Extended Abstract

Ethics in IT Security Research

Sebastian Neuner[1*] and Martin Mulazzani[2] and Sebastian Schrittwieser[3]

[1] SBA Research, Favoritenstraße 16, 1040 Vienna, Austria
[2] SBA Research, Favoritenstraße 16, 1040 Vienna, Austria
[3] Fachhochschule St. Pölten, Matthias Corvinus-Straße 15, 3100 St. Pölten, Austria

E-Mails: sneuner@sba-research.org, mmulazzani@sba-research.org, Sebastian.Schrittwieser@fhstp.ac.at

* Author to whom correspondence should be addressed; Tel.: + 43 (1) 505 36 88; Fax: + 43(1) 505 88 88

Accepted:

Introduction

Research in IT security often comes with decisions and possibilities that may or may not be considered ethical. However, it is often hard for young researchers to estimate the impact of their work, possible consequences and overall morality, as well as to where to draw the line. In some cases it is likely that more than hundreds of thousands of users will be affected, and it is unclear what is in their best interest: removal of a threat? Or rather a deeper analysis of the threat, which could prevent further vulnerabilities or attacks that are similar? In most cases, this decision is then left to the advisor who may have conflicting interests.

Methods

In this talk we will present recent borderline papers from the ethical point of view, and their implications on users. These papers are either directly or indirectly related to our own work, meaning that we will put our own perceptions on morality and ethics in perspective. We will furthermore present fundamental ethical principles which should not only be considered for IT security research, but can be applied to research in general. We argue that the establishment of ethical guidelines or frameworks without prior discussion and consensus in the research community probably would not lead to clarity on which lines in academic research should not be crossed. Especially the world-wide context of IT research

poses challenges which are not easy to overcome: while researchers at US institutions are often forced to go through an IRB approval, nothing comparable exists in the broader context of European research. On the other hand, Europe with its much stronger privacy laws has nothing comparable within the US or Asia.

A good example is the analysis of a botnet [1]: malicious software which is run on thousands of computers, operated by unsuspecting users. These computers are then used for sending spam e-mails, collecting banking information, and many more severe malicious activities. The researcher is now in the dilemma: shut down the botnet which is just one among many more, or conducting a deeper analysis of the botnet? Another option would be to sanitize the computer and fix the underlying vulnerabilities of the computer, to prevent similar and future infections. Or should the user be warned, and made aware of the fact that the system is running malware including instructions on how to get rid of it?

The interest of the user, in particular for experiments where the user is not informed (or asked for consensus), has to be the highest priority. Another priority should be in our opinion that watching users getting harmed is not acceptable. Watching how their personal information is stolen or their computers are being abused for sending spam e-mails is not acceptable. But where to draw the line? What is acceptable, and what isn't? And how can unethical research be punished, when the decision for publication is made by very few, namely the chair of a program committee or the editor of a journal?

Another good example is the Tor network [2], an anonymity network used by thousands of people on a daily basis to stay anonymous on the Internet. Due to the open design of the Tor network and the possibility that everyone can become part of the network and relay traffic for users, it can be tempting to modify traffic in the users best interest, e.g. by blocking malicious domains, inspecting and modifying file transfers or attacking the Tor network by trying to deanonymize its users. But how can deanonymization attacks be evaluated without assessing the impact on real users? In particular for complex systems, simulation is not always possible, and while there are simulation frameworks available for Tor, how can the researcher be sure that all important parameters of the simulation are correct and within expectation?

Conclusions

Neither researchers nor research subjects, e.g. malware authors, users or online services are currently in a position where they can refer to ethical principles that are accepted in the research community as well as on a larger scale. This talk is hopefully another step towards a guideline which is acceptable for all parties. However, much more research is needed in that direction, and the different interests of all stakeholders have to be balanced.

Acknowledgments

This talk extends prior work by the authors published in 2013 [3].

References and Notes

1. Stone-Gross, Brett, et al. “Your botnet is my botnet: analysis of a botnet takeover.” Proceedings of the 16th ACM conference on Computer and communications security. ACM, 2009.
2. Dingledine, Mathewson, et al. “Tor: The Second-generation Onion Router” Proceedings of the 13th Conference on USENIX Security Symposium, 2004
3. Schrittwieser, Sebastian, Martin Mulazzani, and Edgar Weippl. “Ethics in security research which lines should not be crossed?” Security and Privacy Workshops (SPW), 2013 IEEE. IEEE, 2013.