

Building an Active Computer Security Ethics Community

In spite of significant ethical challenges faced by researchers evaluating modern threats, the computer security field has yet to grow its own active ethics community to describe and evaluate the ethical implications of its work.

DAVID
DITTRICH
*University
of Washington*

MICHAEL
BAILEY
*University
of Michigan*

SVEN DIETRICH
*Stevens
Institute of
Technology*

Modern threats such as denial-of-service (DoS) attacks, worms, viruses, phishing, and botnets underscore the need for Internet security research in an increasingly networked and computationally reliant society. Responses to these threats vary from passive observation to calls for the legal right to defend computer systems using aggressive countermeasures.¹ This class of Internet security research is itself at one extreme of a broad spectrum of computer security research that includes embeddable medical devices, automobile and process control systems, electronic voting and payment systems, and personal communication devices. This research involves not only appropriate responses but also difficult issues of privacy and responsible disclosure of vulnerability information.

To better understand the complexities faced by computer security researchers, let's consider two recent case studies (see the "Attack Case Studies" sidebar for a more complete description of these events). The first case involved a 2009 distributed DoS attacks against South Korean and US government and corporate websites. Aggressive actions by Bach Khoa Internetwork Security (BKIS), in which the company remotely retrieved log files and identified IP addresses participating in the botnet, caused significant public dispute and accusations that BKIS violated international law. In the second case study, researchers in Canada investigated a malicious botnet whose victims included the foreign embassies of dozens of countries, development banks, and multinational consulting firms. They used passive monitoring of suspected victim networks

to confirm the intrusions and identify the malware, which

they then reverse-engineered. The researchers gained access to the attackers' command-and-control (C&C) servers to identify the compromised systems.

In the first case, BKIS's actions drew a great deal of attention and controversy in the middle of a media frenzy surrounding the high-profile DDoS attacks. In the second case study, the researchers acted methodically, deliberately, and didn't go public until well after they reported to the victims and the victims' respective law enforcement authorities. Both cases share the following attributes:

- Researchers took active control of malicious botnet C&C servers.
- The attacks targeted high-profile victims, resulting in high-profile news coverage.
- They involved hostile (criminal) activity across international borders.
- The targets included both governmental and non-governmental organizations with ties to sovereign governments in multiple nations.

However, they differ in that the attacks in the first case were fast moving and aggressive (impacting availability), whereas the second involved more subtle and concealed attacks on information and information systems (impacting integrity and confidentiality).

These complementary case studies expose a general

Attack Case Studies

The first case is the 4 July 2009 distributed denial-of-service (DDoS) attacks against South Korean and US government and corporate websites. These attacks drew immediate press attention and concerted efforts to mitigate the damage. On 12 July 2009, Bach Khoa Internetwork Security (BKIS), centered at the Hanoi University of Technology, announced on its blog that it received a request for assistance from the Korean CERT (KrCERT) and information that allowed them to identify eight botnet command-and-control (C&C) servers suspected of controlling the DDoS attacks. BKIS claimed it “fought against C&C servers [and gained] control” of two systems in the UK. It remotely retrieved log files and then counted and geolocated more than 160,000 IP addresses around the world participating in the botnet. Public disputes erupted over BKIS’s actions. The Vietnamese CERT (VNCERT) accused BKIS of violating international law for taking control of the UK-owned servers residing in the US and went public with a complaint it received from KrCERT. BKIS threatened to sue VNCERT for defamation. BKIS claimed it used common tools and practices to discover the vulnerable C&C servers and that accessing those systems remotely “doesn’t require anyone’s permission and anybody can do it.” BKIS justified not reporting to VNCERT during the two-day investigation, citing Article 43 of the Vietnamese government’s Decree 64/2007, which states, “In urgent cases which can cause serious incidents or network terrorism, competent agencies have the right to prevent attacks and report to the coordinating agency later.” The Vietnamese government eventually had to step in.

The second case study involves an Information Warfare Moni-

tor (www.infowar-monitor.net). Between June 2008 and March 2009, researchers in Canada conducted a multiphase investigation of a malicious botnet, the victims of which included the foreign embassies of dozens of countries, the Tibetan government-in-exile, and multinational consulting firms. Initial research using passive monitoring of suspected victim networks confirmed the intrusions and identified the malware, which was then reverse-engineered. Honeypots were then infected and used to collect intelligence on the botnet’s operation and control servers. They “scouted these servers, revealing a wide-ranging network of compromised computers.” Gaining access to the attackers’ C&C front end, they were able to “derive an extensive list of infected systems, and to also monitor the systems operator(s) as the operator(s) specifically instructed target computers.” One year later, a follow-up report was published describing continued investigation of these attacks dubbed the *shadow network*. In this latest report, the researchers describe the principles used to guide their decisions. These include collecting data from compromised computers with the owners’ consent, monitoring the C&C infrastructure enough to understand the attackers’ activities and to enable notification of infected parties at the appropriate time, working with government authorities in multiple jurisdictions to take down the attacker’s C&C infrastructure, and storing and handling data securely. In talking about notification and disclosure of information, the researchers note, “Existing practices in this area are underdeveloped and largely informal. In part, this reflects the fact that global cyber security is still an embryonic field.”

issue in judging computer security research. Even if we can identify the important factors of our research protocols, how do we judge our work’s acceptability? Were the actions in the two case studies justified by their outcomes or were they inappropriate because they performed actions that should never be allowed (that is, remote control of resources the researchers don’t own)? Was the first case any less appropriate because the researchers acted in a way that they were immediate beneficiaries? Fortunately, the field of ethics offers a long history of ethical decision-making that we can rely on to help make sense of precisely these and other issues.²

What Is Ethics?

“The field of ethics (or moral philosophy) involves systematizing, defending, and recommending concepts of right and wrong behavior.”³ Normative ethics is a subfield that seeks to develop a set of morals or guiding principles to influence the conduct of individuals and groups within a population (such as a profession, religion, or society at large). Three main strategies for arriving at these moral standards have emerged over time:³

- *Consequentialism* espouses that the “end justifies the means.” For example, a consequentialist argument regarding torture would evaluate the benefits of the information gained in relation to the loss of an individual’s rights.
- *Deontology*, or duty-based ethics, looks at the rightness or wrongness of the acts themselves and the duty to follow rules. For example, a deontological argument might state that it’s never acceptable to torture anyone, for any reason.
- *Virtue ethics* considers the character of the person making the choice, rather than the act or its consequences. For example, you would consider an individual’s strong moral foundation and history of acting in virtuous ways when evaluating his or her decision to use torture.

The definition of computer ethics has various interpretations in line with these broader definitions.⁴ One of the most oft-cited definitions is from James Moor: “A typical problem in computer ethics arises because there is a policy vacuum about how computer technology should be used. Computers provide us

Table 1. Existing ethical norms. The three frameworks share general issues such as fairness and balancing the benefit of “good” vs. harm or risk, and the protection of innocent parties.

Principle	Question
Societal code	
Defense	Population being protected is identified?
Defense	Looks like use of force?
Defense	Actions are proportional?
Defense	Necessary to repel or prevent harm?
Defense	Benefits of disclosure favor victims over attackers?
Defense	Actions are appropriately directed?
Necessity	Greater moral good defined?
Necessity	No other reasonable options available?
Necessity	Otherwise respectful of rights?
Punishment	Avoids punitive motives?
Retribution	Avoids retributive motives?
Evidentiary	Adequate reason to think preconditions of applying other principles are met?
Professional code	
Do good	Positively impacts human well-being?
Avoid harm	Harms users, public, employees, or employers?
Avoid harm	Efforts made to mitigate or undo negative consequences?
Be honest	Honors property rights?
Be honest	Gives proper credit?
Be honest	Honors confidentiality?
Be fair	Discriminates on basis of race, sex, religion, age, disability, or nationality?
Be fair	Inequities exist between groups?
Privacy	Minimal information collected?
Privacy	Protected from unauthorized access?
Privacy	Data used only for intended purposes?
Academic code	
Respect for persons	Individuals treated as autonomous agents?
Respect for persons	Individuals (or their providers) informed and allowed to consent?
Respect for persons	Individuals with diminished autonomy protected?
Respect for persons	Identities of innocents are protected?
Beneficence	Low potential to inflict harm?
Beneficence	Maximize possible benefits and minimize harms
Beneficence	Risks and benefits systematically evaluated
Justice	Who benefits?
Justice	Fairness (neutrality) of procedures

with new capabilities and these in turn give us new choices for action. Often, either no policies for conduct in these situations exist or existing policies seem inadequate. A central task of computer ethics is to determine what we should do in such cases—that is, to formulate policies to guide our actions.”⁵

Unfortunately, although the rich field of ethics offers us a way to consistently and coherently reason about specific ethical issues, the gap between these approaches and a practical ethical framework is tremendous. In this work we seek to be neither proscriptive nor prescriptive, as we believe it presumptuous to propose such a framework in an area that lacks consensus and shows little active debate. Instead, our goal here is to raise the issue of community involvement. As such our approach is closest to that of Deborah Johnson and Keith Miller⁶ in that we are concerned with building expertise in practical decision-making. Note that we intentionally separate from this discussion the related, but not identical issues surrounding law and computer security (For a brief summary of their relationship, see the “Role of Law Versus Ethics” sidebar.)

Existing Ethical Guidelines

A rich body of research and a long history of ethical decision-making in other fields have resulted in our current set of ethical guidelines for researchers and professionals. Table 1 lists the three existing ethical frameworks we focus on here.

US Academic Standards

In 1947, the Nuremberg Code was the first call for informed consent and voluntary participation in research experiments. The World Medical Association’s Medical Ethics Committee responded in 1954 by writing the Declaration of Helsinki, which was completed and adopted in 1964. This declaration addressed research protocols involving humans in terms of risks and benefits, informed consent, researcher qualifications, and so on, and informed a set of standards, or *good clinical practices* (GCPs). More than a thousand laws, regulations, and guidelines worldwide now protect human research subjects.⁷

In the United States, one of the most well-known cases of medical research abuse involved experiments on low-income African-American men infected with syphilis in Tuskegee, Alabama. These experiments began in 1932 and, although researchers learned in the 1940s that penicillin was an effective treatment, they quietly withheld this information so doctors could see how the disease affected patients as the disease progressed. These experiments continued until they were made public in 1972. In 1974, Congress passed the National Research Act, creating the National Commission for the Protection of Human Subjects of Bio-

medical and Behavioral Research. In 1979, decades after the Nuremberg Code and the Declaration of Helsinki, the National Commission prepared a document known as the Belmont Report.

The Belmont Report describes three basic ethical principles and their application:

- *Respect for persons.* Participation as a research subject is voluntary and follows from informed consent. Individuals should be treated as autonomous agents, and their right to decide about their own best interests respected. Individuals with diminished autonomy, incapable of deciding for themselves, are entitled to protection.
- *Beneficence.* Do not harm. Maximize possible benefits and minimize possible harm. Systematically assess both risk and benefit.
- *Justice.* Each person should receive an equal share in treatments and benefit of research according to individual need, effort, societal contribution, and merit. There should be fairness of procedures and outcomes in selection of subjects.

In 1981, the Department of Health and Human Services (DHHS) issued regulation 45 CFR 46, which, inspired by the Belmont Report, defined requirements for research involving human subjects that apply to individual researchers and their institutions. It also defined the role, responsibilities, and requirement for entities doing grant-funded research to institute and use institutional review boards (IRBs) to oversee this research. In 1991, 15 other US federal departments and agencies adopted 45 CFR 46, Subpart A, in what is now known as the Common Rule, and these rules govern nearly all government-funded scientific research in the US. Table 1 lists academic norms that primarily stem from these human subjects regulations.

While these academic norms govern much of US computer security research, there is no direct analogy between biomedical/behavioral research and computer security research. However, even if the scale or levels of indirection differ, the research in both areas involves real risk of harm. For example, in biomedical research, the researcher might wish to draw blood from a subject to study an experimental drug's effect. The subject is physically present in the research lab and gives consent through a consent form. The number of subjects is typically on the order of hundreds or thousands, and the research proceeds at human speeds (that is, the time necessary to explain the research protocol, read and sign the consent form, draw the blood, and so on). The risk is most often proportional to the number of subjects involved. In computer security research, however, millions of computers—not millions of humans—might be the research subjects. Humans

The Role of Law versus Ethics

The law consists of rules that are recognized by a society and enforceable by some authority. It can impose affirmative obligations to act in certain ways or require people to refrain from certain actions. Although laws are informed by ethics, they are not equivalent and therefore laws aren't entirely congruent with societal ethical norms. For example, we might agree that lying to a friend is unethical, but lying to a friend is not illegal. Lying under oath, on the other hand, is always illegal. Legal and ethical considerations matter to security research in several ways:

- Adherence to ethical principles might be required to meet regulatory or legal requirements (for example, common rule). Conversely, knowing and respecting existing laws might be required by an ethical code (such as ACM).
- A law might identify an individual party's rights and responsibilities, and clarify the line between beneficial acts and harmful ones by defining harm.
- Ethical principals that are adopted by the computer security research community can inform judicial, legislative, and regulatory decisions.
- Where a law is ill-fitting or its interpretation unclear, ethics creates an objective and consistent way for us to reason about the acceptability of our actions.

A full discussion of the relationship between law and ethics for computer security is available elsewhere.^{1,2}

References

1. A.J. Burstein, "Conducting Cybersecurity Research Legally and Ethically," *Proc. 1st Usenix Workshop on Large-Scale Exploits and Emergent Threats (LEET 08)*, Usenix Assoc., 2008, pp. 1–8.
2. D.C. Sicker, P. Ohm, and D. Grunwald, "Legal Issues Surrounding Monitoring During Network Research," *Proc. 7th ACM SIGCOMM Conf. Internet Measurement*, ACM Press, 2007, pp. 141–148.

use those computers, but they might not be the direct research subjects. The researcher likely never sees the humans behind the computers, nor do the humans see the researcher. Yet if the research causes those millions of computers to crash, it could cause a great deal of damage. If news stories of infected computers causing disruption are accurate, planes can crash, power distribution systems can fail, and 9-1-1 emergency-response systems can go silent. It is these not-well-understood potential causes of harm that demand prevention, lest ours, like the previous fields, simply waits until people are harmed before we act, or before legislators write laws that restrict our actions.

Professional Standards and Codes of Conduct

Other bodies have recognized the need to regulate membership and provide guidance on appropriate be-

havior in performing one's duties. Table 1 also lists professional norms that primarily stem from business ethics principles (for example, appropriate handling and use of intellectual property or sensitive information, balancing good versus harm vis-a-vis users or consumers, and ensuring fairness). Unique to this area are intellectual property considerations, such as appropriate credit and respect for information ownership rights.

The three parts of ACM's Code of Ethics and Professional Conduct highlight fundamental ethical considerations, specific professional responsibilities, and leadership imperatives.⁸ Section 1 entreats members to "contribute to society and human well-being" and to "avoid harm to others," along with six other principles (for example, don't discriminate, be honest, respect privacy). Professional responsibilities include calls that "ACM members must obey existing local, state, province, national, and international laws unless there is a compelling ethical basis not to do so," and to "access computing and communication resources only when authorized to do so," along with maintaining competence, accepting review, and so on.

Similarly, IEEE maintains the IEEE Code of Ethics, which, although more abbreviated than the ACM version, contains many of the same imperatives.⁹ Specifically, the code commits members "to the highest ethical and professional conduct." Members agree to avoid conflicts of interest, be honest, engage in responsible decision-making, accept criticism of work, and so on. In addition, IEEE and ACM have approved a joint Software Engineering Code of Ethics governing these practices in both bodies.¹⁰

The professional certification organization (ISC)² requires members to pledge adherence to a code of ethics that includes the following canons: "Protect society, the commonwealth, and the infrastructure; Act honorably, honestly, justly, responsibly, and legally; Provide diligent and competent service to principals, and; Advance and protect the profession."

These are certainly not the only such codes of conduct nor the only works on ethical issues facing computer professionals. Moreover, numerous other professional organizations, whose headquarters are outside the US (such as the Institute for the Management of Information Systems in the UK, the Australian Computer Society, and the Canadian Information Processing Society), as well as individual companies, and academic institutions have their own ethical codes.

International Rules of Engagement

Cybersecurity is often analogized with physical conflict (for example, cyberwarfare), and ethical discussions of responding to computer attacks often focus on "use of force" or self-defense analogies. Actions in war are governed by the Law of Armed Conflict, which

requires involved parties to meet certain prerequisites for lawfully using force. These include military necessity (force is required only to the point of meeting military objectives), distinction (actions must be directed against lawful combatants and military targets, not against civilians and civilian infrastructure), and proportionality (use of force must be less than or equal to the original harm or violation). As a result of international agreements and protocols,¹¹ militaries around the world operate under strict rules of engagement that guide decision-making on the battlefield. David Dittich and Kenneth Himma discuss the legal and ethical frameworks for responding to computer intrusions based on these guidelines.¹ They identify three core ethical principles: defense, necessity, and evidentiary. These principles not only apply to aggressive countermeasures, but they also include actions taken by today's researchers involving accessing or controlling systems outside one's own administrative control. Table 1 lists these principles and related questions.

Limitation of Existing Standards

Although each of the ethical frameworks has something to offer, current conversations at program committees, conferences, and funding panels suggest that these standards fail to bring the clarity required to make important decisions regarding funding, execution, and publication of modern computer security research. We argue that these failures result from three higher-level limitations in computer security research ethical decision-making.

Absence of Shared Community Values

Of all the existing community ethical standards and codes of conduct listed previously, none encompasses all of the issues listed in Table 1. Although the societal and academic codes provide a rich body of knowledge and expertise, nobody envisioned the advent of information and communication technology (ICT) research when they were conceived, let alone the narrower field of computer security research. The professional codes, based on the seminal work of Donn Parker¹² and others who recognized the importance of ethics in information technology, have seen little evolution over time. Another limitation of these codes is that they are often narrow and specific. The ACM code is cited as one guide that program committees can apply in judging academic papers submitted to them for review; however, Mark Allman mentions that interpretations can be varied and application of the code to specific actions difficult.¹³

Lack of Consensus on Enforcement

Even if standards for determining the ethics of various research methods existed, who enforces these stan-

Table 2. The ability of various entities to achieve the goals of ethical decision-making.

Entity	Program committees	Institutional review board	ACM/IEEE	NSF/DARPA
Inspiration	Low	Low	High	Low
Education	Low	Medium	High	Low
Guidance	Medium	Medium	Medium	Medium
Accountability	Medium	High	Low	High
Enforcement	High	High	Low	High

dards and how? The various bodies that could review prospective research protocols, provide guidance to researchers, and enforce adherence are limited. Table 2 lists the comparative strengths and weaknesses of these bodies when compared with the functions of ethical codes and guidelines as defined by Terrell Ward Bynum and Simon Rogerson.⁴

IRBs seem best suited to processing and reviewing applications; however, the existing drawbacks of narrowness of scope, lack of technical expertise, lack of existing regulatory authority, US-centrism, and lack of prospective guidance for researchers limit their usefulness.^{13,14} Funding agencies such as the US National Science Foundation (NSF) and DARPA could similarly provide accountability and enforcement functions, but are also limited by lack of educational resources and prospective guidance for proposers. Moreover, they have authority only over those seeking funding from them. Program committees are more general, international in scope, and already provide a peer-review function. They're not, however, uniform in their membership. In addition, they're typically involved after research is completed, and they perform their function semi-anonymously and in private. Professional associations often have ethics boards, require acceptance of the association's code of conduct as a condition of membership, and provide an educational and inspirational role for their members. However, their authority extends only to members, and enforcement is limited to expulsion from the society.

Limited Individual Expertise

Even though many biomedical and behavioral researchers regularly work with IRBs, computer and information scientists might never have dealt with an IRB. Many universities offer courses in professional ethics and philosophy; however, they might be elective courses not taken by computer science students. Often these classes don't focus on the application of ethical methods to computer science or computer security. Without these experiences or formal training, community members are limited in their ability to have structured debates about ethical issues. All too often, discussions devolve into vague is/isn't state-

ments based on personal beliefs and reviewers fluidly switching between distinct moral theories in an attempt to reason about the applicability of a particular piece of research.

Moving Forward as a Community

Although the limitations we've discussed are daunting, the potential risk of not moving forward—that a single failure could have a chilling effect on all computer security research—motivates the need for action. Whereas this article's primary purpose is to raise awareness of these issues, we can each do several things to help build an active and thriving computer security ethics community.

Building Personal Ethical Decision Making Abilities

Analyzing case studies provides an excellent mechanism for building practical ethical decision making abilities. Bynum and Rogerson suggest a multistaged approach to case study analysis:⁴

- identify key ethical principles,
- detail the case study,
- identify specific ethical issues raised by the case,
- call on personal experience and skills for evaluation and then on the abilities of others, and
- apply a systematic analysis technique.

Clearly, a person's key ethical principles will vary based on such factors as fundamental ethical approach (for example, consequentialism or virtue ethics) and can change depending on the normative values of one's culture. Although the specific ethical issues raised by a case will depend on the combination of these ethical principles and the case in question, we've found the following generic questions to be useful in building ethical decision-making expertise across a wide range of cases and philosophies:

- Does the research aim to protect a specific population, and if so, which population (for example, the owners of infected hosts, the victims of secondary attacks, or the general Internet user)?

- Can studying malicious behavior achieve multiple simultaneous benefits to society (for example, developing new defenses while aiding investigation of criminal acts and assisting victimized network sites)?
- Who benefits more from publication of research findings, and in what order (for example, victims of criminal acts, the researchers themselves, or the criminals perpetrating computer crimes)?
- Is there another way to accomplish the desired research results?
- What is the safest way to disseminate research results without risking improper use by individuals who might not share the researchers' ethical standards?

In an effort to foster case study analysis in the community, we have provided a wide range of useful case studies in a related technical report.²

Consistency

One of the key challenges as we start to build a community of computer security ethics is consistency in how we undertake our ethical analysis. All too often, researchers slide between different sets of ethical philosophies and norms, or focus narrowly on a single issue or benefit perspective, to justify or critique research. To build consensus across a wide range of research and situations, it's advantageous to explore formal methodologies for evaluating these questions.⁴ One method we see frequently is the utilitarian view of consequentialism, which seeks to balance the benefits and harms of any research. In this model, we have found *stakeholder analysis* effectively elicits the potential benefits and harms.

Stakeholder analysis identifies key players in a situation in terms of their interests, involvement, and relationship (that is, producer or recipient) to outcomes such as benefit or harm. In normal use, these stakeholders are involved in the positive outcome of a project:

- Primary stakeholders are “those ultimately affected, either [positively or negatively].” They're typically a computer system's users and consumers of information or information system products or services.
- Secondary stakeholders are “intermediaries in delivery” of the benefits and harms. In the computer security context, they're service providers, operators, or other parties responsible for integrity, availability, and confidentiality of information and information systems.
- Key stakeholders are “those who can significantly influence, or are important to the success [or failure] of the project.” We include the researchers, vendors, designers, and implementers of a system, as well as criminals or attackers.

Although certainly not the only such method for performing ethical analysis, the key here is the use of a systematic approach, consistently applied, and coherently articulated.

Integrity and Accountability

In addition to the ethical questions raised here, researchers should indicate in their publications how they (or others) evaluated their work. Program committees will always need to enforce some amount of discussion about why researchers think their protocols sufficiently account for all the ethical issues brought up in their work. Explicitly discussing the ethical considerations of the higher-risk aspects of proposed research means the program committee members don't need to guess or infer. Indicating review and approval of a submission by an outside oversight organization (such as an IRB) highlights these approvals. A great example of this is The 2011 Symposium on Usable Privacy and Security (SOUPS) that, in its call for papers, indicates, “Papers should mention how the authors addressed any ethical considerations applicable to the research and user studies, such as passing an IRB review.”

Involvement

The recent uptick in conversations about ethics and ethical security practices hasn't gone unnoticed, and a wide variety of opportunities to discuss these issues has emerged. Several recent security conferences (for example, the Workshop on Large-Scale Exploits and Emergent Threats [LEET 2009] and the Network & Distributed System Security [NDSS 2010] conference) offered ethics panels to discuss relevant work and encourage lively (and often entertaining) discussions. The Workshop on Ethics in Computer Security Research (WECSR), now in its second year, offers a venue to discuss the specific problems and ethical application discussed here. Actively serving on an ethical oversight committee (such as IRBs in US academic institutions) offers an excellent opportunity to both learn the process of human subjects protection and make a local difference.^{13,14} Professional organizations, such as the IEEE Ethics and Member Conduct Committee, also offer opportunities to get involved.

Assertion of the Right to Self-Governance

Charles Ess might have been prophetic when he said, “Either we regulate ourselves, or they'll do it for us (and they will do it much worse than we will).” If we don't assert our right to self-governance and follow through by self-policing, we might find regulation forced upon us. One such effort at self-determined standards and enforcement, sponsored by the US Department of Homeland Security's Science and Tech-

nology Directorate, is already underway. DHS hosted a two-day ethics workshop on 26–27 May 2009 in Washington, DC. Inspired by the Belmont Report, the workshop brought together ethicists, IRBs, researchers, and lawyers to discuss these pressing issues. The primary anticipated outcome from this meeting and a wide range of follow-on working group meetings is a set of ethical guidelines, which, though anchored off the original Belmont framework, reflects the unique questions facing ICT researchers. Publically available drafts of these guidelines and companion documents are expected in 2011. Using early drafts of this report, community researchers have also constructed an ethical impact assessment (EIA) tool as a guide to help ICT researchers think about the ethical impact of their work.¹⁵

Take Forward Lessons

In spite of the drawbacks we listed earlier, the existing ethical models offer important lessons when building a new enforcement and oversight mechanism rooted in the computer security community. Although the regulatory nature and unfunded mandate might prevent researchers from directly applying the US IRB model, the committee makeup, review processes, and application mechanism are relevant to any new enforcement mechanism. Any such solution should adapt to integrate with international bodies that cover computer security researchers throughout the world. To include computer security researchers both in academia and industry, any new model would need the authority to review research protocols without any professional membership or university prerequisite. Any new mechanism should ideally evaluate computer security research proposals before research begins. In order to ensure consistency and transparency, any such mechanism should make publically available their study recommendations, when doing so does not impact the researcher or cause other harms.

Reward Ethical Behavior

In its evaluations of research, the computer security community rewards novelty, technical difficulty, and academic rigor. We need to not only police our own community to exclude unacceptable behavior but also reward ethical behavior. Studies that examine particularly thorny security issues, but take great pains to avoid negative consequences, deserve our attention and praise. We shouldn't punish a security study reporting on an event a year after its occurrence if that year was used to assure the security of the affected systems, the privacy or safety of its users, or is otherwise following guidance such as maximizing benefits and minimizing harm for all affected stakeholders. As the pressure to publish is so strong, we might need

to develop new publication venues that let researchers receive credit for their work, but delay full public access for a reasonable amount of time.

Media-driven hype aside, there clearly is a serious threat posed by computer crime and espionage. However, without an active and engaged computer security ethics community, the ability for us to clearly and consistently describe and evaluate the ethical implications of our work is severely impaired, and we risk the worst—that the computer security community might repeat its own version of the abuse and harm that accompanied research in other fields. Given the scale and scope of potential harm in computer security research, can we really risk waiting until there is massive harm before we act? □

Acknowledgments

We thank Doug Maughan and the members of the Menlo Report working group for valuable discussions. We also thank Peter Neumann, Aaron Burstein, Kaiti Carpenter, Erin Kenneally, and the anonymous reviewers for their feedback and assistance in preparing this article.

References

1. D. Dittrich and K.E. Himma, "Active Response to Computer Intrusions," *Handbook of Information Security*, vol. III, John Wiley & Sons, 2005, chap. 182, pp. 664–681.
2. D. Dittrich, M. Bailey, and S. Dietrich, *Towards Community Standards for Ethical Behavior in Computer Security Research*, tech. report CS 2009-01, Stevens Inst. of Technology, 2009; <http://staff.washington.edu/dittrich/papers/dbd2009tr1>.
3. J. Fieser, "Ethics," *Internet Encyclopedia of Philosophy*, 2010; www.iep.utm.edu/ethics.
4. T.W. Bynum and S. Rogerson, *Computer Ethics and Professional Responsibility: Introductory Text and Readings*, Blackwell Publishers, 2003.
5. J.H. Moor, "What Is Computer Ethics?" *Metaphilosophy*, vol. 16, no. 4, 1985, pp. 266–275.
6. D. G. Johnson and K. W. Miller, eds., *Computers Ethics*, Prentice-Hall, 2009.
7. Office for Human Research Protections, Int'l Compilation of Human Research Protections, 2011, www.hhs.gov/ohrp/international/intlcompilation/intlcompilation.html.
8. ACM Council, *Code of Ethics and Professional Conduct*, Oct. 1992; www.acm.org/about/code-of-ethics.
9. IEEE Board of Directors, *IEEE Code of Ethics*, Feb. 2006; www.ieee.org/about/corporate/governance/p7-8.html.
10. D. Gotterbarn, K. Miller, and S. Rogerson, "Software Engineering Code of Ethics," *Comm. ACM*, vol. 40, no. 11, Nov. 1997, pp. 110–118.

11. Int'l Committee of the Red Cross (ICRC), *The Geneva Conventions: The Core of International Humanitarian Law*, Jan. 2006; www.icrc.org/eng/war-and-law/treaties-customary-law/geneva-conventions.
12. D.B. Parker, *Ethical Conflicts in Computer Science and Technology*, AFIPS Press, 1979.
13. M. Allman, "What Ought a Program Committee to Do?" *Proc. Usenix Workshop on Organizing Workshops, Conferences, and Symposia for Computer Systems*, Usenix Assoc., 2008, pp. 1–5.
14. S.L. Garfinkel, "IRBs and Security Research: Myths, Facts, and Mission Creep," *Proc. Usability, Psychology, and Security (UPSEC 08)*, Usenix Assoc., 2008, pp. 13:1–13:5.
15. E. Kenneally, M. Bailey, and D. Maughan, "A Tool for Understanding and Applying Ethical Principles in Network and Security Research," *Proc. Workshop on Ethics in Computer Security Research (WECSR 2010)*, LNCS 6054, Springer, 2010, pp. 240–246.

David Dittrich is a security engineer and information security researcher at the University of Washington Applied Physics Laboratory. His research interests include security operations, focusing on distributed attacker tools and distributed/collaborative responses to such attacks. Dittrich has a BS in computer science from Western Washington University. Contact him at dittrich@u.washington.edu; <http://staff.washington.edu/dittrich>.

Michael Bailey is an assistant research scientist at the University of Michigan, where he studies the security and availability of complex distributed systems. His research interests include the characterization of specific network threats (such as worms, botnets, and spam) and techniques for measuring these threats at scale (such as network anomaly detection and distributed network telescopes). Bailey has a PhD in computer science from the University of Michigan. He is a senior member of IEEE. Contact him at mibailey@eecs.umich.edu.

Sven Dietrich is an assistant professor in the computer science department at the Stevens Institute of Technology. His research interests include computer and network security, denial of service, botnets, anonymity, and privacy. Dietrich has a doctor of arts in mathematics from Adelphi University. He is the vice chair for the IEEE Computer Society Technical Committee on Security and Privacy, and a member of the ACM, the IEEE Computer Society, and the New York Academy of Sciences. Contact him at spock@cs.stevens.edu.



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