Global Information Assurance Certification Paper

*** Northcutt, Accuracy is fine, analysis process is great, love detect 8 with the research, though the paper could have benefited from a bit more. Still what's to complain about, a fine job. 89 *

GIAC Certified Intrusion Analysts (GCIA)
Practical Submission, 10 Detects

John Eisenhauer

April 12, 2000

Detect Background:
All 10 detects are from my network. We have only one Internet connection. A screening router and firewall protect our network. Our web/mail/ftp servers all sit on a separate leg of the firewall. We are fortunate not to have to allow inbound connections through the screening router, or firewall to anything other than a handful of web hosts, and then only on specific ports. Outbound access is limited to a specific set of ports as well.

The detects I am submitting come from two sources. The first is a syslog file that the Cisco screening router logs are written to. We change the file format slightly to make parsing easier, but I will try to strip that out in each trace.

The second is windump. We already have a commercial IDS that runs on NT placed at several key choke points on our network. These IDS boxes use up the monitor port at those locations (we are a 100% switched network.) After seeing the benefit of the additional information that tcpdump/windump is able to provide, I decided to run windump on those boxes in a configuration similar to a shadow sensor. I have been doing this for the last two weeks and even when logging 200mb/hr (no filters) windump doesn't drop packets or interfere with my commercial IDS. Just manually running filters on the windump captures to eliminate 'normal' traffic has revealed some very interesting things. (I am trying to get Shadow to function with these NT boxes as the sensors. If I am successful in this, I will share the process with anyone interested.)

I have sanitized the traces by replacing all of my addresses with hostx.mynet.com, and the source addresses with something else that is made up.

Note: The number I assign to "Network Countermeasures" will vary from detect to detect. This is because I take into consideration the effectiveness of the specific exploit on my screening router / firewall combination when generating this number. I think that this approach will work for me as long as I am consistent across the board. Having only one site, this gives me a broader range of overall severity figures that I can use to weigh problems against one another should I have a "Bad Day."

TRACE ONE

This detect comes from our Cisco Access Router. One of out log parsing filters looks for certain commonly exploited ports including portmap.

Mar 27 15:39:57: list ingressfilter denied tcp host-9.scanner.com(2666) -> host8.mynet.com.(111), 1 packet
Mar 27 16:09:49: list ingressfilter denied tcp host-9.scanner.com(2666) -> host2.mynet.com.(111), 1 packet
Mar 27 16:44:29: list ingressfilter denied tcp host-9.scanner.com(2666) -> host3.mynet.com.(111), 1 packet
Mar 27 20:59:14: list ingressfilter denied tcp host-9.scanner.com(2666) -> host4.mynet.com.(111), 1 packet
Mar 28 17:51:49: list ingressfilter denied tcp host-9.scanner.com(2666) -> host7.mynet.com.(111), 1 packet
Mar 29 05:40:38: list ingressfilter denied tcp host-9.scanner.com(2666) -> host1.mynet.com.(111), 1 packet
Mar 29 07:54:51: list ingressfilter denied tcp host-9.scanner.com(2666) -> host9.mynet.com.(111), 1 packet

Active Targeting: YES
History: No instances of other packets denied from this host (by screening router) in month previous or time since. Also, other traffic to target hosts in this time frame is nonexistent or seems unrelated.
Technique: This is a slow scan for a specific vulnerability over a range of hosts. By spreading the scan out over several days the attacker would most likely hope to evade detection. I do not have enough information make a guess as to what tool was used to conduct the scan.
Intent: I would have to assume that the intent here is hostile. The destination port of 111 most likely means that the attacker was looking for one or more portmapper exploits.
Evaluation: Since all of the destination hosts do not exist on this network, I would have to assume that the attacker has little knowledge of the network layout. The fact that all of the source ports are the same across three days and seven scanned hosts indicates but does not prove that these packets are crafted. The time interval and limited number of hosts scanned does lead me to believe that this is either just one small part of a really big random scan, or a much more targeted and stealthy scan for a specific exploit on my network. I will add this source host to my list of systems to watch out for.
Severity: Criticality = 0 Target hosts do not exist.
Lethality = 5 If successfully exploited, portmapper attacks can give attacker remote root access.
System CM = 5 N/A (I rate nonexistent systems a 5 in this category across the board to because the most secure system is one that doesn't exist!)
Network CM = 5 The stateful screening router did it's job and blocked the packets.
Overall Severity = (0 + 5) – (5 + 5) = -5

The next few detects popped out when I took a look the IP addresses that define both ends of the T-1 that provides our internet feed. The goal here was to improve the accuracy of the filters we run our log files through. After eliminating 'normal' traffic destined for the IP addresses that define the serial interfaces this is what I found. Again, I was only looking 'back' a fixed amount of time.

(One word of note to anyone reading this, when profiling 'normal' traffic on your network, be sure to take the serial interfaces of you access router into account. I didn't do this initially, and consequently had a scare when I saw a huge amount of odd-looking ICMP traffic.)

TRACE TWO

Mar 22 01:19:59: list ingressfilter denied tcp host-31.pornsite-xyz.com (80) -> wanport.mynet.com(59865), 1 packet
Apr 5 11:36:22: list ingressfilter denied tcp host-154.isp-pornsite.com (80) -> wanport.mynet.com(8416), 1 packet

Active Targeting: YES
History: Searching through log files for similar traffic to or from the target host, or the source network yielded much information. On the 22nd and 23 of March 400+ packets destined for random hosts on my network were denied by the screening router from the hosts listed above as well as one other host on the same Class C. Also, there was another cluster of traffic from a few days before to a few days after the April 5 packet. Again, 400+ packets were denied. One more item to note is that the source hosts are not running an HTTP server.
Technique: This looks like a host scan that has fixed it's source port to 80. The destination addresses are comprised of a small random set of my routable addresses and are targeted at random ports between 3 and 63995.
Intent: This looks like an attempt at stealthy network mapping to me. By choosing to use 80 as the source address, and randomizing the destination addresses and ports, at first glance this traffic looks normal.
Evaluation: The problem here is much larger in scope than just what hit the router. (I didn't include those history related log entries because they look similar to the ones above, and would just have required more address sanitation.) If this scan had not hit our access router, it most likely would not have been noticed. The stateful conduits that our screening router sets up on the fly do time out from time to time causing valid packets delayed in transit to be denied. These causes log entries that look just like the ones above. More complex filtering of our screening router logs will catch something like this in the future. Having said that, here is what I think happened. There seem to be porn related streaming video servers on the subnet where the scan originated. I think (based on some browsing that the proxy logs showed) that one or more of my users visited this site and triggered 'something'. That 'something' then initiated a scripted scan (probably using nmap) that proceeded to scan about 400 of my addresses each of the two times that it happened. This scan tried to evade detection by looking like return packets to http requests. The problem with this was that several of the addresses didn't exist.
Severity: Criticality = 3 Even though the scans were random, both servers, and desktops were targeted.
Lethality = 1 This was an attempt at network mapping which could have led to a later attack.
System CM = 3 Had these packets made it through the screening router, many of the systems targeted would have responded in some way that would give out information about their existence, OS, and even possible services that they are running.
Network CM = 2 I consider our firewall to be quite restrictive, BUT the detection system in place missed the bulk of this scan. This has been changed now!
Overall Severity = (3 + 1) – (3 + 2) = -1

TRACE THREE

Apr 3 04:25:13: Sig:1102:Impossible IP Packet - from ISPside.wanport.mynet.com to ISPside.wanport.mynet.com

Active Targeting: YES
History: Searching through log files for similar traffic didn't show much. This particular attack has not shown up in the last month. The router interface that the attack was targeted against does not seem to show an abnormal amount of 'interest.'
Technique: This would have to fall into the category of a denial of service 'send and pray' type of attack.
Intent: The intent here would be to crash the target system.
Evaluation: This is a classic example of the land attack. The source and destination ports are the same. The packet is definitely crafted. A capture of the packet would have been nice, but none was available.
Severity: Criticality = 5 This attack targeted a core router.
Lethality = 1 Most systems are patched against this attack now. In general, Land will fail.
System CM = 4 The majority of systems of systems on our network, and ALL systems outside the firewall are patched against this attack.
Network CM = 5 Both the screening router, and firewall are able to defeat this attack.
Overall Severity = (5 + 1) – (4 + 5) = -3

TRACE FOUR

This detect comes from our Cisco Access Router. The packets stood out because they look like return packets to requests from a host on my network that does not exist.

NOTE: This pattern continued from March 20th through April 11. This end date is significant.

Mar 22 10:45:02: list ingressfilter denied udp letter1.root-servers.net(53) -> unused-9.mynet.com(5029)
Mar 22 10:45:03: list ingressfilter denied udp letter1.root-servers.net(53) -> unused-9.mynet.com(5029)
Mar 22 10:45:09: list ingressfilter denied tcp letter3.root-servers.net(53) -> unused-9.mynet.com(5029)
Mar 22 10:45:12: list ingressfilter denied tcp letter3.root-servers.net(53) -> unused-9.mynet.com(5029)
Mar 22 10:45:35: list ingressfilter denied tcp letter5.root-servers.net(53)-> unused-9.mynet.com(5029)
Mar 22 10:45:38: list ingressfilter denied tcp letter5.root-servers.net(53)-> unused-9.mynet.com(5029)
Mar 22 10:45:40: list ingressfilter denied tcp letter2.root-servers.net(53-> unused-9.mynet.com(5029)
Mar 22 10:45:44: list ingressfilter denied tcp letter5.root-servers.net(53)-> unused-9.mynet.com(5029)
Mar 22 10:45:49: list ingressfilter denied tcp letter2.root-servers.net(53-> unused-9.mynet.com(5029)
Mar 22 10:45:54: list ingressfilter denied tcp letter1.root-servers.net(53) -> unused-9.mynet.com(5029)
Mar 22 10:46:06: list ingressfilter denied tcp letter1.root-servers.net(53) -> unused-9.mynet.com(5029)
Mar 22 10:46:14: list ingressfilter denied tcp letter2.root-servers.net(53-> unused-9.mynet.com(5029)
Mar 22 10:46:29: list ingressfilter denied tcp letter1.root-servers.net(53) -> unused-9.mynet.com(5029)
Mar 22 10:46:50: list ingressfilter denied tcp letter5.root-servers.net(53)-> unused-9.mynet.com(5029)
Mar 22 10:46:54: list ingressfilter denied tcp letter2.root-servers.net(53-> unused-9.mynet.com(5029)
Mar 22 10:47:00: list ingressfilter denied tcp letter1.root-servers.net(53) -> unused-9.mynet.com(5029)
Mar 22 11:21:23: list ingressfilter denied udp letter3.root-servers.net(53) -> unused-9.mynet.com(5029)
Mar 22 11:21:24: list ingressfilter denied udp letter3.root-servers.net(53) -> unused-9.mynet.com(5029)
Mar 22 11:21:27: list ingressfilter denied udp letter3.root-servers.net(53) -> unused-9.mynet.com(5029)
Mar 22 11:21:56: list ingressfilter denied udp letter5.root-servers.net(53)-> unused-9.mynet.com(5029)

TRACE FOUR (CONTINUED)

Active Targeting: YES (Collateral Damage?)
History: A search of previous log files was the key here to everything. This suspicious traffic has been going on for at least since the 20th of March, and possibly longer. In continued through April 11th. The destination address for these packets is not in use, and as far as I know has never been.
Technique: Distributed or not, this was part of a DoS attack or possibly a big smoke screen for something else?
Intent: The intent is unclear.
Evaluation: This traffic stood out because it was directed to a non-existent host. Locating similar traffic and importing the log entries into Excel brought out some interesting patterns. Approximately 3600 packets show up. The destination host is the same in all cases. The destination port is either 877 or 5029. Both TCP and UDP packets are being used in about equal proportions. The source port is always 53. The source addresses do vary slightly. In the context of the logs I have searched there are only five different source addresses.

Four of those source addresses are root name servers. Ex. ?.root-servers.net. The other host I will call host1.suspect.net

The pattern goes like this:
Src host1.suspect.k12.state.us port 53, dest unused-9.mynet.com port 5029,
Mar 28 14:13 to 16:10, 8 packets

src letter1.root-servers.net port 53, dest unused-9.mynet.com port 5029, Mar 20 to Mar 27, ~400 packets
src letter1.root-servers.net port 53, dest unused-9.mynet.com port 877, Mar 28 to April 6, ~400 packets

src letter2.root-servers.net port 53, dest unused-9.mynet.com port 5029, Mar 20 to Mar 27, ~500 packets
src letter2.root-servers.net port 53, dest unused-9.mynet.com port 877, Mar 28 to April 6, ~500 packets

src letter3.root-servers.net port 53, dest unused-9.mynet.com port 5029, Mar 20 to Mar 25, ~150 packets
src letter3.root-servers.net port 53, dest unused-9.mynet.com port 877, Mar 28 to April 6, ~250 packets

src letter4.root-servers.net port 53, dest unused-9.mynet.com port 5029, Mar 20 to Mar 27, ~500 packets
src letter4.root-servers.net port 53, dest unused-9.mynet.com port 877, Mar 28 to April 6, ~500 packets

ALSO, I found these two packets. They are from what looks to be the NAT pool at suspect.k12.state.us.

Mar 28 9:48:54 %IDS-4-ICMP_ECHO_REPLY_SIG: Sig:2000:ICMP Echo Reply - from nat99.suspect.k12.state.us to unused-9.mynet.com
Mar 28 9:49:40 %IDS-4-ICMP_ECHO_REPLY_SIG: Sig:2000:ICMP Echo Reply - from nat115.suspect.k12.state.us to unused-9.mynet.com

Could it be that someone at suspect.net is making sure that our host is still not live so that they can use it as part of a Do attack, or as a decoy? I think it is very significant that these two unsolicited echo replies came just at the point where the ports changed from 5029 to 877.

I did the right thing here. I called the school district in question. After being transferred to the person I think is in change of the IT department, I explained my self. The manager pulled their router person into the room and put us all on speakerphone. I assured them that I was not accusing them of anything, and described the traffic. I emphasized the small amount of traffic I had seen from their network to the unused address on my network. Also, I pointed out that the unsolicited echo replies came just as the port used changed from 5029 to 877.

In talking to them it became clear that they had no outbound access control lists in their router. Also, they had recently put into place a high level of logging in their firewall. They assured me that they would search the logs and get back to me.

They have not called back, BUT the traffic stopped about 15 minutes after I called them. It has not shown back up since. It has been 24hrs now, and in the previous month, the traffic had only stopped for a total of just a few hours.

Severity: Criticality = 5 Our core hosts might not be effected by this, but someone else may very well be.
Lethality = 4 Again, not us, but someone may be subject to a DoS.
System CM = 0 The systems on our end cannot be patched to help with this.
Network CM = 0 We have nothing I know of that could prevent this at it's source.
Overall Severity = (5 + 4) – (0 + 0) = 9 I treat this kind of thing seriously. This attack might not be directed at my site, but I would hope anyone with the same level of information about an attack directed at me would be as diligent.

TRACE FIVE

This detect comes from our Cisco Access Router. The packets stood out because we routinely look for popular Trojan ports. False alarms can usually be eliminated by finding that the detect was only a small part of legitimate traffic and just randomly hit a suspect port. I don't think that was the case here.

```
Apr 10 00:32:20: list ingressfilter denied tcp xyz.abc.jkl.43(42357) -> nat-21.mynet.com(7890)
Apr 10 00:34:31: list ingressfilter denied tcp xyz.abc.jkl.43(3351) -> nat-99.mynet.com(31337)
Apr 10 00:36:22: list ingressfilter denied tcp xyz.abc.jkl.71(36529) -> nat-52.mynet.com(13130)
```

Active Targeting: YES
History: A search of the log files for this suspect subnet produces the entry before and after the 31337. These hosts have not tried to connect to any other hosts on my network in the last month.
Technique: This looks suspiciously like someone I trolling for trojans to me.
Intent: The intent would be to find a system that was open to 'remote administration.'
Evaluation: This kind of thing seems pretty common. I would worry more if a larger number of hosts were scanned for suspect ports, or if a more important system was scanned. I don't know what the 7890, and 13130 ports are. They are not registered with iana or on my trojan list. These may be decoy ports, or unknown Trojan ports.
Severity: Criticality = 1 This is in the range of our NAT pool. It is used by regular desktops.
Lethality = 3 A compromised system would give an attacker user access on the net.
System CM = 4 The workstations are running newer operating systems, but more importantly are all running Antiviral software with signatures less than two weeks old.
Network CM = 5 Both the screening router, and firewall are able to defeat this attack.
Overall Severity = (1 + 3) – (4 + 5) = -5

TRACE SIX

In testing a script meant to detect host scans in out access router logs I found this traffic.

Mar 27 14:37:35 list ingressfilter denied udp not-too-stealthy.scanner99.net(137) -> xxx.yyy.zzz.2(137)
Mar 27 14:37:37 list ingressfilter denied udp not-too-stealthy.scanner99.net(137) -> xxx.yyy.zzz.2(137)
Mar 27 14:37:37 list ingressfilter denied udp not-too-stealthy.scanner99.net(137) -> xxx.yyy.zzz.2(137)
Mar 27 14:37:46 list ingressfilter denied udp not-too-stealthy.scanner99.net(137) -> xxx.yyy.zzz.3(137)
Mar 27 14:37:48 list ingressfilter denied udp not-too-stealthy.scanner99.net(137) -> xxx.yyy.zzz.3(137)
Mar 27 14:37:49 list ingressfilter denied udp not-too-stealthy.scanner99.net(137) -> xxx.yyy.zzz.3(137)

This continued three packets per host, through the entire Class C.

Mar 27 15:22:02 list ingressfilter denied udp not-too-stealthy.scanner99.net(137) -> xxx.yyy.zzz.254(137)
Mar 27 15:22:03 list ingressfilter denied udp not-too-stealthy.scanner99.net(137) -> xxx.yyy.zzz.254 (137)
Mar 27 15:22:05 list ingressfilter denied udp not-too-stealthy.scanner99.net(137) -> xxx.yyy.zzz.254 (137)

Active Targeting: YES
History: This IP does not show up in the previous month's logs, or at any time since.
Technique: This looks like a host scan for MS Networking exploits. I don't think it's fast enough to be nmap, unless the scan interleaved our class c with several others. In that case though, I would have expected the hosts scanned to be random. They were not.
Intent: The intent would be to find a system with unprotected shares to exploit
Evaluation: This looks like an inexperienced attacker that doesn't realize that this kind of activity is logged in many cases. It seems interesting to me that they didn't even randomize the destination IP's. Also, it may be significant that the packets came in pairs of three. This may help identify the tool used.
Severity: Criticality = 0 I'm embarrassed to admin it, but that entire class C is not used.
Lethality = 3 An attacker could get user access at least, or plant a trojan on an unprotected share.
System CM = 2 I would have to say that if an attacker circumvented the firewall, a scan that Large would yield at least a few unprotected shares.
Network CM = 5 Both the screening router, and firewall are able to defeat this attack.
Overall Severity = (0 + 3) – (2 + 5) = -4

TRACE SEVEN

This showed up when testing a windump filter that looks for address leakage. That is, systems connected to my LAN and some other network. These addresses are **not** sanitized because in this case I spent a lot of time chasing a false alarm.

This is the traffic that first caught my attention:
(I removed some lines from the middle. There were 31 to begin with)
14:34:38.799587 169.254.166.1.137 > 169.254.255.255.137: udp 68
14:34:38.802087 169.254.166.1.137 > 169.254.255.255.137: udp 68
14:34:38.818519 169.254.166.1.137 > 169.254.255.255.137: udp 50
14:34:38.820172 169.254.166.1.137 > 169.254.255.255.137: udp 50
14:34:42.544225 169.254.166.1.138 > 169.254.255.255.138: udp 182
14:34:42.546021 169.254.166.1.138 > 169.254.255.255.138: udp 182
14:34:42.551522 169.254.166.1.138 > 169.254.255.255.138: udp 212
14:34:42.553237 169.254.166.1.138 > 169.254.255.255.138: udp 212
14:34:44.061544 169.254.166.1.137 > 169.254.255.255.137: udp 50

I response to the above anomalous traffic, I searched historical data that I had captured using Windump at my default gateway. I was alarmed when I saw 30+ thousand packets going from my PDC, and several BDC's to the address in question in a one-day period. Here is a sample….
22:10:02.864907 mypdc.mynet.com138 > 169.254.166.1.138: udp 238
22:10:05.430314 mypdc.mynet.com138 > 169.254.166.1.138: udp 238
22:46:14.140096 bdc1.mynet.com4172 > 169.254.166.1.1044: udp 8
23:12:27.222224 bdc1.mynet.com4937 > 169.254.166.1.1044: udp 8
00:00:17.200967 mypdc.mynet.com138 > 169.254.166.1.138: udp 226
00:00:17.201298 bdc2.mynet.com138 > 169.254.166.1.138: udp 224
00:00:17.201881 bdc3.mynet.com138 > 169.254.166.1.138: udp 224
00:00:17.263243 bdc1.mynet.com138 > 169.254.166.1.138: udp 230
00:00:17.264096 bdc4.mynet.com138 > 169.254.166.1.138: udp 228
00:00:17.264140 bdc5.mynet.com138 > 169.254.166.1.138: udp 228
00:00:17.264402 bdc7.mynet.com.138 > 169.254.166.1.138: udp 224
00:00:17.266703 bdc8.mynet.com.138 > 169.254.166.1.138: udp 224
00:00:17.275867 bdc9.mynet.com.138 > 169.254.166.1.138: udp 224

Active Targeting: NO, This turns out to be a "feature" in Windows….
History: Unknown, I did not have logging in place that was able to detect this address in the past.
Technique: This looked like address leakage (my net bridged by a workstation to another net.)
Intent: The intent had this not been a false alarm would be to circumvent, intentionally or not, the firewall.
Evaluation: A source and destination address not belonging to me being seen at my default gateway would indicate some sort of address leakage as I call it. Call it whatever you want, but in any case it is a potential back door into my network. This should always be followed up on.

Also, any time that many packets go from my infrastructure systems (in this case a PDC and several BDC's) to an unknown case it's time for alarm. As an astute network administrator pointed out to me, the 169.254.x.x network is used by Microsoft DHCP clients as a default when DHCP requests fail. This was the case this time, and the cause of a false alarm. I included this detect so that others may learn from my mistake. The PDC and BDC's were trying to do some kind of broadcast resolution as well, caused by the original misconfigured windows box.

Severity: Criticality = 5 Core hosts were involved.
Lethality = 0 This was not an attack
System CM = 0 No countermeasures are needed.
Network CM = 5 Both the screening router, and firewall are configured with outbound ACL's to block similar traffic.
Overall Severity = (5 + 0) – (0 + 5) = 0

TRACE EIGHT

This also showed up when testing the same windump filter as above. We were looking for 'address leakage' again. This time we found some.

23:13:04.379538 dhcp-43.cablemodem.isp.net.137 > ns2.isp.net.53: 32+ A? USERDOMAIN.isp.net. (47)
23:13:04.379891 dhcp-43.cablemodem.isp.net.137 > ns2.isp.net.53: 38+ A? USERDOMAIN.isp.net. (47)
23:13:05.879477 dhcp-43.cablemodem.isp.net.137 > ns3.isp.net.53: 26+ A? USERDOMAIN.isp.net. (47)
23:13:05.880187 dhcp-43.cablemodem.isp.net.137 > ns3.isp.net.53: 32+ A? USERDOMAIN.isp.net. (47)
23:13:05.880233 dhcp-43.cablemodem.isp.net.137 > ns2.isp.net.53: 38+ A? USERDOMAIN.isp.net. (47)
23:13:07.276172 dhcp-43.cablemodem.isp.net.137 > ns3.isp.net.53: 32+ A? USERDOMAIN.isp.net. (47)
23:13:07.360750 dhcp-43.cablemodem.isp.net.137 > ns2.isp.net.53: 38+ A? USERDOMAIN.isp.net. (47)
23:13:08.878850 dhcp-43.cablemodem.isp.net.137 > ns3.isp.net.53: 32+ A? USERDOMAIN.isp.net. (47)

Active Targeting: Yes
History: We have suspected for some time that dial-in users may be connected to the Internet via cable modem while concurrently connecting to our network. The implications of this are very bad. I will not go into that aspect of things. We now through our network of sensors have some visibility to this. Matching dial-in logs to the start and stop times of the traffic pointed to one particular user. We have suspected them in the past as being an offender in this area. A manager will talk to them again.
Technique: An uninformed user purchases a cable modem for home, connects their home system to it, and then dials into our corporate network. This effectively opens a backdoor into our network from the Internet.
Intent: The intent would be to circumvent the monitoring and access control that we have on our firewall.
Evaluation: These new 'always on' Internet connections are such a headache. Users want fast Internet access at home, and go out and get a cable modem. Regardless of intent, or whether or not the user knows what they are doing, they then dial into our network with their cable modem connected PC. This makes their unprotected PC a back door into our private network. A policy has been written to ban this, but we now need to find ways to identify the problem when it occurs.

In this case, when the Windows box connected via RAS, it's default gateway changed to our network. For whatever reason, it then began to send out traffic destined for its ISP through our network. This is what we see in the trace.
Severity: Criticality = 5 Core hosts could be exposed.
Lethality = 3 An attacker exploiting this potential back door could gain user access on the network at least.
System CM = 2 This particular kind of thing can expose 'weak' systems to common attacks.
Network CM = 1 The hope we have here is that our detection system catches this early on. The firewall cannot help with this.
Overall Severity = (5 + 3) – (2 + 1) = 5

TRACE NINE

This showed up when a windump filter was run on captured traffic to remove the 'well defined' traffic from everything gathered. Some very interesting things showed up in the 'everything else' lump that was left over. This is one of those items.

Note: These represent only a few of the 50+ packets that made up the original trace. The rest of the original trace items are not included because they looked similar to these.

09:27:51.718624 suspecthost.popular.com > proxy1.mynet.com: (frag 20333:142@1176)
09:27:54.328964 suspecthost.popular.com > proxy1.mynet.com: (frag 20344:143@1176)
09:27:55.779431 suspecthost.popular.com > proxy1.mynet.com: (frag 20345:142@1176)
09:28:02.518749 suspecthost.popular.com > proxy1.mynet.com: (frag 62346:143@1176)
09:28:03.912606 suspecthost.popular.com > proxy1.mynet.com: (frag 62424:142@1176)
09:28:05.305099 suspecthost.popular.com > proxy1.mynet.com: (frag 62462:143@1176)

Active Targeting: Yes
History: The target host in question is a proxy server on our sire. The suspect host seems to send us a few malformed packets every day. Sometimes they have invalid packet headers, SYN FIN, etc, or are invalid fragments.
Technique: This may not be intentional. In any case, invalid fragments such as these can be an attempt at a DoS, or a buffer overflow.
Intent: The intent here is unclear. I have sent email to the hostmaster for the site, and am still looking for a phone number. I really would like to find out whether or not a problem at their site is causing the traffic.
Evaluation: There are several things wrong in this trace. The offset is the same for all of the fragments to start. Also, the sequence numbers are sequential, and then take a big jump. In addition to that, the fragment size changes from packet to packet.

The source host in question is just one of the servers that comprise a rather large, and popular Internet site. Since the site in my opinion is not business related I would like to simply block it at the screening router. However, reality dictates that I need to work with the people in charge of the site to find out why the bad packets are being sent to me. It may even be possible that one of their servers has been compromised.

Severity: Criticality = 3 This is one in a series of redundant outbound proxy servers.
Lethality = 1 The O.S. that the proxy server runs on is up to date, as is the proxy server software. The malformed packets they send us have not caused a problem with the server so far.
System CM = 5 This system countermeasures are what makes the lethality low in this case.
Network CM = 4 The screening router blocks most malformed packets. These fragments still made it past the screening router and were blocked by the firewall.
Overall Severity = (3 + 1) – (5 + 4) = -5

TRACE TEN

This showed up when another windump filter was run on a different sensor's captured traffic to remove the 'well defined' traffic from everything else gathered.

12:04:28.085401 hij.klm.nop.159 > www7.mysite.com: icmp: hij.klm.nop.159 tcp port 12702 unreachable
12:10:41.412041 hij.klm.nop.159 > www7.mysite.com: icmp: hij.klm.nop.159 tcp port 14547 unreachable
12:10:41.666493 hij.klm.nop.159 > www7.mysite.com: icmp: hij.klm.nop.159 tcp port 13438 unreachable
12:10:51.484451 hij.klm.nop.159 > www7.mysite.com: icmp: hij.klm.nop.159 tcp port 15789 unreachable
12:11:55.642138 hij.klm.nop.159 > www7.mysite.com: icmp: hij.klm.nop.159 tcp port 14749 unreachable
12:12:56.923487 hij.klm.nop.159 > www7.mysite.com: icmp: hij.klm.nop.159 tcp port 13548 unreachable
12:13:08.380646 hij.klm.nop.159 > www7.mysite.com: icmp: hij.klm.nop.159 tcp port 15526 unreachable
12:21:40.892271 hij.klm.nop.159 > www7.mysite.com: icmp: hij.klm.nop.159 tcp port 11747 unreachable
12:21:40.950361 hij.klm.nop.159 > www7.mysite.com: icmp: hij.klm.nop.159 tcp port 12699 unreachable

Active Targeting: Yes, but this turns out to be innocent traffic.
History: The destination host in this case is a web server. It serves no other purpose. It doesn't see many ICMP port unreachable messages. This is what got my attention. This particular source host has not been to our site before today. However, further analysis of that day's logs indicates that the host had been 'surfing' our web site off and on for more than one hour.
Technique: This turns out to just be web browsing.
Intent: There is no hostile intent here.
Evaluation: This is a situation where having a windump log file of the previous hour's traffic was very helpful. I was able to find out much more about the 'conversation.'

I won't bore anyone with the trace from the http access except for these two lines.
Here, the source host can be seen gracefully closing the connection.

11:21:39.911013 hij.klm.nop.159.12699 > www7.mynet.com.80: F 979:979(0) ack 5390 win 17940 (DF)
11:21:40.649243 hij.klm.nop.159.11747 > www7.mynet.com.80: . ack 1541 win 17940 (DF)

The source ports on the client side during the surfing seemed to be in line with the ports referenced in the trace.

Between 40 minutes and 1 hr later, the icmp port unreachable messages are received. My best guess is that the packets were delayed in transit back to the browser. By the time they got there, the browsing system was no longer expecting traffic to those ports.

Severity: N/A