Drawing the Line

Carl E. Landwehr
Editor in chief

Ethical behavior has concerned computer security researchers and practitioners for decades. A quarter-century ago, the publication of Fred Cohen's papers defining and describing computer viruses ignited a discussion about the ethics of publishing attack information. Since then, we've seen lengthy discussions of the proper behavior for those who find vulnerabilities in systems—whom they should contact, how long they should wait before making findings public, whether it's ethical to pay people to find vulnerabilities, and so on. But the continuing intertwined evolution of technology and society keeps raising new ethical issues. Consider:

- A system security administrator discovers malware on a server that stores personally identifiable information. It will take a substantial effort to determine whether the malware has actually exfiltrated any of the information, and it might in fact be impossible to determine this with certainty. What are the administrator's responsibilities?
- A computer scientist studying a communications protocol with a new method of analysis discovers a previously unknown vulnerability. The protocol is widely used. Should she write a paper touting the success of the analytic method and submit it for publication, write a research proposal for further development of the analysis method, notify the maintainers of the protocol of the problem, or offer information about the protocol vulnerability for sale on the black market?
- A software engineer tasked with implementing the design for a new application realizes that security and privacy considerations have not been attended to. The problems will occur in the future, and deadlines for delivering the software are looming. What alternatives should the engineer pursue?
- A researcher with access to statistical summaries of data from a social networking site discovers that a new way of looking at the data not only reveals new relationships but also makes it possible to infer the identities of individuals in the study who had been assured of their anonymity. What action should he take? What are the responsibilities of the organization that released the statistical summaries?
- A conference program committee receives a research paper that reports some fascinating new results concerning observed human behavior in cyberspace, but the data were collected from machines whose owners unwittingly downloaded software from a public Web site. Should the committee judge the work solely on the technical merits and accept it, or should it consider whether the data were collected ethically in making its decision? And what does it mean for the data to be ethically collected?

Although these situations raise interesting points for discussion (and Fred Cohen addresses one of them elsewhere in this issue), let's focus on the last one. The security research community today is increasingly asked to provide a scientific basis for its work and quantified evidence of improvements in security. Both of these imperatives lead to a demand for more data, either from controlled experiments or real-world observations. The demand for more usable security functions also motivates data collection involving human subjects.

The rise in botnet activity and financially motivated computer crime in the past few years has led to several research initiatives to study the botnets' structure and operation. But to study a botnet, you need to detect it and even get inside of it. Because botnets generally operate by inserting software on an unwitting user's computer, one tactic is to insinuate the measurement software into the botnet's structure. This could mean inserting software on that unwitting user's computer. Because the user's computer is already compromised, and we expect an ethical researcher at minimum to do no additional harm, perhaps such actions are ethically justified. Some researchers have included explicit discussions of measurement ethics in their papers to address this point.[1] Even so, the compromised user is now also an unwitting experimental subject.

Research funded by the US government that involves human subjects must in general be reviewed and approved by the Institutional Review Board (IRB) of the institution receiving the funding. IRBs were created to prevent the recurrence of the scandalous Tuskegee syphilis medical experiments that were finally terminated in 1972. The scandal led to the creation of a Commission for the Protection of Human Subjects of Biomedical and Behavioral Research. In 1978, that commission produced a 12-page document known as the Belmont Report (US Gov't Printing Office, 1978) that's now the basis for IRBs throughout the USA.

Because of this history, most IRBs have much deeper experience reviewing proposals from biologists and psychologists than from computer scientists and engineers. But as Simson Garfinkel has noted,[2] computer scientists and particularly people working in security research need to become much more familiar with IRBs, the scope of research they cover, and the ethical principles their research proposals need to incorporate.

One of the fundamental principles is that of informed consent. Only adults are deemed capable of providing informed consent—minors can only "assent" and require parental consent. But on the Internet, how does the experimenter establish the subject's age? We might also need to reconsider just what it means to involve a "human subject"—is the analysis of publicly available information (say, a blog post), where authors are clearly identified, a human subjects issue?

Returning to botnet research, some program committees have recently devoted considerable time discussing precisely the questions posed in the last bulleted example. If an IRB has approved the research, should the program committee accept that judgment? Mark Allman[3] has argued that today IRB approval is necessary but not always sufficient, precisely because their staffs might lack the appropriate expertise.

In May 2009, the US Department of Homeland Security's Science and Technology directorate organized a workshop on Basic Ethical Principles for Network Research at which a wide range of potential application areas and experiments were discussed. The workshop report is still in preparation at this writing, but should appear within the next few months.

While I have focused on only one of the examples in my earlier list, the others deserve attention as well. How should our field proceed? Here are three actions we can take:

- Security researchers need to inform themselves about the ethics of human subject experimentation. There is a growing literature on the ethics of data collection from networks, and those proposing and conducting research have an obligation to inform themselves on these issues. For example, the National Academy of Engineering has recently expanded its Online Ethics Center (www.onlineethics.org) which provides case studies and course materials that illuminate several of the examples I presented. That said, the site will benefit from further contributions focused on cybersecurity ethics. An online journal for Internet research ethics published its first issue in 2008 (http://ijire.net).
- IRBs need to develop staffs who are well-informed about the situations and consequences of human subject experimentation in cyberspace. Ultimately, it should be appropriate to rely on an IRB's approval to assure that a given research proposal in fact satisfies ethical guidelines. That is not yet the case.
- Professional societies need to begin developing ethical guidelines for cyberspace monitoring and experimentation. The IEEE, the ACM, and Usenix provide the major publication outlets for research in this field. If these three societies could work together to develop a set of ethical guidelines, the guidelines could have the broad impact we would hope to see. IRBs are required only for US government-funded research. A broadly accepted set of guidelines covering professional society members, who work in many countries and for private firms as well as governments and universities, could have global reach.

Even if it takes time, as it surely will, to establish consensus on the complex issues involved, it's imperative that we keep the discussion of ethics and cybersecurity on the front burner. □

References

1. C. Kanich et al., "Spamalytics: An Empirical Analysis of Spam Marketing Conversion," *Proc. 15th ACM Conf. Computer and Communications Security*, ACM Press, 2008. pp. 3–14.
2. S.L. Garfinkel, "IRBs and Security Research: Myths, Facts, and Mission Creep," *Proc. 1st Conf. Usability, Psychology, and Security* (UPSEC), Usenix Assoc., Apr. 2008; www.usenix.org/event/upsec08/tech/full_papers/garfinkel/garfinkel_html.
3. M. Allman, "What Ought a Program Committee to Do?" *Proc. Workshop on Organizing Workshops, Conferences, and Symposia for Computer Systems* (WOWCS), Usenix Assoc., Apr. 2008; www.usenix.org/event/wowcs08/tech/full_papers/allman/allman_html.