COMBATING CYBERCRIME

Principles, Policies, and Programs

MICHAEL BARRETT, ANDY STEINGRUEBL, BILL SMITH | APRIL 2011

By Michael Barrett, Andy Steingruebl, Bill Smith | April 2011

Introduction

The world has seen a remarkable transformation in commerce with the advent of Internet based companies. Goods and services are routinely purchased and delivered electronically leading to significant changes in industries like journalism, travel, and banking. Online payments (eBilling and ePayments) cut across all industries and are being used by a significant portion of U.S. households. For example, eBay reports that literally millions of people make their entire living solely on the eBay platform.

Clearly, a large portion of the population relies on the Internet, either directly or indirectly, for an ever-increasing set of services. Seemingly nothing can slow this trend, with the possible exception of some catastrophic failure. It is unlikely that the Internet as a whole will experience such a catastrophic technical failure, and in fact the authors believe this borders on the impossible. What *is* possible, and perhaps even likely should current trends continue, is the perception by Internet users that the Internet is unsafe and therefore unsuitable for everyday use. Should this perception become widespread, crowd psychology could take hold and as with the recent world financial crisis, result in a loss of faith in "the system". Certainly there were very tangible and real issues behind the financial crisis, but the long-lasting impact has proven to be the perceptual shift resulting in the Great Recession.

The authors believe that cybercrime, and other cyber issues are the one area that could cause this type of loss of faith in the safety of the Internet.

This paper begins with the premise that cybercrime is slowly getting worse, and that technical measures alone, while necessary and helpful, cannot significantly move the trend line in a positive direction. There is much data to support this position, and we don't think it's helpful to simply regurgitate that here. We believe action is needed to counteract this negative trend, and we present arguments in favor of a multi-faceted regulatory approach to dealing with the problem, as the only viable way to proceed in the long-term.

While this is an issue of global importance and will require international cooperation, this paper takes a fairly U.S.-centric viewpoint in order to demonstrate the range of concrete steps that must be taken to significantly mitigate the impact of cybercrime. It is however written by a company with a global footprint

and a keen understanding of the international issues. We believe that many of the ideas presented herein can be applied globally and that national solutions are not enough. International cooperation is essential for the continued health and vitality of the Internet.

Please note that this document does **not** consistently define acronyms and technical terms inline, but does have a glossary in which all of these definitions may be found.

Distinction between types of cyber issues

One of the many problems with discussing issues in cyberspace, is that a number of issues are conflated under the 'cyber' banner. The difficulty is that the problems grouped into the term 'cyber' are in fact quite different, and the solutions are similarly likely to be different. Without precision about which issue is being discussed, it is very easy to talk at cross-purposes. In this paper, we've chosen to use a particular categorization scheme (largely suggested by Scott Charney[1]), which we believe clearly delineates between distinct issues.

Cybercrime

Cybercrime is criminal activity, using computers and the Internet to steal, whether directly or indirectly, from consumers or businesses. The global scale of the cybercrime problem is not known, although by most estimates it is running at several billion dollars (USD) per year. This paper focuses solely on ensuring that cybercrime is contained. Even within cybercrime, there are different subcategories, such as direct theft of money from bank or credit card accounts, identity theft, as well as theft of intellectual property where the financial costs are indirect (potential loss of sales revenue). In this paper we focus primarily on direct forms of cybercrime, believing that these are the more pressing problem.

Cyber-espionage, by individual or groups

Cyber-espionage is defined as the process of hacking into computer systems in order to steal information. Individuals or group actors might do this either for 'kicks', or because they believe that the information might have commercial value, or should simply be available to the public. There is also a long history of unscrupulous companies spying on competitors and individuals with the term 'industrial espionage' most commonly used in this context.

1 http://go.microsoft.com/?linkid=9746317

Cyber-espionage, by state actors

Cyber-espionage by state actors is extremely difficult to differentiate functionally from cyber-espionage by individuals or groups. The attack techniques are likely to be quite similar. The primary difference is the motivation behind the intrusion (which is most likely to be politically motivated) as well as the possible range of reactions if the intrusion is publicized.

Cyber-terrorism

Cyber-terrorism can be reasonably defined as attacks against one or more parts of the Internet, intended to prevent legitimate users from being able to use Internet-based services, to instill fear that the integrity of services has been compromised, and most importantly to engender fear in the power of the group behind the attack.

Cyber-warfare

The difference between cyber-terrorism and cyber-warfare is threefold: intention, scale, and actor. The intention in a full-scale cyber-war is likely to cripple the target (economy, communications, essential services), or to create confusion prior to or during an actual (kinetic) war. Under this definition, the only cyber-war which has actually occurred to this point, was coincident with the Russian invasion of Georgia. In the case of cyber-warfare, actors are likely to be either under direct state control, or acting with close coordination of the state.

The Problem

Cybercrime, like other forms of crime, is a multi-faceted and ever-changing problem. Our purpose is not a full and thorough examination of the problem space but rather to propose a set of principles and actions which we believe will have significant impact if adhered to and implemented. To give background for our proposals, we introduce several specific issues in the space and briefly discuss them.

Malware and Insecure Computers

Today much of the harm that occurs on the Internet is a direct result of malware-compromised end-user computers. Criminals organize these compromised machines into botnets which are responsible for the growth in spam, phishing, and distributed denial of services (DDoS) attacks.

Malware is typically categorized into two groups; that which exploits a security flaw on a computer, and that which relies on user action to be installed (e.g. rogue Anti-Virus programs). Our technical and policy solutions in this paper focus on malware that exploits technical vulnerabilities rather than that which requires social engineering. While we believe social engineering is a significant problem, non-technical solutions can mitigate the problem. These same solutions prove ineffective against malware that exploits technical vulnerabilities.

Obstacles to Effective Law Enforcement

The difference between the effectiveness of law enforcement in the physical world and on the Internet could not be more striking. Whereas in the real world even minor crimes such as vandalism and burglary resulting in relatively low dollar losses merit at least a visit by a police officer, online crimes exceeding $25,000 (USD) frequently go uninvestigated, much less prosecuted.

We believe there are a number of factors behind this lack of effective law enforcement online:

1. Insufficient funding for cybercrime law enforcement
2. A lack of trained cyber experts within law enforcement
3. A lack of effective international cooperation and data sharing
4. A lack of universality of laws against cybercrime
5. Statutory minimums in cybercrime cases hamper effective enforcement

Obstacles to Private and Public/Private Cooperation

Another set of problems which must be addressed are those laws, regulations, and policies whose net effect is to hinder investigation and policing of cybercrime with no corresponding public benefit. In several jurisdictions, policies such as “Network Neutrality” hinder ISPs and other network providers from acting to eliminate criminal traffic from their networks given the risk of running afoul of network neutrality regimes. Even when laws do not specifically preclude action, conservative interpretation by companies may act to limit otherwise positive behavior.

Privacy laws and policies can also prove an unintentional hurdle to effective private action. Under some interpretations of privacy laws such as the Electronic Communications Privacy Act (ECPA) companies that detect criminal activity on their networks and systems are prohibited from voluntarily sharing information with other entities in order to prevent further criminal activity. For example, some

companies express significant concern with sharing non-redacted spam and phishing mail feeds, for fear of inadvertently violating their customer's privacy rights under ECPA.

As currently interpreted some of these privacy laws serve only to immunize illegal actions from further scrutiny while failing to meaningfully protect any privacy rights.

Individual Rights and Obligations

As mentioned above, malware is a significant threat to the Internet ecosystem and its prevalence is increasing. Fortunately, there are well-known techniques and mechanisms that mitigate its delivery and impact. Unfortunately, there is no requirement that these techniques be employed, either by consumers or other Internet actors (e.g. websites).

Many Internet users are simply not aware that their actions, or lack thereof, have an impact on others and the Internet itself. This lack of awareness can be traced to a lack of education, from grade school through graduate level programs in Computer Science. Computer (Internet) security simply has not been part of our curriculum, even when computer literacy has been.

Additionally, there is a tension between an individual's right to access information and services online and their obligation to employ proper safety techniques designed to ensure their and others ability to access information and services. Similar tensions are seen in the U.S. regarding freedom of movement, speech, and assembly contrasted with public safety requirements such as permits for large public gatherings, and safety inspections for automobiles. Surely a balance can be reached which allows individuals to access the Internet, while still requiring the systems used to perform that access do not put others at risk.

Unreliable Data on the Size and Scope of the Problem

Estimates of the magnitude and scope of Cybercrime vary widely, making it difficult for policymakers and others to determine the level of effort to exert in combating the problem. Individual companies and some industries may have specific information which could prove valuable in answering the following questions:

- How much money is being lost?
- Where is the money going?
- Can the responsible parties be held accountable and prosecuted?

Perhaps surprisingly, to the extent that this information even exists, it has never been aggregated and collated in a helpful way.

Models from prior history of technological innovation

Other forms of technical innovation have followed a maturity curve much like the Internet today, albeit at slower speeds. These innovations were coupled with attendant public policy, self-regulation, and public reaction that are instructive for understanding models that we might choose for better Internet regulation. We have considered the following in formulating our principles and proposals:

- Public Health
- Law of the Sea
- Aviation Law
- Automotive Regulation

We believe these are useful models because they each have numerous actors, jurisdictions, and tension between individual right and collective protection.

A call to Action - Coordinated ecosystem change

We believe that in order to credibly make an impact on cybercrime, a series of changes must be made which impact the problem in multiple dimensions. It is quite unlikely that a single set of changes, however well-intentioned or sweeping, is likely to have the desired effect. In short, there is no single solution that will solve the problem that is cybercrime.

Principles

In developing a series of recommendations for change, we have tried to lay out a realistic set of principles that should underpin them. Some are practical, others are more conceptual, all condition the types of changes that could be made.

1. **Involve the least regulatory change needed to accomplish appropriate levels of safety.**

   We believe that this principle is self-evident: do no more than needed in order to make the Internet safer.

2. **Ensure that laws can be interpreted in ways which credibly allow participants to prioritize safety.**

   Many laws that apply to activities on the Internet, such as the U.S. Electronic Communications Privacy Act (ECPA) were written long before the Internet existed at the scale and scope it does today. As a result laws like ECPA often have basic scope problems that cause unintended consequences. In the specific case of ECPA, we are familiar with cases today, wherein ecosystem participants choose to do the 'wrong thing' and not protect other participants, because their General Counsel's Offices claim that the law does not allow them to act to protect the ecosystem, and even their own customers. Laws that prevent companies from sharing data to protect themselves and their customers from harm with no corresponding increase in customer privacy or security need to be adjusted to take account of modern ecosystem reality.

   Law and regulations need to make it easy for companies to do the right thing to protect themselves and their customers without fear of repercussions. When a law only serves to protect criminal interests in certain cases, it must be adjusted – laws should protect law-abiding citizens from the depredations of criminals, not the other way around.

3. **Make changes which reduce negative externalities in the overall ecosystem.**

   If your device becomes compromised by malware, the negative effects may be felt by other parties rather than you. Consequently you have no direct incentive to properly protect your device. If you plug something into the Internet, you need to make sure that it is up-to-date and reasonably safe.

4. **Accept that the Internet is global - change is needed in every country, using compatible conceptual frameworks.**

   It is a truism, and perhaps a trite one, that the Internet is a completely global communications vehicle. As such, if the law and regulation in any one country or region becomes stronger against cybercrime, it is likely that crime will migrate to jurisdictions which have ineffective laws against cybercrime, and where the chances of detection, prosecution and conviction are low. This has already happened, and while hard data is difficult to come by, it is clear that much cybercrime originates from such jurisdictions.

5. **Avoid attempts to conflate other related issues, such as: intellectual property theft, free speech rights, privacy, etc.**

   We are firm believers in the development of a firm and equitable Internet ecosystem in which the rights of all participants are appropriately protected. However, to this point, the debate over cybercrime has largely been dominated by concerns about theft of intellectual property, rather than direct theft of actual money. While we are sympathetic to ensuring that these rights are protected, we would be greatly concerned if a focus on intellectual property instead of theft of money bogged down progress against cybercrime, while a "grand unified cyber-security" approach is debated.

6. **In general, governments should not mandate nor manage technical controls.**

   While it will be necessary for governments to demand that certain kinds of behaviors are adopted, the authors believe it is quite dangerous to mandate specific technical solutions. While governments have successfully mandated certain types of technology, for example brakes in transportation or vaccines in public health, those mandates addressed relatively static problems.

   In the case of cybercrime, we are facing a dynamic, intelligent, and determined adversary, who has no interest in whether his actions are legal or not, and who has already demonstrated very high levels of adaptability. If governments, through technology mandates, slow down the development and application of new defensive technologies, the attackers (criminals) will adjust their behavior and methods to avoid the mandated technologies, or perhaps exploit vulnerabilities in them to our overall detriment.

7. **Find solutions which improve security, without compromising privacy.**

   Too often, the cybercrime debate is framed in such a way as to imply that privacy is the only goal. In information security circles, it is generally believed that privacy cannot be achieved if a system is insecure. If we design systems to attempt to maximize the privacy of participants, but handcuff the system designers such that they cannot protect participants from criminal actors, then we have not in fact helped the cause of privacy at all.

8. **Accept that full anonymity on the Internet is infeasible for all situations in today's environment.**

   As something of a corollary to principle #7, it is worth explicitly stating that completely anonymized e-commerce is no more possible than completely anonymized "real world" commerce. For example, mail order/telephone order transactions would not be possible (because of excessive fraud) if catalog vendors were not allowed to make fraud decisions based on telephone number and address of the recipients.

9. **Treat data usage for anti-fraud/crime purposes as distinct from data usage for marketing purposes.**

   We are relatively neutral on the question of whether organizations should be able to monitor individuals and their usage of computers (whether by cookie tracking, IP address, FSO or other technique), for marketing purposes. However, we *strongly* believe that if regulation does restrict the use of this data, there should be carefully constructed carve outs for usage of this data in an anti-crime/fraud context[2].

10. **Organizations that perform Internet Governance are part of the solution, not part of the problem.**

    Internet Governance has moved from an obscure issue reserved for discussion by the cognoscenti to a meeting agenda item at a number of international bodies. Certainly, issues exist with the governance of the Internet today, but issues exist with all forms of governance. It is our belief that these issues have been exaggerated or distorted beyond recognition for reasons that are at best unclear. A significant level of largely theoretical angst exists about the U.S. government's residual involvement in the management of the Internet.

    Some parties are playing on this angst to recommend that governance of the Internet should be moved under the auspices of the UN, and specifically the ITU. We believe this is a poor idea, primarily because while the ITU has experience managing stable technologies, that far from describes the Internet with its rapid evolution. Current mechanisms may not be perfect and could be improved, however, it is not clear that centralizing management within a relatively bureaucratic and slow-moving global organization is helpful. Rather, it is our belief that governance of the Internet should be modeled on its architectural underpinnings.

2 http://www.iab.org/about/workshops/privacy/papers/brett_mcdowell-revised.pdf

Initiatives that could improve Internet safety

As we have noted previously, there is no single initiative, whether technical or policy, that can on its own 'secure the Internet'. However, taken holistically, we believe that a package of the following items would make the Internet significantly safer.

Data Collection and Reporting

Build the Internet NTSB

The aviation industry has gone through a number of changes wherein additional regulation has been imposed over time. One of the most significant changes which improved safety was the establishment of the National Transportation Safety Board. While the NTSB has no direct ability to issue regulation, what it has done is to establish a very large database of accidents and their causes, and issue recommendations to the Federal Aviation Administration. The FAA has therefore been able to use the NTSB's data as a very effective method of slowly but continually eliminating various causes of accidents.

We believe that cybercrime is lacking the equivalent of the NTSB database. For example, there is widespread disagreement about the scale of direct losses from cybercrime. While there is good agreement about a range of countries that are known to be troublesome, there is no agreed ranking of which countries take in the most money from cybercrime. Similarly, there has apparently been no assessment of which countries have meaningful anti-cybercrime legislation, or which countries actually attempt to prevent cybercrime.

While the recommendations below will certainly help reduce cybercrime generally, obtaining the kind of data referred to above will make it substantially easier to recommend specific actions and to use the appropriate political pressure on countries that are ignoring the problem.

Law enforcement

Substantially increase investment in law enforcement of cybercrime

Fundamentally, we believe that the relative investment in law enforcement for cybercrime is too low, as compared with the investment in law enforcement for regular crime. While it is undoubtedly true

that there are financial thresholds imposed in prosecutorial decisions for regular crime, these are typically quite low. Thus, in most cases, if an individual steals, say, a smart phone worth $500 USD, he or she is indeed quite likely to get prosecuted. (Assuming the individual is caught or identified, of course.)

However, in cyber-space, if an individual steals $20,000 USD of smart phones online, and is successfully identified, it is questionable as to whether he or she will be arrested/charged/prosecuted. Most agencies which investigate and prosecute cybercrime have informal thresholds, below which cases are unlikely to be actioned. While these vary from agency to agency, in most cases, $20,000 USD represents the lowest level at which action may be taken. This is not because Law Enforcement agencies are unsympathetic, but rather because they are simply overloaded with cases and cannot take on any more work.

We believe that there needs to be a significant increase in the funding of agencies which investigate/prosecute cybercrime offenses. We are uncertain of exactly what 'significant' implies in this context (see 'Build the Internet NTSB'), but we suspect that it is something of the order of a doubling or tripling, at least, of investment in this area. In practice, it is relatively straightforward to tell whether the increased investment is enough, by monitoring the number of cases that agencies are prosecuting where the financial damages are below $1,000 USD.

The authors are occasionally asked how such increased Law Enforcement funding should be paid for. We believe the case for additional resources will be simply made once better data is available regarding the scope of the problem.

Network regulatory oversight

The FCC should incent ISP's to notify customers of malware, modeled on Australia's AISI

The Australian Communications and Media Authority (ACMA), the network regulator in Australia, has developed a voluntary program in which: 1) ISPs sign up for the voluntary Australian Internet Security Initiative (AISI) program[3], 2) ACMA uses various intelligence sources to compile a list of IP addresses which are apparently compromised by malware, 3) ACMA communicates that list to the relevant signed up ISPs, and 4) the ISPs then communicate with the end-customer that there's a problem with one or more of their PCs.

3 http://www.acma.gov.au/WEB/STANDARD/pc=PC_310317

The program has been in place since about 2007, and has been quite effective. Viewed by the number of ISPs which have signed up, it might not seem very impressive, about 80 out of several hundred. However, the 80 ISPS are the large urban ones, and so AISI now helps protect something over 90 percent of Australian consumers who use the Internet. As such, it has been an extremely effective tool at protecting consumers, as well as a great example of public/private cooperation. Moreover, it has been run on an amazingly small budget – the entire ACMA AISI team is about half a dozen people.

We believe that AISI is an excellent model that should be adopted by all network regulators, such as the FCC. Initially, these programs should be voluntary; however, we also feel strongly that once programs like this have been piloted successfully, they should be made mandatory for all ISPs that offer Internet connection services to consumers and small businesses.

The FCC should ensure that transit providers are incented, if not mandated, to screen criminal traffic such as botnet activity

Most transit providers already have network security programs in place by which they monitor the traffic that passes across their network infrastructure. As such, they frequently know which servers are controlling which botnets, and which PCs have been infected by the botnets. Generally speaking, they can tell this simply by looking at the packet headers (i.e. where the network traffic is going from/to), without any inspection of the actual message content.

The transit providers have technical means easily at their disposal to drop this traffic, which is purely criminal in nature. This is an important point. Carrying this traffic simply means that criminals have the advantage, and legitimate Internet users are victimized.

However, as it stands today, the vast majority of Internet backbone transit providers do not interfere with this traffic. The rationale appears to be that there is some uncertainty whether the necessary legal framework exists for them to do so. We believe that a reasonable interpretation of the Commission's current policy, and the regulations which Congress has imposed, in fact do give the transit providers the necessary freedom to block this traffic. However, we suspect that it will be necessary for the FCC to publicly make some kind of statement in this regard.

It is not clear to the authors whether the FCC would be able to force transit providers to take these steps without additional authority from Congress. If the Commission believes that it already has such authority, then we unhesitatingly recommend that it should act to require transit providers to block criminal traffic. If not, then we feel there is a strong case in requesting the additional authority from Congress to do so.

The FCC should ensure that networks do not allow traffic to exit their networks which perform IP Spoofing

Today, many networks do not impose exit controls at their border routers to ensure that the "from" IP address is in fact legitimate for the network that it is originating from. It is technically easy to ensure that this is the case.

When sending traffic over the Internet, data transmission is split into packets of data. The address header contains addresses such as the "to" address (i.e. where the packet should be routed to), and a "from" address (i.e. where the packet purported to come from). The problem is that this header is constructed by the computer which sent the packet, and it can be forged. In the real world, this would be the equivalent of someone putting a return address on an envelope saying "Santa Claus, The North Pole, YZ 98765".

The problem is that computers can be used in Distributed Denial of Service (DDoS) attacks. In this, the criminal simply arranges for as many computers as possible to send as much network traffic as possible toward some particular target. Criminals generally possess large numbers of computers which they can use for this purpose – a botnet. If a particular criminal doesn't possess a botnet, renting them at quite reasonable hourly rates is easy on the black market. The idea is to saturate the target, such that it is unable to service legitimate requests. Usually, these attacks are followed by a short break and a ransom demand – criminals have been able to make quite a good living by preying on small/medium sized businesses in this way.

In principle, DDoS attacks are easy to defend against – simply block the offending IP addresses which are sending the traffic to your web site. That works well – if the attacker isn't doing IP spoofing.

Thus, enforcing a requirement to prevent IP spoofing outside of an originating network will actually make things much, much better for defenders and much harder for attackers. It will not have significant adverse consequences, as IP spoofing has *no* legitimate uses. Network service providers should follow the relevant IETF recommendations for network ingress and egress filtering.

Education

Substantially improve consumer education on cyber-safety

It is clear, from a variety of sources, that most consumers have little idea how to protect themselves online. There has been both government and private sector interest in this topic, with the NCSA's 2010 "Stop. Think. Connect." campaign being an excellent example. It is clear that the problem is much larger than the scope of work happening today. There are many studies that show the majority of Internet users are both afraid of the risk of using the Internet, and simultaneously don't have the information needed to protect themselves online.

While the educational efforts that are occurring today are good, they are simply not at the scale needed to help hundreds of millions of Internet users. This area needs significantly increased investment both from private industry and government. It is hard to know how much additional funding is needed, but it is quite possible that the right answer is "an order of magnitude" higher.

Introduce cyber-safety education curriculum into public schools

New Internet users are coming online at a very young age. There are a number of studies showing how these "digital natives" are in fact more trusting of the Internet. Unfortunately, this is precisely the opposite behavior needed in an environment that is becoming more dangerous over time.

This problem is more tractable than general consumer outreach, as there are formal channels – i.e. schools – by which this group of users can be reached. All that is needed is the development of a formal safety and security curriculum, and an insistence that this topic becomes one of the core areas taught to students.

Targeted action

Charter an organization to directly attack botnets

Botnets have been mentioned a number of times in this white paper. They are a critical resource to criminals, and it is no coincidence that the rise of cybercrime and the size and number of botnets have been inextricably linked. To the extent that government policy can directly impact the size and number of botnets, it will impede the ability of cybercriminals to operate cost effectively.

In the information security world, there is a statement – “security is hard”. This means that it is hard to design a system that is secure against attack, and it is this characteristic that criminals exploit to get their malware onto their victims PCs. However, botnets are also computer systems and thus are difficult to design to resist attack over their network interfaces. Indeed security researchers have done investigation that shows that much malware, and by consequence the botnets that they bring to life, is deeply flawed. Unfortunately, computer researchers are prohibited from taking any action, by law, as they would be violating statutes which are designed to protect the public from harm.

Certainly we are not proposing that such statutes be abandoned – they are the very basis by which civil society generally ensures that order prevails. But, there perhaps should be ways in which these weaknesses in botnets could be legitimately exploited, for example, to take control of the zombie PCs and remove the malware from them.

We are not proposing that this responsibility should belong to individuals, or private enterprises, or indeed academia. Rather, we suggest that this is the kind of power that should be reserved to a state run organization, and indeed that the state is the only legitimate entity to charter such a function. By not doing so today, nations are abrogating their responsibility to protect their citizens.

In a promising sign, as the authors ready this document for publication in April 2011, the U.S. Department of Justice has started to take action against botnets, and has recently taken action against the Coreflood botnet. It did so via a direct attack against the Command and Control servers which controlled Coreflood, rather than against the clients, but nonetheless we regard this as a very promising start.

International legal considerations

Drag MLATs out of the 19th century, and into the 21st

The European convention on cybercrime has represented an extremely important framework for dealing with cybercrime. However, there are two ways in which it has fallen short.

The convention allows nations to cooperate with each other in investigating cases of cybercrime. It permits one state to request that a second state preserves and supplies the necessary data needed to support some particular investigation.

However, the mechanisms used to request the data are positively antediluvian: MLATs (Multi Lateral Assistance Treaties), and “Letters Rogatory”. In all of the cases where we have worked with multi country

investigations, we have *never* seen a case in which the data has been returned to the requesting law enforcement agency in under three months. Six months is more common, and we have heard of cases where the data has been returned more than two years after it was originally requested.

Given the speed at which cyberattacks move, this effectively hobbles the investigating law enforcement agency and frequently cripples investigations. During this time, the criminals keep victimizing citizens and law abiding organizations.

Clearly, there needs to be some level of supervision, and approval, such that rogue officers (or worse) cannot request arbitrary information from another state, without good purpose. But, in the age of the Internet, most workflow functions can be highly automated. The technology to do this exists, and is readily available. Politicians and law makers should consult their law enforcement organizations who best understand how to fix current practices.

Create an international law enforcement model that allows for prosecution without requiring extradition

The cybercrime convention, at least as it is structured today, suffers from a significant flaw. It presupposes that criminals will be extradited from the country in which they reside to the country which has been investigating the alleged crime, and where the victims reside.

Superficially, this seems reasonable. Alas, there are many countries, most notably, Russia, which do not believe in extradition of their citizens under any circumstances. One can debate the rightness or wrongness of this particular position. However, the authors of this paper are realists, and do not think it is prudent to attempt to convince these countries of the errors of their ways. Rather, we think it better to find ways to work with them.

The authors believe the cybercrime convention should support a dual track model – the current one in which the case is researched by the law enforcement agency in the country where the victims are, and the alleged criminal is extradited to that country for the trial; and a second (new) one in which the case is tried in the country in which the alleged criminal resides.

Today, the likelihood of this occurring seems slim, as some Western politicians currently seem determined to prove that their method is the right way, and that other countries are wrong. The only people they are hurting by this position are their own citizens.

Commercial considerations

There are a number of ways in which private industry can help make the Internet safer than it is today. Many of the concepts described here are not particularly revolutionary, but simply need to be applied at scale. This has never occurred, simply because the necessary pressure to do so has never been exerted. The attached list of ideas is by no means a complete list, but implementing even half of them would have a significant positive effect.

Force manufacturers to enable devices to 'fail safe'

Today, many devices have undesirable security characteristics. As numerous devices are now becoming "Internet enabled", this problem will become radically worse quite quickly.

In general, there are some quite well understood principles by which such devices can instead be made safe:

- Ensure that there is a mechanism by which manufacturers can publish security updates for devices, and have the devices auto-update themselves so that the security issues cannot be exploited.
- Ensure that the default settings of the device are secure.

The second of these is very simple to explain. For example, many residential broadband routers permit remote maintenance. This is a desirable feature for a small minority of users, but if it is enabled with a weak default password ("password", or "admin", for example), then it is incredibly easy for criminals to abuse. Indeed, this is not academic, many such remote control features have been so utilized by criminals. But, it is also very easy to prevent – simply ensure that the remote maintenance feature is turned off by default, and cannot be activated unless the user sets an acceptably strong (and non-default) password.

Today, there is no mechanism to force device manufacturers to do this. This is not true for many areas of commerce where product liability concerns have forced manufacturers to develop safety standards, and to uniformly test their products against these. We need a similar model for Internet connected devices, whether it be consumer routers, TVs, refrigerators, or indeed other consumer devices which are becoming Internet connected.

Force unsupported devices off the Internet

Today, a small but depressingly constant set of users access the Internet through PCs that are simply unsafe. The problem is a little complex, and requires explanation.

On a regular basis, security vulnerabilities are found in software, such as operating systems, browsers, etc. Generally, the vendors of that software then come up with a "patch" that fixes the vulnerability. On most occasions, the vulnerability is specific to a particular version of the software. However, less frequently, structural vulnerabilities are found that apply to all versions of some particular piece of software. While this is not a problem if these versions are supported by the vendor, it is very serious if some version is not supported. If the vulnerability enables remote code execution, it is then catastrophic. This means that the world then knows how to compromise a particular machine, but there is no patch available to protect it.

For example, there are a number of remote code execution vulnerabilities which have been found in various browser functions, but which have never been fixed in Windows 98, because these vulnerabilities were discovered after Microsoft took Windows 98 out of support. Basically, this means that anyone who uses a Windows 98 PC has a very strong likelihood of having it compromised by malware.

Sadly, these users have no idea how much risk they carry by using such PCs on the Internet. If this only affected them, then perhaps one could argue that their risk is their own business. However, just as in public health, this is not the case – malware is a public threat, because infected machines can be used to inflict significant damage on other people and organizations.

We believe that there should be a general principle that machines running unsupported software should not be connected to the Internet. We don't view this as being any different from the notion that authorities should have the ability to ban unsafe or highly polluting vehicles from everyday road use – a fairly noncontroversial position.

Enforcement action against "bulletproof hosters"

There is ample evidence that a small number of network hosting companies target their services directly to the criminal black market. These hosting companies tend to have extremely high numbers of spoof websites, botnet command and control servers, malware download sites, and child pornography sites. Many may solely host sites in these various categories, indicating that they serve

only the criminal marketplace. We do make a distinction between those companies that are simply poor at monitoring/managing their customers, and who are taken advantage of by criminals vs. the bulletproof hosting services which solely serve criminals.

The name itself is indicative of the market these companies serve – their services are "bulletproof" in the sense that they will actively not cooperate with requests to take websites down. We believe the very existence of bulletproof hosting is anathema to the civilized operation of the Internet and is a telling "canary in the coalmine" about the state of affairs on the Internet today.

Unfortunately, these services are commonplace. Historically, their operation has been somewhat tolerated by law enforcement, not because they approve of them, but rather because in a certain number of cases they have been waiting to get evidence from such hosting companies. In general, these companies have not appeared within the scope of consumer protection agencies – who appear to be either blissfully unaware of their existence, or in a regulatory funk, trying to determine whether they have the authority to act. The only exception to this is the FTC in the U.S., which has acted against precisely one company. Private industry has acted successfully against a small handful of them (less than half a dozen), by dint of getting these hosting companies "de-peered" by their transit providers.

We worry greatly that private enterprises have been forced to act in this way. While it is quite defensible on the one hand (as a legitimate act of self-defense), it is also perilously close to "vigilante justice" and if misdirected could cause significant inadvertent harm. Instead, we believe that this goal – providing consumer protection agencies with clear legal tools, and a clear policy mandate to shut down criminal hosting companies – should be one that governments embrace.

Have SLAs for hosting companies to remove phish/malware sites

Today, on a daily basis, our operational teams run into sites that have been compromised and used as phishing "spoof" sites, or malware delivery sites. In order to protect our customers, we ask the site owner or ISP (depending on the circumstances), to get their sites decontaminated. We find extreme variability in responsiveness – anywhere from 30 minutes, with good, competent and responsive ISPs, to never. In most cases, response times are measured in days.

In the cases of the sites which we can never get shut down, they often fall into the bulletproof hosting category discussed earlier. For the sites which take a long time to get shut down, there are a number of root causes:

- Inadequately staffed support – not 24x7.
- Low priority given to abuse complaints (it's the thing that gets looked at last, when technicians aren't busy on other problems).
- Onerously high burden of proof demanded by some ISPs (court orders, for example).

We believe that the hosting industry should develop a reasonable set of principles around what level of proof is acceptable in order to request site shutdown/cleanup, and the timeframe in which it should be realistically expected to be accomplished (we suggest 12 hours as the acceptable maximum).

If the industry isn't prepared to come up with a voluntary set of guidelines, we believe regulators should impose one via statute. Monitoring for compliance with agreed to SLAs is an important component of such as program. Independent reporting and analysis of SLA compliance will quickly show which firms are complying with these requirements, and which are not.

Create safe ways for companies to share information about compromised customers, which are exempt from normal rules

One of the topics which is endlessly discussed in information security circles is "enhanced information sharing" as well as "private/public data sharing". It is fair to say that there are two conflicting points of view on this topic:

1. The criminals routinely share information and data with each other, and it substantially enhances their effectiveness.
2. The "good guys" have been discussing data sharing for literally years, and we collectively have not made a lot of forward progress.

There is evidence that the problems stem from several different factors:

1. Lack of a good privacy framework enabling companies with significant privacy requirements to participate without inadvertently risking the privacy of customer data.
2. Potential issues with vetting participants in the ecosystem, to ensure that only legitimate enterprises are able to participate. Of necessity, these policies need to be fairly exclusionary, which raises various issues.
3. There is a genuine tension between government interests in data sharing (which mostly involve a desire for a one way flow, private->government) and the private sector's (which are more bidirectional).
4. Concerns by companies that data they share with governments will be subject to discovery by Freedom of Information Act (FOIA) requests.

At least some of this could be dealt with by a more well thought out approach to data sharing, and potentially chartering a new organization whose sole aim is to facilitate it. However, we suspect that privacy laws may have to be changed to allow the flow of information into such a data exchange, such that it doesn't expose members to significant liability. The reason that this may be needed is that today, many organizations refuse to share data because their General Counsel has advised that doing so would place the company at legal risk should that data later then be exposed. While we think this view is likely incorrect, it has had a chilling effect on the practical ability of companies to cooperate with each other.

Governance

Ensure that ICANN properly enforces the ecosystem safety initiatives that it is contractually obligated to do

ICANN, the Internet Corporation for Assigned Names and Numbers, already has a set of obligations that it is supposed to enforce as it manages the domain name system. However, ICANN has historically been very reluctant to in fact enforce the obligations it already has on registries and registrars, let alone attempt to impose additional constraints on that community.

For example, given the lack of investment in cybercrime law enforcement, companies are forced to trawl through much Internet data in order to identify criminals prior to handing the cases to law enforcement. To the extent that WHOIS data is inaccurate, it substantially impedes these investigations, and simply enables more crime to be committed. One favorite example of the authors: "<Company>, Ho Chi Minh

City, OH, 12345", which is clearly incorrect and indeed turned out to be criminally motivated. It is also self-evidently false and wouldn't pass automated address verification tests - there's no Ho Chi Minh City in the U.S., and zip code 12345 is in New York State, not Ohio.

"Know Your Customer" is a phrase used extensively by financial services regulators, and it is a principle which applies well in other industries too. Basically, it is reasonable to assert that there should be a positive responsibility on the part of ISPs, registries/registrars to know who their customers are, and ensure that the data about them appears to be reasonable correct/up-to-date – within the limits of "commercial reasonableness", of course.

Historically, ICANN has been reluctant to engage on this topic, although there's evidence of a shift in policy under the current leadership. While this change is refreshing, and definitely a step in the right direction, it is not yet clear that ICANN is indeed properly fulfilling its role in managing the safety of the overall ecosystem.

Conclusion

In this white paper, the authors lay out an entire framework of practical actions that could be taken to reduce the impact of cybercrime, and substantially make the Internet safer. Even if only some of these recommendations are implemented, it will make a significant improvement in Internet safety.

While we're hesitant to name any of these initiatives as being more important than any other, we are occasionally asked "list the three things you want us to do". In general, we list:

- Increase investment in cybercrime law enforcement.
- Start the Internet NTSB.
- Fix the Cybercrime Convention.

Just doing those three things would make a big difference, albeit it would be - to paraphrase the punchline of many a joke - merely "a good start".

We expect this paper to be a first step in a multi-stakeholder and iterative process and approach to making substantial progress against cybercrime. We welcome feedback on our proposals.

Glossary

ACMA: Australian Communications and Media Authority. Roughly analogous to the FCC in the US.

AISI: Australian Internet Security Initiative. An initiative managed by ACMA. It is designed to help ISPs communicate with citizens who are using PCs which have been compromised by malware.

Botnet: a large number of compromised PCs, all running the same malware, and controlled by a single criminal or criminal organization.

C&C server: Command & Control server. The central server(s) which control the operation of a botnet.

DDoS attack: Distributed Denial of Service attack. An attack in which an attacker uses many computers to overwhelm a target by sending as much traffic to that target as possible.

De-peering: The act of removing connectivity from one network to another (literally, its "peer"). In general, this happens because of changes in commercial relationships, but in some cases it can be directed toward "bullet proof" hosting services.

FCC: Federal Communications Commission.

Flash: A programming environment, offered by Adobe Systems, which allows developers to build highly interactive web-based applications.

FSO: Flash Shared Object: a parallel cookie-like mechanism that is managed by Flash executables. Until 2010, FSOs did not follow the same cookie behavior as cookies themselves. This has now been changed, and they now follow the same privacy settings.

FTC: Federal Trade Commission.

ICANN: Internet Corporation for Assigned Names and Numbers. ICANN manages the processes by which a large portion of the Internet 'namespace' is allocated. The actual assignment is performed by registries and registrars, but ICANN oversees that ecosystem. ICANN is now relatively autonomous, but there is still some slight level of oversight by NTIA, which exists within the US Department of Commerce.

ISP: Internet Service Provider. This is a loosely used term within the IT industry. In this white paper, we use it in a single meaning – an Internet connectivity provider to residential consumers and small

businesses. We do not include transit providers, who provide connectivity to large businesses, and who operate the Internet's backbone. Nor do we include hosting providers, which provide hosting services (and typically access to connectivity services).

IP address: See TCP/IP address.

LE agency: Law enforcement agency.

Malware: Software which the user of the PC did not knowingly install (or which was deliberately misleading such that the end-user would install), and which is not acting on the behalf of the end-user.

Patch: An update to a piece of software which will fix a vulnerability.

PC: Personal computer.

Remote code execution: By exploiting some particular vulnerability, it is sometimes possible to force a remote computer to run arbitrary code. It is by far the most serious kind of vulnerability, because in essence if you can do this, you can completely subvert what the user thinks that the computer is doing.

SLA: Service Level Agreement: roughly, an agreement that describes the scope of services that one party provides to another; it tends to be contained as part of a binding legal contract. Closely related to SLO.

SLO: Service Level Objective. A set of characteristics that in combination describe the scope of services being provided. Somewhat related to SLAs, SLOs differ in merely being objectives, and there is often little or no commercial consequence if these objectives are not met.

TCP/IP address: the numerical addresses by which traffic on the Internet is routed. While a hopelessly inadequate analogy, it can be thought of as somewhat akin to either a telephone number or a street address. Internet traffic is generally routed from one computer (with a 'source IP address') to a second computer (with a 'destination IP address'). The topic is extremely complex, and the subject of a large number of books. Interested readers wanting more details are referred to such primers.

TCP/IP: One of a very small number of core protocols on which the Internet is based. We have chosen not to define the acronym here, as it is meaningless without the technical background to understand it.

Transit provider: A company which provides network connectivity to the Internet, generally to businesses and other organizations.

Vulnerability: A bug which has been discovered in either a piece of software, or much worse, a protocol, in which the behavior is outside of what the specification envisaged, but where the bug can be abused in repeatable ways to do something that the software was never intended to do.