



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
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Network Abuse – A Backbone Perspective

Levels of Authority

A backbone provider is a company that sells connectivity and (in the case of Genuity) other web services such as web hosting, VPN (Virtual Private Networks), VOIP (Voice Over IP) and managed firewalls to large and small businesses and ISP's who then resell the service(s). We manage and monitor the high speed lines that cross the country and the sea. The services we sell allow our customers access to this global backbone, that they in turn resell. Although this paper is written from the perspective of one provider in particular, I believe that most major backbone providers would encounter similar problems. The biggest problem is that since our customers are resellers, we rarely have any direct authority over the end users. Our customers are bound to our Acceptable use Policy (AUP) and in turn so are their users. It is our job to see that our customers enforce our AUP to their customers. Outside of the problems posed by selling to resellers, I believe any abuse desk will eventually encounter all the problems, and benefit from the solutions I have listed below.

Problems we face

A backbone Abuse Desk has many more responsibilities than that of an abuse desk for a small ISP. Being at the top of the internet food chain means many things. For one, it means we get to deal with reports of abuse from all the different types of services we sell, but we rarely communicate to the actual end user at fault. Since most of our customers use our IP space, all the complaints are directed towards us instead of them. Although this creates more work for us, it also lets us know when we have a serious problem.

Each different service type comes complete with it's own set of vulnerabilities which will be exploited or abused sooner or later. We need to constantly be on top of all the latest Unsolicited Commercial Email (UCE) trends, hacker exploit trends, and copyright/child porn laws, and how they apply internationally. It also means that most of the time we have no direct business relationship with the abusive end user. We are required to notify our customer that our AUP has been violated, but we do not have control over the individual account. The second biggest problem we face is the misconception that complainants have regarding what we can do to solve their problem. Many times if someone on our network is doing something abusive, we cannot just shut him or her off, we have to get in touch with his or her ISP and have them handle the problem. This can sometimes lead to a resolution to the problem taking much longer.

Types of abuse

The dial and DSL network present us with some of the most challenging problems since all the IP's are dynamically assigned. We sell connectivity to our dial network to many ISP's. This means that many ISP's will have the same dial access numbers, and be assigned IP's from the same IP subnet ranges. It is not until we research our RADIUS (Remote Authentication Dial-In user Service) logs that we are able to know which ISP the offending user is connected to

us through. Once we have identified the ISP, we notify them that their user is in violation of the AUP. The problem with this is that the user will almost always then get a new account (especially if they are free) with the same ISP, or in many cases another ISP who is also on our network. The user then continues to abuse the network sourcing more complaints to our abuse desk. Most complainants believe that we have done nothing since they may still see the same abuse, from the same IP ranges. This same type of thing happens with DSL also. This is because the IP addresses are dynamically assigned from the same pools of subnets. There are some proactive steps that can be taken which will be discussed later in this paper.

On the DSL and Dial network the type of abuse usually varies pretty widely. The most common form of abuse relates to UCE. Some times the downstream user is sourcing the UCE, or sometimes the UCE references a site on our network. Either case is against our AUP, and is dealt with immediately with a priority set by how many complaints we receive. This usually involves analyzing email headers or decoding obfuscated, or CGI laden URLs. The other common complaints we get are regarding scans from our networks. These complaints are more complicated because not all logs look the same, and we have to know how to decipher many different types. On top of that we need to be up on all the latest exploits so that when we read the logs we know whether the complaint is regarding a harmless DNS response packet, or a scan for a back door trojan. The end users guilty of the scans are not always at fault either. Most of the time we discover that they had been infected with a worm, that was merely looking for other vulnerable people to infect. Regardless, it is up to our abuse desk to pass this on to our customers, and make sure they are aware of the possibilities, so that they are well informed when they contact their customer.

Other service types to consider are our web hosting and managed connectivity offerings. We see a few instances where our direct customer is sending out UCE to advertise their struggling product. If we determine this to be the case our first response is to try and work with our customer. Many claim innocence because someone in their marketing department told them that this was completely legal. We do our best to re-educate these people as to the best practices they should be using to do this work. If it is determined that these abusive practices are built into their business model, and they refuse to work with us to change that, then we have no choice but to escalate to our legal department for review of the contract. Due to customer education, and a greater awareness of our AUP, the frequency of this happening is significantly lower than it used to be.

Reports of scans from these customers are also a bit more serious since they are usually from production servers instead of end users PCs. The challenge there is trying to determine if this machine is supposed to be generating the reported traffic, or if they have been compromised. Many times a personal firewall set on "paranoid" will generate alarms, on non-threatening traffic. After a full review of the logs we normally contact the customer, and give them our analysis of the situation, or if we are unsure, we request an answer from them. If it appears that they have been compromised, we point them in the right direction to clean the infected machine.

The type of abuse our VOIP service generates is almost always a legal issue. End users connect to our VOIP network through many different providers. They can place a phone call from their PC, to another telephone or computer. All of the complaints we have received so far

are regarding harassment. People think this is an anonymous service and that they can harass anyone without being identified. This varies from constant hang-ups, to verbal threats and harassment. Abuse of this sort, though covered under our AUP, is considered a legal issue because it also falls under federal law. Once we have located the logs showing the abuse, they are forwarded to our customer for action. This usually results in the termination of the guilty end user. We have not yet received complaints that our VOIP network is being used for unsolicited commercial purposes, but like all abuse, it is probably only a matter of time.

Tools

There are many tools we use to combat network abuse. They vary from the investigative tools we use to take action reactively, to measures implemented on the network, to fight abuse proactively. Previously almost all the abuse on our network was fought reactively. Since we have made an attempt to increase our visibility into the company, we have had an easier time of convincing upper management to implement filters on the network to help out proactively. We have seen a drastic decrease in abuse complaints since we have started our proactive steps.

First and foremost, the most important tool is our AUP. We reference it directly every day because every complaint received must be processed in accordance with the AUP. I feel we have a very strong AUP (available at <http://www.genuity.com/aup/aup.htm>) because so far, we have yet to have legitimate abuse reported to us that was not covered by our AUP. The AUP is waved in the face of any customer that does not take our notifications of abuse seriously. We point out the section they are violating, and then point to the section that says we will review their contract with legal, if they do not feel like taking action. This will usually get the attention of any customer.

The customer contract is also considered a tool, and it comes into use for two main reasons. The first reason is related to the above paragraph. If a customer refuses to comply with our AUP policy, we turn to the contract to see about legal ramifications, and possibly disconnection. Compliance with our AUP is written into every contract. We have a well-established relationship with several lawyers for our company so that we may consult with them when ever necessary. The other time the contract comes up is when a complainant is demanding information regarding our customer. A privacy policy is written into every contract that keeps us from giving out any information regarding that customer, or their end users. We often get requests from disgruntled complainants who would like to know who scanned them.

We have several "Best Practices" documents, which we use to combat abuse also. The main one is our best practices for mailing lists. We forward this to any customer who has been identified as directly sending SPAM to advertise their product. Our best practices document outlines how to use a confirmed opt-in mailing list to reach an audience. This is the only type of mailing list we recommend. The idea behind it is that everyone on your list has asked to receive the information being sent to them. To make sure that no one is signed up unknowingly, you must send each address an email confirming the subscription; they are not added to the list, until they respond to that confirmation email. This prevents people from signing up someone else for a mailing list out of spite.

In an effort to be more proactive, we have started to implement filters on our dynamic networks. This came about because our abuse desk found itself playing a game with the abusive end users. We would speak to their ISP, and get the account turned off. They would come back under a new name with the same ISP, or a different one on our network. This was becoming a losing battle, since they could create new accounts faster than we could get them terminated. It was decided that proactive steps needed to be taken. Since the filters must be setup with each customer individually, the process is slow going. In an effort to be more effective, we have started filtering by looking at which of our customers are being abused the worst. The following paragraphs discuss the various filtering techniques we have used.

The first filter implemented was SMTP (simple mail transport protocol) on port 25. The filter blocks all outgoing connections on port 25 unless they are being made to a mail server designated by the filter. This essentially means that the end user can ONLY use their ISP's mail servers to send outbound mail. This keeps them from being able to abuse servers on other networks that are configured as open-relays. An open relay is a machine that will forward mail from anyone on the internet, to anyone else on the internet. A smart SPAMMER will use open relays in conjunction with free anonymous accounts to completely avoid getting into any trouble. Theoretically the end user could just SPAM through their ISP's mail servers, however these servers are monitored closely, and any excessive load is noticed and stopped almost immediately. It is still possible to send SPAM from the filtered networks if the email is web based. This is due to the fact that the traffic from the end users machine travels on port 80 instead of port 25. So far we have not seen this to be a huge issue yet. Filtering on port 25 caused a dramatic drop in the number of complaints we were receiving regarding abuse on our network.

Port 119 NNTP (Network News Transport Protocol) filtering is something we are currently trying to push. This is the protocol for posting an article to Usenet. Many news servers are open for use by anyone similar to an open mail relay. Posting of a commercial nature, and multiple, excessive postings, are all against our AUP, and most news group charters. Port 119 filters prohibit a user from posting to any news server other than their own ISP's. We have started to implement this filter for our customers, but it has not been as effective as the port 25 filters. This is because many news services are web based. This means that the traffic travels outbound from the end users machine on port 80 instead of port 119. Once we have finished port 119 filters on all our customers, we will be able to make a better judgement on the impact it is having on Usenet related abuse complaints.

A new form of abuse that has become very popular on our dial network recently, has caused us to start implementing directional filters on port 80. SPAMMERS are setting up web servers and running them through their dial-up connections. They are then sending out UCE that references the URL, which traces to our dial network. Until we can get the account name cancelled we attempt to bump the users connection off. It appears that they have the DNS zone information set to expire very quickly. This means that if the zone file your local DNS server is caching is more than a couple minutes old, it will go to the primary DNS server for the domain to lookup the information. This enables the SPAMMER to dial back in, get a new IP, and have the site backup within minutes of us bumping him offline. This service is called "dynamic DNS" and there are several providers specializing in this. We have started directional port 80 filtering

so that no inbound connections can be made to our dial network to port 80. This is a great example of how an abuse desk needs to constantly be up to date on the latest trends.

Being up-to-date on trends is a tool in itself. When we see complaints regarding scans from our network we always keep our eyes out for trends. Once we are able to identify the exploit that the scans are looking for, we can send out a mailing to educate our customers. We explain to them that this is a vulnerability that is being scanned for in large quantities, and we tell them how to protect themselves against it. I like to think we are being reactively proactive, by taking the information that is sent to us, and using it to prevent further damage.

One of the most important things for any abuse desk is to have a good connection within the company. We are lucky in that some people very high up in the management tree feel that both security and SPAM issues are important topics. They understand our position, and help us get the tools we need to do the job. Without their support, we would not have most of the influence, and contacts that we do. It is very important that the abuse group has enough visibility into the company to communicate effectively with all the groups. Problems can arise if many parts of the business are unaware of abuse issues. Sales people who use “pink contracts” to bring a customer onboard a prime example of this. A “pink contract” is an addendum allowing a customer to do something that would not be allowed under a normal contract. Usually they are receiving permission to SPAM freely on their ISPs network. Our legal staff is on the lookout for this so that no contract like this should ever get past them.

Due to the large amount of complaints we receive, we developed an in-house tool that attempts to sort the complaints by topic such as UCE, Security, or Usenet, and also group like complaints together. It attempts to take all complaints referencing the same piece of UCE, and group them together for us. This saves us a lot of time, and manual parsing. It also helps by associating the complaints to ticket numbers in our ticketing system. This is what we use to track every single complaint. We have to do this because at any point, if we are accused of not parsing a complaint we need to be able to produce proof that we have.

What does this mean to the rest of the internet?

There are several bills currently in various stages of development and revision that may outlaw UCE in various forms. The opinions on them vary, and delving into that would be out of the scope of this paper. The important thing to note, is that when a bill finally does pass, it will most likely include all or some of the following: requirements that senders of UCE do not falsify any header information, UCE must contain valid reply-to address, sender must honor all opt-out requests, some form of postage to be paid by the sender so the bulk of the cost is not going to the end user or ISP, requirements that ISPs take action against SPAMMERS, protection for ISPs who may retransmit UCE unknowingly, and outlawing “spamware”. Although this list is not by any means complete, any one of these things will mean extra work for abuse desks. Any ISP that does not have someone handling complaints in a timely manner could face punishment under some of the proposed laws.

Every ISP, large or small is eventually going to need a devoted abuse desk. As the number of people on the internet grows, so will the number of people that are using it for no good. All of the problems, procedures, and tools I have listed above will be necessary for any ISP. When it becomes apparent that for one reason or another, an ISP is abuse friendly, the first people to find out, will be the ones ready to abuse it, the second ones to find out, will be the customers. No ISP should get caught without a decent abuse desk, or they will face a lot of damage in the end. Although this paper is focused on the problems faced by a backbone abuse desk, I feel that many of the same things apply to any abuse desk.

© SANS Institute 2000 - 2002, Author retains full rights

References

Festa, Paul. "AT&T admits spam offense after contract exposed." 3 Nov. 2000. URL: <http://news.cnet.com/news/0-1005-200-3369773.html> (28 May 2001)

Festa, Paul. "PSINet cans spammer, pledges reforms" 8 Nov. 2000. URL: <http://news.cnet.com/news/0-1005-202-3585163.html> (28 May 2001)

Smith, David E. "Dynamic DNS." 23 May 2001. URL: <http://www.technopagan.org/dynamic/> (28 May 2001)

"107th Congress – Currently Pending Legislation." 26 Apr. 2001. URL: <http://www.cauce.org/legislation/index.shtml> (28 May 2001)

"Acceptable Use Policy." Version 1.1, 11 Feb. 1998. URL: <http://www.genuity.com/aup/aup.htm> (28 May 2001)

"DNS Hosting for Domains." URL: <http://www.dtdns.com/info/domains.cfm> (28 May 2001)

"Net Policies and Procedures – How to Report an Abuse." URL: <http://www.genuity.com/aup/how.htm> (28 May 2001)

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