A Systematic Study of the Control Failures in the Equifax Cybersecurity Incident

Ilya Kabanov, Stuart Madnick

Working Paper CISL# 2020-19

September 2020

Cybersecurity Interdisciplinary Systems Laboratory (CISL)
Sloan School of Management, Room E62-422
Massachusetts Institute of Technology
Cambridge, MA 02142

A Systematic Study of the Control Failures in the Equifax Cybersecurity Incident

Ilya Kabanov
Cybersecurity at MIT Sloan,
MIT Sloan School of Management
ilya.kabanov@sloan.mit.edu

Stuart Madnick
Cybersecurity at MIT Sloan,
MIT Sloan School of Management &
MIT School of Engineering
smadnick@mit.edu

Abstract

The Equifax data breach, announced in September 2017, occupies the top 10 charts of the largest data breaches in history. While the 148 million affected persons do not bring the event to the top, but the sensitivity of the data stolen makes it one of the most impactful incidents for consumers. The data breach has been investigated by various federal and state agencies that collected and analyzed over 45,000 pages of related documents. This research is built upon those findings and identifies the root causes of the control failures of Equifax's cyber defense system, and determines the improvements to reduce the likelihood of future similar incidents. We reconstructed the Equifax hierarchical cyber safety control system, identified what parts failed and why, and determined the necessary improvements by applying our Cybersafety method, inspired by Causal Analysis using Systems Theory (CAST). This work demonstrates how to discover reasons for the failure of safety and security mechanisms and compose improvement actions. It also provides a set of points individuals can evaluate in their organizations.

1. INTRODUCTION

On September 7, 2017, Equifax Inc. (NYSE: EFX) announced a cybersecurity incident that affected circa 143 million consumers in the United States. In this incident, cybercriminals exploited a vulnerability found in a U.S. website application and obtained access to certain files. Based on the company's investigation, the unauthorized access occurred from mid-May through July 2017 [1]. When the news broke, Equifax's stock dropped by 13% to $123.23 and continued falling until it hit a low at $92.98, wiping out 34% of the company's $17.5 billion market value [2].

The incident was investigated by various federal and state agencies that collected and reviewed over 45,000 pages of related documents. Based on the evidence, the Permanent Subcommittee on Investigations of the United States Senate Committee on Homeland Security and Governmental Affairs and the U.S. House of Representatives Committee on Oversight and Government Reform published their reports on the Equifax data breach [3], [4]. The investigation resulted in a $575 million Equifax settlement with the Federal Trade Commission (FTC), the Consumer Financial Protection Bureau (CFPB), and 50 U.S. states, which was announced on July 22, 2019 (Settlement) [5]. Although the reports describe what transpired, and the Settlement discloses the monetary and security recovery obligations of Equifax, these documents did not focus on why the failures and shortcomings occurred, and what should be done to prevent such failures. This report aims to identify the root causes of the deficiencies of Equifax's cyber defense system and propose improvements to reduce the likelihood of future such cybersecurity attacks. These findings and recommendations are likely to be very relevant to any large organization, not just Equifax.

1.1 Equifax Corporate Profile

Equifax began operations in 1899 and became a publicly owned corporation in 1965. The company is a leading global provider of information solutions and business processes, outsourcing services for businesses, governments, and consumers. Since 2006, Equifax had embarked on an ambitious growth

strategy. In 10 years, it made 18 acquisitions, making Equifax one of the largest private credit-tracking firms in the world [6]. The company reported $3.4 billion in revenue in 2017, which grew by 7% from 2016 [7]. The company's business is based on detailed consumer and business information derived from numerous sources and includes credits, financial assets, utility payments, employment records, income, demographics, and other sensitive information. According to the Equifax annual report in 2016, "[T]he company organizes, assimilates and analyzes data on more than 820 million consumers and more than 91 million businesses worldwide, and its database includes employee data contributed from more than 7,100 employers" [8]. That massive amount of sensitive information is what made Equifax a prime target for cybercriminals [9].

1.2 The U.S. Consumer Reporting Business

One of Equifax's main businesses is consumer credit reporting and scoring. Three major credit bureaus dominate the U.S. market, also termed Credit Reporting Agencies (CRAs)—Equifax, Experian, and TransUnion—that gather and maintain financial and personal information about individuals, as depicted in Figure 1 [10]. CRAs collect information from furnishers such as banks, thrifts, credit unions, and many other institutions. The details may include historical data about credit repayment, rent payment, employment, insurance claims, arrests, bankruptcies, check writing, and account management. Besides, CRAs procure information from public record providers and correlate it with the financial data obtained from the furnishers. In conclusion, CRAs compile all the data, analyze it, and create products, such as credit scores, credit reports, and marketing analytics data. These products are then sold to businesses such as lenders, landlords, employers, utility providers, and others from whom the individuals seek credit or services. For example, lenders evaluate risks associated with granting a loan to determine an interest rate. Another aspect is that insurance companies decide whether to issue an insurance policy and what premium to set based on this information. Individual consumers — data subjects — do not voluntarily provide data to the CRAs and cannot "opt-out" of the data collection process. However, they are entitled to a free annual report from each of the CRAs according to the Fair Credit Reporting Act (FCRA) and can purchase products that CRAs offer [11].

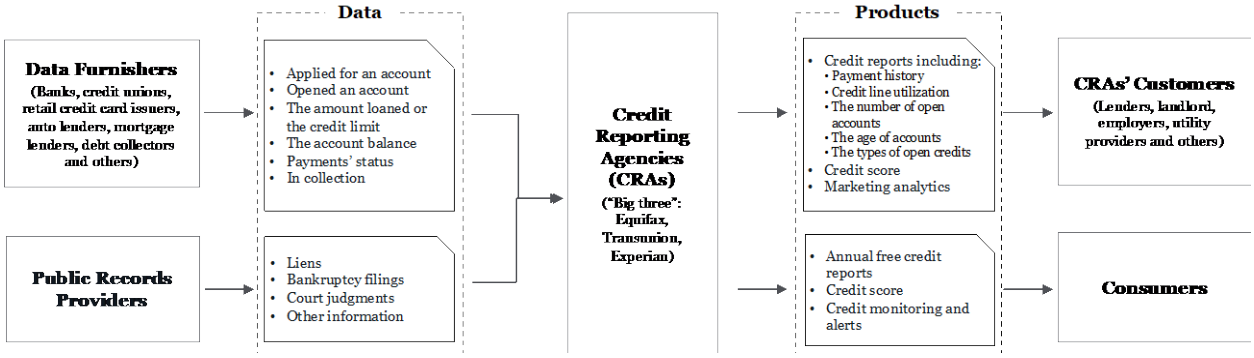

Figure 1. How the CRA business works.

This business model provides CRAs with significant power over consumers' lives because they discern every detail of an individual's financials, employment, immigration status, spending habits, and other aspects of life. The loss of confidentiality or integrity of these data may bring significant harm to individuals. Therefore, because of the massive specifics processed by CRAs and its high level of sensitivity, they are expected to safeguard the data in the best possible manner while applying the most effective and robust data protection cybersecurity practices.

1.3 The History of Cyberattacks on CRAs

The amount and value of information held by the CRAs make them lucrative targets for cybercriminals. Before the Equifax incident of 2017, two massive data theft occurred at Experian, another major CRA. The first incident occurred in 2013 when an Experian subsidiary, purchased in 2012, was compromised, exposing personal and financial data of more than 200 million consumers [12]. Later, in 2015, Experian

disclosed another breach of its computer systems where intruders stole circa 15 million Social Security Numbers (SSNs) and other data on individuals who applied for financing from the wireless provider T-Mobile [13].

Equifax had also experienced data breaches before the main incident in 2017. The complaint for civil penalties, injunctive relief, and restitution filed by the State of Indiana (Complaint) revealed that Equifax suffered a data breach almost every year from 2010 to 2017. In 2010, 2012, and 2014, the company notified the Indiana Attorney General's office on the breaches of the consumers' information [14]. In 2013, Equifax experienced "fraudulent and unauthorized access" of the personal financial information of U.S. First Lady Michelle Obama, Vice President Joe Biden, and celebrities, including vocalist Beyonce and actor Ashton Kutcher [15]. In 2015, the company experienced a data breach of its workforce solution, which impacted almost 100 companies and their employees [14]. Later in 2o15, the Wall Street Journal reported that "a former Equifax employee, believed to be an operative for Chinese entity, had accessed and stolen proprietary information from its systems" [16]. In 2016 and 2017, the company experienced two more data breaches of its workforce solution [14].

The history of the data breaches of CRAs, and Equifax in particular, confirms the fact that the company was a potential target for cyberattacks. Furthermore, it is also apparent that Equifax was aware that its information systems were susceptible to external and internal cyberattacks.

2. CYBERSAFETYMETHOD OF ANALYSIS

Concerning the Equifax data breach of 2017, we utilized the Cybersafety method of analysis, which is inspired by Causal Analysis using System Theory (CAST) [17], to analyze the cybersecurity protections that were in place at the time, examined why they failed, and create recommendations for improvements concerning cybersecurity. Cybersafety is a technique for cybersecurity analysis of complex systems, which was proposed by Salim and Madnick in [18] and later used by Khan, Madnick, and Moulton in [19]. Traditional causality models utilized for safety analysis attribute incidents to an initial component failure or human error that cascades through other components. Such models are adequate for systems with limited complexity, or systems that exhibit linear interactions and simple cause-and-effect linkages [20]. Furthermore, CAST was initially intended to study accidents, which are, of course, accidental. A cyberattack is not an accident. The results of [18] confirmed that the Cybersafety method could produce better outcomes than the Chain-of-Events Model, Fault Tree Analysis, Control Objectives for Information and Related Technology (COBIT) 5 for Information Security, and ISO/IEC[1] 27002 standard, which, are traditionally used for incident analysis. For brevity, although they may be differences from the original CAST, in this paper we will refer to our analysis as CAST.

The objective of CAST is to go beyond the traditional analysis of assigning blame on individuals or detecting a single failure. Instead, the focus is on identifying what controls were in place, why they were unsuccessful in preventing the incident, and on defining improvements that would enforce necessary safety constraints to prevent similar incidents [17].

The CAST-based analysis includes five major steps [21]:

1. Collect the necessary information about the system, including the physical system and processes involved in the loss; hazards (system states), which along with specific environmental conditions, can lead to an incident or loss; and safety requirements and constraints to prevent hazards.
2. Model the existing safety control structure starting from a high-level abstract control structure and then refine it.
3. Examine the components of the control structure to determine why they were not effective in preventing the loss.
4. Identify flaws in the control structure as a whole (the systemic factors).
5. Create recommendations for improvements to the safety control structure to prevent a similar loss in the future.

[1] The International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC).

In this case study, we applied the CAST method to the Equifax cyber incident, which led to the largest data breach at the time and affected close to half of the U.S. population.

3. DETAILS ABOUT THE INCIDENT

On July 29, 2017, Equifax identified, and a day later confirmed, a cyberattack on its Automated Consumer Interview System (ACIS). The detection came to light after Equifax's Information Technology (IT) team updated a Secure Sockets Layer (SSL) certificate on the SSL visibility appliance that monitored the encrypted inbound and outbound network traffic between Equifax IT systems, including ACIS and the Internet. The SSL certificate had expired nine months earlier, in November 2016. After updating the certificate, Equifax employees detected suspicious Internet traffic exiting ACIS that contained image files related to consumer credit investigations. They traced the traffic to an IP address in China—a country where Equifax did not operate. After blocking the IP address, Equifax noticed suspicious traffic from a second IP address owned by a German Internet Service Provider (ISP) that was leased to a Chinese ISP. Based on those findings, Equifax decided to shut down ACIS temporarily [3]. The timeline of the key events associated with the Equifax cybersecurity incident is displayed in Figure 2 below.

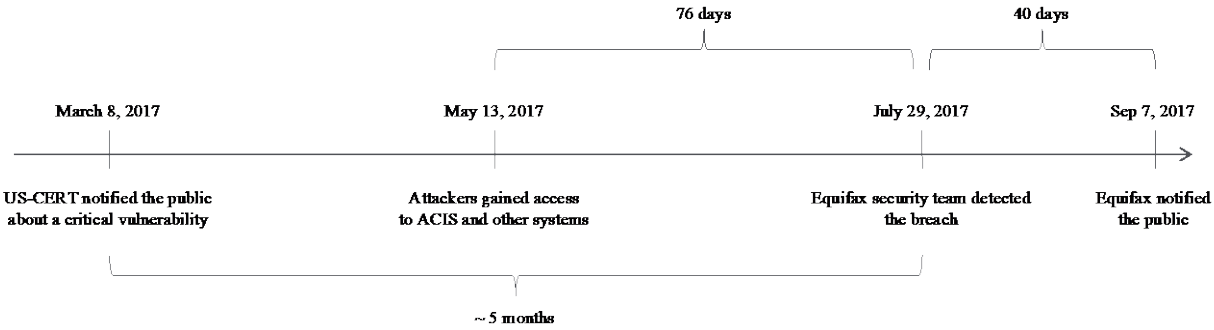

Figure 2. The Equifax cybersecurity incident timeline.

Further analysis revealed that the suspicious traffic was the result of a successful cyberattack on ACIS that started on May 13, 2017. The attackers gained access to ACIS and databases containing consumers' Personal Identifiable Information (PII) and then exfiltrated the data while remaining undetected in the Equifax network for 78 days.

On September 7, 2017, Equifax publicly announced a cybersecurity incident, potentially impacting circa 143 million U.S. consumers. The announcement detailed, "the information accessed [by the attackers] primarily includes names, SSNs, birth dates, addresses, and, in some instances, driver's license numbers" [1]. On October 2, 2017, Equifax announced the completion of a forensic investigation and disclosed that the attackers accessed information of an extra 2.5 million U.S. consumers [22]. Later, on March 1, 2018, the company admitted that "Equifax was able to identify approximately 2.4 million U.S. consumers whose names and partial driver's license details were stolen, but who were not included in the affected population discussed in the company's prior disclosures," thus totaling the number of affected consumers to 148 million [23].

Before the incident, on March 8, 2017, the United States' Computer Emergency Readiness Team (US-CERT) had notified the public about a vulnerability involving the Apache Struts2 Web Framework[2] that would allow an attacker to execute commands on affected systems [24]. Thatvulnerability was present in ACIS and was exploited by the attackers to gain access to the Equifax network. Noteworthy, a working exploit that allowed to attack vulnerable web sites was available to the public on March 11, 2017, on GitHub [25].

[2] Apache Struts is a free, open-source, Model-View-Controller (MVC) framework for creating modern Java web applications and is extensible using a plugin architecture, and ships with plugins to support REST, AJAX and JSON.

4. APACHE STRUTS2 VULNERABILITY

The vulnerability in Apache Struts2, exploited during the attack on Equifax, is listed in the National Vulnerability Database (NVD) as CVE-2017-5638, was assigned the severity score of 10 out of 10. Incorrect exception handling and error-message generation by Jakarta Multipart parser in Apache Struts 2 2.3.x before 2.3.32 and 2.5.x before 2.5.10.1 allowed remote attackers to execute arbitrary commands via a crafted Content-Type, Content-Disposition, or Content-Length Hypertext Transfer Protocol (HTTP) header [26].

The nature of the vulnerability was that Apache Struts2, by default, uses Jakarta's Multipart parser when the Content-Type header is set to multipart / form-data. When an invalid Content-Type header, with a payload written as an Object-Graph Navigation Language (OGNL) expression, passes into Jakarta's Multipart parser, an exception is triggered, which is then passed into a function that builds error messages. The function then evaluates the OGNL payload, resulting in arbitrary code execution.

The solution provided by Apache suggested upgrading to Apache Struts version 2.3.32 or 2.5.10.1. As an alternative, switching to a different implementation of the multipart parser was also proposed as a remediation measure [27].

Although the Apache Struts2 vulnerability is often noted, there were a series of control failures that both created this vulnerability as well as made the impact much more severe that are usually overlooked. That is the focus of this paper.

5. CAST ANALYSIS OF THE EQUIFAX CYBERSECURITY INCIDENT

The CAST analysis provided in this paper is predominantly based on the details of the Equifax cyber incident documented in [3], [4], and [14].

5.1 The System, Hazards and the Safety Constraints

The CAST begins with the identification of the system for analysis, key hazards, and the safety constraints that should have been considered at the design and operation of the system to prevent those hazards [21]. In the context of this analysis, the system is ACIS, and the hazards are the states of the system that, in conjunction with other conditions, can lead to an incident. In the Equifax case, an incident is defined as the high-volume loss of confidential consumer information, and the operation of the system with an exploitable vulnerability is an example of a hazard.

We used the Cyber Kill Chain framework introduced by Lockheed Martin Corporation to identify hazards. The application of a uniformed framework for the hazards' identification streamlines the analysis of cyber incidents with CAST and ensures consistency across various CAST analyses performed by different parties.

The Cyber Kill Chain consists of the seven phases associated with the typical steps taken by individuals who launch a computer attack [28]:

1. *Reconnaissance*. Identification and selection of targets are often made with the help of network scanners and social media research.
2. *Weaponization*. Compilation of a remote access Trojan virus with an exploit into a deliverable payload, ready to be delivered into the victim's computer system.
3. *Delivery*. Transmission of the weapon to the targeted environment (e.g., through a phishing[3] email or USB removable media).
4. *Exploitation*. Execution of the attacker's code, which can target a vulnerability in one or several elements of the software stack.
5. *Installation*. Deployment of a remote access tool on the victim's system to establish a presence inside the system.
6. *Command and Control* (C2). Formation of a C2 channel within the attacker's server to control the victim's system.

[3] Phishing refers to criminal activity that attempts to obtain sensitive information by fraudulent means.

7. *Actions on Objectives*. Execution of the intruders' objectives (e.g., data collection and exfiltration from the victim's environment).

During the analysis of the attack on Equifax, we identified seven significant hazards associated with phases 1 and 3–7 of the Cyber Kill Chain. The Weaponization phase often occurs outside of the victim's system and is not related to direct hazards. In Table 3 below, we displayed the hazards, the Cyber Kill Chain phases related to them, and the associated safety controls violated.

#	Hazard	Cyber Kill Chain Phase	Safety Constraint Violated
H-1	Publicly accessible information about vulnerabilities in the company's IT systems.	Reconnaissance	Non-disclosure of unnecessary details about software used in IT systems.[4]
H-2	Malicious exploit delivered to a system.	Delivery	Blockage of malicious requests sent by the attackers.
H-3	Operation of a system with an exploitable vulnerability.	Exploitation	Elimination of critical vulnerabilities in IT systems through patching and vulnerability management.
H-4	The spread of attack beyond its entry point.	Installation	System's secure design (isolation, authentication, least privilege).
H-5	The attacker's covert actions in the network.	Command and Control (2C)	Detection of suspicious traffic in the network with the Intrusion Detection and Prevention Systems (IDS/IPS).
H-6	Unauthorized access to unencrypted data.	Actions on Objectives	System's secure design (encryption of PII).
H-7	The bulk exfiltration sensitive data.		Identification and potential blockage of unauthorized data exfiltration.

Table 3. The system hazards and system safety constraints.

Hazard **H-1** is the exposure of details about software systems and their components to the public. The authors hypothesize the attackers could have identified that Equifax ACIS utilized a susceptible version of Apache Struts2 using Open Source Intelligence (OSINT)[5] techniques. While the authors have no details about the specific approach, the attackers used for Reconnaissance, some websites, which used the vulnerable version of Apache Struts2, could be found through "Google Dorking."[6] For example, the Google search "intitle:"Struts Problem Report" intext:"development mode is enabled"" returned a list of web sites that used Apache Struts2. The authors advised that as a general safety constraint, a system must not reveal its technical details, and should instead prevent sensitive information from being indexed by search engines. In particular, private access websites that do not rely on publicly-indexable Internet content should deny all crawling from any directory within the website.

Hazard **H-2** is the malicious exploit delivered to the system. The vulnerability in Jakarta's Multipart parser used by Apache Struts2 was exploitable through sending a malicious Content-Type header in an HTTP request. At the time of the attack, ACIS was open to accepting malicious requests as it was not protected by a Web Application Firewall (WAF). In 5.2, we hypothesize on the origins of this fact and elaborate on how a WAF could have prevented the incident.

[4] The authors are not privy to the exact method the attackers used to identify an attack vector.
[5] Open Source Intelligence (OSINT) is the collection and analysis of publicly available information that can be used for planning and executing a cyberattack.
[6] "Google Dorking" is a technique that relies on powerful Google search engine data and can be used to find vulnerable web applications and servers in the internet.

Hazard **H-3** is the operation of a system with a critical vulnerability. The malicious request sent to ACIS exploited an unidentified and unpatched vulnerability in Apache Struts2 and led to the remote execution of unauthorized code delivered in the payload.

Hazard **H-4** is the spread of the attack in the Equifax systems. Other Equifax systems permitted access to their sensitive data through ACIS. The investigation revealed that the attackers gained access to a data repository that contained unencrypted application credentials to other sensitive Equifax databases not isolated from ACIS, and that stored a potentially excessive amount of confidential information and PII [3]. In 5.2, we discuss how the system's insecure design and security architecture contributed to the cyber incident.

Hazard **H-5** is the covert actions of the attackers in the Equifax network. Equifax's Intrusion Detection and Prevention Systems (IDS/IPS) failed to identify and block intruders from establishing control over the affected system and exfiltrating the data. At the time of the Equifax incident, its IDS/IPS was not capable of analyzing the encrypted network traffic entering and leaving Equifax IT systems, including ACIS, because of the SSL certificates installed in the IDS/IPS had expired. We analyze the causes of this failure in 5.3.

Hazard **H-6** is unauthorized access to unencrypted data. The investigation revealed that none of the PII contained in the stolen data was encrypted [4]. Data encryption and the proper management of encryption keys can protect the confidentiality of data, even in a case when attackers succeed in obtaining unauthorized access to computer systems. Moreover, since the ACIS was the subject of the Payment Card Industry Data Security Standard (PCI DSS), at least credit cardholder data should have been protected. In 5.2, we explore the root causes of the lack of encryption of data at rest in the context of secure system design.

Hazard **H-7** is the exfiltration of voluminous sensitive data. We estimated that the attackers were able to exfiltrate at least 14-16Gb[7] of data they harvested from the Equifax databases during the 76 days they remained concealed in the Equifax network. At the time of the incident, Equifax did not have a Data Loss Prevention (DLP) system, which could have detected and blocked the bulk transfer of sensitive data outside of the network.

All seven hazards materialized during the Equifax cyber incident, and the safety constraints that should have been enforced through the hierarchical safety control structure, and the system design and architecture failed to prevent them. In the following two sections, we analyze the reasons for the safety constraint violations as part of the system design and operations.

5.2 Flaws in Design and Security Architecture that Contributed to the Cyber Incident

In the present section, we analyze how the safety constraints could have prevented the cyber incident had they been introduced as part of the ACIS' system design and the company's security architecture. We begin with a set of software engineering design principles that ensure the necessary security properties of a software system and act as its safety constraint. The safety constraints based on these principles create a layered defense, also known as defense in depth system, to safeguard a software system and the information it processes. The investigation reports and other publicly available sources helped us identify four significant violations of the secure design principles, as well as two security architecture elements that were lacking, and hence, contributed to the event.

First, the fundamental principle of secure software design is *isolation (compartmentalization)*. Isolation requires computer sub-systems to be compartmentalized by separating them from each other using physical devices and/or security controls [29]. It can minimize an attack surface by keeping sub-systems partitioned or grouped in self-contained subsets, so the compromise of one will not jeopardize another. The investigation report revealed that ACIS lacked isolation, disclosing that "[T]he ACIS application was not segmented from other, unrelated databases." Furthermore, the Sun systems, which hosted ACIS, had a shared file system across the environment that allowed for access to administrator files across all systems. Notes or configuration files from one system could be accessed from other systems. In that unsegmented network, the attackers were able to move laterally throughout the networks and reach systems beyond the first compromised ACIS [4].

[7] The estimation is based on the number of records stolen, an approximate record size, and a 30% overhead coefficient.

The second crucial secure design principle is called *authentication*. A software system should accept that the environment where the system operates is hostile and requires authentication before granting access to its data. Two major authentication issues were discovered in ACIS and other Equifax IT systems. First, the company employed a weak authentication policy. The system privilege accounts were protected with modest passwords. For instance, one of the databases accessed by the attackers was protected with a four lower-case letter password, which matched the name of the database [14]. The use of weak passwords was common in Equifax IT systems. According to the filing in the U.S. District Court for the Northern District of Georgia, Atlanta Division, Equifax "employed the username "admin" and the password "admin" to protect a portal used to manage credit disputes" [30]. Another example of using a weak password was provided in the KrebsOnSecurity blog, sharing that the Equifax Argentina employee portal's administrative account was protected by "admin/admin" login/password combination [31]. Another improper authentication practice employed by Equifax was storing the application credentials in an unencrypted format on a file share. Russ Ayres, Interim Chief Security Officer of Equifax, testified that "[I]f Equifax had limited access to sensitive files across its systems, the attackers may not have found the stored application credentials used to access sensitive databases outside the ACIS environment" [4].

The third principle is the *least privilege*, which restricts the rights and access of a user and system to only the necessities needed to execute its task. This protective measure guards against the spread of an attack, limiting its potential impact. ACIS had excessive access to other systems not required for its operations, thus violating the least privilege principle. The testimony of Russ Ayres states that "ACIS only needed access to three databases to function, but it was unnecessarily connected to many more," thus confirming the disregard of the least privilege principle. The systems remained open for attackers to access their data using the stolen credentials [4].

The fourth principle is *encryption*. One of the most effective, universal data protection controls—encryption—was also missing in ACIS and other systems accessed by the attackers [4]. Encryption ensures the confidentiality and integrity of data even in the case of a successful attack. If the sensitive and PII stored in the Equifax databases had been encrypted with effective management of encryption keys, the consequences of the attack would be minimized or even non-existent. In contrast, all data exfiltrated from the Equifax systems were stored in the systems without encryption; therefore, compromising the confidentiality of the stolen PII and other sensitive information. As mentioned earlier, ACIS was the subject of PCI DSS, and "Equifax was in the process of making the ACIS application [PCI DSS] compliant when the data breach occurred" [4]. The PCI DSS Requirement 3 mandates the protection of stored cardholder data, which includes a primary account number, expiration date, cardholder name, and a service code [32]. The techniques may include encryption, truncation, masking, or hashing the critical components of cardholder data. Although this requirement is applied to the cardholder data, but not to all the PII and sensitive data stored in ACIS and other impacted systems, its proper application could have reduced the impact. The investigation materials do not reveal the root causes of the lack of encryption of PII and sensitive data. Still, the authors hypothesize it was a mindful company-wide architectural decision or the lack thereof since none of the IT systems impacted by the attack employed encryption.

Graeme Payne, a former Senior Vice President and Chief Information Officer for Global Corporate Platforms at Equifax, shared another relevant aspect of system design choice. He supervised ACIS at the time of the incident. Payne reflected on the importance of an adequate data retention policy: "[T]here was one other factor that did not get a lot of coverage in the incident report is data retention. [W]as it necessary to have that much data on these systems at all?" A similar pattern of collecting and storing customer information that was not required to make a purchase or a return (e.g., driver's license number) led to the broader exposure of customer information during the attack on TJX in 2006 [18].

The company's security architecture had two main elements intended to protect ACIS and other IT systems from external attacks: a network perimeter firewall, and an Intrusion Detection and Prevention Systems (IDS/IPS). The report shows that "the ACIS environment was comprised of two web servers and two application servers, with a network-level firewall(s) set up at the perimeter of the web servers" [4]. We discuss later in 5.3 how Equifax's IDS/IPS failed to identify and block an attack. In the current section, we will focus on the missing safety elements in Equifax's security architecture: a WAF and a Data Loss Prevention (DLP) system.

From the report, and because there is no reference to a WAF in the Equifax network in other publicly available sources related to the Equifax incident, we can conclude that while the network perimeter firewall was in place, the system was not protected by a WAF. A WAF is a layer of protection in addition to a network-level firewall, where the objective is to prevent successful exploitations of security vulnerabilities in web applications. The most common types of these attacks include SQL injection, cross-site scripting (XSS), and Distributed Denial of Service (DDoS) attacks. As described earlier, the exploitation of Apache Struts2 vulnerability was possible by placing an invalid value in the Content-Type header, thus causing the error message function to execute the OGNL payload passed along in the header. Memon confirmed in his analysis that even an open-source WAF (e.g., ModSecurity[8]) could have protected against the attack by detecting the malicious strings passed in the HTTP headers and blocking such traffic should it be tasked to whitelist valid content types or blacklist OGNL expressions [33]. The United States Senate Permanent Subcommittee on Investigations confirmed that the other two large CRAs, TransUnion and Experian, had a WAF in place configured to block attacks on Apache Struts [3]. Therefore, we can conclude the lack of a WAF as part of Equifax's security architecture contributed to the success of the cyberattack.

Another missing safety mechanism in the security architecture was a Data Loss Prevention (DLP) system. Equifax was in the business of processing large volumes of sensitive data, where leakage could impact the lives of consumers and undermine businesses. The data loss as a result of a cyberattack could be one of many possible causes of losing the confidentiality of sensitive data, and data loss prevention is the process of safeguarding it. The process is often supported by a DLP system, where the objective is to identify and monitor the moves of sensitive data to the outside of the network through various channels. The presence of a DLP system in the Equifax security architecture could have prevented the massive exfiltration of the data. Also, having a DLP system is recommended for successful compliance with the PCI DSS requirement A3.2.6, which demands the existence of mechanisms for detecting and preventing clear-text payment card numbers from leaving the controlled environment via unauthorized channels [32].

Once we concluded reviewing the evidence, we had to put them into the context of the environment and time-frame in which the system was developed and implemented to understand the impact on the absence of secure, design-enabled safety constraints. Graeme Payne testified that "ACIS was the dispute and disclosure system that was built in (...) the late 1970s to address the requirements of the [Fair Credit Reporting Act]" [4]. The year of the system's creation can explain the lack of data encryption in the system in the first place since, at that time, the encryption was not considered as an obligatory mechanism to protect consumer information and the cryptographic technologies were not available for general commercial use. However, the legacy nature of the system could not justify the ultimate lack of secure design controls, since Sun Solaris, the operating system that hosted ACIS, could provide proper isolation of individual systems and support the principle of least privilege. The investigation report reveals that "by requiring [network file system (NFS)] requests come from privileged source ports, the server can potentially avert attacks from systems on which the attacker does not have full administrative access" [4]. Another detail is that the web application part of ACIS used Apache Struts2, the component, which was compromised, and that had been released in 2006. Hence, it was not as timeworn as the original core of the system created in the 1970s. Therefore, the incident cannot be purely attributed to the legacy nature of ACIS. The original system design and security architecture decisions were made at the time when cyberattacks were almost non-existent, and many of the protective measures deemed mandatory nowadays were not necessary. However, the attack landscape transformed from the 1970s to 2016, and the essential safety constraints should have been introduced to address the evolving threats that emerged since the initial development of the system.

5.3 The Hierarchical Safety Control Structure

In the previous sections, we analyzed the safety constraints that were neglected as part of the system's design and Equifax's security architecture. In the current section, we center on the hierarchical safety control structure related to system operations and analyze how it was supposed to prevent the cyber incident, and why it failed.

[8] https://www.modsecurity.org/.

Figure 4 shows the hierarchical safety control structure where the safety controls are organized hierarchically across four levels:

- *Level 1* – Equifax Intrusion Detection and Prevention Process (IDPP), where the objective was to monitor inbound and outbound Internet traffic entering and exiting Equifax's IT systems for malicious activities.
- *Level 2* – Equifax Information Technology (IT) and Information Security (IS) teams that operated vulnerability identification and PCI DSS compliance processes.
- *Level 3* – Equifax management and its board of directors that oversaw the company's strategy and operations, including risk management.
- *Level 4* – Federal agencies (FTC and CFPB) and the states, which had enforcement authority over the CRAs, and the payment card brands that mandated compliance with PCI DSS.

The hierarchical safety control structure includes the roles and responsibilities of each component, controls for executing those tasks, and feedback to gauge the effectiveness of controls [17]. We identified nineteen unique safety control loops that aimed at protecting Equifax IT systems from the risks, and we examined the roles that each of the safety mechanisms played in the incident.

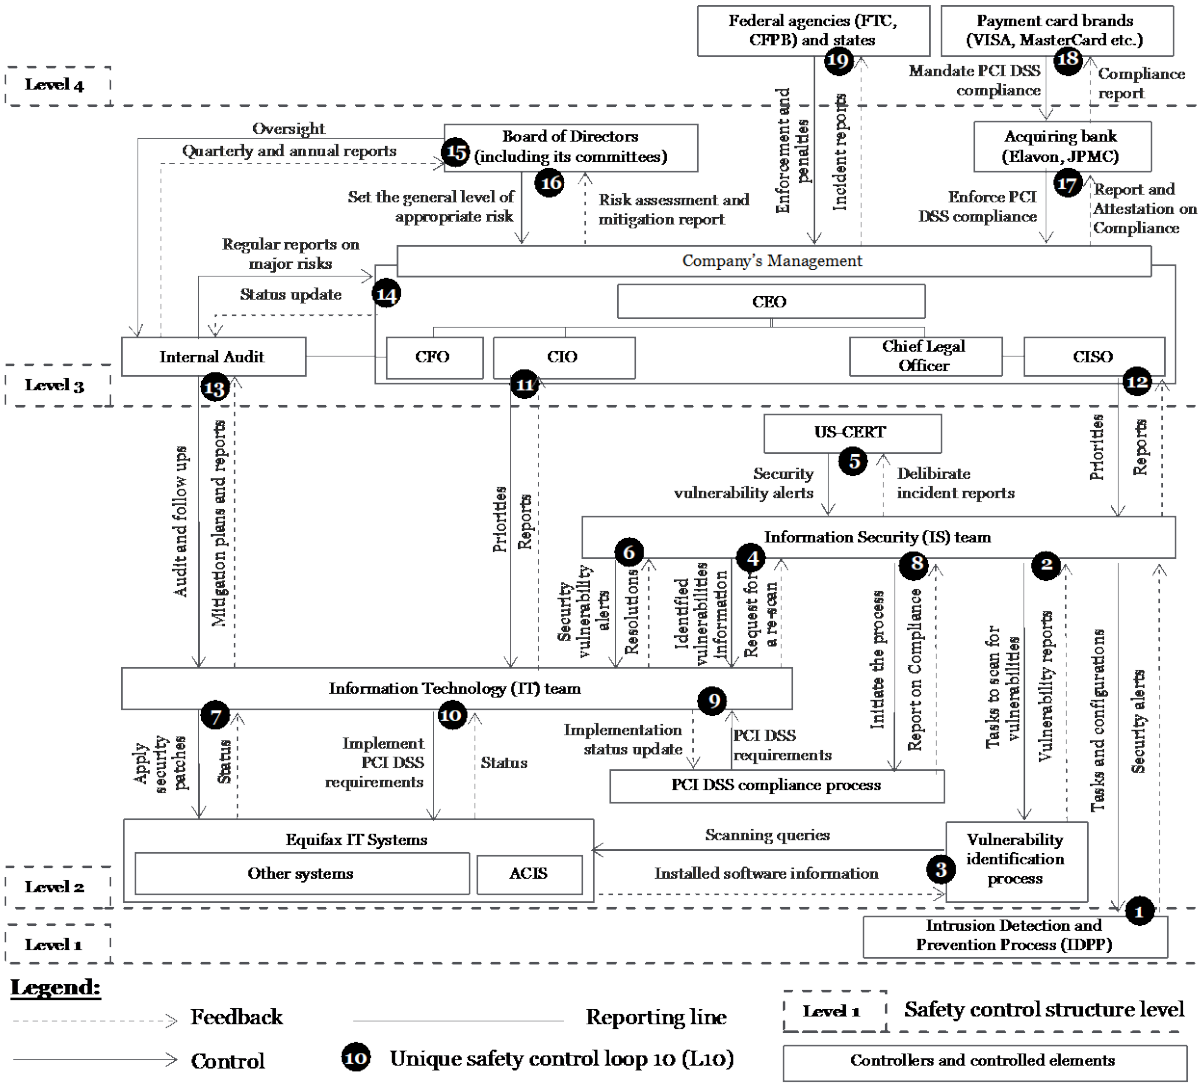

Figure 4. The Equifax hierarchical safety control structure.

At **Level 1**, the objective of Equifax Intrusion Detection and Prevention Process (IDPP) was to monitor and analyze the inbound and outbound Internet traffic to identify and block malicious activities within Equifax's IT systems, including ACIS. The process was powered by the Equifax IDS/IPS. The IDPP was operated by the Equifax IS team (safety control loop L1). With the support of the data from the Equifax data breach report, in Figure 6, we indicated the most probable architecture of the IDS/IPS used at Equifax to fulfill this control loop [4]. The solution included an open-source traffic analyzer Snort[9], which had the objective of identifying suspicious activities in the flowing network traffic and deter them or raise an alert. It also had an SSL visibility appliance, which decrypted transit traffic between Equifax internal servers and external Internet users and systems.

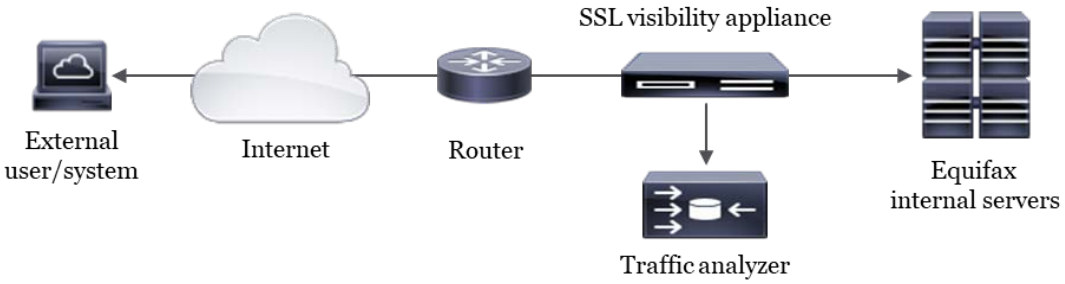

Figure 6. The Equifax IDS/IPS most likely architecture.

However, the SSL certificates necessary to decrypt traffic and installed in the SSL visibility appliance expired in November 2016. Therefore, the IDS/IPS could not analyze such traffic and identify malicious activities; instead, it just let traffic flow through uninterrupted. At the same time, the IS team did not realize that the system was incapable of identifying and preventing intrusions and continued relying on its protective capability against attacks on Apache Struts2. The investigation reports confirm this and disclose that "[T]he Equifax Countermeasures team installed the Snort rule on its IDS/IPS written to detect Apache Struts exploitation attempts on March 14" [4]. The failure of the safety control loop L1 had two major factors. The first was the Equifax manual process for tracking and updating the several hundred SSL certificates. The investigation revealed that the process was error-prone, as it depended on individuals responsible for IT systems. Our conclusion is also established by the fact the Equifax had recognized the deficiency of the manual SSL renewal process and began the deployment of an automated SSL certification management tool in 2016, but did not complete it before the incident [3]. The second concern was the lack of alerts to the IS team that IDS/IPS was nonoperational for almost nine months (until July 29, 2017). The lack of warnings was possible because the system settings allowed traffic to bypass IDS/IPS in case the appliance fails to decrypt it [3].

At **Level 2**, the objective of the IT and IS teams was to remediate vulnerabilities in the Equifax IT systems through the control loops L2–L7 and achieve compliance with PCI DSS through L8–10. The control loops L2 and L3 were responsible for initiating vulnerability scans of the Equifax IT systems, including ACIS, performing such scans, and then reporting them to the IS team. The Equifax IS team scanned the IT systems exposed to the Internet using an unspecified scanner and the McAfee Vulnerability Manager tool. However, "the scan was on the root directory, not the subdirectory where the Apache Struts was listed," and, therefore, it left the vulnerability CVE-2017-5638 in Apache Struts2 used by ACIS unnoticed [4]. The foremost cause of the failure was that the IS team relied on the assumption the vulnerability scanners guaranteed the identification of vulnerabilities and lacked the awareness that vulnerability scanners had difficulty detecting the Apache Struts2 vulnerability. The use of two different scanners, and the limited understanding of how Apache Struts2 works, had further reinforced their confidence. It was known that vulnerability scanners had blind spots, and the detection rate of the popular scanners did not exceed 80% even for the most authenticated scans [34]. Vulnerability scanners identify vulnerabilities through checking the running services and installed versions of software on the system against the publicly known vulnerabilities, and through scanning target servers for specific vulnerabilities. The primary challenge with identifying the Apache Struts2 vulnerability was the fact that Apache Struts2 was not managed via a package manager but

[9] https://www.snort.org/

distributed as a standalone JAR file, which could reside anywhere on a server and was often embedded inside application 'WAR' and 'EAR' files. The identification of the Apache Struts2 vulnerability through scanning a web application also required understanding the application's URL path. We recreated a testing environment with the vulnerable version of Apache Struts2 installed and attempted to identify the vulnerability using *Nmap*[10], an open-source tool for network discovery and security auditing, and *struts-pwn*[11], an exploit for Apache Struts CVE-2017-5638. Both tools failed to identify the Apache Struts2 vulnerability without providing them with a specific path to an application. Furthermore, Beaupre, a cybersecurity expert and principal instructor at SANS, tested Nessus and Nexpose, popular commercial vulnerability scanners, and confirmed they were not finding the vulnerability if not ran in the directory brute-forcing mode [35]. These particulars prove that the control loops L2 and L3 had a known limitation for finding the Apache Struts2 flaw. Still, the IS team, which operated the vulnerability identification process, was not aware of this limitation, thus leading the control loops L2 and L3 to the failure. The IS team accepted the false-negative results as confirmation that the Equifax IT systems do not have that exposure. As a result, the IT team was not informed of the vulnerability in ACIS, and the control loop L4 was never activated.

The control loops L5 and L6 formed another safety mechanism that aimed at alerting IT teams about recent critical vulnerabilities in various software systems, which may or may not be present in the Equifax IT environment. The IS team received the alerts from US-CERT (L5) and then disseminated the message among more than 400 recipients across the organization, including members of the IT team (L6). The investigation revealed that the IS team received an alert from US-CERT on March 8, 2017 (before the time of the compromise on May 13, 2017) regarding a new remote execution vulnerability and was advised to review the Apache Security Bulletin and upgrade Apache Struts to 2.3.32 or Struts 2.5.10.1 [24]. Therefore, we can conclude that L5 functioned as it should. Nevertheless, the notification of relevant members of the IS team about a critical vulnerability (L6) failed since the person who was operating ACIS did not receive the alert as he/she was not part of the recipients' list [3]. The Equifax former CEO, in his testimony to the Senate Banking Committee, also pointed at a failure of a Senior Vice President, whose team was responsible for ACIS to inform the team members about the vulnerability in Apache Struts2 that he was made aware of through the alert he had received from the IS team [36]. We argue that the overall design of the control loop was error-prone and did not guarantee the fulfillment of its objectives with confidence. The design had three significant issues. First, the notification mechanism was based on a group email, which was not a call for action, but rather an information newsletter, thus lacked a feedback loop. For example, the investigation report indicated that of 400 plus recipients who received the alert about the Apache Struts2 vulnerability sent by the IS team, only one replied to the email for clarification [3]. Second, the email alert did not consider its relevance for the recipients based on the software they were operating. The investigation report explained that the company had lacked the comprehensive inventory of its software. Therefore, the IS team could not correlate the warning as part of L5 with the impacted software installed in Equifax and request the IT team to patch it [3]. As a final point, L6 was based on an unrealistic assumption that the managers who were the primary recipients of vulnerability alerts can and will correlate vulnerabilities with the versions and types of software systems used by their teams and demand their patching.

In conclusion, the control loop L7 was responsible for the prompt patching[12] of the Equifax IT systems to ensure that its components execute their function and do not contain publicly known security vulnerabilities. However, the Equifax patching process had a reactive nature and relied on the impractical hypothesis that the vulnerability scanners and alerts guaranteed vulnerability identification and notification. As a result, the patching process was not triggered for the Apache Struts2 vulnerability in ACIS, leaving the system vulnerable. We conclude that the reactive design of the patching process caused the failure of the safety control loop L7. The findings of internal audit on the company's configuration and patch management, conducted in October 2015, confirmed our theory. The internal audit determined that the systems were only patched after notifications from the Equifax IS team about specific vulnerabilities instead of applying security patches proactively. The audit had recommended that management shall implement a proactive patching process, and the management responded with an action plan and the estimated

[10] https://nmap.org/
[11] https://github.com/mazen160/struts-pwn
[12] Patching is a process of applying a set of changes to a computer program, often called a "patch" that is designed to update or improve the program and includes fixing security vulnerabilities.

remediation date of December 31, 2016 [3]. Nonetheless, at the time of the incident, the patching approach was still reactive, and its upgrade was not a priority for the IT team. Another remarkable detail was that the contractual patching requirements for third parties that operated some Equifax IT systems, including ACIS, contradicted Equifax's patch management policy. For instance, the Complaint revealed that the patching of ACIS was outsourced to a third-party firm that had a contractual obligation to apply patches to ACIS within six months, thus breaching Equifax's policy that required to patch critical vulnerabilities within 48 hours [14].

The objective of the safety control loops L8–10 was to ensure the compliance of ACIS with PCI DSS. As learned from the Complaint, ACIS stored credit card payment details, and therefore, must have been compliant with PCI DSS. In particular, PCI DSS has 12 requirement categories, which include the use of a change detection mechanism, strong access control measures, vulnerability identification, and installation of patches, as well as the maintenance of an inventory of the system components in the scope of PCI DSS compliance [32]. The post-incident PCI DSS forensic report summary revealed that ACIS failed all twelve PCI DSS requirement categories [14]. Three vital factors contributed to the failure of the loops L8-10. First, the Complaint stated that Equifax purposely excluded ACIS from the scope of PCI DSS compliance as it could not meet the requirements. For example, the same document reflected that ACIS contained more than 23 million unique Primary Account Numbers (PANs)[13] in plain text, which is prohibited by the standard [14]. Second, the PCI DSS reports misrepresented the actual compliance status. For instance, the Complaint stated that "[I]n 2016, Equifax employee [name] completed the Report on Compliance for Equifax's Global Consumer Solutions business" and included the company's patch management policy, which demanded the remediation of critical vulnerabilities within 48 hours and high vulnerabilities in 30 days, while it was known that Equifax had failed to comply with its policy and had more than 1,000 open externally exposed vulnerabilities [14], [3]. In conclusion, compliance with PCI DSS was not a priority for the IT and IS teams. The Equifax former CISO testified, "[T]he PCI preparation started about a year before [the incident]," but the "plan fell behind, and these items did not get addressed" [4]. We later discuss the factors that led to the failure of the company's management to prioritize mandatory compliance of ACIS with PCI DSS with the IT and IS teams.

Level 3 of the safety control loops hierarchy comprises the safety control loops L11–17. We will analyze the role that the company's management, the board of directors, the internal audit, and the acquiring banks played in prioritizing the work of the IT and IS teams, and how the control loops at Level 3 contributed into the failure of Level 2 safety control loops, which later led to the cyber incident. We begin our analysis by identifying the factors that contributed to the inadequacy of the vulnerability identification and patching processes at Level 2. Then we analyze the role of the company's board of directors in setting and monitoring the company's overall risk appetite. We will conclude our analysis by understanding the role of the acquiring banks and the company's management in the failure of the PCI DSS compliance control loop.

The safety control loops L13 and L14 formed a governing mechanism powered by the Equifax internal audit team over the state of the Level 2 safety control mechanisms and associated risks. The investigation report revealed that in 2015 the Equifax's internal audit team evaluated the company's configuration and patch management practices, thus imposing safety constraints as part of L13. The auditors identified critical deficiencies, stating that "current patch and configuration management controls are not adequately designed to ensure Equifax systems are securely configured and patched in a timely manner" [3]. Also, the auditors recognized the error-prone patching management as a risk and debriefed the audit report to the company's management as part of L14. The auditors also obtained a formal commitment from the IT leadership to address the issues by December 31, 2016. These facts signal that the company's management was aware of the patching process issues. However, the improvements had not been completed by the time of the incident, making it clear that the management failed to prioritize the modernization of the patching system through L11 and L12 despite their former commitment to do so. Moreover, the internal audit team did not perform a follow-up or formal re-audit to confirm the execution of the commitments extended by the management (L14).

We accessed the company's proxy statement 2017 to further understand why the mechanism that was supposed to control the internal audit failed. The proxy statement revealed that the internal audit was

[13] Primary Account Number (PAN) is a unique payment card number that identifies the issuer and the particular cardholder account.

overseen by the audit committee of the board of directors through quarterly and annual reporting to the full board, which had the objective of holistically overseeing risk management at Equifax. However, security-related risks were not in the scope of the board's audit committee, but "[Technology committee, which] focuses on technology-related risks and opportunities, including data security" [37]. We theorized that this misalignment between the responsibilities of the committees of the board could have created a gap in the board's cyber risk oversight. This hypothesis is somewhat confirmed by the statements in the Equifax consent order (Order) from June 2018, in which Equifax was demanded "[W]ithin 30 days from the effective date of this Order, the Board or Audit Committee shall improve the oversight of the Audit function" as a response to the incident [38].

We would also like to refer to other evidence, which confirms that loops L11 and L12 responsible for setting priorities for the company's IT and IS teams, in turn, experienced systematic failures. The investigation reports indicated that the company's management realized the risks of operating Equifax IT systems on the legacy Sun servers when the migration into a new data center begun in 2015. That was necessary because "threat vectors were changing too quickly and this [was] one way to mitigate risk" [4]. In addition to the infrastructure migration, since 2014, ACIS had been planned to be replaced with a new system. Yet, this initiative was "enduring multiple delays as the company prioritized the completion of other initiatives," as Graeme Payne testified. He also stated that "[S]ecurity was probably also a risk, but it was not the primary driver [of the replacement]" [4]. Likewise, Equifax's management failed to prioritize the deployment of an automated SSL certificates management system, thus keeping the SSL update process manual and error-prone. These specifics confirm the systematic issues with the control safety loops L11 and L12, failing to provide adequate priorities to the IT and IS teams despite the management's awareness about the existing deficiencies of the control safety loops at Level 2.

As shown in the hierarchical safety control structure, effective remediation of vulnerabilities required great collaboration, synchronization, and communication between the IT and IS teams. Nonetheless, the organizational structure of Equifax at the time of the incident contributed to the disconnect between the IT and IS teams, and led to an accountability and communication gap. The Chief Information Officer (CIO) reported to the Equifax Chief Executive Officer (CEO), while the Chief Security Officer (CSO) reported to the Chief Legal Officer (CLO). The investigation report highlighted that the reporting structure created a siloed environment where "[I]nformation rarely flowed from one group to the other. Collaboration between IT and Security mostly occurred when required" [4]. In parallel, the CSO was not considered part of the senior leadership team. As a result, she was not often invited to the CEO quarterly senior leadership team meetings, where CIO was always present. After the incident, the company admitted the inefficiency of the IT and IS reporting structure and established the role of the Chief Information Security Officer who reports to the CEO, thus ensuring a productive approach to security. It is interesting to note that the lack of such a CISO was identified as a major factor in our analysis of the 2007 TJX attack [15] but had still not been learned by many companies, including Capital One, almost a decade later.

As discussed earlier, the company's board of directors was responsible for approving its priorities and determining an acceptable level of cyber risk through the safety control loop L16. The proxy statement of 2017 confirms the board's responsibility to establish the general risk appetite level, including data security risks (L16). In detail, the proxy statement discloses that "[O]ur Board oversees risk management at the Company. The Board then sets the general level of risk appropriate for the Company" [37]. We argue that the control safety loop L16 failed to fulfill its objective to establish an appropriate risk level. Three factors steered to this conclusion. First is the high cost of the incident that cannot be called an "acceptable risk level" for any enterprise. In the 2020 proxy statement, Equifax reported $663.5M in costs related to the 2017 cybersecurity incident in addition to the $575-700 million that Equifax agreed to pay as part of the Settlement [39], [5]. The overall related expenses of about $1,24 to $1.36 billion, equals circa 35% of the company's annual revenue. Second is the lack of the acceptable risk level's adjustments made by the board based on the previous data breaches that occurred at the company in 2010, 2012-2017, and a growing number of data security breaches that hit a new record of 1,579 in 2017 jumping by 44.7% from 2016 [40], [14]. The third is the executive compensation rules approved by the board of directors that were focused on profit, thus "motivating them [executives] to prioritize revenue above all other considerations, including information security" [14]. These factors confirm the inference that the control loop L16 failed to set an acceptable risk threshold, but prioritized the revenue growth for the company's executives without setting limits on cyber risk.

As a result, the company presumed its effort in prioritizing cybersecurity as adequate or beyond. The company's proxy 2017 stated that the Chief Financial Officer (CFO) was rated "Distinguished" on his annual 2016 objectives, which included "[C]ontinuing to advance and execute global enterprise risk management processes, including directing increased investment in data security." Similarly, the CLO was recognized for "[C]ontinuing to refine and build out the Company's global security organization" [37]. Noteworthy, the CFO, CLO, and the head of the internal audit who prioritized risk management activities for the company and oversaw the internal audit, one of the foremost corporate risk management instruments, all have remained in the company following the incident[14]. At the same time, the CIO and the CSO were terminated.

To close, we will analyze how the Level 3 safety control structure contributed to the failure of the company to obtain compliance for ACIS with PCI DSS. As we learned from the Complaint, Equifax management was aware that the company was qualified as a Level 1 merchant as it processed more than 6 million credit card transactions a year [14]. As a Level 1 merchant, Equifax was obliged to perform an annual on-site assessment by a Qualified Security Assessor (QSA), prepare a Report on Compliance (ROC), conduct a quarterly network scan by an Approved Scan Vendor (ASV), and submit an Attestation of Compliance (AOC) form to a bank acquirer [32]. The same Complaint reassured that the company's management was aware of the mandatory PCI DSS compliance for ACIS. Still, instead of prioritizing these activities to the IS and IT teams through L11 and L12, they excluded ACIS and Automated Credit Report On-line (ACRO), Equifax's primary credit reporting database as they did not meet the majority of PCI DSS requirements. The Complaint confirmed this hypothesis by concluding that the multiple ROCs and AOCs filed by Equifax "contained false and misleading information" [14]. As a result, the IT and IS teams did not receive appropriate prioritization from management, and ACIS remained incompliant with PCI DSS at the time of the incident. Moreover, the state Attorney General elaborated that "Equifax did so to keep its current customers and data furnishers who required PCI compliance and to acquire new customers and data furnishers who required PCI compliance" [14]. Based on those facts, we can conclude that the loops L11 and L12 failed to accomplish their objective to prioritize PCI DSS compliance actions to the IT and IS teams.

Acquiring banks or acquirers[15] play a significant role in enforcing the compliance of credit card merchants with PCI DSS. Acquirers allow merchants to accept credit card payments from the payment card brands (e.g., Visa, MasterCard, Discover, American Express, etc.), and they must ensure that their merchants satisfy PCI DSS requirements during the entire contract. Acquirers usually demand PCI DSS compliance as a contractual requirement and impose potential fines on a merchant for the failure to comply with PCI DSS. Equifax partnered with two acquiring banks, JP Morgan Chase and Elavon. The safety control loop L17 was intended to impose PCI DSS compliance requirements on Equifax by those acquirers. In our analysis, we referred to the case of the PCI DSS incompliance of TJX at the time of the incident, and it helped us identify three potential sources of the loop's failure [18]. Similar to the TJX case, a potential conflict of interest/role between a bank acquirer and merchant made enforcement a challenge, as the compliance with PCI DSS is neither required by federal law in the U.S. nor by the state, except the State of Nevada. Second, it was problematic for an acquirer to determine the merchant's scope of PCI DSS compliance. The Equifax case revealed that the complexity of the IT systems and their integrations allowed the company to misrepresent the genuine compliance status in its attestations. Third, an acquirer is not empowered to perform technical validations or verifications of the merchants' compliance but rely on the merchant's AOC supported by the ROC[16] prepared by a Qualified Security Assessor (QSA) hired by the merchant. But, as noted above, ACIS was specifically excluded from the security assessment, thus was not reported in any AOC or ROC. At the same time, according to payment card brands' rules, the acquirers hold accountability for the compliance of their merchants, while they rely on verification instruments that cannot guarantee proper validations. In summary, we can conclude that the safety control loops L11, L12, and L17 failed at meeting their objectives of prioritizing the work of IT and IS teams to achieve compliance with the PCI DSS at Level 2 of the safety control structure.

We reached the top **Level 4** of the control structure, which includes two safety control loops, L18, which was responsible for the enforcement of PCI DSS compliance by the payment card brands, and L19, aimed at enforcing the federal and state regulations over CRAs, including Equifax.

[14] At the time of writing the case in June 2020, the mentioned leadership roles were occupied by the same executives.
[15] An acquiring bank is a bank or financial institution that processes credit or debit card payments on behalf of a merchant.
[16] ROC and AOC are mandatory for Level 1 merchants.

The payment card brands play an essential role in mandating acquirers to enforce PCI DSS compliance over their merchants. Although the payment brands demand it, the actual compliance and its comprehensiveness are not guaranteed, as we discovered earlier. Interesting fact: following the data breach public announcement, the payment card brands required Elavon, one of the acquirers, to demand Equifax to comply with PCI DSS. However, the Complaint revealed that "Continuing to the present day, Equifax has continued to accept credit card transactions (...) knowing that its systems are not compliant" [14]. All these facts ratify that L18 failed to mandate the acquirers to enforce PCI DSS compliance over their customer Equifax. The failure of the payment card brands to ensure the compliance of Equifax with PCI DSS was not unique and the unsatisfactory compliance of merchants with PCI DSS has been known for more than a decade. Back to 2006, Salim and Madnick identified that payment card brands failed to ensure TJX's compliance with PCI DSS and confirmed that "the lack of full compliance with PCI DSS also contributed to the cyberattack" [18]. Verizon also confirmed that although the percentage of merchants that maintain full compliance with PCI DSS has increased since the introduction of PCI DSS in 2004, it has reached a mere 55.4% in 2016 and even reduced to 36.7% in 2018 [41]. Based on these and the previously discussed facts, we can conclude the control loop L18 had an error-prone design as it established, and continues to rely on the acquirers-driven enforcement mechanism that fails at achieving its objectives.

While PCI DSS compliance is not demanded by law, Equifax and other CRAs are the subjects of various regulations at a federal and state level, and specific data security obligations imposed by CRAs' partners through their contracts. The overview of the CRA's regulatory landscape is displayed in Figure 7 below. The FCRA, the Financial Services Modernization Act of 1999, called the Gramm-Leach-Bliley Act (GLBA) and the Dodd-Frank Wall Street Reform and Consumer Protection Act (Dodd-Frank Act) are the three primary federal laws regulating the CRAs. Although the FCRA's focus is on the accuracy of consumers' credit-related records, it mandates reasonable confidentiality of the information and requires CRAs to provide consumer reports only to the entities that have an acceptable purpose for receiving it. The GLBA, Section 501(b) mandates that CRAs ensure the protection and confidentiality of personal information. It requires the FTC to promulgate rules, which the FTC articulated in the Safeguards Rule requiring financial institutions to have a comprehensive information security program, including administrative, technical, and physical safeguards to address internal and external risks to the security, confidentiality, and integrity of customer information [42]. In addition to the Safeguards Rule, the FTC has general authority under Section 5 of the FTC Act to combat unfair or deceptive acts or practices [43]. The Dodd-Frank Act grants the authority to the Consumer Financial Protection Bureau (CFPB) to enforce federal consumer laws against CRAs, compose the regulations, oversee, and examine them [44]. As part of its general supervisory authority, the CFPB shall require reports and sometimes conduct examinations for purposes of (1) assessing compliance with the requirements of federal consumer financial law, including the Dodd-Frank Act's prohibition of unfair, deceptive, and abusive acts or practices; (2) obtaining information about the activities and compliance systems of the examined institution; and (3) detecting and assessing risks of consumer financial products and services to consumers and markets [45].

In addition to the federal agencies, the CRAs are also subject to multiple state regulations enforced by state attorneys general and aim at protecting the personal information of consumers. All states have data breach laws that require the CRAs to notify consumers. Almost all states have strong consumer protection laws that prohibit unfair and deceptive acts and practices similar to the FTC Act. At least 13 states have laws that require information safeguards similar to those in the GLBA, and three states have enacted specific information security duties that apply to CRAs. A majority of states require organizations to protect the confidentiality of SSNs [46]. As emphasized earlier, the CRAs also shall comply with PCI DSS and with the security requirements of their partners, financial institutions that have legal requirements published by the Federal Financial Institutions Council (FFIEC), and enforced by federal regulators.

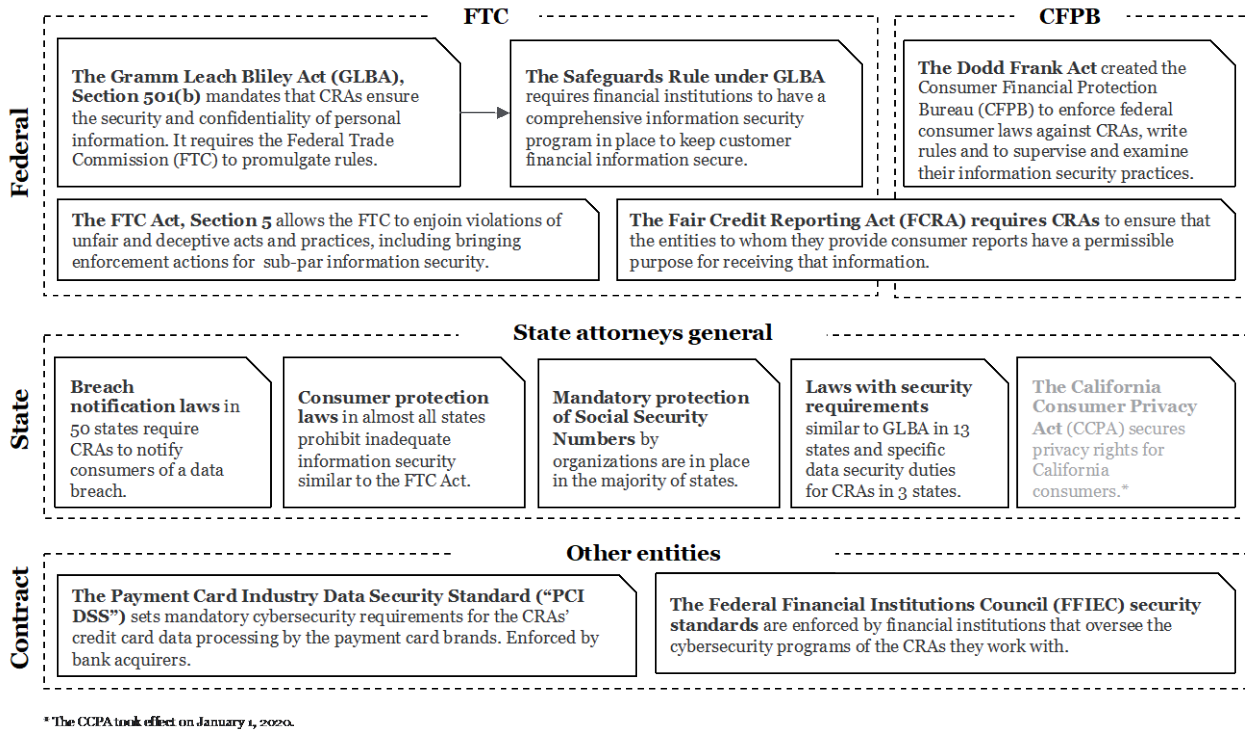

Figure 7. The CRA's regulatory landscape.

The control safety loop L19 was responsible for enforcing the federal and state regulations described above over the CRAs. The FTC reported that "[S]ince 2002, FTC has brought more than 70 cases against companies that have engaged in unfair or deceptive practices involving inadequate protection of consumers' data" [47]. However, Michael Clements, Director at the United States Government Accountability Office (GAO) explained in his testimony before Congress on March 26, 2019, that "FTC lacks a practical enforcement tool for imposing civil money penalties that could help to deter companies, including CRAs, from violating data security provisions of GLBA and its implementing regulations," because "FTC did not have civil penalty authority for violations of requirements under GLBA and must identify affected consumers and any monetary harm they may have experienced" [48]. Also, the GAO report GAO-18-559 highlights that "[FTC] does not have proactive, supervisory authority to examine CRAs' compliance with the FTC Act" [49]. CFPB also had issues with enforcing data security measures over the CRAs according to its mandate under the Dodd-Frank Act. In 2019, the CFPB staff admitted that "[they] do not consider data security risks during their examination prioritization method and have not reassessed the process to determine how to incorporate such risks going forward." The CFPB staff also noted that "unless the bureau finds that the institution has violated the Dodd-Frank Act's prohibition on unfair, deceptive, or abusive acts or practices, or another provision of federal consumer financial law over which CFPB has authority, the bureau cannot take enforcement action, and can only make supervisory recommendations" [48]. However, after the Equifax breach, CFPB used its existing supervisory authority to develop internal guidelines for examining data security and conducted the data security examinations of some CRAs. Similar to the situation with FTC and CFPB, publicly accessible information about proactive actions taken by the states before the incident to enforce adequate data security practices over Equifax was not available.

We can summarize the three root causes of the failure of L19 to identify deficient security practices at Equifax and enforce data security provisions of federal consumer protection regulations. First, both FTC and CFPB lacked the legal basis for taking enforcement actions. Second, FTC did not have supervisory authority to examine the CRAs' compliance. Third, CFPB did not proceed with the CRAs' proactive examination as their outcomes were limited to recommendations and could not result in enforcement actions. Appositively to the lack of proactive data security examinations at federal and state levels, FTC, CFPB, and states took massive enforcement actions after the incident. Those actions resulted in more than

60 government investigations from federal agencies and U.S. state attorneys general and resulted in the Settlement in which Equifax agreed to pay at least $575 million, and up to $700 million [5].

In summary, we can conclude that the Equifax incident ensued as a result of the systematic failures of the safety control loops at all four levels of the hierarchical safety control structure. As observed from the CAST analyses of other major incidents, most parts of the safety control structure contributed in some way to the propagation of the occurrence and, thus, can be improved [50]. Most control elements of Equifax safety structure were sub-par in design based on invalid assumptions, the lack of support from the controls placed higher in the hierarchy, and had feeble controls at the highest level of the structure, which failed to establish a necessary operating momentum in the entire hierarchical safety control system.

6. IMPROVEMENTS

Earlier, we reviewed the failed safety control loops and analyzed the sources of the failures. In this section, we will provide systematic improvements to strengthen the safety control structure and reduce the probability of similar cyber incidents in the future. The CAST analysis relied on the information from official incident reports; therefore, some of the proposed enhancements will echo the actions demanded in the Settlement and the Order issued in response to the Equifax incident. We provided the key elements of the hierarchical safety control system, including the design and operations of the IT system, the organizational composition of the information technology and security leadership, cybersecurity controls at the board level, and the regulations and enforcement of CRAs.

Organize defense in depth of the IT system. Viega and McGraw proposed that "[I]f more than one protective measure exists in a system, then the level of security is not necessarily determined by the weakest part" [51]. Protection mechanisms layered one over another increase the difficulty of an attack as it can be prevented at various layers of defense. We identified seven design principles and protection mechanisms that can improve the protection of IT systems in addition to IDS/IPS and a network perimeter firewall that already existed at Equifax.

1. *The compartmentalization of* the system's components minimizes how an attack on one element can debilitate the rest of the system. Viega and McGraw recommend segmenting a system into several components that can be protected in parallel to prevent the spread of an attack other parts of the system [51].
2. *The principle of least privilege* commands that every entity in a system should be granted a minimum set of permissions to perform its designated tasks [52]. In conjunction with compartmentalization, the least privilege principle can provide a finer grain for granting access to system components and prevent attackers from entering other parts of the system via elevated privileges on the compromised element.
3. *Zero Trust* security model proposes that the components of a system must mistrust all requests from other parts of the system until they are authenticated and authorized.
4. *The encryption* of PII, sensitive, and confidential information at rest and in transit have been recognized as good practices by the FTC and European commissions. Even in the case of encrypted data being stolen and the encryption keys were not compromised, the event is not considered a data breach.
5. *Data retention* policy prevents storing sensitive information beyond the time necessary to provide service.
6. *Web Application Firewall (WAF)* protects systems exposed to the Internet from attacks based on the exploitations of web application-based security vulnerabilities.
7. *Data Loss Protection (DLP)* system which objective is to identify, monitor, and potentially avert unauthorized moves of sensitive data inside and to the outside of the network through various channels.

Organizational composition of information technology and security leadership. It is necessary to have a senior executive in charge of security in the organization. The reporting structure of the security and IT teams remains a debatable topic, but among various organizational designs, the most common are

reporting to Chief Risk Officer, Chief Executive Officer or Chief Information Officer. The key is to ensure that security is represented at an executive leadership level to advise on cyber-related risks and make sure that cyber risks are considered during decision-making, resource allocation, and prioritization. Another significant factor to consider when designing the organizational structure of IS and IT teams is effective communication, which should be enabled through an organizational design and priorities alignment.

Cybersecurity at a board level. The board must prioritize cybersecurity to the company's management. The cyber-risk oversight handbook for corporate boards highlights that the board must strike the appropriate balance between protecting the security of the organization and ensuring the company's growth and profitability [53]. The handbook establishes the five principles which must be adhered to by the board to create a cyber-resilient organization, and we would like to propose them as one of the recommendations.

1. Understand and approach cybersecurity as an enterprise-wide risk management issue, not just an IT issue integrating cyber-risk topics into full-board discussions regarding all significant aspects of the company's business. These topics may include new product offerings, mergers, and acquisitions, new market entry, deployment of new technologies, major capital investment decisions such as facility expansions, or IT system upgrades.
2. Understand the legal implications of cyber risks. As we learned earlier, a company can be the subject of a growing number of regulations and laws in conjunction with the lack of coordination among rule makers and regulators. Therefore, the board should assess whether the organization evaluates and addresses cyber risks from the legal perspective.
3. Obtain and maintain access to cybersecurity expertise and hold systematic discussions with management about cyber risks. The knowledge about cybersecurity risks helps the director understand various threats and vulnerabilities. Similar to financial literacy, "Not everyone on the board is an auditor, but everyone should be able to read a financial statement and understand the financial language of business" [53].
4. Request management to have a company-wide cyber-risk management framework and ensure its adequate staffing and budget. Cybercriminals do not distinguish between organizational and reporting structures and use gaps in the digital interdependencies. Therefore, board directors should ensure that management is taking an enterprise-wide approach to protecting the company's computer systems.
5. Establish cyber risk management and discuss with management risk appetite and what perils must be avoided or mitigated, which can be accepted or transferred through cyber insurance. As it discoursed earlier in the case, the lack of such conversations and establishing risk tolerance can be bought in error by the company's management as openness to unlimited risks in the wake of pure business growth. The handbook proposes that "[D]iscussions of risk tolerance will help identify the level of cyber risk the organization is willing to accept as a practical business consideration" [53].

CRA regulations and enforcement. Following the Equifax data breach, GAO examined deficiencies in the oversight of CRAs by FTC and CFPB and the lack of proactive enforcement of the federal consumer protection law over CRAs [54]. We would like to include three recommendations proposed by GAO as they can strengthen the oversight of CRAs to prioritize the security of the consumers' data they store and process.

1. Consider the FTC civil penalties authority to enforce GLBA's safeguarding provisions.
2. CFPB should identify and track all CRAs that meet thresholds of the larger market participant CRAs[17] through registering CRAs or leveraging state information.
3. CFPB should evaluate its examination process to ensure that it includes data security risks and leads to needed steps taken by CRAs to incorporate the risks identified during examination.

7. CONCLUSIONS

This section completes the CAST analysis of the Equifax cyberattack. The CAST method allowed us to discover insights from various levels of the safety control system and identify the failures of the safety control loops inside and outside the company, which contributed to the incident. The CAST method also

[17] CRAs with more than $7 million in annual receipts from consumer reporting.

provided broad recommendations across all four levels of the hierarchical safety control structure leading to holistic improvements. We identified the system hazards, the hierarchical safety control structure and the safety loops putting a specific emphasis on the origins of their collective failure. The findings highlighted in the document can be used by a broad audience to explore the gaps in the cyber defense systems on all levels of their organizations, beginning from technical security controls to the regulatory compliance the organizations may be the subject of. Also, we introduced a new approach to using an industry-standard Cyber Kill Chain framework for identifying cyber hazards. The broader adoption of a standardized framework will contribute to the uniformity and comparability of the analysis of future cyber incidents using CAST. Finally, we discussed the necessary improvements to reduce the likelihood of similar incidents from occurring.

8. ACKNOWLEDGEMENTS

The comprehensiveness of the analysis was made possible through the release of the incident details to the public in various publications: The Data Protection Report on the actions taken by Equifax and Federal Agencies in response to the 2017 breach (published by the United States Government Accountability Office in August 2018), the Staff Report on how Equifax neglected cybersecurity and suffered a devastating data breach (produced by the United States Senate Permanent Subcommittee on Investigations at the Committee on Homeland Security and Governmental Affairs), and the Majority Staff Report on the Equifax data breach (published by the U.S. House of Representatives Committee on Oversight and Government Reform in December 2018) [49], [3], [4]. We would like to acknowledge Shaharyar Khan for his advice on applying CAST to cyber incident analysis. We extend our appreciation to Graeme Payne, a former Senior Vice President and Chief Information Officer for Global Corporate Platforms at Equifax, for his kind review of this report and the insights he provided regarding the incident and antecedent events. His feedback and the details he provided helped us to conduct our analysis with an extraordinary level of granularity regarding the pertinent technical and organizational aspects. We also appreciate Mr. Payne's opinions on the measures we suggested.

This research was supported, in part, by funds from the members of the Cybersecurity at MIT Sloan (CAMS) consortium.

9. REFERENCES

[1] Equifax, "Equifax Announces Cybersecurity Incident Involving Consumer Information," 7 September 2017. [Online]. Available: https://investor.equifax.com/news-and-events/news/2017/09-07-2017-213000628. [Accessed 28 August 2019].

[2] Equifax, "Historic Price Lookup," 7 September 2017. [Online]. Available: https://investor.equifax.com/stock-information/historic-price-lookup. [Accessed 28 August 2019].

[3] Permanent Subcommittee on Investigations, United States Senate, "How Equifax Neglected Cybersecurity and Suffered a Devastating Data Breach," United States Senate, 2019.

[4] Reform, U.S. House of Representatives Committee on Oversight and Government, "The Equifax Data Breach," 2018.

[5] FTC, "Equifax to Pay $575 Million as Part of Settlement with FTC, CFPB, and States Related to 2017 Data Breach," 22 July 2019. [Online]. Available: https://www.ftc.gov/news-events/press-releases/2019/07/equifax-pay-575-million-part-settlement-ftc-cfpb-states-related. [Accessed 26 May 2020].

[6] Equifax, "Equifax releases fourth quarter 2018 results," 1 March 2018. [Online]. Available: https://investor.equifax.com/news-and-events/news/2019/02-20-2019-215733514. [Accessed 28 August 2019].

[7] Equifax, "Equifax 2017 Annual Report (Form 10-K)," 1 March 2018. [Online]. Available: https://investor.equifax.com/~/media/Files/E/Equifax-IR/Annual%20Reports/2017-annual-report.pdf. [Accessed 28 August 2019].

[8] Equifax, "Equifax Annual Report 2016," 22 February 2017. [Online]. Available: https://investor.equifax.com/~/media/Files/E/Equifax-IR/Annual%20Reports/2016-annual-report.pdf. [Accessed 5 July 2020].

[9] Equifax, "Equifax Releases Details on Cybersecurity Incident, Announces Personnel Changes," 15 September 2017. [Online]. Available: https://investor.equifax.com/news-and-events/news/2017/09-15-2017-224018832. [Accessed 28 August 2019].

[10] The Board of Governors of the Federal Reserve System, "Credit Reports and Credit Scores," 2018. [Online]. Available: https://www.federalreserve.gov/creditreports/pdf/credit_reports_scores_2.pdf. [Accessed 28 August 2019].

[11] The Federal Trade Commission, "The Fair Credit Reporting Act," September 2018. [Online]. Available: https://www.ftc.gov/system/files/documents/statutes/fair-credit-reporting-act/545a_fair-credit-reporting-act-0918.pdf. [Accessed 11 August 2020].

[12] B. Krebs, "Experian Lapse Allowed ID Theft Service Access to 200 Million Consumer Records," 10 March 2014. [Online]. Available: https://krebsonsecurity.com/2014/03/experian-lapse-allowed-id-theft-service-to-access-200-million-consumer-records/.. [Accessed 29 August 2019].

[13] B. Krebs, "Experian Breach Affects 15 Million Consumers," October 2015. [Online]. Available: https://krebsonsecurity.com/2015/10/experian-breach-affects-15-million-consumers/. [Accessed 29 August 2019].

[14] The State of Indiana, "The Complaint for Civil Penalties, Injunctive Relief, and Restitution," 6 May 2019. [Online]. Available: https://calendar.in.gov/site/oag/event/ag-curtis-hill-files-lawsuit-against-equifax-over-2017-data-breach/. [Accessed 22 June 2020].

[15] Reuters, "Equifax Says Probing Unauthorized Access of U.S. Credit Reports," 12 March 2013. [Online]. Available: https://mobile.reuters.com/article/amp/idUSL1N0C4AH020130312. [Accessed 05 July 2020].

[16] The Wall Street Journal, "Before It Was Hacked, Equifax Had a Different Fear: Chinese Spying," 12 September 2018. [Online]. Available: https://www.wsj.com/articles/before-it-was-hacked-equifax-had-a-different-fear-chinese-spying-1536768305. [Accessed 05 July 2020].

[17] N. G. Leveson, Engineering a Safer World, Cambridge: The MIT Press, 2011.

[18] H. Salim and S. Madnick, "Cyber Safety: A Systems Thinking and Systems Theory Approach to Managing Cyber Security Risks," Composite Information Systems Laboratory (CISL), Cambridge, 2014.

[19] S. Khan, S. Madnick and A. Moulton, "Cybersafety Analysis of Industrial Control System for Gas Turbines," Cybersecurity Interdisciplinary Systems Laboratory (CISL), Cambridge, 2018.

[20] N. G. Leveson, STPA Handbook, Cambridge: Massachusetts Institute of Technology, 2018.

[21] Massachusetts Institute of Technology, "CAST Tutorial. Causal Analysis using System Theory. STAMP approach to accident analysis.," 26 March 2013. [Online]. Available: http://psas.scripts.mit.edu/home/wp-content/uploads/2013/04/CAST_Tutorial-2013.pdf. [Accessed 29 August 2019].

[22] Equifax, "Equifax Announces Cybersecurity Firm Has Concluded Forensic Investigation of Cybersecurity Incident," 2 October 2017. [Online]. Available: https://www.equifaxsecurity2017.com/2017/10/02/equifax-announces-cybersecurity-firm-concluded-forensic-investigation-cybersecurity-incident/. [Accessed 14 June 2020].

[23] Equifax, "Equifax Releases Updated Information on 2017 Cybersecurity Incident," 1 March 2018. [Online]. Available: https://investor.equifax.com/news-and-events/news/2018/03-01-2018-140531340. [Accessed 29 August 2019].

[24] US-CERT, "Apache Software Foundation Releases Security Updates," 8 March 2017. [Online]. Available: https://www.us-cert.gov/ncas/current-activity/2017/03/08/Apache-Software-Foundation-Releases-Security-Updates. [Accessed 7 September 2019].

[25] mazen160, "An Exploit for Apache Struts CVE-2017-5638," 11 March 2017. [Online]. Available: https://github.com/mazen160/struts-pwn. [Accessed 29 August 2019].

[26] National Institute of Standards and Technology, "CVE-2017-5638 Detail," 12 March 2017. [Online]. Available: https://nvd.nist.gov/vuln/detail/CVE-2017-5638. [Accessed 7 August 2019].

[27] Apache, "Apache Security Bulletins," Apache, 19 March 2017. [Online]. Available: https://cwiki.apache.org/confluence/display/WW/S2-045. [Accessed 24 May 2020].

[28] M. J. C. R. M. A. Eric M. Hutchins, "Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains," in *The 6th International Conference on i-Warfare and Security*, Washington, D.C., 2011.

[29] M. J. A. Mardjan, "Open Security and Privacy Reference Architecture," 2020. [Online]. Available: https://security-and-privacy-reference-architecture.readthedocs.io/en/latest/. [Accessed 11 August 2020].

[30] THE UNITED STATES DISTRICT COURT FOR THE NORTHERN DISTRICT OF GEORGIA ATLANTA DIVISION, "OPINION AND ORDER," 28 January 2019. [Online]. Available: http://securities.stanford.edu/filings-documents/1063/EI00_15/2019128_r01x_17CV03463.pdf. [Accessed 23 June 2020].

[31] Krebs on Security blog, "Ayuda! (Help!) Equifax Has My Data!," 12 September 2017. [Online]. Available: https://krebsonsecurity.com/2017/09/ayuda-help-equifax-has-my-data/. [Accessed 23 June 2020].

[32] PCI Security Standards Council, "PCI DSS Quick Reference Guide 3.2.1," July 2018. [Online]. Available: https://www.pcisecuritystandards.org/documents/PCI_DSS-QRG-v3_2_1.pdf?agreement=true&time=1592886797608. [Accessed 22 June 2020].

[33] F5, "F5 tech blog," F5, 22 January 2018. [Online]. Available: https://www.nginx.com/blog/modsecurity-apache-struts-cve-2017-5638/. [Accessed 25 May 2020].

[34] H. Hannes, T. Sommestad, J. Almroth and M. Persson, "A quantitative evaluation of vulnerability scanning," *Information Management & Computer Security*, vol. 19, no. 4, pp. 231-247, 2011.

[35] SANS, "Enterprise threat and vulnerability assessment: Equifax and Struts2," 12 June 2018. [Online]. Available: https://www.sans.org/webcast/recording/citrix/108120/143015. [Accessed 22 June 2020].

[36] PBS NewsHour, "Former Equifax CEO testifies before Senate Banking Committee," 4 October 2017. [Online]. Available: https://www.youtube.com/watch?v=11Ft3Ts3mfY. [Accessed 7 July 2020].

[37] Equifax, "Notice of 2017 annual meeting and proxy statement," 24 March 2017. [Online]. Available: https://investor.equifax.com/~/media/Files/E/Equifax-IR/Annual%20Reports/2017-proxy-statement.pdf. [Accessed 22 June 2020].

[38] The Georgia Department of banking and finance , "Equifax final consent order," 25 June 2018. [Online]. Available: https://dbf.georgia.gov/document/publication/equifax-final-consent-order-dated-6-25-2018/download. [Accessed 22 June 2020].

[39] Equifax, "Equifax 2020 proxy statement," 27 March 2020. [Online]. Available: https://investor.equifax.com/~/media/Files/E/Equifax-IR/Annual%20Reports/2020-proxy-statement.pdf. [Accessed 22 June 2020].

[40] Identity Theft Resource Center, "Data Breach Reports," Identity Theft Resource Center, 2019.

[41] Verizon, "2019 Payment Security Report," Verizon, 2019.

[42] The Federal Trade Commission, "Electronic Code of Federal Regulations. Part 314— Standards for safeguarding customer information.," 23 May 2002. [Online]. Available: https://www.ecfr.gov/cgi-bin/text-idx?c=ecfr&sid=1e9a81d52a0904d70a046d0675d613b0&rgn=div5&view=text&node=16%3A1.0.1.3.38&idno=16. [Accessed 5 August 2020].

[43] The Federal Trade Commission, "Section 5 of the FTC Act: principles of navigation," 18 October 2013. [Online]. Available: https://www.ftc.gov/system/files/documents/public_statements/section-5-ftc-act-principles-navigation/131018section5.pdf. [Accessed 5 August 2020].

[44] 111th Congress, "H.R.4173 - Dodd-Frank Wall Street Reform and Consumer Protection Act," 21 July 2010. [Online]. Available: https://www.congress.gov/bill/111th-congress/house-bill/4173/text. [Accessed 6 August 2020].

[45] U.S. Code, *12 U.S. Code § 5514. Supervision of nondepository covered persons. (b)(1),* 2010.

[46] The Consumer Data Industry Association, "Data security regulation of consumer reporting agencies," 2019. [Online]. Available: https://www.cdiaonline.org/resources/for-reporting-industry-professionals/data-security-regulation-of-consumer-reporting-agencies/. [Accessed 1 August 2020].

[47] The Federal Trade Commission , "Privacy and Data Security Update: 2019," December 2019. [Online]. Available: https://www.ftc.gov/system/files/documents/reports/privacy-data-security-update-2019/2019-privacy-data-security-report-508.pdf. [Accessed 1 August 2020].

[48] M. Clements, "Testimony Before the Subcommittee on Economic and Consumer Policy, Committee on Oversight and Reform, House of Representatives," 26 March 2019. [Online]. Available: https://www.congress.gov/116/meeting/house/109175/witnesses/HHRG-116-GO05-Wstate-ClementsM-20190326.pdf. [Accessed 1 August 2020].

[49] United States Government Accountability Office, *Actions Taken by Equifax and Federal Agencies in Response to the 2017 Breach*, Washington, DC: United States Government Accountability Office, 2018.

[50] Nancy G. Leveson, "CAST Analysis of the Shell Moerdijk Accident," Massachusetts Institute of Technology.

[51] J. Viega and G. McGraw, Building Secure Software - How to Avoid Security Problems the Right Way., Addison-Wesley, 2002.

[52] J. H. Saltzer and M. D. Schroeder, "The protection of information in computer systems.," in *Fourth ACM Symposium on Operating System Principles*, 1973.

[53] National Association of Corporate Directors, "Cyber-Risk Oversight," National Association of Corporate Directors, Washington DC, 2017.

[54] U.S. Government Accountability Office, "Actions Needed to Strengthen Oversight of Consumer Reporting Agencies," U.S. Government Accountability Office, Washington, DC, 2019.