All.Net Analyst Report and Newsletter

Welcome to our Analyst Report and Newsletter

Ethics in security research

The Menlo Report[1] proposes a framework for ethical guidelines for computer and information security research. As such, it may have far-reaching and longstanding effect on the field and advances within it. Regardless of your views on this subject, if you are interested in the future of the field and its impacts on society, you should read and comment publicly on it.[2] This short report summarizes and discusses my public comments, submitted in January of 2012, during the key comment period.

Some general comments on my review and the document as a whole

My review of such documents is usually cursory at best. I take a brief read and comment on what I see as I see it. Like, I think, most such commenters, I don't have a lot of spare time for unfunded reviews of government documents. And when I do get the time to do it, the form of such interactions is quite limited. I think the online comment section has 1K or so of space per comment. So I make several. The alternative is to write a more comprehensive document on the subject and submit it, but this is atypical.

The group that worked on this document seemed to me to be sincere and to have a desire to help move the process forward. They thought of some issues and apparently ignored or failed to think of others. Perhaps the process, objectives, time frames, or other constraints limited their effort, and they came up with a 15 page document which is probably plenty long, but seems to lack the necessary references and research to support such an important potential result.

I have performed human subjects research in the information protection arena, including a series of experiments involving psychological effects of deception on attackers as a defensive approach.[3] In our research, I found that there are potentially serious negative consequences to the subjects of such experiments, and ultimately that we need to be deeply concerned about such experiments in terms of their effects. Of course these were experiments intended to have psychological effects on individuals and groups attacking computer systems. But I have also read and attended many publications and talks in which ethical issues with human experimentation were not adequately considered, and I consider these to be very problematic. I am currently working on a research grant in which human subject experiments are also a key component, and it is turning out to be problematic for the overall program as well. These issues can be, have been, and will be resolved, but I still think it is vital to seriously consider and understand them, and for you to comment on your views of these approaches before they get codified into a situation where they have some level of legal standing.

Here then are my specific comments regarding the Menlo Report, with modification and extension to what I posted on the public comment site.

---

1 "The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research", DHS-2011-0074, September 15, 2011.
2 http://www.regulations.gov/#!searchResults;rpp=10;po=0;s=DHS-2011-0074
3 See http://all.net/ → "Research" → "Deception for Protection" for details.

Specific comments

Comment 1: *"Respect for Law and Public Interest"* (the section)

> fails to indicate the seemingly obvious requirement not to break laws. Many researchers disagree with laws, and while social disobedience is a fine ethical stand, it is not appropriate for researchers to apply this in their research with human subjects or other persons. This has been done repeatedly in the recent past in the computer security area.

Statements like: *"Impinging on the privacy of an organization may be ethically justified if it yields substantial social benefit through an increase in cybersecurity,"* are problematic in many ways. For example, and without limit;

> (1) Why does a researcher who does not know the internal details of the potential impact on a company or its workers, shareholders, or other indirect equities get to choose to violate the confidentiality of that company's information and potentially put others at risk?
>
> (2) Given that legally, corporations are people (in the US), why is it less justifiable to violate a corporate person's privacy than a non-corporate person's? And
>
> (3) What qualifies a researcher to make the judgment regarding the benefit to society vs. the harm to an organization?

It seems to me that this is granting a researcher the right to make judgments that they are not qualified to make (in most cases). The rules of the road should be such that the knowledge and expertise of the researcher is adequate to make the necessary judgments, and this sort of judgment seems beyond the scope of the typical computer researcher.

Added commentary: I think that there is a deep need for those considering such research to undergo the same sort of education and training required of those who do psychological experiments, but in the context of their own field. I also think there is a need for similar sorts of formal protection against the publication of experiments that are not properly vetted through institutional review boards (IRBs), that those boards should have the proper background to understand the nature of experiments as they do in other fields, and that there should be clear standards applied as they are in other fields.

Comment 2: *"As intermediaries between a research and end users, they may be in a position of authority to serve as proxies for consent on behalf of their customers when it is otherwise impracticable for the researcher to individually obtain informed consent from end users."*

> This - to me - is an outrageous statement. The only person who may be informed in order to consent is the actual human subject. The notion that a company that has generically informed users and/or employees of possible future research in a generic way can grant third parties the right to harm those individuals in arbitrary and undisclosed ways is adverse to all of the foundations of ethical research.
>
> Furthermore, informed consent also implies that the subjects be informed of the results of the research and/or deceptions involved in performing the research, which the researcher is responsible for, and not the organization(s) they contacted for data use.

I understand that it makes it easier and more convenient to do research, but it seems to me to be highly unethical, for example, for researchers to get Google's permission to use any and all information about all of its users for studies of the correlation between usage patterns and diseases (similar studies have been done for early influenza detection using Internet search terms) with the potential risk of then releasing details on which people have what diseases. The myriad claims that statistical results don't reveal individual subject details have been repeatedly shown to be unsupported by the facts through public disclosures from such study results about public (and private) individuals. The state of the art is inadequate to protect privacy in this context today.

Comment 3: *"Where feasible, researchers should obtain informed consent to collect, use, or disclose sensitive identifying data, or to interact with information systems in ways that could negatively affect those systems or their users."*

Where NOT feasible, researchers should not be permitted to perform this research.

It is too easy to claim it is infeasible or not spend the time and money necessary to get informed consent. But today, in practical terms, this means that researchers proceed nonetheless. This is the wrong default for researchers. The rule must be clear that without the informed consent of each and every individual human subject who might be impacted, the research may not proceed.

Also, the term *"identifying information"* in this context implies a major lapse in coverage. Research may adversely affect human subjects and third parties even if no identifying information is revealed or gathered. For example, deceptions in experiments can do harm to subjects and effect their lives even if no personal information whatsoever is collected or revealed.

Informed consent should not be permitted by a company for its employees, as this implies coercion, and in particular, such things as notices used at logins and other similar terms of employment should not be allowed to be used as a basis for the claim of informed consent. This should be made explicit and explicitly prohibited.

The same is true of contracts (e..g, with ISPs) students (e.g., those who attend classes should be able to learn without being test subjects) and other legal mechanisms that result in coercion through indirect means.

Comment 4: *"There are justifiable reasons why it may be impracticable to obtain informed consent. In ICTR the frequency of this occurring may be greater than in traditional human-centered research. Of the three components in the informed consent process – notice, comprehension, and voluntariness – providing notice may be particularly challenging given the scale and scope of many operational ICTR environments. It may be impracticable, it may not be technically feasible to identify subjects, or it may interfere with scientific integrity of the results."*

In these cases, since informed consent is NOT obtained, the research should not be permitted to go ahead. Otherwise, we will simply see lots of research claiming technical infeasibility of informed consent going ahead and harming people.

"It may be infeasible to identify, or obtain consent from millions of users whose everyday communication generates traffic across a heavily aggregated backbone link in a traffic modeling study."

This is a mistaken confluence of the notion of a "human subject".

“Or it can be onerous to attempt to inform the owners of hundreds of thousands of compromised home computers that are being used as a single instrument of criminal activity (i.e., a botnet) under study."

So the "ethical" thing to do is to violate constitutional rights and digitally enter the homes of millions of people without their consent or a search warrant - potentially funded by the government - law enforcement no less (acting as an agent) - so you can get information for your research on criminal activities? This is likely illegal, certainly unethical, and outrageous. It's like saying that it's too inconvenient to get permission for DNA tests, so we will covertly collect skin samples from everyone on the subway. This is not ethical research.

Comment 5: *"Even with reasonable measures to detect and reduce potential harm, the malicious software being studied could still accidentally infect other computers. However, regardless of researcher actions, those computers would have been infected when the malicious software propagates at the attacker’s direction"*

I think a researcher looking into biological weapons has a special requirement to prevent their release, even if others might be infected anyway. A researcher who intentionally gets a computer infected has a special responsibility to not allow others to be harmed by it. Thus the researcher has a standard beyond that of the ordinary person and the claim that others "would have been infected" is outrageous - we cannot predict that future. If the researcher is not competent to protect from the spread of the disease to others, they should not be permitted to do such research. That would be highly unethical.

Comment 6: The document fails to adequately address the harm associated with social spreading of information. Research involving deception, for example, which is quite common today and likely to become more so in the future, brings with it the potential to spread rumors and misinformation, to cause psychological and other harm to individuals, and has the potential to cause operational faults and failures at a wide range of scales and with long-term indirect implications. Privacy is not the only issue to be addressed, even though it is the primary issue addressed in the document provided. Effects on integrity are poorly addressed (e.g., the infection of others by researchers), availability is given only minimal attention (e.g., denial of services produced by research and its indirect effects), accountability is not addressed (i.e., the need to have a complete record of all activities of he researchers available so that others can assess and understand the harm done to them and others in retrospective), use control issues are ignored (e.g., the researcher gains information like access codes, which are illegal to possess without authorization, and altered software in user systems which might have unpredicted effects depending on the mechanisms in place in those systems). There is little of a basis within this document on "ethics". I see essentially no ethical understanding displayed or discussed. It is merely a list of conclusions without the reasoning behind them. As such, it does a poor job of justifying itself and fails to address the real practical needs of the community.

