



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
"Hacker Tools, Techniques, and Incident Handling (Security 504)"
at <http://www.giac.org/registration/gcih>

Looking at your firewall logs

William Farnsworth


September 3, 2000

Executive Summary

We could title this discussion: “A funny thing happened on the way to find the port scanners”... and be right on target. I recently came back from a security conference scared of all the black hats that might be attacking our firewall and even worse, penetrating our defenses... what was happening? Who was attacking? Was it a big deal? What was going on? All great questions that we should all ask ourselves.

But... what we found instead was most likely a piece of mis-configured software probably from our own company... but that's what this story is about... and how we followed the basic incident handling procedures just in case it was a real attack.

The basic plan was to find port scanners and other attempts to intrude into the company's rather large network organization. We are not the size of an e-Bay, or amazon.com, but we do have a world wide WAN and thousands of computers connected to our corporate network.



We host our own web servers and do work with other companies, so our presence on the web is not big... but we are visible.

We maintain a border router, DMZ and Axent's Raptor firewall as our primary defenses against attack. The Raptor system creates logs of daily activity and perhaps we can use these logs as the input data into a post-event analysis and possible intrusion detection of attacks.

Let's quickly review the standard firewall log analysis tools and look for a way to see individual data items or perhaps use “standard office tools” to look at raw data from the firewall log. This is a way to look for attacker signatures and try and get an overall perspective of “What was going on.”

Preparation

Firewall analysis tools

The first step was to look for some standard tool that will slice and dice the firewall logs to look for events of interest.

© 2000, William Farnsworth

Review of some of the available systems

A quick browse of the web looking for analysis tools indicates that they generally fall into one of two categories, either usage tools or intrusion detection systems. The *dmoz* "Open Directory" site lists 47 commercial log analysis tools [\(8\)](#). The listed tools appear to be concentrating on analysis of usage instead of intrusion detection.

What they do

These firewall analysis tools are primarily focused on two types of analysis... web usage and bandwidth usage and abuse.

Analysis of web usage is important if you host a web site and need to see how many hits particular pages get hit each day. The marketing department uses this kind of information to identify what customers are interested in; and what marketing approaches are not useful. An ISP or other application service provider might use this type of analysis for charging particular customers for the use of their web system.

Another use of this type of analysis is for an internal service organization to distribute charges for Internet access. Individual organizations identified by IP address can be charged for the amount of time they are connected through various ports.

For example, "Lance" asked a couple of years ago in the Firewall Wizards mailing list... *"Specifically I'm looking for the ability to "sniff" out bandwidth hogs, and visitors to questionable sites. I'd also like the ability to generate reports for any particular internal user (based on IP address)."* [\(3\)](#)

Although it may be hard to believe, some users *actually abuse their Internet access privilege in a typical corporate environment!* Some folks think that it is fine for them to view sites of questionable content including gambling, pornography and auctioning sites during business hours. Or how about listening to their favorite radio station from their hometown? Don't y'a just love Internet radio? Bandwidth hog.

But that is another discussion... we are not going to focus on establishing appropriate standards of employee use of corporate resources here.

In addition to the commercial firewall log analysis tools, there are a number of downloadable applications available for a minimal fee and sometimes for free.

For a list of analysis tools, Peter Davis and Associates [\(4\)](#) has a web page that lists many firewall analysis tools. FreshMeat.Net [\(5\)](#) includes a wide variety of log analysis tools including things as diverse as monitoring Quake logs to analyzing web traffic.

These analysis tools will perform the high level analysis of the logs... but they do not give us a view into the details of the packets that reach our firewalls. We can't see all the minute details of each packet. All we get to see is the detection of predictable patterns.

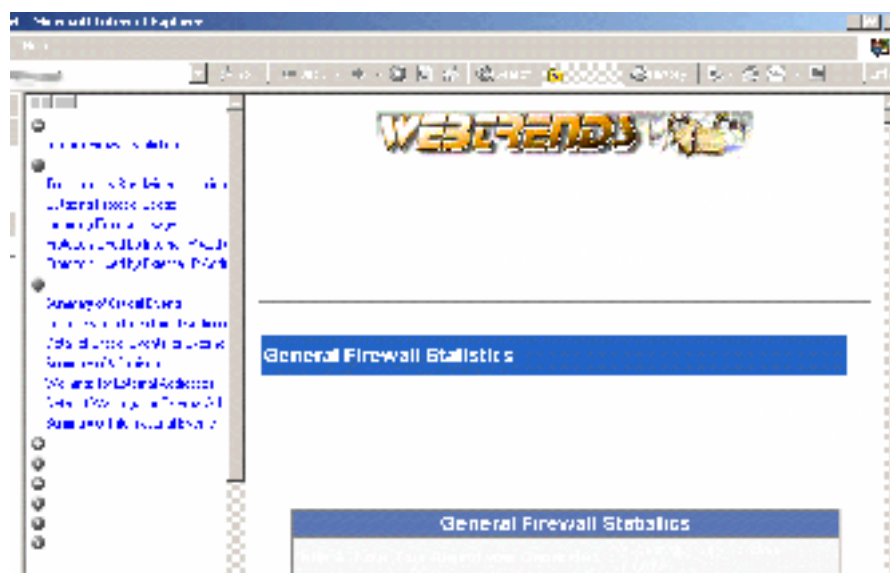
Looking at your firewall logs...

WebTrends is one of these products to monitor and review firewall logs. The results are typically generated each day and presented as web pages.

The standard reports give a good overall idea of what kinds of traffic flow through the firewall...

including web, email, FTP and telnet protocols are displayed in tables. In addition,

summary statistics of critical and warning events are included. If we look at the default listings, there is very little detail information about the exact causes of the warnings and critical events. PC Magazine recently had a complete review of WebTrends and comparable firewall log analysis tools [\(16\)](#).



Return Value	# of Events	% of Total
000 - Possible port scans	298561	100.0%
001 - Possible port scans	298561	100.0%
002 - Possible port scans	298561	100.0%
003 - Possible port scans	298561	100.0%
004 - Possible port scans	298561	100.0%
005 - Possible port scans	298561	100.0%
006 - Possible port scans	298561	100.0%
007 - Possible port scans	298561	100.0%
008 - Possible port scans	298561	100.0%
009 - Possible port scans	298561	100.0%
010 - Possible port scans	298561	100.0%
011 - Possible port scans	298561	100.0%
012 - Possible port scans	298561	100.0%
013 - Possible port scans	298561	100.0%
014 - Possible port scans	298561	100.0%
015 - Possible port scans	298561	100.0%
016 - Possible port scans	298561	100.0%
017 - Possible port scans	298561	100.0%
018 - Possible port scans	298561	100.0%
019 - Possible port scans	298561	100.0%
020 - Possible port scans	298561	100.0%
021 - Possible port scans	298561	100.0%
022 - Possible port scans	298561	100.0%
023 - Possible port scans	298561	100.0%
024 - Possible port scans	298561	100.0%
025 - Possible port scans	298561	100.0%
026 - Possible port scans	298561	100.0%
027 - Possible port scans	298561	100.0%
028 - Possible port scans	298561	100.0%
029 - Possible port scans	298561	100.0%
030 - Possible port scans	298561	100.0%
031 - Possible port scans	298561	100.0%
032 - Possible port scans	298561	100.0%
033 - Possible port scans	298561	100.0%
034 - Possible port scans	298561	100.0%
035 - Possible port scans	298561	100.0%
036 - Possible port scans	298561	100.0%
037 - Possible port scans	298561	100.0%
038 - Possible port scans	298561	100.0%
039 - Possible port scans	298561	100.0%
040 - Possible port scans	298561	100.0%
041 - Possible port scans	298561	100.0%
042 - Possible port scans	298561	100.0%
043 - Possible port scans	298561	100.0%
044 - Possible port scans	298561	100.0%
045 - Possible port scans	298561	100.0%
046 - Possible port scans	298561	100.0%
047 - Possible port scans	298561	100.0%
048 - Possible port scans	298561	100.0%
049 - Possible port scans	298561	100.0%
050 - Possible port scans	298561	100.0%
051 - Possible port scans	298561	100.0%
052 - Possible port scans	298561	100.0%
053 - Possible port scans	298561	100.0%
054 - Possible port scans	298561	100.0%
055 - Possible port scans	298561	100.0%
056 - Possible port scans	298561	100.0%
057 - Possible port scans	298561	100.0%
058 - Possible port scans	298561	100.0%
059 - Possible port scans	298561	100.0%
060 - Possible port scans	298561	100.0%
061 - Possible port scans	298561	100.0%
062 - Possible port scans	298561	100.0%
063 - Possible port scans	298561	100.0%
064 - Possible port scans	298561	100.0%
065 - Possible port scans	298561	100.0%
066 - Possible port scans	298561	100.0%
067 - Possible port scans	298561	100.0%
068 - Possible port scans	298561	100.0%
069 - Possible port scans	298561	100.0%
070 - Possible port scans	298561	100.0%
071 - Possible port scans	298561	100.0%
072 - Possible port scans	298561	100.0%
073 - Possible port scans	298561	100.0%
074 - Possible port scans	298561	100.0%
075 - Possible port scans	298561	100.0%
076 - Possible port scans	298561	100.0%
077 - Possible port scans	298561	100.0%
078 - Possible port scans	298561	100.0%
079 - Possible port scans	298561	100.0%
080 - Possible port scans	298561	100.0%
081 - Possible port scans	298561	100.0%
082 - Possible port scans	298561	100.0%
083 - Possible port scans	298561	100.0%
084 - Possible port scans	298561	100.0%
085 - Possible port scans	298561	100.0%
086 - Possible port scans	298561	100.0%
087 - Possible port scans	298561	100.0%
088 - Possible port scans	298561	100.0%
089 - Possible port scans	298561	100.0%
090 - Possible port scans	298561	100.0%
091 - Possible port scans	298561	100.0%
092 - Possible port scans	298561	100.0%
093 - Possible port scans	298561	100.0%
094 - Possible port scans	298561	100.0%
095 - Possible port scans	298561	100.0%
096 - Possible port scans	298561	100.0%
097 - Possible port scans	298561	100.0%
098 - Possible port scans	298561	100.0%
099 - Possible port scans	298561	100.0%
100 - Possible port scans	298561	100.0%
Total number of Warnings for log file	298561	100.0%

If we look at the warnings category, the largest number of events is "Possible port scans"... Ah ha! We knew that there was something going on!

Intrusion detection systems

Lance Spitzner has a very interesting overview introduction to Intrusion Detection Systems [\(6\)](#). Much of his discussion is about a perl script for analysis of Checkpoint FireWall-1 logs [\(17\)](#). This script looks for probes on ports that are not normally active.

Logsurfer [\(7\)](#) by Wolfgang Ley and Uwe Ellerman is a similar Perl based script tool that will analyze your logs looking for surreptitious behavior and announce it to you.

What they don't

But these firewall log-monitoring tools don't let you get a good handle on the exact details of everything that is actually happening on the outside surface of your firewall. They aren't full blown intrusion detection systems. With both these sets of tools, you get pre-digested information that another person has determined meets your needs. You really can't get a good feeling about what the data actually looks like.

Now, to be fair, this is a difficult job. If you want to see that Johnny is hogging the bandwidth of your firewall listening to his hometown radio station, or looking at the latest Braves scores when he should have been working, these tools will help.


Scientists will frequently just look at the raw data to assimilate some sort of patterns or anomalies. Clifford Stoll did this in The Cuckoo's Egg (10). He used a printout to look for strange patterns of computer usage in raw data to find an attacker. We have little hope of going on an adventure like Mr. Stoll's, but perhaps we should just look at the logs.

Looking at the logs

What you can see

Well... you can see a bunch of data. Lots and lots. We pick July 11th of this year as a "typical" day for the firewall activity. A Tuesday; not a weekend day with a reduced load, nor a Monday with an increased email load. Just your average, typical day. Hmm... the log file from Axent's Raptor firewall is huge!! Well over a million records. So much for looking at the raw data and finding the attacker needle in the firewall log haystack.

Looking at your firewall logs...



Garbage

So what exactly is taking up that almost 300Mb of space?

Mostly junk... mostly just plain successful SMTP connections, lots of folks surfing the web. Not junk for the users. They want it to work, and the firewall helps make it happen. But for us, we want to see attackers, so most of the data we see is unimportant for us.

Content abuse

Every once in a while, we see users going to sites that they shouldn't. "Shouldn't" means they violate corporate policy that says users should not utilize corporate assets for non-business use. In most cases, going to a golf site is not consistent with our corporate use. But in some cases, folks designing golf courses are corporate customers.

We can look for obvious sites that are not business use, and a short email to the employee's manager most times solves the problem.

Hacker evidence and forensics

A good overview for the beginning firewall log viewer is the "FAQ: Firewall Forensics (What am I seeing?" [\(2\)](#) It answers many questions including, "What does this port mean?" "This document explains what you see in firewall logs, especially what port numbers are used by what services. You can use this information to help figure out what hackers are up to.

What you can't do in the firewall logs

Easily search

We could use tools like grep to find individual text strings but even with these tools it is hard to *relate all the information in a log*. Multiple records contain similar information and it is hard if not impossible to relate the large amount of information. We have a real overflow of data with no little information.

Identification

Using alternative tools for analysis

Databases

So our problem remains, how do we get a feeling for the data when there is such a large amount? How do we cut the data into smaller pieces without losing the needle in the haystack?

Databases in general have been used for years to manage large amounts of data and to select data for further analysis.

IBM recommends using SQL queries for finding information in firewall logs [\(1\)](#). Their technique looks for denied packets from the FILTER MATCH table. This table is evidently one of the methods of storing the firewall log in the database.

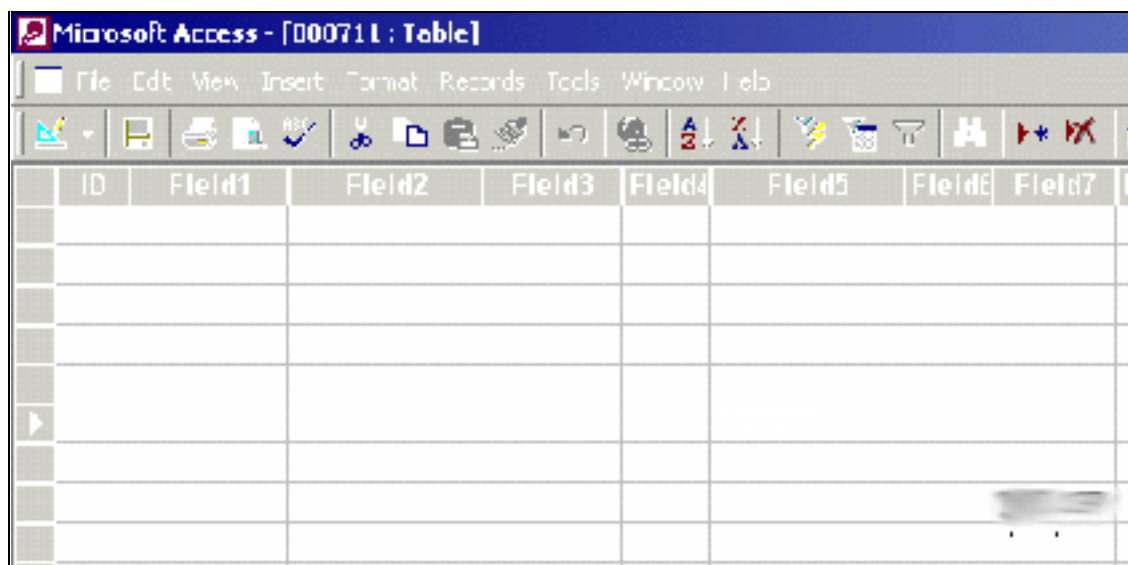
Microsoft Access

We don't have IBM's databases, but we do have a couple of Microsoft's databases. Microsoft SQL is a fairly powerful database that could easily handle the million plus records in our firewall logs. But we could also use Microsoft's Access. Access is a desktop database, but with a little coaxing we can use it for the first phase in the analysis.

The trick here is to browse the data and look for obvious patterns and then cut the data down to size if at all possible. Spreadsheets are great for analysis, but they are limited in the amount of data that they can absorb and manipulate at one time. We use the database to extract the interesting data for analysis with the spreadsheet.

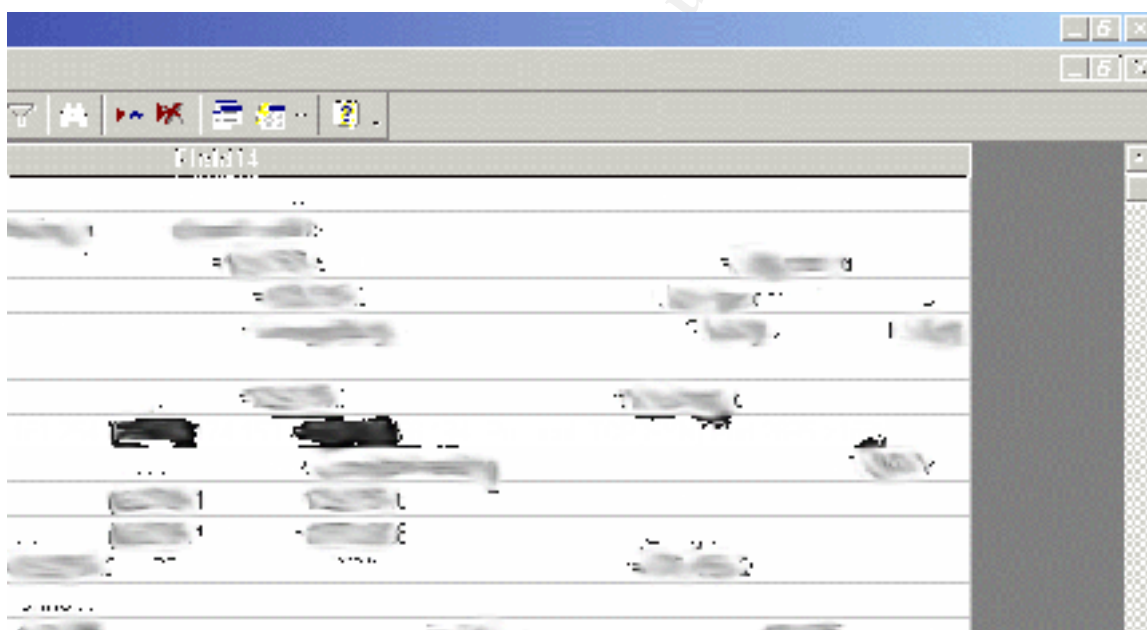
Sure enough, we can now use the firewall log itself to indicate possible problems. If we look at the severity level of the event, we can eliminate those events that probably not a problem.

Looking at your firewall logs...



The screenshot shows the Microsoft Access application window titled "Microsoft Access - [000711 : Table]". The menu bar includes File, Edit, View, Insert, Format, Records, Tools, Window, and Help. The toolbar contains various icons for database operations. The table view shows a grid with columns labeled ID, Field1, Field2, Field3, Field4, Field5, Field6, Field7, and Field8. The first row is highlighted, and the data in the first few columns is visible.

But the first part of the record is not super interesting... we can see the time, we already said the date was 7/11/2000. Field5 indicates the severity level of the event; in this case it is a "Warning". But the really cool stuff is in Field14:



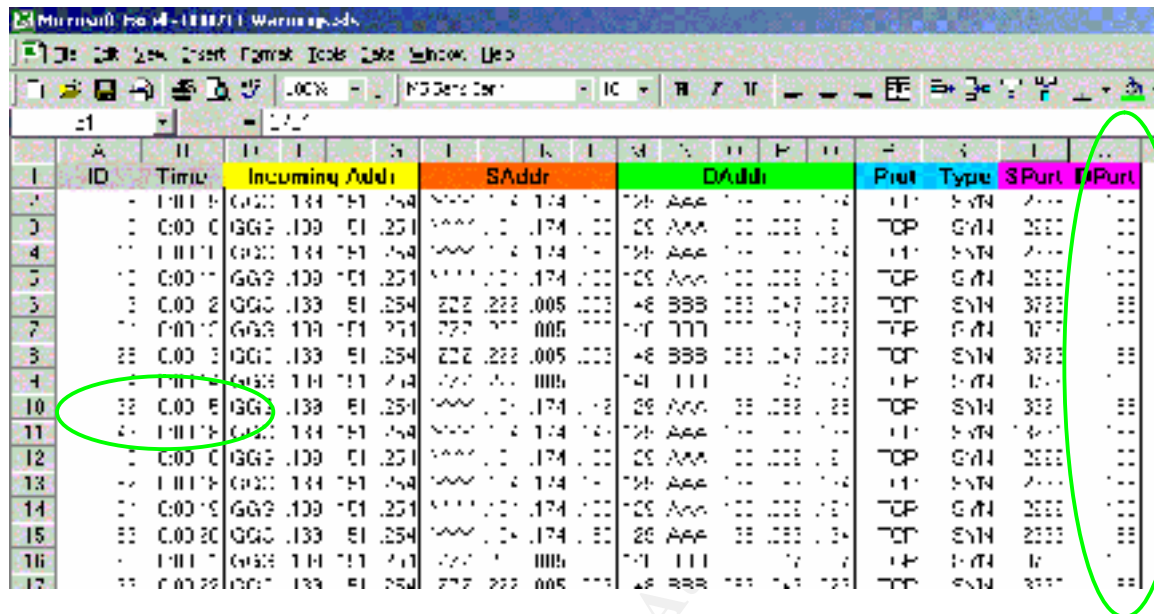
The screenshot shows a detailed view of a record in Field14. The record displays port, protocol, and address information. The data is presented in a structured format with multiple rows of information. The text is somewhat blurred, but the structure is clear.

Here is the indicated port, protocol and address information. (Please excuse the blurred parts of the image, all the addresses here have been modified or obscured for security purposes.)

So, let's cut out the "Warning" level records and put them into a spreadsheet. One of the limitations of Excel is that it is capable of reading only 65,536 rows of data. We actually have more warnings than that. We look at "only" the first 64K. Not quite as good as our million records in Access, but we can analyze the data more easily.

Looking at your firewall logs...

Spreadsheets



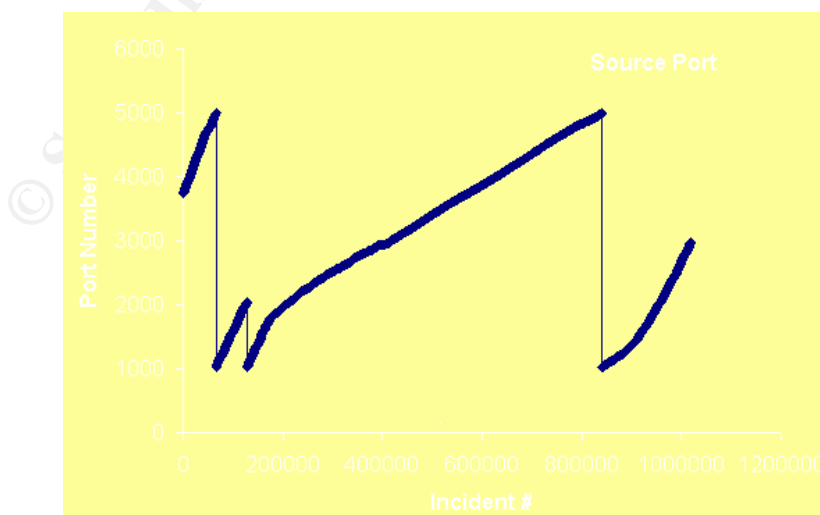
ID	Time	Incoming Addr	SAddr	DAddr	Port	Type	SPort	DPort
1	0:00:00	GGG.133.109	104.174.150	222.86.194	158	SYN	158	158
2	0:00:30	GGG.133.109	104.174.150	222.86.194	158	SYN	158	158
3	0:01:00	GGG.133.109	104.174.150	222.86.194	158	SYN	158	158
4	0:01:30	GGG.133.109	104.174.150	222.86.194	158	SYN	158	158
5	0:02:00	GGG.133.109	104.174.150	222.86.194	158	SYN	158	158
6	0:02:30	GGG.133.109	104.174.150	222.86.194	158	SYN	158	158
7	0:03:00	GGG.133.109	104.174.150	222.86.194	158	SYN	158	158
8	0:03:30	GGG.133.109	104.174.150	222.86.194	158	SYN	158	158
9	0:04:00	GGG.133.109	104.174.150	222.86.194	158	SYN	158	158
10	0:04:30	GGG.133.109	104.174.150	222.86.194	158	SYN	158	158
11	0:05:00	GGG.133.109	104.174.150	222.86.194	158	SYN	158	158
12	0:05:30	GGG.133.109	104.174.150	222.86.194	158	SYN	158	158
13	0:06:00	GGG.133.109	104.174.150	222.86.194	158	SYN	158	158
14	0:06:30	GGG.133.109	104.174.150	222.86.194	158	SYN	158	158
15	0:07:00	GGG.133.109	104.174.150	222.86.194	158	SYN	158	158
16	0:07:30	GGG.133.109	104.174.150	222.86.194	158	SYN	158	158
17	0:08:00	GGG.133.109	104.174.150	222.86.194	158	SYN	158	158

Now this data is interesting. We first start to see a pattern. Look at the destination port (DPort). All set to 158! And... look at the timing... each one is about 30 seconds apart. We have incoming traffic twice a minute, all going to one port. Interesting that the source address appears to be from two sites: YYY.104.174.150 and ZZZ.222.86.194 (remember these are not real addresses). I wonder if this continues throughout the day.

We might have a denial of service attack here... a large number of SYN packets to a single port in a day...

But before we jump to any conclusions, let's look a little deeper... Ahh... the source port appears to be incrementing. (Possibly mis-configured software?)

OK, since we are using office automation products to find problems... let's plot the source address port number and look at what's happening through the day.



Looking at your firewall logs...

The source port is incrementing for each attempt at the port. It appears to start around 1,024 and reset at about 5,000.

There is an interesting jump between Incident #128,088 and #129,434. It would appear that the port number reset to the beginning. If we look at the times for these, the first attempt is at 5:47:52 and #129,434 is at 5:52:06. A whole five-minute break. Reboot?

Histogram

Scientific data analysis folks often will look at the number of events that happen at each value. We all remember that teacher that used to grade "on the curve"... she/he was counting the number of students in each grade category and assigning their grade based on most of the students getting a "C."


We can use Excel to do the same... we take the data and sort it by bins, each bin being the source TCP/IP port. Part of the results of this sorting is in this table.

Value	Count	Frequency	Value	Count	Frequency	Value	Count	Frequency	Value	Count	Frequency
158	61683	94.2%	1022	1	0.0%	873	1	0.0%	641	1	0.0%
445	1534	2.3%	1017	1	0.0%	860	1	0.0%	640	1	0.0%
524	777	1.2%	1014	1	0.0%	853	1	0.0%	637	1	0.0%
514	637	1.0%	1010	1	0.0%	852	1	0.0%	624	1	0.0%
135	531	0.8%	1007	1	0.0%	846	1	0.0%	622	1	0.0%
111	88	0.1%	1003	1	0.0%	840	1	0.0%	621	1	0.0%
631	79	0.1%	1002	1	0.0%	834	1	0.0%	615	1	0.0%
161	20	0.0%	997	1	0.0%	797	1	0.0%	601	1	0.0%
109	9	0.0%	992	1	0.0%	787	1	0.0%	588	1	0.0%
1024	4	0.0%	982	1	0.0%	770	1	0.0%	567	1	0.0%
851	3	0.0%	974	1	0.0%	759	1	0.0%	562	1	0.0%
883	2	0.0%	972	1	0.0%	751	1	0.0%	559	1	0.0%
855	2	0.0%	965	1	0.0%	747	1	0.0%	537	1	0.0%
754	2	0.0%	963	1	0.0%	729	1	0.0%	534	1	0.0%
737	2	0.0%	954	1	0.0%	727	1	0.0%	530	1	0.0%
597	2	0.0%	942	1	0.0%	725	1	0.0%	528	1	0.0%
553	2	0.0%	938	1	0.0%	708	1	0.0%	525	1	0.0%
488	2	0.0%	934	1	0.0%	701	1	0.0%	523	1	0.0%
483	2	0.0%	922	1	0.0%	695	1	0.0%	522	1	0.0%
421	2	0.0%	919	1	0.0%	692	1	0.0%	507	1	0.0%
339	2	0.0%	907	1	0.0%	688	1	0.0%	505	1	0.0%
288	2	0.0%	899	1	0.0%	685	1	0.0%	497	1	0.0%
182	2	0.0%	897	1	0.0%	680	1	0.0%	487	1	0.0%
146	2	0.0%	896	1	0.0%	679	1	0.0%	486	1	0.0%
2	2	0.0%	878	1	0.0%	678	1	0.0%	470	1	0.0%

As we saw before, the number of warning events attempting to reach port 158 is HUGE!! More than 61,000 attempts to connect of the 65, 535 samples we can look at... 94% from that one port... hmm... there really is something wrong...


Since we are looking at the data in a new way, how about another graph? Cool..

Looking at your firewall logs...



The graph is a plot of the data in the table we just looked at. Not much new, we see Port 158 as a huge portion of the events at the firewall, but wait. If we look closely, there are a couple of small spikes along the lower part of the graph, indicating that there are a few other ports that are getting hit pretty hard as well. Nowhere near the number at 158, but still a significant number.

Let's plot it with the vertical axis not going way up to 70,000 events, but just to ten. All the ports that are above that will appear as vertical lines up to (and beyond) ten, but this way we can see what else is going on down in the lower number of events. Remember, we still are interested in possible port scanners. These events will probably appear as a few probes at a large number of ports, not a lot of probes at one port (thank you Mr. 158)...



© 2000, William Farnsworth

Looking at your firewall logs...

Sure enough, we see a bunch of single probes at a large number of ports... but there are also a large number of probes against about 25 ports. Hmm... wonder what they are?

To find out, let's quickly look at the rest of the ports in the purple from the table above and check the "well known ports" assignments. These "well known ports" are available at thousands of web sites. Seems that everyone and their brother feels the need to add these port definitions to their web site. [\(9\)](#)

Anyway, let's look at the ports that appear to be hit very hard...

Value	Count	Frequency	Well known port information	
158	61683	94.2%	pcmail-srv	158/tcp PCMail Server
445	1534	2.3%	microsoft-ds	445/tcp
524	777	1.2%		???
514	637	1.0%	syslog	514/udp
135	531	0.8%	loc-srv	135/tcp Location Service
111	88	0.1%	sunrpc	111/udp rpcbind SUN Remote Procedure Call
631	79	0.1%		???
161	20	0.0%	snmp	161/udp

There is probably some more investigation necessary here. A significant number of events on the unknown ports 524 and 631; we will have to look at the source addresses and try and understand what is happening; and follow-up with a more pointed search of the Internet for information about these ports.

Ports 111 and 135 appear to be possible portmapper connections [\(2\)](#). We do run a significantly large Microsoft shop, so 135 may be valid RPC connections, but we need to follow up. UDP connections to the SNMP port (161) are a flag that we probably are being probed as well [\(2\)](#).

But wait? What about the 142 events that happened only once? One has to suspect that since these are single events they possibly represent real probes. Is that a big deal? Depends on what policy your organization has. We look at this as normal probing that happens all the time. After all, that is why we have a firewall and protection against the big bad Internet. But in any case, since we started looking for real probes, this is where we should start to look for patterns in the incoming source addresses and destination addresses. And then start looking across different days to see if somebody is really trying hard to get in, or it is just casual port mapping.

But, since our investigation dragged us into a possible performance issue, we really need to find out what is happening to port 158... and see what the best way is to get rid of it!

Is it mis-configured software (my best bet) or is it some unknown DOS attack? (Pretty puny DOS attack, huh?) We need to see what happens at Port 158, and try and see what's in the packets. Are they just simple SYN packets indicating an initial attempt to

Looking at your firewall logs...

complete the three way handshake? Or do they actually contain some other kind of information?

DMZ captures and analysis


In order to really find out what is happening to your network outside your firewall you have to... to... look at it!!

Install a neutral machine for analysis of the packet traffic outside the firewall and look at the details of each packet in the DMZ. This is useful for full analysis of packets that will get filtered by the firewall. You then can correlate the time packets are rejected by the firewall with the capture.

Either tcpdump for a UNIX machine or netmon for a Microsoft NT or Windows 2000 can be used for this purpose. In either case, you should capture a few seconds worth of the entire packet. "tcpdump -s 1581 -w savefile" will save all 1581 bytes of the packet into the "savefile" for later analysis. Tcpdump can then be run later (possibly in combination with tcpshow) to provide a full analysis of the traffic in the DMZ. Netmon will capture the full packet and will provide a nifty graphical interface for detailed analysis / forensics.

Note that the netmon product that allows your monitoring machine to enter promiscuous mode is the one included with the SMS Back Office product from Microsoft. The netmon included in the Windows NT resource kit will only capture traffic to and from your machine. You will need the SMS version to capture traffic in and through your DMZ.

The netmon trace indicates that there is not other content in the SYN packet... so it really does appear to be misconfigured software trying to establish a connection.



Looking at your firewall logs...

Let's look more on the Internet for more pmail information. A search at the Microsoft site indicates that the "PC Mail Server" uses this address.

Internet Research

The Microsoft resource kit web site [\(12\)](#) indicates that port 158 is used for the pmail-srv repository.

158 TCP pmail-srv repository PC Mail Server

The PC Mail program was used years ago as the default mail system from Microsoft. Our company has been a Microsoft shop for many years and has installed mail systems of various kinds, including PC Mail, for many years.

PC Mail introduced the concept that all computers are not connected to the network at all times. The "repository" was a place on the mail server to store messages for clients that connected for possibly brief times. In addition, users may have more than one workstation or PC and will need to access their mail from various locations [\(14\)](#).

In the Fall1997 semester at Carnegie Mellon University, Tridas Mukhopadhyay, Instructor and Brian Butler, Teaching Assistant held a class 70-456: *Telecommunications for Business* that had a project assignment. Hadrian D'Souza presented a project titled "*Email and Electronic Communication*" which was an overview of the then current email architectures [\(13\)](#). It included a discussion on DMSP, the generic name for the Microsoft implementation of PC Mail.

"A DMSP session proceeds as follows: a client begins the session with the repository by opening a connection to the repository's machine. The client then authenticates both itself and its user to the repository with a "login" operation. If the authentication is successful, the user performs an arbitrary number of DMSP operations before ending the session with a "logout" operation, at which time the connection is closed by the repository."

This section indicates that we may actually be witnessing a client "opening a connection to the repository's machine"... or better said, trying to open a connection with the mail repository.

Continued searching on the Microsoft site turns up an article on Apple Macintosh PC Mail interconnectivity issues... seems that the default configuration of the PC Mail Macintosh application attempts to connect every 30 seconds. This sounds very familiar to what we have seen. Attempted connections at about 30 second intervals. [\(15\)](#)

Containment

OK... it is pretty clear that we have a mis-configured piece of software, very likely an old implementation of a Macintosh client accessing a non-existent PC Mail repository.

There is a possible performance issue here as well. Our firewall is an application level firewall that does a reasonable amount of processing on each packet received. We have no other evidence of performance problems, but more than 60% of the packets hitting the outside face of the firewall are noise, and they should not be there.

Step 1: contact the potential "offenders." A quick search of the whois records at Internic identifies the two locations. One appears to be a major corporation in the United States,

Looking at your firewall logs...

and the other is an ISP at a European country where we maintain our European headquarters operation. We poked around a little bit and someone in our NOC indicates that we may have done some contract work for the U.S. corporation at question. And... the destination address is in the building where our contractors work.

We have a fairly large remote work force and all their traffic from the Internet come through our firewall in our US headquarters. Possibly a remote user in the European country? Sure enough, the destination address is a server in our European headquarters.

But neither of the two destination addresses is currently active.

Step 2: A couple of pieces of Email to the NOC (network operations center) at each of the source addresses turns up one return email indicating that this "might be a difficult problem to solve." This is not very comforting.

We are still left with the question "what to do?"

Recovery

Step 3: Remember that we have a border router that provides our external interface to the Internet on the outside of our DMZ. This router has some processing and it is fairly easy to establish a rule to drop all attempts to reach port 158.

We configure a rule on the border router that drops all access to port 158 on any internal machine. There is a little concern that we will disable somebody's valid access on port 158 but all our investigation points toward antique software. We will monitor any requests or problems reported to our NOC or help desk.

Problem not solved but deflected for the moment.

Is this enough? I am not sure. One of the great things about the Internet is that there is not a single police force demanding compliance. The bad thing about the Internet is that there is not a single police force demanding compliance. Emails to the offending site's NOC are about all you can do. In this case we will wait and see. And then we will go from there. We have recovered.

Lessons Learned

Have to view the data to get the whole picture

Firewalls are a trade off. We can't hope to capture all the details about each packet that arrives at the outside. We have a typical performance vs. detailed analysis problem. Firewalls are not meant to be true intrusion detection machines and as such they do not capture all the details necessary for complete analysis.

Firewalls are good at protecting our internal resources from external attackers. We can use either state machine algorithms or true application level firewalls... you choose what's best for you... perhaps the best combination is a state machine firewall at the outside with a true application firewall behind it.

In most cases you must use an aggressive monitoring plan to look at the firewall logs for spurious traffic that can overload your firewall.

Looking at your firewall logs...

If you want to monitor bandwidth abusers and folks doing inappropriate surfing the firewall log analysis tools are great.

A good way to look at this data is to use the "standard" office type applications including personal database software in combination with spreadsheets to look for trends and spurious traffic. This let's you get a feeling for the data and traffic outside the firewall and find strange things... Like port 158....

Now on to the rest of the traffic...

Reference:

- (1) "SQL queries for firewall log database tables", IBM AS/400 Information Center. 18 July 1999.
URL: <http://as400bks.rochester.ibm.com/pubs/html/as400/v4r4/ic2924/info/RZAGY521FWSQLEXAMPLES.HTM> (20 July 2000).
- (2) Graham, Robert. "Firewall Seen FAQ." RobertGraham.Com > infosec. 1 June 2000.
URL: <http://www.robertgraham.com/pubs/firewall-seen.html> (18 July 2000).
- (3) "Proxy Log Analysis Software." Network Flight Recorder, Inc. Firewall Wizards Mailing list Archive. 9 Feb 1998. URL: <http://www.nfr.net/firewall-wizards/mail-archive/1998/Feb/0042.html>. (10 July 2000).
- (4) Peter Davis and Associates.; "Logging and Monitoring Software." URL: <http://www.pdaconsulting.com/Log.htm> (20 July 2000).
- (5) FreshMeat.Net. "Console/Log Analyzers." 9 July 2000. URL: <http://freshmeat.net/appindex/console/log%20analyzers.html> (20 Jul 2000).
- (6) Spitzner, Lance. "Intrusion Detection for FW-1." 23 May 2000. URL: <http://www.enteract.com/~lspitz/intrusion.html> . (20 July 2000).
- (7) DFN-CERT.de. "Logsurfer Homepage." 22 Dec 1999. URL: <http://www.cert.dfn.de/eng/logsurf/>. (20 Jul 2000).
- (8) dmoz Open Directory Project. "Top: Computers:" Software: Internet: SiteManagement: Log Analysis: Commercial" 18 July 2000. URL: http://www.dmoz.org/Computers/Software/Internet/Site_Management/Log_Analysis/Commercial/. (20 Jul 2000).
- (9) Brick, Doug. "Stuff." 21 Jun 2000. URL: <http://kom.net/~dbrick/>. (6 Aug 2000).
- (10) Stoll, Clifford, The Cuckoo's Egg, Simon & Schuster, NY, 1995
- (11) Allen, Julia; etal ; "State of the Practice of Intrusion Detection Technologies". Software Engineering Institute, Carnegie Mellon Institute. 31 July 2000. URL: <http://www.sei.cmu.edu/publications/documents/99.reports/99tr028/99tr028abstract.html> (8 Aug 2000)
- (12) URL: http://www.microsoft.com/WINDOWS2000/library/resources/reskit/samplechapters/cnfc/cnfc_por_zqyu.asp

Looking at your firewall logs...

- (13) D'Souza, Hadrian; "Email and Electronic Communication", Project for 70-456: "Telecommunications for Business", Carnegie Mellon University, Fall 1997. URL: <http://www.gsia.cmu.edu/bb26/70-456/projects/email/technical.html>; (2 Sep 2000)
- (14) Wobus, John; "Serving Desktop Computers Using a Central Mail Server on an Internet", Posting on the Usenet mailing lists comp.mail. misc,news.answers, comp.answers, 1 May 1999, URL: <http://www.cs.ruu.nl/wais/html/na-dir/LANs/mail-protocols.html>; (2 Sept 2000)
- (15) Microsoft Corporation; "SFM: Optimizing Polling Rate for Macintosh Clients for PCMail"; Article Q150978, Microsoft Knowledge Base; 18 Feb 1999, URL: <http://support.microsoft.com/support/kb/articles/Q150/9/78.asp> (2 Sept 2000)
- (16) Herel, Heath; "WebTrends Log Analyzer 5.0", ZDNet, PCMagazine, 24 May 2000, URL: <http://www.zdnet.com/pcmag/stories/pipreviews/0,9836,141363,00.html> (3 Sept 2000)
- (17) Check Point Software Technologies, "Check Point FireWall-1", © 2000, URL: <http://www.checkpoint.com/products/firewall-1/>, 3 Sept 2000

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