



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>

*** Northcutt, number one is a fun read and a good reason to keep sweeping your net work Win2k boxes you don't know you have. Good set of detects, sometime the analysis was hard to decode. 76 *

Practical for GIAC Certified Intrusion Analyst Certification

By Suzanne Hernandez

April 23, 2000

These detects were done on my own network and the destination addresses have been changed.

Detect 1 – Vulnerability in DNS

Apr 03 02:05:33 ns1 named[1337]: unapproved update from [13.1.1.1].1025 for our.domain

This message came from our syslog running on our dns server. I had never seen this message before and investigated it in DNS and BIND. It turns out this is a feature in 8.x version of bind that allows you to remotely add to or delete records from a domain. We run a second level domain on our network, so further investigation was necessary even though the update was unsuccessful. This feature is necessary to allow a dhcp server to add records to a dns server whenever a machine registers itself with the dhcp server. The nsupdate works on udp port 53 just like a regular dns query.

To recreate this, I created a bogus domain with an mx record and added the following to the options section in the dns configuration file:

```
allow-update { any; };
```

This option is off by default.

I then typed the following from another server:

```
[root@server] # nsupdate
➤ update delete bogus.domain. in mx
➤ update add bogus.domain. 99999999 in mx 0 hacker.exchange.server.
```

When this was completed, I did an mx lookup and got the following result:

```
Nslookup q=mx bogus.domain
```

Bogus.domain preference = 0, mail exchanger = hacker.exchange.server

Now I have effectively hijacked bogus.domain's mail and redirected it to my exchange server.

Evidence of active targeting: Absolutely. They specifically targeted our second level domain. Therefore they targeted the dns server. The nsupdate command figures out what the primary domain name server is for the domain its trying to update. I was unable to catch what record or records they were trying to add or delete and in my recreation chose what I would do as a malicious attacker. I would hijack the email of an important network.

Existence: The network was AT&T ITS. They own the entire 12 network and when I called up they told me they are an ISP.

History: I searched through all our archived syslogs for our dns servers and did not find any previous attempts to update our DNS servers. We were already blocking a small portion of the 12 network, so we have seen malicious behavior before, but I expect to see port and ip scans from ISP's., especially such large ones.

Severity: Medium. I would have given this a severity of high, but did not for several reasons. High because there is one and only one level of defense for this vulnerability and that is in the configuration file on the dns server. Since this uses udp port 53 and we allow dns queries from anyone, we do not block this port. Similarly, tcp wrappers would not help for the same reason. The reason I gave this a medium for my network is because I know how well our dns servers are tightened down, I know how few people administer the dns servers and I have made them aware of this vulnerability. Also, we have no dhcp servers and no reason to open this up to even one IP address. I do worry about hurried installations at other locations where stress and time and not thinking might cause someone to open this up to anyone when the dhcp server is not adding records to the domain. I do know the hacker world knows about this and sees it as an opportunity; otherwise we would not have seen a user on an ISP attempting to update our domain.

Detect 2 – Random scan for existence of hosts

Mar 6 01:55:57	tcp dont.tell.anyone.im.an. telnet	x.y.157.72	1082
Mar 6 01:56:23	tcp dont.tell.anyone.im.an. telnet	x.y.182.18	1653
Mar 6 01:57:08	tcp dont.tell.anyone.im.an. telnet	x.y.1.12	1026

dont.tell.anyone.im.an.xconvict.com 3

This log is from my firewall and means the intruder got past the router. Any packet with the ACK bit set gets through the router as the return packet in an established session. Since the router does not allow tcp ports 1082, 1653 or 1026 through, this hacker had to have crafted the packet and was just searching for the existence of these three hosts. He

used port 23 to attempt to blend in with our other telnet ports. My router also does not allow port 23 through so his attempt to “blend” was unsuccessful. These packets hit the firewall because none of these subnets exist on my network.

Evidence of active targeting: Yes, but not targeting us alone. Because of the variation of the source ports, and the time between each attempt, I would say he was scanning several class B networks. Also, he only attempted three machines so I don’t believe he was actively targeting my network in particular.

Existence: The name don’t.tell.anyone.im.an.xconvict.com resolves to 216.95.145.61. The 216.95.145.0 network belongs to mynofee.com out of Calgary, AB and the domain xconvict.com does exist. Just his chosen domain and hostname says something about whether this is someone who would do official business with my network.

History: I have not seen any additional activity from this network in the recent past.

Severity: Low. The intruder was not doing a scan based on any knowledge of our network. He did make it through our internal router, but not through our firewall. He could conceivably find hosts on our external subnets by setting the ACK bit.

Detect 3 – Scan on Network and Broadcast address

```
Apr 8 18:46:03 151.20.138.250 No_Src_Port No_DNS_Name x.y.0.255 No_Dst_Port No_DNS_Name  
No_Flags Echo_Request ping  
Apr 8 18:46:14 buster1 BUSTED: 18:46:03 151.20.138.250 No_Src_Port No_DNS_Name x.y.255.0  
No_Dst_Port No_DNS_Name No_Flags Echo_Request ping  
Apr 8 18:46:14 buster1 BUSTED: 18:46:03 151.20.138.250 No_Src_Port No_DNS_Name x.y.255.255  
No_Dst_Port No_DNS_Name No_Flags Echo_Request ping
```

This was seen on our network intrusion detection system that runs outside of our external router.

Evidence of Active Targeting: Yes, he was targeting the broadcast and network address of our Class B Network.

Existence: This is network 151.20.0.0 which is IUNet out of Italy.

History: We have seen other similar ports scans, ip scans from IUNet and have blocked portions of the 151.20 network.

Severity: Low. Our router does not allow broadcast or network addresses of 0 or 255 through. Our router also does not allow echo requests inbound.

Detect 4 – Denial of Service

Mar 9 17:10:27	tcp	secure1.comned.com	https	x.y.246.7	1536
Mar 9 17:15:20	tcp	secure1.comned.com	https	x.y.246.7	1536
Mar 9 17:18:10	tcp	212.108.4.40	http	x.y.246.7	1536
Mar 9 17:18:16	tcp	212.108.4.40	http	x.y.246.7	1536
Mar 9 17:18:39	tcp	212.108.4.40	http	x.y.246.7	1536
Mar 9 17:18:42	tcp	212.108.4.95	http	x.y.246.7	1536
Mar 9 17:18:45	tcp	212.108.4.95	http	x.y.246.7	1536
Mar 9 17:18:52	tcp	212.108.4.95	http	x.y.246.7	1536
Mar 9 17:19:24	tcp	212.108.4.40	http	x.y.246.7	1536
Mar 9 17:20:15	tcp	secure1.comned.com	https	x.y.246.7	1536
Mar 9 17:23:05	tcp	212.108.4.40	http	x.y.246.7	1536
Mar 9 17:25:10	tcp	secure1.comned.com	https	x.y.246.7	1536
Mar 9 17:27:58	tcp	212.108.4.40	http	x.y.246.7	1536
Mar 9 17:27:59	tcp	212.108.4.63	http	x.y.246.7	1536
Mar 9 17:28:03	tcp	212.108.4.63	http	x.y.246.7	1536
Mar 9 17:28:09	tcp	212.108.4.63	http	x.y.246.7	1536
Mar 9 17:30:03	tcp	secure1.comned.com	https	x.y.246.7	1536
Mar 9 17:32:53	tcp	212.108.4.40	http	x.y.246.7	1536
Mar 9 17:34:57	tcp	secure1.comned.com	https	x.y.246.7	1536
Mar 9 17:37:47	tcp	212.108.4.40	http	x.y.246.7	1536
Mar 9 17:39:52	tcp	160.79.30.211	https	x.y.246.7	1536
Mar 9 17:39:52	tcp	160.79.30.221	http	x.y.246.7	1536
Mar 9 17:39:52	tcp	212.108.4.234	8080	x.y.246.7	1536
Mar 9 17:39:52	tcp	212.108.4.132	https	x.y.246.7	1536
Mar 9 17:39:52	tcp	secure1.comned.com	https	x.y.246.7	1536
Mar 9 17:39:54	tcp	160.79.30.211	https	x.y.246.7	1536
Mar 9 17:39:55	tcp	212.108.4.132	https	x.y.246.7	1536
Mar 9 17:39:55	tcp	212.108.4.234	8080	x.y.246.7	1536
Mar 9 17:39:55	tcp	160.79.30.221	http	x.y.246.7	1536
Mar 9 17:40:01	tcp	160.79.30.211	https	x.y.246.7	1536
Mar 9 17:40:01	tcp	160.79.30.221	http	x.y.246.7	1536
Mar 9 17:40:13	tcp	160.79.30.211	https	x.y.246.7	1536
Mar 9 17:40:13	tcp	212.108.4.132	https	x.y.246.7	1536
Mar 9 17:40:14	tcp	212.108.4.234	8080	x.y.246.7	1536
Mar 9 17:40:14	tcp	160.79.30.221	http	x.y.246.7	1536
Mar 9 17:40:38	tcp	212.108.4.234	8080	x.y.246.7	1536
Mar 9 17:41:27	tcp	212.108.4.234	8080	x.y.246.7	1536
Mar 9 17:43:03	tcp	212.108.4.234	8080	x.y.246.7	1536
Mar 9 17:45:04	tcp	212.108.4.234	8080	x.y.246.7	1536

The address x.y.246.7 does not exist on my network nor does the subnet 246. This was collected by the firewall, so it made it through the router, which means the ACK bit was set. I believe this intruder sourced himself as x.y.246.7 and then hit these three networks. Hard to call this a denial of service attack since the times are so long between hit, sometimes up to 4 minutes. The destination port stays the same, so he must be using some kind of utility that allows him to decide what address and what source port to use. He did this for 35 minutes.

Evidence of active targeting: Yes. Assuming this is a denial of service, he was actively targeting my whole network.

Existence: The three networks he targeted in order to target us were secur1.comned.com, address is 194.151.209.16 and the network is Comned Network in Amsterdam. The second network 160.79.30.0 is also Comned Network and the 212.108.4 is American Global Network in Mississippi.

History: It is impossible to tell if this has happened before since he is spoofing an address on my network.

Severity: Low. None of these packets made it past the firewall although they did make it past our router.

Detect 5 - Misconfiguration

Apr 7 03:41:23	tcp	apache.our.domain	netbio	x.y.44.99	3052
Apr 7 03:41:29	tcp	apache.our.domain	netbio	x.y.44.99	3052
Apr 7 03:41:41	tcp	apache.our.domain	netbio	x.y.44.99	3052
Apr 7 03:53:59	tcp	apache.our.domain	netbio	x.y.44.99	3054
Apr 7 03:54:02	tcp	apache.our.domain	netbio	x.y.44.99	3054
Apr 7 03:54:08	tcp	apache.our.domain	netbio	x.y.44.99	3054
Apr 7 03:54:20	tcp	apache.our.domain	netbio	x.y.44.99	3054
Apr 7 04:19:47	tcp	apache.our.domain	netbio	x.y.44.99	3059
Apr 7 04:19:50	tcp	apache.our.domain	netbio	x.y.44.99	3059
Apr 7 04:19:56	tcp	apache.our.domain	netbio	x.y.44.99	3059

The address x.y.44.99 does not exist and neither does the subnet 44. These packets were seen on the internal side of the firewall and a sniffer trace showed that this was the sync-ack packet of the tcp handshake. My first thought was a possible spoofed address. I did a sniffer trace on the internal side of the firewall and was able to see that the source was also inside our firewall. Further sniffer traces of specific ports on our internal switch allowed me to narrow this down to one particular port. The device plugged into this port was a UPS which had a connection to the network. Apache is our Primary Domain Controller and apparently, the UPS was attempting to reach this PDC. This was constant all day and all night. After speaking to members of my team, I discovered that when the UPS was installed and it asked for an address, the address was mistyped.

Evidence of active targeting: Yes. The target in this case was our PDC.

Existence. The “intruder” existed on our own network.

History. This began apparently when the UPS was installed and never stopped until we unplugged it from our network and then reconfigured it.

Severity: Low, since this was only a misconfiguration, but it sure had me stumped and assuming maliciousness when I first saw it.

Detect 6 – Ring0

```
Mar 31 14:29:43 buster1 14:29:42 164.138.233.50 3660 Lyon-19-50.abo.wanadoo.fr x.y.220.164 80 No_DNS_Name S www-http tcp
Mar 31 14:29:48 buster1 14:29:48 164.138.233.50 3660 Lyon-19-50.abo.wanadoo.fr x.y.220.164 80 No_DNS_Name S www-http tcp
Mar 31 14:30:00 buster1 14:30:00 164.138.233.50 3660 Lyon-19-50.abo.wanadoo.fr x.y.220.164 80 No_DNS_Name S www-http tcp
Mar 31 14:30:24 buster1 14:30:24 164.138.233.50 3720 Lyon-19-50.abo.wanadoo.fr x.y.220.164 8080 No_DNS_Name S 8080 tcp
Mar 31 14:30:27 buster1 14:30:27 164.138.233.50 3720 Lyon-19-50.abo.wanadoo.fr x.y.220.164 8080 No_DNS_Name S 8080 tcp
Mar 31 14:30:33 buster1 14:30:33 164.138.233.50 3720 Lyon-19-50.abo.wanadoo.fr x.y.220.164 8080 No_DNS_Name S 8080 tcp
Mar 31 14:30:45 buster1 14:30:45 164.138.233.50 3720 Lyon-19-50.abo.wanadoo.fr x.y.220.164 8080 No_DNS_Name S 8080 tcp
Mar 31 14:31:09 buster1 14:31:09 164.138.233.50 3764 Lyon-19-50.abo.wanadoo.fr x.y.220.164 3128 No_DNS_Name S 3128 tcp
Mar 31 14:31:12 buster1 14:31:12 164.138.233.50 3764 Lyon-19-50.abo.wanadoo.fr x.y.220.164 3128 No_DNS_Name S 3128 tcp
Mar 31 14:31:18 buster1 14:31:18 164.138.233.50 3764 Lyon-19-50.abo.wanadoo.fr x.y.220.164 3128 No_DNS_Name S 3128 tcp
Mar 31 14:31:30 buster1 14:31:30 164.138.233.50 3764 Lyon-19-50.abo.wanadoo.fr x.y.220.164 3128 No_DNS_Name S 3128 tcp
```

I see these scans all day long. Netscape proxy server uses port 8080, but we change most of our default ports to make these less likely to succeed. We do not block these anymore since there are so many all day long, but port 8080 and port 3128 cannot make it past our router with this sync packets. Port 80 is allowed only to legitimate web servers and these are limited to only certain networks and domains.

Evidence of active targeting: Yes. They were targeting our network, even though the address was a non-existent address. They were searching for a web server or proxy server.

Existence: This IP address belongs to France telecom interactive and they own the whole 164.138.

History: We have seen other port scans from this network but nothing that was successful. We have blocked some parts of this network.

Severity: Low. This would not make it through our router. Only a port 80 packet to a real web server would make it, but this domain could never make it to that server. It would be stopped at our firewall for not being a valid source allowed through to that web server. We have no web servers on 95 or 98.

Detect 7 – Attempt to pull zones

```
Apr 8 22:20:03 209.78.48.225 2961 p33-max4.wlg.ihug.co.nz x.y.50.1 53 dns1.our.domain S domain tcp
Apr 8 22:20:08 209.78.48.225 2961 p33-max4.wlg.ihug.co.nz x.y.50.1 53 dns1.our.domain S domain tcp
Apr 8 22:20:11 209.78.48.225 2961 p33-max4.wlg.ihug.co.nz x.y.50.1 53 dns1.our.domain S domain tcp
Apr 8 22:20:23 209.78.48.225 2961 p33-max4.wlg.ihug.co.nz x.y.50.1 53 dns1.our.domain S domain tcp
```

Evidence of active targeting: Yes. The intruder is targeting a dns server and attempting to pull zones.

Existence: This is a New Zealand domain. The IP address is 209.78.48.225 which is on a network owned by The Internet Group out of Napa, California.

History: I have seen this network doing port scans and ip scans but none successful.

Severity: Low. Our dns server zone transfers are limited down to secondaries in multiple locations.

Detect 8 - IP Scan

Jan 28 17:04:11	tcp	race.swip.net	6667	x.y.232.109	39426
Jan 28 17:04:15	tcp	race.swip.net	6667	x.y.232.109	39426
Jan 28 17:05:11	tcp	race.swip.net	6667	x.y.216.68	39426
Jan 28 17:05:14	tcp	race.swip.net	6667	x.y.216.68	39426
Jan 28 17:06:58	tcp	race.swip.net	6667	x.y.93.30	39426
Jan 28 17:07:01	tcp	race.swip.net	6667	x.y.93.30	39426
Jan 28 17:18:40	tcp	race.swip.net	6667	x.y.27.79	39426
Jan 28 17:18:45	tcp	race.swip.net	6667	x.y.27.79	39426
Jan 28 17:22:14	tcp	race.swip.net	6667	x.y.235.108	39426
Jan 28 17:22:19	tcp	race.swip.net	6667	x.y.235.108	39426
Jan 28 17:34:08	tcp	race.swip.net	6667	x.y.102.0	39426
Jan 28 17:34:12	tcp	race.swip.net	6667	x.y.102.0	39426
Jan 28 17:34:18	tcp	race.swip.net	6667	x.y.102.0	39426

This was seen on our firewall. They are scanning our network for valid IP addresses and the ack bit is obviously set because they got through our router. They scan each address twice and then wait quite a while before scanning the next address. I would assume we are not the only targets.

Evidence of active targeting: Yes. Our network and others were probably the targets.

Existence: The hostname rac.swip.net resolves to 130.244.126.146 which is Swipnet out of Sweden.

History: The only additional packets I have seen from this Class B Network are 2 udp/53 packets to our dns servers.

Severity: Low. They never made it past our firewall.

Detect 9 - ?

Apr 11 01:23:23	tcp	0.1.2.3.4.5.6.AM	54566	x.y.63.178	10747
-----------------	-----	------------------	-------	------------	-------

This packet was detected on the firewall and was addressed to a non-existent address on our network. The hostname caught my eye. This packet had to have the ack bit set to get through our router so it was either a spoofed address x.y.63.178 sent to 0.1.2.3 which forced 0.1.2.3 to send a syn ack packet to my network. Or the hacker sent a packet to our network and spoofed his source address in order to send resets back to 0.1.2.3.4.5.6.AM. Or he could simply have been sending sync-ack packets to various random addresses. It's impossible to tell with only this one packet.

Evidence of active targeting: Yes, he was probably targeting multiple networks including my own.

Evidence: This hostname will not resolve, but the domain 6.AM is using two dns servers on the kuwait.net domain. This domain is on the network 209.85.227 which is owned by KuwaitNet out of Montreal, Quebec, Canada.

History: I have not seen this hostname or domain before. I would have remembered this one.

Severity: Low. Even though this packet made it through our router, it did not make it past the firewall and the intruder was not aware of our network topology since he scanned an invalid subnet.

Detect 10 – New application?

Apr 3 17:10:53	udp s00c04fa35d4c.domain1	1026	x.y.63.255	41524
Apr 4 15:50:02	udp s00c04f58b627.domain1	1025	x.y.63.255	41524
Apr 4 16:57:08	udp s00c04fa35d4c.domain2	1026	x.y.63.255	41524
Apr 5 17:05:28	udp s00c04fa35d4c.domain2	1026	x.y.63.255	41524
Apr 7 08:25:03	udp s00c04fa35d4c.domain2	1026	x.y.63.255	41524
Apr 8 00:25:23	udp a.b.202.161	3218	x.y.23.255	41524
Apr 8 00:39:13	udp a.b.202.161	3239	x.y.23.255	41524
Apr 10 00:33:21	udp a.b.202.161	3578	x.y.23.255	41524
Apr 10 09:04:00	udp AFCIS.domain3	1031	x.y.22.255	41524
Apr 10 09:11:35	udp AFCIS.domain3	1034	x.y.22.255	41524

I started seeing these packets hit the firewall from three different trusted networks. They are always a broadcast. Two of the broadcast are to subnets that do exist and one is to a non-existent subnet. All are external servers running SCO Unix. My feeling is that an application was installed on multiple networks and this is in the default configuration. There must have been a time when subnet 63 existed and this application was configured to talk back to our network. On the other hand, it could be three boxes that were hacked, but the fact the two of the subnets are real makes me think otherwise. I work for a large organization that runs many servers that these trusted networks talk to. I have contacted all three networks and am still in the process of working this with them. The times seem to indicate beginnings and ends of shifts as if these hosts are being turned on at these

times and they broadcast to find whatever it is they are looking for on our network. At the same time I also blocked all broadcasts on our external router so these will no longer show up on my firewall.

Evidence of active targeting: Yes. They are targeting two legitimate subnets on my network.

Existence: These are all three trusted networks that we deal with and expect to see traffic from.

History: This has been going on for over a month. I do not otherwise see malicious traffic from these networks.

Severity: Low since these are trusted networks but I am hoping to resolve this issue and find out what is causing these packets.

© SANS Institute 2000 - 2002, Author retains full rights