DOI:10.1145/3029589

Ulf Lindqvist and Peter G. Neumann

Inside Risks
The Future of the Internet of Things

The IoT can become ubiquitous worldwide—if the pursuit of systemic trustworthiness can overcome the potential risks.

As suggested in the previous *Communications* Inside Risks column ("Risks of Automation," October 2016[8]), the Internet of Things (IoT) has the potential to encompass and instrument an enormous range of connected devices—including home appliances and utilities, wearables, homes and corporate buildings, industrial processes, medical devices, law-enforcement devices, military equipment, and other connected applications that today might be barely imaginable. In the present context, "Things" are simply those computerized and networked devices that become part of the IoT. Some of those Things will be directly accessible over the Internet, whereas others would be supposedly hidden in local networks behind firewalls and address-translating routers.

There are already many risks recognizably associated with the IoT. Some risks are old and well known, but exacerbated by the unprecedented scale of the IoT; estimates for the next few years suggest tens of billions of Things. Other risks may be new, stemming from the nature of how these Things are designed, what they are used for, how they are deployed and managed (or not managed), and how market forces will influence the development. In this column, we outline some of those risks and what might need to happen if the IoT is to deliver the benefits envisioned for it—with a reasonable level of trustworthiness. Our message is intended as a wake-up call for computer professionals, but is also relevant to everyone involved as a user.

Security and privacy are both extremely important in the IoT, because the potential consequences of successful attacks could impact human lives and safety, and cause death and destruction—directly or indirectly. Privacy violations that let criminals exploit information about potential victims can also constitute threats to safety.

Things Turning Evil

A recent distributed denial-of-service (DDoS) attack[7] has demonstrated the ubiquitousness of vulnerabilities in the current still-primitive Internet of Things. Many devices including closed-circuit TV cameras, cable set-top boxes, and digital video recorders (DVRs) were compromised and used as unwitting botnet zombies. This significant event used malware (Mirai) that searches for vulnerable victims, and whose source code had been freely published. By targeting the DNS services provided by Dyn, this attack seriously interfered with user access to major services such as Twitter, Amazon, Tumblr, Reddit, Spotify, and Netflix. In one fell swoop, it exposed the tip of just one of many hazardous icebergs. While earlier DDoS attacks using Mirai had exploited hundreds of thousands of devices, this attack appeared to involve tens of millions of compromised devices—according to a statement from Dyn.[13] The attack illustrates some of the risks associated with having very large numbers of inadequately protected Things connected to the Internet—particularly Things that are simple enough to be vulnerable to compromise, but sufficiently capable to be part of a distributed attack that floods the victims' sites with seemingly legitimate requests. Note that the own-

ers or users of compromised devices are often not aware their devices are being used to attack other systems.

Vulnerabilities

Evidently, many of these devices that unwittingly contributed to that DDoS attack were not actually behind any sort of firewall, or else had weak default firewall configurations that were easily exploited. Furthermore, some of the Things infected by Mirai were themselves small-office or home-office routers. While Mirai specifically exploited hardcoded passwords for Telnet/SSH services that users could not disable, it is generally foolish to put all the blame on any one weak link, when almost everything is a potential weak link.

Today, almost every computer-related system is likely to be already compromised, or else easily misused. We have *weakness in depth and breadth*, not *strength in depth*. Therefore, many problems will need to be overcome to make the IoT viable. We consider some of those problems, and some possible remediations. Ultimately, we need a total-system perspective that addresses the potential vulnerabilities in the devices, the alleged firewall security, the network connections, the cloud services (some not even known to the users), and the Internet itself, as well as all its users and would-be malfeasors. The IoT is not an entity per se—it encompasses all of these entities and inevitably depends on them.

We suggest this recent DDoS botnet episode is merely a harbinger of events to come. IoT risks in the future will be pervasive, including potential compromises of requirements relating to trustworthiness. Such requirements must address networkwide issues such as human safety, security, reliability, robustness, resilience, functional interoperability, seamless ease of installation and use, rapid automated remediation of serious flaws, personal as well as institutional privacy, human well-being, and much more.

Some Illustrative IoT Risks

Denial-of-service attacks are damaging, but the ability to subvert Things remotely for arbitrary manipulation must be considered particularly threatening. Here are just a few examples of application areas where the use of IoT devices brings inherent risks:

- Hospitals and healthcare establishments tend to use devices that are already remotely controlled or accessible Things: patient monitors, body scanners, pacemakers, defibrillators, infusion pumps, main and auxiliary power, lighting, air conditioning, and much more.
- Critical infrastructure sectors such as electric power, oil, natural gas, manufacturing, and transportation use IoT devices as sensors and actuators for automation and remote monitoring and control. The controllers themselves may be Internet accessible.
- Self-driving and automation-assisted interconnected automobiles must clearly be considered as Things, especially in automated highways of the future. Recent demonstrations of the ability to remotely take over critical vehicle controls illustrate just a few of the risks.[5]

Unlike general-purpose computers, IoT devices may be more closely associated with the physical world. While there have so far been relatively

IMAGE BY ALICIA KUBISTA/ANDRIJ BORYS ASSOCIATES

few cases where physical destruction has been intentionally caused through computer compromise, this is likely to be a risk of serious concern for the IoT. From the known cases of programs in the 1960s that could exercise disk arms to cause the drives to self-destruct, to the 2007–2010 Stuxnet attack that appeared to be designed to damage nuclear enrichment centrifuges (and reportedly succeeded), cyberphysical attacks have exploited vulnerabilities that are features rather than flaws. In addition to the Things that control switches, valves, and motors, many Things have batteries—which suggests the potential ability to remotely cause certain devices to overheat enough to cause a fire or explosion. If vehicles or medical devices are remotely taken over by malicious attackers, people could be injured or killed by someone clicking from anywhere on the Internet. Manipulation of sensors or insertion of misinformation could indirectly cause other health hazards by inducing chemical spills, disrupting energy systems, or misrouting vehicles. Thus, human safety must be a fundamental issue for many types of Things.

Another critical difference between IoT devices and general-purpose computers involves management. For a desktop computer, laptop, tablet, or smartphone, there are rich interactions between users and devices. Some notion of management also must exist: for corporate devices there are system administrators in important designated roles, while for personal devices the user is typically also the administrator. However, for IoT devices, there may be very little room for user interaction, and the concept of 'management' is unclear.

While operating systems and applications for general-purpose computers in desktop, laptop, tablet, or smartphone form factor tend to be easy to keep updated, many IoT devices are difficult or impossible for users to update. Some devices will remain in use for their entire lifetimes, precisely as delivered—unless they are recalled, discarded, or just forgotten. In some of those cases, security updates will be essentially impossible or extremely difficult. In other cases, devices may be directly accessible remotely over the Internet; any update mechanisms must be secured so that attackers cannot subvert them and insert their own updates or attacks.

For Things that necessarily have interactions with human users, their small size typically will not allow for touchscreens or keyboards. Thus, they must either rely on another device such as a tablet or smartphone for interaction, or else use other emerging modes of interaction such as voice inputs. For voice interfaces, there are problems with linguistic ambiguities, and obvious privacy risks associated with ubiquitous devices that continuously record and process voice conversations, as well as interesting opportunities for replay or synthesized voice-command attacks from one device to another device. As already evident in advertising applications, audio interfaces could also be used for covert ultrasound communication, inaudible to humans.[9]

Whereas botnet attacks may typically be stopped by blocking the command and control servers that orchestrate the attacks, the individual IoT devices are still compromised, and could be pulled into a new botnet at any time. We are left with many questions. For example, who is responsible for fixing these devices? What incentive would the owner of a connected camera have for going through the trouble of updating its firmware if it seems to work just fine as it is? Who is liable when major disruptions occur? Is it the manufacturer, the vendor, the person or organization who deployed the device, the cloud or back-end communications provider, or the unwitting user of the device? Each of these alternatives entails its own set of risks.

Until recently, consideration of most of these risks has been dominated by the competitive rush to market, with very few concerns for trustworthiness. This reality tends to cause security and privacy to be sadly neglected. Clearly, that must change, suggesting the advent of some serious far-sighted systemic considerations—especially where the risks might be greatest.

Another critical difference between IoT devices and general-purpose computers involves management.

Confronting the Risks

We next attempt to outline some steps that might be desirable. As has been noted in past Inside Risks columns, we have a serious need for considering risks in the context of total systems. The Internet of Things requires a much deeper concern for total-system trustworthiness, in which the security of Things is only one aspect—especially because at the moment there is essentially no real security in computer systems and networks. This reality is clearly making the problems of assuring trustworthiness much more difficult.

We enumerate here just a few of the steps that might be helpful for developers, administrators, and users. However, we explicitly caution that this summary is only an essential beginning, and inherently incomplete. It may not be surprising that what is needed is more or less consistent with the series of National Academies' Computer Science and Technology Board reports over the past several decades, including most recently.[2] In addition, NIST's Special Publication 800-160, Computer Security Resource (Nov. 2016; https://doi.org/10.6028/NIST.SP.800-160), addresses important engineering aspects. Also, in the context of the IoT, we need to reemphasize many topics that have been discussed in the Inside Risks series more generally and that are highly relevant here.

Some IoT devices will have simple applications running on bare metal, that is, without general-purpose operating systems. Other Things might need simple operating systems focusing just on specialized requirements such as real-time guarantees, while yet others may require full-fledged operating systems. Thus, scalable hardware and software are likely to be useful for economic reasons and operational effectiveness. Implementations are likely to range from micro-operating systems on small processors to larger reprogrammable environments for centralized

control of Things for entire enterprises. Similarly, a range of development support is needed—from totally embedded as-delivered hardware with no possibility of software changes (except perhaps for recalls and possible remote updates) up to Things with flexible development environments and programming-language support. Thus, programming languages and compilers might need to encompass the very simple and the much more complex. Concerns for greater trustworthiness will be important, especially for embedding potentially unsecure applications into a nevertheless trustworthy environment.

Users generally lack expertise and patience, have limited ability to cope with complexity, and are unaware of corner cases. Consequently, the design and implementation of user interfaces for Things and their controllers will require special attention and care. These interfaces need to be seamlessly easy to use, intuitively self-evident, and friendly for those who are technologically impaired, as well as adequately configurable by everyone. Particularly problematical are easily managed Things that exist today (conventional light bulbs, toasters, and so on) whose computerization might render them completely unusable when they fail. Even worse might be mechanically fail-safe devices today that might no longer work manually. One such example might be a fully automated automobile whose doors cannot be opened from the inside if the battery dies or the car is under water, or perhaps a refrigerator door that cannot be opened because its Thing controller has crashed—or been hacked. Fail-sensible techniques will be essential.

The needs for seamless installation and integration are critical from the customers' viewpoint, but this should not be a motivation for ignoring security. One of the major risks here is the prevailing quest for simplicity—for example, just barely meeting the bar for compliance with standards and expectations, as well as poorly addressing needs for ease of installation and ease of use. Standards are needed to facilitate interoperable installations involving many different vendors' devices. Connection protocols should not be as simplistic and unsecure as they often are today.

Today's supposedly sage advice about how to deal with safety and security needs to be significantly upgraded.

Any local networks within a home or enterprise must be suitably isolated from the Internet and other outside connections—except where interactions are explicitly desired and adequate protection can be assured. Certain systems and Things will to some extent have to be resilient and resistant to insider misuse, although that may be less important to friendly homes than corporate entities. On the other hand, Internet firewalls must be much more impervious to outsider misuse than today. Ideally, fixed passwords and default encryption keys should be eschewed in devices—although they are far too common today, and indeed were exploited by the Mirai malware (as noted previously). Nevertheless, there will be cases where trustworthy updates cannot be achieved and recalls might be the only alternative. To enforce recalls, firewalls may need to recognize traffic from recalled and/or compromised Things, and block the communication to protect systems on the rest of the Internet.

Also, we must consider needs for oversight, consumer protection, regulation, and liability for flagrant violations that result in serious risks. As software makes its rapid transfer into our physical world through "smart" Things, we cannot afford to simply transfer the notion that software tends to be provided "as is"—without liability for the consequences of flaws. Electronic products that have the potential to hurt or kill people are typically subject to some form of government regulation and testing to protect consumers. When the safe operation of a product is dependent on its software being secure and reliable, regulation will need to address those aspects of product safety. Also, the responsibilities of everyone involved need to be established and made clear. For example, if your home burns down because of a hacking attack on your IoT installation, or your negligence in failing to protect your technological devices, could your insurance companies deny coverage for known but unaddressed vulnerabilities, or even preexisting conditions?

In summary, we will need some meaningfully trustworthy hardware and software components, and much better development and deployment practices than we have at present—to enable the IoT to provide adequate human safety, security, reliability, usability, and satisfied users.

Some Specific Efforts

It is highly desirable to study a few types of Things as developing prototypes in research and development, and attempt to ensure that all reasonable risks have at least been addressed. We would benefit from a few very successful cases to pave the way for how this could be done in the future. The combination of system engineering, hardware and software engineering, and careful application development—perhaps with some formal analyses to provide better assurance—would be extremely valuable to everyone else competing in the IoT marketplace. Thus, a few well-designed, well-developed, and trustworthy systems that are well documented would provide wonderful examples for other developers. A step in that direction is the documented example of principled security design for a fictitious wearable fitness-tracking system that was produced by the IEEE's Center for Secure Design under the auspices of the IEEE Cybersecurity Initiative.[12]

It would also be very important to provide developers with the tools and knowledge to build security, privacy, reliability, and other aspects of trustworthiness into the systems that they build. This is particularly important for developers of IoT systems who may have even less security expertise than traditional software developers. We have recognized this need, and are involved in several efforts to address the situation—including the new IEEE Cy-

bersecurity Development Conference (IEEE SecDev),[3] and a strategic independent R&D initiative at SRI International on IoT security and privacy.

Some Thoughts for the Future

Today's population displays a wide range when it comes to understanding computer technology, ability to use it, and to have access to it. We can't deny access to essential services to portions of the population by ignoring their inability to correctly use certain technologies. Above all, we have serious needs for better computer literacy in the entire population.

Many of the risks and needs discussed here are not just specific to the Internet of Things, and have commonalities with more general uses of computers. However, we must also consider self-driving vehicles as Things in the evolving automated highways, as well as automated airplanes—and treat them similarly in the same basic context. The very concept of the IoT brings us to a much more personal and visceral focus in its manifestations in homes, vehicles, and wearables, and in that sense it touches everyone to some extent. Even those who are unwilling may eventually be forced to buy IoT-enabled appliances, simply because there are no longer any alternatives.

Today's supposedly sage advice about how to deal with safety and security needs to be significantly upgraded. For example, while we are familiar with admonitions such the following, not everyone follows them: Beware of social engineering, hucksters, and easy solutions! Don't click on suspicious links! Don't display your most personal information on social media! Adhere to (or better yet, exceed) best practices for security! The new risks will be much more pervasive, and we will need to determine what reasonable caution and common sense will look like in the world of the IoT. Indeed, the IoT is likely to become very contentious unless serious coordinated efforts are made proactively by governments, standards committees, purveyors of Things and Thing infrastructures (including the Internet itself) and user communities. For considerable further background, please see recent testimony before the U.S. Congress.[4,10] Also, some so-called best practices are considered in recommendations from the Department of Homeland Security[11] and BITAG.[6]

The future may be very murky unless proactive attention is paid to decide which Things can realistically be implemented wisely—and which might be simply too risky.

However, as we have noted in earlier Inside Risks columns, best practices are generally nowhere near good enough.

Considering the Keys Under Doormats report,[1] the prospect of billions of sensor-equipped and Internet-connected IoT devices would be tempting to any organization that wants to collect information for intelligence or evidence, or to exploit the devices for propagating DDoS attacks, or other nefarious purposes. The risks of dumbing down cybersecurity and cryptography for such purposes would be enormous—especially with respect to the IoT.

There is much more on this topic than could be written here. However, this column is only an initial stake in the ground. Overall, there are no easy answers, but the time to begin asking the incisive questions is now.

Conclusion

We have described problems and potential risks that are associated with the evolving Internet of Things. It remains to be seen whether the IoT and its Things can *burgeon* (grow and flourish, as the way of the future), or *sturgeon* (sometimes surviving competitively for up to two decades if not caught), or be more like the female salmon (with very short lives once they spawn). In any case, we need much more than a *surgeon* to fix things (and Things). Incremental change is not likely to succeed (indeed, it has been ineffective for so many years), and some sort of radical change may be needed. The future may be very murky unless proactive attention is paid to decide which Things can realistically be implemented wisely—and which might be simply too risky. We must then ensure that those beneficial Things can be integrated into the necessary total-system trustworthiness (which we do not yet have). Thus, we need to *urge on* to make the IoT truly usable, and then *surge on* to ensure that it happens with appropriate trustworthiness. C

References

1. Abelson, H. et al. Keys under doormats: Mandating insecurity by requiring government access to all data and communications. *Journal of Cybersecurity 1*, 1 (Nov. 17, 2015); Oxford University Press; http://www.cybersecurity.oxfordjournals.org/content/1/1/69
2. Computer Science and Technology Board, National Academies of Science, Engineering, and Medicine. *Foundational Science for Cybersecurity*, final report, 2017.
3. Cunningham, R. et al. IEEE SecDev 2016: Prioritizing Secure Development. IEEE Security and Privacy (July–Aug. 2016), 82–84. https://www.computer.org/csdl/mags/sp/2016/04/msp2016040082.pdf
4. Fu, K. Infrastructure Disruption: Internet of Things Security, Testimony before the U.S. House of Representatives Committee on Energy and Commerce, Subcommittee on Communications and Technology and Subcommittee on Commerce, Manufacturing, and Trade (Nov. 16, 2016); https://energycommerce.house.gov/hearings-and-votes/hearings/understanding-role-connected-devices-recent-cyber-attacks
5. Greenberg, A. Hackers remotely kill a Jeep on the highway—With me in it. *Wired* (July 21, 2015); http://www.wired.com/2015/07/hackers-remotely-kill-jeep-highway/
6. Internet of Things (IoT) Security and Privacy Recommendations, BITAG Broadband Internet Technical Advisory Group, November 2016: http://www.bitag.org/documents/BITAG_Report_-_Internet_of_Things_(IoT)_Security_and_Privacy_Recommendations.pdf
7. Krebs on Security (Oct. 21, 2016); https://krebsonsecurity.com/2016/10/hacked-cameras-dvrs-powered-todays-massive-internet-outage/
8. Neumann, P.G. Risks of automation: A cautionary total-system perspective of our cyberfuture. *Commun. ACM 59*, 10 (Oct. 2016); http://www.csl.sri.com/neumann/insiderisks.html#240
9. Newman, L.H. How to block the ultrasonic signals you didn't know were tracking you. *Wired* (Nov. 3, 2016); https://www.wired.com/2016/11/block-ultrasonic-signals-didnt-know-tracking/
10. Schneier, B. Testimony before the U.S. House of Representatives Committee on Energy and Commerce, Subcommittee on Communications and Technology and Subcommittee on Commerce, Manufacturing, and Trade (Nov. 16, 2016); https://energycommerce.house.gov/hearings-and-votes/hearings/understanding-role-connected-devices-recent-cyber-attacks
11. Strategic Principles for Securing the Internet of Things. Department of Homeland Security, along with an IoT Fact Sheet (Nov. 15, 2016); https://www.dhs.gov/securingtheIoT
12. West, J. et al. WearFit: Security Design Analysis of a Wearable Fitness Tracker, February 2016; http://cybersecurity.ieee.org/blog/2016/02/17/wearfit-security-design-analysis-of-a-wearable-fitness-tracker/
13. York, K. Dyn Statement on 10/21/2016 DDoS Attack (Oct. 22, 2016); http://dyn.com/blog/dyn-statement-on-10212016-ddos-attack/

Ulf Lindqvist (ulf.lindqvist@sri.com) is a Program Director in the Computer Science Lab at SRI International.

Peter G. Neumann (peter.neumann@sri.com) is Senior Principal Scientist in the Computer Science Lab at SRI International, and moderator of the ACM Risks Forum.