Editors: Richard Ford, rford@se.fit.edu
Deborah Frincke, deborah.frincke@pnl.gov

Ethics in Security Vulnerability Research

Andrea M. Matwyshyn
University of Pennsylvania

Ang Cui, Angelos D. Keromytis, and Salvatore J. Stolfo
Columbia University

Debate has arisen in the scholarly community, as well as among policymakers and business entities, regarding the role of vulnerability researchers and security practitioners as sentinels of information security adequacy. (The exact definition of vulnerability research and who counts as a "vulnerability researcher" is a subject of debate in the academic and business communities. For purposes of this article, we presume that vulnerability researchers are driven by a desire to prevent information security harms and engage in responsible disclosure upon discovery of a security vulnerability.) Yet provided that these researchers and practitioners do not themselves engage in conduct that causes harm, their conduct doesn't necessarily run afoul of ethical and legal considerations. We advocate crafting a code of conduct for vulnerability researchers and practitioners, including the implementation of procedural safeguards to ensure minimization of harm.

Why Vulnerability Research Matters

The computer and network technologies that we've come to depend on in every part of our life are imperfect. During the past decade, the practice of finding and exploiting such imperfections has matured into a highly lucrative industry. To combat this escalating threat, security researchers find themselves in a perpetual race to identify and eliminate vulnerabilities before attackers can exploit them and to educate and train practitioners to test for known vulnerabilities in deployed systems. While nefarious parties operate in secrecy without fear of law, ethics, or public scrutiny, legitimate security researchers operate in the open and are subject to these constraints. Hence, researchers are sometimes hesitant to explore important information security issues owing to concern about their ethical and legal implications.

Provided that vulnerability research is done ethically, researchers perform an important social function: they provide information that closes the information gap between the creators, operators, or exploiters of vulnerable products and the third parties who will likely be harmed because of them. A culture war is currently under way in the cybersecurity industry and research communities regarding the value of investment in vulnerability analysis of products and operations. On one hand, many data security champions argue that maintaining best practices in information security, which includes diligent analysis of products for vulnerabilities and flaws, is "the right thing to do" both for the system operator and society as a whole. Yet skeptics (including some security professionals) argue that short-term expenditures on such "nonessential" items as analysis should be curtailed, and that the results of any analyses should be kept secret. The return on investment in security isn't visible in the short term, and, therefore, detractors feel empowered to ignore the well-known long-term costs of vulnerability, which include negative effects on the value of intangible assets and goodwill. They argue that investment in security is squandering corporate assets that could be better utilized to generate strong short-run returns for shareholders.

Unfortunately, corporate information security skeptics currently have a firm hold inside many enterprises. In particular, empirical data indicates that companies aren't successfully anticipating and managing information risk. For example, in the 2008 PricewaterhouseCoopers annual information security survey of more than 7,000 respondents—comprising CEOs, CFOs, CIOs, CSOs, vice presidents, and directors of IT and information security from 119 countries—at least three of 10 respondents couldn't answer basic questions about their organizations' information security practices. Thirty-five percent didn't know how many security incidents occurred in the past year; 44 percent didn't know what types of security incidents presented the greatest threats to the company; 42 percent couldn't identify the source of security incidents; and, finally, 67 percent said their

organization didn't audit or monitor compliance with the corporate information security policy—whether the attack was most likely to have originated from employees (either current or former), customers, partners or suppliers, hackers, or others. According to this annual longitudinal research, many company leaders lack a well-rounded view of their information security compliance activities: "business and IT executives may not have a full picture of compliance lapses ... Fewer than half of all respondents say their organization audits and monitors user compliance with security policies (43 percent)" and "only 44 percent conduct compliance testing" (www.pwc.com/extweb/insights.nsf/docid/0E50FD887E3DC70F852574DB005DE509/$File/Safeguarding_the_new_currency.pdf). Rampant data breaches of millions of records in 2009 further speak for themselves, demonstrating widespread inadequacies in corporate information handling (www.privacyrights.org/ar/ChronDataBreaches.htm). Meanwhile, each of those breached records is attached to a company or a consumer potentially harmed by the disclosure.

It's undisputable that lax information security and vulnerable products erode commercial trust and impose costs on third parties—business partners, shareholders, consumers, and the economic system as a whole. The reason for this arises from the nature of information risk: its impact is inherently transitive. This means that if a company fails to secure another company's information, the negative effects to the shared data are similar to those that would have occurred if the original company had been breached itself (for example, banks affected by data breaches have argued that they can't continue to absorb the downstream costs of other companies' information security mistakes[1]). In practice, this means that negative financial externalities are imposed on individuals and companies not responsible for the data loss. Furthermore, information stolen about individual consumers is sometimes used for identity theft. Harms to social institutions also occur. The social security system, for example, has been threatened in part due to rampant social security number vulnerability.[2] Similarly, the integrity of social structures, such as law enforcement and the criminal justice system, is negatively affected by information crime. For instance, identity thieves sometimes identify themselves using a victim's personal information when charged with a crime.

Increasingly, ethics scholars are recognizing the existence of core ethics standards that apply to all commercial activities.

The proper calculus with respect to information security adequacy should turn on the simple ethical question: "Have we verified that our products and operations don't cause avoidable harm to others?" This "duty not to harm" can be operationalized in information security practices in at least two ways. First, it involves timely, fair, and accurate disclosure of the existence of security vulnerabilities that put consumers, business partners, and the social system at risk, thereby enabling these affected parties to mitigate their exposure to information risk. Second, it involves due care in research and development, as well as auditing and updating information security practices to stay in step with the state of the art. To date, neither of these practices are a universal norm of corporate conduct. Further current legal regimes aren't robust; the law is currently inadequate to enforce this duty not to harm.[3] An impactful information gap exists, which vulnerability researchers help to close. Without this intermediation, it's unlikely that meaningful improvements in information security will occur in a timely manner. Meanwhile, the consequences of widespread vulnerability carry heavy social costs.

Vulnerability Research: Neither Unethical nor Illegal

Increasingly, ethics scholars are recognizing the existence of core ethics standards that apply to all commercial activities. They point to factors such as acting honestly and in good faith, warn against conflicts of interest, require the exercise of due care, and emphasize fairness and just results. (In "Confronting Morality in Markets," Thomas Dunfee and N.E. Bowie argue that morality is expressed within markets and could result in pressures on organizations to respond.[4]) Perhaps the most basic of benign moral concerns in ethics is the duty to avoid knowingly or recklessly harming others—that is, the "duty not to harm."

Some critics of vulnerability research assert that it's inherently unethical, presumably because it involves testing systems and analyzing products created and maintained by someone other than the researcher (http://searchsecurity.techtarget.com/magazineFeature/0,296894,sid14_gci1313268,00.html). If we apply the ethics principle of the duty not to harm, however, a strong argument exists that at least a portion of vulnerability research is ethical and, in fact, ethically desirable. Provided that vulnerability research is technologically nondisruptive, doesn't damage

the functionality of the products and systems it tests or otherwise harm third parties, the ethical duty not to harm appears to be met. Additionally, with some vulnerability research, its goal is explicitly to prevent or mitigate harm occurring to third parties because of vulnerable products and operations whose creator has failed to disclose the danger. As such, we can argue that the ethical duty not to harm might even mandate vulnerability research in some cases: the community of researchers possessing special skills to protect society from vulnerable products could have a moral obligation to use these skills not only for personal gain but also for the benefit of society as a whole. (For some ethicists, corporations might have ethical obligations to exercise unique competencies for societal good.[5])

Perhaps the most superficially potent objections vulnerability research skeptics have raised involve the law. First, critics assert that vulnerability research is unnecessary and, second, that all such research is, by definition, "illegal" because it violates state and federal computer intrusion statutes or intellectual property rights. On the point of vulnerability research being superfluous, critics state that acting responsibly for a business entity in the area of information security simply means complying with the law and that the law defines what constitutes good business practices. This objection fundamentally misunderstands the relationship between responsible corporate conduct and legal regulation. Law is merely a floor of conduct, not a marker of best practices or ethical conduct. Leading ethicists have explicitly rejected the idea that law and business ethics necessarily converge.[6] Furthermore, although both US and international regulators are beginning to take action in the realm of information security regulation, legally speaking, the field is still in its infancy. To date, the information security legal regime adopted in the US to address issues of vulnerability is an imperfect patchwork of state and federal laws, widely critiqued in legal scholarship;[7] it's also barely a decade old, doctrinally inconsistent, and in a state of flux.[3] A need for timely, fair, and accurate disclosure of the existence of information security problems arises from the ethical duty not to harm, regardless of the state of the law. By the time disclosure is legally mandated, irreparable harm has usually occurred. In fact, we can view the law as creating negative incentives for correcting security vulnerabilities: because contract law has allowed technology producers to disclaim essentially all liability associated with their products, there are limited financial incentives for these producers to disclose the existence of vulnerabilities and fix products promptly so as to avoid lawsuits. Vulnerability research fills an information void the law doesn't adequately address.

Although it's likely that a court would construe some forms of vulnerability research to be in violation of state or federal computer intrusion statutes, it's equally likely that some forms of this research would be deemed legally permissible. Even intellectual property rights have recognized limits at which concerns of consumer harm exist. In fact, Congress has encouraged vulnerability research in certain instances—for example, in the Digital Millennium Copyright Act, Congress explicitly protects research designed to test the privacy-invading potential and security implications of particular digital rights management technology.[8] Furthermore, the exact construction of the Computer Fraud and Abuse Act, the leading federal computer intrusion statute, is a subject of much debate and dissention, even among federal appellate courts. (For example, critics have analyzed the 7th Circuit[9] and the 9th Circuit[10] to stand in direct contradiction of each other with regard to whether an employee who accesses employer files and uses that information for his own purposes has committed a violation of the Computer Fraud and Abuse Act.) Its interpretation and meaning are far from clear. It's not obvious, for example, which forms of vulnerability research are prohibited by law. In the absence of clear legal guidance, however, it's essential that the community of vulnerability researchers commence a dialogue on self-regulation, best practices, and the boundaries of ethical conduct.

Crafting Norms of Vulnerability Research

Using a case study from research conducted at Columbia University, we propose several possible "best practices" in vulnerability research that we believe should be incorporated into a vulnerability researchers' code of conduct. Research demonstrates that the existence of corporate codes of conduct on ethical behavior are significantly related to such behavior or to whether employees behave ethically.[11] In particular, codes that clearly stipulate standards for information security conduct and sanctions for data mishandling are likely to generate more ethical conduct.[12] Because inappropriately done vulnerability research can cause significant harm to systems and the people who rely on them, this type of research should be undertaken with care.

Our work at Columbia University looked at vulnerabilities in routers and other embedded networked devices as they are deployed across the Internet rather than strictly confined to an isolated laboratory. Such embedded networked devices have become a ubiquitous fixture in the modern home and office as well as in the global communication infrastruc-

ture. Devices like routers, NAS appliances, home entertainment appliances, Wi-Fi access points, webcams, voice-over-IP appliances, print servers, and video conferencing units reside on the same networks as our personal computers and enterprise servers and together form our world-wide communication infrastructure. Widely deployed and often misconfigured, they constitute highly attractive targets for exploitation.

We conducted a vulnerability assessment of embedded network devices within the world's largest ISPs and civilian networks spanning North America, Europe, and Asia. Our goal was to identify the degree of vulnerability of the overall networking infrastructure and, having devised some potential defenses, to determine their practicality and feasibility as a reactive defense. To give a sense of the problem's scale, we provide some quantitative data. In our vulnerability assessment, we scanned 486 million IP addresses, looking for a trivial vulnerability: embedded systems with a default password setting to their telnet or Web server interface. Out of the 3 million Web servers and 2.8 million telnet servers discovered, 102,896 embedded devices were openly accessible with default administrative credentials (username and password). Some of these devices were routers or devices managing or controlling the connectivity of hundreds (or thousands) of other devices. Other unprotected devices such as video conferencing units, IP telephony devices, and networked monitoring systems can be exploited to extract vast amounts of highly sensitive textual, audio, and visual data.

In trying to devise defenses for such devices, however, we're forced to acknowledge and think about the ethical questions this technical reality raises: such devices constitute "network plumbing," which few people want to think about or spend much time tinkering with, except when it visibly fails. Even with the active support of router manufacturers, expecting that users would update their embedded systems (which aren't as straightforward to update as the typical desktop or laptop operating system) isn't an optimal strategy from the standpoint of minimizing harm. In fact, anecdotal evidence suggests that publicizing the vulnerabilities we knew about in the form of a vendor-approved or even vendor-supplied software patch would likely cause more damage—such a move would attract the attention of previously unaware attackers and create a sense of urgency in exploiting these vulnerabilities before they "disappear." Reactive defenses, on the other hand, could sidestep these issues by hardening those systems without any action by the device owners, in response to a detected attack (whether against a specific device or the network as a whole).

However, this entire line of research raises a slew of ethical, moral, and even legal questions. Is our vulnerability assessment of the Internet (or a large fraction of it) ethical? Is our disclosure of the assessment and its results ethical? What of the contemplated defenses? Although proactively deploying our defenses without owners' consent across the Internet would likely be viewed as unethical, there's also a reasonable expectation that in the event of a major cybersecurity incident, an organization such as the US Department of Homeland Security would choose to employ such means to defend the critical infrastructure. Where, then, do qualified security professionals lie on this spectrum? What about someone who discovers a weakness in such an attack and rapidly develops and deploys a countermeasure that uses the same attack vector to install itself on still-vulnerable systems? Crude attempts along these lines could be seen in the CodeRed/CodeGreen engagement and Li0n/Cheese worms in 2001, and the Santy/anti-Santy Web worms in 2004.

Based on our experiences and discussions conducting this research, we propose the following suggestions for best practices for ethical vulnerability research.

Disclose Intent and Research

As a first step, the research's intent should be publicly announced, including details about the methods involved in acquiring data or testing devices or products for vulnerabilities. Open communication of this information can be easily accomplished through a well-publicized Web site, such as Columbia University's home router vulnerability assessment Web page at www.hacktory.cs.columbia.edu.

Seek Legal Counsel Prior to Starting Research

Active debate is under way in the courts and legal academic community regarding the appropriate construction of computer intrusion and intellectual property statutes, and researchers should consult advice of counsel prior to commencing a project whenever practicable. Possible sources of inexpensive legal advice for researchers include intellectual property law clinics operated by law schools, university general counsel, law firm pro bono initiatives, and nonprofit organizations concerned about issues of civil liberty and consumer protection such as the Electronic Frontier Foundation.

Be Proactive about Data Protection

At every stage of the research, the team must be informed about the nature and need to safeguard data. There are several important considerations.

- All members of the research team should receive, review, and

preferably sign a "best data practices policy" that states the rules of research conduct and information handling that the principal investigator sets forth for the lab or project at hand. Because graduate students and other individuals working on the research team might be unfamiliar with the data collection, handling, and use practices the principal investigator expects, obtaining the entire team's agreement on the rules of data protection for the project prior to starting will help prevent misunderstandings and careless errors. In the unlikely event that a team member engages in unethical behavior, the existence of this policy demonstrates that the unethical conduct was indeed a transgression, even if the conduct falls in an ethical "gray area."

- Access to any sensitive data used as part of or obtained during the research should be carefully safeguarded with limited access on a "need-to-know" basis only. Certainly, security practitioners must safeguard the customer's confidential information about its own security posture.
- Finally, data should be anonymized to the greatest extent possible in accordance with current capabilities.

Further Knowledge in the Field

Basic research should further the state of knowledge in the field of information security. General results of a scientific or technical nature revealing the scale and scope of significant vulnerabilities should be widely published in the scientific literature; publications should reveal sufficient detail to help other researchers devise new vulnerability assessment methods.

Report Serious Vulnerabilities

Any significant findings of harmful vulnerabilities should be reported, directly or indirectly, to the people who can best correct the problems. The optimal channels for this disclosure are currently a matter of debate in both the legal and information security community. At present, each principal investigator should assess the unique facts and circumstances of the vulnerability and apply the duty not to harm. In other words, an ethical vulnerability researcher will determine which methods of notification are most likely to limit harm to third parties, in particular, users of the vulnerable product and those who rely on its use. In short, the goal of an ethical security researcher's disclosure is always to minimize harm, to reinforce through conduct the goals of the research stated prior to commencement, and improve the overall state of information security.

Prepare the Next Generation of Professionals and Researchers

To better train the next generation of security professional and researchers to combat information security harms, the computer science curriculum should include penetration testing and real-world exploitation techniques. Building and auditing systems effectively to minimize harm requires knowledge parity in practical skills between malicious actors and the security champions who seek to protect innocent third parties.

The technical considerations of any security professional's basic training are quite challenging. The practice of security professionals is perhaps far more complex when we consider the moral and ethical challenges that confront each of us when we apply our knowledge and skills to protect the systems on which we depend. ☐

References

1. S. Gaudin, "Banks Hit T.J. Maxx Owner with Class-Action Lawsuit," *Information Week*, 25 Apr. 2007; www.informationweek.com/news/internet/showArticle.jhtml?articleID=199201456&queryText=Banks%20Hit%20T.J.%20Maxx%20Owner%20with%20Class-Action%20Lawsuit.
2. A.M. Matwyshyn, "Material Vulnerabilities: Data Privacy, Corporate Information Security and Securities Regulation," *Berkeley Business Law J.*, vol. 3, 2005, pp. 129–203.
3. A.M. Matwyshyn, "Techno-consen(t)sus," *Wash. Univ. Law. Rev.*, vol. 85, 2007, pp. 529–574.
4. N.E. Bowie and T.W. Dunfee, "Confronting Morality in Markets," *J. Business Ethics*, vol. 38, no. 4, 2002, pp. 381–393.
5. T.W. Dunfee, "Do Firms with Unique Competencies for Rescuing Victims of Human Catastrophes Have Special Obligations?" *Business Ethics Q.*, vol. 16, no. 2, 2006, pp. 185–210.
6. T.W. Dunfee, "The World is Flat in the Twenty-First Century: A Response to Hasnas," *Business Ethics Q.*, vol. 17, no. 3, 2007, pp. 427–431.
7. P.M. Schwartz, "Notifications of Data Security Breaches," *Michigan Law Rev.* vol. 105, 2007, pp. 913–971.
8. *Digital Millennium Copyright Act*, *US Code*, Title 12, section 1201(i)–(j).
9. *Int'l Airport Centers, LLC et al. v. Jacob Citrin*, *Federal Supplement, 3rd Series*, vol. 440, 2006, p. 419 (US Court of Appeals for the 7th Circuit).
10. *LVRC Holdings v Brekka*, *Federal Supplement, 3rd Series*, vol. 581, 2009, p. 1127, 1137 (US Court of Appeals for the 9th Circuit).
11. R.C. Ford and W.D. Richardson, "Ethical Decision Making: A Review of the Empirical Literature," *J. Business Ethics*, vol. 13, 1994, 205–221.
12. M.S. Schwartz, T.W. Dunfee,

and M.J. Kline, "Tone at the Top: An Ethics Code for Directors?" *J. Business Ethics*, vol. 58, no. 1, 2005, pp. 79–100.

Andrea M. Matwyshyn *is an assistant professor of Legal Studies and Business Ethics at the Wharton School at the University of Pennsylvania. Her research focuses on the intersection of information security and privacy regulation, corporate law, and technology policy. She is the editor of* Harboring Data: Information, Security, Law, and the Corporation *(Stanford Press 2009). Contact her at amatwysh@wharton.upenn.edu.*

Ang Cui *is a graduate research assistant with the Department of Computer Science at Columbia University. His research interests include next-generation botnets and the defense and exploitation of routers and other network embedded devices. Contact him at ang@cs.columbia.edu.*

Angelos D. Keromytis *is an associate professor with the Department of Computer Science at Columbia University and director of the Network Security Lab. His research interests revolve around most aspects of security, with particular interest in systems and software security, cryptography, and access control. Keromytis has a PhD in computer science from the University of Pennsylvania. He's a senior member of the ACM and IEEE. Contact him at angelos@cs.columbia.edu.*

Salvatore J. Stolfo *is a professor of computer science at Columbia University. His research interests include computer security, intrusion detection, machine learning, and parallel computing. Stolfo has a PhD in computer science from New York University's Courant Institute. Contact him at sal@cs.columbia.edu.*

Selected CS articles and columns are also available for free at http://ComputingNow.computer.org.

Interested in writing for this department? Please contact editors Richard Ford (rford@se.fit.edu) and Deborah Frincke (deborah.frincke@pnl.gov).