Analysis of Target Corporation's 2013 Data Breach Via a Deontological Framework

STS Research Paper
Presented to the Faculty of the
School of Engineering and Applied Science
University of Virginia

By

Yonathan Fisseha

March 1, 2020

On my honor as a University student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments.

Signed: _____

Approved: _____ Date _____
Benjamin J. Laugelli, Assistant Professor, Department of Engineering and Society

1 Introduction

On December 19th of 2013, Target Corporation announced that it has suffered from a cyber attack between November 27th and December 15th of the same year, exposing approximately 40 million credit card and debit card information ("Target Corp.", 2013). Almost a month later, after a lengthy forensic investigation in collaboration with law enforcement agencies, Target published an update stating that the attack impacted up to 70 millions customers and exposed personal data of customers including names, addresses, phone numbers, and emails. Target experienced a 2%-6% sales decline and paid 57.9 million USD in settlements with financial institutions and US States ("Target Corp.", 2014). The Target data breach is in the top 10 data breaches in history in terms of the number of customers impacted by the breach (Rivero, 2018).

The existing discussions on the case focus on the technical shortcomings that enabled the attack, and generally avoid discussing the ethical violations. A rigorous normative discussion of the ethical violations involved using a deontological ethical framework can improve our understanding of corporations' ethical responsibility in parallel with their well understood legal and technical responsibilities. Such an analysis can highlight gaps between the ethical responsibilities we expect corporations to uphold and the current subpar status quo.

Specifically, I will argue that although Target met the minimum technical industry standards, Target is still morally liable. Towards this goal, I will first argue that while it is challenging to assign blame to individual employees, Target, as a collective, can be held morally accountable. Second, I will analyze the case via a deontological framework against a set of prima facie duties for computing professionals: 1) design and implement systems that are robustly and usably secure, 2) maintain high standards of professional competence, conduct, and ethical practice, and, 3) articulate and apply organizational policies that reflect the principles of the Code. I will demonstrate that the 2013 data breach renders Target morally irresponsible via these two arguments.

2 Background

The Target data breach attack was sophisticated and carefully executed. The attackers first infiltrated Fazio Mechanical's network using a phishing attack vector and installed a Trojan on Fazio's system. The Trojan stole Fazio's credentials to Target's external billing system, and acted as a point-of-entry to Target's network. Once in Target's network, the attackers successfully navigated from the business side of the network to more sensitive portions by exploiting vulnerable and out-of-date systems. The attackers then installed a customized point-of-sale malware that stole credit and debit card information from cash registers. The stolen data was then encrypted by the same compromised hosts in the network and exported out to drop sites in Russia and Brazil (Shu, Tian, Ciambrone, & Yao, 2017; US Senate, 2014). Figure 1 summarizes the attack paths and the involved parties.

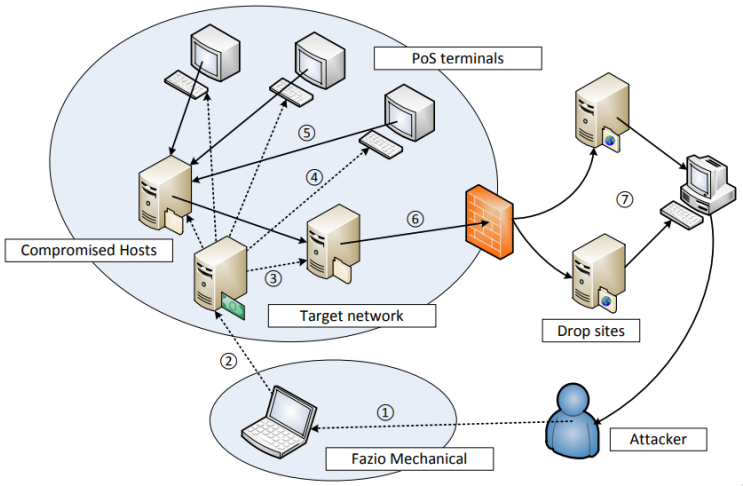

Figure 1: An overview of the data breach attack on Target. Attackers infiltrated Fazio Mechanical's network to enter Target's network. Then used vulnerable servers to install malware on point-of-sale terminals and extract debit and credit card information. Adapted from Shu et al. (2017).

3 Literature Review

The existing literature on the Target data breach is centered around the technical flaws that were exploited by the attacker, and how organizations can better protect themselves from

such attacks in the future. Shu et al. outline how the attack was executed and list four ways Target could have detected or prevented the attack, including appropriately responding to the cybersecurity system's automated warnings, properly segmenting the business network from the sensitive user data stores, hardening point-of-sale terminal security, applying proper access control on third-party partners (2017). A similar technical report by the US Senate's Committee on Commerce carefully analyzes the techniques used in the attack, and catalogs similar key points where Target failed to detect or prevent the attack."Target gave network access to a third-party vendor...which did not appear to follow broadly accepted information security practices" then once the attacker was in the network, "Target appears to have failed to respond to multiple automated warnings." Moreover, the attackers' ability to successfully navigate the network after the intrusion suggests that, "Target failed to properly isolate its most sensitive network assets." Finally, Target failed to respond to the anti-intrusion software's warnings regarding the attackers' escape route (US Senate, 2014). These technical discussions, while worthwhile in their own right, fail to consider the ethical dimension of the case and at most identifying the points of failure. They do not give neither a descriptive nor normative evaluation of the ethical standing of Target.

More broadly, there is a diverse literature on the capabilities of organizations to act as moral agents. Buttrick, Davidson, and McGowan emphasize the lack of literature on a "business' moral responsibility for data breach" and instead perform their moral analysis based on the ethics of marketing with regards to trust and responsibility (2016). They recognize three positions one can take in terms of businesses' responsibility to their customers: the contractual view, the due care theory, and social costs view. Since consumers lack the knowledge that the producer has, "the due care position recognizes the imbalance and the vulnerable position of the consumer by placing additional duties on the business" (Buttrick et al., 2016). Culnan and Williams develop the ethics of data breaches based the notion of vulnerability as well, "vulnerability explains many of our widely held moral intuitions...it exists because the disadvantaged party suffers a deficit of information and control" (2009). Customers or

stakeholders become vulnerable when they share their private data with businesses and, consequently, businesses receive the additional duties of protecting the customer. In the case of a data breach, the business causes harm to the customer by exposing the customer's private data which can be used for fraud and identity theft, and also cause the customer anxiety of future harm (Solove & Citron, 2017).

While the literature on business' moral capabilities is a necessary foundation for the analysis of this case via a deontological framework, it does not pass a normative judgment for the specific case either. Rather, it asserts that businesses have moral responsibility to their customers. I will analyze the Target 2013 data breach specifically against a set of duties accepted by the computing community to demonstrate that Target was morally irresponsible in the 2013 data breach.

4 Conceptual Framework

The analysis of Target's moral status with regards to the 2013 data breach draws on two frameworks: collective responsibility and deontological ethics. Collective responsibility enables the analysis to take a wider scope and provide a normative judgment on Target as a collective because it does not seem possible to pass judgment on specific employees due to the limited public information on the case. The analysis draws on deontological ethics because it provides a structured and rigorous framework to analyze a non-human agent's, i.e. Target's, ethics by appealing to duties that are already well established.

Collective responsibility is defined as, "the responsibility of a collective of people" to capture the intuition that there is "more to responsibility in complex cases than just the sum of the responsibility of the individuals considered in isolation" (Poel & Royakkers, 2011). It is likely that some of the individuals can be held accountable to some extent, but the distribution of moral fault among the individuals might be hard to determine. In such situations, it is challenging to attribute moral fault to individuals, but the collective can still be held morally accountable. This phenomena is called the *problem of many hands* per Poel

and Royakkers.

Deontological ethics is a normative ethical theory based on rules and principles that guide actions (Donaldson & Werhane, 2002). There are two major approaches to deontological ethics: Kantian deontology and social contract theory. Kantian deontology focuses on the individual's decision making based on duties and universal rules while social contract focuses on "general social principles that rational persons in certain ideal situations would agree upon and adopt" (2002). The analysis will use the social contract approach since it relates more closely to the case via the concept of collective responsibility– a social contract between Target and its customers. Therefore, a deontological analysis requires a predefined set of moral duties against which a given action can be evaluated. What these principles or duties should be is a topic of much discussion in theoretical and applied ethics. For this case study, however, I have selected the Association for Computing Machinery's (ACM) Code of Ethics. ACM Code of Ethics is designed to "guide the ethical conduct of all computing professionals...and anyone who uses computing technology in an impactful way"(Anderson, 1992). The Code is cited in courts, popular news, and taught both at the high school and college level to computing students (Brinkman & Carter, 2017). The Code has a broad scope so it can cover most practical ethical situations computing professionals might face, but the following are most relevant:

1. Design and implement systems that are robustly and usably secure
2. Maintain high standards of professional competence, conduct, and ethical practice
3. Articulate, apply, and support policies and processes that reflect the principles of the Code

Note that from the perspective of both deontological frameworks above, a failure to act in accord with a given principle is sufficient ground to say an agent acted immorally. This follows from the core deontological theory that the ethics of an action is determined solely based on its relation to a given set of prima facie duties. Consequently, it is enough for Target to violate any one of the principle in ACM's Code of Ethics to be considered immoral. The

analysis in Section 5 demonstrates how Target violated the three duties above, leading up to the 2013 data breach.

5 Analysis

As highlighted in Section 1 and 4, before analyzing the case using ACM's Code of Ethics, I will advance the argument that Target, as a collective, can be held morally accountable although individual employees cannot be. That is, I will argue that Target has a moral responsibility as a collective because the data breach occurred because of multiple failures in the organization as whole.

First, Target's leadership failed to institute and enforce security policies for their third-party vendors. Second, Target's security team failed to respond to multiple warnings from their security systems. Third, Target's network architecture was poorly designed by the network engineers. The data breach attack could have been prevented if any one of these three were not true. Assigning all the moral fault to any one of these large groups is not appropriate, although each one of them contributed to the data breach in some manner. However, it is very clear that collectively, as Target, they are morally responsible for the data breach. The lack of detailed public information on the data breach exacerbates the challenge in distributing moral responsibility. For example, it is possible that a security engineer reported the automated warnings or even investigate them, but failed to identify the ongoing attack– a lack of competency on the employee's part. Alternatively, the employee could have investigated and reported the attack successfully but an overtly bureaucratic system could have ignored the report– an organizational communication failure. It is impractical to attempt to narrow down the scope of the analysis to the level of an individual or even a group without much more specific information on the breach. The following analysis thus applies ACM's Code of Ethics to Target as a whole and demonstrates that Target was morally irresponsible.

5.1 Poor System Design and Implementation

The goal is to protect network resources by restricting communications which, in turn, has several beneficial security effects including: (i) reducing the number of entry points into a network, (ii) limiting the network access of an attacker who has penetrated the network, (iii) hindering the attacker's ability to move to other network devices and (iv) increasing the defender's ability to detect and remediate cyber intrusions (NSA IAD 2013). Segmentation is typically implemented by firewalls, network egress and ingress filters, application-

Target violated the first principle from ACM's Code of Ethics: design and implement systems that are robustly and usably secure. Target's system was not designed and implemented to be robustly and usably secure. More specifically, the network architecture was poorly designed and implemented from a security perspective. It requires more technical background in network segmentation to properly analyze how the poor network architecture enabled the data breach. "Network segmentation is concerned with partitioning a network into segments and controlling communications between segments and between segments and the Internet (where we assume cyber attackers preside)" (Wagner, Şahin, Pena, Riordan, & Neumayer, 2017). Each network segment can then strictly control the incoming and outgoing traffic based on custom business rules. For example, a business can have a `general-internal` segment and a `sensitive-internal` segment. The `general-internal` network can be shared by all systems that process insensitive data and might allow access from all employees. On the other hand, the `sensitive-internal` segment can be shared by systems that process sensitive user data, and consequently deploy a much more strict configuration on employees who can access that network. The network segment could deploy a whitelist configuration, a restricted list of entities who *can* access the network segment; in contrast, a blacklist configuration is a list of entities that *cannot* access the system which is often used by less sensitive systems. Traffic between network segments is then carefully monitored and filtered (Wagner et al., 2017).

Network segmentation is widely used by businesses to protect sensitive data, as it adds

a layer of protection between the attackers and the sensitive data since the attack path is typically in escalating direction of network security as was the case with the Target attack (Wagner et al., 2017). In fact, network segmentation is a requirement for retail businesses by the Payment Card Industry Data Security Standard, which Target successfully passed on their last inspection in 2013 (Gikas, 2010; US Senate, 2014). The Target network was well secured from outside attacks, but security within the internal network was very poor. The lack of internal network security is apparent from the path of the attack, as it allowed the attackers to navigate from less secure network areas to more secure areas (Shu et al., 2017). However, it is difficult to precisely identify the specific flaw due to the lack of public data on the breach. The fact that Target passed the security inspection but still suffered from data breach might seem contradictory at face value. However, the PCI standard provides only a minimum standard and thus leaves room for vulnerabilities. There are at least two possibilities for how the attackers were able to navigate across the network: erroneous configurations or other network security related vulnerabilities. In the case of an erroneous network configuration or design, it is a clear sign of insecure implementation and thus a violation of the first principle of ACM's Code of Ethics.

Another alternative worth considering in reconciling the seemingly contradictory reality is that at least a part of the network infrastructure implementation had a vulnerability that was exploited by the attackers to defy the network segmentation. Unconfirmed sources reported by Plachkinova and Maurer indicate that an audit team, after the attack, followed a similar attack path by starting from the business portion of the network and gained access to the exploited segment of the network without any authentication. The audit team achieved this by exploiting out-of-date software hosted on Target's servers (2019). If this was indeed how the attackers managed to navigate through the internal network, Target's lack of proper software and system maintenance, which is a significant portion of software implementation, indicates that Target was morally irresponsible. ACM's Code of Ethics reinforces this principle, "as threats can arise and change after a system is deployed, computing professionals

should integrate mitigation techniques and policies, such as monitoring, patching, and vulnerability reporting," which Target failed to do (Anderson, 1992). It is worth noting that under deontological ethics, the analysis does not need to prove that a given fault led to the data breach– the burden of proof is only in showing that Target was in violation of at least one of the principles. Even if the attackers did not use these techniques to traverse through the network, there is sufficient evidence that Target still designed and implemented an insecure system and thus should be held morally accountable.

5.2 Lack of High Standards of Professional Competence

Target violated the second principle from ACM's Code of Ethics: maintain high standards of professional competence, conduct, and ethical practice. Target's violation of this principle comes from their lack of response to early security warnings of the attack. Target invested in a, "well-known and reputable intrusion and malware detection service named FireEye, which was guided by the CIA during its early development" just six months prior to the data breach (Shu et al., 2017). FireEye detected malicious activity in the Target system as early as November 30th, 18 days after the attackers breached the network (US Senate, 2014). Additionally, according to the same report from the US Senate, "Target's Symantec antivirus software also detected malicious behavior around November 28, implicating the same server flagged by FireEye's software" (2014). Both of these warning notifications were received by Target's round-the-clock security team in Bangalore, India which reported the warnings to Target headquarters in Minneapolis, Minnesota (Plachkinova & Maurer, 2019). At this point in the attack path, the attackers had tested and installed the point-of-sale data extraction malware. Three days later, the attackers installed the data exfiltration malware preparing the escape route, which triggered more warnings from FireEye. Both of these warnings were ignored; moreover, "some prevention functionalities were turned off by the administrators who were not familiar with the FireEye system" (Shu et al., 2017). This is a sign of incompetence in the proper use of the new FireEye system. Target only began investigating

on December 12, when the U.S. Justice Department warned them about suspicious activity involving payment cards (Finkle, 2014).

Consequently, Target missed multiple opportunities to terminate the attack and avoid the data breach. Target was in clear violation of the principle of competence in ACM's Code of Ethics. The Code elaborates, "professional competence starts with technical knowledge....and requires skill in communication, in reflective analysis, and in recognizing and navigating ethical challenges" (Anderson, 1992). Target's violation is two folds: first, it failed to properly communicate the warnings in a timely manner and competently take an action to stop the attack; second, it performed work in an area where it was not competent resulting in the misconfiguration of FireEye. The Code explicitly requires one to perform work "only in areas of competence" and if there is a lack of expertise one must, "disclose this to the employer or client" (1992).

I have argued that Target's failure to address the automated warnings and their decision to disable the malware removal feature of FireEye demonstrates a lack of professional competence and thus makes Target morally irresponsible. However, proponents of Target, and more specifically Target's security team, defend Target's actions because FireEye has a high false positive rate for security warnings, "security personnel typically do not get excited about such generic alerts because FireEye does not provide much information about those threats" which consequently, "would have made it tough to have singled out that threat as being particularly malicious" (Finkle, 2014). The proponents are advancing the argument that Target has acted in good faith and done their best, and the problem is only apparent in hindsight. This, however, does not absolve Target of moral responsibility. If FireEye's false positive rate was too high for Target, they should have pursued other replacement malware detection systems or explored different security techniques all together. Additionally, Target could have increased the number of security engineers such that each warning from FireEye could receive professional and competent attention. Indeed, Target's settlement with the state of Illinois enumerates various organizational and technical changes Target must imple-

ment as part of the agreement (Illinois Attorney General, 2017). For example, it required Target to include file integrity monitoring as part of a major overhaul of their existing security system, "including, but not limited to, a file integrity monitoring solution, designed to notify personnel of unauthorized modifications to critical applications or operating system files within the Cardholder Data Environment" (2017). This indicates that Target's practices could have been better and the breach could have been prevented. In any case, simply ignoring warnings from FireEye, a system designed to prevent exactly such attacks, is not a morally acceptable action.

5.3 Poor Third-Party Security Policy

Target violated the third principle of ACM's Code of Ethics: articulate, apply, and support policies and processes that reflect the principles of the Code. Particularly, Target did not have policies that articulate and apply the Code's principle of designing and implementing systems that are robustly and usably secure. As shown in Section 2, the attack was initiated through a third-party vendor called Fazio Mechanical Services. Fazio is a supplier of refrigeration devices and services, and it began working with Target to support the expansion of fresh food offerings (Plachkinova & Maurer, 2019). Fazio was a victim of a phishing attacks which infected their system with a credential stealing Trojan, and "due to the poor security training and security system...the Trojan gave the attackers full range of power over the Fazio's system" (Shu et al., 2017). Although Fazio claims their system and security measures were in full compliance with industry best practices, there are allegations that they used a free antivirus software that does not provide real-time protection (Plachkinova & Maurer, 2019).

Target is in violation of the third principle from ACM's Code of Ethics. The Code states, "leaders should pursue clearly defined organizational policies that are consistent with the Code and effectively communicate them to relevant stakeholders" (Anderson, 1992). Target did not have clearly defined organizational policies setting standards for third-party vendor

security consistent with the Code, "it is not clear whether Target enforced any ongoing security reviews of its vendors to ensure compliance with security best practices" (Plachkinova & Maurer, 2019). The settlement agreement with the State of Illinois reaffirms the need for strong third-party vendor policies, "Target shall develop, implement, and revise as necessary written, risk-based policies and procedures for auditing vendor compliance with Target's Information Security Program" (Illinois Attorney General, 2017). Target's poor third-party vendor security policy consequently makes Target collectively morally irresponsible.

6 Conclusion

The 2013 Target data breach is a case of collective responsibility and the problem of many hands. While it is challenging to fairly and precisely distribute moral blame for the data breach, it is possible to rigorously analyze the case where Target stands as a collective moral entity. The analysis via a deontological ethical framework shows that Target was morally irresponsible. More specifically, Target violated the following three principles from ACM's Code of Ethics: 1) design and implement systems that are robustly and usably secure, 2) maintain high standards of professional competence, conduct, and ethical practice, and 3) articulate and apply organizational policies that reflect the principles of the Code. I have demonstrated that Target was morally irresponsible by carefully analyzing specific technical and organizational failures that led to the data breach and by descriptively showing how these failure violate the ethical principles of ACM's Code of Ethics.

In a world that is becoming increasingly interconnected and data-intensive, data breaches are also becoming an unfortunately common phenomenon. These unfortunate events encourage engineers to reflect and improve the technical processes, designs, and implementations of the systems. Our understanding of the ethical dimension of data breaches needs to evolve along side. Corporations and businesses get a moral free-pass based on the argument that data breaches are too complicated to ethically analyze, however this analysis shows that we can still pass a normative judgment for the specific case. A better understanding of the

ethical dimension can lead to a more secure and trustworthy relationship between businesses and customers.

Word count: 3813

References

Anderson, R. E. (1992). ACM code of ethics and professional conduct. *Communications of the ACM*, *35*(5), 94–99.

Brinkman, B., & Carter, K. (2017, 03). The ACM code of ethics and professional conduct: Teaching strategies and the coming update (abstract only). In (p. 721-721). doi: 10.1145/3017680.3022340

Buttrick, H. G., Davidson, J., & McGowan, R. J. (2016). The skeleton of the data breach: The ethical and legal concerns. *Rich. JL & Tech.*, *23*, 1.

Culnan, M. J., & Williams, C. C. (2009). How ethics can enhance organizational privacy: Lessons from the choicepoint and tjx data breaches. *MIS Quarterly*, *33*(4), 673–687. Retrieved from `http://www.jstor.org/stable/20650322`

Donaldson, T., & Werhane, P. H. (2002). Introduction to ethical reasoning. *Ethical Issues in Business. A Philosophical Approach. New Jersey: Prentice May*, 1–11.

Finkle, J. (2014, Mar). *Target says it declined to act on early alert of cyber breach.* Thomson Reuters. Retrieved from `https://www.reuters.com/article/us-target-breach/target-says-it-declined-to-act-on-early-alert-of-cyber-breach-idUSBREA2C14F20140313`

Gikas, C. (2010). A general comparison of fisma, hipaa, iso 27000 and pci-dss standards. *Information Security Journal: A Global Perspective*, *19*(3), 132–141.

Illinois Attorney General. (2017, May). Attorney general madigan announces $18.5 million settlement with target over data breach. *Illinois Attorney General*. Retrieved from `https://illinoisattorneygeneral.gov/pressroom/2017_05/20170523b.html`

Plachkinova, M., & Maurer, C. (2019). Security breach at target. *Journal of Information Systems Education*, *29*(1), 7.

Poel, I. R., & Royakkers, L. M. M. (2011). *Ethics, technology and engineering: An introduction.* Wiley-Blackwell.

Rivero, N. (2018, Nov). *The biggest data breaches of all time, ranked.* Quartz. Retrieved from

https://qz.com/1480809/the-biggest-data-breaches-of-all-time-ranked/

Shu, X., Tian, K., Ciambrone, A., & Yao, D. (2017). Breaking the target: An analysis of target data breach and lessons learned. *arXiv preprint arXiv:1701.04940*.

Solove, D. J., & Citron, D. K. (2017). Risk and anxiety: A theory of data-breach harms. *Tex. L. Rev.*, *96*, 737.

Target corp. (2013, Dec). *Target.com*. Retrieved from https://corporate.target.com/press/releases/2013/12/target-confirms-unauthorized-access-to-payment-car

Target corp. (2014, Jan). *Target.com*. Retrieved from https://corporate.target.com/press/releases/2014/01/target-provides-update-on-data-breach-and-financia

US Senate, T. D. (2014). A "kill chain" analysis of the 2013 target data breach.

Wagner, N., Şahin, C. Ş., Pena, J., Riordan, J., & Neumayer, S. (2017). Capturing the security effects of network segmentation via a continuous-time markov chain model. In *Proceedings of the 50th annual simulation symposium* (p. 17).