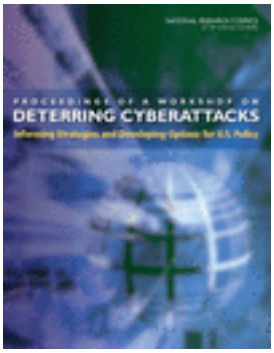

Proceedings of a Workshop on Deterring CyberAttacks: Informing Strategies and Developing Options for U.S. Policy

Committee on Deterring Cyberattacks: Informing Strategies and Developing Options; National Research Council

ISBN: 0-309-16086-3, 400 pages, 8 1/2 x 11, (2010)

Targeting Third-Party Collaboration

Geoff A. Cohen
Elysium Digital

INTRODUCTION

On its way from an attacker to a target, a cyberattack will pass through the control of many third parties. Packets will flow across wires owned by network operators; hacked machines owned by private individuals will use Internet service providers (ISPs) to transmit additional attacks; and Domain Name System (DNS) servers will resolve domain names by querying independent registrars for the current IP address. Packets will be processed by operating system and application software written by independent software vendors, configured by and installed by system integrators on hardware built by yet different companies.[1]

At the same time, it is difficult—perhaps prohibitively so—to reliably attribute attacks to a specific actor. Even when malicious actors are identified with some confidence, they are most often beyond the reach of law enforcement or other retaliation.[2]

These two factors lead to the conclusion that if we are to reduce the level of cyberattacks against private and government facilities, we need to focus efforts on changing the behavior of third parties, in particular making them less willing to cooperate with cyberattackers. These parties may be cooperating with the attacker, may be operating willfully blindly to the uses of their system, or may in fact be unaware that their network has been suborned (such as in the case of corporate networks infected by a botnet). Levying a sufficiently high price on these third parties for enabling attacks, for example by allowing mislabeled IP datagrams, hosting infected machines, or refusing to perform due diligence on clients of domain registration or hosting service, will dissuade these providers from cooperating. Such a price might include economic sanctions, law enforcement measures, and increased scrutiny from technical elements of the military or intelligence communities; they may also include more aggressive measures such as updating global routing tables to isolate the network.

These measures need not—probably can not—be unilateral on the part of the U.S. Government, but can be carried out in cooperation with network operators, the Internet community, foreign law enforcement, etc. However, it does not require complete global cooperation to achieve some measure of success.

[1]Goldsmith, Jack Landman and Tim Wu. *Who Controls the Internet? Illusions of a Borderless World* (Oxford University Press 2006).
[2]Menn, Joseph. *Fatal System Error* (Public Affairs, New York, 2010).

Focusing on third-party enablers has a number of advantages. First, over time it reduces the freedom with which hostile actors can operate. Second, it reduces the amount of damage done over time to private entities by infected networks. Third, it allows law enforcement or intelligence assets to focus attention on a smaller number of recalcitrant actors. Fourth, by limiting the number of routes through which attacks can be sent, it slows the speed at which attackers learn the precise vulnerabilities of U.S. assets and test various attacks.

COLLABORATION AND ITS DISCONTENTS

In English usage, the word "collaboration" has ambiguous connotations; it can be positive when referring, for example, to colleagues collaborating on a project.[3] In wartime, however, collaboration refers to working with the enemy, and is a crime.[4]

We mean to encompass that ambiguity by using the term "collaboration" here. We don't care whether the security failures are caused maliciously by bad actors, or only through inattention (e.g. by allowing malware to infect computers under their control). In legal terms, this distinction might go toward punitive damages rather than compensatory damages. But in either case, damage was caused in part through the actions or inactions of the particular party, and it should be held responsible.

The scope is meant to encompass any way in which a third party (intentionally or not) enables an attacker to launch an attack on a defender. This includes Tier 1 ISPs across whose wires (or even through the sky on a wireless signal) an attack might flow. It also includes more distant but still relevant factors such as software vendors that provided software containing the exploited vulnerability; domain name registrars that host a particular malware-containing site; or operators whose networks contain zombie client machines infected by malware. Other examples include software vendors that introduced vulnerabilities; ISPs that failed to perform adequate ingress filtering; hosting services that turned a blind eye to illicit activity or invalid addresses; sovereign nations that neglect to enforce cybercrime laws or refuse to offer cooperation with foreign investigative services, and so on.

Further, we consider the problem of such attacks as they happen in peacetime; we don't consider the appropriate reaction to attacks during declared armed conflict, or in the escalatory stage immediately before major conflict erupts. Nor do we consider the use of information warfare as an adjunct to traditional kinetic operations. The focus here is on the actions that the U.S. government should employ in the face of attacks or evidence of attack capability (such as malware infection) during "normal" peacetime.

A goal of this paper is to suggest a range of appropriate incentives and disincentives to make third parties less likely to collaborate with attacks.

THE URGENT CASE FOR ACTION

The distinction between "cyberwar" and "cybercrime" is a misleading one, leading one to believe that these are fundamentally different, even if related, activities. But they are far more similar: they are attacks on U.S. national interest, using the same techniques, launched by the same computers and using the same infrastructure to carry out the attack, planned and executed by the same people. The only differences are whether the target is U.S. military/government or private individuals or corporations, and whether the attack is happening in peacetime or during a wider conflict. Important distinctions, to be sure, but surely suggesting that an inability to deter cybercrime is a sign of a wider failure in our system of deterrence. And to the extent that we allow cybercrime to continue essentially unchecked, we are giving our adversaries rich opportunities to shape their attacks and render them increasingly

[3]cooperation. Thesaurus.com. *Roget's 21st Century Thesaurus,* Third Edition, Philip Lief Group, 2009, http://www.thesaurus.com/browse/cooperation.

[4]Compare with the usage of "cooperation" (always positive) and "collusion" (always negative). Ibid.

effective, sophisticated, and profitable; these improved attacks (again, most likely launched by the same personnel) are the ones we will face in a time of direct conflict. In the meantime, U.S. and allied interests are suffering significant economic losses from cybercrime, resources transferred to hostile parties and used to further develop attacks. We are, in a real sense, paying our adversaries to build weapons with which to attack us; a curious twist on the prediction on the death of capitalism often attributed to Lenin: "They will sell us the rope with which we will hang them."

Viewed more positively, cybercrime and minor cyberattacks are actually providing an important service to U.S. interests, by testing our defenses and revealing vulnerabilities and complicit third parties. If the United States takes advantage of these opportunities by improving defenses and reforming or shutting down third parties, then perhaps cyberattacks against U.S. interests in an actual high-stakes conflict might be considerably less destructive.

Most cyberattacks are minor, and happen hundreds of times a day across the globe. As such, we all live the day after these attacks, and so we worry more—or ought to, anyway—about how to clean up, whether to disclose the attack publicly, how to repair, and, most important for our purposes here, what we might do to prevent a future attack.

INTRODUCTION TO THIRD-PARTY INTERVENTION

The idea of going after third parties when it is difficult to identify or target actual malefactors is not a new one. Classic examples include "dram shop" laws that attempt to reduce drunk driving by threatening sanctions against proprietors of bars. More recent examples include the Unlawful Internet Gambling Enforcement Act (UIGEA), which intended to clamp down on Internet gambling in the United States by forbidding U.S. financial institutions from transferring assets on behalf of customers from U.S. accounts to overseas gambling institutions. Here, a more traditional target would be the casinos, but they were overseas and thus out of the reach of direct regulation. An alternative target would be the individual users who were gambling, but any attempt to monitor and regulate this behavior would be extremely difficult because of the relatively inobvious nature of it and the extreme degree to which the behavior is distributed. By targeting banks, the government identified a point in the business chain that had a fairly small number of actors (and ones quite accustomed, if not happily, to government regulation).

Such efforts have had limited success in the past. Even in the face of dram shop laws, there is still drunk driving (although some analysts have identified a small reduction in drunk driving in the presence of such laws).[5] Shortly after UIGEA went into effect, there is still online gambling, still accessed by U.S. citizens.[6]

A number of existing U.S. or international laws (or proposed laws) provide potential platforms for third-party intervention for cyberattacks. These include the Council of Europe's Convention on Cybercrime, the currently proposed U.S. Cybersecurity Act, in addition to common law or international norms, such as liability law. This section provides an introduction to some of these.

Communications Decency Act

The Communications Decency Act (CDA) was an attempt by Congress in 1996 to regulate pornography and "indecent material" on the Internet. Section 230 of the Act provides a limited immunity for "providers or users of interactive computer services" if they are not the originator of the allegedly defamatory, indecent, or other actionable material. In other words, a website operator is not liable for defamatory statements made by one of its users, even if that user made the statement on the website.

[5] *Drinkers, Drivers, and Bartenders: Balancing Private Rights and Public Accountability,* F. A. Sloan, ed., Chicago: University of Chicago Press, 2000.

[6] "On Poker: UIGEA law now official, but effects remain unclear," Chuck Blount, http://www.mysanantonio.com/sports/On_Poker_UIGEA_law_now_official_but_effects_remain_unclear_95472109.html, accessed July 9, 2010.

This immunity has explicit limits: it does not extend to intellectual property (which is dealt with through the DMCA) or criminal matters.

CALEA

The Communications Assistance for Law Enforcement Act (CALEA) of 1994 requires telecommunications carriers—which later court decisions and executive branch guidance make clear include broadband Internet service providers and Tier 1 carriers—to cooperate in enabling law enforcement access to their networks for monitoring communication, that is, "wiretapping."

This statute is only questionably wide enough to allow investigation into security breaches such as botnet activity or hacking, but provides an important precedent and template for future statutory language requiring networking infrastructure companies to cooperate with law enforcement.

The technological measures that complied with such language, however, might create just as much a problem as it was intended to solve. Government-approved backdoors into Tier 1 providers or hosting services would no doubt become attractive targets for hackers and would make the situation worse. However, if compliance required only that carriers and services allow investigative agencies efficient and standardized access to security audit and logging data, then perhaps tech-savvy investigators could more effectively identify hacked machines and trace the exploit back another level.

DMCA

The Digital Millennium Copyright Act (DMCA) contained a set of amendments to Title 17 (Copyrights) passed by Congress in 1996 and designed, in part, to adjust copyright law to respond to technological innovation such as digital computers and networking. Two sections provide potential analogs to mechanisms to combat malicious activity on the Internet.

First, DMCA provides a coupled takedown notice/safe harbor mechanism. Rights-holders can send a "Section 500" takedown notice to an Internet content provider or ISP that copyrighted material is being made available in an unauthorized manner. If the recipient promptly removes the material, then the recipient is in the "safe harbor;" it cannot then be held liable for infringement.

The recent case of Viacom v. YouTube provides some guidance to interpret the scope of this safe harbor. In the case, Viacom, working with many other rights-holders, sued YouTube, claiming that copyright infringement was common on YouTube. YouTube responded that they had complied promptly with all (proper) takedown requests; they won on summary judgment.

A possible parallel with cybercrime would create a similar dynamic. Third parties that had evidence that they were harmed by cybercrime activity, such as a DDoS attack from a botnet hosted in an ISPs network, could send a takedown request to that ISP. If the ISP promptly responded in some way that remediated the threat, then they would have safe harbor protection from further civil or criminal suits under CFAA or similar laws. The cybercrime safe harbor should possibly be narrower than in DMCA: general knowledge of botnet activity and negligence in performing industry-standard countermeasures should certainly block a party from receiving safe-harbor protection, even if it does comply promptly with cybercrime takedown requests.

The second relevant mechanism created by DMCA is the anti-circumvention restrictions. Under DMCA, it is illegal to create or distribute instruments capable of circumventing technological controls intended to limit access to copyrighted material. With appropriate carve-outs for academic and security research (carve-outs unfortunately missing from DMCA), similar language could apply to the construction and distribution of software intended for illegal use in hacking, even if this code is written overseas (as it very often is); see DeCSS case—websites that even link to the offending code were found liable.

In general, it's worth noting that our society and legal code have tried much, much harder to control and regulate culture through copyright law than it has to control the threat of computer crime or computer war. There are multiple legal recourses for private entities to bring down offending material that

violates copyright, but very few options for entities under attack that could actually destroy a business. It is perhaps time to rethink our priorities.

CFAA

The Computer Fraud and Abuse Act (CFAA), strictly speaking 18 USC. Sec. 1030, sets out a number of scenarios of malicious and unauthorized access to the computers that trigger criminal sanctions.

Section 1030 was created by the Comprehensive Crime Control Act of 1984, which criminalized for the first time unauthorized access of computers (as a misdemeanor). It was amended in 1986 by the Computer Fraud and Abuse Act, and amended again eight further times, including by the USA PATRIOT Act in 2001 and by the Identity Theft Enforcement and Restitution Act in 2008. It criminalizes a number of actions covering unauthorized access to certain protected computers, defrauding, causing damage to computer systems (or extortion involving threats to perform damage), trafficking in passwords, etc.

A particular crime is to "knowingly cause the transmission of a program, information, code, or command, and as a result of such conduct, intentionally [cause] damage without authorization, to a protected computer." This would seem to cover many forms of attacks, including denial-of-service attacks, transmission of viruses/worms/etc.

Normally, crimes under CFAA are investigated by the U.S. Secret Service (except for espionage-related investigations, which are handled by the FBI) Under certain circumstances, private parties can bring civil suits under CFAA.[7]

Liability

The current state of U.S. law shields software vendors from product liability claims, as software is usually sold as granting a license to use a product, rather than as the sale of a product. This legal regime could be changed by legislation, creating an enormous economic pressure on software vendors to improve the security of their software. The argument for exposing vendors is that the software marketplace has failed to provide adequate security, due to many factors including the illiquidity of the software market (especially the market for operating systems) and the fact that security failures are an externality the cost of which is borne by neither the vendor nor, in many cases, the purchaser. Even under the current legal framework, software vendors are potentially vulnerable to negligence suits.[8]

National Security Measures

There are a host of U.S. Government powers flowing from national security requirements, including the use of National Security Letters,[9] the war powers ability of the President to shut down wired and wireless networks through Section 606 of the Telecommunications Act, and potentially additional powers granted by the Cybersecurity Act of 2010 or other future legislation.

These mechanisms are undoubtedly powerful tools, and may be appropriate in responding to extremely large and time-sensitive attacks. However, they don't scale well; if the President or Federal judges needed to intercede to combat every computer attack, they would quickly have no time to do anything else. Also, these measures raise concerns about civil liberties, the constitutionality of executive

[7] *Prosecuting Computer Crimes*, Chapter 1, U.S. Department of Justice, Computer Crime and Intellectual Property Section, Criminal Division, Scott Eltringham, Editor in Chief, February 2007, downloaded from http://www.justice.gov/criminal/cybercrime/ccmanual/.

[8] Chandler, Jennifer A., Improving Software Security: A Discussion of Liability for Unreasonably Insecure Software. *Securing Privacy in the Internet Age*, Stanford University Press, 2006. Available at http://ssrn.com/abstract=610041.

[9] Doyle, Charles. "National Security Letters in Foreign Intelligence Investigations: A Glimpse of the Legal Background and Recent Amendments," Congressional Research Service, September 2009, downloaded from http://www.fas.org/sgp/crs/intel/RS22406.pdf.

branch interference in protected speech, and the possibility of using such mechanisms to attack political enemies rather than actual adversaries.[10]

LEGAL PRINCIPLES

Government Role in Private Security

In general, the private sector is best equipped and structured to respond to an evolving cyber threat.
—National Strategy to Secure Cyberspace, 2003.

Security is a nearly textbook example of a market failure. The cost of security failure is an enormous externality: it is most often not borne by the organization that failed (for example, if the organization is penetrated and its computers turned into a botnet to send out spam, or if the organization makes a great deal of profit selling insecure software). The knowledge on how to best respond to foreign attacks is rare and dispersed; private corporations are notoriously bad at spending money to combat long-term threats that may not materialize.

Even when solutions are known, coordination costs often prevent deployment of the solution for many years. For example, the use of the Domain Name System as an attack vector to hack machines has been known since the early 1990s,[11] and yet the proposed solution, DNSSEC, was still in the process of being deployed in 2010.[12]

Government can efficiently counter coordination cost (by, for example, setting mandatory standards and timetables for deployment), is a center for expertise, and is the agent of the collective public for acknowledging externalities by levying taxes or fines against the initiator and compensating the damaged. As such, it can serve an important, if not critical role in addressing security failures in the private sector. A more extreme version of this claim is a legal principle that the government has an affirmative responsibility for assuring the security of private infrastructure.

Shared Affirmative Responsibility for Security

A broader principle is that all actors in the software marketplace, including vendors, network infrastructure operators, resellers, and corporate and individual users, have an affirmative responsibility to take reasonable and timely measures to assure the secure and legal operation of their systems and services. This includes the duties to monitor their systems for security, detect vulnerabilities or actual exploits, respond to notice from third parties that a security problem exists; inform relevant users, operators, and authorized clearinghouses; and resolve such vulnerabilities and if relevant distributing such patches.

This legal principle, if adopted, would run directly against the legal framework that affords common carrier rights, and would have to be carefully crafted to not create undue burdens on entities such as Tier 1 ISPs that could not possibly be aware of all activity on their network. See below.

Too often cybercrime investigations have been met with studied airs of ignorance; hosting providers are often shocked to hear that they may have been hosting malware.

[10]See, for example, Fellow, Avery, "Experts Caution Against Federal Web Regulations," Courthouse News Service, http://www.courthousenews.com/2010/08/26/29911.htm.

[11]See, for example, RFC 3833, message by Jim Galvin posted to dns-security@tis.com on November 19, 1993, and Bellovin, S., "Using the Domain Name System for System Break-Ins," Proceedings of the Fifth Usenix Unix Security Symposium, June 1995.

[12]See, for example, "Status Update, 2010-07-16," http://www.root-dnssec.org/2010/07/16/status-update-2010-07-16/.

Presumption of Harm

Another proposed legal principle is that malware, exploits, vulnerabilities, etc. create a presumption of harm and are presumed to be counter to the public interest. The legal principle proposed here, somewhat parallel to the Supreme Court's standards for considering injunctions,[13] considers the public interest, the possibility of irreparable harm, and the balance of hardships of discovering the security violation, addressing it after the fact, and the (potential) damage inflicted. In situations where security falls behind exploit technology (as it seems we are in today), the balance of hardships surely indicates that actors are more responsible for cleaning up their own house. In situations where security vulnerabilities are rarer and more isolated, perhaps the balance of hardships would give some more lenience to operators.

The Balance of Incentives

To sum up, the central idea in this approach is that all operators in the network system have a responsibility for securing their system, which includes active monitoring for vulnerabilities and exploitations. This responsibility carries with it the potential for serious legal consequences if actors fail to meet their expected standards of behavior.

Third parties can gain immunity from such consequences—enter a safe harbor, in legal terms—through a variety of ways, which include either remediating the problem or providing proper notice to some other actor up or downstream of the problem. Once notified, those other actors are now on the hook to remediate or otherwise respond.

The idea is to balance these incentives appropriately: the threat of punishment, and positive incentives to fix the problem or shift the burden elsewhere.

A scale of degree of punishment would start with merely failing to meet minimal standards of behavior, progress to "knew or should have known" of vulnerabilities, continue to willful negligence, and end with actual collaboration with malicious actors. At each stage, well-defined (by government or neutral parties such as IETF) standards of behavior define exactly the duties and responsibilities of the actor, as well as the mechanisms to inform others of their own responsibilities.

INSTRUMENTS OF POLICY

Here, we discuss potential actions against third parties—the "retaliation" for allowing cyberattacks to take place. Of course, the presumption is that these entities are not actually malicious actors, or they would be facing criminal investigation instead. As such, the penalties are not designed to be overly punitive. However, at the very least they ought to compensate society for the actual, and some fraction of the potential, damage incurred by their failure to secure their system.

Further, such penalties can't merely be part of the cost of business, as long-haul truckers understand that invariably they will get speeding tickets on the highway. The penalties must be significant enough to drive behavioral changes.

One question is whether actual exploitation must happen in order to trigger the penalty, or is merely having an open vulnerability enough? The first case is easier to sell politically and is linked more closely to existing legal and policy concepts of negligence, contributory infringement, and standards of behavior and responsibility. It would also be easier to specify damages, since in fact some damage must have occurred.

However, this narrower range of response would leave out the much larger range of vulnerabilities that have yet to be exploited. Targeting these actors, while much more politically challenging, and more difficult to fairly assess the potential damage, would have a much broader and more significant policy

[13] *eBay Inc. and Half.com, v. MercExchange, L.L.C.,* 126 S. Ct. 1837.

impact. First, it would at least in some cases avoid the actual damage from occurring, since vulnerabilities would be patched and addressed before a malicious actor could take advantage of it. Second, it would create incentives for more widespread changes in behavior, since many more people and organizations would be swept up in the regulation, and the chances of getting hit with a fine would be much larger if you can be fined for simply failing to secure a network.

This approach has more in common with public health campaigns against epidemics than traditional security research. Such an approach emphasizes protection of as-yet unattacked populations through isolation, inoculation of key at-risk groups, and only secondarily the treatment of infected groups. And although the applicability of this approach is most obvious for the biological metaphor attacks of viruses and worms, it can apply more broadly to other malware infection or even more general vulnerabilities.

Market Intervention

As computer security represents a market failure, one solution is for the government to try to address the problem by intervening in the market to achieve a better societal outcome. This could include mandating computer security insurance, with higher rates applied to users of known-insecure systems.

In contrast, systems or services that met higher standards of behavior and have good or excellent track records could be certified by the government as being especially secure, a sort of "Energy Star" program for security. Government agencies could be directed to purchase only those products and services that met this certification; certainly, much of the private sector would do so as well. Targeted tax incentives could also accomplish a similar goal of putting a thumb on the balance in the market.

Collective Action

Many counter-malware activities on the Internet require information to be gathered at many points throughout the network, and information to be distributed. By its very nature, some detection and analysis on the Internet requires the examination of thousands, or millions, of sites. Responding appropriately may also require a distributed response. For example, one technique to fight spam is for many users to pool together the knowledge of originators of spam, and then agree to disallow mail from any of those sources (so-called "blacklisting").

The Internet itself enables this kind of loosely coordinated collective action of interested parties (sometimes also called "crowdsourcing"). Prominent examples include spam-tracking services such as SpamHaus, a non-profit which accepts information from any user regarding spam and then makes it available to system operators.

Another example is the Knujon project to examine the complete DNS records. Previous projects to try to inspect the Domain Name System ran against the sheer scope of the task: nearly 120 million registered domain names in the top five generic top-level domains. By distributing the task to volunteers, Knujon was able to identify hundreds of registrars that were violating their registry accreditation agreement with ICANN.[14]

Of course, information must be the base of action to be useful. In SpamHaus' case, the list of suspected spam-sites is made available for other network operators to (voluntarily) use to mark mail as spam and potentially discard it.

It's clear that collective action will be a key technique in fighting malware and cybercrime and cyberwarfare. Especially as other aspects of this report are adopted, there will be an increasing amount of information collected on the relative degree of proper or rogue behavior on the part of networks, hosts, registrars, etc. This information should be widely disseminated and used with discretion in determining

[14] *Knujon Internet Security Report: Audit of the gTLD Internet Structure, Evaluation of Contractual Compliance, and Review of Illicit Activity by Registrar -2010.* http://www.knujon.com/knujon_audit0610.pdf.

whether other entities should trust them. For example, BGP route updates from a rogue ISP that had issued faulty or malicious updates in the past should not be automatically accepted, but require system administrator action. Similarly, SSL certificates signed by a certificate authority in an untrusted nation or from an untrusted ISP should be flagged to users.

Collective action is not without risk. Distributed monitoring and blacklisting could metastasize into mob rule or vigilante justice. Smoothly working procedures could be subverted by bad actors. Principles that could guard against such failures include making the operations and decision making process clearly transparent; providing well-documented appeals and grievances procedures; and some notion of outside accountability or authority, even if tenuous. Non-governmental bodies such as the IETF may be able to serve this role, for example by supporting the drafting and publishing of Best Current Practices documents that describe varieties of collective action.

A second variety of problem with collective action is illustrated by the litigation between e360Insight (a bulk emailer) and SpamHaus. Lawsuits or even the potential threat may be enough to shut down many collective efforts. What's needed is clear statutory guidance that establishes a legal presumption of validity for such activity in the face of negligence or other lawsuits. Such a presumption should only be overcome only by evidence of willfulness or malicious intent.

RESPONSIBILITIES OF THIRD PARTIES

We consider any actor in the network to be a third party that may be a target of this policy. This can include elements of the Domain Name System, registrars, ISPs (Tier 1, smaller, and corporate networks), hosting services, hardware and software vendors, and even sovereign nations. Of course, each of these has different vulnerabilities, presents a different threat model to outsiders, and is responsive to different levers of power. Registrars must comply with rules set by ICANN, while smaller ISPs are dependent on Tier 1 ISPs for connectivity. Sovereign nations may be the hardest case, but even here, diplomatic and economic pressure can be brought to change behavior. For example, Romania was traditionally a haven for cybercrime, but under pressure from the EU and threatened with the withdrawal of technological assistance from richer countries, it became a signatory to the Convention on Cybercrime and is reportedly cooperating more with foreign investigations.[15]

Network Operators

Network operators play a special and challenging role, as they are both in the best position to monitor for security failures and to intervene rapidly and effectively when they occur, but at the same time do not want to be in the business of individually inspecting and approving customers or traffic. To require network operators to monitor the content of traffic would be a significant burden, although it may be that they are doing much of this already for law enforcement purposes.

An important goal of this project is to encourage entities to detect malicious activity and pass the information closer to a location that can take action. For example, local ISPs are well-positioned to identify particular IP addresses (or MAC addresses) as the source of port-scans or spam indicating compromise by malware. They are also best positioned to determine the actual subscriber that was using that particular address at that particular time. The next logical step is for the ISP to communicate its suspicion to the user (presumably with offers of assistance in cleaning up the problem). By providing safe harbor to ISPs that communicate and offer to assist, this policy would encourage ISPs to do that.

Corporate ISPs—that is, companies that provide Internet access to their employees at work, but own and operate their computers—should receive no such safe harbor. They're responsible for their own messes, and any statutory changes in liability should reflect the difference in ISPs.

[15]Menn, Joseph, *Fatal System Error*, Public Affairs, New York, 2010.

A basic step is to perform ingress and egress filtering, that is, to check at network boundaries for incoming or outgoing packets that are marked as having impossible origin addresses. IP address spoofing is a key ingredient of the difficulty in attribution: if attacking packets (such as SYN requests used in a DDoS attack) are marked as having come from a false originating machine, then obviously it will be extremely difficult to identify the actual source. Network operators can make such spoofing much more difficult by discarding incoming packets that are marked as having originated at a machine that could not have actually been routed through the peered network. Similarly, outgoing packets should be dropped if they are not marked as having originated within the network or a network with a valid route. A further useful step would be to log and investigate such events. This is not a panacea, of course; it merely restricts the freedom of malevolent actors to spoof. However, it will at least provide some degree of clues to investigators trying to determine the source of attack.

Such filtering is not a new technique—RFC 2827 described it in May 2000.[16] It is considered a "Best Current Practice," but is not required. If a network attack occurs which uses spoofed IP addresses, then any network that transmitted such networks and failed to apply ingress or egress filtering, then it is potentially a target of a retaliatory action by the government.

Such security requirements are, for example, required by the Payment Card Industry Data Security Standard.[17] (More broadly, this is a useful example/model for how private industry could develop stronger security practice requirements, and create a supporting audit/certification regime). Inter-ISP terms of service ought to make sending on unfiltered packets a violation that allows termination.

A similar attack is to publish invalid routes. For example, in March and April of 2010, a small Chinese ISP published invalid routes to more than 30,000 networks using the Border Gateway Protocol (BGP), the Internet's protocol for communicating routes between networks. Many ISPs accepted these routes, thus in effect isolating those networks.[18] Whether a result of a sloppy system administration or an intentional experiment, this clearly demonstrates a critical, current vulnerability in the architecture of the Internet; it's one that can be addressed best by ISPs.

Team Cymru has identified a number of measures by which ISPs can filter BGP messages, including throwing out routing information apparently coming from invalid or valid but unallocated IP blocks.[19]

To reiterate, these technological and administrative defenses are not new; the question is whether the U.S. government can require organizations to follow them, or conversely if a failure to follow them creates enough liability that the organizations could face punitive measures if they fail to follow them and as a result an attack occurs.

Registrars and Hosting Providers

Currently, much of the domain name registration information (that is, the names and contact information associated with a particular registration) is invalid. And yet registrars face no repercussions for allowing false information, nor do they perform any checking. Often, payment is accepted using stolen credit cards. The situation is little different with hosting providers. When investigators tracing an attack come to a hosting provider or seek to determine who registered a domain name, they often hit a dead end.

While it would be overly burdensome and intrusive to require these sorts of companies to police the content being hosted, they certainly should be able to verify the identity—or existence—of the users submitting the form. Those who negligently allow invalid data should face consequences when such a

[16] http://tools.ietf.org/html/rfc2827. It was updated in May 2004 by RFC 3704 (see http://tools.ietf.org/html/rfc3704).
[17] https://www.pcisecuritystandards.org/security_standards/pci_dss.shtml.
[18] http://www.computerworld.com/s/article/9175081/A_Chinese_ISP_momentarily_hijacks_the_Internet_again_.
[19] http://www.team-cymru.org/Services/Bogons/.

site is used in an attack. This approach is analogous to the "Know Your Customer" regulations applied to financial institutions after 9/11 by the Bank Secrecy Act and the USA PATRIOT Act.[20]

Individuals

Ultimately, the vast majority of the unsecured aspect of computing is the individuals who operate or own personal computers. Any other efforts by security vendors, operating system authors, ISPs, or even police, intelligence agencies, or the military to secure the computing infrastructure—or even merely erect minimal defenses against attack—will be an uphill struggle if the endpoints remain compromised. For all the malevolence of malware authors or criminal website operators, in most situations there's no damage until someone's machine becomes infected. (Even in snooping or man-in-the-middle attack, it's often because the intermediate node has itself become infected first.)

Without yet raising the question of legal responsibility, it's straight-forward and well-documented that the security burden on individuals comprises ensuring that they're running the most recent patched version of system software and applications; running an updated virus scanner and security suite; and avoiding risky behavior such as installing software from untrusted sources or visiting shady websites.

This is all too well known and has been for years; clearly it is insufficient to merely announce these burdens. It is no more effective than fighting the obesity epidemic by publishing another diet book or exhorting people to eat less and exercise more. Clearly, if endpoints are to be secured, then some more proactive and regulatory action on behalf of the government is necessary.

Can individuals be punished if their infected computers are found to have been used in an attack? Even determining such data could raise significant privacy and civil liberties concerns. Instituting mechanisms to detect and legal infrastructure to, for example, punish citizens for visiting certain cites is, even if well-intentioned, a mechanism for oppression. Counterintuitively, a broader and stronger mandate that applied to all users would be a more effective while encroaching less on liberty.

No doubt an individual mandate requiring individuals who used networked computers to buy, install, and maintain security software would be as controversial as the mandate to purchase health insurance—indeed, the issues are quite similar both in legal analysis and in the way an individual's choices affect global outcomes. Unflattering parallels will undoubtedly be drawn with China's mandate for computer manufacturers to install Green Dam, its Internet filter that would prevent access to sites containing content on, for example, pornography and democracy. (The parallels may be deeper: some security researchers hypothesized that due to security flaws in Green Dam, it would become compromised by malicious actors, and the population of all connected PCs in China could become an enormous botnet.)

And yet this is no different than requiring drivers to have bought insurance from the market. A requirement for security software is even less burdensome than one for auto insurance, since effective free options are available.

It's certainly true that the aggregate cost of maximally securing every end-point is enormous and probably prohibitive; and yet it is not necessary to reach that goal to be effective. As long as the number of vulnerable nodes decreases to such an extent that certain types of hacking become unprofitable (relative to other activities), then this approach will have some value. This is not to mention the reduced pain and damage accruing to those individuals that will not be harmed given this higher—even if not maximal—level of security.

WILL THIS STRATEGY BE EFFECTIVE?

As described above, efforts to prevent bad behavior by controlling third parties has had mixed results, from gambling and drunk driving to counterfeit goods. More immediately, successful efforts

[20]See http://en.wikipedia.org/wiki/Know_your_customer.

to shut down known malicious ISPs have resulted in dramatic drops in spam rates in the short run (estimated up to 70% reductions in world-wide spam[21]); however, spammers quickly develop new techniques and find (or establish) new ISPs willing to convey spam.[22]

For these reasons, it is perhaps overly optimistic to consider that such a technique could work. However, much of the problem in computer security is that there is little learning; each round is played as if it is the entire game. Repeated instances of the same security hole don't seem to result in changed behavior by other actors. For example, consider the Chinese ISP that was responsible twice in two months for hijacking global routing tables. As the saying goes, fool me once, shame on you; fool me twice, shame on me. And yet knowledge of bad actors isn't encoded into the system in any way; there is, for example, no easy way for end-users to know that they have visited a site which is known to host malware, that is registered through a registrar that is known to work with spammers, that is using a certificate generated by a certifying authority in a country known for protecting hackers, and so on.

A legal/technical framework that consistently applied standards of liability for negligence and safe harbors for responsible actions could conceivably raise the level of secure behavior across the system, rather than simply playing whack-a-mole against particularly extreme bad actors such as McColo or Waledac.

Future analysis might be able to more precisely predict effects of improved security measures. For example, game theoretic or complex adaptive system simulation models of threats and responses, tied with white-hat/black-hat exercises, could suggest ranges of effectiveness, or more narrowly describe the regimes in which particular approaches might be effective. Such models might also suggest the adversary's countermoves.

Distant Third Parties

It is easy enough to imagine a broad range of regulatory regimes to apply to companies acting within the U.S., and most likely such an approach could also be introduced into close allies in the industrialized West. But how to address software vendors, hosting providers, or other services in other nations? This is the problem of the Russian Business Network (RBN), a so-called "bulletproof" network of services that has been implicated in major hacking schemes and cyberattacks. And yet due to its complicated, interlocking structure with shifting identifies (mostly forged), it is extremely difficult for agents in the West to identify a particular malefactor with any confidence. And given the lack of cooperation provided by Russian law enforcement, the RBN operates with impunity.[23] Can the overflight approach advocated here help with this situation?

It can in three ways, although of course no domestic peace-time regulatory regime will be able to fully shut down offensive capability or even action in an uncooperative nation. First, it reduces the ability of a malicious actor's ability to launch attacks from within the defensive perimeter of the United States. Most recent offensive information operations, such as the Korean attack, involved U.S.-located assets in at least some stage, whether zombie-compromised PCs on U.S. ISPs, U.S. registrars or hosting providers, botnet control communication servers (e.g. IRC servers), or even compromised machines used as botnet controllers. Increased regulatory retaliation for security failures against U.S. agencies will drive these activities overseas, making attacks easier to defend against because they will come through a much smaller number of nodes.

Second, from the other point of view, we would be denying the advantages of the use of U.S. infrastructure to malevolent actors. There's a reason that world businesses—legitimate and otherwise—want to use hosting services in the U.S.: high-quality, comparatively cheap, universally available, 24/7 electric

[21] "Web Provider Busted, Spam Drops," Stefanie Hoffman, CRN, Nov. 13, 2008, http://www.crn.com/security/212002482.

[22] "Analyzing the Aftermath of the McColo Shutdown," Steve DiBenedetto, Dan Massey, Christos Papadopoulos, Patrick Walsh, Workshop on Trust and Security in the Future Internet, (FIST'09), Seattle, WA, July 2009, downloaded from http://www.cs.colostate.edu/~christos/papers/DiBenedetto09a.pdf.

[23] *Inside Cyber Warfare: Mapping the Cyber Underworld*, Jeffrey Carr, O'Reilly Media, December 2009.

power; convenient access to electronic replacement parts (including highly available parcel delivery); and in general an uncorrupted polity in which business expenses are reasonable and predictable, and relatively safe from gang warfare. In general, the U.S. is a friendly location for Internet-based business. Denying these advantages to attackers forces them to rely on inferior infrastructure in non-cooperating nations. This is no mere coincidence: nearly by definition, those nations that refuse to prosecute cybercrimes (because in fact cybercrime gangs are an extension of state power) are the ones where criminal elements define the rules, an inherently dangerous place for business.

Improved security will also have spill-over benefits to other nations. If large operating system vendors, for example, improve the security of their product for U.S. consumption, then the security of endpoints around the world will improve. If Tier 1 ISPs accelerate adoption of improved protocols, such as IPv6, then that will further develop the international marketplace in products and software that support those protocols, as well as validating standards of behavior that will encourage worldwide adoption. Individually these measures may not result in large improvements in security. IPv6, in particular, is no guarantee of security. However, increasing U.S. spending (public and private) and activity in the realm of securing computing systems should have an aggregate benefit in reducing the cost and increasing the availability of secure software, as well as increasing the number of skilled professionals world-wide. Perhaps most importantly, it will help establish and strengthen norms of behavior among software and networking professionals and researchers that securing their products and services is a necessary element to allowing the Internet to continue to grow and thrive.