

Cybertrust in the IoT Age



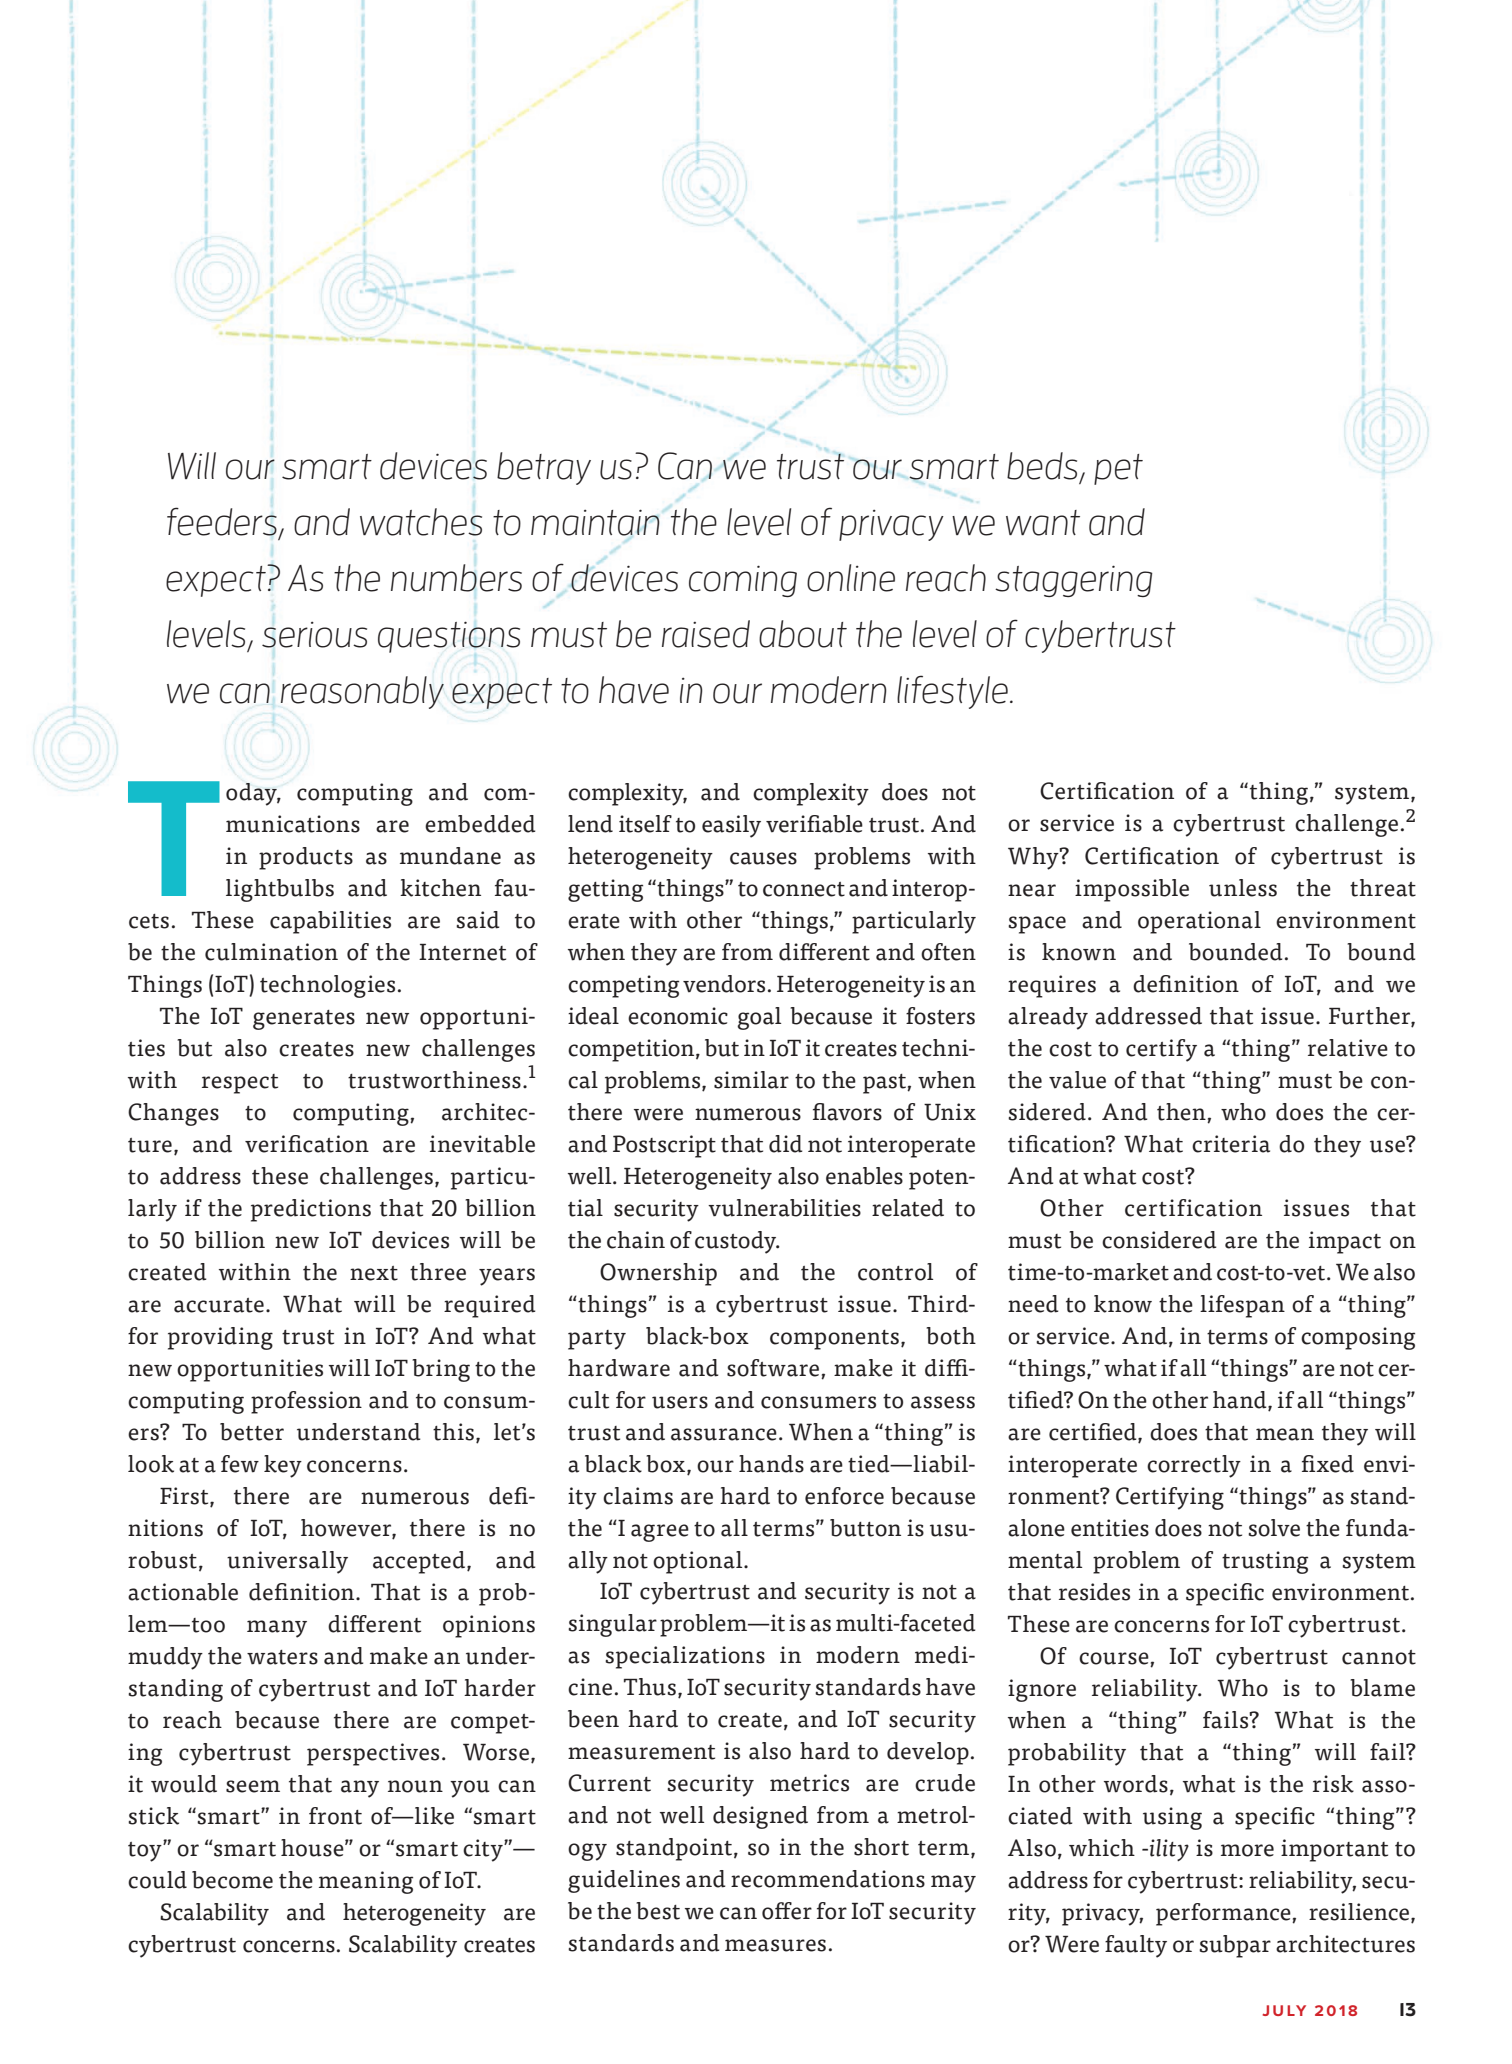
Jeffrey Voas, US National Institute of Standards and Technology

Rick Kuhn, US National Institute of Standards and Technology

Constantinos Kolias, George Mason University

Angelos Stavrou, George Mason University

Georgios Kambourakis, University of the Aegean



Will our smart devices betray us? Can we trust our smart beds, pet feeders, and watches to maintain the level of privacy we want and expect? As the numbers of devices coming online reach staggering levels, serious questions must be raised about the level of cybertrust we can reasonably expect to have in our modern lifestyle.

Today, computing and communications are embedded in products as mundane as lightbulbs and kitchen faucets. These capabilities are said to be the culmination of the Internet of Things (IoT) technologies.

The IoT generates new opportunities but also creates new challenges with respect to trustworthiness.¹ Changes to computing, architecture, and verification are inevitable to address these challenges, particularly if the predictions that 20 billion to 50 billion new IoT devices will be created within the next three years are accurate. What will be required for providing trust in IoT? And what new opportunities will IoT bring to the computing profession and to consumers? To better understand this, let's look at a few key concerns.

First, there are numerous definitions of IoT, however, there is no robust, universally accepted, and actionable definition. That is a problem—too many different opinions muddy the waters and make an understanding of cybertrust and IoT harder to reach because there are competing cybertrust perspectives. Worse, it would seem that any noun you can stick “smart” in front of—like “smart toy” or “smart house” or “smart city”—could become the meaning of IoT.

Scalability and heterogeneity are cybertrust concerns. Scalability creates

complexity, and complexity does not lend itself to easily verifiable trust. And heterogeneity causes problems with getting “things” to connect and interoperate with other “things,” particularly when they are from different and often competing vendors. Heterogeneity is an ideal economic goal because it fosters competition, but in IoT it creates technical problems, similar to the past, when there were numerous flavors of Unix and Postscript that did not interoperate well. Heterogeneity also enables potential security vulnerabilities related to the chain of custody.

Ownership and the control of “things” is a cybertrust issue. Third-party black-box components, both hardware and software, make it difficult for users and consumers to assess trust and assurance. When a “thing” is a black box, our hands are tied—liability claims are hard to enforce because the “I agree to all terms” button is usually not optional.

IoT cybertrust and security is not a singular problem—it is as multi-faceted as specializations in modern medicine. Thus, IoT security standards have been hard to create, and IoT security measurement is also hard to develop. Current security metrics are crude and not well designed from a metrology standpoint, so in the short term, guidelines and recommendations may be the best we can offer for IoT security standards and measures.

Certification of a “thing,” system, or service is a cybertrust challenge.² Why? Certification of cybertrust is near impossible unless the threat space and operational environment is known and bounded. To bound requires a definition of IoT, and we already addressed that issue. Further, the cost to certify a “thing” relative to the value of that “thing” must be considered. And then, who does the certification? What criteria do they use? And at what cost?

Other certification issues that must be considered are the impact on time-to-market and cost-to-vet. We also need to know the lifespan of a “thing” or service. And, in terms of composing “things,” what if all “things” are not certified? On the other hand, if all “things” are certified, does that mean they will interoperate correctly in a fixed environment? Certifying “things” as stand-alone entities does not solve the fundamental problem of trusting a system that resides in a specific environment. These are concerns for IoT cybertrust.

Of course, IoT cybertrust cannot ignore reliability. Who is to blame when a “thing” fails? What is the probability that a “thing” will fail? In other words, what is the risk associated with using a specific “thing”? Also, which *-ility* is more important to address for cybertrust: reliability, security, privacy, performance, resilience, or? Were faulty or subpar architectures

employed? Were the “things” that were employed defective? Were the best “things” available at that time used? Was the IoT system over-engineered and too much money spent? And would it be foolish to discount the importance of the expected operational usage profile? Did the engineers know the environment and context their IoT system would exist in? And is the system designed with respect to the expected operational usage profile? The point is that cybertrust in IoT cannot ignore reliability.

IoT testing is also a cybertrust concern. This is partially due to scalability and heterogeneity, but more importantly, it is the massive number of combinations of potential inputs and the fact that many IoT systems control actuators and have binary or very small output spaces.³

Data is the lifeblood of IoT systems. Where data originates from has an impact on cybertrust. Leased data originates from vendors at the time of their choosing and with the integrity of their choosing. The possibility of data-tampering cannot be dismissed. Data integrity is pivotal, so knowing how secure data is from accidental problems or malicious tampering, delay, or theft is a cybertrust concern. Further, what about faulty interfaces, faulty communication protocols, unreliable clouds or clouds that leak data, and unreliable wireless service providers? These too are IoT cybertrust concerns.

And finally, artificial intelligence (AI) and “smart” are of a piece, there is not one without the other. With access to the computing power offered by clouds and the refinement of machine learning and other AI techniques, AI is a mainstay in automation, robotics, and the Industrial Internet of Things (IIoT). But how do you trust the AI

algorithms and implementations? Must you be a quant to do so?

This special issue is devoted to such questions, and thanks to incredibly high-quality submissions from around the world, we are able to explore the current research through the following five research papers.

IN THIS ISSUE

In “Evolution and Trends in IoT Security,” Rodrigo Román-Castro, Javier López, and Stefanos Gritzalis introduce some changes in information security that accompanied the adoption of IoT over time. In many areas, the challenges of IoT security have been met by adapting existing approaches, but for some problems, advances have been limited. These include forensics, and human-factors aspects of security and usability. Additionally, it is not clear how to do security engineering for IoT, due to the significant differences between these systems and traditional client-server environments. Although research has lagged in some of these areas, some interesting new developments could dramatically change the way security can be engineered for IoT systems. These include physically unclonable functions, which are hardware elements that work like one-way functions, with fingerprints that are easy to evaluate but hard to predict. The authors survey these and other developments, charting where progress is being made and where significant hurdles remain for protecting IoT systems.

In “IoT as a Land of Opportunity for DDoS Hackers,” Natalija Vlajic and Daiwei Zhou investigate the changes to DDoS risks being brought about by IoT equipment and search engines specifically focused on these small devices. By gathering all the information a hacker would need for a DDoS attack

using webcams, the authors show the ease with which an actual attack could be organized and carried out. They suggest that the attack potential of these devices could differ in both degree and type, compared with conventional attacks on servers. In particular, the existence of search engines that make it possible to identify large numbers of potentially vulnerable IoT devices, through highly specific search parameters, makes it possible for attackers to skip the reconnaissance step in gathering devices for a botnet. At the same time, IoT devices studied in the paper generally had little or no adherence to industry standards for anti-DDoS protection, as outlined in IETF RFC 4987. Recommending this RFC and other specific provisions, the authors outline a series of criteria for both vendors and users of IoT devices to reduce the DDoS threat.

Detecting attacks in IoT systems is particularly challenging because of the rapidly changing size and configuration of subnets. In “Intrusion Detection in the Era of IoT: Building Trust via Traffic Filtering and Sampling,” Weizhi Meng provides a case study showing how this problem can be addressed. By illustrating the application of packet filtering and sampling in a hierarchical IoT network using both Bayesian and blacklist-based filtering and two different sampling methods, the article identifies conditions under which intrusion detection techniques can be reasonably effective, and factors that reduce this effectiveness. Challenges remain in applying filtering and sampling, and the author summarizes aspects of attack models and limitations of Bayesian approaches in the IoT environment, demonstrating the need for different detection methods and to find the right balance among multiple techniques.

DISCLAIMER

The identification of any commercial product or trade name does not imply endorsement or recommendation by the National Institute of Standards and Technology, nor is it intended to imply that the materials or equipment identified are necessarily the best available for the purpose.

One approach to securing privacy is the use of attribute-based credentials. Jose Maria de Fuentes, Lorena Gonzalez-Manzano, Agusti Solanas, and Fatbardh Veseli describe this method in “Attribute-Based Credentials for Privacy-Aware Smart Health Services in IoT-Based Smart Cities.” The authors provide illustrative scenarios in which smart city technologies can improve the life of citizens, while simultaneously protecting their privacy. They advocate the use of attribute-based credentials, where users obtain some credentials or assured attributes from an issuer. With these, users can then create tokens proving possession of the credentials without revealing any other information, using zero-knowledge proofs possibly with blind signatures. Commercially available systems are analyzed in the context of realistic smart health scenarios, and the authors recommend one technology as the most effective for the hypothetical smart-city applications.

Beyond technical issues, IoT also presents new complications for public policy, the legal system, and regulatory agencies, as a result of some of the same factors that bring technical challenges. Jatinder Singh, Christopher Millard, Chris Reed, Jennifer

Cobbe, and Jon Crowcroft highlight some of the many legal issues of IoT in “Accountability in the IoT: Systems, Law, and Ways Forward.” In particular, IoT components might be owned and operated by different organizations, separated by management and geography. Additionally, the dynamic nature of IoT systems means that relationships among responsible parties are ever changing, and individual devices may be used simultaneously by multiple parties for different purposes. Yet accountability is critical to the success of the IoT industry, as it is in any IT field. This article gives insights into accountability aspects including governance and responsibility; privacy and surveillance; and safety and security issues, providing a valuable background that is essential but often not well understood by technologists.

ABOUT THE AUTHORS


JEFFREY VOAS is a computer scientist at the US National Institute of Standards and Technology (NIST). His research interests include IoT, blockchain, and software testing. Voas received a PhD in computer science from the College of William and Mary. He is a Fellow of IEEE. Contact him at jeff.voas@nist.gov.

RICK KUHN is a computer scientist in the Computer Security Division at NIST. His research interests include combinatorial methods in software testing/assurance, access control, and empirical studies of software failure. Kuhn received an MS in computer science from the University of Maryland College Park. He is a Fellow of IEEE. Contact him at kuhn@nist.gov.

CONSTANTINOS KOLIAS is a research assistant professor at George Mason University. His research interests include IoT security, intrusion detection systems, and wireless networks security. Koliass received a PhD in wireless intrusion detection systems from the University of the Aegean. Contact him at kkolias@gmu.edu.

ANGELOS STAVROU is a professor and the director of the CARE Center at George Mason University. His research interests include large systems security & survivability, intrusion detection systems and security for IoT and mobile devices. Stavrou received a PhD in computer science from Columbia University. He is a Senior Member of IEEE. Contact him at astavrou@gmu.edu.

GEORGIOS KAMBOURAKIS, is an associate professor in the Department of Information and Communication Systems Security and director of the Laboratory of Information and Communication Systems Security (Info Sec Lab) at the University of the Aegean. Contact him at gkamb@aegean.gr.

The articles in this special issue were selected to explore the state of cybertrust in the age of IoT. We hope that readers will find these articles to be an interesting and informative introduction to the challenges of developing trust in IoT-based systems. 

REFERENCES

1. I. Bojanova and J. Voas, “Trusting the Internet of Things,” *IT Professional*, vol. 19, no. 5, 2017, pp. 16–19.
2. J. Voas and P. Laplante, “IoT’s Certification Quagmire,” *Computer*, vol. 51, no. 4, 2018, pp. 86–89.
3. J. Voas, R. Kuhn, and P. Laplante, “Testing IoT-based Systems,” *Proc. 12th IEEE Int’l Symp. Service-Oriented System Engineering (SOSE 18)*, 26–29 Mar. 26–29, 2018, pp. 48–52.