Ethical Implications of Security Vulnerability Research for Critical Infrastructure Protection

Livinus Obiora Nweke[1] and Stephen D. Wolthusen[1,2]

[1]Information Security and Communication Technology
Norwegian University of Science and Technology (NTNU)
Gjøvik, Norway
[2]School of Mathematics and Information Security
Royal Holloway, University of London
Egham, United Kingdom

Abstract. Security vulnerability research (SVR) involves searching for security flaws in a system. Such activity is likely to raise ethical concerns which need to be considered. For example, if a security researcher discovers a vulnerability in a critical infrastructure that can be exploited by an attacker; what is the right thing to do? Based on 'duty of care' principle and the fact that a public disclosure would force the critical infrastructure operator to quickly address the issue; going public with the discovery seems to be the right course of action. However, based on 'do not cause harm to others' principle, a public disclosure could badly affect the reputation of the critical infrastructure operator. Also, there is the possibility that the disclosed vulnerability could be exploited by an attacker before the operator is able to resolve the problem. The question would then be: is public disclosure still the right thing to do? This type of situation raises an ethical dilemma because a critical infrastructure is a system that is essential for the maintenance of vital societal functions and any attack against such an infrastructure would have a devastating effect. In this paper, we examine the ethical implications of SVR for critical infrastructure protection using the three normative ethical theories. First, we review the state-of-the-art of ethics in SVR. Then, we investigate how the three different normative ethical frameworks would respond to a hypothetical scenario relating to security vulnerability in a critical infrastructure in order to provide guidance for security researchers involved in SVR. Finally, we present a discussion on how a security researcher would make an ethical decision when confronted with an ethical dilemma. We observe from this study that a security researcher could rely on the three different normative ethical frameworks to reason about the best course of action during SVR for critical infrastructure protection.

Keywords: Security Vulnerability Research, Ethics, Ethical Implications, Critical Infrastructure Protection

15th International Conference on Wirtschaftsinformatik,
March 08-11, 2020, Potsdam, Germany

1 Introduction

In the battle of attackers and defenders as seen in the cyber space, defenders are always at a disadvantage considering that they are required to get it right every time. For attackers, they just need to get it right once and they have additional leverage to try as many times as possible. The situation is further exacerbated with the observation made by Rescorla in [17] that the possibility of a defender discovering a vulnerability before an attacker is very small. Thus, it has led to companies encouraging external security researchers to come forward about discovered vulnerabilities through vulnerability reward programs also called bug bounty programs.

Companies are increasingly recognizing the value of running vulnerability reward programs, where security vulnerability researchers are rewarded for reporting security flaws they discover in their systems. Huawei recently launched a bug bounty program that would pay up to $220,000 for demonstrating critical weakness in one of its Android devices [7]. In fact, a Gartner report has shown that approximately 50% of companies are likely to adopt vulnerability reward program by 2022 [9]. However, searching for security vulnerabilities still raises ethical concerns that need to be investigated. Although SVR has some legal ramifications, the current trends in the discussion of SVR has been centred on ethical issues as there have been very little-known legal actions against security researchers involved in SVR.

Ethics seeks to define what is the right or wrong course of action when confronted with an issue. When considering the legal aspects of an issue, we apply the principles laid down in the laws. Equally, when we consider the ethical aspects, we refer to principles enshrined in ethical and moral values. Ethical actions differ from legal actions; as there are legal actions that are not ethically acceptable [21]. Whilst practising law involves anticipating how and what a judge might decide when presented with an issue, ethics appears as a superior and stable reference to which laws can refer to [8].

In this paper, we consider the ethical implications of SVR using the three ethical philosophies, namely: deontological, consequentialist and virtue ethics. We first review the state-of-the-art of ethics in SVR. Using this understanding, we then investigate how the three different ethical frameworks would respond to SVR which involves a critical infrastructure. A critical infrastructure is a system or an asset that is essential for the maintenance of vital societal functions. Communication networks, electric power, water and wastewater, transportation, oil and gas infrastructure, healthcare, and satellite communications are defined as critical infrastructure. For instance, if a security researcher discovers a vulnerability in a critical infrastructure like the electric power station, that can be exploited by an attacker; what is the right thing to do? This is to provide guidance for security researchers involved in SVR for critical infrastructure protection. Finally, we present a discussion on how a security researcher would make an ethical decision relying on the right or wrong course of action as described by the different ethical frameworks.

The rest of this paper is organised as follows. Section 2 examines the three traditional normative ethical theories, the main contrasts between the three ethical theories and the process security vulnerability researchers could rely upon to make an ethical decision. Section 3 presents a literature review of ethics in SVR starting with a review of the efforts by several organizations in developing an ethical framework and a review of other works related to ethics in SVR. Section 4 presents a hypothetical ethical dilemma and what different ethical frameworks described would agree to be the right or wrong course of action. Section 5 discusses how a security researcher would make an ethical decision relying on the likely right or wrong course of action as presented by the different ethical frameworks. Section 6 concludes the paper and presents future work.

2 Background

There are three main approaches to ethics and they include: meta-ethics, which deals with the nature of moral judgements; normative ethics, where the interest is in the content of moral judgements and the yardstick for what is right or wrong; and applied ethics, that considers controversial topics like voluntary euthanasia, animal rights and capital punishment [2]. We are concerned with normative ethics in this paper.

Normative ethics not only provide us with a way of evaluating moral judgements; it also facilitates the formulation of a decision procedure to direct moral action [5].This is in contrast with other main approaches to ethics where for instance, meta-ethics is concerned with the origin and meaning of ethical concepts and applied ethics only concentrates on the analysis of specific, controversial moral issues [6]. Both approaches do not offer suggestions that may guide security researchers in the decision-making process when confronted with an ethical dilemma. Hence, we utilize the likely decision procedure that can guide moral action, which normative ethics offers to reason about the ethical implications of SVR for critical infrastructure protection and to present a discussion on how a security researcher would make an ethical decision when confronted with an ethical dilemma. The traditional normative ethical theories are deontological, consequentialist and virtue ethics.

Deontological ethics is also known as duty ethics because it focuses on the duties and obligations that we have in any given situation [2]. It is a type of normative ethics that guide and assess our choices of what we ought to do [1]. An individual is required to always perform their duty regardless of the outcome. The right conduct is given as doing one’s duties and doing the right thing despite the consequences. Deontological ethics provides a system of rules with consistent expectations for all people. This implies that if an action is ethically right, it would apply to every person in that given circumstance.

In consequentialist ethics, the focus is only on the consequences of an action [2]. It is usually seen as the opposite of deontological ethics because most deontological theories are best understood in contrast to consequentialist ones [1]. The right or wrong course of action in any situation is one that produces the best consequences. For example, there is a type of consequentialist theory which suggests that the right course of action is one that benefits the greatest number of people while an action that benefits the least number of people is said to be the wrong course of action. Consequentialist ethics is touted to be a more pragmatic approach because of its focus on the results of an action. However, it is very difficult to predict the consequences of an action due to insufficient information and as such, the use of consequentialist ethics may produce undesirable outcomes.

Virtue ethics views ethics in terms of the character traits (either positive or negative) that might motivate one in a given situation [2]. It is usually identified as an ethics that focuses on virtues or moral character, different from the approach that focuses on duties or rules (deontology) or that focuses on the consequences of actions (consequentialism) [12]. Here, the concern is what kind of person we should be and what our actions indicate about our character. The right or wrong course of action in any situation is whatever a virtuous person would do in that situation. The implication of this is that several types of behaviour may be called ethical, considering there might be many different types of good characters and many ways of developing them.

Generally, when confronted with an ethical dilemma during SVR, security researchers may rely on the following process to make an ethical decision [2]:

- Recognize an ethical issue;
- consider the parties involved;
- gather all the relevant information;
- formulate actions and consider alternatives;
- make a decision and consider it;
- act; and
- reflect on the outcome.

The above steps are based on the three broad frameworks to guide ethical decision making that were derived from the normative ethical theories, namely: the consequentialist framework, which focus on the effects of the possible courses of action taking into account those that will be affected directly and indirectly; the duty framework, where the focus is on the duties and obligations that we have in a given situation; and the virtue framework, where our goal is to identify the character traits that might motivate us in a given scenario [2]. These three frameworks are important in the decision-making process because an understanding of the pros and cons of each framework would enable a security researcher to decide the most appropriate framework to adopt considering the given situation that is presented.

3 Literature Review

Several efforts have been made in the past to develop an ethical framework for security researchers. In this section, we present a literature review of the ethics in SVR starting with the existing ethical frameworks and the need for an ethical framework for SVR. Also, we review other works related to ethics in SVR to provide justification for our choice of case study analysis in examining the right or wrong course of action when confronted with an ethical dilemma in SVR.

3.1 Existing Ethical Frameworks

There is currently no single ethical framework that guides the conduct of security researchers. However, organizations such as Association for Computing Machinery (ACM), Institute of Electrical and Electronics Engineers (IEEE), and Information

Systems Security Association, Inc. (ISSA) have codes of ethics that guides the behaviour of its members in the discharge of their professional duties. We provide a review of these codes of ethics in this subsection to give a general overview of existing ethical frameworks and the need for an ethical framework for SVR.

ACM updated its code of ethics and professional conduct referred to as “the Code” in 2018. As observed in [11], the revised Code of Ethics aims to address the significant progresses made in computing technology since the 1992 version, in addition to the growing reliance on computing in every aspect of human endeavour. The general objective of the code is to inspire and guide the ethical conduct of all computing professionals and anyone who uses computing technology in an impactful way [11]. In a similar way, IEEE Code of Ethics seeks to commit its members to the highest ethical and professional conduct [13]. Furthermore, a joint work by ACM and IEEE published the Software Engineering Code of Ethics and Professional Practice [10]. It commits software engineers in both bodies to making the analysis, specification, design, development, testing and maintenance of software a beneficial and respected profession [10]. Lastly, ISSA aims to promote practices that will ensure the confidentiality, integrity, and availability of organizational information resources [14]. It achieves this by requiring that its members reflect the highest standards of ethical conduct.

It is easy to observe that the ethical frameworks presented in the preceding paragraphs have some drawbacks in that they fail to offer a clear decision-making process when confronted with an ethical issue. The authors in [4] opine that the three drawbacks include: absence of shared community values, lack of consensus on enforcement, and limited individual expertise. They argue that none of the existing ethical standards envisioned the multi-facet nature of the narrow field of security research. Hence, it is important to review other works related to ethics in SVR to serve as a basis for selecting the right ethical framework for making ethical decision in SVR and further advancing the field of ethics in SVR.

3.2 Ethics in Security Vulnerability Research

In the absence of a clear ethical decision-making process for security researchers when confronted with an ethical dilemma, we consider the different approaches adopted by other works related to ethics in SVR in this subsection. The review was conducted in IEEE Explore, Google Scholar, Science Direct and Springer with keyword phrase; “SVR”. Related literatures were assessed in relation to their suitability for evaluating ethics in SVR. The goal was to provide us with an overview of the state-of-the-art of ethics in SVR and suggestions on the best approach to employ given the consensus among the related works.

One of the earliest and most influential works related to ethics in SVR is presented in [15]. The authors first distinguish between a hacker, which they referred to as a person breaking into computer systems or committing into other such activities and information security personnel, which they term as those responsible for protecting systems. They observe that information security personnel tend to categorically state what is an ethical behaviour and what is not (deontological ethics). In contrast, they suggest that a typical reasoning among hackers (which in our case are security researchers) is that their activities provide good outcome for the greatest number of people (consequentialist ethics). These two approaches as they noted are in strong conflicts, which raises the need for a unified ethical framework.

Carle in [3] explores different ethical frameworks and then applied them to situations where security researchers would have to make an ethical decision. The goal of this work is to provide a common ethical framework that security researchers could rely upon in deciding the right or wrong course of action. The paper presented case studies of different scenarios that a security researcher may be confronted with. Several ethical frameworks described in the work are then employed to suggest a common ethical framework that would hold security researchers to an even higher ethical standard.

Some authors have argued that SVR is ethical. The author in [19] reasons that the question should not be if it is ethical to do SVR but rather, if someone possess the skill to analyse and provide better insights into the problem; whether it is ethical for the person not to do SVR. In the case of authors in [16], they maintain that security vulnerability researchers perform an essential social function as they provide information gap between the creators, or exploiters of vulnerable systems and the third parties who will likely be harmed because of them. Also, they use case study analysis to recommend best practices in SVR.

An ethical guideline for security researchers is presented in [18]. The paper discusses some of the ethical dilemmas security researchers have been confronted with in recent years. Specifically, the work employs case studies analysis of ethical failings as a means of demonstrating the problems that could arise when the right or wrong course of action is not perfectly clear. In the same way, authors in [4] use case studies analysis to recommend a process of building personal ethical decision-making abilities. They argue that analysing case studies provide an excellent approach for building practical ethical decision-making abilities.

Schrittwieser et al. in [20] echo the need for a discussion among practitioners before the development of ethical guidelines or frameworks. This is to ensure clarity on which lines academic security researchers should not cross in the course of their researcher activities. The work is mainly concerned with how to ensure that research activities in the field of information security do not harm others and to facilitate how those research activities can be evaluated from an ethical point of view. Using case studies of controversially discussed research, the authors discuss fundamental ethical principles and compare them in the paper with their justification regarding research ethics. The analysis shows how difficult it is to stipulate generally accepted and universally valid ethical standards.

As can be noted, case study analysis is the recurrent approach that is deployed in most of the works related to ethics in SVR as presented above. It involves the use of hypothetical scenarios, or real-life examples of situations that security vulnerability researchers may face. We employ a similar approach in the next section to consider what different ethical frameworks we have presented so far would argue to be the right or wrong course of action for a hypothetical scenario. Finally, we present in Section 5, an ethical decision-making process that security vulnerability researchers may rely upon in deciding the right or wrong course of action when confronted with an ethical dilemma.

4 Ethical Implications of Security Vulnerability Research

SVR is likely to present a security researcher with an ethical dilemma that requires the researcher to reason about the right or wrong course of action. For example, if a security researcher discovers a vulnerability in a critical infrastructure that can be exploited by an attacker: what is the right thing to do? It may seem apparent that the right thing to do is to inform the critical infrastructure operators to fix the flaw. However, there have been situations in which a security researcher informs the critical infrastructure operators about a vulnerability and still nothing is done to address the issue (which also present potential risk to the general public). In addition, disclosing the vulnerability could damage the reputation of the critical infrastructure operator and may expose the critical infrastructure to the danger of being exploited by an attacker before the issue has been addressed. This type of situation presents an ethical dilemma for security researchers and they would have to decide the right thing to do.

There are several other questions security researchers would have to consider when deciding the right course of action in the scenario presented in the previous paragraph. First, should they make the discovery public? This is based on 'duty of care' principle and the fact that public disclosure would force the critical infrastructure operator to quickly address the issue. Next, going public with the disclosure could badly affect the reputation of the critical infrastructure operator and there is also the possibility of an attacker exploiting the vulnerability before the issue has been addressed. Based on 'do not cause harm to others' principle, is going public with the disclosure the right thing to do for security researchers?

These and many other ethical dilemmas are faced by security researchers, which they would have to decide the right course of action. By considering the likely course of action of different ethical philosophies using the frameworks presented in Section 2, we aim to provide security researchers with different approaches which they could rely on to reason about the right course of action when confronted with an ethical dilemma. Thus, using the ethical dilemma faced by a security researcher as presented in the introductory paragraph of this section, let us consider what the different ethical decision-making frameworks (discussed in Section 2) would agree to be the right or wrong course of action.

The duty framework would argue that the security researcher should never had considered searching for vulnerabilities without the consent of the critical infrastructure operator in the first place. This is because it violates the duty of the security researcher towards the critical infrastructure operator. As of now, there are very few companies that allow external security researchers to search for flaws in their systems. Although the situation is likely to change in the future, it follows that in the scenario we depicted, the security researcher may had searched and discovered the security flaws without the consent of the critical infrastructure operator. Hence, deontological ethics would consider the action of the security researcher to be wrong as it violates the rights of the critical infrastructure operator.

Unlike the duty framework, consequentialist framework would rather consider the likely outcome of disclosing the vulnerability in relation to the number of people that could benefit from such disclosure. The consequentialist framework would argue that the public who are at risk as a result of the vulnerable critical infrastructure represents the greatest number of people against the reputation of the critical infrastructure operator. Therefore, the right course of action for consequentialist framework would be to disclose the vulnerability to the public in order to force the critical infrastructure operator to fix the vulnerability in the critical infrastructure.

For the virtue framework, there are no strict or simple rules (as found in deontological and consequentialist frameworks) that may be employed to ascertain the right course of action. However, we can apply Aristotle's doctrine of the mean to decide the likely course of action in the scenario we depicted. Fieser in [6] suggests that Aristotle's doctrine of the mean views most virtues to be the mean between more extreme character traits. Both the consequentialist and the deontological frameworks presented above could be seen as exemplification of two extreme character traits. We could label one-character trait as extreme pragmatism" and the other as "being dismissive of consequences". We could imagine a character trait that is concerned with both the consequences of his actions and omissions and with non-consequentialist aspect of his behavior as a mean. How such a person would be disposed to act should be the kind of virtue we are trying to imagine. In the case in question, the critical infrastructure operator should first be informed about the vulnerability and should be given certain period to fix the security flaw after which the vulnerability should be disclosed to the public.

5 Discussion

Having considered the likely course of action for the different normative ethical frameworks we have studied in this paper, we present an ethical decision-making process that security vulnerability researchers may rely upon in deciding the right or wrong course of action when confronted with an ethical dilemma using the decision-making process, we described in Section 2 that is based on the three broad frameworks to guide ethical decision making, namely: consequentialist framework; the duty framework; and the virtue framework.

When applying the frameworks to make an ethical judgement regarding a specific situation, the first step is for the security researcher to recognize that the situation is an ethical issue. For example, in the scenario we depicted in the previous section, the security researcher would have to recognize that disclosing or not disclosing the discovered vulnerability in a critical infrastructure is an ethical issue before the ethical decision-making process can proceed. This is an important step as without recognizing the concerns an issue we are faced with raises, it is practically impossible to follow through the remaining steps of the ethical decision-making process. Hence, the security researcher would have to acknowledge that the situation presented raises ethical concerns before the next step of the process can be considered.

With the recognition of the ethical issue that disclosing or not disclosing the discovered vulnerability in a critical infrastructure poses, the security researcher would have to consider the parties involved. The relevant stakeholders in this case are the critical infrastructure operator, the public that uses the critical infrastructure, and the security researcher. Other stakeholders may include, the organization the security researcher represents, the professional bodies the security researcher belong to, and so on. After the relevant parties involved have been identified, the next step is to gather all the relevant information.

Unfortunately, it is impossible to gather all the required information for any given situation. However, the security researcher faced with our hypothetical scenario would have to make effort to obtain as much relevant information as possible. For example, the amount of time it would take to fix the vulnerability, the number of people that would be affected in the case of a successful exploit, the potential damage to the users, the likelihood that an attacker would exploit the vulnerability, whether or not the vulnerability has already been exploited, etc., are some of the relevant information the security researcher may have to consider. This is a very important activity as it is a prerequisite for the next step in the ethical decision-making process. Without enough information, it would be difficult to formulate actions and alternatives.

In formulating actions and alternatives, the security researcher may rely on the likely course of actions for the different ethical frameworks that were described in the previous section. Those likely course of actions should form the basis of the possible actions the security researcher would have to take and the available alternatives. It follows that for our hypothetical scenario, the security researcher would consider each of the likely course of actions for the three ethical frameworks and decide which best addresses the given situation. Then, the security researcher would have to examine how they feel about their decision. Finally, the security researcher would have to act and reflect on the outcome of their decision.

6 Conclusion

In this paper, we have considered the ethical implications of SVR using the three normative ethical philosophies. We started by describing the three traditional normative ethical theories namely: deontological ethics, consequential ethics and virtue ethics. Then, we presented a state-of-the-art review of ethics in SVR. We have employed case study analysis to examine the likely course of actions of the different normative ethical frameworks we studied using a hypothetical scenario. Lastly, a discussion on how a security research would make an ethical decision is presented.

The field of information security is constantly evolving, and we believe that active discussion on the different ethical dilemmas' security researchers may face is needed in the future. For example, a situation where a security researcher employed by a company and have signed a non-disclosure agreement, discovers a vulnerability in a system that may result in death. What would be the right course of action? Current discussions have been centred on external researchers but there have been cases where an internal researchers' duty to the employer conflicts with the researchers' duty to the public. This type of discussions will ensure that security researchers are aware of the implications of their actions and will foster more responsible behaviour in the cyber space.

References

1. Alexander, L., Moore, M.: Deontological Ethics. *The Stanford Encyclopedia of Philosophy*, https://plato.stanford.edu/archives/win2016/entries/ethics-deontological (2016)
2. Bonde, S., Firenze P. A framework for making ethical decisions, (2013), https://www.brown.edu/academics/scienceand-technology-studies/framework-making-ethical-decisions
3. Carle, S.: Crossing the line: Ethics for the security professional (2003), https://www.sans.org/reading-room/whitepapers/hackers/crossing-line-ethicssecurity-professional-890
4. Dittrich, D., Bailey, M., Dietrich, S.: Building an active computer security ethics community. IEEE Security & Privacy 9(4), 32–40 (2010)
5. Driver J.: Normative Ethics. *The Oxford Handbook of Contemporary Philosophy* (Sep 2009), https://doi.org/10.1093/oxfordhb/9780199234769.003.0002
6. Fieser, J.: “Ethics”, *The Internet Encyclopedia of Philosophy*, ISSN 2161-0002, https://www.iep.utm.edu/ethics/, (2020).
7. Forbes: (Nov 2018), https://www.forbes.com/sites/thomasbrewster/2019/11/18/huaweibeats-google-in-offering-220000-to-hackers-who-find-android-backdoors/
8. Fuster, G.G., Gutwirth, S.: Ethics, law and privacy: Disentangling law from ethics in privacy discourse. In: Proc. Technology and Engineering 2014 IEEE Int. Symp. Ethics in Science. pp. 1–6 (May 2014). https://doi.org/10.1109/ETHICS.2014.6893376
9. Gartner: Emerging technology analysis: Bug bounties and crowdsourced security testing (2018)
10. Gotterbarn, D., Miller, K., Rogerson, S.: Software engineering code of ethics. Communications of the ACM 40(11), 110–118 (1997)
11. Gotterbarn, D., Brinkman, B., Flick, C., Kirkpatrick, M.S., Miller, K., Vazansky, K., Wolf, M.J.: Acm code of ethics and professional conduct. ACM (2018)
12. Hursthouse, R., Pettigrove, G.: Virtue Ethics. *The Stanford Encyclopedia of Philosophy*, https://plato.stanford.edu/archives/win2018/entries/ethics-deontological (2018)
13. IEEE: Ieee code of ethics (Feb 2014), https://www.ieee.org/about/corporate/governance/p7-8.html
14. ISSA: Issa code of ethics, https://www.members.issa.org/page/CodeofEthics (2020)
15. Leiwo, J., Heikkuri, S.: An analysis of ethics as foundation of information security in distributed systems. In: Proceedings of the Thirty-First Hawaii International Conference on System Sciences. vol. 6, pp. 213–222. IEEE (1998)
16. Matwyshyn, A.M., Cui, A., Keromytis, A.D., Stolfo, S.J.: Ethics in security vulnerability research. IEEE Security & Privacy 8(2), 67–72 (2010)
17. Rescorla, E.: Is finding security holes a good idea? IEEE Security Privacy 3(1), 14–19 (Jan 2005). https://doi.org/10.1109/MSP.2005.17
18. Sassaman, L.: Ethical guidelines for computer security researchers: “be reasonable”. Lecture Notes in Computer Science 6054, 250 (2010)
19. Schneier, B.: The ethics of vulnerability research (2008), https://www.schneier.com/blog/archives/2008/05/the-ethics-of-v.html
20. Schrittwieser, S., Mulazzani, M., Weippl, E.: Ethics in security research which lines should not be crossed? In: 2013 IEEE Security and Privacy Workshops. pp. 1–4. IEEE (2013)
21. Shou, D.: Ethical considerations of sharing data for cybersecurity research. In: International Conference on Financial Cryptography and Data Security. pp. 169– 177. Springer (2011)