Making IoT Worthy of Human Trust

Hilda Hadan*
hhadan@iu.edu
Nicolas Serrano*
nicserra@iu.edu
Sanchari Das*
sancdas@indiana.edu
L Jean Camp*
ljcamp@indiana.edu
* School of Informatics, Computing, & Engineering
Indiana University Bloomington

***Abstract*— The Public Key Infrastructure (PKI) is the foundation which enables secure and trusted transactions across the Internet. PKI is subject to both continuous attacks and regular improvements; for example, advances in cryptography have led to rejections of previously trusted algorithms (i.e., SHA1, MD5). Yet there have also been organizational failures and malicious acts by trusted parties. In this work, we focus on the socio-technical components of the current X.509 PKI with the goals of better understanding its vulnerabilities, and ideally informing the implementation of future PKIs. We begin with a taxonomy of failure modes encompassing chronic, catastrophic, high impact, and frequent PKI failures. This categorization was informed by a survey of non-expert perceptions of PKI and an interdisciplinary workshop addressing the future of security in the Internet of Things. To evaluate the failure modes, we conducted qualitative interviews with policy scholars and experts in applied cryptography. We summarize the results of the survey and workshop and detail the expert interviews. Our findings indicate that there are significant failure types which neither the technical nor policy community can deeply engage separately. The underlying assumptions about the rate and severity of failure differ between these communities. Yet there is a common awareness of the vulnerabilities of the end-users: the people who are required to trust PKI to interact and engage with the Internet. We identify an urgency in mitigating such critical issues, in part because of the increasing adoption of cyberphysical systems and the Internet of Things (IoT). We conclude that there is a need for integrated organizational, policy, and technical coordination to address the chronic and potentially catastrophic risks. We introduce possible economic or regulatory solutions and highlight the key takeaways.[0]**

I. INTRODUCTION

Public Key Infrastructure (PKI) is comprised of issuers of public key certificates (Certificate Authorities) [42], software manufacturers who determine which issuers will be trusted by default (root program operators) [57], and the software developers [51] who implement the code that interprets the resulting cryptographic instances [16]. The key management and certificate management is done by in PKI which should enable an organization to maintain a trustworthy network and attest to its identity [24]. The final decision to trust a certificate or not appears to be with the clients on the edge; interacting with a website; however, simply clicking on a site with a standard browser entails trusting the code that is loaded.

This trust model is similar to other infrastructures of trust, in automobiles, food, and medical supplies where the attestation occurs before the final consumer engages with the product. However, in these markets, people are not expected to identify fraudulent foodstuffs, poisonous medications, or dangerous vehicles without expert guidance. Phishing [17], [18], pharming [49], drive-by downloads [23], man-in-the-middle-attack [41], and even control plane attacks [14], [25], are not considered as threats in the domain of consumer responsibility in these infrastructures.[1]

With the increasing usage of Internet in our daily life, massive expansion of the Internet of Things (IoT) devices in the home [52], and the increasing integration of Internet connectivity in real-time control of physical infrastructures [19], it is worthwhile to reconsider the PKI governance model. Based on our research, we found *buyer beware* to be a consensus model of the state of PKI. We use a taxonomy of failure modes in PKI to begin to ask if all of these failures are appropriately situated in the consumer domain. We ground this taxonomy in examples of past failures and evaluated it with a series of interviews addressed to experts in the field. These initial pilot interviews are designed both to improve the interview instrument and to elicit any failure mode missing from the taxonomy.

We identify and illuminate threats and weaknesses as well as addressing risk perceptions in PKI, particularly those related to IoT. We discuss how these may impinge the architecture, design, usage, and policy of IoT devices. The goal is not to define the final state of appropriate PKI governance but rather to identify concerns from multiple perspectives. Our goal is to define the risk from PKI by defining failure modes and addressing the trust issues which can lead to a false sense of security among end-users [46]. We include an initial investigation on the perceptions of risks in PKI from policymakers and technologists.

Our study reveals that the current PKI is grounded in an assumption that individuals are responsible for ensuring their own safety by careful examination of context and source. Our set of investigations also illustrates that even

[0]This initial draft should be referenced as: Hilda Hadan, Nicolas Serrano, Sanchari Das, L. Jean Camp. *Making IoT Worthy of Human Trust*. Telecommunications Policy Research Conference, 20-21 Sept. 2019; Washington, DC.

[1]The history of the regulation of food, medicine, and transit form an implicit backdrop of this discussion but are beyond the scope of the work.

on the web, non-technical people do not have the necessary understanding to implement that examination. Further, the failure modes identified include those beyond the ability of an end-user to detect or defend against. We argue that the next generation PKI should integrate a holistic understanding of human and organizational assumptions by design. Our goal is to contribute to better options for PKI in ways that protect privacy and address the contextual nature of trust, taking human behaviors and practices into account. We describe the previous research in this domain in section II and describe our methodology in section III where in the first section we describe the motivation, including exemplar PKI failures, leading to the taxonomy. We described how we used a survey to obtained end-users' understanding of the nature of certificate, how we developed our failure modes from a workshop, and how we used a set of pilot interviews to inquire about the perceptions of PKI by technical and policy elites and determine their areas of agreement and difference. The results from the survey, workshop, and interviews are shown in section IV, which helped us in discussing the current and future scope of such research while providing concluding remarks in section V.

II. RELATED WORK

Our motivation for this research is to contribute to a longer interdisciplinary research agenda toward PKI which lead us in answering the following questions:

- *RQ1. What are the individual, organizational, technical, and broader social/political factors that shape the frameworks of PKI design, adoption, and use?*
- *RQ2. What are the factors that translate to practices and expectations of PKI? What are the means to enhance and improve these PKI practises to address the current security needs?*
- *RQ3.How can we utilize design changes to impact the organizational and end-user needs guided by the trust models?*

These questions are of significant interest for the academicians, security practitioners, and policymakers. Understanding how these contexts interrelate and how different players in the PKI infrastructure perceive these interactions is critical to generate secure and functional cyber systems of the future. In our work, we evaluated the failure modes of PKI and how these might be addressed for IoT. PKI, in its current role essentially verifies domain names and enables key generation. From necessity, it is an infrastructure that policymakers assume to be secure. Yet the cryptographic attestations that, for example, phishing sites from *.RU* should be trusted indicate that this is a challenge beyond the technological.

A. Certificates

X.509 certificates are mathematically gracious and subtle, but they can fail in the larger context in which they are deployed. There are four ways that certificates typically fail:

First, the set of facts embedded in the signature is somehow incorrect, either because of changes over time or incorrect issuance. This stems from organizational failure. Second, the cryptography (including the digital signature or hash value) could be flawed. This is a mathematical failure, but the fact that broken algorithms continue to be used despite their vulnerability is arguably an organizational failure. Similarly, the fact that weak keys are chosen is also a result of management or organization failure; excluding those cases where weaknesses were unknown at the time of issuance. Third, the software that is supposed to confirm the authenticity of the certificate is subverted, and thus authenticates false and incorrect information. This is a failure of human developers who are trying to get a job done. Finally, individuals could perceive that the certificate means something quite different than the intended issuance and implications. This is a usability failure. The need for appropriate warnings for end-users is a usability challenge and is not the focus of our research.

In our past research of bank certificates, we found problems with the first and the second (incorrect facts and flawed cryptography) [22]. Other researchers have documented serious problems in terms of the third, i.e., bad evaluation software [11], [15], [34]. Failures resulting from flawed cryptography, including bad choices of keys, will continue to be identified as part of the natural development of research in mathematics and computer science. Yet when vulnerabilities are discovered, organizations may not take action. In a particularly illustrative example, the fragility of one core algorithm (SHA1) resulted in unanimous rejection of future certificates using the algorithm (again SHA1) by the governing CA/Browser Forum and the root program operators for any certificate signed after January 1, 2016.

In the larger marketplace, the result was not immediate compliance but rather a CA business opportunity for back-dated mathematically weak non-compliant certificates, specially, WoSign and StartCom quietly offered back-dated certificates rather than upgrading [30]. By September of 2017, the newest version of Chrome rejected all WoSign and StartCom certificates, simply removing them as a trusted authority. SSL Labs did the same. Microsoft instituted a softer landing, not trusting WoSign and SartCom certificates not previously validated before September 2016. (Most certificates issued from WoSign and SartCom lasted one year, so this was a slower but also certain removal.) That this mathematical advance was an institutional issue is further illustrated by the discovery by Mozilla that WoSign (a Chinese company) had purchased StartCom (an Israeli company) in April of 2016 through two Hong Kong subsidiaries, and in 2016 this was not known by StartCom customers.

Yet this is not simply a story of international business intrigue. Similarly in 2017, Symantec was found to be untrustworthy in part due to mis-issued SSL certificates [3], [4], [20], [21]. Symantec managed multiple widely-used certificate authorities including Thawte, VeriSign, Equifax, GeoTrust, and RapidSSL. Symantec was forced to sell the PKI business to DigiCert in part because they were simply too large a component of the infrastructure to remove; but

they were simultaneously untrustworthy agents for certificate issuance. This illustrates shows that a mathematical weakness– specifically advances in cryptography showing effective attacks on SHA1 – was not isolated from market behaviors.

B. Trust Model

The decision to trust a centralized authority, who can delegate trust to other authorities, was identified by as fundamentally flawed. Fellow of the ACM and Member of the National Academy, Martin Abadi noted early on that the concentration of trust "leads to many opportunities for incorrect trust decisions, which can be encouraged by various economic or political incentives" [6]. He proposed entirely local governance, with individuals and organizations selecting their own roots of trust through risk analysis and strategic decision-making.

Just as mathematical advances and resulting vulnerabilities have business implications, we understand software failures and incorrect information as partially human, organizational and economically driven failures, as well as technical flaws. Code that is theoretically examined by thousands of eyes, may still have a simple typo, meaning that all claims of a certificate will be interpreted as valid (e.g., the GOTO fail described by Wheeler, 2016 [58]; Langley, 2014 [37]). Specifically in that case, the OpenSSL library, perhaps the most used in the world, had a logical failure consisting of the repetition of the phrase "GOTO FAIL" at the end of the certification loop [53]. The compilation of redundant code resulted in the acceptance of any cryptographically valid certificate as having a legitimate attestation for any domain. That is, as long as the math was correct (i.e., this public key signature of the hash value) the domain of the website did not have to match the domain of the certificate. All valid certificates were judged as equally valid for any domain with this error.

This is the most widespread cause of straight-forward technical coding failures which are also a chronic problem with HTTPS. Yet these are not only straight forward technical issues but also human problems and organizational problems. From the human side, a study of the source of such failures found Stack Overflow and search results are a major source of insecure "solutions" to coding problems [7]. There is a strong social component to sharing of code and solutions; didactic perfection in standards will not work without integrating these human and organizational components [10]. Until 2013 there were no full-time employees of the OpenSSL foundation, a core infrastructure was managed by a few people working part-time with an unstable set of contributions, which is an organizational weakness. Currently, the infrastructure relies on contributions from people otherwise full-time, corporate, and foundation support.[2]

[2]https://www.openssl.org/blog/blog/2018/12/20/20years/

C. Usability Issues

To date, there has been limited documentation and analysis of how (and if) these socio-technical failures vary across domains and how to intervene to correct them. In one notable example, Park addressed the citizen-facing adoption of PKI in South Korean and concludes that the success was despite technical choices (e.g., ActiveX) not because of technical excellence [44]. Another researcher proposed a socio-technical solution to the challenges of TLS by building a model that embeds the existence of possible warnings and potential user responses [26]. As with the work of Abadi et al., philosopher Helen Nissenbaum and peers identified the model of trust in PKI as quite nearly perverse in human terms, creating models that are in direct opposition of the understanding of trust in human interactions and organizations [13]. Yet the need for an examination of large scale PKI governance have not been seriously engaged. More recent investigations has been subsumed by faith that a single source of all trust – Google's Certificate Transparency [31] – will address the challenges of excessive centralization previously identified by Nissenbaum, Abadi, and their co-authors.

III. METHODOLOGY

There were three components to this research. The following, we describe the method and timing of each component. We show how each informed the other. All research components described in this section were reviewed and approved by the university IRB.

A. Survey

The first component is a set of surveys with hundreds of non-technical people about the nature of certificates. The question was embedded in a larger survey and read, 'If you know, please describe what a "security certificate" is in the context of the Internet, otherwise write "Don't know."'. The initial goal of that investigation was to evaluate indicators of skills and knowledge; yet the answers were more varied that expected. That "Don't know" was the most popular answer is not surprising. As security professionals, we expected responses primarily addressing security and authentication but many concerned the rule of law and privacy.

The survey provided a quantitative picture of the understanding of the role of PKI on the web. The survey was broad and participants without expertise were particularly sought. The result is summarized in section IV-A. More detailed survey methodology is available in the analysis of the factors that impinged the understanding of phishing and certificates [45].

B. Workshop

The second was a Chatham House Rule workshop that addressed issues of failures in the PKI for IoT. That workshop *The Best Practice for Living in the Internet of Things* took place August third and fourth in 2017 in Seattle, Washington.

The goal of the workshop was to chart a path from the current state of IoT to a secure state. The workshop was

not a series of predetermined presentations, rather there was an introductory ground-setting keynote followed by intensive collaboration. Breakout groups discussed their own areas of expertise the first day (cryptography, usability, secure code, recovery) and then were placed in interdisciplinary groups the second day.

The results that most strongly influenced this work are summarized in section IV-B.

C. Interview

The final component was a set of interviews of highly expert populations. This overlaps with the two communities whose representatives, at the workshop, seemed to have the most substantive disagreements about what can be expected from the IoT end-user: policymakers and applied cryptographers. The policy proposal of disposing of all IoT devices upon failure or vulnerability was embraced by investors in start-ups and some policymakers. Cryptographers, recognizing that a mathematical advance could simultaneously result in vulnerabilities of entire classes of goods, rejected this. The core underlying disagreement appeared to be grounded in assumptions of frequency of failure and mode of failure. Thus we engaged in a set of pilot interviews with highly expert communities of technology: policy professionals and applied cryptographers.

The pilot interviews were held during the 46th TPRC conference and RWC 2019 symposium. The interview developed as a means of collecting opinions from experts to understand the perceptions of risk from a congress of the people most likely to have thought deeply about the role of the citizen or consumer in IoT.

The questionnaires focused on both technical and non-technical aspects, including but not limited to identification of potential failure modes of PKI and the role of domain names in IoT, as well as personal experiences with malicious websites and certificate warnings. We closed by asking for advice on the optimal path forward, towards a secure IoT. An additional aim of interviews was also to improve our questionnaires, the participants' opinions on the questionnaires were also asked at the end of each interview. As only minor changes were suggested by only 3 of the 12 participants, we used the same interview at Real World Crypto (RWC). We report our results in section IV-C.

We interviewed 9 technical researchers in information policy, 1 policy journalist, and 2 Ph.D. students from American University who were attending the TPRC 46th conference and 6 expert technologists who were attending the RWC 2019 symposium. All interviews were audio recorded. All the interview records were transcribed anonymously with only the day and a participant number of the recording stored.

IV. RESULTS AND DISCUSSION

Here we describe the results of each component of the research separately and also discuss the results as a whole.

A. Survey Findings

We implemented a survey to evaluate non-technical people's awareness of security and recruited a range of participants. Some participants in the survey were very optimistic about the scope of security and privacy indicated by a lock icon. Many answers indicated a belief in a form of governance very different than that provided by the CA/Browser Forum: an assumption of local legal accountability of the site. It is not possible to be certain without more research that these survey participants do not actually examine the details of certificates, trusting only those with issuers and subjects in their own jurisdictions. However, we are highly confident based on previous research and uniformity of personal experience that this is not the case. The codebook developed to quantify the qualitative data covered the answers with a high degree of certainty. The codes were as follows, with the most accurate describe on top.

Table I presents the list of codes that were used to bin the responses to security certificate questions arising from four coders. We expected a range of answers addressing privacy and security, yet multiple participants responded that X.509 certificates conveyed information about the legal accountability of the site. The geographic proximity of the issuer and the subject may imply the ability to seek redress. We cannot assert the likelihood that non-technical individuals actually viewed these certificate fields (or even know they existed). The second surprising result was the optimism for the scope and the function of a security certificate.

The survey was of 822 participants, ranging from 18 to over 70. The participants skewed younger, more educated, and the median annual income reported was between twenty and thirty thousand. The group was dominated by MTurk participants (696). After rejecting participants based on timing or response and attention check questions, there were 399 valid responses for our initial survey MTurk post, which was not associated with any other task. In addition, we embedded the question in a card-sorting experiment to evaluate if mental models of participants corresponded to perceptions of certificates. (There was a correlation between the correctness of answers and participants' mental models, please see [12]).

There were 294 MTurk participants in that second study, answering identical expertise and certificate questions before engaging in the task. There were 49 participants recruited from Mini University. Mini University is an event organized for active retirees by the IU Alumni Association for Lifelong Learning, so these participants were all over 65 and either IU alumni or spouses of IU alumni. We also recruited 23 participants from the Grace Hopper Celebration of Women in Computing. There were also 27 participants from the Bloomington Farmers' Market. Finally, we attended DashCon to access younger, Internet-active fandom communities and obtained 106 valid responses. DashCon was a one-time convention of Tumblr enthusiasts who skewed younger, less technical, and with very few males.

The total distribution of answers is shown in Table I. Our initial observations indicated that there were patterns within

Code	Count	Corresponding Qualitative Code
13	2	DNS: the certificate is associated with the domain name
12	70	Identity Verification: The certificate confirms that "I am who I say that I am" authentication
11	47	Cryptography: The certificate is association of information protection using cryptography, https
10	13	Information security: The certificate ensures a secure connection between the browser and the website
9	27	Website registration: When a website has to register or be certified and the certificate checks this certification/registration
8	34	Validation: The certificate states the site is valid, not a fake website
7	8	Information access by website: The certificate makes sure that only the website has access to the stored information
6	14	Protection: The certificate actively protects against attacks, including hackers/unauthorized people/virus, it indicates technical competence
5	3	Accountability: The certificate expresses that there is legal accountability, that there is a reliable underlying agreement between the user and website of accountability for information
4	55	Security/safety: The certificate says that the website is generally safe and secure
3	77	Trustworthiness of website: The website can be trusted to be benevolent (morally, legally, or ethically upstanding), not necessarily competent
2	49	Other
1	422	Don't know
0	73	N/A

TABLE I
QUALITATIVE CODES FOR CERTIFICATES ORDERED BY LEVEL OF ACCURACY.

the populations. For example, the second description on the use of encryption was dominated by MTurk participants who responded to the request for a survey. DashCon participants were far more likely to generate answers that were coded into categories 3 (trustworthiness), 7 (active prevention of attacks), and 12 (identification of the website).

Analysis of the additional questions – relating to computer experience, security knowledge, and expertise – found that the differences between groups were a function of expertise and not cohort nor gender. Details of the analysis can be found in the work by Rajivan, Kelley, Moriano, and Camp [45].

B. Workshop Findings

A discussion at the workshop identified the set of failure modes participants predicted for the IoT. These proposed failure modes are summarized in subsection IV-C.3. Another discussion focused on usability and concluded that individuals should be empowered to decide whom to trust in a transparent and usable manner. Together, these discussions inevitable and unavoidably lead to PKI.

The exploration in the workshop was focusing on defining the ground that must be covered to reach a secure IoT. The requirements included (i) an agile cryptographic infrastructure, (ii) building secure and robust code on that infrastructure, (iii) communicating to end-users about how to manage the resulting coded artifacts, and (iv) enabling recovery when failures occur (particularly attacks). These four fundamental challenges are echoes of the five ways in which a certificate validation can produce a false positive: bad design, bad math, bad code, failed revocation, and human errors.

The requirement for a robust cryptographic infrastructure requires agile cryptography and an understanding of the failure modes that may subvert the trust of cryptographic attestations.

C. Interview Findings

The 18 interview participants included three women and fifteen men; nine were academic researchers in information policy, one was a policy journalist, two were Ph.D. students from American University, and six were expert technologists with at least a decade of experience in applied cryptography in addition to doctorates in related fields.

1) Defining IoT: In order to ground our observations of vulnerabilities, our first goal was to determine if people in the discussions in these two communities had shared definitions. For example, when asked to define the IoT, Would one group answer the query thinking of automobiles, another thinking of phones, and a third focusing on industrial control systems? IoT can refer to a wide range of technologies. We found that there was a common understanding of the IoT domain. In both conferences, nearly all participants mentioned electronic devices and not systems, while many participants responded that IoT devices have the ability to connect and communicate through the Internet. Technical participants were more explicitly focused on the home; both groups assumed that living spaces are impinged by the IoT.

Conversely, three participants believed that 'IoT' is a nonsense term because everything is made up of 'Things'. One such participant was a technology journalist, one a cryptographer, and the other a policy professional. They expressed the opinion that the Internet of Things could also be the Internet, and the Internet could also be the Internet of Things.

The usage of IoT devices was most commonly associated with data collection, and less often with actuating on the physical environment. With this in mind, it is perhaps unsurprising that a number of policymakers in the TPRC conference highlighted 'privacy' protection as the primary challenge associated with deploying and managing IoT devices. In addition to concerns about privacy, the functionality of the device and the implications for data handling were central among RWC participants. For example, there was a focus on voice-activated home assistants in RWC, in part because of the requirement for constant audio surveillance. The difference was that TPRC participants focused on the resulting data flow and risks; while RWC attendees considered the capture and retention of data.

The most commonly cited pain points for all participants are the low 'level of trust' of the network and insecure 'data handling' methods in IoT devices. Additional concerns include the consequence of connecting to the network

and the resulting types of data collected by device and shared through the network. In addition, technologists also expressed concerns about the trustworthiness of the device manufacturer, the security of the device itself, the security of device updates, and the trustworthiness of certificates.

All participants agreed on the importance of transparency. Currently, IoT devices do not provide necessary security nor privacy information. Without considerable improvements in transparency, people will be unable to know to what sites their devices are connecting and what data are compiled. Without considerable improvements in documentation and isolation, the manufacturers and network service providers may be similarly challenged in trying to determine what sites the devices should be connecting to in normal operation.

A common goal of providing trust is that owners should have an awareness of the type, timing, use, and management of data compiled about them. Yet making digital footprints visible on the web is an unsolved problem; this is exacerbated with mobile devices, and IoT devices may share more data yet provide no interface to the consumers. There was a strong agreement between participants that there is a role for cryptography in that people should be able to verify what data is being collected and should have the right to opt-in, making a purposeful decision to share data.

2) Failures of Public Key Certificates: Given the intimacy, ubiquity, and diffusion of IoT devices, there is a requirement for a fully functional and trustworthy model of trust. PKI is a component of creating that trust. In practice, this has meant stapling on the X.509 model used on the web. Yet nearly all participants agreed that the current model of public key certificates is marginally adequate for the web, and not adequate in the IoT. Policymakers specifically highlighted "key management" and "untrusted CAs" as existing issues. Beyond that, technologists also indicated the lack of a centralized certificate store in an IoT environment. They were also concerned about the trustworthiness of service providers and operators. One technologist said, "even when I'm connecting to a legitimate website on a legitimate network, I still worry about whether the people operating the website will do the right thing with my data."

3) Failure Modes: The failure modes of PKI in the web are well understood. Here we map these to the IoT. The most common situation for a device is that *the device has no certificate* and thus no attestations about security at all.

The second most frequent case is a *false negative*. In this case, usually there is an expired certificate, but the facts remain valid so it is trustworthy in practice. Another possible source of errors is that a manufacturer self-signed the certificate. Given the attestation is of the manufacturer's identity in a device the issuer manufactured, we are confident in also classifying this as a false negative – the trusted is arbitrarily identified as untrustworthy.

The third case is when the untrustworthy is identified as trustworthy, or a *false positive*. This may be because a valid certificate was issued, but the facts have changed so it is not trustworthy. Alternatively it could be the case that issuer did not properly verify the subject; or the issuer may have knowingly issued a rogue certificate.

4) Expectations and Experiences of Failure: Having enumerated the ways in which the infrastructure is vulnerable and certificates are vulnerable, we now move to risk perceptions of PKI.

Technologists tended to describe their interactions with websites as existing at two tiers of trust. A lower tier of trust exists when they are browsing content online from a reputable party such as a newspaper. In this case, they would not confirm a domain or validate the connection. However, when there is a higher risk context they indicated a greater degree of personal investment of time and attention to secure themselves. Recommended actions include double checking the URL, clicking to evaluate the certificate information, and seeking specific security markers on the page. (The practice of experts in seeking specific security markers on web pages has been validated in laboratory experiments. Ironically, this trust in the presence of the lock made more expert participants less resilient to https-enabled phishing attacks than non-experts, illustrating a subtle risk of trust indicators [35].)

Most noted that their trust is to some extent confidence in themselves. For example, participants noted that they would not engage in personal banking or other finance on others' devices because they worry about "cross-contamination". In the best case scenario, they would like to have a separate machine for online banking or do everything offline by going to a branch in-person. In practice this is often implemented by using few trusted networks, VPNs, and virtual machines.

Every participant has had the experience of false negatives in website warnings. That is, all have experienced certificate warnings on valid websites. Most reported seeing these warnings frequently, especially warnings on expired certificates. They believe those warnings came from maintenance errors or the laziness of certificate maintainers. In addition to that, technologists also noted the inconsistency between browsers, the inconsistency between devices, and the server configuration issues. Although warnings keep occurring, interviews clearly show that nearly half of policymakers chose to ignore them. Based on their familiarity with the websites, they believe the websites they visited were legitimate. On the contrary, half of the technologists are curious and will explore the putatively malicious sites where the warning occurred. They also tend to believe the site was compromised rather than thinking it is legitimate. However, it is worrying to see that, for non-technical users, there is no way to determine whether warnings came from maliciousness or just maintenance failures. This reflects the fact that warnings are ignored unless they block content, and when they block content, network functionality is limited so functionality trumps security.

Additionally, three technologists and six policymakers said they have gone to malicious websites accidentally, but no warnings occurred. Policymakers reported that this situation happened to them frequently. The other three technologists and remaining six policymakers felt they had had the same experience but remained uncertain. It is important to consider that many 'threats' cannot simply be detected. As one policymaker said, "the goal of attacks is to avoid detection."

Therefore, there is a greater chance that malicious activities happen invisibly than that false negatives are unobservable. There are different types of malicious websites, such as sites that deliver malware, phishing sites, and sites that are simply operating fraudulent business scams. All participants agreed that there is no single solution or purely technical solution of all these types of malice. Well crafted malicious websites will not trigger any warnings because they can get legitimate certificates from free Certificate Authorities (CAs) such as *letsencrypt.org*. By matching the domains and the certificates, the browsers will recognize these websites as valid. One technologist said, "PKI certificates only tells you the operator of the site is using a certificate that matches the domain that you attempted to visit, it does not tell you the intent of the people operating that domain or the content in the domain" Strategies against these threats include but are not limit to trustworthiness assessments and risk mitigation. In addition, technologists suggested that there could be a government body or non-governmental organization (NGO) to provide a certification process to make the PKI become a secure attestation. In practice, due to the goal of HTTPS Everywhere, *letsencrypt.org* offers free certificates making these certificates ever less meaningful in terms of authentication and identity attestation.

Nearly half of policymakers pointed to strengthening existing regulations since the usage of IoT devices is usually associated with sensitive data. However, as indicated previously, users do not have transparency and control over what is being collected so consent may not be feasible. Data is widely used not directly for the IoT but to support other commercial activities, such as behavioral advertising. Both policymakers and technologists advocated for making the person and their privacy more central to the design and operation of IoT. Both camps expressed that people should have control over their data, rather than being powerless. Addressing the current state of security and privacy, policymakers mentioned the term "buyer beware."

All policymakers agreed that there is a critical need for risk mitigation and a clear role for regulations; that the *Framework for Electronic Commerce* model with its focus on innovation and acceptance of "buyer beware" from the web is a dangerous model for the IoT [1]. Yet the way forward on regulation is unclear, and a technological solution is seen as ideal. In a complementary thread, technologists focused on out of the box security. All participants uniformly identified issues of usability, which in this context could be considered risk communication. One technologist said, "X.509 assumes its user has a user interface and can inspect and understand things, but that is not the case in IoT." Technologists assert that manufacturers should take the responsibilities to ensure the devices are secure out of the box without any user involvement. When pressed to explain how, technologists assumed some form of incentives, liability, or other regulation. Similarly, policymakers struggled with how to make risk mitigation happen in practice: incentives, ex-post or ex-ante regulation all have serious limitations in computer security and privacy [29]. When pressed to explain how to secure the IoT, policymakers pointed to the promise of technological innovation.

5) Certificate Authority Requirements: All participants agreed that there should be some characteristics a certification provider must prove in order to be trustworthy, but neither policymakers nor technologists provided a clear list of such requirements. Technologists believed that trust is not a technical problem. Rather than focusing on the characteristics of CAs, they felt that technology providers that are going to rely on PKI should have a clear public statement about what principles they choose to rely on. (This is supported in the current version of X.509 as Extended Key Usage (EKU), a field which specifies limitation and context of use. EKU is complimentary to the proposed list of approved connections and trustworthy domains embedded in a Manufacturers Usage Description [38].) Further, one policymaker proposed that there should be a process to verifying any such set of obligations through auditing. In a group that shares the same set of criteria, it could be possible in theory for members to trust each other because they are all verified and audited according to the same principles. Currently, there are auditing requirements under all root programs and in the CA/Browser Forum [50]. This has not prevented systematic CA failures and abuse, e.g. VeriSign, Wosign.

6) Security for Consumers: Businesses are eager to introduce new technologies to the marketplace and to be the first to market. Consumer protection and long-term device maintenance can easily be neglected. For example, as one of the biggest e-commerce platforms, eBay supports vendors who are selling products where there is no reasonable expectation that they will continue to exist in the next quarter. It is unreasonable to expect those devices to have strong security.

All participants agreed that at least some of the burden for ensuring security should be removed from users and be placed on more capable institutions. In fact, as noted in the survey results, the word "security" has no consistent definition and may be nonsensical to non-technical consumers. From the perspective of the interviewees, it seems that people don‘t care about security, they don‘t understand the risk of security, and even products that seem simple to experts confuse most people. Interviewees felt that consumers make purchase decisions by trusting third-party brands, such as Google, Amazon, and Apple. In addition, good online reviews, low prices, and discounts are attractive to consumers. A number of policymakers suggested that some basic training or a public awareness campaign could be provided. At least there should be an awareness that the components and devices purchased from Taobao, Alibaba, or eBay are likely to be insecure and untrustworthy. Both populations identified the need for a third-party evaluation to evaluate the devices for users. Beyond that, technologists also pointed out that there could be security indicators, such as a label on the package so users can see if a device is secure out of the box. Consumer Reports Digital Standards [2] is beginning to address the first; AllJoyn is targeted at the second [27]. Yet industry adoption is slow.

When devices fail, possible solutions provided by policymakers are to 1) discard, 2) isolate, and 3) white hat

hacking. The suitability of answers depends on at least three factors: participants' income level, expertise, and device quality. Most policymakers have stable and relatively high incomes. Rather than spending time to find a way to maintain the device, replacing it may be an easier choice. Others, including graduate students, suggested isolating those devices so they can remain in use.

Students' and technologists' arguments can be fairly summarized as follows using the language of the interview responses: *devices built by random manufacturers by grabbing crappy scripts online and sticking them together into the processors will fail*.

Policymakers also mentioned that white hat hacking could be a solution. On a complementary note, technologists suggested code escrow (not to be confused with key escrow). Technologists prefer to fix the problem from the foundation suggesting that when building a device the update system should be built before any other elements. If a vendor goes out of business, there should be some requirements that the marketed code be escrowed by some trustworthy third-party or NGO. Beyond this, technologists also indicated the warranty for some categories of devices so there can be some protection even after the company departs, consumers can still get support. 'Finding responsible stakeholders' is the main challenge of this issue. Essentially, the technologists identified all the challenges of ex-post and ex-ante regulation. None of the participants supported 'exceptional access' (i.e., the current version of the perennial key escrow proposal) as a solution because it will bring in more vulnerabilities while solving no problems.

In terms of white hat hacking, only two participants thought it could potentially be useful and then only if the hacker were the legitimate owner of the device. Yet bug bounties are essentially white hat hacking and common in the industry, and law enforcement hacking continues unabated.

7) The Future of Domain Names in the IoT: X.509 is fundamentally an attestation of identity via domain name. Although any holder of a random domain name is implicitly not trusted, almost all participants agreed that DNS vulnerabilities will be extended into the IoT environment and under the same governance as with the Internet today. Thus, the potential risks will shift to the IoT environment as well. Naturally, IoT devices are going to connect to the outside world. By looking at domain names, users can theoretically know where the device is connecting to. However even this small measure may be unavailable. Three participants mentioned that most IoT devices don't have a user interface. Without an interface, people might not be able to know if the device is connecting to the manufacturer or to a malicious entity.

8) The Roles of Marketplace, Technologists, and Policymakers: Obviously, the marketplace, technologists, and policymakers all have roles in resolving the risks of IoT. In theory, the marketplace directly engages with consumers. One benefit from the marketplace is that, by paying some premium, consumers could ensure their devices were more secure and offered greater longevity. Yet there is currently no way for a consumer to make this judgment.

One problem with the marketplace is the abuse of information. Security does not ensure privacy. For example, knowing that Google has internal security and that information is encrypted does not resolve the privacy issues from the face that a subcontractor in Ukraine could surveil the homes that contain a Nest Camera [28]. Beyond this, one technologist noted that the marketplace as a whole is often dominated by short-term players. Such players do not have any long-term interest in maintaining the security of devices. Therefore, technologists and policymakers should step in and enforce standards of better security and integrity.

The role of technologists, participants suggested, is to identify the potential risks and communicate the risks with the other parties. Policymakers argued that technologists should also be careful to not create more problems while innovating. Neither of the above parties can resolve the problems independently. For the technologists, the role of the policymakers is to push the regulations to force cooperation by all stakeholders. Policymakers should lead in this space with some regulatory and legal liability rules that support the technologists in their professional roles in making secure and reliable devices.

Overall, technologists cannot solve the problems alone and policies are needed. All stakeholders (e.g., vendors, policymakers, technologists) should collaborate. Policies and regulations should be enforced and such enforcement supports security engineers seeking to mitigate risk by design.

9) Messages to Technologists, Policymakers, and Consumers: At the end of the interview, we collected the messages the participants would share with other stakeholders: policymakers, consumers, and technologists.

Policy experts recommended that technologists should strengthen identity validation, provide consumer education, and redesign the current model of public key certificates to improve the trust and reliability of IoT devices. One policymaker also suggested that technologists should all get a degree in sociology because, as mentioned previously, not all problems have engineering solutions.

Technologists suggested that policymakers should create regulatory frameworks for security indicators, for code escrow, for consumer rights, to mandate secure software updates, and to better protect the consumer experience. There should also be liability regulations for damage, so if the data are leaked, someone will be responsible for recovery and there will be a corresponding cost.

All participants agreed that non-technical users should be aware of the (un)trustworthiness of IoT devices. AllJoyn indicates in its documentation that consumers should take due care, and not to trade sensitive data for lower prices or convenience. Policymakers suggested that average consumers should be educated about basic knowledge of IoT and PKI. Yet there is not an option to use these services in a privacy-preserving manner regardless of willingness to pay. And in fact, attempting to refuse to engage with technology may result in only losing the benefits while not only gaining no additional privacy but also paying the price of not participating [43]. For example, if a person decides to stop using their Facebook account they are still tracked in detail through the

Facebook advertisement APK embedded in more than the Facebook app. The Facebook APK is included in more than one IoT mobile app with neither indication nor notification of the extent of the data compilation. The Facebook APK has sent financial, personal, and health information including not only the location but AI-driven estimates of health status including menstrual cycles [54]. This is in addition to the ubiquitous tracking provided to Facebook by cookies, web fingerprinting, and other tracking technologies.

The industry standards appear to embrace the buyer beware model. The following is from the AllJoyn web page in March of 2019 [27], as requirements for consumers:

- Users should be trained with some basic security knowledge.
- They should be aware of the risks of using IoT devices.
- They should be careful to not trade sensitive data for low price or convenience.

This is simply not feasible. Privacy policies are unreadable, change frequently, and it would be a life's work to read all of these [33], [40], [56].

Technologists recognized this and essentially advocated for avoiding classes of technology. In particular, many refuse to purchase home assistants (e.g., Alexa) because voice-response technologies continually record and transmit daily conversation. The technologists suggested that IoT devices should not be trusted and had consistent security recommendations. One of these was setting up the devices on their own 'guest' network. Yet even technologists have had information exposed; for example, one highly expert participant purchased a Nest thermostat without being aware that there were an always-on motion detection and voice recordings sent to the Nest operations center in Ukraine [9].

To cryptographers, policymakers suggested that the key management system should be refined. Together with technologists, they should build a model to develop consumer protections. One TPRC interviewee would simply say, "Do a good job!" to technologists. Similarly, one cryptographer said nearly the same when asked if he wanted to communicate to policymakers: " Be careful!".

The spectrum of IoT is broad and is expanding. One technologist expressed concern over the risks of devices with a wide range of processing capabilities, power, connectivity, and exposure to the network. No matter how powerful the devices are, they all get updates. No matter how powerless, they can all be subverted. There should be different regulations and policies around different devices. At some point, scale will become an issue. We need to think about which classes of devices require updates and which do not. Technologists indicated that the IoT is a place that needs a huge amount of regulation, but it is lacking in the current environment.

V. CONCLUSION

In his seminal work, Graham Allison [8] noted that when the State Department seeks a military solution and the Department of Defense seeks a diplomatic solution then the challenge is essentially seen as insurmountable by each. A similar observation might apply to the discussion of applied cryptographers and policy professionals in developing a foundation of meaningful trust for the IoT. PKI is meeting the need for transit encryption but systematically fails to provide identity authentication or risk information. End-users' reliance on PKI as a guarantee of behavior, quality, or governance indicates that it is both trusted and untrustworthy.

Both policymakers and technologists are unsatisfied with the current X.509 trust model. All agreed that PKI failures are not simply technical problems, yet neither community had clear advice for the other. In the small set of people we interviewed, there was some agreements about the risks, but less about the solutions. There is a consensus that this is an interdisciplinary problem, and a complete agreement about the critical importance of PKI. However, there is no agreement about what solutions can be engineered or constructed, nor what policies should be adopted. While no one opposed Certificate Transparency (CT) [31], none of the participants believed it would solve the core trust challenges. It is simply another level of indirection. We requested an interview with a Certificate Transparency author; however, the request wasn't accepted. It is the perception of the authors that this is because CT designers believe that the system resolved issues of certificate reliability. This may have created a bias among technical participants in that only those who perceive a problem participated in the conversation.

In both communities, there was an awareness of the vulnerability and lack of underlying trust provided, with an overall view of the state of the Internet as "buyer beware". Recall the focus by technologists on manufacturers' responsibilities and policymakers' identification of the need for risk mitigating rules and technologies. Participants both expressed a hope or even proposed a requirement that users could be trained to carefully validate the operation of PKI in the IoT; and simultaneously recognized the futility of this. The challenges to regulation on the Internet can be seen as exacerbated by or mitigated by the physical nature of the IoT because devices are physical and each has a jurisdiction, which addresses some of the challenges of Internet governance. Yet the problems of firms too large and too small remain. There are judgment-proof firms that are too small to litigate (as they simply disappear). Conversely, there are too big to fail firms. This is best illustrated by VeriSign, now DigiCert, that is a large part of the infrastructure – too powerful to remove despite cavalier, irresponsible, and dangerous behaviors in the issuance of tens of thousands of certificates.

The two major points of diversion were that the Real World Crypto (RWC) participants focused on the out of the box experience and the security assumptions communicated at the moment of installation. Technologists accept PKI as the infrastructure that is available and seek to improve it through agility and isolation. The second core difference was that members of the policy community indicated that disposal of insecure IoT or cyberphysical devices was a feasible way forward.

Specifically, in the policy realm, there seemed to be an acceptance of the potential of disposing of failed devices or built-in expiration dates. An assumption that vulnerability

is an unusual state of affairs is required for this to be reasonable. Yet device vulnerabilities are endemic and with a mathematical advance (i.e., factoring), could be catastrophic. As we consider PKI that is deeply embedded in the operations of automobiles or medical equipment, the infeasibility of this becomes clear. Imagine a situation in which every vehicle with AllStar would simply be considered inoperable as a matter of policy and practice, as opposed to the actual situation in which the vehicles were carefully updated [36].

In terms of recovery, disposal was dismissed by those in the security community at the previous workshop. In communications at the preceding workshop, it was identified as a dangerous idea, creating perverse incentives, risking catastrophic disruptions, and excessive carbon impacts as embedded (but otherwise useful products) are made inoperable. However, the solution was supported at the workshop by multiple policymakers and participants from start-ups, so it is relevant for discussion. Depending on the value and quality of the IoT device and the interviewees' income level different solutions are applicable. Isolation, code escrow, or white hacking were all mentioned as solutions. There is some hope for improved isolation as NIST moves forward with the Manufactures Usage Description standard. [3] AllJoyn is an additional effort identified as a positive step forward specifically referenced at RWC. AllJoyn also aligns with the transparency efforts at Commerce under the Software Bill of Materials [5]. The requirement of having an update system in place is included in the IEEE Building Code for IoT [39] and the Consumer Reports' Digital Standard [2].

Code escrow or open-source requirements were mentioned by the technical participants, yet not by the policymakers. The distinction between code escrow and key escrow is critical here.

While there was a partial disagreement about the definition of the IoT, there was an agreement for the need for a robust trust infrastructure as IoT becomes pervasive. Specifically, three participants stated that making the IoT separate from the "I" was "nonsense" and "a marketing term".

All agreed that current practices are inadequate. The goal of each stakeholder interviewed was to create a trustworthy environment. Yet there was no consensus of the type of regulation nor the path forward.

Many of the different challenges as defined by failure modes were either seen as beyond the scope of technology by security experts, and as technical problems by policy experts.

The failure modes were found to be comprehensive. The feedback identified these as having adequate coverage of identifiable risks. The particular solutions identified by the participants were, when coded, found to fall into seven basic categories:

- improved code for evaluating certificates,
- fewer trusted roots,
- quicker and more effective responses to rogue events,
- more nuanced and fewer warnings,
- warnings that include indicators about why the warning was issued,
- audits to support these indicators, and
- improved alignment between user understanding and warning implications.

In terms of **improving code**, there is both current research and considerable industry efforts. Specifically, Acar et al. [7] suggested promoting security through increasing the security awareness of developers. In addition, they found that developers are unlikely to give up using convenient resources and that security will remain a secondary concern. Acar et al. also suggested a need for rewriting resources, such as OpenSSL, to make them more usable. Yet there is a lack of resources to do this; a standards to know when it is done; and of liability for even egregious code to motivate the investment.

In terms of **fewer trusted roots**, this is an essential matter of governance. Since trust relationships vary from different users and organizations, the authentication decisions shouldn't be global and absolute [6].

The new standards embed potential for identification in the IoT but this requires depending on domain names and associated information. Necessary limits on the scope of trust and more consumer protection than that provided by the CA/Browser Forum are needed. Given the *browser* is literally in the name of the forum, that it is not an ideal fit for IoT is not a surprise. Yet without purposeful action, that will be the natural path of progression.

Notably, **improved alignment between user understanding and warning implications** assumes that the burden will remain on users. This also implies the existence of a reliable interaction between the person and the (Io)things on a platform that can provide timely warnings. This is inherent to web browsing but problematic in the IoT. What platform is appropriate for issuing warnings? How can users respond? While there is research in tactile and ambient interactions that may address this problem, it is certainly unsolved. (e.g., [32], [55], [59].)

Perhaps the most important conclusion is that since security is never entirely perfect and the moral hazard of disposal will create greater harm (as well as carbon impact) then isolation and not disposal in the face of subversion is the policy choice for the IoT.

Following from this, we explicitly note that the limits of human cognition and expertise require the minimization of human interaction out of the box. This requires the coordination of offerings of different providers in different industries to make isolation out of the box a reality. The need for such security is a focus of current efforts (recall NIST, MUD, Alljoyn) yet basic research for the role of the human in these standards is needed.

Current structures for governance of domain names and PKI have created a vacuum, one that is partially but inadequately filled by large tech players concentrating power and then declaring temporary victory (e.g., Transparency, DNS over HTTP). Yet in three investigations over three years, this has not been shown to be adequate nor do technologists expect these victories to be permanent.

[3]Consider the April 9, 2019 NIST event *Mitigating IoT-Based DDoS Industry Day* which aligns with investment in developing operational MUD standards `https://www.nccoe.nist.gov/events/mitigating-iot-based-ddos-industry-day`

VI. LIMITATION AND FUTURE RESEARCH

There are questions beyond the scope of a single project and we hope that the research encourages reconsideration of these: Regulation, but what regulation? Redesign, into what and by whom? Inadequate governance, but how to improve it?

Although we have been able to observe a general demand for addressing the failure modes in PKI for IoT, there's no agreement about what regulation should exist nor even the relevant set of stakeholders.

Our specific near-term research goal is to examine proposals, make explicit the requirements on the end user, and if possible define interventions for these in IoT. Our focus within the NIST, AllJoyn, or MUDs frameworks includes device isolation, limits on trust, and risk communication.

A clear challenge is that CAs are part of the problem. Under the current structure, CAs are effectively judgment-proof. For the smaller players, where providing certificates is the core business, then the business must cease to exist if it is untrustworthy. In addition, many CAs are too big to fail. The sale of Symantec's certificate practice to Digicert – because Symantec is in the Internet infrastructure yet proven to be untrustworthy – is an example of this [4], [47]. Thus a post-hoc or harm-based regulatory regime based on liability has potential but must be approached with caution. It is clear from past events (e.g., WoSign backdating certificates [30] as well as VeriSign and Symantec's chronic abuse of trust) that risk-based or ex-ante is not effective in the current infrastructure, as CAs continue to issue harmful certificates. WoSign is thriving in the Chinese market [4], and the certificate infrastructure of Symantec remains embedded in systems across the globe (albeit with a new operator). Conversely, Mozilla and Google rejected United Arab Emirates' mercenary group DarkMatter as a trusted root suggests that the pattern of including any minimally valid root certificate until after harm has been proven is no longer a standard practice [48]. (Microsoft does not appear to have ever included DarkMatter in its Root Program.) This is an example of ex ante regulation.

An additional challenge is to create meaningful and effective warnings. The research in this domain is lead by Google, and thus is focused on global solutions. Contextual or localized risk communication is an open and under-researched challenge.

A core challenge in PKI is that it has functioned as an architecture of global trust. Thus the emerging questions of creating local measures of trust, identifying where these are appropriate, and implementing this in practice remain open. Because these problems integrate technical nuance, human interactions, warning science, and permissible harms across jurisdictions, the design of the solutions (and then the challenges of implementation and adoption) require coordination at a larger scope than previously seen.

[4]Wosign Chinese Consumers: `https://www.wosign.com/english/Who_uses_WoSign.htm`

VII. ACKNOWLEDGEMENT

This research was supported in part by the National Science Foundation under CNS 1565375, CNS 1814518; Cisco Research Support, and the Comcast Innovation Fund. We would like to acknowledge the assistance of Shakthidhar Gopavaram, Jacob Abbott, and Andrew Kim who provided critical technical contributions in implementation of the experiment through coding. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the the US Government, the National Science Foundation, Cisco, Comcast, nor Indiana University.

REFERENCES

[1] Framework for electronic commerce.
[2] Consumer reports and disconnect and ranking digital rights and the cyber independent testing lab and aspiration digital standard, May 2017. `https://www.thedigitalstandard.org/`.
[3] Misissued/suspicious symantec certificates, 2017.
[4] Symantec sells its ca business to digicert, 2017.
[5] Ntia software component transparency, 2019.
[6] Martín Abadi, Andrew Birrell, Ilya Mironov, Ted Wobber, and Yinglian Xie. Global authentication in an untrustworthy world. In *Proceedings of the 14th USENIX Conference on Hot Topics in Operating Systems*, HotOS'13, pages 19–19, Berkeley, CA, USA, 2013. USENIX Association.
[7] Y. Acar, M. Backes, S. Fahl, D. Kim, M. L. Mazurek, and C. Stransky. You get where you're looking for: The impact of information sources on code security. In *2016 IEEE Symposium on Security and Privacy (SP)*, pages 289–305, May 2016.
[8] G.T. Allison and P. Zelikow. *Essence of Decision: Explaining the Cuban Missile Crisis*. Alternative Etext Formats. Longman, 1999.
[9] Nick Bastone. After a big privacy backlash, google's nest explains which of its products have microphones and why, 02 2019.
[10] A. Begel, J. Bosch, and M. Storey. Social networking meets software development: Perspectives from github, msdn, stack exchange, and topcoder. *IEEE Software*, 30(1):52–66, Jan 2013.
[11] Chad Brubaker, Suman Jana, Baishakhi Ray, Sarfraz Khurshid, and Vitaly Shmatikov. Using frankencerts for automated adversarial testing of certificate validation in ssl/tls implementations. In *Proceedings of the 2014 IEEE Symposium on Security and Privacy*, SP '14, pages 114–129, Washington, DC, USA, 2014. IEEE Computer Society.
[12] Krishna C Bathina, L Camp, and Shakthidhar Gopavaram. Stopping the big bad wolf: Measuring online risk behavior. 01 2015.
[13] L Jean Camp, Helen Nissenbaum, and Cathleen McGrath. Trust: A collision of paradigms. In *International Conference on Financial Cryptography*, pages 91–105. Springer, 2001.
[14] D. Chasaki and T. Wolf. Attacks and defenses in the data plane of networks. *IEEE Transactions on Dependable and Secure Computing*, 9(6):798–810, Nov 2012.
[15] Yuting Chen and Zhendong Su. Guided differential testing of certificate validation in ssl/tls implementations. In *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering*, ESEC/FSE 2015, pages 793–804, New York, NY, USA, 2015. ACM.
[16] David Cooper, Stefan Santesson, Stephen Farrell, Sharon Boeyen, Russell Housley, and William Polk. Internet x. 509 public key infrastructure certificate and certificate revocation list (crl) profile. Technical report, 2008.
[17] Sanchari Das, Andrew Kim, Zachary Tingle, and L Jean Camp. All about phishing exploring user research through a systematic literature review. In *Proceedings of the Thirteenth International Symposium on Human Aspects of Information Security & Assurance (HAISA 2019)*, 2019.
[18] Sanchari Das, DongInn Kim, Timothy Kelley, and L Jean Camp. Grifting in the digital age.
[19] US Department of Homeland Security. The future of smart cities: Cyber-physical infrastructure risk, 08 2015.
[20] Andrew Whalley Chrome Security Devon O'Brien, Ryan Sleevi. Chrome's plan to distrust symantec certificates, 2017.
[21] Emily Stark Chrome security team Devon O'Brien, Ryan Sleevi. Distrust of the symantec pki: Immediate action needed by site operators, 2018.

[22] Zheng Dong, Kevin Kane, and L. Jean Camp. Detection of rogue certificates from trusted certificate authorities using deep neural networks. *ACM Trans. Priv. Secur.*, 19(2):5:1–5:31, September 2016.

[23] Manuel Egele, Peter Wurzinger, Christopher Kruegel, and Engin Kirda. Defending browsers against drive-by downloads: Mitigating heap-spraying code injection attacks. In *International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment*, pages 88–106. Springer, 2009.

[24] Thales eSecurity. What is public key infrastructure (pki).

[25] M. Fecko, K. Manousakis, K. Young, Jaewon Kang, A. Pachulski, and W. Phoel. Mitigation of control plane attacks at the network layer. In *MILCOM 2015 - 2015 IEEE Military Communications Conference*, pages 444–449, Oct 2015.

[26] A. Ferreira, R. Giustolisi, J. Huynen, V. Koenig, and G. Lenzini. Studies in socio-technical security analysis: Authentication of identities with tls certificates. In *2013 12th IEEE International Conference on Trust, Security and Privacy in Computing and Communications*, pages 1553–1558, July 2013.

[27] Open Connectivity Foundation. Alljoyn open source project, 2019.

[28] Geoffrey A. Fowler. The doorbells have eyes: The privacy battle brewing over home security cameras, 2019.

[29] Vaibhav Garg and L Jean Camp. Spare the rod, spoil the network security? economic analysis of sanctions online. In *2015 APWG Symposium on Electronic Crime Research (eCrime)*.

[30] M Gervase. Wosign and startcom. *Mozilla*, 2016.

[31] Google. Certificate transprency.

[32] Julie A Jacko. *Human computer interaction handbook: Fundamentals, evolving technologies, and emerging applications*. CRC press, 2012.

[33] Carlos Jensen and Colin Potts. Privacy policies as decision-making tools: an evaluation of online privacy notices. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, pages 471–478. ACM, 2004.

[34] David Kaloper-Meršinjak, Hannes Mehnert, Anil Madhavapeddy, and Peter Sewell. Not-quite-so-broken tls: Lessons in re-engineering a security protocol specification and implementation. In *Proceedings of the 24th USENIX Conference on Security Symposium*, SEC'15, pages 223–238, Berkeley, CA, USA, 2015. USENIX Association.

[35] Timothy Kelley and Bennett I Bertenthal. Attention and past behavior, not security knowledge, modulate users' decisions to login to insecure websites. *Information & Computer Security*, 24(2):164–176, 2016.

[36] K. Koscher, A. Czeskis, F. Roesner, S. Patel, T. Kohno, S. Checkoway, D. McCoy, B. Kantor, D. Anderson, H. Shacham, and S. Savage. Experimental security analysis of a modern automobile. In *2010 IEEE Symposium on Security and Privacy*, pages 447–462, May 2010.

[37] Adam Langley. Apple's ssl/tls bug, 2014.

[38] Eliot Lear, Ralph Droms, and Dan Romascanu. Manufacturer Usage Description Specification. RFC 8520, March 2019.

[39] Ulf Lindqvist and Michael Locasto. Building code for the Internet of Things. September 2017.

[40] Aleecia M McDonald and Lorrie Faith Cranor. The cost of reading privacy policies. *Isjlp*, 4:543, 2008.

[41] Ulrike Meyer and Susanne Wetzel. A man-in-the-middle attack on umts. In *Proceedings of the 3rd ACM workshop on Wireless security*, pages 90–97. ACM, 2004.

[42] Michael Myers, Rich Ankney, Ambarish Malpani, Slava Galperin, and Carlisle Adams. X. 509 internet public key infrastructure online certificate status protocol-ocsp. Technical report, 1999.

[43] Greg Norcie and L Camp. The price of privacy: An examination of the economic costs of abstention from social networks. 10 2015.

[44] D. Park. Social life of pki: Sociotechnical development of korean public-key infrastructure. *IEEE Annals of the History of Computing*, 37(2):59–71, Apr 2015.

[45] Prashanth Rajivan, Pablo Moriano, Timothy Kelley, and L Jean Camp. Factors in an end user security expertise instrument. *Information & Computer Security*, 25(2):190–205, 2017.

[46] Fahmida Y. Rashid. Phishing sites exploit trust in valid ssl certificates.

[47] Fahmida Y. Rashid. Google to symantec: We don't trust you anymore, 2017.

[48] Timothy Richter. Google prohibits darkmatter certificates from android and chrome, 07 2019.

[49] Gustav Rydstedt, Baptiste Gourdin, Elie Bursztein, and Dan Boneh. Framing attacks on smart phones and dumb routers: tap-jacking and geo-localization attacks. In *Proceedings of the 4th USENIX conference on Offensive technologies*, pages 1–8. USENIX Association, 2010.

[50] Nicolas Serrano, Hilda Hadan, and L. Jean Camp. A complete study of p.k.i. (pki's known incidents), 2019.

[51] Tom St Denis. *Cryptography for developers*. Elsevier, 2006.

[52] Joshua Streiff, Olivia Kenny, Sanchari Das, Andrew Leeth, and L Jean Camp. Who's watching your child? exploring home security risks with smart toy bears. In *2018 IEEE/ACM Third International Conference on Internet-of-Things Design and Implementation (IoTDI)*, pages 285–286. IEEE, 2018.

[53] Synopsys Editorial Team. Understanding the apple 'goto fail;' vulnerability, 2014.

[54] Lisa Vaas. Facebook apps secretly sending sensitive data back to the mothership, 2 2019.

[55] Sanchari Das Vafa Andalibi, Joshua Streiff and Jean Camp. Securtle for cyber defense. 2019.

[56] Tony Vila, Rachel Greenstadt, and David Molnar. Why we can't be bothered to read privacy policies models of privacy economics as a lemons market. pages 403–407, 01 2003.

[57] Kenneth M Walker, Daniel F Sterne, M Lee Badger, Michael J Petkac, David L Shermann, and Karen A Oostendorp. Confining root programs with domain and type enforcement (dte). In *Proceedings of the 6th USENIX Security Symposium*, volume 10, 1996.

[58] David A. Wheeler. The apple goto fail vulnerability: lessons learned, 2017.

[59] Z Zimmerman and L Jean Camp. Elder-friendly design's effects on acceptance of novel technologies. *Elderly Interaction Design CHI*, 2010.