

Internet and Network Service Provider Network Abuse Management

David Schwartzburg

Abstract—Network abuse is a pertinent issue for network and Internet service providers alike. Establishing an acceptable use policy and allocating resources to process and respond to evidence of network abuse is essential for service providers. The general structure of the abuse management process will be discussed.

Index Terms—ISP, NSP, Abuse Policy, Network Abuse, Complaint Handling, AUP, and TOS.

I. INTRODUCTION

INTERNET service providers (ISP) and Network Service Providers (NSP) of all sizes must possess a well defined, documented, and supported abuse process. As the world becomes increasingly dependent on information technology, “significant opportunities now exist for both mischievous and malicious abuse via IT systems” [1]. There are many types of harmful unsolicited activities [2], some intentional, others unintended, that require a response, action, or some form of mitigation from the responsible service provider. A few examples of these activities include worm propagation, email viruses, denial of service attacks, unsolicited bulk email (UBE or spam), phishing scams, child pornography, and copyright infringement. All types of network abuse are costly for victim organizations. For example, handling spam can cost organizations large sums of money and can result in lost employee productivity [3].

For these reasons, and to maintain utility of the public Internet it is imperative that ISPs and NSPs take proactive measures to prevent the above and other types of inappropriate conduct from originating from their networks. If preventative measures are neither feasible nor entirely effective, then the importance of having staff that are capable of responding to complaints that run the gamut of network abuse issues cannot be underestimated.

In order to appropriately deal with abusive actions from connected network customers, the first requirement is the establishment of an acceptable use policy (AUP). This policy is used to define acceptable network behavior and also specifically identifies prohibited behaviors.

In addition to the AUP, a provider should have an individual or a team of professionals capable of reviewing and interpreting complaints received about abusive activity

originating from the users of a provider's networks. In order for a provider's abuse group to be aware of issues stemming from their customers' use of the network they must be obtain abuse data from internal and external sources.

This paper will examine the need for effective abuse process management by ISPs and NSPs. Methods of organizing and handling complaint data will be discussed, as well as the sources of relevant data that are indicative of abusive activity. The goal is to demonstrate the importance of the collection and response to evidence of network abuse in the maintenance of safe, efficient networks.

II. NETWORK ABUSE POLICY

One of the most critical elements of an effective abuse management process is the establishment of the AUP. The AUP may often be stipulated in a provider's terms of service (TOS) which is an agreement between the provider and the customer that identifies both parties' responsibilities [4]. It is difficult to hold a network user responsible for illicit activity if they are not made aware, in advance, of which types of activity are appropriate and which types are prohibited. For example, on their website, Spamhaus emphasizes the importance of an effective AUP on the prevention of spam [5]. They also provide examples of the AUPs of a number of service providers. A review of these policies indicates that a well-defined AUP should have a number of basic components.

The first section of an AUP is the introduction. This section provides a definition of the terms and parties that will be referenced in the remainder of the policy. The introduction may specify the purpose and scope of the policy, or they may be defined in separate sections. An essential section of an effective AUP is one which identifies specific types of prohibited network activity. This section will generally cover: spam, virus, worm, trojan, fraud, copyright, legal, harassment, and security topics. Regardless of how they are organized, it is important that a provider's AUP address these topics. The AUP should be revised as new issues or types of network abuse surface. Therefore, a section or provision is usually included that will address the terms under which the provider will revise the policy and in what fashion notification of revision will be made.

An effective AUP will generally contain information on how infractions of the policy should be reported. This provides victims of network abuse with a clearly identified

Manuscript received December 2, 2005.

D. Schwartzburg is with the abuse group of a NSP/ISP. (e-mail: djs0605@mail.ecu.edu)

process to properly report evidence of AUP infractions that should be brought to the attention of the provider. Another topic that the AUP may frequently address is the responsibilities that the provider assigns to the users of their network. This section could hold the customer responsible for maintaining the security of their resources connected to the provider's network. It may also state that all of the customer's end users are bound by the AUP. An AUP should specify by what means the provider will notify their customers of any infractions of the policy.

Finally, the AUP should stipulate that the provider has the right to monitor their networks for abuse and respond accordingly. This provision should allow the provider to perform scanning of systems and networks in response to complaints or evidence received that indicates that a network abuse has occurred.

The AUP is a policy to which all customers should be bound and made aware. It is important that the AUP be carefully constructed to effectively communicate what activities are prohibited. It should clearly define the scope of the policy, the consequences of violations, and the potential actions that the provider may take in response to infractions of the policy.

III.SOURCES OF EVIDENCE OF NETWORK ABUSE

A.External Sources

RFC 2142 specifies the email address <abuse@domain> as the initial point of contact for abuse issues stemming from a particular organization [6]. For external parties that have evidence of abusive network activity originating from a particular organization, this address is their means of relaying evidence and notification of policy infractions to the responsible provider.

In addition to maintaining an abuse mailbox in compliance with RFC 2142, a provider should also add their abuse contact to the data maintained by the registrar for their IP block assignments. This will make it easier for users to identify the appropriate abuse contact when they perform a WHOIS query for a particular IP address. A provider may also register their desired abuse contact preferences with an organization such as abuse.net which maintains a WHOIS server of abuse contact data that is searchable by domain [7]

It is important that abuse email be processed in order to rectify network abuse that may be originating from a provider's network. If email complaints are not processed and investigated, then it is possible that a network provider might experience "blocking, filtering and blacklisting ... with other IP related Service providers" [8]. This could lead to disruption of service due to email realtime blackhole lists (RBL) [9] or a network configuration, such as a null route (black hole) that prevents traffic from routing to or from a particular resource as is normally expected.

There are a number of different categories of senders that

regularly send complaints about AUP infractions to the abuse address. The first type is a regular user. This may be a user on the provider's own network, or perhaps a user from a different network. In either case, they most likely will be reporting a single incident that has occurred. This may be referred to as a one-off due to the fact that these incidents are usually manually reported by a user. Another category of complainant is referred to as a trusted sender. This category is comprised of individuals or organizations that frequently submit reports. These reports are likely to have a consistent format, and there is a high degree of veracity in the allegations made in previous complaints. A sub-category of the trusted sender would be the aggregation complainants. This would include sources such as Dshield, myNetWatchman, Spamcop, and other firewall log and spam complaint aggregation services. For example, myNetWatchman is a free service that uses a software agent to collect firewall data, submit it to a centralized system, and then correlate it and generate alerts to the responsible parties of unsolicited activities generally indicative of a compromised system [10]. This data can prove highly useful as there are usually numerous sources within a single aggregated report that are receiving the same type of unsolicited activity on disparate networks, which adds to the credibility of such evidence.

Another external source of evidence is law enforcement agencies. These agencies may require a response from the provider in regards to legal issues or investigations. With any size provider, there is always the possibility of illegal activity occurring, and the result of that activity may be a subpoena or other legal instrument used to obtain evidence from the provider.

B.Internal Sources

A provider may take a more active role in the abuse management process by seeking internal sources of evidence of network abuse. One such source would be firewall logs that describe suspicious activity originating from a particular user. If this data can be accurately identified and processed by the abuse group, then it may be possible to isolate infected and compromised systems that are scanning for other systems to compromise. Other similar sources would include data obtained from IDS, honeypot [11], honeynet [12], or monitored dark IP space. All of these sources may be useful in identifying systems that have security concerns such as viruses, worms, or trojans or systems that have been compromised by an attacker. Data that originates from monitored dark IP space can be particularly useful in identifying compromised systems as by definition, dark IP space should not see any traffic destined for it [13].

IV.CLASSIFYING NETWORK ABUSE EVIDENCE

It may be necessary to categorize evidence and distribute it to various groups or individuals. For example, evidence

might be divided up into a number of categories such as security, legal, and spam. The security category would be comprised of evidence of viruses, worms, trojans, hacking activity, unsolicited network activity, and phishing scams, while the legal category would include issues involving copyright infringement, harassment, death threats, child pornography, and other legal issues. The legal classification would also apply to court orders, subpoenas, summonses, discovery requests, and warrants [14]. Evidence that is delegated to the spam category would include complaints about UBE that originated from the provider's network and spamvertised web sites [15] hosted on the providers network. UBE evidence may be additionally classified into subcategories based on whether the messages appear to have originated from a compromised system or directly from the systems of intentional spammers.

Depending on the size of the provider, the volume of complaints that they receive, and the amount of resources the provider can allocate to abuse management, these evidence classifications may be managed in a number of different ways. They may all be handled by one individual or by a group of individuals. It is also possible that evidence classified as security or spam will be handled by the same group, while legal issues, which may require a security clearance be handled by another group or individual. For large providers, each classification may have its own group of agents to process evidence of the respective classification.

V. PROCESSING EVIDENCE OF NETWORK ABUSE

A. Email Complaints

Most abuse evidence will likely be in the form of email complaints. One possible exception is evidence that originates from a provider's internal sources. For a small provider, it may be feasible for a single individual to manually process abuse complaints using a simple email client configured to download or view email for the abuse mailbox. If a group of individuals are all tasked to process abuse evidence, then a more complex system that allows for the tracking of issues and emails will most likely be required.

B. System and Software Requirements

For large providers that have a high volume of abuse evidence, a team of individuals responding to and working on abuse incidents, or the desire to maintain a more secure record of their abuse management process, additional systems will most likely be required. Securely storing and archiving incoming email evidence, customer responses, and outgoing email responses from the provider will be a priority. The method of storing this data is not necessarily important, as long as the data is accessible by all individuals within the abuse group who are authorized and have a need to view it. A well-defined backup plan should be in place to protect the integrity and availability of the stored email history. A database may be an ideal solution for the storage of the abuse

email history as it allows for the indexing of values, such as the sender, recipient, IP addresses, and URLs. For providers with a large volume of abuse email data to retain, a database may be the best option.

For more advanced abuse functions, customized tools may be used by or developed for the abuse group. These tools might include freely available open source tools, commercial software packages, or customized tools developed specifically for or by the abuse group. Some of the functions performed by these tools may include incident tracking, vulnerability testing, network scanning, and customer lookup tools.

C. Automation

For providers that have large volumes (thousands of messages per day or week) of network abuse evidence to process, it becomes very inefficient to have abuse agents manually processing evidence. Many types of abuse evidence will have a consistent format. Complaints that originate from firewall log or spam aggregation services or internal sources are frequently formatted consistently. When the format of evidence is consistent, it becomes relatively simple to implement a program or script that can process the evidence and pull the relevant data from the evidence.

For example, 135, 139, and 445 are examples of TCP ports associated with a few Microsoft Windows services that are commonly exploited by worms and other malicious types of software [16, 17]. When credible evidence exists that a system is scanning on one or more of these ports, then it is probable that the system, or systems if address translation is involved, is infected with a worm. In these cases of known types of malicious activity, it is possible to automate the process of customer warning and suspension. Thresholds can be set in the programs or scripts that control how many warnings a customer should receive, how frequently they should receive them, and whether a customer's service should be suspended after a certain number of warnings. A key to automating the abuse process is being able to identify the customer responsible for a network abuse incident.

D. Identifying Customers

One essential ability an abuse group must have in order to function effectively is access to tools that allow the group to accurately identify the customers that are actually responsible for generating the network abuse incidents originating from the provider's networks. Misidentification of customers can have grave consequences, including legal ramifications, so it is essential that these tools are accurate [18]. This is especially true when identification of dynamic customers is required. If different customers have been assigned the same IP address over the course of a relatively short span of time, then it becomes easier to misidentify customers, especially if dynamic leases and network abuse evidence are logged in different time zones. It is important to correct the times of leases and abuse incidents into a common time zone before a customer lookup is performed.

If the abuse group does not have the ability to associate a specific customer with a particular network abuse incident,

then they will likely be powerless to take any action in response to that abuse to prevent it from recurring.

E. Investigations

The AUP should disclose to customers that the provider retains the right to investigate alleged network abuses originating from the provider's networks. In the case of legal issues or investigations that involve warrants, discovery requests, or subpoenas, the provider will most likely be capturing data or disclosing information in cooperation with the investigation of a law enforcement agency.

The provider may perform investigations of its own in response to evidence of abuse related to non-legal issues, such as activity that is generated by worms or viruses. For systems that are generating unsolicited connection attempts to other systems and networks, a port scan of the system in question is often revealing. Nmap is an excellent, open source tool that can be used to scan a system and help identify potential abnormalities [19]. Once a scan has been run in Nmap, or a different port scanning utility, information may be returned that indicates what type of system has been scanned, and what type of services it may be running. Research from anti-virus vendor Sophos has shown that there is a fifty percent chance of an exposed, unpatched Microsoft Windows system becoming infected in the first twelve minutes it is connected to an unprotected network [20]. If Nmap results indicate that the system is a Microsoft Windows system that does not appear to have firewall protection, then the probability is high that the system is infected. This is especially true in the cases of systems that are sending unsolicited network traffic.

Other investigation may include the use of commercial, open source, or custom developed tools to test if systems are vulnerable to various exploits, have unsecured Windows file shares, are an open proxy or relay, or display evidence of being compromised by an attacker.

In addition to investigations that result from evidence of what is most likely worm or virus related activity, investigations may also be performed in response to evidence of spam or phishing abuse that has occurred. This is especially the case if this evidence has been classified as being likely to have originated from infected or compromised systems.

It is important in conducting investigations that the abuse group be careful not to use techniques that would actually exploit, penetrate, or otherwise cause harm to the customer's systems. This would cross the line of what a responsible provider should ethically do in response to evidence of abuse in order to protect their customers, and would likely place the provider in an undesirable legal position.

F. Response to Abuse

Once evidence of abuse has been received, classified, and investigated, then a provider takes an action in response to the evidence. The provider may determine that the evidence is not accurate or reliable and that no infraction of their AUP has occurred, and take no action. In the case in which it has been determined that there has been an infraction of the

provider's AUP, the provider may decide to not take any action at the present time other than to archive the evidence for future reference. The abuse group may decide to issue a warning or notification to the customer. If the infraction is severe enough in nature, then the customer's service may be suspended with little or no prior warning for their own protection and the protection of others. The provider should have a defined plan for acting on valid evidence and may want to set thresholds for the number of warnings that are issued prior to a suspension of service. This information may be published in the AUP, but it may also be kept confidential by the provider in order to prevent intentional subversion of the policy.

Whenever the provider takes an action against the customer, whether it be notification, warning, or suspension, they will need agents that are capable of dealing with customer responses to their actions. The provider will need to determine if this function can be performed in whole or in part by their support organization, or whether the function will be delegated entirely to the abuse group. Regardless of which group this function is assigned to, the responsible agents will need to be capable of dealing with the customers in the same fashion that they are notified, whether that be email, telephone, or by other means.

VI. CONCLUSION

Providers of network and Internet access must have policy in place and resources designated to investigate and respond to network abuse incidents. Without these resources in place it will be more probable that ISP and NSP networks will be abused and that providers will face legal repercussions or connectivity sanctions from other providers or RBLs. Failure to implement or enforce an AUP will lead a provider to be considered a good candidate for those that are intentionally generating network abuse such as UBE, phishing, and network attacks. The likely end result in such a situation is the degradation of the provider's reputation in the online community.

REFERENCES

- [1] P. S. Dowland, S. M. Furnell, H. M. Illingworth, and P. L. Reynolds. 1999. Computer crime and abuse: A survey of public attitudes and awareness. *Computers & Security* 18 (8): 715-726.
- [2] P. Sandford, D. J. Parish, and J. M. Sandford. Identifying Internet Abuse in ISP Networks. In: *Proceedings of Safety and Security in a Networked World: Balancing Cyber-Rights and Responsibilities*. September 8-10, 2005, Oxford, United Kingdom.
- [3] B. Hancock. 1999. First Internet E-Mail Corporate Usage Report; concludes E-mail Abuse at Epidemic Levels. *Computers & Security* 18 (3): 188.
- [4] Wikipedia. Terms of Service. Available at http://en.wikipedia.org/wiki/Terms_of_service. Retrieved November 20, 2005.
- [5] Spamhaus. Responsible ISP Acceptable Use Policies. Available at <http://www.spamhaus.org/aups.html>. Retrieved November 20, 2005.
- [6] D. Crocker. 1997. Mailbox Names For Common Services, Roles and Functions. Request For Comments: 2142. Available from <http://www.ietf.org/rfc/rfc2142.txt>. Retrieved November 20, 2005.
- [7] Network Abuse Clearinghouse. 2005. Available from <http://www.abuse.net/>. Retrieved November 21, 2005.
- [8] AT&T Acceptable Use Policy. 2004. Available from <http://www.att.com/aup/>. Retrieved November 21, 2005.
- [9] Webopedia. 2005. Realtime Blackhole List. Available from <http://www.webopedia.com/TERM/R/RBL.html>. Retrieved November 22, 2005.
- [10] myNetWatchman. What is myNetWatchman. Available from <http://www.mynetwatchman.com/about.asp>. Retrieved November 22, 2005.
- [11] Wikipedia. Honeybot. Available from <http://en.wikipedia.org/wiki/Honeybot>. Retrieved November 23, 2005.
- [12] Wikipedia. Honeynet. Available from <http://en.wikipedia.org/wiki/Honeynet>. Retrieved November 23, 2005.
- [13] V. Yegneswaran, P. Barford, and D. Plonka. 2004. On the design and use of Internet sinks for network abuse monitoring. *Recent Advances in Intrusion Detection, Proceedings Lecture Notes in Computer Science* 3224: 146-165.
- [14] LFC Hosting. Acceptable Use Policy. Available from http://lfchosting.com/scripts/frameset.cgi?/corp_policies_acceptable.shtml. Retrieved November 27, 2005.
- [15] AbuseButler. Spamvertized Sites. Available from <http://spamvertized.abusebutler.com/>. Retrieved November 27, 2005.
- [16] myNetWatchman. Ports Being Attacked Most. Available from <http://www.mynetwatchman.com/tp.asp>. Retrieved November 29, 2005.
- [17] SANS. 2005. The Twenty Most Critical Internet Security Vulnerabilities (Updated) - The Experts Consensus. Available from <http://www.sans.org/top20/#w1>. Retrieved November 29, 2005.
- [18] R. Sylvester. 2005. Mistaken child-porn raid leads to lawsuit. Available from http://www.kansas.com/mld/eagle/news/local/crime_courts/12620843.htm. Retrieved November 29, 2005.
- [19] Nmap. Free Security Scanner For Network Exploration & Security Audits. Available from <http://www.insecure.org/nmap/index.html>. Retrieved November 29, 2005.
- [20] Sophos. 2005. Virus writing on the up as average time to infection spirals down. Available from http://www.sophos.com/pressoffice/news/articles/2005/07/pr_uk_midyearroundup2005.html. Retrieved November 29, 2005.