

Going to the Source: Reporting Security Incidents to ISPs

James C. Slora Jr. 2002-03-12

Going to the Source: Reporting Security Incidents to ISPs

by James C. Slora Jr.

last updated March 12, 2002

Introduction

My interest in abuse notifications began when Warez pirates started using my trustingly anonymous FTP server as their personal playground. I realized that my system needed to be locked against this type of intrusion and that I had failed to provide adequate safeguards. But I still felt violated – these people were intruding into a place where they knew they had no business.

I complained to the service providers of several of the pirates, but did not get very satisfying results - just a few form letters with no follow-up. It seemed like they did not care, and that I had wasted my time.

Other types of attacks afterwards nevertheless convinced me to continue to submit abuse reports, and over time I realized that much of the failure of my early reports was caused by my own poor technique. Since then I have modified my original log submission methods to cope with the realities of getting help from abuse handlers, and have greatly improved my rate of successful resolution.

My technique continues to evolve with experience, but I'd like to share some of the things I think I have learned so far from many successful and unsuccessful abuse reports.

Why File Abuse Reports?

The first and best reason to file an abuse report is to help protect the systems for which you are responsible.

Firewalls, bastion routers and intrusion detection systems play an important role in protecting our networks, but they fail to address the source of break-in attempts - the vast worldwide interconnection of hostile users, compromised systems, software bugs and configuration glitches.

The Internet is studded with automated intruders like Nimda and Sadmin that find new victims every day. They heckle us persistently in our intrusion detection and firewall logs – as do the Warez pirates, script kiddies, spammers, and other dark forces. We think that we understand most of these attackers. We think that we know how to prevent them from accessing our systems, but it would be an arrogant mistake to believe that they are not a threat. If systems administrators take action against the attacks they know about, there is some chance that they may prevent other problems. Their reports may help eliminate a compromise on a system that poses no threat when it is clean. The report could help a casual experimental abuser of network scanners to rethink their habits. It could even help an ISP or law enforcement agency to identify a source of serious illegal activity.

Most end-users are happy to learn about a problem with their system, and most providers have an interest in reducing abuse of their services. Many, if not most, of the intrusion attempts against my networks have been performed through compromised systems - not directly by abusive users. Thus when I file a report, I am not necessarily filing a complaint against a hacker. I am more likely letting someone know that they have a configuration issue that leaves them open to exploitation by another person, and that this exploitation is causing unwanted traffic to my network.

When to Report Abuse

I report nearly all serious intrusion attempts, and as many minor ones as I can make time for.

Reports are worthwhile even for one-time intrusion attempts. My reports have helped schools, universities, businesses, and government agencies to detect and recover from system compromises that might have caused much larger problems later. My reports were not special, they just caught the attention of administrators so they could apply their own expertise to repairing their systems.

Correspondents generally respect that if the person filing a report considers traffic to be an intrusion, then it is. They will usually investigate the complaint, and will at least provide the individual filing the report with an explanation of the traffic if they consider it to be legitimate. I have never received a hostile response. I assume that this means that hostile users don't care about my little complaints, but I also do my best not to make a complaint to a potentially hostile user.

Is the Abuse Contact Trustworthy?

When deciding whether to file a report and to whom I will send it, I always try to assess the trustworthiness of the abuse contact for the problem system. Trustworthiness refers to the likelihood that an abuse report will be given to a person with hostile intent, which may inadvertently help an attacker instead of preventing further problems. There is no absolute knowledge of the motivations of admins of other networks, but I use a few rules of thumb to make my best guess.

If I decide the abuse contact is trustworthy, I report to them. If they do not appear trustworthy I report to their upstream provider, or I do not report at all. If the contact is not trustworthy and the abuse continues, I report it to law enforcement agencies and block traffic from the offending address. Law enforcement generally does not assign a high priority to these reports, but it may help them to catch someone who also abuses other networks. It may also contribute to the detection of someone who is involved in something much worse than the activity you noticed.

The following contact people are likely to be trustworthy:

- Major ISPs;
- Reputable companies;
- School administrators in generally lawful countries; and,
- Government agencies of a generally lawful country.

The following contacts should be considered not likely to be trustworthy:

- Admin contact with @hotmail.com or other unverifiable address outside of the target domain;
- Unknown businesses that cannot be identified as safe with minor research;
- Individual dial-up or broadband users;
- Contacts from unknown institutions in countries that have reputations for widespread corruption or disrespect for international law; and,
- Any contact I do not think I understand.

Creating Successful Abuse Reports

Reports can generally be handled via e-mail, but this may not always be the best choice if an immediate response is required. If the unwanted traffic is causing urgent problems for your network, the telephone is sometimes the best way to achieve quick resolution. Whether by phone or e-mail, you will need to start by conveying the basic information to the responsible handler.

Keep initial reports professional, courteous, short, and above all informative. In order to assess whether or not a report satisfies these criteria, pretend you are the recipient, that the report is one of three hundred newly arrived messages, and that you have five seconds to scan the piece of e-mail to decide whether it is worth your attention. Within that five seconds, you should be able to establish the following:

- What is the complaint?
- Who caused the problem?
- Who was bothered by this?
- What do you expect us to do about it?

After addressing these key points in the opening paragraph, get straight to the details that will allow the handler to perform their own analysis of what your logs show. Avoid practices that delay response, cause problems with automated analysis, communicate unnecessary hostility, or are simply irrelevant.

Techniques that Help a Report to Succeed

The following list of dos and don'ts will help the filer of the report to ensure that the reports will succeed.

- Do include the full log entries or packet captures in the body of the message;
- Do spell out the time zone and offset from Greenwich Mean Time (GMT) used in your logs;
- Do include log field definitions;
- Do summarize the source and destination IPs of the abuse and destination port;
- Do use friendly and professional language;
- Do assess the trustworthiness of the people you are considering notifying;
- Do notify law enforcement if there is a successful break-in, or if there is strong evidence of serious efforts to compromise your system;

- Do CC the upstream provider of small networks;
- Do use plain text rather than HTML e-mail;
- If requested, do include alternate means of contacting you – particularly the appropriate telephone information;
- Do "obfuscate" or hide your own IP address information in copies sent to unconcerned parties (like security mailing lists);
- Do include the name of the WHOIS service that associated the offending IP address with the administrator you are contacting. WHOIS information is sometimes incorrect; and,
- Do file a separate report for each source address that is sending unwanted traffic.

Techniques to avoid include:

- Don't remove or obfuscate target IP addresses in the message to the abuse handler;
- Don't quote from the provider's abuse policy;
- Don't include analysis of the abuse unless it is critical for the handler's understanding of the situation;
- Don't include legal opinions or quotes of statutes;
- Don't include attachments unless requested;
- Don't threaten prosecution or retaliation;
- Don't include WHOIS lookups - the provider is aware of their netblocks; and,
- Don't report Nimda or Code Red infections to major providers unless they are causing serious traffic problems for you. Large providers have either automated their approach to controlling the worms or they have simply given up.

In some cases, people may be issuing reports to ISPs that do not communicate in the same language. It is important that you approach these ISPs in a manner that communicates respect and clearly communicates the situation. This can be tricky. In order to be respectful, you should communicate the report in the language of the ISP. If you do not speak the language, this may require that you either have a friend who speaks the language check it or, as a last case scenario, use an on-line translator.

In order to minimize the risk of miscommunication, you should also include the report in your native tongue or in English - for better or for worse, English is the *lingua franca* of the Internet, and more people working on the Internet understand English than any other language. It will also be helpful to introduce the translation by stating that you do not speak the language of the ISP and that the translation has been run through an automated translating service.

Don't be Discouraged

The first response to an abuse report is often an automated letter indicating that your complaint has been received. It may feel like your complaint has been ignored, but this is usually not the case.

Read the response carefully - it will often contain very useful information such as "all abuse complaints must be filed by visiting the following URL" or "your complaint has been forwarded to the administrator of the system in question." In the first case, if you did not read the response, you would not know that your report would be ignored unless it was submitted according to the handler's procedure. In the second case, you might miss the fact that XYZ ISP had inexplicably forwarded your entire complaint, with full source and target information, to the evil blackhat who had been attacking your system - giving him a roadmap of your ability to detect his activity against your system.

Follow-Up and Escalation

Sometimes an abuse handler will request more information, such as any new log entries that may show continued abuse. They may also request that you re-provide your logs in a different format, or may appear to ask for exactly what you gave them in the first report.

Do your best to give them exactly what they are asking for. If they have taken the trouble to ask you for anything at all, it means that they have read your report and that they have some interest in acting upon it. Some handlers appear to (and probably should) use these requests as a screening tool. If you do not provide the information they request, they may rightly decide that you have not attached much importance to the notification.

Typically, abuse handlers will respond within a day or two with their summary of the action taken as a result of your report. If you receive no response within one week, send a follow-up notice. If there is still no response, you may get better action by notifying their upstream provider. Upstream providers can be extremely helpful in pushing a smaller provider to resolve complaints, because they do not want their abuse departments to be burdened with complaints that pertain to a network that should be managing its own problems.

If upstream ISPs are not helpful enough, a CERT (Computer Emergency Response Team) may be able to help you escalate your request. National CERTs are operated by many countries

throughout the world. They know the resolution procedures that work best within their culture, and can often relate your problem to other ongoing incidents. They also offer expertise in local law, languages and customs. Your own nation's CERT can also be a valuable resource, even if the unwanted network traffic comes from another country.

When All Else Fails

Sometimes an address or range of addresses may continue to send unwanted traffic to your network despite repeated contacts with the responsible abuse handler, upstream handlers, and all other resources at your disposal. In some of these cases, you may decide that your best choice is to block all traffic from the offending netblock. This is generally considered to be an extreme action - your network will lose any legitimate traffic from that address range, including contact from the administrators of the offending network. Firewall or router performance can be significantly degraded by blocking many address ranges, possibly resulting in self-inflicted denials of service.

Before making this choice, be sure to carefully consider whether you have performed with due diligence and eliminated other more reasonable options.

Some questions to consider are:

- Is the abuse contact trying (but failing) to help you, or is the contact ignoring or aggravating the problem?
- Are there any other administrative contacts who could be helpful in reducing the unwanted traffic?
- Are you positive that the traffic was generated by hostile activity?
- How much danger does the traffic pose to your network?
- Is escalation through law enforcement warranted?
- What business or personal contact will be lost if traffic is blocked?
- Should the addresses be permanently or temporarily blocked?
- What is the smallest range of addresses that should be blocked to protect your systems?

A temporary block can help to protect your network from the immediate threat of hostile traffic, but can complicate any effort to work with the abuse handler to eliminate the threat at its source. Make sure to provide alternate contact information to the handler if you have deemed them to be trustworthy.

If you decide to permanently block all traffic from a particular address range, it is usually best NOT to notify the administrator of the offending network. Blocking is considered an insult to a friendly administrator, and a challenge to a hostile administrator.

What Can Be Considered a Successful Conclusion?

Small providers and end-users will often tell you the specific action that they have taken to resolve the cause of your complaint. Large providers rarely provide much feedback about actions taken as a result of your report, because of subscriber privacy policy or because of legal concerns. This does not reflect on the relative success of the reports to either one. Success should be measured primarily in terms of the protection of your own systems: if you have improved your security without losing important functionality, you have succeeded. Some specific successful resolutions are:

- The handler thanks you and says that the offending computer has been taken off-line while a configuration problem or system compromise is repaired. You receive no further unwanted traffic from the system.
- You never hear anything at all, but you receive no further unwanted traffic from the offending system.
- You decide to block all traffic from a particular system or network because abuse continues despite your reports. You feel confident that you are blocking hostile traffic without losing important legitimate communication.

Ideally a successful abuse report has made your network, and a little piece of the Internet, a slightly safer place. But don't expect hostile traffic on your network to be eliminated by one report, or by a hundred, or by a thousand. The Internet is a big world and not all of it is friendly.

The author is Assistant Director of Information Technology and security officer for Patton Harris Rust & Associates, pc, a civil engineering firm headquartered in Chantilly, Virginia.

[Privacy Statement](#)

Copyright 2006, SecurityFocus