Why the Equifax Breach Should Not Have Mattered

Marten Lohstroh

Department of Electrical Engineering and Computer Science
University of California, Berkeley
Email: marten@berkeley.edu

***Abstract*—Data security, which is concerned with the prevention of unauthorized access to computers, databases, and websites, helps protect digital privacy and ensure data integrity. It is extremely difficult, however, to make security watertight, and security breaches are not uncommon. The consequences of stolen credentials go well beyond the leakage of other types of information because they can further compromise other systems. This paper criticizes the practice of using clear-text identity attributes, such as Social Security or driver's license numbers—which are in principle not even secret—as acceptable authentication tokens or assertions of ownership, and proposes a simple protocol that straightforwardly applies public-key cryptography to make identity claims verifiable, even when they are issued remotely via the Internet. This protocol has the potential of elevating the business practices of credit providers, rental agencies, and other service companies that have hitherto exposed consumers to the risk of identity theft, to where identity theft becomes virtually impossible.**

***Index Terms*—Computer Security, Authentication, Identity management systems, Technology social factors, privacy.**

I. INTRODUCTION

A recent data breach at consumer credit reporting agency Equifax has compromised sensitive personal information such as Social Security numbers, birth dates, and driver's license numbers of as many as 143 million US consumers and up to 44 million British residents [16]. Consumers are outraged and concerned about the looming threat of falling victim to identity theft. Dozens of people whose personal data were leaked have filed lawsuits against Equifax. Politicians decry the lack of control that consumers have over their information as it is collected and sold by companies. Senator Warren and others were the first to respond with legislation they introduced as the Freedom from Equifax Exploitation (FREE) Act [29], which would let consumers "freeze" and "unfreeze" their accounts at no cost. In the meantime, the Federal Trade Commission (FTC) has started an investigation into the hack, scrutinizing the company's data security practices for possible neglect.

While all these responses address legitimate concerns regarding the business practices of credit reporting agencies and emphasize the necessity of companies taking responsibility for the protection of personal information that is stored on their servers, one very important question remains unasked: Why, in this day and age, do we allow plain-text, copyable strings of characters to serve as a means of authentication?

This paper provides answers to this question, and it suggests a path forward by providing a practical solution towards improving the state of affairs. In particular, it shows how well-established cryptographic algorithms can be leveraged to drastically improve the verifiability of identity claims. The result is a secure and verifiable format for the storage and exchange of identity claims, which, once it establishes itself as an industry standard, would render plain-text personal information useless for the purpose of committing identity fraud. Had this kind of technology been deployed, the Equifax breach would have still been problematic from a privacy standpoint, but it would not have subjected millions of consumers to a substantial risk of falling victim to identity fraud.

II. BACKGROUND

The problem of identity theft is multifaceted, and it can be studied through many different lenses. While the technical execution of an attack can be explained in terms of computer science, engineering, or even psychology (social engineering [14]), the phenomenon also can be treated as a societal problem, and its incidence be explained in the frameworks of law and economics. FTC consumer complaint statistics indicate that identity theft has been steadily on the rise throughout the 2000s [12]. Two very different explanations for this have gained traction in the legal academic literature. On the one hand, Lopucki [20] argues that the decline of public life and the gradual removal of contact information from public registers has lead to an inability of businesses to authenticate clients. Solove [26], on the other hand, attributes the problem to a lack of control that consumers have over their personal information. This lack of control increases the likelihood that one's personal information ends up in the hands of criminals.

The thesis of this article is that neither Lopucki nor Solove is right. Of course, impersonation is easier when it can be done over the phone or from behind a computer screen, but publicly available personal information—precisely because it is *publicly* available—can be no useful aid in identifying fraudsters. The lack of control over personal information, however, *is* problematic, but particularly because it is violating of privacy, not because it necessarily weakens security. Besides, given

This work was supported in part by the TerraSwarm Research Center, one of six centers supported by the STARnet phase of the Focus Center Research Program (FCRP) a Semiconductor Research Corporation program sponsored by MARCO and DARPA.

the mega security breaches that we have seen over the recent years (Home Depot, Target, and JPMorgan have had high-profile security breaches, all of which occurred fairly recently, in 2014 [28]), so much data has already been leaked that there is little left to control. The solution to identity theft must therefore not be sought in measures that limit or increase *access* to personal data, but in a mechanism that protects against the improper *usage* of personal data.

III. Incentives and Risk

Commercial businesses optimize their activities to maximize net profit. As such, fraud is a risk like any other that needs to be managed. One way to manage risk is to invest in measures that reduce the risk, another is to insure against damages. Companies will make a trade-off between these strategies that they expect is most profitable [5], [2]. Preventative measures can reduce insurance expenses and foster consumer trust (which potentially leads to higher revenue), but the cost of implementing preventative measures may outweigh the reduction in insurance cost, in which case the question remains: How pertinent are those measures to the customer experience? If a risk poses no threat to loss of revenue and it is more effective to insure against damages, a company has no incentive to allocate budget for prevention. The only other factors that will tip the scale towards prevention are laws, rules, and regulations which, if not abided by, could hurt business via lawsuits, license revocations, or fines.

Credit card companies are notorious for the poor security of their payment methods: particularly in the US, transactions normally do not require the customer to enter a PIN-code, and in-person transactions typically require no identification; sometimes not even a signature. Purchases made through the Internet often do not involve any authentication at all; it is typically sufficient to enter the card holder name and billing address, along with the credit card number and CCV. It is easy to explain why credit card companies continue to use such primitive technology: it is more profitable to do so. Sophisticated security measures are costly to implement, but, perhaps more importantly, they potentially make it more complicated (or impossible) to make purchases, particularly online, via phone, or in remote areas where direct authentication of a transaction may not be feasible. Hence, improved security, aside from the cost of implementing it, could lead to a reduction in revenue. Apparently, these costs are expected to outweigh the cost of credit card fraud. Instead of improving security to prevent fraud, credit card companies have shifted their focus to just-in-time fraud *detection*, leveraging statistical models (e.g., [27]) and machine learning techniques (e.g., [21]) to identify suspicious transactions in order to block them or verify their legitimacy with the customer. This turns out to be an adequate solution for the merchant, the consumer, and credit card issuer; the financial damage of credit card fraud is mitigated with minimal impact on the customer experience, and the financial liability is absorbed by the card issuer or the merchant.

As shown in the credit card example, when there exists a direct relationship between a business and its customer, damage control may indeed be a more pragmatic answer to fraud than it is to try to prevent it. The problem is: there might not be such relationship. In the case of Equifax, other businesses, not consumers, are the company's customers. It would be more accurate to say that consumers (or their data) are Equifax's product. Hence, Equifax has little incentive to align its interest with the interests of consumers. On the contrary, safeguarding consumer data is a costly undertaking. If companies like Equifax were more invested in protecting consumer data breaches may be less likely to occur, but impenetrable security is very difficult (if not impossible) to achieve. And when a breach does occur, where a credit card company can simply issue new credit cards and block and revoke compromised ones, data brokers and credit reporting agencies have no such power with respect to the data they handle. This leaves the burden of having to deal with the aftermath of a potential abuse of any stolen information to rest entirely on the affected consumers, even though they never entrusted the company that leaked their information with storing their personal data to begin with.

An example of such an aftermath is the more insidious variant of a credit card fraud that entails the illegitimate creation of a new account, which can take months before it is discovered. Again, it is a misalignment of incentives between businesses and their customers that ends up facilitating identity theft. In [17], Hoofnagle exposes just how embarrassingly easy it is for impostors to obtain credit or medical services. The paper discusses sixteen cases of identity fraud, and in almost all of these cases credit or services were granted on the basis of applications that were rife with errors that should have suggested fraud. Yet, because most of the cost of identity theft is externalized, and businesses make a trade off that optimizes their own profit, they choose to issue credit or provide services even in marginal situations. Companies do not want more rigorous screening because it will cost them revenue. This textbook example of externalizing cost will not disappear unless careless application screening are to become penalized by law, or the injury that results of negligent screening is to be considered a legal basis for a tort claim.

Of course, Equifax is now in a world of legal trouble as it faces lawsuits that are seeking class action status [1]. The legal battle that shall unfold in the years to come may lead the industry to re-evaluate its priorities with respect to the protection of consumer data, but that will only address part of the problem. The remainder of this article discusses a technical solution to the data breach problem that, instead of focusing on improvement of data confidentiality, aims to diminish the practical value of plain-text personal information by providing a reliable method for proofs of identity. It is clear, however, that the industry cannot be expected to take the lead in adopting this technology unless a change is incentivised or enforced by law.

IV. Certified Identity Claims

The purpose of a Certified Identity Claim (CIC) is for a relying party (RP) e.g., an airline, to obtain from a subject (S), e.g., a traveler, an attestation of the relationship between S and some set of attributes (e.g., last name, date of birth, and passport number)—information similar to that which was compromised by the Equifax breach. The use of CICs is a solution to the problem of identity theft, which entails the fraudulent use of attributes that relate to *another* individual than the person using them, for example, to gain unwarranted access to resources. It is the ability to verify the relationship between an identity claim and the entity that issues it that makes identity theft virtually impossible.

Sidebar: Encryption and Digital Signatures

Public-key cryptography algorithms use a separate key for encryption and decryption. A key pair consists of a *private* and a *public* key. The private key must be kept secret by its user, while the public key is safe to share with others. Once something is encrypted with a private key, it can only be decrypted with the corresponding public key, and vice versa.

Example: Alice (A) and Bob (B) have the following key pairs, respectively: $\{A_{pub}, A_{prv}\}$ and $\{B_{pub}, B_{prv}\}$. A and B can now securely exchange messages between each other by encrypting them with each others' public key. A prepares the secret message for B as follows: $B_{pub}(S)$. B (and only B) can read S by decrypting the message: $B_{prv}(B_{pub}(S))$. The same machinery can be used to issue signatures, simply by swapping the order in which the keys are used. B can apply a signature to a contract C as follows: $B_{prv}(C)$, and anyone, including A, can read it using: $B_{pub}(B_{prv}(C))$, and be assured that the message truly originates from B. If it were signed with another key than B_{prv}, the contents of C would not be readable.

The certification of identity attributes is carried out by a trusted third party, referred to as the attribute authority (AA). In this system, identity theft would require an attacker to either compromise the security between S and AA or steal the private key of AA. Both are very difficult to achieve, and both are reparable, respectively by resetting S's authentication credentials or by revocation of AA's certificate.

A. Verification of CICs

The use of a CIC involves a particular sequence of actions and message exchanges. Let us examine the message sequence that is presumably the most common, where RP requests S to prove ownership of some attribute, and S obliges by requesting AA to provide a CIC, which S finally relays to RP. This sequence, illustrated in Fig. 1, would substitute for requests for *uncertified* identity claims which, on the Web, are ordinarily solicited via a form that is rendered in the browser and filled out manually with keyboard input from the user. An important difference is that while an HTML form provides a description that is understandable by humans, the request sent from RP to S has to be passed along to a third party, AA, which, after authenticating S, will have to parse the request and compile an appropriate answer. All of the latter should be carried out automatically, so the format of the request must be machine readable. Importantly, the request must *also* be rendered in a human-readable form to enable the user operating S to determine whether it would indeed like to grant the request and forward it to AA, or deny it and cancel the exchange.

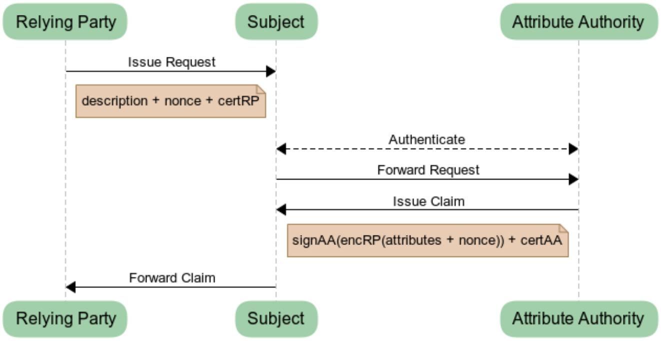

Figure 1. A UML sequence diagram that describes issuance of a CIC.

If we examine Fig. 1 a bit closer, we see that the request consists of three parts; a description, a nonce, and a certificate. The description encodes what particular information is requested from S. The nonce is a randomly generated bit string that is only used once, never across subsequent requests. The nonce is required to prevent replay attacks. Because AA bundles the requested attributes with the nonce that is provided by RP, it is evident that the certification must have occurred in response to RP's request. Without the nonce, S could reuse a CIC that it had previously obtained from AA, even if AA would no longer certify the claim at the present time (e.g., because S no longer has an active account).

After authenticating S, AA obtains the request and has to formulate a valid response. For instance, if RP is a lender and S is a potential borrower, the request may look like `[name, credit_score]`, and the matching attributes would look something like `['John Davis', 589]`. The next step is to encrypt the attributes together with the nonce using the public key of RP. The resulting cipher text is then signed by AA (using its private key). The extra round of encryption that is applied, before AA does its signing, is critical because it prevents RP from being able to stage a man-in-the-middle attack. Without it, RP would be able to retrieve a request from *another* relying party, say X, forward it to S (along with the same nonce), obtain a CIC from S and submit it to X. Neither X or S would be able to detect the attack. Because AA encrypts the attributes with the public key of RP, X will not be able to read them. It is important that the signature is applied after the encryption and not the other way around. If the keys were to be used in reversed order, RP would be able to decrypt the signed attributes and encrypt them with X's public key before forwarding them to S. Alternatively, if RP were to present X's certificate to S instead of its own, the attack would be detectable by S.

B. Trust

Before RP can verify a claim, it first has to decide whether AA is authoritative to assert the veracity of the claim. The authenticity of the certification provided by AA may be derived from the certificate that RP receives from S. This certificate ties the "common name" of the AA (e.g., bankofamerica.com) to a public key, attested by a chain of trust that should lead back to an intermediary or root CA that RP also trusts. The most trivial way for RP to keep track of relationships between AAs and the types of claims they are trusted to certify, is to maintain a white list. Simply put, for each claims request that RP makes, it will need to *a priori* decide whose certifications it will accept as material. For example, for a claim that establishes holdership of a bank account, all well-established banks could be acceptable AAs. Another method would be for RP to request a claim from AA and expect it to be certified by the American Bankers Association. What this mechanism basically gives rise to is a dynamic variant of a Public-Key Infrastructure (PKI) that is guided by ad-hoc inquisition rather than deliberate top-down one-off design.

C. Privacy

CICs can be viewed as a privacy-enhancing technology. Apart from preventing misuse of personal information, the use of CICs obviates the need for the collection of personal information beyond what is strictly necessary. Because identity claims are independently verifiable, there is no need to rely on cross referencing with other information in order to gain confidence in the authenticity of the claims. Better yet, a relying party will be less inclined to rely on omni-directional identifiers [7] (i.e., information that uniquely identifies an entity across multiple contexts, such as a Social Security Number) if it can increase confidence in the veracity of the identity claims it actually cares about. For instance, assume that a service provider needs to determine eligibility on the basis of an applicant's age, taxable income, and marital status. Instead of requesting copies of the applicant's passport, tax return, and marriage certificate, it could simply request CICs that attest the required information, without even revealing as much as the applicant's name.

The decoupling between RP and AA is deliberate. The fact that all communication between them goes through S puts the user in maximal control of the information flow. In this way, the user controls when and how often a third party can access their personal information. Of course, modulo legal provisions in a privacy statement, a user has no control over what a third party does with the personal information it has already acquired. However, particularly for transient claims (e.g., account balance), it is useful that the protocol gives S the power to throttle how often RP receives updates. Of course, the decisions of S may well be performed automatically on the basis of some predefined policy rather than having a human on the loop. The avoidance of direct communication between RP and AA could potentially also prevent AA from tracking the activities of S, but the protocol in Fig. 1 reveals RP's identify to AA by sharing its public key.

Although a relying party is capable of storing the contents a CIC in clear text, the sequence of cryptographic operations on the information guarantees that RP cannot extend the ability to verify the CIC to a third party without also handing over its private key, which would compromise RP's own security. To inspect the CIC, it must first be decrypted with AA's public key and then once again with RP's private key. Without the latter step, the contents of the CIC is unreadable, and if the claim is shared in clear text it is no longer verifiable.

D. Vulnerabilities

The security of CICs is predicated on some basic assumptions which, albeit reasonable, under certain circumstance could be broken. Let us discuss them one by one.

1) Compromised keys: It is possible that the private key of RP or AA gets stolen. This would impair the ability of S to authenticate them, thus making sessions susceptible to man-in-the-middle attacks. When AA has its private key stolen, the trustworthiness of its signature is also compromised.

2) Faulty or dishonest parties: The cryptography prevents attackers from tampering with claims, but it has no way to prevent false information from ending up in a claim. Claims are trusted on the basis of authority. In principle, AA could certify a false claim, or it could impersonate S. However, this sort of malicious behavior would violate trust and would quickly lead to the demise of AA.

3) Broken random number generators: It is important that a sufficiently random nonce is produced by RP. If it fails to do so, S may be able to successfully "replay" a previously issued CIC.

4) Quantum computers: The strength of public-key cryptography hinges on the computational intractability of mathematical problems like prime factorization of large integers. Quantum computers hold the promise of performing factorization in exponentially less time than classical or stochastic methods could [10]. Thus far, quantum computers are not that capable, but there are reasonable expectations that quantum supremacy is achievable in the not-too-distant future [6].

E. Adoptability

The proposed protocol can be implemented using off-the-shelf technology and existing web standards (RSA [25] for the signing and verification of claims, X.509 [18] for the certificates that authenticate the signatures, HTTPS [24] for end-point verification and secure communication), and would require only a modest standardization effort. Specifically, the syntax of requests and claims needs to be agreed upon, and there must be a standardized conceptualization of the attributes that can be featured in them: all parties involved will need to understand the contents of a CIC. SAML [19] is an XML-based standard that is already equipped with the ability to encode assertions regarding identity attributes, but the idea of certified identity claims could be extended beyond the assertion of simple attributes. A more elaborate protocol could potentially leverage the W3C Web Ontology Language (OWL) [4] and allow for the formulation of more expressive assertions that involve relationships between multiple entities and certifications by multiple authorities.

An important feature of CICs is that they do not require session-oriented protocols; CICs are self-contained and are tied to a particular request by means of a nonce and a particular RP via encryption. This means that the exchange of CICs can be conducted entirely in a RESTful [11] fashion. Secure connections are assumed between S and RP and S and AA. If these connections are not secure, a man-in-the-middle attack could be staged by an attacker who is interested in letting the victim unwittingly submit claims that are not theirs but the attacker's. For instance, an attacker could hijack an online purchase and reroute the delivery of the purchased item by injecting a false shipping address. An obvious way to secure the message exchange is to use HTTPS, which could also facilitate authentication of RP and AA with respect to S. The methods of authentication of S by RP and AA, respectively, are intentionally left unspecified. Where one organization may want rely on token-based authentication, another may want to use key authentication, a password, biometrics, or some sort of two-factor authentication. We can assume that the level security implemented by RP and AA is proportional to the sensitivity of the kinds of activities that it protects (e.g., money transfers, access to medical records, etc.).

V. Related Work

Public-key cryptography has been widely adopted as a means to authenticate Web pages (X.509 certificates) and encrypt Web traffic (TLS/SSL). CICs and X.509 certificates are similar, but where a X.509 certificate is designed specifically to tie a public key to a hostname, organization, or individual, a CIC may bear any sort of identity claim. Another difference is that an X.509 certificate is signed by a Certificate Authority (CA), whereas a CIC is signed by an Attribute Authority. This distinction is significant, because the trustworthiness of a CA is determined by its relationship with other CAs, whereas the trustworthiness of an AA is determined with respect to a particular set of attributes, and stands or falls by the relationship that the AA has with the subject. For example, the non-profit organization Let's Encrypt[1] provides free SSL/TLS certificates for any organization or individual that can demonstrate that it controls the domain of the hostname it wishes to tie their public key to. Verification of control is simple and can be automated; the proof requires no more than the creation of a provisioned DNS record and HTTP-accessible resource. But how would just any AA verify whether an individual has any outstanding traffic tickets, for instance? Indeed, it requires an authority with access to very specific information about the subject in order to veritably certify such identity claim. To vouch for the absence of outstanding traffic tickets, the Department of Motor Vehicles would be an bona fide authority, but the American Automobile Association would not.

The concept of claims-based identity was pioneered by Cameron [7] at Microsoft in the mid-2000s when he laid out seven "laws" of identity, which provide useful guidelines for the design of systems that cope with identity on the Internet. Cameron proposes a definition of identity in terms of claims rather than identifiers. With this definition, he shifts the paradigm from identifying individuals to proving the relationship between digital identity and real-world objects. By choosing the term "claim" instead of "assertion," he emphasizes that this relationship is always imbued with some level of uncertainty, and could manifest itself as a potential weakness in any secure system. The protocol described in this paper implements the first four of Cameron's laws, which boil down to: "user consent," "minimal disclosure," "justifiable use," and "context isolation." The fifth law states that an identity system needs to allow for pluralism of operators and technologies, which the protocol also does, granted that it leaves the method for end-point authentication unspecified. The sixth and seventh laws are strictly concerned with user experience, which is entirely outside of the scope of this paper.

The notion of *federated* identity captures the fact that identity attributes are often stored across multiple distinct identity management systems. Federated identity management (FIM) [8] entails the exchange of identity information between identity management systems. A subset of FIM is concerned with single sign-on (SSO), in which a user's single authentication ticket, or token, is trusted across multiple systems or even organizations. Open standards for FIM are WS-Security and its extension WS-Trust [3]. The goal of FIM is somewhat different from CIC. Where FIM is designed to facilitate long-term collaboration between businesses to provide services to an overlapping customer base, CIC facilitates exchanges between entirely unrelated stakeholders on a per-request basis. Where FIM is aimed at providing tight integration between federations, CIC provides a loose coupling. Popular SSO protocols are OpenID [23] and OAuth [15]. Online platforms like Google[2] and Facebook[3] also allow third parties to rely on their user authentication mechanism and provide limited access to their users' account details.

Governments have also attempted to design systems for identity management and authentication. Some have become quite successful nationwide identity systems, and are leveraged for purposes as proof of age, proof of citizenship, and for generating digital signatures. These government-run systems, however, tend to poorly interact with commercially-run systems. The National Strategy for Trusted Identities in Cyberspace (NSTIC) that President Obama signed in 2011, sought to change that by creating an "Identity Ecosystem" for the improvement of online transactions [22]. The program went through dozens of pilots, but the idea never materialized

[1] https://letsencrypt.org

[2] https://developers.google.com/identity/

[3] https://developers.facebook.com/docs/facebook-login/

in any way, shape, or form that was ready for adoption. Large and complex identity systems come with many risks and challenges [9] such as over-centralization, lack of interoperability, possible privacy violations, etc. Basically, the more problems a system attempts to solve, the more challenges it presents to widespread adoption. Highly successful and widely-adopted security mechanisms (e.g., HTTPS) tend to be simple, only address a single issue, and tend not to have too many dependencies on infrastructure.

VI. Conclusions

The problem of identity theft is a solvable problem; the issue was not prevalent prior to the advent of Web technology and e-commerce, and its current heavy manifestation indicates a problem with the technology we rely on. A lack of sophistication in the way we verify personal information is to blame. While some suggest that "data is the new currency," [13], the way we pass around personal information is the monetary equivalent of exchanging plain-paper banknotes with handwritten denominations on them. This primitive state of affairs is not a reason to scrutinize the security of sites that retain our personal information in order to prevent theft; it is an argument for finding ways to verify the *authenticity* of personal information—similar to how security features on paper money help distinguish genuine bank notes from counterfeit ones.

The CIC protocol presented in this paper provides a simple solution to the inexcusable lack of verifiability of identity claims that leads to the tens of millions of cases of identity theft and tens of billions of dollars of fraud damages that are suffered, year after year. The protocol can be understood by anyone with basic knowledge of public-key cryptography, it is inherently distributed, it is fully interoperable with any existing authentication infrastructure, and it could be seamlessly integrated into the online user experience as a substitute of ordinary HTML forms.

Once the use of CICs becomes widespread, the use of cleartext identity attributes can be abandoned altogether, which would render security breaches, like the one that happened to Equifax, considerably less detrimental. However, it will be up to governments or regulators like the Federal Trade Commission to impose restrictions on the acceptance of unverifiable user-provided electronic data, particularly for the purpose of entering legally binding agreements, because most businesses have very little incentive to impose such restrictions themselves.

Acknowledgments

I would like to thank Prof. Chris Hoofnagle for his encouragement and critical feedback on this work, and for the fruitful discussions I had with him about identity fraud, privacy law, and consumer protections. I am also thankful to Antonio Iannopollo for his helpful suggestions that greatly enhanced the readability of this paper.

References

[1] After the breach, Equifax now faces the lawsuits. http://www.chicagotribune.com/business/ct-equifax-data-breach-lawsuits-20170922-story.html. (Accessed on 11/26/2017).

[2] R. Anderson and T. Moore. The Economics of Information Security. *Science*, 314(5799):610–613, 2006.

[3] S. Anderson, J. Bohren, T. Boubez, M. Chanliau, G. Della-Libera, B. Dixon, P. Garg, M. Gudgin, P. Hallam-Baker, M. Hondo, et al. Web Services Trust Language (WS-Trust), 2004.

[4] S. Bechhofer. OWL: Web Ontology Language. In *Encyclopedia of Database Systems*, pages 2008–2009. Springer, 2009.

[5] R. Böhme. Security Metrics and Security Investment Models. In *IWSEC*, pages 10–24. Springer, 2010.

[6] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, J. M. Martinis, and H. Neven. Characterizing Quantum Supremacy in Near-term Devices. *arXiv preprint arXiv:1608.00263*, 2016.

[7] K. Cameron. The Laws of Identity, May 2005. *Microsoft Corporation*, 2005.

[8] D. Chadwick. Federated Identity Management. *Foundations of Security Analysis and Design V*, pages 96–120, 2009.

[9] N. R. Council. *IDs – Not That Easy: Questions About Nationwide Identity Systems*. The National Academies Press, Washington, DC, 2002.

[10] D. Deutsch and R. Jozsa. Rapid Solution of Problems by Quantum Computation. In *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, volume 439, pages 553–558. The Royal Society, 1992.

[11] R. T. Fielding and R. N. Taylor. *Architectural Styles and the Design of Network-based Software Architectures*. University of California, Irvine Doctoral dissertation, 2000.

[12] K. M. Finklea. *Identity Theft: Trends and Issues*. DIANE Publishing, 2010.

[13] C. Gates and P. Matthews. Data is the new currency. In *Proceedings of the 2014 workshop on New Security Paradigms Workshop*, pages 105–116. ACM, 2014.

[14] S. Granger. Social Engineering Fundamentals, Part I: Hacker Tactics. *Security Focus, December*, 18, 2001.

[15] D. Hardt. The OAuth 2.0 Authorization Framework. 2012.

[16] A. Hern. Equifax told to inform Britons whether they are at risk after data breach, September 2017. Retrieved 2017-09-11.

[17] C. J. Hoofnagle. Internalizing Identity Theft. *UCLA Journal of Law & Technology*, 2009.

[18] R. Housley, W. Ford, W. Polk, and D. Solo. Internet X. 509 Public Key Infrastructure Certificate and CRL profile. Technical report, 1998.

[19] J. Hughes and E. Maler. Security Assertion Markup Language (SAML) v2.0 Technical Overview. *OASIS SSTC Working Draft sstc-saml-tech-overview-2.0-draft-08*, pages 29–38, 2005.

[20] L. M. LoPucki. Did Privacy Cause Identity Theft? *The Hastings Law Journal*, 54:1277, 2002.

[21] S. Maes, K. Tuyls, B. Vanschoenwinkel, and B. Manderick. Credit Card Fraud Detection Using Bayesian and Neural Networks. In *Proceedings of the 1st International NAISO Congress on Neuro Fuzzy Technologies*, pages 261–270, 2002.

[22] K. N. Megas, P. Lam, E. M. Nadeau, and C. Soutar. NSTIC Pilots: Catalyzing the Identity Ecosystem. *NIST Interagency/Internal Report (NISTIR)-8054*, 2015.

[23] D. Recordon and D. Reed. OpenID 2.0: a Platform for User-centric Identity Management. In *Proceedings of the Second ACM Workshop on Digital Identity Management*, pages 11–16. ACM, 2006.

[24] E. Rescorla. HTTP over TLS. RFC 2818, 2000.

[25] R. L. Rivest, A. Shamir, and L. Adleman. A Method for Obtaining Digital Signatures and Public-key Cryptosystems. *Communications of the ACM*, 21(2):120–126, 1978.

[26] D. J. Solove. Identity Theft, Privacy, and the Architecture of Vulnerability. *The Hastings Law Journal*, 54:1227, 2002.

[27] A. Srivastava, A. Kundu, S. Sural, and A. Majumdar. Credit Card Fraud Detection Using Hidden Markov Model. *IEEE Transactions on Dependable and Secure Computing*, 5(1):37–48, Jan 2008.

[28] R. Walters. Cyber Attacks on US Companies in 2014. *The Heritage Foundation*, 4289:1–5, 2014.

[29] E. Warren et al. Freedom from Equifax Exploitation Act, 115th Congress, September 2017.