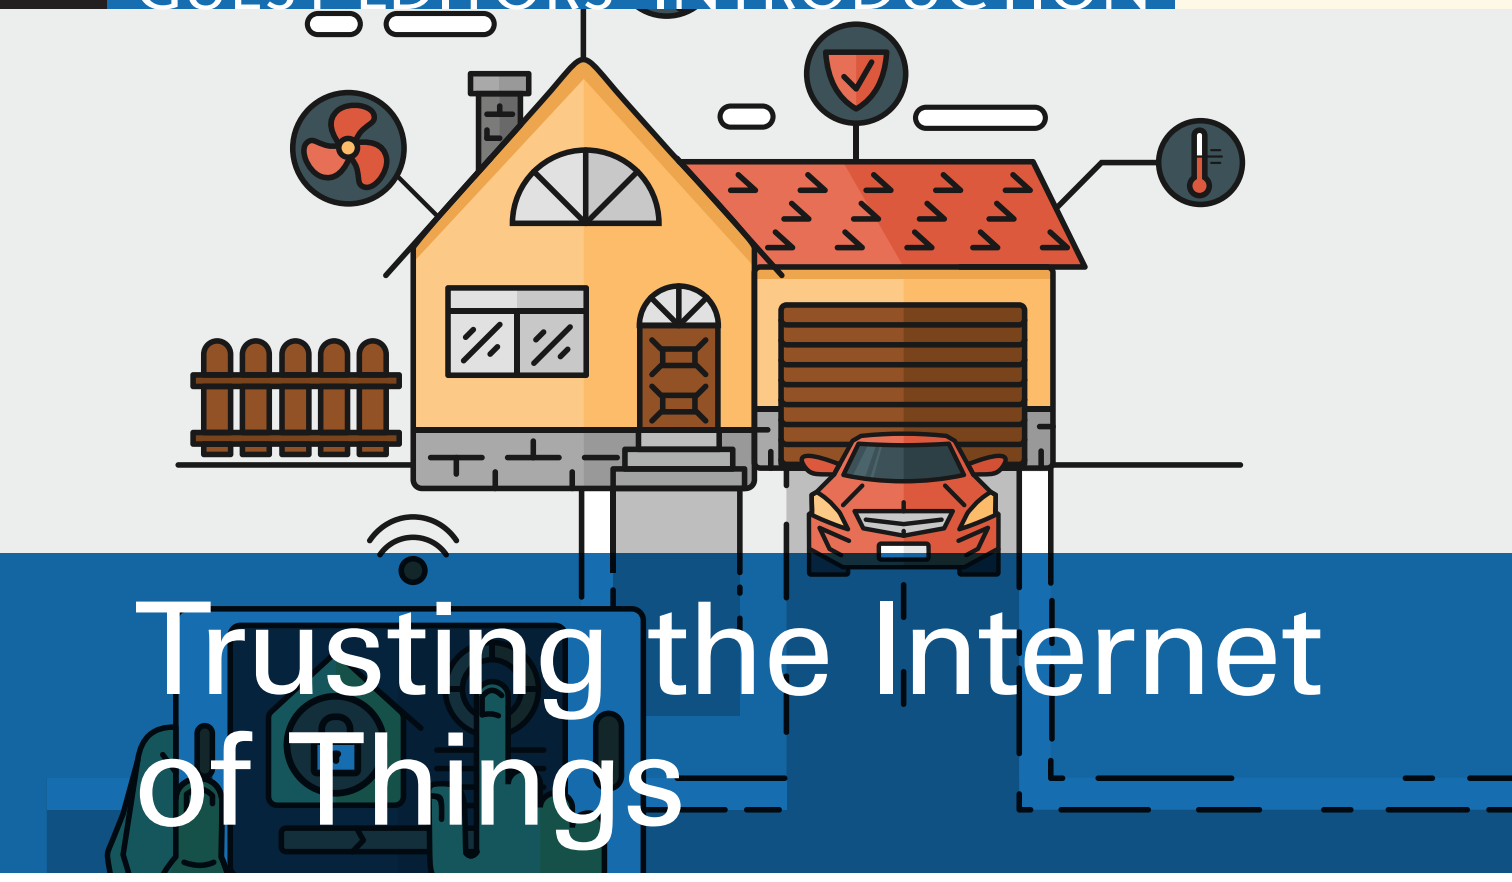

Trusting the Internet of Things

Irena Bojanova and Jeffrey Voas, *NIST*

The Internet of Things (IoT) is a technology revolution. It is expected to affect everything from healthcare delivery, to how food is produced, to how we work, to all forms of transportation and communication, and to virtually all forms of automation. With that said, the IoT will impact everyone, and in multiple ways. With a technology advance with such large effects on society, it is imperative that IoT-based systems can be trusted. This means that they should exhibit some level of secure, reliable, and private behaviors, as well as other attributes generally associated with quality.

Challenges Facing the IoT

The above-mentioned desirable behaviors are not new to the IoT—they have been around for previous generations of IT and physical systems for decades. Here, we highlight eight key differences between former IT systems and future IoT systems. Previous IT systems benefited from having fewer of the following concerns.

No Actionable and Universally Accepted Definition

What is the IoT? It would be easier to address trust-related issues in the IoT if we all agreed as to what it is. IT systems such as cloud computing were around for years before achieving a "generally agreed upon" definition.[1] The IoT is likely to follow the same path, but in the meantime, it will be defined by de facto standards and interpretations. This is simply the result of a vacuum that gets filled only by such de facto interpretations.

Large Scalability

Because of the important role that sensors play in the IoT, their large number—as data generators into a system—can quickly overwhelm the ability of a network of things to handle the workflow and dataflow needed to achieve system goals. The IoT will likely be a strong provider of so-called big data. The bottom line here is that data and information overload is not a recipe for improving trust.

 Published by the IEEE Computer Society 1520-9202/17/$33.00 © 2017 IEEE

Large Heterogeneity

Because things in a system can be acquired from many different vendors of differing integrity and reliability, a potential system can have many pieces and parts connected that would not have been imaginable had the supply of things come from a smaller vendor pool. Supply chain issues, counterfeit parts, and a real lack of understanding of "thing pedigree," along with high heterogeneity, become IoT trust concerns.

Lack of Ownership and Component Control

Not all things in a system can be under your control. You might use a public cloud, a commercial Internet and wireless provider, and numerous third-party components. You might lease data from other providers that own the sensors that create the data. In short, you might have a system of things for which you have near zero control except for the system's architecture, if even that.

Few Rules on Composability and Interoperability

This is a byproduct of the first four concerns. Consider a simple example—you have fish A in a saltwater aquarium and fish B in a freshwater aquarium. You mix both aquariums into a new one. What will happen? The answer is twofold: whether the two fish will get along (for example, not kill each other) and whether the new hybrid (brackish) water will allow them to survive. This highlights the heterogeneity of the two water environments along with the interoperability (maliciousness or lack thereof) of the fish. This simple example showcases why trust concerns, under suspect environmental composability expectations, cannot be discounted.

Little Regulatory Oversight and Governance

For systems deemed safety-critical, oversight and governance is expected, accepted, and understood. Such systems are expensive and slow to create. There will likely be different vertical domains and criticality domains in which the IoT can be used: consumer grade, military grade, industrial grade (the Industrial Internet), classified networks, and so on. Those involved in policy, governance, and regulation, including nongovernmental organizations (NGOs) and standards organizations, have yet to step up to the challenge here. What challenge? Well, defining the IoT would be a start. Understanding how regulation and governance fit into these different grades will be a challenge for all users, given the previous concerns we've discussed.

Lack of Standards and Certification

Building standards takes years, and the IoT community is far more likely to wind up with de facto standards than prescriptive standards. The same holds for certification—we are more likely to wind up with a catalogue of a few products that become "gold standards" due to their rapid overtaking of market share than products that might have been certified in a traditional manner but delayed in market delivery. Those involved in policy, governance, and regulation, including NGOs and standards organizations, have yet to step up to the challenge here. This trust concern relates to the aforementioned lack of oversight and governance.

Lack of Specific and Appropriate Testing Approaches

Given the first six concerns we describe, where do we begin when talking about appropriate ways to test things and systems of things? The notion of agile programming and producing products quickly to meet market demand or beat competitors to the market compromises trust. Systems take time to test. There is no agreed-upon notion of how to test networks of things.

(More information on these trust concerns and others is available in prior work.[2])

Forces Acting on the IoT

In addition to the eight concerns we've listed, we believe there are three significant additional forces at work here that play a significant role in determining the trustworthiness of IoT-based systems.

The first is speed.[3] The speed at which computations and data generation can occur in a system of things is increasing rapidly. This speed affects the systems' ability to log and audit these transactions in a manner that makes forensics and recovery from faults and failures less likely.[4] That is, there are fewer ways to "put on the brakes," undo incorrect computations, and fix internal and external data anomalies. Computing faster to the wrong outcome offers no trust.

The second force is data and information overload. Do you remember years ago when you could ask colleagues if they had read a recent paper and they said "yes"? Today, we are all bombarded with publications and blogs nonstop, and that answer is more likely "no" because your colleagues probably read different content than you did. The idea of seemingly limitless sensors and the data associated with them lends itself to an inability to take the time necessary to weed out corrupt data and then perform data analytics on sound data. The recent advent of the term "fake news" applies here. If you are performing analytics or computations on corrupt data, trust is a foolish expectation.

Finally, the third force is autonomy, robotics, and artificial intelligence (AI). Most agree that we might have reached the point at which AI has finally come into its own. Many promises were made back in the 1980s about AI that did not come true then, but now with the computing power of clouds and the refinement of machine learning and other AI algorithms, AI is becoming a key player in automation, robotics, and the Industrial Internet. But how do you trust the algorithms? Must you be a quant to do so? This is a terrific question for research.

In This Issue

The three forces and eight concerns we've discussed create a perfect storm against achieving trust in the IoT. Now, having given you our opinion on some of the trust challenges, we turn our focus to introducing you to the three articles selected for this special issue to see what our authors think are key trust issues in the IoT.

Cybersecurity risk assessment approaches provide a platform for better protection against pertinent risks. In "Security Risk Assessment in Internet of Things Systems," Jason R.C. Nurse, Sadie Creese, and David De Roure from the University of Oxford, UK, argue about the need for new approaches to assess risk and build trust as the complexity, pervasiveness, and automation of technology systems increases, particularly with the IoT. New risks could be arising in this ecosystem related to the high degrees of connectivity present or the coupling of digital, cyber-physical, and social systems. The article seeks to make a case for new methodologies to assess risk in this context that consider the dynamics and uniqueness of the IoT while maintaining the rigor of best practice in risk assessment. The authors discuss the current cybersecurity risk assessment paradigm, the relevant dynamics of the IoT, where current risk assessment methods fail in the IoT, and the need for new approaches to assess IoT system risk.

Authentication and authorization are essential parts of basic security processes. In "Authentication and Authorization for the Internet of Things," Hokeun Kim and Edward A. Lee from the University of California, Berkeley, remind us that these processes are sorely needed in the IoT. Their article focuses on how the emergence of edge computing (fog computing) creates new opportunities for security and trust management in the IoT. The authors discuss IoT security challenges and ways of building trust in networked systems, and introduce a network architecture using local authentication and authorization entities. They also examine the challenges for a more secure authorization infrastructure.

The IoT should support the demanding smart systems of the 21st century, such as smart cities, smart transportation, smart healthcare, and the smart power grid. In "Security and Privacy for a Green Internet of Things," Ted H. Szymanski from McMaster University, Canada, presents an approach to achieving exceptional performance, cybersecurity, and privacy in an Industrial-Tactile IoT, in datacenters, and in green cloud computing systems. Deterministic communications, a software-defined networking (SDN) control plane, and lightweight layer-2 encryption are combined to achieve various benefits. These include an SDN control plane that embeds millions of deterministic virtual networks in layer 2; the removal of congestion, interference, and denial-of-service (DOS) attacks and targeted cyberattacks in layer 2; the detection within microseconds of unauthorized packets from a cyberattacker in layer 2; the reduction of IoT delays to the speed of light; and the achievement of exceptional privacy using lightweight encryption with long keys. The author has developed a field-programmable gate array hardware testbed to illustrate these concepts.

In summary, systems composed from things with no verifiable trust certainly make arguments of achieved system-level trust difficult

to defend. In many cases, we will be reduced to alternatives such as predictive methods that attempt to quantify confidence levels in trust. These confidence levels can be created using a variety of evidence, such as from testing or even qualitative methods. While there are those who champion formal methods, we wonder whether formal methods will become practical enough to be employed for IoT-based systems. We hope you enjoy this theme issue. And as always, we appreciate feedback from the readers and suggestions for other topics related to trust and the IoT that might turn into future theme issues of *IT Pro*.

References

1. P. Mell and T. Grance, *The NIST Definition of Cloud Computing*, NIST Special Publication 800-145, Sept. 2011; bit.ly/24r0tMW.
2. J. Voas, *Networks of "Things,"* NIST Special Publication 800-183, July 2016; dx.doi.org/10.6028/NIST.SP.800-183.
3. L. Barroso et al., "Attack of the Killer Microseconds," *Comm. ACM*, vol. 60, no. 4, 2017, pp. 48–54.
4. A. Stavrou and J. Voas, "Verified Time," *Computer*, vol. 50, no. 3, 2017, pp. 78–82.

Irena Bojanova *is a computer scientist at NIST, a member at large of the IEEE CS Publications Board, and an associate editor in chief of* IT Professional. *Read her blogs "Sensing IoT" and "A Cloud Blog" on Computing Now. She is a senior IEEE member. Contact her at irena.bojanova@computer.org.*

Jeffrey Voas *is a computer scientist at NIST, a cofounder of Cigital, and* Computer*'s Cybertrust column editor. He's an IEEE Fellow. Contact him at j.voas@ieee.org.*

myCS Read your subscriptions through the myCS publications portal at **http://mycs.computer.org**