, Denial of service in WINS with malformed data to port 137.

5. Attack Mechanism

This attack is a malformed name query to UDP port 137, the Windows Name Service (WINS / NETBIOS-NS). This service, detailed in RFC 1001 and roughly equivalent to the Domain Name Service (DNS), translates computer names into IP addresses. The WINS service resolves the *NetBios names* traditionally used by Windows systems rather than the *Fully Qualified Domain Names* used by DNS. Logically, any packet related to this service should either contain a request for a computer name, or the answer to such a request – not both. Reviewing the various advisory lists did not turn up anything very similar to this attack. However, this detect is alarming because it was obviously crafted with an abnormal payload.

Each of these packets is corrupt in several respects. To better illustrate this, the following is an example of a normal name request.

Example of Valid NETBIOS-NS Traffic

```

18:11:11.235359 MY.NET.1.137 > MY.NET.1.255.137:
>>> NBT UDP PACKET(137): QUERY; REQUEST; BROADCAST
TrnID=0xC918
OpCode=0
NmFlags=0x11
Rcode=0
QueryCount=1
AnswerCount=0
AuthorityCount=0
AddressRecCount=0

                                QuestionRecords:
Name=IL06EXP01      NameType=0x00 (Workstation)
QuestionType=0x20
QuestionClass=0x1

(ttl 128, id 60492)
0x0000  4500 004e ec4c 0000 8011 f911 81bc a8c8      E..N.L.....
0x0010  81bc a8ff 0089 0089 003a 819a c918 0110      .....:.....
0x0020  0001 0000 0000 0000 2045 4a45 4d44 4144      .....EJEMDAD
0x0030  4745 4646 4946 4144 4144 4243 4143 4143      GEFFIFADBCACAC
0x0040  4143 4143 4143 4141 4100 0020 0001      ACACACAAA.....

```

A few things are obviously different between the valid packet and the crafted one. The real packet is quite a bit longer, as a NetBios name is 32 bits long, padded with spaces (hex encoded as 43 41) where needed. Where the NetBios name should exist, the corrupted packet has three groupings of '55 55'. Valid hex values in the NetBios name field span from 43 41 to 48 4F, with several omissions for invalid characters. Beyond the inadequate size and values of the NetBios name, the corrupt packets have an erroneous NetBios header.

The NetBios header determines the function of the packet. The headers of these packets indicate that the packets contain both a name query and a response to a name query. This is not a valid occurrence. In addition to the formatting errors, the fact that the packet contains the text 'hello, world' strongly suggests that the packet is artificial. This is doubly so as those fields should contain low numerical values - the numerical values that correspond to these ASCII characters are too large to be valid.

As a side note, this attack was determined to be more-or-less a false positive. This attack was a real attack in that an incorrectly formatted packet was repeatedly sent to an entire subnet. This had the side effect of disrupting several printers and at least one Windows 2000 Professional workstation. This attack is also a false positive in the sense that there was no hostile intent. The traffic analyzed in this detect was caused by an in-house engineer learning java-based network programming. For reasons that remain unknown, the software engineer testing Java, on Mac OSX Beta, decided to send UDP broadcast traffic to port 137. The engineer discovered that testing networking code on a production network is a good way to meet several security people, and promised to keep his future development on the test network.

More information on the NetBios Name Service can be found in the Microsoft support article 'How WINS Lookup Works from Windows NT DNS', available at

<http://support.microsoft.com/support/kb/articles/Q173/1/61.ASP>? Instructions for translating computer names in a Netbios packet to an English representation can be found at Microsoft support article Q194203, available at <http://support.microsoft.com/support/kb/articles/Q194/2/03.ASP>. More information about name resolution in a Windows environment is available at <http://msdn.microsoft.com/library/backgrnd/html/Dnsnt4.htm>.

Related CVE issues:

CVE-1999-0288 Denial of service in WINS with malformed data to port 137 (NETBIOS Name Service).

CVE-2000-0673 The NetBIOS Name Server (NBNS) protocol does not perform authentication, which allows remote attackers to cause a denial of service by sending a spoofed Name Conflict or Name Release datagram, aka the "NetBIOS Name Server Protocol Spoofing" vulnerability.

CVE-1999-0810 Denial of service in Samba NETBIOS name service daemon (nmbd).

CVE-1999-0153 Windows 95/NT out of band (OOB) data denial of service through NETBIOS port, aka WinNuke.

6. Correlations

This particular attack has not been otherwise detected as it was a unique false positive. However, a cursory search through the daily SANS detects (<http://www.sans.org/searchsans?p=1&lang=en&mode=all&q=udp+137>) shows numerous scans to this port. Windows itself has a long history of security problems, lending credibility to the thought that anything unusual could cause a problem, intentionally or not. See the aforementioned CVE entries for additional references.

7. Evidence of Active Targeting

This appears to have been purposefully targeted directly to this particular subnet since the target subnet was the attacking host's own segment and no sensors outside of this subnet detected a similar event.

8. Severity

Severity is measured by (Criticality + Lethality) – (Network Countermeasures + Host Countermeasures), on a 1 to 5 point scale.

Criticality equals 3.

The attack was a subnet that mostly contains standard desktops. However, it is possible that very important systems are also located on this subnet.

Lethality equals 2.

It is hard to say the effect this attack could have on a system. One of the systems that received this packet locked up shortly thereafter, possibly as a reaction to this attack. In theory though, the request is not badly malformed, and should be handled normally..

Network Countermeasures equal 1.

Existing network countermeasures are ineffective, as nothing stopped this attack from appearing on the network in general.

Host Countermeasures equal 3.

The variety of hosts on the subnet makes this hard to gauge, but most hosts are reasonably patched, and should not suffer from this attack.

The severity is $(3+2) - (1+3) = 1$.

9. Defensive Recommendation

Check the Vendor (Microsoft), and vulnerability lists such as Bugtraq, for security advisories and patches. Enforce policies requiring network applications be developed on isolated (non-production) networks.

10. Multiple Choice Question

The packets shown almost certainly crafted because:

- a. The source MAC address is all "f"s
- b. The UDP payload is the string "hello, world"
- c. The source port is the same for all three packets
- d. The Windows Name Service runs on TCP port 137, not UDP 137

Correct Answer: **b**

Detect # 2 – Web Server Folder Traversal

Detect IDS Alert and Packet Capture

Time	Target	Attacker	Attack	Information
4/2/2001 9:36	MY.NET.8.71	host213-123-20-190.bntinternet.com	HTTP UTF8 backtick	URL=/scripts/../../winnt/system32/cmd.exe

10:36:56.462875 213.123.20.190.2731 > MY.NET.8.71.80: P 62608328:62608394(66)
ack 370849747 win 8576 (DF)

```
0x0000      4500 006a abd9 4000 6a06 f077 d57b 14be E..j..@.j..w.{..  
0x0010      81bc 0847 0aab 0050 03bb 53c8 161a b7d3 ...G...P..S.....  
0x0020      5018 2180 b8ca 0000 4745 5420 2f73 6372 P.!.....GET./scr  
0x0030      6970 7473 2f2e 2e25 6330 2561 662e 2e2f ipts/..%c0%af../  
0x0040      7769 6e6e 742f 7379 7374 656d 3332 2f63 winnt/system32/c  
0x0050      6d64 2e65 7865 3f2f 632b 6469 7220 4854 md.exe?/c+dir.HT  
0x0060      5450 2f31 2e30 0d0a 0d0a TP/1.0....
```

1. Source Of Attack

The source for this detect is a hosts in a DMZ on my network. The targeted system is a well-known corporate/e-commerce web server.

2. Detect was Generated By

Black ICE IDS/Personal Firewall, Server Version 2.5en/ej. Black ICE is primarily a host-based IDS that combines IDS and firewall functionality. The 'Evidence' file, a tcpdump-style network capture, was retrieved from the sensor and provided the trace above. The output above was created, by processing the Black ICE evidence file with Windump 2.1 beta.

3. Probability Source IP was Spoofed

Minimal. The attack requires the attacker to receive the results. No other attacks, to any of our web servers, have been seen from this or its subnet. If an attacker is going to go to the trouble of using a spoofed IP address, the attacker will probably attack more than one web server – or at least use a more advanced attack. If the attacker were spoofing the source and sending more attacks, those attacks would have had to come from hosts ‘upstream’ of the attacker – most likely from the same subnet. This is not the case. Additionally, no correlating attacks from this IP address have been reported to SANS or the Security Focus incidents list.

4. Description of Attack

The attacker is trying to list the files on the server using the Windows command **cmd /c dir** in the **c:\winnt\system32** directory. If this attack were successful, it would indicate that the targeted server were vulnerable to the Web Server Folder Traversal bug, as described in the Microsoft Security Bulletin MS00-078, found at <http://www.microsoft.com/technet/security/bulletin/MS00-078.asp>.

5. Attack Mechanism

This is a Unicode attack, which takes advantage of how MS IIS processes URLs containing Unicode. Unicode is the UTF-8 standard, which is why Black ICE designates this as an UTF8 attack.

This bug was originally discussed in the Packetstorm forums in October of 2000. Rain Forrest Puppy (RFP), the author of Whisker, was the first to be able to reliably recreate this issue. His description of the “IIS %c1%1c bug” is available on his web page at <http://www.wiretrip.net/rfp/p/doc.asp?id=57>.

If this attack had been successful, the attacker would be able to run quite a few commands on the targeted web server. However, as pointed out by RFP, IIS runs under the IUSR_machine context, limiting the access/permissions that this service/account has.

Also, there is a CVE entry for the web folder transversal attack. **CVE-2000-0884** describes the problem where *"IIS 4.0 and 5.0 allows remote attackers to read documents outside of the web root, and possibly execute arbitrary commands, via malformed URLs that contain UNICODE encoded characters, aka the "Web Server Folder Traversal" vulnerability.* "

This attack is also described in depth in the SANS document "Web Server Folder Traversal" vulnerability (MS00-078), by Steven Shields. This document is available at <http://www.sans.org/infosecFAQ/threats/traversal.htm>.

6. Correlations

This is a fairly common attack, and has been out in the wild for approximately 6 months. As it very simple, and has a large target audience – all un-patched IIS servers, it is a common scan. As previously mentioned, no other attacks from the attacking IP address could be found. However, there are many examples of Unicode attacks.

SANS Daily Detects, March 7th, 2001, by Laurie@edu.

<http://www.sans.org/y2k/030701-1500.htm>

SANS Daily Detects, March 1st, 2001, by Gary Portnoy.

<http://www.sans.org/y2k/030101.htm>

SANS Daily Detects, March 21th, 2001, by Security@auckland

<http://www.sans.org/y2k/032101-1100.htm>

7. Evidence of Active Targeting

The target server was actively targeted, as this does not seem to have been a wide scale scan, nor is it likely to be an accident. No other web server – with a similar IP address, or a similar name, received this attack.

8. Severity

Severity is measured by (Criticality + Lethality) – (Network Countermeasures + Host Countermeasures), on a 1 to 5 point scale.

Criticality is 4.

This system is a very important corporate web server. It also handles e-commerce. If this system were defaced, it would be a significant PR hit, and our company would lose brand image. If customer information were taken... well, it would be a *Bad Thing*TM.

Lethality is 5.

An attacker could do anything from deface the website to modify system files and, with the help of a few commands, gain administrator-level access to the system.

System Countermeasures are 5.

This system is fully patched, and reviewed on a regular basis for security problems.

Network Countermeasures are 4.

This system is running Black ICE, from Network Ice. This product performs IDS and will detect, and block, an attack like this from ever reaching the web server on which it is running. As this software intercepts this activity before it ever reaches the windows TCP stack, I am counting this as a network countermeasure. It is only a 4, rather than a 5, because the web port is open, and it can potentially receive an attack of this nature (HTTP. over port 80).

Severity = (5+4) – (5+4) = 0. This system is in no particular danger.

9. Defensive Recommendation

Vulnerability scans should be run against this system on a regular basis, from a machine that Black Ice trusts, so as to still test the underling OS and applications on the system. There should be (and are) documented procedures in place to ensure that the server receives all applicable patches within a short amount of time. In particular, patch MS00-057, for "File permission canonicalization, will correct this issue.

10. Multiple Choice Question

The danger in this attack is that:

- Windows allows unauthenticated users to execute commands on the IIS web server by default.
- Windows does not parse UTF8 correctly
- The Windows Web server (IIS) has a built in back-door password
- The SSL encryption used by IIS is faulty

Answer: b

Detect # 3 – Socks Proxy Port Probe, Port 1080**Detect Alert Log**

Start Time	End Time	Target IP	Attacker IP	Issue
2/13/2001 2:25	2/13/2001 2:25	my.net.57.121	AsiaOffice.1.5.71	SOCKS port probe
2/18/2001 16:35	2/19/2001 15:22	my.net.57.121	AsiaOffice.60.7.19	SOCKS port probe

2/18/2001 19:12	2/18/2001 19:12	my.net.57.121	AsiaOffice.1.5.71	SOCKS port probe
2/19/2001 8:34	2/19/2001 8:34	my.net.79.71	AsiaOffice.60.7.19	SOCKS port probe
2/19/2001 8:50	2/19/2001 8:51	my.net.147.53	AsiaOffice.60.7.19	SOCKS port probe
2/19/2001 8:56	2/19/2001 8:56	my.net.168.2	AsiaOffice.60.7.19	SOCKS port probe
2/19/2001 8:56	2/19/2001 8:56	my.net.168.71	AsiaOffice.60.7.19	SOCKS port probe
2/19/2001 8:56	2/19/2001 8:56	my.net.168.246	AsiaOffice.60.7.19	SOCKS port probe
2/19/2001 8:58	2/19/2001 8:58	my.net.168.89	AsiaOffice.60.7.19	SOCKS port probe
2/19/2001 9:09	2/19/2001 9:09	my.net.168.86	AsiaOffice.60.7.19	SOCKS port probe
2/19/2001 9:09	2/19/2001 9:09	my.net.168.111	AsiaOffice.60.7.19	SOCKS port probe
2/19/2001 21:49	2/19/2001 21:49	my.net.147.53	AsiaOffice.1.5.71	SOCKS port probe
2/19/2001 22:01	2/19/2001 22:01	my.net.168.89	AsiaOffice.1.5.71	SOCKS port probe
2/19/2001 22:09	2/19/2001 22:09	my.net.168.111	AsiaOffice.1.5.71	SOCKS port probe
2/19/2001 22:10	2/19/2001 22:10	my.net.168.147	AsiaOffice.1.5.71	SOCKS port probe
2/19/2001 22:10	2/19/2001 22:10	my.net.168.133	AsiaOffice.1.5.71	SOCKS port probe
2/19/2001 22:10	2/19/2001 22:10	my.net.168.158	AsiaOffice.1.5.71	SOCKS port probe
2/19/2001 22:11	2/19/2001 22:11	my.net.168.63	AsiaOffice.1.5.71	SOCKS port probe
2/19/2001 22:12	2/19/2001 22:12	my.net.168.57	AsiaOffice.1.5.71	SOCKS port probe
2/19/2001 22:12	2/19/2001 22:12	my.net.168.13	AsiaOffice.1.5.71	SOCKS port probe
2/19/2001 22:12	2/19/2001 22:12	my.net.168.21	AsiaOffice.1.5.71	SOCKS port probe
2/19/2001 22:12	2/19/2001 22:12	my.net.168.93	AsiaOffice.1.5.71	SOCKS port probe
2/19/2001 22:12	2/19/2001 22:12	my.net.168.86	AsiaOffice.1.5.71	SOCKS port probe
2/19/2001 22:12	2/19/2001 22:12	my.net.168.71	AsiaOffice.1.5.71	SOCKS port probe
2/19/2001 23:18	2/19/2001 23:18	my.net.5.4	AsiaOffice.1.5.71	SOCKS port probe
2/25/2001 0:13	2/25/2001 3:06	my.net.57.121	AsiaOffice.60.6.196	SOCKS port probe
2/25/2001 0:57	2/25/2001 3:49	my.net.147.53	AsiaOffice.60.6.196	SOCKS port probe
2/25/2001 1:09	2/25/2001 4:01	my.net.168.89	AsiaOffice.60.6.196	SOCKS port probe
2/25/2001 1:22	2/25/2001 4:13	my.net.168.86	AsiaOffice.60.6.196	SOCKS port probe
3/5/2001 19:01	3/5/2001 19:01	my.net.57.121	AsiaOffice.60.7.54	SOCKS port probe
3/5/2001 19:01	3/5/2001 19:27	my.net.57.45	AsiaOffice.60.7.54	SOCKS port probe
3/5/2001 20:01	3/5/2001 20:01	my.net.168.89	AsiaOffice.60.7.54	SOCKS port probe
3/5/2001 20:04	3/5/2001 20:04	my.net.187.52	AsiaOffice.60.7.54	SOCKS port probe
3/5/2001 20:07	3/5/2001 20:07	my.net.168.130	AsiaOffice.60.7.54	SOCKS port probe
3/5/2001 20:07	3/5/2001 20:07	my.net.168.1	AsiaOffice.60.7.54	SOCKS port probe
3/5/2001 20:07	3/5/2001 20:07	my.net.168.20	AsiaOffice.60.7.54	SOCKS port probe
3/5/2001 20:07	3/5/2001 20:07	my.net.168.28	AsiaOffice.60.7.54	SOCKS port probe
3/5/2001 20:07	3/5/2001 20:07	my.net.168.111	AsiaOffice.60.7.54	SOCKS port probe
3/5/2001 20:07	3/5/2001 20:07	my.net.168.59	AsiaOffice.60.7.54	SOCKS port probe
3/5/2001 20:07	3/5/2001 20:07	my.net.168.37	AsiaOffice.60.7.54	SOCKS port probe
3/5/2001 20:07	3/5/2001 20:07	my.net.168.94	AsiaOffice.60.7.54	SOCKS port probe
3/5/2001 20:07	3/5/2001 20:07	my.net.168.63	AsiaOffice.60.7.54	SOCKS port probe
3/5/2001 20:07	3/5/2001 20:07	my.net.168.95	AsiaOffice.60.7.54	SOCKS port probe
3/5/2001 20:08	3/5/2001 20:08	my.net.168.77	AsiaOffice.60.7.54	SOCKS port probe
3/5/2001 20:08	3/5/2001 20:08	my.net.168.66	AsiaOffice.60.7.54	SOCKS port probe
3/5/2001 20:08	3/5/2001 20:08	my.net.168.65	AsiaOffice.60.7.54	SOCKS port probe
3/5/2001 20:08	3/5/2001 20:08	my.net.168.76	AsiaOffice.60.7.54	SOCKS port probe
3/5/2001 20:08	3/5/2001 20:08	my.net.168.86	AsiaOffice.60.7.54	SOCKS port probe

3/5/2001 20:14	3/5/2001 20:14	my.net.168.237	AsiaOffice.60.7.54	SOCKS port probe
3/5/2001 20:14	3/5/2001 20:14	my.net.168.242	AsiaOffice.60.7.54	SOCKS port probe
3/5/2001 20:14	3/5/2001 20:14	my.net.168.194	AsiaOffice.60.7.54	SOCKS port probe
3/5/2001 20:14	3/5/2001 20:14	my.net.168.202	AsiaOffice.60.7.54	SOCKS port probe
3/5/2001 20:14	3/5/2001 20:14	my.net.168.196	AsiaOffice.60.7.54	SOCKS port probe
3/5/2001 20:14	3/5/2001 20:14	my.net.168.199	AsiaOffice.60.7.54	SOCKS port probe

1. Source Of Attack

The attack originated from multiple systems on the internal network, all from the Asia subnets of our company. The internal network is that of a large enterprise, with excess of 100,000 hosts.

2. Detect was Generated By

This detect was generated by Black ICE, Enterprise version, and was flagged as a "SOCKS Port Probe". The data was imported from the Black ICE SQL database into an Excel spread sheet for easier analysis.

3. Probability Source IP was Spoofed

Minimal. Multiple scans from different internal addresses were found in a short time period. Nothing about the scans is remotely stealthy. As internal intrusion detection is new to our enterprise, it is unlikely that the individuals conducting these scans, if the scans are intentional, would bother to go to the trouble of spoofing IPs to avoid detection, as the perceived chances of getting caught would be dismally low. (Before now!)

4. Description of Attack

The attack is multiple probes to the internal network, all from one geographically clustered branch of the internal network, trying to locate SOCKS proxies on TCP port 1080. Internet access from within the internal network is only available via authenticated proxy access.

There were four individual source addresses in this detect. Each source IP address was confirmed to be a unique host, and not a single host changing IP addresses via DHCP. The targeted networks in these scans overlap between source hosts. This indicates that, supposing the scans are intentional, the people doing the scans are probably working independently of each other, with minimal collaboration. As enterprise IDS sensor placement is still rather rare, and largely confined to a few particular areas of the network, it is almost certain that the detected scans are only a small percentage of the actual scans conducted. The scans started at approximately the same time period, suggesting a common motivator or root cause.

Investigation into the scans confirmed that the scans were a result of users attempting to find additional web proxies to gain unrestricted, and uncharged web access. Shortly before the scans began, our company implemented a policy causing departments to be charged, per megabyte, for employee Internet usage. The charge per megabyte increased substantially for users based in foreign countries. As a result of this, several people in one of the pricier regions of the world attempted to find new ways to access the Internet that did not incur large usage costs for their department. Finding new web proxies would also mean that the web usage of that individual would be anonymous, and unfiltered by the enterprise proxy software, providing additional motivation to find new proxies.

5. Attack Mechanism

The attack consists of a search for new (non-public) proxies within the internal network. The proxies could then be used as an anonymous jump point to the other networks, or to hide the true source of a particular activity. This ‘attack’ does not attempt to take advantage of particular software vulnerability per se, but instead tries to find unsecured or misconfigured network access points to assist unauthorized activity.

6. Correlations

Internally, multiple similar scans were detected in a short time frame, suggesting the search for proxies is not a new, and theoretically complex, endeavor. A quick bit of research shows that proxy probes are common enough that SANS analyst Christopher Misra wrote a document on the prevalence of SOCKS probes. This document is available at <http://www.sans.org/y2k/socks.htm>. The SOCKS port also happens to be on a list of important ports to block in the SANS document **Top Ten Blocking Recommendations Using Cisco ACLs** available at www.sans.org/infosecFAQ/firewall/blocking_cisco.htm. [Searching SANS for port 1080](#) shows searching for Socks proxies is a rather common activity, as reported in the daily detects.

7. Evidence of Active Targeting

This attack does not appear to have been actively targeted. Considering the range of hosts targeted by the probe, and the large sections of the network that are unmonitored by IDS, it is very likely that the majority of the enterprise was scanned for proxies, not any specific subset of the company.

8. Severity

Severity is measured by (Criticality + Lethality) – (Network Countermeasures + Host Countermeasures), on a 1 to 5 point scale.

Criticality of this attack is 4

If successful, this attack would locate all systems acting as a proxy to other networks, the Internet in particular. This is dangerous as it opens a very large, and possibly hidden, doorway into and out of an assumed secure environment.

Lethality of this attack is 1

No direct harm can come from this probe. It will not compromise a system, or harm it in any way.

The Network Countermeasures rating is 4

There are strict procedures in place regulating each physical connection in to, or out of, the network. Each perimeter system has very strict configuration policies and restrictions. However, when an organization is very large and diverse, things will probably slip through, or be done incorrectly somewhere.

Host countermeasures rating is 4

Since this attack does not target a particular vulnerability, but rather a misconfiguration, the host countermeasures relate to the likelihood of a host being victim (an open network proxy) to this probe.

Severity = $(5+4) - (4+4) = \text{one}$. Our network should be fairly safe.

9. Defensive Recommendation

Damage assessment is needed. The attackers have brought up a very interesting question – how many proxy servers are running in our environment? Are they configured correctly? A full scan of the network should be done as soon as possible, simply to scan for proxy servers. Each proxy server that is found should receive a thorough security review to ensure system integrity.

10. Multiple Choice Question

A SOCKS proxy can be used to:

- a. Create a VPN by negotiating IPSEC tunnels
- b. Tunnel a protocol securely through a firewall
- c. Detect port scans
- d. Reconfigure firewall rule sets

Answer: **b.**

Detect # 4 – RPC Port Probes

Detect Alert Log

Date / Time	Source Address	Source Name	Destination	Attack
2001-03-28 22:44:48	24.17.38.227		MY.NET.1.2	RPC TCP port probe
2001-03-28 21:55:37	61.13.119.91	c91.h061013119.is.net.tw	MY.NET.1.2	RPC TCP port probe
2001-03-28 16:26:10	213.51.213.57	cp79460-a.mill1.lb.nl.home.com	MY.NET.1.2	RPC TCP port probe
2001-03-28 12:12:31	211.43.98.4		MY.NET.1.2	RPC TCP port probe
2001-03-28 08:06:29	128.2.244.157	KIWI01.CNBC.CMU.EDU	MY.NET.1.2	RPC TCP port probe
2001-03-28 00:16:10	210.65.21.238		MY.NET.1.2	RPC TCP port probe
2001-03-27 21:21:36	211.22.3.196		MY.NET.1.2	RPC TCP port probe
2001-03-27 03:09:23	24.1.236.195	c68757-a.lvrmr1.sfb.a.home.com	MY.NET.1.2	RPC TCP port probe
2001-03-27 02:20:23	210.62.171.14	sim-ppp14.my.net.tw	MY.NET.1.2	RPC TCP port probe
2001-03-27 01:51:23	205.215.42.19		MY.NET.1.2	RPC TCP port probe
2001-03-26 23:34:28	24.147.73.152	h00a0246b47c7.ne.mediaone.net	MY.NET.1.2	RPC TCP port probe
2001-03-26 19:49:51	211.184.130.2		MY.NET.1.2	RPC TCP port probe
2001-03-26 17:32:05	200.186.216.4		MY.NET.1.2	RPC TCP port probe
2001-03-26 01:56:51	134.140.112.209	queen.simmons.edu	MY.NET.1.2	RPC TCP port probe
2001-03-26 01:34:53	24.68.2.24	24.68.2.24.on.wave.home.com	MY.NET.1.2	RPC TCP port probe
2001-03-26 01:24:26	210.167.238.17	imap.yes.ne.jp	MY.NET.1.2	RPC TCP port probe
2001-03-25 20:06:49	211.34.136.193		MY.NET.1.2	RPC TCP port probe
2001-03-25 13:52:23	194.102.225.156		MY.NET.1.2	RPC TCP port probe
2001-03-25 10:38:07	211.210.2.251		MY.NET.1.2	RPC TCP port probe
2001-03-25 02:54:27	130.251.188.42		MY.NET.1.2	RPC TCP port probe
2001-03-24 03:55:39	63.68.194.14	UNIX.fullport.com	MY.NET.1.2	RPC TCP port probe
2001-03-24 02:23:34	203.133.10.61		MY.NET.1.2	RPC TCP port probe
2001-03-23 02:07:05	192.16.148.126		MY.NET.1.2	RPC TCP port probe
2001-03-22 22:03:25	196.23.186.3		MY.NET.1.2	RPC TCP port probe
2001-03-22 16:00:51	211.185.1.2		MY.NET.1.2	RPC TCP port probe
2001-03-21 16:19:57	203.233.237.131		MY.NET.1.2	RPC TCP port probe
2001-03-21 01:01:29	202.131.132.231		MY.NET.1.2	RPC TCP port probe
2001-03-20 22:29:47	202.234.13.2	ns.lips.co.jp	MY.NET.1.2	RPC TCP port probe
2001-03-20 03:08:57	213.10.1.146	ipd50a0192.speed.planet.nl	MY.NET.1.2	RPC TCP port probe
2001-03-19 23:49:37	24.226.49.149	d226-49-149.home.cgocable.net	MY.NET.1.2	RPC TCP port probe
2001-03-19 15:07:38	209.133.49.198		MY.NET.1.2	RPC TCP port probe
2001-03-19 12:21:48	211.185.118.61		MY.NET.1.2	RPC TCP port probe
2001-03-19 04:25:51	211.15.220.151		MY.NET.1.2	RPC TCP port probe
2001-03-19 01:24:14	211.111.165.136		MY.NET.1.2	RPC TCP port probe
2001-03-19 00:10:20	211.181.92.5		MY.NET.1.2	RPC TCP port probe
2001-03-18 23:09:28	200.213.49.9		MY.NET.1.2	RPC TCP port probe
2001-03-18 08:15:10	211.33.37.94	s211-33-37-94.thrunet.ne.kr	MY.NET.1.2	RPC TCP port probe
2001-03-18 00:42:39	63.228.120.6	backup.vertebraedesign.com	MY.NET.1.2	RPC TCP port probe
2001-03-17 08:46:49	24.27.29.245	cs2729-245.austin.rr.com	MY.NET.1.2	RPC TCP port probe

2001-03-17 03:34:02	24.5.96.22	cc700113-a.vron1.nj.home.com	MY.NET.1.2	RPC TCP port probe
2001-03-17 03:29:08	62.110.84.135		MY.NET.1.2	RPC TCP port probe
2001-03-17 03:15:46	202.107.35.32		MY.NET.1.2	RPC TCP port probe
2001-03-17 01:28:26	210.96.114.61		MY.NET.1.2	RPC TCP port probe
2001-03-16 21:59:40	211.169.220.51		MY.NET.1.2	RPC TCP port probe
2001-03-16 07:08:47	194.225.41.65		MY.NET.1.2	RPC TCP port probe
2001-03-16 03:23:48	24.113.42.86	cr1001896-a.rchmd1.bc.wave.home.com	MY.NET.1.2	RPC TCP port probe
2001-03-15 20:03:26	210.204.3.61		MY.NET.1.2	RPC TCP port probe
2001-03-15 19:46:51	24.27.38.237	cs2738-237.austin.rr.com	MY.NET.1.2	RPC TCP port probe
2001-03-15 15:57:09	211.57.228.53		MY.NET.1.2	RPC TCP port probe
2001-03-15 13:20:03	210.106.81.150		MY.NET.1.2	RPC port probe
2001-03-15 11:04:11	210.232.2.28		MY.NET.1.2	RPC port probe
2001-03-15 04:57:49	211.225.96.68		MY.NET.1.2	RPC port probe
2001-03-15 03:37:57	210.179.202.231		MY.NET.1.2	RPC port probe
2001-03-14 21:17:39	210.204.3.61		MY.NET.1.2	RPC port probe
2001-03-14 17:13:31	211.54.236.83		MY.NET.1.2	RPC port probe
2001-03-14 17:12:22	211.57.228.53		MY.NET.1.2	RPC port probe
2001-03-14 12:19:06	142.59.127.138	s142-59-127-138.ab.hsia.telus.net	MY.NET.1.2	RPC port probe
2001-03-13 21:47:39	202.5.197.126		MY.NET.1.2	RPC port probe
2001-03-13 00:43:19	24.234.60.241	cm241.60.234.24.lvcn.com	MY.NET.1.2	RPC port probe
2001-03-13 00:02:52	212.110.133.67	epiline.so-com.net	MY.NET.1.2	RPC port probe
2001-03-12 22:44:30	202.13.5.140		MY.NET.1.2	RPC port probe
2001-03-12 16:26:22	211.251.177.199		MY.NET.1.2	RPC port probe
2001-03-12 10:47:23	211.251.148.1		MY.NET.1.2	RPC port probe
2001-03-12 05:07:43	24.5.72.120	c340877-a.sttl1n1.wa.home.com	MY.NET.1.2	RPC port probe
2001-03-11 20:13:21	211.251.177.199		MY.NET.1.2	RPC port probe
2001-03-11 14:04:25	210.201.119.173	119-173.kntech.com.tw	MY.NET.1.2	RPC port probe
2001-03-11 14:00:46	211.41.167.35		MY.NET.1.2	RPC port probe
2001-03-11 00:28:02	193.226.2.208		MY.NET.1.2	RPC port probe
2001-03-10 22:12:18	210.227.23.140		MY.NET.1.2	RPC port probe
2001-03-10 16:51:05	63.92.203.19	staff.wtaccess.com	MY.NET.1.2	RPC port probe
2001-03-10 03:16:51	12.37.238.254		MY.NET.1.2	RPC port probe
2001-03-09 17:01:08	211.106.160.232		MY.NET.1.2	RPC port probe
2001-03-08 23:43:21	24.5.72.120	c340877-a.sttl1n1.wa.home.com	MY.NET.1.2	RPC port probe
2001-03-08 19:46:54	200.47.115.51	line115-51.iplanisp.com	MY.NET.1.2	RPC port probe
2001-03-07 21:58:52	210.99.151.35		MY.NET.1.2	RPC port probe
2001-03-07 03:10:34	210.126.141.99		MY.NET.1.2	RPC port probe
2001-03-06 17:06:54	194.29.174.176		MY.NET.1.2	RPC port probe
2001-03-06 06:25:55	210.205.66.161		MY.NET.1.2	RPC port probe
2001-03-06 03:29:37	210.68.82.7	cef7.cef.org.tw	MY.NET.1.2	RPC port probe
2001-03-05 18:51:27	211.38.254.51		MY.NET.1.2	RPC port probe
2001-03-05 17:52:41	62.30.207.35	pc-62-30-207-35-so.blueyonder.co.uk	MY.NET.1.2	RPC port probe
2001-03-05 17:24:03	206.221.244.82	h206-221-244-82.central.grouptelecom.net	MY.NET.1.2	RPC port probe
2001-03-05 16:46:40	202.89.196.221		MY.NET.1.2	RPC port probe

2001-03-05 09:30:35	208.230.168.107	webhost.fycloud.com	MY.NET.1.2	RPC port probe
2001-03-05 02:58:40	24.6.173.12	cx363772-c.chnd1.az.home.com	MY.NET.1.2	RPC port probe
2001-03-05 00:00:45	63.89.102.121	121.102-89-63.adsl.directlink.net	MY.NET.1.2	RPC port probe
2001-03-04 21:48:22	24.30.110.47	we-24-30-110-47.we.mediaone.net	MY.NET.1.2	RPC port probe
2001-03-04 20:55:28	212.68.202.69	mail.osiplus.com	MY.NET.1.2	RPC port probe
2001-03-04 17:17:53	64.184.123.91	064-184-123-091.inaddr.vitts.com	MY.NET.1.2	RPC port probe
2001-03-04 06:45:56	211.57.228.53		MY.NET.1.2	RPC port probe
2001-03-04 05:26:42	211.57.229.2		MY.NET.1.2	RPC port probe
2001-03-03 21:22:35	211.115.216.34		MY.NET.1.2	RPC port probe
2001-03-03 00:27:08	24.234.60.241	cm241.60.234.24.lvcn.com	MY.NET.1.2	RPC port probe
2001-03-02 13:44:39	211.8.31.81		MY.NET.1.2	RPC port probe
2001-03-01 22:27:46	212.242.74.141		MY.NET.1.2	RPC port probe
2001-03-01 22:12:28	211.100.116.254		MY.NET.1.2	RPC port probe
2001-02-28 10:53:37	203.115.24.36		MY.NET.1.2	RPC port probe
2001-02-28 01:52:19	202.101.228.103		MY.NET.1.2	RPC port probe
2001-02-27 20:01:09	63.112.228.50		MY.NET.1.2	RPC port probe
2001-02-27 04:15:08	63.161.151.225	owt-63-161-151-225.owt.com	MY.NET.1.2	RPC port probe
		adsl-63-198-149-36.dsl.isan03.pacbell.net	MY.NET.1.2	RPC port probe
2001-02-26 03:46:46	63.198.149.36		MY.NET.1.2	RPC port probe
2001-02-26 02:24:35	207.137.100.251		MY.NET.1.2	RPC port probe
2001-02-25 02:39:34	210.182.173.157		MY.NET.1.2	RPC port probe
2001-02-25 02:30:24	64.35.57.156	ns.centralmedica.com	MY.NET.1.2	RPC port probe
2001-02-24 14:04:05	211.20.31.206		MY.NET.1.2	RPC port probe
2001-02-24 11:11:03	213.236.131.24	oslo.dhcp-24.wan.no	MY.NET.1.2	RPC port probe
2001-02-24 09:00:22	203.239.104.40		MY.NET.1.2	RPC port probe
2001-02-24 05:25:27	200.207.217.174		MY.NET.1.2	RPC port probe
		adsl-208-191-154-64.dsl.hstntx.swbell.net	MY.NET.1.2	RPC port probe
2001-02-24 01:34:27	208.191.154.64		MY.NET.1.2	RPC port probe
2001-02-24 01:24:51	210.200.114.13	mail.kingnet.net.tw	MY.NET.1.2	RPC port probe
2001-02-23 00:10:26	211.179.51.112		MY.NET.1.2	RPC port probe
2001-02-22 21:09:46	24.218.214.129	taveren.ne.mediaone.net	MY.NET.1.2	RPC port probe
2001-02-22 08:22:35	129.7.129.7	ping.CC.UH.EDU	MY.NET.1.2	RPC port probe
2001-02-22 02:57:07	24.234.60.241	dhcp241.60.lvcn.com	MY.NET.1.2	RPC port probe
2001-02-22 01:34:30	216.135.151.34	user-vc8f5p2.biz.mindspring.com	MY.NET.1.2	RPC port probe
2001-02-22 00:08:54	24.5.224.70	c1069409-d.mntp1.il.home.com	MY.NET.1.2	RPC port probe
2001-02-21 14:58:25	200.202.38.27		MY.NET.1.2	RPC port probe
2001-02-21 04:22:20	63.65.232.3		MY.NET.1.2	RPC port probe
2001-02-21 00:59:47	216.13.12.8		MY.NET.1.2	RPC port probe
2001-02-20 01:08:23	24.240.175.251	24-240-175-251.hsacorp.net	MY.NET.1.2	RPC port probe
2001-02-19 21:11:22	24.168.62.191	24-168-62-191.nyc.rr.com	MY.NET.1.2	RPC port probe
2001-02-19 03:36:27	24.6.61.68	cc41550-a.mtpls1.sc.home.com	MY.NET.1.2	RPC port probe
2001-02-19 02:58:07	200.15.46.70		MY.NET.1.2	RPC port probe
2001-02-18 12:18:10	211.63.91.110		MY.NET.1.2	RPC port probe
2001-02-17 15:33:18	216.217.51.250	ATHM-216-217-xxx-250.home.net	MY.NET.1.2	RPC port probe
		cmldme-cmt1-c3-24-25-176-199.maine.rr.com	MY.NET.1.2	RPC port probe
2001-02-16 20:48:23	24.25.176.199		MY.NET.1.2	RPC port probe
2001-02-16 20:29:20	210.74.122.94		MY.NET.1.2	RPC port probe

2001-02-16 19:13:06	216.13.12.8		MY.NET.1.2	RPC port probe
2001-02-16 02:30:59	216.55.6.170	melodigrafik.com	MY.NET.1.2	RPC port probe
2001-02-16 01:35:12	211.43.176.179		MY.NET.1.2	RPC port probe
2001-02-15 04:53:16	64.76.126.46		MY.NET.1.2	RPC port probe
2001-02-15 00:59:57	211.51.63.46		MY.NET.1.2	RPC port probe

1. Source Of Attack

These logs are from a system on the AT&T @home network. The RPC Port probe was the most common attack detected in approximately one and half months. The @home network is a well-known address range that both receives and sends large numbers of attacks.

2. Detect was Generated By

Black ICE Defender, consumer version 2.1. Logs were taken from Black ICE's Attacklist.csv file, which stores a list of all attacks. The attack was sorted by attack type, and the most common attack was used for this analysis.

3. Probability Source IP was Spoofed

Minimal, but varies. The scans are almost certainly targeted to a large number of systems in the @Home network space. This address space is highly populated with home users (actual individual consumer accounts of the general public). These systems tend to run windows 9x, or improperly administered Linux installations, and could be considered a good place to find 'low hanging fruit' to attack. Almost by definition of such a wide-scoped, obvious scan, the attackers are unlikely to bother using spoofed IP addresses.

4. Description of Attack

This attack is a probe to locate the RPC portmapper service on port 111. The RPC service can provide an attacker with a list of services running on a Unix system, letting the attacker know what vulnerabilities the system might have. The RPC service itself is often the target, as it has been found to have quite a few vulnerabilities. The 'Ramen' Linux worm that has plagued badly configured Linux systems lately takes advantage of vulnerabilities in RPC (rpc.statd), and probably represents quite a few of these RPC probes. There are also several CVE entries that relate to attacks that target RPC.

[CVE-1999-0320](#): SunOS rpc.cmsd allows attackers to obtain root access by overwriting arbitrary files

[CVE-1999-0493](#): rpc.statd allows remote attackers to forward RPC calls to the local operating system via the SM_MON and SM_NOTIFY commands, which in turn could be used to remotely exploit other bugs such as in automountd.

CVE-1999-0003: Execute commands as root via buffer overflow in Tooltalk database server (rpc.ttdbserverd)

CVE-1999-0353: rpc.pcnfsd in HP gives remote root access by changing the permissions on the main printer spool directory.

This attack is also well described by the Network ICE documentation at <http://advice.networkice.com/advice/intrusions/2003102/default.htm> and at <http://advice.networkice.com/advice/intrusions/2003016/default.htm>. Their description, designed with the home user in mind, also has the following to say: *Probes like this result from "script-kiddies", hackers just above the skill level of trained monkeys. They download attack programs (called "scripts") from various sites on the net, then run them against millions of machines.*

5. Attack Mechanism

The attacker – be it an individual or a Ramen worm infected system – scans large subnets in search of systems running the RPC service. Depending on the attacker's intent, the scan might be a normal TCP connection request, or it might use crafted packets in an attempt to avoid intrusion detection systems. These methods would include *stealth* and *null* scans.

6. Correlations

This is a very common, widespread scan – which is the reason it was analyzed. Because of this, it is easy to find evidence of similar scans on the various incident reporting lists. A search of SANS for port 111, for example, shows 480 matches, most of which are reported incidents involving RPC scans. This search is available at <http://www.sans.org/searchsans?p=1&lang=en&mode=all&q=111>. Specific examples of these reports are as follows.

SANS Detects Analyzed March 21st, 2001

<http://www.sans.org/y2k/032101-1500.htm>

SANS Detects Analyzed March 17th, 2001

<http://www.sans.org/y2k/031701-1400.htm>

SANS Detects Analyzed January 17th, 2001

<http://www.sans.org/y2k/011701.htm>

Additionally, the SANS document, A breakdown of SANS Top Ten Threats, by Mary Chaddock list an RPC vulnerability as number 3. This document is available at [http://www.sans.org/infosecFAQ/threats/top_ten.htm-3.%20Remote%20Procedure%20Call%20\(RPC](http://www.sans.org/infosecFAQ/threats/top_ten.htm-3.%20Remote%20Procedure%20Call%20(RPC)

7. Evidence of Active Targeting

None. This scan is targeting large swaths of the @Home network, blindly probing for Unix systems that happen to have the RPC service running and Internet

accessible. If nothing else, the target system is a firewalled Windows 2000 system, which is quite unlikely to have a Unix portmapper service running on TCP port 111 – it was just one of the many scanned.

8. Severity

Severity is measured by (Criticality + Lethality) – (Network Countermeasures + Host Countermeasures), on a 1 to 5 point scale.

Criticality equals 2.

The target system is a generic home PC of no particular strategic value to anyone.

Lethality equals 2.

This scan, in and of itself, has a very low Lethality. At worst, the attacker, if successful, will find that the RPC service is running on the system, and possibly a list of services running. This information alone cannot harm the actual system.

Host Countermeasures are 5.

The target is a Windows host, not running a Unix RPC service. The Windows Host also has all relevant patches applied.

Network Countermeasures equal 5.

Since the target host is also running a proactive firewall/IDS (Black ICE), the OS – and thus the host – would never see the attack, even if the RPC service were running. The firewall portion of Black ICE blocks this port completely. Because Black ICE operates on traffic before it ever gets to the Windows TCP stack, I am considering it a network countermeasure.

Severity is $(2+2) - (5+5)$, equaling -6 . There is no threat to the targeted host.

9. Defensive Recommendation

The user needs to ensure that the firewall/IDS is updated on a regular basis, and unneeded ports are blocked. As this is a home system, not providing any services to the Internet, all ports on the external/Internet interface should be turned off.

10. Multiple Choice Question

Question: The Sun RPC service on port 111...

- A. Is used to share files with Windows systems
- B. Is a secure replacement for SSH
- C. Lists the services running on a Unix system, and each service's port
- D. Confirms that a computer is using Sun Solaris 7.0 as an OS

Answer: c

Detect # 5 – DNS Probes

Alert Log

Snort logs:

Mar 31 04:49:35 takahe snort[64292]: IDS277 - NAMED Iquery Probe:

211.178.63.4:1293 -> 130.216.1.1:53

Mar 31 04:49:36 takahe snort[64292]: IDS278 - SCAN -named Version probe:

211.178.63.4:1293 -> 130.216.1.1:53

Mar 31 04:59:17 takahe snort[64292]: IDS277 - NAMED Iquery Probe:

211.178.63.4:1733 -> 130.216.1.1:53

Mar 31 04:59:17 takahe snort[64292]: IDS278 - SCAN -named Version probe:

211.178.63.4:1733 -> 130.216.1.1:53

Mar 31 05:07:06 takahe snort[64292]: IDS277 - NAMED Iquery Probe:

211.178.63.4:1053 -> 130.216.93.1:53

Mar 31 05:07:06 takahe snort[64292]: IDS278 - SCAN -named Version probe:

211.178.63.4:1053 -> 130.216.93.1:53

Mar 31 05:15:26 takahe snort[64292]: IDS277 - NAMED Iquery Probe:

211.178.63.4:1361 -> 130.216.191.1:53

Mar 31 05:15:26 takahe snort[64292]: IDS278 - SCAN -named Version probe:

211.178.63.4:1361 -> 130.216.191.1:53

Mar 31 05:45:47 takahe snort[64292]: IDS277 - NAMED Iquery Probe:

211.178.63.4:1600 -> 130.216.38.3:53

Mar 31 06:01:30 takahe snort[64292]: IDS277 - NAMED Iquery Probe:

211.178.63.4:2029 -> 130.216.223.3:53

Mar 31 06:01:30 takahe snort[64292]: IDS278 - SCAN -named Version probe:

211.178.63.4:2029 -> 130.216.223.3:53

Mar 31 06:04:18 takahe snort[64292]: IDS277 - NAMED Iquery Probe:

211.178.63.4:1803 -> 130.216.1.4:53

Mar 31 06:04:19 takahe snort[64292]: IDS278 - SCAN -named Version probe:

211.178.63.4:1803 -> 130.216.1.4:53

1. Source Of Attack

This attack occurred on Friday, March 31st, and was reported to SANS on April 4th. <http://www.sans.org/y2k/040401-1200.htm>.

2. Detect was Generated By

This alert was generated by Snort, the “lightweight intrusion detection system”. Snort is available at www.snort.org

3. Probability Source IP was Spoofed

Possible. Per the report to SANS, the activity of this nature ceased as soon as this IP address was blocked. It is unclear if only this IP address was blocked, or the entire C class address range. If only the IP address was blocked, and the attacks stopped, then the address was almost certainly not spoofed, as the attacker would have just switched addresses. As this is a simple scan, using a spoofed IP address would be fairly simple.

4. Description of Attack

The attacker is launching two separate probing attacks against a large network space, apparently against the entire 130.216.255.255 address range. The first of these is a probe to determine if the system supports the DNS Inverse Query (IQUERY) function. The second probe type determines what version of bind the system is running.

5. Attack Mechanism

This attack is a scan to a large number of systems, trying to get a response from the named service. A response to either of these queries will inform the attacker that the target system is a DNS server and will narrow down the list of possible exploits to use on that DNS server. As a result of this scan, an attacker will have a list of DNS servers, even if none of them are vulnerable to the targeted vulnerabilities. Knowing the system is a DNS server; the attacker could also follow up by attempting a *DNS Zone Transfer*. This could give the attacker a lot of information about an organization’s internal network.

The first probe is checking to see if the targeted system supports a DNS Inverse Query. This will tell the attacker that the targeted system is a DNS server and if the DNS server might be vulnerable to a particular vulnerability. Originally reported in April of 1998, this vulnerability exploits a bounds checking error in *named*, the Bind daemon that runs DNS on many Unix platforms.

CVE entry CVE-1999-0009 describes this particular vulnerability as “*Inverse query buffer overflow in BIND 4.9 and BIND 8 Releases.*” A bugtraq report, number 134, is also available at <http://www.securityfocus.com/bid/134.html>. More information on this attack, and signature, is available from whitehats.com at <http://www.whitehats.com/IDS/277>.

The second probe is simply an attempt to determine what version of Bind is running. Knowing this, the attacker can narrow down the list of vulnerabilities a system might have; and thus what exploits to use. Even if the system is running a

fully patched version of Bind, as soon as a new vulnerability is found for that version of Bind the attacker will have a list of newly-vulnerable servers. More information on this request is available from [whitehats.com](http://www.whitehats.com/info/IDS278) at <http://www.whitehats.com/info/IDS278> and from Network Ice at <http://advice.networkice.com/Advice/Intrusions/2000417/default.htm>.

6. Correlations

This scan is coming from an IP address in Korea. This particular address/class C cannot be found in any of the online incidents lists. Either this host/address has not attacked a lot of hosts, or the victims haven't reported/noticed it. Korea tends to have a large number of systems, often compromised hosts that launch attacks/scans. In early march a different Korean address was doing a lot of scanning. A brief discussion on this can be found on SecurityFocus's Incidents list at <http://www.securityfocus.com/archive/75/167892>.

DNS is a very popular target in general. In late January, CERT released an advisory regarding a new vulnerability in Bind version 8. This caused an increase in DNS probes, and script-kiddies looked for easy targets. This, and several other older BIND problems are the stimulus for this type of scan. The new Lion (1i0n) worm, <http://www.sans.org/y2k/lion.htm>, and the Adore worm, <http://www.sans.org/y2k/adore.htm>, spread through DNS vulnerabilities.

There are no available in-depth records on the victim hosts, so long term trending for this type of attack is not possible.

7. Evidence of Active Targeting

Most likely not targeted. It is unknown if all the targeted systems are DNS servers. If these servers are not DNS servers, than the attacker is blindly scanning an address range looking for DNS servers. If the targeted systems are DNS servers, then the attacker has already done some basic reconnaissance on the victim, and this is a highly targeted attack. Most likely, this is a random scan.

8. Severity

Severity is measured by (Criticality + Lethality) – (Network Countermeasures + Host Countermeasures), on a 1 to 5 point scale.

Criticality is rated 3.

The functions of the targets system are not noted. As such, this cannot be adequately judged.

Lethality is rated 2.

These probes only gather information, and do not attempt to harm the target systems in any way. However, they are precursors of an attack, and will enable an attacker to better gauge how to attack the victim.

Host Countermeasures are rated 3.

Data about the targeted hosts is not available. It is assumed that since someone is monitoring attacks, some level of thought has been put into the configuration of the Internet accessible hosts to ensure they are not easily victimized. Additionally, the first probe targets a very old attack, which *should* have been patched years ago.

Network Countermeasures are rated 1.

It is unknown if any device blocked this traffic from reaching these hosts.

Severity is $(3+2)-(3+1)$, the severity is one.

9. Defensive Recommendation

Any DNS servers within this range should be checked to ensure they have all applicable patches. The gateway device should have ACLs applied to ensure that communication to DNS (tcp/udp 53) is only possible to the intended DNS servers.

10. Multiple Choice Question

The IQUERY vulnerability, probed for above, is based on

- a. A blank default administrator password
- b. The ability to 'bounce' traffic off the DNS server to another address
- c. A bounds checking problem
- d. A flaw allowing an attacker to perform a DNS Zone transfer

Answer: **c.**

Assignment 2 – Describe the State of Intrusion Detection Reconnaissance Techniques using Spoofed IP Addresses

Overview

This paper describes methods that can be used to gather information from remote systems using false, or ‘spoofed’, IP addresses. It is important for an intrusion detection analyst to understand the methods used by attackers to take advantage of spoofed IP addresses, in order to detect those methods – or at least consider them when making an analysis.

While the focus of this paper is the use of spoofed IP addresses for reconnaissance, it is important to note that the same methods can also be used to facilitate attacks. The emphasis of this paper is on how spoofed IP addresses can be used to gain information about a targeted system, not on the actual tools, which use these techniques. For consistency and simplicity, I will refer to the reconnaissance method as the ‘attack’, and to the person performing the attack as the ‘attacker’.

Introduction

Systems that communicate via the Internet Protocol (IP) do so by exchanging small messages called packets. These packets use both a source and a destination address to determine where the IP packet came from, and where it is going. By forging an artificial source IP address, an attacker can make an IP packet appear to have come from a completely different source than it actually did. This is known as ‘spoofing’ the IP address. The benefit, to a hacker/cracker of using a spoofed IP address, is that it makes the attack very difficult, if not impossible, to track back to attacker. Because, by definition, the spoofed packet does not return to the attacker’s system, spoofed packets are often overlooked as methods of reconnaissance.

A review of numerous [SANS GCIA practicals](#), showed that many analysts found scans which ‘originated from’ IP addresses that were invalid – such as those addresses registered to IANA - as being unsuited to reconnaissance efforts. The Internet Assigned Numbers Authority (IANA) addresses are specified in RFC 1918, “Address Allocation For Private Internets”. These are addresses that any group can use internally, but will not route externally. For example, in Detect 4 of the GCIA practical done by [Graham Stork](#), he states the following: *“The attack is a stimulus of some kind, it is not a scan as scans are not effective when reserved addresses are used, because the information gained by the scan is not returned to the attacker.”* However, given the right conditions, such an IP address range is perfect for use by some reconnaissance methods, and can even be used to support a full two-way TCP connection that is nearly untraceable.

Spoofed IP Addresses As Background Noise

Perhaps the simplest use of spoofed IP addresses is to create ‘background noise’. An attacker can use spoofed IP addresses to create suspicious traffic that cannot easily be tracked down to the actual attacker. The intent here is not to leverage data from the actual spoofed packets, but to allow the attacker’s real activity, or identity, to be hidden among the false packets.

[Nmap](#), perhaps the most common network scanner at the moment, allows the use of numerous ‘decoy’ addresses. Using the `-D` option in Nmap, such as `nmap -O -D 10.1.1.1, 10.1.1.2, actual.attacker.ip.address, 10.1.1.3 10.2.2.1` will allow an attacker to determine the operating system of the host at 10.2.2.1 while making it appear that the system is being scanned by four simultaneous hosts, only one of which (the 3rd sequentially) is the attacker.

Although this technique is certainly not quiet, it is effective. If ten decoy hosts are used, and all are valid (reachable) hosts, the target will have to investigate all 11 hosts to determine which host was actually the sender. The difficulty in detecting the true attacker increases as larger numbers of decoy addresses are used.

One way to help determine which hosts did not send the packets (and therein which host did) is to search firewall and router logs for incoming error messages from the ten hosts that were spoofed, as those hosts react to the packets sent by the target in response to the stimulus from the attacker. Of course, this depends on the target and the decoys having responded, as well as the packet logging being enabled and accessible to the analyst.

Indirect Reconnaissance of a Target by Observation of the Spoofed Host

A much more stealthy method of reconnaissance is to monitor the host that is being spoofed and detect state changes, if any, caused by the response of the targeted host. The basic process is illustrated below in figure 1. This process is significantly stealthier than using Nmap to do a scan, as the Attacker can gain the desired information without having the Target ever seeing the Attacker’s IP address at all.

Before the attack can actually be performed, a predictable condition that is both observable and can be manipulated must be found. This could be any criterion of a host that changes predictably when a known event occurs. Then a host with that condition must be located. To better describe this method of reconnaissance, we will review the tool *Idlescan* that implements this method. This particular tool is a port scanner that uses predictable IP Identification Numbers (IP IDs) as the observable state. However, the same method can be used against any other observable and modifiable state on a host. The IP ID is the two-byte number (bytes 4 and 5) in an IP packet header that is used as the unique identifier for that packet.

In December of 1999, a person using the alias LiquidK [released a tool called Idlescan](#) based on a [bugtraq post](#) one year earlier by a person using the alias Antirez. Idlescan has


been thoroughly reviewed in the [GCIA practical](#), and several related bugtraq postings, by Teri Bidwell. [An additional paper](#), entitled ‘Spoof Bounce’ was written on the principles behind Idlescan by Kevin Dixon. Idlescan makes use of three tendencies pointed out in Antirez’s bugtraq posting:

- (1) * hosts reply SYN|ACK to SYN if tcp target port is open,
reply RST|ACK if tcp target port is closed.*
- (2) * You can know the number of packets that hosts are sending
using id ip header field.*
- (3) * hosts reply RST to SYN|ACK, reply nothing to RST.*

The significance of this is that due to predictable IP IDs, it is possible to remotely determine if a particular host is sending traffic to a third party. This is accomplished by observing the IP IDs of traffic between an attacker and the sensor host. If the IP IDs in two consecutive packets from the sensor host have incremented by an amount greater than the known increment rate for one packet, then the sensor host has sent an additional packet. For this to work for reconnaissance the sensor host must not receive any additional network traffic. A typical home computer with a DSL, or cable modem, connection would work perfectly as the sensor host in the middle of the night, or perhaps the middle of the workday.

Using another of the described tendencies, it is also possible to predict how a host will react to a port scan. If a host is listening on a port, a probe (SYN) to that port will result in a SYN/ACK. If that port is not listening, the host will respond with a RST/ACK. Furthermore, the final tendency shows that a host does not respond to a RST – or a RST/ACK. Using this knowledge we can step through this attack, and explain why it works. The attack is illustrated in figure 1.

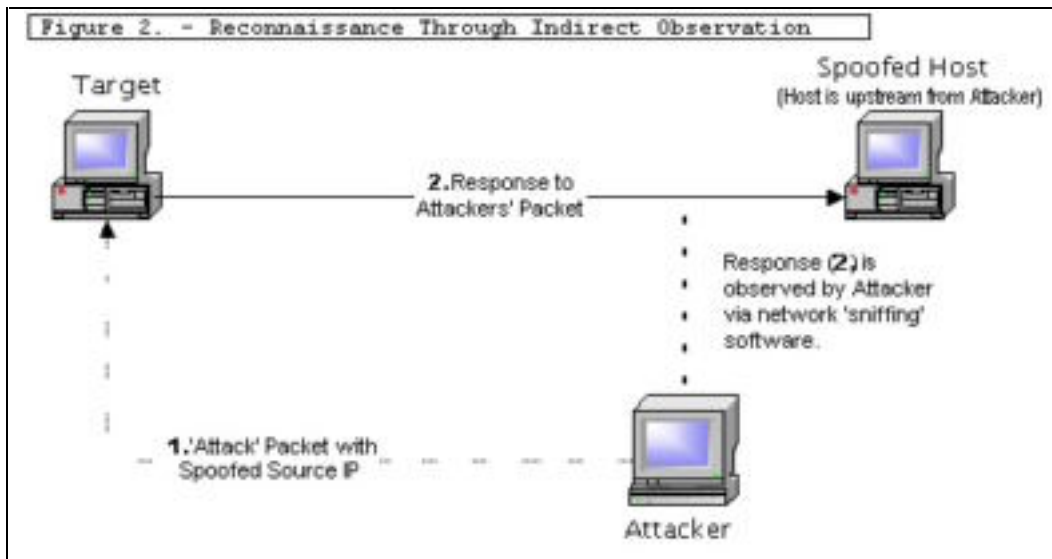
Preparation: Before the attack, the attacker needs to find a qualifying sensor host. In this case the sensor host must (1) have a TCP/IP stack that produces predictable IP IDs, (2) is not sending or receiving other network traffic, and (3) is capable of receiving network traffic from both the Attacker and the Target.



- Stage 1:** The Attacker communicates with the Sensor host to determine the Sensors' current IP ID number. A 'ping' would suffice for this need.
- Stage 2:** The Attacker sends a port probe to the Target, spoofing the IP address of the Sensor host as the source IP address.
- Stage 3:** The Target host responds to the spoofed packet. If the port being probed is open (listening), the Target sends a SYN/ACK to the Sensor. If the probed port is not listening, the Target sends a RST/ACK to the Sensor. If the sensor receives a SYN/ACK from the target the Sensor will try to establish a TCP connection with its own SYN/ACK to the Target. The Sensor will have just sent a packet, and its IP Id will have incremented. On the other hand, if the Sensor receives a RST/ACK from the Target, it takes no action and its IP Id does not increment.
- Stage 4:** The Attacker queries the Sensor again to determine if the IP ID (the known, predictable state) has changed since the initial contact by the Attacker by an amount great enough to indicate a packet has been sent from the Sensor. The Attacker can completely port scan the Target without sending his real IP address to the Target even once!

Reconnaissance Through Indirect Observation

Another sneaky way that attackers can use spoofed IP addresses to hide their trail is through indirect observation of the responses to their attack. In this scenario, detailed in figure 2, the attacker sends spoofed packets to the target, and then observes the responses via a promiscuous network monitor, or 'sniffer'. This attack scheme can be implemented in several different ways, but has a few basic requirements.



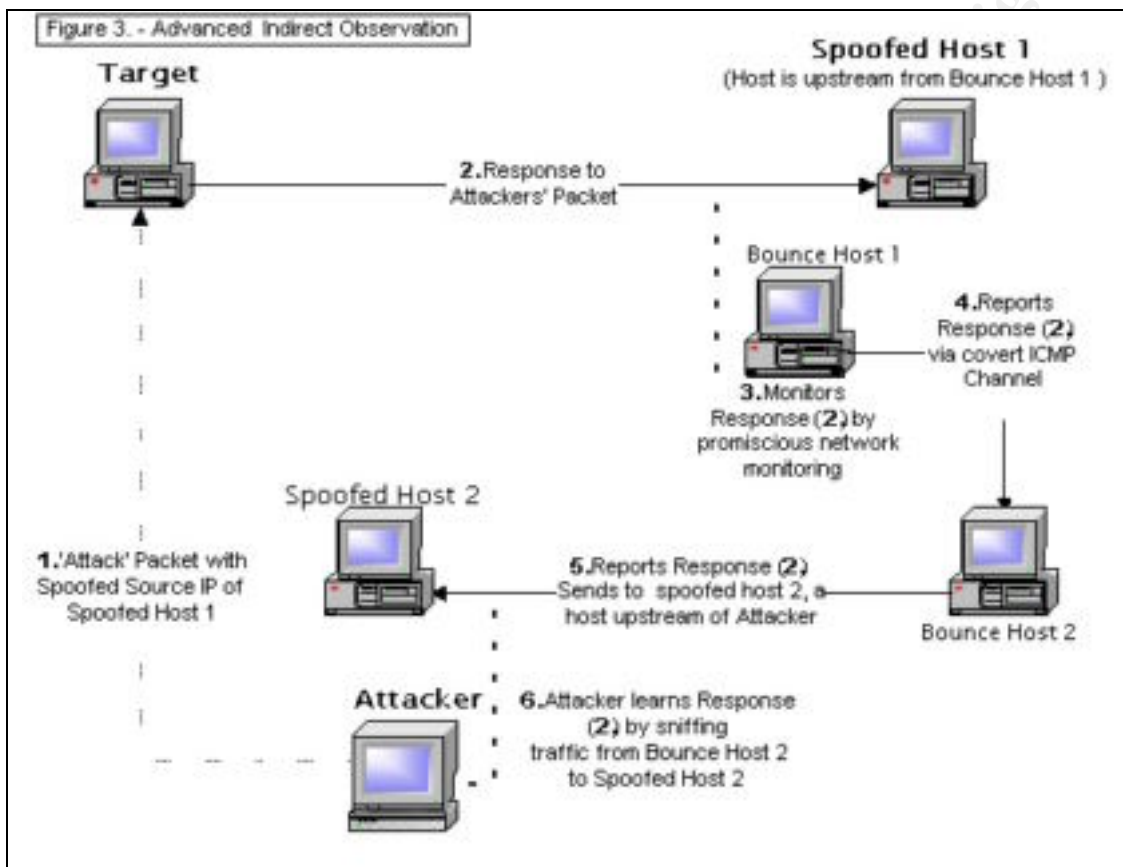
The most significant requirement is that the Spoofed Host must be 'upstream' of the Attacker – meaning you must be on the Target or Spoofed Hosts network segment, or along the path the packet will take. This can be done most easily by spoofing the address of a host on the same local network segment as the attacker. This type of scanning was reported by CERT, in November of 1998, with [Incident Note IN-98-05](#). This report details IMAP scans using compromised hosts on the same subnet as the spoofed host. Similarly, the attacker can achieve the same results by compromising a host on the same subnet as the target.

Depending on the nature of the attack, responses by the Spoofed Host may interfere with the reconnaissance process. For example, beyond just being able to observe the results of a port scan, a talented hacker/cracker could use the sniffed replies to monitor the response to a connection request and artificially build and maintain a connection between the target and the spoofed host. To do this, the Attacker would send a SYN packet to the Target spoofing the source address of an upstream host. The Target would reply to the Spoofed Host with a SYN/ACK. The Attacker would observe the SYN/ACK and respond with an artificially crafted ACK to match the Target's SYN/ACK. In this manner, a full TCP connection could be established between the Target and the Spoofed Host all without any actual participation by the Spoofed Host. However, this type of activity requires that the spoofed address be silent -one that will not send error messages back to the target. A firewall that drops denied packets without returning an ICMP error message would work perfectly for a spoofed host. The IANA reserved addresses are often routable internally within an organization, and could also be used, provided that no responding host is at the address, and no router interferes by responding with an ICMP error such as 'Host Unreachable' or 'time exceeded'.

Advanced Reconnaissance Through Indirect Observation

Very similar to the previous scenario is the use of additional compromised hosts to further hide the true return path of the observed data. At Defcon 8, in July of 2000,

Simple Nomad [discussed](#) an improvement on the attack model mentioned above. He advocated the use of multiple hosts that would observe a response and forward it, or an encoded version of it, through (or past) one or more additional hosts before sending the result to the Attackers' actual host. These systems could communicate directly, via sniffing more spoofed traffic, or by a covert channel. This scheme is illustrated in Figure 3.



The attacker can add as many layers of additional bounce hosts as desired, at the cost of additional complexity and latency. The attacker could also send the attacks using the same type of covert host communication that is used to report the response. The possibilities are unlimited.

Summary

A number of ways have been found to use Spoofed IP addresses to get more information than would be expected. Although the more advanced techniques are complex, and require a high degree of skill to implement, automated tools are always making the difficult tasks easier. For example, many of the various Distributed Denial of Service (DDoS) attacks use a layered architecture that is similar to the advanced spoofing reconnaissance methods mentioned here. Beyond the implementation details mentioned here, it is also important to remember that ARP table and router table manipulation can further obfuscate the true data path. It is important to assume that the sender of a spoofed

packet was able to see the results of any activity, and not all the traffic seen is really going where we think it is.

References

- [1] Stork, Graham. GCIA Certification Practical, December 2000. Detect 4, p10-14.
URL http://www.sans.org/y2k/practical/graham_stork_GCIA.doc (3/21/2001)
- [2] Fyodor. Nmap Man Page.
URL http://www.sans.org/y2k/practical/graham_stork_GCIA.doc (3/21/2001)
- [3] LiquidK. Bugtraq Post: "idlescan (ip.id portscanner)", December 03, 1999.
URL <http://www.securityfocus.com/archive/1/37272> (3/21/2001)
- [4] Antirez. Bugtraq post: "new tcp scan method", December 18, 1998.
URL <http://www.securityfocus.com/archive/1/11581> (3/21/2001)
- [5] Bidwell, Teri. GIAC Network Intrusion Detection GCIA Practical, October 2000.
URL http://www.sans.org/y2k/practical/Teri_Bidwell_GCIA.doc (3/21/2001)
- [6] CERT, IN-98-05: "Probes with Spoofed IP Addresses", November 1998.
URL http://www.cert.org/incident_notes/IN-98-05.html (3/21/2001)
- [7] Simple Nomad, Decfon 8 presentation, July 2000.
Presentation notes online: <http://www.nmrc.org/lab/defcon2000.ppt> (3/21/2001)
- [8] Smetannikov, Max. "Specter of Web Attacks looms again", Interactive Week.
August 2000. <http://www.zdnet.co.uk/news/2000/31/ns-17143.html> (3/21/2001)
- [9] Dixon, Kevin. Spoof Bounce, SANS Reading Room. February 19, 2001
URL <http://www.sans.org/infosecFAQ/intrusion/spoof.htm> (3/22/2001)
- [10] GIAC Certification Program Sample Citations and Quote
URL http://www.sans.org/giactc/GIACTC_citations.htm (3/22/2001)

Assignment 3 – Analyze This

Overview

This assignment is to review the provided intrusion logs for an environment and provide an analysis of the situation within the network from a security standpoint. The environment is that of GIAC Enterprises, a maker of fortune cookies. The GIAC Enterprises network will also be referred to as ‘MY.NET’, as this is its designation within the alert logs.

The logs are approximately 150 megabytes of data collected between the end of November 2000 and the beginning of January 2001, by the Snort Intrusion Detection System. Snort is a free IDS tool available at [www.Snort.org](http://www.snort.org). This paper will provide a quick description of the steps taken to provide the analysis, an in-depth discussion of the attacks detected, and then will summarize the results of the analysis.

Executive Summary

Based upon roughly two months of log data, the GIAC Enterprises network is under attack from a large number of different attackers, both internal and external. There is evidence that some internal machines may have been successfully compromised (“hacked”), and need an in-depth examination. GIAC employees engaging in inappropriate use of GIAC computers could also cause some of the anomalous activity. The fact that some internal hosts have responded to some of the basic network reconnaissance methods from the external addresses shows that perimeter security is insufficient. Externally, a large number of systems, from several different nations, seem to be launching a variety of attacks against the GIAC Enterprises network.

Analysis Process

The logs were contained in four Zip files; one for Snort Alerts, two for Snort detected scans, and one for OOS (Out Of Spec), or abnormal, packets. Each Zip file contained a number of individual files, each comprised of a daily emailed list of activity.

To process the large amount of files in a more efficient manner, all the provided data was loaded into a SQL database. This allows fast, flexible, and granular queries with minimal customized programming, and no dependency on publicly available free tools – such as SnortSnarf (<http://www.silicondefense.com/software/Snortsnarf/>) – that did not seem to scale well to large scale log analysis. Importing the log files into a SQL database was done using the following basic steps.

Step One: Several files of each type were visually inspected to determine data content and formatting. This allowed the design of the SQL database. The MS SQL 7.0 database was created at this time.

- Step Two:** Each grouping of files was consolidated into one large file, using built-in commands on a Windows 2000 Professional system.
- Step Three:** Examination of each sub-file showed an extraneous header at the beginning of each file. The Extraneous data was quickly parsed out of each master file with another Perl script.
- Step Four:** Another simple Perl script was used to take each log entry and parse it into a database-friendly format, based upon the design selected in Step one.
- Step Five:** A format file for the SQL BCP file import utility was created and run against the data file, created in Step four, to import the data into the SQL database. The newly created tables were indexed to speed query time.
- Step Six:** SQL queries were run and used to import data into Excel, providing easy-to-interpret data.

Information on attacking IP addresses was gathered through a variety of sources. The registered owner of many IP address blocks can be found at flumps.org. The NetGeo Internet mapping project, conducted by Caida.org, is a very valuable resource for correlating an IP address with a geographical point. The Internic Registry Whois is a great starting point to resolve DNS names into IP Addresses and owners. The ARIN database is convenient to determine ownership of an IP address range. These data sources are summarized in the following table.

Organization	URL
Flumps.org	http://www.flumps.org/ip
Caida.org	http://netgeo.caida.org/perl/netgeo.cgi
Internic	http://www.internic.net/whois.html
ARIN	http://www.arin.net/whois/

Alerts And Basic Statistics

The Snort alerts offer the dataset that best indicates the overall security condition of GIAC Enterprises. These alerts often indicate simple port-scan activity, but also indicate the more serious events that may indicate advanced/threatening attacks. There were twenty-four unique types of attacks; these are represented in the table below.

All Alerts From Snort Intrusion Logs	Alert Count
Watchlist 000220 IL-ISDNNET-990517	105918
SYN-FIN scan!	51192
DNS udp DoS attack described on unisog	16146
Tiny Fragments – Possible Hostile Activity	5340
connect to 515 from outside	4238
Watchlist 000222 NET-NCFC	2401
WinGate 1080 Attempt	2239
Attempted Sun RPC high port access	2053
Null scan!	826

Queso fingerprint	710
SNMP public access	591
NMAP TCP ping!	558
Russia Dynamo - SANS Flash 28-jul-00	546
SMB Name Wildcard	515
SUNRPC highport access!	204
connect to 515 from inside	159
Broadcast Ping to subnet 70	154
TCP SMTP Source Port traffic	100
Back Orifice	77
External RPC call	59
Probable NMAP fingerprint attempt	8
SITE EXEC - Possible wu-ftpd exploit - GIAC000623	3
Happy 99 Virus	1
STATDX UDP attack	1

The listing of alerts tells us a lot, but not everything. Consideration of the top destination ports will tell us what services are being targeted. In the 194,039 logged alerts, there are 701 unique ports targeted. The top 15 of these represent 90% of the total alerts. The following chart lists the top 15 ports.

Top 15 Destination Port In All Alerts		
Port	Service	Alert Count
6688	Napster	37785
53	DNS	35136
6699	Napster	29329
21	FTP	21619
4876	Unknown traffic from Watchlist 000220	9525
4967	Unknown, traffic from Watchlist 000220	9315
109	POP2 mail protocol	9099
0	NULL, Tiny fragments and syn-fin scans	5494
515	LPD printing daemon	4397
1525	Oracle	4191
6346	Gnutella	2351
32771	Sun RPC	2257
1080	SOCKS Proxy	2240
25	SMTP	1755
443	HTTPS / SSL – Secure Web Traffic	1673

There are a variety of attacks being used against GIAC Enterprises. The threat of an attack is relative to where the attack is coming from. A quick review of the origin of these attacks is necessary to put the attack analysis in the proper context. The following chart shows us that the majority of the detected attacks are external, but a disturbing number also originate internally.

Total Alert Sources		
Attack Source	Attack Count	Percent
Internal Source	1366	0.70%
External Source	192673	99.30%
Total Attacks	194039	100.00%


One caveat should be discussed – the ‘internal’ attacks, might not be internal, they could be spoofed. However, without an understanding of Snort sensor placement, it is not possible to determine if this could be the case. Without that data, it is assumed that proper border filtering is done to block/deny inbound packets (those from external networks) with internal source addresses.

The Top Ten Alerts Defined

To better understand the threat posed by these attacks, it is critical to understand each attack. Accordingly, a description of each of the top ten attacks will be presented, as they account for 98% of the Snort alerts. Each of these ten attacks will be discussed within the context of the activity logged within GIAC Enterprises. The final 14 alerts will then be defined in general terms.

Alert 1. Watchlist 000220 IL-ISDNNET-990517 (105,918 Alerts)

This traffic is flagged because it comes from Israel (212.179.x.x), generally considered a ‘hot spot’ of malicious activity on the Internet. The Snort rule for this alert is not available for review, so it is uncertain if traffic going to Israel would also trigger an alert. Based on the pattern of the traffic causing the alerts, it is assumed that traffic going to Israel would not trigger an alert, and that such traffic is most likely occurring. If GIAC Enterprises is located in Israel, or has many customers there, than this level of activity may be normal. Otherwise, there is a reason to be concerned. The following table lists the both the top 10 source IP addresses, and, when available, the registered owner of these subnets.



Top 10 Sources For Watchlist 000220 (Israel)		
Count	Source Address	Registered Owner
48786	212.179.79.2	Elonet - www.elonet.com - web development
39015	212.179.27.111	BIRANIT-GOREN
4563	212.179.95.5	Cable-Modem-Experiment
2353	212.179.77.20	KIBOTZ-SAAR - www.saar.com - hiking equipment
1517	212.179.44.105	GIVAT-BRENER - http://www.gbrener.org.il/
1387	212.179.42.102	Pablikum
1221	212.179.38.135	Infomall - see http://www.sans.org/y2k/100500.htm
1054	212.179.58.12	Spinoff-II
1002	212.179.45.241	Unknown
926	212.179.56.5	Unknown

So now we see where most (96%) of the Israeli traffic is coming from, but what is it? These alerts are based on the country of origin, not anything that specifically indicates an attack. Review of the destination ports can help determine what activity is actually occurring – and if we should be concerned. The following table lists the top ten destination ports.

Watchlist 000220 Top 10 Destination Ports	
Port	Alert Count
6688	37765
6699	29194
4876	9525
4967	9315
1525	4191
6346	1914
2209	1517
443	1388
4078	1221
41033	1062

The bulk of the traffic is going to source ports 6688 and 6699. These ports are most widely used for Napster, a music sharing utility, and for Gnutella, a file sharing utility. This may be more of a policy issue than an attack. However, both of these programs can be used to transfer any type of data, and can present a security risk.


It is important to note that source IP 212.179.38.135 belongs to Infomall. A different Infomall source IP was part of a distributed IMAP/POP scan reported to SANS in

October of 2000. This report is available at <http://www.sans.org/y2k/100500.htm>. Analysis of this traffic shows that all traffic from 212.179.38.135, port 1108, going to MY.NET.98.114, port 4078. This activity occurs on December 31, 2000 between 3:12am and 3:34 am, logged time. Neither of these ports is well known for a specific service or trojan. No activity was logged from MY.NET.98.114, so further investigation of this activity and host is needed.

Alert 2. SYN-FIN scan! (51,192 Alerts)

This traffic is caused when a TCP packet is seen with both the SYN and FIN flags set. These flags are used to start and end, respectively, a TCP Session. These two flags should never occur together under normal conditions. These packets are often used for network scans, as older intrusion detection systems would not log them. More information on this attack can be found on the Whitehats website, at <http://www.whitehats.com/IDS/198>.

None of these attacks originated within GIAC Enterprises – a good sign! Externally, the source IP addresses are very geographically distributed. The following chart shows the distribution of the top ten source IP addresses. These addresses are responsible for 99% of the SYN –FIN Scans.



Interestingly, but unsurprisingly, the address 194.234.48.26 was reported to Sans on Jan 6, 2001 as the source of an FTP scan. This report is available at <http://www.sans.org/y2k/011601-1530.htm>. It is also worthwhile to consider the target of all these port scans, to see if any particular system is being scanned more often than the others. As the following table shows, only MY.NET.253.112 is receiving more scans than the others – three times as many! This host will be reviewed in a later section of this security analysis.

Syn-Fin Alerts Top 10 Target Hosts	
Destination Address	Alerts
MY.NET.253.112	19
MY.NET.21.15	8
MY.NET.11.212	7
MY.NET.5.125	7
MY.NET.11.230	6
MY.NET.18.143	6
MY.NET.21.208	6
MY.NET.11.177	6
MY.NET.21.229	6
MY.NET.17.196	6

We now know what computers were scanning GIAC Enterprises, and what systems were being scanned... But what were the scanners looking for?

Alert Count	Top 5 SYN-FIN Scan Destination Ports
21604	21 - FTP
18863	53 - DNS
9099	109 - POP2
1580	9055 - unknown
18	259 - Efficient Short Remote Operations (RFC-2188)

FTP is the most common service targeted by the SYN-FIN scans. There have traditionally been a large number of vulnerabilities regarding FTP, making it a common target. Many CVE entries detail these problems, including FTP port bouncing, [CVE-1999-0017](#), and several root vulnerabilities in [CVE-1999-0080](#), [CVE-1999-0219](#), and [CVE-1999-0368](#).

DNS is the next most popular port. It too has had a long history of vulnerabilities, with a new BIND DNS vulnerability very recently. A list of DNS related CVE entries can be found at <http://cve.mitre.org/cgi-bin/cvekey.cgi?keyword=dns>.

Port 109 is the home of POP2, a much older, and now seldom used, protocol. However, some mail servers install this service by default, so it is possible that this scan is trying to find default mail server installations. A wide POP2 scan was reported in early November of 2000. See <http://archives.neohapsis.com/archives/incidents/2000-11/0049.html> for additional details. A review of the attack detects posted to SANS shows that port 109 was a fairly common target for scans from mid to early-late 2000. These detects can be found at <http://www.sans.org/searchsans/perfect/search/search.pl?lang=en&mode=all&q=109+detects+analyzed+>. Mail servers should be checked to ensure that this service is not running, and it should be blocked at the egress router.

The scans to port 259 are targeting the Efficient Short Range Operations Protocol, specified in RFC 2188. This is a service similar to RPC, designed to be efficient for wireless applications. This RFC is available at <ftp://ftp.isi.edu/in-notes/rfc2188.txt>. One SANS detect analysis, available at <http://www.sans.org/y2k/020701.htm>, suggests that a search for this protocol is part of an OS fingerprinting technique.


Alert 3. DNS udp DoS attack described on unisog (16146 Alerts)

This alert is generated by UDP DNS rests that appear to originate from 209.67.50.203. This is a Domain Name System (DNS) server run by Exodus Communications Inc, in California. The explanation of this alert can be found in a discussion at <http://www.theorygroup.com/Archive/Unisog/2001/msg00027.html>, and at <http://www.theorygroup.com/Archive/Unisog/2001/msg00028.html>. In January of 2001, a large number of Denial of Service (DoS) attacks were seen 'from' 209.67.50.203 to several DNS servers. This attacks tries to overwhelm a server by sending too many requests for the server to respond to. In this case, the system 209.67.50.203 is the actual victim/target of the DoS. An attacker is sending spoofed packets to multiple systems, including those within MY.NET, to cause a large number of DNS responses to 209.67.50.203. This is not a serious threat to GIAC Enterprises specifically, just an attempt to use GIAC Enterprises' resources to attack an outside party.

Alert 4. Tiny Fragments - Possible Hostile Activity (5340 Alerts)

This alert is created by the Snort pre-processor mini-frag, and is generated when an unusually small packet is detected. Attackers sometimes try to hide their activities from Intrusion Detection systems by breaking each IP packet of the attack into very small sections, called fragments. The goal of fragmentation is to evade intrusion detection systems that do not check fragmented packets, or that reassemble fragmented packets differently than the target system does. More information on how fragmentation is used to bypass intrusion detection can be found in the SANS paper http://www.sans.org/infosecFAQ/intrusion/net_id.htm. More information on the mini-frag preprocessor can be found at <http://www.dpo.uab.edu/~andrewb/Snort/preprocessors.html>.

The only information logged about most of these events is source and destination addresses, port numbers are often not included. As a result, the true threat of these events cannot be adequately determined. An analysis of who the attackers are is important to understand these events.



It is immediately clear that although the top single source IP address is from an @Home user, the bulk of these alerts are generated from hosts within China. As GIAC Enterprises specializes in selling fortune cookies, this activity may be coming from business partners who are using software that does unusual fragmentation, or through network connections that only accept very small packets. In this scenario, the traffic from China would be false positives. The China address range does not seem well documented, so the exact origins for many of these addresses cannot be determined. However, several of these addresses seem to be general access ISPs and educational facilities, neither of which is very likely to be a business partner. It is recommended that these addresses be reviewed for valid business purposes. To put this traffic in the correct perspective, a review of the target systems is also needed. The following table shows the top 10 destinations for the Tiny fragment alerts.

Tiny Fragments	
Top 10 Destinations	Alert Count
MY.NET.1.8	3148
MY.NET.1.10	1264
MY.NET.217.162	727
MY.NET.60.11	168
MY.NET.1.9	8
208.162.62.208	7
MY.NET.202.18	6
MY.NET.100.230	5
MY.NET.98.123	2
MY.NET.215.106	2

Two things are most notable about this data; the system MY.NET.1.8 is receiving the majority of this traffic, and that an external address, 208.162.62.208, is also receiving quite a few of these abnormal packets. Analysis of the host MY.NET.1.8 can be found later in this document, in the Hosts section. This host is receiving a large amount of attention, and needs to be evaluated for proper configuration and/or checked for a compromise. The traffic to 208.162.62.208 is from MY.NET.219.122 on port 4000, a port that is commonly known for ICQ 'chat' software. Analysis of the hosts MY.NET.219.122 and 208.162.62.208 can be found in the Hosts section of this paper.

Alert 5. Connect to 515 from outside (4238 Alerts)

These alerts are generated in response to computers outside of the network searching for internal UNIX/Linux/BSD systems with the LPR (printer) service running. A vulnerability for this service was announced in November of 2000. This was followed by an increase in the number of scans for this service. More information can be found in the SANS report found at <http://www.sans.org/newlook/alerts/port515.htm>. A CVE entry for a related attack is CVE-1999-0032, and can be found at <http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-1999-0032>. It is recommended that GIAC enterprises review network requirements, and consider blocking external access to this port, as external use of internal printers probably is not necessary. The following chart shows the top ten source IP addresses for this alert.

Connect to 515	
Source IP Address	Scan Count
141.211.176.99	2236
216.119.15.88	1273
209.217.166.69	713
192.118.36.9	7
62.46.70.175	4
172.161.186.125	1
207.173.179.18	1
128.61.36.117	1
24.160.143.196	1
24.4.196.167	1

The majority of the attacks of this type came from 141.211.176.99, on December 15th, 2000, between 12:24am and 12:55 am. This IP address is registered to the University of Michigan. The computer at 216.119.15.88 is also responsible for a large percentage of these alerts. It scanned GIAC Enterprises on December 20th, 2000, between 11:13 pm and 11:39 pm. This IP address is registered to JPS.net, an Earthlink company in California. No correlating scans, for either of these two sources, can be found on any of the common incident reporting lists.

Alert 6. Watchlist 000222 NET-NCFC (2401 Alerts)

These alerts are generated in response to any traffic originating from the 159.226.255.255 address range. These addresses belong to The Computer Network Center Chinese Academy of Sciences, part of the Institute of Computing Technology Chinese Academy of Sciences, in Beijing, China. As with Israel, this address range is a common source for hostile activity. These alerts were generated by traffic from 31 unique source addresses within the Academy, and chronologically spread out with no single time frame. These attacks targeted 19 unique hosts within MY.NET, with 83% of the attacks directed at five particular hosts. This is shown in the following table.

Watchlist 000222 Alerts	
Top 5 Destinations	Alerts
MY.NET.100.230	789
MY.NET.6.7	540
MY.NET.253.41	278
MY.NET.5.29	275
MY.NET.253.42	112
All Other Destinations	407

Consideration of the destination ports is important to properly determine the nature of the traffic related to these alerts. The following table shows the top 10 destination ports.

Watchlist 000222 Alerts	
Top 10 Dest. Ports	Alerts
25	1486
143	505
443	275
113	81
21	10
51221	5
53677	4
49574	4
49255	2
7187	2

Destination port 25 is the target in the majority (62%) of this traffic. This port is primarily known for the Simple Mail Transfer Protocol (SMTP), but is also known for quite a few Trojans. Reports of similar scans have been previously reported to SANS – see <http://www.sans.org/y2k/043000.htm> and <http://www.sans.org/y2k/052800-1100.htm> for details.


Port 443, best known for HTTPS – used for secure web transactions – is also a common destination port. If the Academy purchases fortune cookies from GIAC Enterprises, a large amount of this traffic might be legitimate email and secure web traffic. For example, most of the traffic to the top destination, MY.NET.100.230, consists of repeated connections to 25 (SMTP), and 113 (IDENTD) from a few (6) particular hosts. This traffic could easily be considered as legitimate. However, much of remaining traffic

originates from a few source IP addresses to multiple destination addresses, fitting the profile of scanning rather than of legitimate activity.

Another odd source of activity is the traffic from 159.226.115.1 to MY.NET.253.41. 159.226.115.1 sent 152 packets/alerts to MY.NET.253.41, port 25 (SMTP) over three days, with the majority (95%) of these being between 9:57 pm on December 6th, 2000, and 1:53 am on December 7th. All 62 packets on the 6th came from source port 32866, suddenly changing to port 32905 for all traffic on the 7th. Then there is no further communication until January 4th. The source port typically changes each time a new connection is established. This activity may be quite normal, or it could be questionable. In either case, it does not fit the profile of a scan. Customer analysis needs to be done, to determine if the Academy is a current, or prospective customer, and the address range may need to be blocked if not.

Alert 7. WinGate 1080 Attempt (2239 Alerts)

These alerts indicate a reconnaissance effort to find SOCKS (wingate) proxies. More information can be found at Max Visions' Whitehats.com Snort Rule database, located at <http://www.whitehats.com/info/IDS175>. Socks proxies can be used to redirect traffic, allowing, for example, attacks to be routed through the proxy, making the proxy computer appear to be the real source of the attack. None of these attacks originate internally. The external sources are shown in the following chart.



The preceding table shows that the sources of these scans are quite diverse, with no single source being very dominant. This is either a very coordinated distributed scan, or a common attack. It is worth noting that two separate IP addresses are registered to Webmaster Incorporated, so activity from this source should be tracked in the future.

Alert 8. Attempted Sun RPC high port access (2050 Alerts)

This is a fairly common probe, wherein the attacker is searching for the Sun RPC Portmapper service on port 32771. Connection to this service will allow the attacker to take advantage of any exploits for this service, as well as providing the attacker with a list of all services running on the system. More data on this alert, and signature, can be found at <http://www.whitehats.com/info/IDS429>. There is also a CVE candidate (CAN-1999-0632) entry under review for this issue. This CVE candidate is available at <http://cve.mitre.org/cgi-bin/cvename.cgi?name=CAN-1999-0632>.

The majority (97%) of these alerts come from systems on the 205.188.153.255 subnet. This subnet is registered to America Online. Additionally, almost all of the traffic pertaining to these alerts comes from source port 4000, even between different hosts. This commonality suggests a few possibilities. It is possible that there is a misconfiguration with some software, or hardware, at America Online that is accidentally sending packets to GIAC Enterprises. This theory is supported by the fact that some of the source IP addresses are repeatedly querying the same destination systems. For example, 205.188.153.108 contacts MY.NET.222.218 several hundred times at narrow, but erratic, intervals for several days. This is not typical behavior for a scan, and a sample is shown in the following table.

Excerpt from Attempted Sun RPC high port access Alerts					
Date	Time	Source IP Address	Port	Destination IP	Port
12/9/2000	07:07:54.167	205.188.153.108	4000	MY.NET.222.218	32771
12/9/2000	07:12:39.200	205.188.153.108	4000	MY.NET.222.218	32771
12/9/2000	07:14:53.860	205.188.153.108	4000	MY.NET.222.218	32771
12/9/2000	07:22:53.530	205.188.153.108	4000	MY.NET.222.218	32771
12/9/2000	07:24:53.447	205.188.153.108	4000	MY.NET.222.218	32771
12/9/2000	07:25:53.420	205.188.153.108	4000	MY.NET.222.218	32771
12/9/2000	07:28:53.280	205.188.153.108	4000	MY.NET.222.218	32771

A different theory would be that a new script-kiddie tool has been released, so the 31337 hackers on AOL are scanning for easy targets. If this were the case, however, correlating data should be found on various incidents lists, such as SANS. This does not seem to be the case.

The final obvious answer is that a system on the AOL subnet has been compromised by a sneaky hacker/cracker. The hacker could then launch a full scan of MY.NET, spoofing IP addresses of other systems on the local network. The attacker would then monitor (“sniff”) the network for replies to the scan. An example of this type of scanning is

detailed in a CERT notice, IN-98-05, available at http://www.cert.org/incident_notes/IN-98-05.html. This theory is supported by the overall pattern of the alerts. All 13 of the source IP addresses are sequential, which is a bit unusual. The traffic pattern seems to be such that a given source host will query a given target several times – as would be expected with a non-reliable observation system and then a different source IP address, will query a different target system, using the same port. An attacker that is trying to be somewhat subtle, and nearly untraceable, could conceivably exhibit this traffic pattern.

Alert 9. Null scan! (826 Alerts)

These alerts are generated when a packet with a sequence number of zero, and no flags set, is seen. Similar to SYN-FIN packets, these packets are illegally formatted, and should not occur under normal circumstances. The object of this traffic is to gather information about foreign systems while trying to evade logging and intrusion detection systems. This is a fairly simple, and common, scan. This is supported by the following table, listing the top 10 source IP addresses, which account for only 12% of the NULL scans detected.


Null Scan!	
Top 10 Source IP Addresses	Alert Count
63.253.110.142 - US - McLeod USA	19
24.112.150.20 - Canada - @Home ISP	16
63.253.98.172 - US - McLeod USA	11
63.253.98.171 - US - McLeod USA	10
63.252.95.21 - US - McLeod USA	9
63.253.110.157 - US - McLeod USA	9
24.113.198.51 - Canada - @Home ISP	8
63.252.96.36 - US - McLeod USA	8
63.252.92.159 - US - McLeod USA	7
63.253.110.147 - US - McLeod USA	7

It can be seen that eight of the top ten sources are all from within the McLeod USA address range. Note, this range is registered to Splitrock, a subsidiary of McLeod USA. This address space accounts for an impressive 475 – 57% - of the 826 logged NULL Scans. Even more impressive is that 258 unique, but clustered, IP addresses within the McLeod address range are responsible for these 475 scans, with most of these systems scanning only one or two of GIAC Enterprises' systems. It is quite likely that an attacker (or group) has compromised a few systems on the McLeod network. The attacker is then performing the scan by spoofing other systems on the same physical network in order to observe the responses. This would allow an attacker to scan a large number of systems using a large number of systems, and possibly not triggering any intrusion detection systems. This is the same methodology suggested as the third possibility for the cause of the Attempted RPC Portmapper Connections above. Please see CERT notice IN-08-05 for more information. Because each scanning computer is only trying to connect to one or two targets, most intrusion detection systems will not notice this as an attack – it is only the malformed packet signature that gives this scan away. Another possibility is that

McLeod USA has a malfunctioning piece of equipment that is malfunctioning and sending out malformed packets.

Alert 10. Queso fingerprint (710 Alerts)

These alerts are caused by Queso, the OS 'Fingerprinting' tool. Queso is a predecessor to the common utility Nmap, and determines the Operating System (OS) of the computer it is used against. An attacker uses Queso as a reconnaissance tool to determine the OS of the target machine(s), as a way to determine what vulnerabilities may exist on that system. This saves the attacker from the embarrassment of running windows exploits on a Unix machine, or visa versa This tool is available at www.apostols.org/projectz/queso/. The use of this tool means that someone wants to know more about the GIAC computers so that they are better prepared to break in. On the bright side, none of the Queso scans originated internally.



Eighty Percent of the attacks were committed by top ten source IP addresses, 6 of which were from the address space of TU-Dresden in Germany, which is responsible for 227, or 31%, of the Queso scans. The TU-Dresden hosts are involved in no alerts other than these Queso scans. No other correlating attacks have been found from these hosts on any of the common incident lists.

Queso Alerts	
Top 10 Destinations	Alert Count
MY.NET.219.114	204
MY.NET.201.130	127
MY.NET.201.62	51
MY.NET.204.38	39
MY.NET.223.226	38
MY.NET.224.242	37
MY.NET.201.66	28

MY.NET.202.46	20
MY.NET.53.108	16
MY.NET.60.8	10

The list of the top ten destination IP addresses shows that a disproportionate number of Queso scans were targeted at two specific hosts. Review of the top host, MY.NET.219.114, can be found in the Hosts section of this document. This host should be reviewed for proper security configurations, as it particularly popular.

The Other Alerts

This section describes the characteristics of the remaining 14 alerts, the bottom 2% statistically, that were found within the Snort logs. These attacks are only defined in general terms.

Alert 11. SNMP public access

This alert indicates that the 'Public' community was queried on a Simple Network Management Protocol (SNMP) enabled device. A full description can be found at <http://advice.networkice.com/advice/reference/networking/snmp/default.htm>. Discussion of the security implications for this alert is covered in the Internal Alerts section of this document.

Alert 12. NMAP TCP ping!

This alert is caused when the port scanning, and operating system fingerprinting, tool Nmap is used to map a network. This tool will give an attacker a significant amount of information about a target host or network. The NMAP TCP Ping! Alert is specifically generated when Nmap is used to query hosts to see if they are up and responding ('alive'). Nmap is described in depth at <http://advice.networkice.com/advice/intrusions/2001526/default.htm>. The security implications of these scans for GIAC Enterprises are discussed in the Internal Alerts section of this document.

Alert 13. Russia Dynamo - SANS Flash 28-jul-00

The cause of this alert is a Windows Trojan, discovered in January of 2000, that sends information to a computer in Russia, in the 194.87.6.255 range. More information on this alert can be found at <http://www.sans.org/y2k/072818.htm> and <http://archives.neohapsis.com/archives/sans/2000/0068.html>. The security implications of these scans for GIAC Enterprises are discussed in the Internal Alerts section of this document.

Alert 14. SMB Name Wildcard

This alert is generated when a query is sent to a windows system, or a *nix system running Samba, to enumerate the available shares on that system. More information on this activity is available from Network Ice as a 'MS Share Dump' at <http://advice.networkice.com/advice/intrusions/2002802/default.htm>.

Alert 15. SUNRPC highport access!

This alert indicates that an internal host was queried for a Solaris RPC 'portmapper' port. This will allow an attack to enumerate the 'high-port' services running on the target system.

Alert 16. Connect to 515 from inside

This alert indicates that a system on the internal network is attempting to initiate a connection with the Unix-based LPD (printer) service on port 515. In GIAC Enterprises, some of the internal hosts are attempting to initiate this connection with external hosts. This is probably not normal behavior.

Alert 17. Broadcast Ping to subnet 70

This alert is generated when a host sends a 'ping' to an entire subnet. A ping is an ICMP Echo request packet that will cause the target system(s) to send an ICMP Echo Reply, showing that the host is up. An attacker will send a ping to a subnet for a variety of reasons. This is a quick way of mapping a network, as one packet will generate a response from every listening target system. This can also be used to fingerprint the operating system of a host, as a MS Windows system will not respond to a broadcast ping. A broadcast ping can also be used as a Denial of Service (DoS) attack, by sending out many broadcast pings to multiple subnets using a fake ('spoofed') source IP address. As each *nix host on each subnet responds, the system at the faked IP address will be overwhelmed by ICMP Echo Response messages.

Alert 18. TCP SMTP Source Port traffic

This alert is generated in response to any TCP traffic with a source port of 25. SMTP traffic usually has a high source port (above 1024) and a destination port of 25. Traffic that originates on TCP port 25 is unusual, and therein suspect.

Alert 19. Back Orifice

This is an alert when traffic is sent to port 31337 in search of the Back Orifice Trojan. This traffic may be a simple scan for an infected machine, or it might be communication with an infected machine. More Information on Back Orifice can be found at <http://www.networkice.com/advice/phauna/rats/back%5Forifice/default.htm>. In the case of GIAC Enterprises, the Back Orifice traffic fits the profile of a scan, with no activity appearing to come from a GIAC system.

Alert 20. External RPC call

This alert is generated when an external host attempts to access the portmapper port on a Unix system. The portmapper service provides information on the services running on a computer, and can give an attacker a lot of information about the target system.

Alert 21. Probable NMAP fingerprint attempt

Similar to the NMAP TCP Ping! Alert, this alert indicates that someone is attempting to determine the operating system (OS) of a remote computer by examining the responses of a few specific network packets. Each OS has a different TCP/IP stack, with it's own 'quirks' that respond differently to various stimuli. More information on how OS fingerprinting works is available at <http://www.insecure.org/nmap/nmap-fingerprinting-article.txt>.

Alert 22. SITE EXEC - Possible wu-ftpd exploit - GIAC000623

This alert is generated when the string 'site exec' is sent to an FTP service. 'Site exec' is part of an exploit on the wu-ftpd ftp service, and does not occur often under normal circumstances.

Alert 23. Happy 99 Virus

This alert is caused by the propagation of the email Happy99 worm. This alert was generated when an external host tried to infect an internal host via email. More information on this worm can be found at <http://www.symantec.com/avcenter/venc/data/happy99.worm.html>.

Alert 24. STATDX UDP attack

This alert is caused by an attempt to compromise a system via vulnerability in the RedHat Linux rpc.stad service. This attack originated externally. More information is available at <http://www.kulua.org/Archives/kulua-l/200008/msg00159.html>.

Alerts Originating Internally

The previous alert analysis has shown a variety of external threats. Unfortunately, there are quite a few internal ones as well. Some of the alerts logged can be attributed to false positives – normal behavior that happens to fit an attack profile – but some cannot be attributed to false positives. The following chart shows the breakdown of all internally originating alerts.

All Alerts With Internal Source	Alert Count
Russia Dynamo - SANS Flash 28-jul-00	442

SNMP public access	418
NMAP TCP ping!	262
Connect to 515 from inside	159
SMB Name Wildcard	78
Tiny Fragments - Possible Hostile Activity	7

Internal Alert 1. Russia Dynamo – SANS Flash 28-jul-00

The most common attack from the list of internal alerts is the Russia Dynamo alert. The cause of this alert is a Windows Trojan, discovered in January of 2000, that sends information to a computer in Russia, in the 194.87.6.255 range. More information on this alert can be found at <http://archives.neohapsis.com/archives/sans/2000/0068.html> and at <http://www.sans.org/y2k/072818.htm>. All the logged activity is from MY.NET.208.138, going to 194.87.6.38, a host in Russia. As this Trojan is almost a year old, antiviral software will be able to detect and remove it. This activity indicates that either there is no policy in place mandating antivirus software be installed, operating and updated on all computers, or the policy is being ignored. This computer should be taken offline and corrected at once. It would be prudent to check all internal systems to verify that antiviral software is installed and current.

Internal Alert 2. SNMP public access

This alert indicates that the 'Public' community was queried on a Simple Network Management Protocol (SNMP) enabled device. A full description can be found at <http://advice.networkice.com/advice/reference/networking/snmp/default.htm>. SNMP is listed in the top ten threats by sans, as of January 18, 2001. This list is available at <http://www.sans.org/topten.htm>.

Devices that report sensitive information should have their community names changed to something other than the defaults of 'public' and 'private'. The majority of this traffic is from internal host to internal host, and is probably legitimate queries. Some of the SNMP queries, however, are externally initiated. For example, 128.46.156.231, a system at Purdue University in Indiana, was able to query several internal machines.

Internal SNMP Public Access		
Source Address	Destination Address	Alert Count
128.46.156.231	MY.NET.100.143	36
128.46.156.231	MY.NET.100.206	21
128.46.156.231	MY.NET.100.99	104

These machines should be reviewed to determine what data the attacker at 128.46.156.231 was able to retrieve. Also, border routers/firewalls should be modified to block all externally originated SNMP traffic unless there is a strong business justification for this open access.

Internal Alert 3. NMAP TCP ping!

This alert is caused when the port scanning, operating system fingerprinting tool Nmap is used to map a network. This tool will provide an attacker with a significant amount of information about a target host or network. There is very seldom a use for this tool legitimately within an organization, unless if it is being used by that organizations' computer security department. Nmap is described in depth at <http://advice.networkice.com/advice/intrusions/2001526/default.htm>. All internal nmap scanning was done from one host, MY.NET.70.38, scanning all hosts between MY.NET.0.0 and MY.NET.0.144. This system should be reviewed, and users questioned, to determine if has been compromised, or is being inappropriately used by an internal user. A review of this host can be found in the Hosts section later in this document.

Internal Alert 4. SMB Name Wildcard

This alert is generated when a query is sent to a windows system, or a *nix system running Samba, to enumerate the available shares on that system. This activity is quite normal at times, yet it can reveal a lot of information to an attacker. Some viruses may also show this behavior as they try to infect shared drives on a network. The logged Snort alerts show a recurring one-to-one relationship between four unique source hosts and four separate destination hosts. This fits the profile of normal network activity, and should be considered a false positive.

Internal Alert 5. Tiny Fragments – Possible Hostile Activity

As discussed above, this alert is generated when an unusually small packet is detected. All 7 of these alerts represent traffic from MY.NET.219.122, port 4000, to 208.162.62.208, NULL port. This host should sound familiar – it is already mentioned above as being individually reviewed later in this document. This is most likely *not a good thing*TM. Traffic in tiny fragments is unusual. Traffic to Null ports is unusual. Doubly unusual traffic from an internal host to an external host is disconcerting. Host MY.NET.219.122 is reviewed in the Hosts section of this document.

General Port Scans

In addition to the scans mentioned in the alerts above, there were an additional 38269 unique port scans between November 24, 2000 and January 18th, 2001. These consisted of two separate types of scans. The first of these is known as a stealth port scan. This type of scan attempts to avoid detection by starting, but not completing, a full TCP connection. There were 21059 of these. The second type of scan is very simple – a system just tries to connect to as many ports and systems as possible. This is referred to as threshold scans as they are logged when a system connects to more hosts in a given time than a specified threshold allows for. There were 17210 of these scans.

Internally Originating Scans

There were a disturbing number – 30599 – of scans that originated from hosts within GIAC Enterprises. This is 79% of the total scans. If there is network-mapping software running within GIAC enterprises, then there is no problem. Otherwise the scanning systems need to be investigated for compromise or improper use by the employees of GIAC Enterprises. The following chart shows the top ten internal sources, accounting for 57% of the total scans, and 71% of the internally originating scans.

Internally Originating Port Scans		
Scan Type	Internal Source Address	Scan Count
Stealth	MY.NET.217.150	6256
Stealth	MY.NET.217.158	4926
Threshold	MY.NET.100.230	3009
Stealth	MY.NET.219.126	2193
Threshold	MY.NET.253.24	2002
Stealth	MY.NET.217.126	1461
Stealth	MY.NET.217.182	1328
Threshold	MY.NET.140.21	498
Threshold	MY.NET.1.3	330

It is disturbing that so many of the internal scans are stealth scans, a fact that suggests the scans are neither accidental nor innocent. Each of these systems needs to be thoroughly reviewed for compromise or inappropriate usage.

Externally Originating Scans

Only 21% of the port scans logged originated from external hosts. The following chart shows the top ten external port scanning IP addresses.

Externally Originating Port Scans		
Scan Type	External Source Address	Scan Count
Threshold	212.64.74.169	336
Threshold	24.7.86.215	302
Stealth	24.113.198.51	272
Threshold	216.99.200.242	270
Threshold	24.3.0.36	195
Threshold	152.163.206.134	133
Stealth	63.78.39.192	126
Threshold	24.189.31.228	112
Threshold	62.227.243.120	91
Threshold	164.67.22.71	63

The port scans are not particularly complicated, and come from a diverse number of sources. The only commonality between the source IP addresses is that several originate from the @Home network, which is a very common source for attacks.

Individual Host Analysis

This section is for all the hosts, both internal and external, that stuck out as unusual in the alert analysis.

MY.NET.219.122, 208.162.62.208, 128.2.166.68

MY.NET.219.122 was flagged for review because it is an internal host – one that is on the network of GIAC Enterprises - and was logged sending tiny fragments to an external host. The hosts 208.162.62.208 and 128.2.166.68 are being reviewed because they are external hosts receiving questionable traffic from an internal host.

Host	Owner
MY.NET.219.122	Internal (GIAC Enterprises)
208.162.62.208	ALAWEB.COM (ISP). AI, US
128.2.166.68	Carnegie Mellon University. Pa, US

The only other alert caused by MY.NET.219.122 is a ‘Connect to 515 from Inside’ going to 128.2.166.68. There are no other alerts to/from 208.162.62.208 or to/from 128.2.166.68. However, this simply means that none were logged. It is possible that the logging sensors failed to log some activity because of high traffic load (dropped packets), because the communication did not fit any of the attack ‘signatures’ being used, or because of a temporary logging system outage. The internal system should be immediately quarantined and reviewed for compromise. If a compromise is likely, then administrators for the other two systems should be notified. It is possible that the administrators for the other two systems will have additional logs pertaining to MY.NET.219.122, which may be quite helpful.

MY.NET.1.8

This host is being individually reviewed as it is receiving substantially more tiny fragments than any other host.

Alerts to MY.NET.1.8	
Alert Type	Alert Count
Tiny Fragments - Possible Hostile Activity	3148
NMAP TCP ping!	63
DNS udp DoS attack	6
SYN-FIN scan!	1
connect to 515 from outside	1

There are a total of 3219 individual alerts, comprising of 5 separate alert types, pertaining to traffic going to host MY.NET.1.8 from 31 separate attacking hosts. The top 10 of these are depicted in the table below.

Top 10 Sources For Alerts to MY.NET.1.8	Alert Count
202.205.5.10 - China - cernet.edu.cn	521
202.101.43.222 - China - sldt.com.cn	344
61.134.9.133 - China - public.xa.sn.cn	317
61.140.75.3 - China - chinanet.cn.net	289
202.96.96.3 - China - chinanet.cn.net	265
202.108.43.152 - China - chinanet.cn.net	261
210.12.160.130 - Beijing, China - chinagb.net	254
202.108.43.151 - China - chinanet.cn.net	225
61.140.75.4 - China - chinanet.cn.net	157
61.140.75.5 - China - chinanet.cn.net	153

The majority of attacks to MY.NET.1.8 are from the China address space. It should be considered whether there is sufficient justification to block traffic from some of these IP address ranges at the border router/firewall. The host MY.NET.1.8 should be checked for compromise. There are no alerts logged with a source of MY.NET.1.8, which suggests that the system is still secure. Security-related configurations should be double checked, and unnecessary services should be disabled.

MY.NET.253.112

This host is being reviewed because it is receiving a disproportionate amount of SYN-FIN Scans – twice as many as any other system. The SYN-FIN scans are distributed fairly evenly through most of the GIAC hosts, so this host sticks out as unusually popular. Looking at the alerts that it is involved in, 262 of its 294 alerts involve the Watchlist 000220 – hosts in Israel. All of this traffic is to port 443. If MY.NET.253.113 is expected to process HTTPS/SSL secure web traffic from systems in Israel, then there probably is not much of a problem with this host.

MY.NET.219.114

This host is being reviewed as it received almost twice as many Queso scans than any other host. All alerts involving this host are Queso scans. All but one of the 205 Queso scans targeting this system originated from 206.65.191.129, a system on the UUnet network. The only logged traffic from that host were the Queso scans, occurring on November 28th, 2000, between 12:02 pm and 12:44 pm. This is very unusual for a Queso scan. The external host only targeted one system, and it did so over 200 times in clustered groupings in less than 45 minutes. It is quite likely that this is a false positive – or the world's dumbest script-kiddie.

Analysis Summary

The network of GIAC Enterprises is the target of a large amount of suspicious, and often hostile, activity. This activity can be divided into two main groupings, based on origin; internal and external.

Externally, a large number of geographically distributed attackers are probing the GIAC systems for weaknesses. Without deeper knowledge of the IDS sensor placement, it is not possible to determine how many of these attempts were successful. If the sensors are placed outside of GIAC's firewalls, then these attacks are significantly less worrisome. If the sensors are placed on the internal network, then GIAC's network defenses are very inadequate, and much of this activity is potentially quite dangerous. Some traffic from internal hosts to external hosts seems to be in response to the external probing. This suggests that the network defenses are, at least to some extent, inadequate as the probes are getting through. However, some of this activity could be attributed to external business partners who are expected to access the internal hosts. Network perimeter defenses (firewalls) should be reviewed, and modified to be more effective at blocking the activity noted within this report.

Many of the externally originating alerts come from a few clustered IP address ranges. Examples of these are Israel, the Chinese Computer Academy, ChinaNet, and TU Dresden in Germany. These groups should be reviewed to determine if customer and business partners exist within those ranges, to determine if those addresses can be blocked at the perimeter.

There were several internal hosts that received a disproportionate number of attacks from external sources. These hosts should be investigated on an individual basis to determine why they are singled out. It should be verified that each of these systems is properly patched. If all of the MY.NET systems do not need to be directly Internet accessible, then a proxy/NAT architecture should be implemented to further protect each host from external hostilities.

Internally, there are some very obvious problems. These range from computer virus infection, to large amounts of port scanning. The fact that most of the internal scanning activity is done using 'stealth' scans, using crafted packets, strongly suggests that either several internal hosts are compromised, or those hosts are being inappropriately used by GIAC employees. Both scenarios should be considered likely, with the scanning hosts being reviewed by qualified personnel, and the system operators questioned about usage policies. If an 'appropriate use' policy does not exist regarding computer resources, then such a policy should be immediately developed and distributed to employees.