ISC Internet Systems Consortium

Blackhat DC 2010

Title: Whose Internet Is It, Anyway?

Presenters:
Andrew Fried, Researcher, Internet Systems Consortium
Ben Butler, Director of Abuse, GoDaddy
Richard Cox, Chief Information Officer, Spamhaus

Introduction

In 1939 the movie “The Wizard of Oz” was released. Many of you will remember this dialog:

> **Dorothy**: Do you think there could be wild animals in here?
> **Tin Woodsman**: Perhaps.
> **Scarecrow**: Even ones that, that eat... straw?
> **Tin Woodsman**: Some, but mostly lions and tigers and bears.
> **Dorothy**: Lions?
> **Scarecrow**: And tigers?
> **Tin Woodsman**: And bears.
> **Dorothy**: Oh my!

Fast forward 71 years. Its now 2010, and when going out into cyberland, we once again find ourselves having to worry about lions, and tigers, and bears. Yes, the Internet has become a very dangerous place. Every person stepping out onto the Internet gets exposed to viruses, trojans, denial of service attacks, child pornography, corporate espionage, identity theft and bank fraud. It’s a cycle that occurs twenty-four hours a day, seven days a week, three hundred and sixty five days a year.

So how did things get this bad? Who’s doing something about it? Why isn’t the government protecting us? Who really is responsible for the Internet? These are good questions that we’re going to discuss.

Who Owns the Internet?

The Internet began as a U.S. government project that was primarily focused on building a network that was resilient to physical disruptions caused by downed links. It later migrated to Universities and finally became available to the general public in the early 1990’s. By the mid

1990's the Internet experienced explosive growth worldwide. Today, the numbers of people with Internet access is in the billions. We speak different languages, live in different countries, and yet can transparently communicate with others thanks to the magic of TCP/IP.

There is no one "internet". The Internet actually consists of tens of thousands of independently owned and operated inter-connected networks. The various networks talk to one another through private peering arrangements and commercial interchanges. Telecoms, cable companies, Internet providers, government agencies and corporations own the various components, but no one government or corporation owns the "Internet".

Due to the international aspect of the Internet, no one country's laws govern its use. So the Internet really isn't owned or managed by any one government or company - it's components are privately owned, privately managed and privately run. It's an excellent example of a working democratic anarchy.

Standards (RFCs)

Since no one group or country controls the Internet, how to all the various components seem to work seamlessly with one another? Standards! RFCs, which stands for "Request for Comments", to be precise.

According to Wikipedia:

> *In computer network engineering, a* ***Request for Comments*** *(**RFC**) is a memorandum published by the Internet Engineering Task Force (IETF) describing methods, behaviors, research, or innovations applicable to the working of the Internet and Internet-connected systems.*

In order for one network to successfully talk to another, standardization was necessary. This standardization was based upon RFC's, which essentially consists of hundreds of different "standards", all reached by consensus of engineers, scientists and researchers. Failure to adhere to the standards results in a "failure to communicate". If you want to be part of the community you must conform. The RFCs form the basis for the Internet's "Common Law", dictating everything from the format of a packet to the standards for email. There are even RFCs for processes and best practices.

The RFCs are the only globally agreed upon rules governing how the Internet I supposed to work.

Law & Order

Let's digress for a moment and discuss the law in nebulous terms. One of the first concepts many criminology students learn is that law and order are two diametrically opposed concepts. Order relates to the ability of the government to deal with proscribed behavior, while law restricts how the government goes about enforcing that order.

A second criminology concept theorizes that laws must be generally accepted before they can be enforceable. Lets take a simple analogy. When you drive down the road, you stop at stop signs and red lights because the law says you must. And because of that, you do not have to stop at every intersection, because we assume that anyone approaching upcoming intersections will also obey the law and stop, yielding right of way.

But, suppose ten percent of the population ignored those laws and routinely ran red lights and stop signs. You could no longer assume that it was safe to drive down the road without stopping at almost every intersection for fear that one of ten cars would plow into the road and cause a collision. The result is that traffic becomes arduous, congested and “order” breaks down. There would be so much disorder that law enforcement would be unable to deal with each and every violation. The “system” would get overwhelmed, order would diminish and enforcement would become highly selective.

Within the Internet world, we find that the concept of law and order is further blurred. What happens when someone from one country commits an act against the resident of another country, and the act they commit is not illegal in their country?

If someone sitting in New York city accesses a website in the UK that contains sexually explicit material, no law is broken. But if someone were sitting in certain Middle Eastern countries and did the same thing, they would have broken their local laws. What happens when someone in NY bounces off a server in that same Middle Eastern country and accesses the same porn site. Did the person in NY commit a crime?

Sometimes what’s right and wrong on the Internet is subject to “local interpretation”.

This reminds me of an old story. Late one evening a lady calls the police to report that a man is exposing himself in public. An officer is dispatched to the lady’s’ residence. Upon arrival he enters the house and is led to her bedroom. He asks the women where she saw the man. She opens her bedroom window and points to a faint light in what appears to be a wooded area behind the home. The officer advises that he doesn’t see anything out back other than the lights from a house in the far distance. The woman hands the officer high powered binoculars and says, “Try looking now!”

Investigations involving illegal computer acts require massive investigatory investments in time and resources. Some countries simply don’t have the technology or resources to combat online crime, while others turn a blind eye due to the fact that the computer crimes are bringing in wealth to the community.

When we go onto the Internet, we assume that everyone will obey the laws we were taught and assume that those that break the laws would be arrested and prosecuted. Unfortunately, this is not the reality of the Internet. While there have been successful investigations and prosecutions of those that have committed crimes on the Internet, the percentage of prosecution to crimes would be several decimal places to the left of one percent.

From a statistical standpoint, the probability that a miscreant will be identified, investigated, arrested, prosecuted, convicted and sentenced for committing crimes against others on the Internet is so small that is poses absolutely no deterrent effect whatsoever.

Clearly, the laws have failed to stem the tide of criminal behavior on the Internet – a trend that is not expected to change any time soon. There is little law on the Internet, but there is an increasing effort by non-government entities to take measures to maintain order on the Internet.   Private industry, researchers, Internet service providers, universities, community activists, network vigilantes and a myriad of others conduct these efforts.

Less encumbered by the “law and order” rules that governments must adhere to, security researchers and industry have begun to band together to fill the void created by the lack of effective “order” on the Internet.

Who’s In Charge?

Since no one government owns the Internet, and no one government has enacted laws that are globally enforceable, and no one government controls the RFC process, who’s in charge? The short answers are “no one” and “everyone”.

There is no single omnipotent ruling body that can exert domination over the Internet. There are organizations, such as ICANN (Internet Corporation for Assigned Names and Numbers), that have been formed by private sector initiatives to provide some degree of standardization and conformance. And there are regionally designated organizations charged with divvying up address space and domains. But there is no one organization that can globally enforce conformance with RFCs, much less laws governing identity theft, fraud, theft, etc. That’s the “no one’s in charge” part of my answer.

However, when you connect to the Internet, you take on a certain responsibility and vested interest for protecting your own health and well-being. You have the right to decide which sites you visit, which emails to read, where you shop, what you download and what you send out, so long as you don’t violate the laws of the country in which you are located in. Well, that holds true for most places other than China and a few other oppressive regimes.

Your failure to take due care while perusing the Internet can result in dire consequences – loss of personal identity information, theft of funds from your credit card and bank account, your system can become infected and cease working properly – the list goes on. Unfortunately, this happens all too often. And it happens on a magnanimous scale, all the time, throughout the “connected” world. So the onus of protecting one’s self ultimately falls onto each and every Internet user. You are the one in charge of what you do and what gets done to you. That’s the “everyone’s in charge” part of my answer.

The Concept of Policing

Prior to the early 1800’s, protection of the citizenry was primarily the responsibility of the citizenry itself. If you wanted to feel safe on your property you had to guard it. At some point in time, citizens came up with the idea that they could relinquish that authority to a government organization and have them take on the responsibility for guarding their property and enforcing order. The good thing about that was that people could actually sleep at night knowing that someone else was out there looking out for them.

Sir Robert Peel formed the first modern, professional police force of its kind in London in 1829. An Internet police force is yet to be established with both the responsibility and authority to protect citizens of the Internet.

Security Community

Despite the vastness of the Internet, the security community tends to be a small, close knit collective group of individuals that all seem to know one another. There isn't a lot of inbreeding, but it's a small group nonetheless.

An extraordinarily significant percentage of the researchers are volunteers that tend to have a real $dayjob. Once the sun goes down they emerge from their caves wearing their super hero research capes and begin to seek out badness.

Each researcher tends to focus their primary efforts in a certain direction; specialization is the norm. There are researchers that focus on reverse engineering malware, detecting and blocking spam, looking for sql injection attacks, cross site scripting attacks, phishing, vishing, money mule schemes, dns cache poisoning, etc. My specific contribution to the community focuses on the analysis of lots and lots of dns data.

Hardly anything that goes across the Internet doesn't leave a trail of some sort, and the amount of data generated by "the Internet" is staggering. But researchers are often quite adept at working with large data sets in search of the proverbial "needle in a haystack". What I've always found to be amazing is the variety of perspectives a bunch of researchers can have looking at the same data.

Another phenomenon of the security community is that they tend to work closely with one another, often sharing data, tips and techniques. This enables them to solve issues far faster as a group than each could do individually. It's a team effort despite there being no teams.

The results of their efforts are often transparent, yet affect every Internet user. For instance, a very significant percentage of all emails that are sent pass through RBLS created and maintained by Spamhaus. What many don't realize is that those lists are built through some amazing data analytics that require the processing and analysis of tens of millions of emails that are cultivated from a variety of sources.

While some research is used for investigative leads, the vast majority is used to identify problems and document situations that must be mitigated.

Mitigation

As threats have evolved, so has the methodology needed to mitigate them. In the "early days", most threats were handled by the system administrators of compromised boxes. Next up the list were the registered domain name owners. With the demise of accurate whois records, whether overly fraudulent or "anonymous", mitigation efforts were now brought to the registered netblock owner and/or the Internet service provider.

Fast flux botnets changed all of that. Almost overnight the effective point of mitigation changed from the ISP to the registrars of the fast flux domains. Many of the registries and registrars were initially ill equipped to handle timely takedowns of bad domains – they were 8 – 5 operations, with no night, weekend or holiday abuse desk coverage. That has begun to change, but remains a significant problem at the present time.

Additional steps can often be taken to further reduce threats once they are identified. If the threat is advertised via email, the mail server sending the spam can be listed in a number of realtime blackhole lists (RBLs) such as Spamhaus. Links advertised within the spam that direct the victim to a fraudulent site can be listed in RBLs like SURBL, which are used to block emails advertising blocked URLs. Content filter services like Websense can be provided with links that they can add to their block lists. Anti-virus vendors are sent malware samples for analysis and inclusion in their signature files. Email accounts and Phone numbers associated with scams can be suspended or deactivated.

Service providers can deactivate accounts used to distribute files on file sharing sites, network operators can be called upon to null route connections to servers and, in some rare instances, entire networks.

The Art of Effective Takedowns

What's the good of looking for badness if you can't do something about it? One of the very few satisfactions many of us get is to successfully whack a bad site or domain, hopefully preventing additional people from falling victim. Having taken thousands of phishing sites down over the past five years, I've learned a few tricks to expedite that process.

Lets begin with what I hope is a common sense rule here – make sure what you're trying to shut down is involved in illegal activity. You'd think this would be a "duhhh". Be sure you can articulate what they're doing, what kind law they're breaking and what criteria you used to determine the site or domain is bad.

Secondly, realize that a cooperating ISP or registrar incurs some liability by shutting down sites, and that with some exceptions, they are not legally required to shut down a bad site or domain. The fact that they're actually accepting emails to their abuse@ email account and not content filtering that email is a good sign. It's pretty disillusioning to report a phishing site to the abuse@ email address only to have the email bounce because it contained a phishing link. Been there, done that.

Due to liability issues, the ISP or registrar will have to document what they see and justify any actions they take. By doing much of that for them you can significantly speed up their ability to move from the complaint to mitigation stage.

When reporting incidents, assume they really don't look forward to dealing with issues like these, so make it as painless on them as possible. Keep your complaints short and sweet, sticking to the 5 w's – who, what, when, were and why. Omit your opinions, small talk, extraneous observations, etc. Don't cut and paste lots of useless/meaningless data that will cause them to grit their teeth.

If you have a working relationship with a particular IPS or registrar, the next step is possible not relevant – explain who you are and why you're complaining. For instance, "I received this email", or "my customer is a financial organization that is being targeted by a phishing scheme". If you deal with an organization enough times you'll find that your complaints enter the "fast track" processing lane. If the ISP or registrar does not know you, they will certainly need to independently validate any complaints you're raising against their customers.

Documenting a complaint is generally rather straightforward. Begin by providing a link or URL to the illegal content. Provide whois data for the domain, DNS address and PTR records for the server's IP. Attach a screen shot of the offending site preferable in png or jpg format. Do *not* send your complaints in word processing documents – stick to text based emails and attachments. Almost all abuse desks store complaints in databases, and they may not be able to retrieve and view the complaint using your word processor.

If the bad site is related to previous bad sites you referred to the same ISP or registrar, include a list of those sites. Quite often once they look at past tickets and determine the current complaint is related they take immediate action and move on.

If you report a site and it's subsequently taken down, it's always a nice gesture to write back a one-line email thanking them for their efforts. Sometimes we tend to lose our social skills and forgo little amenities like saying thank you. Be nice. Be appreciative.

Conclusion

The Internet has no single owner, no single authority, no single body of laws, no single law enforcement agency and answers to no single governing body. Absent all that, the primary responsibility for protecting Internet travelers from the massive number of attacks that take place every day are the firewall manufacturers, content filter developers, anti-virus companies, developers, security researchers, universities, Internet service providers and organizations like SURBL and Spamhaus. Yes, missing from that list are the various governments; security is a grass roots effort, at best.

As the bad guys have continued advancing their skills, so have the security researchers. It's become a "Spy vs Spy" race to outdo one another. Or, as I like to describe it, it's a cat and mouse game, and we're the mice.

Security researchers have made tremendous progress over the past several years in identifying sources of data and fine-tuning their data mining techniques. It is becoming increasing difficult for the bad guys to conduct business illegal schemes on the Internet without being detected.

Once schemes are detected efforts are directed towards mitigation, much of which centers around the proverbial "whack a mole" method; bad sites are taken down, domains get "whacked", email addresses get locked. In some cases, victims are notified, but anyone having done that a few times quickly learns that the victims almost always believe the person that's calling them is the bad guy. Many of us no longer attempt to contact individual victims.

There have been some larger success stories such as the de-peering of McColo, an ISP that catered to criminals and failed to act on abuse complaints. But most of what security researchers strive for the detection and mitigation of the immediate threat, followed by the thunder of law enforcement rushing to investigate and prosecute. Yes, we're an optimistic bunch with unrealistic expectations.

Significant advances in mitigation capabilities are needed, particularly when schemes originate from so-called "weakest-link" countries, the two most prominent of which are Russia and China. As the problems from those countries continue unabated there is increasing consensus among those in the security industry that nothing short of completely blocking network traffic in and out of those countries will ultimately bring about the cooperation needed from those governments.

While there have been criminal prosecutions for various crimes that have been committed on the Internet, the percentage of successful prosecutions in relation to the overall number of crimes committed is so small that enforcement does not represent a deterrence whatsoever. The miscreants operate with the perception of complete impunity. Statistically speaking, they have good reason for that perception.

The only winning strategy for curbing online crime would necessitate significant changes in the way credit cards are used and accepted, along with better protections for ACH transfers. Shutting down the "carding" industry would instantly disrupt a significant percentage of the schemes being perpetrated on the Internet every day.