Have we Crossed the Line? The Growing Ethical Debate in Modern Computer Security Research

David Dittrich
University of Washington
dittrich@u.washington.edu

Michael Bailey
University of Michigan
mibailey@eecs.umich.edu

Sven Dietrich
Stevens Institute of Technology
spock@cs.stevens.edu

ABSTRACT

If someone has the ability to take control of a botnet, can they just clean up all the infected hosts? Can we deceive users, if our goal is to better understand how they are deceived by attackers? Can we demonstrate the need for better methods, by breaking something that people rely on today? To be effective, we must find ways to balance societal needs and ethical issues surrounding our research, lest we drift to the extremes—becoming the very thing we deplore, or ceding the Internet to the miscreants because we fear to act. In this paper, we advocate for a community dialogue on the ethical issues in computer security and the ethical standards that we intend to enforce as a community.

1. INTRODUCTION

Modern threats such as Denial of Service Attacks, Worms, Viruses, Phishing, and Botnets underscore the need for security research in an increasingly networked and computationally reliant society. Unfortunately, as our understanding of these phenomena have grown, so has the uncertainty in the computer security research community on the appropriate ways in which to observe and address these problems.

For example, consider the area of botnet research, which centers around the detection and mitigation of large numbers of infected hosts, or *bots*, networked into a single distributed system, or *botnet*. We have recently seen a steady increase in criminal activity using botnets. In response, we see an increase in academic research and federal funding to counter this threat. This criminal activity is compounded by the emergence of politically motivated attacks, such as those against elements of the cyber-infrastructure of Estonia. Responses to these threats vary, from passive measurement and observation, to calls for the legal right to defend computer systems from attack using aggressive countermeasures.

Unfortunately, the structured public discussion of an ethical framework to guide decision making about actions taken while researching and countering botnet attacks, and indeed in a broader set of computer security research, has not kept pace. Existing structures for determining the ethical behavior (e.g., Institutional Review Boards (IRB), Professional Codes of Conduct) fail to provide detailed actionable guidance due to the absence of technical expertise in this specific domain and a lack community shared values [1]. There is growing frustration expressed by researchers, program committees, and professional organizations about the limits of ethical research and who has responsibility to enforce them [1, 4].

Our primary goal in this work is *to encourage a continuing dialogue on the ethics of computer security research.* Through this dialogue, we hope to build both an expertise that can be used in various policy enforcement bodies (e.g., program committees, IRBs) and will help us arrive at a form of community consensus.

2. UNDERSTANDING THE NEED

LxLabs, a company based in Bangalore, India, markets a web server virtualization system called *HyperVM*, which uses an administration interface named *Kloxo*. One company who uses HyperVM and Kloxo is UK-based Vacert.com. On Sunday, June 7, 2009, Vacert.com suffered a compromise of their web hosting system, resulting in over 100,000 accounts being deleted from the system. On Monday, June 8, 2009, LxLabs' CEO, 32 year old K T Ligesh, was found dead in his apartment of an apparent suicide. [6]

Just a few days before, on June 6, 2009, an analysis of "several dozen vulnerabilities in kloxo" with complete details on how to exploit these vulnerabilities was posted anonymously to the web site *milw0rm*. The time line in this analysis describes an attempt by the unknown security researcher to correspond with staff at LxLabs about the vulnerabilities, which includes such problems as file permission bypass, cross-site scripting, symbolic link exploitation, denial of service, and arbitrary command execution at elevated privilege through unclean user input. The posting claims an initial report was sent two weeks prior and that resources demonstrating the vulnerabilities were made available to assist LxLabs in confirming and fixing the problems. After two email exchanges with an unnamed LxLabs employee, no further communication as promised from LxLabs, and no observed attempts by LxLabs to even review the resources, the researcher posted the full analysis and exploit details. Within days, multiple sites using Kloxo (including Vacert.com) were attacked by unknown parties.

It is not known whether there is any relationship between the person(s) who attacked and damaged the web sites and the security researcher who published the vulnerability information, nor the identity of the person who the researcher was in communication with at LxLabs. There is no indication that the security researcher attempted to report these problems to any other organizations. Finally, there is no indication that the researcher considered releasing only partial details in order to warn Kloxo users or the general public and give them a chance to protect themselves prior to release of full details including exploits, as is recommended in various responsible vulnerability disclosure guidelines.

3. BUILDING ETHICAL STANDARDS

While the previous case study may seem extreme it is representative of a growing set of cases [3], which challenge program committees, institutional review boards, and our broad community to evaluate the risks and benefits of security research. Increasingly we require fine-grained guidance in a form that could be evaluated and acted on. Researchers themselves need to be able to include in their publications an indication that they have made the effort to evaluate their work against the ethical concerns. Such internal and external evaluations need to be performed in a way that is uniform across all research situations and topics.

The approach adopted here is close to that of Johnson and Miller [5] in that we are concerned with building expertise in practical decision making. Bynum and Rogerson [2] suggest a multi-staged approach to case study analysis in order to build ethical judgement capabilities. These stages include: identifying key ethical principles, detailing the case study, identifying specific ethical issues raised by the case, calling on your own experience and skills for evaluation, then the abilities of others, and finally, applying a systematic analysis technique. In the following sections we briefly summarize an application of these approaches.

3.1 General Ethical Issues

When considering actions related to research or mitigation of malicious or illegal activity, there are many issues that must be considered. These involve issues of (a) proportionality, (b) targeting, (c) necessity, (d) desired outcome, (e) potential consequences, and (f) the greater moral good to society that is expected to result (and whether it outweighs any potential harm to innocent third parties.) For example, the kinds of questions that researchers should ask themselves include (but are not limited to):

- Are the research results intended to protect a specific population, and if so, which population?
- Is there a way to achieve multiple benefits to society simultaneously when studying criminal behavior?
- Who will benefit more from publication of research findings, and in which order: Victims of criminal acts; authorities responsible for protecting their citizens; the researchers themselves; or the criminals who are perpetrating computer crimes?
- Is there any other way to accomplish the desired research result(s)?
- What is the safest way to disseminate research results without risk of improper use by individuals who may not share the researchers' ethical standards?

3.2 Analyzing Case Studies Systematically

Following Bynum and Rogerson, we have identified a handful of issues that researchers can use to evaluate their research. However, to build consensus across a wide range of research and situations, it is advantageous to explore formal methodologies for evaluating these questions. Such methodologies allow for comparisons to be made across researchers and projects. In the following section we examine some potential methodologies.

3.2.1 Stakeholder Analysis

Stakeholder Analysis identifies the key players in the situation in terms of their interests, involvement, and their relationship (i.e., producer or recipient) of outcomes such as benefit or harm. We will adapt the definitions of stakeholders from `http://www.theasa.org/networks/apply/ethics/analysis/stakeholder.htm` for the purposes of this section.

- **Primary stakeholders** are, "those ultimately affected, either [positively or negatively]." These will typically be the end-users of computer systems, 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, these would be 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 will include the researcher(s), vendor(s), those who design and implement systems, and criminals or attackers.

3.2.2 Roles and Responsibilities Analysis

Roles and Responsibilities Analysis takes the identified Stakeholders, and lists both their role or roles in the situation, as well as their responsibilities towards each other and to society as a whole. Once stakeholders have been identified, and roles and responsibilities mapped out, one can start to define desired outcomes in terms of maximizing benefits and minimizing harms to stakeholders. Alternative actions that fall within the delineated roles and responsibilities can then be weighed against each other in terms of expected outcomes. One of the hardest challenges is in trying to identify potential negative outcomes that may result from alternative actions in order to minimize unintended consequences. This is where involvement of trusted external parties, such as peer-review of proposed actions or protocols, can help.

4. REVISITING KLOXO / HYPERVM

The study of kloxo is interesting and unique in terms of the possible relationship with the suicide of the CEO of the vendor and it brings in many issues of risk/benefit across many parties. In kloxo, we can identify the following stakeholders:

- **Key Stakeholders** The *researcher* who discovered the vulnerabilities. The *programmers* who were responsible for creating the HyperVM system and Kloxo administrative front end. The *corporate management of the vendor* (LxLabs), which includes the *CEO*. The *Criminals and Attackers* who would exploit vulnerabilities for their own purposes.
- **Secondary Stakeholders** The *service providers* who purchased HyperVM / Kloxo.
- **Primary Stakeholders** The *customers* of the service providers who use the virtual servers. The *consumers* who obtain products or services from the customers of the service providers (e.g., the online merchants using virtual storefronts hosted on HyperVM virutal machines.)

The researcher attempted to contact the corporate management of LxLabs, presumably to convince them to make decisions that would direct the programmers to fix the bugs that the researcher identified. Implicitly, we assume the researcher chose to contact the vendor privately to allow

them to fix the problem in order to protect the primary stakeholders (i.e., virtual machine customers and their end consumers.)

The action of the researcher as a key stakeholder to make detailed vulnerability and exploit information to the vendor is intended to assist the vendor in correcting the problems and eliminating the vulnerability. This creates a benefit to the primary stakeholders by protecting their services and accounts, as well as benefiting the secondary stakeholders by improving their product and protecting their customers.

It is the vendor's responsibility as a key stakeholder to use this information to minimize potential harm to the primary stakeholders. While the researcher did not state this explicitly, we can assume that the researcher has taken upon himself/herself the responsibility of assisting in protecting the primary and secondary stakeholders. We can infer that the action of reporting was intended to obtain the outcome of protecting the primary stakeholders by minimizing harm to them that would result from a malicious actor finding and exploiting these vulnerabilities before the vendor corrected them.

The researcher had several optional pathways that could achieve this same goal:

- The researcher could have taken a high-level outline of the vulnerabilities and provided them to a reporter, who could have written a news article disclosing (in general terms) that vulnerabilities in the HyperVM / Kloxo system were discovered and warning the primary stakeholders (i.e., customers and end consumers). The primary stakeholders could then take their own actions to ask questions, harden defenses, ensure they had current backups, or consider moving their storefronts to other service providers.
- The researcher could have identified a representative set of HyperVM / Kloxo customers and warn them (again, in general terms) of the vulnerabilities, and/or provided mitigation details. These customers could have been encouraged to contact LxLabs and put pressure on the vendor to fix the problems. This would also have the same added benefits in terms of minimization of harm as the previous option. This would not be as easy as contacting a single reporter, or reporting to a CERT organization, but would still move towards achieving the goal of protecting the customers and end consumers.
- The researcher could have published a high-level summary of the vulnerabilities, rather than full exploit details. This may well result in calls from full-disclosure advocates to provide more details, and possibly criticism of the researcher for over-stating the significance of their findings. Anyone with the same (or greater) skills would be able to repeat the research and thus possess the same ability to exploit these vulnerable systems, however this would take time that the vendor may be able to use to fix the problems before any harm is done to the primary stakeholders. The researcher thus has to balance personal benefit from first discovery and/or immediate full disclosure, potential harm resulting from criticism for partial disclosure, and potential harm to primary stakeholders from release of exploit information prior to patches being available to fix the bugs.

The use of anonymity by the researcher for unstated reasons leaves open many questions. (i) It may indicate that there is no personal gain to the researcher from disclosure. Then again, it also has the potential of avoiding accountability for any actions that are taken, including unintended consequences that cause harm. (ii) Releasing full details only two weeks after first contacting the vendor is another difficult issue. Because there was no evidence of the vendor even looking at the vulnerability details, could the researcher have been acting with a punitive motive against the vendor? (iii) The researcher may be a disgrunteled current/former employee with a retributive motive.

It is unreasonable for the researcher to anticipate the CEO would commit suicide, nor is it provable that the pressure from disclosure and resulting damage from exploitation of the vulnerabilities contributed to the suicide. It is foreseeable, however, that disclosure of full exploit details without warning would likely result in one or more parties using this information to do anything made possible, up to and including destroying the contents of any servers they could find. The exploit details alone do not help the primary stakeholders in protecting themselves, as there is nothing they can do (short of immediately switching to another virtual machine provider, which would take significant effort and time.) Since the researcher provided no mitigation details, the information that was released reasonably can be seen to benefit attackers more. Thus, if the primary goal of the researcher was to minimize harm to the primary stakeholders, the choice to disclose the vulnerabilities two weeks after first attempting to contact the vendor resulted in the exact opposite result (i.e., increased harm and decreased benefit to both primary and secondary stakeholders.)

5. CONCLUSION

More questions are typically raised about the ethics of computer security research activities than answers are provided. To help understand these issues and define a workable ethical framework, we believe that a more structured series of public discussions are urgently needed. For a glimpse at our efforts to help energize this needed dialog, please see our extended technical report [3].

6. REFERENCES

[1] M. Allman. What ought a program committee to do? In *WOWCS'08: Proceedings of the USENIX Workshop on Organizing Workshops, Conferences, and Symposia for Computer Systems*, pages 1–5, 2008.

[2] T. W. Bynum and S. Rogerson. *Computer Ethics and Professional Responsibility: Introductory Text and Readings.* Blackwell Publishers, Inc., Cambridge, MA, USA, 2003.

[3] D. Dittrich, M. D. Bailey, and S. Dietrich. Towards community standards for ethical behavior in computer security research. Technical Report 2009-01, Stevens Institute of Technology, Hoboken, NJ, USA, April 2009.

[4] S. L. Garfinkel. IRBs and security research: Myths, facts and mission creep. In *Proceedings of UPSEC '08 (Usability, Psychology and Security)*, Apr. 2008.

[5] D. G. Johnson and K. W. Miller, editors. *Computers Ethics.* Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 2009.

[6] J. Leyden. LxLabs boss found hanged after vuln wipes websites. `http://www.theregister.co.uk/2009/06/09/lxlabs_funder_death/`, June 2009.