Global Information Assurance Certification Paper

SANS GCIH PRACTICAL
By Bill J. Reed
12 Feb, 2001

Executive Summary

In the winter of 2000, a Department of Defense (DoD) Agency's Incident Handling Team was notified by the Defense Criminal Investigative Service (DCIS) of a potential compromise of a server registered under their domain. Above and beyond the potential compromise was the fact that the IP address associated with this system was suspected of attacking another system, a certain company's server. File systems were erased on the company's server by the attacker and the system was actually down for over one month. This system being down resulted in a substantial loss of revenue for the company.

The DoD Agency's Incident Handling Team recognized their IP address as belonging to a firewall that was protecting one of its networks. There was no easy way to disconnect the firewall from the network and perform the investigation. Obviously, the firewall is a production security system that must remain connected and operational. Since the Incident Handling Team was not "on site" where this firewall was, they decided to access it remotely using a secure encrypted/strong authenticated method. This is the same day-to-day method they use to access all remote firewalls.

Upon accessing the firewall, the lead Incident Handler started looking through the firewall logs for the IP address of the attacked/compromised system at the company. It was soon discovered that certain source IP addresses were coming in to the firewall on the external interface and going back out the same interface to the attacked company's server, as well as to a few other locations around the world. As soon as this was discovered, all of these types of connections were copied from the logs and written to CD.

After analyzing the logs some more and looking over the configuration of this firewall, the lead Incident Handler realized the problem. It was not that the firewall had been compromised per se but it was being used in a manner that was not intended. The firewall administrator had created a rule using a sort of a shortcut feature that was present to this particular firewall vendor. Basically, a rule was created that allowed the external interface of the firewall to connect to the Universe and the telnet, ftp, and http proxies were part of this rule. With this particular firewall vendor, if an interface is either the source or destination defined in a rule, the interface takes on whatever networks are defined to be legal behind it. In this case, the firewall administrator was likely trying to allow proxied web, ftp, and telnet connections out from his firewall but instead configured the rule to allow outsiders to use the proxies to re-connect back out. The lead Incident Handler fixed this rule right away and soon after, it was verified that outsiders were no longer able to use the firewall proxies for connecting back out.

The firewall logs connected to this incident were handed over to the DCIS Agent and were entered in to his evidence system. He continued to investigate this incident and soon traced the attacker to a single Internet Service Provider (ISP). After that, he was able to trace it all the way back to the attacker's house. At this point, a search warrant was obtained and they went in for the search. More than enough evidence was collected at the house. The attacker pleaded guilty about one year after the original incident was reported and was convicted.

Next page includes a diagram (Fig. 1 The Attack Process) showing how the attacker used our firewall to access/attack other's systems, compromising one company's server. A brief explanation of the diagram follows:

1. The attacker dials up an ISP using an unauthorized/stolen account.

2. Once connected in to the ISP, the attacker uses telnet to connect to the telnet proxy on the DoD Agency firewall. Ftp and http were used as well.

3. Since the firewall is misconfigured, the attacker is able to telnet back out. This enables the attacker to mask his real IP address. The attacked then believe it is the DoD Agency firewall that is attacking them. In this case, the attacker root compromised a company's server and deleted files. That drastic move is what started the investigation.

4. This step shows that the attacker and a few others did access other servers around the Internet as well by first going through the misconfigured firewall. Only the attacker specified in this report was ever convicted of any wrongdoing.

The DoD Agency's network behind the firewall was never accessed in any way by the attacker. All attacker connections were to the firewall and then back out.

Fig. 1 The Attack Process

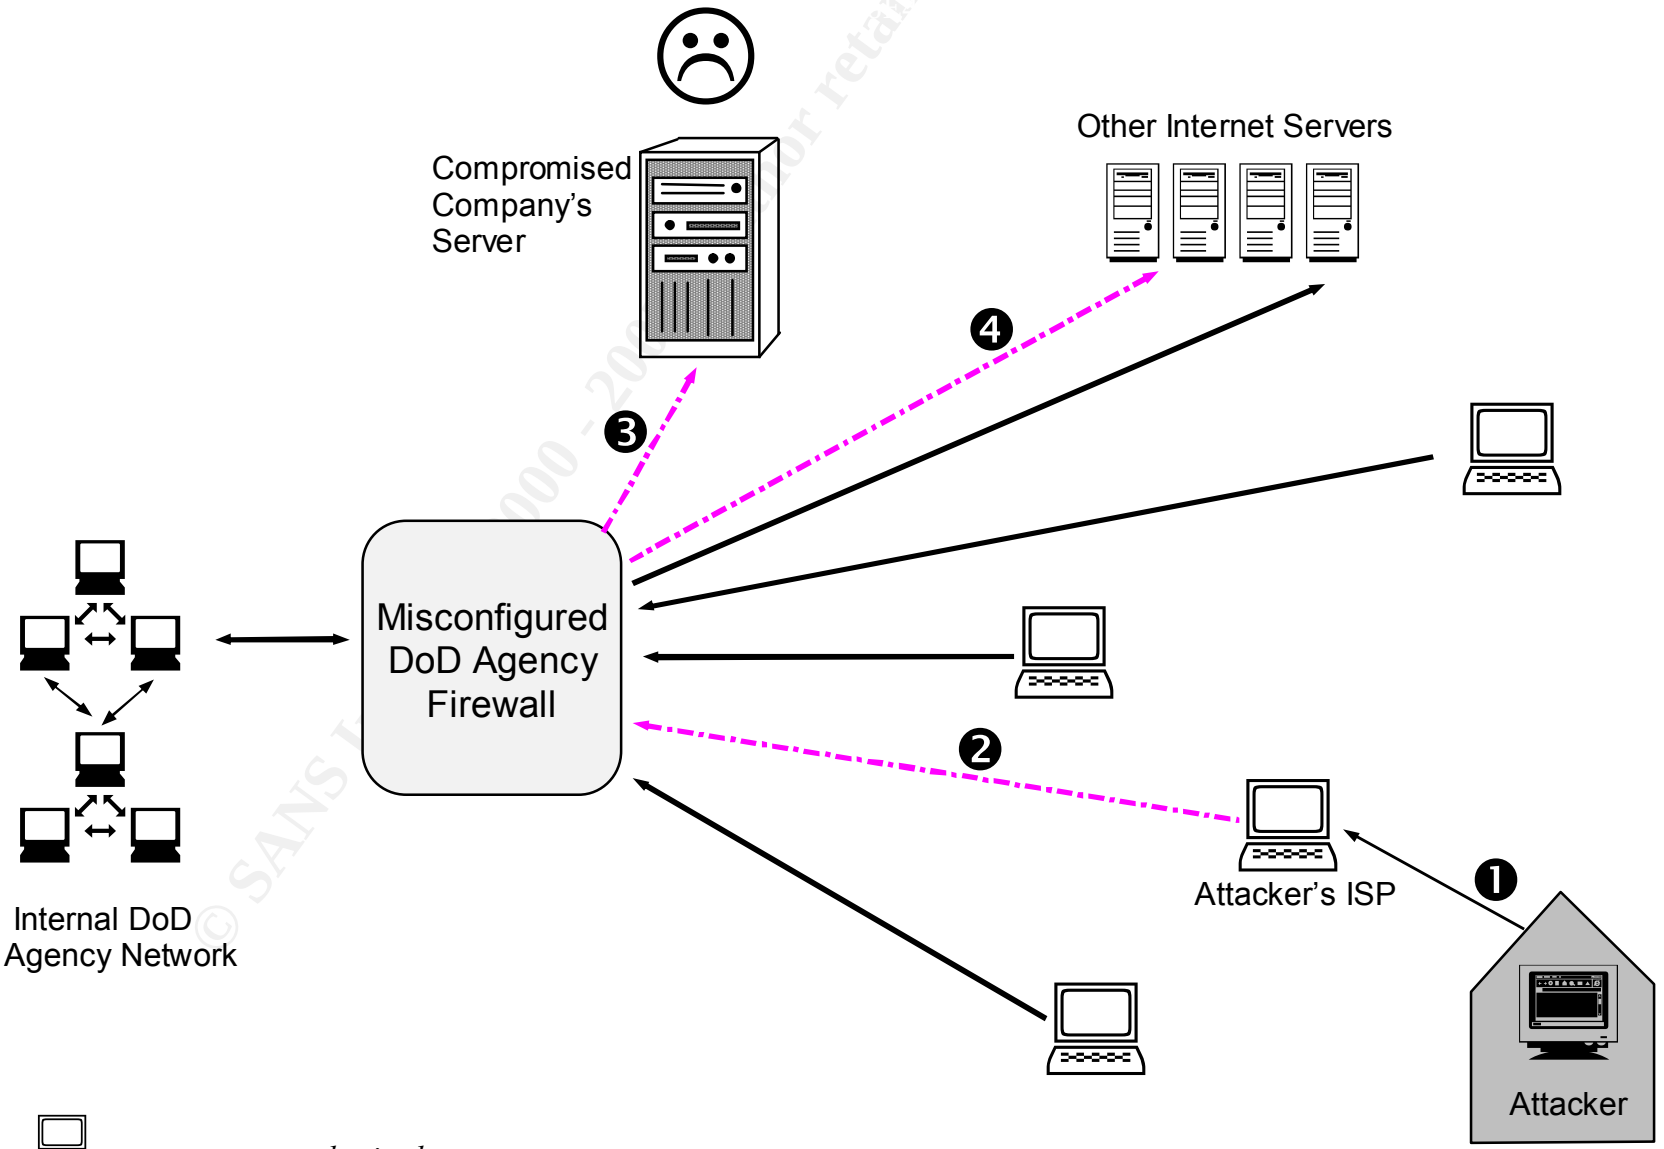

Note: *represents unauthorized user*

Introduction

This incident will be explained using the six stages of Incident Handling taught by SANS. They are Preparation, Identification, Containment, Eradication, Recovery, and Follow Up/Lessons Learned. Please note that I did not handle this particular incident myself but a former co-worker did. I have interviewed my former co-worker and also a representative from the Defense Criminal Investigative Service, DCIS. Both of these individuals were paramount in the successful handling of this incident. I say "successful handling" not only because they helped stop the attacker/intruder and kept him from doing more damage to others, but also because they collected the evidence that helped in his prosecution.

Preparation

Our particular DoD Agency has established it's own Incident Handling Team. This team has high level support within the agency. We also have a certain level of support within DoD. We have strict procedures in place on how to properly handle incidents and report them. The Incident Handling Team members are trained and knowledgeable in different areas. Since our agency is small compared to other parts of DoD, this team not only handles incidents but also fields Information Assurance (IA) initiatives for the agency. We install and support Intrusion Detection Systems (IDS), firewalls, encryption software, and other security solutions. We send out advisories and other information, such as OS hardening tips, router Access Control Lists, and all sorts of ideas/solutions that will improve security in the agency. We developed our own perl and shell scripts to parse the agency's firewall logs and detect scans/attacks. The script results are manually looked over by Incident Handling Team members on a daily basis. Certain detects are acted upon based on our Incident Handling Procedures.

Part of our Incident Handling Procedures is to always notify DCIS of a suspected intrusion. This is true even if the system suspected of being compromised is not our own but it is another DoD entity's. This part of our Incident Handling Process helps to keep us communicating with DCIS. Everyone on the team is well aware of this and because of it, we have a good working relationship with DCIS. We help each other by sharing. We also have similar procedures in place for reporting incidents to DoD CERT[1] (DoD's Computer Emergency Response Team) and to our own HQ.

For this particular incident, one of the key preparation procedures that helped us was that of firewall log backups. We developed a secure method that backs up all firewalls and encrypts the logs. We then receive the logs and store them both online and offline. We keep firewall logs forever offline. Online on a central server, we keep them for as long as we have space, generally a few months. We also keep the logs on the firewall itself for as long as we have space. Depending on the firewall, this could range from a month to over a year. This entire process is automated except for parts of the storing offline.

Also worth noting for this particular incident, is the fact that our Incident Handling Team maintains secure remote access to all agency firewalls. We periodically check this access to make sure we can still get in. That way when emergencies come up, we do not have to depend on the local firewall administrators at each site. We can immediately access the firewall in question ourselves and investigate. We also have access to all IDS systems in the agency. It is another means to investigate incidents.

Our agency has security policies in place. There is a requirement that approved warning banners be placed on all systems, including desktops and web pages. Of course, even with this policy in place, we often find systems without warning banners. Unfortunately, our firewall used by the attacker in this case did not have warning banners setup on the proxies. Luckily, it didn't affect the turnout of the case but we did install warning banners on telnet and ftp firewall proxies after this. Other common sense security policies are in place as well, such as a password policy and firewall policy.

Employees in the agency are taught to report possible security incidents. Each location in the agency has an individual or more responsible for security and they know to report everything to us, the Incident Handling Team. Every so often, we have training sessions, where we teach system and network administrators about security, in particular how to recognize suspicious events that could be incidents. We also meet with the security officers at each site and train them. We usually invite them all to a central location and each member of the Incident Handling team gives a briefing on a particular subject. We try to do this once a year.

We have a secure web page on our Intranet that all security officers and firewall administrators in the agency have access to. On this page, we list all the information that we can that will help them. For example, all of our advisories are kept on this page. Also all of the recommended firewall blocks of "HOTLIST" systems or networks are listed. These are systems or networks that are actively scanning or attacking our systems. We also communicate regularly with site security officers via PGP encrypted email, or sometimes via secure telephone and/or secure fax.

We often do site reviews in our agency. We analyze a site's security and then write up a report. A time period is given in which the site is to address all of the issues. We also sometimes perform "Red Teaming" exercises. This is where one team actively attacks the network while the others defend and detect. This is very good training and preparation for defending against real attacks. It is a very good test of our current Incident Handling Procedures.

All of the methods and procedures mentioned above are many of the ways in which our Incident Handling Team better prepares for handling incidents. As we confront new types of attacks, deal with new circumstances, lose or gain employees, lose or gain customers, learn of new ideas through training such as SANS, and undergo priority changes for our group, some of our methods and procedures will change. Hopefully these changes will be for the better and improve our Incident Handling performance.

Generally speaking, this has been the case for our team. We are learning through experience and recognize the fact that preparation is the key to success here.

Identification

Our Incident Handling Team learned of this incident after being contacted by the local DCIS agent. The DCIS agent learned of this incident after a federal prosecutor notified the DoD CERT. The federal prosecutor was involved because a company's server had been compromised and taken down. The source IP address of the attack was an IP address registered to our agency.

This event pretty much came in to us as an incident, so we really didn't have to prove that it was one. The facts at this point were that a company's server had been root compromised and files were deleted. The server was now offline and the source of the attack was stated as being one of our IP addresses. My manager assigned this incident to our lead Incident Handler. He immediately recognized the IP address as belonging to one of our firewalls.

Obviously, a firewall is not a device that you can easily take off the network. If we had to do that, we could have, but the lead Incident Handler decided to first login to the firewall in question to investigate. This is something that our office does all the time anyway, since we support firewalls in our agency. So really it would be a normal event for us. We have a secure encrypted and strong authenticated method of accessing firewalls remotely as we did this one. With these factors in mind and also the fact that our Incident Handling Team is not really budgeted or setup to travel to the scene of the crime every time, the decision was easily made to login remotely.

Once on the firewall, the lead Incident Handler immediately searched the logs for the attacked Company's IP address. It was discovered in the logs but it was never just the firewall IP address going to the attacked IP address. There was always another IP address connecting in to our firewall first and then going back out, taking on the identity of our firewall. At this point, the lead Incident Handler knew that we had a configuration issue on this particular firewall.

At this point, all the logs involving connections coming in to our firewall and then connecting back out were written to CD. This evidence was turned over to the DCIS agent, who then entered it into his evidence system. The lead Incident Handler produced an initial report based on his notes thus far. This report was sent to management, the DoD CERT, and DCIS. The rest of the Incident Handling team was also aware of what was going on.

Containment

As learned in the Identification phase above, we knew there was a configuration issue on the firewall. The lead Incident Handler started to check the rules and immediately discovered the problem. The local firewall administrator evidently was confused on how to correctly create what this particular firewall vendor calls an "interface based rule". He had created a rule that inadvertently allowed telnet, http, and ftp proxy connections from the outside. When I say the outside, I mean the Internet. The connections were denied for inbound access through the firewall but the rule would allow the outsider to turn around and connect somewhere else on the outside. So in this way, our firewall could be used to mask the identity of the real source with an attacker going on to his or her destination, looking as if they are our firewall.

With this firewall vendor, if an interface is listed as a source or destination in a rule, it takes on the legal networks that are behind the interface. So the external interface is really the Internet. The internal interface is really the legal internal address space. The firewall administrator should have configured the rule with the internal interface as the source and the Universe as the destination. Instead, he configured the rule with the external interface as the source and the Universe as the destination. This is what allowed the connections in and right back out. Evidently, the correct rule was also in place or users would have been complaining about no web services.

The lead Incident Handler immediately removed the incorrect rule and reconfigured the policy on this firewall. Now our recommendation to all sites running this particular brand of firewall is to not use the interface based rules at all but to use specifically defined entities and/or groups of entities instead. This way there can be no confusion over the matter. It also lines up with the way we configure the other firewall vendors used in our agency.

Please see the text below for a run-down of how the lead Incident Handler detected the firewall misconfiguration and how he corrected it. The tools used to assess and contain this incident were our secure remote login feature and the GUI interface of this particular firewall vendor.

Assess and Contain Process

Networks explained:

10.1.0.0/16 - Attacker's network
10.2.0.0/16 - Attacked company's network
10.3.0.0/16 - DoD Agency firewall's external interface
10.4.0.0/16 - DoD Agency firewall's internal interface
10.5.0.0/16 - DoD Agency firewall's internal interface
10.6.0.0/16 - DoD Agency firewall's DMZ interface

From our DoD Agency's CERT HQ, the lead Incident Handler did a secure remote login to the firewall in question. Once on the firewall, he used the *grep* command to look for the attacked's IP address (10.2.1.1) in the firewall log files.

```
# grep 10.2.1.1 logfile > /tmp/incident-logs.1
```

Many entries from /tmp/incident-logs.1 looked like this:

Jan 15 09:08:34 firewall telnetd[1024]: 121 Statistics: interval=40.97 id=abcde sent=26 rcvd=80 **srcif=if1** src=10.1.1.1/3456 srcname=attackers-.ISP.com **dstif=if1** dst=10.2.1.1/23 dstname=attacked-comapny.com proto=telnet rule=10

The lead Incident Handler noticed that the firewall's source interface (if1) was the same as the destination interface (if1). This is not normal. Usually, a source interface will be either the internal or external interface of the firewall and the destination will be the opposite of whatever the source was. The logs have immediately shown the lead Incident Handler that the firewall must be misconfigured.

Next, the lead Incident Handler verified the network configuration of this firewall. This was so he could get a mental picture of which interfaces were internal and which external or DMZ. In order to associate IP addresses to interface names, the *ifconfig* command was issued on the firewall.

```
# ifconfig -a

lo0: flags=849<UP,LOOPBACK,RUNNING,MULTICAST> mtu 8232
        inet 127.0.0.1 netmask ff000000
if1: flags=863<UP,BROADCAST,NOTRAILERS,RUNNING,MULTICAST> mtu 1500
        inet 10.3.1.1 netmask ffff0000 broadcast 10.3.255.255
if2: flags=863<UP,BROADCAST,NOTRAILERS,RUNNING,MULTICAST> mtu 1500
        inet 10.4.1.1 netmask ffffff00 broadcast 10.4.255.255
if3: flags=863<UP,BROADCAST,NOTRAILERS,RUNNING,MULTICAST> mtu 1500
        inet 10.5.1.1 netmask ffff0000 broadcast 10.5.255.255
if4: flags=863<UP,BROADCAST,NOTRAILERS,RUNNING,MULTICAST> mtu 1500
        inet 10.6.1.1 netmask ffffff00 broadcast 10.6.255.255
```

The lead Incident Handler then accessed the firewall in question using the remote management software or GUI. What he knew up to this point was that a rule dealing with telnet was allowing connections in and then allowing a connection to be made back out. These connections were using the external interface of if1.

A quick browse through the rules revealed a rule in place with the source listed as interface if1 and the destination was listed as Universe. Telnet, ftp, and http were part of this rule and logging was turned on thank goodness. Evidently, this rule was an administrator's mistake. The lead Incident Handler deleted this rule and re-installed firewall policy.

Rule ID	Source	Destination	Proxies	Logging
10	if1	Universe	telnet, ftp, http	yes

There was another rule that allowed these same services outbound. This rule was correct as it only allowed internal users to use the proxies going outbound to the Internet. This rule was left in place at this time.

Rule ID	Source	Destination	Proxies	Logging
17	if2	Universe	telnet, ftp, http	yes

There was another rule since this firewall had multiple internal interfaces. It was left alone as well.

Rule ID	Source	Destination	Proxies	Logging
18	if3	Universe	telnet, ftp, http	yes

Our very strong recommendation after dealing with this particular incident was to NOT use interface based rules. The alternative would be something like this.

Rule ID	Source	Destination	Proxies	Logging
18	Internal-nets	Universe	telnet, ftp, http	yes

, where Internal-nets is a group containing all of the legal internal networks at this particular site. It is much easier to understand without having to know all of the interfaces on the firewall.

Better yet would be to drop the telnet proxy from the "outbound to Universe" rule as the telnet service is not really a high usage public service these days. We all know that the big risk is that the account id, password, and data is sent "in-the-clear". We now limit telnet usage in the agency and VPN, or replace it with ssh.

So this was how the lead Incident Handler detected the problem and fixed it, thus putting a stop on the attacker using us to attack others. The only problem was that we were a bit too late. The company had already been root compromised and taken down.

Backup Procedures

With this particular incident, since it was not an intrusion, we didn't back up the system. We do backup our firewall logs on a regular basis so we didn't need to do that because of this incident per say. It was already done automatically. I will explain that process now, including hardware and software used.

1. The firewall platform in this particular case is a popular flavor of Unix. Cron runs a script installed by the firewall vendor, which basically rotates the logs at midnight. Within the vendor script, we have added a call to our own script.

2. Our script will take the old log file and PGP encrypt it using the private key of the firewall and a public key that our Incident Handling Team controls. This log is then sent via sendmail to an account setup on one of our servers that we use for centrally storing and processing firewall logs.

3. Once the mail is received on our end, another script is launched which does some security checks to verify the sending machine, etc. At this point, the log is decrypted using PGP and the signature is verified.

4. If everything goes as normal, our incident detection scripts will then run against this log. A copy of the log is also stored on the system. Once the system disk space reaches a certain percent full, a notification is sent to us, which lets us know to write some logs to CD. We do that step manually every so often. This entire process is done for all firewalls in our agency of this vendor type and it has worked well for us.

Here is our script on the firewall that gets called nightly:

```
#!/bin/ksh
#
# this will mail a logfile to us
#
# Script provided by us
#
HOME=~root
MAILTO=accountname@somewhere.dod.mil
TMP=/var/tmp
PATH=/usr/local/bin:/usr/bin
export PATH TMP HOME

if [ -z "$1" ]; then
  echo "USAGE: log2us logfile"
  exit
 fi

logfile=$1

if [[ $lf = *.gz ]]; then
lf=`basename $logfile.gz`
  export G=1
else
lf=`basename $logfile`
fi
```

```
HN=`hostname`

# Get messages, compress, encrypt and mail
(echo FILE=$HN.$lf.txt.gz;
if [ "$G" = "1" ]; then
  cat $logfile
else
  gzip -c9 $logfile
fi
) | pgp +compress=off -safe "Account Name " account 2>/dev/null|\
mailx -s "$HN $lf" $MAILTO
```

The script(s) on the receiving end are too sensitive to list here. Besides, the code is very long and hard to follow. It would probably bore you to death.

Eradication

We already learned the cause of this incident in the Containment section above. The symptoms were obvious once we were notified of what was going on. The firewall log entries showed that the inbound and outbound interfaces were the same. This was easy to fix through the firewall management interface or GUI, which is what the lead Incident Handler did. So in this particular incident, the cause was easy to remove once it was known.

The Incident Handling Team did take some other actions.

1. A final report was sent to the DoD CERT, DCIS, and HQ.

2. Since other agency proxies may suffer the same type of misconfiguration, we put out an advisory indicating that all proxy servers (firewalls or typical proxy servers such as web) should be checked to make sure they wouldn't allow this kind of connection.

3. We verified the proper configuration of all the firewalls we had access to at the time.

4. We modified our incident detection scripts to check for an interface being the same for inbound communications as well as for outbound communications. If so, an Action incident is flagged for us to see and act upon.

Recovery

Once the lead Incident Handler deleted the bad firewall rule, we were easily able to verify that the firewall was good again. We simply connected in to the same proxies used before and tried to connect back out. The firewall was now blocking all of these types of connections. For a while, we continued to see the attacker and others trying to come back to the firewall. They were now blocked from using the firewall in this manner. We considered our self fully recovered.

The company attacked never did fully recover. They were a small business and this incident appears to have really hurt them as they are now out of business. They lost a significant amount of money after the attacker compromised their system and then deleted files on it. They were down for one month at the time.

Lessons Learned/Follow UP

There are numerous lessons learned after dealing with this particular incident.

1. It is easy to misconfigure software that will permit misuse. In our case, a device meant to increase security was itself misused against an innocent victim. We must be more careful on how we configure firewalls. More training for firewall administrators must be a priority.

2. A system can be misused for a quite a while without anyone even realizing it. In our case, the proxies had been misused for a couple of weeks before we even realized what was going on. Who knows how long it would have been before we noticed this if it weren't for the DCIS tip. We must find ways to better check logs, which usually means some level of manual check.

3. Automated log checking needs a great deal of attention and thought. We must realize that we can not catch everything but we can always add new signatures, or take a new approach. The idea I like is to detect the "not normal". Commercial products to do this function may have to be looked at as well.

4. Since this incident, we have learned to augment our infrastructure with proper placement of Intrusion Detection Systems (IDS). A properly configured IDS would have detected this particular incident on the very first telnet performed from the firewall.

5. DCIS mentioned to me that the attacked server's time was totally wrong. This made it very difficult for him to investigate this incident. We should always make sure that all of our systems have the proper time and point to an authorized timeserver.

6. DCIS also mentioned that a timely response is key. Everyone involved with an incident must act in a timely manner and communicate efficiently with everyone else. Unbelievably, the attacked company was the most difficult to get a hold of in this case.

7. Something else interesting to me in regards to how DCIS tracked down the attacker in our case. They used the comments of the attacker, as this information was available online[2] when he warred with another hacker in public forum. Evidently, the attacker said too much in public. They then tracked him through the web pages he defaced[3]. To make a long story short, everything pointed to one ISP and this is how they caught him. The lesson learned is that there is a ton of information to help us research an incident available online. We only have to look for it.

As far as follow up is concerned, we did go back and make sure all of our firewalls were configured correctly to protect against this type of attack. The advisory we put out to the field (mentioned above in the Eradication section) was followed up on. We also improved our incident detection scripts to capture this type of attack and other types. Above and beyond firewalls, Intrusion Detection Systems have been deployed across our agency. This is another tool for capturing incidents. Site reviews are starting to take place more often, in which we hope we will be able to recognize problems early.

Conclusion

This incident was unique in that we were not compromised but we did have unauthorized use. The attacker bounced off our firewall in order to mask his identity before attacking others. This was the first ever incident that our young Incident Handling Team was involved with where the attacker was physically caught and prosecuted. With all of our other incidents, the attackers did some level of damage and we detected it and cleaned up the mess but nobody was ever caught. This time we were able to help DCIS by supplying them with evidence (firewall logs) and also by letting them know where else the attacker was going. This information helped them to bring the attacker to justice.

This incident, as well as others, did help us improve our security infrastructure. We saw the need for IDSes across the agency and have deployed them. We also continue to improve our firewall configurations.

References

1. DoD CERT - www.cert.mil
2. Antionline - www.antionline.com
3. Attrition.org - www.attrition.org
4. The SANS Institute - Computer Security Incident Handling, Step By Step, version 1.5